



The Impact of Operating Systems on Modern CPU Designs (and Vice Versa)

Chris Schlaeger
Director
Operating System Research Center

February, 2008

Software Visible Features changed the Game

After the Megahertz Race

- 64 bit computing
- AMD-V™ technology
- Power consumption
- Multi-core

What has changed?

- Frequency and cache size changes are invisible to the OS
- Architectural features require changes to the OS!

Degrees of Freedom for CPU designs

Address and instruction width

Memory bus connection

Instruction set

Pipeline stages

Number of execution units


Number of cores

Number of CPUs

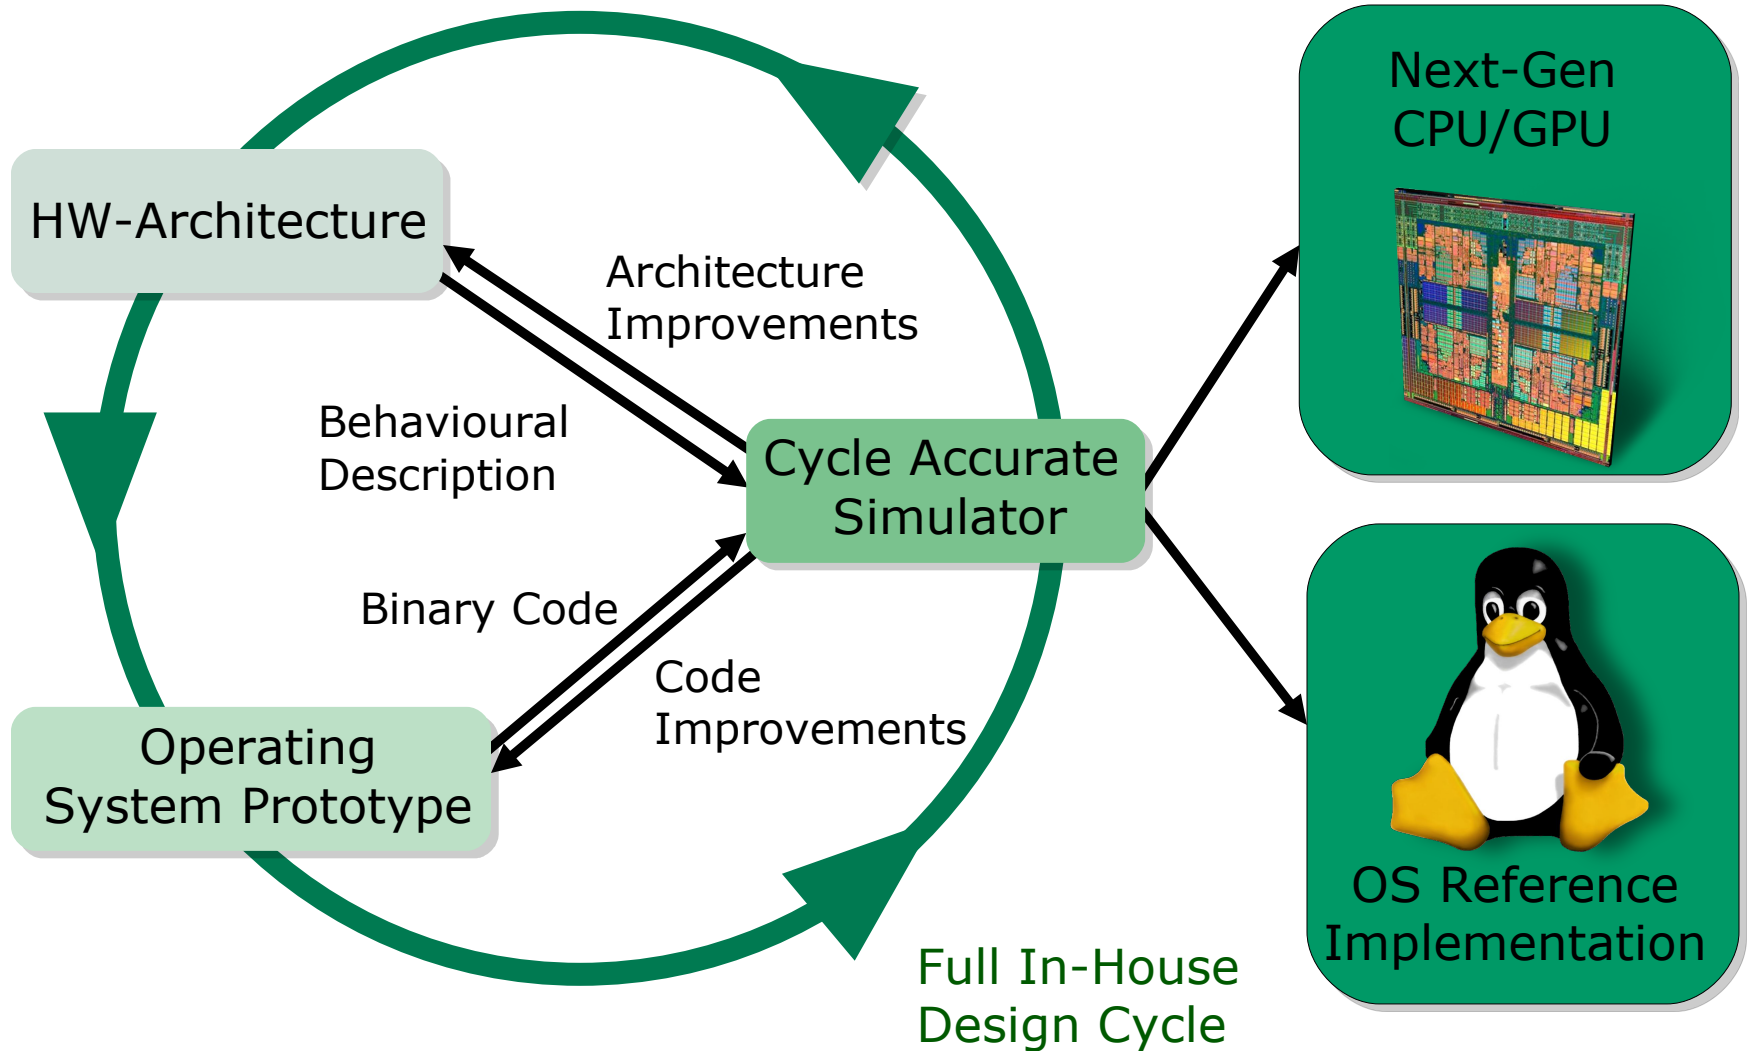
CPU Interconnects

Caches



...



AMD's HW/SW Co-Design Approach



Uniform vs. Non-Uniform Memory



Traditional x86 architecture
 • Frontside bus limits memory bandwidth to a fixed maximum

Direct Connect Architecture
 • Memory bandwidth scales with number of processors

Example: Advanced Synchronization Facility

Advanced Synchronization Facility

- Proposed facility for low-overhead atomic memory modification
 - Change a set of cache lines, mass-commit atomically
- Primitive for higher-level synchronization primitives
 - Roll your own DCAS / LL-SC
 - Highly flexible
- Use almost any x86 instructions in critical sections

Critical section structure

DCAS:

retry:

```

LOCK MOV R8, [mem1] ; Specification begins
LOCK MOV R9, [mem2]
ACQUIRE RCX, 2 ; Atomic block begins
JNZ retry ; Retry in case of abort
XOR RCX, RCX ; DCAS semantics
CMP R8, RAX
JNZ out
CMP R9, RBX
JNZ out
XCHG RDI, [mem1]
XCHG RSI, [mem2]
MOV RCX, 1

```

out:

```

COMMIT ; End of critical section

```


Software Transactional Memory (STM)

- Motivation:
 - Parallel programming is here to stay
 - Parallel programming is hard
- Transactional Memory (TM): New programming paradigm to ease parallel programming
 - Provides transactions (as known from databases)
 - Today's software implementations (STM) incur high overhead
 - Can we accelerate STMs using hardware support?

