



This project has received funding from the European Union's Horizon 2020 Research and Innovation programme under grant agreement No. 801015



EXA2PRO overview

Dr. Lazaros Papadopoulos
ICCS/NTUA

EoCoE EXA2PRO workshop



22 February 2021

EXA2PRO at a Glance

- **Call identifier:** H2020-FETHPC-2017
- **Topic:** FETHPC-02-2017 - Transition to Exascale Computing
- **8 partners from 5 countries**
- **Budget:** 3,475,223 €
- **Duration:** 36 months
- **The project officially ends on 30 July 2021**



EXA2PRO motivation and objectives



Enhancing Programmability and boosting Performance
Portability for Exascale Computing Systems

Motivation



Need to support the growth of HPC applications and their transition to exascale.

Currently, **only few applications** are expected to benefit from upcoming exascale computing systems

Main goal

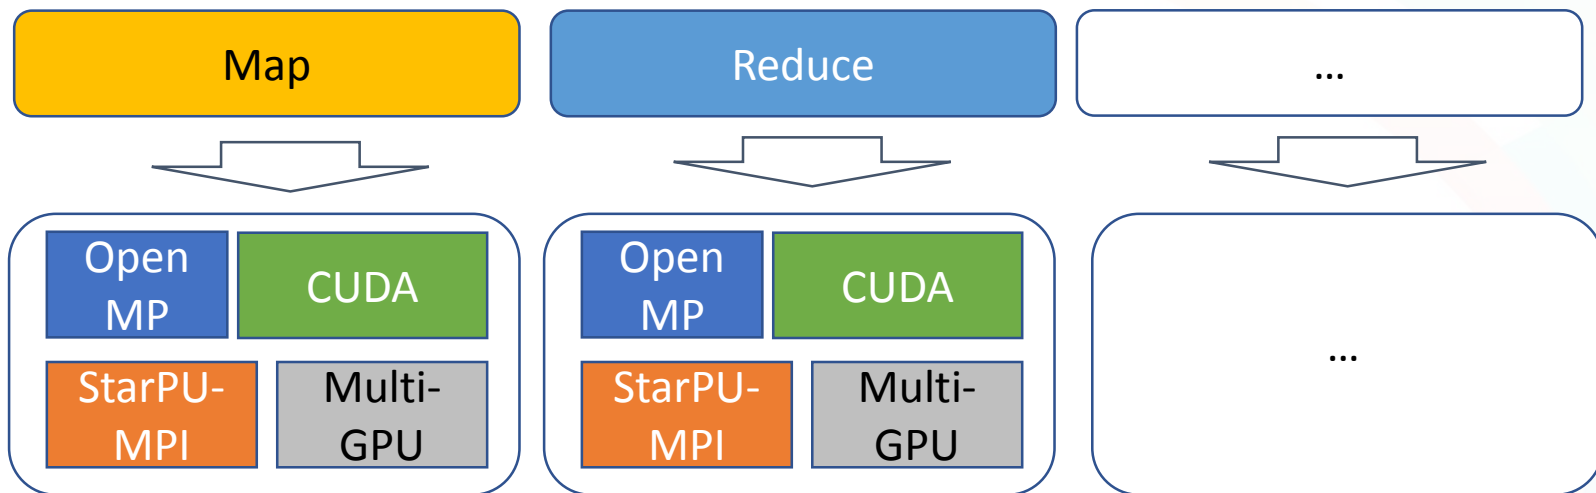
Promote **and lower the barrier of access to exascale computing systems** to scientific community and industry

How?

- Exascale systems are expected to be **heterogeneous**
- EXA2PRO improves the programmability of heterogeneous computing systems by **hiding their complexity** from application developers

Main Approach

- **Objective 1:** Follow a single-source multiple-backend approach

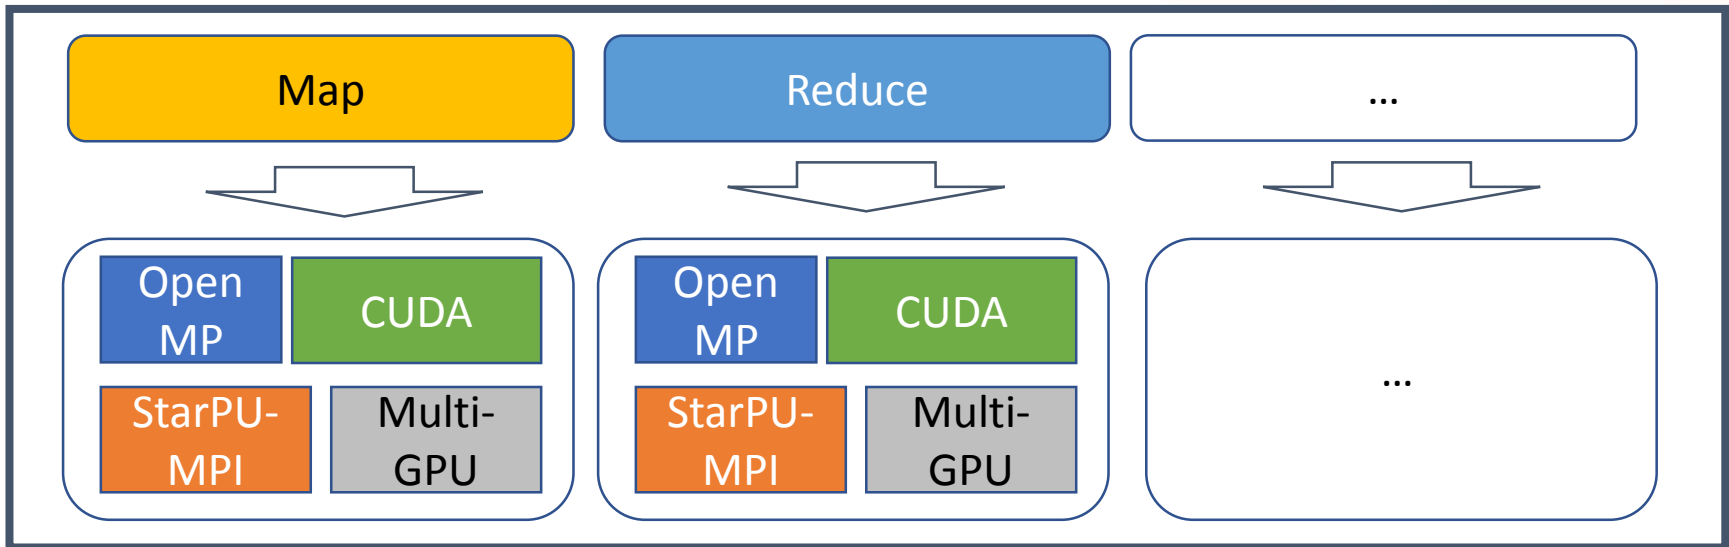


Patterns of computation implemented for a variety of backends (Skeletons)

