

Parallel and Distributed Programming for Data/Computation Intensive Applications

Bilal Jan

Department of Control and Computer Engineering
Politecnico di Torino

Supervisors: Bartolomeo Montrucchio (DAUIN)
Carlo Ragusa (DENERG)

February 27, 2015



Outlines

Introduction

Parallel Sorting of Large Data Sets

Fast Fourier Transforms on GPU

Accelerating Micromagnetic Simulations on GPUs

Conclusion



GPU Computing

- ▶ An assembly of hundreds and thousands of PUs
- ▶ SIMD Processing: Single Instruction on Multiple Data streams simultaneously
- ▶ Well suited for highly parallel numeric applications
- ▶ Best Price/Performance ratio
- ▶ Programming Tools
 - ▶ CUDA (NVIDIA Proprietary)
 - ▶ OpenCL (Open Standard and Heterogeneous)




Figure: CPU vs GPU core ratio




Figure: Performance Comparison CPU vs GPU

Parallel Sorting 1/3: Min-Max Butterfly Sort

- ▶ Finds minimum and maximum in data
- ▶ For data of size N
 - ▶ Total Stages are $\log_2 N$
 - ▶ Complexity in terms of butterflies (comparators) is $(N/2)\log_2 N$
- ▶ All stages are executed sequentially
- ▶ Butterflies inside any stage S_i are executed in parallel
- ▶ After complete run of the Algorithm minimum and maximum values in data are placed at $x(0)$ and $x(N-1)$




Fig : 1 8x8 Min-Max Butterfly

Parallel Sorting 2/3: Full Butterfly Sorting

- ▶ Complete Sorting
- ▶ For data of size **N**
 - ▶ Total Stages are $\log_2 N + \sum_{r=1}^{\log_2 N - 1} (r)$
 - ▶ Complexity in terms of butterflies (comparators) is $(N/2) \times T_Stages$
- ▶ All stages are executed sequentially
- ▶ Butterflies inside any stage **S_i** are executed in parallel






Fig : Size 8 Full Butterfly Sorting.


Parallel Sorting 3/3: Results Butterfly Network Sort




Sorting time of Min-Max Butterfly sort for various input size and devices types.



Sorting time of full Butterfly Sort for various input in the power of 2 and different devices types



Speedup of full Butterfly sort vs Bitonic sort.



Speedup of full Butterfly sort against less parallel Odd-Even sort.

Fourier Transformation on GPU

- ▶ DFT converts a time domain signal into frequency domain.
- ▶ High computation complexity $\mathcal{O}(n^2)$
- ▶ FFT are fast methods for computing the DFT.
- ▶ FFT complexity $\mathcal{O}(n \log n)$.
- ▶ The parallel structure of Cooley-Tukey FFT is well suited for GPU architecture

$$X(k) = \sum_{n=0}^{N-1} x(n) W_N^{nk}$$

$$\begin{bmatrix} X_0 \\ X_1 \\ \vdots \\ X_{N-1} \end{bmatrix} = \begin{bmatrix} W^0 & W^0 & \dots & W^0 \\ W^0 & W^1 & \dots & W^{N-1} \\ \vdots & \vdots & \ddots & \vdots \\ W^0 & W^{N-1} & \dots & W^{(N-1)(N-1)} \end{bmatrix} \times \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_{N-1} \end{bmatrix}$$

$$X[K] = \sum_{n=0}^{\frac{N}{2}-1} x[2n] W_{N/2}^{nk} + W_N^k \sum_{n=0}^{\frac{N}{2}-1} x[2n+1] W_{N/2}^{nk}$$

$$X[K] = A[k] + W_N^k \cdot B[k]$$

The Cooley-Tukey Algorithm for FFT

Twiddle Properties

Symmetry

$$W_N^{k+\frac{N}{2}} = -W_N^k$$

Periodicity

$$W_N^{k+N} = W_N^k$$

Recursion

$$W_N^2 = -W_{N/2}$$


$$\begin{bmatrix} Y_0 \\ Y_1 \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \times \begin{bmatrix} A \\ B \end{bmatrix}$$

