ETH zürich

Supercomputers: instruments for science or dinosaurs that haven't gone extinct yet?

Thomas C. Schulthess

Do you really mean dinosaurs?

We must be in the wrong movie

Not much has changed since the late 1980s

- pick up any book on HPC from the early 1990s, e.g. this on from 1993
- Shared-memory multiprocessors
  - cash coherency …
  - pipelining, SIMD, threading, …
- Explicitly parallel languages (imperative programming)
  - Fortran 90
  - High Performance Fortran
  - Explicitly Parallel Programming Environment … (PVM)
  - …
- No fundamental changes sind then!

Incremental and other changes

- Evolution of Fortran
- Evolution of C++ – you can mention C++ at an HPC conference without being thrown out and template meta-programming is somewhat usable
- PVM to MPI
- OpenMP besides pthreads
- PGAS languages (not sure how broadly accepted)

Only two potentially disruptive changes (in terms of broader acceptance)

(1) **Python** – but not in HPC please!

(2) **CUDA with GPU** – are you kidding me?

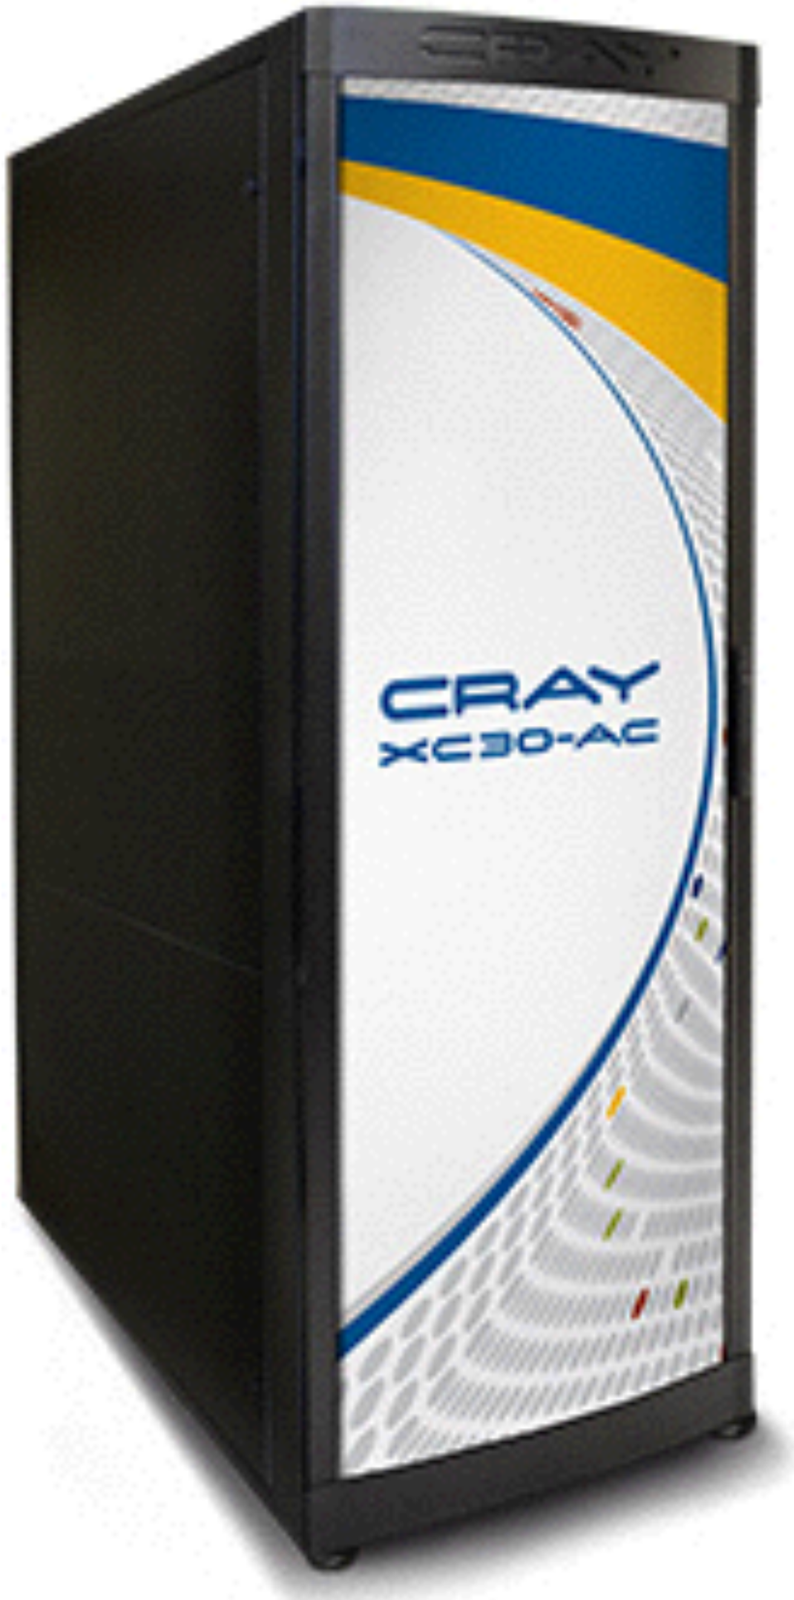

intel inside

Pete's Top 5 HPC initiatives or technologies to watch in 2013 (www.hpcwire.com):

- Big Data solutions
- Industrial use of petascale supercomputers
- Accelerator and Heterogeneous Computing
- The fusion of HPC technologies into big data computing applications
- Anything Cray is doing!

Pete Ungaro
CEO, Cray Inc.

The future just keeps getting brighter at Cray these days. It's in no small part due to the strategic genius of its charismatic CEO, Peter Ungaro. Ungaro is largely credited for masterminding Cray's entry into the multi-billion dollar "Big Data" market in early 2012 with the introduction of YARC Data. Following Cray's recent acquisition of Appro, shareholders are experiencing a fresh dose of investor confidence, as evidenced by the steady increase in Cray's stock value over the past few years. So what's next on the radar from Cray this year? Only time will tell, but you can bet we are following Mr. Ungaro again this year.

source: www.hpcwire.com

Raj's few worth considerations (www.hpcwire.com)

- Heterogeneous computing becomes the next step, not a detour – not just in high-end, but mainstream HPC
- Big Data transforms HPC; First real business value propositions from use of analytics beyond niche using HPC
- Domain optimised HPC Clouds – the promise of HPC with the value of the cloud model
- Integration – driving fundamental changes in system architecture
- Everyone wants their own – … not just investing in "having" the best of class, but also in "making" the best of class

Raj Hazra
Vice President and General Manager of High Performance Computing, Intel

According to our analyst friends across the board, Intel is on a hot track. The company is in the process of actively transforming the company and go-to-market plan in a direction that takes it from traditional HPC and towards a new technical computing track. One of the key players behind this strategy is Intel's Raj Hazra, who appears to be Intel's technical computing rising star. Already he's made waves in 2014 with the hiring of former Penguin Computing CEO, Charles Wuischpard as the Vice President of the Data Center Group and General Manager of Workstations and Technical Computing. It will be interesting to see what else Hazra has up his sleeve for 2014. We caught up with him to ask about his thoughts on the year and the direction of supercomputing.

source: www.hpcwire.com

What is "Big Data"?

It is not big storage, rather one should think of "Data Science"

So what is "Data Science"?

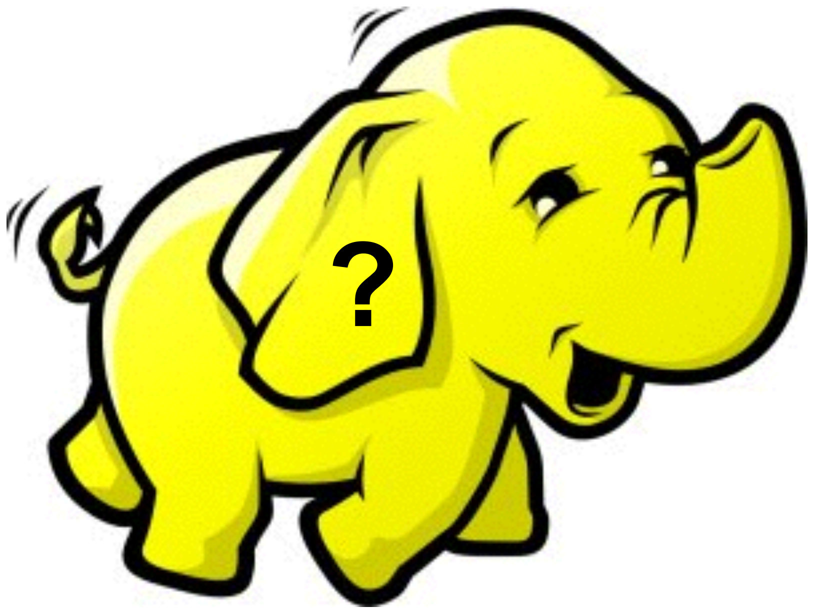

Forget about MapReduce, this is just one of many algorithmic motifs, HDFS is the key: "performance sucks but the API is important"

"Data Science" is not simply engineering of big databases

- Traditional database engineering focuses on database design, software engineering, data flow, data extraction, etc.
- Data science is about
  - **discovery**: "identify[ing] good data sources and metrics. Sometimes request[ing] the data to be created"
  - **access**: "access the data, sometimes via an API, …"
  - **distilling**: "extract from the data the information that leads to …"

source of quotes: Vincent Granville's "Developing Analytic Talent, becoming a data scientists"

Data science is more than statistic …

- Analysis of NASA images to discover new planets, astroids, …
- …
- Computational chemistry to simulate new molecules for … treatment
- …
- Returning highly relevant results to any Google search
- Scoring all credit card transactions (fraud detection)
- Tax fraud detection and detection of terrorism
- Weather forecasts

source: Vincent Granville's "Developing Analytic Talent, becoming a data scientists"

So serious data science is just science!

"I have studied all available charts of the planets and stars and none of them match the others. There are just as many measurements and methods as there are astronomers and all of them disagree. What is needed is a long-term project with the aim of mapping the heavens conducted from a single location over a period of several years."

