



中山大學
SUN YAT-SEN UNIVERSITY



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

Advanced Computer Architecture

高级计算机体系结构

第6讲：DLP and GPU (1)

张献伟

xianweiz.github.io

DCS5367, 11/9/2021



中山大學
SUN YAT-SEN UNIVERSITY



作业 – HW2

- <https://xianweiz.github.io/teach/dcs5637/f2021.html>
- 截止时间: **11.21**, 23:59
- 提交方式: 超算习堂
 - 注册 (<https://easyhpc.net/>)
 - 加入课程 (<https://easyhpc.net/course/133>)
 - 作业列表: HW2



HW-2, review-form.txt

John L. Hennessy | David A. Patterson
COMPUTER
ARCHITECTURE

A Quantitative Approach

DCS5637 - 高级计算机体系结构 (周二)



0人 其他

Advanced Computer Architecture, Fall 2021 课程主页: <https://xianweiz.github.io/teach/dcs5637/f2021.html>

加入课程



中山大學
SUN YAT-SEN UNIVERSITY



Quiz Questions


- Q1: list 3+ design goals in computer architecture?
- Q2: typical pipeline stages of an instruction?
- Q3: list at least three techniques to improve ILP?
- Q4: usage of register renaming?
- Q5: briefly explain 'memory wall'
- Q6: sort GDDR6/DDR4/HBM2 in bandwidth (lower first)

Quiz Questions

- Q1: list 3+ design goals in computer architecture?
Functional, high-perf, reliable, low-cost, low-power, ...
- Q2: typical pipeline stages of an instruction?
Instruction fetch, inst decode, execute, mem access, write back
- Q3: list at least three techniques to improve ILP?
Pipelining, scheduling, unrolling, prediction, multi issue, ...
- Q4: usage of register renaming?
Remove name dependences to improve ILP
- Q5: briefly explain 'memory wall'
Memory is far slower than processors
- Q6: sort GDDR6/DDR4/HBM2 in bandwidth (lower first)
 $\text{DDR4 (64x2.6Gb/s)} < \text{GDDR6 (256x14Gb/s)} < \text{HBM2 (4096x1Gb/s)}$


Serial Computing[串行计算]

- Traditionally, software has been written for serial computation
 - To be run on a single computer having a single CPU
 - A problem is broken into a discrete series of instructions
 - Instructions are executed one after another
 - Only one instruction may execute at any moment



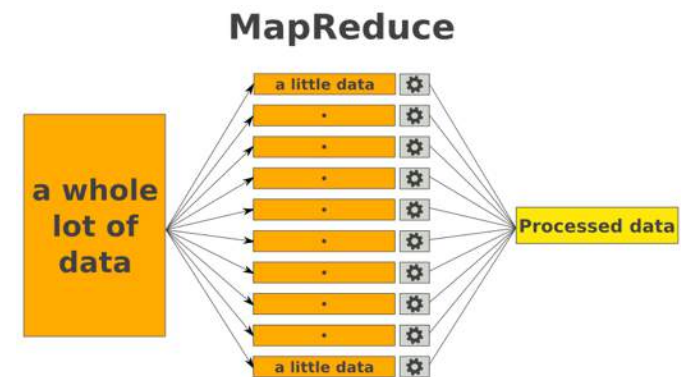


Parallel Computing[并行计算]

- Simultaneously use multiple compute resources to solve a computational problem
 - Typically in high-performance computing (HPC)
- HPC focuses on performance
 - To solve biggest possible problems in the least possible time



Types of Parallel Computing[并行类型]

- **Instruction** level parallelism[指令级并行]
 - Classic RISC pipeline (fetch, ..., write back)
- **Task** parallelism[任务级并行]
 - Different operations are performed concurrently
 - Task parallelism is achieved when the processors execute on the same or different data
- **Data** parallelism[数据级并行]
 - Distribution of data across different parallel computing nodes
 - Data parallelism is achieved when each processor performs the same task on different pieces of the data




Taxonomy[分类]

- **Flynn's Taxonomy** (1966) is widely used to classify parallel computers
 - Distinguishes multi-processor computer architectures according to how they can be classified along the two independent dimensions of **Instruction Stream** and **Data Stream**
 - Each of these dimensions can have only one of two possible states: **Single** or **Multiple**
- 4 possible classifications according to Flynn

S I S D Single Instruction stream Single Data stream	S I M D Single Instruction stream Multiple Data stream
M I S D Multiple Instruction stream Single Data stream	M I M D Multiple Instruction stream Multiple Data stream


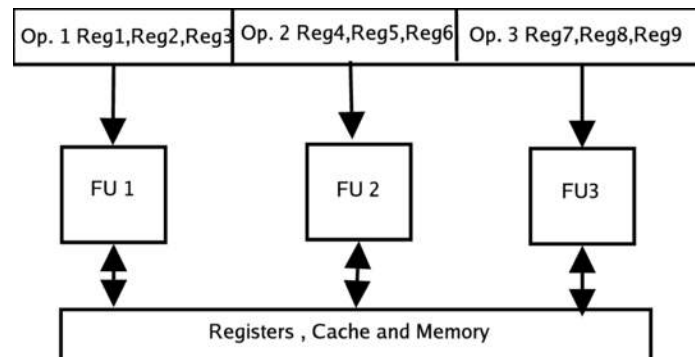
Taxonomy (cont.)

