Research Methods in Computer Architecture

TDT8108 Topics in Information Technology

11. February 2015

Assoc. Prof Magnus Jahre, CARD

Outline

- Introduction: Evaluation Requirements and Possibilities
- Computer Simulation
- Example Paper:
  «A Quantitative Study of Memory System Interference in Chip Multiprocessor Architectures»
- How to get a PhD?
- Final comments

Who am I?

Education

- Master of Technology, NTNU, 2007
- PhD, NTNU, 2010. Title: «Managing Shared Resources in Chip Multiprocessor Memory Systems»

Current

- Associate Professor in Computer Architecture (since 2012)
- Head of CARD research group
- Coordinator of the Energy Efficient Computing Systems «fyrtårn»

Research Interests

- Memory systems for multi-core architectures, heterogeneous computer systems, energy efficiency, computer architecture simulation, compilers and system software

What is Science?

Wikipedia:

Science is a *systematic* enterprise that builds and organizes *knowledge* in the form of *testable explanations* and *predictions* about the universe.

What is Computer Science?

Wikipedia:

Computer science is the *scientific* and *practical* approach to *computation* and its *applications*. It is the *systematic study* of the feasibility, structure, expression, and mechanization of the methodical *processes (or algorithms)* that underlie the acquisition, representation, processing, storage, communication of, and access to information […]

Evaluation Requirements

Validity

- How well are you measuring what you want to measure?

Reliability

- Are the results stable and consistent?

Reproducibility

- Do you provide enough information for independent researchers to verify your findings?

Other issues

- Your implementation must be correct
- Your models must be sound (i.e. crucial components modeled in sufficient detail)

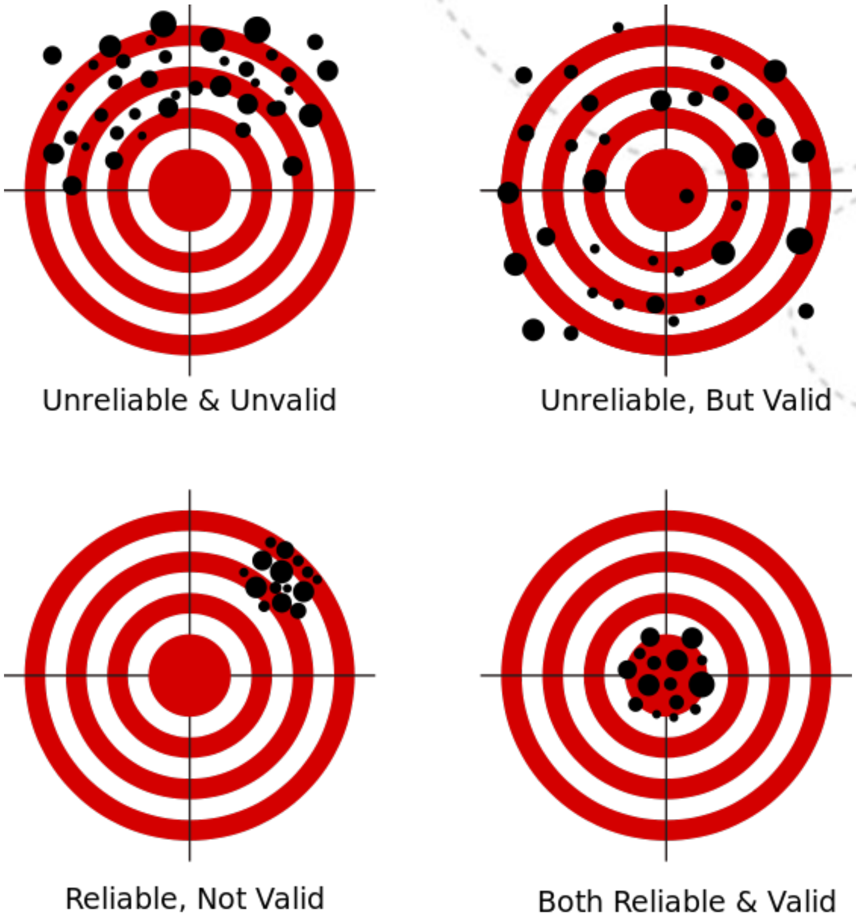

Acknowledgement: Figure from Wikipedia

Evaluation Possibilities

Analytical modeling

- Suitable for early studies and large design spaces
- Concern: Normally validated against simulators if validated at all

Simulation

- Can model complex interactions between independent units
- Rapid implementation and evaluation of new hardware possible
- Possibility for extensive, non-invasive measurements
- Concern: Simulators are rarely validated against real hardware
- Concern: How can we verify validity across the design space?

