HPC Architecture & Cluster Construction: From Novice to ASC Challenger

Yuyang Jin
Department of Computer Science and Technology
Tsinghua University

About Me

- **Assistant Researcher**
- **Department of Computer Science and Technology, Tsinghua University**
- **Research topic: Parallel computing**
- **Advisor of THUSCC Diablo team**

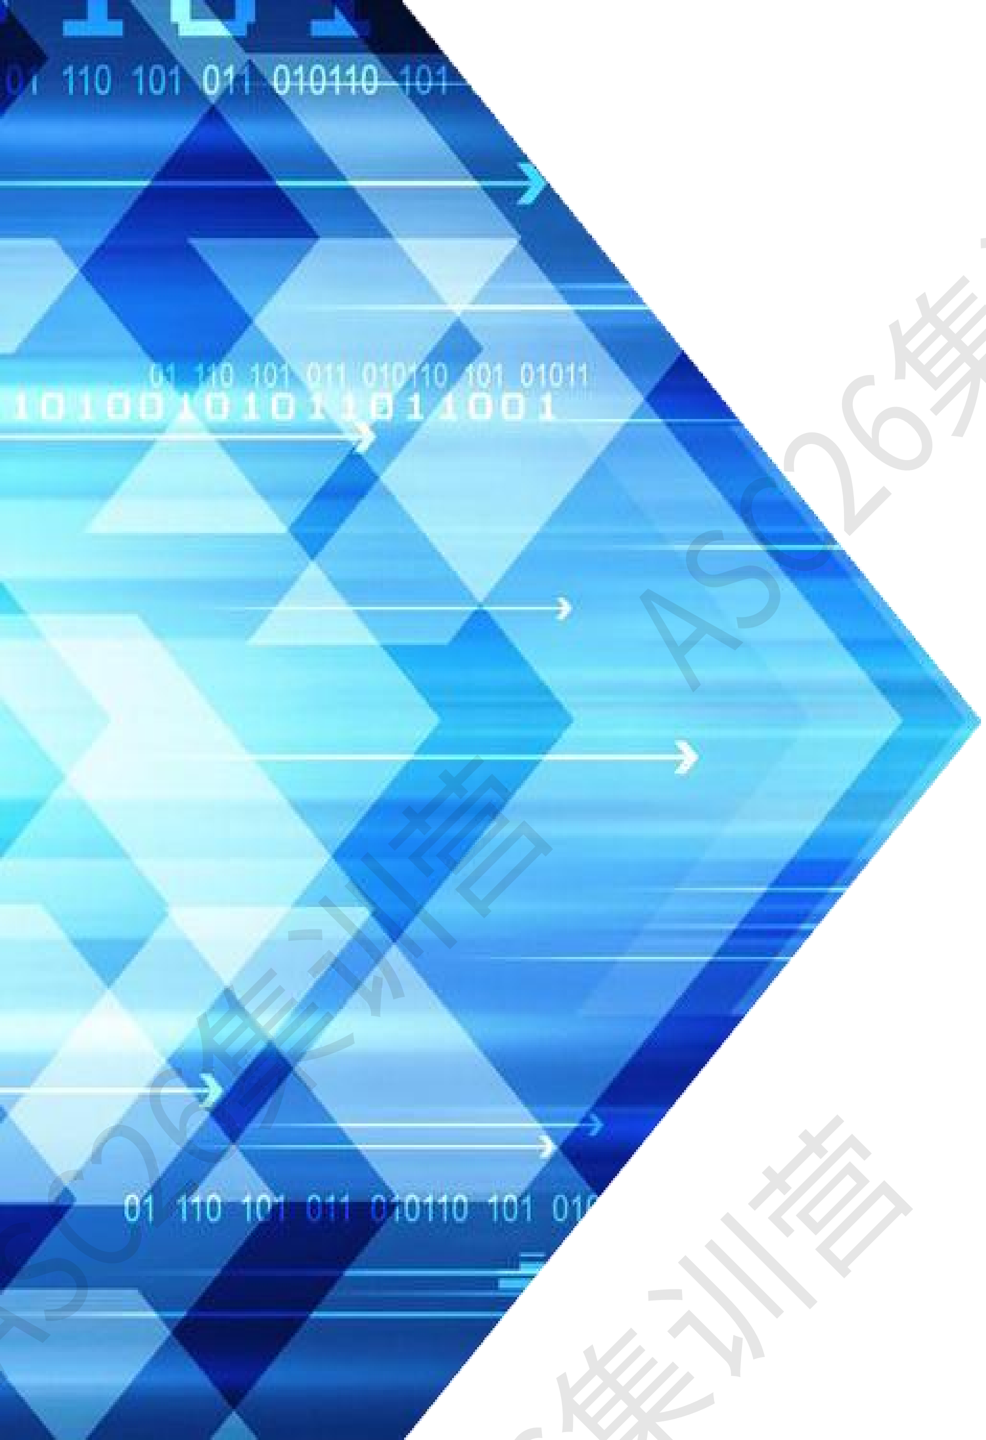

Outline:

- Background
- History
- Deep Dive
  - Processor, Interconnect & Storage
- System Design

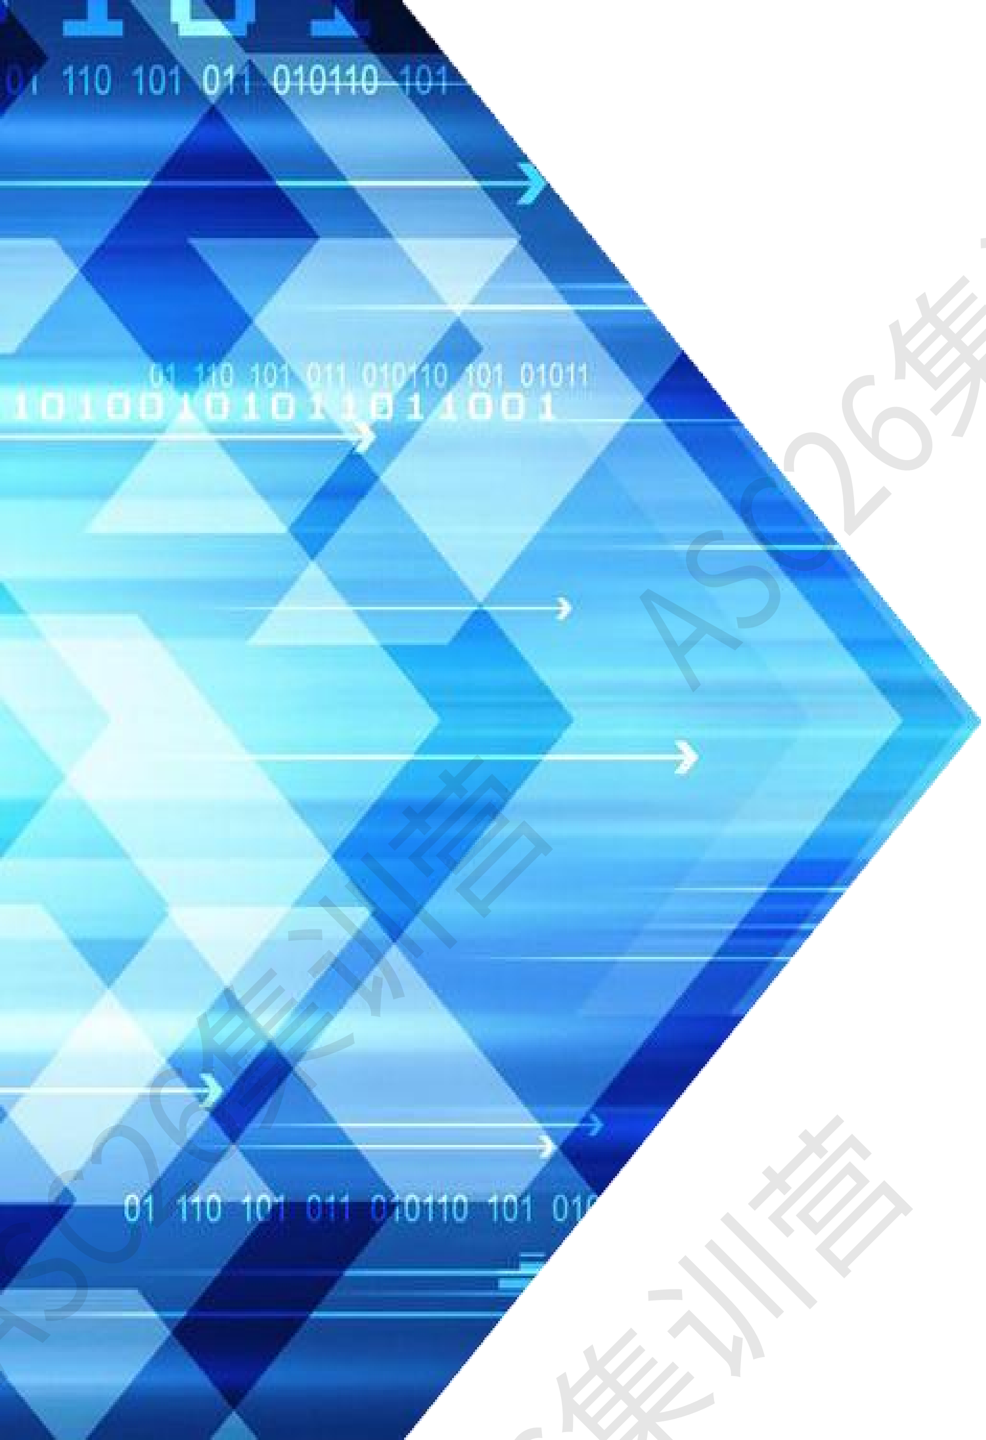

Outline：

- **Background**
- History
- Deep Dive
  - Processor, Interconnect & Storage
- System Design

What is HPC?

- **HPC (High Performance Computing):**
  - Aggregating computing power to solve complex problems
- **Supercomputer**
  - The "Formula 1" of computers
- **Key feature**
  - Parallel computing

Why Do We Need HPC?

(1) Scientific Discovery

- **Simulation**
  - The "Third Pillar" of Science (Theory, Experiment, Simulation)
- **Examples:**
  - Weather Forecasting, Molecular Dynamics, Astrophysics
- **Requirement:**
  - High Precision (FP64), High Reliability

Why Do We Need HPC?

aeronautical

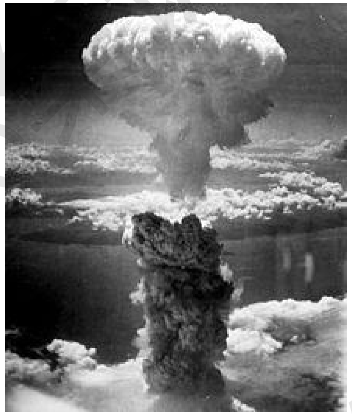

nuclear

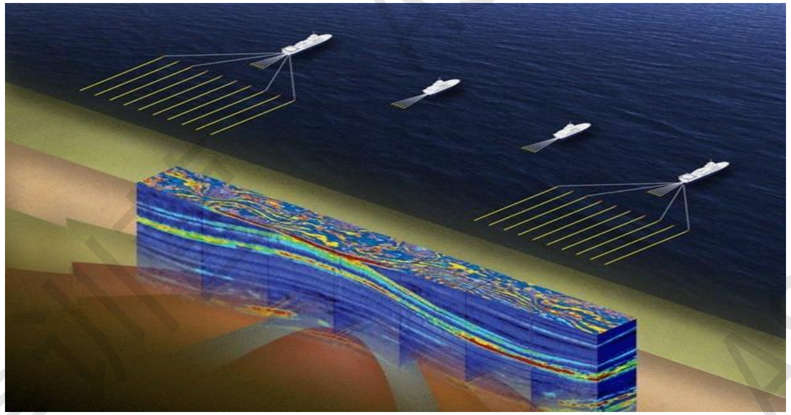

geology

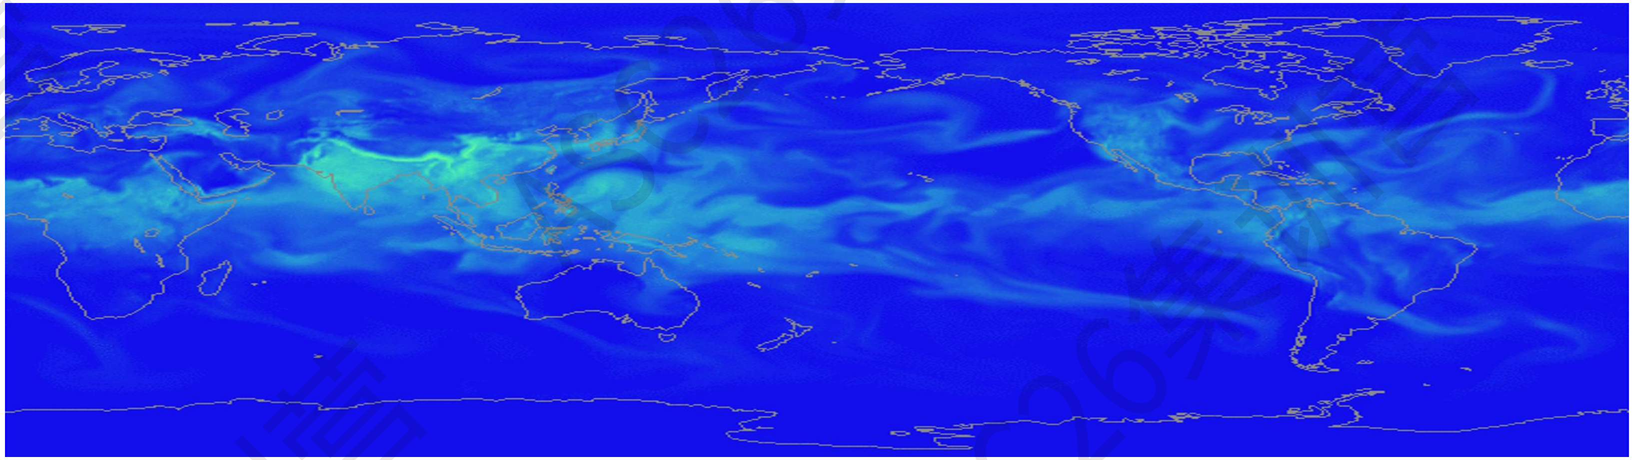

Climate

Why Do We Need HPC?

