QFlex 3.0: Fast and Accurate ARM Server Simulation

Shanqing Lin, Ali Ansari, Ayan Chakraborty, Bugra Eryilmaz, Yuanlong Li

Mohammad Alian, Babak Falsafi

ARM General-Purpose Computing Workshop, ISCA'25

PARSA
PARALLEL SYSTEMS ARCHITECTURE LAB

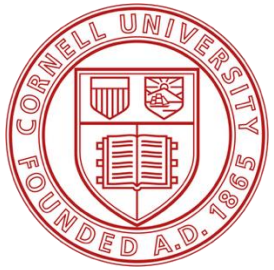

EPFL

The Simulation Wall

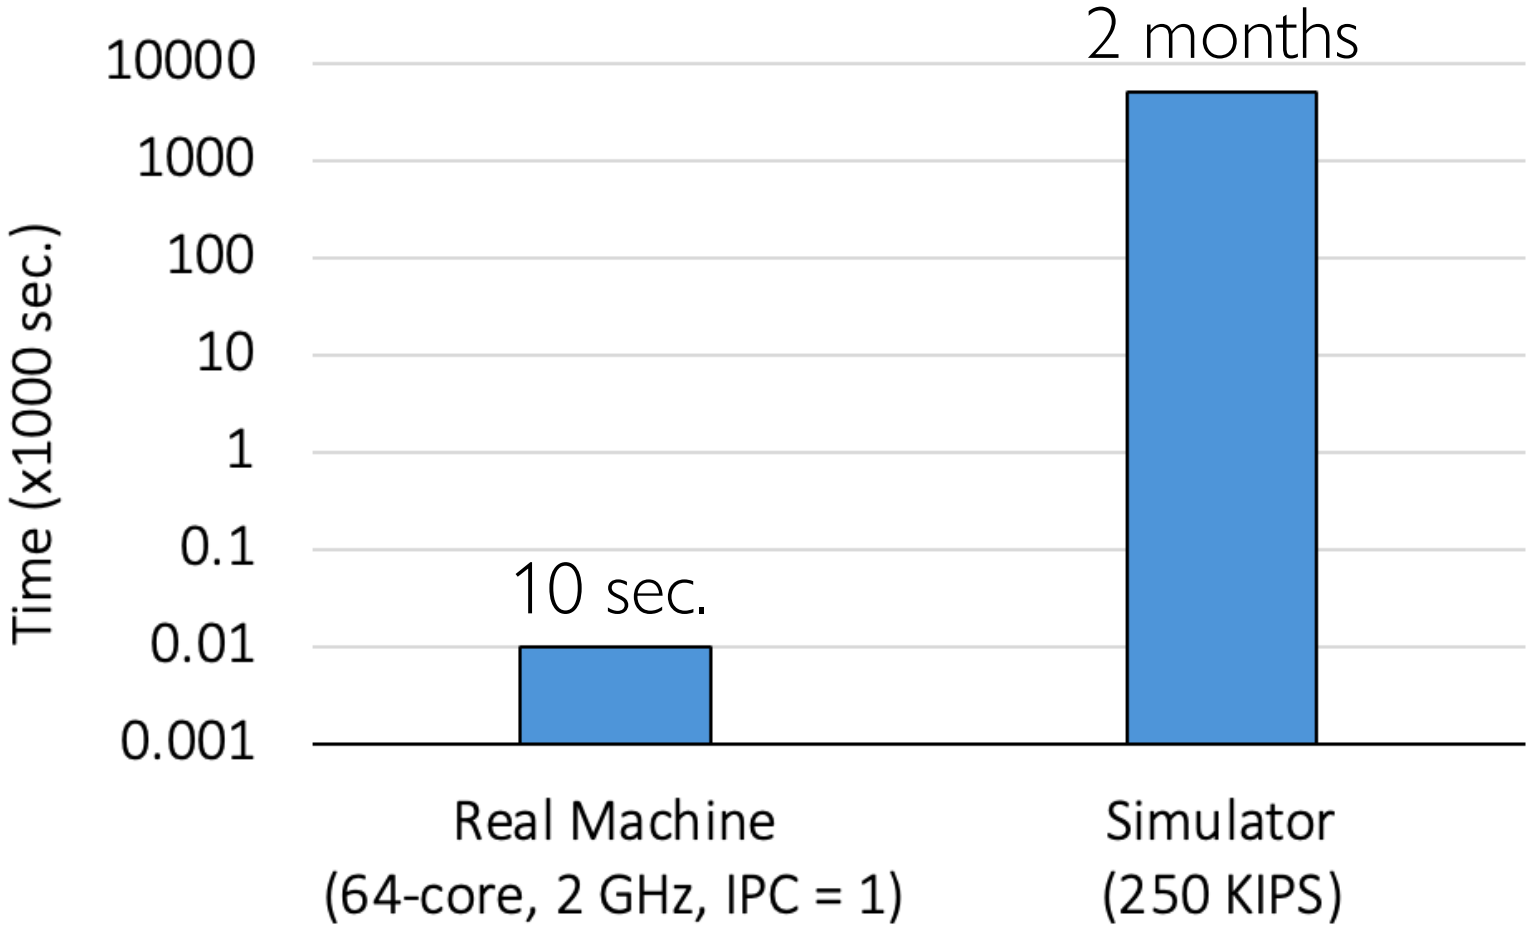

Simulating only a few seconds takes several months

Our Contribution

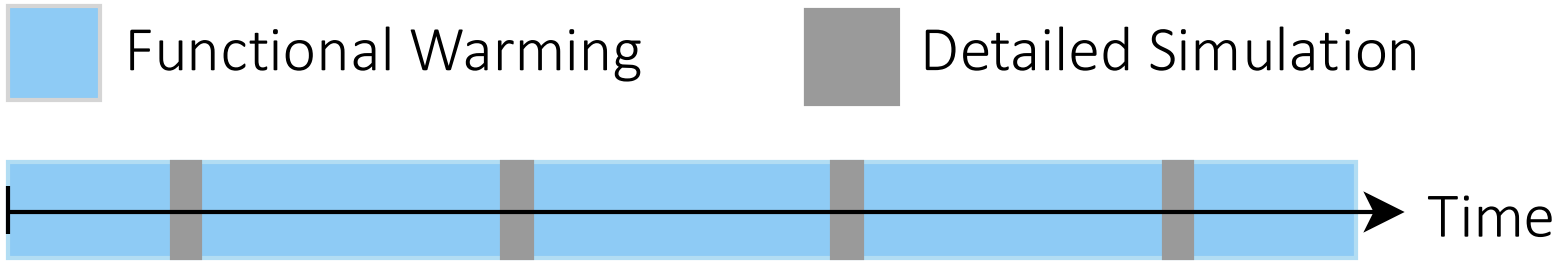

QFlex 3.0: a full-system cycle-accurate ARM server simulator

- Reduces timing requirements by ~500x
- Bounds error with statistical guarantees
- Runs at ~10 MIPS
- 10 s takes roughly 10 hours

QFlex 3.0 provides a fast and accurate server simulator

Outline

PARALLEL SYSTEMS ARCHITECTURE LAB

- ~~Overview~~
- Server simulation requirements
- QFlex 3.0
- Evaluation
- Future directions
- Conclusions

