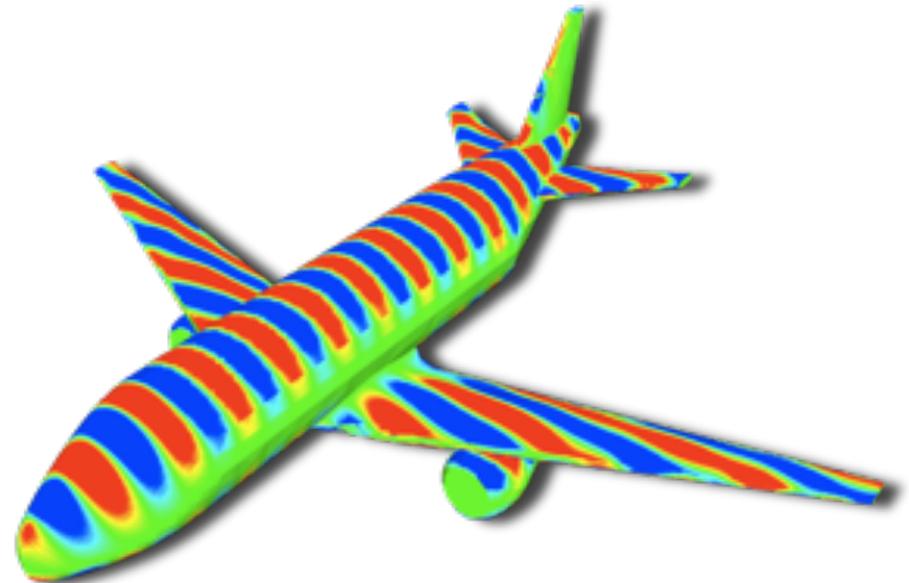
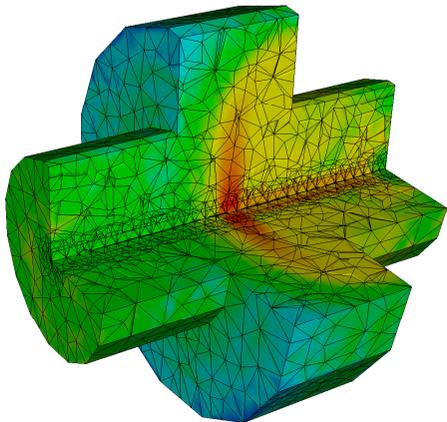

DG-FEM for PDE's

Lecture 8

Jan S Hesthaven
Brown University
Jan.Hesthaven@Brown.edu



A brief overview of what's to come

- Lecture 1: Introduction and DG-FEM in 1D
- Lecture 2: Implementation and numerical aspects
- Lecture 3: Insight through theory
- Lecture 4: Nonlinear problems
- Lecture 5: Extension to two spatial dimensions
- Lecture 6: Introduction to mesh generation
- Lecture 7: Higher order/Global problems
- **Lecture 8: 3D and advanced topics**

Lecture 8

- ✓ Let's briefly recall what we know
- ✓ Part I: 3D problems and extensions
 - ✓ Formulations and examples
 - ✓ Adaptivity and curvilinear elements
- ✓ Part II: The need for speed
 - ✓ Parallel computing
 - ✓ GPU computing
 - ✓ Software beyond Matlab

Lets summarize

We are done with all the basics -- and we have started to see it work for us -- we know how to do

- ✓ 1D/2D problems
- ✓ Linear/nonlinear problems
- ✓ First and higher operators
- ✓ Complex geometries
- ✓ ... and we have insight into theory

All we need is 3D -- and with that comes
the need for speed !

Extension to 3D ?

It is really simple at this stage !

Weak form:

$$\int_{D^k} \left[\frac{\partial u_h^k}{\partial t} \ell_n^k(\mathbf{x}) - \mathbf{f}_h^k \cdot \nabla \ell_n^k(\mathbf{x}) \right] d\mathbf{x} = - \oint_{\partial D^k} \hat{\mathbf{n}} \cdot \mathbf{f}^* \ell_n^k(\mathbf{x}) d\mathbf{x},$$

Strong form:

$$\int_{D^k} \left[\frac{\partial u_h^k}{\partial t} + \nabla \cdot \mathbf{f}_h^k \right] \ell_n^k(\mathbf{x}) d\mathbf{x} = \oint_{\partial D^k} \hat{\mathbf{n}} \cdot \left[\mathbf{f}_h^k - \mathbf{f}^* \right] \ell_n^k(\mathbf{x}) d\mathbf{x},$$

$$\mathbf{f}^* = \{ \{ \mathbf{f}_h(\mathbf{u}_h) \} \} + \frac{C}{2} \llbracket \mathbf{u}_h \rrbracket. \quad C = \max_u \left| \lambda \left(\hat{\mathbf{n}} \cdot \frac{\partial \mathbf{f}}{\partial \mathbf{u}} \right) \right|,$$

Nothing is essential new

Extension to 3D

Apart from the 'logistics' all we need to worry about is to choose our element and how to represent the solution

$$u(\mathbf{r}) \simeq u_h(\mathbf{r}) = \sum_{n=1}^{N_p} \hat{u}_n \psi_n(\mathbf{r}) = \sum_{i=1}^{N_p} u(\mathbf{r}_i) \ell_i(\mathbf{r}),$$

$$\mathbf{u} = \mathcal{V} \hat{\mathbf{u}}, \quad \mathcal{V}^T \boldsymbol{\ell}(\mathbf{r}) = \boldsymbol{\psi}(\mathbf{r}), \quad \mathcal{V}_{ij} = \psi_j(\mathbf{r}_i).$$

We need points

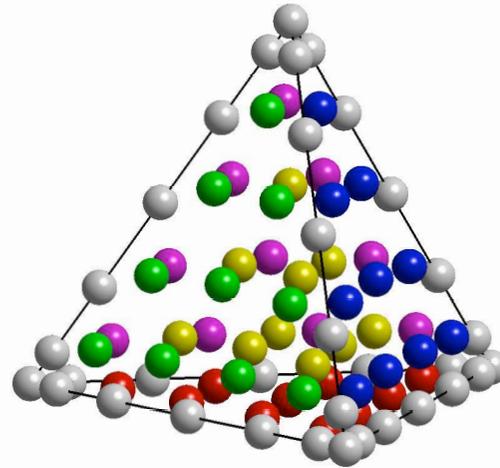
$$N_p = \frac{(N+1)(N+2)(N+3)}{6},$$

We need an orthonormal basis

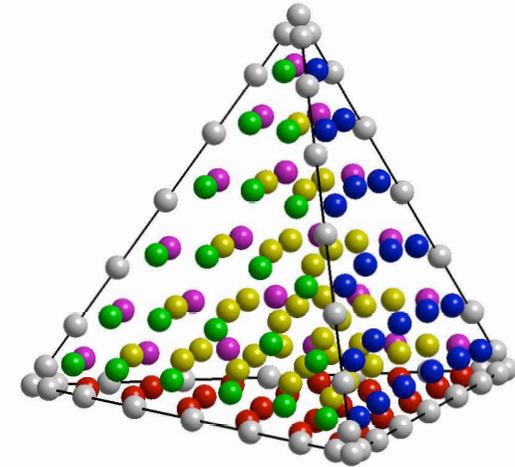
$$\psi_{ijk}(r, s, t) = 2\sqrt{2} P_i^{(0,0)}(a) P_j^{(2i+1,0)}(b) P_k^{(2i+2j+2,0)}(c) (1-b)^i (1-c)^{i+j},$$

Extension to 3D

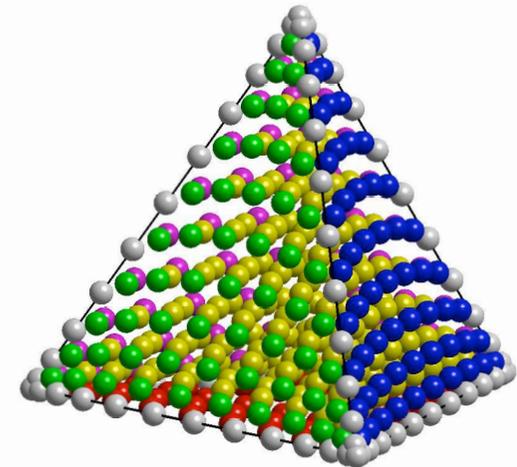
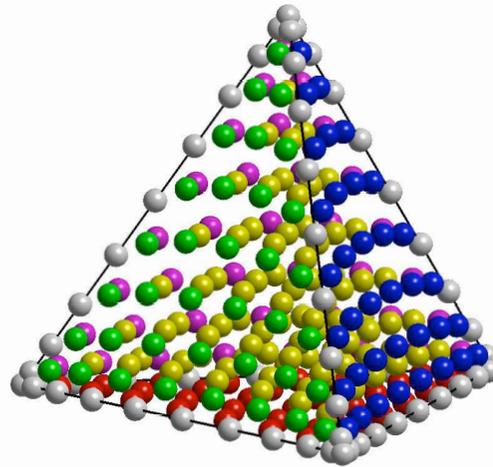
For other element types, one simply need to define nodes and modes for that elements



$N = 6$



$N = 8$



Extension to 3D

Everything is identical in spirit

Mass matrix

$$\mathcal{M}^k = J^k (\mathcal{V}\mathcal{V}^T)^{-1}.$$

Diff matrix

$$\mathcal{D}_r \mathcal{V} = \mathcal{V}_r, \quad \mathcal{D}_s \mathcal{V} = \mathcal{V}_s, \quad \mathcal{D}_t \mathcal{V} = \mathcal{V}_t,$$

Derivative

$$\frac{\partial}{\partial x} = \frac{\partial r}{\partial x} \mathcal{D}_r + \frac{\partial s}{\partial x} \mathcal{D}_s + \frac{\partial t}{\partial x} \mathcal{D}_t,$$

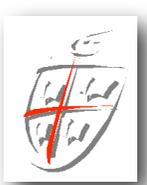
$$\frac{\partial}{\partial y} = \frac{\partial r}{\partial y} \mathcal{D}_r + \frac{\partial s}{\partial y} \mathcal{D}_s + \frac{\partial t}{\partial y} \mathcal{D}_t,$$

$$\frac{\partial}{\partial z} = \frac{\partial r}{\partial z} \mathcal{D}_r + \frac{\partial s}{\partial z} \mathcal{D}_s + \frac{\partial t}{\partial z} \mathcal{D}_t,$$

Stiffness matrix

$$\mathcal{S}_r = \mathcal{M}^{-1} \mathcal{D}_r, \quad \mathcal{S}_s = \mathcal{M}^{-1} \mathcal{D}_s, \quad \mathcal{S}_t = \mathcal{M}^{-1} \mathcal{D}_t.$$

Example - Maxwell's equations



Consider Maxwell's equations

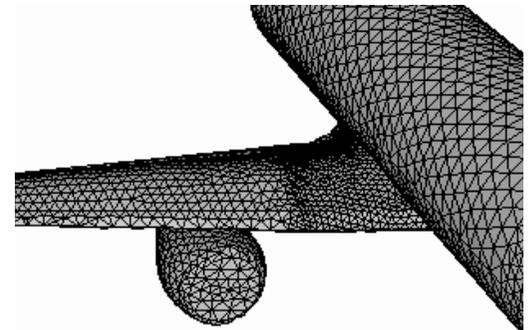
$$\varepsilon \partial_t E - \nabla \times H = -j, \quad \mu \partial_t H + \nabla \times E = 0,$$

Write it on conservation form as

$$\frac{\partial q}{\partial t} + \nabla \cdot F = -J \quad F = \begin{bmatrix} -\hat{e} \times H \\ \hat{e} \times E \end{bmatrix} \quad q = \begin{bmatrix} E \\ H \end{bmatrix}$$

Represent the solution as

$$\Omega = \sum_k D^k \quad q_N = \sum_{i=1}^N q(\mathbf{x}_i, t) L_i(\mathbf{x})$$



and assume

$$\int_D \left(\frac{\partial \mathbf{q}_N}{\partial t} + \nabla \cdot \mathbf{F}_N - \mathbf{J}_N \right) L_i(\mathbf{x}) \, d\mathbf{x} = \oint_{\partial D} L_i(\mathbf{x}) \hat{\mathbf{n}} \cdot [\mathbf{F}_N - \mathbf{F}^*] \, d\mathbf{x}.$$

Example - Maxwell's equations



On each element we then define

$$\hat{M}_{ij} = \int_D L_i L_j \, d\mathbf{x}, \quad \hat{S}_{ij} = \int_D \nabla L_j L_i \, d\mathbf{x}, \quad \hat{F}_{ij} = \oint_{\partial D} L_i L_j \, d\mathbf{x},$$

With the numerical flux given as

$$\hat{\mathbf{n}} \cdot [\mathbf{F} - \mathbf{F}^*] = \begin{cases} \mathbf{n} \times (\gamma \mathbf{n} \times [\mathbf{E}] - [\mathbf{B}]), \\ \mathbf{n} \times (\gamma \mathbf{n} \times [\mathbf{B}] + [\mathbf{E}]), \end{cases} \quad [Q] = Q^- - Q^+$$

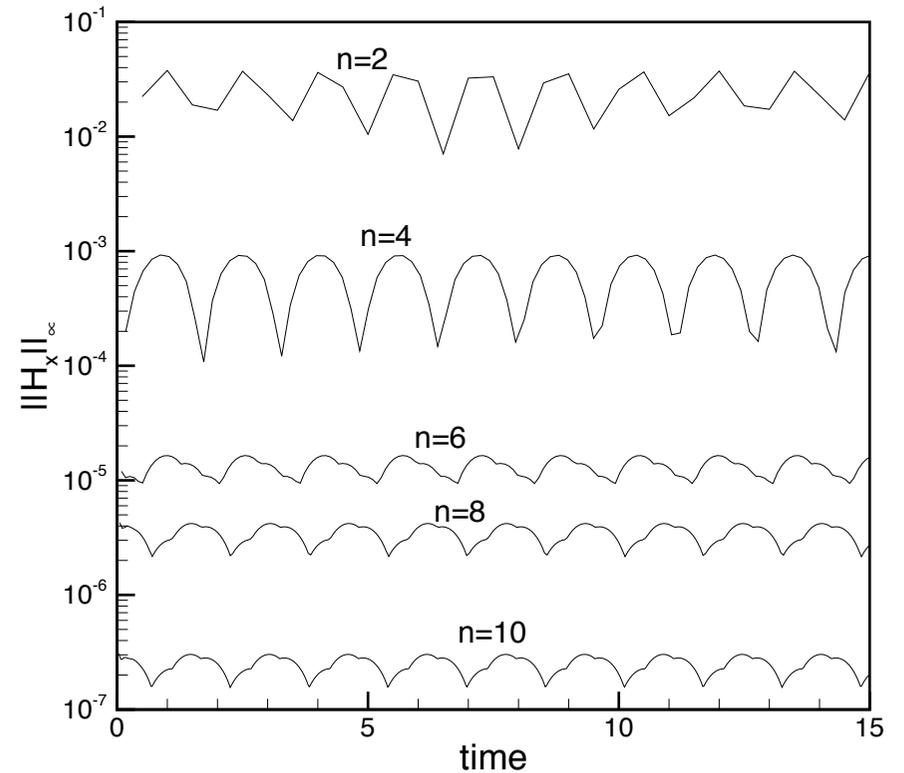
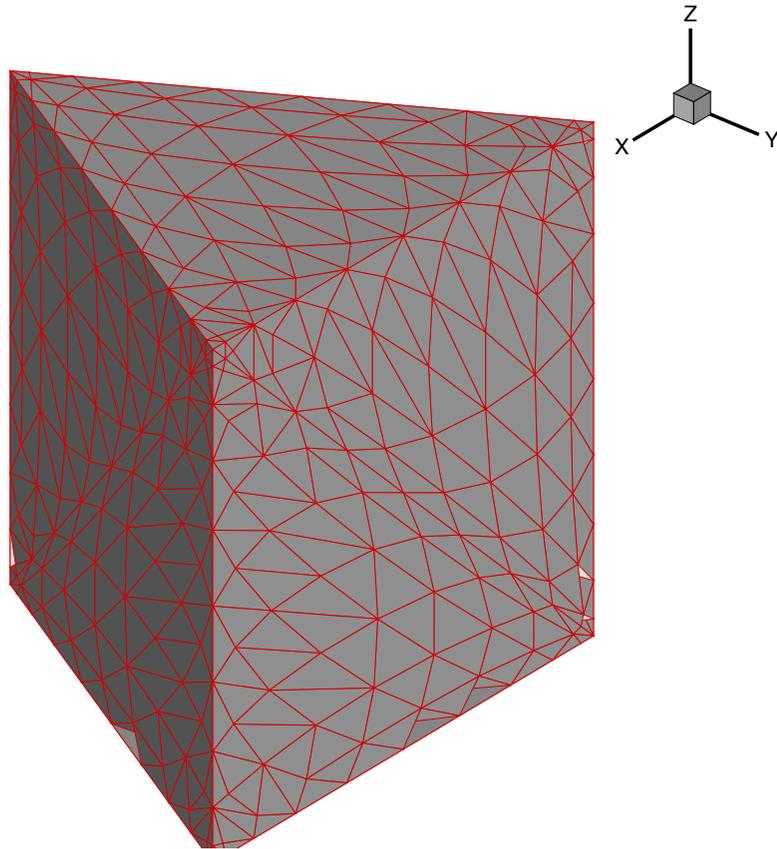
To obtain the local matrix based scheme

$$\hat{M} \frac{d\hat{\mathbf{q}}}{dt} + \hat{S} \cdot \hat{\mathbf{F}} - \hat{M} \hat{\mathbf{J}} = \hat{F} \hat{\mathbf{n}} \cdot [\hat{\mathbf{F}} - \hat{\mathbf{F}}^*],$$

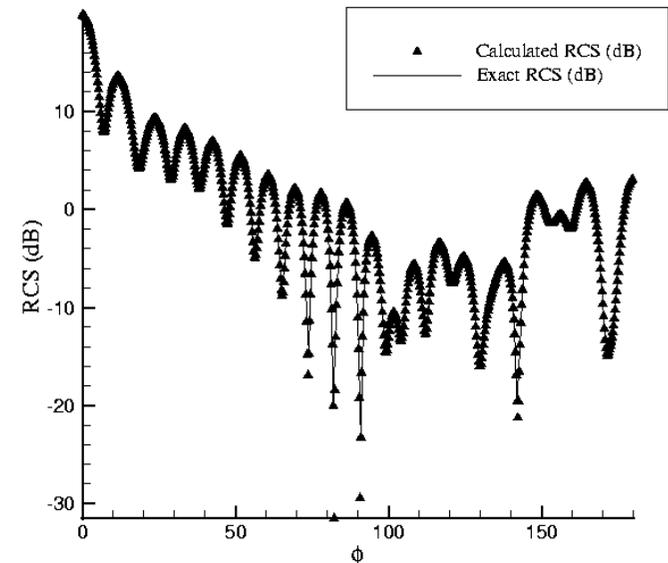
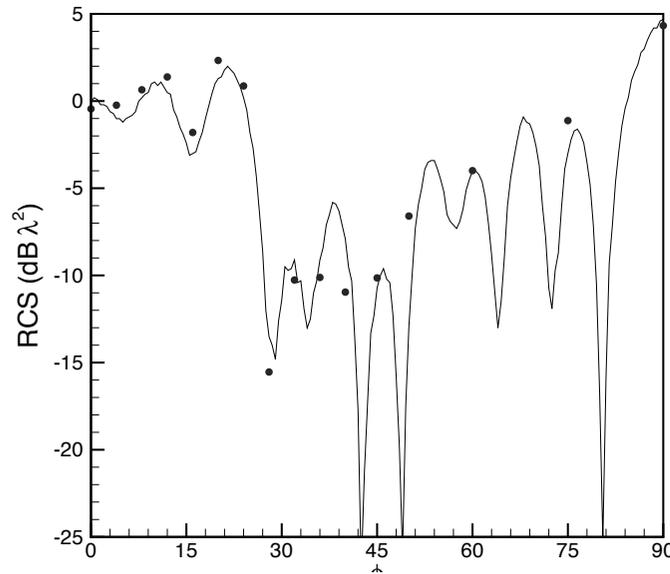
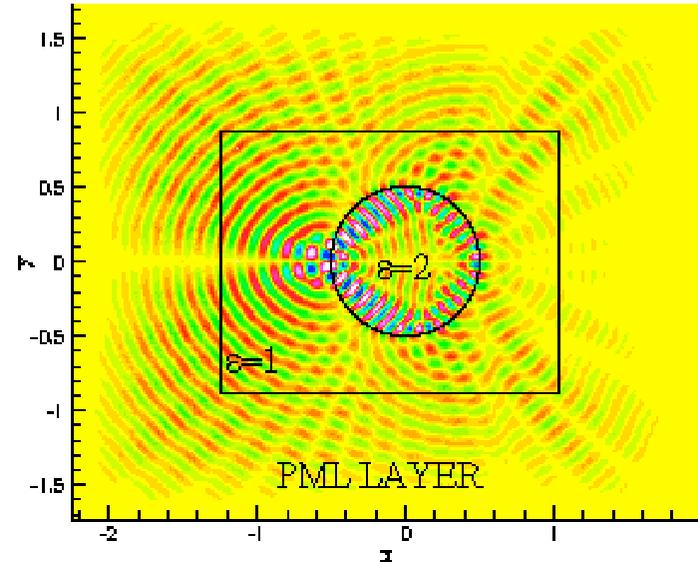
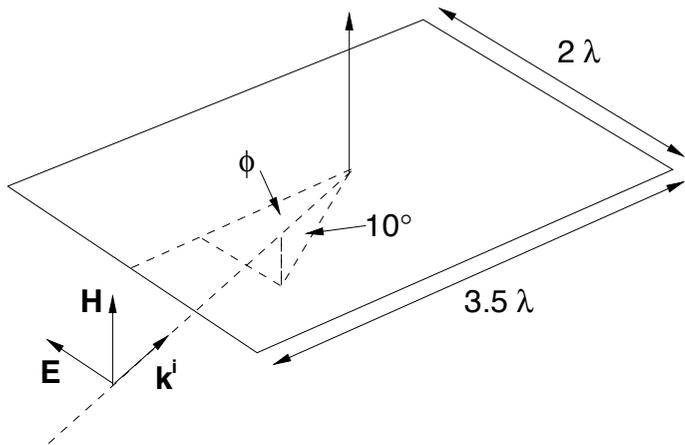
One then typically uses an explicit Runge-Kutta to advance in time - just like 1D/2D.

