

PySpike: Agile Co-Verification of RISC-V SoC's and Beyond

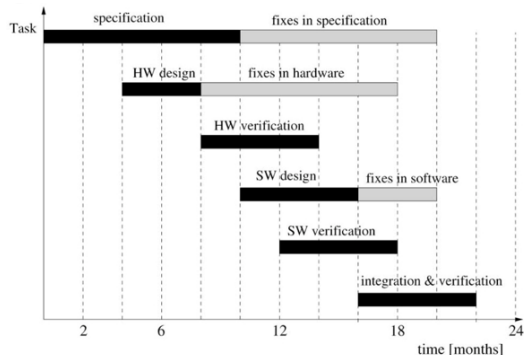
LIU Yu

`<liuyu@huimtlab.org>`




Open Source @ Siemens 2025

Sequential engineering



Collaborative engineering




... the executable specification may be used as a golden *reference model* for the software development process Indeed, such a specification and the use of cosimulation tools may also help to reveal errors in the test and firmware development cycle very early. [1]

- 1 Motivation
- 2 Notable Features
- 3 Real-World Examples
- 4 Our Open-Source Playbook
- 5 Remarks and Outlook

Motivation


De facto standard RISC-V ISA simulator [2]

- functional model of RISC-V processors (harts) with state-of-the-art ISA support;
- originated from UC Berkeley, alongside RISC-V, actively maintained for 15+ years;
<https://github.com/riscv-software-src/riscv-isa-sim>
- **reference model** for hardware/software co-verification, e.g. *differential testing* [3]:



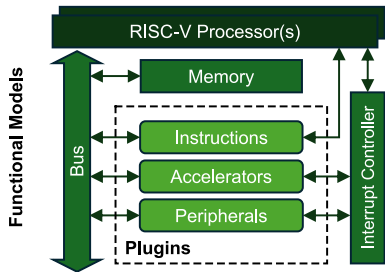
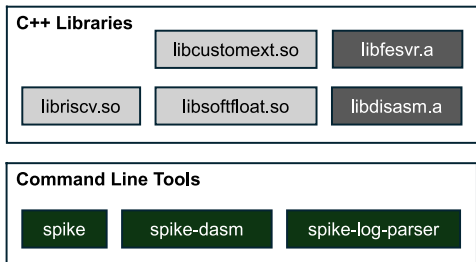
De facto standard RISC-V ISA simulator [2]

- functional model of RISC-V processors (harts) with state-of-the-art ISA support;
- originated from UC Berkeley, alongside RISC-V, actively maintained for 15+ years;
<https://github.com/riscv-software-src/riscv-isa-sim>
- **reference model** for hardware/software co-verification, e.g. *differential testing* [3]:




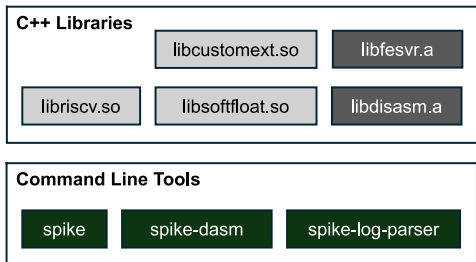
Spike [2], the upstream project at a glance

- C++ code base (40k+ lines), 190+ contributors, 2.9k stars on GitHub (2025/10);
- CLI tools (**spike**, **xspike**, **spike-dasm**, ...) and C++ libs (**libriscv.so**, ...);
- plugin system based on dynamic loading (**dlopen**) of shared objects / libraries;



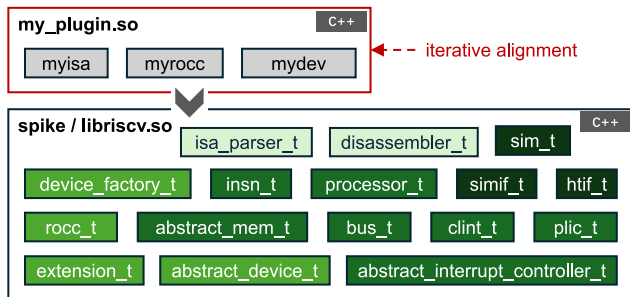
Spike [2], the upstream project at a glance

