

---

# **Hammer**

***Release 1.0.0***

**Berkeley Architecture Research**

**Feb 09, 2023**



## CONTENTS:

<b>1</b>	<b>Hammer Basics</b>	<b>3</b>
1.1	Hammer Overview . . . . .	3
1.1.1	Main Hammer . . . . .	3
1.1.2	Tech Plugins . . . . .	4
1.1.3	Tool Plugins . . . . .	4
1.1.4	Calling Hammer . . . . .	5
1.2	Hammer Setup . . . . .	5
1.2.1	User Setup . . . . .	5
1.2.2	Developer Setup . . . . .	6
1.3	Migration Guide . . . . .	9
1.3.1	Selecting Plugins . . . . .	10
1.3.2	import Statements . . . . .	10
1.3.3	Technology JSON . . . . .	10
1.3.4	Plugin File Structure . . . . .	11
1.3.5	Resource Files . . . . .	12
1.3.6	pyproject.toml . . . . .	12
<b>2</b>	<b>Hammer Technology Plugins</b>	<b>13</b>
2.1	Hammer Technologies . . . . .	13
2.1.1	HammerTechnology Class . . . . .	14
2.2	Hammer Tech JSON . . . . .	14
2.2.1	Technology Install . . . . .	14
2.2.2	DRC/LVS Deck Setup . . . . .	15
2.2.3	Library Setup . . . . .	15
2.2.4	Stackup . . . . .	16
2.2.5	Sites . . . . .	17
2.2.6	Special Cells . . . . .	17
2.2.7	Don't Use, Physical-Only Cells . . . . .	17
2.3	Hammer Tech defaults.yml . . . . .	18
2.4	Sky130 Technology Library . . . . .	19
2.4.1	PDK Structure . . . . .	19
2.4.2	NDA Files . . . . .	19
2.4.3	Prerequisites for PDK Setup . . . . .	20
2.4.4	OpenRAM SRAMs . . . . .	20
2.4.5	PDK Setup . . . . .	20
2.4.6	Known DRC Issues . . . . .	20
2.5	ASAP7 Technology Library . . . . .	21
2.5.1	Setup and Environment . . . . .	21
2.5.2	Dummy SRAMs . . . . .	22
2.5.3	Known Issues . . . . .	22

2.6	Nangate45 Technology Library . . . . .	22
2.6.1	Dummy SRAMs . . . . .	22
2.6.2	Supported Tools . . . . .	23
<b>3</b>	<b>Hammer Tool Plugins</b>	<b>25</b>
3.1	Hammer CAD Tools . . . . .	25
3.2	Setting up a Hammer CAD Tool Plugin . . . . .	25
3.2.1	Tool Class . . . . .	26
3.2.2	Steps . . . . .	26
3.2.3	Getting Settings . . . . .	26
3.2.4	Writing TCL . . . . .	26
3.2.5	Executing the Tool . . . . .	26
3.2.6	Tool Outputs . . . . .	26
3.2.7	defaults.yml . . . . .	27
3.3	OpenROAD Place-and-Route Tool Plugin . . . . .	27
3.3.1	Tool Setup . . . . .	27
3.3.2	Tool Steps . . . . .	27
3.3.3	Step Details . . . . .	27
3.3.4	Issue Archiving . . . . .	28
<b>4</b>	<b>Hammer Flow Steps</b>	<b>29</b>
4.1	Hammer Actions . . . . .	29
4.2	Synthesis . . . . .	29
4.2.1	Synthesis Setup Keys . . . . .	29
4.2.2	Synthesis Input Keys . . . . .	30
4.2.3	Synthesis Inputs . . . . .	30
4.2.4	Synthesis Outputs . . . . .	30
4.2.5	Synthesis Commands . . . . .	30
4.3	Place-and-Route . . . . .	31
4.3.1	P&R Setup Keys . . . . .	31
4.3.2	P&R Input Keys . . . . .	31
4.3.3	P&R Inputs . . . . .	32
4.3.4	P&R Outputs . . . . .	32
4.3.5	P&R Commands . . . . .	32
4.4	DRC . . . . .	32
4.4.1	DRC Setup Keys . . . . .	33
4.4.2	DRC Input Keys . . . . .	33
4.4.3	DRC Inputs . . . . .	33
4.4.4	DRC Outputs . . . . .	33
4.4.5	DRC Commands . . . . .	33
4.5	LVS . . . . .	34
4.5.1	LVS Setup Keys . . . . .	34
4.5.2	LVS Input Keys . . . . .	34
4.5.3	LVS Inputs . . . . .	34
4.5.4	LVS Outputs . . . . .	34
4.5.5	LVS Commands . . . . .	35
4.6	Simulation . . . . .	35
4.6.1	Simulation Setup Keys . . . . .	35
4.6.2	Simulation Input Keys . . . . .	35
4.6.3	Simulation Inputs . . . . .	37
4.6.4	Simulation Outputs . . . . .	37
4.6.5	Simulation Commands . . . . .	37
4.7	Power . . . . .	37
4.7.1	Power Setup Keys . . . . .	38

4.7.2	Simulation Input Keys . . . . .	38
4.7.3	Power Inputs . . . . .	39
4.7.4	Power Outputs . . . . .	39
4.7.5	Power Commands . . . . .	39
4.8	Formal Verification . . . . .	40
4.8.1	Formal Verification Setup Keys . . . . .	40
4.8.2	Formal Verification Input Keys . . . . .	40
4.8.3	Formal Verification Inputs . . . . .	41
4.8.4	Formal Verification Outputs . . . . .	41
4.8.5	Formal Verification Commands . . . . .	41
4.9	Static Timing Analysis . . . . .	42
4.9.1	STA Setup Keys . . . . .	42
4.9.2	STA Input Keys . . . . .	42
4.9.3	STA Inputs . . . . .	42
4.9.4	STA Outputs . . . . .	43
4.9.5	STA Commands . . . . .	43
<b>5</b>	<b>Hammer Use</b>	<b>45</b>
5.1	Hammer IR and Meta Variables . . . . .	45
5.1.1	Hammer IR . . . . .	45
5.1.2	The hammer-config library . . . . .	45
5.1.3	Basics . . . . .	46
5.1.4	Overriding . . . . .	46
5.1.5	Meta actions . . . . .	46
5.1.6	Applying multiple meta actions . . . . .	47
5.1.7	Common meta actions . . . . .	47
5.1.8	Type Checking . . . . .	49
5.1.9	Key History . . . . .	49
5.1.10	Key Description Lookup . . . . .	50
5.1.11	Reference . . . . .	51
5.2	Hammer APIs . . . . .	51
5.2.1	Power Specification . . . . .	52
5.2.2	Timing Constraints . . . . .	52
5.2.3	Floorplan & Placement . . . . .	52
5.2.4	Bumps . . . . .	52
5.2.5	Pins . . . . .	52
5.2.6	Power Straps . . . . .	53
5.2.7	Special Cells . . . . .	53
5.2.8	Submission . . . . .	53
5.3	Flow Control . . . . .	53
5.3.1	Command-line Interface . . . . .	53
5.4	Extending Hammer with Hooks . . . . .	54
5.4.1	Hook Methods . . . . .	54
5.4.2	Including Hooks . . . . .	54
5.4.3	Technology, Tool, and User-Provided Hooks . . . . .	55
5.5	Hammer Buildfile . . . . .	56
5.6	Hierarchical Hammer Flow . . . . .	56
5.6.1	Hierarchical Hammer Config . . . . .	56
5.6.2	Flow Management and Actions . . . . .	57
5.6.3	Cadence Implementation . . . . .	57
5.6.4	Tips for Constraining Hierarchical Modules . . . . .	58
5.6.5	Special Notes & Limitations . . . . .	58
<b>6</b>	<b>Hammer Examples</b>	<b>61</b>

6.1	End-to-End Integration Tests . . . . .	61
6.2	OpenROAD and Sky130 . . . . .	61
6.2.1	Instructions . . . . .	61
<b>7</b>	<b>Indices and tables</b>	<b>63</b>



# HAMMER

Hammer is a physical design framework that wraps around vendor specific technologies and tools to provide a single API to create ASICs. Hammer allows for reusability in ASIC design while still providing the designers leeway to make their own modifications.

Hammer (Highly Agile Masks Made Effortlessly from RTL) is a framework for building physical design generators for digital VLSI flows. It is an evolving set of APIs that enable reuse in an effort to speed up VLSI flows, which have traditionally been entirely rebuilt for different projects, technologies, and tools.

Hammer is able to generate scripts and collateral for a growing range of CAD tools while remaining technology-agnostic using a well-defined set of common APIs. Tool- and technology-specific concerns live inside plugins, implement APIs, and provide a set of working default configurations.

The vision of Hammer is to reduce the cycle time on VLSI designs, enabling rapid RTL design space exploration and allowing a designer to investigate the impact of various parameters like timing constraints and floorplans without needing to worry about low-level details.

For high-level details about Hammer's design principles and capabilities, please refer to our DAC 2022 paper entitled [Hammer: A Modular and Reusable Physical Design Flow Tool](#). We kindly request that this paper be cited in any publications where Hammer was used.





## HAMMER BASICS

This documentation will give an overview of Hammer, its basic setup, its components, and its structure, as well as some typical project setup.

### 1.1 Hammer Overview

Hammer has a set of actions and automatically takes the output of one action and converts it into the input for another. For instance, a synthesis action will output a mapped verilog file which will then automatically be piped to the place-and-route input when a place-and-route action is called.

A user's Hammer environment is typically separated into four different components: core Hammer, one or more tool plugins, a technology plugin, and a set of project-specific Hammer input files. Hammer is meant to expose a set of generalized APIs that are then implemented by tool- and technology-specific plugins.

Hammer is included in a larger project called [Chipyard](#) which is the unified repo for an entire RTL, simulation, emulation, and VLSI flow from Berkeley Architecture Research. There is an in-depth Hammer demo there, and it is a great place to look at a typical Hammer setup.

#### 1.1.1 Main Hammer

Hammer provides the Python backend for a Hammer project and exposes a set of APIs that are typical of modern VLSI flows. These APIs are then implemented by a tool plugin and a technology plugin of the designer's choice. The structure of Hammer is meant to enable re-use and portability between technologies.

Hammer takes its inputs and serializes its state in form of YML and JSON files. The designer sets a variety of settings in the form of keys in different namespaces that are designated in Hammer to control its functionality. These keys are contained in `hammer/src/hammer-vlsi/defaults.yml`. This file shows all of the keys that are a part of main Hammer and provides sensible defaults that may be overridden or are set to null if they must be provided by the designer.

Here is an example of a snippet that would be included in the user's input configuration.

```
vlsi.core.technology: "asap7"
vlsi.inputs.supplies:
  VDD: "0.7 V"
  GND: "0 V"
```

This demonstrates two different namespaces, `vlsi.core` and `vlsi.inputs`, and then two different keys, `technology` and `supplies`, which are set to the `asap7` technology and 0.7 Volts supply voltage, respectively.

Further details about these keys and how they are manipulated is found in the [Hammer IR and Meta Variables](#) section.

### 1.1.2 Tech Plugins

A technology plugin consists of two or more files: a `*.tech.json` and a `defaults.yml`.

The `*.tech.json` contains pointers to relevant PDK files and fundamental technology constants. These values are not meant to be overridden, nor can they be for the time being.

`defaults.yml` sets default technology variables for Hammer to consume, which may be specific to this technology or generic to all. These values may be overridden by design-specific configurations. An example of this is shown in the open-source technology plugins in `hammer/technology/`, such as `asap7`, and how to setup a technology plugin is documented in more detail in the *Hammer Technology Plugins* section.

---

**Note:** Unless you are a UCB BAR or BWRC affiliate or have set up a 3-way technology NDA with us, we cannot share pre-built proprietary technology plugin repositories.

---

### 1.1.3 Tool Plugins

A Hammer tool plugin actually implements tool-specific steps of the VLSI flow in Hammer in a template-like fashion. The TCL commands input to the tool are created using technology and design settings provided by the designer.

There are currently three Hammer tool plugin repositories for commercial tools: `hammer-cadence-plugins`, `hammer-synopsys-plugins`, and `hammer-mentor-plugins`. In them are tool plugin implementations for actions including synthesis, place-and-route, DRC, LVS, and simulation. `hammer-cadence-plugins` and `hammer-synopsys-plugins` are publicly available; however, users must request access to `hammer-mentor-plugins`.

---

**Note:** If you are not a UCB BAR or BWRC affiliate and have access to Mentor Graphics (now Siemens) tools, please email [hammer-plugins-access@lists.berkeley.edu](mailto:hammer-plugins-access@lists.berkeley.edu) with a request for access to the `hammer-mentor-plugins` repository. MAKE SURE TO INCLUDE YOUR GITHUB ID IN YOUR EMAIL AND YOUR ASSOCIATION TO SHOW YOU HAVE LICENSED ACCESS TO THE TOOLS. There will be no support guarantee for the plugin repositories, but users are encouraged to file issues and contribute patches where needed.

---

There are also a set of open-source tools (e.g. Yosys, OpenROAD, Magic, Netgen) provided in `hammer/` under their respective actions.

These plugins implement many of the common steps of a modern physical design flow. However, a real chip flow will require many custom settings and steps that may not be generalizable across technology nodes. Because of this, Hammer has an “escape-hatch” mechanism, called a hook, that allows the designer to inject custom steps between the default steps provided by the CAD tool plugin. Hooks are python methods that emit TCL code and may be inserted before or after an existing step or replace the step entirely. This allows the designer to leverage the APIs built into Hammer while easily inserting custom steps into the flow. Hooks are discussed in more detail in the “Example usage” portion of the Hammer documentation.

### 1.1.4 Calling Hammer

To use Hammer on the command line, the designer will invoke the `hammer-vlsi` utility. This is calling the `__main__()` method of the `CLIDriver` class. An example invocation is below:

```
hammer-vlsi -e env.yml -p config.yml --obj_dir build par
```

Using hooks requires the designer to extend the `CLIDriver` class. A good example exists in the [Chipyard](#) repository (`chipyard/vlsi/example-vlsi`). This would change the invocation to something like the following:

```
example-vlsi -e env.yml -p config.yml --obj_dir build par
```

In both of these commands, an environment configuration is passed to Hammer using a `-e` flag, which in this case is `env.yml`. `env.yml` contains pointers to the required tool licenses and environment variables. These environment settings will not be propagated to the output configuration files after each action.

Any number of other YML or JSON files can then be passed in using the `-p` flag. In this case, there is only one, `config.yml`, and it needs to set all the required keys for the step of the flow being run. Passing in multiple files looks like `-p config1.yml -p config2.yml`. Refer to the [Hammer IR and Meta Variables](#) section for the implications of multiple config files.

