Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra





EvgLAM/SRNWP Meeting 2023





Priority Project IMPACT

Goal: Port ICON to GPUs

- Almost all components for operational ICON configuations ported to GPUs with OpenACC directives
- Running pre-operationally at MeteoSwiss starting Oct 2, 2023
- Project will be completed in December 2023

Main contributors

N. Burgdorfer¹, V. Cherkas¹, R. Dietlicher¹, E. Germann¹, F. Gessler¹, D. Hupp¹, X. Lapillonne¹, C. Müller¹, M. Röthlin¹, M. Stellio¹, G. Van Parys¹, G. Vollenweider¹, C. Osuna¹, A. Walser¹, M. Bettiol³, R. Meli³, A. Gopal³, M. Jacob⁴, A. Jocksch³, J. Jucker², R. Meli³, W. Sawyer³, U. Schättler⁴, D. Alexeev⁵, R. Scatamacchia ¹*MeteoSwiss*, ²C2SM, ³CSCS, ⁴DWD, ⁵*Nvidia*, ⁶*ITAF Met*



completed
in progress
on CPU + data copy and interface ported
not started



Performance of ICON on GPU







4

Lessons learned

- Prioritize investments into software refactoring
- Plan on rewriting / refactoring your software multiple times
- Increase level of abstraction (no separation of concerns in Fortran + X)
- Choose a programming model that developers like
- Provide solutions for integrating with existing legacy code and incremental adoption

ICON-C - Consolidation Project

- ICON community infrastructure project to consolidate the ICON code and make it more modular, i.e. encapsulate the different components with clear/defined interfaces
- Allow for more **scalable development workflow**, in particular through an improved testing of individual components

Main contributors

SMO

Claudia Frauen; Joerg Behrens; Ralf Mueller; P. Samanta; Daniel Reinert; Florian Prill; Thomas Deppisch; Roland Wirth; Sergey Kosukhin; Luis Kornblueh; Reiner Schnur; Claudius Holeksa; Yen-Chen Chen; Terry Cojean; William Sawyer; Andreas Jocksch; Jonas Jucker; Christopher Bignamini; Xavier Lapillonne, Roland Potthast, Daniel Klocke, Hendryk Bockelmann, Gholamali Hoshyaripour, Oliver Fuhrer





What are the aims of ComIn?

- Providing a standardized and stable **public interface** for third party codes ('**plugins**') coupled to ICON
- Significantly **reduced maintenance** for ICON as well as for third party code developers
- Plugins easier to migrate to new ICON releases
- Enables **multi-language support** (Fortran, C/C++, Python)





How does ComIn work in a nutshell?

- ComIn organizes the **data exchange** and **simulation events** between the ICON model and multiple plugins.
- **ComIn Callback Register**: Subroutines of the plugins are called at pre-defined events during the ICON simulation.
- The **ComIn Adapter Library** is included by ICON and the plugins. It contains descriptive data structures and regulates the access and the creation of model variables.

So what about Python?

Machine Learning Scientist



Atmospheric Scientist

$$\vec{u}(\vec{x}) = \sum_{i=1}^{k} \lambda_i \phi_i(\vec{x}) \vec{n}_i$$

Low level of abstraction General purpose language No separation of concerns



IC N

Implementation (Fortran + OpenACC):

```
!SACC PARALLEL &
!$ACC PRESENT( iqidx_d, ..., ptr_vn_d, e_flx_avg_d, vn_d, vt_d, rbf_vec_coeff_e_d )
!$ACC LOOP GANG PRIVATE( i startidx, i endidx, jb )
    DO jb = i_startblk, i_endblk
     IF ( i_startblk == jb ) THEN; i_startidx = e_startidx; ELSE; i_startidx = 1; ENDIF
      IF ( i endblk == jb ) THEN; i endidx = e endidx; ELSE; i endidx = nproma; ENDIF
!$ACC LOOP VECTOR COLLAPSE(2)
    DO je = i startidx, i endidx
        DO jk = 1, nlev
          iqidx 1 = iqidx d(je,jb,1)
          ! Average normal wind components
          ptr_vn_d(je,jk,jb) = e flx avg_d(je,1,jb)*vn_now_d(je,jk,jb)&
            + e flx avg d(je,2,jb)*vn now d(igidx 1,jk,igblk 1) &
            :
          ! RBF reconstruction of tangential wind component
          vt now d(je,jk,jb) = rbf vec coeff e d(1,je,jb) &
            * vn now d(igidx_1,jk,igblk_1) &
            :
        ENDDO
      ENDDO
    ENDDO
                                                                                 10
!SACC END PARALLEL
```

