

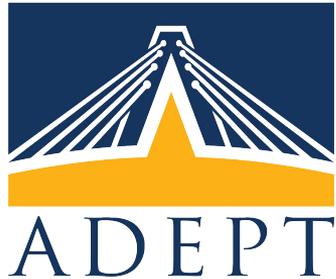
The BOOM Project

Christopher Celio, Jerry Zhao, Abraham Gonzalez, Ben Korpan,
Krste Asanovic, David Patterson



Berkeley
Architecture
Research

Berkeley
UNIVERSITY OF CALIFORNIA

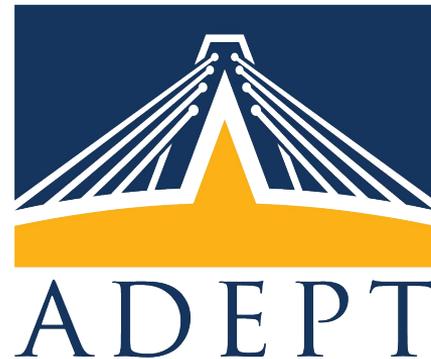


Outline

- **Current state of the project**
- How to get started
- How to contribute
- Things to work on
- Research ideas
- Conclusion and discussion

Project Maintainers

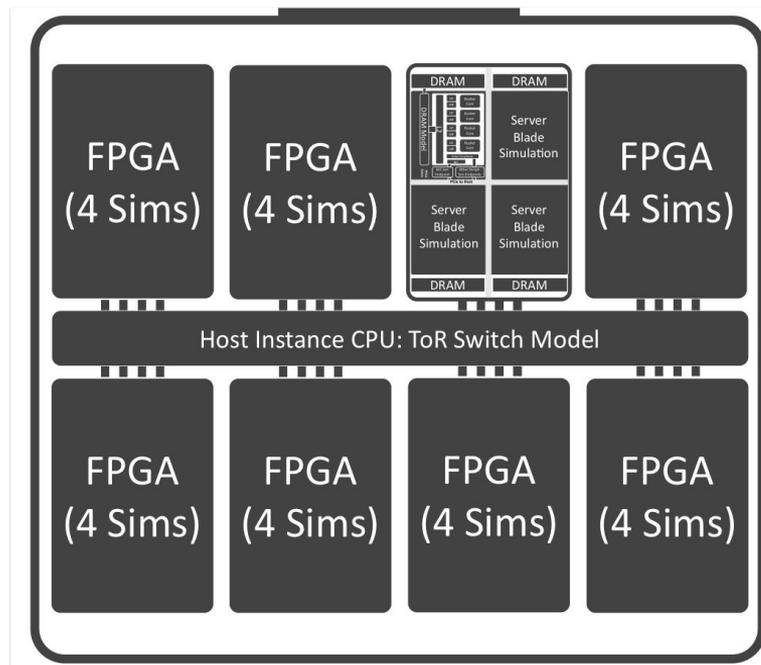
- Chris Celio is now at Esperanto Technologies
 - Works on the ET-Maxion out-of-order core
- Work at Berkeley
 - Jerry Zhao - Senior undergraduate
 - Maintenance + Keystone support + vectors
 - Abe Gonzalez, Ben Korpan - First-year graduate students
 - Ring Scalar design of BOOM
 - Keystone and security additions
 - General maintenance + Documentation updates



**Berkeley
Architecture
Research**

Recent/Ongoing Work

- Taped out as BROOM
- Integrated with FireSim - a open-source FPGA-accelerated simulation tool that runs on Amazon EC2 F1
 - Boot Linux on a multi-core BOOM with 16 GB DDR3, UART, Ethernet NIC in the cloud for 50 cents/hour at ~100 MHz
- Work towards BOOMv3 release
 - Memory model
 - RVC
 - Performance debugging
 - And more!