`--obj_dir build` designates what directory Hammer should use as a working directory. All default action run directories and output files will be placed here.

Finally, `par` designates that this is a place-and-route action.

In this case, Hammer will write outputs to the path `$PWD/build/par-rundir`.

For the full list of Hammer command-line arguments, run `hammer-vlsi --help` or take a peek in the `hammer/vlsi/cli_driver.py` file.

## 1.2 Hammer Setup

Hammer depends on Python 3.9+.

The default technology, ASAP7, has some extra requirements. See its [README](#) for instructions.

### 1.2.1 User Setup

You can install Hammer from PyPI:

```
pip install hammer-vlsi
```

If you are using ASAP7, you need to install `hammer-vlsi` with the `asap7` extra dependency (`gdspy` or `gdstk`). By default, `gdspy` is installed:

```
pip install hammer-vlsi[asap7]
```

If instead, you want to install `gdstk`:

```
pip install hammer-vlsi[asap7-gdstk]
```

After installation, verify that you can run the `hammer-vlsi` script from the command line.

```
hammer-vlsi -h
```

Note: certain tools and technologies will have additional system requirements. For example, LVS with Netgen requires Tcl/Tk 8.6, which is not installed for CentOS7/RHEL7 and below. Refer to each respective tool and technology's documentation for those requirements.

### Installing Hammer as a Source Dependency

In some cases, it is useful to install Hammer as a source dependency. For instance, when developing tool or PDK plugins alongside a new feature or API changes in main Hammer, installing hammer as a source dependency will allow you to make changes in main hammer and see them reflected immediately when running code for your tool/PDK plugin.

#### From Another Poetry Project

Hammer tool (e.g. `hammer-cadence-plugins`) and PDK plugin repositories are poetry projects (with a `pyproject.toml` in their root). To depend on Hammer as a source dependency, first clone Hammer somewhere on your disk. Then add this snippet to the tool/PDK plugin repo's `pyproject.toml` (and remove any PyPI dependency on Hammer):

```
[tool.poetry.dependencies]
#hammer-vlsi = "^1.0.0"
hammer-vlsi = {path = "path/to/hammer", extras = ["asap7"], develop = true}
```

Run `poetry update` and `poetry install`. Do not commit the changes to `pyproject.toml` or `poetry.lock` without first removing the source dependency. You only need to specify `extras` if you need the `asap7` optional dependency (`gdstk`).

#### From a Generic Python Project

Other repos, such as Chipyard, are not poetry projects, but still depend on Hammer. To use Hammer as a source dependency:

1. Remove the PyPI `hammer-vlsi` dependency from the project (e.g. by editing a `conda env.yml` file and rerunning dependency resolution)
2. Clone Hammer somewhere on your disk
3. Activate the virtualenv of the project (e.g. Chipyard)
4. Run `pip install -e .` from the root of Hammer *within the project's virtualenv*

### 1.2.2 Developer Setup

1. Clone Hammer with `git`

```
git clone git@github.com:ucb-bar/hammer
cd hammer
```

2. Install `poetry` to manage the development virtualenv and dependencies

```
curl -sSL https://install.python-poetry.org | python3 -
```

3. Create a poetry-managed virtualenv using the dependencies from `pyproject.toml`

```
# create the virtualenv inside the project folder (in .venv)
poetry config virtualenvs.in-project true
poetry install
```

4. Activate the virtualenv. Within the virtualenv, Hammer is installed and you can access its scripts defined in `pyproject.toml` (in `[tool.poetry.scripts]`)

```
poetry shell
hammer-vlsi -h
```

## Using PyCharm

This project works out of the box with PyCharm. You should install the [Pydantic plugin](#) to enable autocomplete for Pydantic BaseModel.

## Unit Tests with pytest

Within the poetry virtualenv, from the root of Hammer, run the tests (`-v` will print out each test name explicitly)

```
pytest tests/ -v
```

If you want to skip the single long running-test in `test_stackup.py`:

```
pytest tests/ -m "not long" -v
```

If you want to run only a specific test use `-k` with a snippet of the test function you want to run:

```
pytest tests/ -k "lsf" -v
```

```
> tests/test_submit_command.py::TestSubmitCommand::test_lsf_submit[lsf] PASSED
```

By default, `pytest` will only display what a test prints to stdout if the test fails. To display stdout even for a passing test, use `-rA`:

```
pytest tests/test_build_systems.py -k "flat_makefile" -rA -v

> _____ TestHammerBuildSystems.test_flat_makefile _____
> ----- Captured stdout call -----
> [<global>] Loading hammer-vlsi libraries and reading settings
> [<global>] Loading technology 'nop'
> ===== short test summary info =====
> PASSED tests/test_build_systems.py::TestHammerBuildSystems::test_flat_makefile
```

### Type Checking with mypy

There is a [small issue](#) with the `ruamel.yaml` package typechecking which can be hacked around with:

```
touch .venv/lib/python3.10/site-packages/ruamel/py.typed
```

Inside your poetry virtualenv, from the root of Hammer, run:

```
mypy --namespace-packages --warn-unused-ignores -p hammer

Success: no issues found in 146 source files

mypy --namespace-packages --warn-unused-ignores tests

Success: no issues found in 25 source files
```

### Testing Different Python Versions with tox

Hammer is supposed to work with Python 3.9+, so we run its unit tests on all supported Python versions using `tox` and `pyenv`.

#### 1. Install pyenv

```
curl https://pyenv.run | bash
```

Restart your shell and run `pyenv init` (and follow any of its instructions). Then restart your shell again.

#### 2. Install Python versions

See the `.python-version` file at the root of hammer and install those Python versions using `pyenv`.

```
pyenv install 3.9.13
pyenv install 3.10.6
```

Once the Python interpreters are installed, run `pyenv versions` from the root of hammer.

```
pyenv versions
system
* 3.9.13 (set by .../hammer/.python-version)
* 3.10.6 (set by .../hammer/.python-version)
```

#### 3. From within your poetry virtualenv, run tox

```
tox
```

This will run the pytest unit tests using all the Python versions specified in `pyproject.toml` under the `[tool.tox]` key.

You can run tests only on a particular environment with `-e`

```
tox -e py39 # only run tests on Python 3.9
```

You can pass command line arguments to the pytest invocation within a tox virtualenv with `--`

```
tox -e py39 -- -k "lsf" -v
```

## Adding / Updating Dependencies

To add a new Python (pip) dependency, modify `pyproject.toml`. If the dependency is only used for development, add it under the key `[tool.poetry.dev-dependencies]`, otherwise add it under the key `[tool.poetry.dependencies]`. Then run `poetry update` and `poetry install`. The updated `poetry.lock` file should be committed to Hammer.

To update an existing dependency, modify `pyproject.toml` with the new version constraint. Run `poetry update` and `poetry install` and commit `poetry.lock`.

## Building Documentation

- Within your poetry virtualenv, `cd doc`
- Modify any documentation files. You can migrate any `rst` file to Markdown if desired.
- Run `sphinx-build . build`
- The generated HTML files are placed in `build/`
- Open them in your browser `firefox build/index.html`

## Publishing

Build a sdist and wheel (results are in `dist`):

```
poetry build
```

To publish on TestPyPI:

1. Create an account on [TestPyPi](#)
2. Note the source repository `testpypi` in `pyproject.toml` under the key `[tool.poetry.source]`
  - To add another [PyPI repository to poetry](#): `poetry source add <source name> <source url>`
3. Publish: `poetry publish --repository testpypi -u <username> -p <password>`

## 1.3 Migration Guide

Hammer's software infrastructure changed significantly for version 1.0.0. In order to use the latest version of Hammer, you will need to make changes to your existing tool/tech plugins, YAML/JSON input files, and Python files that import Hammer.

This guide is relevant for both plugin developers and general users.

The documentation for old Hammer is [cached here](#).

### 1.3.1 Selecting Plugins

Previously, tech and tool plugins were selected with a combination of tool name and relative plugin path:

```
vlsi.core.technology: <tech_name>
vlsi.core.technology_path: ["hammer-<tech_name>-plugin"]
vlsi.core.technology_path_meta: append

vlsi.core.synthesis_tool: <syn_tool_name>
vlsi.core.synthesis_tool_path: ["hammer-<vendor>-plugin/syn"]
vlsi.core.synthesis_tool_meta: append
```

Now, you simply need to specify the package name:

```
vlsi.core.technology: hammer.technology.<tech_name>

vlsi.core.synthesis_tool: hammer.synthesis.<syn_tool_name>
```

See the *Plugin File Structure* section below for details about plugin package structure.

### 1.3.2 import Statements

When importing Hammer classes and/or methods, replace old statements:

```
import hammer_tech
from hammer_vlsi import HammerTool
```

with:

```
import hammer.tech
from hammer.vlsi import HammerTool
```

The rule of thumb is that underscores are replaced with periods. This will match the package directory structure under the `hammer/` directory.

### 1.3.3 Technology JSON

Previously, the technology JSON file may have contained entries like this:

```
"gds map file": "path/to/layermap.file"
```

Now, keys must not contain spaces in line with JSON syntax:

```
"gds_map_file": "path/to/layermap.file"
```

The fields for `installs` and `tarballs` have also changed. Generally, `path` is now `id` and `base var` is now `path` to remove confusion. For example, previously (ASAP7 example for reference):

```
"installs": [
  {
    "path": "$PDK",
    "base var": "technology.asap7.pdk_install_dir"
  }
]
```

(continues on next page)



(continued from previous page)

```

],
"tarballs": [
  {
    "path": "ASAP7_PDK_CalibreDeck.tar",
    "homepage": "http://asap.asu.edu/asap/",
    "base var": "technology.asap7.tarball_dir"
  }
]

```

is now:

```

"installs": [
  {
    "id": "$PDK",
    "path": "technology.asap7.pdk_install_dir"
  }
],
"tarballs": [
  {
    "root": {
      "id": "ASAP7_PDK_CalibreDeck.tar",
      "path": "technology.asap7.tarball_dir"
    },
    "homepage": "http://asap.asu.edu/asap/",
    "optional": true
  }
]

```

### 1.3.4 Plugin File Structure

Plugins previously did not have a file structure requirement. For example, it could have looked like this:

```

mytech/
  __init__.py
  defaults.yml
  mytech.tech.json
  layer.map          # some tech-specific file
action/
  action_tool/
    __init__.py      # contains action_tool-specific hooks
    tool.options     # some tech-specific file needed by this tool

```

The new structure must follow Python package convention, hence it will look as follows:

```

hammer/
  mytech/
    __init__.py
    defaults.yml
    mytech.tech.json
    layer.map        # some tech-specific file
    action/

```

(continues on next page)

(continued from previous page)

```
action_tool/  
    __init__.py    # contains action_tool-specific hooks  
    tool.options  # some tech-specific file needed for this tool
```

### 1.3.5 Resource Files

Technology plugins will often provide special static files needed by tools, such as the `layer.map` file in the above example. If this file is already specified with the `prependlocal` meta action, the behavior will remain the same:

```
action.tool.layer_map: "layer.map"  
action.tool.layer_map_meta: prependlocal
```

However, within a plugin, if you need to access to a specific file for a tool-specific hook, for example, in `action/action_tool/__init__.py`, replace this:

```
with open(os.path.join(self.tool_dir, "tool.options")) as f:
```

With this and append Pathlib methods like `read_text()` as required:

```
importlib.resources.files(self.package).joinpath("tool.options")
```

### 1.3.6 pyproject.toml

Plugins that are repositories separate from Hammer must now have a `pyproject.toml` file so that they become poetry projects. Refer to *From Another Poetry Project* for details.

## HAMMER TECHNOLOGY PLUGINS

This guide will walk you through how to set up a technology plugin to be used in Hammer.

You may use the included [free ASAP7 PDK](#) or the [open-source Sky130 PDK](#) plugins as reference when building your own technology plugin.

Technology plugins must be structured as Python packages underneath the `hammer` package. The package should contain an class object named `tech` to create an instance of the technology. This object should be a subclass of `HammerTechnology`.

Technology plugins must also have `.tech.json` and `defaults.yml` files. See the following sections for how to write them.

### 2.1 Hammer Technologies

Hammer currently has open-source technology plugins in the `hammer/technology` folder. Other proprietary technology plugins cannot be made public without 3-way NDAs or special agreements with their respective foundries.

The structure of each technology plugin package is as follows:

- `hammer`
  - `TECHNOLOGY_NAME`
    - \* `__init__.py` contains the tech class object and methods for PDK installation/extraction and getting *hooks*.
    - \* `<name>.tech.json` contains the static information about the technology. See the section [Hammer Tech JSON](#) for a guide.
    - \* `defaults.yml` contains the default tech-specific [Hammer IR and Meta Variables](#). See the section [Hammer Tech defaults.yml](#) for a guide.
    - \* `ACTION_NAME`
      - `__init__.py` (optional) can contain a technology's implementation of an action. Commonly, this is used for the SRAM compilation action.
      - `TOOL_NAME`
        - `__init__.py` (optional) contains callable *hook* methods specific to the technology and tool, if there are too many to put in `TECHNOLOGY_NAME/__init__.py`.

Resources that are needed may go in this directory structure where necessary.

### 2.1.1 HammerTechnology Class

The HammerTechnology class is the base class that all technology plugins should inherit from and is defined in `hammer/tech/__init__.py`. Particularly useful methods are:

- `post_install_script`: the plugin subclass should override this method to apply any non-default PDK installation steps or hotfixes in order to set up the technology libraries for use with Hammer.
- `read_libs`: this is a particularly powerful method to read the libraries contained in the `tech.json` file and filter them using the filters also contained the same file. See the [Library Filters](#) section for details.
- `get_tech_<action>_hooks`: these methods are necessary if tech-specific hooks are needed, and must return a list of hooks when a given tool name implementing the relevant action is called. Refer to the [Extending Hammer with Hooks](#) section for details about how to write hooks.

## 2.2 Hammer Tech JSON

The `tech.json` for a given technology sets up some general information about the install of the PDK, sets up DRC rule decks, sets up pointers to PDK files, and supplies technology stackup information. For a full schema that the tech JSON supports, please see the TechJSON Pydantic model class in `hammer/tech/__init__.py`.

