

SCOORE: ***Santa Cruz Out-of-Order RISC Engine***

FPGA Design Issues

Francisco J. Mesa-Martinez, Abhishek Sharma, Andrew W. Hill, Carlos A. Cabrera, Cyrus Bazeghi, Hari Kolakaleti, Joseph Nayfach, Keertika Singh, Kevin S. Halle, Matthew D. Fischler, Melisa Nuñez, Sangeetha Nair, Suraj Narendra Kurapati, Wael Ali Ashmawi, and **Jose Renau**



- **Simulation speed**

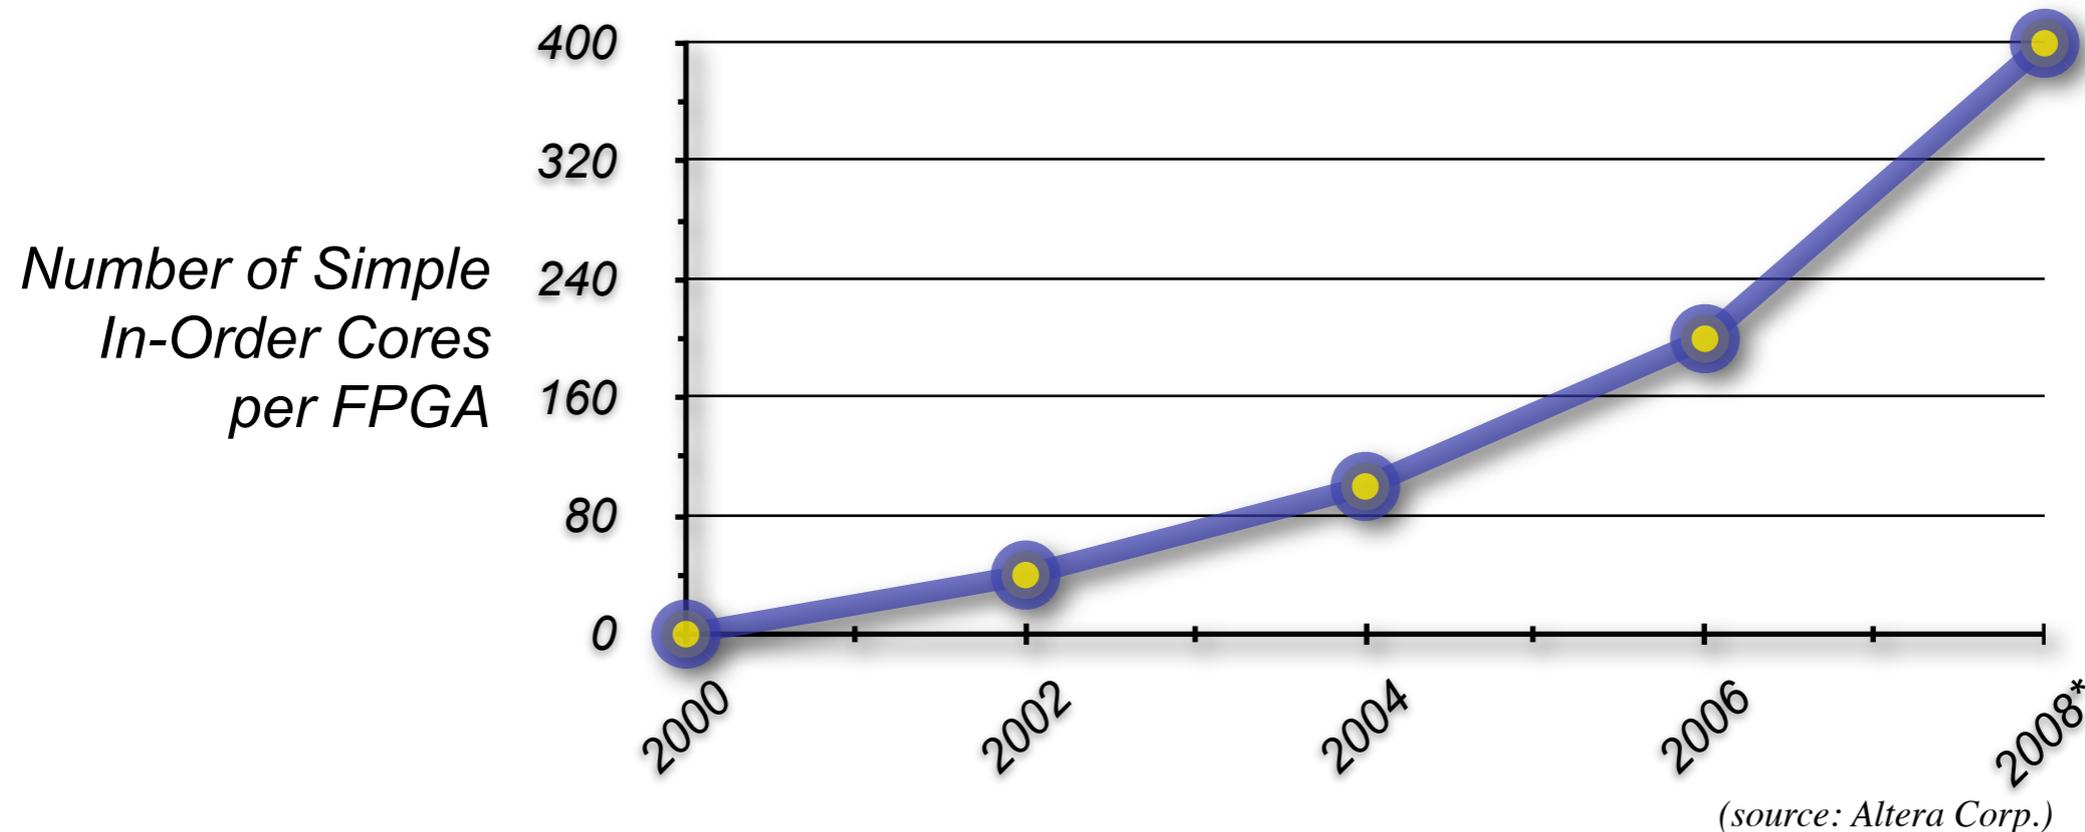
- Most paper evaluations are limited to short execution times
- This gets worse with multiprocessor and multicore simulations

