

# **GPGPU and Manycore Programming**

**M. Sato**

**University of Tsukuba**

# Overview

---

- **Why GPU**
- **CUDA**
- **OpenCL**
- **OpenACC**
  
- **MIC (Manycore**
- **(OpenMP)**
- **OpenMP 4.0 (offloading)**

## reference

---


- **NVIDIAのCUDAの情報 Learn More about CUDA - NVIDIA**
  - [http://www.nvidia.co.jp/object/cuda\\_education\\_jp.html](http://www.nvidia.co.jp/object/cuda_education_jp.html)
  - 正式なマニュアルは、NVIDIA CUDA programming Guide
- **わかりやすいCUDAのスライド**
  - <http://www.sintef.no/upload/IKT/9011/SimOslo/eVITA/2008/seland.pdf>
- **CUDAのコード例**
  - <http://tech.ckme.co.jp/cuda.shtml>
- **OpenCL NVIDIAのページ**
  - [http://www.nvidia.co.jp/object/cuda\\_opengl\\_jp.html](http://www.nvidia.co.jp/object/cuda_opengl_jp.html)
- **後藤弘茂のWeekly海外ニュース**
  - スケーラブルに展開するNVIDIAのG80アーキテクチャ (2007年4月16日) <http://pc.watch.impress.co.jp/docs/2007/0416/kaigai350.htm>
  - KhronosがGDCでGPUやCell B.E.をサポートするOpenCLのデモを公開 (2009年3月30日) <http://pc.watch.impress.co.jp/docs/2009/0330/kaigai497.htm>

# GPU Computing

---


- **GPGPU - General-Purpose Graphic Processing Unit**
  - A technology to make use of GPU for general-purpose computing (scientific applications)
- **CUDA (Compute Unified Device Architecture)**
  - Co-designed Hardware and Software to exploit computing power of NVIDIA GPU for GP computing.
  - (In other words), at the moment, in order to obtain full performance of GPGPU, a program must be written in CUDA language.
- It is attracting many people's interest since GPU enables great performance much more than that of CPU (even multi-core) in some scientific fields.
- **Why GPGPU now? — — price (cost-performance)!!!**

# Applications (From NVIDIA's slides)




146X

容積測定時の白質連結のインタラクティブな視覚化




17X

Matlabでの等方性乱流シミュレーション




100X

天体物理学におけるN体計算



110X

分子動力学におけるイオン配置



149X


スワップションのあるLIBORモデルの金融シミュレーション



30X

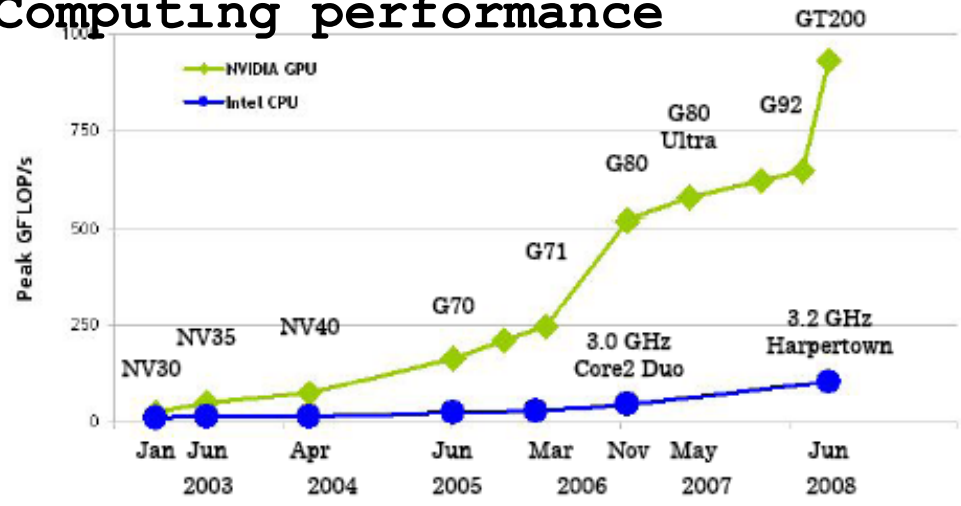
類似タンパク質および遺伝子配列検索の厳密なCmatch文字列照合

# CPU vs. GPU




Connected via PCIexpress

Computing performance



GT200 = GeForce GTX 280	G71 = GeForce 7900 GTX	NV35 = GeForce FX 5950 Ultra
G92 = GeForce 9800 GTX	G70 = GeForce 7800 GTX	NV30 = GeForce FX 5800
G80 = GeForce 8800 GTX	NV40 = GeForce 6800 Ultra	

Memory bandwidth



## NVIDIA GPGPU's architecture

### Many multiprocessor in a chip

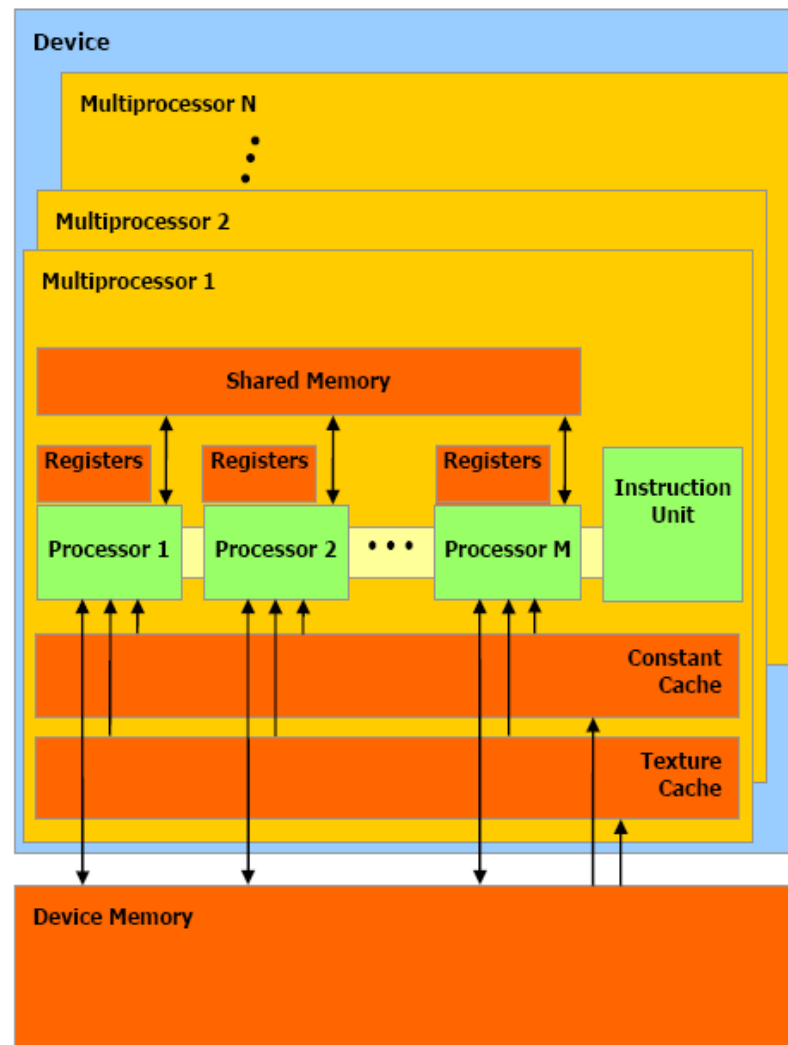
- eight Scalar Processor (SP) cores,
- two special function units for transcendentals
- a multithreaded instruction unit
- on-chip shared Memory

### SIMT (single-instruction, multiple-thread).

- The multiprocessor maps each thread to one scalar processor core, and each scalar thread executes independently with its own instruction address and register state.
- creates, manages, schedules, and executes threads in groups of 32 parallel threads called warps.

### Complex memory hierarchy

- Device Memory (Global Memory)
- Shared Memory
- Constant Cache
- Texture Cache



# CUDA (Compute Unified Device Architecture)


---

- **C programming language on GPUs**
- **Requires no knowledge of graphics APIs or GPU programming**
- **Access to native instructions and memory**
- **Easy to get started and to get real performance benefit**
- **Designed and developed by NVIDIA**
- **Requires an NVIDIA GPU (GeForce 8xxx/Tesla/Quadro)**
- **Stable, available (for free), documented and supported**
- **For both Windows and Linux**



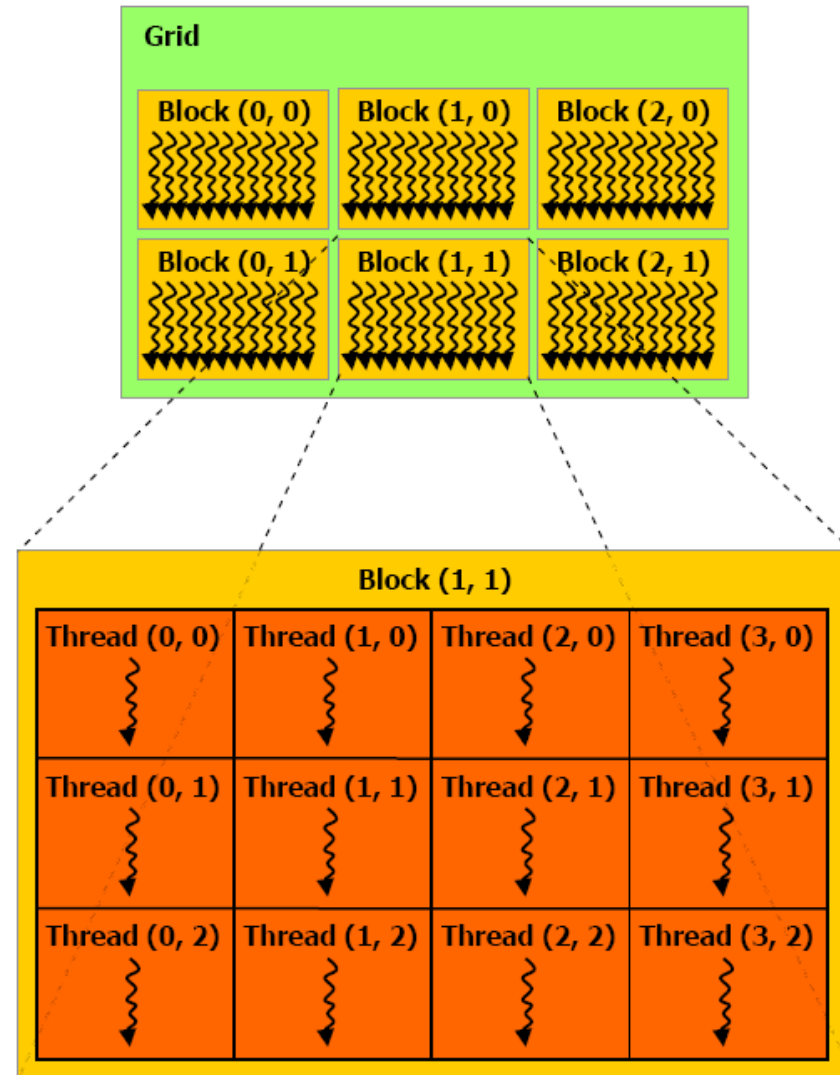
## CUDA Programming model (1/2)