An example - Maxwell's equations

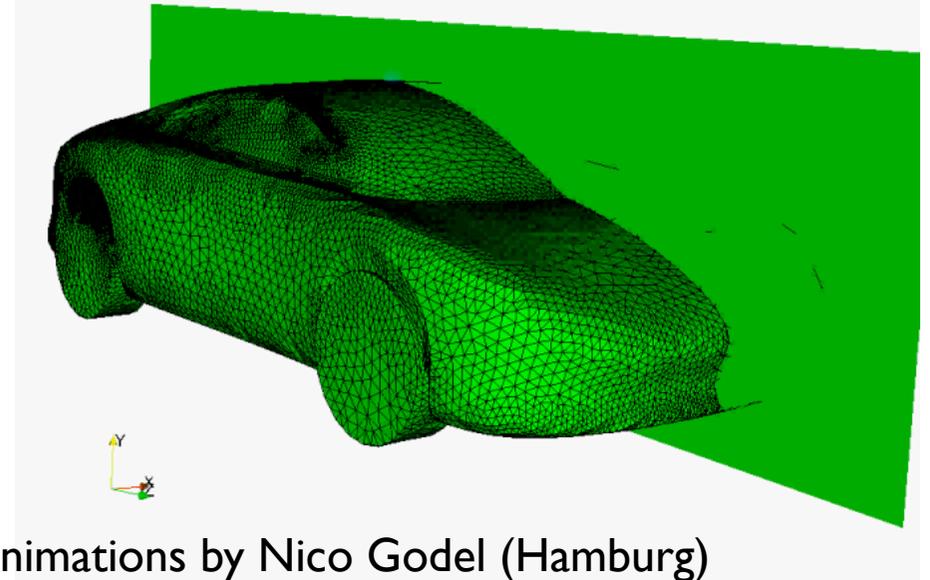
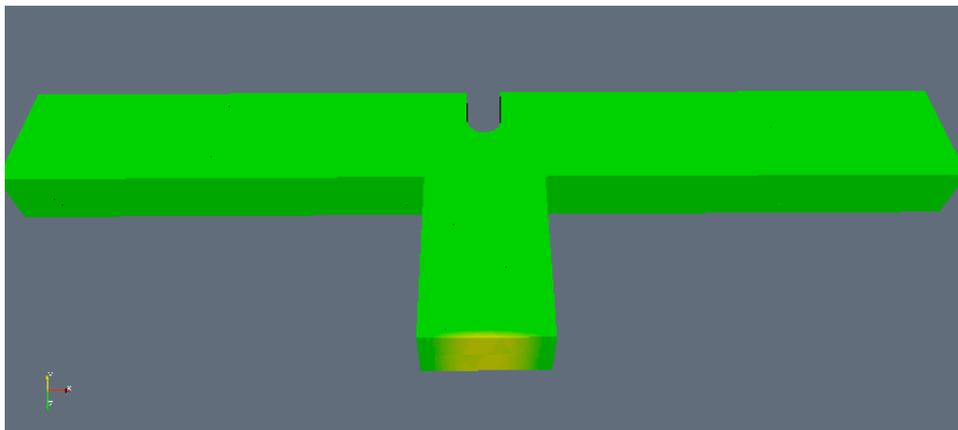
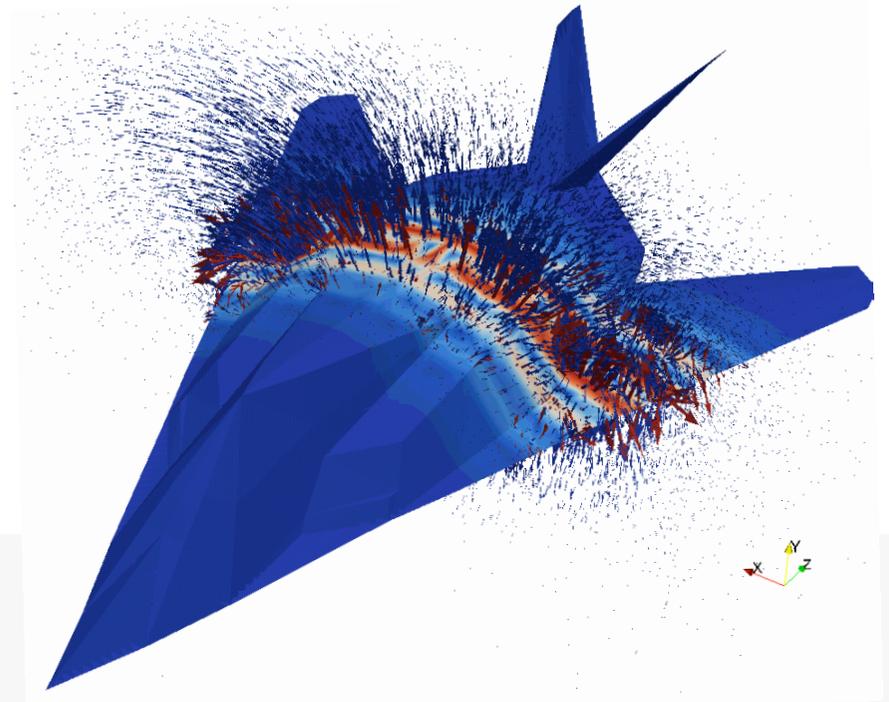
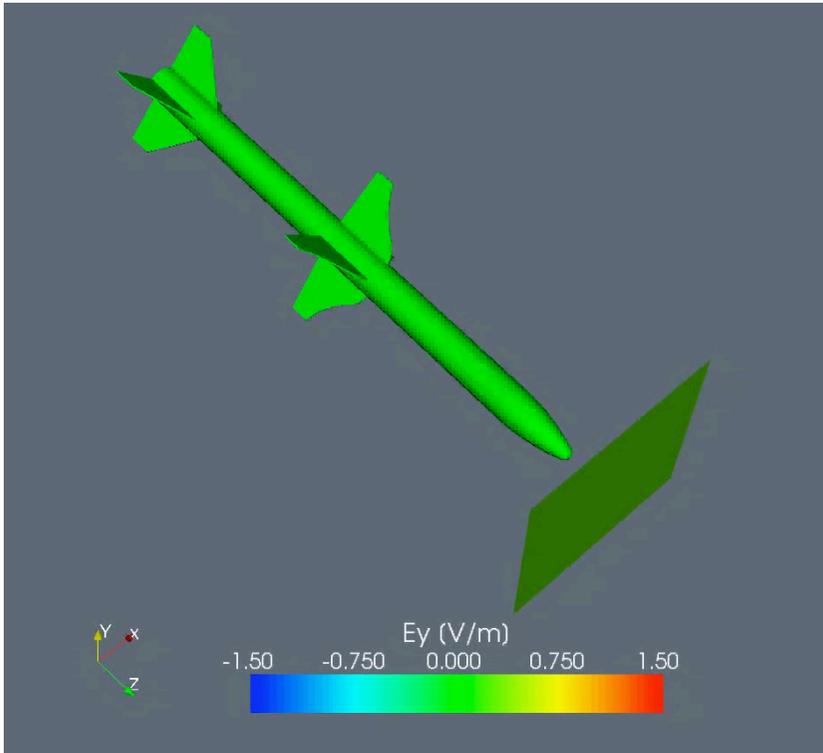
Simple wave propagation



An example - Maxwell's equations



An example - Maxwell's equations

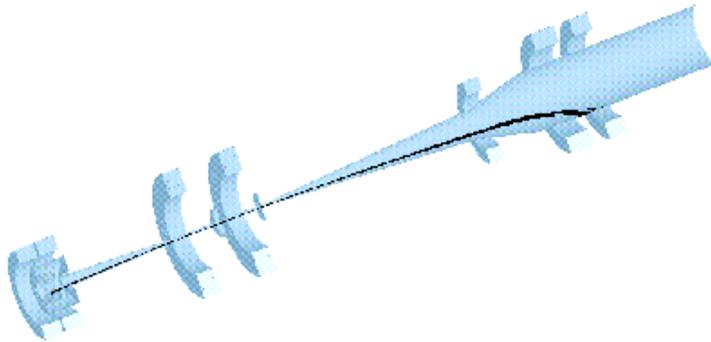


Animations by Nico Godel (Hamburg)

Kinetic Plasma Physics

Important applications

- ✓ High-power/High-frequency microwave generation
- ✓ Particle accelerators
- ✓ Laser-matter interaction
- ✓ Fusion applications, e.g., plasma edge
- ✓ etc



Kinetic Plasma Physics

In high-speed plasma problems dominated by kinetic effects, one needs to solve for $f(x,p,t)$ - $6D+1$

Vlasov/Boltzmann equation

$$\partial_t f + v \cdot \partial_x f + q(E + v \times B) \cdot \partial_p f = \langle \text{Sources} \rangle - \langle \text{Sinks} \rangle.$$

Maxwell's equations

$$\partial_t E - \frac{1}{\epsilon} \nabla \times H = -\frac{j}{\epsilon},$$

$$\partial_t H + \frac{1}{\mu} \nabla \times E = 0,$$

$$\nabla \cdot H = 0, \quad \nabla \cdot E = \frac{\rho}{\epsilon}.$$

Coupled through $\rho := \int f \, dv, \quad j := \int v f \, dv.$

Particle-in-Cell (PIC) Methods

This is an attempt to solve the Vlasov/Boltzmann equation by sampling with P particles

$$f(x, p, t) = \sum_{n=1}^P q_n S(x - x_n(t)) \delta(p - p_n(t)),$$

$$\rho(x, t) = \sum_{n=1}^P q_n S(x - x_n(t)), \quad j(x, t) = \sum_{n=1}^P v_n q_n S(x - x_n(t))$$

Ideally we have

$$S(x) = \delta(x) \quad \longleftarrow \quad \text{a point particle}$$

However, this is not practical, nor reasonable - so $S(x)$ is a **shape-function**

Particle-in-Cell Methods

Maxwell's equations

$$\begin{aligned}\varepsilon \partial_t E - \nabla \times H &= -j, & \mu \partial_t H + \nabla \times E &= 0, \\ \nabla \cdot (\varepsilon E) &= \rho, & \nabla \cdot (\mu H) &= 0,\end{aligned}$$

Particle/Phase dynamics

$$\frac{dx_n}{dt} = v_n(t) \quad \frac{dmv_n}{dt} = q_n(E + v_n \times H) \quad m = \frac{1}{\sqrt{1 - (v_n/c)^2}}$$

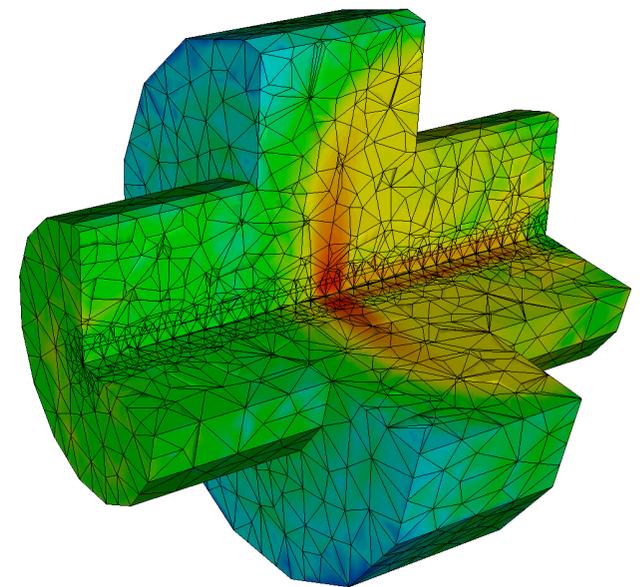
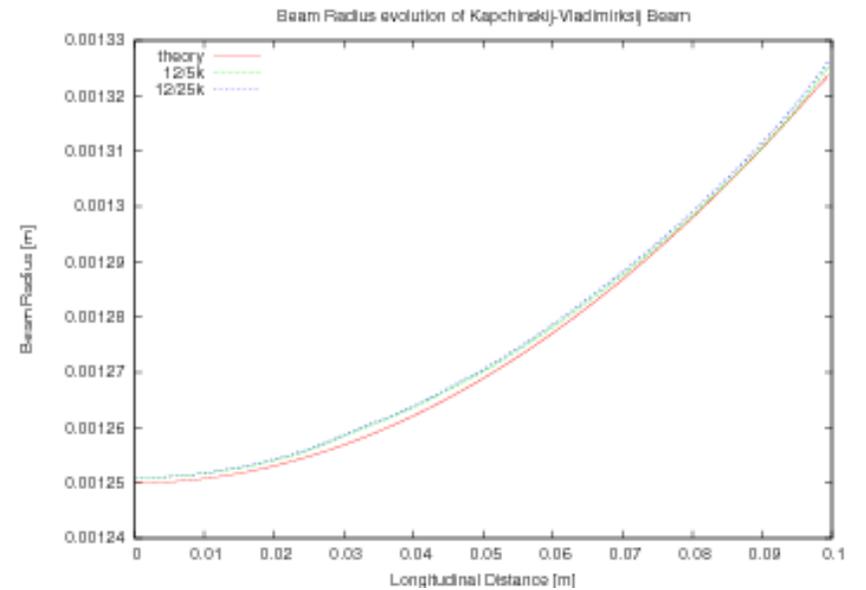
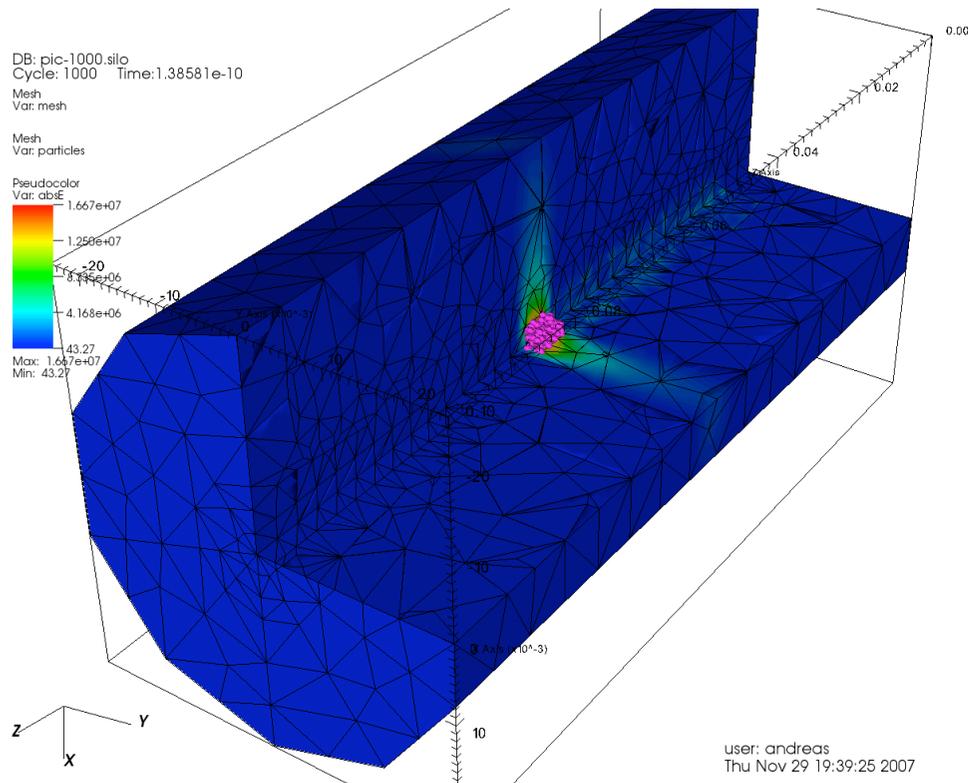
Particles-to-fields

$$\rho(x, t) = \sum_{n=1}^P q_n S(x - x_n(t)), \quad j(x, t) = \sum_{n=1}^P v_n q_n S(x - x_n(t))$$

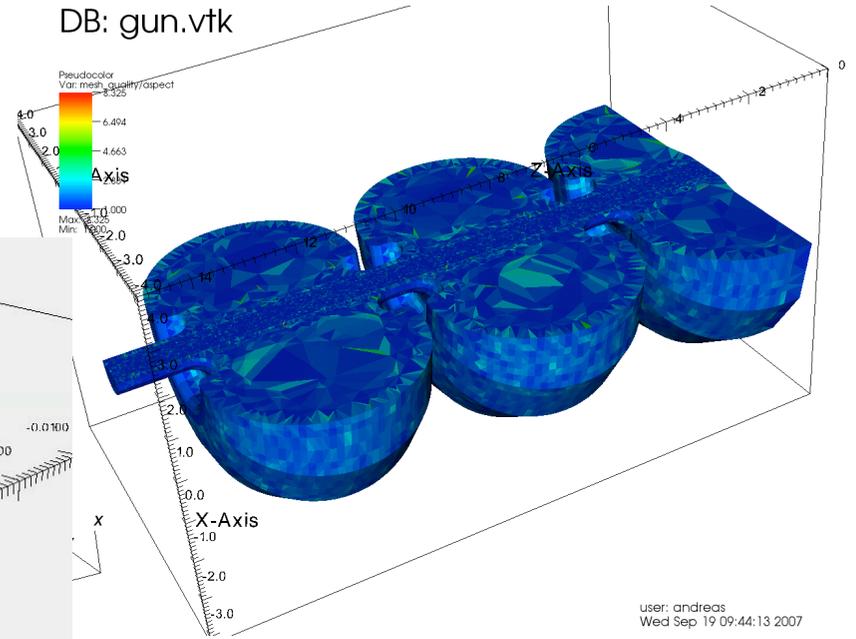
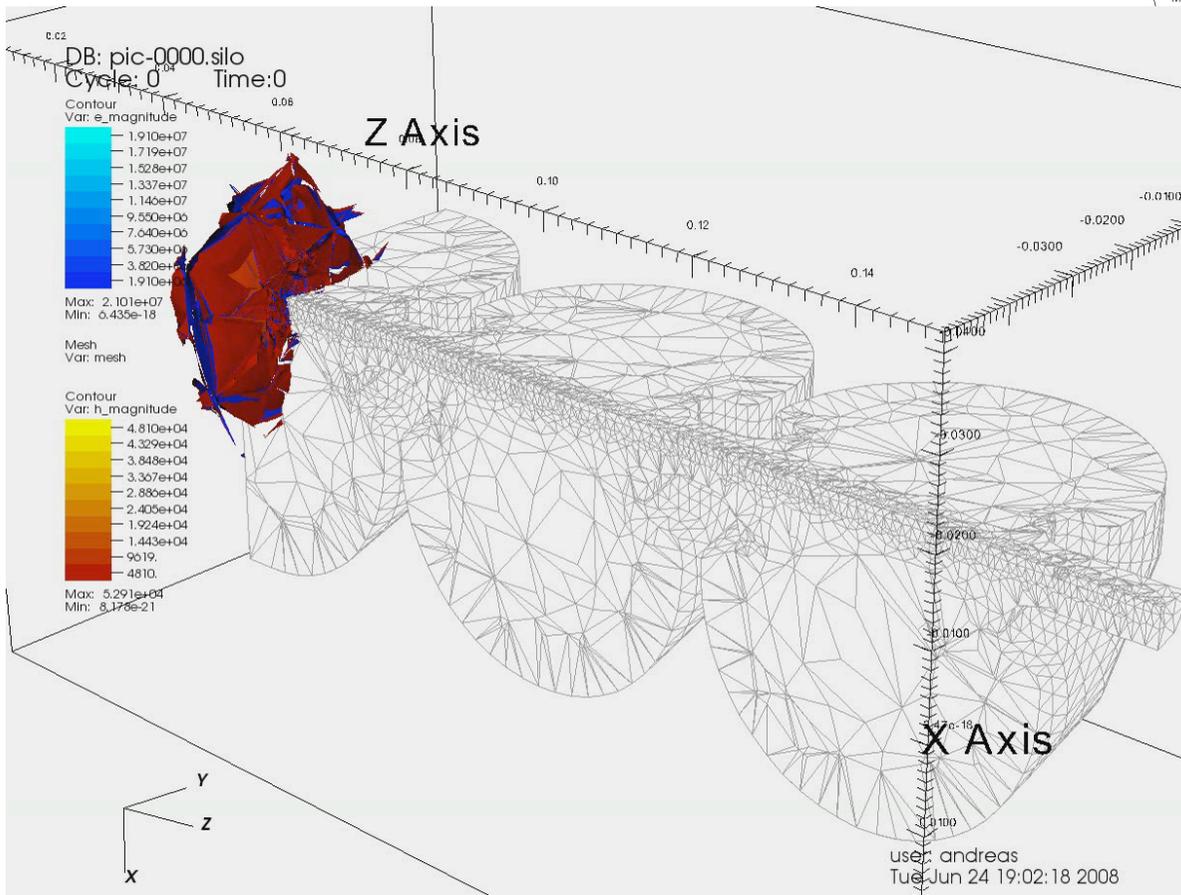
Fields-to-particles

$$E(x_n), H(x_n)$$

Some 3D results



Particle gun



3D Extension

Nothing special !

Everything you have done in 1D/2D you can do in 3D in exactly the same way.

- ✓ Linear/nonlinear problems
- ✓ First order/higher order operators
- ✓ Complex geometries

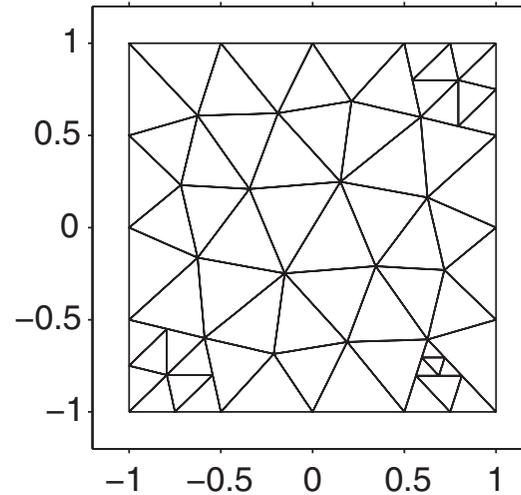
Further extensions

- ✓ Adaptivity/non-conforming elements
- ✓ Curvilinear elements

Adaptivity/non-conformity

Question: Do element faces always have to match ?

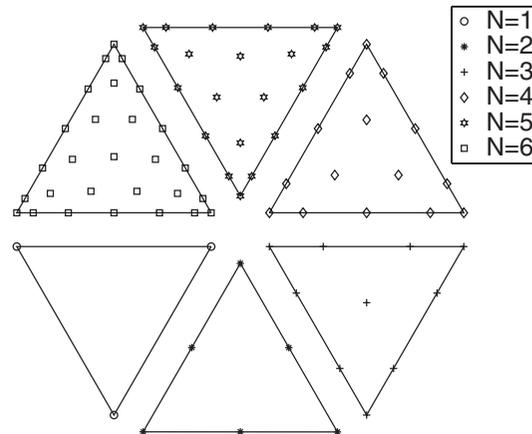
Answer: No



h-nonconform

Question: Can one use different order in each element ?

