Introduction to High Performance Computers

Richard Gerber

NERSC User Services

U.S. DEPARTMENT OF ENERGY | Office of Science

NERSC National Energy Research Scientific Computing Center

BERKELEY LAB Lawrence Berkeley National Laboratory

Why Do You Care About Architecture?

- **To use HPC systems well, you need to understand the basics and conceptual design**
  - Otherwise, too many things are mysterious
- **Programming for HPC systems is hard**
  - To get your code to work properly
  - To make it run efficiently (performance)
- **You want to efficiently configure the way your job runs**
- **The technology is just cool!**

Outline

- Terminology
- 5 main parts of an HPC system
- CPUs
- Nodes
- Interconnect
- Data Storage
- HPC Systems

What are the main parts of a computer?

Boy Scouts of America
Offer a Computers Merit Badge

Merit Badge Requirements

…

4. Explain the following to your counselor:

a. The five major parts of a computer.

…

What are the “5 major parts”?

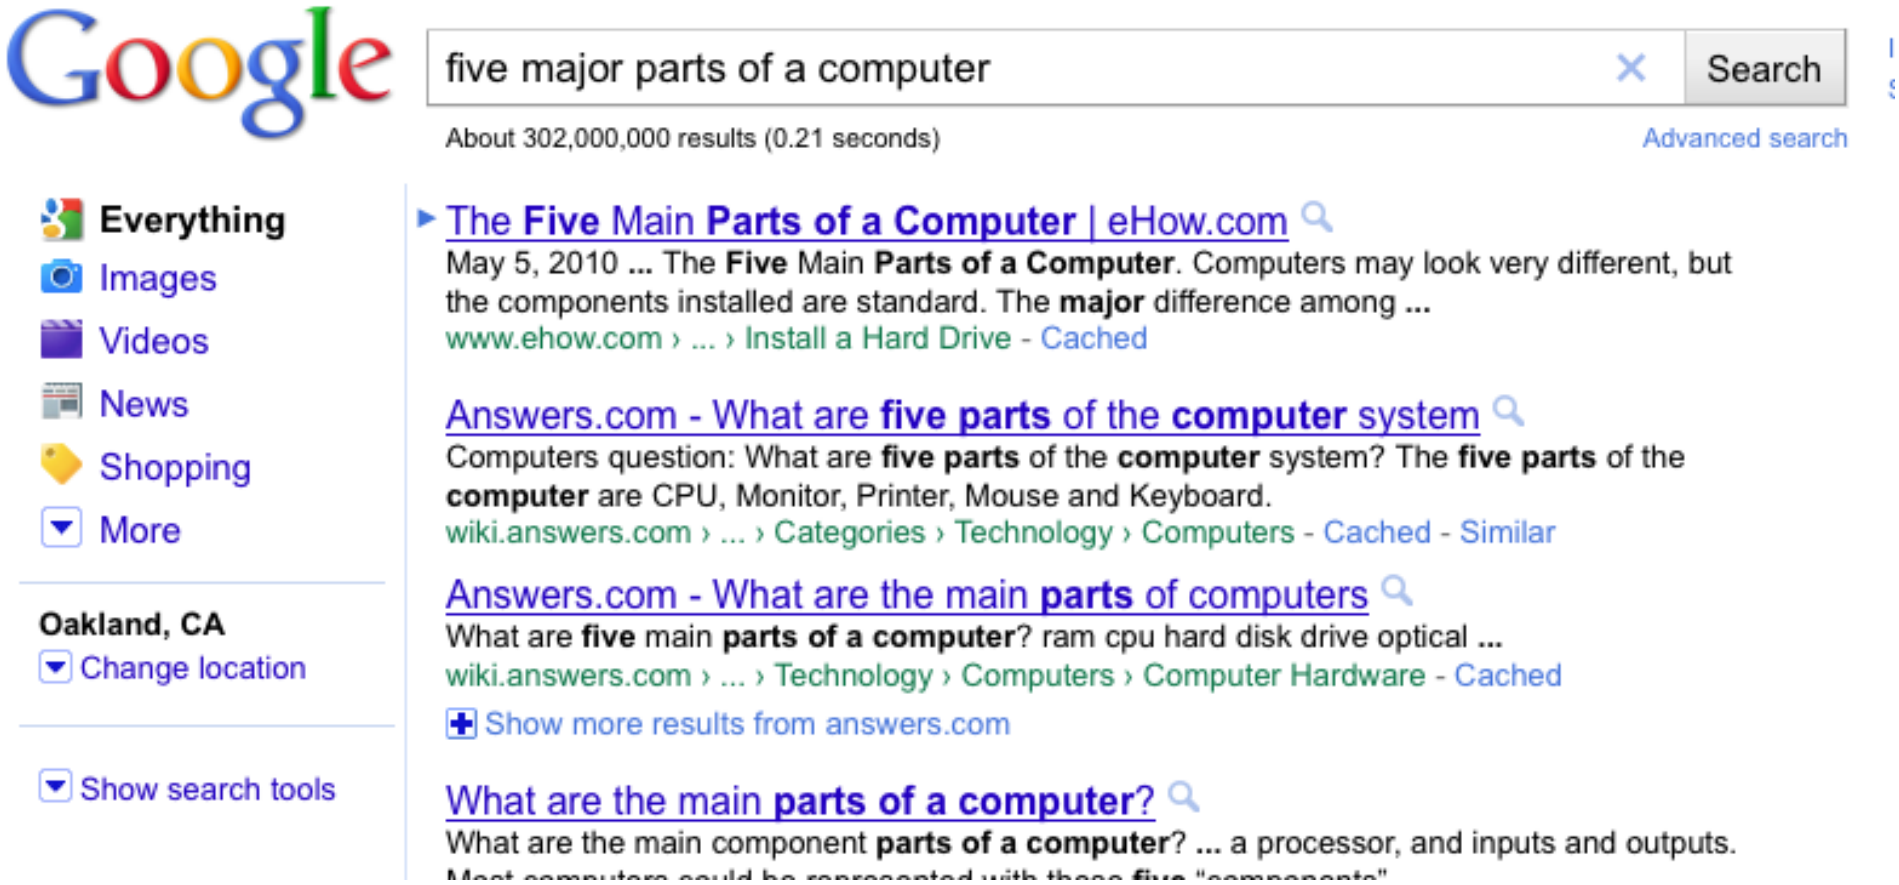

Five Major Parts

eHow.com	Answers.com	Fluther.com	Yahoo!	Wikipedia
CPU	CPU	CPU	CPU	Motherboard
RAM	Monitor	RAM	RAM	Power Supply
Hard Drive	Printer	Storage	Power Supply	Removable Media
Video Card	Mouse	Keyboard/ Mouse	Video Card	Secondary Storage
		Monitor		
Motherboard	Keyboard	Motherboard	Motherboard	Sound Card
		Case / Power Supply		IO Peripherals

It Depends on Your Perspective

- What is a computer?
  - It depends what you are interested in.
    - CPU, memory, video card, motherboard, ...
    - Monitor, mouse, keyboard, speakers, camera, …
- We'll take the perspective of an application programmer or a scientist running a code on an HPC system
- What features of an HPC system are important for you to know about?

5 Major Parts of an HPC System

1. CPUs
2. Memory (volatile)
3. Nodes
4. Inter-node network
5. Non-volatile storage (disks, tape)

Definitions & Terminology

- **HPC**
  - High Performance Computing
  - Scientific computing at scale
- **CPU**
  - Central Processing Unit
  - Now ambiguous terminology
  - Generic for “some unit that computes”
  - Context-sensitive meaning
- **Core (Intel has this reversed with CPU)**
  - Hardware unit that performs arithmetic operations
  - A CPU may have more than one core
- **Die**
  - An integrated circuit manufactured as a unit
  - Many cores may be included on a die
- **Socket**
  - A physical package that connects to a computer board
  - A socket package may be composed of multiple dies

Definitions & Terminology

- **FLOP: Floating Point Operation**
  - e.g., `a+b`, `a*b+c`
  - FLOPs/sec is a common performance metric
- **SMP**
  - Defn: Symmetric Multiprocessing
  - Common usage: Collection of processors that have (approx equal) access to a shared pool of memory in a single memory address space
- **FORTRAN**
  - Programming language popular with scientists, esp. in HPC
  - Unpopular with Computer Scientists (who may make fun of you)
- **MPP**
  - Massively parallel processing
  - You must be running an MPP code on an MPP computer to be considered doing HPC
- **Interconnect**
  - A high-performance data network that connects nodes to each other and possibly other devices

Definitions & Terminology

- **Memory**
  - Volatile storage of data or computer instructions
- **Bandwidth**
  - The rate at which data is transferred between destinations (typically GB/s)
- **Latency**
  - The time needed to initialize a data transfer (ranges from 10^{-9} to 10^{-6} secs or more)
- **SRAM: Static RAM (random access memory)**
  - Fast
  - 6 transistors store a bit
    - Expensive
    - Limits storage density
- **DRAM: Dynamic RAM (random access memory)**
  - Slower
  - 1 transistor, 1 capacitor stores a bit
    - Higher density, cheaper
  - Capacitor voltage needs to be refreshed
    - Additional power

