ETH zürich

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

Thomas C. Schulthess

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

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!

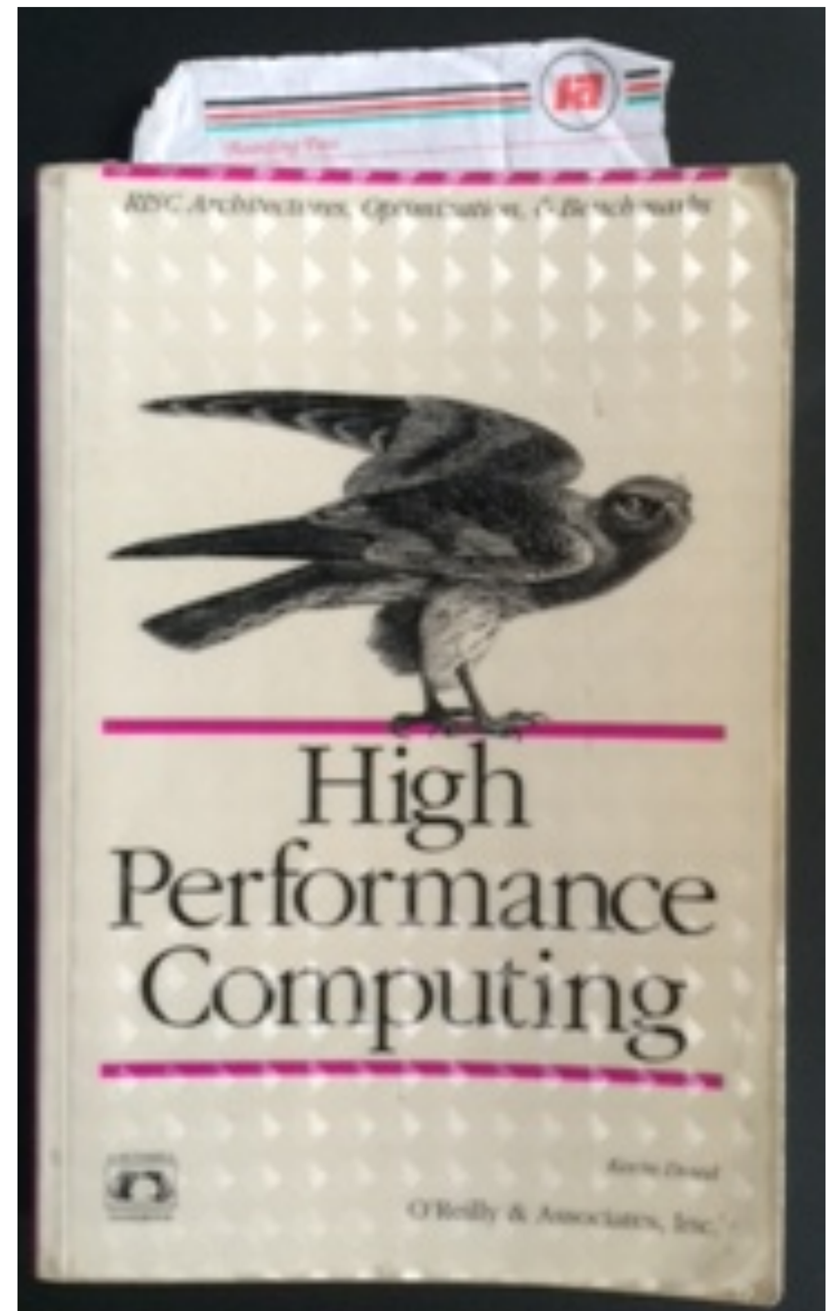

Incremental changes …

- Evolution of Fortran
- Evolution of C++ – you can mention C++ at HPC conferences 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?

Maybe it is time to turn away from HPC, to Big Data?

Computer scientists tell us:

data volume doubles every 2 years! | analysis of unstructured data | performance really matters!

Big Data = Volume, Variety, Velocity, Veracity

In science, are you kidding?

Many Big Data applications in science:
> **Life science**: biological and genomic data, personalised health
> **Environmental science**: remote and in-situ sensing, simulations
> **Social sciences**: digital humanities, social media, economics & finance
> **Computer science**: search, data integration, unstructured data analytics

