

Nuitka Changelog

In this document, we track the per version changes and comments. This becomes a document on the website, as well as individual posts on the Nuitka blog.

Nuitka Release 1.7

This release is focused on adding plenty of new features in Nuitka, with the new isolated mode for standalone being headliners, but there are beginnings for including functions as not compiled, and really a lot of new anti-bloat new features for improved handling, and improving user interaction.

Also many packages were improved specifically to use less unnecessary stuff, some of which are commonly used. For some things, e.g. avoiding tkinter, this got also down to polishing modules that have GUI plugins to avoid those if another GUI toolkit is used.

In terms of bug fixes, it's also a lot, and macOS got again a lot of improvements that solve issues in our dependency detection. But also a long standing corruption for code generation of cell variables of contractions in loops has finally been solved.

Bug Fixes

- Python3.11: The MSVC compiler for Windows will not work before 14.3 (Visual Studio 2022) if used in conjunction with Python 3.11, point it out to the user and ignore older versions. Fixed in 1.6.1 already.
- Standalone: Added support for the `pint` package. Fixed in 1.6.1 already.
- Standalone: Added missing standard library dependency for `statistics`. Fixed in 1.6.1 already.
- Compatibility: Fix, the `transformers` auto models were copying invalid bytecode from compiled functions. Added workaround to use compiled function `.clone()` method. Fixed in 1.6.1 already.
- Compatibility: Added workaround for `scipy.optimize.cobyla` package. Fixed in 1.6.1 already.
- Anaconda: Detect Anaconda package from `conda install` vs. PyPI package from `pip install`, the specifics should only be applied to those. Adapted our configurations to make the difference. Fixed in 1.6.1 already.
- Anaconda: Do not search DLLs for newer `shapely` versions. Fixed in 1.6.1 already.
- Standalone: Add new implicit dependencies for `pycryptodome.ECC` module. Fixed in 1.6.1 already.
- Standalone: Fix `tls_client` for Linux by not non-Linux DLLs. Fixed in 1.6.1 already.
- MacOS: When using `--macos-app-name`, the executable name of a bundle could become wrong and prevent the launch of the program. Now uses the actual executable name. Fixed in 1.6.1 already.
- Multidist: The docs didn't properly state the option name to use which is `--main` and also it didn't show up in help output. Fixed in 1.6.2 already.
- Standalone: Added support for `polars` package. Fixed in 1.6.3 already.
- Standalone: Added implicit imports for `apscheduler` triggers. Fixed in 1.6.3 already.
- Standalone: Add data files to AXML parser packages. Added in 1.6.4 already.
- Fix, `exec` nodes didn't annotate their exception exit. Fixed in 1.6.4 already.
- Standalone: Added data files for `open_clip` package. Fixed in 1.6.4 already.
- Standalone: Avoid data files warning with old `pendulum` package. Fixed in 1.6.4 already.
- Standalone: Added implicit dependencies for `faker` module. Fixed in 1.6.4 already.

- Added workaround for `opentele` exception raising trying to look at the exception frame before its raised. Fixed in 1.6.4 already.
- Nuitka-Python: Do not check for unknown built-in modules. Fixed in 1.6.4 already.
- Scons: Fix, the total `ccache` file number given could be wrong. Ignored messages were counted still as compiled, leading to larger sum of files than actually there was. Fixed in 1.6.5 already.
- Fix, multiprocessing resource tracker was not properly initialized. On at least macOS this was causing it to work relatively badly, because it could fail to actually use it. Fixed in 1.6.5 already.
- Standalone: Added support for `cassandra-driver` package. Fixed in 1.6.5 already.
- Onefile: Have Python process suicide when bootstrap surprisingly died, respecting the provided grace time for shutdown. Fixed in 1.6.5 already.
- Plugins: Fix, package versions for at least Ubuntu packages can be broken, such that at least `pkg_resources` rejects them. Handle that and use fallback to next version detection method. Fixed in 1.6.5 already.
- Onefile: Handle `SIGTERM` and `SIGQUIT` just like `SIGINT` on non-Windows. The Python code with see `KeyboardInterrupt` for all 3 signals, so it's easier to implement. Previously onefile would exit without cleanup being performed. Fixed in 1.6.5 already.
- Standalone: Fix, need to add more implicit dependencies for `pydantic` because we do no longer include e.g. `decimal` and `uuid` automatically.
- Standalone: Added missing implicit dependencies for `fiona` package. Added in 1.6.6 already.
- Standalone: Added missing implicit dependencies for `rasterio` package. Added in 1.6.6 already.
- Standalone: Fix, need to add more implicit dependencies for `pydantic`. Added in 1.6.6 already.
- Fix, the data composer used a signed value for encoding constant blob sizes, limiting it needlessly to half the size possible.
- Windows: Avoid dependency on API not available on all versions, specifically Windows 7 didn't work anymore. With this, symlinks are only resolved where they actually exist, and MinGW64 does it too now.
- Standalone: Added support for `.location` attribute for `pkg_resources` distribution objects.
- Anti-Bloat: Avoid using `dask` and `numba` in the `tsfresh` package.
- Fix, outline cell variables must be re-initialized on entry. The code would be crashing for for outlines used in a loop, since the cleanup code for these cell variables would release the cell that was created during containing scope setup.
- Standalone: Added missing dependency of `pygeos` package.
- Standalone: Added `sqlalchemy` implicit dependency.
- Standalone: Added data files for `mnemonic` package.
- Fix, attribute checks could cause corruption when used on objects that raise exceptions during `__getattr__`.
- Python2: Fix, wasn't making sure instance attribute lookups were actually only done with `str` attributes.
- macOS: Fix, need to allow versioned DLL dependency from un-versioned DLLs packaged.
- Standalone: Added DLLs for `rtree` package.
- Standalone: Added support for newer `skimage` package.
- Standalone: Added support for newer `matplotlib` package.
- Standalone: Fix, our `numpy.testing` replacement, was lacking a function `assert_array_almost_equal` used in at least the `pytransform3d` package.

New Features

- Added support for `--python-flag=isolated` mode. In this mode, packages are not expandable via environment variable provided paths and `sys.path` is emptied which makes imports from the file system not work.
- The options for forcing outputs were renamed to `--force-stdout-spec` and `force-stderr-spec` to force output to files and now work on non-Windows as well. They kind of were before, but e.g. `%PROGRAM%` was not implemented for all OSes yet.
- Capturing of all outputs now extends beyond the Python level outputs is now attempting to capture C level outputs as well. These can be traces of Nuitka itself, but also messages from C libraries. On Windows, with MinGW64 this does not work, and it still only captures MinGW64, due to limitations of using different C run-times. With MSVC it works for the compiled program and C, but DLLs can have their own C runtime outputs that are still not caught.
- Added new spec value `%PROGRAM_BASE%` which will avoid the suffix `.exe` or `.bin` of binaries that `%PROGRAM%` will still give.
- Plugins: Added ability to query if a package in an Anaconda package or not, with the new `is_conda_package()` function in Nuitka package configuration. Added in 1.6.1 already.
- Plugins: Provide control tags during plugin startup with new interface, such that these become globally visible.
- Plugins: Allow to give `--include-qt-plugins` options of Qt binding plugins to be given multiple times. This is for consistency with other options. These now expand the list of plugins rather than replacing it.
- Added experimental code to include functions decorated in certain ways to be included as bytecode. Prepare the inclusion as source code in a similar fashion. This was used to make example PyQt5 code work properly with timers where it doesn't normally work, but is still in development before it will be generally useful. For that it reacts to `@pyqtSlot` decorators.
- Plugins: Make anti-bloat not warn when bloating modules include their group. This helps when e.g. `distributed` is going to use `dask`, then we warn about `distributed`, but not anymore, when that then uses `dask`. And that intention to avoid `dask` is now in the warning given for `distributed`.
- Plugins: Added ability to decide module inclusion based on using module name and not only the used name. This will be super useful to make some imports not count per se for inclusion.
- Plugins: Added new `no-auto-follow` Yaml configuration for anti-bloat, that makes imports from one module not automatically included. That can make optional import removal much easier.
- Plugins: Added new function for when clauses, such that it now can be tested if this Python version has a certain built-in name, e.g. when: `'not has_builtin_module("_socket")'` will not apply configuration `_socket` is an extension module rather than built-in. This can be used to avoid unnecessary changes.

Optimization

- Optimization: Better `hasattr` handling. Added ability for generated expression base class to monitor the attribute name for becoming constant and then calling a new abstract method due to `auto_compute_handling` saying `wait_constant:name`.
- Optimization: Added type shapes for `setattr` and `hasattr` built-ins as well as the attribute check node for better code generation.
- Optimization: Added dedicated nodes for `importlib.resources.files` to allow including the used package automatically.
- Standalone: Include only platform DLLs for `tls_client` rather than all DLLs for all platforms. Added in 1.6.1 already.

- Anti-Bloat: Avoid including `sympy.testing` for `sympy` package. Added in 1.6.3 already.
- Anti-Bloat: Avoid `IPython` in `transformers` package. Added in 1.6.3 already.
- Anti-Bloat: Avoid `transformers.testing_util` inclusion for `transformers` package as it will trigger `pytest` inclusion.
- Anti-Bloat: Added missing method to our `numpy.testing` stub, so it can be used with more packages. Added in 1.6.4 already.
- Anti-Bloat: Avoid `numba` usage from parts of `pandas`. Added in 1.6.4 already.
- Anti-Bloat: Avoid `pytest` usage in `patsy` more completely. Added in 1.6.4 already.
- Standalone: Added data files needed for `pycountry` package. Added in 1.6.4 already.
- Anti-Bloat: Avoid `unittest` usage in `numpy` package. Added in 1.6.4 already.
- Anti-Bloat: Avoid using `pytest` in `statsmodels` package. Added in 1.6.4 already.
- Anti-Bloat: Avoid including `PIL.ImageQt` when `no-qt` plugin is used. Added in 1.6.4 already.
- Anti-Bloat: Avoid `IPython` usage in `dask`. We do not cover bloat with `dask` allowed well yet, more like this should be added. Added in 1.6.5 already.
- Anti-Bloat: Avoid `dask` via `distributed` in `fsspec` package. Added in 1.6.5 already.
- Anti-Bloat: Avoid `IPython` in `patsy` package. Added in 1.6.5 already.
- Anti-Bloat: Avoid `setuptools` in newer `torch` as well. Added in 1.6.5 already.
- Anti-Bloat: Avoid `tkinter` inclusion in `PIL` and `matplotlib` if another GUI plugin is active. This is using the control tags made available by GUI plugins.
- Anti-Bloat: Avoid warning for `from unittest import mock` imports. These are common, and not considered actual usage of `unittest` anymore.
- Anti-Bloat: Avoid `pandas` usage in `tqdm`. This uses the new `no-auto-follow` feature that will enable the optional integration of `tqdm` if `pandas` is included by other means only.
- Anti-Bloat: Better method of avoiding `socket` in `email.utils`. With changing the source code to delay the import of `socket` to the only function using it. `Socket` is now included only if used elsewhere. These changes however, are only done if `_socket` if is not a built-in module, because only then they really matters. And using a simple `--include-module=socket` will restore this. This approach is more robust and less invasive.

Organisational

- Added `run-inside-nuitka-container` for use in CI scripts. With this, dependencies of package building and testing from correct system installation should go away.
- Release: Add CI container for use with `run-inside-nuitka-container` to make Debian package releases. This provides a more stable and flexible environment rather than building through ansible maintained environments, since different branches can more easily use different versions, or new features for the container handling.
- Release: Use upload tokens rather than PyPI password in uploads, and secure the account with 2FA.
- UI: Avoid duplicate warnings for anti-bloat detected imports. In case of `from unittest import mock` there were 2 warnings given, for `unittest` and `unittest.mock` but that is superfluous.
- macOS: More beginner friendly version of Apple Python standalone error. They won't know why it is, and where to get a working Python version, so we explain more and added a download link.
- Scons: Consider only 5 minutes slow for a module compilation in backend. Many machines are busy or slow by nature, so don't warn that much.

- GitHub: Actions no longer work (easily) with Python2, so we removed those and need to test it elsewhere.
- UI: Output the filename of the XML node dump from `--xml` as well.
- UI: Make `--edit-module-code` work with onefile outputs as well.
- Debugging: Allow yaml condition traceback to go through in `--debug` mode, so exception causes are visible.
- Plugins: Make more clear what is the forbidden module user, such that it is possible to debug it.
- UI: Inform user about slow linking, and `--lto=no` choice in case `auto` was used. This should make this option more obvious for new users that somehow victim of not defaulting to `no`, but still having a slow link.
- Debugging: Include PDBs for DLLs in unstripped mode already. Previously this was only done for debug mode, but that's a bit high of a requirement, and we sometimes need to debug where things do not happen in debug mode.
- User Manual: Added typical problem with `python -m compiled_module` execution not working and why that is so.
- Debian: Do not include PDF files in packages. These are probably not used that much, but they cause issues at times, that are likely not worth the effort.

Cleanups

- Moved OS error reporting as done in onefile binary to common code for easier reuse in plugins.
- Moved helper codes for expanding paths and for getting the path to the running executable to file path common code for clearer code structure.
- Removed `x-bits` from files that do not need them. For `__main__` files, they are not needed, and for some files they were outright wrong.
- Python3.12: Avoid usage of `distutils.utils` which were using to disable bytecode compilation for things we expect to not work.
- Solve TODO and use more modern git command `git branch --show-current` to detect branch, our CI will have this for sure.
- In our Yaml configuration prefer the GUI toolkit control tags, e.g. `use_pyside6` over the `plugin("pyside6")` method.

Tests

- Release: Use CI container for linter checks, so different branches can use different versions with less pain involved.
- macOS: Allow all system library frameworks to be used, not just a few selected ones, there is many of them and they should all exist on every system. Added in 1.6.1 already.
- Made the `pendulum` test actually useful to cover new and old `pendulum` actually working properly.

Summary

This release really polished `anti-bloat` to the point where we now have all the tools needed. Also `torch` in newest version is now working nicely again with it, and a few rough edges of what we did with 1.6 for not including extension modules were removed. This polishing will go on, but has reached really high levels. More and more people are capable of helping with PRs here.

The optimization work outside of `anti-bloat` was really minor, with only the two attribute built-in nodes being worked on, and only `hasattr` seeing real improvements. However, this was more of a structural

thing. The `wait_constant` technique will not get applied more often, but it also will need a `wait_all_constant` companion, before we can expect scalability improvements.

Restoring Windows 7 is important to many people deploying to old systems, and the like.

However, in the coming release, we need to attack loop tracing. The only bugs currently remaining are related to wrong tracing of items, and it also is a limitation for hard imports to work. So scalability from doing more of the `wait_constant` work, and from more clever loop tracing shall be the focus of the 1.8 release.

Nuitka Release 1.6

This release bumps the much awaited 3.11 support to full level. This means Nuitka is now expected to behave identical to CPython3.11 for the largest part.

There is plenty of new features in Nuitka, e.g. a new testing approach with reproducible compilation reports, support for including the metadata if an distribution, and more.

In terms of bug fixes, it's also huge, and esp. macOS got a lot of improvements that solve issues with prominent packages in our dependency detection. And then for PySide we found a corruption issue, that got workarounds.

Bug Fixes

- The new dict in optimization was compile time crashing on code where the dictionary shaped value checked for a key was actually an conditional expression

```
# Was crashing
"value" in ``some_dict if condition else other_dict``
```

Fixed in 1.5.1 already.

- Standalone: Added support for `openvino`. This also required to make sure to keep used DLLs and their dependencies in the same folder. Before they were put on the top level. Fixed in 1.5.1 already.
- Android: Convert `RPATH` to `RUNPATH` such that standalone binaries need no `LD_LIBRARY_PATH` guidance anymore. Fixed in 1.5.1 already.
- Standalone: Added support for newer `skimage`. Fixed in 1.5.1 already.
- Standalone: Fix, new data file type `.json` needed to be added to the list of extensions used for the Qt plugin bindings. Fixed in 1.5.2 already.
- Standalone: Fix, the `nuitka_types_patch` module using during startup was released, which can have bad effects. Fixed in 1.5.2 already.
- Android: More reliable detection of the Android based Python Flavor. Fixed in 1.5.2 already.
- Standalone: Added data files for `pytorch_lightning` and `lightning_fabric` packages. Added in 1.5.2 already.
- Windows: Fix, the preservation of `PATH` didn't work on systems where this could lead to encoding issues due to reading a MBCS value and writing it as a unicode string. We now read and write the environment value as `unicode` both. Fixed in 1.5.3 already.
- Plugins: Fix, the `scons` report values were not available in case of removed `--remove-output` deleting it before use. It is now read in case if will be used. Fixed in 1.5.3 already.
- Python3.11: Added support for `ExceptionGroup` built-in type. Fixed in 1.5.4 already.
- Anaconda: Fix, using `numpy` in a virtualenv and not from conda package was crashing. Fixed in 1.5.4 already.

- Standalone: Added support for `setuptools`. Due to the anti-bloat work, we didn't notice that if that was not sufficiently usable, the compiled result was not usable. Fixed in 1.5.4 already.
- Distutils: Added support for `pyproject` with `src` folders. This supports now `tool.setuptools.packages.find` with a `where` value with `pyproject` files, where it typically is used like this:

```
[tool.setuptools.packages.find]
where = ["src"]
```

- Windows: Fix, the `nuitka-run` batch file was not working. Fixed in 1.5.4 already.
- Standalone: Add `pymoo` implicit dependencies. Fixed in 1.5.5 already.
- macOS: Avoid deprecated API, this should fix newer Xcode being used. Fixed in 1.5.5 already.
- Fix, the multiprocessing in spawn mode didn't handle relative paths that become invalid after process start. Fixed in 1.5.5 already.
- Fix, spec `%CACHE_DIR%` was not given the correct folder on non-Windows. Fixed in 1.5.5 already.
- Fix, special float values like `nan` and `inf` didn't properly generate code for C values. Fixed in 1.5.5 already.
- Standalone: Add missing DLL for `onnxruntime` on Linux too. Fixed in 1.5.5 already.
- UI: Fix, illegal python flags value could enable `site` mode. by mistake and were not caught. Fixed in 1.5.6 already.
- Windows: Fix, user names with spaces failed with MinGW64 during linking. Fixed in 1.5.6 already.
- Linux: Fix, was not excluding all libraries from `glibc`, which could cause crashes on newer systems. Fixed in 1.5.6 already.
- Windows: Fix, could still pickup SxS libraries distributed by other software when found in `PATH`. Fixed in 1.5.6 already.
- Windows: Fix, do not use cache DLL dependencies if one the files listed there went missing. Fixed in 1.5.6 already.
- Onefile: Reject path spec that points to a system folder. We do not want to delete those when cleaning up clearly. Added in 1.5.6 already.
- Plugins: Fix, the `dill-compat` was broken by code object changes. Fixed in 1.5.6 already.
- Standalone: Added workaround for `networkx` decorator issues. Fixed in 1.5.7 already.
- Standalone: Added workaround for `PySide6` problem with disconnecting signals from methods. Fixed in 1.5.7 already.
- Standalone: Added workaround for `PySide2` problem with disconnecting signals.
- Fix, need to make sure the `yaml` package is located absolutely or else case insensitive file systems can confuse things. Fixed in 1.5.7 already.
- Standalone: Fix, extra scan paths were not considered in caching of module imports, breaking the feature in many cases. Fixed in 1.5.7 already.
- Windows: Fix, avoid system installed `appdirs` package as it is frequently broken. Fixed in 1.5.7 already.
- Standalone: The bytecode cache check needs to handle re-checking relative imports found in the cache better. Otherwise some standard library modules were always recompiled due to apparent import changes. Fixed in 1.5.7 already.
- Nuitka-Python: Fix, do not insist on `PYTHONHOME` making it to `os.environ` in order to delete it again. Fixed in 1.5.7 already.

- Nuitka-Python: Allow builtin modules of all names. This is of course what it does. Fixed in 1.5.7 already.
- Nuitka-Python: Ignore empty extension module suffix. Was confusing Nuitka to consider every file an extension module potentially. Fixed in 1.5.7 already.
- Plugins: Properly merge code coming from distinct plugins. The `__future__` imports need to be moved to the start. Added in 1.5.7 already.
- Standalone: Added support for `opentele` package. Fixed in 1.5.7 already.
- Standalone: Added support for newer `pandas` and `pyarrow` usage. Fixed in 1.5.7 already.
- Standalone: Added missing implicit dependency for PySide6. Fixed in 1.5.7 already.
- Fix, the `pyi`-file parser didn't handle doc strings, and could be crash for comment contents not conforming to be import statement code. Fixed in 1.5.8 already.
- Standalone: Added support for `pyqtlet2` data files.
- Python2: Fix, `PermissionError` doesn't exist on that version, which could lead to issues with retries for locked files e.g. but was also observed with symlinks.
- Plugins: Recognize the error given by `upx` if a file is already compressed.
- Fix, so called "fixed" imports were not properly tracking their use, such that they then didn't show up in reports, and didn't cause dependencies on the module, which could e.g. impact `importlib` to not be included even if still being used.
- Windows: Fix, retries for payload attachment were crashing when maximum number of retries were reached. Using the common code for retries solves that, since that code handles it just fine.
- Standalone: Added support for the `av` module.
- Distutils: Fix, should build from files in `build` folder rather than `source` files. This allows tools like `versioneer` that integrate with `setuptools` to do their thing, and get the result of that to compilation rather than the original source files.
- Standalone: Added support for the `Equation` module.
- Windows/macOS: Avoid problems with case insensitive file systems. The `nuitka.Constants` module and `nuitka.constants` package could collide, so we now avoid that package, there was only what is now `nuitka.Serialization` in there anyway. Also similar problem with `nuitka.utils.Json` and `json` standard library module.
- Standalone: Added support `transformers` package.
- Standalone: Fix for `PyQt5` which needs a directory to exist.
- macOS: Fix, was crashing with `PyQt6` in standalone mode when trying to register plugins to non-default path. We now try to skip the need, which also makes it work.
- Fix, recursion error for complex code that doesn't happen in `ast` module, but during conversion of the node tree it gives to our own tree, were not handled, and crashed with `RecursionError`. This is now also handled, just like the error from `ast`.
- Standalone: Added support for `sqlfluff`.
- Standalone: Added support for PySide 6.5 on macOS solving DLL dependency issues.
- Scons: Recognize more `ccache` outputs properly, their logging changed and provided irrelevant states, and ones not associated so far.
- Onefile: Fix, could do random exit codes when failing to fork for whatever reason.
- Standalone: Added support for `pysnmp` package.
- Standalone: Added support for `torchaudio` and `tensorflow` on macOS. These contain broken DLL dependencies as relative paths, that are apparently ignored by macOS, so we do that too now.

- Onefile: Use actual rather than guessed standalone binary name for `multiprocessing` spawns. Without this, a renamed onefile binary, didn't work.
- Fix, side effect nodes, that are typically created when an expression raises, were use in optimization contexts, where they do not work.
- Standalone: Added missing implicit dependency for `sentence_transformers` package.
- macOS: Fix, added missing dependency for `platform` module.

New Features

- Support for Python 3.11 is finally there. This took very long, because there were way more core changes than with previous releases. Nuitka integrates close to that core, and is as such very affected by this. Also a lot of missed opportunities to improve 3.7 or higher, 3.9 or higher, and 3.10 or higher were implemented right away, as they were discovered on the way. Those had core changes not yet taken advantage of and as a result got faster with Nuitka too.
- Reports: Added option `--report-diffable` to make the XML report created with `--report` become usable for comparison across different machine installations, users compiling, etc. so it can be used to compare versions of Nuitka and versions of packages being compiled for changes. Also avoid short names in reports, and resolve them back to long names, so they become more portable too.
- Reports: Added option to provide custom data from the user. We use it in our testing to record the pipenv state used with things like `--report-user-provided=pipenv-lock-hash=64a5e4` with this data ending up inside of reports, where tools like the new testing tool `nuitka-watch` can use it to decide if upstream packages changed or not. These are free form, just needs to fit XML rules.
- Plugins: Added `include-pyi-file` flag to data-files section. If provided, the `.pyi` file belonging to a specific module is included. Some packages, e.g. `skimage` depend at runtime on them. For data file options and configuration, these files are excluded, but this is now the way to force their inclusion. Added in 1.5.1 already.
- Compatibility: Added support for including distribution metadata with new option `--include-distribution-metadata`.
This allows generic walks over distributions and their entry points to succeed, as well as version checks with the metadata packages that are not compile time optimized.
- Distutils: Handle extension modules in build tasks. Also recognize if we built it ourselves, in which case we remove it for rebuild. Added in 1.5.7 already.
- Linux: Detect DLL like filenames that are Python extension modules, and ignore them when listing DLLs of a package with `--list-package-dlls` option. So far, this was a manual task to figure out actual DLLs. This will of course improve the Yaml package configuration tooling .
- Onefile: Allow forcing to use no compression for the onefile payload, useful for debugging, to avoid long compression times and for test coverage of the rare case of not compressing if the bootstrap handles that correctly too.
- Need to resolve symlinks that were used to call the application binary in some places on macOS at least. We therefore implemented the previously experimental and Windows only feature for all platforms.
- Standalone: Added support including symlinks on non-Windows in standalone distribution, if they still point to a path that is inside the distribution. This can save a bunch of disk space used for some packages that e.g. distribute DLL links on Linux.
- Onefile: Added support for including symlinks from the standalone distribution as such on non-Windows. Previously they were resolved to complete copies.

- UI: Respect code suffixes in package data patterns. With this e.g. `--include-package-data=package_name:*.py` is doing what you say, even if of course, that might not be working.
- UI: Added option `--edit-module-code` option.
To avoid manually locating code to open it in Visual Code replaced old `find-module` helper to be a main Nuitka option, where it is more accessible. This also goes beyond it, such that it resolves standalone file paths to module names to make debugging easier, and that it opens the file right away.
- Standalone: Added support for handling missing DLLs. Needed for macOS PySide6.5.0 from PyPI, which contains DLL references that are broken. With this feature, we can exclude DLLs that wouldn't work anyway.

Optimization

- Anti-Bloat: Remove `IPython` usage in `huggingface_hub` package versions. Added in 1.5.2 already.
- Anti-Bloat: Avoid `IPython` usage in `tokenizers` module. Added in 1.5.4 already.
- Added support for module type as a constant value. We want to add all types we have shapes for to allow better `type(x)` optimization. This is only the start.
- Onefile: During payload unpacking the memory mapped data was copied to an input buffer. Removing that avoids memory copying and reduces usage.
- Onefile: Avoid repeated directory creations. Without it, the bootstrap was creating already existing directories up to the root over and over, making many unnecessary file system checks. Added in 1.5.5 already.
- Anti-Bloat: Remove usage of `IPython` in `trio` package. Added in 1.5.6 already.
- Onefile: Use resource for payload on Win32 rather than overlay. This integrates better with signatures, removing the need to check for original file size. Changed in 1.5.6 already.
- Onefile: Avoid using `zstd` input buffer, but using the memory mapped contents directly avoiding to copy uncompressed payload data. Changed in 1.5.6 already.
- Onefile: Avoid double slashes in expanded onefile temp spec paths, they are just ugly.
- Anti-Bloat: Remove usage of `pytest` and `IPython` for some packages used by newer `torch`. Added in 1.5.7 already.
- Anti-Bloat: Avoid `triton` to use `setuptools`. Added in 1.5.7 already.
- Anti-Bloat: Avoid `pytest` in newer `networkx` package. Added in 1.5.7 already.
- Prepare optimization for more built-in types with experimental code, but we need to disable it for now as it requires more completeness in code generation to cover them all. We did some, e.g. module type, but many more will be missing still.
- Prepare optimization of class selection at compile time, by having a helper function rather than a dedicated node. This work is not complete though, and cannot be activated yet.
- Windows: Cache short path name resolutions. Esp. for reporting, we now do a lot more of these than before, and this avoids they can become too time consuming.
- Faster constant value handling for float value checks by avoiding module lookups per value.
- Minimize size for hello world distribution such that no unused extension modules are included, by excluding even more modules and using modules from automatic inclusion of standard library.
- Anti-Bloat: Catch `pytest` namespaces `py` and `_pytest` sooner, to point to the actual uses more directly.

- Anti-Bloat: Usage of `doctest` equals usage of "unittest" so cover it too, to point to the actual uses more directly.
- Ever more spelling fixes in code and tests were identified and fixed.
- Make sure side effect nodes indicate properly that they are raising, allowing exceptions to fully bubble up. This should lead to more dead code being recognized as such.

Organisational

- GitHub: Added marketplace action designed to cross platform build with Nuitka on GitHub directly. Usable with both standard and commercial Nuitka versions, and pronouncing it as officially supported.

Check out out at [Nuitka-Action](#) repository.

- Windows: When MSVC doesn't have WindowsSDK, just don't use it, and proceed, to e.g. allow fallback to winlibs gcc.
- User Manual: The code to update benchmark numbers as giving was actually wrong. Fixed in 1.5.1 already.
- UI: Make it clear that partially supported versions are considered experimental, not unsupported. Fixed in 1.5.2 already.
- Plugins: Do not list deprecated plugins with `plugin-list`, they do not have any effect, but listing them, makes people use them still. Fixed in 1.5.4 already.
- Plugins: Make sure all plugins have descriptions. Some didn't have any yet, and sometimes the wording was improved. Fixed in 1.5.4 already.
- UI: Accept `y` as a shortcut for `yes` in prompts. Added in 1.5.5 already.
- Reports: Make sure the DLL dependencies for Linux are in a stable order. Added in 1.5.6 already.
- Plugins: Check for latest fixes in PySide6. Added in 1.5.6 already.
- Windows XP: For Python3.4 make using Python2 scones work again, we cannot have 3.5 or higher there. Added in 1.5.6 already.
- Quality: Updated to latest PyLint. With Python 3.11 the older one, was not really working, and it was about time. Due to its many changes, we included it in the hotfix, so those can still be done. Changed in 1.5.7 already.
- Release: Avoid broken `requires.txt` in source distribution. This apparently breaks poetry. Changed in 1.5.7 already.
- GitHub: Enhanced issue template for more clarity, esp. to avoid unnecessary options, e.g. using `--onefile` for issues that show up with `--standalone` already, to report factory branch issues rather on Discord, and give a quick tip for a likely reproducer if a package fails to import.
- User Manual: Added instructions on how to add a DLL or executable to a standalone distribution.
- User Manual: Example paths in the table for path specs, meant for Windows were not properly escaping the backslashes and therefore rendered incorrectly.
- Visual Code: Python3.11 is now the default configuration for C code editing.
- Developer Manual: Updated descriptions for adding test suite. While added the Python 3.11 test suite, these instructions were further improved.
- Debugging: Make it easier to fully deactivate free lists. Now only need to set max size to 0 and the free list will not be used.
- Debugging: Added more assertions, added corrections to feature disables, check args after function calls for validity, check more types to be as expected.

- Plugins: Enhanced plugin error messages generally, with `--debug` exceptions become warning messages with the original exception being raised instead, making debugging during development much easier.
- UI: Make it clear what not using `ccache` actually means. Not everybody is familiar with the design of Nuitka there or what the tool can actually do.
- UI: Do not warn about not found distributions but merely inform of them.

Since Nuitka is fully compatible with these, no need to consider those a warning, for some packages they also are given really a lot.

- UI: Catch user error of wrong cases plugin names

This now points out the proper name rather than denying the existence outright. We do not want to accept wrong case names silently.

Cleanups

- Use proper API for setting `PyConfig` values during interpreter initialization. There is otherwise always the risk of crashes, should these values change during runtime. Fixed in 1.5.2 already.
- For our reformulations have a helper function that build release statements for multiple variables at once. This removed a bunch of repetitive code from re-formulations.
- Move the pyi-file parser code out of the module nodes and to source handling, where it is more closely related.

Tests

- Adding a `nuitka-watch` tool, which is still experimental and for use with the [Nuitka-Watch](#) repository.
- Refined macOS standalone exceptions further to cover more normal usages of files on that OS and for frameworks that applications typically use from the system.
- Detect and consider onefile mode if given in project options as well.
- Was not really applying import check in programs tests. Added in 1.5.6 already.
- Added coverage of testing the signing of Windows binaries with the commercial plugin.
- Added coverage of version information to hello world onefile test, so we can use it for Virus tools checks.
- Added tests to cover PyQt6 and PySide6 plugin availability, we so far only had that for PyQt5, which is of course not relevant, and totally different code anyway.
- Cleanup distutils tests case to use common test case scanning. We now decide version skips based on names, and had to get away from number suffixes, so they are now in the middle.

Summary

The class bodies optimization has made some progress in this release, going to a re-formulation of the metaclass selection, so as to allow its future optimization. We are not yet at "compiled objects", but this is a promising road. We need to make some optimization improvements for inlining constant value calls, then this can become really important, but by itself these changes do not yield a lot of improvement.

For macOS again a bunch of time was spent to improve and complete the detection of DLL dependencies. More corner cases are covered now and more packages just work fine as a result.

The most important is to become Python3.11 compatible, even if attribute lookups, and other things, and not yet optimized. We will get to that in future releases. For now, compatibility is the first step to take.

For GitHub users, the Nuitka-Action will be interesting. But it's still in develop. We keep adding missing options of Nuitka for a while it seems, but for most people it should be usable already.

The new `nuitka-watch` ability, should allow us to detect breaking PyPI releases, that need a new tweak in Nuitka sooner. But it will probably grow in the coming releases to full value only. For now the tool itself is not yet finished.

From here, a few open ends in the CPython 3.11 test suite will have to be addressed, and maybe some of the performance tricks that it now will enable, e.g. with repeated attribute lookups.

Nuitka Release 1.5

This release contains the long awaited 3.11 support, even if only on an experimental level. This means where 3.10 code is used, it is expected to work equally well, but the Python 3.11 specific new features have yet been done.

There is plenty of new features in Nuitka, e.g. much enhanced reports, Windows ARM native compilation support, and the usual slew of anti-bloat updates, and newly supported packages.

Bug Fixes

- Standalone: Added implicit dependencies for `charset_normalizer` package. Fixed in 1.4.1 already.
- Standalone: Added platform DLLs for `sounddevice` package. Fixed in 1.4.1 already.
- Plugins: The info from Qt bindings about other Qt bindings being suppressed for import, was spawning multiple lines, breaking tests. Merged to a single line until we do text wrap for info messages as well. Fixed in 1.4.1 already.
- Plugins: Fix `removeDllDependencies` was broken and could not longer be used to remove DLLs from inclusion. Fixed in 1.4.1 already.
- Fix, assigning methods of lists and calling them that way could crash at runtime. The same was true of dict methods, but had never been observed. Fixed in 1.4.2 already.
- Standalone: Added DLL dependencies for `onnxruntime`. Fixed in 1.4.2 already.
- Standalone: Added implicit dependencies for `textual` package. Fixed in 1.4.2 already.
- Fix, boolean tests of lists could be optimized to wrong result when list methods got recognized, due to not annotating the escape during that pass properly. Fixed in 1.4.3 already.
- Standalone: Added missing implicit dependency of `apsw`. Fixed in 1.4.3 already.

Note

Currently `apsw` only works with manual workarounds and only in limited ways, there is an import level incompatible with `__init__` being an extension module, that Nuitka does not yet handle.

- Python3: Fix, for range arguments that fail to divide there difference, the code would have crashed. Fixed in 1.4.3 already.
- Standalone: Fix, added support for newer `pkg_resources` with another vendored package. Fixed in 1.4.4 already.
- Standalone: Fix, added support for newer `shapely` 2.0 versions. Fixed in 1.4.4 already.

- Plugins: Fix, some yaml package configurations with DLLs by code didn't work anymore, notably old `shapely 1.7.x` versions were affected. Fixed in 1.4.4 already.
- Fix, for onefile final result the `--output-dir` option was ignored. Fixed in 1.4.4 already.
- Standalone: Added `mozilla-ca` package data file. Fixed in 1.4.4 already.
- Standalone: Fix, added missing implicit dependency for newer `gevent`. Fixed in 1.4.4 already.
- Scons: Accept an installed Python 3.11 for Scons execution as well. Fixed in 1.4.4 already.
- Python3.7: Some `importlib.resource` nodes asserted against use in 3.7, expecting it to be 3.8 or higher, but this interface is present in 3.7 already. Fixed in 1.4.5 already.
- Standalone: Fix, Python DLLs installed to the Windows system folder were not included, causing the result to be not portable. Fixed in 1.4.5 already.
- Python3.9+: Fix, `metadata.resources` files method `joinpath` in some contexts is expected to accept variable number of arguments. Fixed in 1.4.5 already.
- Standalone: Workaround for `customtkinter` data files on non-Windows. Fixed in 1.4.5 already.
- Standalone: Added support for `overrides` package. Fixed in 1.4.6 already.
- Standalone: Added data files for `strawberry` package. Fixed in 1.4.7 already.
- Fix, anti-bloat plugin caused crashes when attempting to warn about packages coming from `--include-package` by the user. Fixed in 1.4.7 already.
- Windows: Fix, main program filenames with an extra dot apart from the `.py` suffix, had the part beyond that wrongly trimmed. Fixed in 1.4.7 already.
- Fix, list methods didn't properly annotated value escape during their optimization, which could lead to wrong optimization for boolean tests. Fixed in 1.4.7 already.
- Standalone: Added support for `imagej`, `scyjava`, `jpy` packages. Fixed in 1.4.8 already.
- Fix, using `--include-package` on extension module names was not working. Fixed in 1.4.8 already.
- Standalone: Added support for `tensorflow.keras` namespace as well.
- Distutils: Fix namespace packages were not including their contained modules properly with regards to `__file__` properties, making relative file access impossible.
- Onefile: On Windows the onefile binary did lock itself, which could fail with certain types of AV software. This is now avoided.
- Accessing files using the top level `metadata.resources` files object was not working properly, this is now supported too.
- MSYS2: Make sure mixing POSIX and Windows slashes causes no issues by hard-coding the onefile archive to use the subsystem slash rather than what MSYS prefers to use internally.
- Standalone: Added missing dependencies of newer `imageio`.
- Fix, side effect nodes didn't annotate their non-exception raising nature properly, if that was the case.

New Features

- Added experimental support for Python 3.11, for 3.10 language level code it should be fully usable, but the `CPython311` test suite has not even been started to check newly added or changed features.
- Windows: Support for native Python on Windows ARM64, which needs 3.11 or higher, but standalone and therefore onefile do not yet work, due to lack of any form of binary dependency analysis tool.

This platform is relatively new in Python and generally. For the time being standalone and onefile should be done with Intel based Python, they would also be ARM64 only, whereas 32/64 Bit binaries can be run on all Windows ARM platforms.

- Reports: Write compilation report even in case of Nuitka being interrupted or crashing. This then includes the exception, and a status like `completed` or `interrupted`. At this time this happens only when `--report=` was specified, but in the future we will likely write one in case of Nuitka crashes.
- Reports: Now the details of the used Python version, its flavor, the OS and the architecture are included. This is crucial information for analysis and can make `--version` output unnecessary.
- Reports: License reports now handle `UNKNOWN` license by falling back to checking the classifiers, and therefore include the correct license e.g. with `setuptools`. Also in case no license text is found, do not create an empty block. Added in 1.4.4 already.
- Reports: In case the distribution name and the contained package names differ, output the list of packages included from a distribution. Added in 1.4.4 already.
- Reports: Include data file sizes in report. Added in 1.4.7 already.
- Reports: Include memory usage into the compilation report as well.
- macOS: Add support for downloading `ccache` on arm64 (M1/M2) too. Added in 1.4.4 already.
- UI: Allow `--output-filename` for standalone mode again. Added in 1.4.3 already.
- Standalone: Improved isolation with Python 3.8 or higher. Using new init mechanisms of Python, we now achieve that the scan for `pyvenv.cfg` on in current directory and above is not done, using it will be unwanted.
- Python2: Expose `__loader__` for modules and register with `pkg_resources` too which expects these to be present for custom resource handling.
- Python3.9+: The `metadata.resources` files objects method `iterdir` was not implemented yet. Fixed in 1.4.5 already.
- Python3.9+: The `metadata.resources` files objects method `absolute` was not implemented yet.
- Added experimental ability to create virtualenv from an existing compilation report with new `--create-environment-from-report` option. It attempts to create a requirements file with the used packages and their versions. However, sometimes it seems not to be possible to due to conflicts.

Optimization

- Onefile: Use memory mapping for calculating the checksum of files on all platforms. This is faster and simpler code. So far it had only be done this way on Windows, but other platforms also benefit a lot from it.
- Onefile: Use memory mapping for accessing the payload rather than file operations. This avoids differences to macOS payload handling and is much faster too.
- Anti-Bloat: Avoid using `dask` in `joblib`.

Note

Newer versions of `joblib` do not currently work yet due to their own form of multiprocessing spawn not being supported yet.

- Anti-Bloat: Adapt for newer `pandas` package.

- Anti-Bloat: Remove more `IPython` usages in newer tensorflow.
- Use dedicated class bodies for Python2 and Python3, with the former has a static dict type shape, and with Python3 this needs to be traced in order to tell what the meta class put in there.
- Compile time optimize dict `in/not in` and `dict.has_key` operations statically where the keys of a dict are known. As a result, the class declarations of Python3 no longer created code for both branches, the one with `metaclass =` in the class declaration and without. That means also a big scalability improvement.
- For the Python3 class bodies, the usage of `locals()` was not recognized as not locally escaping all the variables, leading to variable traces where each class variable was marked as escaped for no good reason.
- Added support for `dict.fromkeys` method, making the code generation understand and handle static methods as well.
- Added support for `os.listdir` and `os.path.basename`. Added in 1.4.5 already for use in implementing the `iterdir` method, but they are also now optimized by themselves.
- Added support for trusted constant values of the `os` module. These are `curdir`, `pardir`, `sep`, `extsep`, `altsep`, `pathsep`, `linesep` which may enable some minor compile time optimization to happen and completes this aspect of the `os` module.
- Faster digit size checks during float code generation for better compile time performance.
- Faster list operations due to using `PyList_CheckExact` everywhere this is applicable, this mostly makes debug operations faster, but also deep copying list values, or extending lists with iterables, etc.
- Optimization: Collect module usages of the given module during its abstract execution. This avoids a full tree visit afterwards only to find them. It is much cheaper to collect them while we go over the tree. This enhances the scalability of large compilations by ca. 5%.
- Optimization: Faster determination of loop variables. Rather than using a generic visitor, we use the children having generator codes to add traversal code that emits relevant variables to the user directly.
- Cache extra search paths in order to avoid repeated directory operations as these are known to be slow at times.
- Standalone: Do not include `py.typed` data files, these indicator files are for IDEs, but not needed at run time ever.
- Make sure that the generic attribute code optimization is also effective in cases where a Python DLL is used. Previously this was only guaranteed to be used with static libpython.
- Faster list constant usage

Small immutable constants get their own code that is much faster for small sizes.

Medium sized lists get code that just is hinted the size, but takes items from a source list, still a lot faster.

For repeated lists where all elements are the same, we use a dedicated helper for all sizes, that is even faster except for small ones with LTO enabled, where the C compiler may already do that effectively.

- Added optimization for `os.path.abspath` and `os.path.isabs` which of course have not as much potential for compile time optimization, but we needed them for providing `.absolute()` for the meta path loader files implementation.
- Faster class dictionary propagation decision. Instead of checking for trace types, let the trace object decide. Also abort immediately on first inhibit, rather than checking all variables. This improves Python2 compile time, and Python3 where this code is now starting to get used when the class dictionary is shown to have dict type.

- Specialize type method `__prepare__` which is used in the Python3 re-formation of class bodies to initialize the class dictionary. Where the metaclass is resolved, we can use this to decide that the standard empty dictionary is used statically, enabling class dictionary propagation for best scalability.

At this time this only happens with classes without bases, but we expect to soon do this with all compile time known base classes. At this time, these optimization to become effective, we need to optimize meta class selection from bases classes, as well as modification of base classes with `__mro_entries__` methods.

- The `bool` built-in on boolean values is now optimized away.

Since it's used also for conditions being extracted, this is actually somewhat relevant, since it could keep code alive in side effects at least for no good reason and this allows a proper reduction.

Organisational

- Project: Require the useful stuff for installation of Nuitka already. These are things we cannot inline really, but otherwise will frequently be warned about, e.g. `zstandard` for onefile and `ordered-set` for fast operation, but we do not require packages that might fail to install.
- User Manual: Added section about virus scanners and how to avoid false reports.
- User Manual: Enhanced description for plugin module loading, the old code was too complicated and actually working only for a mode of including plugin code that is discouraged.
- User Manual: Fix section for standalone finding files on wrong level.

- Windows: Using the console on Python 3.4 to 3.7 is not working very well with e.g. many Asian systems. Nuitka fails to setup the encoding for stdin and stdout on this platform. It can then produce exceptions on input or output of unicode data, that doesn't overlap with UTF-8.

We now inform the user of these older Python with a warning and mnemonic, to either disable the console or to upgrade to Python 3.8 or higher, which normally won't be much of an issue for most users. Added in 1.4.1 already.

- Debugging: Fixup debugging reference count output with Python3.4. For Python 3.11 compatibility tests, actually it was useful to compare with a version that doesn't have coroutines yet. Never tell me, supporting old versions is not good.
- Deprecating support for Python 3.3, there is no apparent use of this version, and it has gained specific bugs, that are indeed not worth our time. Python 2.6 and Python 2.7 will continue to be supported probably indefinitely.
- Recommend `ordered-set` for Python 3.7 to 3.9 as well, as not only for 3.10+ because on Windows, to install `orderset` MSVC needs to be installed, whereas `ordered-set` has a wheel for ready use.
- Actually `zstandard` requirement is for a minimal version, added that to the requirement files.
- Debugging: Lets not reexecute Nuitka in case if we are debugging it from Visual Code.
- Debugging: Include the `.pdb` files in Windows standalone mode for proper C tracebacks should that be necessary.
- UI: Detect the GitHub flavor of Python as well.
- Quality: Check the `clang-format` version to avoid older ones with bugs that made it switch whitespace for one file. Using the one from Visual Code C extension is a good idea, since it will often be available. Running the checks on newer Ubuntu GitHub Actions runner to have the correct version available.
- Quality: Updated the version of `rstfmt` and `isort` to the latest versions.
- GitHub: Added commented out section for enabling `ssh` login, which we occasionally need to git bisect problems specific to GitHub Python flavor.

- Plugins: Report problematic plugin name with module name or DLL name when these raise exceptions.
- Use `ordered-set` package for Python3.7+ rather than only Python3.10+ because it doesn't need any build dependency on Windows.
- UI: When showing source changes, also display the module name with the changed code.
- UI: Use function intended for user query when asking about downloads too.
- UI: Do not report usage of `ccache` for linking from newer version, that is not relevant.
- Onefile: Make sure we have proper error codes when reporting IO errors.
- MSVC: Detect a version for developer prompts too. This version is needed for use in enabling version specific features.
- Started UML diagrams with `plantuml` that will need to be completed before using them in then new and more visual parts of Nuitka documentation.
- UI: Check icon conversion capability at start of compilation rather than error exiting at the very end informing the user about required `imageio` packages to convert to native icons.
- Quality: Enhanced autoformat on Windows, which was susceptible to tools introducing Windows new lines before other steps were performed, that then could be confused, also enforcing use of UTF-8 encoding when working with Nuitka source code for formatting.

Cleanups

- The `delvewheel` plugin was still using a `zmq` class name from its original implementation, adapted that.
- Use common template for generator frames as well. This made them also work with 3.11, by avoiding duplication.
- Applied code formatting to many more files in `tests`, etc.
- Removed a few micro benchmarks that are instead to be covered by construct based tests now.
- Enhanced code generation for specialized in-place operations to avoid unused code for operations that do not have any shortcuts where the operation would be actual in-place of a reference count 1 object.
- Better code generation for module variable in-place operations with proper indentation and no repeated calls.
- Plugins: Use the `namedtuple` factory that we created for informational tuples from plugins as well.
- Make details of download utils module more accessible for better reuse.
- Remove last remaining Python 3.2 version check in C code, for us this is just Python3 with 3.2 being unsupported.
- Cleanup, name generated call helper file properly, indicating that it is a generated file.

Tests

- Made the CPython3.10 test suite largely executable with Python 3.11 and running that with CI now.
- Allow measuring constructs without writing the code diff again. Was crashing when no filename was given.
- Make Python3.11 test execution recognized by generally accepting partially supported versions to execute the tests with.
- Handle also `newfstat` directory checks in file usage scan. This are used on newer Linux systems.

- GitHub: In actions use `--report` for coverage and upload the reports as artifacts.
- Use `no-qt` plugin to avoid warnings in `matplotlib` test rather than disabling the warnings about Qt bindings.
- macOS: Detect if the machine can take runtime traces, which on Apple Silicon by default it cannot.
- macOS: Cover all APIs for file tracing, rather than just one for extended coverage.
- Fix, `distutils` test was not installing the built wheels, but source archive and therefore compiling that second time.
- For the `pyproject.toml` using tests, Nuitka was always downloaded from PyPI rather than using the version under test.
- Ignore `ld` info output about mismatching architecture libraries being ignored. Fixed in 1.4.1 already.

Summary

With this release an important new avenue for scalability has been started. While for Python2 class bodies were very often reduced to just that dictionary creation, with Python3 that was not the case, due to the many new complexities, and while this release makes a start, we will be able to continue this path towards much more scalable class creation codes. And while the performance does not really matter all that much for these, knowing these, will ultimately lead us to "compiled classes" as our own type, and "compiled objects" that may well perform much faster.

Already now, the enhancements to class creation codes will result in smaller binaries, but much more is expected the more this is completed.

The majority of the work was of course to become Python3.11 compatible, and unfortunately the attribute lookups are not as optimized as for 3.10 yet, which may cause disappointing results for performance initially. We will need to complete that before benchmarks will make much sense.

For the next release, full Python 3.11 support is planned. I believe it should be usable. Problems with 3.11 may get hotfixes, but ultimately the develop version is probably the one to recommend when using 3.11 with Nuitka, as there will be the whole set of fixes, since not everything will be ported back.

The new reports should be used in bug reporting soon. We foresee that for issue reports, these may well become mandatory. Together with the ability to create a `virtualenv` from the reports, this may make reproducing issues a breeze, but first tries on complex projects were also highlighting that it may not be as simple.

Nuitka Release 1.4

This release contains a large amount of performance work, where specifically Python versions 3.7 or higher see regressions in relative performance to CPython fixed. Many cases of macros turned to functions have been found and resolved. For 3.10 specifically we take advantage of new opportunities for optimization. And generally avoiding DLL calls will benefit execution times on platform where the Python DLL is used, most prominently Windows.

Then this also adds new features, specifically custom reports. Also tools to aid with adding Nuitka package configuration input data, to list DLLs and data files.

With `multidist` we see a brand new ability to combine several programs into one, that will become very useful for packaging multiple binaries without the overhead of multiple distributions.

Bug Fixes

- Standalone: Added implicit dependencies for `dependency_injector` package. Fixed in 1.3.1 already.
- Fix, the generated metadata nodes for distribution queries had an error in their generated children handling that could cause crashes at compile time. Fixed in 1.3.2 already.

- Standalone: Added implicit dependencies for `passlib.apache` package. Fixed in 1.3.2 already.
- Windows: Fix, our shortcut to find DLLs by analyzing loaded DLLs stumbled in a case of a DLL loaded into the compiling Python that had no filename associated, while strange, we need to handle this as well. Fixed in 1.3.3 already.
- Standalone: Also need to workaround more decorator tricks for `networkx`. Fixed in 1.3.3 already.
- Scons: Fix, was not updating `PATH` environment variable anymore, which could lead to externally provided compilers and internal winlibs gcc clashing on Windows, but should be a general problem. Fixed in 1.3.4 already.
- Standalone: Added support for `cefpython3` package. Fixed in 1.3.4 already.
- Standalone: Added support for newer `webview` package versions. Fixed in 1.3.4 already.
- Standalone: Fix, some extension modules set their `__file__` to `None` during multi phase imports, which we then didn't update anymore, however that is necessary. Fixed in 1.3.4 already.
- Python3.10+: Fix, was not supporting `match` cases where an alternative had no condition associated. Fixed in 1.3.5 already.
- Windows: Identify Windows ARM architecture Python properly. We do not yet support it, but we should report it properly and some package configurations are already taking it already into account. Fixed in 1.3.5 already.
- Fix, the Nuitka meta path based loader, needs to expose a `__module__` attribute because there is code out there, that identifies standard loaders through looking at this value, but crashes without it. Fixed in 1.3.5 already.
- Fix, very old versions of the `importlib_metadata` backport were using themselves to load their `__version__` attribute. Added a workaround for it, since in Nuitka it doesn't work until after loading the module.
- Fix, value escapes for attribute and subscript assignments sources were not properly annotated. This could cause incorrect code execution. Fixed in 1.3.6 already.
- Fix, "pure" functions, which are currently only our complex call helper functions, were not visited in all cases. This lead to a crash in code generation after modules using them got demoted to bytecode. After use from cache, this didn't happen again. Fixed in 1.3.6 already.
- Standalone: Added more implicit dependencies of crypto packages. Fixed in 1.3.6 already.
- Standalone: Added implicit dependencies of `pygments.styles` module. Fixed in 1.3.6 already.
- Fix, was falsely encoding `Ellipsis` too soon during tree building. It is not quite like `True` and `False`. Fixed in 1.3.6 already.
- Standalone: Fix, `numpy` on macOS didn't work inside an application bundle anymore. Fixed in 1.3.7 already.
- Python3.8+: Fix, need to follow change for extension module handling, otherwise some uses of `os.add_dll_directory` fail to work. Fixed in 1.3.8 already.
- Standalone: Added missing implicit dependencies of `sqlalchemy`. Fixed in 1.3.8 already.
- Python3.9+: Fix, resource reader files was not fully compatible and needed to register with `importlib.resources.as_file` to work well with it. Fixed in 1.3.8 already.
- Fix, the version check for `cv2` was not working with the `opencv-python-headless` variant. Package name and distribution name is not a 1:1 mapping for all things. Fixed in 1.3.8 already.
- Standalone: Added DLLs needed for `tls_client` package.
- Fix, imports of resolved names should be modified for runtime too. Where Nuitka recognizes aliases, as e.g. the `requests` module does, it only adding a dependency on the resolved name, but not `requests` itself. The import however was still done at runtime on `requests` which then didn't work. This was only visible if only these aliases to other modules were used.

- Onefile: Fix, do not send duplicate CTRL-C to child process. Our test only send it to the bootstrap process, rather than the process group, as it normally is working, therefore misleading us into sending it to the child even if not needed.
- Onefile: When not using cached mode, on Windows the temporary folder used sometimes failed to delete after the executable stopped with CTRL-C. This is due to races in releasing of locks and process termination and AV tools, so we now retry for some time, to make sure it is always deleted.
- Standalone: Fix, was not ignoring `.dylib` when scanning for data files unlike all other DLL suffixes.
- Standalone: Added missing implicit dependency of `mplcairo`.
- Standalone: The main binary name on non-Windows didn't have a suffix `.bin` unlike in accelerated mode. However, this didn't work well for packages which have binaries colliding with the package name. Therefore now the suffix is added in this case too.
- macOS: Workaround bug in `platform_utils.paths`. It is guessing the wrong path for included data files with Nuitka.
- Standalone: Added DLLs of `sound_lib`, selecting by OS and architecture.
- Fix, for package metadata as from `importlib.metadata.metadata` for use at runtime we need to use both package name and distribution name to create it, or else it failed to work. Packages like `opencv-python-headless` can now with this too.
- Standalone: Added support for `tkinterweb` on Windows. Other platforms will need work to be done later.

New Features

- UI: Added new option to listing package data files. This is for use with analyzing standalone issues. And will output all files that are data files for a given package name.

```
python -m nuitka --list-package-data=tkinterweb
```

- UI: Added new option to listing package DLL files. This is also for use with analyzing standalone issues.

```
python -m nuitka --list-package-dlls=tkinterweb
```

- Reports: The usages of modules, successful or not, are now included in the compilation report. Checking out which ones are `not-found` might help recognition of issues.
- Multidist: You can now experimentally create binaries with multiple entry points. At runtime one of multiple `__main__` will be executed. The option to use is multiple `--main=some_main.py` arguments. If then the binary name is changed, on execution you get a different variant being executed.

Note

Using it with only one replaces the previous use of the positional argument given and is not using multidist at all.

Note

Multidist is compatible with onefile, standalone, and mere acceleration. It cannot be used for module mode obviously.

For deployment this can solve duplication.

Note

For wheels, we will probably change those with multiple entry points to compiling multidist executables, so we do avoid Python script entry points there. But this has not yet been done.

- Onefile: Kill non-cooperating child processes on CTRL-C after a grace period, that can be controlled at compile time with `--onefile-child-grace-time` the hard way. This avoids hangs of processes that fail to properly shutdown.
- Plugins: Add support for extra global search paths to mimic `sys.path` manipulations in the Yaml configuration with new `global-sys-path` import hack.
- Reports: Include used distributions of compiled packages and their versions.
- Reports: Added ability to generate custom reports with `--report-template` where the user can provide a Jinja2 template to make his own reports.
- Anti-Bloat: Added support for checking python flags. There are `no_asserts`, `no_docstrings` and `no_annotations` now. These can be used to limit rules to be only applied when these optional modes are active.

Not all packages will work in these modes, but often can be enhanced to work with relatively little patching. This allows to limit these patches to only where they are necessary.

Optimization

- Anti-Bloat: Avoid using `sparse` and through that Numba in the `scipy` package, reducing its distribution footprint. Part of 1.3.3 already.
- Anti-Bloat: Avoid IPython and Numba in `trimesh` package. Part of 1.3.3 already.
- Anti-Bloat: Avoid Numba in `shap` package. Part of 1.3.8 already.
- Anti-bloat: Removed `xgboost` `docstring` dependencies, such that `--python-flag=no_docstrings` can be used with this package.
- For guided deep copy `frozenset` and empty `tuple` need no copies

This also speeds up copies of non-empty tuples by avoiding that size checking branch in construction with Python 3.10 or higher.

- For node construction, avoid keyword argument style calls of the base class, where there is only a single argument. They don't really help readability, but cost compile time.
- Determine guard mode of frames dynamically and avoid frame preservation checks where they are not needed.

For Python2 this is necessary, but not for Python3, so make the function avoid finding the parent frame for that version entirely, which should speed up compilation as well.

By not hard coding frame guard mode at creation time, and instead determine it at compile time, after optimization, so this now allows to use the "once" mode more often. This affects contractions and also classes on the module level right now. They do not need a cached frame, since their code is only executed once.

By avoiding that useless code, the C compiler also has a slightly better scalability, since the classes are all created in one function that then has less code.

- The bytecode cache is now checking if the used modules or attempted to be used modules are available or not in just the same way. Previously it was very dependent on the file system to contain the same things, which was not giving cache hits even after only creating a new folder near a binary, since that affected importable modules. With the new check it should be much more directly hitting even across different virtual environments, but with same code.
- Generate base classes or mixins for all kinds of expression, statements and statement sequences. The previous code had a dedicated variant for single child, to allow faster operation in a common case, but still a lot of `hasattr/getattr/setattr` on dynamic attribute names were done. This was making the tree traversal during optimization slower than necessary.

Another shortcoming was that for some nodes, some values are optional, where for others, they are not. Some values are a `tuple` actually, while most are nodes only. However, dealing with this generically was also slower than necessary.

The new code now enforces children types during creation and updated, it rejects unexpected `None` values for non-optional children, and it provides generated code to do this in the fastest way possible, although surely some more improvements will come here.

Also when abstract executing the tree, rather than generically visiting all children, this now just unrolls this, and there are even some modes added, where a node can indicate properties, e.g. `auto_compute_handling = "final,no_raise"` will tell the code generator that this expression never raises in the computation, and is final, i.e. doesn't have any code to evaluate, because it cannot be optimized any further.

Also the way `checkers` previously worked, for every node creation, for every child update, a dictionary lookup had to be done. This is now hard coded for the few nodes that actually want to convert values on the fly and we might make a difference in the future for optional checkers, such that these are only run in debug mode.

These changes brought about much faster compilation, however the big elephant in the room will still be merging value traces, and scalability problems remain there.

- Attribute node generation for method specs like `dict.update`, etc. now provide type shapes. From these type shapes, mixins for the result value type are picked automatically. Previously these shapes were added manually. In some cases, they were even missing. In a few cases, where the type is dependent on the Python version, we do not currently do this though, so this needs more work, but expanding the coverage got easier in this way.
- Determining the used modules of a module requires a tree visit operations, that then asked for node types and used different APIs. This has been unified to be able to call a virtual method instead, which saves some compile time.
- After scanning for a module, we then determined the module kind even after we previously knew it during the scan. Also, this was checking `os.path.isdir` which was making it relatively slow and wasting 5% compile time on the IO being done. The check got enhanced and most often replaced with using the knowledge from the original import scan eliminating this time.
- Already most helper code of Nuitka was included from `.c` files, but compiled generators and compiled cells codes were not yet done like this, making life unnecessarily harder for the compiler and linker. This should also allow more optimization for some codes.
- Cache the plugin decisions about recursion for a module name. When a module is imported multiple times plugins were each asked again and again, which is not a good thing to do.
- Avoid usage of `PyObject_RichCompareBool` API, as we have our own comparison functions that are faster and faster to call without crossing of DLL barrier.
- Python3.8+: Avoid usage of `PyIndex_Check` which has become an API in 3.8, and was as a result not inlined anymore with a DLL barrier was to be crossed, making all kinds of multiplication and subscript/index operations slower.

- Replace `PyNumber_Index` API with our own code. As of 3.10 it enforces a conversion to `long` that for Nuitka is not a good thing to do in all places. But also due to DLL barrier it was potentially slow to call, and is used a lot, and we can drop the checks that are useless for Nuitka.
- Python3.7+: Avoid the use of `PyImport_GetModule` for looking up imported modules from `sys.modules`, rather look it up from interpreter internals, also this was using subscript functions, when this is always a dictionary.
- Avoid using `PyImport_GetModuleDict` and instead have our own API to get this quicker.
- Faster exception match checks and sub type checks.

This solves a `TODO` about inlining the API function used, so we can be faster in a relatively common operation. For every exception handler, we had to do one API call there.

- Faster subtype checks.

These are common in binary operations on non-identical types, but also needed for the exception checks, and object creation through class type calls. With our own `PyType_IsSubType` replacement these faster to use and avoid the API call.

- Faster Python3 `int` value startup initialization.

On Python 3.9 or higher we can get small int values directly from the interpreter, and with 3.11 they are accessible as global values.

Also we no longer de-duplicate small int values through our cache, since there is no use in this, saving a bunch of startup time. And we can create the values with our own API replacement, that will work during startup already and save API calls as these can be relatively slow. And esp. for the small values, this benefits from not having to create them.

- Faster Python3 `bytes` value startup initialization.

On Python 3.10 or higher, we can create these values ourselves without an API call, avoiding its overhead.

Also we no longer de-duplicate small bytes values through our cache, because that is already done by the API and our replacement, so this was just wasting time.

- Faster `slice` object values with Python 3.10 or higher

On Python 3.10 or higher, we can create these values ourselves without an API call, avoiding its overhead.

These are important for Python3, because `a[x:y]` in the general case has to use `a[slice(x,y)]` on that version, making this somewhat relevant to performance in some cases.

- Faster `str` built-in with API calls

For common cases, this avoids API calls. We mostly have this such that `print` style tests do not have this as API calls where we strive to remove all API calls for given programs.

- Faster exception normalization.

For the common case, we have our own variant of `PyErr_NormalizeException` that will avoid the API call. It may still call the `PyObject_IsSubclass` API, for which we only have started replacement work, but this is already a step ahead in the right direction.

- Faster object releases

For Python3.8 or higher when our code released objects, it was doing that with an API call, due to a macro change in Python headers. We revert that and do it still on our own which avoids the performance penalty.

- Enable Python threading during extension module DLL loading

We now release the GIL for Python3.8 or higher when loading the DLL, following a change in that version.

- Faster variable handling in trace collection. The code was doing checks for variable types, to decide what to do e.g. when control flow escapes for a variable. However, this is faster if solved with a virtual method in those variable classes, shifting the responsibility to inside there.
- For call codes the need to check the return value was not perfectly annotated in all cases. This is now driven by the expression rather than passed, and will result in better code generated in some corner cases.

Organisational

- Release: Make clear we require `wheel` and `setuptools` to install by adding a `pyproject.toml` that addresses a warning of `pip`. Part of 1.3.6 release already.
- Debugging: When plugins evaluate `when` conditions that raise, output which it was exactly. Part of 1.3.3 already.
- Anti-Bloat: Added a mnemonic and more clear message for the case of unwanted imports being encountered. Also do not warn about IPython itself using IPython packages, that must of course be considered normal. Now it also lists the module that does the unwanted usage immediately. Previously this was not as clear.
- UI: More clear output for not yet supported Python version. Make it more clear in the message, what is the highest supported version, and what version is Nuitka and what is Python in this.
- UI: Make sure data files have normalized paths. Specifically on Windows, otherwise a mix of slashes could appear. Part of 1.3.6 release already.
- UI: Make it clear that disabling the console harms your debugging when we suggest the `--disable-console` for GUI packages. Otherwise using that, they just deprive themselves of ways to get error information.
- UI: The ordering of scons `ccache` report was not enforced. Part of 1.3.7 release already.
- Quality: Use proper temporary filename during autoformat, so as to avoid flicker in Visual Code, e.g. search results.
- User Manual: Was still using old option name for `--onefile-tempdir-spec` that has since been made not OS specific, with even the OS specific name being removed.
- Standalone: Do not include data files scanned with `site-packages` or `__pycache__` folders. This should make it easier to use `--include-data-file=./**.qml:` when you have a virtualenv living in the same folder.
- Onefile: Added check for compression ability before starting the compilation to inform the user immediately.
- Release: Mark macOS as supported in PyPI categories. This is of course true for a long time already.
- Release: Mark Android as supported in PyPI categories as well. With some extra work, it can be used.
- User Manual: Added section pointing to and explaining compilation reports. This has become extremely useful even if still somewhat work in progress.
- User Manual: Added table with included custom reports, at this time only the license reports, which is very rough shape and needs contributors for good looks and content.

Cleanups

- Plugins: Moved parts of the `pywebview` plugin that pertain to the DLLs and data files to package configuration.

- Made the user query code a dedicated function, so it can be reused and more consistent across its uses in Nuitka. With a default that is proposed to a user, and a default that applies if used non-interactively. We will switch all prompts to using this.
- Code generation for module, class and function frames is now unified, removing duplication while also becoming more flexible. For generators this work has been started, but is not yet completed.
- Nodes exposing used modules now implement the same virtual method providing a list of them.
- Make sure to pass `tuple` values rather than `list` values from the tree building stage and node optimization creating new nodes. This allows us to drop conversions previously done inside of nodes.

Tests

- Do not enable deprecated plugins, the warnings about them break tests.
- Ignore Qt binding warnings in tests, some are less supported than `PySide6` or commercial `PySide2`.

Summary

The focus of this release was first a major restructuring of how children are handled in the node tree. The generated code opens up the possibility of many more scalability improvements in the coming releases. The pure iteration speed for the node tree will make compile times for the Python part even shorter in coming releases. Scalability will be a continuous focus for some releases.

Then the avoiding of API calls is a huge benefit for many platforms that are otherwise at a disadvantage. This is also only started. We will aim at getting more complex programs to do next to none of these, so far only some tests are working after program start without them, which is of course big progress. We will progress there with future releases as well.

Catching up on problems that previous migrations have not discovered is also a huge step forward to restoring the performance supremacy, that was not there anymore in extreme cases.

The Yaml package configuration work is showing its fruits. More people have been able to contribute changes for `anti-bloat` or missing dependencies than ever before.

Some part of the Python 3.11 work have positively influenced things, e.g. with the frame cleanup. The focus of the next release cycle shall be to add support for it. Right now, generator frames need a cleanup to be finished, to also become better and working with 3.11 at the same time. Where possible, work to support 3.11 was also conducted as a cleanup action, or reduction of the technical debts.

All in all, it is fair to say that this release is a big leap forward in all kinds of ways.

Nuitka Release 1.3

This release contains a large amount of performance work, that should specifically be useful on Windows, but also generally. A bit of scalability work has been applied, and as usual many bug fixes and small improvements, many of which have been in hotfixes.

Bug Fixes

- macOS: Framework build of `PySide6` were not properly supporting the use of `WebEngine`. This requires including frameworks and resources in new ways, and actually some duplication of files, making the bundle big, but this seems to be unavoidable to keep the signature intact.
- Standalone: Added workaround for `dotenv`. Do not insist on compiled package directories that may not be there in case of no data files. Fixed in 1.2.1 already.
- Python3.8+: Fix, the `ctypes.CDLL` node attributes the `winmode` argument to `Python2`, which is wrong, it was actually added with 3.8. Fixed in 1.2.1 already.

- Windows: Attempt to detect corrupt object file in MSVC linking. These might be produced by `cl.exe` crashes or `clcache` bugs. When these are reported by the linker, it now suggests to use the `--clean-cache=ccache` which will remove it, otherwise there would be no way to cure it. Added in 1.2.1 already.
- Standalone: Added data files for `folium` package. Added in 1.2.1 already.
- Standalone: Added data files for `branca` package. Added in 1.2.1 already.
- Fix, some forms `try` that had exiting `finally` branches were tracing values only assigned in the `try` block incorrectly. Fixed in 1.2.2 already.
- Alpine: Fix, Also include `libstdc++` for Alpine to not use the system one which is required by its other binaries, much like we already do for Anaconda. Fixed in 1.2.2 already.
- Standalone: Added support for latest `pytorch`. One of our workarounds no longer applies. Fixed in 1.2.2 already.
- Standalone: Added support for webcam on Windows with `opencv-python`. Fixed in 1.2.3 already.
- Standalone: Added support for `pytorch_lightning`, it was not finding metadata for `rich` package. Fixed in 1.2.4 already.

For the release we found that `pytorch_lightning` may not find `rich` installed. Need to guard `version()` checks in our package configuration.

- Standalone: Added data files for `dash` package. Fixed in 1.2.4 already.
- Windows: Retry replace `clcache` entry after a delay, this works around Virus scanners giving access denied while they are checking the file. Naturally you ought to disable those for your build space, but new users often don't have this. Fixed in 1.2.4 already.
- Standalone: Added support for `scipy` 1.9.2 changes. Fixed in 1.2.4 already.
- Catch corrupt object file outputs from `gcc` as well and suggest to clean cache as well. This has been observed to happen at least on Windows and should help resolve the `ccache` situation there.
- Windows: In case `clcache` fails to acquire the global lock, simply ignore that. This happens sporadically and barely is a real locking issue, since that would require two compilations at the same time and for that it largely works.
- Compatibility: Classes should have the `f_locals` set to the actual mapping used in their frame. This makes Nuitka usable with the `multidispatch` package which tries to find methods there while the class is building.
- Anaconda: Fix, newer Anaconda versions have TCL and Tk in new places, breaking the `tk-inter` automatic detection. This was fixed in 1.2.6 already.
- Windows 7: Fix, onefile was not working anymore, a new API usage was not done in a compatible fashion. Fixed in 1.2.6 already.
- Standalone: Added data files for `lark` package. Fixed in 1.2.6 already.
- Fix, `pkgutil.iter_modules` without arguments was given wrong compiled package names. Fixed in 1.2.6 already.
- Standalone: Added support for newer `clr` DLLs changes. Fixed in 1.2.7 already.
- Standalone: Added workarounds for `tensorflow.compat` namespace not being available. Fixed in 1.2.7 already.
- Standalone: Added support for `tkextrafont`. Fixed in 1.2.7 already.
- Python3: Fix, locals dict test code testing if a variable was present in a mapping could leak references. Fixed in 1.2.7 already.
- Standalone: Added support for `timm` package. Fixed in 1.2.7 already.

- Plugins: Add `tls` to list of sensible plugins. This enables at least `pyqt6` plugin to do networking with SSL encryption.
- Standalone: Added implicit dependencies of `sklearn.cluster`.
- FreeBSD: Fix, `fcopyfile` is no longer available on newest OS version, and include files for `sendfile` have changed.
- MSYS2: Add back support for MSYS Posix variant. Now `onefile` works there too.
- Fix, when picking up data files from command line and plugins, different exclusions were applied. This has been unified to get better coverage for avoiding to include DLLs and the like as data files. DLLs are not data files and must be dealt with differently after all.

New Features

- UI: Added new option for cache disabling `--disable-cache` that accepts `all` and cache names like `ccache`, `bytecode` and on Windows, `dll-dependencies` with selective values.

Note

The `clcache` is implied in `ccache` for simplicity.

- UI: With the same values as `--disable-cache` `Nuitka` may now also be called with `--clean-cache` in a compilation or without a filename argument, and then it will erase those caches current data before making a compilation.
- macOS: Added `--macos-app-mode` option for application bundles that should run in the background (`background`) or are only a UI element (`ui-element`).
- Plugins: In the `Nuitka` package configuration files, the `when` allows now to check if a plugin is active. This allowed us to limit console warnings to only packages whose plugin was activated.
- Plugins: Can now mark a plugin as a GUI toolkit responsible with the consequence that other toolkit detector plugins are all disabled, so when using `tk-inter` no longer will you be asked about `PySide6` plugin, as that is not what you are using apparently.
- Plugins: Generalized the GUI toolkit detection to include `tk-inter` as well, so it will now point out that `wx` and the Qt bindings should be removed for best results, if they are included in the compilation.
- Plugins: Added ability to provide data files for macOS `Resources` folder of application bundles.
- macOS: Fix, Qt WebEngine was not working for framework using Python builds, like the ones from PyPI. This adds support for both `PySide2` and `PySide6` to distribute those as well.
- MSYS2: When asking a CPython installation to compress from the POSIX Python, it crashed on the main filename being not the same.
- Scons: Fix, need to preserve environment attached modes when switching to winlibs gcc on Windows. This was observed with MSYS2, but might have effects in other cases too.

Optimization

- Python3.10+: When creating dictionaries, lists, and tuples, we use the newly exposed dictionary free list. This can speedup code that repeatedly allocates and releases dictionaries by a lot.

- Python3.6+: Added fast path to dictionary copy. Compact dictionaries have their keys and values copied directly. This is inspired by a Python3.10 change, but it is applicable to older Python as well, and so we did.
- Python3.9+: Faster compiled object creation, esp. on Python platforms that use a DLLs for libpython, which is a given on Windows. This makes up for core changes that went unnoticed so far and should regain relative speedups to standard Python.
- Python3.10+: Faster float operations, we use the newly exposed float free list. This can speed up all kinds of float operations that are not doable in-place by a lot.
- Python3.8+: On Windows, faster object tracking is now available, this previously had to go through a DLL call, that is now removed in this way as it was for non-Windows only so far.
- Python3.7+: On non-Windows, faster object tracking is now used, this was regressed when adding support for this version, becoming equally bad as all of Windows at the time. However, we now managed to restore it.
- Optimization: Faster deep copy of mutable tuples and list constants, these were already faster, but e.g. went up from 137% gain factor to 201% on Python3.10 as a result. We now use guided a deep copy, which then has the information, what types it is going to copy, removing the need to check through a dictionary.
- Optimization: Also have own allocator function for fixed size objects. This accelerates allocation of compiled cells, dictionaries, some iterators, and lists objects.
- More efficient code for object initialization, avoiding one DLL calls to set up our compiled objects.
- Have our own `PyObject_Size` variant, that will be slightly faster and avoids DLL usage for `len` and size hints, e.g. in container creations.
- Avoid using non-optimal `malloc` related macros and functions of Python, and instead of the fasted form generally. This avoids Python DLL calls that on Windows can be particularly slow.
- Scalability: Generated child mixins are now used for the generated package metadata hard import nodes calls, and for all instances of single child tuple containers. These are more efficient for creation and traversal of the tree, directly improving the Python compile time.
- Scalability: Slightly more efficient compile time constant property detections. For `frozenset` there was not need to check for hashable values, and some branches could be replaced with e.g. defining our own `EllipsisType` for use in short paths.
- Windows: When using MSVC and LTO, the linking stage was done with only one thread, we now use the proper options to use all cores. This is controlled by `--jobs` much like C compilation already is. For large programs this will give big savings in overall execution time. Added in 1.2.7 already.
- Anti-Bloat: Remove the use of `pytest` for `dash` package compilation.
- Anti-Bloat: Remove the use of `IPython` for `dotenv`, `pyvista`, `python_utils`, and `trimesh` package compilation.
- Anti-Bloat: Remove `IPython` usage in `rdkit` improving compile time for standalone by a lot. Fixed in 1.2.7 already.
- Anti-Bloat: Avoid `keras` testing framework when using that package.

Organisational

- Plugins: The `numpy` plugin functionality was moved to `Nuitka` package configuration, and as a result, the plugin is now deprecated and devoid of functionality. On non-Windows, this removes unused duplications of the `numpy.core` DLLs.
- User Manual: Added information about macOS entitlements and Windows console. These features are supported very well by `Nuitka`, but needed documentation.

- UI: Remove alternative options from `--help` output. These are there often only for historic reasons, e.g. when an option was renamed. They should not bother users reading them.
- Plugins: Expose the `mnemonics` option to plugin warnings function, and use it for `pyside2` and `pyqt5` plugins.
- Quality: Detect trailing/leading spaces in Nuitka package configuration `description` values during their automatic check.
- UI: Detect the CPython official flavor on Windows by comparing to registry paths and consider real prefixes, when being used in `virtualenv` more often, e.g. when checking for CPython on Apple.
- UI: Enhanced `--version` output to include the C compiler selection. It is doing that respecting your other options, e.g. `--clang`, etc. so it will be helpful in debugging setup issues.
- UI: Some error messages were using `%r` rather than `'%s'` to output file paths, but that escaped backslashes on Windows, making them look worse, so we changed away from this.
- UI: Document more clearly what `--output-dir` actually controls.
- macOS: Added options hint that the `Foundation` module requires bundle mode to be usable.
- UI: Allow using both `--follow-imports` and `--nofollow-imports` on command line rather than erroring out. Simply use what was given last, this allows overriding what was given in project options tests should the need arise.
- Reports: Include plugin reasons for pre and post load modules provided. This solves a TODO and makes it easier to debug plugins.
- UI: Handle `--include-package-data` before compilation, removing the ability to use pattern. This makes it easier to recognize mistakes without a long compilation and plugins can know them this way too.
- GitHub: Migration workflows to using newer actions for Python and checkout. Also use newer Ubuntu LTS for Linux test runner.
- UI: Catch user error of running Nuitka with the `pythonw` binary on Windows.
- UI: Make it clear that MSYS2 defaults to `--mingw64` mode. It had been like this, but the `--help` output didn't say so.
- GitHub: Updated contribution guidelines for better readability.
- GitHub: Use organisation URLs everywhere, some were still pointing to the personal rather than the organisation one. While these are redirected, it is not good to have them like this.
- Mastodon: Added link to <https://fosstodon.org/@kayhayan> to the PyPI package and User Manual.