- **Disconnect between theory and implementation**

- System simulators are an abstraction
- Novel architectural ideas may have hidden implementation complexity

Are FPGAs a Way Out ?

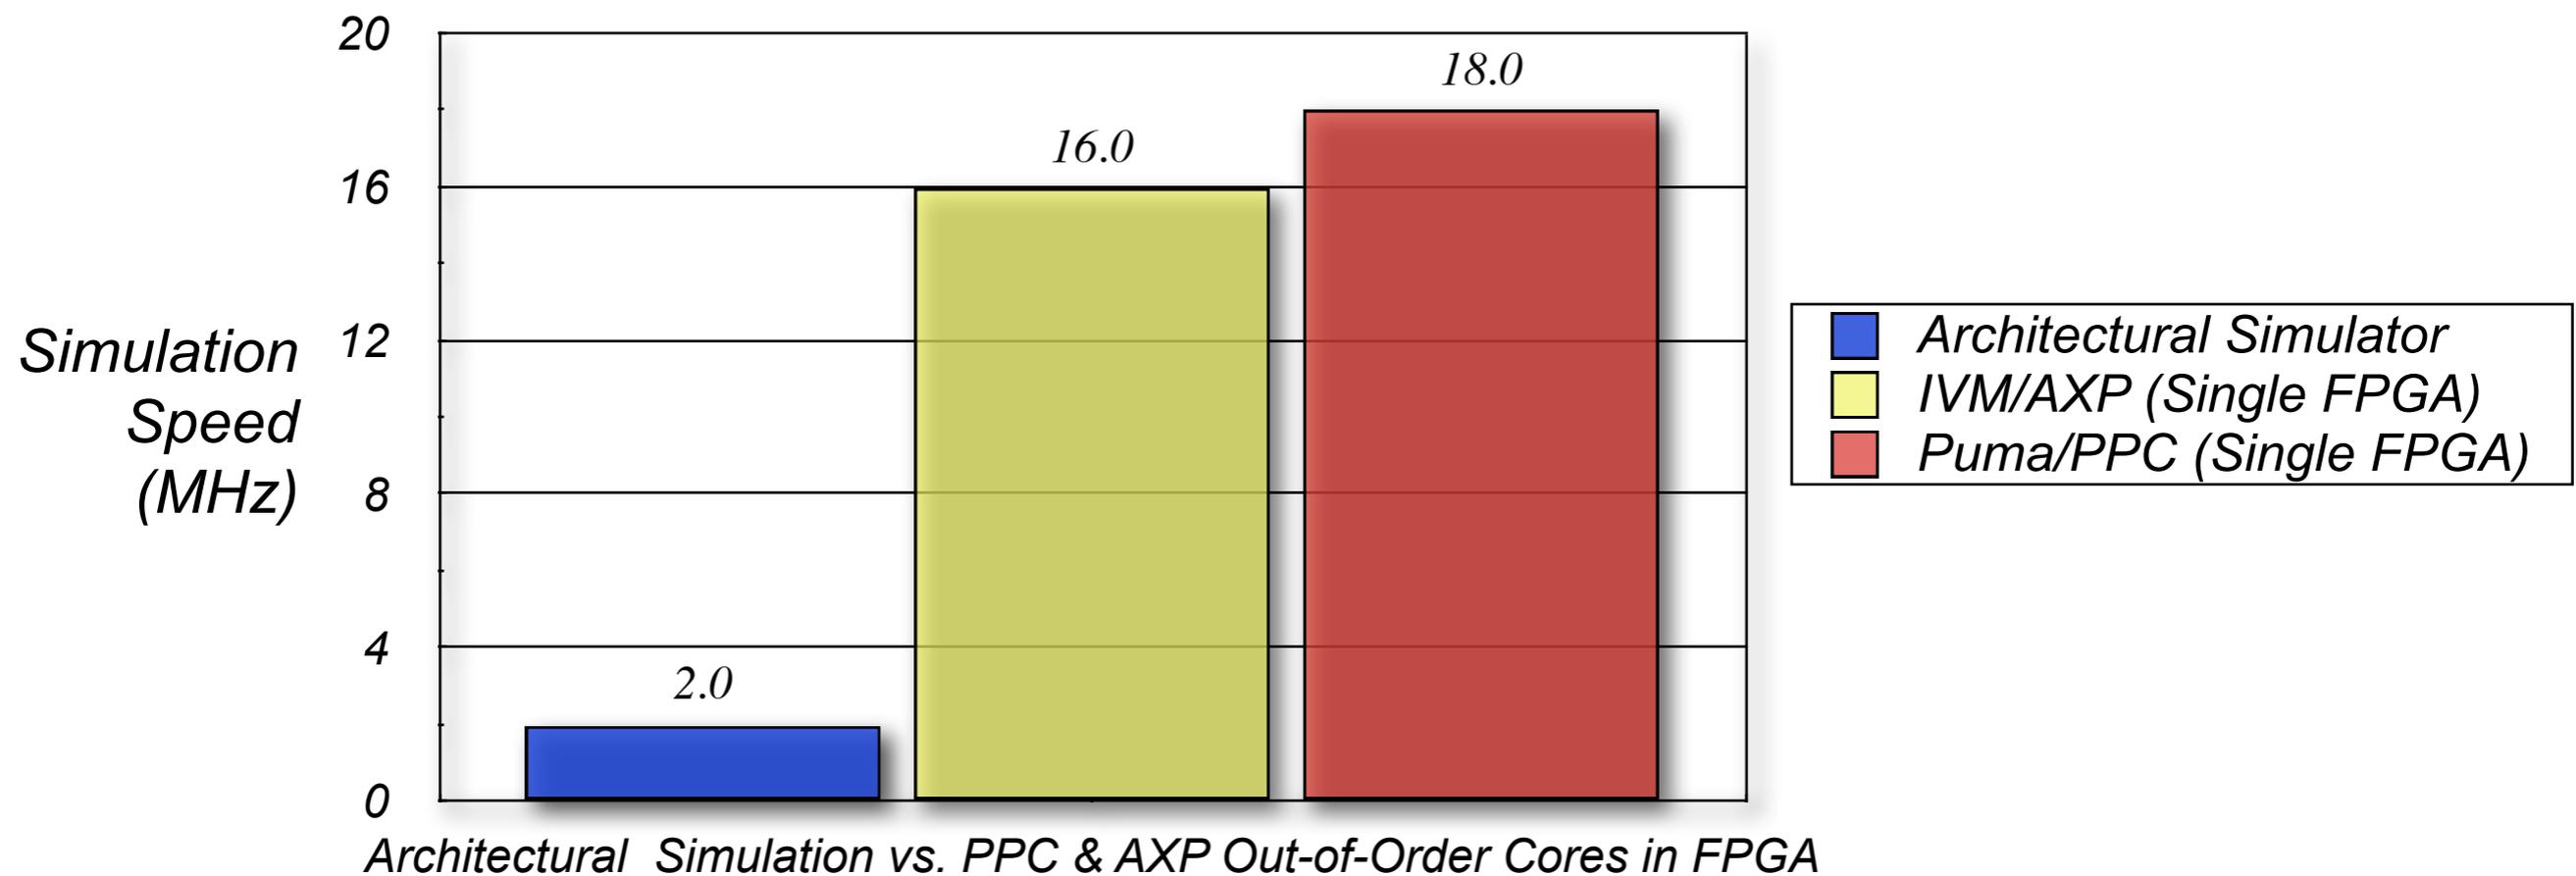
- FPGAs are getting faster, larger, cheaper...