–Tycho Brahe, 1563

The first "BigData" project in history

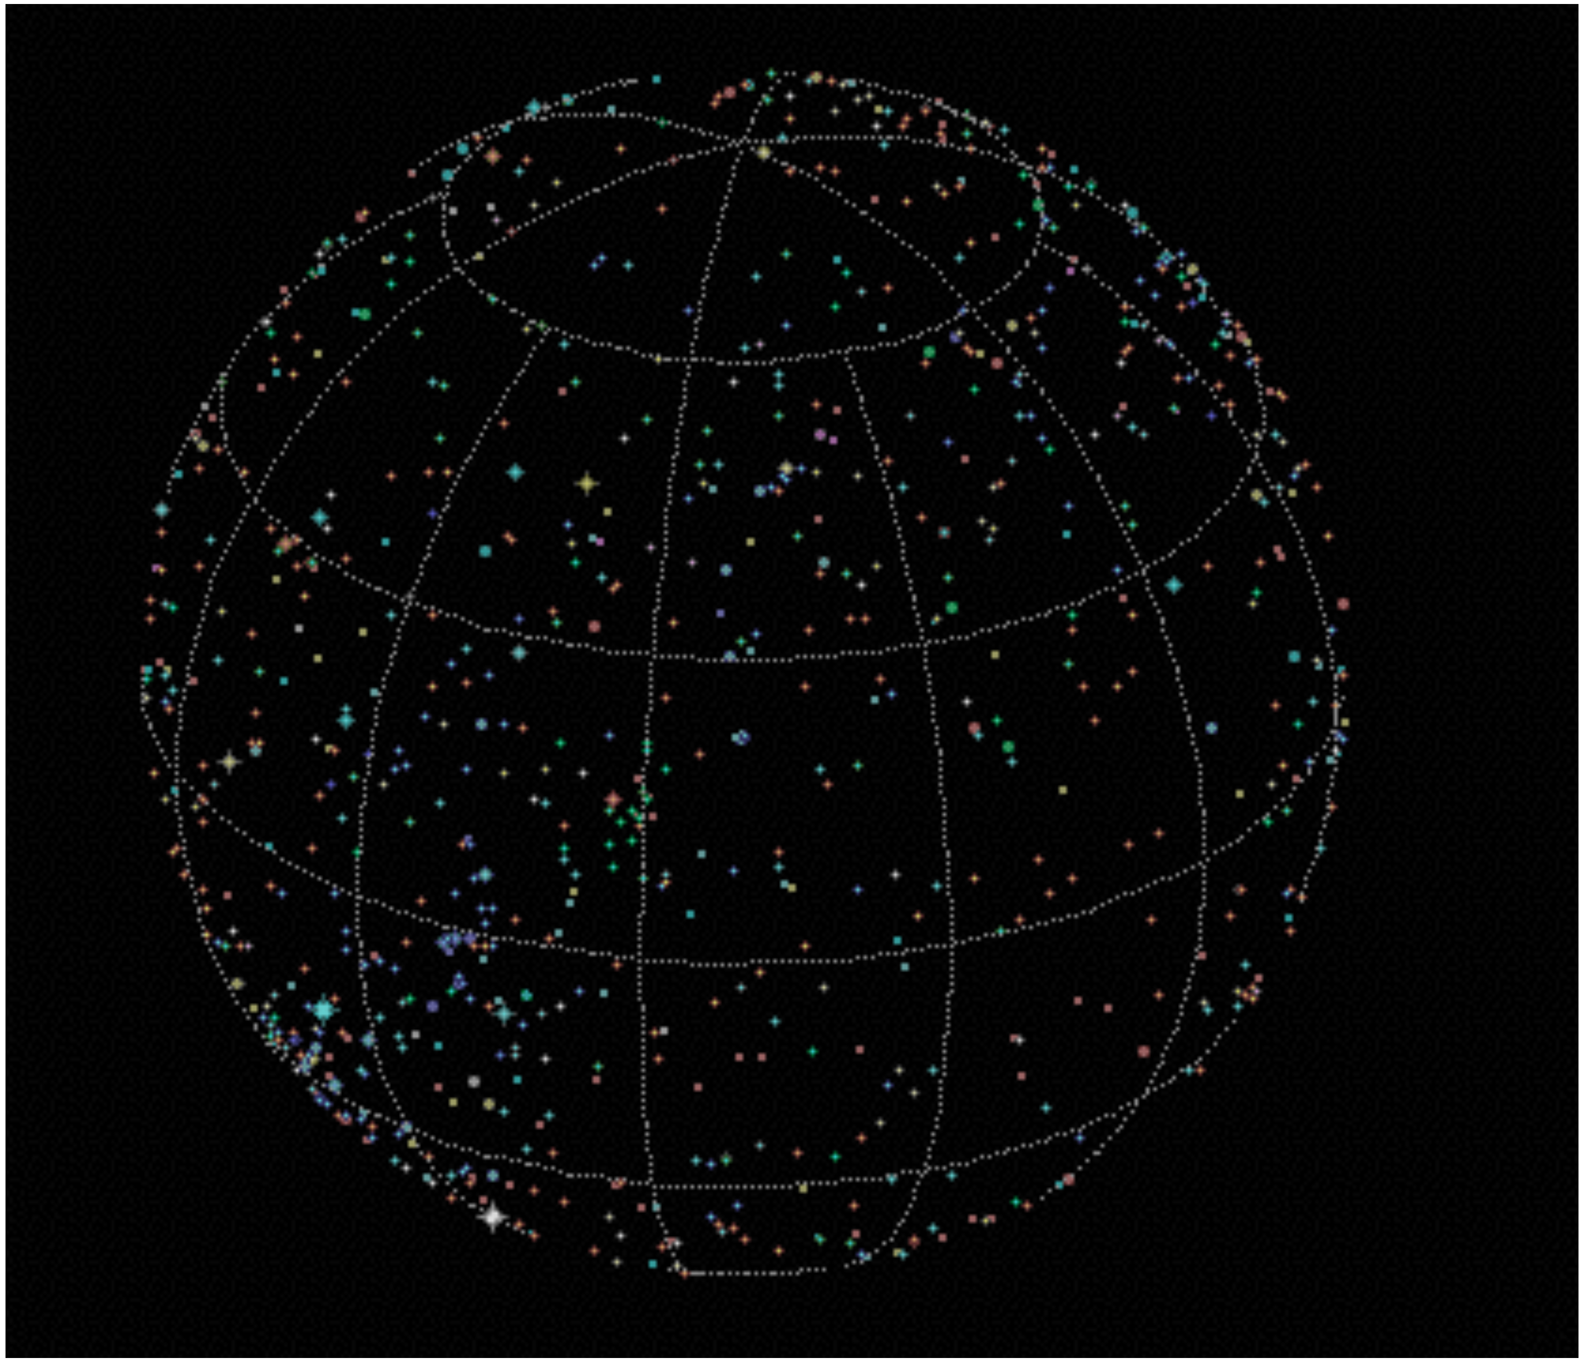

source: www.pafko.com/tycho/

ETH zürich

Data given to Johannes Kepler

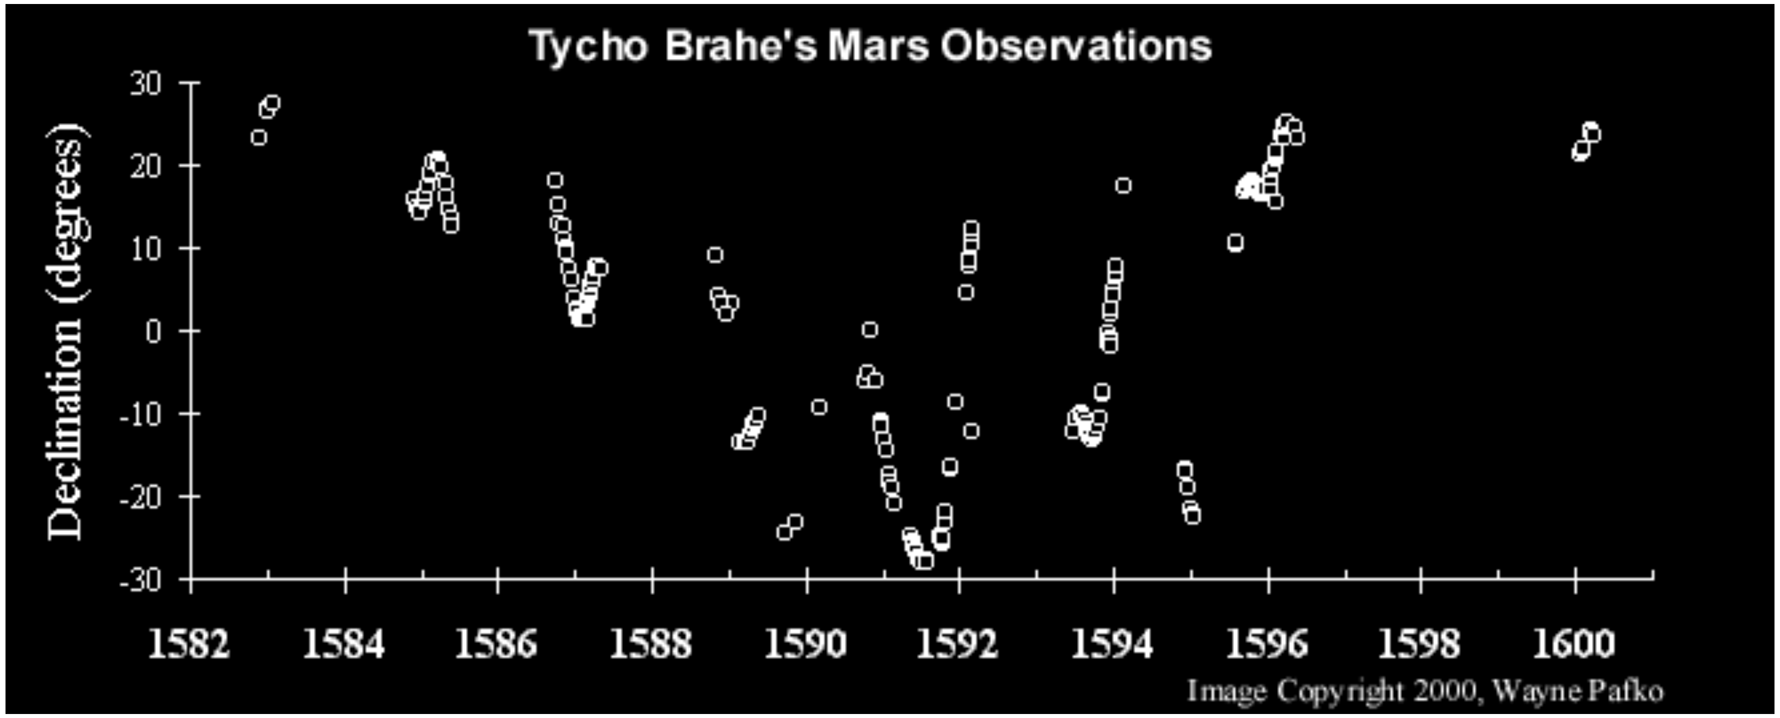

source: www.pafko.com/tycho/

Kepler's modelling and simulations

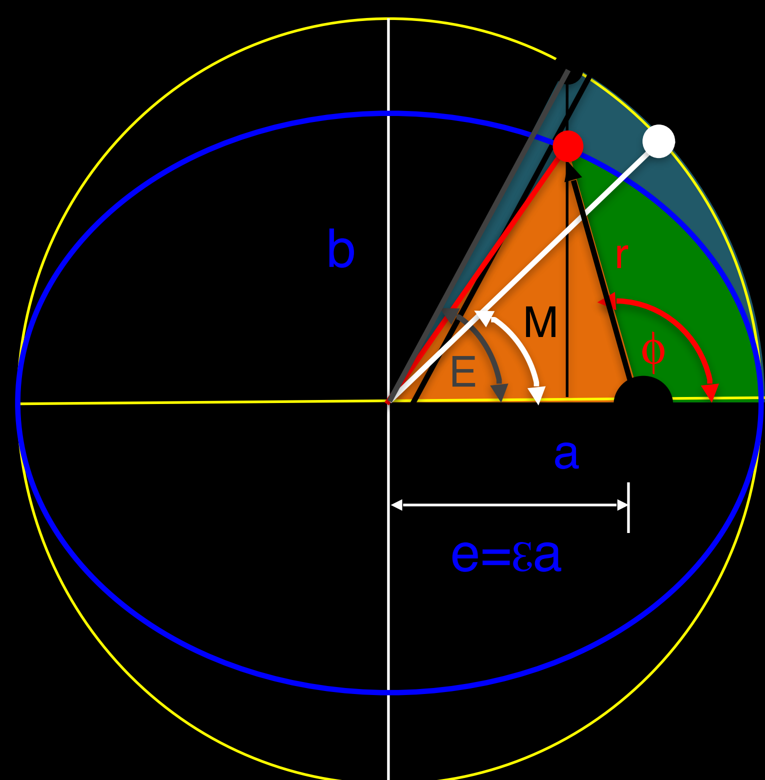

$$M = E - \varepsilon \sin E$$

1. Solve $E(M)$ (Numerics)
2. Solve $\phi(E)$ (Geometry)

After Kepler's analysis and Newton's theory …

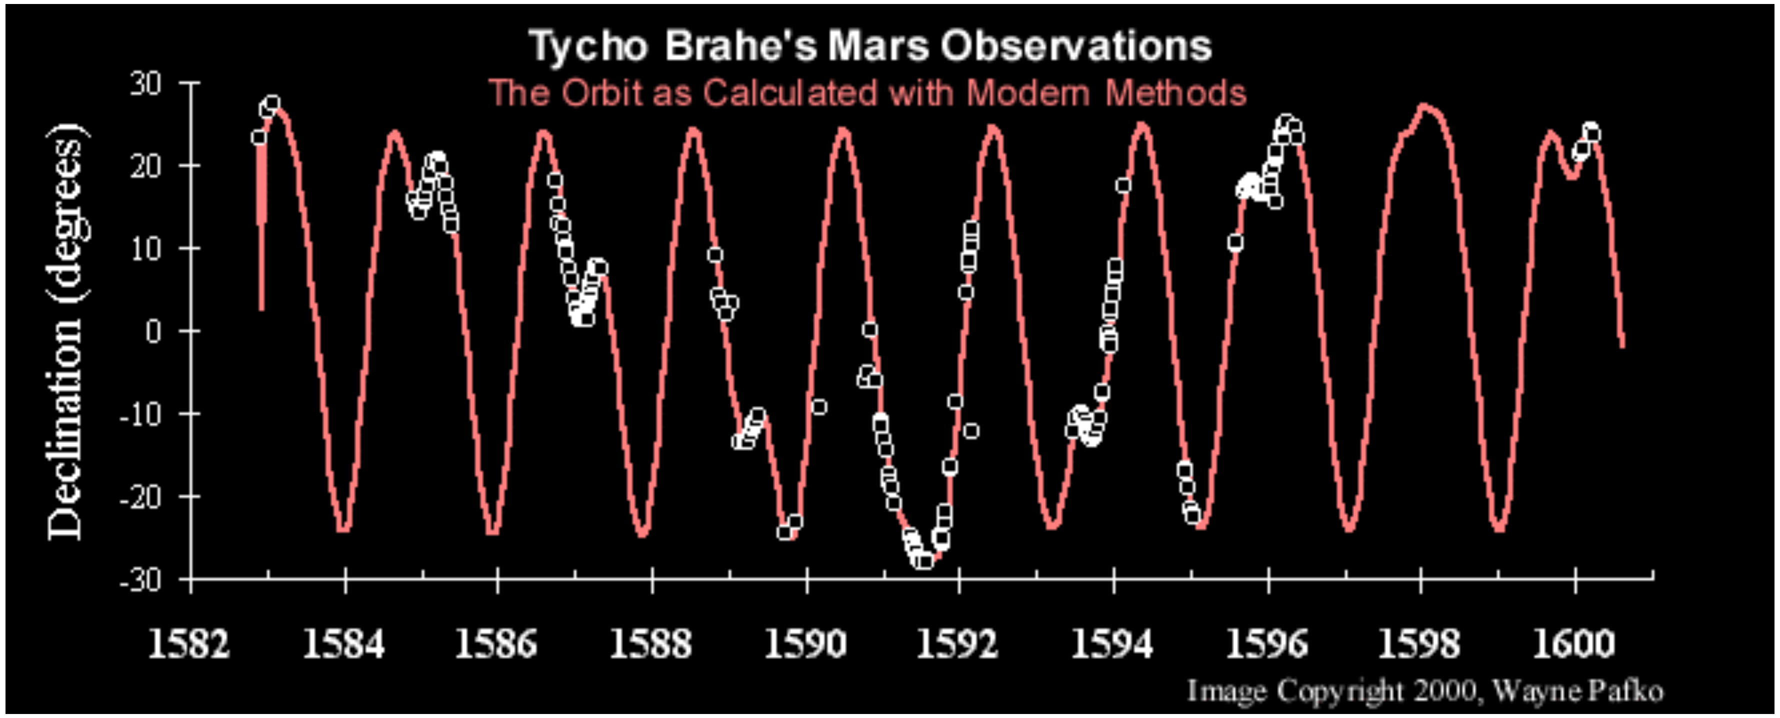

source: www.pafko.com/tycho/

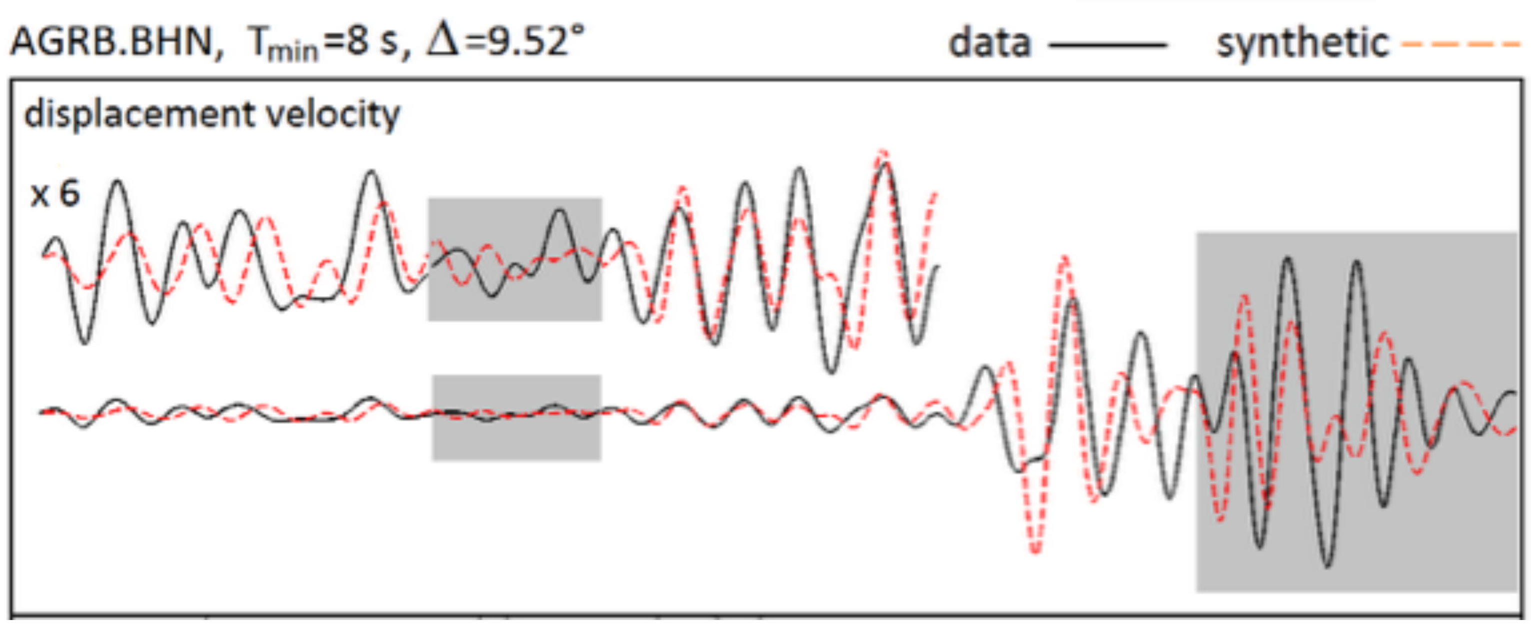