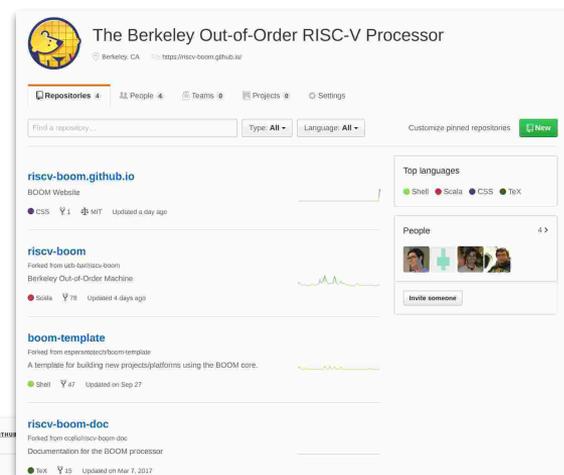
FireSim 32 Node Rack (128 Cores)

Outline

- Current state of the project
- **How to get started**
- How to contribute
- Things to work on
- Research ideas
- Conclusion and discussion

Getting Started

- Central location of information is the BOOM website and GitHub organization
 - Website: riscv-boom.github.io
 - GitHub: github.com/riscv-boom
- Website contents
 - Link to GitHub organization
 - BOOM v2 Specification
 - 70 pages of documentation on microarchitecture, future work, etc
 - Publications/Talks/Reports
 - Current team members



BOOM GitHub Organization

- riscv-boom
 - The core source code
 - Note: That elements of RocketChip are reused
- boom-template
 - Template for building new projects with the BOOM core
 - Includes all necessary materials to build/run the core
- riscv-boom.github.io
 - Project website
- riscv-boom-doc
 - Documentation for the core, project tools, etc

The screenshot shows the GitHub organization page for 'The Berkeley Out-of-Order RISC-V Processor'. The page lists several repositories:

- riscv-boom.github.io**: BOOM Website (CSS, MIT, updated a day ago)
- riscv-boom**: Forked from sbt-barracuda-boom (Scala, updated a day ago)
- boom-template**: Forked from esperesoch/boom-template (Shell, updated on Sep 27)
- riscv-boom-doc**: Forked from rafalriscv-boom-doc (Text, updated on Mar 2, 2017)

Below the repository list, a snippet of the README for 'The Berkeley Out-of-Order RISC-V Processor' is shown. It describes BOOM as a synthesizable core for the RV64G RISC-V superscalar Berkeley Out-of-Order Machine. A table lists features supported by BOOM:

| Feature | BOOM |
|--------------------------------|----------------|
| ISA | RISC-V (RV64G) |
| Synthesizable | ✓ |
| FPGA | ✓ |
| Parameterized | ✓ |
| Floating Point (IEEE 754-2008) | ✓ |
| Atomic Memory Op Support | ✓ |

BOOM v2 Specification

- Documentation on
 - Core units/stages
 - Most detailed implementation resource outside of code
 - Repository layout and reused resources from RocketChip
 - Verification
 - Tools and tests
 - Using gem5 pipeline viewer
 - Torture tests and RISC-V tests
 - Physical realization information
 - Future work
 - Includes preliminary design ideas for each piece
- Will look to change to ReadTheDocs format

The Berkeley Out-of-Order Machine (BOOM) Design Specification

Christopher Colla, David Patterson, and Kostas Amanatidis
University of California, Berkeley, California 94720-1770
colla@eecs.berkeley.edu

December 4, 2016

This draft is a work-in-progress.

The information in this publication is subject to change without notice. The document is available on <https://riscv.org>.

Copyright © 2016 Christopher Colla

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

Contents

| | |
|--|----|
| 1 Introduction & Overview | 5 |
| 1.1 The BOOM Pipeline | 5 |
| 1.2 The RISC-V ISA | 7 |
| 1.3 The Core Hardware Constructing Language | 7 |
| 1.4 Quickstart | 8 |
| 1.5 The BOOM Repository | 8 |
| 1.6 The RocketChip Repository Layout | 8 |
| 1.6.1 The Rocket Core - a Library of Reusable Components | 10 |
| 2 Instruction Fetch | 11 |
| 2.1 The Rocket L-Cache | 11 |
| 2.2 The Fetch Buffer | 12 |
| 3 Branch Prediction | 13 |
| 3.1 The Rocket Next-Line Predictor (NLP) | 13 |
| 3.1.1 NLP Positioning | 14 |
| 3.1.2 NLP Updates | 14 |
| 3.2 The Branch Predictor (BPT) | 15 |
| 3.2.1 Making Predictions | 15 |
| 3.2.2 Updating the Branch Predictor | 16 |
| 3.2.3 Managing the Global History Register | 18 |
| 3.2.4 The Branch Branch Buffer (BBB) | 19 |
| 3.2.5 Branch Snapshot State | 20 |
| 3.2.6 The Branch Branch Predictor Case | 20 |
| 3.2.7 The Two-bit Counter Tables | 21 |
| 3.2.8 The Global Predictor | 22 |
| 3.2.9 The L2FP Predictor | 22 |
| 3.2.10 Other Predictors | 25 |
| 3.3 Branch Prediction Configuration | 26 |
| 3.3.1 Custom Configuration Options | 26 |
| 3.3.2 TAGE Configuration | 26 |
| 4 The Decode Stage | 27 |