(2) The AI Revolution

- **Artificial Intelligence:**
  - Deep Learning, LLMs (ChatGPT, DeepSeek)
- **Requirement:**
  - Massive Parallelism, Mixed Precision (FP16/BF16/INT8)
- **Data Explosion:**
  - Training data has grown exponentially

Why Do We Need HPC?

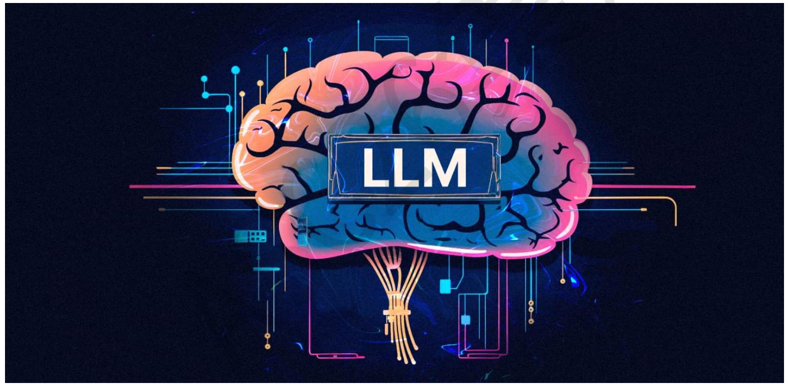

Large language model

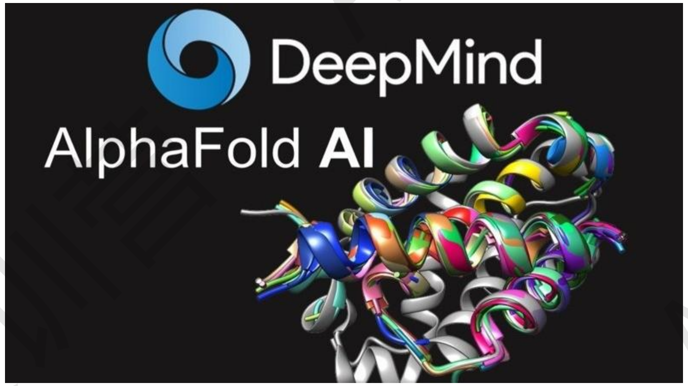

AI for science

Agentic AI

Artificial general intelligence

Why Do We Need HPC?

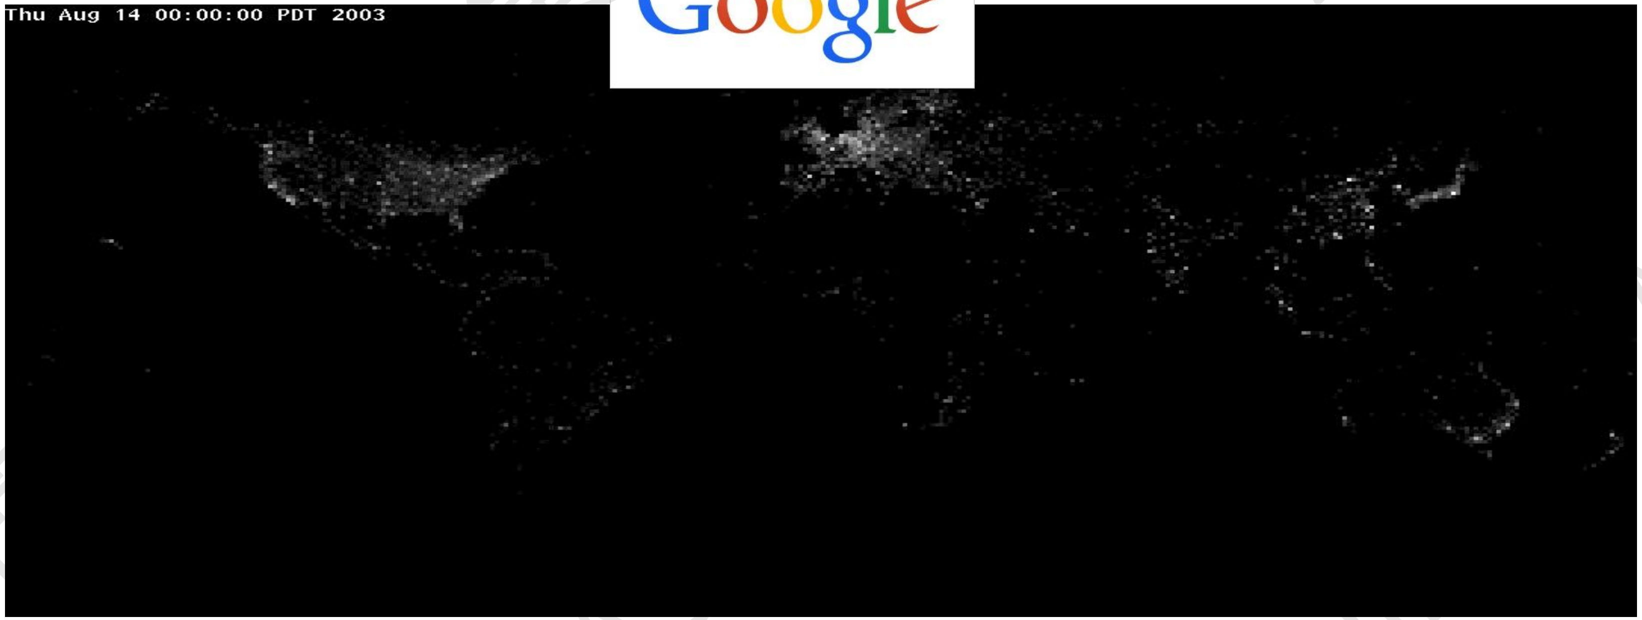

- ☐ Google query attributes (~2003)
  - 150M queries/day (2000/second)
  - 100 countries
  - 8.0B documents in the index
- ☐ Google infrastructure (~2003)
  - >200,000 Linux servers
  - Storage capacity >5 petabytes
  - Traffic growth 20-30%/month

Why Do We Need HPC?

- Science
  - **Global climate modeling**
  - Astrophysical modeling
  - Biology: **genomics**; protein folding; drug design
  - Computational Chemistry
  - Computational Material Sciences and **Nanosciences**
- Business
  - Financial and economic modeling
  - Transaction processing, web services and search engines
- Engineering
  - **Crash simulation**
  - Semiconductor design
  - Earthquake and structural modeling
  - Computation fluid dynamics (airplane design)
  - Combustion (engine design)
  - **Oil field applications**
- Defense
  - **Nuclear weapons** -- test by simulations
  - Cryptography

Serial vs. Parallel Computing

- **Serial:** Tasks executed one after another

- **Parallel:** Multiple tasks executed simultaneously

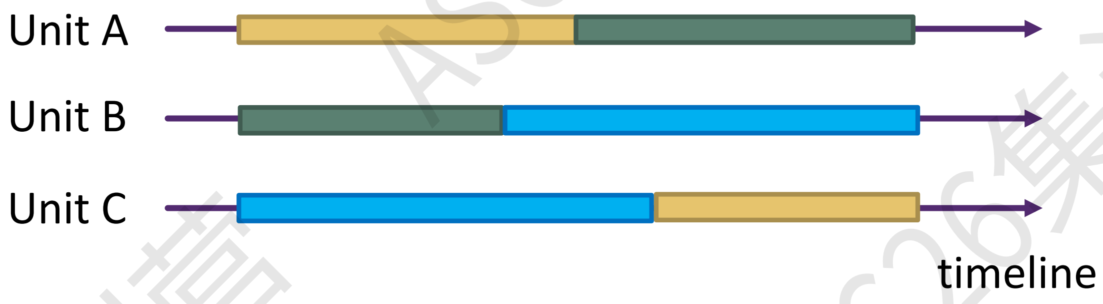

Scale-up vs. Scale-out

- **Scale-up (Vertical):**
  - Bigger, faster single processor (Mainframe style)
  - Limitations: Heat, Physics, Cost
- **Scale-out (Horizontal):**
  - Adding more nodes (Cluster style)
  - The HPC Way

Key Metric: FLOPS

Computing ability (FLOPS)		
1Mflop/s	1Megaflop/s	10^{6}flop/s
1Gflop/s	1Gigaflop/s	10^{9}flop/s
1Tflop/s	1Teraflop/s	10^{12}flop/s
1Pflop/s	1Petaflop/s	10^{15}flop/s
1Eflop/s	1Exaflop/s	10^{18}flop/s
1Zflop/s	1Zettaflop/s	10^{21}flop/s

Storage ability		
1MB	1Megabytes	2^{20}Bytes
1GB	1 Gigabytes	2^{30}Bytes
1TB	1 Terabytes	2^{40}Bytes
1PB	1 Petabytes	2^{50}Bytes
1EB	1 Exabytes	2^{60}Bytes
1ZB	1 Zettabytes	2^{70}Bytes

World-lead (~2.7Eflop/s FP64, 5.6PB)

About FLOPS

One paper with 1000 numbers. Read one paper per second
= 1KFlop/s

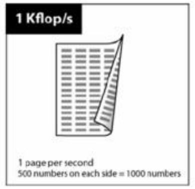

10cm

100m

100km

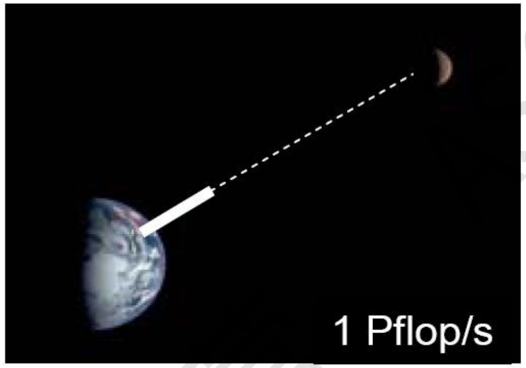

100, 000km
¼ distance between the Earth and the Moon

The Top500 List

- **Ranking:**
  - The bi-annual list of the world's fastest supercomputers (ISC and SC)
- **Benchmark:**
  - HPL (High Performance Linpack)
- **Current Era:**
  - Exascale (10^18 FLOPS) - e.g., Frontier

www.top500.org

The Top500 List

TOP500 2025.11	Machine Name	Peak (PFLOPS)	HPL efficiency	Heterogeneity	Accelerator	Power (MW)
1	El Capitan（USA）	**2,746.38**	63.4%	yes	AMD Instinct MI300A	**29.6**
2	Frontier（USA）	**2,055.72**	65.8%	yes	AMD Instinct MI250X	**24.6**
3	Aurora（USA）	**1,980.01**	51.1%	yes	Intel Data Center GPU Max	**38.7**
4	JUPITER Booster（Germany）	**930.00**	85.3%	yes	NVIDIA GH200	**13.1**
5	Eagle（USA）	**846.84**	66.3%	yes	NVIDIA H100	
6	HPC6（Italy）	**606.97**	78.7%	yes	AMD Instinct MI250X	**8.5**
7	Fugaku（Japan）	**537.21**	82.3%	no	SVE	**29.9**
8	Alps（Swi）	**574.84**	75.7%	yes	NVIDIA GH200	**7.1**