Variable Performance in Server Workloads

PARALLEL SYSTEMS ARCHITECTURE LAB

Web Serving from CloudSuite 4.0, running on Neoverse N1 @ 3 GHz

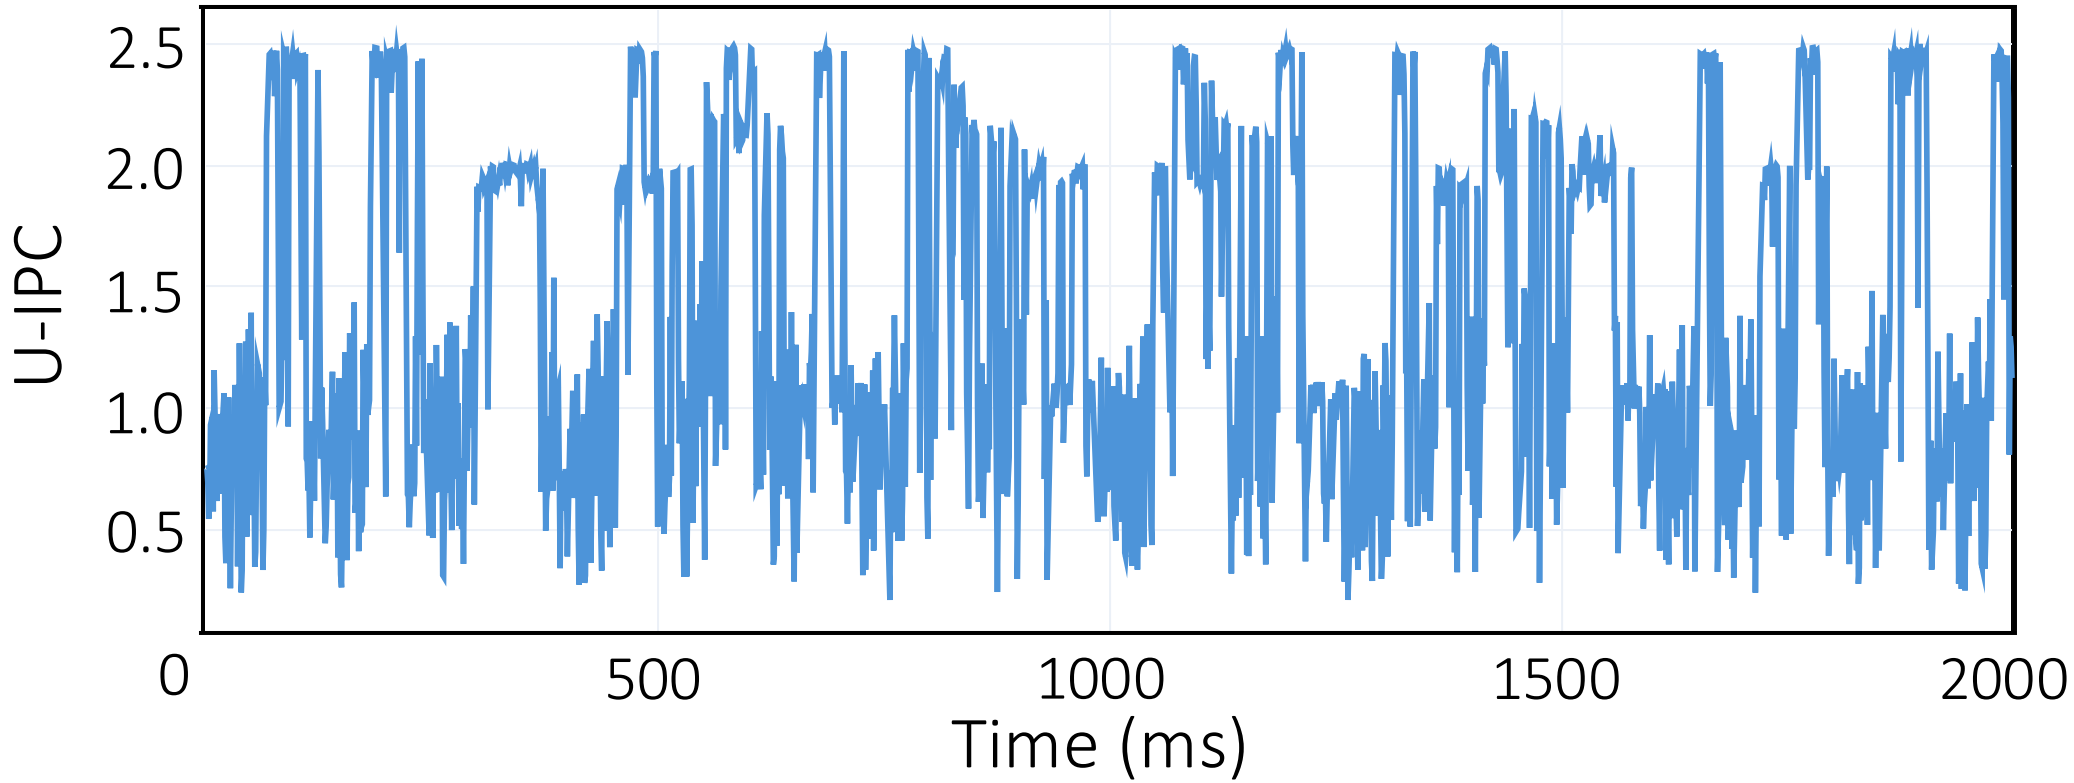

Workload characteristics change within a few milliseconds