Answer: Yes



p-nonconform

Example - Adaptive solution

We consider a standard test case

$$\nabla^2 u(\mathbf{x}) = f(\mathbf{x}) \quad u = 0, \mathbf{x} \in \partial\Omega$$

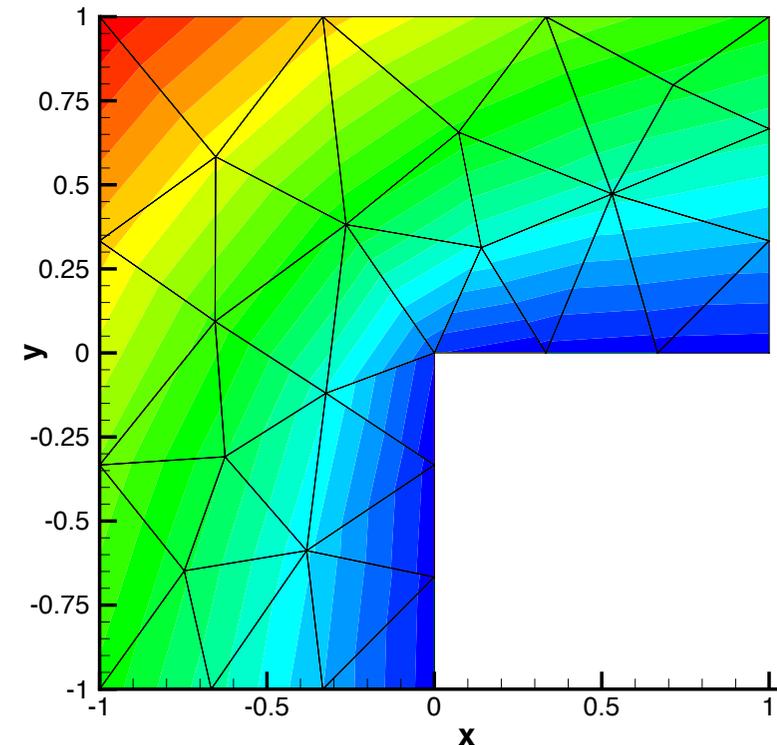
Domain is L-shaped

RHS so that the exact solution is

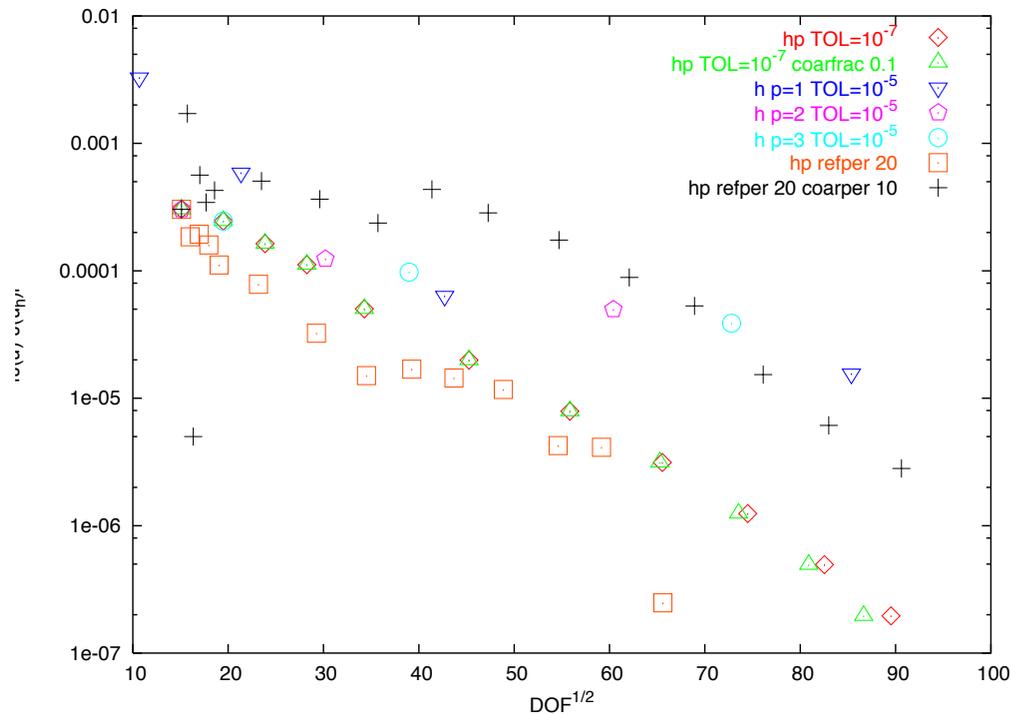
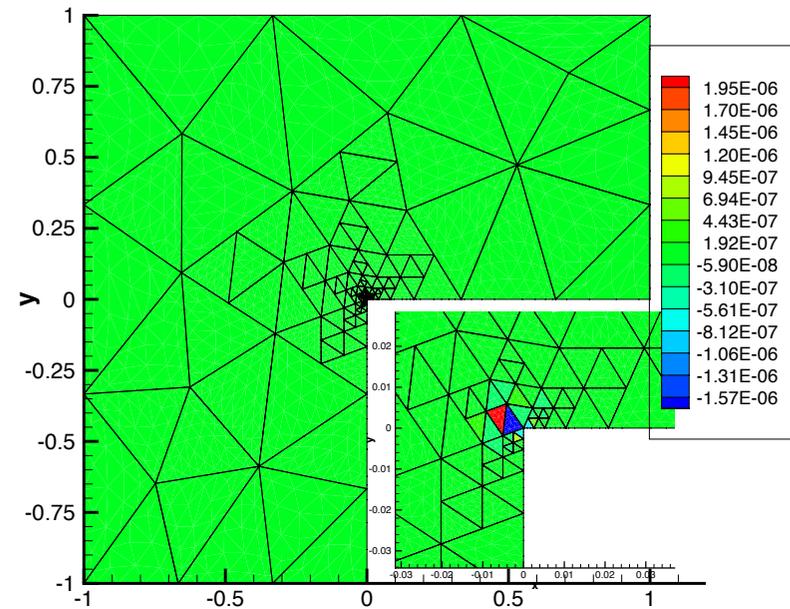
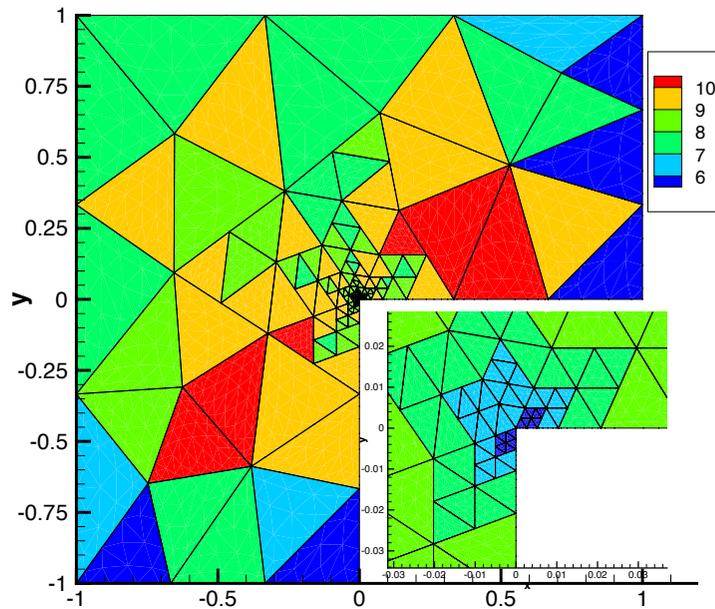
$$u(r, \theta) = r^{2/3} \sin(2\pi/3\theta)$$

Solution is singular !

Solved using full hp-adaptive solution



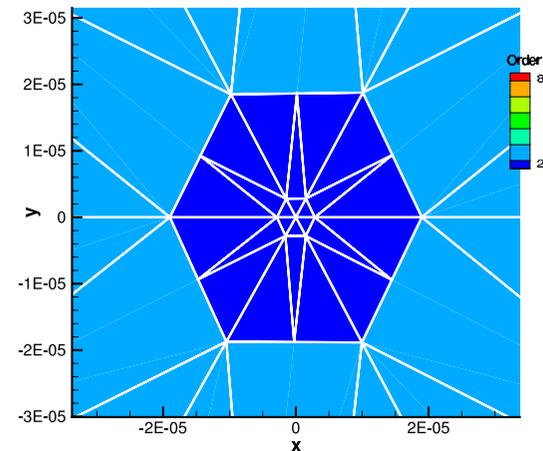
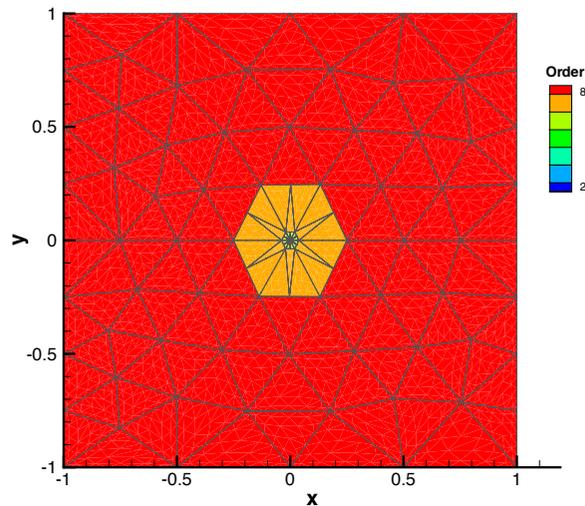
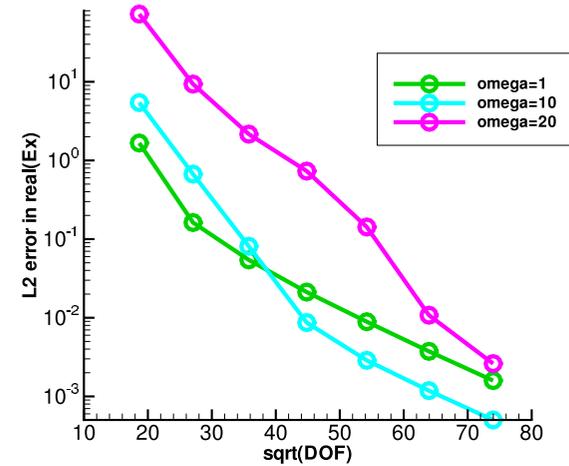
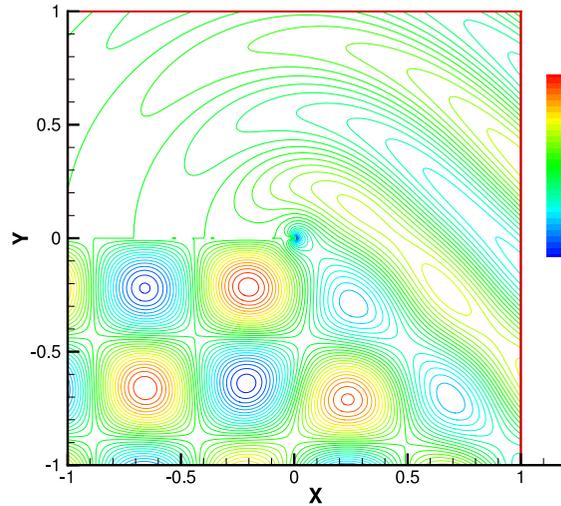
Example - Adaptive solution



Spectral convergence
even for a singular
solution

Example - Adaptive solution - Maxwell's

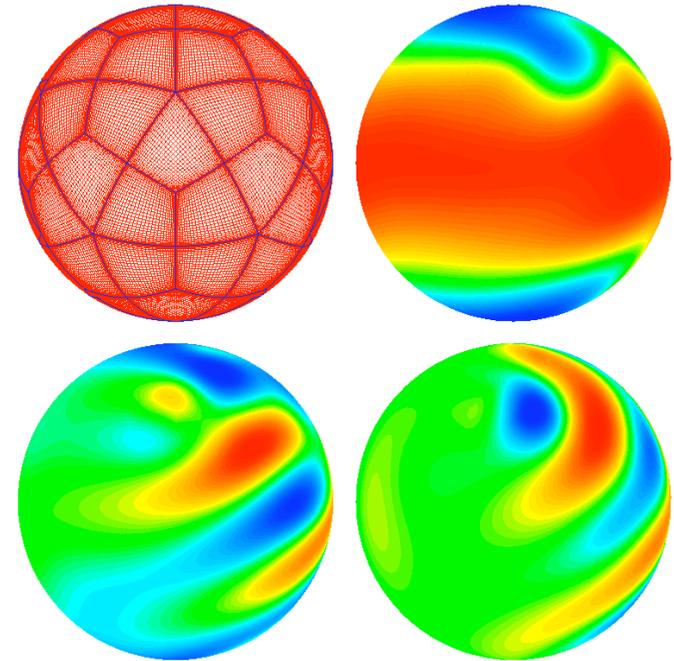
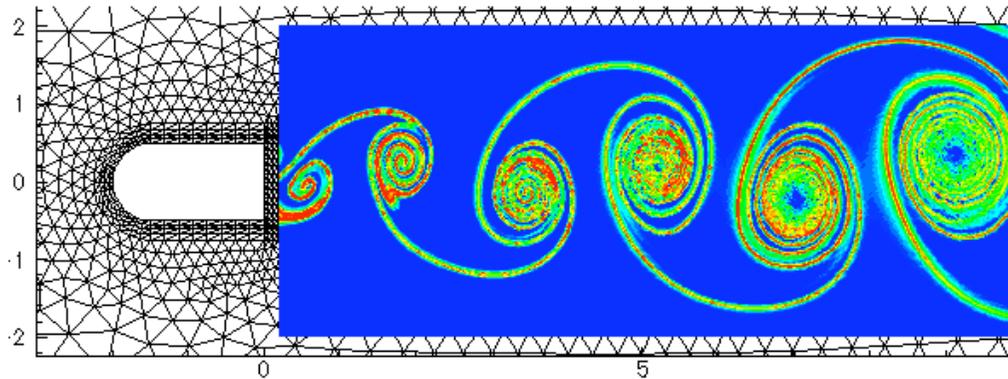
$$\nabla \times \nabla \times \mathbf{E} + \omega^2 \mathbf{E} = \mathbf{f}, \mathbf{n} \times \mathbf{E} = 0, \mathbf{x} \in \Omega$$



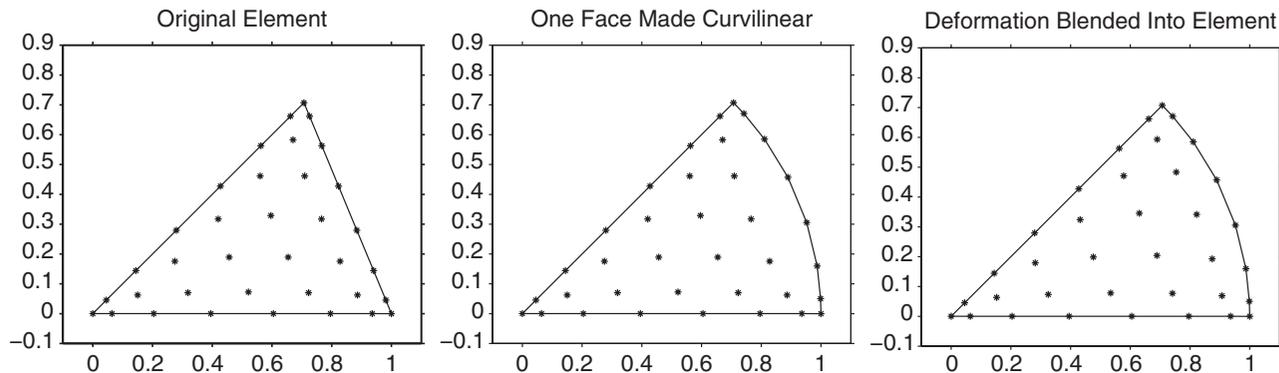
Curvilinear elements

What: Elements that conform exactly to a curved boundary

Why: Accuracy !



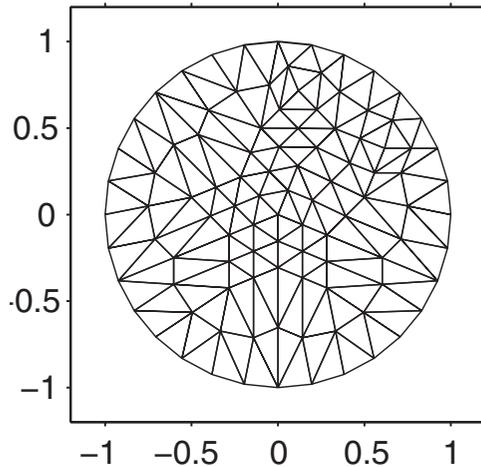
This is a unique feature to high-order elements



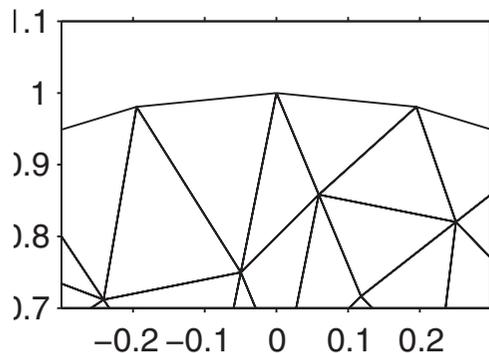
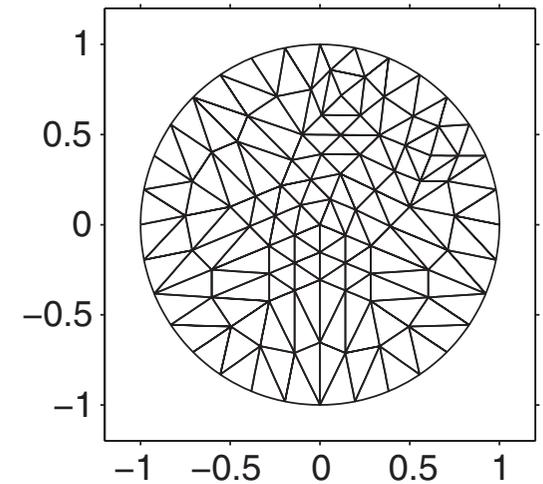
Example - Maxwell's equations

$$H^x(x, y, t = 0) = 0, \quad H^y(x, y, t = 0) = 0,$$

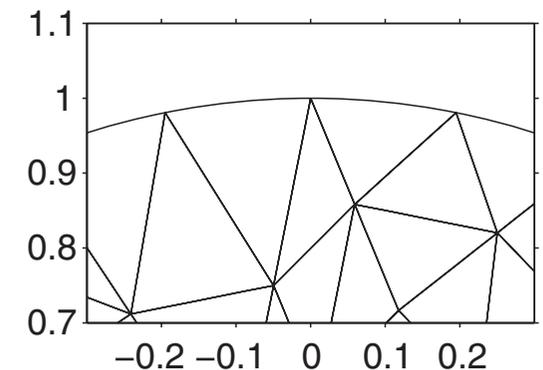
$$E^z(x, y, t = 0) = J_6(\alpha_6 r) \cos(6\theta) \cos(\alpha_6 t),$$



N	h	$h/2$	$h/4$	Rate
1	1.09E-01	3.78E-02	8.96E-03	1.80
2	2.37E-02	2.70E-03	3.58E-04	3.02
3	5.77E-03	1.09E-03	2.70E-04	2.21
4	4.39E-03	1.09E-03	2.72E-04	2.01
5	4.38E-03	1.09E-03	2.72E-04	2.00
6	4.40E-03	1.09E-03	2.72E-04	2.01
7	4.41E-03	1.09E-03	2.73E-04	2.01



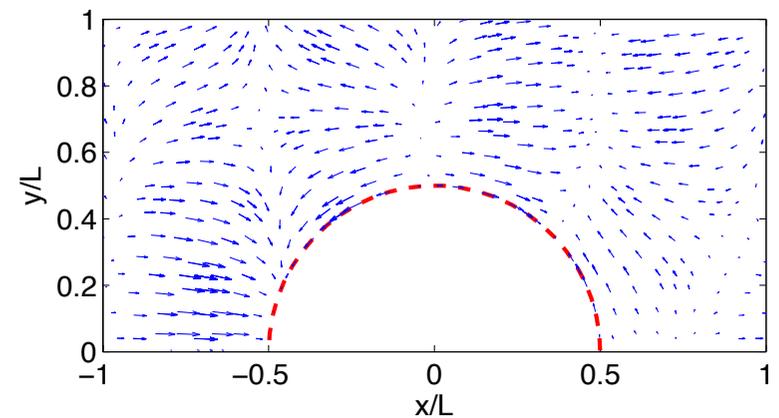
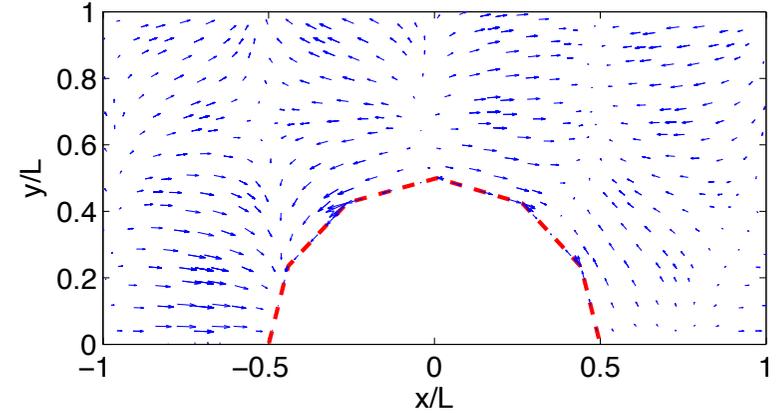
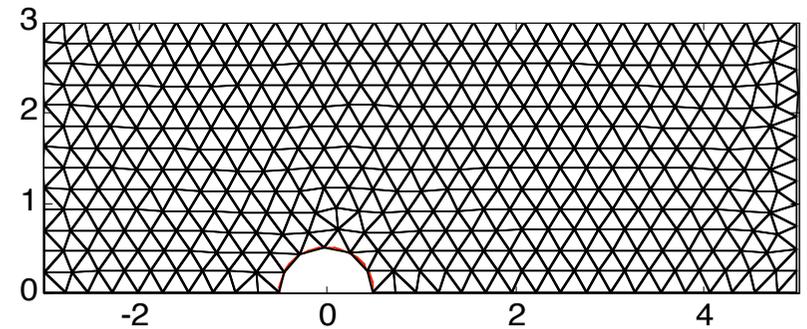
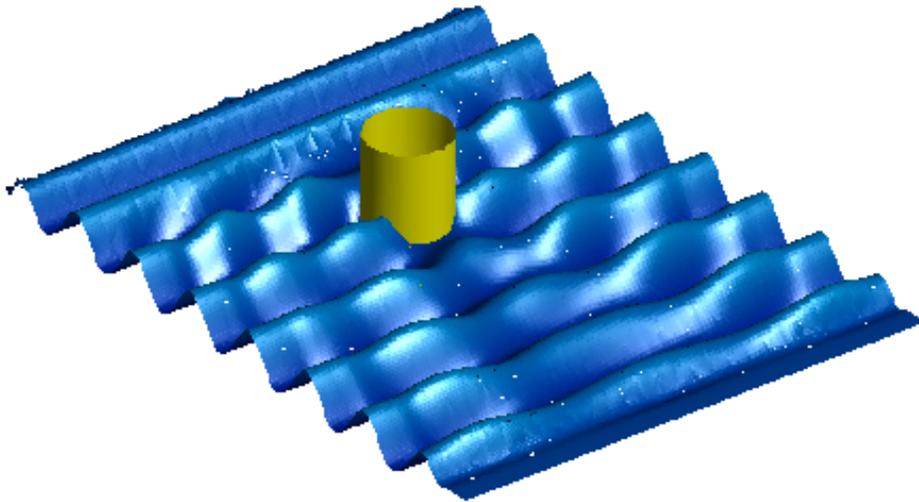
N	h	$h/2$	$h/4$	Rate
1	1.09E-01	3.78E-02	8.96E-03	1.80
2	2.21E-02	2.23E-03	2.05E-04	3.38
3	3.12E-03	1.92E-04	1.28E-05	3.97
4	6.01E-04	1.95E-05	5.88E-07	5.00
5	9.89E-05	1.69E-06	2.72E-08	5.92
6	1.74E-05	1.31E-07	9.81E-10	7.06
7	2.08E-06	8.97E-09	7.93E-11*	7.34*



