Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

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

ENIAC Project : ICON port to GPU for climate application using OpenACC and CLAW-DSL

Xavier Lapillonne, Valentin Clement, Philippe Marti, Remo Dietlicher, Will Sawyer, Marco Girogetta, Luis Kornblueh, and the ENIAC team
MeteoSwiss, C2SM, ETHZ, CSCS, MPI-M

Overview

- The ENIAC project
- Porting strategy
- JSBACH soil model with the CLAW-dsl
- Outlook

ICON ENIAC (Enabling ICON on GPU)

Focus on global climate mode

- Adapt the ICON model to run on GPU and many-core architecture
  - Base line GPU port with OpenACC compiler directives
- Prepare the ICON model for actual use-cases in weather and climate on emerging HPC systems
- Achieving a high degree of performance portability by:
  - using source-to-source translation tool : CLAW
  - using alternative approaches such as the GridTools domain-specific language

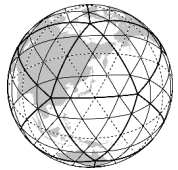

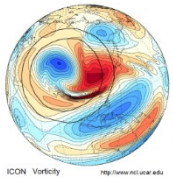

- Participants : ETHZ, C2SM, MeteoSwiss, MPI-M, DWD, DKRZ
- For NWP : IMPACT project, see Remo Dietlicher's tals

Application : QUBICC (MPI-M)

- "Quasi-Biennial Oscillation in a Changing Climate (QUBICC)“
- The QBO winds as observed for ~60 years are fairly regular, except for the breakdown in early 2016. How is QBO affected in a changing climate
- PRACE proposal submitted to run global simulations at ca 2 km on Piz Daint GPU system

Modell configuration

- Dynamics/transport on a 2.5 km grid (R2B10) with ca. 200 layers up to ca. 80 km.
- Radiation: Option A for GPUs: RRTMGP; option B for CPUs: concurrent PSRAD
- Vertical diffusion implicitly coupled to JSBACH-lite"Graupel" cloud microphysics & saturation adjustment

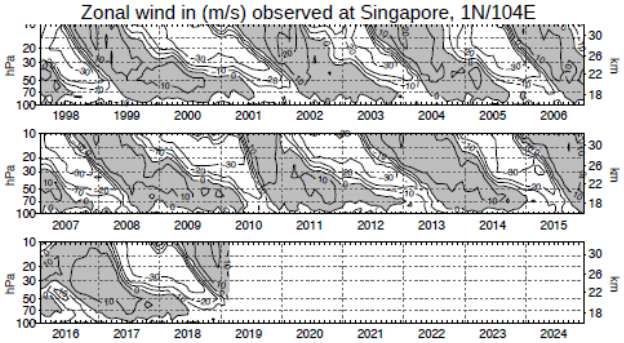

Credit: Marco Giorgetta MPI-M

Porting Strategy

- First base version OpenACC:
  - Dynamics+Climate physics+interface code : OpenACC
  - JSBACH (soil model) : claw-dsl
- Required code adaptation and changes
- Performance for the full model, cannot not be evaluated yet until a consistent set is ported. Should be availble in October 2019
- Report performance for separate component
- Only supported compiler : PGI (requirement for OpenACC 2.6)

Refactoring of the Dynamics

Original OpenACC

```
SUBROUTINE solve_nonhydrostatic_eqns
  …
  !$ACC PARALLEL LOOP GANG
  DO jb = 1, nblocks
    CALL get_indices_c( )
    !$ACC LOOP VECTOR COLLAPSE(2)
    DO jk = 1, nlev
      DO jc = 1, nproma
        prog_var(jc,jk,jb) = f(jc,jk,jb)
      END DO
    END DO
  END DO
  !$ACC END PARALLEL
   :
  !$ACC PARALLEL LOOP GANG
  DO jb = 1, nblocks
   :
  END DO
  !$ACC END PARALLEL
   :
END SUBROUTINE solve_nonhydrostatic_eqns
```

