

# Discontinuous Galerkin solver design on hybrid computers.

Inria Tonus & IRMA Université de Strasbourg

September 2, 2014

- 1 Vlasov and hyperbolic systems of conservation laws.
- 2 Efficient conservation laws solver.
- 3 Discontinuous Galerkin (DG) solver design.

1) The  $(x, v)$  Vlasov equation is an  $x$ -only hyperbolic system of conservation laws.

# Vlasov reduction

We consider the one-dimensional Vlasov-Poisson model

$$\partial_t f + v \partial_x f + E \partial_v f = 0, \quad (1)$$

$$\partial_x E = -1 + \int_v f dv, \quad (2)$$

where  $f(x, v, t)$  is the distribution function,  $E(x, t)$  is the electric field.

We consider the space-periodic case

$$f(0, v, t) = f(L, v, t), \quad \frac{1}{L} \int_x \int_v f(x, v, 0) = 1,$$

$$\int_{x=0}^L E dx = 0.$$

We also suppose that

$$v \in ]-\infty; \infty[.$$

We consider a finite number of independent velocity functions  $\{\varphi_k(v), k = 1 \dots P\}$  and expand  $f$  on this basis

$$\begin{aligned} f(x, v, t) &\simeq \sum_{j=1}^P w^j(x, t) \varphi_j(v) \\ &= w^j(x, t) \varphi_j(v) \text{ (sum on repeated indices).} \end{aligned} \quad (3)$$

The unknown scalar  $f(x, v, t)$  is replaced by the unknown vector

$$W(x, t) = \left( w^1(x, t), w^2(x, t), \dots, w^P(x, t) \right)^T.$$

Armstrong, 1967 [1], Tang & *al.* 1992, 1993 [12, 13], Schumer, Holloway, 1998 [11], le Bourdieu, de Vuyst, Jacquet 2006 [7]...

We introduce the expansion (3) in the Vlasov equation (1), multiply by  $\varphi_i$  and integrate with respect to  $v$ . We obtain

$$M\partial_t W + A\partial_x W + B(E)W = 0, \quad (4)$$

where

$$M_{ij} = \int_v \varphi_i \varphi_j, \quad A_{ij} = \int_v v \varphi_i \varphi_j, \quad B(E)_{ij} = E \int_v \varphi_i \partial_v \varphi_j.$$

$M$  is symmetric positive,  $A$  is symmetric. The system (4) is thus a hyperbolic system of conservation laws ( $M^{-1}A$  has real eigenvalues).

$$\partial_t W + \partial_x F(W) + S(W) = 0.$$

$$F(W) = M^{-1}AW, \quad S(W) = M^{-1}B(E)W.$$

# General hyperbolic system

The approach can be extended to gyrokinetic models, higher dimensions, etc. [5]. Hyperbolic system theory is also useful for MHD, Maxwell, fluids, etc.

General case:  $W(X, t)$ ,  $X = (x_1, \dots, x_d)$ .

$$\partial_t W + \partial_{x_i} F^i(W) + S(W) = 0.$$

Flux:  $n = (n_1 \cdots n_d)$ ,  $F(W, n) = F^i(W)n_i$ . Hyperbolicity: the eigenvalues of  $D_W F(W, n)$  are real.

II) Comparison of several implementations of a very simple conservation laws solver.



# A simple conservation laws solver

Solution  $W(x, y, t) \in \mathbb{R}^m$  of Maxwell/fluids/MHD/Vlasov equations  $X = (x, y)$ . System of conservation laws.

$$\partial_t W + \partial_x F^x(W) + \partial_y F^y(W) = 0.$$

Approximation  $W_{i,j}^n$  of  $W(i\Delta x, j\Delta y, n\Delta t)$ . Finite volume method + Strang splitting

$$\frac{W_{i,j}^* - W_{i,j}^n}{\Delta t} + \frac{F_{i+1/2,j}^{x,n} - F_{i-1/2,j}^{x,n}}{\Delta x} = 0,$$

$$\frac{W_{i,j}^{n+1} - W_{i,j}^*}{\Delta t} + \frac{F_{i,j+1/2}^{y,n} - F_{i,j-1/2}^{y,n}}{\Delta y} = 0.$$

Numerical flux:  $F_{i+1/2,j}^{x,n} = F_{\text{num}}^x(W_{i,j}^n, W_{i+1,j}^n)$ ,  
 $F_{i,j+1/2}^{y,n} = F_{\text{num}}^y(W_{i,j}^n, W_{i,j+1}^n)$ .