Cleanups

- Nodes for hard import calls of package meta data now have their base classes fully automatically created, replacing what was previously manual code. This aims at making them more consistent and easier to add.
- When adding the new Scons file for C compiler version output, more values that are needed for both onefile and backend compilation were moved to centralized code, simplifying these somewhat again.
- Remove unused `main_module` tag. It cannot happen that a module name matches, and still thinks of itself as `__main__` during compilation, so that idea was unnecessary.
- Generate the dictionary copy variants from template code rather than having manual duplications. For `dict.copy()`, for deep copy (needed e.g. when there are escaping mutable keyword argument constant values in say a function call), and for `**kw` value preparation in the called function (checking argument types), we have had diverged copies, that are now unified in a single Jinja2 template for optimization.

- Plugins: Also allow providing generators for providing extra DLLs much like we already do for data files.
- Naming of basic tests now makes sure to use a `Test` suffix, so in Visual Code selector they are more distinct from Nuitka code modules.
- Rather than populating empty dictionaries, helper code now uses factory functions to create them, passing keys and values, and allowing values to be optional, removing noisy `if` branches at call side.
- Removed remaining `PyDev` annotations, we don't need those anymore for a long time already.
- Cleanup, avoid lists objects for ctypes defined functions and their `arglist`, actually tuples are sufficient and naturally better to use.
- Spelling cleanups were resumed, as an ongoing action.

Tests

- Added construct test that demonstrates the mutable constant argument passing for lists to see the performance gains in this area too.
- Made construct runner `--diff` output usable for interactive usage.
- Repaired Nuitka Speedcenter, but it's not yet too usable for general consumption. More work will be needed there, esp. to make comparisons more accessible for the general public.

Summary

The major achievement of this release was the removal of the long lived `numpy` plug-in, replacing it with package based configuration, that is even more optimal and works perfectly across all platforms on both important package installation flavors.

This release has a lot of consolidation efforts, but also as a result of 3.11 investigations, addresses a lot of issues, that have crept in over time with Python3 releases since 3.7, each time, something had not been noticed. There will more need for investigation of relative performance losses, but this should address the most crucial ones, and also takes advantage of optimization that had become with 3.10 already.

There is also some initial results from cleanups with the composite node tree structure, and how it is managed. Generated "child(ren) having" mixins, allow for faster traversal of the node tree.

Some technical things also improved in Scons. Using multiple cores in LTO with MSVC with help this a lot, although for big compilations `--lto=no` probably has to be recommended still.

More `anti-bloat` work on more packages rounds up the work.

For macOS specifically, the WebEngine support is crucial to some users, and the new `--macos-app-mode` with more GUI friendly default resolves long standing problems in this area.

And for MSYS2 and FreeBSD, support has been re-activated, so now 4 OSes work extremely well (others too likely), and on those, most Python flavors work well.

The performance and scalability improvements are going to be crucial. It's a pity that 3.11 is not yet supported, but we will be getting there.

Nuitka Release 1.2

This release contains a large amount of new compatibility features and a few new optimization, while again consolidating what we have. Scalability should be better in many cases.

Bug Fixes

- Standalone: Added implicit dependency of `thinc` backend. Fixed in 1.1.1 already.

- Python3.10: Fix, `match` statements with unnamed star matches could give incorrect results. Fixed in 1.1.1 already.

```
match x:
    case [*_, y]:
        ... # y had wrong value here.
```

- Python3.9+: Fix, file reader objects must convert to `str` objects. Fixed in 1.1.1 already.

```
# This was the `repr` rather than a path value, but it must be usable
# like that too.
str(importlib.resources.files("package_name").joinpath("lala"))
```

- Standalone: Added data file of `echopype` package. Fixed in 1.1.1 already.
- Anti-Bloat: Remove non-sense warning of compiled `pyscf`. Fixed in 1.1.1 already.
- macOS: Fix, in LTO mode using `incbin` can fail, switch to source mode for constants resources. Fixed in 1.1.2 already.
- Standalone: Add support for `sv_ttk` module. Fixed in 1.1.2 already.
- macOS: Fix, was no longer correcting `libpython` path, this was a regression preventing CPython for creating properly portable binary. Fixed in 1.1.2 already.
- macOS: Fix, main binary was not included in signing command. Fixed in 1.1.3 already.
- Standalone: Added implicit dependency of `orjson`. Due to `zoneinfo` not being automatically included anymore, this was having a segfault. Fixed in 1.1.3 already.
- Standalone: Added support for new `shapely`. Fixed in 1.1.4 already.
- macOS: Ignore extension module of non-matching architecture. Some wheels contain extension modules for only `x86_64` arch, and others contain them only for `arm64`, preventing the standalone build. Fixed in 1.1.4 already.
- Standalone: Added missing `sklearn` dependencies. Fixed in 1.1.4 already.
- Fix, packages available through relative import paths could be confused with the same ones imported by absolute paths. This should be very hard to trigger, by normal users, but was seen during development. Fixed in 1.1.4 already.
- Standalone: Apply import hacks for `pywin32` modules only on Windows, otherwise it can break e.g. macOS compilation. Fixed in 1.1.4 already.
- Windows: More robust DLL dependency caching, otherwise e.g. a Windows update can break things. Also consider plugin contribution, and Nuitka version, to be absolutely sure, much like we already do for bytecode caching. Fixed in 1.1.4 already.
- Standalone: Fix, `seaborn` needs the same workaround as `scipy` for corruption with MSVC. Fixed in 1.1.4 already.
- UI: Fix, the `options-nanny` was no longer functional and therefore failed to warn about non working options and package usages. Fixed in 1.1.5 already.
- macOS: Do not use extension modules of non-matching architecture. Fixed in 1.1.5 already.
- Windows: Fix, resolving symlinks could fail for spaces in paths. Fixed in 1.1.6 already.
- Standalone: Added missing DLL for `lightgbm` module. Fixed in 1.1.6 already.
- Compatibility: Respect `super` module variable. It is now possible to have a module level change of `super` but still get compatible behavior with Nuitka. Fixed in 1.1.6 already.
- Compatibility: Make sure we respect `super` overloads in the builtin module. Fixed in 1.1.6 already.

- Fix, the anti-bloat replacement code for `numpy.testing` was missing a required function. Fixed in 1.1.6 already.
- Fix, `importlib.import_module` static optimization was mishandling a module name of `.` with a package name given. Fixed in 1.1.6 already.
- macOS: Fix, some extension modules use wrong suffixes in self references, we need to not complain about this kind of error. Fixed in 1.1.6 already.
- Fix, do not make `ctypes.wintypes` a hard import on non-Windows. Nuitka asserted against it failing, where some code handles it failing on non-Windows platforms. Fixed in 1.1.6 already.
- Standalone: Added data files for `vedo` package. Fixed in 1.1.7 already.
- Plugins: Fix, the `gi` plugin did always set `GI_TYPELIB_PATH` even if already present from user code. Also it did not handle errors to detect its value during compile time. Fixed in 1.1.7 already.
- Standalone: Added missing dependencies for `sqlalchemy` to have all SQL backends working. Fixed in 1.1.7 already.
- Added support Nixpkgs's default non-writable `HOME` directory. Fixed in 1.1.8 already.
- Fix, distribution metadata name and package name need not align, need to preserve the original looked up name from `importlib.metadata.distribution` call. Fixed in 1.1.8 already.
- Windows: Fix, catch usage of unsupported `CLCACHE_MEMCACHED` mode with MSVC compilation. It is just unsupported.
- Windows: Fix, file version was spoiled from product version if it was the only version given.
- Windows: The default for file description in version information was not as intended.
- Plugins: Workaround for PyQt5 as contained in Anaconda providing wrong paths from the build machine.
- macOS: After signing a binary with a certificate, compiling the next one was crashing on a warning about initially creating an ad-hoc binary.
- Fix, detect case of non-writable cache path, make explaining error exit rather than crashing attempting to write to the cache.
- macOS: Added support for `pyobjc` in version 9.0 or higher.

New Features

- Python3.11: For now prevent the execution with 3.11 and give a warning to the user for a not yet supported version. This can be overridden with `--experimental=python311` but at this times will not compile anything yet due to required and at this time missing core changes.
- macOS: Added option `--macos-sign-notarization` that signs with runtime signature, but requires a developer certificate from Apple. As its name implies, this is for use with notarization for their App store.
- DLLs used via `delvewheel` were so far only handled in the `zmq` plugin, but this has been generalized to cover any package using it. With that, e.g. `shapely` just works. This probably helps many other packages as well.
- Added `__compiled__` and `__compiled_constant__` attributes to compiled functions.
With this, it can be decided per function what it is and bridges like `pyobjc` can use it to create better code on their side for constant value returning functions.
- Added `support_info` check to Nuitka package format. Make it clear that `pyobjc` is only supported after 9.0 by erroring out if it has a too low version. It will not work at all before that version added support in upstream. Also using this to make it clear that `opencv-python` is best supported in version 4.6 or higher. It seems e.g. that video capture is not working with 4.5 at this time.

- Added `--report-template` which can be used to provide Jinja2 templates to create custom reports, and refer to built-in reports, at this time e.g. a license reports.

Optimization

- Trust the absence of a few selected hard modules and convert those to static raises of import errors.
- For conditional nodes where only one branch exits, and the other does not, no merging of the trace collection should happen. This should enhance the scalability and leads to more static optimization being done after these kinds of branches on variables assigned in such branches.

```

if condition1:
    a = 1
else:
    raise KeyError

if condition2:
    b = 1

# Here, "a" is known to be assigned, but before it was only "maybe"
# assigned, like "b" would have to be since, the branch may or may
# not have been taken.

```

- Do not merge tried blocks that are aborting with an exception handler that is not aborting. This is very similar to the change for conditional statements, again there is a control flow branch, that may have to be merged with an optional part, but sometimes that part is not optional from the perspective of the code following.

```

try:
    ... # potentially raising, but not aborting code
    return something() # this aborts
except Exception:
    a = 1

try:
    ... # potentially raising, but not aborting code
except Exception:
    b = 1

# Here, "a" is known to be assigned, but before it was only "maybe"
# assigned, like "b" would have to be since, the branch may or may
# not have been taken.

```

- Exception matches were annotating a control flow escape and an exception exit, even when it is known that no error is possible to be happening that comparison.

```

try:
    ...
except ImportError: # an exception match is done here, that cannot raise
    ...

```

- Trust `importlib.metadata.PackageNotFoundError` to exist, with this some more metadata usages are statically optimized. Added in 1.1.4 already.

- Handle constant values from trusted imports as trusted values. So far, trusted import values were on equal footing to regular variables, which on the module level could mean that no optimization was done, due to control flow escapes happening.

```
# Known to be False at compile time.
from typing import TYPE_CHECKING
...

if TYPE_CHECKING:
    from something import normally_unused_unless_type_checking
```

In this code example above, the static optimization was not done because the value may be changed on the outside. However, for trusted constants, this is no longer assumed to be happening. So far only `if typing.TYPE_CHECKING:` using code had been optimized.

- macOS: Use sections for main binary constants binary blob rather than C source code (which we started in a recent hotfix due to LTO issues with incbin) and onefile payload. The latter enables notarization of the onefile binary as well and makes it faster to unpack as well.
- Windows: Do not include DLLs from SxS. For PyPI packages these are generally unused, and self compiled modules won't be SxS installations either. We can add it back where it turns out needed. This avoids including `comctl32` and similar DLLs, which ought to come from the OS, and might impede backward compatibility only.
- Disabled C compilation of very large `azure` modules.
- The per module usage information of other modules was only updated in first pass was used in later passes. But since they can get optimized away, we have to update every time, avoiding to still include unused modules.
- Anti-Bloat: Fight the use of `dask` in `xarray` and `pint`, adding a mode controlling its use. This is however still incomplete and needs more work.
- Fix, the anti-bloat configuration for `rich.pretty` introduced a `SyntaxError` that went unnoticed. In the future compilation will abort when this happens.
- Standalone: Added support for including DLLs of `llvmlite.binding` package.
- Anti-Bloat: Avoid using `pywin32` through `pkg_resources` import. This seems rather pointless and follows an optimization done for the inline copy of `Nuitka` already, the `ctypes` code path works just fine, and this may well be the only reason why `pywin32` is included, which is by itself relatively large.
- Cache directory contents when scanning for modules. The `sys.path` and package directories were listed over and over, wasting time during the import analysis.
- Optimization: Was not caching not found modules, but retrying every time for each usage, potentially wasting time during import analysis.
- Anti-Bloat: Initial work to avoid `pytest` in `patsy`, it is however incomplete.

Organisational

- User Manual: Explain how to create 64/32 bits binaries on Windows, with there being no option to control it, this can otherwise be a bit unobvious that you have to just use the matching Python binary.
- UI: Added an example for a cached onefile temporary location spec to the help output of `--onefile-tempdir-spec` to make cached more easy to achieve in the proper way.
- UI: Quote command line options with space in value better, no need to quote an affected command line option in its entirety, and it looks strange.

- macOS: Catch user error of disabling the console without using the bundle mode, as it otherwise it has no effect.
- macOS: Warn about not providing an icon with disabled console, otherwise the dock icon is empty, which just looks bad.
- Debian: Also need to depend on `glob2` packages which the yaml engine expects to use when searching for DLLs.
- Debian: Pertain inline copies of modules in very old builds, there is e.g. no `glob2` for older releases, but only recent Debian releases need very pure packages, our backport doesn't have to do it right.
- macOS: More reliable detection of Homebrew based Python. Rather than checking file system via its `sitecustomize` contents. The environment variables are only present to some usages.
- Installations with pip did not include all license, README files, etc. which however was intended. Also the attempt to disable bytecode compilation for some inline copies was not effective yet.
- Renamed `pyzmq` plugin to `delvewheel` as it is now absolutely generic and covers all uses of said packaging technique.
- Caching: Use cache directory for cached downloads, rather than application directory, which is just wrong. This will cause all previously cached downloads to become unused and repeated.
- Quality: Updated development requirements to latest `black`, `isort`, `yamllint`, and `tqdm`.
- Visual Code: Added recommendation for extension for Debian packaging files.
- Added warning for `PyQt5` usage, since its support is very incomplete. Made the `PyQt6` warning more concrete. It seems only Qt threading does not work, which is of course still significant. Instead `PySide2` and `PySide6` are recommended.
- UI: Have dedicated options group for onefile, the spec for the temporary files is not a Windows option at all anymore. Also move the warnings group to the end, people need to see the inclusion related group first.
- User Manual: Explain how to create 64/32 bits binaries on Windows, which is not too obvious.

Cleanups

- Moved `PySide` plugins DLL search extra paths to the `Yaml` configuration. In this way it is not dependent on the plugin being active, avoiding cryptic errors on macOS when they are not found.
- Plugins: Avoid duplicate link libraries due to casing. We are now normalizing the link library names, which avoids e.g. `Shell132` and `shell132` to be in the result.
- Cleanups to prepare a `PyLint` update that so otherwise failed due to encountered issues.
- Plugins: Pass so called build definitions for generically. Rather than having dedicated code for each, and plugins can now provide them and pass their index to the `scons` builds.
- Onefile: Moved file handling code to common code reducing code duplication and heavily cleaning up the bootstrap code generally.
- Onefile: The `CRC32` checksum code was duplicated between constants `blob` and `onefile`, and has moved to shared code, with an actual interface to take the checksum.
- Spelling cleanups resumed, e.g. this time more clearly distinguishing between `run time` and `runtime`, the first is when the program executes, but the latter can be an environment provided by a C compiler.

Tests

- Tests: Added test that applies anti-bloat configuration to all found modules.

- Tests: Tests: Avoid including unused `nuitka.tools` code in reflected test, which should make it faster. The compiler itself doesn't use that code.

Summary

This release is again mainly a consolidation of previous release, as well as finishing off existing features. Optimization added in previous releases should have all regressions fixed now, again with a strong series of hotfixes.

New optimization was focused around findings with static optimization not being done, but still resulting in general improvements. Letting static optimization drive the effort is still paying off.

Scalability has seen improvements through some of the optimization, this time a lot less anti-bloat work has been done, and some things are only started. The focus will clearly now shift to making this a community effort. Expect postings that document the Yaml format we use.

For macOS specifically, with the sections work, onefile should launch faster, should be more compatible with signing, and can now be used in notarization, so for that platform, things are more round.

For Windows, a few issues with version information and onefile have been addressed. We should be able to use memory mapped view on this platform too, for faster unpacking of the payload, since it doesn't have to go through the file anymore.

Nuitka Release 1.1

This release contains a large amount of new compatibility features, while consolidating what we have. Scalability should be better in some cases.

Bug Fixes

- Standalone: Enhanced dependency scan of dependent DLLs to forward the containing package, so it can be searched in as well. This fixed at least PySide on macOS. Fixed in 1.0.1 already.
- macOS: Enhanced dependency detection to use normalized paths and therefore to be more stable. Fixed in 1.0.1 already.
- Standalone: Added support for the `networkx` package which uses new support for a function decorator trying to copy function default values. Fixed in 1.0.1 already.
- Standalone: Include data files for `pandas.io.format` package. This one has Jinja2 template files that will be needed when using this package.
- Python3.10: Fix, could crash in case a class was not giving `match` arguments, but the user did attempt to match them. This happened e.g. with `range` objects. Fixed in 1.0.2 already.
- Standalone: Added data files needed for `pyenchant` package. Fixed in 1.0.2 already.
- Python3.10: Fix, matching sequence with `as` assignments in them didn't check for sub-pattern given. Fixed in 1.0.2 already.
- Standalone: Fix, do not attempt to list non-existent `PATH` entries on Windows, these can crash the dependency detection otherwise. Fixed in 1.0.2 already.
- Standalone: Fix, on newer Linux, `linux-vdso.so.1` appears in output of `ldd` in a way that suggests it may exist, which of course it does not, this is a kernel virtual DLL. Fixed in 1.0.3 already.
- Fix, comparison expressions could give wrong results as a regression of the new release. Fixed in 1.0.3 already.
- Fix, on older Python (before 3.6), it could crash on data files defined in the Yaml config. Fixed in 1.0.4 already.
- Fix, binary operations could give wrong results as a regression of the new release. Fixed in 1.0.4 already.

- Standalone: Added support for `pyzbar` package. Fixed in 1.0.5 already.
- Standalone: Fix, empty directory structures were not working anymore due to a regression in the last release. Fixed in 1.0.5 already.
- Windows: Fix, detected Pythons from Windows registry may of course fail to execute, because they were e.g. manually deleted. This would show e.g. in onefile compression. Fixed in 1.0.5 already.
- Onefile: Fix, using a too old `zstandard` without finding another Python with a suitable one, lead to run time unpacking errors. Fixed in 1.0.6 already.
- Fix, the inline copy of Jinja2 imported `logging` for no good reason, which lead to errors for users who have a module of the same name, that it was then using instead. Fixed in 1.0.6 already.
- Fix, disable LTO mode for Anaconda Python, it is known to not work. Fixed in 1.0.6 already.
- Linux: Fix, no need to insist on icon path for onefile anymore. Fixed in 1.0.6 already.
- Standalone: Fix, the new version `certifi` was not working on Windows and 3.10 anymore. Fixed in 1.0.7 already.
- Standalone: Added support for more `rapidfuzz` implicit dependencies. Fixed in 1.0.8 already.
- Standalone: Added support for `vibora`. Fixed in 1.0.8 already.
- Fix, must not expose module name objects to Python import hooks. Fixed in 1.0.8 already.
- Fix, calls to bound methods of string values generated incorrect calls. Fixed in 1.0.8 already.
- Fix, do not crash in version detection on `gcc` error exit querying of its version.
- Standalone: Added back support for older versions of the `pyzmq` package.
- Standalone: Ignore `PATH` elements that fail to be listed as a directory. It appears e.g. on Windows, folders can exist, despite being unusable in fact. These can then cause errors in DLL dependency scan. Also avoid having `PATH` set when executing dependency walker, it appears to use it even if not asked to.
- Standalone: Added support for `tzlocal` package.
- Python3.10: Fix, `complex` literals were not working for mappings in `match` statements.
- Fix, `bool` built-in expressions were not properly annotating exception raises, where the value cannot raise on truth check.
- Standalone: Added support for the `agscheduler` package. Plugins must be done manually still with explicit `--include-module` calls.
- Standalone: Added support for using `shapely` in Anaconda as well.
- Debian: Fix, versioned dependency for `libzstd` should also be in package, this should restore Nuitka package builds for Debian Jessie.
- Standalone: Added support for `vtk` package.
- Windows: Fix, avoid using `pywin32` in our `appdirs` usage, it might be a broken installation and is optional to `appdirs` anyway, which then will fallback to using `ctypes` to make the lookups.
- Standalone: Added support for more `pandas` versions.
- Standalone: Adding support for `mk1` implicit DLL usage in Anaconda.
- Standalone: Added support for `jsonschema` with Python 3.10.
- Standalone: Added support for `pyfiglet` fonts data files.
- Scons: Avoid `gcc` linker command line length limits for module mode too.
- Standalone: Added data file of `distributed.config`.
- Standalone: Add support for `cv2` GUI on Linux, the Qt platform plugin is now included.

- Fix, the anti-bloat configuration for `numpy.testing` tools exposed an incomplete `suppress_warnings` replacement that could lead to errors in some functions of `numpy`.
- Standalone: Fix DLL dependency caching on Windows, need to consider DLL content of course too.
- Standalone: Added missing dependency for `torchvision`.
- Standalone: Added support for `torchvision` on Anaconda as well.
- Standalone: Added support for `panda3d`.
- Windows: Fix, need to make sure to use UTF-8 encoding for define values like company name. Otherwise the local system encoding is used, but the C compiler expects UTF-8 in wide literals. This may crash or give wrong results.
- Standalone: Added `facenet_torch` data files.
- Anaconda: Include `libstdc++.so` on Linux or else e.g. `cv2` will not work with system library.
- Windows: Fix, can have file version without a company name.

New Features

- Python3.10: Added support for assignments in `match` alternatives | syntax.
- Compatibility: Register Nuitka meta path based loader with `pkg_resources` such that checking resource presence with `has_resource` works too. This should also add support for using `jinja2.PackageLoader`, previously only `jinja2.FileSystemLoader` worked. Fixed in 1.0.1 already.
- Compatibility: Make function `__defaults__` attribute size changeable. For a long time, this was a relatively big issue for some packages, but now this is supported as well.
- Compatibility: Added support for `importlib.metadata.distribution` and `importlib_metadata.distribution` functions as well `importlib.metadata.metadata` and `importlib_metadata.metadata` functions.
- Onefile: Added support for including other binaries than the main executable in the payload. So far on non-Windows, we only made the main binary executable, hard coded, and nothing else. But Some things, e.g. Qt web engine, do require binaries to be used, and these no longer have the issue of missing x-bit on macOS and Linux now.
- Standalone: Resolve executable path when called through symbolic link, which makes file resolutions work properly for it, for this type of installing it in `%PATH%`.
- Python3.9+: Added support for `importlib.resources.files` with compiled modules.
It returns traversable objects, which can be used to open files, check them, etc. and this e.g. allows `jsonschema` to work with Python 3.10, despite bugs in CPython's compatibility layer.
- UI: Added interface method to specify filename patterns with package data inclusion option, making `--include-package-data` usable in many more cases, picking the only files or file types you want. You can now use `--include-package-data=package_name=*.txt` and select only a subset of package data files in this way. Before this, it included everything and `--noinclude-data-files` would have to be used.
- macOS: Make `runtime` signing an experimental option.
- Consistently allow `when` conditions for all package configuration elements, e.g. also DLLs.
- Plugins: Added method to overload to work on standalone binary specifically. This makes it easier to only modify that specific binary.
- Plugins: Added support for regular expressions in anti-bloat replacements, with new `replacements_re` code.

Optimization

- Add support for `os.path` hard module imports along with specialized nodes for file tests `os.path.exists`, `os.path.isfile`, and `os.path.isdir` aiming at tracking used files, producing warnings about missing files in the future.
- Standalone: Do not include `concurrent` standard library package automatically. This avoids the inclusion of `multiprocessing` which we essentially had reverted during last release cycle.
- Standalone: Do not include `zoneinfo` standard library package automatically. It has many data files and is often not used (yet).
- Standalone: Do not include `asyncio` standard library package automatically anymore.
- Avoid compilation of large generated codes in the `asynqua` package.
- Compile time optimize `pkg_resources.iter_entry_points` too, such that these can be used to resolve plugin modules, which helps with adding support for `agscheduler` package plugins. Note that these still need to be manually included with `--include-module` but now that works.
- For known truth values of the right hand side of `and` or `or` conditions, reduce the expression as far as possible.
- Added dedicated assignment node for hard imports, which then are propagated in classes as well, allowing for more static optimization for code on the class level.
- Added linker options to make static `--static-libpython` work with clang on Linux as well.
- macOS: Make sure `libpython` is loaded relative to the executable. This is needed for at least Anaconda Python.
- macOS: Fix, need to search environment specific DLL paths and only then global paths, otherwise mixed Python versions will not work correctly.
- Anti-Bloat: Remove IPython usage in `rich` package.
- Anti-Bloat: Avoid `doctest` dependency when using `pyrect`.
- Anti-Bloat: Some `unittest` removals from `pytorch` using libraries.
- Keep the Scons report items sorted, or else it varies for the hashing of dependencies with Python versions before 3.6, causing cache misses without need.

Organisational

- UI: Output the `.cmd` file created (if any) on Windows, e.g. when run in a virtualenv or for uninstalled Python versions, it will otherwise not run in accelerated mode, but previously the output suggested to run the executable directly.
- UI: Enhanced command line option description of `--include-plugin-directory` which is frequently misunderstood. That option barely does what people want it to do. Point them to using the other options that are easy to use and will work.
- UI: Specified needed Python version for use in `--python-for-scons` so users can know ahead of time what versions are suitable.
- Reports: Added information about data files including, optimization times per module, active plugins.
- Debugging: Repaired offline DLL dependency listing tool, such that it can be used during Windows DLL analysis.
- Make `--xml` accept a filename for the node tree dump, and change it so it can be executed in addition to actual compilation. This way we need not be super-robust about keeping stdout clean, to not break XML parsing.
- Plugins: Avoid useless warning about PySide2 plugin usage if another Qt plugin is actually selected.

- UI: Catch error of directories being used as data files where plain files are expected and point out that other options must be used.
- User Manual: Added section about accessing files in standalone mode too, so people can make sure it works properly.
- Onefile: Using `%TEMP%` folder should not by itself prevent cached onefile mode, only really variable paths should. People may want to have this as some kind of temporary cache still.
- UI: Catch user error of using elements, that resolve to absolute values in the middle of path specs, so using e.g. `something/%PROGRAM%` is now a mistake caught at compile time. These values can only be at the start of spec values naturally.
- Quality: Updated to newer version of `rstfmt`.
- UI: Nicer error message when a forbidden import is requested as an implicit import by a plugin.
- Python3.11: Adapted to allocator and exception state changes, but more will be needed to compile at all.
- Visual Code: Find `clang-format` from the recommended C++ extension of Visual Code, which makes it finally available on macOS easily too.
- UI: Quote command line argument values as necessary when stating them in the logging. Otherwise they are not directly usable on the shell and also less readable.
- Debian: Do not list fake modules as used debian packages codes, which could e.g. happen with the pre-load code of `pkg_resources` if that is from a Debian package. Fake packages should not be mentioned for these lists though.
- Nuitka-Python: Added support to set link time flags coming from statically included packages.
- For our `isort` trick of splitting files in two parts (mostly to setup import paths for `nuitka` package), make sure the second parts starts with a new line.
- Added more usable form `--output-filename` to specify the output filename, the short form has become barely usable after we switched to enforcing no space separation for command line arguments.
- UI: Check if output filename's directory exists ahead of time, and error exit if not, otherwise compilation crashed only in the very end, trying to create the final result.
- UI: When exiting with no error code, do not use red color or `FATAL` error annotation, that is not justified.
- Quality: Make sure the Yaml auto-format does not change effective contents.
- Quality: Added ability to limit autoformat by file type, which can be handy when e.g. only the yaml files should be scanned.
- UI: Removed `--windows-onefile-tempdir-spec` alias of `--onefile-tempdir-spec` option, it is no longer Windows specific.

Cleanups

- Prefer single quotes rather than double quotes in our package configuration Yaml files, otherwise esp. regular expressions with escapes become very confusing.
- Move import hacks to general mechanism in Yaml package configuration files. This is for extra paths from package names or from directory paths relative to the package. This removes special purpose code from core code paths and allows their re-use.
- Again more spelling cleanups have been done, to make the code cleaner to read and search.
- Unified how plugins treat iteration over their value list, and how the `when` condition is applied for the various kinds of sections.

- Output compilation command that failed during coverage taking, which makes it unnecessary to attempt to reconstruct what happened from test modes.

Tests

- Added coverage for comparisons that need argument swaps.
- Allow more time in onefile keyboard signal test, otherwise it can be a race on slow machines, e.g. emulated machines.
- Tests: Added support for running a local web server.

Summary

This release is mainly a consolidation of previous release. Optimization added in previous release did in fact introduce regressions, that needed to be addressed and were cause for relatively many hotfixes.

The Yaml nuitka package configuration feature is getting ever more powerful, but is not one bit more documented, such that the community as a whole is not yet capable of adding missing dependencies, data files, DLLs, and even anti-bloat patches.

New optimization was focused around compatibility with very few exceptions, where the non-automatic standard library work is standing out, and allows for smaller binaries in many cases.

Scalability has seen improvements through a few optimization, but mainly again with anti-bloat work being done. This is owed to the fact that consolidation was the name of the game.

For Anaconda specifically, a lot more software is covered, and generally, `cv2` and `torch` related tools are now working better, but it seems DLL handling will remain problematic in many instances.

The compilation report contains much more information and is getting there in terms of completeness. At some point, we should ask for it in bug reports.

Nuitka Release 1.0

This release contains a large amount of new features, while consolidating what we have with many bug fixes. Scalability should be dramatically better, as well as new optimization that will accelerate some code quite a bit. See the summary, how this release is paving the way forward.

Bug Fixes

- Python3: Fix, `bytes.decode` with only `errors` argument given was not working. Fixed in 0.9.1 already.
- MSYS2: Fix, the accelerate mode `.cmd` file was not working correctly. Fixed in 0.9.1 already.
- Onefile: Fix, the bootstrap when waiting for the child, didn't protect against signals that interrupt this call. This only affected users of the non-public `--onefile-tempdir` option on Linux, but with that becoming the default in 1.0, this was discovered. Fixed in 0.9.1 already.
- Fix, `pkg_resources` compile time generated `Distribution` values could cause issues with code that put it into calls, or in tried blocks. Fixed in 0.9.1 already.
- Standalone: Added implicit dependencies of `xlib` package. Fixed in 0.9.1 already.
- macOS: Fix, the package configuration for `wx` had become invalid when restructuring the Yaml with code and schema disagreeing on allowed values. Fixed in 0.9.1 already.
- Fix: The `str.format` with a single positional argument didn't generate proper code and failed to compile on the C level. Fixed in 0.9.1 already.
- Fix, the type shape of `str.count` result was wrong. Fixed in 0.9.1 already.

- UI: Fix, the warning about collision of just compiled package and original package in the same folder hiding the compiled package should not apply to packages without an `__init__.py` file, as those do **not** take precedence. Fixed in 0.9.2 already.
- Debugging: Fix, the fallback to `lldb` from `gdb` when using the option `--debugger` was broken on anything but Windows. Fixed in 0.9.2 already.
- Python3.8: The module `importlib.metadata` was not recognized before 3.9, but actually 3.8 already has it, causing the compile time resolution of package versions to not work there. Fixed in 0.9.3 already.
- Standalone: Fix, at least on macOS we should also scan from parent folders of DLLs, since they may contain sub-directories in their names. This is mostly the case, when using frameworks. Fixed in 0.9.2 already.
- Standalone: Added package configuration for `PyQt5` to require onefile bundle mode on macOS, and recommend to disable console for `PyQt6`. This is same as we already do for `PySide2` and `PySide6`. Fixed in 0.9.2 already.
- Standalone: Removed stray macOS onefile bundle package configuration for `pickle` module which must have been added in error. Fixed in 0.9.2 already.
- UI: Catch user error of attempting to compile the `__init__.py` rather than the package directory. Fixed in 0.9.2 already.
- Fix, hard name import nodes failed to clone, causing issues in optimization phase. Fixed in 0.9.2 already.
- Fix, avoid warnings given with gcc 11. Fixed in 0.9.2 already.
- Fix, dictionary nodes where the operation itself has no effect, e.g. `dict.copy` were not properly annotating that their dictionary argument could still cause a raise and have side effects, triggering an assertion violation in Nuitka. Fixed in 0.9.2 already.
- Standalone: Added `pynput` implicit dependencies on Linux. Fixed in 0.9.2 already.
- Fix, boolean condition checks on variables converted immutable constant value assignments to boolean values, leading to incorrect code execution. Fixed in 0.9.2 already.
- Python3.9: Fix, could crash on generic aliases with non-hashable values. Fixed in 0.9.3 already.

```
dict[str:any]
```

- Python3: Fix, an iteration over `sys.version_info` was falsely optimized into a tuple, which is not always compatible. Fixed in 0.9.3 already.
- Standalone: Added support for `xgboost` package. Fixed in 0.9.3 already.
- Standalone: Added data file for `text_unidecode` package. Fixed in 0.9.4 already.
- Standalone: Added data files for `swagger_ui_bundle` package. Fixed in 0.9.4 already.
- Standalone: Added data files for `connexion` package. Fixed in 0.9.4 already.
- Standalone: Added implicit dependencies for `sklearn.utils` and `rapidfuzz`. Fixed in 0.9.4 already.
- Python3.10: Fix, the reformulation of `match` statements could create nodes that are used twice, causing code generation to assert. Fixed in 0.9.4 already.
- Fix, module objects removed from `sys.modules` but still used could lack a reference to themselves, and therefore crash due to working on a released module variables dictionary. Fixed in 0.9.5 already.
- Fix, the MSVC compiles code generated for SciPy 1.8 wrongly. Added a workaround for that code to avoid triggering it. Fixed in 0.9.6 already.

- Fix, calls to `str.format` where the result is not used, could crash the compiler during code generation. Fixed in 0.9.6 already.
- Standalone: For DLLs on macOS and Anaconda, also consider the `lib` directory of the root environment, as some DLLs are otherwise not found.
- Fix, allow `nonlocal` and `global` for `__class__` to be used on the class level.
- Fix, `xrange` with large values didn't work on all platforms. This affected at least Python2 on macOS, but potentially others as well.
- Windows: When scanning for installed Pythons to e.g. run Scons or onefile compression, it was attempting to use installations that got deleted manually and could crash.
- macOS: Fix, DLL conflicts are now resolved by checking the version information too, also all cases that previously errored out after a conflict was reported, will now work.
- Fix, conditional expressions whose statically decided condition picking a branch will raise an exception could crash the compilation.

```
# Would previously crash Nuitka during optimization.
return 1/0 if os.name == "nt" else 1/0
```

- Windows: Make sure we set C level standard file handles too
At least newer subprocess was affected by this, being unable to provide working handles to child processes that pass their current handles through, and also this should help DLL code to use it as level.
- Standalone: Added support for `pyqtgraph` data files.
- Standalone: Added support for `dipy` by anti-bloat removal of its testing framework that wants to do unsupported stuff.
- UI: Could still give warnings about modules not being followed, where that was not true.
- Fix, `--include-module` was not working for non-automatic standard library paths.

New Features

- Onefile: Recognize a non-changing path from `--onefile-tempdir-spec` and then use cached mode. By default a temporary folder is used in the spec value, make it delete the files afterwards.
The cached mode is not necessarily faster, but it is not going to change files already there, leaving the binaries there intact. In the future it may also become faster to execute, but right now checking the validity of the file takes about as long as re-creating it, therefore no gain yet. The main point, is to not change where it runs from.
- Standalone: Added option to exclude DLLs. You can now use `--noinclude-dlls` to exclude DLLs by filename patterns.
They may e.g. come from Qt plugins, where you know, or experimented, that it is not going to be used in your specific application. Use with care, removing DLLs will lead to very hard to recognize errors.
- Anaconda: Use `CondaCC` from environment variables for Linux and macOS, in case it is installed. This can be done with e.g. `conda install gcc_linux-64` on Linux or `conda install clang_osx-64` on macOS.
- Added new option `--nowarn-mnemonic` to disable warnings that use mnemonics, there is currently not that many yet, but it's going to expand. You can use this to acknowledge the ones you accept, and not get that warning with the information pointer anymore.
- Added method for resolving DLL conflicts on macOS too. This is using version information and picks the newer one where possible.

- Added option `--user-package-configuration-file` for user provided Yaml files, which can be used to provide package configuration to Nuitka, to e.g. add DLLs, data files, do some anti-bloat work, or add missing dependencies locally. The documentation for this does not yet exist though, but Nuitka contains a Yaml schema in the `misc/nuitka-package-config-schema.json` file.
- Added `nuitka-project-else` to avoid repeating conditions in Nuitka project configuration, this can e.g. be used like this:

```
# nuitka-project-if: os.getenv("TEST_VARIANT", "pyside2") == "pyside2":
#   nuitka-project: --enable-plugin=no-qt
# nuitka-project-else:
#   nuitka-project: --enable-plugin=no-qt
#   nuitka-project: --noinclude-data-file=*.svg
```

Previously, the inverted condition had to be used in another `nuitka-project-if` which is no big deal, but less readable.

- Added support for deep copying uncompiled functions. There is now a section in the User Manual that explains how to clone compiled functions. This allows a workaround like this:

```
def binder(func, name):
    try:
        result = func.clone()
    except AttributeError:
        result = types.FunctionType(func.__code__, func.__globals__, name=func.__name__, argdefs=func.__defaults__, closure=func.__closure__)
        result = functools.update_wrapper(result, func)
        result.__kwdefaults__ = func.__kwdefaults__

    result.__name__ = name
    return result
```

- Plugins: Added explicit deprecation status of a plugin. We now have a few that do nothing, and are just there for compatibility with existing users, and this now informs the user properly rather than just saying it is not relevant.
- Fix, some Python installations crash when attempting to import modules, such as `os` with a `ModuleName` object, because we limit string operations done, and e.g. refuse to do `.startswith` which of course, other loaders that your installation has added, might still use.
- Windows: In case of not found DLLs, we can still examine the run time of the currently compiling Python process of Nuitka, and locate them that way, which helps for some Python configurations to support standalone, esp. to find CPython DLL in unusual spots.
- Debian: Workaround for `lib2to3` data files. These are from `stdlib` and therefore the patched code from Debian needs to be undone, to make these portable again.

Optimization

- Scalability: Avoid merge traces of initial variable versions, which came into play when merging a variable used in only one branch. These are useless and only made other optimization slower or impossible.
- Scalability: Also avoid merge traces of merge traces, instead flatten merge traces and avoid the duplication doing so. There were pathological cases, where this reduced optimization time for functions from infinite to instant.
- For comparison helpers, switch comparison where possible, such that there are only 3 variants, rather than 6. Instead the boolean result is inverted, e.g. changing `>=` into `not <` effectively. Of course this can only be done for types, where we know that nothing special, i.e. no method overloads of `__gte__` is going on.
- For binary operations that are commutative with the selected types, in mixed type cases, swap the arguments during code generation, such that e.g. `long_a + float_b` is actually computed as

`float_b + long_a`. This again avoids many helpers. It also can be done for `*` with integers and container types.

- In cases, where a comparison (or one of the few binary operation where we consider it useful), is used in a boolean context, but we know it is impossible to raise an exception, a C boolean result type is used rather than a `nuitka_bool` which is now only used when necessary, because it can indicate the exception result.
- Anti-Bloat: More anti-bloat work was done for popular packages, covering also uses of `setuptools_scm`, `nose` and `nose2` package removals and warnings. There was also a focus on making `mmvc`, `tensorflow` and `tiff file` compile well, removing e.g. the uses of the `tensorflow` testing framework.
- Faster comparison of `int` values with constant values, this uses helpers that work with C `long` values that represent a single "digit" of a value, or ones that use the full value space of C `long`.
- Faster comparison of `float` values with constant values, this uses helpers that work with C `float` values, avoiding the useless Python level constant objects.
- Python2: Comparison of `int` and `long` now has specialized helpers that avoids converting the `int` to a `long` through coercion. This takes advantage of code to compare C `long` values (which are at the core of Python2 `int` objects, with `long` objects).
- For binary operation on mixed types, e.g. `int * bytes` the slot of the first function was still considered, and called to give a `Py_NotImplemented` return value for no good reason. This also applies to mixed operations of `int`, `long`, and `float` types, and for `str` and `unicode` values on Python2.
- Added missing helper for `**` operation with floats, this had been overlooked so far.
- Added dedicated nodes for `ctypes.CDLL` which aims to allow us to detect used DLLs at compile time in the future, and to move closer to support its bindings more efficiently.
- Added specialized nodes for `dict.popitem` as well. With this, now all of the dictionary methods are specialized.
- Added specialized nodes for `str.expandtabs`, `str.translate`, `str.ljust`, `str.rjust`, `str.center`, `str.zfill`, and `str.splitlines`. While these are barely performance relevant, this completes all `str` methods, except `removeprefix` and `removesuffix` that are Python3.9 or higher.
- Added type shape for result of `str.index` operation as well, this was missing so far.
- Optimize `str`, `bytes` and `dict` method calls through variables.
- Optimize calls through variables containing e.g. mutable constant values, these will be rare, because they all become exceptions.
- Optimize calls through variables containing built-in values, unlocking optimization of such calls, where it is assigned to a local variable.
- For generated attribute nodes, avoid local doing import statements on the function level. While these were easier to generate, they can only be slow at run time.
- For the `str` built-in annotate its value as derived from `str`, which unfortunately does not allow much optimization, since that can still change many things, but it was still a missing attribute.
- For variable value release nodes, specialize them by value type as well, enhancing the scalability, because e.g. parameter variable specific tests, need not be considered for all other variable types as well.

Organisational

- Plugins: Major changes to the Yaml file content, cleaning up some of the DLL configuration to more easy to use.

The DLL configuration has two flavors, one from code and one from filename matching, and these got separated into distinct items in the Yaml configuration. Also how source and dest paths get provided got simplified, with a relative path now being used consistently and with sane defaults, deriving the destination path from where the module lives. Also what we called patterns, are actually prefixes, as there is still the platform specific DLL file naming appended.

- Plugins: Move mode checks to dedicated plugin called `options-nanny` that is always enabled, giving also much cleaner Yaml configuration with a new section added specifically for these. It controls advice on the optional or required use of `--disable-console` and the like. Some packages, e.g. `wx` are known to crash on macOS when the console is enabled, so this advice is now done with saner configuration.
- Plugins: Also for all Yaml configuration sub-items where is now a consistent `when` field, that allows checking Python version, OS, Nuitka modes such as standalone, and only apply configuration when matching this criterion, with that the anti-bloat options to allow certain bloat, should now have proper effect as well.
- The use of `AppImage` on Linux is no more. The performance for startup was always slower, while having lost the main benefit of avoiding IO at startup, due to new cached mode, so now we always use the same bootstrap binary as on macOS and Windows.
- UI: Do not display implicit reports reported by plugins by default anymore. These have become far too many, esp. with the recent `stdlib` work, and often do not add any value. The compilation report will become where to turn to find out why a module is included.
- UI: Ask the user to install the ordered set package that will actually work for the specific Python version, rather than making him try one of two, where sometimes only one can work, esp. with Python 3.10 allowing only one.
- GitHub: More clear wording in the issue template that `python -m nuitka --version` output is really required for support to be given.
- Attempt to use Anaconda `ccache` binary if installed on non-Windows. This is esp. handy on macOS, where it is harder to get it.
- Windows: Avoid byte-compiling the inline copy of Scons that uses Python3 when installing for Python2.
- Added experimental switches to disable certain optimization in order to try out their impact, e.g. on corruption bugs.
- Reports: Added included DLLs for standalone mode to compilation report.
- Reports: Added control tags influencing plugin decisions to the compilation report.
- Plugins: Make the `implicit-imports` dependency section in the Yaml package configuration a list, for consistency with other blocks.
- Plugins: Added checking of tags such from the package configuration, so that for things dependent on python version (e.g. `python39_or_higher`, `before_python39`), the usage of Anaconda (`anaconda`) or certain OS (e.g. `macos`), or modes (e.g. `standalone`), expressions in `when` can limit a configuration item.
- Quality: Re-enabled string normalization from black, the issues with changes that are breaking to Python2 have been worked around.
- User Manual: Describe using a minimal virtualenv as a possible help low memory situations as well.
- Quality: The yaml auto-format now properly preserves comments, being based on `ruamel.yaml`.
- Nuitka-Python: Added support for the Linux build with Nuitka-Python for our own CPython fork as well, previously only Windows was working, and macOS will follow later.

- The commit hook when installed from git bash was working, but doing so from `cmd.exe` didn't find a proper path for shell from the `git` location.
- Debugging: A lot of experimental toggles were added, that allow control over the use of certain optimization, e.g. use of dict, list, iterators, subscripts, etc. internals, to aid in debugging in situations where it's not clear, if these are causing the issue or not.
- Added support for Fedora 36, which requires some specific linker options, also recognize Fedora based distributions as such.
- Removed long deprecated option `--noinclude-matplotlib` from numpy plugin, as it hasn't had an effect for a long time now.
- Visual Code: Added extension for editing Jinja2 templates. This one even detects that we are editing C or Python and properly highlights accordingly.

Cleanups

- Standalone: Major cleanup of the dependency analysis for standalone. There is no longer a distinction between entry points (main binary, extension modules) and DLLs that they depend on. The OS specific parts got broken out into dedicated modules as well and decisions are now taken immediately.
- Plugins: Split the Yaml package configuration files into 3 files. One contains now Python2 only stdlib configuration, and another one general stdlib.
- Plugins: Also cleanup the `zmq` plugin, which was one the last holdouts of now removed plugin method, moving parts to the Yaml configuration. We therefore no longer have `considerExtraDlls` which used to work on the standalone folder, but instead only plugin code that provides included DLL or binary objects from `getExtraDlls` which gives Nuitka much needed control over DLL copying. This was a long lasting battle finally won, and will allow many new features to come.
- UI: Avoid changing whitespace in warnings, where we have intended line breaks, e.g. in case of duplicate DLLs. Went over all warnings and made sure to either avoid new-lines or have them, depending on wanted output.
- Iterator end check code now uses the same code as rich comparison expressions and can benefit from optimization being done there as well.
- Solved TODO item about code generation time C types to specify if they have error checking or not, rather than hard coding it.
- Production of binary helper function set was cleaned up massively, but still needs more work, comparison helper function set was also redesigned.
- Changing the spelling of our container package to become more clear.
- Used `namedtuple` objects for storing used DLL information for more clear code.
- Added spellchecker ignores for all attribute and argument names of generated fixed attribute nodes.
- In auto-format make sure the imports float to the top. That very much cleans up generated attribute nodes code, allowing also to combine the many ones it makes, but also cleans up some of our existing code.
- The package configuration Yaml files are now sorted according to module names. This will help to avoid merge conflicts during hotfixes merge back to develop and automatically group related entries in a sane way.
- Moved large amounts of code producing implicit imports to Yaml configuration files.
- Changed the `tensorflow` plugin to Yaml based configuration, making it a deprecated do nothing plugin, that only remains there for a few releases, to not crash existing build scripts.

- Lots of spelling cleanups, e.g. renaming `nuitka.codegen` to `nuitka.code_generation` for clarity.

Tests

- Added generated test to cover `bytes` method. This would have found the issue with `decode` potentially.
- Enhanced standalone test for `ctypes` on Linux to actually have something to test.

Summary

This release improves on many things at once. A lot of work has been put into polishing the Yaml configuration that now only lacks documentation and examples, such that the community as a whole should become capable of adding missing dependencies, data files, DLLs, and even anti-bloat patches.

Then a lot of new optimization has been done, to close the missing gaps with `dict` and `str` methods, but before completing `list` which is already a work in progress pull request, and `bytes`, we want to start and generate the node classes that form the link or basis of dedicated nodes. This will be an area to work on more.

The many improvements to existing code helpers, and them being able to pick target types for the arguments of comparisons and binary operations, is a pre-cursor to universal optimization of this kind. What is currently only done for constant values, will in the future be interesting for picking specific C types for use. That will then be a huge difference from what we are doing now, where most things still have to use `PyObject *` based types.

Scalability has again seen very real improvements, memory usage of Nuitka itself, as well as compile time inside Nuitka are down by a lot for some cases, very noticeable. There is never enough of this, but it appears, in many cases now, large compilations run much faster.

For macOS specifically, the new DLL dependency analysis, is much more capable or resolving conflicts all by itself. Many of the more complex packages with some variants of Python, specifically Anaconda will now be working a lot better.

And then, of course there is the big improvement for Onefile, that allows to use cached paths. This will make it more usable in the general case, e.g. where the firewall of Windows hate binaries that change their path each time they run.

Future directions will aim to make the compilation report more concise, and given reasons and dependencies as they are known on the inside more clearly, such that it can be a major tool for testing, bug reporting and analysis of the compilation result.

Nuitka Release 0.9

This release has a many optimization improvements, and scalability improvements, while also adding new features, with also some important bug fixes.

Bug Fixes

- Fix, hard module name lookups leaked a reference to that object. Fixed in 0.8.1 already.
- Python2: Fix, `str.decode` with `errors` as the only argument wasn't working. Fixed in 0.8.1 already.
- Fix, could corrupt created uncompiled class objects `__init__` functions in case of descriptors being used.
- Standalone: Added support for newer `torch`. Fixed in 0.8.1 already.
- Standalone: Added support for newer `torchvision`. Fixed in 0.8.1 already.

- Fix, could compile time crash during initial parsing phase on constant dictionary literals with non-hashable keys.

```
{ { }:1, }
```

- Fix, hard imported sub-modules of hard imports were falsely resolved to their parent. Fixed in 0.8.3 already.

```
importlib.resources.read_bytes() # gave importlib has no attribute...
```

- Windows: Fix, outputs with `--force-stdout-spec` or `--force-stderr-spec` were created with the file system encoding on Python3, but they need to be `utf-8`.
- Fix, didn't allow zero spaces in Nuitka project options, which is not expected.
- Modules: Fix, the `del __file__` in the top level module in module mode caused crashes at run time, when trying to restore the original `__file__` value, after the loading CPython corrupted it.
- Python2.6: Fixes for installations without `pkg_resources`.
- Windows: Fix for very old Python 2.6 versions, these didn't have a language assigned that could be used.
- Security: Fix for [CVE-2022-2054](#) where environment variables used for transfer of information between Nuitka restarting itself, could be used to execute arbitrary code at compile time.
- Anaconda: Fix, the torch plugin was not working on Linux due to missing DLL dependencies.
- Fix, static optimization of `importlib.import_module` with a package given, for an absolute import was optimized into the wrong import, package was not ignored as it should be.
- Windows: Installed Python scan could crash on trying installation paths from registry that were manually removed in the mean time, but not through an uninstaller.
- Standalone: Added missing implicit dependency for `pyreadstat` because parts of standard library it uses are no more automatically included.
- Windows: Could still crash when no `powershell` is available with symlinks, handle this more gracefully.
- Standalone: Added more missing Plotly dependencies, but more work will be needed to complete this.
- Standalone: Add missing `stdlib` dependency on `multiprocessing` by `concurrent.futures.process`.
- Standalone: Fix, implicit dependencies assigned to `imageio` on PIL plugins should actually be assigned to `PIL.Image` that actually loads them, so it works outside of `imageio` too.

New Features

- UI: Added new option `--user-package-configuration-file` to allow users to provide extra Yaml configuration files for the Nuitka plugin mechanism to add hidden dependencies, anti-bloat, or data files, for packages. This will be useful for developing PRs to the standard file of Nuitka. Currently the schema is available, but it is not documented very well yet, so not really ready for end users just yet.
- Standalone: Added new `no-qt` plugin as an easy way to prevent all of the Qt bindings from being included in a compilation.
- Include module search path in compilation report.

Optimization

- Faster dictionary iteration with our own replacement for `PyDict_Next` that avoids the DLL call overhead (in case of non-static libpython) and does less unnecessary checks.
- Added optimization for `str.count` and `str.format` methods as well, this should help in some cases with compile time optimization.
- The node for `dict.update` with only an iterable argument, but no keyword arguments, was in fact unused due to wrongly generated code. Also the form with no arguments wasn't yet handled properly.
- Scalability: Use specialized nodes for pair values, i.e. the representation of `x = y` in e.g. dictionary creations. With constant keys, and values, these avoid full constant value nodes, and therefore save memory and compile time for a lot of code.
- Anti-Bloat: Added more scalability work to avoid including modules that make compilation unnecessarily big.
- Python3.9+: Faster calls in case of mixed code, i.e. compiled code calling uncompiled code.
- Removing duplicates and non-existent entries from modules search path should improve performance when locating modules.
- Optimize calls through variables as well. These are needed for the package resource nodes to properly resolve at compile time from their hard imports to the called function.
- Hard imported names should also be considered very trusted themselves, so they are e.g. also optimized in calls.
- Anti-Bloat: Avoid more useless imports in Pandas, Numba, Plotly, and other packages, improving the scalability some more.
- Added dedicated nodes for `pkg_resources.require`, `pkg_resources.get_distribution`, `importlib.metadata.version`, and `importlib_metadata.version`, so we can use compile time optimization to resolve their argument values where possible.
- Avoid annotating control flow escape for all release statements. Sometimes we can tell that `__del__` will not execute outside code ever, so this then avoids marking values as escaped, and taking the time to do so.
- Calls of methods through variables on `str`, `dict`, `bytes` that have dedicated nodes are now also optimized through variables.
- Boolean tests through variables now also are optimized when the original assignment is a compile time constant that is not mutable. This is only basic, but will allow tests on `TYPE_CHECKING` coming from a `from typing import TYPE_CHECKING` statement to be optimized, avoiding this overhead.

Cleanups

- Changed to `torch` plugin to Yaml based configuration, making it obsolete, it only remains there for a few releases, to not crash existing build scripts.
- Moved older package specific hacks to the Yaml file. Some of these were from hotfixes where the Yaml file wasn't yet used by default, but now there is no need for them anymore.
- Removed most of the `pkg-resources` plugin work. This is now done during optimization phase and rather than being based on source code matches, it uses actual value tracing, so it immediately covers many more cases.
- Continued spelling improvements, renaming identifiers used in the source that the `cspell` based extension doesn't like. This aims at producing more readable and searchable code.
- Generated attribute nodes no longer do local imports of the operation nodes they refer to. This also avoids compile time penalties during optimization that are not necessary.

- Windows: Avoid useless bytecode of inline copy used by Python3 when installing for Python2, this spams just a lot of errors.

Organisational

- Removed MSI installers from the download page. The MSI installers are discontinued as Python has deprecated their support for them, as well as Windows 10 is making it harder for users to install them. Using the PyPI installation is recommended on Windows.
- Merged our Yaml files into one and added schema description, for completion and checking in Visual Code while editing. Also check the schema in `check-nuitka-with-yamllint` which is now slightly misnamed. The schema is in no way final and will see improvements in future releases.
- UI: Nicer progress bar layout that avoids flicker when optimizing modules.
- UI: When linking, output the total number of object files used, to have that knowledge after the progress bar for C compilation is gone.
- Quality: Auto-format the package configuration Yaml file for anti-bloat, implicit dependencies, etc.
- GitHub: Point out the commit hook in the PR template.
- UI: Nicer output in case of no commercial version is used.
- Updated the MinGW64 winlibs download used on Windows to the latest version based on gcc 11, the gcc 12 is not yet ready.
- Git: Make sure we are not affected by `core.autocrlf` setting, as it interferes with auto-format enforcing Unix newlines.
- Removed the MSI downloads. Windows 10 has made them harder to install and Python itself is discontinuing support for them, while often it was only used by beginners, for which it was not intended.
- Anaconda: Make it more clear how to install static libpython with precise command.
- UI: Warn about using Debian package contents. These can be non-portable to other OSes.
- Quality: The auto-format now floats imports to the top for consistency. With few exceptions, it was already done like this. But it makes things easier for generated code.

Tests

- The reflected test was adapted to preserve `PYTHONPATH` now that module presence influences optimization.

Summary

This release marks a point, that will allow us to open up the compatibility work for implicit dependencies and anti-bloat stuff even further. The Yaml format will need documentation and potentially more refinement, but will open up a future, where latest packages can be supported with just updating this configuration.

The scalability improvements really make a difference for many libraries and are a welcome improvement on both memory usage and compile time. They are achieved by an accord of static optimization of

One optimization aimed at optimizing tuple unpacking, was not finished in time for this release, but will be subject of a future release. It has driven many other improvements though.

Generally, also from the UI, this is a huge step forward. With links to the website for complex topics being started, and the progress bar flicker being removed, the tool has yet again become more user friendly.

Nuitka Release 0.8

This release has a massive amount of bug fixes, builds on existing features, and adds new ones.

Bug Fixes

- Windows: Fix, need to ignore cases where shorting in path for external use during compilation gives an permission error. Fixed in 0.7.1 already.
- Compatibility: Added workaround for `scipy.stats` function copying. Fixed in 0.7.1 already.
- Windows: Fix, detect ARM64 arch of MSVC properly, such that we can give a proper mismatch for the Python architecture. Fixed in 0.7.1 already.
- Standalone: Fix, using `importlib.metadata` module failed to include `email` from standard library parts no longer included by default. Fixed in 0.7.1 already.
- macOS: Fix, the dependency parser was using normalized paths where original paths must be used. Fixed in 0.7.1 already.
- Standalone: Fix, using `shiboken6` module (mostly due to `PySide6`) failed to include `argparse` from the standard library from standard library parts no longer included by default. Fixed in 0.7.1 already.
- Onefile: Fix, the detection of a usable Python for compression could crash. Fixed in 0.7.2 already.
- Onefile: Adding the payload on Windows could run into locks still being held, need to wait in that case. Fixed in 0.7.2 already.
- Fix, need to include `pkg_resources` as well, we need it for when we use Jinja2, which is more often now. For Python3 this was fixed in 0.7.3 already. Later a version to use with Python2 was added as well.
- Release: The wheels built for Nuitka when installed through URLs were not version specific, but due to different inline copies per OS and Python version, they must not be reused. Therefore we now pretend to contain an extension module, which handles that. Fixed in 0.7.3 already.
- Standalone: Fix, using `urllib.requests` module failed to include `http.client` from standard library parts no longer included by default. Fixed in 0.7.3 already. Later `http.cookiejar` was added too.
- Standalone: Do not compress MSVC run time library when using `upx` plugin, that is not going to work. Fixed in 0.7.4 already.
- Standalone: Fix, on Windows more files should be included for TkInter to work with all software. Fixed in 0.7.5 already.
- Distutils: Added support for `package_dir` directive to specify where source code lives. Fixed in 0.7.6 already.
- Standalone: Fix, using `shelve` module failed to include `dbm` from standard library parts no longer included by default. Fixed in 0.7.6 already.
- Standalone: Added support for `arcade` data files. Fixed in 0.7.7 already.
- Standalone: Fix, bytecode demotions should use relative filenames rather than original ones. Fixed in 0.7.7 already.
- Standalone: Fix, must remove extension module objects from `sys.modules` before executing an extension module that will create it. This fixes cases of cyclic dependencies from modules loaded by the extension module.
- Windows: In case of an exception, `clcache` was itself triggering one during its handling, hiding the real exception behind a `TypeError`.

- Windows: Improved `clcache` locking behavior, avoiding a race. Also increase characters used for key from 2 to 3 chars, making collisions far more rare.
- Standalone: Added support for `persistent` package.
- Standalone: Added support for newer `tensorflow` package.
- Module: Fix, need to restore the `__file__` and `__spec__` values of top level module. It is changed by CPython after import to an incompatible file name, and not our loader, preventing package resources to be found.
- Standalone: Added support for `Crpytodome.Cipher.PKCS1_v1_5`.
- Fix, `pkgutil.iter_modules` without arguments was not working properly with our meta path based loader.
- Windows: Fix, could crash on when the Scons report was written due to directories in `PATH` that failed to encode.
- Compatibility: Fix, positive `divmod` and modulo `%` with negative remainders of positive floats was not correct.
- Fix, `str.encode` with only `errors` value, but default value for encoding was crashing the compilation.
- Python3.10+: Fix, `match` statement sliced values must be lists, not tuples.
- Standalone: Added support for newer `glfw` and `OpenGL` packages.
- Python3: Fix, failed to read bytecode only `stdlib` files. This affect mostly Fedora Python which does this for encodings.
- Python3.5+: Fix, two phase loading of modules could release it immediately.
- Standalone: Added missing dependency for `pydantic`.
- Fix, the `str.split` rejected default `sep` value with only `maxsplit` given as a keyword argument.
- Standalone: Added missing dependency of `wsgiref` module.
- Standalone: Added support for `falcon` module.
- Standalone: Added support for `eliot` module.
- Fix, need to mark assigned from variables as escaped. Without it, some aliased loop variables could be misunderstood and falsely statically optimized.
- Standalone: Added support for newer `uvicorn` package.
- Standalone: Added data files for the `accessible_output2`, `babel`, `frozendict`, and `sound_lib` package.
- Standalone: Added support for newer `sklearn` package.
- Standalone: Added support for `tkinterdnd2` package.
- Python3.7+: Fix, the error message wasn't fully compatible for unsubscriptable type exception messages.
- Standalone: Fix, `idlelib` from `stdlib` was always ignored.
- Python3.4+: Fix, the `__spec__.origin` as produced by `find_spec` of our meta path based loader, didn't have the correct `origin` attribute value.
- Standalone: Disable QtPDF plugin of PySide 6.3.0, because it's failing dependency checks. On macOS this was blocking, we will change it to detection if that is necessary in a future release.
- Standalone: Added support for `orderedmultidict`.
- Standalone: Added support for `clr` module.

- Standalone: Added support for newer `cv2` package.

New Features

- Added new option `--module-name-choice` to select what value `__name__` and `__package__` are going to be. With `--module-name-choice=runtime` (default for `--module` mode), the created module uses the parent package to deduce the value of `__package__`, to be fully compatible. The value `--module-name-choice=original` (default for other modes) allows for more static optimization to happen.
- Added support for `get_resource_reader` to our meta path based loader. This allows to avoid useless temporary files in case `importlib.resources.path` is used, due to a bad interaction with the fallback implementation used without it.
- Added support for `--force-stdout-spec`` and ``--force-stderr-spec` on all platforms, this was previously limited to Windows.
- Added support for requiring and suggesting modes. In part this was added to 0.7.3 already, and is used e.g. to enforce that on macOS the `wx` will only work as a GUI program and crash unless `--disable-console` is specified. These will warn the user or outright error the compilation if something is known to be needed or useful.
- Debian: Detect version information for "Debian Sid". Added in 0.7.4 already, and also improved how Debian/Ubuntu versions are output.
- Added new option `--noinclude-data-files` to instruct Nuitka to not include data files matching patterns given. Also attached loggers and tags to included data file and include them in the compilation report.
- Distutils: When using `pyproject.toml` without `setup.py` so far it was not possible to pass arguments. This is now possible by adding a section like this.

```
[nuitka]
# options without an argument are passed as boolean value
show-scons = true

# options with single values, e.g. enable a plugin of Nuitka
enable-plugin = pyside2

# options with several values, e.g. avoiding including modules, accepts
# list argument.
nofollow-import-to = ["*.tests", "*.distutils"]
```

The option names are the same, but without leading dashes. Lists are only needed when passing multiple values with the same option.

- macOS: Add support for specifying signing identity with `--macos-sign-identity` and access to protected resources `--macos-app-protected-resource`.
- Included data files are now reported in the compilation report XML as well.
- Accept absolute paths for searching paths of binaries. This allows e.g. the `upx` plugin to accept both a folder path and the full path including the binary name to work when you specify the binary location with `--upx-binary` making it more user friendly.
- Python3.10: Added support for positional matching of classes, so far only keyword matching was working.
- Added support for path spec values `%CACHE_DIR`, `%COMPANY%`, `%PRODUCT%`, `%VERSION%`, and `%HOME` in preparation of onefile once again being able to be cached and not unpacked repeatedly for each execution.

- Standalone: Detect missing `tk-inter` plugin at run time. When TCL fails to load, it then outputs a more helpful error. This ought to be done for all plugins, where it's not clear if they are needed.
- Anti-Bloat: Added support for plain replacements in the `anti-bloat.yml` file. Before with `replacement``, the new value had to be produced by an `eval`, which makes for less readable values due to extra quoting. for plain values.

Optimization

- Python3.10+: Added support for `ordered-set` PyPI package to speed up compilation on these versions too, adding a warning if no accelerated form of `OrderedSet` is used, but believed to be usable.
- Optimization: Added `bytes.decode` operations. This is only a start and we needed this for internal usage, more should follow later.
- Much more `anti-bloat` work was added. Avoiding `ipython`, `unittest`, and sometimes even `doctest` usage for some more packages.
- The `ccache` was not always used, sometimes it believed to catch a potential race, that we need to tell it to ignore. This will speed up re-compilation of the C side in many cases.
- Do not compile the meta path based loader separate, which allows us to not expose functions and values only used by it. Also spares the C compiler one file.
- Added various dedicated nodes for querying package resources data, e.g. `pkgutil.get_data`. This will make it easier to detect cases of missing data files in the future.
- Added more hard imports, some of which help scalability in the compilation, because these are then known to exist in standalone mode, others are used for package resource specific operations.
- Onefile: Releasing decompression buffers avoiding unnecessary high memory usage of bootstrap binary.
- Standalone: Avoid proving directories with no DLLs (e.g. from packages) towards `ldd`, this should avoid exceeding command line limits.
- For `clcache` remove writing of the stats file before Scons has completed, which avoids IO and locking churn.
- Standalone: Avoid including `wsgiref` from `stdlib` by default.

Cleanups

- Removed references to `chrrpath` and dead code around it, it was still listed as a dependency, although we stopped using it a while ago.
- Removed uses of `anti-bloat` in examples and tests, it is now enabled by default.
- Made standard plugin file naming consistent, their name should be `*Plugin.py`.
- Cleaned up `tensorflow` plugin. The source modification was moved to `anti-bloat` where it is easy to do. The implicit dependencies are now in the config file of `implicit-imports` plugin.
- Massive cleanups of data file handling in plugins. Adding methods for producing the now required objects.
- The Scons file handling got further cleaned up and unified, doing more things in common code.
- Avoid `#ifdefs` usages with new helper function `Nuitka_String_FromFormat` that implies them for more readable code.
- Split the allowance check from the encountering. Allow plugins and options all to say if an import should be followed, and only when that is decided, to complain about it. Previously the attempt was causing an error, even if another plugin were to decide against it later.

- Python2.6: Avoid warnings from MSVC for out specialized `long` code. In testing it worked correctly, but this is more explicit and doesn't rely on C implementation specific behavior, although it appears to be universal.

Organisational

- UI: Warning tests are now wrapped to multiple lines if necessary. That makes it more accessible for larger messages that contain more guiding information.
- Documented how to use local Nuitka checkout with `pyproject.toml` files, this makes debugging Nuitka straightforward in these setups.
- Added instructions on how to pass extra C and linker flags and to the User Manual.
- Made our auto-format usable for the Nuitka website code too.
- Removed dependencies on `chrpath` and the now dead code that would use it, we are happy with `patchelf`.
- Updated to latest versions of requirements for development, esp. `black` and `pylint`.
- Renamed `--macos-onefile-icon` to `--macos-app-icon` because that is what it is really used for.
- Unified how dependencies are installed in GitHub actions.
- Updated man page contents for option name changes from last releases.
- Updated the MinGW64 winlibs download used on Windows to the latest version.
- Updated the `ccache` binary used on Windows with MinGW64. This is in preparation of using it potentially for MSVC as well.
- Updated Visual Code C config to use Python3.10 and MSVC 2022 include paths.

Tests

- Better outputs from standalone library compilation test, esp. when finding a problem, present the script to reproduce it immediately.
- Enhanced generated tests to cover `str` methods to use keyword arguments.
- Added automatic execution of `pyproject.toml` driven test case.
- Enhanced output in case of `optimization` test failures, dumping what value is there that has not become a compile time constant.

Summary

This release has seen a lot of consolidation. The plugins layer for data files is now all powerful, allowing much nicer handling of them by the plugins, they are better reported in normal output, and they are also part of the report file that Nuitka can create. You may now also inhibit their inclusion from the command line, if you decide otherwise.

The `pyproject.toml` now supporting Nuitka arguments is closing an important gap.

Generally many features got more polished, e.g. non-automatic inclusion of `stdlib` modules has most problems fixed up.

An important area of improvement, are the hard imports. These will be used to replace the source based resolution of package requirements with ones that are proper nodes in the tree. Getting these hard imports to still retain full compatibility with existing imports, that are more or less `__import__` uses only, was revealing quite a bit of technical debt, that has been addressed with this release.

For onefile, the cached mode is being prepared with the variables added, but will only be in a later release.

Also a bunch of new or upgraded packages are working now, and the push for `anti-bloat` work has increased, making many compilations even more lean, but scalability is still an issue.

Nuitka Release 0.7

This release is massively improving macOS support, esp. for M1 and the latest OS releases, but it also has massive improvements for usability and bug fixes in all areas.

Bug Fixes

- Fix, `set` creation wasn't annotating its possible exception exit from hashing values and is not as free of side effects as `list` and `tuple` creations are. Fixed in 0.6.19.1 already.
- Windows: Fix, `--experimental` option values got lost for the C compilation when switching from MSVC to MinGW64, making them have no effect. Fixed in 0.6.19.1 already.
- Windows: Fix, Clang from MinGW64 doesn't support LTO at this time, therefore default to `no` for it. Fixed in 0.6.19.1 already.
- Debian: Fix, failed to detect Debian unstable as suitable for linking, it doesn't have the release number. Fixed in 0.6.19.1 already.
- Standalone: Added data files for `pygsheets` package. Fixed in 0.6.19.2 already.
- Fix, paths from plugin related file paths need to be made absolute before used internally, otherwise the cache can fail to deduplicate them. Fixed in 0.6.19.2 already.
- Python3: With gcc before version 5, e.g. on CentOS 7, where we switch to using g++ instead, the gcc version checks could crash. Fixed in 0.6.19.2 already.
- Windows: Disable MinGW64 wildcard expansion for command line arguments. This was breaking command lines with arguments like `--filename *.txt`, which under `cmd.exe` are left alone by the shell, and are to be expanded by the program. Fixed in 0.6.19.2 already.
- Standalone: Added missing implicit dependency needed for `--follow-stdlib` with Python for some uses of the `locale` module. Fixed in 0.6.19.2 already.
- Standalone: Added workarounds for newest `numpy` that wants to set `__code__` objects and required improvements for macOS library handling. Fixed in 0.6.19.3 already.
- Windows: Caching of DLL dependencies for the main programs was not really working, requiring to detect them anew for every standalone compilation for no good reason. Fixed in 0.6.19.3 already.
- Windows: Fix, CTRL-C from a terminal was not propagated to child processes on Windows. Fixed in 0.6.19.4 already.
- Standalone: With `certifi` and Python3.10 the `importlib.resource` could trigger Virus scanner inflicted file access errors. Fixed in 0.6.19.4 already.
- Python3.10: Reverted error back iteration past end of generator change for Python 3.10.2 or higher to become compatible with that too. Fixed in 0.6.19.5 already.
- Standalone: Added support for `anyio` and by proxy for Solana. Fixed in 0.6.19.5 already.
- Fix, compilation with resource mode `incbin` and `--debugger` was not working together. Fixed in 0.6.19.5 already.
- Fix, format optimization of known `str` objects was not properly annotating an exception exit when being optimized away, causing consistency checks to complain. Fixed in 0.6.19.5 already.
- Windows: Fix, `clcache` didn't work for non-standard encoding source paths due to using the direct mode, where wrong filenames are output by MSVC. Fixed in 0.6.19.5 already.
- Windows: Fix, `ccache` cannot handle source code paths for non-standard encoding source paths. Fixed in 0.6.19.5 already.

- Python2.6: Fix, calls to `iteritems` and `iterkeys` on known dictionary values could give wrong values. Fixed in 0.6.19.5 already.
- Fix, the value of `__module__` if set by the metaclass was overwritten when creating types. Fixed in 0.6.19.6 already.
- Plugins: Add support for the latest version of `pkg_resources` that has "vendored" even more packages. Fixed in 0.6.19.6 already.
- Onefile: The onefile binary was locked during run time and could not be renamed, preventing in-place updates. This has been resolved and now on Windows, the standard trick for updating a running binary of renaming it, then placing the new file works.
- Fix, wasn't checking the `zstandard` version and as a result could crash if too old versions of it. This is now checked.
- macOS: Large amounts of bug fixes for the dependency scanner. It got cleaned up and now handles many more cases correctly.
- Windows: Fix, was not properly detecting wrong ClangCL architecture mismatch with the Python architecture. This could result in strange errors during C compilation in this setup.
- Standalone: Added implicit dependencies for the `asynccpg` module.
- Linux: Detect Debian or Ubuntu base and distribution name more reliably. This helps esp. with static libpython optimization being recognized automatically.

New Features

- We now disallow options that take arguments to be provided without using `=`.
- Previously `--lto no` worked just as well as `--lto=no` did. And that was the cause of problems when `--lto` first became a choice.
Recently similar, but worse problems were observed, where e.g. `--include-module` could swallow trailing other arguments when users forgot to specify the name by accident. Therefore this style of giving options is now explicitly rejected.
- Compiled types of Nuitka now inherit from uncompiled types. This should allow easier and more complete compatibility, making even code in extension modules that uses `PyObject_IsInstance` work, e.g. `pydantic`.
- macOS: Added signing of application bundles and standalone binaries for deployment to newer macOS platforms and esp. M1 where these are mandatory for execution.
- macOS: Added support for selecting the single macOS target arch to create a binary for. The `universal` architecture is not yet supported though, but will be added in a future release.
- Added support for compression in onefile mode through the use of an other Python installation, that has the `zstandard` module installed. With this it will work with 2.6 or higher, but require a 3.5 or higher Python with it installed in either `PATH` or on Windows in the registry alternatively.
- Added UPX plugin to compress created extension modules and binaries and for standalone mode, the included DLLs. For onefile, the compression is not useful since it has the payload already compressed.
- Added a more explicit way to list usable MSVC versions with `--msvc=list` rather than requiring an invalid value. Check values given in the same way that Scons will do.
- Added support for `--python-flag=-u` which disabled outputs buffers, so that these outputs are written immediately.
- Plugins: Always on plugins now can have command line options. We want this for the `anti-bloat` plugin that is enabled by default in this release.

- Plugins: Added ability for plugin to provide fake dependencies for a module. We want this for the `multiprocessing` plugin, that is now enabled by default in this release too.
- Plugins: Added ability for plugins to modify DLLs after copy for standalone. We will be using this in the new `upx` plugin.
- Added retry for file copies that fail due to still running program. This can happen on Windows with DLLs in standalone mode. For interactive compilation, this allows a retry to happen after prompting the user.
- UI: Added ability to list MSVC versions with `--msvc=list`, and detect illegal values given to `--msvc=` before Scons sees them, it also crashes with a relative unhelpful error message.
- UI: When linking, close the C compilation progress bar and state that that linking is going on. For some very large LTO compilations, it was otherwise at 100% and still taking a long time, confusing users.
- Plugins: Added new plugin that is designed to handle DLL dependencies through a configuration file that can both handle filename patterns as well as code provided DLL locations.
- Optimization: Exclude parts of the standard library by default. This allows for much smaller standalone distributions on modules, that can be expected to never be an implicit dependency of anything, e.g. `argparse` or `pydoc`.

Optimization

- Standalone: Do not include `encodings.bz2_codec` and `encodings.idna` anymore, these are not file system encodings, but require extension modules.
- Make sure we use proper `(void)` arguments for C functions without arguments, as for C functions, that makes a real difference, they are variable args functions and more expensive to call otherwise.
- For standalone, default to using `--python-flag=no_site` to avoid the overhead that the typically unused `site` module incurs. It often includes large parts of the standard library, which we now want to be more selective about. There is new Python flag added called `--python-flag=site` that restores the inclusion of `site` module.
- Standalone: Exclude non-critical codec modules from being technical, i.e. have to be available at program startup. This removes the need for e.g. `bz2` related extension modules previously included.
- In reformulations, use dictionary methods directly, we have since introduced dictionary specific methods, and avoid the unnecessary churn during optimization.
- The complex call helper could trigger unnecessary passes in some cases. The pure functions were immediately optimized, but usages in other modules inside loops sometimes left them in incomplete states.
- Windows: Avoid repeated hashing of the same files over and over for `clcache`.
- Cache dependencies of bytecode demoted modules in first compile and reuse that information in subsequent compilations.
- Linux: Added option for switching compression method for onefile created with `AppImage`. The default is also now `gzip` and not `xz` which has been observed to cause much slower startup for little size gains.
- Standalone: For failed relative imports, during compiled time absolute imports were attempted still and included if successful, the imports would not be used at run time, but lead to more modules being included than necessary.

Organisational

- There is now a [Discord server for Nuitka community](#) where you can hang out with the developers and ask questions. It is mirrored with the Gitter community chat, but offers more features.
- The `anti-bloat` is now on by default. It helps scalability by changing popular packages to not provide test frameworks, installation tools etc. in the resulting binary. This oftentimes reduces the compilation by thousands of modules.
- Also the `multiprocessing` plugin is now on by default. Detecting its need automatically removes a source of problems for first time users, that didn't know to enable it, but had all kinds of strange crashes from multiprocessing malfunctioning. This should enhance the out of the box experience by a lot.
- With this release, the version numbering scheme will be changed. For a long time we have used 4 digits, where one is a leading zero. That was initially done to indicate that it's not yet ready. However, that is just untrue these days. Therefore, we switch to 3 digits, and a first hotfix with now be 0.7.1 rather than 0.6.19.1, which is too long.

