



Modern ChimeraX

Zach Pearson
RBVI General Meeting
26 May 2022

Background



What is ChimeraX, really?

Four Pillars of the ChimeraX Architecture:

- An isolated, internally consistent distribution of Python...
- Including a namespace package for research computing distributed with that runtime...
- Distributed alongside useful 3rd party binaries...
- Launched by a small C binary that presents a friendly facade to end users



What is ChimeraX, really?

- An isolated, internally consistent distribution of Python...
- Including a namespace package for research computing distributed with that runtime...
- Distributed alongside useful 3rd party binaries...
- Launched by a small C binary that presents a friendly facade to end users

This is a fundamentally good architecture; this presentation does not argue for changing it.



What is ChimeraX, really?

- An isolated, internally consistent distribution of Python...
- Including a namespace package for research computing distributed with that runtime...
- Distributed alongside useful 3rd party binaries..
- Launched by a small C binary that presents a friendly facade to end users

This is a fundamentally good architecture; this presentation does not argue for changing it.

Today I want to focus on the namespace package



What is ChimeraX, really?

- An isolated, internally consistent distribution of Python...
- Including a namespace package for research computing distributed with that runtime...
- Distributed alongside useful 3rd party binaries..
- Launched by a small C binary that presents a friendly facade to end users

This is a fundamentally good architecture; this presentation does not argue for changing it.

Today I want to focus on the namespace package, and the C binary that calls it.



What is ChimeraX, really?

- An isolated, internally consistent distribution of Python...
- Including a namespace package for research computing distributed with that runtime...
- Distributed alongside useful 3rd party binaries..
- Launched by a small C binary that presents a friendly facade to end users

This is a fundamentally good architecture; this presentation does not argue for changing it.

Today I want to focus on the namespace package, and the C binary that calls it.

Ultimate goal: merge `native-packaging` into `develop` as the basis for 1.5

Background: Environment



1.5 Objectives

- ChimeraX in PyPi ([Ticket #1470](#)) – 3 years old
- Report problematic version number checking to pip ([Ticket #3403](#)) – 2 years old
- Automate uploading wheels to the Toolshed ([Ticket #3058](#)) – 2 years old
- Move pip dependencies into bundles ([Ticket #3890](#)) – 19 months old
- Make native Apple M1 CPU distribution ([Ticket #4663](#)) - 1 year old
- Install PyPi packages directly within ChimeraX ([Ticket #4762](#)) – 11 months old
- Why doesn't devel install install PyPi dependency? ([Ticket #5157](#)) – 8 months old
- Support pyproject.toml ([Ticket #6434](#)) – 7 weeks old

We are already moving in this direction, and we are already committed to *some* work. I am here to explain why in my opinion this vision fits best. Why *this* work is worthwhile in addition to, but also often lieu of, some of the work we would otherwise be committed to.

ChimeraX Binary



Currently

- We use ChimeraX to bootstrap itself: a small subset of ChimeraX python packages are used to build the larger ChimeraX
- **ChimeraX** the C binary depends somewhat less-than-obviously on **chimerax.core** the Python package
 - **ChimeraX_main.py** is really either **chimerax.__main__.py** or **chimerax.core.__main__.py**, but it is in **src/apps/ChimeraX**
 - Depends on **mkinfo.py** and **geticoninfo.py**, neither of which are useful without some attempt to install the core

As a result:

- Core **must** be made on Windows before ChimeraX, to get the build info to make an icon
- Core is initially made on macOS to generate a plist, which is then replaced later

ChimeraX Namespace Package



Currently

- Core and bundles are organized as separate packages, when they're really all subpackages
- When they reach their final destination, though, they are all under the ChimeraX namespace
- Two internal bundles excepted, building a bundle involves massaging a lot of ChimeraX metadata into a format acceptable for Python tooling

As a result, ChimeraX itself depends on ChimeraX.



Currently

- To build a bundle:
 - Define bundle_info.xml
 - Add metadata for setup.py and ChimeraX's custom metadata
 - Bundle builder will generate a setup.py and execute it, building a plain Python wheel
 - This could go on PyPi, but my one of my arguments is it would be somewhere between “quite” and “extremely” annoying for the library to depend on the built ChimeraX distribution and not vice versa
- ChimeraX will use metadata we place in dist_info/METADATA to set up bundles at runtime
- Some packages do not have the same installed folder name as their source folder names
 - Linters, test suites, IDEs, all of which operate on source directories and expect a 1:1 correspondence, don't know where our code is – this handicaps us.
- Language servers have no idea where our code is. If you have no prior knowledge of ChimeraX code IDEs cannot come to the rescue to help you discover it.