On large grids ( $> 1024 \times 1024$ ). We compare:

- a naive C implementation without optimization on a CPU single core;
- the same program, but compiled with optimizations;
- the same program with an additional optimization (tiling for optimizing data locality);
- the same program with OpenMP parallelization on a 16-core CPU.

# Comparison

Implementation	Time	Speedup
Naive code	30 days	1
Naive code + optim. compil.	146 h	5
Naive code + optim. compil. + tiling	97 h	8
OpenMP version (16 cores)	6.2 h	116

OpenCL model: an accelerator device (GPU or CPU) is made of

- Global memory (typically 2 GB);
- Compute units (typically 30).

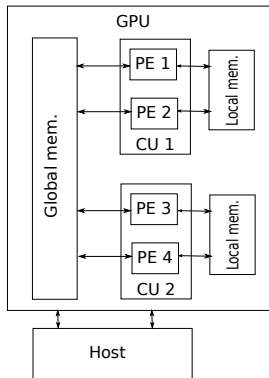
Each compute unit is made of

- Processing elements (typically 8);
- Local memory (typically 32 kb).

The same program (kernel) can be executed on all the processing elements at the same time.

- All the processing elements have access to the global memory.
- The processing elements have only access to the local memory of their compute unit.
- If two processing elements write at the same location at the same time, only one wins...
- The access to the global memory is slow while the access to the local memory is fast (generally...)

A (virtual) GPU with 2 Compute Units and 4 Processing Elements



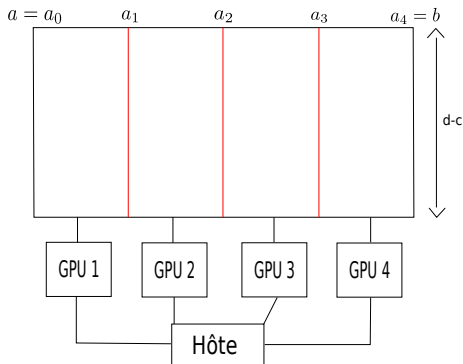
- OpenCL: “Open Computing Language”. Library of C functions, called from the host, in order to drive the GPU + C-like language for the kernels. Available since 2009. Specification managed by the Khronos Group (OpenGL).
- Virtually, it allows to have as many compute units (work-groups) and processing elements (work-items) as needed.
- The threads are sent to the GPU thanks to a mechanism of command queues on the real compute units and processing elements. OpenCL manages events and a task graph for asynchronous out-of-order operations.
- Portable: the same program can run on a multicore CPU or a GPU. Drivers exist for: AMD CPU and GPU, Intel CPU and GPU, MIC, ARM, IBM Power7, StarPU, etc.

# Implementation of the splitting scheme

We organize the data in a  $(x, y)$  grid and for each time step:

- we associate a work-item to each cell of the grid and a work-group to each row.
- we compute the fluxes balance in the  $x$ -direction for each cell of each row of the grid.
- we transpose the grid (exchange  $x$  and  $y$ ) with an optimized memory transfer algorithm [10] (see also [9]).
- we compute the fluxes balance in the  $y$ -direction for each row of the transposed grid. Memory access are optimal.
- we transpose again the grid.

We apply a subdomain decomposition for multi-GPU computing





# Comparison

Implementation	Time	Speedup
Naive code	30 days	1
Naive code + optim. compil.	146 h	5
Naive code + optim. compil. + tiling	97 h	8
OpenMP (CPU Intel 8x2 cores)	6.2 h	116
OpenCL (CPU Intel 6x2 cores)	18 h	40
OpenCL (NVIDIA Tesla K20)	20 min	2160
OpenCL (AMD Radeon HD 7970)	16 min	2650
OpenCL + MPI (4 x AMD Radeon HD 7970)	5 min	7848

Essential: no test, optimized transposition (10x faster than the naive memory access)


## Well known conclusions

- GPU computing can be really faster,
- But: no test in OpenCL or CUDA kernels,
- Memory access are essentials.