It has been observed that people disregard differences in the third digit, but actually for Nuitka these have oftentimes been very important updates. This change is to rectify it, and a new release will be 0.8, and there will be a 1.0 release after 0.9.

- Added a new section to User Manual that explains how to manually load files, such that it is cleaner and compatible code. Using paths relative to current directory is not the right way, but there are nice helpers that make it very simple and correct with all kinds of contexts.
- Report the MSVC version in Scons output during compilation. The 2022 version is required, but we support everything back to 2008, to work on very old systems as well. This will help identifying differences that arise from there.
- Quality: Find Clang format from MSVC 2022 too. We use in auto format of Nuitka source code, but need to also search that as a new path.
- Added a spellchecker extension for Visual Code, resulting in many spelling fixes in all kinds of documentation and code. This finds more things than `codespell`, but also has a lot of false alarms.
- Check value of `--onefile-tempdir-spec` for typical user errors. It cannot be `.` as that would require to overwrite the onefile binary on Windows, and will generally behave very confusing. Warn about absolute or relative paths going outside of where the binary lives. Can be useful in controlled setups, but not generally. Also warn about using no variables, making non-unique paths.
- macOS: Flavor detection was largely expanded. The `Apple` flavor is recognized on more systems. `Homebrew` was newly added, and we actually can detect `CPython` reliably as a first.
- Added a tool from leo project to create better `.pyi` files for modules. We will make use of it in the future to enhance the files created by Nuitka to not only contain hidden dependencies, but optionally also module signatures.
- Plugins: Clearer information from `pyside2` that patched wheels might be mandatory and workarounds only patches cannot be done for older Python.
- Added progress bars for DLL dependency detection and DLL copying for standalone. These both can end up using take a fair bit of time depending on project size, and it's nice to know what's going on.
- macOS: Added support for using both `--onefile` and `--macos-create-app-bundle` as it is needed for PySide2 due to issues with signing code now.
- Added warning when attempting to include extension modules in an accelerated compilation.
- Modules: Catch the user error of following all imports when creating a module. This is very unlikely to produce usable results.
- Start integrating [Sourcery](#) for improved Nuitka code. It will comment on PRs and automatically improve Nuitka code as we develop it.

- Debugging: Added command line tool `find-module` that outputs how Nuitka locates a module in the Python environment it's ran with. That removes the need to use Python prompt to dump `__file__` of imported modules. Some modules even hide parts of their namespace actively during run-time, the tool will not be affected by that.

Cleanups

- Refactored Python scan previously used for Scons usage on versions that need to run in with another Python to be more generally usable.
- Use explicit `nuitka.utils.Hashing` module that allows the core to perform these operations with simpler code.
- macOS: Use `isPathBelow` for checking if something is a system library for enhanced robustness and code clarity.
- macOS: Make sure to use our proper error checking wrappers for command execution when using tools like `otool` or `codesign`.
- Standalone: Avoid a temporary file with a script during technical import detection. These have been observed to potentially become corrupted, and this avoids any chance of that happening, while also being simpler code.
- Avoid naming things `shlib` and call them `extension` instead. Inspired by the spell checker disliking the former term, which is also less precise.
- Removed the dead architecture selection option for Windows, it was unused for a long time.
- Moved Windows `SxS` handling of DLLs to a more general place where also macOS specific tasks are applied, to host standard modification of DLLs during their copying.

Tests

- Better matching of relative filenames for search modes of the individual test suite runners.
- Debugger outputs on segfaults were no longer visible and have been restored.

Summary

This release is tremendous progress for macOS. Finally biting the bullet and paying obscene amounts of money to rent an M1 machine, it was possible to enhance the support for this platform. Currently typical packages for macOS are being made compatible as well, so it can now be expected to perform equally well.

On the quality side, the spell checker has had some positive effects, finding typos and generally misspelled code, that `codespell` does not, due to it being very conservative.

The trend to enhance plugins has continued. The copying of DLLs is very nearly finalized. Making more plugins enabled by default is seeing a lot of progress, with 2 important ones addressed.

Work on the size of distributions has seen a lot of positive results, in that now standalone distributions are often very minimal, with many extension modules from standard library no longer being present.

Nuitka Release 0.6.19

This release adds support for 3.10 while also adding very many new optimization, and doing a lot of bug fixes.

Bug Fixes

- Calls to `importlib.import_module` with expressions that need releases, i.e. are not constant values, could crash the compilation. Fixed in 0.6.18.1 already.
- After a fix for the previous release, modules that fail to import are attempted again when another import is executed. However, during this initialization for top level module in `--module` mode, this was done repeatedly, and could cause issues. Fixed in 0.6.18.1 already.
- Standalone: Ignore warning given by `patchelf` on Linux with at least newer OpenSUSE. Fixed in 0.6.18.1 already.
- Fix, need to avoid computing large values out of `<<` operation as well. Fixed in 0.6.18.2 already.

```
# This large value was computed at run time and then if used, also
# converted to a string and potentially hashed, taking a long time.
1 << sys.maxint
```

- Standalone: Ignore warning given by `patchelf` on Linux about a workaround being applied.
- Fix, calls to `importlib.import_module` were not correctly creating code for dynamic argument values that need to be released, causing the compilation to report the error. Fixed in 0.6.18.1 already.
- MSYS2: Fix, the console scripts are actually good for it as opposed to CPython, and the batch scripts should not be installed. Fixed in 0.6.18.2 already.
- Setuptools: Added support older version of `setuptools` in meta build integration of Nuitka.
- Fix, calls to `importlib.import_module` with 2 arguments that are dynamic, were not working at all. Fixed in 0.6.18.2 already.
- Windows: Compiling with MinGW64 without `ccache` was not working due to issues in Scons. Fixed in 0.6.18.2 already.
- Fix, the `repr` built-in was falsely annotated as producing a `str` value, but it can be also derived or `unicode` in Python2.
- Fix, attribute nodes were not considering the value they are looking up on. Now that more values will know to have the attributes, that was causing errors. Fixed in 0.6.18.2 already.
- Fix, left shifting can also produce large values and needs to be avoided in that case, similar to what we do for multiplications already. Fixed in 0.6.18.2 already.
- UI: The new option `--disable-ccache` didn't really have the intended effect. Fixed in 0.6.18.3 already.
- UI: The progress bar was causing tearing and corrupted outputs, when outputs were made, now using proper `tqdm` API for doing it, this has been solved. Fixed in 0.6.18.4 already.
- Fix, the constant value `sys.version_info` didn't yet have support for its type to be also a compile time constant in e.g. tuples. Fixed in 0.6.18.4 already.
- Onefile: Assertions were not disabled, and on Windows with MinGW64 this lead to including the C filenames of the `zstd` inline copy files and obviously less optimal code. Fixed in 0.6.18.4 already.
- Standalone: Added support for `bottle.ext` loading extensions to resolve at compile time. Fixed in 0.6.18.5 already.
- Standalone: Added support for `seedir` required data file. Fixed in 0.6.18.5 already.
- MSYS2: Failed to link when using the static `libpython`, which is also now the default for MSYS2. Fixed in 0.6.18.5 already.
- Python3.6+: Fix, the intended finalizer of compiled `asyncgen` was not present and in fact associated to help type. This could have caused corruption, but that was also very unlikely. Fixed in 0.6.18.5 already.
- Python3: Fix, need to set `__file__` before executing modules, as some modules, e.g. newer PyWin32 use them to locate things during their initialization already.

- Standalone: Handle all PyWin32 modules that need the special DLLs and not just a few.
- Fix, some `.pth` files create module namespaces with `__path__` that does not exist, ignore these in module importing.
- Python2.6-3.4: Fix, modules with an error could use their module name after it was released.
- Distutils: When providing arguments, the method suggested in the docs is not compatible with all other systems, e.g. not `setuptools_rust` for which a two elements tuple form needs to be used for values. Added support for that and documented its use as well in the User Manual.
- Python3.7+: Do no longer allow deleting cell values, this can lead to corruption and should be avoided, it seems unlikely outside of tests anyway.
- Standalone: Added support for more ciphers and hashes with `pycryptodome` and `pycryptodomex`, while also only including Ciphers when needed.
- Distutils: Was not including modules or packages only referenced in the entry point definition, but not in the list of packages. That is not compatible and has been fixed.
- Fix, must not expose the constants blob from extension modules, as loading these into a compiled binary can cause issues in this case.
- Standalone: Added support for including OpenGL and SSL libraries with `PySide2` and `PySide6` packages.
- Windows: Fix, the `cmd` files created for uninstalled Python and accelerated programs to find the Python installation were not passing command line arguments.
- Windows: Executing modules with `--run` was not working properly due to missing escaping of file paths.
- Fix, parsing `.pyi` files that make relative imports was not resolving them correctly.
- Python3: Fix, when disabling the console on Windows, make sure the file handles still work and are not `None`.
- Windows: Fix, need to claim all OS versions of Windows as supported, otherwise e.g. high DPI features are not available.

New Features

- Programs that are to be executed with the `-m` flag, can now be compiled with `--python-flag=-m` and will then behave in a compatible way, i.e. load the containing package first, and have a proper `__package__` value at run time.
- We now can write XML reports with information about the compilation. This is initially for use in PGO tests, to decide if the expected forms of inclusions have happened and should grow into a proper reporting tool over time. At this point, the report is not very useful yet.
- Added support for Python 3.10, only `match` statements are not completely supported. Variations with `|` matches that also assign are not allowed currently.
- Windows: Allow using `--clang` with `--mingw64` to e.g. use the `clang.exe` that is contained in the Nuitka automatic download rather than `gcc.exe`.
- Added support for Kivy. Works through a plugin that is automatically enabled and needs no other inputs, detecting everything from using Kivy at compile time.
- Added initial support for Haiku OS, a clone of BeOS with a few differences in their Python installation.
- Added experimental plugin `trio` that works around issues with that package.

Optimization

- Also trust hard imports made on the module level in function level code, this unlocks many more static optimization e.g. with `sys.version_info` when the import and the use are not on the same level.
- For the built-in type method calls with generic implementation, we now do faster method descriptor calls. These avoid creating a temporary `PyCFunction` object, that the normal call slot would, this should make these calls faster. Checking them for compiled function, etc. was only wasteful, so this makes it more direct.
- Loop and normal merge traces were keeping assignments made before the loop or inside a branch, that was otherwise unused alive. This should enable more optimization for code with branches and loops. Also unused loop traces are now recognized and removed as well.
- Avoiding merges of escaped traces with the unescaped trace, there is no point in them. This was actually happening a lot and should mean a scalability improvement and unlock new optimization as well.
- Avoid escaping un-init traces. Unset values need not be considered as potentially modified as that cannot be done.
- The `str` shape is now detected through variables, this enables many optimization on the function level.
- Added many `str` operation nodes.

These are specifically all methods with no arguments, as these are very generic to add, introduced a base class for them, where we know they all have no effect or raise, as these functions are all guaranteed to succeed and can be served by a common base class.

This covers the `str.capitalize`, `str.upper`, `str.lower`, `str.swapcase`, `str.title`, `str.isalnum`, `str.isalpha`, `str.isdigit`, `str.islower`, `str.isupper`, `str.isspace`, and `str.istitle` functions.

For static optimization `str.find` and `str.rfind` were added, as they are e.g. used in a `sys.version.find(...)` style in the `os` module, helping to decide to not consider `OS/2` only modules.

Then, support for `str.index` and `str.rindex` was added, as these are very similar to `str.find` forms, only that these may raise an exception.

Also add support for `str.split` and `str.rsplit` which will be used sometimes for code needed to be compile time computed, to e.g. detect imports.

Same goes for `endswith` and `startswith`, the later is e.g. popular with `sys.platform` checks, and can remove a lot of code from compilation with them now being decided at compile time.

Note

A few `str` methods are still missing, with time we will achieve all of them, but this will take time.

- Added trust for `sys.builtin_module_names` as well. The `os` module is using it to make platform determinations.
- When writing constant values, esp. `tuple`, `list`, or `dict` values, an encoding of "last value" has been added, avoiding the need to repeat the same value again, making many values more compact.
- When starting Nuitka, it usually restarts itself with information collected in a mode without the `site` module loaded, and with hash randomization disabled, for deterministic behaviour. There is a option

to prevent this from happening, where the goal is to avoid it, e.g. in testing, say for the coverage taking, but that meant to parse the options twice, which also loads a lot of code.

Now only a minimal amount of code is used, and the options are parsed only on the restart, and then an error is raised when it notices, it was not allowed to do so. This also makes code a lot cleaner.

- Specialized comparison code for Python2 `long` and Python3 `int` code, making these operations much faster to use.
- Specialized comparison code for Python2 `unicode` and Python3 `str` code, making these operations much faster to use, currently only `==` and `!=` are fully accelerated, the other comparisons will follow.
- Enable static `libpython` with Python3 Debian packages too. As with Python2, this will improve the performance of the created binary a lot and reduce size for standalone distribution.
- Comparisons with `in` and `not in` also consider value traces and go through variables as well where possible. So far only the rich comparisons and `is` and `is not` did that.
- Create fixed import nodes in re-formulations rather than `__import__` nodes, avoiding later optimization doing that, and of course that's simpler code too.
- Python 3.10: Added support for `union` types as compiled time constants.
- Modules are now fully optimized before considering which modules they are in turn using, this avoids temporary dependencies, that later turn out unused, and can shorten the compilation in some cases by a lot of time.
- On platforms without a static link library, in LTO mode, and with `gcc`, we can use the `-O3` mode, which doesn't work for `libpython`, but that's not used there. This also includes fake static `libpython`, as used by MinGW64 and Anaconda on Windows.
- The `anti-bloat` plugin now also handles newer `sklearn` and knows more about the standard library, and its runners which it will exclude from compilation if use for it. Currently that is not the default, but it should become that.

Organisational

- Migrated the Nuitka blog from Nikola to Sphinx based ABlog and made the whole site render with Sphinx, making it a lot more usable.
- Added a small presentation about Nuitka on the Download page, to make sure people are aware of core features.
- The `gi` plugin is now always on. The copying of the `typelib` when `gi` is imported is harmless and people can disable the plugin if that's not needed.
- The `matplotlib` plugin is new and also always on. It previously was part of the `numpy` plugin, which is doing too many unrelated things. Moving this one out is part of a plan to split it up and have it on by default without causing issues.
- MSYS2: Detecting `MinGW` and `POSIX` flavors of this Python. For the `MinGW` flavor of MSYS2, the option `--mingw64` is now the default, before it could attempt to use `MSVC`, which is not going to work for it. And also the `Tcl` and `Tk` installations of it are being detected automatically for the `tk-inter` plugin.
- Added Windows version to Nuitka version output, so we have this for bug reports.
- User Manual: Added example explaining how to access values from your code in Nuitka project options.
- UI: For Python flavors where we expect a static `libpython`, the error message will now point out how to achieve it for each flavor.
- UI: Disable progress bar when `--show-scons` is used, it makes capturing the output from the terminal only harder.

- UI: Catch error of specifying both `--msvc=` and `--mingw64` options.
- Distutils: Improved error messages when using `setuptools` or `build` integration and failing to provide packages to compile.
- Plugins: Removed now unused feature to rename modules on import, as it was only making the code more complex, while being no more needed after recently adding a place for meta path based importers to be accounted for.
- Twitter: Use embedded Tweet in Credits, and regular follow button in User Manual.
- Warnings about imports not done, are now only given when optimization can not remove the usage, and no options related to following have been given.
- Added Windows version to `--version` output of Nuitka. This is to more clearly recognize Windows 10 from Windows 11 report, and also the odd Windows 7 report, where tool chain will be different.
- In Visual Code, the default Python used is now 3.9 in the "Linux" C configuration. This matches Debian Bullseye.
- Nicer outputs from check mode of the auto-format as run for CI testing, displays problematic files more clearly.
- Remove broken links to old bug tracker that is no longer online from the Changelog.
- UI: When hitting CTRL-C during initial technical import detection, no longer ask to submit a bug report with the exception stack, instead exit cleanly.
- Windows: Enable LTO mode for MinGW64 and other gcc by default. We require a version that can do it, so take advantage of that.
- For cases, where code generation of a module takes long, make sure its name is output when CTRL-C is hit.
- Windows: Splash screen only works with MSVC, added error indicator for MinGW64 that states that and asks for porting help.

Cleanups

- Generate all existing C code for generic builtin type method calls automatically, and use those for method attribute lookups, making it easier to add more.
- Changed `TkInter` module to data file providing interface, yielding the 2 directories in question, with a filter for `demos`.
- The importing code got a major overhaul and no longer works with relative filenames, or filenames combined with package names, and module names, but always only with module names and absolute filenames. This cleans up some of the oldest and most complex code in Nuitka, that had grown to address various requirements discovered over time.
- Major cleanup of Jinja2 template organisation.
Renamed all C templates from `.j2` to `.c.j2` for clarity, this was not done fully consistent before. Also move all C templates to `nuitka.codegen` package data, it will be confusing to make a difference between ones used during compile time and for the static generation, and the lines are going to become blurry.
Added Jinja2 new macro `CHECK_OBJECTS` to avoid branches on argument count in the call code templates. More of these things should be added.
Cleanup of code that generates header declarations, there was some duplication going on, that made it hard to generate consistent code.
- Removed `nuitka.finalizatiions.FinalizationBase`, we only have one final visitor that does everything, and that of course makes a lot of sense for its performance.

- Major cleanup of the Scons C compiler configuration setup. Moved things to the dedicate function, and harmonized it more.
- Resolved deprecation warnings given by with `--python-debug` for Nuitka.

Tests

- Started test suite for Python PGO, not yet completely working though, it's not yet doing what is needed though.
- Added generated test that exercises str methods in multiple variations.
- Revived `reflected` test suite, that had been removed, because of Nuitka special needs. This one is not yet passing again though, due to a few details not yet being as compatible as needed.
- Added test suite for CPython 3.10 and enable execution of tests with this version on GitHub actions.

Summary

This release is another big step forward.

The amount of optimization added is again very large, some of which yet again unlocks more static optimization of module imports, that previously would have to be considered implicit. Now analyzing these on the function level as well, we can start searching for cases, where it could be done, but is not done yet.

After starting with `dict`, method optimization has focused on `str` which is esp. important for static optimization of imports. The next goal will here be to cover `list` which are important for run time performance and currently not yet optimized. Future releases will progress there, and also add more types.

The C type specialization for Python3 has finally progressed, such that it is also covering the `long` and `unicode` and as such not limited to Python2 as much. The focus now needs to turn back to not working with `PyObject *` for these types, but e.g. with `+= 1` to make it directly work with `CLONG` rather than `LONG` for which structural changes in code generation will be needed.

For scalability, the `anti-bloat` work has not yet progressed as much as to be able to enable it by default. It needs to be more possible to disable it where it causes problems, e.g. when somebody really wants to include `pytest` and test frameworks generally, that's something that needs to be doable. Compiling without `anti-bloat` plugin is something that is immediately noticeable in exploding module amounts. It is very urgently recommended to enable it for your compilations.

The support for Windows has been further refined, actually fixing a few important issues, esp. for the Qt bindings too.

This release adds support for 3.10 outside of very special `match` statements, bringing Nuitka back to where it works great with recent Python. Unfortunately `orderedset` is not available for it yet, which means it will be slower than 3.9 during compilation.

Overall, Nuitka is closing many open lines of action with this. The `setuptools` support has yet again improved and at this point should be very good.

Nuitka Release 0.6.18

This release has a focus on new features of all kinds, and then also new kinds of performance improvements, some of which enable static optimization of what normally would be dynamic imports, while also polishing plugins and adding also many new features and a huge amount of organisational changes.

Bug Fixes

- Python3.6+: Fixes to `asyncgen`, need to raise `StopAsyncIteration` rather than `StopIteration` in some situations to be fully compatible.

- Onefile: Fix, LTO mode was always enabled for onefile compilation, but not all compilers support it yet, e.g. MinGW64 did not. Fixed in 0.6.17.1 already.
- Fix, `type` calls with 3 arguments didn't annotate their potential exception exit. Fixed in 0.6.17.2 already.
- Fix, trusted module constants were not working properly in all cases. Fixed in 0.6.17.2 already.
- Fix, `pkg-resources` exiting with error at compile time for unresolved requirements in compiled code, but these can of course still be optional, i.e. that code would never run. Instead give only a warning, and run time fail on these. Fixed in 0.6.17.2 already.
- Standalone: Prevent the inclusion of `drm` libraries on Linux, they have to come from the target OS at run time. Fixed in 0.6.17.2 already.
- Standalone: Added missing implicit dependency for `ipcqueue` module. Fixed in 0.6.17.3 already.
- Fix, Qt webengine support for everything but `PySide2` wasn't working properly. Partially fixed in 0.6.17.3 already.
- Windows: Fix, bootstrap splash screen code for Windows was missing in release packages. Fixed in 0.6.17.3 already.
- Fix, could crash on known implicit data directories not present. Fixed in 0.6.17.3 already.
- macOS: Disable download of `ccache` binary for M1 architecture and systems before macOS 10.14 as it doesn't work on these. Fixed in 0.6.17.3 already.
- Standalone: The `pendulum.locals` handling for Python 3.6 was regressed. Fixed in 0.6.17.4 already.
- Onefile: Make sure the child process is cleaned up even after its successful exit. Fixed in 0.6.17.4 already.
- Standalone: Added support for `xmlschema`. Fixed in 0.6.17.4 already.
- Standalone: Added support for `curses` on Windows. Fixed in 0.6.17.4 already.
- Standalone: Added support for `coincurve` module. Fixed in 0.6.17.5 already.
- Python3.4+: Up until Python3.7 inclusive, a workaround for stream encoding (was ASCII), causing crashes on output of non-ASCII, other Python versions are not affected. Fixed in 0.6.17.5 already.
- Python2: Workaround for LTO error messages from older gcc versions. Fixed in 0.6.17.5 already.
- Standalone: Added support for `win32print`. Fixed in 0.6.17.6 already.
- Fix, need to prevent usage of static `libpython` in module mode or else on some Python versions, linker errors can happen. Fixed in 0.6.17.6 already.
- Standalone: Do not load `site` module early anymore. This might have caused issues in some configurations, but really only would be needed for loading `inspect` which doesn't depend on it in standalone mode. Fixed in 0.6.17.6 already.
- Fix, could crash with generator expressions in finally blocks of tried blocks that return. Fixed in 0.6.17.7 already.

```
try:
    return 9
finally:
    "".join(x for x in b"some_iterable")
```

- Python3.5+: Compatibility of comparisons with `types.CoroutineType` and `types.AsyncGeneratorType` types was not yet implemented. Fixed in 0.6.17.7 already.

```

# These already worked:
assert isinstance(compiledCoroutine(), types.CoroutineType) is True
assert isinstance(compiledAsyncgen(), types.AsyncGeneratorType) is True

# These now work too:
assert type(compiledCoroutine()) == types.CoroutineType
assert type(compiledAsyncgen()) == types.AsyncGeneratorType

```

- Standalone: Added support for `ruamel.yaml`. Fixed in 0.6.17.7 already.
- Distutils: Fix, when building more than one package, things could go wrong. Fixed in 0.6.17.7 already.
- Fix, for module mode filenames are used, and for packages, you can specify a directory, however, a trailing slash was not working. Fixed in 0.6.17.7 already.
- Compatibility: Fix, when locating modules, a package directory and an extension module of the same name were not used according to priority. Fixed in 0.6.17.7 already.
- Standalone: Added workaround `importlib_resources` insisting on Python source files to exist to be able to load datafiles. Fixed in 0.6.17.7 already.
- Standalone: Properly detect usage of hard imports from standard library in `--follow-stdlib` mode.
- Standalone: Added data files for `opensapi_spec_validator`.
- MSYS2: Fix, need to normalize compiler paths before comparing.
- Anaconda: For accelerated binaries, the created `.cmd` file wasn't containing all needed environment.
- macOS: Set minimum OS version derived from the Python executable used, this should make it work on all supported platforms (of that Python).
- Standalone: Added support for automatic inclusion of `xmlschema` package datafiles.
- Standalone: Added support for automatic inclusion of `eel` package datafiles.
- Standalone: Added support for `h5py` package.
- Standalone: Added support for `phonenumbers` package.
- Standalone: Added support for `feedparser` package, this currently depends on the `anti-bloat` plugin to be enabled, which will become enabled by default in the future.
- Standalone: Added `gi` plugin for said package that copies `typelib` files and sets the search path for them in standalone mode.
- Standalone: Added necessary datafiles for `eel` package.
- Standalone: Added support for `QtWebEngine` to all Qt bindings and also make it work on Linux. Before only `PySide2` on Windows was supported.
- Python3: Fix, the `all` built-in was wrongly assuming that bytes values could not be false, but in fact they are if they contain `\0` which is actually false. The same does not happen for string values, but that's a difference to be considered.
- Windows: The LTO was supposed to be used automatically on with MSVC 14.2 or higher, but that was regressed and has been repaired now.
- Standalone: Extension modules contained in packages, depending on their mode of loading had the `__package__` value set to a wrong value, which at least impacted new `matplotlib` detection of Qt backend.
- Windows: The `python setup.py install` was installing binaries for no good reason.

New Features

- Setuptools support. Documented `bdist_nuitka` and `bdist_wheel` integration and added support for Nuitka as a build package backend in `pyproject.toml` files. Using Nuitka to build your wheels is supposed to be easy now.
- Added experimental support for Python 3.10, there are however still important issues with compatibility with the CPython 3.9 test suite with at least `asyncgen` and `coroutines`.
- macOS: For app bundles, version information can be provided with the new option `--macos-app-version`.
- Added Python vendor detection of `Anaconda`, `pyenv`, `Apple Python`, and `pyenv` and output the result in version output, this should make it easier to analyse reported issues.
- Plugins: Also handle the usage of `__name__` for metadata version resolution of the `pkg-resources` standard plugin.
- Plugins: The `data-files` standard plugin now reads configuration from a Yaml file that `data-files.yml` making it more accessible for contributions.
- Windows: Allow enforcing usage of MSVC with `--msvc=latest`. This allows you to prevent accidental usage of MinGW64 on Windows, when MSVC is intended, but achieves that without fixing the version to use.
- Windows: Added support for LTO with MinGW64 on Windows, this was previously limited to the MSVC compiler only.
- Windows: Added support for using `--debugger` with the downloaded MinGW64 provided `gdb.exe`.

Note

It doesn't work when executed from a Git bash prompt, but e.g. from a standard command prompt.

- Added new experimental flag for compiled types to inherit from uncompiled types. This should allow easier and more complete compatibility, making even code in extension modules that uses `PyObject_IsInstance` work, providing support for packages like `pydantic`.
- Plugins: The Qt binding plugins now resolve `pyqtgraph` selection of binding by hard coding `QT_LIB`. This will allow to resolve its own dynamic imports depending on that variable at compile time. At this time, the compile time analysis is not covering all cases yet, but we hope to get there.
- macOS: Provide `minOS` for standalone builds, derived from the setting of the Python used to create it.
- UI: Added new option `--disable-ccache` to prevent Nuitka from injecting `ccache` (Clang, gcc) and `clcache` (MSVC) for caching the C results of the compilation.
- Plugins: Added experimental support for `PyQt6`. While using `PySide2` or `PySide6` is very much recommended with Nuitka, this allows its use.
- UI: Added option `--low-memory` to allow the user to specify that the compilation should attempt to use less memory where possible, this increases compile times, but might enable compilation on some weaker machines.

Optimization

- Added dedicated attribute nodes for attribute values that match names of dictionary operations. These are optimized into dedicated nodes for methods of dictionaries should their expression have an exact dictionary shape. These in turn optimize calls on them statically into dictionary operations. This is done for all methods of `dict` for both Python2 and Python3, namely `get`, `items`, `iteritems`, `itervalues`, `iterkeys`, `viewvalues`, `viewkeys`, `pop`, `setdefault`, `has_key`, `clear`, `copy`, `update`.

The new operation nodes also add compile time optimization for being used on constant values where possible.

- Also added dedicated attribute nodes for string operations. For operations, currently only part of the methods are done. These are currently only `join`, `strip`, `lstrip`, `rstrip`, `partition`, `rpartition`. Besides performance, this subset was enough to cover compile time evaluation of module name computation for `importlib.import_module` as done by SWIG bindings, allowing these implicit dependencies to be discovered at compile time without any help, marking a significant improvement for standalone usage.
- Annotate type shape for dictionary `in/not in` nodes, this was missing unlike in the generic `in/not in` nodes.
- Faster processing of "expression only" statement nodes. These are nodes, where a value is computed, but then not used, it still needs to be accounted for though, representing the value release.

```
something() # ignores return value, means statement only node
```

- Windows: Enabled LTO by default with MinGW64, which makes it produce much faster results. It now yield faster binaries than MSVC 2019 with `pystone`.
- Windows: Added support for C level PGO (Profile Guided Optimization) with MSVC and MinGW64, allowing extra speed boosts from the C compilation on Windows as well.
- Standalone: Better handling of `requests.packages` and `six.moves`. The old handling could duplicate their code. Now uses a new mechanism to resolve metapath based importer effects at compile time.
- Avoid useless exception checks in our dictionary helpers, as these could only occur when working with dictionary overloads, which we know to not be the case.
- For nodes, have dedicated child mixin classes for nodes with a single child value and for nodes with a tuple of children, so that these common kind of nodes operate faster and don't have to check at run time what type they are during access.
- Actually make use of the egg cache. `Nuitka` was unpacking eggs in every compilation, but in wheel installs, these can be quite common and should be faster.
- Star arguments annotated their type shape, but the methods to check for dictionary exactly were not affected by this preventing optimization in some cases.
- Added `anti-bloat` configuration for main programs present in the modules of the standard library, these can be removed from the compilation and should lower dependencies detected.
- Using static `libpython` with `pyenv` automatically. This should give both smaller (standalone mode) and faster results as is the case when using this feature..
- Plugins: Added improvements to the `anti-bloat` plugin for `gevent` to avoid including its testing framework.
- Python3.9+: Faster calls into uncompiled functions from compiled code using newly introduced API of that version.
- Statically optimize `importlib.import_module` calls with constant args into fixed name imports.
- Added support for `sys.version_info` to be used as a compile time constant. This should enable many checks to be done at compile time.

- Added hard import and static optimization for `typing.TYPE_CHECKING`.
- Also compute named import lookup through variables, expanding their use to more cases, e.g. like this:

```
import sys
...
if sys.version_info.major >= 3:
    ...
```

- Also optimize compile time comparisons through variable names if possible, i.e. the value cannot have changed.
- Faster calls of uncompiled code with Python3.9 or higher avoiding DLL call overhead.

Organisational

- Commercial: There are `Buy Now` buttons available now for the direct purchase of the [Nuitka Commercial](#) offering. Finally Credit Card, Google Pay, and Apple Pay are all possible. This is using Stripe. Get in touch with me if you want to use bank transfer, which is of course still best for me.
- The main script runners for Python2 have been renamed to `nuitka2` and `nuitka2-run`, which is consistent with what we do for Python3, and avoids issues where `bin` folder ends up in `sys.path` and prevents the loading of `nuitka` package.
- Windows: Added support for Visual Studio 2022 by updating the inline copy of Scons used for Windows to version 4.3.0, on non Windows, the other ones will keep being used.
- Windows: Requiring latest MinGW64 with version 11.2 as released by winlibs, because this is known to allow LTO, where previous releases were missing needed binaries.
- Reject standalone mode usage with Apple Python, as it works only with the other supported Pythons, avoiding pitfalls in attempting to distribute it.
- Move hosting of documentation to Sphinx, added Changelog and some early parts of API documentation there too. This gives much more readable results than what we have done so far with Nikola. More things will move there.
- User Manual: Add description how to access code attributes in `nuitka-project` style options.
- User Manual: Added commands used to generate performance numbers for Python.
- User Manual: List other Python's for which static linking is supposed to work.
- Improved help for `--include-package` with a hint how to exclude some of the subpackages.
- Started using Jinja2 in code templates with a few types, adding basic infrastructure to do that. This will be expanded in the future.
- Updated plugin documentation with more recent information.
- Added Python flavor as detected to the `--version` output for improved bug reports.
- Linux: Added distribution name to `--version` output for improved bug reports.
- Always enable the `gevent` plugin, we want to achieve this for all plugins, and this is only a step in that direction.
- Added project URLs for PyPI, so people looking at it from there have some immediate places to checkout.
- Debian: Use common code for included PDF files, which have page styles and automatic corrections for `rst2pdf` applied.
- Updated to latest `black`, `isort`, `pylint` versions.

- The binary names for Python2 changed from `nuitka` and `nuitka-run` to `nuitka2` and `nuitka2-run`. This harmonizes it with Python2 and avoids issues, where the `bin` folder in `sys.path` can cause issues with re-execution of Nuitka finding those to import.

Note

You ought to be using `python -m nuitka` style of calling Nuitka anyway, as it gives you best control over what Python is used to run Nuitka, you can pick `python2` there if you want it to run with that, even with full path. Check the relevant section in the User Manual too.

- Added support for Fedora 34 and Fedora 35.

Cleanups

- In a change of mind `--enable-plugin` has become the only form to enable a plugin used in documentation and tests.
- Massive cleanup of `numpy` and Qt binding plugins, e.g. `pyside2`. Data files and DLLs are now provided through proper declarative objects rather than copied manually. The handling of PyQt5 from the plugin should have improved as a side effect.
- Massive cleanups of all documentation in ReST format. Plenty of formatting errors were resolved. Many typos were identified and globally fixed. Spellings e.g. of "Developer Manual" are now enforced with automatic replacements. Also missing or wrong quotes were turned to proper methods. Also enforce code language for shell scripts to be the same everywhere.
- Removed last usages of `getPythonFlags()` and made the function private, replacing their use with dedicated function to check for individual flags.
- Avoid string comparison with `nuitka.utils.getOS()` and instead add accessors that are more readable, e.g. `nuitka.utils.isMacOS()` and put them to use where it makes sense.
- Replaced usages of string tests in list of python flags specified, with functions that check for a specific name with a speaking function name.
- Added mixin for expressions that have no side effect outside of their value, providing common method implementation more consistently.
- Remove code geared to using old PyLint and on Python2, we no longer use that. Also removed annotations only used for overriding Python2 builtins from Nuitka code.
- The PDF specific annotations were moved into being applied only in the PDF building step, avoiding errors for raw PDF directives.
- Apply Visual Code auto-format to our Yaml files. This is unfortunately not an automatic formatting yet.
- Introduce dedicated `nuitka.utils.Json` module, as we intend to expand its usage, e.g. for caching.
- Replacing remaining usages of `print` functions with uses of `nuitka.Tracing` instead.
- Massive cleanup of the `gevent` plugin, use proper method to execute code after module load, rather than source patching without need. The plugin no longer messes with inclusions that other code already provides for standalone.
- Using own helper to update `sys` module attributes, to avoid errors from old C compilers, and also cleaning up using code to not have to cast on string constants.

- More consistent naming of plugin classes, and enforce a relationship of detector class names to the names of detected plugins. The new naming consistency is now enforced.

Tests

- Added CPython 3.10 test suite, it needs more work though.
- Added generated test that exercises dictionary methods in multiple variations.
- Test suite names were specified wrongly in a few of them.

Summary

This release is again a huge step forward. It refines on PGO and LTO for C level to work with all relevant compilers. Internally Python level PGO is prepared, but only a future release will feature it. With that, scalability improvements as well as even more performance improvements will be unlocked.

The amount of optimization added this time is even bigger, some of which unlocks static optimization of module imports, that previously would have to be considered implicit. This work will need one extra step, namely to also trace hard imports on the function level, then this will be an extremely powerful tool to solve these kinds of issues in the future. The next release will have this and go even further in this area.

With the dictionary methods, and some string methods, also a whole new kind of optimization has been started. These will make working with `dict` containers faster, but obviously a lot of ground is to cover there still, e.g. `list` values are a natural target not yet started. Future releases will progress here.

Type specialization for Python3 has not progressed though, and will have to be featured in a future releases though.

For scalability, the `anti-bloat` work has continued, and this should be the last release, where this is not on by default. Compiling without it is something that is immediately noticeable in exploding module amounts. It is very urgently recommended to enable it for your compilations.

The support for macOS has been refined, with version information being possible to add, and adding information to the binary about which OSes are supported, as well as rejecting Apple Python, which is only a trap if you want to deploy to other OS versions. More work will be needed to support `pyenv` or even Homebrew there too, for now CPython is still the recommended platform to use.

This release achieves major compatibility improvements. And of course, the experimental support for 3.10 is not the least. The next release will strive to complete the support for it fully, but this should be usable at least, for now please stay on 3.9 if you can.

Nuitka Release 0.6.17

This release has a focus on performance improvements, while also polishing plugins and adding many new features.

Bug Fixes

- Fix, plugins were not catching being used on packages not installed. Fixed in 0.6.16.2 already.
- macOS: Fix weaknesses in the `otool` parsing to determine DLL dependency parsing. Fixed in 0.6.16.2 already.
- Linux: Allow onefile program args with spaces contained to be properly passed. Fixed in 0.6.16.3 already.
- Windows: Avoid using less portable C function for `%PID%` formatting, which restores compilation on Windows 7 with old toolchains. Fixed in 0.6.16.3 already.
- Standalone: Added support for `fstrings` package. Fixed in 0.6.16.3 already.

- Compatibility: Fix, need to import `.pth` files after `site` module, not before. This was causing crashes on CentOS7 with Python2. Fixed in 0.6.16.3 already.
- Compatibility: Fix, when extension modules failed to load, in some cases the `ImportError` was lost to a `KeyError`. Fixed in 0.6.16.3 already.
- Fix, linker resource modes `code` and `linker` were not working anymore, but are needed with LTO mode at least. Fixed in 0.6.16.3 already.
- Standalone: Bytecode modules with null bytes in standard library, typically from disk corruption, were not handled properly. Fixed in 0.6.16.3 already.
- Fix, failed `.throw()` into generators could cause corruption. Fixed in 0.6.16.4 already.
- Python2: Fix, the bytecode compilation didn't respect the `--python-flag=no_asserts` mode. Fixed in 0.6.16.4 already.
- Fix, calls were not annotating their arguments as escaped, causing corruption of mutable in static optimization. Fixed in 0.6.16.5 already.
- Fix, some sequence objects, e.g. `numpy.array` actually implement in-place add operations that need to be called. Fixed in 0.6.16.5 already.
- Windows: Fix, onefile binaries were not working after being signed. This now works.
- Standalone: Added missing implicit dependency for `sklearn`.
- Compatibility: Modules giving `SyntaxError` from source were not properly handled, giving run time `ImportError`. Now they are giving `SyntaxError`.
- Fix, the LTO mode has issues with `incbin` usage on older gcc, so use `linker` mode when it is enabled.
- Python3: Fix, locals dict codes were not properly checking errors that the mapping might raise when setting values.
- Fix, modules named `entry` were causing compile time errors in the C stage.
- macOS: Never include files from OS private frameworks in standalone mode.
- Fix, the python flag `--python-flag=no_warning` wasn't working on all platforms.
- Compatibility: Fix, the main code of the `site` module wasn't executing, so that its added builtins were not there. Of course, you ought to use `--python-flag=no_site` to not have it in the normal case.
- Python2: Added code path to handle edited standard library source code which then has no valid bytecode file.
- Anaconda: In module mode, the CondaCC wasn't recognized as form of gcc.
- Fix, bytecode modules could shadow compiled modules of the same name.
- Onefile: Fix, expansion of `%PID%` wasn't working properly on non-Windows, making temp paths less unique. The time stamp is not necessarily enough.
- Fix, `multiprocessing` error exits from slave processes were not reporting tracebacks.
- Standalone: Added `xcbglinTEGRATIONS` to the list of sensible Qt plugins to include by default, otherwise rendering will be inferior.
- Standalone: Added `platformthemes` to the list of sensible Qt plugins to include by default, otherwise file dialogs on non-Windows would be inferior.
- Fix, created `.pyi` files were not ordered deterministically.
- Standalone: Added support for `win32file`.
- Fix, namespace packages were not using run time values for their `__path__` value.
- Python3.7+: Fix, was leaking `AttributeError` exceptions during name imports.

- Fix, standard library detection could fail for relative paths.

New Features

- Added experimental support for C level PGO (Profile Guided Optimization), which runs your program and then uses feedback from the execution. At this time only gcc is supported, and only C compiler is collecting feedback. Check the User Manual for a table with current results.
- macOS: Added experimental support for creating application bundles. For these, icons can be specified and console can be disabled. But at this time, onefile and accelerated mode are not yet usable with it, only standalone mode works.
- Plugins: Add support for `pkg_resources.require` calls to be resolved at compile time. These are not working at run time, but this avoids the issue very nicely.
- Plugins: Massive improvements to the anti-bloat plugin, it can now make `numpy`, `scipy`, `skimage`, `pywt`, and `matplotlib` use much less packages and has better error handling.
- Plugins: Added anti-bloat ability to append code to a module, which might get used in the future by other plugins that need some sort of post load changes to be applied.
- Plugins: Added ability to replace code of functions at parse time, and use this in anti-bloat plugin to replace functions that do unnecessary stuff with variants that often just do nothing. This is illustrated here.

```
gevent._util:
description: "remove gevent release framework"
change_function:
    "prereleaser_middle": "'(lambda data: None)'"
    "postreleaser_before": "'(lambda data: None)'"
```

This example is removing `gevent` code that loads dependencies used for their CI release process, that need not be part of normal programs.

- Added ability to persist source code changes done by plugins in the Python installation. This is considered experimental and needs write access to the Python installation, so this is best done in a virtualenv and it may confuse plugins.
- Added support for `multiprocessing.tracker` and spawn mode for all platforms. For non-default modes outside of Windows, you need to `--enable-plugin=multiprocessing` to use these.
- Plugins: Allow multiple entry points to be provided by one or several plugins for the same modules. These are now merged into one automatically.
- Standalone: Fix for `numpy` not working when compiling with `--python-flag=no_docstrings`.
- Fix, method calls were not respecting descriptors provided by types with non-generic attribute lookups.
- Windows: Add support for using self-compiled Python3 from the build folder too.
- Added support for Nuitka-Python 2.7, which will be our faster Python fork.
- Colorized output for error outputs encountered in Scons, these are now yellow for better recognition.

Optimization

- Faster threading code was used for Python3.8 or higher, and this has been extended to 3.7 on Windows, but we won't be able to have it other platforms and not on earlier Python3 versions.

- Faster calls esp. with keyword arguments. Call with keywords no longer create dictionaries if the call target supports that, and with 3.8 or higher, non-compiled code that allows vectorcall is taken advantage of.
- Faster class creation that avoids creation of argument tuples and dictionaries.
- Faster attribute check code in case of non-present attributes.
- Faster unbound method calls, unlike bound methods calls these were not optimized as well yet.
- Type shapes for star arguments are now known and used in optimization.

```
def f(*args, **kwargs):
    type(args) # Statically known to be tuple
    type(kwargs) # Statically known to be dict
```

- Python2: Faster old-style class creation. These are classes that do not explicitly inherit from `object`.
- Python2: Faster string comparisons for Python by specializing for the `str` type as well.
- Python3: Added specialization for `bytes` comparisons too. These are naturally very much the same as `str` comparisons in Python2.
- Added specialization for `list` comparisons too. We had them for `tuples` only so far.
- Faster method calls when called from Python core, our `tp_call` slot wasn't as good as it can be.
- Optimization: Faster deep copies of constants. This can speed up constant calls with mutable types. Before it was checking the type too often to be fast.
- Allow using static linking with Debian Python giving much better performance with the system Python. This is actually a huge improvement as it makes things much faster. So far it's only automatically enabled for Python2, but it seems to work for Python3 on Debian too. Needs more tweaking in the future.
- Optimization: Added `functools` module to the list of hard imports in preparation of optimizing `functools.partial` to work better with compiled functions.
- Python2: Demote to `xrange` when iterating over `range` calls, even for small ranges, they are always faster. Previously this was only done for values with at least 256 values.
- Enable LTO automatically for Debian Python, this also allows more optimization.
- Enable LTO automatically for Anaconda with CondaCC on non-Windows, also allowing more optimization.

Organisational

- Added section in the User Manual on how to deal with memory issues and C compiler bugs. This is a frequent topic and should serve as a pointer for this kind of issue.
- The `--lto` option was changed to require an argument, so that it can also be disabled. The default is `auto` which is the old behaviour where it's enabled if possible.
- Changed `--no-progress` to `--no-progressbar` in order to make it more clear what it's about. Previously it was possible to relate it to `--show-progress`.
- No longer require specific versions of dependencies in our `requirements.txt` and relegate those to only being in `requirements-devel.txt` such that by default Nuitka doesn't collide with user requirements on those same packages which absolutely all the time don't really make a difference.
- Added ability to check all unpushed changes with `pylint` with a new `./bin/check-nuitka-with-pylint --unpushed` option. Before it was only possible to make the check (quickly) with `--diff`, but that stopped working after commits are made.
- Revived support for `vmp` based analysis of compiled programs, but it requires a fork of it now.

- Make Windows specific compiler options visible on all platforms. There is no point in them being errors, instead warnings are given when they are specified on non-Windows.
- Added project variable `Commercial` for use in Nuitka project syntax.
- Consistent use of metavaris for nicer help output should make it more readable.
- Avoid `ast` tree dumps in case of `KeyboardInterrupt` exceptions, they are just very noisy. Also not annotate where Nuitka was in optimization when a plugin is asking to `sys.exit`.

Cleanups

- Encoding names for UTF-8 in calls to `.encode()` were used inconsistent with and without dashes in the source code, added cleanup to auto-format that picks the one blessed.
- Cleanup taking of run time traces of DLLs used in preparation for using it in main code eventually, moving it to a dedicated module.
- Avoid special names for Nuitka options in test runner, this only adds a level of confusion. Needs more work in future release.
- Unify implementation to create modules into single function. We had 3 forms, one in recursion, one for main module, and one for plugin generated code. This makes it much easier to understand and use in plugins.
- Further reduced code duplication between the two Scons files, but more work will be needed there.
- Escaped variables are still known to be assigned/unassigned rather than unknown, allowing for many optimizations to still work on them., esp. for immutable value
- Enhanced auto-format for rest documents, bullet list spacing is now consistent and spelling of organisational is unified automatically.
- Moved icon conversion functionality to separate module, so it can be reused for other platforms more easily.

Tests

- Removed `reflected` test, because of Nuitka special needs to restart with variable Python flags. This could be reverted though, since Nuitka no longer needs anything outside inline copies, and therefore no longer loads from site packages.
- Use `anti-bloat` plugin in standalone tests of Numpy, Pandas and tests to reduce their compile times, these have become much more manageable now.
- Enhanced checks for used files to use proper below path checks for their ignoring.
- Remove `reflected` test, compiling Nuitka with Nuitka has gotten too difficult.
- Verify constants integrity at program end in debug mode again, so we catch corruption of them in tests.

Summary

This release is one of the most important ones in a long time. The PGO and LTO, and static libpython work make a big different for performance of created binaries.

The amount of optimization added is also huge, calls are much faster now, and object creations too. These avoiding to go through actual dictionaries and tuples in most cases when compiled code interacts gives very significant gains. This can be seen in the increase of `pystone` performance.

The new type specializations allow many operations to be much faster. More work will follow in this area and important types, `str` and `int` do not have specialized comparisons for Python3, holding it back somewhat to where our Python2 performance is for these things.

For scalability, the `anti-bloat` work is extremely valuable, and this plugin should become active by default in the future, for now it must be strongly recommended. It needs more control over what parts you want to deactivate from it, in case of it causing problems, then we can and should do it.

The support for macOS has been enhanced a lot, and will become perfect in the next release (currently develop). The bundle mode is needed for all kinds of GUI programs to not need a console. This platform is becoming as well supported as the others now.

Generally this release marks a huge step forward. We hope to add Python level PGO in the coming releases, for type knowledge retrofitted without any annotations used. Benchmarks will become more fun clearly.

Nuitka Release 0.6.16

This release is mostly polishing and new features. Optimization looked only at threading performance, and LTO improvements on Windows.

Bug Fixes

- Fix, the `pkg-resources` failed to resolve versions for `importlib.metadata` from its standard library at compile time. Fixed in 0.6.15.1 already.
- Standalone: Fix, `--include-module` was not including the module if it was an extension modules, but only for Python modules. Fixed in 0.6.15.1 already.
- Standalone: Added missing implicit dependencies for `gi.overrides`. Fixed in 0.6.15.1 already.
- Python3.9: Fix, could crash when using generic aliases in certain configurations. Fixed in 0.6.15.2 already.
- Fix, the tensorflow plugin needed an update due to changed API. Fixed in 0.6.15.3 already.
- When error exiting Nuitka, it now closes any open progress bar for cleaner display.
- Standalone: Added missing dependency for `skimage`.
- Standalone: The `numpy` plugin now automatically includes Qt backend if any of the Qt binding plugins is active.

New Features

- Python3.5+: Added support for onefile compression. This is using `zstd` which is known to give very good compression with very high decompression, much better than e.g. `zlib`.
- macOS: Added onefile support.
- FreeBSD: Added onefile support.
- Linux: Added method to use `tempdir` onefile support as used on other platforms as an alternative to `AppImage` based.
- Added support for recursive addition of files from directories with patterns.
- Attaching the payload to onefile now has a progress bar too.
- Windows: Preliminary support for the yet unfinished Nuitka-Python that allows static linking and higher performance on Windows, esp. with Nuitka.
- Windows: In acceleration mode, for uninstalled Python, now a `CMD` file is created rather than copying the `DLL` to the binary directory. That avoids conflicts with architectures and of course useless file copies.

- New abilities for plugin `anti-bloat` allow to make it an error when certain modules are imported. Added more specific options for usual trouble makes, esp. `setuptools`, `pytest` are causing an explosion for some programs, while being unused code. This makes it now easier to oversee this.
- It's now possible to override `appdirs` decision for where cache files live with an environment variable `NUITKA_CACHE_DIR`.
- The `-o` option now also works with `onefile` mode, it previously rejected anything but `acceleration` mode. Fixed in 0.6.15.3 already.
- Plugins: It's now possible for multiple plugins to provide pre or post load code for the same module.
- Added indications for compilation modes `standalone` and `onefile` to the `__compiled__` attribute.
- Plugins: Give nicer error message in case of colliding command line options.

Optimization

- Faster threading code is now using for Python 3.8 or higher and not only 3.9, giving a performance boost, esp. on Windows.
- Using `--lto` is now the default with MSVC 2019 or higher. This will give smaller and faster binaries. It has been available for some time, but not been the default yet.

Cleanups

- Using different progress bar titles for C compilation of Python code and C compilation of `onefile` bootstrap.
- Moved platform specific detections, for FreeBSD/OpenBSD/macOS out of the Scons file and to common Nuitka code, sometimes eliminating duplications with one version being more correct than the other.
- Massive cleanup of `datafile` plugin, using pattern descriptions, so more code duplication can be removed.
- More cleanup of the `scons` files, sharing more common code.

Organisational

- Under the name `Nuitka-Python` we are now also developing a fork of CPython with enhancements, you can follow and join it at <https://github.com/Nuitka/Nuitka-Python> but at this time it is not yet ready for prime time.
- `Onefile` under Windows now only is temporary file mode. Until we figure out how to solve the problems with locking and caching, the mode where it installs to the AppData of the user is no longer available.
- Renamed the plugin responsible for PyQt5 support to match the names of others. Note however, that at this time, `PySide2` or `PySide6` are to be recommended.
- Make it clear that `PySide 6.1.2` is actually going to be the supported version of `PySide6`.
- Use MSVC in GitHub actions.

Summary

This release had a massive focus on expanding existing features, esp. for `onefile`, and plugins API, such that we can now configure `anti-bloat` with `yaml`, have really nice `datafile` handling options, and have `onefile` on all OSes practically.

Nuitka Release 0.6.15

This release polished previous work with bug fixes, but there are also important new things that help make Nuitka more usable, with one important performance improvement.

Bug Fixes

- Fix, hard imports were not automatically used in code generation leading to errors when used. Fixed in 0.6.14.2 already.
- Windows: Fix, `clcache` was disabled by mistake. Fixed in 0.6.14.2 already.
- Standalone: Added data files for `jsonschema` to be copied automatically.
- Standalone: Support for `pendulum` wasn't working anymore with the last release due to plugin interface issues.
- Retry downloads without SSL if that fails, as some Python do not have working SSL. Fixed in 0.6.14.5 already.
- Fix, the `ccache` path wasn't working if it contained spaces. Fixed in 0.6.14.5 already.
- Onefile: For Linux and ARM using proper download off appimage. Fixed in 0.6.14.5 already.
- Standalone: Added support for `pyreadstat`. Fixed in 0.6.14.5 already.
- Standalone: Added missing dependencies for `pandas`. Fixed in 0.6.14.6 already.
- Standalone: Some preloaded packages from `.pth` do not have a `__path__`, these can and must be ignored.
- Onefile: On Linux, the `sys.argv[0]` was not the original file as advertised.
- Standalone: Do not consider `.mesh` and `.frag` files as DLLs in the Qt bindings when including the qml support. This was causing errors on Linux, but was generally wasteful.
- Fix, project options could be injected twice, which could lead to errors with options that were only allowed once, e.g. `--linux-onefile-icon`.
- Windows: When updating the resources in created binaries, treat all kinds of `OSError` with information output.
- Onefile: Remove onefile target binary path at startup as well, so it cannot cause confusion after error exit.
- Onefile: In case of error exit from `AppImage`, preserve its outputs and attempt to detect if there was a locking issue.
- Standalone: Scan package folders on Linux for DLLs too. This is necessary to properly handle `PyQt5` in case of Qt installed in the system as well.
- Standalone: On Linux, standard QML files were not found.
- Standalone: Enforce C locale when detecting DLLs on Linux, otherwise whitelisting messages didn't work properly on newer Linux.
- Standalone: Added support for `tzdata` package data files.
- Standalone: Added support for `exchangelib`.
- Python3.9: Fix, type subscripts could cause optimization errors.
- UI: Project options didn't properly handle quoting of arguments, these are normally removed by the shell.
- Linux: The default icon for Python in the system is now found with more version specific names and should work on more systems.

- Standalone: Do not include `libstdc++` as it should come from the system rather.

New Features

- Added plugin `anti-bloat` plugin, intended to fight bloat. For now it can make including certain modules an error, a warning, or force `ImportError`, e.g. `--noinclude-setuptools-mode=nofollow` is very much recommended to limit compilation size.
- The `pkg-resources` builtin now covers `metadata` and `importlib_metadata` packages for compile time version resolution as well.
- Added support for `PySide2` on Python version before 3.6, because the patched code needs no workarounds. Fixed in 0.6.14.3 already.
- Windows: Convert images to icon files on the fly. So now you can specify multiple PNG files, and Nuitka will create an icon out of that automatically.
- macOS: Automatically download `ccache` binary if not present.
- Plugins: New interface to query the main script path. This allows plugins to look at its directory.
- UI: Output the versions of Nuitka and Python during compilation.
- UI: Added option to control static linking. So far this had been enabled only automatically for cases where we are certain, but this allows to force enable or disable it. Now an info is given, if Nuitka thinks it might be possible to enable it, but doesn't do it automatically.
- UI: Added `--no-onefile` to disable `--onefile` from project options.

Optimization

- Much enhanced GIL interaction with Python3.9 giving a big speed boost and better threading behaviour.
- Faster conversion of iterables to `list`, if size can be know, allocation ahead saves a lot of effort.
- Added support for `GenericAlias` objects as compile time constants.

Organisational

- Enhanced GitHub issue raising instructions.
- Apply `rstfmt` to all documentation and make it part of the commit hook.
- Make sure to check Scons files as well. This would have caught the code used to disable `clcache` temporarily.
- Do not mention Travis in PR template anymore, we have stopped using it.
- Updated requirements to latest versions.
- Bump requirements for development to 3.7 at least, tools like black do not work with 3.6 anymore.
- Started work on Nuitka Python, a CPython fork intended for enhanced performance and standalone support with Nuitka.

Cleanups

- Determine system prefix without `virtualenv` outside of Scons, such that plugins can share the code. There was duplication with the `numpy` plugin, and this will only be more complete using all approaches. This also removes a lot of noise from the scons file moving it to shared code.

- The Qt plugins now collect QML files with cleaner code.

Tests

- Nicer error message if a wrong search mode is given.
- Windows: Added timeout for determining run time traces with dependency walker, sometimes this hangs.
- Added test to cover the zip importer.
- Making use of project options in onefile tests, making it easier to execute them manually.

Summary

For Windows, it's now easier than ever to create an icon for your deployment, because you can use PNG files, and need not produce ICO files anymore, with Nuitka doing that for you.

The onefile for Linux had some more or less severe problems that got addressed, esp. also when it came to QML applications with PySide.

On the side, we are preparing to greatly improve the caching of Nuitka, starting with retaining modules that were demoted to bytecode. There are changes in this release, to support that, but it's not yet complete. We expect that scalability will then be possible to improve even further.

Generally this is mostly a maintenance release, which outside of the threading performance improvement has very little to offer for faster execution, but that actually does a lot. Unfortunately right now it's limited to 3.9, but some of the newer Python's will also be supported in later releases.

Nuitka Release 0.6.14

This release has few, but important bug fixes. The main focus was on expanding standalone support, esp. for PySide2, but also and in general with plugins added that workaround `pkg_resources` usage for version information.

Also an important new features was added, e.g. the project configuration in the main file should prove to be very useful.

Bug Fixes

- Compatibility: Fix, modules that failed to import, should be retried on next import.
So far we only ever executed the module body once, but that is not how it's supposed to be. Instead, only if it's in `sys.modules` that should happen, which is the case after successful import.
- Compatibility: Fix, constant `False` values in right hand side of `and/or` conditions were generating wrong code if the left side was of known `bool` shape too.
- Standalone: Fix, add `styles` Qt plugins to list of sensible plugins.
Otherwise no mouse hover events are generated on some platforms.
- Compatibility: Fix, relative `from` imports beyond level 1 were not loadingg modules from packages if necessary. Fixed in 0.6.13.3 already.
- Standalone: The `crypto` DLL check for Qt bindings was wrong. Fixed in 0.6.13.2 already.
- Standalone: Added experimental support for PySide6, but for good results, 6.1 will be needed.
- Standalone: Added support for newer matplotlib. Fixed in 0.6.12.1 already.
- Standalone: Reverted changes related to `pkg_resources` that were causing regressions. Fixed in 0.6.13.1 already.

- Standalone: Adding missing implicit dependency for `cytoolz` package. Fixed in 0.6.13.1 already.
- Standalone: Matching for package names to not suggest recompile for was broken and didn't match. Fixed in 0.6.13.1 already.

New Features

- Added support for project options.

When found in the filename provided, Nuitka will inject options to the commandline, such that it becomes possible to do a complex project with only using

```
python -m nuitka filename.py
```

```
# Compilation mode, support OS specific.
# nuitka-project-if: {OS} in ("Windows", "Linux"):
#     nuitka-project: --onefile
# nuitka-project-if: {OS} not in ("Windows", "Linux"):
#     nuitka-project: --standalone

# The PySide2 plugin covers qt-plugins
# nuitka-project: --enable-plugin=pyside2
# nuitka-project: --include-qt-plugins=sensible,qml

# The pkg-resources plugin is not yet automatic
# nuitka-project: --enable-plugin=pkg-resources

# Nuitka Commercial only features follow:

# Protect the constants from being readable.
# nuitka-project: --enable-plugin=data-hiding

# Include datafiles for Qt into the binary directory.
# nuitka-project: --enable-plugin=datafile-inclusion
# nuitka-project: --qt-datadir={MAIN_DIRECTORY}
# nuitka-project: --qt-datafile-pattern=*.js
# nuitka-project: --qt-datafile-pattern=*.qml
# nuitka-project: --qt-datafile-pattern=*.svg
# nuitka-project: --qt-datafile-pattern=*.png
```

Refer to the User Manual for a table of directives and the variables allowed to be used.

- Added option to include whole data directory structures in standalone.

The new option `--include-data-dir` was added and is mostly required for onefile mode, but recommended for standalone too.

- Added `pkg-resources` plugin.

This one can resolve code like this at compile time without any need for pip metadata to be present or used.

```
pkg_resources.get_distribution("module_name").version
pkg_resources.get_distribution("module_name").parsed_version
```

- Standalone: Also process early imports in optimization.

Otherwise plugins cannot work on standard library modules. This makes it possible to handle them as well.

Optimization

- Faster binary operations.

Applying lessons learnt during the enhancements for in-place operations that initially gave worse results than some manual code, we apply the same tricks for all binary operations, which speeds them up by significant margins, e.g. 30% for float addition, 25% for Python int addition, and still 6% for Python int addition.

- More direct optimization of unary operations on constant value.

Without this, `-1` was not directly a constant value, but had to go through the unary `-` operation, which it still does, but now it's done at tree building time.

- More direct optimization for `not` in branches.

Invertible comparisons, i.e. `is/is not` and `in/not in` do not have to be done during optimization. This mainly avoids noise during optimization from such unimportant steps.

- More direct optimization for constant slices.

These are used in Python3 for all subscripts, e.g. `a[1:2]` will use `slice(1,2)` effectively. For Python2 they are used less often, but still. This also avoids a lot of noise during optimization, mostly on Python3

- Scons: Avoid writing database to disk entirely.

This saves a bit of disk churn and makes it unnecessary to specify the location such that it doesn't collide between Python versions.

- For optimization passes, use previous max total as minimum for next pass. That will usually be a more accurate result, rather than starting from 1 again. Part of 0.6.13.1 already.
- Enhancements to the branch merging improve the scalability of Nuitka somewhat, although the merging itself is still not very scalable, there are some modules that are very slow to optimize still.
- Use `orderiset` if available over the inline copy for `OrderedSet` which is much faster and improves Nuitka compile times.
- Make `pkgutil` a hard import too, this is in preparation of more optimization for its functions.

Organisational

- Upstream patches for `PySide6` have been contributed and merged into the development branch `dev`. Full support should be available once this is released as part of 6.1 which is waiting for Qt 6.1 naturally.
- Patches for `PySide2` are available to commercial customers, see [PySide2 support](#) page.
- Formatted all documents with `rstfmt` and made that part of the commit hook for Nuitka. It now works for all documents we have.
- Updated inline copy of `tqdm` to 4.59.0 which ought to address spurious errors given.
- User Manual: Remove `--show-progress` from the tutorial. The default progress bar is then disabled, and is actually much nicer to use.
- Developer Manual: Added description of how context managers should be named.
- Cleanup language for some warnings and outputs.

It was still using obsolete "recursion" language rather than talking about "following imports", which is the new one.

Cleanups

- Remove dead code related to constants marshal, the data composer has replaced this.
- Avoid internal API usage for loading extension modules on Linux, there is a function in `sys` module to get the ld flags.

Tests

- Fix, the `only` mode wasn't working properly.
- Use new project options feature for specific options in basic tests allowing to remove them from the test runner.

Summary

For PySide2 things became more perfect, but it takes upstream patches unfortunately such that only PySide6.1 will be working out of the box outside of the commercial offering. We will also attempt to provide workarounds, but there are some things that cannot be done that way.

This release added some more scalability to the optimization process, however there will be more work needed to make efficient branch merges.

For onefile, a feature to include whole directories had been missing, and could not easily be achieved with the existing options. This further rounds this up, now what's considered missing is compression and macOS support, both of which should be coming in a future release.

For the performance side of things, the binary operator work can actually yield pretty good gains, with double digit improvements, but this covers only so much. Much more C types and better type tracing would be needed, but there was no progress on this front. Future releases will have to revisit the type tracing to make sure, we know more about loop variables, etc. so we can achieve the near C speed we are looking for, at least in the field of `int` performance.

This release has largely been driven by the [Nuitka Commercial](#) offering and needs for compatibility with more code, which is of course always a good thing.

Nuitka Release 0.6.13

This release follows up with yet again massive improvement in many ways with lots of bug fixes and new features.

Bug Fixes

- Windows: Icon group entries were not still not working properly in some cases, leading to no icon or too small icons being displayed. Fixed in 0.6.12.2 already.
- Windows: Icons and version information were copied from the standalone executable to the onefile executable, but that failed due to race situations, sometimes reproducible. Instead we now apply things to both independently. Fixed in 0.6.12.2 already.
- Standalone: Fixup scanning for DLLs with `ldconfig` on Linux and newer versions making unexpected outputs. Fixed in 0.6.12.2 already.
- UI: When there is no standard input provided, prompts were crashing with `EOFError` when `--assume-yes-for-downloads` is not given, change that to defaulting to "no" instead. Fixed in 0.6.12.2 already.
- Windows: Detect empty strings for company name, product name, product and file versions rather than crashing on them later. Them being empty rather than not there can cause a lot of issues in other places. Fixed in 0.6.12.2 already.

- Scons: Pass on exceptions during execution in worker threads and abort compilation immediately. Fixed in 0.6.12.2 already.
- Python3.9: Still not fully compatible with typing subclasses, the enhanced check is now closely matching the CPython code. Fixed in 0.6.12.2 already.
- Linux: Nicer error message for missing `libfuse` requirement.
- Compatibility: Lookups on dictionaries with `None` value giving a `KeyError` exception, but with no value, which is not what CPython does.
- Python3.9: Fix, future annotations were crashing in debug mode. Fixed in 0.6.12.3 already.
- Standalone: Corrections to the `glfw` were made. Fixed in 0.6.12.3 already.
- Standalone: Added missing implicit dependency for `py.test`. Fixed in 0.6.12.3 already.
- Standalone: Adding missing implicit dependency for `pyreadstat`.
- Windows: Added workaround for common `clcache` locking problems. Since we use it only inside a single Scons process, we can avoid using Windows mutexes, and use a process level lock instead.
- Plugins: Fix plugin for support for `eventlet`. Fixed in 0.6.12.3 already.
- Standalone: Added support for latest `zmq` on Windows.
- Scons: the `--quiet` flag was not fully honored yet, with Scons still making a few outputs.
- Standalone: Added support for alternative DLL name for newer `PyGTK3` on Windows. Fixed in 0.6.12.4 already.
- Plugins: Fix plugin for support for `gevent`. Fixed in 0.6.12.4 already.
- Standalone: Added yet another missing implicit dependency for `pandas`.
- Plugins: Fix, the `qt-plugins` plugin could stumble over `.mesh` files.
- Windows: Fix, dependency walker wasn't properly working with unicode `%PATH%` which could e.g. happen with a `virtualenv` in a home directory that requires them.
- Python3: Fixed a few Python debug mode warnings about unclosed files that have sneaked into the codebase.

New Features

- Added new options `--windows-force-stdout-spec` and `--windows-force-stderr-spec` to force output to files. The paths provided at compile time can resolve symbolic paths, and are intended to e.g. place these files near the executable. Check the User Manual for a table of the currently supported values. At this time, the feature is limited to Windows, where the need arose first, but it will be ported to other supported OSes eventually. These are most useful for programs run as `--windows-disable-console` or with `--enable-plugin=windows-service`.

Note

These options have since been renamed to `--force-stdout` and `--force-stderr` and have been made to work on all OSes.

- Windows: Added option `--windows-onefile-tempdir-spec` (since renamed to `--onefile-tempdir-spec`) to provide the temporary directory used with `--windows-onefile-tempdir` in onefile mode, you can now select your own pattern, and e.g. hardcode a base directory of your choice rather than `%TEMP%`.

- Added experimental support for `PySide2` with workarounds for compiled methods not being accepted by its core. There are known issues with `PySide2` still, but it's working fine for some people now. Upstream patches will have to be created to remove the need for workarounds and full support.

Optimization

- Use binary operation code for their in-place variants too, giving substantial performance improvements in all cases that were not dealt with manually already, but were covered in standard binary operations. Until now only some string, some float operations were caught sped up, most often due to findings of Nuitka being terribly slower, e.g. not reusing string memory for inplace concatenation, but now all operations have code that avoids a generic code path, that is also very slow on Windows due calling to using the embedded Python via API being slow.
- For mixed type operations, there was only one direction provided, which caused fallbacks to slower forms, e.g. with `long` and `float` values leading to inconsistent results, such that `a - 1` and `1 - a` being accelerated or not.
- Added C boolean optimization for a few operations that didn't have it, as these allow to avoid doing full computation of what the object result would have to do. This is not exhausted fully yet.
- Python3: Faster `+/-/+=/-=` binary and in-place operations with `int` values providing specialized code helpers that are much faster, esp. in cases where no new storage is allocated for in-place results and where not a lot of digits are involved.
- Python2: The Python3 `int` code is the Python2 `long` type and benefits from the optimization of `+/-/+=/-=` code as well, but of course its use is relatively rare.
- Improved `__future__` imports to become hard imports, so more efficient code is generated for them.
- Counting of instances had a run time impact by providing a `__del__` that was still needed to be executed and limits garbage collection on types with older Python versions.
- UI: Avoid loading `tqdm` module before it's actually used if at all (it may get disabled by the user), speeding up the start of Nuitka.
- Make sure to optimize internal helpers only once and immediately, avoiding extra global passes that were slowing down Python level compilation by of large programs by a lot.
- Make sure to recognize the case where a module optimization can provide no immediate change, but only after a next run, avoiding extra global passes originating from these, that were slowing down compilation of large programs by a lot. Together with the other change, this can improve scalability by a lot.
- Plugins: Remove implicit dependencies for `pkg_resources.extern` and use aliases instead. Using one of the packages, was causing all that might be used, to be considered as used, with some being relatively large. This was kind of a mistake in how we supported this so far.
- Plugins: Revamped the `eventlet` plugin, include needed DNS modules as bytecode rather than as source code, scanning them with `pkgutil` rather than `filesystem`, with much cleaner code in the plugin. The plugin is also now enabled by default.

Organisational

- Removed support for `pefile` dependency walker choice and inline copy of the code. It was never as good giving incomplete results, and after later improvements, slower, and therefore has lost the original benefit over using Dependency Walker that is faster and more correct.
- Added example for onefile on Windows with the version information and with the temporary directory mode.

- Describe difference in file access with onefile on Windows, where `sys.argv[0]` and `os.path.dirname(__file__)` will be different things.
- Added inline copy of `tqdm` to make sure it's available for progress bar output for 2.7 or higher. Recommend having it in the Debian package.
- Added inline copy of `colorama` for use on Windows, where on some terminals it will give better results with the progress bar.
- Stop using old PyLint for Python2, while it would be nice to catch errors, the burden of false alarms seems to high now.
- UI: Added even more checks on options that make no sense, made sure to do this only after a possible restart in proper environment, so warnings are not duplicated.
- For Linux onefile, keep appimage outputs in case of an error, that should help debugging it in case of issues.
- UI: Added traces for plugin provided implicit dependencies leading to inclusions.
- Added inline copy of `zstd` C code for use in decompression for the Windows onefile bootstrap, not yet used though.
- Added checks to options that accept package names for obvious mistakes, such that `--include-package-data --mingw64` will not swallow an option, as that is clearly not a package name, that would hide that option, while also not having any intended effect.
- Added ignore list for decision to recompile extension modules with available source too. For now, Nuitka will not propose to recompile `Cython` modules that are very likely not used by the program anyway, and also not for `lxml` until it's clear if there's any benefit in that. More will be added in the future, this is mostly for cases, where `Cython` causes incompatibilities.
- Added support for using abstract base classes in plugins. These are not considered for loading, and allow nicer implementation of shared code, e.g. between `PyQt5` and `PySide2` plugins, but allow e.g. to enforce the provision of certain overloads.
- User Manual: Remove the instruction to install `clcache`, since it's an inline copy, this makes no sense anymore and that was obsolete.
- Updated PyLint to latest versions, and our requirements in general.

Cleanups

- Started removal of PyLint annotations used for old Python2 only. This will be a continuous action to remove these.
- Jinja2 based static code generation for operations was cleaned up, to avoid code for static mismatches in the result C, avoiding language constructs like `if (1 && 0)` with sometimes larger branches, replacing it with Jinja2 branches of the `{% if ... %}` form.
- Jinja2 based Python2 `int` code was pioniering the use of macros, but this was expanded to allow kinds of types for binary operations, allow their reuse for in-place operation, with these macros making returns via `goto` exits rather than `return` statements in a function. Landing pads for these exits can then assign target values for in-place different from what those for binary operation result return do.
- Changed the interfacing of plugins with DLL dependency detection, cleaning up the interactions considerably with more unified code, and faster executing due to cached plugin decisions.
- Integrate manually provided slot function for `unicode` and `str` into the standard static code generation. Previously parts were generated and parts could be generated, but also provided with manual code. The later is now all gone.
- Use a less verbose progress bar format with less useless infos, making it less likely to overflow.

- Cleanup how payload data is accessed in Windows onefile bootstrap, preparing the addition of decompression, doing the reading from the file in only one dedicated function.
- When Jinja2 generated exceptions in the static code, it is now done via proper Jinja2 macros rather than Python code, and these now avoid useless Python version branches where possible, e.g. because a type like `bytes` is already Python version specific, with the goal to get rid of `PyErr_Format` usage in our generated static code. That goal is future work though.
- Move safe strings helpers (cannot overflow) to a dedicated file, and remove the partial duplication on the Windows onefile bootstrap code.
- The Jinja2 static code generation was enhanced to track the usage of labels used as goto targets, so that error exits, and value typed exits from operations code no longer emitted when not used, and therefore labels that are not used are not present.
- For implicit dependencies, the parsing of the `.pyi` file of a module no longer emits a dependency on the module itself. Also from plugins, these are now filtered away.

Tests

- Detect if onefile mode has required downloads and if there is user consent, otherwise skip onefile tests in the test runner.
- Need to also allow accesses to files via short paths on Windows if these are allowed long paths.
- The standalone tests on Windows didn't actually take run time traces and therefore were ineffective.
- Added standalone test for `glfw` coverage.
- Construct based tests for in-place operations are now using a value for the first time, and then a couple more times, allowing for real in-place usage, so we are sure we measure correctly if that's happening.

Summary

Where the big change of the last release were optimization changes to reduce the global passes, this release addresses remaining causes for extra passes, that could cause these to still happen. That makes sure, Nuitka scalability is very much enhanced in this field again.

The new features for forced outputs are at this time Windows only and make a huge difference when it comes to providing a way to debug Windows Services or programs in general without a console, i.e. a GUI program. These will need even more specifiers, e.g. to cover program directory, rather than exe filename only, but it's a very good start.

On the tooling side, not a lot has happened, with the clcache fix, it seems that locking issues on Windows are gone.

The plugin changes from previous releases had left a few of them in a state where they were not working, but this should be restored. Interaction with the plugins is being refined constantly, and this releases improved again on their interfaces. It will be a while until this becomes stable.

Adding support for PySide2 is a headline feature actually, but not as perfect as we are used to in other fields. More work will be needed, also in part with upstream changes, to get this to be fully supported.

For the performance side of things, the in-place work and the binary operations work has taken proof of concept stuff done for Python2 and applied it more universally to Python3. Until we cover all long operations, esp. `*` seems extremely important and is lacking, this cannot be considered complete, but it gives amazing speedups in some cases now.

Future releases will revisit the type tracing to make sure, we know more about loop variables, to apply specific code helpers more often, so we can achieve the near C speed we are looking for in the field of `int` performance.

Nuitka Release 0.6.12

This release is yet again a massive improvement in many ways with lots of bug fixes and new features.

Bug Fixes

- Windows: Icon group entries were not working properly in some cases, leading to no icon or too small icons being displayed.
- Standalone: The PyQt implicit dependencies were broken. Fixed in 0.6.11.1 already.
- Standalone: The datafile collector plugin was broken. Fixed in 0.6.11.3 already.
- Standalone: Added support for newer forms of `matplotlib` which need a different file layout and config file format. Fixed in 0.6.11.1 already.
- Plugins: If there was an error during loading of the module or plugin, it could still be attempted for use. Fixed in 0.6.11.1 already.
- Disable notes given by gcc, these were treated as errors. Fixed in 0.6.11.1 already.
- Windows: Fix, spaces in gcc installation paths were not working. Partially fixed in 0.6.11.4 already.
- Linux: Fix, missing onefile icon error message was not complete. Fixed in 0.6.11.4 already.
- Standalone: Workaround `zmq` problem on Windows by duplicating a DLL in both expected places. Fixed in 0.6.11.4 already.
- Standalone: The `dill-compat` plugin wasn't working anymore. Fixed in 0.6.11.4 already.
- Windows: Fix mistaken usage of `sizeof` for wide character buffers. This caused Windows onefile mode in temporary directory. Fixed in 0.6.11.4 already.
- Windows: Fix, checking subfolder natured crashed with different drives on Windows. Fixed in 0.6.11.4 already.
- Windows: Fix, usage from MSVC prompt was no longer working, detect used SDK properly. Fixed in 0.6.11.4 already.
- Windows: Fix, old clcache installation uses pth files that prevented our inline copy from working, workaround was added.
- Windows: Also specify stack size to be used when compiling with gcc or clang.
- Fix, claim of Python 3.9 support also in PyPI metadata was missing. Fixed in 0.6.11.5 already.
- Python3.9: Subscripting `type` for annotations wasn't yet implemented.
- Python3.9: Better matching of types for metaclass selection, generic aliases were not yet working, breaking some forms of type annotations in base classes.
- Windows: Allow selecting `--msvc-version` when a MSVC prompt is currently activated.
- Windows: Do not fallback to using gcc when `--msvc-version` has been specified. Instead it's an error if that fails to work.
- Python3.6+: Added support for `del ()` statements, these have no effect, but were crashing Nuitka.

```
del a # standard form
del a, b # same as del a; del b
del (a, b) # braces are allowed
del ( ) # allowed for consistency, but wasn't working.
```

- Standalone: Added support for `glfw` through a dedicated plugin.
- macOS: Added support for Python3 from system and CPython official download for latest OS version.