Currently

- To build a bundle:
 - Define bundle_info.xml
 - Add metadata for setup.py and ChimeraX's custom metadata
 - Bundle builder will generate a setup.py and execute it, building a plain Python wheel
 - This could go on PyPi, but my one of my arguments is it would be somewhere between “quite” and “extremely” annoying for the library to depend on the built ChimeraX distribution and not vice versa
- ChimeraX will use metadata we place in dist_info/METADATA to set up bundles at runtime
- Some packages do not have the same installed folder name as their source folder names
 - Linters, test suites, IDEs, all of which operate on source directories and expect a 1:1 correspondence, don't know where our code is – this handicaps us.
- Language servers have no idea where our code is. IDEs cannot come to the rescue.

The native-packaging branch



native-packaging from 10,000ft

- Resolve the dependency between the C binary and the namespace package
 - Each one can optionally be developed, distributed separately
- Preserve the core's ability to be built independently from the distribution build system
- Convert bundles to native tooling: pyproject.toml and naked setup.py
 - No more bundle_info.xml *by default* – it was converted to setup.py anyway
- Provide a metadata writer in **core** sufficient to build compliant wheels...
- ...then force bundles to include setuptools, wheel, and at least (but not necessarily only*) **core** as build dependencies in pyproject.toml

As a result, the main build dependency would be core instead of bundle builder.

*If you need to link against some package's shared object then make it a build dependency



Immediate advantages of the branch

Two tickets become more or less the same ticket:

- ChimeraX in PyPi ([Ticket #1470](#)) reduces to uploading wheels to PyPi
 - This can be automated with Python's own `cibuildwheel` package
- Automate uploading wheels to the Toolshed ([Ticket #3058](#)) reduces to an “upload it here too” clause at the end of 1470's solution

Two tickets are naturally completed through the routine work of native packaging:

- Move pip dependencies into bundles ([Ticket #3890](#))
- Support pyproject.toml ([Ticket #6434](#))



Immediate advantages, cont'd

One ticket nearly solves itself:

- Make native Apple M1 CPU distribution ([Ticket #4663](#))
 - Wheels get their tags programmatically. We can multi-arch compile, and still get a specific tag: arm64/x64
 - To build a universal2 wheel requires a universal2 version of Python
 - We can make a virtual environment from the Python prereq, then Python does the heavy lifting for us



Immediate advantages, cont'd

We will be forced to deal with one crucial ticket:

- Why doesn't devel install install PyPi dependency? ([Ticket #5157](#))

And also one stale ticket

- Report problematic version number checking to pip ([Ticket #3403](#))

Both of which I surmise will require more group consensus than code to close.

Speaking of #3403: catch my next presentation “~= specifier considered harmful” because of its total disregard for `import this`'s simplicity clause and the nasty gotchas it carries with it.



What native-packaging does

- Move `ChimeraX_main.py` → `core/__main__.py`
- Point the ChimeraX binary at `chimerax.core`
- Read `argv` from `sys.argv` as the first call in `__main__:init` vs passing in `argv`
- Define `CX_DIST_VER` as a variable that can be used by `make/C` preprocessor
- On macOS: delay installing the plist until the very end, do it once
- On Windows: Get icon info from the build system!



What native-packaging does, cont'd

- Use the ChimeraX binary to modify the ChimeraX module at runtime using the Python C API
 - Cleanly inject the buildinfo variables when it's run from the distribution

```
Py_Initialize(); // blank Python environment
// "chimeraX" module is empty because it contains no code
PyObject* py_cx = PyImport_ImportModule("chimeraX");
PyModule_AddStringConstant(py_cx, "_CHIMERAX_C_DIST_BUILD_DATE", __DATE__)
```

Could also use `PyImport_AddModule` so as not to touch the top-level ChimeraX module



What native-packaging does, cont'd

- How do we know what kind of build we're running with variable injection?

