

#### **GPU Consideration for Next Generation Weather (and Climate) Simulations**

### Oliver Fuhrer<sup>1</sup>, Tobias Gisy<sup>2</sup>, Xavier Lapillonne<sup>3</sup>, Will Sawyer<sup>4</sup>, Ugo Varetto<sup>4</sup>, Mauro Bianco<sup>4</sup>, David Müller<sup>2</sup>, and Thomas C. Schulthess<sup>4,5,6</sup>

<sup>1</sup>Swiss Federal Office of Meteorology and Climatology; <sup>2</sup>Supercomputing Systems; <sup>3</sup>Center for Climate Systems Modeling, ETH Zurich; <sup>4</sup>Swiss National Supercomputing Center <sup>5</sup>Institute for Theoretical Physics, ETH Zurich; <sup>6</sup>Computer Science and Mathematics Division, ORNL





### PLEASE NOTE THAT ...

1. I am a computational physicist and have no research background in atmospheric science – please do not expect to receive deep insights in climate science or meteorology during my presentation

2. I report on work in progress of a project that runs through 12/2012 – many results I present are preliminary



# Today's state of the art climate simulation (resolution T85 ~ 148 km)

Surface temperature change relative to 1870-1899 baseline CCSM3 IPCC AR4



GTC Asia, Beijing





# Experimental climate running at higher resolution (resolution T341 ~ 37 km)







#### Why resolution is such an issue for Switzerland







Source: Oliver Fuhrer, MeteoSwiss

Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich

#### **Prognostic uncertainty**

The weather system is chaotic

 $\rightarrow$  rapid growth of small perturbations (butterfly effect)





Ensemble method: compute distribution over many simulations

Source: Oliver Fuhrer, MeteoSwiss

GTC Asia, Beijing





#### WHAT NEEDS TO BE IMPROVED?

### Higher resolution Larger ensemble Can we have both?



# Increasing resolution (weak scaling): larger parallel computers only solve part of the problem





Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich

#### Applications running at scale on Jaguar @ ORNL (Spring 2011)

CSCS

Swiss National Supercomputing Centre

Domain area	Code name	Institution	# of cores	Performance	Notes	
Materials	DCA++	ORNL	213,120	1.9 PF	2008 Gordon Bell Prize Winner	
Materials	WL-LSMS	ORNL/ETH	223,232	1.8 PF	2009 Gordon Bell Prize Winner	
Chemistry	NWChem	PNNL/ORNL	224,196	1.4 PF	2008 Gordon Bell Prize Finalist	
Materials	DRC	ETH/UTK	186,624	1.3 PF	2010 Gordon Bell Prize Hon. Mention	
Nanoscience	OMEN	Duke	222,720	> 1 PF	2010 Gordon Bell Prize Finalist	
Biomedical	МоВо	GaTech	196,608	780 TF	2010 Gordon Bell Prize Winner	
Chemistry	MADNESS	UT/ORNL	140,000	550 TF		
Materials	LS3DF	LBL	147,456	442 TF	2008 Gordon Bell Prize Winner	
Seismology	SPECFEM3D	USA (multiple)	149.784	165 TF	2008 Gordon Bell Prize Finalist	
Combustion	S3D	SNL	147,456	83 TF		
Weather	WRF	USA (multiple)	150,000	50 TF		



### 

 $\mathbf{G}_{c}(\{s_{i},l\}_{k+1}) = \mathbf{G}_{c}(\{s_{i},l\}_{0}) + [\mathbf{a}_{0}|\mathbf{a}_{1}|...|\mathbf{a}_{k}] \times [\mathbf{b}_{0}|\mathbf{b}_{1}|...|\mathbf{b}_{k}]^{t}$ 

Complexity for *k* updates remains  $O(kN_t^2)$ 

But we can replace *k* rank-1 updates with one matrix-matrix multiply plus some additional bookkeeping.





#### Arithmetic intensity of a computation

$$\alpha = \frac{\text{floating point operations}}{\text{data transferred}}$$

$$\mathbf{G}_{c}(\{s_{i}, l\}_{k+1}) = \mathbf{G}_{c}(\{s_{i}, l\}_{k}) + \mathbf{a}_{k} \times \mathbf{b}_{k}^{t}$$
$$\alpha \approx o(1)$$

$$\begin{aligned} \mathbf{G}_c(\{s_i, l\}_{k+1}) &= \mathbf{G}_c(\{s_i, l\}_0) + [\mathbf{a}_0 | \mathbf{a}_1 | \dots | \mathbf{a}_k] \times [\mathbf{b}_0 | \mathbf{b}_1 | \dots | \mathbf{b}_k]^t \\ \alpha &\approx o(k) \end{aligned}$$



