Environnement et Changement climatique Canada
Environment and Climate Change Canada

Canada

Towards a new dynamical kernel in GEM

Vivian Lee and Stéphane Gaudreault

18th Workshop on High Performance Computing in Meteorology

ECMWF, Reading, UK

September 24 - 28, 2018

The Evolution of High Performance Computing (HPC)

- Dynamics Research group : develops the dynamical core of the GEM Model and undertakes research on numerical techniques.
- HPC is evolving rapidly : the best numerical algorithms on today's supercomputer could be suboptimal in the future
  - Computation speed limited by the data movement rate
- Are we ready for such architectures ?
  - GEM will be improved
    - Improved semi-Lagrangian advection
    - Optimized schwarz method **(Qaddouri & Gaudreault, in preparation)**
    - Multigrid preconditioner
  - Investigation of alternative numerical methods

The drive for a new kernel...

- This project is driven by **Stéphane Gaudreault** who has developed a parallel implementation (MPI/OpenMP) of the exponential propagation iterative (EPI) method in GEM with the Yin-Yang overset grid.

- All slides are borrowed from **Stéphane Gaudreault's** presentations

- Collaborators :

  ECCC : Michel Desgagné, André Plante, Rabah Aider, Abdessamad Qaddouri, Monique Tanguay, Martin Charron

  University of California, Merced : Mayya Tokman, Valentin Dallerit, Tomasso Buvoli, Greg Rainwater

Criteria for a good numerical scheme

- Stable
  - even for a large timestep size
- Precise
  - small errors, avoid phase error and spurious dispersion effects

- High ratio floating-point operations/data movement
- Parallel scalability
- Fast

ODEs, Implicit, explicit ...

$$\frac{dq}{dt} = F(q(t)),\, q(0) = q_0$$

- **Explicit** : calculate the state of a system at a later time from the state of the system at the current time

$$q_{n+1} = q_n + F(q_n) \cdot \Delta t$$

- **Implicit** : find a solution by solving an equation involving a later state of the system.

$$q_{n+1} = q_n + F(q_{n+1}) \cdot \Delta t$$

ODEs, Implicit, explicit, stability ...

- The error made at one time step do not cause the errors to be magnified as the computation are continued ...

It is well known that the **timestep size of explicit methods needs to be very small** to ensure stability.

Is it true for all explicit schemes ?

NO!

Unconditionally stable explicit eulerian scheme do exist!

Exponential time integrators

- Innovative approach previously studied in the icosahedral model of Janusz Pudykiewicz

  2013 : Colm Clancy and Janusz A. Pudykiewicz

  On the use of exponential time integration methods in atmospheric models
- Method of lines
  - Discretize in space to obtain a system of ODEs, then in time
- The method is general. It could be used with
  - any Eulerian discretization scheme
  - any global grid
- Recent progress in the computational linear algebra has led to more efficient algorithms like the PHIPM method of Niesen and Wright
- Considering the results obtained with exponential integration methods in various areas of science, it is justified to investigate them in the context of numerical weather prediction where the selection of a time integration scheme is a key element

  2016 : Stéphane Gaudreault and Janusz A. Pudykiewicz

  An efficient exponential time integration method for the numerical solution of the shallow water equations on the sphere

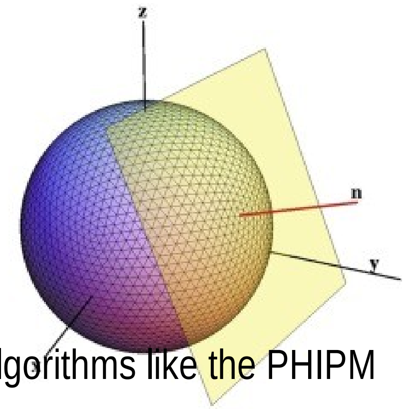

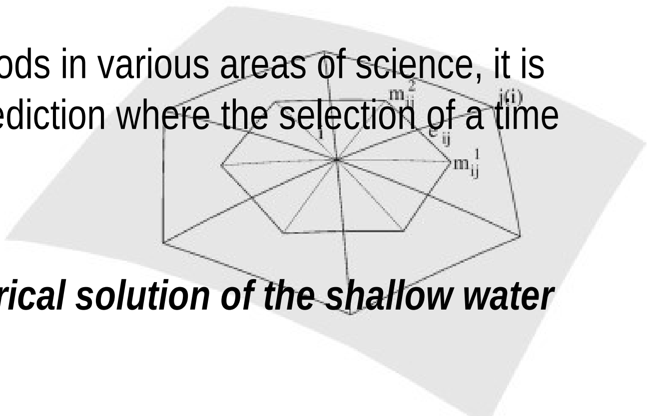

Exponential Propagation Iterative (EPI)

- EPI3 (3rd order scheme) is :

$$q_{n+1} = q_n + \varphi_1(\mathcal{J}_n \Delta t) F_n \, \Delta t + \frac{2}{3} \varphi_2(\mathcal{J}_n \Delta t) R_{n-1} \, \Delta t$$

where

$$R_{n-1} = F(q_{n-1}) - F(q_n) - \mathcal{J}_n \cdot (q_{n-1} - q_n)$$

is evaluated from previous time levels. **(Tokman, 2006)**

$$\mathcal{J}_n = \frac{\partial}{\partial q} F(q(t)) \Big|_{t=t_n}$$

say $A = \mathcal{J}_n \Delta t$ and $\varphi_0(A) = e^A$

- **Challenge : compute the φ-functions** :

$$\varphi_p(A) = \sum_{n=0}^{\infty} \frac{1}{(n+p)!} A^n$$

- Promising approach to overcome accuracy drawbacks related to the large time-step choice while still correctly simulating all relevant wave dispersion relations.

The PHIPM Solver (Niesen and Wright, 2012)

- The good
  - Attractive option when little or no information about the spectrum or norm of the Jacobian matrix is known *a priori*. **(Tokman et al. 2012) (Clancy and Pudykiewicz, 2013)**
- The bad
  - The convergence of the phipm algorithm is often inconsistent
  - Poor parallel scaling : Arnoldi procedure requires $O(m^2)$ calls to MPI_AllReduce.
  - Important details related to a specific computer architecture are not taken into account in the adaptive procedure
- The ugly
  - Can be sensitive to rounding errors as p increases.

The KIOPS Solver

(Gaudreault, Rainwater and Tokman, 2018)

- **KIOPS :Krylov with Incomplete Orthogonalization Procedure Solver**

- **The good**
  - Attractive option when little or no information about the spectrum or norm of the Jacobian matrix is known *a priori*.
- **Even better**
  - Efficient calculation (O(m) instead of O(m^2)), good parallel scaling, and consistent convergence

- **Publication**

KIOPS: A fast adaptive Krylov subspace solver for exponential integrators. Journal of Computational Physics.

Gaudreault, S., Rainwater, G. and Tokman, M., 2018.

