Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Eidgenössisches Departement des Innern EDI
**Bundesamt für Meteorologie und Klimatologie MeteoSchweiz**

# IMPACT NWP and plan for ICON on GPU

**Remo Dietlicher** (2), V. Clement (2), O. Fuhrer (1), X. Lapillonne (1),
P. Marti (2), K. Osterried (2), C. Osuna (1), W. Sawyer (3)

(1) Federal Office of Meteorology and Climatology MeteoSwiss
(2) Center for Climate Systems Modeling, ETH Zürich
(3) Swiss National Supercomputing Centre (CSCS)

# Motivation: why use new hardware?

- They are already there!
- 8 out of the top 10 supercomputers (Top500 November 2018) are based on accelerator/many-core technologies
- Fastest supercomputer system Summit at ORNL is GPU based
- Piz Daint at CSCS, the largest machine in Europe (#5) is GPU based

→ ENIAC: ENabling the Icon model on heterogeneous ArciteCtures

→ PP IMPACT: Icon on Massively Parallel ArchiteCtures

Piz Daint hybrid system at CSCS

# COSMO PP IMPACT

Focus on NWP and LAM mode

- Future HPC will show ever more parallelism than actual GPU & CPU:
  - further adapt our models and programming tools
  - Investigate task parallelism
- IMPACT will investigate these aspects for **ICON**
- First implementation of a baseline **OpenACC** version :
  - Initial version of the dycore is available
  - Physical parametrizations require a full port
- Investigate appicability of **OpenMP for accelerators**
- Need to invest in **software engineering** beyond OpenACC in order to make support of different platforms maintainable in the long term:
  - Implement ICON Dycore based on **GridTools-DSL (Domain Specific Language)**
  - Evaluate use of **CLAW-DSL** abstractions for the physical parameterisation
  - **Improve modularity** of the model structure