This is essential to fully benefit for complex problems

Example - Boussinesq equations

The correct representation of the boundary is essential for accuracy and speed



Example - Spherical Shallow Water equ

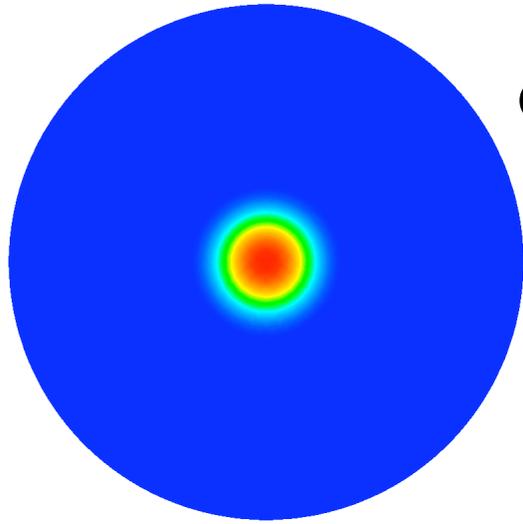
Dynamics of a thin layer of fluids on a sphere

$$\frac{\partial}{\partial t} \begin{bmatrix} \varphi \\ \varphi u \\ \varphi v \\ \varphi w \end{bmatrix} + \frac{\partial}{\partial x} \begin{bmatrix} \varphi u \\ \varphi u^2 + \frac{1}{2} \varphi^2 \\ \varphi uv \\ \varphi uw \end{bmatrix} + \frac{\partial}{\partial y} \begin{bmatrix} \varphi v \\ \varphi v^2 + \frac{1}{2} \varphi^2 \\ \varphi vu \\ \varphi vw \end{bmatrix} + \frac{\partial}{\partial z} \begin{bmatrix} \varphi w \\ \varphi w^2 + \frac{1}{2} \varphi^2 \\ \varphi wv \\ \varphi wu \end{bmatrix} = \begin{bmatrix} 0 \\ -\frac{f}{a} (y\varphi w - z\varphi v) + \mu x \\ -\frac{f}{a} (z\varphi u - x\varphi w) + \mu y \\ -\frac{f}{a} (x\varphi v - y\varphi u) + \mu z \end{bmatrix}$$

$$\frac{\partial \bar{\varphi}}{\partial t} + \nabla \cdot \bar{F} = S(\bar{\varphi})$$

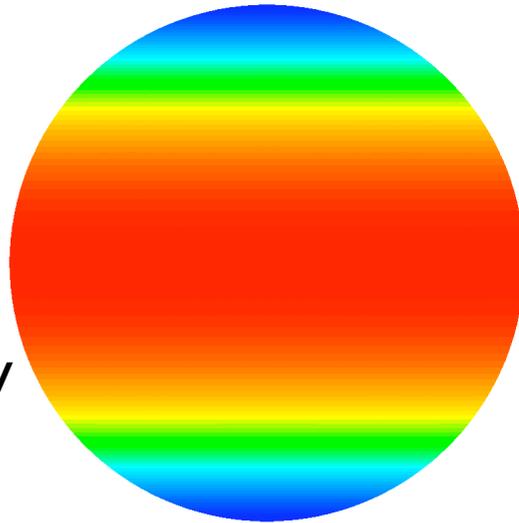
Standard benchmark (Williamsson) in geophysical flow modeling

Example - Spherical Shallow Water equ

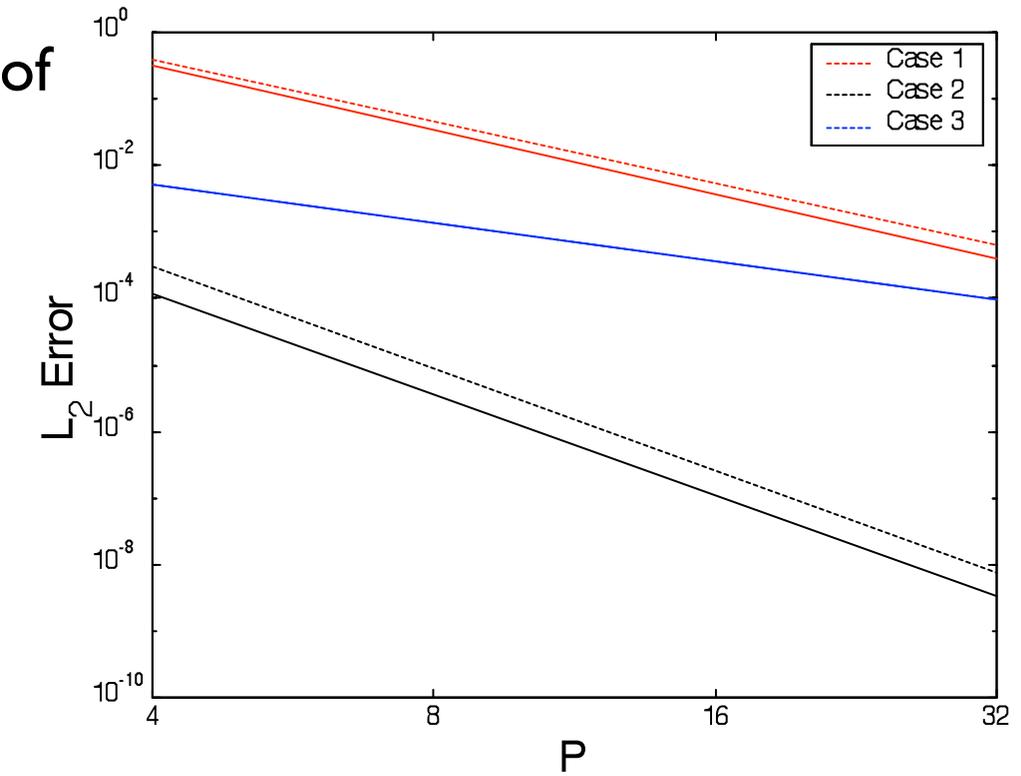


Case I: Rotation of cosine wave

Case II: Zonal Geostrophic Flow



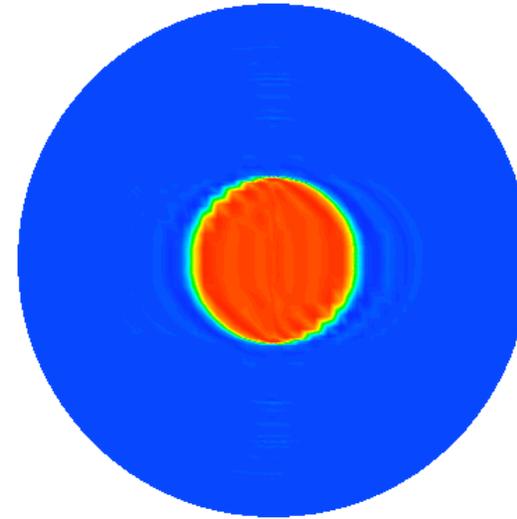
Case III: Zonal flow w/ compact support



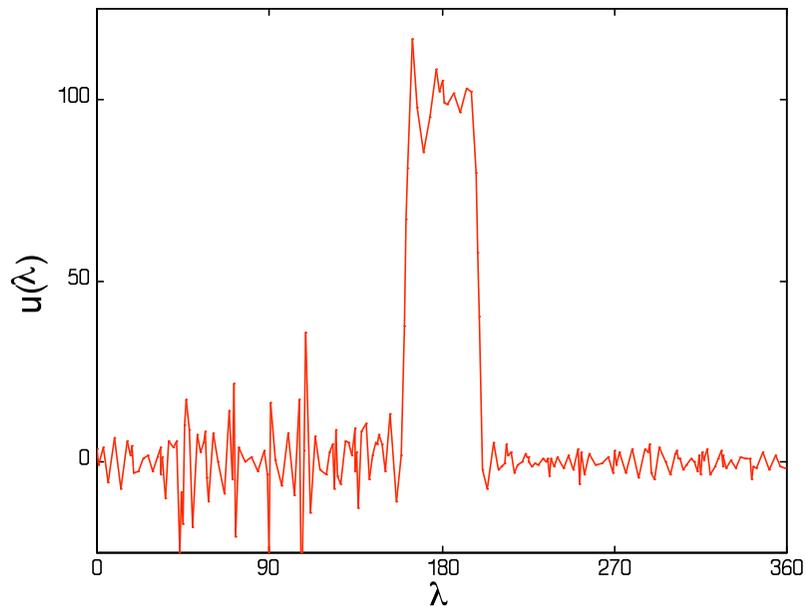
Example - Spherical Shallow Water equ

Rotation of cylinder

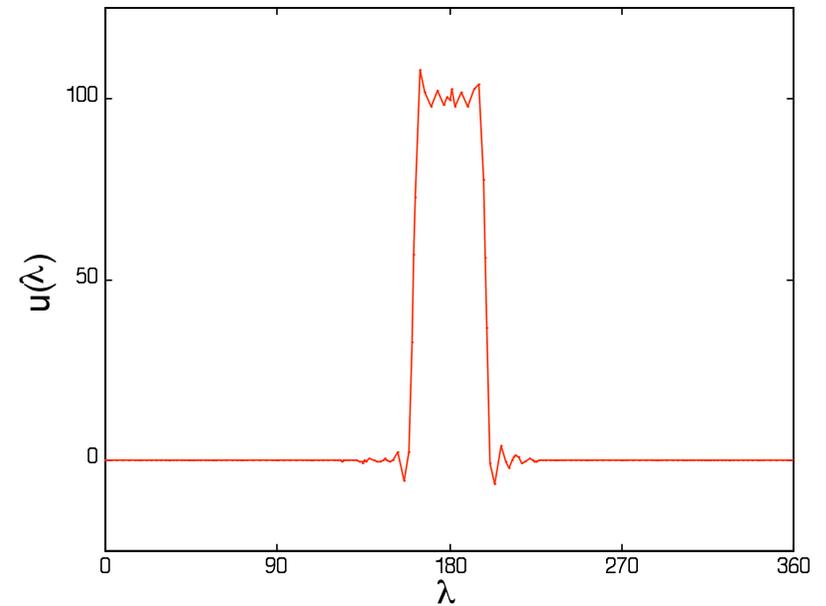
$N=8$



SEM



DG-FEM



Summary of Part I

We have generalized everything to 3D

- ✓ Linear/nonlinear problems
- ✓ First order/higher order operators
- ✓ Complex geometries
- ✓ Apaptivity
- ✓ Curvilinear elements

There is only one significant obstacle to solving large problems

SPEED !

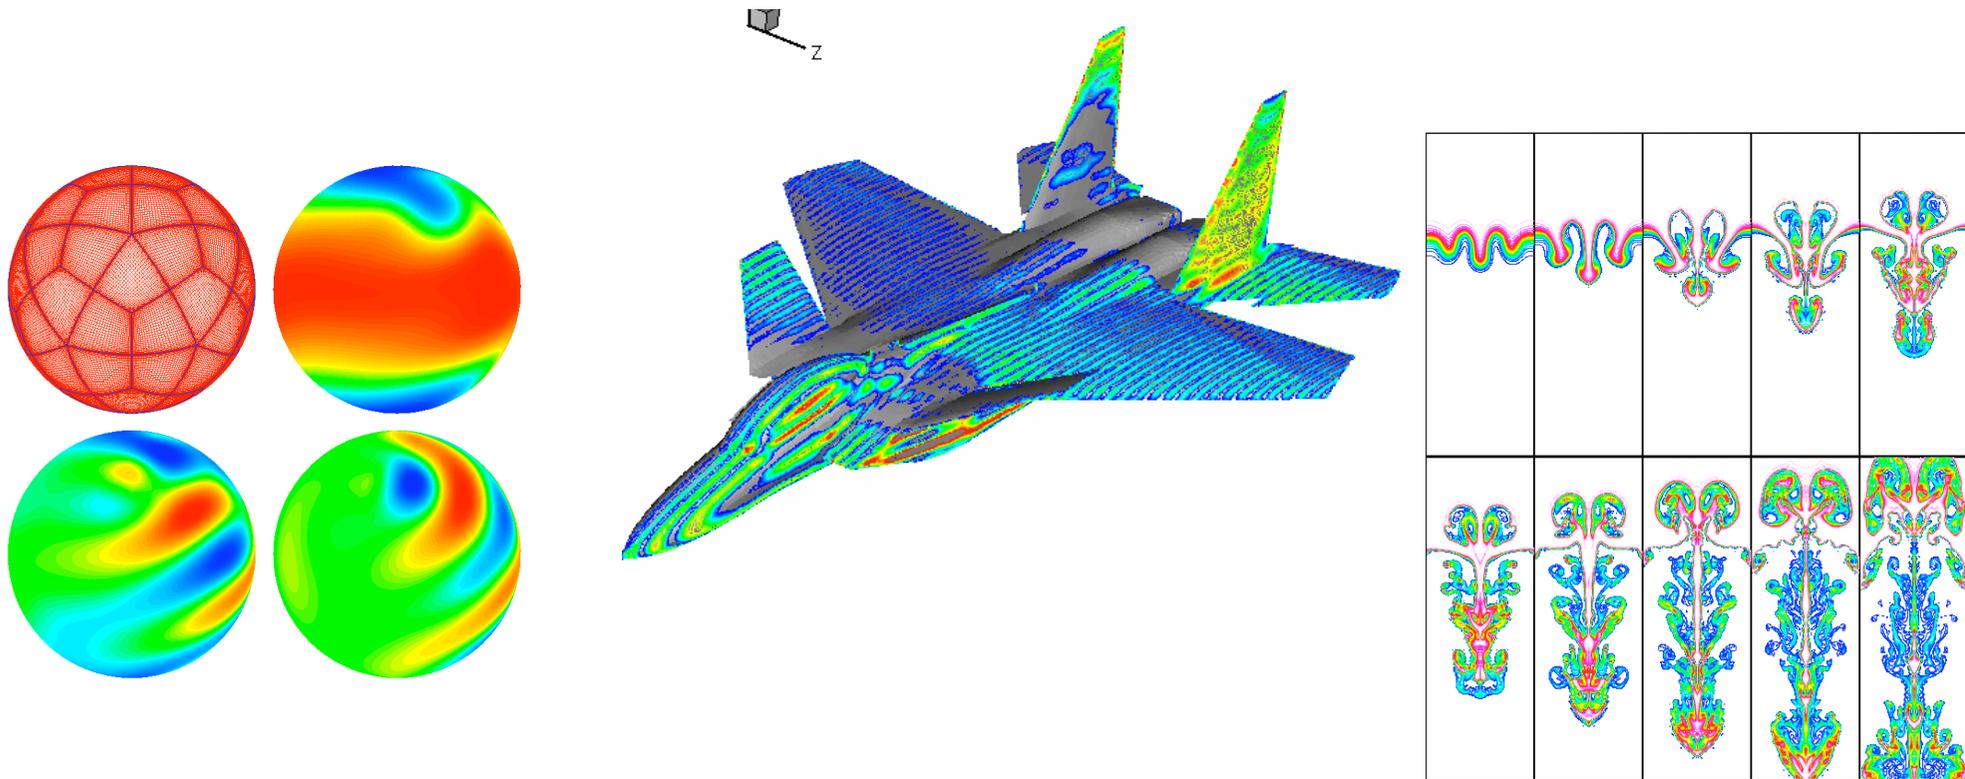
Lecture 8

- ✓ Let's briefly recall what we know
- ✓ Part I: 3D problems and extensions
 - ✓ Formulations and examples
 - ✓ Adaptivity and curvilinear elements
- ✓ Part II: The need for speed
 - ✓ Parallel computing
 - ✓ GPU computing
 - ✓ Software beyond Matlab

The need for speed !

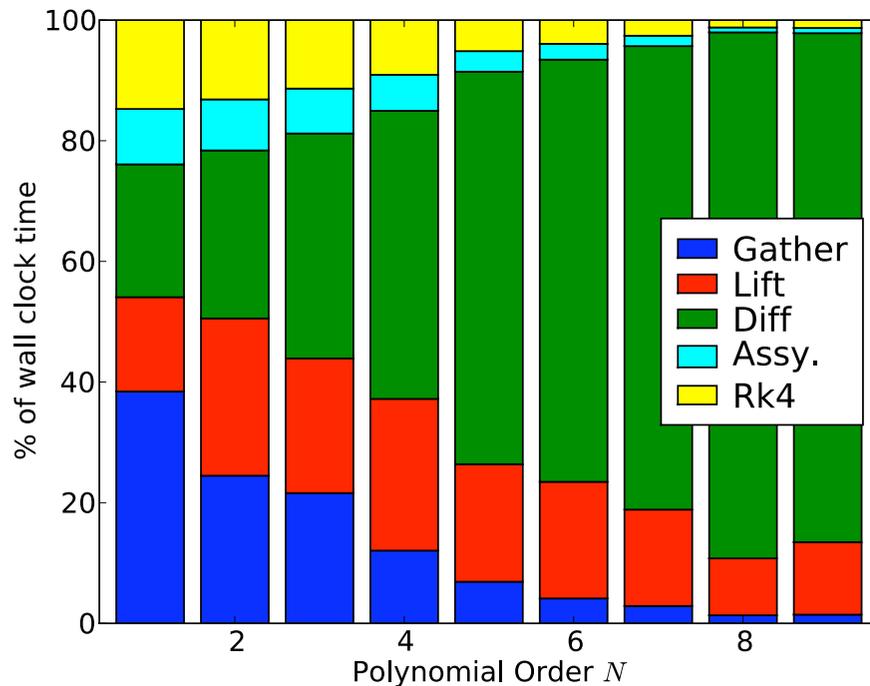
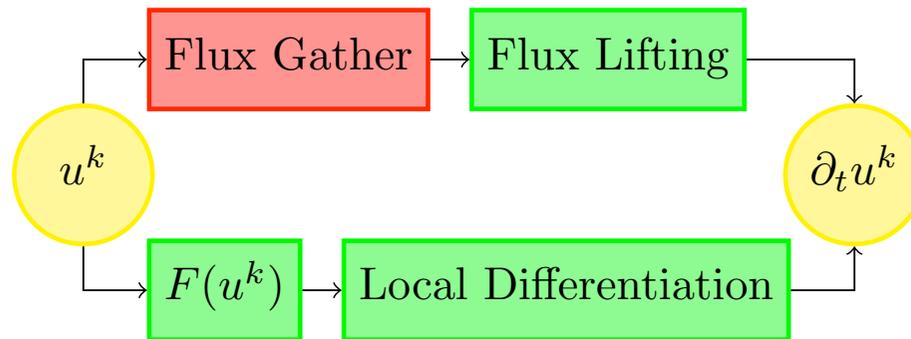
So far, we have focused on 'simple' serial computing using Matlab based model.

However, this will not suffice for many applications



The need for speed

Let us first understand where we spend the time



Test case is
3D Maxwells

The majority of
work is **local**

The need for speed

The locality suggest that parallel computing will be beneficial

- ✓ Using OpenMP, the local work can be distributed over elements through loops.
- ✓ Using MPI the locality ensures a surface communication model.
- ✓ Mixed OpenMP/MPI models also possible
- ✓ A similar line of arguments can be used for iterative solvers.

Parallel performance

# Processors	64	128	256	512
Scaled RK time	1.00	0.48	0.24	0.14
Ideal time	1.00	0.50	0.25	0.13

High performance is achieved through -

- ✓ Local nature of scheme
- ✓ Pure matrix-matrix operations
- ✓ Local bandwidth minimization
- ✓ Very efficient on-chip performance (~75%)

Challenges -

- ✓ Efficient parallel preconditioning

Parallel computing

DG-FEM maps very well to classic multi-processor computing clusters and result in excellent speed-up.

... but such machines are expensive to buy and run.

Ex: To get on the Top500 list, requires about \$3m to purchase a cluster with 16-18Tflop/s performance.

What we need is supercomputing on the desktop

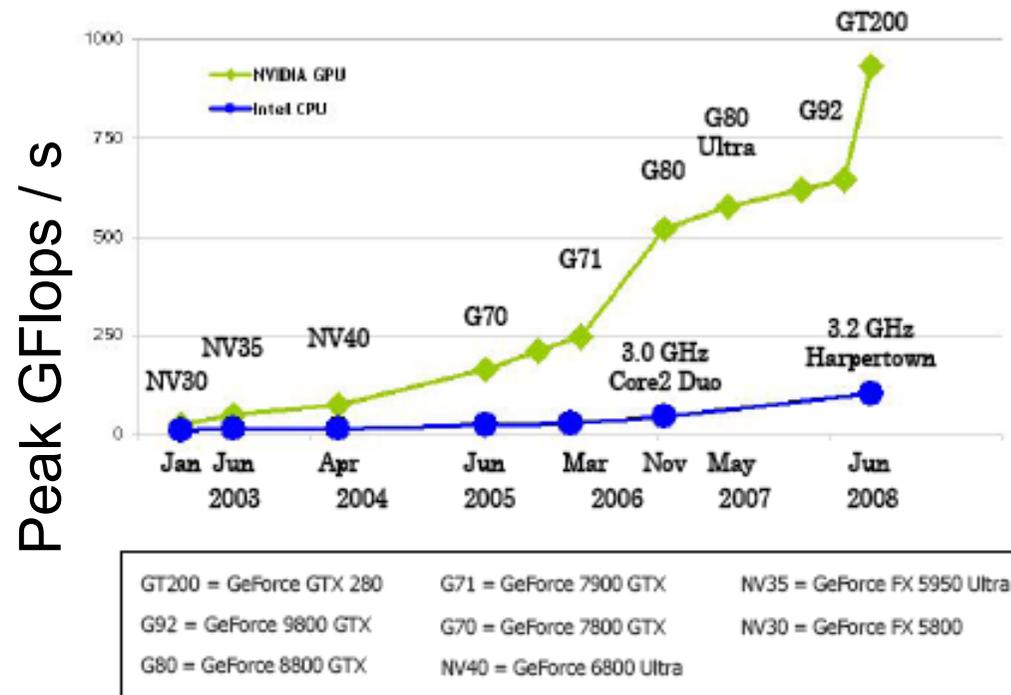
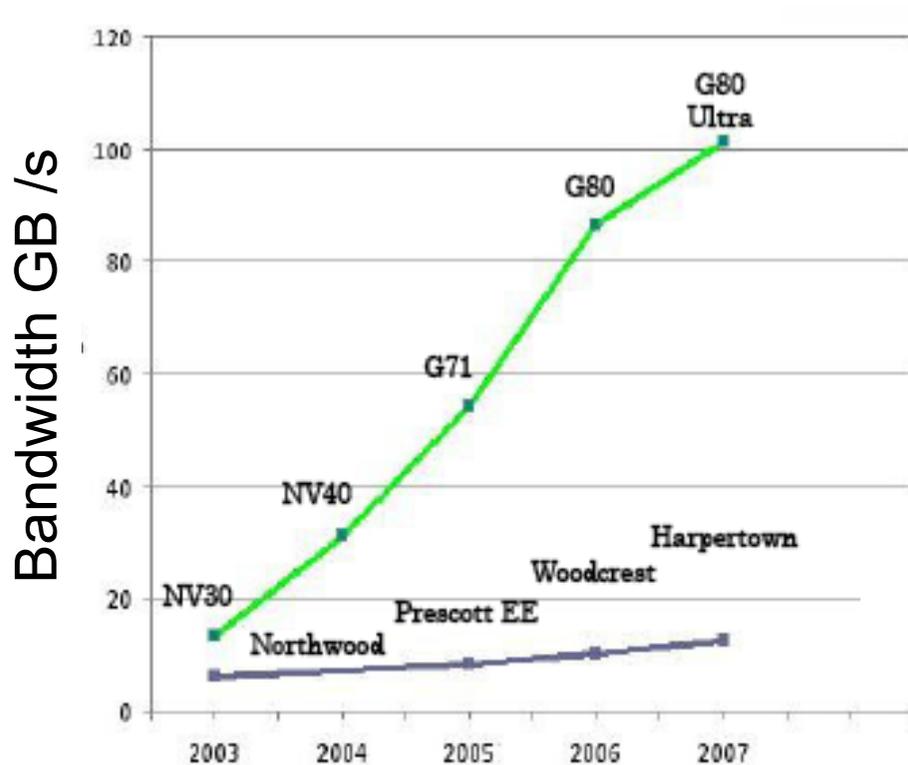
For FREE !