Advanced Synchronization Facility Evaluation

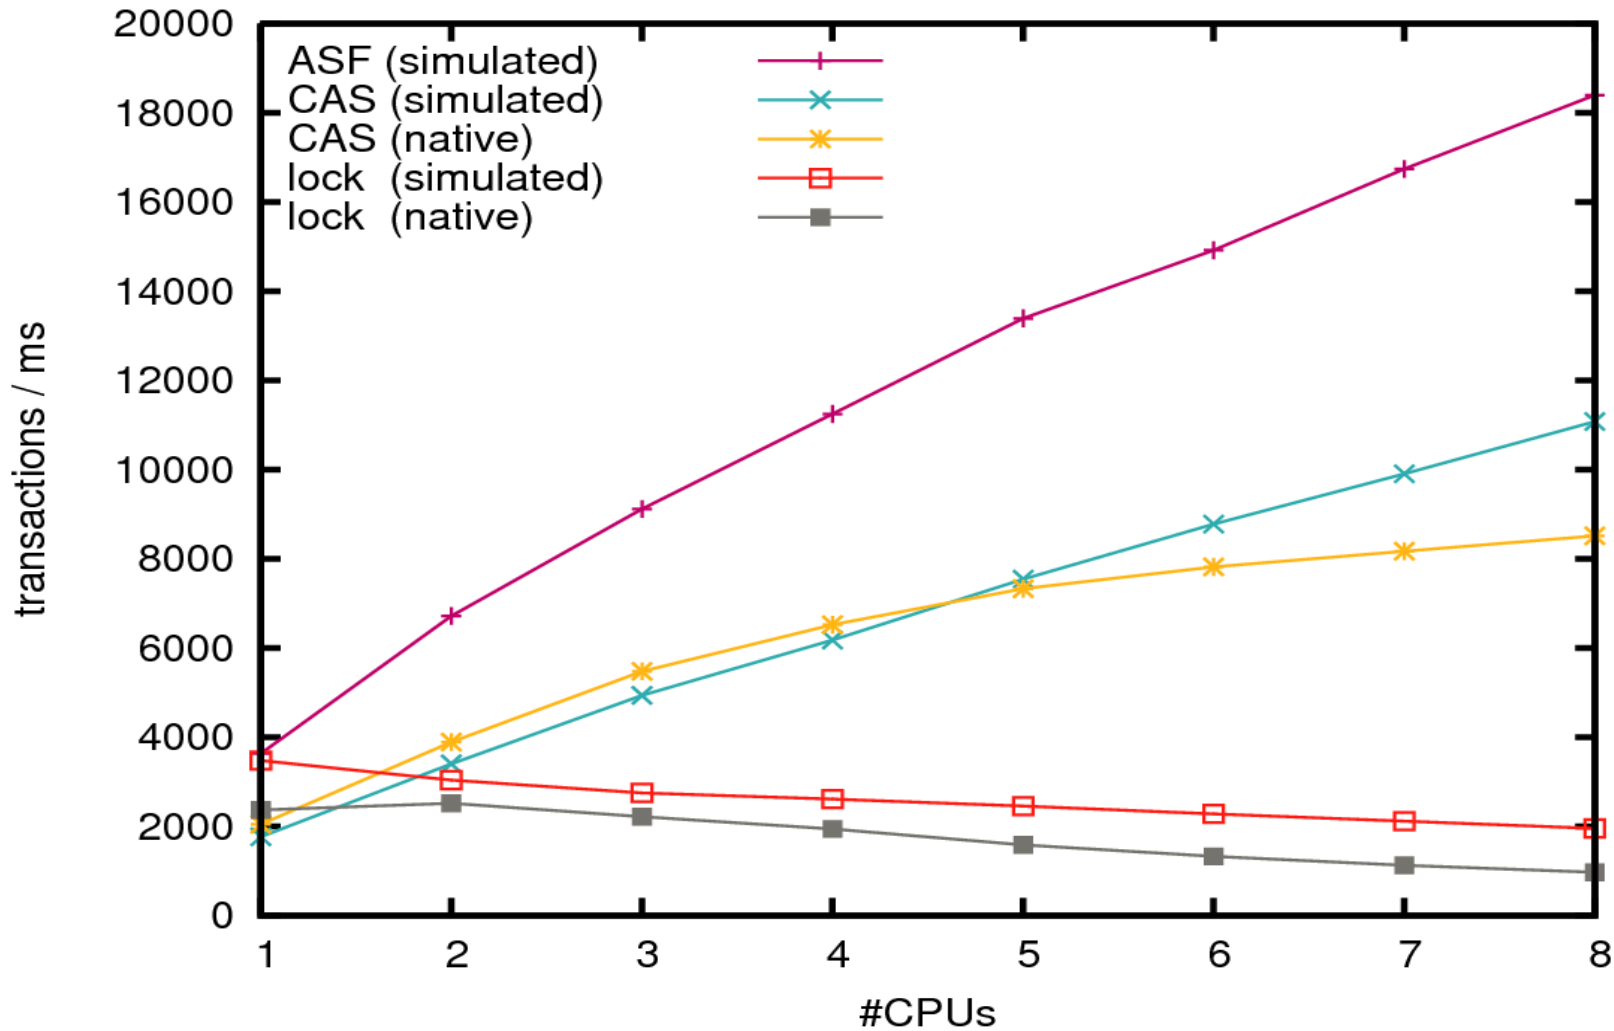
Can we accelerate STMs using Advanced Synchronization Facility?

