

Data Parallel Execution and CUDA Memories

Abhijit Bendale 03/11/2014











Simple Processing Flow





Simple Processing Flow





Simple Processing Flow







float *h_A, *h_B, *h_C float *d_A
malloc(); // assign memory cudaMalloc
free(); // free memory cudaFree()





We need to assign memory in the device (GPU) for the Variables that we wish to use in the device





All the blocks within the device have access to global memory of the device

+Grid, Block, Thread, Kernel..

int i = threadIdx.x + blockDim.x * blockIdx.x;



Execution of a CUDA program.

+Grid, Block, Thread, Kernel..

int i = threadIdx.x + blockDim.x * blockIdx.x;



+Grid, Block, Thread, Kernel..

int i = threadIdx.x + blockDim.x * blockIdx.x;



int i =	threadIdx.x	+	blockDim.x	*	<pre>blockIdx.x;</pre>
258 =	2	+	256 * 1		

A more complete version of vecAdd()







Memory Spaces

CPU and GPU have separate memory spaces

- Data is moved across PCIe bus
- Use functions to allocate/set/copy memory on GPU
 - Very similar to corresponding C functions

Pointers are just addresses

- Can't tell from the pointer value whether the address is on CPU or GPU
- Must exercise care when dereferencing:
 - Dereferencing CPU pointer on GPU will likely crash
 - Same for vice versa



+ Data Parallel Execution Model

- Fine-grained, data-parallel threads are fundamental means of parallel execution in CUDA
- Each thread uses a unique co-ordinate given by threadId {x,y,z}
- We will now study
 - Organization of threads
 - Resource Assignment to threads
 - Synchronization of threads
 - Scheduling of threads in a grid



Block IDs and Thread IDs





- Grid is a 3D array of blocks. Each block is a 3D array of threads.
- The exact organization of a grid is determined by configuration parameters <<< >>> of kernel launch statement
- <<< No of Blocks in grid, Number of threads per block >>>
- E.g. if we have to execture 4096 threads
 - dim3 dimBlock(128, 1, 1) \rightarrow Create 1D grid of 128 blocks
 - dim3 dimGrid(32, 1, 1) \rightarrow each block has 32 threads
 - vecAddKernel <<< dimBlock, dimGrid >>> (....);
 - 128*32 = 4096

Multi-Dimensional Data

MEDICAL IMAGING

BIOINFORMATICS

SUPERCOMPUTING CENTERS







COMPUTATIONAL FINANCE











+ Processing Multi-Dimensional Data

- The choice of 1D, 2D, 3D thread organization is done based on the nature of data
- ID data : vector manipulation
- 2D Data: Image processing
- 3D Data: MRI Scans



4 unused infeads in A difectio

Process an image of size 76 x 62 pixels. Block size: 16x16, Total blocks needed = 5x4 = 20

Goal : Double the value of each pixel

+A2DExample



```
pictureKernel<<< dimGrid, dimBlock>> (d Pin, d Pout, n, m)
```

gridDim.x = 5, gridDim.y = 4, blockDim.x = 16, blockDim.y = 16 Total number of threads generated = $76 \ge 62 = 4712$

+ Picture Kernel Code

}

__global___ void PictureKernel(float* d_pin, float* d_Pout, int m, int n){

Source code of PictureKernel() showing 2D thread mapping to a data pattern

if ((Row < m) && (Col < n))

rows and cols both in range



Some cols Out of range

Some rows And cols out Of range

Some rows out of range

+ Row Major indexing



index = row * width + col



Helpful in dynamic memory allocation. Address are considered in continuous locations based on datatype E.g. int = 4 bytes

+ Handling Data in 3D Arrav RGB 689 0.706 0.118 0.884 0.535 0.532 0.653 0.925 0.159 0.101 ... 0.633 0.528 0.493 intensitv 0.465 0.512 0.512 647 0.515 0.816 Paae 2 0.300 0.205 0.526 ... 0.219 0.328 712 0.929 ... 0.128 0.133 ... intensity values 0.100 0.121 0.113 0.175 0.986 0.234 0.432 .. Page 1 0.760 0.531 ... 0.765 0.128 0.863 0.521 ... 0.997 0.910 ... red 1.000 0.985 0.761 0.698 ... 0.995 0.726 ... intensity 0.455 0.783 0.224 0.395 ... values 0.021 0.500 0.311 0.123 ... 1.000 1.000 0.867 0.051 ... ₽lane 1.000 0.945 0.998 0.893 ... 0.990 0.941 1.000 0.876 ... row 0.902 0.867 0.834 0.798 ...