... or at least at a fraction of the price

CPU's vs GPU's



Notice the following



The memory bandwidth and the peak performance on Graphics cards (GPU's) is developing MUCH faster than on CPU's

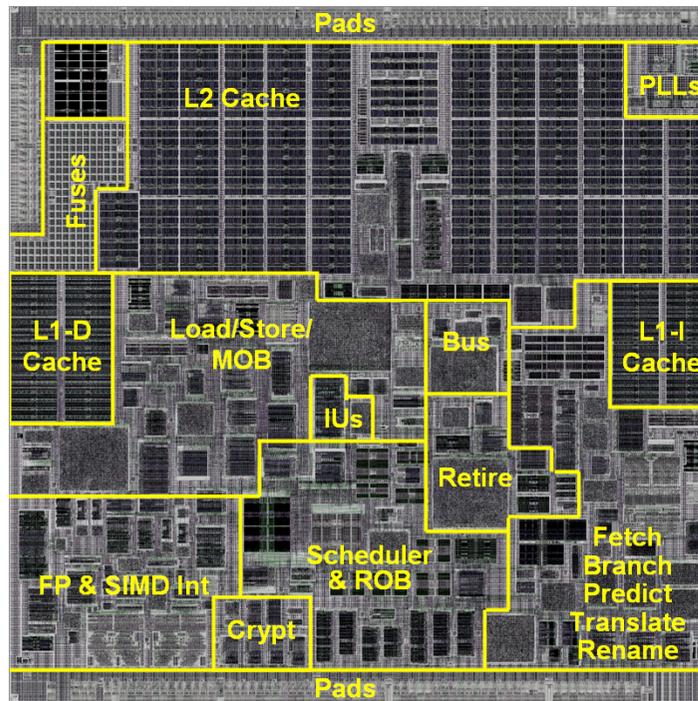
At the same time, the mass-marked for gaming drives the prices down -- we have to find a way to exploit this !

But why is this ?



Target for CPU:

- ✓ Single thread very fast
- ✓ Large caches to hide latency
- ✓ Predict, speculate etc



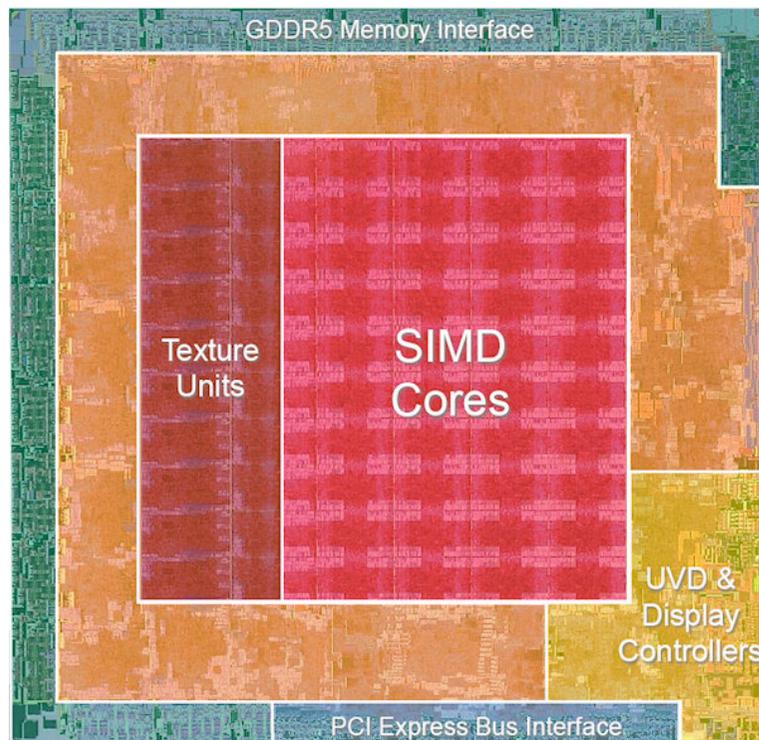
Lots of very complex logic
to predict behavior

But why is this ?



For streaming/graphics cards it is different

- ✓ Throughput is what matters
- ✓ Hide latency through parallelism
- ✓ Push hierarchy onto programmer

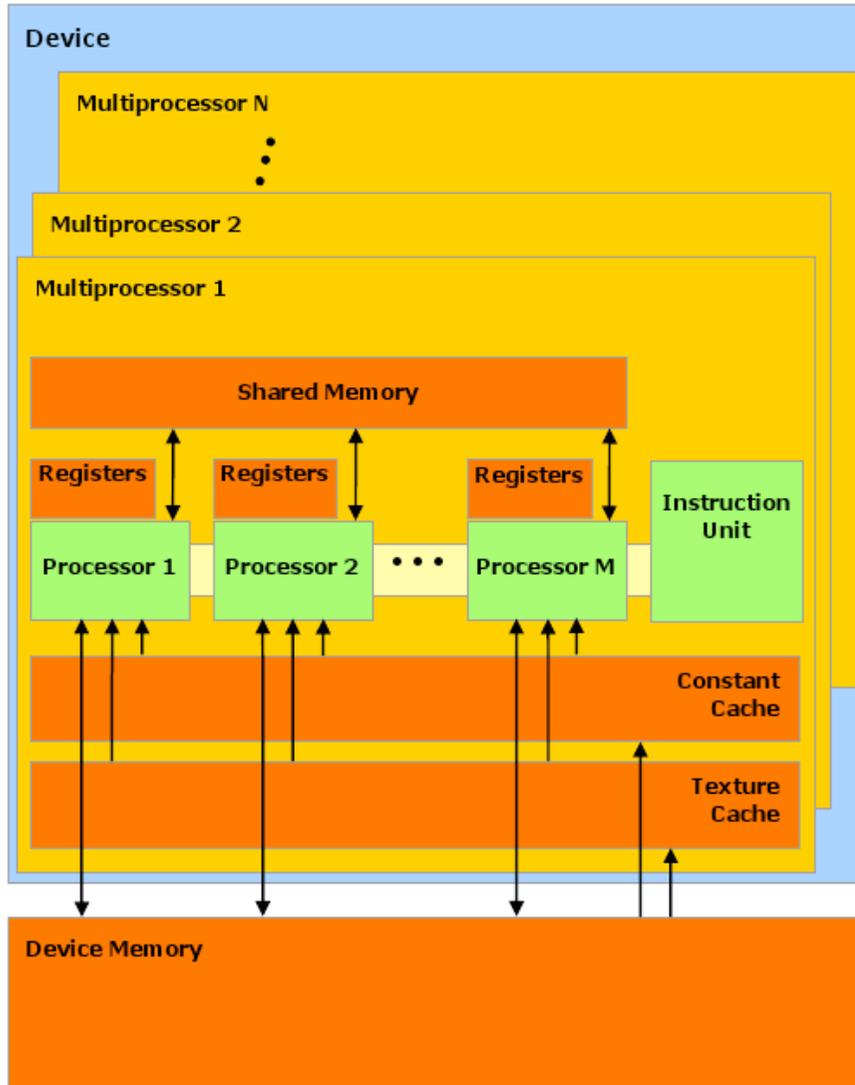


Much simpler logic with a focus on performance

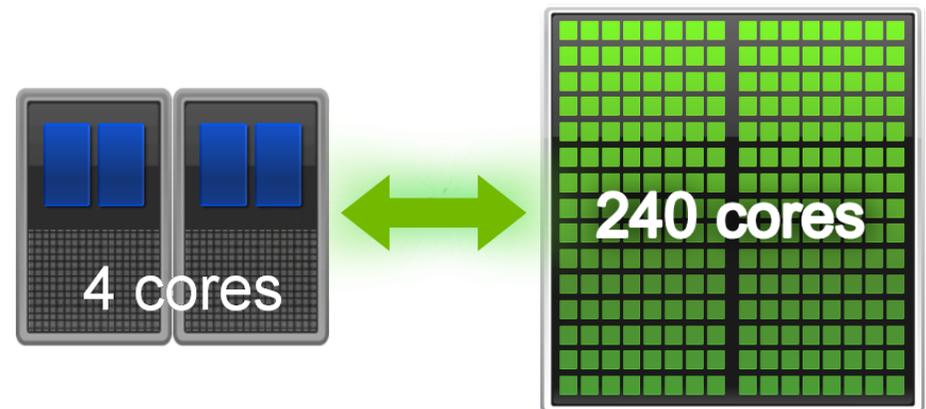
GPUs 101



GPU layout



- ✓ 1 GPU = 30 MPs
- ✓ 1 MP has 1 IU, 8 SP, 1 DP
- ✓ 1 MP has 16KiB shared and 32 KiB Register memory
- ✓ 240 (512) threads
- ✓ Dedicated RAM at 140GB/s
- ✓ Limited caches





Gains

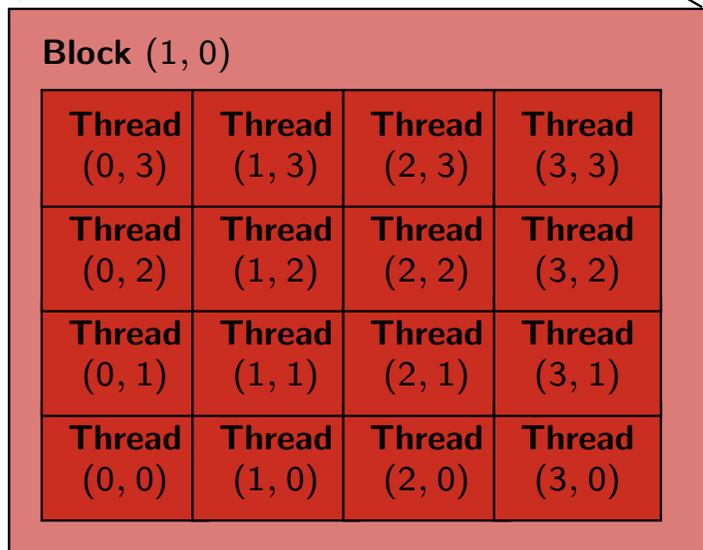
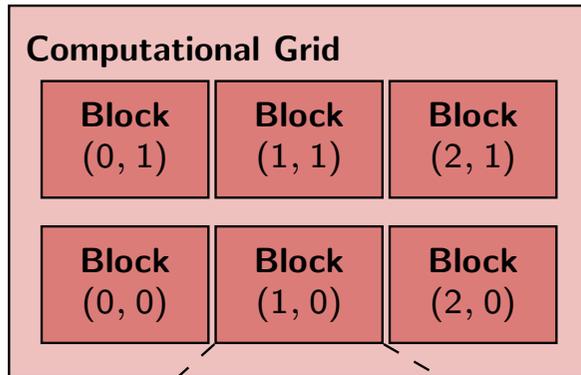
- ⊕ Memory Bandwidth
(140 GB/s vs. 12 GB/s)
- ⊕ Compute Bandwidth
(Peak: 1 TF/s vs. 50 GF/s,
Real: 200 GF/s vs. 10 GF/s)

Losses

- ⊖ Recursion
- ⊖ Function pointers
- ⊖ Exceptions
- ⊖ IEEE 754 FP compliance
- ⊖ Cheap branches (i.e. ifs)

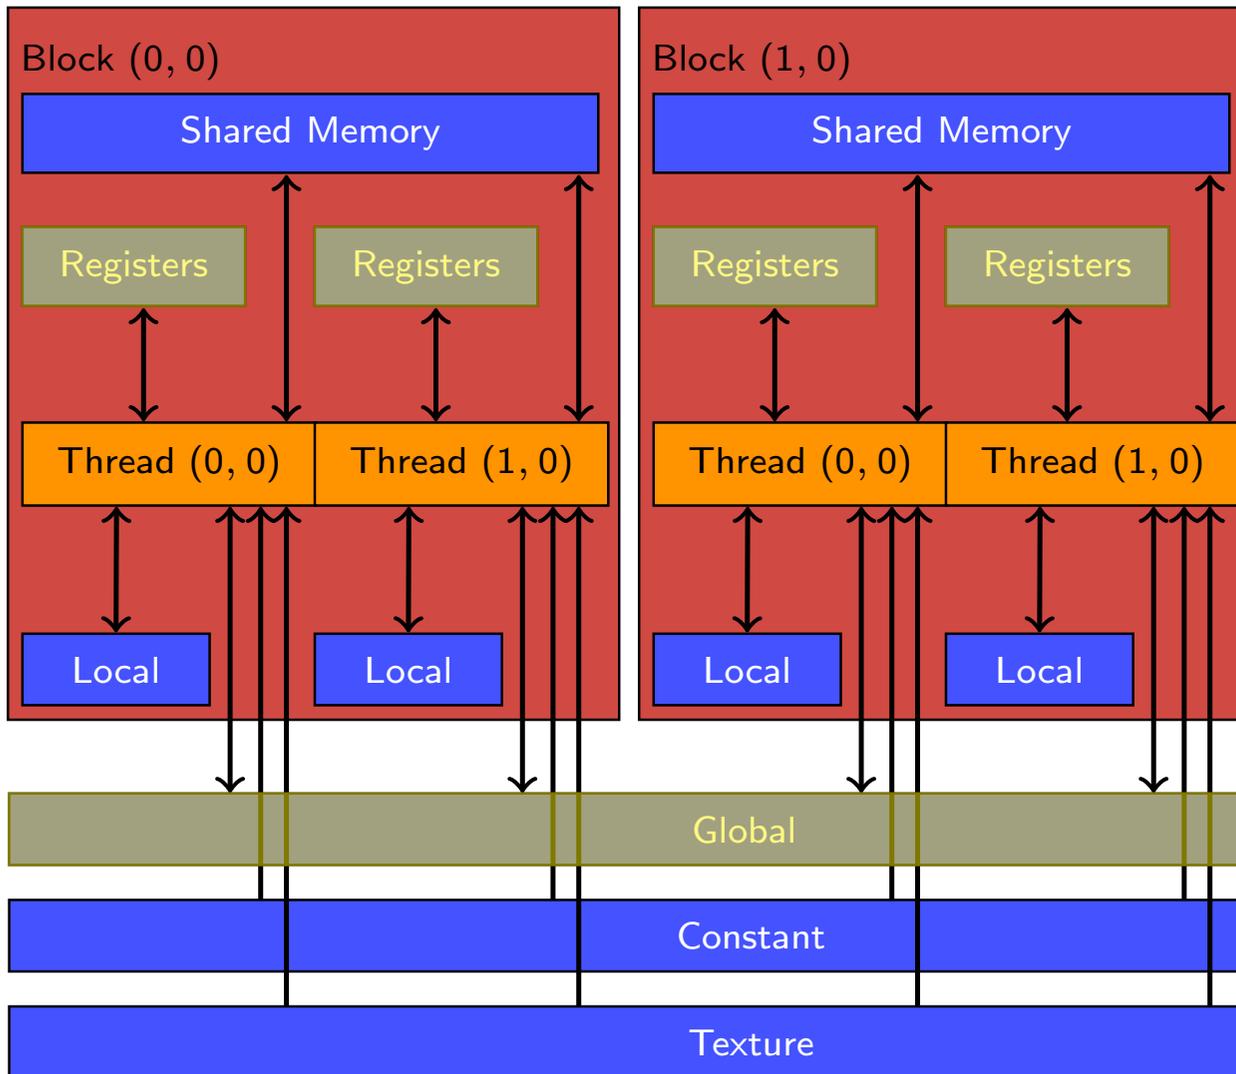
Already here it is clear that programming models/codes may have to undergo substantial changes -- and that not all will work well

GPUs 101



- ✓ Genuine multi-tiered parallelism
 - ✓ Grids
 - ✓ blocks
 - ✓ threads
- ✓ Only threads within a block can talk
 - ✓ Blocks must be executed in order
- ✓ Grids/blocks/threads replace loops
- ✓ Until recently, only single precision
- ✓ Code-able with CUDA (C-extension)

GPUs 101



Memory model:

- ✓ Registers
- ✓ Local shared
- ✓ Global

GPUs 101

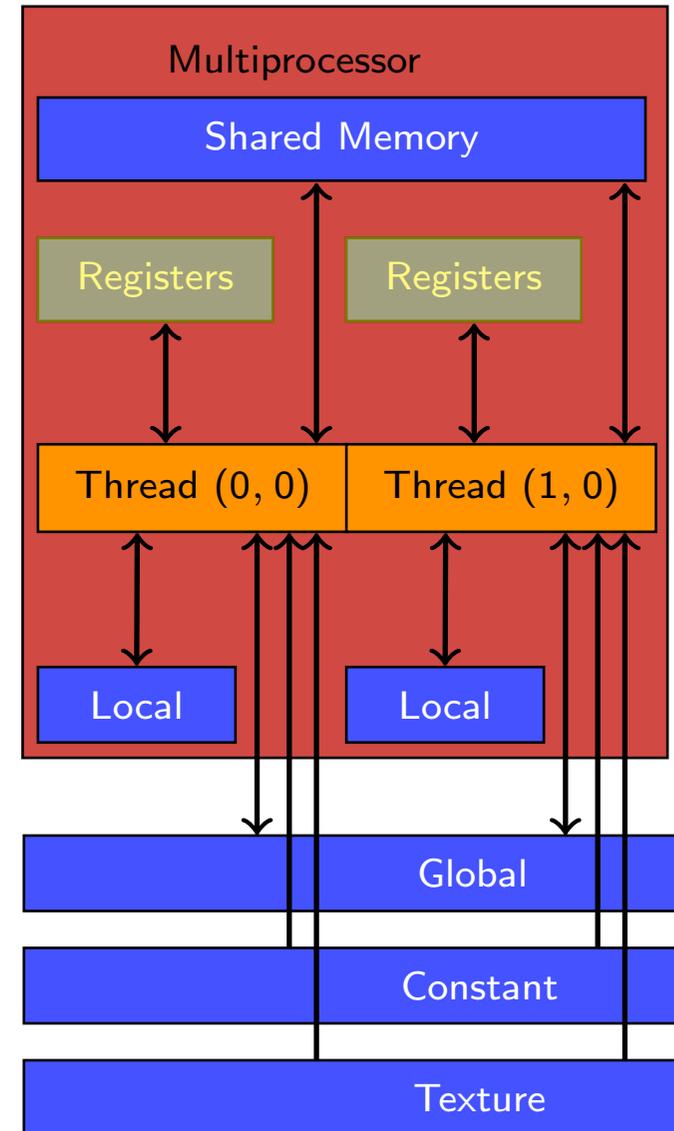


✓ Lots of multi-processors (about 30)

... communicate through global mem

✓ Registers, shared memory, and threads communicate with low latency

... but memory is limited (16-32 KiB)



GPUs 101



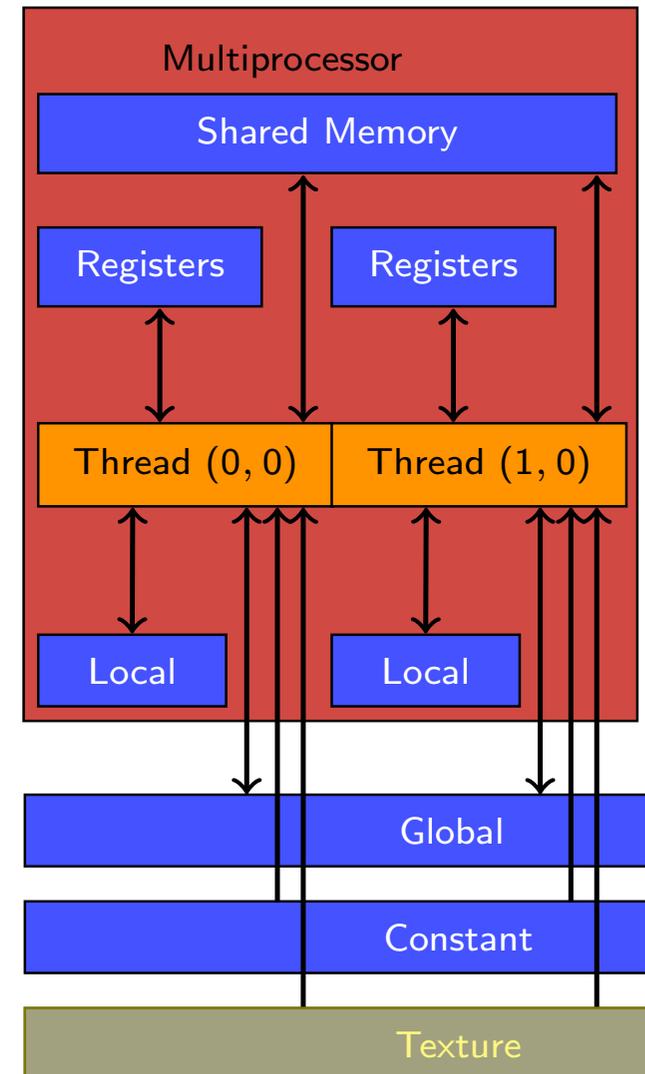
✓ Global memory (4GiB/PU) is plentiful

... but latency is high (512 bit bus)
... and stride one is preferred

✓ Texture is similar to global memory

... allows more general access patterns
... but it is read only

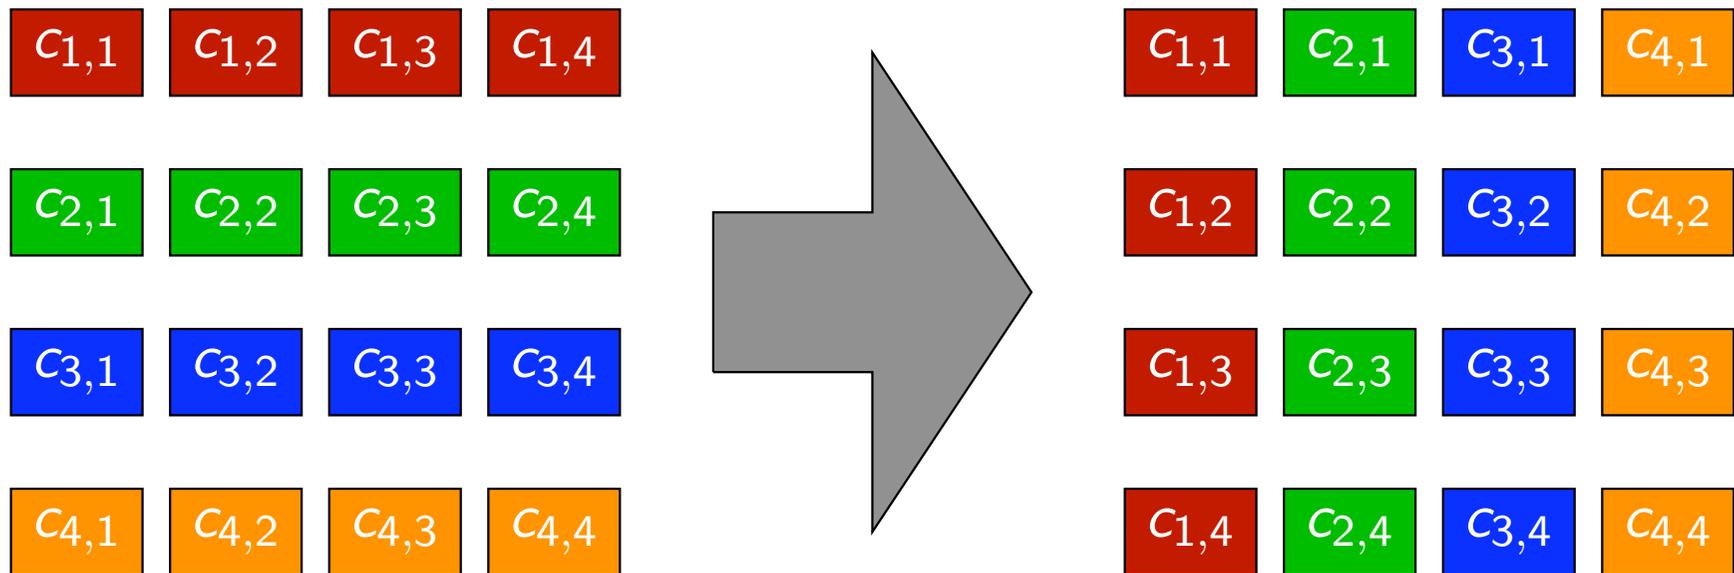
Type	Per	Access	Latency
Registers	thread	R/W	1
Local	thread	R/W	1000
Shared	block	R/W	2
Global	grid	R/W	1000
Constant	grid	R/O	1-1000
Texture	grid	R/O	1000