EPIC (Exponential Propagation Integrators Collection)

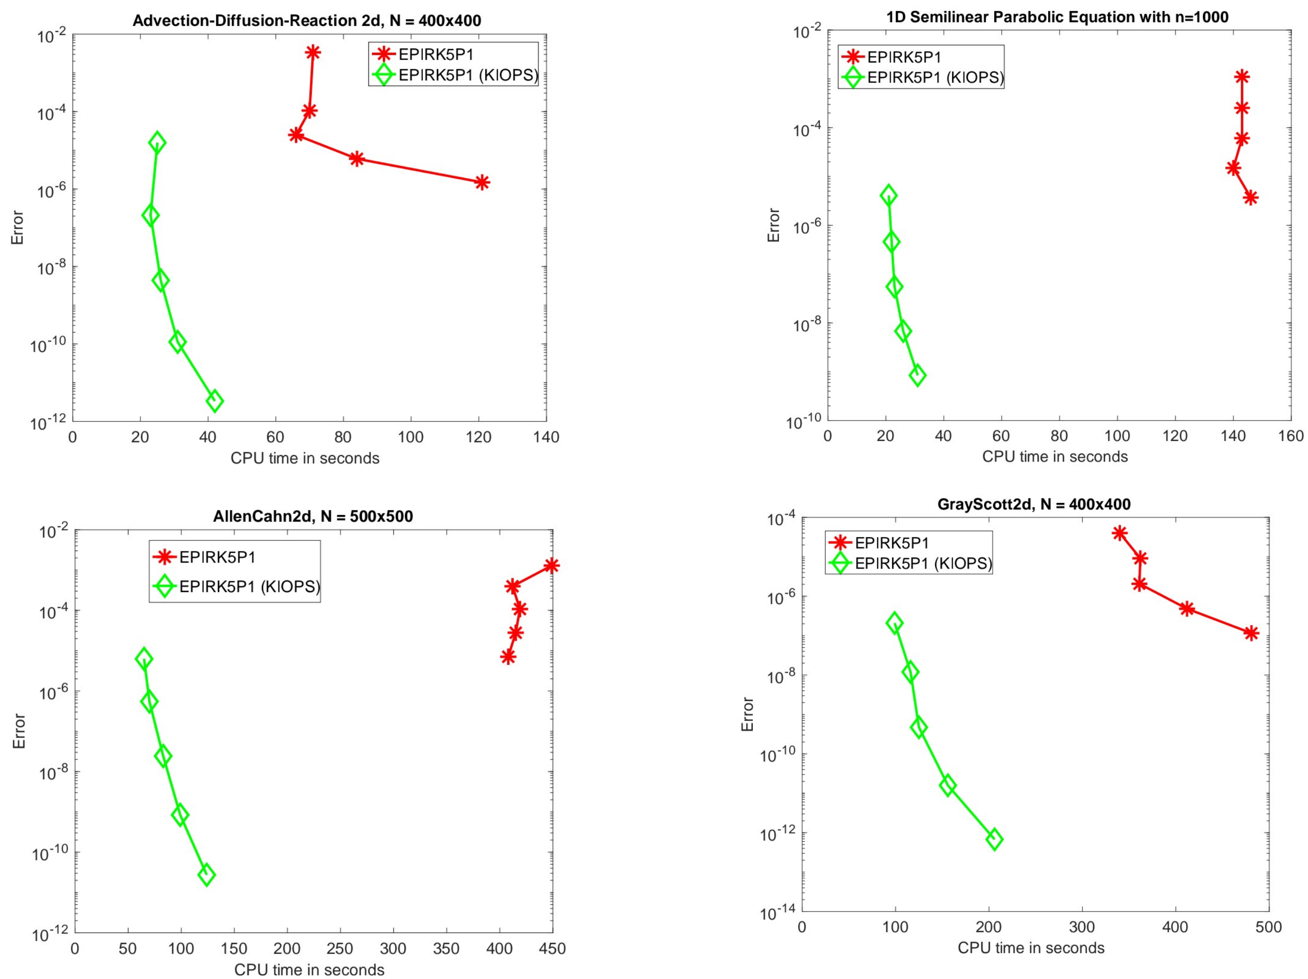

http://faculty.ucmerced.edu/mtokman/#software

Shallow Water Model on the Yin-Yang Grid

$$\frac{dh}{dt} = -\frac{\overline{u}}{a\cos\phi}\,\delta^a_\lambda h^* - \frac{\overline{v}}{a}\,\delta^a_\phi h^* - \frac{h^*}{a\cos\phi}\,(\delta_\lambda u + \delta_\phi v\cos\phi) \qquad (1)$$

$$\frac{du}{dt} = -\frac{u}{a\cos\phi}\,\delta^a_\lambda u - \frac{\overline{v}}{a}\,\delta^a_\phi u + \left(f + u\frac{\tan\phi}{a}\right)\overline{v} - \frac{g}{a\cos\phi}\,\delta_\lambda h \qquad (2)$$

$$\frac{dv}{dt} = -\frac{\overline{u}}{a\cos\phi}\,\delta^a_\lambda v - \frac{v}{a}\,\delta^a_\phi v - \left(f + \overline{u}\frac{\tan\phi}{a}\right)\overline{u} - \frac{g}{a}\,\delta_\phi h \qquad (3)$$

- Arakawa C-grid
- Low numerical diffusion upwind scheme

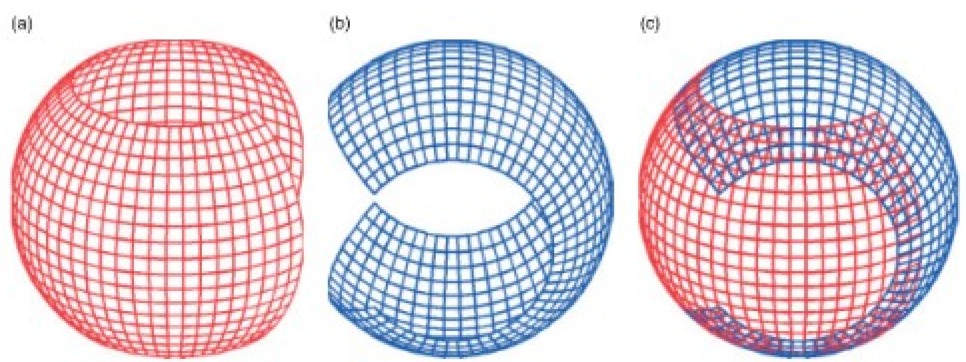

Image : A. Qaddouri

Model configurations (v4.8-LTS.4)

- FISL = Fully Implicit Semi-Lagrangian
- EXPO = EPI3
- Test cases from Williamson et al, (1992)

	FISL	EXPO
Equations	Euler autobarotropic	Shallow water (In french : equations de Barré de Saint-Venant)
Resolution	1 degree	1 degree
timestep	2160s (case1) 450s (other)	2160 (case1) 450s (other)
Grd_maxcfl	3	
Grd_overlap	2 degree	0
Schm_itcn, Schm_itnlh	2, 2	
Tolerance		1e-7
Angle alpha	0	0