# Shock-bubble interaction

Two-phase flow conservation laws system. Density  $\rho$ , velocity  $(u, v)$ , internal energy  $e$ , gas mass fraction  $\phi$ .

$$W = (\rho, \rho u, \rho v, \rho(e + u^2/2 + v^2/2), \rho\phi).$$



<http://www.youtube.com/watch?v=c8hcqihJzbu>


# Numerical results

$$t_{\max} = 0.45 \text{ ms}$$


Grid:  $40,000 \times 20,000$  (4 billions unknowns for each time step) [6]

GPU time: 30 h ( $10 \times$  NVIDIA K20)


# Density



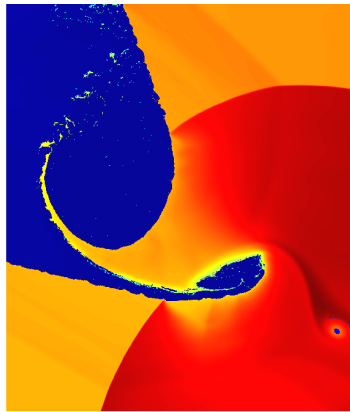
# Zoom 1



# Zoom 2



RHO  
693  
1.27 1.38e+03 0.978




RHO  
693  
1.27 1.39e+03

# MagnetoHydroDynamics (MHD)

$W = (\rho, \rho u, \rho v, \rho w, B_x, B_y, B_z, \rho(e + u^2/2 + v^2/2))$ . Velocity  $(u, v, w)$ , magnetic field  $(B_x, B_y, B_z)$ .

Orszag-Tang vortex (grid size up to  $15000 \times 15000$ ) [8].





III) A more versatile approach: the Discontinuous Galerkin (DG) method.

Generalization of the FV method, DG method in a 3D space,  $X = (x, y, z)$ . We consider a mesh of the computational domain. In a cell  $L$  of the mesh, the field is approximated by polynomial basis functions (sum on repeated indices)

$$W(X, t) = W_L^j(t)\varphi_j^L(X), \quad X \in L.$$

The numerical solution satisfies the DG approximation scheme

$$\forall L, \forall i \quad \int_L \partial_t W \varphi_i^L - \int_L F(W, \nabla \varphi_i^L) + \int_{\partial L} F(W_L, W_R, n_{LR}) \varphi_i^L = 0.$$

- $R$  denotes the neighbor cells along  $\partial L$ .
- $n_{LR}$  is the unit normal on  $\partial L$  oriented from  $L$  to  $R$ .
- $F(W_L, W_R, n)$  is the numerical flux (which satisfies  $F(W, W, n) = F^x(W)n_x + F^y(W)n_y$ ).
- Time integration of a system of ordinary differential equations.  
Mass matrix  $M_{ij}^L = \int_L \varphi_i^L \varphi_j^L$

advantages:

- varying order, mesh refinement.
- local stencil.
- high order  $\Rightarrow$  high amount of local computations.
- many optimizations for hexahedrons meshes.
- MIMD/SIMD parallelism. Subdomains (MPI), elementary computations (OpenMP, CUDA, OpenCL) [3].

possible issues:

- memory access (unstructured mesh).
- branch tests in kernels (physical models, boundary conditions, etc.)
- MPI communications imply GPU $\leftrightarrow$ Host memory transfers.

CLAC means Conservation Laws Approximation on many Cores. It is a C++ library developed with AxesSim company in Strasbourg. It is actually used for actual electromagnetic simulations.

- “Reasonable” C++: a few templates, almost no inheritance.
- Google coding rules <http://google-styleguide.googlecode.com/svn/trunk/cppguide.xml>
- Git on Inria GForge.
- Doxygen.
- scons.
- Boost: unit tests, graphs.

