



中山大學
SUN YAT-SEN UNIVERSITY



国家超级计算广州中心
NATIONAL SUPERCOMPUTER CENTER IN GUANGZHOU

Computer Architecture

计算机体系结构

第0讲：课程介绍、概述

张献伟

xianweiz.github.io

DCS3013, 8/31/2022



中山大學
SUN YAT-SEN UNIVERSITY



About[关于课程]

- 年级专业

- 教学班1：20级保密管理、20信计、20网安、20计科（大数据）、20计科（系统结构-限选）
- 教学班2：20级计算机科学与技术1、2班【吴迪老师】

- 相关课程

- 计算机组成原理、编译原理
- 程序设计、数据结构、操作系统等

- 计算机体系结构

- 如何设计一台符合系统设计目标的计算设备？
- 使用量化分析方法，对计算机瓶颈问题定位、描述、分析、评估
 - 包括指令流水线、并行、存储层次结构、多核结构等

Textbook[课程教材]



- 主要教材

- John L. Hennessy and David A. Patterson, 计算机体系结构：量化研究方法（英文版·原书第6版）

- ▣ 非必须购买

- 参考资料

- ECE 447, Onur Mutlu (Carnegie Mellon U.)
- CIS 501, Milo Martin (U Penn)
- Computer Organization and Design RISC-V Edition: The Hardware Software Interface (2nd Edition), Hennessy and Patterson
- Computer Systems: A Programmer's Perspective (CSAPP), Bryant and O'Hallaron
- Introduction to Computing Systems: From Bits and Gates to C and Beyond, Patt and Patel



Textbook[课程教材] (cont.)



John Hennessy and David Patterson circa 1991, upon publication of their groundbreaking book on computer architecture. (Shane Harvey photo)

Pioneers of Modern Computer Architecture Receive ACM A.M. Turing Award

Hennessy and Patterson's Foundational Contributions to Today's Microprocessors Helped Usher in Mobile and IoT Revolutions

Groundbreaking Textbook

Hennessy and Patterson presented new scientifically-based methodologies in their 1990 textbook *Computer Architecture: a Quantitative Approach*. The book has influenced generations of engineers and, through its dissemination of key ideas to the computer architecture community, is credited with significantly increasing the pace of advances in microprocessor design. In *Computer Architecture*, Hennessy and Patterson encouraged architects to carefully optimize their systems to allow for the differing costs of memory and computation. Their work also enabled a shift from seeking raw performance to designing architectures that take into account issues such as energy usage, heat dissipation, and off-chip communication. The book was groundbreaking in that it was the first text of its kind to provide an analytical and scientific framework, as well as methodologies and evaluation tools for engineers and designers to evaluate the net value of microprocessor design.

Instructor[任课教师]



博士，2011 – 2017， University of Pittsburgh



学士，2007 – 2011， 西北工业大学

中山大學

副教授，2020.10 – 今



工程师/研究员，2017.08 – 2020.09



实习研究员，2016.05 – 2016.08

计算机体系结构



高性能及智能计算

编程及编译优化



本科：编译原理/实验（2021s/2022s）

研究生：高级计算机体系结构（2021f/2022f）

Slides[课件及答疑]

- 课件
 - <https://xianweiz.github.io/teach/dcs3013/f2022.html>
 - 课后上传更新
- **课程QQ群: 419 655 749**
- 超算习堂: <https://easyhpc.net/course/157>
- 教师
 - 张献伟 (超算中心)
 - Email: zhangxw79@mail.sysu.edu.cn
 - 课前课间答疑, 其他时间需预约
- 助教
 - TBD (tbd@mail2.sysu.edu.cn)
 - TBD (tbd@mail2.sysu.edu.cn)



Grading[课时及考核]

