COMP/ELEC 525
Advanced Microprocessor Architecture

Prof. Scott Rixner
Duncan Hall 3028
rixner@rice.edu

January 9, 2007

Goals

- **Introduction to research topics in processor design**
- **Solid understanding of processor microarchitecture**
- **Exposure to systems simulation**

Logistics

Lectures:	T/Th 10:50-12:05 DH 1075
Lecturer:	Prof. Scott Rixner
Labby:	TBA
Grading:	20% Discussion Participation 20% Discussion Leading 60% Project
Readings:	Selected research papers from the literature Available on the web

Information

URL: http://www.owlnet.rice.edu/~elec525/

Check the web page for announcements

E-mail: rixner@rice.edu

Readings

- Each student will lead a class discussion
- Prepare a presentation of the papers
  - Overview
  - Critique
  - Open discussion questions
- Lead the ensuing discussion (this will be much easier if the discussion questions are well thought out!)
- All students expected to read all of the papers and contribute to the discussion after the presentation

Discussions

- Not a book report of the papers
  - Everyone should already have read them!
- What are the most important points?
  - What are the main ideas?
  - What are the papers' contributions?
  - What interesting things did you learn from them?
- How do these papers relate to everything else we've read?
- Critique
  - Is there reason to believe/disbelieve the conclusions of the authors?
  - Is their methodology reasonable?
  - Would you use their ideas in a real microprocessor?
  - etc.

Project

- **Work in groups**
- **Extension of issues discussed in class**
- **Requires a well-thought out hypothesis and evaluation**
- **Final result**
  - **Paper**
  - **Presentation in class**
- **Multiple project checkpoints and presentations**
- **Resulting papers/presentations will be posted on the web**

What Is Computer Architecture?

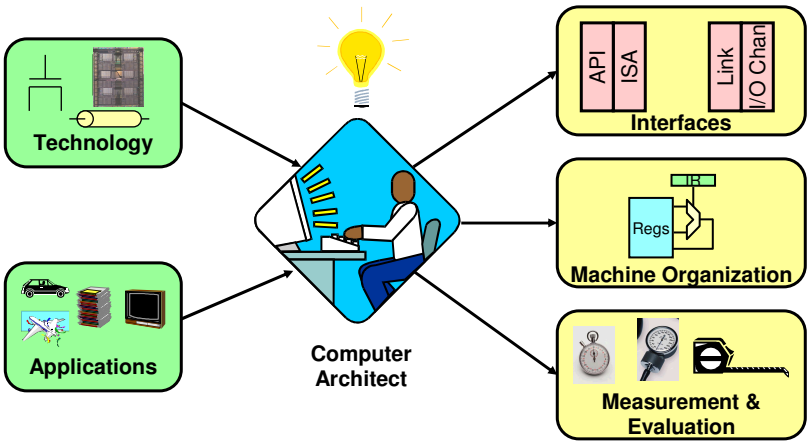

Applications Drive Machine Balance

- **Numerical simulations**
  - **Floating-point performance**
  - **Main memory bandwidth**
- **Transaction processing**
  - **I/Os per second**
  - **Integer CPU performance**
- **Decision support**
  - **I/O bandwidth**
- **Embedded control**
  - **I/O timing**
- **Media processing**
  - **Low-precision 'pixel' arithmetic**


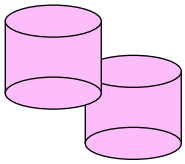

Processor Microarchitecture

- **Iterative process**
  - **Generate proposed architecture**
  - **Estimate cost**
  - **Measure performance**
- **Current emphasis is on overcoming sequential nature of programs**
  - **Deep pipelining**
  - **Multiple issue**
  - **Dynamic scheduling**
  - **Branch prediction/speculation**
  - **Multiple cores**

1971: Intel 4004

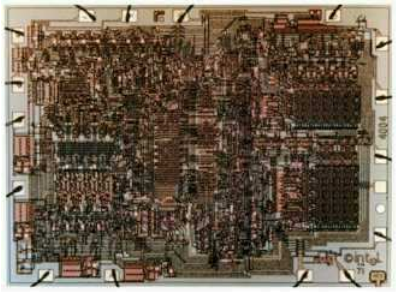

- **First microprocessor**
- **10 micron process**
- **2,300 transistors**
- **12 mm^2 die**
- **4-bit bus**
- **640 bytes of addressable memory**
- **750 KHz**