$$A + W^k B = Y[k]$$

$$A - W^k B = Y[k+1]$$

2 Point DFT

Cooley-Tukey DIT-FFT, Radix-2, N=8



S=Memory Load/Store Stride

ToPe: An OpenCL based multidimensional FFT Library


- ▶ Almost Arbitrary length transform size
- ▶ Complex-to-Complex Transform type
- ▶ Multi-Radix ($N = r^n$, where $r = 2 - 8, 10, 15, 16$)
- ▶ Algorithms (Cooley-Tukey DIT, Modified Cooley-Tukey, Mixed Radix FFT)
- ▶ Dimension supported up to 3D
- ▶ Precision (single and double)
- ▶ Auto tuning for multiple GPU with Static Load Balancing (GPU+Thread Level)
- ▶ Open Source (<http://code.google.com/p/tope-fft>)

Speedup over FFTW $\sim 50\times$


Speedup over cuFFT $\sim 5\times$

Results 1/2: ToPe FFT


Running Time of FFT Libraries for 1D Radix-5 S-Precision on GTX-260




Running Time of FFT Libraries for 1D Radix-6 S-Precision on GTX-260



Running Time of FFT Libraries for 1D Radix-7 S-Precision on GTX-260




Running Time of FFT Libraries for 1D Radix-10 S-Precision on GTX-260




Results 2/2: ToPe FFT

Running Time of ToPe 2D-FFT on GTX-260 for D-Precision C2C
Input Size $N = N_x \cdot N_y$ ($N_x=1024$, N_y changes)




2D FFT

Running Time of ToPe 2D-FFT on GTX-260 for D-Precision C2C
Input Size $N = N_x \cdot N_y$ ($N_x=2048$, N_y changes)




2D FFT

Speedup ToPe vs FFTW on multiple GPUs for radix-2



1D Speedup against FFTW on Multiple GPUs

Speedup ToPe vs cuFFT on multiple GPUs for radix-2 on GTX-295



1D Speedup against cuFFT on Multiple GPUs

Micromagnetics 1/3: Accelerating Magnetostatic Field Computation using GPUs

The equation for Magnetostatic field \mathbf{H} at a cell position is


$$\mathbf{H}(\mathbf{r}) = - \sum_{\mathbf{r}'}^n \mathbf{N}(\mathbf{r} - \mathbf{r}') \cdot \mathbf{m}(\mathbf{r}')$$

\mathbf{N} is 3×3 geometric tensor and \mathbf{m} is magnetization


$$\tilde{\mathbf{H}} = -\tilde{\mathbf{N}} \cdot \tilde{\mathbf{M}}$$

$$\tilde{\mathbf{M}}(k'_x, k'_y, k'_z) = \sum_{r'_z=0}^{2N_z-1} \sum_{r'_y=0}^{2N_y-1} \sum_{r'_x=0}^{2N_x-1} \mathbf{m}(r'_x, r'_y, r'_z) \exp \left[\frac{-2\pi j r'_x k'_x}{2N_x} \right] \\ \exp \left[\frac{-2\pi j r'_y k'_y}{2N_y} \right] \exp \left[\frac{-2\pi j r'_z k'_z}{2N_z} \right]$$