- Can we use these new FPGAs to:
 - Speed up simulation time?
 - Bridge the divide between theory and implementation?

Not so Fast!

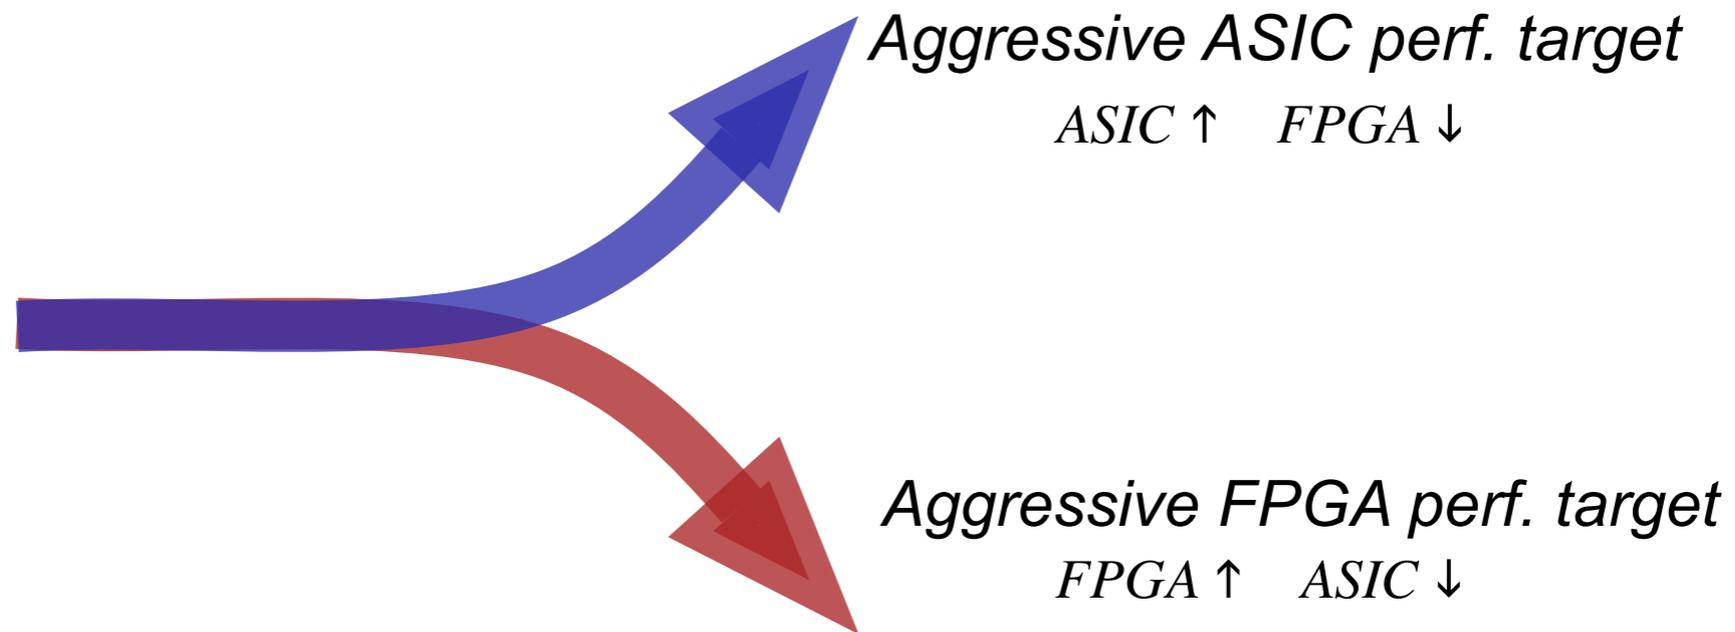
- Existing Out-of-Order designs do not map well to FPGAs
 - Most are designed with a focus on ASIC implementation