- 课时（3学分，54学时）

- 1-18周，周三 5/6节（14:20-16:00）
- 10-18周，周一5/6节（14:20-16:00）
- 地点：教学大楼 C105

- 考核

- 课堂参与（20%）- 点名、提问、测试
- 平时作业（40%）- 课下
 - 习题、实践、paper review。。。。
- 课程项目（40%）
 - simulator

- 课堂

- 随机点名
 - 缺席优先
- 随机提问
 - 后排优先
- 随机测试
 - 不定时间

- 实验/作业

- 个人完成
 - 杜绝抄袭
- 按时提交
 - 超算习堂


Schedule[上课安排]

周次	课程内容	周次	课程内容
第1周 (8.31)	三：课程介绍、概述	第10周 (10.30/11.2)	一：Data Level Parallelism (1) 三：Data Level Parallelism (2)
第2周 (9.7)	三：量化分析方法 (1)	第11周 (11.7/11.9)	一：Data Level Parallelism (3) 三：Thread Level Parallelism (1)
第3周 (9.14)	三：量化分析方法 (2)	第12周 (11.14/11.16)	一：Thread Level Parallelism (2) 三：Thread Level Parallelism (3)
第4周 (9.21)	三：ISA/ILP (1)	第13周 (11.21/11.23)	一：Interconnect (1) 三：Interconnect (2)
第5周 (9.28)	三：ISA/ILP (2)	第14周 (11.28/11.30)	一：Interconnect (3) 三：Warehouse Scale Computer (1)
第6周 (10.5)	三：ISA/ILP (3)	第15周 (12.5/12.7)	一：Warehouse Scale Computer (2) 三：Warehouse Scale Computer (3)
第7周 (10.12)	三：Memory System (1)	第16周 (12.12/12.14)	一：Domain Specific Arch (1) 三：Domain Specific Arch (2)
第8周 (10.19)	三：Memory System (2)	第17周 (12.19/12.21)	一：Domain Specific Arch (3) 三：Advanced Topics (1)
第9周 (10.26)	三：Memory System (3)	第18周 (12.26/12.28)	一：Advanced Topics (2) 三：Advanced Topics (3)

Outline[内容安排]

- Overview and Fundamentals[概览与基础]
- Instruction Set Architecture[指令集架构]
 - Quick intro/review
- Instruction Level Parallelism[指令级并行]
 - Pipelining, Branch Prediction, Instruction Scheduling
- Memory Hierarchy[存储层级]
 - Memory, Cache, Virtual Memory
- Data/Thread-Level Parallelism[数据/线程级并行]
 - SIMD, GPU
- Domain-specific Architectures[领域专用架构]
- And more ...

Examples



CPU



implicitly managed

GPU



mixed

'TPU'



explicitly managed

Goals[课程目标]

- This course covers HW/SW, and the interface[内容覆盖]
 - We will focus on performance analysis and design tradeoffs
- Two key goals of this course are[主要目标]
 - To understand how **hardware** components works with the **software** layer and how decisions made in hardware affect the software/programmer
 - To enable you to be comfortable in making **design and optimization** decisions that cross the boundaries of different layers and system components
- Two other goals of this course[额外目标]
 - Enable you to think **critically**
 - Enable you to think **broadly**

一些问题？ ？ ？

“摩尔定律”真的要终结吗？量子计算是未来吗？

GPU怎么就从图形专用到了计算通用？

NPU, TPU, 加速器？异构计算？

Intel/AMD/Nvidia市值：2012 ($I=12A=25N$) vs. 2022 ($I=0.94A=0.34N$)

Intel挤牙膏？ AMD YES!

Meltdown, Spectre漏洞为什么会发生？

DRAM、DDR、HBM、LPDDR，什么关系？

华为被制裁，卡脖子卡的是什么？ x86, ARM, RISC-V?