$$\begin{aligned} \tilde{\mathbf{H}}_{x(i,j,k)} &= \tilde{\mathbf{N}}_{xx'(i,j,k)} \tilde{\mathbf{M}}_{x(i,j,k)} + \tilde{\mathbf{N}}_{xy'(i,j,k)} \tilde{\mathbf{M}}_{y(i,j,k)} + \tilde{\mathbf{N}}_{xz'(i,j,k)} \tilde{\mathbf{M}}_{z(i,j,k)} \\ \tilde{\mathbf{H}}_{y(i,j,k)} &= \tilde{\mathbf{N}}_{yx'(i,j,k)} \tilde{\mathbf{M}}_{x(i,j,k)} + \tilde{\mathbf{N}}_{yy'(i,j,k)} \tilde{\mathbf{M}}_{y(i,j,k)} + \tilde{\mathbf{N}}_{yz'(i,j,k)} \tilde{\mathbf{M}}_{z(i,j,k)} \\ \tilde{\mathbf{H}}_{z(i,j,k)} &= \tilde{\mathbf{N}}_{zx'(i,j,k)} \tilde{\mathbf{M}}_{x(i,j,k)} + \tilde{\mathbf{N}}_{zy'(i,j,k)} \tilde{\mathbf{M}}_{y(i,j,k)} + \tilde{\mathbf{N}}_{zz'(i,j,k)} \tilde{\mathbf{M}}_{z(i,j,k)} \end{aligned}$$



A magnetized body with non-periodic \mathbf{m} . The size of the magnetic body is doubled along each axis to remove aliasing effects and to make a periodic signal in Fourier space.



(a) Firstly $N_y N_z$ transforms are carried out along the x-axis. (b) In the second stage $(N_x + 1) N_z$ transforms are carried out along the y-axis. The +1 due to conjugate property of FFT's. (c) Finally, $(N_x + 1) 2N_y$ transforms are carried out along the z-axis.

Micromagnetics 2/3: Flowchart of our GPU Field Solver

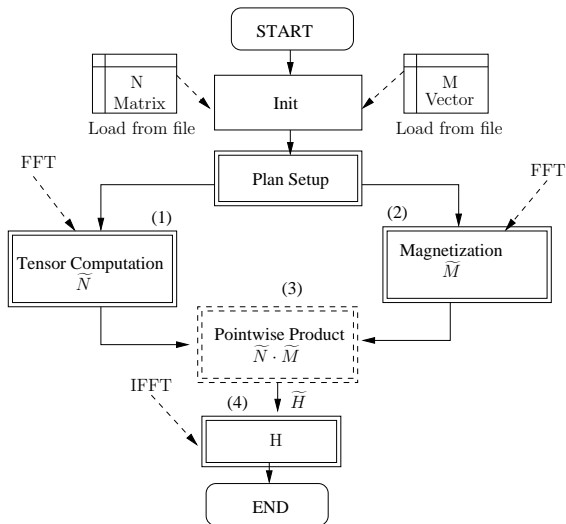


Figure: Design Overview of GPU Magnetostatic Field solver. The double line rectangles show processes performed in parallel by the GPU. The dotted line rectangles show the dot product performed by parallel threads on CPU.

Micromagnetics 3/3: Results and Comparisons




Fig: Total Demag field time GPU-GTX 260 S-Precision with MTT and point-wise multiplication time on CPU




Fig: Total Demag Field Computation speedup against OOMMF

Accuracy

Computation Cells	ϵ_{mean}	ϵ_{max}	η_{mean}
1 Million	$7.59e-11$	$3.52e-10$	$1.01e-14$
2 Million	$4.04e-11$	$1.68e-10$	$5.33e-15$
4 Million	$7.56e-11$	$3.65e-10$	$9.96e-15$
8 Million	$7.91e-11$	$3.38e-10$	$1.04e-14$

Table: Validation of our simulation computing \mathbf{H} against μ -mag Standard problem 4 with an S-state initial magnetization. Here \mathbf{H}' is OOMMF computed. Here, $\epsilon = \|\mathbf{H} - \mathbf{H}'\|_2$, and $\eta = \epsilon / \|\mathbf{H}'\|$

Time Stepping Technique 1/2: Runge Kutta like scheme for the Integration of LLG

- ▶ Numerical solution for Landau-Lifshitz equation of motion for magnetization
- ▶ The fundamental LLG equation:

$$\frac{\partial \mathbf{m}}{\partial t} = -\mathbf{m} \times (\mathbf{h}_{\text{eff}} + \alpha (\mathbf{m} \times \mathbf{h}_{\text{eff}}))$$