- SISD: single instruction, single data
 - A serial (non-parallel) computer
- **SIMD**: single instruction, multiple data
 - Best suited for specialized problems characterized by a high degree of regularity, such as graphics/image processing
- MISD: multiple instruction, single data
 - Few (if any) actual examples of this class have ever existed
- MIMD: multiple instruction, multiple data
 - Examples: supercomputers, multi-core PCs, VLIW



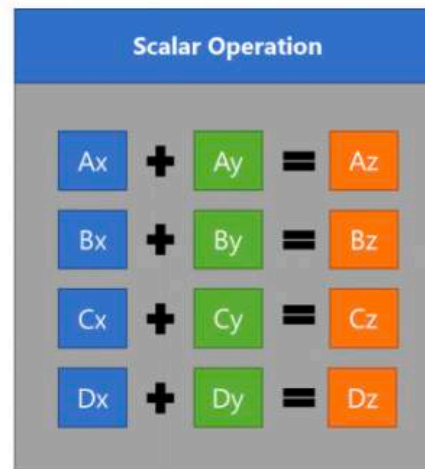
SIMD: vs. superscalar and VLIW[对比]

- SIMD performs the same operation on multiple data elements with **one single instruction**
 - Data-level parallelism
- Superscalar dynamically issues **multi insts** per clock[超标量]
 - Instruction level parallelism (ILP)
- VLIW receives **long instruction words**, each comprising a field (or opcode) for each execution unit[超长指令字]
 - Instruction level parallelism (ILP)

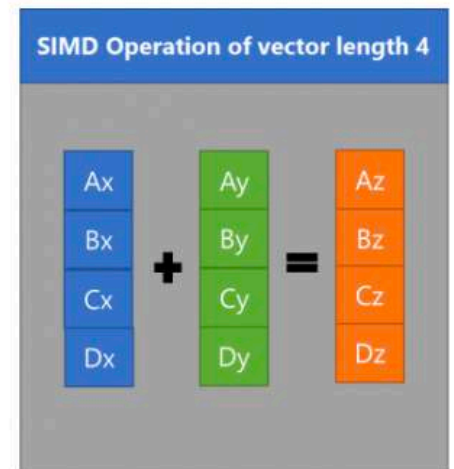


SIMD: Vector Processors[向量处理器]

- Vector processor (or array processor)[处理器]
 - CPU that implements an instruction set containing instructions that operate on one-dimensional arrays (vectors)
- People use vector processing in many areas[应用]
 - Scientific computing
 - Multimedia processing (compression, graphics, image processing, ...)
- Instruction sets[指令集]
 - MMX
 - SSE
 - AVX
 - NEON
 - ...



Single Instruction Single Data:



Single Instruction Multiple Data:

SIMD: MMX


- MMX is officially a meaningless initialism trademarked by Intel; unofficially,
 - MultiMedia eXtension
 - Multiple Math eXtension
 - Matrix Math eXtension
- Introduced on the “Pentium with MMX Technology” in 1998
- SIMD computation processes multiple data in parallel with a single instruction
 - MMX gives 2 x 32-bit computations at once
 - MMX defined 8 “new” 64-bit integer registers (mm0 ~ mm7)
 - 3DNow! was the AMD extension of MMX




SIMD: SSE

- Streaming SIMD Extensions


- SSE defines 8 new 128-bit registers (xmm0 ~ xmm7) for FP32 computations
 - Since each register is 128-bit long, we can store total 4 FP32 numbers
- 4 simultaneous 32-bit computations



`mulss xmm1, xmm0`




`mulps xmm1, xmm0`




SIMD: AVX

- Advanced Vector Extensions (AVX)
 - A new-256 bit instruction set extension to SSE
 - 16-registers available in x86-64
 - Registers renamed from XMMi to YMMi
 - Yet a proposed extension is AVX-512
 - A 512-bit extension to the 256-bit XMM
 - Supported in from Intel's Xeon Phi x200 (Knights Landing) and Skylake-SP, and onwards



