Dataflow based Rapid Prototyping activities

Jean-François NEZAN - Ophélie RENAUD

SEAM
Sustainable & Energy Aware Methods for SKA

DARK ERA

Rising STARS

Karol DESNOS	Mickael DARDAILLON	Nicolas GAC	Jacques MORIN

Square Kilometre Array (SKA)

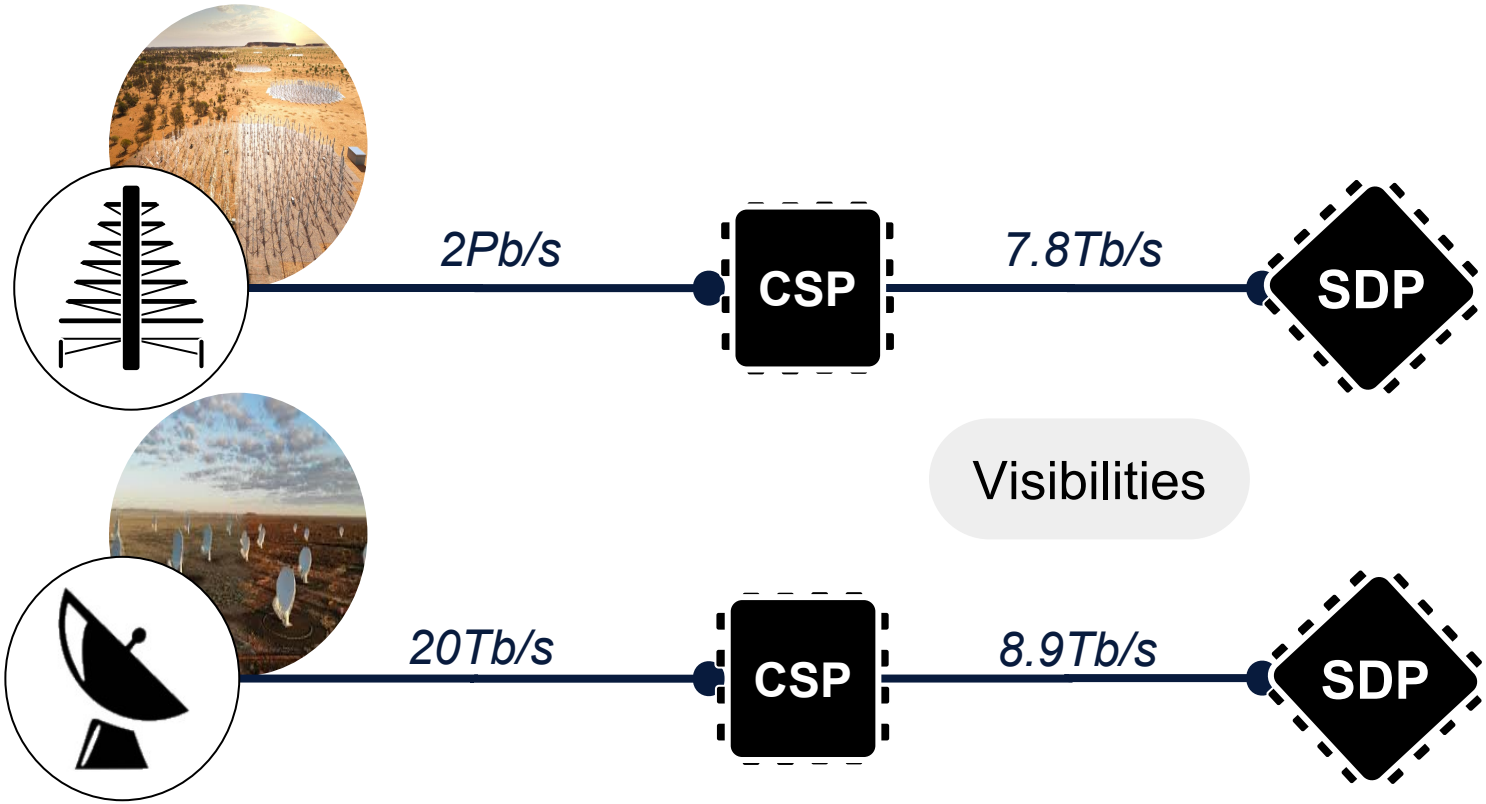

Rapid Prototyping with PREESM

Save Time and Money

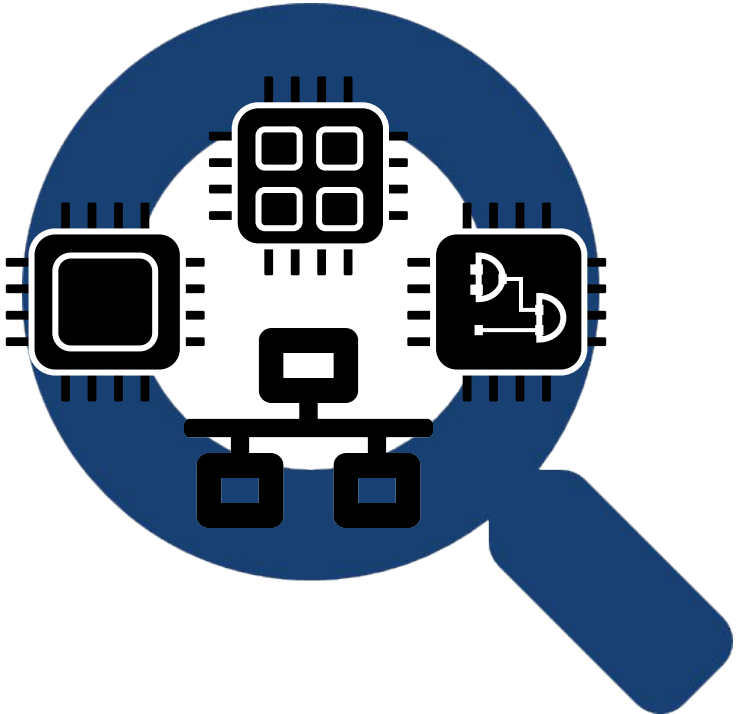

Graph based Algorithm-Architecture Adequation (AAA)

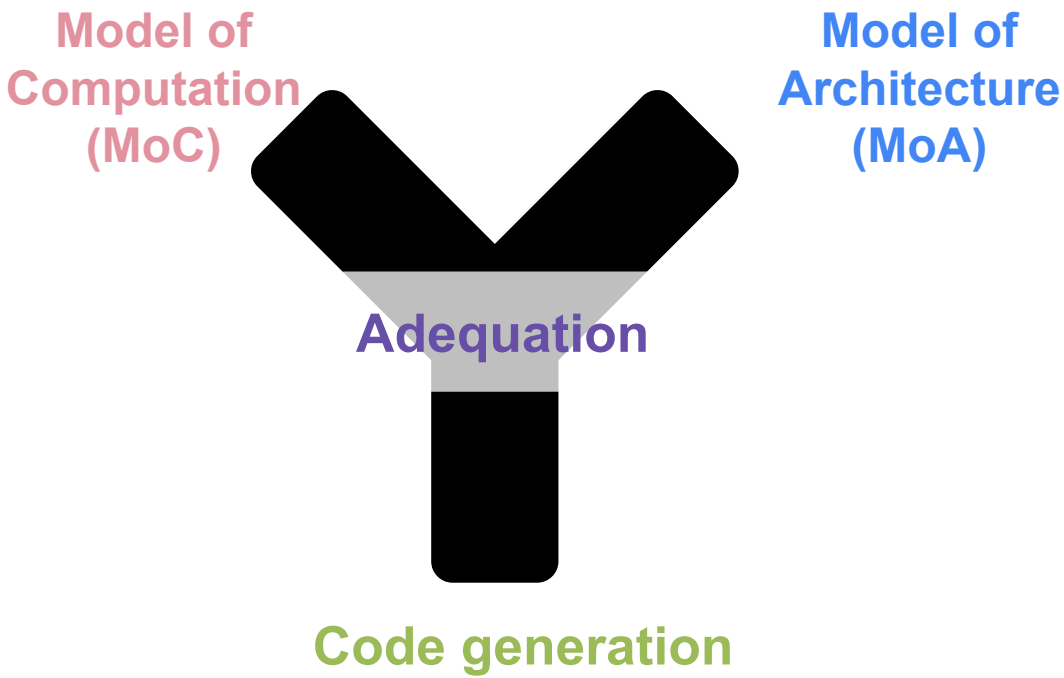

[r] C. Erbas, et. al, « Multi objective optimization and evolutionary algorithms for the application mapping problem in multiprocessor system-on-chip design »

