03 March 2026

Performance Progress of CVA6-CHERI

Converging on Competitive Performance

Author
Jonathan Woodruff – Capabilities Limited

CVA6-CHERI Overview

- CVA6 is a leading open-source application core
  - (Also used in embedded applications)
- Capabilities Limited has picked up a research project from Zero-day Labs that added the CHERI research ISA to CVA6 and improved it:
  - Implemented the RV64Y draft standard
  - Improved area and timing
  - Performed much greater verification
  - Brought up OpenSBI, CheriBSD, CheriLinux, and Cheri-seL4
  - Rebased on current CVA6 head, inheriting many recent improvements

CVA6-CHERI Overview

- Pushing on performance features
  - Superscalar CVA6-CHERI
    - The rebase brought in mature dual-issue superscalar support
    - We enabled the superscalar build to work with CHERI
  - High-Performance Data Cache
    - The CVA6 project is moving to the new HPDCache
    - The HPDCache efficiently handles large (or even just 64-byte) cache lines

Superscalar CVA6

1. **Register Renaming:** To mitigate Write-After-Write (WAW), register renaming tracks the latest instruction writing to integer or floating-point registers, ensuring correct operand forwarding.
2. **Branch Prediction:** A two-level branch predictor with private history per entry (128 entries, 3-bit history) is integrated, reducing misprediction penalties by 30% across the Embench-IoT suite compared to the simpler bimodal predictor used in CVA6 and CVA6S.
3. **ALU-to-ALU Operand Forwarding:** To reduce execution latency, lightweight operand forwarding is introduced for cases where two ALU instructions are issued in the same cycle. This allows the second instruction to directly use the result of the first one without delay.
4. **FPU Integration:** The floating-point unit (FPU) is integrated into the dual-issue CVA6S+, as CVA6S lacks this feature. To minimize area overhead, it shares the write-back (WB) port with the secondary ALU, requiring additional hazard logic to prevent contention.