A distributed-memory HPC system

A distributed-memory HPC system

A distributed-memory HPC system

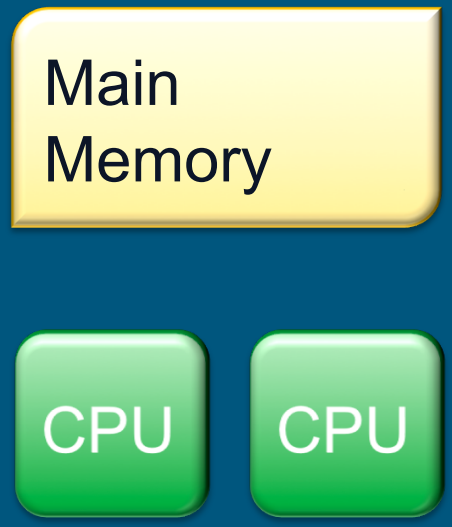

A distributed-memory HPC system

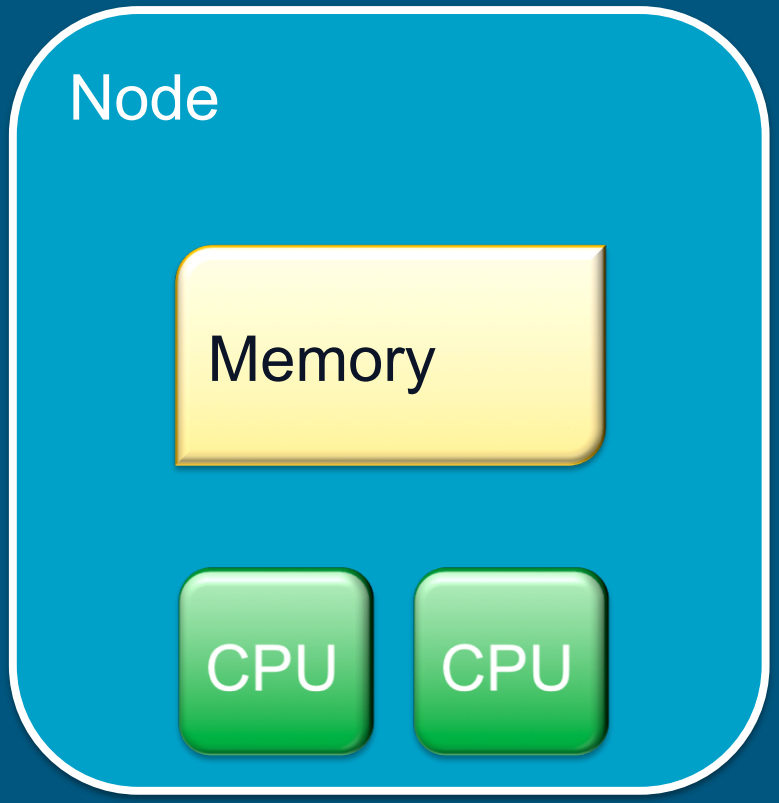

A distributed-memory HPC system

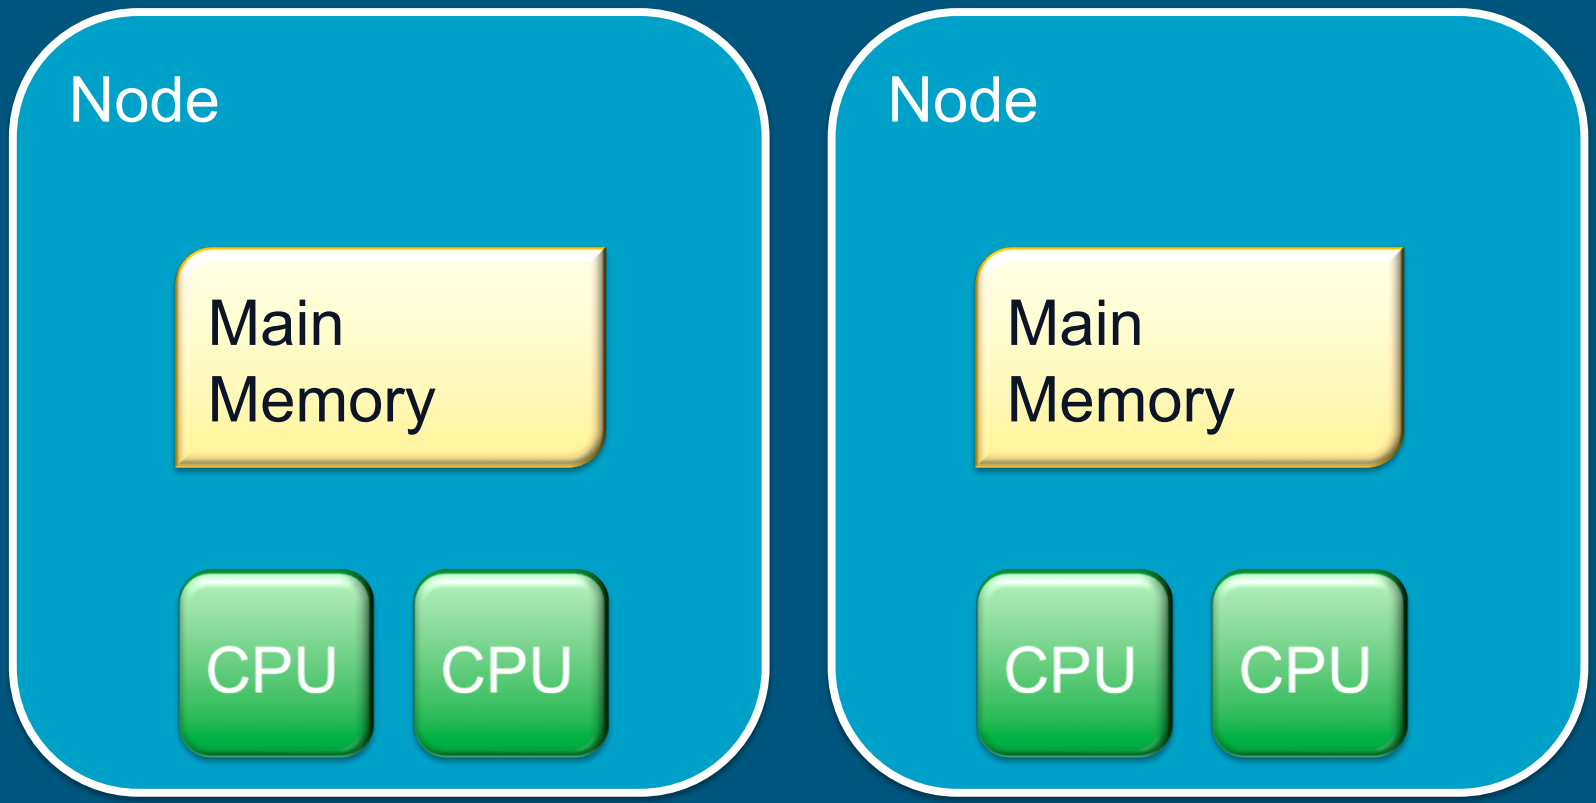

A distributed-memory HPC system

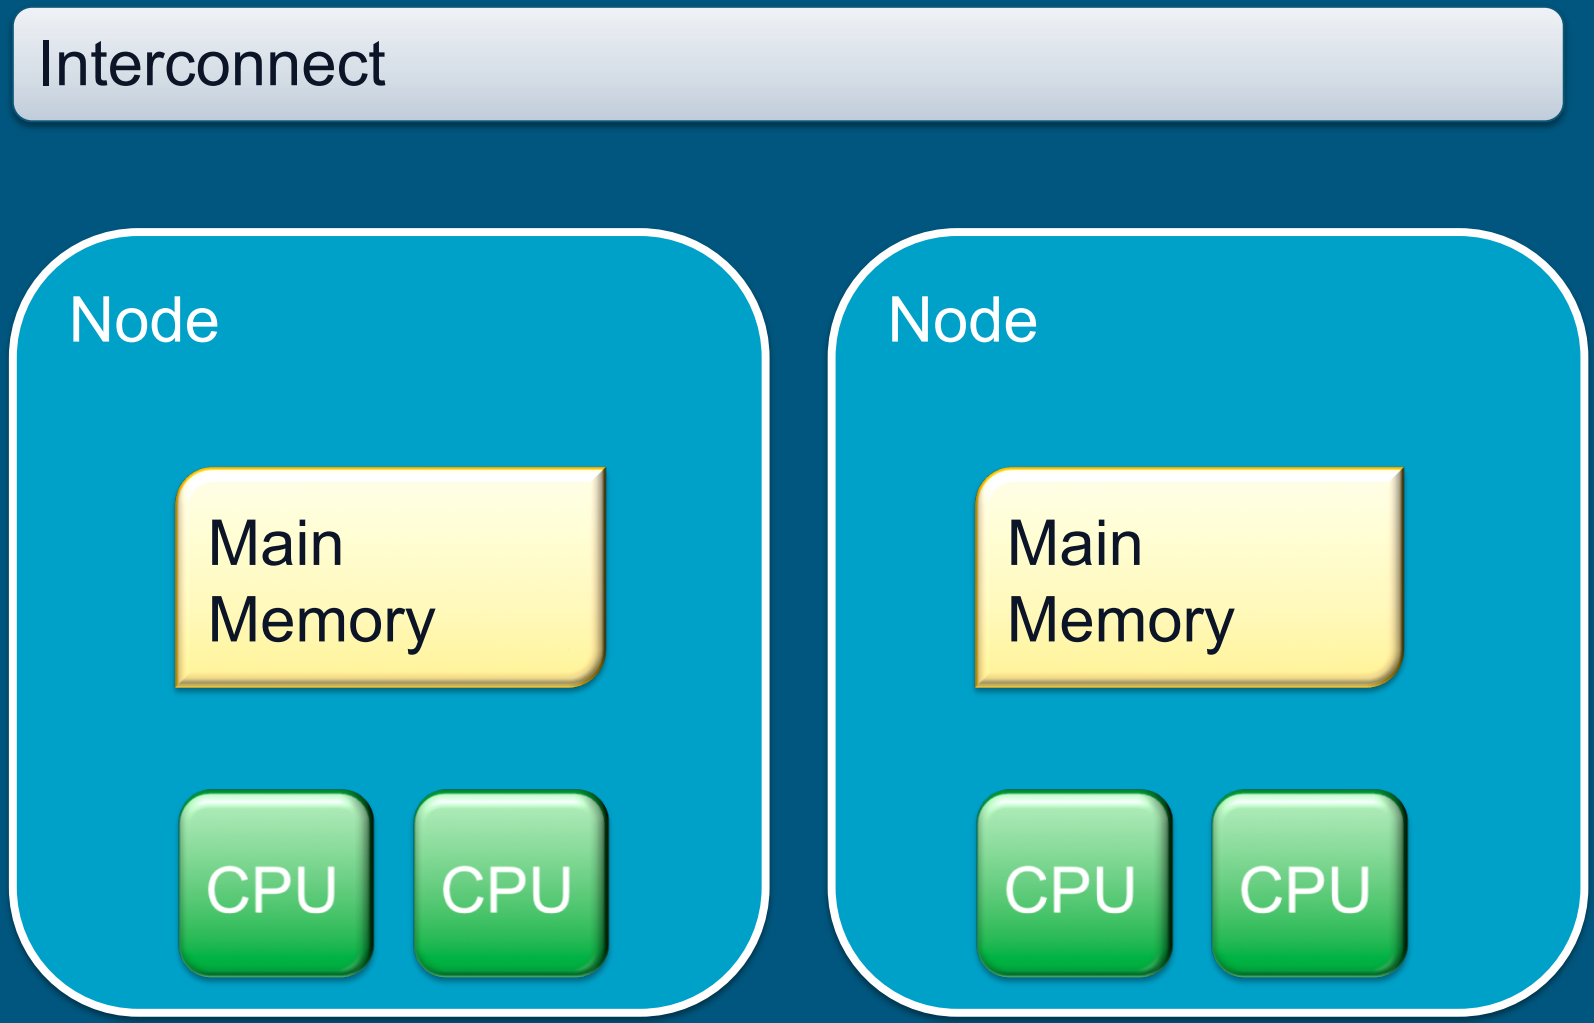

A distributed-memory HPC system

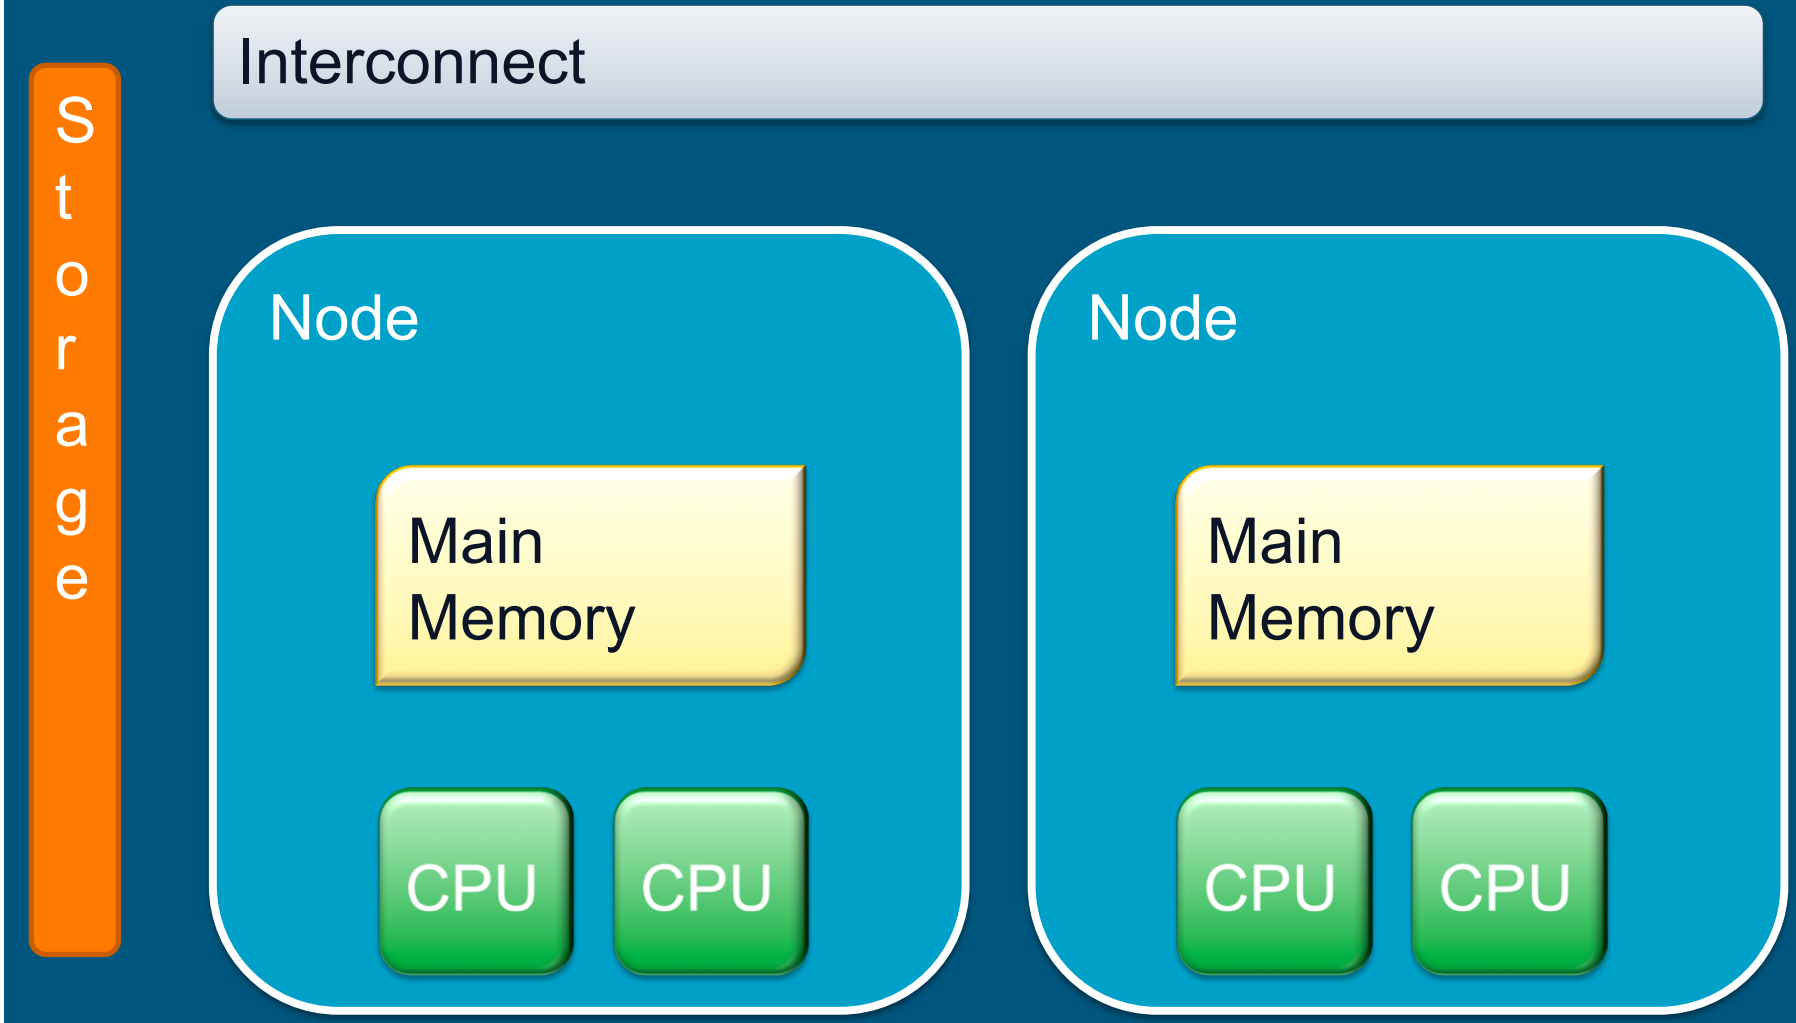

A distributed-memory HPC system

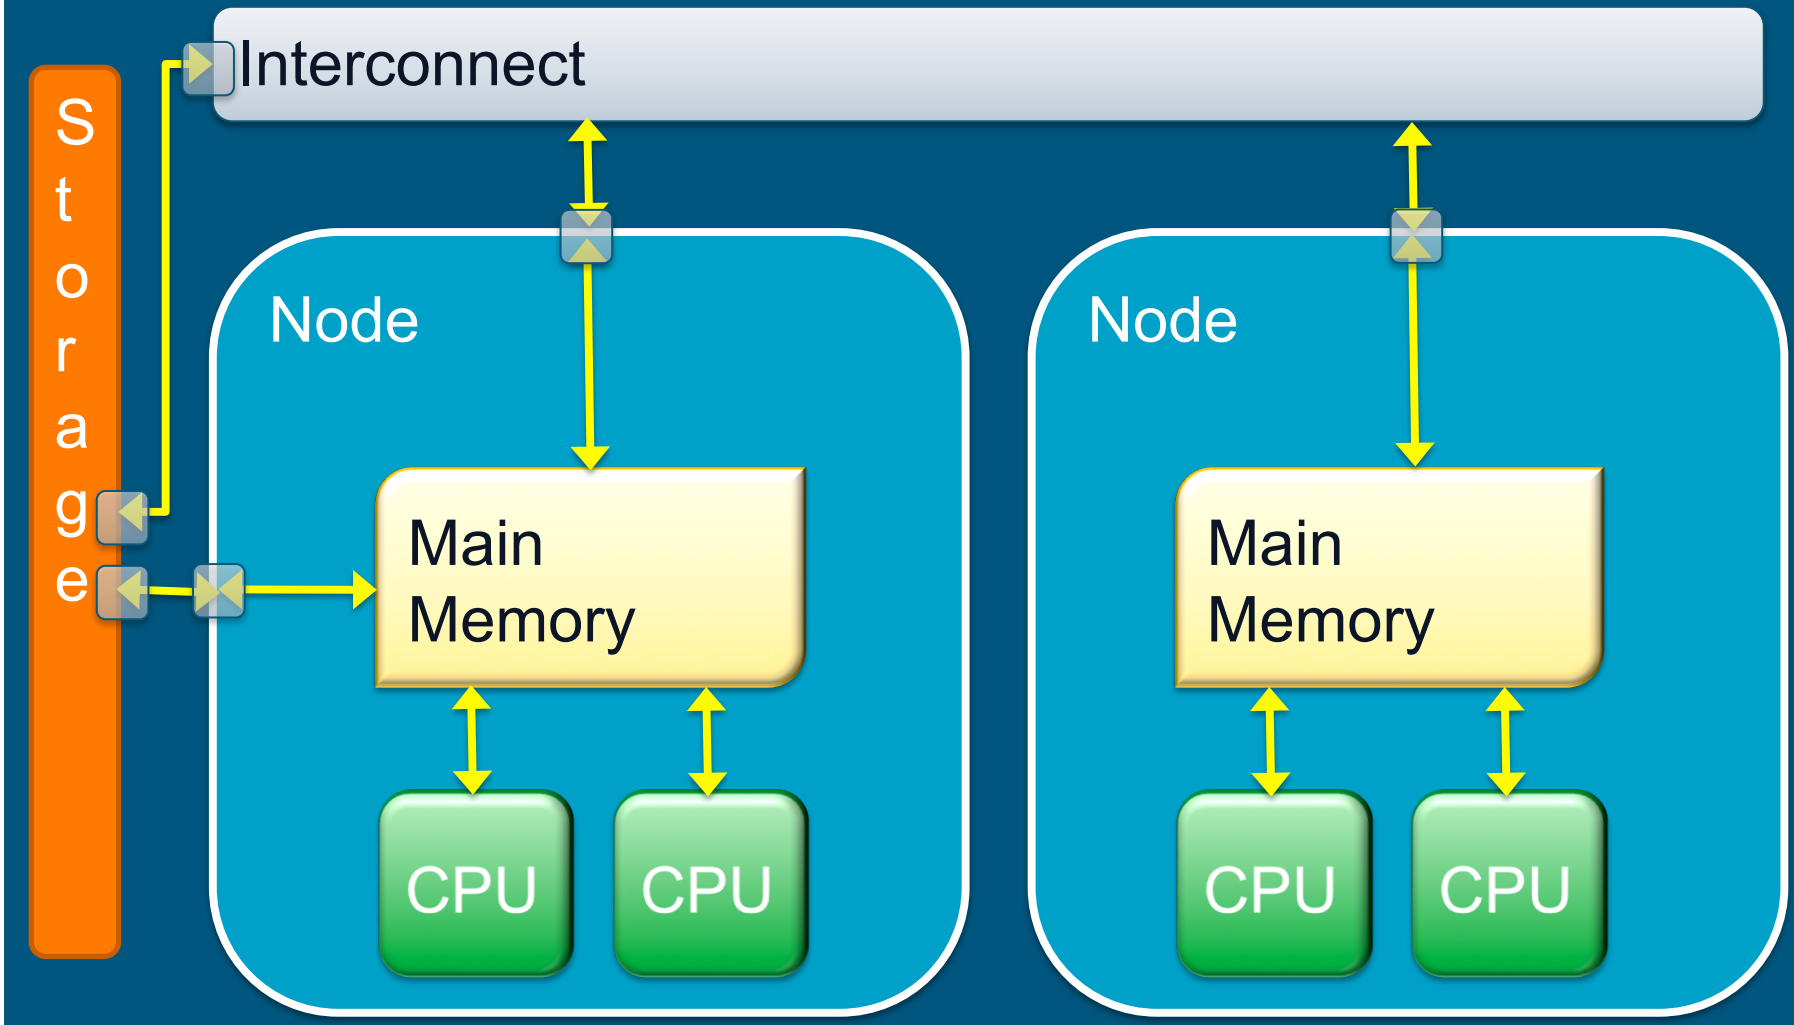

5 Major Parts of an HPC System

1. **CPUs**
2. **Memory (volatile)**
3. Nodes
4. Inter-node network
5. Non-volatile storage (disks, tape)

Stored Program Computers

- **Modern computers are "stored program computers"**
  - Conceived by Turing in 1936
  - Implemented in 1949 (EDVAC)
- **Instructions are stored as data in memory**
  - Read and executed by control units
- **Arithmetic and logic**
  - Performed by functional units separate from instruction control units

CPU (1 core)

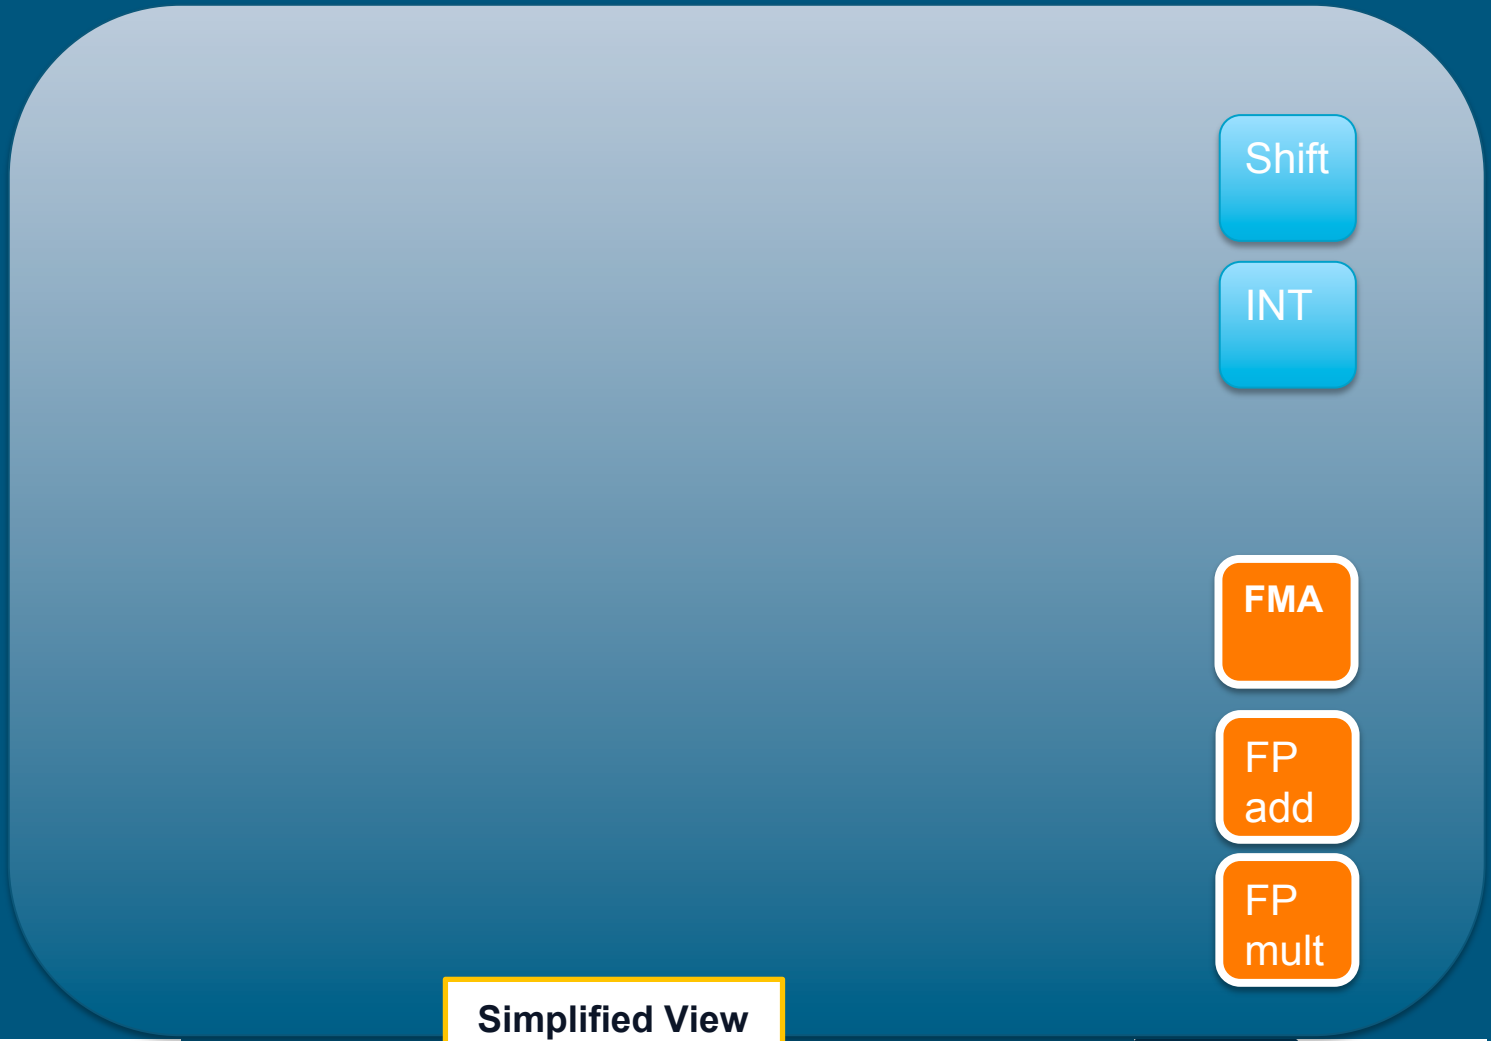

Simplified View

CPU (1 core)

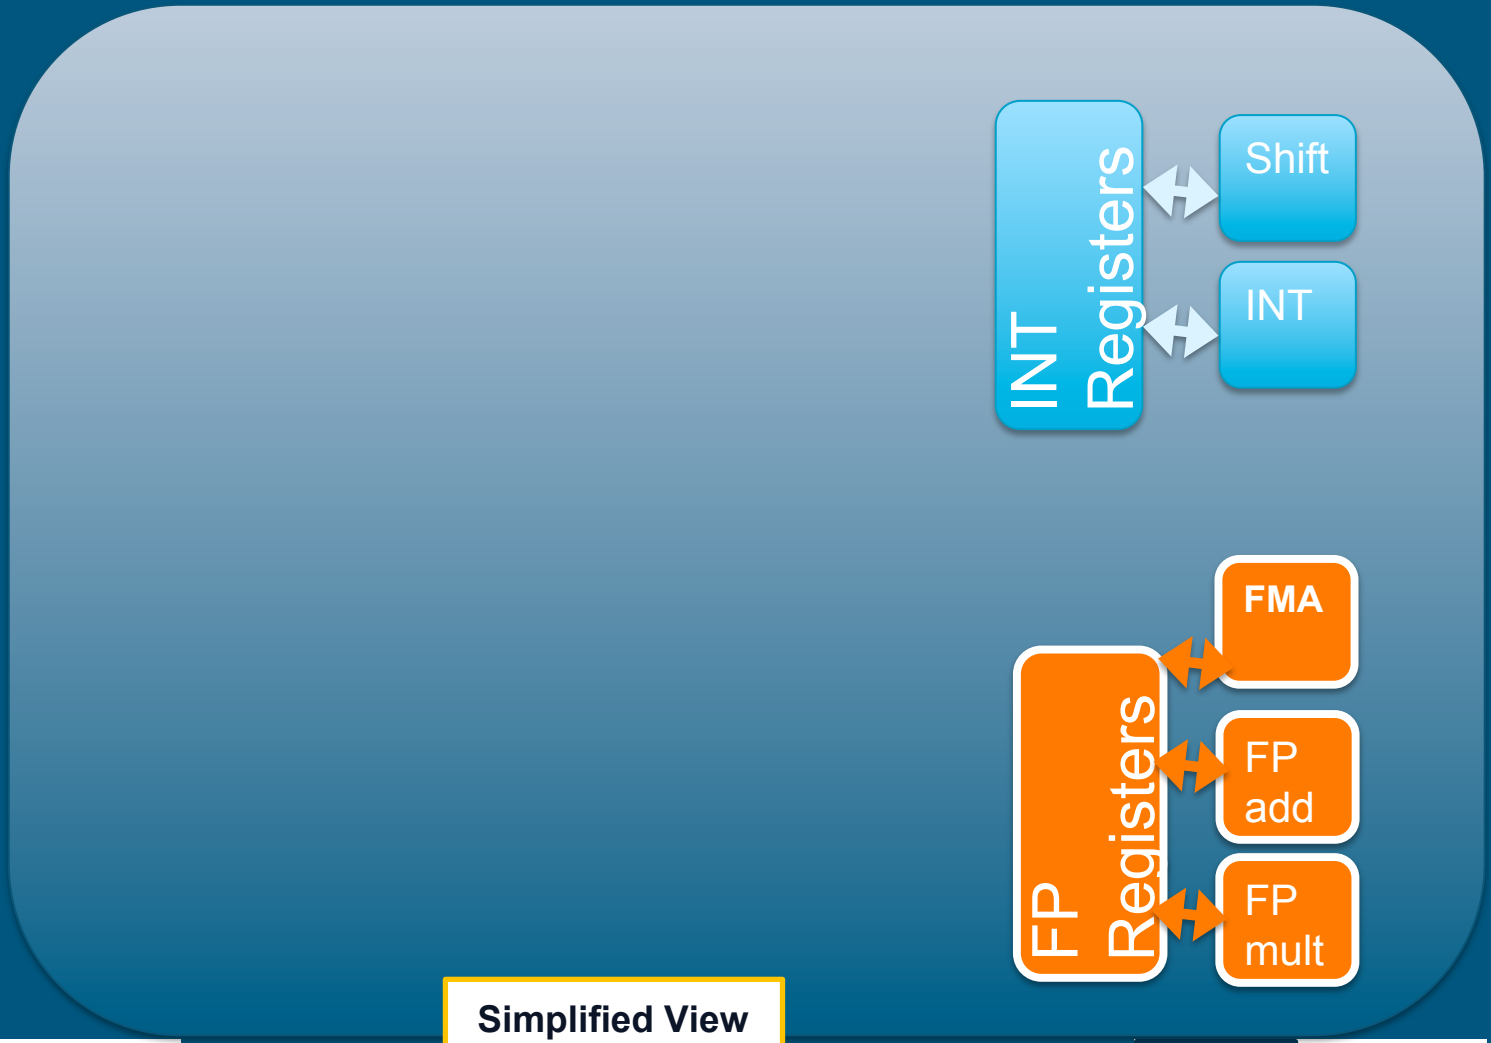

Simplified View

CPU (1 core)

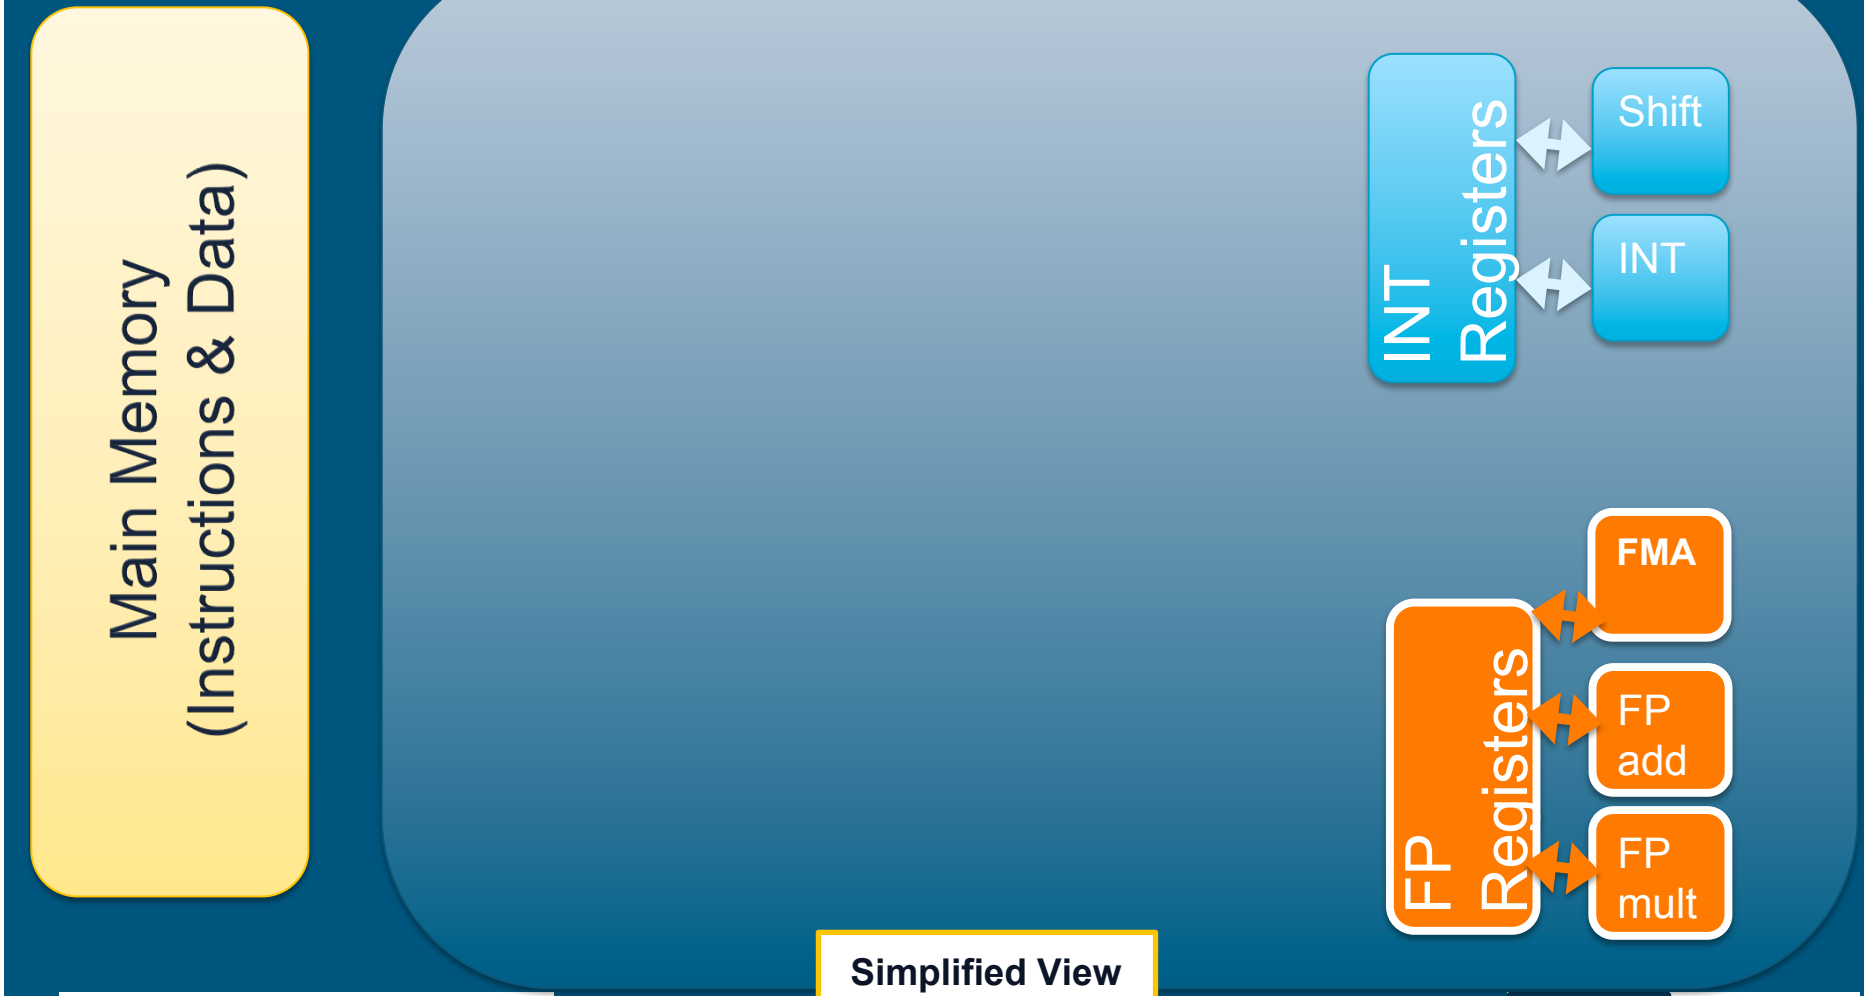

Simplified View

CPU (1 core)