Advection of a cosine bell

- Video
- 12 days simulation

Advection of a cosine bell : absolute error

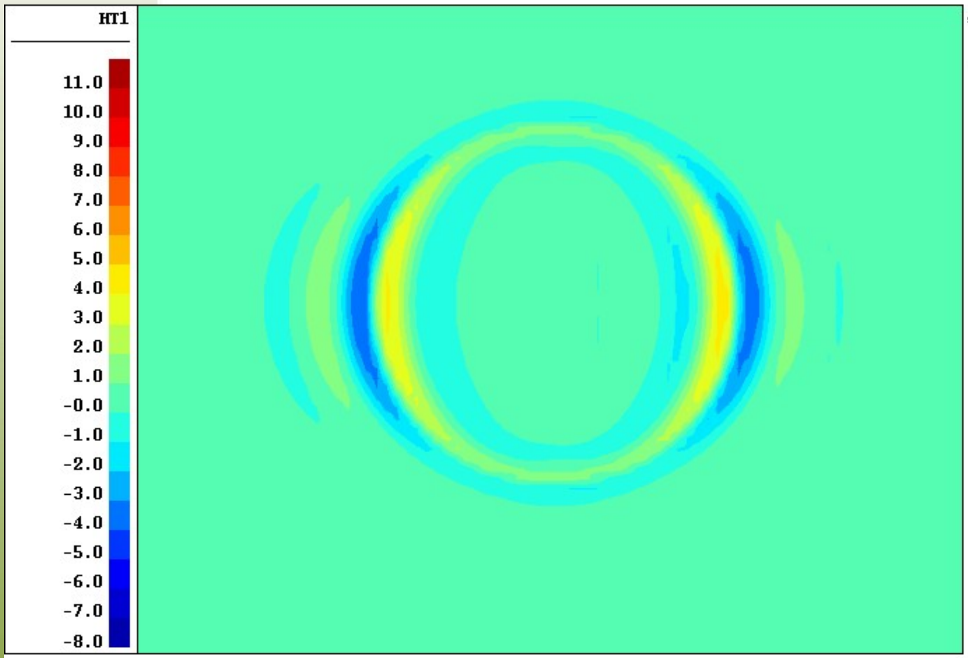

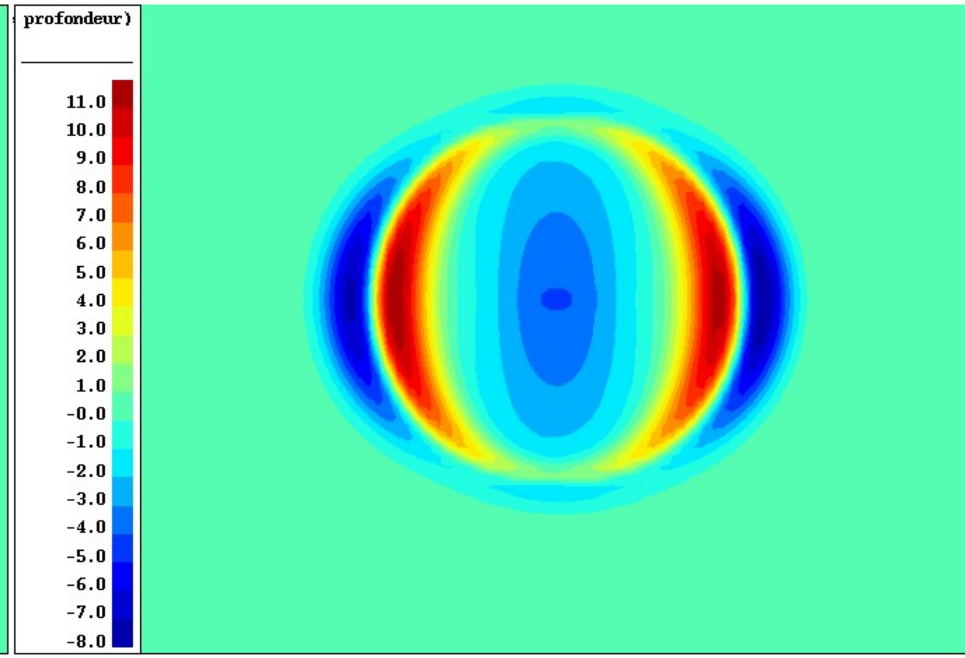

EXPO	Max of	FISL
[-4.24, 4.28]	[negative, positive] Error values	[-8.3, 12]