- GPU is programmed as a compute device working as co-processor from CPU(host).
  - Codes for data-parallel, compute intensive part are offloaded as functions to the device
  - Offload hot-spot in the program which is frequently executed on the same data
    - For example, data-parallel loop on the same data
  - Call “kernel” a code of the function compiled as a function for the device
  - Kernel is executed by multiple threads of device.
    - Only one kernel is executed on the device at a time.
- Host (CPU) and device(GPU) has its owns memory, host memory and device memory
- Data is copied between both memory.



# CUDA Programming model (2/2)

- computational Grid is composed of multiple thread blocks
- thread block includes multiple threads
- Each thread executes kernel
  - A function executed by each thread called “kernel”
  - Kernel can be thought as one iteration in parallel loop
- computational Grid and block can have 1,2,3 dimension
- The reserved variable, `blockID` and `threadID` have ID of threads.




## Example: Element-wise Matrix Add

```
void add_matrix
( float* a, float* b, float* c, int N ) {
  int index;
  for ( int i = 0; i < N; ++i )
  for ( int j = 0; j < N; ++j ) {
    index = i + j*N;
    c[index] = a[index] + b[index];
  }
}

int main() {
  add_matrix( a, b, c, N );
}
```

### CPU program

The nested for-loops are replaced with an implicit grid



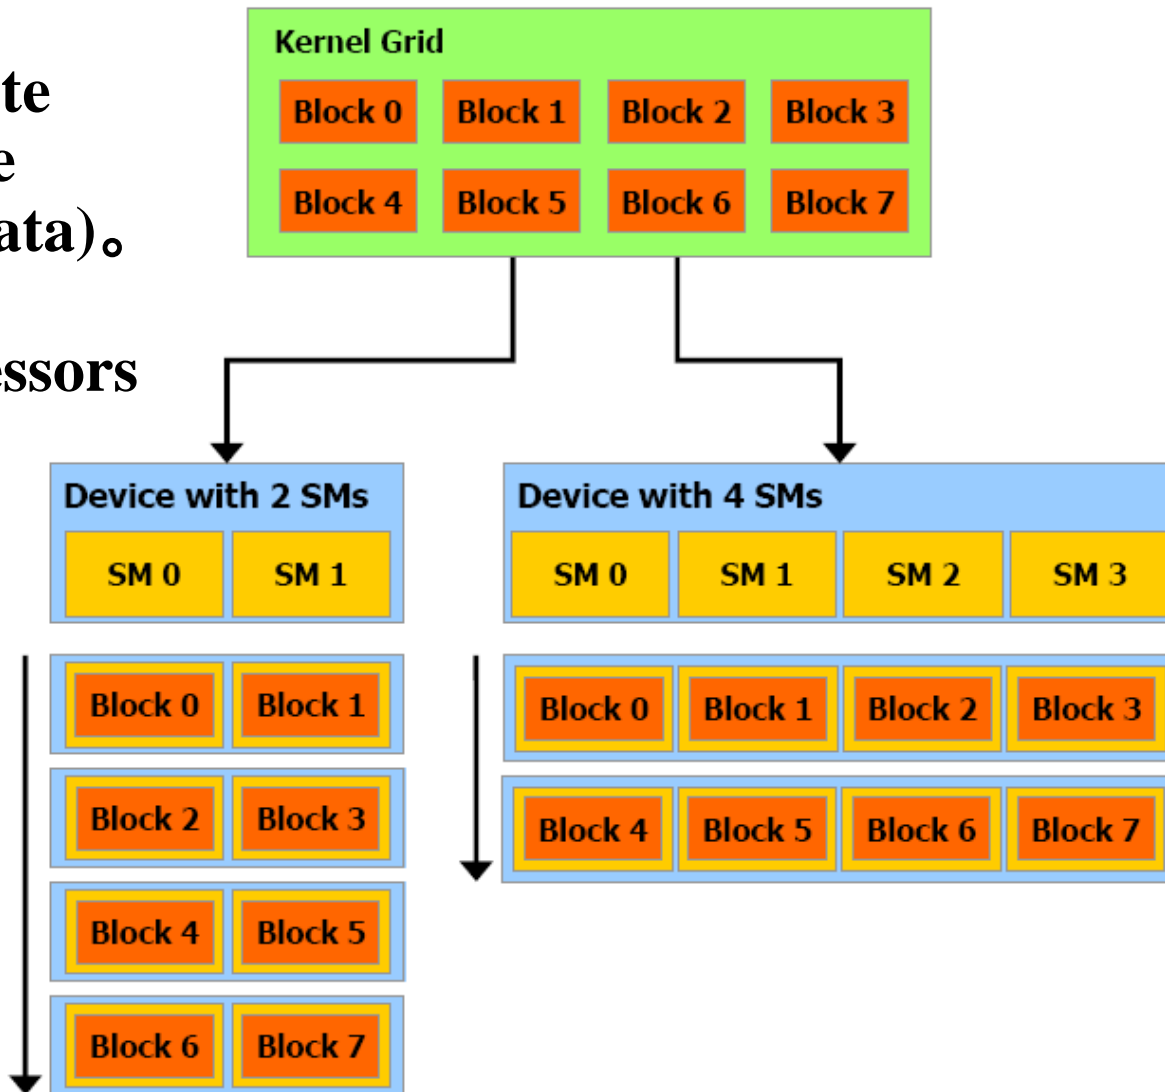
### CUDA program

```
__global__ add_matrix
( float* a, float* b, float* c, int N ) {
  int i = blockIdx.x * blockDim.x + threadIdx.x;
  int j = blockIdx.y * blockDim.y + threadIdx.y;
  int index = i + j*N;
  if ( i < N && j < N )
    c[index] = a[index] + b[index];
}

int main() {
  dim3 dimBlock( blocksize, blocksize );
  dim3 dimGrid( N/dimBlock.x, N/dimBlock.y );
  add_matrix<<<dimGrid, dimBlock>>>( a, b, c, N );
}
```

# How to be executed

- SM (Streaming Multiprocessor) execute blocks in SIMD (single instruction/multiple data).
- SM consists of 8 processors




# An example of GPGPU configuration



## 10シリーズアーキテクチャ

- 240個のスレッドプロセッサがカーネルスレッドを処理
- 30個のマルチプロセッサ、それぞれが次のユニットを内蔵
  - 8個のスレッドプロセッサ
  - 1個の倍精度ユニット
  - スレッド協調のための共有メモリ



# Lecture on Programming Environment



	Number of Multiprocessors (1 Multiprocessor = 8 Processors)	Compute Capability
GeForce GTX 295	2x30	1.3
GeForce GTX 285, GTX 280	30	1.3
GeForce GTX 260	24	
GeForce 9800 GX2	2x16	
GeForce GTS 250, GTS 150, 9800 GTX, 9800 GTX+, 8800 GTS 512	16	
GeForce 8800 Ultra, 8800 GTX	16	
GeForce 9800 GT, 8800 GT, GTX 280M, 9800M GTX	14	
GeForce GT 130, 9600 GSO, 8800 GS, 8800M GTX, GTX 260M, 9800M GT	12	

**Tesla C1060**  
 コア数: 240コア  
 プロセッサ周波数: 1.3GHz  
 搭載メモリ: 4GB  
 単精度浮動小数点演算性能: 933GFlops (ピーク)  
 倍精度浮動小数点演算性能: 78GFlops (ピーク)  
 メモリ帯域: 102GB/sec  
 標準電力消費量: 187.8W  
 浮動小数点演算: IEEE 754 単精度/倍精度  
 ホスト接続: PCI Express x16 (PCI-E2.0対応)

Tesla S1070	4x30	1.3
<b>Tesla C1060</b>	30	1.3
Tesla S870	4x16	1.0
Tesla D870	2x16	1.0
Tesla C870	16	1.0
Quadro Plex 2200 D2	2x30	1.3
Quadro Plex 2100 D4	4x14	1.1
Quadro Plex 2100 Model S4	4x16	1.0

# Invoke (Launching) Kernel

---

- Host processor invoke the execution of kernel in this form similar to function call:

```
kernel<<<dim3 grid, dim3 block, shmem_size>>>(...)
```

- Execution Configuration ( “<<< >>>”)
  - Dimension of computational grid : x and y
  - Dimension of thread block: x、 y、 z

```
dim3 grid(16 16);  
dim3 block(16,16);  
kernel<<<grid, block>>>(...);  
kernel<<<32, 512>>>(...);
```