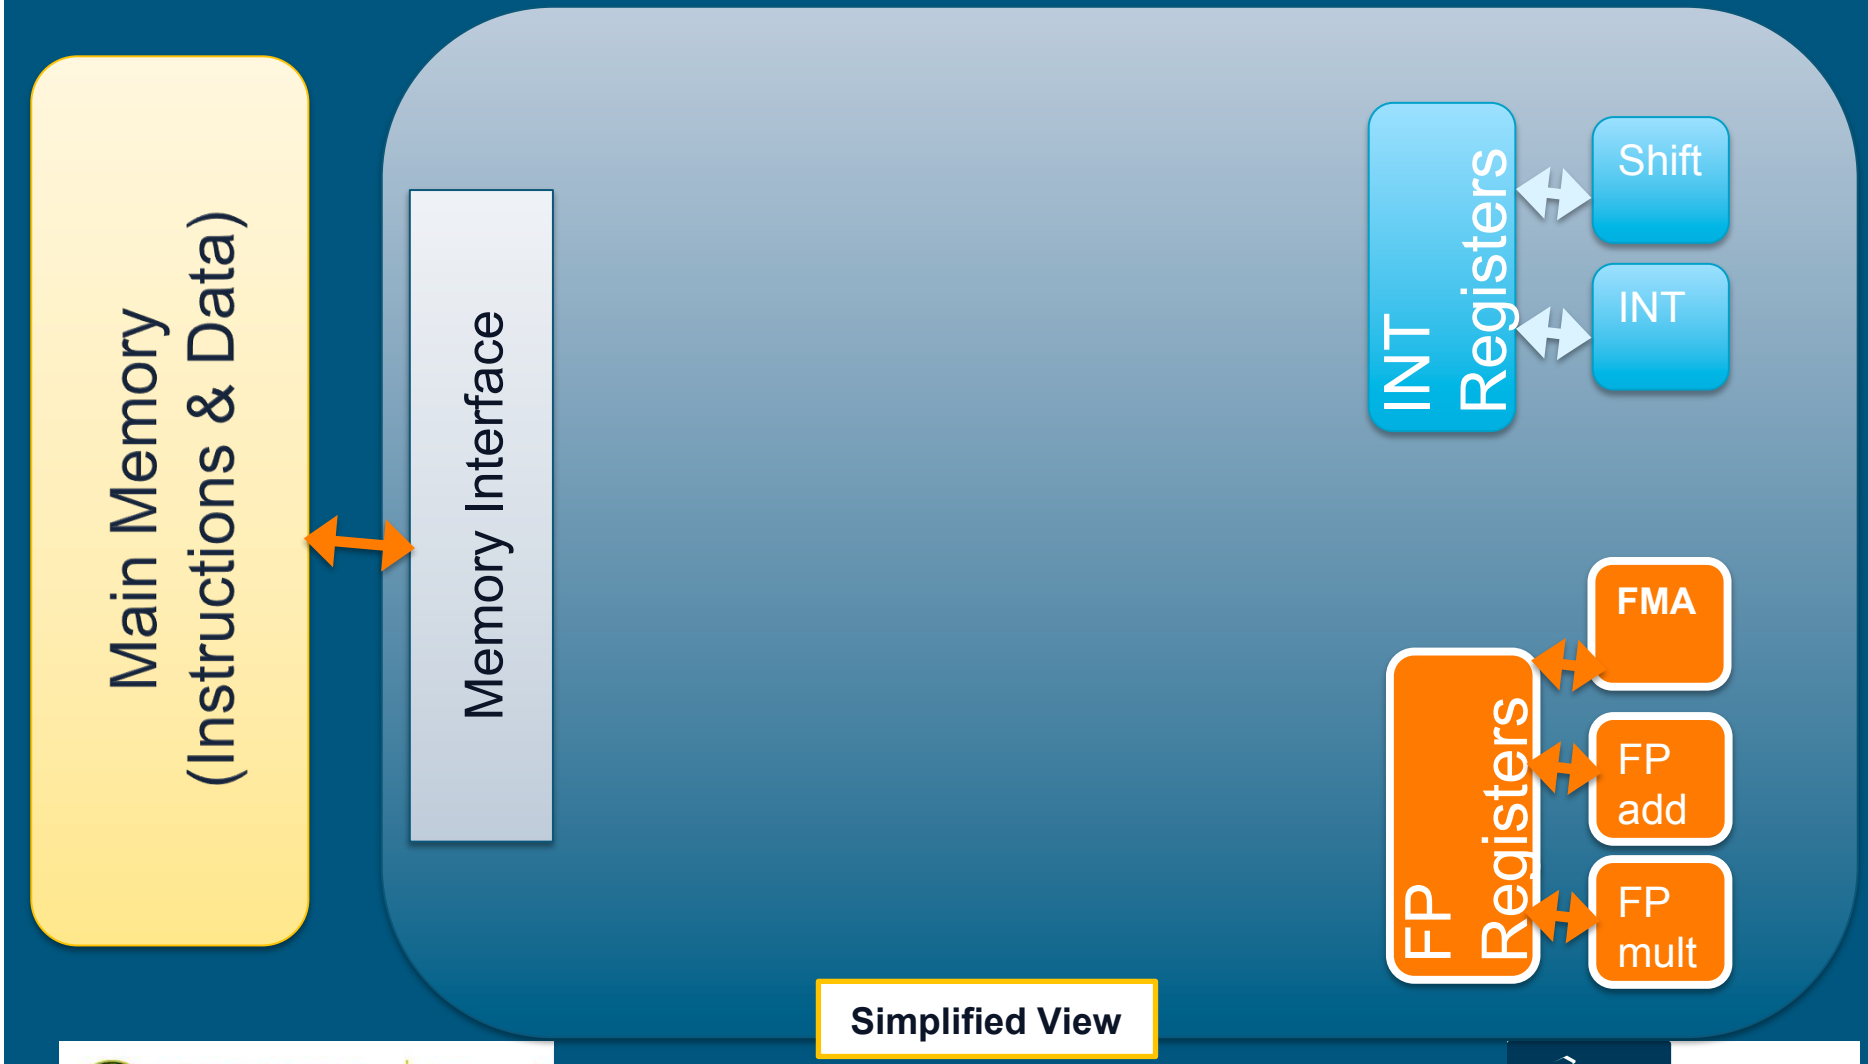

Simplified View

CPU (1 core)

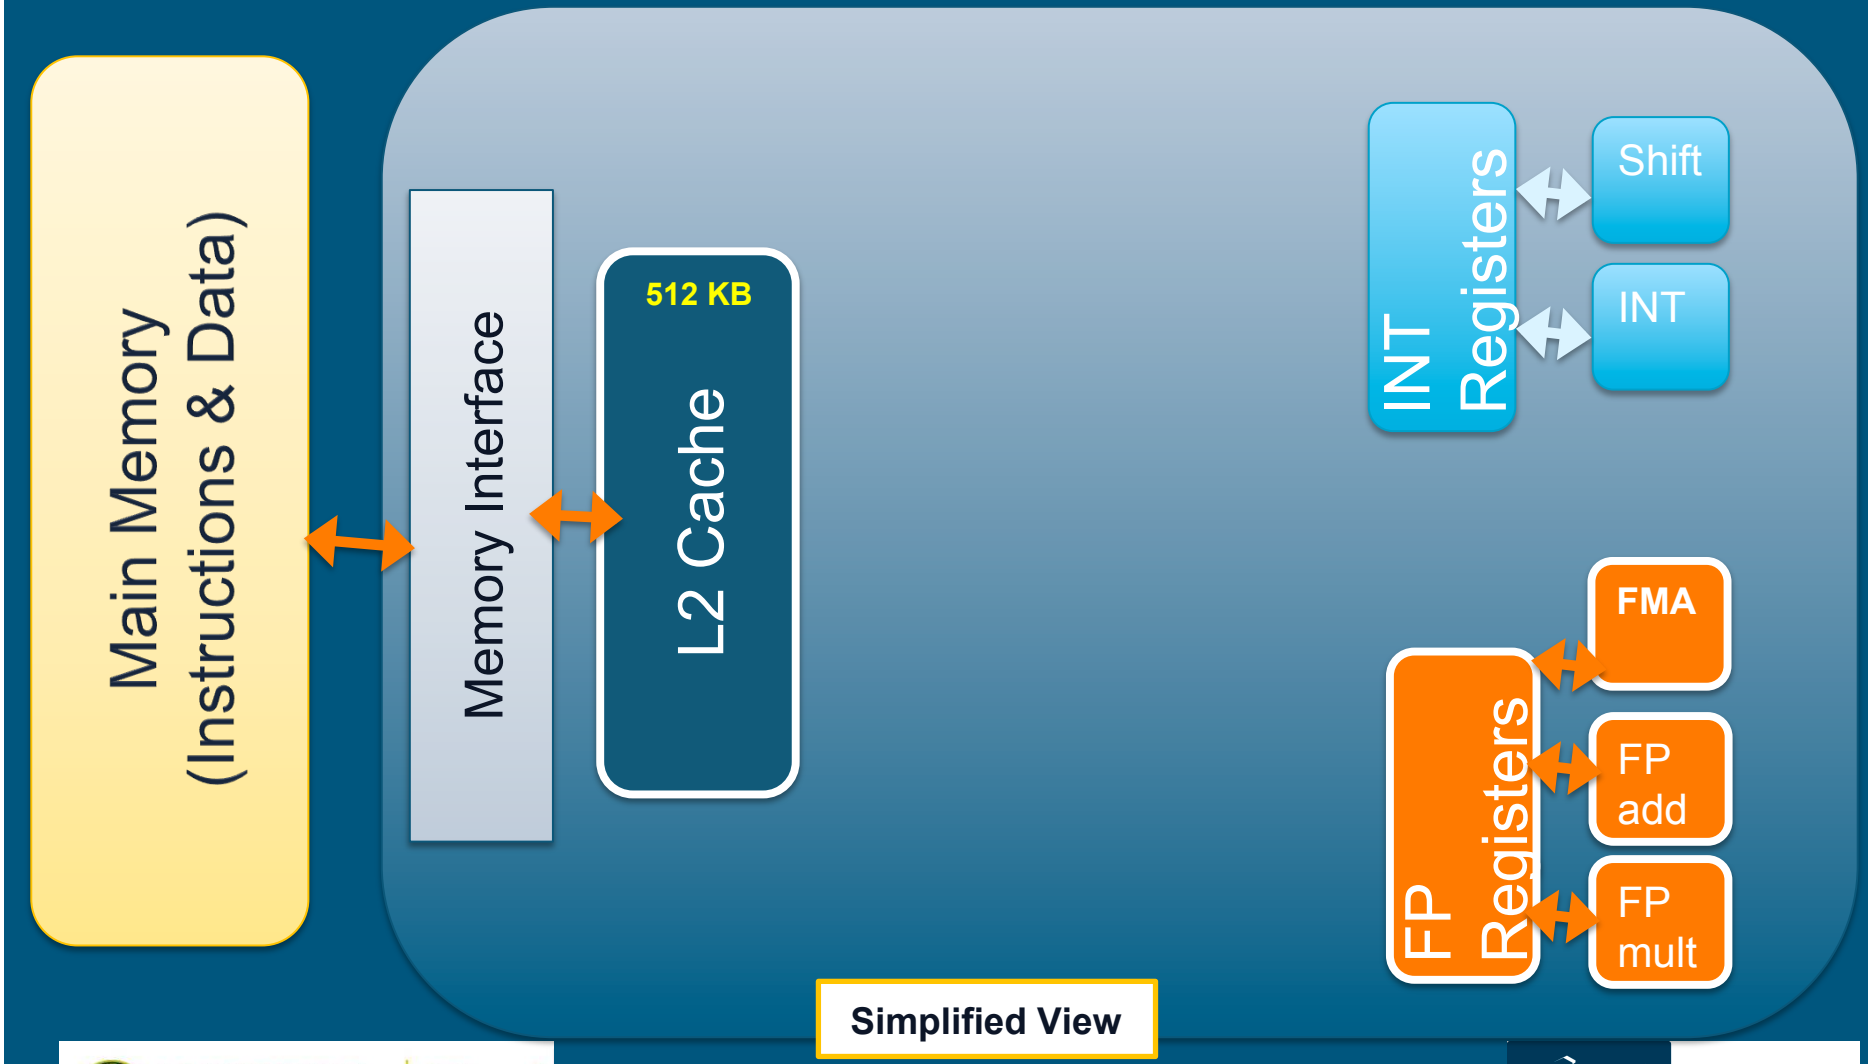

Simplified View

CPU (1 core)

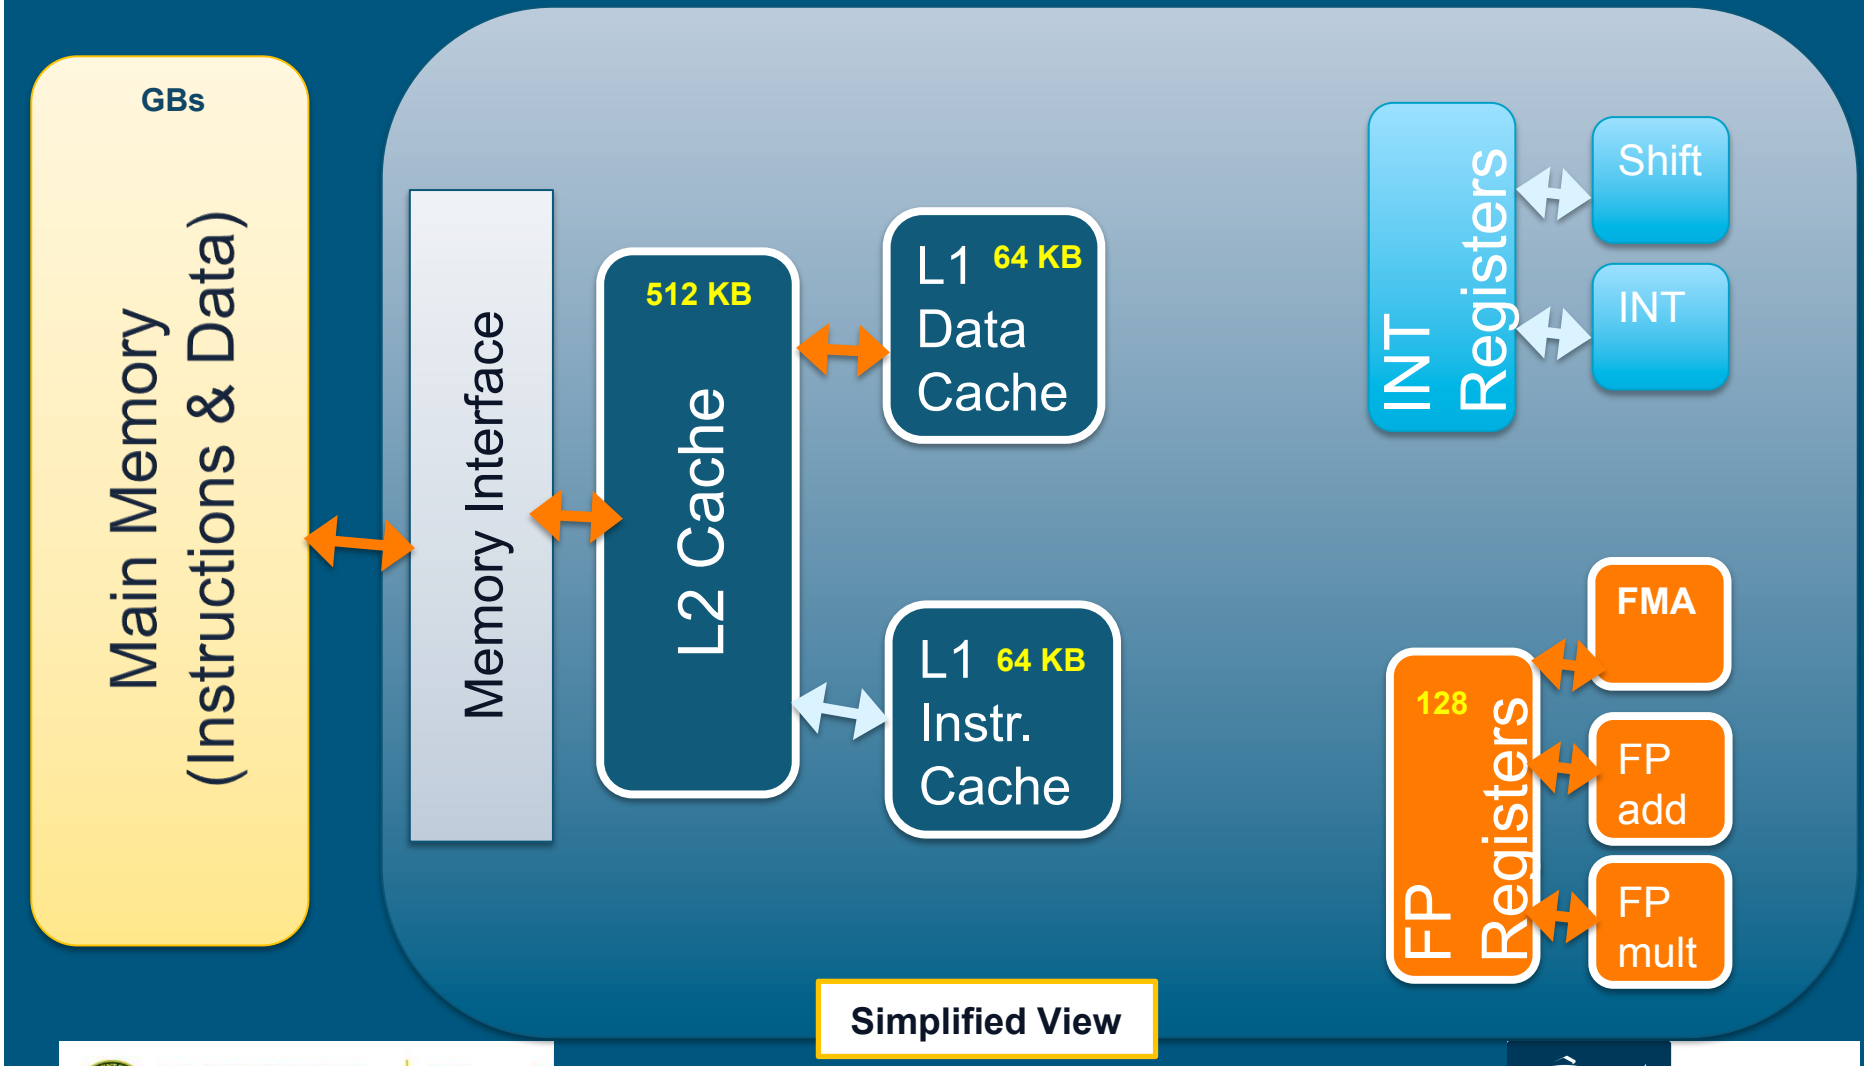
GBs
Main Memory
(Instructions & Data)
Memory Interface
512 KB
L2 Cache
L1 64 KB
Data
Cache
L1 64 KB
Instr.
Cache
INT
Registers
Shift
INT
128
FP
Registers
FMA
FP
add
FP
mult
Simplified View

CPU (1 core)

