



# AN EVALUATION OF HIGH-LEVEL MECHANISTIC CORE MODELS

TREVOR E. CARLSON, WIM HEIRMAN,  
STIJN EYERMAN, IBRAHIM HUR, LIEVEN EECKHOUT





[HTTP://WWW.SNIPERSIM.ORG](http://www.snipersim.org)  
MONDAY, JANUARY 19TH 2015  
HIPEAC CONFERENCE 2015, AMSTERDAM

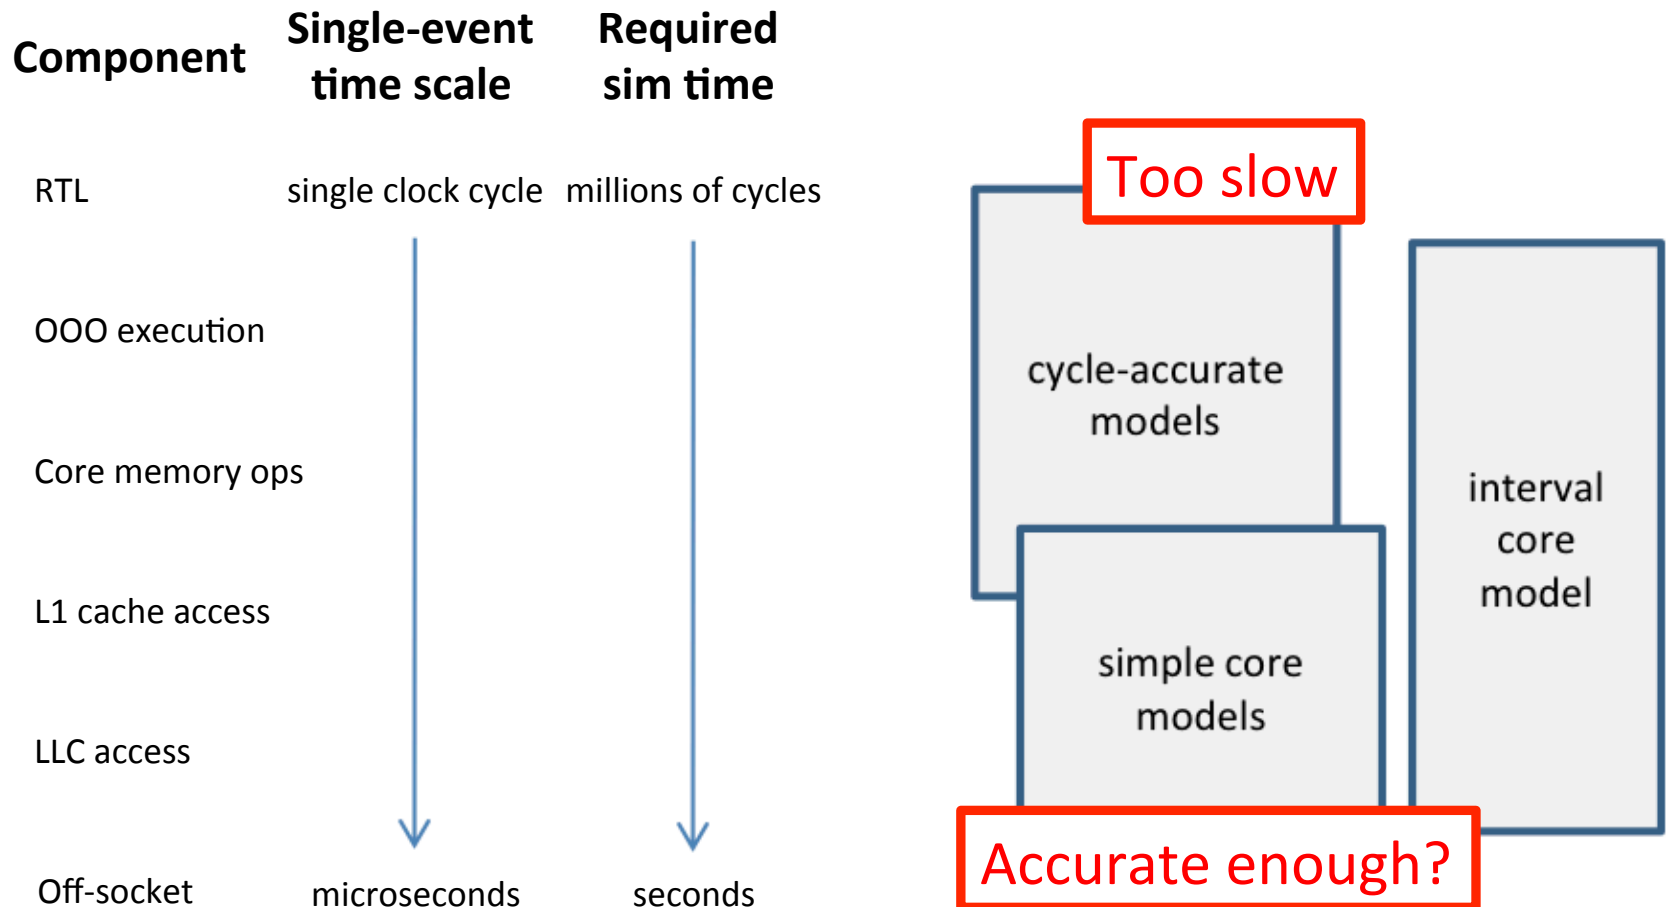
# DESIGN FUTURE HARDWARE AND SOFTWARE

---

- Design the processor of tomorrow
- Optimize next-gen software
- How can we design and evaluate large numbers of design options?
  - Can we do it in a fast way?



# NEEDED DETAIL DEPENDS ON FOCUS



# ONE-IPC MODELING

---

- A simple high-abstraction model, often used in memory hierarchy studies
- Alternative for memory access traces
  - Allows for timing feedback to software (ex. work stealing)
- Drawbacks
  - ILP and MLP is not modeled
  - Memory request rates are not accurate
  - Number of outstanding misses are not correct
    - Underestimation of required queue sizes
  - Does not properly hide latency like an out-of-order core
    - Overestimate runtime improvements

# CORE MODEL CASE STUDY

---


- Do the drawbacks previously listed actually impact accuracy?
- Evaluate this core with 2 cache configurations
  - Which configuration has better performance?