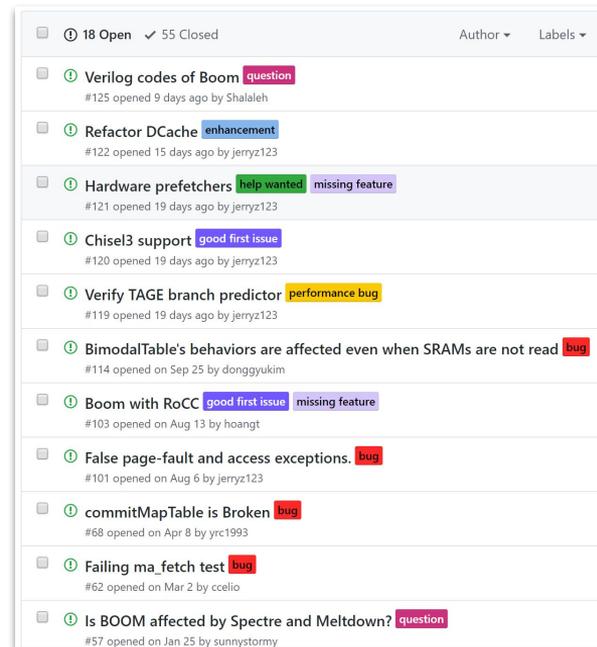
2

Outline

- Current state of the project
- How to get started
- **How to contribute**
- Things to work on
- Research ideas
- Conclusion and discussion

Contribute to the project

- Bug reports
 - Verify against Spike ISA simulator
 - Add issue with reproducible code snippets
- Feature Additions
 - Reach out to the Berkeley team to discuss major changes
 - Keep BOOM consistent
 - Avoid conflicting work
 - Bounce around ideas, get help, etc
 - Submit a pull request (PR)
 - Detailed description of changes
 - Don't forget to update documentation alongside code

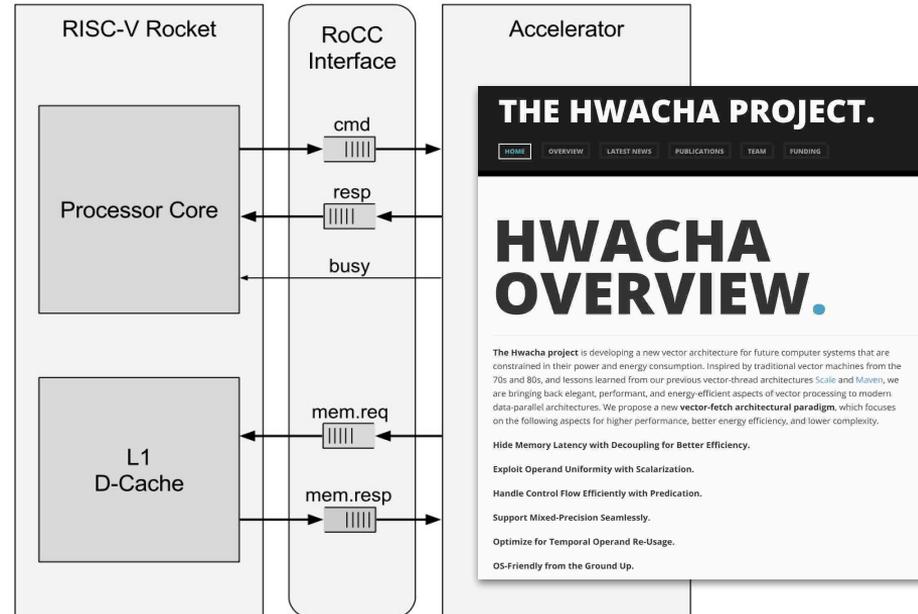


Outline

- Current state of the project
- How to get started
- How to contribute
- **Things to work on**
- Research ideas
- Conclusion and discussion

RoCC Support

- Enable support for co-processors and accelerators
 - Hwacha vector accelerator
 - Crypto units, DMA engines
 - + Your custom block
- Not trivial to implement
 - Adapt RoCC to an OoO machine
 - Where do instruction bits live?
 - One-at-a-time, stall until RoCC instruction is ready to commit
 - BoCC?