Focus of this session

Experiments on real hardware

- Accounts for all details of underlying implementation
- Concern: Hard to create an experimental setup that is reliable
- Concern: Hard to verify the validity of results
- Limitation: Significant effort needed to implement new hardware

From: K. Skadron et al. «Challenges in Computer Architecture Evaluation», Computer, 36, 2003.

Amdahl's Law

- *"If a fraction s of a (uniprocessor) computation is inherently serial, the speedup is at most 1/s"*

- Total work in computation
  - serial fraction s
  - parallel fraction p
  - s + p = 1 (100%)

- S(n) = Time(1) / Time(n)
  = (s + p) / [s +(p/n)]
  = 1 / [s + (1-s) / n]
  = n / [1 + (n - 1)s]

- **Relevance to this talk:**
  If your scheme is applicable to a small part of the total, the overall impact is severely limited

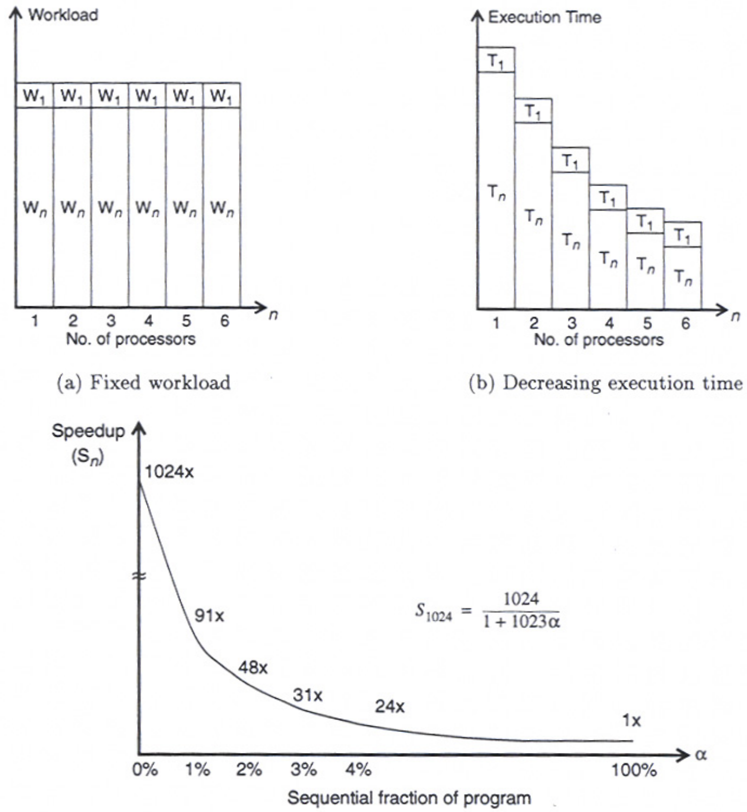

(a) Fixed workload

(b) Decreasing execution time

(c) Speedup with a fixed load

Figure 3.8 Fixed-load speedup model and Amdahl's law.

Outline

Introduction: Evaluation Requirements and Possibilities

Computer Simulation

Example Paper:
«A Quantitative Study of Memory System Interference in Chip Multiprocessor Architectures»

How to get a PhD?