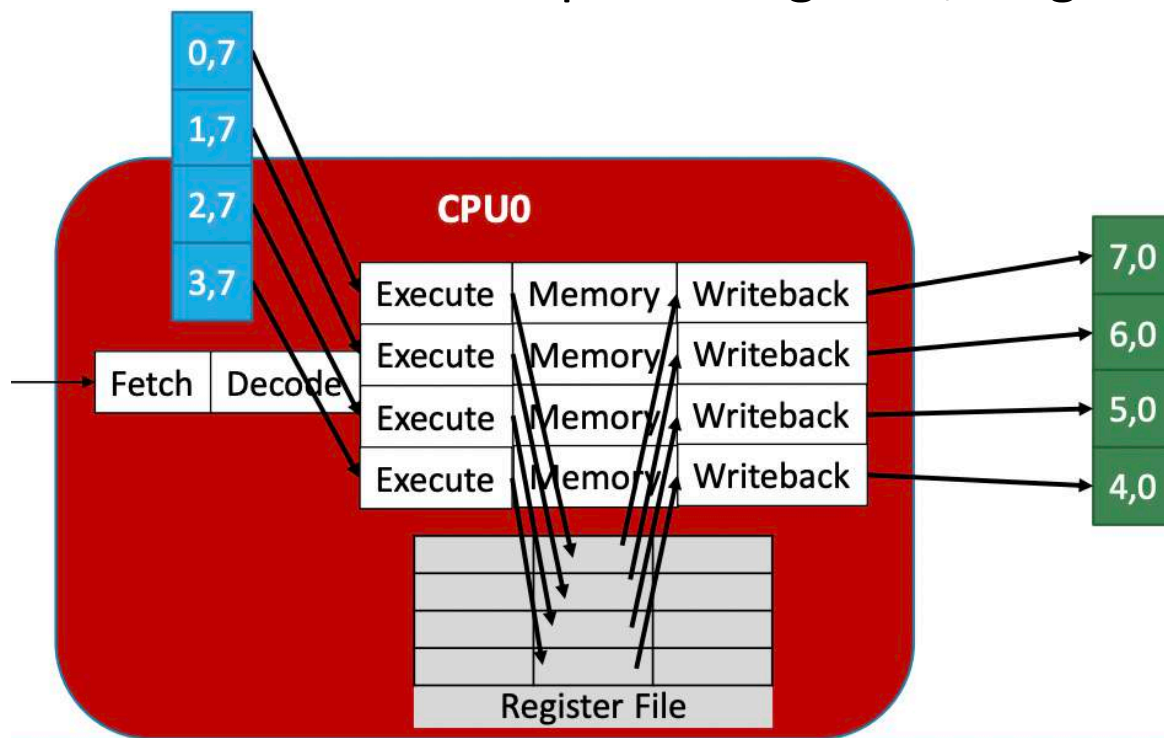
SIMD: NEON

- ARM Advanced SIMD Extensions
 - Introduced by ARM in 2004 to accelerate media and signal processing
 - NEON can for example execute MP3 decoding on CPUs running at 10 MHz
 - 128-bit SIMD Extension for the ARMv7 & ARMv8
 - Data types can be: signed/unsigned 8-bit, 16-bit, 32-bit or 64-bit




Data Parallelism: SIMD

- Single Instruction Multiple Data
 - Split identical, independent work over multiple execution units (lanes)
 - More efficient: eliminate redundant fetch/decode
 - One Thread + Data Parallel Ops → Single PC, single register file




Data Parallelism: SIMT

- Single Instruction Multiple Thread
 - Split identical, independent work over multiple threads
 - Multiple Threads + Scalar Ops → One PC, multiple register files
 - \approx SIMD + multithreading
 - Each thread has its own registers



Execution Model[执行模型]

MIMD



Multiple
independent threads

Multicore CPUs

SIMD



One thread with wide
execution datapath

x86 SSE/AVX

SIMT



Multiple lockstep
threads

GPUs

- SI(MD/MT)

- Broadcasting the same instruction to multiple execution units
- Replicate the execution units, but they all share the same fetch/decode hardware

SIMD and SIMT are used interchangeably



SIMD: GPU vs. CPU/Traditional

- Traditional SIMD contains a **single thread**
 - Programming model is SIMD (no threads)
 - SW needs to know vector length
 - ISA contains vector/SIMD instructions
- GPU SIMD consists of **multiple scalar threads** executing in a SIMD manner (i.e., same instruction executed by all threads)
 - Each thread can be treated individually (i.e., placed in a different warp) → programming model not SIMD
 - SW does not need to know vector length
 - Enables memory and branch latency tolerance
 - ISA is scalar → vector instructions formed dynamically
- Essentially, it is SPMD programming model implemented on SIMD hardware

Example: add two vectors

C:

```
for(i=0;i<n;++i) a[i]=b[i]+c[i];
```

Matlab:

```
a=b+c;
```

SIMD:

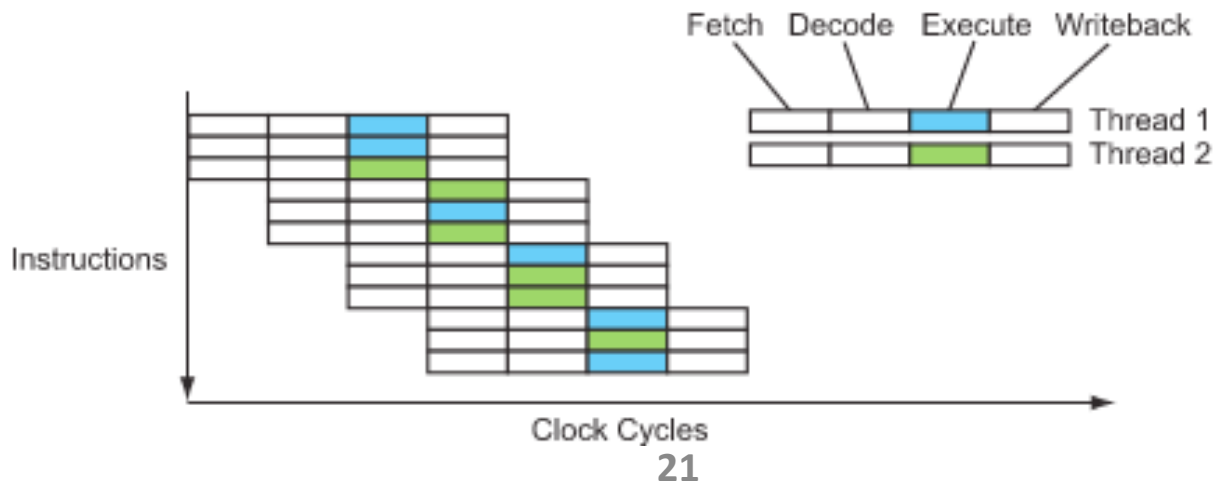
```
void add(uint32_t *a, uint32_t *b, uint32_t *c, int n) {  
    for(int i=0; i<n; i+=4) {  
        //compute c[i], c[i+1], c[i+2], c[i+3]  
        uint32x4_t a4 = vld1q_u32(a+i);  
        uint32x4_t b4 = vld1q_u32(b+i);  
        uint32x4_t c4 = vaddq_u32(a4,b4);  
        vst1q_u32(c+i,c4);  
    }  
}
```

