

# AMATH 483/583

# High Performance Scientific Computing

## **Lecture 1: Introduction and Overview**

Xu Tony Liu, PhD

Paul G. Allen School of Computer Science & Engineering

University of Washington

Seattle, WA

# Overview



- Hello class!
- Course administration and mechanics
- HPC: past, present, future
- Tour of course topics
- Code development
  - C++
  - Docker
  - bash

# Course Essentials

- AMATH 483 A
- AMATH 583 A B C(EDGE/online) E
- Tu/Th 12:00pm-1:20pm
- MUE 153
- <https://amath583.github.io/sp22/>
  
- Prerequisites: AMATH 301 or CSE 142
  - Some experience programming (C, C++, Python, Matlab)
- Course text: Parallel Programming: Concepts and Practice, Bertil Schmidt, Jorge Gonzalez-Dominguez, Christian Hundt, Moritz Schlarb.

# Suggested Course Texts

- Course texts: Schmidt et al, Mattson et al
- Links in “Resources”



# Canvas

This will take you to Github course site

## AMATH 483 A Sp 22: High-Performance Scientific Computing

Jump to Today

### Course website

The bulk of the course materials, including lectures, notes, and assignments, are hosted on the course website. If you find any broken links or missing content, please let us know. Your feedback is really appreciated.

Currently, the Panopto lectures from the previous years (Spring 2019) is available in the Github course website. The Panopto lectures will be updated accordingly.

Follow [this link](#) to access the Github Course Site" item in the left-hand column. Complete materials from [Spring 2019](#) and from [Spring 2021](#) are also available.

Recorded lectures

Also via podcast

Sign up!

More about this in a minute

- W canvas
- Account
- Dashboard
- Courses
- Calendar
- Inbox
- People

- AMATH 483 A > Syllabus
- Spring 2022
- Home
- Syllabus
- Panopto Recordings
- Class Notebook
- Piazza
- Grades
- Discussions
- Collaborations
- People

# Course Materials on Github

- PDF versions of the slides will be posted in advance of lecture
- Recordings of lecture are available online 90 minutes after lecture (via panopto / canvas – links will also be on course website)
- Subscribe to the podcast!

# Your Instructional Team

- Xu Tony Liu, [x0@uw.edu](mailto:x0@uw.edu)
  - Saba Heravi, [heravi@uw.edu](mailto:heravi@uw.edu)
  - Michael Kupperman, [kupperma@uw.edu](mailto:kupperma@uw.edu)
- 
- Office hours (All times PDT) and locations post on course website
  - Mon 3:00pm - 5:00pm @Zoom, link posted on Canvas
  - TTh 2:00pm - 3:00pm LEW 315
  - WF 4:00pm - 5:00pm @Zoom, link posted on Canvas

# More about me

- PhD at EECS Washington State U, MS at CS Indiana University
- Distinguished Graduate Research Fellow at Pacific Northwest National Lab
- Research Scientist & Lecturer at UW
- Research interests: High Performance Computing, scalable graph analytics, graph and hypergraph theory
  
- To learn with you
- To learn from you

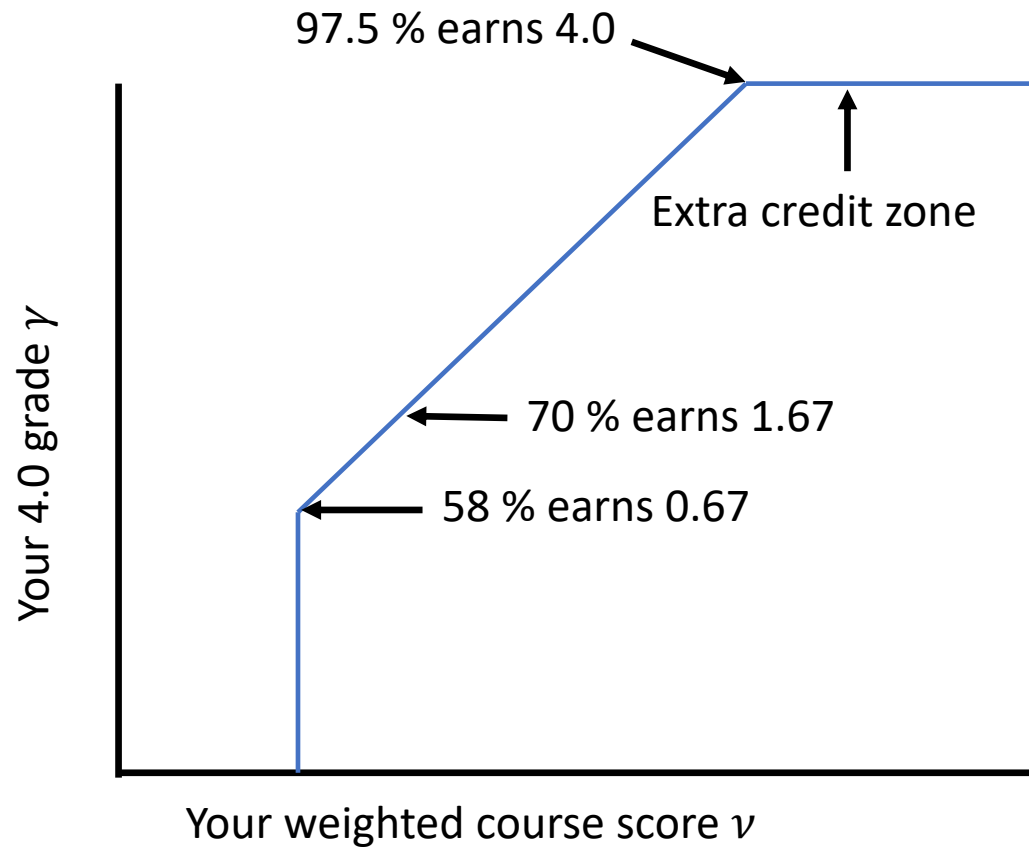


# Course Mechanics

- 8 problem sets (60% of your grade, lowest score dropped)
- 2 take home exams (mid-term and final 20% of your grade each)
- 20% penalty per late day (with 4 grace days)
- One “challenge flag”
- Piazza for course discussions, Q/A
- Gradescope for grading
  
- See the course syllabus linked on the course web site
  
- When in doubt – ask!

# Grades

How your 4.0 grade is computed



$$\gamma = \frac{4.2\nu}{50} - 4.2$$

# Computing Resources

- Your laptop
- Linux or Linux-like development environment
  - Docker (supported)
  - Mac OS X
  - Windows subsystem for Linux
- Hyak
  - GPU and Message Passing Interface (MPI)
- (See course web page for more info)

# Academic Integrity

- You are being evaluated in this course for how much **you** learn
- Not for someone else's work
- You may not claim someone else's work as your own (plagiarism)
- You may use any source you like for your work (with limits on AMATH 483/583 classmates)
- But ***you must cite your sources***
- Penalty for plagiarism is zero score on entire problem set
  - Copying something if you say you copied it is not plagiarism
  - (Though you may not get full credit, you won't get the plagiarism zero)

# What's wrong with this picture?



# What's wrong with this picture?



# Technology

- Laptop use permitted in class (provisionally)
  - **MUTE** your laptop
- **PROVIDED**
- The class creates and maintains a course notebook via OneNote (cf. course canvas page)

# Course Philosophy

- Most of your learning will take place doing problem sets
- Learner-centered approach (learning outcomes)

**Hardware**



**Software**





# What This Course is About

- How algorithms, data, software, and hardware interact to affect performance (and how to orchestrate them to get high performance)
- At the completion of this course, you will be able to
  - Write software that fully utilizes hardware performance features
  - Describe the principal architecture mechanisms for high performance and algorithmic and software techniques to take advantage of them
  - Recognize opportunities for performance improvement in extant code
  - Describe a strategy for tuning HPC code
- Today and years from now