Final comments

Simulator-based Research Workflow

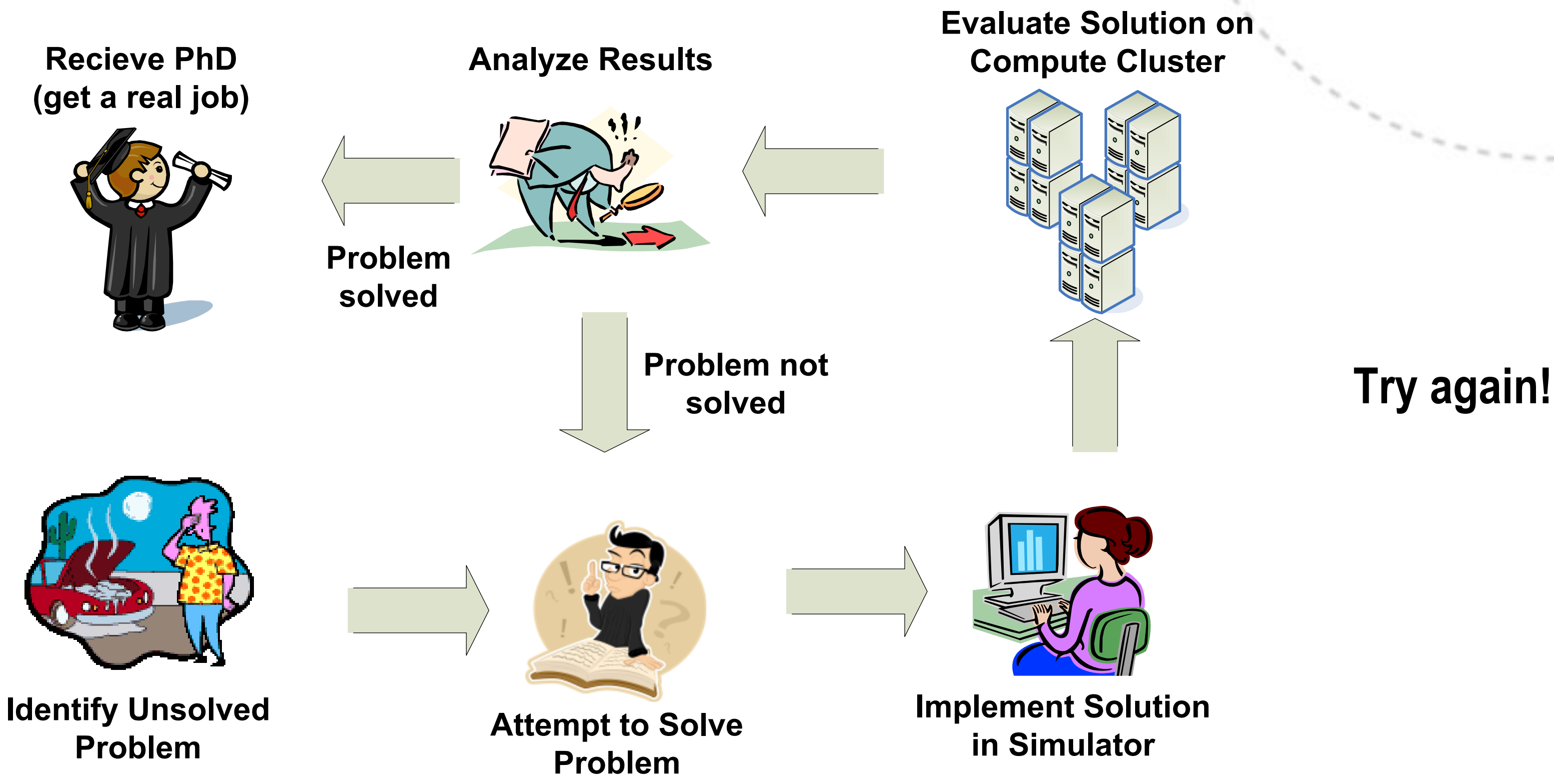

What is a simulator?

Model of a system

- Model the interesting parts with high accuracy
- Model the rest of the system with **sufficient** accuracy

One-to-one correspondence with actual hardware?

- Not necessary
- Try to model **behavior**
- Simplify your code wherever possible

"All models are wrong but some are useful"
(G. Box, 1979)

Benchmarks

Benchmark types

- Kernels, toy programs and synthetic benchmarks
  - Widely discredited
  - Too easy for compiler writers and computer architects to cheat
- Embedded benchmarks: EEMBC, MiBench, etc.
- Desktop benchmarks: SPEC 2006, Parsec, etc.
- Server benchmarks: TPC-C, TPC-H, etc.

Computers should work well for a collection of programs

- *Average* performance is the key metric
- Benchmarks are often assembled into *suites*

Key question: How representative are the selected benchmarks of the programs run by actual users in the target domain?