Dataflow MoC

✓ Expression of several types of parallelism

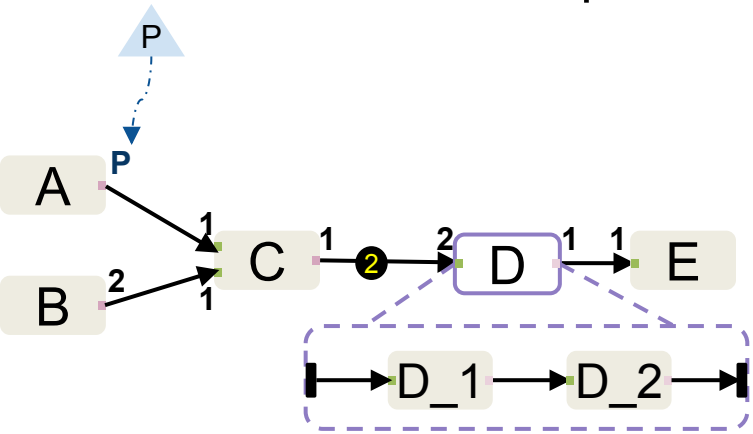

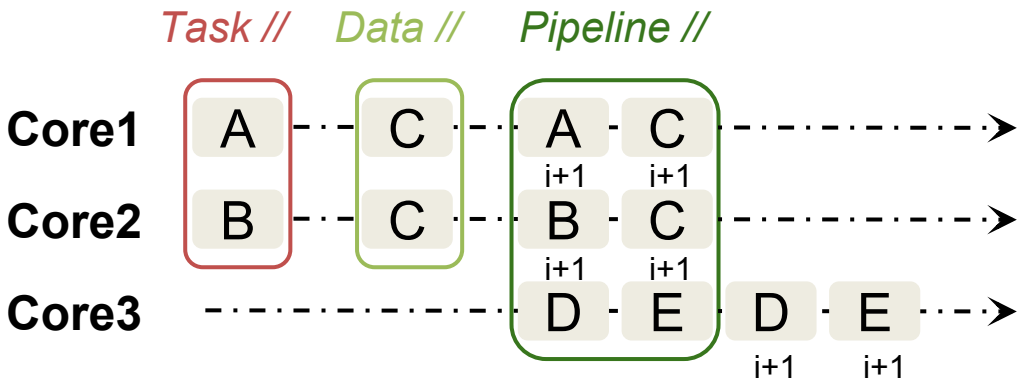

✓ Architecture Independence
✓ Prédictability

✓ Encapsulation (Hierarchy)
✓ Flexibility (Parameters)

SDF: Synchronous Dataflow
PiSDF: Parameterized and Interfaced SDF

[r] E. Lee and D. Messerschmitt. "Static scheduling of synchronous data flow programs for digital signal processing".
[r] K. Desnos et. al. "Pimm : Parameterized and interfaced dataflow meta-model for mpsocs runtime reconfiguration"

S-LAM MoA

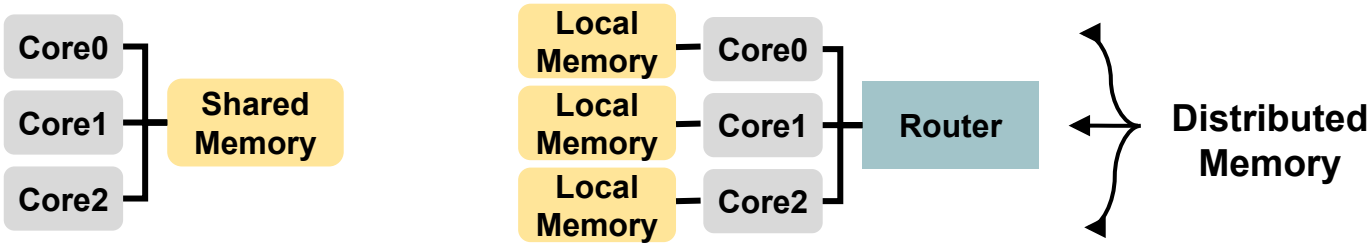

At least 2 types of memory architecture modeling

S-LAM: System-Level Architecture Model
[r] Maxime Pelcat, et. al, « A system-level architecture model for rapid prototyping of heterogeneous multicore embedded systems »