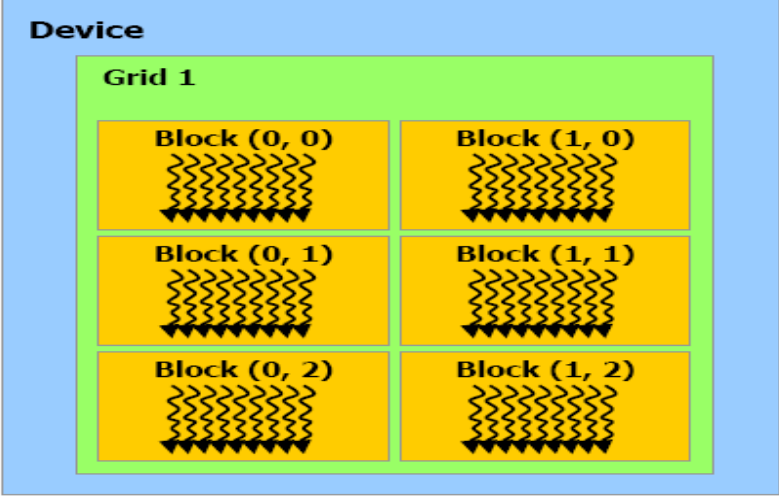


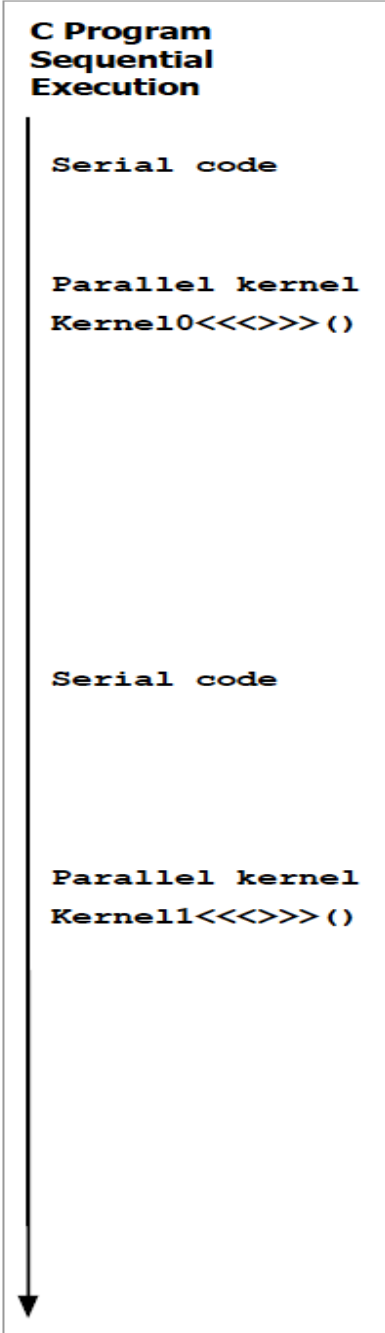
# CUDA kernel and thread

---

- **Parallel part of applications are executed as a kernel of CUDA on the device**
  - One kernel is executed at a time
  - Many threads execute kernel function in parallel.
- **Difference between CUDA thread and CPU thread**
  - **CUDA thread is a very light-weight thread**
    - Overhead of thread creation is very small
    - Thread switching is also very fast since it is supported by hardware.
  - **CUDA exploit its performance and efficient execution by a thousands of threads.**
    - Conventional Multicore supports only a few threads (by software)

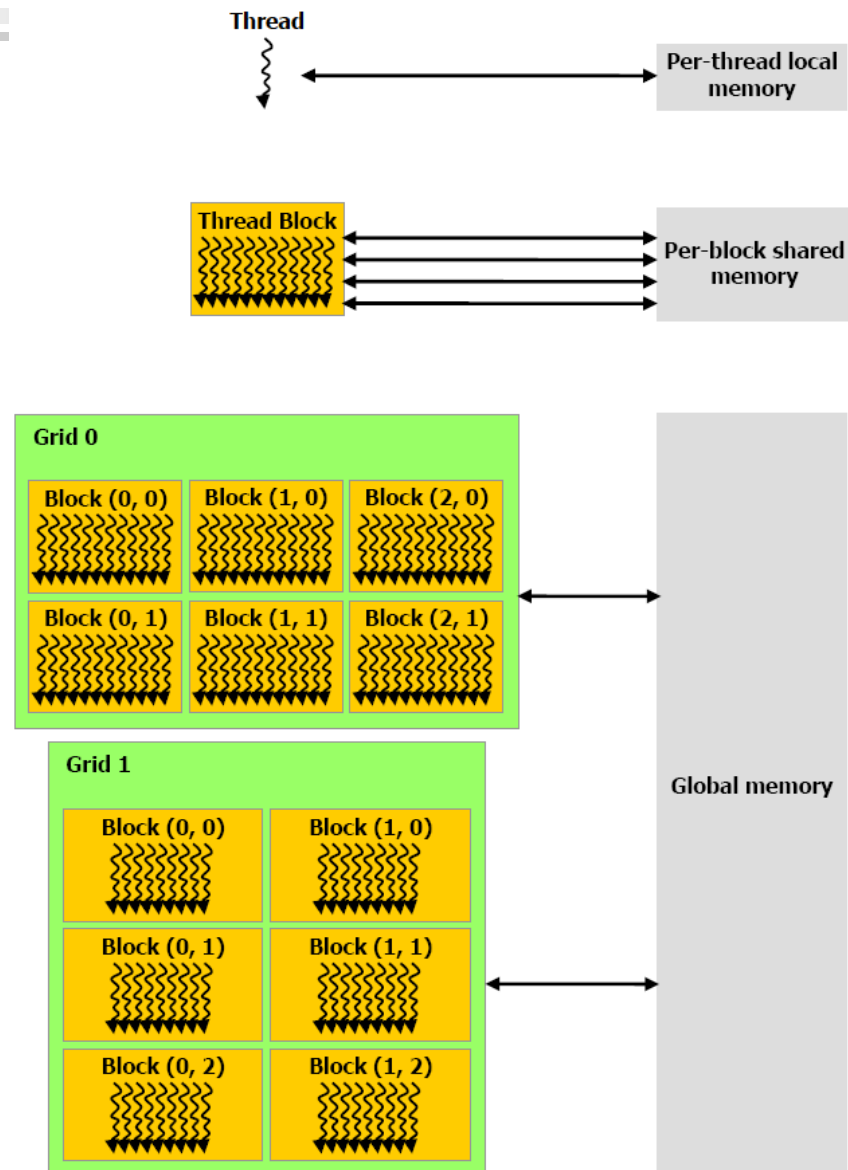


## Execution of CPU Code and Kernel code by Device



## Grid, Block, thread and Memory hierarchy

- Thread can access local memory (per-thread)
- Thread can access “shared memory” on chip, which is attached for each thread block (SM).
- Thread in Computational Grid access and share a global memory.



# Memory management (1/2)

---

- CPU and GPU have different memory space.
- Hosts (CPU) manages device (GPU) memory
- **Allocation and Deallocation of GPU memory**
  - `cudaMalloc(void ** pointer, size_t nbytes)`
  - `cudaMemset(void * pointer, int value, size_t count)`
  - `cudaFree(void* pointer)`

```
int n = 1024;
int nbytes = 1024*sizeof(int);
int *d_a = 0;
cudaMalloc( (void**) &d_a, nbytes );
cudaMemset( d_a, 0, nbytes );
cudaFree(d_a);
```

## Memory management (2/2)

---

- **Data copy operation between CPU and device**
  - `cudaMemcpy(void *dst, void *src, size_t nbytes, enum cudaMemcpyKind direction);`
    - Direction specifies how to copy from src to dst , see below
    - Block a caller of CPU thread (execution) until the memory transfer completes.
    - Copy operation starts after previous CUDA calls.
  - `enum cudaMemcpyKind`
    - `cudaMemcpyHostToDevice`
    - `cudaMemcpyDeviceToHost`
    - `cudaMemcpyDeviceToDevice`

# Executing Code on the GPU

---

- **Kernels are C functions with some restrictions**
  - Can only access GPU memory
  - Must have void return type
  - No variable number of arguments (“varargs”)
  - Not recursive
  - No static variables
  - Function arguments
- **Function arguments automatically copied from CPU to GPU memory**

# Function Qualifiers

---

- **\_\_global\_\_** : invoked from within host (CPU) code, cannot be called from device (GPU) code must return void
- **\_\_device\_\_** : called from other GPU functions, cannot be called from host (CPU) code
- **\_\_host\_\_** : can only be executed by CPU, called from host
- **\_\_host\_\_** and **\_\_device\_\_** can be combined.
  - Sample use: overloading operators
  - Compiler will generate both CPU and GPU code

# CUDA Built-in Device Variables

---

- **\_\_global\_\_** and **\_\_device\_\_** functions have access to these automatically defined variables
  - **dim3 gridDim;**
    - Dimensions of the grid in blocks (at most 2D)
  - **dim3 blockDim;**
    - Dimensions of the block in threads
  - **dim3 blockIdx;**
    - Block index within the grid
  - **dim3 threadIdx;**
    - Thread index within the block

## A simple example

---

```
__global__ void minimal( int* d_a)
{
    *d_a = 13;
}
```

```
__global__ void assign( int* d_a, int value)
{
    int idx = blockDim.x * blockIdx.x + threadIdx.x;
    d_a[idx] = value;
}
```



## A simple example

```
__global__ void assign2D(int* d_a, int w, int h, int value)
{
    int iy = blockDim.y * blockIdx.y + threadIdx.y;
    int ix = blockDim.x * blockIdx.x + threadIdx.x;
    int idx = iy * w + ix;
    d_a[idx] = value;
}
...
assign2D<<<dim3(64, 64), dim3(16, 16)>>>(...);
```

