ZSIM TUTORIAL

Validation

Outline

- Introduction
- Methodology
- Single-threaded results
- Multi-threaded results
- Contention models
- Conclusion

Introduction

- How accurate is a simulator?
- What are the sources of inaccuracies?
- What kind of workloads and studies is a simulator intended for?
- Important to do validation before using a simulator.

Tony Nowatzki et.al, Architectural Simulators Considered Harmful, IEEE MICRO 2015

Validation in ZSim

- Micro-benchmarks that stress different micro-architectural structures and events.
  - Ex. Time taken to do integer add, multiply.
  - Lets us catch even minor modeling inaccuracies.

- Wide range of workloads from different benchmark suites
  - Single threaded – SPECCPU2006
  - Multi threaded – PARSEC, SPLASH2, SPECOMP 2001

Comparison to other simulators

- ZSim has an average error of 10% for both single-threaded and multi-threaded workloads.

- MARSS
  - Cycle accurate OOO x86 model
  - Performance differences range from -59% to 50% with only 5 benchmarks being within 10%
- Sniper
  - Approximate OOO model
  - Absolute errors over 50% on SPLASH2 benchmarks
- Graphite, Hornet, SlackSim – no known validation study

Methodology

- Zsim models an x86 core model.
  - It is possible to validate against real hardware system.
- We run each application on the real machine and also simulate it on zsim.
- We record several relevant performance counters on the real machine.
  - Compare them against zsim's results.
- We perform multiple profiling and simulation runs to avoid noisy comparisons.

System Configuration

We validate ZSim against a Westmere system.

HW	Xeon L5640 (6-core Westmere), 24GB DDR3-1333, no hyperthreading, turbo/DVFS disabled
SW	Linux 3.5 x86-64, gcc 4.6.2, Pin 2.12
Bound-weave	1000-cycle intervals, 6 weave threads
Cores	6 x86-64 OOO cores at 2.27GHz
L1I caches	32KB, 4-way, LRU, 3-cycle latency
L1D caches	32KB, 8-way, LRU, 4-cycle latency
L2 caches	256KB, 8-way, LRU, 7-cycle latency, private
L3 cache	12MB, 16-way, hashed, 6 2MB banks, 14-cycle bank lat, shared, inclusive, MESI coherence w/ in-cache directory, 16 MSHRs
Network	Ring, 1-cycle/hop, 5-cycle injection latency
Mem ctrl	1 controller, 3 DDR3 channels, closed page, FCFS scheduling, fast powerdown with threshold timer = 15 mem cycles
DRAM	24GB, DDR3-1333, 2 4GB RDIMMs per channel

Hardware and Software Configuration of the real system and the corresponding ZSim configuration

Single-threaded validation

- Validate OOO core model with the full SPEC CPU2006 suite.
- Run each application for 50 billion instructions using ref(largest) input set.