Col

The concept easily extends in 3D. Just have to keep track of Addition dimension

int Plane = blockIdx.z * blockDim.z + threaIdx;

The linearized access to array P will be in the form P[Plane * m * n + Row*n + Col]

Thus we have to keep track of 3 variables: Plane, row, col

+ Matrix Multiplication Example



We will consider square matrices only for clarity

Matrix Multiplication Using **Multiple Blocks**

- Break-up Pd into tiles •
- Each block calculates one • tile
 - Each thread calculates one element
 - Block size equal to tile size



bx

Tile Width = Block width (difference betn book editions)

Matrix Multiplication : Thread to Data Mapping

global void MatrixMulKernel(float* d_M, float* d_N, int m, int n){

}



- Divide the computation in tiles/blocks.
- Some block dimensions might be better than others (remember our 76 x 62 image example ?)
- Find the optimal Block/Tile Sizes: "Autotuning" to gain maximum performance gain
- E.g. if we wanted to process matrix of size 1000 x 1000
 - #define BLOCK_SIZE 16 // will generate 64 x 64 blocks
 - #define BLOCK_SIZE 32 // will generate 32 x 32 blocks
- Which one to use of the above 2 configuration? (Determined by other parameters as well (like number of streaming multiprocessors etc)

CUDA Thread Block

- All threads in a block execute the same kernel program (SPMD)
- Programmer declares block:
 - Block size 1 to **512** concurrent threads
 - Block shape 1D, 2D, or 3D
 - Block dimensions in threads
- Threads have thread id numbers within block
 - Thread program uses thread id to select work and address shared data
- Threads in the same block share data and synchronize while doing their share of the work
- Threads in different blocks cannot cooperate
 - Each block can execute in any order relative to other blocs!

Thread Id #: 0 1 2 3 ... m Thread program

CUDA Thread Block

Courtesy: John Nickolls, NVIDIA

© David Kirk/NVIDIA and Wen-mei W. Hwu, 2007-2009 ECE498AL, University of Illinois, Urbana-Champaign



+ Thread Synchronization

- To ensure that all threads in a block have completed a phase of their execution of the kernel before any of them can move on to the next phase
- CUDA follows barrier synchronization: Wait till all threads from the block have completed execution before context/ task switch

+ Barrier Synchronization



Device Runtime Component: Synchronization Function



void __syncthreads();

Synchronizes all threads in a block

- Once all threads have reached this point, execution resumes normally
- Used to avoid RAW / WAR / WAW hazards when accessing shared

Allowed in conditional code only if the conditional is uniform across the entire thread block

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 completes
- copy starts after all previous CUDA calls have completed

cudaThreadSynchronize()

blocks until all previous CUDA calls complete





Host Synchronization Example

// copy data from host to device cudaMemcpy(a_d, a_h, numBytes, cudaMemcpyHostToDevice);

// execute the kernel
inc_gpu<<<<ceil(N/(float)blocksize), blocksize>>>(a_d, N);

// run independent CPU code
run_cpu_stuff();

// copy data from device back to host cudaMemcpy(a_h, a_d, numBytes, cudaMemcpyDeviceToHost);

© 2008 NVIDIA Corporation

Basics
Thread Synchronization: Points to understand

- CUDA allows thread synchronization within the block but not across blocks
- This means blocks do not have time-dependency on one another: Can be executed in any order
- This flexibility allows scalable implementations
- Enables execution of same code at wide range of speeds (hence same code can be applied to different hardware) e.g. different wait times for resources
- The ability to execute the same application code on hardware with a different number of execution resouces is referred to as transparent scalability

Transparent Scalability

- Hardware is free to assign blocks to any processor at any time
 - A kernel scales across any number of parallel processors



© David Kirk/NVIDIA and Wen-mei W. Hwu, 2007-2009 ECE498AL, University of Illinois, Urbana-Champaign

+Assigning Resources



+ Executing thread blocks



© David Kirk/NVIDIA and Wen-mei W. Hwu, 2007-2009 ECE498AL, University of Illinois, Urbana-Champaign 17

SMs execute the

Execution resources are organized into streaming multiprocessors SMs

Scheduling Blocks onto SMs

- HW Schedules thread blocks onto available SMs
 - No guarantee of ordering among thread blocks
 - HW will schedule thread blocks as soon as a previous thread block finishes

+.