New Features

- UI: With `tqdm` installed alongside `Nuitka`, experimental progress bars are enabled. Do not use `--show-progress` or `--verbose` as these might have to disable it.
- Plugins: Added APIs for final processing of the result and onefile post processing.
- Onefile: On Windows, the Python process terminates with `KeyboardInterrupt` when the user sends CTRL-break, CTRL-C, shutdown or logoff signals.
- Onefile: On Windows, in case of the launching process terminating unexpectedly, e.g. due to Taskmanager killing it, or a `os.sigkill` resulting in that, the Python process still terminates with `KeyboardInterrupt`.
- Windows: Now can select icons by index from files with multiple icons.

Optimization

- Avoid global passes caused by module specific optimization. The variable completeness is now traced per module and function scope, allowing a sooner usage. Unused temporary variables and closure variables are removed immediately. Recognizing possible auto releases of parameter variables is also instantly.

This should bring down current passes from 5-6 global passes to only 2 global passes in the normal case, reducing frontend compile times in some cases massively.

- Better unary node handling. Dedicated nodes per operation allow for more compact memory usage and faster optimization.
- Detect flow control and value escape for the repr of node based on type shape.
- Enhanced optimization of caught exception references, these never raise or have escapes of control flow.
- Exception matching operations are more accurately annotated, and may be recognized to not raise in more cases.
- Added optimization for the `issubclass` built-in.
- Removed `scons` caching as used on Windows entirely. We should either be using `clcache` or `ccache` automatically now.
- Make sure to use `__slots__` for all node classes. In some cases, mixins were preventing the feature from being it. We now enforce their correct specification of slots, which makes sure we can't miss it anymore. This should again gain more speed and save memory at frontend compile time.
- Scons: Enhanced gcc version detection with improved caching behavior, this avoids querying the same gcc binary twice.

Organisational

- The description of `Nuitka` on PyPI was absent for a while. Added back by adding long description of the project derived from the README file.
- Avoid terms `blacklist`, `whitelist` and `slave` in the `Nuitka` code preferring `blocklist`, `ignorelist` and `child` instead, which are actually more clear anyway. We follow a general trend to do this.
- Configured the inline copy of `Scons` so `pylance` has an easier time to find it.
- The git commit hook had stopped applying diffs with newest git, improved that.
- Updated description for adding new CPython test suite.

- Using https URLs for downloading dependency walker, for it to be more secure.
- The commit hook can now be disabled, it's in the Developer Manual how to do it.

Cleanups

- Moved unary operations to their own module, the operators module was getting too crowded.
- The scones files for Python C backend and Windows onefile got cleaned up some more and moved more common code to shared modules.
- When calling external tools, make sure to provide null input where possible.
- Unified calling `install_name_tool` into a single method for adding rpath and name changes both at the same time.
- Unified how tools like `readelf`, `ldconfig` etc. are called and error exits and outputs checked to make sure we don't miss anything as easily.

Tests

- Adapted for some openSUSE specific path usages in standalone tests.
- Basic tests for onefile operation and with termination signal sent were added.

Summary

The big changes in this release are the optimization changes to reduce the global passes and the memory savings from other optimization. These should again make Nuitka more scalable with large projects, but there definitely is work remaining.

Adding nice stopping behaviour for the Onefile mode on Windows is seemingly a first, and it wasn't easy, but it will make it more reliable to users.

Also tooling of gcc and MSVC on Windows got a lot more robust, covering more cases, and macOS support has been renewed and should be a lot better now.

The progress bar is a nice touch and improves the overall feel of the compilation process, but obviously we need to aim at getting faster overall still. For projects using large dependencies, e.g. Pandas the compilation is still far too slow and that will need work on caching frontend results, and better optimization and C code generation for the backend.

Nuitka Release 0.6.11

This release is a massive improvement in many ways with lots of bug fixes and new features.

Bug Fixes

- Fix, the `.pyi` file parser didn't handle relative imports. Fixed in 0.6.10.1 already.
- Windows: Fix, multiprocessing plugin was not working reliable following of imports from the additional entry point. Fixed in 0.6.10.1 already.
- Pipenv: Workaround parsing issue with our `setup.py` to allow installation from GitHub. Fixed in 0.6.10.1 already.
- Merging of branches in optimization could give nondeterministic results leading to more iterations than necessary. Fixed in 0.6.10.1 already.
- Windows: Avoid profile powershell when attempting to resolve symlinks. Fixed in 0.6.10.1 already.

- Windows: Fix, always check for `stdin`, `stdout`, and `stderr` presence. This was so far restricted to gui mode applications, but it seems to be necessary in other situations too. Fixed in 0.6.10.1 already.
- Python2: Fix, `--trace-execution` was not working for standalone mode but can be useful for debugging. Fixed in 0.6.10.1 already.
- Windows: Onefile could run into path length limits. Fixed in 0.6.10.3 already.
- Windows: The winlib gcc download link became broken and was updated. Fixed in 0.6.10.3 already.
- Plugins: The "`__main__`" module was not triggering all plugin hooks, but it needs to for completeness.
- Standalone: Fix, symlinked Python installations on Windows were not working, with dependency walker being unable to look into these. Fixed in 0.6.10.4 already.
- Standalone: Fix support for numpy on Windows and macOS, the plugin failed to copy important DLLs. Fixed in 0.6.10.4 already.
- Python3: For versions before 3.7, the symlink resolution also needs to be done, but wasn't handling the bytes output yet. Fixed in 0.6.10.4 already.
- Fix, folder based inclusion would both pick up namespace folders and modules of the same name, crashing the compilation due to conflicts. Fixed in 0.6.10.4 already.
- Fix, the `--lto` wasn't used for clang on non-Windows yet.
- Fix, the order of locals dict releases wasn't enforced, which could lead to differences that break caching of C files potentially. Fixed in 0.6.10.5 already.
- Fix, `hash` nodes didn't consider if their argument was raising, even if the type of the argument was `str` and therefore the operation should not. Fixed in 0.6.10.5 already.
- Fix, need to copy type shape and escape description for the replacement inverted comparisons when used with `not`, otherwise the compilation can crash as these are expected to be present at all times. Fixed in 0.6.10.5 already.
- Fix, some complex constant values could be confused, e.g. `-0j` and `0j`. These corner cases were not properly considered in the constant loading code, only for `float` so far.
- Standalone: Fix, bytecode only standard library modules were not working. This is at least used with Fedora 33.
- Linux: Fix, extension modules compiled with `--lto` were not working.
- Windows: Retry if updating resources fails due to Virus checkers keeping files locked.
- Plugins: Pre- and postload code of modules should not be allowed to cause `ImportError`, as these will be invisible to the other parts of optimization, instead make them unraisable error traces.
- Standalone: Adding missing import for SciPy 1.6 support.
- Windows: Fix, only export required symbols when using MinGW64 in module mode.

New Features

- Python3.9: Added official support for this version.
- Onefile: Added command line options to include data files. These are `--include-package-data` which will copy all non-DLLs and non-Python files of package names matching the pattern given. And `--include-data-file` takes source and relative target file paths and copies them. For onefile this is the only way to include files, for standalone mode they are mostly a convenience function.
- Onefile: Added mode where the file is unpacked to a temporary folder before running instead of doing it to `appdata`.

- Onefile: Added linux specific options `--linux-onefile-icon` to allow provision of an icon to use in onefile mode on Linux, so far this was only available as the hard coded path to a Python icon, which also didn't exist on all platforms.
- UI: Major logging cleanup. Everything is now using our tracing classes and even error exits go through there and are therefore colored if possible.
- Plugins: Make it easier to integrate commercial plugins, now only an environment variable needs to point to them.
- UI: Enhanced option parsing gives notes. This complains about options that conflict or that are implied in others. Trying to catch more usage errors sooner.
- Plugins: Ignore exceptions in buggy plugin code, only warn about them unless in debug mode, where they still crash Nuitka.
- Scons: More complete scons report files, includes list values as well and more modes used.
- Windows: The `clcache` is now included and no longer used from the system.
- Output for `clcache` and `ccache` results got improved.
- Enhanced support for `clang`, on Windows if present near a `gcc.exe` like it is the case for some winlibs downloads, it will be used. To use it provide `--mingw64 --clang` both. Without the first one, it will mean `clangcl.exe` which uses the MSVC compiler as a host.

Optimization

- Some modules had very slow load times, e.g. if they used many list objects due to linear searches for memory deduplication of objects. We now have dictionaries of practically all constant objects loaded, making these more instant.
- Use less memory at compile time due using `__slots__` for all node types, finally figured out, how to achieve this with multiple inheritance.
- Use hedley for compiler macros like `unlikely` as they know best how to do these.
- Special case the merging of 2 branches avoiding generic code and being much faster.
- Hard imports have better code generated, and are being optimized into for the few standard library modules and builtin modules we handle, they also now annotate the type shape to be module.
- No longer annotate hard module import attribute lookups as control flow escapes. Not present attributes are changed into static raises. Trust for values is configured for a few values, and experimental.
- Avoid preloaded packages for modules that have no side effects and are in the standard library, typically `.pth` files will use e.g. `os` but that's not needed to be preserved.
- Use `incbin` for including binary data through inline assembly of the C compiler. This covers many more platforms than our previous linker option hacks, and the fallback to generated C code. In fact everything but Windows uses this now.

Organisational

- Windows: For Scons we now require a Python 3.5 or higher to be installed to use it.
- Windows: Removed support for gcc older than version 8. This specifically affects CondaCC and older MinGW64 installations. Since Nuitka can now download the MinGW64 10, there is no point in having these and they cause issues.
- We took over the maintenance of `clcache` as Nuitka/clcache which is not yet ready for public consumption, but should become the new source of clache in the future.
- Include an inline copy of `clcache` in Nuitka and use it on Windows for MSVC and ClangCL.

- Removed compatibility older aliases of follow option, `--recurse-*` and require `--follow-*` options to be used instead.
- For pylint checking, the tool now supports a `--diff` mode where only the changed files get checked. This is much faster and allows to do it more often before commit.
- Check the versions of isort and black when doing the auto-format to avoid using outdated versions.
- Handling missing pylint more gracefully when checking source code quality.
- Make sure to use the codespell tool with Python3 and make sure to error exit when spelling problems were found, so we can use this in GitHub actions too.
- Removed Travis config, we now only use GitHub actions.
- Removed landscape config, it doesn't really exist anymore.
- Bumped all PyPI dependencies to their latest versions.
- Recommend ccache on Debian, as we now consider the absence of ccache something to warn about.
- Plugins: The DLLs asked for by plugins that are not found are no longer warned about.
- Allow our checker and format tools to run on outside of tree code. We are using that for Nuitka/clcache.
- Added support for Fedora 33 and openSUSE 15.3, as well as Ubuntu Groovy.
- Windows: Check if Windows SDK is installed for MSVC and ClangCL.
- Windows: Enhanced wording in case no compiler was found. No longer tell people how to manually install MinGW64, that is no longer necessary and `pywin32` is not needed to detect MSVC, so it's not installed if not found.
- Detect "embeddable Python" by missing include files, and reject it with proper error message.
- Added onefile and standalone as a use case to the manual and put also the DLL and data files problems as typically issues.

Cleanups

- Avoid decimal and string comparisons for Python versions checks, these were lazy and are going to break once 3.10 surfaces. In testing we now use tuples, in Nuitka core hexacimal values much like CPython itself does.
- Stop using subnode child getters and setters, and instead only use subnode attributes. This was gradually changed so far, but in this release all remaining uses have migrated. This should also make the optimization stage go faster.
- Change node constructors to not use a decorator to resolve conflicts with builtin names, rather handle these with manual call changes, the decorator only made it difficult to read and less performant.
- Move safe string helpers to their own dedicated helper file, allowing for reuse in plugin code that doesn't want to use all of Nuitka C helpers.
- Added utils code for inline copy imports, as we use that for quite a few things now.
- Further restructured the Scons files to use more common code.
- Plugins: The module name objects now reject many `str` specific APIs that ought to not be used, and the code got changed to use these instead, leading to cleaner and more correct usages.
- Using named tuples to specify included data files and entry points.
- Use `pkgutil` in plugins to scan for modules rather than listing directories.

Tests

- New option to display executed commands during comparisons.
- Added test suite for onefile testing.

Summary

This release has seen Python3.9 and Onefile both being completed. The later needs compression added on Windows, but that can be added in a coming release, for now it's fully functional.

The focus clearly has been on massive cleanups, some of which will affect compile time performance. There is relatively little new optimization otherwise.

The adoption of clcache enables a very fast caching, as it's now loaded directly into the Scons process, avoiding a separate process fork.

Generally a lot of polishing has been applied with many cleanups lowering the technical debt. It will be interesting to see where the hard module imports can lead us in terms of more optimization. Static optimization of the Python version comparisons and checks is needed to lower the amount of imports to be processed.

Important fixes are also included, e.g. the constants loading performance was too slow in some cases. The `multiprocessing` on Windows and `numpy` plugins were regressed and finally everything ought to be back to working fine.

Future work will have to aim at enhanced scalability. In some cases, Nuitka still takes too much time to compile if projects like Pandas include virtually everything installed as an option for it to use.

Nuitka Release 0.6.10

This release comes with many new features, e.g. onefile support, as well as many new optimization and bug fixes.

Bug Fixes

- Fix, was memory leaking arguments of all complex call helper functions. Fixed in 0.6.9.6 already.
- Plugins: Fix, the dill-compatible code needs to follow API change. Fixed in 0.6.9.7 already.
- Windows: Fixup for multiprocessing module and complex call helpers that could crash the program. Fixed in 0.6.9.7 already.
- Fix, the frame caching could leak memory when using caching for functions and generators used in multiple threads.
- Python3: Fix, importing an extension module below a compiled module was not possible in accelerated mode.
- Python3: Fix, keyword arguments for `open` built-in were not fully compatible.
- Fix, the scons python check should also not accept directories, otherwise strange misleading error will occur later.
- Windows: When Python is installed through a symbolic link, MinGW64 and Scons were having issues, added a workaround to resolve it even on Python2.
- Compatibility: Added support for `co_freevars` in code objects, e.g. newer matplotlib needs this.
- Standalone: Add needed data files for gooey. Fixed in 0.6.9.4 already.
- Scons: Fix, was not respecting `--quiet` option when running Scons. Fixed in 0.6.9.3 already.
- Scons: Fix, wasn't automatically detecting Scons from promised paths. Fixed in 0.6.9.2 already.

- Scons: Fix, the clcache output parsing wasn't robust enough. Fixed in 0.6.9.1 already.
- Python3.8: Ignore all non-strings provided in doc-string fashion, they are not to be considered.
- Fix, `getattr`, `setattr` and `hasattr` could not be used in finally clauses anymore. Fixed in 0.6.9.1 already.
- Windows: For Python3 enhanced compatibility for Windows no console mode, they need a `sys.stdin` or else e.g. `input` will not be compatible and raise `RuntimeError`.

New Features

- Added experimental support for Python 3.9, in such a way that the CPython3.8 test suite passes now, the 3.9 suite needs investigation still, so we might be missing new features.
- Added experimental support for Onefile mode with `--onefile` that uses `AppImage` on Linux and our own bootstrap binary on Windows. Other platforms are not supported at this time. With this, the standalone folder is packed into a single binary. The Windows variant currently doesn't yet do any compression yet, but the Linux one does.
- Windows: Added downloading of `ccache.exe`, esp. as the other sources so far recommended were not working properly after updates. This is taken from the official project and should be good.
- Windows: Added downloading of matching MinGW64 C compiler, if no other was found, or that was has the wrong architecture, e.g. 32 bits where we need 64 bits.
- Windows: Added ability to copy icon resources from an existing binary with new option `--windows-icon-from-exe`.
- Windows: Added ability to provide multiple icon files for use with different desktop resolutions with new option `--windows-icon-from-ico` that got renamed to disambiguate from other icon options.
- Windows: Added support for requesting UAC admin right with new option `--windows-uac-admin`.
- Windows: Added support for requesting "uiaccess" rights with yet another new option `--windows-uac-uiaccess`.
- Windows: Added ability to specify version info to the binary. New options `--windows-company-name`, `--windows-product-name`, `--windows-file-version`, `--windows-product-version`, and `--windows-file-description` have been added. Some of these have defaults.
- Enhanced support for using the Win32 compiler of MinGW64, but it's not perfect yet and not recommended.
- Windows: Added support for LTO mode for MSVC as well, this seems to allow more optimization.
- Plugins: The numpy plugin now handles matplotlib3 config files correctly.

Optimization

- Use less C variables in dictionary created, not one per key/value pair. This improved scalability of C compilation.
- Use common code for module variable access, leading to more compact code and enhanced scalability of C compilation.
- Use error exit during dictionary creation to release the dictionary, list, tuple, and set in case of an error occurring while they are still under construction. That avoids releases of it in error exists, reducing the generated code size by a lot. This improves scalability of C compilation for generating these.
- Annotate no exception raise for local variables of classes with know dict shape, to avoid useless error exits.

- Annotate no exception exit for `staticmethod` and `classmethod` as they do not check their arguments at all. This makes code generated for classes with these methods much more compact, mainly improving their scalability in C compilation.
- In code generation, prefer `bool` over `nuitka_bool` which allows to annotate exception result, leading to more compact code. Also cleanup so that code generation always go through the C type objects, rather than doing cases locally, adding a C type for `bool`.
- Use common code for C code handling `const None` return only, to cases where there is any immutable constant value returned, avoid code generation for this common case. Currently mutable constants are not handled, this may be added in the future.
- Annotate no exception for exception type checks in handlers for Python2 and no exception if the value has exception type shape for Python3. The exception type shape was newly added. This avoids useless exception handlers in most cases, where the provided exception is just a built-in exception name.
- Improve speed of often used compile time methods on nodes representing constant values, by making their implementation type specific to improve frontend compile time speed, we check e.g. mutable and hashable a lot.
- Provide truth value for variable references, enhancing loop optimization and merge value tracing, to also decide this correctly for values only read, and then changed through attribute, e.g. `append` on lists. This allows many more static optimization.
- Use `staticmethod` for methods in Nuitka nodes to achieve faster frontend compile times where possible.
- Use dedicated helper code for calls with single argument, avoiding the need have a call site local C array of size one, just to pass a pointer to it.
- Added handling of `hash slot`, to predict hashable keys for dictionary and sets.
- Share more slot provision for built-in type shapes from mixin classes, to get them more universally provided, even for special types, where their consideration is unusual.
- Trace "user provided" flag only for constants where it really matters, i.e. for containers and generally potentially large values, but not for every number or boolean value.
- Added lowering of `bytearray` constant values to `bytes` value iteration, while handling constant values for this optimization with dedicated code for improved frontend compilation speed.
- The dict built-in now annotates the dictionary type shape of its result.
- The wrapping side-effects node now passes on the type shape of the wrapped value, allowing for optimization of these too.
- Split `slice` nodes into variants with 1, 2 or 3 arguments, to avoid the overhead of determining which case we have, as well as to save a bit of memory, since these are more frequently used on Python3 for subscript operations. Also annotate their type shape, allowing more optimization.
- Faster dictionary lookups, esp. in cases where errors occur, because we were manually recreating a `KeyError` that is already provided by the dict implementation. This should also be faster, as it avoids a CPython API call overhead on the DLL and they can provide a reference or not for the returned value, simplifying using code.
- Faster dictionary containment checks, with our own dedicated helper, we can use code that won't create an exception when an item is not present at all.
- Faster hash lookups with our own helper, separating cases where we want an exception for non-hashable values or not. These should also be faster to call.
- Avoid acquiring thread state in exception handling that checks if a `StopIteration` occurred, to improved speed on Python3, where it involves locking, but this needs to be applied way more often.

- Make sure checks to debug mode and full compatibility mode are done with the variables introduced, to avoid losing performance due to calls for Nuitka compile time enhancements. This was so far only done partially.
- Split constant references into two base classes, only one of them tracking if the value was provided by the user. This saves compile time memory and avoids the overhead to check if sizes are exceeded in cases they cannot possibly be so.
- The truth value of container creations is now statically known, because the empty container creation is no longer a possibility for these nodes, allowing more optimization for them.
- Optimize the bool built-in with no arguments directory, allow to simplify the node for single argument form to avoid checks if an argument was given.
- Added iteration handles for `xrange` values, and make them faster to create by being tied to the node type, avoiding shared types, instead using the mixin approach. This is in preparation to using them for standard iterator tracing as well. So far they are only used for `any` and `all` decision.
- Added detection if a iterator next can raise, using existing iterator checking which allows to remove needless checks and exception traces. Adding a code variant for calls to next that cannot fail, while tuning the code used for `next` and unpacking next, to use faster exception checking in the C code. This will speed up unpacking performance for some forms of unpacking from known sizes.
- Make sure to use the fastest tuple API possible in all of Nuitka, many place e.g. used `PyTuple_Size`, and one was in a performance critical part, e.g. in code that used when compiled functions as called as a method.
- Added optimized variant for `_PyList_Extend` for slightly faster unpacking code.
- Added optimized variant for `PyList_Append` for faster list contractions code.
- Avoid using `RemoveFileSpec` and instead provide our own code for that task, slightly reducing file size and avoiding to use the `Shlapi` link library.

Tests

- Made reflected test use common cleanup of test folder, which is more robust against Windows locking issues.
- Only output changed CPython output after the forced update of cached value was done, avoiding duplicate or outdated outputs.
- Avoid complaining about exceptions for in-place operations in case they are lowered to non-inplace operations and then raise unsupported, not worth the effort to retain original operator.
- Added generated test for subscript operations, also expanding coverage in generated tests by making sure, conditional paths are both taken by varying the `cond` value.
- Use our own code helper to check if an object has an attribute, which is faster, because it avoids creating exceptions in the first place, instead of removing them afterwards.

Cleanups

- Make sure that code generation always go through the C type objects rather than local `elif` casing of the type. This required cleaning up many of the methods and making code more abstract.
- Added base class for C types without reference counting, so they can share the code that ignores their handling.
- Remove `getConstant` for constant value nodes, use the more general `getCompileTimeConstant` instead, and provide quick methods that test for empty tuple or dict, to use for checking concrete values, e.g. with call operations.

- Unified container creation into always using a factory function, to be sure that existing container creations are not empty.
- Stop using `@calledWithBuiltinArgumentNamesDecorator` where possible, and instead make explicit wrapping or use correct names. This was used to allow e.g. an argument named `list` to be passed from built-in optimization, but that can be done in a cleaner fashion. Also aligned no attributes and the argument names, there was inconsistency there.
- Name mangling was done differently for attribute names and normal names and with non-shared code, and later than necessary, removing this as a step from variable closure taking after initial tree build.
- As part of the icon changes, now handled in Python code, we stop using the `rc` binary and handle all resources ourselves, allowing to remove that code from the Scons side of things.
- Moved file comparison code of standalone mode into file utils function for use in plugins as well.
- Unified how path concatenation is done in Nuitka helper code, there were more or less complete variants, this is making sure, the most capable form is used in all cases.
- Massive cleanup to our scons file, by moving out util code that only scons uses, hacks we apply to speed up scons, and more to separate modules with dedicated interfaces.
- When using `enumerate` we now provide start value of 1 where it is appropriate, e.g. when counting source code lines, rather than adding `count+1` on every usage, making code more readable.

Organisational

- Do not recommend Anaconda on Windows anymore, it seems barely possible to get anything installed on it with a fresh download, due to the resolver literally working for days without finishing, and then reporting conflicts, it would only be usable when starting with Miniconda, but that seems less interesting to users, also gcc 5.2 is way too old these days.
- The commit hook should be reinstalled, since it got improved and adapted for newer git versions.
- Added link to donations to funding document, following a GitHub standard.
- Bumped requirements for development to the latest versions, esp. newer isort.
- Added a rough description of tests to do to add a new CPython test suite, to allow others to take this task in the future.
- Updated the git hook so that Windows and newest git works.
- Make it more clear in the documentation that Microsoft Appstore Python is not supported.

Summary

This is the big release in terms of scalability. The optimization in this release mostly focused on getting things that cause increased compile times sorted out. A very important fix avoids loop optimization to leak into global passes of all modules unnecessarily, but just as important, generated code now is much better for the C compiler to consume in observed problematic cases.

More optimization changes are geared towards reducing Nuitka frontend compile time, which could also be a lot in some cases, ending up specializing more constant nodes and how they expose themselves to optimization.

Other optimization came from supporting Python 3.9 and things come across during the implementation of that feature, e.g. to be able to make differences with unpacking error messages, we provide more code to handle it ourselves, and to manually optimize how to interact with e.g. `list` objects.

For Windows, the automatic download of `ccache` and a matching MinGW64 if none was found, is a new step, that should lower the barrier of entry for people who have no clue what a C compiler is. More

changes are bound to come in this field with future releases, e.g. making a minimum version requirement for gcc on Windows that excludes unfit C compilers.

All in all, this release should be taken as a major cleanup, resolving many technical debts of Nuitka and preparing more optimization to come.

Nuitka Release 0.6.9

This releases contains important bug fixes for regressions of the 0.6.8 series which had relatively many problems. Not all of these could be addressed as hotfixes, and other issues were even very involved, causing many changes to be necessary.

There are also many general improvements and performance work for tracing and loops, but the full potential of this will not be unlocked with this release yet.

Bug Fixes

- Fix, loop optimization sometimes didn't determinate, effectively making Nuitka run forever, with no indication why. This has been fixed and a mechanism to give up after too many attempts has been added.
- Fix, closure taking object allowed a brief period where the garbage collector was exposed to uninitialized objects. Fixed in 0.6.8.1 already.
- Python3.6+: Fix corruption for exceptions thrown into `asyncgen`. Fixed in 0.6.8.1 already.
- Fix, deleting variables detected as C type `bool` could raise an `UnboundLocalError` that was wrong. Fixed in 0.6.8.1 already.
- Python3.8.3+: Fix, future annotations parsing was using hard coded values that were changed in CPython, leading to errors.
- Windows: Avoid encoding issues for Python3 on more systems, by going from wide characters to unicode strings more directly, avoiding an encoding as UTF-8 in the middle. Fixed in 0.6.8.2 already.
- Windows: Do not crash when warning about uninstalled MSVC using Python3. This is a Scons bug that we fixed. Fixed in 0.6.8.3 already.
- Standalone: The output of dependency walker should be considered as "latin1" rather than UTF-8. Fixed in 0.6.8.3 already.
- Standalone: Added missing hidden dependencies for `flask`. Fixed in 0.6.8.1 already.
- Standalone: Fixed `win32com.client` on Windows. Fixed in 0.6.8.1 already.
- Standalone: Use `pkgutil` to scan encoding modules, properly ignoring the same files as Python does in case of garbage files being there. Fixed in 0.6.8.2 already.
- Plugins: Enabling a plugin after the filename to compile was given, didn't allow for arguments to the passed, causing problems. Fixed in 0.6.8.3 already.
- Standalone: The `certifi` data file is now supported for all modules using it and not only some.
- Standalone: The bytecode for the standard library had filenames pointing to the original installation attached. While these were not used, but replaced at run time, they increased the size of the binary, and leaked information.
- Standalone: The path of `sys.executable` was not `None`, but pointing to the original executable, which could also point to some temporary virtualenv directory and therefore not exist, also it was leaking information about the original install.
- Windows: With the MSVC compiler, elimination of duplicate strings was not active, causing even unused strings to be present in the binary, some of which contained file paths of the Nuitka installation.

- Standalone: Added support for pyglet.
- Plugins: The command line handling for Pmw plugin was using wrong defaults, making it include more code than necessary, and to crash if it was not there.

New Features

- Windows: Added support for using Python 2.7 through a symlink too. This was already working for Python3, but a scon's problem prevented this from working.
- Caching of compiled C files is now checked with `ccache` and `clcache`, and added automatically where possible, plus a report of the success is made. This can accelerate the re-compile very much, even if you have to go through Nuitka compilation itself, which is not (yet) cached.
- Added new `--quiet` option that will disable informational traces that are going to become more.
- The Clang from MSVC installation is now picked up for both 32 and 64 bits and follows the new location in latest Visual Studio 2019.
- Windows: The `ccache` from Anaconda is now supported as well as the one from `msys64`.

Optimization

- The value tracing has become more correct with loops and in general less often inhibits optimization. Escaping of value traces is now a separate trace state allowing for more appropriate handling of actual unknowns.
- Memory used for value tracing has been lowered by removing unnecessary states for traces, that we don't use anymore.
- Windows: Prevent scon's from scanning for MSVC when asked to use MinGW64. This avoids a performance loss doing something that will then end up being unused.
- Windows: Use function level linking with MSVC, this will allow for smaller binaries to be created, that don't have to include unused helper functions.

Cleanups

- The scon's file now uses Nuitka utils functions and is itself split up into several modules for enhanced readability.
- Plugin interfaces for providing extra entry points have been cleaned up and now named tuples are used. Backward compatibility is maintained though.

Organisational

- The use of the logging module was replaced with more of our custom tracing and we now have the ability to write the optimization log to a separate file.
- Old style plugin options are now detected and reported as a usage error rather than unknown plugin.
- Changed submodules to use git over https, so as to not require ssh which requires a key registered and causes problems with firewalls too.
- More correct Debian copyright file, made formatting of emails in source code consistent.
- Added repository for Ubuntu focal.

Summary

The main focus of this release has been bug fixes with only a little performance work due to the large amount of regressions and other findings from the last release.

The new constants loading for removes a major scalability problem. The checked and now consistently possible use of `ccache` and `clcache` allows for much quicker recompilation. Nuitka itself can still be slow in some cases, but should have seen some improvements too. Scalability will have to remain a focus for the next releases too.

The other focus, was to make the binaries contain no original path location, which is interesting for standalone mode. Nuitka should be very good in this area now.

For optimization, the new loop code is again better. But it was also very time consuming, to redo it, yet again. This has prevented other optimization to be added.

And then for correctness, the locals scope work, while very invasive, was necessary, to handle the usage of locals inside of contractions, but also will be instrumental for function inlining to become generally available.

So, ultimately, this release is a necessary intermediate step. Upcoming releases will be able to focus more clearly on run time performance again as well as on scalability for generated C code.

Nuitka Release 0.6.8

This releases contains important general improvements and performance improvements and enhanced optimization as well as many bug fixes that enhance the Python 3.8 compatibility.

Bug Fixes

- Python3.5+: Fix, coroutines and `asyncgen` could continue iteration of awaited functions, even after their return, leading to wrong behaviour.
- Python3.5+: Fix, absolute imports of names might also refer to modules and need to be handled for module loading as well.
- Fix, the `fromlist` of imports could loose references, potentially leading to corruption of contained strings.
- Python3.8: Fix, positional only arguments were not enforced to actually be that way.
- Python3.8: Fix, complex calls with star arguments that yielded the same value twice, were not yet caught.
- Python3.8: Fix, evaluation order for nested dictionary contractions was not followed yet.
- Windows: Use short paths, these work much better to load extension modules and TCL parts of TkInter cannot handle unicode paths at all. This makes Nuitka work in locations, where normal Python cannot.
- Windows: Fixup dependency walker in unicode input directories.
- Standalone: Use frozen module loader only at `libpython` initialisation and switch to built-in bytecode loader that is more compatible afterwards, increasing compatibility.
- Standalone: Fix for `pydantic` support.
- Standalone: Added missing hidden dependency of `uvicorn`.
- Fix, the parser for `.pyi` files couldn't handle multiline imports.
- Windows: Derive linker arch of Python from running binary, since it can happen that the Python binary is actually a script.

- Fixup static linking with `libpython.a` that contains `main.o` by making our colliding symbols for `Py_GetArgcArgv` weak.
- Python3.7: Fix misdetection as `asyncgen` for a normal generator, if the iterated value is `async`.
- Distutils: Fix `build_nuitka` for modules under nested namespaces.
- OpenBSD: Follow usage of `clang` and other corrections to make accelerated mode work.
- macOS: Fixup for standalone mode library scan.
- Fix, the logging of `--show-modules` was broken.
- Windows: Enable `/bigobj` mode for MSVC for large compilations to work.
- Windows: Fixup crash in warning with `pefile` dependency manager.
- Windows: Fixup `win32com` standalone detection of other Python version `win32com` is in `system PATH`.
- Fix, the python flag for static hashes didn't have the intended effect.
- Fix, generators may be resurrected in the cause of their destruction, and then must not be released.
- Fix, method objects didn't implement the methods `__reduce__` and `__reduce_ex__` necessary for pickling them.
- Windows: Fix, using a Python installation through a symlink was not working.
- Windows: Fix, icon paths that were relative were not working anymore.
- Python3.8: Detect duplicate keywords yielded from star arguments.
- Fix, methods could not be pickled.
- Fix, generators, coroutines and `asyncgen` might be resurrected during their release, allow for that.
- Fix, frames need to traverse their attached locals to be released in some cases.

New Features

- Plugin command line handling now allows for proper `optparse` options to be used, doing away with special parameter code for plugins. The arguments now also become automatically passed to the instantiations of plugins.
Loading and creation of plugins are now two separate phases. They are loaded when they appear on the command line and can add options in their own group, even required ones, but also with default values.
- Started using logging with name-spaces. Applying logging per plugin to make it easier to recognize which plugin said what. Warnings are now colored in red.
- Python3.5+: Added support for two step module loading, making Nuitka loading even more compatible.
- Enhanced import tracing to work on standalone binaries in a useful manner, allow to compare with normal binaries.
- Fix, the `setattr` built-in was leaking a reference to the `None` value.

Optimization

- Proper loop SSA capable of detecting shapes with an incremental initial phase and a final result of alternatives for variables written in the loop. This detects shapes of manual integer incrementing loops correctly now, it doesn't see through iterators yet, but this will come too.

- Added type shapes for all operations and all important built-in types to allow more compile time optimization and better target type selection.
- Target type code generation was expanded from manual usage with conditions to all operations allowing to get at bool target values more directly.
- For in-place operations, there is the infrastructure to generate them for improved performance, but so far it's only used for Python2 int, and not for the many types normal operations are supported.
- Force usage of C boolean type for all indicator variables from the re-formulation. In some cases, we are not yet there with detections, and this gives instant benefit.
- Complex constants didn't annotate their type shape, preventing compile time optimization for them.
- Python3.8: Also support vectorcall for compiled method objects. These are rarely used in new Python, but can make a difference.
- Remove loops that have only a final break. This happens in static optimization in some cases, and allows more optimization to be done.
- Avoid using a preparing a constant tuple value for calls with only constant arguments.
- Avoid using `PyErr_Format` where it's not necessary by adding specialized helpers for common cases.
- Detect `del` statements that will raise an exception and replace with that.
- Exception matching is boolean shape, allowing for faster code generation.
- Disable recursion checks outside of full compat mode.
- Avoid large blocks for conditional statements that only need to enclose the condition evaluation.
- Added shortcuts for interactions between compiled generator variants, to avoid calls to their C methods with argument passing, etc.

Organisational

- Updated Developer Manual with changes that happened, removing the obsolete language choice section.
- Added 3.8 support mentions in even more places.
- The mailing list has been deleted. We now prefer Gitter chat and GitHub issues for discussions.
- Visual Code recommended extensions are now defined as such in the project configuration and you will be prompted to install them.
- Visual Code environments for `Py38` and `Py27` were added for easier switch.
- Catch usage of Python from the Microsoft App Store, it is not supported and seems to limit access to the Python installation for security reasons that make support impossible.
- Make it clear that `--full-compat` should not be used in help output.
- Added instructions for MSVC runtimes and standalone compilation to support Windows 7.
- More complete listing of copyright holders for Debian.
- Updated to newer black and PyLint.
- Enhanced gcc version check, properly works with gcc 10 and higher.

Tests

- Pylint cleanups for some of the tests.
- Added test for loading of user plugins.

- Removed useless outputs for `search` mode skipping non-matches.

Cleanups

- Limit command line handling for multiprocessing module to when the plugin is actually used, avoiding useless code of Windows binaries.
- Pylint cleanup also foreign code like `oset` and `odict`.
- In preparation of deprecating the alternative, `--enable-plugin` has become the only form used in documentation and tests.
- Avoid numeric pylint symbols more often.
- Distutils: Cleanup module name for distutils commands, these are not actually enforced by distutils, but very ugly in our coding conventions.
- The "cannot get here" code to mark unreachable code has been improved and no longer needs an identifier passed, but uses the standard C mechanism for that.
- Removed accessors for lookup sources from nodes, allowing for faster usage and making sure, lookups are only done where needed.

Summary

This release is huge in terms of bugs fixed, but also extremely important, because the new loop SSA and type tracing, allows for many more specialized code usages. We now can trace the type for some loops to be specifically an integer or long value only, and will become able to generate code that avoids using Python objects, in these cases.

Once that happens, the performance will make a big jump. Future releases will have to consolidate the current state, but it is expected that at least an experimental addition of C type `float` or `long` can be added, add to that `iterator` type shape and value analysis, and an actual jump in performance can be expected.

Nuitka Release 0.6.7

This release contains bug fixes and improvements to the packaging, for the RPM side as well as for Debian, to cover Python3 only systems as they are now becoming more common.

Bug Fixes

- Compatibility: The value of `__module__` for extension modules was not dependent into which package the module was loaded, it now is.
- Anaconda: Enhanced detection of Anaconda for Python 3.6 and higher.
- CentOS6: Detect gcc version to allow saving on macro memory usage, very old gcc didn't have that.
- Include Python3 for all Fedora versions where it works as well as for openSUSE versions 15 and higher.
- Windows: Using short path names to interact with Scons avoids problems with unicode paths in all cases.
- macOS: The usage of `install_name_tool` could sometimes fail due to length limits, we now increase it at link time.
- macOS: Do not link against `libpython` for module mode. This prevented extension modules from actually being usable.
- Python3.6: Follow coroutine fixes in our asyncgen implementation as well.

- Fix, our version number handling could overflow with minor versions past 10, so we limited it for now.

New Features

- Added support for Python 3.8, the experimental was already there and pretty good, but now added the last obscure features too.
- Plugins can now provide C code to be included in the compilation.
- Distutils: Added targets `build_nuitka` and `install_nuitka` to complement `bdist_nuitka`, so we support software other than wheels, e.g. RPM packaging that compiles with Nuitka.
- Added support for `lldb` the Clang debugger with the `--debugger` mode.

Optimization

- Make the file prefix map actually work for gcc and clang, and compile files inside the build folder, unless we are running in debugger mode, so we use `ccache` caching across different compilations for at least the static parts.
- Avoid compilation of `__frozen.c` in accelerated mode, it's not used.
- Prefer using the inline copy of `scons` over systems `scons`. The later will only be slower. Use the fallback to external `scons` only from the Debian packages, since there we consider it forbidden to include software as a duplicate.

Organisational

- Added recommended plugins for Visual Code, replacing the list in the Developer Manual.
- Added repository for Fedora 30 for download.
- Added repository for CentOS 8 for download.
- Updated inline copy of `Scons` used for Python3 to 3.1.2, which is said to be faster for large compilations.
- Removed Eclipse setup from the manual, it's only inferior at this point and we do not use it ourselves.
- Debian: Stop recommending `PyQt5` in the package, we no longer use it for built-in GUI that was removed.
- Debian: Bumped the standards version and modernized the packaging, solving a few warnings during the build.

Cleanups

- `Scons`: Avoid to add Unix only include paths on Windows.
- `Scons`: Have the static source code in a dedicated folder for clarity.

Tests

- Added tests to GitHub Actions, for the supported Python versions for all of Linux, macOS and Windows, covering the later publicly for the first time. We use Anaconda on macOS for the tests now, rather than Homebrew.
- Enable IO encoding to make sure we use UTF-8 for more test suites that actually need it in case of problems.

- Comparing module outputs now handles segfaults by running in the debugger too.

Summary

This release adds full support for Python 3.8 finally, which took us a while, and it cleans up a lot on the packaging side. There aren't that many important bug fixes, but it's still nice to this cleaned up.

We have important actual optimization in the pipeline that will apply specialization to target types and for comparison operations. We expect to see actual performance improvements in the next release again.

Nuitka Release 0.6.6

This release contains huge amounts of crucial bug fixes all across the board. There is also new optimization and many organisational improvements.

Bug Fixes

- Fix, the top level module must not be bytecode. Otherwise we end up violating the requirement for an entry point on the C level.
- Fix, avoid optimizing calls with default values used. This is not yet working and needed to be disabled for now.
- Python3: Fix, missing keyword only arguments were not enforced to be provided keyword only, and were not giving the compatible error message when missing.
- Windows: Find `win32com` DLLs too, even if they live in sub folders of `site-packages`, and otherwise not found. They are used by other DLLs that are found.
- Standalone: Fixup for problem with standard library module in most recent Anaconda versions.
- Scons: Fix, was using `CXXFLAGS` and `CPPFLAGS` even for the C compiler, which is wrong, and could lead to compilation errors.
- Windows: Make `--clang` limited to `clang-cl.exe` as using it inside a MinGW64 is not currently supported.
- Standalone: Added support for using `lib2to2.pgen`.
- Standalone: Added paths used by openSUSE to the Tcl/Tk plugin.
- Python3.6+: Fix, the `__main__` package was `None`, but should be `" "` which allows relative imports from itself.
- Python2: Fix, compile time optimization of floor division was using normal division.
- Python3: Fix, some run time operations with known type shapes, were falsely reporting error message with `unicode` or `long`, which is of course not compatible.
- Fix, was caching parent package, but these could be replaced e.g. due to bytecode demotion later, causing crashes during their optimization.
- Fix, the value of `__compiled__` could be corrupted when being deleted, which some modules wrappers do.
- Fix, the value of `__package__` could be corrupted when being deleted.
- Scons: Make sure we can always output the compiler output, even if it has a broken encoding. This should resolve MSVC issues on non-English systems, e.g. German or Chinese.
- Standalone: Support for newest `sklearn` was added.
- macOS: Added resolver for run time variables in `otool` output, that gets PyQt5 to work on it again.
- Fix, floor division of run time calculations with float values should not result in `int`, but `float` values instead.

- Standalone: Enhanced support for `boto3` data files.
- Standalone: Added support for `osgeo` and `gdal`.
- Windows: Fix, there were issues with spurious errors attaching the constants blob to the binary due to incorrect C types provided.
- Distutils: Fix, need to allow `/` as separator for package names too.
- Python3.6+: Fix reference losses in `asyncgen` when throwing exceptions into them.
- Standalone: Added support for `dill`.
- Standalone: Added support for `scikit-image` and `skimage`.
- Standalone: Added support for `weasyprint`.
- Standalone: Added support for `dask`.
- Standalone: Added support for `pendulum`.
- Standalone: Added support for `pytz` and `pytzdata`.
- Fix, `--python-flags=no_docstrings` no longer implies disabling the assertions.

New Features

- Added experimental support for Python 3.8, there is only very few things missing for full support.
- Distutils: Added support for packages that are in a namespace and not just top level.
- Distutils: Added support for single modules, not only packages, by supporting `py_modules` as well.
- Distutils: Added support for distinct namespaces.
- Windows: Compare Python and C compiler architecture for MSVC too, and catch the most common user error of mixing 32 and 64 bits.
- Scons: Output variables used from the outside, so the debugging is easier.
- Windows: Detect if clang installed inside MSVC automatically and use it if requested via `--clang` option. This is only the 32 bits variant, but currently the easy way to use it on Windows with Nuitka.

Optimization

- Loop variables were analysed, but results were only available on the inside of the loop, preventing many optimization in these cases.
- Added optimization for the `abs` built-in, which is also a numerical operator.
- Added optimization for the `all` built-in, adding a new concept of iteration handle, for efficient checking that avoids looking at very large sequences, of which properties can still be known.

```
all(range(1, 100000)) # no need to look at all of them
```

- Added support for optimizing `ImportError` construction with keyword-only arguments. Previously only used without these were optimized.

```
raise ImportError(path="lala", name="lele") # now optimized
```

- Added manual specialization for single argument calls, solving a TODO, as these will be very frequent.
- Memory: Use single child form of node class where possible, the general class now raises an error if used with used with only one child name, this will use less memory at compile time.

- Memory: Avoid list for non-local declarations in every function, these are very rare, only have it if absolutely necessary.
- Generate more compact code for potential `NameError` exceptions being raised. These are very frequent, so this improves scalability with large files.
- Python2: Annotate comparison of `None` with `int` and `str` types as not raising an exception.
- Shared empty body functions and generators.
One shared implementation for all empty functions removes that burden from the C compiler, and from the CPU instruction cache. All the shared C code does is to release its arguments, or to return an empty generator function in case of generator.
- Memory: Added support for automatic releases of parameter variables from the node tree. These are normally released in a try finally block, however, this is now handled during code generation for much more compact C code generated.
- Added specialization for `int` and `long` operations `%`, `<<`, `>>`, `|`, `&`, `^`, `**`, `@`.
- Added dedicated nodes for representing and optimizing based on shapes for all binary operations.
- Disable gcc macro tracing unless in debug mode, to save memory during the C compilation.
- Restored Python2 fast path for `int` with unknown object types, restoring performance for these.

Cleanups

- Use dedicated `ModuleName` type that makes the tests that check if a given module name is inside a namespace as methods. This was hard to get right and as a result, adopting this fixed a few bugs and or inconsistent results.
- Expand the use of `nuitka.PostProcessing` to cover all actions needed to get a runnable binary. This includes using `install_name_tool` on macOS standalone, as well copying the Python DLL for acceleration mode, cleaning the `x` bit for module mode. Previously only a part of these lived there.
- Avoid including the definitions of dynamically created helper functions in the C code, instead just statically declare the ones expected to be there. This resolves Visual Code complaining about it, and should make life also easier for the compiler and caches like `ccache`.
- Create more helper code in closer form to what `clang-format` does, so they are easier to compare to the static forms. We often create hard coded variants for few arguments of call functions, and generate them for many argument variations.
- Moved setter/getter methods for Nuitka nodes consistently to the start of the node class definitions.
- Generate C code much closer to what `clang-format` would change it to be.
- Unified calling `install_name_tool` on macOS into one function that takes care of all the things, including e.g. making the file writable.
- Debug output from scons should be more consistent and complete now.
- Sort files for compilation in scons for better reproducible results.
- Create code objects version independent, avoiding python version checks by pre-processor, hiding new stuff behind macros, that ignore things on older Python versions.

Tests

- Added many more built-in tests for increased coverage of the newly covered ones, some of them being generic tests that allow to test all built-ins with typical uses.
- Many tests have become more PyLint clean as a result of work with Visual Code and it complaining about them.

- Added test to check PyPI health of top 50 packages. This is a major GSoC 2019 result.
- Output the standalone directory contents for Windows too in case of a failure.
- Added generated tests to fully cover operations on different type shapes and their errors as well as results for typical values.
- Added support for testing against installed version of Nuitka.
- Cleanup up tests, merging those for only Python 3.2 with 3.3 as we no longer support that version anyway.
- Execute the Python3 tests for macOS on Travis too.

Organisational

- The donation sponsored machine called `donatix` had to be replaced due to hardware breakage. It was replaced with a Raspberry-Pi 4.
- Enhanced plugin documentation.
- Added description of the git workflow to the Developer Manual.
- Added checker script `check-nuitka-with-codespell` that reports typos in the source code for easier use of `codespell` with Nuitka.
- Use newest PyLint and clang-format.
- Also check plugin documentation files for ReST errors.
- Much enhanced support for Visual Code configuration.
- Trigger module code is now written into the build directory in debug mode, to aid debugging.
- Added deep check function that descends into tuples to check their elements too.

Summary

This release comes after a long time of 4 months without a release, and has accumulated massive amounts of changes. The work on CPython 3.8 is not yet complete, and the performance work has yet to show actual fruit, but has also progressed on all fronts. Connecting the dots and pieces seems not far away.

Nuitka Release 0.6.5

This release contains many bug fixes all across the board. There is also new optimization and many organisational improvements.

Bug Fixes

- Python3.4+: Fixed issues with modules that exited with an exception, that could lead to a crash, dealing with their `__spec__` value.
- Python3.4+: The `__loader__` method `is_package` had the wrong signature.
- Python3.6+: Fix for `async with` being broken with uncompiled generators.
- Python3.5+: Fix for `coroutines` that got their awaited object closed behind their back, they were complaining with `RuntimeError` should they be closed themselves.
- Fix, constant values `None` in a bool target that could not be optimized away, lead to failure during code generation.

```
if x() and None:
    ...
```

- Standalone: Added support for sha224, sha384, sha512 in crypto package.
- Windows: The icon wasn't properly attached with MinGW64 anymore, this was a regression.
- Windows: For compiler outputs, also attempt preferred locale to interpret outputs, so we have a better chance to not crash over MSVC error messages that are not UTF-8 compatible.
- macOS: Handle filename collisions for generated code too, Nuitka now treats all filesystems for all OS as case insensitive for this purpose.
- Compatibility: Added support for tolerant `del` in class exception handlers.

```
class C:
    try:
        ...
    except Exception as e:
        del e

    # At exception handler exit, "e" is deleted if still assigned
```

We already were compatible for functions and modules here, but due to the special nature of class variables really living in dictionaries, this was delayed. But after some other changes, it was now possible to solve this TODO.

- Standalone: Added support for Python3 variant of Pmw.
- Fix, the NumPy plugin now handles more installation types.
- Fix, the qt plugin now handles multiple library paths.
- Fix, need `libm` for some Anaconda variants too.
- Fix, left over bytecode from plugins could crash the plugin loader.
- Fix, `pkgutil.iter_packages` is now working for loaded packages.

New Features

- Python3.8: Followed some of the changes and works with beta2 as a Python 3.7, but none of the new features are implemented yet.
- Added support for Torch, Tensorflow, Gevent, Sklearn, with a new Nuitka plugin.
- Added support for "hinted" compilation, where the used modules are determined through a test run.
- Added support for including TCL on Linux too.

Optimization

- Added support for the `any` built-in. This handles a wide range of type shapes and constant values at compile time, while also having optimized C code.
- Generate code for some `CLONG` operations in preparation of eventual per expression C type selection, it then will allow to avoid objects in many instances.
- Windows: Avoid creating link libraries for MinGW64 as these have become unnecessary in the mean time.

- Packages: Do not export entry points for all included packages, only for the main package name it is importable as.

Organisational

- Added support for Visual Studio 2019 as a C compiler backend.
- Improved plugin documentation describing how to create plugins for Nuitka even better.
- There is now a mode for running the tests called `all` which will execute all the tests and report their errors, and only fail at the very end. This doesn't avoid wasting CPU cycles to report that e.g. all tests are broken, but it allows to know all errors before fixing some.
- Added repository for Fedora 30 for download.
- Added repository for openSUSE 15.1 for download.
- Ask people to compile hello world program in the GitHub issue template, because many times, they have setup problems only.
- Visual Studio Code is now the recommended IDE and has integrated configuration to make it immediately useful.
- Updated internal copy of Scons to 3.1.0 as it incorporates many of our patches.
- Changed wordings for optimization to use "lowering" as the only term to describe an optimization that simplifies.

Cleanups

- Plugins: Major refactoring of Nuitka plugin API.
- Plugins: To locate module kind, use core Nuitka code that handles more cases.
- The test suite runners are also now auto-formatted and checked with PyLint.
- The Scons file is now PyLint clean too.
- Avoid `build_definitions.h` to be included everywhere, in that it's only used in the main program part. This makes C linter hate us much less for using a non-existent file.

Tests

- Run the tests using Travis on macOS for Python2 too.
- More standalone tests have been properly whitelisting to cover openssl usage from local system.
- Disabled PySide2 test, it's not useful to fail and ignore it.
- Tests: Fixups for coverage testing mode.
- Tests: Temporarily disable some checks for constants code in reflected tests as it only exposes `marshal` not being deterministic.

Summary

This release is huge again. Main points are compatibility fixes, esp. on the coroutine side. These have become apparently very compatible now and we might eventually focus on making them better.

Again, GSoC 2019 is also showing effects, and will definitely continue to do so in the next release.

Many use cases have been improved, and on an organisational level, the adoption of Visual Studio Code seems an huge improvement to have a well configured IDE out of the box too.

In upcoming releases, more built-ins will be optimized, and hopefully the specialization of operations will hit more and more code with more of the infrastructure getting there.

Nuitka Release 0.6.4

This release contains many bug fixes all across the board. There is also new optimization and many organisational improvements.

Bug Fixes

- When linking very large programs or packages, with gcc compiler, Scons can produce commands that are too large for the OS. This happens sooner on the Windows OS, but also on Linux. We now have a workaround that avoids long command lines by using `@sources.tmp` syntax.
- Standalone: Remove temporary module after its use, instead of keeping it in `sys.modules` where e.g. `Quart` code tripped over its `__file__` value that is illegal on Windows.
- Fixed non-usage of our enhanced detection of gcc version for compilers if given as a full path.
- Fixed non-detection of `gnu-cc` as a form of gcc compiler.
- Python3.4: The `__spec__` value corrections for compiled modules was not taking into account that there was a `__spec__` value, which can happen if something is wrapping imported modules.
- Standalone: Added implicit dependencies for `passlib`.
- Windows: Added workaround for OS command line length limit in compilation with MinGW64.
- Python2: Revive the `enum` plugin, there are backports of the buggy code it tries to patch up.
- Windows: Fixup handling of SxS with non zero language id, these occur e.g. in Anaconda.
- Plugins: Handle multiple PyQt plugin paths, e.g. on openSUSE this is done, also enhanced finding that path with Anaconda on Windows.
- Plugins: For `multiprocessing` on Windows, allow the `.exe` suffix to not be present, which can happen when ran from command line.
- Windows: Better version checks for DLLs on Python3, the `ctypes` helper code needs more definitions to work properly.
- Standalone: Added support for both `pycryptodome` and `pycryptodomex`.
- Fix, the `chr` built-in was not giving fully compatible error on non number input.
- Fix, the `id` built-in doesn't raise an exception, but said otherwise.
- Python3: Proper C identifiers for names that fit into `latin-1`, but are not `ascii` encodings.

New Features

- Windows: Catch most common user error of using compiler from one architecture against Python from another. We now check those and compare it, and if they do not match, inform the user directly. Previously the compilation could fail, or the linking, with cryptic errors.
- Distutils: Using `setuptools` and its runners works now too, not merely only pure distutils.
- Distutils: Added more ways to pass Nuitka specific options via distutils.
- Python3.8: Initial compatibility changes to get basic tests to work.

Organisational

- Nuitka is participating in the GSoC 2019 with 2 students, Batakrishna and Tommy.

- Point people creating PRs to using the `pre-commit` hook in the template. Due to making the style issues automatic, we can hope to encounter less noise and resulting merge problems.
- Many improvements to the `pre-commit` hook were done, hopefully completing its development.
- Updated to latest `pylint`, `black`, and `isort` versions, also added `codespell` to check for typos in the source code, but that is not automated yet.
- Added description of how to use experimental flags for your PRs.
- Removed mirroring from Bitbucket and Gitlab, as we increasingly use the GitHub organisation features.
- Added support for Ubuntu Disco, removed support for Ubuntu Artful packages.

Optimization

- Windows: Attach data blobs as Windows resource files directly for programs and avoid using C data files for modules or MinGW64, which can be slow.
- Specialization of helper codes for `+` is being done for more types and more thoroughly and fully automatic with Jinja2 templating code. This does replace previously manual code.
- Added specialization of helper codes for `*` operation which is entirely new.
- Added specialization of helper codes for `-` operation which is entirely new.
- Dedicated nodes for specialized operations now allow to save memory and all use type shape based analysis to predict result types and exception control flow.
- Better code generation for boolean type values, removing error checks when possible.
- Better static analysis for even more type operations.

Cleanups

- Fixed many kinds of typos in the code base with `codespell`.
- Apply automatic formatting to more test runner code, these were previously not done.
- Avoid using `shutil.copytree` which fails to work when directory already exists, instead provide `nuitka.util.FileOperations.copyTree` and use that exclusively.

Tests

- Added new mode of operation to test runners, `only` that executes just one test and stops, useful during development.
- Added new mechanism for standalone tests to expression modules that need to be importable, or else to skip the test by a special comment in the file, instead of by coded checks in the test runner.
- Added also for more complex cases, another form of special comment, that can be any expression, that decides if the test makes sense.
- Cover also `setuptools` in our `distutils` tests and made the execution more robust against variable behavior of `distutils` and `setuptools`.
- Added standalone test for `Urllib3`.
- Added standalone test for `rsa`.
- Added standalone test for `Pmw`.
- Added standalone test for `passlib`.

Summary

Again this release is a sign of increasing adoption of Nuitka. The GSoC 2019 is also showing effects, definitely will in the next release.

This release has a lot of new optimization, called specialization, but for it to really used, in many instances, we need to get away from working on C types for variables only, and get to them being used for expressions more often. Otherwise much of the new special code is not used for most code.

The focus of this release has been again to open up development further and to incorporate findings from users. The number of fixes or new use cases working is astounding.

In upcoming releases, new built-ins will be optimized, and specialization of operations will hit more and more code now that the infrastructure for it is in place.

Nuitka Release 0.6.3

This has a focus on organisational improvements. With more and more people joining Nuitka, normal developers as well as many GSoC 2019 students, the main focus was to open up the development tools and processes, and to improve documentation.

That said, an impressive amount of bug fixes was contributed, but optimization was on hold.

Bug Fixes

- Windows: Added support for running compiled binaries in unicode path names.
- Standalone: Added support for crytodomex and pycparser packages.
- Standalone: Added support for OpenSSL support in PyQt on Windows.
- Standalone: Added support for OpenGL support with QML in PyQt on Windows.
- Standalone: Added support for SciPy and extended the NumPy plugin to also handle it.
- UI: The option `--plugin-list` still needed a positional argument to work.
- Make sure `sys.base_prefix` is set correctly too.
- Python3: Also make sure `sys.exec_prefix` and `sys.base_exec_prefix` are set correctly.
- Standalone: Added platform plugins for PyQt to the default list of sensible plugins to include.
- Fix detection of standard library paths that include `. .` path elements.

Optimization

- Avoid using static C++ runtime library when using MinGW64.

New Features

- Plugins: A plugin may now also generate data files on the fly for a given module.
- Added support for FreeBSD/PowerPC arch which still uses `gcc` and not `clang`.

Organisational

- Nuitka is participating in the GSoC 2019.
- Added documentation on how to create or use Nuitka plugins.
- Added more API doc to functions that were missing them as part of the ongoing effort to complete it.

- Updated to latest PyLint 2.3.1 for checking the code.
- Scons: Using newer Scons inline copy with Python 2.7 as, the old one remains only used with Python 2.6, making it easier to know the relevant code.
- Auto-format was very much enhanced and handles C and ReST files too now. For Python code it does pylint comment formatting, import statement sorting, and blackening.
- Added script `misc/install-git-hooks.py` that adds a commit hook that runs auto-format on commit. Currently it commits unstaged content and therefore is not yet ready for prime time.
- Moved adapted CPython test suites to [GitHub repository under Nuitka Organisation](#).
- Moved Nuitka-website repository to [GitHub repository under Nuitka Organisation](#).
- Moved Nuitka-speedcenter repository to [GitHub repository under Nuitka Organisation](#).
- There is now a [Gitter chat for Nuitka community](#).
- Many typo and spelling corrections on all the documentation.
- Added short installation guide for Nuitka on Windows.

Cleanups

- Moved commandline parsing helper functions from common code helpers to the main program where of course their only usage is.
- Moved post processing of the created standalone binary from main control to the freezer code.
- Avoid using `chmod` binary to remove executable bit from created extension modules.
- Windows: Avoid using `rt.exe` and `mt.exe` to deal with copying the manifest from the `python.exe` to created binaries. Instead use new code that extracts and adds Windows resources.
- Fixed many `ResourceWarnings` on Python3 by improved ways of handling files.
- Fixed deprecation warnings related to not using `collections.abc`.
- The runners in `bin` directory are now formatted with `black` too.

Tests

- Detect Windows permission errors for two step execution of Nuitka as well, leading to retries should they occur.
- The salt value for CPython cached results was improved to take more things into account.
- Tests: Added more trick assignments and generally added more tests that were so far missing.

Summary

With the many organisational changes in place, my normal work is expected to resume for after and yield quicker improvements now.

It is also important that people are now enabled to contribute to the Nuitka web site and the Nuitka speedcenter. Hope is to see more improvements on this otherwise neglected areas.

And generally, it's great to see that a community of people is now looking at this release in excitement and pride. Thanks to everybody who contributed!

Nuitka Release 0.6.2

This release has a huge focus on organisational things. Nuitka is growing in terms of contributors and supported platforms.

Bug Fixes

- Fix, the Python flag `--python-flag=-O` was removing doc strings, but that should only be done with `--python-flag=-OO` which was added too.
- Fix, accelerated binaries failed to load packages from the `virtualenv` (not `venv`) that they were created and ran with, due to not propagating `sys.prefix`.
- Standalone: Do not include `plat-*` directories as frozen code, and also on some platforms they can also contain code that fails to import without error.
- Standalone: Added missing implicit dependency needed for newer NumPy versions.

New Features

- Added support for Alpine Linux.
- Added support for MSYS2 based Python on Windows.
- Added support for Python flag `--python flag=-OO`, which allows to remove doc strings.
- Added experimental support for `pefile` based dependency scans on Windows, thanks to Orsirir for this contribution.
- Added plugin for proper Tkinter standalone support on Windows, thanks to Jorj for this contribution.
- There is now a `__compiled__` attribute for each module that Nuitka has compiled. Should be like this now, and contains Nuitka version information for you to use, similar to what `sys.version_info` gives as a `namedtuple` for your checks.

```
__nuitka_version__(major=0, minor=6, micro=2, releaselevel="release")
```

Optimization

- Experimental code for variant types for `int` and `long` values, that can be plain C value, as well as the `PyObject *`. This is not yet completed though.
- Minor refinements of specialized code variants reducing them more often the actual needed code.

Organisational

- The Nuitka GitHub Organisation that was created a while ago and owns the Nuitka repo now, has gained members. Check out <https://github.com/orgs/Nuitka/people> for their list. This is an exciting transformation for Nuitka.
- Nuitka is participating in the GSoC 2019 under the PSF umbrella. We hope to grow even further. Thanks to the mentors who volunteered for this important task. Check out the [GSoC 2019 page](#) and thanks to the students that are already helping out.
- Added Nuitka internal [API documentation](#) that will receive more love in the future. It got some for this release, but a lot is missing.
- The Nuitka code has been `black`-ened and is formatted with an automatic tool now all the way, which makes contributors lives easier.
- Added documentation for questions received as part of the GSoC applications and ideas work.
- Some proof reading pull requests were merged for the documentation, thanks to everybody who addresses these kinds of errors. Sometimes typos, sometimes broken links, etc.
- Updated inline copy of Scons used for Python3 to 3.0.4, which hopefully means more bugs are fixed.

Summary

This release is a sign of increasing adoption of Nuitka. The GSoC 2019 is showing early effects, as is more developers joining the effort. These are great times for Nuitka.

This release has not much on the optimization side that is user visible, but the work that has begun is capable of producing glorious benchmarks once it will be finished.

The focus on this and coming releases is definitely to open up the Nuitka development now that people are coming in as permanent or temporary contributors in (relatively) high numbers.

Nuitka Release 0.6.1

This release comes after a relatively long time, and contains important new optimization work, and even more bug fixes.

Bug Fixes

- Fix, the options `--[no]follow-import-to=package_name` was supposed to not follow into the given package, but the check was executed too broadly, so that e.g. `package_name2` was also affected. Fixed in 0.6.0.1 already.
- Fix, wasn't detecting multiple recursions into the same package in module mode, when attempting to compile a whole sub-package. Fixed in 0.6.0.1 already.
- Fix, constant values are used as C boolean values still for some of the cases. Fixed in 0.6.0.1 already.
- Fix, referencing a function cannot raise an exception, but that was not annotated. Fixed in 0.6.0.2 already.
- macOS: Use standard include of C bool type instead of rolling our own, which was not compatible with newest Clang. Fixed in 0.6.0.3 already.
- Python3: Fix, the `bytes` built-in type actually does have a `__float__` slot. Fixed in 0.6.0.4 already.
- Python3.7: Types that are also sequences still need to call the method `__class_getitem__` for consideration. Fixed in 0.6.0.4 already.
- Python3.7: Error exits from program exit could get lost on Windows due to `__spec__` handling not preserving errors. Fixed in 0.6.0.4 already.
- Windows: Negative exit codes from Nuitka, e.g. due to a triggered assertion in debug mode were not working. Fixed in 0.6.0.4 already.
- Fix, conditional `and` expressions were mis-optimized when not used to not execute the right hand side still. Fixed in 0.6.0.4 already.
- Python3.6: Fix, generators, coroutines, and `asyncgen` were not properly supporting annotations for local variables. Fixed in 0.6.0.5 already.
- Python3.7: Fix, class declarations had memory leaks that were untestable before 3.7.1 fixed reference count issues in CPython. Fixed in 0.6.0.6 already.
- Python3.7: Fix, `asyncgen` expressions can be created in normal functions without an immediate awaiting of the iterator. This new feature was not correctly supported.
- Fix, star imports on the module level should disable built-in name optimization except for the most critical ones, otherwise e.g. names like `all` or `pow` can become wrong. Previous workarounds for `pow` were not good enough.
- Fix, the `scons` for Python3 failed to properly report build errors due to a regression of the `Scons` version used for it. This would mask build errors on Windows.

- Python3.4: Fix, packages didn't indicate that they are packages in their `__spec__` value, causing issues with `importlib_resources` module.
- Python3.4: The `__spec__` values of compiled modules didn't have compatible `origin` and `has_location` values preventing `importlib_resources` module from working to load data files.
- Fix, packages created from `.pth` files were also considered when checking for sub-packages of a module.
- Standalone: Handle cases of conflicting DLLs better. On Windows pick the newest file version if different, and otherwise just report and pick randomly because we cannot really decide which ought to be loaded.
- Standalone: Warn about collisions of DLLs on non-Windows only as this can happen with wheels apparently.
- Standalone: For Windows Python extension modules `.pyd` files, remove the SxS configuration for cases where it causes problems, not needed.
- Fix: The `exec` statement on file handles was not using the proper filename when compiling, therefore breaking e.g. `inspect.getsource` on functions defined there.
- Standalone: Added support for OpenGL platform plugins to be included automatically.
- Standalone: Added missing implicit dependency for `zmq` module.
- Python3.7: Fix, using the `-X utf8` flag on the calling interpreter, aka `--python-flag=utf8_mode` was not preserved in the compiled binary in all cases.

Optimization

- Enabled C target type `void` which will catch creating unused stuff more immediately and give better code for expression only statements.
- Enabled in-place optimization for module variables, avoiding write back to the module dict for unchanged values, accelerating these operations.
- Compile time memory savings for the `yield` node of Python2, no need to track if it is in an exception handler, not relevant there.
- Using the single child node for the `yield` nodes gives memory savings at compile time for these, while also making them operate faster.
- More kinds of in-place operations are now optimized, e.g. `int += int` and the `bytes` ones were specialized to perform real in-place extension where possible.
- Loop variables no longer loose type information, but instead collect the set of possible type shapes allowing optimization for them.

Organisational

- Corrected download link for Arch AUR link of develop package.
- Added repository for Ubuntu Cosmic (18.10) for download.
- Added repository for Fedora 29 for download.
- Describe the exact format used for `clang-format` in the Developer Manual.
- Added description how to use CondaCC on Windows to the User Manual.

Cleanups

- The operations used for `async for`, `async with`, and `await` were all doing a look-up of an awaitable, and then executing the `yield from` from that awaitable as one thing. Now this is split into two parts, with a new `ExpressionYieldFromAwaitable` as a dedicated node.
- The `yield` node types, now 3 share a base class and common computation for now, enhancing the one for awaitable, which was not fully annotating everything that can happen.
- In code generation avoid statement blocks that are not needed, because there are no local C variables declared, and properly indent them.

Tests

- Fixups for the manual Valgrind runner and the UI changes.
- Test runner detects lock issue of `clcache` on Windows and considers it a permission problem that causes a retry.

Summary

This addresses even more corner cases not working correctly, the out of the box experience should be even better now.

The push towards C level performance for integer operation was held up by the realization that loop SSA was not yet there really, and that it had to be implemented, which of course now makes a huge difference for the cases where e.g. `bool` are being used. There is no C type for `int` used yet, which limits the impact of optimization to only taking shortcuts for the supported types. These are useful and faster of course, but only building blocks for what is to come.

Most of the effort went into specialized helpers that e.g. add a `float` and an `int` value in a dedicated fashion, as well as comparison operations, so we can fully operate some minimal examples with specialized code. This is too limited still, and must be applied to ever more operations.

What's more is that the benchmarking situation has not improved. Work will be needed in this domain to make improvements more demonstrable. It may well end up being the focus for the next release to improve Nuitka speedcenter to give more fine grained insights across minor changes of Nuitka and graphs with more history.

Nuitka Release 0.6.0

This release adds massive improvements for optimization and a couple of bug fixes.

It also indicates reaching the mile stone of doing actual type inference, even if only very limited.

And with the new version numbers, lots of UI changes go along. The options to control recursion into modules have all been renamed, some now have different defaults, and finally the filenames output have changed.

Bug Fixes

- Python3.5: Fix, the awaiting flag was not removed for exceptions thrown into a coroutine, so next time it appeared to be awaiting instead of finished.
- Python3: Classes in generators that were using built-in functions crashed the compilation with C errors.
- Some regressions for XML outputs from previous changes were fixed.
- Fix, `hasattr` was not raising an exception if used with non-string attributes.
- For really large compilations, MSVC linker could choke on the input file, line length limits, which is now fixed for the inline copy of Scons.

- Standalone: Follow changed hidden dependency of `PyQt5` to `PyQt5.sip` for newer versions
- Standalone: Include certificate file using by `requests` module in some cases as a data file.

Optimization

- Enabled C target type `nuitka_bool` for variables that are stored with boolean shape only, and generate C code for those
- Using C target type `nuitka_bool` many more expressions are now handled better in conditions.
- Enhanced `is` and `is not` to be C source type aware, so they can be much faster for them.
- Use C target type for `bool` built-in giving more efficient code for some source values.
- Annotate the `not` result to have boolean type shape, allowing for more compile time optimization with it.
- Restored previously lost optimization of loop break handling `StopIteration` which makes loops much faster again.
- Restore lost optimization of subscripts with constant integer values making them faster again.
- Optimize in-place operations for cases where left, right, or both sides have known type shapes for some values. Initially only a few variants were added, but there is more to come.
- When adjacent parts of an f-string become known string constants, join them at compile time.
- When there is only one remaining part in an f-string, use that directly as the result.
- Optimize empty f-strings directly into empty strings constant during the tree building phase.
- Added specialized attribute check for use in re-formulations that doesn't expose exceptions.
- Remove locals sync operation in scopes without local variables, e.g. classes or modules, making `exec` and the like slightly leaner there.
- Remove `try` nodes that did only re-raise exceptions.
- The `del` of variables is now driven fully by C types and generates more compatible code.
- Removed useless double exception exits annotated for expressions of conditions and added code that allows conditions to adapt themselves to the target shape `bool` during optimization.

New Features

- Added support for using `.egg` files in `PYTHONPATH`, one of the more rare uses, where Nuitka wasn't yet compatible.
- Output binaries in standalone mode with platform suffix, on non-Windows that means no suffix. In accelerated mode on non-Windows, use `.bin` as a suffix to avoid collision with files that have no suffix.
- Windows: It's now possible to use `clang-cl.exe` for CC with Nuitka as a third compiler on Windows, but it requires an existing MSVC install to be used for resource compilation and linking.
- Windows: Added support for using `ccache.exe` and `clcache.exe`, so that object files can now be cached for re-compilation.
- For debug mode, report missing in-place helpers. These kinds of reports are to become more universal and are aimed at recognizing missed optimization chances in Nuitka. This features is still in its infancy. Subsequent releases will add more like these.

Organisational

- Disabled comments on the web site, we are going to use Twitter instead, once the site is migrated to an updated Nikola.
- The static C code is now formatted with `clang-format` to make it easier for contributors to understand.
- Moved the construct runner to top level binary and use it from there, with future changes coming that should make it generally useful outside of Nuitka.
- Enhanced the issue template to tell people how to get the `develop` version of Nuitka to try it out.
- Added documentation for how use the object caching on Windows to the User Manual.
- Removed the included GUI, originally intended for debugging, but XML outputs are more powerful anyway, and it had been in disrepair for a long time.
- Removed long deprecated options, e.g. `--exe` which has long been the default and is no more accepted.
- Renamed options to include plugin files to `--include-plugin-directory` and `--include-plugin-files` for more clarity.
- Renamed options for recursion control to e.g. `--follow-imports` to better express what they actually do.
- Removed `--python-version` support for switching the version during compilation. This has only worked for very specific circumstances and has been deprecated for a while.
- Removed `--code-gen-no-statement-lines` support for not having line numbers updated at run time. This has long been hidden and probably would never gain all that much, while causing a lot of incompatibility.

Cleanups

- Moved command line arguments to dedicated module, adding checks was becoming too difficult.
- Moved rich comparison helpers to a dedicated C file.
- Dedicated binary and unary node bases for clearer distinction and more efficient memory usage of unary nodes. Unary operations also no longer have in-place operation as an issue.
- Major cleanup of variable accesses, split up into multiple phases and all including module variables being performed through C types, with no special cases anymore.
- Partial cleanups of C type classes with code duplications, there is much more to resolve though.
- Windows: The way `exec` was performed is discouraged in the `subprocess` documentation, so use a variant that cannot block instead.
- Code proving information about built-in names and values was using not very portable constructs, and is now written in a way that PyPy would also like.

Tests

- Avoid using `2to3` for basic operators test, removing test of some Python2 only stuff, that is covered elsewhere.
- Added ability to cache output of CPython when comparing to it. This is to allow CI tests to not execute the same code over and over, just to get the same value to compare with. This is not enabled yet.

Summary

This release marks a point, from which on performance improvements are likely in every coming release. The C target types are a major milestone. More C target types are in the work, e.g. `void` is coming for expressions that are done, but not used, that is scheduled for the next release.

Although there will be a need to also adapt optimization to take full advantage of it, progress should be quick from here. There is a lot of ground to cover, with more C types to come, and all of them needing specialized helpers. But as soon as e.g. `int`, `str` are covered, many more programs are going to benefiting from this.

Nuitka Release 0.5.33

This release contains a bunch of fixes, most of which were previously released as part of hotfixes, and important new optimization for generators.

Bug Fixes

- Fix, nested functions with local classes using outside function closure variables were not registering their usage, which could lead to errors at C compile time. Fixed in 0.5.32.1 already.
- Fix, usage of built-in calls in a class level could crash the compiler if a class variable was updated with its result. Fixed in 0.5.32.1 already.
- Python 3.7: The handling of non-type bases classes was not fully compatible and wrong usages were giving `AttributeError` instead of `TypeError`. Fixed in 0.5.32.2 already.
- Python 3.5: Fix, `await` expressions didn't annotate their exception exit. Fixed in 0.5.32.2 already.
- Python3: The `enum` module usages with `__new__` in derived classes were not working, due to our automatic `staticmethod` decoration. Turns out, that was only needed for Python2 and can be removed, making `enum` work all the way. Fixed in 0.5.32.3 already.
- Fix, recursion into `__main__` was done and could lead to compiler crashes if the main module was named like that. This is not prevented. Fixed in 0.5.32.3 already.
- Python3: The name for list contraction's frames was wrong all along and not just changed for 3.7, so drop that version check on it. Fixed in 0.5.32.3 already.
- Fix, the hashing of code objects has creating a key that could produce more overlaps for the hash than necessary. Using a `C1` on line 29 and `C` on line 129, was considered the same. And that is what actually happened. Fixed in 0.5.32.3 already.
- macOS: Various fixes for newer Xcode versions to work as well. Fixed in 0.5.32.4 already.
- Python3: Fix, the default `__annotations__` was the empty dict and could be modified, leading to severe corruption potentially. Fixed in 0.5.32.4 already.
- Python3: When an exception is thrown into a generator that currently does a `yield from` is not to be normalized.
- Python3: Some exception handling cases of `yield from` were leaking references to objects. Fixed in 0.5.32.5 already.
- Python3: Nested namespace packages were not working unless the directory continued to exist on disk. Fixed in 0.5.32.5 already.
- Standalone: Do not include `icuuc.dll` which is a system DLL. Fixed in 0.5.32.5 already.
- Standalone: Added hidden dependency of newer version of `sip`. Fixed in 0.5.32.5 already.
- Standalone: Do not copy file permissions of DLLs and extension modules as that makes deleting and modifying them only harder. Fixed in 0.5.32.6 already.

- Windows: The multiprocessing plugin was not always properly patching the run time for all module loads, made it more robust. Fixed in 0.5.32.6 already.
- Standalone: Do not preserve permissions of copied DLLs, which can cause issues with read-only files on Windows when later trying to overwrite or remove files.
- Python3.4: Make sure to disconnect finished generators from their frames to avoid potential data corruption. Fixed in 0.5.32.6 already.
- Python3.5: Make sure to disconnect finished coroutines from their frames to avoid potential data corruption. Fixed in 0.5.32.6 already.
- Python3.6: Make sure to disconnect finished asyncgen from their frames to avoid potential data corruption. Fixed in 0.5.32.6 already.
- Python3.5: Explicit frame closes of frames owned by coroutines could corrupt data. Fixed in 0.5.32.7 already.
- Python3.6: Explicit frame closes of frames owned by asyncgen could corrupt data. Fixed in 0.5.32.7 already.
- Python 3.4: Fix threaded imports by properly handling `_initializing` in compiled modules `__spec__` attributes. Before it happen that another thread attempts to use an unfinished module. Fixed in 0.5.32.8 already.
- Fix, the options `--include-module` and `--include-package` were present but not visible in the help output. Fixed in 0.5.32.8 already.
- Windows: The multiprocessing plugin failed to properly pass compiled functions. Fixed in 0.5.32.8 already.
- Python3: Fix, optimization for in-place operations on mapping values are not allowed and had to be disabled. Fixed in 0.5.32.8 already.
- Python 3.5: Fixed exception handling with coroutines and asyncgen `throw` to not corrupt exception objects.
- Python 3.7: Added more checks to class creations that were missing for full compatibility.
- Python3: Smarter hashing of unicode values avoids increased memory usage from cached converted forms in debug mode.

Organisational

- The issue tracker on GitHub is now the one that should be used with Nuitka, winning due to easier issue templating and integration with pull requests.
- Document the threading model and exception model to use for MinGW64.
- Removed the `enum` plug-in which is no longer useful after the improvements to the `staticmethod` handling for Python3.
- Added Python 3.7 testing for Travis.
- Make it clear in the documentation that `pyenv` is not supported.
- The version output includes more information now, OS and architecture, so issue reports should contain that now.
- On PyPI we didn't yet indicated Python 3.7 as supported, which it of course is.