SIMT:

```
__global__ void add(float *a, float *b, float *c) {  
    int i = blockIdx.x * blockDim.x + threadIdx.x;  
    a[i]=b[i]+c[i]; //no loop!  
}
```


SMT[多线程]

- SMT: simultaneous multithreading
 - Instructions from multiple threads issued on the same cycle
 - Use register renaming and dynamic scheduling facility of multi-issue architecture
 - Needs more hardware support
 - Register files, PC's for each thread
 - Support to sort out which threads to get results from which instructions
 - Thread scheduling, context switching
 - Maximize utilization of execution units



SMT vs. SIMT[比较]

- SMT: maximize the chances of an instruction to be issued without having to switch to another thread
 - superscalar execution
 - out-of-order execution
 - register renaming
 - branch prediction
 - speculative execution
 - cache hierarchy
 - speculative prefetching
- SIMT: keep massive threads to achieve high throughput
 - Hardware becomes simpler and cheaper
 - No OoO, no prefetching, ...





CPU vs. GPU[比较]

- CPU

- Low compute density
- Complex control logic
- Fewer cores optimized for serial operations
 - Fewer execution units (ALUs)
 - Higher clock speeds
- Low latency tolerance

- GPU


- High compute density
- Simple control logic
- 1000s cores optimized for parallel operations
 - Many parallel execution units (ALUs)
 - Lower clock speeds
- High latency tolerance



GPU Overview


Software

```
__global__ void scale(float a, float * X)
{
    unsigned int tid;
    tid = blockIdx.x * blockDim.x
        + threadIdx.x;
    X[tid] = a * X[tid];
}
```



Architecture: **multi-thread** programming model

SIMT microarchitecture



Hardware datapaths: **SIMD** execution units



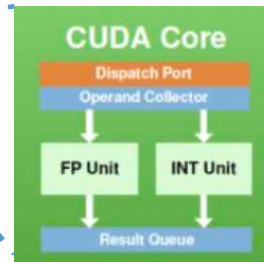
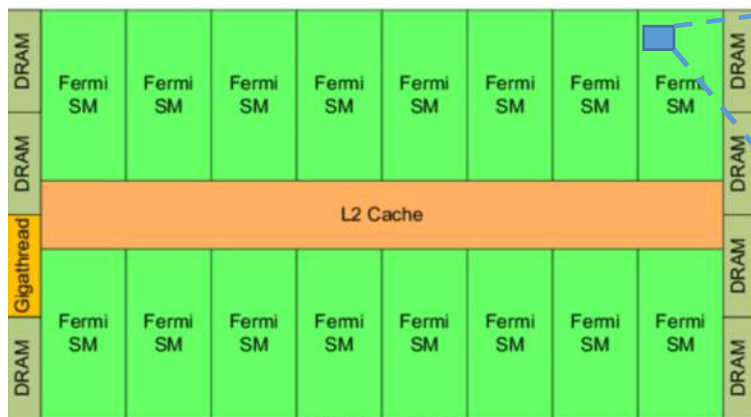
Hardware

3

GPU Overview(cont.)

- A GPU contains several largely independent processors called "**Streaming Multiprocessors**" (SMs)
 - Each SM hosts multiple "**cores**", and each "core" runs a thread
 - For instance, Fermi(2010) has up to 16 SMs w/ 32 cores per SM
 - So up to 512 threads can run in parallel
- Some SIMT threads are grouped to execute in lockstep
 - One warp contains 32 threads
- Multiple '**groups**' can be executed simultaneously
 - For Fermi, up to 48 warps per SM