Let's consider an example



Matrix transpose

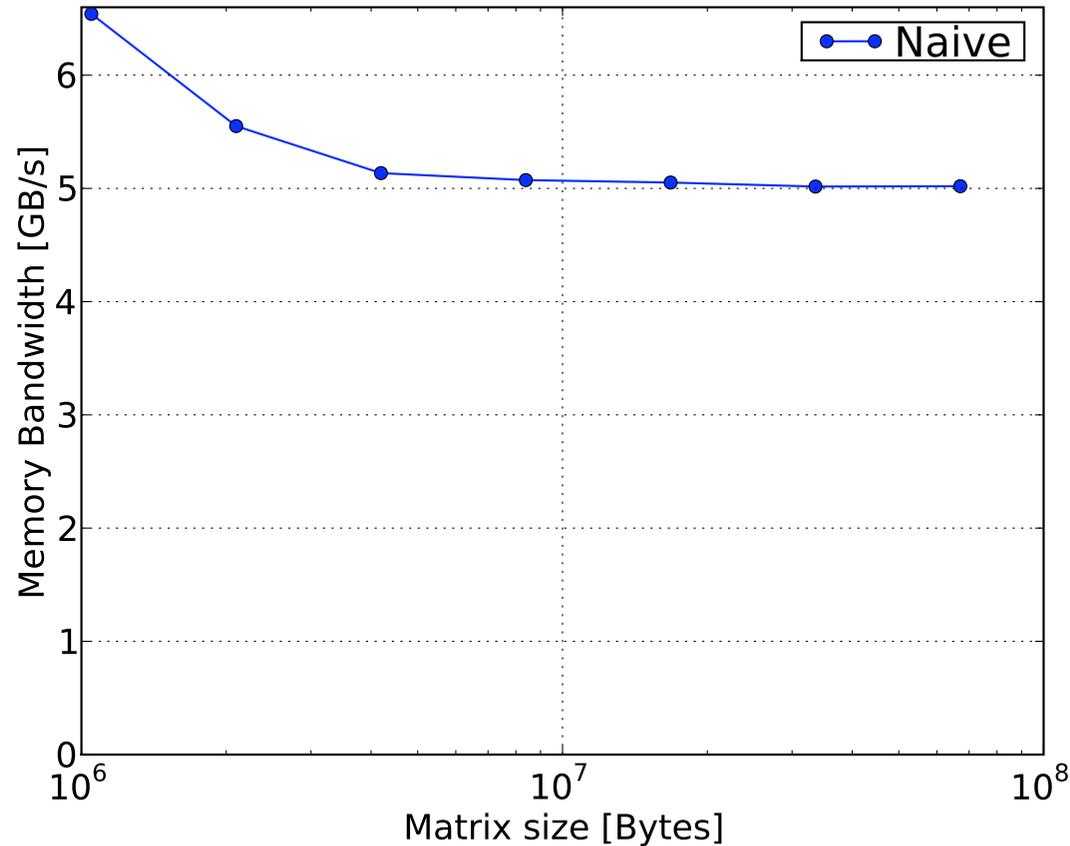


Memory bandwidth will be a limit here

Let's consider an example



Using just global memory



As CPU

Reading from global mem:



stride: 1 → one mem.trans.

Writing to global mem:

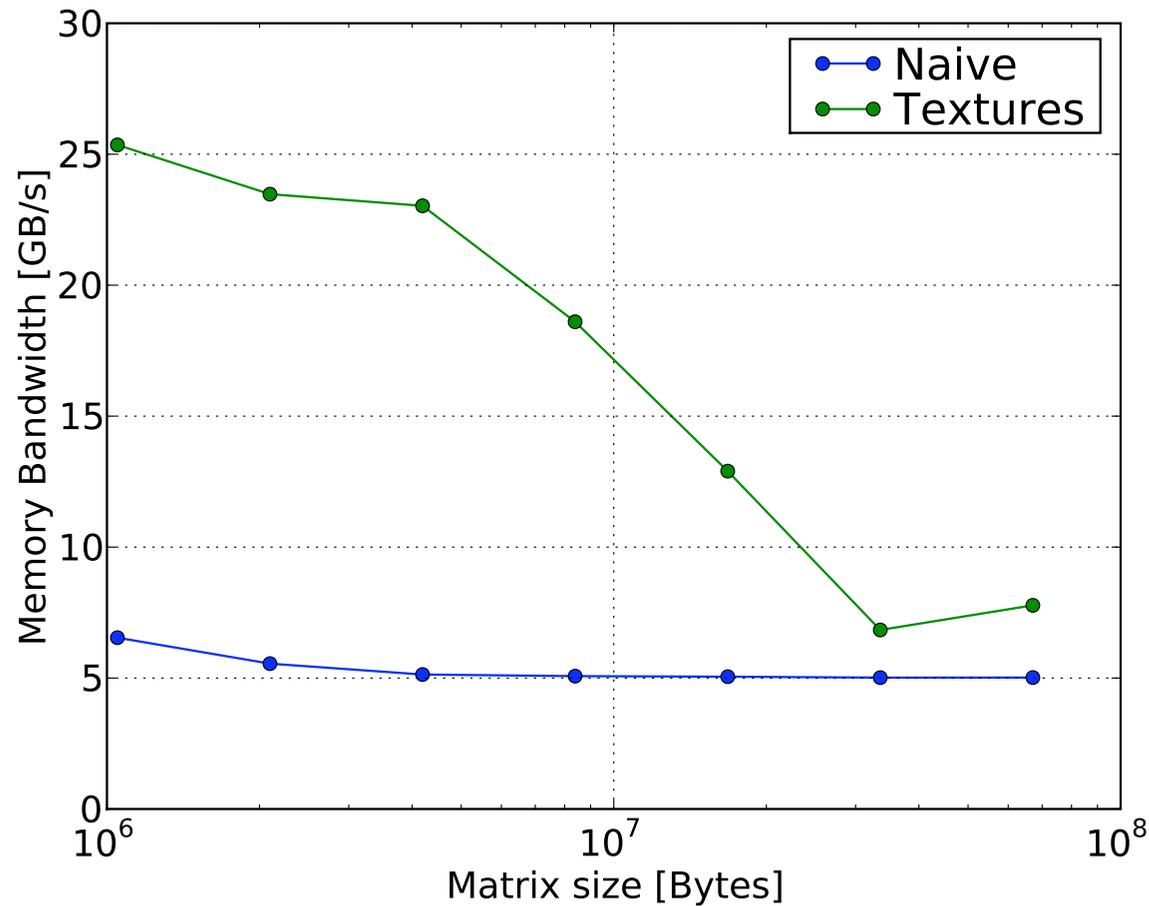


stride: 16 → **16 mem.trans.!**

Let's consider an example



Using just texture(read)+global(write) memory

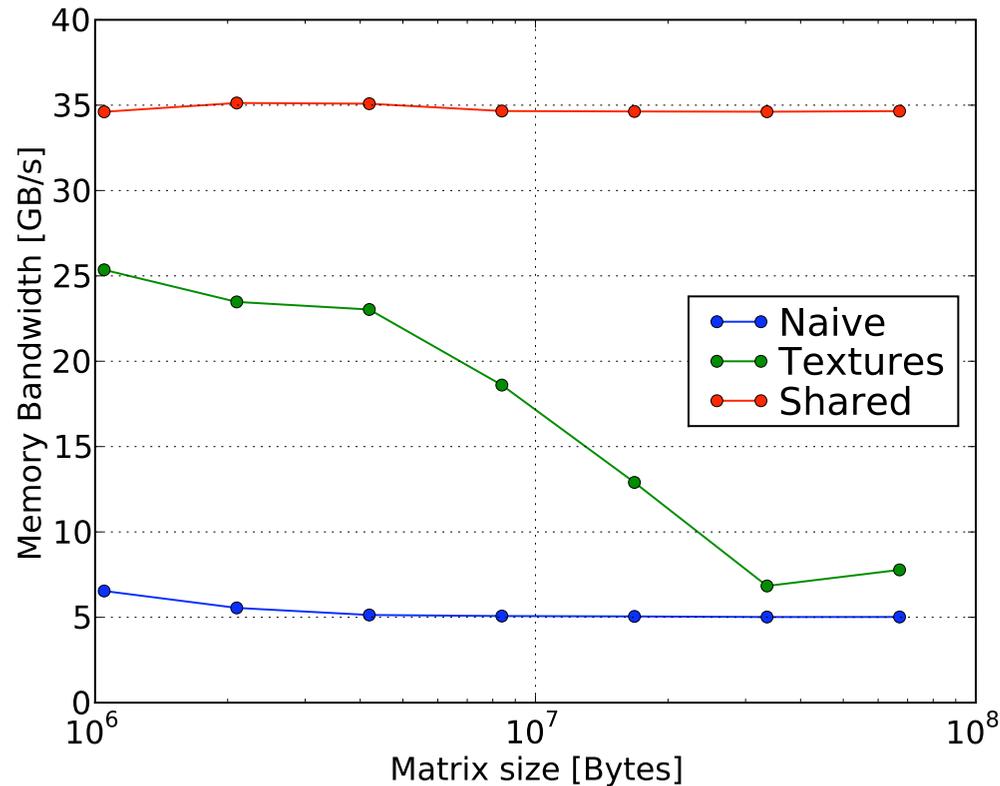


Getting better

Let's consider an example



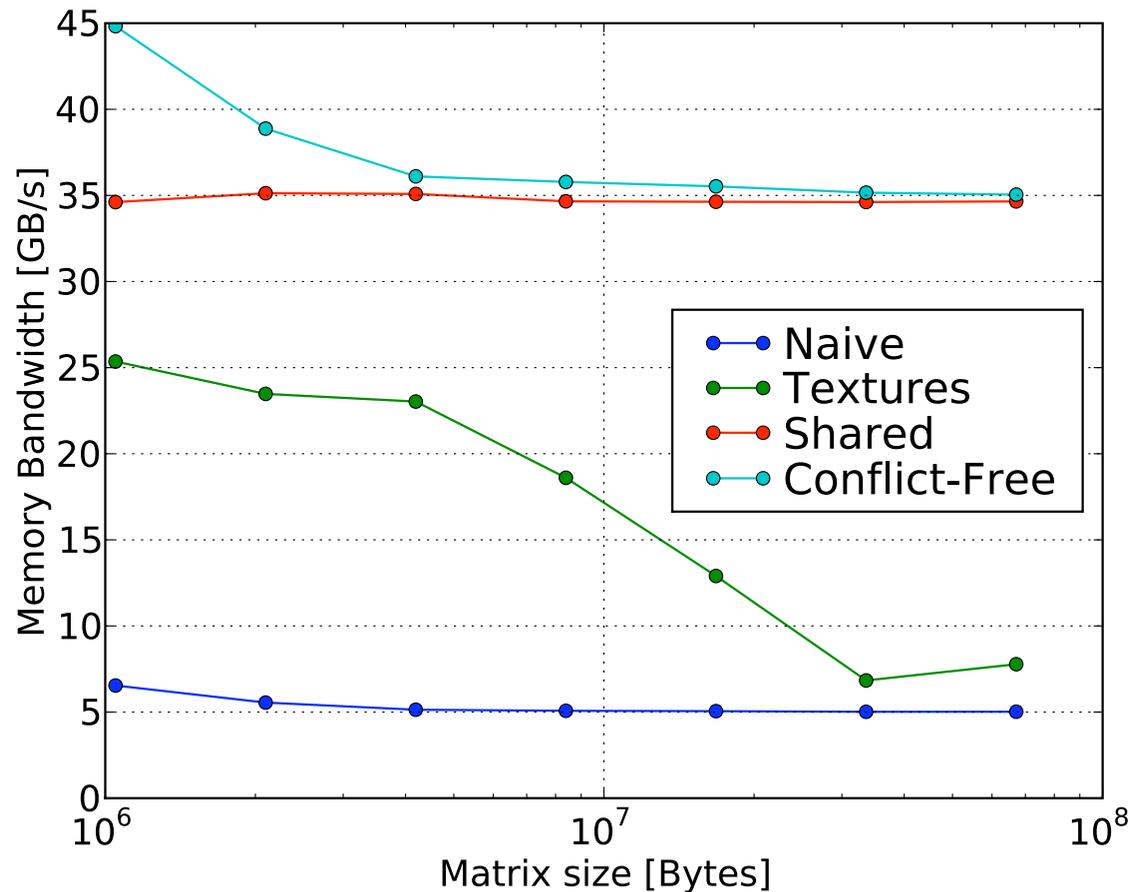
Transpose block-by-block in shared memory -
this does not care about strides



Let's consider an example



Additional improvements are possible for small matrices - bank conflicts in shared memory



A factor of 7-8 over CPU

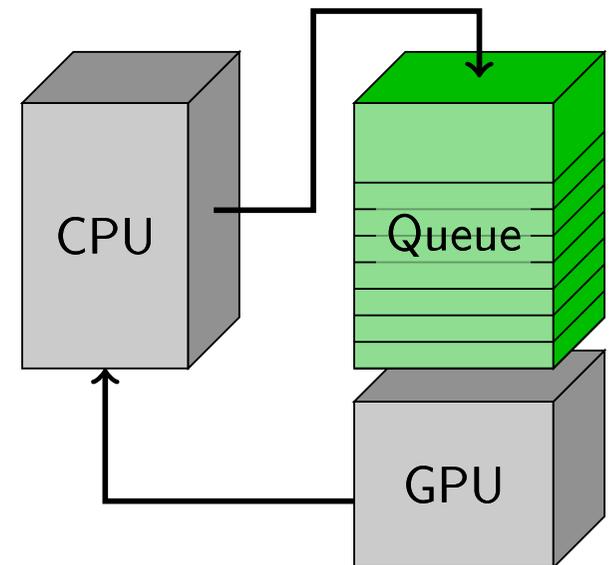
CPU vs GPU



As T. Warburton said in his talk:

*The CPU is mainly the traffic controller
... although it need not be*

- ✓ The CPU and GPU runs asynchronously
- ✓ CPU submits to GPU queue
- ✓ CPU synchronizes GPUs
- ✓ Explicitly controlled concurrency is possible



GPUs overview



- ✓ GPUs exploit multi-layer concurrency
 - ✓ The memory hierarchy is deep
 - ✓ Memory padding is often needed to get optimal performance
 - ✓ Several types of memory must be used for performance
-

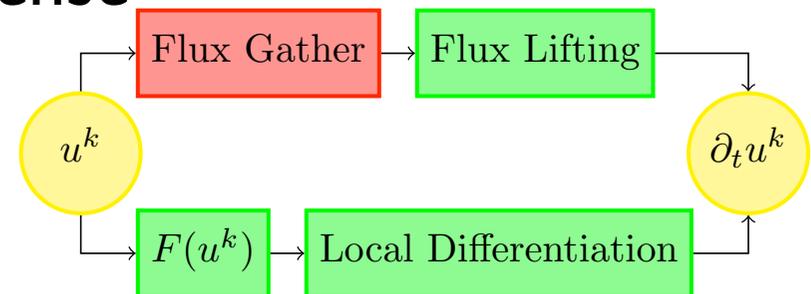
- ✓ First factor of 5 is not too hard to get
- ✓ Next factor of 5 requires quite some work
- ✓ Additional factor of 2-3 requires serious work

Nodal DG on GPU's



So what does all this mean ?

- ✓ GPU's has deep memory hierarchies so local is good
 - ➔ The majority of DG operations are local
- ✓ Compute bandwidth \gg memory bandwidth
 - ➔ High-order DG is arithmetically intense
- ✓ GPU global memory favors dense data
 - ➔ Local DG operators are all dense

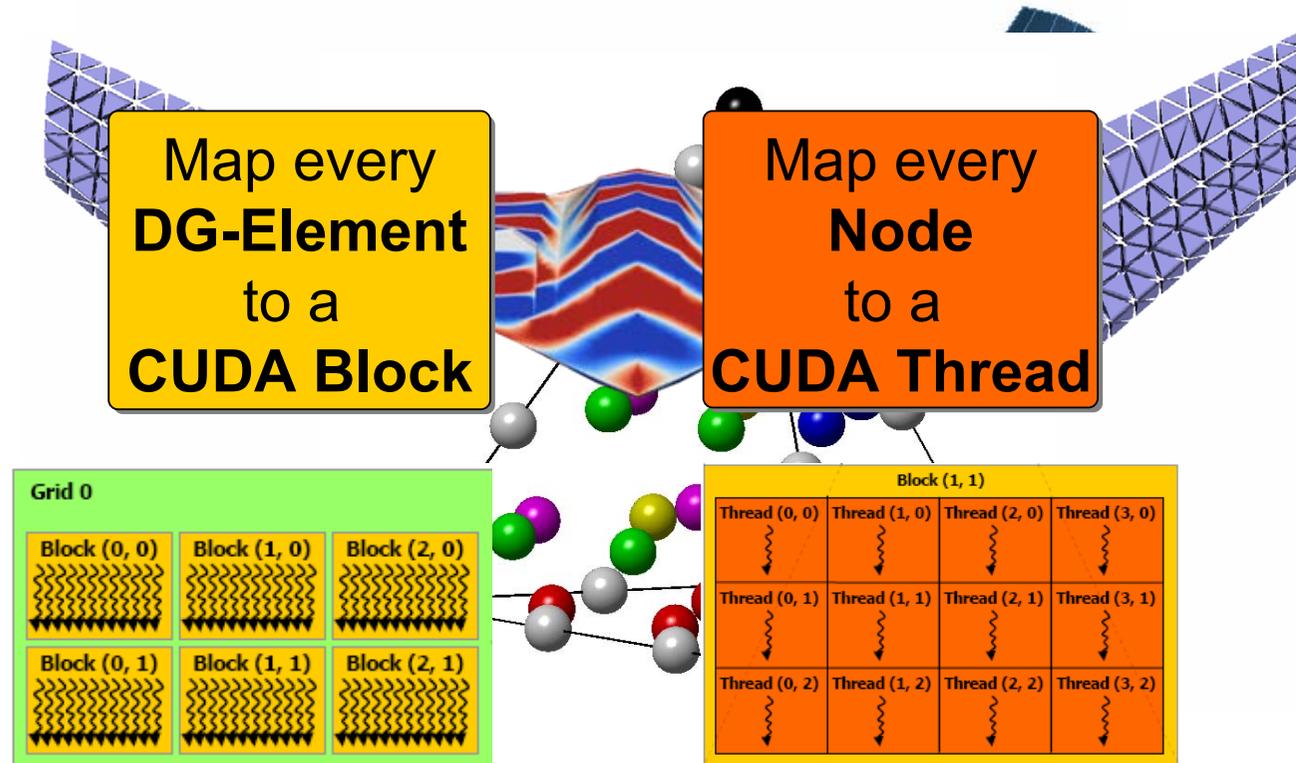


With proper care we should be able to obtain excellent performance for DG-FEM on GPU's

Nodal DG on GPU's



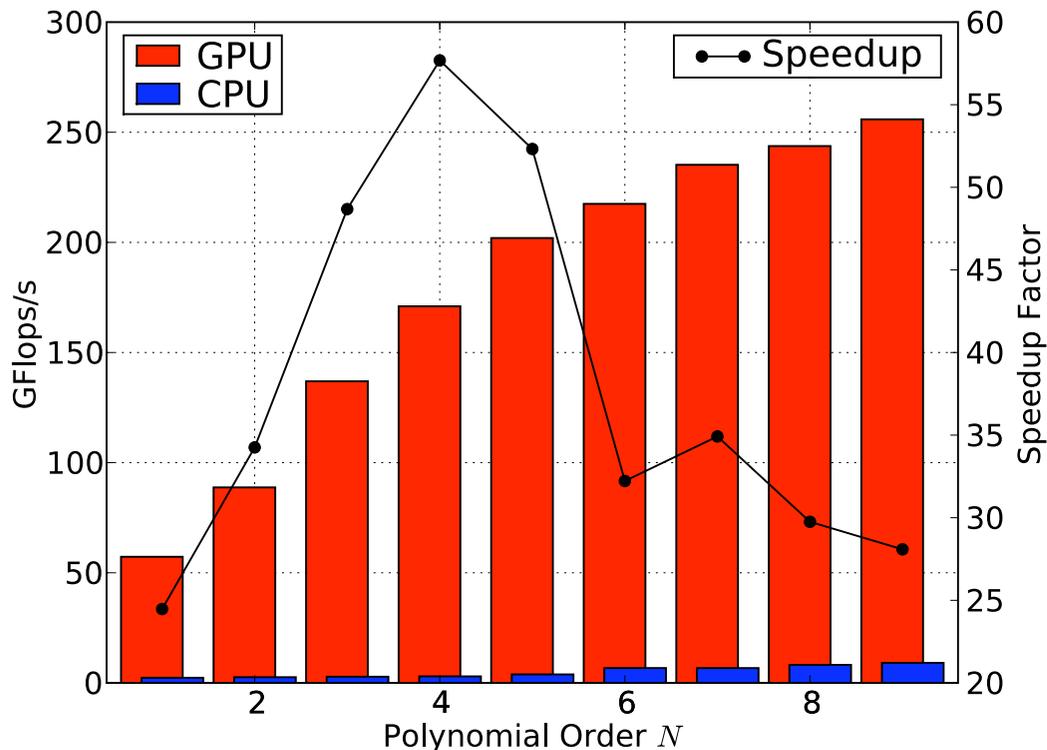
Nodes in threads, elements for blocks



Other choices:

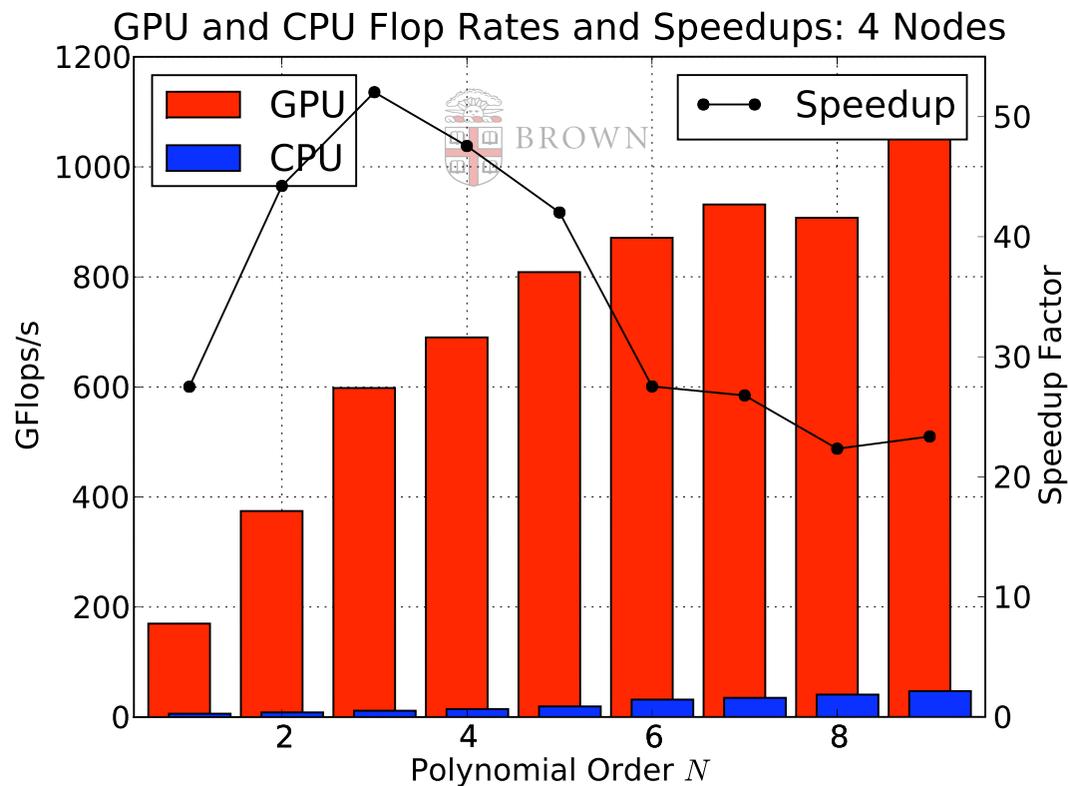
- ✓ D-matrix in shared, data in global (small N)
- ✓ Data in shared, D-matrix is global (large N)

