### Advanced Numerical Methods on GPUs

#### Dominik Göddeke and Stefan Turek

Institut für Angewandte Mathematik (LS3) TU Dortmund

dominik.goeddeke,stefan.turek@math.tu-dortmund.de
http://www.mathematik.tu-dortmund.de/~goeddeke
http://www.mathematik.tu-dortmund.de/LS3

ENUMATH 2011 Mini-Symposium Leicester, UK, September 7





fakultät für

#### Problems with current hardware

- Memory wall: Data movement cost prohibitively expensive
- Power wall: Nuclear power plant for each machine (in the cloud)?
- ILP wall: 'Automagic' maximum resource utilisation?
- Memory wall + power wall + ILP wall = brick wall

#### Inevitable paradigm shift: Parallelism and heterogeneity

- In a single chip: singlecore  $\rightarrow$  multicore, manycore, ...
- In a workstation (cluster node): NUMA, CPUs and GPUs, ...
- In a big cluster: different nodes, communication characteristics, ...

#### This is our problem as mathematicians

- Affects standard workstations and even laptops
- In most cases, cannot be hidden from us properly

#### Without respecting parallelism

- Impossible to exploit ever increasing peak performance
- Sequential codes even run slower on newer hardware (!)

#### Challenges

- Technical: Compilers can't solve these problems, libraries are limited
- Numerical: Traditional methods often contrary to hardware trends
- Goal: Redesign existing numerical schemes (and invent new ones) to work well in the fine-grained parallel setting

#### GPUs ('manycore') are forerunners of this development

- 10 000s of simultaneously active threads
- Promises of significant speedups
- Focus of this mini-symposium



#### Raw marketing numbers

- $\blacksquare>2~{\rm TFLOP/s}$  peak floating point performance
- Lots of papers claim  $> 100 \times$  speedup

#### Looking more closely

- Single or double precision? Same on both devices?
- Sequential CPU code vs. parallel GPU implementation?
- Standard operations' or many low-precision graphics constructs?

#### Reality

- GPUs are undoubtedly fast, but so are CPUs
- Quite often: CPU codes significantly less carefully tuned
- Anything between  $5 30 \times$  speedup is realistic (and worth the effort)

#### This mini-symposium

- Brief introduction to GPU computing
- Discussion of advanced numerical methods on GPUs
- State-of-the-art examples covering a wide range of numerical methods and applications

#### Session 1 (today): Introduction and toolkits

- 11:00–11:30: Dominik Göddeke: Mini-symposium welcome & introduction to GPU computing
- 11:30–12:00: Dominik Göddeke: Mixed-precision GPU-multigrid solvers with strong smoothers and applications in CFD and CSM
- 12:00–12:30: Mike Giles: OP2: An open-source library for unstructured grid applications
- 12:30: open discussion (or beating the lunch queue)

#### Session 2 (tomorrow): Applications in CFD and CSM

- 11:00–11:30: Martin Geier: EsoStripe – An aligned data-layout for efficient CFD simulations on GPUs using the Lattice Boltzmann Method
- 11:30–12:00: Allan Peter Engsig-Karup: On the development of a GPU-accelerated nonlinear free-surface model for coastal engineering
- 12:00–12:30: Martin Lilleeng Sætra: Shallow water simulations on implicitly defined global grids
- 12:30–13:00: Christian Dick: CUDA FE multigrid with applications in flow/solid mechanics

#### Session 3 (tomorrow): Solvers and preconditioners

- 14:00–14:30: Jan-Philipp Weiß: Fine-grained parallel preconditioners on GPUs and beyond
- 14:30–15:00: Robert Strzodka: GPU bandwidth optimization of preconditioners
- 14:30–15:00: Stephan Kramer: Parallel preconditioning strategies for decoupled indoor air flow simulation

■ 15:00–15:30: Hans Knibbe:

GPU implementation of a Krylov solver preconditioned by a shifted Laplace multigrid method for the Helmholtz equation

## **Introduction to GPU Computing**

Programming Languages and Existing Libraries Pseudorandomly Chosen 'Didactical' Examples

#### Obviously the most general approach

- Often unavoidable when programming for performance
- Not necessarily optimal in terms of programming effort
- Main focus of the work presented in this mini-symposium
- Rationale: When developing new numerical methods, you don't want some 'arbitrary' layer of abstraction hiding things from you

#### Two environments

- CUDA: More mature, bigger 'ecosystem', NVIDIA only
- OpenCL: Vendor-independent, open industry standard
- Interfaces to C/C++, Fortran, Python, .NET, ...
- Important: Hardware abstraction and 'expressiveness' are identical

#### Compilers

- PGI Accelerator Compiler: OpenMP-like code annotations for Fortran and C
- New: Ongoing work to extend/generalise OpenMP to GPUs (!)