Adequation - Standard dataflow resource allocation

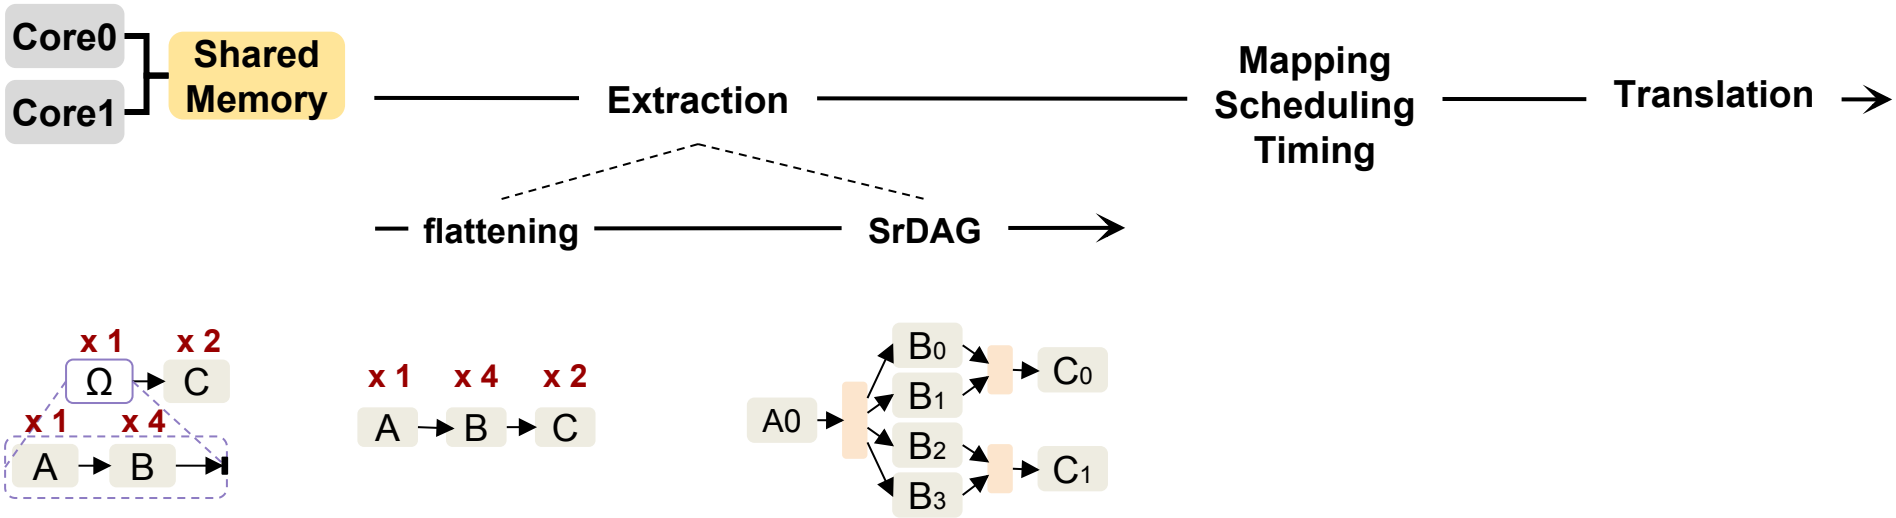

Adequation - Standard dataflow resource allocation

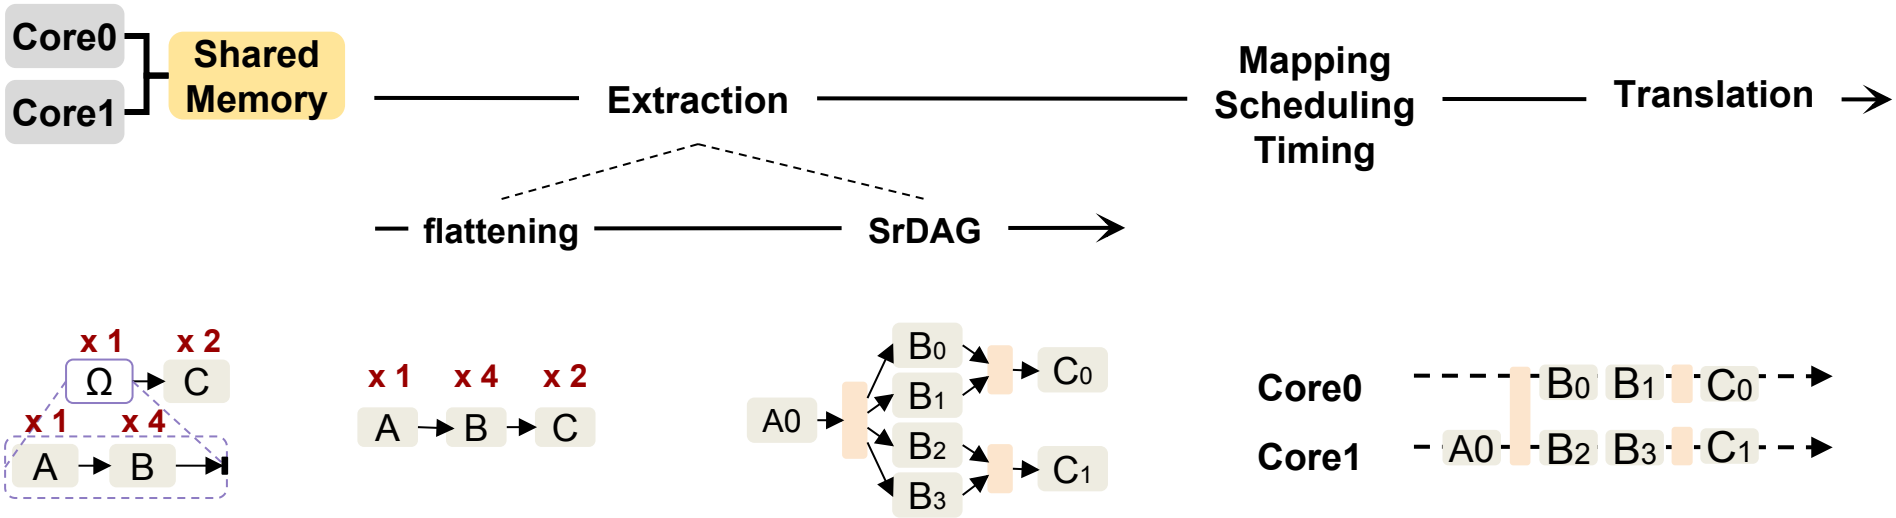

Adequation - Standard dataflow resource allocation

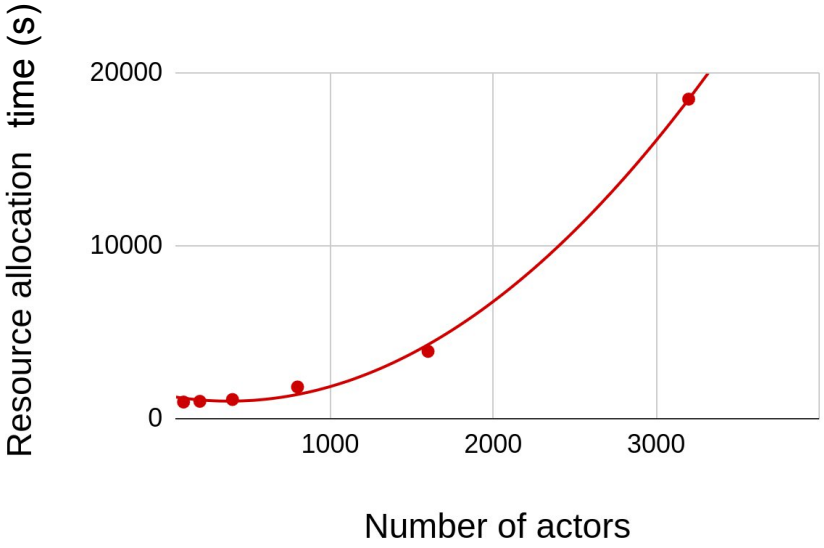

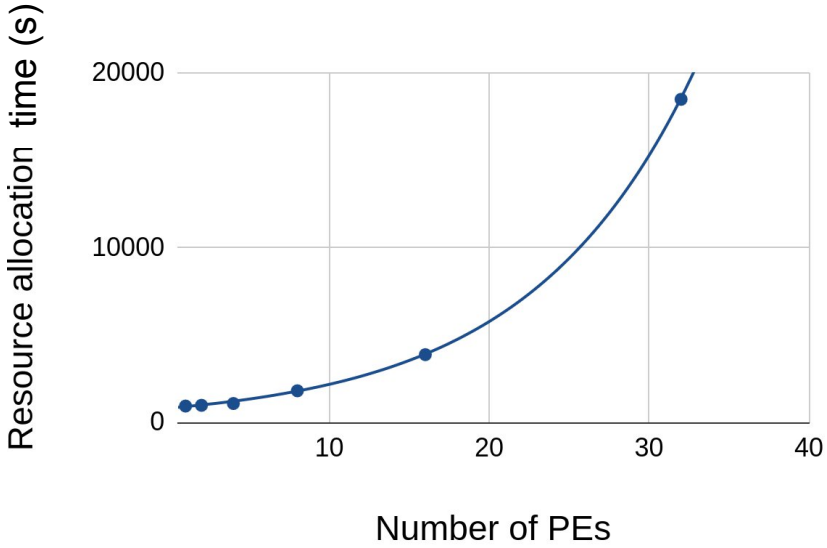

: O(**A**.log(A)+**P**.(A+E))