Problem with Abbreviated Measurement

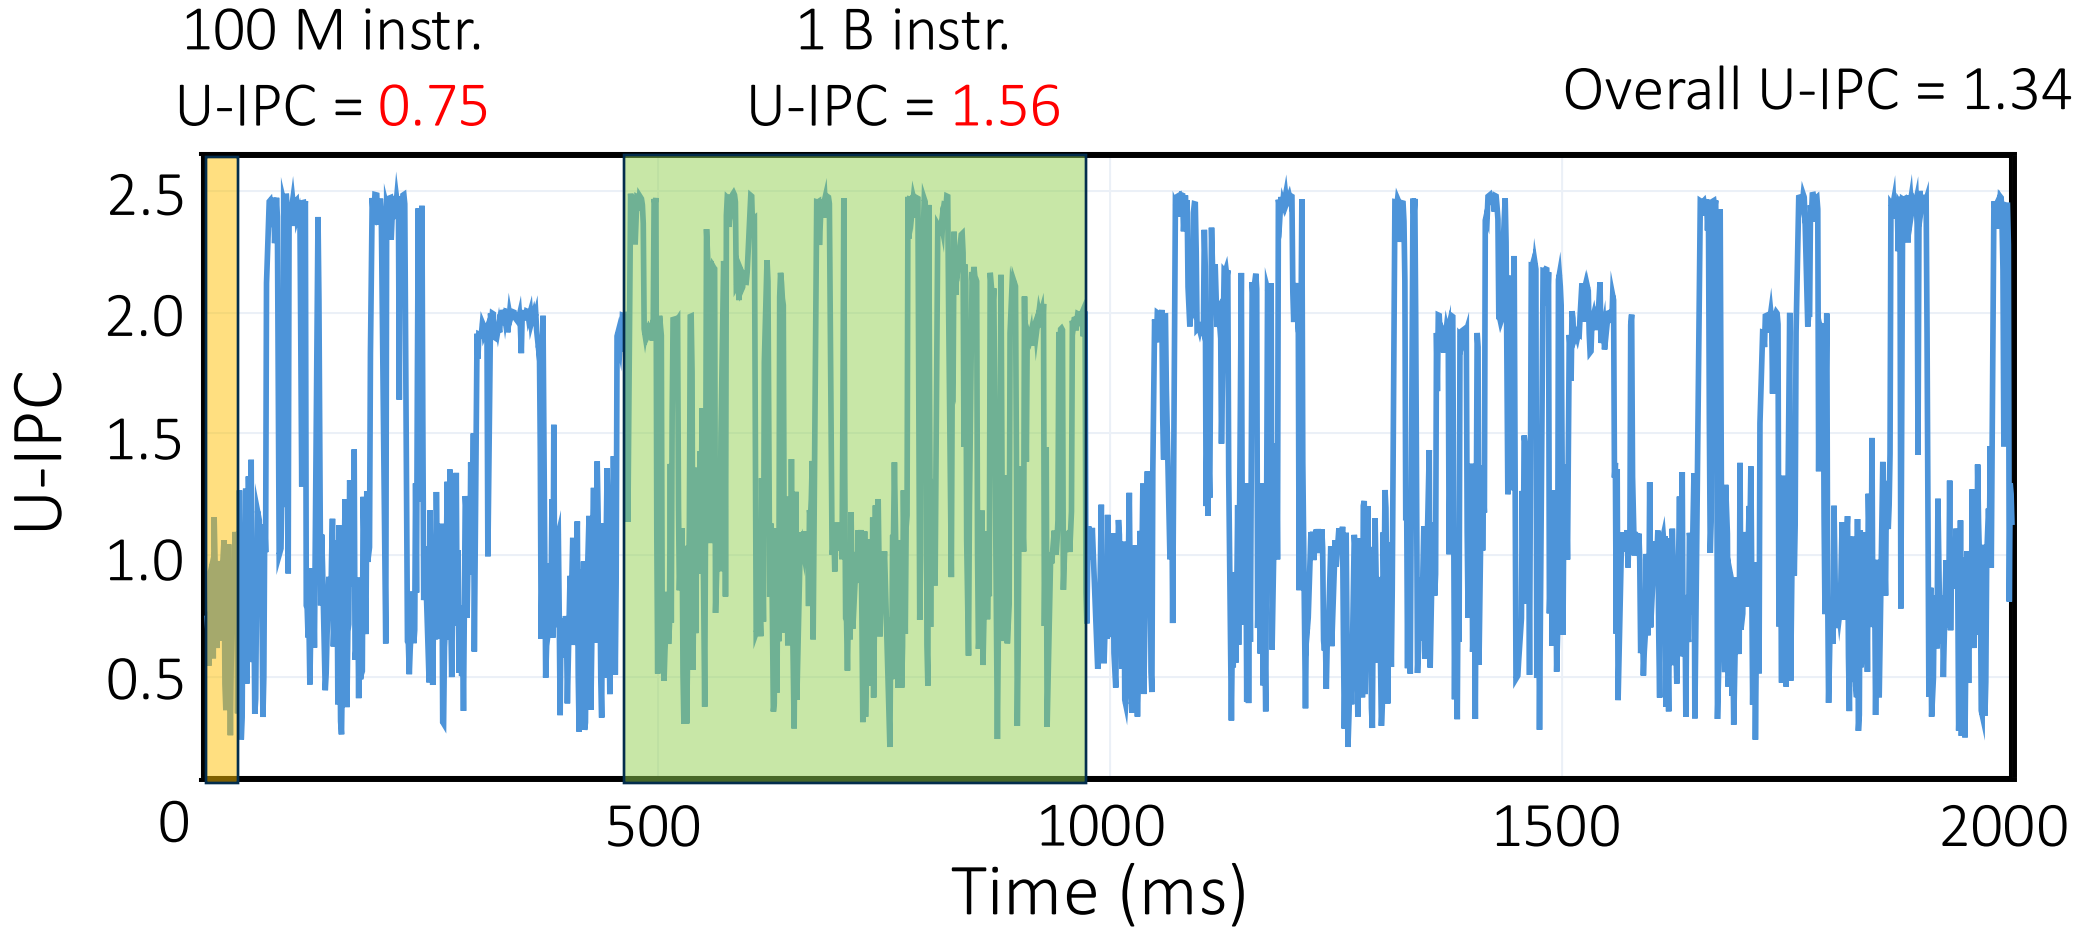

Results of current practices are not representative

Statistical Sampling

PARALLEL SYSTEMS ARCHITECTURE LAB

Report average of U-IPCs across uniformly distributed sampling units [ISCA'03]