太阳位置会影响计算机稳定？

GPU任务怎么被执行的？ profiling怎么实现的？

为什么要有cuBLAS基础库？

数据中心、超算、云，什么关系？

“冯·诺伊曼”架构要被淘汰了？


CA?[什么是体系结构]

- Computer Architecture

- The science and art of designing, selecting, and interconnecting **hardware components** and designing the **hardware/software interface** to create a computing system that meets functional, performance, energy consumption, cost, and other specific goals[硬件设计 – 软硬件接口 – 系统目标]

- Computer architecture is the glue that binds SW and HW

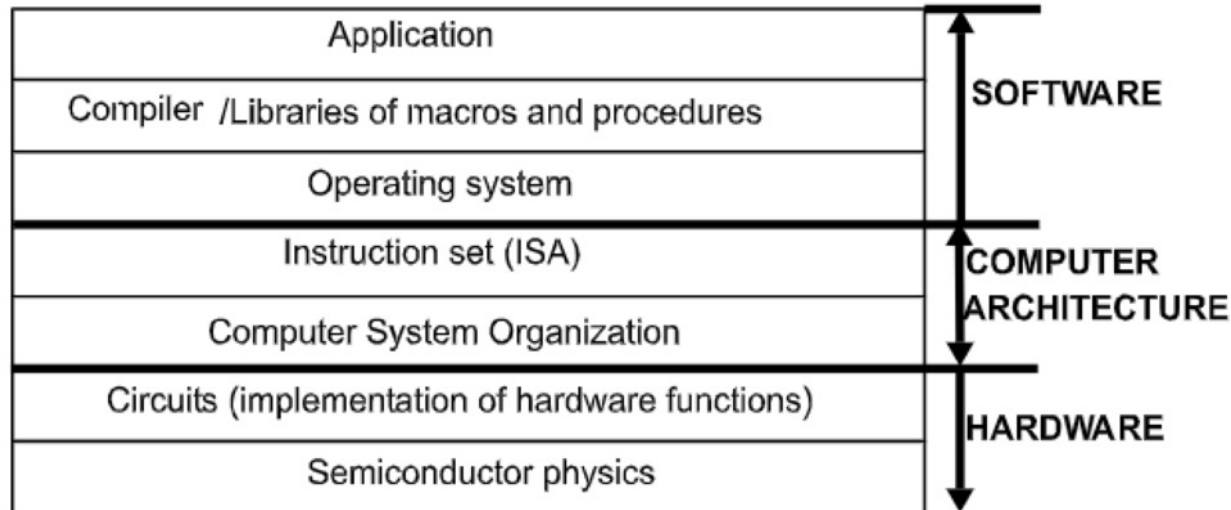
- Inter-disciplinary in nature[连接软硬件、跨学科]



建筑 == 体系结构

CA?[什么是体系结构] (cont.)

- Computer organization[组成]
 - How to build computers?
- Computer architecture[架构]
 - How to design computer system?
- Layered view of computer systems[分层角度]



体系结构 = 指令集体系结构 + 计算机组成

Computer Design

Instruction Set Design

■ Machine Language

- Compiler View
- Computer Architecture

"Building Architect"

“建筑设计师”

功能定义与设计

例如：

- **指令集体系结构 (ISA)** 设计考虑：是否提供“乘法指令”？
- **组成 (Organization)** 考虑：如何实现乘法指令(用专门的乘法器还是用一个加法器+移位器实现)
- **计算机实现技术 (Technology)** 考虑：如何布线、用什么材料、工艺？

Computer Hardware Design

- Machine Implementation
 - Logic Designer's View
 - Computer Organization


"Construction Engineer"

“建造工程师”

考虑用什么材料，如何布线等

Layering[计算机分层与抽象]

- Layering to deal with complex systems[分层]
 - Computers are complex, built in **layers**
 - Software: applications, compiler, operating system
 - Hardware: logic, circuits, electrons
- Layering brings abstraction[抽象]
 - A higher level only needs to know about the interface to the lower level, not how the lower level is implemented
 - 99% of users don't know hardware layers implementation
 - 90% of users don't know implementation of any layer
 - That's okay, world still works just fine
- Then, **why** would you want to know what goes on underneath or above?[为什么你要知道?]
 - Because you are professional