Outline

Introduction: Evaluation Requirements and Possibilities

Computer Simulation

Example Paper:
«A Quantitative Study of Memory System Interference in Chip Multiprocessor Architectures»

How to get a PhD?

Final comments

Chip Multiprocessor Interference

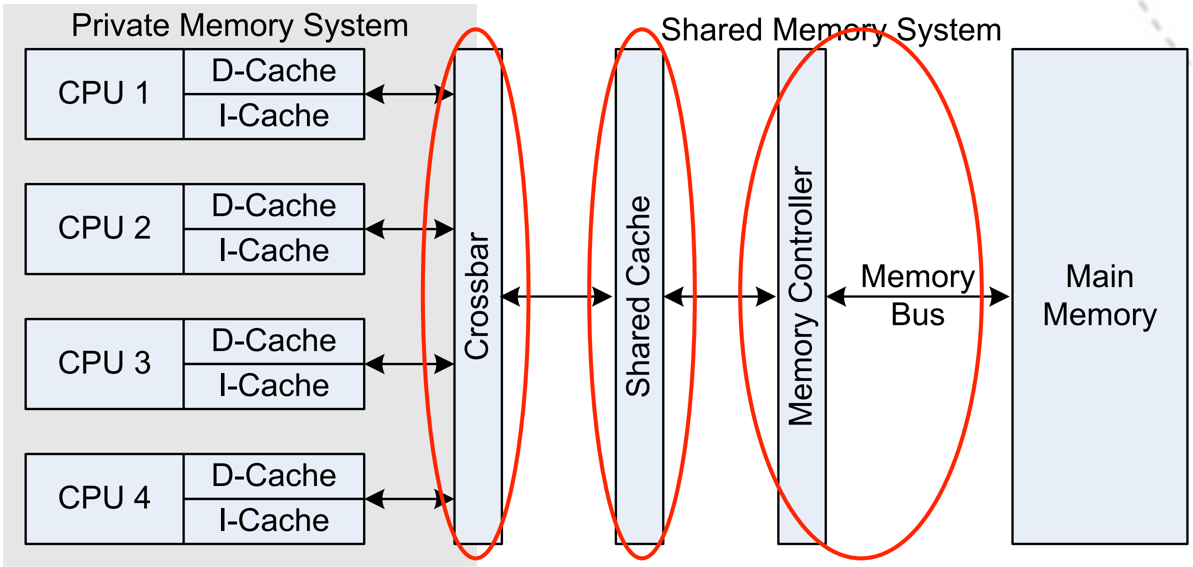

- Hardware-controlled, shared resources
  - Crossbar bandwidth
  - Shared cache capacity
  - Memory bus bandwidth
  - Memory capacity is allocated by the operating system