## Example code to increment array elements

### CPU code

```
void inc_cpu(int*a, intN)
{
    int idx;
    for (idx =0;idx<N;idx++)
        a[idx]=a[idx] + 1;
}

voidmain()
{
    ...
    inc_cpu(a, N);
}
```

### CUDA codes

```
__global__ void
inc_gpu(int*a_d, intN){
    int idx = blockIdx.x* blockDim.x
            +threadIdx.x;
    if (idx < N)
        a_d[idx] = a_d[idx] + 1;
}

void main()
{
    ...
    dim3dimBlock (blocksize);
    dim3dimGrid(ceil(N/
                    (float)blocksize));
    inc_gpu<<<dimGrid,
            dimBlock>>>(a_d, N);
}
```

## Example (host-side program)

```
// allocate host memory
int numBytes = N * sizeof(float)
float* h_A = (float*) malloc(numBytes);

// allocate device memory
// float* d_A = 0;
cudaMalloc((void**) &d_A, numBytes);

// Copy data from host to device
cudaMemcpy(d_A, h_A, numBytes, cudaMemcpyHostToDevice);

// Execute kernel
increment_gpu<<< N/blockSize, blockSize>>>(d_A, b);

// copy back data from device to host
cudaMemcpy(h_A, d_A, numBytes, cudaMemcpyDeviceToHost);

// Free device memory
cudaFree(d_A);
```

```
int main() {
    float *a = new float[N*N];
    float *b = new float[N*N];
    float *c = new float[N*N];

    for ( int i = 0; i < N*N; ++i ) {
        a[i] = 1.0f; b[i] = 3.5f; }

    float *ad, *bd, *cd;
    const int size = N*N*sizeof(float);
    cudaMalloc( (void**)&ad, size );
    cudaMalloc( (void**)&bd, size );
    cudaMalloc( (void**)&cd, size );

    cudaMemcpy( ad, a, size, cudaMemcpyHostToDevice );
    cudaMemcpy( bd, b, size, cudaMemcpyHostToDevice );

    dim3 dimBlock( blockSize, blockSize );
    dim3 dimGrid( N/dimBlock.x, N/dimBlock.y );
    add_matrix<<<dimGrid, dimBlock>>>( ad, bd, cd, N );

    cudaMemcpy( c, cd, size, cudaMemcpyDeviceToHost );

    cudaFree( ad ); cudaFree( bd ); cudaFree( cd );
    delete[] a; delete[] b; delete[] c;
    return EXIT_SUCCESS;
}
```

# CUDA Qualifiers for variable

---

- **\_\_device\_\_**
  - Allocated in device global memory (Large, high-latency, no cache)
  - Allocated by `cudaMalloc` ( `__device__` is default)
  - Access by every thread.
  - extent: during execution of application
- **\_\_shared\_\_**
  - Stored in on-chip “shared memory” (SRAM, low latency)
  - Allocated by execution configuration or at compile time
  - Accessible by all threads in the same thread block
- **Unqualified variables**
  - Scalars and built-in vector types are stored in registers
  - Arrays may be in registers or local memory (*registers are not addressable*)

# How to use/specify shared memory

## Compile time

```
__global__ void kernel(...)  
{  
...  
__shared__ float sData[256];  
...  
}  
int main(void)  
{  
...  
kernel<<<nBlocks, blockSize>>> (...);  
}
```

## Invocation time

```
__global__ void kernel(...)  
{  
...  
extern __shared__ float sData[];  
...  
}  
  
int main(void)  
{  
...  
smBytes =  
blockSize*sizeof(float);  
kernel<<<nBlocks, blockSize,  
smBytes>>> (...);  
...  
}
```

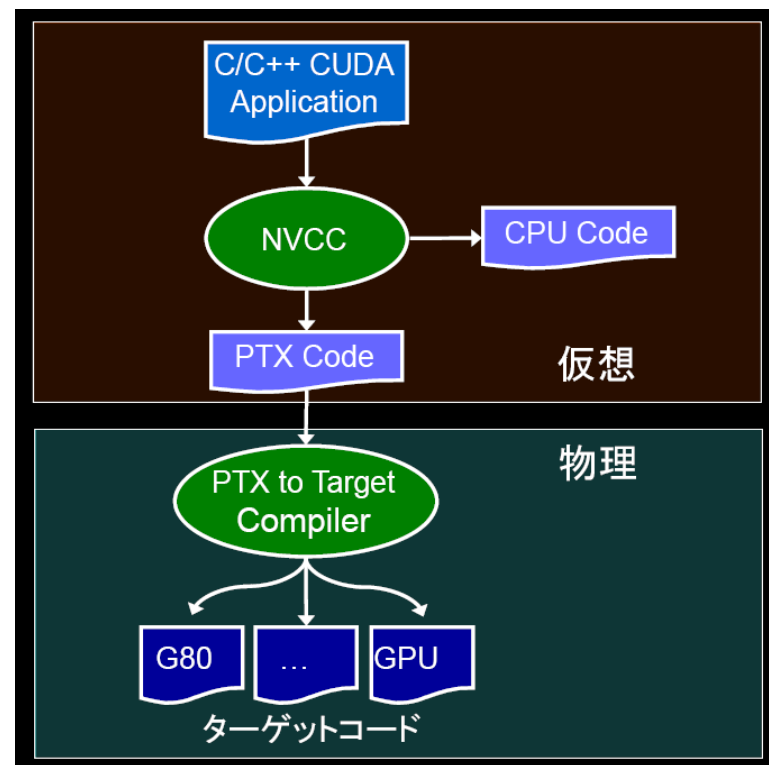
# GPU Thread Synchronization

---

- **`void __syncthreads () ;`**
  - Synchronizes all threads in a block
  - Generates barrier synchronization instruction
  - No thread can pass this barrier until all threads in the block reach it
  - Used to avoid RAW / WAR / WAW hazards when accessing shared memory
- **Allowed in conditional code only if the conditional is uniform across the entire thread block**
- **Synchronization between blocks is not supported**
  - Done by host-side

## Compiler

- C Source program with CUDA is compiled by **nvcc**.
- **Nvcc** is a **ccompile-driver**:
  - Execute required tools and **udacc**、**g++**、**cl**
- **Nvcc** generates following codes:
  - C object code (CPU code)
  - **PTX** code for GPU
  - **Glue code** to call GPU from CPU
- **Objects** required to execute **CUDA** program
  - **CUDA** core library (**cuda**)
  - **CUDA** runtime library (**cudart**)





# Optimization of GPU Programming

---

- **Maximize parallel using GPGPU**
- **Optimize/ avoid memory access to global memory**
  - Rather than storing data, re-computation may be cheaper in some cases
  - Coalescing memory access
  - Use cache in recent NVIDIA GPGPU
- **Optimize/avoid communication between CPU(host) and GPU (Device)**
  - Communication through PCI Express is expensive
  - Re-computing (redundant computing) may be cheaper than communications.

# Optimization of Memory access

---

- **Coalescing global memory access**
  - Combine memory access to contiguous area
- **Make use of shared memory**
  - Much faster than global memory (several x 100 times faster)
    - On-chip Memory
    - Low latency
  - Threads in block share the memory.
  - All threads can share the data computed by other threads.
  - To load shared memory from global memory, coalesce the memory and use them
- **Use cache (shared memory) as in conventional CPU**
  - Recent GPGPU has a cache at the same level of shared memory

## How to make use of different kinds of memory

---

- **Constant memory:**
  - Quite small, < 20K
  - As fast as register access if all threads in a warp access the same location
- **Texture memory:**
  - Spatially cached
  - Optimized for 2D locality
  - Neighboring threads should read neighboring addresses
  - No need to think about coalescing
- **Constraint:**
  - These memories can only be updated from the CPU

## Access to Global memory

---

- **4 cycles to issue on memory fetch**
- **but 400-600 cycles of latency**
  - The equivalent of 100 MADs
- **Likely to be a performance bottleneck**
- **Order of magnitude speedups possible**
  - Coalesce memory access (結合メモリアクセス)
- **Use shared memory to re-order non-coalesced addressing (共有メモリの利用)**

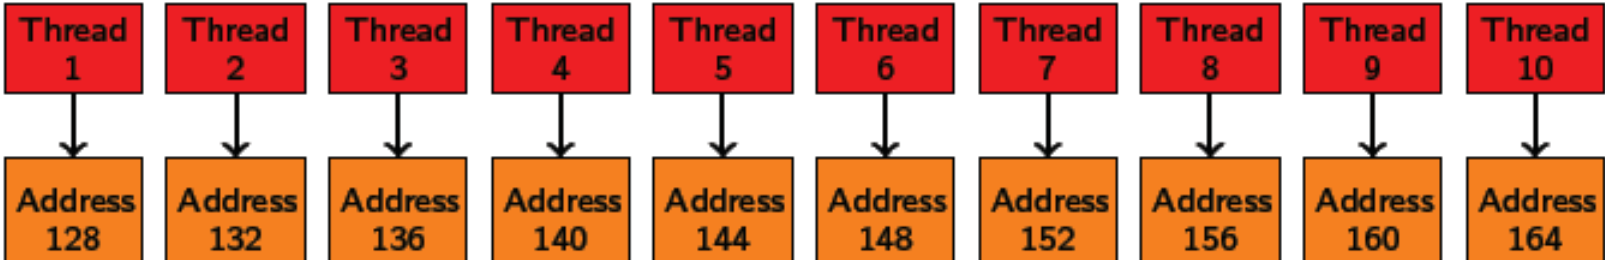
# Coalesced Memory Access

---


To exploit performance, global memory access should be coalesced (combined).