Component	L2 Private Config	L2 Shared Config
Size	256 KiB / core	1 MiB / 4 cores
Associativity	8-way	16-way
Access latency	8 cycles	30 cycles

# CORE MODEL CASE STUDY

The shared L2 reduces misses in almost all cases

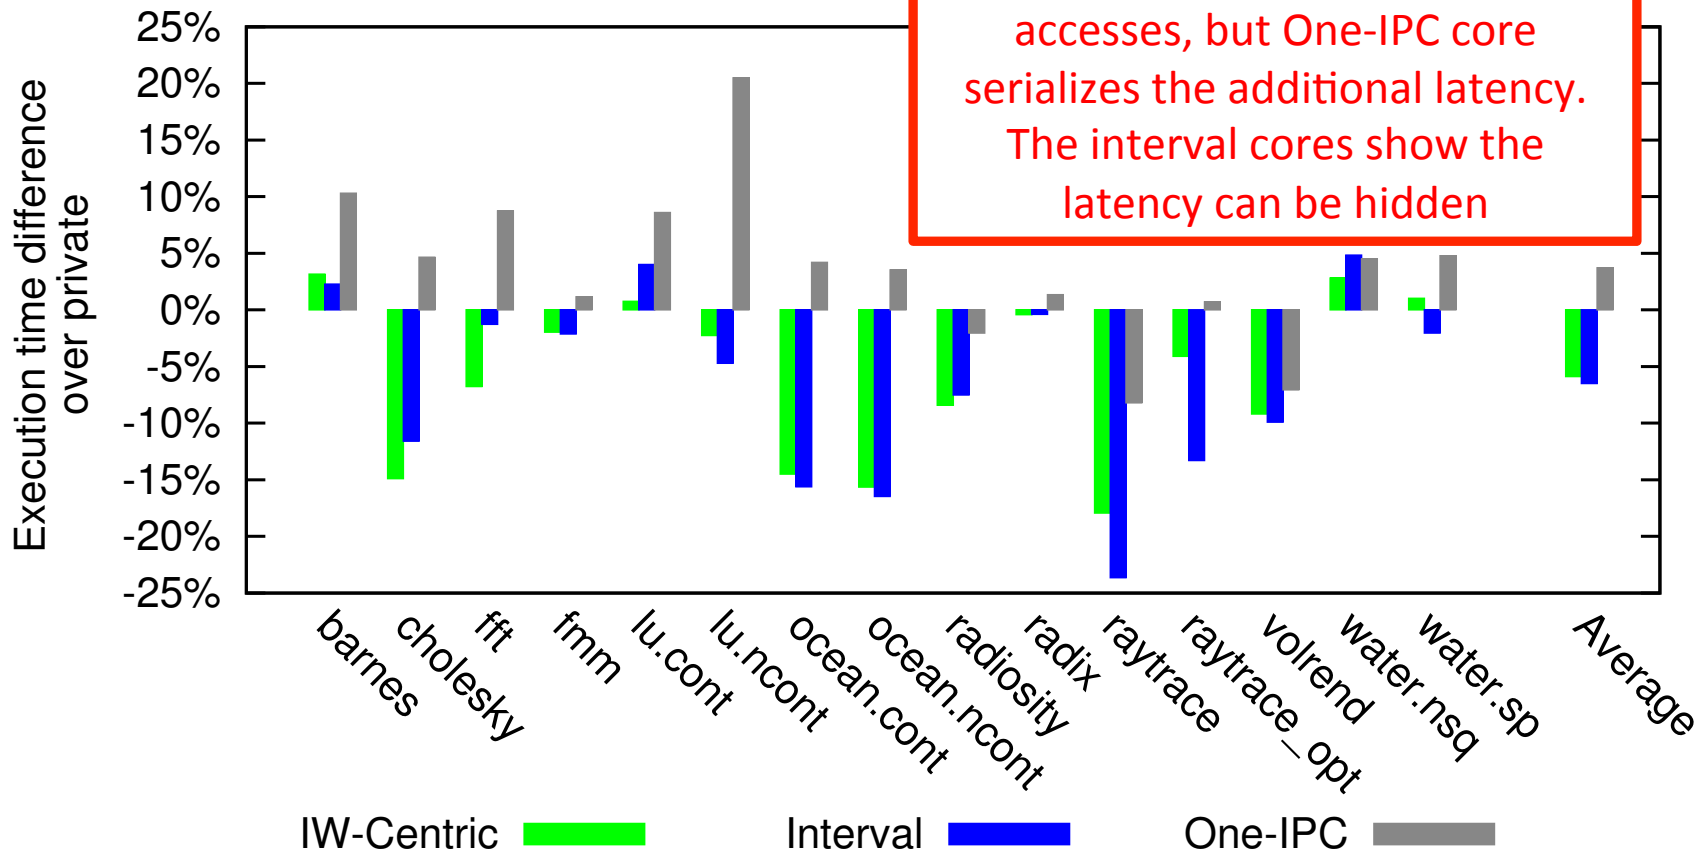
One-IPC and interval core models agree



# CORE MODEL CASE STUDY


L2: 8 cycle latency -> 30 cycles

There is a reduction in DRAM accesses, but One-IPC core serializes the additional latency. The interval cores show the latency can be hidden



# ANALYTICAL MODELING

- Interval model is an approximation of the ideal, balanced out-of-order processor




- Does not currently support multi-core processors



# BRIDGE ANALYTICAL MODELING, SIMULATION

---


- Interval Simulation
  - In-order stream of instructions (from Pin, QEMU)