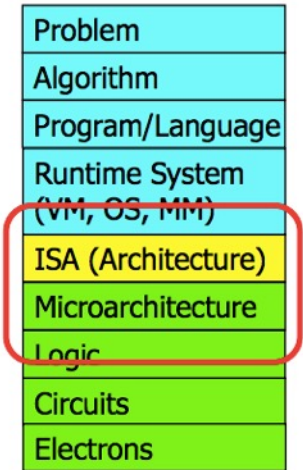


Layering[计算机分层与抽象](cont.)

- As long as everything **goes well**, not knowing what happens in the underlying level (or above) is **not a problem**
- What if
 - The program you wrote is running slow?[慢]
 - The program you wrote does not run correctly?[错误]
 - The program you wrote consumes too much energy?[高能耗]
- What if
 - The hardware you designed is too hard to program?[难用]
 - The hardware you designed is too slow because it does not provide the right primitives to the software?[性能差]
- What if
 - You want to design a much more efficient and higher performance system?[效率更高、性能更高]

Layering[计算机分层与抽象](cont.)

- How would you solve the problem?[怎样解决]
- What is the right place to solve the problem?
[哪里解决]
 - Programmer?
 - System software? Compiler?
 - Hardware? Circuits?



- Breaking the abstraction layers and knowing what is underneath enables you to solve problems and design better future systems[打破抽象]
- Cooperation between multiple components and layers enable more effective solutions and systems[协同考虑]

Architect[架构和架构师]

- Computer architect[计算机架构师]
 - To make **design trade-offs** across the hw/sw interface to meet functional, performance and cost requirements
- Being an architect is not easy[并不容易]
 - You need to consider many things in designing a **new** system + have good **intuition/insight** into ideas/tradeoffs
- But, it is fun and can be very technically rewarding[会很有意思]
- And, enables a great future[影响未来]
 - Advancement of computer architecture is vital to all other areas of computing
 - E.g., IoT, Embedded, Mobile, Data centers, HPC

Role of [Computer] Architect[职责]

- Look **backward** (to the past)
 - Understand tradeoffs and designs, upsides/downsides, past workloads. Analyze and evaluate the past.
- Look **forward** (to the future)
 - Be the dreamer and create new designs. Listen to dreamers.
 - Push the state of the art. Evaluate new design choices.
- Look **up** (towards problems in the computing stack)
 - Understand important problems and their nature.
 - Develop architectures and ideas to solve important problems.
- Look **down** (towards device/circuit technology)
 - Understand the capabilities of the underlying technology.
 - Predict and adapt to the future of technology (you are designing for N years ahead). Enable the future technology.

Why Study CA?[为什么学习体系结构]

- Understand why computers work the way they do
 - Source code --> instructions --> executions
 - Processors --> cache/memory --> storage
- Understand where computers are going
 - Future capabilities drive the (computing) world
 - Real-world impact: no computer architecture -> no computers!
- Understand high-level design concepts and performance
 - The best architects understand all the levels (hw, OS, apps, alg)
 - Need to understand hardware to write fast software
- Job?
 - Your MS/PhD research may involve CA
 - Your job positions (design or research, hw or sw) may need CA
 - You may manage a team working on systems

Positions???[工作职位]

CAREERS AT NVIDIA

Deep Learning Architect – New College Grad

US, CA, Santa Clara

Apply

NVIDIA is seeking computer architects to help design processor and system architectures that will enable compelling Deep Learning performance, architecture and efficiency improvements. This role offers the opportunity to directly impact the future hardware roadmap in a fast-growing technology company that leads the AI revolution. If you are obsessed with improving deep learning performance beyond anything possible with today's hardware and software, this is the place to be.

