Principles of Parallel Computer Architecture

Fall 2014

Guang R. Gao

University of Delaware
http://www.udel.edu

Computer Architecture and Parallel Systems Laboratory
http://www.capsl.udel.edu

Outline

- Course Information
  - **Contact Information**
  - Class Schedule
  - Homework
  - Grading
  - Course Project
  - Course Requirements and Additional Materials
- Course Introduction

Contact Information

Instructors

Prof. Guang R. Gao

ggao.capsl(at)gmail.com

Office Hours:

TBA

or by appointment

Instructors

Stephane Zuckerman

szuckerman(at)udel.edu

Office Hours:

TBA

Or by appointment

Course Webpage:

http://www.capsl.udel.edu/courses/eleg652/2014/main.php?p=home

Contact Information

Teaching Assistant

Jaime

jaime(at)udel.edu

office Hours:

TBA

Or by appointment

Coordinator and Invited Lecturer

Dr. Haitao Wei

Whtaohust(at)gmail.com

TBA

Or by appointment

Course Webpage:
http://www.capsl.udel.edu/courses/eleg652/2014/main.php?p=home

Other Volunteers

Jose Monsalve

Aaron Landwehr

Acknowledgment

All the material of this course have been created under several years of hard work by many people. Special thanks to:

Dr. Joseph Manzano

Dr. Daniel Orozco

Dr. Elkin Garcia

And …

Outline

- Course Information
  - Contact Information
  - **Class Schedule**
  - Homework
  - Grading
  - Course Project
  - Course Requirements and Additional Materials
- Course Introduction

BEFORE WE BEGIN:

When is the right time to learn about something?

Now!

Approximate Class Schedule

Topic	Week	Date
Introduction + Vector Processing	1	08/26
Vector Processing	2	09/02
Parallel Programming	3	09/09
Instruction Level Parallelism and Superscalar Architectures	4	09/16
Shared Memory Architectures	5, 6, 7	09/23, 09/30, 10/07
Networks	8	10/14
GPUs	9	10/21
Execution Models + Dataflow	10, 11	10/28, 11/04
Manycore and Multicore Architecture Era	12, 13	11/11, 11/25
Final Project + Final Exam	14	12/04

Outline

- Course Information
  - Contact Information
  - Class Schedule
  - **Homework**
  - Grading
  - Course Project
  - Course Requirements and Additional Materials
- Course Introduction

We like Homework!

Linpack: The de facto standard to measure supercomputer performance.

We will measure the performance of 3 machines:

- Any parallel machine that you have access to.
- cluster.capsl.udel.edu
- Your personal machine.

You may start now. All the required information is publicly available.

Learning Dynamics

Start homework as early as possible.

The assignments are designed to broaden the horizon of your understanding. Research any knowledge you lack to do the homework.

- Example: Homework 1 requires knowledge on certain topics. It is your responsibility to do the required reading.

Know your responsibilities, do not expect reminders.

Outline

- Course Information
  - Contact Information
  - Class Schedule
  - Homework
  - **Grading**
  - Course Project
  - Course Requirements and Additional Materials
- Course Introduction

Grades

Activity	Percentage
Homework	25%
Quiz 1 (Middle of the semester)	20%
Quiz 2 (near or on Last class)	25%
Project	25%
Participation	05%

Outline

- Course Information
  - Contact Information
  - Class Schedule
  - Homework
  - Grading
  - **Course Project**
  - Course Requirements and Additional Materials
- Course Introduction

The project

An interesting question will be presented.

The objective is to propose solutions and to design experiments to address the question.

The results will be presented at the end of the semester.

The project topics will be discussed soon.

Outline

- Course Information
  - Contact Information
  - Class Schedule
  - Homework
  - Grading
  - Course Project
  - **Course Requirements and Additional Materials**
- Course Introduction

What do I need? (Besides what you learn here)

Access to a computer, with internet access, where you can install software.

Understanding on how to run and compile a C program.

Understanding on how to use a Linux/Unix system and SSH

This is the moment to seek help!
If you do not have the tools, let us know NOW!