Average U-IPC of sampling units = 1.37, workload's U-IPC = 1.34

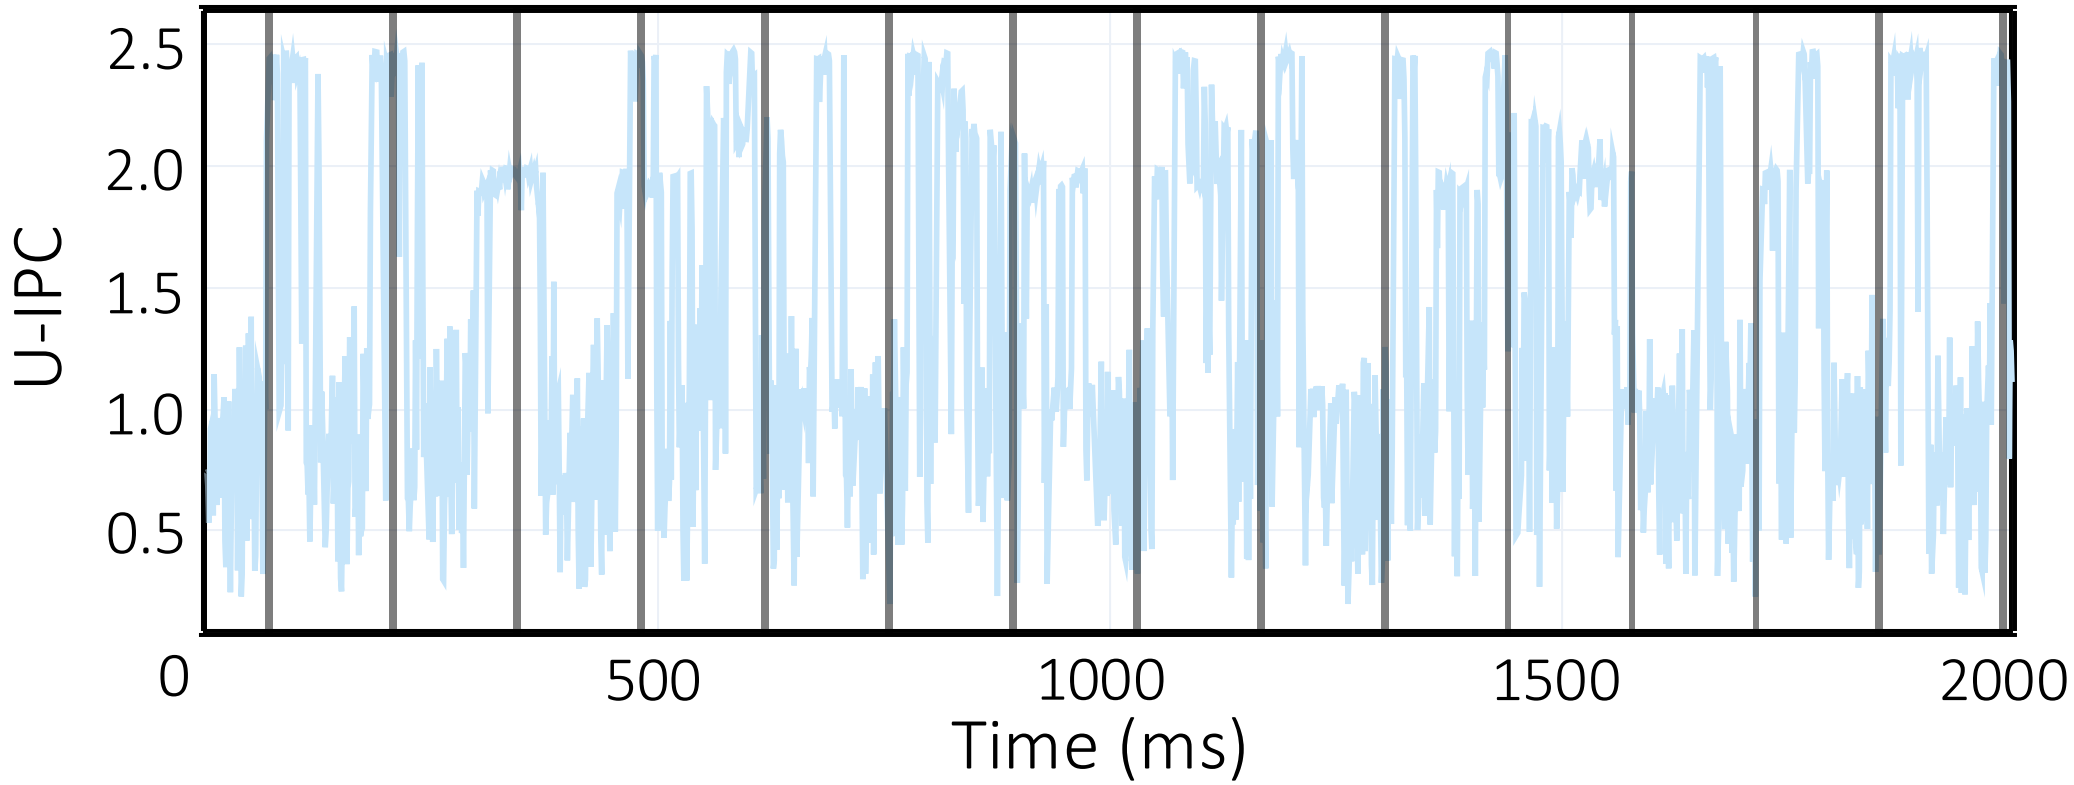

Statistical sampling provides accurate result

Initial State of Sampling Units

PARALLEL SYSTEMS ARCHITECTURE LAB

- Taxonomy
  - Architectural state (registers, memory, devices) created by ISA emulation
  - Short-term µArch state (queues, buffers): warmed by detailed simulation
  - Long-term µArch state (caches, TLBs, branch predictors)
- Problem: how to warm up long-term µArch state?
  - Community: by longer detailed simulation → no guarantee for completeness
  - Our solution [ISCA'03]: by continuous functional warming