Memory System Changes

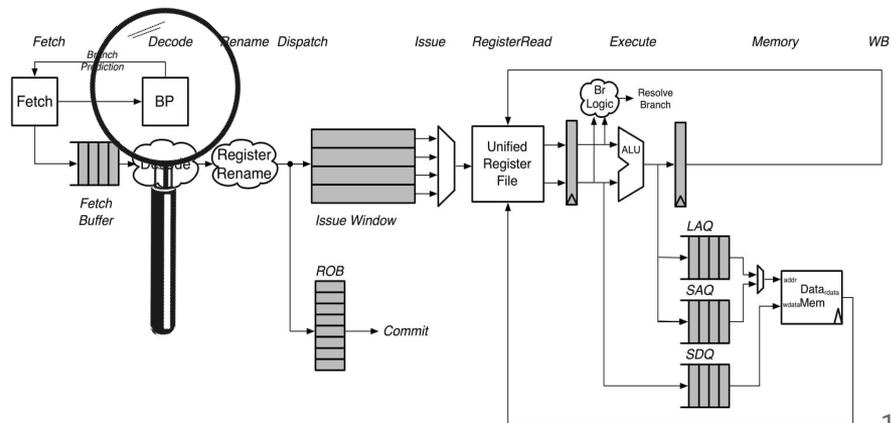
- Total Store Ordering and Weak Memory Ordering
 - Currently implements the relaxed consistency model
 - Update to support RVWMO, order loads to the same address
- Prefetchers
- Memory Disambiguation Predictor
 - Predict dependencies between different memory accesses
- Multi-ported data-cache
 - Commit multiple loads per cycle

RISC-V Compressed Instructions

- Enables smaller 16b encodings of common instructions
 - All 16b instructions map to a 32b longer instruction
 - 32b instructions can now be aligned to the half-word boundary
- Challenges
 - A few 16b instructions exhibit different behaviour than 32b instructions
 - Ex. JAL may return PC+2 or PC+4 depending on which version of the instruction it was
 - Must track if it was a 16b/32b instruction through the pipeline
 - Require code changes in several structures
 - Reorder Buffer (ROB)
 - Branch Predictor
 - More!

Verification

- Tests
 - Unit tests for particular sub-components/stages
 - More coverage on pipeline
- Add co-simulation of commit trace against an ISA simulator
- Initial things to verify
 - Specification compliance
 - Branch Prediction
 - TAGE
 - GShare
 - Much more!
- Work on improving simulation visibility



General maintenance

- Documentation additions
 - Update documentation for new additions
 - Transition to ReadTheDocs
- Centralized GitHub organization updates
 - Move all relevant BOOM repo's here for easy access and visibility
- Refactoring
 - Update code for clarity, concision and maintainability

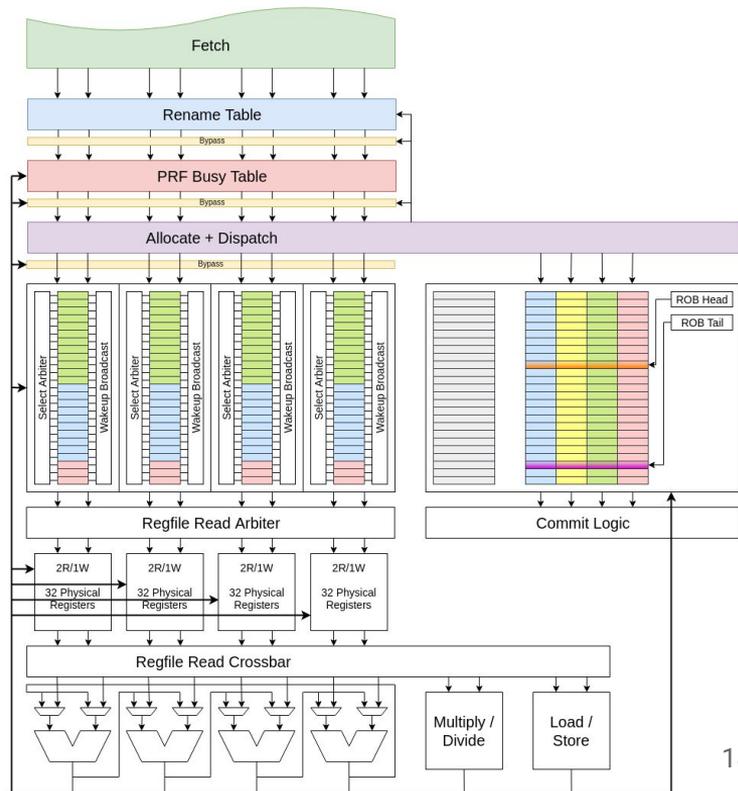
| Feature | BOOM |
|---------------|----------------|
| ISA | RISC-V (RV64C) |
| Synthesizable | ✓ |
| FPGA | ✓ |