Main Approach

- **Objective 1:** Follow a single-source multiple-backend approach

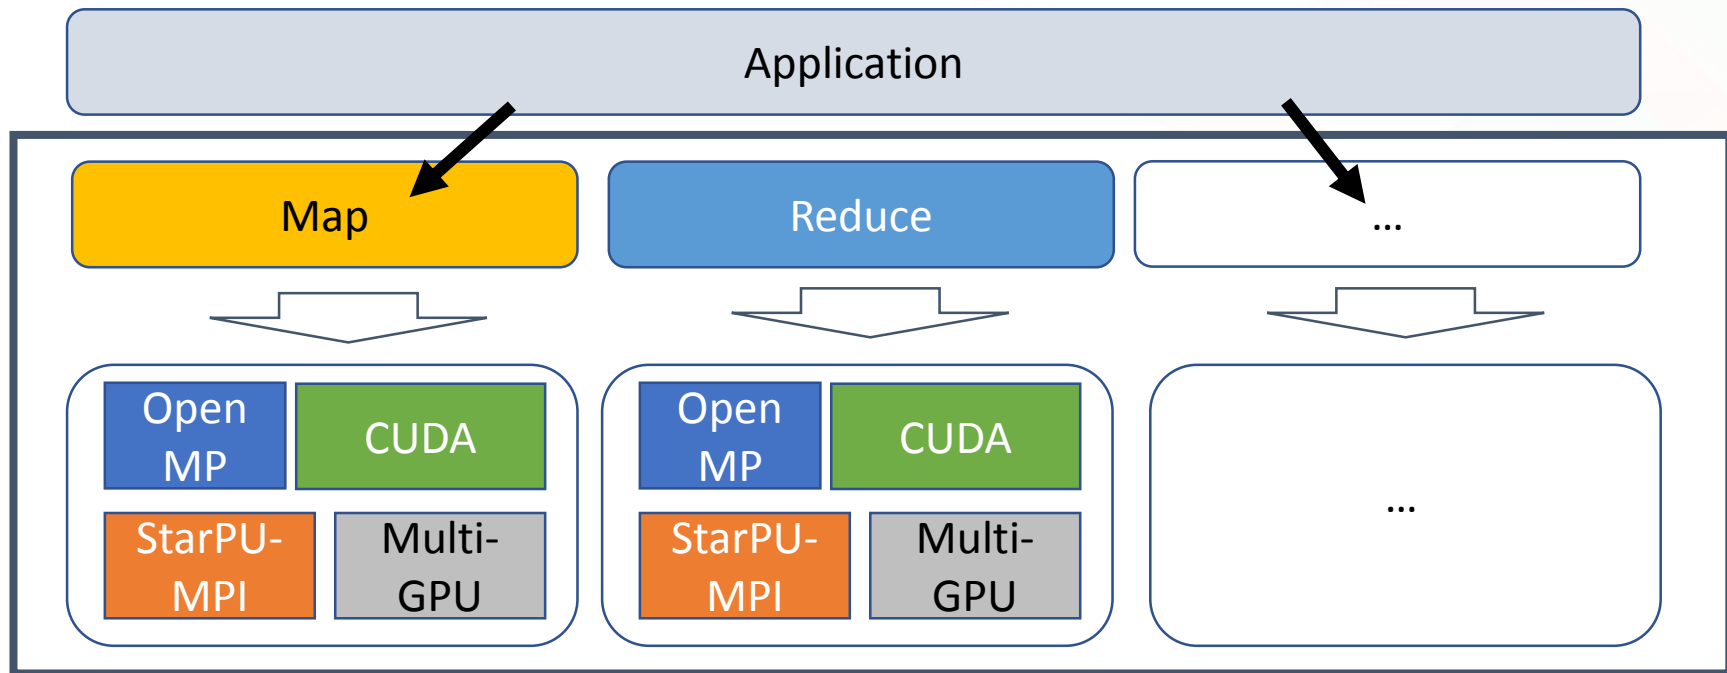
SkePU EXA2PRO high-level interface



Patterns of computation implemented for a variety of backends (Skeletons)

Main Approach

- **Objective 1:** Follow a single-source multiple-backend approach



- Patterns of computation implemented for a variety of backends (Skeletons)
- Developers are expected to implement computations of their application through the SkePU EXA2PRO tool