### 2.2.1 Technology Install

The user may supply the PDK to Hammer as an already extracted directory and/or as a tarball that Hammer can automatically extract. Setting `technology.TECH_NAME`, `install_dir` and/or `tarball_dir` (key is setup in the `defaults.yml`) will fill in as the path prefix for paths supplied to PDK files in the rest of the `tech.json`. `install-example` shows an example of the installs and tarballs from the ASAP7 plugin.

```
"name": "ASAP7 Library",
"grid_unit": "0.001",
"time_unit": "1 ps",
"installs": [
  {
    "id": "$PDK",
    "path": "technology.asap7.pdk_install_dir"
  },
  {
    "id": "$STDCELLS",
    "path": "technology.asap7.stdcell_install_dir"
  }
],
"tarballs": [
  {
    "root": {
      "id": "ASAP7_PDK_CalibreDeck.tar",
      "path": "technology.asap7.tarball_dir"
    },
    "homepage": "http://asap.asu.edu/asap/",
    "optional": true
  }
],
```

The `id` field is used within the file listings further down in the file to prefix `path`, as shown in detail below. If the file listing begins with `cache`, then this denotes files that exist in the tech cache, which are generally placed there by the tech plugin's post-installation script (see ASAP7's `post_install_script` method). Finally, the encrypted Calibre decks are provided in a tarball and denoted as optional.

## 2.2.2 DRC/LVS Deck Setup

As many DRC & LVS decks for as many tools can be specified in the `drc decks` and `lvs decks` keys. Additional DRC/LVS commands can be appended to the generated run files by specifying raw text in the `additional_drc_text` and `additional_lvs_text` keys. `deck-example` shows an example of an LVS deck from the ASAP7 plugin.

```
"lvs_decks": [
  {
    "tool_name": "calibre",
    "deck_name": "all_lvs",
    "path": "ASAP7_PDK_CalibreDeck.tar/calibredecks_r1p7/calibre/ruledirs/lvs/lvsRules_
↳calibre_asap7.rul"
  }
],
"additional_lvs_text": "LVS SPICE EXCLUDE CELL \*SRAM*RW*\nLVS BOX \*SRAM*RW*\nLVS_
↳FILTER \*SRAM*RW*\n OPEN",
```

The file pointers, in this case, use the tarball prefix because Hammer will be extracting the rule deck directly from the ASAP7 tarball. The additional text is needed to tell Calibre that the dummy SRAM cells need to be filtered from the source netlist and boxed and filtered from the layout.

## 2.2.3 Library Setup

The `libraries` key also must be setup in the JSON plugin. This will tell Hammer where to find all of the relevant files for standard cells and other blocks for the VLSI flow. `library-example` shows an example of the start of the library setup and one entry from the ASAP7 plugin.

```
"libraries": [
  {
    "lef_file": "$STDCELLS/techlef_misc/asap7_tech_4x_201209.lef",
    "provides": [
      {
        "lib_type": "technology"
      }
    ]
  },
  {
    "nldm_liberty_file": "$STDCELLS/LIB/NLDM/asap7sc7p5t_SIMPLE_RVT_TT_nldm_201020.lib.gz
↳",
    "verilog_sim": "$STDCELLS/Verilog/asap7sc7p5t_SIMPLE_RVT_TT_201020.v",
    "lef_file": "$STDCELLS/LEF/scaled/asap7sc7p5t_27_R_4x_201211.lef",
    "spice_file": "$STDCELLS/CDL/LVS/asap7sc7p5t_27_R.cdl",
    "gds_file": "$STDCELLS/GDS/asap7sc7p5t_27_R_201211.gds",
    "qrc_techfile": "$STDCELLS/qrc/qrcTechFile_typ03_scaled4xV06",
    "spice_model_file": {
      "path": "$PDK/models/hspice/7nm_TT.pm"
    }
  },
]
```

(continues on next page)

(continued from previous page)

```

"corner": {
  "nmos": "typical",
  "pmos": "typical",
  "temperature": "25 C"
},
"supplies": {
  "VDD": "0.70 V",
  "GND": "0 V"
},
"provides": [
  {
    "lib_type": "stdcell",
    "vt": "RVT"
  }
]
},

```

The file pointers, in this case, use the \$PDK and \$STDCELLS prefix as defined in the installs. The `corner` key tells Hammer what process and temperature corner that these files correspond to. The `supplies` key tells Hammer what the nominal supply for these cells are. The `provides` key has several sub-keys that tell Hammer what kind of library this is (examples include `stdcell`, `fiducials`, `io pad cells`, `bump`, and `level shifters`) and the threshold voltage flavor of the cells, if applicable. Adding the tech LEF for the technology with the `lib_type` set as `technology` is necessary for place and route.

## Library Filters

Library filters are defined in the `LibraryFilter` class in `hammer/tech/__init__.py`. These allow you to filter the entire set of libraries based on specific conditions, such as a file type or corner. Additional functions can be used to extract paths, strings, sort, and post-process the filtered libraries.

For a list of pre-built library filters, refer to the properties in the `LibraryFilterHolder` class in the same file, accessed as `hammer.tech.filters.<filter_method>`

## 2.2.4 Stackup

The `stackups` sets up the important metal layer information for Hammer to use. `stackups-example` shows an example of one metal layer in the `metals` list from the ASAP7 example tech plugin.

```

{"name": "M3", "index": 3, "direction": "vertical", "min_width": 0.072, "pitch": 0.144,
  "offset": 0.0, "power_strap_widths_and_spacings": [{"width_at_least": 0.0, "min_spacing": 0.072}], "power_strap_width_table": [0.072, 0.36, 0.648, 0.936, 1.224, 1.512]}

```

All this information is typically taken from the tech LEF and can be automatically filled in with a script. The metal layer name and layer number is specified. `direction` specifies the preferred routing direction for the layer. `min_width` and `pitch` specify the minimum width wire and the track pitch, respectively. `power_strap_widths_and_spacings` is a list of pairs that specify design rules relating to the widths of wires and minimum required spacing between them. This information is used by Hammer when drawing power straps to make sure it is conforming to some basic design rules.

## 2.2.5 Sites

The `sites` field specifies the unit standard cell size of the technology for Hammer.

```
"sites": [
  {"name": "asap7sc7p5t", "x": 0.216, "y": 1.08}
]
```

This is an example from the ASAP7 tech plugin in which the `name` parameter specifies the core site name used in the tech LEF, and the `x` and `y` parameters specify the width and height of the unit standard cell size, respectively.

## 2.2.6 Special Cells

The `special_cells` field specifies a set of cells in the technology that have special functions. `special-cells-example` shows a subset of the ASAP7 tech plugin for 2 types of cells: `tapcell` and `stdfiller`.

```
"special_cells": [
  {"cell_type": "tapcell", "name": ["TAPCELL_ASAP7_75t_L"]},
  {"cell_type": "stdfiller", "name": ["FILLER_ASAP7_75t_R", "FILLER_ASAP7_75t_L",
  ↳ "FILLER_ASAP7_75t_SL", "FILLER_ASAP7_75t_SRAM", "FILLERxp5_ASAP7_75t_R", "FILLERxp5_
  ↳ ASAP7_75t_L", "FILLERxp5_ASAP7_75t_SL", "FILLERxp5_ASAP7_75t_SRAM"]},
```

There are 8 `cell_type`s supported: `tiehicell`, `tielocell`, `tiehilocell`, `endcap`, `iofiller`, `stdfiller`, `decap`, and `tapcell`. Depending on the tech/tool, some of these cell types can only have 1 cell in the name list.

There is an optional `size` list. For each element in its corresponding name list, a `size` (type: `str`) can be given. An example of how this is used is for `decap` cells, where each listed cell has a typical capacitance, which a place and route tool can then use to place decaps to hit a target total decapacitance value. After characterizing the ASAP7 decaps using Voltus, the nominal capacitance is filled into the `size` list:

```
{ "cell_type": "decap", "name": ["DECAPx1_ASAP7_75t_R", "DECAPx1_ASAP7_75t_L", "DECAPx1_
  ↳ ASAP7_75t_SL", "DECAPx1_ASAP7_75t_SRAM", "DECAPx2_ASAP7_75t_R", "DECAPx2_ASAP7_75t_L",
  ↳ "DECAPx2_ASAP7_75t_SL", "DECAPx2_ASAP7_75t_SRAM", "DECAPx2b_ASAP7_75t_R", "DECAPx2b_
  ↳ ASAP7_75t_L", "DECAPx2b_ASAP7_75t_SL", "DECAPx2b_ASAP7_75t_SRAM", "DECAPx4_ASAP7_75t_R
  ↳ ", "DECAPx4_ASAP7_75t_L", "DECAPx4_ASAP7_75t_SL", "DECAPx4_ASAP7_75t_SRAM", "DECAPx6_
  ↳ ASAP7_75t_R", "DECAPx6_ASAP7_75t_L", "DECAPx6_ASAP7_75t_SL", "DECAPx6_ASAP7_75t_SRAM",
  ↳ "DECAPx10_ASAP7_75t_R", "DECAPx10_ASAP7_75t_L", "DECAPx10_ASAP7_75t_SL", "DECAPx10_
  ↳ ASAP7_75t_SRAM"], "size": ["0.39637 fF", "0.402151 fF", "0.406615 fF", "0.377040 fF",
  ↳ "0.792751 fF", "0.804301 fF", "0.813231 fF", "0.74080 fF", "0.792761 fF", "0.804309 fF",
  ↳ "0.813238 fF", "0.75409 fF", "1.5855 fF", "1.6086 fF", "1.62646 fF", "1.50861 fF",
  ↳ "2.37825 fF", "2.4129 fF", "2.43969 fF", "2.26224 fF", "3.96376 fF", "4.02151 fF", "4.
  ↳ 06615 fF", "3.7704 fF"]},
```

## 2.2.7 Don't Use, Physical-Only Cells

The `dont_use` list is used to denote cells that should be excluded due to things like bad timing models or layout. The `physical_only_cells` list is used to denote cells that contain only physical geometry, which means that they should be excluded from netlisting for simulation and LVS. Examples from the ASAP7 plugin are below:

```
"dont_use_list": [
  "ICGx*DC*",
  "AND4x1*",
```

(continues on next page)

(continued from previous page)

```

    "SDFLx2*",
    "A021x1*",
    "XOR2x2*",
    "OAI31xp33*",
    "OAI221xp5*",
    "SDFLx3*",
    "SDFLx1*",
    "A0I211xp5*",
    "OAI322xp33*",
    "OR2x6*",
    "A201A101Ixp25*",
    "XNOR2x1*",
    "OAI32xp33*",
    "FAx1*",
    "OAI21x1*",
    "OAI31xp67*",
    "OAI33xp33*",
    "A021x2*",
    "A0I32xp33*"
  ],
  "physical_only_cells_list": [
    "TAPCELL_ASAP7_75t_R", "TAPCELL_ASAP7_75t_L", "TAPCELL_ASAP7_75t_SL", "TAPCELL_ASAP7_
    ↪ 75t_SRAM",
    "TAPCELL_WITH_FILLER_ASAP7_75t_R", "TAPCELL_WITH_FILLER_ASAP7_75t_L", "TAPCELL_WITH_
    ↪ FILLER_ASAP7_75t_SL", "TAPCELL_WITH_FILLER_ASAP7_75t_SRAM",
    "FILLER_ASAP7_75t_R", "FILLER_ASAP7_75t_L", "FILLER_ASAP7_75t_SL", "FILLER_ASAP7_75t_
    ↪ SRAM",
    "FILLERxp5_ASAP7_75t_R", "FILLERxp5_ASAP7_75t_L", "FILLERxp5_ASAP7_75t_SL", "FILLERxp5_
    ↪ ASAP7_75t_SRAM"
  ],

```

## 2.3 Hammer Tech defaults.yml

The defaults.yml for a technology specifies some technology-specific *Hammer IR and Meta Variables* that should be left as default unless you desire to override them. Some of the them work directly with the keys in the tech.json.

Most of the keys in the defaults.yml are a part of the vlsi and technology namespaces. An example of the setup of the defaults.yml is located in hammer/technology/asap7/defaults.yml and certain important keys should be common to most technology plugins:

- vlsi.core.node defines the node that the place-and-route tool expects. It affects what kind of licenses are needed.
- vlsi.inputs should at least have the nominal supplies and a typical pair of characterized setup & hold corners.
- vlsi.technology needs to specify a placement\_site as defined in the technology LEF, a bump\_block\_cut\_layer to set blockages under bumps, and optional tap\_cell\_interval and tap\_cell\_offset for placing well taps.
- technology.core needs to specify the stackup to use, which layer the standard cell power rails are on, and a reference cell to draw the lowest layer power rails over.

Tool environment variables (commonly needed for DRC/LVS decks) and other necessary default options specific to the



technology should be set in this file. This file can also define technology-specific Hammer IR, in the the namespace `technology.<tech_name>.<key_name>`. As always, they can be overridden by other snippets of Hammer IR.

The data types for all keys in `defaults.yml` can be found in `defaults_types.yml`. When adding or overriding to `defaults.yml`, make sure that said data types are updated accordingly to prevent problems with the type checker.

## 2.4 Sky130 Technology Library

Hammer now supports the Skywater 130nm Technology process. The [SkyWater Open Source PDK](<https://skywater-pdk.readthedocs.io/>) is a collaboration between Google and SkyWater Technology Foundry to provide a fully open source Process Design Kit (PDK) and related resources, which can be used to create manufacturable designs at SkyWater's facility. The Skywater 130nm PDK files are located in a repo called [skywater-pdk](<https://github.com/google/skywater-pdk/>) A tool called [Open-PDKs (open\_pdk)]([https://github.com/RTimothyEdwards/open\\_pdk/](https://github.com/RTimothyEdwards/open_pdk/)) was developed to generate all the files typically found in a PDK. Open-PDKs uses the contents in *skywater-pdk*, and outputs files to a directory called *sky130A*.

[OpenLANE](<https://github.com/efabless/openlane/>) is an open-source RTL to GDSII VLSI flow that supports Sky130.

Before setting up the Sky130 PDK, we recommend familiarizing yourself with the process's documentation [summarized below](#Resources).

### 2.4.1 PDK Structure

OpenLANE expects a certain file structure for the Sky130 PDK. We recommend adhering to this file structure for Hammer as well. All the files reside in a root folder (named something like *skywater* or *sky130*). The environment variable `$PDK_ROOT` should be set to this folder's path:

```
export PDK_ROOT=<path-to-skywater-directory>
```

`$PDK_ROOT` contains the following:

- *skywater-pdk*
  - Original PDK source files
- *open\_pdk*
  - install of Open-PDKs tool
- *share/pdk/sky130A*
  - output files from Open-PDKs compilation process

### 2.4.2 NDA Files

Using commercial VLSI tools with this process requires files that are currently under NDA. If you have access to these, you will be able to use the Hammer VLSI flow out-of-the-box with the Sky130 process. Some of these NDA files, such as the Calibre DRC/LVS decks, rely on the environment variable `$PDK_HOME` containing the path to the version of the NDA files you are using. You may set it in the Hammer plugin, or as follows:

```
export PDK_HOME=<path-to-skywater-src-nda>/s8/<version#>
```

This plugin has been tested with PDK version *V2.0.1*; your mileage may vary for other versions.