"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

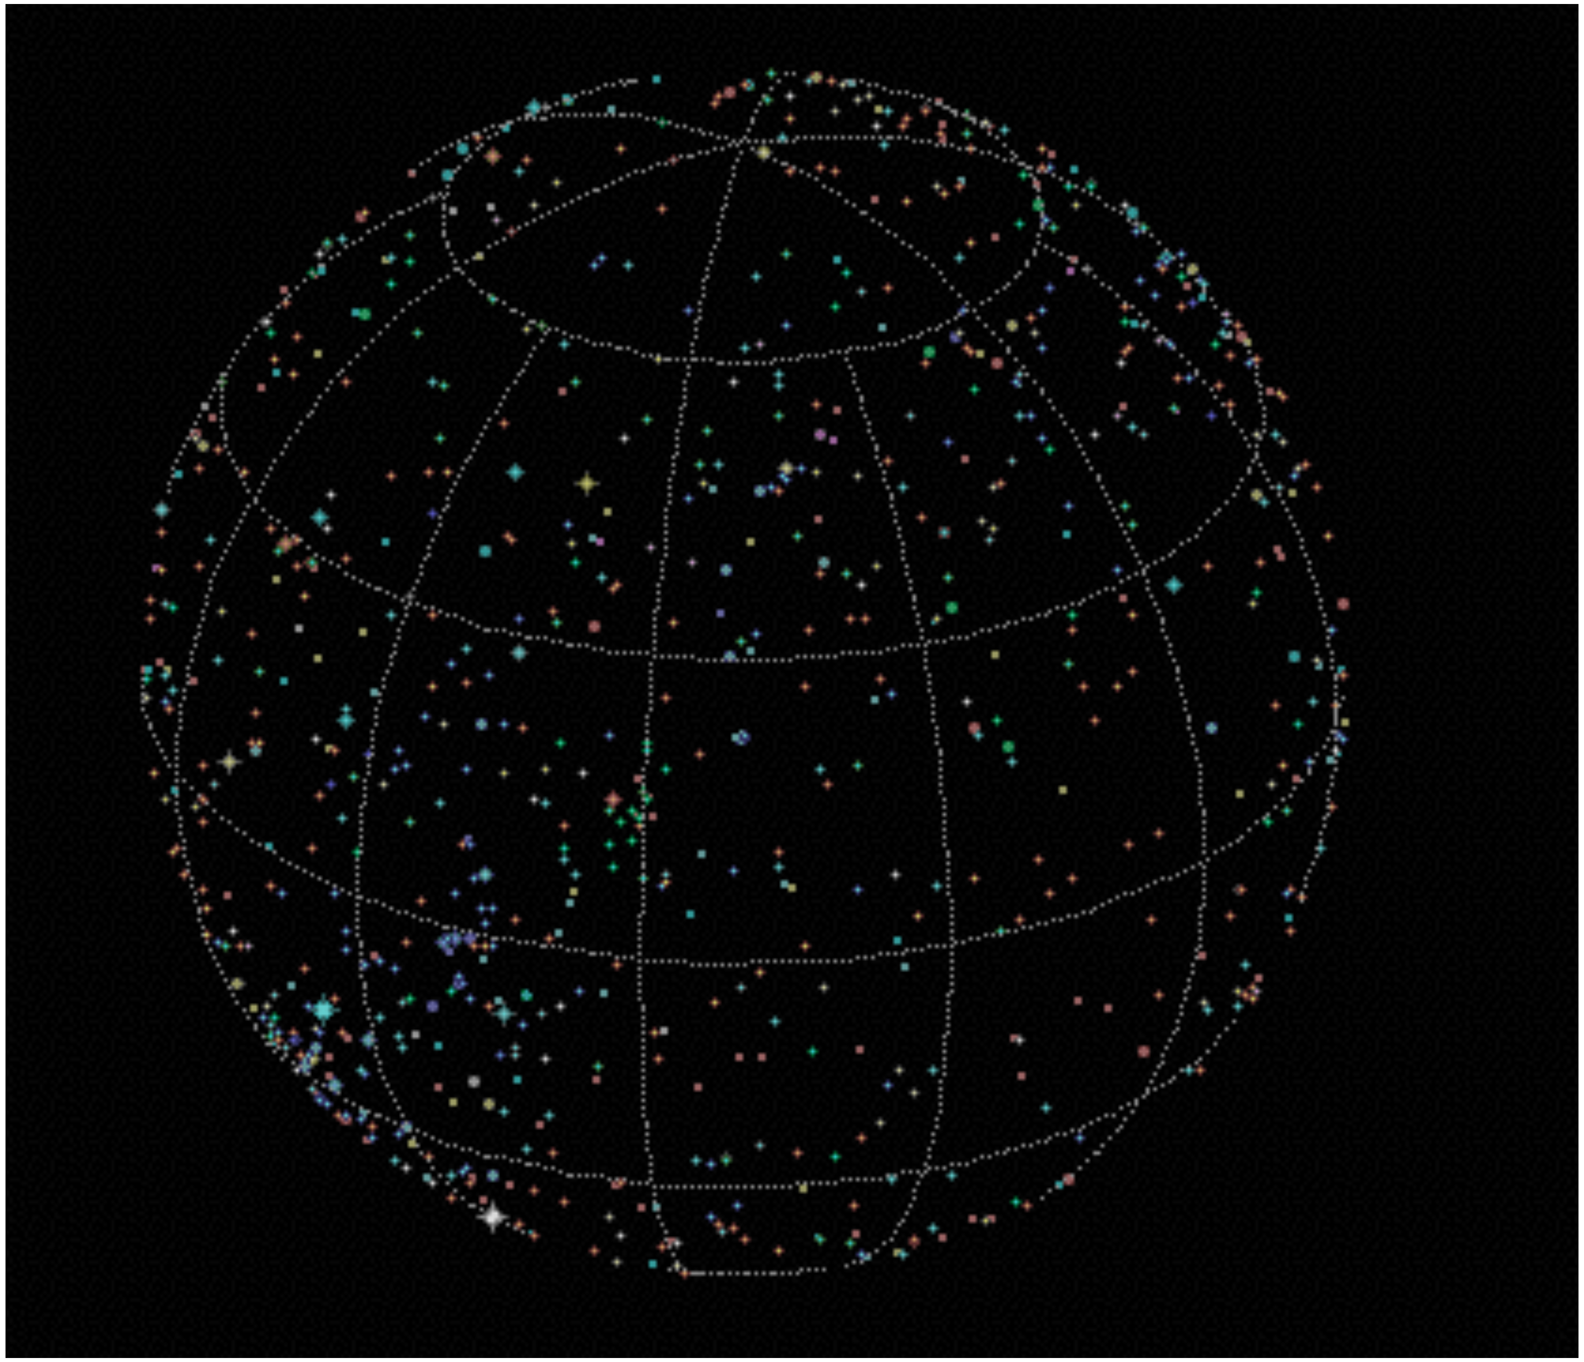

source: www.pafko.com/tycho/

Data given to Johannes Kepler

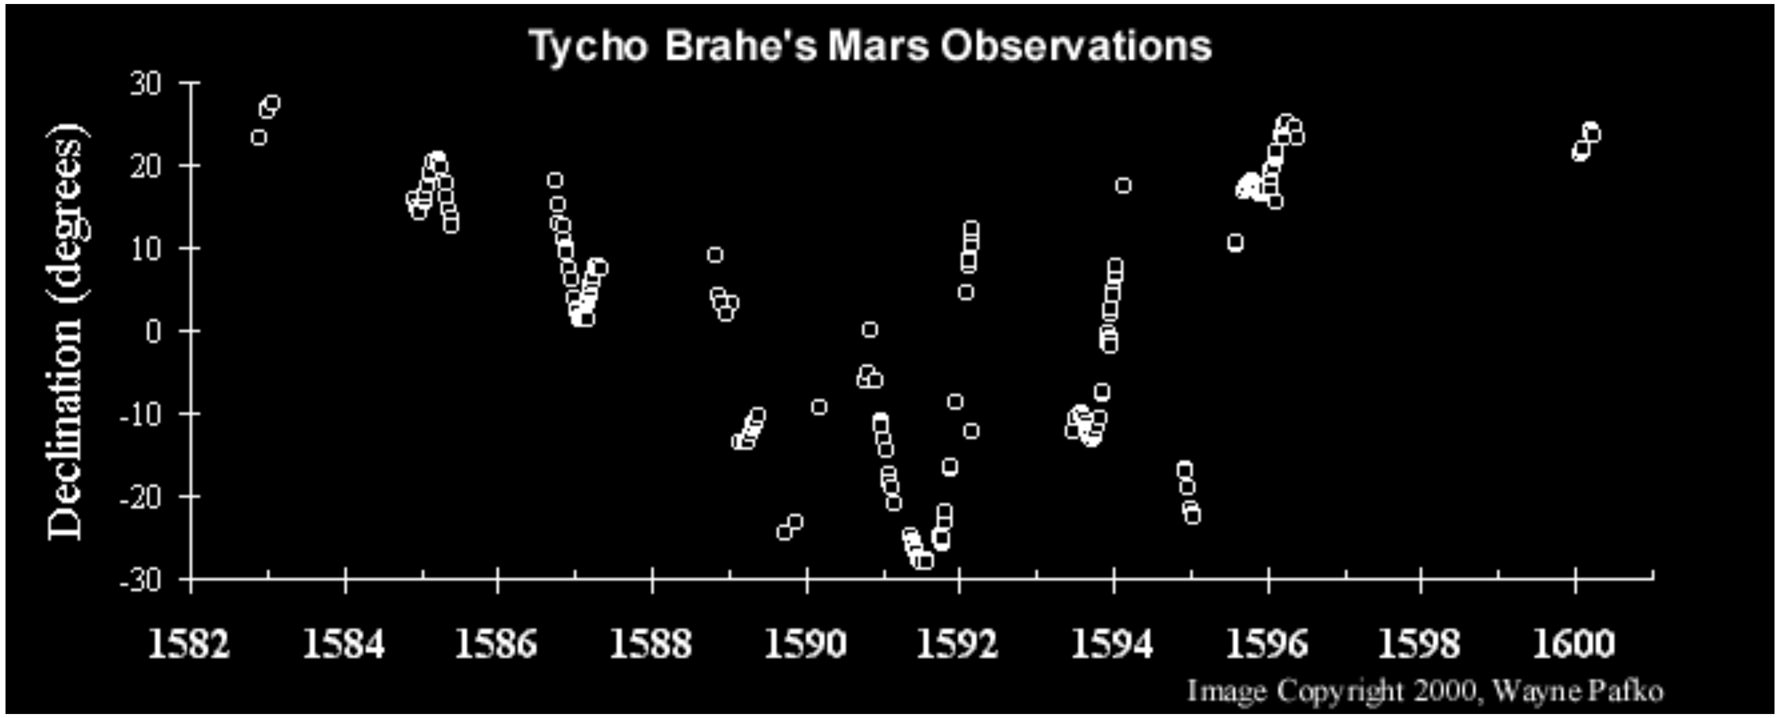

source: www.pafko.com/tycho/

Kepler’s modelling and simulations

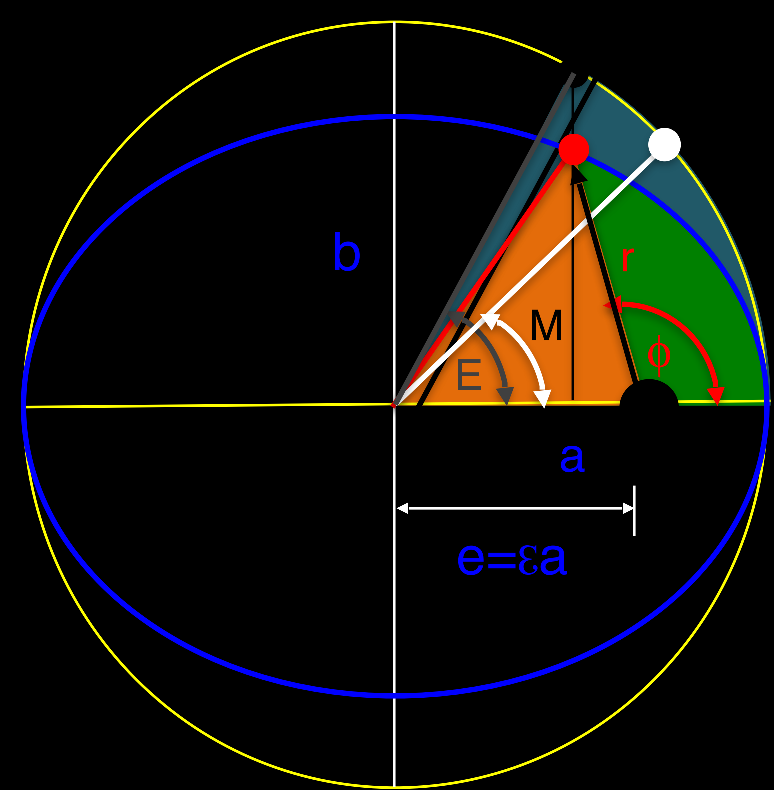

$$M = E - \varepsilon \sin E$$