DARK-ERA achievements

Dataflow modeling : Radio-Interferometric Imaging Pipeline

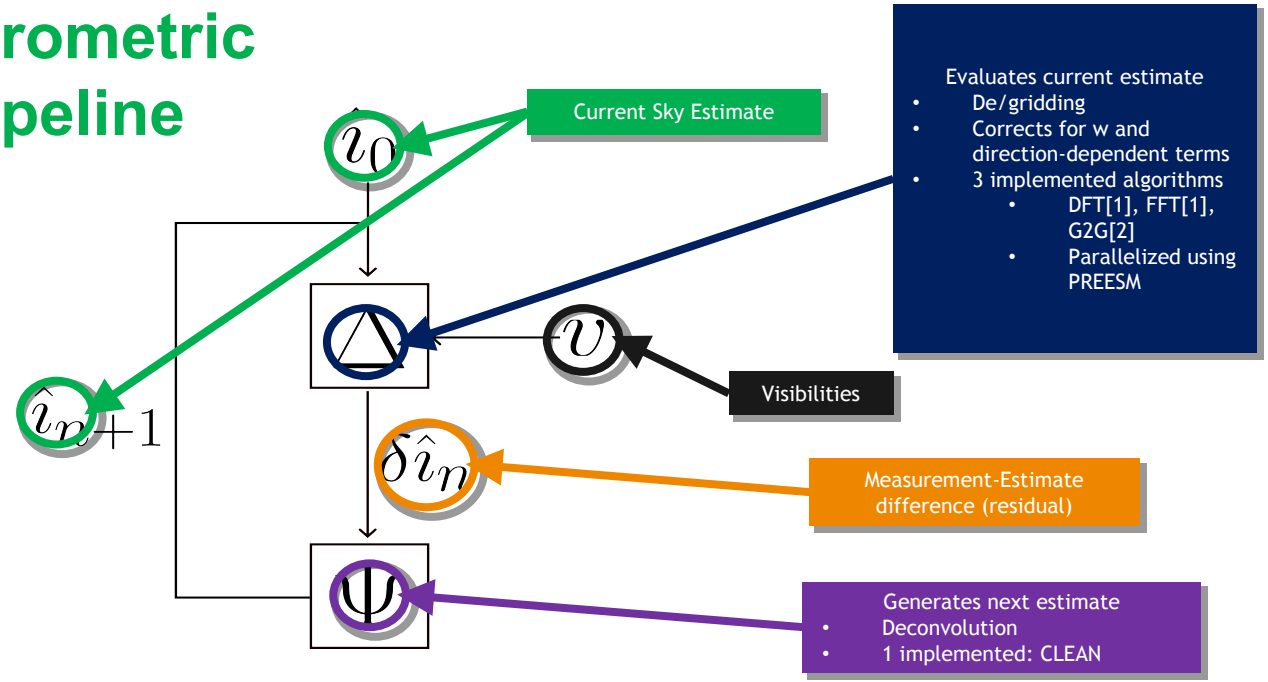

DARK-ERA achievements

Split/Merge for automatic parallelism

Generic dataflow model of the Imaging Pipeline

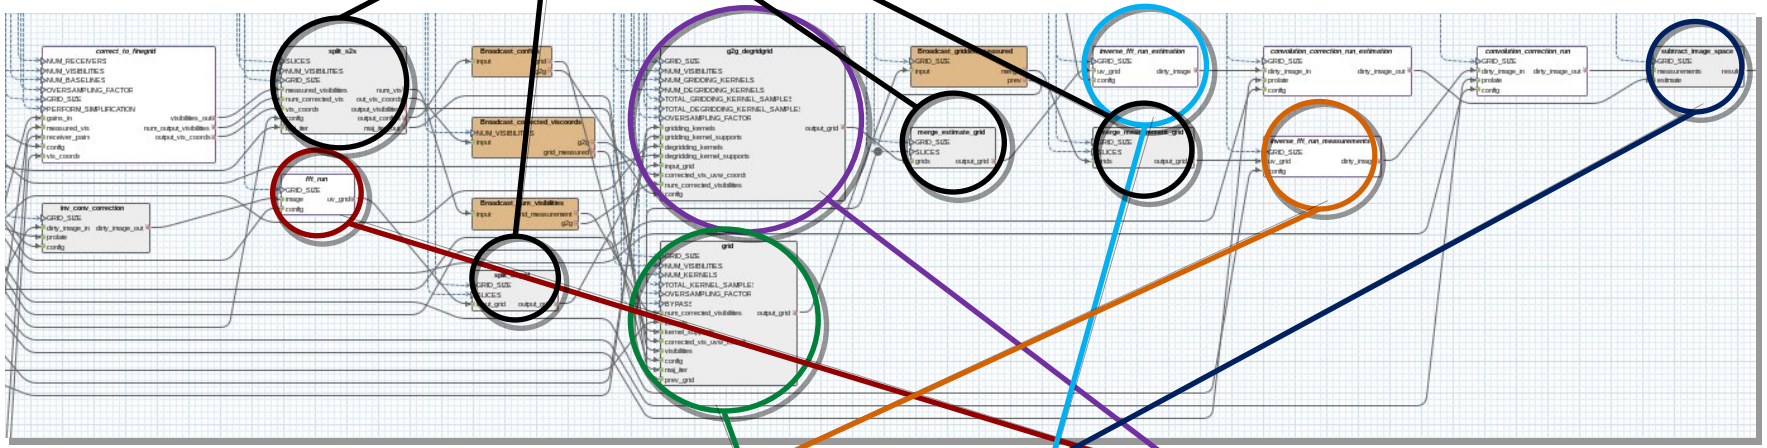

$$\delta\hat{i}_n = F^\dagger G^\dagger v - F^\dagger G^\dagger G F \hat{i}_n$$

DARK-ERA achievements

- Application modeling : SEP, DifX correlator
- Complexity decrease : Granularity optimization for multicore architecture
- Clustered memory architectures :
  - Multi-node HPC resource allocation
  - S-LAM Model of HPC Network
- Code generation
  - Grid 5000
  - GPU

l'Observatoire de Paris | PSL

OBSERVATOIRE DE LA CÔTE D'AZUR
UNIVERSITÉ CÔTE D'AZUR

IETR

EPFL EcoCloud Embedded Systems Laboratory

From Energy efficient compute technologies

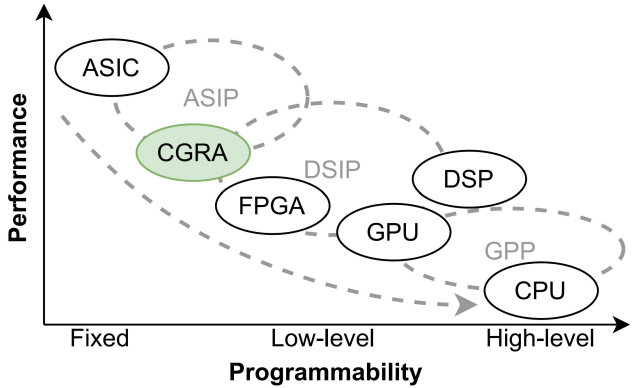

To Low Power HPC

Integration of power-efficient accelerators in HPC systems

- **User Interface**
  - Graphical Interface, C, Python, OpenMP, etc.
- **Models of Computations**
  - Adapted to the HW-SW co-design step
- **Design Space Exploration**
  - Early design strategies
- **Ressource allocation & code generation**
  - Fast and efficient HPC SW integration

Multi-scale co-design process

The expected results of SEAMS will include:

- the design of accelerators optimised to efficiently process complex SKA pipelines
- new methods for scheduling continuous execution at low power and high throughput.

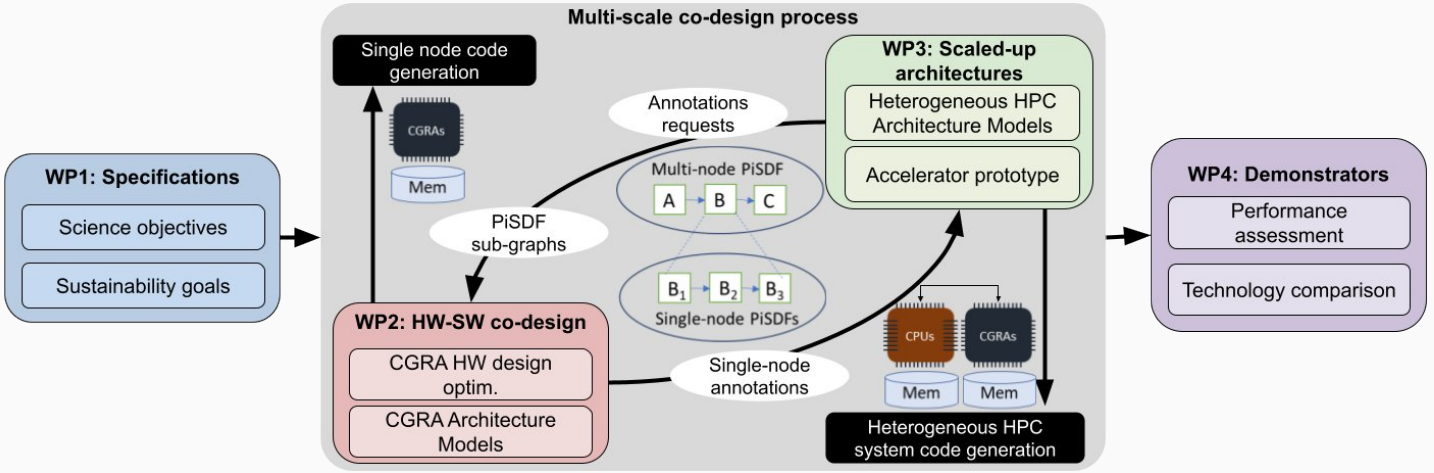

Others activities

- Embedded AI
- Design tools interoperability : ONNx, MLIR, LLVM
- SPML : Signal Processing and Machine Learning

INSA IETR VAADER DARK ERA

PhD contributions

Model Based Granularity Optimization for High Performance Computing Systems in Astronomy

Defended on October 9, 2024

Supervised by Jean-François Nezan & Karol Desnos

Outline

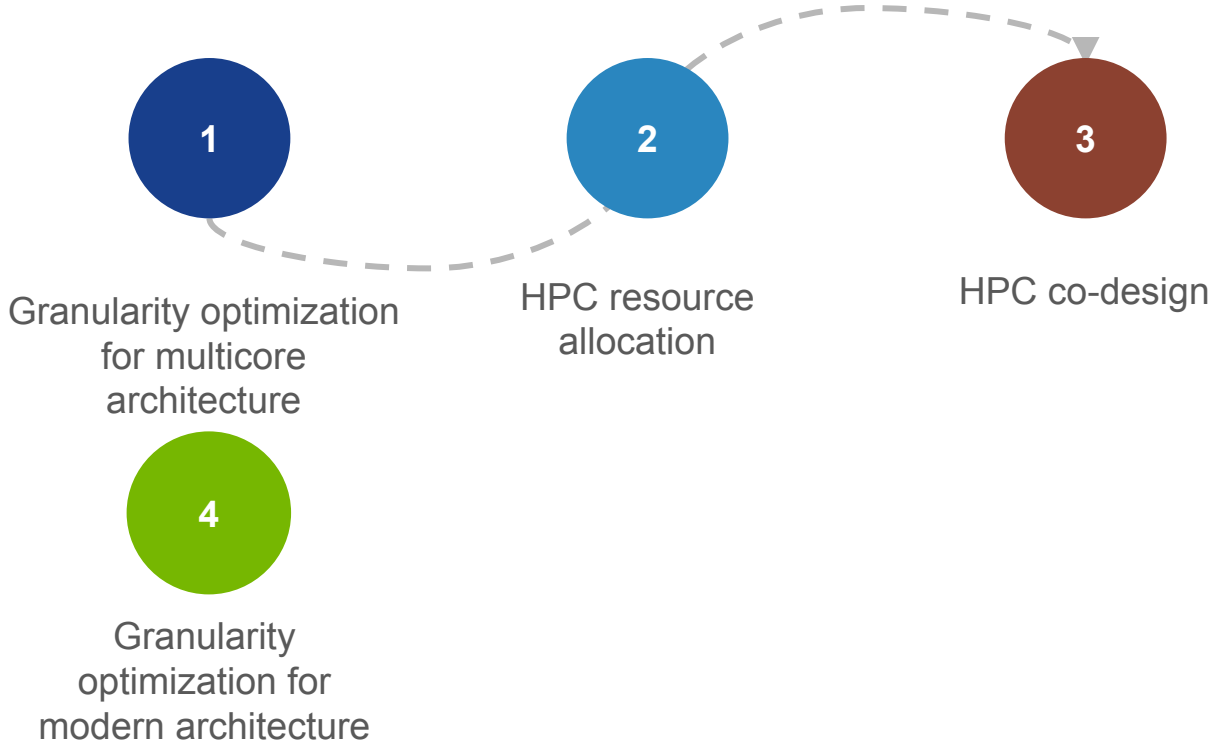
1
2
3
Granularity optimization for multicore architecture
HPC resource allocation
HPC co-design
4
Granularity optimization for modern architecture