- C++ code base (40k+ lines), 190+ contributors, 2.9k stars on GitHub (2025/10);
- CLI tools (**spike**, **xspike**, **spike-dasm**, ...) and C++ libs (**libriscv.so**, ...);
- plugin system based on dynamic loading (**dlopen**) of shared objects / libraries;



Spike [2], the upstream project at a glance

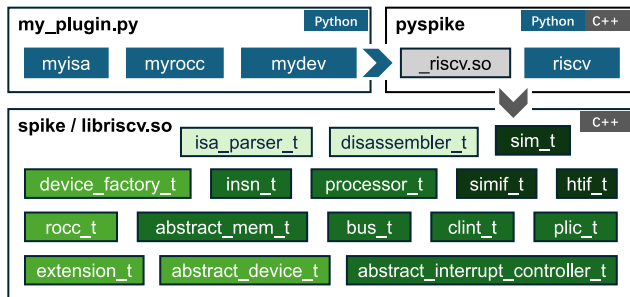
- C++ code base (40k+ lines), 190+ contributors, 2.9k stars on GitHub (2025/10);
- CLI tools (**spike**, **xspike**, **spike-dasm**, ...) and C++ libs (**libriscv.so**, ...);
- plugin system based on dynamic loading (**dlopen**) of shared objects / libraries;



```
$ spike \  
  --isa=rv64gc_xmyisa \  
  --priv=msu \  
  --extlib=my_plugin.so \  
  --extension=myrocc \  
  --device=mydev,0x20000000 \  
  program.elf
```

Agile co-verification calls for the feasibility of ...

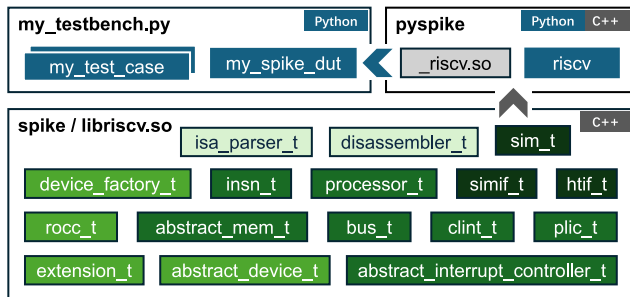
- fast-prototyping custom instructions, accelerators, peripherals, ...;
- programmatically manipulating Spike for interoperation with testbenches;
- reusing small goodies like ISA-string parser, RISC-V disassembler, ...;



```
$ pyspike \  
  --isa=rv64gc_xmyisa \  
  --priv=msu \  
  --extlib=my_plugin.py \  
  --extension=myrocc \  
  --device=mydev,0x20000000 \  
  program.elf
```

Agile co-verification calls for the feasibility of ...

- fast-prototyping custom instructions, accelerators, peripherals, ...;
- programmatically manipulating Spike for interoperation with testbenches;
- reusing small goodies like ISA-string parser, RISC-V disassembler, ...;



```
$ python -q
>>> from riscv.sim import *
>>> help(sim_t)
Help on class sim_t in module
↳ riscv._riscv.sim:

class sim_t(riscv._riscv.htif.htif_t,
↳ riscv._riscv.simif.simif_t)
...
```

Notable Features

Install PyPI package: **spike**

1 setup virtual environment (optional):

```
$ python3 -m venv venv  
$ source venv/bin/activate
```

2 install **spike** package with **pip**:

```
(venv) $ pip install --pre spike  
Collecting spike  
  Downloading spike-0.0.5.dev...
```

3 check availability of **pyspike**:

```
(venv) $ pyspike --help  
Spike RISC-V ISA Simulator 1.1.1-dev  
...  
usage: spike [host options] <target  
  ↪ program> [target options]  
...
```

CLI Tools

Python

spike

pyspike

Source Code

Python

examples

tests

Facade Package (riscv)