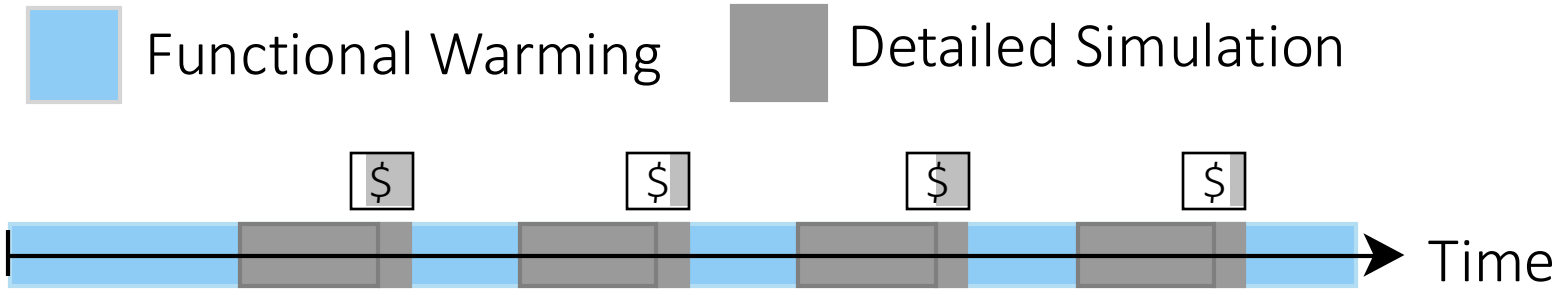

Continuous functional warming guarantees complete warmup

Full-System Simulation

PARALLEL SYSTEMS ARCHITECTURE LAB

Server workloads require OS support for devices and multitasking

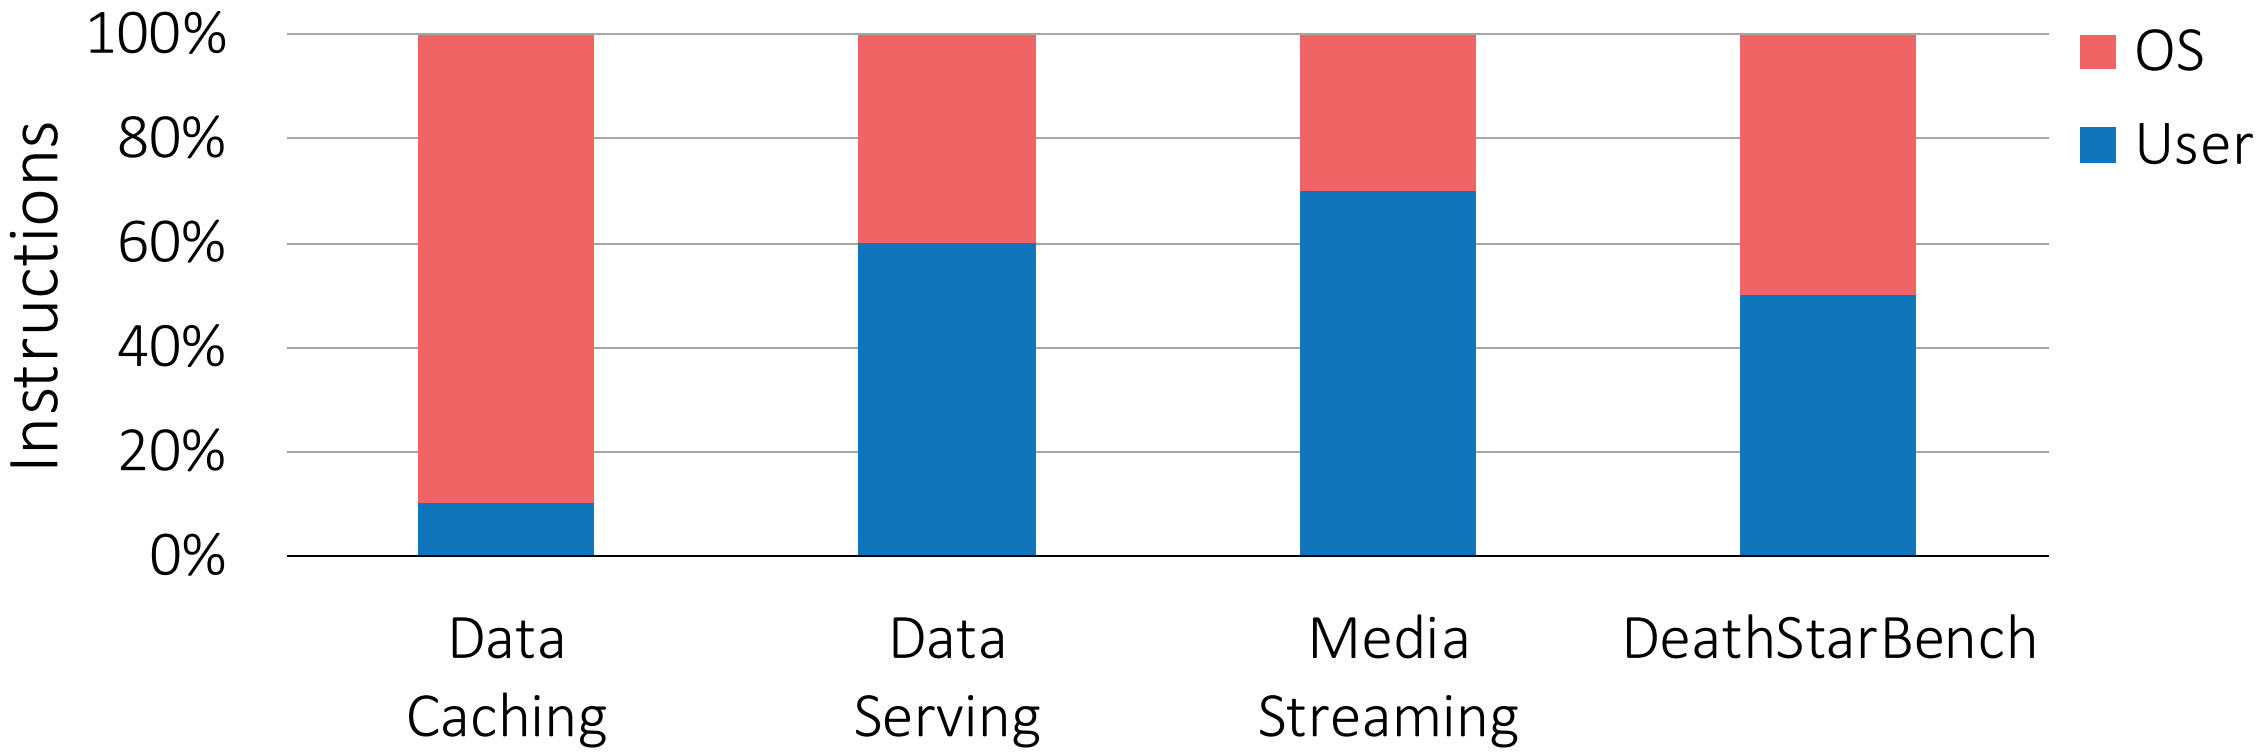

Full-system simulation is necessary for server simulation

Existing Tools: Either Fast or Full-System

PARALLEL SYSTEMS ARCHITECTURE LAB

Simulators	Single-core Speed	Full-system	ARM Support
Sniper	F: 10s MIPS D: 1 MIPS	No	Yes
gem5	0.1s MIPS	Yes	Yes
ZSim	10s MIPS	No	No
Goal	10s MIPS	Yes	Yes

No practical simulator to design ARM servers

QFlex 3.0

PARALLEL SYSTEMS ARCHITECTURE LAB

- Added statistical sampling support to QFlex 2.0
- Rewritten functional simulator in Rust
  - Enhanced simulation speed and simplicity

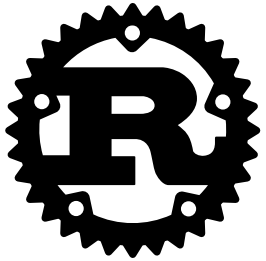

- Expanded ARMv8 support

arm

- Updated µArch components (prefetching, branch predictors, etc.)
- Pre-built CloudSuite 4.0 & DeathStarBench images