Outline

- Current state of the project
- How to get started
- How to contribute
- Things to work on
- **Research ideas**
- Conclusion and discussion

Microarchitectural Exploration

- Chisel + RISC-V ecosystem enables agile architecture research
 - Don't settle on modelling; build and evaluate fully functional systems!
- BOOM is an important piece of groundwork
 - Can be leveraged to research high performance microarchitecture
 - 15,000+ lines of concise, partitioned code: feasible for one person to implement features
- Current ideas:
 - “RingScalar” integer execution core
 - Support for RISC-V vector extension



Hardware Security

- Good news? Not vulnerable to Meltdown or Foreshadow
- Bad news? Vulnerable to general Spectre style attacks
 - Actually good news
- BOOM - baseline for security research
- Leverage Chisel + RocketChip productivity to explore mitigations
 - Secure caches
 - Cache partitioning
 - Timing obfuscation
- A secure cache proof-of-concept is in the works
 - 1 man-week to block transient cache refills
- Enclaves are cool too

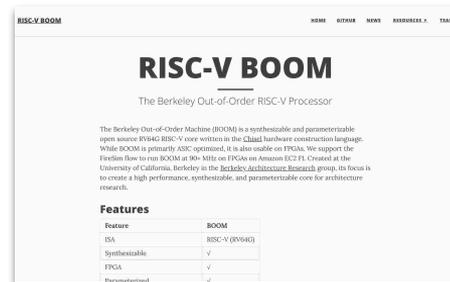


Outline

- Current state of the project
- How to get started
- How to contribute
- Things to work on
- Research ideas
- **Conclusion and discussion**

Conclusion

- Berkeley Out-of-Order Machine
 - Concisely expressed in ~15k LOC Chisel
 - RV64G + Priv v1.11
 - Synthesizable
 - Extensible
 - OPEN SOURCE!
- A lot of work to be done
 - Verification
 - Useful features
- Great opportunity to do make an impact
 - Architecture research
 - Work on an open source project



What do you want to do?

What would you like to see in BOOM and how can we help?

Questions and thanks!

- Get started today
 - Website: riscv-boom.github.io
 - GitHub: [riscv-boom/riscv-boom](https://github.com/riscv-boom/riscv-boom)
- Contact us
 - Google group: <https://groups.google.com/forum/#!forum/riscv-boom>
- Thanks to all who have worked on BOOM
 - Pi-Feng Chiu, Rimas Avizienis, Jonathan Bachrach, Scott Beamer, David Biancolin, Henry Cook, Palmer Dabbelt, John Hauser, Adam Izraelevitz, Sagar Karandikar, Ben Keller, Donggyu Kim, Jack Koenig, Jim Lawson, Yunsup Lee, Richard Lin, Eric Love, Martin Maas, Chick Markley, Albert Magyar, Howard Mao, Miquel Moreto, Quan Nguyen, Albert Ou, Brian Richards, Colin Schmidt, Wenyu Tang, Stephen Twigg, Huy Vo, Andrew Waterman, Angie Wang, and more

Acknowledgements

- Research partially funded by DARPA Award Number HR0011-12-2-0016, the Center for Future Architecture Research, a member of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA, and ASPIRE and ADEPT Lab industrial sponsors and affiliates Intel, Google, Siemens, SK Hynix, Seagate, Huawei, Nokia, NVIDIA, Oracle, and Samsung. Approved for public release; distribution is unlimited. The content of this presentation does not necessarily reflect the position or the policy of the US government and no official endorsement should be inferred. Any opinions, findings, conclusions, or recommendations in this paper are solely those of the authors and does not necessarily reflect the position or the policy of the sponsors.
- Projects mentioned
 - FireSim: Website: <https://fires.im/> Paper: sagark.org/assets/pubs/firesim-isca2018.pdf
 - Hwacha: Website: <http://hwacha.org/>
 - RingScalar: Paper: <https://dspace.mit.edu/handle/1721.1/34012>