Mapping of Thread Blocks

- Each thread block is mapped to one or more warps
- The hardware schedules each warp independently

Thread Scheduling Example

- SM implements zero-overhead warp scheduling
 - At any time, only one of the warps is executed by SM
 - Warps whose next instruction has its inputs ready for consumption are eligible for execution
 - Eligible warps are selected for execution on a prioritized scheduling policy
 - All threads in a warp execute the same instruction when selected

+ Thread Scheduling

- Each Block is executed as 32-thread Warps
 - An implementation decision, not part of the CUDA programming model
 - Warps are scheduling units in SM
- If 3 blocks are assigned to an SM and each block has 256 threads, how many Warps are there in an SM?
 - Each Block is divided into 256/32 = 8 Warps
 - There are 8 * 3 = 24 Warps

© David Kirk/NVIDIA and Wen-mei W. Hwu, 2007-2009 ECE498AL, University of Illinois, Urbana-Champaign

G80 Example: Thread Scheduling (Cont.)

- SM implements zero-overhead warp scheduling
 - Warps whose next instruction has its operands ready for consumption are eligible for execution
 - Eligible Warps are selected for execution on a prioritized scheduling policy
 - All threads in a warp execute the same instruction when selected

© David Kirk/NVIDIA and Wen-mei W. Hwu, 2007-2009 ECE498AL, University of Illinois, Urbana-Champaign

G80 Block Granularity Considerations

- For Matrix Multiplication using multiple blocks, should I use 8X8, 16X16 or 32X32 blocks?
 - For 8X8, we have 64 threads per Block. Since each SM can take up to 768 threads, there are 12 Blocks. However, each SM can only take up to 8 Blocks, only 512 threads will go into each SM!
 - For 16X16, we have 256 threads per Block. Since each SM can take up to 768 threads, it can take up to 3 Blocks and achieve full capacity unless other resource considerations overrule.
 - For 32X32, we have 1024 threads per Block. Not even one can fit into an SM!

© David Kirk/NVIDIA and Wen-mei W. Hwu, 2007-2009 ECE498AL, University of Illinois, Urbana-Champaign

+Recap

Grid: Total number of threads Block: Organization for threads Streaming Processors: Resource allocation for Threads Grids, blocks, threads can have multiple dimensions

+ Device and Resource Query

- How to find device configurations?
- How many SMs? How many threads per block?

// cudaDeviceProp is a C structure

		int	maxTexture2DLayered [3]
		int	maxTexture2DLinear [3]
ata	Fields	int	maxTexture3D [3]
int	asyncEngineCount	int	maxTextureCubemap
int	canManHostMemory	int	maxTextureCubemapLayered [
int	clockPate	_int_	maxThreadsDim [3]
int	computeMode	int	maxThreadsPerBlock
int	concurrentKernels	int	maxThreadsPerMultiProcessor
int	deviceOverlap	int	memoryBusWidth
int	ECCEnabled	int	memoryClockRate
int	integrated	size_t	memPitch
int	kernelExecTimeoutEnabled	int	minor
int	I2CacheSize	int	multiProcessorCount
int	major	char	name [256]
int	maxGridSize [3]	int	pciBusID
int	maxSurface1D	int	pciDeviceID
int	maxSurface1DLayered [2]	int	pciDomainID
int	maxSurface2D [2]	int	regsPerBlock
int	maxSurface2DLayered [3]	size_t	sharedMemPerBlock
int	maxSurface3D [3]	size_t	surfaceAlignment
int	maxSurfaceCubemap	int	tccDriver
int	maxSurfaceCubemapLayered [2]	size t	textureAlignment
int	maxTexture1D	size t	texturePitchAlignment
int	maxTexture1DLayered [2]	size t	totalConstMem
int	maxTexture1DLinear	size t	totalGlobalMem
int	maxTexture2D [2]	int	unifiedAddressing
int	maxTexture2DGather [2]	int	warpSize