Python

__init__

__main__

dev

isa

_utils

Core Module (_riscv.so)

C++

py_bridge

cfg

csrs

decode

devices

disasm

extension

fesvr

processor

sim

Bundled Spike Distribution

C++

spike

libriscv.so

libdisasm.a


libfesvr.a

C++ headers (fdt / fesvr / riscv / softfloat)

(Supports Linux / macOS; requires **Python 3.8+**)

Feature #1: Python in Spike (PIS)

- write instruction / peripheral models in Python, and plug them into Spike;




```
$ pyspike --isa=rv32gc_xmyisa
  ↳ --extlib=myisa.py program.elf
```

```
1 # @file myisa.py
2 from typing import List
3 from riscv import isa
4 from riscv.csr import *
5 from riscv.disasm import *
6 from riscv.processor import *
7
8 @isa.register("myisa")
9 class MyISA(isa.ISA):
10     def __init__(self): ...
11     def get_instructions(self, proc)
12         ↳ -> List[insn_desc_t]: ...
13     def get_disasms(self, proc) ->
14         ↳ List[disasm_insn_t]: ...
15     def get_csrs(self, proc) ->
16         ↳ List[csr_t]: ...
17     def reset(self, proc) -> None: ...
```

Feature #1: Python in Spike (PIS)

- write instruction / peripheral models in Python, and plug them into Spike;




```
$ pyspike --isa=rv32gc --extlib=mydev.py
↳ --device=mydev,0x20000000
↳ program.elf
```

```
1 # @file mydev.py
2 from typing import Optional
3 from riscv import dev
4 from riscv.sim import sim_t
5
6 @dev.register("mydev")
7 class MyDEV(dev.MMIO):
8     def __init__(self, sim: sim_t,
9         ↳ args: Optional[str]): ...
9     def load(self, addr: int, size:
10        ↳ int) -> bytes: ...
10    def store(self, addr: int, data:
11        ↳ bytes) -> None: ...
11    def size(self) -> int: ...
12    def tick(self, rtc_ticks: int) ->
13        ↳ None:
```

Feature #2: Spike in Python (SIP)

- instantiate and seamlessly integrate Spike DUT into Python testbenches;




```
$ PYSPIKE_LIBS=mydev.py \
  pytest -v testbench.py
```

```
1 # @file conftest.py
2 import pytest
3 from riscv.cfg import *
4 from riscv.debug_module import *
5 from riscv.sim import sim_t
6
7 @pytest.fixture
8 def mydut():
9     yield sim_t(
10         cfg=cfg_t(isa="rv32gc", priv="m",
11                 ↪ mem_layout=[mem_cfg_t(
12                 ↪ 0x90000000, 0x40000)]),
13         halted=False,
14         plugin_device_factories=[("mydev",
15                                   ↪ ("0x20000000", ))],
16         args=["program.elf"],
17         dm_config=debug_module_config_t())
```


Feature #2: Spike in Python (SIP)

- reuse Spike's internal gadgets in Python, i.e., RISC-V disassembler;