#### Frameworks

- PetSc and Trilinos: GPU support in some (important) sub-packages
- HMPP, StarPU, Quark: Load-balancing in heterogeneous systems

#### Standard software with GPU backends

- Matlab: GPU backends for plain Matlab and some toolboxes
- And many more: Mathematica, Ansys, OpenFOAM, ....

#### **Fourier Transforms**

- CUFFT: NVIDIA, part of the CUDA toolkit
- APPML (formerly ACML-GPU): AMD Accelerated Parallel Processing Math Libraries

#### Dense linear algebra

- CUBLAS: NVIDIA's basic linear algebra subprograms
- APPML (formerly ACML-GPU): AMD Accelerated Parallel Processing Math Libraries
- CULA: Third-party LAPACK, matrix decompositions and eigenvalue problems
- MAGMA and PLASMA: BLAS/LAPACK for multicore and manycore (ICL, Tennessee)

#### Sparse linear algebra and solvers

- CUSPARSE: CSR-SpMV (part of the CUDA toolkit)
- CUDPP: Building blocks for some important operations (NVIDIA and UC Davis, open-source)
- CUSP: Krylov subspace methods with simple preconditioners (NVIDIA, open-source)
- Next version of CUSPARSE: ILU(k) preconditioner
- PARDISO: sparse direct solvers

#### My personal two cents

- Structured case 'solved', unstructured case is the challenging one!
- As always in the sparse world: Little to no standardisation

## **GPU Programming Model**

#### **Step 1: Simplification**

Remove caches and hard-wired logic (branch prediction, ...)

#### Step 2: Invest transistors into compute

Add as many of these 'stripped-down' cores to the chip as price/performance/power budgets allow

#### Step 3: 'Beef up' cores by increasing SIMD width

- 16–64 functional units per core (CPUs: 2–4) execute the same instruction in each cycle, one hardware thread per ALU
- Add local shared scratchpad memory and register file

#### Step 4: Tidy up

Add several memory controllers (and graphics-specific circuits)

#### Main difference between CPUs and GPUs

- So far, this design is not spectacularly different
- CPUs are optimised for latency of an individual task
- GPUs are optimised for throughput of many similar tasks

## Key design feature: Hardware scheduler switches (groups of) threads in zero time as soon as one stalls

- Reason for stalls: Off-chip memory transfers (1000+ cycles), instructions mapping to many μ-ops, ...
- Thread creation and management entirely in hardware

#### Leads to programming model

 Code written for one thread, SIMD-isation done by the hardware, with some parameterisation to enable mapping of threads to data

#### High level view

Data parallelism with limited synchronisation and data sharing

#### Key concept: Thread blocks = virtualised multiprocessors

- Note: CUDA terminology, similar in OpenCL
- Batch computations into 'thread blocks'
- Thread blocks resident in one multiprocessor
- Blocks are independent, no guarantee of execution order
- Threads per block specified by the user (32–1024), problem-specific tunable parameter
- Threads in one block may cooperate via cheap barrier synchronisation and shared memory
- Threads from different blocks may only coordinate via global memory, synchronisation only at kernel scope

#### Execution: Warps = SIMD granularity

- Threads in one block are executed in 'warps' of 32, enumerated in natural order
- One instruction per warp (SIMD granularity)
- Warps are independent, no guaranteed execution order
- Scheduler switches to next available warp in case of stall (availability due to finished memory transaction, entire block reaches barrier, ...)
- Threads in one warp may follow different execution paths ('divergence'), resolved by serialisation and thus performance penalty

#### Limited resources

- Register file (32 K 4-byte entries) and shared memory (16–48 kB) are partitioned among all blocks
- Rule of thumb: Ensure at least two resident blocks per multiprocessor for good throughput ('occupancy')

#### Caches

- Small global L2 cache, 768 kB currently
- Tiny L1 cache per multiprocessor, 16–48 kB
- Tiny 'texture cache' per memory controller, optimised for 2D locality

#### Parallel memory system

- 6–10 partitions, round-robin assignment in small chunks of 256 kB
- Access granularity: half-warp, i.e. 16 threads access 16 values
- Hardware may 'coalesce' these parallel accesses into as few as one bulk memory transaction ⇒ crucial for performance
- Requires adhering to strict rules for memory access patterns of neighbouring threads
- Avoid 'partition camping', i.e. data layouts which map accesses to only one partition