# COSMO-1 6h Performance gains

CPU: 32x Intel® Xeon® E5-2690 v3
GPU: 32x NVIDIA® Tesla® P100

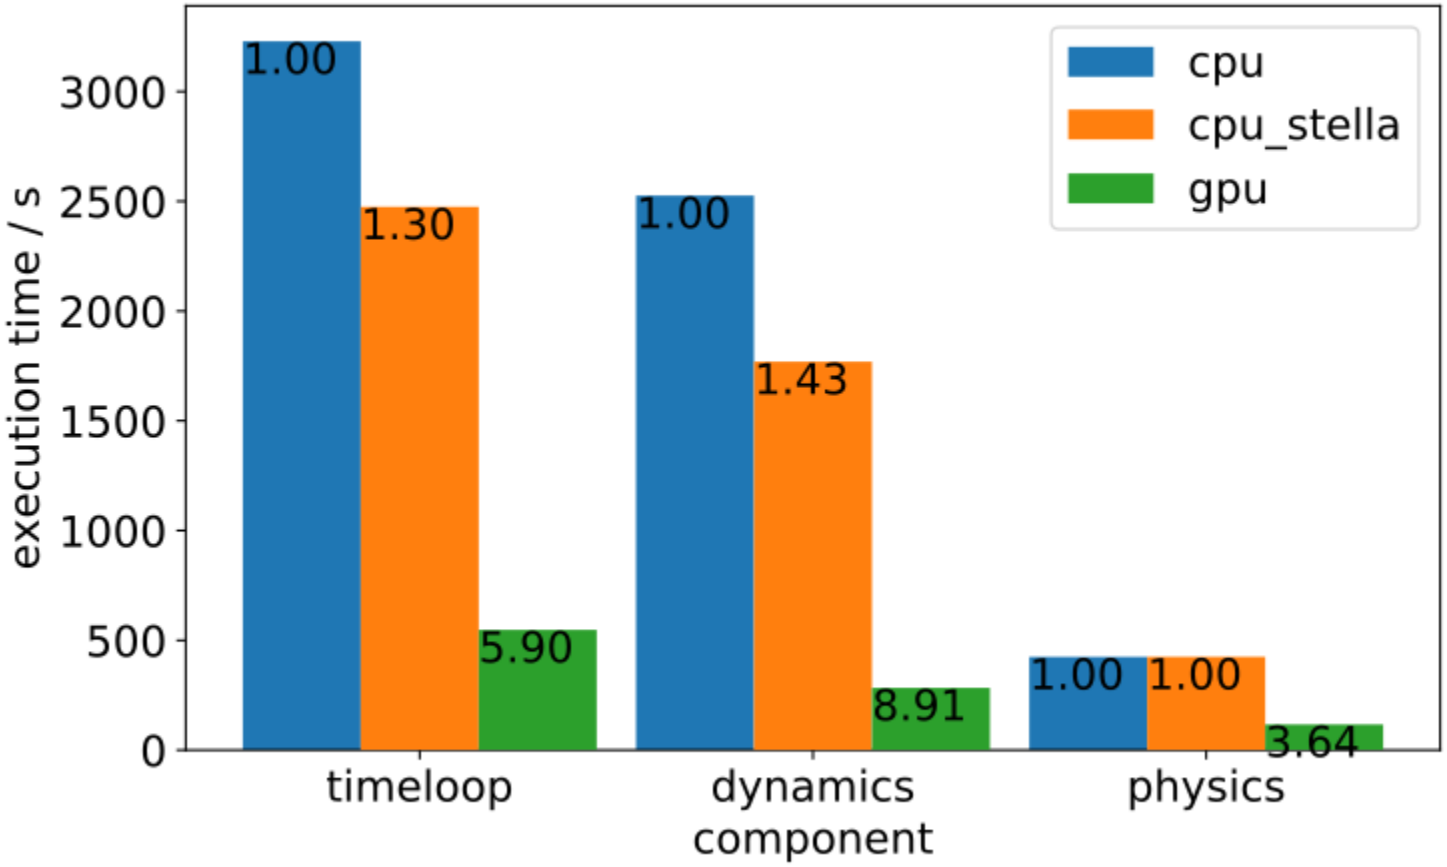


COSMO 5.05b_rc compiled with cray

# Exposing parallelism with OpenACC

```
 SUBROUTINE diag_pres (pt_prog, pt_diag, p_metrics,        &
                       jb, i_startidx, i_endidx, slev, nlev, lacc)
...
    !$ACC DATA PRESENT(pt_diag%tempv, pt_diag%pres_sfc, pt_diag%pres_ifc, pt_diag%pres, &
    !$ACC              pt_diag%dpres_mc, pt_prog%theta_v, pt_prog%exner, p_metrics%ddqz_z_full)

    !$ACC PARALLEL LOOP GANG VECTOR DEFAULT(PRESENT) &
    !$ACC          PRIVATE(dz1, dz2, dz3)
    DO jc = i_startidx, i_endidx
      ! Height differences between surface and third-lowest main level
      dz1 = p_metrics%ddqz_z_full(jc,nlev,jb)
      dz2 = p_metrics%ddqz_z_full(jc,nlev-1,jb)
      dz3 = 0.5_wp*p_metrics%ddqz_z_full(jc,nlev-2,jb)

      ! Compute surface pressure starting from third-lowest level; this is done in order to avoid
      ! contamination by sound-wave activity in the presence of strong latent heating
      pt_diag%pres_sfc(jc,jb) = p0ref * EXP( cpd_o_rd*LOG(pt_prog%exner(jc,nlev-2,jb)) + &
        grav_o_rd*(dz1/pt_diag%tempv(jc,nlev,jb) + dz2/pt_diag%tempv(jc,nlev-1,jb) +     &
        dz3/pt_diag%tempv(jc,nlev-2,jb)) )

      pt_diag%pres_ifc(jc,nlev+1,jb) = pt_diag%pres_sfc(jc,jb)
    ENDDO
...
    !$ACC END DATA
  END SUBROUTINE diag_pres
```

# Status of porting ICON-LAM for NWP to GPU

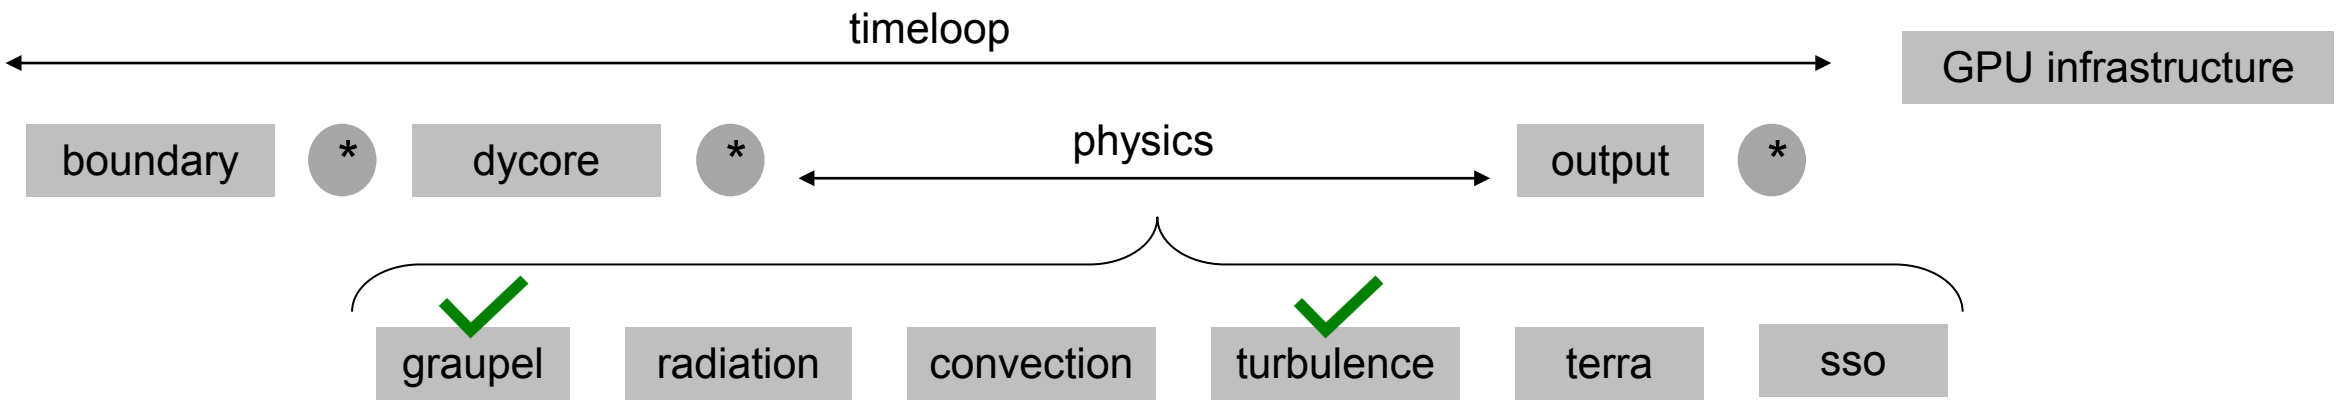


Ported as part of the ENIAC project and CSCS

Can be adapted from COSMO

ecRad? RRTMGP?