New Features

- Added support for MiniConda Python.

Optimization

- Using goto based generators that return from execution and resume based on heap storage. This makes tests using generators twice as fast and they no longer use a full C stack of 2MB, but only 1K instead.
- Conditional `a if cond else b`, `a and b`, `a or b` expressions of which the result value is are now transformed into conditional statements allowing to apply further optimizations to the right and left side expressions as well.
- Replace unused function creations with side effects from their default values with just those, removing more unused code.
- Put all statement related code and declarations for it in a dedicated C block, making things slightly more easy for the C compiler to re-use the stack space.
- Avoid linking against `libpython` in module mode on everything but Windows where it is really needed. No longer check for static Python, not needed anymore.
- More compact function, generator, and `asyncgen` creation code for the normal cases, avoid `qualname` if identical to `name` for all of them.
- Python2 class dictionaries are now indeed directly optimized, giving more compact code.
- Module exception exits and thus its frames have become optional allowing to avoid some code for some special modules.
- Uncompiled generator integration was backported to 3.4 as well, improving compatibility and speed there as well.

Cleanups

- Frame object and their cache declarations are now handled by the way of allocated variable descriptions, avoid special handling for them.
- The interface to "forget" a temporary variable has been replaced with a new method that skips a number for it. This is done to keep expression use the same indexes for all their child expressions, but this is more explicit.
- Instead of passing around C variables names for temporary values, we now have full descriptions, with C type, code name, storage location, and the init value to use. This makes the information more immediately available where it is needed.
- Variable declarations are now created when needed and stored in dedicated variable storage objects, which then in can generate the code as necessary.
- Module code generation has been enhanced to be closer to the pattern used by functions, generators, etc.
- There is now only one spot that creates variable declaration, instead of previous code duplications.
- Code objects are now attached to functions, generators, coroutines, and `asyncgen` bodies, and not anymore to the creation of these objects. This allows for simpler code generation.
- Removed fiber implementations, no more needed.

Tests

- Finally the `asyncgen` tests can be enabled in the CPython 3.6 test suite as the corrupting crash has been identified.
- Cover ever more cases of spurious permission problems on Windows.

- Added the ability to specify specific modules a comparison test should recurse to, making some CPython tests follow into modules where actual test code lives.

Summary

This release is huge in many ways.

First, finishing "goto generators" clears an old scalability problem of Nuitka that needed to be addressed. No more do generators/coroutines/asyncgen consume too much memory, but instead they become as lightweight as they ought to be.

Second, the use of variable declarations carrying type information all through the code generation, is an important pre-condition for "C types" work to resume and become possible, what will be 0.6.0 and the next release.

Third, the improved generator performance will be removing a lot of cases, where Nuitka wasn't as fast, as its current state not using "C types" yet, should allow. It is now consistently faster than CPython for everything related to generators.

Fourth, the fibers were a burden for the debugging and linking of Nuitka on various platforms, as they provided deprecated interfaces or not. As they are now gone, Nuitka ought to definitely work on any platform where Python works.

From here on, C types work can take it, and produce the results we are waiting for in the next major release cycle that is about to start.

Also the amount of fixes for this release has been incredibly high. Lots of old bugs esp. for coroutines and asyncgen have been fixed, this is not only faster, but way more correct. Mainly due to the easier debugging and interface to the context code, bugs were far easier to avoid and/or find.

Nuitka Release 0.5.32

This release contains substantial new optimization, bug fixes, and already the full support for Python 3.7. Among the fixes, the enhanced coroutine work for compatibility with uncompiled ones is most important.

Bug Fixes

- Fix, was optimizing write backs of attribute in-place assignments falsely.
- Fix, generator stop future was not properly supported. It is now the default for Python 3.7 which showed some of the flaws.
- Python3.5: The `__qualname__` of coroutines and asyncgen was wrong.
- Python3.5: Fix, for dictionary unpackings to calls, check the keys if they are string values, and raise an exception if not.
- Python3.6: Fix, need to check assignment unpacking for too short sequences, we were giving `IndexError` instead of `ValueError` for these. Also the error messages need to consider if they should refer to "at least" in their wording.
- Fix, outline nodes were cloned more than necessary, which would corrupt the code generation if they later got removed, leading to a crash.
- Python3.5: Compiled coroutines awaiting uncompiled coroutines was not working properly for finishing the uncompiled ones. Also the other way around was raising a `RuntimeError` when trying to pass an exception to them when they were already finished. This should resolve issues with `asyncio` module.
- Fix, side effects of a detected exception raise, when they had an exception detected inside of them, lead to an infinite loop in optimization. They are now optimized in-place, avoiding an extra step later on.

New Features

- Support for Python 3.7 with only some corner cases not supported yet.

Optimization

- Delay creation of `StopIteration` exception in generator code for as long as possible. This gives more compact code for generators, which now pass the return values via compiled generator attribute for Python 3.3 or higher.
- Python3: More immediate re-formulation of classes with no bases. Avoids noise during optimization.
- Python2: For class dictionaries that are only assigned from values without side effects, they are not converted to temporary variable usages, allowing the normal SSA based optimization to work on them. This leads to constant values for class dictionaries of simple classes.
- Explicit cleanup of nodes, variables, and local scopes that become unused, has been added, allowing for breaking of cyclic dependencies that prevented memory release.

Tests

- Adapted 3.5 tests to work with 3.7 coroutine changes.
- Added CPython 3.7 test suite.

Cleanups

- Removed remaining code that was there for 3.2 support. All uses of version comparisons with 3.2 have been adapted. For us, Python3 now means 3.3, and we will not work with 3.2 at all. This removed a fair bit of complexity for some things, but not all that much.
- Have dedicated file for import released helpers, so they are easier to find if necessary. Also do not have code for importing a name in the header file anymore, not performance relevant.
- Disable Python warnings when running scones. These are particularly given when using a Python debug binary, which is happening when Nuitka is run with `--python-debug` option and the inline copy of Scons is used.
- Have a factory function for all conditional statement nodes created. This solved a TODO and handles the creation of statement sequences for the branches as necessary.
- Split class reformulation into two files, one for Python2 and one for Python3 variant. They share no code really, and are too confusing in a single file, for the huge code bodies.
- Locals scopes now have a registry, where functions and classes register their locals type, and then it is created from that.
- Have a dedicated helper function for single argument calls in static code that does not require an array of objects as an argument.

Organisational

- There are now `requirements-devel.txt` and `requirements.txt` files aimed at usage with scones and by users, but they are not used in installation.

Summary

This releases has this important step to add conversion of locals dictionary usages to temporary variables. It is not yet done everywhere it is possible, and the resulting temporary variables are not yet propagated in the all the cases, where it clearly is possible. Upcoming releases ought to achieve that most Python2 classes will become to use a direct dictionary creation.

Adding support for Python 3.7 is of course also a huge step. And also this happened fairly quickly and soon after its release. The generic classes it adds were the only real major new feature. It breaking the internals for exception handling was what was holding back initially, but past that, it was really easy.

Expect more optimization to come in the next releases, aiming at both the ability to predict Python3 metaclasses `__prepare__` results, and at more optimization applied to variables after they became temporary variables.

Nuitka Release 0.5.31

This release is massive in terms of fixes, but also adds a lot of refinement to code generation, and more importantly adds experimental support for Python 3.7, while enhancing support for Pyt5 in standalone mode by a lot.

Bug Fixes

- Standalone: Added missing dependencies for `PyQt5.Qt` module.
- Plugins: Added support for `PyQt5.Qt` module and its `qml` plugins.
- Plugins: The sensible plugin list for PyQt now includes that platforms plugins on Windows too, as they are kind of mandatory.
- Python3: Fix, for uninstalled Python versions wheels that linked against the `Python3` library as opposed to `Python3X`, it was not found.
- Standalone: Prefer DLLs used by main program binary over ones used by wheels.
- Standalone: For DLLs added by Nuitka plugins, add the package directory to the search path for dependencies where they might live.
- Fix, the `vars` built-in didn't annotate its exception exit.
- Python3: Fix, the `bytes` and `complex` built-ins needs to be treated as a slot too.
- Fix, consider if `del` variable must be assigned, in which case no exception exit should be created. This prevented Tkinter compilation.
- Python3.6: Added support for the following language construct:

```
d = {"metaclass": M}

class C(**d):
    pass
```

- Python3.5: Added support for cyclic imports. Now a `from` import with a name can really cause an import to happen, not just a module attribute lookup.
- Fix, `hasattr` was never raising exceptions.
- Fix, `bytearray` constant values were considered to be non-iterable.
- Python3.6: Fix, now it is possible to `del __annotations__` in a class and behave compatible. Previously in this case we were falling back to the module variable for annotations used after that which is wrong.

- Fix, some built-in type conversions are allowed to return derived types, but Nuitka assumed the exact type, this affected `bytes`, `int`, `long`, `unicode`.
- Standalone: Fix, the `_socket` module was insisted on to be found, but can be compiled in.

New Features

- Added experimental support for Python 3.7, more work will be needed though for full support. Basic tests are working, but there are at least more coroutine changes to follow.
- Added support for building extension modules against statically linked Python. This aims at supporting manylinux containers, which are supposed to be used for creating widely usable binary wheels for Linux. Programs won't work with statically linked Python though.
- Added options to allow ignoring the Windows cache for DLL dependencies or force an update.
- Allow passing options from distutils to Nuitka compilation via setup options.
- Added option to disable the DLL dependency cache on Windows as it may become wrong after installing new software.
- Added experimental ability to provide extra options for Nuitka to `setuptools`.
- Python3: Remove frame preservation and restoration of exceptions. This is not needed, but leaked over from Python2 code.

Optimization

- Apply value tracing to local dict variables too, enhancing the optimization for class bodies and function with `exec` statements by a lot.
- Better optimization for "must not have value", wasn't considering merge traces of uninitialized values, for which this is also the case.
- Use 10% less memory at compile time due to specialized base classes for statements with a single child only allowing `__slots__` usage by not having multiple inheritance for those.
- More immediately optimize branches with known truth values, so that merges are avoided and do not prevent trace based optimization before the pass after the next one. In some cases, optimization based on traces could fail to be done if there was no next pass caused by other things.
- Much faster handling for functions with a lot of `eval` and `exec` calls.
- Static optimization of `type` with known type shapes, the value is predicted at compile time.
- Optimize containers for all compile time constants into constant nodes. This also enables further compile time checks using them, e.g. with `isinstance` or `in` checks.
- Standalone: Using threads when determining DLL dependencies. This will speed up the un-cached case on Windows by a fair bit.
- Also remove unused assignments for mutable constant values.
- Python3: Also optimize calls to `bytes` built-in, this was so far not done.
- Statically optimize iteration over constant values that are not iterable into errors.
- Removed Fortran, Java, LaTeX, PDF, etc. stuff from the inline copies of Scons for faster startup and leaner code. Also updated to 3.0.1 which is no important difference over 3.0.0 for Nuitka however.
- Make sure to always release temporary objects before checking for error exits. When done the other way around, more C code than necessary will be created, releasing them in both normal case and error case after the check.
- Also remove unused assignments in case the value is a mutable constant.

Cleanups

- Don't store "version" numbers of variable traces for code generation, instead directly use the references to the value traces instead, avoiding later lookups.
- Added dedicated module for `complex` built-in nodes.
- Moved C helpers for integer and complex types to dedicated files, solving the TODOs around them.
- Removed some Python 3.2 only codes.

Organisational

- For better bug reports, the `--version` output now contains also the Python version information and the binary path being used.
- Started using specialized exceptions for some types of errors, which will output the involved data for better debugging without having to reproduce anything. This does e.g. output XML dumps of problematic nodes.
- When encountering a problem (compiler crash) in optimization, output the source code line that is causing the issue.
- Added support for Fedora 28 RPM builds.
- Remove more instances of mentions of 3.2 as supported or usable.
- Renovated the graphing code and made it more useful.

Summary

This release marks important progress, as the locals dictionary tracing is a huge step ahead in terms of correctness and proper optimization. The actual resulting dictionary is not yet optimized, but that ought to follow soon now.

The initial support of 3.7 is important. Right now it apparently works pretty well as a 3.6 replacement already, but definitely a lot more work will be needed to fully catch up.

For standalone, this accumulated a lot of improvements related to the plugin side of Nuitka. Thanks to those involved in making this better. On Windows things ought to be much faster now, due to parallel usage of dependency walker.

Nuitka Release 0.5.30

This release has improvements in all areas. Many bug fixes are accompanied with optimization changes towards value tracing.

Bug Fixes

- Fix, the new `setuptools` runners were not used by `pip` breaking the use of Nuitka from PyPI.
- Fix, imports of `six.moves` could crash the compiler for built-in names. Fixed in 0.5.29.2 already.
- Windows: Make the `nuitka-run` not a symlink as these work really bad on that platform, instead make it a full copy just like we did for `nuitka3-run` already. Fixed in 0.5.29.2 already.
- Python3.5: In module mode, `types.coroutine` was monkey patched into an endless recursion if including more than one module, e.g. for a package. Fixed in 0.5.29.3 already.
- Python3.5: Dictionary unpackings with both star arguments and non star arguments could leak memory. Fixed in 0.5.29.3 already.

```
c = {a: 1, **d}
```

- Fix, distutils usage was not working for Python2 anymore, due to using `super` for what are old style classes on that version.
- Fix, some method calls to C function members could leak references.

```
class C:
    for_call = functools.partial

    def m(self):
        self.for_call() # This leaked a reference to the descriptor.
```

- Python3.5: The bases classes should be treated as an unpacking too.

```
class C(*D): # Allowed syntax that was not supported.
    pass
```

- Windows: Added back batch files to run Nuitka from the command line. Fixed in 0.5.29.5 already.

New Features

- Added option `--include-package` to force inclusion of a whole package with the submodules in a compilation result.
- Added options `--include-module` to force inclusion of a single module in a compilation result.
- The multiprocessing plug-in got adapted to Python 3.4 changes and will now also work in accelerated mode on Windows.
- It is now possible to specify the Qt plugin directories with e.g. `--enable-plugin==qt_plugins=imageformats` and have only those included. This should avoid dependency creep for shared libraries.
- Plugins can now make the decision about recursing to a module or not.
- Plugins now can get their own options passed.

Optimization

- The re-raising of exceptions has gotten its own special node type. This aims at more readability (XML output) and avoiding the overhead of checking potential attributes during optimization.
- Changed built-in `int`, `long`, and `float` to using a slot mechanism that also analyses the type shape and detects and warns about errors at compile time.
- Changed the variable tracing to value tracing. This meant to cleanup all the places that were using it to find the variable.
- Enable must have / must not value value optimization for all kinds of variables including module and closure variables. This often avoids error exits and leads to smaller and faster generated code.

Tests

- Added burn test with local install of pip distribution to virtualenv before making any PyPI upload. It seems pip got its specific error sources too.

- Avoid calling `2to3` and prefer `<python> -m lib2to3` instead, as it seems at least Debian Testing stopped to provide the binary by default. For Python 2.6 and 3.2 we continue to rely on it, as they don't support that mode of operation.
- The PyLint checks have been made more robust and even more Python3 portable.
- Added PyLint to Travis builds, so PRs are automatically checked too.
- Added test for `distutils` usage with Nuitka that should prevent regressions for this new feature and to document how it can be used.
- Make coverage taking work on Windows and provide the full information needed, the rendering stage is not there working yet though.
- Expanded the trick assignment test cases to cover more slots to find bugs introduced with more aggressive optimization of closure variables.
- New test to cover multiprocessing usage.
- Generating more code tests out of doctests for increased coverage of Nuitka.

Cleanups

- Stop using `--python-version` in tests where they still remained.
- Split the forms of `int` and `long` into two different nodes, they share nothing except the name. Create the constants for the zero arg variant more immediately.
- Split the output comparison part into a dedicated testing module so it can be re-used, e.g. when doing `distutils` tests.
- Removed dead code from variable closure taking.
- Have a dedicated module for the metaclass of nodes in the tree, so it is easier to find, and doesn't clutter the node base classes module as much.
- Have a dedicated node for `raise` statements instead of checking for all the arguments to be non-present.

Organisational

- There is now a pull request template for GitHub when used.
- Deprecating the `--python-version` argument which should be replaced by using `-m nuitka` with the correct Python version. Outputs have been updated to recommend this one instead.
- Make automatic import sorting and auto-format tools properly executable on Windows without them changing new lines.
- The documentation was updated to prefer the call method with `-m nuitka` and manually providing the Python binary to use.

Summary

This release continued the `distutils` integration adding first tests, but more features and documentation will be needed.

Also, for the locals dictionary work, the variable tracing was made generic, but not yet put to use. If we use this to also trace dictionary keys, we can expect a lot of improvements for class code again.

The locals dictionary tracing will be the focus before resuming the work on C types, where the ultimate performance boost lies. However, currently, not the full compatibility has been achieved even with currently using dictionaries for classes, and we would like to be able to statically optimize those better anyway.

Nuitka Release 0.5.29

This release comes with a lot of improvements across the board. A lot of focus has been given to the packaging side of Nuitka, but also there is a lot of compatibility work.

Bug Fixes

- Windows: When using Scons for Python3 and Scons for Python2 on the same build directory, a warning would be given about the need to migrate. Make the Scons cache directory use the Python ABI version as a key too, to avoid these issues. Fixed in 0.5.28.1 already.
- Windows: Fixup for Python3 and Scons no more generating the MinGW64 import library for Python anymore properly. Was only working if cached from a previous install of Nuitka. Fixed in 0.5.28.1 already.
- Plugins: Made the data files plugin mandatory and added support for the scrapy package needs.
- Fix, added implicit dependencies for `pkg_resources.external` package. Fixed in 0.5.28.1 already.
- Fix, an import of `x.y` where this was not a package didn't cause the package `x` to be included.
- Standalone: Added support for `six.moves` and `requests.packages` meta imports, these cause hidden implicit imports, that are now properly handled.
- Standalone: Patch the `__file__` value for technical bytecode modules loaded during Python library initialization in a more compatible way.
- Standalone: Extension modules when loaded might actually raise legit errors, e.g. `ImportError` of another module, don't make those into `SystemError` anymore.
- Python3.2: The `__package__` of sub-packages was wrong, which could cause issues when doing relative imports in that sub-package.
- Python3: Contractions in a finally clause could crash the compiler.
- Fix, unused closure variables could lead to a crash in they were passed to a nested function.
- Linux: Standalone dependency analysis could enter an endless recursion in case of cyclic dependencies.
- Python3.6: Async generation expressions need to return a `None` value too.
- Python3.4: Fix, `__spec__` is a package attribute and not a built-in value.

New Features

- It is now possible to run Nuitka with `some_python_you_choose -m nuitka ...` and therefore know exactly which Python installation is going to be used. It does of course need Nuitka installed for this to work. This mechanism is going to replace the `--python-version` mechanism in the future.
- There are dedicated runners for Python3, simply use `nuitka3` or `nuitka3-run` to execute Nuitka if your code is Python3 code.
- Added warning for implicit exception raises due to mismatch in unpacking length. These are statically detected, but so far were not warned about.
- Added cache for `depends.exe` results. This speeds up standalone mode again as some of these calls were really slow.
- The import tracer is more robust against recursion and works with Python3 now.
- Added an option to assume yes for downloading questions. The currently only enables the download of `depends.exe` and is intended for CI servers.

- There is now a report file for scons, which records the values used to run things, this could be useful for debugging.
- Nuitka now registers with distutils and can be used with `bdist_wheel` directly, but this lacks documentation and tests. Many improvements in the distutils build.

Optimization

- Forward propagate compile time constants even if they are only potential usages. This is actually the case where this makes the most sense, as it might remove its use entirely from the branches that do not use it.
- Avoid extra copy of `finally` code. The cloning operation takes time and memory, and this shaved of 0.3% of Nuitka memory usage, as these can also become dangling.
- Class dictionaries are now proper dictionaries in optimization, using some dedicated code for name lookups that are transformed to dedicated locals dictionary or mapping (Python3) accesses. This currently does not fully optimize, but will in coming releases, and saves about 25% of memory compared to the old code.
- Treating module attributes `__package__`, `__loader__`, `__file__`, and `__spec__` with dedicated nodes, that allow or forbid optimization dependent on usage.
- Python3.6: Async generator expressions were not working fully, become more compatible.
- Fix, using `super` inside a contraction could crash the compiler.
- Fix, also accept `__new__` as properly decorated in case it's a `classmethod` too.
- Fix, removed obsolete `--nofreeze-stdlib` which only complicated using the `--recurse-stdlib` which should be used instead.

Organisational

- The `nuitka` Python package is now installed into the public namespace and used from there. There are distinct copies to be installed for both Python2 and Python3 on platforms where it is supported.
- Using `twine` for upload to PyPI now as recommended on their site.
- Running `pylint` on Windows became practical again.
- Added RPM packages for Fedora 26 and 27, these used to fail due to packaging issues.
- Added RPM packages for openSUSE Leap 42.2, 42.3 and 15.0 which were simply missing.
- Added RPM packages for SLE 15.
- Added support for PyLint 1.8 and its new warnings.
- The RPM packages no longer contain `nuitka-run3`, it will be replaced by the new `nuitka3-run` which is in all packages.
- The runners used for installation are now easy install created, but patched to avoid overhead at run time.
- Added repository for Ubuntu Artful (17.10) for download, removed support for Ubuntu Yakkety, Vivid and Zesty (no more supported by them).
- Removed support for Debian Wheezy and Ubuntu Precise (they are too old for modern packaging used).
- There is now a issue template for GitHub when used.

Tests

- Windows: Standalone tests were referencing an old path to `depends.exe` that wasn't populated on new installs.
- Refinements for CPython test suites to become more stable in results. Some tests occasionally fail to clean up, or might do nondeterministic outputs, or are not relevant at all.
- The tests don't use the runners, but more often do `-m nuitka` to become executable without having to find the proper runner. This improves usage during the RPM builds and generally.
- Travis: Do not test development versions of CPython, even for stable release, they break too often.

Summary

This release consolidates a lot of what we already had, adding hopeful stuff for distutils integration. This will need tests and documentation though, but should make Nuitka really easy to use. A few features are still missing to make it generally reliable in that mode, but they are going to come.

Also the locals dictionary work is kind of incomplete without a proper generic tracing of not only local variables, but also dictionary keys. With that work in place, a lot of improvements will happen.

Nuitka Release 0.5.28

This release has a focus on compatibility work and contains bug fixes and work to enhance the usability of Nuitka by integrating with distutils. The major improvement is that contractions no longer use pseudo functions to achieve their own local scope, but that there is now a dedicated structure for that representing an in-lined function.

Bug Fixes

- Python3.6: Fix, `async for` was not yet implemented for async generators.
- Fix, functions with keyword arguments where the value was determined to be a static raise could crash the compiler.
- Detect using MinGW64 32 bits C compiler being used with 64 bits Python with better error message.
- Fix, when extracting side effects of a static raise, extract them more recursively to catch expressions that themselves have no code generation being used. This fixes at least static raises in keyword arguments of a function call.
- Compatibility: Added support for proper operation of `pkgutil.get_data` by implementing `get_data` in our meta path based loader.
- Compatibility: Added `__spec__` module attribute was previously missing, present on Python3.4 and higher.
- Compatibility: Made `__loader__` module attribute set when the module is loading already.
- Standalone: Resolve the `@rpath` and `@loader_path` from `otool` on macOS manually to actual paths, which adds support for libraries compiled with that.
- Fix, nested functions calling `super` could crash the compiler.
- Fix, could not use `--recurse-directory` with arguments that had a trailing slash.
- Fix, using `--recurse-directory` on packages that are not in the search crashed the compiler.
- Compatibility: Python2 `set` and `dict` contractions were using extra frames like Python3 does, but those are not needed.
- Standalone: Fix, the way `PYTHONHOME` was set on Windows had no effect, which allowed the compiled binary to access the original installation still.
- Standalone: Added some newly discovered missing hidden dependencies of extension modules.

- Compatibility: The name mangling of private names (e.g. `__var`) in classes was applied to variable names, and function declarations, but not to classes yet.
- Python3.6: Fix, added support for list contractions with `await` expressions in async generators.
- Python3.6: Fix, `async for` was not working in async generators yet.
- Fix, for module tracebacks, we output the module name `<module name>` instead of merely `<module>`, but if the module was in a package, that was not indicated. Now it is `<module package.name>`.
- Windows: The cache directory could be unicode which then failed to pass as an argument to `scons`. We now encode such names as UTF-8 and decode in `Scons` afterwards, solving the problem in a generic way.
- Standalone: Need to recursively resolve shared libraries with `ldd`, otherwise not all could be included.
- Standalone: Make sure `sys.path` has no references to CPython compile time paths, or else things may work on the compiling machine, but not on another.
- Standalone: Added various missing dependencies.
- Standalone: Wasn't considering the DLLs directory for standard library extensions for freezing, which would leave out these.
- Compatibility: For `__future__` imports the `__import__` function was called more than once.

Optimization

- Contractions are now all properly inlined and allow for optimization as if they were fully local. This should give better code in some cases.
- Classes are now all building their locals dictionary inline to the using scope, allowing for more compact code.
- The dictionary API was not used in module template code, although it helps to generate more compact code.

New Features

- Experimental support for building platform dependent wheel distribution.

```
python setup.py --command-packages=nuitka.distutils clean -a bdist_nuitka
```

Use with caution, this is incomplete work.

- Experimental support for running tests against compiled installation with `nose` and `py.test`.
- When specifying what to recurse to, now patterns can be used, e.g. like this `--recurse-not-to=*.tests` which will skip all tests in submodules from compilation.
- By setting `NUITKA_PACKAGE_package_name=/some/path` the `__path__` of packages can be extended automatically in order to allow and load uncompiled sources from another location. This can be e.g. a `tests` sub-package or other plugins.
- By default when creating a module, now also a `module.pyi` file is created that contains all imported modules. This should be deployed alongside the extension module, so that standalone mode creation can benefit from knowing the dependencies of compiled code.
- Added option `--plugin-list` that was mentioned in the help output, but still missing so far.

- The import tracing of the `hints` module has achieved experimental status and can be used to test compatibility with regards to import behavior.

Cleanups

- Rename tree and codegen `Helper` modules to unique names, making them easier to work with.
- Share the code that decides to not warn for standard library paths with more warnings.
- Use the `bool` enum definition of Python2 which is more elegant than ours.
- Move quality tools, auto-format, isort, etc. to the `nuitka.tools.quality` namespace.
- Move output comparison tool to the `nuitka.tools.testing` namespace.
- Made frame code generation capable of using nested frames, allowing the real inline of classes and contraction bodies, instead of "direct" calls to pseudo functions being used.
- Proper base classes for functions that are entry points, and functions that are merely a local expression using return statements.

Tests

- The search mode with pattern, was not working anymore.
- Resume hash values now consider the Python version too.
- Added test that covers using test runners like `nose` and `py.test` with Nuitka compiled extension modules.

Organisational

- Added support for Scons 3.0 and running Scons with Python3.5 or higher. The option to specify the Python to use for scons has been renamed to reflect that it may also be a Python3 now. Only for Python3.2 to Python3.4 we now need another Python installation.
- Made recursion the default for `--recurse-directory` with packages. Before you also had to tell it to recurse into that package or else it would only include the top level package, but nothing below.
- Updated the man pages, correct mentions of its C++ to C and don't use now deprecated options.
- Updated the help output which still said that standalone mode implies recursion into standard library, which is no longer true and even not recommended.
- Added option to disable the output of `.pyi` file when creating an extension module.
- Removed Ubuntu Wily package download, no longer supported by Ubuntu.

Summary

This release was done to get the fixes and new features out for testing. There is work started that should make generators use an explicit extra stack via pointer, and restore instruction state via goto dispatchers at function entry, but that is not complete.

This feature, dubbed "goto generators" will remove the need for fibers (which is itself a lot of code), reduce the memory footprint at run time for anything that uses a lot of generators, or coroutines.

Integrating with `distutils` is also a new thing, and once completed will make use of Nuitka for existing projects automatic and trivial to do. There is a lot missing for that goal, but we will get there.

Also, documenting how to run tests against compiled code, if that test code lives inside of that package, will make a huge difference, as that will make it easier for people to torture Nuitka with their own test cases.

And then of course, nested frames now mean that every function could be inlined, which was previously not possible due to collisions of frames. This will pave the route for better optimization in those cases in future releases.

The experimental features will require more work, but should make it easier to use Nuitka for existing projects. Future releases will make integrating Nuitka dead simple, or that is the hope.

And last but not least, now that Scons works with Python3, chances are that Nuitka will more often work out of the box. The older Python3 versions that still retain the issue are not very widespread.

Nuitka Release 0.5.27

This release comes a lot of bug fixes and improvements.

Bug Fixes

- Fix, need to add recursed modules immediately to the working set, or else they might first be processed in second pass, where global names that are locally assigned, are optimized to the built-in names although that should not happen. Fixed in 0.5.26.1 already.
- Fix, the accelerated call of methods could crash for some special types. This had been a regress of 0.5.25, but only happens with custom extension types. Fixed in 0.5.26.1 already.
- Python3.5: For `async def` functions parameter variables could fail to properly work with in-place assignments to them. Fixed in 0.5.26.4 already.
- Compatibility: Decorators that overload type checks didn't pass the checks for compiled types. Now `isinstance` and as a result `inspect` module work fine for them.
- Compatibility: Fix, imports from `__init__` were crashing the compiler. You are not supposed to do them, because they duplicate the package code, but they work.
- Compatibility: Fix, the `super` built-in on module level was crashing the compiler.
- Standalone: For Linux, BSD and macOS extension modules and shared libraries using their own `$ORIGIN` to find loaded DLLs resulted in those not being included in the distribution.
- Standalone: Added more missing implicit dependencies.
- Standalone: Fix, implicit imports now also can be optional, as e.g. `_tkinter` if not installed. Only include those if available.
- The `--recompile-c-only` was only working with C compiler as a backend, but not in the C++ compatibility fallback, where files get renamed. This prevented that edit and test debug approach with at least MSVC.
- Plugins: The PyLint plug-in didn't consider the symbolic name `import-error` but only the code `F0401`.
- Implicit exception raises in conditional expressions would crash the compiler.

New Features

- Added support for Visual Studio 2017.
- Added option `--python2-for-scons` to specify the Python2 execute to use for calling Scons. This should allow using Anaconda Python for that task.

Optimization

- References to known unassigned variables are now statically optimized to exception raises and warned about if the according option is enabled.

- Non-hashable keys in dictionaries are now statically optimized to exception raises and warned about if the according option is enabled.
- Enable forward propagation for classes too, resulting in some classes to create only static dictionaries. Currently this never happens for Python3, but it will, once we can statically optimize `__prepare__` too.
- Enable inlining of class dictionary creations if they are mere return statements of the created dictionary. Currently this never happens for Python3, see above for why.
- Python2: Selecting the metaclass is now visible in the tree and can be statically optimized.
- For executables, we now also use a free list for traceback objects, which also makes exception cases slightly faster.
- Generator expressions no longer require the use of a function call with a `.0` argument value to carry the iterator value, instead their creation is directly inlined.
- Remove "pass through" frames for Python2 list contractions, they are no longer needed. Minimal gain for generated code, but more lightweight at compile time.
- When compiling Windows x64 with MinGW64 a link library needs to be created for linking against the Python DLL. This one is now cached and re-used if already done.
- Use common code for `NameError` and `UnboundLocalError` exception code raises. In some cases it was creating the full string at compile time, in others at run time. Since the later is more efficient in terms of code size, we now use that everywhere, saving a bit of binary size.
- Make sure to release unused functions from a module. This saves memory and can be decided after a full pass.
- Avoid using `OrderedDict` in a couple of places, where they are not needed, but can be replaced with a later sorting, e.g. temporary variables by name, to achieve deterministic output. This saves memory at compile time.
- Add specialized return nodes for the most frequent constant values, which are `None`, `True`, and `False`. Also a general one, for constant value return, which avoids the constant references. This saves quite a bit of memory and makes traversal of the tree a lot faster, due to not having any child nodes for the new forms of return statements.
- Previously the empty dictionary constant reference was specialized to save memory. Now we also specialize empty set, list, and tuple constants to the same end. Also the hack to make `is` not say that `{ } is { }` was made more general, mutable constant references and now known to never alias.
- The source references can be marked internal, which means that they should never be visible to the user, but that was tracked as a flag to each of the many source references attached to each node in the tree. Making a special class for internal references avoids storing this in the object, but instead it's now a class property.
- The nodes for named variable reference, assignment, and deletion got split into separate nodes, one to be used before the actual variable can be determined during tree building, and one for use later on. This makes their API clearer and saves a tiny bit of memory at compile time.
- Also eliminated target variable references, which were pseudo children of assignments and deletion nodes for variable names, that didn't really do much, but consume processing time and memory.
- Added optimization for calls to `staticmethod` and `classmethod` built-in methods along with type shapes.
- Added optimization for `open` built-in on Python3, also adding the type shape `file` for the result.
- Added optimization for `bytearray` built-in and constant values. These mutable constants can now be compile time computed as well.
- Added optimization for `frozenset` built-in and constant values. These mutable constants can now be compile time computed as well.

- Added optimization for `divmod` built-in.
- Treat all built-in constant types, e.g. `type` itself as a constant. So far we did this only for constant values types, but of course this applies to all types, giving slightly more compact code for their uses.
- Detect static raises if iterating over non-iterables and warn about them if the option is enabled.
- Split of `locals` node into different types, one which needs the updated value, and one which just makes a copy. Properly track if a functions needs an updated locals dict, and if it doesn't, don't use that. This gives more efficient code for Python2 classes, and `exec` using functions in Python2.
- Build all constant values without use of the `pickle` module which has a lot more overhead than `marshal`, instead use that for too large `long` values, non-UTF-8 `unicode` values, `nan` float, etc.
- Detect the linker arch for all Linux platforms using `objdump` instead of only a hand few hard coded ones.

Cleanups

- The use of `INCREASE_REFCOUNT` got fully eliminated.
- Use functions not vulnerable for buffer overflow. This is generally good and avoids warnings given on OpenBSD during linking.
- Variable closure for classes is different from all functions, don't handle the difference in the base class, but for class nodes only.
- Make sure `maybeNone` doesn't return `None` which means normally "unclear", but `False` instead, since it's always clear for those cases.
- Comparison nodes were using the general comparison node as a base class, but now a proper base class was added instead, allowing for cleaner code.
- Valgrind test runners got changed to using proper tool namespace for their code and share it.
- Made construct case generation code common testing code for re-use in the speedcenter web site. The code also has minor beauty bugs which will then become fixable.
- Use `appdirs` package to determine place to store the downloaded copy of `depends.exe`.
- The code still mentioned C++ in a lot of places, in comments or identifiers, which might be confusing readers of the code.
- Code objects now carry all information necessary for their creation, and no longer need to access their parent to determine flag values. That parent is subject to change in the future.
- Our import sorting wrapper automatically detects imports that could be local and makes them so, removing a few existing ones and preventing further ones on the future.
- Cleanups and annotations to become Python3 PyLint clean as well. This found e.g. that source code references only had `__cmp__` and need rich comparison to be fully portable.

Tests

- The test runner for construct tests got cleaned up and the constructs now avoid using `xrange` so as to not need conversion for Python3 execution as much.
- The main test runner got cleaned up and uses common code making it more versatile and robust.
- Do not run test in debugger if CPython also segfaulted executing the test, then it's not a Nuitka issue, so we can ignore that.
- Improve the way the Python to test with is found in the main test runner, prefer the running interpreter, then `PATH` and registry on Windows, this will find the interesting version more often.
- Added support for "Landscape.io" to ignore the inline copies of code, they are not under our control.

- The test runner for Valgrind got merged with the usage for constructs and uses common code now.
- Construct generation is now common code, intended for sharing it with the Speedcenter web site generation.
- Rebased Python 3.6 test suite to 3.6.1 as that is the Python generally used now.

Organisational

- Added inline copy of `appdirs` package from PyPI.
- Added credits for RedBaron and isort.
- The `--experimental` flag is now creating a list of indications and more than one can be used that way.
- The PyLint runner can also work with Python3 pylint.
- The Nuitka Speedcenter got more fine tuning and produces more tags to more easily identify trends in results. This needs to become more visible though.
- The MSI files are also built on AppVeyor, where their building will not depend on me booting Windows. Getting these artifacts as downloads will be the next step.

Summary

This release improves many areas. The variable closure taking is now fully transparent due to different node types, the memory usage dropped again, a few obvious missing static optimizations were added, and many built-ins were completed.

This release again improves the scalability of Nuitka, which again uses less memory than before, although not an as big jump as before.

This does not extend or use special C code generation for `bool` or any type yet, which still needs design decisions to proceed and will come in a later release.

Nuitka Release 0.5.26

This release comes after a long time and contains large amounts of changes in all areas. The driving goal was to prepare generating C specific code, which is still not the case, but this is very likely going to change soon. However this release improves all aspects.

Bug Fixes

- Compatibility: Fix, for star imports didn't check the values from the `__all__` iterable, if they were string values which could cause problems at run time.

```
# Module level
__all__ = (1,)

# ...
# other module:
from module import *
```

- Fix, for star imports, also didn't check for values from `__all__` if they actually exist in the original values.
- Corner cases of imports should work a lot more precise, as the level of compatibility for calls to `__import__` went from absurd to insane.

- Windows: Fixed detection of uninstalled Python versions (not for all users and DLL is not in system directory). This of course only affected the accelerated mode, not standalone mode.
- Windows: Scan directories for `.pyd` files for used DLLs as well. This should make the PyQt5 wheel work.
- Python3.5: Fix, coroutines could have different code objects for the object and the frame using by it.
- Fix, slices with built-in names crashed the compiler.

```
something[id:len:range]
```

- Fix, the C11 via C++ compatibility uses symlinks to C++ filenames where possible instead of making a copy from the C source. However, even on Linux that may not be allowed, e.g. on a DOS file system. Added fallback to using full copy in that case.
- Python3.5: Fix coroutines to close the "yield from" where an exception is thrown into them.
- Python3: Fix, list contractions should have their own frame too.
- Linux: Copy the "rpath" of compiling Python binary to the created binary. This will make compiled binaries using uninstalled Python versions transparently find the Python shared library.
- Standalone: Add the "rpath" of the compiling Python binary to the search path when checking for DLL dependencies on Linux. This fixes standalone support for Travis and Anaconda on Linux.
- Scons: When calling scons, also try to locate a Python2 binary to overcome a potential Python3 virtualenv in which Nuitka is running.
- Standalone: Ignore more Windows only encodings on non-Windows.

New Features

- Support for Python 3.6 with only few corner cases not supported yet.
- Added options `--python-arch` to pick 32 or 64 bits Python target of the `--python-version` argument.
- Added support for more kinds of virtualenv configurations.
- Uninstalled Python versions such as Anaconda will work fine in accelerated mode, except on Windows.

Optimization

- The node tree children are no longer stored in a separate dictionary, but in the instance dictionary as attributes, making the tree more lightweight and in principle faster to access. This also saved about 6% of the memory usage.
- The memory usage of Nuitka for the Python part has fallen by roughly 40% due to the use of new style classes, and slots where that is possible (some classes use multiple inheritance, where they don't work), and generally by reducing useless members e.g. in source code references. This of course also will make things compiled faster (the C compilation of course is not affected by this.)
- The code generation for frames was creating the dictionary for the raised exception by making a dictionary and then adding all variables, each tested to be set. This was a lot of code for each frame specific, and has been replaced by a generic "attach" mechanism which merely stores the values, and only takes a reference. When asked for frame locals, it only then builds the dictionary. So this is now only done, when that is absolutely necessary, which it normally never is. This of course makes the C code much less verbose, and actual handling of exceptions much more efficient.

- For imports, we now detect for built-in modules, that their import cannot fail, and if name lookups can fail. This leads to less code generated for error handling of these. The following code now e.g. fully detects that no `ImportError` or `AttributeError` will occur.

```
try:
    from __builtin__ import len
except ImportError:
    from builtins import len
```

- Added more type shapes for built-in type calls. These will improve type tracing.
- Compiled frames now have a free list mechanism that should speed up frames that recurse and frames that exit with exceptions. In case of an exception, the frame ownership is immediately transferred to the exception making it easier to deal with.
- The free list implementations have been merged into a new common one that can be used via macro expansion. It is now type agnostic and be slightly more efficient too.
- Also optimize "true" division and "floor division", not only the default division of Python2.
- Removed the need for statement context during code generation making it less memory intensive and faster.

Cleanups

- Now always uses the `__import__` built-in node for all kinds of imports and directly optimizes and recursion into other modules based on that kind of node, instead of a static variant. This removes duplication and some incompatibility regarding defaults usage when doing the actual imports at run time.
- Split the expression node bases and mixin classes to a dedicated module, moving methods that only belong to expressions outside of the node base, making for a cleaner class hierarchy.
- Cleaned up the class structure of nodes, added base classes for typical compositions, e.g. expression with and without children, computation based on built-in, etc. while also checking proper ordering of base classes in the metaclass.
- Moved directory and file operations to dedicated module, making also sure it is more generally used. This makes it easier to make more error resilient deletions of directories on e.g. Windows, where locks tend to live for short times beyond program ends, requiring second attempts.
- Code generation for existing supported types, `PyObject *`, `PyObject **`, and `struct Nuitka_Cellobject *` is now done via a C type class hierarchy instead of `elif` sequences.
- Closure taking is now always done immediately correctly and references are take for closure variables still needed, making sure the tree is correct and needs no finalization.
- When doing variable traces, initialize more traces immediately so it can be more reliable.
- Code to setup a function for local variables and clean it up has been made common code instead of many similar copies.
- The code was treating the `f_executing` frame member as if it were a counter with increases and decreases. Turn it into a mere boolean value and hide its usage behind helper functions.
- The "maybe local variables" are no more. They were replaced by a new locals dict access node with a fallback to a module or closure variable should the dictionary not contain the name. This avoids many ugly checks to not do certain things for that kind of variable.
- We now detect "exec" and "unqualified exec" as well as "star import" ahead of time as flags of the function to be created. We no longer need to mark functions as we go.
- Handle "true", "floor" and normal division properly by applying future flags to decide which one to use.

- We now use symbolic identifiers in all PyLint annotations.
- The release scripts started to move into `nuitka.tools.release` so they get PyLint checks, auto-format and proper code re-use.
- The use of `INCREASE_REFCOUNT_X` was removed, it got replaced with proper `Py_XINCRREF` usages.
- The use of `INCREASE_REFCOUNT` got reduced further, e.g. no generated code uses it anymore, and only a few compiled types do. The function was once required before "C-ish" lifted the need to do everything in one single function call.

Tests

- More robust deletion of directories, temporary stages used by CPython test suites, and standalone directories during test execution.
- Moved tests common code into `nuitka.tools.testing` namespace and use it from there. The code now is allowed to use `nuitka.utils` and therefore often better implementations.
- Made standalone binaries robust against GTK theme access, checking the Python binary (some `site.py` files do that),

Organisational

- Added repository for Ubuntu Zesty (17.04) for download.
- Added support for testing with Travis to complement the internal Buildbot based infrastructure and have pull requests on GitHub automatically tested before merge.
- The `factory` branch is now also on GitHub.
- Removed MSI for Python3.4 32 bits. It seems impossible to co-install this one with the 64 bits variant. All other versions are provided for both bit sizes still.

Summary

This release marks huge progress. The node tree is now absolutely clean, the variable closure taking is fully represented, and code generation is prepared to add another type, e.g. for `bool` for which work has already started.

On a practical level, the scalability of the release will have increased very much, as this uses so much less memory, generates simpler C code, while at the same time getting faster for the exception cases.

Coming releases will expand on the work of this release.

Frame objects should be allowed to be nested inside a function for better re-formulations of classes and contractions of all kinds, as well as real inline of functions, even if they could raise.

The memory savings could be even larger, if we stopped doing multiple inheritance for more node types. The `__slots__` were and the child API change could potentially make things not only more compact, but faster to use too.

And also once special C code generation for `bool` is done, it will set the stage for more types to follow (`int`, `float`, etc). Only this will finally start to give the C type speed we are looking for.

Until then, this release marks a huge cleanup and progress to what we already had, as well as preparing the big jump in speed.

Nuitka Release 0.5.25

This release contains a huge amount of bug fixes, lots of optimization gains, and many new features. It also presents many organisational improvements, and many cleanups.

Bug Fixes

- Python3.5: Coroutine methods using `super` were crashing the compiler. Fixed in 0.5.24.2 already.
- Python3.3: Generator return values were not properly transmitted in case of `tuple` or `StopIteration` values.
- Python3.5: Better interoperability between compiled coroutines and uncompiled generator coroutines.
- Python3.5: Added support to compile in Python debug mode under Windows too.
- Generators with arguments were using two code objects, one with, and one without the `CO_NOFREE` flag, one for the generator object creating function, and one for the generator object.
- Python3.5: The duplicate code objects for generators with arguments lead to interoperability issues with between such compiled generator coroutines and compiled coroutines. Fixed in 0.5.24.2 already.
- Standalone: On some Linux variants, e.g. Debian Stretch and Gentoo, the linker needs more flags to really compile to a binary with `RPATH`.
- Compatibility: For set literal values, insertion order is wrong on some versions of Python, we now detect the bug and emulate it if necessary, previous Nuitka was always correct, but incompatible.

```
{1, 1.0}.pop() # the only element of the set should be 1
```

- Windows: Make the batch files detect where they live at run time, instead of during `setup.py`, making it possible to use them for all cases.
- Standalone: Added package paths to DLL scan for `depends.exe`, as with wheels there now sometimes live important DLLs too.
- Fix, the clang mode was regressed and didn't work anymore, breaking the macOS support entirely.
- Compatibility: For imports, we were passing for `locals` argument a real dictionary with actual values. That is not what CPython does, so stopped doing it.
- Fix, for raised exceptions not passing the validity tests, they could be used after free, causing crashes.
- Fix, the environment `CC` wasn't working unless also specifying `CXX`.
- Windows: The value of `__file__` in module mode was wrong, and didn't point to the compiled module.
- Windows: Better support for `--python-debug` for installations that have both variants, it is now possible to switch to the right variant.

New Features

- Added parsing for shebang to Nuitka. When compiling an executable, now Nuitka will check if the `#!` portion indicates a different Python version and ask the user to clarify with `--python-version` in case of a mismatch.
- Added support for Python flag `--python-flag=-O`, which allows to disable assertions and remove doc strings.

Optimization

- Faster method calls, combining attribute lookup and method call into one, where order of evaluation with arguments doesn't matter. This gives really huge relative speedups for method calls with no arguments.

- Faster attribute lookup in general for `object` descendants, which is all new style classes, and all built-in types.
- Added dedicated `xrange` built-in implementation for Python2 and `range` for Python3. This makes those faster while also solving ordering problems when creating constants of these types.
- Faster `sum` again, using quick iteration interface and specialized quick iteration code for typical standard type containers, `tuple` and `list`.
- Compiled generators were making sure `StopIteration` was set after their iteration, although most users were only going to clear it. Now only the `send` method, which really needs that does it. This speed up the closing of generators quite a bit.
- Compiled generators were preparing a `throw` into non-started compilers, to be checked for immediately after their start. This is now handled in a generic way for all generators, saving code and execution time in the normal case.
- Compiled generators were applying checks only useful for manual `send` calls even during iteration, slowing them down.
- Compiled generators could duplicate code objects due to handling a flag for closure variables differently.
- For compiled frames, the `f_trace` is not writable, but was taking and releasing references to what must be `None`, which is not useful.
- Not passing `locals` to import calls make it less code and faster too.

Organisational

- This release also prepares Python 3.6 support, it includes full language support on the level of CPython 3.6.0 with the sole exception of the new generator coroutines.
- The improved mode is now the default, and full compatibility is now the option, used by test suites. For syntax errors, improved mode is always used, and for test suites, now only the error message is compared, but not call stack or caret positioning anymore.
- Removed long deprecated option "--no-optimization". Code generation too frequently depends on not seeing unoptimized code. This has been hidden and broken long enough to finally remove it.
- Added support for Python3.5 numbers to Speedcenter. There are now also tags for speedcenter, indicating how well "develop" branch fares in comparison to the stable branch.
- With a new tool, source code and Developer Manual contents can be kept in sync, so that descriptions can be quoted there. Eventually a full Sphinx documentation might become available, but for now this makes it workable.
- Added repository for Ubuntu Yakkety (16.10) for download.
- Added repository for Fedora 25 for download.

Cleanups

- Moved the tools to compare CPython output, to sort import statements (`isort`) to auto-format the source code (Redbaron usage), and to check with PyLint to a common new `nuitka.tools` package, runnable with `__main__` modules and dedicated runners in `bin` directory.
- The tools now share code to find source files, or have it for the first time, and other things, e.g. finding needed binaries on Windows installations.
- No longer patch traceback objects `dealloc` function. Should not be needed anymore, and most probably was only bug hiding.
- Moved handling of ast nodes related to import handling to the proper reformulation module.

- Moved statement generation code to helpers module, making it accessible without cyclic dependencies that require local imports.
- Removed deprecated method for getting constant code objects in favor of the new way of doing it. Both methods were still used, making it harder to analyse.
- Removed useless temporary variable initializations from complex call helper internal functions. They worked around code generation issues that have long been solved.
- The ABI flags are no longer passed to Scons together with the version.

Tests

- Windows: Added support to detect and to switch debug Python where available to also be able to execute reference counting tests.
- Added the CPython 3.3 test suite, after cleaning up the worst bits of it, and added the brandnew 3.6 test suite with a minimal set of changes.
- Use the original 3.4 test suite instead of the one that comes from Debian as it has patched quite a few issues that never made it upstream, and might cause crashes.
- More construct tests, making a difference between old style classes, which have instances and new style classes, with their objects.
- It is now possible to run a test program with Python3 and Valgrind.

Summary

The quick iteration is a precursor to generally faster iteration over unknown object iterables. Expanding this to general code generation, and not just the `sum` built-in, might yield significant gains for normal code in the future, once we do code generation based on type inference.

The faster method calls complete work that was already prepared in this domain and also will be expanded to more types than compiled functions. More work will be needed to round this up.

Adding support for 3.6.0 in the early stages of its release, made sure we pretty much have support for it ready right after release. This is always a huge amount of work, and it's good to catch up.

This release is again a significant improvement in performance, and is very important to clean up open ends. Now the focus of coming releases will now be on both structural optimization, e.g. taking advantage of the iterator tracing, and specialized code generation, e.g. for those iterations really necessary to use quick iteration code.

Nuitka Release 0.5.24

This release is again focusing on optimization, this time very heavily on the generator performance, which was found to be much slower than CPython for some cases. Also there is the usual compatibility work and improvements for Pure C support.

Bug Fixes

- Windows: The 3.5.2 coroutine new protocol implementation was using the wrapper from CPython, but it's not part of the ABI on Windows. Have our own instead. Fixed in 0.5.23.1 already.
- Windows: Fixed second compilation with MSVC failing. The files renamed to be C++ files already existed, crashing the compilation. Fixed in 0.5.23.1 already.
- Mac OS: Fixed creating extension modules with `.so` suffix. This is now properly determined by looking at the importer details, leading to correct suffix on all platforms. Fixed in 0.5.23.1 already.

- Debian: Don't depend on a C++ compiler primarily anymore, the C compiler from GNU or clang will do too. Fixed in 0.5.23.1 already.
- Pure C: Adapted scons compiler detecting to properly consider C11 compilers from the environment, and more gracefully report things.

Optimization

- Python2: Generators were saving and restoring exceptions, updating the variables `sys.exc_type` for every context switch, making it really slow, as these are 3 dictionary updates, normally not needed. Now it's only doing it if it means a change.
- Sped up creating generators, coroutines and coroutines by attaching the closure variable storage directly to the object, using one variable size allocation, instead of two, once of which was a standard `malloc`. This makes creating them easier and avoids maintaining the closure pointer entirely.
- Using dedicated compiled cell implementation similar to `PyCellObject` but fully under our control. This allowed for smaller code generated, while still giving a slight performance improvement.
- Added free list implementation to cache generator, coroutines, and function objects, avoiding the need to create and delete this kind of objects in a loop.
- Added support for the built-in `sum`, making slight optimizations to be much faster when iterating over lists and tuples, as well as fast `long` sum for Python2, and much faster `bool` sums too. This is using a prototype version of a "qiter" concept.
- Provide type shape for `xrange` calls that are not constant too, allowing for better optimization related to those.

Tests

- Added workarounds for locks being held by Virus Scanners on Windows to our test runner.
- Enhanced constructs that test generator expressions to more clearly show the actual construct cost.
- Added construct tests for the `sum` built-in on various types of `int` containers, making sure we can do all of those really fast.

Summary

This release improves very heavily on generators in Nuitka. The memory allocator is used more cleverly, and free lists all around save a lot of interactions with it. More work lies ahead in this field, as these are not yet as fast as they should be. However, at least Nuitka should be faster than CPython for these kind of usages now.

Also, proper pure C in the Scons is relatively important to cover more of the rarer use cases, where the C compiler is too old.

The most important part is actually how `sum` optimization is staging a new kind of approach for code generation. This could become the standard code for iterators in loops eventually, making `for` loops even faster. This will be for future releases to expand.

Nuitka Release 0.5.23

This release is focusing on optimization, the most significant part for the users being enhanced scalability due to memory usage, but also break through structural improvements for static analysis of iterators and the debut of type shapes and value shapes, giving way to "shape tracing".

Bug Fixes

- Fix support Python 3.5.2 coroutine changes. The checks got added for improved mode for older 3.5.x, the new protocol is only supported when run with that version or higher.
- Fix, was falsely optimizing away unused iterations for non-iterable compile time constants.

```
iter(1) # needs to raise.
```

- Python3: Fix, `eval` must not attempt to `strip` memoryviews. The was preventing it from working with that type.
- Fix, calling `type` without any arguments was crashing the compiler. Also the exception raised for anything but 1 or 3 arguments was claiming that only 3 arguments were allowed, which is not the compatible thing.
- Python3.5: Fix, follow enhanced error checking for complex call handling of star arguments.
- Compatibility: The `from x import x, y` re-formulation was doing two `__import__` calls instead of re-using the module value.

Optimization

- Uses only about 66% of the memory compared to last release, which is very important step for scalability independent of re-loading. This was achieved by making sure to break loop traces and their reference cycle when they become unused.
- Properly detect the `len` of multiplications at compile time from newly introduces value shapes, so that this is e.g. statically optimized.

```
print(len(" *" * 10000000000))
```

- Due to newly introduced type shapes, `len` and `iter` now properly detect more often if values will raise or not, and warn about detected raises.

```
iter(len(something)) # Will always raise
```

- Due to newly introduced "iterator tracing", we can now properly detect if the length of an unpacking matches its source or not. This allows to remove the check of the generic re-formulations of unpackings at compile time.

```
a, b = b, a # Will never raise due to unpacking
a, b = b, a, c # Will always raise, 3 items cannot unpack to 2
```

- Added support for optimization of the `xrange` built-in for Python2.
- Python2: Added support for `xrange` iterable constant values, pre-building those constants ahead of time.
- Python3: Added support and `range` iterable constant values, pre-building those constants ahead of time. This brings optimization support for Python3 ranges to what was available for Python2 already.
- Avoid having a special node variange for `range` with no arguments, but create the exception raising node directly.
- Specialized constant value nodes are using less generic implementations to query e.g. their length or iteration capabilities, which should speed up many checks on them.
- Added support for the `format` built-in.
- Python3: Added support for the `ascii` built-in.

Organisational

- The movement to pure C got the final big push. All C++ only idioms of C++ were removed, and everything works with C11 compilers. A C++03 compiler can be used as a fallback, in case of MSVC or too old gcc for instance.
- Using pure C, MinGW64 6x is now working properly. The latest version had problems with `hypot` related changes in the C++ standard library. Using C11 solves that.
- This release also prepares Python 3.6 support, it includes full language support on the level of CPython 3.6.0b1.
- The CPython 3.6 test suite was run with Python 3.5 to ensure bug level compatibility, and had a few findings of incompatibilities.

Cleanups

- The last holdouts of classes in Nuitka were removed, and many idioms of C++ were stopped using.
- Moved range related helper functions to a dedicated include file.
- Using `str is not bytes` to detect Python3 `str` handling or actual `bytes` type existence.
- Trace collections were using a mix-in that was merged with the base class that every user of it was having.

Tests

- Added more static optimization tests, a lot more has become feasible to decide at run time, and is now done. These are to detect regressions in that domain.
- The CPython 3.6 test suite is now also run with CPython 3.5 which found some incompatibilities.

Summary

This release marks a huge step forward. We are having the structure for type inference now. This will expand in coming releases to cover more cases, and there are many low hanging fruits for optimization. Specialized codes for variable versions of certain known shapes seems feasible now.

Then there is also the move towards pure C. This will make the backend compilation lighter, but due to using C11, we will not suffer any loss of convenience compared to "C-ish". The plan is to use continue to use C++ for compilation for compilers not capable of supporting C11.

The amount of static analysis done in Nuitka is now going to quickly expand, with more and more constructs predicted to raise errors or simplified. This will be an ongoing activity, as many types of expressions need to be enhanced, and only one missing will not let it optimize as well.

Also, it seems about time to add dedicated code for specific types to be as fast as C code. This opens up vast possibilities for acceleration and will lead us to zero overhead C bindings eventually. But initially the drive is towards enhanced `import` analysis, to become able to know the precise module expected to be imported, and derive type information from this.

The coming work will attack to start whole program optimization, as well as enhanced local value shape analysis, as well specialized type code generation, which will make Nuitka improve speed.

Nuitka Release 0.5.22

This release is mostly an intermediate release on the way to the large goal of having per module compilation that is cacheable and requires far less memory for large programs. This is currently in progress, but required many changes that are in this release, more will be needed.

It also contains a bunch of bug fixes and enhancements that are worth to be released, and the next changes are going to be more invasive.

Bug Fixes

- Compatibility: Classes with decorated `__new__` functions could miss out on the `staticmethod` decorator that is implicit. It's now applied always, unless of course it's already done manually. This corrects an issue found with Pandas. Fixed in 0.5.22.1 already.
- Standalone: For at least Python 3.4 or higher, it could happen that the locale needed was not importable. Fixed in 0.5.22.1 already.
- Compatibility: Do not falsely assume that `not` expressions cannot raise on boolean expressions, since those arguments might raise during creation. This could lead to wrong optimization. Fixed in 0.5.22.2 already.
- Standalone: Do not include system specific C libraries in the distribution created. This would lead to problems for some configurations on Linux in cases the glibc is no longer compatible with newer or older kernels. Fixed in 0.5.22.2 already.
- The `--recurse-directory` option didn't check with decision mechanisms for module inclusion, making it impossible to avoid some things.

Optimization

- Introduced specialized constant classes for empty dictionaries and other special constants, e.g. "True" and "False", so that they can have more hard coded properties and save memory by sharing constant values.
- The "technical" sharing of a variable is only consider for variables that had some sharing going in the first place, speeding things up quite a bit for that still critical check.
- Memory savings coming from enhanced trace storage are already visible at about 1%. That is not as much as the reloading will mean, but still helpful to use less overall.

Cleanups

- The global variable registry was removed. It was in the way of unloading and reloading modules easily. Instead variables are now attached to their owner and referenced by other users. When they are released, these variables are released.
- Global variable traces were removed. Instead each variable has a list of the traces attached to it. For non-shared variables, this allows to sooner tell attributes of those variables, allowing for sooner optimization of them.
- No longer trace all initial users of a variable, just merely if there were such and if it constitutes sharing syntactically too. Not only does this save memory, it avoids useless references of the variable to functions that stop using it due to optimization.
- Create constant nodes via a factory function to avoid non-special instances where variants exist that would be faster to use.
- Moved the C string functions to a proper `nuitka.utils.CStrings` package as we use it for better code names of functions and modules.
- Made `functions` and explicit child node of modules, which makes their use more generic, esp. for re-loading modules.
- Have a dedicated function for building frame nodes, making it easier to see where they are created.

Summary

This release is the result of a couple of months work, and somewhat means that proper re-loading of cached results is becoming in sight. The reloading of modules still fails for some things, and more changes will be needed, but with that out of the way, Nuitka's footprint is about to drop and making it then absolutely scalable. Something considered very important before starting to trace more information about values.

This next thing big ought to be one thing that structurally holds Nuitka back from generating C level performance code with say integer operations.

Nuitka Release 0.5.21

This release focused on scalability work. Making Nuitka more usable in the common case, and covering more standalone use cases.

Bug Fixes

- Windows: Support for newer MinGW64 was broken by a workaround for older MinGW64 versions.
- Compatibility: Added support for the (unofficial) C-Python API `Py_GetArgcArgv` that was causing `prctl` module to fail loading on ARM platforms.
- Compatibility: The proper error message template for complex call arguments is now detected as compile time. There are changes coming, that are already in some pre-releases of CPython.
- Standalone: Wasn't properly ignoring `Tools` and other directories in the standard library.

New Features

- Windows: Detect the MinGW compiler arch and compare it to the Python arch. In case of a mismatch, the compiler is not used. Otherwise compilation or linking gives hard to understand errors. This also rules out MinGW32 as a compiler that can be used, as its arch doesn't match MinGW64 32 bits variant.
- Compile modules in two passes with the option to specify which modules will be considered for a second pass at all (compiled without program optimization) or even become bytecode.
- The developer mode installation of Nuitka in develop mode with the command `pip install -e nuitka_git_checkout_dir` is now supported too.

Optimization

- Popular modules known to not be performance relevant are no longer C compiled, e.g. `numpy.distutils` and many others frequently imported (from some other module), but mostly not used and definitely not performance relevant.

Cleanups

- The progress tracing and the memory tracing and now more clearly separate and therefore more readable.
- Moved RPM related files to new `rpm` directory.
- Moved documentation related files to `doc` directory.
- Converted import sorting helper script to Python and made it run fast.

Organisational

- The Buildbot infrastructure for Nuitka was updated to Buildbot 0.8.12 and is now maintained up to date with Ansible.
- Upgraded the Nuitka bug tracker to Roundup 1.5.1 to which I had previously contributed security fixes already active.
- Added SSL certificates from Let's Encrypt for the web server.

Summary

This release advances the scalability of Nuitka somewhat. The two pass approach does not yet carry all possible fruits. Caching of single pass compiled modules should follow for it to become consistently fast.

More work will be needed to achieve fast and scalable compilation, and that is going to remain the focus for some time.

Nuitka Release 0.5.20

This release is mostly about catching up with issues. Most address standalone problems with special modules, but there are also some general compatibility corrections, as well as important fixes for Python3.5 and coroutines and to improve compatibility with special Python variants like Anaconda under the Windows system.

Bug Fixes

- Standalone Python3.5: The `_decimal` module at least is using a `__name__` that doesn't match the name at load time, causing programs that use it to crash.
- Compatibility: For Python3.3 the `__loader__` attribute is now set in all cases, and it needs to have a `__module__` attribute. This makes inspection as done by e.g. `flask` working.
- Standalone: Added missing hidden dependencies for `Tkinter` module, adding support for this to work properly.
- Windows: Detecting the Python DLL and EXE used at compile time and preserving this information use during backend compilation. This should make sure we use the proper ones, and avoids hacks for specific Python variants, enhancing the support for Anaconda, WinPython, and CPython installations.
- Windows: The `--python-debug` flag now properly detects if the run time is supporting things and error exits if it's not available. For a CPython3.5 installation, it will switch between debug and non-debug Python binaries and DLLs.
- Standalone: Added plug-in for the `Pwm` package to properly combine it into a single file, suitable for distribution.
- Standalone: Packages from standard library, e.g. `xml` now have proper `__path__` as a list and not as a string value, which breaks code of e.g. `PyXML`.
- Standalone: Added missing dependency of `twisted.protocols.tls`.
- Python3.5: When finalizing coroutines that were not finished, a corruption of its reference count could happen under some circumstances.
- Standalone: Added missing DLL dependency of the `uuid` module at run time, which uses `ctypes` to load it.

New Features

- Added support for Anaconda Python on this Linux. Both accelerated and standalone mode work now.
- Added support for standalone mode on FreeBSD.

- The plug-in framework was expanded with new features to allow addressing some specific issues.

Cleanups

- Moved memory related stuff to dedicated utils package `nuitka.utils.MemoryUsage` as part of an effort to have more topical modules.
- Plugins now have a dedicated module through which the core accesses the API, which was partially cleaned up.
- No more "early" and "late" import detections for standalone mode. We now scan everything at the start.

Summary

This release focused on expanding plugins. These were then used to enhance the success of standalone compatibility. Eventually this should lead to a finished and documented plug-in API, which will open up the Nuitka core to easier hacks and more user contribution for these topics.

Nuitka Release 0.5.19

This release brings optimization improvements for dictionary using code. This is now lowering subscripts to dictionary accesses where possible and adds new code generation for known dictionary values. Besides this there is the usual range of bug fixes.

Bug Fixes

- Fix, attribute assignments or deletions where the assigned value or the attribute source was statically raising crashed the compiler.
- Fix, the order of evaluation during optimization was considered in the wrong order for attribute assignments source and value.
- Windows: Fix, when `g++` is the path, it was not used automatically, but now it is.
- Windows: Detect the 32 bits variant of MinGW64 too.
- Python3.4: The finalize of compiled generators could corrupt reference counts for shared generator objects. Fixed in 0.5.18.1 already.
- Python3.5: The finalize of compiled coroutines could corrupt reference counts for shared generator objects.

Optimization

- When a variable is known to have dictionary shape (assigned from a constant value, result of `dict` built-in, or a general dictionary creation), or the branch merge thereof, we lower subscripts from expecting mapping nodes to dictionary specific nodes. These generate more efficient code, and some are then known to not raise an exception.

```
def someFunction(a, b):  
    value = {a: b}  
    value["c"] = 1  
    return value
```

The above function is not yet fully optimized (dictionary key/value tracing is not yet finished), however it at least knows that no exception can raise from assigning `value["c"]` anymore and creates more efficient code for the typical `result = {}` functions.

- The use of "logical" sharing during optimization has been replaced with checks for actual sharing. So closure variables that were written to in dead code no longer inhibit optimization of the then no more shared local variable.
- Global variable traces are now faster to decide definite writes without need to check traces for this each time.

Cleanups

- No more using "logical sharing" allowed to remove that function entirely.
- Using "technical sharing" less often for decisions during optimization and instead rely more often on proper variable registry.
- Connected variables with their global variable trace statically avoid the need to check in variable registry for it.
- Removed old and mostly unused "assume unclear locals" indications, we use global variable traces for this now.

Summary

This release aimed at dictionary tracing. As a first step, the value assign is now traced to have a dictionary shape, and this then used to lower the operations which used to be normal subscript operations to mapping, but now can be more specific.

Making use of the dictionary values knowledge, tracing keys and values is not yet inside the scope, but expected to follow. We got the first signs of type inference here, but to really take advantage, more specific shape tracing will be needed.

Nuitka Release 0.5.18

This release mainly has a scalability focus. While there are few compatibility improvements, the larger goal has been to make Nuitka compilation and the final C compilation faster.

Bug Fixes

- Compatibility: The nested arguments functions can now be called using their keyword arguments.

```
def someFunction(a, (b, c)):
    return a, b, c

someFunction(a=1, **{" .1": (2, 3)})
```

- Compatibility: Generators with Python3.4 or higher now also have a `__del__` attribute, and therefore properly participate in finalization. This should improve their interactions with garbage collection reference cycles, although no issues had been observed so far.
- Windows: Was outputting command line arguments debug information at program start. Fixed in 0.5.17.1 already.

Optimization

- Code generated for parameter parsing is now a *lot* less verbose. Python level loops and conditionals to generate code for each variable has been replaced with C level generic code. This will speed up the backend compilation by a lot.

- Function calls with constant arguments were speed up specifically, as their call is now fully prepared, and yet using less code. Variable arguments are also faster, and all defaulted arguments are also much faster. Method calls are not affected by these improvements though.
- Nested argument functions now have a quick call entry point as well, making them faster to call too.
- The `slice` built-in, and internal creation of slices (e.g. in re-formulations of Python3 slices as subscripts) cannot raise.
- Standalone: Avoid inclusion of bytecode of `unittest.test`, `sqlite3.test`, `distutils.test`, and `ensurepip`. These are not needed, but simply bloat the amount of bytecode used on e.g. macOS.
- Speed up compilation with Nuitka itself by avoid to copying and constructing variable lists as much as possible using an always accurate variable registry.

Cleanups

- Nested argument functions of Python2 are now re-formulated into a wrapping function that directly calls the actual function body with the unpacking of nested arguments done in nodes explicitly. This allows for better optimization and checks of these steps and potential in-lining of these functions too.
- Unified slice object creation and built-in `slice` nodes, these were two distinct nodes before.
- The code generation for all statement kinds is now done via dispatching from a dictionary instead of long `elif` chains.
- Named nodes more often consistently, e.g. all loop related nodes start with `Loop` now, making them easier to group.
- Parameter specifications got simplified to work without variables where it is possible.

Organisational

- Nuitka is now available on the social code platforms gitlab as well.

Summary

Long standing weaknesses have been addressed in this release, also quite a few structural cleanups have been performed, e.g. strengthening the role of the variable registry to always be accurate, is groundlaying to further improvement of optimization.

However, this release cycle was mostly dedicated to performance of the actual compilation, and more accurate information was needed to e.g. not search for information that should be instant.

Upcoming releases will focus on usability issues and further optimization, it was nice however to see speedups of created code even from these scalability improvements.

Nuitka Release 0.5.17

This release is a major feature release, as it adds full support for Python3.5 and its coroutines. In addition, in order to properly support coroutines, the generator implementation got enhanced. On top of that, there is the usual range of corrections.

Bug Fixes

- Windows: Command line arguments that are unicode strings were not properly working.
- Compatibility: Fix, only the code object attached to exceptions contained all variable names, but not the one of the function object.

- Python3: Support for virtualenv on Windows was using non-portable code and therefore failing.
- The tree displayed with `--display-tree` duplicated all functions and did not resolve source lines for functions. It also displayed unused functions, which is not helpful.
- Generators with parameters leaked C level memory for each instance of them leading to memory bloat for long running programs that use a lot of generators. Fixed in 0.5.16.1 already.
- Don't drop positional arguments when called with `--run`, also make it an error if they are present without that option.

New Features

- Added full support for Python3.5, coroutines work now too.

Optimization

- Optimized frame access of generators to not use both a local frame variable and the frame object stored in the generator object itself. This gave about 1% speed up to setting them up.
- Avoid having multiple code objects for functions that can raise and have local variables. Previously one code object would be used to create the function (with parameter variable names only) and when raising an exception, another one would be used (with all local variable names). Creating them both at start-up was wasteful and also needed two tuples to be created, thus more constants setup code.
- The entry point for generators is now shared code instead of being generated for each one over and over. This should make things more cache local and also results in less generated C code.
- When creating frame codes, avoid working with strings, but use proper emission for less memory churn during code generation.

Organisational

- Updated the key for the Debian/Ubuntu repositories to remain valid for 2 more years.
- Added support for Fedora 23.
- MinGW32 is no more supported, use MinGW64 in the 32 bits variant, which has less issues.

Cleanups

- Detecting function type ahead of times, allows to handle generators different from normal functions immediately.
- Massive removal of code duplication between normal functions and generator functions. The later are now normal functions creating generator objects, which makes them much more lightweight.
- The `return` statement in generators is now immediately set to the proper node as opposed to doing this in variable closure phase only. We can now use the ahead knowledge of the function type.
- The `nonlocal` statement is now immediately checked for syntax errors as opposed to doing that only in variable closure phase.
- The name of contraction making functions is no longer skewed to empty, but the real thing instead. The code name is solved differently now.
- The `local_locals` mode for function node was removed, it was always true ever since Python2 list contractions stop using pseudo functions.
- The outline nodes allowed to provide a body when creating them, although creating that body required using the outline node already to create temporary variables. Removed that argument.

- Removed PyLint false positive annotations no more needed for PyLint 1.5 and solved some TODOs.
- Code objects are now mostly created from specs (not yet complete) which are attached and shared between statement frames and function creations nodes, in order to have less guess work to do.

Tests

- Added the CPython3.5 test suite.
- Updated generated doctests to fix typos and use common code in all CPython test suites.

Summary

This release continues to address technical debt. Adding support for Python3.5 was the major driving force, while at the same time removing obstacles to the changes that were needed for coroutine support.

With Python3.5 sorted out, it will be time to focus on general optimization again, but there is more technical debt related to classes, so the cleanup has to continue.

Nuitka Release 0.5.16

This is a maintenance release, largely intended to put out improved support for new platforms and minor corrections. It should improve the speed for standalone mode, and compilation in general for some use cases, but this is mostly to clean up open ends.

Bug Fixes

- Fix, the `len` built-in could give false values for dictionary and set creations with the same element.

```
# This was falsely optimized to 2 even if "a is b and a == b" was true.
len({a, b})
```

- Python: Fix, the `gi_running` attribute of generators is no longer an `int`, but `bool` instead.
- Python3: Fix, the `int` built-in with two arguments, value and base, raised `UnicodeDecodeError` instead of `ValueError` for illegal bytes given as value.
- Python3: Using `tokenize.open` to read source code, instead of reading manually and decoding from `tokenize.detect_encoding`, this handles corner cases more compatible.
- Fix, the PyLint warnings plug-in could crash in some cases, make sure it's more robust.
- Windows: Fix, the combination of Anaconda Python, MinGW 64 bits and mere acceleration was not working.
- Standalone: Preserve not only namespace packages created by `.pth` files, but also make the imports done by them. This makes it more compatible with uses of it in Fedora 22.
- Standalone: The extension modules could be duplicated, turned this into an error and cache finding them during compile time and during early import resolution to avoid duplication.
- Standalone: Handle "not found" from `ldd` output, on some systems not all the libraries wanted are accessible for every library.
- Python3.5: Fixed support for namespace packages, these were not yet working for that version yet.
- Python3.5: Fixes lack of support for unpacking in normal `tuple`, `list`, and `set` creations.

```
[*a] # this has become legal in 3.5 and now works too.
```

Now also gives compatible `SyntaxError` for earlier versions. Python2 was good already.

- Python3.5: Fix, need to reduce compiled functions to `__qualname__` value, rather than just `__name__` or else pickling methods doesn't work.
- Python3.5: Fix, added `gi_yieldfrom` attribute to generator objects.
- Windows: Fixed harmless warnings for Visual Studio 2015 in `--debug` mode.

Optimization

- Re-formulate `exec` and `eval` to default to `globals()` as the default for the locals dictionary in modules.
- The `try` node was making a description of nodes moved to the outside when shrinking its scope, which was using a lot of time, just to not be output, now these can be postponed.
- Refactored how freezing of bytecode works. Uncompiled modules are now explicit nodes too, and in the registry. We only have one or the other of it, avoiding to compile both.

Tests

- When `strace` or `dtruss` are not found, given proper error message, so people know what to do.
- The doc tests extracted and then generated for CPython3 test suites were not printing the expressions of the doc test, leading to largely decreased test coverage here.
- The CPython 3.4 test suite is now also using common runner code, and avoids ignoring all Nuitka warnings, instead more white listing was added.
- Started to run CPython 3.5 test suite almost completely, but coroutines are blocking some parts of that, so these tests that use this feature are currently skipped.
- Removed more CPython tests that access the network and are generally useless to testing Nuitka.
- When comparing outputs, normalize typical temporary file names used on posix systems.
- Coverage tests have made some progress, and some changes were made due to its results.
- Added test to cover too complex code module of `idna` module.
- Added Python3.5 only test for unpacking variants.

Cleanups

- Prepare plug-in interface to allow suppression of import warnings to access the node doing it, making the import node is accessible.
- Have dedicated class function body object, which is a specialization of the function body node base class. This allowed removing class specific code from that class.
- The use of "win_target" as a scons parameter was useless. Make more consistent use of it as a flag indicator in the scons file.
- Compiled types were mixing uses of `compiled_` prefixes, something with a space, sometimes with an underscore.

Organisational

- Improved support for Python3.5 missing compatibility with new language features.
- Updated the Developer Manual with changes that SSA is now a fact.
- Added Python3.5 Windows MSI downloads.

- Added repository for Ubuntu Wily (15.10) for download. Removed Ubuntu Utopic package download, no longer supported by Ubuntu.
- Added repository with RPM packages for Fedora 22.

Summary

So this release is mostly to lower the technical debt incurred that holds it back from supporting making more interesting changes. Upcoming releases may have continue that trend for some time.

This release is mostly about catching up with Python3.5, to make sure we did not miss anything important. The new function body variants will make it easier to implement coroutines, and help with optimization and compatibility problems that remain for Python3 classes.

Ultimately it will be nice to require a lot less checks for when function in-line is going to be acceptable. Also code generation will need a continued push to use the new structure in preparation for making type specific code generation a reality.

Nuitka Release 0.5.15

This release enables SSA based optimization, the huge leap, not so much in terms of actual performance increase, but for now making the things possible that will allow it.

This has been in the making literally for years. Over and over, there was just "one more thing" needed. But now it's there.

The release includes much stuff, and there is a perspective on the open tasks in the summary, but first out to the many details.

Bug Fixes

- Standalone: Added implicit import for `reportlab` package configuration dynamic import. Fixed in 0.5.14.1 already.
- Standalone: Fix, compilation of the `ctypes` module could happen for some import patterns, and then prevented the distribution to contain all necessary libraries. Now it is made sure to not include compiled and frozen form both. Fixed in 0.5.14.1 already.
- Fix, compilation for conditional statements where the boolean check on the condition cannot raise, could fail compilation. Fixed in 0.5.14.2 already.
- Fix, the `__import__` built-in was making static optimization assuming compile time constants to be strings, which in the error case they are not, which was crashing the compiler.

```
__import__(("some.module",)) # tuples don't work
```

This error became only apparent, because now in some cases, Nuitka forward propagates values.

- Windows: Fix, when installing Python2 only for the user, the detection of it via registry failed as it was only searching system key. This was [a github pull request](#). Fixed in 0.5.14.3 already.
- Some modules have extremely complex expressions requiring too deep recursion to work on all platforms. These modules are now included entirely as bytecode fallback.
- The standard library may contain broken code due to installation mistakes. We have to ignore their `SyntaxError`.
- Fix, pickling compiled methods was failing with the wrong kind of error, because they should not implement `__reduce__`, but only `__deepcopy__`.
- Fix, when running under wine, the check for scons binary was fooled by existence of `/usr/bin/scons`.