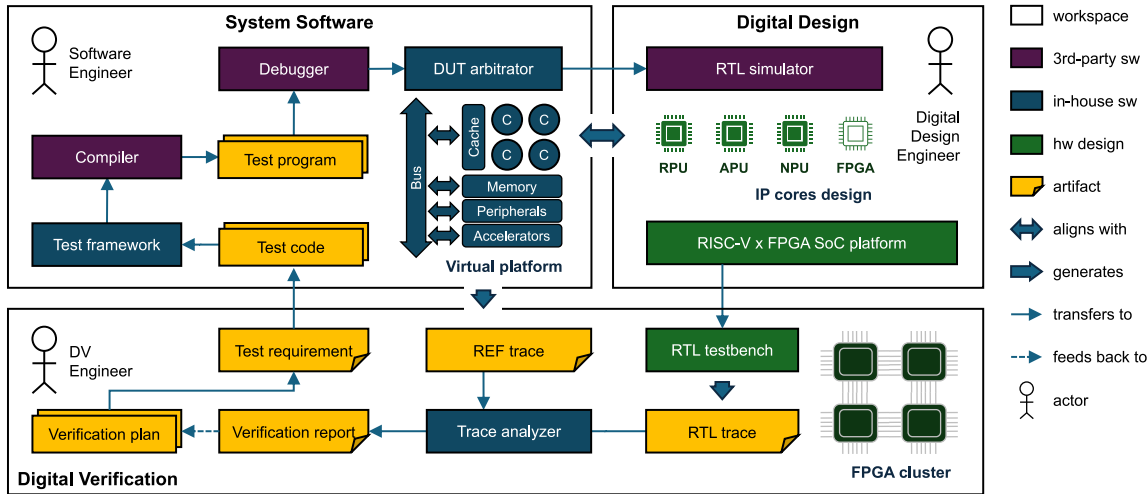


```
$ PYSPIKE_LIBS=myisa.py \
python mydasm.py program.elf
```

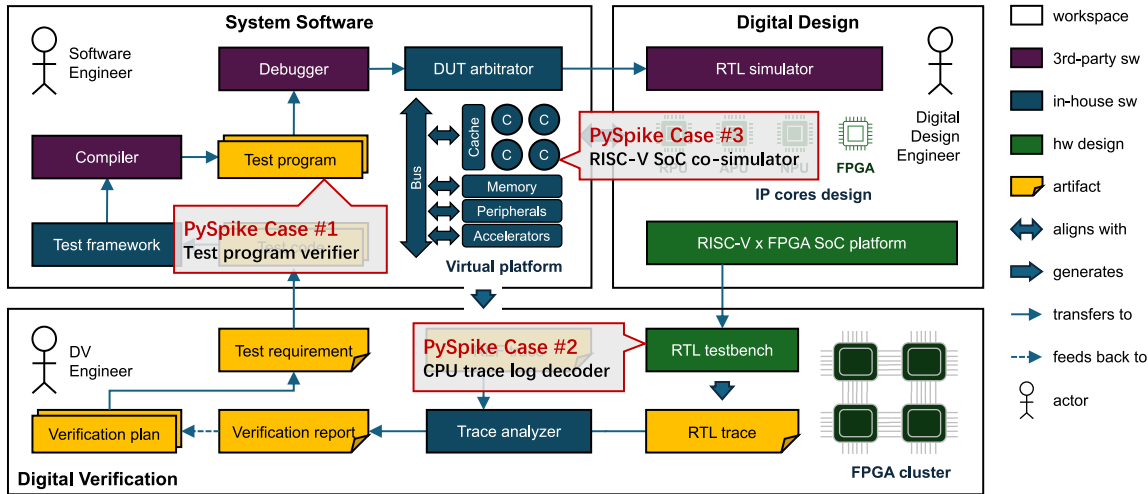
```
1  #!/usr/bin/env python3
2  # @file mydasm.py
3  from sys import argv
4  from elftools.elf.elffile import *
5  from riscv.decode import *
6  from riscv.disasm import *
7  from riscv.isa_parser import *
8
9  d = disassembler_t(isa_parser_t(
10     ↪ "rv64gc_xmyisa", "msu"))
11 e = ELFFile.load_from_path(argv[1])
12 s = e.get_section_by_name(".text")
13 pc = s.header.sh_addr
14 for op in insn_fetch_all(s.data()):
15     print(f"0x{pc:08x} (0x{op:08x}): ",
16         ↪ d.disassemble(x:=insn_t(op)))
17     pc += len(x)
```