Steady-state nonlinear zonal geostrophic flow

After 12 days

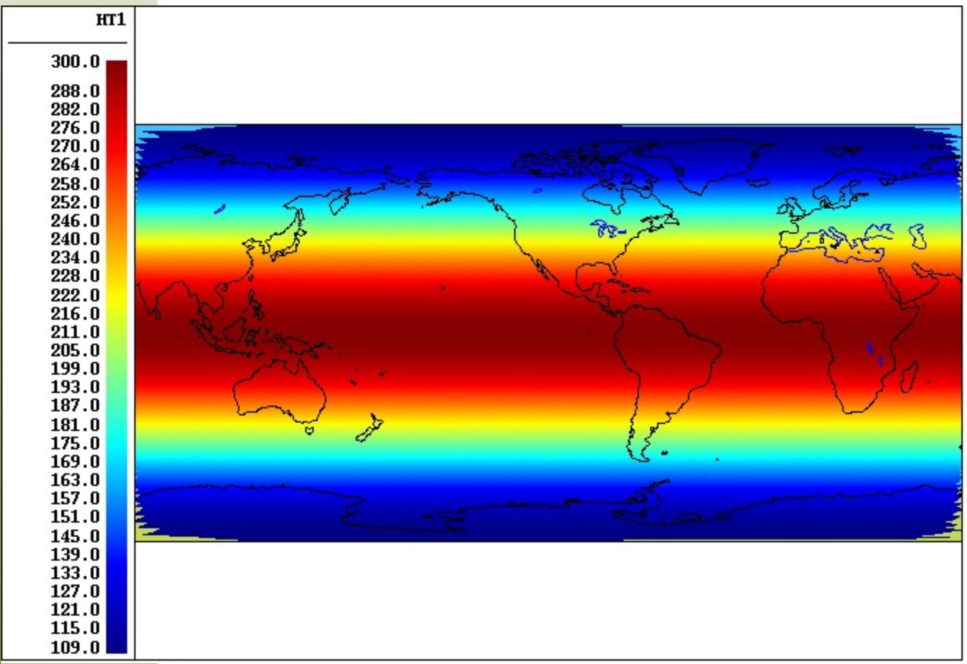

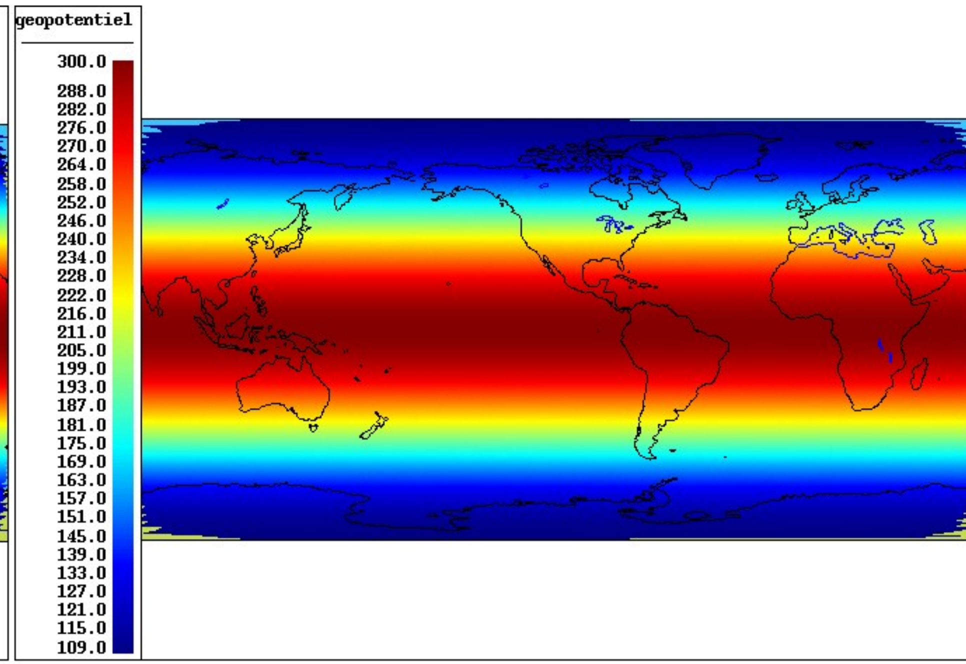

EXPO

FISL

Steady-state nonlinear zonal geostrophic flow : absolute error

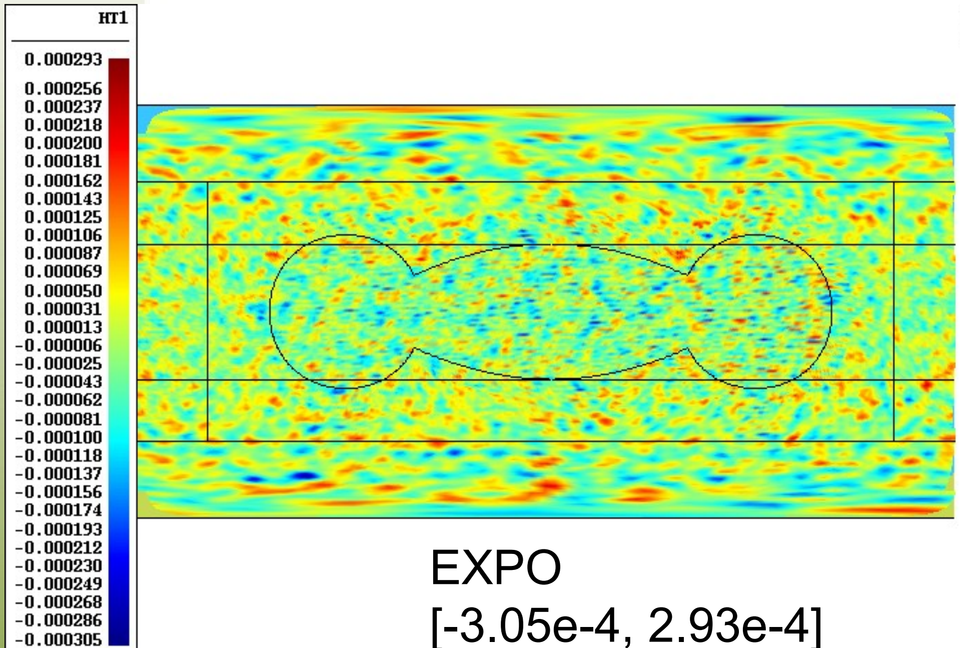

EXPO
[-3.05e-4, 2.93e-4]

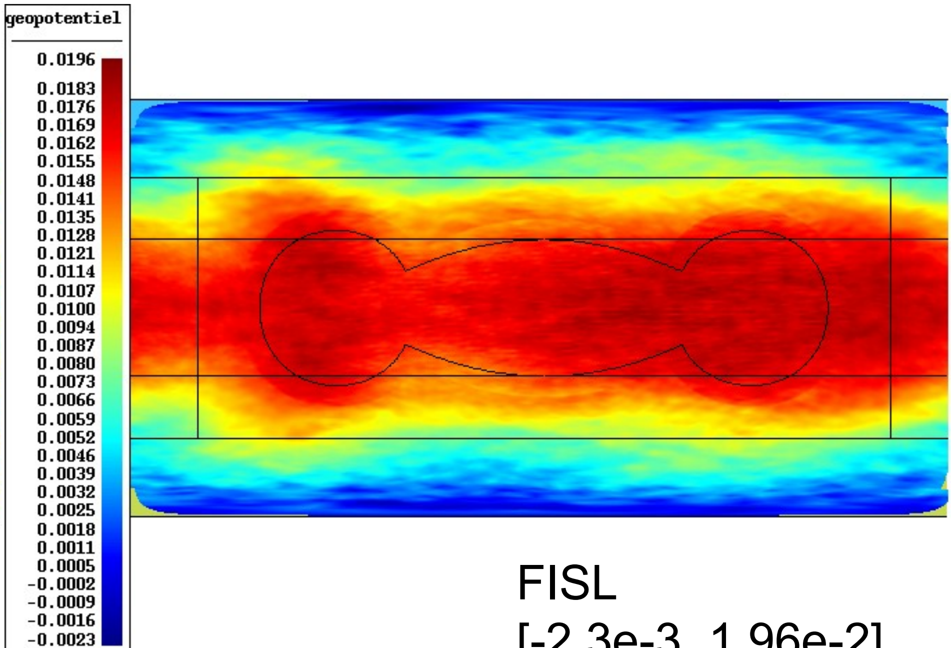

FISL
[-2.3e-3, 1.96e-2]

