



Co-sponsored by:

**COURSE ON ADVANCED COMPUTER ARCHITECTURE***Prof. Yale Patt, The University of Texas at Austin*

March 10-17, 2011, room C0.110, Science Park 904, University of Amsterdam

Using computers to solve problems requires starting with a natural language formulation of the problem and systematically transforming it until one has a machine language (ISA) specification of the problem (i.e., a program). This then is executed on the implementation hardware. As Moore's Law continues to provide more and more transistors on a chip (50 billion transistors running at a clock speed of greater than 10 GHz in a few years), application specialists continue to think up more and more applications that require these additional processing capability. The ISA is the interface between the software that produces a program and the hardware that carries it out. The ISA is implemented by a microarchitecture that is constrained by trade-offs such as performance, power consumption, cost, reliability, availability, etc. In this course, we will examine some of these tradeoffs, in light of the increasing transistor count and the fundamentals of microarchitecture. Our study will be grounded in fundamental principles of microarchitecture, but will involve current state-of-the-art approaches.

Course detailed program:**I. March 10 (13-17h):***Fundamental principles and tradeoffs.*

- A science of tradeoffs
- The transformation hierarchy
- The microarchitecture view
- The algorithm, the compiler, the microarchitecture
- The physical view
- Design points
- Numbers
- Thinking outside the box
- Evolution of the Process

II. March 11 (13-17h):*Emphasis on Tradeoffs*

- Alternatives in the ISA
- Alternatives in the microarchitecture
- Alternatives at the system level

III. March 14 (15-19h):*Run-time mechanisms*

- Branch Prediction
- Trace Cache
- SMT, SSMT
- L2 miss activity

Compile-time mechanisms

- Block-structured ISA
- Predication
- Fast track, slow track
- Wish branches (part by compiler,

part by microarchitecture)

- Braids
- Multiple levels of cache

IV. March 16 (15-19h):*Approaches to concurrency (Uniprocessor)*

- SIMD
- VLIW
- DAE
- HPS
- Data Flow

Approaches to concurrency (Multiprocessor)


- Tightly coupled vs. loosely coupled
- Metrics (speedups, efficiency, redundancy)
- Amdahl's Law
- Interconnection structures
- Memory Consistency
- Cache Coherence

Case studies

- Cell
- Power
- Niagara
- Larrabee vs Fusion vs GPUs

V. March 17 (13-17h):*Impact of multi-core*

- How we got here
- Multi-nonsense
- Where we go from here



Yale N. Patt received the BS degree from Northeastern and the MS and PhD degrees from Stanford, all in electrical engineering. He holds the Ernest Cockrell, Jr. Centennial Chair in Engineering and is professor of ECE at The University of Texas at Austin. He continues to thrive on teaching the large (400 students) freshman introductory course in computing and advanced graduate courses in microarchitecture, directing the research of eight PhD students, and consulting in the microprocessor industry. He is the co-author of Introduction to Computing Systems: from bits and gates to C and beyond (McGraw-Hill, 2nd edition, 2004), his preferred approach to introducing serious students to computing. He has received the highest honors in his field for both his research (the 1996 IEEE/ACM Eckert-Mauchly Award) and teaching (the 2000 ACM Karlstrom Outstanding Educator Award). More details can be found on his web page www.ece.utexas.edu/~patt. He is a Fellow of both the IEEE and ACM.

Interested? : please send an e-mail to Erik Hitipeuw (F.B.Hitipeuw@uva.nl) to register.