Would be useful: continuous integration...

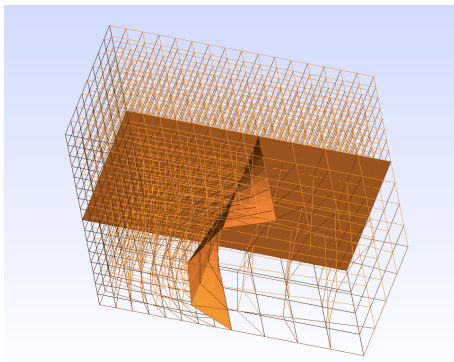
- Subdomain decomposition: each domain is associated to a MPI node. Each MPI node is associated to an OpenCL device (CPU or GPU).
- Zone decomposition: each subdomain is split into volume zones and interface zones. A zone possess identical elements (same order, same geometry, same physical model). A computation kernel is compiled for each zone (for avoiding branch tests).
- (Simple) non-conformity between zones is allowed.
- Geometry and interpolation are separated (possibility to replace memory access by computations).

# CLAC data structure

Example of a domain made of two subdomains, three volume zones and three interface zones. the mesh is non-conforming.

Subdomain 0: only one big refined volume zone. Two interface zones.

Subdomain 1: two small volume zones (coarse and refined). Three interface zones.



- Numerical integration: Gauss-Legendre integration points  $G_k^L$ ,  $G_k^{\partial L}$  and weights  $\omega_k^L$ ,  $\omega_k^{\partial L}$  on hexahedrons

$$\int_L h(X) dX \simeq \sum_k \omega_k^L h(G_k^L).$$

Nodal basis function  $\varphi_i^L(G_k^L) = \delta_{i,k}$ .

- Several optimizations: diagonal mass matrix, complexity  $(d+1)^3 \rightarrow 3(d+1)$ , etc. [4]

# Implementation details

- A single kernel for  $\partial L$  and  $L$  integration steps. Intermediate results stored in the cache of the compute unit. One processor per Gauss point. The number of Gauss points is different on  $\partial L$  and  $L \Rightarrow$  some processors are idling in the volume integration step.
- A function class (pointer to the actual function + headers and sources: needed for the OpenCL compilation at runtime). We generally hide memory access into function calls.
- Customized kernels are assembled and compiled for each zone.



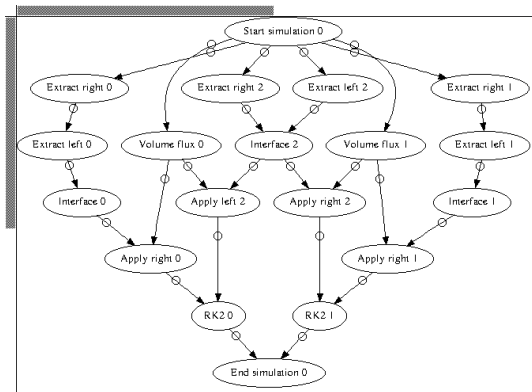
Important kernels:

- Volume zone: internal fluxes and sources assembly (“Volume flux”).
- Interface zone: field extraction from right and left volume zones (“Extract left”, “Extract right”). The extraction may imply MPI communications.
- Interface zone: flux computations (“Interface”).
- Interface zone: boundary fluxes assembly on the left and right volume zones (“Apply right” or “Apply left”).
- Volume zones: RK2 integration step (“RK2”).

GPU-host transfers occur only in the extraction task at an interface between two subdomains.

# Task dependency graph

For the moment, we use Boost Graph.




Problem: how to express the dependency between MPI and OpenCL operations ?

- We decided to rely only on the OpenCL events management.
- The beginning of a task depends on the completions of a list of OpenCL events. The task is itself associated to an OpenCL event.
- At an interface zone between two subdomains, an extraction task contains a GPU to host memory transfer, a MPI send/receive communication and a host to GPU transfer.
- we create an OpenCL user event, and launch a MPI blocking sendrecv in a thread. At the end of the communication, in the thread, the OpenCL event is marked as completed. Using threads avoids blocking the main program flow.
- Efficiency: test on 4 NVidia GTX 780. Synchronous mode, speedup  $\simeq 2.6$ . Asynchronous mode, speedup  $\simeq 4.15$


$W = (E_x, E_y, E_z, H_x, H_y, H_z)$ : electric and magnetic field. Maxwell equations.