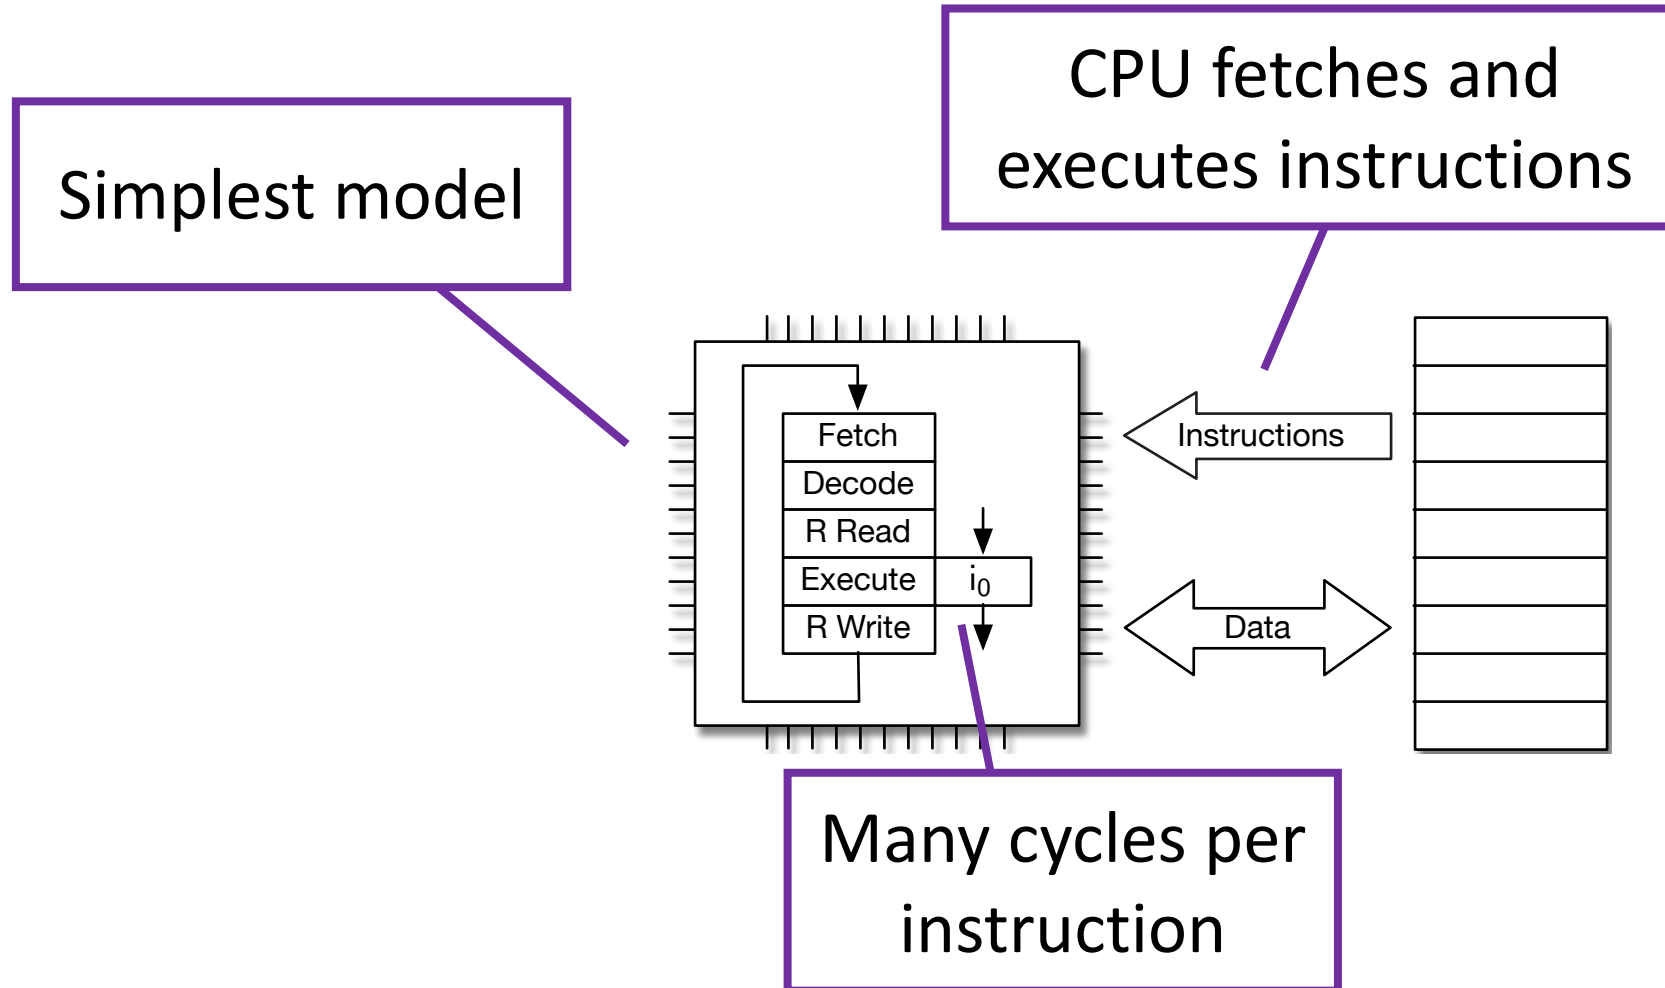
# What this Course is not About

- Not a software engineering course (but you will learn some basics)
- Not a programming course (ditto)
- Not an architecture course (but you will learn essential models)
- Not a parallel programming course
- Not an operating system course
  
- But you **will** learn essentials in each of these areas – and more importantly, how they **interact** to affect (and effect) performance
- (There are entire courses on each of these topics)

# The HPC Canon (as of 2022)

Technology	Paradigm	Hammer
CPU (single core)	Sequential	C compiler
SIMD/Vector (single core)	Data parallel	Intrinsic
Multicore	Threads	pthread library
NUMA shared memory	Threads	pthread library
GPU	GPU	CUDA
Clusters	Message passing	MPI