Max of
[negative, positive]
Error values

Steady-state nonlinear zonal geostrophic flow : mass conservation error

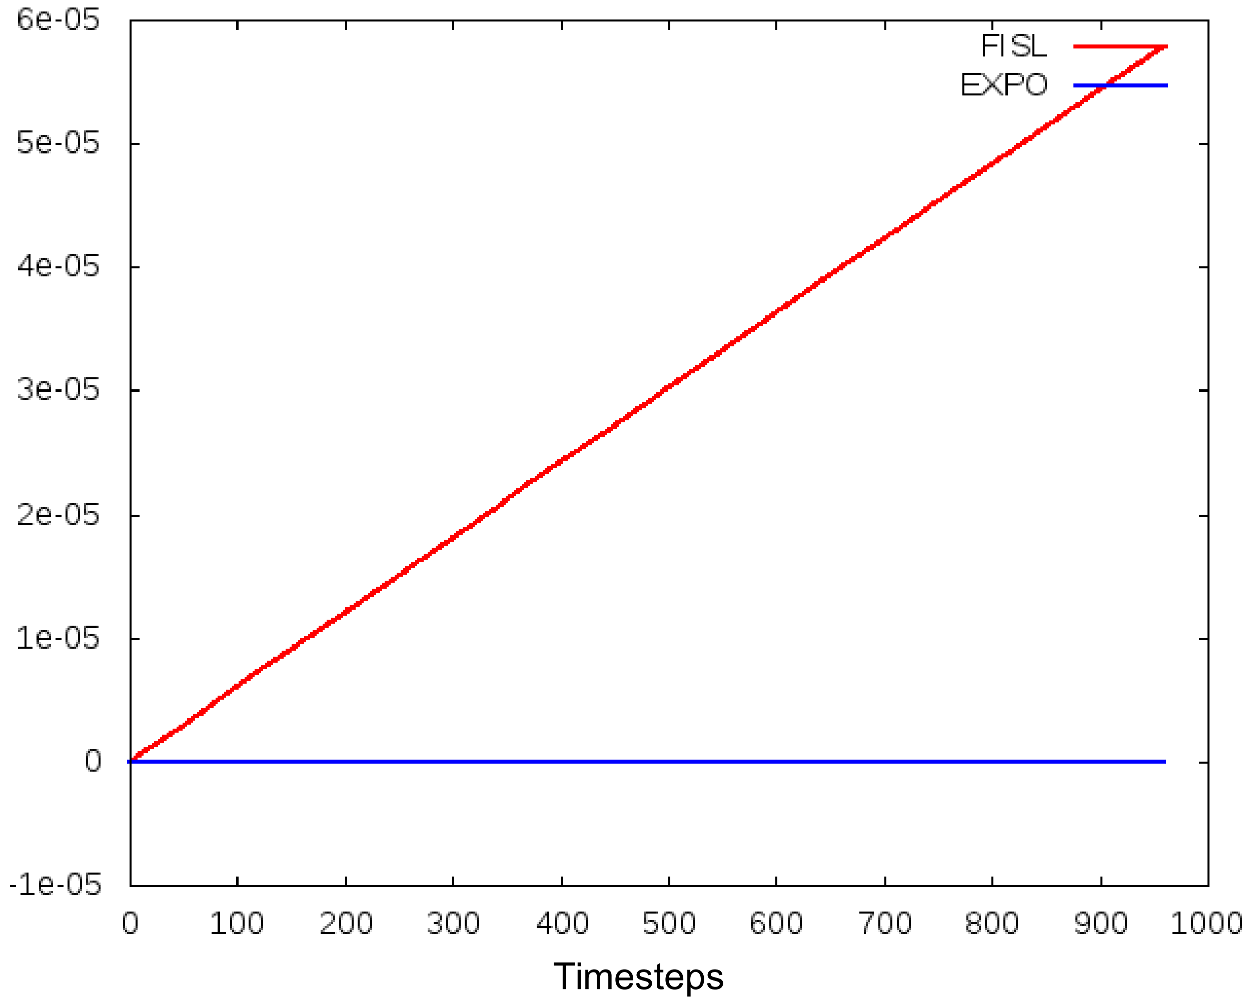

Zonal flow over an isolated mountain

- Video
- 15 days simulation

Zonal flow over an isolated mountain

After 15 days

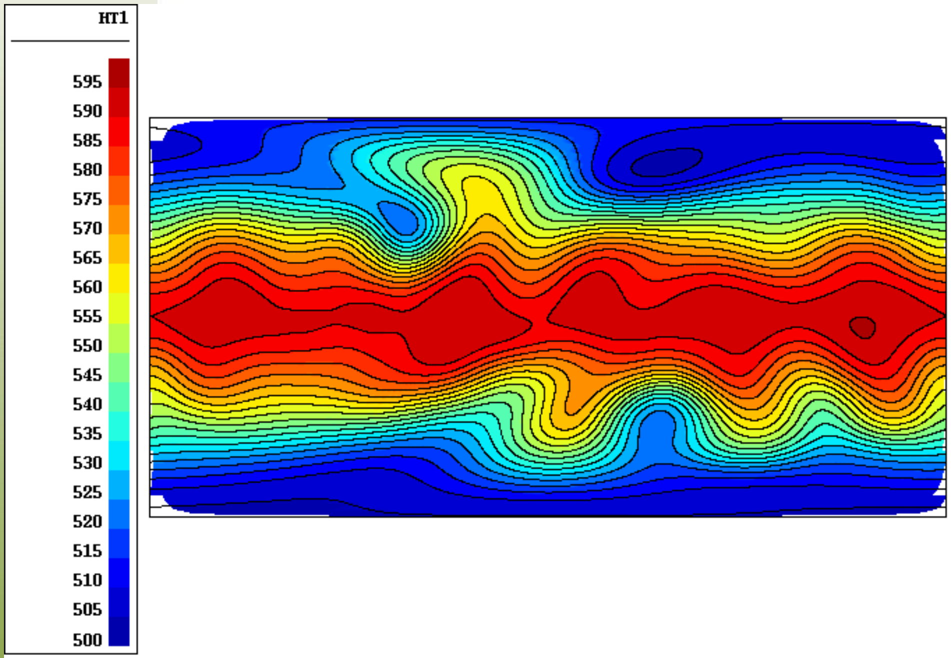

EXPO

