Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Federal Department of Home Affairs FDHA
Federal Office of Meteorology and Climatology MeteoSwiss

Achievement and learnings from the PP POMPA
Submission of a new HPC project IMPACT

Xavier Lapillonne, MeteoSwiss

V. Clement[3], O. Fuhrer[1], C. Osuna[1], K. Osterried[3], H. Vogt[2], A. Walser[1], C. Charpilloz[1], P. Spoerri[3], T. Wicky, P. Marti[3], R. Scatamacchia, U. Schaettler[5] and all PP POMPA contributors

[1]MeteoSwiss, [2]CSCS, [3]C2SM ETH, [4]COMET, [5]DWD

PP POMPA main achievements

- Project duration 2010-2018
- GPU capable version of the COSMO model (official release Q4 2018)
- Support for single and double precision
- Performance portable implementation of the Dynamical core using STELLA library
- OpenACC port of Physics, Assimilation and organization code (7500 directives)
- 7 publications

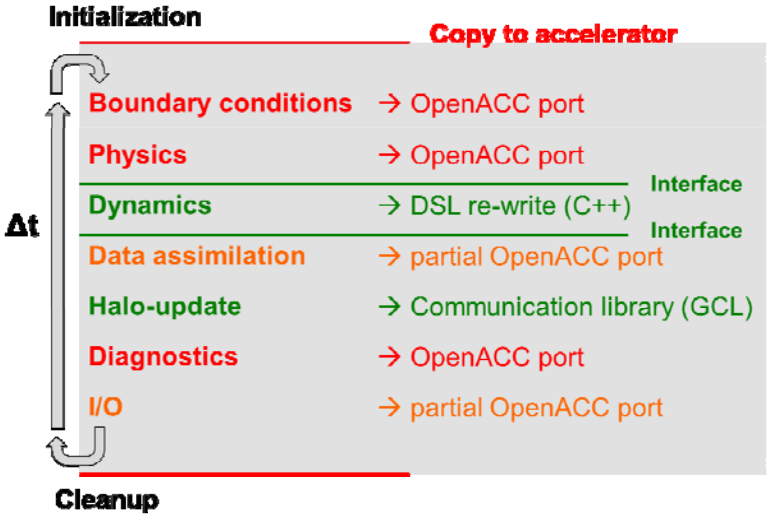

COSMO on GPU

! Thanks to all contributors !