Luckily we already have a way to do this: **make**! We can even use our automated system!

```
$ make -D$(BUILD_TYPE) apps/ChimeraX
```

```
//launcher.c
#ifdef daily
PyModule_AddStringConstant(py_cx, "_CHIMERAX_C_DIST_VERSION", "1.5-daily")
#else
PyModule_AddStringConstant(py_cx, "_CHIMERAX_C_DIST_VERSION", "1.5")
#endif
int result = PyMain(new_argc, new_argv);
```



What native-packaging does, cont'd

- Move src/apps to /apps – they are not part of the namespace package
- Move src/examples to /examples - Not sure where else it should go
- Move **all** bundles to src/
- Merge src/Makefile and src/bundles/Makefile
 - Both get **much** simpler with native packaging, but with some coercion worked on the old system too
- Move metadata from bundle_info.xml to pyproject.toml: `[tools.chimerax]`
- Make classifiers a dynamic category in pyproject.toml



What native-packaging does, cont'd

- Move src/apps to /apps – they are not part of the namespace package
- Move src/examples to /examples - Not sure where else it should go
- Move **all** bundles to src/
- Merge src/Makefile and src/bundles/Makefile
 - Both get **much** simpler with native packaging, but with some coercion worked on the old system too
- Move metadata from bundle_info.xml to pyproject.toml: `[tools.chimerax]`
- Make classifiers a dynamic category in pyproject.toml



What native-packaging does, cont'd

- Add a toml-to-classifier converter to `chimerax.core` that can be used standalone without *any* other bundle
- Use that module to write `classifiers.txt` which `setup.py` reads
- Other bundles should include “`chimerax.core >= 1.5.0`” in `[build-system.requires]`
- `setup.py` should `from chimerax.core.write_bundle_classifiers import write_bundle_classifiers`, then call it

Wheels are then built which have *all* of the correct ChimeraX metadata.

I'm not saying it's the best solution we can find, it's just my proof-of-concept.



Tricking 'python -m build' into using local packages

- First, make a virtual environment somewhere and activate it
- Install **pypiserver**
- Make a directory somewhere called “**packages**”
- “**python -m build --wheel**” in the core source directory
- Copy the built wheel into the packages directory
- Run pypiserver and point it at the packages directory
- Put this in `~/.config/pip/pip.conf`:

```
[global]
```

```
extra-package-index=http://localhost:8080/simple/
```

Can't wait to retire this hack – which the maintainers say is a bug – when **core** is on PyPi.



What a true solution enables

Ideally, decoupling ChimeraX the C binary from `chimerax` the namespace package should let us use native tooling to build every part of ChimeraX, including core.

We will eventually *have* to change the bundle builder to support `setuptools`'s deprecation of direct invocations of `setup.py`. We once thought that would be `setup.cfg` for pure python packages; however...

We **must** support `pyproject.toml`, as `setuptools` will eventually deprecate `setup.cfg` too

ChimeraX bundles will probably always require a `setup.py*`, but we can store package metadata, ChimeraX metadata, and most importantly *read* ChimeraX's metadata from `pyproject.toml` files


So why don't we just use the native config files instead of `bundle_info.xml`?

We can read the bundle name for the toolshed, for example, from `[tool.chimerax.bundle-name]`

Classifiers come from `[tool.chimerax.classifiers]`

*`build.py` in `poetry`'s case

Other Miscellaneous Changes

- 
- mk/os.make sets \$(CC) and \$(CXX) on macOS to clang **PLUS** what should be \$(CFLAGS) and \$(CXXFLAGS) but macOS already knows it should use clang as it's Apple's preferred compiler
 - This caused issues compiling md_crds/ because a dependency thinks all the \$(CC) or \$(CXX) definition is the compiler and not just the first space-separated field of the definition
 - To fix: remove \$(CC) and \$(CXX) definitions from macOS, make them \$(CFLAGS) and \$(CXXFLAGS)
 - We do not need two sync directories
 - Uploaded slightly modified segger tarball with extension only recognized by my branch to change its internal \$(TOP) value
 - Adjust spacenavigator makefile to use the \$(TOP) variable, and the new path to copyright.txt
 - **chimerax.core.main**: It's now two more folders deep so needs two more iterations of the **dn ()** basedir finding hack



But let's not get carried away...