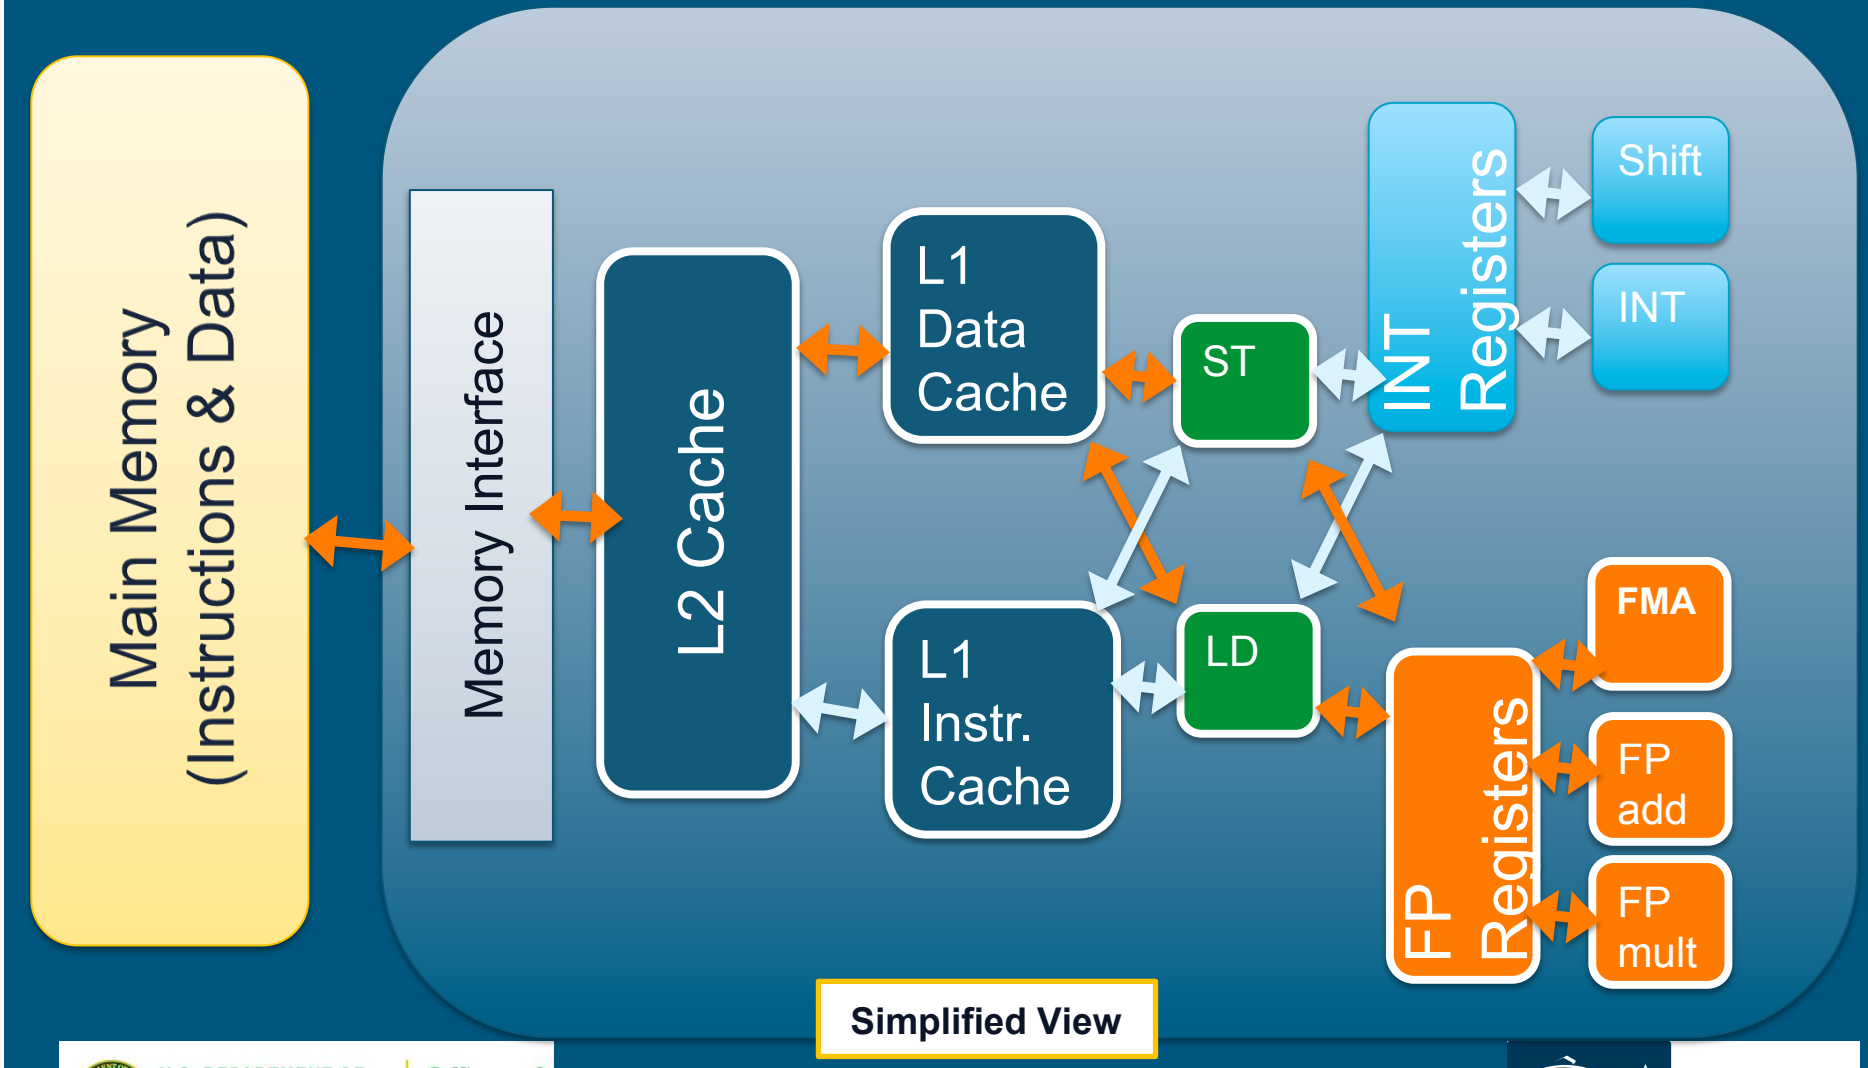
Main Memory
(Instructions & Data)
Memory Interface
L2 Cache
L1
Data
Cache
L1
Instr.
Cache
ST
LD
INT
Registers
Shift
INT
FP
Registers
FMA
FP
add
FP
mult
Simplified View

Additional Functional Units

- **There's a lot more on the CPU than shown previously, e.g.**
  - L3 cache (~10 MB)
  - SQRT/Divide/Trig FP unit
  - "TLB" to cache memory addresses
  - Instruction decode
  - …

On-Chip Parallelism

- **Chip designers have added lots of complexity to increase performance**
- **Instruction Parallelism**
  - Pipelined functional units (e.g. FPU)
  - Superscalar processors
- **Data Parallelism**
  - SIMD
- **Out of order & speculative execution**

Example: Pipeline Stages in an FPU

- Separate mantissa / exponent
- Multiply mantissas
- Add exponents
- Normalize result
- Insert sign

Pipelining

Dave Patterson's Laundry example: 4 people doing laundry

wash (30 min) + dry (40 min) + fold (20 min) = 90 min

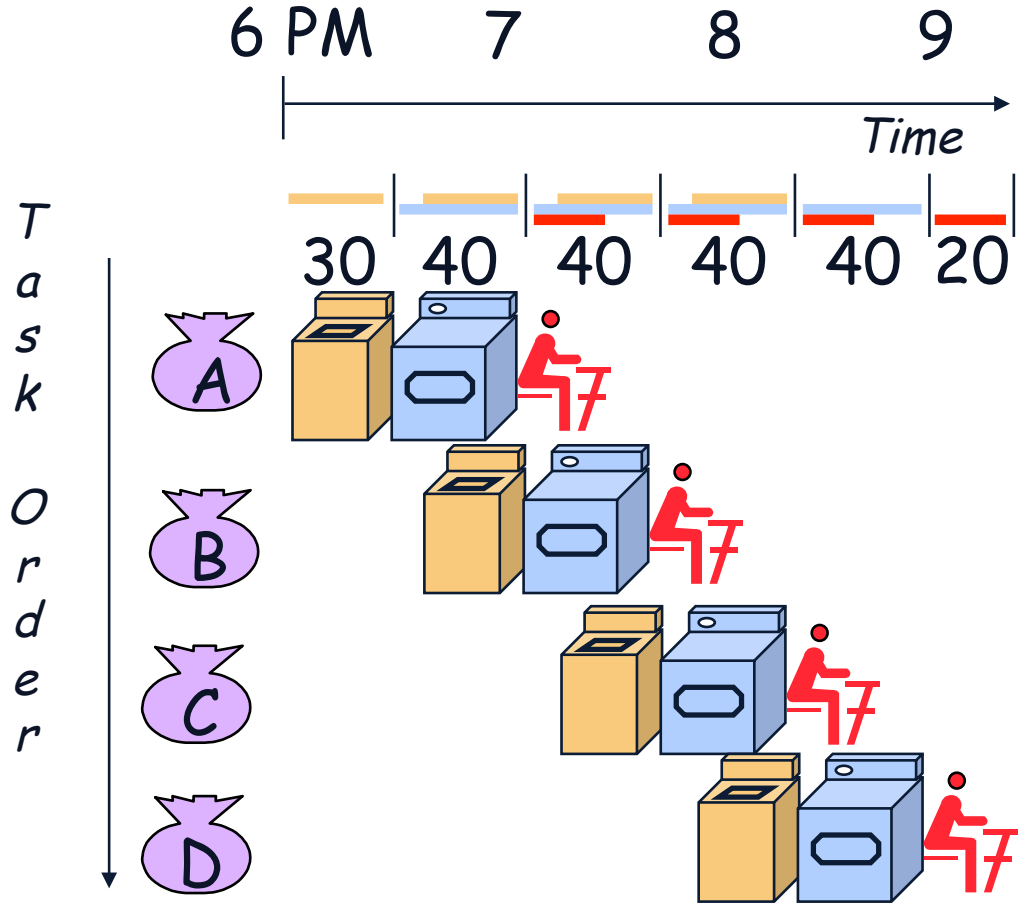

- In this example:
  - Sequential execution takes 4 * 90min = 6 hours
  - Pipelined execution takes 30+4*40+20 = 3.5 hours
- Bandwidth = loads/hour
- BW = 4/6 l/h w/o pipelining
- BW = 4/3.5 l/h w pipelining
- BW <= 1.5 l/h w pipelining, more total loads
- Pipelining helps bandwidth but not latency (90 min)
- Bandwidth limited by slowest pipeline stage
- Potential speedup = Number pipe stages

Superscalar Processors

- **Superscalar processors can execute more than one instruction per clock cycle**
- **Example**
  - 2 FMAs
  - 2 integer ops
  - Multiple LOAD & STORE

SIMD

- Special registers can hold multiple words of data
- A single instruction (e.g. floating point multiply) is applied to all the data at once
- "SSE[2-4]" : Streaming SIMD Extension instruction set for x86
- aka "Vectorization"

Latencies

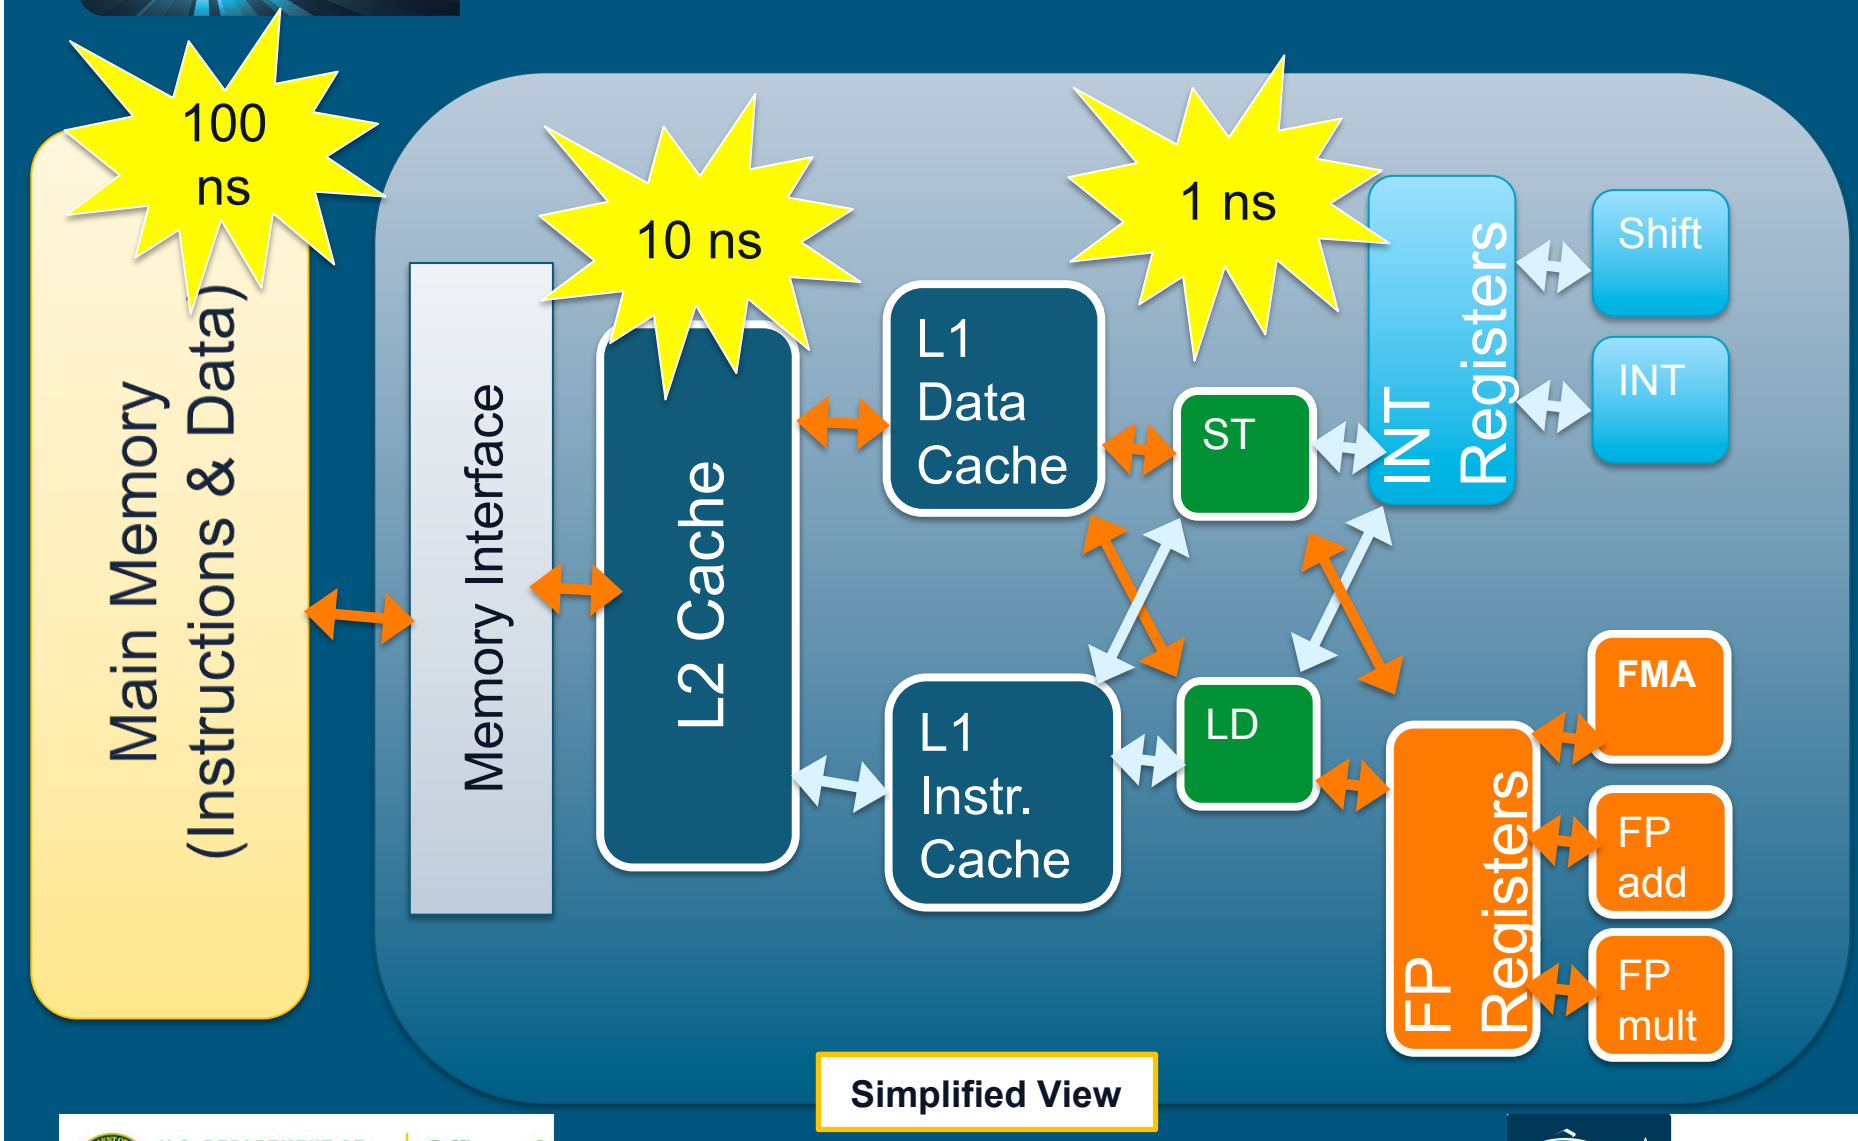

Simplified View

Memory Bottleneck

- **Modern HPC CPUs can achieve ~5 Gflop/sec per compute core**
  - 2 8-byte data words per operation
  - 80 GB/sec of data needed to keep CPU busy
- **Memory interfaces provide a few GB/ sec per core from main memory**
- **Memory latency – the startup time to begin fetching data from memory – is even worse**

Memory Bottleneck

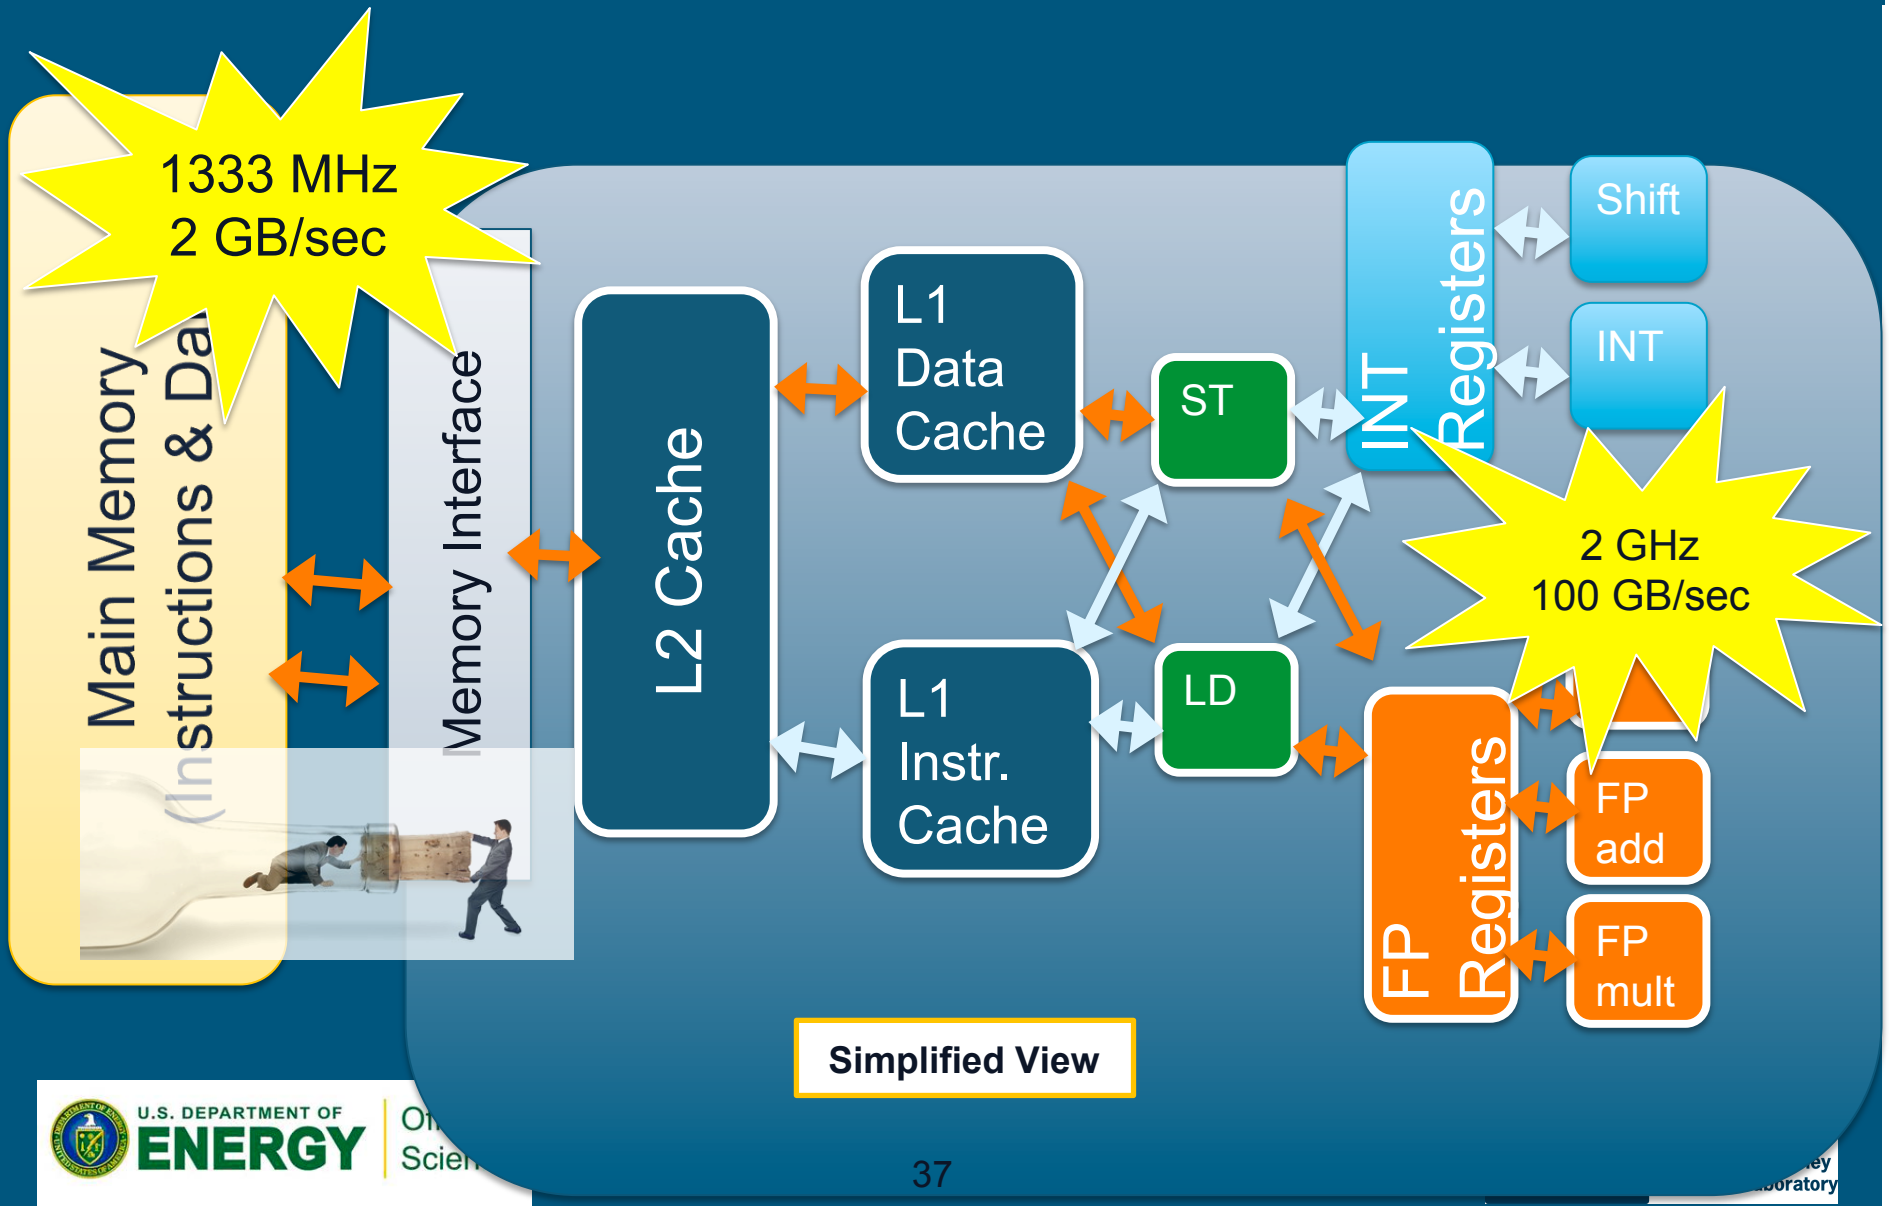

Multicore Processors

- There are no more single-core CPUs (processors) as just described
- All CPUs (processors) now consist of multiple compute "cores" on a single "chip" or "die" with possibly multiple chips per "socket" (the unit that plugs into the motherboard)
- May not be "symmetric" wrt cores and functional units
- Increased complexity
- The trend is for ever-more cores per die

Moore's Law

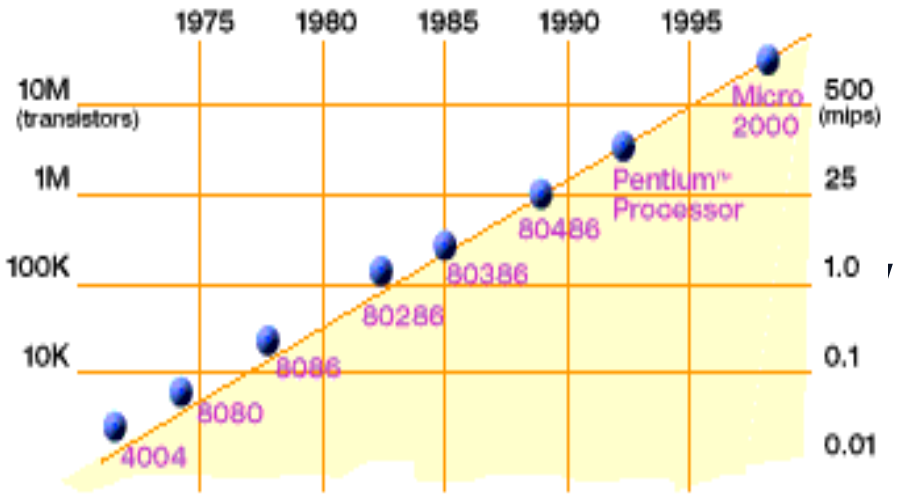

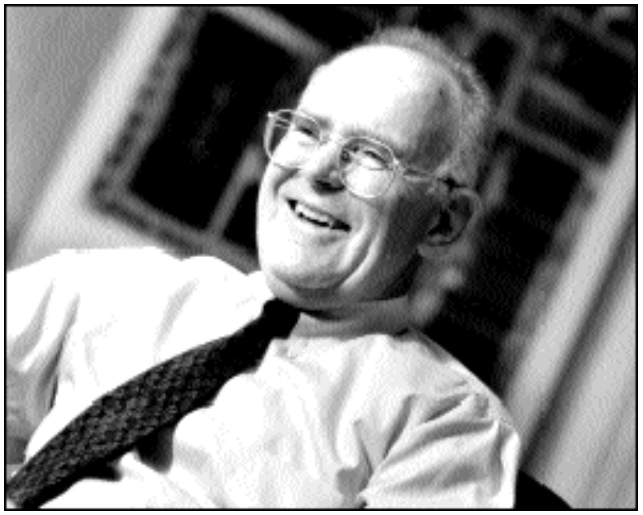

2X transistors/Chip Every 1.5 years

Called "Moore's Law"

Microprocessors have become smaller, denser, and more powerful.