Performance results

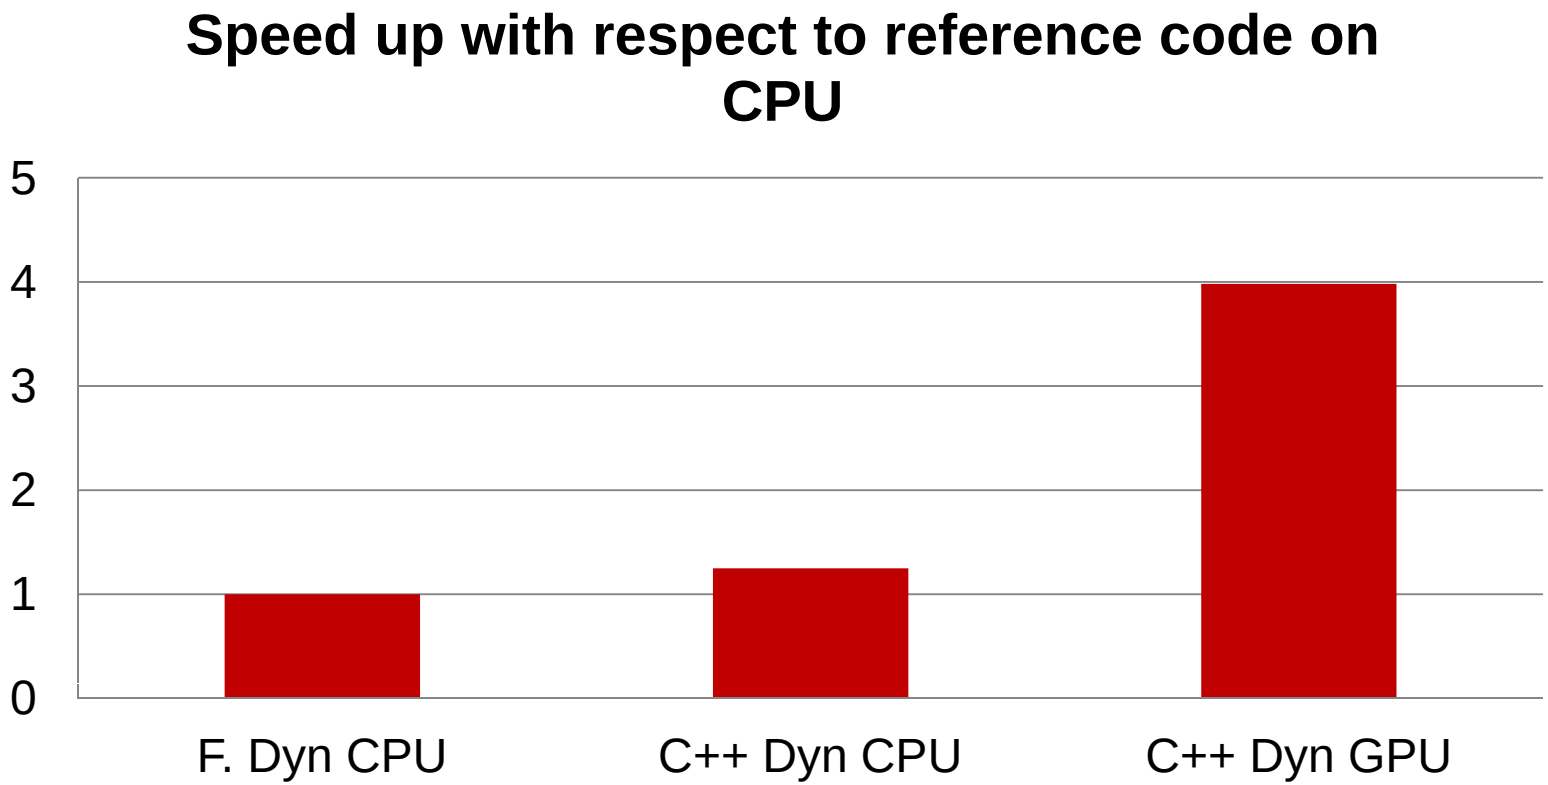

Results for 1 COSMO-E member 2h simulation using 8 GPU sockets (Nvidia P100) or 8 CPU sockets (8 Intel Haswell CPUs with 12 cores each = 96 cores) . Measured on the Piz Daint system at CSCS with COSMO 5.5-prerelease, branch full_gpu3

One application : quasi-global simulation on Piz Daint

Configuration

- 80°S to 80°N covering 98.4% of Earth's surface
- Regular lat/lon grid
- Idealized baroclinic wave test (Jablonowski and Williamson, 2006)
- 10 day simulation
- Minimal I/O
- Metric : SYPD = Simulated years per wallclock day

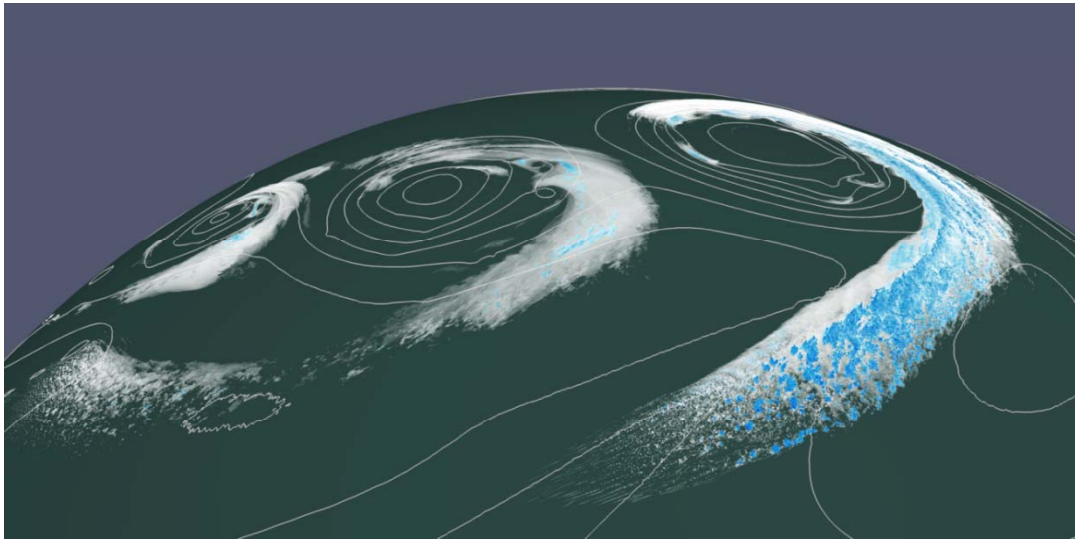

Visualization by Tarun Chadha (C2SM): clouds > 10^{-3} g/kg (white) and precipitation > 4 10^{-2} g/kg (blue)
Fuhrer et al., GMD 2018

Piz Daint

~5000 hybrid nodes with P100 GPUs (#6 Top500 6.2018)

MeteoSwiss

Fuhrer et al., 2017, GMD, in rev.