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

After Kepler's analysis and Newton's theory …

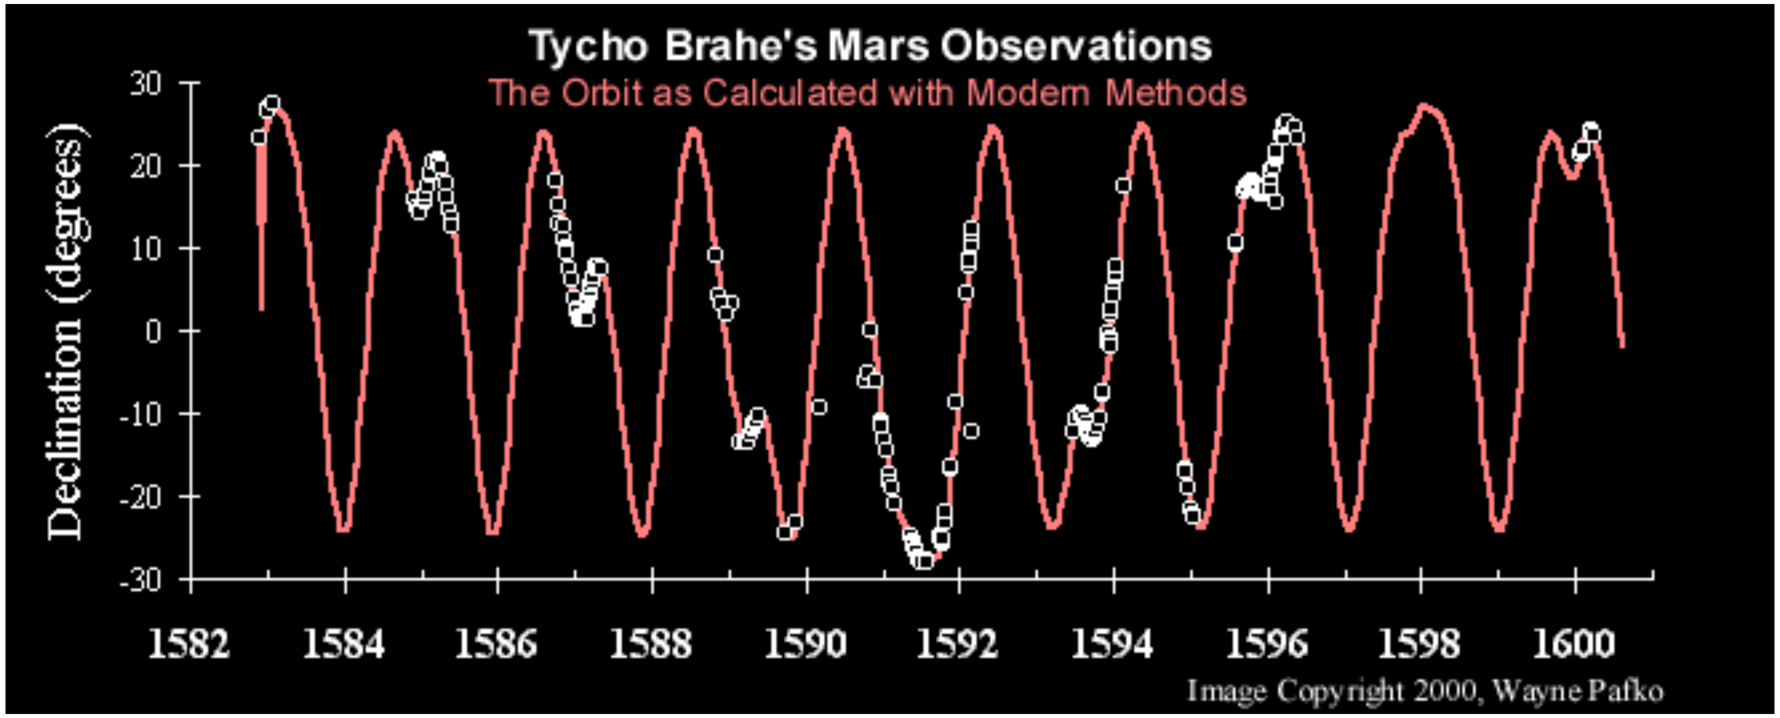

source: www.pafko.com/tycho/

"Piz Daint", CSCS ' new flagship system and one of Europe's most powerful petascale supercomputers

Presently the world's most energy efficient petascale supercomputer!