Real-World Examples

PySpike in the big picture of co-verification

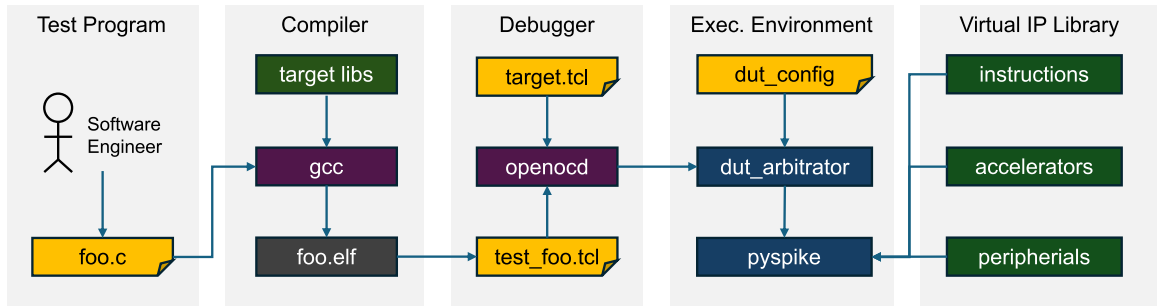


PySpike in the big picture of co-verification



Case #1: Continuous test program verifier

- PySpike provides a reference DUT for verifying test-programs.

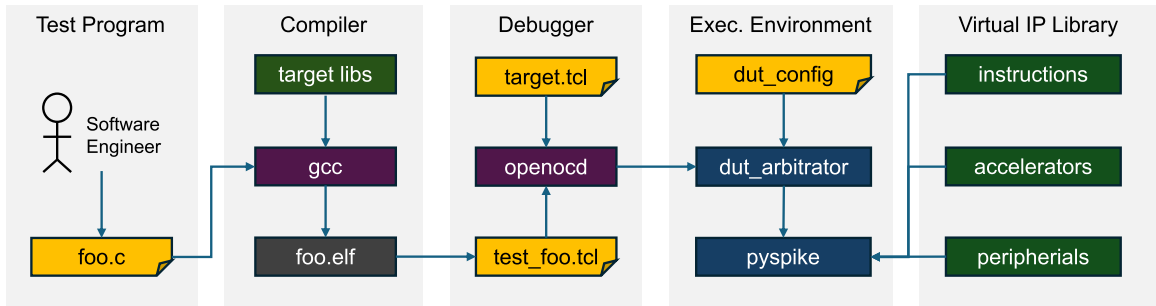


```
$ dut_arbitrator.py target-pyspike
Listening for remote bitbang
↪ connection on port 12345.
```

```
$ openocd \
  --file target.tcl \
  --file test_foo.tcl
...
```

Case #1: Continuous test program verifier

- PySpike provides a reference DUT for verifying test-programs.

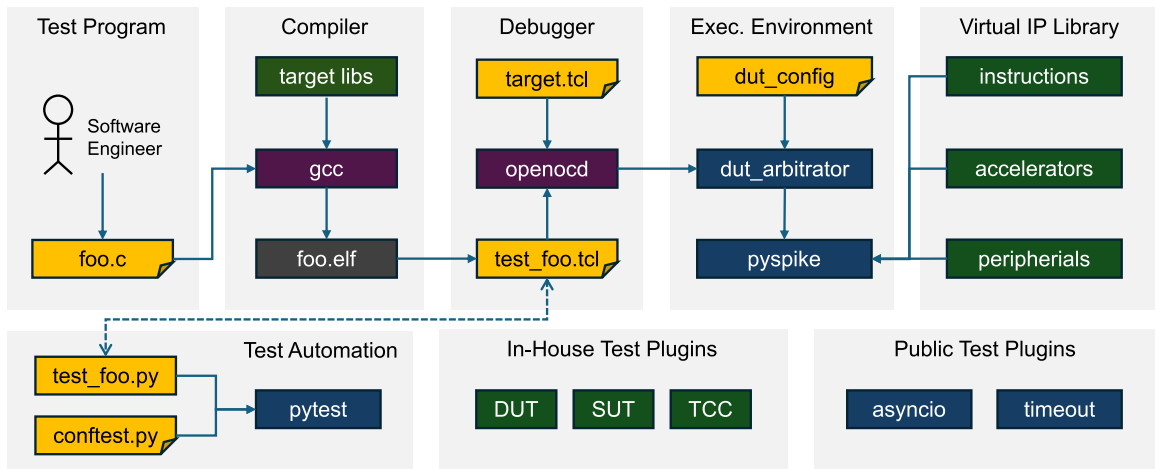


```
$ dut_arbitrator.py target-pyspike
Listening for remote bitbang
↪ connection on port 12345.
CPU1> Hello World!
```

```
$ openocd \
--file target.tcl \
--file test_foo.tcl
...
```