Supercomputer

USA El Capitan

Europe Jupiter Booster

Japan Fugaku

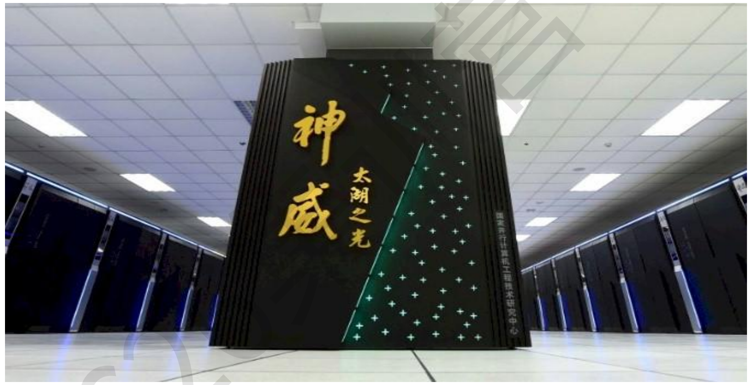

Chinese Sunway TaihuLight

The Top500 Trend

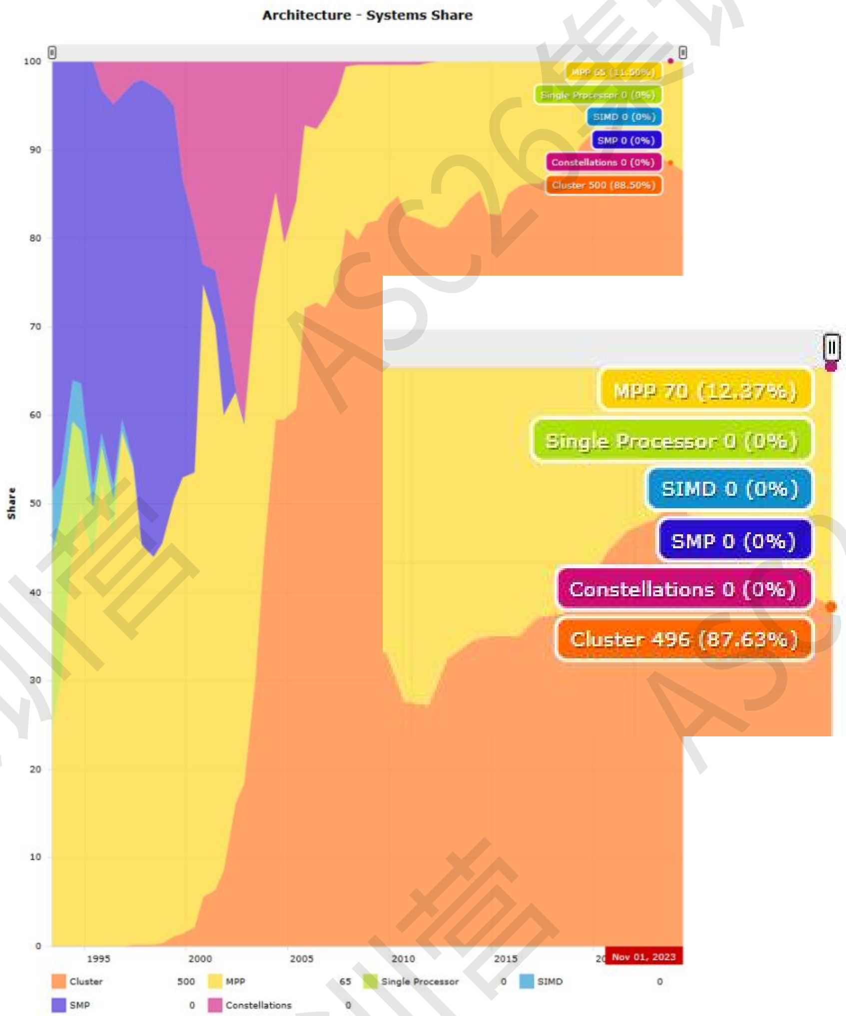

The Top500 Trend

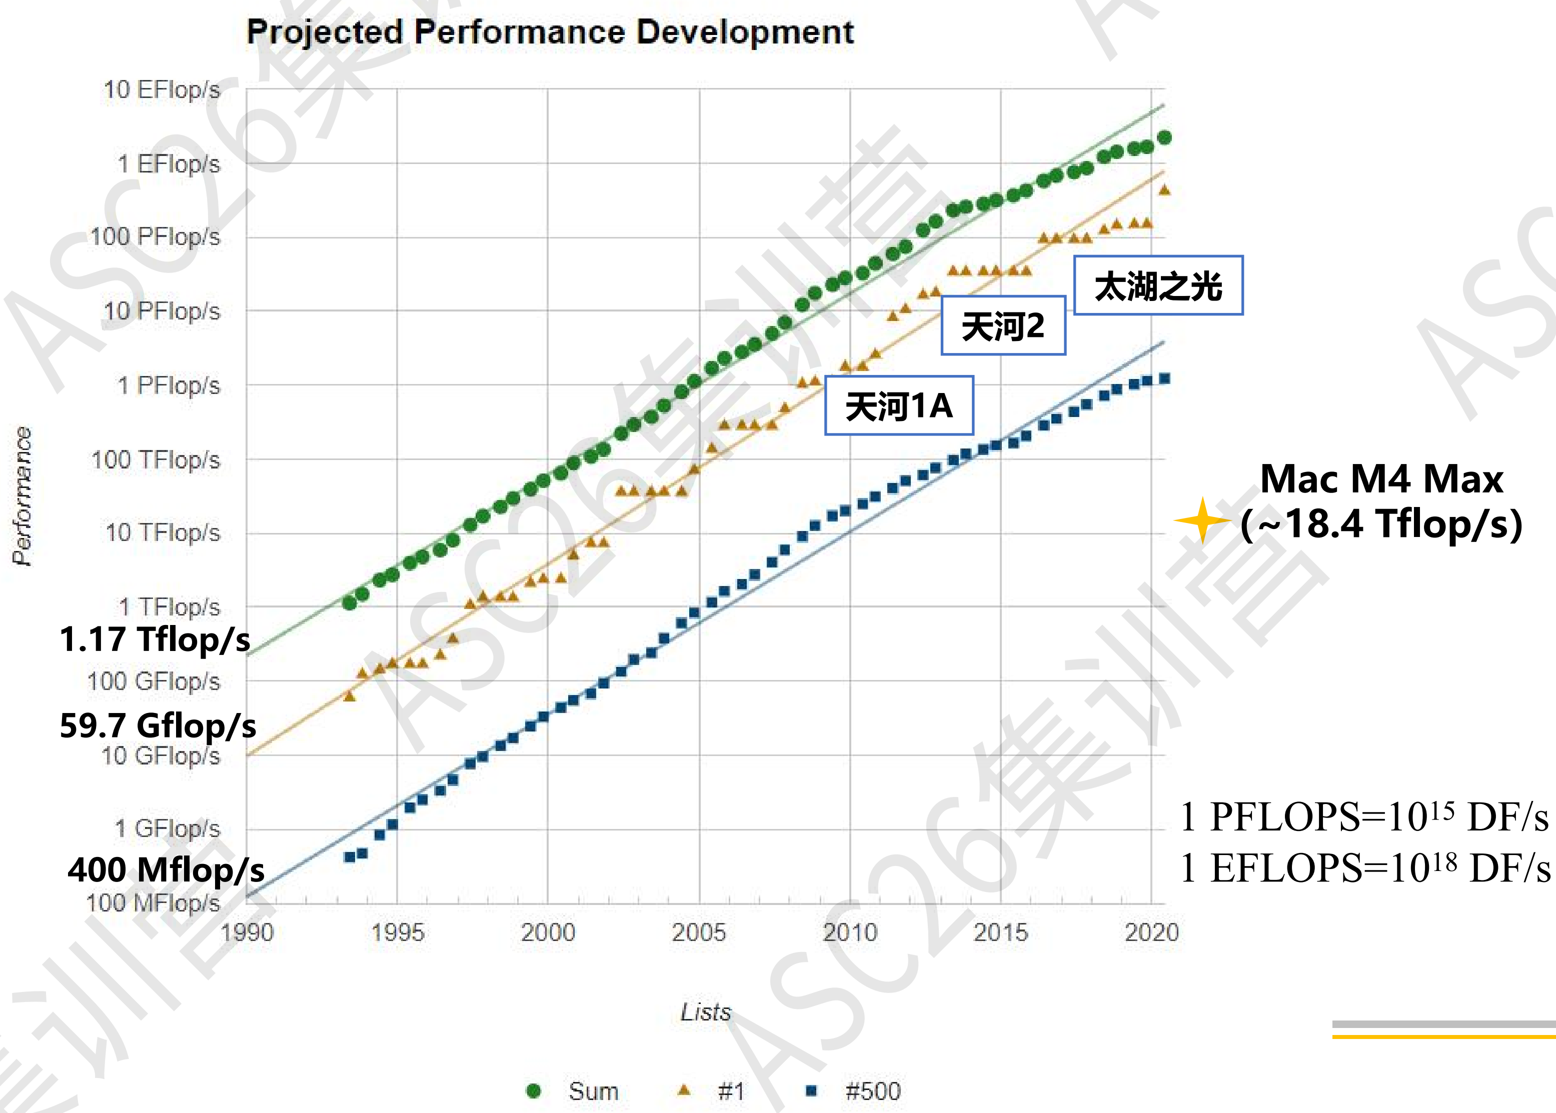

1 PFLOPS=10^{15} DF/s
1 EFLOPS=10^{18} DF/s

Top500: Chinese Supercomputer

www.top500.org

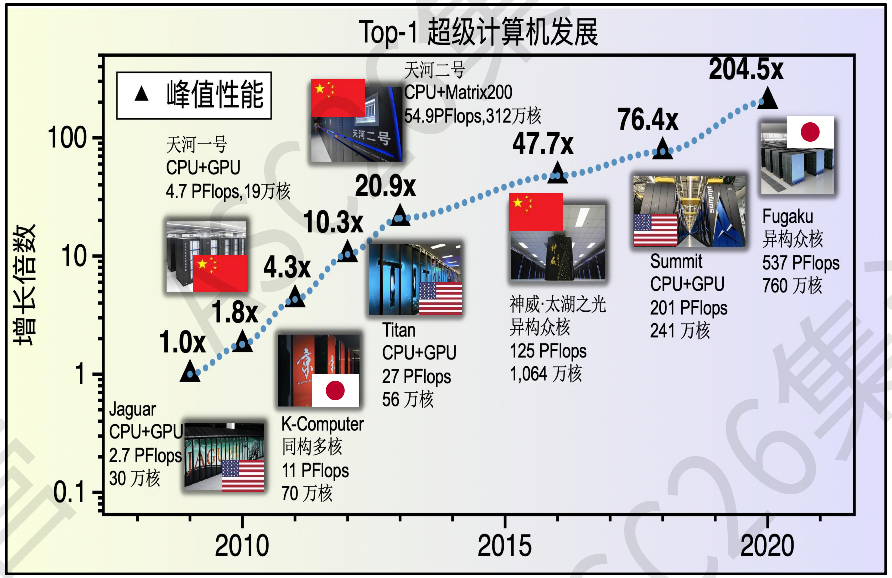

Performance share over Counties/Regions 2024.06

(171)	United States	4,408,382,050	Japan	672,085,510	Finland	391,388,310
(80)	China	356,860,483	Italy	332,618,330	Switzerland	312,137,040
(40)	Germany	274,903,370	France	232,747,500	Spain	221,872,600

China is world-leading in the development of top-tier supercomputing systems

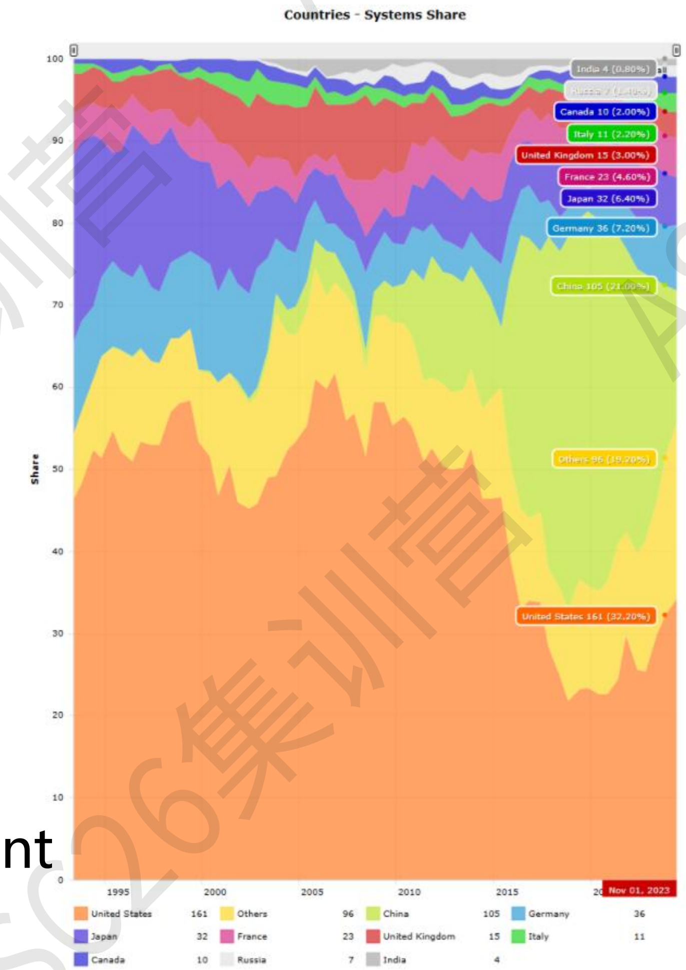

Top500 No.21 Sunway TaihuLight

- The world's **FIRST** 10-million-core computing system, employing heterogeneous many-core processors, with a LINPACK efficiency of 74.16%
- High-density assembly, occupying 605 square meters (FIT building system 100Tflop/s, ~300 square meters).
- Supernode full-switching and fat-tree two-level network, single-port communication bandwidth 56Gbps
- Water cooling system