Comparing instruction range: 0-50 B

App	Real/Sim IPC		Real/Sim UPC		PerfErr	R/S	BrMPKI	UopErr	ApprxIns	Real/Sim L1I/L1D/L2/L3 MPKI	Real L1D TLB MPKI	KrnCyc	KrnIns	MIPS	Slwdwn
434.zeusmp-ref	1.05	1.05	1.22	1.25	0.00	0.18	0.32	0.02	0.0000	0.01 / 23.71 / 7.37 / 5.61	0.00 / 23.71 / 7.35 / 5.69	0.24	0.0060	0.0006	16.2	147.0
459.GemsFDTD-ref	0.66	0.66	0.85	0.89	0.00	0.19	0.24	0.04	0.0000	0.17 / 35.63 / 29.25 / 20.39	0.10 / 35.38 / 28.32 / 20.62	1.10	0.0077	0.0023	9.1	165.6
473.astar-ref	0.56	0.56	0.63	0.62	0.00	38.82	45.65	-0.01	0.0003	0.01 / 31.96 / 22.63 / 3.07	0.00 / 29.97 / 23.08 / 3.12	1.85	0.0040	0.0033	13.3	95.7
410.bwaves-ref	1.49	1.48	1.78	1.77	0.01	4.64	4.14	0.00	0.0000	0.01 / 4.95 / 3.76 / 3.32	0.00 / 5.70 / 4.02 / 3.52	0.01	0.0143	0.0016	24.2	139.7
433.milc-ref	0.64	0.63	0.74	0.73	0.01	0.10	0.15	-0.00	0.0041	0.01 / 19.80 / 19.72 / 19.06	0.00 / 20.89 / 20.77 / 20.55	0.01	0.0117	0.0030	12.4	117.3
464.h264ref-ref	1.85	1.87	2.17	2.18	-0.01	1.56	1.56	-0.01	0.0000	0.56 / 3.55 / 2.06 / 0.02	0.49 / 3.46 / 2.02 / 0.02	0.01	0.0027	0.0006	21.8	192.1
454.calculix-ref	1.71	1.74	1.81	1.84	-0.01	1.23	1.51	0.00	0.0000	0.14 / 5.71 / 3.32 / 0.22	0.08 / 5.71 / 3.64 / 0.28	0.05	0.0060	0.0020	29.1	133.6
456.hmmer-ref	1.82	1.79	2.63	2.59	0.02	0.19	0.21	-0.00	0.0001	0.03 / 8.01 / 3.62 / 0.04	0.01 / 5.65 / 3.42 / 0.06	0.00	0.0019	0.0005	17.9	230.0
445.gobmk-ref	1.13	1.11	1.30	1.27	0.02	16.29	23.20	-0.01	0.0000	12.00 / 5.03 / 2.48 / 0.45	11.62 / 4.48 / 2.23 / 0.46	0.07	0.0023	0.0005	15.1	169.8
401.bzip2-ref	1.21	1.29	1.40	1.48	-0.06	11.67	12.75	-0.00	0.0001	0.00 / 14.33 / 8.14 / 0.08	0.00 / 14.13 / 8.52 / 0.11	0.04	0.0044	0.0007	20.0	137.8
416.gamess-ref	1.88	1.77	2.26	2.12	0.06	1.36	1.72	-0.01	0.0000	2.03 / 2.97 / 0.19 / 0.00	1.65 / 2.56 / 0.20 / 0.00	0.00	0.0016	0.0003	19.1	223.8
453.povray-ref	1.59	1.70	2.09	2.16	-0.06	3.92	2.56	-0.03	0.0000	2.25 / 10.28 / 0.20 / 0.00	1.80 / 10.51 / 0.01 / 0.00	0.61	0.0015	0.0007	15.8	228.9
462.libquantum-ref	0.55	0.59	0.48	0.52	-0.06	0.00	0.00	-0.00	0.0000	0.00 / 32.89 / 32.90 / 32.84	0.00 / 32.89 / 32.89 / 32.84	0.04	0.0027	0.0009	11.8	105.7
481.wrf-ref	1.84	2.00	1.78	1.84	-0.08	2.12	2.46	-0.05	0.0007	0.22 / 3.13 / 2.00 / 0.32	0.19 / 3.40 / 2.22 / 0.51	0.01	0.0178	0.0050	29.5	141.5
403.gcc-ref	1.10	1.02	1.31	1.19	0.08	7.43	13.47	-0.02	0.0005	6.79 / 13.13 / 6.39 / 0.65	6.63 / 12.62 / 6.32 / 0.65	1.17	0.0046	0.0019	14.9	167.6
444.namd-ref	1.62	1.77	1.81	1.97	-0.08	1.86	1.78	-0.00	0.0000	0.00 / 6.94 / 0.19 / 0.03	0.00 / 7.13 / 0.18 / 0.04	0.00	0.0022	0.0007	27.1	136.1
465.tonto-ref	1.80	1.97	2.25	2.39	-0.08	1.10	1.80	-0.03	0.0000	1.75 / 6.45 / 1.29 / 0.00	1.47 / 6.43 / 1.19 / 0.00	0.03	0.0022	0.0007	20.8	197.1
458.sjeng-ref	1.46	1.32	1.60	1.51	0.10	9.79	17.23	0.03	0.0007	1.83 / 1.61 / 0.45 / 0.29	2.32 / 1.39 / 0.57 / 0.45	0.16	0.0036	0.0005	19.1	173.0
447.dealII-ref	1.52	1.69	1.91	2.12	-0.10	3.09	3.07	-0.00	0.0000	0.24 / 5.84 / 2.25 / 0.30	0.16 / 5.86 / 2.33 / 0.40	0.60	0.0070	0.0030	21.0	164.3
400.perlbench-ref	1.62	1.47	1.87	1.69	0.11	5.87	5.68	-0.00	0.0000	6.59 / 4.05 / 1.88 / 0.04	4.40 / 3.62 / 1.65 / 0.08	0.20	0.0086	0.0023	18.3	200.9
450.soplex-ref	0.58	0.66	0.67	0.74	-0.12	5.20	5.00	-0.03	0.0011	0.10 / 26.14 / 19.39 / 13.01	0.27 / 22.76 / 18.37 / 12.36	4.29	0.0072	0.0034	12.7	103.1
482.sphinx3-ref	1.33	1.52	1.63	1.85	-0.12	3.33	2.99	-0.01	0.0001	0.07 / 16.66 / 12.50 / 0.28	0.04 / 14.91 / 12.52 / 0.47	0.31	0.0034	0.0011	17.8	169.9
435.gromacs-ref	1.12	1.31	1.37	1.59	-0.14	2.00	2.10	-0.00	0.0001	0.01 / 9.10 / 2.38 / 0.00	0.00 / 8.68 / 2.00 / 0.01	0.00	0.0014	0.0006	21.2	120.2
437.leslie3d-ref	0.69	0.60	0.86	0.80	0.15	0.13	0.72	0.07	0.0000	0.04 / 38.83 / 27.84 / 19.33	0.01 / 38.59 / 27.96 / 19.42	0.06	0.0031	0.0010	9.2	169.7
429.mcf-ref	0.44	0.38	0.50	0.43	0.17	15.23	18.29	0.01	0.0000	0.01 / 86.73 / 70.43 / 26.67	1.76 / 77.63 / 72.03 / 21.40	9.23	0.0085	0.0015	6.6	151.7
436.cactusADM-ref	0.86	1.03	1.30	1.50	-0.17	0.02	0.03	-0.04	0.0000	3.85 / 14.68 / 10.18 / 5.76	5.22 / 14.48 / 9.99 / 4.96	7.91	0.0074	0.0016	11.2	173.5
483.xalancbmk-ref	0.86	1.05	0.95	1.14	-0.18	1.98	5.74	-0.02	0.0000	5.83 / 29.48 / 28.45 / 0.33	4.20 / 27.72 / 26.90 / 0.37	3.61	0.0054	0.0027	14.0	139.1
471.omnetpp-ref	0.71	0.94	0.76	1.01	-0.25	3.75	5.14	0.01	0.0003	1.38 / 39.49 / 35.52 / 5.17	0.52 / 33.58 / 31.62 / 5.09	13.42	0.0039	0.0023	10.4	154.6
470.lbm-ref	0.64	0.87	0.73	0.99	-0.26	0.05	0.04	0.01	0.0000	0.01 / 51.68 / 31.44 / 16.06	0.00 / 53.16 / 34.25 / 16.03	0.00	0.0043	0.0008	11.5	127.0