A100: 128 SMs w/ 64 cores per SM



25

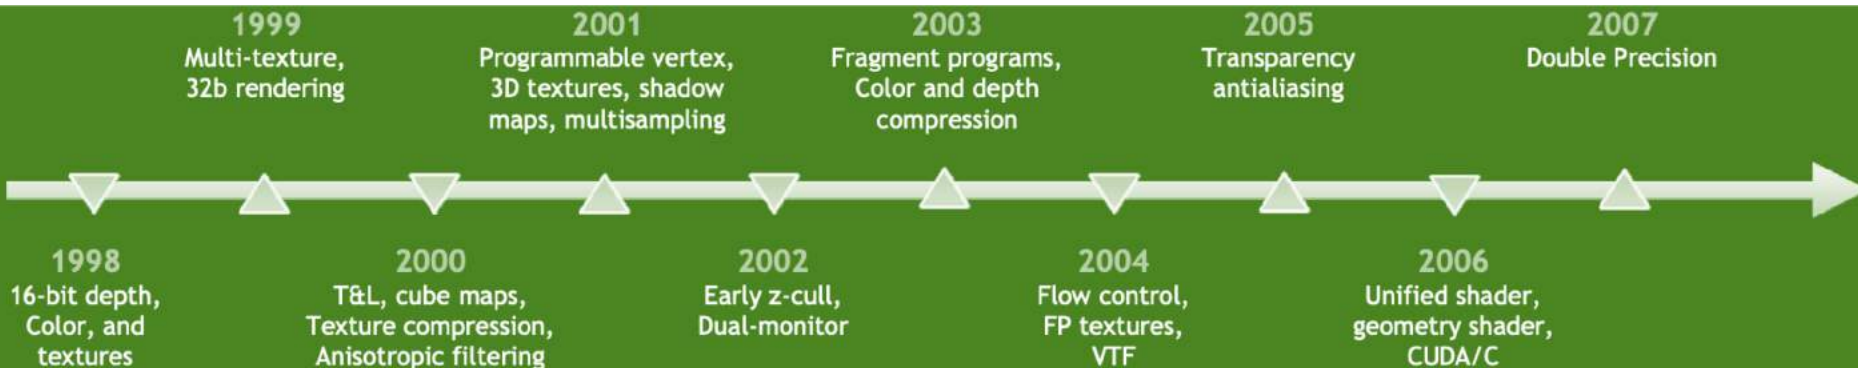
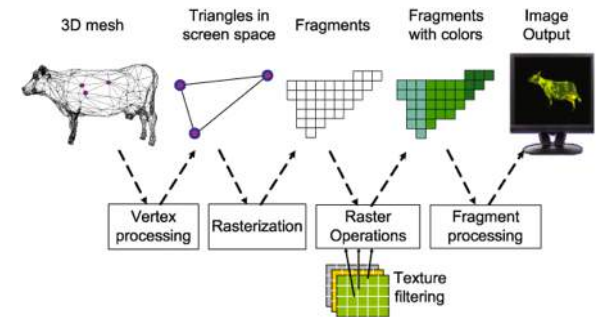
Example of SIMT Execution

Assume 32 threads are grouped into one warp.



GPU Evolution[演进]

- Arcade boards and display adapters (1951 - 1995)
 - ATI: founded in 1985
 - Nvidia: founded in 1993
- 3D revolution (1995 – 2006)
 - Term “graphics processing unit”: 1999
 - Nvidia GeForce 256
 - Rivalry between ATI and Nvidia
- General purpose GPU (2006 - present)
 - AI , data analytics, scientific computing, graphics rendering, etc



GPGPU History[简史]

Year	AMD	Nvidia	Note
2006	AMD acquired ATI	Tesla (CUDA Launch)	Unified shader model
2007	TeraScale		Unified shader uarch
2009	TeraScale 2		
2010	TeraScale 3	Fermi / GTX580	First compute GPU
2011	GCN 1.0 / gfx6		VLW → SIMD
2012		Kepler / GTX680	CUDA cores: 512 → 1536
2013	GCN 2.0 / gfx7		
2014	GCN 3.0 / gfx8	Maxwell / GTX980	Energy efficiency
2016	GCN 4.0 / gfx8	Pascal / GTX1080	
2017	GCN 5.0 / gfx9	Volta / GV100	First chip with Tensor cores
2018	GCN 5.1 / gfx9	Turing / RTX2080	
2019	RDNA 1.0 / gfx10		
2020	RDNA 2.0 / gfx10 CDNA 1.0 / gfx9	Ampere / RTX3090	First chip with Matrix cores

TFLOPS[衡量算力]

- A100 Tensor Core GPU

- 108 SMs

- GA100 Full GPU with 128 SMs

- Base clock: 1065 MHz

- Boost clock: 1410 MHz

- Performance

- FP64: 9.7 TFLOPS

- FP32: 19.5 TFLOPS

- Calculate TFLOPS

- FP64: $1410 \text{ MHz} \times (32 \times 2) \text{ ops/clock} \times 108 \text{ SMs}$



GPUs in Supercomputer[超算中的GPU]

- Exascale: 50 GFLOPS/Watt (goal) → **51.7** GFLOPS/Watt

System	Titan (2012)	Summit (2017)	Frontier (2021)
Peak	27 PF	200 PF	> 1.5 EF
# nodes	18,688	4,608	> 9,000
Node	1 AMD Opteron CPU 1 NVIDIA Kepler GPU	2 IBM POWER9™ CPUs 6 NVIDIA Volta GPUs	1 AMD EPYC CPU 4 AMD Radeon Instinct GPUs 40+ TFLOPS
Memory		2.4 PB DDR4 + 0.4 HBM + 7.4 PB On-node storage	4.6 PB DDR4 + 4.6 PB HBM2e + 36 PB On-node storage, 75 TB/s Read 38 Write
On-node interconnect	PCI Gen2 No coherence across the node	NVIDIA NVLINK Coherent memory across the node	AMD Infinity Fabric Coherent memory across the node
System Interconnect	Cray Gemini network 6.4 GB/s	Mellanox Dual-port EDR IB 25 GB/s	Four-port Slingshot network 100 GB/s
Topology	3D Torus	Non-blocking Fat Tree	Dragonfly
Storage	32 PB, 1 TB/s, Lustre Filesystem	250 PB, 2.5 TB/s, IBM Spectrum Scale™ with GPFS™	695 PB HDD+11 PB Flash Performance Tier, 9.4 TB/s and 10 PB Metadata Flash. Lustre
Power	9 MW	13 MW	29 MW

