


CUDA Programming

Aiichiro Nakano

*Collaboratory for Advanced Computing & Simulations
Department of Computer Science
Department of Physics & Astronomy
Department of Quantitative & Computational Biology
University of Southern California*


Email: anakano@usc.edu

**Goal: Multithreading on graphics processing units (GPUs);
heterogenous device concept**



Graphics Processing Unit (GPU)

- **GPU:** A specialized processor that offloads 3D graphics rendering from the central processing unit (CPU).
- **GPGPU:** General-purpose computing on GPU, by using a GPU to perform computation traditionally handled by the CPU; GPU is considered as a multithreaded, massively data parallel co-processor (accelerator).
- **NVIDIA Quadro, Tesla & newer GPUs** are capable of general-purpose computing with the use of Compute Unified Device Architecture (CUDA).



Tesla A100 (6,912 CUDA cores)

CUDA

How to program GPGPU?

- **Compute Unified Device Architecture**
- **Integrated host (CPU) + device (GPU) application programming interface based on C language, developed at NVIDIA**
- **CUDA homepage**
http://www.nvidia.com/object/cuda_home.html
- **Widely used in the deep-learning community**
<https://www.deeplearningbook.org/contents/applications.html>

Using CUDA on Discovery

- **Add the following commands in .bashrc in your home directory**

```
module purge
module load usc
module load cuda/10.1.243
```

- **Compilation**

```
nvcc -o pi pi.cu
```

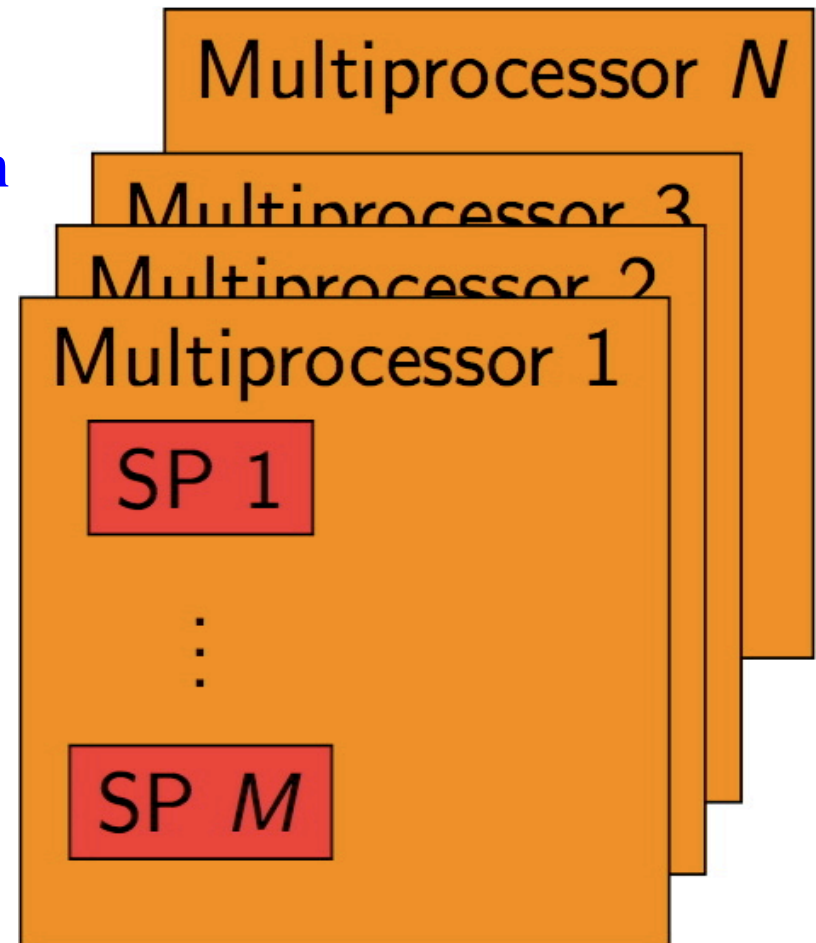
- **Submit a Slurm script**

```
#!/bin/bash
#SBATCH --nodes=1
#SBATCH --ntasks-per-node=1
#SBATCH --gres=gpu:1
#SBATCH --time=00:00:59
#SBATCH --output=pi.out
#SBATCH -A anakano_429
./pi
```

