
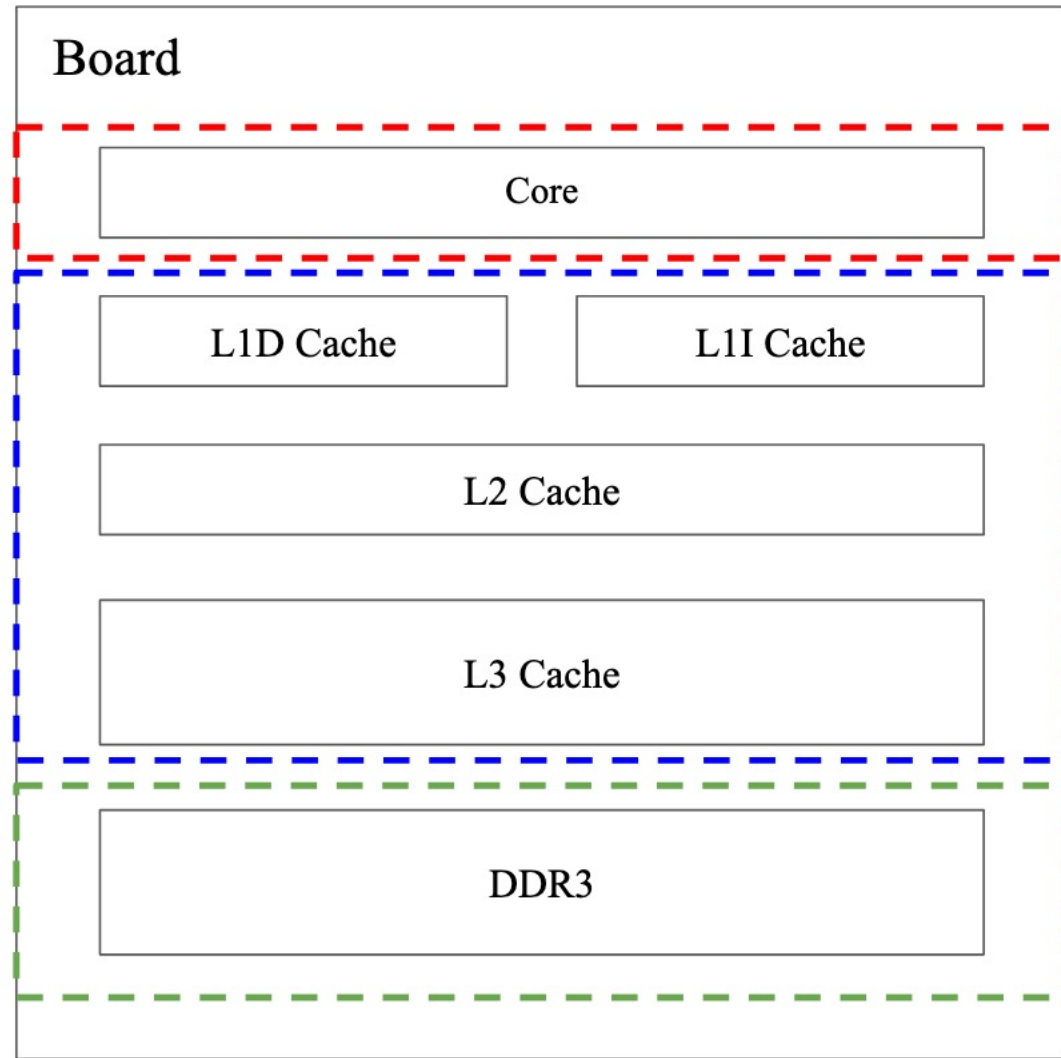


Potential and Limitation of High-Frequency Cores and Caches

Kunal Pai, Anusheel Nand & Jason Lowe-Power

Motivation

- Model emerging tech with very high clocks
 - Cryogenic computing
 - Superconducting circuits
- We show gem5 is a viable modeling framework
 - Run realistic workloads
 - Show high speedups for some workloads
 - Show cost of data movement