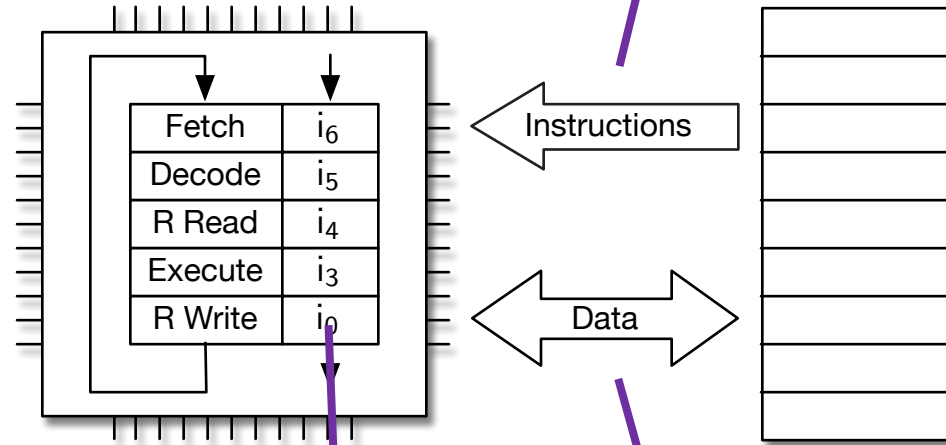
# Scaling progression of CPUs



# Pipelining

Pipelining

Instructions are fetched in a stream

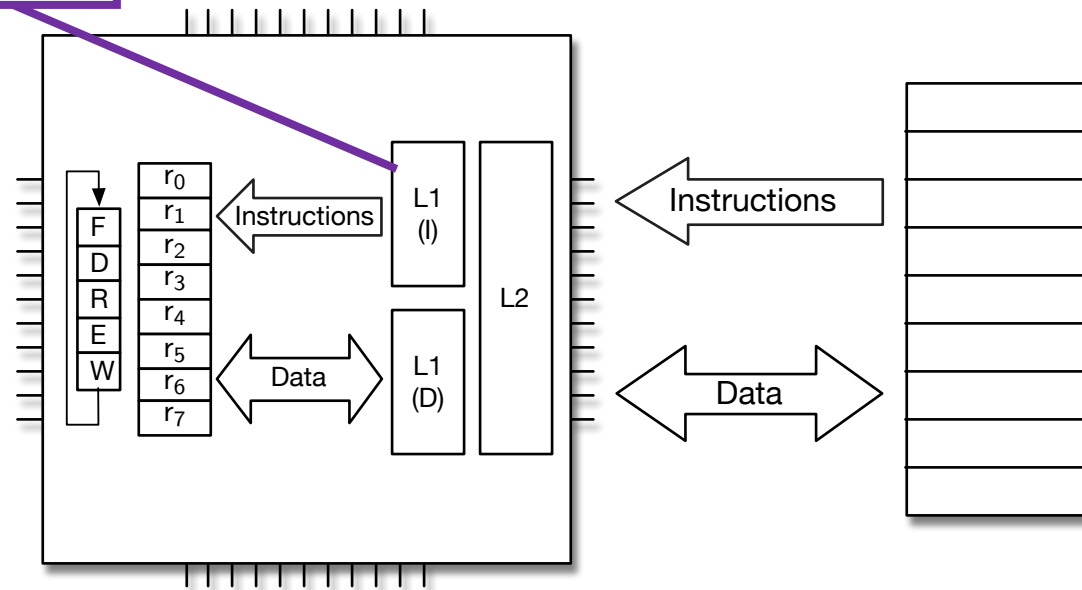


Processed in a pipeline

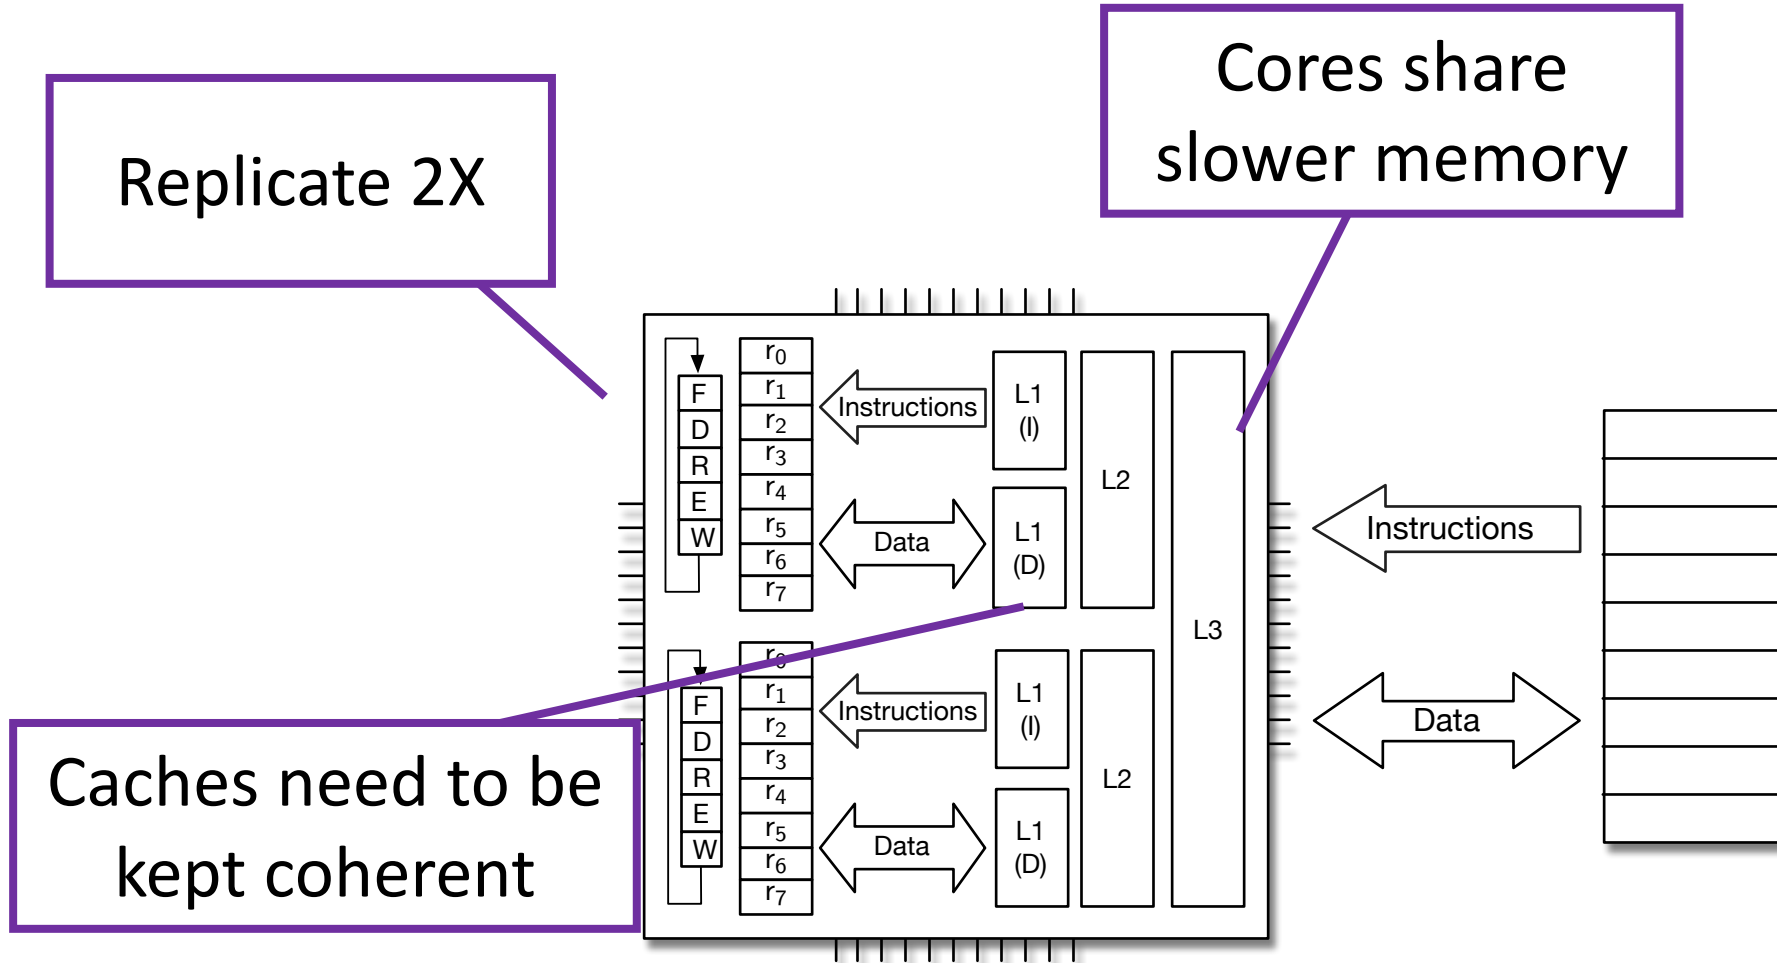
A long trip from memory

# Hierarchical Memory

Use special, fast memory to keep data and instructions close



# Multicore CPUs



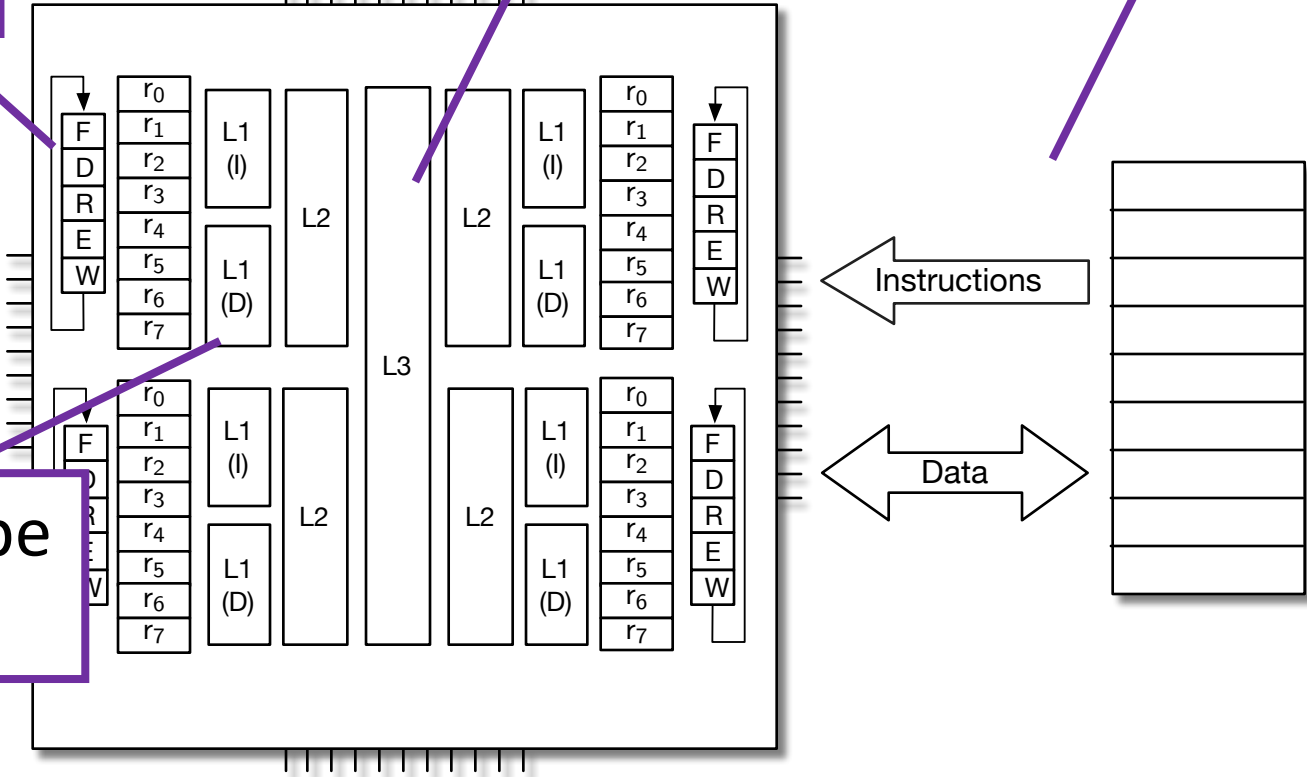
# Even more cores

Replicate 4X

Cores share slower memory

Include super-slow DRAM

Caches need to be kept coherent



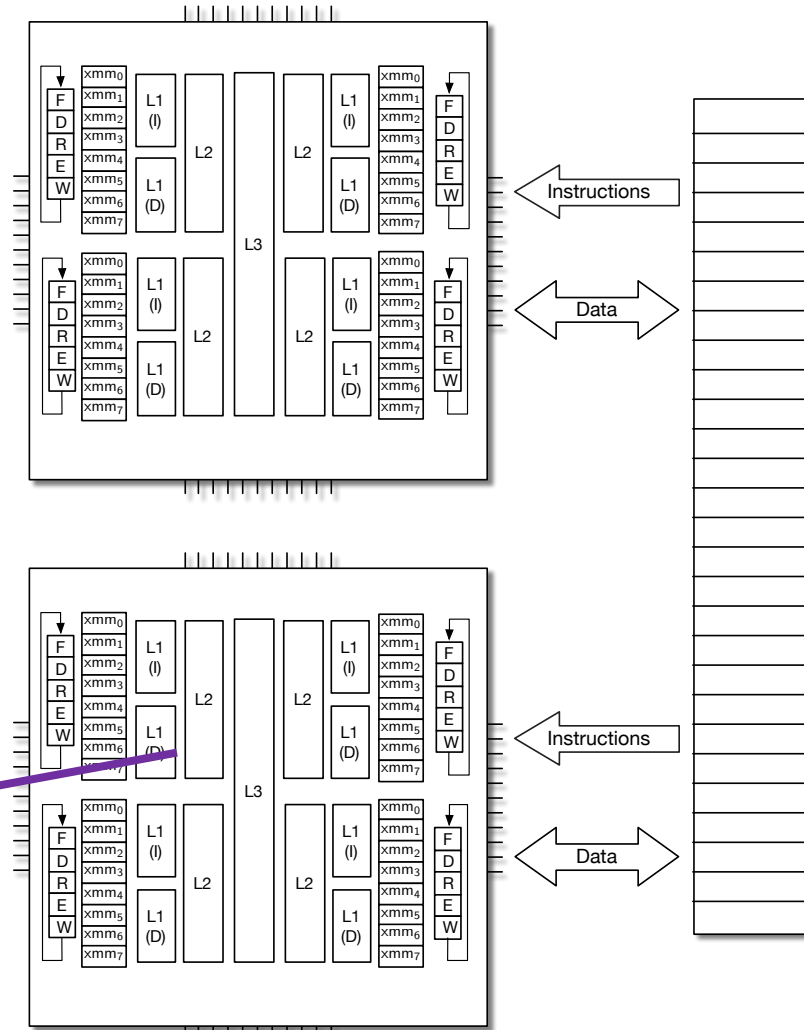


# Symmetric Multi-Processor (SMP)

Multiple CPU chips

AKA "sockets"

Caches still need to be kept (somewhat) coherent



Memory may be uniformly shared among sockets

Uniform memory access (UMA)

# Asymmetric Multi-Processor


Multiple CPU chips

AKA "sockets"


Caches still need to be kept (somewhat) coherent: CC-NUMA

Memory may be non-uniformly shared among sockets

Non-uniform memory access (NUMA – most common)



# GPU



# The Next Step


Put sockets on a blade

Put blades in a chassis

Put chassis in a rack

Put racks in a center

Put centers in the cloud



# Then you have a supercomputer

But how do you use it?




The HP (Build on each of 2022)

Build on each of other

We will be following this path this quarter

Technology	Paradigm	Language
CPU (single core)	Sequential	C++
SIMD/Vector (single core)	Data parallel	
Multicore	Threads	
NUMA shared memory	Threads	
GPU	GPU	
Clusters	Message passing	MPI



Order of evolution (more or less)

Technology and paradigm

# Tour of the Course (HPC hardware)

- Basic CPU machine model
  - Hierarchical memory (registers, cache, virtual memory)
  - Instruction level parallelism
  - Multicore processors
  - Shared memory parallelism
  - GPU
  - Distributed memory parallelism
- 
- Use running examples



By Hteink.min - commons:File:Louvre Pyramid.jpg, CC BY-SA 3.0, <https://en.wikipedia.org/w/index.php?curid=38292385>

# Tour of the Course (HPC Software)

- Elements of C++
- Elements of software organization
- Elements of software practice
- Elements of performance measurement and optimization

**Hardware**



**Software**



# Computing is Indispensable to Science and Engineering

- The 3<sup>rd</sup> (and 4<sup>th</sup>?) pillar(s)