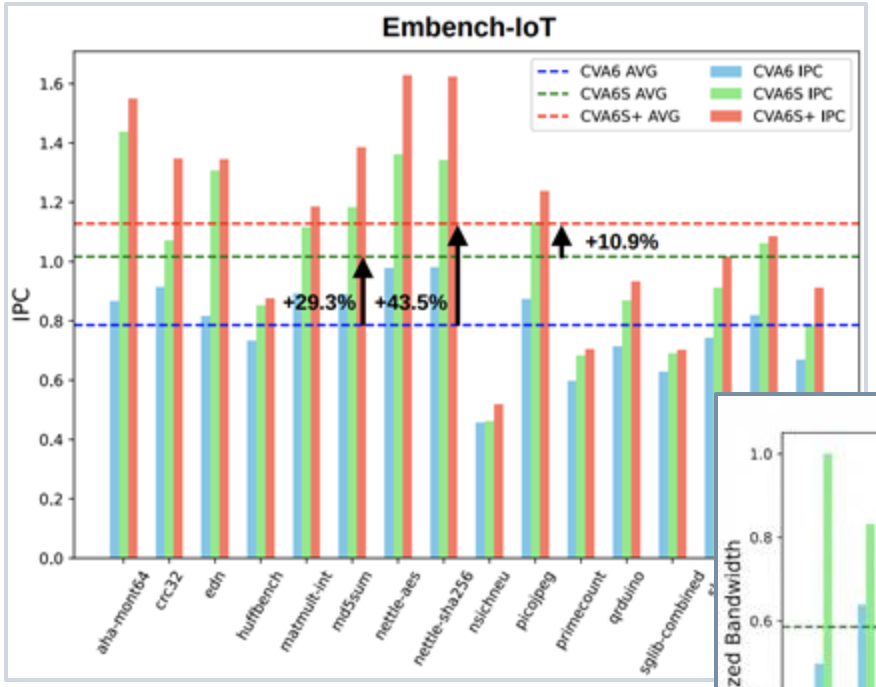

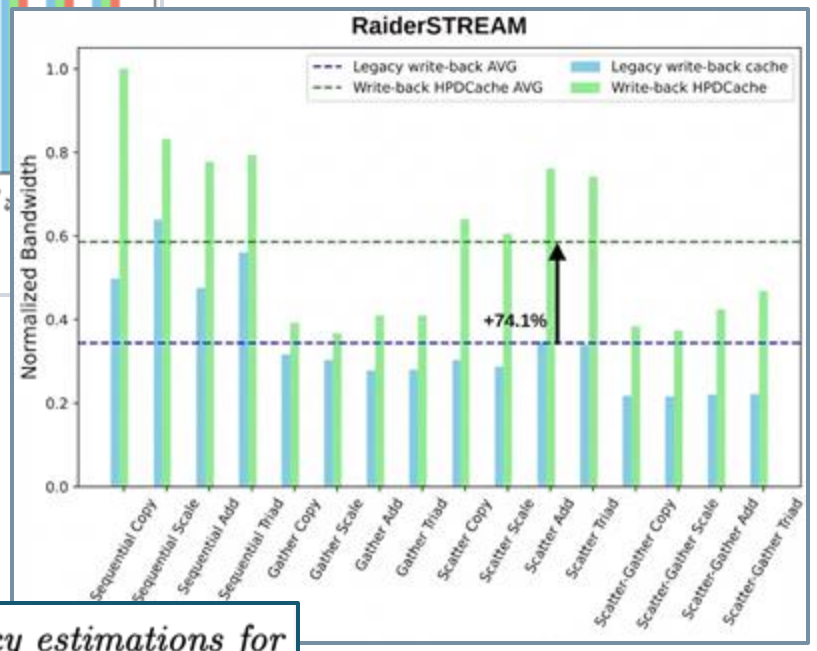

Table 1: *Area at 900 MHz and frequency estimations for the different cores and cache subsystems (16 kiB 4 way ICache, 32 kiB 8 way DCache).*

Core	DCache	Pipeline [mm2]	ICache [mm2]	DCache [mm2]	Max. Freq. [MHz]
CVA6	HPDCache	0.057	0.041	0.077	1090
CVA6S+	Legacy	0.070	0.041	0.095	960
CVA6S+	HPDCache	0.073	0.041	0.077	1095

Tedeschi, Riccardo, Gianmarco Ottavi, Côme Allart, Nils Wistoff, Zexin Fu, Filippo Grillotti, Fabio De Ambroggi et al. "**CVA6S+: A superscalar RISC-V core with high-throughput memory architecture**." arXiv preprint arXiv:2505.03762 (2025).

CHERI

Superscalar CVA6-CHERI

- We added the CHERI CLU
  - A single CHERI CLU to save area and timing
  - As opposed to two ALUs
- Branches and Capability operations share ALU0, and so cannot be issued together
- Otherwise, we don't add any other scheduling restrictions