Device Management

CPU can query and select GPU devices

- cudaGetDeviceCount(int* count)
- cudaSetDevice(int device)
- cudaGetDevice(int *current_device)
- cudaGetDeviceProperties(cudaDeviceProp* prop,

int device)

cudaChooseDevice(int *device, cudaDeviceProp* prop)

- Multi-GPU setup:
 - device 0 is used by default
 - one CPU thread can control one GPU
 - multiple CPU threads can control the same GPU
 - calls are serialized by the driver

Kepler GK110 supports the new CUDA Compute Capability 3.5. (For a brief overview of CUDA see *Appendix A - Quick Refresher on CUDA*). The following table compares parameters of different Compute Capabilities for Fermi and Kepler GPU architectures:

	FERMI GF100	FERMI GF104	KEPLER GK104	KEPLER GK110
Compute Capability	2.0	2.1	3.0	3.5
Threads / Warp	32	32	32	32
Max Warps / Multiprocessor	48	48	<mark>64</mark>	<mark>64</mark>
Max Threads / Multiprocessor	1536	1536	2048	2048
Max Thread Blocks / Multiprocessor	8	8	16	16
32-bit Registers / Multiprocessor	32768	32768	65536	65536
Max Registers / Thread	63	<mark>63</mark>	<mark>63</mark>	255
Max Threads / Thread Block	1024	1024	1024	1024
Shared Memory Size Configurations (bytes)	16K	16K	16K	16K
	48K	48K	32K	32K
			48K	48K
Max X Grid Dimension	2^16-1	2^16-1	2^32-1	2^32-1
Hyper-Q	No	No	No	Yes
Dynamic Parallelism	No	No	No	Yes

Compute Capability of Fermi and Kepler GPUs

Specifications

Note: The below specifications represent this GPU as incorporated into NVIDIA's reference graphics card design. Graphics card specifications may vary by Add-in-card manufacturer. Please refer to the Add-in-card manufacturers' website for actual shipping specifications.

GTX 780 GPU Engine Specs:

CUDA Cores	230
Base Clock (MHz)	863
Boost Clock (MHz)	900
Texture Fill Rate (billion/sec)	160.

GTX 780 Memory Specs:

Memory Speed	6.0 Gbps
Standard Memory Config	3072 MB
Memory Interface	GDDR5
Memory Interface Width	384-bit
Memory Bandwidth (GB/sec)	288.4

GTX 780 Support:

Important Technologies	GPU Boost 2.0, PhysX, TXAA, NVIDIA G-SYNC-ready, SHIELD-ready
Other Supported Technologies	3D Vision, CUDA, DirectX 11, Adaptive VSync, FXAA, 3D Vision Surround, SLI-ready
OpenGL	4.3
Bus Support	PCI Express 3.0
Certified for Windows 7, Windows 8, Windows Vista, or Windows XP	Yes
3D Vision Ready	Yes
3D Gaming	Yes
Blu Ray 3D	Yes
3D Vision Live (Photos and Videos)	Yes

We need to assign memory in the device (GPU) for the Variables that we wish to use in the device

Hardware Implementation of CUDA Memories

- 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

Importance of Memory Access Efficiency

Every iteration has 2 global memory access for one floating point addition and one floating point multiplication. Thus it has *Compute to global memory access ratio (CGMA) is 1:1*

It has major performance implications: Eg: Memory Bandwidth: 200 GB/s Floating point size: 4 Bytes. Therefore 50 Gigs single precision operands/sec i.e. it will execute at the max 50 GFLOPS.

Peak performance usually at 1500 GFLOPS (1.5 TFLOPS) Only way to get around this is to increase CGMA ratio i.e

REDUCE MEMORY ACCESS

Lets understand threads in detail

- Thread is a virtualized von Neuman processor
- In von Neuman model, code of program is stored in memory, PC keeps track of particular point of the program, IR has instructions, Registers and memory holds value of variables and data structure

Processing Units and Threads

- Modern processors are designed to allow context switching, where multiple threads can time-share processor
- During context switch, intermediate values are saved in registers/memory
- GPUs allow multiple processors, single instruction i.e all processors execute same instructions. Hence, resource sharing between threads is important.
- The reason why threads are organized into blocks/warps

CUDA Variable Type Qualifiers

Variable declaration	Memory	Scope	Lifetime
int var;	register	thread	thread
<pre>int array_var[10];</pre>	local	thread	thread
<pre>shared int shared_var;</pre>	shared	block	block
device int global_var;	global	grid	application
<pre>constant int constant_var;</pre>	constant	grid	application