Motivation

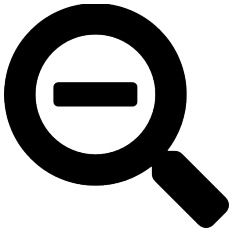

Resource Allocation Time

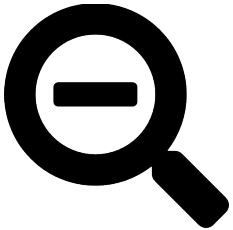

Execution Time

Clustering to reduce graph complexity

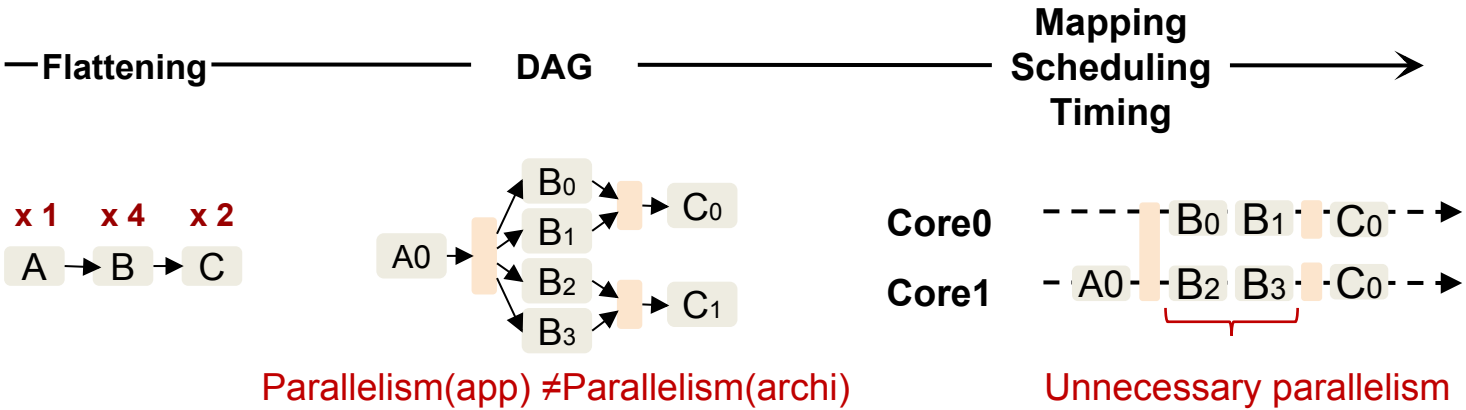

Grain-Adaptive Dataflow Resource Allocation

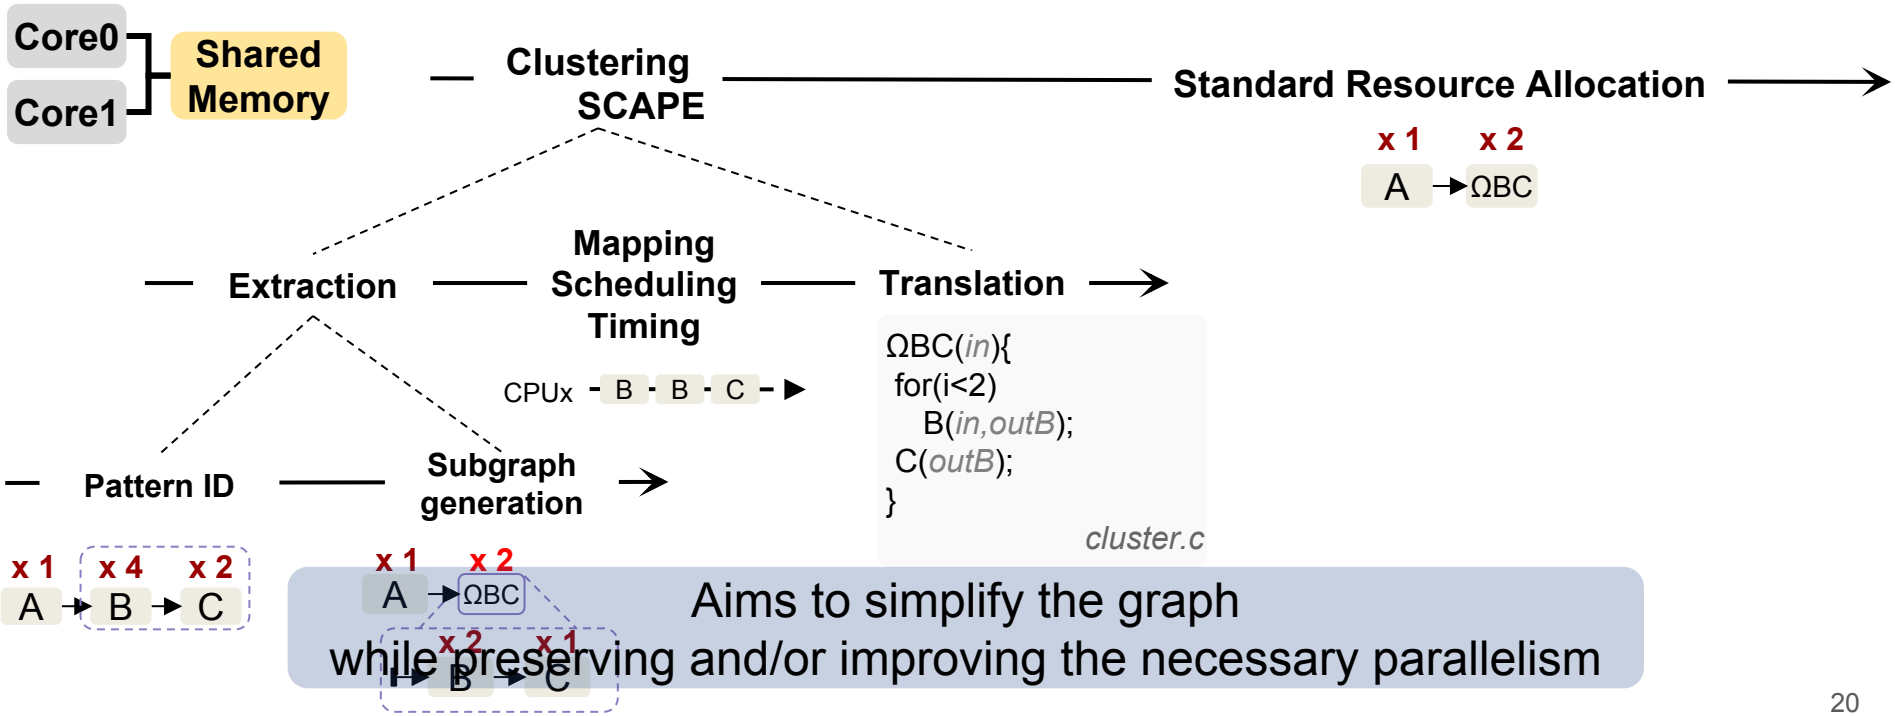

SCAPE: Scaling up of Clusters of Actors on Processing Element

Comparative Analysis: SCAPE Methods vs. Without on Stereo Application Deployment

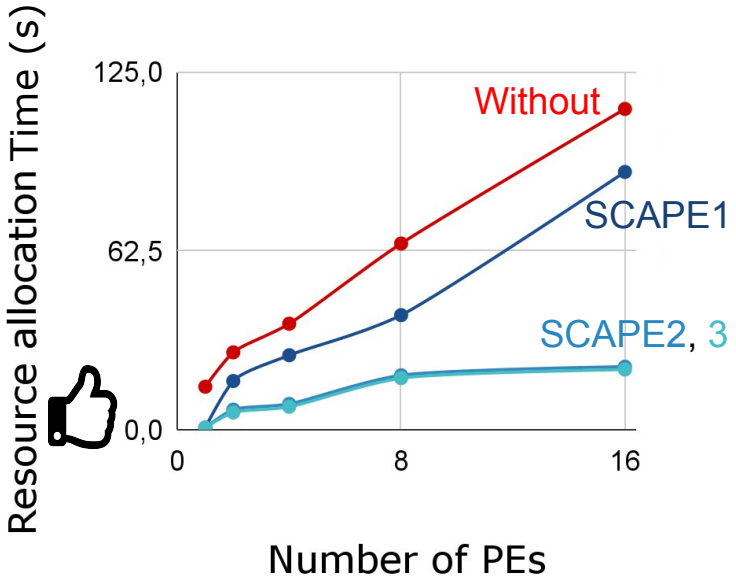

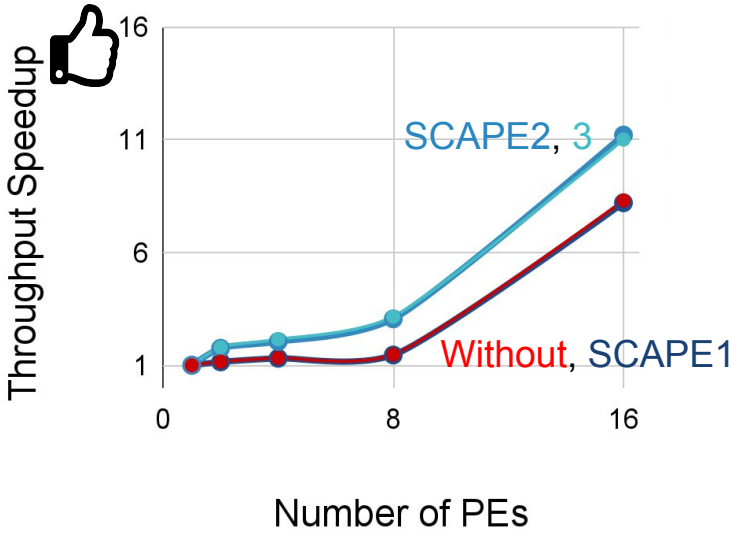

SCAPEs always faster than without
SCAPE2 = SCAPE3 → same configuration

SCAPE1 = without → keep data //