- Can carry out investigations where physical experiments would be too fast, too slow, too hot, too cold, too costly, too dangerous, etc
- Examples: Weather, climate, fusion, crash testing, etc.
- HPC means more and better scientific discovery
- Better world, survival of the planet



# How to Learn





- Ask yourself at the beginning of this quarter
  - What scientific problem you want to solve?
  - What do you want to learn from this course that can prepare you?
  - Can you write a sequential program to solve it?
  - What is the performance of it?
  - ...
  
- Ask yourself at the end of this quarter
  - Do you master the skillset to solve your scientific problem?
  - Can you write a high-performance program to solve it?
  - What is the performance of it?
  - What is the speedup?
    - compare your sequential program with your high-performance program

# Editorial Comment







- The most exciting phrase to hear in science, the one that heralds new discoveries, is not “Eureka!” (I found it) but “That’s funny”
  - Attributed to Isaac Asimov (and others)

# Discovery Science (DOE)



# Shock Wave Processing of Advanced Reactive Materials



# Uses of HPC (a sample)


- Cosmology
- Earthquake
- Weather
- Climate modeling
- Automobile crash testing
- Aircraft design
- Jet engine design
- Stockpile stewardship
- Nuclear fusion
- Protein folding
- Modeling the brain
- Modeling bloodstream
- Epidemiology
- Rendering (CGI)
- Sigint
- Block chains
- Gene sequencing
- Etc

# Name this Famous Person



# Historical Trends in Computing


## Where Does High Performance Come From?



By AI.Graphic (Own work) [CC BY-SA 3.0 (<http://creativecommons.org/licenses/by-sa/3.0>)], via Wikimedia Commons





# Name This Famous Person





# Supercomputers Then and Now



# End of Moore's Law



# Where Does High Performance Come From?



By AI.Graphic (Own work) [CC BY-SA 3.0 (<http://creativecommons.org/licenses/by-sa/3.0>)], via Wikimedia Commons

# Name this Famous Person



# Supercomputers Then and Now



# Then (20 years ago)

Rank	System	Cores	Rmax (GFlop/s)	Rpeak (GFlop/s)	Power (kW)	
1	ASCI Red Intel	7,264	1,068.0	1,453.0		
2	Center for Computations of Tsukuba Japan	6/2048 Tsukuba	2,048	368.2	614.4	
	of Japan	Numerical Wind Tunnel Fujitsu	167	229.0	281.3	498
		SR2201/1024 Hitachi	1,024	220.4	307.2	