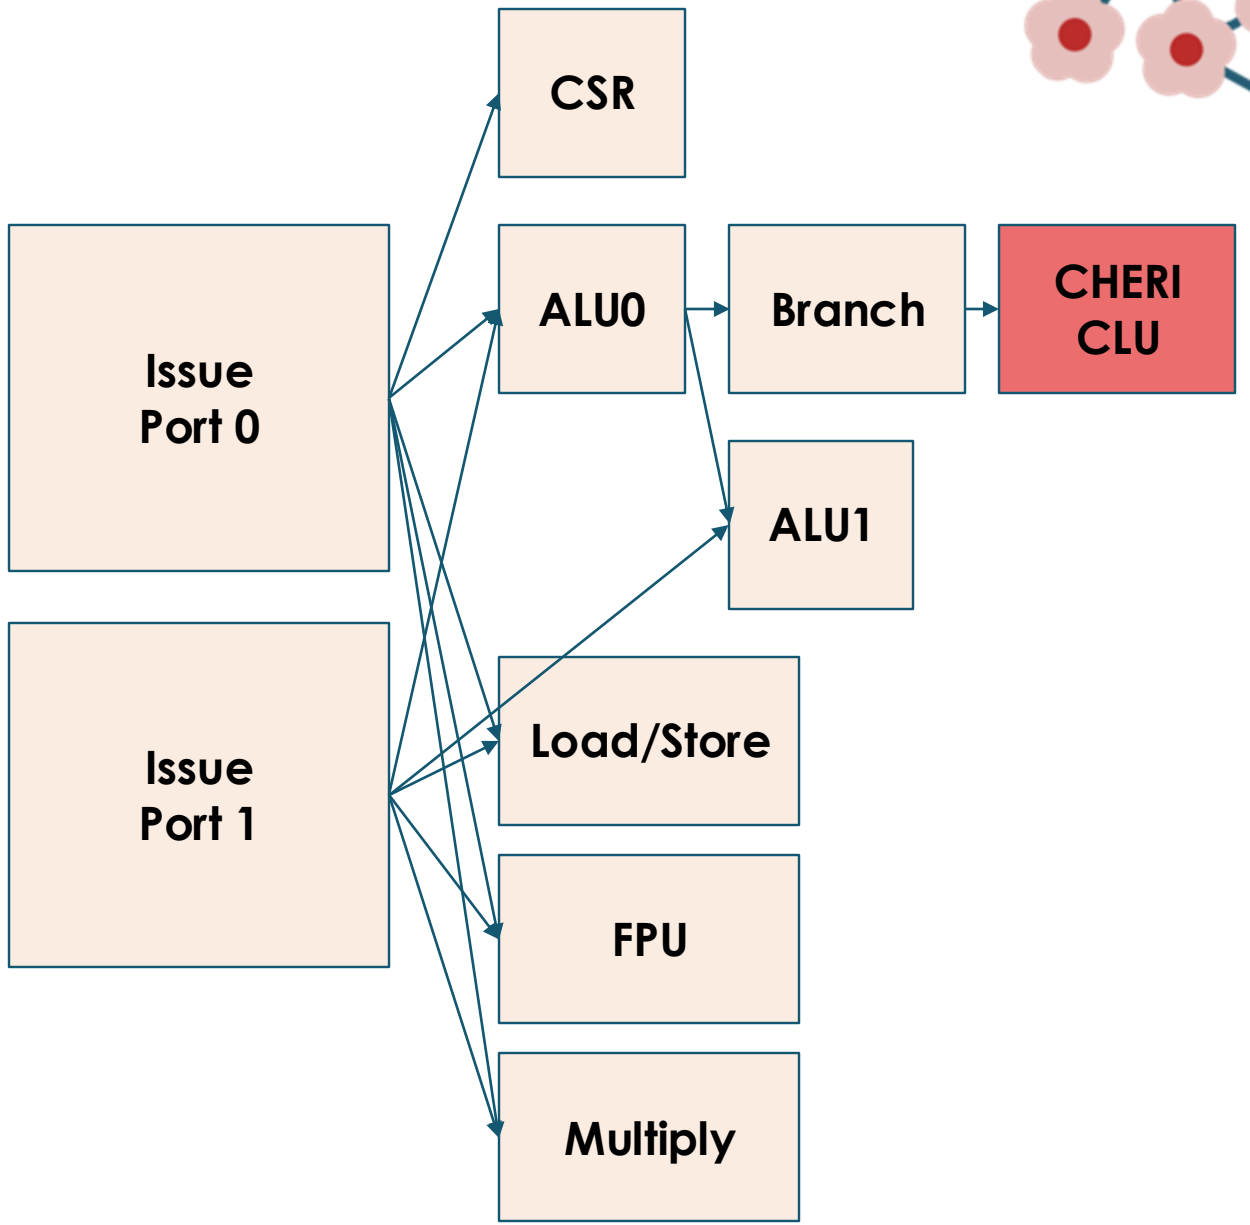

CVA6-CHERI Superscalar Performance

- IPC 1.09 for rv64
- IPC 1.02 for rv64y
- Opportunities for parity
  - Tie CLU to ALU2?
  - Enable CADD on ALU2?