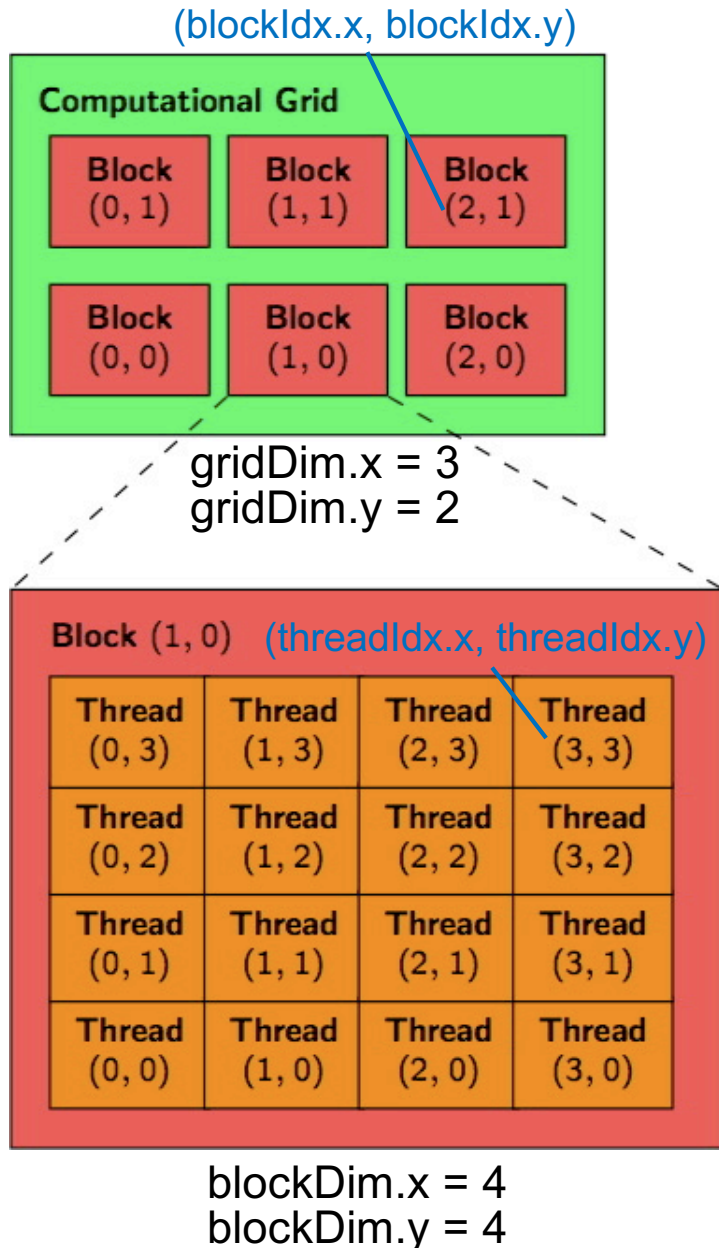
<https://aiichironakano.github.io/cs596/src/cuda/pi.cu>

Example of NVIDIA GPU at CARC

- **Host (CPU)**
 - > **Dual octacore ($2 \times 8 = 16$) Intel Xeon**
 - > **Clock rate: 2.4 GHz**
 - > **Memory: 64 GB**
- **Device (GPU): Dual NVIDIA Tesla K20m**
 - > **Number of streaming multiprocessors (SMs) per GPU: 13**
 - > **Number of cores (or streaming processors, SPs) per SM: 192**
 - > **Total number of cores: $13 \times 192 = 2496$**
 - > **Clock rate: 706 MHz**
 - > **Global memory: 5 GB**
 - > **Shared memory per SM: 48 KB**



Grid, Blocks & Threads



- **Computational grid** = a 1 or 2D grid of thread blocks (*cf.* SMs); each block = a 1, 2 or 3D array of ≤ 512 threads (*cf.* SPs); the application specifies the grid & block dimensions
 - **gridDim** provides dimension of grid; 1 or 2 element struct: “.x” & “.y”
 - **blockDim** provides dimension of block; 1, 2 or 3 element struct: “.x”, “.y” & “.z”
- All threads within a block execute the same kernel (SPMD) & cooperate *via* shared memory, atomic operations & barrier synchronization
- Each block has a unique block ID
 - **blockIdx** is 1 or 2 element struct
- Each thread has a unique ID within the block
 - **threadIdx** is a struct with up to 3 elements: “.x”, “.y” (in 2 or 3D) & “.z” (in 3D) for the innermost, intermediated & outermost index
- Each thread uses the block & thread IDs to decide what data to work on (*i.e.*, SPMD)

cf. `vproc[3]`, `vthrd[3]`, `vid[3]`, `vtd[3]` in `hmd.c`

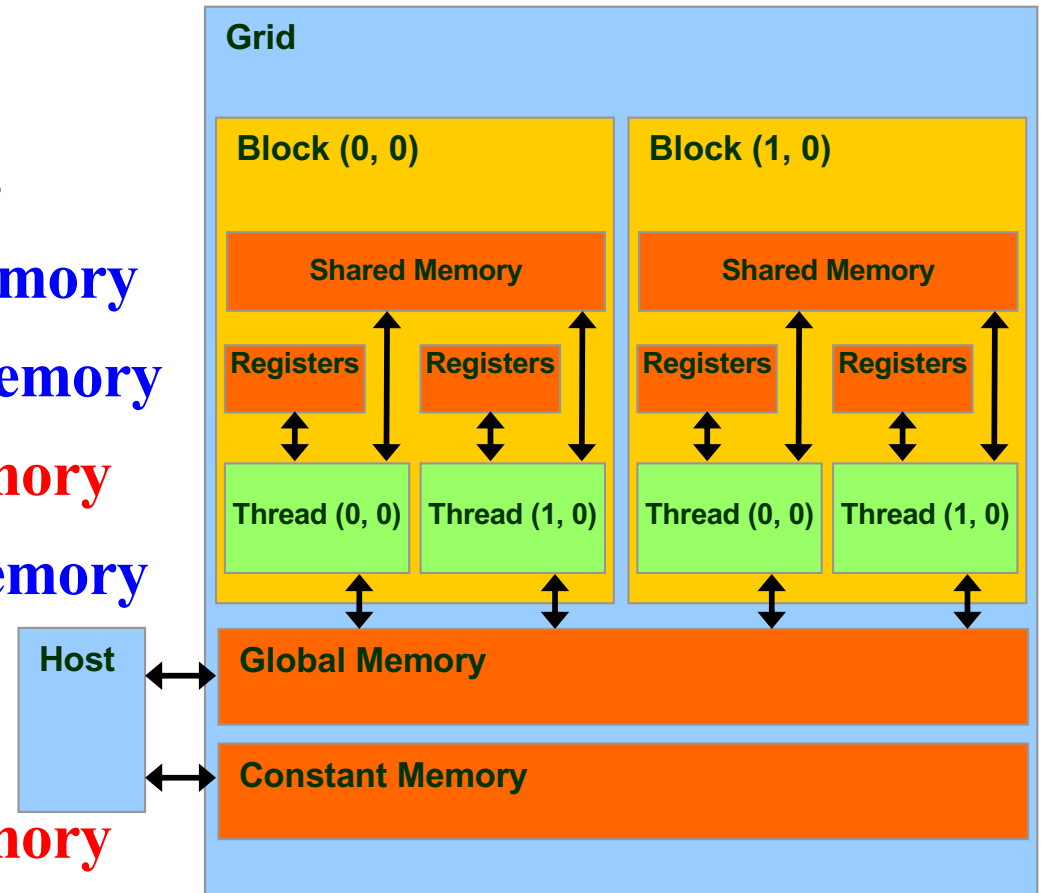
Hierarchical Device Memory

Each thread can:

- Read/write **per-thread** registers
- Read/write **per-thread** local memory
- Read/write **per-block** shared memory
- Read/write **per-grid** global memory
- Read only **per-grid** constant memory

Host code can:

- Read/write **per-grid** global memory
- Read/write **per-grid** constant memory



We will only use global device memory in assignment

Device Memory Allocation

cudaMalloc()

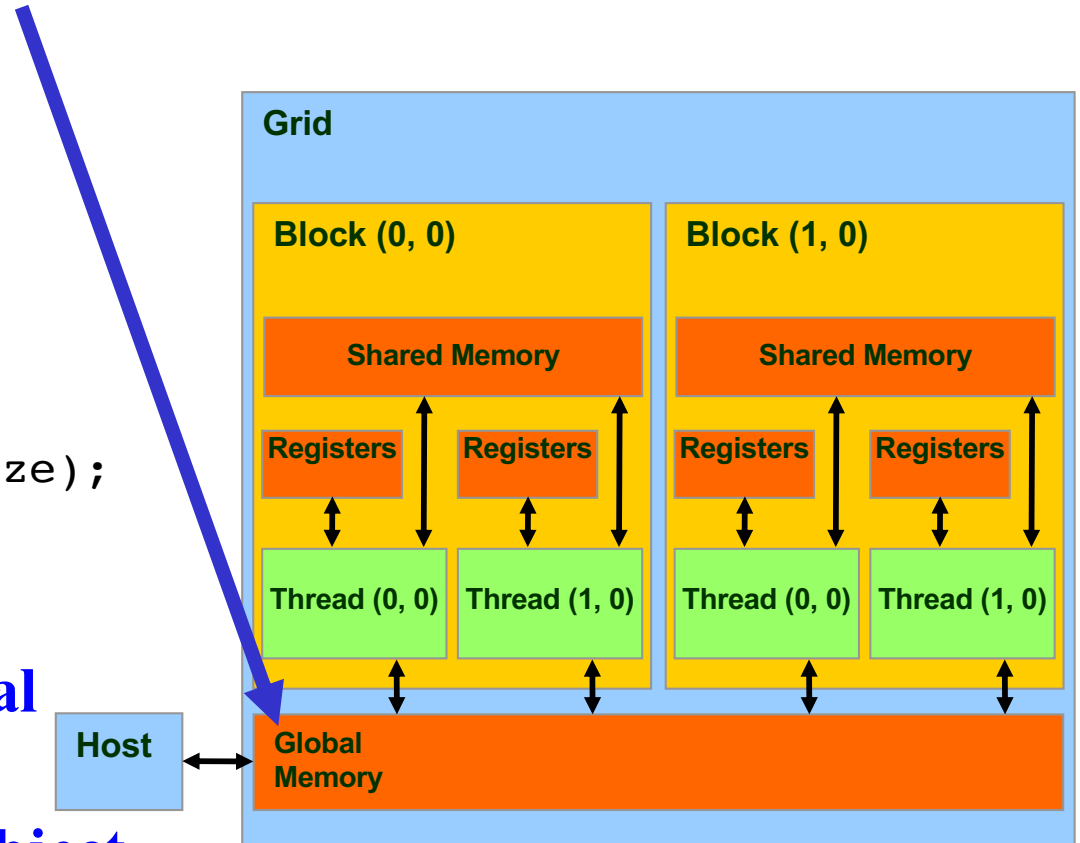
- Allocates object in the device global memory
- Requires two parameters:
 - Address of a pointer to the allocated object
 - Size of of allocated object

```
cudaMalloc((void **)&sumDev, size);
```

cudaFree()

- Frees object from device global memory
- Parameter: Pointer to freed object

```
cudaFree(sumDev);
```