4,500 Gflops



700 Gflops

# Top500 as of Nov 2021 (top500.org)

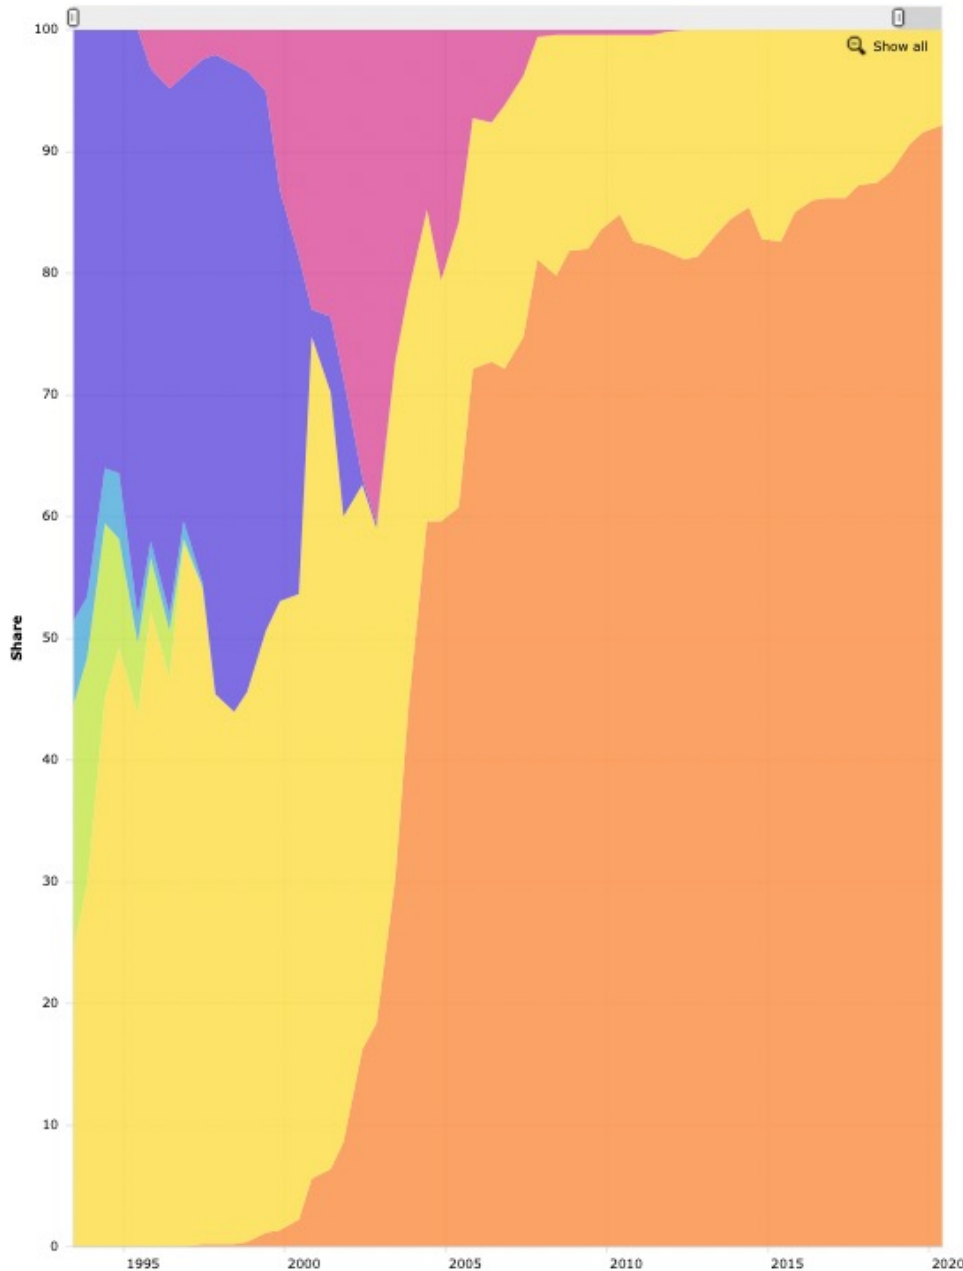
Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	<b>Supercomputer Fugaku</b> - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442,010.0	537,212.0	29,899
2	<b>Summit</b> - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148,600.0	200,794.9	10,096
3	<b>Sierra</b> - IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94,640.0	125,712.0	7,438
4	<b>Sunway TaihuLight</b> - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway, NRCPC National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371



# Top500 as of Nov 2021 (top500.org)

5	<b>Perlmutter</b> - HPE Cray EX235n, AMD EPYC 7763 64C 2.45GHz, NVIDIA A100 SXM4 40 GB, Slingshot-10, HPE DOE/SC/LBNL/NERSC United States	761,856	70,870.0	93,750.0	2,589
6	<b>Selene</b> - NVIDIA DGX A100, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100, Mellanox HDR Infiniband, Nvidia NVIDIA Corporation United States	555,520	63,460.0	79,215.0	2,646
7	<b>Tianhe-2A</b> - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000, NUDT National Super Computer Center in Guangzhou China	4,981,760	61,444.5	100,678.7	18,482
8	<b>JUWELS Booster Module</b> - Bull Sequana XH2000 , AMD EPYC 7402 24C 2.8GHz, NVIDIA A100, Mellanox HDR InfiniBand/ParTec ParaStation ClusterSuite, Atos Forschungszentrum Juelich (FZJ) Germany	449,280	44,120.0	70,980.0	1,764

Architecture - Systems Share



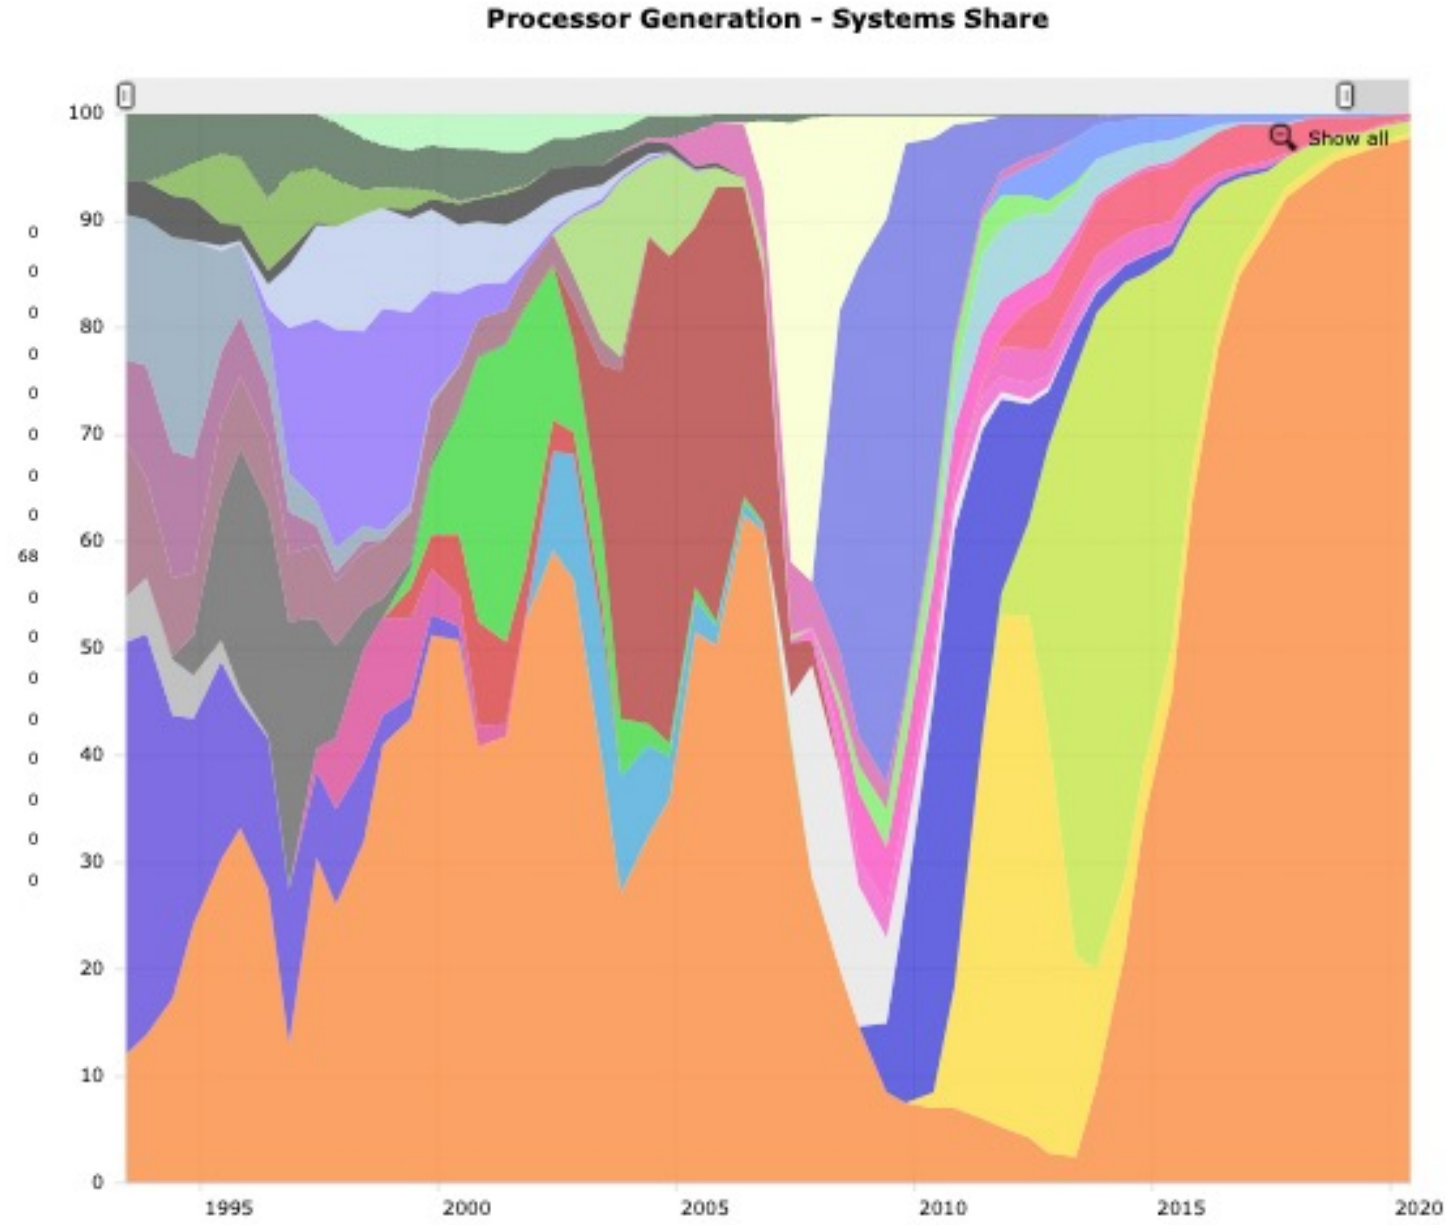
# Top 500

- MPP: Massively Parallel Processor







Generated from top500.org

# Top 500



Generated from top500.org

# Shock Wave Processing of Advanced Reactive Materials



# Multiphysics Solver

$$\psi = \psi_e(\mathbf{F}_e, T) + \psi_p(\chi, T) + \psi_T(T)$$

$$-\rho_0 T \left( \frac{\partial \dot{\psi}_T}{\partial T} \right) + \text{Div} \mathbf{Q} = Q_p + Q_e + Q_c + \rho_0 r$$