# BRIDGE ANALYTICAL MODELING, SIMULATION

- Interval Simulation

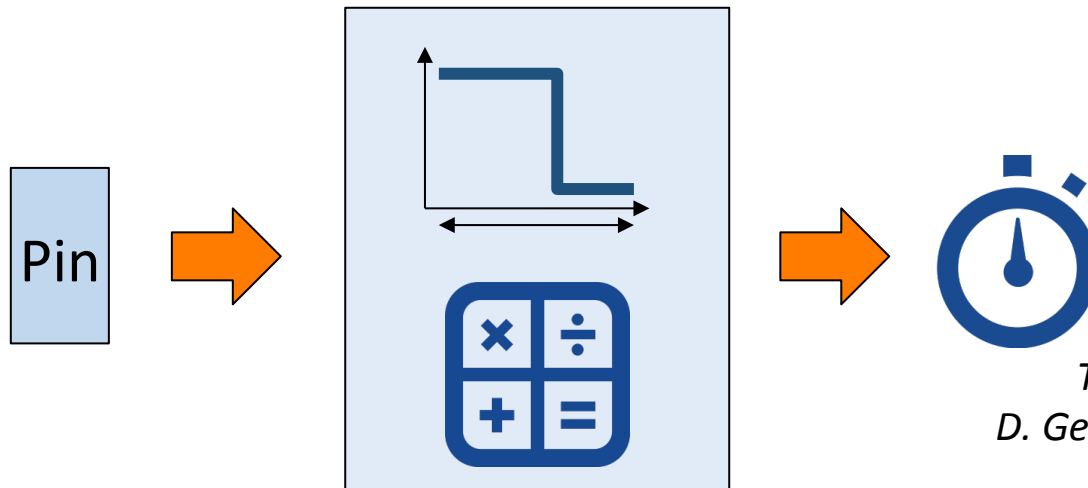
- In-order stream of instructions (from Pin, QEMU)
- Intervals (and their corresponding delays) are formed from miss events (e.g. branch pred., etc.)



# BRIDGE ANALYTICAL MODELING, SIMULATION

- Interval Simulation


- In-order stream of instructions (from Pin, QEMU)
- Intervals (and their corresponding delays) are formed from miss events (e.g. branch pred., etc.)
- Steady-state performance is determined with queueing theory (Little's Law)



*T. E. Carlson et al., SC 2011*  
*D. Genbrugge et al., HPCA 2010*

# INTERVAL SIMULATION


- Little's Law and interval simulation assume that a processor is balanced with respect to the dispatch width
  - Front end (ex. fetch units) are designed to provide instructions at the dispatch width of the processor
  - The execution units can handle any type of instruction at any time
  - Commit width is sufficiently large to prevent stalling
- Predict the progress of the upcoming instructions by using the most recent instructions



# MODERN PROCESSORS ARE NOT IDEAL

- Intel Nehalem issue ports

- One branch port
- One load and one store port
- Limited ports for FP operations
- 3 general-purpose (integer operation) ports
  - Shared with floating point ports



- Nehalem (and other modern cores) are unable to issue the dispatch-width ( $W=4$ ) of instructions of the same type each cycle

- Some applications (fp, memory-intensive, etc.) and micro-benchmarks do

# MAIN INTERVAL ENHANCEMENTS

---

- **Interval simulation issue contention**
  - Takes into account the maximum execution rate possible for each port
  - For example:
    - One issue port per load
    - An application with 100% of instruction as loads
    - Results in a maximum IPC of 1.0
    - Traditional interval simulation would not take this into account and only use dependencies to derive performance
- **Instruction-Window Centric (IW-Centric)**
  - Replaces Little's Law calculations with detailed cycle-level wake-up and issue logic
  - Slower code as the expense of higher fidelity