### 2.4.3 Prerequisites for PDK Setup

- [Magic](<http://opencircuitdesign.com/magic/>)
  - required for *open\_pdk*s file generation
  - tricky to install, closely follow the directions on the *Install* page of the website
    - \* as the directions indicate, you will likely need to manually specify the location of the Tcl/Tk package installation using *–with-tcl* and *–with-tk*

If using conda, these installs alone caused the Magic install to work:

```
conda install -c intel tcl conda install -c anaconda tk conda install -c anaconda libglu
```

### 2.4.4 OpenRAM SRAMs

TODO: add overview here

```
git clone https://github.com/efabless/sky130\_sram\_macros.git
```

### 2.4.5 PDK Setup

In *\$PDK\_ROOT*, clone the skywater-pdk repo and generate the liberty files for each library:

```
git clone https://github.com/google/skywater-pdk.git cd skywater-pdk # Expect a large download! ~7GB
at time of writing. git submodule init libraries/*/latest git submodule update # Regenerate liberty files
make timing
```

Again in *\$PDK\_ROOT*, clone the open\_pdk repo and run the install process to generate the *sky130A* directory:

```
git clone https://github.com/RTimothyEdwards/open\_pdk.git cd open_pdk ./configure
--enable-sky130-pdk=$PDK_ROOT/skywater-pdk/libraries --prefix=$PDK_ROOT --enable-sram-
sky130=$PDK_ROOT/sky130_sram_macros # if using SRAMs make make install
```

This generates all the Sky130 PDK files and installs them to *\$PDK\_ROOT/share/pdk/sky130A*

### 2.4.6 Known DRC Issues

The Hammer flow generates DRC-clean digital logic when run with the Calibre DRC deck and SRAMs excluded.

TODO: summarize issues with SRAMs and open-source Magic DRC checks.

## Resources The good thing about this process being open-source is that most questions about the process are answerable through a google search. The tradeoff is that the documentation is a bit of a mess, and is currently scattered over a few pages and Github repos. We try to summarize these below.

Git repos:

- [SkyWater Open Source PDK](<https://github.com/google/skywater-pdk/>)
  - Git repo of the main Skywater 130nm files
- [Open-PDKs]([https://github.com/RTimothyEdwards/open\\_pdk/](https://github.com/RTimothyEdwards/open_pdk/))
  - Git repo of Open-PDKs tool that compiles the Sky130 PDK
- [Git repos on foss-eda-tools](<https://foss-eda-tools.googleusercontent.com/>)
  - Additional useful repos, such as Berkeley Analog Generator (BAG) setup

Documentation:

- [SkyWater SKY130 PDK's documentation](https://skywater-pdk.readthedocs.io)
  - Main documentation site for the PDK
- [Join the SkyWater PDK Slack Channel](https://join.skywater.tools/)
  - By far the best way to have questions about the process answered, with 80+ channels for different topics
- [Skywater130 Standard Cell and Primitives Overview](http://diychip.org/sky130/)
  - Additional useful documentation for the PDK
- [FOSSi Foundation YouTube Channel](https://www.youtube.com/c/FOSSiFoundation/)
  - Contains informational videos on the Sky130 process, OpenLANE flow, and Efabless Shuttle program

SRAMs:

- [OpenRAM](https://github.com/VLSIDA/OpenRAM/)
  - Open-source static random access memory (SRAM) compiler
- [OpenRAM pre-compiled macros](https://github.com/efabless/sky130\_sram\_macros/)
  - Precompiled sizes are 1kbytes, 2kbytes and 4kbytes

OpenLANE flow:

- [OpenLANE](https://github.com/efabless/openlane/)
  - Open-source RTL to GDSII flow
- [Magic](http://opencircuitdesign.com/magic/)
  - Open-source VLSI layout tool that handles DRC checks in OpenLANE

## 2.5 ASAP7 Technology Library

Hammer's default demonstration PDK is [ASAP7](http://asap.asu.edu/asap/). There are some special setup and known issues with this open PDK.

### 2.5.1 Setup and Environment

In addition to requirements for *hammer-vlsi*, using ASAP7 also requires: - ASAP7 PDK version 1p7, available [here on GitHub](https://github.com/The-OpenROAD-Project/asap7). We recommend downloading an archive of or shallow cloning the repository. - The Calibre deck tarball (downloaded separately from the website) must not be extracted. It should be placed in the directory specified by *technology.asap7.tarball\_dir*. For ease, this can be same directory as the repository. - Either the *gdstk* or *gdspy* GDS manipulation utility is required for 4x database downscaling. *gdstk* (available [here on GitHub](https://github.com/heitzmann/gdstk)), version >0.6) is highly recommended; however, because it is more difficult to install, *gdspy* (available [here on Github](https://github.com/heitzmann/gdspy/releases)), specifically version 1.4 can also be used instead, but it is much slower. - Calibre must be the DRC/LVS tool. The rule decks only support 2017-year Calibre versions.

\*At this moment, for BWRC affiliates, the environment needed for a *gdstk* or *gdspy* install is setup only on the LSF cluster machines. To install it, first enable the Python dev environment: `` scl enable rh-python36 bash `` *gdstk* must be built from source, with dependencies such as CMake, LAPACK, and ZLIB. All must be installed from source into your own environment. For the ZLIB build, make sure the *-fPIC* CFLAG is used. Then, in the *gdstk* source, add the following CMake flags at the top of *setup.py*: *-DZLIB\_LIBRARY=/path/to/your/libz.a* and *-DZLIB\_INCLUDE\_DIR=/path/to/your/include*. Then, in the *gdstk* source folder: `` python3 setup.py install`

--user ` For *gds Spy*, the installation is much easier, depending just on *pip* and *numpy*: ` `python -m pip install gds Spy` --user ` Or, replace the pip installation with installation from source in the *hammer/src/tools/gds Spy* submodule.

### 2.5.2 Dummy SRAMs

The ASAP7 plugin comes with a set of dummy SRAMs, which are **NOT** used by default (not included in the default tech.json).

They are **completely blank** (full obstructions on layers M1-M3, will not pass DRC & LVS). All pins are on M4, with the signal all on the left side and the power stripes running across. The M5 power stripes are able to connect up.

**All SRAMs are scaled up by 4x, so they are subject to the scaling script.**

*sram-cache-gen.py* generates *sram-cache.json* using *srams.txt*, which contains a list of available SRAMs in Hammer IR. *sram-cache.json* memories is used by MacroCompiler to insert these memories into the design.

Finally, the SRAMCompiler in *sram\_compiler/\_\_init\_\_.py* is used to generate the ExtraLibrarys (including .lib, .lef, .gds) needed by the particular design.

### 2.5.3 Known Issues

1. *ICG\*DC\** cells are set as don't use due to improper LEF width.
2. Many additional cells are not LVS clean, and are set as don't use.
3. Innovus tries to fix non-existent M3 and M5 enclosure violations, unfortunately lengthening violation fixing time. Ignore these when reviewing the violations in Innovus.
4. If you specify core margins in the placement constraints, they left and bottom margins should be a multiple of 0.384 to avoid DRC violations. Layer offsets for M4-M7 are adjusted manually to keep all wires on-grid.
5. Common expected DRC violations (at reasonable, <70% local density): - M(4,5,6,7).AUX.(1,2) only if the floor-plan size requirement above is not satisfied - V7.M8.AUX.2 and V2.M2.EN.1 due to incomplete via defs for V2 and V7 in power grid - FIN.S.1 appears to be incorrect, standard cell fins are indeed on the right pitch - LVT.W.1 caused by 0.5-width isolated-VT filler cells due to lack of implant layer spacing rules - LISD.S.3, LIG.S.4 due to some combinations of adjacent cells - V0.S.1 in ASYNC\_DFFHx1 - Various M4, GATE violations in/around dummy SRAMs

## 2.6 Nangate45 Technology Library

This is a PDK that comes packaged with [OpenROAD-flow](<https://github.com/The-OpenROAD-Project/OpenROAD-flow>). Right now, it is only validated against the OpenROAD-flow backend of hammer.

### 2.6.1 Dummy SRAMs

The Nangate45 plugin comes with fake SRAMs without behavioral models or GDS files. They only have timing libs for the TT corner (same as the pdk's stcell lib).

## **2.6.2 Supported Tools**

Only the OpenROAD-flow tools are supported, but compatibility between the two plugins is no longer maintained.



## HAMMER TOOL PLUGINS

This guide discusses the use and creation of CAD tool plugins in Hammer. A CAD tool plugin provides the actual implementation of Hammer APIs and outputs the TCL necessary to control its corresponding CAD tool.

Tool plugins must be structured as Python packages under the `hammer.action` package hierarchy. For example, if `dc` is a synthesis tool, it must be contained under `hammer/synthesis/dc`. The package should contain an class object named `'tool'` to create an instance of the tool. `tool` should be a class object of an appropriate subclass of `HammerTool` (e.g. `HammerSynthesisTool`).

### 3.1 Hammer CAD Tools

Hammer currently has open-source CAD tool plugins in the `hammer/<action>` folders and three repos for CAD tools from commercial vendors: `hammer-cadence-plugins`, `hammer-synopsys-plugins`, and `hammer-mentor-plugins`. `hammer-mentor-plugins` is a private repo since it contains tool-specific commands not yet cleared for public release. Access to them may be granted for Hammer users who already have licenses for those tools. See the note about [plugins access](#) for instructions for how to request access.

The structure of each repository is as follows:

- hammer
  - ACTION
    - \* TOOL\_NAME
      - `__init__.py` contains the methods needed to implement the tool
      - `defaults.yml` contains the default *Hammer IR and Meta Variables* needed by the tool

ACTION is the Hammer action name (e.g. `par`, `synthesis`, `drc`, etc.). TOOL\_NAME is the name of the tool, which is referenced in your configuration. For example, if ACTION is `par` and TOOL\_NAME is `par_tool_foo`, the configuration would reference it as `hammer.par.par_tool_foo`.

### 3.2 Setting up a Hammer CAD Tool Plugin

This guide will discuss what a Hammer user may do if they want to implement their own CAD tool plugin or extend the current CAD tool plugins. There are some basic mock-up examples of how this can be done in the `par` and `synthesis` directories inside `hammer/`.

### 3.2.1 Tool Class

Writing a tool plugin starts with writing the tool class. Hammer already provides a set of classes and mixins for a new tool to extend. For example, the Hammer Innovus plugin inherits from `HammerPlaceAndRouteTool` and `CadenceTool`.

### 3.2.2 Steps

Each tool implements a steps method. For instance, a portion of the steps method for a place-and-route tool may look like:

```
@property
def steps(self) -> List[HammerToolStep]:
    steps = [
        self.init_design,
        self.floorplan_design,
        self.route_design
    ]
    return self.make_steps_from_methods(steps)
```

Each of the steps are their own methods in the class that will write TCL that will execute with the tool.

### 3.2.3 Getting Settings

Hammer provides the method `get_setting("KEY_NAME")` for the tool to actually grab the settings from the user's input YML or JSON files. One example would be `self.get_setting("par.blockage_spacing")` so that Hammer can specify to the desired P&R tool what spacing to use around place and route blockages.

### 3.2.4 Writing TCL

Hammer provides two main methods for writing TCL to a file: `append` and `verbose_append`. Both do similar things but `verbose_append` will emit additional TCL code to print the command to the terminal upon execution.

### 3.2.5 Executing the Tool

When all the desired TCL has been written by various step methods, it is time to execute the tool itself. Hammer provides the method `run_executable(args, cwd=self.run_dir)` to do so. `args` is a Python list of flags to be run with the tool executable. `cwd=self.run_dir` sets the "current working directory" and allows the plugin to specify in what directory to execute the command.

### 3.2.6 Tool Outputs

After execution, the Hammer driver will emit a copy of the Hammer IR database in JSON format to the run directory as well as specific new fields created by the activity. The name of the output JSON files will be related to the activity type (e.g. `par-output.json` and `par-output-full.json` for the `par` activity). The `-full` version contains the entire Hammer IR database, while the other version contains only the output entries created by this activity. The individual fields are created when the `export_config_outputs` method is called. Each implementation of this tool must override this method with a new one that calls its `super` method and appends any additional output fields to the output dictionary, as necessary.



### 3.2.7 defaults.yml

Each tool may optionally implement a set of default *Hammer IR and Meta Variables*. These defaults may set different values for global Hammer IR configuration, or add new tool-specific keys. Tool-specific keys must have a namespace in the format `<action>.<tool_name>.<key_name>`. This file should have a companion `defaults_types.yml` that defines the required types as well.

## 3.3 OpenROAD Place-and-Route Tool Plugin

### 3.3.1 Tool Setup

First, follow [these directions](<https://github.com/The-OpenROAD-Project/OpenROAD#build>) to build and install OpenROAD.

### 3.3.2 Tool Steps

See `__init__.py` for the implementation of these steps.

```
# INITIALIZE init_design
# FLOORPLAN floorplan_design place_bumps macro_placement place_tapcells power_straps
# PLACE initial_global_placement io_placement global_placement resize detailed_placement
# CTS clock_tree add_fillers ## ROUTING global_route detailed_route
# FINISHING extraction write_design
```

### 3.3.3 Step Details

A few of the steps are detailed below, mainly to document some details that aren't present in the comments within the plugin. The plugin is still the best reference for what each step does, the explanations below are just more nuanced.

#### ### Macro Placement

Currently, there are two ways to place macros.

If `par.openroad.floorplan_mode == generate`, use the `vlsi.inputs.placement_constraints` key to specify the path and (x,y) coordinate of each macro in the design. The placement constraints for ALL macros must be specified, or OpenROAD will eventually throw an error about an unplaced instance. Note that the (x,y) position is fixed to the bottom left corner of the macro, then the orientation is applied. This means that for rotations besides `r0`, the (x,y) position WILL NOT correspond to the bottom left corner of the final placed macro (e.g. the (x,y) position for a macro with orientation `mx` will correspond to the top left corner of the final placed macro).

If `par.openroad.floorplan_mode == auto`, the macros will be placed automatically using OpenROAD's `macro_placement` command, leading to poor-quality but sane results.

#### ### Detailed Route

This is by far the most time-consuming step. For context, on a 16-core machine running one thread per core, a small RISC-V CPU with caches and a few digital peripherals, from [this tutorial in Chipyard](<https://chipyard.readthedocs.io/en/stable/VLSI/Sky130-OpenROAD-Tutorial.html>), takes several hours.