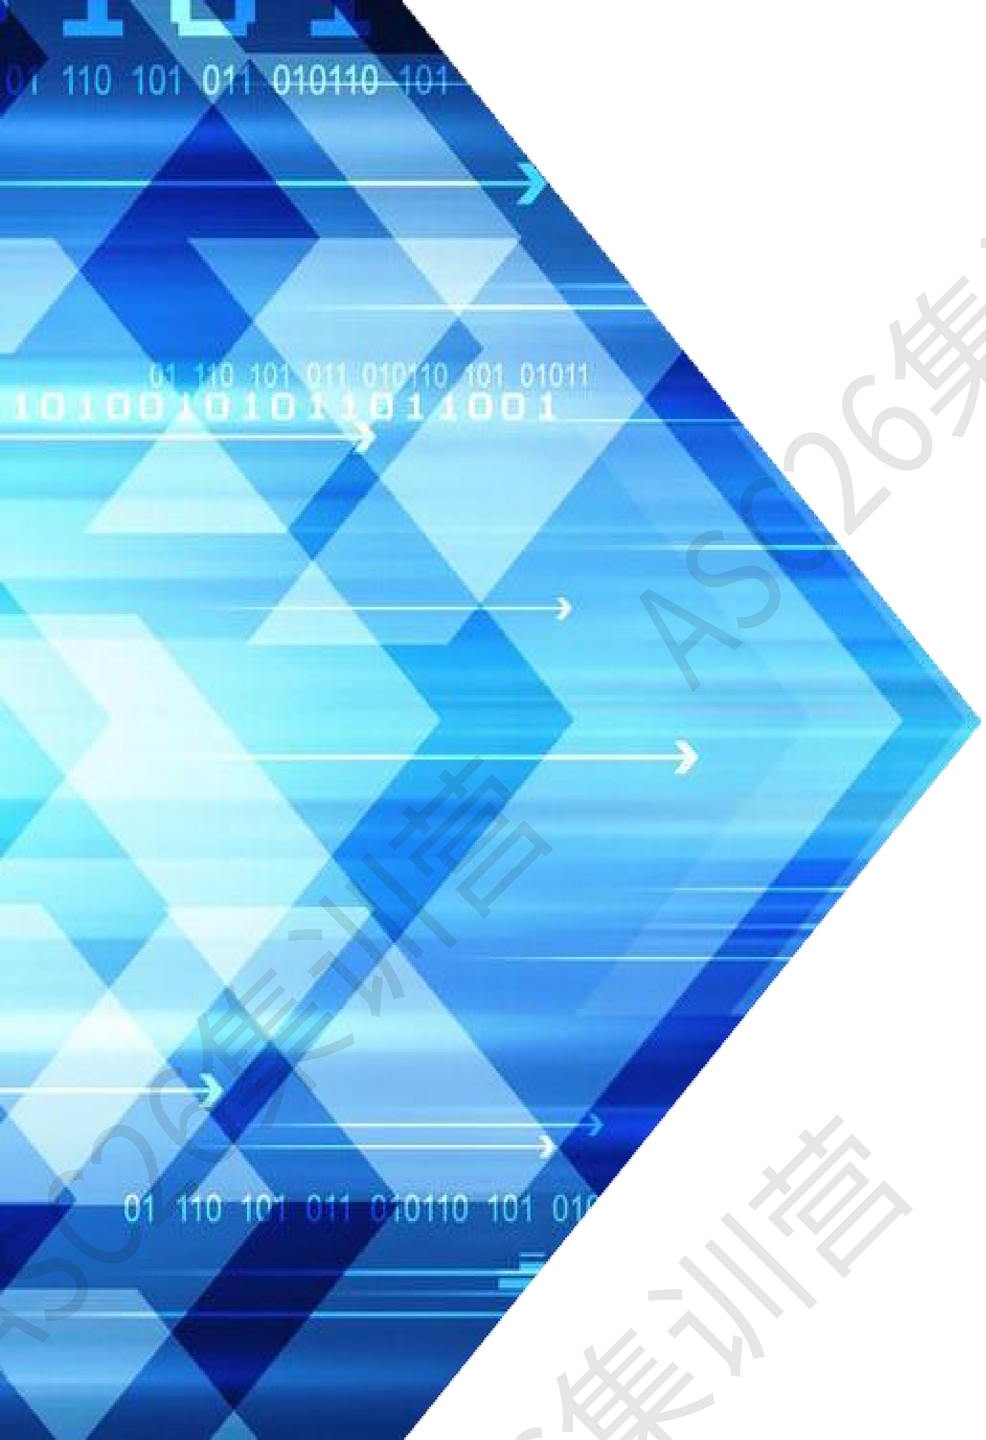

Outline:

- Background
- **History**
- Deep Dive
  - Processor, Interconnect & Storage
- System Design

Evolution of HPC

- **Phase 1: The Vector Era (70s-80s)**
- **Phase 2: MPP & Clusters (90s-00s)**
- **Phase 3: Heterogeneous Era (2010s-Present)**
- **Next Phase: Quantum computing? Neuromorphic computing?**

Evolution Phase 1

The Vector Era (70s-80s)

- **Representative:**
  - Cray-1
- **Architecture:**
  - Custom Vector Processors
- **Feature:**
  - Expensive, Specialized, Cool design (C-shape)

Evolution Phase 1

The Vector Era (70s-80s)

Cray-1, 1976
Integrated circuits used
136 megaFLOPS

Evolution Phase 2

MPP & Clusters (90s-00s)

- **MPP:**
  - Massively Parallel Processing.
- **Beowulf Clusters:**
  - Building supercomputers from COTS (Commercial Off-The-Shelf) hardware.
- **Impact:**
  - Lower cost, standardized Linux + Ethernet/Infiniband

Evolution Phase 2

MPP & Clusters (90s-00s)

1996, ASCI Red
First to 1 teraFLOPS

The final system had a total of 9298 Pentium II OverDrive processors, each clocked at 333 MHz. The system consisted of 104 cabinets, taking up about 2500 square feet (230 m²). The system was designed to use commodity mass-market components and to be very scalable.

Evolution Phase 3

Heterogeneous Era (2010s-Present)

- **The Wall:**
  - CPU frequency scaling stalled (Power/Heat)
- **The Solution:**
  - Accelerators (GPU, MIC)
- **Architecture:**
  - CPU (Host) + GPU (Device)
  - Tianhe-1A led this trend

Evolution Phase 3

Heterogeneous Era (2010s-Present)

- **CPU:**
  - Low core count, complex logic, branch prediction
- **GPU:**
  - Thousands of cores, simple logic, high throughput
- **Analogy:**
  - CPU = Ferrari (Fast individual transport)
  - GPU = Bus (High capacity transport)

Evolution Phase 3

Heterogeneous Era (2010s-Present)

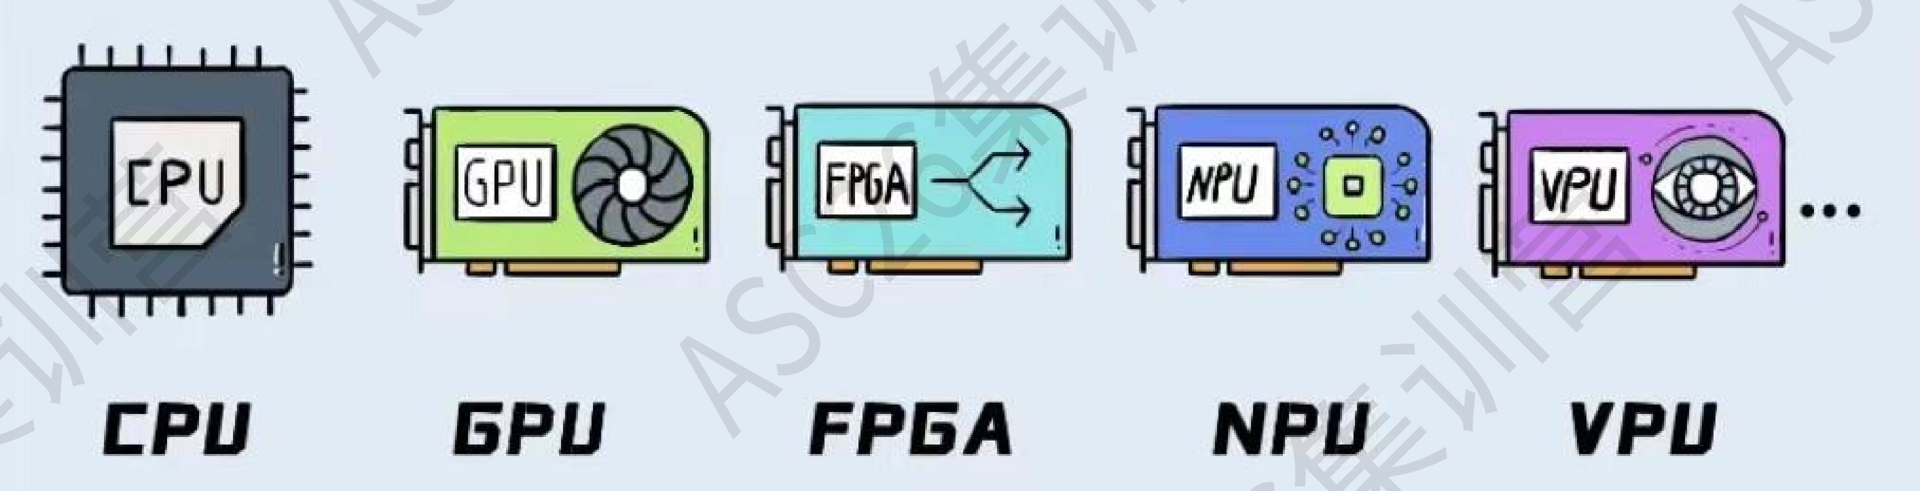

The Rise of AI Architecture

- **New Workloads:**
  - Matrix Multiplications (GEMM)
- **New Units:**
  - Tensor Cores (NVIDIA), Matrix Cores
- **Precision Shift:**
  - From FP64 to Mixed Precision (FP16/BF16)

AI4Science: Convergence

- **Trend:**
  - Using AI to accelerate Scientific Simulation
- **Examples:**
  - AI for Protein Folding (AlphaFold), AI for Weather
- **Hardware Need:**
  - Chips that handle BOTH high precision (Science) and tensors (AI)

Summary of Evolution

- Custom Hardware
- → Commodity Clusters
- → Heterogeneous Accelerators
- → AI Convergence
- → Next?