Optimal GPU nproma: 1000 — 4000

Refactored OpenACC

```
SUBROUTINE solve_nonhydrostatic_eqns
  …
  DO jb = 1, nblocks
    CALL get_indices_c( )
  !$ACC PARALLEL
    !$ACC LOOP GANG VECTOR COLLAPSE(2)
    DO jk = 1, nlev
      DO jc = 1, nproma
        prog_var(jc,jk,jb) = f(jc,jk,jb)
      END DO
    END DO
  !$ACC END PARALLEL
  END DO
  :
  DO jb = 1, nblocks
    CALL get_indices_c( )
  !$ACC PARALLEL
     :
  !$ACC END PARALLEL
  END DO
   :
END SUBROUTINE solve_nonhydrostatic_eqns
```

Optimal nproma: as big as possible

Difficult patterns for GPUs

Index lists

```
DO jk = slev, elev
  DO je = i_startidx, i_endidx
    IF ( condition(je,jk) ) THEN
      icnt_c1 = icnt_c1 + 1
      idxlist_c1(icnt_c1,jb) = je
      levlist_c1(icnt_c1,jb) = jk
```

Array of structs

```
fl_e1(je,jk)%p1%lon = arrival_pts(je,1,1,jk,jb)
```

Struct of arrays

```
fl_e1%p1%lon(je,jk) = arrival_pts(je,1,1,jk,jb)
```

Array syntax

```
IF (lclean_mflx) THEN
  prep_adv%mass_flx_me(:,:,jb) = 0._wp
  prep_adv%vn_traj    (:,:,jb) = 0._wp
ENDIF
```

Innermost functions

```
DO jk = slev, elev
  DO je = i_startidx, i_endidx
    z_gauss_pts(je,3,1) = DOT_PRODUCT(shape_func(1:4,3),z_x(je,1:4))
```