The Reference Books

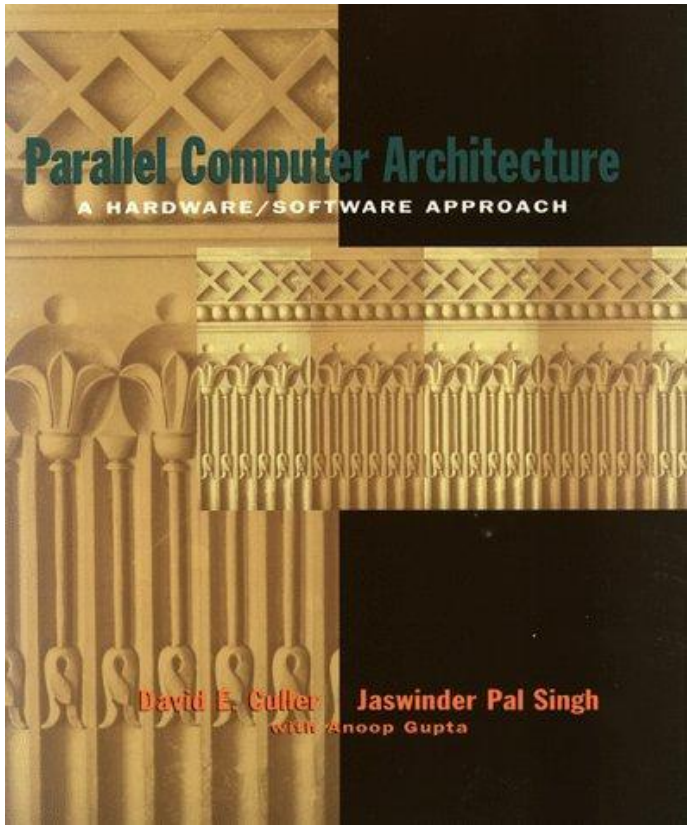

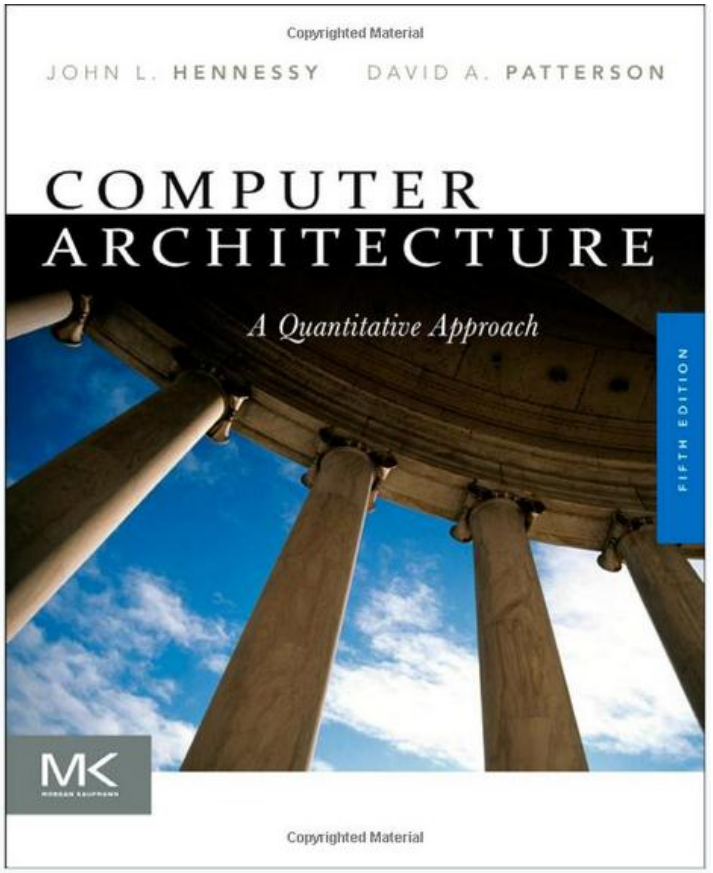

Textbooks complement the class material, they are excellent for better understanding.

However, the class material comes from many sources such as experience of the instructors, research papers, research done by other groups and so on.