SCAPE2-3 > without → add pipeline //

Motivation

Collaboration

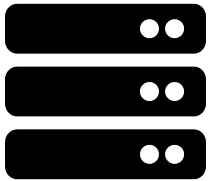

HPC Multi-Node Resource Allocation

Multi-node S-LAM modeling

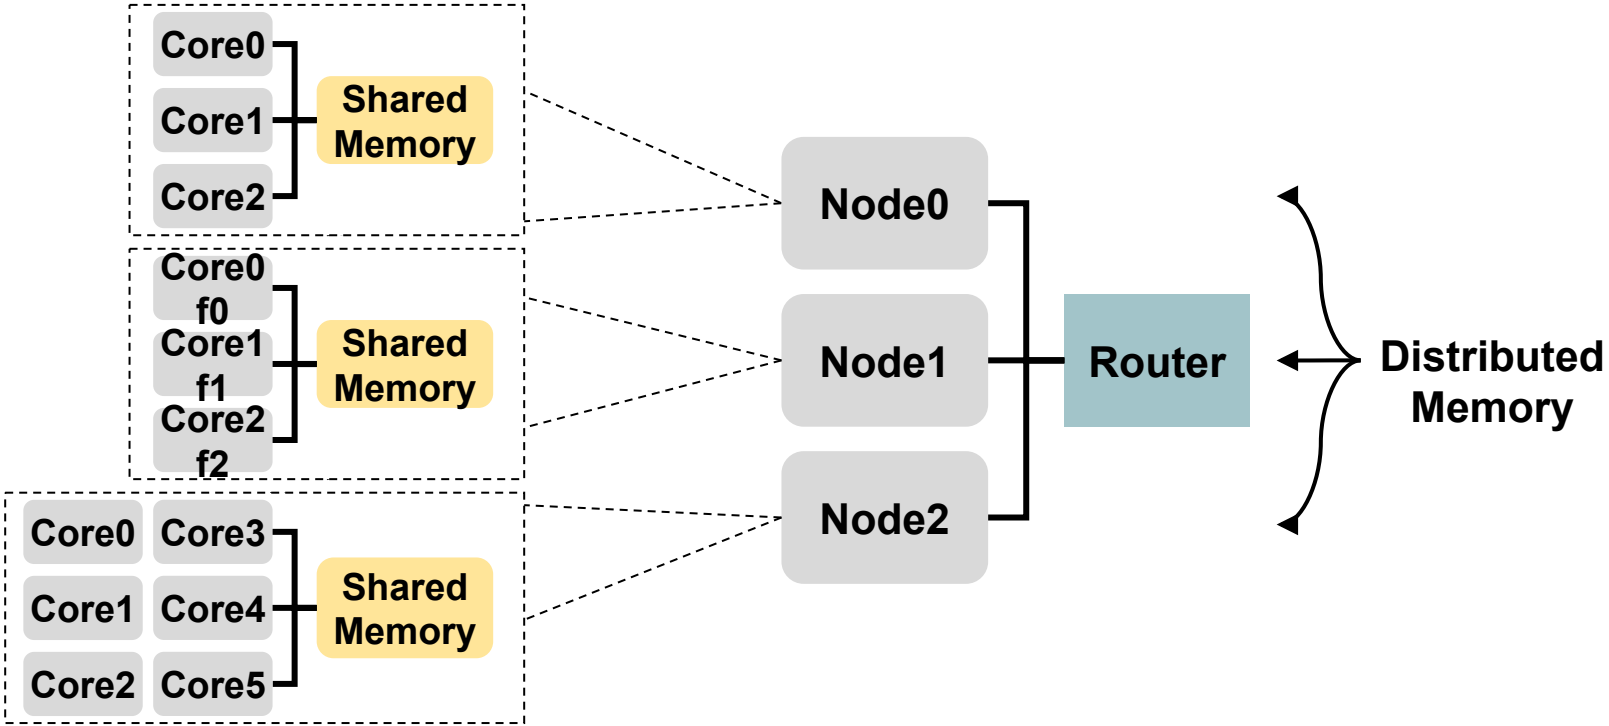

Multi-Node CPU resource allocation

Dataflow

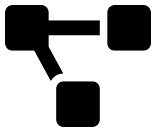

S-LAM

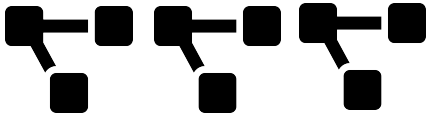

SimSDP

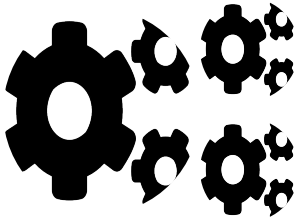

Balanced-workload
Between nodes

Code

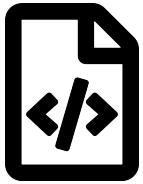

Multi-Node CPU resource allocation

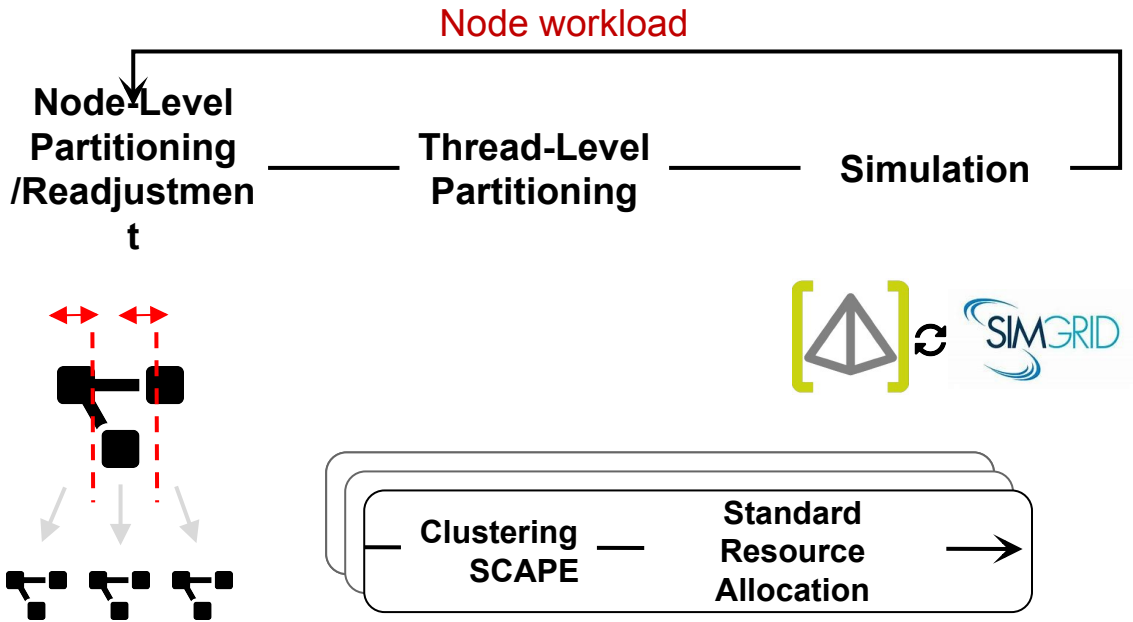

Iterative because 2 levels of resource allocation: within nodes and on the overall system