* 2-way nesting infrastructure, physics interfaces, output diagnostics, interpolation, …

# The challenge of keeping up
## Porting to GPU

Feature workflow
generic case

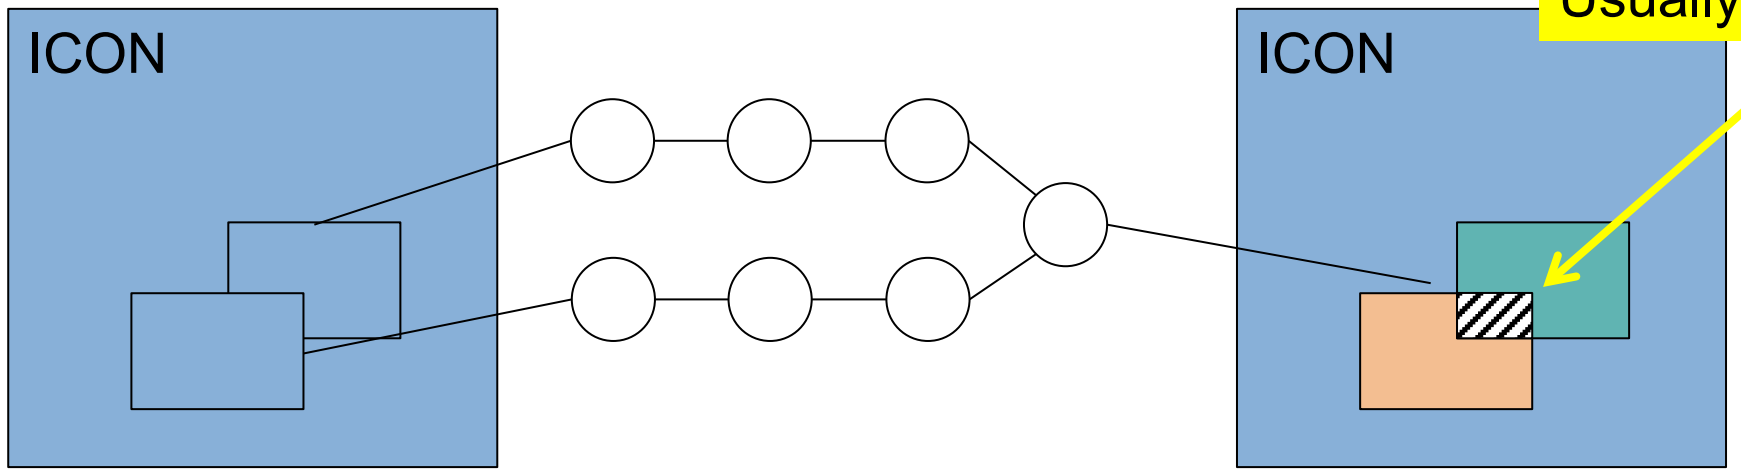


"OpenACC" workflow

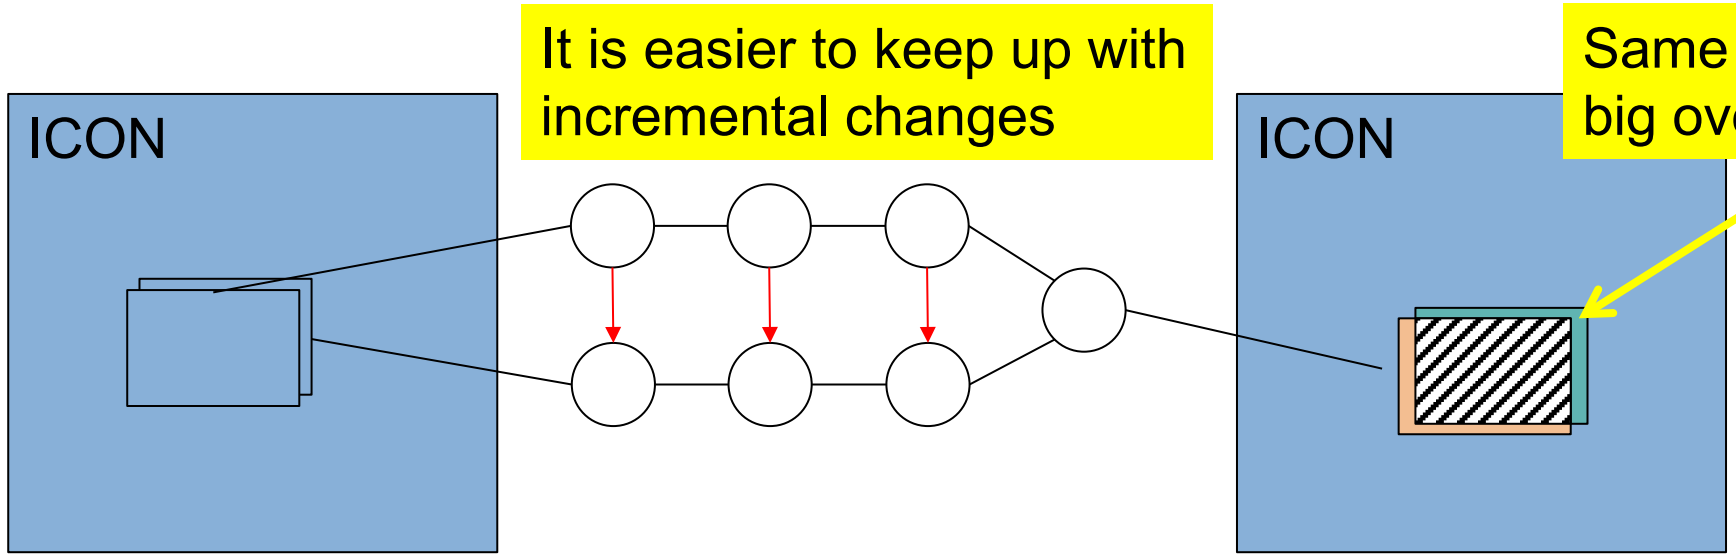