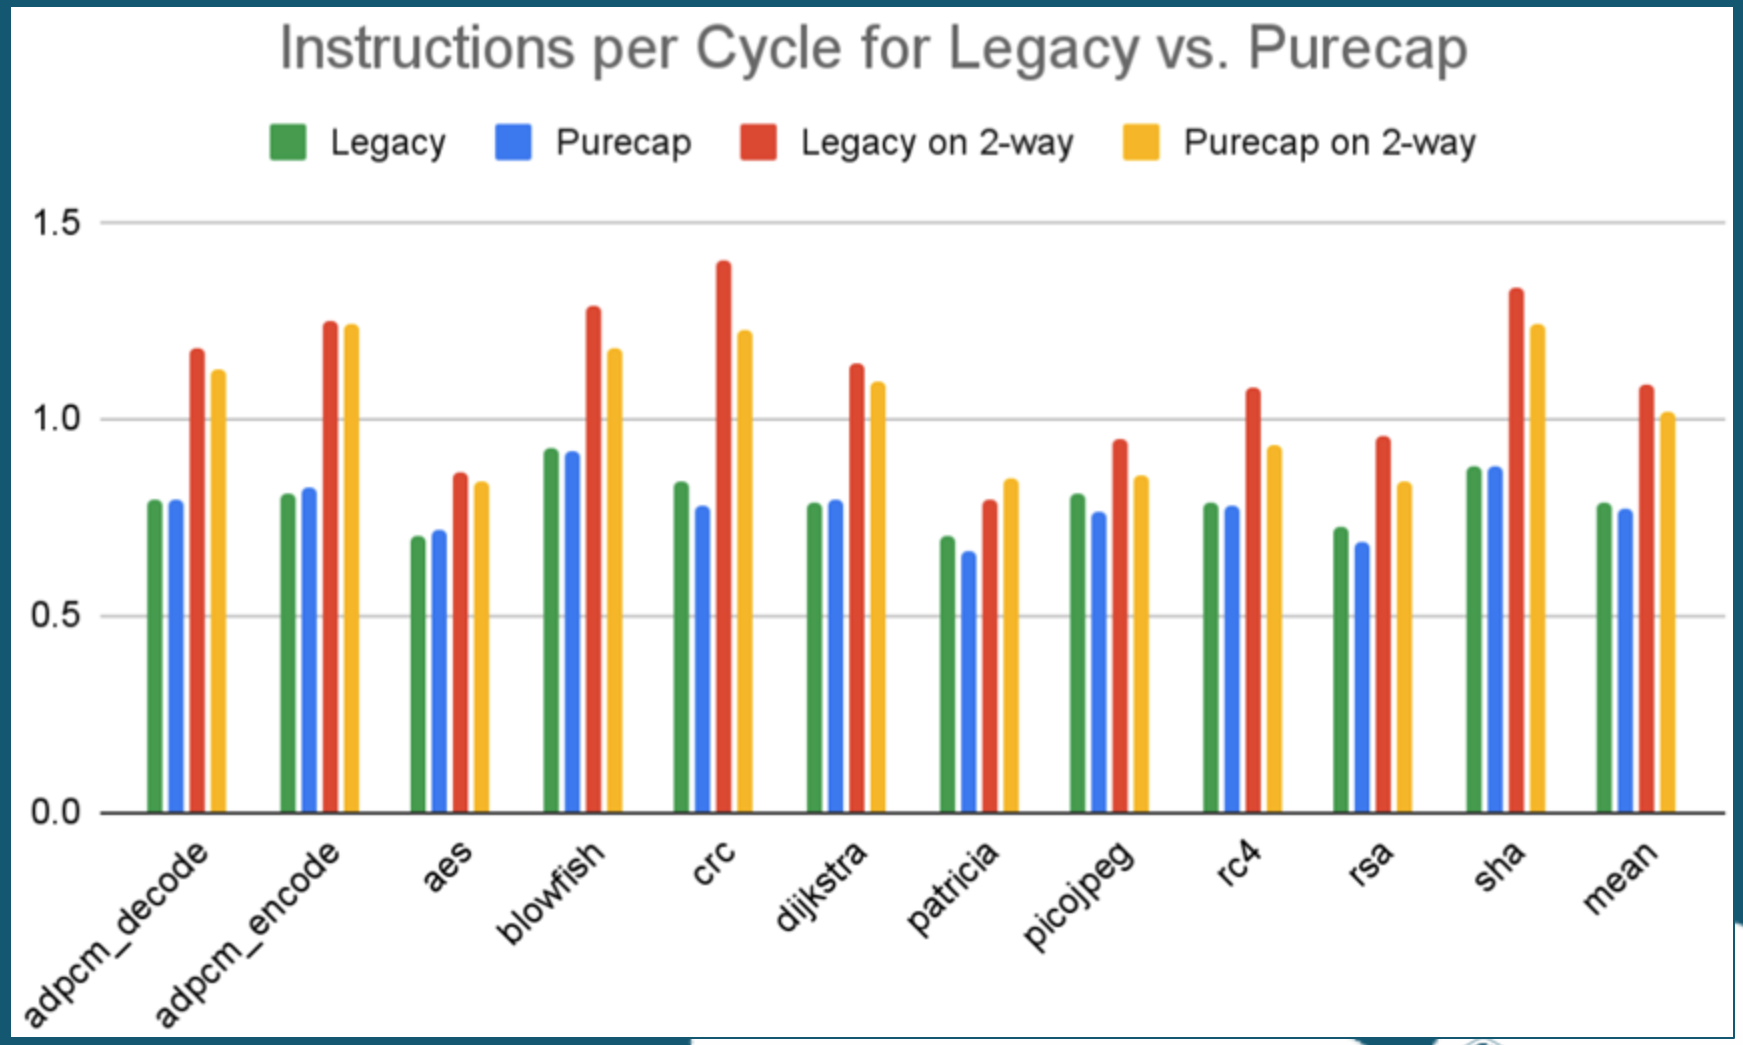

CHERI

High-Performance Data Cache

- Set-associative cache with configurable (at compilation) number of ways and sets;
- Multiple ports for requests with support for wide data;
- Write-Through Policy with write-buffer supporting write-coalescing and multiple outstanding write requests;
- Pipelined micro-architecture for high clock frequencies;
- Single-cycle latency for read hit and write requests;
- Non-blocking pipeline with read under multiple misses;
- OoO execution of requests to prevent head-of-line blocking;
- Support for atomic memory operations (AMOs);
- Support for Cache Management Operations (CMOs) [2];