Frontier: 1.5 EFLOPS, How???[E级超算]

- Per node[单节点]

- Custom EPYC HPC-optimized CPU

- "zen 3" milan w/ 64-core

- Four Instinct GPUs

- CDNA MI200 w/ 256 CUs

- Full-rate FP64 (128 ops/clock/CU)

- 9000+ nodes[整体系统]


- CPU: $9000 \times 4 \text{ TFLOPS/CPU} = 36 \text{ PFLOPS}$

- GPU: $9000 \times 4 \times 42.2 \text{ TFLOPS/GPU} = 1519 \text{ PFLOPS}$

- Per GPU: $128 \text{ ops/clock} \times 1.5\text{G} \times 220 = 42.2 \text{ TFLOPS}$

- GPU provides **97.7%** computation power

- $1519 / (1519 + 36)$



A100: 9.75 TFLOPS
MI100: 11.54 TFLOPS

OLCF spock training: AMD hardware and software, 05/2021,

<https://www.olcf.ornl.gov/wp-content/uploads/2021/04/Spock-MI100-Update-5.20.21.pdf>

<https://www.hpcwire.com/2021/03/15/amd-launches-epyc-milan-with-19-skus-for-hpc-enterprise-and-hyperscale/>

天河超算

- 2009, 天河-1

- CPU + **ATI GPU**

240 GFLOPS

- 2 * Xeon E5540/E5450, 1 ATI Radeon HD 4870 X2 (TeraScale)

- 实测/峰值563.1T/1206.2T FLOPS

- 2009.11 TOP500第五

- 2010, 天河-1A

- CPU + **Nvidia GPU**

515 GFLOPS

- 2 * Intel Xeon X5670, 1 Nvidia Tesla M2050 (Fermi)

- 2048 Galaxy "FT-1000" 1 GHz 8-core processors

- 实测/峰值2.566P/4.7P FLOPS

- 2010.11 TOP500第一

Tianhe-1, <https://www.top500.org/system/176546/>

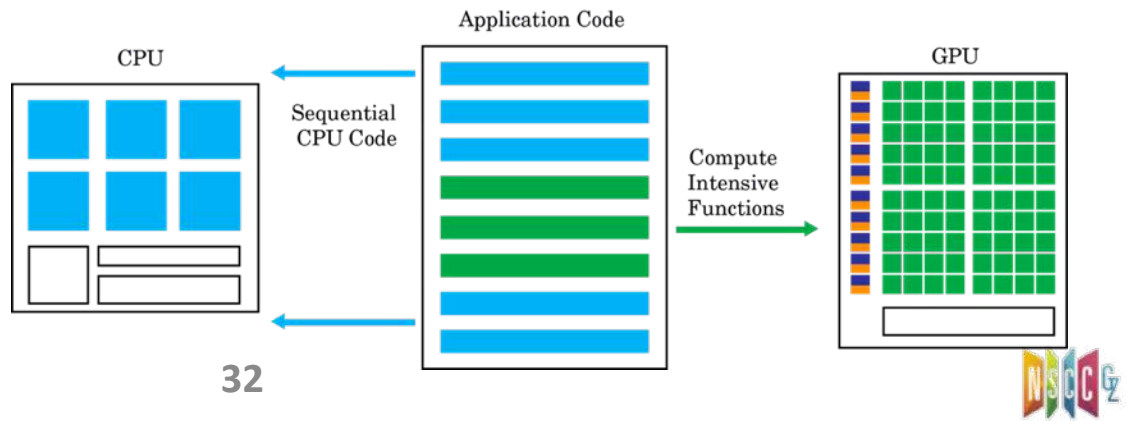

Tianhe-1A, <https://top500.org/system/176929/>

Tianhe-1A, <http://blog.zorinaq.com/introducing-tianhe-1a-4702-tflops-of-gpu-power-made-in-china-and/>




GPU Programming Model[编程模型]

- GPU is viewed as a compute **device** that
 - Is a coprocessor to CPU (**host**)
 - Has its own main memory called **device memory**
 - Runs **many threads** in parallel
- Data-parallel parts of an application are executed on the device as **kernels**, which run in parallel on many threads
- CPU thread vs. GPU thread
 - GPU threads are very lightweight
 - A few vs. thousands for full efficiency



Thread Organization[线程组织]

- A kernel is executed as a **grid** of thread blocks
- A thread **block** is a batch of threads that can cooperate with each other by
 - Synchronizing their execution
 - Efficiently sharing data through low-latency shared memory
- The grid and its associated blocks are just organizational constructs
 - The **threads** are the things that do the work