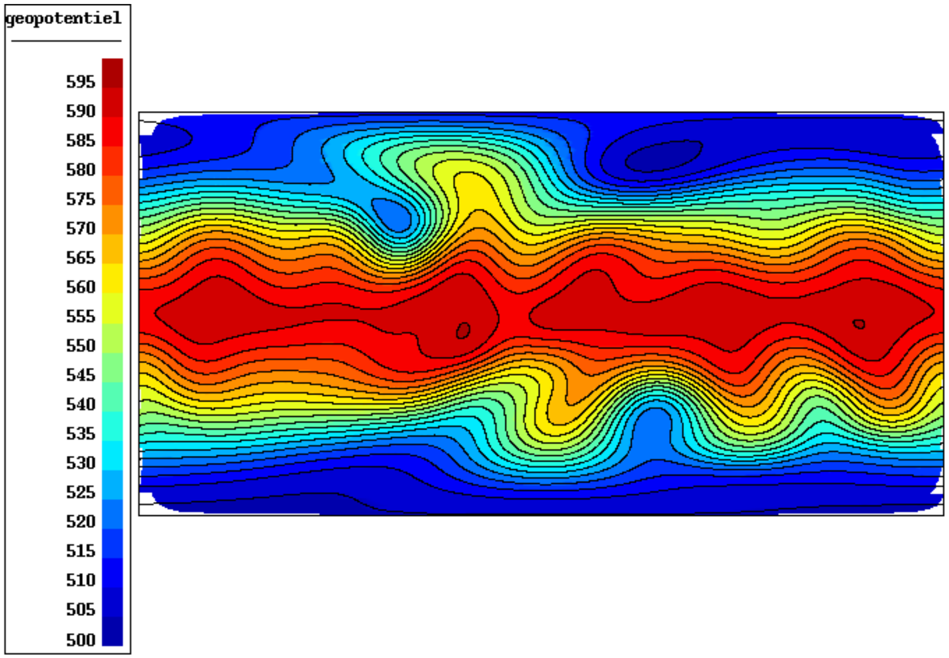

FISL

Zonal flow over an isolated mountain
Comparison with a high-resolution model

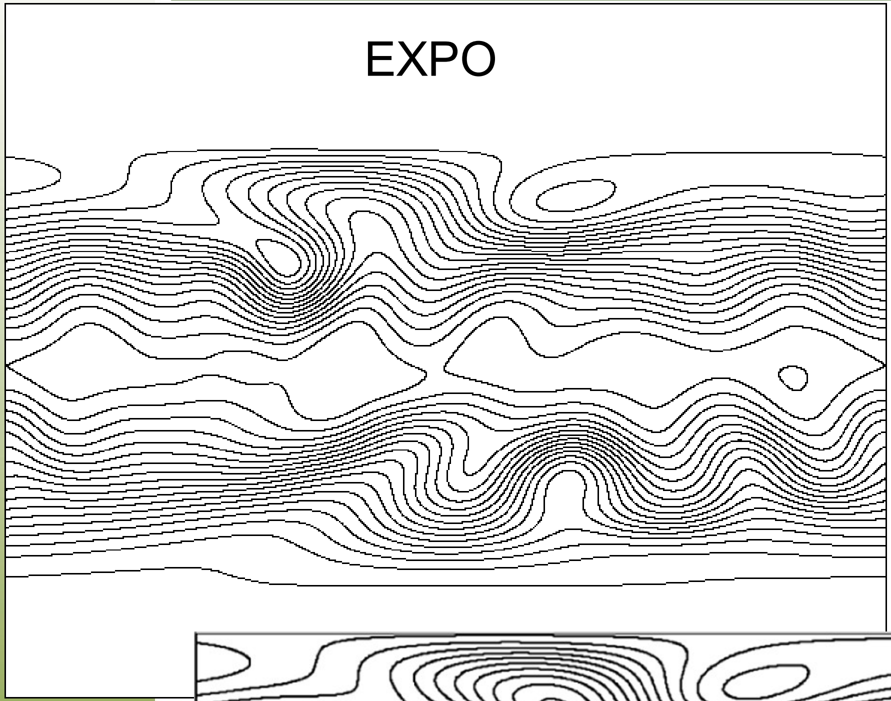

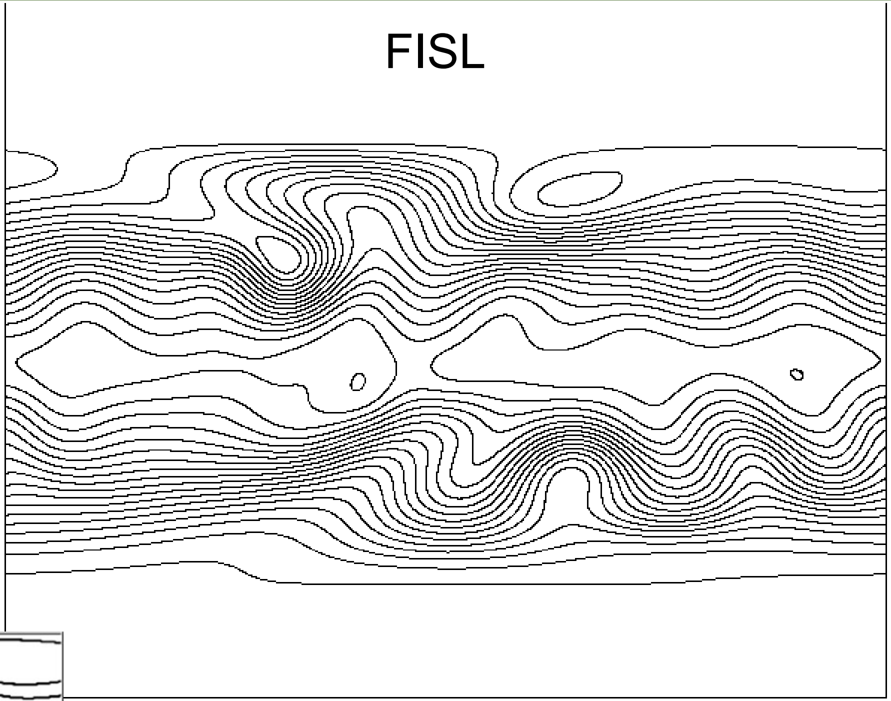

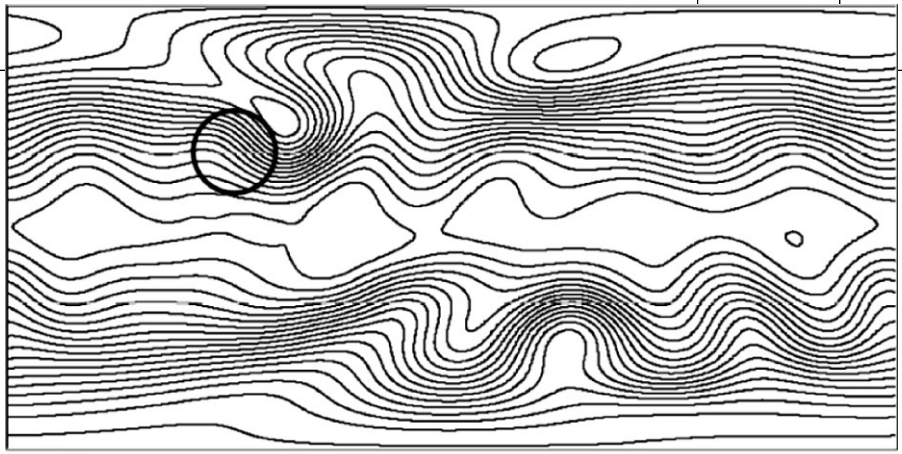

Qaddouri et al. 2012
spectral model
1080 × 540 lin. Gaussian grid (T539)
Timestep : 240 s