Strong scaling

- Near-global simulations at a fixed horizontal resolution

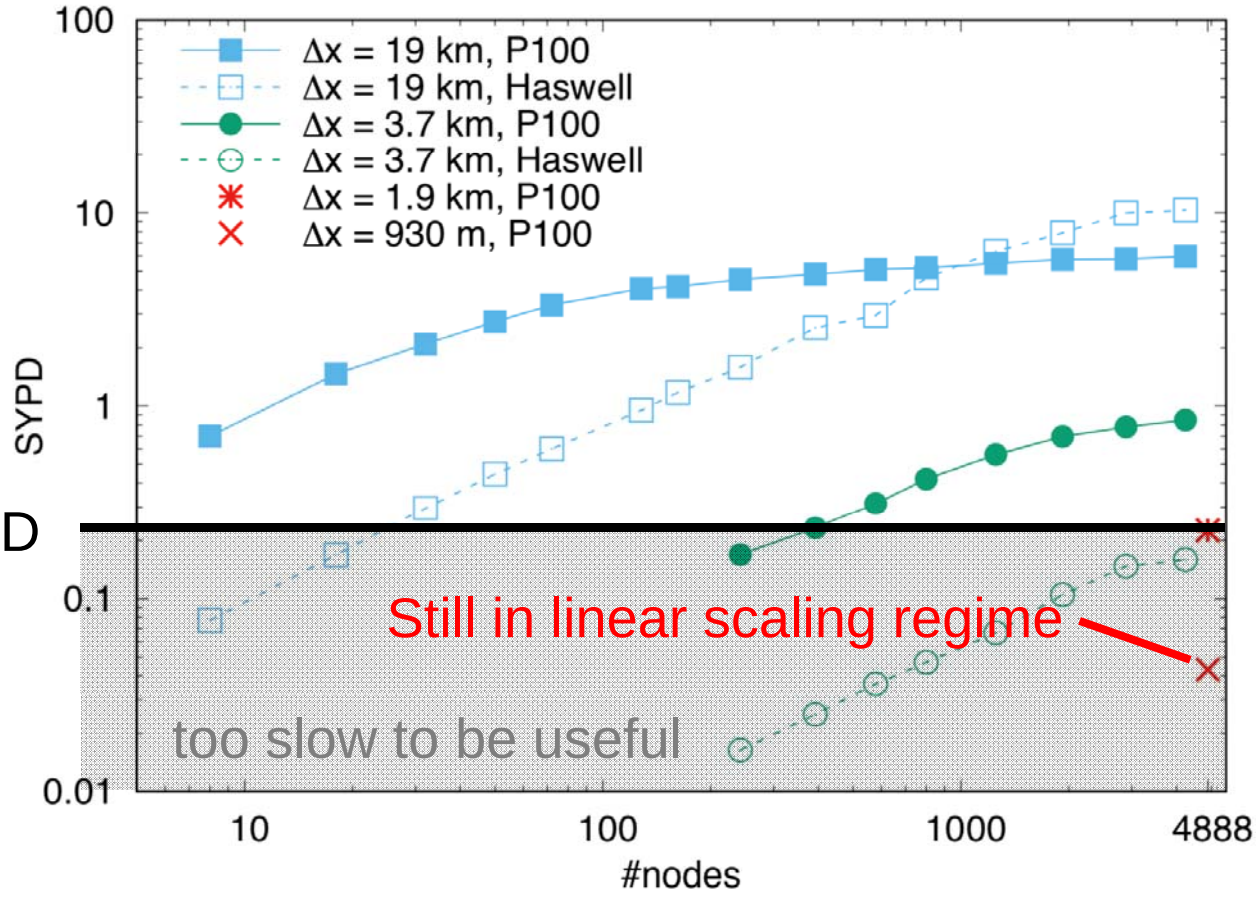

19 km

3.7 km

1.9 km

0.93 km

0.25 SYPD
Decadal climate simulation possible

Some learnings (1/2)