$$Q_p = \mathbf{F}_e^T \mathbf{P} \left[ \mathbf{F}_T^T - T \left( \frac{\partial \mathbf{F}_T}{\partial T} \right)^T \right] : \dot{\mathbf{F}}_p - \frac{\partial \psi_p}{\partial \chi} \cdot \dot{\chi} -$$

$$- T \left[ \mathbf{P} \left( \mathbf{F}_p^{-T} : \dot{\mathbf{F}}_p \right) - J_p J_T \frac{\partial W_e}{\partial \mathbf{F}_e} \mathbf{F}_p^{-T} \dot{\mathbf{F}}_p^T \mathbf{F}_p^{-T} \mathbf{F}_T^{-T} \right] : \mathbf{F}_e \mathbf{F}_p \frac{\partial \mathbf{F}_T}{\partial T} + T \rho_0 \left( \frac{\partial \dot{\psi}_p}{\partial T} \right)$$

$$Q_e = -T \left[ \mathbf{P} \left( \frac{\partial \mathbf{F}_T}{\partial T} \right)^T \mathbf{F}_p^T : \dot{\mathbf{F}}_e + \left( J_p J_T \frac{\partial^2 W_e}{\partial \mathbf{F}_e \partial \mathbf{F}_e} : \dot{\mathbf{F}}_e \mathbf{F}_p^{-T} \mathbf{F}_T^{-T} \right) : \mathbf{F}_e \mathbf{F}_p \frac{\partial \mathbf{F}_T}{\partial T} \right] + T \rho_0 \left( \frac{\partial \dot{\psi}_e}{\partial T} \right)$$

$$Q_c = -T \left\{ \left[ \mathbf{P} \left( \mathbf{F}_T^{-T} : \frac{\partial \mathbf{F}_T}{\partial T} \right) - \mathbf{P} \left( \frac{\partial \mathbf{F}_T}{\partial T} \right)^T \mathbf{F}_T^{-T} + J_p J_T \frac{\partial^2 W_e}{\partial \mathbf{F}_e \partial T} \mathbf{F}_p^{-T} \mathbf{F}_T^{-T} \right] : \mathbf{F}_e \mathbf{F}_p \frac{\partial \mathbf{F}_T}{\partial T} + \mathbf{P} : \mathbf{F}_e \mathbf{F}_p \frac{\partial^2 \mathbf{F}_T}{\partial T \partial T} \right\} \dot{T}$$


$$\mathbf{P} \Big|_{\Omega_0} = \rho_0 \frac{\partial \psi_e}{\partial \mathbf{F}_e} \mathbf{F}_p^{-T} \mathbf{F}_T^{-T} = J_p J_T \frac{\partial W_e}{\partial \mathbf{F}_e} \mathbf{F}_p^{-T} \mathbf{F}_T^{-T}$$

$$\eta = - \frac{\partial \psi}{\partial T} \Big|_{\mathbf{F}} = \frac{1}{\rho_0} \mathbf{P} : \mathbf{F}_e \mathbf{F}_p \frac{\partial \mathbf{F}_T}{\partial T} - \frac{\partial \psi_e}{\partial T} - \frac{\partial \psi_p}{\partial T} - \frac{\partial \psi_T}{\partial T}$$

$$\mathbf{Q} = -\mathbf{F}^{-1} \boldsymbol{\kappa} \mathbf{F}^{-1} \nabla_0 T$$

## ► Multiplicative split

$$\mathbf{F} = \mathbf{F}_e \mathbf{F}_p \mathbf{F}_T$$



# Physics: Systems of Partial Differential Equations (PDEs)

Elliptic	Parabolic	Hyperbolic
$\begin{aligned} \nabla \cdot \mathbf{P} &= \mathbf{f}_0 \text{ in } \Omega_0 \\ [[\mathbf{P} \cdot \mathbf{N}_0]] &= [[\mathbf{t}_c]] \text{ on } S_0 \\ \mathbf{P} \cdot \mathbf{N}_0 &= \mathbf{t}_0 \text{ on } \partial\Omega_{t_0} \\ \mathbf{u} &= \mathbf{u}_p \text{ on } \partial\Omega_{u_0} \end{aligned}$	$\begin{aligned} \rho_0 c \dot{T} - \nabla_X \mathbf{Q} &= \mathbf{Q}_f \text{ on } \Omega_0 \\ [[\mathbf{Q} \cdot \mathbf{N}_0]] &= 0 \text{ on } S_0 \\ \mathbf{Q} &= \mathbf{Q}_p \text{ on } \partial\Omega_{Q_0} \\ T &= T_p \text{ on } \partial\Omega_{T_0} \end{aligned}$	$\begin{aligned} \rho_0 \frac{\partial^2 \mathbf{u}}{\partial t^2} - \nabla \cdot \mathbf{P} &= \mathbf{f}_0 \text{ in } \Omega_0 \\ [[\mathbf{P} \cdot \mathbf{N}_0]] &= [[\mathbf{t}_c]] \text{ on } S_0 \\ \mathbf{P} \cdot \mathbf{N}_0 &= \mathbf{t}_0 \text{ on } \partial\Omega_{t_0} \\ \mathbf{u} &= \mathbf{u}_p \text{ on } \partial\Omega_{u_0} \\ \mathbf{u}(0) &= \mathbf{u}_0 \text{ in } \Omega_0 \\ \dot{\mathbf{u}}(0) &= \mathbf{v}_0 \text{ in } \Omega_0 \end{aligned}$
<ul style="list-style-type: none"> <li>constitutive law hyperelastic multiscale</li> </ul>	<ul style="list-style-type: none"> <li>constitutive law Fourier's law</li> </ul>	<ul style="list-style-type: none"> <li>constitutive law hyperelastic multiscale</li> </ul>
<ul style="list-style-type: none"> <li>state variables damage visco-elastic</li> </ul>	<ul style="list-style-type: none"> <li>state variables porosity chemical reactions</li> </ul>	<ul style="list-style-type: none"> <li>state variables damage visco-elastic</li> </ul>
$\rho_0 = J\rho$	<ul style="list-style-type: none"> <li>mass conservation based on physics</li> </ul>	$\rho_0 = J\rho$
<ul style="list-style-type: none"> <li>solution strategy sparse iterative solver dual domain dec.</li> </ul>	<ul style="list-style-type: none"> <li>solution strategy sparse iter. solver <math>\alpha</math>-method integrator</li> </ul>	<ul style="list-style-type: none"> <li>solution strategy sparse iterative solver dual domain dec. MD-AVI</li> </ul>

Courtesy Karel Matous,  
U. Notre Dame

# Computational Science

System of Partial  
Differential Eqns

$$\begin{aligned}\nabla \cdot P &= f_0 \quad \text{in } \Omega_0 \\ \llbracket P \cdot N_0 \rrbracket &= \llbracket t_c \rrbracket \quad \text{on } S_0 \\ P \cdot N_0 &= t_0 \quad \text{on } \partial\Omega_{t_0} \\ u &= u_p \quad \text{on } \partial\Omega_{u_0}\end{aligned}$$

Find P that  
satisfies this

(too hard)

Find x that  
satisfies this

(too hard)

$$F(x) = 0$$

discretize

System of  
Nonlinear Eqns

System of Linear  
Eqns

$$Ax = b$$

Find x that  
satisfies this

linearize

All scientific  
computing is this

A problem we  
can solve

# Computational Science

- The fundamental computation at the core of many (incl. ML) computational science programs is solving

$$Ax = b$$

Requirements for machine learning are changing this

- Assume  $x, b \in \mathcal{R}^N$  and  $A \in \mathcal{R}^{N \times N}$

We will see this a lot

- i.e.,  $x$  and  $b$  are vectors with  $N$  real elements and  $A$  is a matrix with  $N$  by  $N$  real elements

This is what computers can do

- Solution process only requires basic arithmetic operations



# Problem Solving

- Software development is difficult
- How do humans attack complex problems?
- Apply the same principles to software
  
- Modular / reusable
- Well defined interfaces and functionality
- Understandable



# First basic truth of code





- Code is a communication medium with other developers
- And with a future version of yourself