source: A. Fichtner, ETH Zurich

Data from many stations and earthquakes

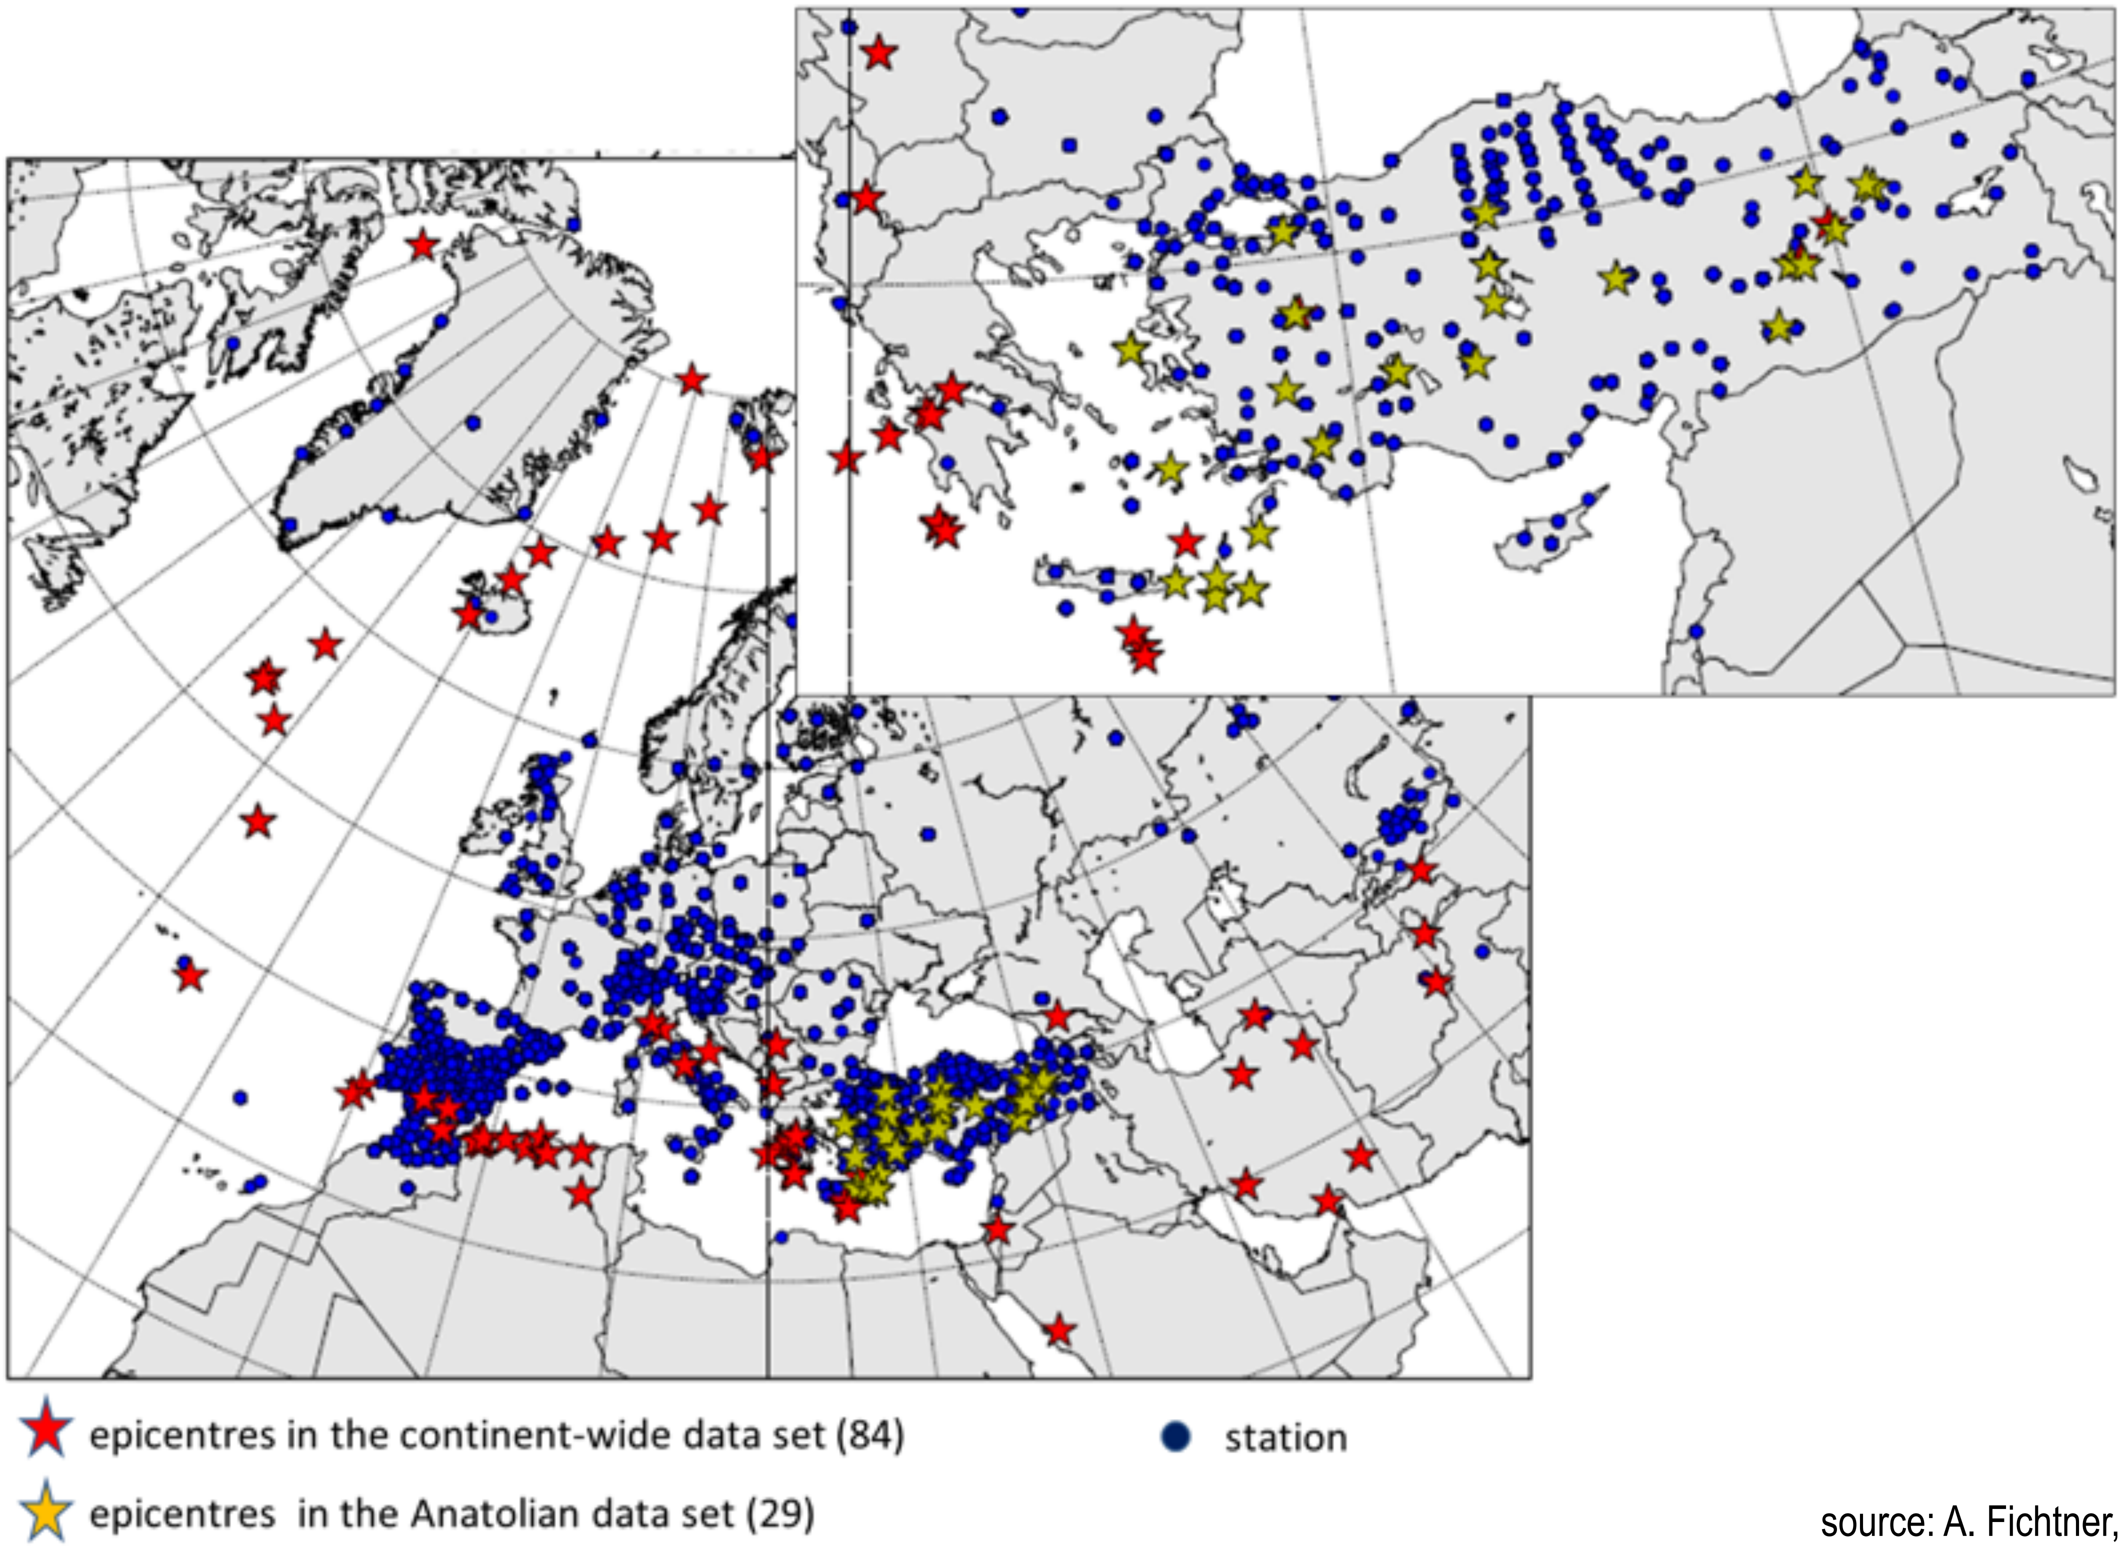

source: A. Fichtner, ETH Zurich

Very large simulations allow inverting large data sets in oder to generate high-resolution models of the earth's mantle

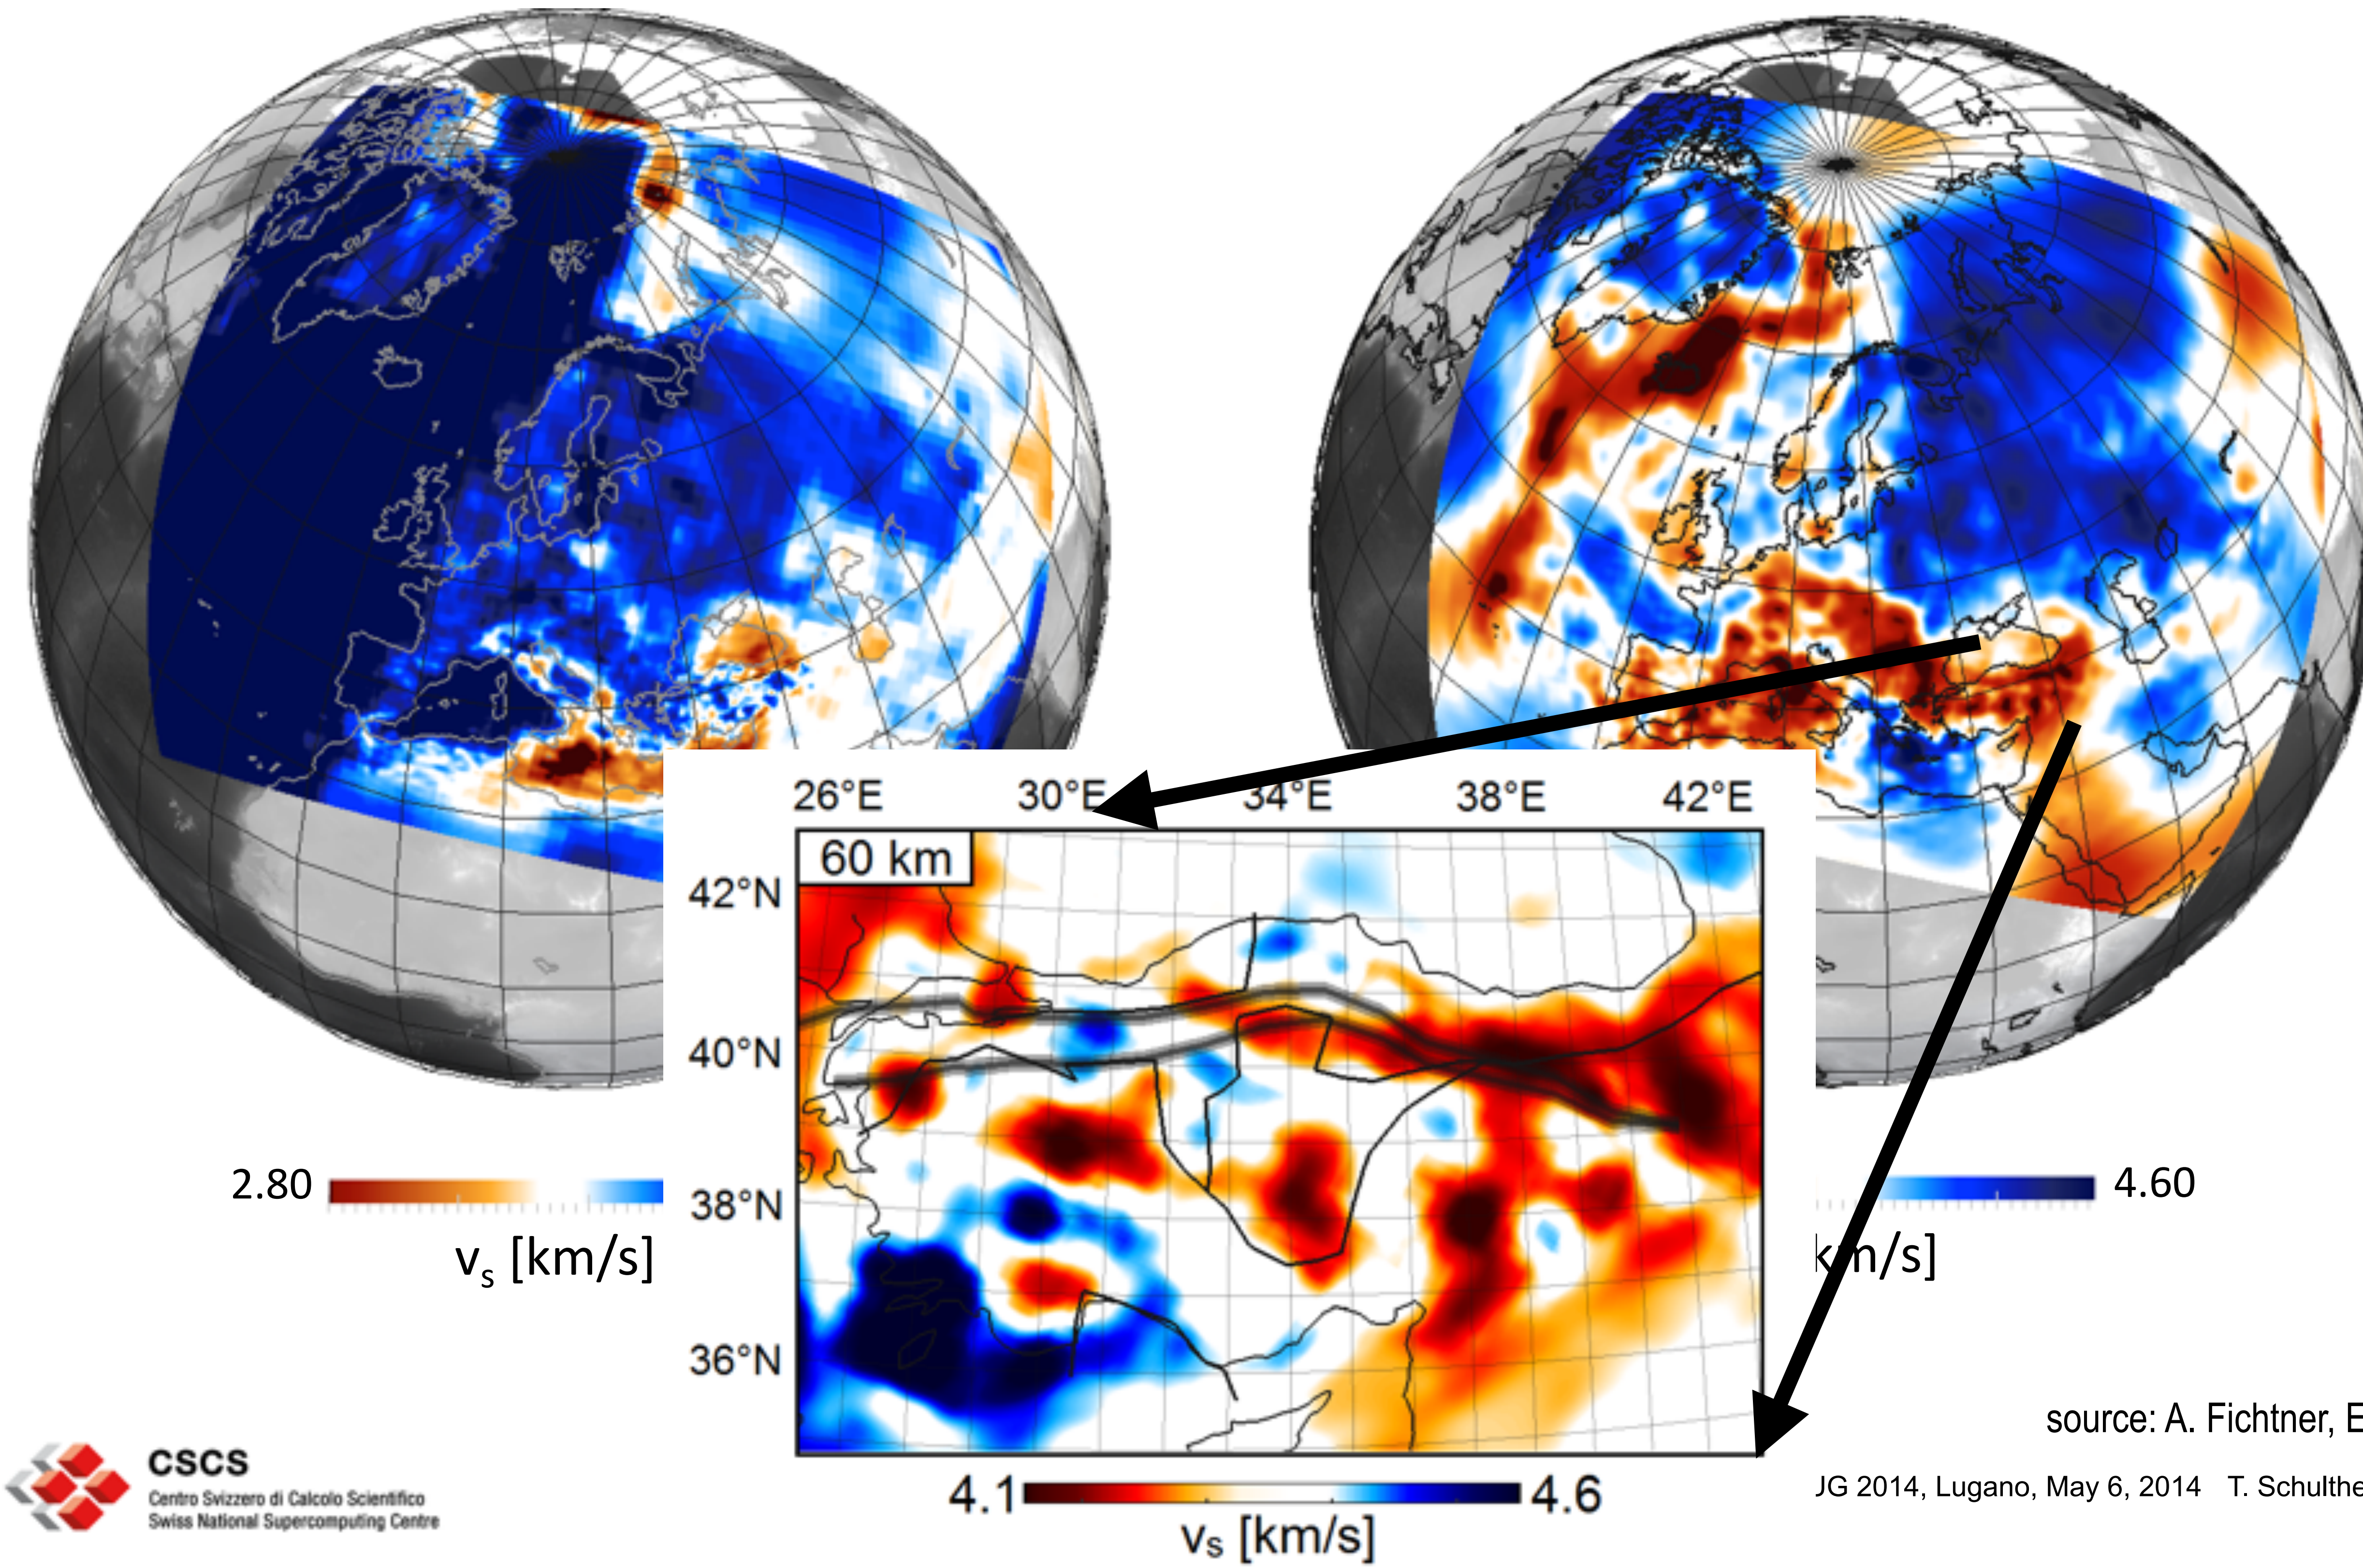

source: A. Fichtner, ETH Zurich

Ice floats!

first early science result on “Piz Daint” (Joost VandeVondele, Jan. 2014)

high-accuracy quantum simulation produce correct results for water

ETH zürich

Jacob's ladder to heaven of chemical accuracy (J. Perdew)