- **Driver:** Application needs dictate Hardware design

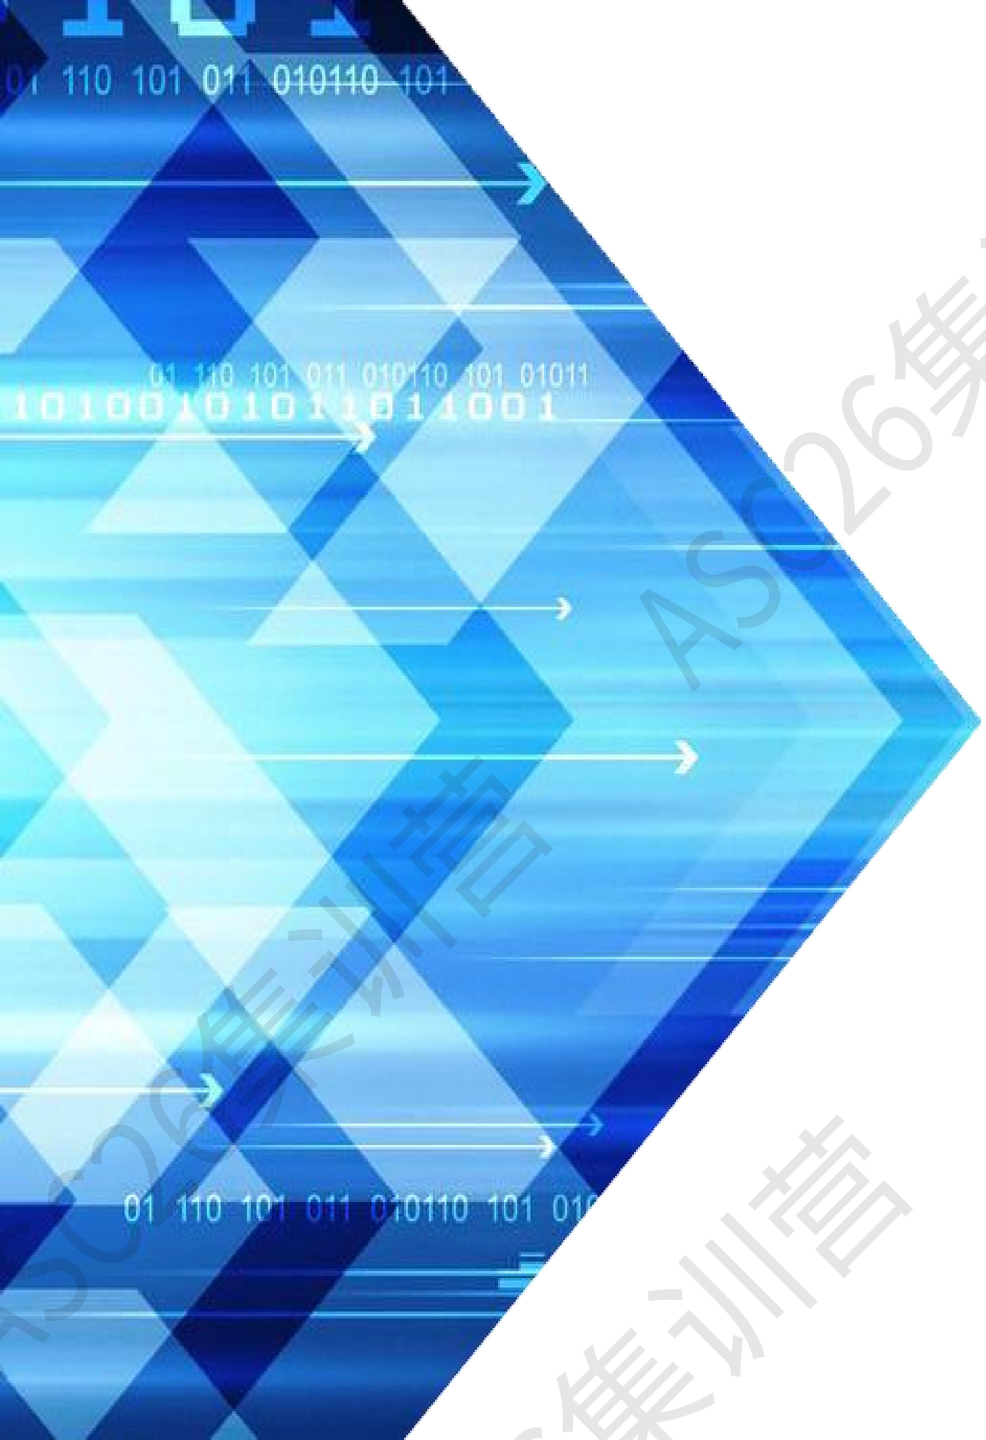

Outline:

- Background
- History
- **Deep Dive**
  - **Processor, Interconnect & Storage**
- System Design

A Server Node

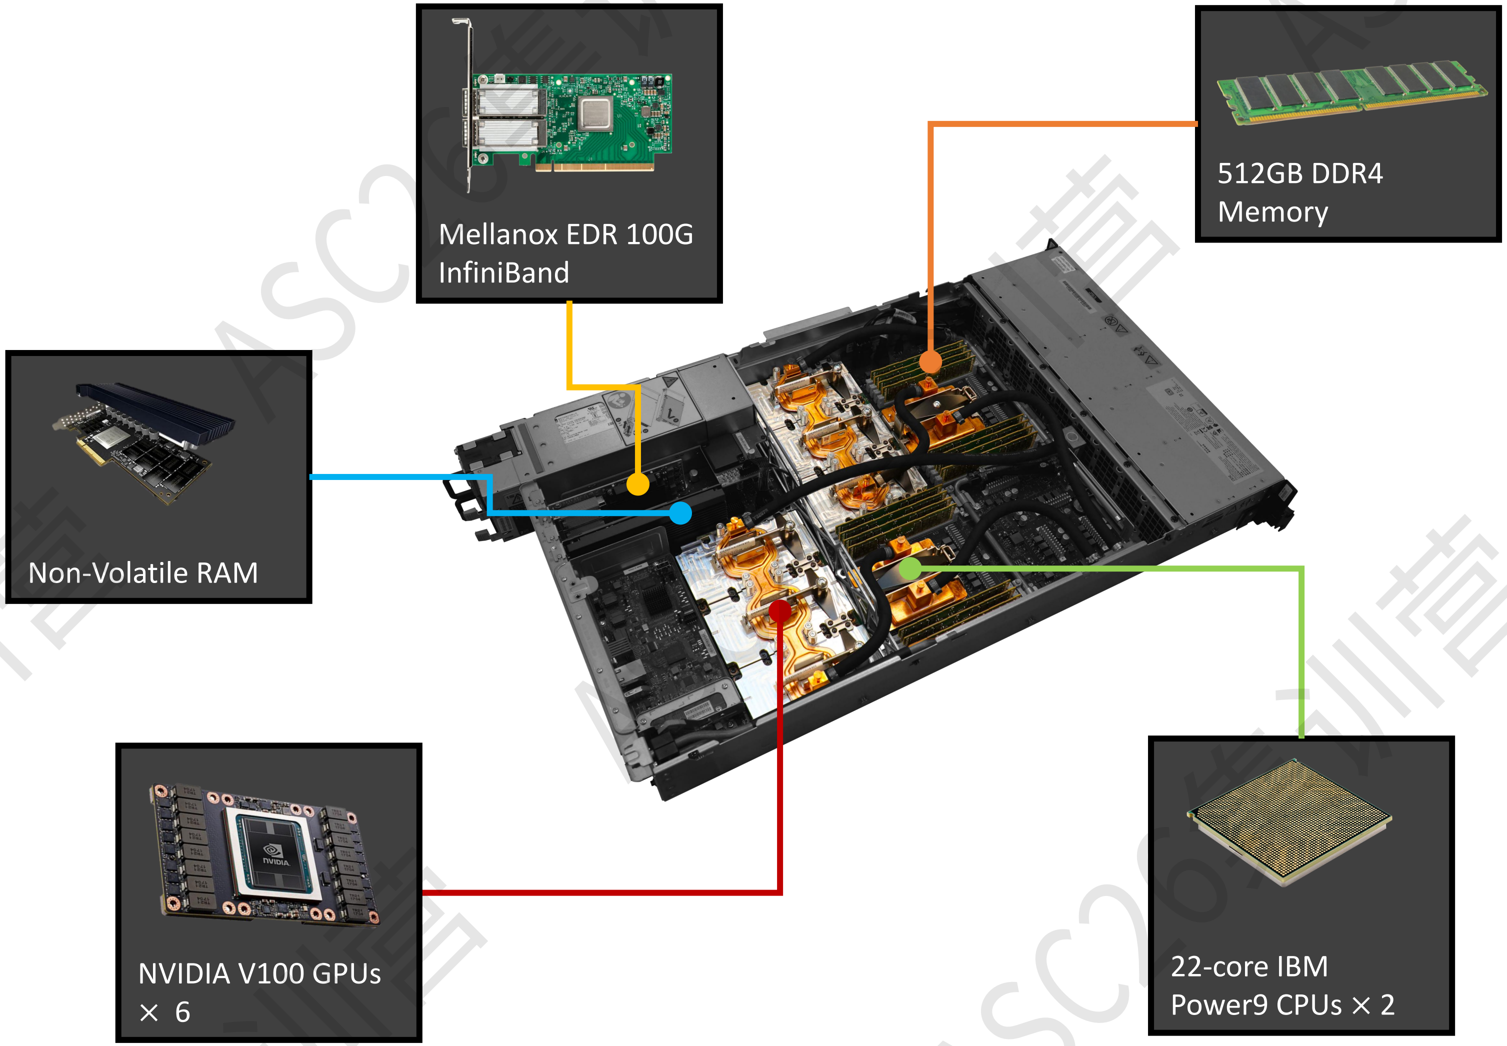

CPU Microarchitecture

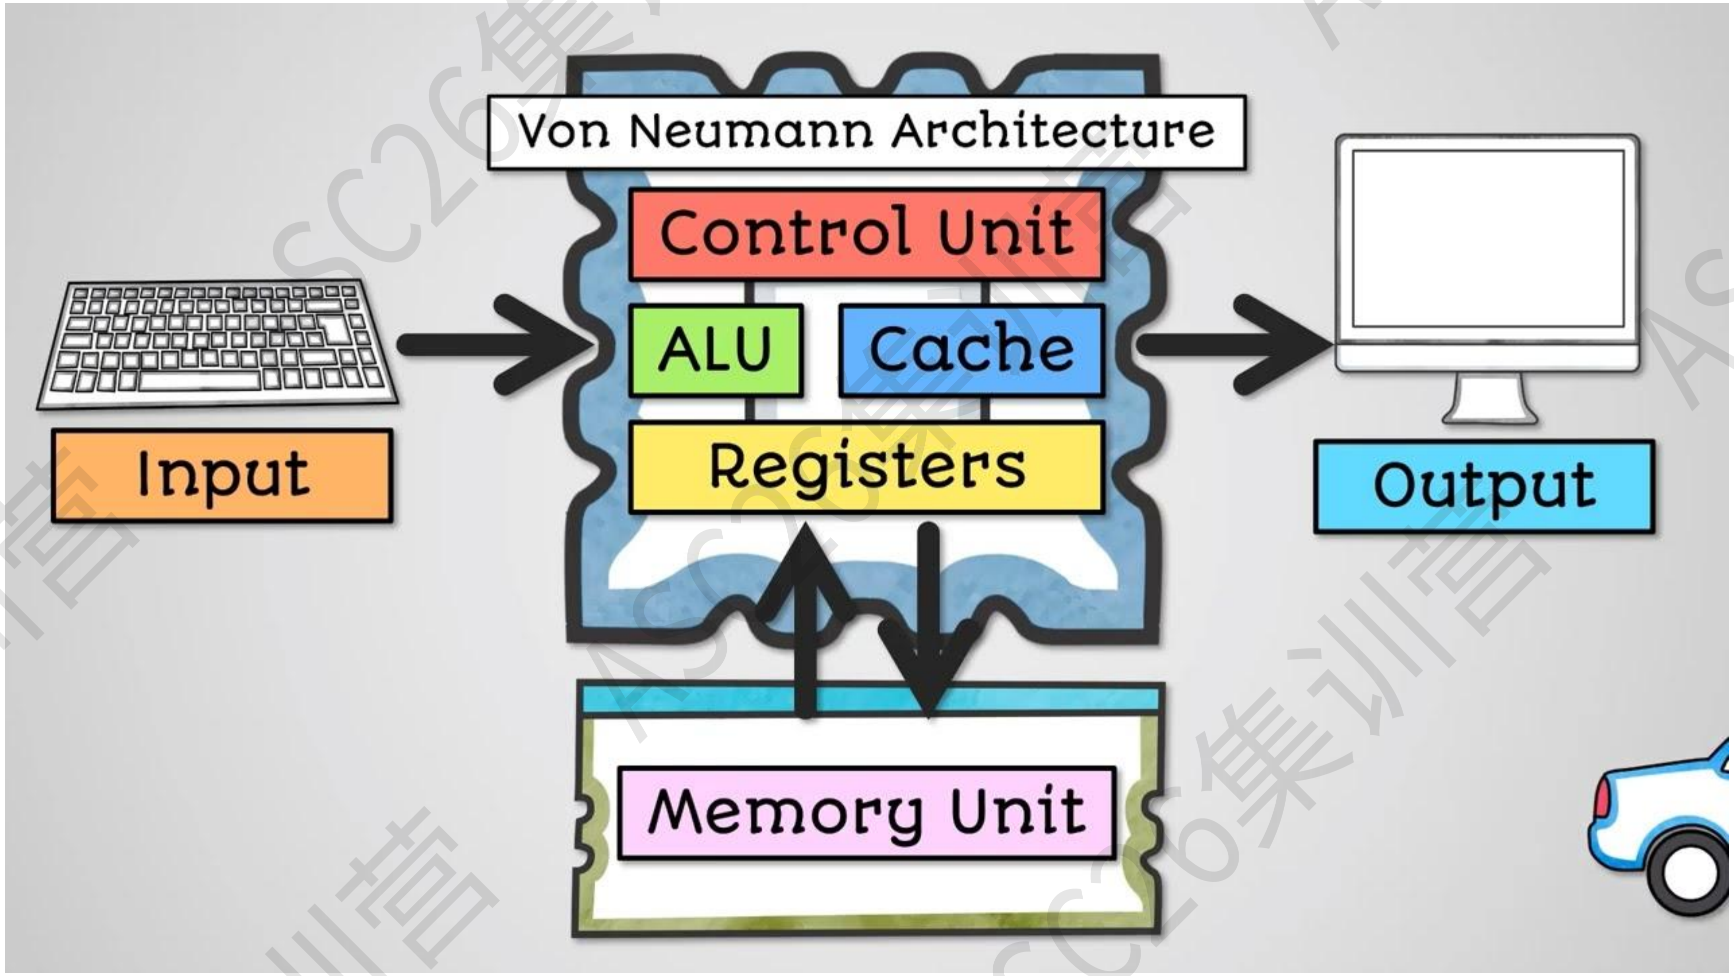

CPU Microarchitecture

For High Performance:

- **SIMD:**
  - Single Instruction Multiple Data (AVX-512)
- **Core Count:**
  - Increasing (64, 96, 128 cores per socket)
- **Memory Hierarchy:**
  - L1/L2/LL Cache, NUMA

CPU Microarchitecture

- **SIMD:**
  - Single Instruction Multiple Data (AVX-512)

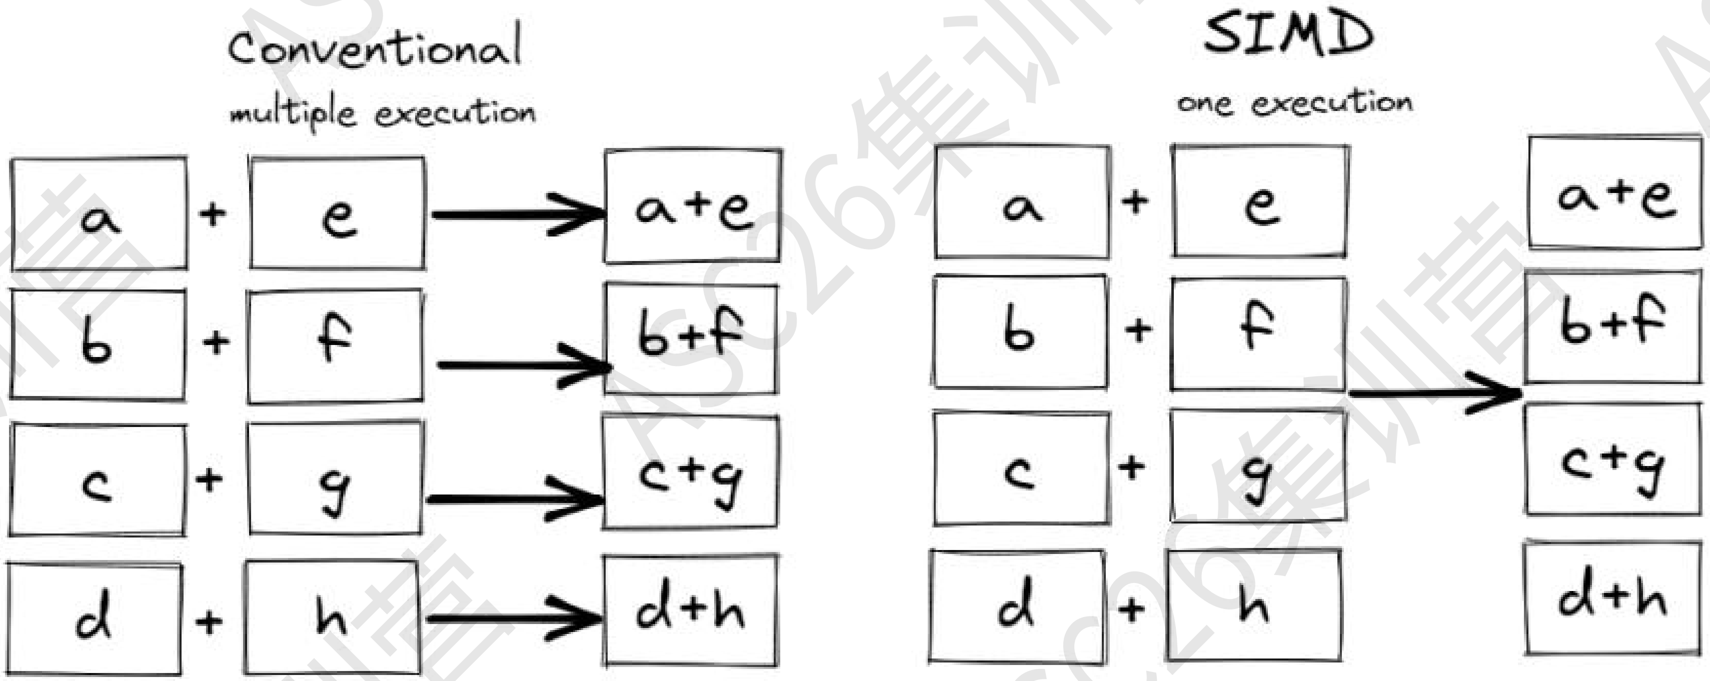

CPU Microarchitecture

- **Core Count:**
  - Increasing (64, 96, 128 cores per socket)

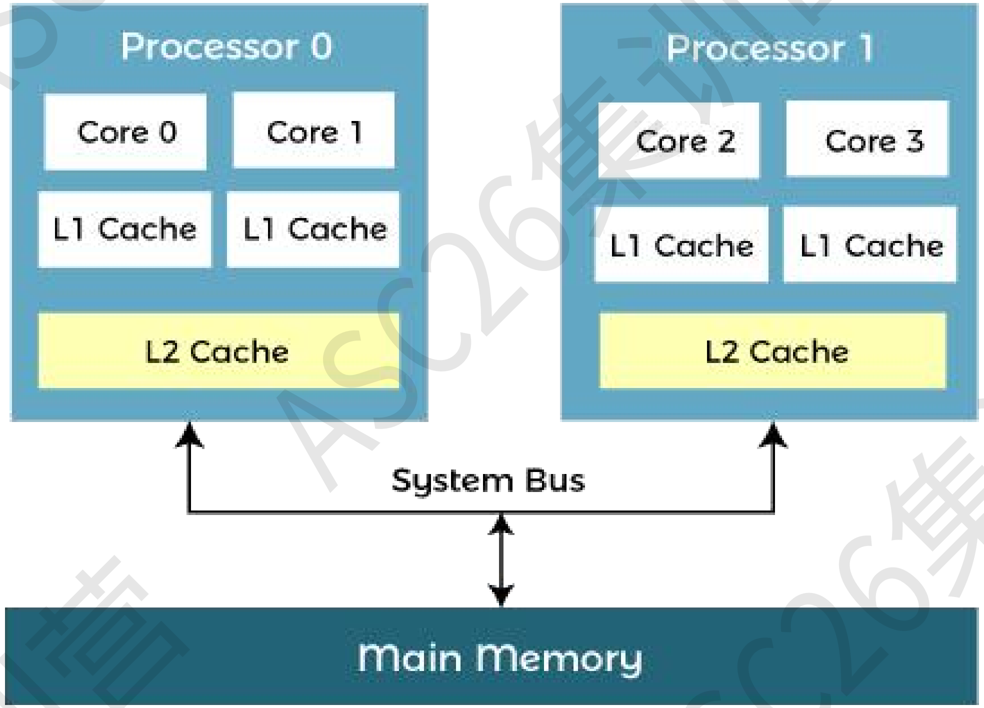

CPU Microarchitecture

- **Memory Hierarchy:**
  - L1/L2/LL Cache, NUMA

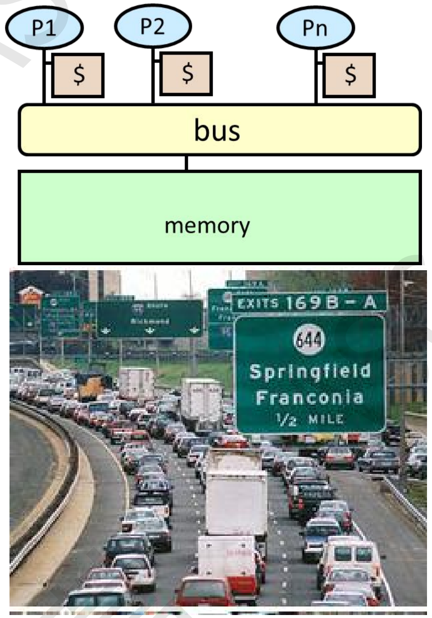

Uniform Memory Access (UMA)

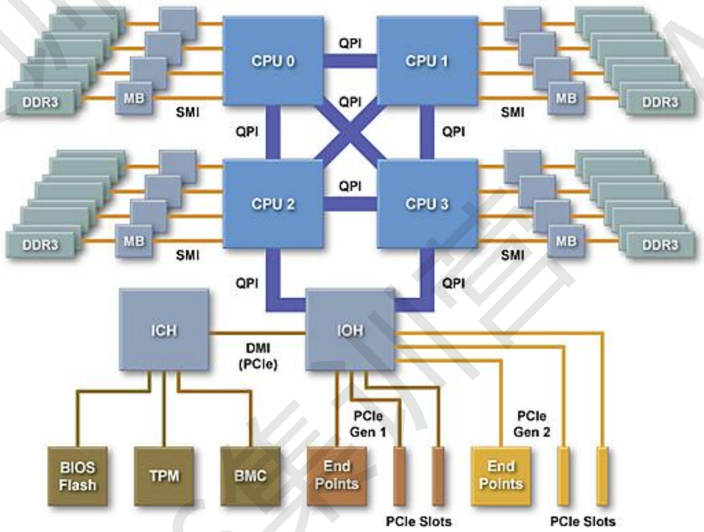

Non-Uniform Memory Access (NUMA)

GPU Architecture

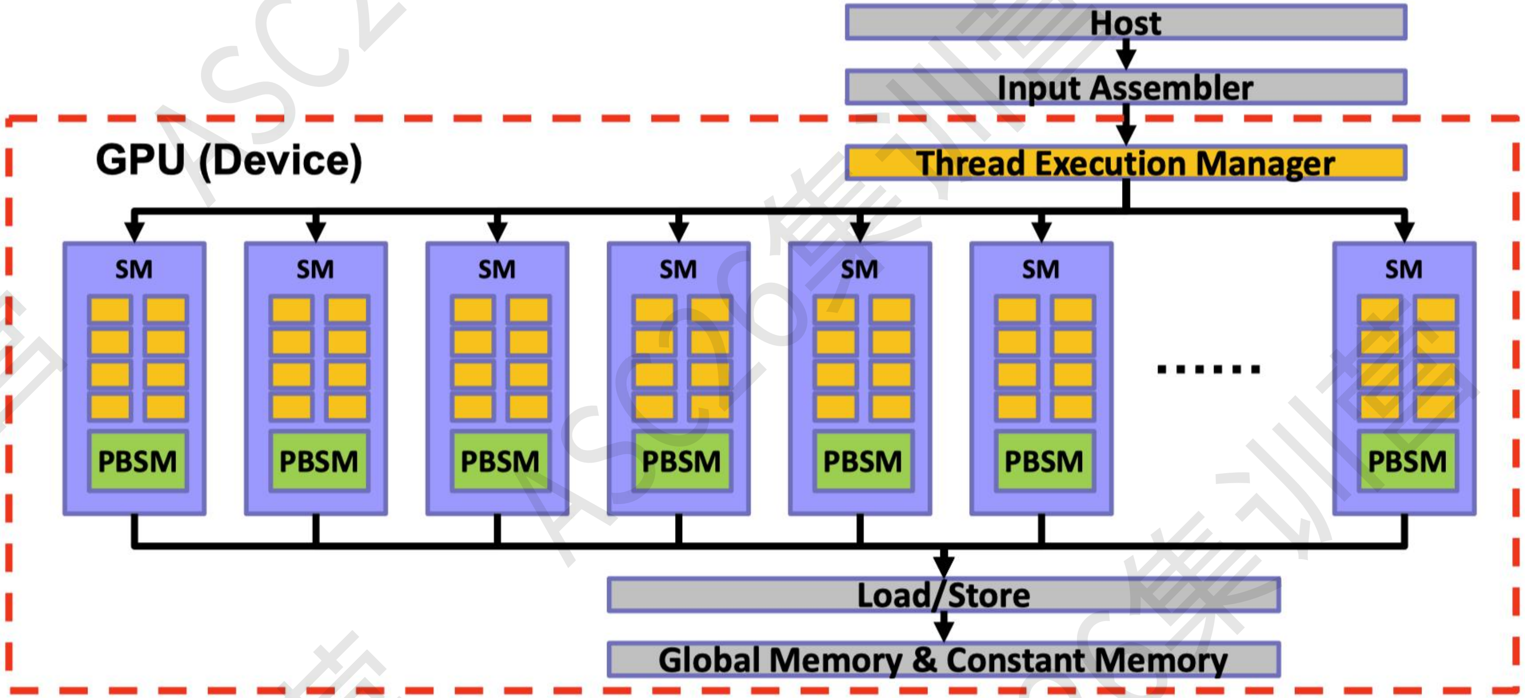

GPU Architecture

- **SM (Streaming Multiprocessors):**
  - The building blocks
- **Memory Type:**
  - HBM (High Bandwidth Memory) - Critical for speed
- **Bandwidth:**
  - TB/s level (much faster than DDR)