CloudSuite

Statistical Sampling Support

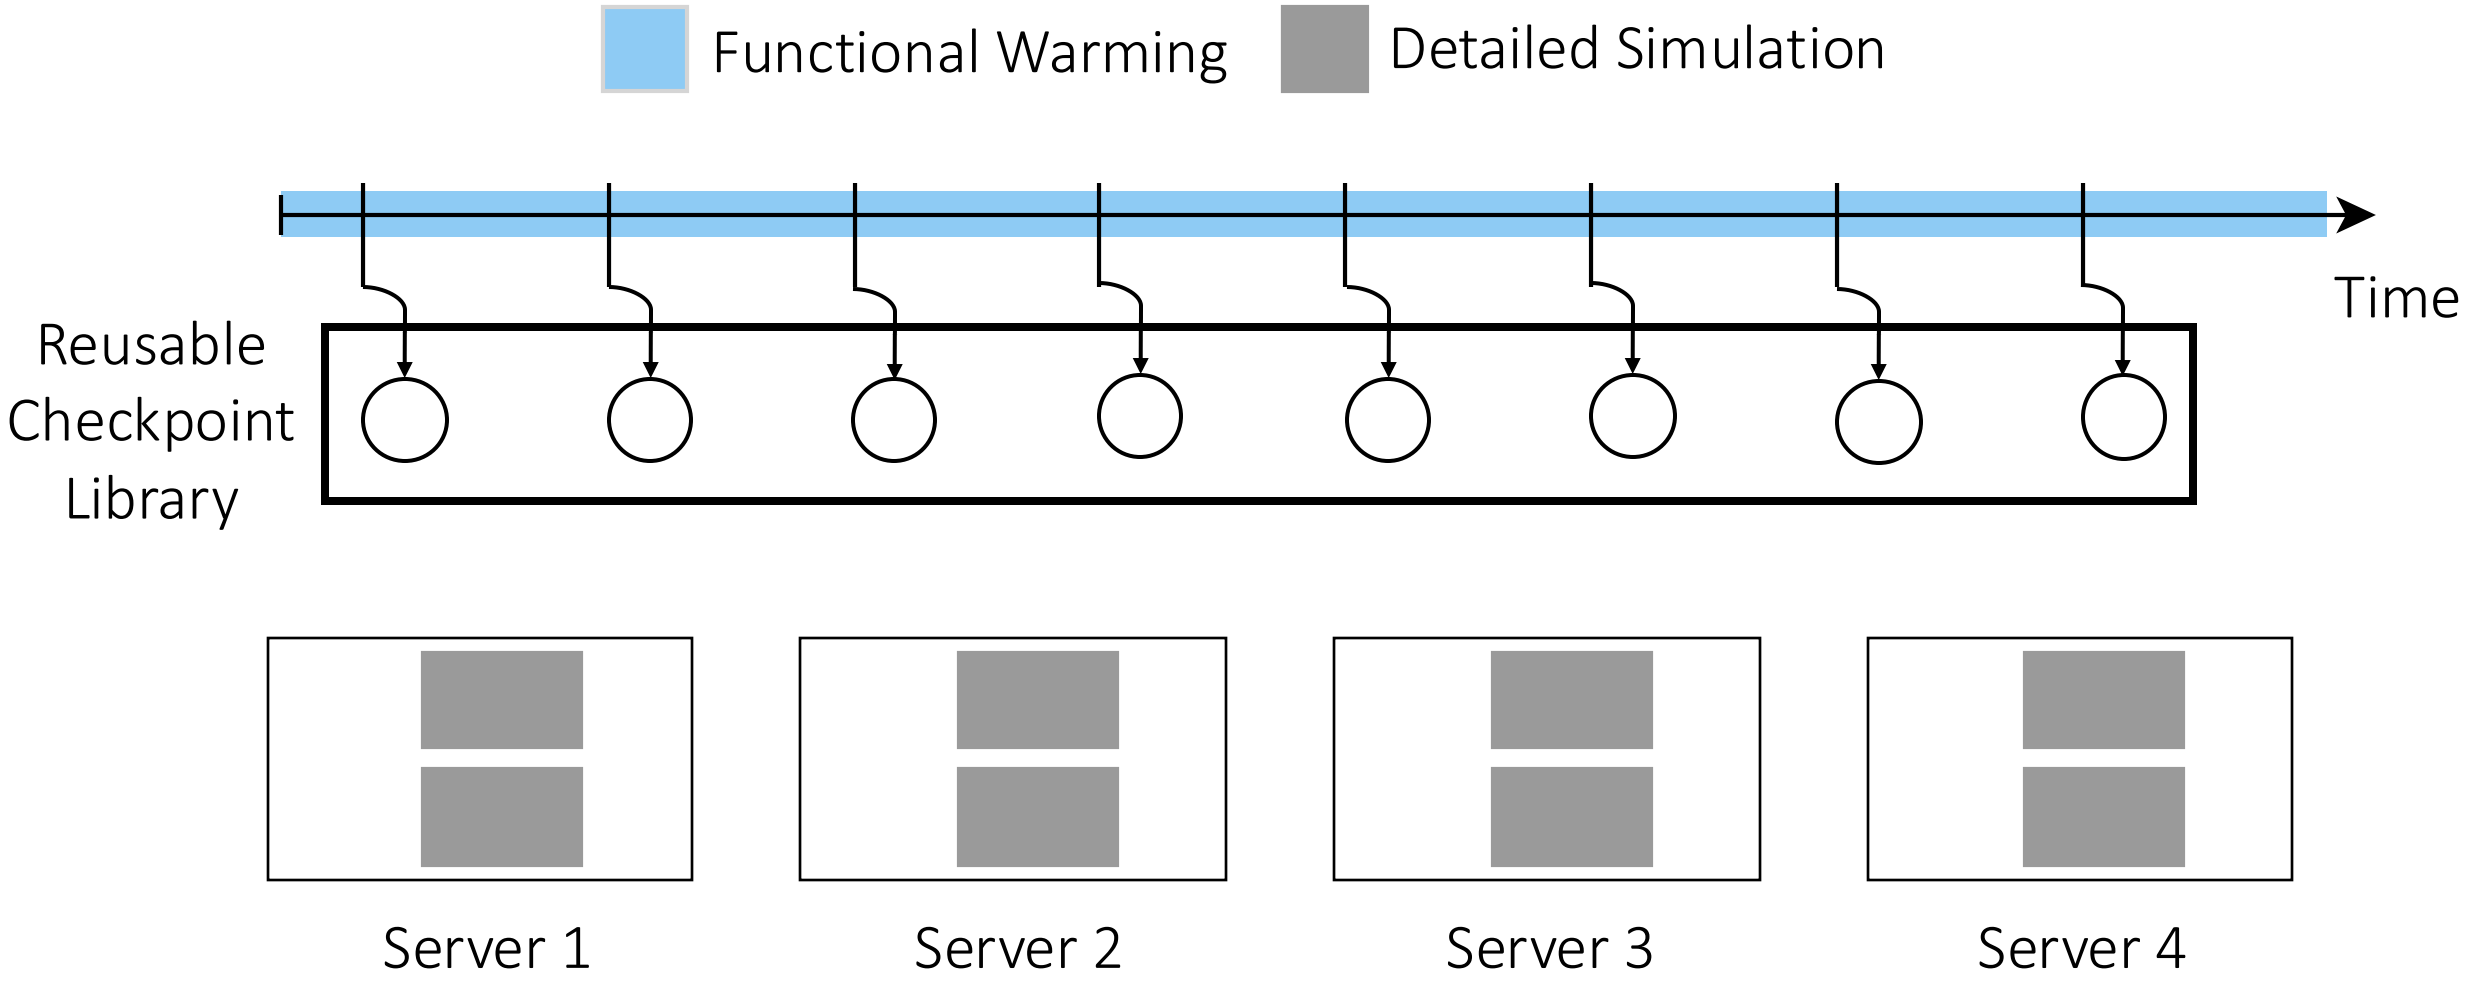

Detailed simulation is embarrassingly parallel

Methodology

PARALLEL SYSTEMS ARCHITECTURE LAB

Workloads:

- CloudSuite 4.0
- DeathStarBench (not in this talk)

Host:

- Dual-socket Xeon Gold 5520+ (56 cores in total) with 2 TB of memory

Target machine:

- 4-wide OoO cores
- TAGE and ITTAGE branch predictors
- 64 KB 4-way L1-I and L1-D
- 2 MB/tile LLC
- MESI directory, 4x overprovisioned