Need to workaround for OpenACC

Performance dynamic component

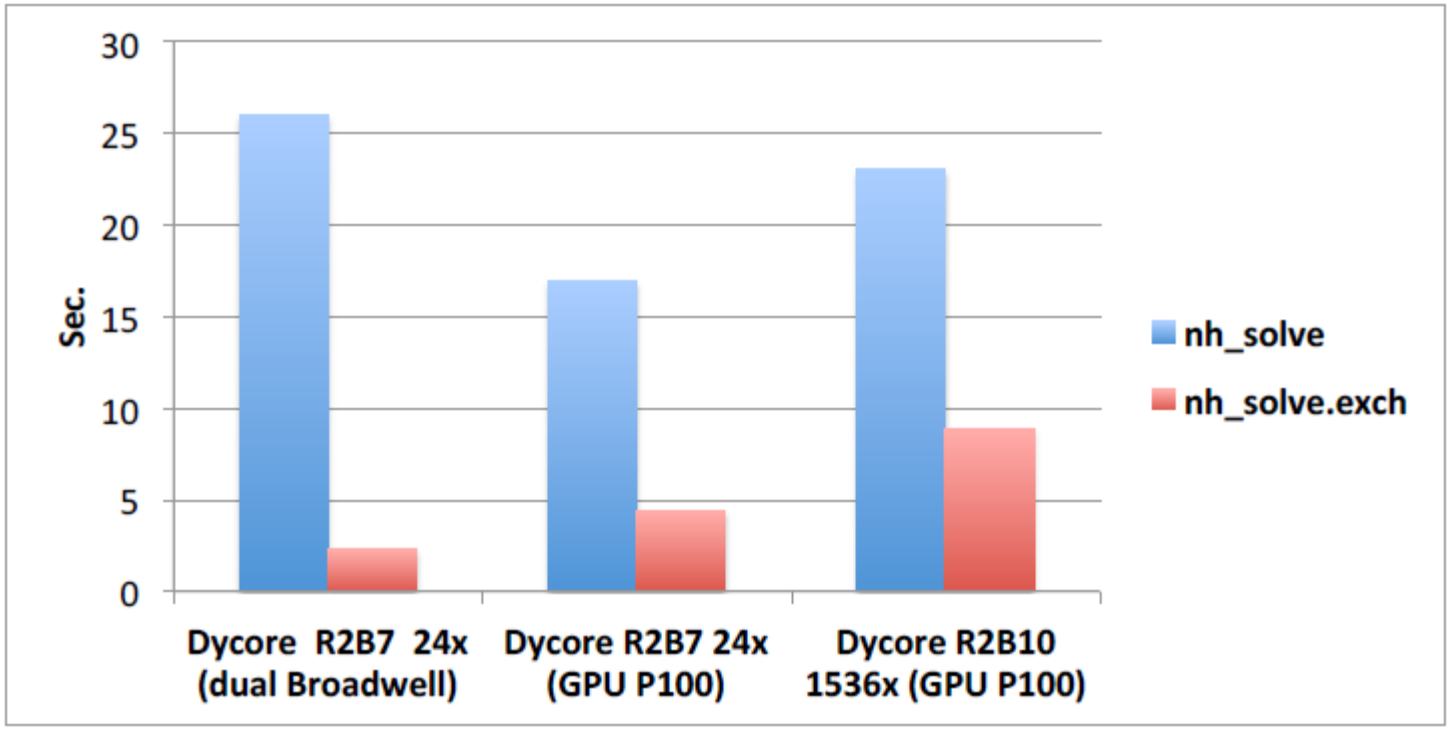

- R2B7 : Global, ca 20km
- Communication needs to be optimized

Source : Will Sawyer ,CSCS

Status GPU port using OpenACC

Focus on a reduce set of the climate physics

Component	Status	Comments
Dynamics	Done	- new implementation to fit strategy in the physics (large nproma) - Still need to remove updates
ECHAM Physics : surf_update, vdiff, …	Done,	- Stil neet to remove updates - Optimization on ongoing
Physics: JSBACH	Done	- using CLAW DSL, no plain OpenACC port planned at this stage - Validation need to be finalized
Physics: Radiation (RRTMGT)	On going	- RRTMGP already exist with OpenACC - Need to port interface
Physics:Graupel microphys. (Seifert)	On going	- Almost complete, need to be merged with latest change from MPI (tendency)
Physics: 2 moments microphys. (HAM)	Not started	
aerosol module (HAM-light)	Not started	
Organization code	On going	- Close to completion for QUBICC case. Should be completed in november 2019

Experience porting with OpenACC

- Difficulties, Compiler issues (deep copy, atomics, …)
- Difficult pattern to port, e.g. large derived type, Object oriented Fortran
- Limited compiler support for newer standard version >= 2.6 (only PGI)
- Validation of GPU vs. CPU results is time consuming
- ICON is VERY large : ~800 Fortran module files, ~1mio LOC
- Rapidly evolving code

JSBACH - ICON soil

- Fortran 2003/2008 code style - ~25k LOC
- Special DSL to define storage and access, Pipeline of tasks on different tiles
- 80% of compute code written as ELEMENTALs
- Elementals function are a special case of single column abstraction with no vertical loop : use claw-dsl
- CLAW port : ~140 EXPAND kernels automatically generated, ~90 ELEMENTAL (SCA) kernels automatically generated, ~100 files parsed by the CLAW Compiler, ~20 files with deep transformation
- Port is now 99% complete

CLAW Single Column Abstraction (SCA)