Zonal flow over an isolated mountain
Mass conservation error

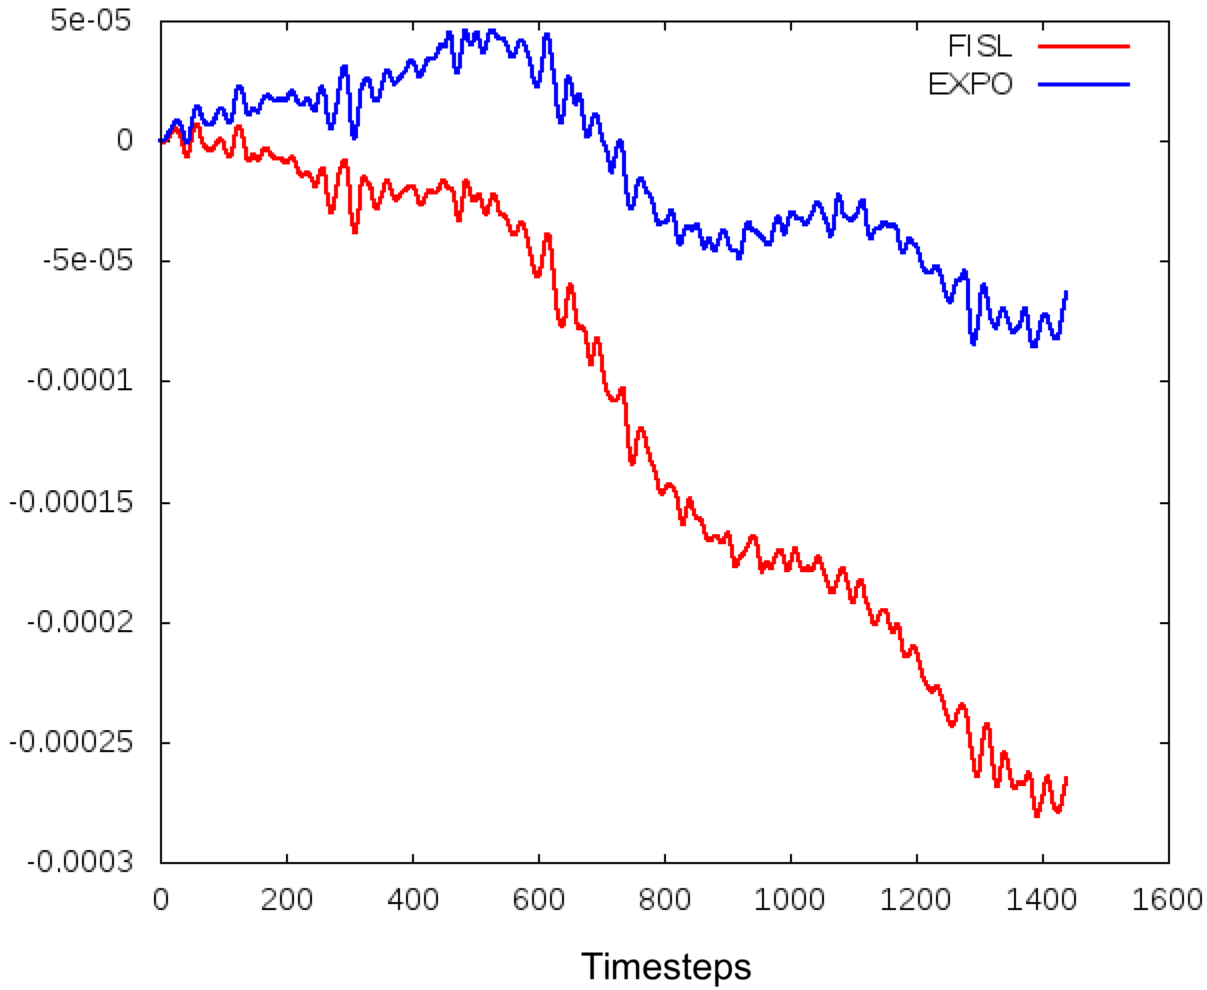

Timestep length : 900 s

dt=1 hour (LONG TIME STEP)

EXPO

FISL
Grd_maxcfl=4

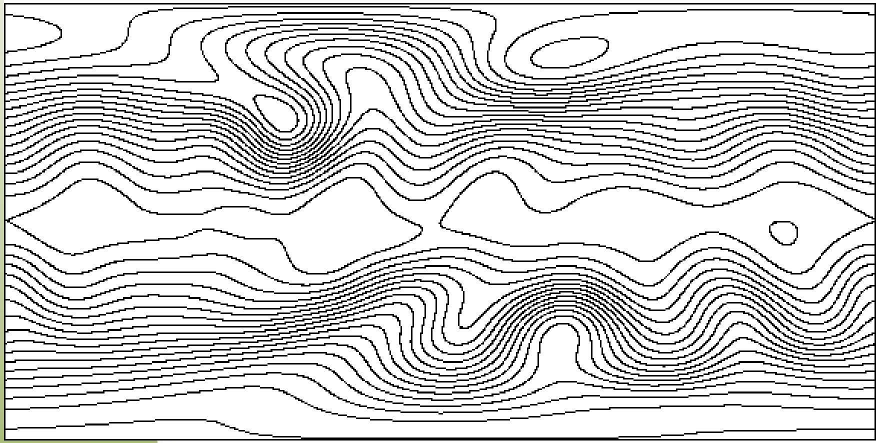

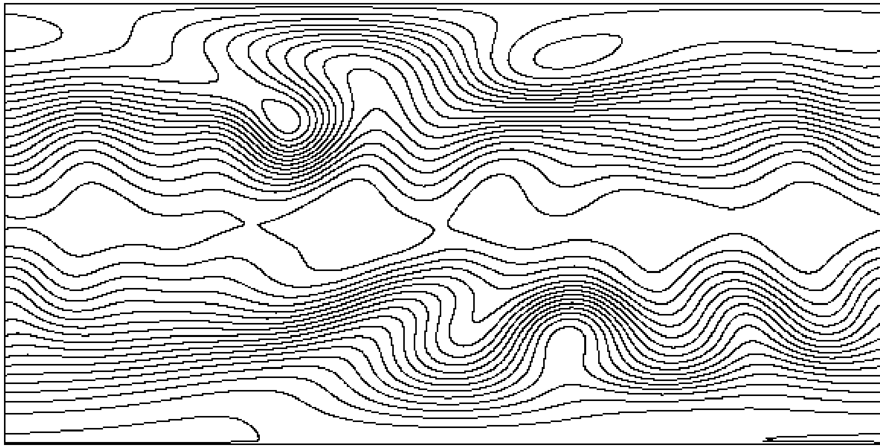

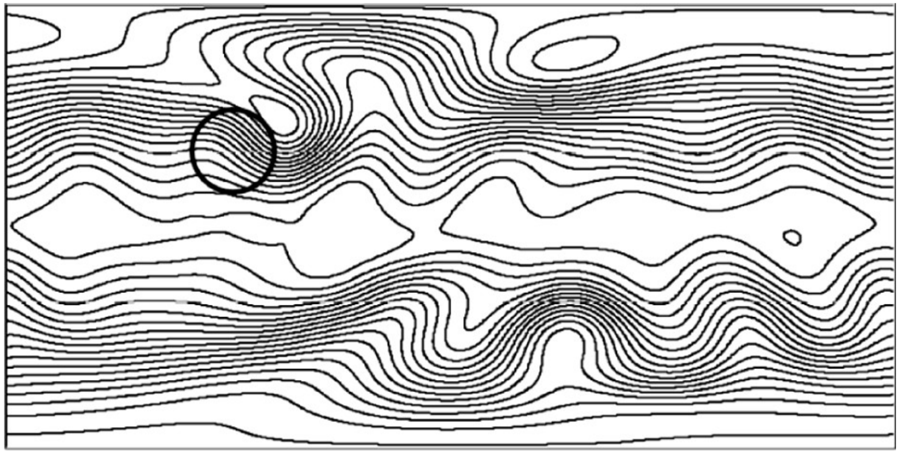

REF

The Non-hydrostatic Euler Equations

- 5 dynamical variables : $\mathrm{u}^i,\ \hat{\Pi},\ \theta_v$

- Momentum $$\frac{\partial u^i}{\partial t} + u^j u^i_{,j} + 2\Gamma^i_{j0} u^j + \Gamma^i_{jk} u^j u^k = -h^{ij}\left(\theta_v\ \hat{\Pi}_{,j} + \psi_{,j}\right)$$