Don't do this

- You can easily write code that no one (including you) can understand

# Two simple rules for writing software



# C++ development philosophy

P.1: Express ideas directly in code

P.2: Write in ISO Standard C++

P.3: Express intent

P.4: Ideally, a program should be statically type safe

P.5: Prefer compile-time checking to run-time checking

P.6: What cannot be checked at compile time should be checkable at run time

P.7: Catch run-time errors early

P.8: Don't leak any resources

P.9: Don't waste time or space

P.10: Prefer immutable data to mutable data

P.11: Encapsulate messy constructs, rather than spreading through the code

P.12: Use supporting tools as appropriate

P.13: Use support libraries as appropriate

From C++ Core  
Guidelines

Only one rule  
about C++

Many follow  
from the two  
simple rules


# Developing your code



- That includes (especially) mental labor
- Use productivity tools
- **VS code** (rec'd), Atom, Eclipse

```
www.ts - node-express-ts
1  import app from './app';
2  import debugModule = require('debug');
3  import http = require('http');
4
5  const debug = debugModule('node-express-typescript:server');
6
7  // Get port from environment and store in Express.
8  const port = normalizePort(process.env.PORT || '3000');
9  app.set('port', port);
10
11 // create
12 const ser
13 server.li
14 server.on
15 server.on
16
17 /**
18  * Normal
19  */
20 function normalizePort(val: any): number|string|boolean {
21   let port = parseInt(val, 10);
22
23   if (isNaN(port)) {
24     // named pipe
25     return val;
```

# What about ...?



- Muscle memory for typing is not the same as productivity (know the difference)
  - Stretch yourself
- Use any environment where you are most productive
- We can only support one (VS code + clang + Linux)
- Assignments must work with autograder

# HPC Legacy

- Command-line and text based (tty)
- Fortran (or “C-tran”)

# Programming


```
int main() {  
    int a = 1;  
    double x = 0.3;  
    foo(x, a);  
}
```

?