# The challenge of keeping up
## Updating GPU capable code – Status quo

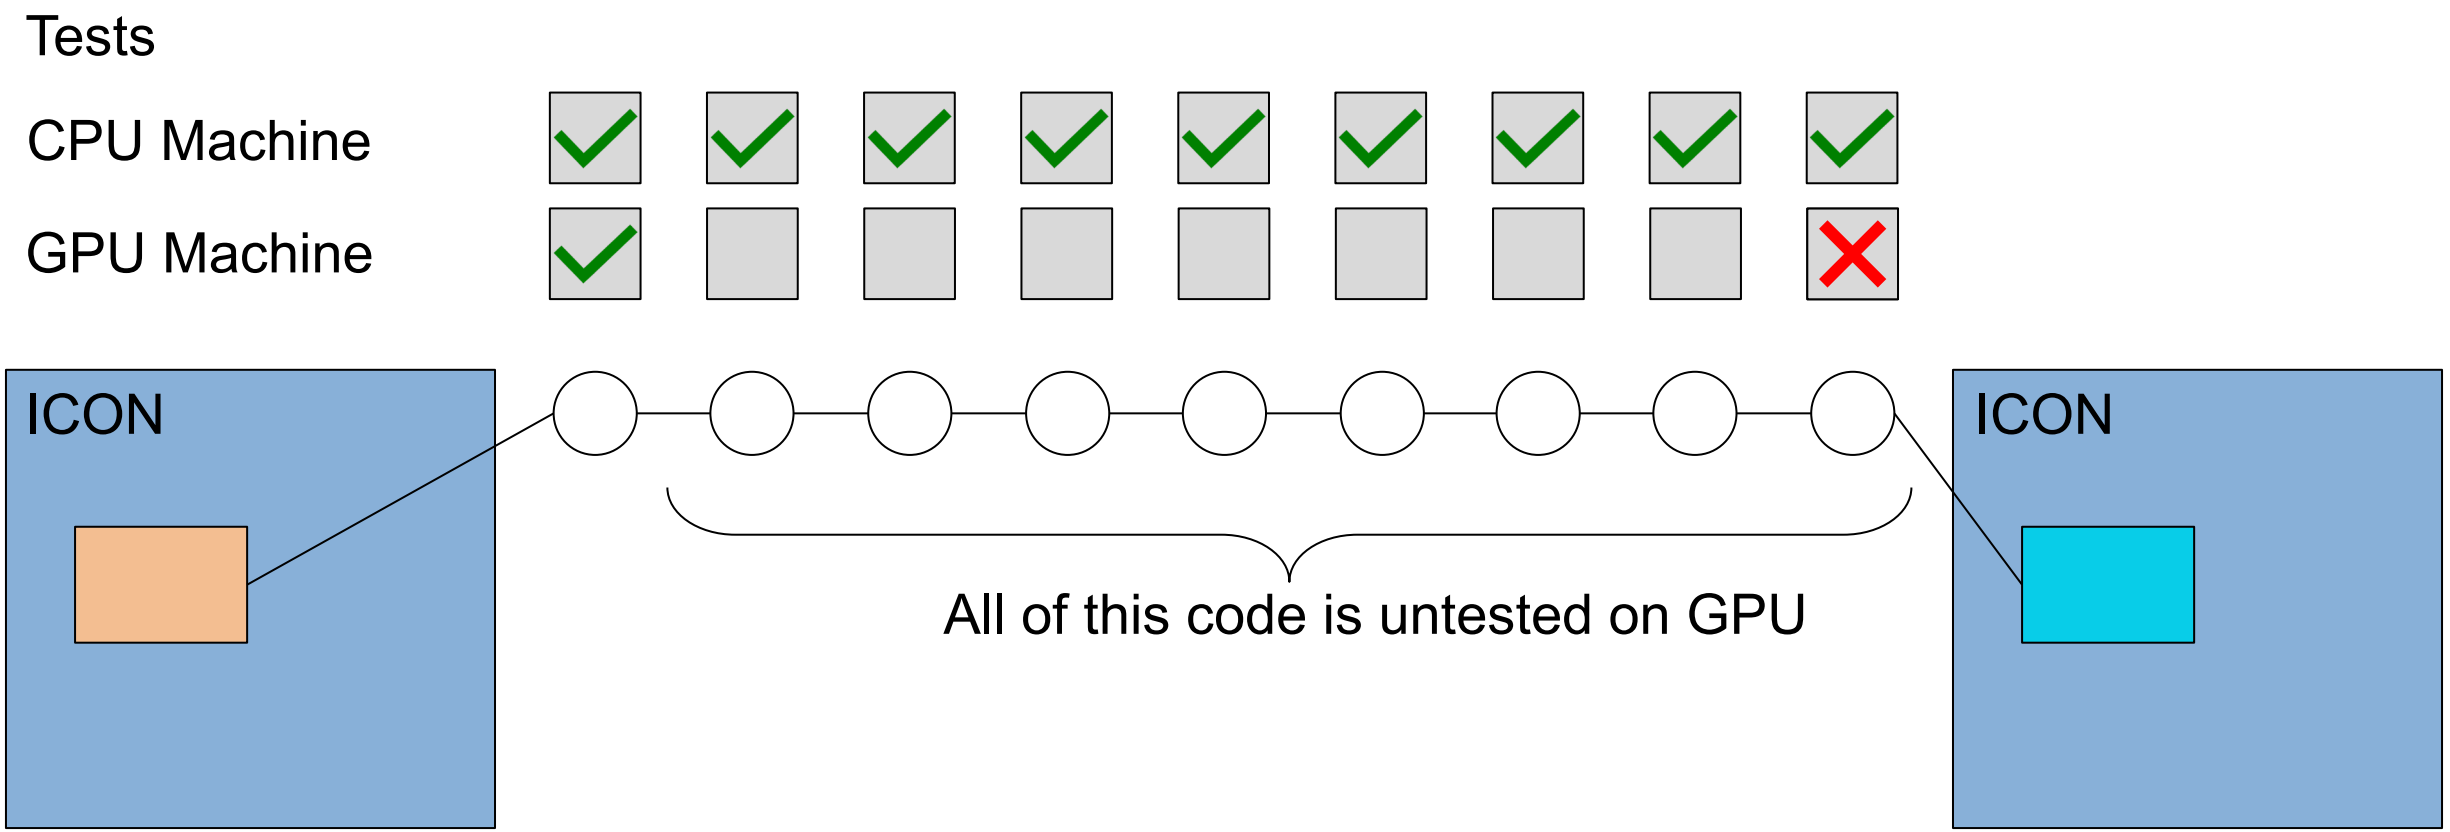

# The challenge of keeping up
## Updating GPU capable code – Continuous integration

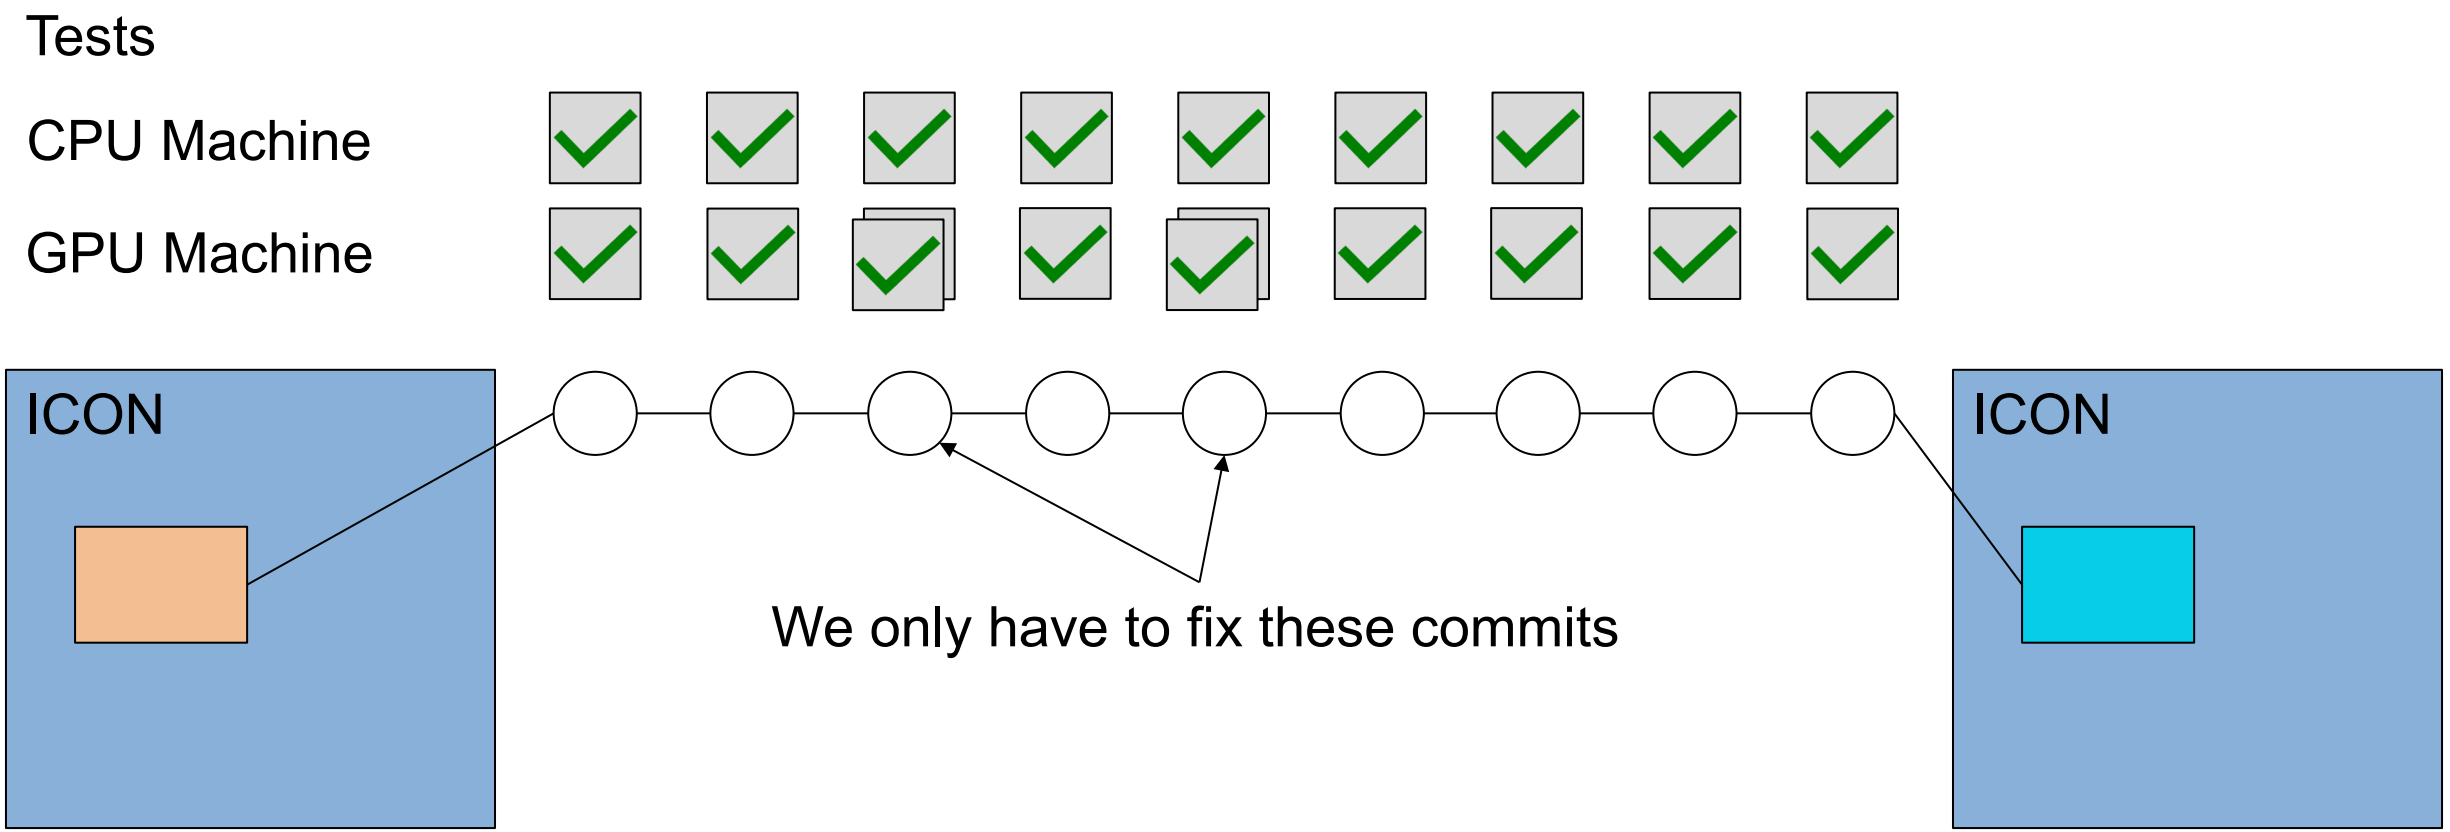

# Continuous integration in ICON

We are missing both of the aspects outlined above:

- Testing
  - There is (almost) no testing for the GPU code on Buildbot.
  - Recently added test for climate physics.
  - A test for NWP physics still pending.
  - The tests that are run are branch dependent.
  - Can be resolved in front of a Computer.
- Frequent merges
  - Each institute has their own repository and target application (MCH: weather for Switzerland on GPUs, DWD: weather for Germany on CPUs, MPI-M: climate projections, …). Therefore merging takes a lot of time and brings limited gain.
  - This point needs to be resolved at a round table.