- Dedicated Configuration-and-Status Registers (CSR) address space to access performance counters and runtime configuration registers;
- Stride-based, programmable, hardware memory prefetcher;
- Native FIFO-like (ready-valid) memory interfaces and adapter for AMBA AXI5 interface;
- Different bus widths are supported on the memory interface, from 64 to 512 bits.

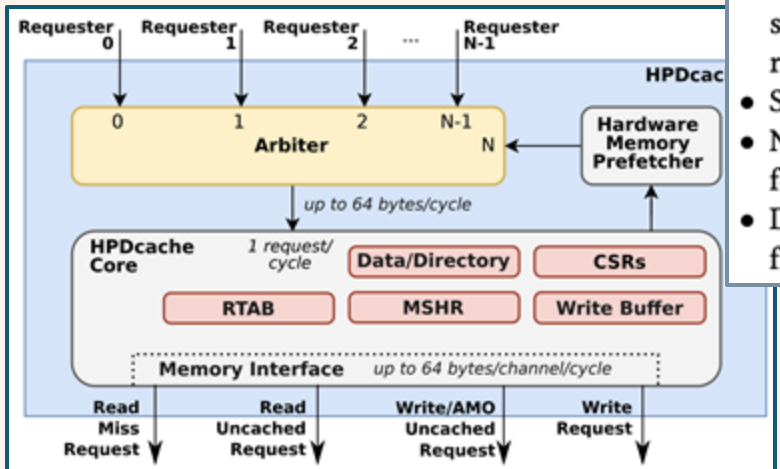

Figure 1: HPDcache Micro-architecture Overview

Fuguet, César. "HPDcache: Open-source high-performance L1 data cache for RISC-V cores." In Proceedings of the 20th ACM International Conference on Computing Frontiers, pp. 377-378. 2023.