What you'll be doing:

- Understand, analyze, profile, and optimize deep learning training workloads on state-of-the-art hardware and software platforms.
- Guide development of future generations of deep learning processors and computing platforms.
- Develop detailed performance models and simulators for computing systems accelerating DL training.
- Collaborate across the company to guide the direction of machine learning at NVIDIA; spanning teams from hardware to software and research to production.
- Drive HW/SW co-design of NVIDIA's full deep learning platform stack, from silicon to DL frameworks.

What we'd like to see:

- You are pursuing a PhD or MS or have equivalent in CS, EE or CSEE (or equivalent experience).
- Strong background in computer architecture, preferably with focus on high-performance parallel processors.
- Background in machine learning and neural networks, in particular training.
- Experience analyzing and tuning application performance.
- Experience with processor and system-level performance modelling.
- Programming skills in C++ and Python.
- Familiarity with GPU computing (CUDA, OpenCL).

达摩院-芯片性能架构师-北京

发布时间:	2021-04-20	工作地点:	北京	工作年限:	五年以上
所属部门:	阿里集团	学历:	硕士	招聘人数:	若干

团队介绍:

阿里巴巴集团达摩院旗下的数据计算-计算技术实验室致力于前瞻性研究,探索异构计算,存储,和互联的系统架构,软硬件协同设计,体系结构,电路设计,编译和编程环境等方面的技术问题,研制高性能、低功耗的异构计算系统,人工智能计算芯片,以及其他芯片架构及系统。通过自上而下基于应用驱动和自下而上基于新技术的研究方法,利用计算机体系结构设计优化, VLSI等领域的技术积累与合作伙伴在计算资源优化、新计算体系方向等构建创新系统,提升阿里巴巴集团计算能力并复用于国民经济的各行各业中。

岗位描述:

参与新型计算芯片的开发,作用包括:

- 分析业务软件的性能特征,在系统和仿真环境中分析性能如何,以帮助芯片架构的设计决策。
- 开发和维护性能分析工具,实现数据分析。
- 参与SOC模块的性能建模,和设计团队一起优化芯片的性能,并和验证。
- 开发性能测试用例,测试系统的微架构特征和设计参数。
- 在新兴计算中,领域(机器学习,视频计算等),分析业务特点,推动软硬件协同设计。

岗位要求:

- 热爱芯片架构,愿意学习。
 - 软件编程能力(如C++, python)
 - 具有计算机体系结构背景,熟悉SOC组件。
- 职位加分:
- 有相关的计算机体系结构研究经历。
 - 有性能模型建模的经验。