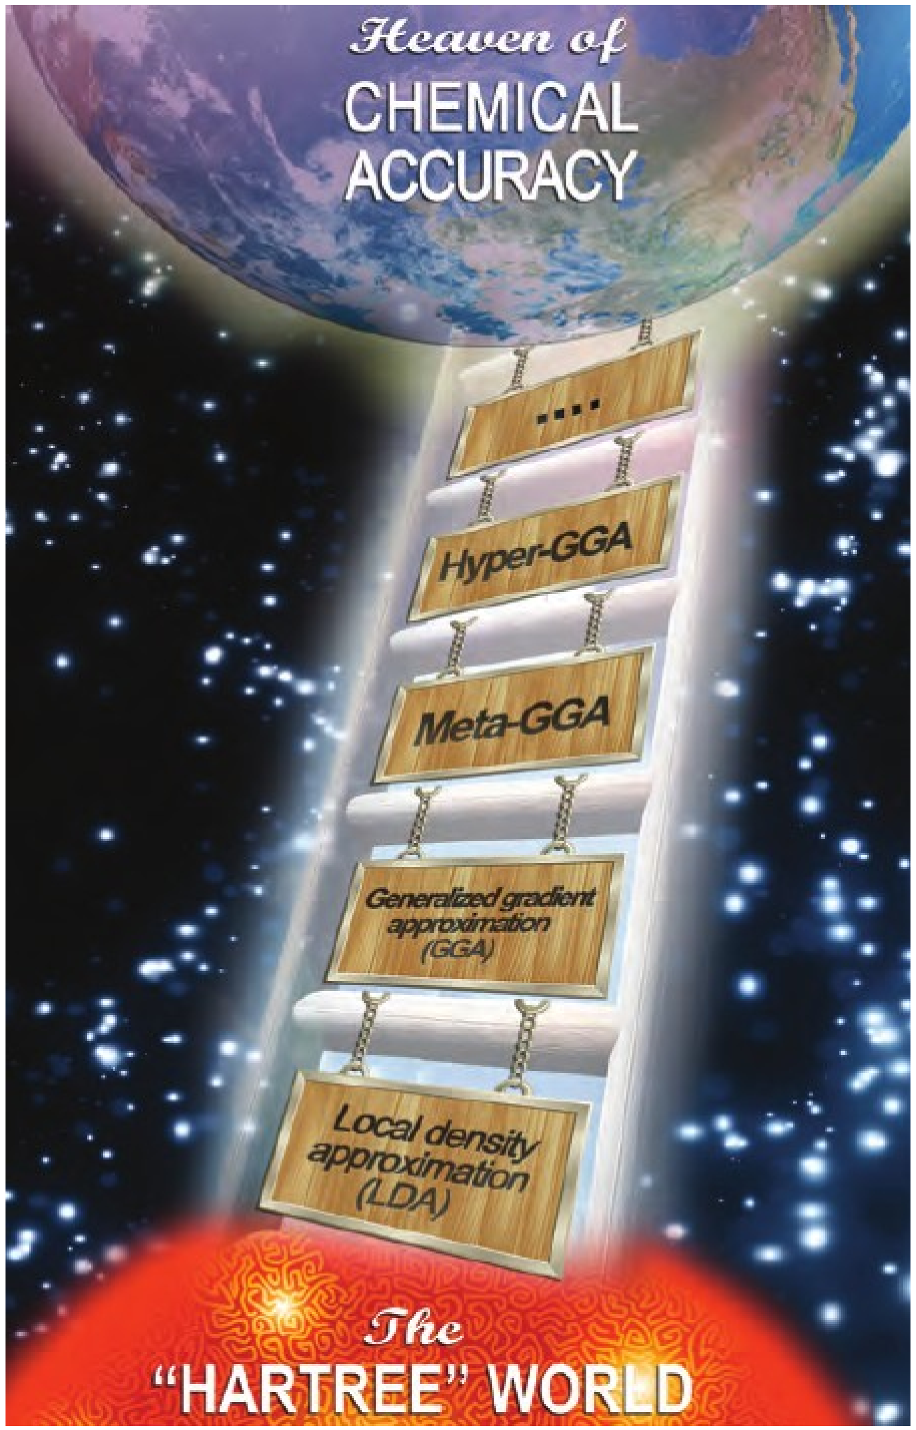

Full many-body Schrödinger Equation:

$$(\mathbb{H} - E)\Psi(\xi_1, ..., \xi_N) = 0$$

$$\mathbb{H} = \sum_{i=1}^{N} \left(-\frac{\hbar^2}{2m}\nabla_i^2 + v(\xi_i) \right) + \frac{1}{2} \sum_{i \neq j = 1}^{N} w(\xi_i, \xi_j)$$

$$w(\xi_i, \xi_j) = \frac{e^2}{|\vec{r}_i - \vec{r}_j|}$$

$$\xi = \{\vec{r}, \sigma\}$$

Ice floats with new "rung 5" simulations using MP2-based simulations with CP2K (VandeVondele & Hutter)

Ice didn't float with previous simulations using "rung 4" hybrid functionals

Kohn-Sham Equation with Local Density Approximation:

$$\left(-\frac{\hbar^2}{2m}\nabla^2 + v_{\mathrm{LDA}}(\vec{r}) \right) \psi_i(\vec{r}) = \epsilon_i \psi_i(\vec{r})$$

Modelling interactions between water molecules requires quantum simulations with extreme accuracy

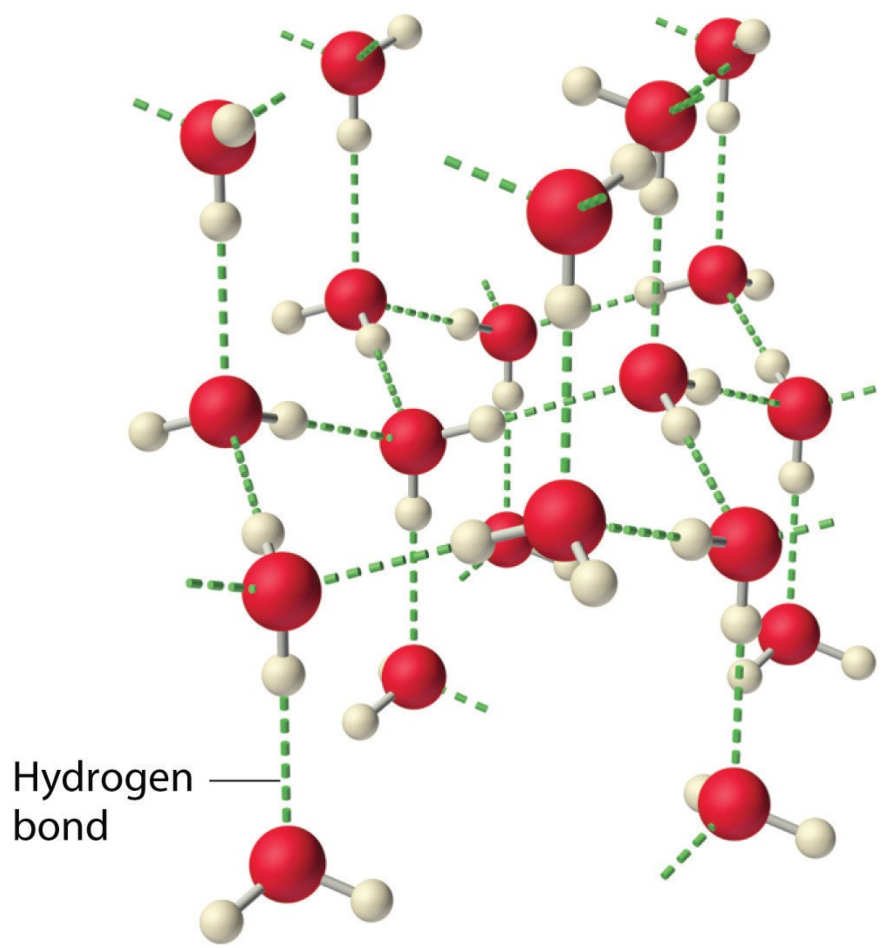

Energy scales

total energy: -76.438… a.u.
hydrogen bonds: ~0.0001 a.u.

required accuracy: 99.9999%

If you don’t get this right, then …

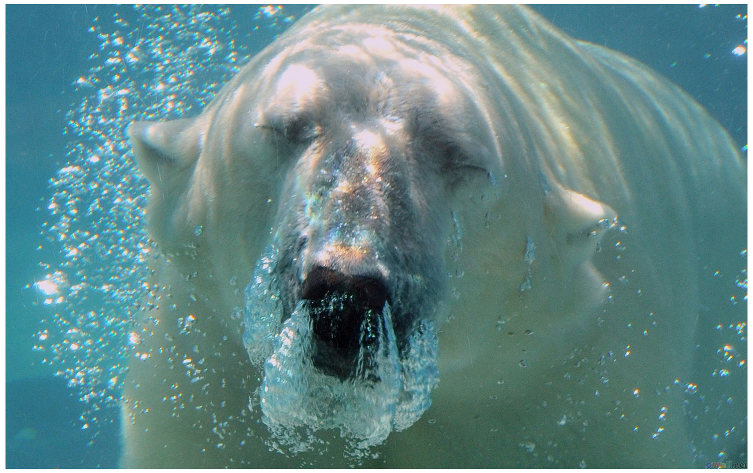

ab initio simulation prior to Joost’s MP2-based runs with CP2K

Pillars of the scientific method

Mathematics / Simulation

(1) Synthesis of models and data: recognising characteristic features of complex systems with calculations of limited accuracy (e.g. inverse problems)

(2) Solving theoretical problems with high precision: complex structures emerge from simple rules (natural laws), more accurate predictions from "beautiful" theory (in the Penrose sense)

Theory (models)

Experiment (data)

Pillars of the scientific method

Mathematics / Simulation

(1) Synthesis of models and data: recognising characteristic features of complex systems with calculations of limited accuracy (e.g. inverse problems)

(2) Solving theoretical problems with high precision: complex structures emerge from simple rules (natural laws), more accurate predictions from "beautiful" theory

Note the changing role of high-performance computing: HPC is now an essential tool for science, used by all scientists (for better or worse), rather than being limited to the domain of applied mathematics and providing numerical solution to theoretical problems only few understand

Given all this, here is a bold prediction for what Cray's roadmap might look like – the essence of BigData vs. HPC

BigData version:
> HDFS
> Python and other dynamic languages

HPC version:
> Fortran
> MPI+OpenMP

“We need fundamentally new programming models, but won’t accept anything other than ~~Fortran~~ MPI+OpenMP”

–the HPC community discussing exascale challenges

$$V(r) = \sum_{\text{bonds}} k_b(b-b_0)^2 + \sum_{\text{angles}} k_\theta(\theta-\theta_0)^2 + \sum_{\text{dihedrals}} k_\phi(1+\cos(n\phi-\phi_0)) + \sum_{\text{impropers}} k_\psi(\psi-\psi_0)^2 + \sum_{\substack{\text{non-bonded}\\ \text{pairs}(i,j)}} 4\varepsilon_{ij}\left[\left(\frac{\sigma_{ij}}{r_{ij}}\right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}}\right)^{6}\right] + \sum_{\substack{\text{non-bonded}\\ \text{pairs}(i,j)}} \frac{q_i q_j}{\varepsilon_D r_{ij}}$$

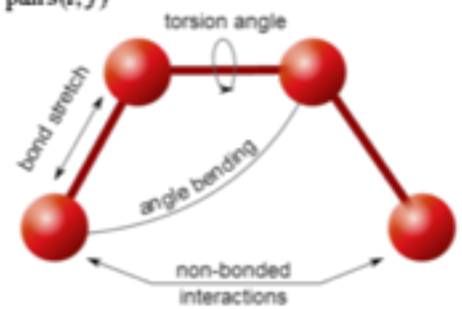

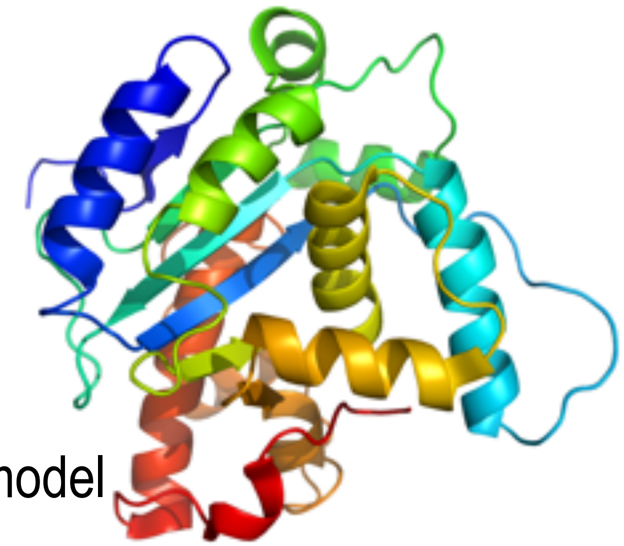

Physical model

Mathematical description

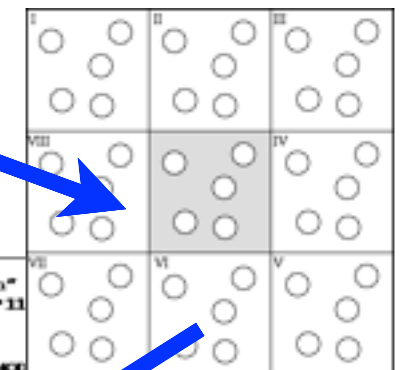

Discretisation / algorithm

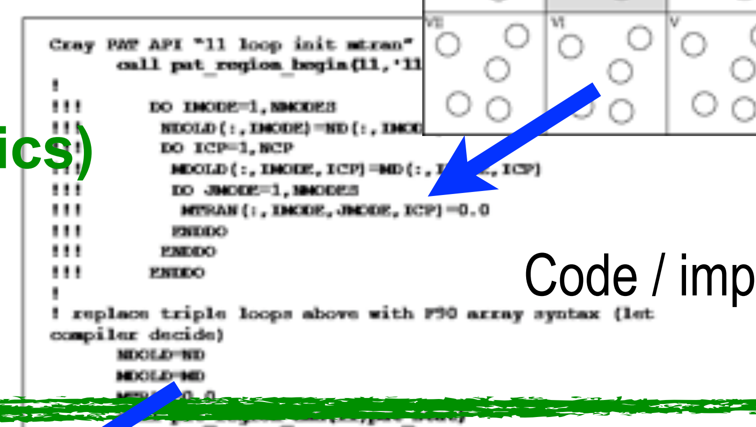

Code / implementation

Domain science (incl. applied mathematics)

Code compilation

A given supercomputer

“Port” serial code to supercomputers

- **> vectorize**
- **> parallelize**
- **> petascaling**
- **> exascaling**
- **> ...**

Computer engineering (& computer science)

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

velocities

$$\frac{\partial u}{\partial t} = -\left\{\frac{1}{a\cos\varphi}\frac{\partial E_h}{\partial\lambda} - vV_a\right\} - \dot{\zeta}\frac{\partial u}{\partial\zeta} - \frac{1}{\rho a\cos\varphi}\left(\frac{\partial p'}{\partial\lambda} - \frac{1}{\sqrt{\gamma}}\frac{\partial p_0}{\partial\lambda}\frac{\partial p'}{\partial\zeta}\right) + M_u$$

$$\frac{\partial v}{\partial t} = -\left\{\frac{1}{a}\frac{\partial E_h}{\partial\varphi} + uV_a\right\} - \dot{\zeta}\frac{\partial v}{\partial\zeta} - \frac{1}{\rho a}\left(\frac{\partial p'}{\partial\varphi} - \frac{1}{\sqrt{\gamma}}\frac{\partial p_0}{\partial\varphi}\frac{\partial p'}{\partial\zeta}\right) + M_v$$

$$\frac{\partial w}{\partial t} = -\left\{\frac{1}{a\cos\varphi}\left(u\frac{\partial w}{\partial\lambda} + v\cos\varphi\frac{\partial w}{\partial\varphi}\right)\right\} - \dot{\zeta}\frac{\partial w}{\partial\zeta} + \frac{g}{\sqrt{\gamma}}\frac{\rho_0}{\rho}\frac{\partial p'}{\partial\zeta} + M_w + g\frac{\rho_0}{\rho}\left\{\frac{(T-T_0)}{T} - \frac{T_0 p'}{T p_0} + \left(\frac{R_v}{R_d} - 1\right)q^v - q^l - q^f\right\}$$

pressure

$$\frac{\partial p'}{\partial t} = -\left\{\frac{1}{a\cos\varphi}\left(u\frac{\partial p'}{\partial\lambda} + v\cos\varphi\frac{\partial p'}{\partial\varphi}\right)\right\} - \dot{\zeta}\frac{\partial p'}{\partial\zeta} + g\rho_0 w - \frac{c_{pd}}{c_{vd}}pD$$

temperature

$$\frac{\partial T}{\partial t} = -\left\{\frac{1}{a\cos\varphi}\left(u\frac{\partial T}{\partial\lambda} + v\cos\varphi\frac{\partial T}{\partial\varphi}\right)\right\} - \dot{\zeta}\frac{\partial T}{\partial\zeta} - \frac{1}{\rho c_{vd}}pD + Q_T$$

water

$$\frac{\partial q^v}{\partial t} = -\left\{\frac{1}{a\cos\varphi}\left(u\frac{\partial q^v}{\partial\lambda} + v\cos\varphi\frac{\partial q^v}{\partial\varphi}\right)\right\} - \dot{\zeta}\frac{\partial q^v}{\partial\zeta} - (S^l + S^f) + M_{q^v}$$

$$\frac{\partial q^{l,f}}{\partial t} = -\left\{\frac{1}{a\cos\varphi}\left(u\frac{\partial q^{l,f}}{\partial\lambda} + v\cos\varphi\frac{\partial q^{l,f}}{\partial\varphi}\right)\right\} - \dot{\zeta}\frac{\partial q^{l,f}}{\partial\zeta} - \frac{g}{\sqrt{\gamma}}\frac{\rho_0}{\rho}\frac{\partial P_{l,f}}{\partial\zeta} + S^{l,f} + M_{q^{l,f}}$$

turbulence

$$\frac{\partial e_t}{\partial t} = -\left\{\frac{1}{a\cos\varphi}\left(u\frac{\partial e_t}{\partial\lambda} + v\cos\varphi\frac{\partial e_t}{\partial\varphi}\right)\right\} - \dot{\zeta}\frac{\partial e_t}{\partial\zeta} + K_m^v\frac{g\rho_0}{\sqrt{\gamma}}\left\{\left(\frac{\partial u}{\partial\zeta}\right)^2 + \left(\frac{\partial v}{\partial\zeta}\right)^2\right\} + \frac{g}{\rho\theta_v}F^{\theta_v} - \frac{\sqrt{2}e_t^{3/2}}{\alpha_M l} + M_{e_t}$$