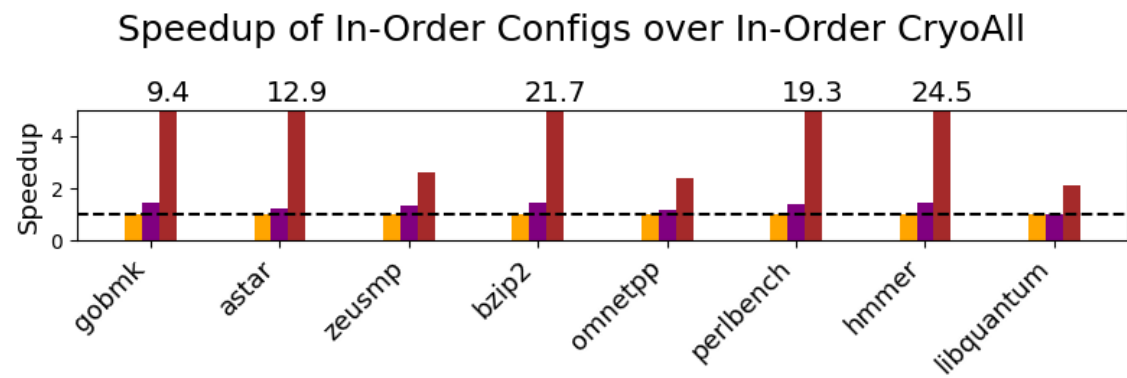
Experimental Set-Up

- Cryogenic components in gem5: BOOM (OOO), HiFive Unmatched (In-Order), 3-level cache hierarchy
- Clock freq.: **4 GHz**.
- Superconducting variants: same μ arch as cryo components
- Clock freq.: **100 GHz**.
- Simulated configurations:
 - **CryoAll**: Cryo BOOM + Cryo Cache.
 - **SuperCryo**: Super BOOM + Cryo Cache.
 - **SuperAll**: Super BOOM + Super Cache.
 - **In-Order CryoAll**: Cryo HiFive Unmatched + Cryo Cache.
 - **In-Order SuperCryo**: Super HiFive Unmatched + Cryo Cache.
 - **In-Order SuperAll**: Super HiFive Unmatched + Super Cache.
- Ran **full-sized workloads**: SimPoints, SPEC 2006 ref size.



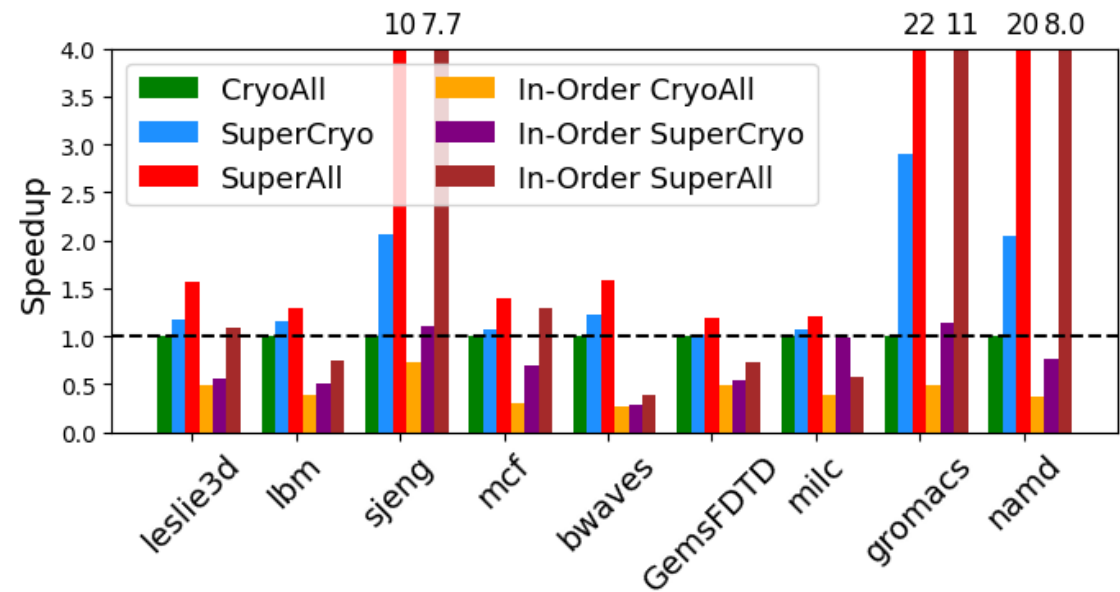

Performance Improvement

- More impact on in-order
 - Latency hiding less important
- High potential bar:
 - but low freq. caches are bottleneck
- Memory-intensive workloads:
 - minimal improvement
- Main bottleneck:
 - “Room Temp” DRAM



In-Order vs. Out-of-Order

- Out-of-order: more speedup
- Out-of-order CryoAll:
 - faster than In-Order CryoAll & In-Order SuperCryo
- Take away: Big potential benefits, but only for some workloads: **Accelerator**



Data Movement

- Superconducting memory is hard to scale up: CryoAll and SuperCryo (in-order and OOO) – realistic configs.
- Need high bandwidth to the cache for speedups.
- Maximum 500 GB/s for L1D Cache in SuperCryo configuration.
 - Reasonable for optics!