Other sources of information

Journals

IEEE Computer

IEEE Transactions in Computers

IEEE Transaction on Parallel and Distributed Systems

Conference Proceedings

PACT	Parallel Architectures and Compilation Techniques
MICRO	ACM/IEEE Symposium on Micro-Architectures
HPCA	ACM/IEEE Symposium High Performance Computer Architecture
ISCA	International Symposium on Computer Architectures
PLDI	International Symposium on Parallel Language Design and Implementation

Others ?

Outline

- Course Information
- **Course Introduction**
  - Parallel Architecture
  - Moore's Law and Architectural Trends
  - References

INTRODUCTION

Course Introduction

The Role of a Computer Architect

Maximize *Productivity* and *Performance*

Productivity = Programmability and a reduction in development time

Performance = "Reasonable" Throughput given technology and cost limitations

Parallelism

Two or more tasks may execute at the same time

Alternative to higher frequency clocks

Applies to all levels of computer design

Importance has been constantly raising since several "walls" were hit

In the near future, it will be become the paradigm on all aspects of computing

Outline

- Course Information
- Course Introduction
  - **Parallel Architecture**
  - Moore's Law and Architectural Trends
  - References

Parallel Computer Architecture (?)

Kraken Supercomputer at OakRidge National Laboratory.

Answers the question:

How do we **build** and **use** a Parallel Digital Computer?

Parallel Computer: A system with several processing elements.

Why Parallel Architectures?

Q: How do you make your program run faster?

Traditional Answer:

Wait for a faster processor to be released!

For 60 years, the average programmer has studied and produced **serial** programs. In most cases, with no knowledge at all of what parallel programming is….

However, no significantly faster processors are being released any more.

There is no free lunch!

If you want your programs to be faster, you have to go parallel.

10 Years with the same clock rate!

The happy times: 1950 to 2000 where every 10 years a 10X speedup was experienced.

Processor	Clock Speed	Date
4040	740kHz	1974
386	20Mhz	1985
486	100Mhz	1989
Pentium	300MHz	1993
Pentium II	450MHz	1997
Pentium III	1.4GHz	1999
Pentium IV	3.8GHz	2000
Core Duo	2.33GHz	2006
Core 2	3.33GHz	2006
Core i7	3.47GHz	2008

There is a limit to serial performance

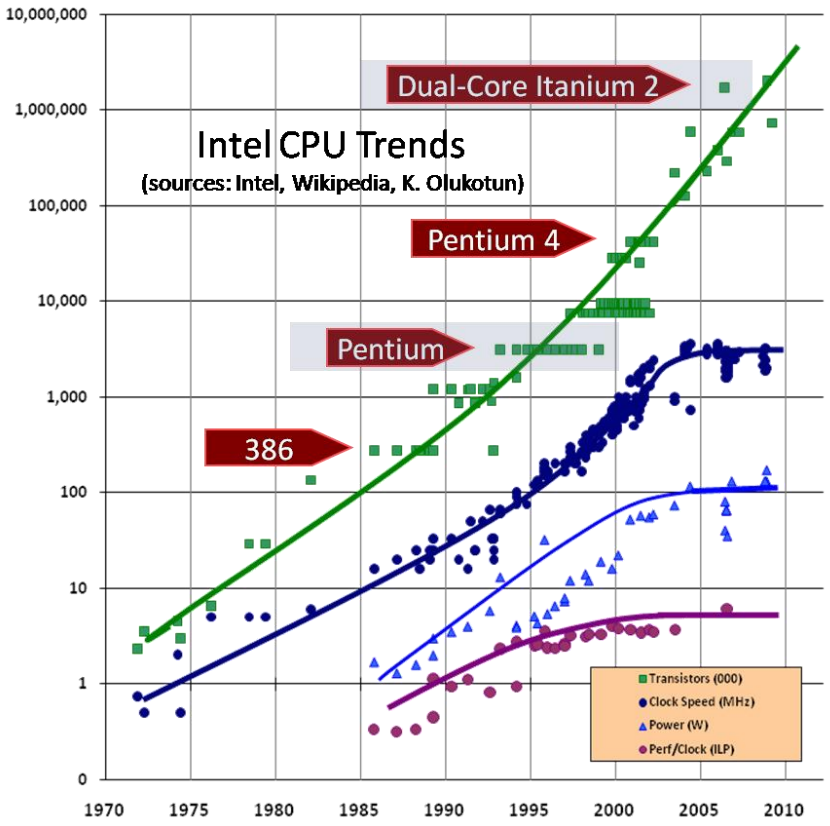

We face serious problems, called "Walls":

Frequency Wall

Memory Wall

Power Wall

ILP Wall

We will talk about each one.