New Features

- Added experimental support for Python3.5, coroutines don't work yet, but it works perfectly as a 3.4 replacement.
- Added experimental Nuitka plug-in framework, and use it for the packaging of Qt plugins in standalone mode. The API is not yet stable nor polished.
- New option `--debugger` that makes `--run` execute directly in `gdb` and gives a stack trace on crash.
- New option `--profile` executes compiled binary and outputs measured performance with `vmprof`. This is work in progress and not functional yet.
- Started work on `--internal-graph` to render the SSA state into diagrams. This is work in progress and not too functional yet.
- Plug-in framework added. Not yet ready for users. Working `PyQt4` and `PyQt5` plug-in support. Experimental Windows `multiprocessing` support. Experimental PyLint warnings disable support. More to come.
- Added support for Anaconda accelerated mode on macOS by modifying the `rpath` to the Python DLL.
- Added experimental support for `multiprocessing` on Windows, which needs monkey patching of the module to support compiled methods.

Optimization

- The SSA analysis is now enabled by default, eliminating variables that are not shared, and can be forward propagated. This is currently limited mostly to compile time constants, but things won't remain that way.
- Code generation for many constructs now takes into account if a specific operation can raise or not. If e.g. an attribute look-up is known to not raise, then that is now decided by the node the looked is done to, and then more often can determine this, or even directly the value.
- Calls to C-API that we know cannot raise, no longer check, but merely assert the result.
- For attribute look-up and other operations that might be known to not raise, we now only assert that it succeeds.
- Built-in loop-ups cannot fail, merely assert that.
- Creation of built-in exceptions never raises, merely assert that too.
- More Python operation slots now have their own computations and some of these gained overloads for more compile time constant optimization.
- When taking an iterator cannot raise, this is now detected more often.
- The `try/finally` construct is now represented by duplicating the final block into all kinds of handlers (`break`, `continue`, `return`, or `except`) and optimized separately. This allows for SSA to trace values more correctly.
- The `hash` built-in now has dedicated node and code generation too. This is mostly intended to represent the side effects of dictionary look-up, but gives more compact and faster code too.
- Type `type` built-in cannot raise and has no side effect.
- Speed improvement for in-place float operations for `+=` and `*=`, as these will be common cases.

Tests

- Made the construct based testing executable with Python3.

- Removed warnings using the new PyLint warnings plug-in for the reflected test. Nuitka now uses the PyLint annotations to not warn. Also do not go into PyQt for reflected test, not needed. Many Python3 improvements for cases where there are differences to report.
- The optimization tests no longer use 2to3 anymore, made the tests portable to all versions.
- Checked more in-place operations for speed.

Organisational

- Many improvements to the coverage taking. We can hope to see public data from this, some improvements were triggered from this already, but full runs of the test suite with coverage data collection are yet to be done.

Summary

The release includes many important new directions. Coverage analysis will be important to remain certain of test coverage of Nuitka itself. This is mostly done, but needs more work to complete.

Then the graphing surely will help us to debug and understand code examples. So instead of tracing, and reading stuff, we should visualize things, to more clearly see, how things evolve under optimization iteration, and where exactly one thing goes wrong. This will be improved as it proves necessary to do just that. So far, this has been rare. Expect this to become end user capable with time. If only to allow you to understand why Nuitka won't optimize code of yours, and what change of Nuitka it will need to improve.

The comparative performance benchmarking is clearly the most important thing to have for users. It deserves to be the top priority. Thanks to the PyPy tool `vmp`, we may already be there on the data taking side, but the presenting and correlation part, is still open and a fair bit of work. It will be most important to empower users to make competent performance bug reports, now that Nuitka enters the phase, where these things matter.

As this is a lot of ground to cover. More than ever. We can make this compiler, but only if you help, it will arrive in your life time.

Nuitka Release 0.5.14

This release is an intermediate step towards value propagation, which is not considered ready for stable release yet. The major point is the elimination of the `try/finally` expressions, as they are problems to SSA. The `try/finally` statement change is delayed.

There are also a lot of bug fixes, and enhancements to code generation, as well as major cleanups of code base.

Bug Fixes

- Python3: Added support assignments trailing star assignment.

```
*a, b = 1, 2
```

This raised `ValueError` before.

- Python3: Properly detect illegal double star assignments.

```
*a, *b = c
```

- Python3: Properly detect the syntax error to star assign from non-tuple/list.

```
*a = 1
```

- Python3.4: Fixed a crash of the binary when copying dictionaries with split tables received as star arguments.
- Python3: Fixed reference loss, when using `raise a from b` where `b` was an exception instance. Fixed in 0.5.13.8 already.
- Windows: Fix, the flag `--disable-windows-console` was not properly handled for MinGW32 run time resulting in a crash.
- Python2.7.10: Was not recognizing this as a 2.7.x variant and therefore not applying minor version compatibility levels properly.
- Fix, when choosing to have frozen source references, code objects were not use the same value as `__file__` did for its filename.
- Fix, when re-executing itself to drop the `site` module, make sure we find the same file again, and not according to the `PYTHONPATH` changes coming from it. Fixed in 0.5.13.4 already.
- Enhanced code generation for `del variable` statements, where it's clear that the value must be assigned.
- When pressing CTRL-C, the stack traces from both Nuitka and Scons were given, we now avoid the one from Scons.
- Fix, the dump from `--xml` no longer contains functions that have become unused during analysis.
- Standalone: Creating or running programs from inside unicode paths was not working on Windows. Fixed in 0.5.13.7 already.
- Namespace package support was not yet complete, importing the parent of a package was still failing. Fixed in 0.5.13.7 already.
- Python2.6: Compatibility for exception check messages enhanced with newest minor releases.
- Compatibility: The `NameError` in classes needs to say `global name` and not just `name` too.
- Python3: Fixed creation of XML representation, now done without `lxml` as it doesn't support needed features on that version. Fixed in 0.5.13.5 already.
- Python2: Fix, when creating code for the largest negative constant to still fit into `int`, that was only working in the main module. Fixed in 0.5.13.5 already.
- Compatibility: The `print` statement raised an assertion on unicode objects that could not be encoded with `ascii` codec.

New Features

- Added support for Windows 10.
- Followed changes for Python 3.5 beta 2. Still only usable as a Python 3.4 replacement, no new features.
- Using a self compiled Python running from the source tree is now supported.
- Added support for Anaconda Python distribution. As it doesn't install the Python DLL, we copy it along for acceleration mode.
- Added support for Visual Studio 2015. Fixed in 0.5.13.3 already.
- Added support for self compiled Python versions running from build tree, this is intended to help debug things on Windows.

Optimization

- Function in-lining is now present in the code, but still disabled, because it needs more changes in other areas, before we can generally do it.

- Trivial outlines, result of re-formulations or function in-lining, are now in-lined, in case they just return an expression.
- The re-formulation for `or` and `and` has been giving up, eliminating the use of a `try/finally` expression, at the cost of dedicated boolean nodes and code generation for these.
This saves around 8% of compile time memory for Nuitka, and allows for faster and more complete optimization, and gets rid of a complicated structure for analysis.
- When a frame is used in an exception, its locals are detached. This was done more often than necessary and even for frames that are not necessary our own ones. This will speed up some exception cases.
- When the default arguments, or the keyword default arguments (Python3) or the annotations (Python3) were raising an exception, the function definition is now replaced with the exception, saving a code generation. This happens frequently with Python2/Python3 compatible code guarded by version checks.
- The SSA analysis for loops now properly traces "break" statement situations and merges the post-loop situation from all of them. This significantly allows for and improves optimization of code following the loop.
- The SSA analysis of `try/finally` statements has been greatly enhanced. The handler for `finally` is now optimized for exception raise and no exception raise individually, as well as for `break`, `continue` and `return` in the tried code. The SSA analysis for after the statement is now the result of merging these different cases, should they not abort.
- The code generation for `del` statements is now taking advantage should there be definite knowledge of previous value. This speed them up slightly.
- The SSA analysis of `del` statements now properly decided if the statement can raise or not, allowing for more optimization.
- For list contractions, the re-formulation was enhanced using the new outline construct instead of a pseudo function, leading to better analysis and code generation.
- Comparison chains are now re-formulated into outlines too, allowing for better analysis of them.
- Exceptions raised in function creations, e.g. in default values, are now propagated, eliminating the function's code. This happens most often with Python2/Python3 in branches. On the other hand, function creations that cannot are also annotated now.
- Closure variables that become unreferenced outside of the function become normal variables leading to better tracing and code generation for them.
- Function creations cannot raise except their defaults, keyword defaults or annotations do.
- Built-in references can now be converted to strings at compile time, e.g. when printed.

Organisational

- Removed gitorious mirror of the git repository, they shut down.
- Make it more clear in the documentation that Python2 is needed at compile time to create Python3 executables.

Cleanups

- Moved more parts of code generation to their own modules, and used registry for code generation for more expression kinds.

- Unified `try/except` and `try/finally` into a single construct that handles both through `try/except/break/continue/return` semantics. Finally is now solved via duplicating the handler into cases necessary.

No longer are nodes annotated with information if they need to publish the exception or not, this is now all done with the dedicated nodes.

- The `try/finally` expressions have been replaced with outline function bodies, that instead of side effect statements, are more like functions with return values, allowing for easier analysis and dedicated code generation of much lower complexity.
- No more "tolerant" flag for release nodes, we now decide this fully based on SSA information.
- Added helper for assertions that code flow does not reach certain positions, e.g. a function must return or raise, aborting statements do not continue and so on.
- To keep cloning of code parts as simple as possible, the limited use of `makeCloneAt` has been changed to a new `makeClone` which produces identical copies, which is what we always do. And a generic cloning based on "details" has been added, requiring to make constructor arguments and details complete and consistent.
- The re-formulation code helpers have been improved to be more convenient at creating nodes.
- The old `nuitka.codegen` module `Generator` was still used for many things. These now all got moved to appropriate code generation modules, and their users got updated, also moving some code generator functions in the process.
- The module `nuitka.codegen.CodeTemplates` got replaced with direct uses of the proper topic module from `nuitka.codegen.templates`, with some more added, and their names harmonized to be more easily recognizable.
- Added more assertions to the generated code, to aid bug finding.
- The auto-format now sorts pylint markups for increased consistency.
- Releases no longer have a `tolerant` flag, this was not needed anymore as we use SSA.
- Handle CTRL-C in `scons` code preventing per job messages that are not helpful and avoid tracebacks from `scons`, also remove more unused tools like `rpm` from out in-line copy.

Tests

- Added the CPython3.4 test suite.
- The CPython3.2, CPython3.3, and CPython3.4 test suite now run with Python2 giving the same errors. Previously there were a few specific errors, some with line numbers, some with different `SyntaxError` be raised, due to different order of checks.
This increases the coverage of the exception raising tests somewhat.
- Also the CPython3.x test suites now all pass with debug Python, as does the CPython 2.6 test suite with 2.6 now.
- Added tests to cover all forms of unpacking assignments supported in Python3, to be sure there are no other errors unknown to us.
- Started to document the reference count tests, and to make it more robust against SSA optimization. This will take some time and is work in progress.
- Made the compile library test robust against modules that raise a syntax error, checking that Nuitka does the same.
- Refined more tests to be directly executable with Python3, this is an ongoing effort.

Summary

This release is clearly major. It represents a huge step forward for Nuitka as it improves nearly every aspect of code generation and analysis. Removing the `try/finally` expression nodes proved to be necessary in order to even have the correct SSA in their cases. Very important optimization was blocked by it.

Going forward, the `try/finally` statements will be removed and dead variable elimination will happen, which then will give function inlining. This is expected to happen in one of the next releases.

This release is a consolidation of 8 hotfix releases, and many refactorings needed towards the next big step, which might also break things, and for that reason is going to get its own release cycle.

Nuitka Release 0.5.13

This release contains the first use of SSA for value propagation and massive amounts of bug fixes and optimization. Some of the bugs that were delivered as hotfixes, were only revealed when doing the value propagation as they still could apply to real code.

Bug Fixes

- Fix, relative imports in packages were not working with absolute imports enabled via future flags. Fixed in 0.5.12.1 already.
- Loops were not properly degrading knowledge from inside the loop at loop exit, and therefore this could have lead missing checks and releases in code generation for cases, for `del` statements in the loop body. Fixed in 0.5.12.1 already.
- The `or` and `and` re-formulation could trigger false assertions, due to early releases for compatibility. Fixed in 0.5.12.1 already.
- Fix, optimization of calls of constant objects (always an exception), crashed the compiler. Fixed in 0.5.12.2 already.
- Standalone: Added support for `site.py` installations with a leading `def` or `class` statement, which is defeating our attempt to patch `__file__` for it.
- Compatibility: In full compatibility mode, the tracebacks of `or` and `and` expressions are now as wrong as they are in CPython. Does not apply to `--improved` mode.
- Standalone: Added missing dependency on `QtGui` by `QtWidgets` for `PyQt5`.
- macOS: Improved parsing of `otool` output to avoid duplicate entries, which can also be entirely wrong in the case of Qt plugins at least.
- Avoid relative paths for main program with file reference mode `original`, as it otherwise changes as the file moves.
- MinGW: The created modules depended on MinGW to be in `PATH` for their usage. This is no longer necessary, as we now link these libraries statically for modules too.
- Windows: For modules, the option `--run` to immediately load the modules had been broken for a while.
- Standalone: Ignore Windows DLLs that were attempted to be loaded, but then failed to load. This happens e.g. when both `PySide` and `PyQt` are installed, and could cause the dreaded conflicting DLLs message. The DLL loaded in error is now ignored, which avoids this.
- MinGW: The resource file used might be empty, in which case it doesn't get created, avoiding an error due to that.
- MinGW: Modules can now be created again. The run time relative code uses an API that is WinXP only, and MinGW failed to find it without guidance.

Optimization

- Make direct calls out of called function creations. Initially this applies to lambda functions only, but it's expected to become common place in coming releases. This is now 20x faster than CPython.

```
# Nuitka avoids creating a function object, parsing function arguments:  
(lambda x: x)(something)
```

- Propagate assignments from non-mutable constants forward based on SSA information. This is the first step of using SSA for real compile time optimization.
- Specialized the creation of call nodes at creation, avoiding to have all kinds be the most flexible form (keyword and plain arguments), but instead only what kind of call they really are. This saves lots of memory, and makes the tree faster to visit.
- Added support for optimizing the `slice` built-in with compile time constant arguments to constants. The re-formulation for slices in Python3 uses these a lot. And the lack of this optimization prevented a bunch of optimization in this area. For Python2 the built-in is optimized too, but not as important probably.
- Added support for optimizing `isinstance` calls with compile time constant arguments. This avoids static exception raises in the `exec` re-formulation which tests for `file` type, and then optimization couldn't tell that a `str` is not a `file` instance. Now it can.
- Lower in-place operations on immutable types to normal operations. This will allow to compile time compute these more accurately.
- The re-formulation of loops puts the loop condition as a conditional statement with `break`. The `not` that needs to apply was only added in later optimization, leading to unnecessary compile time efforts.
- Removed per variable trace visit from optimization, removing useless code and compile time overhead. We are going to optimize things by making decision in assignment and reference nodes based on forward looking statements using the last trace collection.

New Features

- Added experimental support for Python 3.5, which seems to be passing the test suites just fine. The new `@ matrix` multiplier operators are not yet supported though.
- Added support for patching source on the fly. This is used to work around a (now fixed) issue with `numexpr.cpython` making type checks with the `is` operation, about the only thing we cannot detect.

Organisational

- Added repository for Ubuntu Vivid (15.04) for download. Removed Ubuntu Saucy and Ubuntu Raring package downloads, these are no longer supported by Ubuntu.
- Added repository for Debian Stretch, after Jessie release.
- Make it more clear in the documentation that in order to compile Python3, a Python2 is needed to execute Scons, but that the end result is a Python3 binary.
- The PyLint checker tool now can operate on directories given on the command line, and whitelists an error that is Windows only.

Cleanups

- Split up standalone code further, moving `depends.exe` handling to a separate module.

- Reduced code complexity of scons interface.
- Cleaned up where trace collection is being done. It was partially still done inside the collection itself instead in the owner.
- In case of conflicting DLLs for standalone mode, these are now output with nicer formatting, that makes it easy to recognize what is going on.
- Moved code to fetch `depends.exe` to dedicated module, so it's not as much in the way of standalone code.

Tests

- Made `BuiltinsTest` directly executable with Python3.
- Added construct test to demonstrate the speed up of direct lambda calls.
- The deletion of `@test` for the CPython test suite is more robust now, esp. on Windows, the symbolic links are now handled.
- Added test to cover `or` usage with in-place assignment.
- Cover local relative `import from .` with `absolute_import` future flag enabled.
- Again, more basic tests are now directly executable with Python3.

Summary

This release is major due to amount of ground covered. The reduction in memory usage of Nuitka itself (the C++ compiler will still use much memory) is very massive and an important aspect of scalability too.

Then the SSA changes are truly the first sign of major improvements to come. In their current form, without eliminating dead assignments, the full advantage is not taken yet, but the next releases will do this, and that's a major milestone to Nuitka.

The other optimization mostly stem from looking at things closer, and trying to work towards function in-lining, for which we are making a lot of progress now.

Nuitka Release 0.5.12

This release contains massive amounts of corrections for long standing issues in the import recursion mechanism, as well as for standalone issues now visible after the `__file__` and `__path__` values have changed to become run time dependent values.

Bug Fixes

- Fix, the `__path__` attribute for packages was still the original filename's directory, even in file reference mode was `runtime`.
- The use of `runtime` as default file reference mode for executables, even if not in standalone mode, was making acceleration harder than necessary. Changed to `original` for that case. Fixed in 0.5.11.1 already.
- The constant value for the smallest `int` that is not yet a `long` is created using `1` due to C compiler limitations, but `1` was not yet initialized properly, if this was a global constant, i.e. used in multiple modules. Fixed in 0.5.11.2 already.
- Standalone: Recent fixes around `__path__` revealed issues with PyWin32, where modules from `win32com.shell` were not properly recursed to. Fixed in 0.5.11.2 already.

- The importing of modules with the same name as a built-in module inside a package falsely assumed these were the built-ins which need not exist, and then didn't recurse into them. This affected standalone mode the most, as the module was then missing entirely.

```
# Inside "x.y" module:
import x.y.exceptions
```

- Similarly, the importing of modules with the same name as standard library modules could go wrong.

```
# Inside "x.y" module:
import x.y.types
```

- Importing modules on Windows and macOS was not properly checking the checking the case, making it associate wrong modules from files with mismatching case.
- Standalone: Importing with `from __future__ import absolute_import` would prefer relative imports still.
- Python3: Code generation for `try/return expr/finally` could loose exceptions when `expr` raised an exception, leading to a `RuntimeError` for `NULL` return value. The real exception was lost.
- Lambda expressions that were directly called with star arguments caused the compiler to crash.

```
(lambda *args: args)(*args) # was crashing Nuitka
```

Optimization

- Focusing on compile time memory usage, cyclic dependencies of trace merges that prevented them from being released, even when replaced were removed.
- More memory efficient updating of global SSA traces, reducing memory usage during optimization by ca. 50%.
- Code paths that cannot and therefore must not happen are now more clearly indicated to the backend compiler, allowing for slightly better code to be generated by it, as it can tell that certain code flows need not be merged.

New Features

- Standalone: On systems, where `.pth` files inject Python packages at launch, these are now detected, and taking into account. Previously Nuitka did not recognize them, due to lack of `__init__.py` files. These are mostly pip installations of e.g. `zope.interface`.
- Added option `--explain-imports` to debug the import resolution code of Nuitka.
- Added options `--show-memory` to display the amount of memory used in total and how it's spread across the different node types during compilation.
- The option `--trace-execution` now also covers early program initialisation before any Python code runs, to ease finding bugs in this domain as well.

Organisational

- Changed default for file reference mode to `original` unless standalone or module mode are used. For mere acceleration, breaking the reading of data files from `__file__` is useless.
- Added check that the in-line copy of `scons` is not run with Python3, which is not supported. Nuitka works fine with Python3, but a Python2 is required to execute `scons`.

- Discover more kinds of Python2 installations on Linux/macOS installations.
- Added instructions for macOS to the download page.

Cleanups

- Moved `oset` and `odict` modules which provide ordered sets and dictionaries into a new package `nuitka.container` to clean up the top level scope.
- Moved `SyntaxErrors` to `nuitka.tree` package, where it is used to format error messages.
- Moved `nuitka.Utils` package to `nuitka.utils.Utils` creating a whole package for utils, so as to better structure them for their purpose.

Summary

This release is a major maintenance release. Support for namespace modules injected by `*.pth` is a major step for new compatibility. The import logic improvements expand the ability of standalone mode widely. Many more use cases will now work out of the box, and less errors will be found on case insensitive systems.

There is aside of memory issues, no new optimization though as many of these improvements could not be delivered as hotfixes (too invasive code changes), and should be out to the users as a stable release. Real optimization changes have been postponed to be next release.

Nuitka Release 0.5.11

The last release represented a significant change and introduced a few regressions, which got addressed with hot fix releases. But it also had a focus on cleaning up open optimization issues that were postponed in the last release.

New Features

- The filenames of source files as found in the `__file__` attribute are now made relative for all modes, not just standalone mode.

This makes it possible to put data files along side compiled modules in a deployment.

Bug Fixes

- Local functions that reference themselves were not released. They now are.

```
def someFunction():
    def f():
        f() # referencing 'f' in 'f' caused the garbage collection to fail.
```

Recent changes to code generation attached closure variable values to the function object, so now they can be properly visited. Fixed in 0.5.10.1 already.

- Python2.6: The complex constants with real or imaginary parts `-0.0` were collapsed with constants of value `0.0`. This became more evident after we started to optimize the `complex` built-in. Fixed in 0.5.10.1 already.

```
complex(0.0, 0.0)
complex(-0.0, -0.0) # Could be confused with the above.
```

- Complex call helpers could leak references to their arguments. This was a regression. Fixed in 0.5.10.1 already.
- Parameter variables offered as closure variables were not properly released, only the cell object was, but not the value. This was a regression. Fixed in 0.5.10.1 already.
- Compatibility: The exception type given when accessing local variable values not initialized in a closure taking function, needs to be `NameError` and `UnboundLocalError` for accesses in the providing function. Fixed in 0.5.10.1 already.
- Fix support for "venv" on systems, where the system Python uses symbolic links too. This is the case on at least on Mageia Linux. Fixed in 0.5.10.2 already.
- Python3.4: On systems where `long` and `Py_ssize_t` are different (e.g. Win64) iterators could be corrupted if used by uncompiled Python code. Fixed in 0.5.10.2 already.
- Fix, generator objects didn't release weak references to them properly. Fixed in 0.5.10.2 already.
- Compatibility: The `__closure__` attributes of functions was so far not supported, and rarely missing. Recent changes made it easy to expose, so now it was added.
- macOS: A linker warning about deprecated linker option `-s` was solved by removing the option.
- Compatibility: Nuitka was enforcing that the `__doc__` attribute to be a string object, and gave a misleading error message. This check must not be done though, `__doc__` can be any type in Python.

Optimization

- Variables that need not be shared, because the uses in closure taking functions were eliminated, no longer use cell objects.
- The `try/except` and `try/finally` statements now both have actual merging for SSA, allowing for better optimization of code behind it.

```
def f():
    try:
        a = something()
    except:
        return 2

    # Since the above exception handling cannot continue the code flow,
    # we do not have to invalidate the trace of "a", and e.g. do not have
    # to generate code to check if it's assigned.
    return a
```

Since `try/finally` is used in almost all re-formulations of complex Python constructs this is improving SSA application widely. The uses of `try/except` in user code will no longer degrade optimization and code generation efficiency as much as they did.

- The `try/except` statement now reduces the scope of tried block if possible. When no statement raised, already the handling was removed, but leading and trailing statements that cannot raise, were not considered.

```
def f():
    try:
        b = 1
        a = something()
```

```
c = 1
except:
    return 2
```

This is now optimized to.

```
def f():
    b = 1
    try:
        a = something()
    except:
        return 2
    c = 1
```

The impact may on execution speed may be marginal, but it is definitely going to improve the branch merging to be added later. Note that `c` can only be optimized, because the exception handler is aborting, otherwise it would change behaviour.

- The creation of code objects for standalone mode and now all code objects was creating a distinct filename object for every function in a module, despite them being same content. This was wasteful for module loading. Now it's done only once.

Also, when having multiple modules, the code to build the run time filename used for code objects, was calling import logic, and doing lookups to find `os.path.join` again and again. These are now cached, speeding up the use of many modules as well.

Cleanups

- Nuitka used to have "variable usage profiles" and still used them to decide if a global variable is written to, in which case, it stays away from doing optimization of it to built-in lookups, and later calls.

The have been replaced by "global variable traces", which collect the traces to a variable across all modules and functions. While this is now only a replacement, and getting rid of old code, and basing on SSA, later it will also allow to become more correct and more optimized.

- The standalone now queries its hidden dependencies from a plugin framework, which will become an interface to Nuitka internals in the future.

Testing

- The use of deep hashing of constants allows us to check if constants become mutated during the run-time of a program. This allows to discover corruption should we encounter it.
- The tests of CPython are now also run with Python in debug mode, but only on Linux, enhancing reference leak coverage.
- The CPython test parts which had been disabled due to reference cycles involving compiled functions, or usage of `__closure__` attribute, were reactivated.

Organisational

- Since Google Code has shutdown, it has been removed from the Nuitka git mirrors.

Summary

This release brings exciting new optimization with the focus on the `try` constructs, now being done more optimal. It is also a maintenance release, bringing out compatibility improvements, and important bug fixes, and important usability features for the deployment of modules and packages, that further expand the use cases of Nuitka.

The git flow had to be applied this time to get out fixes for regression bug fixes, that the big change of the last release brought, so this is also to consolidate these and the other corrections into a full release before making more invasive changes.

The cleanups are leading the way to expanded SSA applied to global variable and shared variable values as well. Already the built-in detect is now based on global SSA information, which was an important step ahead.

Nuitka Release 0.5.10

This release has a focus on code generation optimization. Doing major changes away from "C++-ish" code to "C-ish" code, many constructs are now faster or got looked at and optimized.

Bug Fixes

- Compatibility: The variable name in locals for the iterator provided to the generator expression should be `.0`, now it is.
- Generators could leak frames until program exit, these are now properly freed immediately.

Optimization

- Faster exception save and restore functions that might be in-lined by the backend C compiler.
- Faster error checks for many operations, where these errors are expected, e.g. instance attribute lookups.
- Do not create traceback and locals dictionary for frame when `StopIteration` or `GeneratorExit` are raised. These tracebacks were wasted, as they were immediately released afterwards.
- Closure variables to functions and parameters of generator functions are now attached to the function and generator objects.
- The creation of functions with closure taking was accelerated.
- The creation and destruction of generator objects was accelerated.
- The re-formulation for in-place assignments got simplified and got faster doing so.
- In-place operations of `str` were always copying the string, even if was not necessary.

```
a += b # Was not re-using the storage of "a" in case of strings
```

- Python2: Additions of `int` for Python2 are now even faster.
- Access to local variable values got slightly accelerated at the expense of closure variables.
- Added support for optimizing the `complex` built-in.
- Removing unused temporary and local variables as a result of optimization, these previously still allocated storage.

Cleanup

- The use of C++ classes for variable objects was removed. Closure variables are now attached as `PyCellObject` to the function objects owning them.
- The use of C++ context classes for closure taking and generator parameters has been replaced with attaching values directly to functions and generator objects.
- The indentation of code template instantiations spanning multiple was not in all cases proper. We were using emission objects that handle it new lines in code and mere `list` objects, that don't handle them in mixed forms. Now only the emission objects are used.
- Some templates with C++ helper functions that had no variables got changed to be properly formatted templates.
- The internal API for handling of exceptions is now more consistent and used more efficiently.
- The printing helpers got cleaned up and moved to static code, removing any need for forward declaration.
- The use of `INCREASE_REFCOUNT_X` was removed, it got replaced with proper `Py_XINCREf` usages. The function was once required before "C-ish" lifted the need to do everything in one function call.
- The use of `INCREASE_REFCOUNT` got reduced. See above for why that is any good. The idea is that `Py_INCREF` must be good enough, and that we want to avoid the C function it was, even if in-lined.
- The `assertObject` function that checks if an object is not `NULL` and has positive reference count, i.e. is sane, got turned into a preprocessor macro.
- Deep hashes of constant values created in `--debug` mode, which cover also mutable values, and attempt to depend on actual content. These are checked at program exit for corruption. This may help uncover bugs.

Organisational

- Speedcenter has been enhanced with better graphing and has more benchmarks now. More work will be needed to make it useful.
- Updates to the Developer Manual, reflecting the current near finished state of "C-ish" code generation.

Tests

- New reference count tests to cover generator expressions and their usage got added.
- Many new construct based tests got added, these will be used for performance graphing, and serve as micro benchmarks now.
- Again, more basic tests are directly executable with Python3.

Summary

This is the next evolution of "C-ish" coming to pass. The use of C++ has for all practical purposes vanished. It will remain an ongoing activity to clear that up and become real C. The C++ classes were a huge road block to many things, that now will become simpler. One example of these were in-place operations, which now can be dealt with easily.

Also, lots of polishing and tweaking was done while adding construct benchmarks that were made to check the impact of these changes. Here, generators probably stand out the most, as some of the missed optimization got revealed and then addressed.

Their speed increases will be visible to some programs that depend a lot on generators.

This release is clearly major in that the most important issues got addressed, future releases will provide more tuning and completeness, but structurally the "C-ish" migration has succeeded, and now we can reap the benefits in the coming releases. More work will be needed for all in-place operations to be accelerated.

More work will be needed to complete this, but it's good that this is coming to an end, so we can focus on SSA based optimization for the major gains to be had.

Nuitka Release 0.5.9

This release is mostly a maintenance release, bringing out minor compatibility improvements, and some standalone improvements. Also new options to control the recursion into modules are added.

Bug Fixes

- Compatibility: Checks for iterators were using `PyIter_Check` which is buggy when running outside of Python core, because it's comparing pointers we don't see. Replaced with `HAS_ITERNEXT` helper which compares against the pointer as extracting for a real non-iterator object.

```
class Iterable:
    def __init__(self):
        self.consumed = 2

    def __iter__(self):
        return Iterable()

iter(Iterable()) # This is suppose to raise, but didn't with Nuitka
```

- Python3: Errors when creating class dictionaries raised by the `__prepare__` dictionary (e.g. enum classes with wrong identifiers) were not immediately raised, but only by the `type` call.

This was not observable, but might have caused issues potentially.

- Standalone macOS: Shared libraries and extension modules didn't have their DLL load paths updated, but only the main binary. This is not sufficient for more complex programs.
- Standalone Linux: Shared libraries copied into the `.dist` folder were read-only and executing `chrpath` could potentially then fail. This has not been observed, but is a conclusion of macOS fix.
- Standalone: When freezing standard library, the path of Nuitka and the current directory remained in the search path, which could lead to looking at the wrong files.

Organisational

- The `getattr` built-in is now optimized for compile time constants if possible, even in the presence of a `default` argument. This is more a cleanup than actually useful yet.
- The calling of `PyCFunction` from normal Python extension modules got accelerated, especially for the no or single argument cases where Nuitka now avoids building the tuple.

New Features

- Added the option `--recurse-pattern` to include modules per filename, which for Python3 is the only way to not have them in a package automatically.
- Added the option `--generate-cplusplus-only` to only generate the C++ source code without starting the compiler.

Mostly used for debugging and testing coverage. In the later case we do not want the C++ compiler to create any binary, but only to measure what would have been used.

Organisational

- Renamed the debug option `--c++-only` to `--recompile-c++-only` to make its purpose more clear and there now is `--generate-c++-only` too.

Tests

- Added support for taking coverage of Nuitka in a test run on a given input file.
- Added support for taking coverage for all Nuitka test runners, migrating them all to common code for searching.
- Added uniform way of reporting skipped tests, not generally used yet.

Summary

This release marks progress towards having coverage testing. Recent releases had made it clear that not all code of Nuitka is actually used at least once in our release tests. We aim at identifying these.

Another direction was to catch cases, where Nuitka leaks exceptions or is subject to leaked exceptions, which revealed previously unnoticed errors.

Important changes have been delayed, e.g. the closure variables will not yet use C++ objects to share storage, but proper `PyCellObject` for improved compatibility, and to approach a more "C-ish" status. These is unfinished code that does this. And the forward propagation of values is not enabled yet again either.

So this is an interim step to get the bug fixes and improvements accumulated out. Expect more actual changes in the next releases.

Nuitka Release 0.5.8

This release has mainly a focus on cleanups and compatibility improvements. It also advances standalone support, and a few optimization improvements, but it mostly is a maintenance release, attacking long standing issues.

Bug Fixes

- Compatibility Windows macOS: Fix importing on case insensitive systems.
It was not always working properly, if there was both a package `Something` and `something`, by merit of having files `Something/__init__.py` and `something.py`.
- Standalone: The search path was preferring system directories and therefore could have conflicting DLLs.
- Fix, the optimization of `getattr` with predictable result was crashing the compilation. This was a regression, fixed in 0.5.7.1 already.
- Compatibility: The name mangling inside classes also needs to be applied to global variables.
- Fix, providing `clang++` for `CXX` was mistakenly thinking of it as a `g++` and making version checks on it.
- Python3: Declaring `__class__` global is now a `SyntaxError` before Python3.4.
- Standalone Python3: Making use of module state in extension modules was not working properly.

New Features

- The filenames of source files as found in the `__file__` attribute are now made relative in standalone mode.

This should make it more apparent if things outside of the distribution folder are used, at the cost of tracebacks. Expect the default ability to copy the source code along in an upcoming release.

- Added experimental standalone mode support for PyQt5. At least headless mode should be working, plugins (needed for anything graphical) are not yet copied and will need more work.

Cleanup

- No longer using `imp.find_module` anymore. To solve the casing issues we needed to make our own module finding implementation finally.
- The name mangling was handled during code generation only. Moved to tree building instead.
- More code generation cleanups. The compatible line numbers are now attached during tree building and therefore better preserved, as well as that code no longer polluting code generation as much.

Organisational

- No more packages for openSUSE 12.1/12.2/12.3 and Fedora 17/18/19 as requested by the openSUSE Build Service.
- Added RPM packages for Fedora 21 and CentOS 7 on openSUSE Build Service.

Tests

- Lots of test refinements for the CPython test suites to be run continuously in Buildbot for both Windows and Linux.

Summary

This release brings about two major changes, each with the risk to break things.

One is that we finally started to have our own import logic, which has the risk to cause breakage, but apparently currently rather improved compatibility. The case issues were not fixable with standard library code.

The second one is that the `__file__` attributes for standalone mode is now no longer pointing to the original install and therefore will expose missing stuff sooner. This will have to be followed up with code to scan for missing "data" files later on.

For SSA based optimization, there are cleanups in here, esp. the one removing the name mangling, allowing to remove special code for class variables. This makes the SSA tree more reliable. Hope is that the big step (forward propagation through variables) can be made in one of the next releases.

Nuitka Release 0.5.7

This release is brings a newly supported platform, bug fixes, and again lots of cleanups.

Bug Fixes

- Fix, creation of dictionary and set literals with non-hashable indexes did not raise an exception.

```
{[]: None} # This is now a TypeError
```

Optimization

- Calls to the `dict` built-in with only keyword arguments are now optimized to mere dictionary creations. This is new for the case of non-constant arguments only of course.

```
dict(a=b, c=d)
# equivalent to
{"a": b, "c": d}
```

- Slice `del` with indexable arguments are now using optimized code that avoids Python objects too. This was already done for slice look-ups.
- Added support for `bytearray` built-in.

Organisational

- Added support for OpenBSD with fiber implementation from library, as it has no context support.

Cleanups

- Moved slicing solutions for Python3 to the re-formulation stage. So far the slice nodes were used, but only at code generation time, there was made a distinction between Python2 and Python3 for them. Now these nodes are purely Python2 and slice objects are used universally for Python3.

Tests

- The test runners now have common code to scan for the first file to compile, an implementation of the `search` mode. This will allow to introduce the ability to search for pattern matches, etc.
- More tests are directly executable with Python3.
- Added `recurse_none` mode to test comparison, making using extra options for that purpose unnecessary.

Summary

This solves long standing issues with slicing and subscript not being properly distinguished in the Nuitka code. It also contains major bug fixes that really problematic. Due to the involved nature of these fixes they are made in this new release.

Nuitka Release 0.5.6

This release brings bug fixes, important new optimization, newly supported platforms, and important compatibility improvements. Progress on all fronts.

Bug Fixes

- Closure taking of global variables in member functions of classes that had a class variable of the same name was binding to the class variable as opposed to the module variable.
- Overwriting compiled function's `__doc__` attribute more than once could corrupt the old value, leading to crashes. Fixed in 0.5.5.2 already.
- Compatibility Python2: The `exec` statement `execfile` were changing `locals()` was given as an argument.

```
def function():
    a = 1

    exec code in locals() # Cannot change local "a".
    exec code in None # Can change local "a"
    exec code
```

Previously Nuitka treated all 3 variants the same.

- Compatibility: Empty branches with a condition were reduced to only the condition, but they need in fact to also check the truth value:

```
if condition:
    pass
# must be treated as
bool(condition)
# and not (bug)
condition
```

- Detection of Windows virtualenv was not working properly. Fixed in 0.5.5.2 already.
- Large enough constants structures are now unstreamed via `marshal` module, avoiding large codes being generated with no point. Fixed in 0.5.5.2 already.
- Windows: Pressing CTRL-C gave two stack traces, one from the re-execution of Nuitka which was rather pointless. Fixed in 0.5.5.1 already.
- Windows: Searching for virtualenv environments didn't terminate in all cases. Fixed in 0.5.5.1 already.
- During installation from PyPI with Python3 versions, there were errors given for the Python2 only scons files. Fixed in 0.5.5.3 already.
- Fix, the arguments of `yield from` expressions could be leaked.
- Fix, closure taking of a class variable could have in a sub class where the module variable was meant.

```
var = 1

class C:
    var = 2

    class D:
        def f():
            # was C.var, now correctly addressed top level var
            return var
```

- Fix, setting `CXX` environment variable because the installed gcc has too low version, wasn't affecting the version check at all.
- Fix, on Debian/Ubuntu with `hardening-wrapper` installed the version check was always failing, because these report a shortened version number to Scons.

Optimization

- Local variables that must be assigned also have no side effects, making use of SSA. This allows for a host of optimization to be applied to them as well, often yielding simpler access/assign code, and discovering in more cases that frames are not necessary.
- Micro optimization to `dict` built-in for simpler code generation.

Organisational

- Added support for ARM "hard float" architecture.
- Added package for Ubuntu 14.10 for download.
- Added package for openSUSE 13.2 for download.
- Donations were used to buy a Cubox-i4 Pro. It got Debian Jessie installed on it, and will be used to run an even larger amount of tests.
- Made it more clear in the user documentation that the `.exe` suffix is used for all platforms, and why.
- Generally updated information in User Manual and Developer Manual about the optimization status.
- Using Nikola 7.1 with external filters instead of our own, outdated branch for the web site.

Cleanups

- PyLint clean for the first time ever. We now have a Buildbot driven test that this stays that way.
- Massive indentation cleanup of keyword argument calls. We have a rule to align the keywords, but as this was done manually, it could easily get out of touch. Now with a auto-format tool based on RedBaron, it's correct. Also, spacing around arguments is now automatically corrected. More to come.
- For `exec` statements, the coping back to local variables is now an explicit node in the tree, leader to cleaner code generation, as it now uses normal variable assignment code generation.
- The `MaybeLocalVariables` became explicit about which variable they might be, and contribute to its SSA trace as well, which was incomplete before.
- Removed some cases of code duplication that were marked as TODO items. This often resulted in cleanups.
- Do not use `replaceWith` on child nodes, that potentially were re-used during their computation.

Summary

The release is mainly the result of consolidation work. While the previous release contained many important enhancements, this is another important step towards full SSA, closing one loop whole (class variables and `exec` functions), as well as applying it to local variables, largely extending its use.

The amount of cleanups is tremendous, in huge part due to infrastructure problems that prevented release repeatedly. This reduces the technological debt very much.

More importantly, it would appear that now eliminating local and temporary variables that are not necessary is only a small step away. But as usual, while this may be easy to implement now, it will uncover more bugs in existing code, that we need to address before we continue.

Nuitka Release 0.5.5

This release is finally making full use of SSA analysis knowledge for code generation, leading to many enhancements over previous releases.

It also adds support for Python3.4, which has been longer in the making, due to many rather subtle issues. In fact, even more work will be needed to fully solve remaining minor issues, but these should affect no real code.

And then there is much improved support for using standalone mode together with virtualenv. This combination was not previously supported, but should work now.

New Features

- Added support for Python3.4

This means support for `clear` method of frames to close generators, dynamic `__qualname__`, affected by `global` statements, tuples as `yield from` arguments, improved error messages, additional checks, and many more detail changes.

Optimization

- Using SSA knowledge, local variable assignments now no longer need to check if they need to release previous values, they know definitely for the most cases.

```
def f():
    a = 1 # This used to check if old value of "a" needs a release
    ...
```

- Using SSA knowledge, local variable references now no longer need to check for raising exceptions, let alone produce exceptions for cases, where that cannot be.

```
def f():
    a = 1
    return a # This used to check if "a" is assigned
```

- Using SSA knowledge, local variable references now are known if they can raise the `UnboundLocalError` exception or not. This allows to eliminate frame usages for many cases. Including the above example.
- Using less memory for keeping variable information.
- Also using less memory for constant nodes.

Bug Fixes

- The standalone freezing code was reading Python source as UTF-8 and not using the code that handles the Python encoding properly. On some platforms there are files in standard library that are not encoded like that.
- The fiber implementation for Linux amd64 was not working with glibc from RHEL 5. Fixed to use now multiple `int` to pass pointers as necessary. Also use `uintptr_t` instead of `intptr_t` to transport pointers, which may be more optimal.
- Line numbers for exceptions were corrupted by `with` statements due to setting line numbers even for statements marked as internal.
- Partial support for `win32com` by adding support for its hidden `__path__` change.
- Python3: Finally figured out proper chaining of exceptions, given proper context messages for exception raised during the handling of exceptions.
- Corrected C++ memory leak for each closure variable taken, each time a function object was created.

- Python3: Raising exceptions with tracebacks already attached, wasn't using always them, but producing new ones instead.
- Some constants could cause errors, as they cannot be handled with the `marshal` module as expected, e.g. `(int,)`.
- Standalone: Make sure to propagate `sys.path` to the Python instance used to check for standard library import dependencies. This is important for virtualenv environments, which need `site.py` to set the path, which is not executed in that mode.
- Windows: Added support for different path layout there, so using virtualenv should work there too.
- The code object flag "optimized" (fast locals as opposed to locals dictionary) for functions was set wrongly to value for the parent, but for frames inside it, one with the correct value. This lead to more code objects than necessary and false `co_flags` values attached to the function.
- Options passed to `nuitka-python` could get lost.

```
nuitka-python program.py argument1 argument2 ...
```

The above is supposed to compile `program.py`, execute it immediately and pass the arguments to it. But when Nuitka decides to restart itself, it would forget these options. It does so to e.g. disable hash randomization as it would affect code generation.

- Raising tuples exception as exceptions was not compatible (Python2) or reference leaking (Python3).

Tests

- Running `2to3` is now avoided for tests that are already running on both Python2 and Python3.
- Made XML based optimization tests work with Python3 too. Previously these were only working on Python2.
- Added support for ignoring messages that come from linking against self compiled Pythons.
- Added test case for threaded generators that tortures the fiber layer a bit and exposed issues on RHEL 5.
- Made reference count test of compiled functions generic. No more code duplication, and automatic detection of shared stuff. Also a more clear interface for disabling test cases.
- Added Python2 specific reference counting tests, so the other cases can be executed with Python3 directly, making debugging them less tedious.

Cleanups

- Really important removal of "variable references". They didn't solve any problem anymore, but their complexity was not helpful either. This allowed to make SSA usable finally, and removed a lot of code.
- Removed special code generation for parameter variables, and their dedicated classes, no more needed, as every variable access code is now optimized like this.
- Stop using C++ class methods at all. Now only the destructor of local variables is actually supposed to do anything, and there are no methods anymore. The unused `var_name` got removed, `setVariableValue` is now done manually.
- Moved assertions for the fiber layer to a common place in the header, so they are executed on all platforms in debug mode.
- As usual, also a bunch of cleanups for PyLint were applied.

- The `locals` built-in code now uses code generation for accessing local variable values instead having its own stuff.

Organisational

- The Python version 3.4 is now officially supported. There are a few problems open, that will be addressed in future releases, none of which will affect normal people though.
- Major cleanup of Nuitka options.
 - Windows specific stuff is now in a dedicated option group. This includes options for icon, disabling console, etc.
 - There is now a dedicated group for controlling backend compiler choices and options.
- Also pickup `g++44` automatically, which makes using Nuitka on CentOS5 more automatic.

Summary

This release represents a very important step ahead. Using SSA for real stuff will allow us to build the trust necessary to take the next steps. Using the SSA information, we could start implementing more optimizations.

Nuitka Release 0.5.4

This release is aiming at preparatory changes to enable optimization based on SSA analysis, introducing a variable registry, so that variables no longer trace their references to themselves.

Otherwise, MinGW64 support has been added, and lots of bug fixes were made to improve the compatibility.

Optimization

- Using new variable registry, now properly detecting actual need for sharing variables. Optimization may discover that it is unnecessary to share a variable, and then it no longer is. This also allows `--debug` without it reporting unused variable warnings on Python3.
- Scons startup has been accelerated, removing scans for unused tools, and avoiding making more than one gcc version check.

Bug Fixes

- Compatibility: In case of unknown encodings, Nuitka was not giving the name of the problematic encoding in the error message. Fixed in 0.5.3.3 already.
- Submodules with the same name as built-in modules were wrongly shadowed. Fixed in 0.5.3.2 already.
- Python3: Added implementations of `is_package` to the meta path based loader.
- Python3.4: Added `find_spec` implementation to the meta path based loader for increased compatibility.
- Python3: Corrections for `--debug` to work with Python3 and MSVC compiler more often.
- Fixed crash with `--show-scons` when no compiler was found. Fixed in 0.5.3.5 already.
- Standalone: Need to blacklist `lib2to3` from standard library as well. Fixed in 0.5.3.4 already.
- Python3: Adapted to changes in `SyntaxError` on newer Python releases, there is now a `msg` that can override `reason`.

- Standalone Windows: Preserve `sys.executable` as it might be used to fork binaries.
- Windows: The caching of Scons was not arch specific, and files could be used again, even if changing the arch from `x86` to `x86_64` or back.
- Windows: On 32 bit Python it can happen that with large number of generators running concurrently (>1500), one cannot be started anymore. Raising an `MemoryError` now.

Organisational

- Added support for MinGW64. Currently needs to be run with `PATH` environment properly set up.
- Updated internal version of Scons to 2.3.2, which breaks support for VS 2008, but adds support for VS 2013 and VS 2012. The VS 2013 is now the recommended compiler.
- Added RPM package and repository for RHEL 7.
- The output of `--show-scons` now includes the used compiler, including the MSVC version.
- Added option `--msvc` to select the MSVC compiler version to use, which overrides automatic selection of the latest.
- Added option `-python-flag=no_warnings` to disable user and deprecation warnings at run time.
- Repository for Ubuntu Raring was removed, no more supported by Ubuntu.

Cleanups

- Made technical and logical sharing decisions separate functions and implement them in a dedicated variable registry.
- The Scons file has seen a major cleanup.

Summary

This release is mostly a maintenance release. The Scons integrations has been heavily visited, as has been Python3 and esp. Python3.4 compatibility, and results from the now possible debug test runs.

Standalone should be even more practical now, and MinGW64 is an option for those cases, where MSVC is too slow.

Nuitka Release 0.5.3

This release is mostly a follow up, resolving points that have become possible to resolve after completing the C-ish evolution of Nuitka. So this is more of a service release.

New Features

- Improved mode `--improved` now sets error lines more properly than CPython does in many cases.
- The `-python-flag=-S` mode now preserves `PYTHONPATH` and therefore became usable with `virtualenv`.

Optimization

- Line numbers of frames no longer get set unless an exception occurs, speeding up the normal path of execution.
- For standalone mode, using `--python-flag-S` is now always possible and yields less module usage, resulting in smaller binaries and faster compilation.

Bug Fixes

- Corrected an issue for frames being optimized away where in fact they are still necessary. Fixed in 0.5.2.1 already.
- Fixed handling of exception tests as side effects. These could be remainders of optimization, but didn't have code generation. Fixed in 0.5.2.1 already.
- Previously Nuitka only ever used the statement line as the line number for all the expression, even if it spawned multiple lines. Usually nothing important, and often even more correct, but sometimes not. Now the line number is most often the same as CPython in full compatibility mode, or better, see above.
- Python3.4: Standalone mode for Windows is working now.
- Standalone: Undo changes to `PYTHONPATH` or `PYTHONHOME` allowing potentially forked CPython programs to run properly.
- Standalone: Fixed import error when using PyQt and Python3.

New Tests

- For our testing approach, the improved line number handling means we can undo lots of changes that are no more necessary.
- The compile library test has been extended to cover a third potential location where modules may live, covering the `matplotlib` module as a result.

Cleanups

- In Python2, the list contractions used to be re-formulated to be function calls that have no frame stack entry of their own right. This required some special handling, in e.g. closure taking, and determining variable sharing across functions.
This now got cleaned up to be properly in-lined in a `try/finally` expression.
- The line number handling got simplified by pushing it into error exits only, removing the need to micro manage a line number stack which got removed.
- Use `intptr_t` over `unsigned long` to store fiber code pointers, increasing portability.

Organisational

- Providing own Debian/Ubuntu repositories for all relevant distributions.
- Windows MSI files for Python 3.4 were added.
- Hosting of the web site was moved to metal server with more RAM and performance.

Summary

This release brings about structural simplification that is both a follow-up to C-ish, as well as results from a failed attempt to remove static "variable references" and be fully SSA based. It incorporates changes aimed at making this next step in Nuitka evolution smaller.

Nuitka Release 0.5.2

This is a major release, with huge changes to code generation that improve performance in a significant way. It is the result of a long development period, and therefore contains a huge jump ahead.

New Features

- Added experimental support for Python 3.4, which is still work in progress.
- Added support for virtualenv on macOS.
- Added support for virtualenv on Windows.
- Added support for macOS X standalone mode.
- The code generation uses no header files anymore, therefore adding a module doesn't invalidate all compiled object files from caches anymore.
- Constants code creation is now distributed, and constants referenced in a module are declared locally. This means that changing a module doesn't affect the validity of other modules object files from caches anymore.

Optimization

- C-ish code generation uses less C++ classes and generates more C-like code. Explicit temporary objects are now used for statement temporary variables.
- The constants creation code is no more in a single file, but distributed across all modules, with only shared values created in a single file. This means improved scalability. There are remaining bad modules, but more often, standalone mode is now fast.
- Exception handling no longer uses C++ exception, therefore has become much faster.
- Loops that only break are eliminated.
- Dead code after loops that do not break is now removed.
- The `try/finally` and `try/except` constructs are now eliminated, where that is possible.
- The `try/finally` part of the re-formulation for `print` statements is now only done when printing to a file, avoiding useless node tree bloat.
- Tuples and lists are now generated with faster code.
- Locals and global variables are now access with more direct code.
- Added support for the anonymous `code` type built-in.
- Added support for `compile` built-in.
- Generators that statically return immediately, e.g. due to optimization results, are no longer using frame objects.
- The complex call helpers use no pseudo frames anymore. Previous code generation required to have them, but with C-ish code generation that is no more necessary, speeding up those kind of calls.
- Modules with only code that cannot raise, need not have a frame created for them. This avoids useless code size bloat because of them. Previously the frame stack entry was mandatory.

Bug Fixes

- Windows: The resource files were cached by Scons and re-used, even if the input changed. The could lead to corrupted incremental builds. Fixed in 0.5.1.1 already.
- Windows: For functions with too many local variables, the MSVC failed with an error "C1026: parser stack overflow, program too complex". The rewritten code generation doesn't burden the compiler as much.
- Compatibility: The timing deletion of nested call arguments was different from C++. This shortcoming has been addressed in the rewritten code generation.

- Compatibility: The `__future__` flags and `CO_FREECELL` were not present in frame flags. These were then not always properly inherited to `eval` and `exec` in all cases.
- Compatibility: Compiled frames for Python3 had `f_restricted` attribute, which is Python2 only. Removed it.
- Compatibility: The `SyntaxError` of having a `continue` in a finally clause is now properly raised.
- Python2: The `exec` statement with no locals argument provided, was preventing list contractions to take closure variables.
- Python2: Having the ASCII encoding declared in a module wasn't working.
- Standalone: Included the `idna` encoding as well.
- Standalone: For `virtualenv`, the file `orig-prefix.txt` needs to be present, now it's copied into the "dist" directory as well. Fixed in 0.5.1.1 already.
- Windows: Handle cases, where Python and user program are installed on different volumes.
- Compatibility: Can now finally use `execfile` as an expression. One of our oldest issues, no 5, is finally fixed after all this time thanks to C-ish code generation.
- Compatibility: The order or call arguments deletion is now finally compatible. This too is thanks to C-ish code generation.
- Compatibility: Code object flags are now more compatible for Python3.
- Standalone: Removing "rpath" settings of shared libraries and extension modules included. This makes standalone binaries more robust on Fedora 20.
- Python2: Wasn't falsely rejecting `unicode` strings as values for `int` and `long` variants with base argument provided.
- Windows: For Python3.2 and 64 bits, global variable accesses could give false `NameError` exceptions. Fixed in 0.5.1.6 already.
- Compatibility: Many `exec` and `eval` details have become more correctly, the argument handling is more compatible, and e.g. future flags are now passed along properly.
- Compatibility: Using `open` with no arguments is now giving the same error.

Organisational

- Replying to email from the issue tracker works now.
- Added option name alias `--xml` for `--dump-xml`.
- Added option name alias `--python-dbg` for `--python-debug`, which actually might make it a bit more clear that it is about using the CPython debug run time.
- Remove option `--dump-tree`, it had been broken for a long time and unused in favor of XML dumps.
- New digital art folder with 3D version of Nuitka logo. Thanks to Juan Carlos for creating it.
- Using "README.rst" instead of "README.txt" to make it look better on web pages.
- More complete whitelisting of missing imports in standard library. These should give no warnings anymore.
- Updated the Nuitka GUI to the latest version, with enhanced features.
- The builds of releases and update of the [downloads page](#) is now driven by Buildbot. Page will be automatically updated as updated binaries arrive.

Cleanups

- Temporary keeper variables and the nodes to handle them are now unified with normal temporary variables, greatly simplifying variable handling on that level.
- Less code is coming from templates, more is actually derived from the node tree instead.
- Releasing the references to temporary variables is now always explicit in the node tree.
- The publishing and preservation of exceptions in frames was turned into explicit nodes.
- Exception handling is now done with a single handle that checks with branches on the exception. This eliminates exception handler nodes.
- The `dir` built-in with no arguments is now re-formulated to `locals` or `globals` with their `.keys()` attribute taken.
- Dramatic amounts of cleanups to code generation specialities, that got done right for the new C-ish code generation.

New Tests

- Warnings from MSVC are now error exits for `--debug` mode too, expanding the coverage of these tests.
- The outputs with `python-dbg` can now also be compared, allowing to expand test coverage for reference counts.
- Many of the basic tests are now executable with Python3 directly. This allows for easier debug.
- The library compilation test is now also executed with Python3.

Summary

This release would deserve more than a minor number increase. The C-ish code generation, is a huge body of work. In many ways, it lays ground to taking benefit of SSA results, that previously would not have been possible. In other ways, it's incomplete in not yet taking full advantage yet.

The release contains so many improvements, that are not yet fully realized, but as a compiler, it also reflects a stable and improved state.

The important changes are about making SSA even more viable. Many of the problematic cases, e.g. exception handlers, have been stream lined. A whole class of variables, temporary keepers, has been eliminated. This is big news in this domain.

For the standalone users, there are lots of refinements. There is esp. a lot of work to create code that doesn't show scalability issues. While some remain, the most important problems have been dealt with. Others are still in the pipeline.

More work will be needed to take full advantage. This has been explained in a [separate post](#) in greater detail.

Nuitka Release 0.5.1

This release brings corrections and major improvements to how standalone mode performs. Much of it was contributed via patches and bug reports.

Bug Fixes

- There was a crash when using `next` on a non-iterable. Fixed in 0.5.0.1 already.
- Module names with special characters not allowed in C identifiers were not fully supported. Fixed in 0.5.0.1 already.

- Name mangling for classes with leading underscores was not removing them from resulting attribute names. This broke at `__slots__` with private attributes for such classes. Fixed in 0.5.0.1 already.
- Standalone on Windows might need "cp430" encoding. Fixed in 0.5.0.2 already.
- Standalone mode didn't work with `lxml.etree` due to lack of hard coded dependencies. When a shared library imports things, Nuitka cannot detect it easily.
- Wasn't working on macOS 64 bits due to using Linux 64 bits specific code. Fixed in 0.5.0.2 already.
- On MinGW the constants blob was not properly linked on some installations, this is now done differently (see below).

New Features

- Memory usages are now traced with `--show-progress` allowing us to trace where things go wrong.

Optimization

- Standalone mode now includes standard library as bytecode by default. This is workaround scalability issues with many constants from many modules. Future releases are going to undo it.
- On Windows the constants blob is now stored as a resource, avoiding compilation via C code for MSVC as well. MinGW was changed to use the same code.

New Tests

- Expanded test coverage for "standalone mode" demonstrating usage of "hex" encoding, PySide, and PyQt packages.

Summary

This release is mostly an interim maintenance release for standalone. Major changes that provide optimization beyond that, termed "C-ish code generation" are delayed for future releases.

This release makes standalone practical which is an important point. Instead of hour long compilation, even for small programs, we are down to less than a minute.

The solution of the scalability issues with many constants from many modules will be top priority going forward. Since they are about how even single use constants are created all in one place, this will be easy, but as large changes are happening in "C-ish code generation", we are waiting for these to complete.

Nuitka Release 0.5.0

This release breaks interface compatibility, therefore the major version number change. Also "standalone mode" has seen significant improvements on both Windows, and Linux. Should work much better now.

But consider that this part of Nuitka is still in its infancy. As it is not the top priority of mine for Nuitka, which primarily is intended as an super compatible accelerator of Python, it will continue to evolve nearby.

There is also many new optimization based on structural improvements in the direction of actual SSA.

Bug Fixes

- The "standalone mode" was not working on all Redhat, Fedora, and openSUSE platforms and gave warnings with older compilers. Fixed in 0.4.7.1 already.
- The "standalone mode" was not including all useful encodings. Fixed in 0.4.7.2 already.

- The "standalone mode" was defaulting to `--python-flag=-S` which disables the parsing of "site" module. That unfortunately made it necessary to reach some modules without modifying `PYTHONPATH` which conflicts with the "out-of-the-box" experience.
- The "standalone mode" is now handling packages properly and generally working on Windows as well.
- The syntax error of having an all catching except clause and then a more specific one wasn't causing a `SyntaxError` with Nuitka.

```
try:
    something()
except:
    somehandling()
except TypeError:
    notallowed()
```

- A corruption bug was identified, when re-raising exceptions, the top entry of the traceback was modified after usage. Depending on `malloc` this was potentially causing an endless loop when using it for output.

New Features

- Windows: The "standalone" mode now properly detects used DLLs using [Dependency Walker](#) which it offers to download and extra for you.

It is used as a replacement to `ldd` on Linux when building the binary, and as a replacement of `strace` on Linux when running the tests to check that nothing is loaded from the outside.

Optimization

- When iterating over `list`, `set`, this is now automatically lowered to `tuples` avoiding the mutable container types.

So the following code is now equivalent:

```
for x in [a, b, c]:
    ...

# same as
for x in (a, b, c):
    ...
```

For constants, this is even more effective, because for mutable constants, no more is it necessary to make a copy.

- Python2: The iteration of large `range` is now automatically lowered to `xrange` which is faster to loop over, and more memory efficient.
- Added support for the `xrange` built-in.
- The statement only expression optimization got generalized and now is capable of removing useless parts of operations, not only the whole thing when it has not side effects.

```
[a, b]

# same as
```

```
a  
b
```

This works for all container types.

Another example is `type` built-in operation with single argument. When the result is not used, it need not be called.

```
type(a)  
  
# same as  
a
```

And another example `is` and `is not` have no effect of their own as well, therefore:

```
a is b  
  
# same as  
a  
b
```

- Added proper handling of conditional expression branches in SSA based optimization. So far these branches were ignored, which only acceptable for temporary variables as created by tree building, but not other variable types. This is preparatory for introducing SSA for local variables.

Organisational

- The option `--exe` is now ignored and creating an executable is the default behavior of `nuitka`, a new option `--module` allows to produce extension modules.
- The binary `nuitka-python` was removed, and is replaced by `nuitka-run` with now only implies `--execute` on top of what `nuitka` is.
- Using dedicated [Buildbot](#) for continuous integration testing and release creation as well.
- The [Downloads](#) now offers MSI files for Win64 as well.
- Discontinued the support for cross compilation to Win32. That was too limited and the design choice is to have a running CPython instance of matching architecture at Nuitka compile time.

New Tests

- Expanded test coverage for "standalone mode" demonstrating usage of "hex" encoding, and PySide package.

Summary

The "executable by default" interface change improves on the already high ease of use. The new optimization do not give all that much in terms of numbers, but are all signs of structural improvements, and it is steadily approaching the point, where the really interesting stuff will happen.

The progress for standalone mode is of course significant. It is still not quite there yet, but it is making quick progress now. This will attract a lot of attention hopefully.

As for optimization, the focus for it has shifted to making exception handlers work optimal by default (publish the exception to `sys.exc_info()` and create traceback only when necessary) and be based on

standard branches. Removing special handling of exception handlers, will be the next big step. This release includes some correctness fixes stemming from that work already.

Nuitka Release 0.4.7

This release includes important new features, lots of polishing cleanups, and some important performance improvements as well.

Bug Fixes

- The RPM packages didn't build due to missing in-line copy of Scons. Fixed in 0.4.6.1 already.
- The recursion into modules and unfreezing them was not working for packages and modules anymore. Fixed in 0.4.6.2 already.
- The Windows installer was not including Scons. Fixed in 0.4.6.3 already.
- Windows: The immediate execution as performed by `nuitka --execute` was not preserving the exit code.
- Python3.3: Packages without `__init.py__` were not properly embedding the name-space package as well.
- Python3: Fix, modules and packages didn't add themselves to `sys.modules` which they should, happened only for programs.
- Python3.3: Packages should set `__package__` to their own name, not the one of their parents.
- Python3.3: The `__qualname__` of nested classes was corrected.
- For modules that recursed to other modules, an infinite loop could be triggered when comparing types with rich comparisons.

New Features

- The "standalone" mode allows to compile standalone binaries for programs and run them without Python installation. The DLLs loaded by extension modules on Windows need to be added manually, on Linux these are determined automatically already.

To achieve running without Python installation, Nuitka learned to freeze bytecode as an alternative to compiling modules, as some modules need to be present when the CPython library is initialized.

- New option `--python-flag` allows to specify flags to the compiler that the "python" binary normally would. So far `-S` and `-v` are supported, with sane aliases `no_site` and `trace_imports`.

The recommended use of `--python-flag=-S` is to avoid dependency creep in standalone mode compilations, because the `site` module often imports many useless things that often don't apply to target systems.

Optimization

- Faster frame stack handling for functions without `try/except` (or `try/finally` in Python3). This gives a speed boost to "PyStone" of ca. 2.5% overall.
- Python2: Faster attribute getting and setting, handling special cases at compile time. This gives a minor speed boost to "PyStone" of ca. 0.5% overall.
- Python2: Much quicker calls of `__getattr__` and `__setattr__` as this is now using the quicker call method avoiding temporary tuples.
- Don't treat variables usages used in functions called directly by their owner as shared. This leads to more efficient code generation for contractions and class bodies.

- Create `unicode` constants directly from their UTF-8 string representation for Python2 as well instead of un-streaming. So far this was only done for Python3. Affects only program start-up.
- Directly create `int` and `long` constants outside of 2^{31} and $2^{32}-1$, but only limited according to actual platform values. Affects only program start-up.
- When creating `set` values, no longer use a temporary `tuple` value, but use a properly generated helper functions instead. This makes creating sets much faster.
- Directly create `set` constants instead of un-streaming them. Affects only program start-up.
- For correct line numbers in traceback, the current frame line number must be updated during execution. This was done more often than necessary, e.g. loops set the line number before loop entry, and at first statement.
- Module variables are now accessed even faster, the gain for "PyStone" is only 0.1% and mostly the result of leaner code.

Organisational

- The "standalone mode" code (formerly known as "portable mode" has been redone and activated. This is a feature that a lot of people expect from a compiler naturally. And although the overall goal for Nuitka is of course acceleration, this kind of packaging is one of the areas where CPython needs improvement.
- Added package for Ubuntu 13.10 for download, removed packages for Ubuntu 11.04 and 11.10, no more supported.
- Added package for openSUSE 13.1 for download.
- Nuitka is now part of Arch and can be installed with `pacman -S nuitka`.
- Using dedicated [Buildbot](#) for continuous integration testing. Not yet public.
- Windows: In order to speed up repeated compilation on a platform without `ccache`, added Scons level caching in the build directory.
- Disabled hash randomization for inside Nuitka (but not in ultimately created binaries) for a more stable output, because dictionary constants will not change around. This makes the build results possible to cache for `ccache` and Scons as well.

Tests

- The `programs` tests cases now fail if module or directory recursion is not working, being executed in another directory.
- Added test runner for packages, with initial test case for package with recursion and sub-packages.
- Made some test cases more strict by reducing `PYTHONPATH` provision.
- Detect use of extra flags in tests that don't get consumed avoiding ineffective flags.
- Use `--execute` on Windows as well, the issue that prevented it has been solved after all.

Cleanups

- The generated code uses `const_`, `var_`, `par_` prefixes in the generated code and centralized the decision about these into single place.
- Module variables no longer use C++ classes for their access, but instead accessor functions, leading to much less code generated per module variable and removing the need to trace their usage during code generation.

- The test runners now share common code in a dedicated module, previously they replicated it all, but that turned out to be too tedious.
- Massive general cleanups, many of which came from new contributor Juan Carlos Paco.
- Moved standalone and freezer related codes to dedicated package `nuitka.freezer` to not pollute the `nuitka` package name space.
- The code generation use variable identifiers and their accesses was cleaned up.
- Removed several not-so-special case identifier classes because they now behave more identical and all work the same way, so a parameters can be used to distinguish them.
- Moved main program, function object, set related code generation to dedicated modules.

Summary

This release marks major technological progress with the introduction of the much sought standalone mode and performance improvements from improved code generation.

The major break through for SSA optimization was not yet achieved, but this is again making progress in the direction of it. Harmonizing variables of different kinds was an important step ahead.

Also very nice is the packaging progress, Nuitka was accepted into Arch after being in Debian Testing for a while already. Hope is to see more of this kind of integration in the future.

Nuitka Release 0.4.6

This release includes progress on all fronts. The primary focus was to advance SSA optimization over older optimization code that was already in place. In this domain, there are mostly cleanups.

Another focus has been to enhance Scons with MSVC on Windows. Nuitka now finds an installed MSVC compiler automatically, properly handles architecture of Python and Windows. This improves usability a lot.

Then this is also very much about bug fixes. There have been several hot fixes for the last release, but a complicated and major issue forced a new release, and many other small issues.

And then there is performance. As can be seen in the [performance graph](#), this release is the fastest so far. This came mainly from examining the need for comparison slots for compiled types.

And last, but not least, this also expands the base of supported platforms, adding Gentoo, and self compiled Python to the mix.

Bug Fixes

- Support Nuitka being installed to a path that contains spaces and handle main programs with spaces in their paths. Fixed in 0.4.5.1 already.
- Support Python being installed to a path that contains spaces. Fixed in 0.4.5.2 already.
- Windows: User provided constants larger than 65k didn't work with MSVC. Fixed in 0.4.5.3 already.
- Windows: The option `--windows-disable-console` was not effective with MSVC. Fixed in 0.4.5.3 already.
- Windows: For some users, Scons was detecting their MSVC installation properly already from registry, but it didn't honor the target architecture. Fixed in 0.4.5.3 already.
- When creating Python modules, these were marked as executable ("x" bit), which they are of course not. Fixed in 0.4.5.3 already.
- Python3.3: On architectures where `Py_ssize_t` is not the same as `long` this could lead to errors. Fixed in 0.4.5.3 already.

- Code that was using nested mutable constants and changed the nested ones was not executing correctly.
- Python2: Due to list contractions being re-formulated as functions, `del` was rejected for the variables assigned in the contraction.

```
[expr(x) for x in iterable()]

del x # Should work, was gave an unjustified SyntaxError.
```

New Features

- Compiled types when used in Python comparison now work. Code like this will work:

```
def f():
    pass

assert type(f) == types.FunctionType
```

This of course also works for `in` operator, and is another step ahead in compatibility, and surprising too. And best of all, this works even if the checking code is not compiled with Nuitka.

- Windows: Detecting MSVC installation from registry, if no compiler is already present in PATH.
- Windows: New options `--mingw64` to force compilation with MinGW.

Optimization

- Rich comparisons (`==`, `<`, and the like) are now faster than ever before due to a full implementation of its own in Nuitka that eliminates a bit of the overhead. In the future, we will aim at giving it type hints to make it even faster. This gives a minor speed boost to PyStone of ca. 0.7% overall.
- Integer comparisons are now treated preferably, as they are in CPython, which gives 1.3% speed boost to CPython.
- The SSA based analysis is now used to provide variable scopes for temporary variables as well as reference count needs.

Cleanups

- Replaced "value friend" based optimization code with SSA based optimization, which allowed to remove complicated and old code that was still used mainly in optimization of `or` and `and` expressions.
- Delayed declaration of temp variables and their reference type is now performed based on information from SSA, which may given more accurate results. Not using "variable usage" profiles for this anymore.
- The Scons interface and related code got a massive overhaul, making it more consistent and better documented. Also updated the internal copy to 2.3.0 for the platforms that use it, mostly Windows.
- Stop using `os.system` and `subprocess.call(..., shell = True)` as it is not really portable at all, use `subprocess.call(..., shell = False)` instead.
- As usual lots of cleanups related to line length issues and PyLint.

Organisational

- Added support for Gentoo Linux.
- Added support for self compiled Python versions with and without debug enabled.
- Added use of Nuitka fonts for headers in manuals.
- Does not install in-line copy of Scons only on systems where it is not going to be used, that is mostly non-Windows, and Linux where it is not already present. This makes for cleaner RPM packages.

Summary

While the SSA stuff is not yet bearing performance fruits, it starts to carry weight. Taking over the temporary variable handling now also means we can apply the same stuff to local variables later.

To make up for the delay in SSA driven performance improvements, there is more traditional code acceleration for rich comparisons, making it significant, and the bug fixes make Nuitka more compatible than ever.

So give this a roll, it's worth it. And feel free to join the mailing list (since closed) or [make a donation](#) to support Nuitka.

Nuitka Release 0.4.5

This release incorporates very many bug fixes, most of which were already part of hot fixes, usability improvements, documentation improvements, new logo, simpler Python3 on Windows, warnings for recursion options, and so on. So it's mostly a consolidation release.

Bug Fixes

- When targeting Python 3.x, Nuitka was using "python" to run Scons to run it under Python 2.x, which is not good enough on systems, where that is already Python3. Improved to only do the guessing where necessary (i.e. when using the in-line copy of Scons) and then to prefer "python2". Fixed in 0.4.4.1 already.
- When using Nuitka created binaries inside a "virtualenv", created programs would instantly crash. The attempt to load and patch `inspect` module was not making sure that `site` module was already imported, but inside the "virtualenv", it cannot be found unless. Fixed in 0.4.4.1 already.
- The option `--recurse-directory` to include plugin directories was broken. Fixed in 0.4.4.2 already.
- Python3: Files with "BOM" marker causes the compiler to crash. Fixed in 0.4.4.2 already.
- Windows: The generated code for `try/return/finally` was working with gcc (and therefore MinGW), but not with MSVC, causing crashes. Fixed in 0.4.4.2 already.
- The option `--recurse-all` did not recurse to package `__init__.py` files in case `from x.y import z` syntax was used. Fixed in 0.4.4.2 already.
- Python3 on macOS: Corrected link time error. Fixed in 0.4.4.2 already.
- Python3.3 on Windows: Fixed crash with too many arguments to a `kwnonly` argument using function. Fixed in 0.4.4.2 already.
- Python3.3 on Windows: Using "yield from" resulted in a link time error. Fixed in 0.4.4.2 already.
- Windows: Added back XML manifest, found a case where it is needed to prevent clashes with binary modules.
- Windows: Generators only worked in the main Python threads. Some unusual threading modules therefore failed.
- Using `sys.prefix` to find the Python installation instead of hard coded paths.

New Features

- Windows: Python3 finds Python2 installation to run Scons automatically now.
Nuitka itself runs under Python3 just fine, but in order to build the generated C++ code into binaries, it uses Scons which still needs Python2.
Nuitka will now find the Python2 installation searching Windows registry instead of requiring hard coded paths.
- Windows: Python2 and Python3 find their headers now even if Python is not installed to specific paths.
The installation path now is passed on to Scons which then uses it.
- Better error checking for `--recurse-to` and `--recurse-not-to` arguments, tell the user not to use directory paths.
- Added a warning for `--recurse-to` arguments that end up having no effect to the final result.

Cleanups

- Import mechanism got cleaned up, stopped using "PyImport_ExtendInittab". It does not handle packages, and the `sys.meta_path` based importer is now well proven.
- Moved some of the constraint collection code mess into proper places. It still remains a mess.

Organisational

- Added `LICENSE.txt` file with Apache License 2.0 text to make it more immediately obvious which license Nuitka is under.
- Added section about Nuitka license to the [User Manual](#).
- Added [Nuitka Logo](#) to the distribution.
- Use Nuitka Logo as the bitmap in the Windows installer.
- Use Nuitka Logo in the documentation ([User Manual](#) and [Developer Manual](#)).
- Enhanced documentation to number page numbers starting after table of contents, removed header/footer from cover pages.

Summary

This release is mostly the result of improvements made based on the surge of users after Europython 2013. Some people went to extents and reported their experience very detailed, and so I could aim at making e.g. their misconceptions about how recursion options work, more obvious through warnings and errors.

This release is not addressing performance improvements. The next release will be able to focus on that. I am taking my claim of full compatibility very serious, so any time it's broken, it's the highest priority to restore it.

Nuitka Release 0.4.4

This release marks the point, where Nuitka for the first time supports all major current Python versions and all major features. It adds Python 3.3 support and it adds support for threading. And then there is a massive amount of fixes that improve compatibility even further.

Aside of that, there is major performance work. One side is the optimization of call performance (to CPython non-compiled functions) and to compiled functions, both. This gave a serious improvement to performance.

Then of course, we are making other, long term performance progress, as in "--experimental" mode, the SSA code starts to optimize unused code away. That code is not yet ready for prime time yet, but the trace structure will hold.

New Features

- Python3.3 support.

The test suite of CPython3.3 passes now too. The `yield from` is now supported, but the improved argument parsing error messages are not implemented yet.

- Tracing user provided constants, now Nuitka warns about too large constants produced during optimization.

- Line numbers of expressions are now updates as evaluation progresses. This almost corrects.

Now only expression parts that cannot raise, do not update, which can still cause difference, but much less often, and then definitely useless.

- Experimental support for threads.

Threading appears to work just fine in the most cases. It's not as optimal as I wanted it to be, but that's going to change with time.

Optimization

- Previous corrections for `==`, `!=`, and `<=`, caused a performance regression for these operations in case of handling identical objects.

For built-in objects of sane types (not `float`), these operations are now accelerated again. The overreaching acceleration of `>=` was still there (bug, see below) and has been adapted too.

- Calling non-compiled Python functions from compiled functions was slower than in CPython. It is now just as fast.
- Calling compiled functions without keyword arguments has been accelerated with a dedicated entry point that may call the implementation directly and avoid parameter parsing almost entirely.
- Making calls to compiled and non-compiled Python functions no longer requires to build a temporary tuple and therefore is much faster.
- Parameter parsing code is now more compact, and re-uses error raises, or creates them on the fly, instead of hard coding it. Saves binary size and should be more cache friendly.

Bug Fixes

- Corrected false optimization of `a >= a` on C++ level.

When it's not done during Nuitka compile time optimization, the rich comparison helper still contained short cuts for `>=`. This is now the same for all the comparison operators.

- Calling a function with default values, not providing it, and not providing a value for a value without default, was not properly detecting the error, and instead causing a run time crash.

```
def f(a, b=2):  
    pass
```

```
f(b=2)
```

This now properly raises the `TypeError` exception.

- Constants created with `+` could become larger than the normally enforced limits. Not as likely to become huge, but still potentially an issue.
- The `vars` built-in, when used on something without `__dict__` attribute, was giving `AttributeError` instead of `TypeError`.
- When re-cursing to modules at compile time, script directory and current directory were used last, while at run time, it was the other way around, which caused overloaded standard library modules to not be embedded.

Thanks for the patch to James Michael DuPont.

- Super without arguments was not raising the correct `RuntimeError` exception in functions that cannot be methods, but `UnboundLocalError` instead.

```
def f():  
    super() # Error, cannot refer to first argument of f
```

- Generators no longer use `raise StopIteration` for return statements, because that one is not properly handled in `try/except` clauses, where it's not supposed to trigger, while `try/finally` should be honored.
- Exception error message when throwing non-exceptions into generators was not compatible.
- The use of `return` with value in generators is a `SyntaxError` before Python3.3, but that was not raised.
- Variable names of the `"__var"` style need to be mangled. This was only done for classes, but not for functions contained in classes, there they are now mangled too.
- Python3: Exceptions raised with causes were not properly chaining.
- Python3: Specifying the file encoding corrupted line numbers, making them all off by one.