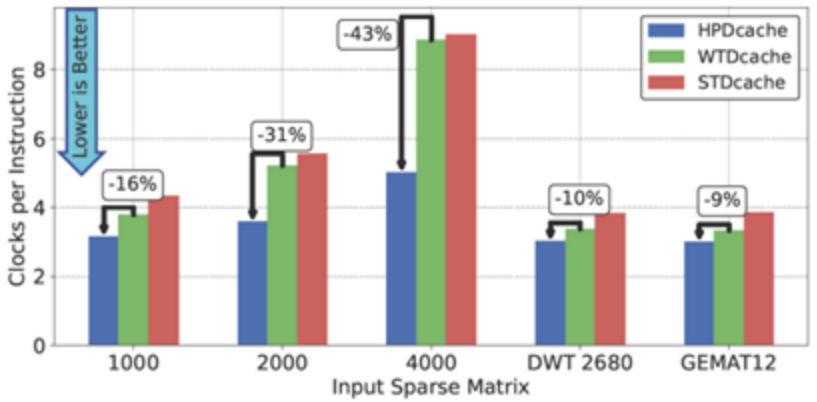

Fig. 1. Performance Comparison Based on Clocks per Instruction for SpMV

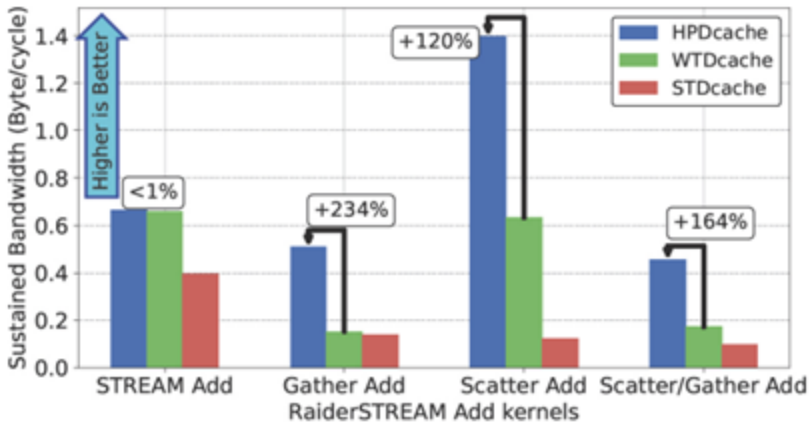

Fig. 2. Bandwidth Measurements using Add Kernels from RaiderSTREAM

Million, Davy, Noelia Oliete-Escuín, and César Fuguet. "Breaking the memory wall with a flexible open-source l1 data-cache." In 2024 Design, Automation & Test in Europe Conference & Exhibition (DATE), pp. 1-2. IEEE, 2024.

CHERI

CVA6-CHERI High-Performance Data Cache

- We have added CHERI Tag management to the HPDCache
- This allows us flexibility for choosing line sizes larger than 32-bytes without further increasing the width of the SRAM memories
- We will likely use 64-byte lines, but would like to try 128-byte lines to see if it improves pure-capability performance
  - 64-byte lines were standardised for 32-bit architectures and could hold 16, 32-bit pointers
  - 64-byte lines only hold 4 capability pointers

Dual-issue, HPDCache and SPECInt2006 Benchmarks

- Average of 31% cycles eliminated in SPECInt2006 benchmarks
- Average of 53% Data cache load misses eliminated
  - 64-byte vs 32-byte cache lines
  - Write-back vs. Write-through