Host-Device Data Transfer

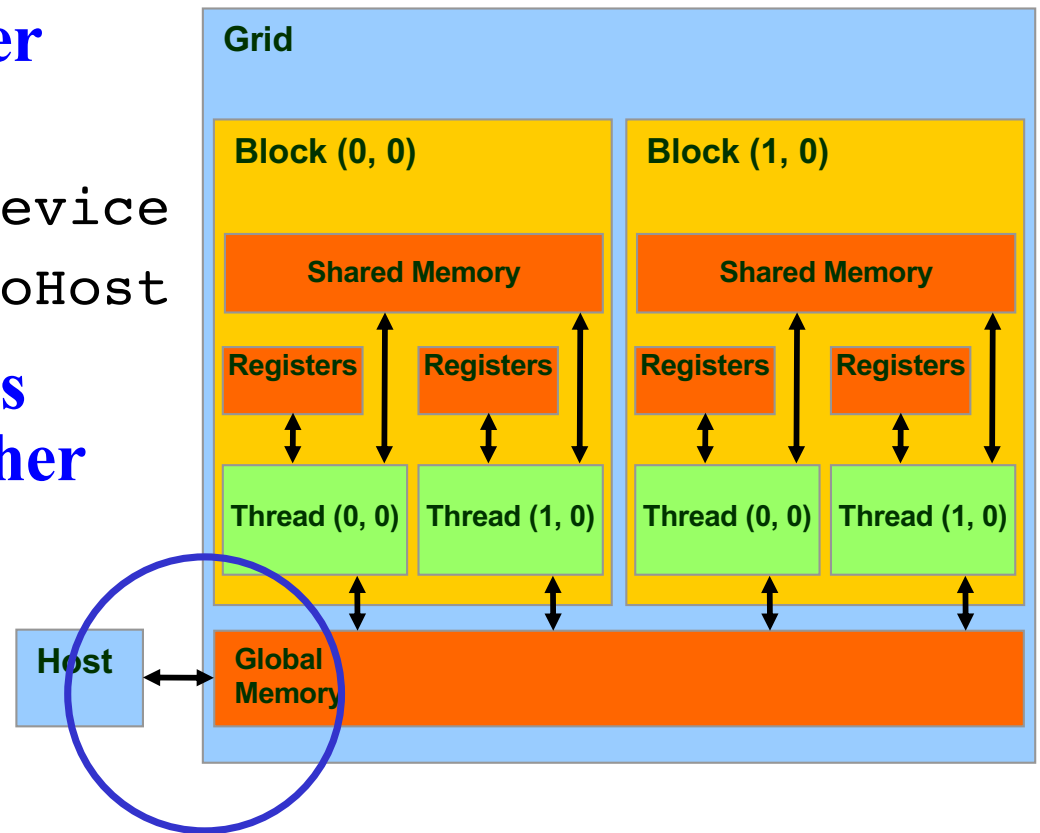
`cudaMemcpy(dest, src, size, cmd)`

- Arguments

- `dest` = pointer to array to receive data
- `src` = pointer to array to source data
- `size` = # of bytes to transfer
- `cmd` = transfer direction

- > `cudaMemcpyHostToDevice`
- > `cudaMemcpyDeviceToHost`

- Transfer specified # of bytes from one memory to the other in direction specified



```
cudaMemcpy(sumHost, sumDev, size, cudaMemcpyDeviceToHost);
```

Kernel Program for Device

- Set of threads triggered by invocation of a single kernel

- **Definition** 

```
__global__ void kernel_fun(argument_list)
```

Kernel that can be called from a host function

- **Invocation**

```
kernel_fun <<<execution configuration>>> (operands)
```

– Range specifies set of values for thread grid

```
host_fun() {  
    ...  
    dim3 dimGrid(4, 2, 1)  
    dim3 dimBlock(2, 2, 2)  
    kernel_fun <<<dimGrid, dimBlock>>> (operands)  
    ...  
}
```

4×2 grid (3rd dimension not used)

2×2×2 block

3-element struct accessed by dimGrid.x, dimGrid.y, dimGrid.z

Built-in Variables

- `dim3 gridDim;`

Dimensions of the grid in blocks (`gridDim.z` **unused**)

- `dim3 blockDim;`

Dimensions of the block in threads

cf. `vproc[3]` & `vthrd[3]` in `hmd.c`

- `dim3 blockIdx;`

Block index within the grid

- `dim3 threadIdx;`

Thread index within the block

cf. `vid[3]` & `vtd[3]` in `hmd.c`

Calculate Pi with CUDA: pi.cu (1)

```
// Using CUDA device to calculate pi
#include <stdio.h>
#include <cuda.h>

#define NBIN 10000000 // Number of bins
#define NUM_BLOCK 13 // Number of thread blocks
#define NUM_THREAD 192 // Number of threads per block
int tid;
float pi = 0;


// Kernel that executes on the CUDA device
__global__ void cal_pi(float *sum, int nbin, float step, int nthreads, int nblocks) {
    int i;
    float x;
    int idx = blockIdx.x*blockDim.x+threadIdx.x; // Sequential thread index across blocks
    for (i=idx; i< nbin; i+=nthreads*nblocks) { // Interleaved bin assignment to threads
        x = (i+0.5)*step;
        sum[idx] += 4.0/(1.0+x*x); // Data privatization
    }
}
```

blockIdx.x:	0					1						2
threadIdx.x:	0	1	2	...	191	0	...	191	0	...		
idx:	0	1	2	...	191	192	...	383	384	...		

1D grid & block

gridDim.x|y = 13|1
blockDim.x|y|z = 192|1|1

Total number of threads = 13×192 = 2,496



Calculate Pi with CUDA: pi.cu (2)