"automatic" scalar variables without qualifier reside in a register

- compiler will spill to thread local memory
- "automatic" array variables without qualifier reside in thread-local memory

CUDA Variable Type Performance

Variable declaration	Memory	Penalty
int var;	register	1x
<pre>int array_var[10];</pre>	local	100x
<pre>shared int shared_var;</pre>	shared	1x
device int global_var;	global	100x
<pre>constant int constant_var;</pre>	constant	1x

- scalar variables reside in fast, on-chip registers
- shared variables reside in fast, on-chip memories
- thread-local arrays & global variables reside in uncached off-chip memory
- constant variables reside in cached off-chip memory

- Global memory resides in device memory (DRAM)
 - Much slower access than shared memory
- Tile data to take advantage of fast shared memory:
 - Generalize from adjacent_difference example
 - Divide and conquer

Partition data into subsets that fit into shared memory

© 2008 NVIDIA Corporation

Handle each data subset with one thread block

Load the subset from global memory to shared memory, using multiple threads to exploit memorylevel parallelism

© 2008 NVIDIA Corporation

Perform the computation on the subset from shared memory

© 2008 NVIDIA Corporation

Copy the result from shared memory back to global memory

- Carefully partition data according to access patterns
- Read-only
 Constant memory (fast)
- R/W & shared within block → _____shared____ memory (fast)
- R/W within each thread registers (fast)
- Indexed R/W within each thread -> local memory (slow)

Tiled Matrix Multiplication

	Phase 1		Phase 2			
Т _{0,0}	Md _{0,0}	Nd _{0,0}	PValue _{0,0} +=	Md _{2,0}	Nd _{0,2}	PValue _{0,0} +=
	↓	↓	Mds _{0,0} *Nds _{0,0} +	↓	↓	Mds _{0,0} *Nds _{0,0} +
	Mds _{0,0}	Nds _{0,0}	Mds _{1,0} *Nds _{0,1}	Mds _{0,0}	Nds _{0,0}	Mds _{1,0} *Nds _{0,1}
T _{1,0}	Md _{1,0}	Nd _{1,0}	PValue _{1,0} +=	Md _{3,0}	Nd _{1,2}	PValue _{1,0} +=
	↓	↓	Mds _{0,0} *Nds _{1,0} +	↓	↓	Mds _{0,0} *Nds _{1,0} +
	Mds _{1,0}	Nds _{1,0}	Mds _{1,0} *Nds _{1,1}	Mds _{1,0}	Nds _{1,0}	Mds _{1,0} *Nds _{1,1}
Т _{0,1}	Md _{0,1}	Nd _{0,1}	PdValue _{0,1} +=	Md _{2,1}	Nd _{0,3}	PdValue _{0,1} +=
	↓	↓	Mds _{0,1} *Nds _{0,0} +	↓	↓	Mds _{0,1} *Nds _{0,0} +
	Mds _{0,1}	Nds _{0,1}	Mds ₁₁ *Nds _{0,1}	Mds _{0,1}	Nds _{0,1}	Mds _{1,1} *Nds _{0,1}
T _{1,1}	Md _{1,1}	Nd _{1,1}	PdValue _{1,1} +=	Md _{3,1}	Nd _{1,3}	PdValue _{1,1} +=
	↓	↓	Mds _{0,1} *Nds _{1,0} +	↓	↓	Mds _{0,1} *Nds _{1,0} +
	Mds _{1,1}	Nds _{1,1}	Mds _{1,1} *Nds _{1,1}	Mds _{1,1}	Nds _{1,1}	Mds _{1,1} *Nds _{1,1}

time

Tiled Multiply

 Break up the execution of the kernel into phases so that the data accesses in each phase is focused on one subset (tile) of Md and Nd

0

2

by 1

Md

TILE_WIDTH TILE_WIDTH

WIDTH

0

12

ty

TILE_WIDTH-1

A Small Example: Multiplication

© David Kirk/NVIDIA and Wen-mei W. Hwu, 2007-2009 ECE498AL, University of Illinois, Urbana-Champaign