Iterative impact of the method over various granularity

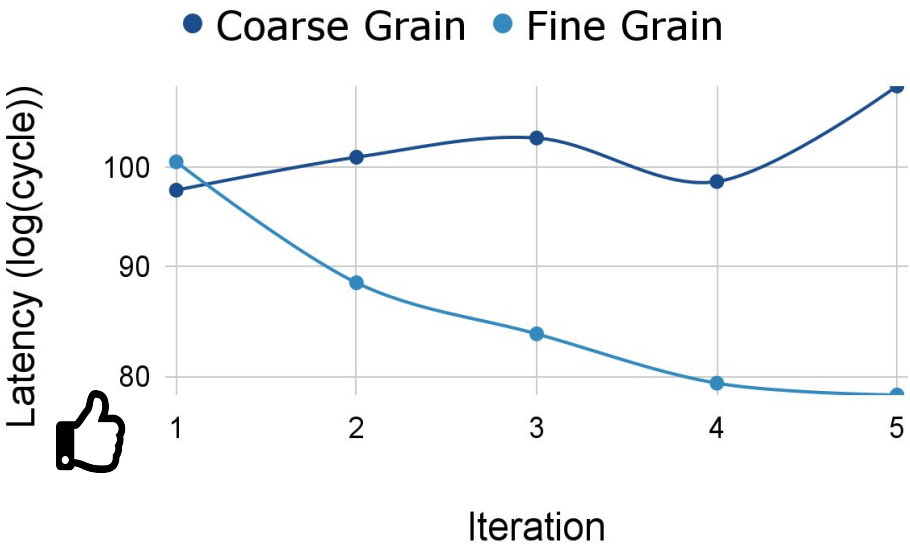

T(com)<<T(A) → directly reach optimum
T(com)>>T(A) → iteration relevancy

***O. Renaud**, A. Gougeon, K. Desnos, C. Phillips, J. Tuthill, M. Quinson, J.-F. Nezan. SimSDP: Automatic Workload-Balancing on Multi-Node & Multi-Core HPC Architectures based on dataflow models, TPDS, under submission*

Motivation

Collaboration

Multi-node Architecture

Λ(**Node**, **Core**, **Network**) ⇒ρ()

Modeling HPC Network

Cluster with crossbar

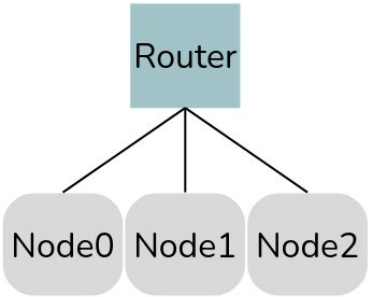

Cluster with a shared backbone

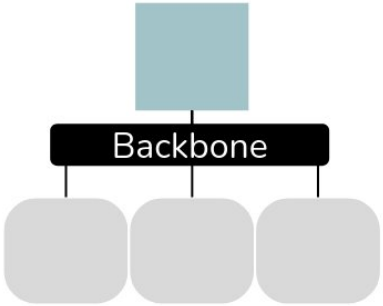

Torus cluster

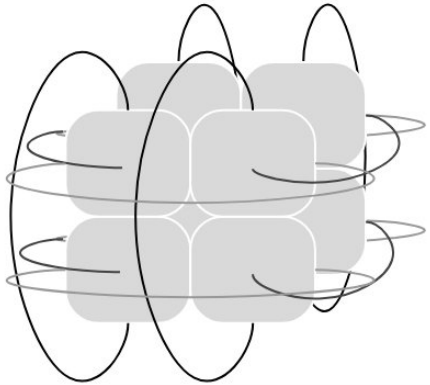

Fat tree cluster

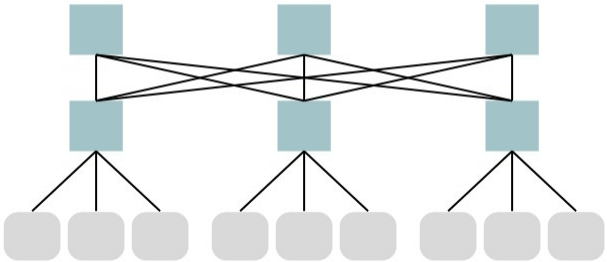

Dragonfly cluster

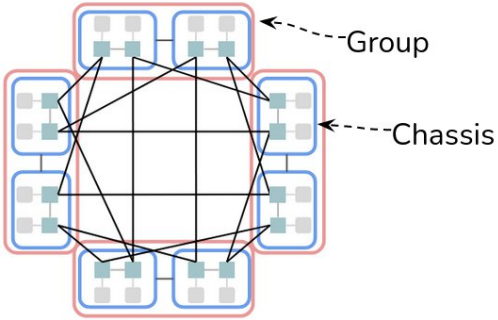

Iterative Simulations for Optimizing Parallelism

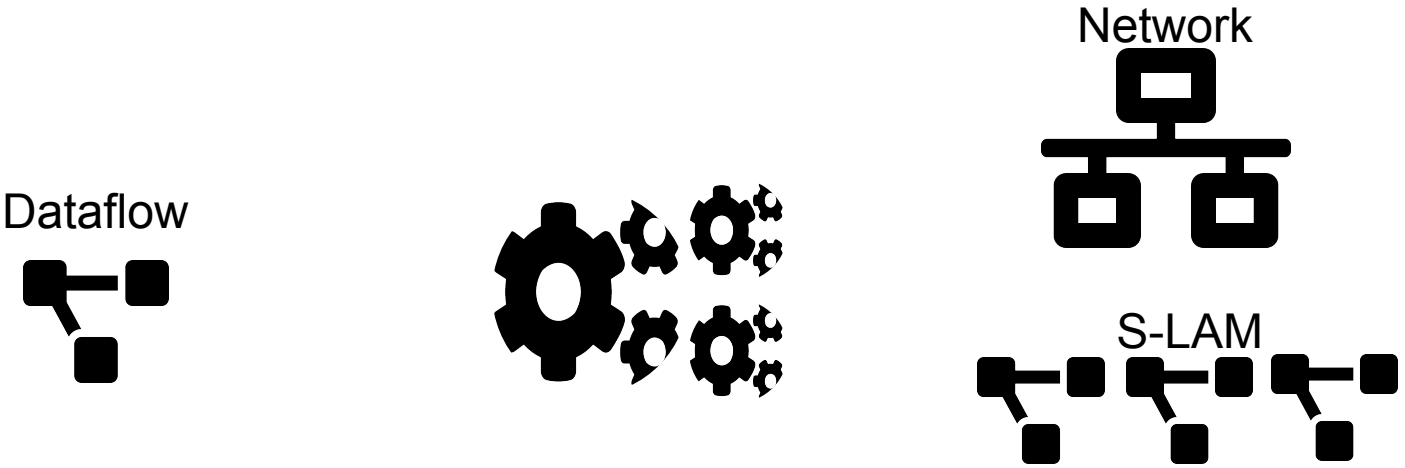

$$P_{\max} = \max\{N \times C \mid \exists i \geq \delta_I : L_{\text{final}}(P, i) \leq L_{\text{final}}(S_{\max})\}$$

***O. Renaud**, K. Desnos, E. Raffin, J.-F. Nezan. SimSDP: Dataflow Application Distribution on Heterogeneous Multi-Node Multi-Core Architectures, EUSIPCO, 2024*

Motivation

Collaboration

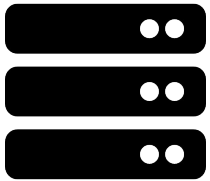

CPU-GPU Resource Allocation

Automated dataflow optimization for NVIDIA GPUs

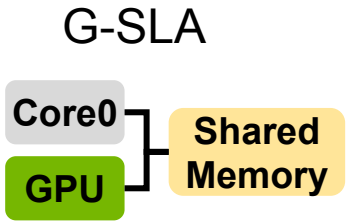

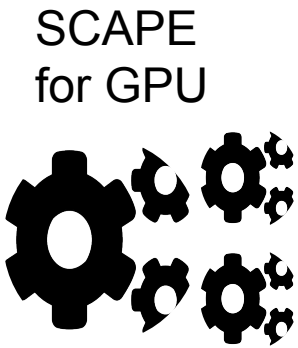

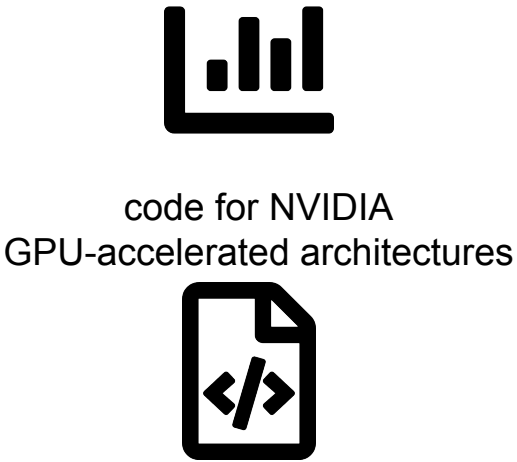

G-SLA: GPU-oriented System-Level Architecture Model

[r]Saman Payvar, Maxime Pelcat, and Timo D Hämäläinen, « A model of architecture for estimating GPU processing performance and power »

*This work carried out by the intern I supervised, **Ewen Michel,** at Swinburne University and IETR as part as the Rising STAR project

Summary conclusion

Achievement:

- ✓ Granularity optimization for multicore architecture.
- ✓ Granularity optimization for HPC multinode & multicore architecture (SimSDP).
- ✓ Method to retrieve the best architecture for a given application.
- ✓ Granularity optimization for GPU accelerated architecture.
- ✓ Deployed and validated on Grid5000.
- ✓ Implemented into PREESM and available on GitHub.