2000: Intel Pentium 4

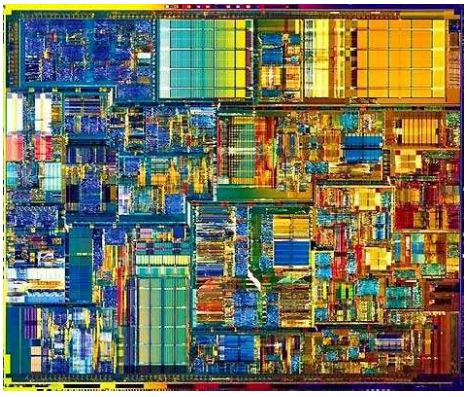

- **Issues up to 5 ops per cycle**
- **MMX, SSE, and SSE2**
- **0.18 micron process**
- **42 million transistors**
- **217 mm^2 die**
- **64-bit bus**
- **8KB data cache**
- **~12000 µop trace cache**
- **256KB L2 cache**
- **1.4 GHz**

2006: Intel Dual Core Itanium 2

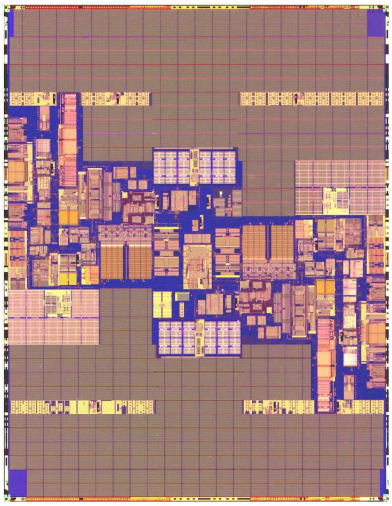

- 2 Itanium 2 processors
- Each core
  - 2-way multithreading
  - Issues up to 8 ops per cycle
  - 16KB inst. & data L1 caches
  - 1MB inst. & 256KB data L2 caches
  - 12MB L3 cache
- Virtualization technology
- 0.09 micron process
- 1.72 billion transistors
  - Cores: 57M
  - L1/L2 caches: 106.5M
  - L3 caches: 1550M
  - Bus logic and I/O: 6.7M
- 596 mm^2 die
- 128-bit bus
- 1.6 GHz

2015: Projected

???

- 0.025 micron process
- 7 billion transistors
- 310 mm^2 die
- 33 GHz (local on-chip)
- 29 GHz (off-chip)

Comparison

	1971	2006	2015
Technology (nm)	10000	90	25
Transistors (millions)	0.0023	1720	7000
Area (mm^2)	12	596	310
Clock Speed (MHz)	0.75	1600	33400

- **What were the trends in the past?**
  - **How did that affect microprocessor design?**
- **What do we expect in the future?**
  - **How will this affect microprocessor design?**