- ▶ Discretization using mid-point rule:

$$\mathbf{m}^{k+1}(i) - \mathbf{m}^k(i) = \tau/2 [\mathbf{m}^{k+1}(i) + \mathbf{m}^k(i)] \times \mathcal{H}$$

- ▶ Time-stepping scheme based on the properties of mid-point rule and Runge-Kutta method.
- ▶ Properties of magnetization dynamics are preserved in all steps.
- ▶ At each time-step a 3×3 linear system of equations is solved instead of $3N \times 3N$ in the previous cases
- ▶ It reduces the computations of the \mathbf{h}_{eff} per time-step.

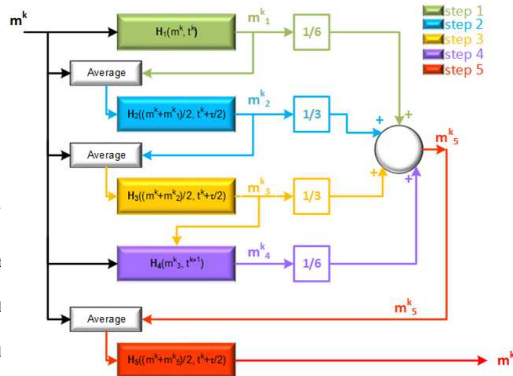




Fig: Flow of different steps.


Time Stepping Technique 2/2: Results Runge Kutta like scheme ...




Comparison with OOMMF using muMAG Standard Problem 4. Plot of spatially average magnetization $\langle \mathbf{m} \rangle$ versus time with constant applied field at an angle of 190° off x-axis and time-step of 1.25 ps.



Plot of $1-\dot{m}_y$ as a function of time for various time-steps, confirming the conservation of magnetization.



Plot of the relative error $\epsilon_\alpha = (\tilde{\alpha} - \alpha)/\alpha$ as a function of time for a 5 nm cell size and various time-steps



Plot of the relative energy for conservative case $\alpha = 0$ as a function of time at various time-steps. The energy is preserved for all time steps.

Conclusion

- ▶ Developed and implemented new sorting algorithms
- ▶ Developed generic FFT library on GPUs
- ▶ GPU accelerated Magnetostatic Field Solver
- ▶ Designed and developed new time integration method for LLG equation
- ▶ Designed new load balancing scheme on multiple GPUs

Publications

Proceedings

- ▶ (2013) B. Jan, B. Montrucchio, C. Ragusa, F.G. KHAN, O.U. KHAN, “Parallel Butterfly Sorting Algorithm”. Proceeding, IASTED-Parallel and Distributed Computing and Networks PDCN-2013 Innsbruck, Austria. DOI:10.2316/P.2013.795-026.
- ▶ (2014) O. Khan, B. Jan, C. Ragusa, A. Rahim, F. Khan, B. Montrucchio, “Optimization of a Mult-Dimensional FFT Library for Accelerating Magnetostatic Field Computations”. In: 10th European Conference on Magnetic Sensors and Actuators, Vienna, Austria, July 6-9, 2014. p.250, ISBN 9783854650218.

Journals

- ▶ (2012) B. Jan, B. Montrucchio, C. Ragusa, F. Khan, O. Khan, “Fast parallel sorting algorithms on GPUs”. In International Journal of Distributed and Parallel Systems, vol 3, pp.107-118. ISSN 2229-3957, DOI:10.5121/ijdp.2012.3609.
- ▶ (2014) A. Rahim, C. Ragusa, B. Jan, O. Khan, “ A mixed mid-point Runge-Kutta like scheme for the integration of LLG”, in Journal of Applied Physics, vol. 115 n.17, 17D101, ISSN 0021-8979.
- ▶ (2015) F.Khan, B. Jan, C.Ragusa, B. Montrucchio and O. Khan, "An Optimized Magnetostatic Field Solver on GPU using OpenCL", in Elsevier journal on "Computers and Electrical Engineering".

Thank You for Your Attention.