Serial is dead. Long live Parallel!

Most consumer electronics have some form of parallel architecture inside of them

Reasons for the Change

An evolutionary change in computing due to:

Technology		Architecture	
Decrease in feature size Allowing more components into a chip	More and more performance and power hungry applications	Effectively organizing components to maximize uses of resources and minimizing damaging side effects	Find Cost Effective ways to get the desired performance out of the given Hardware / Software combo
	Applications		**Economics**

Outline

- Course Information
- Course Introduction
  - Parallel Architecture
  - **Moore's Law and Architectural Trends**
  - References

Moore's Law

The complexity for minimum component costs has increased at a rate of roughly a factor of two per year...

Certainly over the short term this rate can be expected to continue, if not to increase.

Over the longer term, the rate of increase is a bit more uncertain, although there is no reason to believe it will not remain nearly constant for at least 10 years.

That means by 1975, the number of components per integrated circuit for minimum cost will be 65,000.

I believe that such a large circuit can be built on a single wafer.

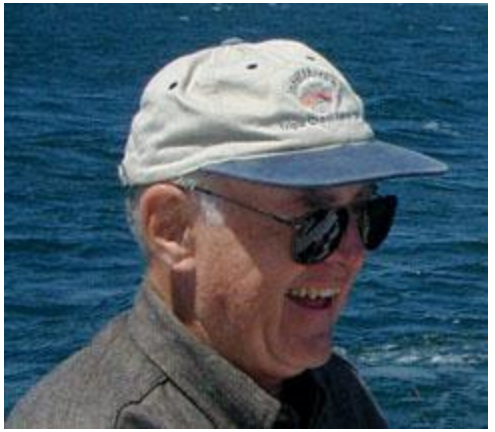

The experts look ahead

Cramming more components onto integrated circuits

With unit cost falling as the number of components per circuit rises, by 1975 economics may dictate squeezing as many as 65,000 components on a single silicon chip

By Gordon E. Moore

Director, Research and Development Laboratories, Fairchild Semiconductor division of Fairchild Camera and Instrument Corp.

The future of integrated electronics is the future of electronics itself. The advantages of integration will bring about a proliferation of electronics, pushing this science into many new areas.

Integrated circuits will lead to such wonders as home computers—or at least terminals connected to a central computer—automatic controls for automobiles, and personal portable communications equipment. The electronic wristwatch needs only a display to be feasible today.

But the biggest potential lies in the production of large systems. In telephone communications, integrated circuits in digital filters will separate channels on multiplex equipment. Integrated circuits will also switch telephone circuits and perform data processing.

Computers will be more powerful, and will be organized in completely different ways. For example, memories built of integrated electronics may be distributed throughout the machine instead of being concentrated in a central unit. In addition, the improved reliability made possible by integrated circuits will allow the construction of larger processing units. Machines similar to those in existence today will be built at lower costs and with faster turn-around.

Present and future

By integrated electronics, I mean all the various technologies which are referred to as microelectronics today as well as any additional ones that result in electronics functions supplied to the user as irreducible units. These technologies were first investigated in the late 1950's. The object was to miniaturize electronics equipment to include increasingly complex electronic functions in limited space with minimum weight. Several approaches evolved, including microassembly techniques for individual components, thin-film structures and semiconductor integrated circuits.

Each approach evolved rapidly and converged so that each borrowed techniques from another. Many researchers believe the way of the future to be a combination of the various approaches.