# Finite difference stencil to solve PDE on structured grid – climate, meteorology, seismic imaging, combustion



Partial differential equations on structured grid

$$\frac{\partial u}{\partial t} - \alpha \nabla^2 u = 0$$





#### Algorithmic motifs and their arithmetic intensity

Arithmetic intensity: number of operations per word of memory transferred

![](_page_12_Figure_4.jpeg)

![](_page_13_Picture_0.jpeg)

#### Numerical weather predictions for Switzerland (MeteoSwiss)

![](_page_13_Picture_3.jpeg)

![](_page_14_Picture_0.jpeg)

![](_page_14_Picture_1.jpeg)

Assimilation

Dynamics

#### Performance profile of (original) COSMO-CCLM

Runtime based 2 km production model of MeteoSwiss

![](_page_14_Figure_4.jpeg)

Source: Oliver Fuhrer, MeteoSwiss

![](_page_15_Picture_0.jpeg)

![](_page_15_Picture_1.jpeg)

#### Example of motif in the physics code

do j = 1, niter
 do i = 1, nwork
 c(i) = a(i) \*\* b(i) + sin(b(i)) \* log(a(i))
 end do
end do

$$c = a^b + sin(b) \cdot log(a)$$

(niter=48; nwork=4096000)

Processor	Barcelona	Istambul	Magny-Cours	Interlagos
# cores	4 @ 2.6GHz	6 @ 2.4 GHz	12 @ 2.2 GHz	16 @ 2.1 GHz
Single core	17.8 s	17.3 s	18.8 s	19.9 s
All cores	4.5 s	2.9 s	1.61 s	1.31 s
Speedup	3.96	5.97	11.7	15.19
	(1.38)	(1.83)	(1.27)	

![](_page_16_Picture_0.jpeg)

![](_page_16_Picture_1.jpeg)

#### **Example of motif in the dynamics code**

 $\frac{\partial a}{\partial t} = k \frac{\partial^2 a}{\partial x^2}$ 

Processor	Barcelona	Istambul	Magny-Cours	Interlagos
# cores	4 @ 2.6GHz	6 @ 2.4 GHz	12 @ 2.2 GHz	16 @ 2.1 GHz
Single core	0.80 s	0.84 s	0.63 s	0.65 s
All cores	0.56 s	0.46 s	0.18 s	0.16 s
Speedup	1.4	1.8 2.56	3.5 0.13	4.1
	(1.38)	(1.83)	(1.27)	

![](_page_17_Picture_0.jpeg)

![](_page_17_Picture_1.jpeg)

#### Analyzing the two examples – how are they different?

![](_page_17_Figure_3.jpeg)

- Arithmetic throughput is a per core resource that scale with number of cores and frequency
- Memory bandwidth is a shared resource between cores on a socket

GTC Asia, Beijing

![](_page_18_Picture_2.jpeg)

#### Looking at the entire COSMO code

- Running COSMO-2, keeping compute power constant, but varying the available memory bandwidth
- Keep number of cores constant and vary number of cores used per socket

![](_page_18_Figure_7.jpeg)

1 cores/CPU

4 cores/CPU

2 cores/CPU

![](_page_18_Picture_8.jpeg)

100.00%

![](_page_19_Figure_2.jpeg)

![](_page_20_Figure_2.jpeg)

![](_page_21_Picture_1.jpeg)

#### Running the simple examples on the Cray XK6

#### Compute bound (physics) problem

Machine	Interlagos	Fermi (2090)	GPU+transfer
Time	1.31 s	0.17 s	1.9 s
Speedup	1.0 (REF)	7.6	0.7

#### Memory bound (dynamics) problem

Machine	Interlagos	Fermi (2090)	GPU+transfer
Time	0.16 s	0.038 s	1.7 s
Speedup	1.0 (REF)	4.2	0.1

#### The simple lesson: leave data on the GPU!

![](_page_22_Picture_0.jpeg)

![](_page_22_Picture_1.jpeg)

Assimilation

Dynamics

#### Performance profile of (original) COSMO-CCLM

Runtime based 2 km production model of MeteoSwiss

![](_page_22_Figure_4.jpeg)

![](_page_23_Picture_0.jpeg)

![](_page_23_Picture_1.jpeg)

#### Approach to refactoring of COSMO

- Physics
  - 40k lines, 20% of runtime
  - Many developers
  - "Plug-ins" from other models
  - More compute bound

- Port to GPUkeep sourcedirectives

![](_page_24_Picture_1.jpeg)

#### **Refactoring Physics in COSMO: OpenACC directives**

```
!$acc update device(a,b)
do j=1,niter
 !$acc region do kernel
  do i=1,n
    c(i) = a(i) + b(i) * ( a(i+1) - 2.0d0*a(i) + a(i-1) )
  end do
 !$acc update host(c)
end do
!$acc update host(c)
```

Unfortunately, things are not always that easy and code restructuring is required. But this usually helps improving performance if the code ran on a CPU.