Gordon Moore (co-founder of Intel) predicted in 1965 that the transistor density of semiconductor chips would double roughly every 18 months.

Slide source: Jack Dongarra

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB Lawrence Berkeley National Laboratory

Power Density Limits Serial Performance

- Concurrent systems are more power efficient
  - Dynamic power is proportional to V^2fC
  - Increasing frequency (f) also increases supply voltage (V) → cubic effect
  - Increasing cores increases capacitance (C) but only linearly
  - Save power by lowering clock speed

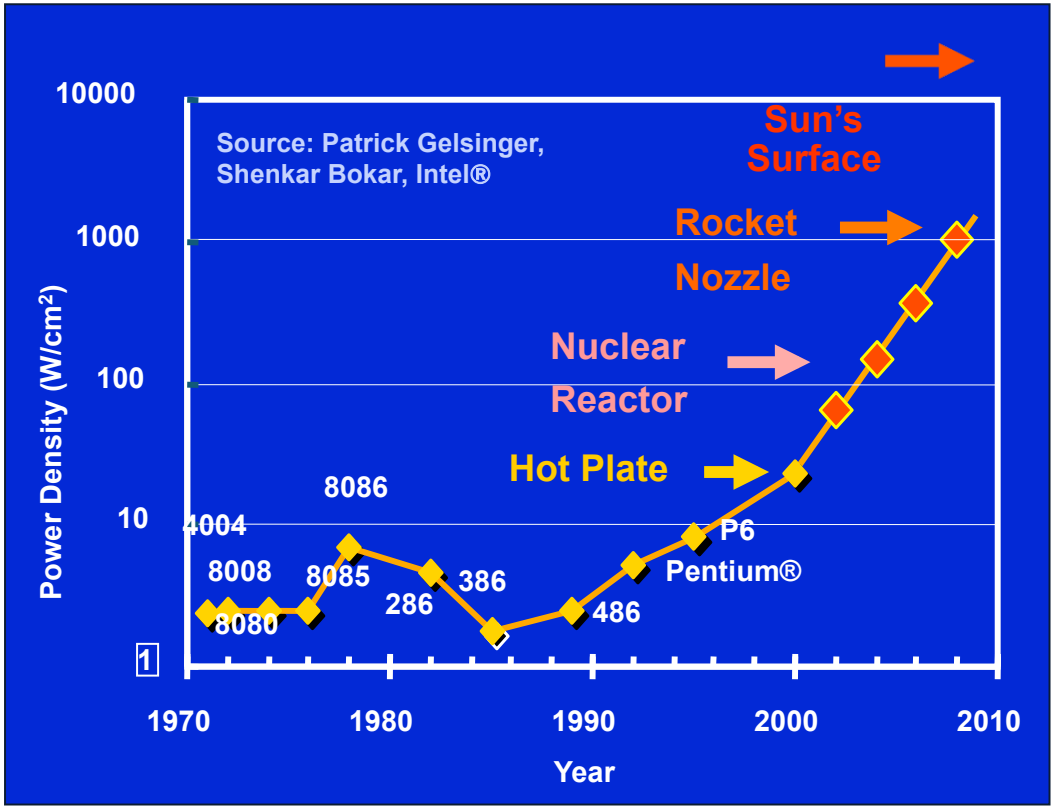

- High performance serial processors waste power
  - Speculation, dynamic dependence checking, etc. burn power
  - Implicit parallelism discovery
- More transistors, but not faster serial processors

Revolution in Processors

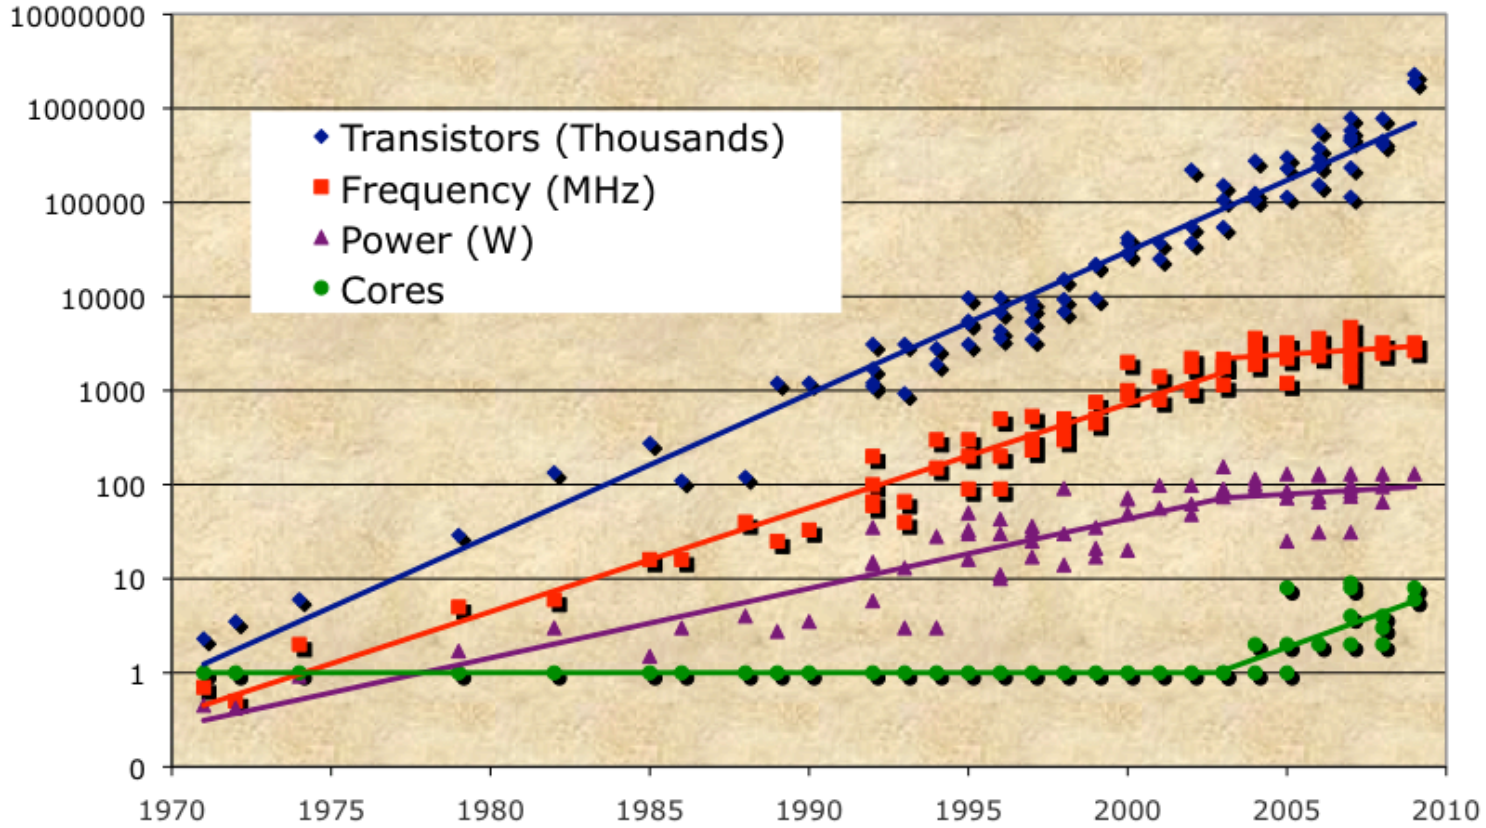

- Chip density is continuing increase ~2x every 2 years
- Clock speed is not
- Number of processor cores may double instead
- Power is under control, no longer growing

Moore's Law reinterpreted

- Number of cores per chip will double every two years
- Clock speed will not increase (possibly decrease)
- Need to deal with systems with millions of concurrent threads
- Need to deal with inter-chip parallelism as well as intra-chip parallelism

Example HPC CPU (2 core)

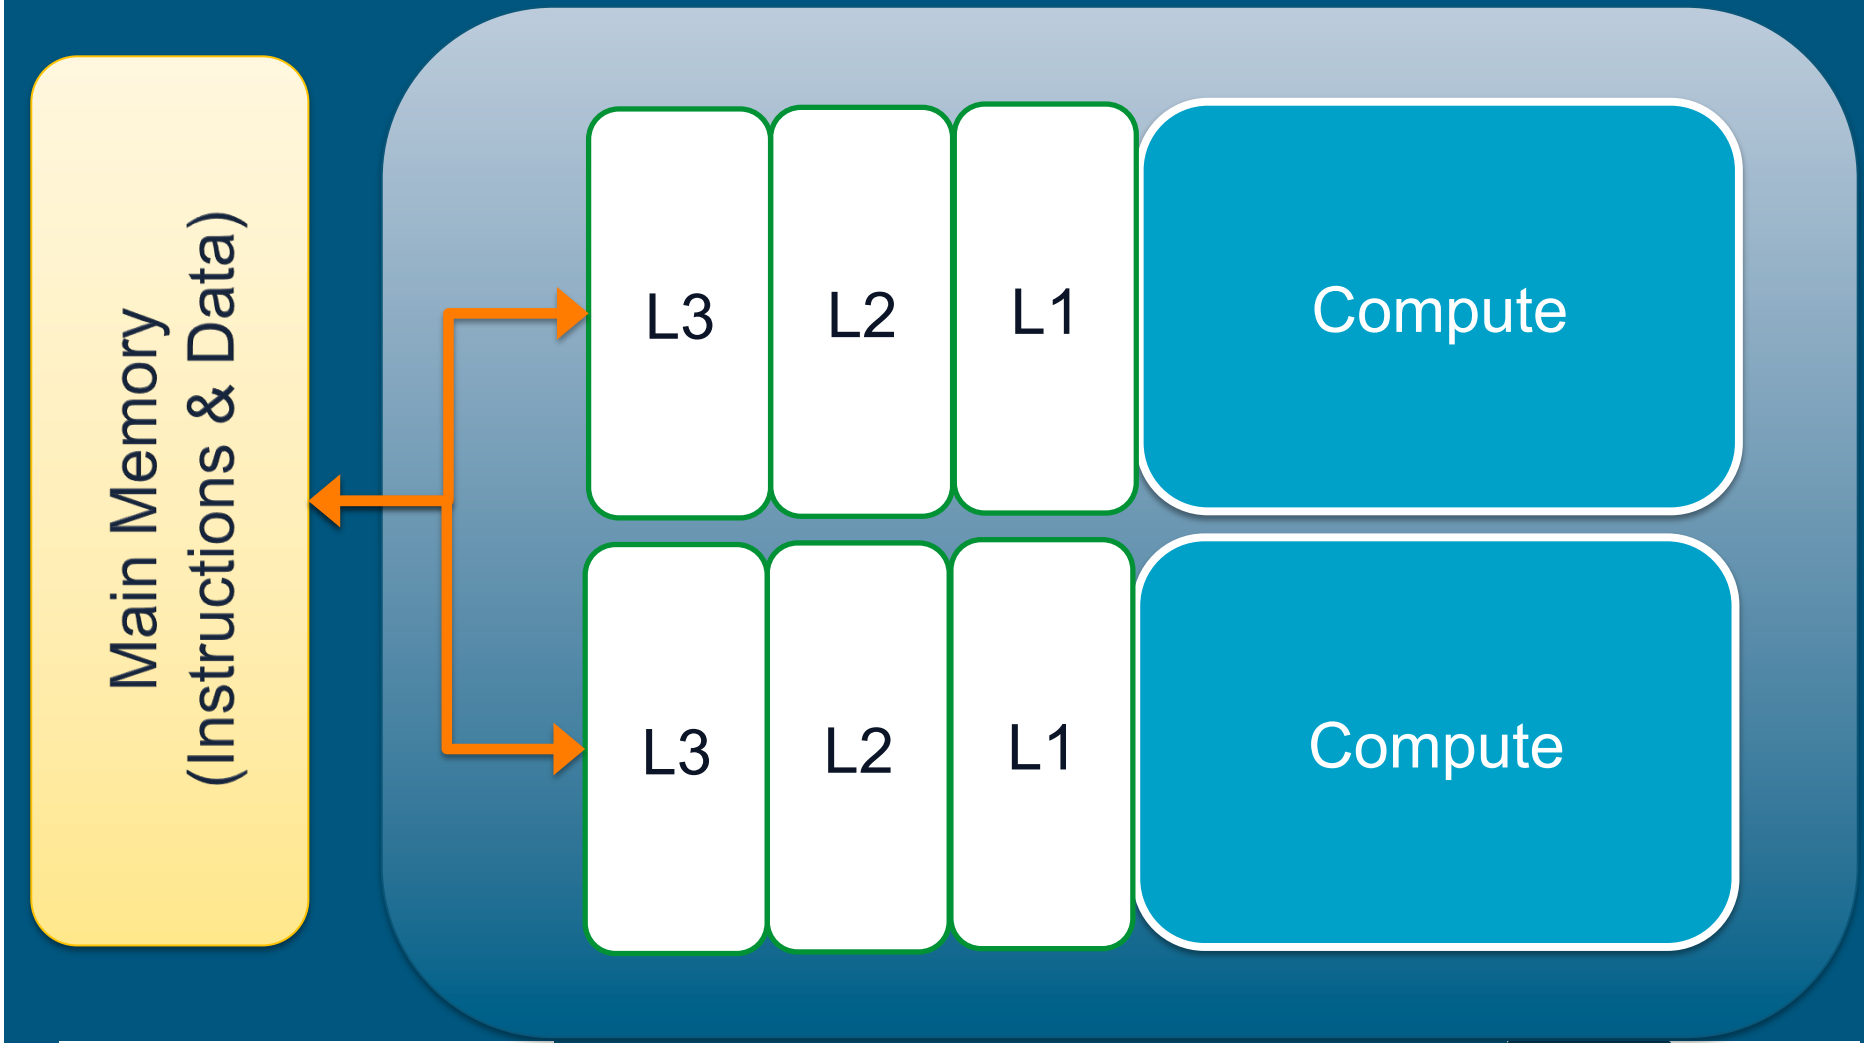

Example HPC CPU (4 core)

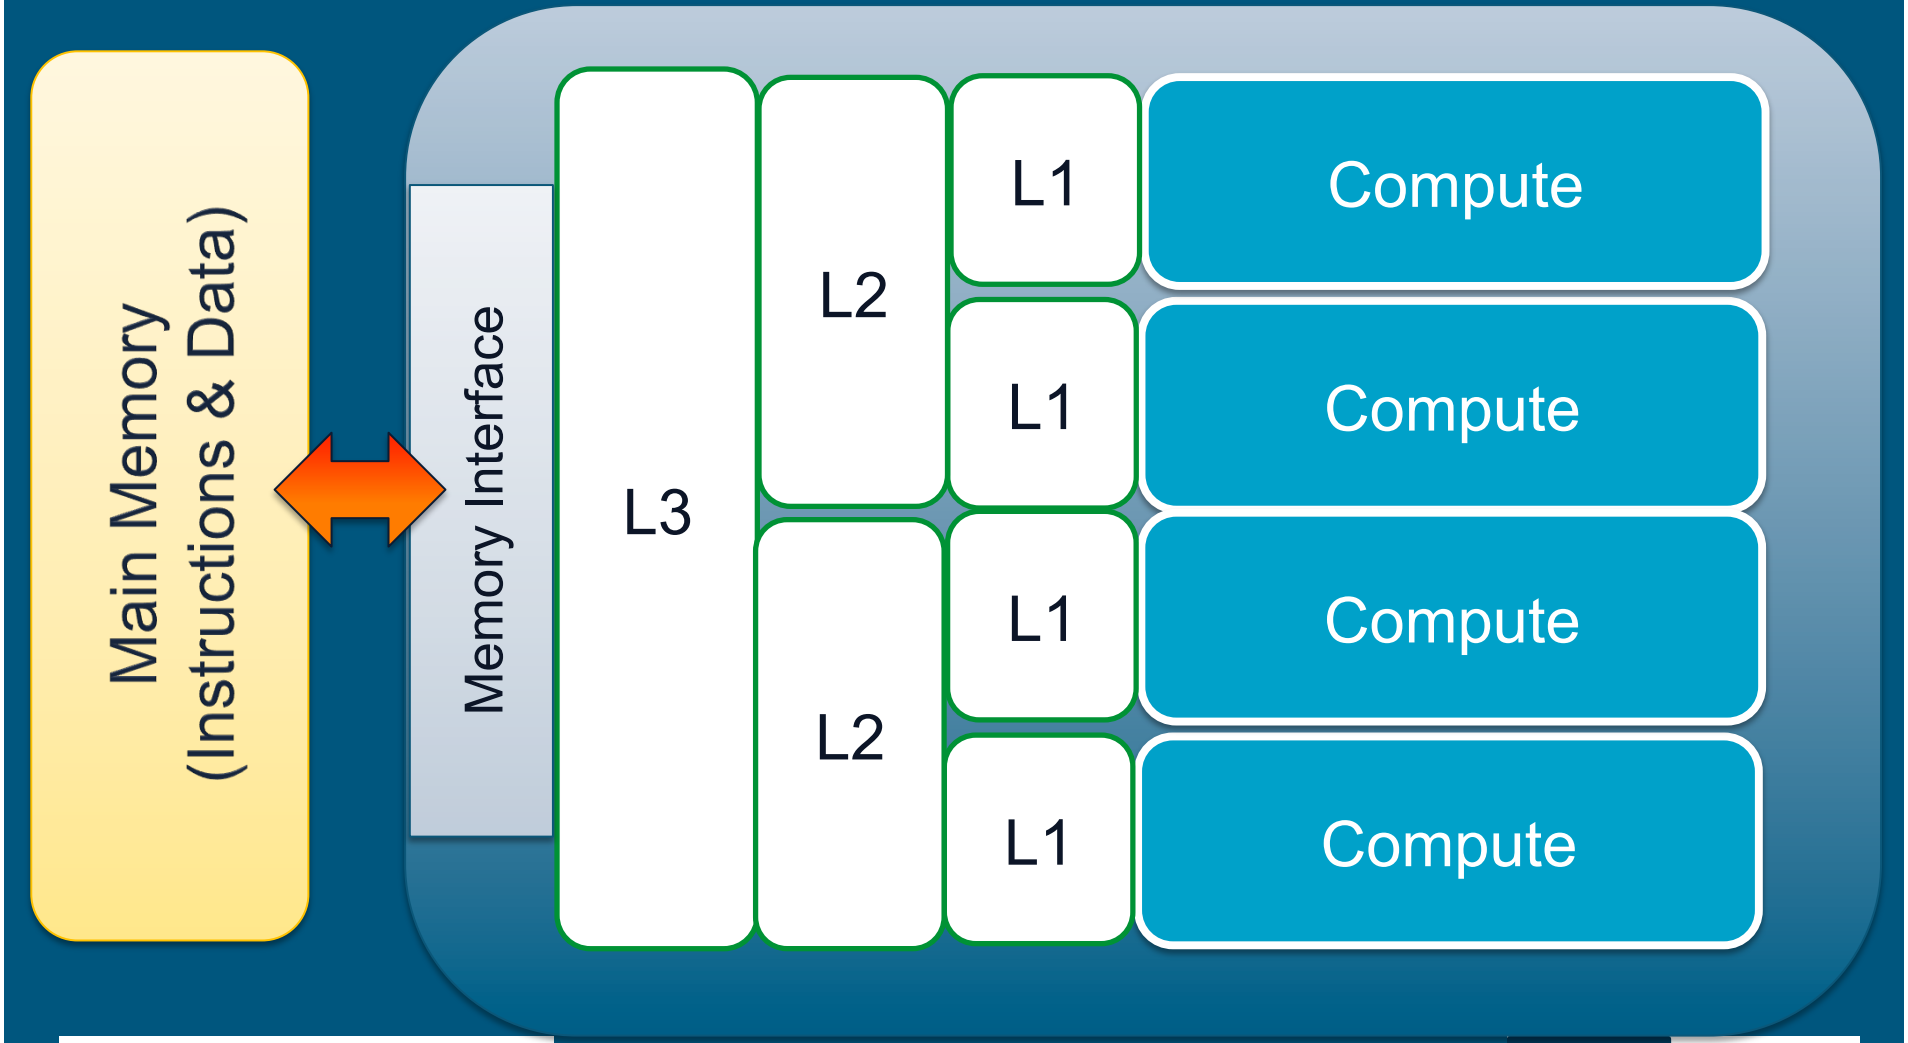

Hypothetical Socket (8 core)