IPC Error

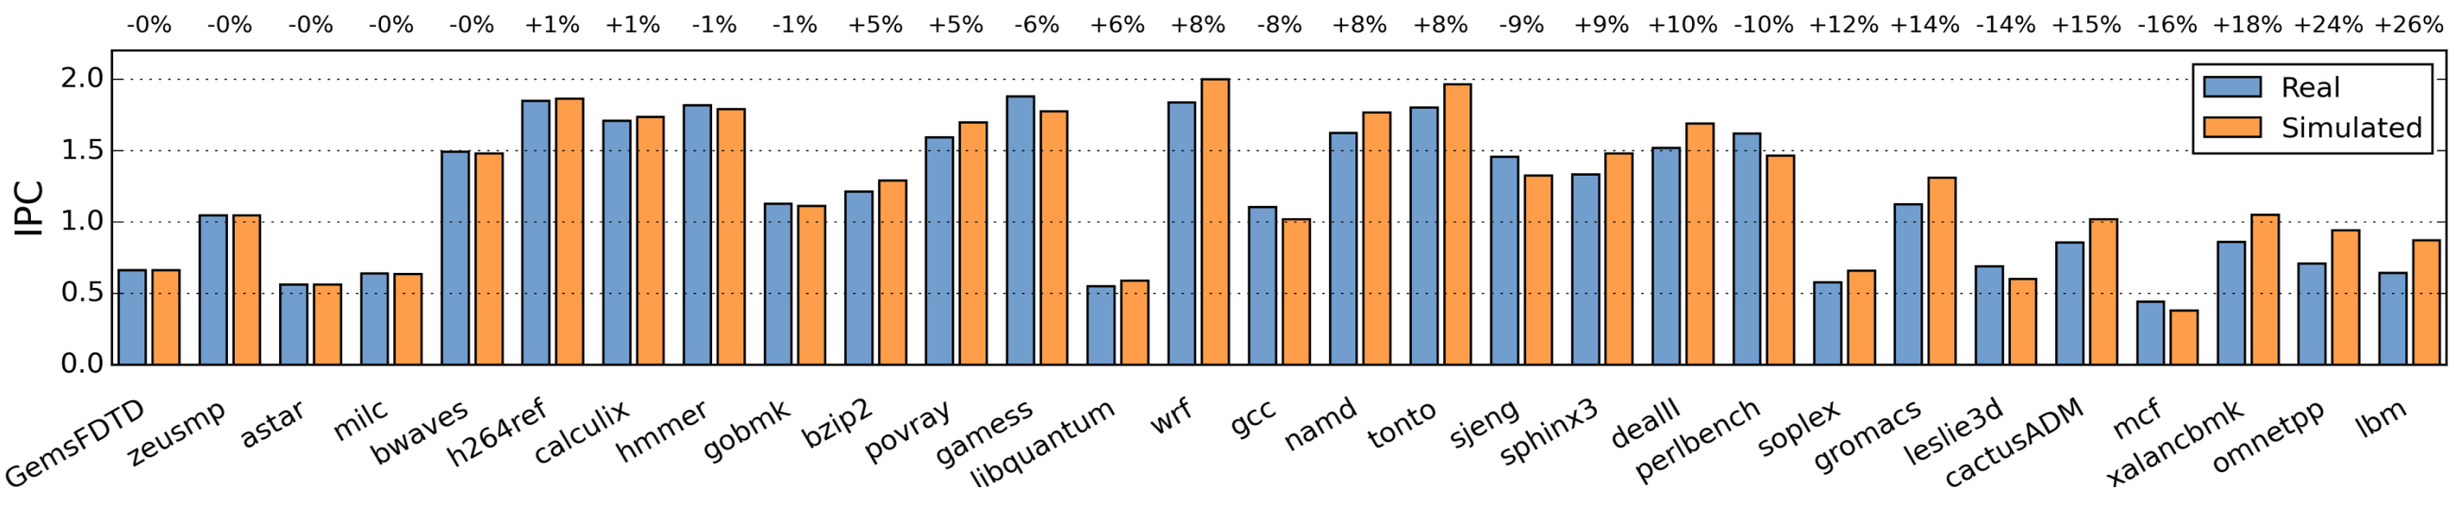

- Average absolute IPC error is 8.5%.
- Max error is 26%
- In 21 out of the 29 benchmarks, error is less than 10%.

