Evaluating Graph Coloring on GPUs

Final Project for the GPU class - Spring 2010 submitted as a poster to PPoPP 2011

Pascal Grosset, Peihong Zhu, Shusen Liu, Suresh Venkatasubramanian, and Mary Hall





Existing Parallel Solutions

- · We did not find any relevant related works for graph coloring on GPUs
- Main inspiration:
- Gebremedhin and Manne
- (G-M) algorithm for shared memory architectures 4 stages:
 - Partitioning
 - Pseudo-coloring

 - Conflict detection
 Conflict Resolution

Proposed Framework

- · Adapt existing framework to GPUs
- Phase 1: Graph Partitioning
 o Decide how the Graph will be partitioned into subgraphs
- Phase 2: Graph Coloring & Conflict Identification
 Graph coloring using one of the heuristics
 First Fit, SDO & LDO, Max In, Max Out Conflict Identification
- Phase 3: Sequential Conflict Resolution

 To definitely remove all conflicts

Max In & Max Out

Two new heuristics o Decision parameter: number of vertices having neighbors outside the subgraph while Num_Colored < N (Number of vertices in subgraph) do
 max = -1
for i = 1 to N do
 if !colored(Vi) then
 no = Number of neighbors outside partition
 if no > max then
 max = no
 index = i
 if no == max then
 if d(Vi) > d(Vindex) then
 index = i Color Vindex Num_Colored ++

Phase 2: Graph Coloring & Conflict Identification Main part

- Transfer data from CPU to GPU
- 1. Graph Coloring Run Graph Coloring: 1 thread per subgraph - cudaEventSynchronize -2. Conflicts Indentification
- sets color of conflicted nodes to 0: 1 thread per node
 cudaEventSynchronize -
- Transfer Conflicts to CPU
- Count Conflicts
 - If conflicts < threshold
 - exit $_{\circ}$ Else
 - Repeat from 1

Data Storage

- Adjacency Matrix (Initial) • • Too big
- Adjacency List
- Very compact
 - Bad memory access pattern
 - bad performance

"Sort of" Adjacency List

- o Size of each list: max degree o Good balance between performance and size
 - can still be optimized

.

Phase 2: Graph Coloring & Conflict Identification Main part

- · Transfer data from CPU to GPU
- 1. Graph Coloring Run Graph Coloring: 1 thread per subgraph cudaEventSynchronize -2. Conflicts Indentification
- sets color of conflicted nodes to 0: 1 thread per node - cudaEventSynchronize -• Transfer Conflicts to CPU
- Count Conflicts
 o If conflicts < threshold
- exit
- 。Else Repeat from 1

	Test	Data	а			
Data source:	Jniversity of Flo	orida S	parse N	/latrix (Colle	ection
	Name	n	m	Density	Δ	Avg Degree
	Structural Engineering ct20stif nasasrb pwtk	52,329 54,870 217,918	1,375,396 1,366,097 5,926,171	0.00100 0.00091 0.00025	206 275 179	50 47 52
nasasrb	Civil Engineering pkustk10 pkustk11	81,920 87,804	2,114,154 2,565,054	0.00063 0.00067	89 131	52 55
	pkustk13 Automotive	94,893	3,260,967	0.00072	299	68
	hood	220,542	5,273,947	0.00020	76	47
	ldoor msdoor	958,464 417,792	22,785,136 9,912,536	0.00005 0.00011	76 76	46 46
	Ship section	140.874	2 926 265	0.00020	101	51
	supsect	190,874	4,066,619	0.00039	101	54
pwtk	shipsec8	114,919	3,269,240	0.00050	125	55 54
	Linear Car Analysis bmw3_2	227,362	5,539,634	0.00021	335	49

Benchmarks Sequential Algorithms First-Fit SDO & LDO Implementation direct from H. Al-Omari and K. E. Sabri, "New Graph Coloring Algorithms • O (n³) • up to 1000x speedups!!! Optimized (our) implementation (as a red black Tree) O (m log n) 20x - 40x speedup n: number of vertices, m: number of edges

Implementation

· Some Details:

Г

- $_{\circ}\,$ Tests were carried out on a Tesla S1070 and Fermi GTX 480
- with Caching
 Memory transfer time included in all results
 All results are the average of 10 runs

• Detailed times: Graph = hood; Algo = First Fit

Coloring	22.560	15.0256	2 10456
Coloring	23.305	15.0250	2.19430
Detect	7.1281	6.6528	6.58186
Count Time	0.003488	0.003712	0.00336
			I











Results: Metis vs Non-Metis

- · Naive partitioning is most of the time faster and more efficient than using Metis
 - Metis is not effective for such small partitions
 - Yields unbalanced partitions at such small graph sizes
 Unbalanced is bad for GPU
- Metis was slower despite not taking into account the time for Metis to run! •

Conclusion

- Set of guidelines for graph coloring

 Fastest: Parallel First Fit
 Much better colors than Sequential First Fit
 Slightly slower than Sequential First Fit

 - Best Results (if you do not care about speed)
 Sequential SDO & LDO implemented as a Red-Black Tree
 Balance of Speed and Colors
 - Parallel Max Out or Min Out average of 20X speedup over sequential SDO & LDO
- Use small subgraph Sizes
- Naive partitioning is good
 CUDA does not only makes calculations faster but can also be used to improve results - First Fit!