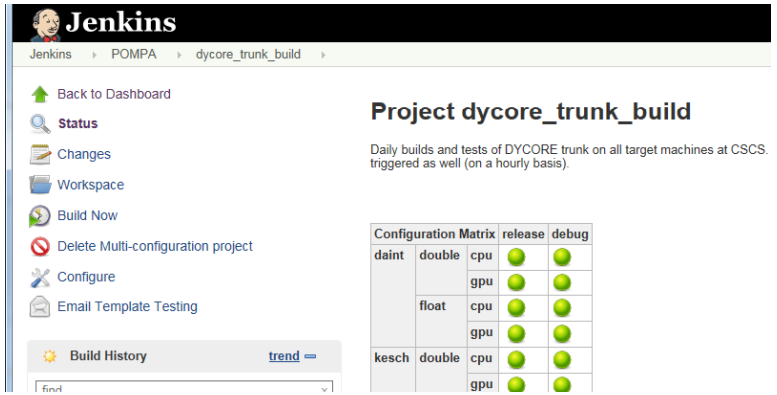

- Increased complexity :
  - GPU support
  - Single and double precision capability

➡ Improve software development process : code review with dedicated tool (Github Pull Request – thanks to the SCA for using this!), design review …

➡ Automatic on demand and nightly build and tests (Jenkins)

- **Moving target** : we ported COSMO to GPU 3 times (4.18, 5.0 and 5.5)
- Synchorinization with model developer and incremental step is key

Some learnings (2/2)

OpenACC directives

- Incremental insertion in existing code, ideal for porting large components
- Not always performance portable, only for GPUs

C++ embeded DSL

- Separation of concerns between user code and optimized backend
- Performance portable, future proof : can be extended to new architecture
- Steep learning curve, domain scientist did not like the language : we now have 2 dynamical cores!

Laplacian with Stella-DSL

```
template<typename TEnv>
struct Divergence {
  STENCIL_STAGE(TEnv)

  STAGE_PARAMETER(FullDomain, phi)
  STAGE_PARAMETER(FullDomain, lap)
  STAGE_PARAMETER(FullDomain, flx)

  static void Do(Context ctx, FullDomain) {
    ctx[div::Center()] = ctx[phi::Center()] -
      ctx[alpha::Center()] * (ctx[flx::Center()
      ctx[flx::At(iminus1)] + ctx[fly::Center()
      ctx[fly::At(jminus1)] )
  }
};
```

Future plans with COSMO on GPU

- MeteoSwiss will switch to the official COSMO version in Q4 2018
- Complete GridTools implementation of the C++ Dynamic Q4 2018, proposition to integration in official code Q1 2019
- COMET : ENS COSMO-ME (40 members), and ENS COSMO-IT (20 members) Q4 2018

Why investing a new HPC project ?

Physical limits , Moore's law is over

- 2004 end of frequency scaling → multi-core
- 2012 end of rapid cost decline → constant $/transistor
- 2019 heat dissipation constraints → massively parallel architecture
- 2021 end of reduction in feature size → ?

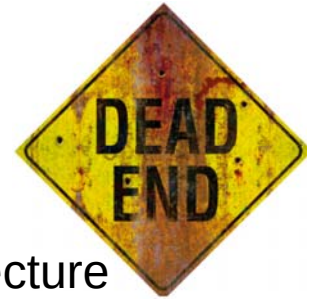

K. Flamm 2017, IEEE Computing in Science & Engineering

Specialized hardware, large diversity
increased on chip parallelism

GPUs

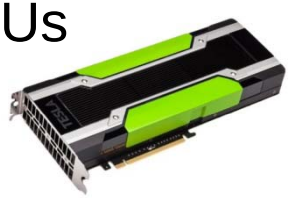

FPGAs

Vector

Many-core processors

Unknown

AI Accelerators
(e.g. Google Tensor Flow, Intel Nervana)

COSMO PP IMPACT (submitted)

ICON on Massively Parallel ArChiTecture

- Build on Know-How from the POMPA project
- Implement baseline GPU capable version **with OpenACC directives**
- Investigate OpenMP 4.5 for accelerator (currently not compilers)
- **Improve modularity** of the model structure
- Achieve performance portability on all architectures
  - Develop **new DSL approach**
  - Implement ICON Dycore based on **DSL (Domain Specific Language)**
  - Evaluate use of **CLAW-DSL** abstractions for the physical parameterisation
  - Investigate task parallelism

Low risk, known technologies

High risk
High potential
Future proof