### Write Design OpenROAD uses an external tool called KLayout, which executes a custom Python script (`def2stream.py`) to write the final design to a GDS file. This is the same process used in the default VLSI flow from OpenROAD.

### 3.3.4 Issue Archiving

This plugin supports generating a tar.gz archive of the current build/par-rundir directory with a runme.sh script to reproduce the results. This feature is enabled with the `par.openroad.create_archive_mode` YAML key, which can take the following values: \* `after_error` - if OpenROAD errors, create an archive of the run \* `always` - create an archive after every par run (regardless of whether OpenROAD errors) \* `latest_run` - create an archive of latest par run (don't run OpenROAD, useful if OpenROAD gets stuck on endless optimization iterations but never actually "errors") \* `none` - never create an archive

Creating an archive does the following: \* Generates a issue directory based on the date/time, top module, OS platform \* Dumps the logger output \* Generates a runme.sh script \* Copies .tcl, .sdc, .pdn, and .lef files in par-rundir to the issue dir \* Copies all input Verilog to issue dir \* Copies all input LEFs to issue dir (some may not be used since they are hacked in read\_lef()) \* Copies all input LIBs to issue dir \* Hacks copied par.tcl script to remove all abspaths \* Hacks copied par.tcl script to comment out write\_gds block w/ KLayout \* Tars up issue dir

## HAMMER FLOW STEPS

This documentation will walk through the currently supported steps of the Hammer flow: synthesis, place-and-route, DRC, LVS, and simulation. Two other action types are supported, but not discussed: PCB collateral generation, and SRAM generation.

### 4.1 Hammer Actions

Hammer has a set of actions including synthesis, place-and-route, DRC, LVS, simulation, SRAM generation, and PCB collateral generation. All of the Hammer actions and their associated key inputs can be found in `hammer/config/defaults.yml` and are documented in detail in each action's documentation.

Hammer will automatically pass the output files of one action to subsequent actions if those actions require the files. Hammer does this with conversion steps (e.g. `syn-to-par` for synthesis outputs to place-and-route inputs), which map the outputs from one tool into the inputs of another (e.g. `synthesis.outputs.output_files` maps to `par.outputs.input_files`). The Hammer make infrastructure builds these conversion rules automatically, so using the Make infrastructure is recommended. See the *Hammer Buildfile* section for more details.

Hammer actions are implemented using tools. See the *Hammer CAD Tools* section for details about how these tools are set up.

### 4.2 Synthesis

Hammer supports synthesizing Verilog-based RTL designs to gate-level netlists. This action requires a tool plugin to implement `HammerSynthesisTool`.

#### 4.2.1 Synthesis Setup Keys

- Namespace: `vlsi.core`
  - `synthesis_tool`
    - \* Module of the synthesis tool, e.g. `hammer.synthesis.genus`

## 4.2.2 Synthesis Input Keys

- Namespace: `synthesis`
  - `inputs.input_files` ([])
    - \* A list of file paths to source files to be passed to the synthesis tool. The paths may be relative to the directory in which `hammer-vlsi` is called.
  - `inputs.top_module` (str)
    - \* Name of the top level verilog module of the design.
  - `clock_gating_mode` (str)
    - \* `auto`: turn on automatic clock gating inference in CAD tools
    - \* `empty`: do not do any clock gating

## 4.2.3 Synthesis Inputs

There are no prerequisites to running synthesis other than setting the keys that are described above.

## 4.2.4 Synthesis Outputs

- Mapped verilog file: `obj_dir/syn-rundir/{TOP_MODULE}.mapped.v`
- Mapped design SDF: `obj_dir/syn-rundir/{TOP_MODULE}.mapped.sdf`
- Synthesis output Hammer IR is contained in `obj_dir/syn-rundir/syn-output.json`
- Synthesis reports for gates, area, and timing are output in `obj_dir/syn-rundir/reports`

The synthesis output Hammer IR is converted to inputs for the P&R tool and the simulation tool by the Hammer `syn-to-par` and `syn-to-sim` commands, respectively.

## 4.2.5 Synthesis Commands

- Synthesis Command
  - `hammer-vlsi -e env.yml -p config.yml --obj_dir OBJ_DIR syn`
- Synthesis to Place-and-route
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/syn-rundir/syn-output.json --obj_dir OBJ_DIR syn-to-par`
- Synthesis to Simulation
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/syn-rundir/syn-output.json --obj_dir OBJ_DIR syn-to-sim`

## 4.3 Place-and-Route

Hammer has an action for placing and routing a synthesized design. This action requires a tool plugin to implement `HammerPlaceAndRouteTool`.

### 4.3.1 P&R Setup Keys

- Namespace: `vlsi.core`
  - `par_tool`
    - \* Module of the P&R tool, e.g. `hammer.par.innovus`

### 4.3.2 P&R Input Keys

- Namespace: `vlsi`
  - These are built-in Hammer APIs covered in *Hammer APIs*
- Namespace: `par`
  - `inputs.input_files ([])`
    - \* List of paths to post-synthesis netlists. Auto-populated after syn-to-par.
  - `inputs.top_module (str)`
    - \* Name of top RTL module to P&R. Auto-populated after syn-to-par.
  - `inputs.post_synth_sdc (str)`
    - \* Post-synthesis generated SDC. Auto-populated after syn-to-par.
  - `inputs.gds_map_mode (str)`
    - \* Specify which GDS layermap file to use. `auto` uses the technology-supplied file, whereas `manual` requires a file to specified via `inputs.gds_map_file`.
  - `inputs.gds_merge (bool)`
    - \* True tells the P&R tool to merge all library & macro GDS before streamout. Otherwise, only references will exist and merging needs to be done later, by a tool such as Calibre, gdstk, or gdspy.
  - `inputs.physical_only_cells_mode (str)`
    - \* Specifies which set of cells to exclude from SPICE netlist because they have no logical function. `auto` uses the technology-supplied list, whereas `manual` and `append` overrides and appends to the supplied list, respectively.
  - `submit (dict)`
    - \* Can override global settings for submitting jobs to a workload management platform.
  - `power_straps_mode (str)`
    - \* Power straps configuration. `generate` enables Hammer's power straps API, whereas `manual` requires a TCL script in `power_straps_script_contents`.
  - `blockage_spacing (Decimal)`
    - \* Global obstruction around every hierarchical sub-block and hard macro
  - `generate_power_straps_options (dict)`

- \* If `generate_power_straps_method` is `by_tracks`, this struct specifies all the options for the power straps API. See [Hammer APIs](#) for more detail.

### 4.3.3 P&R Inputs

There are no other prerequisites to running place & route other than setting the keys described above.

### 4.3.4 P&R Outputs

- Hierarchical (e.g. Cadence ILMs) and between-step snapshot databases in `OBJ_DIR/par-rundir`
- GDSII file with final design: `{OBJ_DIR}/par-rundir/{TOP_MODULE}.gds`
- Verilog gate-level netlist: `{OBJ_DIR}/par-rundir/{TOP_MODULE}.lvs.v`
- SDF file for post-par simulation: `{OBJ_DIR}/par-rundir/{TOP_MODULE}.sdf`
- Timing reports: `{OBJ_DIR}/par-rundir/timingReports`
- A script to open the final chip: `{OBJ_DIR}/par-rundir/generated_scripts/open_chip`

P&R output Hammer IR `{OBJ_DIR}/par-rundir/par-output.json` is converted to inputs for the DRC, LVS, and simulation tools by the `par-to-drc`, `par-to-lvs`, and `par-to-sim` actions, respectively.

### 4.3.5 P&R Commands

- P&R Command (after `syn-to-par` is run)
  - `hammer-vlsi -e env.yml -p {OBJ_DIR}/par-input.json --obj_dir OBJ_DIR par`
- P&R to DRC
  - `hammer-vlsi -e env.yml -p config.yml -p {OBJ_DIR}/par-rundir/par-output.json --obj_dir OBJ_DIR par-to-drc`
- P&R to LVS
  - `hammer-vlsi -e env.yml -p config.yml -p {OBJ_DIR}/par-rundir/par-output.json --obj_dir OBJ_DIR par-to-lvs`
- P&R to Simulation
  - `hammer-vlsi -e env.yml -p config.yml -p {OBJ_DIR}/par-rundir/par-output.json --obj_dir OBJ_DIR par-to-sim`

## 4.4 DRC

Hammer has an action for running design rules check (DRC) on a post-place-and-route GDS. This action requires a tool plugin to implement `HammerDRCTool`.

### 4.4.1 DRC Setup Keys

- Namespace: `vlsi.core`
  - `drc_tool`
    - \* Module of the DRC tool e.g. `hammer.drc.calibre`

### 4.4.2 DRC Input Keys

- Namespace: `drc`
  - `inputs.top_module` (str)
    - \* Name of top RTL module to run DRC on. Auto-populated after `par-to-drc`.
  - `inputs.layout_file` (str)
    - \* GDSII file from place-and-route. Auto-populated after `par-to-drc`.
  - `inputs.additional_drc_text_mode` (str)
    - \* Chooses what custom DRC commands to add to the run file. `auto` selects the one provided in the *Hammer Tech JSON*.
  - `inputs.drc_rules_to_run` ([str])
    - \* This selects a subset of the rules given in the technology's Design Rule Manual (DRM). The format of these rules will be technology- and tool-specific.
  - `submit` (dict)
    - \* Can override global settings for submitting jobs to a workload management platform.

### 4.4.3 DRC Inputs

There are no other prerequisites to running DRC other than setting the keys described above.

### 4.4.4 DRC Outputs

- DRC results report and database in `{OBJ_DIR}/drc-rundir`
- A run file: `{OBJ_DIR}/drc-rundir/drc_run_file`
- A script to interactively view the DRC results: `{OBJ_DIR}/drc-rundir/generated_scripts/view_drc`

### 4.4.5 DRC Commands

- DRC Command (after `par-to-drc` is run)
  - `hammer-vlsi -e env.yml -p {OBJ_DIR}/drc-input.json --obj_dir OBJ_DIR drc`

## 4.5 LVS

Hammer has an action for running layout-versus-schematic (LVS) on a post-place-and-route GDS and gate-level netlist. This action requires a tool plugin to implement `HammerLVSTool`.

### 4.5.1 LVS Setup Keys

- Namespace: `vlsi.core`
  - `lvs_tool`
    - \* Module of the lvs tool, e.g. `hammer.lvs.calibre`

### 4.5.2 LVS Input Keys

- Namespace: `lvs`
  - `inputs.top_module` (str)
    - \* Name of top RTL module to run LVS on. Auto-populated after `par-to-lvs`.
  - `inputs.layout_file` (str)
    - \* GDSII file from P&R. Auto-populated after `par-to-lvs`.
  - `inputs.schematic_files` ([str])
    - \* Netlists from P&R'd design and libraries. Auto-populated after `par-to-lvs`.
  - `inputs.additional_lvs_text_mode` (str)
    - \* Chooses what custom LVS commands to add to the run file. `auto` selects the one provided in the *Hammer Tech JSON*.
  - `submit` (dict)
    - \* Can override global settings for submitting jobs to a workload management platform.

### 4.5.3 LVS Inputs

There are no other prerequisites to running LVS other than setting the keys described above.

### 4.5.4 LVS Outputs

- LVS results report and database in `{OBJ_DIR}/lvs-rundir`
- A run file: `{OBJ_DIR}/lvs-rundir/lvs_run_file`
- A script to interactively view the LVS results: `{OBJ_DIR}/lvs-rundir/generated_scripts/view_lvs`



### 4.5.5 LVS Commands

- LVS Command (after par-to-lvs is run)
  - `hammer-vlsi -e env.yml -p {OBJ_DIR}/lvs-input.json --obj_dir OBJ_DIR lvs`

## 4.6 Simulation

Hammer supports RTL, post-synthesis, and post-P&R simulation. It provides a simple API to add flags to the simulator call and automatically passes in collateral to the simulation tool from the synthesis and place-and-route outputs. This action requires a tool plugin to implement `HammerSimTool`.

### 4.6.1 Simulation Setup Keys

- Namespace: `vlsi.core`
  - `sim_tool`
    - \* Module of the simulation tool, e.g. `hammer.sim.vcs`

### 4.6.2 Simulation Input Keys

- Namespace: `sim.inputs`
  - `input_files ([])`
    - \* A list of file paths to source files (verilog sources, testharness/testbench, etc.) to be passed to the synthesis tool (both verilog and any other source files needed). The paths may be relative to the directory in which `hammer-vlsi` is called.
  - `top_module (str)`
    - \* Name of the top level module of the design.
  - `options ([str])`
    - \* Any options that are passed into this key will appear as plain text flags in the simulator call.
  - `defines ([str])`
    - \* Specifies define options that are passed to the simulator. e.g. when using VCS, this will be added as `+define+{DEFINE}`.
  - `compiler_cc_opts ([str])`
    - \* Specifies C compiler options when generating the simulation executable. e.g. when using VCS, each `compiler_cc_opt` will be added as `-CFLAGS {compiler_cc_opt}`.
  - `compiler_ld_opts ([str])`
    - \* Specifies C linker options when generating the simulation executable. e.g. when using VCS, each `compiler_ld_opt` will be added as `-LDFLAGS {compiler_ld_opt}`.
  - `timescale (str)`
    - \* Plain string that specifies the simulation timescale. e.g. when using VCS, `sim.inputs.timescale: '1ns/10ps'` would be passed as `-timescale=1ns/10ps`
  - `benchmarks ([str])`

- \* A list of benchmark binaries that will be run with the simulator to test the design if the testbench requires it. For example, this may be RISC-V binaries. If unspecified or left empty, the simulation will execute as normal.
- `parallel_runs` (int)
  - \* Maximum number of simulations to run in parallel. -1 denotes all in parallel. 0 or 1 denotes serial execution.
- `tb_name` (str)
  - \* The name of the testbench/test driver in the simulation.
- `tb_dut` (str)
  - \* Hierarchical path to the top level instance of the “dut” from the testbench.
- `level` ("rtl" or "gl")
  - \* This defines whether the simulation being run is at the RTL level or at the gate level.
- `all_regs` (str)
  - \* Path to a list of all registers in the design, typically generated from the synthesis or P&R tool. This is used in gate level simulation to initialize register values.
- `seq_cells` (str)
  - \* Path to a list of all sequential standard cells in the design, typically generated from the synthesis or P&R tool. This is used in gate level simulation.
- `sdf_file` (str)
  - \* Path to Standard Delay Format file used in timing annotated simulations.
- `gl_register_force_value` (0 or 1)
  - \* Defines what value all registers will be initialized to for gate level simulations.
- `timing_annotated` (false or true)
  - \* Setting to false means that the simulation will be entirely functional. Setting to true means that the simulation will be time annotated based on synthesis or P&R results.
- `saif.mode` ("time", "trigger", "trigger\_raw", or "full")
  - \* "time": pair with `saif.mode.start_time` and `saif.mode.end_time` (TimeValue) to dump between 2 timestamps
  - \* "trigger": inserts a trigger into the simulator run script
  - \* "trigger\_raw": inserts a given start/end trigger tcl script into the simulator run script. Specify scripts with `saif.mode.start_trigger_raw` and `saif.mode.end_trigger_raw` (str)
  - \* "full": dump the full simulation
- `execution_flags` ([str])
  - \* Each string in this list will be passed as an option when actually executing the simulation executable generated from the previous arguments.
  - \* Can also use `execution_flags_prepend` and `execution_flags_append` for additional execution flags
- `execute_sim` (true or false)
  - \* Determines whether or not the simulation executable that is generated with the above inputs with the given flags or if the executable will just be generated.