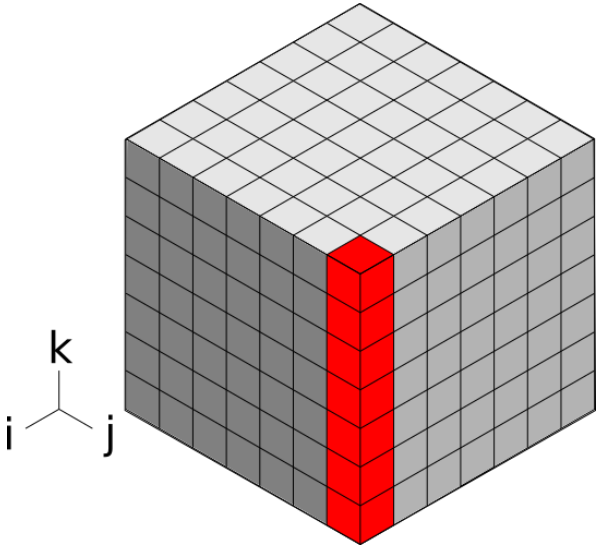

Targets physical parameterization

- No notion of independent horizontal dimension
  - No DO statements over horizontal
  - Arrays are demoted to get rid of the horizontal

Separation of concerns

- Domain scientists focus on their problem (1 column, 1 box)
- CLAW Compiler produce code for each target architecture and directive languages

“The CLAW DSL: Abstractions for Performance Portable Weather and Climate Models.” PASC’18 Proceedings. https://doi.org/10.1145/3218176.3218226

What is the CLAW Compiler?

- Source-to-source translator
- Based on the OMNI Compiler Project
- Fortran 2008
- Open source under the BSD license
- Available on GitHub with the specifications
- High-level transformation framework