Cleanups

- For containers (`tuple`, `list`, `set`, `dict`) defined on the source code level, Nuitka immediately created constant references from them.

For function calls, class creations, slice objects, this code is now re-used, and its dictionaries and tuples, may now become constants immediately, reducing noise in optimization steps.

- The parameter parsing code got cleaned up. There were a lot of relics from previously explored paths. And error raises were part of the templates, but now are external code.
- Global variable management moved to module objects and out of "Variables" module.
- Make sure, nodes in the tree are not shared by accident.

This helped to find a case of duplicate use in the complex call helpers functions. Code generation will now notice this kind of duplication in debug mode.

- The complex call helper functions were manually taking variable closure, which made these functions inconsistent to other functions, e.g. no variable version was allocated to assignments.

Removing the manual setting of variables allowed a huge reduction of code volume, as it became more generic code.

- Converting user provided constants to create containers into constants immediately, to avoid noise from doing this in optimization.
- The `site` module is now imported explicitly in the `__main__` module, so it can be handled by the recursion code as well. This will help portable mode.
- Many line length 80 changes, improved comments.

New Tests

- The CPython3.3 test suite was added, and run with both Python3.2 and Python3.3, finding new bugs.
- The `doctest` to code generation didn't successfully handle all tests, most notably, "test_generators.py" was giving a `SyntaxError` and therefore not actually active. Correcting that improved the coverage of generator testing.

Organisational

- The portable code is still delayed.
Support for Python3.3 was a higher priority, but the intention is to get it into shape for Europython still. Added notes about it being disabled it in the [User Manual](#) documentation.

Summary

This release is in preparation for Europython 2013. Wanted to get this much out, as it changes the status slides quite a bit, and all of that was mostly done in my Cyprus holiday a while ago.

The portable code has not seen progress. The idea here is to get this into a development version later.

Nuitka Release 0.4.3

This release expands the reach of Nuitka substantially, as new platforms and compilers are now supported. A lot of polish has been applied. Under the hood there is the continued and in-progress effort to implement SSA form in Nuitka.

New Features

- Support for new compiler: Microsoft Visual C++.
You can now use Visual Studio 2008 or Visual Studio 2010 for compiling under Windows.
- Support for NetBSD.
Nuitka works for at least NetBSD 6.0, older versions may or may not work. This required fixing bugs in the generic "fibers" implementation.
- Support for Python3 under Windows too.
Nuitka uses Scons to build the generated C++ files. Unfortunately it requires Python2 to execute, which is not readily available to call from Python3. It now guesses the default installation paths of CPython 2.7 or CPython 2.6 and it will use it for running Scons instead. You have to install it to `C:\Python26` or `C:\Python27` for Nuitka to be able to find it.
- Enhanced Python 3.3 compatibility.
The support the newest version of Python has been extended, improving compatibility for many minor corner cases.
- Added warning when a user compiles a module and executes it immediately when that references `__name__`.

Because very likely the intention was to create an executable. And esp. if there is code like this:

```
if __name__ == "__main__":  
    main()
```

In module mode, Nuitka will optimize it away, and nothing will happen on execution. This is because the command

```
nuitka --execute module
```

is behavioral more like

```
python -c "import module"
```

and that was a trap for new users.

- All Linux architectures are now supported. Due to changes in how evaluation order is enforced, we don't have to implement for specific architectures anymore.

Bug Fixes

- Dictionary creation was not fully compatible.

As revealed by using Nuitka with CPython3.3, the order in which dictionaries are to be populated needs to be reversed, i.e. CPython adds the last item first. We didn't observe this before, and it's likely the new dictionary implementation that finds it.

Given that hash randomization makes dictionaries item order undetermined anyway, this is more an issue of testing.

- Evaluation order for arguments of calls was not effectively enforced. It is now done in a standards compliant and therefore fully portable way. The compilers and platforms so far supported were not affected, but the newly supported Visual Studio C++ compiler was.
- Using a `__future__` import inside a function was giving an assertion, instead of the proper syntax error.
- Python3: Do not set the attributes `sys.exc_type`, `sys.exc_value`, `sys.exc_traceback`.
- Python3: Annotations of function worked only as long as their definition was not referring to local variables.

Optimization

- Calls with no positional arguments are now using the faster call methods.

The generated C++ code was using the `()` constant at call site, when doing calls that use no positional arguments, which is of course useless.

- For Windows now uses OS "Fibers" for Nuitka "Fibers".

Using threads for fibers was causing only overhead and with this API, MSVC had less issues too.

Organisational

- Accepting [Donations](#) via Paypal, please support funding travels, website, etc.

- The [User Manual](#) has been updated with new content. We now do support Visual Studio, documented the required LLVM version for clang, Win64 and modules may include modules too, etc. Lots of information was no longer accurate and has been updated.
- The Changelog has been improved for consistency, wordings, and styles.
- Nuitka is now available on the social code platforms as well
 - Bitbucket (since removed)
 - [GitHub](#)
 - Gitorious (since discontinued)
 - Google Code (since discontinued)
- Removed `clean-up.sh`, which is practically useless, as tests now clean up after themselves reasonably, and with `git clean -dfx` working better.
- Removed "create-environment.sh" script, which was only setting the `PATH` variable, which is not necessary.
- Added `check-with-pylint --emacs` option to make output its work with Emacs compilation mode, to allow easier fixing of warnings from PyLint.
- Documentation is formatted for 80 columns now, source code will gradually aim at it too. So far 90 columns were used, and up to 100 tolerated.

Cleanups

- Removed useless manifest and resource file creation under Windows.
Turns out this is no longer needed at all. Either CPython, MinGW, or Windows improved to no longer need it.
- PyLint massive cleanups and annotations bringing down the number of warnings by a lot.
- Avoid use of strings and built-ins as run time pre-computed constants that are not needed for specific Python versions, or Nuitka modes.
- Do not track needed tuple, list, and dict creation code variants in context, but e.g. in `nuitka.codegen.TupleCodes` module instead.
- Introduced an "internal" module to host the complex call helper functions, instead of just adding it to any module that first uses it.

New Tests

- Added basic tests for order evaluation, where there currently were None.
- Added support for "2to3" execution under Windows too, so we can run tests for Python3 installations too.

Summary

The release is clearly major step ahead. The new platform support triggered a whole range of improvements, and means this is truly complete now.

Also there is very much polish in this release, reducing the number of warnings, updated documentation, the only thing really missing is visible progress with optimization.

Nuitka Release 0.4.2

This release comes with many bug fixes, some of which are severe. It also contains new features, like basic Python 3.3 support. And the [performance diagrams](#) got expanded.

New Features

- Support for FreeBSD.

Nuitka works for at least FreeBSD 9.1, older versions may or may not work. This required only fixing some "Linuxisms" in the build process.

- New option for warning about compile time detected exception raises.

Nuitka can now warn about exceptions that will be raised at run time.

- Basic Python3.3 support.

The test suite of CPython3.2 passes and fails in a compatible way. New feature `yield from` is not yet supported, and the improved argument parsing error messages are not implemented yet.

Bug Fixes

- Nuitka already supported compilation of "main directories", i.e. directories with a "`__main__.py`" file inside. The resulting binary name was "`__main__.exe`" though, but now it is "`directory.exe`"

```
# ls directory
__main__.py

# nuitka --exe directory
# ls
directory directory.exe
```

This makes this usage more obvious, and fixes an older issue for this feature.

- Evaluation order of binary operators was not enforced.

Nuitka already enforces evaluation order for just about everything. But not for binary operators it seems.

- Providing an `# coding: no-exist` was crashing under Python2, and ignored under Python3, now it does the compatible thing for both.
- Global statements on the compiler level are legal in Python, and were not handled by Nuitka, they now are.

```
global a # Not in a function, but on module level. Pointless but legal!
a = 1
```

Effectively these statements can be ignored.

- Future imports are only legal when they are at the start of the file.

This was not enforced by Nuitka, making it accept code, which CPython would reject. It now properly raises a syntax error.

- Raising exceptions from context was leaking references.

```
raise ValueError() from None
```

Under CPython3.2 the above is not allowed (it is acceptable starting CPython3.3), and was also leaking references to its arguments.

- Importing the module that became `__main__` through the module name, didn't recurse to it.

This also gives a warning. PyBench does it, and then stumbles over the non-found "pybench" module. Of course, programmers should use `sys.modules["__main__"]` to access main module code. Not only because the duplicated modules don't share data.

- Compiled method `repr` leaked references when printed.

When printing them, they would not be freed, and subsequently hold references to the object (and class) they belong to. This could trigger bugs for code that expects `__del__` to run at some point.

- The `super` built-in leaked references to given object.

This was added, because Python3 needs it. It supplies the arguments to `super` automatically, whereas for Python2 the programmer had to do it. And now it turns out that the object lost a reference, causing similar issues as above, preventing `__del__` to run.

- The `raise` statement didn't enforce type of third argument.

This Python2-only form of exception raising now checks the type of the third argument before using it. Plus, when it's `None` (which is also legal), no reference to `None` is leaked.

- Python3 built-in exceptions were strings instead of exceptions.

A gross mistake that went uncaught by test suites. I wonder how. Them being strings doesn't help their usage of course, fixed.

- The `-nan` and `nan` both exist and make a difference.

A older story continued. There is a sign to `nan`, which can be copied away and should be present. This is now also supported by Nuitka.

- Wrong optimization of `a == a`, `a != a`, `a <= a` on C++ level.

While it's not done during Nuitka optimization, the rich comparison helpers still contained short cuts for `==`, `!=`, and `<=`.

- The `sys.executable` for `nuitka-python --python-version 3.2` was still `python`.

When determining the value for `sys.executable` the CPython library code looks at the name `exec` had received. It was `python` in all cases, but now it depends on the running version, so it propagates.

- Keyword only functions with default values were losing references to defaults.

```
def f(*, a=X()):
    pass

f()
f() # Can crash, X() should already be released.
```

This is now corrected. Of course, a Python3 only issue.

- Pressing CTRL-C didn't generate `KeyboardInterrupt` in compiled code.

Nuitka never executes "pending calls". It now does, with the upside, that the solution used, appears to be suitable for threading in Nuitka too. Expect more to come out of this.

- For `with` statements with `return`, `break`, or `continue` to leave their body, the `__exit__` was not called.

```
with a: # This called a.__enter__().
    return 2 # This didn't call a.__exit__(None, None, None).
```

This is of course quite huge, and unfortunately wasn't covered by any test suite so far. Turns out, the re-formulation of `with` statements, was wrongly using `try/except/else`, but these ignore the problematic statements. Only `try/finally` does. The enhanced re-formulation now does the correct thing.

- Starting with Python3, absolute imports are now the default.

This was already present for Python3.3, and it turns out that all of Python3 does it.

Optimization

- Constants are now much less often created with `pickle` module, but created directly.

This esp. applies for nested constants, now more values become `is` identical instead of only `==` identical, which indicates a reduced memory usage.

```
a = ("something_special",)
b = "something_special"

assert a[0] is b # Now true
```

This is not only about memory efficiency, but also about performance. Less memory usage is more cache friendly, and the `"=="` operator will be able to shortcut dramatically in cases of identical objects.

Constants now created without `pickle` usage, cover `float`, `list`, and `dict`, which is enough for PyStone to not use it at all, which has been added support for as well.

- Continue statements might be optimized away.

A terminal `continue` in a loop, was not optimized away:

```
while 1:
    something
    continue # Now optimized away
```

The trailing `continue` has no effect and can therefore be removed.

```
while 1:
    something
```

- Loops with only `break` statements are optimized away.

```
while 1:
    break
```

A loop immediately broken has of course no effect. Loop conditions are re-formulated to immediate `"if ... : break"` checks. Effectively this means that loops with conditions detected to be always false to see the loop entirely removed.

New Tests

- Added tests for the found issues.

- Running the programs test suite (i.e. recursion) for Python3.2 and Python3.2 as well, after making adaptation so that the absolute import changes are now covered.
- Running the "CPython3.2" test suite with Python3.3 based Nuitka works and found a few minor issues.

Organisational

- The [Downloads](#) page now offers RPMs for RHEL6, CentOS6, F17, F18, and openSUSE 12.1, 12.2, 12.3. This large coverage is thanks to openSUSE build service and "ownssh" for contributing an RPM spec file.

The page got improved with logos for the distributions.

- Added "ownssh" as contributor.
- Revamped the [User Manual](#) in terms of layout, structure, and content.

Summary

This release is the result of much validation work. The amount of fixes the largest of any release so far. New platforms, basic Python3.3 support, consolidation all around.

Nuitka Release 0.4.1

This release is the first follow-up with a focus on optimization. The major highlight is progress towards SSA form in the node tree.

Also a lot of cleanups have been performed, for both the tree building, which is now considered mostly finished, and will be only reviewed. And for the optimization part there have been large amounts of changes.

New Features

- Python 3.3 experimental support
 - Now compiles many basic tests. Ported the dictionary quick access and update code to a more generic and useful interface.
 - Added support for `__qualname__` to classes and functions.
 - Small compatibility changes. Some exceptions changed, absolute imports are now default, etc.
 - For comparison tests, the hash randomization is disabled.
- Python 3.2 support has been expanded.

The Python 3.2 on Ubuntu is not providing a helper function that was used by Nuitka, replaced it with out own code.

Bug Fixes

- Default values were not "is" identical.

```
def defaultKeepsIdentity(arg="str_value"):
    print arg is "str_value"

defaultKeepsIdentity()
```

This now prints "True" as it does with CPython. The solution is actually a general code optimization, see below.

- Usage of `unicode` built-in with more than one argument could corrupt the encoding argument string. An implementation error of the `unicode` was releasing references to arguments converted to default encoding, which could corrupt it.
- Assigning Python3 function annotations could cause a segmentation fault.

Optimization

- Improved propagation of exception raise statements, eliminating more code. They are now also propagated from all kinds of expressions. Previously this was more limited. An assertion added will make sure that all raises are propagated. Also finally, raise expressions are converted into raise statements, but without any normalization.

```
# Now optimizing:
raise TypeError, 1 / 0
# into (minus normalization):
raise ZeroDivisionError, "integer division or modulo by zero"

# Now optimizing:
(1 / 0).something
# into (minus normalization):
raise ZeroDivisionError, "integer division or modulo by zero"

# Now optimizing:
function(a, 1 / 0).something
# into (minus normalization), notice the side effects of first checking
# function and a as names to be defined, these may be removed only if
# they can be demonstrated to have no effect.
function
a
raise ZeroDivisionError, "integer division or modulo by zero"
```

There is more examples, where the raise propagation is new, but you get the idea.

- Conditional expression nodes are now optimized according to the truth value of the condition, and not only for compile time constants. This covers e.g. container creations, and other things.

```
# This was already optimized, as it's a compile time constant.
a if ("a",) else b
a if True else b

# These are now optimized, as their truth value is known.
a if (c,) else b
a if not (c,) else b
```

This is simply taking advantage of infrastructure that now exists. Each node kind can overload "getTruthValue" and benefit from it. Help would be welcome to review which ones can be added.

- Function creations only have side effects, when their defaults or annotations (Python3) do. This allows to remove them entirely, should they be found to be unused.
- Code generation for constants now shares element values used in tuples.

The general case is currently too complex to solve, but we now make sure constant tuples (as e.g. used in the default value for the compiled function), and string constants share the value. This should reduce memory usage and speed up program start-up.

Cleanups

- Optimization was initially designed around visitors that each did one thing, and did it well. It turns out though, that this approach is unnecessary, and constraint collection, allows for the most consistent results. All remaining optimization has been merged into constraint collection.
- The names of modules containing node classes were harmonized to always be plural. In the beginning, this was used to convey the information that only a single node kind would be contained, but that has long changed, and is unimportant information.
- The class names of nodes were stripped from the "CPython" prefix. Originally the intent was to express strict correlation to CPython, but with increasing amounts of re-formulations, this was not used at all, and it's also not important enough to dominate the class name.
- The re-formulations performed in tree building have moved out of the "Building" module, into names "ReformulationClasses" e.g., so they are easier to locate and review. Helpers for node building are now in a separate module, and generally it's much easier to find the content of interest now.
- Added new re-formulation of `print` statements. The conversion to strings is now made explicit in the node tree.

New Tests

- Added test to cover default value identity.

Organisational

- The upload of [Nuitka to PyPI](#) has been repaired and now properly displays project information again.

Summary

The quicker release is mostly a consolidation effort, without much actual performance progress. The progress towards SSA form matter a lot on the outlook front. Once this is finished, standard compiler algorithms can be added to Nuitka which go beyond the current peephole optimization.

Nuitka Release 0.4.0

This release brings massive progress on all fronts. The big highlight is of course: Full Python3.2 support. With this release, the test suite of CPython3.2 is considered passing when compiled with Nuitka.

Then lots of work on optimization and infrastructure. The major goal of this release was to get in shape for actual optimization. This is also why for the first time, it is tested that some things are indeed compile time optimized to spot regressions easier. And we are having performance diagrams, [even if weak ones](#):

New Features

- Python3.2 is now fully supported.
- Fully correct `metaclass` = semantics now correctly supported. It had been working somewhat previously, but now all the corner cases are covered too.
 - Keyword only parameters.
 - Annotations of functions return value and their arguments.

- Exception causes, chaining, automatic deletion of exception handlers as values.
- Added support for starred assigns.
- Unicode variable names are also supported, although it's of course ugly, to find a way to translate these to C++ ones.

Bug Fixes

- Checking compiled code with `instance(some_function, types.FunctionType)` as "zope.interfaces" does, was causing compatibility problems. Now this kind of check passes for compiled functions too.
- The frame of modules had an empty locals dictionary, which is not compatible to CPython which puts the globals dictionary there too.
- For nested exceptions and interactions with generator objects, the exceptions in `sys.exc_info()` were not always fully compatible. They now are.
- The `range` builtin was not raising exceptions if given arguments appeared to not have side effects, but were still illegal, e.g. `range([], 1, -1)` was optimized away if the value was not used.
- Don't crash on imported modules with syntax errors. Instead, the attempted recursion is simply not done.
- Doing a `del` on `__defaults` and `__module__` of compiled functions was crashing. This was noticed by a Python3 test for `__kwdefaults__` that exposed this compiled functions weakness.
- Wasn't detecting duplicate arguments, if one of them was not a plain arguments. Star arguments could collide with normal ones.
- The `__doc__` of classes is now only set, where it was in fact specified. Otherwise it only polluted the name space of `locals()`.
- When `return` from the tried statements of a `try/finally` block, was overridden, by the final block, a reference was leaked. Example code:

```
try:
    return 1
finally:
    return 2
```

- Raising exception instances with value, was leaking references, and not raising the `TypeError` error it is supposed to do.
- When raising with multiple arguments, the evaluation order of them was not enforced, it now is. This fixes a reference leak when raising exceptions, where building the exception was raising an exception.

Optimization

- Optimizing attribute access to compile time constants for the first time. The old registry had no actual user yet.
- Optimizing subscript and slices for all compile time constants beyond constant values, made easy by using inheritance.
- Built-in references now convert to strings directly, e.g. when used in a print statement. Needed for the testing approach "compiled file contains only prints with constant value".
- Optimizing calls to constant nodes directly into exceptions.

- Optimizing built-in `bool` for arguments with known truth value. This would be creations of tuples, lists, and dictionaries.
- Optimizing `a is b` and `a is not b` based on aliasing interface, which at this time effectively is limited to telling that `a is a` is true and `a is not a` is false, but this will expand.
- Added support for optimizing `hasattr`, `getattr`, and `setattr` built-ins as well. The `hasattr` was needed for the `class` re-formulation of Python3 anyway.
- Optimizing `getattr` with string argument and no default to simple attribute access.
- Added support for optimizing `isinstance` built-in.
- Was handling "BreakException" and "ContinueException" in all loops that used `break` or `continue` instead of only where necessary.
- When catching "ReturnValueException", was raising an exception where a normal return was sufficient. Raising them now only where needed, which also means, function need not catch them ever.

Cleanups

- The handling of classes for Python2 and Python3 have been re-formulated in Python more completely.
 - The calling of the determined "metaclass" is now in the node tree, so this call may possible to in-line in the future. This eliminated some static C++ code.
 - Passing of values into dictionary creation function is no longer using hard coded special parameters, but temporary variables can now have closure references, making this normal and visible to the optimization.
 - Class dictionary creation functions are therefore no longer as special as they used to be.
 - There is no class creation node anymore, it's merely a call to `type` or the metaclass detected.
- Re-formulated complex calls through helper functions that process the star list and dict arguments and do merges, checks, etc.
 - Moves much C++ code into the node tree visibility.
 - Will allow optimization to eliminate checks and to compile time merge, once in-line functions and loop unrolling are supported.
- Added "return None" to function bodies without a an aborting statement at the end, and removed the hard coded fallback from function templates. Makes it explicit in the node tree and available for optimization.
- Merged C++ classes for frame exception keeper with frame guards.
 - The exception is now saved in the compiled frame object, making it potentially more compatible to start with.
 - Aligned module and function frame guard usage, now using the same class.
 - There is now a clear difference in the frame guard classes. One is for generators and one is for functions, allowing to implement their different exception behavior there.
- The optimization registries for calls, subscripts, slices, and attributes have been replaced with attaching them to nodes.
 - The ensuing circular dependency has been resolved by more local imports for created nodes.
 - The package "nuitka.transform.optimization.registries" is no more.
 - New per node methods "computeNodeCall", "computeNodeSubscript", etc. dispatch the optimization process to the nodes directly.

- Use the standard frame guard code generation for modules too.
 - Added a variant "once", that avoids caching of frames entirely.
- The variable closure taking has been cleaned up.
 - Stages are now properly numbered.
 - Python3 only stage is not executed for Python2 anymore.
 - Added comments explaining things a bit better.
 - Now an early step done directly after building a tree.
- The special code generation used for unpacking from iterators and catching "StopIteration" was cleaned up.
 - Now uses template, Generator functions, and proper identifiers.
- The `return` statements in generators are now re-formulated into `raise StopIteration` for generators, because that's what they really are. Allowed to remove special handling of `return` nodes in generators.
- The specialty of CPython2.6 yielding non-None values of lambda generators, was so far implemented in code generation. This was moved to tree building as a re-formulation, making it subject to normal optimization.
- Mangling of attribute names in functions contained in classes, has been moved into the early tree building. So far it was done during code generation, making it invisible to the optimization stages.
- Removed tags attribute from node classes. This was once intended to make up for non-inheritance of similar node kinds, but since we have function references, the structure got so clean, it's no more needed.
- Introduced new package `nuitka.tree`, where the building of node trees, and operations on them live, as well as recursion and variable closure.
- Removed `nuitka.transform` and move its former children `nuitka.optimization` and `nuitka.finalization` one level up. The deeply nested structure turned out to have no advantage.
- Checks for Python version was sometimes "> 300", where of course ">= 300" is the only thing that makes sense.
- Split out helper code for exception raising from the handling of exception objects.

New Tests

- The complete CPython3.2 test suite was adapted (no `__code__`, no `__closure__`, etc.) and is now passing, but only without "--debug", because otherwise some of the generated C++ triggers (harmless) warnings.
- Added new test suite designed to prove that expressions that are known to be compile time constant are indeed so. This works using the XML output done with `--dump-xml` and then searching it to only have print statements with constant values.
- Added new basic CPython3.2 test "Functions32" and "ParameterErrors32" to cover keyword only parameter handling.
- Added tests to cover generator object and exception interactions.
- Added tests to cover `try/finally` and `return` in one or both branches correctly handling the references.
- Added tests to cover evaluation order of arguments when raising exceptions.

Organisational

- Changed my email from GMX over to Gmail, the old one will still continue to work. Updated the copyright notices accordingly.
- Uploaded [Nuitka to PyPI](#) as well.

Summary

This release marks a milestone. The support of Python3 is here. The re-formulation of complex calls, and the code generation improvements are quite huge. More re-formulation could be done for argument parsing, but generally this is now mostly complete.

The 0.3.x series had a lot releases. Many of which brought progress with re-formulations that aimed at making optimization easier or possible. Sometimes small things like making "return None" explicit. Sometimes bigger things, like making class creations normal functions, or getting rid of `or` and `and`. All of this was important ground work, to make sure, that optimization doesn't deal with complex stuff.

So, the 0.4.x series begins with this. The focus from now on can be almost purely optimization. This release contains already some of it, with frames being optimized away, with the assignment keepers from the `or` and `and` re-formulation being optimized away. This will be about achieving goals from the "ctypes" plan as discussed in the Developer Manual.

Also the performance page will be expanded with more benchmarks and diagrams as I go forward. I have finally given up on "codespeed", and do my own diagrams.

Nuitka Release 0.3.25

This release brings about changes on all fronts, bug fixes, new features. Also very importantly Nuitka no longer uses C++11 for its code, but mere C++03. There is new re-formulation work, and re-factoring of functions.

But the most important part is this: Mercurial unit tests are working. Nearly. With the usual disclaimer of me being wrong, all remaining errors are errors of the test, or minor things. Hope is that these unit tests can be added as release tests to Nuitka. And once that is done, the next big Python application can come.

Bug Fixes

- Local variables were released when an exception was raised that escaped the local function. They should only be released, after another exception was raised somewhere.
- Identifiers of nested tuples and lists could collide.

```
a = ((1, 2), 3)
b = ((1,), 2, 3)
```

Both tuples had the same name previously, not the end of the tuple is marked too. Fixed in 0.3.24.1 already.

- The `__name__` when used read-only in modules in packages was optimized to a string value that didn't contain the package name.
- Exceptions set when entering compiled functions were unset at function exit.

New Features

- Compiled frames support. Before, Nuitka was creating frames with the standard CPython C/API functions, and tried its best to cache them. This involved some difficulties, but as it turns out, it is actually possible to instead provide a compatible type of our own, that we have full control over.

This will become the base of enhanced compatibility. Keeping references to local variables attached to exception tracebacks is something we may be able to solve now.

- Enhanced Python3 support, added support for `nonlocal` declarations and many small corrections for it.
- Writable `__defaults__` attribute for compiled functions, actually changes the default value used at call time. Not supported is changing the amount of default parameters.

Cleanups

- Keep the functions along with the module and added "FunctionRef" node kind to point to them.
- Reformulated `or` and `and` operators with the conditional expression construct which makes the "short-circuit" branch.
- Access `self` in methods from the compiled function object instead of pointer to context object, making it possible to access the function object.
- Removed "OverflowCheck" module and its usage, avoids one useless scan per function to determine the need for "locals dictionary".
- Make "compileTree" of "MainControl" module to only do what the name says and moved the rest out, making the top level control clearer.
- Don't export module entry points when building executable and not modules. These exports cause MinGW and MSVC compilers to create export libraries.

Optimization

- More efficient code for conditional expressions in conditions:

```
if a if b else c:
    ...
```

See above, this code is now the typical pattern for each `or` and `and`, so this was much needed now.

Organisational

- The remaining uses of C++11 have been removed. Code generated with Nuitka and complementary C++ code now compile with standard C++03 compilers. This lowers the Nuitka requirements and enables at least g++ 4.4 to work with Nuitka.
- The usages of the GNU extension operation `a ? : b` have replaced with standard C++ constructs. This is needed to support MSVC which doesn't have this.
- Added examples for the typical use cases to the [User Manual](#).
- The "compare_with_cpython" script has gained an option to immediately remove the Nuitka outputs (build directory and binary) if successful. Also the temporary files are now put under `/var/tmp` if available.
- Debian package improvements, registering with `doc-base` the [User Manual](#) so it is easier to discover. Also suggest `mingw32` package which provides the cross compiler to Windows.

- Partial support for MSVC (Visual Studio 2008 to be exact, the version that works with CPython2.6 and CPython2.7).

All basic tests that do not use generators are working now, but those will currently cause crashes.

- Renamed the `--g++-only` option to `--c++-only`.

The old name is no longer correct after clang and MSVC have gained support, and it could be misunderstood to influence compiler selection, rather than causing the C++ source code to not be updated, so manual changes will be used.

- Catch exceptions for `continue`, `break`, and `return` only where needed for `try/finally` and loop constructs.

New Tests

- Added CPython3.2 test suite as "tests/CPython32" from 3.2.3 and run it with CPython2.7 to check that Nuitka gives compatible error messages. It is not expected to pass yet on Python3.2, but work will be done towards this goal.
- Make CPython2.7 test suite runner also execute the generated "doctest" modules.
- Enabled tests for default parameters and their reference counts.

Summary

This release marks an important point. The compiled frames are exciting new technology, that will allow even better integration with CPython, while improving speed. Lowering the requirements to C++03 means, we will become usable on Android and with MSVC, which will make adoption of Nuitka on Windows easier for many.

Structurally the outstanding part is the function as references cleanup. This was a blocker for value propagation, because now functions references can be copied, whereas previously this was duplicating the whole function body, which didn't work, and wasn't acceptable. Now, work can resume in this domain.

Also very exciting when it comes to optimization is the remove of special code for `or` and `and` operators, as these are now only mere conditional expressions. Again, this will make value propagation easier with two special cases less.

And then of course, with Mercurial unit tests running compiled with Nuitka, an important milestone has been hit.

For a while now, the focus will be on completing Python3 support, XML based optimization regression tests, benchmarks, and other open ends. Once that is done, and more certainty about Mercurial tests support, I may call it a 0.4 and start with local type inference for actual speed gains.

Nuitka Release 0.3.24

This release contains progress on many fronts, except performance.

The extended coverage from running the CPython 2.7 and CPython 3.2 (partially) test suites shows in a couple of bug fixes and general improvements in compatibility.

Then there is a promised new feature that allows to compile whole packages.

Also there is more Python3 compatibility, the CPython 3.2 test suite now succeeds up to "test_builtin.py", where it finds that `str` doesn't support the new parameters it has gained, future releases will improve on this.

And then of course, more re-formulation work, in this case, class definitions are now mere simple functions. This and later function references, is the important and only progress towards type inference.

Bug Fixes

- The compiled method type can now be used with `copy` module. That means, instances with methods can now be copied too. Fixed in 0.3.23.1 already.
- The `assert` statement as of Python2.7 creates the `AssertionError` object from a given value immediately, instead of delayed as it was with Python2.6. This makes a difference for the form with 2 arguments, and if the value is a tuple. Fixed in 0.3.23.1 already.
- Sets written like this didn't work unless they were predicted at compile time:

```
{value}
```

This apparently rarely used Python2.7 syntax didn't have code generation yet and crashed the compiler. Fixed in 0.3.23.1 already.

- For Python2, the default encoding for source files is `ascii`, and it is now enforced by Nuitka as well, with the same `SyntaxError`.
- Corner cases of `exec` statements with nested functions now give proper `SyntaxError` exceptions under Python2.
- The `exec` statement with a tuple of length 1 as argument, now also gives a `TypeError` exception under Python2.
- For Python2, the `del` of a closure variable is a `SyntaxError`.

New Features

- Added support creating compiled packages. If you give Nuitka a directory with an `__init__.py` file, it will compile that package into a `.so` file. Adding the package contents with `--recurse-dir` allows to compile complete packages now. Later there will be a cleaner interface likely, where the later is automatic.
- Added support for providing directories as main programs. It's OK if they contain a `__main__.py` file, then it's used instead, otherwise give compatible error message.
- Added support for optimizing the `super` built-in. It was already working correctly, but not optimized on CPython2. But for CPython3, the variant without any arguments required dedicated code.
- Added support for optimizing the `unicode` built-in under Python2. It was already working, but will become the basis for the `str` built-in of Python3 in future releases.
- For Python3, lots of compatibility work has been done. The Unicode issues appear to be ironed out now. The `del` of closure variables is allowed and supported now. Built-ins like `ord` and `chr` work more correctly and attributes are now interned strings, so that monkey patching classes works.

Organisational

- Migrated `bin/benchmark.sh` to Python as `misc/run-valgrind.py` and made it a bit more portable that way. Prefers `/var/tmp` if it exists and creates temporary files in a secure manner. Triggered by the Debian "insecure temp file" bug.
- Migrated `bin/make-dependency-graph.sh` to Python as `misc/make-dependency-graph.py` and made a more portable and powerful that way.

The filtering is done a more robust way. Also it creates temporary files in a secure manner, also triggered by the Debian "insecure temp file" bug.

And it creates SVG files and no longer PostScript as the first one is more easily rendered these days.

- Removed the `misc/gist` git sub-module, which was previously used by `misc/make-doc.py` to generate HTML from [User Manual](#) and [Developer Manual](#).

These are now done with Nikola, which is much better at it and it integrates with the web site.

- Lots of formatting improvements to the change log, and manuals:
 - Marking identifiers with better suited ReStructured Text markup.
 - Added links to the bug tracker all Issues.
 - Unified wordings, quotation, across the documents.

Cleanups

- The creation of the class dictionaries is now done with normal function bodies, that only needed to learn how to throw an exception when directly called, instead of returning `NULL`.

Also the assignment of `__module__` and `__doc__` in these has become visible in the node tree, allowing their proper optimization.

These re-formulation changes allowed to remove all sorts of special treatment of `class` code in the code generation phase, making things a lot simpler.

- There was still a declaration of `PRINT_ITEMS` and uses of it, but no definition of it.
- Code generation for "main" module and "other" modules are now merged, and no longer special.
- The use of raw strings was found unnecessary and potentially still buggy and has been removed. The dependence on C++11 is getting less and less.

New Tests

- Updated CPython2.6 test suite "tests/CPython26" to 2.6.8, adding tests for recent bug fixes in CPython. No changes to Nuitka were needed in order to pass, which is always good news.
- Added CPython2.7 test suite as "tests/CPython27" from 2.7.3, making it public for the first time. Previously a private copy of some age, with many no longer needed changes had been used by me. Now it is up to par with what was done before for "tests/CPython26", so this pending action is finally done.
- Added test to cover Python2 syntax error of having a function with closure variables nested inside a function that is an overflow function.
- Added test "BuiltinSuper" to cover `super` usage details.
- Added test to cover `del` on nested scope as syntax error.
- Added test to cover `exec` with a tuple argument of length 1.
- Added test to cover `barry_as_FLUFL` future import to work.
- Removed "Unicode" from known error cases for CPython3.2, it's now working.

Summary

This release brought forward the most important remaining re-formulation changes needed for Nuitka. Removing class bodies, makes optimization yet again simpler. Still, making function references, so they can be copied, is missing for value propagation to progress.

Generally, as usual, a focus has been laid on correctness. This is also the first time, I am release with a known bug though, one which I believe now, may be the root cause of the mercurial tests not yet passing.

The solution will be involved and take a bit of time. It will be about "compiled frames" and be a (invasive) solution. It likely will make Nuitka faster too. But this release includes lots of tiny improvements, for Python3 and also for Python2. So I wanted to get this out now.

As usual, please check it out, and let me know how you fare.

Nuitka Release 0.3.23

This release is the one that completes the Nuitka "sun rise phase".

All of Nuitka is now released under [Apache License 2.0](#) which is a very liberal license, and compatible with basically all Free Software licenses there are. It's only asking to allow integration, of what you send back, and patent grants for the code.

In the first phase of Nuitka development, I wanted to keep control over Nuitka, so it wouldn't repeat mistakes of other projects. This is no longer a concern for me, it's not going to happen anymore.

I would like to thank Debian Legal team, for originally bringing to my attention, that this license will be better suited, than any copyright assignment could be.

Bug Fixes

- The compiled functions could not be used with `multiprocessing` or `copy.copy`. Fixed in 0.3.22.1 already.
- In-place operations for slices with not both bounds specified crashed the compiler. Fixed in 0.3.22.1 already.
- Cyclic imports could trigger an endless loop, because module import expressions became the parent of the imported module object. Fixed in 0.3.22.2 already.
- Modules named `proc` or `func` could not be compiled to modules or embedded due to a collision with identifiers of CPython2.7 includes. Fixed in 0.3.22.2 already.

New Features

- The function copying fix also makes pickling of compiled functions available. As it is the case for non-compiled functions in CPython, no code objects are stored, only names of module level variables.

Organisational

- Using the Apache License 2.0 for all of Nuitka now.
- Speedcenter has been re-activated, but is not yet having a lot of benchmarks yet, subject to change.

Update

We have given up on this version of speedcenter meanwhile, and generate static pages with graphs instead. We can this still speedcenter.

New Tests

- Changed the "CPython26" tests to no longer disable the parts that relied on copying of functions to work as that is now supported.
- Extended in-place assignment tests to cover error cases of we had issues with.
- Extended compile library test to also try and compile the path where `numpy` lives. This is apparently another path, where Debian installs some modules, and compiling this would have revealed issues sooner.

Summary

The release contains bug fixes, and the huge step of changing [the license](#). It is made in preparation to PyCON EU.

Nuitka Release 0.3.22

This release is a continuation of the trend of previous releases, and added more re-formulations of Python that lower the burden on code generation and optimization.

It also improves Python3 support substantially. In fact this is the first release to not only run itself under Python3, but for Nuitka to *compile itself* with Nuitka under Python3, which previously only worked for Python2. For the common language subset, it's quite fine now.

Bug Fixes

- List contractions produced extra entries on the call stack, after they became functions, these are no more existent. That was made possible by making frame stack entries an optional element in the node tree, left out for list contractions.
- Calling a compiled function in an exception handler cleared the exception on return, it no longer does that.
- Reference counter handling with generator `throw` method is now correct.
- A module "builtins" conflicted with the handling of the Python `builtins` module. Those now use different identifiers.

New Features

- New `metaclass` syntax for the `class` statement works, and the old `__metaclass__` attribute is properly ignored.

```
# Metaclass syntax in Python3, illegal in Python2
class X(metaclass=Y):
    pass
```

```
# Metaclass syntax in Python2, no effect in Python3
class X:
    __metaclass__ = Y
```

Note

The way to make a use of a metaclass in a portable way, is to create a based class that has it and then inherit from it. Sad, isn't it. Surely, the support for `__metaclass__` could still live.

```
# For Python2/3 compatible source, we create a base class that has the
# metaclass used and doesn't require making a choice.

CPythonNodeMetaClassBase = NodeCheckMetaClass("CPythonNodeMetaClassBase", (object,), {})
```

- The `--dump-xml` option works with Nuitka running under Python3. This was not previously supported.
- Python3 now also has compatible parameter errors and compatible exception error messages.
- Python3 has changed scope rules for list contractions (assignments don't affect outside values) and this is now respected as well.
- Python3 has gained support for recursive programs and stand alone extension modules, these are now both possible as well.

Optimization

- Avoid frame stack entries for functions that cannot raise exceptions, i.e. where they would not be used.

This avoids overhead for the very simple functions. And example of this can be seen here:

```
def simple():
    return 7
```

- Optimize `len` built-in for non-constant, but known length values.

An example can be seen here:

```
# The range isn't constructed at compile time, but we still know its
# length.
len(range(100000000))

# The string isn't constructed at compile time, but we still know its
# length.
len(" *" * 1000)

# The tuple isn't constructed, instead it's known length is used, and
# side effects are maintained.
len((a(), b()))
```

This new optimization applies to all kinds of container creations and the `range` built-in initially.

- Optimize conditions for non-constant, but known truth values.

At this time, known truth values of non-constants means `range` built-in calls with known size and container creations.

An example can be seen here:

```
if (a,):
    print "In Branch"
```

It's clear, that the tuple will be true, we just need to maintain the side effect, which we do.

- Optimize `or` and `and` operators for known truth values.

See above for what has known truth values currently. This will be most useful to predict conditions that need not be evaluated at all due to short circuit nature, and to avoid checking against constant values. Previously this could not be optimized, but now it can:

```
# The access and call to "something()" cannot possibly happen
0 and something()
```

```

# Can be replaced with "something()", as "1" is true. If it had a side
# effect, it would be maintained.
1 and something()

# The access and call to "something()" cannot possibly happen, the value
# is already decided, it's "1".
1 or something()

# Can be replaced with "something()", as "0" is false. If it had a side
# effect, it would be maintained.
0 or something()

```

- Optimize print arguments to become strings.

The arguments to print statements are now converted to strings at compile time if possible.

```
print 1
```

becomes:

```
print "1"
```

- Combine print arguments to single ones.

When multiple strings are printed, these are now combined.

```
print "1+1=", 1 + 1
```

becomes:

```
print "1+1= 2"
```

Organisational

- Enhanced Python3 support, enabling support for most basic tests.
- Check files with PyLint in deterministic (alphabetical) order.

Cleanups

- Frame stack entries are now part of the node tree instead of part of the template for every function, generator, class or module.
- The try/except/else has been re-formulated to use an indicator variable visible in the node tree, that tells if a handler has been executed or not.
- Side effects are now a dedicated node, used in several optimization to maintain the effect of an expression with known value.

New Tests

- Expanded and adapted basic tests to work for Python3 as well.
- Added reference count tests for generator functions throw, send, and close methods.

- Cover calling a function with `try/except` in an exception handler twice. No test was previously doing that.

Summary

This release offers enhanced compatibility with Python3, as well as the solution to many structural problems. Calculating lengths of large non-constant values at compile time, is technically a break through, as is avoiding lengthy calculations. The frame guards as nodes is a huge improvement, making that costly operational possible to be optimized away.

There still is more work ahead, before value propagation will be safe enough to enable, but we are seeing the glimpse of it already. Not for long, and looking at numbers will make sense.

Nuitka Release 0.3.21

This releases contains some really major enhancements, all heading towards enabling value propagation inside Nuitka. Assignments of all forms are now all simple and explicit, and as a result, now it will be easy to start tracking them.

Contractions have become functions internally, with statements use temporary variables, complex unpacking statement were reduced to more simple ones, etc.

Also there are the usual few small bug fixes, and a bunch of organisational improvements, that make the release complete.

Bug Fixes

- The built-in `next` could causes a program crash when iterating past the end of an iterator. Fixed in 0.3.20.1 already.
- The `set` constants could cause a compiler error, as that type was not considered in the "mutable" check yet. Fixed in 0.3.20.2 already.
- Performance regression. Optimize expression for exception types caught as well again, this was lost in last release.
- Functions that contain `exec`, are supposed to have a writable locals. But when removing that `exec` statement as part of optimization, this property of the function could get lost.
- The so called "overflow functions" are once again correctly handled. These once were left behind in some refactoring and had not been repaired until now. An overflow function is a nested function with an `exec` or a star import.
- The syntax error for `return` outside of a function, was not given, instead the code returned at run time. Fixed to raise a `SyntaxError` at compile time.

Optimization

- Avoid `tuple` objects to be created when catching multiple exception types, instead call `exception match` check function multiple times.

- Removal of dead code following `break`, `continue`, `return`, and `raise`. Code that follows these statements, or conditional statements, where all branches end with it.

Note

These may not actually occur often in actual code, but future optimization may produce them more frequently, and their removal may in turn make other possible optimization.

- Detect module variables as "read only" after all writes have been detected to not be executed as removed. Previously the "read only indicator" was determined only once and then stayed the same.
- Expanded conditional statement optimization to detect cases, where condition is a compile time constant, not just a constant value.
- Optimize away assignments from a variable to the same variable, they have no effect. The potential side effect of accessing the variable is left intact though, so exceptions will be raised still.

Note

An exception is where `len = len` actually does have an impact, because that variable becomes assignable. The "compile itself" test of Nuitka found that to happen with `long` from the `nuitka.__past__` module.

- Created Python3 variant of quick `unicode` string access, there was no such thing in the CPython C/API, but we make the distinction in the source code, so it makes sense to have it.
- Created an optimized implementation for the built-in `iter` with 2 parameters as well. This allows for slightly more efficient code to be created with regards to reference handling, rather than using the CPython C/API.
- For all types of variable assigned in the generated code, there are now methods that accept already taken references or not, and the code generator picks the optimal variant. This avoids the drop of references, that e.g. the local variable will insist to take.
- Don't use a "context" object for generator functions (and generator expressions) that don't need one. And even if it does to store e.g. the given parameter values, avoid to have a "common context" if there is no closure taken. This avoids useless `malloc` calls and speeds up repeated generator object creation.

Organisational

- Changed the Scons build file database to reside in the build directory as opposed to the current directory, not polluting it anymore. Thanks for the patch go to Michael H Kent, very much appreciated.
- The `--experimental` option is no longer available outside of checkouts of git, and even there not on stable branches (`master`, `hotfix/...`). It only pollutes `--help` output as stable releases have no experimental code options, not even development version will make a difference.
- The binary "bin/Nuitka.py" has been removed from the git repository. It was deprecated a while ago, not part of the distribution and served no good use, as it was a symbolic link only anyway.
- The `--python-version` option is applied at Nuitka start time to re-launch Nuitka with the given Python version, to make sure that the Python run time used for computations and link time Python

versions are the same. The allowed values are now checked (2.6, 2.7 and 3.2) and the user gets a nice error with wrong values.

- Added `--keep-pythonpath` alias for `--execute-with-pythonpath` option, probably easier to remember.
- Support `--debug` with clang, so it can also be used to check the generated code for all warnings, and perform assertions. Didn't report anything new.
- The contents environment variable `CXX` determines the default C++ compiler when set, so that checking with `CXX=g++-4.7 nuitka-python ...` has become supported.
- The `check-with-pylint` script now has a real command line option to control the display of `TODO` items.

Cleanups

- Changed complex assignments, i.e. assignments with multiple targets to such using a temporary variable and multiple simple assignments instead.

```
a = b = c
```

```
_tmp = c
b = _tmp
a = _tmp
```

In CPython, when one assignment raises an exception, the whole thing is aborted, so the complexity of having multiple targets is no more needed, now that we have temporary variables in a block.

All that was really needed, was to evaluate the complete source expression only once, but that made code generation contain ugly loops that are no more needed.

- Changed unpacking assignments to use temporary variables. Code like this:

```
a, b = c
```

Is handled more like this:

```
_tmp_iter = iter(c)
_tmp1 = next(_tmp_iter)
_tmp2 = next(_tmp_iter)
if not finished(_tmp_iter):
    raise ValueError("too many values to unpack")
a = _tmp1
b = _tmp2
```

In reality, not really `next` is used, as it wouldn't raise the correct exception for unpacking, and the `finished` check is more condensed into it.

Generally this cleanup allowed that the `AssignTargetTuple` and associated code generation was removed, and in the future value propagation may optimize these `next` and `iter` calls away where possible. At this time, this is not done yet.

- Exception handlers assign caught exception value through assignment statement.

Previously the code generated for assigning from the caught exception was not considered part of the handler. It now is the first statement of an exception handler or not present, this way it may be optimized as well.

- Exception handlers now explicitly catch more than one type.

Catching multiple types worked by merits of the created tuple object working with the Python C/API function called, but that was not explicit at all. Now every handler has a tuple of exceptions it catches, which may only be one, or if None, it's all.

- Contractions are now functions as well.

Contractions (list, dict, and set) are now re-formulated as function bodies that contain for loops and conditional statements. This allowed to remove a lot of special code that dealt with them and will make these easier to understand for optimization and value propagation.

- Global is handled during tree building.

Previously the global statement was its own node, which got removed during the optimization phase in a dedicated early optimization that applied its effect, and then removed the node.

It was determined, that there is no reason to not immediately apply the effect of the global variable and take closure variables and add them to the provider of that `global` statement, allowing to remove the node class.

- Read only module variable detection integrated to constraint collection.

The detection of read only module variables was so far done as a separate step, which is no more necessary as the constraint collection tracks the usages of module variables anyway, so this separate and slow step could be removed.

New Tests

- Added test to cover order of calls for complex assignments that unpack, to see that they make a fresh iterator for each part of a complex assignment.
- Added test that unpacks in an exception catch. It worked, due to the generic handling of assignment targets by Nuitka, and I didn't even know it can be done, example:

```
try:
    raise ValueError(1, 2)
except ValueError as (a, b):
    print "Unpacking caught exception and unpacked", a, b
```

Will assign a=1 and b=2.

- Added test to cover return statements on module level and class level, they both must give syntax errors.
- Cover exceptions from accessing unassigned global names.
- Added syntax test to show that star imports do not allow other names to be imported at the same time as well.
- Python3 is now also running the compile itself test successfully.

Summary

The progress made towards value propagation and type inference is very significant, and makes those appears as if they are achievable.

Nuitka Release 0.3.20

This time there are a few bug fixes and some really major cleanups, lots of new optimization and preparations for more. And then there is a new compiler clang and a new platform supported. macOS X appears to work mostly, thanks for the patches from Pete Hunt.

Bug Fixes

- The use of a local variable name as an expression was not covered and lead to a compiler crash. Totally amazing, but true, nothing in the test suite of CPython covered this. Fixed in release 0.3.19.1 already.
- The use of a closure variable name as an expression was not covered as well. And in this case corrupted the reference count. Fixed in release 0.3.19.1 already.
- The `from x import *` attempted to respect `__all__` but failed to do so. Fixed in release 0.3.19.2 already.
- The `from x import *` didn't give a `SyntaxError` when used on Python3. Fixed in release 0.3.19.2 already.
- The syntax error messages for "global for function argument name" and "duplicate function argument name" are now identical as well.
- Parameter values of generator function could cause compilation errors when used in the closure of list contractions. Fixed.

New Features

- Added support for disabling the console for Windows binaries. Thanks for the patch go to Michael H Kent.
- Enhanced Python3 support for syntax errors, these are now also compatible.
- Support for macOS X was added.
- Support for using the clang compiler was added, it can be enforced via `--clang` option. Currently this option is mainly intended to allow testing the "macOS X" support as good as possible under Linux.

Optimization

- Enhanced all optimization that previously worked on "constants" to work on "compile time constants" instead. A "compile time constant" can currently also be any form of a built-in name or exception reference. It is intended to expand this in the future.
- Added support for built-ins `bin`, `oct`, and `hex`, which also can be computed at compile time, if their arguments are compile time constant.
- Added support for the `iter` built-in in both forms, one and two arguments. These cannot be computed at compile time, but now will execute faster.
- Added support for the `next` built-in, also in its both forms, one and two arguments. These also cannot be computed at compile time, but now will execute faster as well.
- Added support for the `open` built-in in all its form. We intend for future releases to be able to track file opens for including them into the executable if data files.
- Optimize the `__debug__` built-in constant as well. It cannot be assigned, yet code can determine a mode of operation from it, and apparently some code does. When compiling the mode is decided.

- Optimize the `Ellipsis` built-in constant as well. It falls in the same category as `True`, `False`, `None`, i.e. names of built-in constants that are singletons.
- Added support for anonymous built-in references, i.e. built-ins which have names that are not normally accessible. An example is `type(None)` which is not accessible from anywhere. Other examples of such names are `compiled_method_or_function`.

Having these as represented internally, and flagged as "compile time constants", allows the compiler to make more compile time optimization and to generate more efficient C++ code for it that won't e.g. call the `type` built-in with `None` as an argument.

- All built-in names used in the program are now converted to "built-in name references" in a first step. Unsupported built-ins like e.g. `zip`, for which Nuitka has no own code or understanding yet, remained as "module variables", which made access to them slow, and difficult to recognize.
- Added optimization for module attributes `__file__`, `__doc__` and `__package__` if they are read only. It's the same as was done for `__name__` so far only.
- Added optimization for slices and subscripts of "compile time constant" values. These will play a more important role, once value propagation makes them more frequent.

Organisational

- Created a "change log" from the previous release announcements. It's as ReStructured Text and converted to PDF for the release as well, but I chose not to include that in Debian, because it's so easy to generate the PDF on that yourself.
- The posting of release announcements is now prepared by a script that converts the ReStructured Text to HTML and adds it to Wordpress as a draft posting or updates it, until it's release time. Simple, sweet and elegant.

Cleanups

- Split out the `nuitka.nodes.Nodes` module into many topic nodes, so that there are now `nuitka.nodes.BoolNodes` or `nuitka.nodes.LoopNodes` to host nodes of similar kinds, so that it is now cleaner.
- Split `del` statements into their own node kind, and use much simpler node structures for them. The following blocks are absolutely the same:

```
del a, b.c, d
```

```
del a
del b.c
del d
```

So that's now represented in the node tree. And even more complex looking cases, like this one, also the same:

```
del a, (b.c, d)
```

This one gives a different parse tree, but the same bytecode. And so Nuitka need no longer concern itself with this at all, and can remove the tuple from the parse tree immediately. That makes them easy to handle. As you may have noted already, it also means, there is no way to enforce that two things are deleted or none at all.

- Turned the function and class builder statements into mere assignment statements, where defaults and base classes are handled by wrapping expressions.

Previously they are also kind of assignment statements too, which is not needed. Now they were reduced to only handle the `bases` for classes and the `defaults` for functions and make optional.

- Refactored the decorator handling to the tree building stage, presenting them as function calls on "function body expression" or class body expression".

This allowed to remove the special code for decorators from code generation and C++ templates, making decorations easy subjects for future optimization, as they practically are now just function calls.

```
@some_classdecorator
class C:
    @staticmethod
    def f():
        pass
```

It's just a different form of writing things. Nothing requires the implementation of decorators, it's just functions calls with function bodies before the assignment.

The following is only similar:

```
class C:
    def f():
        pass

    f = staticmethod(f)

C = some_classdecorator(C)
```

It's only similar, because the assignment to an intermediate value of `C` and `f` is not done, and if an exception was raised by the decoration, that name could persist. For Nuitka, the function and class body, before having a name, are an expression, and so can of course be passed to decorators already.

- The in-place assignments statements are now handled using temporary variable blocks

Adding support for scoped temporary variables and references to them, it was possible to re-formulate in-place assignments expressions as normal look-ups, in-place operation call and then assignment statement. This allowed to remove static templates and will yield even better generated code in the future.

- The for loop used to have has a "source" expression as child, and the iterator over it was only taken at the code generation level, so that step was therefore invisible to optimization. Moved it to tree building stage instead, where optimization can work on it then.
- Tree building now generally allows statement sequences to be `None` everywhere, and pass statements are immediately eliminated from them immediately. Empty statement sequences are now forbidden to exist.
- Moved the optimization for `__name__` to compute node of variable references, where it doesn't need anything complex to replace with the constant value if it's only read.
- Added new bases classes and mix-in classes dedicated to expressions, giving a place for some defaults.
- Made the built-in code more reusable.

New Tests

- Added some more diagnostic tests about complex assignment and `del` statements.
- Added syntax test for star import on function level, that must fail on Python3.
- Added syntax test for duplicate argument name.
- Added syntax test for global on a function argument name.

Summary

The decorator and building changes, the assignment changes, and the node cleanups are all very important progress for the type inference work, because they remove special casing that previously would have been required. Lambdas and functions now really are the same thing right after tree building. The in-place assignments are now merely done using standard assignment code, the built functions and classes are now assigned to names in assignment statements, much *more* consistency there.

Yet, even more work will be needed in the same direction. There may e.g. be work required to cover `with` statements as well. And assignments will become no more complex than unpacking from a temporary variable.

For this release, there is only minimal progress on the Python3 front, despite the syntax support, which is only minuscule progress. The remaining tasks appear all more or less difficult work that I don't want to touch now.

There are still remaining steps, but we can foresee that a release may be done that finally actually does type inference and becomes the effective Python compiler this project is all about.

Nuitka Release 0.3.19

This time there are a few bug fixes, major cleanups, more Python3 support, and even new features. A lot of things in this are justifying a new release.

Bug Fixes

- The man pages of `nuitka` and `nuitka-python` had no special layout for the option groups and broken whitespace for `--recurse-to` option. Also `--g++-only` was only partially bold. Released as 0.3.18.1 hot fix already.
- The command line length improvement we made to Scons for Windows was not portable to Python2.6. Released as 0.3.18.2 hot fix already.
- Code to detect already considered packages detection was not portable to Windows, for one case, there was still a use of `/` instead of using a `joinpath` call. Released as 0.3.18.3 already.
- A call to the range built-in with no arguments would crash the compiler, see Released as 0.3.18.4 already.
- Compatibility Fix: When rich comparison operators returned false value other `False`, for comparison chains, these would not be used, but `False` instead, see .
- The support for `__import__` didn't cover keyword arguments, these were simply ignored. Fixed, but no warning is given yet.

New Features

- A new option has been added, one can now specify `--recurse-directory` and Nuitka will attempt to embed these modules even if not obviously imported. This is not yet working perfectly yet, but will receive future improvements.

- Added support for the `exec` built-in of Python3, this enables us to run one more basic test, `GlobalStatement.py` with Python3. The test `ExecEval.py` nearly works now.

Optimization

- The no arguments `range()` call now optimized into the static CPython exception it raises.
- Parts of comparison chains with constant arguments are now optimized away.

Cleanups

- Simplified the `CPythonExpressionComparison` node, it now always has only 2 operands.
If there are more, the so called "comparison chain", it's done via `and` with assignments to temporary variables, which are expressed by a new node type `CPythonExpressionTempVariableRef`. This allowed to remove `expression_temps` from C++ code templates and generation, reducing the overall complexity.
- When executing a module (`--execute` but not `--exe`), no longer does Nuitka import it into itself, instead a new interpreter is launched with a fresh environment.
- The calls to the variadic `MAKE_TUPLE` were replaced with calls the `MAKE_TUPLExx` (where `xx` is the number of arguments), that are generated on a as-needed basis. This gives more readable code, because no `EVAL_ORDERED_xx` is needed at call site anymore.
- Many node classes have moved to new modules in `nuitka.nodes` and grouped by theme. That makes them more accessible.
- The choosing of the debug python has moved from Scons to Nuitka itself. That way it can respect the `sys.abiflags` and works with Python3.
- The replacing of `.py` in filenames was made more robust. No longer is `str.replace` used, but instead proper means to assure that having `.py` as other parts of the filenames won't be a trouble.
- Module recursion was changed into its own module, instead of being hidden in the optimization that considers import statements.
- As always, some PyLint work, and some minor `TODO` were solved.

Organisational

- Added more information to the [Developer Manual](#), e.g. documenting the tree changes for `assert` to become a conditional statement with a `raise` statement, etc.
- The Debian package is as of this version verified to be installable and functional on to Ubuntu Natty, Maverick, Oneiric, and Precise.
- Added support to specify the binary under test with a `NUITKA` environment, so the test framework can run with installed version of Nuitka too.
- Made sure the test runners work under Windows as well. Required making them more portable. And a workaround for `os.execl` not propagating exit codes under Windows.
- For windows target the MinGW library is now linked statically. That means there is no requirement for MinGW to be in the `PATH` or even installed to execute the binary.

New Tests

- The `basic`, `programs`, `syntax`, and `reflected` were made executable under Windows. Occasionally this meant to make the test runners more portable, or to work around limitations.

- Added test to cover return values of rich comparisons in comparison chains, and order of argument evaluation for comparison chains.
- The `Referencing.py` test was made portable to Python3.
- Cover no arguments `range()` exception as well.
- Added test to demonstrate that `--recurse-directory` actually works. This is using an `__import__` that cannot be predicted at run time (yet).
- The created source package is now tested on pbuilder chroots to be pass installation and the basic tests, in addition to the full tests during package build time on these chroots. This will make sure, that Nuitka works fine on Ubuntu Natty and doesn't break without notice.

Summary

This releases contains many changes. The "temporary variable ref" and "assignment expression" work is ground breaking. I foresee that it will lead to even more simplifications of code generation in the future, when e.g. in-place assignments can be reduced to assignments to temporary variables and conditional statements.

While there were many improvements related to Windows support and fixing portability bugs, or the Debian package, the real focus is the optimization work, which will ultimately end with "value propagation" working.

These are the real focus. The old comparison chain handling was a big wart. Working, but no way understood by any form of analysis in Nuitka. Now they have a structure which makes their code generation based on semantics and allows for future optimization to see through them.

Going down this route is an important preparatory step. And there will be more work like this needed. Consider e.g. handling of in-place assignments. With an "assignment expression" to a "temporary variable ref", these become the same as user code using such a variable. There will be more of these to find.

So, that is where the focus is. The release now was mostly aiming at getting involved fixes out. The bug fixed by comparison chain reworking, and the `__import__` related one, were not suitable for hot fix releases, so that is why the 0.3.19 release had to occur now. But with plugin support, with this comparison chain cleanup, with improved Python3 support, and so on, there was plenty of good stuff already, also worth to get out.

Nuitka Release 0.3.18

This is to inform you about the new stable release of Nuitka. This time there are a few bug fixes, and the important step that triggered the release: Nuitka has entered Debian Unstable. So you if want, you will get stable Nuitka releases from now on via `apt-get install nuitka`.

The release cycle was too short to have much focus. It merely includes fixes, which were available as hot fixes, and some additional optimization and node tree cleanups, as well as source cleanups. But not much else.

Bug Fixes

- Conditional statements with both branches empty were not optimized away in all cases, triggering an assertion of code generation. Released as 0.3.17a hot fix already.
- Nuitka was considering directories to contain packages that had no `__init__.py` which could lead to errors when it couldn't find the package later in the compilation process. Released as 0.3.17a hot fix already.
- When providing `locals()` to `exec` statements, this was not making the `locals()` writable. The logic to detect the case that default value is used (`None`) and be pessimistic about it, didn't consider the actual value `locals()`. Released as 0.3.17b hot fix already.

- Compatibility Fix: When no defaults are given, CPython uses `None` for `func.func_defaults`, but Nuitka had been using `None`.

Optimization

- If the condition of assert statements can be predicted, these are now optimized in a static raise or removed.
- For built-in name references, there is now dedicated code to look them up, that doesn't check the module level at all. Currently these are used in only a few cases though.
- Cleaner code is generated for the simple case of `print` statements. This is not only faster code, it's also more readable.

Cleanups

- Removed the `CPythonStatementAssert` node.

It's not needed, instead at tree building, assert statements are converted to conditional statements with the asserted condition result inverted and a raise statement with `AssertionError` and the assertion argument.

This allowed to remove code and complexity from the subsequent steps of Nuitka, and enabled existing optimization to work on assert statements as well.

- Moved built-in exception names and built-in names to a new module `nuitka.Builtins` instead of having in other places. This was previously a bit spread-out and misplaced.
- Added cumulative `tags` to node classes for use in checks. Use it to annotate which node kinds to visit in e.g. per scope finalization steps. That avoids kinds and class checks.
- New node for built-in name lookups

This allowed to remove tricks played with adding module variable lookups for `staticmethod` when adding them for `__new__` or module variable lookups for `str` when predicting the result of `type('a')`, which was unlikely to cause a problem, but an important `TODO` item still.

Organisational

- The "[Download](#)" page is now finally updated for releases automatically.

Up to this release, I had to manually edit that page, but now mastered the art of upload via XMLRPC and a Python script, so that don't lose as much time with editing, checking it, etc.

- The Debian package is backportable to Ubuntu Natty, Maverick, Oneiric, I expect to make a separate announcement with links to packages.
- Made sure the test runners work with bare `python2.6` as well.

New Tests

- Added some tests intended for type inference development.

Summary

This release contains not as much changes as others, mostly because it's the intended base for a Debian upload.

The `exec` fix was detected by continued work on the `branch/feature/minimize_CPython26_tests_diff` branch, but that work is now complete.

It is being made pretty (many git rebase iterations) with lots of Issues being added to the bug tracker and referenced for each change. The intention is to have a clean commits repository with the changed made.

But of course, the real excitement is the "type inference" work. It will give a huge boost to Nuitka. With this in place, new benchmarks may make sense. I am working on getting it off the ground, but also to make us more efficient.

So when I learn something. e.g. `assert` is not special, I apply it to the `develop` branch immediately, to keep the differences as small as possible, and to immediately benefit from such improvements.

Nuitka Release 0.3.17

This is to inform you about the new stable release of Nuitka. This time there are a few bug fixes, lots of very important organisational work, and yet again improved compatibility and cleanups. Also huge is the advance in making `--deep` go away and making the recursion of Nuitka controllable, which means a lot for scalability of projects that use a lot of packages that use other packages, because now you can choose which ones to embed and which ones one.

The release cycle had a focus on improving the quality of the test scripts, the packaging, and generally to prepare the work on "type inference" in a new feature branch.

I have also continued to work towards CPython3.2 compatibility, and this version, while not there, supports Python3 with a large subset of the basic tests programs running fine (of course via `2to3` conversion) without trouble. There is still work to do, exceptions don't seem to work fully yet, parameter parsing seems to have changed, etc. but it seems that CPython3.2 is going to work one day.

And there has been a lot of effort, to address the Debian packaging to be cleaner and more complete, addressing issues that prevented it from entering the Debian repository.

Bug Fixes

- Fixed the handling of modules and packages of the same name, but with different casing. Problem showed under Windows only. Released as 0.3.16a hot fix already.
- Fixed an error where the command line length of Windows was exceeded when many modules were embedded, Christopher Tott provided a fix for it. Released as 0.3.16a hot fix already.
- Fix, avoid to introduce new variables for where built-in exception references are sufficient. Released as 0.3.16b hot fix already.
- Fix, add the missing `staticmethod` decorator to `__new__` methods before resolving the scopes of variables, this avoids the use of that variable before it was assigned a scope. Released as 0.3.16b hot fix already.

New Features

- Enhanced compatibility again, provide enough `co_varnames` in the code objects, so that slicing them up to `code_object.co_argcount` will work. They are needed by `inspect` module and might be used by some decorators as well.
- New options to control the recursion:
`--recurse-none` (do not warn about not-done recursions) `--recurse-all` (recurse to all otherwise warned modules) `--recurse-to` (confirm to recurse to those modules) `--recurse-not-to` (confirm to not recurse to those modules)

Optimization

- The optimization of constant conditional expressions was not done yet. Added this missing constant propagation case.

- Eliminate near empty statement sequences (only contain a pass statement) in more places, giving a cleaner node structure for many constructs.
- Use the pickle "protocol 2" on CPython2 except for `unicode` strings where it does not work well. It gives a more compressed and binary representation, that is generally more efficient to un-stream as well. Also use the cPickle protocol, the use of `pickle` was not really necessary anymore.

Organisational

- Added a [Developer Manual](#) to the release. It's incomplete, but it details some of the existing stuff, coding rules, plans for "type inference", etc.
- Improved the `--help` output to use `metavar` where applicable. This makes it more readable for some options.
- Instead of error message, give help output when no module or program file name was given. This makes Nuitka help out more convenient.
- Consistently use `#!/usr/bin/env python` for all scripts, this was previously only done for some of them.
- Ported the PyLint check script to Python as well, enhancing it on the way to check the exit code, and to only output changes things, as well as making the output of warnings for `TODO` items optional.
- All scripts used for testing, PyLint checking, etc. now work with Python3 as well. Most useful on Arch Linux, where it's also already the default for `Python`.
- The help output of Nuitka was polished a lot more. It is now more readable and uses option groups to combine related options together.
- Make the tests run without any dependence on `PATH` to contain the executables of Nuitka. This makes it easier to use.
- Add license texts to 3rd party file that were missing them, apply `licensecheck` results to cleanup Nuitka. Also removed own copyright statement from in-line copy of Scons, it had been added by accident only.
- Release the tests that I own as well as the Debian packaging I created under "Apache License 2.0" which is very liberal, meaning every project will be able to use it.
- Don't require copyright assignment for contributions anymore, instead only "Apache License 2.0", the future Nuitka license, so that the code won't be a problem when changing the license of all of Nuitka to that license.
- Give contributors listed in the [User Manual](#) an exception to the GPL terms until Nuitka is licensed under "Apache License 2.0" as well.
- Added an `--experimental` option which can be used to control experimental features, like the one currently being added on branch `feature/ctypes_annotation`, where "type inference" is currently only activated when that option is given. For this stable release, it does nothing.
- Check the static C++ files of Nuitka with `cppcheck` as well. Didn't find anything.
- Arch Linux packages have been contributed, these are linked for download, but the stable package may lag behind a bit.

Cleanups

- Changed `not` boolean operation to become a normal operator. Changed `and` and `or` boolean operators to a new base class, and making their interface more similar to that of operations.
- Added cumulative `tags` to node classes for use in checks. Use it to annotate which node kinds to visit in e.g. per scope finalization steps. That avoids kinds and class checks.

- Enhanced the "visitor" interface to provide more kinds of callbacks, enhanced the way "each scope" visiting is achieved by generalizing it as "child has not tag 'closure_taker'" and that for every "node that has tag 'closure_taker'".
- Moved `SyntaxHighlighting` module to `nuitka.gui` package where it belongs.
- More white listing work for imports. As recursion is now the default, and leads to warnings for non-existent modules, the CPython tests gave a lot of good candidates for import errors that were white listed.
- Consistently use `nuitka` in test scripts, as there isn't a `Nuitka.py` on all platforms. The later is scheduled for removal.
- Some more PyLint cleanups.

New Tests

- Make sure the basic tests pass with CPython or else fail the test. This is to prevent false positives, where a test passes, but only because it fails in CPython early on and then does so with Nuitka too. For the syntax tests we make sure they fail.
- The basic tests can now be run with `PYTHON=python3.2` and use `2to3` conversion in that case. Also the currently not passing tests are not run, so the passing tests continue to do so, with this run from the release test script `check-release`.
- Include the syntax tests in release tests as well.
- Changed many existing tests so that they can run under CPython3 too. Of course this is via `2to3` conversion.
- Don't fail if the CPython test suites are not there.

Currently they remain largely unpublished, and as such are mostly only available to me (exception, `feature/minimize_CPython26_tests_diff` branch references the CPython2.6 tests repository, but that remains work in progress).

- For the compile itself test: Make the presence of the Scons in-line copy optional, the Debian package doesn't contain it.
- Also make it more portable, so it runs under Windows too, and allow to choose the Python version to test. Check this test with both CPython2.6 and CPython2.7 not only the default Python.
- Before releasing, test that the created Debian package builds fine in a minimal Debian `unstable chroot`, and passes all the tests included in the package (`basics`, `syntax`, `programs`, `reflected`). Also many other Debian packaging improvements.

Summary

The "git flow" was used again in this release cycle and proved to be useful not only for hot fix, but also for creating the branch `feature/ctypes_annotation` and rebasing it often while things are still flowing.

The few hot fixes didn't require a new release, but the many organisational improvements and the new features did warrant the new release, because of e.g. the much better test handling in this release and the improved recursion control.

The work on Python3 support has slowed down a bit. I mostly only added some bits for compatibility, but generally it has slowed down. I wanted to make sure it doesn't regress by accident, so running with CPython3.2 is now part of the normal release tests.

What's still missing is more "hg" completeness. Only the `co_varnames` work for `inspect` was going in that direction, and this has slowed down. It was more important to make Nuitka's recursion more accessible with the new options, so that was done first.

And of course, the real excitement is the "type inference" work. It will give a huge boost to Nuitka, and I am happy that it seems to go well. With this in place, new benchmarks may make sense. I am working on getting it off the ground, so other people can work on it too. My idea of `ctypes` native calls may become true sooner than expected. To support that, I would like to add more tools to make sure we discover changes earlier on, checking the XML representations of tests to discover improvements and regressions more clearly.

Nuitka Release 0.3.16

This time there are many bug fixes, some important scalability work, and again improved compatibility and cleanups.

The release cycle had a focus on fixing the bug reports I received. I have also continued to look at CPython3 compatibility, and this is the first version to support Python3 somewhat, at least some of the basic tests programs run (of course via `2to3` conversion) without trouble. I don't know when, but it seems that it's going to work one day.

Also there has an effort to make the Debian packaging cleaner, addressing all kinds of small issues that prevented it from entering the Debian repository. It's still not there, but it's making progress.

Bug Fixes

- Fixed a packaging problem for Linux and x64 platform, the new `swapFiber.S` file for the fiber management was not included. Released as 0.3.15a hot fix already.
- Fixed an error where optimization was performed on removed unreachable code, which lead to an error. Released as 0.3.15b hot fix already.
- Fixed an issue with `__import__` and recursion not happening in any case, because when it did, it failed due to not being ported to new internal APIs. Released as 0.3.15c hot fix already.
- Fixed `eval()` and `locals()` to be supported in generator expressions and contractions too. Released as 0.3.15d hot fix already.
- Fixed the Windows batch files `nuitka.bat` and `nuitka-python.bat` to not output the `rem` statements with the copyright header. Released as 0.3.15d hot fix already.
- Fixed re-raise with `raise`, but without a current exception set. Released as 0.3.15e hot fix already.
- Fixed `vars()` call on the module level, needs to be treated as `globals()`. Released as 0.3.15e hot fix already.
- Fix handling of broken new lines in source files. Read the source code in "universal line ending mode". Released as 0.3.15f hot fix already.
- Fixed handling of constant module attribute `__name__` being replaced. Don't replace local variables of the same name too. Released as 0.3.15g hot fix already.
- Fixed assigning to `True`, `False` or `None`. There was this old `TODO`, and some code has compatibility craft that does it. Released as 0.3.15g hot fix already.
- Fix constant dictionaries not always being recognized as shared. Released as 0.3.15g hot fix already.
- Fix generator function objects to not require a return frame to exist. In finalize cleanup it may not.
- Fixed non-execution of cleanup codes that e.g. flush `sys.stdout`, by adding `Py_Finalize()`.
- Fix `throw()` method of generator expression objects to not check arguments properly.
- Fix missing fallback to subscript operations for slicing with non-indexable objects.
- Fix, in-place subscript operations could fail to apply the update, if the intermediate object was e.g. a list and the handle just not changed by the operation, but e.g. the length did.
- Fix, the future spec was not properly preserving the future division flag.

Optimization

- The optimization scales now much better, because per-module optimization only require the module to be reconsidered, but not all modules all the time. With many modules recursed into, this makes a huge difference in compilation time.
- The creation of dictionaries from constants is now also optimized.

New Features

- As a new feature functions now have the `func_defaults` and `__defaults__` attribute. It works only well for non-nested parameters and is not yet fully integrated into the parameter parsing. This improves the compatibility somewhat already though.
- The names `True`, `False` and `None` are now converted to constants only when they are read-only module variables.
- The `PYTHONPATH` variable is now cleared when immediately executing a compiled binary unless `--execute-with-pythonpath` is given, in which case it is preserved. This allows to make sure that a binary is in fact containing everything required.

Organisational

- The help output of Nuitka was polished a lot more. It is now more readable and uses option groups to combine related options together.
- The in-line copy of Scons is not checked with PyLint anymore. We of course don't care.
- Program tests are no longer executed in the program directory, so failed module inclusions become immediately obvious.
- The basic tests can now be run with `PYTHON=python3.2` and use `2to3` conversion in that case.

Cleanups

- Moved `tags` to a separate module, make optimization emit only documented tags, checked against the list of allowed ones.
- The Debian package has seen lots of improvements, to make it "lintian clean", even in pedantic mode. The homepage of Nuitka is listed, a watch file can check for new releases, the git repository and the gitweb are referenced, etc.
- Use `os.path.join` in more of the test code to achieve more Windows portability for them.
- Some more PyLint cleanups.

New Tests

- There is now a `Crasher` test, for tests that crashed Nuitka previously.
- Added a program test where the imported module does a `sys.exit()` and make sure it really doesn't continue after the `SystemExit` exception that creates.
- Cover the type of `__builtins__` in the main program and in imported modules in tests too. It's funny and differs between module and dict in CPython2.
- Cover a final `print` statement without newline in the test. Must still receive a newline, which only happens when `Py_Finalize()` is called.
- Added test with functions that makes a `raise` without an exception set.

- Cover the calling of `vars()` on module level too.
- Cover the use of `eval` in contractions and generator expressions too.
- Cover `func_defaults` and `__default__` attributes for a function too.
- Added test function with two `raise` in an exception handler, so that one becomes dead code and removed without the crash.

Summary

The "git flow" was really great in this release cycle. There were many hot fix releases being made, so that the bugs could be addressed immediately without requiring the overhead of a full release. I believe that this makes Nuitka clearly one of the best supported projects.

This quick turn-around also encourages people to report more bugs, which is only good. And the structure is there to hold it. Of course, the many bug fixes meant that there is not as much new development, but that is not the priority, correctness is.

The work on Python3 is a bit strange. I don't need Python3 at all. I also believe it is that evil project to remove cruft from the Python core and make developers of all relevant Python software, add compatibility cruft to their software instead. Yet, I can't really stop to work on it. It has that appeal of small fixups here and there, and then something else works too.

Python3 work is like when I was first struggling with Nuitka to pass the CPython2 unit tests for a first time. It's fun. And then it finds real actual bugs that apply to CPython2 too. Not doing `Py_Finalize` (but having to), the slice operations shortcomings, the bug of subscript in-place, and so on. There is likely more things hidden, and the earlier Python3 is supported, the more benefit from increased test covered.

What's missing is more "hg" completeness. I think only the `raise` without exception set and the `func_defaults` issue were going into its direction, but it won't be enough yet.

Nuitka Release 0.3.15

This is to inform you about the new stable release of Nuitka. This time again many organisational improvements, some bug fixes, much improved compatibility and cleanups.

This release cycle had a focus on packaging Nuitka for easier consumption, i.e. automatic packaging, making automatic uploads, improvement documentation, and generally cleaning things up, so that Nuitka becomes more compatible and ultimately capable to run the "hg" test suite. It's not there yet, but this is a huge jump for usability of Nuitka and its compatibility, again.

Then lots of changes that make Nuitka approach Python3 support, the generated C++ for at least one large example is compiling with this new release. It won't link, but there will be later releases.

And there is a lot of cleanup going on, geared towards compatibility with line numbers in the frame object.

Bug Fixes

- The main module was using `__main__` in tracebacks, but it must be `<module>`. Released as 0.3.14a hot fix already.
- Workaround for "execfile cannot be used as an expression". It wasn't possible to use `execfile` in an expression, only as a statement.

But then there is crazy enough code in e.g. mercurial that uses it in a lambda function, which made the issue more prominent. The fix now allows it to be an expression, except on the class level, which wasn't seen yet.

- The in-line copy of Scons was not complete enough to work for "Windows" or with `--windows-target` for cross compile. Fixed.

- Cached frames didn't release the "back" frame, therefore holding variables of these longer than CPython does, which could cause ordering problems. Fixed for increased compatibility.
- Handle "yield outside of function" syntax error in compiled source correctly. This one was giving a Nuitka traceback, now it gives a `SyntaxError` as it needs to.
- Made syntax/indentation error output absolutely identical to CPython.
- Using the frame objects `f_lineno` may fix endless amounts bugs related to traceback line numbers.

New Features

- Guesses the location of the MinGW compiler under Windows to default install location, so it need not be added to `PATH` environment variable. Removes the need to modify `PATH` environment just for Nuitka to find it.
- Added support for "lambda generators". You don't want to know what it is. Lets just say, it was the last absurd language feature out there, plus that didn't work. It now works perfect.