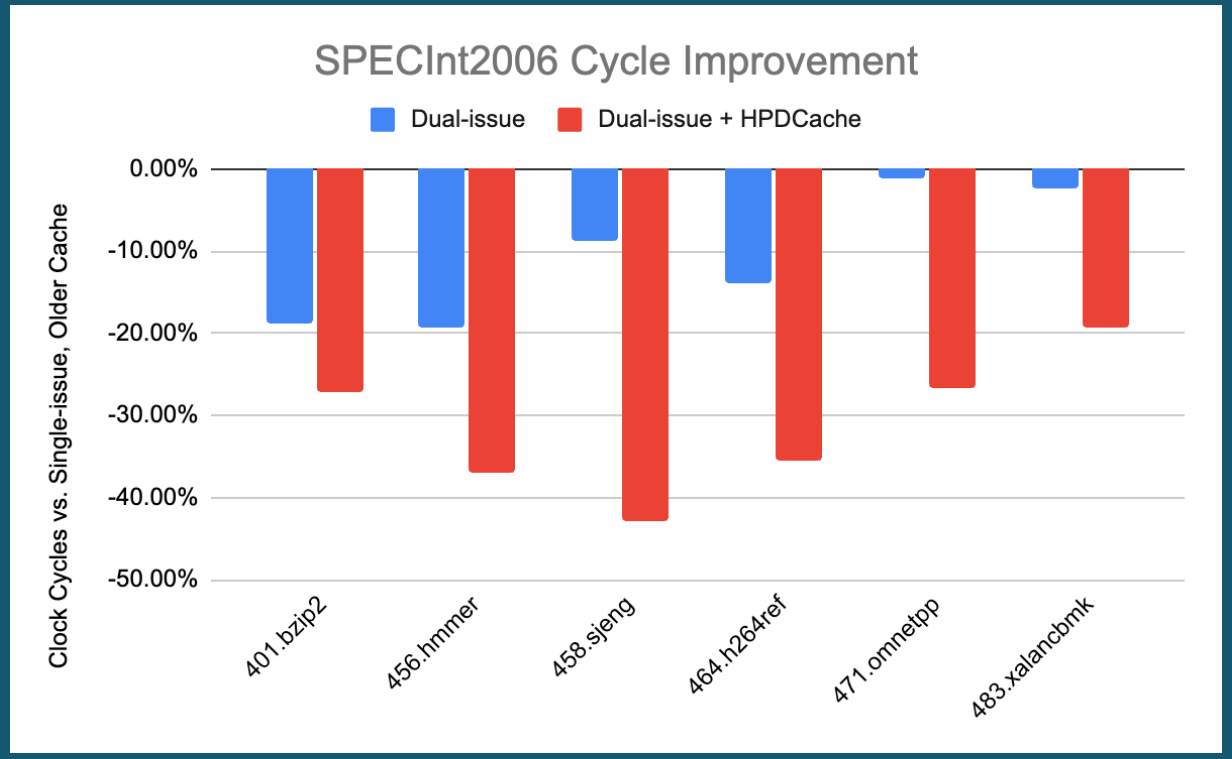

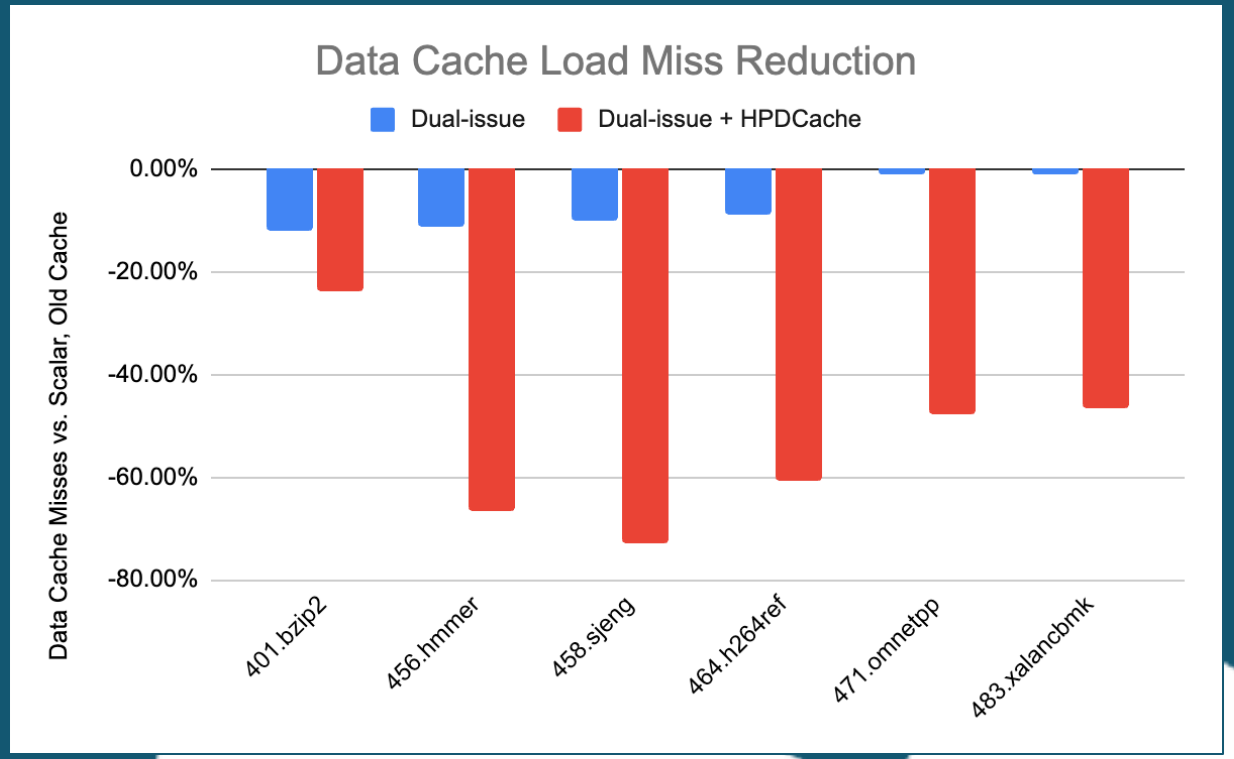

CHERI

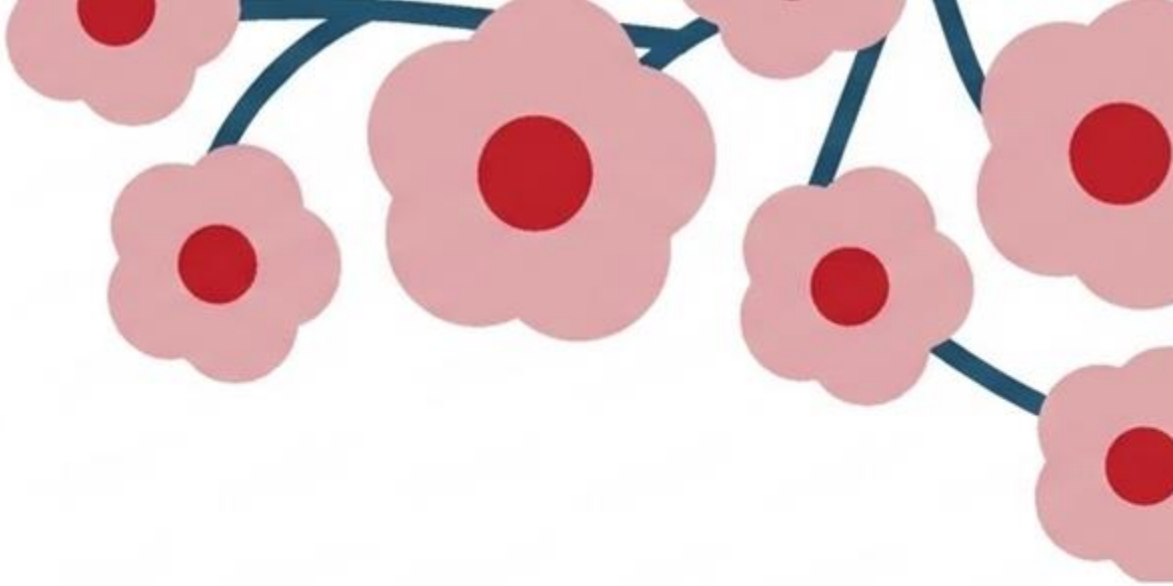

CHERI

Questions?

Contact: Jonathan.Woodruff@capabilitieslimited.co.uk
Website: https://www.capabilitieslimited.co.uk/current-projects/cheri-cva6