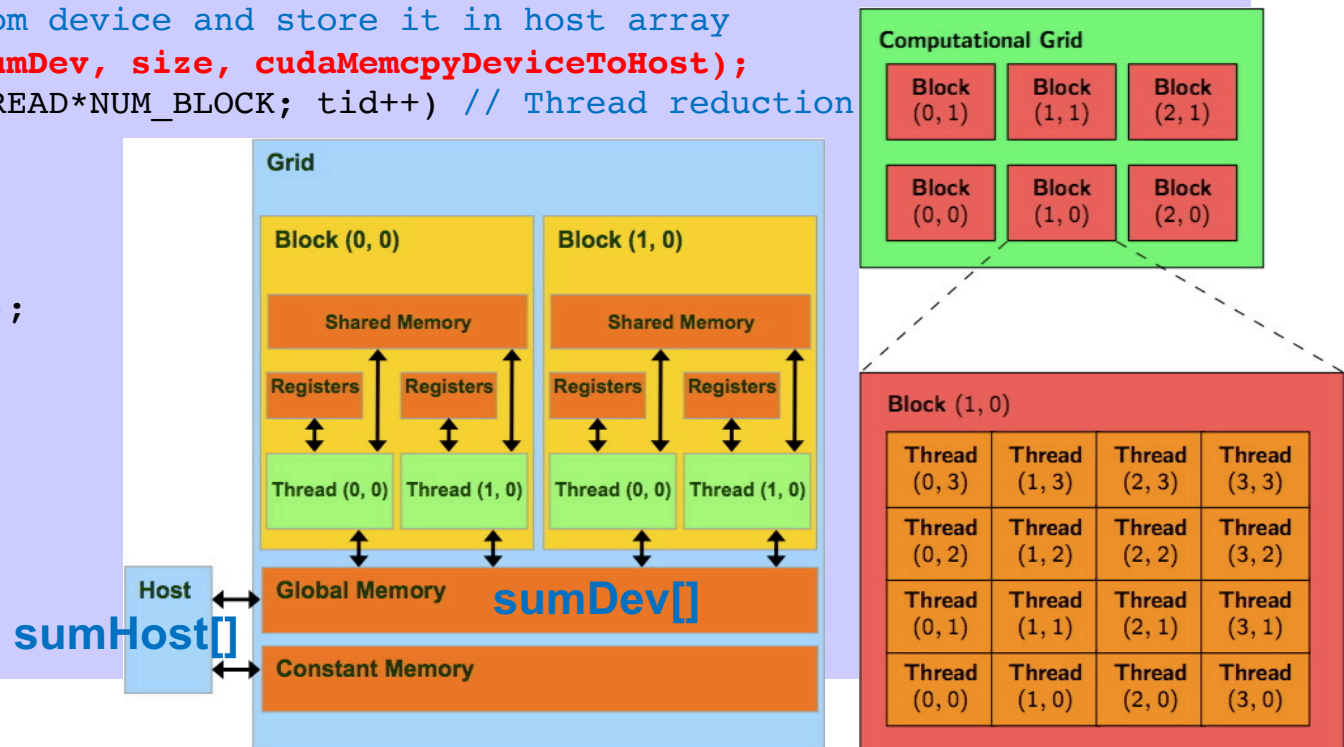
```
// Main routine that executes on the host
int main(void) {
    dim3 dimGrid(13, 1, 1); // Grid dimensions
    dim3 dimBlock(192, 1, 1); // Block dimensions
    float *sumHost, *sumDev; // Pointer to host & device arrays

    float step = 1.0/NBIN; // Step size
    size_t size = NUM_BLOCK*NUM_THREAD*sizeof(float); //Array memory size
    sumHost = (float *)malloc(size); // Allocate array on host
    cudaMalloc((void **) &sumDev, size); // Allocate array on device
    // Initialize array in device to 0
    cudaMemset(sumDev, 0, size);
    // Do calculation on device by calling CUDA kernel
    cal_pi <<<dimGrid, dimBlock>>> (sumDev, NBIN, step, NUM_THREAD, NUM_BLOCK);
    // Retrieve result from device and store it in host array
    cudaMemcpy(sumHost, sumDev, size, cudaMemcpyDeviceToHost);
    for(tid=0; tid<NUM_THREAD*NUM_BLOCK; tid++) // Thread reduction
        pi += sumHost[tid];
    pi *= step;

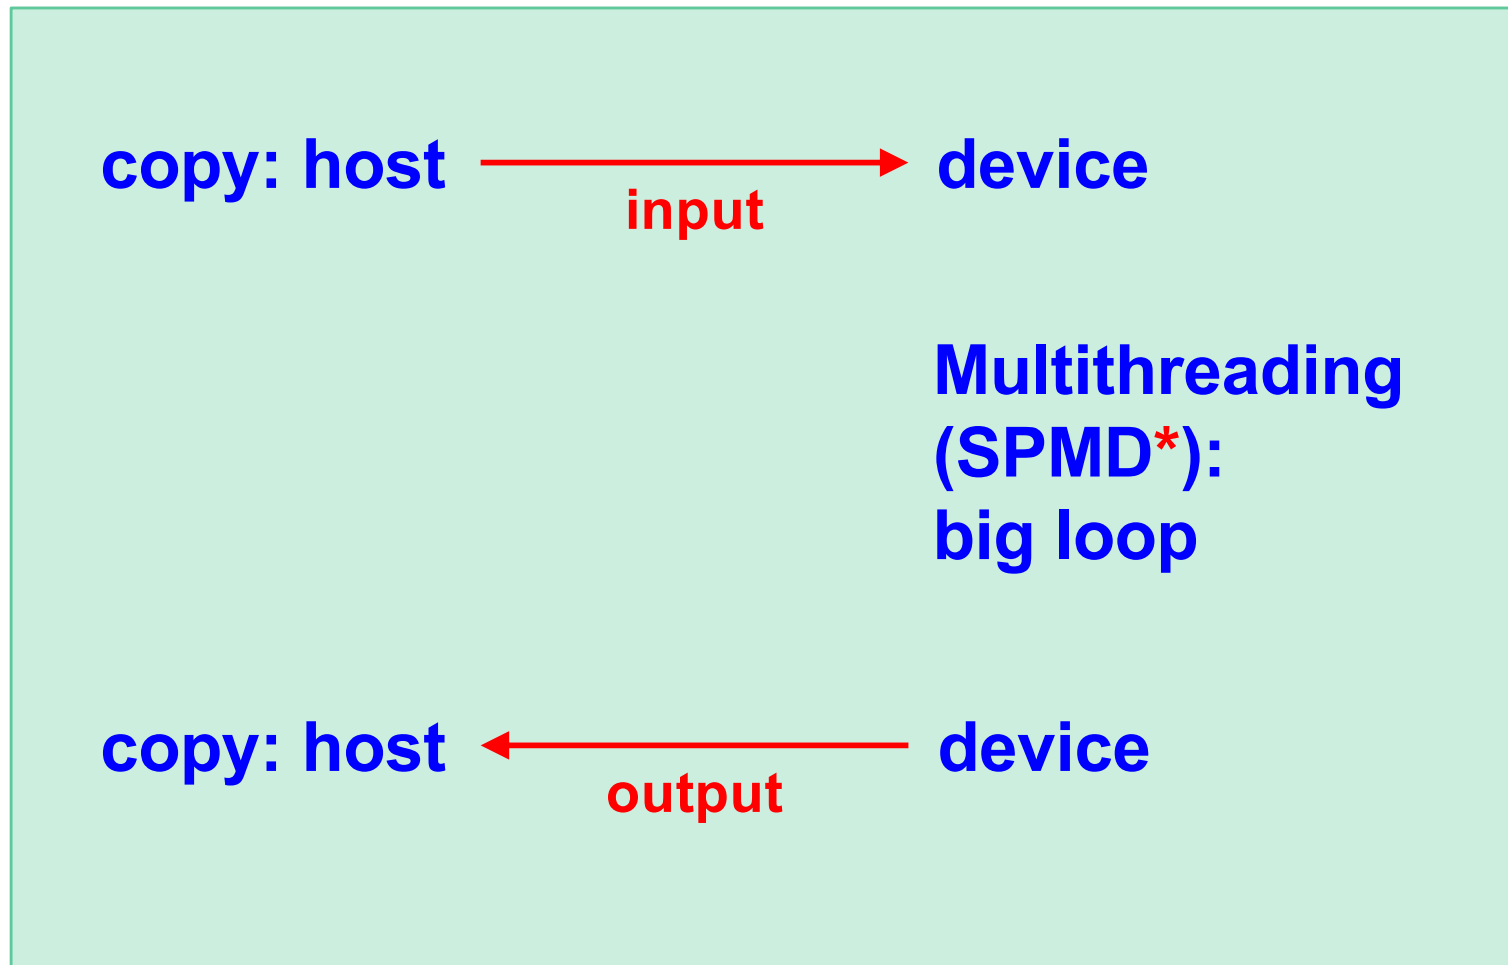
    // Print results
    printf("PI = %f\n",pi);

    // Cleanup
    free(sumHost);
    cudaFree (sumDev);

    return 0;
}
```