CPU vs. GPU

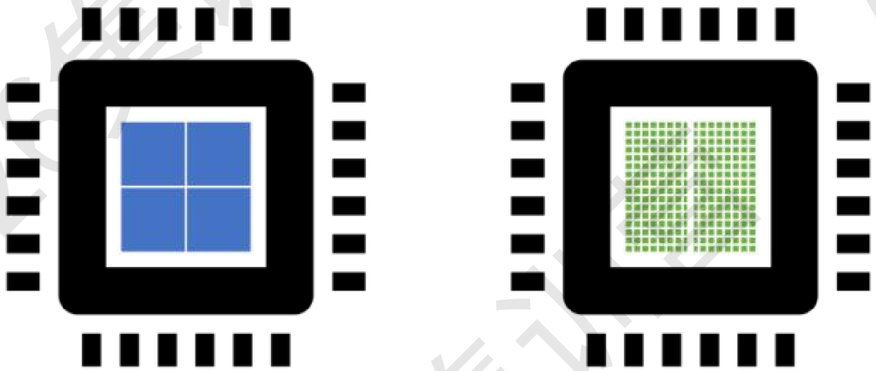

CPU	GPU
Central Processing Unit	Graphics Processing Unit
4-8 Cores	100s or 1000s of Cores
Low Latency	High Throughput
Good for Serial Processing	Good for Parallel Processing
Quickly Process Tasks That Require Interactivity	Breaks Jobs Into Separate Tasks To Process Simultaneously
Traditional Programming Are Written For CPU Sequential Execution	Requires Additional Software To Convert CPU Functions to GPU Functions for Parallel Execution

Interconnects: The Nervous System

- **Role:**
  - Connecting nodes to act as one system
- **Protocols:**
  - InfiniBand (IB), RoCE (Ethernet), Omni-Path
- **Standard:**
  - InfiniBand (HDR 200Gb, NDR 400Gb)

Key Network Metrics

- **Bandwidth:**
  - Throughput (Gbps).
- **Latency:**
  - Delay (Microseconds).
- **Significance:**
  - Low latency is vital for synchronization in parallel apps (MPI).

InfiniBand

- SDR - Single Data Rate
- DDR - Double Data Rate
- QDR - Quad Data Rate
- FDR - Fourteen Data Rate
- EDR - Enhanced Data Rate
- HDR - High Data Rate
- NDR - Next Data Rate

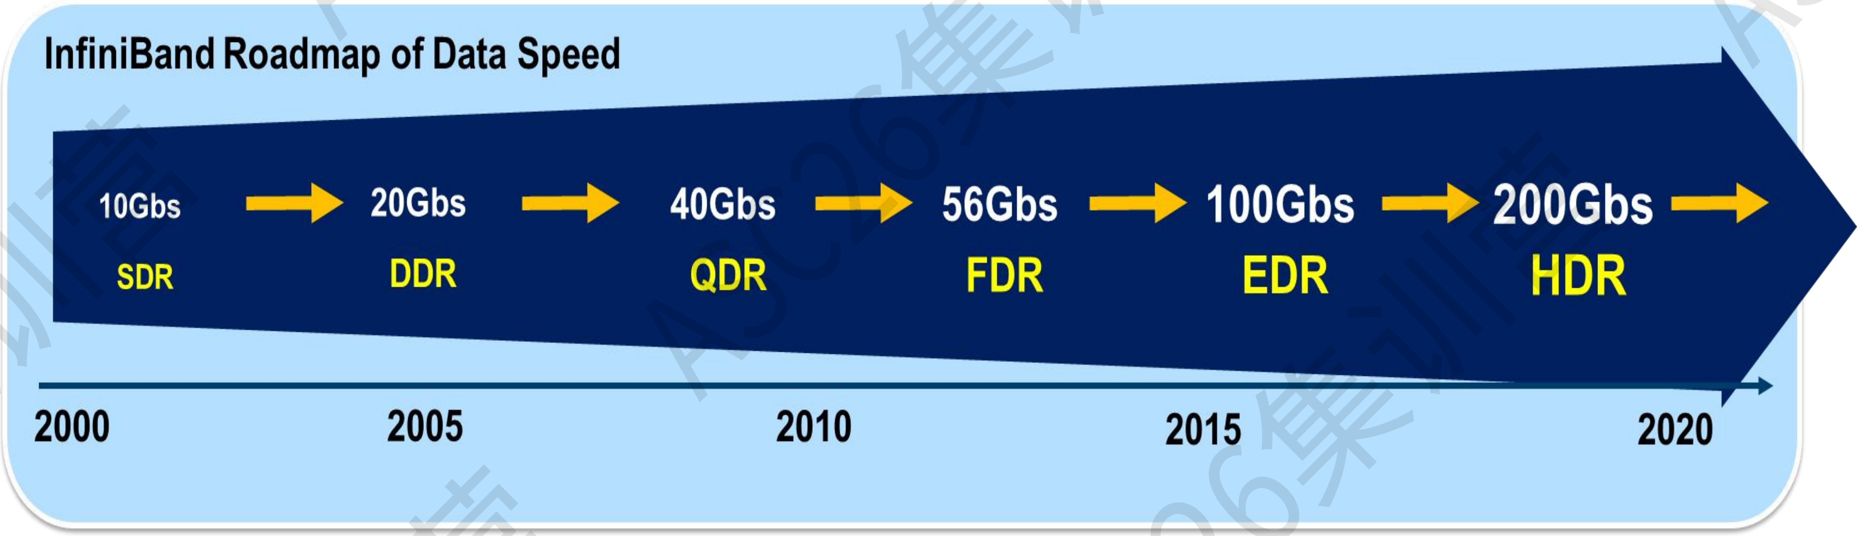

NVLink

8 GPUs, inter-GPU 900GB/s, 7.2TB/s in total

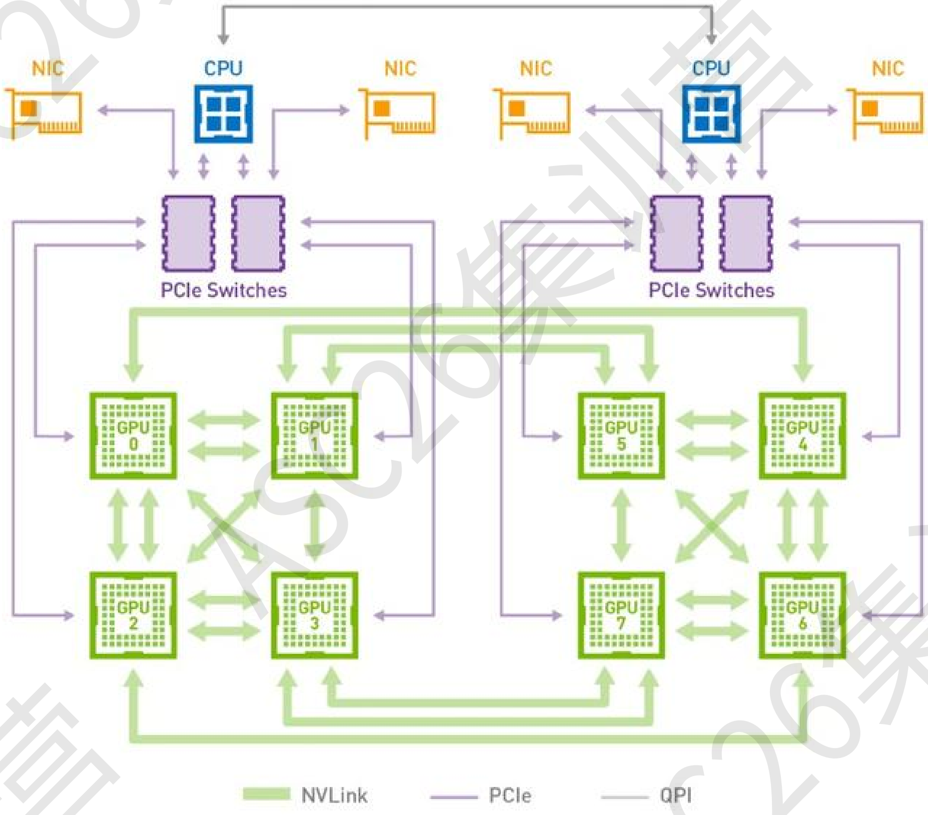

Network Topology

- **Fat-Tree:**
  - Non-blocking, most common in competitions.
- **Dragonfly / Torus:**
  - For massive scale systems.
- **Switch Hierarchy:**
  - Spine and Leaf switches.

Storage Hierarchy

- **Pyramid:**
  - Registers -> L1/L2/L3 Cache -> HBM -> RAM -> NVMe SSD -> HDD
- **Speed vs. Capacity:**
  - The lower you go, the slower and larger it gets.

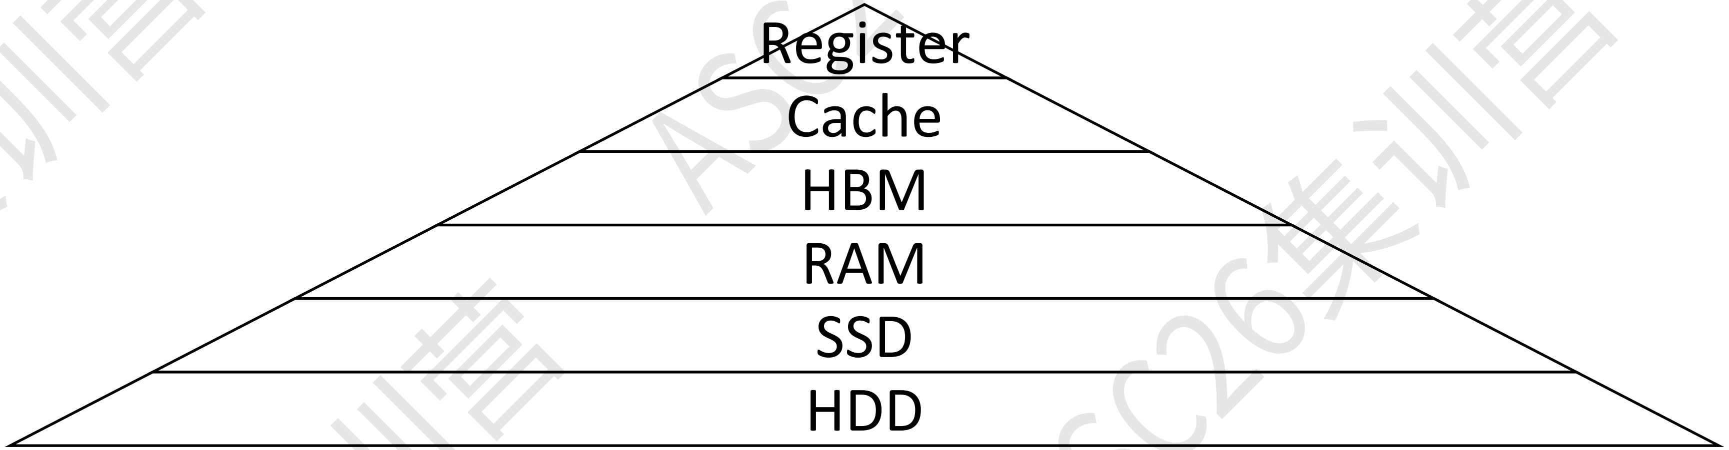

Parallel File Systems

- **Problem:**
  - NFS is too slow for clusters.
- **Solution:**
  - Lustre, BeeGFS, GPFS.
- **Mechanism:**
  - Striping data across multiple Object Storage Targets (OSTs).