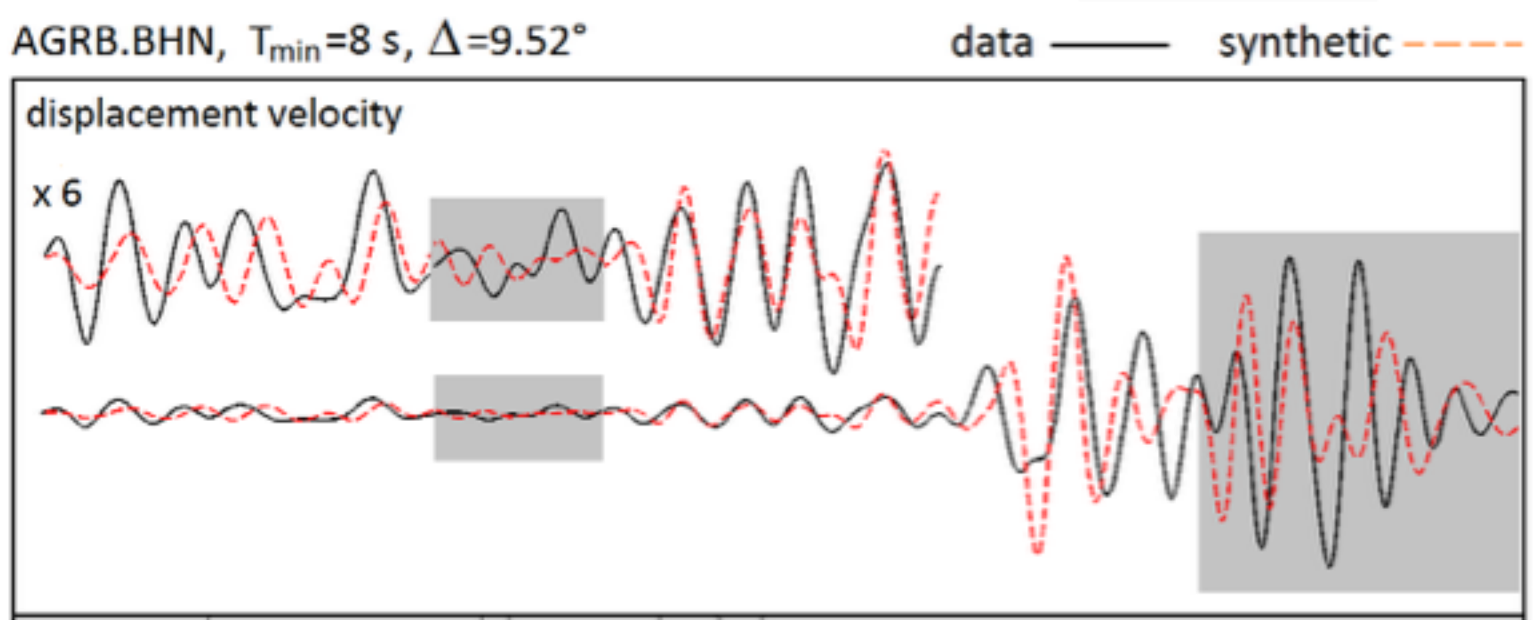

source: A. Fichtner, ETH Zurich

Data from many stations and earthquakes

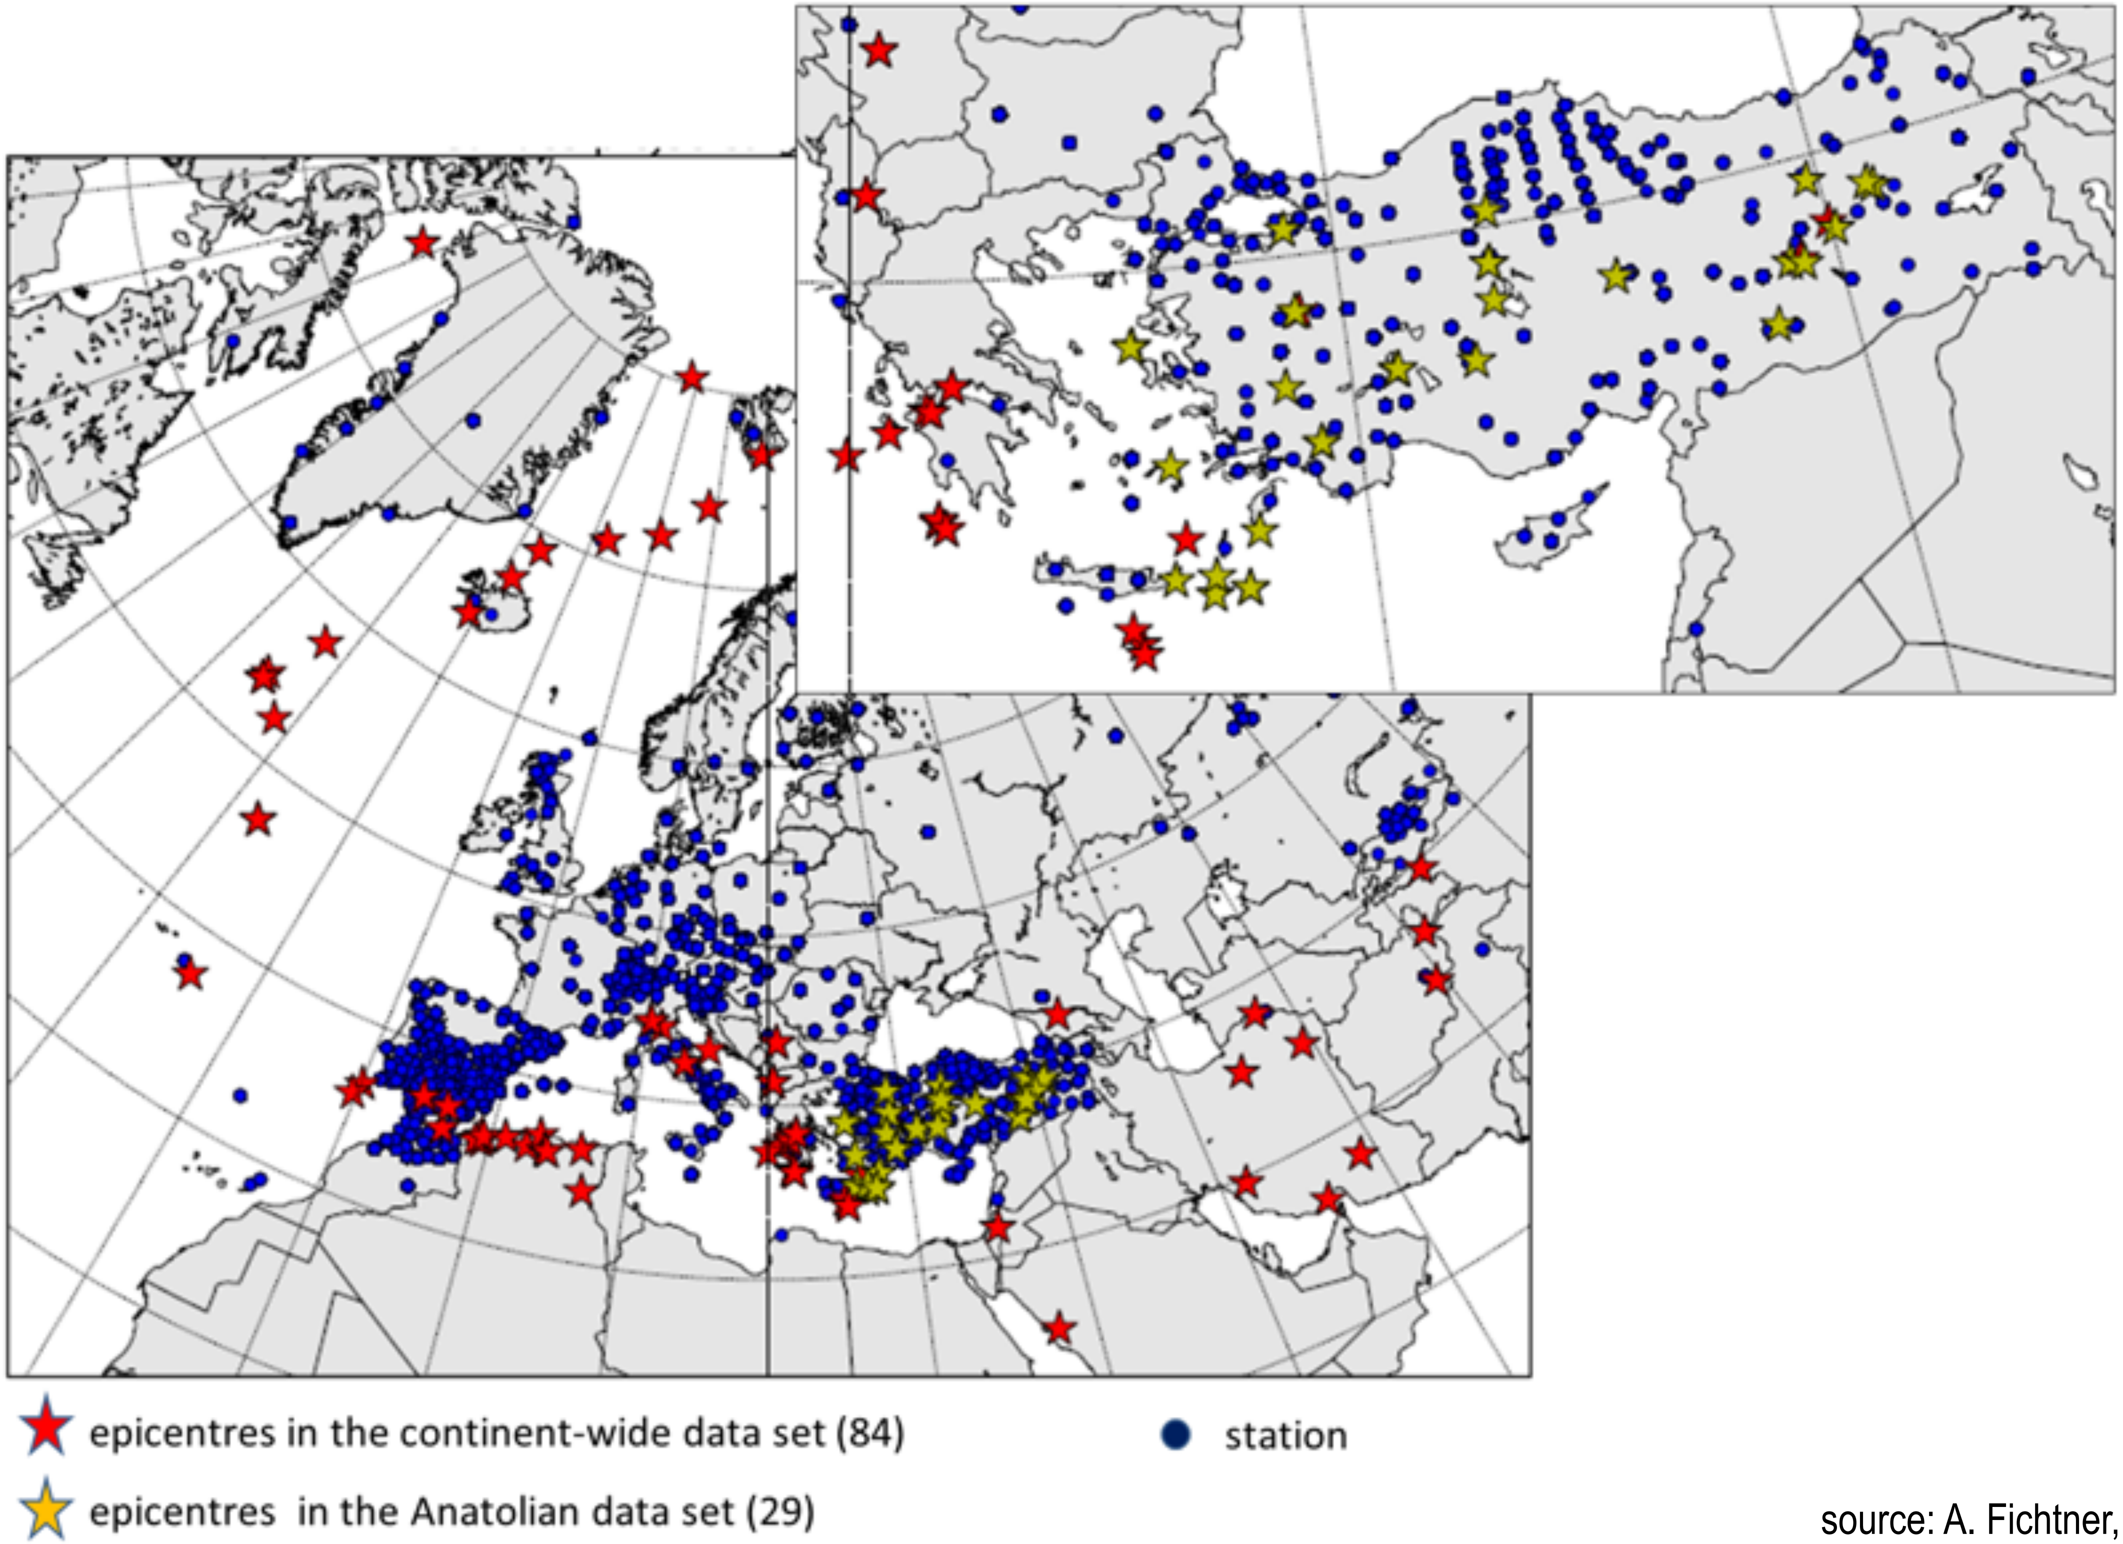

source: A. Fichtner, ETH Zurich

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

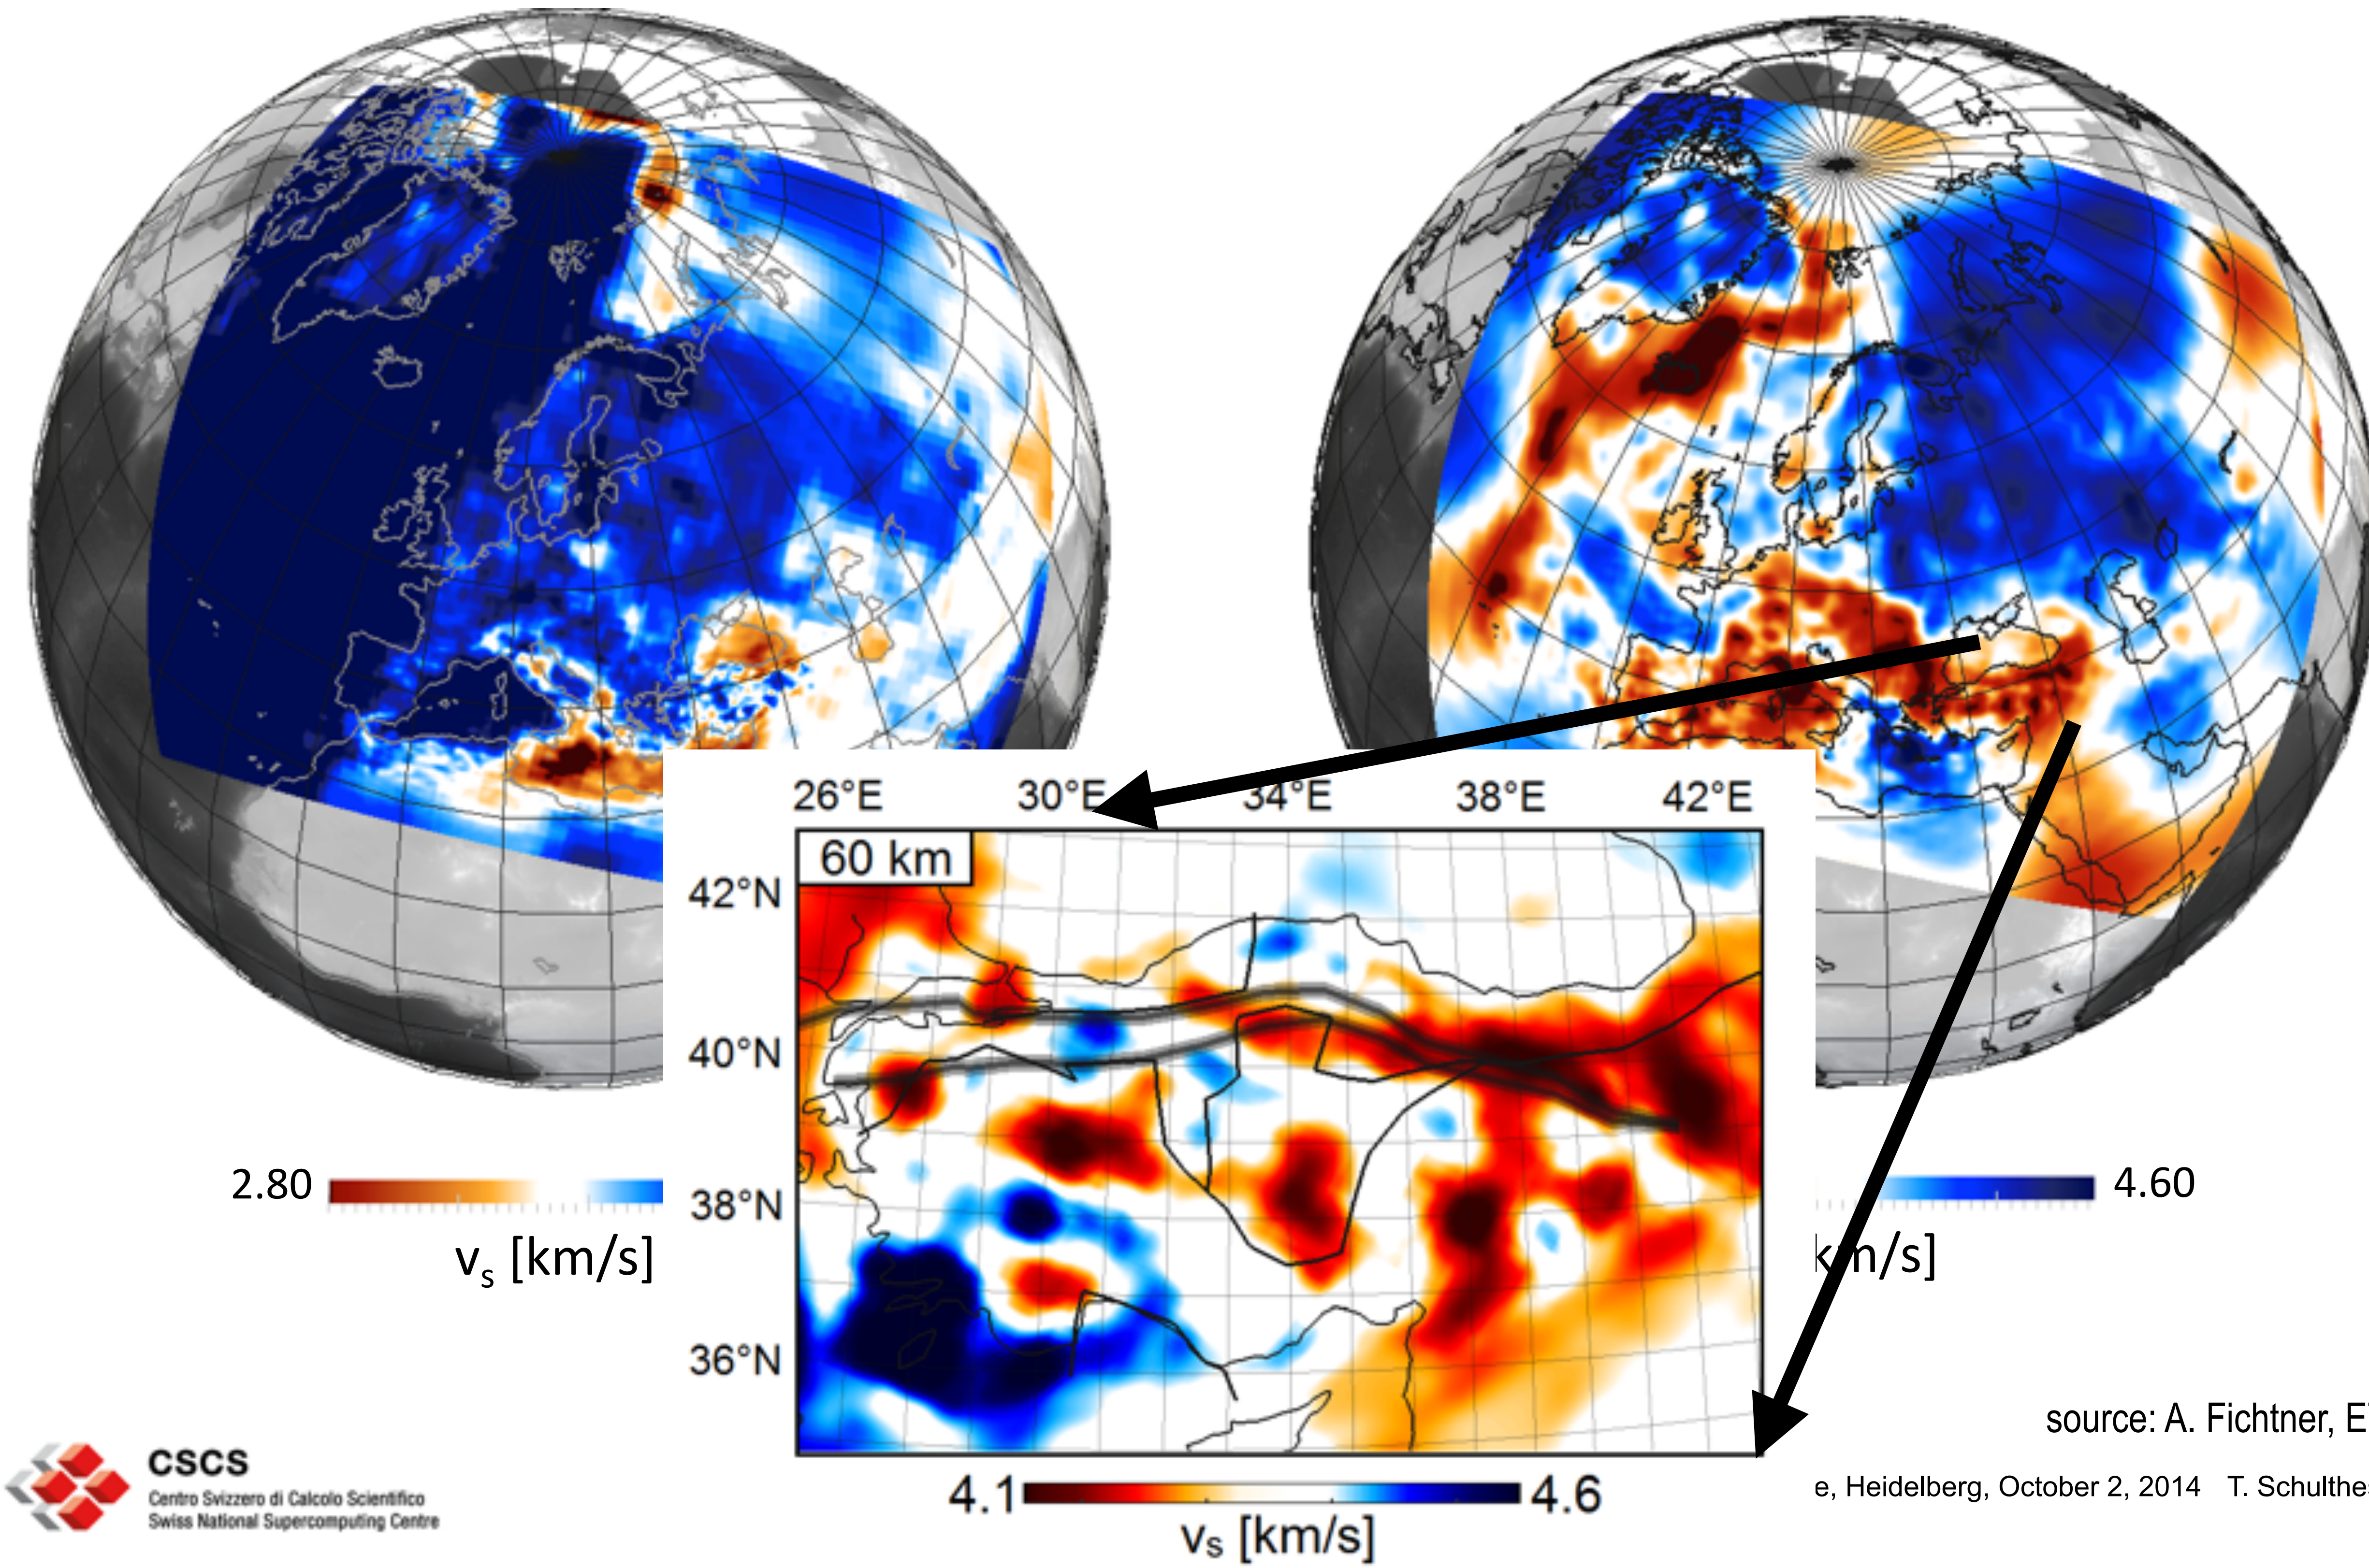

source: A. Fichtner, ETH Zurich

Pillars (paradigms?) 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 (paradigms?) 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

ETH zürich

The performance metric we should care about are

Energy & Time

How we should optimize

Time to solution (TTS):

- do we have to minimise time to solution?
- no, it just needs to be good enough to meet operational constrains

Energy to solution (ETS):

- energy is directly proportional to cost (power = energy / time)
- given all operational constraints, energy should be minimised

Cloud-Resolving Simulation of Winter Storm Kyrill

David Leutwyler, Oliver Fuhrer, Christoph Schär, Andrea Arteaga, Isabelle Bey, Mauro Bianco, Ben Cumming, Tobias Gysi, Xavier Lapillonne, Daniel Lüthi, Carlos Osuna, Anne Roches, Thomas Schulthess

COSMO-2 in production at Meteo Swiss

- Problem size (COSMO-2): 2km resolution on 540x314x60 grid points
- Time to solution: 33h forecast must run in 39 minutes
- Energy to solution: minimise (since always proportional to cost)