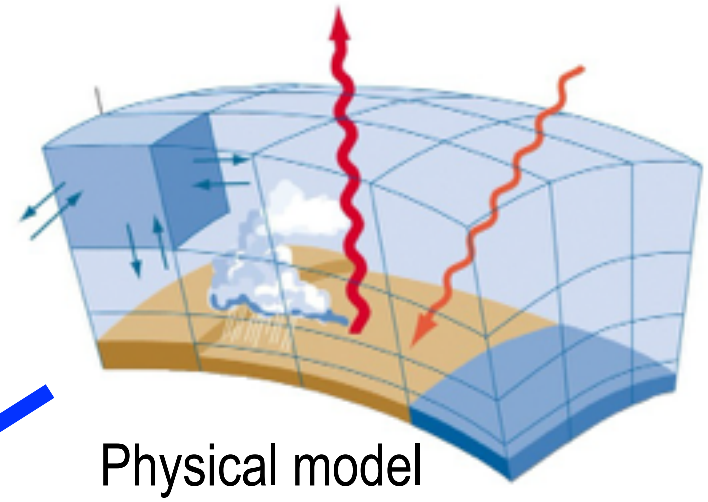

Physical model

Mathematical description

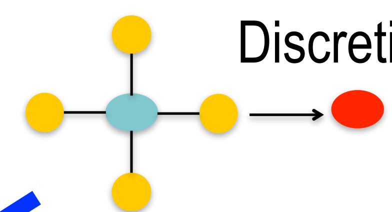

Discretisation / algorithm

Domain science (incl. applied mathematics)

```
lap(i,j,k) = -4.0 * data(i,j,k) +
    data(i+1,j,k) + data(i-1,j,k) +
    data(i,j+1,k) + data(i,j-1,k);
```

Code / implementation

Code compilation

"Port" serial code to supercomputers
- > vectorize
- > parallelize
- > petascaling
- > exascaling
- > ...

Computer engineering (& computer science)

Architectural options / design

A given supercomputer

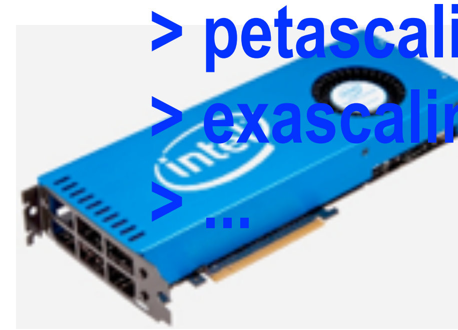

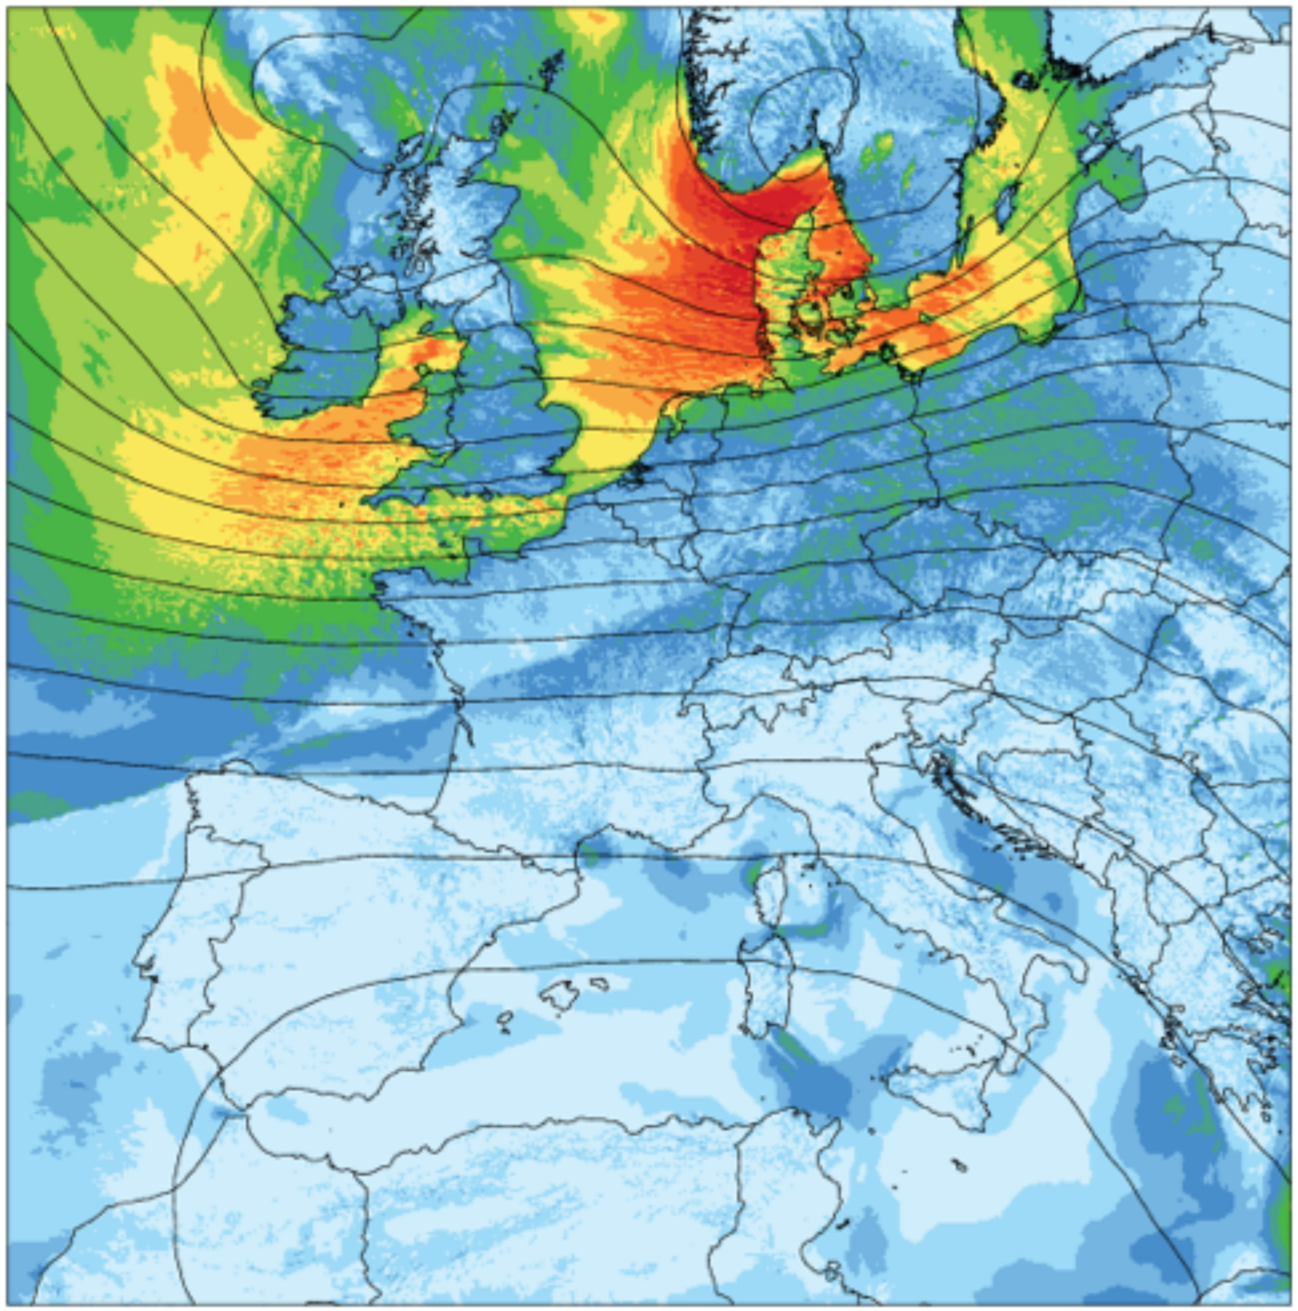

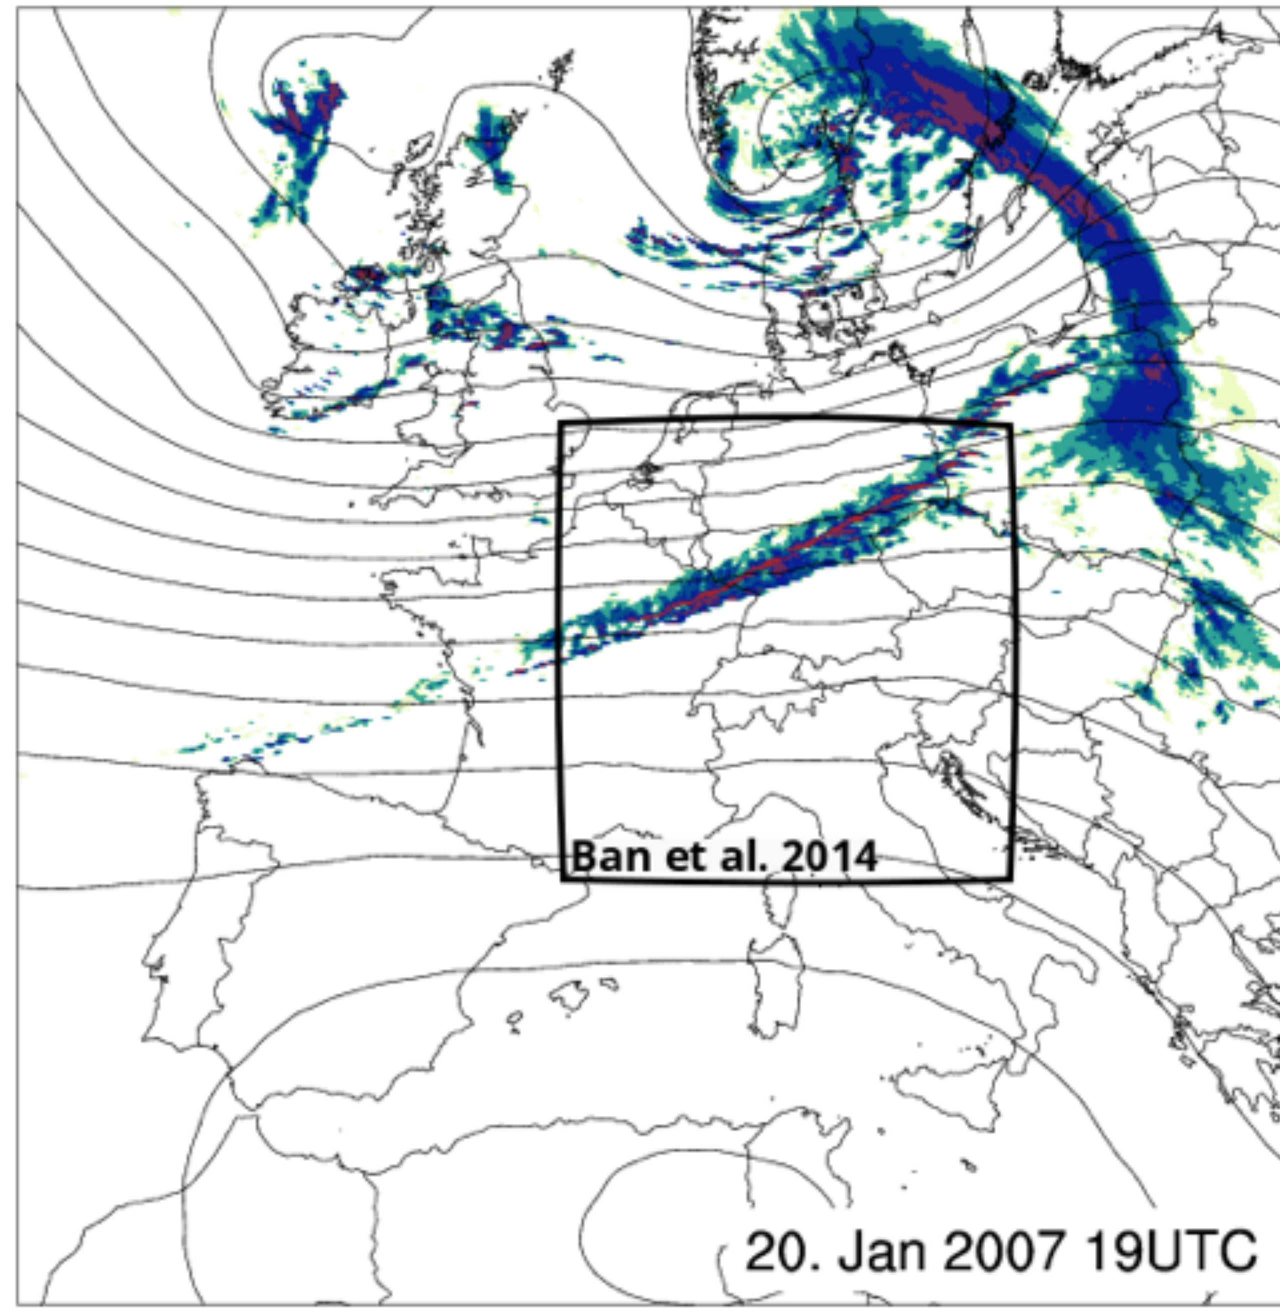

source: David Leutwyler

ETH zürich

Speedup of COSMO-2 production problem – apples to apples comparison with 33h forecast of Meteo Swiss