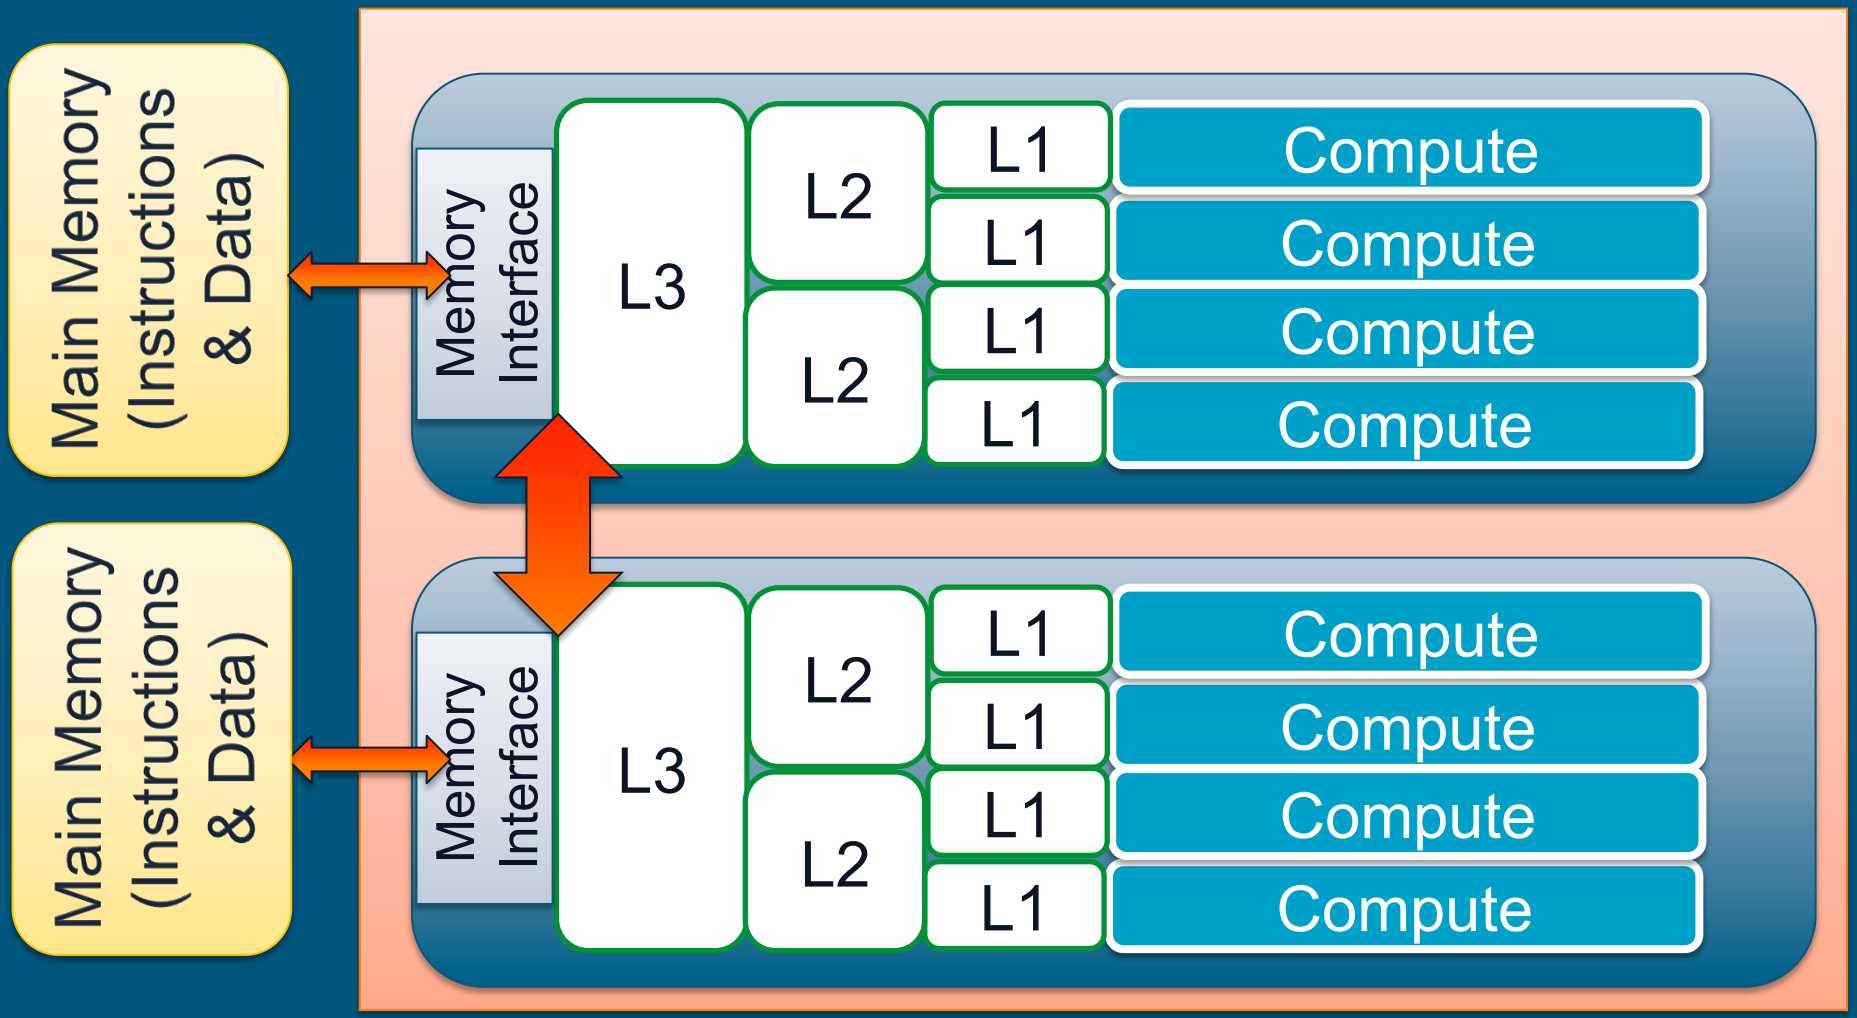

5 Major Parts of an HPC System

1. CPUs
2. Memory (volatile)
3. **Nodes**
4. Inter-node network
5. Non-volatile storage (disks, tape)

HPC Node

- **A "node" is a (physical) collection of CPUs, memory, and interfaces to other nodes and devices.**
  - Single memory address space
  - Shared memory pool
  - Memory access "on-node" is significantly faster than "off-node" memory access
  - Often called an "SMP node" for "Shared Memory Processing"
    - Not necessarily "symmetric" memory access as in "Symmetric Multi-Processing"

Example SMP Node

SMP Node – Each core has equal access to memory and cache

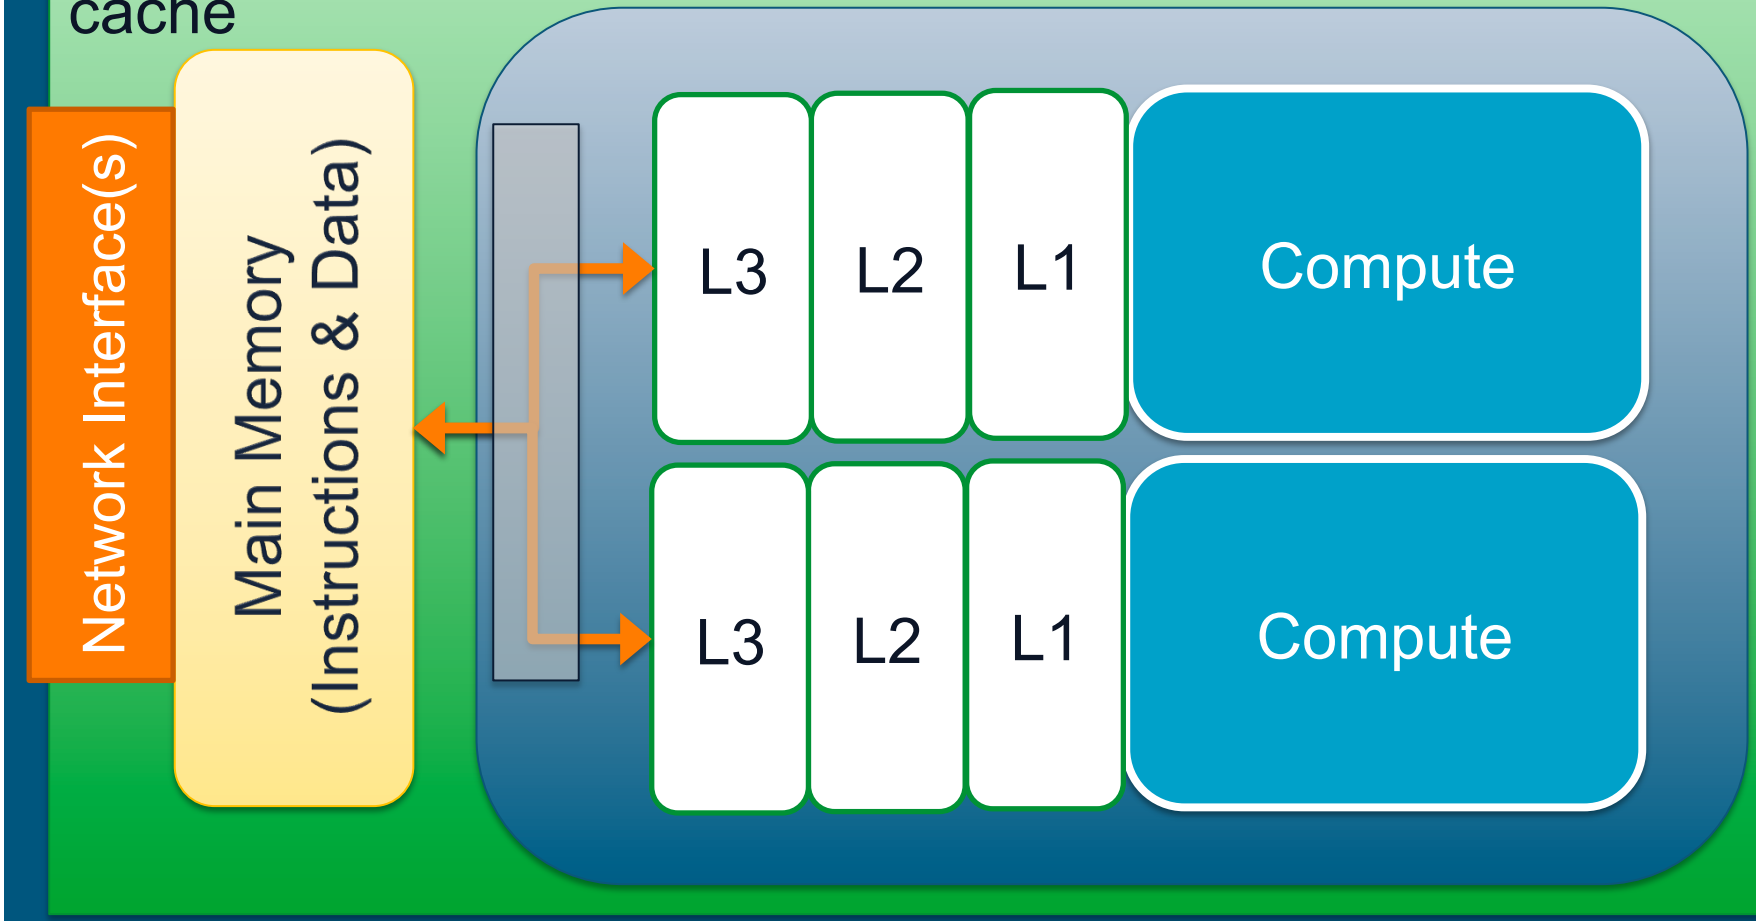

Example NUMA Node

NUMA Node – Non-Uniform Memory Access

Single address space

Example NUMA Node

NUMA Node – Non-Uniform Memory Access

Single address space

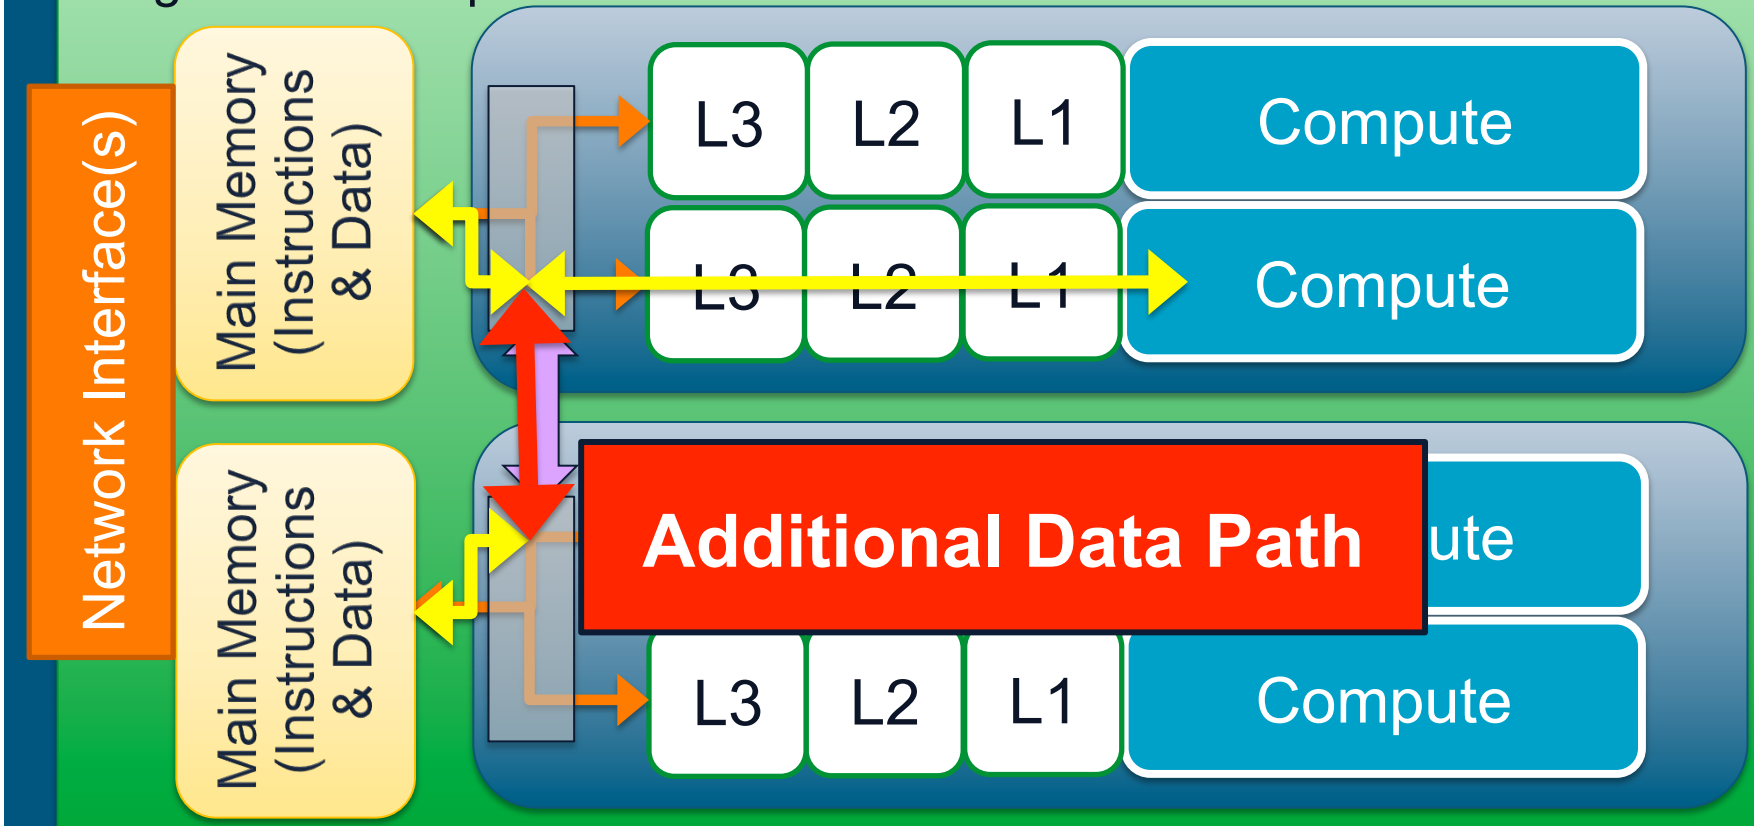

NUMA node (4 dies, 2 sockets)

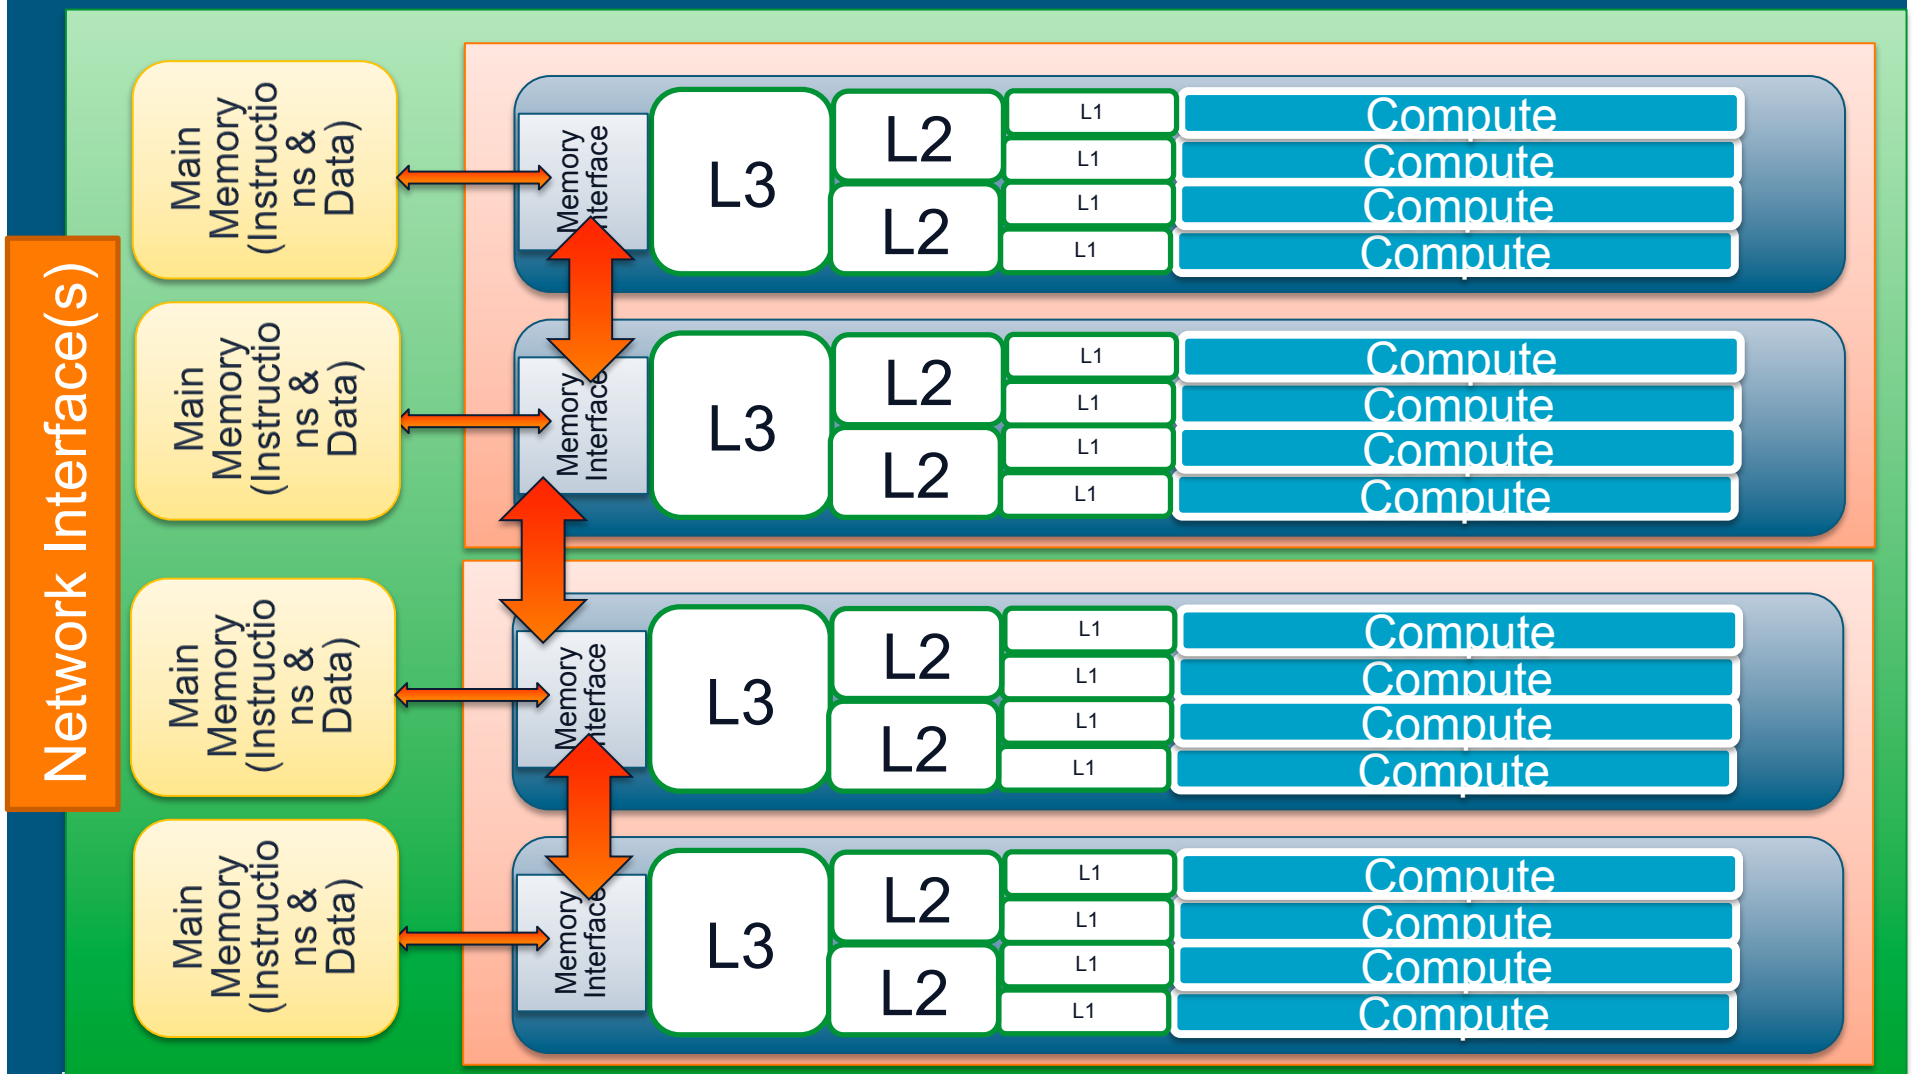

5 Major Parts of an HPC System

1. CPUs
2. Memory (volatile)
3. Nodes
4. **Inter-node network**
5. Non-volatile storage (disks, tape)

Distributed Memory Systems

- **Most HPC systems are "distributed memory"**
  - Many nodes, each with its own local memory and distinct memory space
  - Nodes communicate over a specialized high-speed, low-latency network
  - SPMD (Single Program Multiple Data) is the most common model
    - Multiple copies of a single program (tasks) execute on different processors, but compute with different data
    - Explicit programming methods (MPI) are used to move data among different tasks

Interconnect Characteristics

- **Latency**
  - The startup-time needed to initiate a data transfer between nodes (time to send a zero-byte message)
  - Latencies between different nodes may be different
  - Typically ~ a few μsec
- **Bandwidth**
  - Data transfer rate between nodes
  - May be quoted as uni- or bi-directional
  - Typically ~ a few GB/sec in/out of a node
- **Bisection Bandwidth**
  - If a network is divided into two equal parts, the bandwidth between them is the bisection bandwidth

Examples

Network	Bandwidth (GB/s)	Latency (µs)
Arista 10GbE(stated)	1.2	4.0
BLADE 10GbE(measured)	1.0	4.0
Cray SeaStar2+ (measured)	6.0	4.5
Cray Gemini (measured)	6.1	1.0
IBM (Infiniband) (measured)	1.2	4.5
SGI NumaLink 5(measured)	5.9	0.4
Infiniband (measured)	1.3	4.0
Infinipath (measured)	0.9	1.5
Myrinet 10-G (measured)	1.2	2.1