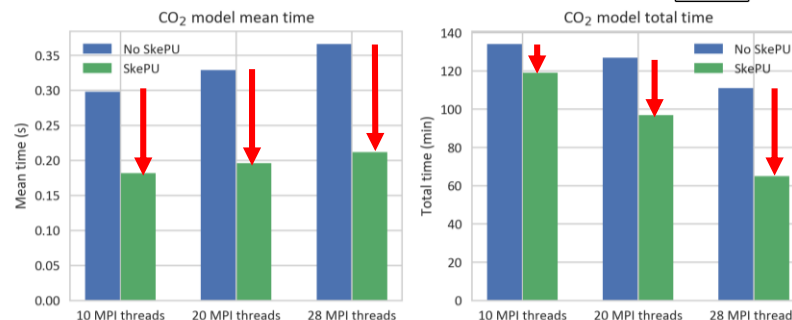
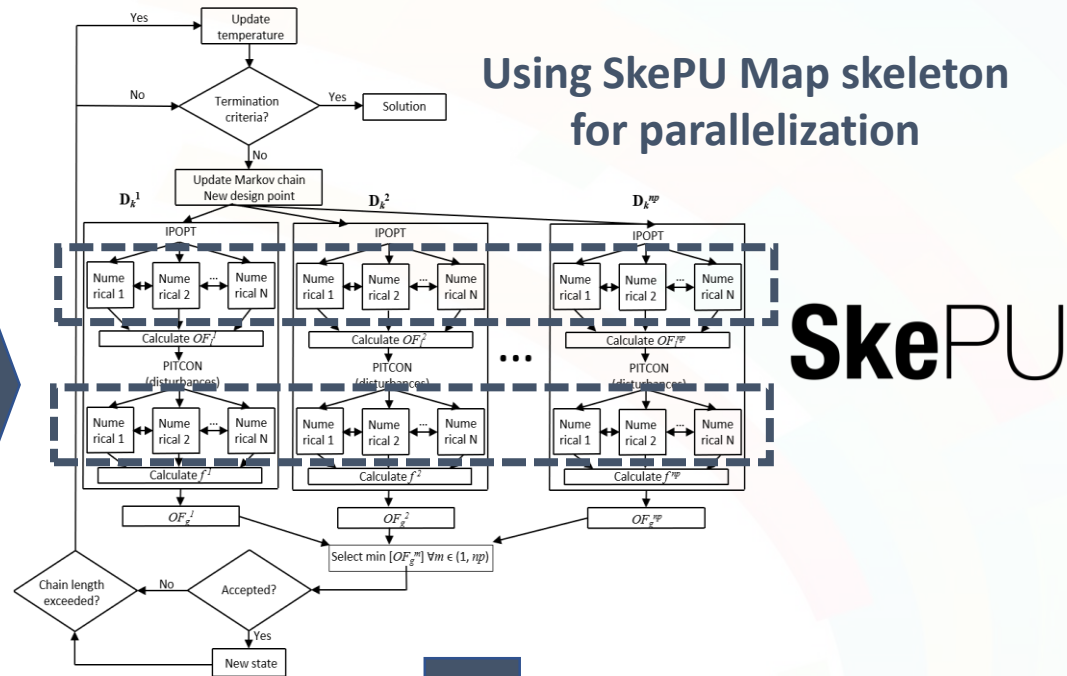
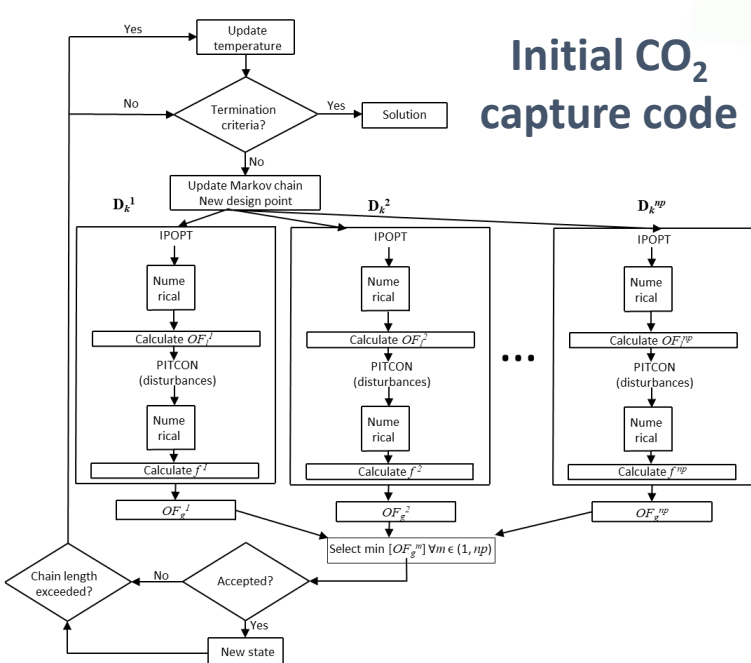
Main Approach

- **Objective 1: Single-source multiple-backend approach**
- **Advantages for application developers:**
 - Experiment with different platforms and accelerators (OpenMP, CUDA, OpenCL-enabled)

Main Approach

- **Objective 1: Single-source multiple-backend approach**
- **Advantages for application developers:**
 - Experiment with different platforms and accelerators (OpenMP, CUDA, OpenCL-enabled)
 - Can directly use the backend implementations already provided by EXA2PRO SkePU tool. → Productivity

Using EXA2PRO tools to advance the design of material and processes for CO₂ capture