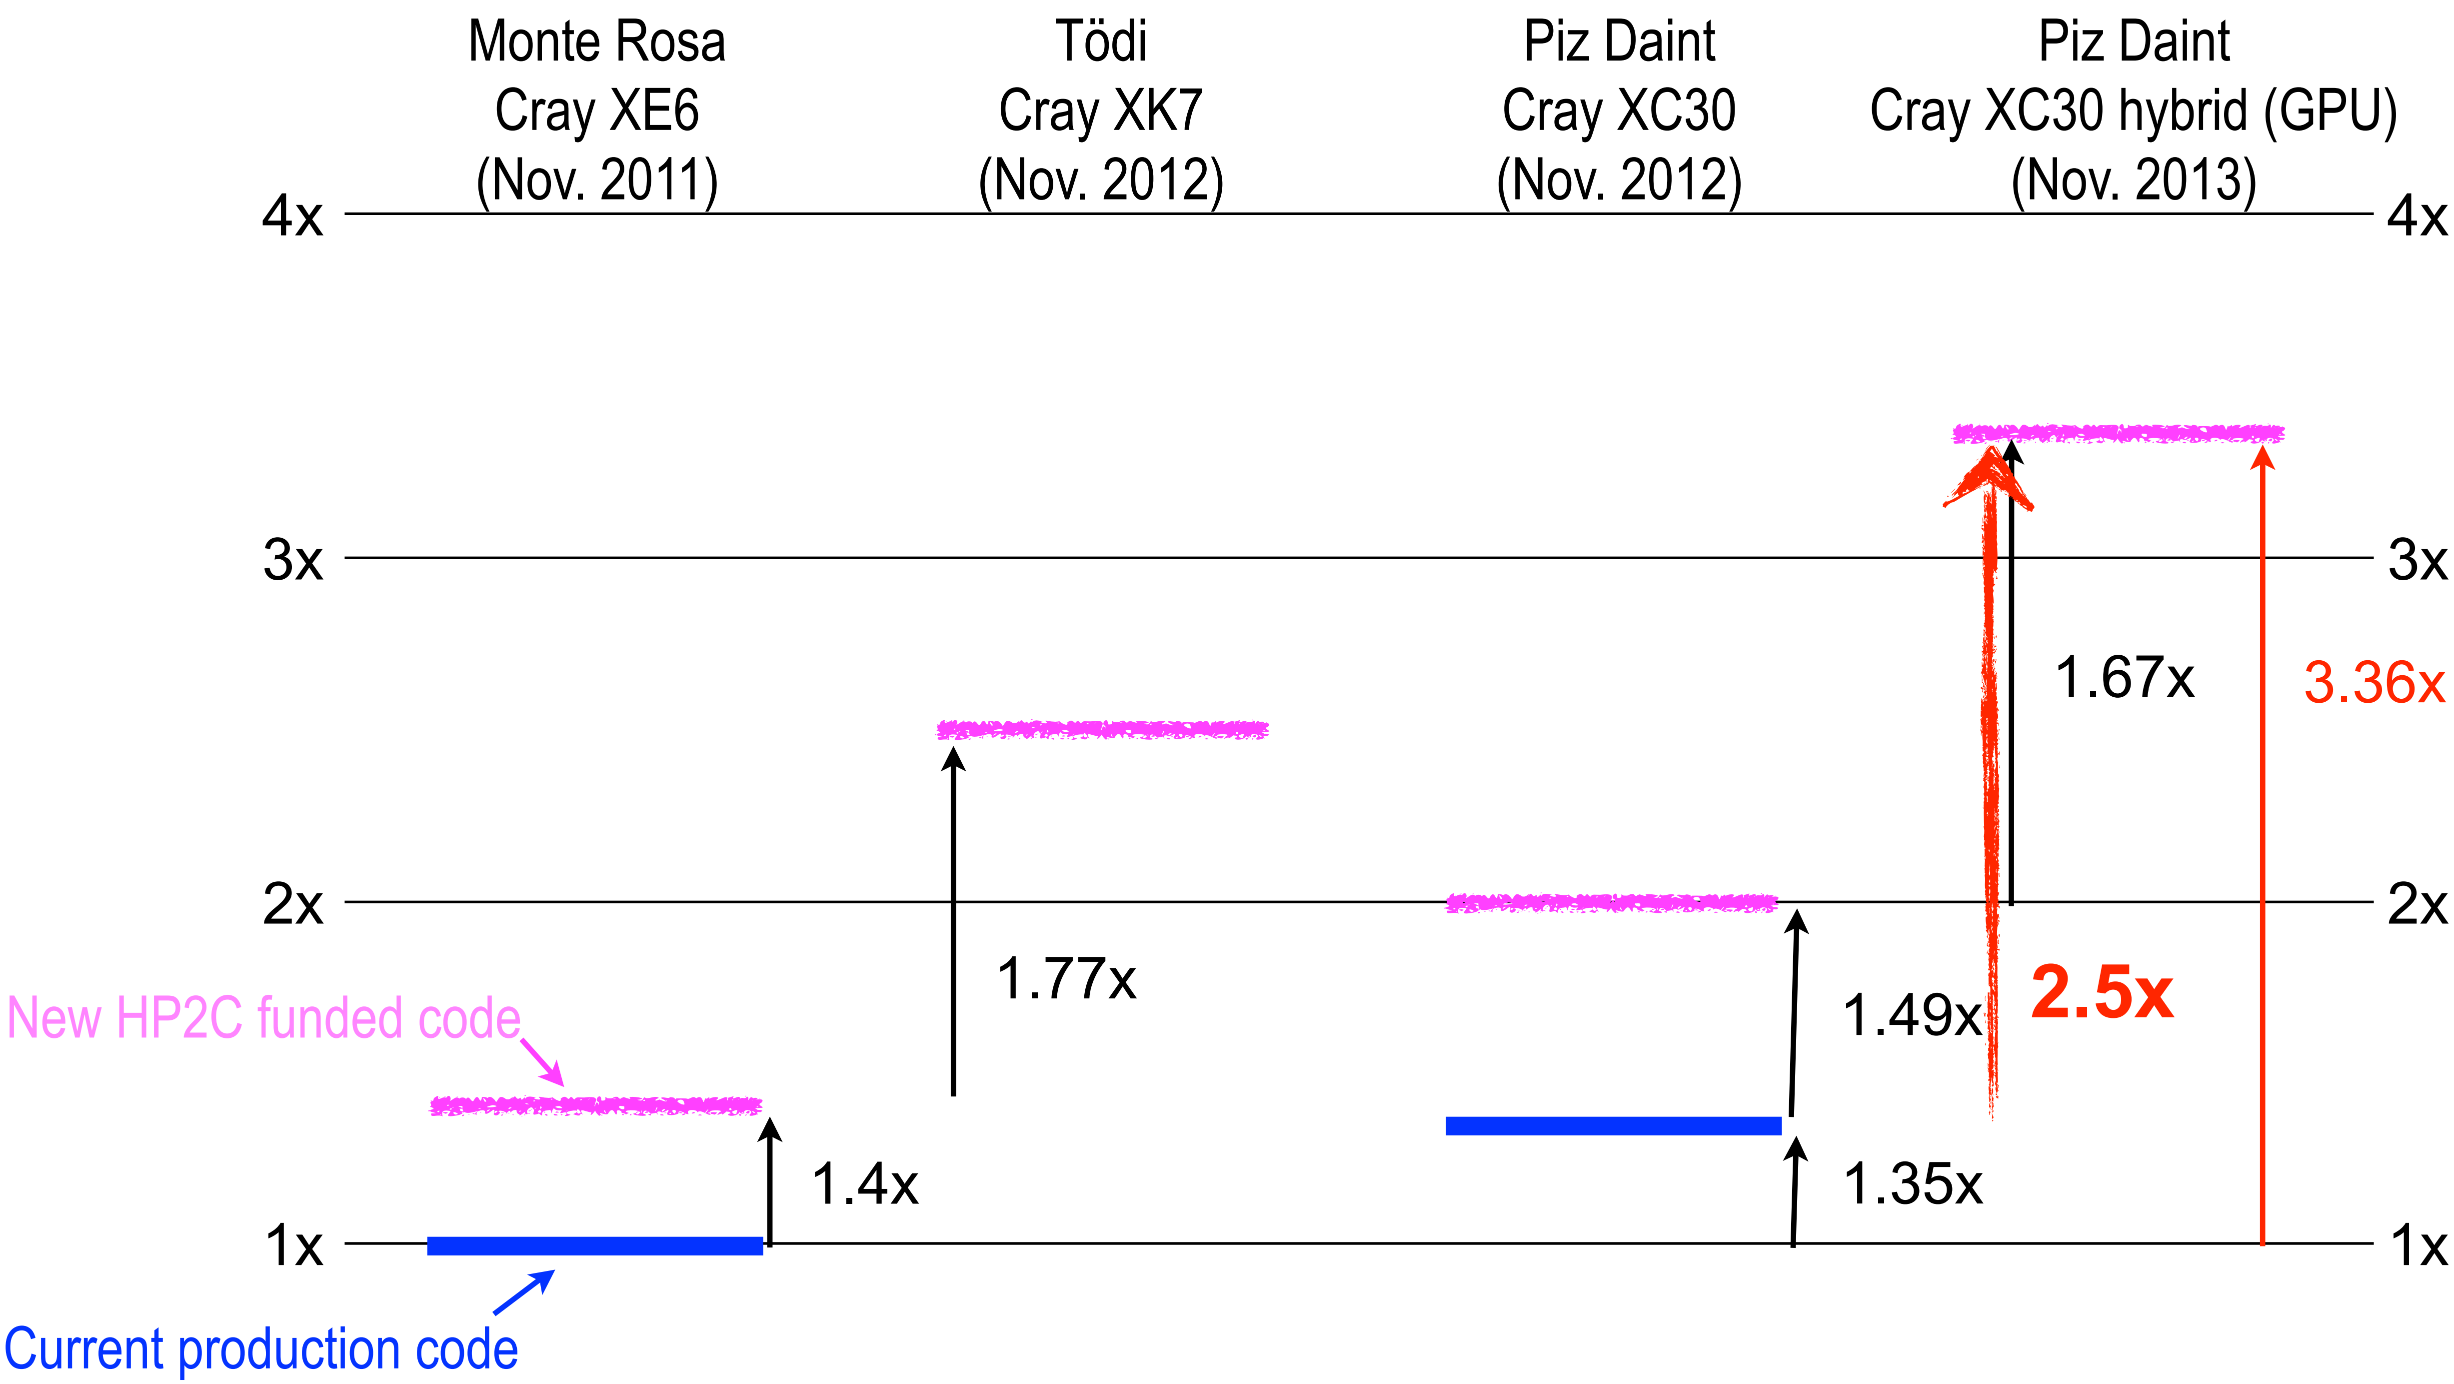

Energy to solution (kWh / ensemble member)