MPKI Errors for different caches

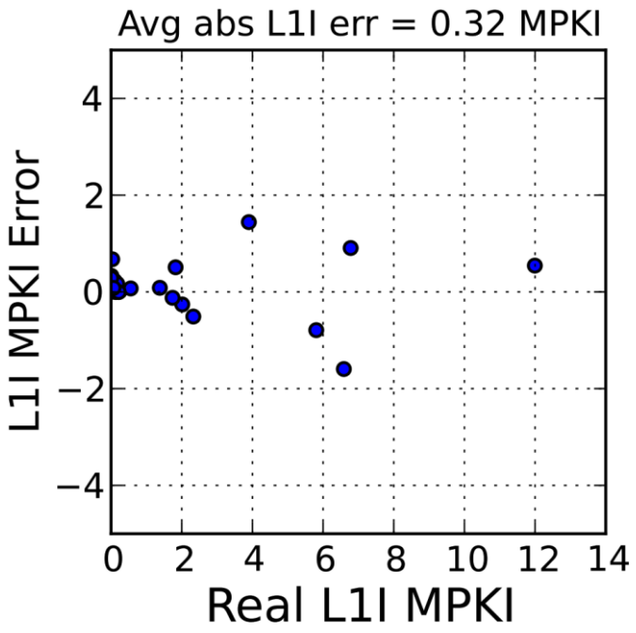

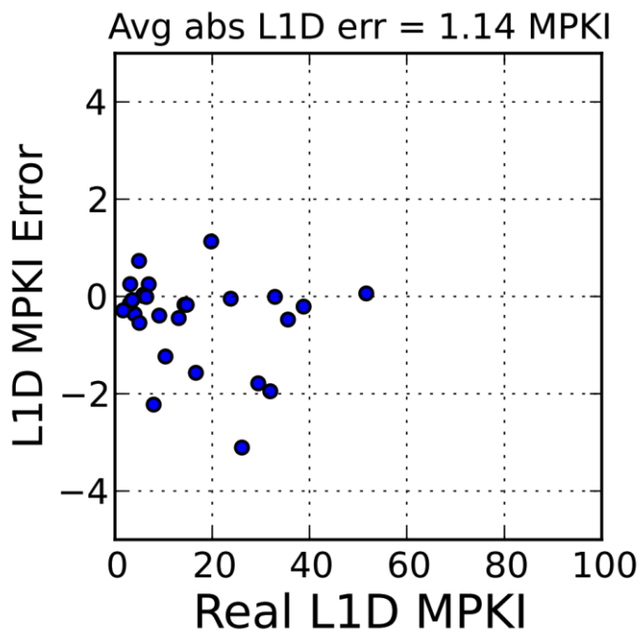

Average Absolute MPKI errors

L1i - 0.32

L1d - 1.14

L2 - 0.59

L3 - 0.30

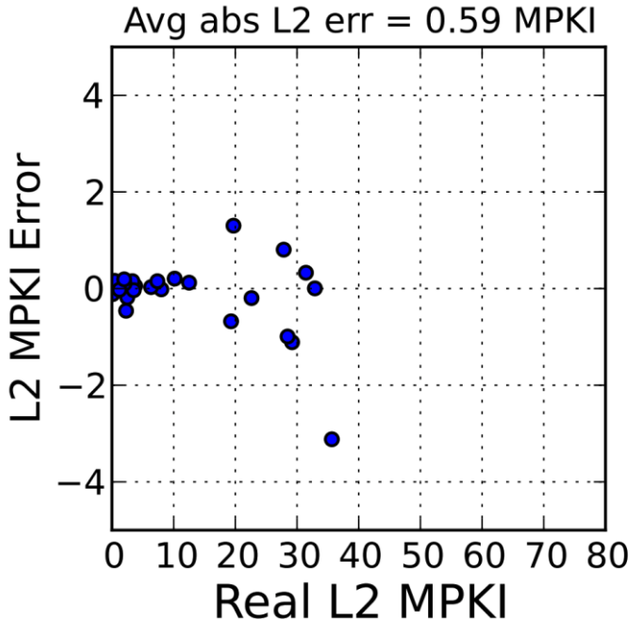

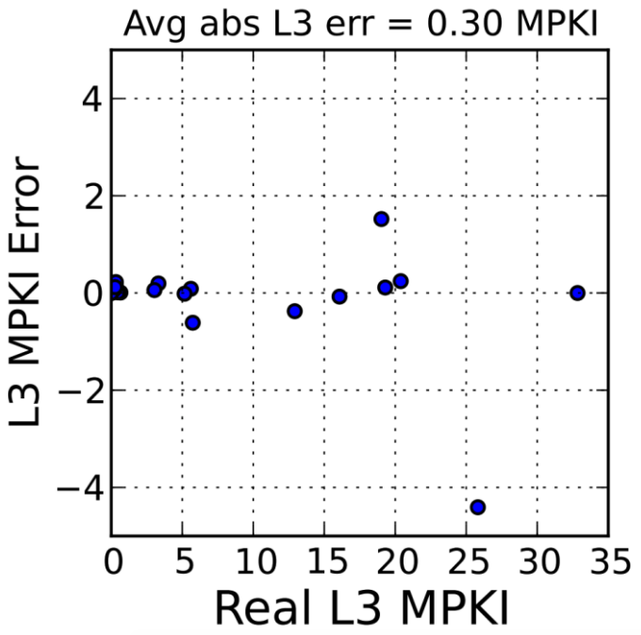

Traces

IPC Trace

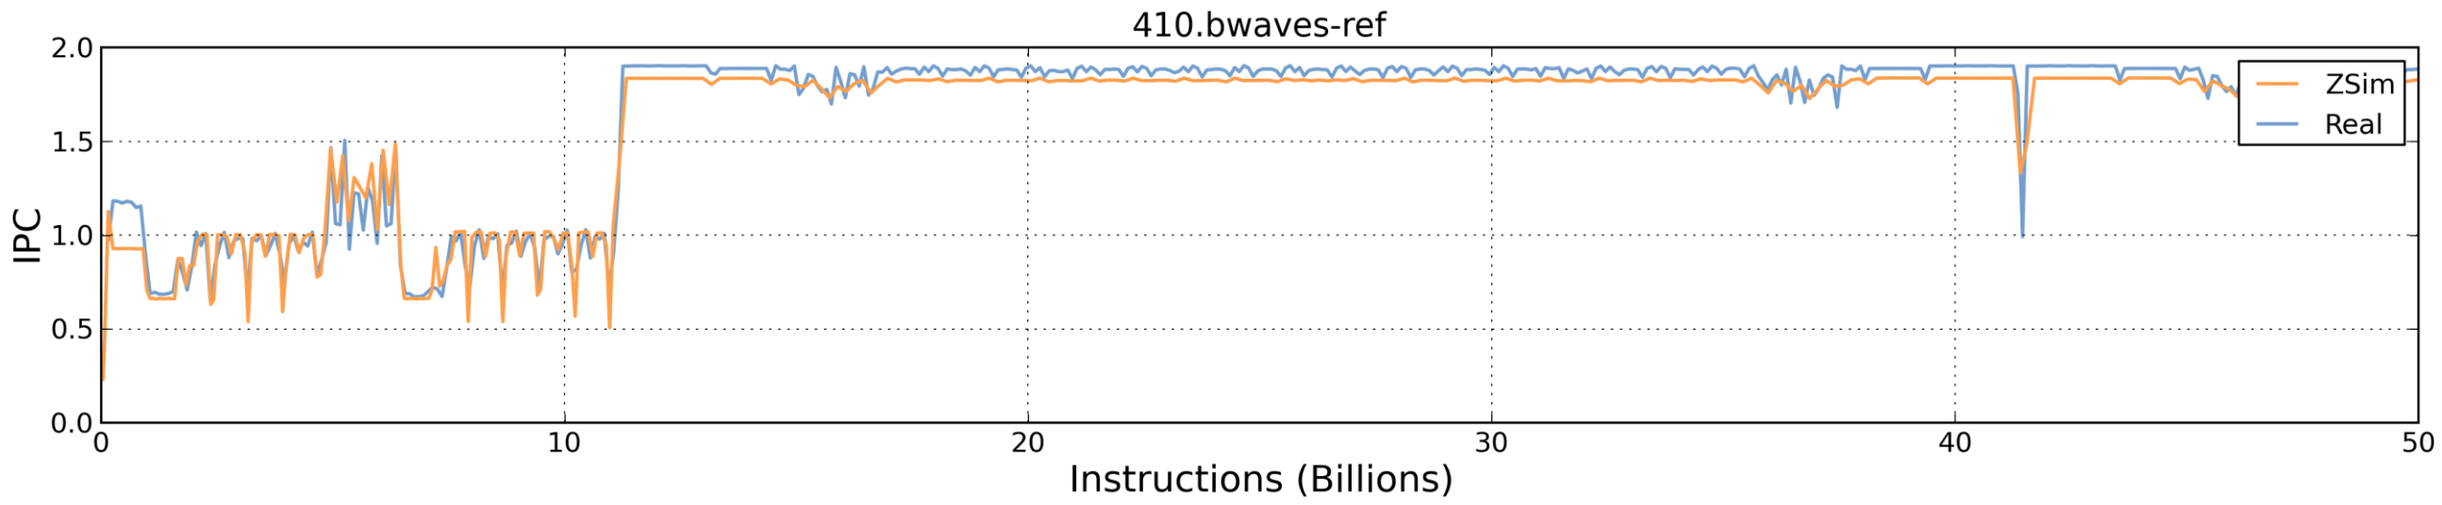

L3 MPKI Trace

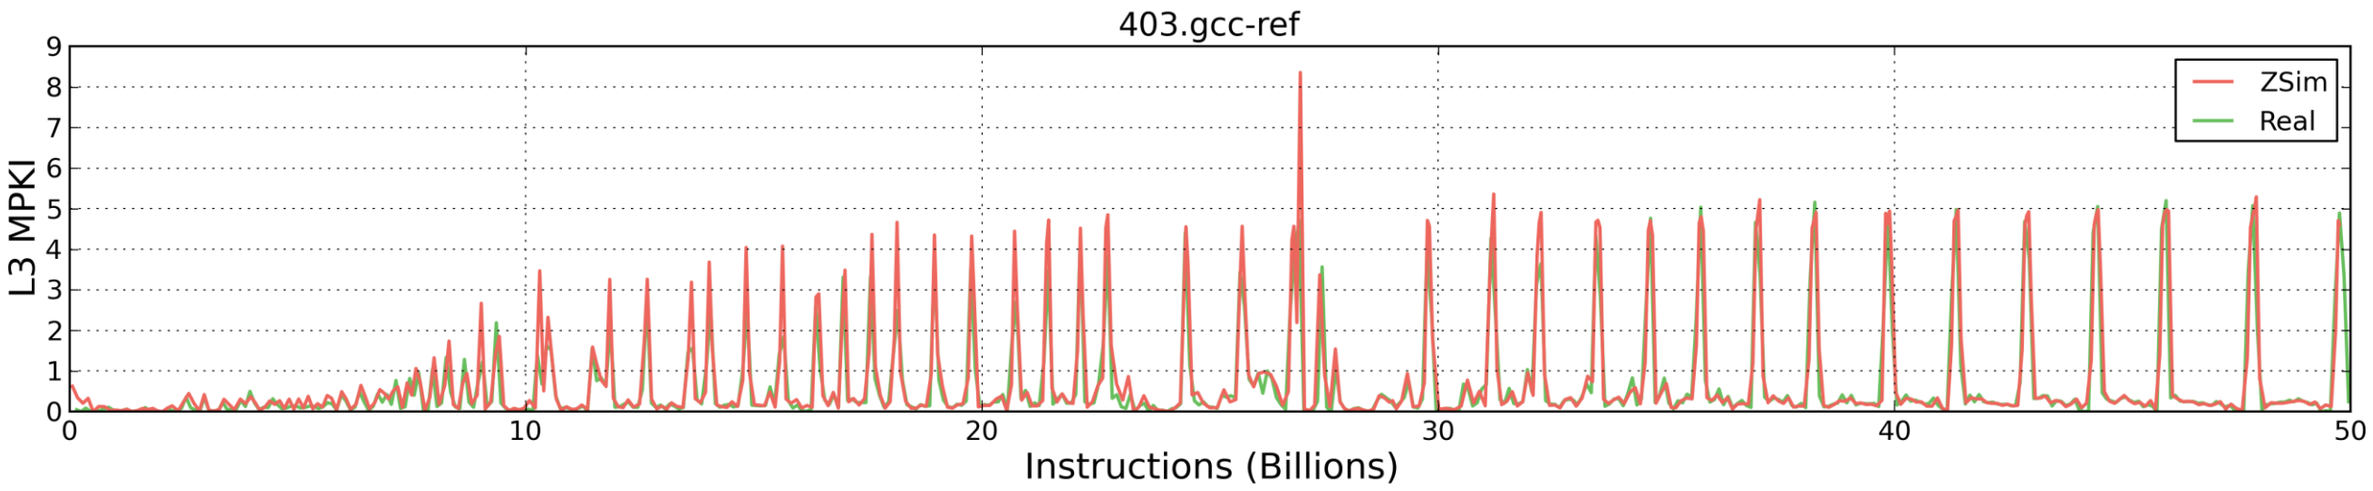

Major sources of error

- Does not model TLB and page table walkers.
- Inaccuracies in the front end model.
  - The modeled 2-level branch predictor with an idealized BTB has significant errors in some cases.
- Most of the errors are observed in benchmarks that have non-negligible TLB misses.
- It is difficult to figure out the exact details of a processor's architecture.

µop coverage

- ZSim implements decoding for the most frequently used op-codes.
- Only 0.01% of executed instructions have an approximate dataflow decoding
- Modern compilers only produce a fraction of the x86 ISA.
- Ignores micro-sequenced instructions.
- Uop error = (uop real – uop zsim)/uop real
- Average µop error is 1.3%.

Multithreaded validation

- 22 applications from different benchmark suites
  - 6 from PARSEC, 7 from SPLASH2, 9 from SPEC OMP2001
- Run most workloads at 6 threads
  - Those that need power of 2 threads run with 4 threads
- Measure performance as **1/(time to completion)** and not IPC.