- A half warp (16 threads) memory access is coalesced.
- Contiguous memory access
  - 64 bytes – each thread reads a single word (int、floatなど)
  - 128bytes- each thread reads a double word (int2、float2など)
  - 256バイト- each thread reads a quad word (int4、float4など)
  - Float3 is not aligned ! ! !
- その他の制限
  - The start address of the contiguous area (Warp base address (WBA)) must be aligned the boundary of multiple of 数 $16 * \text{sizeof}(\text{type})$
  - The k-th thread in half warp must access the k-th element of the block
  - All threads in half warp may not be access.

# Coalesced Memory Access




Coalesced memory access:  
 Thread  $k$  accesses  $WBA + k$




Coalesced memory access:  
 Thread  $k$  accesses  $WBA + k$   
 Not all threads need to participate


# Case not coalesced



Non-Coalesced memory access:  
Misaligned starting address




Non-Coalesced memory access:  
Non-sequential access



Non-Coalesced memory access:  
Wrong size of type

# Example of memory optimization :

## Matrix Transpose



```

__global__ void
transpose_naive( float *out, float *in, int w, int h ) {
    unsigned int xIdx = blockDim.x * blockIdx.x + threadIdx.x;
    unsigned int yIdx = blockDim.y * blockIdx.y + threadIdx.y;

    if ( xIdx < w && yIdx < h ) {
        unsigned int idx_in = xIdx + w * yIdx;
        unsigned int idx_out = yIdx + h * xIdx;

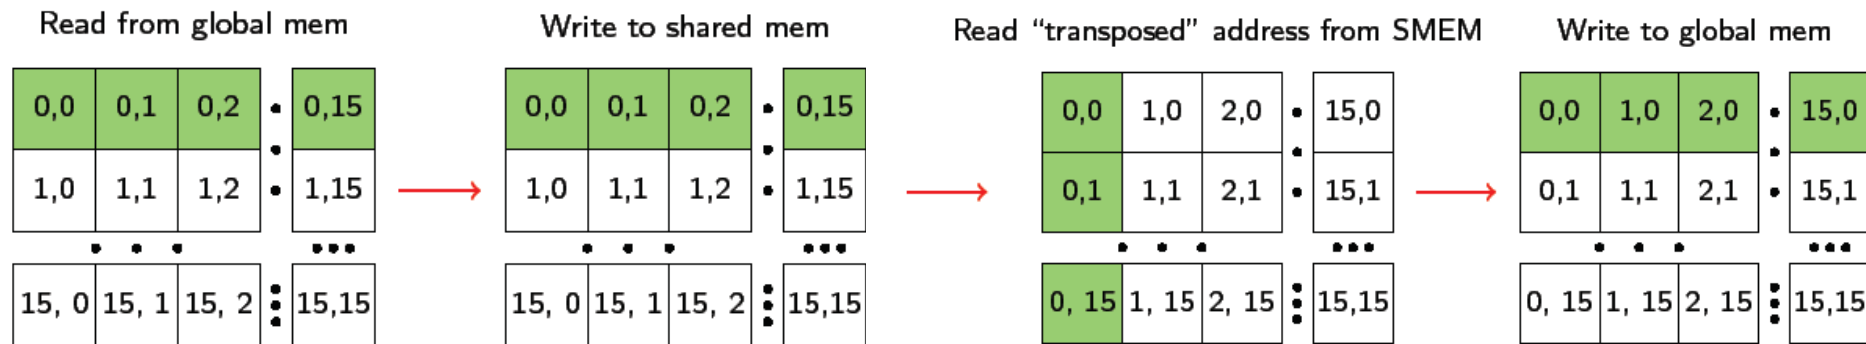
        out[idx_out] = in[idx_in];
    }
}

```

read側 (in) は、結合されるが、  
write側 (out) 側は結合されない。



# Optimization of memory access



- **By blocking, fetch block of data from shared memory, and store back the block of data to shared memory.**
- **The above example, thread block of 16 x 16 execute.**
- **Matrix is read and write for each 16 x 16 block**
- **When write back, write access is coalesced by contiguous memory address.**

## Optimized code (Coalesced)

```
__global__ void
transpose( float *out, float *in, int w, int h ) {
    __shared__ float block[BLOCK_DIM*BLOCK_DIM];
    unsigned int xBlock = blockDim.x * blockIdx.x;
    unsigned int yBlock = blockDim.y * blockIdx.y;
    unsigned int xIndex = xBlock + threadIdx.x;
    unsigned int yIndex = yBlock + threadIdx.y;
    unsigned int index_out, index_transpose;
    if ( xIndex < width && yIndex < height ) {
        unsigned int index_in = width * yIndex + xIndex;
        unsigned int index_block = threadIdx.y * BLOCK_DIM + threadIdx.x;
        block[index_block] = in[index_in];
        index_transpose = threadIdx.x * BLOCK_DIM + threadIdx.y;
        index_out = height * (xBlock + threadIdx.y) + yBlock + threadIdx.x;
    }
    __syncthreads();
    if ( xIndex < width && yIndex < height ) {
        out[index_out] = block[index_transpose];
    }
}
```

- **Example results**

Grid Size	Coalesced	Non-coalesced	Speedup
128 × 128	0.011 ms	0.022 ms	2.0×
512 × 512	0.07 ms	0.33 ms	4.5×
1024 × 1024	0.30 ms	1.92 ms	6.4×
1024 × 2048	0.79 ms	6.6 ms	8.4×

<http://www.sintef.no/upload/IKT/9011/SimOslo/eVITA/2008/seland.pdf>

## Optimization of Host-device communication

---

- **The bandwidth between host and device is very narrow compared with the bandwidth of device memory.**
  - Peak bandwidth 4GB/s (PCIe x16 1.0) vs. 76 GB/s (Tesla C870)
- **Minimize the communication between host-device**
  - Intermediate results must be kept in device memory to avoid communications
- **Grouping communication**
  - Large chunk of communication is more efficient than several small chunk of communications
- **Asynchronous communication**
  - Make use of stream
  - `cudaMemcpyAsync(dst, src, size, direction, 0);`

# Host Synchronization

---

- All kernel launches are *asynchronous*
  - control returns to CPU immediately
  - kernel executes after all previous CUDA calls have completed
- `cudaMemcpy ()` is *synchronous*
  - control returns to CPU after copy complete
  - copy starts after all previous CUDA calls have completed
- `cudaThreadSynchronize ()`
  - blocks until all previous CUDA calls complete

# OpenCL


---

- **Programming language for general purpose GPU computing.**
- **While C for CUDA is proprietary by NVIDIA, OpenCL is targeting cross-platform environments.**
  - **Only only for GPU such as NVIDIA and AMD(ATI), but also for conventional multicore CPU and many-core, such as Cell Broadband Engine(Cell B.E) and Intel MIC**
- **The point is that it targets for data parallel program by GPU and also for task-parallel of multi-core.**
- **What is different from CUDA? : Similar programming mode for kernel, but different in execution environment.**

# Kernel and Memory model

## OpenCL Memory Model

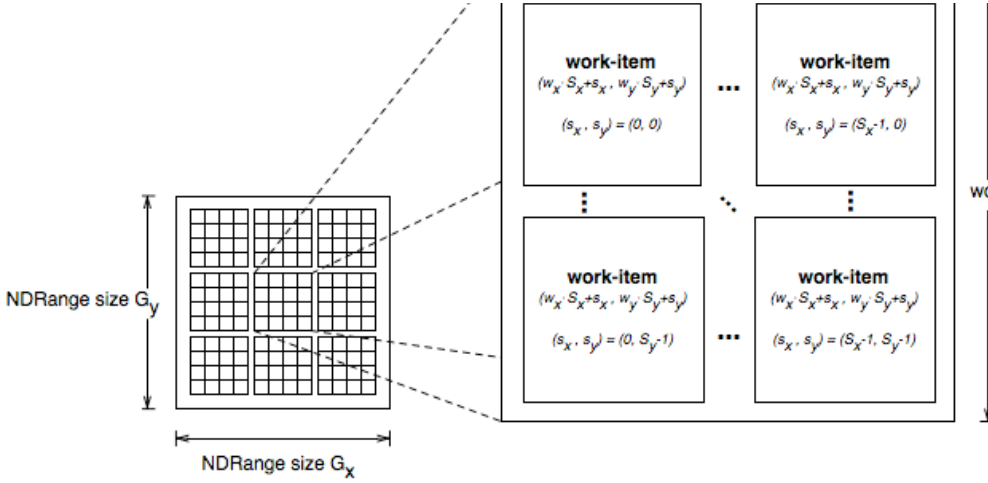
- **Private Memory**
  - Per work-item
- **Local Memory**
  - Shared within a workgroup (16Kb)
- **Local Global/Constant Memory**
  - Not synchronized
- **Host Memory**
  - On the CPU