GPU Programming Choices[编程选择]

- **CUDA** – Compute Unified Device Architecture

- Developed by Nvidia – proprietary
- First serious GPGPU language/environment



- **OpenCL** – Open Computing Language

- From makers of OpenGL
- Wide industry support: AMD, Apple, Qualcomm, Nvidia (begrudgingly), etc



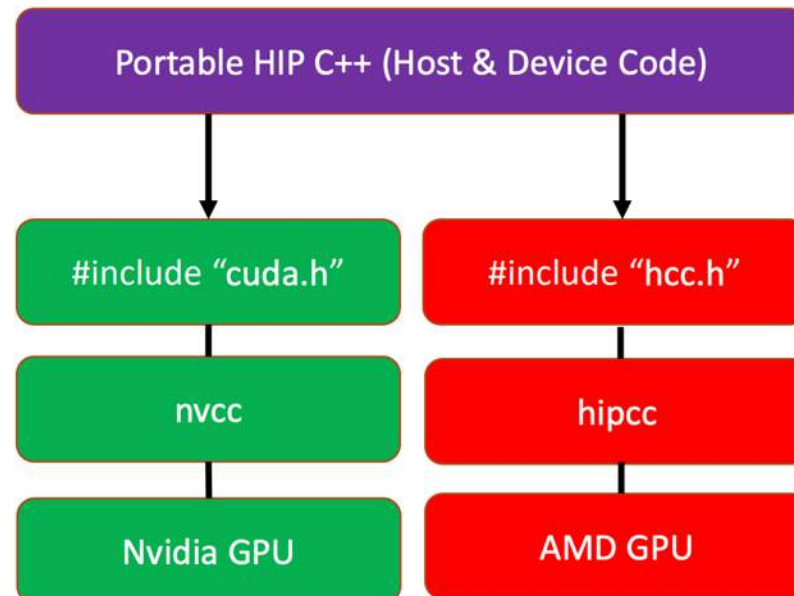
- **HIP** - Heterogeneous-compute Interface for Portability

- Owned by AMD
- A C++ runtime API and kernel language that allows developers to create portable applications that can run on AMD's accelerators as well as CUDA devices



HIP

- Is open-source
- Provides an API for an application to leverage GPU acceleration for both AMD and Nvidia devices
- Syntactically similar to CUDA. Most CUDA API calls can be converted in place: `cuda --hipify--> hip`
- Supports a strong subset of CUDA runtime functionality



HIP vs. CUDA

- Kernel declare
 - Syntactically the same
- APIs

```
cudaMalloc(&d_x, N*sizeof(double));  
  
cudaMemcpy(d_x, x, N*sizeof(double),  
            cudaMemcpyHostToDevice);  
  
cudaDeviceSynchronize();
```

```
hipMalloc(&d_x, N*sizeof(double));  
  
hipMemcpy(d_x, x, N*sizeof(double),  
           hipMemcpyHostToDevice);  
  
hipDeviceSynchronize();
```

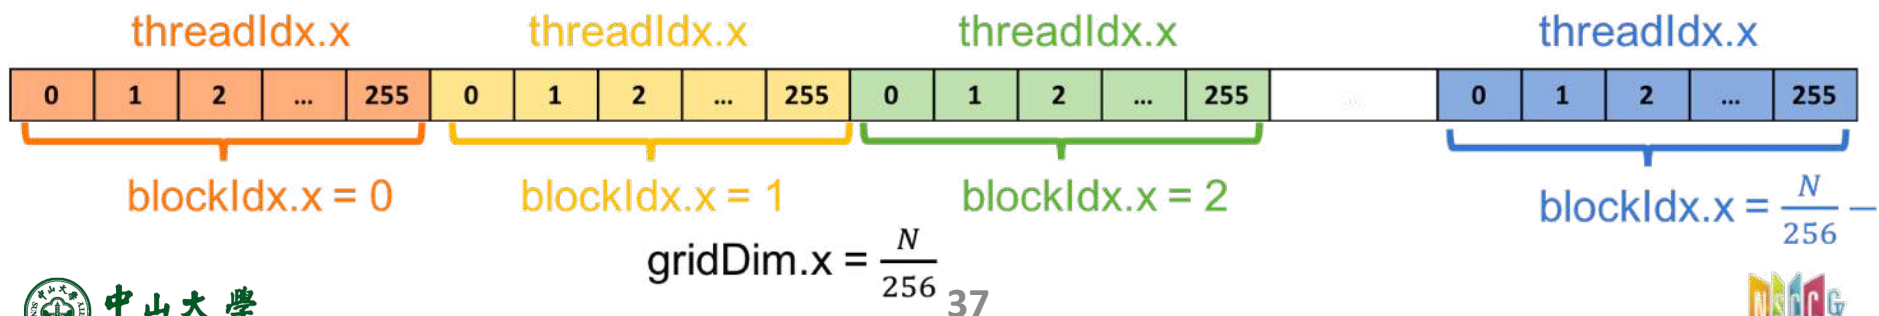
- Kernel launch

```
some_kernel<<<gridsize, blocksize,  
            shared_mem_size, stream>>>  
(arg0, arg1, ...);
```

```
hipLaunchKernelGGL(some_kernel,  
                   gridsize, blocksize,  
                   shared_mem_size, stream,  
                   arg0, arg1, ...);
```