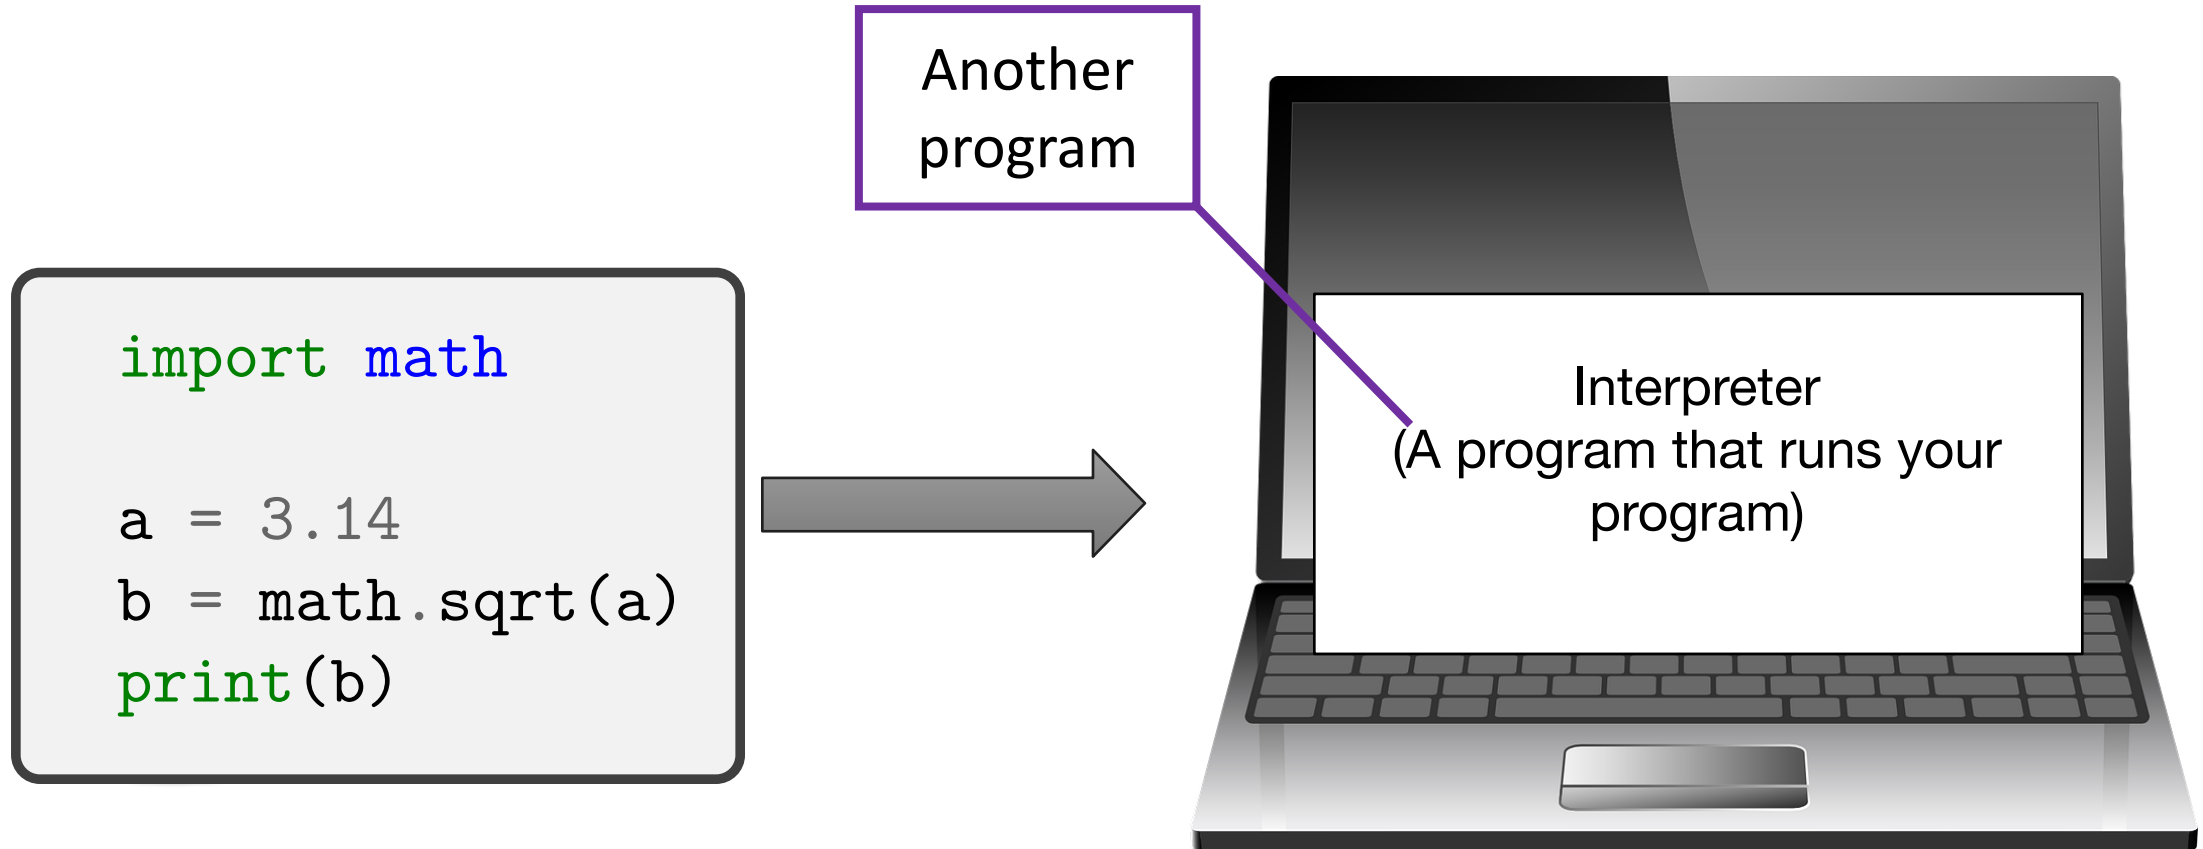


# Programming

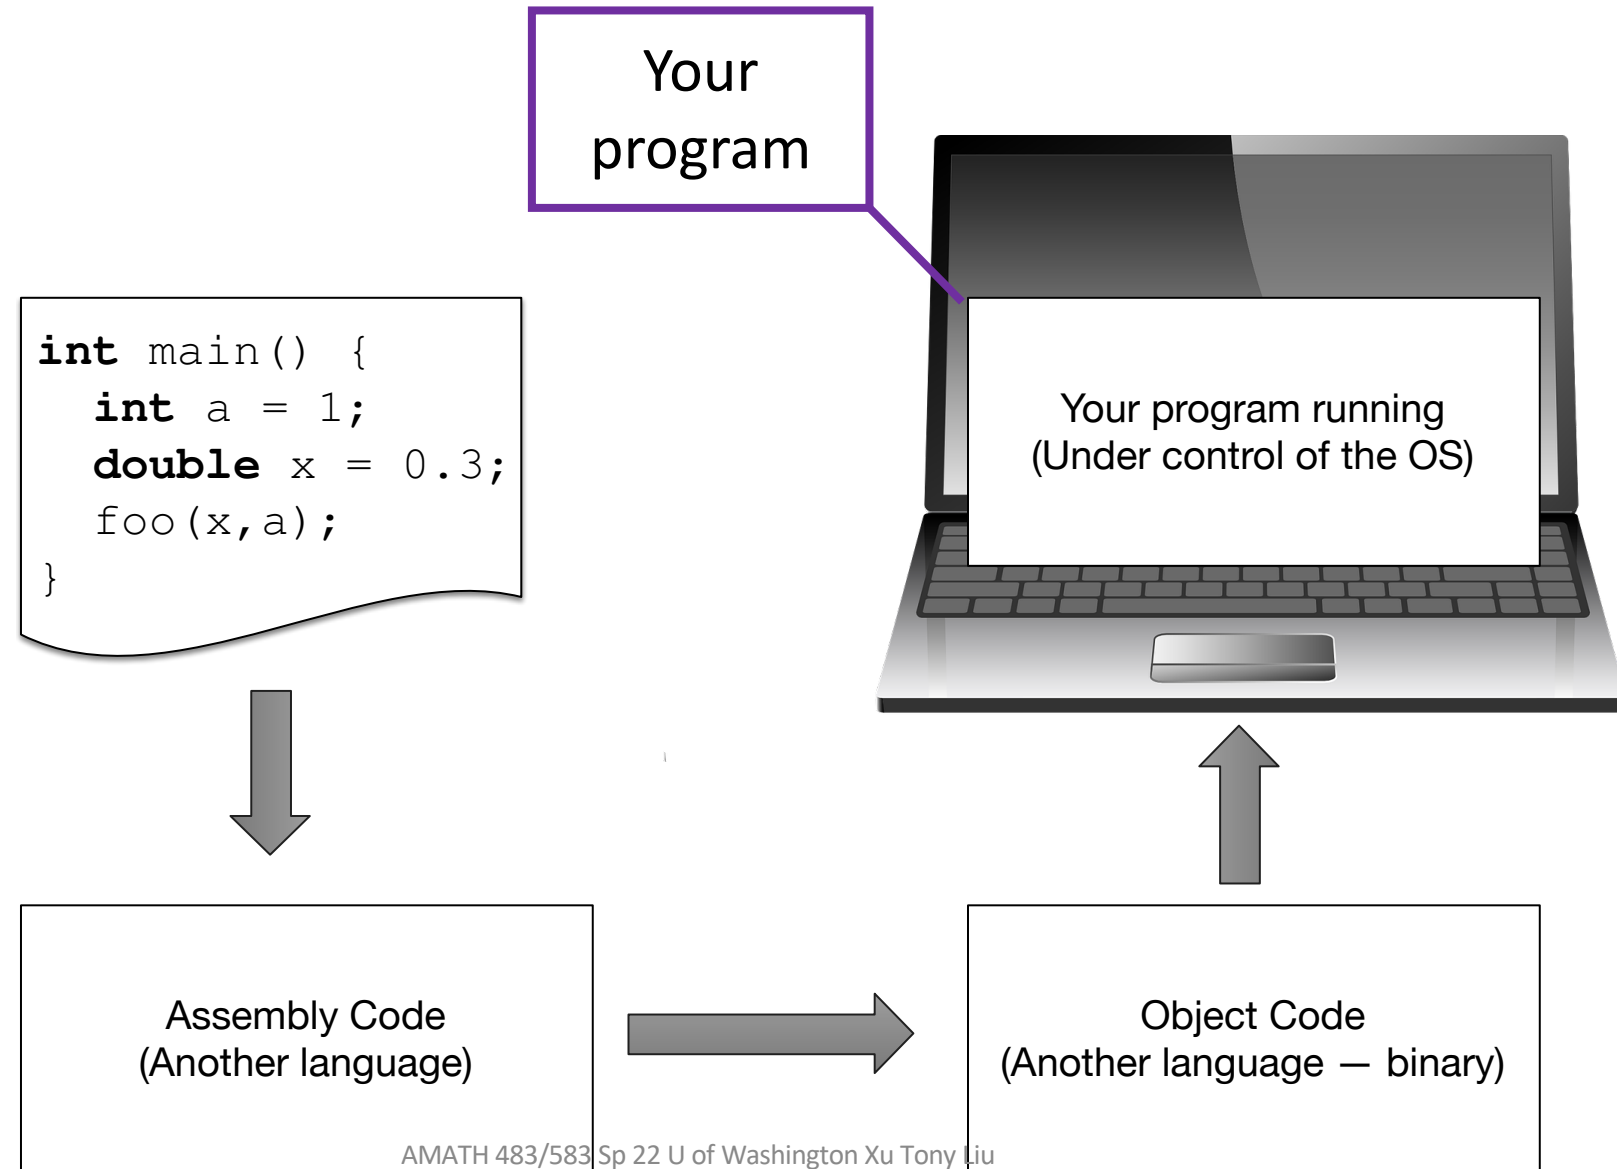


© S. Harris

# Interpreted language (Python)



# Compiled language



# Interpreted vs compiled

Use math library

Call function from math library

Use functions from iostream library

Use math library

```
import math  
  
a = 3.14  
b = math.sqrt(a)  
print(b)
```

Curly braces for code blocks

Code must be in a function

```
#include <cmath>  
#include <iostream>
```

Call function

```
int main() {  
  
    double a = 3.14;  
    double b = std::sqrt(a);  
    std::cout << b << std::endl;  
  
    return 0;  
}
```

Declare variables

Print result

Variables are typed

IO also in std

“std” rather than “math”

# Compilation

```
#include <cmath>
#include <iostream>

int main() {

    double a = 3.14;
    double b = std::sqrt(a);
    std::cout << b << std::endl;

    return 0;
}
```

You can't run  
this code

It needs to be  
turned into code  
that can run

An  
"executable"

Multi-step  
process

Compile to  
**object file**

Then link in  
libraries for  
sqrt and IO

Bits just for  
this code

# Compiling

- To compile one source file to an executable
  - `$ g++ filename.cpp`
  - (What is the name of the executable?)
- To compile multiple source files to an executable
  - `$ g++ one.cpp two.cpp three.cpp`
- To create an object file
  - `$ g++ -c one.cpp -o one.o`
- To create an executable from multiple object files
  - `$ g++ one.o two.o three.o -o myexecutable`

# Slice of C++

- C++11 (C++14, C++17, C++20) are quite modern languages
- But C++11 (et al.) and libraries are *huge*
- We will use a focused slice of C++11
- Use some modern features
- Avoid legacy features (such as pointers)
- Avoid modern features (Object Oriented)

```
#include <cmath>
#include <iostream>


int main() {

    double a = 3.14;
    double b = std::sqrt(a);
    std::cout << b << std::endl;

    return 0;
}
```

# The Base Environment

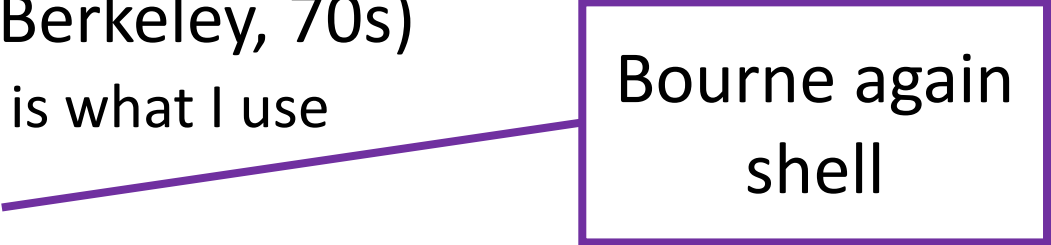
- We will run a pseudo-Linux (a bash shell) in a Docker container
- Provides a uniform environment for everyone to use (compiler etc)
- We can much more effectively support one environment
- Documentation in problem set and online



docker




# shells

- sh: “Bourne shell” (Stephen Bourne, Bell Labs c.1977)
- ksh: Korn shell (David Korn, Bell Labs, c. 1983)
- csh: C shell (Bill Joy, UC Berkeley, 70s)
  - and cousin tcsh – which is what I use
- bash (Brian Fox, 1989) 
  - who knows what this stands for (without searching)
- All are Linux (Unix) processes with read-eval-print loops
- But also complete systems scripting language for dealing with Unix
  - Unix philosophy: data in text format, small programs using text I/O

# Thank You!

# Creative Commons BY-NC-SA 4.0 License



© Andrew Lumsdaine, 2017-2022

Except where otherwise noted, this work is licensed under

<https://creativecommons.org/licenses/by-nc-sa/4.0/>

Cuda and Thrust programming examples © Nvidia