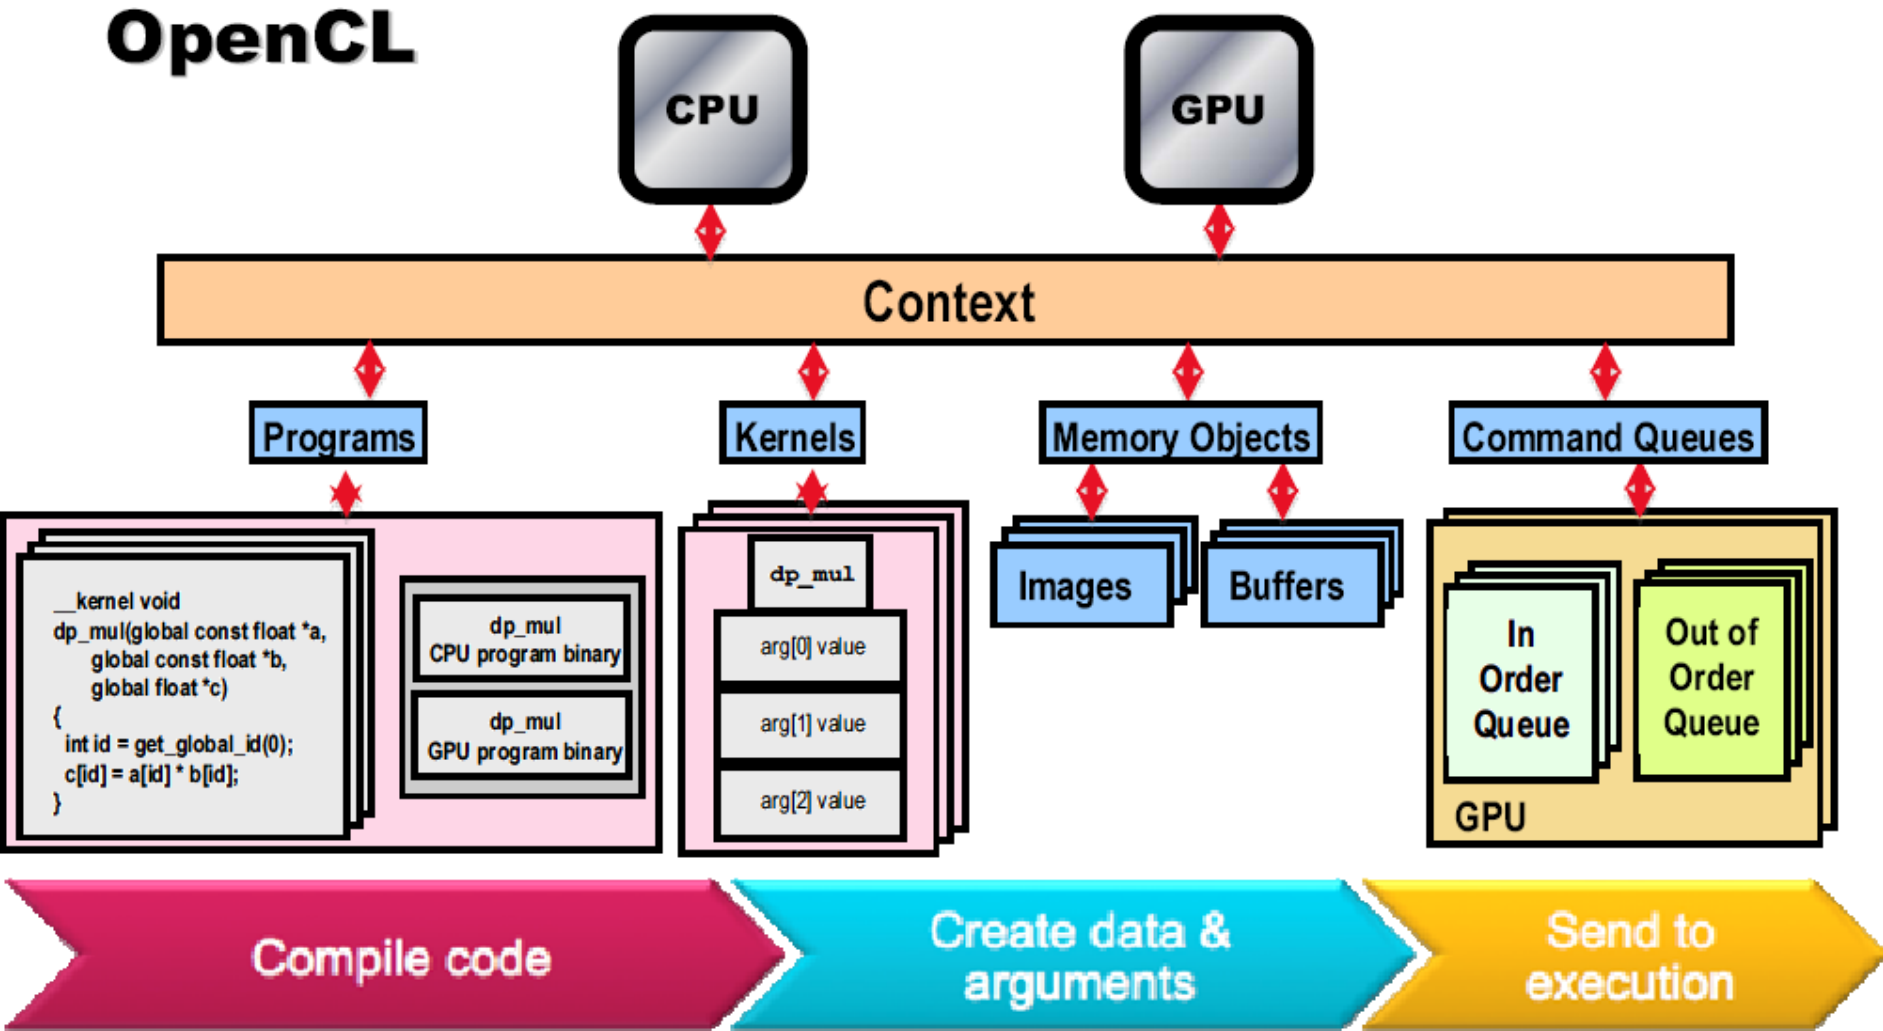
### Data Parallel

```
kernel void
dp_mul(global const float *a,
       global const float *b,
       global float *result)
{
  int id = get_global_id(0);

  result[id] = a[id] * b[id];
}
// execute dp_mul over "n" work-items
```



# Execution Environment of OpenCL





## Different Programming Styles



- **C for CUDA**
  - C with parallel keywords
  - C runtime that abstracts driver API
  - Memory managed by C runtime
  - Generates PTX
- **OpenCL**
  - Hardware API - similar to OpenGL
  - Programmer has complete access to hardware device
  - Memory managed by programmer
  - Generates PTX

現状のC for CUDAとOpenCLでは、位置付けがずれる。  
OpenCLがミドルウェアの土台としての色彩が濃いローレベルAPIであるのに対して、C for CUDAの方が抽象化の度合いが高くアプリケーションを書きやすい

# OpenACC

---

- **A spin-off activity from OpenMP ARB for supporting accelerators such as GPGPU and MIC**
- **NVIDIA, Cray Inc., the Portland Group (PGI), and CAPS enterprise**
- **Directive to specify the code offloaded to GPU.**




# A simple example

```

#define N 1024
int main(){
int i;
int a[N], b[N],c[N];
#pragma acc data copyin(a,b) copyout(c)
{
  #pragma acc parallel
  {
    #pragma acc loop
    for(i = 0; i < N; i++){
      c[i] = a[i] + b[i];
    }
  }
}
}


```

direction	copy	copyin	copyout
Host->device	○	○	
Device->Host	○		○



# A simple example

```
#define N 1024
int main(){
int i;
int a[N], b[N], c[N];
#pragma acc data copyin(a,b) copyout(c)
{
  #pragma acc parallel
  {
    #pragma acc loop
    for(i = 0; i < N; i++){
      c[i] = a[i] + b[i];
    }
  }
}
```



execute iterations  
like CUDA kernel

# Matrix Multiply in OpenACC

```
#define N 1024

void main(void)
{
    double a[N][N], b[N][N], c[N][N];
    int i,j;
    // ... setup data ...
#pragma acc parallel loop copyin(a, b) copyout(c)
    for(i = 0; i < N; i++){
#pragma acc loop
        for(j = 0; j < N; j++){
            int k;
            double sum = 0.0;
            for(k = 0; k < N; k++){
                sum += a[i][k] * b[k][j];
            }
            c[i][j] = sum;
        }
    }
}
```


## Stencil Code (Laplace Solver) in OpenACC

```
#define XSIZE 1024
#define YSIZE 1024
#define ITER 100
int main(void){
    int x, y, iter;
    double u[XSIZE][YSIZE], uu[XSIZE][YSIZE];
    // setup ...
#pragma acc data copy(u, uu)
    {
        for(iter = 0; iter < ITER; iter++){
            //old <- new
#pragma acc parallel loop
                for(x = 1; x < XSIZE-1; x++){
#pragma acc loop
                    for(y = 1; y < YSIZE-1; y++)
                        uu[x][y] = u[x][y];
                }
            //update
#pragma acc parallel loop
                for(x = 1; x < XSIZE-1; x++){
#pragma acc loop
                    for(y = 1; y < YSIZE-1; y++)
                        u[x][y] = (uu[x-1][y] + uu[x+1][y]
                                + uu[x][y-1] + uu[x][y+1]) / 4.0;
                }
            } //acc data end
    }
```