Interference can occur in all shared units

Current CMP implementations do not take interference into account

Offline Interference Measurement

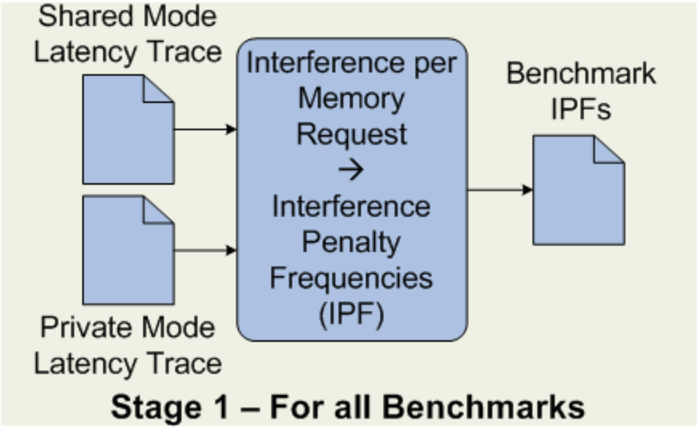

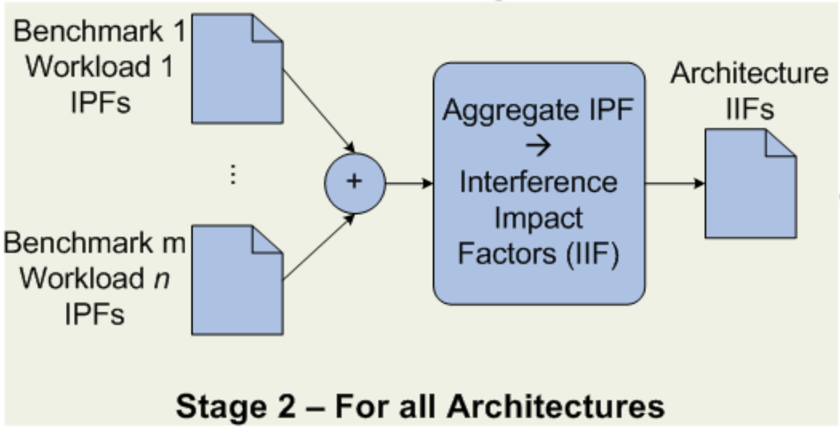

Interference Penalty Frequency (IPF) counts the number requests that experienced an interference latency of *i* cycles

Interference Impact Factor (IIF) is the interference latency times the probability of it arising, i.e. $IIF(i) = i \cdot P(i)$

Performance Variability (Fairness)