Nodal DG on GPU's



DG-FEM on four GPU
one card

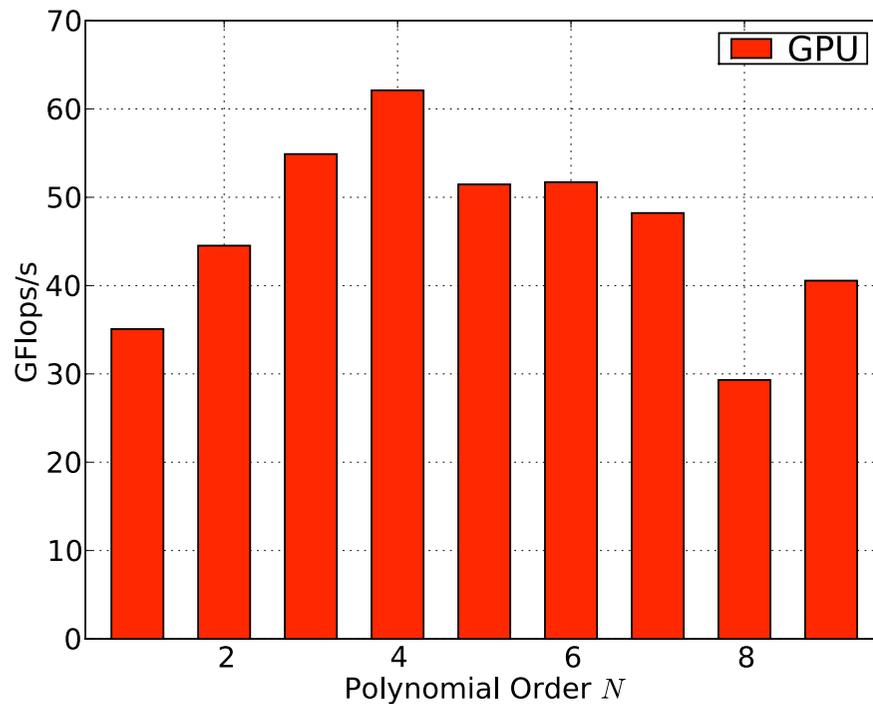
DG-FEM on one GPU





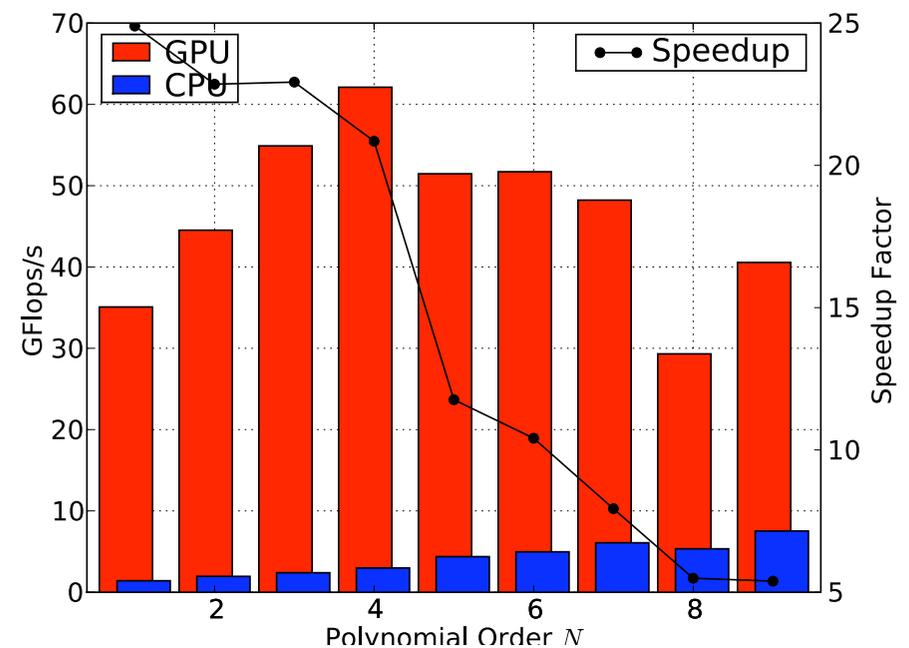
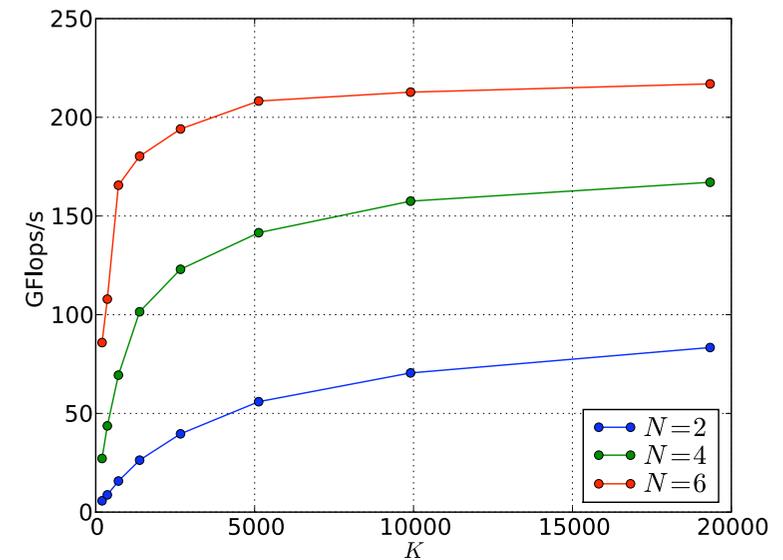
Nodal DG on GPU's

Where you need it most



Also in double precision

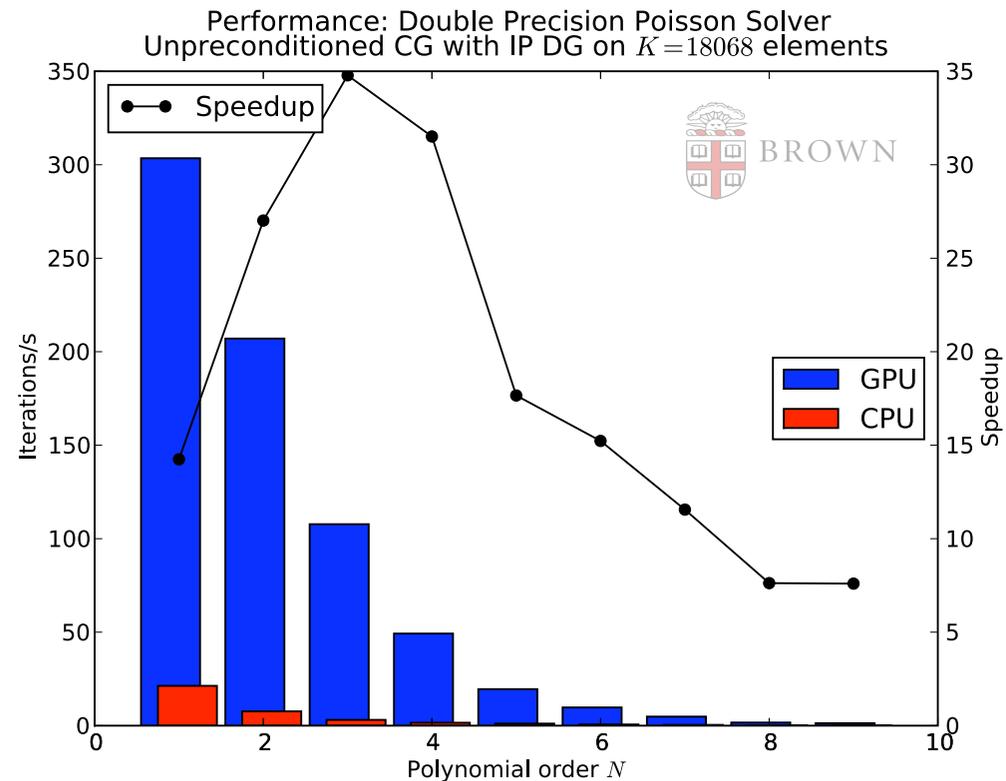
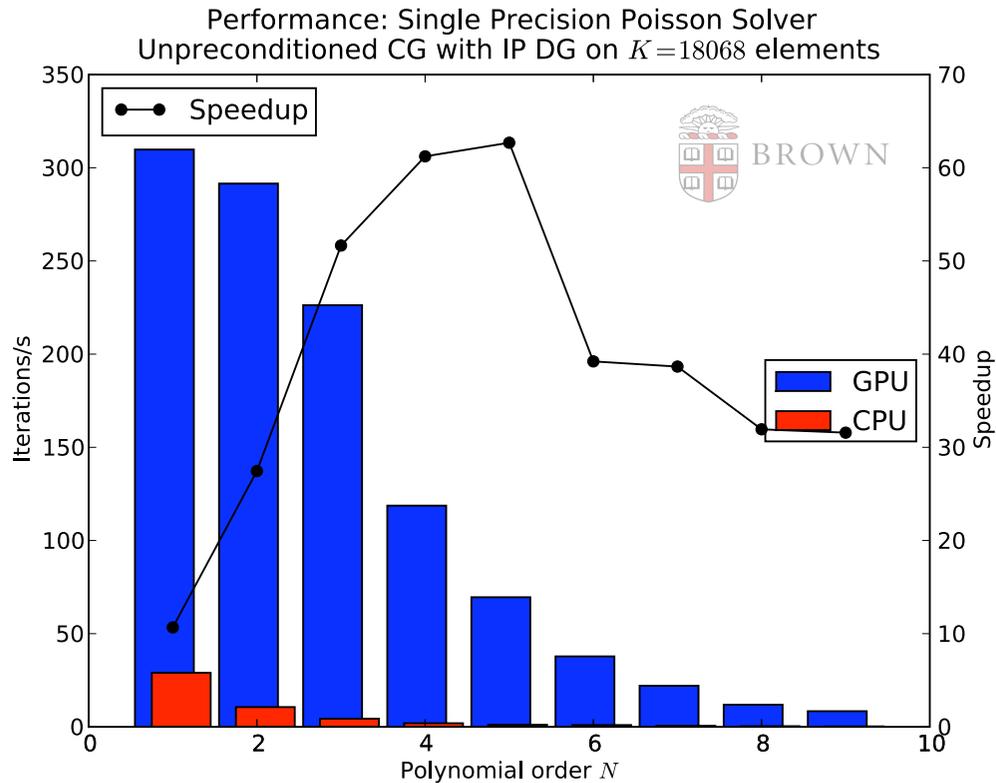
... and for larger and larger grids



Nodal DG on GPU's

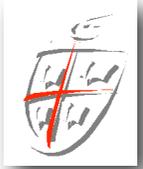


Similar results for DG-FEM Poisson solver with CG



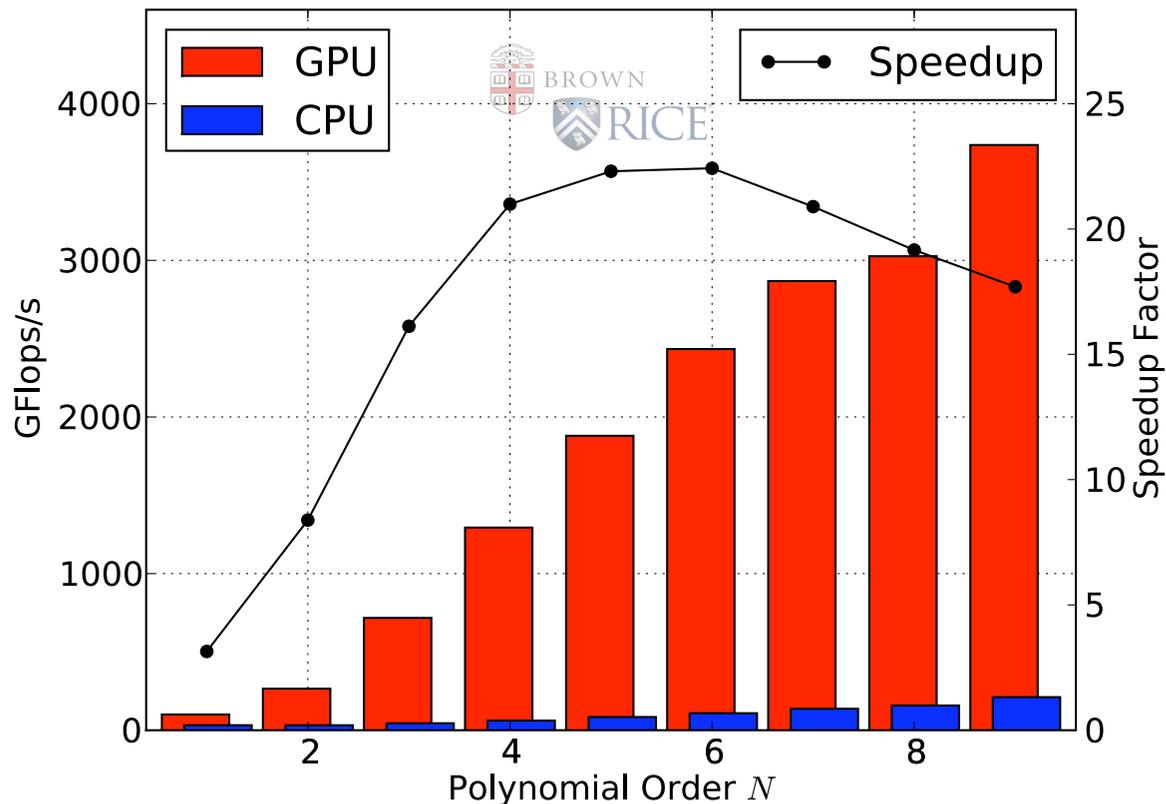
Note: No preconditioning

Combined GPU/MPI solution



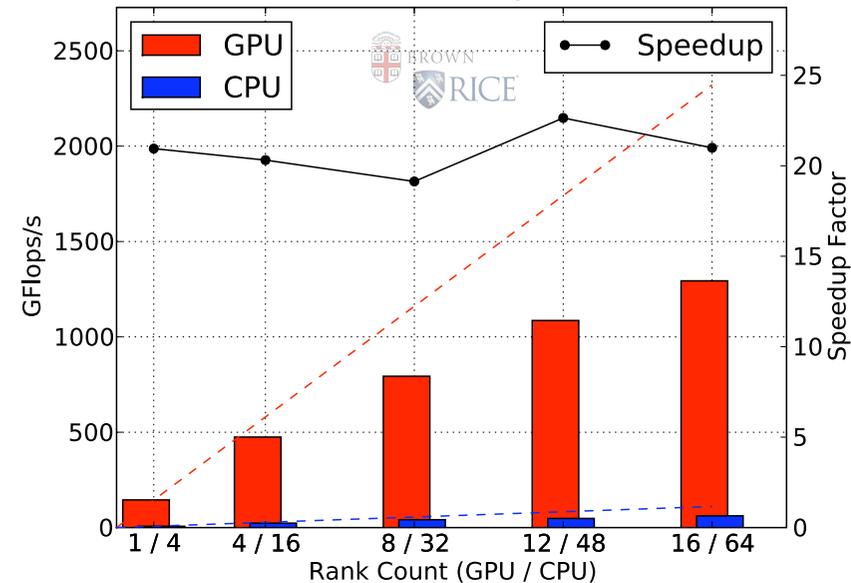
MPI across network

Flop Rates and Speedups: 16 GPUs vs 64 CPU cores

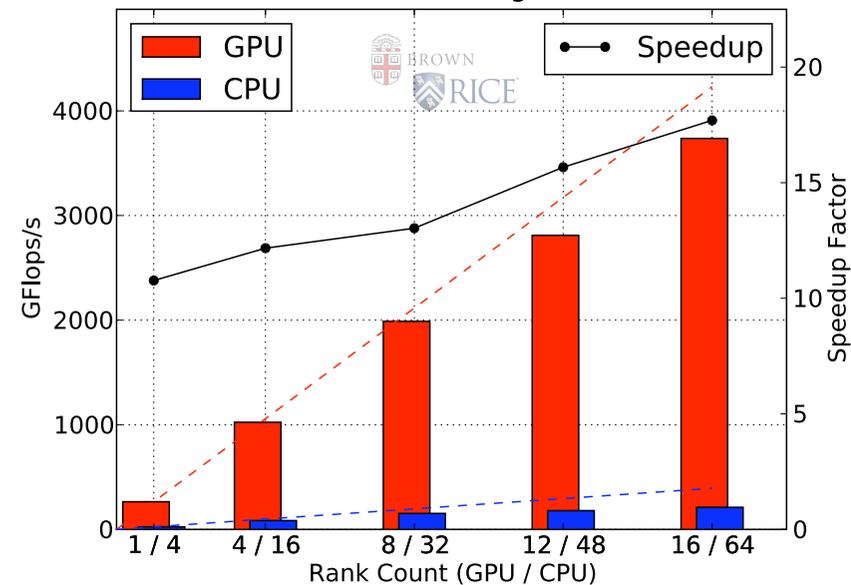


Good scaling when problem is large

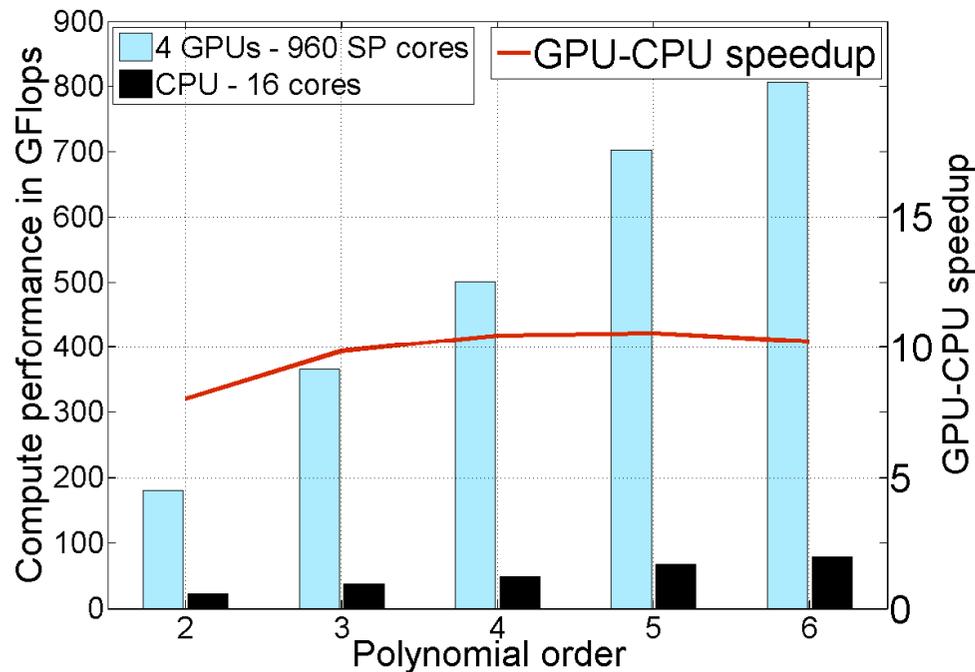
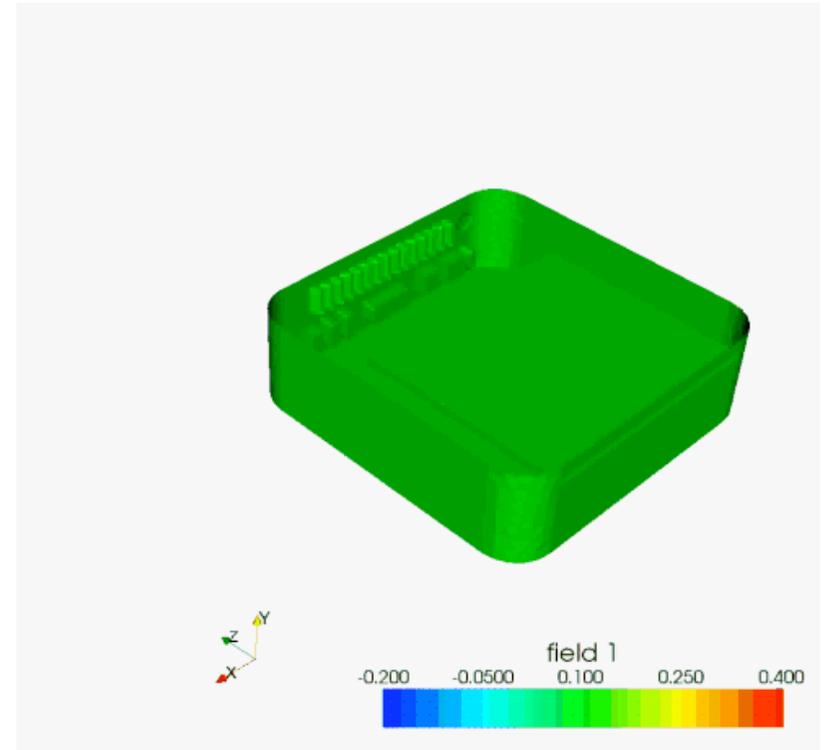
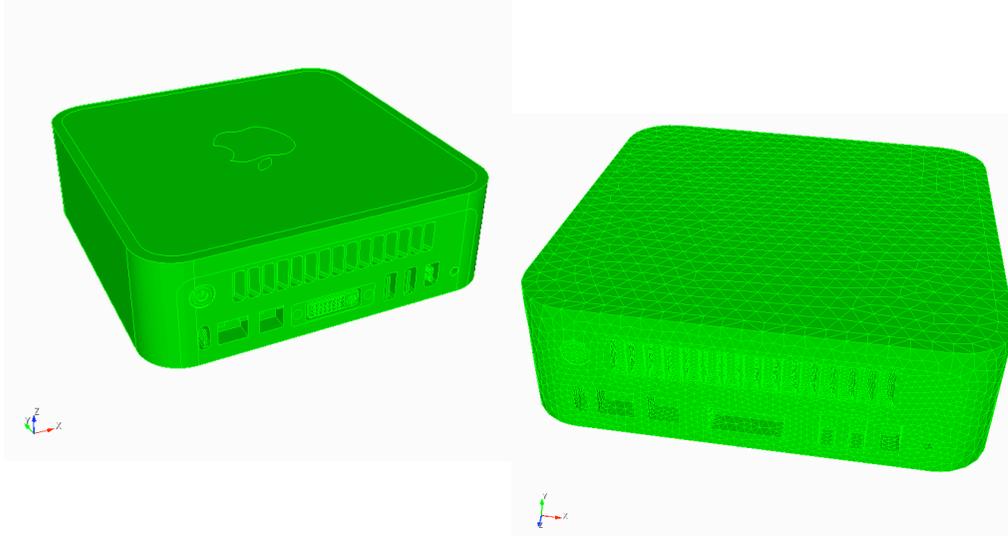
GPU/CPU Weak Scaling: DG Order 4