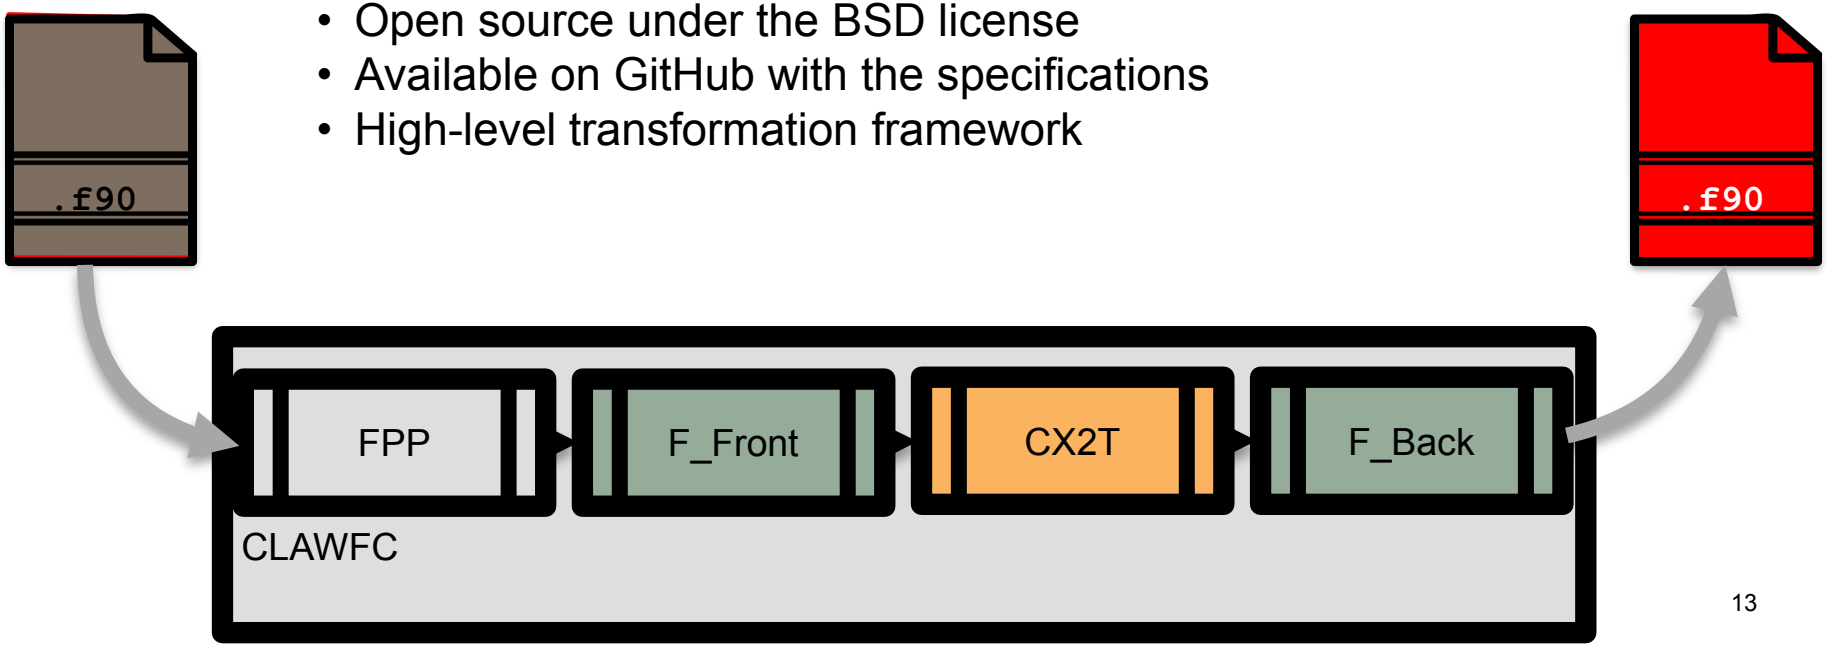

Code example (original + CLAW directive)

CLAW directives

```
PURE ELEMENTAL SUBROUTINE calc_radiation_surface_net(swvis_down, swnir_down, &
  alb_vis, alb_nir, lw_down, t, rad_net, swvis_net, swnir_net, sw_net, lw_net)

  USE mo_phy_schemes, ONLY: lwnet_from_lwdown

  !$claw model-data
  REAL(wp), INTENT(in) :: swvis_down, swnir_down, alb_vis,  alb_nir, lw_down, t
  REAL(wp), INTENT(out) :: rad_net
  REAL(wp), INTENT(out), OPTIONAL :: swvis_net, swnir_net, sw_net, lw_net
  !$claw end model-data
  !$claw sca

  REAL(wp) :: zswvis_net, zswnir_net, zsw_net, zlw_net

  ! Compute net SW radiation from downward SW and albedo
  zswvis_net = swvis_down * (1._wp - alb_vis)
...
```

Code example (transformed)

clawfc -D__ICON__ --model-config=icon_jsbach.toml --target=gpu --directive=acc -o code_transformed.f90 original.f90

```
SUBROUTINE calc_radiation_surface_net ( swvis_down , swnir_down , &
 alb_vis,  &
 alb_nir , lw_down , t , rad_net , swvis_net , swnir_net , sw_net , lw_net)
 USE mo_phy_schemes , ONLY: lwnet_from_lwdown
 INTEGER , INTENT(IN) :: kproma

 REAL ( KIND= wp ) , INTENT(IN) :: swvis_down ( : )
 REAL ( KIND= wp ) , INTENT(IN) :: swnir_down ( : )
 REAL ( KIND= wp ) , INTENT(IN) :: alb_vis ( : )
 REAL ( KIND= wp ) , INTENT(IN) :: alb_nir ( : )
 REAL ( KIND= wp ) , INTENT(IN) :: lw_down ( : )
 REAL ( KIND= wp ) , INTENT(IN) :: t ( : )
 REAL ( KIND= wp ) , INTENT(OUT) :: rad_net ( : )
…
```

```
…
!$acc data &
!$acc present(swvis_down,swnir_down,alb_vis,alb_nir,&
!$acc ,  lw_down,t,rad_net &
!$acc ,swvis_net,swnir_net,sw_net,lw_net)
!$acc parallel
!$acc loop gang vector
 DO horizontal = 1 , kproma , 1
  zswvis_net = swvis_down ( horizontal ) * ( 1._wp - alb_vis …
  zswnir_net = swnir_down ( horizontal ) * ( 1._wp - alb_nir  …
  zsw_net = zswvis_net + zswnir_net
  zlw_net = lwnet_from_lwdown ( lw_down ( horizontal ) , …
  rad_net ( horizontal ) = zsw_net + zlw_net
…
```