- Methodology:
 - Add facility to performance simulator
 - Enhance an STM with facility
 - Run facility-enabled STM on top of simulator
- In collaboration with TU Dresden
 - Diploma thesis
 - Foundation of VELOX project

Simulating Advanced Synchronization Facility Performance

- PTLsim
 - “Cycle-accurate” K8 simulator (open source)
 - Extended with models of SMP, cache coherency
 - Added Advanced Synchronization Facility
 - Worked around stability issues
- Why PTLsim?
 - PTLsim has extremely fast turnaround (native execution mode based on Xen hypervisor)
 - PTLsim is easy to share
 - Experience with Xen


Simulation Results – lock-free data structure



STM Acceleration using the Facility

- Using TinySTM
 - State-of-the-art lock-based STM system
- Tried several optimization approaches
 - Most were dead ends
 - Most hot paths had been optimized already
- Final approach:
 - Use facility to reduce bookkeeping
 - Monitor up to 8 memory lines using the facility

Simulation Results – STM (list)



Simulation Results — Conclusion

- Simulation has sufficiently high fidelity
- Advanced Synchronization Facility has very good potential for lock-free programming:
 - Best single-thread performance
 - Excellent scalability
- Advanced Synchronization Facility has potential for accelerating STM
 - Competitive single-thread performance
 - Excellent scalability
 - Not all effects understood yet

7 Challenges for the Future

Challenges for the Future

- OS and architecture support for parallelization
- OS support for noncoherent memory architectures
- Message Passing vs. Shared Memory vs. ?
- OS support for accelerators
- Power Management
- Architecture and operating-system support for secure computing
- Fault Tolerance

Trademark Attribution

AMD, the AMD Arrow logo, AMD Opteron and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. Linux® is a registered trademark of Linus Torvalds. Other names used in this presentation are for identification purposes only and may be trademarks of their respective owners. Linux 2.0 Penguin courtesy of lewing@isc.tamu.edu

©2008 Advanced Micro Devices, Inc. All rights reserved.