SimSDP Tutorial available on the PREESM website:
SimSDP: Multinode Design Space Exploration - Preesm

Post-doctorate overview

Optimizing Radio Astronomy Imaging Algorithms on HPC Systems with SimSDP and Fine-Grained Descriptions

Started on October 1, 2024

Supervised by Nicolas Gac & Martin Quinson

Generic imaging pipeline

Objective: Develop a proof of concept for SimSDP, enabling simulation of radio astronomy imaging pipelines on HPC architectures.

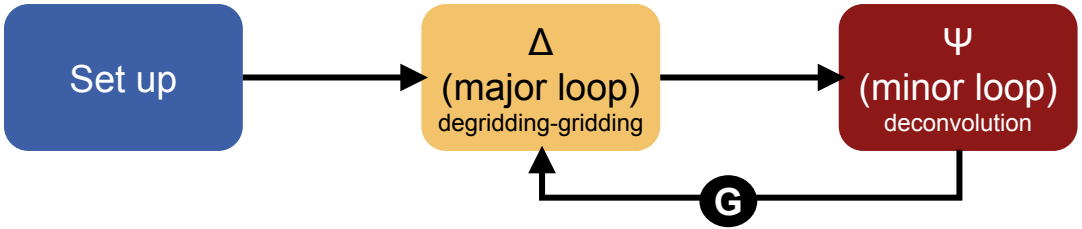

Direct Fourier Transform (DFT) *(simple but long)*

Fast Fourier Transform (FFT) *(faster on big grid)*

Grid to Grid (G2G) *(faster, same O)* *N. Monnier*

Högbom CLEAN *(simple but long)*

🚧 Smax = 2.8 → no sense to deploy this application on more than 3 cores, and even less on HPC systems

[r] S. Wang, N. Gac, H. Miomandre, J.-F. Nezan, K. Desnos, F. Orieux « An Initial Framework for Prototyping Radio-Interferometric Imaging Pipelines»

G2G dataflow pipeline with a finer dataflow description

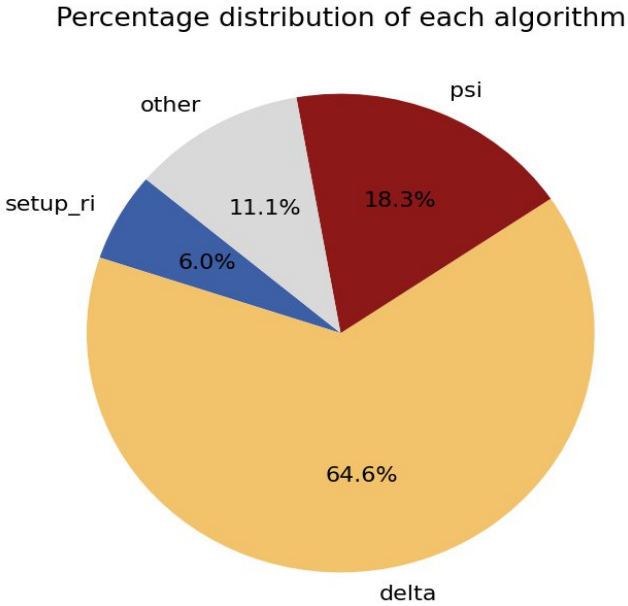

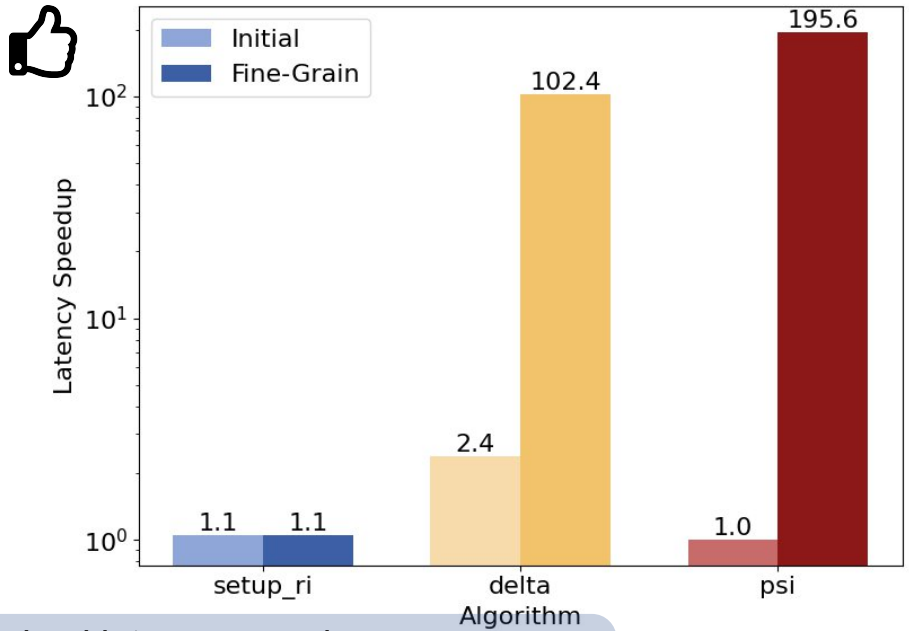

Improve theoretical maximal latency speedup
Far from SKA HPC parallelism [50_000 - 100_000] nodes
We expect SimSDP to benefit more from parallelism

Post-doctorate outline

Phase 1: Automating Fine-Grained Transformation
→ Dataflow models require fine granularity for scalability and GPU resource allocation.

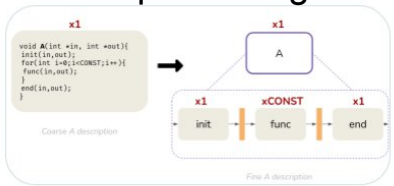

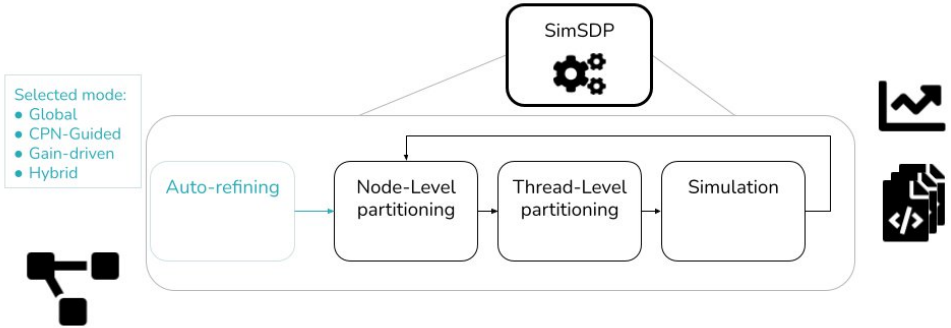

Phase 2: GPU Integration in SimSDP
→ Extend co-design approaches to GPU-accelerated architectures using SimSDP.
Phase 3: Comparison with DASK and DALiuGE
→Develop dataflow graph translators

Thank you

***O. Renaud**, D. Gageot, K. Desnos, J.-F. Nezan. SCAPE: HW-Aware Clustering of Dataflow Actors for Tunable Scheduling Complexity, **DASIP**, 2023*

***O. Renaud**, N. Haggui, K. Desnos, J.-F. Nezan. Automated Clustering and Pipelining of Dataflow Actors for Controlled Scheduling Complexity, **EUSIPCO**, 2023*

***O. Renaud**, H. Miomandre, K. Desnos, J.-F. Nezan. Automated Level-Based Clustering of Dataflow Actors for Controlled Scheduling Complexity, **JSA**, 2024*

***O. Renaud**, A. Gougeon, K. Desnos, C. Phillips, J. Tuthill, M. Quinson, J.-F. Nezan. SimSDP: Dataflow Application Distribution on Heterogeneous Multi-Node Multi-Core Architectures, **TPDS**, under submission*

***O. Renaud**, K. Desnos, E. Raffin, J.-F. Nezan. SimSDP: Dataflow Application Distribution on Heterogeneous Multi-Node Multi-Core Architectures, **EUSIPCO**, 2024*

*E. Michel, **O. Renaud**, K. Desnos, A. Deller, C. Phillips, J.-F. Nezan. Automated Deployment of Radio Astronomy Pipeline on CPU-GPU Processing Systems: DiFX as a Case Study, **ADASS**, 2024*