# Making sure results are correct on GPU

These are results for a global test (atm_amip)

We need a reference output and tolerances for each test

References and tolerances need to be updated for each change in model physics

There is a script to do that, needs to be incorporated in ICON workflow

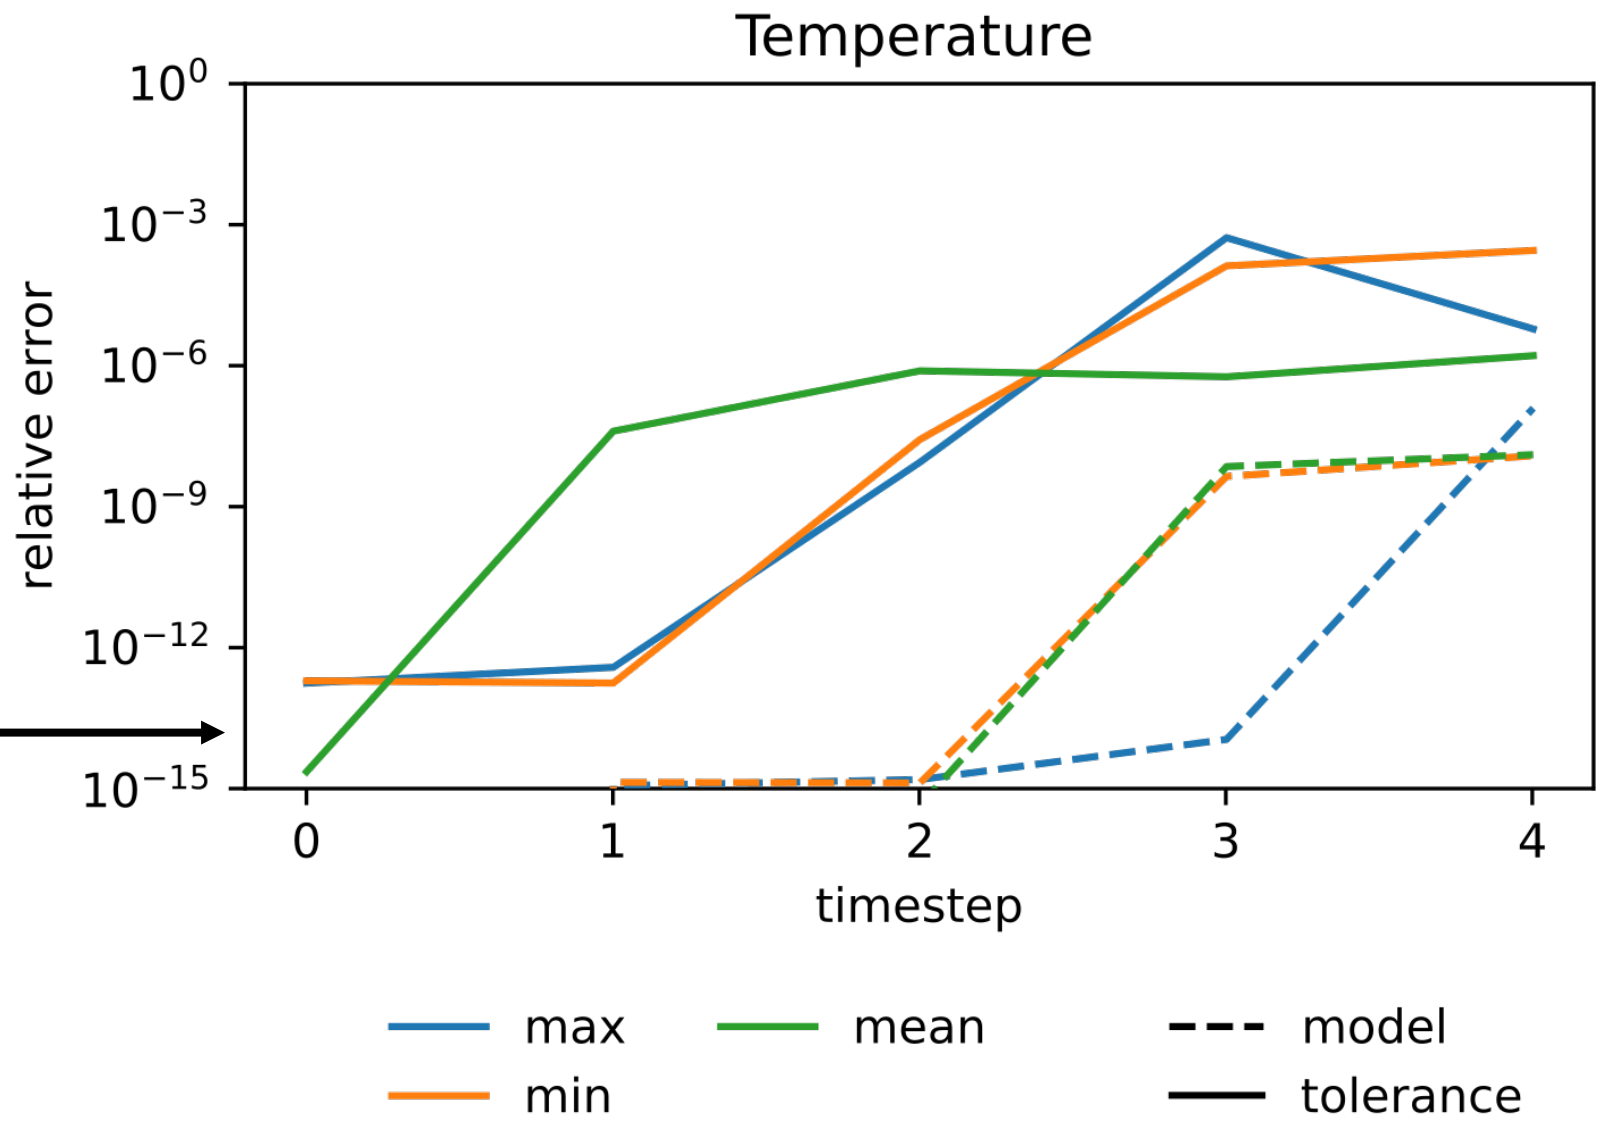


Perturb initial conditions by 10^-14 (relative)