Main Networks Types

- **Switched**
  - Network switches connect and route network traffic over the interconnect
- **Mesh**
  - Each node sends and receives its own data, and also relays data from other nodes
  - Messages hop from one node to another until they reach their destination (must deal with routing around down nodes)

Fat Tree Switched Network

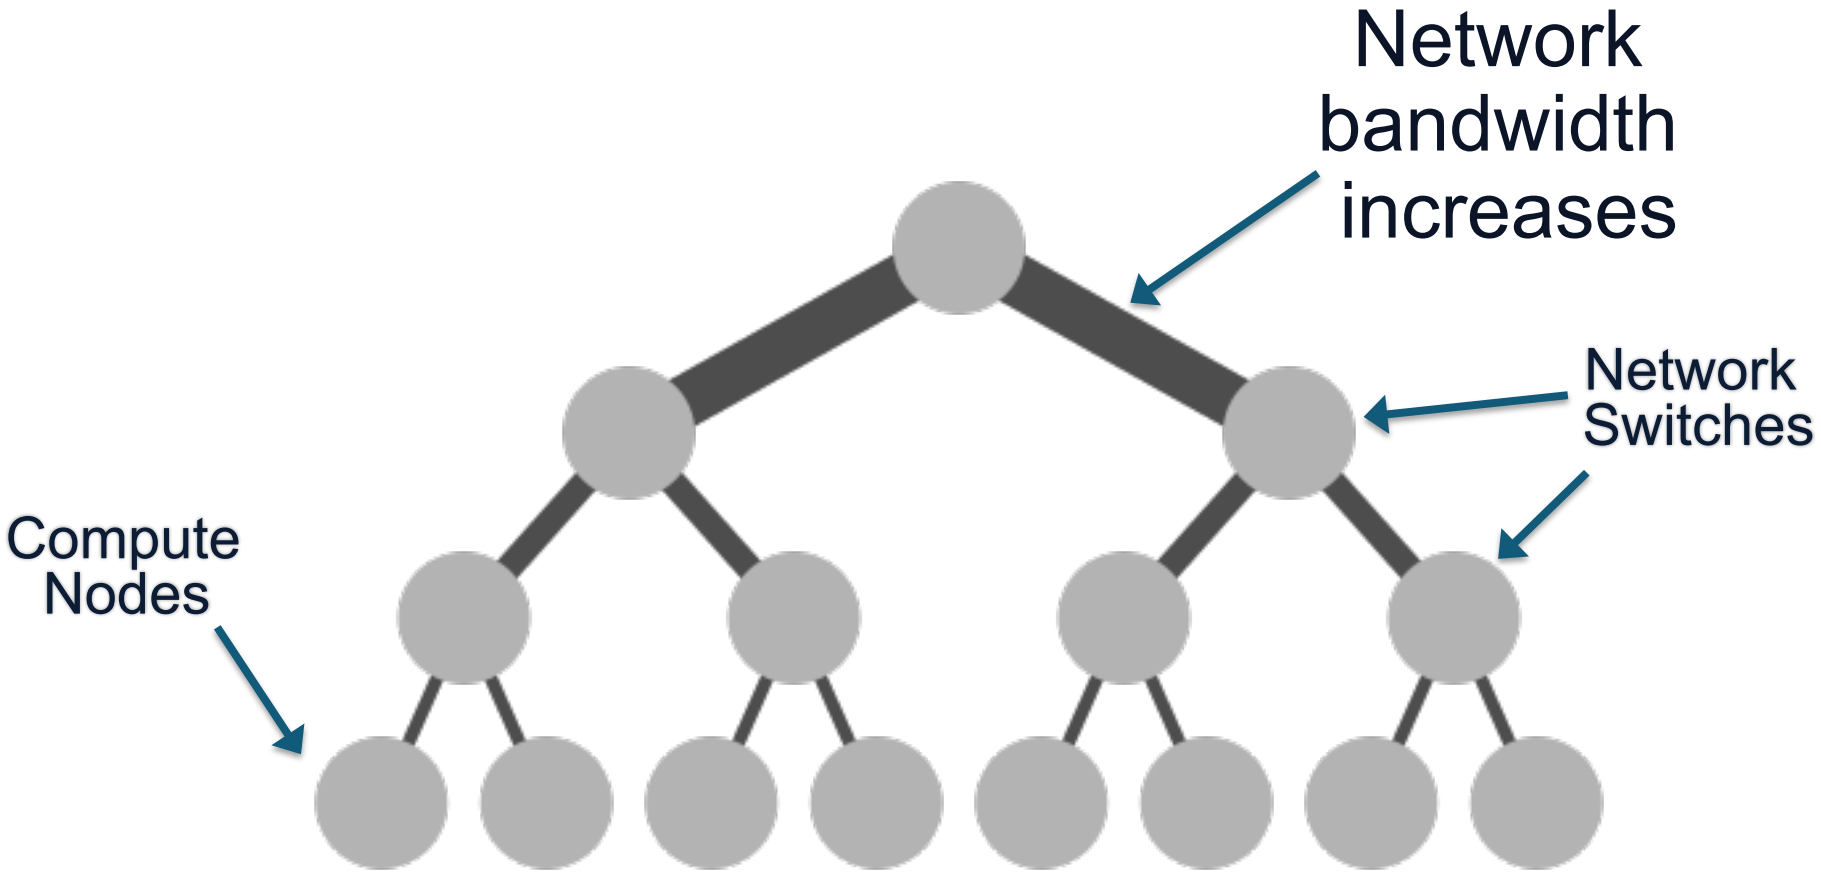

Implementation Can Be Complex

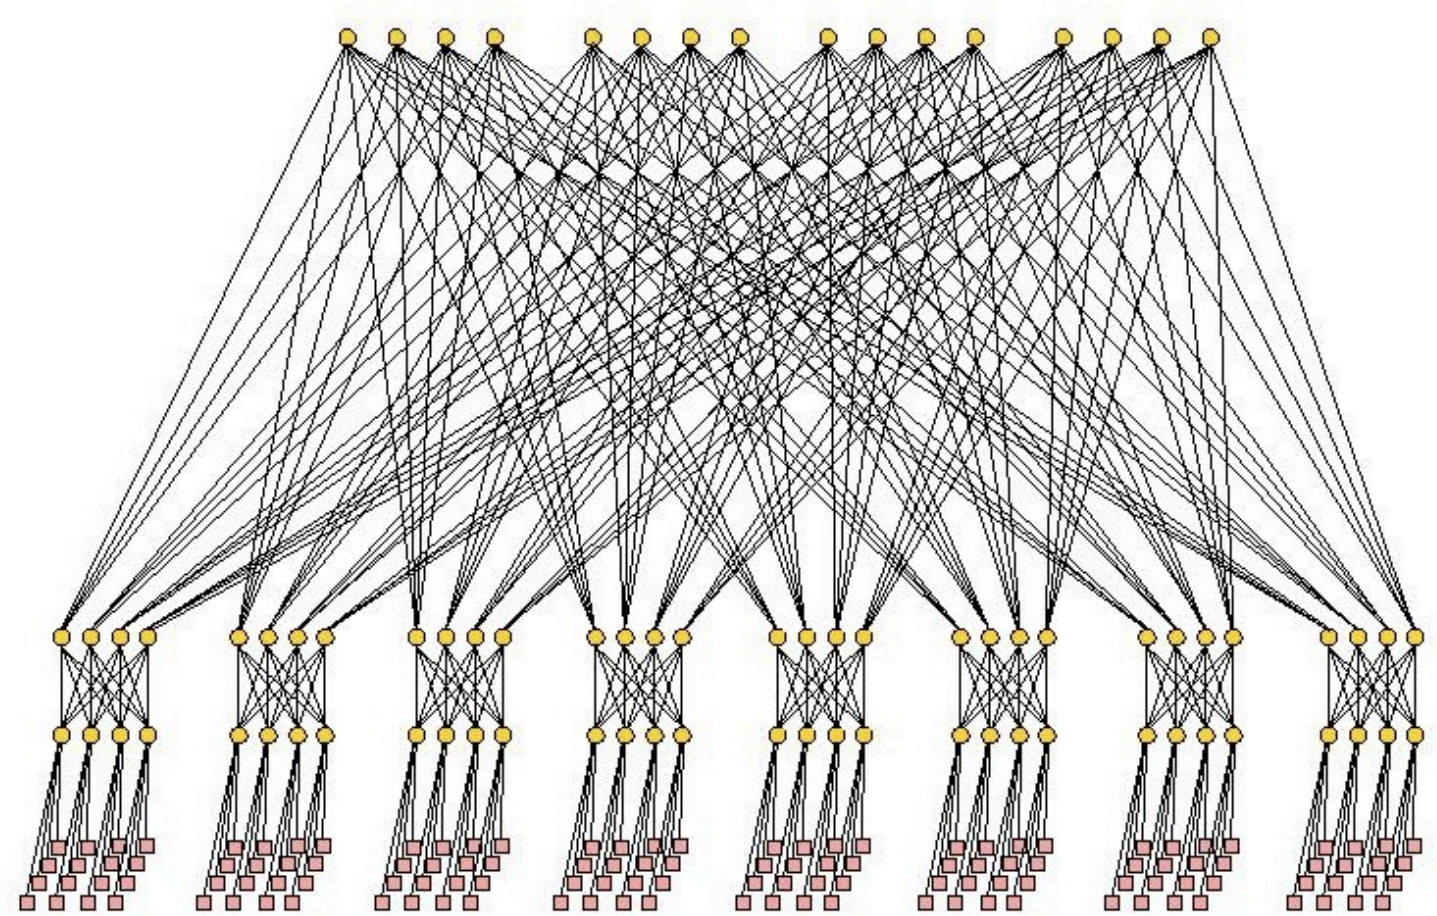

128-way fat tree

Mesh Networks

Mesh network topologies can be complex

Grids
Cubes
Hypercubes
Tori

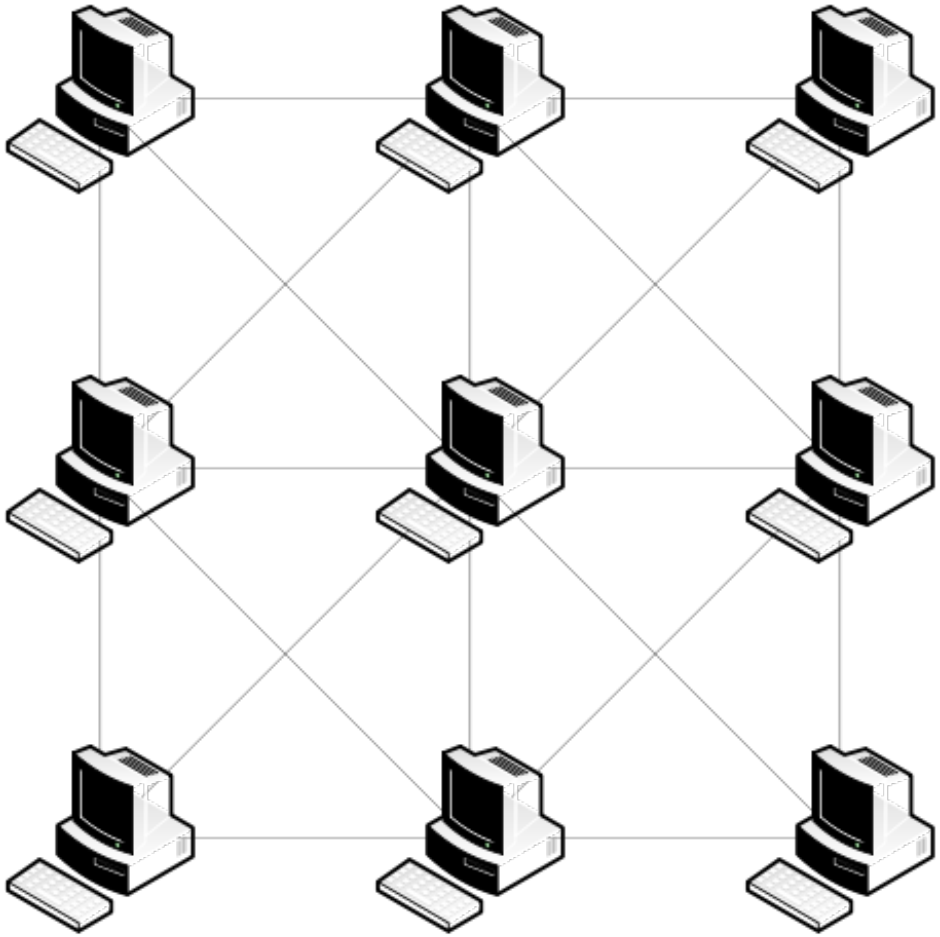

4-D Hypercube

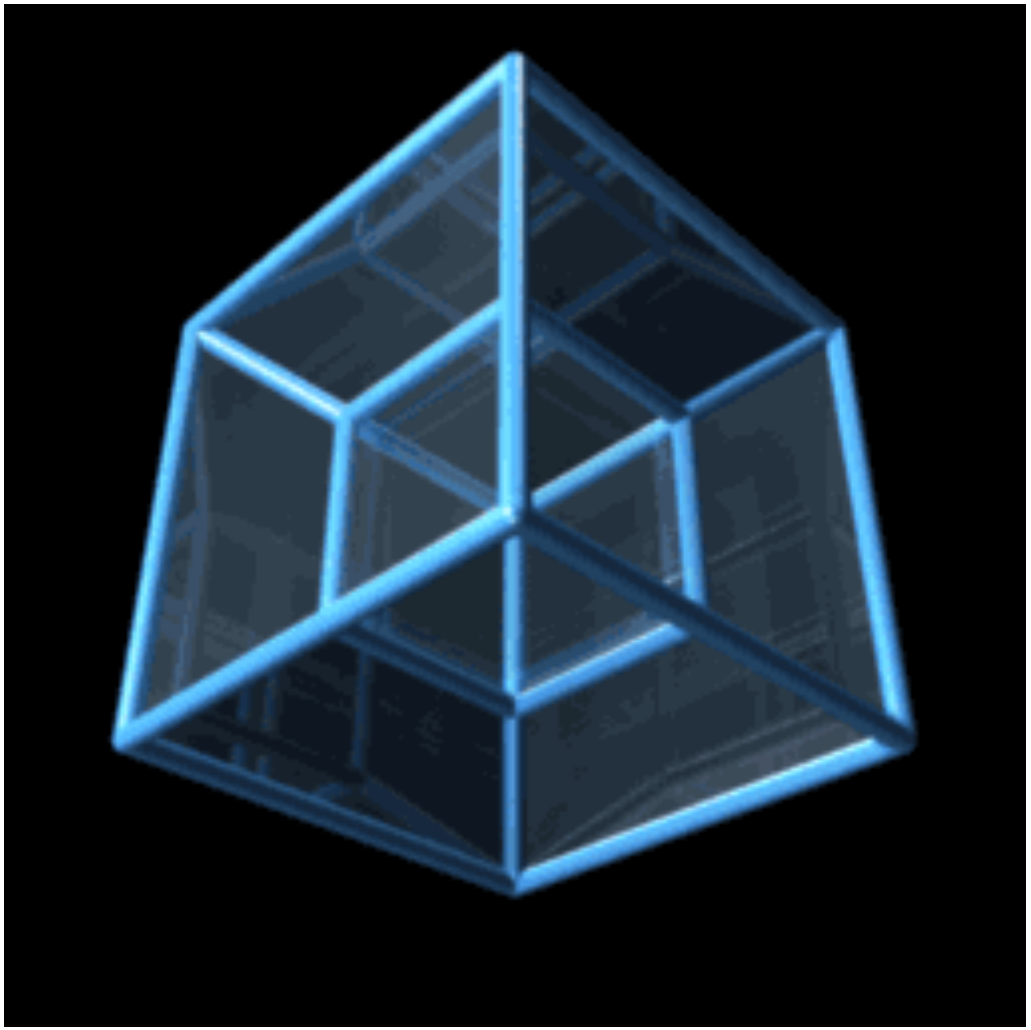

5 Major Parts of an HPC System

1. CPUs
2. Memory (volatile)
3. Nodes
4. Inter-node network
5. **Non-volatile storage (disks, tape)**

Data is getting bigger all the time

- User I/O needs growing each year in scientific community
- For our largest users I/O parallelism is mandatory
- I/O remains a bottleneck for many users
- Early 2011 – Hopper: 2 PB /scratch (we thought that was huge!)
- New systems at TACC and NCAR have ~ 18 PB / scratch!!!!

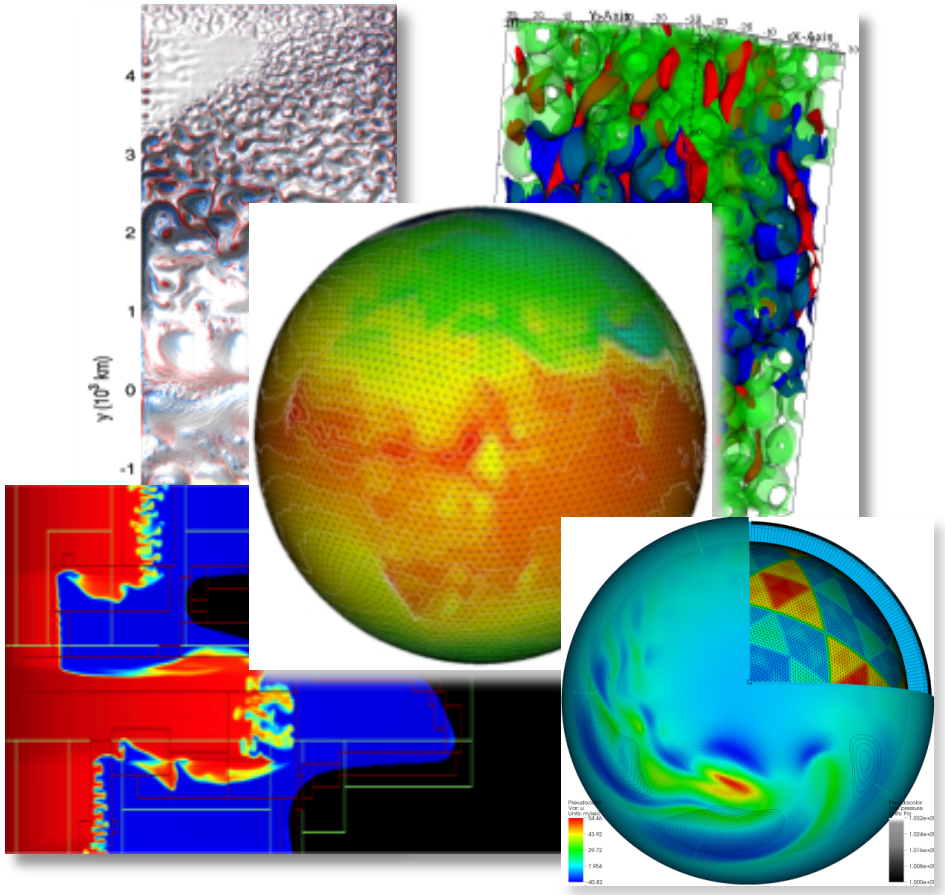

Images from David Randall, Paola Cessi, John Bell, T. Scheibe

Largest and Slowest Memory

- **File storage is the slowest level in the data memory hierarchy**
  - But it's permanent
  - Not uncommon for checkpoints / memory dumps to be taking a large fraction of total run time (>50%?)
  - NERSC users say they want no more than 10% of time to be IO
- **FLASH**
  - Non-volatile solid-state memory
  - Fast
  - Expensive
  - Some experimental systems use FLASH for fast IO

Latencies

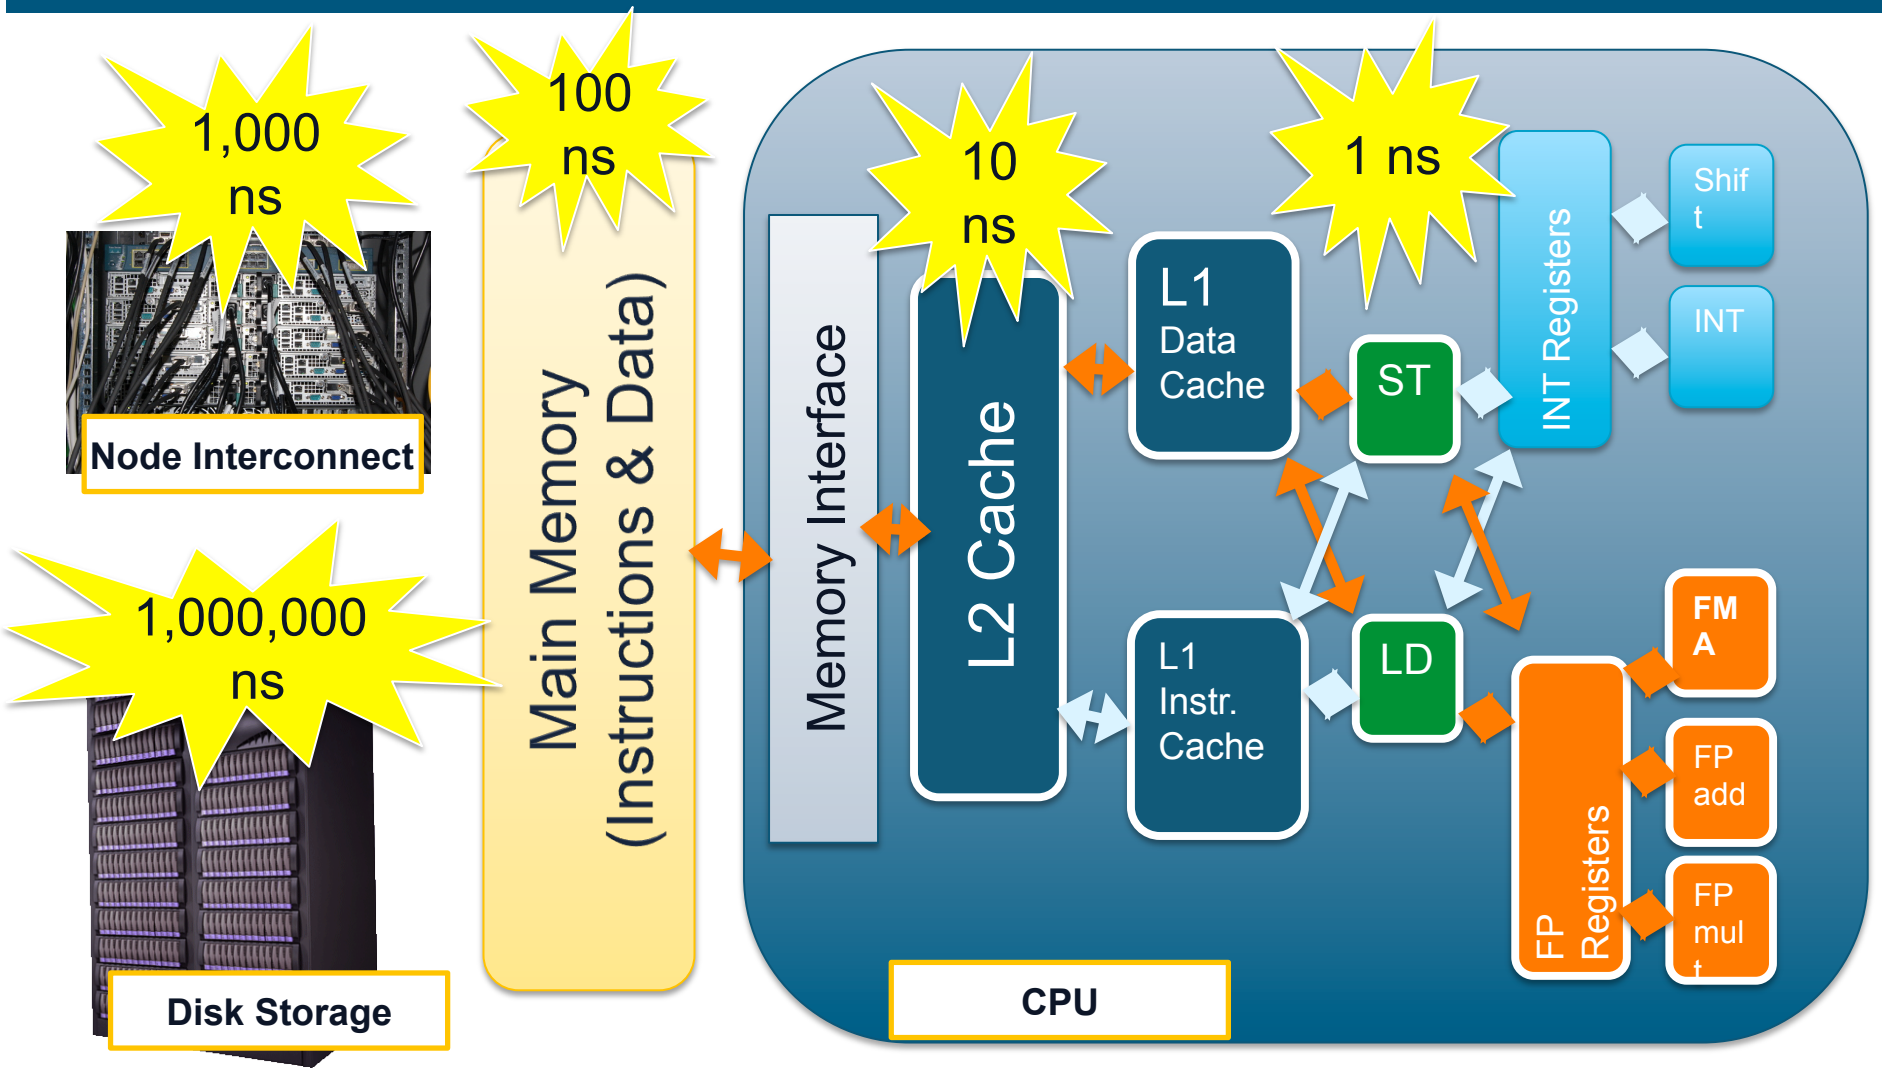
1,000 ns
Node Interconnect
100 ns
Main Memory (Instructions & Data)
1,000,000 ns
Disk Storage
10 ns
Memory Interface
L2 Cache
1 ns
L1 Data Cache
L1 Instr. Cache
ST
LD
INT Registers
Shift
INT
FP Registers
FMA
FP add
FP mult
CPU