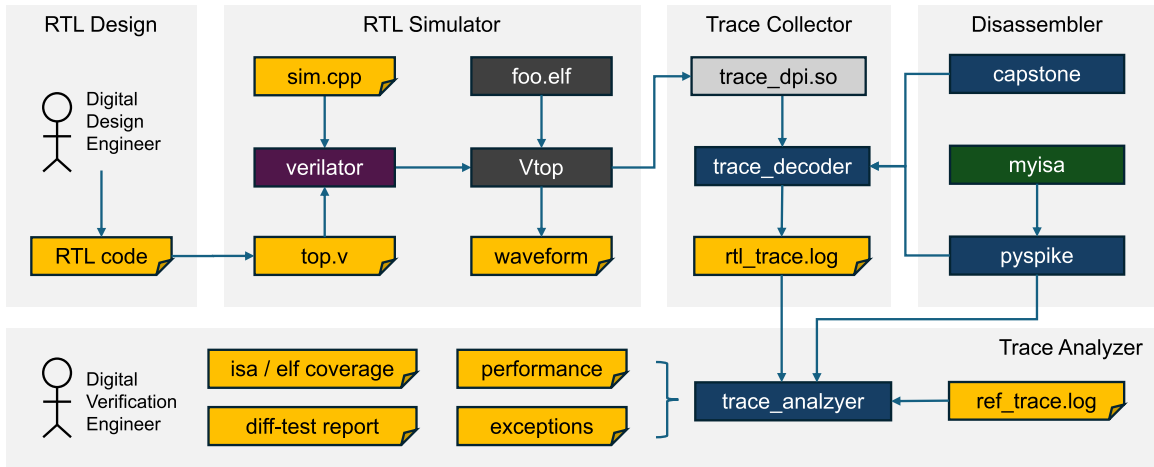
Case #1: Continuous test program verifier

- PySpike provides a reference DUT for verifying test-programs.



Case #2: CPU trace log decoder and beyond

- PySpike provides extensible RISC-V disassembler with top-norch ISA support.



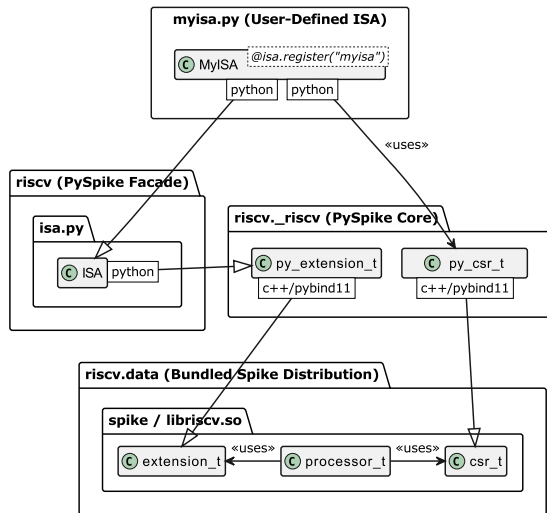
Our Open-Source Playbook

Concept

- design exclusively through Spike's **dlopen**-based plugin system, avoid intrusive patches to Spike's code;
- bundle Spike by compiling directly from upstream source at pinned commit;

Rationale

- zero-effort adoption for existing Spike users, w/ or w/o custom plugins;
- minimal maintenance overhead when syncing with upstream changes;



Concept

- design exclusively through Spike's **dlopen**-based plugin system, avoid intrusive patches to Spike's code;
- bundle Spike by compiling directly from upstream source at pinned commit;