Organisational

- You can now download a Windows installer and a Debian package that works on Debian Testing, current Ubuntu and Mint Linux.
- New release scripts give us the ability to have hot fix releases as download packages immediately. That means the "git flow" makes even more beneficial to the users.
- Including the generated "README.pdf" in the distribution archives, so it can be read instead of "README.txt". The text file is fairly readable, due to the use of ReStructured Text, but the PDF is even nicer to read, due to e.g. syntax highlighting of the examples.
- Renamed the main binaries to `nuitka` and `nuitka-python`, so that there is no dependency on case sensitive file systems.
- For Windows there are batch files `nuitka.bat` and `nuitka-python.bat` to make Nuitka directly executable without finding the `Python.exe`, which the batch files can tell from their own location.
- There are now man pages of `nuitka` and `nuitka-python` with examples for the most common use cases. They are of course included in the Debian package.
- Don't strip the binary when executing it to analyse compiled binary with `valgrind`. It will give better information that way, without changing the code.

Optimization

- Implemented `swapcontext` alike (`swapFiber`) for x64 to achieve 8 times speedup for Generators. It doesn't do useless syscalls to preserve signal masks. Now Nuitka is faster at frame switching than CPython on x64, which is already good by design.

Cleanups

- Using the frame objects to store current line of execution avoids the need to store it away in helper code at all. It ought to also help a lot with threading support, and makes Nuitka even more compatible, because now line numbers will be correct even outside tracebacks, but for mere stack frame dumps.
- Moved the `for_return` detection from code generation to tree building where it belongs. Yield statements used as return statements need slightly different code for Python2.6 difference. That solved an old `TODO`.

- Much Python3 portability work. Sometimes even improving existing code, the Python compiler code had picked up a few points, where the latest Nuitka didn't work with Python3 anymore, when put to actual compile.

The test covered only syntax, but e.g. meta classes need different code in CPython3, and that's now supported. Also helper code was made portable in more places, but not yet fully. This will need more work.

- Cleaned up uses of debug defines, so they are now more consistent and in one place.
- Some more PyLint cleanups.

New Tests

- The tests are now executed by Python scripts and cover `stderr` output too. Before we only checked `stdout`. This unveiled a bunch of issues Nuitka had, but went unnoticed so far, and triggered e.g. the frame line number improvements.
- Separate syntax tests.
- The scripts to run the tests now are all in pure Python. This means, no more MinGW shell is needed to execute the tests.

Summary

The Debian package, Windows installer, etc. are now automatically updated and uploaded. From here on, there can be such packages for the hot fix releases too.

The exception tracebacks are now correct by design, and better covered.

The generator performance work showed that the approach taken by Nuitka is in fact fast. It was fast on ARM already, but it's nice to see that it's now also fast on x64. Programs using generators will be affected a lot by this.

Overall, this release brings Nuitka closer to usability. Better binary names, man pages, improved documentation, issue tracker, etc. all there now. I am in fact now looking for a sponsor for the Debian package to upload it into Debian directly.

Update

The upload to Debian happened for 0.3.18 and was done by Yaroslav Halchenko.

What's missing is more "hg" completeness. The frame release issue helped it, but `inspect.getargs()` doesn't work yet, and is a topic for a future release. Won't be easy, as `func_defaults` will be an invasive change too.

Nuitka Release 0.3.14

This is to inform you about the new stable release of Nuitka. This time it contains mostly organisational improvements, some bug fixes, improved compatibility and cleanups.

It is again the result of working towards compilation of a real program (Mercurial). This time, I have added support for proper handling of compiled types by the `inspect` module.

Bug Fixes

- Fix for "Missing checks in parameter parsing with star list, star dict and positional arguments". There was whole in the checks for argument counts, now the correct error is given. Fixed in 0.3.13a already.
- The simple slice operations with 2 values, not extended with 3 values, were not applying the correct order for evaluation. Fixed in 0.3.13a already.
- The simple slice operations couldn't handle `None` as the value for lower or upper index. Fixed in 0.3.11a already.
- The in-place simple slice operations evaluated the slice index expressions twice, which could cause problems if they had side effects. Fixed in 0.3.11a already.

New Features

- Run time patching the `inspect` module so it accepts compiled functions, compiled methods, and compiled generator objects. The `test_inspect` test of CPython is nearly working unchanged with this.
- The generator functions didn't have `CO_GENERATOR` set in their code object, setting it made compatible with CPython in this regard too. The `inspect` module will therefore return correct value for `inspect.isgeneratorfunction()` too.

Optimization

- Slice indexes that are `None` are now constant propagated as well.
- Slightly more efficient code generation for dual star arg functions, removing useless checks.

Cleanups

- Moved the Scons, static C++ files, and assembler files to new package `nuitka.build` where also now `SconsInterface` module lives.
- Moved the Qt dialog files to `nuitka.gui`
- Moved the "unfreezer" code to its own static C++ file.
- Some PyLint cleanups.

New Tests

- New test `Recursion` to cover recursive functions.
- New test `Inspection` to cover the patching of `inspect` module.
- Cover `execfile` on the class level as well in `ExecEval` test.
- Cover evaluation order of simple slices in `OrderCheck` too.

Organisational

- There is a new issue tracker available (since migrated and removed)

Please register and report issues you encounter with Nuitka. I have put all the known issues there and started to use it recently. It's Roundup based like <https://bugs.python.org> is, so people will find it familiar.

- The `setup.py` is now apparently functional. The source releases for download are made it with, and it appears the binary distributions work too. We may now build a windows installer. It's currently in testing, we will make it available when finished.

Summary

The new source organisation makes packaging Nuitka really easy now. From here, we can likely provide "binary" package of Nuitka soon. A windows installer will be nice.

The patching of `inspect` works wonders for compatibility for those programs that insist on checking types, instead of doing duck typing. The function call problem, was an issue found by the Mercurial test suite.

For the "hg.exe" to pass all of its test suite, more work may be needed, this is the overall goal I am currently striving for. Once real world programs like Mercurial work, we can use these as more meaningful benchmarks and resume work on optimization.

Nuitka Release 0.3.13

This release is mostly the result of working towards compilation of a real programs (Mercurial) and to merge and finalize the frame stack work. Now Nuitka has a correct frame stack at all times, and supports `func_code` and `gi_code` objects, something previously thought to be impossible.

Actually now it's only the "bytecode" objects that won't be there. And not attributes of `func_code` are meaningful yet, but in theory can be supported.

Due to the use of the "git flow" for Nuitka, most of the bugs listed here were already fixed in on the stable release before this release. This time there were 5 such hot fix releases, sometimes fixing multiple bugs.

Bug Fixes

- In case of syntax errors in the main program, an exception stack was giving that included Nuitka code. Changed to make the same output as CPython does. Fixed in 0.3.12a already.
- The star import (`from x import *`) didn't work for submodules. Providing `*` as the import list to the respective code allowed to drop the complex lookups we were doing before, and to simply trust CPython C/API to do it correctly. Fixed in 0.3.12 already.
- The absolute import is *not* the default of CPython 2.7 it seems. A local `posix` package shadows the standard library one. Fixed in 0.3.12 already.
- In `--deep` mode, a module may contain a syntax error. This is e.g. true of "PyQt" with `port_v3` included. These files contain Python3 syntax and fail to be imported in Python2, but that is not to be considered an error. These modules are now skipped with a warning. Fixed in 0.3.12b already.
- The code to import modules wasn't using the `__import__` built-in, which prevented `__import__` overriding code to work. Changed import to use the built-in. Fixed in 0.3.12c already.
- The code generated for the `__import__` built-in with constant values was doing relative imports only. It needs to attempt relative and absolute imports. Fixed in 0.3.12c already.
- The code of packages in "`__init__.py`" believed it was outside of the package, giving problems for package local imports. Fixed in 0.3.12d already.
- It appears that "Scons", which Nuitka uses internally and transparent to you, to execute the compilation and linking tasks, was sometimes not building the binaries or shared libraries, due to a false caching. As a workaround, these are now erased before doing the build. Fixed in 0.3.12d already.
- The use of `in` and `not in` in comparison chains (e.g. `a < b < c` is one), wasn't supported yet. The use of these in comparison chains `a in b in c` is very strange.
Only in the `test_grammar.py` it was ever used I believe. Anyway, it's supported now, solving this `TODO` and reducing the difference. Fixed in 0.3.12e already.
- The order of evaluation for `in` and `not in` operators wasn't enforced in a portable way. Now it is correct on "ARM" too. Fixed in 0.3.12e already.

Optimization

- The built-ins `GeneratorExit` and `StopIteration` are optimized to their Python C/API names where possible as well.

Cleanups

- The `__file__` attribute of modules was the relative filename, but for absolute filenames these become a horrible mess at least on Linux.
- Added assertion helpers for sane frame and code objects and use them.
- Make use of `assertObject` in more places.
- Instead of using `os.path.sep` all over, added a helper `Utils.joinpath` that hides this and using `os.path.join`. This gives more readable code.
- Added traces to the "unfreezer" guarded by a `define`. Helpful in analyzing import problems.
- Some PyLint cleanups removing dead code, unused variables, useless pass statement, etc.

New Tests

- New tests to cover `SyntaxError` and `IndentationError` from `--deep` imports and in main program.
- New test to cover evaluation order of `in` and `not in` comparisons.
- New test to cover package local imports made by the `"__init__.py"` of the package.

Organisational

- Drop `"compile_itself.sh"` in favor of the new `"compile_itself.py"`, because the later is more portable.
- The logging output is now nicer, and for failed recursions, outputs the line that is having the problem.

Summary

The frame stack work and the `func_code` are big for compatibility.

The `func_code` was also needed for "hg" to work. For Mercurial to pass all of its test suite, more work will be needed, esp. the `inspect` module needs to be run-time patched to accept compiled functions and generators too.

Once real world programs like Mercurial work, we can use these as more meaningful benchmarks and resume work on optimization.

Nuitka Release 0.3.12

This is to inform you about the new release of Nuitka many bug fixes, and substantial improvements especially in the organisational area. There is a new [User Manual \(PDF\)](#), with much improved content, a `sys.meta_path` based import mechanism for `--deep` mode, git flow goodness.

This release is generally also the result of working towards compilation of a real programs (Mercurial) and to get things work more nicely on Windows by default. Thanks go to Liu Zhenhai for helping me with this goal.

Due to the use of the "git flow", most of the bugs listed here were already fixed in on the stable release before this release. And there were many of these.

Bug Fixes

- The order of evaluation for base classes and class dictionaries was not enforced.

Apparently nothing in the CPython test suite did that, I only noticed during debugging that Nuitka gave a different error than CPython did, for a class that had an undefined base class, because both class body and base classes were giving an error. Fixed in 0.3.11a already.

- Method objects didn't hold a reference to the used class.

The effect was only noticed when `--python-debug` was used, i.e. the debug version of Python linked, because then the garbage collector makes searches. Fixed in 0.3.11b already.

- Set `sys.executable` on Linux as well. On Debian it is otherwise `/usr/bin/python` which might be a different version of Python entirely. Fixed in 0.3.11c already.
- Embedded modules inside a package could hide package variables of the same name. Learned during PyCON DE about this corner case. Fixed in 0.3.11d already.
- Packages could be duplicated internally. This had no effect on generated code other than appearing twice in the list if frozen modules. Fixed in 0.3.11d already.
- When embedding modules from outside current directory, the look-up failed. The embedding only ever worked for the compile itself and programs test cases, because they are all in the current directory then. Fixed in 0.3.11e already.
- The check for ARM target broke Windows support in the Scons file. Fixed in 0.3.11f already.
- The star import from external modules failed with an error in `--deep` mode. Fixed in 0.3.11g already.
- Modules with a parent package could cause a problem under some circumstances. Fixed in 0.3.11h already.
- One call variant, with both list and dict star arguments and keyword arguments, but no positional parameters, didn't have the required C++ helper function implemented. Fixed in 0.3.11h already.
- The detection of the CPU core count was broken on my hexacore at least. Gave 36 instead of 6, which is a problem for large programs. Fixed in 0.3.11h already.
- The in-line copy of Scons didn't really work on Windows, which was sad, because we added it to simplify installation on Windows precisely because of this.
- Cleaning up the build directory from old sources and object files wasn't portable to Windows and therefore wasn't effective there.
- From imports where part of the imported were found modules and parts were not, didn't work. Solved by the feature branch `meta_path_import` that was merged for this release.
- Newer MinGW gave warnings about the default visibility not being possible to apply to class members. Fixed by not setting this default visibility anymore on Windows.
- The `sys.executable` gave warnings on Windows because of backslashes in the path. Using a raw string to prevent such problems.
- The standard library path was hard coded. Changed to run time detection.

Cleanups

- Version checks on Python runtime now use a new define `PYTHON_VERSION` that makes it easier. I don't like `PY_VERSION_HEX`, because it is so unreadable. Makes some of the checks a lot more safe.
- The `sys.meta_path` based import from the `meta_path_import` feature branch allowed the cleanup the way importing is done. It's a lot less code now.
- Removed some unused code. We will aim at making Nuitka the tool to detect dead code really.

- Moved `nuitka.Nodes` to `nuitka.nodes.Nodes`, that is what the package is intended for, the split will come later.

New Tests

- New tests for import variants that previously didn't work: Mixed imports. Imports from a package one level up. Modules hidden by a package variable, etc.
- Added test of function call variant that had no test previously. Only found it when compiling "hg". Amazing how nothing in my tests, CPython tests, etc. used it.
- Added test to cover the partial success of import statements.
- Added test to cover evaluation order of class definitions.

Organisational

- Migrated the "README.txt" from org-mode to ReStructured Text, which allows for a more readable document, and to generate a nice [User Manual](#) in PDF form.
- The amount of information in "README.txt" was increased, with many more subjects are now covered, e.g. "git flow" and how to join Nuitka development. It's also impressive to see what code blocks and syntax highlighting can do for readability.
- The Nuitka git repository has seen multiple hot fixes.

These allowed to publish bug fixes immediately after they were made, and avoided the need for a new release just to get these out. This really saves me a lot of time too, because I can postpone releasing the new version until it makes sense because of other things.

- Then there was a feature branch `meta_path_import` that lived until being merged to `develop` to improve the import code, which is now released as part of the main branch. Getting that feature right took a while.
- And there is the feature branch `minimize_CPython26_tests_diff` which has some success already in documenting the required changes to the "CPython26" test suite and in reducing the amount of differences, while doing it. We have a frame stack working there, albeit in too ugly code form.
- The release archives are now built using `setuptools`. You can now also download a zip file, which is probably more Windows friendly. The intention is to work on that to make `setup.py` produce a Nuitka install that won't rely on any environment variables at all. Right now `setup.py` won't even allow any other options than `sdist` to be given.
- Ported "compile_itself.sh" to "compile_itself.py", i.e. ported it to Python. This way, we can execute it easily on Windows too, where it currently still fails. Replacing `diff`, `rm -rf`, etc. is a challenge, but it reduces the dependency on MSYS tools on Windows.
- The compilation of standard library is disabled by default, but `site` or `dist` packages are now embedded. To include even standard library, there is a `--really-deep` option that has to be given in addition to `--deep`, which forces this.

Summary

Again, huge progress. The improved import mechanism is very beautiful. It appears that little is missing to compile real world programs like "hg" with Nuitka. The next release cycle will focus on that and continue to improve the Windows support which appears to have some issues.

Nuitka Release 0.3.11

This is to inform you about the new release of Nuitka with some bug fixes and portability work.

This release is generally cleaning up things, and makes Nuitka portable to ARM Linux. I used to host the Nuitka homepage on that machine, but now that it's no longer so, I can run heavy compile jobs on it. To my surprise, it found many portability problems. So I chose to fix that first, the result being that Nuitka now works on ARM Linux too.

Bug Fixes

- The order of slice expressions was not correct on x86 as well, and I found that with new tests only. So the porting to ARM revealed a bug category, I previously didn't consider.
- The use of `linux2` in the Scons file is potentially incompatible with Linux 3.0, although it seems that at least on Debian the `sys.platform` was changed back to `linux2`. Anyway, it's probably best to allow just anything that starts with `linux` these days.
- The `print` statement worked like a `print` function, i.e. it first evaluated all printed expressions, and did the output only then. That is incompatible in case of exceptions, where partial outputs need to be done, and so that got fixed.

Optimization

- Function calls now each have a dedicated helper function, avoiding in some cases unnecessary work. We will may build further on this and in-line `PyObject_Call` differently for the special cases.

Cleanups

- Moved many C++ helper declarations and in-line implementations to dedicated header files for better organisation.
- Some dependencies were removed and consolidated to make the dependency graph sane.
- Multiple decorators were in reverse order in the node tree. The code generation reversed it back, so no bug, yet that was a distorted tree.

Finding this came from the ARM work, because the "reversal" was in fact just the argument evaluation order of C++ under x86/x64, but on ARM that broke. Correcting it highlighted this issue.

- The deletion of slices, was not using `Py_ssize` for indexes, disallowing some kinds of optimization, so that was harmonized.
- The function call code generation got a general overhaul. It is now more consistent, has more helpers available, and creates more readable code.
- PyLint is again happier than ever.

New Tests

- There is a new basic test `OrderChecks` that covers the order of expression evaluation. These problems were otherwise very hard to detect, and in some cases not previously covered at all.
- Executing Nuitka with Python3 (it won't produce correct Python3 C/API code) is now part of the release tests, so non-portable code of Nuitka gets caught.

Organisational

- Support for ARM Linux. I will make a separate posting on the challenges of this. Suffice to say now, that C++ leaves way too much things unspecified.
- The Nuitka git repository now uses "git flow". The new git policy will be detailed in another [separate posting](#).

- There is an unstable `develop` branch in which the development occurs. For this release ca. 40 commits were done to this branch, before merging it. I am also doing more fine grained commits now.
- Unlike previously, there is `master` branch for the stable release.
- There is a script "`make-dependency-graph.sh`" (Update: meanwhile it was renamed to "`make-dependency-graph.py`") to produce a dependency graphs of Nuitka. I detected a couple of strange things through this.
- The Python3 `__pycache__` directories get removed too by the cleanup script.

Numbers

We only have "PyStone" now, and on a new machine, so the numbers cannot be compared to previous releases:

python 2.6:

```
Pystone(1.1) time for 50000 passes = 0.48
This machine benchmarks at 104167 pystones/second
```

Nuitka 0.3.11 (driven by python 2.6):

```
Pystone(1.1) time for 50000 passes = 0.19
This machine benchmarks at 263158 pystones/second
```

So this a speedup factor of 258%, last time on another machine it was 240%. Yet it only proves that the generated and compiled are more efficient than bytecode, but Nuitka doesn't yet do the relevant optimization. Only once it does, the factor will be significantly higher.

Summary

Overall, there is quite some progress. Nuitka is a lot cleaner now, which will help us later only. I wanted to get this out, mostly because of the bug fixes, and of course just in case somebody attempts to use it on ARM.

Nuitka Release 0.3.10

This new release is major milestone 2 work, enhancing practically all areas of Nuitka. The focus was roundup and breaking new grounds with structural optimization enhancements.

Bug Fixes

- Exceptions now correctly stack.

When you catch an exception, there always was the exception set, but calling a new function, and it catching the exception, the values of `sys.exc_info()` didn't get reset after the function returned.

This was a small difference (of which there are nearly none left now) but one that might effect existing code, which affects code that calls functions in exception handling to check something about it.

So it's good this is resolved now too. Also because it is difficult to understand, and now it's just like CPython behaves, which means that we don't have to document anything at all about it.

- Using `exec` in generator functions got fixed up. I realized that this wouldn't work while working on other things. It's obscure yes, but it ought to work.
- Lambda generator functions can now be nested and in generator functions. There were some problems here with the allocation of closure variables that got resolved.

- List contractions could not be returned by lambda functions. Also a closure issue.
- When using a mapping for globals to `exec` or `eval` that had a side effect on lookup, it was evident that the lookup was made twice. Correcting this also improves the performance for the normal case.

Optimization

- Statically raised as well as predicted exceptions are propagated upwards, leading to code and block removal where possible, while maintaining the side effects.

This is brand new and doesn't do everything possible yet. Most notable, the matching of raised exception to handlers is not yet performed.

- Built-in exception name references and creation of instances of them are now optimized as well, which leads to faster exception raising/catching for these cases.
- More kinds of calls to built-ins are handled, positional parameters are checked and more built-ins are covered.

Notable is that now checks are performed if you didn't potentially overload e.g. the `len` with your own version in the module. Locally it was always detected already. So it's now also safe.

- All operations and comparisons are now simulated if possible and replaced with their result.
- In the case of predictable true or false conditions, not taken branches are removed.
- Empty branches are now removed from most constructs, leading to sometimes cleaner code generated.

Cleanups

- Removed the lambda body node and replaced it with function body. This is a great win for the split into body and builder. Regular functions and lambda functions now only differ in how the created body is used.
- Large cleanup of the operation/comparison code. There is now only use of a simulator function, which exists for every operator and comparison. This one is then used in a prediction call, shared with the built-in predictions.
- Added a `Tracing` module to avoid future imports of `print_function`, which annoyed me many times by causing syntax failures for when I quickly added a print statement, not noting it must have the braces.
- PyLint is happier than ever.

New Tests

- Enhanced `OverflowFunctions` test to cover even deeper nesting of overflow functions taking closure from each level. While it's not yet working, this makes clearer what will be needed. Even if this code is obscure, I would like to be that correct here.
- Made `Operators` test to cover the ```` operator as well.
- Added to `ListContractions` the case where a contraction is returned by a lambda function, but still needs to leak its loop variable.
- Enhanced `GeneratorExpressions` test to cover lambda generators, which is really crazy code:

```
def y():
    yield ((yield 1), (yield 2))
```

- Added to `ExecEval` a case where the `exec` is inside a generator, to cover that too.

- Activated the testing of `sys.exc_info()` in `ExceptionRaising` test.

This was previously commented out, and now I added stuff to illustrate all of the behavior of CPython there.

- Enhanced `ComparisonChains` test to demonstrate that the order of evaluations is done right and that side effects are maintained.
- Added `BuiltinOverload` test to show that overloaded built-ins are actually called and not the optimized version. So code like this has to print 2 lines:

```
from __builtin__ import len as _len

def len(x):
    print x

return _len(x)

print len(range(9))
```

Organisational

- Changed "README.txt" to no longer say that "Scons" is a requirement. Now that it's included (patched up to work with `ctypes` on Windows), we don't have to say that anymore.
- Documented the status of optimization and added some more ideas.
- There is now an option to dump the node tree after optimization as XML. Not currently use, but is for regression testing, to identify where new optimization and changes have an impact. This make it more feasible to be sure that Nuitka is only becoming better.
- Executable with Python3 again, although it won't do anything, the necessary code changes were done.

Summary

It's nice to see, that I some long standing issues were resolved, and that structural optimization has become almost a reality.

The difficult parts of exception propagation are all in place, now it's only details. With that we can eliminate and predict even more of the stupid code of "pybench" at compile time, achieving more infinite speedups.

Nuitka Release 0.3.9

This is about the new release of Nuitka which some bug fixes and offers a good speed improvement.

This new release is major milestone 2 work, enhancing practically all areas of Nuitka. The main focus was on faster function calls, faster class attributes (not instance), faster unpacking, and more built-ins detected and more thoroughly optimizing them.

Bug Fixes

- Exceptions raised inside with statements had references to the exception and traceback leaked.
- On Windows the binaries `sys.executable` pointed to the binary itself instead of the Python interpreter. Changed, because some code uses `sys.executable` to know how to start Python scripts.

- There is a bug (fixed in their repository) related to C++ raw strings and C++ "trigraphs" that affects Nuitka, added a workaround that makes Nuitka not emit "trigraphs" at all.
- The check for mutable constants was erroneous for tuples, which could lead to assuming a tuple with only mutable elements to be not mutable, which is of course wrong.

Optimization

This time there are so many new optimization, it makes sense to group them by the subject.

Exceptions

- The code to add a traceback is now our own, which made it possible to use frames that do not contain line numbers and a code object capable of lookups.
- Raising exceptions or adding to tracebacks has been made way faster by reusing a cached frame objects for the task.
- The class used for saving exceptions temporarily (e.g. used in `try/finally` code, or with statement) has been improved.

It now doesn't make a copy of the exception with a C++ `new` call, but it simply stores the exception properties itself and creates the exception object only on demand, which is more efficient.

- When catching exceptions, the addition of tracebacks is now done without exporting and re-importing the exception to Python, but directly on the exception objects traceback, this avoids a useless round trip.

Function Calls

- Uses of `PyObject_Call` provide `NULL` as the dictionary, instead of an empty dictionary, which is slightly faster for function calls.
- There are now dedicated variants for complex function calls with `*` and `**` arguments in all forms.
These can take advantage of easier cases. For example, a merge with star arguments is only needed if there actually were any of these.
- The check for non-string values in the `**` arguments can now be completely short-cut for the case of a dictionary that has never had a string added. There is now code that detects this case and skips the check, eliminating it as a performance concern.

Parameter Parsing

- Reversed the order in which parameters are checked.
Now the keyword dictionary is iterated first and only then the positional arguments after that is done. This iteration is not only much faster (avoiding repeated lookups for each possible parameter), it also can be more correct, in case the keyword argument is derived from a dictionary and its keys mutate it when being compared.
- Comparing parameter names is now done with a fast path, in which the pointer values are compared first. This can avoid a call to the comparison at all, which has become very likely due to the interning of parameter name strings, see below.
- Added a dedicated call to check for parameter equality with rich equality comparison, which doesn't raise an exception.
- Unpacking of tuples is now using dedicated variants of the normal unpacking code instead of rolling out everything themselves.

Attribute Access

- The class type (in executables, not yet for extension modules) is changed to a faster variant of our own making that doesn't consider the restricted mode a possibility. This avoids very expensive calls, and makes accessing class attributes in compiled code and in non-compiled code faster.
- Access to attributes (but not of instances) got in-lined and therefore much faster. Due to other optimization, a specific step to intern the string used for attribute access is not necessary with Nuitka at all anymore. This made access to attributes about 50% faster which is big of course.

Constants

- The bug for mutable tuples also caused non-mutable tuples to be considered as mutable, which lead to less efficient code.
- The constant creation with the g++ bug worked around, can now use raw strings to create string constants, without resorting to un-pickling them as a work around. This allows us to use `PyString_FromStringAndSize` to create strings again, which is obviously faster, and had not been done, because of the confusion caused by the g++ bug.
- For string constants that are usable as attributes (i.e. match the identifier regular expression), these are now interned, directly after creation. With this, the check for identical value of pointers for parameters has a bigger chance to succeed, and this saves some memory too.
- For empty containers (set, dict, list, tuple) the constants created are now are not unstreamed, but created with the dedicated API calls, saving a bit of code and being less ugly.
- For mutable empty constant access (set, dict, list) the values are no longer made by copying the constant, but instead with the API functions to create new ones. This makes code like `a = []` a tiny bit faster.
- For slice indices the code generation now takes advantage of creating a C++ `Py_ssize_t` from constant value if possible. Before it was converting the integer constant at run time, which was of course wasteful even if not (very) slow.

Iteration

- The creation of iterators got our own code. This avoids a function call and is otherwise only a small gain for anything but sequence iterators. These may be much faster to create now, as it avoids another call and repeated checks.
- The next on iterator got our own code too, which has simpler code flow, because it avoids the double check in case of NULL returned.
- The unpack check got similar code to the next iterator, it also has simpler code flow now and avoids double checks.

Built-ins

- Added support for the `list`, `tuple`, `dict`, `str`, `float` and `bool` built-ins along with optimizing their use with constant parameter.
- Added support for the `int` and `long` built-ins, based on a new "call spec" object, that detects parameter errors at compile time and raises appropriate exceptions as required, plus it deals with keyword arguments just as well.

So, to Nuitka it doesn't matter now if you write `int(value)` or `int(x = value)` anymore. The `base` parameter of these built-ins is also supported.

The use of this call spec mechanism will be expanded, currently it is not applied to the built-ins that take only one parameter. This is a work in progress as is the whole built-ins business as not all the built-ins are covered yet.

Cleanups

- In 0.3.8 per module global classes were introduced, but the `IMPORT_MODULE` kept using the old universal class, this got resolved and the old class is now fully gone.
- Using `assertObject` in more cases, and in more places at all, catches errors earlier on.
- Moved the addition to tracebacks into the `_PythonException` class, where it works directly on the contained traceback. This is cleaner as it no longer requires to export exceptions to Python, just to add a traceback entry.
- Some `PyLint` cleanups were done, reducing the number of reports a bit, but there is still a lot to do.
- Added a `DefaultValueIdentifier` class that encapsulates the access to default values in the parameter parsing more cleanly.
- The module `CodeTemplatesListContractions` was renamed to `CodeTemplatesContractions` to reflect the fact that it deals with all kinds of contractions (also set and dict contractions), not just list contractions.
- Moved the with related template to its own module `CodeTemplatesWith`, so its easier to find.
- The options handling for g++ based compilers was cleaned up, so that g++ 4.6 and MinGW are better supported now.
- Documented more aspects of the Scons build file.
- Some more generated code white space fixes.
- Moved some helpers to dedicated files. There is now `calling.hpp` for function calls, an `importing.cpp` for import related stuff.
- Moved the manifest generation to the scons file, which now produces ready to use executables.

New Tests

- Added a improved version of "pybench" that can cope with the "0 ms" execution time that Nuitka has for some if its sub-tests.
- Reference counting test for with statement was added.
- Micro benchmarks to demonstrate try finally performance when an exception travels through it.
- Micro benchmark for with statement that eats up exceptions raised inside the block.
- Micro benchmarks for the read and write access to class attributes.
- Enhanced `Printing` test to cover the trigraphs constant bug case. Output is required to make the error detectable.
- Enhanced `Constants` test to cover repeated mutation of mutable tuple constants, this covers the bug mentioned.

Organisational

- Added a credits section to the "README.txt" where I give credit to the people who contributed to Nuitka, and the projects it is using. I will make it a separate posting to cite these.
- Documented the requirements on the compiler more clearly, document the fact that we require scons and which version of Python (2.6 or 2.7).
- There is now a codespeed implementation up and running with historical data for up to Nuitka 0.3.8 runs of "PyStone" and with pybench. It will be updated for 0.3.9 once I have the infrastructure in place to do that automatically.
- The cleanup script now also removes .so files.

- The handling of options for g++ got improved, so it's the same for g++ and MinGW compilers, plus adequate errors messages are given, if the compiler version is too low.
- There is now a `--unstripped` option that just keeps the debug information in the file, but doesn't keep the assertions.

This will be helpful when looking at generated assembler code from Nuitka to not have the distortions that `--debug` causes (reduced optimization level, assertions, etc.) and instead a clear view.

Nuitka Release 0.3.8

This is to inform you about the new release of Nuitka with some real news and a slight performance increase. The significant news is added "Windows Support". You can now hope to run Nuitka on Windows too and have it produce working executables against either the standard Python distribution or a MinGW compiled Python.

There are still some small things to iron out, and clearly documentation needs to be created, and esp. the DLL hell problem of `msvcr90.dll` vs. `msvcrt.dll`, is not yet fully resolved, but appears to be not as harmful, at least not on native Windows.

I am thanking Khalid Abu Bakr for making this possible. I was surprised to see this happen. I clearly didn't make it easy. He found a good way around `ucontext`, identifier clashes, and a very tricky symbol problems where the CPython library under Windows exports less than under Linux. Thanks a whole lot.

Currently the Windows support is considered experimental and works with MinGW 4.5 or higher only.

Otherwise there have been the usual round of performance improvements and more cleanups. This release is otherwise milestone 2 work only, which will have to continue for some time more.

Bug Fixes

- Lambda generators were not fully compatible, their simple form could yield an extra value. The behavior for Python 2.6 and 2.7 is also different and Nuitka now mimics both correctly, depending on the used Python version
- The given parameter count cited in the error message in case of too many parameters, didn't include the given keyword parameters in the error message.
- There was an `assert False` right after warning about not found modules in the `--deep` mode, which was of course unnecessary.

Optimization

- When unpacking variables in assignments, the temporary variables are now held in a new temporary class that is designed for the task specifically.

This avoids the taking of a reference just because the `PyObjectTemporary` destructor insisted on releasing one. The new class `PyObjectTempHolder` hands the existing reference over and releases only in case of exceptions.

- When unpacking variable in for loops, the value from the iterator may be directly assigned, if it's to a variable.

In general this would be possible for every assignment target that cannot raise, but the infrastructure cannot tell yet, which these would be. This will improve with more milestone 3 work.

- Branches with only `pass` inside are removed, `pass` statements are removed before the code generation stage. This makes it easier to achieve and decide empty branches.
- There is now a global variable class per module. It appears that it is indeed faster to roll out a class per module accessing the `module *` rather than having one class and use a `module **`, which is quite disappointing from the C++ compiler.

- Also `MAKE_LIST` and `MAKE_TUPLE` have gained special cases for the 0 arguments case. Even when the size of the variadic template parameters should be known to the compiler, it seems, it wasn't eliminating the branch, so this was a speedup measured with valgrind.
- Empty tried branches are now replaced when possible with `try/except` statements, `try/finally` is simplified in this case. This gives a cleaner tree structure and less verbose C++ code which the compiler threw away, but was strange to have in the first place.
- In conditions the `or` and `and` were evaluated with Python objects instead of with C++ `bool`, which was unnecessary overhead.
- List contractions got more clever in how they assign from the iterator value.

It now uses a `PyObjectTemporary` if it's assigned to multiple values, a `PyObjectTempHolder` if it's only assigned once, to something that could raise, or a `PyObject *` if an exception cannot be raised. This avoids temporary references completely for the common case.

Cleanups

- The `if`, `for`, and `while` statements had always empty `else` nodes which were then also in the generated C++ code as empty branches. No harm to performance, but this got cleaned up.
- Some more generated code white space fixes.

New Tests

- The CPython 2.7 test suite now also has the `doctests` extracted to static tests, which improves test coverage for Nuitka again.

This was previously only done for CPython 2.6 test suite, but the test suites are different enough to make this useful, e.g. to discover newly changed behavior like with the lambda generators.

- Added Shed Skin 0.7.1 examples as benchmarks, so we can start to compare Nuitka performance in these tests. These will be the focus of numbers for the 0.4.x release series.
- Added a micro benchmark to check unpacking behavior. Some of these are needed to prove that a change is an actual improvement, when its effect can go under in noise of in-line vs. no in-line behavior of the C++ compiler.
- Added "pybench" benchmark which reveals that Nuitka is for some things much faster, but there are still fields to work on. This version needed changes to stand the speed of Nuitka. These will be subject of a later posting.

Organisational

- There is now a "tests/benchmarks/micro" directory to contain tiny benchmarks that just look at a single aspect, but have no other meaning, e.g. the "PyStone" extracts fall into this category.
- There is now a `--windows-target` option that attempts a cross-platform build on Linux to Windows executable. This is using "MingGW-cross-env" cross compilation tool chain. It's not yet working fully correctly due to the DLL hell problem with the C runtime. I hope to get this right in subsequent releases.
- The `--execute` option uses wine to execute the binary if it's a cross-compile for windows.
- Native windows build is recognized and handled with MinGW 4.5, the VC++ is not supported yet due to missing C++0x support.
- The basic test suite ran with Windows so far only and some adaptations were necessary. Windows new lines are now ignored in difference check, and addresses under Windows are upper case, small things.

Numbers

python 2.6:

```
Pystone(1.1) time for 50000 passes = 0.65  
This machine benchmarks at 76923.1 pystones/second
```

Nuitka 0.3.8 (driven by python 2.6):

```
Pystone(1.1) time for 50000 passes = 0.27  
This machine benchmarks at 185185 pystones/second
```

This is a 140% speed increase of 0.3.8 compared to CPython, up from 132% compared to the previous release.

Nuitka Release 0.3.7

This is about the new release with focus on performance and cleanups. It indicates significant progress with the milestone this release series really is about as it adds a `compiled_method` type.

So far functions, generator function, generator expressions were compiled objects, but in the context of classes, functions were wrapped in CPython `instancemethod` objects. The new `compiled_method` is specifically designed for wrapping `compiled_function` and therefore more efficient at it.

Bug Fixes

- When using `Python` or `Nuitka.py` to execute some script, the exit code in case of "file not found" was not the same as CPython. It should be 2, not 1.
- The exit code of the created programs (`--deep` mode) in case of an uncaught exception was 0, now it an error exit with value 1, like CPython does it.
- Exception tracebacks created inside `with` statements could contain duplicate lines, this was corrected.

Optimization

- Global variable assignments now also use `assign0` where no reference exists.
The assignment code for module variables is actually faster if it needs not drop the reference, but clearly the code shouldn't bother to take it on the outside just for that. This variant existed, but wasn't used as much so far.
- The instance method objects are now Nuitka's own compiled type too. This should make things slightly faster by itself.
- Our new compiled method objects support dedicated method parsing code, where `self` is passed directly, allowing to make calls taking a fast path in parameter parsing.
This avoids allocating/freeing a `tuple` object per method call, while reduced 3% ticks in "PyStone" benchmark, so that's significant.
- Solved a `TODO` of `BUILTIN_RANGE` to change it to pre-allocating the list in the final size as we normally do everywhere else. This was a tick reduction of 0.4% in "PyStone" benchmark, but the measurement method normalizes on loop speed, so it's not visible in the numbers output.
- Parameter variables cannot possibly be uninitialized at creation and most often they are never subject to a `del` statement. Adding dedicated C++ variable classes gave a big speedup, around 3% of "PyStone" benchmark ticks.

- Some abstract object operations were re-implemented, which allows to avoid function calls e.g. in the `ITERATOR_NEXT` case, this gave a few percent on "PyStone" as well.

Cleanups

- New package `nuitka.codegen` to contain all code generation related stuff, moved `nuitka.templates` to `nuitka.codegen.templates` as part of that.
- Inside the `nuitka.codegen` package the `MainControl` module now longer reaches into `Generator` for simple things, but goes through `CodeGeneration` for everything now.
- The `Generator` module uses almost no tree nodes anymore, but instead gets information passed in function calls. This allows for a cleanup of the interface towards `CodeGeneration`. Gives a cleaner view on the C++ code generation, and generally furthers the goal of other than C++ language backends.
- More "PyLint" work, many of the reported warnings have been addressed, but it's not yet happy.
- Defaults for `yield` and `return` are `None` and these values are now already added (as constants) during tree building so that no such special cases need to be dealt with in `CodeGeneration` and future analysis steps.
- Parameter parsing code has been unified even further, now the whole entry point is generated by one of the function in the new `nuitka.codegen.ParameterParsing` module.
- Split variable, exception, built-in helper classes into separate header files.

New Tests

- The exit codes of CPython execution and Nuitka compiled programs are now compared as well.
- Errors messages of methods are now covered by the `ParameterErrors` test as well.

Organisational

- A new script "benchmark.sh" (now called "run-valgrind.py") script now starts "kcache-grind" to display the valgrind result directly.

One can now use it to execute a test and inspect valgrind information right away, then improve it. Very useful to discover methods for improvements, test them, then refine some more.

- The "check-release.sh" script needs to unset `NUITKA_EXTRA_OPTIONS` or else the reflection test will trip over the changed output paths.

Numbers

python 2.6:

```
Pystone(1.1) time for 50000 passes = 0.65
This machine benchmarks at 76923.1 pystones/second
```

Nuitka 0.3.7 (driven by python 2.6):

```
Pystone(1.1) time for 50000 passes = 0.28
This machine benchmarks at 178571 pystones/second
```

This is a 132% speed of 0.3.7 compared to CPython, up from 109% compare to the previous release. This is a another small increase, that can be fully attributed to milestone 2 measures, i.e. not analysis, but purely more efficient C++ code generation and the new compiled method type.

One can now safely assume that it is at least twice as fast, but I will try and get the PyPy or Shedskin test suite to run as benchmarks to prove it.

No milestone 3 work in this release. I believe it's best to finish with milestone 2 first, because these are quite universal gains that we should have covered.

Nuitka Release 0.3.6

The major point this for this release is cleanup work, and generally bug fixes, esp. in the field of importing. This release cleans up many small open ends of Nuitka, closing quite a bunch of consistency `TODO` items, and then aims at cleaner structures internally, so optimization analysis shall become "easy". It is a correctness and framework release, not a performance improvement at all.

Bug Fixes

- Imports were not respecting the `level` yet. Code like this was not working, now it is:

```
from .. import something
```

- Absolute and relative imports were e.g. both tried all the time, now if you specify absolute or relative imports, it will be attempted in the same way than CPython does. This can make a difference with compatibility.
- Functions with a "locals dict" (using `locals` built-in or `exec` statement) were not 100% compatible in the way the locals dictionary was updated, this got fixed. It seems that directly updating a dict is not what CPython does at all, instead it only pushes things to the dictionary, when it believes it has to. Nuitka now does the same thing, making it faster and more compatible at the same time with these kind of corner cases.
- Nested packages didn't work, they do now. Nuitka itself is now successfully using nested packages (e.g. `nuitka.transform.optimizations`)

New Features

- The `--lto` option becomes usable. It's not measurably faster immediately, and it requires g++ 4.6 to be available, but then it at least creates smaller binaries and may provide more optimization in the future.

Optimization

- Exceptions raised by pre-computed built-ins, unpacking, etc. are now transformed to raising the exception statically.

Cleanups

- There is now a `getVariableForClosure` that a variable provider can use. Before that it guessed from `getVariableForReference` or `getVariableForAssignment` what might be the intention. This makes some corner cases easier.
- Classes, functions and lambdas now also have separate builder and body nodes, which enabled to make `getSameScopeNodes()` really simple. Either something has children which are all in a new scope or it has them in the same scope.
- Twisted workarounds like `TransitiveProvider` are no longer needed, because class builder and class body were separated.

- New packages `nuitka.transform.optimizations` and `nuitka.transform.finalizations`, where the first was `nuitka.optimizations` before. There is also code in `nuitka.transform` that was previously in a dedicated module. This allowed to move a lot of displaced code.
- `TreeBuilding` now has fast paths for all 3 forms, things that need a "provider", "node", and "source_ref"; things that need "node" and "source_ref"; things that need nothing at all, e.g. pass.
- Variables now avoid building duplicated instances, but instead share one. Better for analysis of them.

New Tests

- The Python 2.7 test suite is no longer run with Python 2.6 as it will just crash with the same exception all the time, there is no `importlib` in 2.6, but every test is using that through `test_support`.
- Nested packages are now covered with tests too.
- Imports of upper level packages are covered now too.

Organisational

- Updated the "README.txt" with the current plan on optimization.

Numbers

python 2.6:

```
Pystone(1.1) time for 50000 passes = 0.65
This machine benchmarks at 76923.1 pystones/second
```

Nuitka 0.3.6 (driven by python 2.6):

```
Pystone(1.1) time for 50000 passes = 0.31
This machine benchmarks at 161290 pystones/second
```

This is 109% for 0.3.6, but no change from the previous release. No surprise, because no new effective new optimization means have been implemented. Stay tuned for future release for actual progress.

Nuitka Release 0.3.5

This new release of Nuitka is an overall improvement on many fronts, there is no real focus this time, likely due to the long time it was in the making.

The major points are more optimization work, largely enhanced import handling and another improvement on the performance side. But there are also many bug fixes, more test coverage, usability and compatibility.

Something esp. noteworthy to me and valued is that many important changes were performed or at least triggered by Nicolas Dumazet, who contributed a lot of high quality commits as you can see from the gitweb history. He appears to try and compile Mercurial and Nuitka, and this resulted in important contributions.

Bug Fixes

- Nicolas found a reference counting bug with nested parameter calls. Where a function had parameters of the form `a, (b,c)` it could crash. This got fixed and covered with a reference count test.

- Another reference count problem when accessing the locals dictionary was corrected.
- Values `0.0` and `-0.0` were treated as the same. They are not though, they have a different sign that should not get lost.
- Nested contractions didn't work correctly, when the contraction was to iterate over another contraction which needs a closure. The problem was addressing by splitting the building of a contraction from the body of the contraction, so that these are now 2 nodes, making it easy for the closure handling to get things right.
- Global statements in function with local `exec()` would still use the value from the locals dictionary. Nuitka is now compatible to CPython with this too.
- Nicolas fixed problems with modules of the same name inside different packages. We now use the full name including parent package names for code generation and look-ups.
- The `__module__` attribute of classes was only set after the class was created. Now it is already available in the class body.
- The `__doc__` attribute of classes was not set at all. Now it is.
- The relative import inside nested packages now works correctly. With Nicolas moving all of Nuitka to a package, the compile itself exposed many weaknesses.
- A local re-raise of an exception didn't have the original line attached but the re-raise statement line.

New Features

- Modules and packages have been unified. Packages can now also have code in `"__init__.py"` and then it will be executed when the package is imported.
- Nicolas added the ability to create deep output directory structures without having to create them beforehand. This makes `--output-dir=some/deep/path` usable.
- Parallel build by Scons was added as an option and enabled by default, which enhances scalability for `--deep` compilations a lot.
- Nicolas enhanced the CPU count detection used for the parallel build. Turned out that `multithreading.cpu_count()` doesn't give us the number of available cores, so he contributed code to determine that.
- Support for upcoming g++ 4.6 has been added. The use of the new option `--lto` has been prepared, but right now it appears that the C++ compiler will need more fixes, before we can this feature with Nuitka.
- The `--display-tree` feature got an overhaul and now displays the node tree along with the source code. It puts the cursor on the line of the node you selected. Unfortunately I cannot get it to work two-way yet. I will ask for help with this in a separate posting as we can really use a "python-qt" expert it seems.
- Added meaningful error messages in the "file not found" case. Previously I just didn't care, but we sort of approach end user usability with this.

Optimization

- Added optimization for the built-in `range()` which otherwise requires a module and `builtin` module lookup, then parameter parsing. Now this is much faster with Nuitka and small ranges (less than 256 values) are converted to constants directly, avoiding run time overhead entirely.
- Code for re-raise statements now use a simple re-throw of the exception where possible, and only do the hard work where the re-throw is not inside an exception handler.
- Constant folding of operations and comparisons is now performed if the operands are constants.

- Values of some built-ins are pre-computed if the operands are constants.
- The value of module attribute `__name__` is replaced by a constant unless it is assigned to. This is the first sign of upcoming constant propagation, even if only a weak one.
- Conditional statement and/or their branches are eliminated where constant conditions allow it.

Cleanups

- Nicolas moved the Nuitka source code to its own `nuitka` package. That is going to make packaging it a lot easier and allows cleaner code.
- Nicolas introduced a fast path in the tree building which often delegates (or should do that) to a function. This reduced a lot of the dispatching code and highlights more clearly where such is missing right now.
- Together we worked on the line length issues of Nuitka. We agreed on a style and very long lines will vanish from Nuitka with time. Thanks for pushing me there.
- Nicolas also did provide many style fixes and general improvements, e.g. using `PyObjectTemporary` in more places in the C++ code, or not using `str.find` where `x in y` is a better choice.

- The node structure got cleaned up towards the direction that assignments always have an assignment as a child.

A function definition, or a class definition, are effectively assignments, and in order to not have to treat this as special cases everywhere, they need to have assignment targets as child nodes.

Without such changes, optimization will have to take too many things into account. This is not yet completed.

- Nicolas merged some node tree building functions that previously handled deletion and assigning differently, giving us better code reuse.
- The constants code generation was moved to a `__constants.cpp` where it doesn't make `__main__.cpp` so much harder to read anymore.
- The module declarations have been moved to their own header files.
- Nicolas cleaned up the scripts used to test Nuitka big time, removing repetitive code and improving the logic. Very much appreciated.
- Nicolas also documented a things in the Nuitka source code or got me to document things that looked strange, but have reasons behind it.
- Nicolas solved the `TODO` related to built-in module accesses. These will now be way faster than before.
- Nicolas also solved the `TODO` related to the performance of "locals dict" variable accesses.
- `Generator.py` no longer contains classes. The `Contexts` objects are supposed to contain the state, and as such the generator objects never made much sense.
- Also with the help of Scons community, I figured out how to avoid having object files inside the `src` directory of Nuitka. That should also help packaging, now all build products go to the `.build` directory as they should.
- The vertical white space of the generated C++ got a few cleanups, trailing/leading new line is more consistent now, and there were some assertions added that it doesn't happen.

New Tests

- The CPython 2.6 tests are now also run by CPython 2.7 and the other way around and need to report the same test failure reports, which found a couple of issues.

- Now the test suite is run with and without `--debug` mode.
- Basic tests got extended to cover more topics and catch more issues.
- Program tests got extended to cover code in packages.
- Added more exec scope tests. Currently inlining of exec statements is disabled though, because it requires entirely different rules to be done right, it has been pushed back to the next release.

Organisational

- The `g++-nuitka` script is no more. With the help of the Scons community, this is now performed inside the scons and only once instead of each time for every C++ file.
- When using `--debug`, the generated C++ is compiled with `-Wall` and `-Werror` so that some form of bugs in the generated C++ code will be detected immediately. This found a few issues already.
- There is a new git merge policy in place. Basically it says, that if you submit me a pull request, that I will deal with it before publishing anything new, so you can rely on the current git to provide you a good base to work on. I am doing more frequent pre-releases already and I would like to merge from your git.
- The "README.txt" was updated to reflect current optimization status and plans. There is still a lot to do before constant propagation can work, but this explains things a bit better now. I hope to expand this more and more with time.
- There is now a "misc/clean-up.sh" script that prints the commands to erase all the temporary files sticking around in the source tree.

That is for you if you like me, have other directories inside, ignored, that you don't want to delete.

- Then there is now a script that prints all source filenames, so you can more easily open them all in your editor.
- And very important, there is now a "check-release.sh" script that performs all the tests I think should be done before making a release.
- Pylint got more happy with the current Nuitka source. In some places, I added comments where rules should be granted exceptions.

Numbers

python 2.6:

```
Pystone(1.1) time for 50000 passes = 0.65
This machine benchmarks at 76923.1 pystones/second
```

Nuitka 0.3.5 (driven by python 2.6):

```
Pystone(1.1) time for 50000 passes = 0.31
This machine benchmarks at 161290 pystones/second
```

This is 109% for 0.3.5, up from 91% before.

Overall this release is primarily an improvement in the domain of compatibility and contains important bug and feature fixes to the users. The optimization framework only makes a first showing of with the framework to organize them. There is still work to do to migrate optimization previously present

It will take more time before we will see effect from these. I believe that even more cleanups of `TreeBuilding`, `Nodes` and `CodeGeneration` will be required, before everything is in place for the big jump in performance numbers. But still, passing 100% feels good. Time to rejoice.

Nuitka Release 0.3.4

This new release of Nuitka has a focus on re-organizing the Nuitka generated source code and a modest improvement on the performance side.

For a long time now, Nuitka has generated a single C++ file and asked the C++ compiler to translate it to an executable or shared library for CPython to load. This was done even when embedding many modules into one (the "deep" compilation mode, option `--deep`).

This was simple to do and in theory ought to allow the compiler to do the most optimization. But for large programs, the resulting source code could have exponential compile time behavior in the C++ compiler. At least for the GNU g++ this was the case, others probably as well. This is of course at the end a scalability issue of Nuitka, which now has been addressed.

So the major advancement of this release is to make the `--deep` option useful. But also there have been a performance improvements, which end up giving us another boost for the "PyStone" benchmark.

Bug Fixes

- Imports of modules local to packages now work correctly, closing the small compatibility gap that was there.
- Modules with a "-" in their name are allowed in CPython through dynamic imports. This lead to wrong C++ code created. (Thanks to Li Xuan Ji for reporting and submitting a patch to fix it.)
- There were warnings about wrong format used for `Ssize_t` type of CPython. (Again, thanks to Li Xuan Ji for reporting and submitting the patch to fix it.)
- When a wrong exception type is raised, the traceback should still be the one of the original one.
- Set and dict contractions (Python 2.7 features) declared local variables for global variables used. This went unnoticed, because list contractions don't generate code for local variables at all, as they cannot have such.
- Using the `type()` built-in to create a new class could attribute it to the wrong module, this is now corrected.

New Features

- Uses Scons to execute the actual C++ build, giving some immediate improvements.
- Now caches build results and Scons will only rebuild as needed.
- The direct use of `__import__()` with a constant module name as parameter is also followed in "deep" mode. With time, non-constants may still become predictable, right now it must be a real CPython constant string.

Optimization

- Added optimization for the built-ins `ord()` and `chr()`, these require a module and built-in module lookup, then parameter parsing. Now these are really quick with Nuitka.
- Added optimization for the `type()` built-in with one parameter. As above, using from builtin module can be very slow. Now it is instantaneous.
- Added optimization for the `type()` built-in with three parameters. It's rarely used, but providing our own variant, allowed to fix the bug mentioned above.

Cleanups

- Using `scons` is a big cleanup for the way how C++ compiler related options are applied. It also makes it easier to re-build without Nuitka, e.g. if you were using Nuitka in your packages, you can easily build in the same way than Nuitka does.
- Static helpers source code has been moved to ".hpp" and ".cpp" files, instead of being in ".py" files. This makes C++ compiler messages more readable and allows us to use C++ mode in Emacs etc., making it easier to write things.
- Generated code for each module ends up in a separate file per module or package.
- Constants etc. go to their own file (although not named sensible yet, likely going to change too)
- Module variables are now created by the `CPythonModule` node only and are unique, this is to make optimization of these feasible. This is a pre-step to module variable optimization.

New Tests

- Added "ExtremeClosure" from my Python quiz, it was not covered by existing tests.
- Added test case for program that imports a module with a dash in its name.
- Added test case for main program that starts with a dash.
- Extended the built-in tests to cover `type()` as well.

Organisational

- There is now a new environment variable `NUITKA_SCONS` which should point to the directory with the `SingleExe.scons` file for Nuitka. The `scons` file could be named better, because it is actually one and the same who builds extension modules and executables.
- There is now a new environment variable `NUITKA_CPP` which should point to the directory with the C++ helper code of Nuitka.
- The script "create-environment.sh" can now be sourced (if you are in the top level directory of Nuitka) or be used with `eval`. In either case it also reports what it does.

Update

The script has become obsolete now, as the environment variables are no longer necessary.

- To cleanup the many "Program.build" directories, there is now a "clean-up.sh" script for your use. Can be handy, but if you use git, you may prefer its `clean` command.

Update

The script has become obsolete now, as Nuitka test executions now by default delete the build results.

Numbers

python 2.6:

```
Pystone(1.1) time for 50000 passes = 0.65  
This machine benchmarks at 76923.1 pystones/second
```

Nuitka 0.3.4:

```
Pystone(1.1) time for 50000 passes = 0.34  
This machine benchmarks at 147059 pystones/second
```

This is 91% for 0.3.4, up from 80% before.

Nuitka Release 0.3.3

This release of Nuitka continues the focus on performance. It also cleans up a few open topics. One is "doctests", these are now extracted from the CPython 2.6 test suite more completely. The other is that the CPython 2.7 test suite is now passed completely. There is some more work ahead though, to extract all of the "doctests" and to do that for both versions of the tests.

This means an even higher level of compatibility has been achieved, then there is performance improvements, and ever cleaner structure.

Bug Fixes

Generators

- Generator functions tracked references to the common and the instance context independently, now the common context is not released before the instance contexts are.
- Generator functions didn't check the arguments to `throw()` the way they are in CPython, now they do.
- Generator functions didn't trace exceptions to "stderr" if they occurred while closing unfinished ones in "del".
- Generator functions used the slightly different wordings for some error messages.

Function Calls

- Extended call syntax with `**` allows that to use a mapping, and it is now checked if it really is a mapping and if the contents has string keys.
- Similarly, extended call syntax with `*` allows a sequence, it is now checked if it really is a sequence.
- Error message for duplicate keyword arguments or too little arguments now describe the duplicate parameter and the callable the same way CPython does.
- Now checks to the keyword argument list first before considering the parameter counts. This is slower in the error case, but more compatible with CPython.

Classes

- The "locals()" built-in when used in the class scope (not in a method) now is correctly writable and writes to it change the resulting class.
- Name mangling for private identifiers was not always done entirely correct.

Others

- Exceptions didn't always have the correct stack reported.

- The pickling of some tuples showed that "cPickle" can have non-reproducible results, using "pickle" to stream constants now

Optimization

- Access to instance attributes has become faster by writing specific code for the case. This is done in JIT way, attempting at run time to optimize attribute access for instances.
- Assignments now often consider what's cheaper for the other side, instead of taking a reference to a global variable, just to have to release it.
- The function call code built argument tuples and dictionaries as constants, now that is true for every tuple usage.

Cleanups

- The static helper classes, and the prelude code needed have been moved to separate C++ files and are now accessed "#include". This makes the code inside C++ files as opposed to a Python string and therefore easier to read and or change.

New Features

- The generator functions and generator expressions have the attribute "gi_running" now. These indicate if they are currently running.

New Tests

- The script to extract the "doctests" from the CPython test suite has been rewritten entirely and works with more doctests now. Running these tests created increased the test coverage a lot.
- The Python 2.7 test suite has been added.

Organisational

- One can now run multiple "compare_with_cpython" instances in parallel, which enables background test runs.
- There is now a new environment variable "NUITKA_INCLUDE" which needs to point to the directory Nuitka's C++ includes live in. Of course the "create-environment.sh" script generates that for you easily.

Numbers

python 2.6:

```
Pystone(1.1) time for 50000 passes = 0.65  
This machine benchmarks at 76923.1 pystones/second
```

Nuitka 0.3.3:

```
Pystone(1.1) time for 50000 passes = 0.36  
This machine benchmarks at 138889 pystones/second
```

This is 80% for 0.3.3, up from 66% before.

Nuitka Release 0.3.2

This release of Nuitka continues the focus on performance. But this release also revisits the topic of feature parity. Before, feature parity had been reached "only" with Python 2.6. This is of course a big thing, but you know there is always more, e.g. Python 2.7.

With the addition of set contractions and dict contractions in this very release, Nuitka is approaching Python support for 2.7, and then there are some bug fixes.

Bug Fixes

- Calling a function with `**` and using a non-dict for it was leading to wrong behavior.
Now a mapping is good enough as input for the `**` parameter and it's checked.
- Deeply nested packages "package.subpackage.module" were not found and gave a warning from Nuitka, with the consequence that they were not embedded in the executable. They now are.
- Some error messages for wrong parameters didn't match literally. For example "function got multiple..." as opposed to "function() got multiple..." and alike.
- Files that ended in line with a "#" but without a new line gave an error from "ast.parse". As a workaround, a new line is added to the end of the file if it's "missing".
- More correct exception locations for complex code lines. I noted that the current line indication should not only be restored when the call at hand failed, but in any case. Otherwise sometimes the exception stack would not be correct. It now is - more often. Right now, this has no systematic test.
- Re-raised exceptions didn't appear on the stack if caught inside the same function, these are now correct.
- For `exec` the `globals` argument needs to have "`__builtins__`" added, but the check was performed with the mapping interface.
That is not how CPython does it, and so e.g. the mapping could use a default value for "`__builtins__`" which could lead to incorrect behavior. Clearly a corner case, but one that works fully compatible now.

Optimization

- The local and shared local variable C++ classes have a flag "free_value" to indicate if an "PY_DECREF" needs to be done when releasing the object. But still the code used "Py_XDECREF" (which allows for "NULL" values to be ignored.) when the releasing of the object was done. Now the inconsistency of using "NULL" as "object" value with "free_value" set to true was removed.
- Tuple constants were copied before using them without a point. They are immutable anyway.

Cleanups

- Improved more of the indentation of the generated C++ which was not very good for contractions so far. Now it is. Also assignments should be better now.
- The generation of code for contractions was made more general and templates split into multiple parts. This enabled reuse of the code for list contractions in dictionary and set contractions.
- The with statement has its own template now and got cleaned up regarding indentation.

New Tests

- There is now a script to extract the "doctests" from the CPython test suite and it generates Python source code from them. This can be compiled with Nuitka and output compared to CPython. Without this, the doctest parts of the CPython test suite is mostly useless. Solving this improved test coverage, leading to many small fixes. I will dedicate a later posting to the tool, maybe it is useful in other contexts as well.
- Reference count tests have been expanded to cover assignment to multiple assignment targets, and to attributes.
- The deep program test case, now also have a module in a sub-package to cover this case as well.

Organisational

- The [gitweb interface](#) (since disabled) might be considered an alternative to downloading the source if you want to provide a pointer, or want to take a quick glance at the source code. You can already download with git, follow the link below to the page explaining it.
- The "README.txt" has documented more of the differences and I consequently updated the Differences page. There is now a distinction between generally missing functionality and things that don't work in `--deep` mode, where Nuitka is supposed to create one executable.

I will make it a priority to remove the (minor) issues of `--deep` mode in the next release, as this is only relatively little work, and not a good difference to have. We want these to be empty, right? But for the time being, I document the known differences there.

Numbers

python 2.6:

```
Pystone(1.1) time for 50000 passes = 0.65
This machine benchmarks at 76923.1 pystones/second
```

Nuitka 0.3.2:

```
Pystone(1.1) time for 50000 passes = 0.39
This machine benchmarks at 128205 pystones/second
```

This is 66% for 0.3.2, slightly up from the 58% of 0.3.1 before. The optimization done were somewhat fruitful, but as you can see, they were also more cleanups, not the big things.

Nuitka Release 0.3.1

This release of Nuitka continues the focus on performance and contains only cleanups and optimization. Most go into the direction of more readable code, some aim at making the basic things faster, with good results as to performance as you can see below.

Optimization

- Constants in conditions of conditional expressions (`a if cond else d`), `if/elif` or `while` are now evaluated to `true` or `false` directly. Before there would be temporary python object created from it which was then checked if it had a truth value.

All of that is obviously overhead only. And it hurts the typically `while 1:` infinite loop case badly.

- Do not generate code to catch `BreakException` or `ContinueException` unless a `break` or `continue` statement being in a `try: finally:` block inside that loop actually require this. Even while uncaught exceptions are cheap, it is still an improvement worthwhile and it clearly improves the readability for the normal case.

- The compiler more aggressively prepares tuples, lists and dicts from the source code as constants if their contents is "immutable" instead of building at run time. An example of a "mutable" tuple would be ({},) which is not safe to share, and therefore will still be built at run time.

For dictionaries and lists, copies will be made, under the assumption that copying a dictionary will always be faster, than making it from scratch.

- The parameter parsing code was dynamically building the tuple of argument names to check if an argument name was allowed by checking the equivalent of `name in argument_names`. This was of course wasteful and now a pre-built constant is used for this, so it should be much faster to call functions with keyword arguments.
- There are new templates files and also actual templates now for the `while` and `for` loop code generation. And I started work on having a template for assignments.

Cleanups

- Do not generate code for the else of `while` and `for` loops if there is no such branch. This uncluttered the generated code somewhat.
- The indentation of the generated C++ was not very good and whitespace was often trailing, or e.g. a real tab was used instead of "t". Some things didn't play well together here.

Now much of the generated C++ code is much more readable and white space cleaner. For optimization to be done, the humans need to be able to read the generated code too. Mind you, the aim is not to produce usable C++, but on the other hand, it must be possible to understand it.

- To the same end of readability, the empty `else {}` branches are avoided for `if`, `while` and `for` loops. While the C++ compiler can be expected to remove these, they seriously cluttered up things.
- The constant management code in `Context` was largely simplified. Now the code is using the `Constant` class to find its way around the problem that dicts, sets, etc. are not hashable, or that `complex` is not being ordered; this was necessary to allow deeply nested constants, but it is also a simpler code now.
- The C++ code generated for functions now has two entry points, one for Python calls (arguments as a list and dictionary for parsing) and one where this has happened successfully. In the future this should allow for faster function calls avoiding the building of argument tuples and dictionaries all-together.
- For every function there was a "traceback adder" which was only used in the C++ exception handling before exit to CPython to add to the traceback object. This was now in-lined, as it won't be shared ever.

Numbers

python 2.6:

```
Pystone(1.1) time for 50000 passes = 0.65
This machine benchmarks at 76923.1 pystones/second
```

Nuitka 0.3.1:

```
Pystone(1.1) time for 50000 passes = 0.41
This machine benchmarks at 121951 pystones/second
```

This is 58% for 0.3.1, up from the 25% before. So it's getting somewhere. As always you will find its latest version [here](#).

Nuitka Release 0.3.0

This release 0.3.0 is the first release to focus on performance. In the 0.2.x series Nuitka achieved feature parity with CPython 2.6 and that was very important, but now it is time to make it really useful.

Optimization has been one of the main points, although I was also a bit forward looking to Python 2.7 language constructs. This release is the first where I really started to measure things and removed the most important bottlenecks.

New Features

- Added option to control `--debug`. With this option the C++ debug information is present in the file, otherwise it is not. This will give much smaller ".so" and ".exe" files than before.
- Added option `--no-optimization` to disable all optimization.
It enables C++ asserts and compiles with less aggressive C++ compiler optimization, so it can be used for debugging purposes.
- Support for Python 2.7 set literals has been added.

Performance Enhancements

- Fast global variables: Reads of global variables were fast already. This was due to a trick that is now also used to check them and to do a much quicker update if they are already set.
- Fast `break/continue` statements: To make sure these statements execute the finally handlers if inside a try, these used C++ exceptions that were caught by `try/finally` in while or for loops.
This was very slow and had very bad performance. Now it is checked if this is at all necessary and then it's only done for the rare case where a `break/continue` really is inside the tried block. Otherwise it is now translated to a C++ `break/continue` which the C++ compiler handles more efficiently.
- Added `unlikely()` compiler hints to all errors handling cases to allow the C++ compiler to generate more efficient branch code.
- The for loop code was using an exception handler to make sure the iterated value was released, using `PyObjectTemporary` for that instead now, which should lead to better generated code.
- Using constant dictionaries and copy from them instead of building them at run time even when contents was constant.

New Tests

- Merged some bits from the CPython 2.7 test suite that do not harm 2.6, but generally it's a lot due to some `unittest` module interface changes.
- Added CPython 2.7 tests `test_dictcomps.py` and `test_dictviews.py` which both pass when using Python 2.7.
- Added another benchmark extract from "PyStone" which uses a while loop with break.

Numbers

python 2.6:

```
Pystone(1.1) time for 50000 passes = 0.65
This machine benchmarks at 76923.1 pystones/second
```

Nuitka 0.3.0:

```
Pystone(1.1) time for 50000 passes = 0.52  
This machine benchmarks at 96153.8 pystones/second
```

That's a 25% speedup now and a good start clearly. It's not yet in the range of where i want it to be, but there is always room for more. And the `break/continue` exception was an important performance regression fix.

Nuitka Release 0.2.4

This release 0.2.4 is likely the last 0.2.x release, as it's the one that achieved feature parity with CPython 2.6, which was the whole point of the release series, so time to celebrate. I have stayed away (mostly) from any optimization, so as to not be premature.

From now on speed optimization is going to be the focus though. Because right now, frankly, there is not much of a point to use Nuitka yet, with only a minor run time speed gain in trade for a long compile time. But hopefully we can change that quickly now.

New Features

- The use of `exec` in a local function now adds local variables to scope it is in.
- The same applies to `from module_name import *` which is now compiled correctly and adds variables to the local variables.

Bug Fixes

- Raises `UnboundLocalError` when deleting a local variable with `del` twice.
- Raises `NameError` when deleting a global variable with `del` twice.
- Read of to uninitialized closure variables gave `NameError`, but `UnboundLocalError` is correct and raised now.

Cleanups

- There is now a dedicated pass over the node tree right before code generation starts, so that some analysis can be done as late as that. Currently this is used for determining which functions should have a dictionary of locals.
- Checking the exported symbols list, fixed all the cases where a `static` was missing. This reduces the "module.so" sizes.
- With `gcc` the "visibility=hidden" is used to avoid exporting the helper classes. Also reduces the "module.so" sizes, because classes cannot be made static otherwise.

New Tests

- Added "DoubleDeletions" to cover behaviour of `del`. It seems that this is not part of the CPython test suite.
- The "OverflowFunctions" (those with dynamic local variables) now has an interesting test, `exec` on a local scope, effectively adding a local variable while a closure variable is still accessible, and a module variable too. This is also not in the CPython test suite.
- Restored the parts of the CPython test suite that did local star imports or `exec` to provide new variables. Previously these have been removed.

- Also "test_with.py" which covers PEP 343 has been reactivated, the with statement works as expected.

Nuitka Release 0.2.3

This new release is marking a closing in on feature parity to CPython 2.6 which is an important mile stone. Once this is reached, a "Nuitka 0.3.x" series will strive for performance.

Bug Fixes

- Generator functions no longer leak references when started, but not finished.
- Yield can in fact be used as an expression and returns values that the generator user `send()` to it.

Reduced Differences / New Features

- Generator functions already worked quite fine, but now they have the `throw()`, `send()` and `close()` methods.
- Yield is now an expression as is ought to be, it returns values put in by `send()` on the generator user.
- Support for extended slices:

```
x = d[:42, ..., :24:, 24, 100]
d[:42, ..., :24:, 24, 100] = "Strange"
del d[:42, ..., :24:, 24, 100]
```

Tests Work

- The "test_contextlib" is now working perfectly due to the generator functions having a correct `throw()`. Added that test back, so context managers are now fully covered.
- Added a basic test for "overflow functions" has been added, these are the ones which have an unknown number of locals due to the use of language constructs `exec` or `from bla import *` on the function level. This one currently only highlights the failure to support it.
- Reverted removals of extended slice syntax from some parts of the CPython test suite.

Cleanups

- The compiled generator types are using the new C++0x type safe enums feature.
- Resolved a circular dependency between `TreeBuilding` and `TreeTransforming` modules.

Nuitka Release 0.2.2

This is some significant progress, a lot of important things were addressed.

Bug Fixes

- Scope analysis is now done during the tree building instead of sometimes during code generation, this fixed a few issues that didn't show up in tests previously.
- Reference leaks of generator expressions that were not fishing, but then deleted are not more.
- Inlining of `exec` is more correct now.

- More accurate exception lines when iterator creation executes compiled code, e.g. in a for loop
- The list of base classes of a class was evaluated in the context of the class, now it is done in the context of the containing scope.
- The first iterated of a generator expression was evaluated in its own context, now it is done in the context of the containing scope.

Reduced Differences

- With the enhanced scope analysis, `UnboundLocalError` is now correctly supported.
- Generator expressions (but not yet functions) have a `throw()`, `send()` and `close()` method.
- `Exec` can now write to local function namespace even if `None` is provided at run time.
- Relative imports inside packages are now correctly resolved at compile time when using `--deep`.

Cleanups

- The compiled function type got further enhanced and cleaned up.
- The compiled generator expression function type lead to a massive cleanup of the code for generator expressions.
- Cleaned up namespaces, was still using old names, or "Py*" which is reserved to core CPython.
- Overhaul of the code responsible for `eval` and `exec`, it has been split, and it pushed the detection defaults to the C++ compiler which means, we can do it at run time or compile time, depending on circumstances.
- Made `PyTemporaryObject` safer to use, disabling copy constructor it should be also a relief to the C++ compiler if it doesn't have to eliminate all its uses.
- The way delayed work is handled in `TreeBuilding` step has been changed to use closed functions, should be more readable.
- Some more code templates have been created, making the code generation more readable in some parts. More to come.

New Features

- As I start to consider announcing Nuitka, I moved the version logic so that the version can now be queried with `--version`.

Optimization

- Name lookups for `None`, `True` and `False` and now always detected as constants, eliminating many useless module variable lookups.

New Tests

- More complete test of generator expressions.
- Added test program for packages with relative imports inside the package.
- The built-in `dir()` in a function was not having fully deterministic output list, now it does.

Summary

Overall, the amount of differences between CPython and Nuitka is heading towards zero. Also most of the improvements done in this release were very straightforward cleanups and not much work was required, mostly things are about cleanups and then it becomes easily right. The new type for the compiled generator expressions was simple to create, esp. as I could check what CPython does in its source code.

For optimization purposes, I decided that generator expressions and generator functions will be separate compiled types, as most of their behavior will not be shared. I believe optimizing generator expressions to run well is an important enough goal to warrant that they have their own implementation. Now that this is done, I will repeat it with generator functions.

Generator functions already work quite fine, but like generator expressions did before this release, they can leak references if not finished, and they don't have the `throw()` method, which seems very important to the correct operation of `contextlib`. So I will introduce a dedicated type for these too, possibly in the next release.

Nuitka Release 0.2.1

The march goes on, this is another minor release with a bunch of substantial improvements:

Bug Fixes

- Packages now also can be embedded with the `--deep` option too, before they could not be imported from the executable.
- In-lined `exec` with their own future statements leaked these to the surrounding code.

Reduced Differences

- The future `print` function import is now supported too.

Cleanups

- Independence of the compiled function type. When I started it was merely `PyCFunction` and then a copy of it patched at run time, using increasingly less code from CPython. Now it's nothing at all anymore.
- This lead to major cleanup of run time compiled function creation code, no more `methoddefs`, `PyCObject` holding context, etc.
- PyLint was used to find the more important style issues and potential bugs, also helping to identify some dead code.

Summary

The major difference now is the lack of a `throw` method for generator functions. I will try to address that in a 0.2.2 release if possible. The plan is that the 0.2.x series will complete these tasks, and 0.3 could aim at some basic optimization finally.

Nuitka Release 0.2

Good day, this is a major step ahead, improvements everywhere.

Bug Fixes

- Migrated the Python parser from the deprecated and problematic `compiler` module to the `ast` module which fixes the `d[a,] = b` parser problem. A pity it was not available at the time I started, but the migration was relatively painless now.
- I found and fixed wrong encoding of binary data into C++ literals. Now Nuitka uses C++0x raw strings, and these problems are gone.
- The decoding of constants was done with the `marshal` module, but that appears to not deeply care enough about unicode encoding it seems. Using `cPickle` now, which seems less efficient, but is more correct.
- Another difference is gone: The `continue` and `break` inside loops do no longer prevent the execution of finally blocks inside the loop.

Organisational

- I now maintain the "README.txt" in org-mode, and intend to use it as the issue tracker, but I am still a beginner at that.

Update

Turned out I never mastered it, and used ReStructured Text instead.

- There is a public git repository for you to track Nuitka releases. Make your changes and then `git pull --rebase`. If you encounter conflicts in things you consider useful, please submit the patches and a pull request. When you make your clones of Nuitka public, use `nuitka-unofficial` or not the name `Nuitka` at all.
- There is a now a mailing list (since closed).

Reduced Differences

- Did you know you could write `lambda : (yield something)` and it gives you a lambda that creates a generator that produces that one value? Well, now Nuitka has support for lambda generator functions.
- The `from __future__ import division` statement works as expected now, leading to some newly passing CPython tests.
- Same for `from __future__ import unicode_literals` statement, these work as expected now, removing many differences in the CPython tests that use this already.

New Features

- The `Python` binary provided and `Nuitka.py` are now capable of accepting parameters for the program executed, in order to make it even more of a drop-in replacement to `python`.
- Inlining of `exec` statements with constant expressions. These are now compiled at compile time, not at run time anymore. I observed that an increasing number of CPython tests use `exec` to do things in isolation or to avoid warnings, and many more these tests will now be more effective. I intend to do the same with `eval` expressions too, probably in a minor release.

Summary

So give it a whirl. I consider it to be substantially better than before, and the list of differences to CPython is getting small enough, plus there is already a fair bit of polish to it. Just watch out that it needs gcc-4.5 or higher now.

Nuitka Release 0.1.1

I just have just updated Nuitka to version 0.1.1 which is a bug fix release to 0.1, which corrects many of the small things:

- Updated the CPython test suite to 2.6.6rc and minimized much of existing differences in the course.
- Compiles standalone executable that includes modules (with `--deep` option), but packages are not yet included successfully.
- Reference leaks with exceptions are no more.
- `sys.exc_info()` works now mostly as expected (it's not a stack of exceptions).
- More readable generated code, better organisation of C++ template code.
- Restored debug option `--g++-only`.

The biggest thing probably is the progress with exception tracebacks objects in exception handlers, which were not there before (always `None`). Having these in place will make it much more compatible. Also with manually raised exceptions and assertions, tracebacks will now be more correct to the line.

On a bad news, I discovered that the `compiler` module that I use to create the AST from Python source code, is not only deprecated, but also broken. I created the [CPython bug](#) about it, basically it cannot distinguish some code of the form `d[1,] = None` from `d[1] = None`. This will require a migration of the `ast` module, which should not be too challenging, but will take some time.

I am aiming at it for a 0.2 release. Generating wrong code (Nuitka sees `d[1] = None` in both cases) is a show blocker and needs a solution.

So, yeah. It's better, it's there, but still experimental. You will find its latest version here. Please try it out and let me know what you think in the comments section.

Nuitka Release 0.1 (Releasing Nuitka to the World)

Obviously this is very exciting step for me. I am releasing Nuitka today. Finally. For a long time I knew I would, but actually doing it, is a different beast. Reaching my goals for release turned out to be less far away than I hope, so instead of end of August, I can already release it now.

Currently it's not more than 4% faster than CPython. No surprise there, if all you did, is removing the bytecode interpretation so far. It's not impressive at all. It's not even a reason to use it. But it's also only a start. Clearly, once I get into optimizing the code generation of Nuitka, it will only get better, and then probably in sometimes dramatic steps. But I see this as a long term goal.

I want to have infrastructure in the code place, before doing lots of possible optimization that just make Nuitka unmaintainable. And I will want to have a look at what others did so far in the domain of type inference and how to apply that for my project.

I look forward to the reactions about getting this far. The supported language volume is amazing, and I have a set of nice tricks used. For example the way generator functions are done is a clever hack.

Where to go from here? Well, I guess, I am going to judge it by the feedback I receive. I personally see "constant propagation" as a laudable first low hanging fruit, that could be solved.

Consider this readable code on the module level:

```
meters_per_nautical_mile = 1852

def convertMetersToNauticalMiles(meters):
    return meters / meters_per_nautical_mile

def convertNauticalMilesToMeters(miles):
    return miles * meters_per_nautical_mile
```

Now imagine you are using this very frequently in code. Quickly you determine that the following will be much faster:

```
def convertMetersToNauticalMiles(meters):
    return meters / 1852

def convertNauticalMilesToMeters(miles):
    return miles * 1852
```

Still good? Well, probably next step you are going to in-line the function calls entirely. For optimization, you are making your code less readable. I do not all appreciate that. My first goal is there to make the more readable code perform as well or better as the less readable variant.