Local vs. Global File Systems

- **"Local" – Access from one system**
  - Disk attached to node motherboard (PCI): rare in HPC
  - Network attached TB+ file systems
    - Via high-speed internal network (e.g. IB)
    - Direct from node via high-speed custom network (e.g. FibreChannel)
    - Ethernet
    - Activity by other users can impact performance
- **"Global" – Access from multiple systems**
  - Networked file system
  - Activity on other systems can impact performance
  - Useful for avoiding data replication, movement among systems

Archival Storage

- **Large, Permanent Storage**
  - Many PBs
  - Often tape storage fronted by a disk cache
  - HSM
    - Some systems may have Hierarchical Storage Management in place
    - Data automatically migrated to slower, larger storage via some policy
  - Often accessed via ftp, grid tools, and/or custom clients (e.g. hsi for HPSS)

HPC Operating Systems

- **Most HPC OSs are Linux-Based**
  - IBM AIX on POWER (also offers Linux)
- **"Generic" Cluster Systems**
  - Full Linux OS on each node
- **Specialized HPC Systems (e.g., Cray XT series, IBM Blue Gene)**
  - Full Linux OS on login, "services" nodes
  - Lightweight kernel on compute nodes
    - Helps performance
    - May hinder functionality (DLLs, dynamic process creation, some system calls may not be supported.)

HPC Systems

U.S. DEPARTMENT OF ENERGY | Office of Science

NERSC National Energy Research Scientific Computing Center

BERKELEY LAB Lawrence Berkeley National Laboratory

Outline

- Terminology
- 5 main parts of an HPC system
- CPUs
- Nodes
- Interconnect
- Data Storage
- **HPC Systems**

Top 500

- **Listing the 500 most powerful computers in the world**
- **Yardstick: Rmax of Linpack**
  - Solve Ax=b, dense problem, matrix is random
  - Dominated by dense matrix-matrix multiply
- **Update twice a year:**
  - ISC'xy in June in Germany
  - SCxy in November in the U.S.
- **All information available from the TOP500 web site at: www.top500.org**

June 2011 Top 500 List

Rank	Site	Man.	Computer	Country	Cores	Rmax [Tflops]	Power [MW]
1	RIKEN	Fujitsu	K Computer SPARC 64 2GHz, Tofu	Japan	548,352	8,162	9.89
2	National SuperComputer Center in Tianjin	NUDT	Tianhe-1A NUDT TH MPP, Xeon 6C, NVidia, FT-1000 8C	China	186,368	2,566	4.04
3	Oak Ridge National Laboratory	Cray	Jaguar Cray XT5, HC 2.6 GHz	USA	224,162	1,759	6.95
4	National Supercomputing Centre in Shenzhen	Dawning	Nebulae TC3600 Blade, Intel X5650, NVidia Tesla C2050 GPU	China	120,640	1,271	2.58
5	GSIC, Tokyo Inst of Tech	NEC/HP	TSUBAME-2 HP ProLiant, Xeon 6C, NVidia, Linux/ Windows	Japan	73,278	1,192	1.40
6	DOE/NNSA/ LANL/SN	Cray	Cielo Cray XE6, 8C 2.1 GHz	USA	142,272	1,100	3.98
7	NASA / Ames	SGI	Pleiades SGI Altix ICE, Xenon 2.9 GHz, IB	USA	111,104	1,088	4.10
8	DOE/SC/ LBNL/NERSC	Cray	Hopper Cray XE6, 12C 2.1 GHz	USA	153,408	1.054	2.91
9	Commissariat a l'Energie Atomique (CEA)	Bull	Tera 100 Bull bullx super-node S6010/S6030	France	138.368	1,050	4.59
10	DOE/NNSA/LANL	IBM	Roadrunner BladeCenter QS22/LS21	USA	122,400	1,042	2.34

GPUs

- **3 of top 5 are GPU-accelerated systems**
- **"Graphics Processing Units" are composed of 100s of simple "cores" that increase data-level on-chip parallelism**
- **Yet more low-level complexity to consider**
  - Another interface with the socket (or on socket?)
  - Limited, private memory (for now?)
- **Programmability is currently poor**
- **Legacy codes may have to be rewritten to minimize data movement**
- **Not all algorithms map well to GPUs**
- **What is their future in HPC?????**

Performance Development (Top 500)

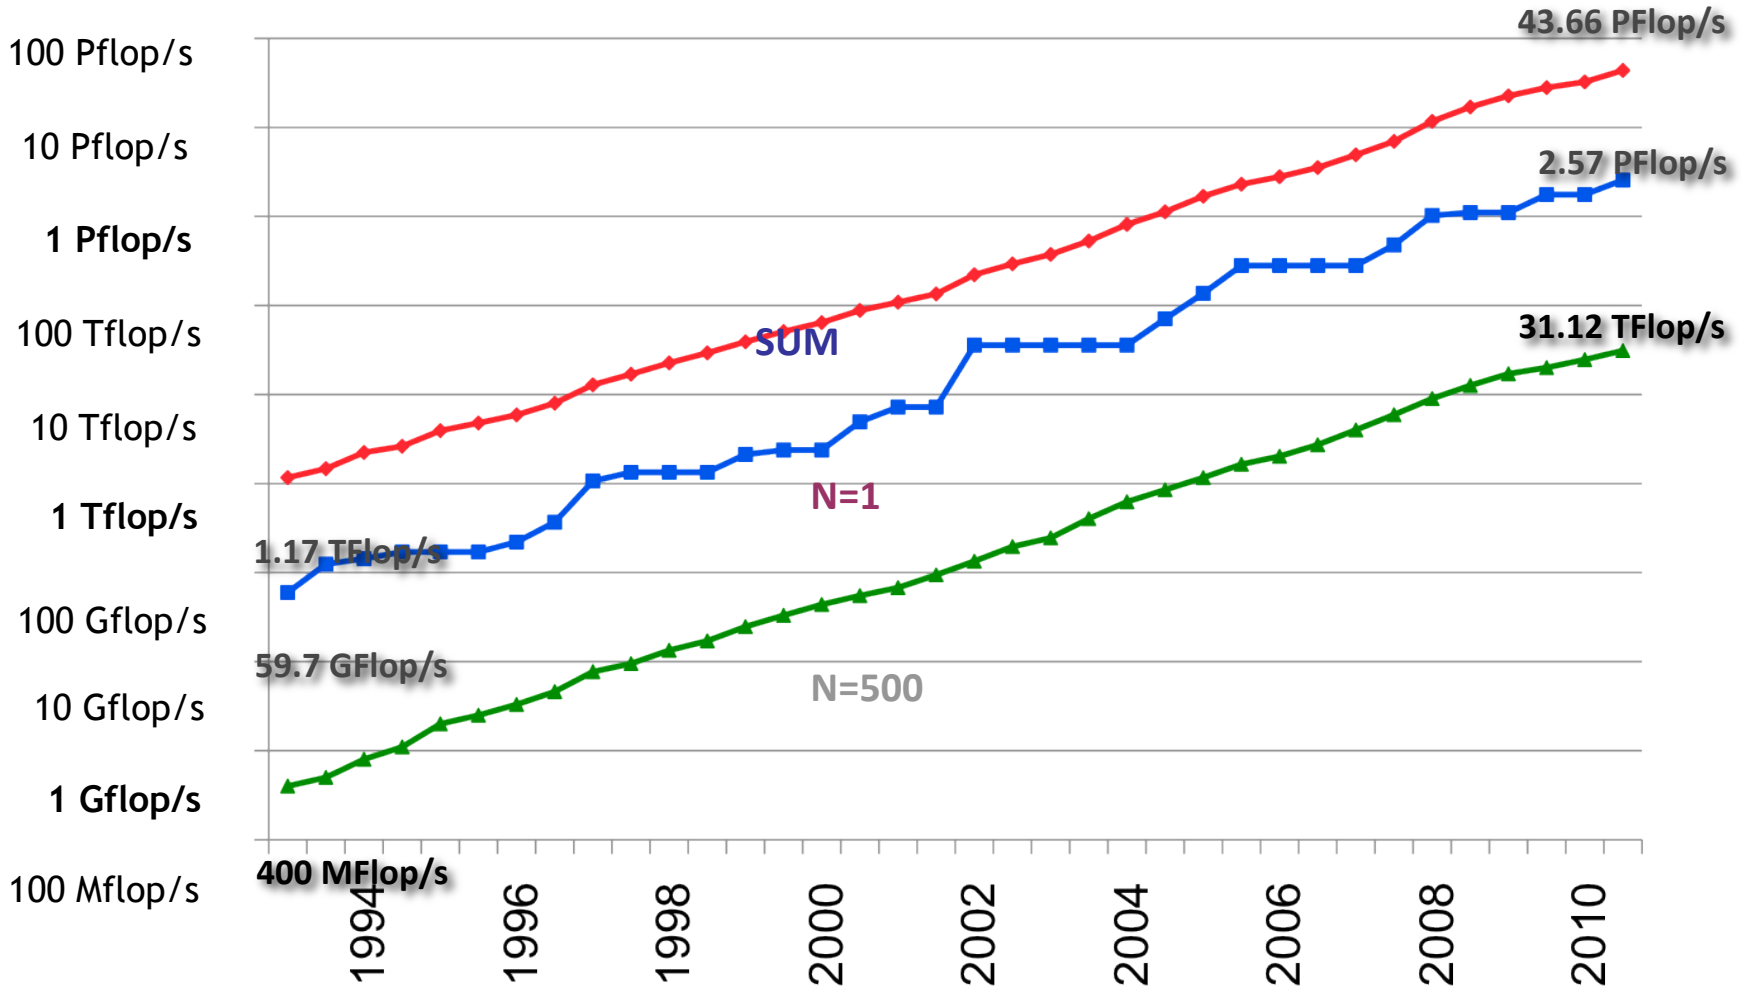

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB Lawrence Berkeley National Laboratory

Projected Performance Development

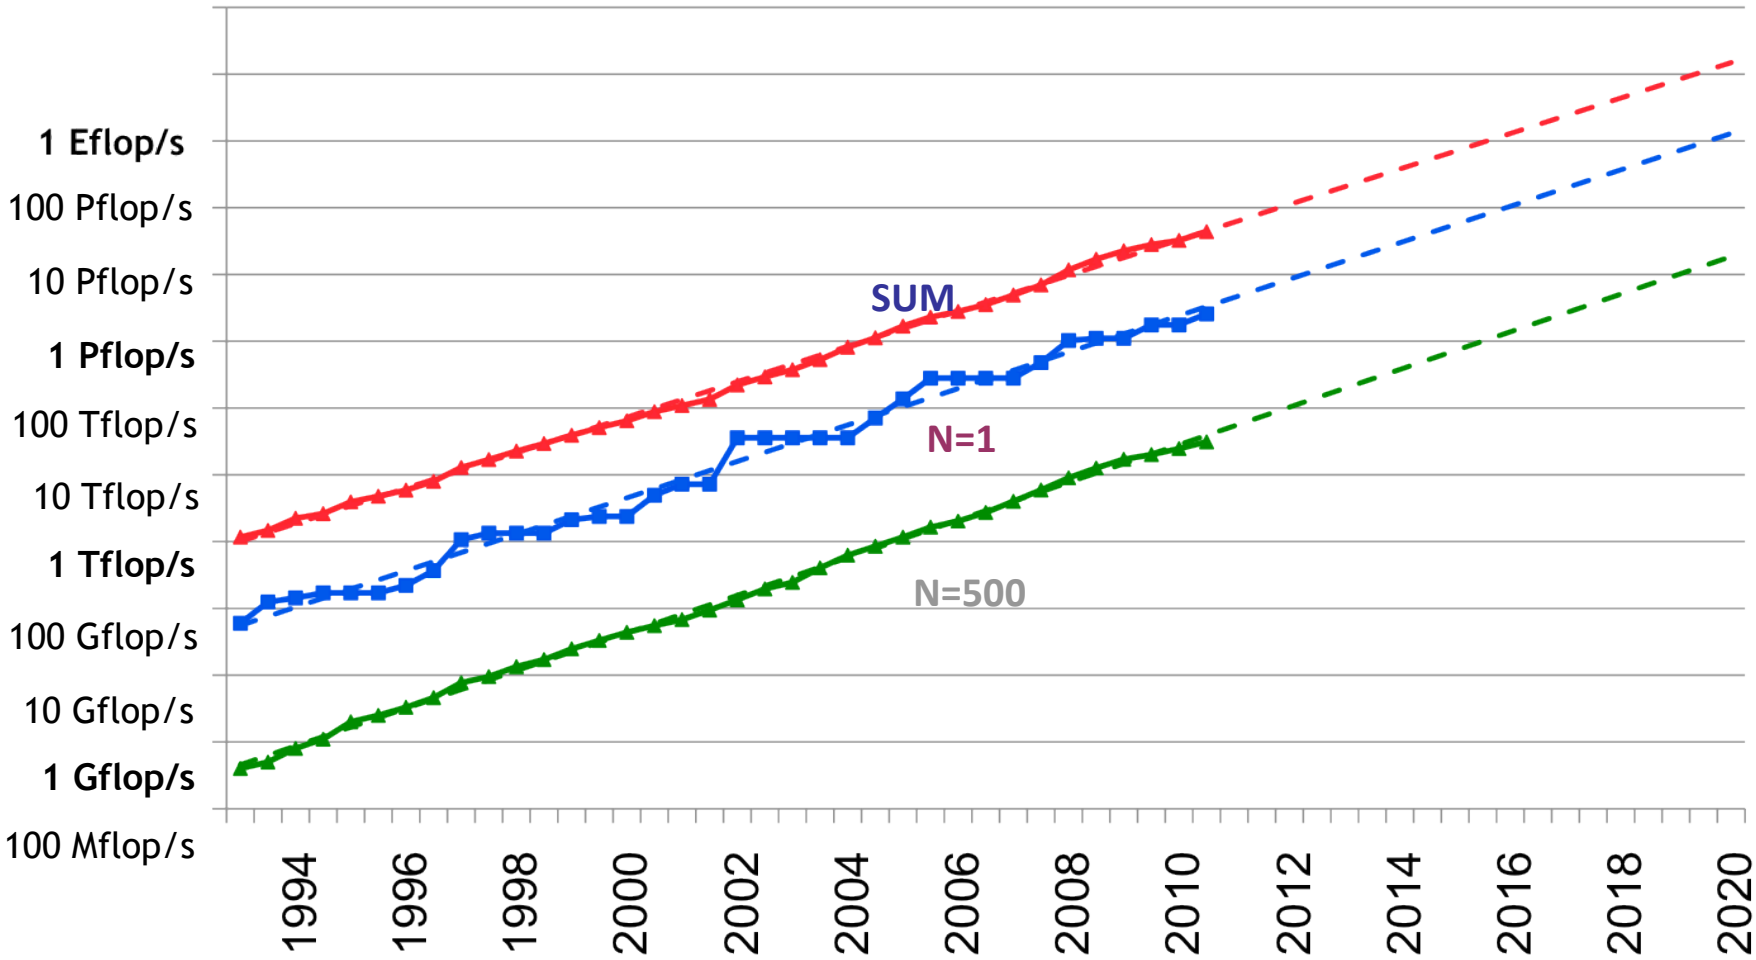

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB Lawrence Berkeley National Laboratory

Core Count

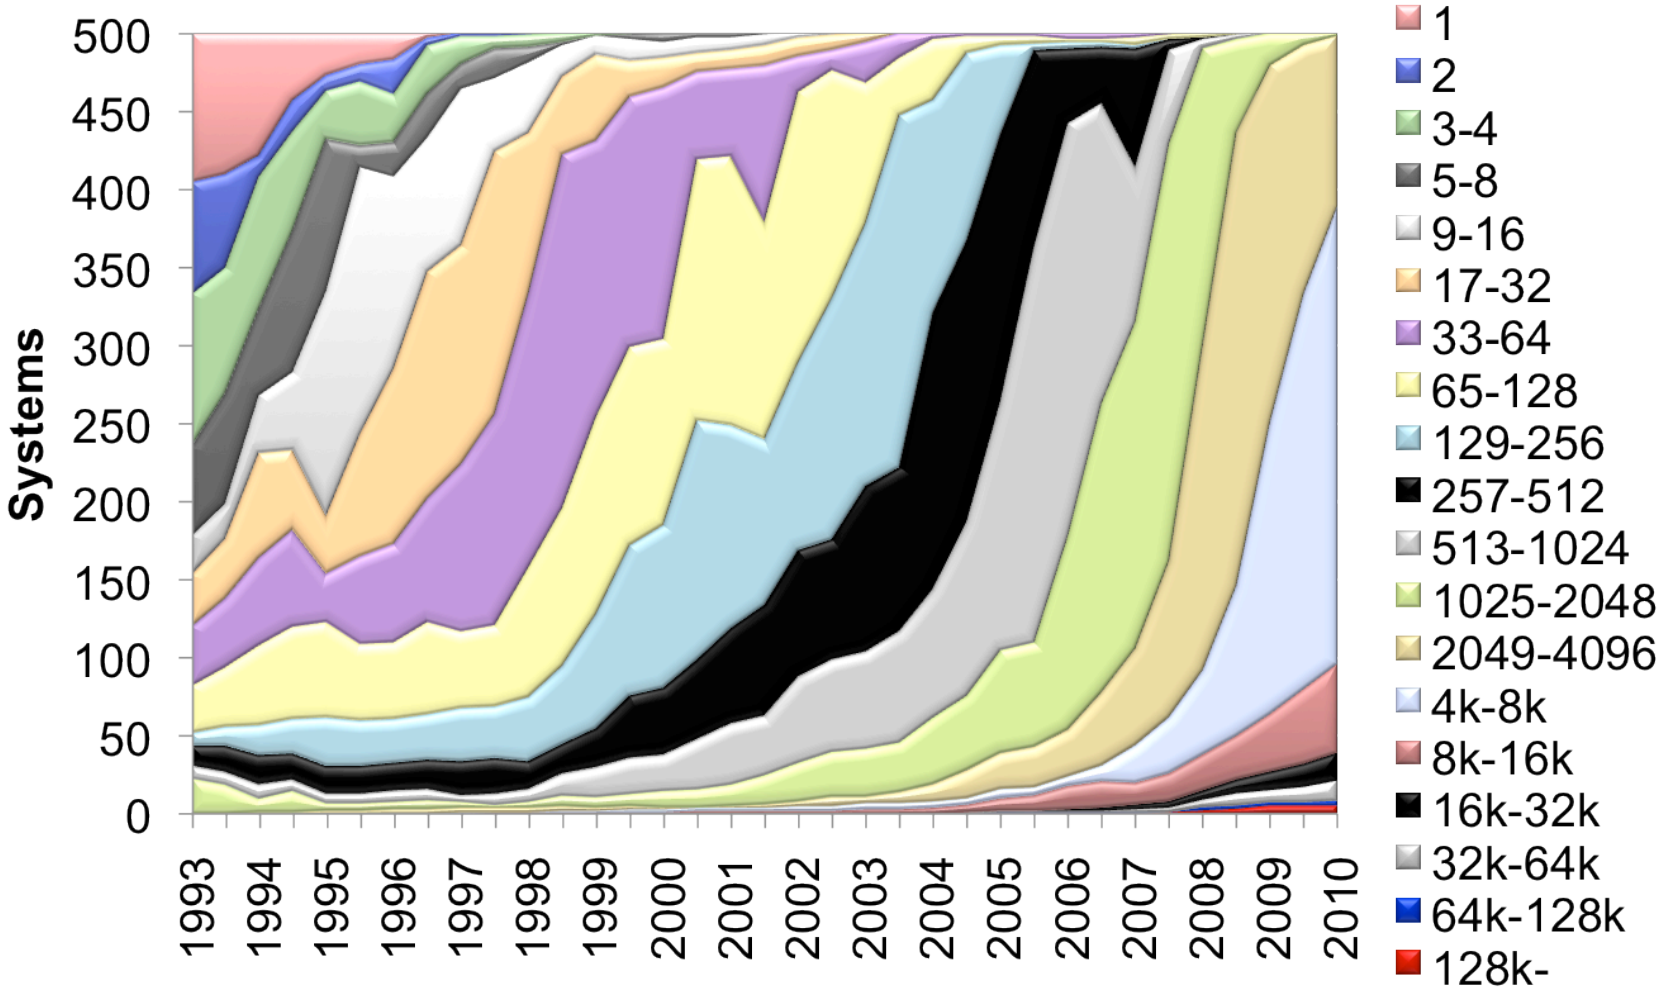
Systems
500
450
400
350
300
250
200
150
100
50
0
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
1
2
3-4
5-8
9-16
17-32
33-64
65-128
129-256
257-512
513-1024
1025-2048
2049-4096
4k-8k
8k-16k
16k-32k
32k-64k
64k-128k
128k-

National Energy Research Scientific Computing Center

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB Lawrence Berkeley National Laboratory