Result: Speed

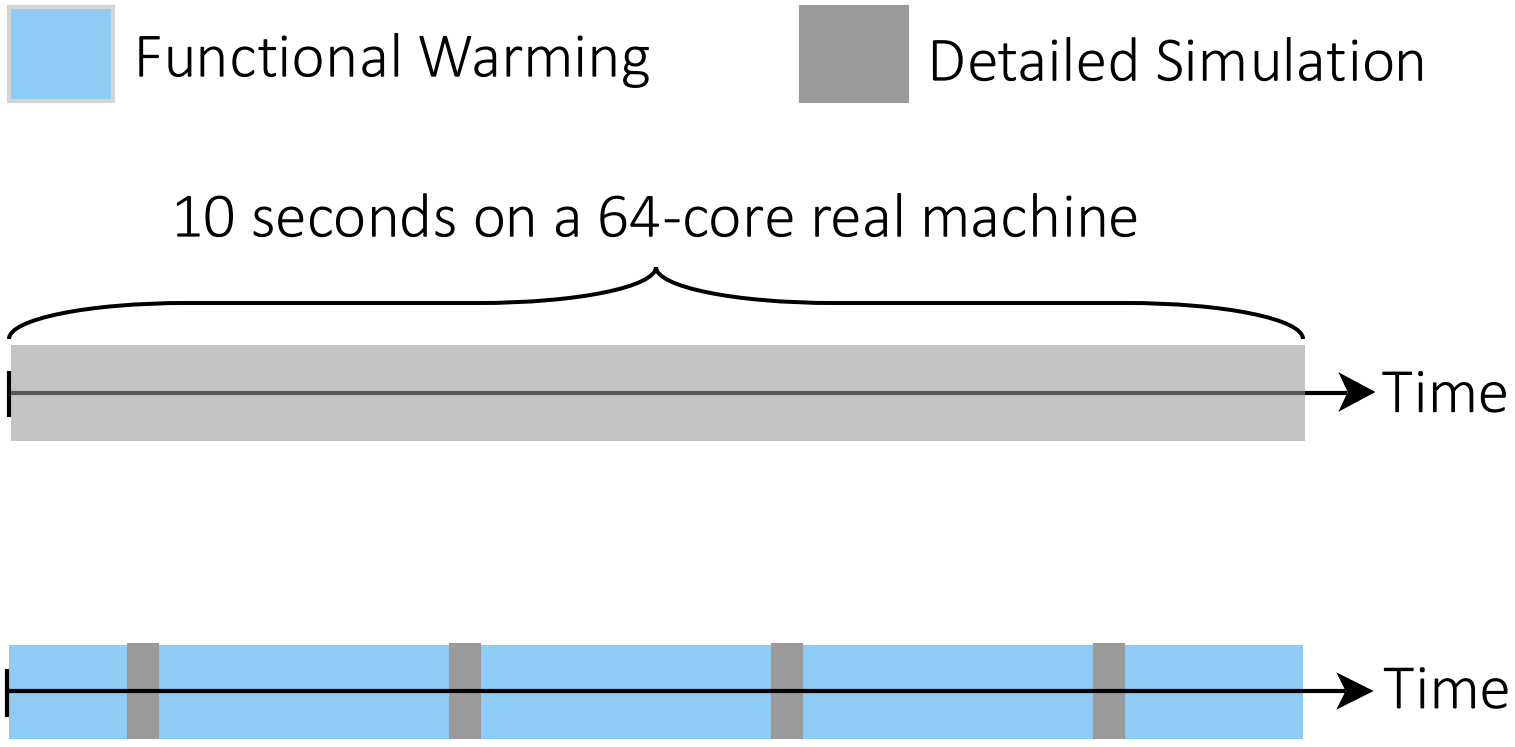

Runs at 250 KIPS
Requires 30 to 68 days

Runs at 10 MIPS
Requires 5 to 34 hours

0.15% to 0.4% cycles simulated in detail and in parallel

QFlex 3.0 runs at 10 MIPS and requires 100x less time

Results: Accuracy

PARALLEL SYSTEMS ARCHITECTURE LAB

Single target core

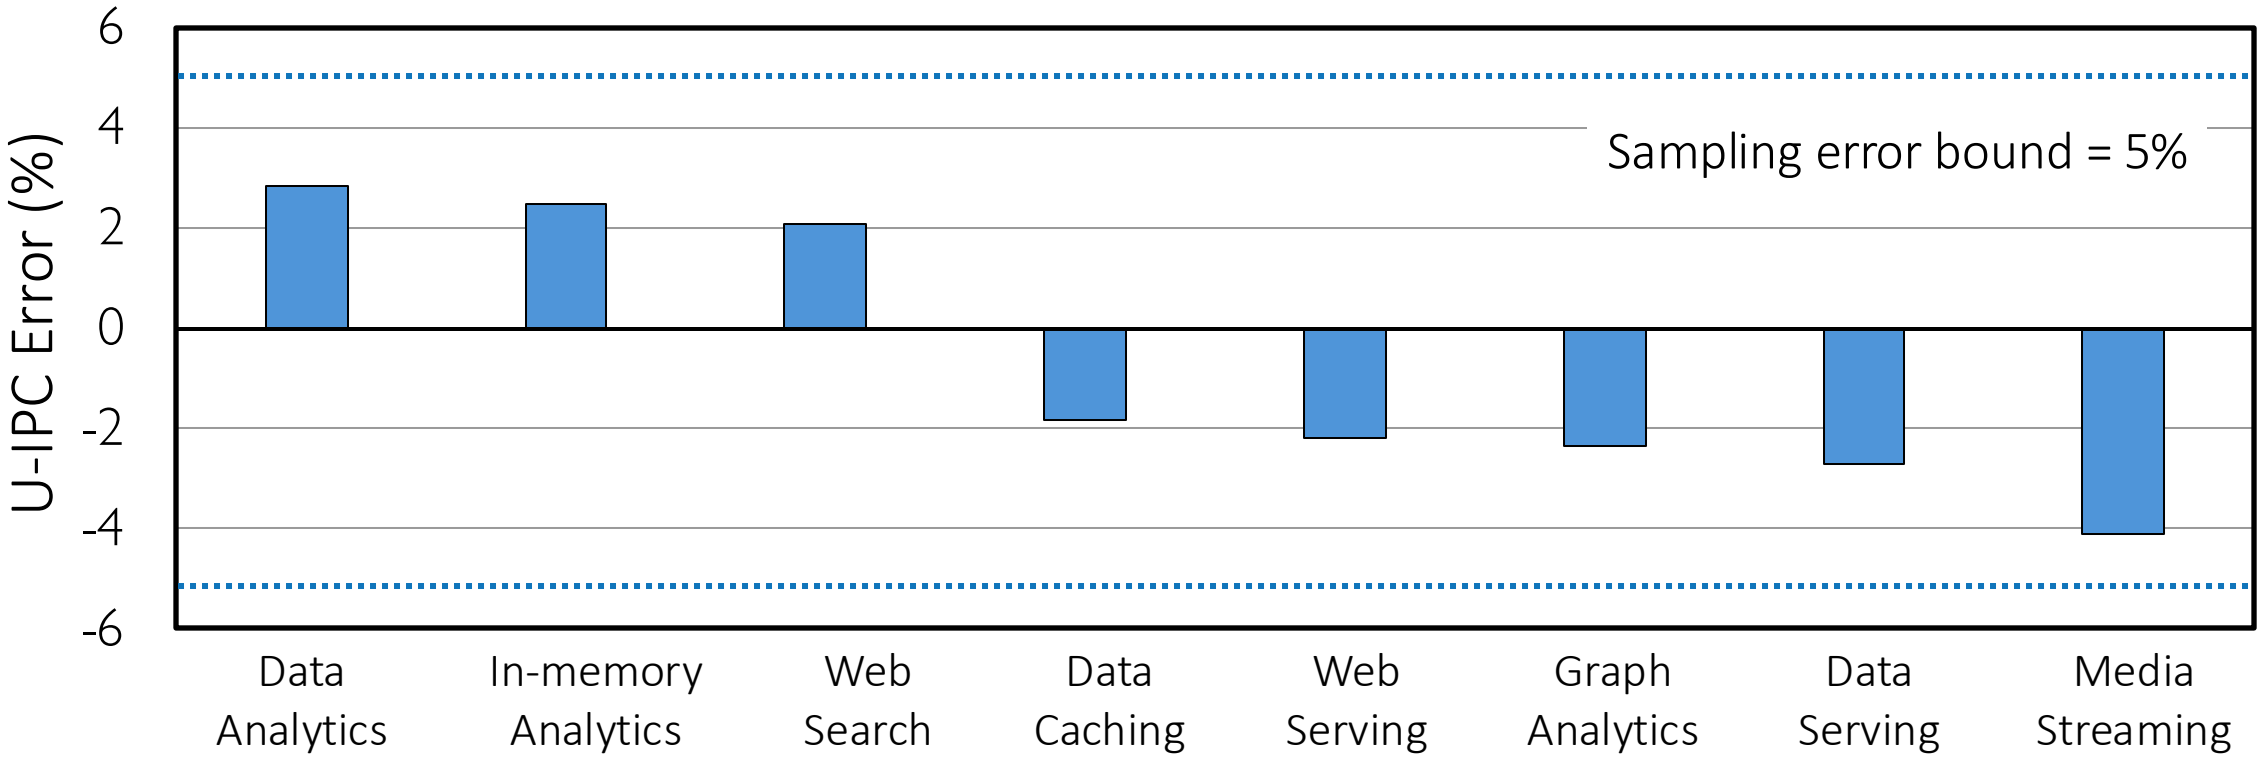

Simulation sampling's results are within 5% of detailed simulation results

Future Directions

PARALLEL SYSTEMS ARCHITECTURE LAB

- Parallel functional warming