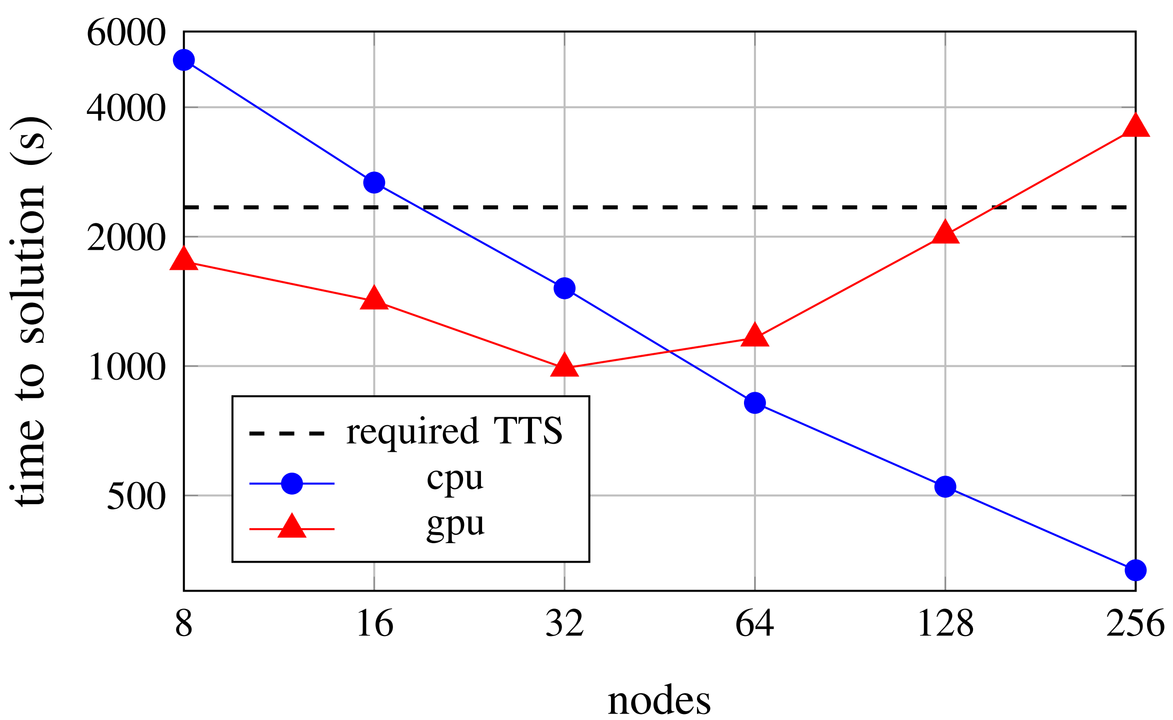

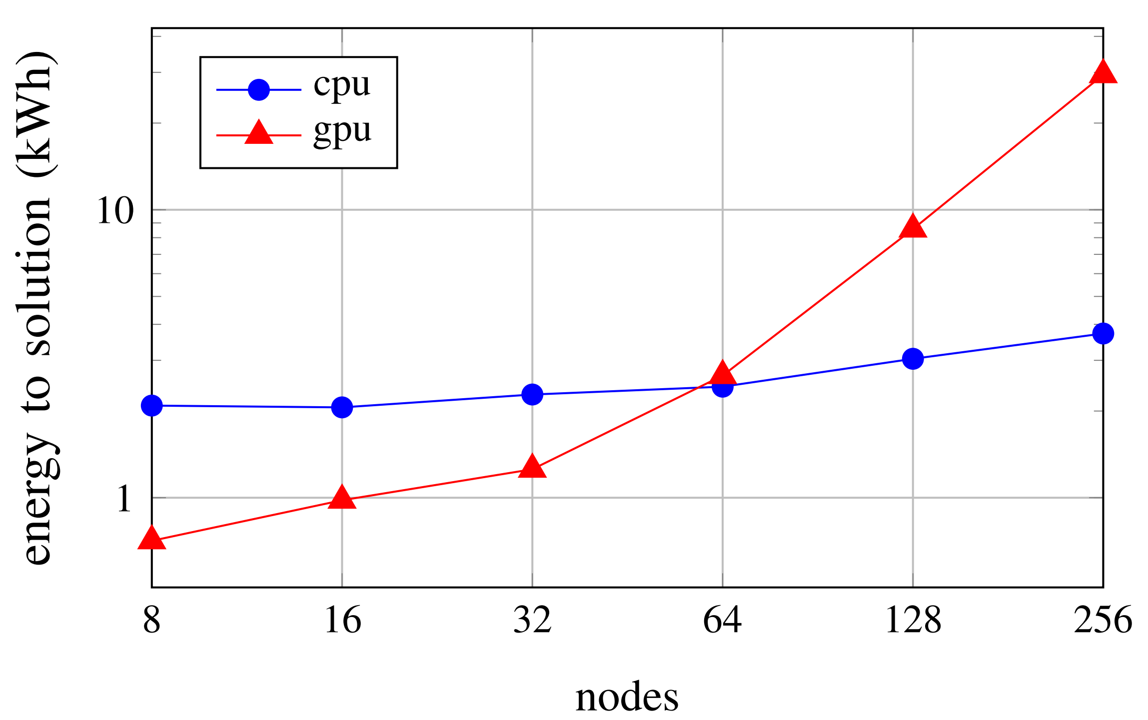

On GPUs: optimal run on 8 nodes with 0.66 kWh energy to solution

On CPUs: optimal run on 20 nodes with 2.2 kWh energy to solution

Relationship between energy and consumed resources

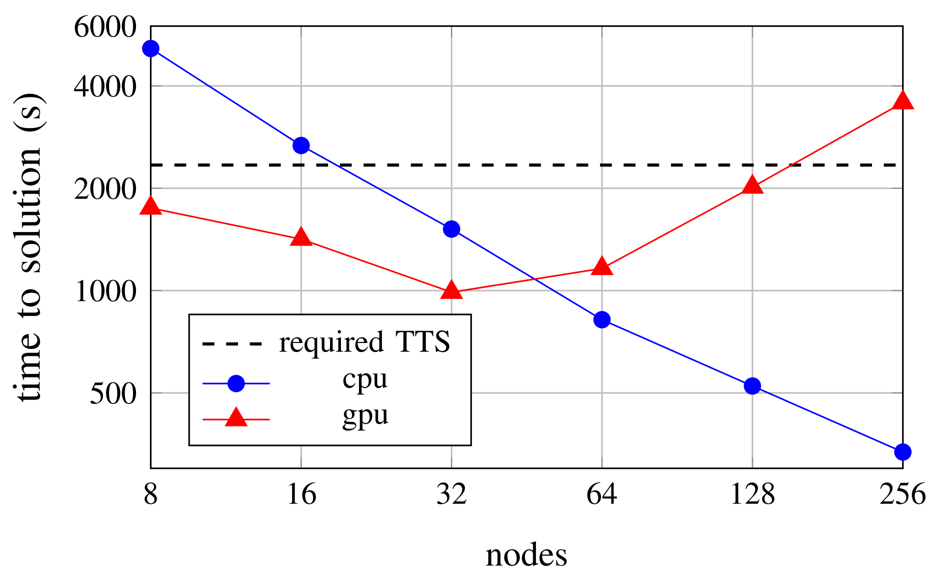

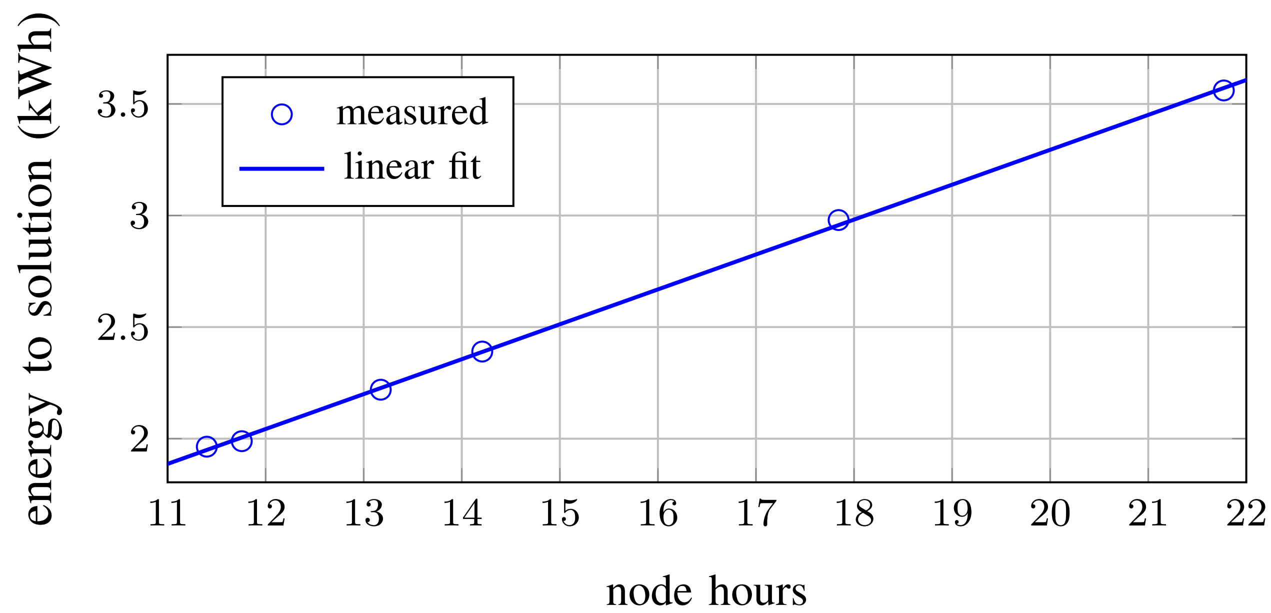

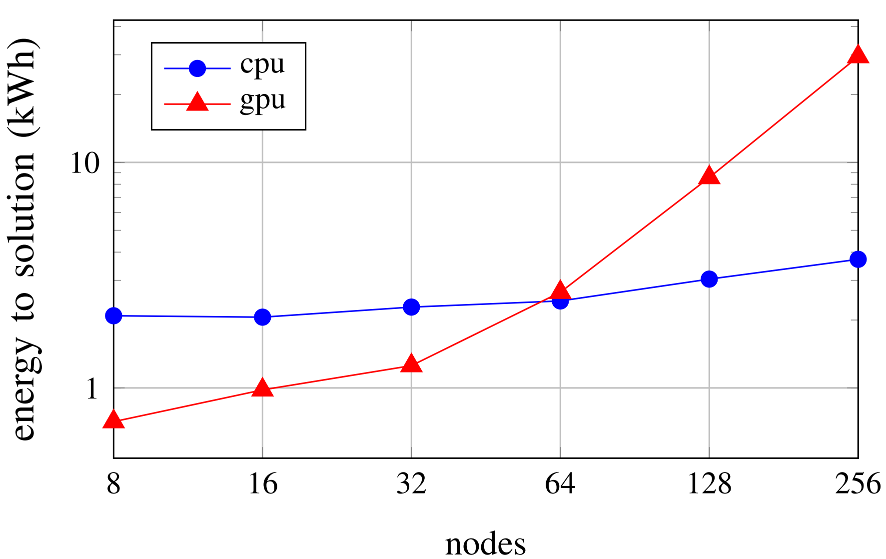

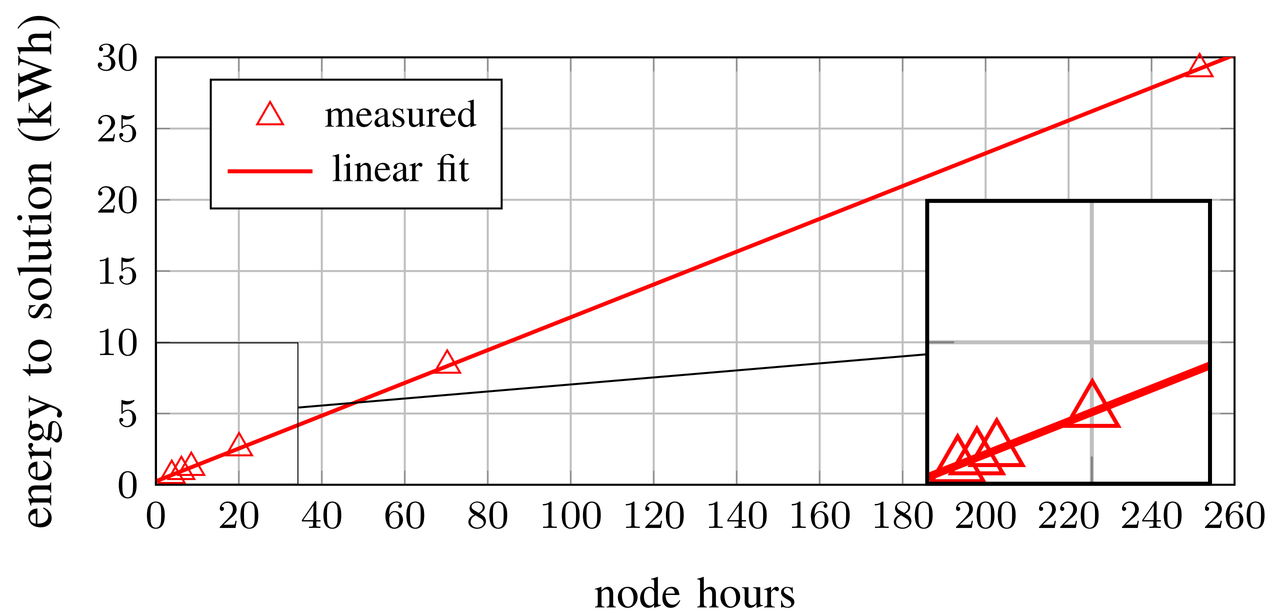

ETH zürich

Energy to solution vs. resources

time to solution on N nodes

$$E_N = \pi_0 N \tau_N + E_0$$

fitting parameters (for now)

arch	π (kW)	E_0 (kJ)	$1 - r^2$
GPU	0.115	880.8	$1.1 \cdot 10^{-5}$
CPU	0.156	599.5	$5.9 \cdot 10^{-4}$

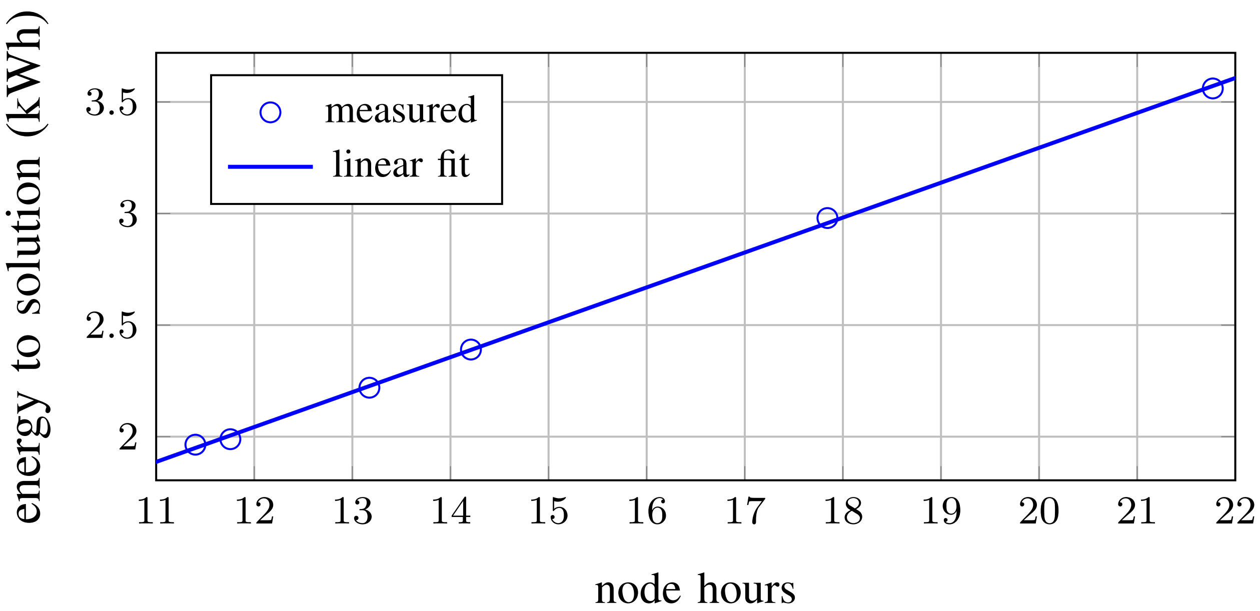

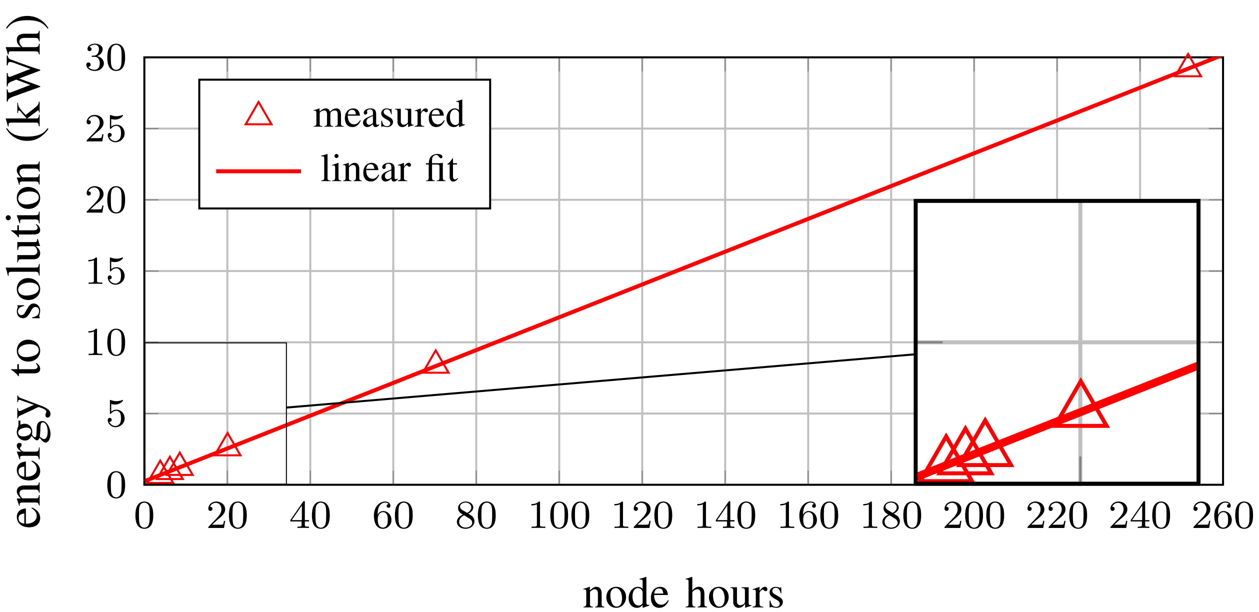

Shedding light on E_0 and π_0

Look at this in terms of single node energy E_1 and time τ_1 to solution

$$E_N = \pi_0 N \tau_N + E_0$$

$$E_0 = E_1 - \pi_0 \tau_1$$

parallel efficiency ↓

$$E_N = E_1 + N\tau_N(1 - P_N)\pi_0$$

Therefore: - E_0 is an effective dynamic energy

- π_0 is a an effective static power related to leakage losses

We can link energy and time to solution of a complex climate/weather simulation to the physics of the underlying computing system

In conclusion

- Science needs high-performance computing and data processing
- Data and computation are not separable in science
- Time and energy to solution are the performance metrics that matter

Finally:

While traditional complexity analysis is useful for time to solution, we need a new mathematical ideas for optimisation of energy to solution

Thank You!

Thomas Lippert for pointing out that Kepler did the first simulations in science

Ben Cumming, Gilles Fourestey, and Raffaele Solcà for help with ETS analysis