Performance Results

- For GPU claw can generate OpenACC or OpenMP for accelerator directives
- OpenACC results obtained using the PGI compiler, timed in full application
- OpenMP for accelerator : Cray. Note : due to current limitation and issues with the compiler the results are obtained with standalone kernels

Performance results

Performance comparison (socket to socket) of JSBACH tasks on Intel Haswell E5-2690v3 and NVIDIA P100. Domain size (horizontal grid points x vertical levels) = 20480 x 47

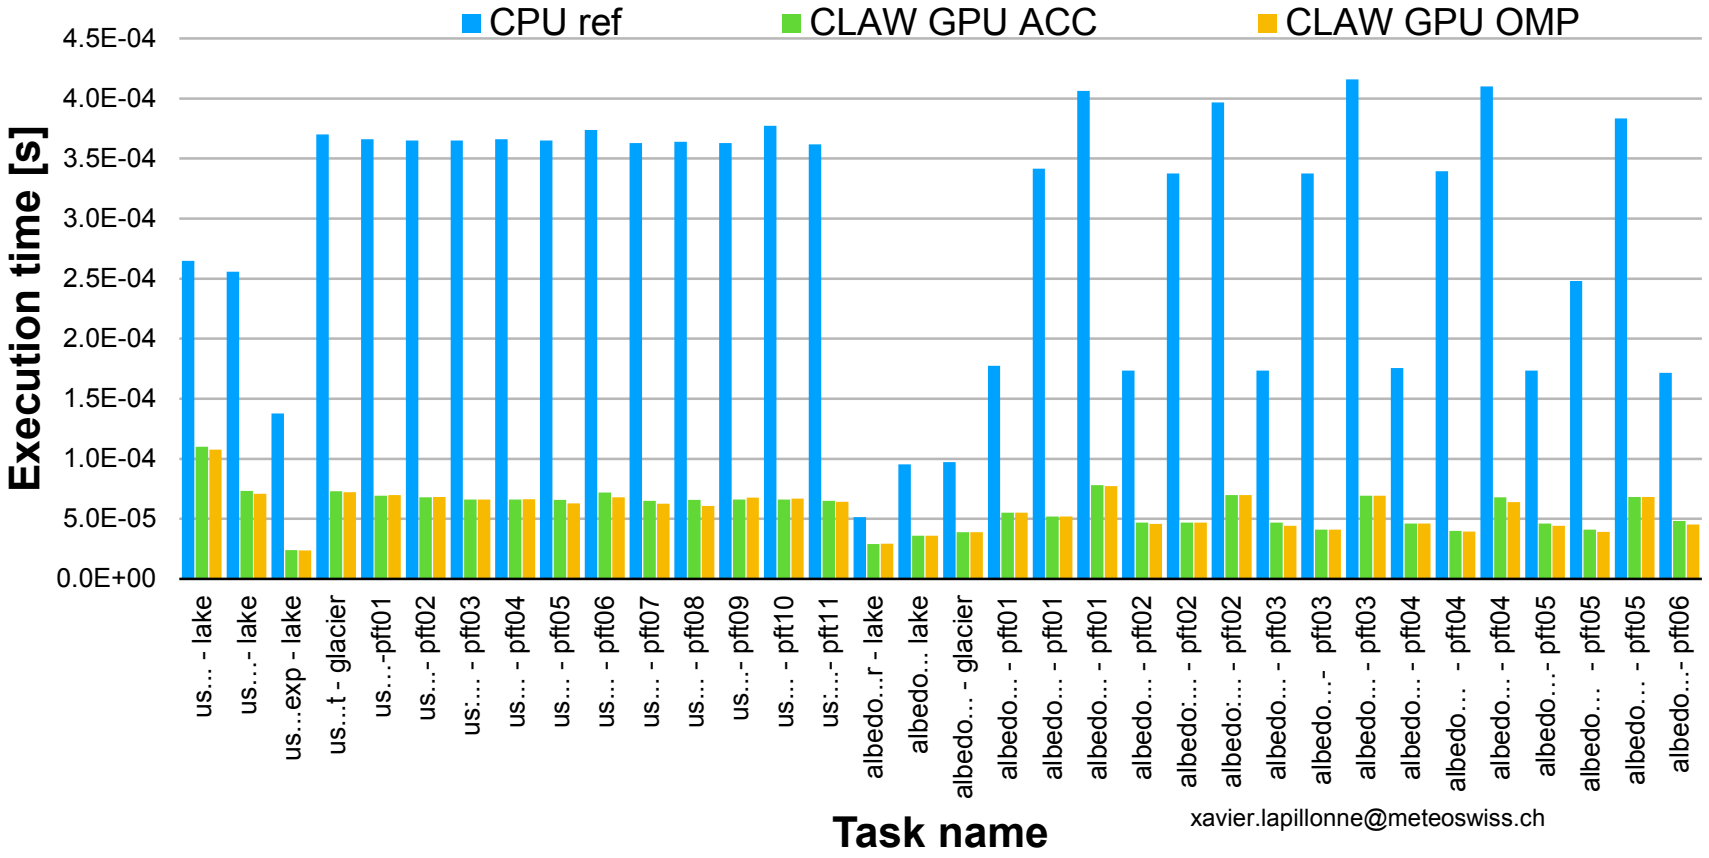

xavier.lapillonne@meteoswiss.ch

JSBACH CLAW-dsl

- Speed up on GPU between 3 to 8 compare to CPU
- Work with CLAW-dsl : ISC'19 Award
- Will be published
  Special issue Supercomputing Frontiers and Innovations
  "Automatic Port to OpenACC/OpenMP for PhysicalParameterizations in Climate and Weather Code Using the CLAW Compiler", accepted with minor modifications