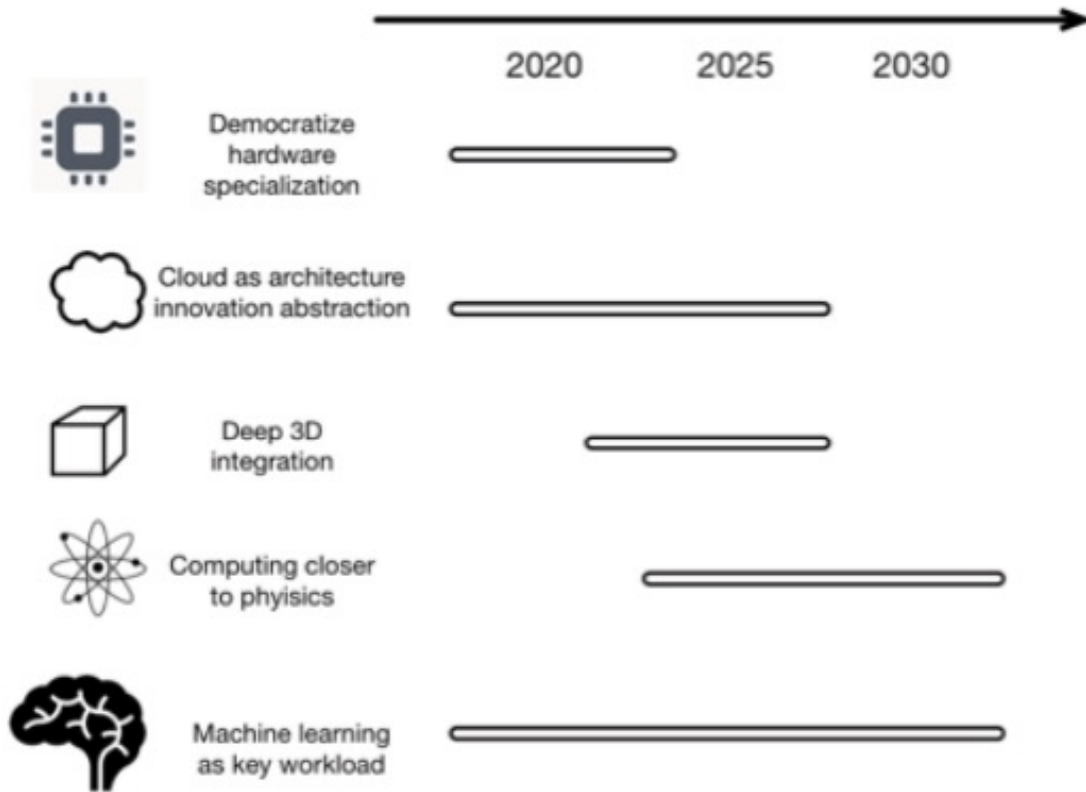
Golden Age for CA[黄金时代]

- **Today** is a very exciting time to study architecture
 - Many new demands from the top
 - Fast changing demands and personalities of users
 - Many new issues at the bottom
- Computing landscape is **very different** from 10-20 years ago (Recall: Intel = 25*Nvidia → 0.34*Nvidia)
 - Every component and its interfaces, as well as entire system designs are being re-examined
 - You can revolutionize the way computers are built, if you understand both the hardware and the software (and change each accordingly)
- **No clear, definitive answers to these problems[有问题, 缺方案]**

Academia[学术界]

Architecture 2030 Workshop @ ISCA 2016

John L. Hennessy, David A. Patterson



- Current challenges[问题]
 - End of Moore's Law and Dennard Scaling
 - Overlooked security
- Future opportunities in computer architecture[机遇]
 - Domain-specific architectures
 - Domain-specific languages
 - Open architectures
 - Agile hardware development

[1] Arch2030, <https://arxiv.org/pdf/1612.03182.pdf> (2016)

[2] [A New Golden Age for Computer Architecture](#) (2019)

Top-tier Conferences[顶级会议]

- **ISCA**

- The International Symposium on Computer Architecture (ISCA)
- 2021: 48th, 76/407 papers (18.6% acceptance rate)

- **MICRO**

- The IEEE/ACM International Symposium on Microarchitecture
- 2021: 54th, 94/430 papers (21.8% acceptance rate)

- **HPCA**

- IEEE International Symposium on High-Performance Computer Architecture
- 2021: 27th, 63/258 papers (24.4% acceptance rate)

- **ASPLOS**

- ACM International Conference on Architectural Support for Programming Languages and Operating Systems
- 2021: 26th, 75/398 (18.8% acceptance rate)

Industry[工业界]

- Intel, 08/2022: Kills Optane Memory^[1]
- Nvidia, 04/2021: Grace, ARM-based data-center CPU^[2]
- Apple, 11/2020: M1, ARM-based SoC^[3]
- AMD, 10/2020: Acquire Xilinx^[4]
- Intel, 09/2020: Xe GPU^[5]
- Samsung, 11/2019: Cease CPU development^[6]
- Amazon, 11/2018: AWS Graviton^[7]
- Intel/IBM/ARM, 01/2018: Meltdown and Spectre^[8]
- Micron, 03/2021: Cease 3D-XPoint, invest CXL^[9]
- AI Chips: Graphcore, Habana Labs, Cerebras, Cambricon...