Tiling Size Effects

© David Kirk/NVIDIA and Wen-mei W. Hwu, 2007-2010 ECE408, University of Illinois, Urbana Champaign

```
__global___void MatrixMulKernel(float* Md, float* Nd, float* Pd, int Width)
{
1. ___shared__float Mds[TILE_WIDTH][TILE_WIDTH];
2. __shared__float Nds[TILE_WIDTH][TILE_WIDTH];
3. int bx = blockIdx.x; int by = blockIdx.y;
4. int tx = threadIdx.x; int ty = threadIdx.y;
// Identify the row and column of the Pd element to work on
5. int Row = by * TILE_WIDTH + ty;
```

```
6. int Col = bx * TILE_WIDTH + tx;
```

```
7. float Pvalue = 0;
// Loop over the Md and Nd tiles required to compute the Pd element
8. for (int m = 0; m < Width/TILE_WIDTH; ++m) {</pre>
```

// Collaborative loading of Md and Nd tiles into shared memory

```
9. Mds[ty][tx] = Md[Row*Width + (m*TILE_WIDTH + tx)];
```

```
10. Nds[ty][tx] = Nd[(m*TILE_WIDTH + ty)*Width + Col];
```

```
11. ____syncthreads();
```

```
12. for (int k = 0; k < TILE_WIDTH; ++k)
```

```
13. Pvalue += Mds[ty][k] * Nds[k][tx];
```

```
14. __syncthreads();
```

```
15. Pd[Row*Width + Col] = Pvalue;
```

Points to Remember about CUDA Memory

- CUDA offers local memories/registers. Using these efficiently reduces access to global memory and improves performance. But it requires algorithm re-design
- Blocks, Registers have limited memory. If data exceeds this shared memory requirement, data has to be split into multiple pieces
- The ability to reason about hardware limitation when developing an application is key aspect of computational thinking
- Tiled algorithms often increase performance. But key to it is to exploit data locality
+ Accelerating MATLAB with CUDA

- Massimiliano Fatica
- NVIDIA
- mfatica@nvidia.com

Won-Ki Jeong University of Utah

wkjeong@cs.utah.edu



MATLAB can be easily extended via MEX files to take advantage of the computational power offered by the latest NVIDIA GPUs (GeForce 8800, Quadro FX5600, Tesla).

Programming the GPU for computational purposes was a very cumbersome task before CUDA. Using CUDA, it is now very easy to achieve impressive speed-up with minimal effort.

This work is a proof of concept that shows the feasibility and benefits of using this approach.





Even though MATLAB is built on many well-optimized libraries, some functions can perform better when written in a compiled language (e.g. C and Fortran).

 MATLAB provides a convenient API for interfacing code written in C and FORTRAN to MATLAB functions with MEX files.

 MEX files could be used to exploit multi-core processors with OpenMP or threaded codes or like in this case to offload functions to the GPU.



- Native MATLAB script cannot parse CUDA code
- New MATLAB script nvmex.m compiles CUDA code (.cu) to create MATLAB function files
- Syntax similar to original mex script:

>> nvmex –f nvmexopts.bat filename.cu –IC:\cuda\include

-LC:\cuda\lib -lcudart

Available for Windows and Linux from:

http://developer.nvidia.com/object/matlab_cuda.html

Mex files for CUDA

A typical mex file will perform the following steps:

- 1.Convert from double to single precision
- 2.Rearrange the data layout for complex data
- 3.Allocate memory on the GPU
- 4. Transfer the data from the host to the GPU
- 5.Perform computation on GPU (library, custom code)
- 6.Transfer results from the GPU to the host
- 7.Rearrange the data layout for complex data
- 8.Convert from single to double
- 9. Clean up memory and return results to MATLAB

Some of these steps will go away with new versions of the library (2,7) and new hardware (1,8)

CUDA MEX example Additional code in MEX file to handle CUDA

/*Parse input, convert to single precision and to interleaved complex format */

/* Allocate array on the GPU */

cufftComplex *rhs_complex_d;

cudaMalloc((void **) &rhs_complex_d,sizeof(cufftComplex)*N*M);

/* Copy input array in interleaved format to the GPU */

cudaMemcpy(rhs_complex_d, input_single, sizeof(cufftComplex)*N*M,

/* Create plan for CUDA FFT NB: transposing dimensions*/

cufftPlan2d(&plan, N, M, CUFFT_C2C);