- Further, these designs require multiple FPGAs!!!

Why Is Out-of-Order So Bad On FPGAs?

- FPGA and ASIC design techniques may be tangential to each other!

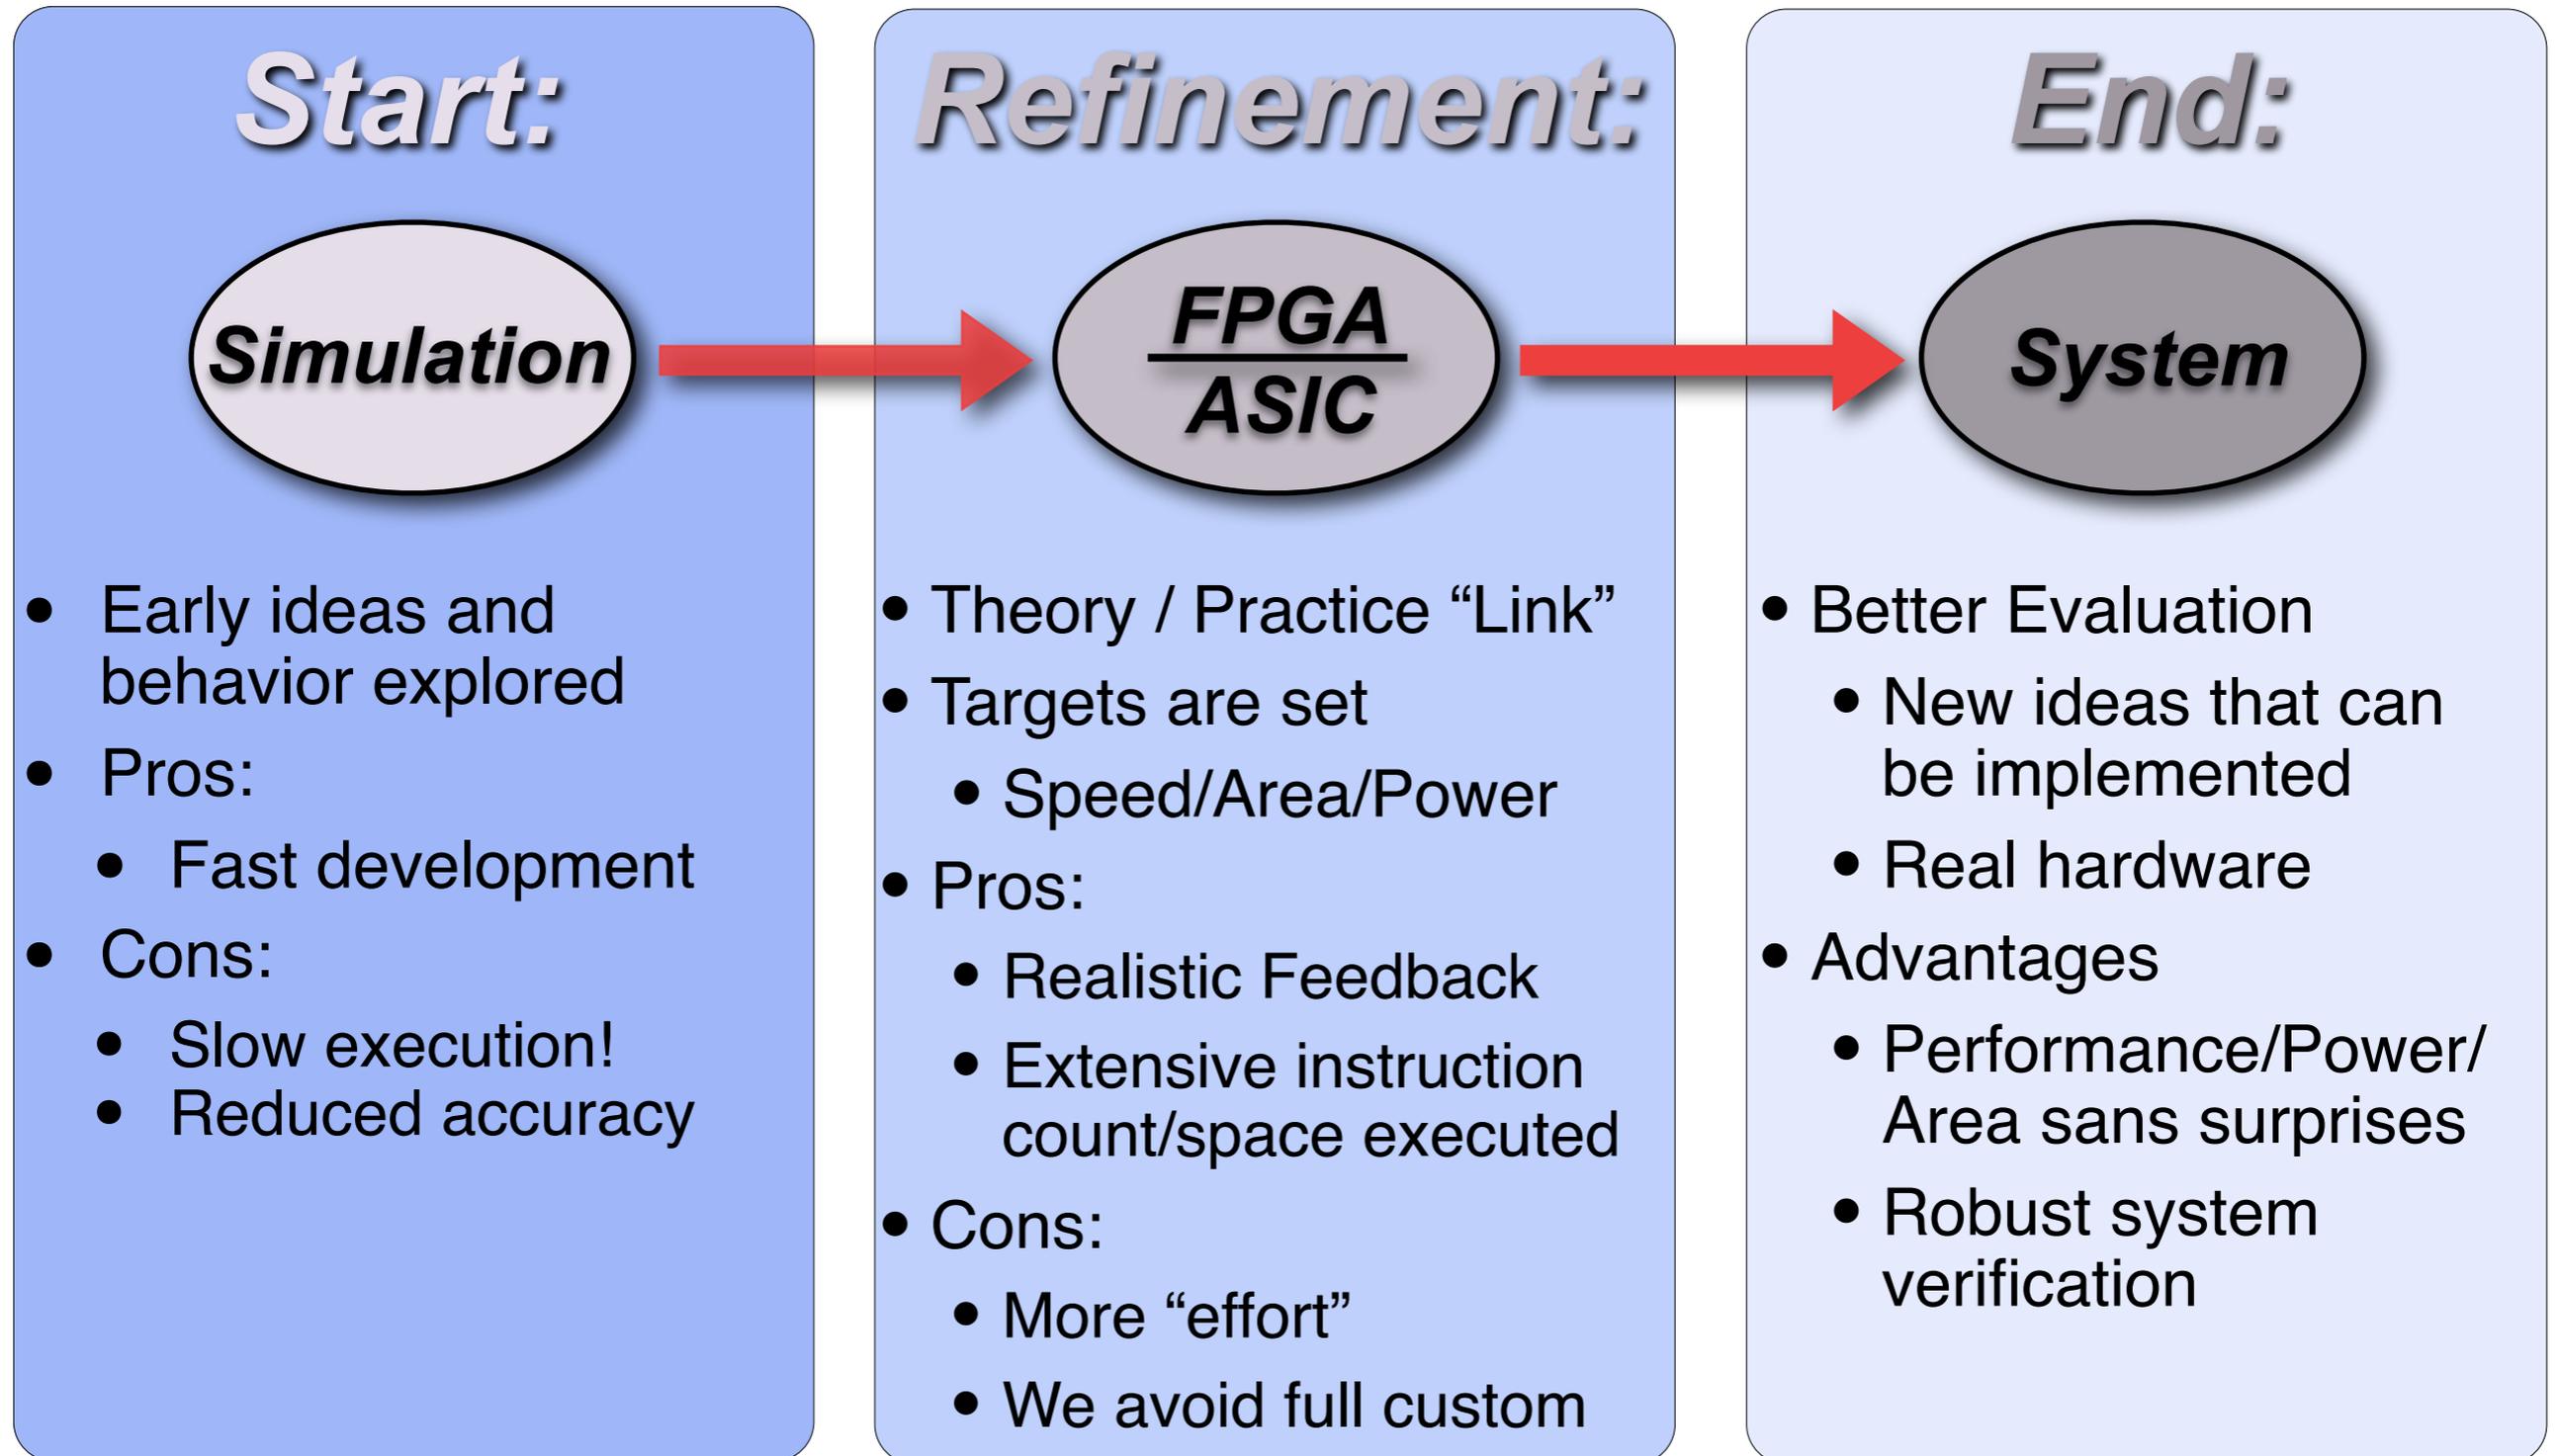


- Out-of-Order aggravates these issues
 - **Strike 1:** Complex structures
 - **Strike 2:** Multiported & fully associative structures
 - **Strike 3:** Wiring & sizing overhead

FPGAs Can Still Be Part Of The Solution

- **Problem:** Simulation speed
- **Solution:**
 - Keep the simulator for quick evaluation
 - Take advantage of an FPGA-aware flow to speed up simulation
 - Use an FPGA-friendly OoO core (SCOORE)
- **Problem:** Disconnect between theory and implementation
- **Solution:**