# Performance of OpenACC code

exec time


matrix multiply



# Performance of OpenACC code

exec time


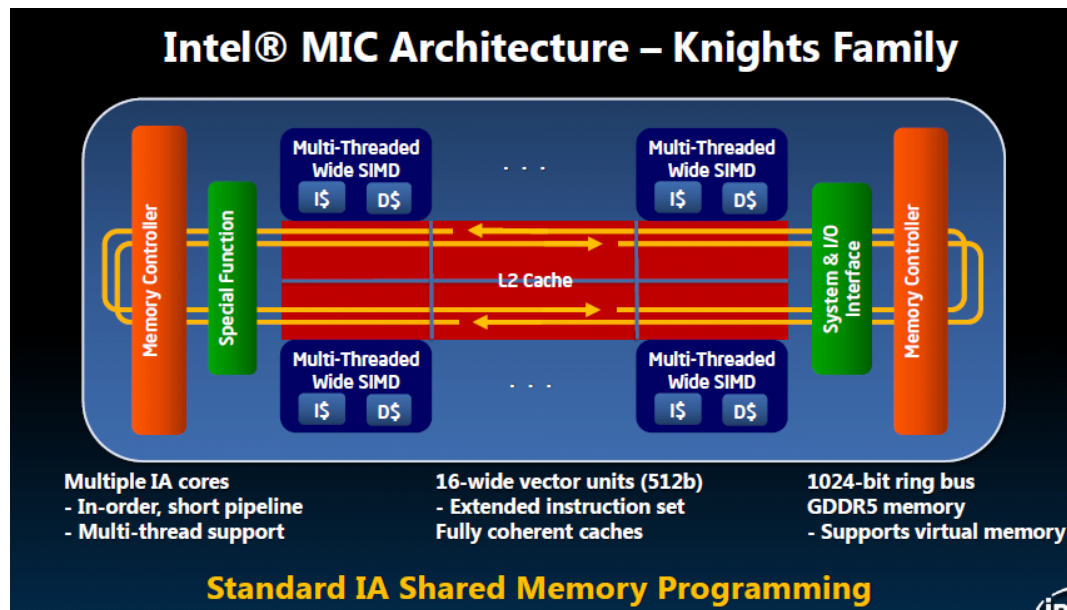
laplace



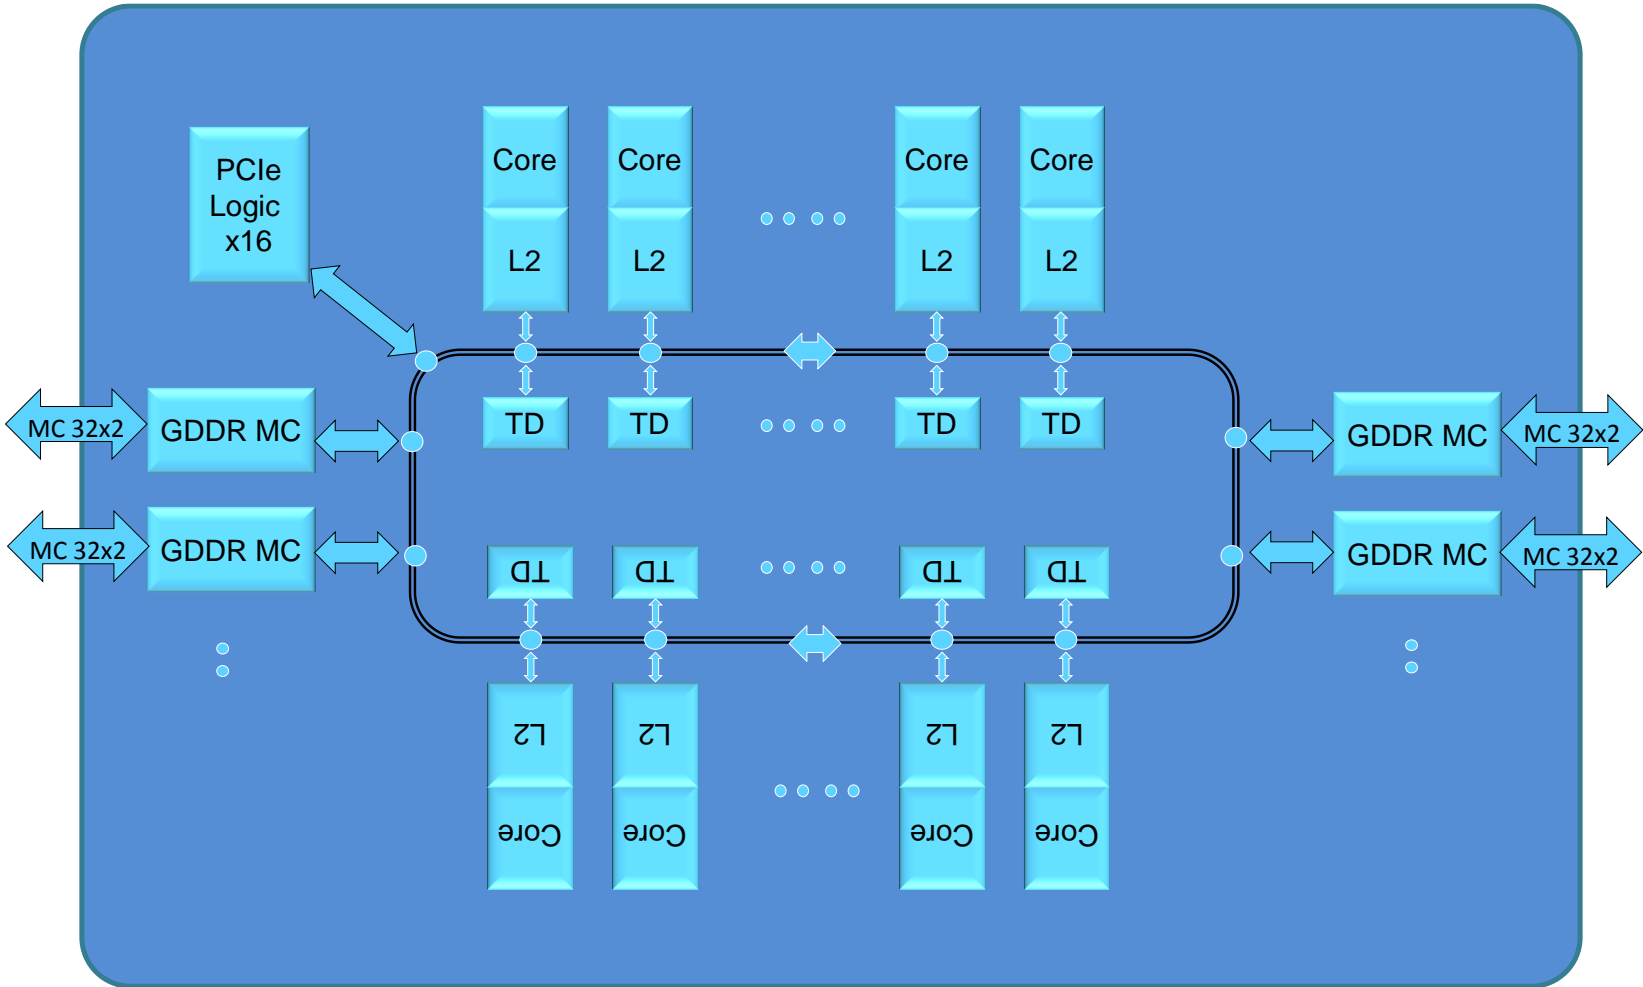
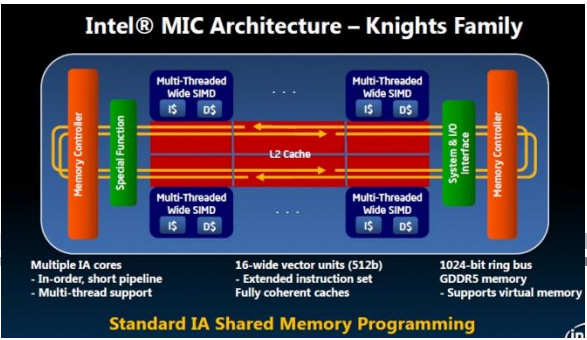


## Xeon Phi (Intel MIC)

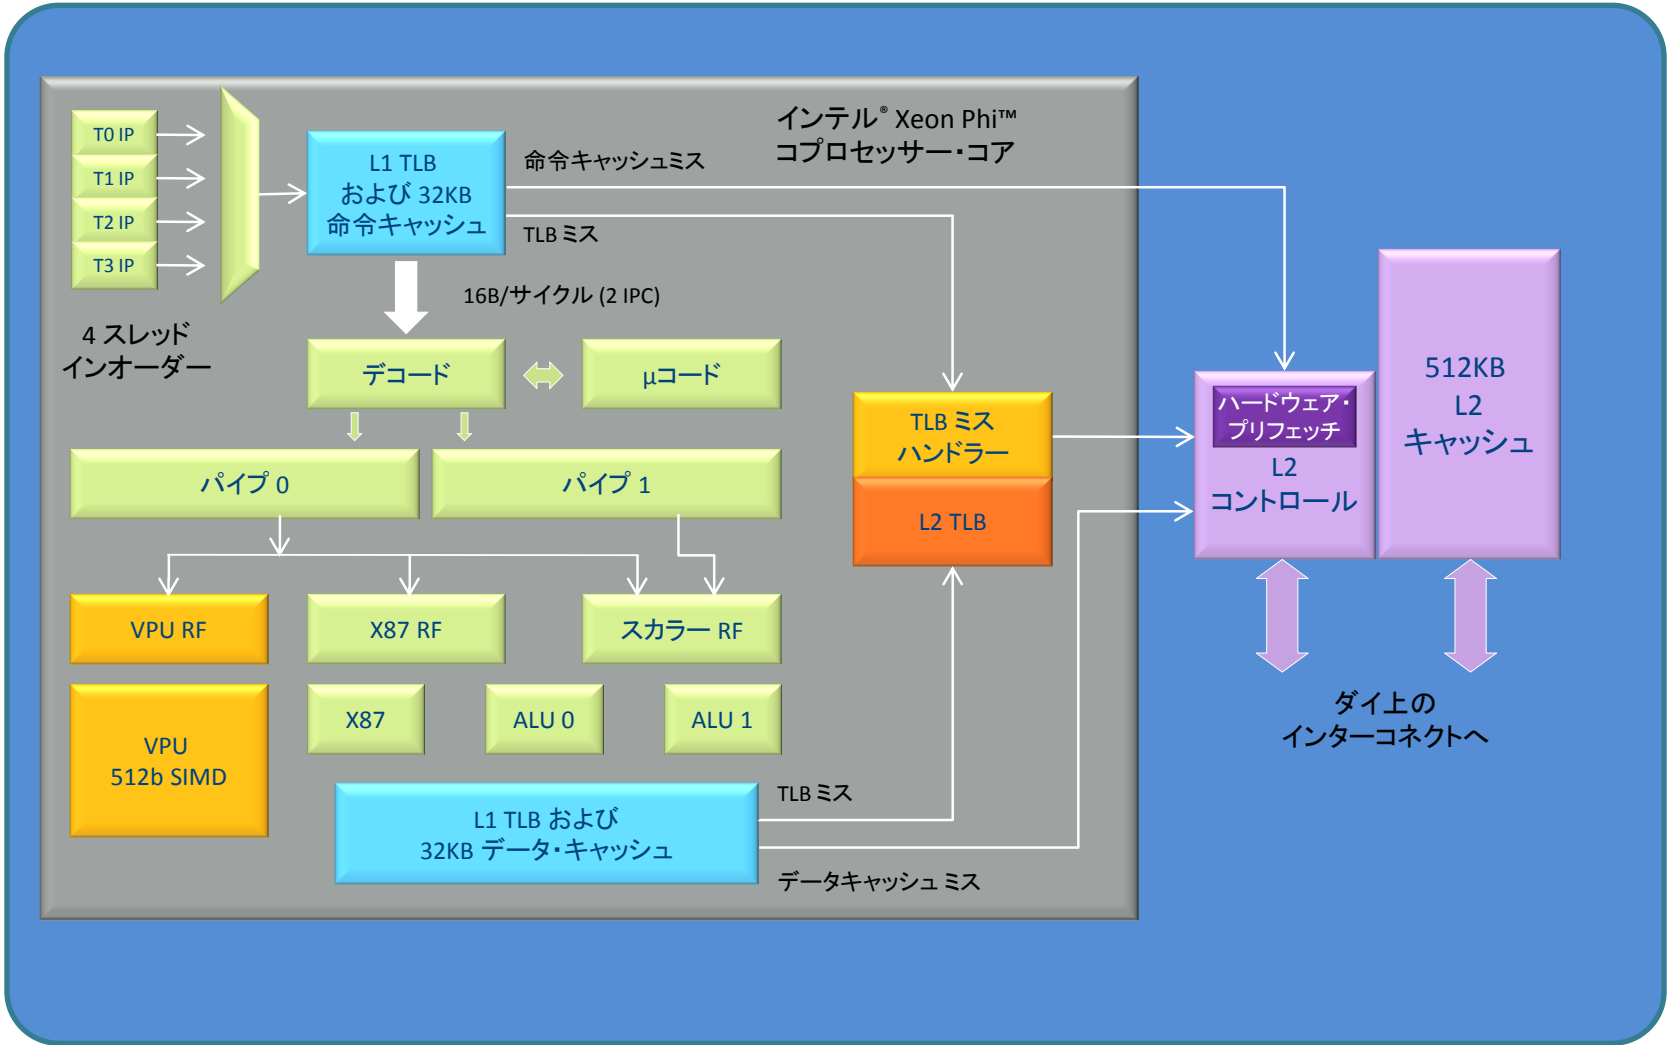
- Intel Manycore architecture
- released as a new lineup, Xeon Phi in Jan 2013
- Many core (> 60) using Intel IA architecture.