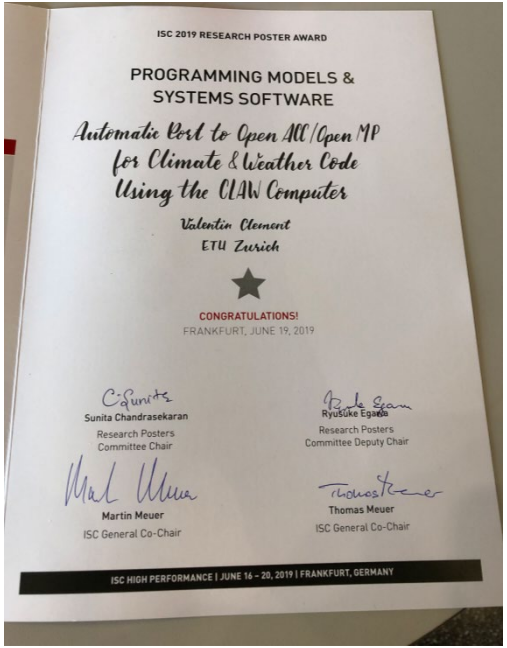

ENIAC outlook and learnings

- First (reduced) GPU version for Climate applications should be ready by Q4 2019 using OpenACC and CLAW-DSL
- MeteoSwiss/C2SM contribution completed in September 2019 (V. Clement + P. Marti)
- Project running until June 2020 (stil funding for MPI-M)
- Potential application : QUBICC, 2km Global on Piz Daint

- Learnings :
  - More difficult and complex to work and port ICON to GPU than COSMO
  - Coordination is critical, regular meeting and video calls with MPI-M,DKRZ,CSCS was key to advance in the project.

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

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

MeteoSwiss
Operation Center 1
CH-8058 Zurich-Airport
T +41 58 460 91 11
www.meteoswiss.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