#### Shared memory

- 16–48 kB 'scratchpad memory' per multiprocessor
- Can be used as a manually controlled cache
- Common use case: Stage off-chip transfers through this memory to achieve better coalescing (different threads load data than compute on it)
- Access granularity: half-warp (16 threads)
- Physically implemented as 16-bank memory, each bank services one request at a time
- 'Bank conflicts': Simultaneous requests map to only one bank, resulting in serialisation and thus up to 16-fold slowdown

#### **GPUs** ...

- are wide-SIMD manycore architectures
- are parallel on all levels (compute and memory)
- operate in a block-threaded way

#### **GPUs** are not

- Vector architectures (rather wide-SIMD+multithread)
- Fully task-parallel (performance stems from data parallelism)
- Easy to program efficiently (getting things running is easy though)

#### GPUs are particularly bad at

- Pointer chasing through memory (serialisation of memory accesses)
- Codes with lots of fine-granular branches
- Codes with lots of synchronisation and huge sequential portions

#### Parallelism and heterogeneity are inevitable

- GPUs are prominent fore-runners of this trend
- Necessary development of novel numerical methods that are better suited for the hardware: hardware-oriented numerics

#### **GPU Architecture**

- Tricky at first, especially in this crash coarse
- But: Learning curve is not so steep if one is familiar with performance tuning for CPUs

#### Active research topic

- 'Structured' cases pretty much solved, irregular and (at first sight) inherently sequential ones are challenging
- Algorithmic research required rather than focus on implementational details

## Mixed-Precision GPU-Multigrid Solvers with Strong Smoothers and Applications in CFD and CSM

#### Dominik Göddeke and Robert Strzodka

Institut für Angewandte Mathematik (LS3), TU Dortmund Max Planck Institut Informatik, Saarbrücken

dominik.goeddeke@math.tu-dortmund.de http://www.mathematik.tu-dortmund.de/~goeddeke

ENUMATH 2011 Mini-Symposium Leicester, UK, September 7







#### **Conflicting situations**

- Existing methods no longer hardware-compatible
- Neither want less numerical efficiency, nor less hardware efficiency

#### Challenge: New algorithmic way of thinking

Balance these conflicting goals

#### Consider short-term hardware details in actual implementations, but long-term hardware trends in the design of numerical schemes

- Locality, locality, locality
- Communication-avoiding (-delaying) algorithms between all flavours of parallelism
- Multilevel methods, hardware-aware preconditioning

# **Grid and Matrix Structures Flexibility** $\leftrightarrow$ **Performance**

#### General sparse matrices (unstructured grids)

- CSR (and variants): General data structure for arbitrary grids
- Maximum flexibility, but during SpMV
  - Indirect, irregular memory accesses
  - Index overhead reduces already low arithm. intensity further
- Performance depends on nonzero pattern (grid numbering)

#### Structured sparse matrices

- Example: Structured grids, suitable numbering  $\Rightarrow$  band matrices
- Important: No stencils, fully variable coefficients
- Direct regular memory accesses, fast independent of mesh
- 'FEAST patches': Exploitation in the design of strong MG components

### Example: Poisson on unstructured mesh



- Nehalem vs. GT200, ≈ 2M bilinear FE, MG-JAC solver
- Unstructured formats highly numbering-dependent
- Multicore 2–3x over singlecore, GPU 8–12x over multicore
- Banded format (here: 8 'blocks') 2–3x faster than best unstructured layout and predictably on par with multicore

## **Strong Smoothers**

Parallelising Inherently Sequential Operations Test case: Generalised Poisson problem with anisotropic diffusion

- $-\nabla \cdot (\mathbf{G} \ \nabla \mathbf{u}) = \mathbf{f}$  on unit square (one FEAST patch)
- $\blacksquare~{\bf G}={\bf I}:$  standard Poisson problem,  ${\bf G}\neq {\bf I}:$  arbitrarily challenging
- Example: G introduces anisotropic diffusion along some vector field



Only multigrid with a strong smoother is competitive

Disclaimer: Not necessarily a good smoother, but a good didactical example.

#### Sequential algorithm

- Forward elimination, sequential dependencies between matrix rows
- Illustrative: Coupling to the left and bottom

#### 1st idea: Classical wavefront-parallelisation (exact)



- Pro: Always works to resolve *explicit* dependencies
- Con: Irregular parallelism and access patterns, implementable?