/* Execute FFT on GPU */

cufftExecC2C(plan, rhs_complex_d, rhs_complex_d, CUFFT_INVERSE);

/* Copy result back to host */

 $cudaMemcpy(\ input_single, rhs_complex_d, sizeof(cufftComplex)*N*M,$

/* Clean up memory and plan on the GPU */

cufftDestroy(plan); cudaFree(rhs_complex_d);

/*Convert back to double precision and to split complex format */

cudaMemcpyHostToDevice);

cudaMemcpyDeviceToHost);



- Focus on 2D FFTs.
- FFT-based methods are often used in single precision (for example in image processing)
- Mex files to overload MATLAB functions, no modification between the original MATLAB code and the accelerated one.
- Application selected for this study:

solution of the Euler equations in vorticity form using a pseudospectral method.

+ Implementation details:

Case A) FFT2.mex and IFFT2.mex

Mex file in C with CUDA FFT functions.

Standard mex script could be used.

Overall effort: few hours

Case B) Szeta.mex: Vorticity source term written in CUDA

Mex file in CUDA with calls to CUDA FFT functions.

Small modifications necessary to handle files with a .cu suffix

Overall effort: ¹/₂ hour (starting from working mex file for 2D FFT)



Hardware:

AMD Opteron 250 with 4 GB of memory

NVIDIA GeForce 8800 GTX

Software:

Windows XP and Microsoft VC8 compiler

RedHat Enterprise Linux 4 32 bit, gcc compiler

MATLAB R2006b

CUDA 1.0



+ FFT2 performance



Vorticity source term

http://www.amath.washington.edu/courses/571-winter-2006/matlab/Szeta.m

function S = Szeta(zeta,k,nu4)
% Pseudospectral calculation of vorticity source term
% S = -(- psi_y*zeta_x + psi_x*zeta_y) + nu4*del^4 zeta
% on a square periodic domain, where zeta = psi_xx + psi_yy is an NxN matrix
% of vorticity and k is vector of Fourier wavenumbers in each direction.

% Output is an NxN matrix of S at all pseudospectral gridpoints

zetahat = fft2(zeta);

[KX KY] = meshgrid(k,k);

% Matrix of (x,y) wavenumbers corresponding

% to Fourier mode (m,n)

del2 = -(KX.^2 + KY.^2);

del2(1,1) = 1; % Set to nonzero to avoid division by zero when inverting

% Laplacian to get psi

psihat = zetahat./del2;

dpsidx = real(ifft2(li*KX.*psihat));

dpsidy = real(ifft2(li*KY.*psihat));



The current CUDA FFT library only supports interleaved format for complex data while MATLAB stores all the real data followed by the imaginary data.

Complex to complex (C2C) transforms used

The accelerated computations are not taking advantage of the symmetry of the transforms.

The current GPU hardware only supports single precision (double precision will be available in the next generation GPU towards the end of the year). Conversion to/from single from/to double is consuming a significant portion of wall clock time.

Advection of an elliptic vortex

256x256 mesh, 512 RK4 steps, Linux, MATLAB file http://www.amath.washington.edu/courses/571-winter-2006/matlab/FS_vortex.m





MATLAB with CUDA (single precision FFTs) 14.9 seconds (11x)

₱seudo-spectral simulation of 2D Isotropic turbulence.

512x512 mesh, 400 RK4 steps, Windows XP, MATLAB file http://www.amath.washington.edu/courses/571-winter-2006/matlab/FS_2Dturb.m







MATLAB with CUDA (single precision FFTs) 93 seconds

- +
 - Power spectrum of vorticity is very sensitive to fine scales. Resul from original MATLAB run and CUDA accelerated one are in excellent agreement





MATLAB run

CUDA accelerated MATLAB run

Timing details

1024x1024 mesh, 400 RK4 steps on Windows, 2D isotropic turbulence

	Runtime Opteron 250	Speed up	Runtime Opteron 2210	Speed up
PCI-e Bandwidth:	1135 MB/s		1483 MB/s	
Host to/from device	1003 MB/S		1223 WB/S	
Standard MATLAB	8098 s		9525s	
Overload FFT2 and IFFT2	4425 s	1.8x	4937s	1.9x
Overload Szeta	735 s	11.x	789s	12.X
Overload Szeta , FFT2 and IFFT2	577 s	14.x	605s	15.7x