Chip Dimensions

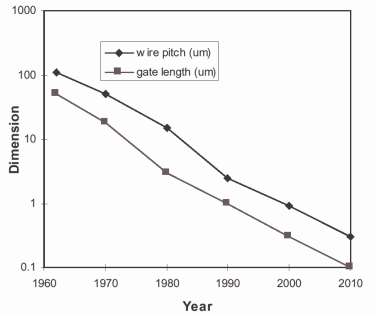

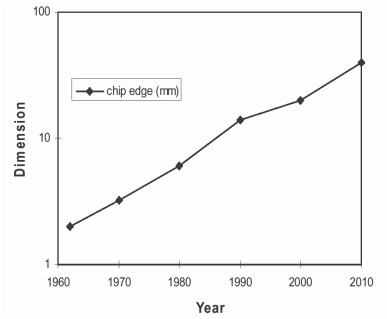

- **Feature size down 13% per year**
- **Edge length up 6% per year**

Available Area and Delay

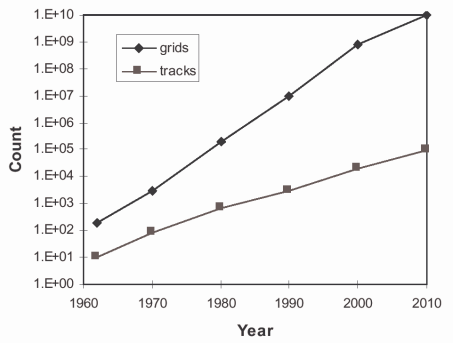

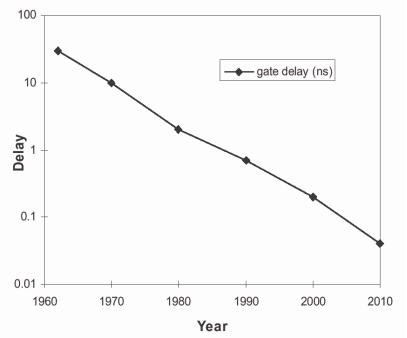

- **Grids up 50% per year**
- **Gate delay down 13% per year**

Chip Capability

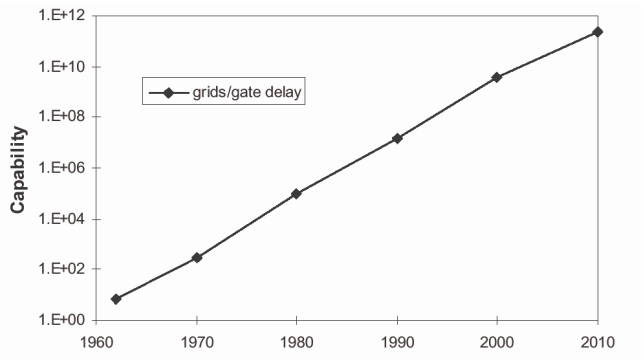

- **Capacity/delay up 70% per year**

Performance Trends

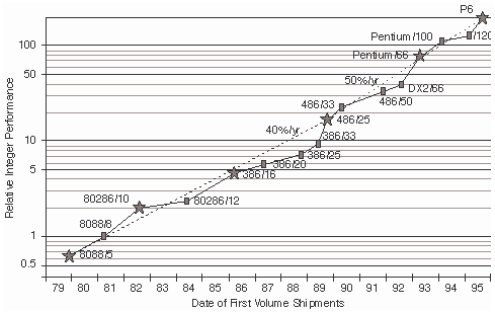

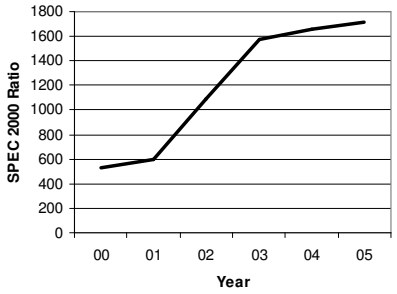

-Microprocessor Report 1/95

Wire Delay

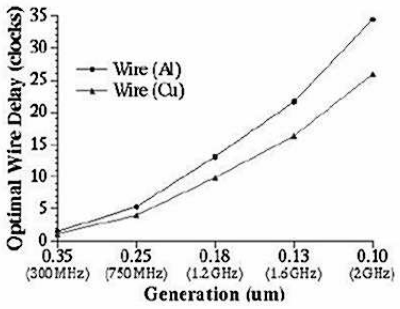

-Microelectronic Research Corporation

- **Wire delay up 100% per generation**
- **10's of clocks to cross chip in 0.1 micron process**