# Architecture of Xeon Phi




# Xeon Phi Core architecture



# Xeon Phi Server

- Current Xeon Phi (Knight Corner) used as a Co-processor



## The "Knights" Family

### Future Knights Products

#### Knights Corner

- 1<sup>st</sup> Intel® MIC product
- 22nm process
- >50 Intel Architecture Cores
- Within PCIe Power Envelope
- Additional Enhancements

#### Knights Ferry

Software Development Platform



Future options subject to change without notice.



# Comparison Xeon(host) and Xeon Phi

factors	Xeon	Xeon Phi
Clock Freq.	2.6 GHz	1.1GHz
Length of Vector	4	8
Peak DP perf /Core	20.8GFLOPS	17.6GFLOPS
Peak DP Scalar Perf/Core	5.2GFLOPS	1.1(0.55)GFLOPS
Number of Core	8	61
Peak DP perf. of Chip	332GFLOPS	1073GFLOPS
Cache size / core	20MB	512KB
Number of Threads /Core	2	4
Memory Freq.	1600MHz	2750MHz
Size of Main Memory	32GB	8GB
Peak Memory Bandwidth	102.4GBS	350GBS

# Programming model

- **offload model: hybrid execution of host and Phi by specifying a region to off-load to Phi**
- **Native model: execution by only Phi (note MPI can be used to communicate between host-Phi and Phi-Phi)**

## Example: Computing PI

```
# define NSET 1000000
int main ( int argc, const char** argv )
{ long int i;
  float num_inside, Pi;
  num_inside = 0.0f;
  #pragma offload target (MIC)
  #pragma omp parallel for reduction(+:num_inside)
  for( i = 0; i < NSET; i++ )
  {
    float x, y, distance_from_zero;
    // Generate x, y random numbers in [0,1)
    x = float(rand()) / float(RAND_MAX + 1);
    y = float(rand()) / float(RAND_MAX + 1);
    distance_from_zero = sqrt(x*x + y*y);
    if ( distance_from_zero <= 1.0f )
      num_inside += 1.0f;
  }
  Pi = 4.0f * ( num_inside / NSET );
  printf("Value of Pi = %f \n",Pi);
}
```

One additional line from the CPU version

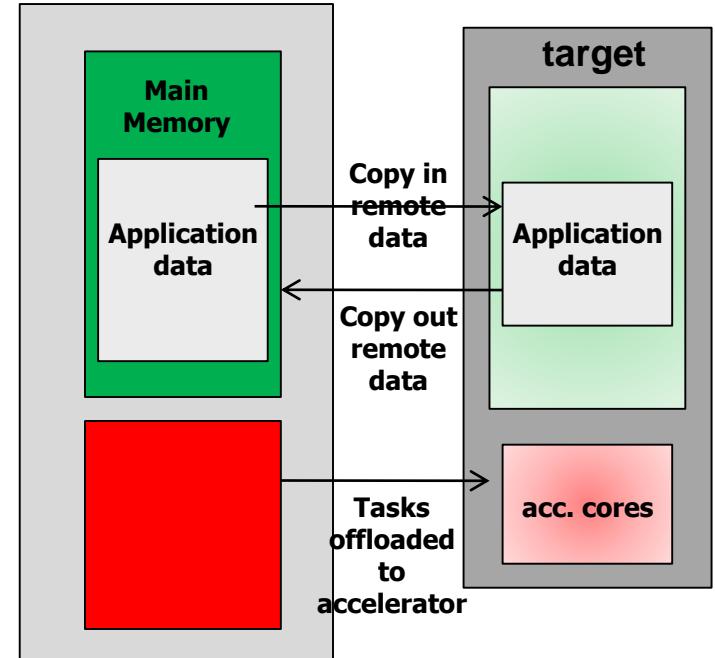
# OpenMP 4.0

---

- **Released July 2013**
  - <http://www.openmp.org/mp-documents/OpenMP4.0.0.pdf>
  - A document of examples is expected to release soon
- **Changes from 3.1 to 4.0 (Appendix E.1):**
  - *Accelerator: 2.9*
  - *SIMD extensions: 2.8*
  - *Places and thread affinity: 2.5.2, 4.5*
  - *Taskgroup and dependent tasks: 2.12.5, 2.11*
  - *Error handling: 2.13*
  - *User-defined reductions: 2.15*
  - *Sequentially consistent atomics: 2.12.6*
  - *Fortran 2003 support*

# Accelerator (2.9): offloading

- **Execution Model: Offload data and code to accelerator**
- *target* construct creates tasks to be executed by devices
- Aims to work with wide variety of accs
  - GPGPUs, MIC, DSP, FPGA, etc
  - A target could be even a remote node, intentionally



```
#pragma omp target
{
  /* it is like a new task
   * executed on a remote device */
  {
```



# Accelerator: explicit data mapping

- Relatively small number of truly shared memory accelerators so far
- Require the user to explicitly *map* data to and from the device memory
- Use array region

```
long a = 0x858;
long b = 0;
int anArray[100]

#pragma omp target data map(to:a) map(tofrom:b,anArray[0:64])
{
    /* a, b and anArray are mapped
     * to the device */

    /* work here */
}
/* b and anArray are mapped
 * back to the host */
```

# Accelerator: hierarchical parallelism

- Organize massive number of threads
  - teams of threads, e.g. map to CUDA grid/block
- Distribute loops over teams

```
#pragma omp target
```

```
#pragma omp teams num_teams(2)  
num_threads(8)
```

```
{  
    //-- creates a “league” of teams  
    //-- only local barriers permitted
```

```
#pragma omp distribute
```

```
for (int i=0; i<N; i++) {
```

```
}
```

Only **target**  
directive makes  
it as accelerator  
region

# target and map examples

```
void vec_mult(int N)
{
    int i;
    float p[N], v1[N], v2[N];
    init(v1, v2, N);
    #pragma omp target map(to: v1, v2) map(from: p)
    #pragma omp parallel for
    for (i=0; i<N; i++)
        p[i] = v1[i] * v2[i];
    output(p, N);
}
```

```
void vec_mult(float *p, float *v1, float *v2, int N)
{
    int i;
    init(v1, v2, N);
    #pragma omp target map(to: v1[0:N], v2[:N]) map(from: p[0:N])
    #pragma omp parallel for
    for (i=0; i<N; i++)
        p[i] = v1[i] * v2[i];
    output(p, N);
}
```

## Final remarks

---

- **GPGPU is a good solution for apps which can be parallelized for GPU.**
  - It can be very good esp. when the app fits into one GPU.
  - If the apps needs more than one GPU, the cost of communication will kill performance.
- **Programming in CUDA is still difficult ...**
  - Performance tuning, memory layout ...
  - OpenACC, and OpenMP 4.0, will help you!
- **Manycore (Xeon Phi) is emerging!**
  - Still offload model ... like GPU
- **For large scale computing, ... we will need multiple GPUs or MIC**
  - Hardware support for communication between GPU/MIC
  - AND, good programming environment