#### 2nd idea: Decouple dependencies via multicolouring (inexact)

Jacobi (red) – coupling to left (green) – coupling to bottom (blue) – coupling to left and bottom (yellow)



#### Analysis

- Parallel efficiency: 4 sweeps with  $\approx N/4$  parallel work each
- Regular data access, but checkerboard pattern challenging for SIMD/GPUs due to strided access
- Numerical efficiency: Sequential coupling only in last sweep

#### 3rd idea: Multicolouring = renumbering

- After decoupling: 'Standard' update (left+bottom) is suboptimal
- Does not include all already available results



- Recoupling: Jacobi (red) coupling to left and right (green) top and bottom (blue) – all 8 neighbours (yellow)
- More computations that standard decoupling
- Experiments: Convergence rates of sequential variant recovered (in absence of preferred direction)

## Tridiagonal smoother (line relaxation)

#### Starting point

- Good for 'line-wise' anisotropies
- 'Alternating Direction Implicit (ADI)' technique alternates rows and columns
- CPU implementation: Thomas-Algorithm (inherently sequential)

	0	1	2	3	4	5	6	7	8	9	10	11
0	х	х				х	х					
1	х	х	х			х	х	х				
2		х	х	х			х	х	х			
3			х	х	х			х	х	х		
4				х	х				х	х		
5	х	х				х	х				х	х
6	х	х	х			х	х	х			х	X
7		х	х	х			х	х	х			X
8			х	х	х			х	х	х		÷.
9				х	х				х	х		
10	)					х	х				х	X
11						х	х	х			х	x
								S. 1				

#### Observations

- One independent tridiagonal system per mesh row
- $\blacksquare$   $\Rightarrow$  top-level parallelisation across mesh rows
- Implicit coupling: Wavefront and colouring techniques not applicable

#### Cyclic reduction for tridiagonal systems

- Exact, stable (w/o pivoting) and cost-efficient
- Problem: Classical formulation parallelises computation but not memory accesses on GPUs (bank conflicts in shared memory)
- Developed a better formulation, 2-4x faster
- Index challenge, general idea: Recursive padding between odd and even indices on all levels



#### **Starting point**

- CPU implementation: Shift previous row to RHS and solve remaining tridiagonal system with Thomas-Algorithm
- Combined with ADI, this is the best general smoother (we have) for this matrix structure

	0	1	2	3	4	5	6	7	8	9	10	11
0	х	х				х	х					
1	х	х	х			х	х	х				
2		х	х	х			х	х	х			
3			х	х	х			х	х	х		
4				х	х				х	х		
5	х	х				х	х				х	х
6	х	х	х			х	х	х			х	Χ.,
7		х	х	х			х	х	х			X
8			х	х	х			х	х	х		2
9				х	х				х	х		
10	)					х	х				х	×
11						х	х	х			х	X

#### **Observations and implementation**

- Difference to tridiagonal solvers: Mesh rows depend sequentially on each other
- Use colouring  $(\#c \ge 2)$  to decouple the dependencies between rows (more colours = more similar to sequential variant)

### Evaluation: Total efficiency on CPU and GPU

Test problem: Generalised Poisson with anisotropic diffusion

- Total efficiency: ( $\mu s$  per unknown per digit)<sup>-1</sup>
- Mixed precision iterative refinement multigrid solver
- Intel Westmere vs. NVIDIA Fermi





#### Summary: Smoother parallelisation

- Factor 10-30 (dep. on precision and smoother selection) speedup over already highly tuned CPU implementation
- Same numerical capabilities on CPU and GPU
- Balancing of numerical and parallel efficiency (hardware-oriented numerics)

CSM and CFD on GPU-Accelerated Clusters

#### Combination of structured and unstructured advantages

- Global macro-mesh: Unstructured, flexible, complex domains
- Local micro-meshes: Structured (logical TP-structure), fast
- Important: Structured ≠ simple meshes!



#### Hybrid multilevel domain decomposition method

- Multiplicative between levels, global coarse grid problem (MG-like)
- Additive horizontally: block-Jacobi / Schwarz smoother (DD-like)
- Local GPU-accelerated MG hides local irregularities

$$\begin{pmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{pmatrix} \begin{pmatrix} \mathbf{u}_1 \\ \mathbf{u}_2 \end{pmatrix} = \mathbf{f}$$

$$\begin{pmatrix} (2\mu+\lambda)\partial_{xx}+\mu\partial_{yy} & (\mu+\lambda)\partial_{xy} \\ (\mu+\lambda)\partial_{yx} & \mu\partial_{xx}+(2\mu+\lambda)\partial_{yy} \end{pmatrix}$$