Summary: CUDA Computing



- * Single program multiple data we have learned is called **single instruction multiple threads (SIMT)** in GPU terminology

New Generations of GPUs



- Running time per molecular dynamics (MD) step on Kepler (K20), Pascal (P100) & Volta (V100) GPUs



New Generations of GPUs (2)


- **A100 has arrived at CARC**

UNIFIED AI ACCELERATION




All results are measured
 BERT Large Training (FP32 & FP16) measures Pre-Training phase, uses PyTorch including (2/3) Phase1 with Seq Len 128 and (1/3) Phase 2 with Seq Len 512,
 V100 is DGX1 Server with 8xV100, A100 is DGX A100 Server with 8xA100, A100 uses TF32 Tensor Core for FP32 training
 BERT Large Inference uses TRT 7.1 for T4/V100, with INT8/FP16 at batch size 256. Pre-production TRT for A100, uses batch size 94 and INT8 with sparsity

ACCELERATING HPC



All results are measured
 Except BerkeleyGW, V100 used is single V100 SXM2, A100 used is single A100 SXM4
 More apps detail: AMBER based on PME-Cellulose, GROMACS with STMV (h-bond), LAMMPS with Atomic Fluid LJ-2.5, NAMD with v3.0a1 STMV_NVE
 Chroma with szscI21_24_128, FUN3D with dpw, RTM with Isotropic Radius 4 1024³, SPECFEM3D with Cartesian four material model
 BerkeleyGW based on ChI Sum and uses 8xV100 in DGX-1, vs 8xA100 in DGX A100

BERT: Bidirectional Encoder Representations from Transformers used in natural language processing (NLP)



cf. [Pytorch GPU engine](#)