I'd like to reiterate:

We do **not** need a new build system, new repo, nor a new fundamental paradigm for reasoning about ChimeraX. For the thousands of files changed by **native-packaging**, over 90% of changes were paths: the location of \$(TOP), where the core Makefile is, and corrections to the gitignore, for example.

This proposed changeset just makes it easy to reason about ChimeraX the distribution and ChimeraX the package at the same time.

Other Miscellaneous Suggestions

- 
- `rg "ioutil"` reveals that `ioutil` C++ code is only referenced in two places:
 - *Defined* in `core`
 - *Included* in `atomic_lib`

Source code that *can* live where it's used needs to live where it's used.

- `rg "pysupport"` reveals that it's used in many places
- `rg "chutil"` reveals that it's used in many places
- `rg "include.*logger"` reveals that it's used in many places

If source code must be abstracted, let's place it in a common top-level "include" directory, or in `src/_lib`, like other Python packages; it's just confusing and not very discoverable that it's polluting the core directory.

Native tooling can then be pointed at that include directory to include code in source distributions, which will then be turned into compiled C/++ extensions.

A New Distribution Scheme



Supporting ChimeraX-Y.Z and PyPi in One Codebase

- With a native-first package, we could have a point release that specifies what versions of our packages are *stable* and *supported for that release*.

Possible commands:

- `make build-pypi`
 - Recursively build individual bundles and upload them to PyPi (except ui)
- `make build-distribution`
 - Could be exactly what we do now

Alternatively:

- Build the ChimeraX environment, then download ChimeraX packages from the toolshed

On the toolshed:

- Host empty package, e.g. `chimerax-1.4-dist` on toolshed
- This package will depend on every bundle that comprises a stable point release



Testing to the Best of Our Abilities

- This way of doing things would enable us to audit ChimeraX via automated testing
 - Each package is buildable, installable, and testable without the ChimeraX environment, as any library package should be, so...
- We could gatekeep uploads to the toolshed behind passing tests, as we already should for PyPi
- ChimeraX-daily *could* become whatever complete suite of mutually compatible bundles managed to pass testing
- Won't be complete at first, but over time will dramatically cut down or possibly eliminate broken daily builds

Alternatively, we could just do what we do now: ship code.



ChimeraX on the Toolshed

- Toolshed is a PyPi repository
- ...and *already* in need of an update, regardless of the adoption of these changes
- Not sure what tech stack Toolshed uses but could probably simplify to simple pypi server with web interface

If we get it into a format that `pip` can handle, that means we can connect to it directly with `pip`

We can add the toolshed to a `pip.conf` in the internal site directory.

- Offer users two options for updating ChimeraX in the GUI
 - Only install stable updates (same minor version)
 - Install any mutually compatible updates
- Allow users to check for updates to specific packages

For users who elect to take the latest stable bundles, it solves the “oh crap we shipped a bug” problem.



What if a user or developer borks their environment?

- `{toolshed,pip} reinstall chimeraX-1.4-py39-none-any.whl`
 - Clear the internal site-packages directory
 - Pull the ChimeraX distribution wheel from the toolshed
 - Reinstall it

The Timing Is Right



Programmmator Cave

- I don't pretend to have all the answers to all the unknown unknowns that this implies...
- ...but I have tried my hardest to anticipate, discover, and provisionally solve as many of them as possible
- Nothing said today is set in stone without the consent of the team
- I used `setuptools` for this talk but `poetry` conversions are easy!
 - (that's the point of `pyproject.toml`!)
 - Poetry lacks dynamic metadata which takes it out of the running *for now*
- While I mentioned some hypothetical possibilities, we should evaluate the talk by how much it reduces already-planned work – the status quo is suboptimal but not *exactly* “broken” – before we get excited about the possibilities

In Summary



Summary

- Native packaging tools can support ChimeraX distribution and ChimeraX library simultaneously with minimal overhead
- Decoupling the binary distribution from the core enables many modern features impossible under our current scheme, both for developers and users
- Modern package architectures make it easier for people who know Python but *not ChimeraX* to apply their knowledge to ChimeraX: we can expect more bundle developers in the future.

Best of all:

- This future is much closer than previously anticipated. An MVP is already mostly done.

Of course, having said that, the 20% that remains will become the new 80% I'm sure.

Thank you!