 - Couple FPGA flow with ASIC targeting to increase accuracy

Proposal: Hybrid ASIC/FPGA Flow



SCOORE: An FPGA-friendly Design

- Design parameters:
 - Design follows hybrid FPGA/ASIC design principles
 - Based on SPARC
 - 12-stage pipeline
 - 4-way superscalar
 - Clustered architecture
 - Out-of-Order (up to 192 in-flight instructions)
 - Memory speculation, load-hit speculation
 - Performance targets:
 - 125MHz FPGA
 - 900MHz ASIC @ 90nm
- We have used SCOORE as the development platform to define a set of guidelines for FPGA-friendly OoO designs

Lesson 1: Don't Cascade SRAM Structures

◆ Problem:

- Cascading has small impact in ASICs, however FPGA implementations suffer greatly
 - Over 50% performance hit!!!
- Cascading between Branch Target Buffer (*BTB*) & Return Address Stack (*RAS*) predictor in IF
 - Critical path which limits speed of IF stage

• Solution:

- Avoid cascading SRAM structures in a single cycle
- FPGA implementation reduces size and clock cycle significantly
 - FPGA and ASIC both benefit from shortened critical path

Lesson II: Avoid Heavily Ported Structures

◆ Problem:

- Heavily multiported structures are not “FPGA-friendly”
 - Limits on the number of Write ports in FPGA SRAM banks
 - The instruction window is a large heavily ported structure
- CAM structures are bad for both ASIC and FPGA
 - ASIC: Large CAM structures are larger and power hungry
 - FPGA: CAM’s are implemented using logic cells!!!