GPU/CPU Weak Scaling: DG Order 9



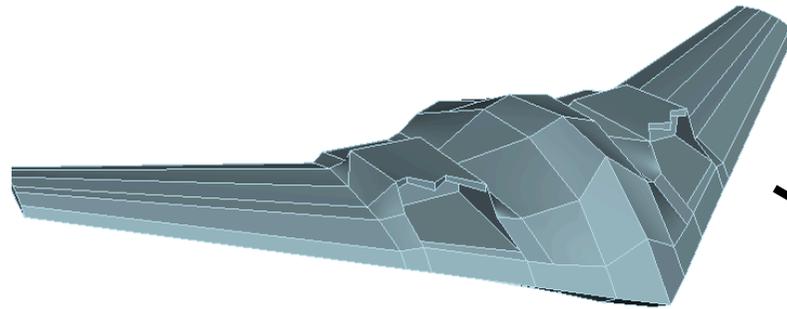
Example - a Mac Mini



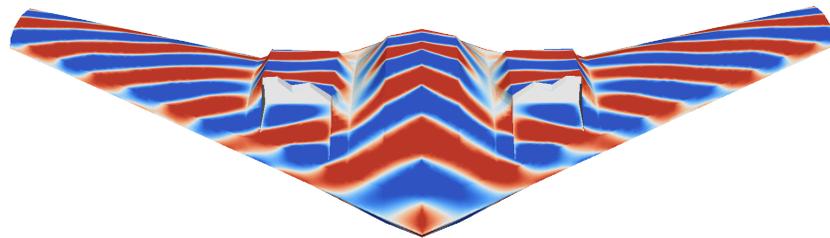
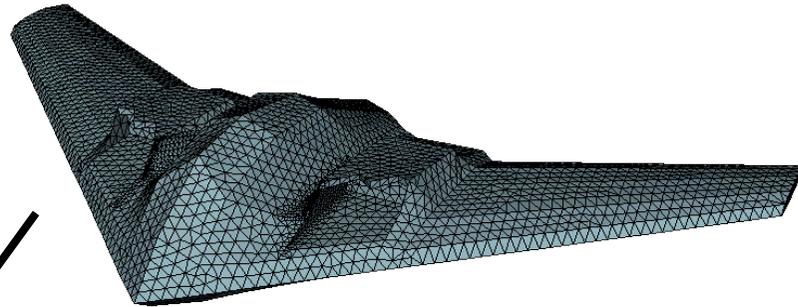
K=201765 elements
3rd order elements

Computation by N. Godel

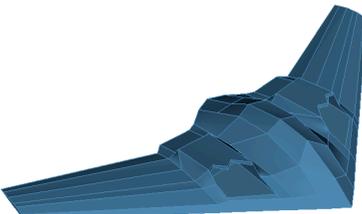
Example: Military aircraft



K=130413 elements
3rd order elements
15.6E6 DoF



Computation by N. Godel

	CPU global	29 h 6 min 46 s	1.0
	GPU global	39 min 1 s	44.8
	GPU multirate	11 min 50 s	147.6

Nodal DG on GPU's



Not just for toy problems

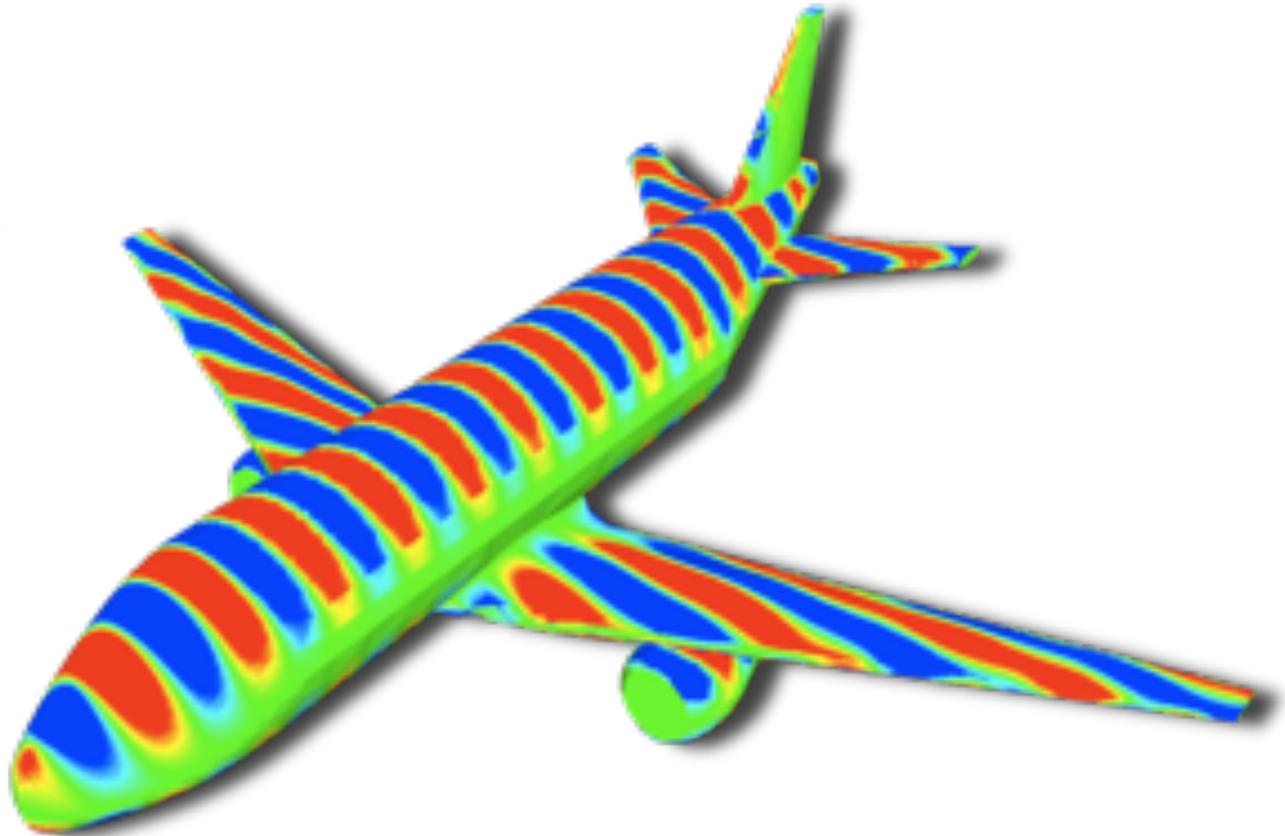
228K elements

5th order elements

78m DOF

68k time-steps

Time ~ 6 hours



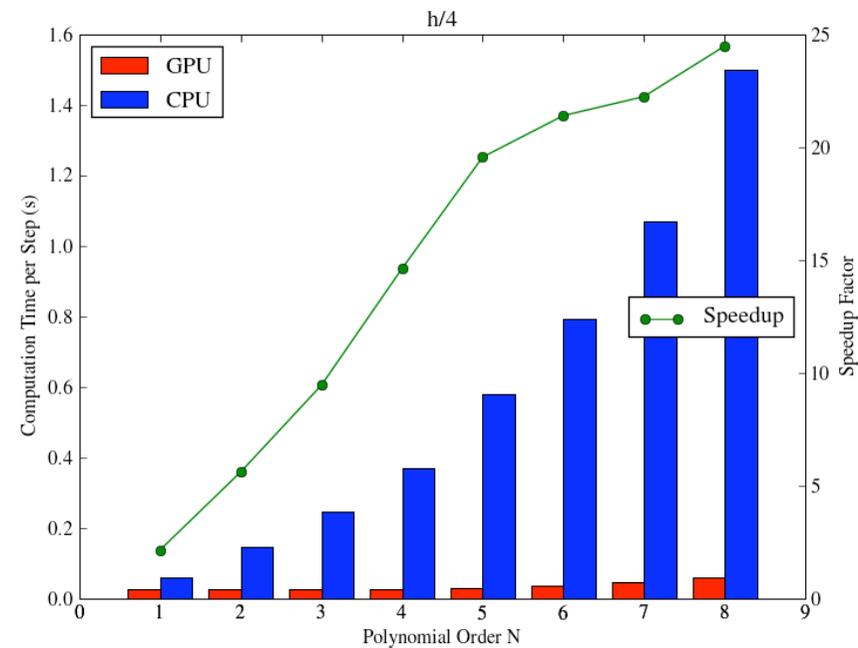
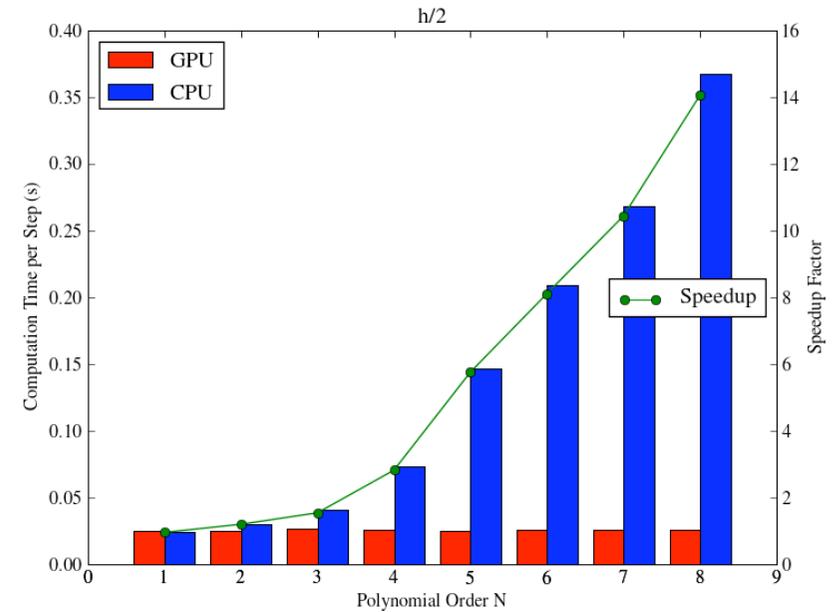
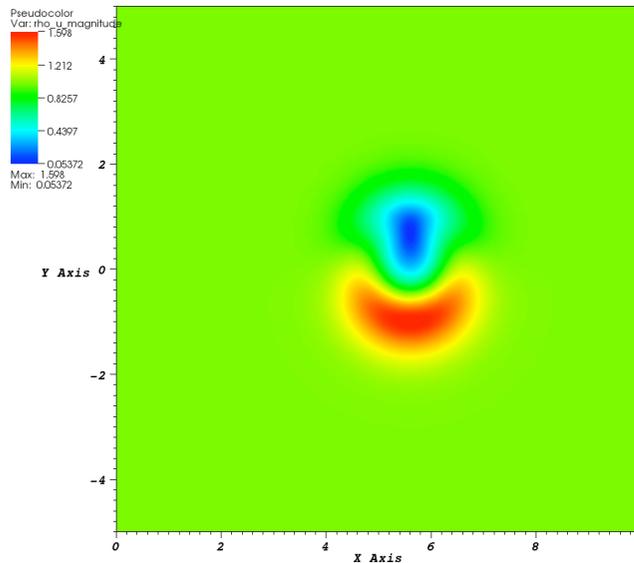
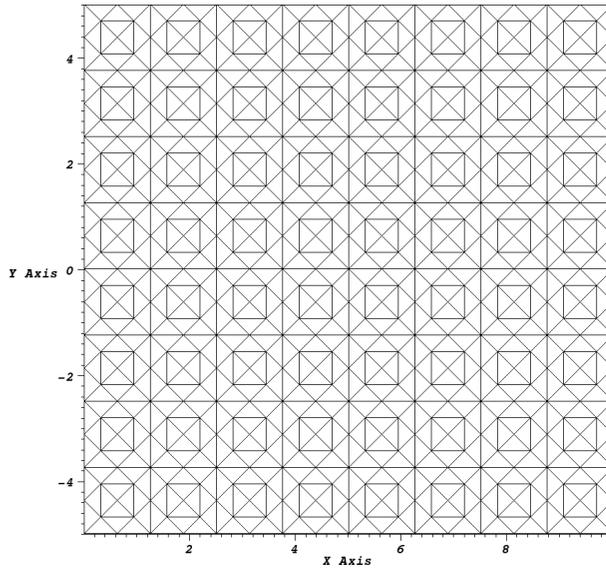
711.9 GFlop/s on one card

Computation by N. Godel

Beyond Maxwell's equations



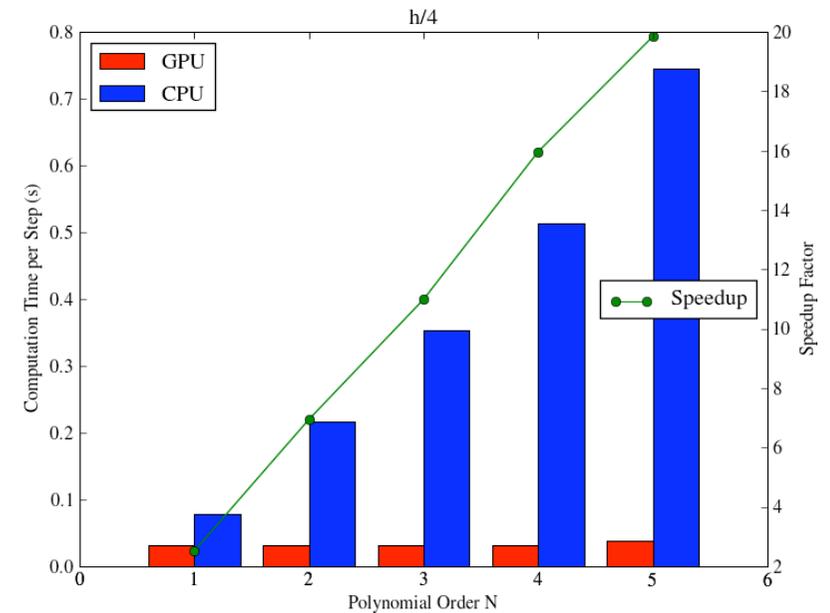
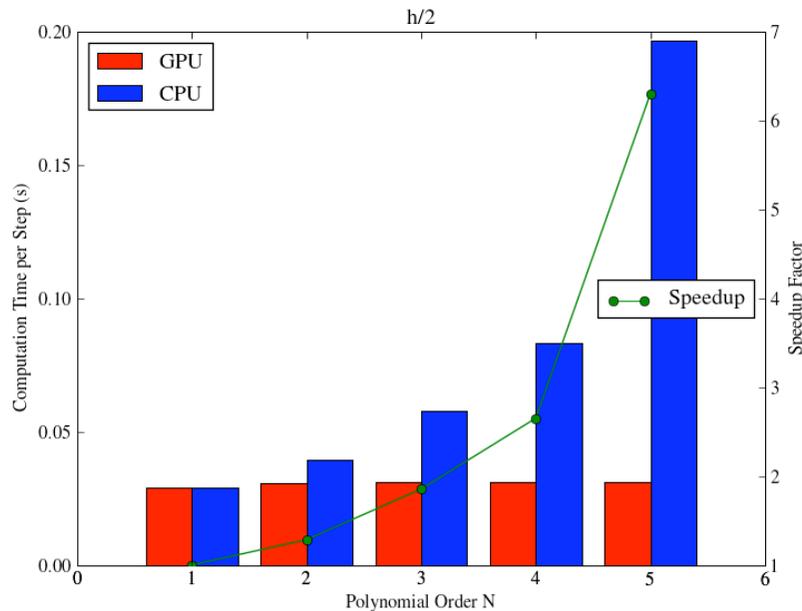
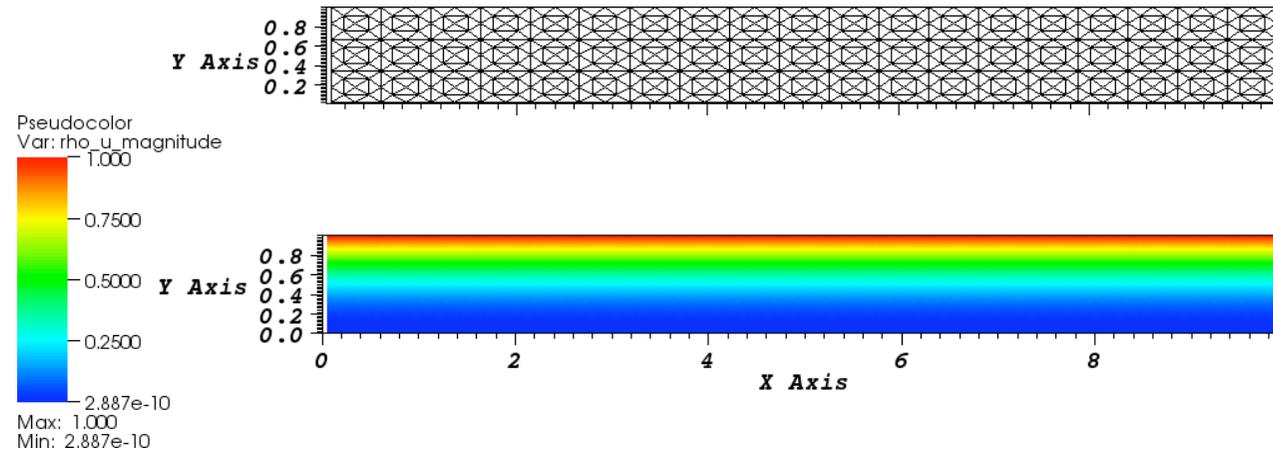
2D Euler test case



Beyond Maxwell's equations



2D Navier-Stokes test case



Want to play yourself ?

The screenshot shows the NVIDIA CUDA Zone website. At the top, there is a navigation bar with the NVIDIA logo, the text "CUDA ZONE", a location dropdown set to "USA - United States", and a search bar. Below the navigation bar are links for "DOWNLOAD CUDA", "WHAT IS CUDA", "DEVELOPING WITH CUDA", "FORUMS", and "NEWS AND EVENTS". The main content area is titled "LATEST CUDA NEWS" and features a grid of 15 article thumbnails. Each thumbnail includes a title, a small image, and a view count. The article "MIDG" is circled in red. The other articles include "Shoot3D Markerless Motion Capture" (70 x), "Concurrent Number Cruncher" (10 x), "GPU4Vision" (GPU4 VISION), "Efficient Computation of Sum Products on GPUs" (270 x), "Low Viscosity Flow Simulations for Animation" (55 x), "Programmer Algorithms-by-Tweak Made easy", "Real-time Digital Holographic Microscopy", "Wait-free Programming for Computations on Graphics Processors", "Real-time Visual Tracker by Stream Processing" (10 x), "Ray Casting Deformable Models", "Molecular Dynamics of DNA and Liquids" (18 x), "GPUGRID.NET", "Mixed Precision Linear Solvers" (27 x), and "Accelerating Density Functional Calculations with GPU" (40 x).

Code MIDG available at <http://nvidia.com/cuda>

Note: Single precision - double to come

Do we have to write it all ?

No :-)

- ✓ Book related codes - all at www.nudg.org
- ✓ Matlab codes
- ✓ NUDG++ - a C++ version of 2D/3D codes (serial)
- ✓ hedge - a Python based meta-programming code.
Support for serial/parallel/GPU
- ✓ MIDG - a bare bones parallel/GPU code for
Maxwell's equations

Do we have to write it all ?

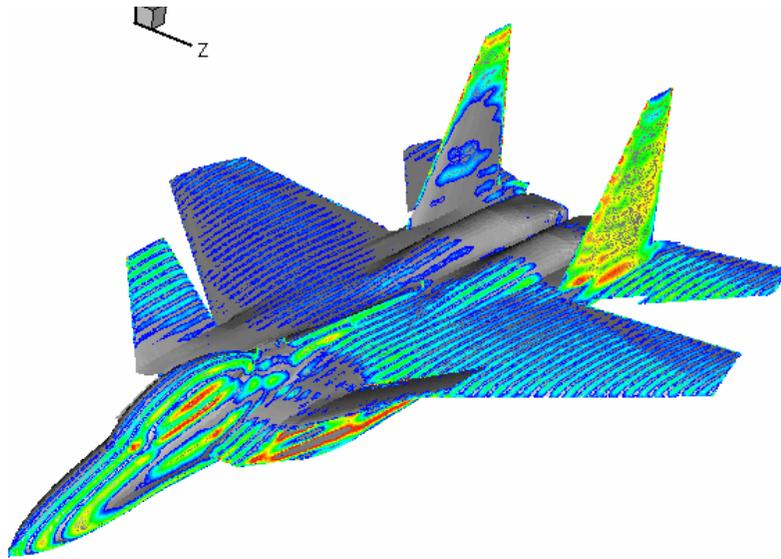
Other codes

✓ **Slegde++** - C++ operator code. Interfaced with parallel solvers (Trilinos and Mumps) and support for adaptivity and non-conformity. Contact Lucas Wilcox (UT Austin/ICES)

✓ **deal.II** - a large code with support for fully non-conforming DG with adaptivity etc. Only for squares/cubes. www.dealii.org

✓ **Nektar++** - a C++ code for both spectral elements/hp and DG. Mainly for CFD. Contact Prof Spencer Sherwin (Imperial College, London)

Progress ?



Year 2001

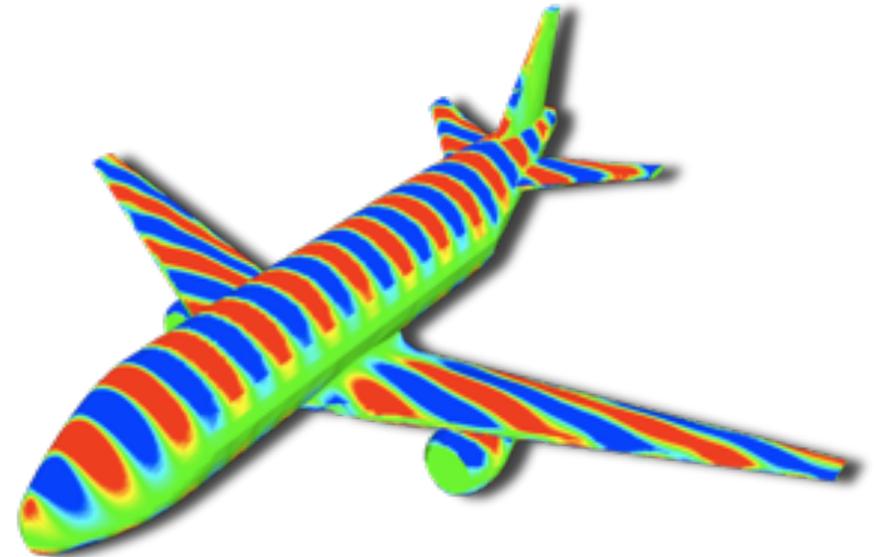
250k tets, 4th order
50m dof, 100k timesteps

24 hours on 512 procs

Year 2008

82k tets, 4th order
17m dof, 60k timesteps

Few hours on GPU



Thanks !

Many people have contributed to this with material, figures, examples etc

- ✓ Tim Warburton (Rice University)
- ✓ Lucas Wilcox (UT Austin)
- ✓ Andreas Kloeckner (Brown)
- ✓ Nico Goedel (Hamburg)
- ✓ Hendrick Riedmann (Stuttgart)
- ✓ Francis Giraldo (NRL, Monterrey)

... and to you for hanging in there !