### 4.6.3 Simulation Inputs

There are no prerequisites to running an RTL simulation other than setting the keys that are described above. Running the `syn-to-sim` action after running synthesis will automatically generate the Hammer IR required to pipe the synthesis outputs to the Hammer simulation tool, and should be included in the Hammer call, as demonstrated in the “Post-Synthesis Gate Level Sim” command below. The same goes for post-place-and-route simulations. The required files for these simulations (SDF, SPEF, etc.) are generated and piped to the simulation tool in the corresponding action’s outputs.

The Hammer simulation tool will initialize register values in the simulation, as that is of particular need when simulating Chisel-based designs, to deal with issues around x-pessimism.

### 4.6.4 Simulation Outputs

The simulation tool is able to output waveforms for the simulation. All of the relevant outputs of the simulation can be found in `OBJ_DIR/sim-rundir/`.

### 4.6.5 Simulation Commands

- RTL Simulation Command
  - `hammer-vlsi -e env.yml -p config.yml --obj_dir OBJ_DIR sim`
- Synthesis to Sim
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/syn-rundir/syn-output.json -o OBJ_DIR/syn-to-sim_input.json --obj_dir OBJ_DIR syn-to-sim`
- Post-Synthesis Gate Level Sim
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/syn-to-sim_input.json --obj_dir OBJ_DIR sim`
- P&R to Simulation
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/par-rundir/par-output.json -o OBJ_DIR/par-to-sim_input.json --obj_dir OBJ_DIR par-to-sim`
- Post-P&R Gate Level Sim
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/par-to-sim_input.json --obj_dir OBJ_DIR sim`

## 4.7 Power

Hammer supports RTL, post-synthesis, and post-P&R power analysis. It provides a simple API to add flags to the power tool call and automatically passes in collateral to the power tool from the other tool steps. This action requires a tool plugin to implement `HammerPowerTool`.

### 4.7.1 Power Setup Keys

- Namespace: `vlsi.core`
  - **power\_tool**
    - \* Module of the power tool, e.g. `hammer.power.voltus`

### 4.7.2 Simulation Input Keys

- Namespace: `power.inputs`
  - **database** (str)
    - \* Path to the place and route database of the design to be analyzed. This path may be relative to the directory in which `hammer-vlsi` is called.
  - **tb\_name** (str)
    - \* The name of the testbench/test driver in the simulation.
  - **tb\_dut** (str)
    - \* Hierarchical path to the to top level instance of the “dut” from the testbench.
  - **spefs** ([str])
    - \* List of paths to all spef (parasitic extraction) files for the design. This list may include a spef file per MMMC corner. Paths may be relative to the directory in which `hammer-vlsi` is called.
  - **waveforms** ([str])
    - \* List of paths to waveforms to be used for dynamic power analysis. Paths may be relative to the directory in which `hammer-vlsi` is called.
  - **start\_times** ([TimeValue])
    - \* List of analysis start times corresponding to each of the **waveforms** used for dynamic power analysis.
  - **end\_times** ([TimeValue])
    - \* List of analysis end times corresponding to each of the **waveforms** used for dynamic power analysis.
  - **saifs** ([str])
    - \* List of paths to SAIF (activity files) for dynamic power analysis. Generally generated by a gate-level simulation. Paths may be relative to the directory in which `hammer-vlsi` is called.
  - **extra\_corners\_only** (bool)
    - \* If overridden to `true`, the power tool will report for only the extra MMMC corners, saving runtime. The typical use case is to only report power and rail analysis for a typical/nominal corner.
  - **input\_files** ([str])
    - \* A list of the paths to the design inputs files (HDL or netlist) for power analysis.
  - **sdh** (str)
    - \* Path to SDC input file.
  - **report\_configs** ([dict])
    - \* List of report configs that specify `PowerReport` structs.
  - **level** (FlowLevel)

- \* Power analysis mode for different levels of the VLSI flow. The available options are `rtl`, `syn`, and `par`.
- **top\_module (str)**
  - \* Top RTL module for power analysis.

### 4.7.3 Power Inputs

Hammer's power analysis can be run with an RTL input, or post-synthesis or post-place-and-route (and with corresponding simulations). Auto-translation of of Hammer IR to the power tool from those outputs are accomplished using the `sim-rtl-to-power`, `syn-to-power`, `sim-syn-to-power`, `par-to-power`, and `sim-par-to-power` actions, as demonstrated below. The required files for power analysis (database, SAIF, SPEF, etc.) are generated and piped to the power tool from the pre-requisite action's outputs.

### 4.7.4 Power Outputs

The power tool outputs static and active power estimations into the `OBJ_DIR/power-rundir/` directory. Exact report format may vary by tool used.

### 4.7.5 Power Commands

RTL Power Analysis:

- RTL Sim
  - `hammer-vlsi -e env.yml -p config.yml --obj_dir OBJ_DIR sim-rtl`
- Simulation to Power
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/sim-rundir/sim-rtl-output.json -o OBJ_DIR/sim-rtl-to-power_input.json --obj_dir OBJ_DIR sim-rtl-to-power`
- Power
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/sim-rtl-to-power_input.json --obj_dir OBJ_DIR power-rtl`

Post-synthesis Power Analysis:

- Syn to Power
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/syn-rundir/syn-output.json -o OBJ_DIR/syn-to-power_input.json --obj_dir OBJ_DIR syn-to-power`
- Syn to Simulation
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/syn-rundir/syn-output.json -o OBJ_DIR/syn-to-sim_input.json --obj_dir OBJ_DIR syn-to-sim`
- Post-Syn Gate Level Sim
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/syn-to-sim_input.json --obj_dir OBJ_DIR sim-syn`
- Simulation to Power
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/sim-rundir/sim-syn-output.json -o OBJ_DIR/sim-syn-to-power_input.json --obj_dir OBJ_DIR sim-syn-to-power`

- Power

```
- hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/syn-to-power_input.json -p  
OBJ_DIR/sim-syn-to-power_input.json --obj_dir OBJ_DIR power-syn
```

Post-P&R Power Analysis:

- P&R to Power

```
- hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/par-rundir/par-output.json -o  
OBJ_DIR/par-to-power_input.json --obj_dir OBJ_DIR par-to-power
```

- P&R to Simulation

```
- hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/par-rundir/par-output.json -o  
OBJ_DIR/par-to-sim_input.json --obj_dir OBJ_DIR par-to-sim
```

- Post-P&R Gate Level Sim

```
- hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/par-to-sim_input.json --obj_dir  
OBJ_DIR sim-par
```

- Simulation to Power

```
- hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/sim-rundir/sim-par-output.json  
-o OBJ_DIR/sim-par-to-power_input.json --obj_dir OBJ_DIR sim-par-to-power
```

- Power

```
- hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/par-to-power_input.json -p  
OBJ_DIR/sim-par-to-power_input.json --obj_dir OBJ_DIR power-par
```

## 4.8 Formal Verification

Hammer supports post-synthesis and post-P&R formal verification. It provides a simple API to provide a set of reference and implementation inputs (e.g. Verilog netlists) to perform formal verification checks such as logical equivalence checking (LEC). This action requires a tool plugin to implement `HammerFormalTool`.

### 4.8.1 Formal Verification Setup Keys

- Namespace: `vlsi.core`

- `formal_tool`

- \* Python module of the formal verification tool e.g. `hammer.formal.conformal`

### 4.8.2 Formal Verification Input Keys

- Namespace: `formal.inputs`

- `check` (str)

- \* Name of the formal verification check that is to be performed. Support varies based on the specific tool plugin. Potential check types/algorithms could be “lec”, “power”, “constraint”, “cdc”, “property”, “eco”, and more. At the moment, only “lec” is supported.

- `input_files` ([])

- \* A list of file paths to implementation source files (verilog, vhd, spice, liberty, etc.) to be passed to the formal verification tool. For a LEC tool, this would be the sources for the “revised” design. The paths may be relative to the directory in which `hammer-vlsi` is called.
- `reference_files` ([])
  - \* A list of file paths to reference source files (verilog, vhd, spice, liberty, etc.) to be passed to the formal verification tool. For a LEC tool, this would be the sources for the “golden” design. The paths may be relative to the directory in which `hammer-vlsi` is called.
- `top_module` (str)
  - \* Name of the top level module of the design to be verified.

### 4.8.3 Formal Verification Inputs

Running the `syn-to-formal` action after running synthesis will automatically generate the Hammer IR required to pipe the synthesis outputs to the Hammer formal verification tool, and should be included in the Hammer call, as demonstrated in the “Post-Synthesis Formal Verification” command below. The same goes for post-place-and-route formal verification.

At this time, only netlists are passed to the formal verification tool. Additional files (SDCs, CPFs, etc.) are needed for more advanced formal verification checks but are not yet currently supported. Similarly, formal verification on the behavioral RTL is also not yet supported.

### 4.8.4 Formal Verification Outputs

The formal verification tool produces reports in `OBJ_DIR/formal-rundir/`. Outputs from formal verification flows such as engineering change order (ECO) patches are not yet supported.

### 4.8.5 Formal Verification Commands

- Synthesis to Formal Verification
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/syn-rundir/syn-output.json -o OBJ_DIR/syn-to-formal_input.json --obj_dir OBJ_DIR syn-to-formal`
- Post-Synthesis Formal Verification
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/syn-to-formal_input.json --obj_dir OBJ_DIR formal`
- P&R to Formal Verification
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/par-rundir/par-output.json -o OBJ_DIR/par-to-formal_input.json --obj_dir OBJ_DIR par-to-formal`
- Post-P&R Formal Verification
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/par-to-formal_input.json --obj_dir OBJ_DIR formal`

## 4.9 Static Timing Analysis

Hammer supports post-synthesis and post-P&R static timing analysis. It provides a simple API to provide a set of design inputs (e.g. Verilog netlists, delay files, parasitics files) to perform static timing analysis (STA) for design signoff. This action requires a tool plugin to implement `HammerTimingTool`.

### 4.9.1 STA Setup Keys

- Namespace: `vlsi.core`
  - `timing_tool`
    - \* Module of the STA tool, e.g. `hammer.timing.tempus`

### 4.9.2 STA Input Keys

- Namespace: `timing.inputs`
  - `input_files` ([])
    - \* A list of file paths to the Verilog gate-level netlist to be passed to the STA tool. The paths may be relative to the directory in which `hammer-vlsi` is called.
  - `top_module` (str)
    - \* Name of the top level module of the design to be timed.
  - `post_synth_sdc` (str)
    - \* Post-synthesis generated SDC. Auto-populated after syn-to-timing.
  - `spefs` ([str])
    - \* List of paths to all spef (parasitic extraction) files for the design. This list may include a spef file per MMMC corner. Paths may be relative to the directory in which `hammer-vlsi` is called.
  - `sdf_file` (str)
    - \* Path to Standard Delay Format file. Auto-populated after syn-to-timing and par-to-timing.
  - `max_paths` (int)
    - \* Maximum number of timing paths to report from the STA tool. Large limits may hurt tool runtime.

### 4.9.3 STA Inputs

Running the `syn-to-timing` action after running synthesis will automatically generate the Hammer IR required to pipe the synthesis outputs to the Hammer STA tool, and should be included in the Hammer call, as demonstrated in the “Post-Synthesis STA” command below. The same goes for post-place-and-route STA.



## 4.9.4 STA Outputs

The STA tool produces reports in OBJ\_DIR/timing-rundir/. Outputs from advanced STA flows such as engineering change order (ECO) patches are not yet supported.

## 4.9.5 STA Commands

- Synthesis to STA
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/syn-rundir/syn-output.json -o OBJ_DIR/syn-to-timing_input.json --obj_dir OBJ_DIR syn-to-timing`
- Post-Synthesis STA
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/syn-to-timing_input.json --obj_dir OBJ_DIR timing`
- P&R to STA
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/par-rundir/par-output.json -o OBJ_DIR/par-to-timing_input.json --obj_dir OBJ_DIR par-to-timing`
- Post-P&R STA
  - `hammer-vlsi -e env.yml -p config.yml -p OBJ_DIR/par-to-timing_input.json --obj_dir OBJ_DIR timing`



## HAMMER USE

This documentation will walk through more advanced features of the Hammer infrastructure. You will learn about Hammer's APIs, flow control, how to write hooks, the supported build infrastructure, and how to set up hierarchical flows.

### 5.1 Hammer IR and Meta Variables

#### 5.1.1 Hammer IR

Hammer IR is the primary standardized data exchange format of Hammer. Hammer IR standardizes physical design constraints such as placement constraints and clock constraints. In addition, the Hammer IR also standardizes communication among and to Hammer plugins, including tool control (e.g. loading tools, etc) and configuration options (e.g. number of CPUs).

#### 5.1.2 The hammer-config library

The [hammer-config library](#) is the part of Hammer responsible for parsing Hammer YAML/JSON configuration files into Hammer IR. Hammer IR is used for the standardization and interchange of data between the different parts of Hammer and Hammer plugins.

---

**Note:** There is a built-in order of precedence, from lowest to highest:

- 1) Hammer's `defaults.yml`
- 2) Tool plugin's `defaults.yml`
- 3) Tech plugin's `defaults.yml`
- 4) User's Hammer IR (using `-p` on the command line)

JSON/YAML files specified with `-p` later on the command line have higher precedence and keys appearing later if duplicated in a given file also take precedence. In the examples below, "Level #" will be used to denote the level of precedence of the configuration snippet.

---

The `get_setting()` method is available to all Hammer technology/tool plugins and hooks (see [Extending Hammer with Hooks](#)).