  - 50x higher simulation speed for a 64-core host simulating a 64-core target
  - Reaching sub-GIPS simulation speed
- Universal checkpoint format
  - Making checkpoints reusable with other simulators (e.g., gem5)
- Multi-node simulation
  - Simulating several standalone servers communicating over network
  - Rack-scale simulation

Conclusions

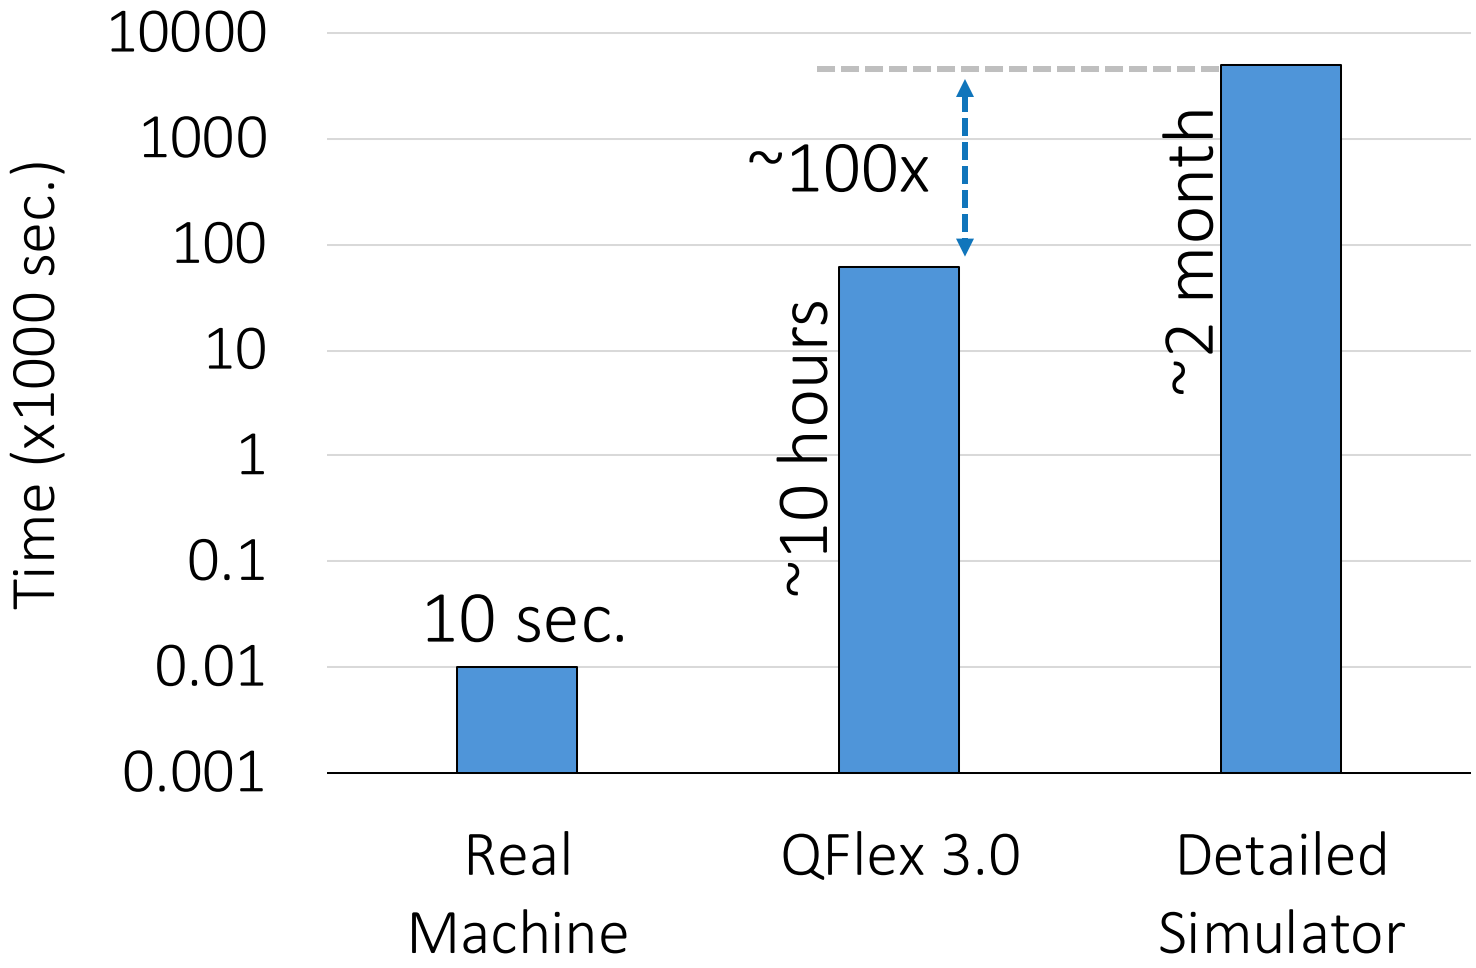

QFlex 3.0: Hour-long simulations, ~100× faster than full detailed simulation

Thank You!

PARSA
PARALLEL SYSTEMS ARCHITECTURE LAB

For more information, please visit us at

parsa.epfl.ch