Performance errors

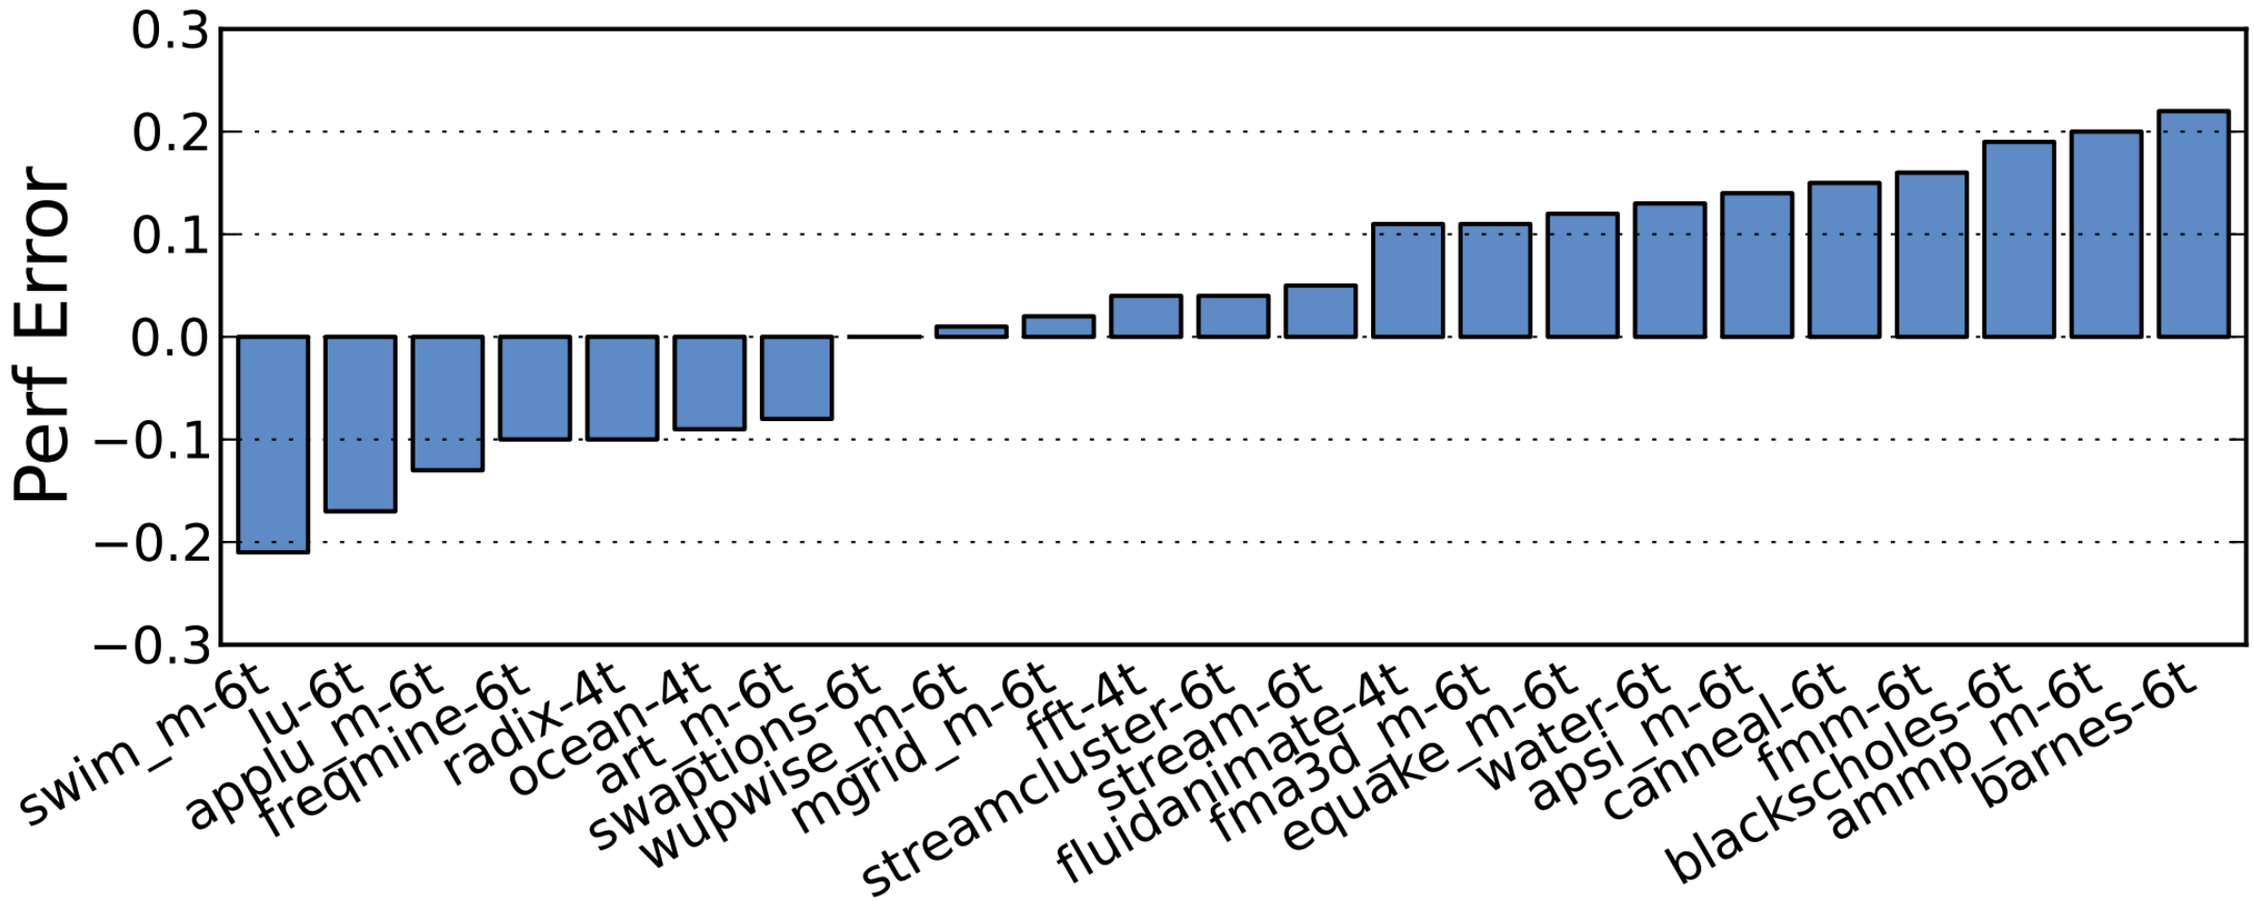

- Average absolute error is 11.2%.
- 10 out of 23 workloads are within 10% error.