### 5.1.3 Basics

```
foo:
  bar:
    adc: "yes"
    dac: "no"
```

The basic idea of Hammer IR in YAML/JSON format is centered around a hierarchically nested tree of YAML/JSON dictionaries. For example, the above YAML snippet is translated to two variables which can be queried in code - `foo.bar.adc` would have `yes` and `foo.bar.dac` would have `no`.

### 5.1.4 Overriding

Hammer IR snippets frequently “override” each other. For example, a technology plugin might provide some defaults which a specific project can override with a project YAML snippet.

For example, if the base snippet contains `foo: 12345` and a snippet in a file of higher precedence contains `foo: 54321`, then `get_setting("foo")` would return `54321`.

### 5.1.5 Meta actions

Sometimes it is desirable that variables are not completely overwritten, but instead modified.

For example, say that the technology plugin provides:

```
vlsi.tech.foobar65.bad_cells: ["NAND4X", "NOR4X"]
```

And let’s say that in our particular project, we find it undesirable to use the `NAND2X` and `NOR2X` cells. However, if we simply put the following in our project YAML, the references to `NAND4X` and `NOR4X` disappear and we don’t want to have to copy the information from the base plugin, which may change, or which may be proprietary, etc.

```
vlsi.tech.foobar65.bad_cells: ["NAND2X", "NOR2X"]
```

The solution is **meta variables**. This lets Hammer’s config parser know that instead of simply replacing the base variable, it should do a particular special action. Any config variable can have `_meta` suffixed into a new variable with the desired meta action.

In this example, we can use the `append` meta action:

```
vlsi.tech.foobar65.bad_cells: ["NAND2X", "NOR2X"]
vlsi.tech.foobar65.bad_cells_meta: append
```

This will yield the desired result of `["NAND4X", "NOR4X", "NAND2X", "NOR2X"]` when `get_setting("vlsi.tech.foobar65.bad_cells")` is called in the end.

### 5.1.6 Applying multiple meta actions

Multiple meta actions can be applied sequentially if the `_meta` variable is an array. Example:

In Level 1:

```
foo.flash: yes
```

In Level 2 (located at `/opt/foo`):

```
foo.pipeline: "CELL_${foo.flash}.lef"
foo.pipeline_meta: ['subst', 'prependlocal']
```

Result: `get_setting("foo.pipeline")` returns `/opt/foo/CELL_yes.lef`.

### 5.1.7 Common meta actions

- `append`: append the elements provided to the base list. (See the above `vlsi.tech.foo65.bad_cells` example.)
- `subst`: substitute variables into a string.

Base:

```
foo.flash: yes
```

Meta:

```
foo.pipeline: "${foo.flash}man"
foo.pipeline_meta: subst
```

Result: `get_setting("foo.flash")` returns `yesman`

- `lazysubst`: by default, variables are only substituted from previous configs. Using `lazysubst` allows us to defer the substitution until the very end.

Example without `lazysubst`:

Level 1:

```
foo.flash: yes
```

Level 2:

```
foo.pipeline: "${foo.flash}man"
foo.pipeline_meta: subst
```

Level 3:

```
foo.flash: no
```

Result: `get_setting("foo.flash")` returns `yesman`

Example with `lazysubst`:

Level 1:

```
foo.flash: yes
```

Level 2:

```
foo.pipeline: "${foo.flash}man"
foo.pipeline_meta: lazysubst
```

Level 3:

```
foo.flash: no
```

Result: `get_setting("foo.flash")` returns `noman`

- **crossref** - directly reference another setting. Example:

Level 1:

```
foo.flash: yes
```

Level 2:

```
foo.mob: "foo.flash"
foo.mob_meta: crossref
```

Result: `get_setting("foo.mob")` returns `yes`

- **transclude** - transclude the given path. Example:

Level 1:

```
foo.bar: "/opt/foo/myfile.txt"
foo.bar_meta: transclude
```

Result: `get_setting("foo.bar")` returns `<contents of /opt/foo/myfile.txt>`

- **prependlocal** - prepend the local path or package resource directory of this config file. Example:

Level 1 (located at `/opt/foo`):

```
foo.bar: "myfile.txt"
foo.bar_meta: prependlocal
```

Result: `get_setting("foo.mob")` returns `/opt/foo/myfile.txt`

- **deepsubst** - like `subst` but descends into sub-elements. Example:

Level 1:

```
foo.bar: "123"
```

Level 2:

```
foo.bar:
  baz: "${foo.bar}45"
  quux: "32${foo.bar}"
foo.bar_meta: deepsubst
```

Result: `get_setting("foo.bar.baz")` returns `12345` and `get_setting("foo.bar.quux")` returns `32123`

### 5.1.8 Type Checking

Any existing configuration file can and should be accompanied with a corresponding configuration types file. This allows for static type checking of any key when calling `get_setting`. The file should contain the same keys as the corresponding configuration file, but can contain the following as values:

- primitive types (`int`, `str`, etc.)
- collection types (`list`)
- collections of key-value pairs (`list[dict[str, str]]`, `list[dict[str, list]]`, etc.) These values are turned into custom constraints (e.g. `PlacementConstraint`, `PinAssignment`) later in the Hammer workflow, but the key value pairs are not type-checked any deeper.
- optional forms of the above (`Optional[str]`)
- the wildcard `Any` type

Hammer will perform the same without a types file, but it is highly recommended to ensure type safety of any future plugins.

### 5.1.9 Key History

With the `ruamel.yaml` package, Hammer can emit what files have modified any configuration keys in YAML format. The file is named `{action}-output-history.yml` and is located in the output folder of the given action.

Example with the file `test-config.yml`:

```
synthesis.inputs:
  input_files: ["foo", "bar"]
  top_module: "z1top.xdc"

vlsi:
  core:
    technology: "hammer.technology.nop"

    synthesis_tool: "hammer.synthesis.nop"
```

`test/syn-rundir/syn-output-history.yml` after executing the command `hammer-vlsi -p test-config.yml --obj_dir test syn`:

```
synthesis.inputs.input_files: # Modified by: test-config.yml
- LICENSE
- README.md
synthesis.inputs.top_module: z1top.xdc # Modified by: test-config.yml

vlsi.core.technology: nop # Modified by: test-config.yml
vlsi.core.synthesis_tool: hammer.synthesis.nop # Modified by: test-config.yml
```

Example with the files `test-config.yml` and `test-config2.yml`, respectively:

```
synthesis.inputs:
  input_files: ["foo", "bar"]
  top_module: "z1top.xdc"

vlsi:
```

(continues on next page)

(continued from previous page)

```

core:
  technology: "hammer.technology.nop"

  synthesis_tool: "hammer.synthesis.nop"

par.inputs:
  input_files: ["foo", "bar"]
  top_module: "z1top.xdc"

vlsi:
  core:
    technology: "${foo.subst}"

    par_tool: "hammer.par.nop"

foo.subst: "hammer.technology.nop2"

```

test/syn-rundir/par-output-history.yml after executing the command `hammer-vlsi -p test-config.yml -p test-config2.yml --obj_dir test syn-par`:

```

foo.subst: hammer.technology.nop2 # Modified by: test-config2.yml
par.inputs.input_files: # Modified by: test-config2.yml
- foo
- bar
par.inputs.top_module: z1top.xdc # Modified by: test-config2.yml
synthesis.inputs.input_files: # Modified by: test-config.yml
- foo
- bar
synthesis.inputs.top_module: z1top.xdc # Modified by: test-config.yml
vlsi.core.technology: hammer.technology.nop2 # Modified by: test-config.yml, ↵
↪ test-config2.yml
vlsi.core.synthesis_tool: hammer.synthesis.nop # Modified by: test-config.yml
vlsi.core.par_tool: hammer.par.nop # Modified by: test-config2.yml

```

### 5.1.10 Key Description Lookup

With the `ruamel.yaml` package, Hammer can execute the `info` action, allowing users to look up the description of most keys. The comments must be structured like so in order to be read properly:

```

foo: bar # this is a comment
# this is another comment for the key "foo"

```

Hammer will take the descriptions from any `defaults.yml` files.

Running `hammer-vlsi -p test-config.yml info` (assuming the above configuration is in `defaults.yml`):

```

foo.bar
Select from the current level of keys: foo.bar

apple

```

(continues on next page)



(continued from previous page)

```

price
Select from the current level of keys: apple
-----
Key: foo.bar.apple
Value: banana
Description: # type of fruit
History: ["defaults.yml"]
-----

Continue querying keys? [y/n]: y

foo.bar
Select from the current level of keys: foo.bar

apple
price
Select from the current level of keys: price
-----
Key: foo.bar.price
Value: 2
Description: # price of fruit
History: ["defaults.yml"]
-----

```

Keys are queried post-resolution of all meta actions, so their values correspond to the project configuration after other actions like `syn` or `par`.

### 5.1.11 Reference

For a more comprehensive view, please consult the `hammer.config` API documentation in its implementation here:

- [https://github.com/ucb-bar/hammer/blob/master/hammer/config/config\\_src.py](https://github.com/ucb-bar/hammer/blob/master/hammer/config/config_src.py)
- [https://github.com/ucb-bar/hammer/blob/master/tests/test\\_config.py](https://github.com/ucb-bar/hammer/blob/master/tests/test_config.py)

In `config_src.py`, most supported meta actions are contained in the `directives` list of the `get_meta_directives` method.

## 5.2 Hammer APIs

Hammer has a growing collection of APIs that use objects defined by the technology plugin, such as `stackups` and `special cells`. They expose useful extracted information from Hammer IR to other methods, such as in tool plugins that will implement this information in a tool-compatible manner.

For syntax details about the Hammer IR needed to use these APIs, refer to the [defaults.yml](#).

### 5.2.1 Power Specification

Simple power specs are specified using the Hammer IR key `vlsi.inputs.supplies`, which is then translated into a Supply object. `hammer_vlsi_impl` exposes the Supply objects to other APIs (e.g. power straps) and can generate the CPF/UPF files depending on which specification the tools support. Multi-mode multi-corner (MMMC) setups are also available by setting `vlsi.inputs.mmmc_corners` and manual power spec definitions are supported by setting the relevant `vlsi.inputs.power_spec...` keys.

### 5.2.2 Timing Constraints

Clock and pin timing constraints are specified using the Hammer IR keys `vlsi.inputs.clocks/output_loads/delays`. These objects can be turned into SDC-style constraints by `hammer_vlsi_impl` for consumption by supported tools.

### 5.2.3 Floorplan & Placement

Placement constraints are specified using the Hammer IR key `vlsi.inputs.placement_constraints`. These constraints are very flexible and have varying inputs based on the type of object the constraint applies to, such as hierarchical modules, hard macros, or obstructions. At minimum, an (x, y) coordinate corresponding to the lower left corner must be given, and additional parameters such as width/height, margins, layers, or orientation are needed depending on the type of constraint. Place-and-route tool plugins will take this information and emit the appropriate commands during floorplanning. Additional work is planned to ensure that floorplans are always legal (i.e. on grid, non-overlapping, etc.).

All Hammer tool instances have access to a method that can produce graphical visualization of the floorplan as an SVG file, viewable in a web browser. To use it, call the `generate_visualization()` method from any custom hook (see *Extending Hammer with Hooks*). The options for the visualization tool are in the Hammer IR key `vlsi.inputs.visualization`.

### 5.2.4 Bumps

Bump constraints are specified using the Hammer IR key `vlsi.inputs.bumps`. Rectangular-gridded bumps are supported, although bumps at fractional coordinates in the grid and deleted bumps are allowed. The place-and-route tool plugin translates the list of bump assignments into the appropriate commands to place them in the floorplan and enable flip-chip routing. The bumps API is also used by the PCB plugin to emit the collateral needed by PCB layout tools such as Altium Designer. This API ensures that the bumps are always in correspondence between the chip and PCB.

The visualization tool mentioned above can also display bump placement and assignments. There are options to view the bumps from the perspective of the ASIC designer or the PCB designer. The views are distinguishable by a reference dot displayed in the left and right corners for the ASIC and PCB perspectives, respectively.

### 5.2.5 Pins

Pin constraints are specified using the Hammer IR key `vlsi.inputs.pin`. PinAssignments objects are generated and passed to the place-and-route tool to place pins along specified block edges on specified metal layers. Preplaced (e.g. hard macros in hierarchical blocks) pins are also supported so that they are not routed. Additional work is planned to use this API in conjunction with the placement constraints API to allow for abutment of hierarchical blocks, which requires pins to be aligned on abutting edges.

## 5.2.6 Power Straps

Power strap constraints are specified using multiple Hammer IR keys in the `par` namespace. The currently supported API supports power strap generation by tracks, which auto-calculates power strap width, spacing, set-to-set distance, and offsets based on basic DRC rules specified in the technology Stackup object. The basic pieces of information needed are the desired track utilization per strap and overall power strap density. Different values can be specified on a layer-by-layer basis by appending `_<layer name>` to the end of the desired option.

## 5.2.7 Special Cells

Special cells are specified in the technology's JSON, but are exposed to provide lists of cells needed for certain steps, such as for fill, well taps, and more. Synthesis and place-and-route tool plugins can grab the appropriate type of special cell for the relevant steps.

## 5.2.8 Submission

Each tool has run submission options given by the Hammer IR key `<tool type>.submit`. Using the `command` and `settings` keys, a setup for LSF or similar workload management platforms can be standardized.

# 5.3 Flow Control

Physical design is necessarily an iterative process, and designers will often require fine control of the flow within any given action. This allows for rapid testing of new changes in the Hammer IR to improve the quality of results (QoR) after certain steps.

Given the flow defined by tool steps and hooks (described in the next section), users can select ranges of these to run. The tool script will then be generated with commands corresponding to only the steps that are to be run.

## 5.3.1 Command-line Interface

Flow control is specified with the following optional Hammer command-line flags, which must always target a valid step/hook:

- `--start_before_step <target>`: this starts the tool from (inclusive) the target step. Alternate flag: `--from_step`
- `--start_after_step <target>`: this starts the tool from (exclusive) the target step. Alternate flag: `--after_step`
- `--stop_after_step`: this stops the tool at (inclusive) the target step. Alternate flag: `--to_step`
- `--stop_before_step`: this stops the tool at (exclusive) the target step. Alternate flag: `--until_step`
- `--only_step`: this only runs the target step

As `hammer-vlsi` is parsing through the steps for a given tool, it will print debugging information indicating which steps are skipped and which are run.