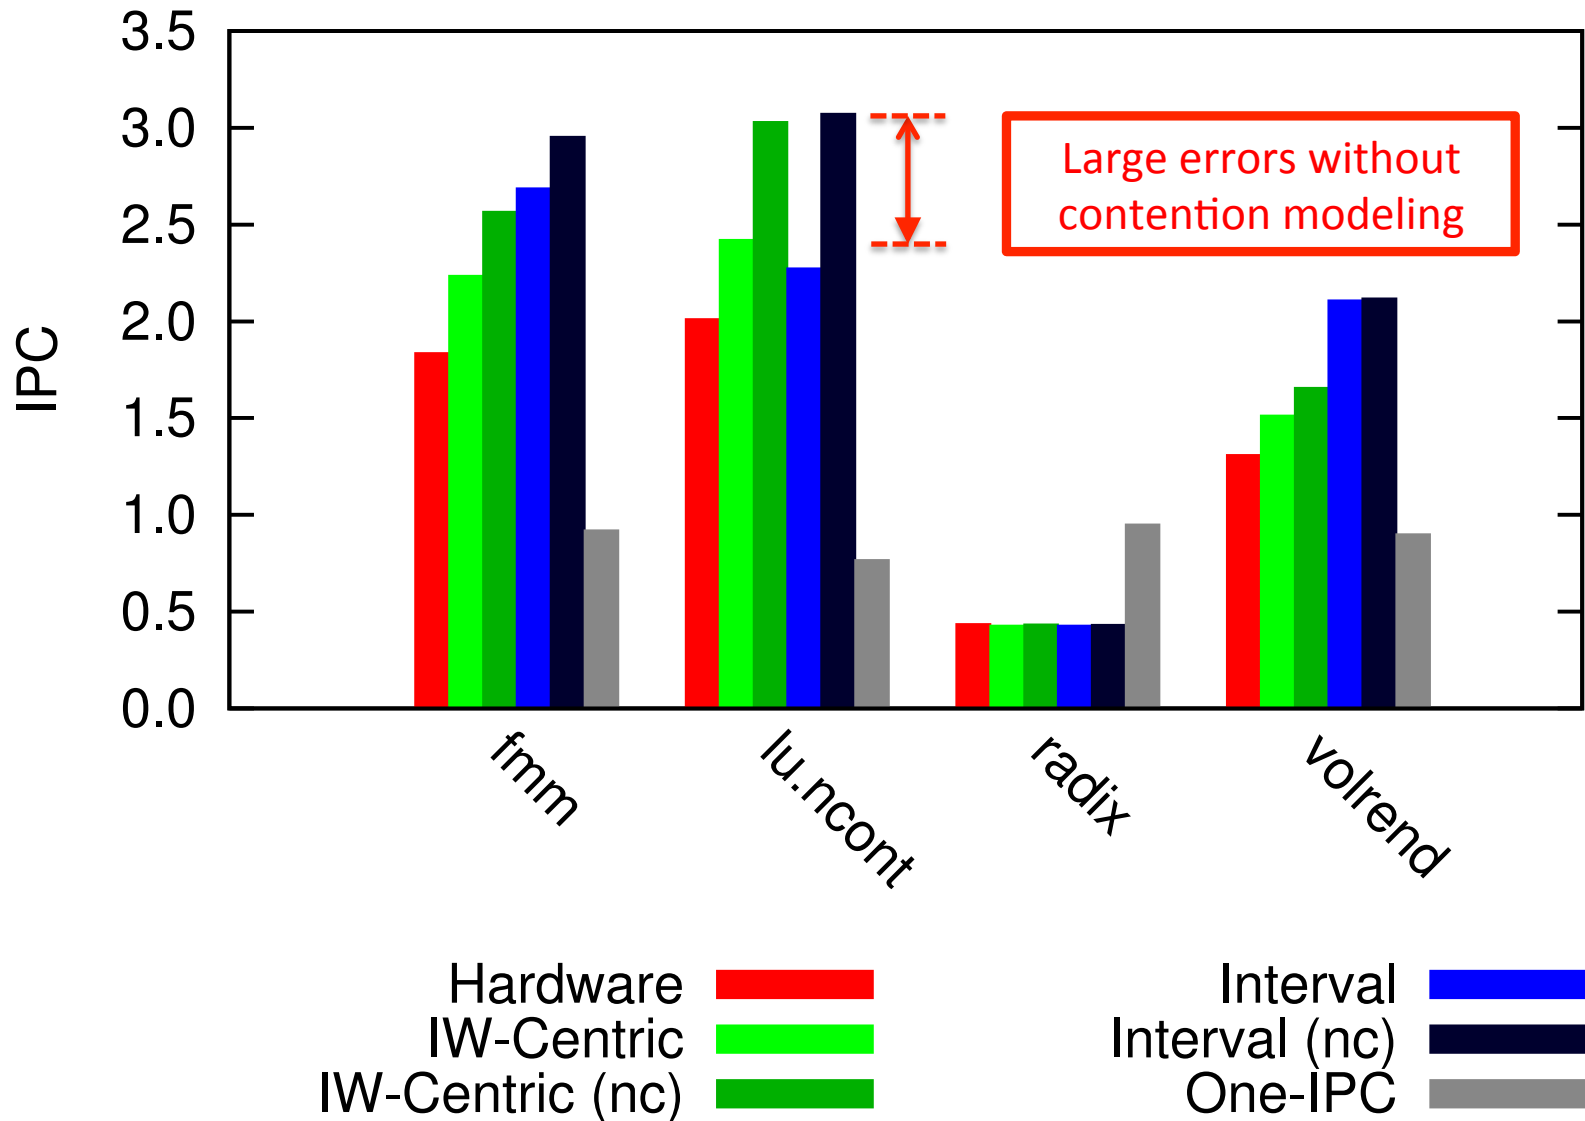
# EXPERIMENTAL SETUP

---

- Sniper Multi-Core Simulator, version 6.0
- SPLASH-2 Benchmark Suite
- Modeled the Intel Xeon X5550 (Nehalem)
- Microarchitectural Configuration
  - 1 and 2 sockets, 4 cores per socket
  - 2.66 GHz, 4-way dispatch, 128-entry ROB
  - Dothan branch predictor
  - L1-I        32 KB, 4 way, 4 cycle access time
  - L1-D        32 KB, 8 way, 4 cycle access time
  - L2           256 KB, 8 way, 8 cycle access time
  - L3           8MB per 4 cores, 16 way, 30 cycle access time
  - DRAM       65 ns access time, 8GB/s per socket
  - Inter-processor Bus QPI, 12.8 GB/s per direction