## Speedup



- USC cluster in Los Alamos, 16 dualcore nodes (Opteron Santa Rosa, Quadro FX5600)
- Problem size 128 M DOF
- Dualcore 1.6x faster than singlecore (memory wall)
- GPU 2.6x faster than singlecore, 1.6x than dualcore

## Speedup analysis

#### Theoretical model of expected speedup

- Integration of GPUs increases resources
- Correct model: Strong scaling within each node
- Acceleration potential of the elasticity solver:  $R_{acc} = 2/3$  (remaining time in MPI and the outer solver)

$$\label{eq:max} {\rm I\hspace{-.1cm} S_{max}} = \frac{1}{1-R_{\rm acc}} \qquad \qquad S_{\rm model} = \frac{1}{(1-R_{\rm acc})+(R_{\rm acc}/S_{\rm local})}$$

#### This example



### Weak scalability

#### Simultaneous doubling of problem size and resources

- Left: Poisson, 160 dual Xeon / FX1400 nodes, max. 1.3 B DOF
- Right: Linearised elasticity, 64 nodes, max. 0.5 B DOF



#### Results

- No loss of weak scalability despite local acceleration
- 1.3 billion unknowns (no stencil!) on 160 GPUs in less than 50 s

## Stationary laminar flow (Navier-Stokes)

$$\begin{pmatrix} \mathbf{A_{11}} & \mathbf{A_{12}} & \mathbf{B_1} \\ \mathbf{A_{21}} & \mathbf{A_{22}} & \mathbf{B_2} \\ \mathbf{B}_1^\mathsf{T} & \mathbf{B}_2^\mathsf{T} & \mathbf{C} \end{pmatrix} \begin{pmatrix} \mathbf{u_1} \\ \mathbf{u_2} \\ \mathbf{p} \end{pmatrix} = \begin{pmatrix} \mathbf{f_1} \\ \mathbf{f_2} \\ \mathbf{g} \end{pmatrix}$$

#### fixed point iteration

assemble linearised subproblems and solve with global BiCGStab (reduce initial residual by 1 digit) Block-Schurcomplement preconditioner

1) approx. solve for velocities with **global MG** (V1+0), additively smoothed by

for all  $\Omega_i$ : solve for  $\mathbf{u}_1$  with local MG

for all  $\Omega_i$ : solve for  $\mathbf{u}_2$  with **local MG** 

2) update RHS: 
$$\mathbf{d}_3 = -\mathbf{d}_3 + \mathbf{B}^{\mathsf{T}}(\mathbf{c}_1, \mathbf{c}_2)^{\mathsf{T}}$$
  
3) scale  $\mathbf{c}_3 = (\mathbf{M}_n^{\mathsf{L}})^{-1}\mathbf{d}_3$ 





magnitude of velocity + coarse grid

#### Solver configuration

- Driven cavity: Jacobi smoother sufficient
- Channel flow: ADI-TRIDI smoother required

#### Speedup analysis

	$R_{i}$	acc	$S_{loc}$	ocal	$S_{total}$	
	L9	L10	L9	L10	L9	L10
DC Re250	52%	62%	9.1x	24.5x	1.63x	2.71x
Channel flow	48%	_	12.5x	-	1.76x	-

FE assembly vs. linear solver, max. problem size

DC F	Re250	Channel				
CPU	GPU	CPU	GPU			
12:88	31:67	38:59	<b>68:28</b>			

## Summary

#### Grid and data layouts

ScaRC approach: locally structured, globally unstructured

#### **GPU** computing

- Parallelising numerically strong recursive smoothers
- More than an order of magnitude speedup

#### Scale-out to larger clusters

- Minimally invasive integration
- Good speedup despite 'Amdahl's Law'
- Excellent weak scalability
- One GPU code to accelerate CSM and CFD applications built on top of ScaRC

#### Collaborative work with

- FEAST group (TU Dortmund): Ch. Becker, S.H.M. Buijssen, M. Geveler, D. Göddeke, M. Köster, D. Ribbrock, Th. Rohkämper, S. Turek, H. Wobker, P. Zajac
- Robert Strzodka (Max Planck Institut Informatik)
- Jamaludin Mohd-Yusof, Patrick McCormick (Los Alamos National Laboratory)

#### Supported by

- DFG: TU 102/22-1, TU 102/22-2
- BMBF: HPC Software f
  ür skalierbare Parallelrechner. SKALB project 01IH08003D

#### Papers

http://www.mathematik.tu-dortmund.de/~goeddeke