#### Speedup GPU vs. CPU

![](_page_25_Figure_3.jpeg)

![](_page_25_Figure_4.jpeg)

![](_page_26_Picture_0.jpeg)

![](_page_26_Picture_1.jpeg)

#### Approach to refactoring of COSMO

- Physics
  - 40k lines, 20% of runtime
  - Many developers
  - "Plug-ins" from other models
  - More compute bound

#### Dynamics

- 20k lines, 60% of runtime
- Few developers
- Strongly memory bandwidth bound

- Port to GPUkeep sourcedirectives

![](_page_26_Figure_15.jpeg)

![](_page_27_Picture_0.jpeg)

#### **Results for dynamical core**

- Fully implemented / verified agains original code
- Speedup CPU (Magny Cours) vs. GPU (Tesla C2075)

Original CPU	Rewrite CPU	Rewrite GPU
1.0	1.8 2.2 (3)	4.0

- CPU speedup du to reduction in memory access
- Further GPU optimization work is ongoing

![](_page_28_Picture_1.jpeg)

#### **Dynamics in COSMO-CCLM**

$$\begin{cases} \frac{\partial u}{\partial t} = \left[ -\left\{ \frac{1}{a \cos \varphi} \frac{\partial E_h}{\partial \lambda} - vV_a \right\} - \left[ \frac{\partial u}{\partial \zeta} \right] \left[ \frac{1}{\rho a \cos \varphi} \left( \frac{\partial p'}{\partial \lambda} - \frac{1}{\sqrt{\gamma}} \frac{\partial p_0}{\partial \rho} \frac{\partial p'}{\partial \zeta} \right) + M_u \right] \\ \frac{\partial v}{\partial t} = \left[ -\left\{ \frac{1}{a} \frac{\partial E_h}{\partial \varphi} + uV_a \right\} - \left[ \frac{\partial v}{\partial \zeta} \right] \left[ \frac{1}{\rho a} \left( \frac{\partial p'}{\partial \varphi} - \frac{1}{\sqrt{\gamma}} \frac{\partial p_0}{\partial \rho} \frac{\partial p'}{\partial \zeta} \right) + M_u \right] \\ \frac{\partial w}{\partial t} = \left[ -\left\{ \frac{1}{a \cos \varphi} \left( u \frac{\partial w}{\partial \lambda} + v \cos \varphi \frac{\partial w}{\partial \varphi} \right) \right\} - \left[ \frac{\partial v}{\partial \zeta} \right] + \frac{g}{\sqrt{\gamma}} \frac{\rho_0}{\rho} \frac{\partial p'}{\partial \zeta} + M_u \right] + g \frac{\rho_0}{\rho} \left\{ \frac{(T - T_0)}{T} - \frac{T_0 p'}{T p_0} + \left( \frac{R_u}{R_d} - 1 \right) q^v - q^l - q^l \right\} \right] \\ \text{pressure} \quad \frac{\partial p'}{\partial t} = \left[ -\left\{ \frac{1}{a \cos \varphi} \left( u \frac{\partial p'}{\partial \lambda} + v \cos \varphi \frac{\partial p'}{\partial \varphi} \right) \right\} - \left[ \frac{\partial p'}{\partial \zeta} \right] + g \rho_0 \varphi - \frac{c_{pd}}{c_{vd}} p D \right] \\ \text{temperature} \quad \frac{\partial T}{\partial t} = \left[ -\left\{ \frac{1}{a \cos \varphi} \left( u \frac{\partial q^v}{\partial \lambda} + v \cos \varphi \frac{\partial q^v}{\partial \varphi} \right) \right\} - \left[ \frac{\partial q^v}{\partial \zeta} \right] - \left[ \frac{\partial q^v}{\partial \zeta} \right] - \left[ \frac{\partial q^v}{\partial \zeta} \right] + v \cos \varphi \frac{\partial q^v}{\partial \varphi} \right] - \left[ \frac{\partial q^v}{\partial \zeta} \right] - \left[ \frac{\partial q^v}{\partial \zeta} \right] + v \cos \varphi \frac{\partial q^v}{\partial \varphi} \right] \\ \text{turbulence} \quad \frac{\partial e_t}{\partial t} = \left[ -\left\{ \frac{1}{a \cos \varphi} \left( u \frac{\partial q^{t/t}}{\partial \lambda} + v \cos \varphi \frac{\partial q^t}{\partial \varphi} \right) \right\} - \left[ \frac{\partial q^{t/t}}{\partial \zeta} \right] - \left[ \frac{g}{\partial Q^t} \right] + \left[ \frac{g}{\partial Q^t} - \left( \frac{\partial q^{t/t}}{\partial \zeta} + V \cos \varphi \frac{\partial q^{t/t}}{\partial \varphi} \right) \right] - \left[ \frac{\partial q^{t/t}}{\partial \zeta} \right] + \left[ \frac{g}{\partial Q^t} - \left( \frac{\partial q^{t/t}}{\partial \zeta} + V \cos \varphi \frac{\partial q^{t/t}}{\partial \varphi} \right] + \left[ \frac{\partial q^{t/t}}{\partial \zeta} \right] + \frac{g}{\partial Q^t} - \left[ \frac{\partial q^{t/t}}{\partial \zeta} + M_{q,t} \right] \\ \text{turbulence} \quad \frac{\partial e_t}{\partial t} = \left[ -\left\{ \frac{1}{a \cos \varphi} \left( u \frac{\partial q^{t/t}}{\partial \lambda} + v \cos \varphi \frac{\partial q^{t/t}}{\partial \varphi} \right) \right] - \left[ \frac{\partial q^{t/t}}{\partial \zeta} \right] + \left[ \frac{g}{\partial Q^t} \right] + \left[ \frac{\partial q^{t/t}}{\partial \zeta} + \frac{g}{\partial Q^t} \right] + \frac{g}{\partial Q^t} \left[ \frac{\partial q^{t/t}}{\partial \zeta} + M_{q,t} \right] \\ \text{turbulence} \quad \frac{\partial e_t}{\partial t} = \left[ -\left\{ \frac{1}{a \cos \varphi} \left( u \frac{\partial q^{t/t}}{\partial \lambda} + v \cos \varphi \frac{\partial q^{t/t}}{\partial \varphi} \right] - \left[ \frac{\partial q^{t/t}}{\partial \zeta} \right] + \frac{g}{\partial Q^t} \right] + \frac{g}{\partial Q^t} \left[ \frac{\partial q^{t/t}}{\partial \zeta} + M_{q,t} \right] \\ \text{turbulence} \quad \frac{\partial e_t}{\partial t} = \left[ -\left\{ \frac{1}{a \cos \varphi} \left( u \frac{\partial q^{t/t}}{\partial \lambda} + v \cos \varphi \frac{\partial q^{t/t}}{\partial \varphi} \right] - \left[ \frac{\partial q^{t/t}}{\partial \zeta} \right] + \frac{g}{\partial Q^t} \left[ \frac{\partial q^{t/t}}{\partial \zeta} + \frac{g}{\partial Q^t} \right] + \frac{g}{\partial Q^t} \left[ \frac{\partial q^{t/t}}{\partial \zeta} \right] + \frac{g}{\partial Q^t} \left[ \frac{$$