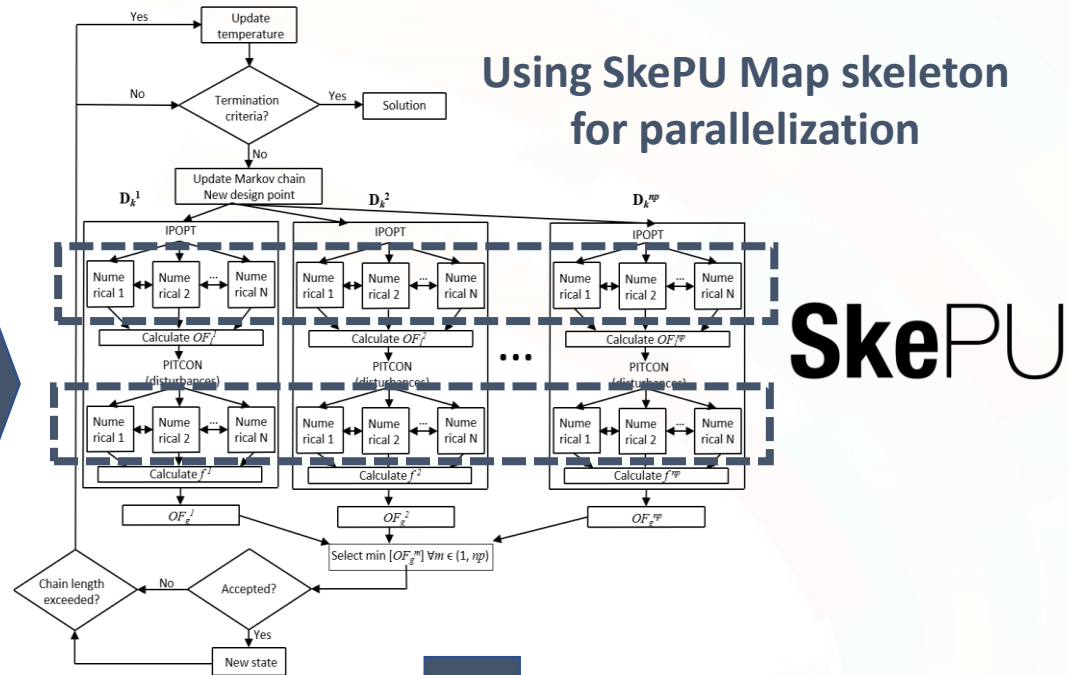
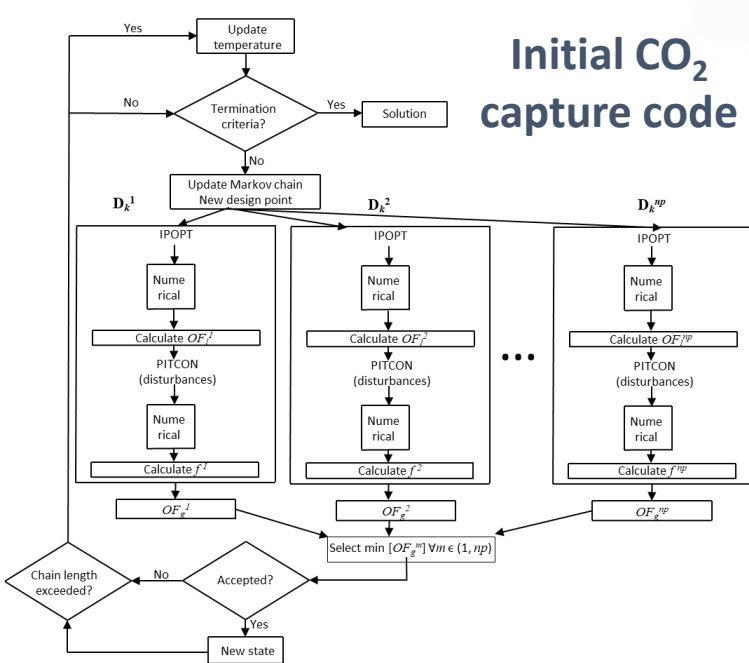


Savings up to 41% of total time by using SkePU version of CO₂ model

N. Vasilas, P. P. Natsiavas, A. Papadopoulos, P. Seferlis, "Process synthesis and controllability assessment of CO₂ capture plants in a parallel environment", In Proc. 22nd Conference on Process Integration, Modelling and Optimisation for Energy Saving and Pollution Reduction, 2019.

H2020, Project No. 801015

Using EXA2PRO tools to advance the design of material and processes for CO₂ capture

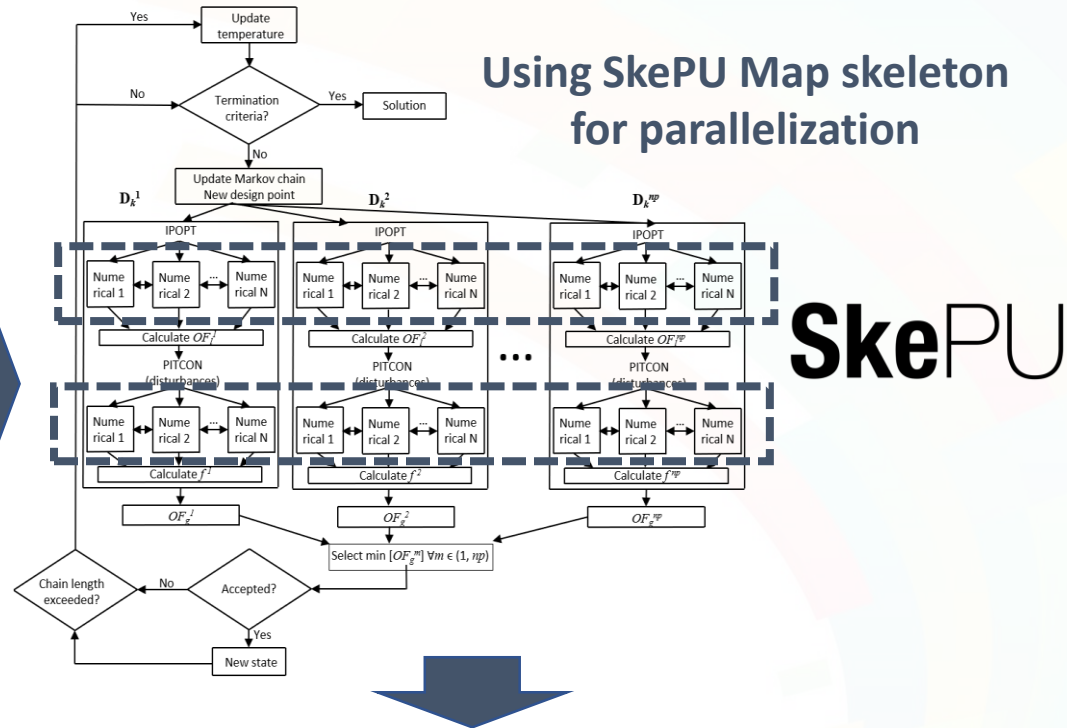
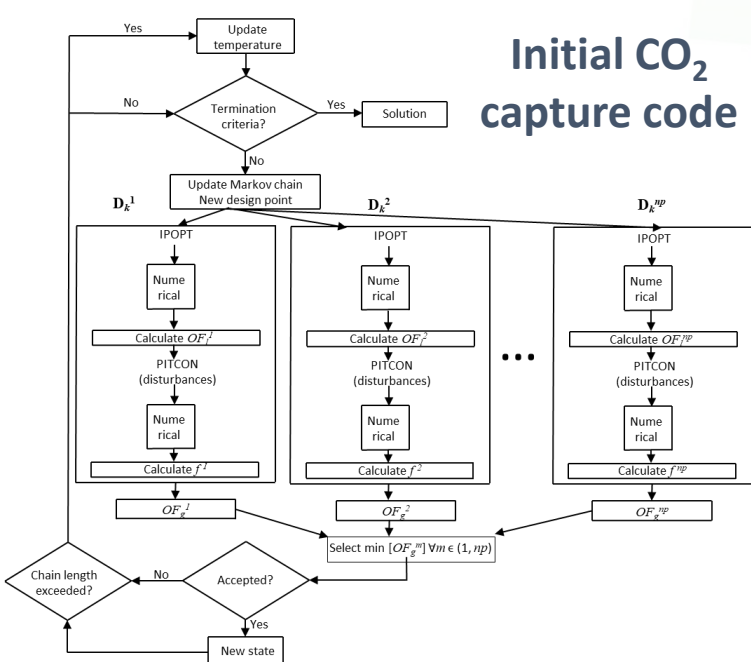