The advocates of semiconductor integrated circuitry are already using the improved characteristics of thin-film resistors by applying such films directly to an active semiconductor substrate. Those advocating a technology based upon films are developing sophisticated techniques for the attachment of active semiconductor devices to the passive film arrays.

Both approaches have worked well and are being used in equipment today.

The author

Dr. Gordon E. Moore is one of the new breed of electronic engineers, schooled in the physical sciences rather than in electronics. He earned a B.S. degree in chemistry from the University of California and a Ph.D. degree in physical chemistry from the California Institute of Technology. He was one of the founders of Fairchild Semiconductor and has been director of the research and development laboratories since 1959.

Electronics, Volume 38, Number 8, April 19, 1965

Moore's Law

More transistors as time goes by.

Does not talk about:
- Speed
- Power
- Parallelism

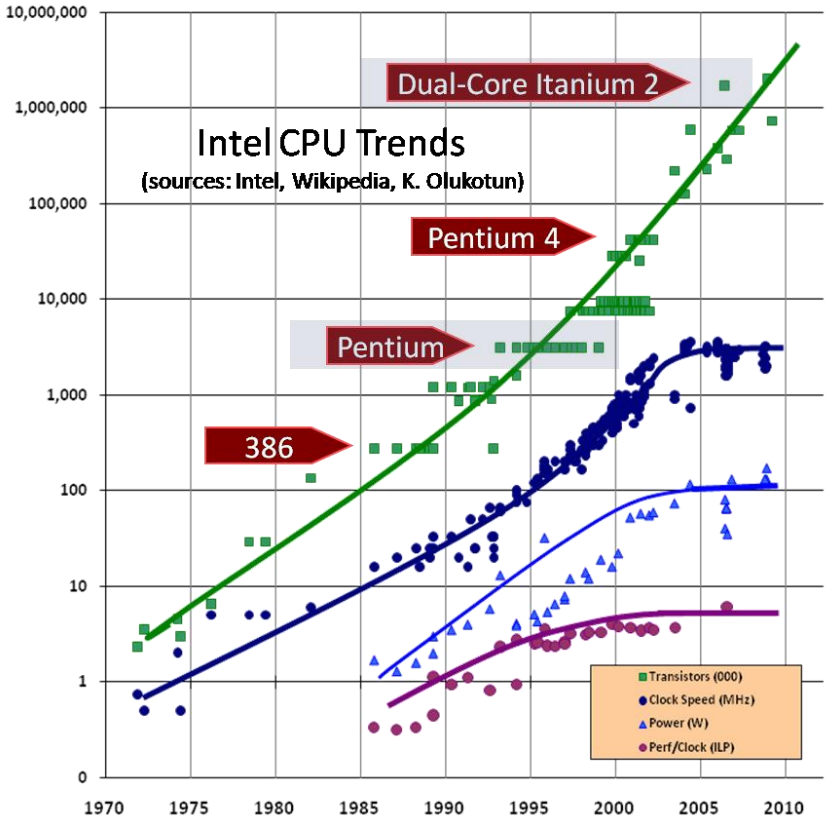

Architectural Trends

Designed for performance

Higher Frequency == Higher Performance **?**

We know many tricks, but they are not enough!

Hide Latencies at all cost! (DMA, scheduling...)
Branch prediction
Prefetching
Dynamic Power Management
Cache Blocking

Past solutions don't solve present problems.

The field is experiencing something similar.
People try to develop the same tricks, but they don't work anymore.

Technology Progress Overview

Computation is almost free
BUT
bandwidth is very expensive

Outline

- Course Information
- Course Introduction
  - Parallel Architecture
  - Moore's Law and Architectural Trends
  - **References**

References

Images from:
http://www.gotw.ca/publications/concurrency-ddj.htm
Wikipedia

Moore's Law Paper:
https://chemicalheritage.org/pubs/moores_law/Moore-Chap-05.pdf

The Kraken figure is in the public domain. Taken from Wikipedia

Moore's Image Copyright: Steve Jurvetson. From
http://www.flickr.com/photos/jurvetson/368370/

References

Moore's Law Figure from Herb Sutter. Taken From http://www.gotw.ca/publications/concurrency-ddj.htm