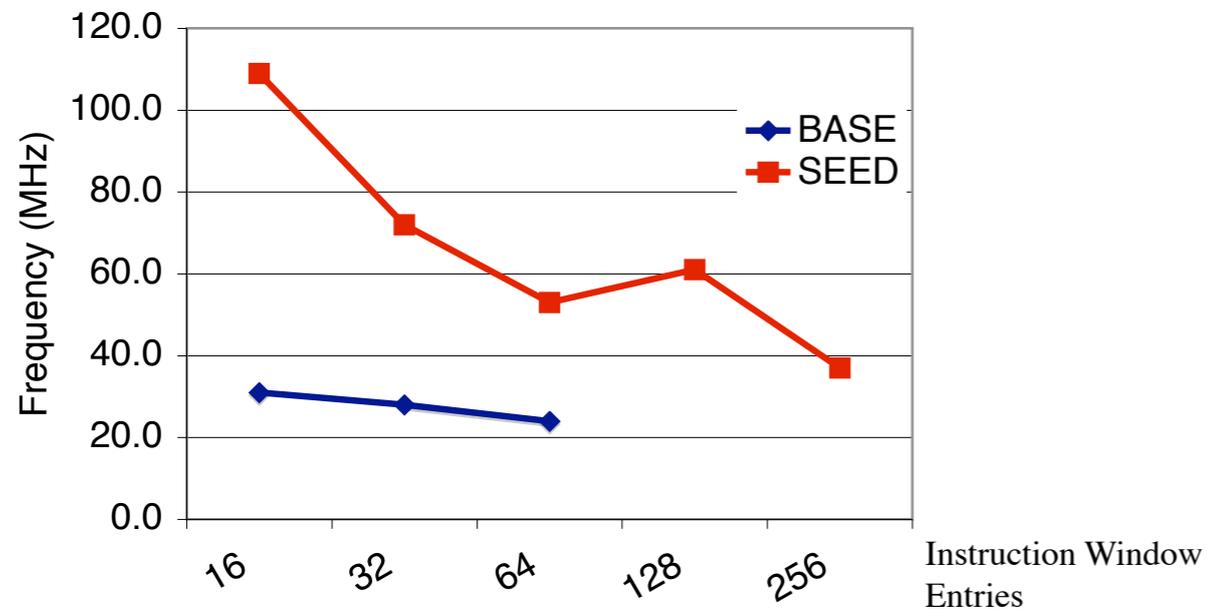
◆ Solution:

- Break monolithic structure into multiple banks
- Ex: SEED (*Scalable Efficient Enforcement of Dependencies*)¹
 - Shows dramatic speedup in FPGA
 - Reduced area/power requirements in ASIC due to reduced port count
 - Allows the implementation of very large instruction windows that were not practical before

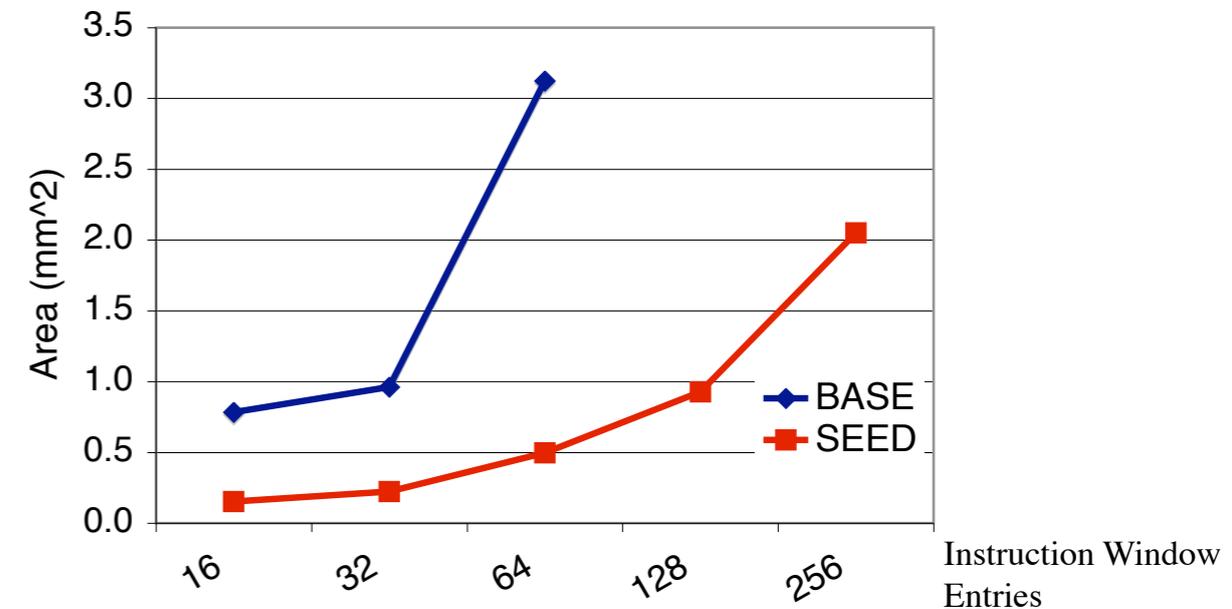
¹ “SEED: Scalable Efficient Enforcement of Dependencies,” F. Mesa-Martinez, M. Huang, and J. Renau, to Appear in PACT06