**Evaluation on GPU
(CUDA & OpenCL)**

N. Vasilas, P. P. Natsiavas, A. Papadopoulos, P. Seferlis, "Process synthesis and controllability assessment of CO₂ capture plants in a parallel environment", In Proc. 22nd Conference on Process Integration, Modelling and Optimisation for Energy Saving and Pollution Reduction, 2019.

H2020, Project No. 801015

Using EXA2PRO tools to advance the design of material and processes for CO₂ capture



Accelerators more accessible to developers → transition to systems that use accelerators

Evaluation on GPU (CUDA & OpenCL)

N. Vasilas, P. P. Natsiavas, A. Papadopoulos, P. Seferlis, "Process synthesis and controllability assessment of CO₂ capture plants in a parallel environment", In Proc. 22nd Conference on Process Integration, Modelling and Optimisation for Energy Saving and Pollution Reduction, 2019.

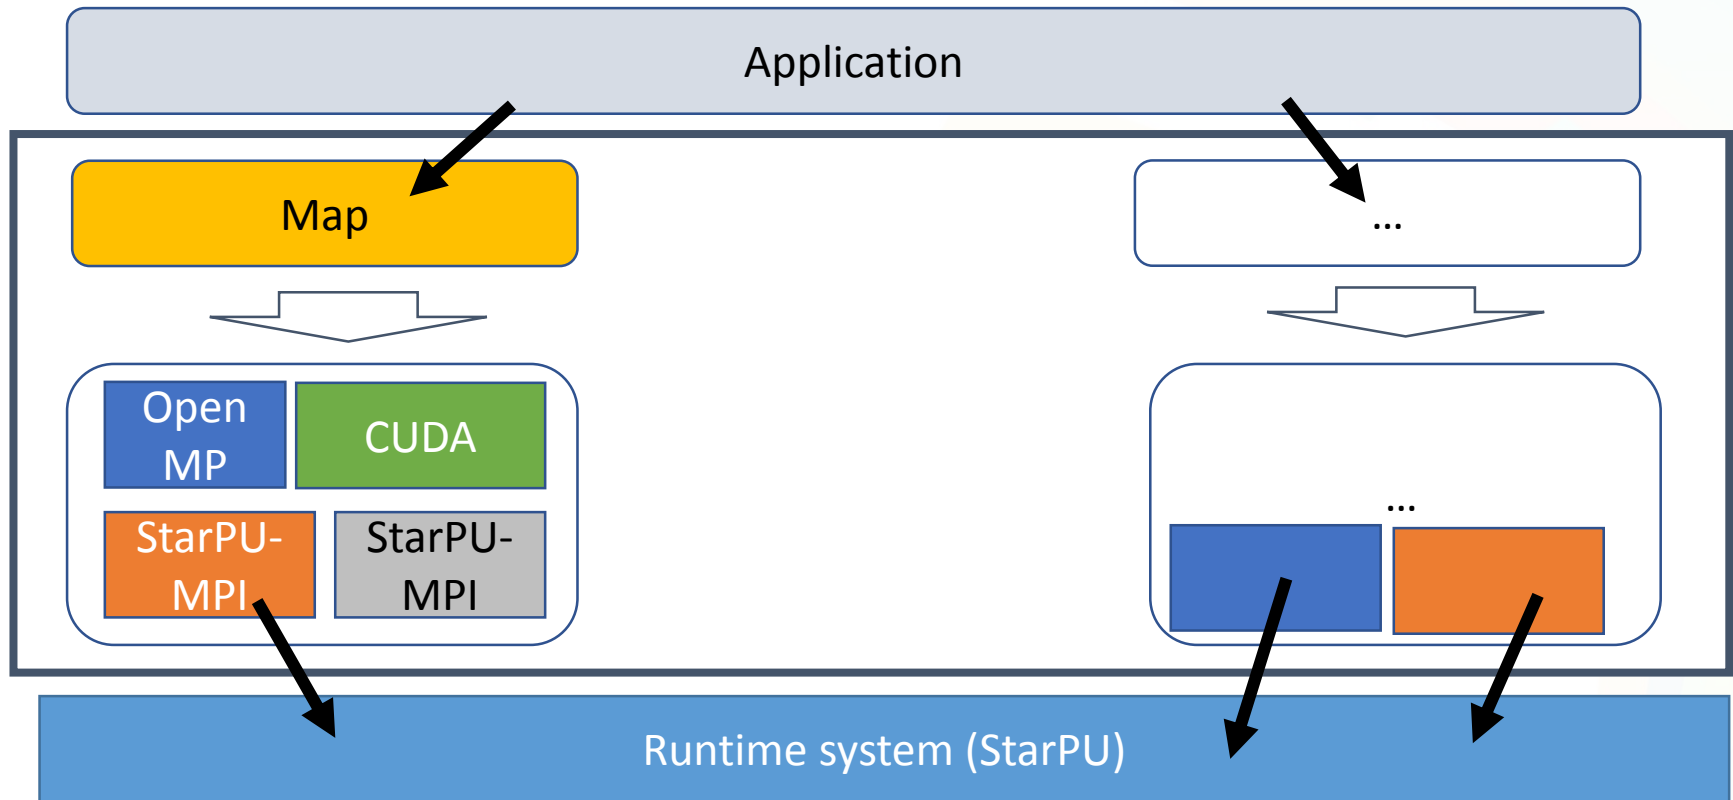
H2020, Project No. 801015

Main Approach

- **Objective 1: Single-source multiple-backend approach**
 - **Advantages for application developers:**
 - Experiment with different platforms and accelerators (OpenMP, CUDA, OpenCL-enabled)
 - Can directly use the backend implementations already provided by EXA2PRO SkePU tool. → Productivity
- **The effort to apply the EXA2PRO interface is lower when:**
 - The application is developed in C/C++
 - The computation kernel is well-isolated