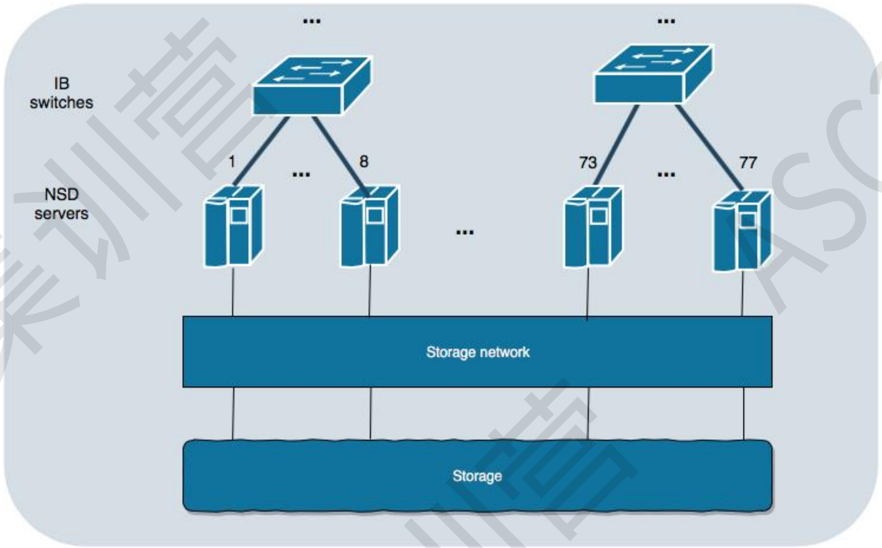

IBM Spectrum Scale Filesystem (Summit)

New architectures

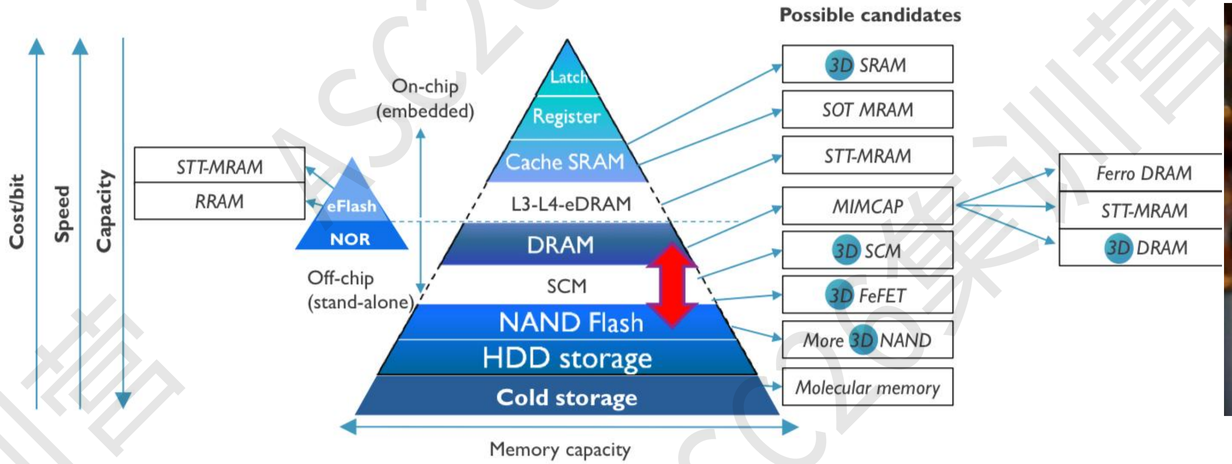

New memory architecture

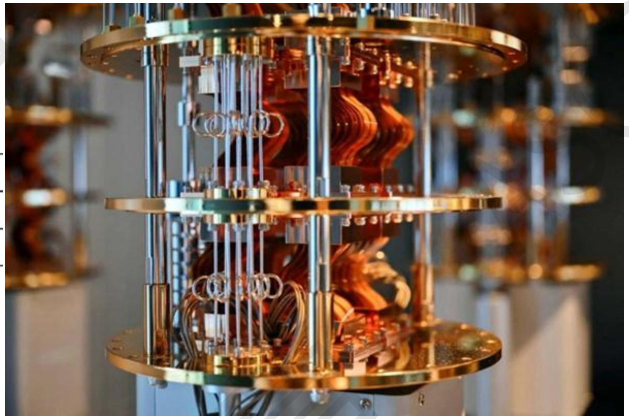

Quantum computer

Summary of HPC Architecture

- **Processor**
  - CPU, GPU, TPU, …
- **Interconnect**
  - InfiniBand, NVLink
  - Topology
- **Storage**
  - Parallel IO

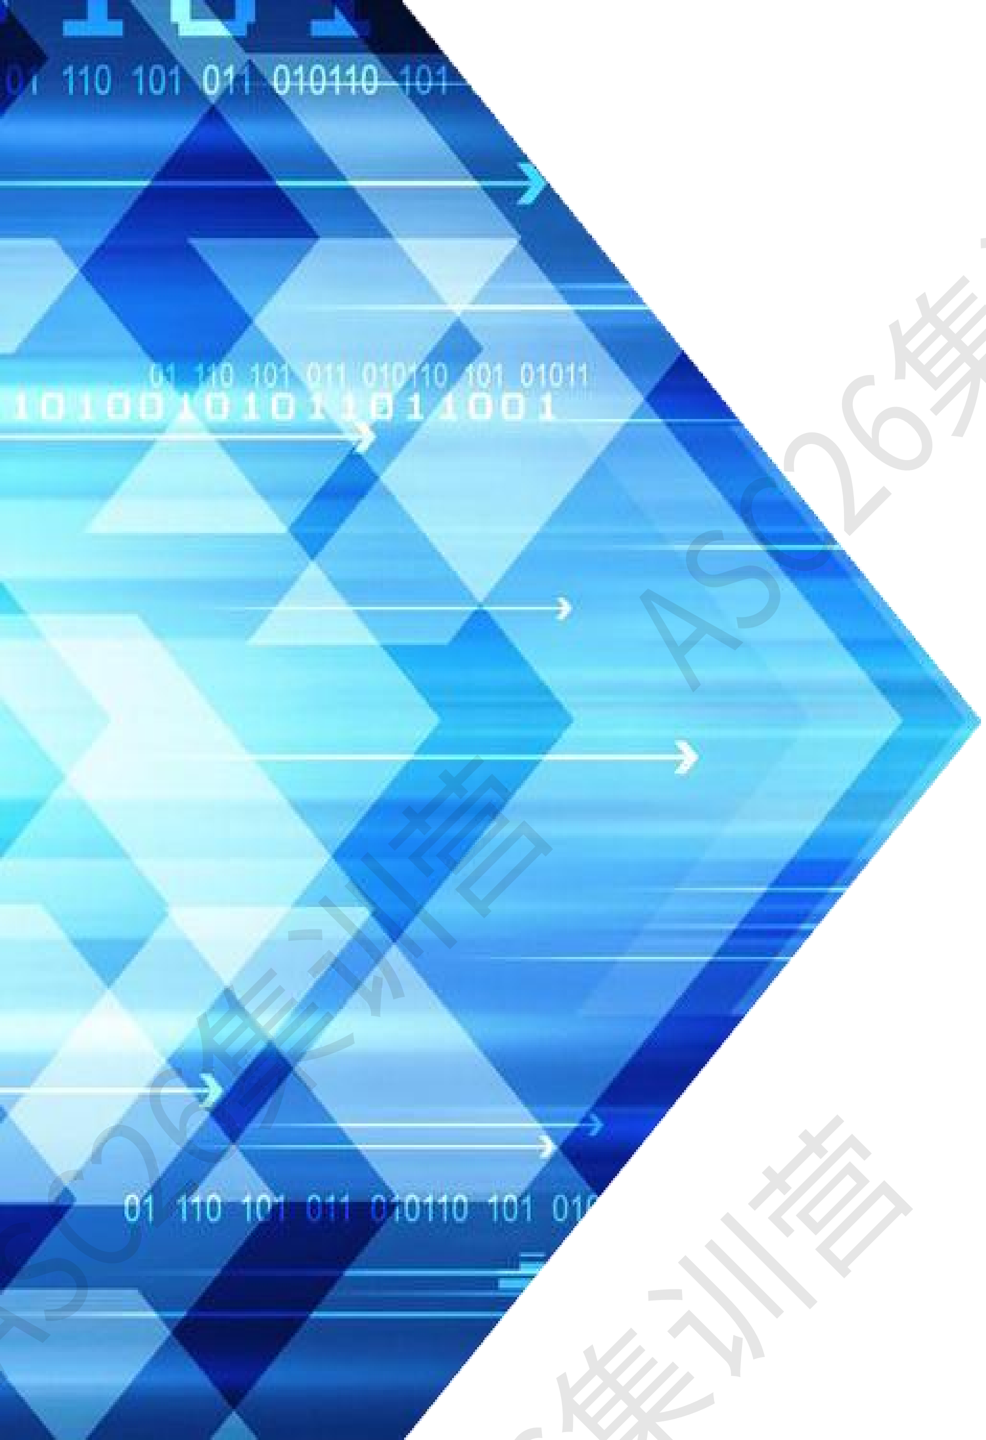

Outline：

- Background
- History
- Deep Dive
  - Processor, Interconnect & Storage
- **System Design**

Intro to SCC

- **Program:**
  - Benchmark + several applications
- **System:**
  - Self-configured hardware/software
  - Limited power
- **Target:**
  - Overall highest performance
  - Highest performance of a single application

Target is Important

- **Applications have different workloads and different hardware requirements**
  - HPL (flops)
  - HPCG (memory bandwidth)
  - Graph500 (memory size)
  - …
- **What is your target?**

Application Workload

- Understand algorithms
- Running & Profiling
- Details are not introduced here

Basic Components

- X server nodes + InfiniBand switch
- Each server node:
  - CPU
  - Accelerators (GPU?)
  - Memory
  - SSD

The Power Challenge

The Power Challenge: Basic

Where Do the Watts Go?

- Switch: ~200W
- Server node:
  - **GPU**: ~350W, The biggest consumer. Highly Controllable.
  - **CPU**: ~250W, Second biggest. Highly Controllable.
  - **Memory**: ~5-10W per DIMM. 16 DIMMs = ~150W! Hard to control dynamic power.
  - **Fans**: Power – $Speed^3$. High speed = Huge power. Up to 330W, Semi-controllable (IPMI).
  - **Baseboard/VRM/SSD**: Static overhead. Fixed.
  - **PSU Loss**: 94%, 2000W PDU ~2127W

The Power Challenge: Scalability

The Efficiency Game: Scaling vs. Power

- More GPUs/nodes are better?

- GPU Power Capping Strategy：
  - Running 4 GPUs at 250W each (Total 1000W)
  - 3 GPUs at 330W each (Total 1000W)

- Fan or CPU/GPU frequency?

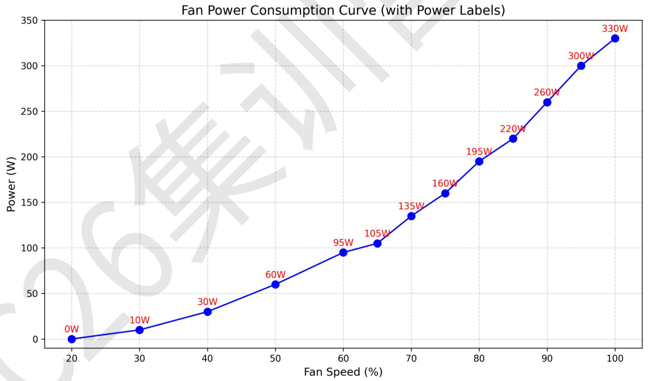

The Power Challenge: Monitor

Monitoring: Power, Temperature & Leakage

- **The Source of Truth:**
  - **PDU** (Power Distribution Unit): The official judge. Always calibrate your internal sensors against the PDU.
- **Monitoring Tools:**
  - Hardware: BMC/IPMI (Out-of-band).
  - Software: nvidia-smi (GPU), turbostat (CPU), ipmitool (System).
  - Visualization: Grafana + Prometheus (Real-time dashboard).
- **The Temperature Vicious Cycle:**
  - High Temp -> Higher Leakage Current -> Higher Power -> Throttling.

The Power Challenge: Monitor

- THUSCC Diablo team's monitor

The Power Challenge: Monitor

- Power of all teams → ?

Summary

- History of cluster computing
- HPC system architecture introduction
- How to design a cluster for SCC

Superpower behind Supercomputers

The superpower is your passion for supercomputing

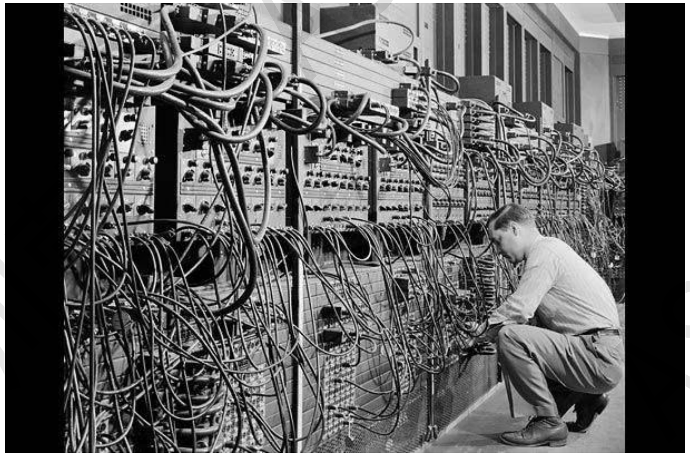

10^{15} X faster

80 years

The first computer in the world
1946, 5KFLOPS

Top-1 supercomputer in the world
2026, 2.7EFLOPS

Thanks!
Good Luck!

Yuyang Jin

jinyuyang@tsinghua.edu.cn