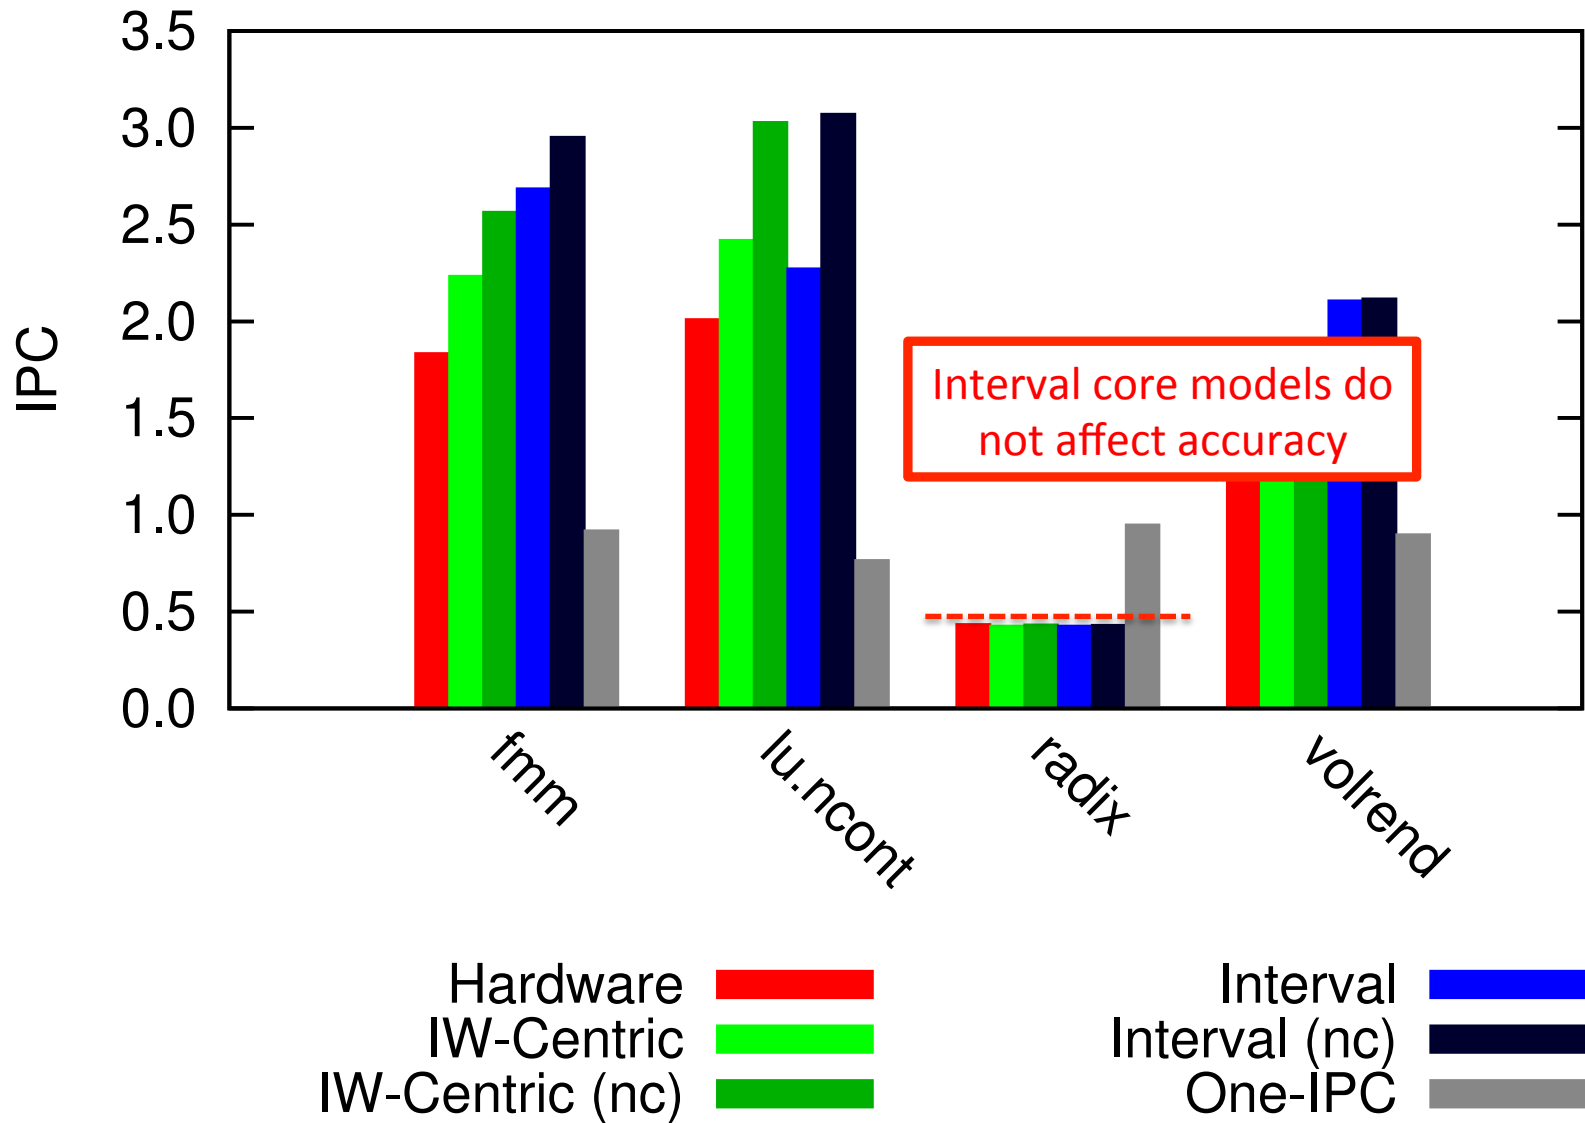


# SIMULATION ACCURACY

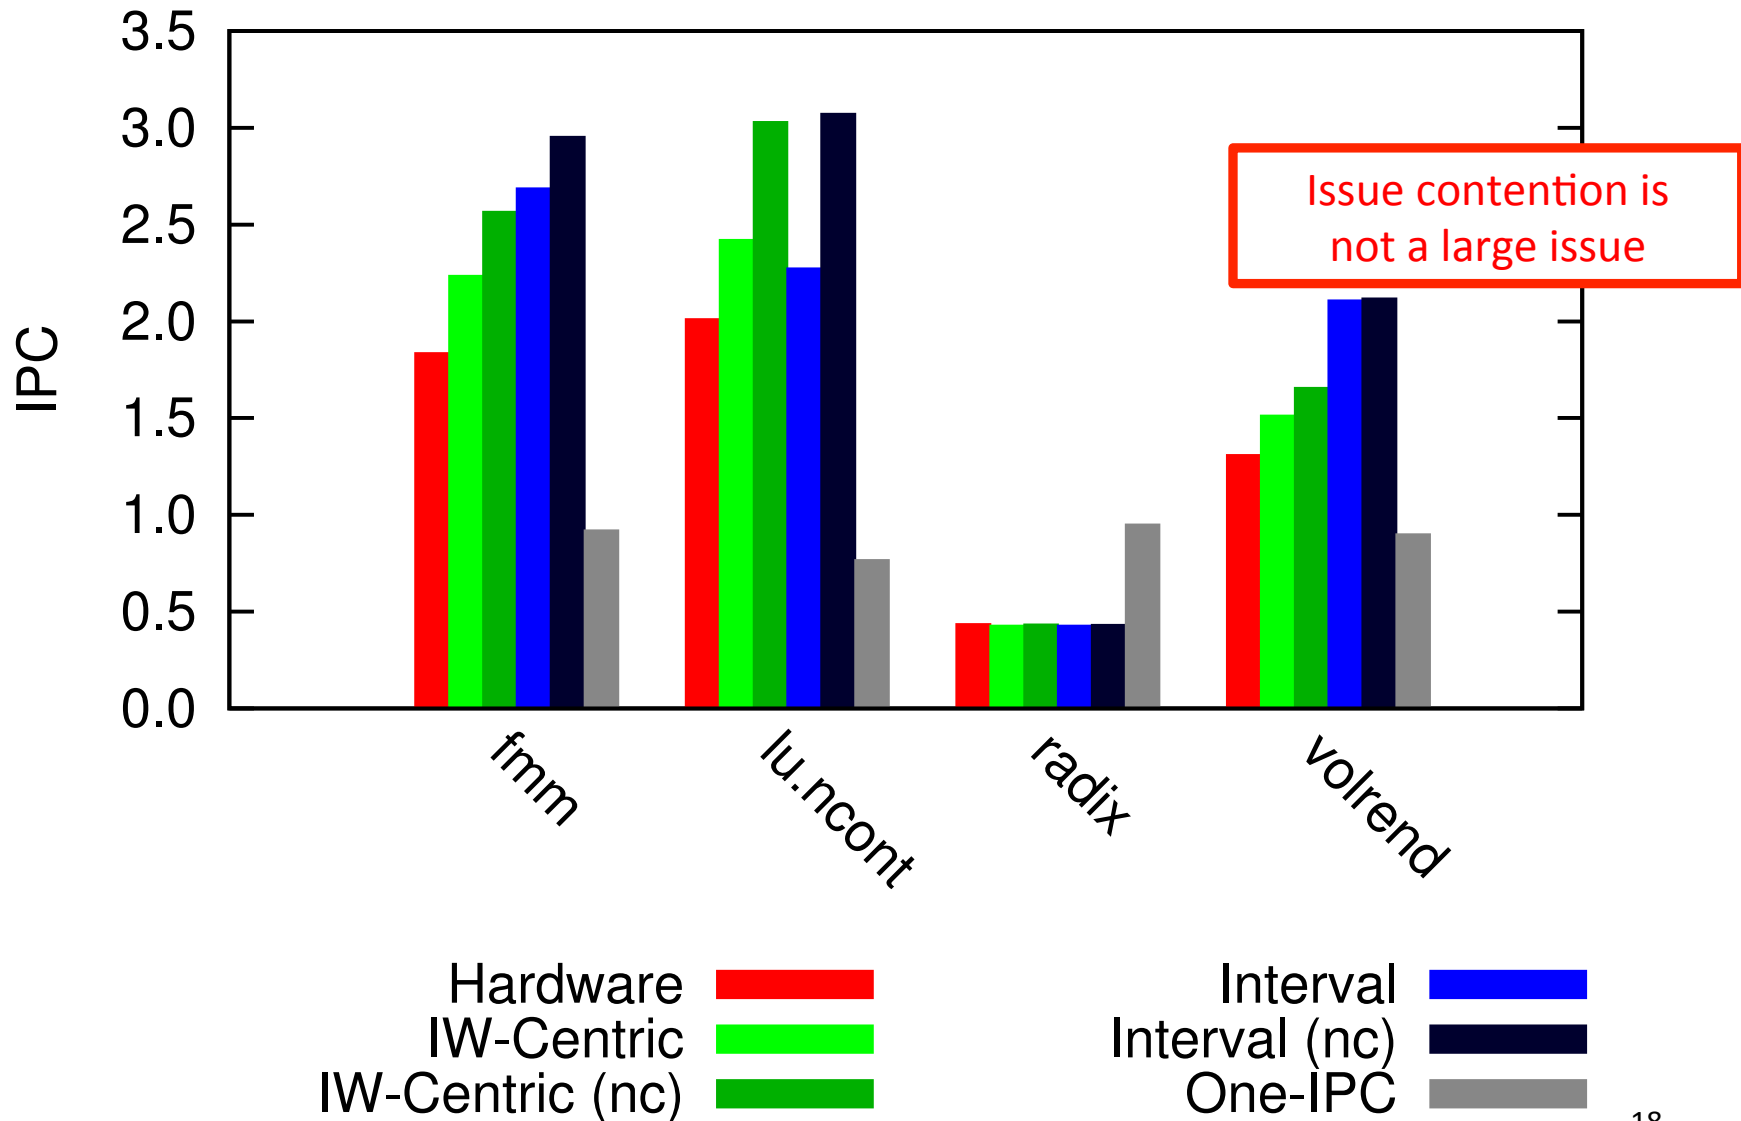




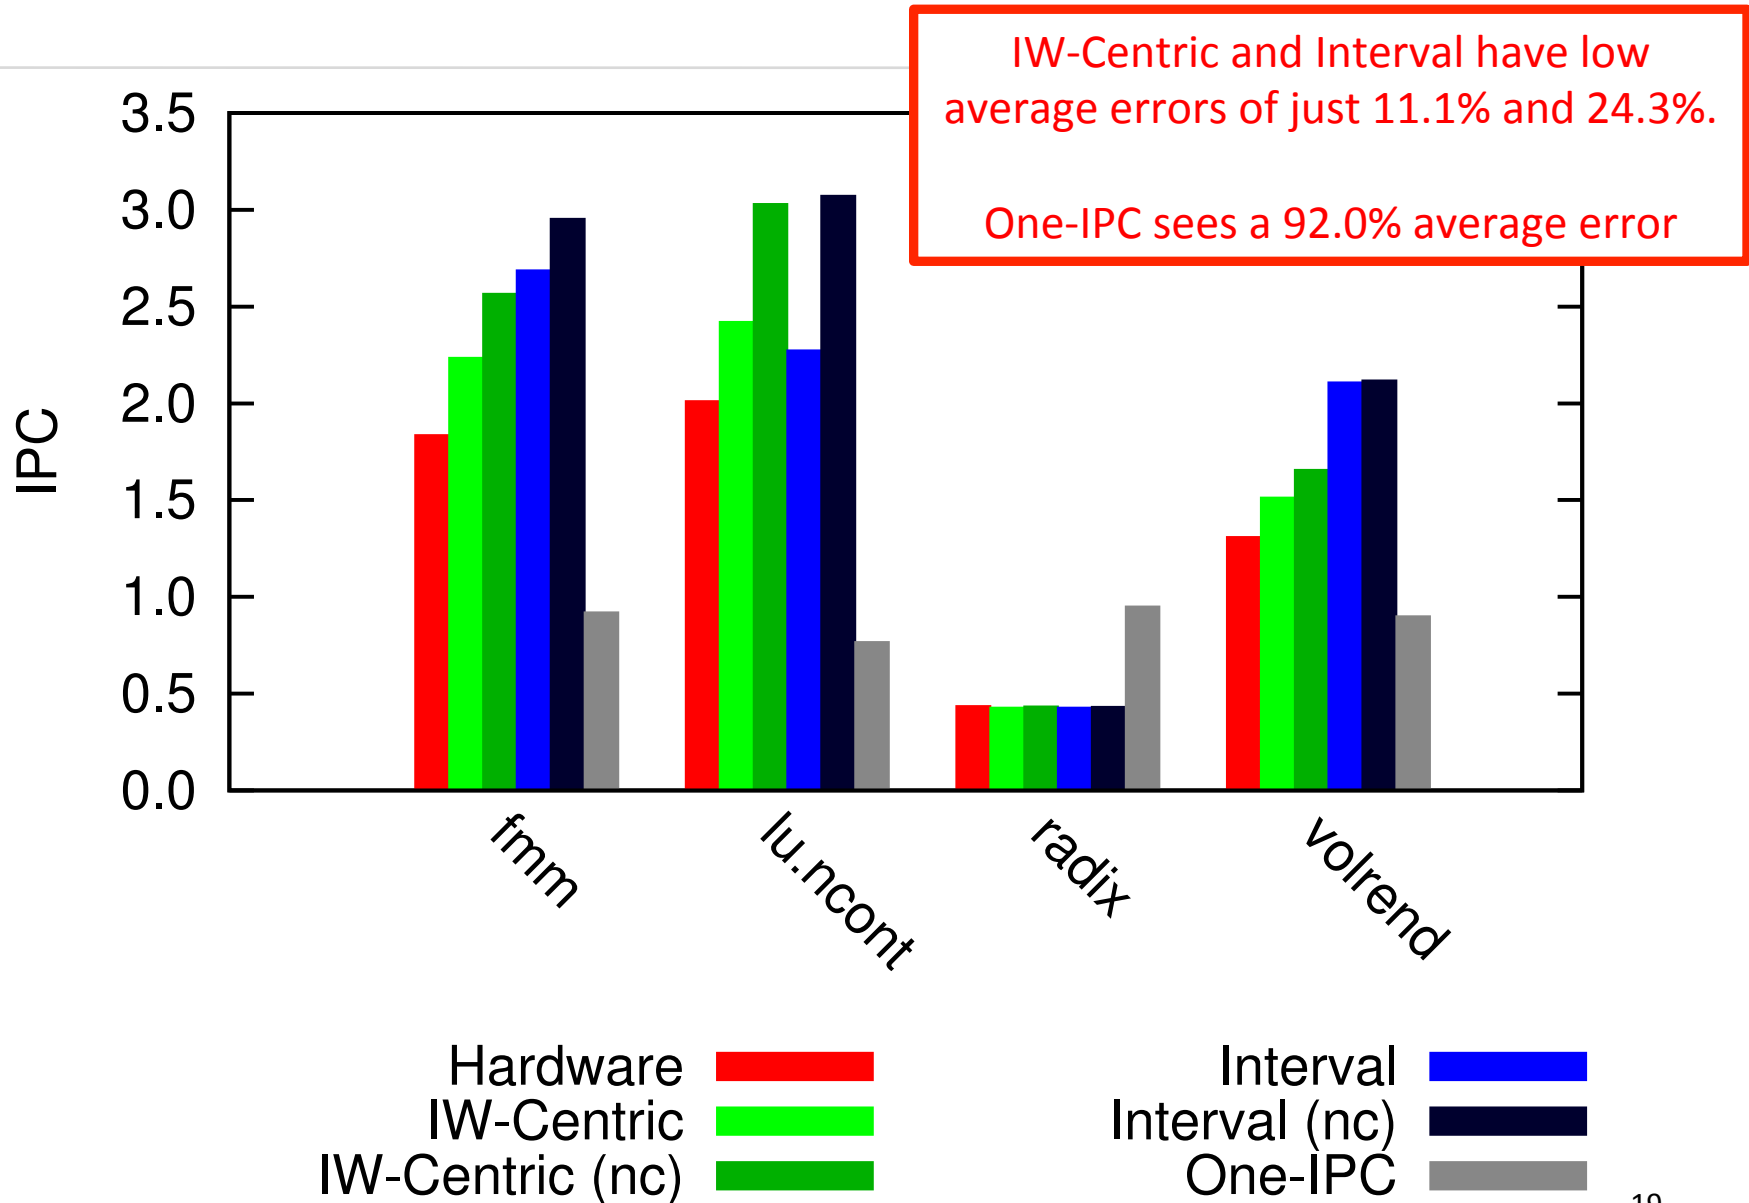
# SIMULATION ACCURACY



# SIMULATION ACCURACY



# SIMULATION ACCURACY




# PERFORMANCE CONCLUSIONS

---


- Interval simulation model has been improved to more accurately reflect limitations of modern processors
- The instruction-window centric model allows us to better understand interval simulation error
  - Interval simulation works with averages over the length of the ROB
  - More precise handling of the dependencies in the ROB provides added accuracy while maintaining performance