Warp & Control Divergence

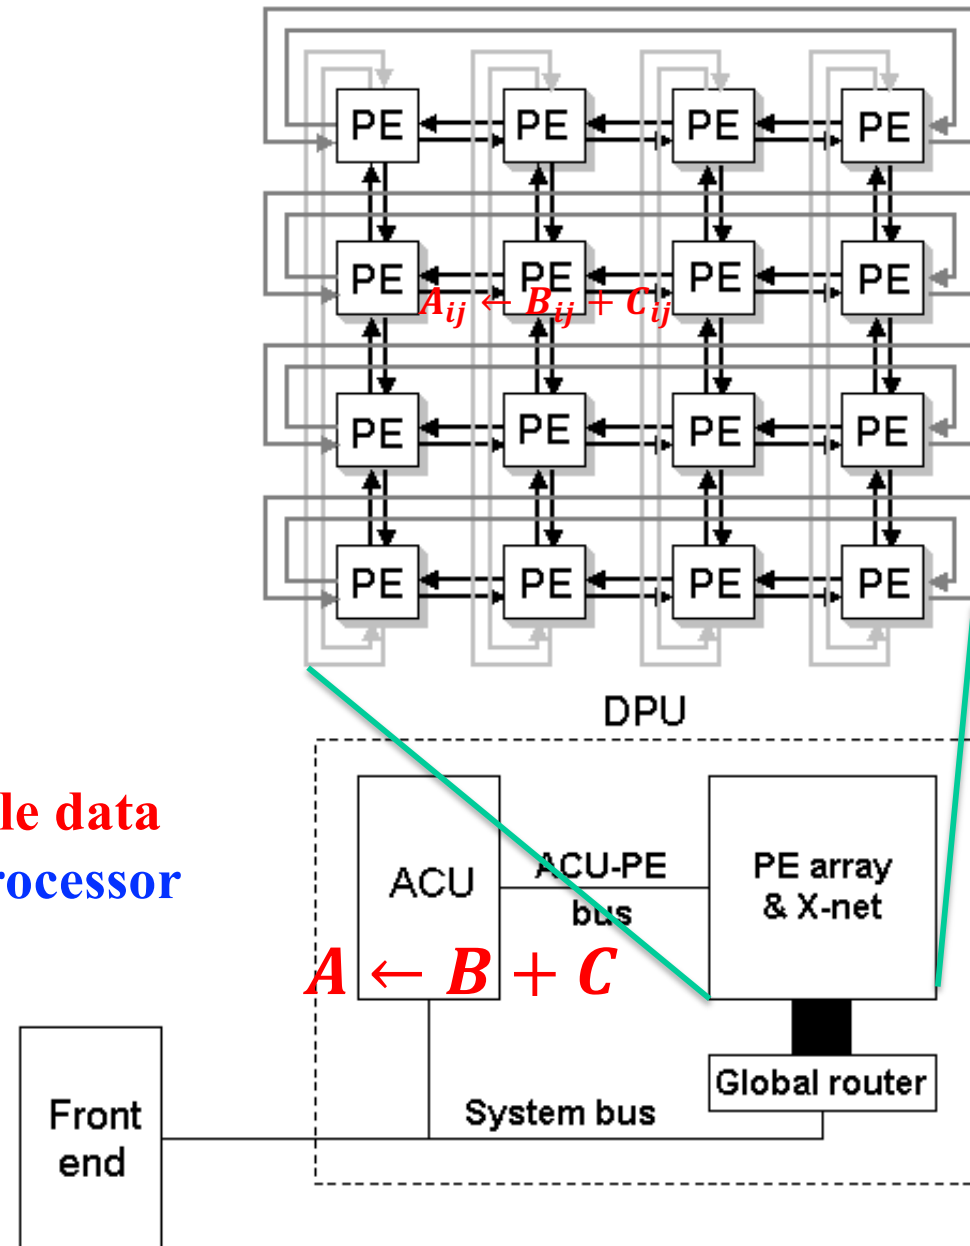
- Threads in a block are subdivided into **warps** (*e.g.* consisting of 32 threads)
- Warps are executed in SIMD (single-instruction multiple-data) fashion, *i.e.*, multiple threads concurrently perform the same operation
- CUDA provides warp-level primitives for efficient warp-level programming
- Single instruction multiple thread (SIMT) execution model penalizes **control divergence**, where different threads execute different instructions
- **Warp voting**: All threads (*e.g.* particles) within a warp vote on which computation to perform, with an overhead of unnecessary computations, for example:
`if (any thread in a warp wants to compute) all threads do`

Massive SIMD Data-Parallel Accelerator



SIMD: single-instruction multiple data
Quantum dynamics on 8,192-processor
(128 × 64) MasPar 1208B

Nakano,
Comput. Phys. Commun.
83, 181 ('94)



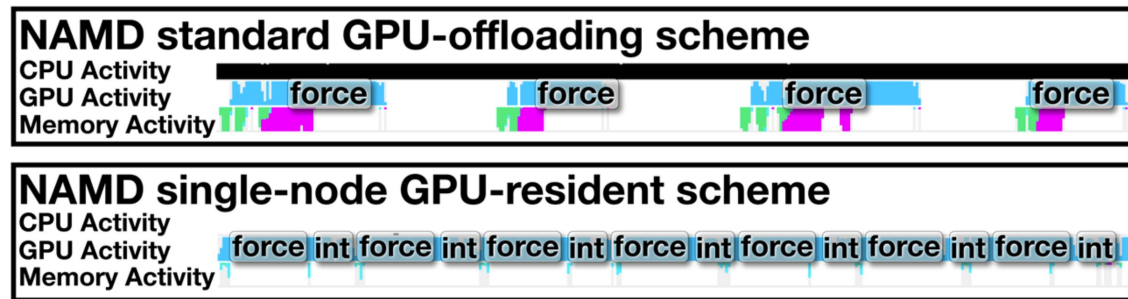
See lecture on [pre-Beowulf parallel computing](#)

Final Projects on GPU

- L. Peng *et al.*, “Parallel lattice Boltzmann flow simulation on emerging multi-core platforms,” *Proc. Euro-Par*, 763 ('08)
- P. E. Small *et al.*, “Acceleration of dynamic n -tuple computations in many-body molecular dynamics,” *Proc. IEEE HPC Asia* ('18)
- Sasan Tavakkol’s final project became a poster in GPU Technology Conference (see nice videos 1 & 2)
- C. Rizzo *et al.*, “PAR2: parallel random walk particle tracking method for solute transport in porous media,” *Comput. Phys. Commun.* **239**, 265 ('19)

Final Project on GPU-MD?

- J. C. Phillips *et al.*, “Quantum-based molecular dynamics simulations using tensor cores,” *J. Chem. Phys* **153**, 044130 ('20)



Persistent GPU kernel

FIG. 5. Standard GPU offload approach compared against new GPU-resident execution scheme for a single-node NAMD simulation of apolipoprotein 1 (ApoA1) in water, consisting of 92 224 atoms. The light blue line tracks GPU activity, while the black strip tracks CPU activity. GPU force calculations are labeled “force,” and GPU integration calculations are labeled “int.”