Certain combinations are not allowed:

- `--only_step` is not compatible with any of the other flags
- `--start_before_step` and `start_after_step` may not be specified together
- `--stop_after_step` and `--stop_before_step` may not be specified together

- `--start_before_step` and `--stop_before_step` may not be the same step
- `--start_after_step` and `--stop_after_step` may not be the same step
- `--start_after_step` and `--stop_before_step` may not be the same or adjacent steps

Logically, a target of `stop_{before|after}_step` before `start_{before|after}_step` will also not run anything (though it is not explicitly checked).

## 5.4 Extending Hammer with Hooks

It is unlikely that using the default Hammer APIs alone will produce DRC- and LVS-clean designs with good QoR in advanced technology nodes if the design is sufficiently complex. To solve that, Hammer is extensible using *hooks*. These hooks also afford power users additional flexibility to experiment with CAD tool commands to tweak aspects of their designs. The hook framework is inherently designed to enable reusability, because successful hook methods that solve a technology-specific problem may be upstreamed into the technology plugin for future designs.

### 5.4.1 Hook Methods

Hooks are fundamentally Python methods that extend a given tool's set of available steps and can inject additional TCL commands into the flow. Hook methods need to take in an instance of a particular `HammerTool`, which provides them with the full set of Hammer IR available to the tool. Hooks (depending on how they are included, see below) get turned into step objects that can be targeted with `--from/after/to/until_step` and other hooks.

Hooks can live in a Python file inside the design root so that it is available to the class that needs to extend the default `CLIDriver`. An example of some skeletons of hooks are found in [Chipyard](#). For more comprehensive examples, refer to the hooks unit tests in the `HammerToolHooksTest` class of [test.py](#).

### 5.4.2 Including Hooks

Hooks modify the flow using a few `HammerTool` methods, such as:

- `make_replacement_hook(<target>, <hook_method>)`: this swaps out an existing target step/hook with the hook method
- `make_pre_insertion_hook(<target>, <hook_method>)`: this inserts the hook method before the target step/hook
- `make_post_insertion_hook(<target>, <hook_method>)`: this inserts the hook method after the target step/hook
- `make_removal_hook(<target>)`: this removes the target step/hook from the flow

Note: `<target>` should be a string (name of step/hook), while `<hook_method>` is the hook method itself. All of the hook methods specified this way are targetable by other hooks.

Sometimes, CAD tools do not save certain settings into checkpoint databases. As a result, when for example a `--from_step` is called, the setting will not be applied when the database from which to continue the flow is read in. To get around this, a concept of “persistence” is implemented with the following methods:

- `make_persistent_hook(<hook_method>)`: this inserts the hook method at the beginning of any tool invocation, regardless of which steps/hooks are run
- `make_pre_persistent_hook(<target>, <hook_method>)`: this inserts the hook method at the beginning of any tool invocation, as long as the target step/hook is located at or after the first step to be run

- `make_post_persistent_hook(<target>, <hook_method>)`: this inserts the hook method at the beginning of any tool invocation if the target step/hook is before the first step to be run, or right after the target step/hook if that step/hook is within the steps to be run.

All persistent hooks are NOT targetable by flow control options, as their invocation location may vary. However, persistent hooks ARE targetable by `make_replacement/pre_insertion/post_insertion/removal_hook`. In this case, the hook that replaces or is inserted pre/post the target persistent hook takes on the persistence properties of the target persistence hook.

Some examples of these methods are found in the Chipyard example, linked above.

A list of these hooks must be provided in an implementation of method such as `get_extra_par_hooks` in the command-line driver. This new file becomes the entry point into Hammer, overriding the default `hammer-vlsi` executable.

### 5.4.3 Technology, Tool, and User-Provided Hooks

---

**Note:** Hooks may be provided by the technology plugin, the tool plugin, and/or the user. The order of step & hook priority is as follows, from lowest to highest:

1. Technology default steps
  2. Technology plugin hooks
  3. Tool plugin hooks
  4. User hooks
- 

A technology plugin specifies hooks in its `__init__.py` (as a method inside its subclass of `HammerTechnology`). It should implement a `get_tech_<action>_hooks(self, tool_name: str)` method. The tool name parameter may be checked by the hook implementation because multiple tools may implement the same action. Technology plugin hooks may only target tool default steps to insert/replace.

The included ASAP7 technology plugin provides an example of how to inject two different types of hooks: 1) a persistent hook invoked anytime after the `init_design` step to set top & bottom routing layers, and 2) two post-insertion hooks, one to modify the floorplan for DRCs and the other to scale down a GDS post-place-and-route using the `gdstk` or `gdspy` GDS manipulation utilities. Note that the persistent hook that is included does not necessarily need to be persistent (Innovus does retain this setting in snapshot databases), but it serves as an example for building your own tech plugin.

A tool plugin specifies hooks in its `__init__.py` (as a method inside its subclass of `HammerTool`). It should implement a `get_tool_hooks(self)` method. In contrast to the tech-supplied hooks, the action name and tool name are not specified because a tool instance can only correspond to a single action. Tool plugin hooks may target its own default steps and technology plugin hooks to insert/replace.

A user specifies hooks in the command-line driver and should implement a `get_extra_<action>_hooks(self)` method. User hooks may target tool default steps, technology plugin hooks, and tool plugin hooks to insert/replace. A good example is the `example-vlsi` file in the Chipyard example, which implements a `get_extra_par_hooks(self)` method that returns a list of hook inclusion methods.

The priority means that if both the technology and user specify persistent hooks, any duplicate commands in the user's persistent hook will override those from the technology's persistent hook.

## 5.5 Hammer Buildfile

Hammer natively supports a GNU Make-based build system to manage build dependencies. To use this flow, `vlsi.core.build_system` must be set to `make`. Hammer will generate a Makefile include in the object directory named `hammer.d` after calling the build action:

```
hammer-vlsi -e env.yml -p config.yml --obj_dir build build
```

`build/hammer.d` will contain environment variables needed by Hammer and a target for each major Hammer action (e.g. `par`, `synthesis`, etc. but not `syn-to-par`, which is run automatically when calling `make par`). For a flat design, the dependencies are created between the major Hammer actions. For hierarchical designs, Hammer will use the hierarchy to build a dependency graph and construct the Make target dependencies appropriately.

`hammer.d` should be included in a higher-level Makefile. While `hammer.d` defines all of the variables that it needs, there are often reasons to set these elsewhere. Because `hammer.d` uses `?=` assignment, the settings created in the top-level Makefile will persist. An example of this setup is found in [Chipyard](#).

To enable interactive usage, `hammer.d` also contains a set of companion “redo” targets (e.g. `redo-par` and `redo-syn`). These targets intentionally have no dependency information; they are for advanced users to make changes to the input config and/or edit the design manually, then continue the flow. Additional arguments can be passed to the “redo” targets with the `HAMMER_EXTRA_ARGS` environment variable, for the following example uses:

- “Patching” the configuration for an action using `HAMMER_EXTRA_ARGS="-p patch.yml"`.
- Flow control using `--to_step/--until_step` and `--from_step/after_step` to stop a run at a particular step or resume one from a previous iteration.

## 5.6 Hierarchical Hammer Flow

Hammer supports a bottom-up hierarchical flow. This is beneficial for very large designs to reduce the computing power by partitioning it into submodules, especially if there are many repeated instances of those modules.

### 5.6.1 Hierarchical Hammer Config

The hierarchal flow is controlled in the `vlsi.inputs.hierarchical` namespace. To specify hierarchical mode, you must specify the following keys. In this example, we have our top module set as `ChipTop`, with a submodule `ModuleA` and another submdule `ModuleAA` below that (these are names of Verilog modules).

```
vlsi.inputs.hierarchical:
  mode: hierarchical
  top_module: ChipTop
  config_source: manual
  manual_modules:
    - ChipTop:
      - ModuleA
    - ModuleA:
      - ModuleAA
  constraints:
    - ChipTop:
      - vlsi.core...
      - vlsi.inputs...
    - ModuleA:
```

(continues on next page)

(continued from previous page)

```

- vlsi.core...
- vlsi.inputs...
- ModuleAA:
- vlsi.core...
- vlsi.inputs...

```

Note how the configuration specific to each module in `vlsi.inputs.hierarchical.constraints` are list items, whereas in a flat flow, they would be at the root level.

Placement constraints for each module, however, are not specified here. Instead, they should be specified in `vlsi.inputs.hierarchical.manual_placement_constraints`. The parameters such as `x`, `y`, `width`, `height`, etc. are omitted from each constraint for clarity. In the bottom-up hierarchical flow, instances of submodules are of type: `hardmacro` because they are hardened from below.

```

vlsi.inputs.hierarchical:
  manual_placement_constraints_meta: append
  manual_placement_constraints:
    - ChipTop:
      - path: "ChipTop"
        type: toplevel
      - path: "ChipTop/path/to/instance/of/ModuleA"
        type: hardmacro
    - ModuleA:
      - path: "ModuleA"
        type: toplevel
      - path: "ModuleA/path/to/instance/of/ModuleAA"
        type: hardmacro
    - ModuleAA:
      - path: "moduleAA"
        type: toplevel

```

## 5.6.2 Flow Management and Actions

Based on the structure in `vlsi.inputs.hierarchical.manual_modules`, Hammer constructs a hierarchical flow graph of dependencies. In this particular example, synthesis and place-and-route of `ModuleAA` will happen first. Synthesis of `ModuleA` will then depend on the place-and-route output of `ModuleAA`, and so forth. These are enumerated in the auto-generated Makefile.

To perform a flow action (`syn`, `par`, etc.) for a module using the auto-generated Makefile, the name of the action is appended with the module name. For example, the generated Make target for synthesis of `ModuleA` is `syn-ModuleA`.

## 5.6.3 Cadence Implementation

Currently, the hierarchical flow is implemented with the Cadence plugin using its Interface Logic Model (ILM) methodology. At the end of each submodule's place-and-route, an ILM is written as the hardened macro, which contains an abstracted view of its design and timing models only up to the first sequential element.

ILMs are similar to LEFs and LIBs for traditional hard macros, except that the interface logic is included in all views. This means that at higher levels of hierarchy, the ILM instances can be flattened at certain steps, such as those that perform timing analysis on the entire design, resulting in a more accurate picture of timing than a normal LIB would afford.

### 5.6.4 Tips for Constraining Hierarchical Modules

In a bottom-up hierarchical flow, it is important to remember that submodules do not know the environment in which they will be placed. This means:

- At minimum, the pins must be placed on the correct edges of the submodule on metal layers that are accessible in the parent level. Depending on the technology, this may interfere with things like power straps near the edge, so custom obstructions may be necessary. If fixed IOs are placed in submodules (e.g. bumps), then in the parent level, those pins must be promoted up using the `preplaced: true` option in the pin assignment.
- Clocks should be constrained more tightly for a submodule compared to its parent to account for extra clock insertion delay, jitter, and skew at increasingly higher levels of hierarchy. Otherwise, you may run into surprise timing violations in submodule instances even if those passed timing in isolation.
- You may need to specify pin delays `vlsi.inputs.delays` for many pins to optimize the partitioning of sequential signals that cross the submodule boundary. By default, without pin delay constraints, the input and output delay are constrained to be coincident with its related clock arrival at the module boundary.
- Custom SDC constraints that originate from a higher level (e.g. false paths from async inputs) need to be specified in submodules as well.
- Custom CPFs will need to be written if differently-named power nets need to be globally connected between submodules. Similarly, hierarchical flow with custom CPFs can also be used to fake a multi-power domain topology until Hammer properly supports this feature.

### 5.6.5 Special Notes & Limitations

1. Hammer IR keys propagate up through the hierarchical tree. For example, if `vlsi.inputs.clocks` was specified in the constraints for `ModuleAA` but not for `ModuleA`, `ModuleA` will inherit `ModuleAA`'s constraints. Take special care of where your constraints come from, especially for a parent module with more than one submodule. To avoid confusion, it is recommended to specify the same set of keys for every module.
2. Hammer IR keys specified at the root level (i.e. outside of `vlsi.inputs.hierarchical.constraints`) do not override the corresponding submodule constraints. However, if you add a Hammer IR file using `-p` on the command line (after the file containing `vlsi.inputs.hierarchical.constraints`), those keys are global and override submodule constraints unless a meta action is specified. To avoid confusion, it is recommended to specify all constraints with `vlsi.inputs.hierarchical.constraints`.
3. Due to the structure of `vlsi.inputs.hierarchical.constraints` as a list structure, currently, there are the following limitations:
  - You must include all of the constraints in a single file. The config parser is unable to combine constraints from different files because most meta actions do not work on list items (advanced users will need to use `deepsubst`). This will make it harder for collaboration, and unfortunately, changes to module constraints at a higher level of hierarchy after submodules are hardened will trigger the Make dependencies, so you will need to modify the generated Makefile or use `redo-targets`.
  - Other issues have been observed, such as the bump API failing (see [this issue](#) at the top module level. This is caused by similar mechanisms as above. The workaround is to ensure that bumps are specified at the root level for only the top module and the bumps step is removed from submodule par actions.
4. Most Hammer APIs are not yet intelligent enough to constrain across hierarchical boundaries. For example:
  - The power straps API is unable to pitch match power straps based on legalized placement of submodule instances or vice versa.
  - The pin placement API does not match the placement of pins that may face each other in two adjacent submodule instances. You will need to either manually place the pins yourself or ensure a sufficient routing channel between the instances at the parent level.



5. Hammer does not support running separate decks for submodule DRC and LVS. Technology plugins may need to be written with Makefiles and/or technology-specific options that will implement different checks for submodules vs. the top level.



## HAMMER EXAMPLES

The following are prebuilt example designs, toolchain and/or flows you can use

### 6.1 End-to-End Integration Tests

In the `e2e/` folder, there is an example of a bare Hammer project setup with tool plugin submodules, Hammer IR snippets in the `configs/`, `env/`, and `pdks/` subfolders, and Verilog files in the `src/` subfolder.

`pyproject.toml` is an example of how to install Hammer as a source dependency, and for Berkeley affiliates, the `envs` are particularly useful for setting up your tool paths and license servers.

Refer to `e2e/README.md` for details of how to run the tests.

### 6.2 OpenROAD and Sky130

The [OpenROAD-flow](#) is an open-source, technology-independent VLSI toolchain. As of this writing, it has the ability to run synthesis, place-and-route, and drc through various open-source tools.

Hammer has the ability to target the OpenROAD toolchain and the Skywater 130nm PDK. So you can now simply push a button in [Chipyard](#) to go from your Chisel design to a somewhat reasonable gds at an example 130-nm node.

#### 6.2.1 Instructions

Follow [these directions in the Chipyard docs](#) to build your own design with OpenROAD and Sky130.



## INDICES AND TABLES

- `genindex`
- `modindex`
- `search`