# RELATIVE SCALING RESULTS

raytrace\_opt



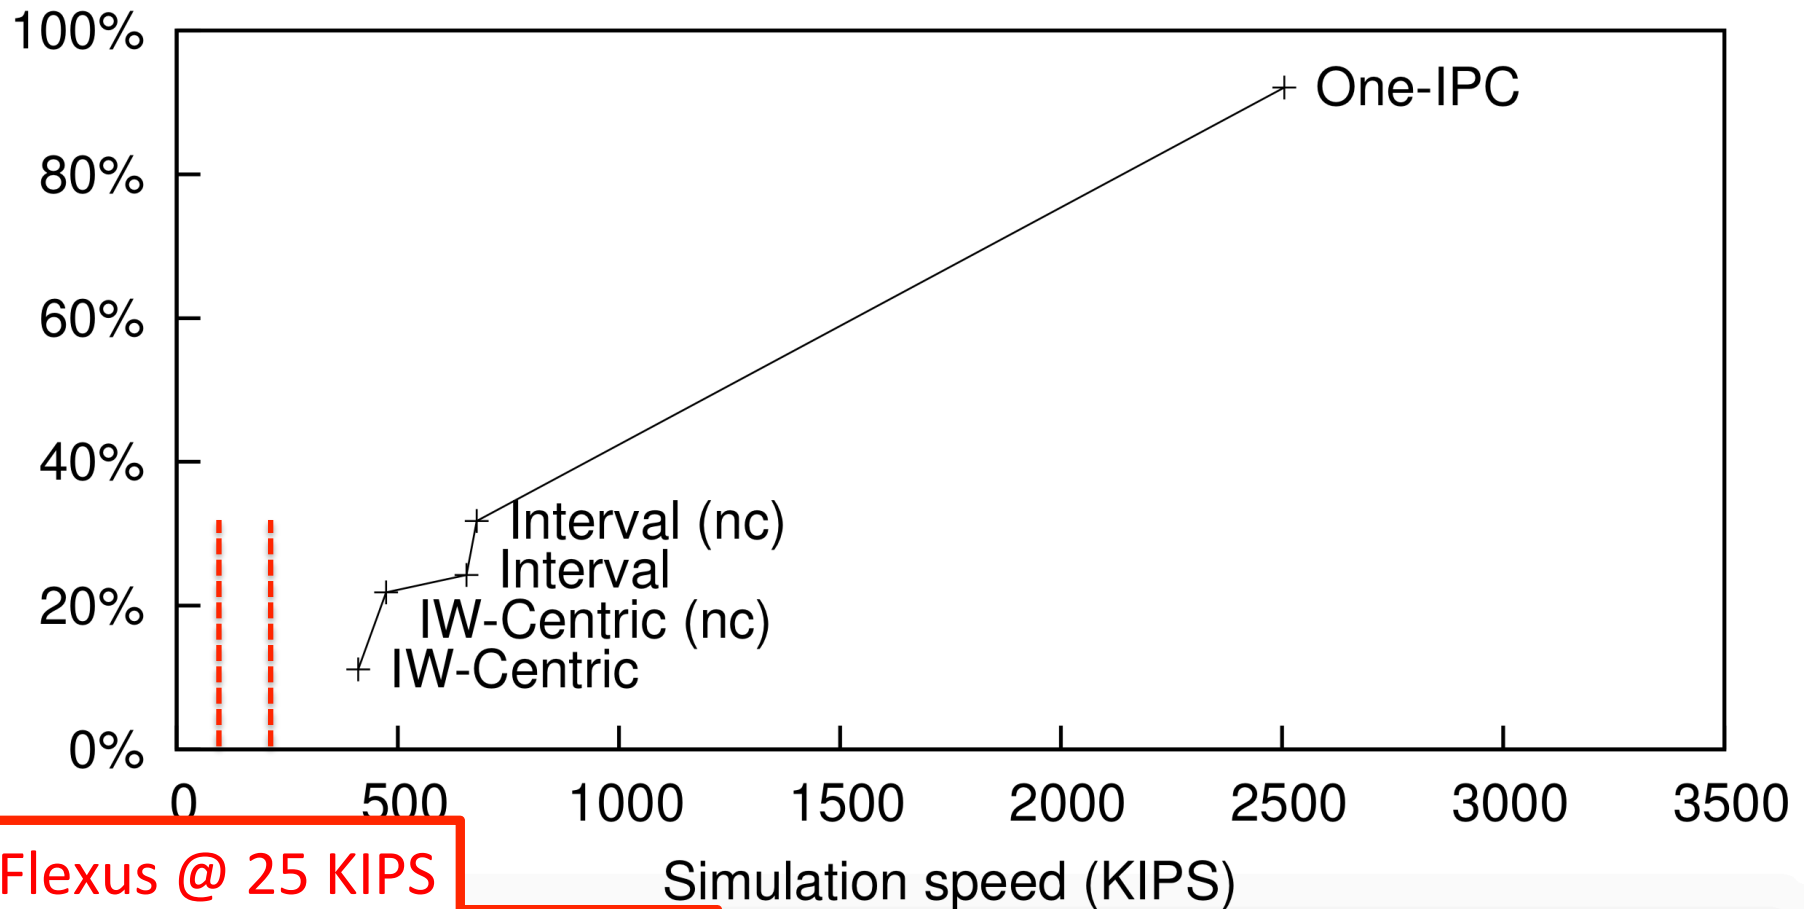
cholesky



One-IPC core models can show large errors, even for relative scaling results

IW-Centric —+— Interval ---\*--- One-IPC ---o---  
 IW-Centric (nc) ---x--- Interval (nc) ---□--- Hardware ---■---

# SIMULATION SPEED COMPARISON



Flexus @ 25 KIPS

gem5 @ 200 KIPS

# ADDITIONAL RESULTS IN THE PAPER

---

- Enhancements to interval simulation
  - Improved modeling of overlapping memory accesses
  - Improve modeling of the front-end miss refill rate
- Additional relative scaling results
- Detailed issue contention example
- Simulation speed and scaling comparison

# SUMMARY

---

- Enhance interval simulation to support issue contention to improve accuracy with very little slowdown
  - 24% avg. absolute error vs. hardware
- Developed a new core model to bridge the performance/accuracy gap with cycle-level simulation
  - 11% avg absolute error vs. hardware
  - Allows for future core designs (in-order, SMT) without the need for analytical models
- Show results that imply that caution must be taken when using simple (One-IPC-style) core models for absolute or even relative studies
- Recently released in Sniper 6.0
  - Interval simulation enhancements
  - IW-Centric core model







# AN EVALUATION OF HIGH-LEVEL MECHANISTIC CORE MODELS

TREVOR E. CARLSON, WIM HEIRMAN,  
STIJN EYERMAN, IBRAHIM HUR, LIEVEN EECKHOUT



[HTTP://WWW.SNIPERSIM.ORG](http://www.snipersim.org)  
MONDAY, JANUARY 19TH 2015  
HIPEAC CONFERENCE 2015, AMSTERDAM