Pin Bandwidth

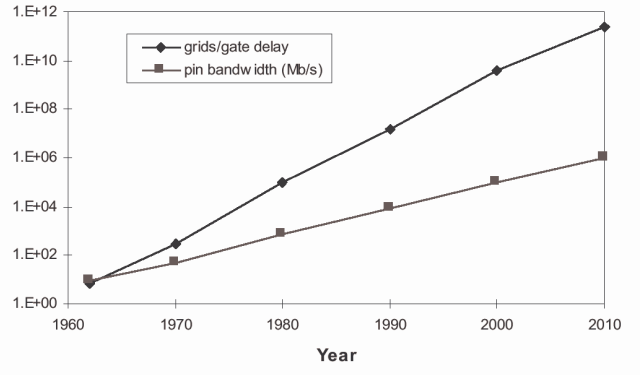

- **Capability up 70% per year**
- **Bandwidth up 28% per year**

Power

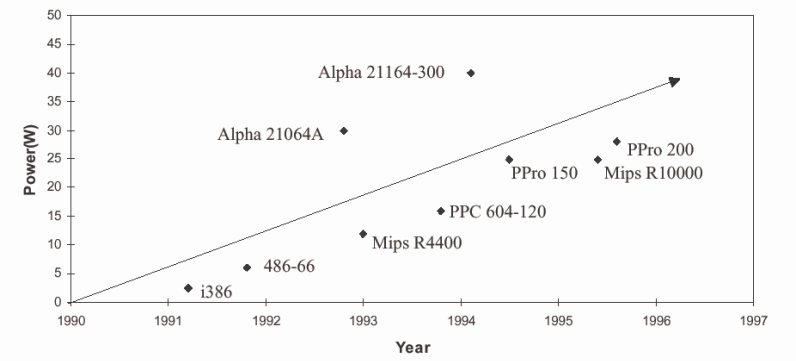

- **Power grows proportionally to performance**

DRAMs

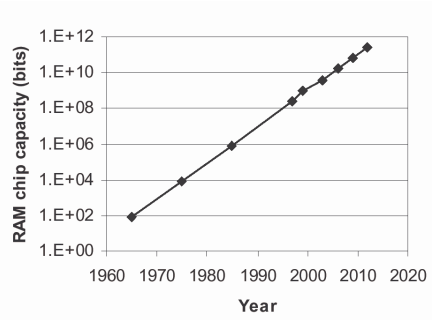

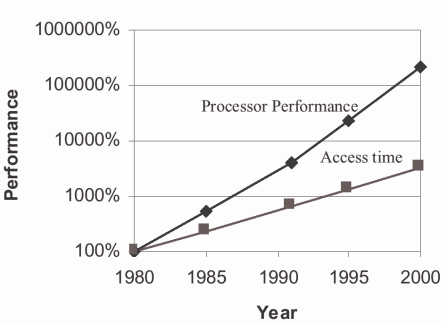

- **Capacity up 60% per year**
- **Access time up 19% per year**

Summary

- **Exponential improvement from technology**
- **Enormous capabilities**
- **Wire delay, pin bandwidth, power limited**
- **Impact on processor architecture?**

Other Influences

- **Compilers**
  - **Should the compiler or hardware do it?**
- **Compatibility**
  - **What software needs to run?**
- **Cost**
  - **Is the performance/functionality worth the cost?**

Areas of Focus

- **Front end**
- **Instruction-level parallelism**
- **Memory system issues**
- **Technology implications**
- **Novel processor architectures**

Front End

- **Delivery of instructions to arithmetic units**
- **Processor speed is irrelevant if you don't have operations to perform**
- **Issues**
  - **Limited memory/cache bandwidth**
  - **Small basic blocks**

Instruction-level Parallelism

- **Execute multiple, independent instructions in parallel**
- **Most commonly exploited type of parallelism**
- **Issues**
  - **Limited architectural state**
  - **Complex scheduling**
  - **Complex bookkeeping**
  - **Limited accessible ILP**

Memory System Issues

- **Storage and delivery of data for/to arithmetic units**
- **Processor speed is irrelevant if you don't have data to operate on**
- **Issues**
  - **Capacity**
  - **Speed**
  - **Bandwidth**

Technology Implications

- **How do technology improvements affect microarchitecture?**
- **Architectural scaling?**
- **Power dissipation?**
- **Clock speed?**
- **Wire delay?**

Modern Processor Architectures

- **What is industry/academia looking at right now?**
- **Do these architectures address technology issues?**
- **Do these architecture address changes in application domains?**

Next Class

- **Lecture/discussion**
- **Papers**
  - **"PC Processor Microarchitecture"**
  - **"2003 Technology Roadmap for Semiconductors"**
  - **"SPEC CPU2000: Measuring CPU Performance in the New Millenium"**