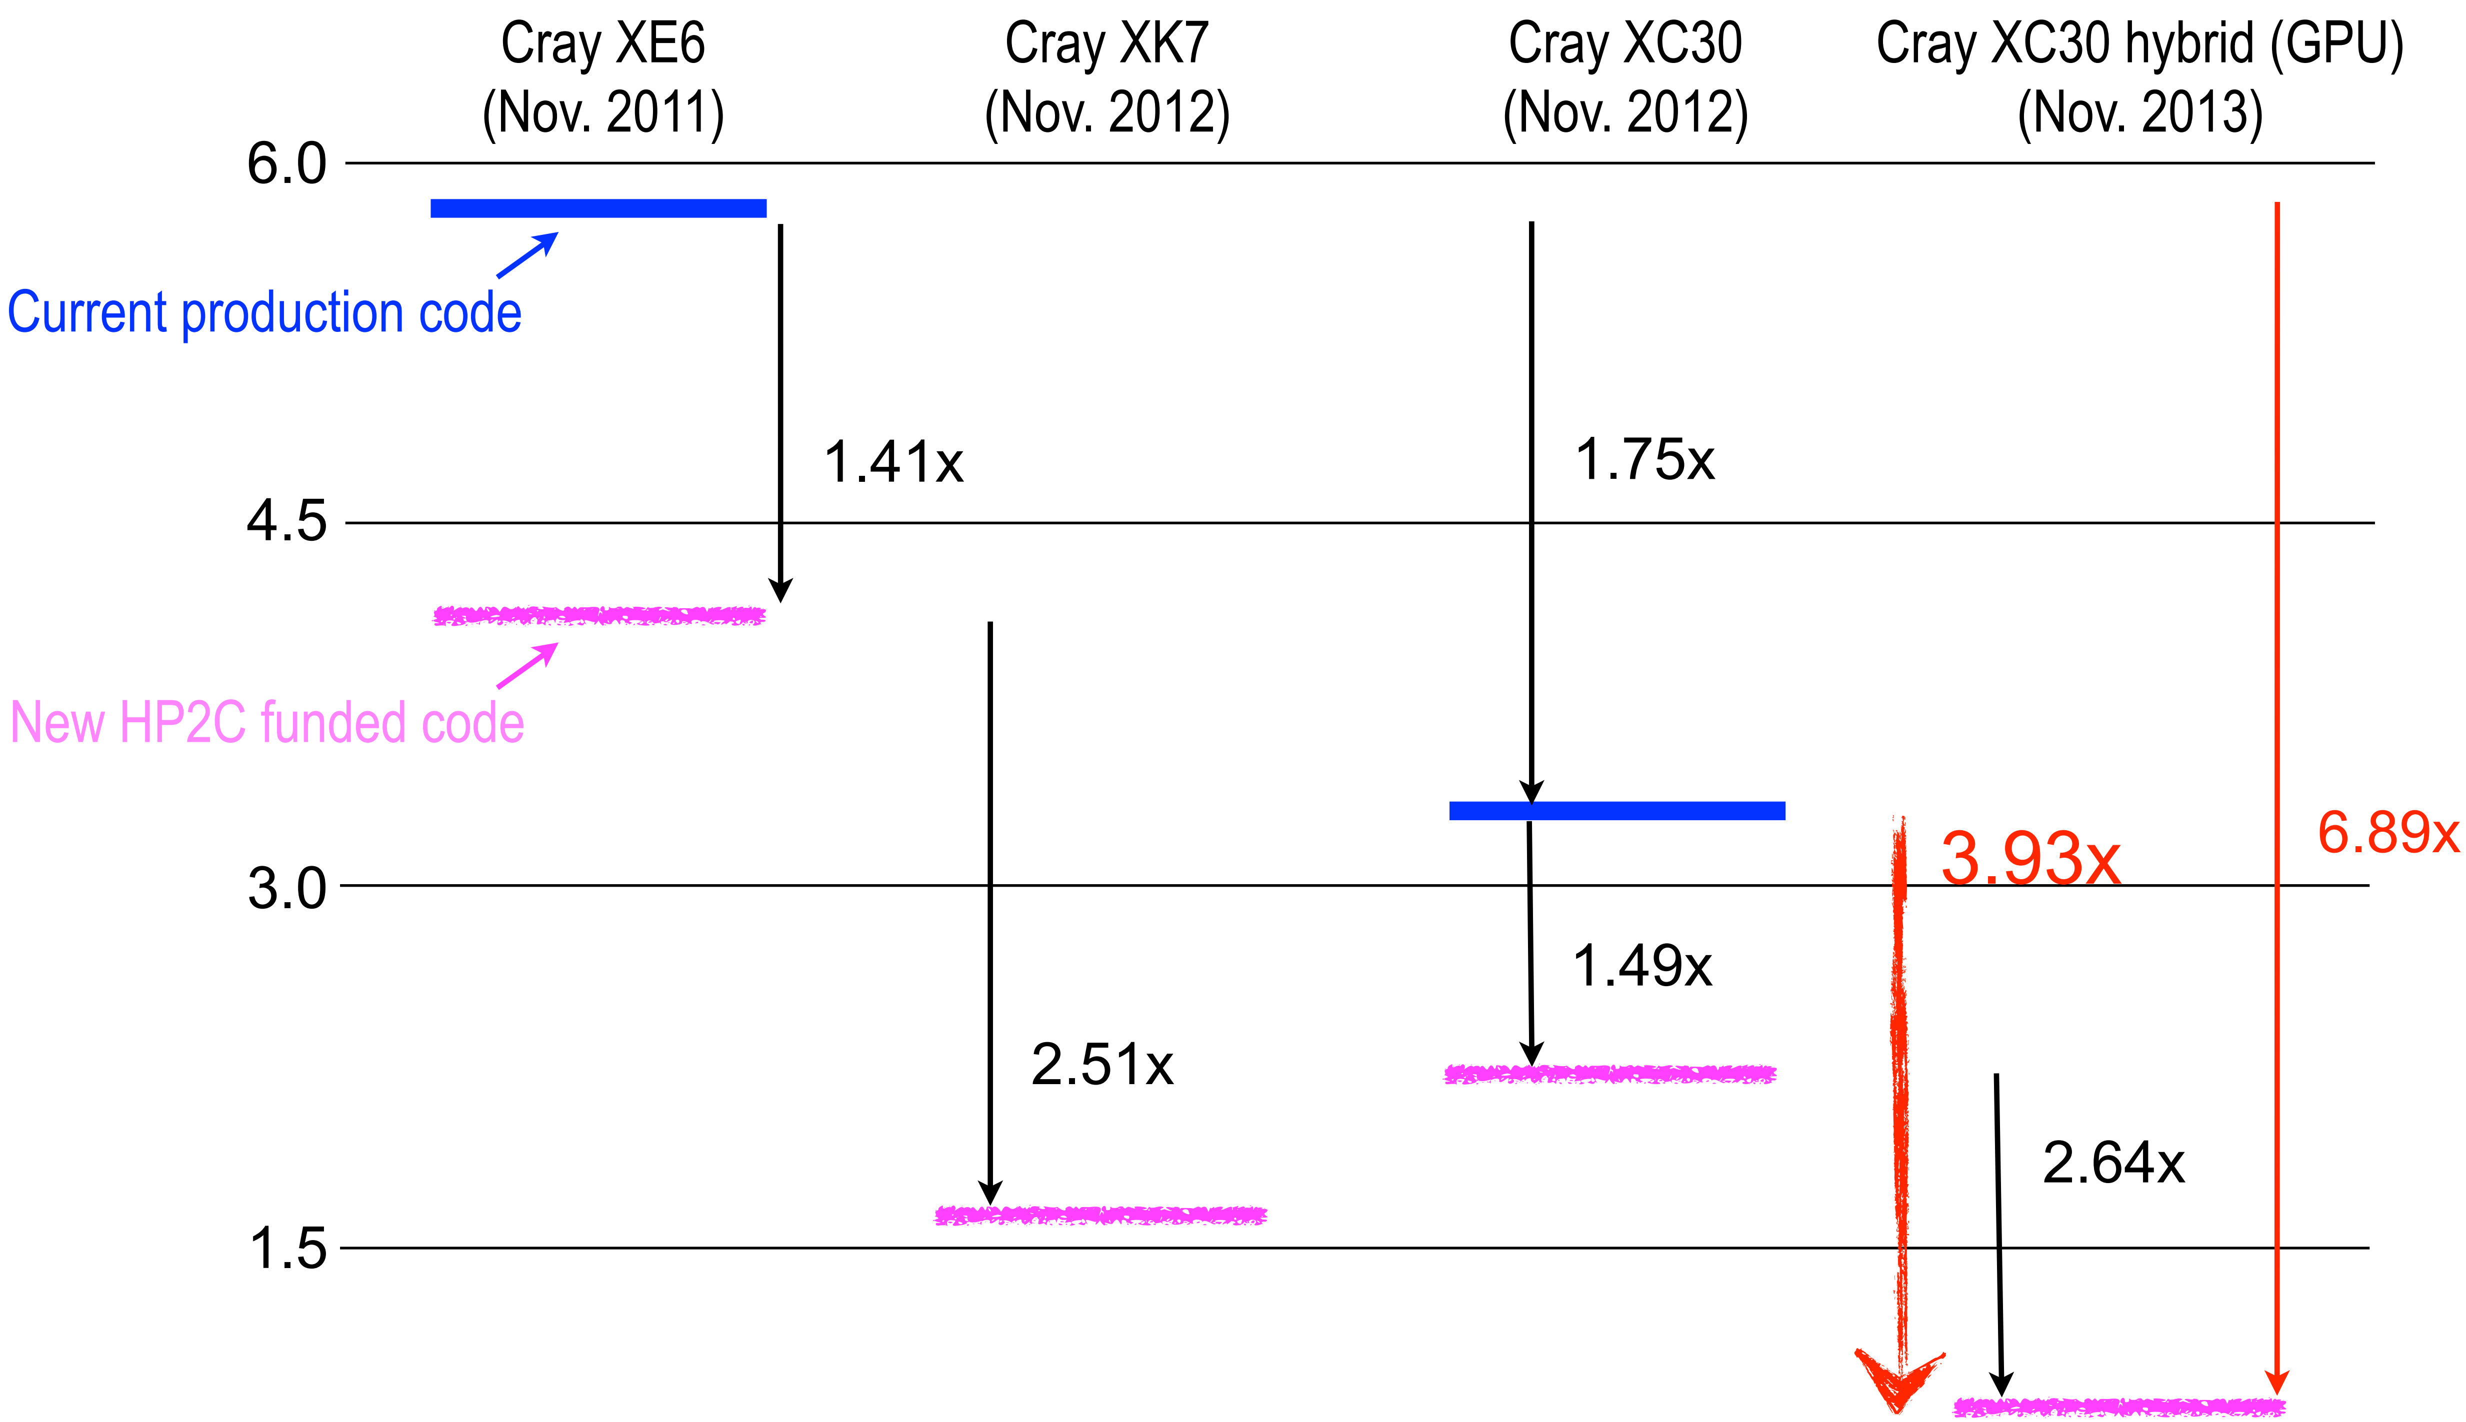

COSMO: current and new (HP2C developed) code

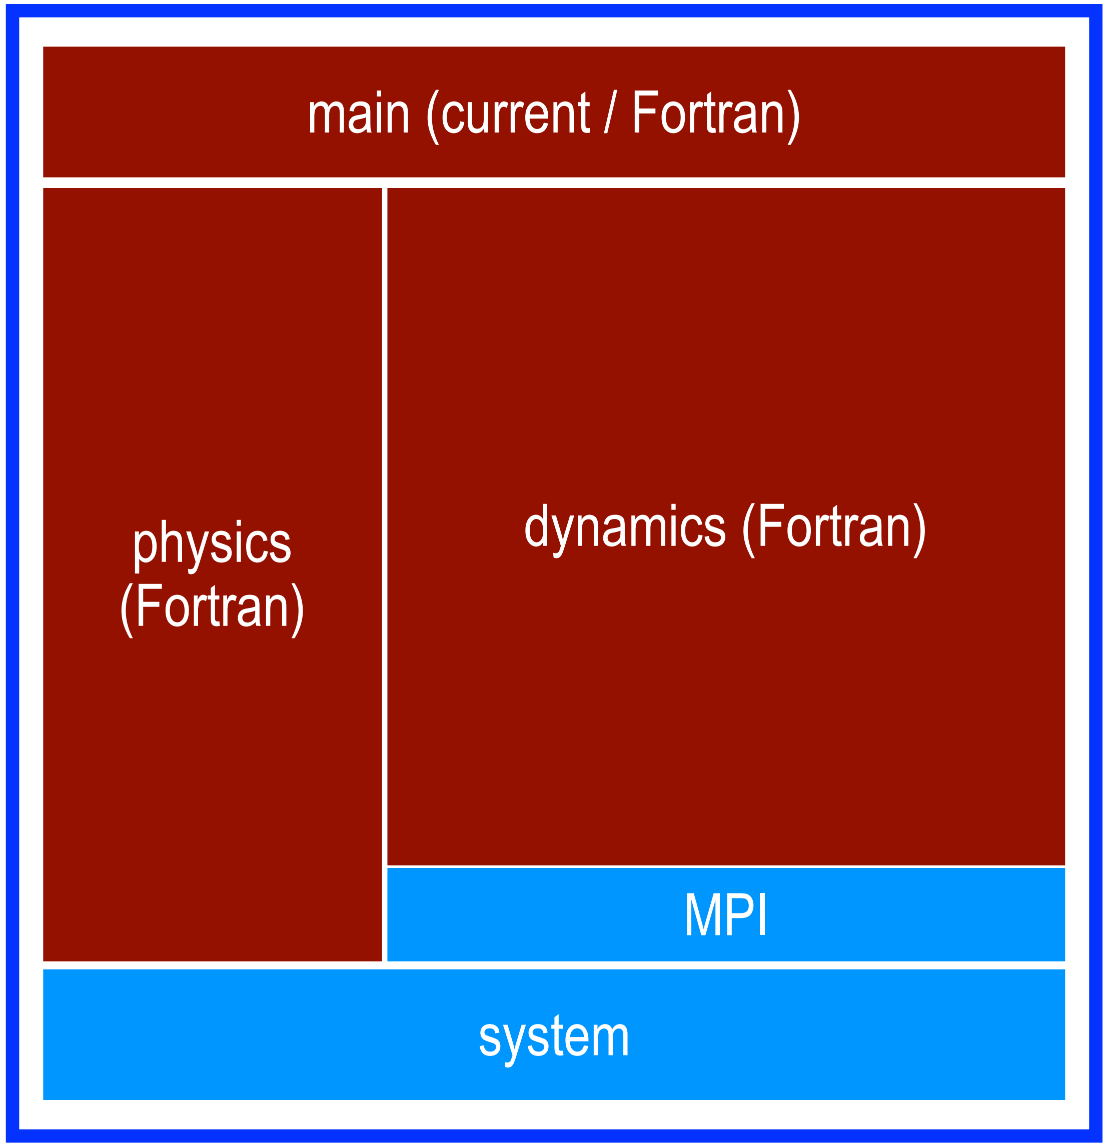

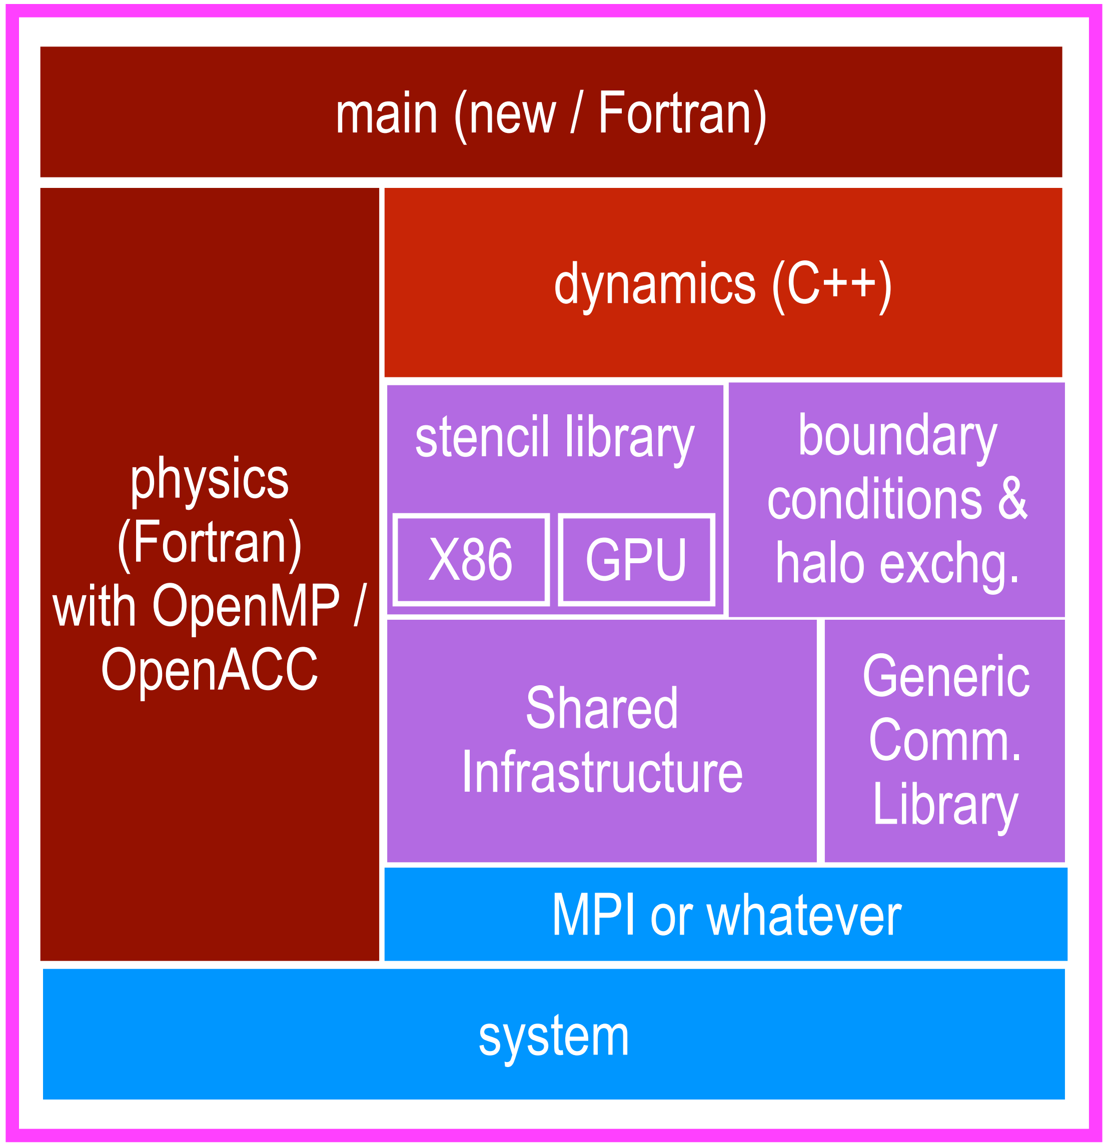

velocities

$$\frac{\partial u}{\partial t} = -\left\{\frac{1}{a\cos\varphi}\frac{\partial E_h}{\partial\lambda} - vV_a\right\} - \dot{\zeta}\frac{\partial u}{\partial\zeta} - \frac{1}{\rho a\cos\varphi}\left(\frac{\partial p'}{\partial\lambda} - \frac{1}{\sqrt{\gamma}}\frac{\partial p_0}{\partial\lambda}\frac{\partial p'}{\partial\zeta}\right) + M_u$$

$$\frac{\partial v}{\partial t} = -\left\{\frac{1}{a}\frac{\partial E_h}{\partial\varphi} + uV_a\right\} - \dot{\zeta}\frac{\partial v}{\partial\zeta} - \frac{1}{\rho a}\left(\frac{\partial p'}{\partial\varphi} - \frac{1}{\sqrt{\gamma}}\frac{\partial p_0}{\partial\varphi}\frac{\partial p'}{\partial\zeta}\right) + M_v$$

$$\frac{\partial w}{\partial t} = -\left\{\frac{1}{a\cos\varphi}\left(u\frac{\partial w}{\partial\lambda} + v\cos\varphi\frac{\partial w}{\partial\varphi}\right)\right\} - \dot{\zeta}\frac{\partial w}{\partial\zeta} + \frac{g}{\sqrt{\gamma}}\frac{\rho_0}{\rho}\frac{\partial p'}{\partial\zeta} + M_w + g\frac{\rho_0}{\rho}\left\{\frac{(T-T_0)}{T} - \frac{T_0 p'}{T p_0} + \left(\frac{R_v}{R_d} - 1\right)q^v - q^l - q^f\right\}$$

pressure

$$\frac{\partial p'}{\partial t} = -\left\{\frac{1}{a\cos\varphi}\left(u\frac{\partial p'}{\partial\lambda} + v\cos\varphi\frac{\partial p'}{\partial\varphi}\right)\right\} - \dot{\zeta}\frac{\partial p'}{\partial\zeta} + g\rho_0 w - \frac{c_{pd}}{c_{vd}}pD$$

temperature

$$\frac{\partial T}{\partial t} = -\left\{\frac{1}{a\cos\varphi}\left(u\frac{\partial T}{\partial\lambda} + v\cos\varphi\frac{\partial T}{\partial\varphi}\right)\right\} - \dot{\zeta}\frac{\partial T}{\partial\zeta} - \frac{1}{\rho c_{vd}}pD + Q_T$$

water

$$\frac{\partial q^v}{\partial t} = -\left\{\frac{1}{a\cos\varphi}\left(u\frac{\partial q^v}{\partial\lambda} + v\cos\varphi\frac{\partial q^v}{\partial\varphi}\right)\right\} - \dot{\zeta}\frac{\partial q^v}{\partial\zeta} - (S^l + S^f) + M_{q^v}$$

$$\frac{\partial q^{l,f}}{\partial t} = -\left\{\frac{1}{a\cos\varphi}\left(u\frac{\partial q^{l,f}}{\partial\lambda} + v\cos\varphi\frac{\partial q^{l,f}}{\partial\varphi}\right)\right\} - \dot{\zeta}\frac{\partial q^{l,f}}{\partial\zeta} - \frac{g}{\sqrt{\gamma}}\frac{\rho_0}{\rho}\frac{\partial P_{l,f}}{\partial\zeta} + S^{l,f} + M_{q^{l,f}}$$

turbulence

$$\frac{\partial e_t}{\partial t} = -\left\{\frac{1}{a\cos\varphi}\left(u\frac{\partial e_t}{\partial\lambda} + v\cos\varphi\frac{\partial e_t}{\partial\varphi}\right)\right\} - \dot{\zeta}\frac{\partial e_t}{\partial\zeta} + K_m^v\frac{g\rho_0}{\sqrt{\gamma}}\left\{\left(\frac{\partial u}{\partial\zeta}\right)^2 + \left(\frac{\partial v}{\partial\zeta}\right)^2\right\} + \frac{g}{\rho\theta_v}F^{\theta_v} - \frac{\sqrt{2}e_t^{3/2}}{\alpha_M l} + M_{e_t}$$

Mathematical description

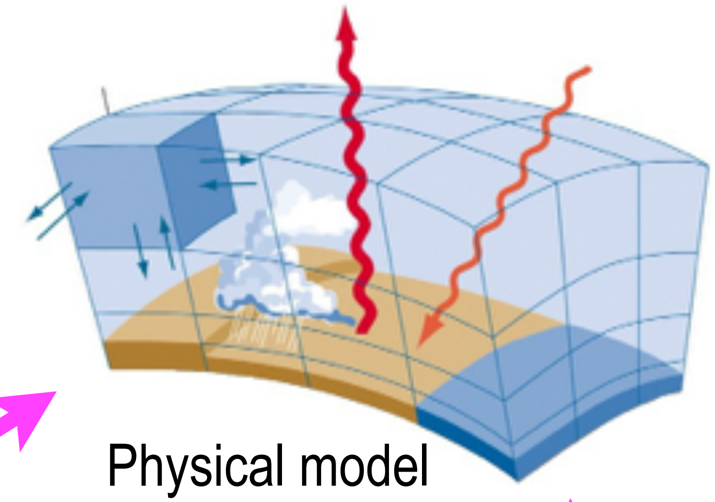

Physical model

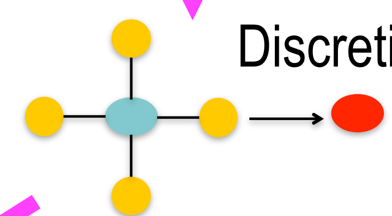

Discretization / algorithm

Tools & Libraries

Domain science (incl. applied mathematics)

```
lap(i,j,k) = -4.0 * data(i,j,k) +
    data(i+1,j,k) + data(i-1,j,k) +
    data(i,j+1,k) + data(i,j-1,k);
```

Code / implementation

Optimal algorithm
Auto tuning

Code compilation

Architectural options / design

Computer engineering (& computer science)