Plane wave with gaussian profile.





- Aircraft geometry described with 3,337,875 hexaedrons ( $\simeq 1$  billion unknowns per time step). Several PML layers at the boundaries.
- We use 8 GPUs to perform the computation. The simulation does not fit into a single GPU memory. 400 Gflops.
- In this test case we spend about 30% of the computation time in the memory transfers between the CPU and the GPU and about 20% in the MPI communications. We expect much better speedups with the asynchronous task graph.

# Aircraft



# Waveguide filter






# Waveguide filter

- 8370 cells, 3 millions unknowns (second order nodal interpolation).
- 10 cPML layers on two sides.
- $dt = 1.36e-13s$  et  $dx = 8.14316e-04m$ .  $t_{max} = 25e-9s$ .  
184076 iterations.  $19.5mm \times 9.525mm$ .
- GPU time 4092s on one NVIDIA GTX 680.








# Waveguide filter

$E_z$  at 11.5 GHz, 12 GHz and 12.4 GHz.



- CLAC: asynchronous hybrid DG solver based on OpenCL and MPI.
- It works...
- Work in progress: Gauss-Lobatto integration, memory transfer optimization (zone transpositions), etc. Summer CEMRACS project  
<http://smai.emath.fr/cemracs/cemracs14/lessiv.pdf>
- Task graph: we are reaching our position of incompetence. StarPU, SOCL ?

-  Thomas P Armstrong.  
Numerical studies of the nonlinear vlasov equation.  
*Physics of Fluids (1958-1988)*, 10(6):1269–1280, 1967.
-  Dominique Aubert and Romain Teyssier.  
Reionization simulations powered by graphics processing units.  
i. on the structure of the ultraviolet radiation field.  
*The Astrophysical Journal*, 724(1):244, 2010.
-  Tristan Cabel, Joseph Charles, and Stephane Lanteri.  
Multi-gpu acceleration of a dgtd method for modeling human  
exposure to electromagnetic waves.  
2011.
-  Gary Cohen, Xavier Ferrieres, and Sébastien Pernet.  
A spatial high-order hexahedral discontinuous galerkin method  
to solve maxwell's equations in time domain.  
*Journal of Computational Physics*, 217(2):340–363, 2006.
-  Bruno Després and Rémy Sart.  
Reduced resistive mhd in tokamaks with general density.

*ESAIM: Mathematical Modelling and Numerical Analysis*,  
46(05):1081–1106, 2012.



Philippe Helluy and Jonathan Jung.

Two-fluid compressible simulations on gpu cluster.  
2014.



S Le Bourdiec, F De Vuyst, and L Jacquet.

Numerical solution of the vlasov–poisson system using  
generalized hermite functions.  
*Computer physics communications*, 175(8):528–544, 2006.



Michel Massaro, Philippe Helluy, and Vincent Loechner.

Numerical simulation for the mhd system in 2d using opencl.  
2013.



David Michéa and Dimitri Komatitsch.

Accelerating a three-dimensional finite-difference wave  
propagation code using gpu graphics cards.  
*Geophysical Journal International*, 182(1):389–402, 2010.



Greg Ruetsch and Paulius Micikevicius.

Optimizing matrix transpose in cuda.

*Nvidia CUDA SDK Application Note, 2009.*



Joseph W Schumer and James Paul Holloway.

Vlasov simulations using velocity-scaled hermite representations.

*Journal of Computational Physics, 144(2):626–661, 1998.*



Tao Tang.

The hermite spectral method for gaussian-type functions.

*SIAM Journal on Scientific Computing, 14(3):594–606, 1993.*



Tao Tang, S McKee, and MW Reeks.

A spectral method for the numerical solutions of a kinetic equation describing the dispersion of small particles in a turbulent flow.

*Journal of computational physics, 103(2):222–230, 1992.*