- Continuity $$\frac{\partial \hat{\Pi}}{\partial t} + u^j \hat{\Pi}_{,j} = -\frac{R_d \hat{\Pi}}{c_{vd}\sqrt{g}}(\sqrt{g}u^j)_{,j}$$

- Thermodynamic $$\frac{\partial \theta_v}{\partial t} + u^j (\theta_v)_{,j} = 0$$

(Charron et al. 2014)

3D Dry Bubble Convection Problem of André Robert (1993)

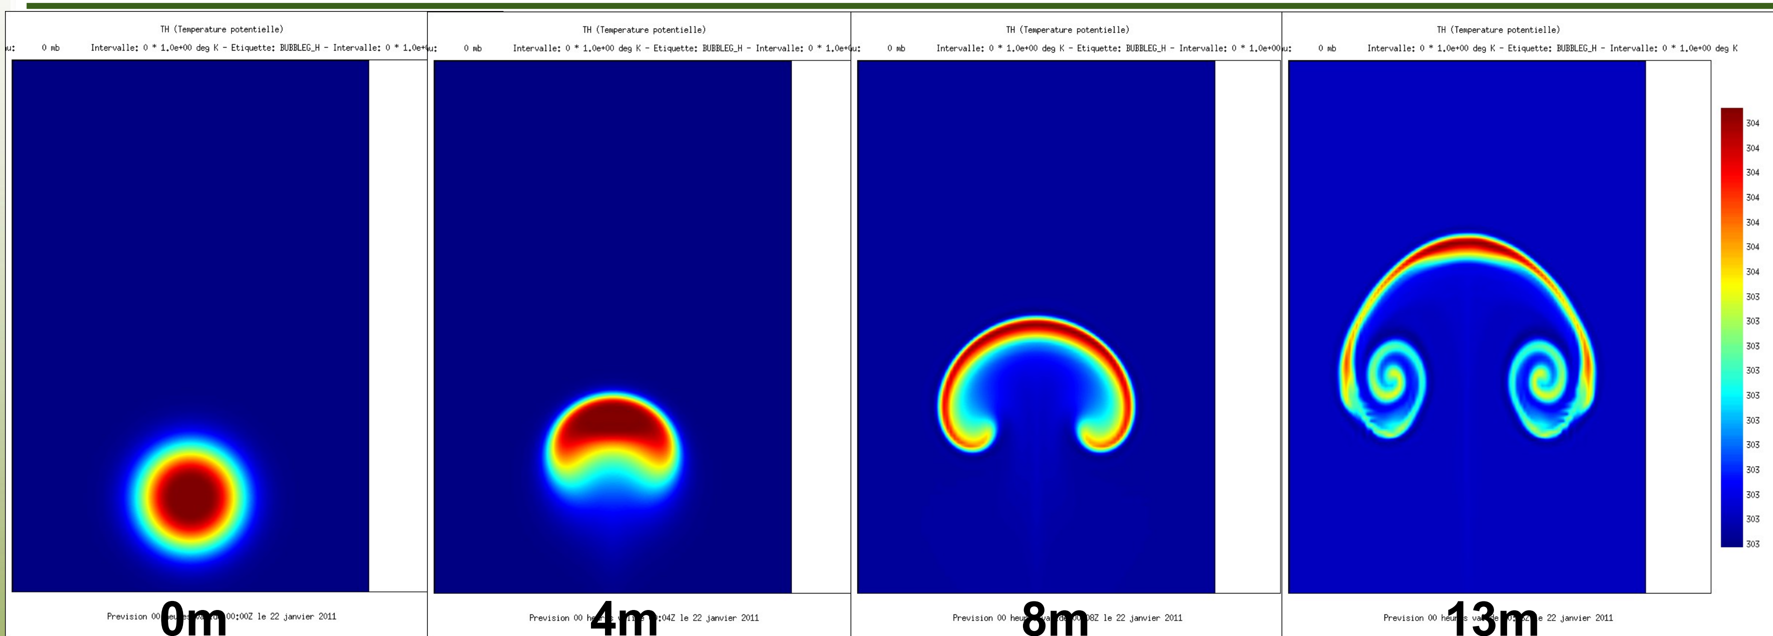

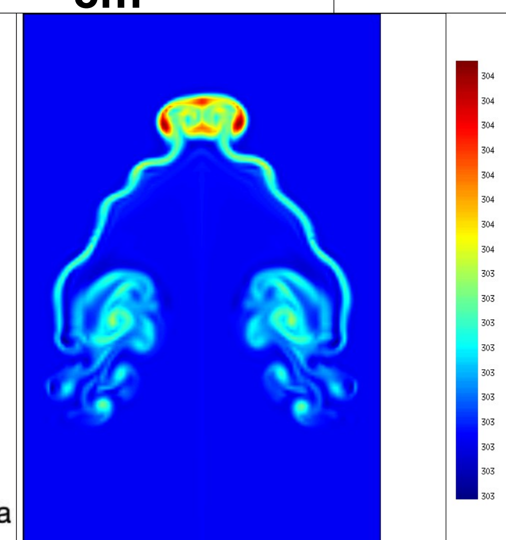

ni = 101, nk = 150,
dx = 10, dz = 10,
dt = 5s, duration = 18m

Conclusions from Experiments

- √ Stable
- √ Precise, mass conservative
- √ High arithmetic intensity
- √ Fast (work in progress)
- √ Parallel scalability (work in progress)
  - A pipelined version has been implemented
- √ Performance degrades for very stiff problems $\frac{\partial y}{\partial t} = F_v(t, y) + F_h(t, y)$
  - Large difference in grid-spacings in the horizontal and vertical directions.
  - Vertical / horizontal separation is being investigated.
  - How to deal with stratospheric polar jets whose speed exceeds 100 m/s?

Towards a new GEM kernel ...

- Compiling current code in different compilers to make it more robust and portable (such as gfortran, pgi)
- Code optimization for performance and parallel scalability
- More work to be done in GEM using the KIOPS solver for performance in stiff problems
- Accelerators (e.g. GPUs)

Thank-you

And many thanks to Stéphane Gaudreault and Monique Tanguay for their support to this presentation.

- Questions?