Main Approach

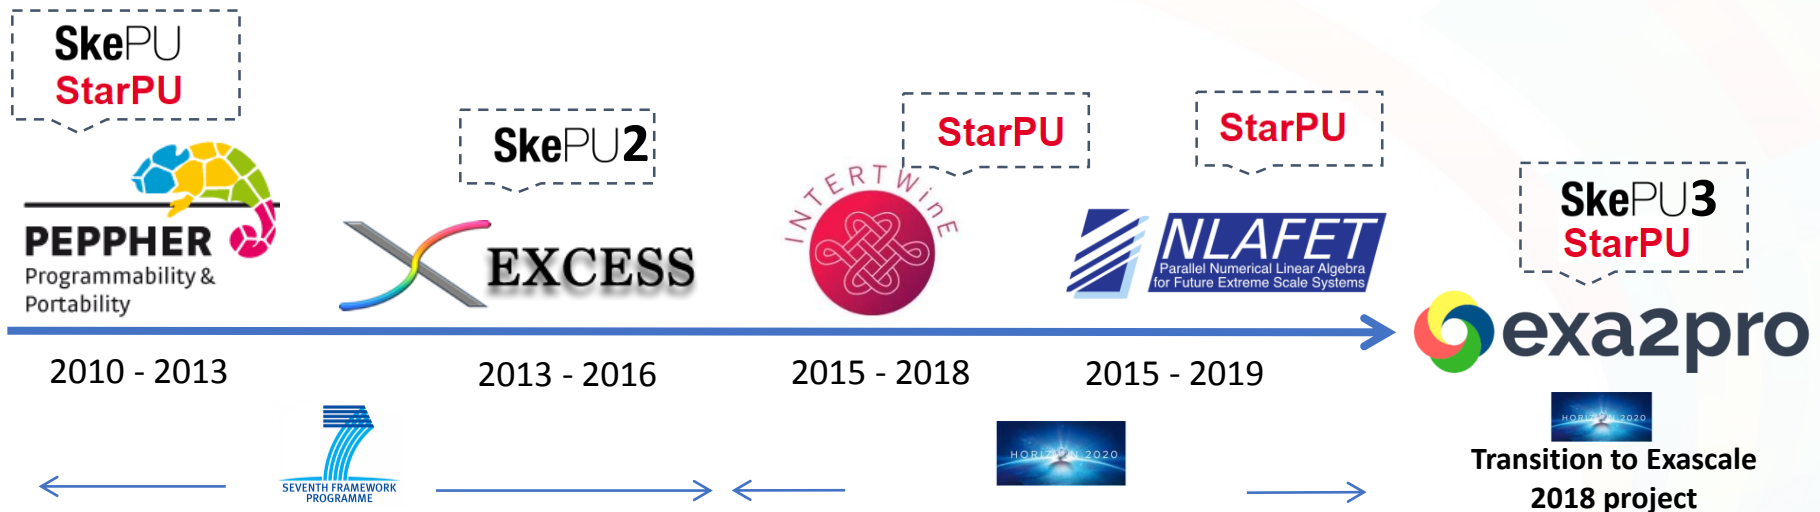
Objective 1: Follow a single-source multiple-backend approach



- Elimination of suboptimal or non-applicable backend implementations at design time
- The runtime system selects the most efficient implementation at runtime considering performance and energy as optimization criteria

Main Approach

Objective 2: Co-design approach between tools and applications



- Tools such as SkePU and StarPU evolve towards exascale computing

Multi-node neural simulations using SkePU

Scaling demonstration across multiple nodes using MatRow SkePU data container (Tetralith cluster: 32 nodes used, 32 cores per node)

Original neural simulation code:

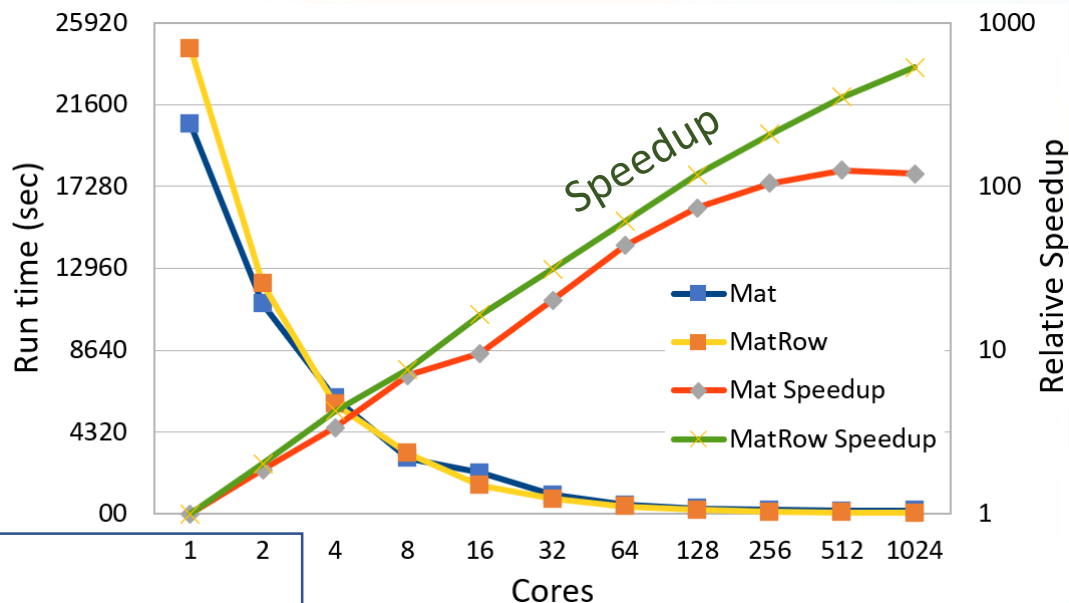
```
for each timestep:
  for each neuron i:
    Get Isyn from adj. Neurons
    ... internal dynamics code ...
    Compute stateNext[i]
```