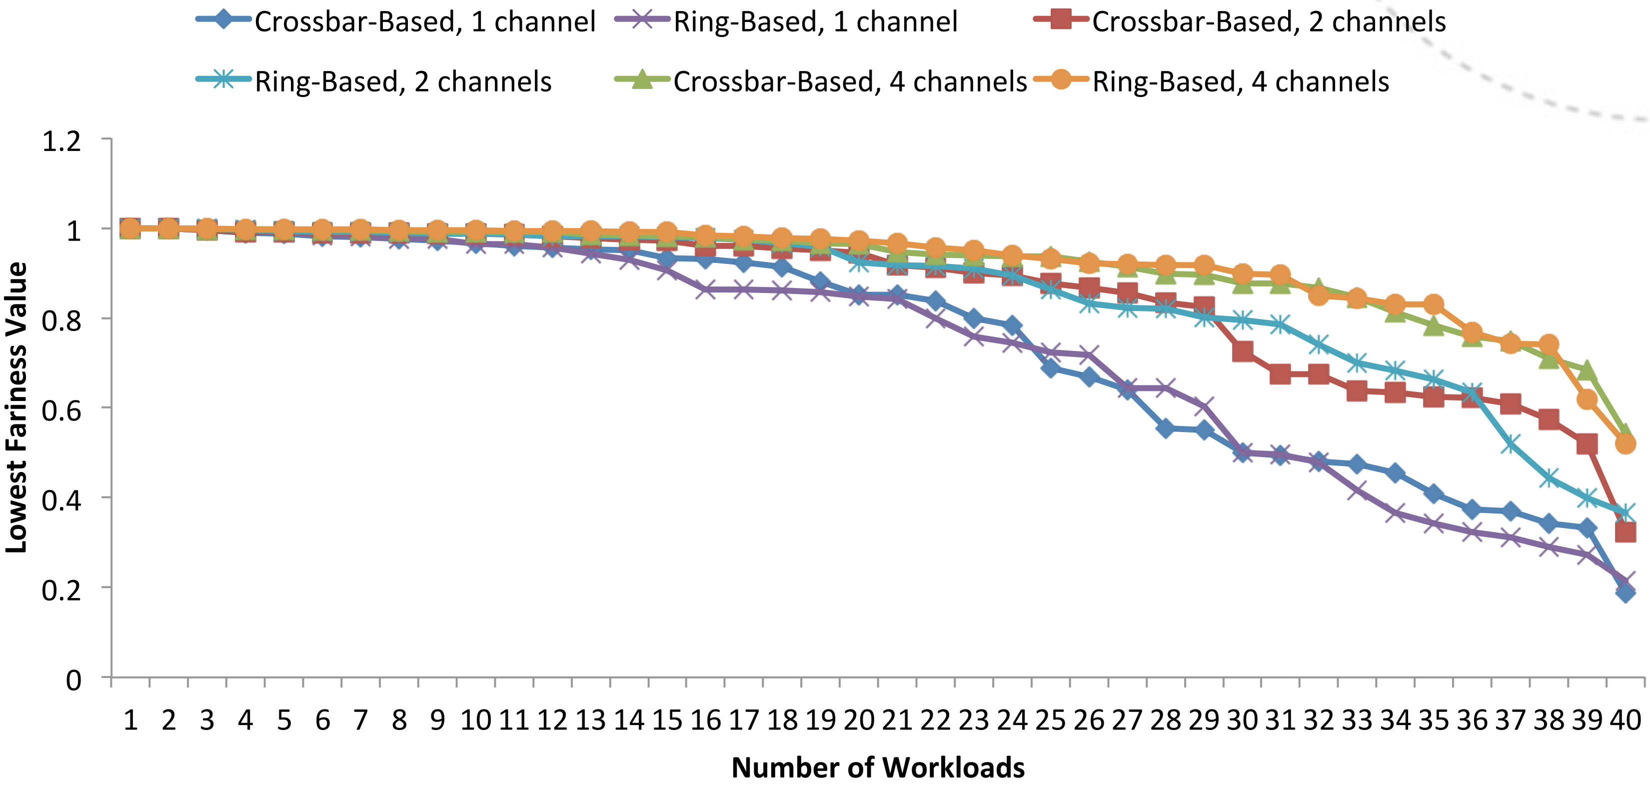

Aggregate Interference Impact

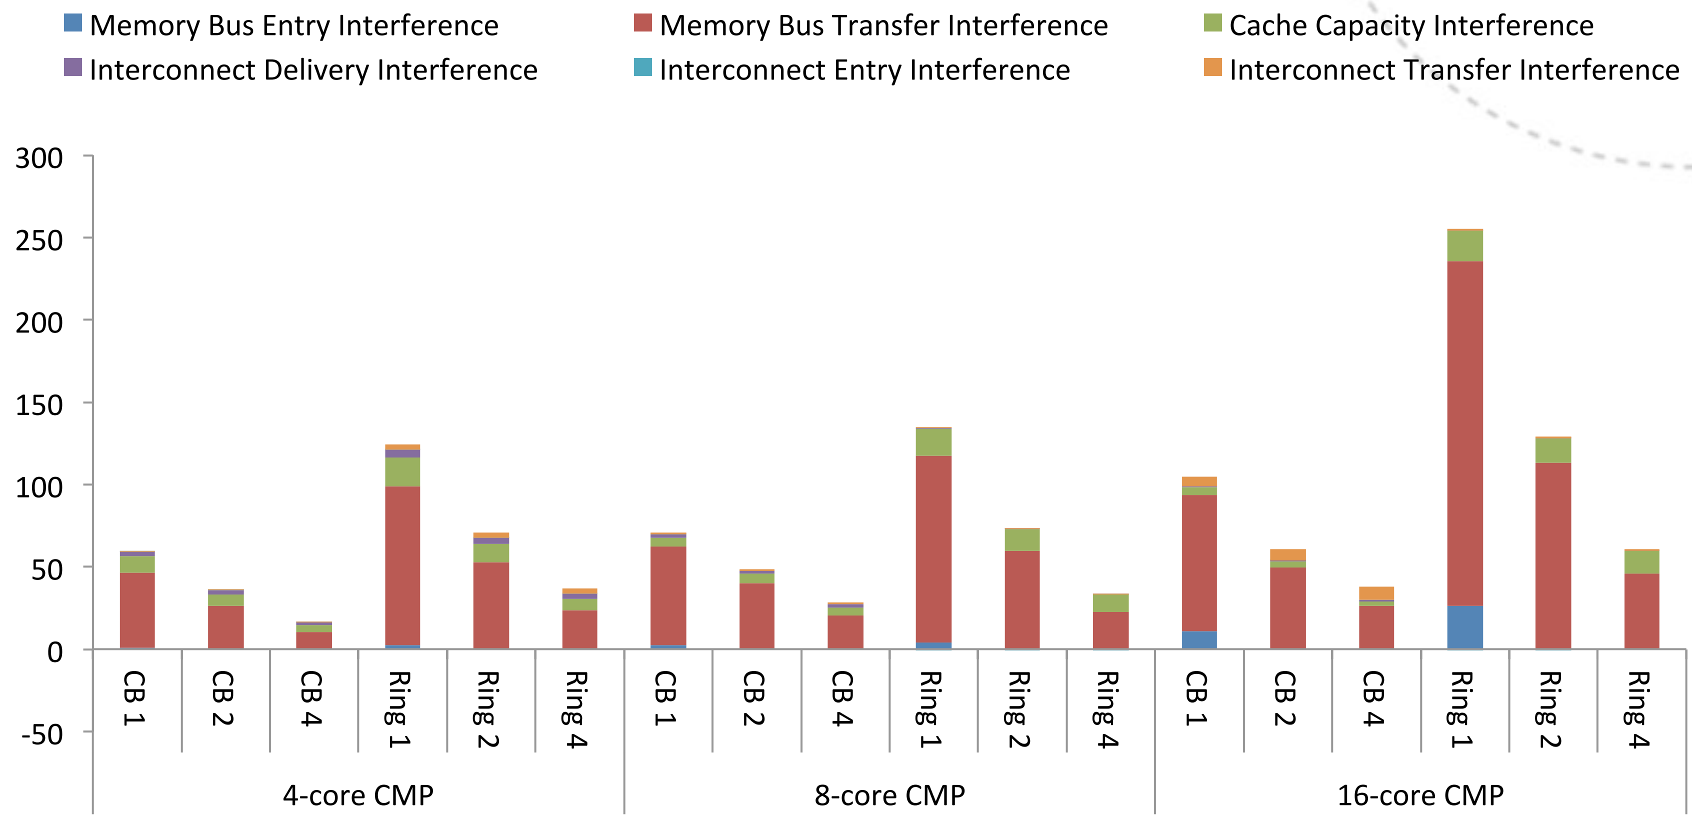

Main Findings

Impact of interference increases as more processor cores are added to CMPs

- Fairly minor point
- Many researchers did not believe this at that time

Competition for the memory bus is the major source of CMP interference

- Accounts for between 63% and 87% of the total interference impact

Other observations

- Competition for cache capacity can account for as little as 5% of the total
- Competition for shared buffers can account for as much as 11% of the total

Discussion points:

- **Was simulation the appropriate method to use for this study?**
- **How would you rate the *validity* of the results?**
- **How would you rate the *reliability* of the results?**
- **Are the experiments *reproducible*?**

Outline

- Introduction: Evaluation Requirements and Possibilities
- Computer Simulation
- Example Paper: «A Quantitative Study of Memory System Interference in Chip Multiprocessor Architectures»
- **How to get a PhD?**
- Final comments

How to get a PhD?

«The purpose of computing is insight, not numbers»
(Richard Hamming, 1915-1998)

Possible "research algorithm"

- Set up an experiment to study a phenomenon including previous work
- If one strategy is superior and better than previous work: Publish
- If trade-off: devise dynamic scheme and publish
- If you cannot beat previous best, improve model detail and start over
- If previous best is too good to beat, write a validation paper, find a new phenomenon and start over

Magnus' 10 Principles for Successfully Completing a PhD

Put in the hours, but manage your time