Kernel Dimensions[维度]

- Built-in variables
 - *blockDim.x*: the size of the block (#threads in the block)
 - *gridDim.x*: the size of the grid (#blocks)
 - *blockIdx.x*: the index of the block within the grid
 - *threadIdx.x*: the index of the thread within the block
- Example: N threads in total, 256 threads per block
 - $\text{blockDim.x} = 256$
 - $\# \text{blocks} = N / 256 \rightarrow \text{gridDim.x}$
 - $\text{blockIdx.x} = [0, 1, \dots, N/256-1]$
 - $\text{threadIdx.x} = [0, 1, \dots, 255]$



Example: Kernel Declare[声明]

- A kernel is declared with the `__global__` attribute
 - Kernels should be declared `void`
 - All pointers passed to kernels must point to device memory
- All threads execute the kernel's body “simultaneously”
 - Each thread uses its unique thread and block IDs to compute a global ID

```
for (int i=0;i<N;i++) {  
    h_a[i] *= 2.0;  
}
```



```
__global__ void myKernel(int N, double *d_a) {  
    int i = threadIdx.x + blockIdx.x*blockDim.x;  
    if (i<N) {  
        d_a[i] *= 2.0;  
    }  
}
```

Example: Kernel Launch[启动]

- Kernels are launched from host

```
dim3 threads(256,1,1);           //3D dimensions of a block of threads
dim3 blocks((N+256-1)/256,1,1);  //3D dimensions the grid of blocks

hipLaunchKernelGGL(myKernel,      //Kernel name (__global__ void function)
                   blocks,        //Grid dimensions
                   threads,       //Block dimensions
                   0,             //Bytes of dynamic LDS space (see extra slides)
                   0,             //Stream (0=NULL stream)
                   N, a);         //Kernel arguments
```

- Analogous to CUDA kernel launch syntax:

```
myKernel<<<blocks, threads, 0, 0>>>(N,a);
```

Example: Memory Allocation[内存分配]

- The host instructs the device to allocate memory and records a pointer to device memory

```
int main() {  
    ...  
    int N = 1000;  
    size_t Nbytes = N*sizeof(double);  
    double *h_a = (double*) malloc(Nbytes);           //Host memory  
  
    double *d_a = NULL;  
    hipMalloc(&d_a, Nbytes);                          //Allocate Nbytes on device  
  
    ...  
  
    free(h_a);                                         //free host memory  
    hipFree(d_a);                                     //free device memory  
}
```


Example: Memory Copy[数据传输]

- The host queues memory transfers
 - hipMemcpyHostToDevice
 - hipMemcpyDeviceToHost
 - hipMemcpyDeviceToDevice

```
//copy data from host to device
```

```
hipMemcpy(d_a, h_a, Nbytes, hipMemcpyHostToDevice);
```

```
//copy data from device to host
```

```
hipMemcpy(h_a, d_a, Nbytes, hipMemcpyDeviceToHost);
```

```
//copy data from one device buffer to another
```

```
hipMemcpy(d_b, d_a, Nbytes, hipMemcpyDeviceToDevice);
```

Example: Putting Together

```
#include "hip/hip_runtime.h"

int main() {
    int N = 1000;
    size_t Nbytes = N*sizeof(double);
    double *h_a = (double*) malloc(Nbytes); //host memory
    double *d_a = NULL;
    HIP_CHECK(hipMalloc(&d_a, Nbytes));

    ...

    HIP_CHECK(hipMemcpy(d_a, h_a, Nbytes, hipMemcpyHostToDevice)); //copy data to device

    hipLaunchKernelGGL(myKernel, dim3((N+256-1)/256,1,1), dim3(256,1,1), 0, 0, N, d_a); //Launch kernel
    HIP_CHECK(hipGetLastError());

    HIP_CHECK(hipMemcpy(h_a, d_a, Nbytes, hipMemcpyDeviceToHost));

    ...

    free(h_a); //free host memory
    HIP_CHECK(hipFree(d_a)); //free device memory
}
```

```
__global__ void myKernel(int N, double *d_a) {
    int i = threadIdx.x + blockIdx.x*blockDim.x;
    if (i<N) {
        d_a[i] *= 2.0;
    }
}
```

```
#define HIP_CHECK(command) { \
    hipError_t status = command; \
    if (status!=hipSuccess) { \
        std::cerr << "Error: HIP reports " \
                    << hipGetErrorString(status) \
                    << std::endl; \
        std::abort(); } }
```

Device Management[管理]

- Host can query *number* of devices visible to system:

```
int numDevices = 0;  
hipGetDeviceCount(&numDevices);
```

- Host tells the runtime to issue instructions to a *particular* device:

```
int deviceId = 0;  
hipSetDevice(deviceID);
```

- Host can query what device is currently *selected*:

```
hipGetDevice(&deviceId);
```

- The host can also query a device's *properties*:

```
hipDeviceProp_t props;  
hipGetDeviceProperties(&props, deviceId);
```

hipDeviceProp_t is a struct that contains useful fields like the device's name, total VRAM, clock speed, and GCN architecture.