- S. Pall *et al.*, “Heterogeneous parallelization and acceleration of molecular dynamics simulations in GROMACS,” *J. Chem. Phys.* **153**, 134110 ('20)


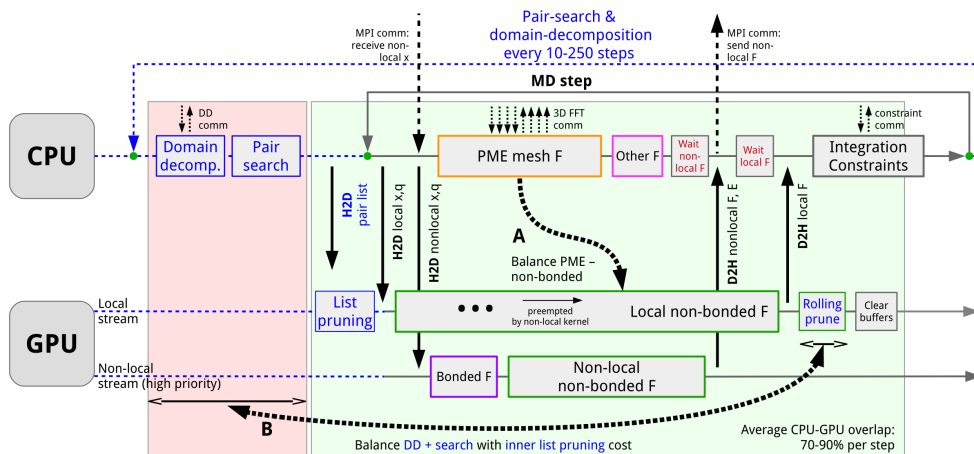



FIG. 4. Cluster pair setups with four particles ($N = 4$ and $M = 4$). Left panel: CPU/SIMD-centric setup. All clusters with solid lines are included in the pair list of cluster i_1 (green). Clusters with filled circles have interactions within the buffered cutoff (green dashed line) of at least one particle in i_1 , while particles in clusters intersected by the buffered cutoff that fall outside of it represent an extra implicit buffer. Right panel: hierarchical super-clusters on GPUs. Clusters i_1 – i_4 (green, magenta, red, and blue) are grouped into a super-cluster. Dashed lines represent buffered cutoffs of each i -cluster. Clusters with any particle in any region will be included in the common pair list. Particles of j -clusters in the joint list are illustrated by discs filled in black to gray; black indicates clusters that interact with all four i -clusters, while lighter gray shading indicates that a cluster only interacts with 1–3 i -cluster(s), e.g., j_m only with i_4 .

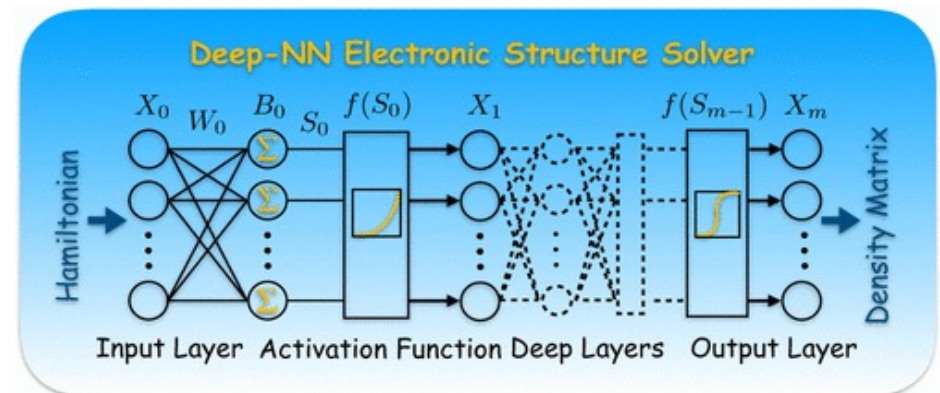
Thread blocking

Final Project on GPU-MD?


- J. Finkelstein *et al.*, “Quantum-based molecular dynamics simulations using tensor cores,” *J. Chem. Theo. Comput.* doi: 10.1021/acs.jctc.1c00726 ('21); Python code for an associated paper is available at https://pubs.acs.org/doi/suppl/10.1021/acs.jctc.1c00057/suppl_file/ct1c00057_si_001.zip



“computational structure naturally takes advantage of the exceptional processing power of the tensor cores (utilizing FP16) and allows for high performance in excess of 100 Tflops on a single Nvidia A100 GPU.”



Aurora: Heterogeneous Future



One of the first
exaflop/s
computers
in the U.S.

Aurora's compute nodes will be equipped with two Intel Xeon Scalable processors and six general-purpose GPUs based on Intel's X^e architecture.


Image: Intel Corporation

GPU Architecture

X^e arch-based "Ponte Vecchio"
GPUtile-based, chipllets, HBM
stack, Foveros 3D integration, 7nm

On-Node Interconnect

CPU-GPU: PCIe
GPU-GPU: X^e Link



Where to Go from Here

- **CUDA is a proprietary language for NVIDIA GPUs**
- **Several open languages are available**
 - > **High-level, directive-based languages**
 - OpenACC:** <https://www.openacc.org>
 - OpenMP 4.5 and later:** <https://www.openmp.org/specifications>
 - > **Low-level, comprehensive languages**
 - OpenCL:** <https://www.khronos.org/opencv>
 - DPC++:** <https://software.intel.com/content/www/us/en/develop/tools/oneapi.html>