Rationale

- zero-effort adoption for existing Spike users, w/ or w/o custom plugins;
- minimal maintenance overhead when syncing with upstream changes;

```
$ pyspike \  
--isa=rv64gc_xmyisa \  
--priv=msu \  
--extlib=myisa.py \  
--device=mydev,0x20000000 \  
program.elf
```



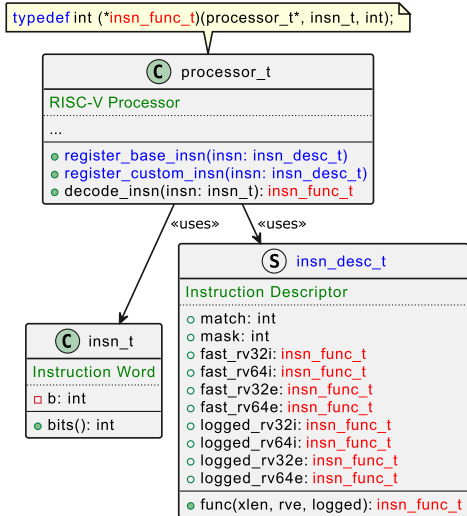
```
$ PYSPIKE_LIBS=myisa.py \  
spike \  
--isa=rv64gc_xmyisa \  
--priv=msu \  
--extlib=libpython3.8.so \  
--extlib=_riscv.so \  
--device=mydev,0x20000000 \  
program.elf
```

Concept

- design exclusively through Spike's **dlopen**-based plugin system, avoid intrusive patches to Spike's code;
- bundle Spike by compiling directly from upstream source at pinned commit;

Rationale

- zero-effort adoption for existing Spike users, w/ or w/o custom plugins;
- minimal maintenance overhead when syncing with upstream changes;



Concept

- proactively engage with upstream: identify hurdles early and resolve them through discussions or pull requests;
- strictly avoid maintaining divergent forks of Spike or PySpike, even internally;

Rationale

- prevent *technical debt* accumulation from conflicts and parallel development;
- strengthen the upstream foundation, which benefits the whole community;

Example (the '--device' args fix)

2023/12: Someone proposed to *fix* Spike's CLI option '**--device**', i.e.,

*Fix Spike **--device** option to pass on args to downstream plugins #1522*

```
$ spike --extlib=my_plugin.so \  
- --device=<name> \  
+ --device=<name>,<args> \  
  program.elf
```

The PR broke our *then-in-house* prototype PySpike for 1) the 'fixed' option assumed ','-separated values, and 2) per-device **args** for the same **name** was impossible.

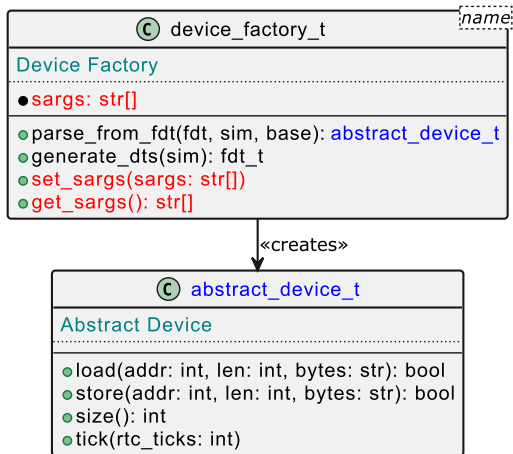
Concept

- proactively engage with upstream: identify hurdles early and resolve them through discussions or pull requests;
- strictly avoid maintaining divergent forks of Spike or PySpike, even internally;

Rationale

- prevent *technical debt* accumulation from conflicts and parallel development;
- strengthen the upstream foundation, which benefits the whole community;

Example (the '--device' args fix)



Concept

- proactively engage with upstream: identify hurdles early and resolve them through discussions or pull requests;
- strictly avoid maintaining divergent forks of Spike or PySpike, even internally;

Rationale

- prevent *technical debt* accumulation from conflicts and parallel development;
- strengthen the upstream foundation, which benefits the whole community;

Example (the '--device' args fix)

2024/04: We raised our concern to the maintainers, received positive feedback, and got our PR merged in ~1 week.