SkePU-ported neural simulation code:

```
Function advanceState(i, stateNow, conn_matrix):
  Get Isyn from adj. Neurons
  ... internal dynamics code ...
  Return stateNext for neuron i

Kernel = Map(get_stateNext)
for each timestep:
  stateNext = Kernel(stateNow, conn_matrix)
```



SkePU

<https://skepu.github.io/>

Presented in **SCOPES 2020** (25-26 May 2020)

<https://scopesconf.org/scopes-20/>

Main Approach

- **Objective 3:** Evaluation the data-flow programming model
 - Used to program the FPGA-based MAXELER data-flow engines (DFEs)



Results: Metalwalls (Simulator for supercapacitors)

Initial results of porting Metalwalls on DFE:

4x performance gain

10x lower energy consumption

50% higher performance/watt than P100 GPU

Results: Metalwalls (Simulator for supercapacitors)

Initial results of porting Metalwalls on DFE:

4x performance gain

10x lower energy consumption

50% higher performance/watt than P100 GPU

Development time: 8 – 12 months → Tools to improve productivity for accessing accelerators

Results: Metalwalls (Simulator for supercapacitors)

Initial results of porting Metalwalls on DFE:

4x performance gain

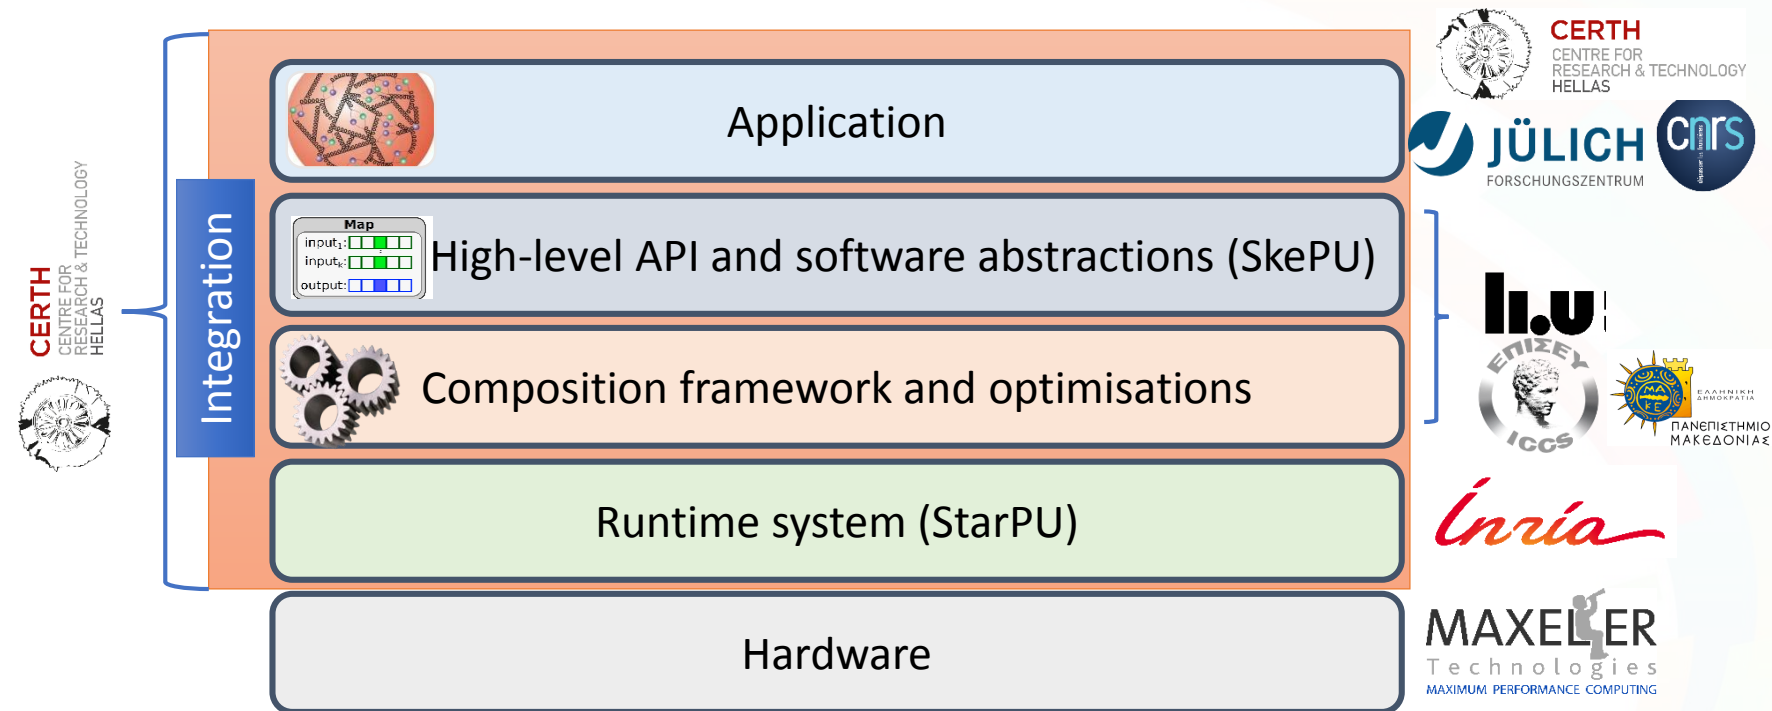
10x lower energy consumption

50% higher performance/watt than P100 GPU

Development time: 8 – 12 months → Tools to improve productivity for accessing accelerators

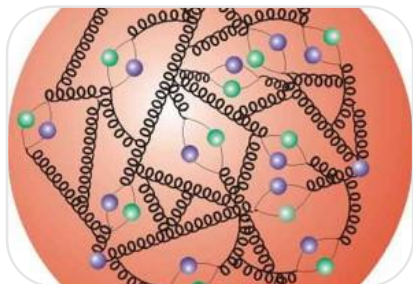
Next step: Use Metalwalls with StarPU for multi-DFE experiments