Be thorough, meticulous and organized in *all* research tasks (this forms the foundation for insight)

Hunt for insight:
"The key observation that led to the contributions in this work is …"

Always know the title of your next paper

Remember that research is not a team sport. Your individual expertise is always evaluated (for PhD, future employment, grant applications, etc.)

Wisely chosen collaborations is the key to significantly improving the quality of your work

Dynamically trade impact for quantity.
Quantity is easy to use to document progress, but impact will realize your career goals.

Always improve your methodology

If you invest significant work, make sure to publish (The Optimal Apple Cake Recipe vs. Towards an Optimal Apple Cake Recipe)

Get a hobby.

You will need something in your life that forces you to think about something else than research

Disclaimer: These views are based on my own experience and may not relate to your field of study. However, you are not allowed to dismiss the statements without articulating ***why*** *the statement does not apply to you.*

Outline

Introduction: Evaluation Requirements and Possibilities

Computer Simulation

Example Paper:
«A Quantitative Study of Memory System Interference in Chip Multiprocessor Architectures»

How to get a PhD?

Final comments

Final Comments

Some great papers are based on a great idea, but all great papers have an excellent methodology

- A methodology at the level of the current best practice is *necessary*
- A great idea/procedure/algorithm is *optional*

There is rarely a single "correct" methodology

- Choose the methodology best suited for you research project
- Research methods improve over time (i.e. if you don't improve, you are left behind)

Thank You!