Ū

GT4Py

- Open-source, co-developed by CSCS, MeteoSwiss, and Al2
- <u>https://github.com/gridtools/gt4py</u>
- Build a "wrapper" around C++ library used to refactor COSMO
- Expose a productive, high-level language
- Inter-operate with Python ecosystem (e.g. NumPy)
- Often referred to as a domainspecific language (DSL)
- Activities at ECMWF (FVM), NOAA (X-SHiELD), NASA (GEOS)

import gt4py.gtscript as gtscript

@gtscript.function

```
def advection_x(dx, u, abs_u, phi):
    adv_phi_x = u[0, 0, 0] / (60. * dx) * (
        + 45. * (phi[1, 0, 0] - phi[-1, 0, 0])
        - 9. * (phi[2, 0, 0] - phi[-2, 0, 0])
        + (phi[3, 0, 0] - phi[-3, 0, 0])
        ) - abs_u[0, 0, 0] / (60. * dx) * (
        + (phi[3, 0, 0] + phi[-3, 0, 0])
        - 6. * (phi[2, 0, 0] + phi[-2, 0, 0])
        + 15. * (phi[1, 0, 0] + phi[-1, 0, 0])
        - 20. * phi[0, 0, 0] )
    return adv_phi_x
```



Code generation

GT4Py code needs to be translated into NumPy or fast code in order to be executed



Similar to ML-frameworks or packages such as Numba, Cython, ...



Demonstration: X-Shield Model

Rewrite X-SHiELD model using GT4Py

- dynamical core
- physical parameterizations
- Infrastructure

Project duration of 2 years









GFDI

Pace demonstrates how a high-level language can insulate us from disruptive changes, provide a more productive development environment, and facilitate the integration with new technologies such as machine learning.

1 Introduction

Current weather and climate models are written in low-level compiled languages such as Fortran for performance, and typically 15 run on high-performance computing (HPC) systems with CPUs. With the end of Moore's law (e.g. Theis and Wong, 2017), HPC

Comparison" on Piz Daint



X-SHIELD @ GT4Py (Dahm et al. 2023) runs at 0.073 SYPD at 1.85 km, 80 levels, 4056 nodes, only microphysics

COSMO (Fuhrer et al. 2018) reported 0.23 SYPD at 1.9 km, 60 levels, 4888 nodes, only microphysics \rightarrow would scale to ~0.12 SYPD for comparable simulation

ICON (Giorgetta et al. 2022) reported 0.13 SYPD at 5 km, 191 levels, 1024 nodes, full physics \rightarrow would scale to ~0.068 SYPD for comparable simulation

IFS (Schulthess et al. 2018) reported 0.088 SYPD at 1.25 km, 62 levels, 4888 nodes, full physics \rightarrow would scale to ~0.20 SYPD for comparable simulation





EXCLAIM Project (ICON-X)



ICON in GT4Py / Python

- Dynamical core, tracer advection and microphysics have been ported.
- Infrastructure for unit-testing and integration back into Fortran codebase
- Next step: Performance improvements (currently on-par with Fortran + OpenACC)

ICON-DSL dycore

Low Level Clouds (R02B09) - 5KM







Extension to unstructured grids

Fortran + OpenMP + OpenACC



GT4Py Declarative

47	dvt_tang = (
48	-(
49	u_vert(E2C2V[0]) * dual_normal_vert_x(E2ECV[0])
50	+ v_vert(E2C2V[0]) * dual_normal_vert_y(E2ECV[0])
51)
52) († (
53	u_vert(E2C2V[1]) * dual_normal_vert_x(E2ECV[1])
54	+ v_vert(E2C2V[1]) * dual_normal_vert_y(E2ECV[1])
55)

Summary

- ICON is available for production on GPU for NWP and climate applications (Fortran + OpenACC)
- On-going effort to improve software engineering, modularization and code maintainability (ICON-C)
- Rewrite of ICON in Python / GT4Py on-going
- Productivity is generally underappreciated



Schweizerische Eidgenossenschaft Confédération suisse Confederazione Svizzera Confederaziun svizra

MeteoSchweiz

÷

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

MeteoSvizzera MétéoSuisse MétéoSuisse Via ai Monti 146 7bis, av. de la Paix Chemin de l'Aérologie CH-6605 Locarno-Monti CH-1211 Genève 2 CH-1530 Payerne T +41 58 460 94 44 T +41 58 460 92 22 T +41 58 460 98 88 www.meteosuisse.ch www.meteosvizzera.ch www.meteosuisse.ch 23 42 ÷ \$ 4 \$ 5 ÷ 국는 슈 수 53 수 ÷ **MeteoSchweiz** 子 令 子 5 슈 4 2