Counter-intuitive MMIO device arguments behavior #1652

Support per-device arguments and device factory reuse #1655

© device_factory_t	name
Device Factory	
• sargs:-str[]	
• parse_from_fdt(fdt, sim, base, sargs): abstract_device_t	
• generate_dts(sim, sargs): fdt_t	
• set_sargs(sargs:-str[])	
• get_sargs(): str[]	

Concept

- share PySpike's development, testing, and packaging workflows openly;
- leverage CI/CD services to automate and enforce quality adherence;

Rationale

- foster transparency and reproducibility for external contributors and users;
- open-source software thrive on community trust, even more so for quality-assuring tools like PySpike;

Example (developing workflow)

1 get source code:

```
$ git clone --recurse-submodules \
  git@github.com:huimtlab/pyspike
$ cd pyspike
```

2 setup virtual environment:

```
$ python3 -m venv venv
$ source venv/bin/activate
```

3 install in *editable* mode:

```
(venv) $ pip install -e '.[dev]'
```

4 lint & test & coverage:

```
(venv) $ pytest -v
...
76 passed, 8 skipped in 9.42s
```


Concept


- share PySpike's development, testing, and packaging workflows openly;
- leverage CI/CD services to automate and enforce quality adherence;

Rationale

- foster transparency and reproducibility for external contributors and users;
- open-source software thrive on community trust, even more so for quality-assuring tools like PySpike;

Example (packaging workflow)

PySpike bundles Spike binaries as Python package data. We migrated internal build & package workflow to GitHub Actions.



Agile chip design methodology is still of limited usage due to lack of tool-chain and developing framework [4]

Remarks and Outlook

Closing remarks and key takeaways

- PySpike opens up Spike's C++ internals for interoperability with Python scripts; we showcase PySpike's potential in boosting agile co-verification with a few examples;
- The development of PySpike is guided by our *open-source playbook*, promising non-intrusive design, upstream engagement, and open engineering process;

Future outlook and call for collaboration

- Distributing embedded toolchains via the Python Package Index (PyPI) grants wider accessibility, and brings in extra benefits like dependency management, environment isolation, multi-version coexistence, . . . ; we hope to explore this area in the future;
- As a language binding, PySpike needs the community's help to keep up with Spike's changes to its evolving feature set and *unstable* C++ API's;

EsionTech (中微亿芯)

- subsidiary of CETC's Research Institute 58, founded in 2013;
- headquarter in Wuxi, R&D centers in Beijing, Shanghai, Wuhan, ...;
- vendor of all-programmable and heterogeneous computing chips;
- new player in the RISC-V ecosystem, since early 2023;



HuiMt Labs (惠山实验室)

- affiliated with EsionTech's Beijing R&D Center;
- small group of open-source software enthusiasts;
- veterans in HW / SW co-design and co-verification;
- contributors to RISC-V tools, i.e. **gcc**, **spike**, **riscv-dv**, ...;



Thank You!

- [1] Jürgen Teich. “Hardware/Software Codesign: The Past, the Present, and Predicting the Future”. In: *Proceedings of the IEEE 100.Special Centennial Issue* (2012), pp. 1411–1430. DOI: 10.1109/JPROC.2011.2182009.
- [2] Andrew Waterman et al. *Spike RISC-V ISA Simulator*. <https://github.com/riscv-software-src/riscv-isa-sim>. 2010-2025. (Visited on 10/29/2025).
- [3] William M. McKeeman. “Differential Testing for Software”. In: *Digital Technical Journal* 10.1 (1998).
- [4] Yinan Xu et al. “Towards Developing High Performance RISC-V Processors Using Agile Methodology”. In: *IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2022, pp. 1178–1199. DOI: 10.1109/MICRO56248.2022.00080.