# Testing OpenACC

```
!$ser verbatim call serialize_input(&
!$ser verbatim          ncall,&
!$ser verbatim          jg,&
!$ser verbatim          p_metrics, &
!$ser verbatim          p_prog, &
!$ser verbatim          p_diag)

  CALL diag_pres (pt_prog, pt_diag, p_metrics,         &
                  jb, i_startidx, i_endidx, slev, nlev)

!$ser verbatim call serialize_output(&
!$ser verbatim          ncall,&
!$ser verbatim          jg,&
!$ser verbatim          p_diag)
```

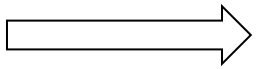

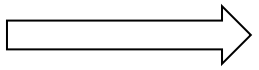

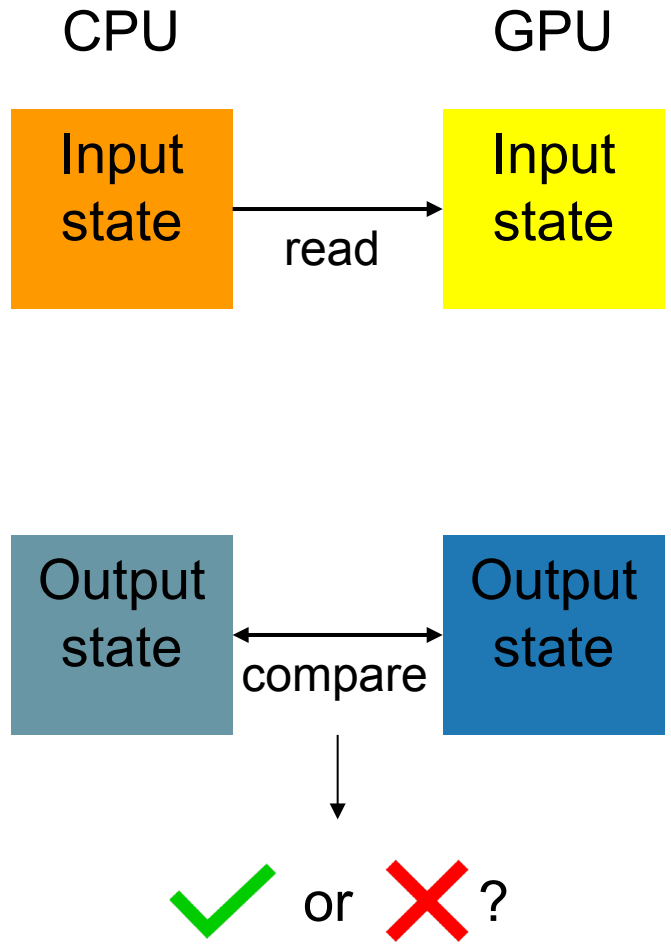

# Summary

- GPU port of ICON is on the way!
- Porting with OpenACC touches almost every line of code
- Keeping up to date with model development makes porting easier
- Technical aspects: improve testing on GPUs
- Collaboration aspects: Can we achieve continuous integration?

**MeteoSchweiz**
Operation Center 1
CH-8058 Zürich-Flughafen
T +41 58 460 91 11 www.meteoschweiz.ch

**MeteoSvizzera**
Via ai Monti 146
CH-6605 Locarno-Monti
T +41 58 460 92 22
www.meteosvizzera.ch

**MétéoSuisse**
7bis, av. de la Paix
CH-1211 Genève 2
T +41 58 460 98 88
www.meteosuisse.ch

**MétéoSuisse**
Chemin de l'Aérologie
CH-1530 Payerne
T +41 58 460 94 44
www.meteosuisse.ch

- Task 2.1: Port the interfaces to the turbdiff, vertdiff and graupel, including testing infrastructure with serialization. Microphysics and turbulence completed.

- Task 3.1: CLAW was applied to microphysics and turbulence for performance portability between CPU and GPU. Code is in master of COSMO physics, need to be imported into ICON

- Task 3.2: We organized a DSL definition workshop in January 2019 with participation of MPI and DWD ICON developers

- Task 2.3: First evaluation of OMP vs OpenACC for standalone kernels.

- Task 2: Baseline performance on CPU and GPU.
  - Performance comparison ICON <-> COSMO on GPU for 1.1km alpine domain

- Task 1.1: Discussions and proposals for an ICON git workflow that follows best practices in CI.

- Task 1.3: Improve testing infrastructure:
  Implement a tolerance check as we have in COSMO that now runs on Buildbot.

- Task 5.1 Organize joint sessions with ENIAC project at the ICON developers meeting, regular weekly meetings and a GPU Hackathon for key ICON developers