![](_page_29_Picture_1.jpeg)

![](_page_29_Figure_2.jpeg)

Tridiagonal solve for matrix ~50-100

Current status of implementation:

- > structure grid motif works well on data parallel, high-memory bandwidth "accelerator"
- > tridiagonal solve fits in L2 cash on CPU and makes good use of single thread performance

#### A crude way to consider the architecture we need for physics and dynamics in COSMO (preliminary)

![](_page_30_Figure_3.jpeg)

Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich

![](_page_31_Picture_1.jpeg)

#### **OPCODE** project: can we run entire MeteoCH production suit on a node with ~8-16 GPUs?

![](_page_31_Picture_3.jpeg)

Cray XT4 (Production Machine at CSCS) 246 AMD Opteron Barcelona

OPCODE

![](_page_31_Picture_6.jpeg)

Server O(20) GPUs Many such Server for ensemble runs

![](_page_31_Picture_9.jpeg)

![](_page_31_Picture_10.jpeg)

![](_page_31_Picture_11.jpeg)

![](_page_31_Picture_12.jpeg)

![](_page_31_Picture_13.jpeg)

![](_page_32_Picture_1.jpeg)

# CONCLUSIONS

We can have both, higher resolution and larger ensembles!

But we have to invest in code refactoring and algorithm re-engineering in order to consider appropriate architectures for our problem\*

(\*) applies to all 12 projects we are studying in the HP2C platform (see <u>www.hp2c.ch</u>)

![](_page_33_Picture_1.jpeg)

### **THANK YOU!**