The EXA2PRO stack



EXA2PRO applications

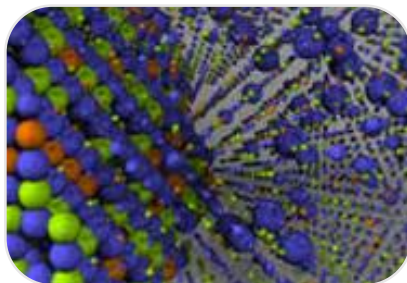
High energy physics



LQCD

- JUELICH
- FORTRAN/C++

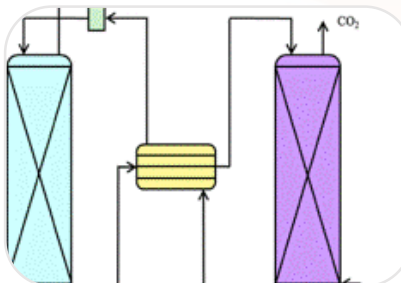
Materials science



KKRNano

- JUELICH
- FORTRAN

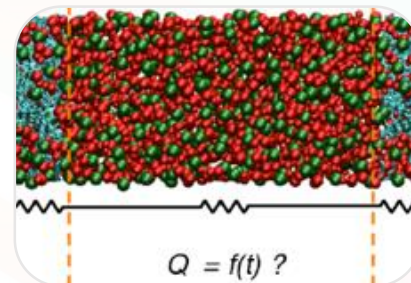
Materials and Processes



Materials and processes for CO₂ Capture

- CERTH
- FORTRAN

Energy



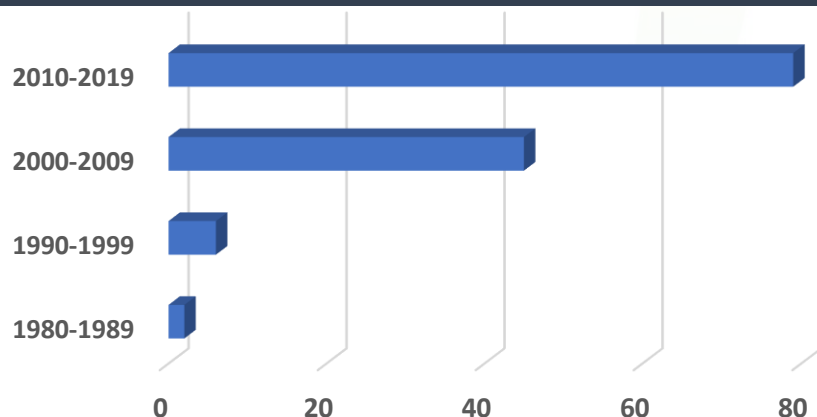
MetalWalls
(Supercapacitors simulator)

- CNRS
- C++

Applications deployed in pre-exascale systems targeting exascale systems

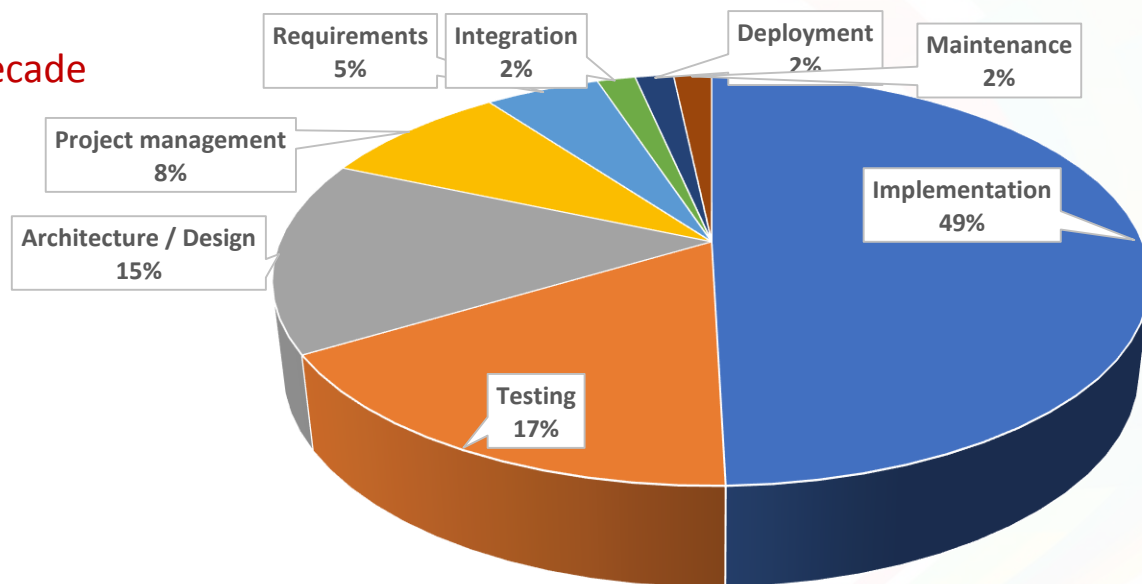
Applications deployed in HPC clusters
→ Target large-scale systems

Other aspects: Software engineering practices in HPC



Number of studies per decade

*Studies in the **HPC domain** that employ **Software Engineering** practices*



Number of studies per development phase

In the EXA2PRO, we developed a tool for improving the maintainability of HPC applications, supporting C/C++ and FORTRAN

<https://exa2pro.eu/#developers>

Conclusions and lessons-learned

- Tool support for accessing accelerators can significantly improve the productivity of application developers
- Interaction between tool developers and application providers was very constructive
 - Contributed to the maturity of the tools and to the application performance improvement
- Data-flow programming is difficult but initial results are very promising

Thank You!

lpapadop@microlab.ntua.gr

 [@exa2pro_h2020](https://twitter.com/exa2pro_h2020)



This project has received funding from the European Union's Horizon 2020 Research and Innovation programme under grant agreement No. 801015

Partners

ICCS, LIU, CETH, INRIA, JUELICH, MAXELER, CNRS, UoM