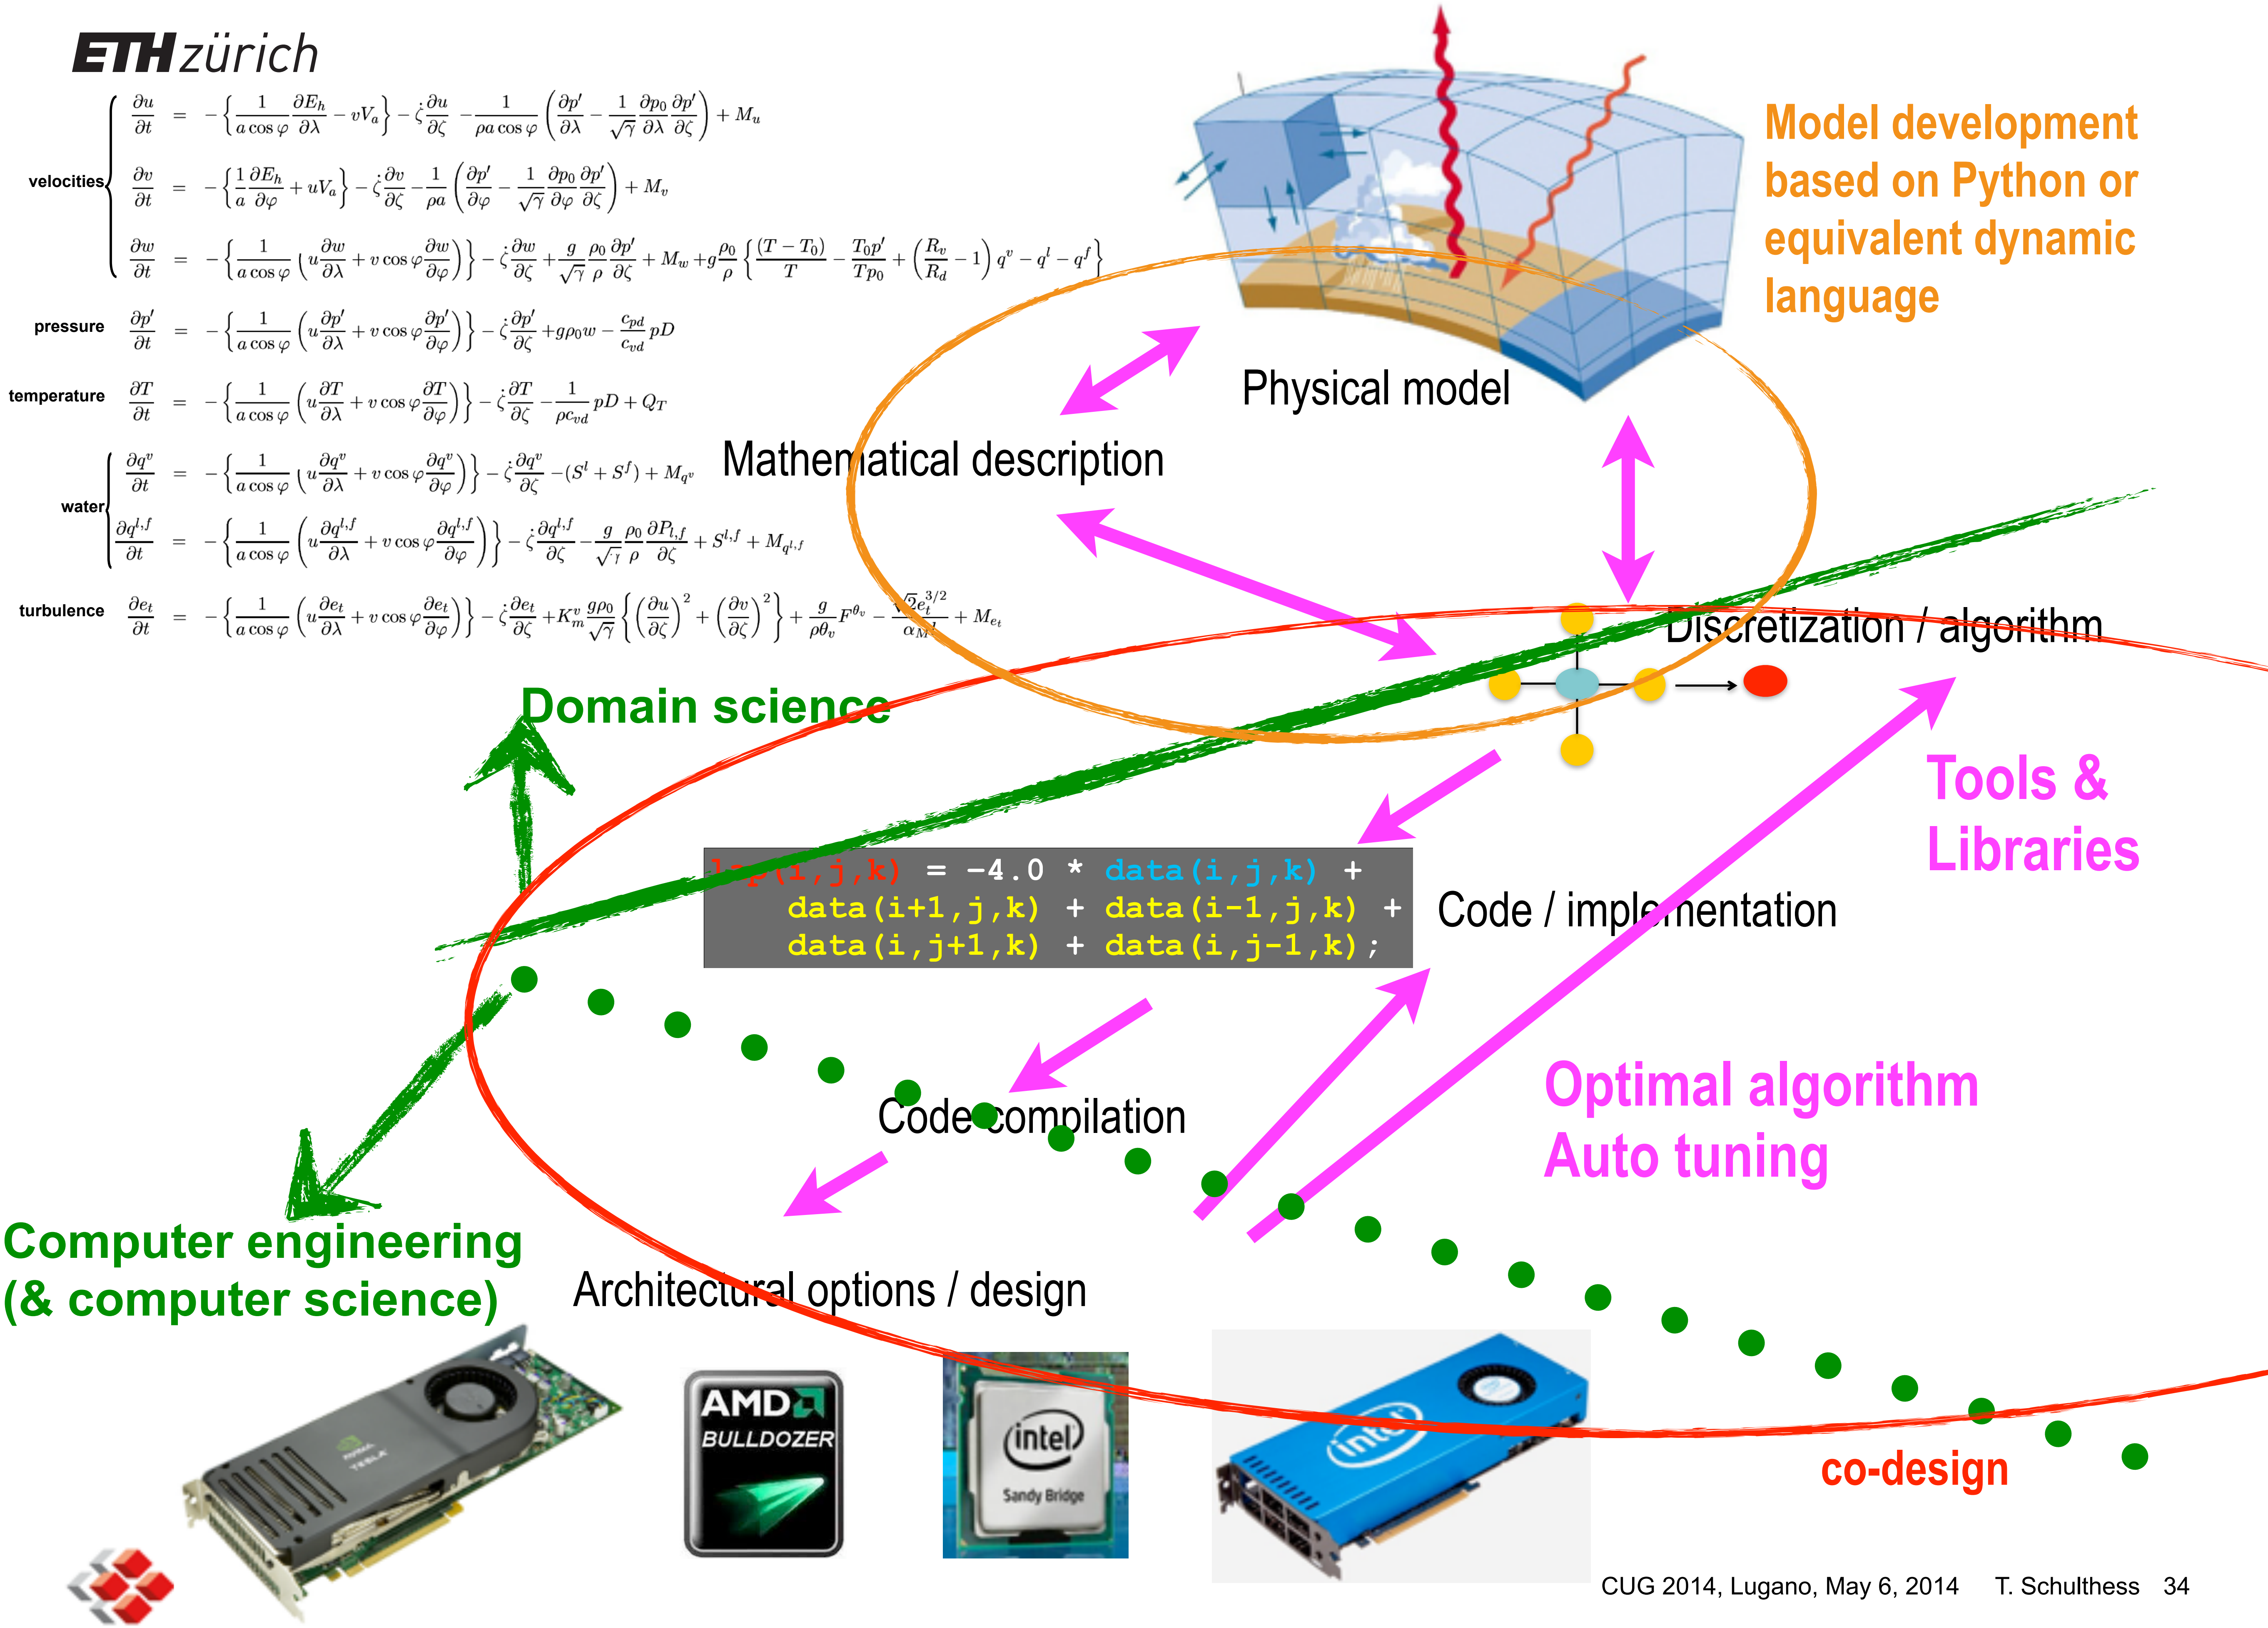

ETH zürich
velocities
pressure
temperature
water
turbulence
Model development based on Python or equivalent dynamic language
Physical model
Mathematical description
Discretization / algorithm
Domain science
Tools & Libraries
lap(i,j,k) = -4.0 * data(i,j,k) +
data(i+1,j,k) + data(i-1,j,k) +
data(i,j+1,k) + data(i,j-1,k);
Code / implementation
Code compilation
Optimal algorithm
Auto tuning
Computer engineering
(& computer science)
Architectural options / design
co-design
AMD
BULLDOZER
intel
Sandy Bridge

ETH zürich

COSMO in five year: current and new (2019) code

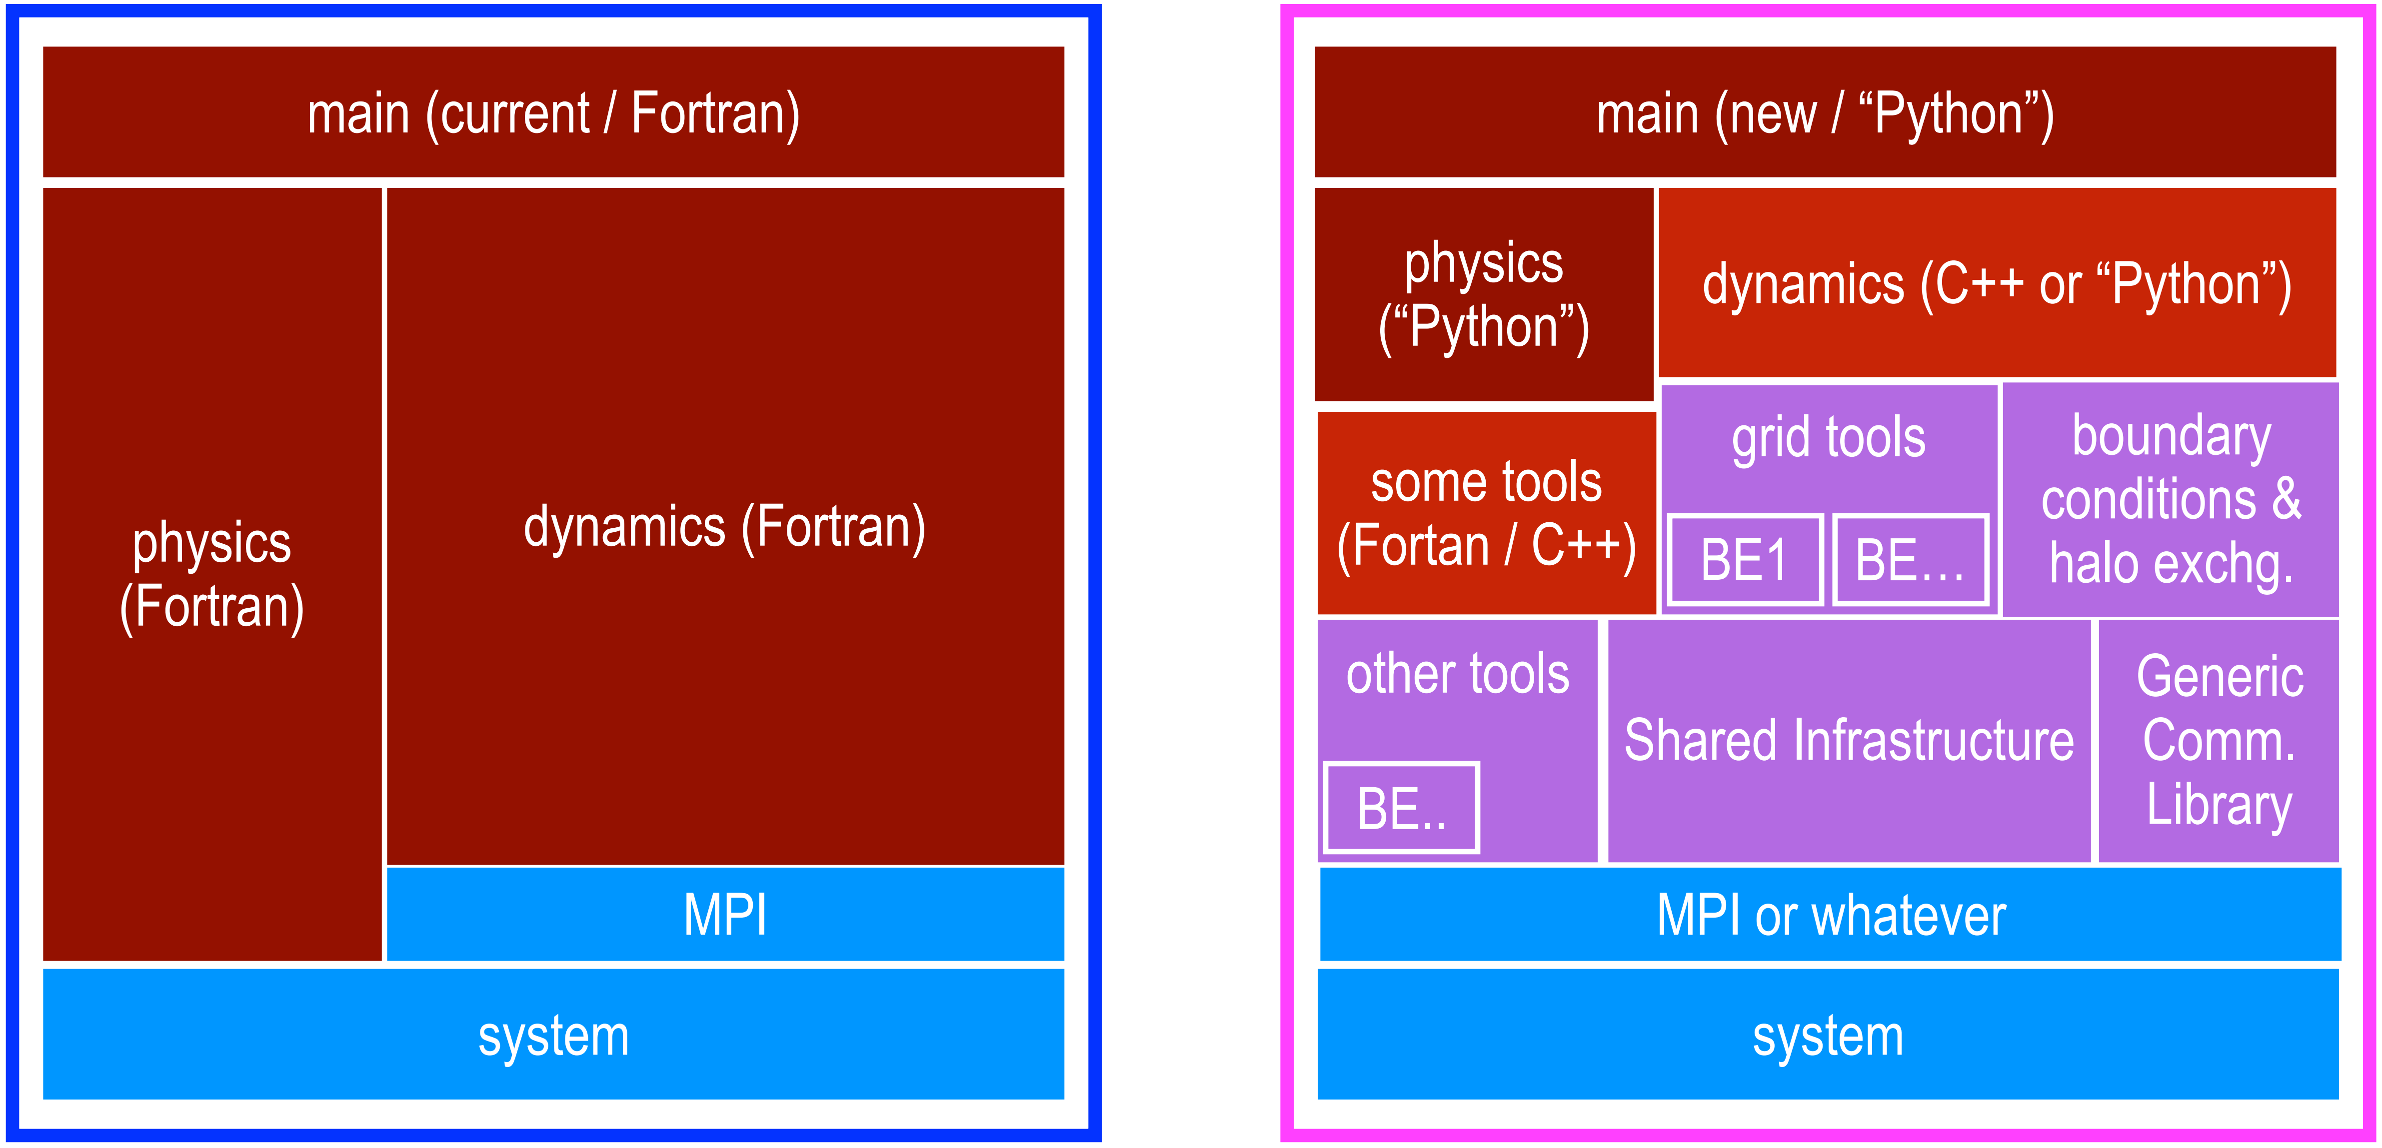

“Software is getting slower more rapidly than hardware becomes faster”

–Niklaus Wirth, 1995

ETH zürich

Thank You!