Contention models

- Many simulators fail to accurately model bandwidth contention.
- ZSim can accurately simulate a real hardware system by using detailed contention models.

- We study the scalability of STREAM benchmark on real machine and simulation with several timing models.
- STREAM saturates memory bandwidth, scaling sub-linearly.

Bandwidth and Scalability

- Without contention, there is no bandwidth limitation and performance scales linearly.
- Approximate Queueing theory model(MD1) is still quite inaccurate.
- Using event-driven model or DRAMSim2 closely approximates real machine.

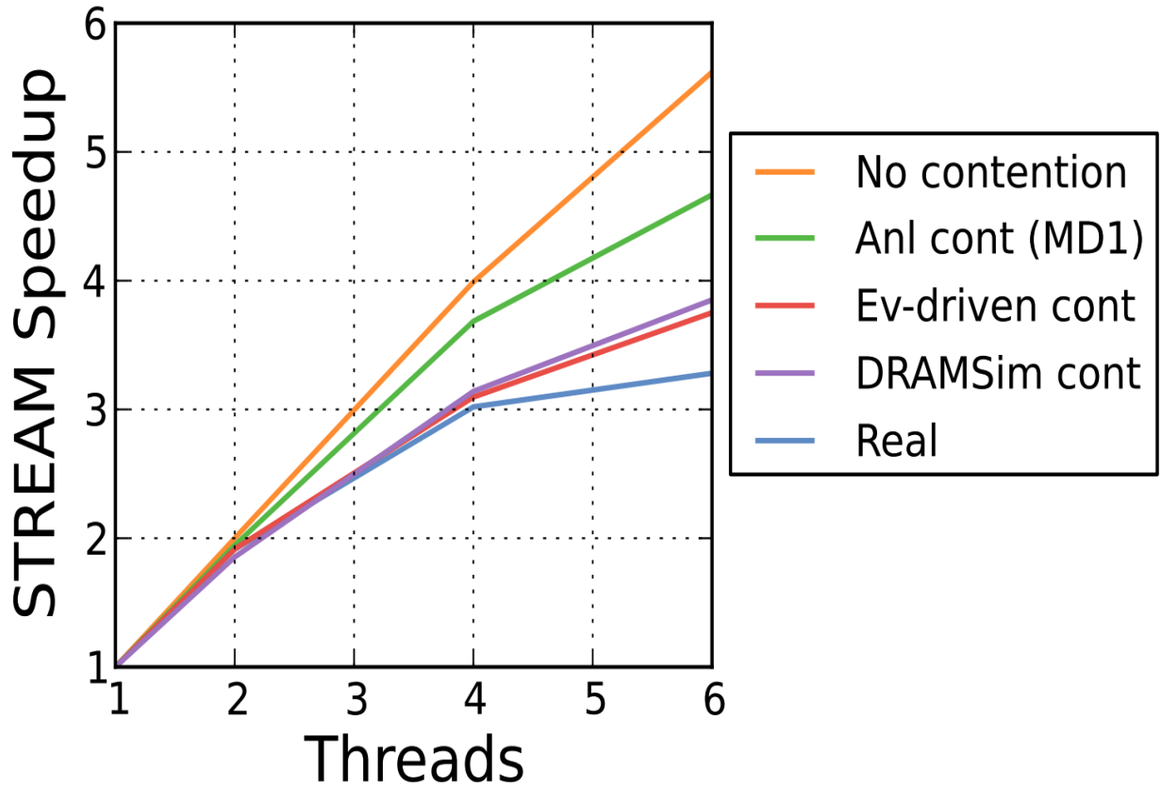

Accuracy vs Speed

- Bound-weave algorithm allows for modeling contention at varying degrees of accuracy.

- Tradeoff between simulation speed and accuracy
  - DRAMSim2 is cycle-accurate – limits ZSim performance to 3 MIPS.
  - Few tens of MIPS with simpler models.

Silvermont validation

- Changed a few parameters to model a silvermont like core.
- Absolute performance error of 20.89%.

- Uop decoding is slightly different.
- Much simpler branch predictor.
- We do not model
  - Differences in backend architecture.
  - Silvermont's prefetcher.

- Possible to reduce the errors by doing more accurate modelling.

Conclusion

- You can trust zsim to be quite accurate, but

'If you are using zsim with workloads or architectures that are significantly different from ours, you should not blindly trust these results'

- Detailed results available at zsim.csail.mit.edu/validation

- Plan to release the complete validation infrastructure in future.

THANK YOU

QUESTIONS?