[1] <https://www.tomshardware.com/news/intel-kills-optane-memory-business-for-good>

[2] <https://www.nvidia.com/en-us/data-center/grace-cpu/>

[3] <https://pdf.wondershare.com/macros/everything-about-apple-m1-chip.html>

[4] <https://www.amd.com/en/corporate/xilinx-acquisition>

[5] <https://www.intel.com/content/www/us/en/products/discrete-gpus/iris-xe-aic.html>

[6] <https://www.anandtech.com/show/15061/samsung-to-cease-custom-cpu-development>

[7] <https://aws.amazon.com/ec2/graviton/>

[8] <https://people.redhat.com/pladd/Meltdown-Spectre-NYRHUG-2018-01.pdf>

[9] <https://investors.micron.com/news-releases/news-release-details/micron-updates-data-center-portfolio-strategy-address-growing>

Exascale Supercomputing[E级超算]

Performance Development



- ECP Director Paul Messina*:
 - Four key challenges: parallelism, memory and storage, reliability and energy consumption.
 - Advances in computer hardware and architecture will contribute to meeting all four challenges.

* <https://www.exascaleproject.org/path-nations-first-exascale-supercomputers-pathforward/>

SC: Aurora

- Aurora @ Argonne National Laboratory*

- 1 exaFLOPS, ≤ 60 MW, 2018-2021-2022

- Compute node:

- 2 Intel Xeon Sapphire Rapids CPUs, 6 Xe GPUs

- First enterprise CPUs to support CXL standard

- GPUs communicates directly to each other via CXL

- Unified memory architecture

- Interconnect


- CPU-GPU: PCIe, GPU-GPU: Xe Link

- System: HPE Slingshot 11; Dragonfly topology with adaptive routing

- Programming models

- Intel oneAPI, MPI, OpenMP, C/C++, Fortran, SYCL/DPC++

- Applications: climate change, cancer, new materials



Aurora: Bringing It All Together

- 2 INTEL XEON SCALABLE PROCESSORS**
"Sapphire Rapids"
- 6 XE ARCHITECTURE BASED GPUs**
"Ponte Vecchio"
- ONEAPI**
Unified programming model
- LEADERSHIP PERFORMANCE**
For HPC, data analytics, AI
- UNIFIED MEMORY ARCHITECTURE**
Across CPU & GPU
- ALL-TO-ALL CONNECTIVITY WITHIN NODE**
Low latency, high bandwidth
- UNPARALLELED I/O SCALABILITY ACROSS NODES**
8 fabric endpoints per node, DAOS

DELIVERED IN 2021



U.S. DEPARTMENT OF ENERGY Argonne intel CRAY

News Under Embargo: November 17, 2019 – 4:00 p.m. Pacific Time

* <https://www.alcf.anl.gov/aurora>

SC: Frontier

- Frontier @ Oak Ridge National Laboratory*
 - 1.5 exaFLOPS, 30 MW, 2021
 - Compute node
 - 1 AMD EPYC CPU (Zen 3) + 4 AMD Radeon Instinct GPU (MI 200)
 - Interconnect
 - Node: AMD Infinity Fabric, coherent memory across the node
 - System:
 - Multiple Slingshot NICs providing 100 GB/s network bandwidth.
 - Slingshot dragonfly network w/adaptive routing
 - Programming models:
 - AMD ROCm, MPI, OpenMP, HIP C/C++, Fortran
 - Applications: modeling and simulation, data analytics, AI



* https://www.olcf.ornl.gov/wp-content/uploads/2019/05/frontier_specsheet.pdf

SC: Tianhe-3

- Tianhe-3 @ Tianjin

- China's native manycore Armv8-based Phytium 2000+ (FTP)
- Matrix 3000 (MTP) Accelerator