SEED: Results

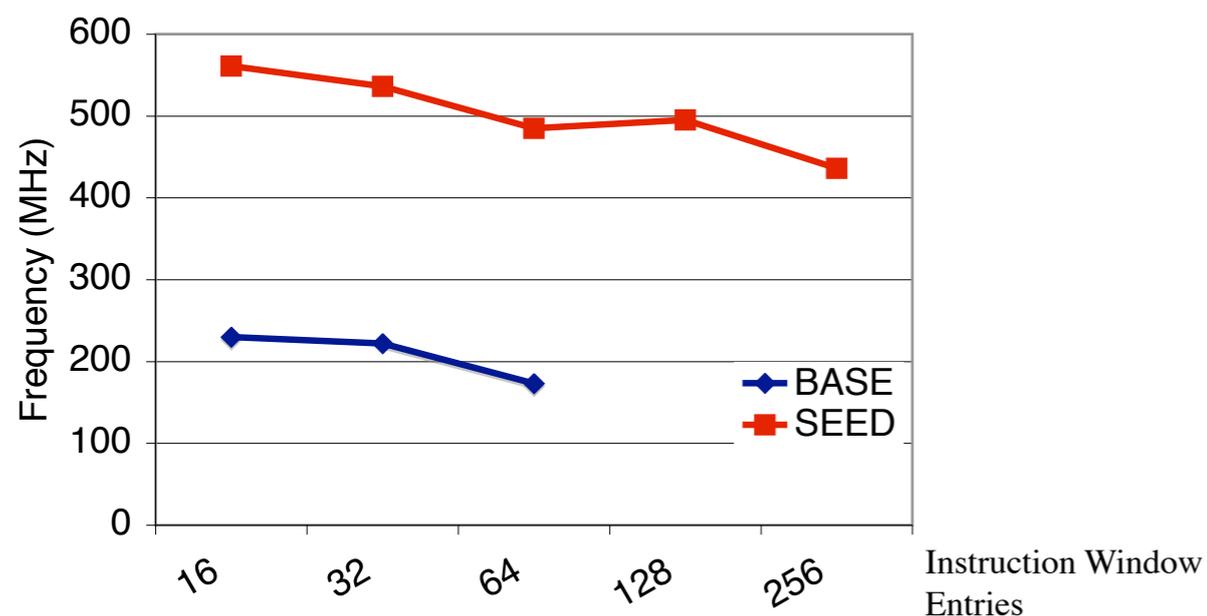
• FPGA Frequency:



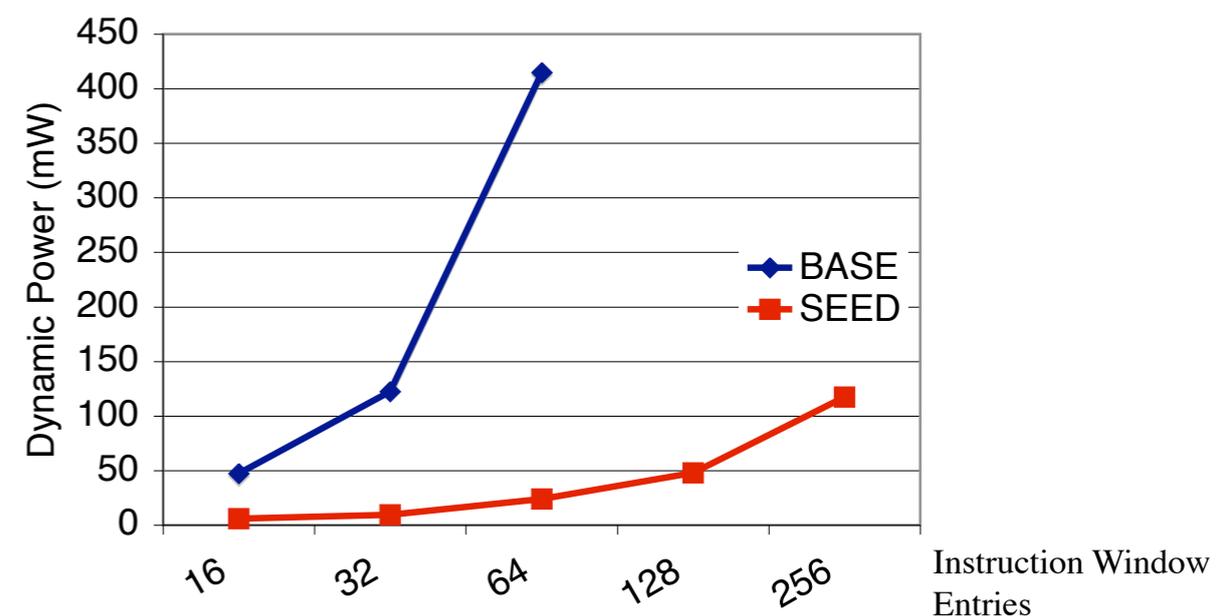
• ASIC Area:



• ASIC Frequency:



• ASIC Dynamic Power:



Lessons Learned: Design Guidelines

● **Wiring and Sizing:**

- FPGA and ASICs require different resource sizing
 - FPGA: Fixed SRAM/MAC/etc structures
 - ASIC: More sensitive to SRAM sizing
- Wire delay is major limiter in ASIC and even bigger in FPGA

● **SRAM Structures:**

- SRAM-like structures in processor should be mapped into the built-in SRAM structures in an FPGA
 - Avoid multiple Write ports: Most FPGAs have a single W port
 - Avoid fully associative structures: FPGAs map such structure into logic
 - Avoid single-cycle cascade of SRAM structures

Acknowledgements

- Puma: Jayakumaran Sivagnaname (University of Michigan)
- IVM: Nicholas Wang (University of Illinois)

Contact Information

Web: <http://masc.soe.ucsc.edu/>

Your's truly: **Francisco J. Mesa-Martinez**

e-mail: javi@soe.ucsc.edu

Advisor: **Jose Renau**

e-mail: renau@soe.ucsc.edu



SEED: Theory of Operation

- Integration of SEED into an out-of-order pipeline:

