



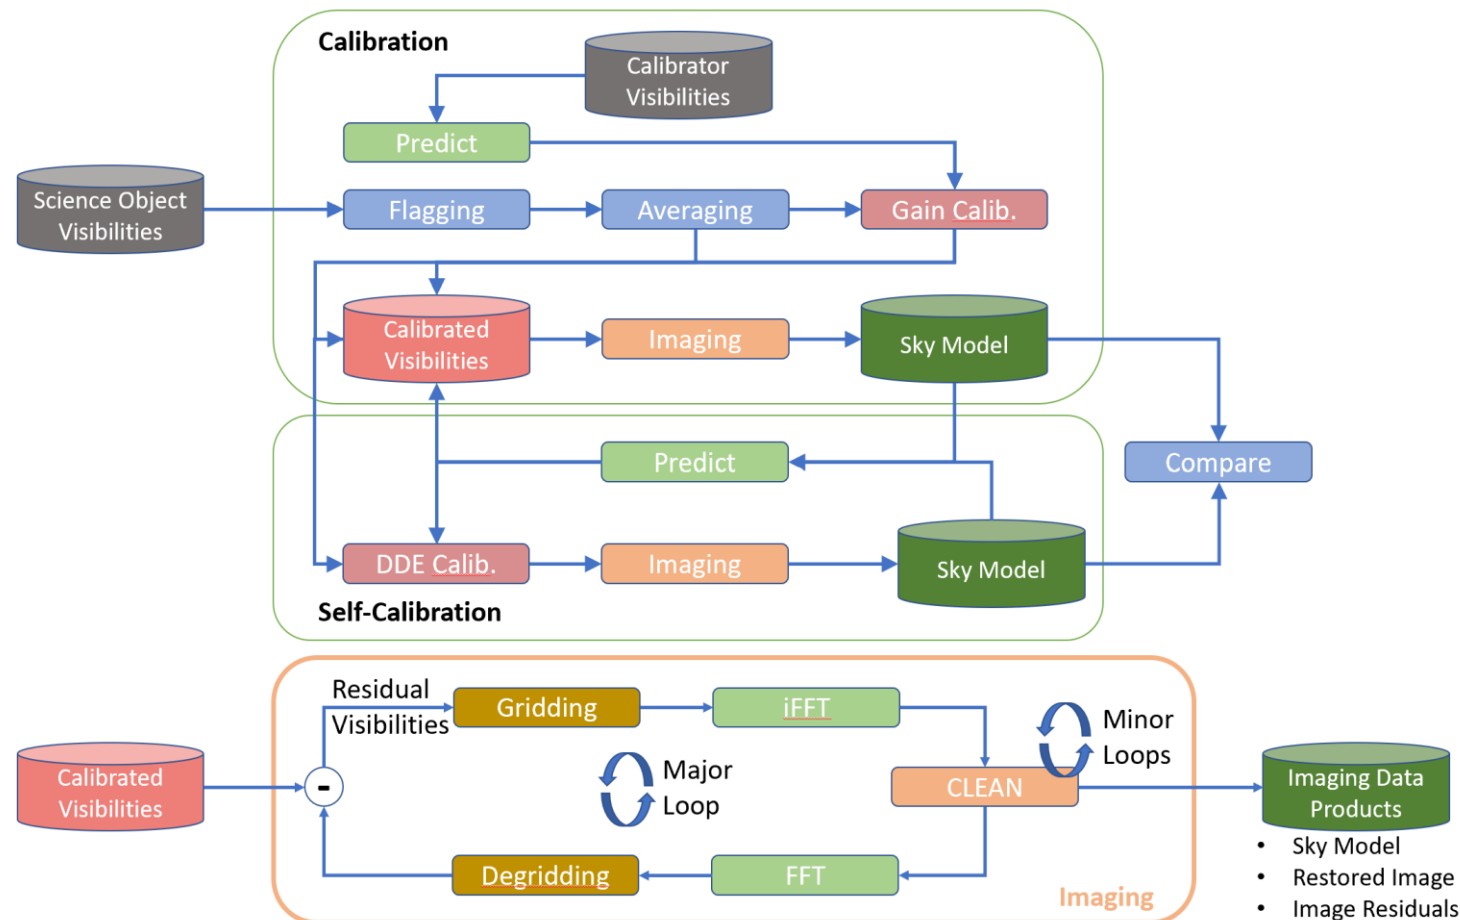
PROGRAMME
DE RECHERCHE
NUMÉRIQUE
POUR L'EXASCALE

ExaDoST: WP4 – SKA Illustrator

Damien Gratadour (OP – CNRS)
Hugo Gaquère, Baptiste Besnard (OP)
Sunrise Wang, Shan Mignot (OCA)
Iheb Becher, Arnaud Collioud, Marwane Dalal, Etienne
Bonnassieux (LAB)

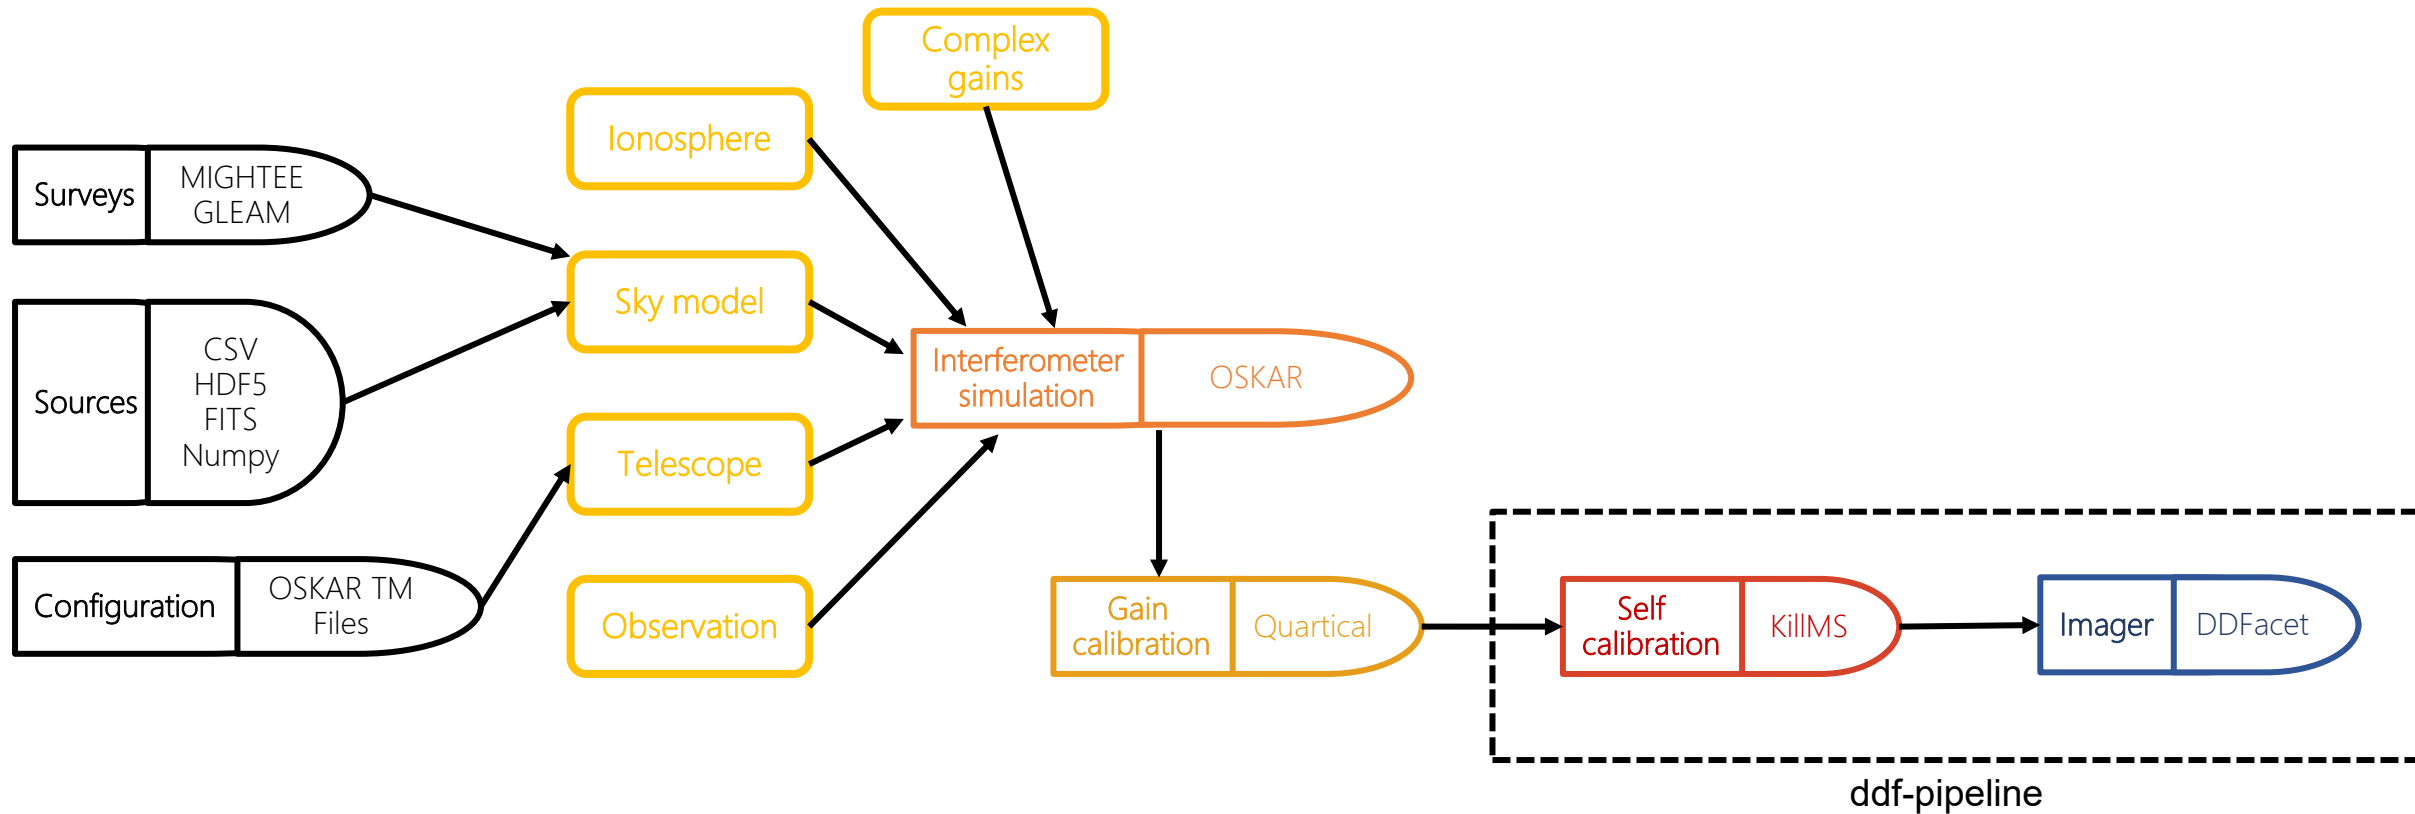
SKA illustrator

Complex, iterative workflow with pipelining opportunities



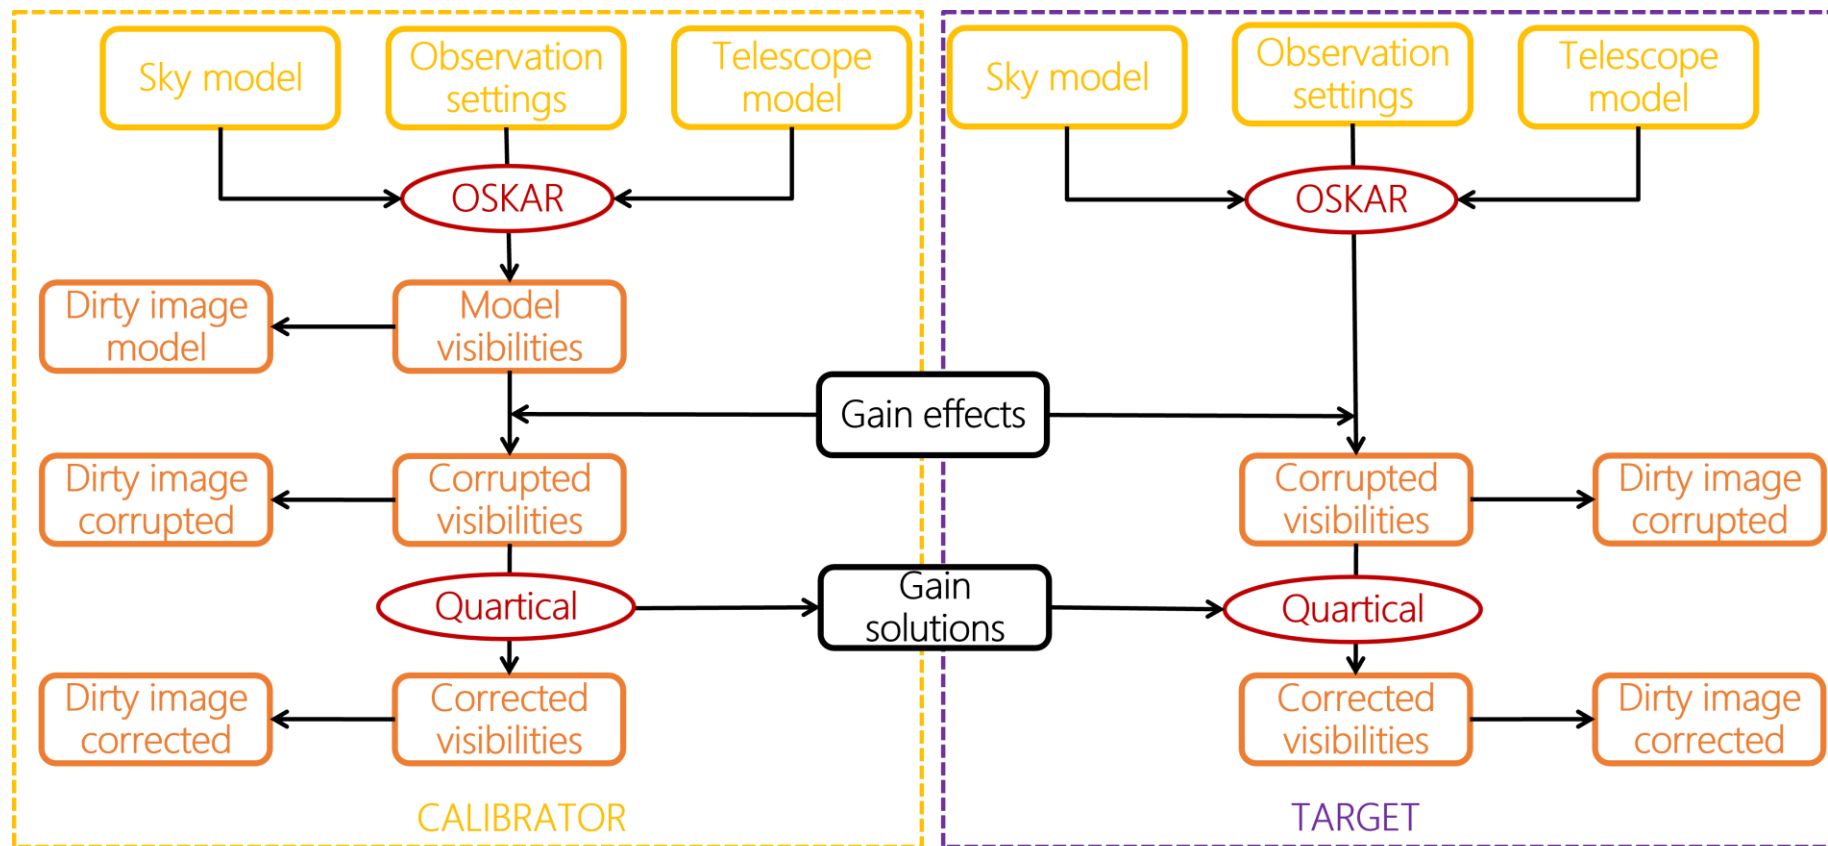
SKA illustrator

Existing codebase



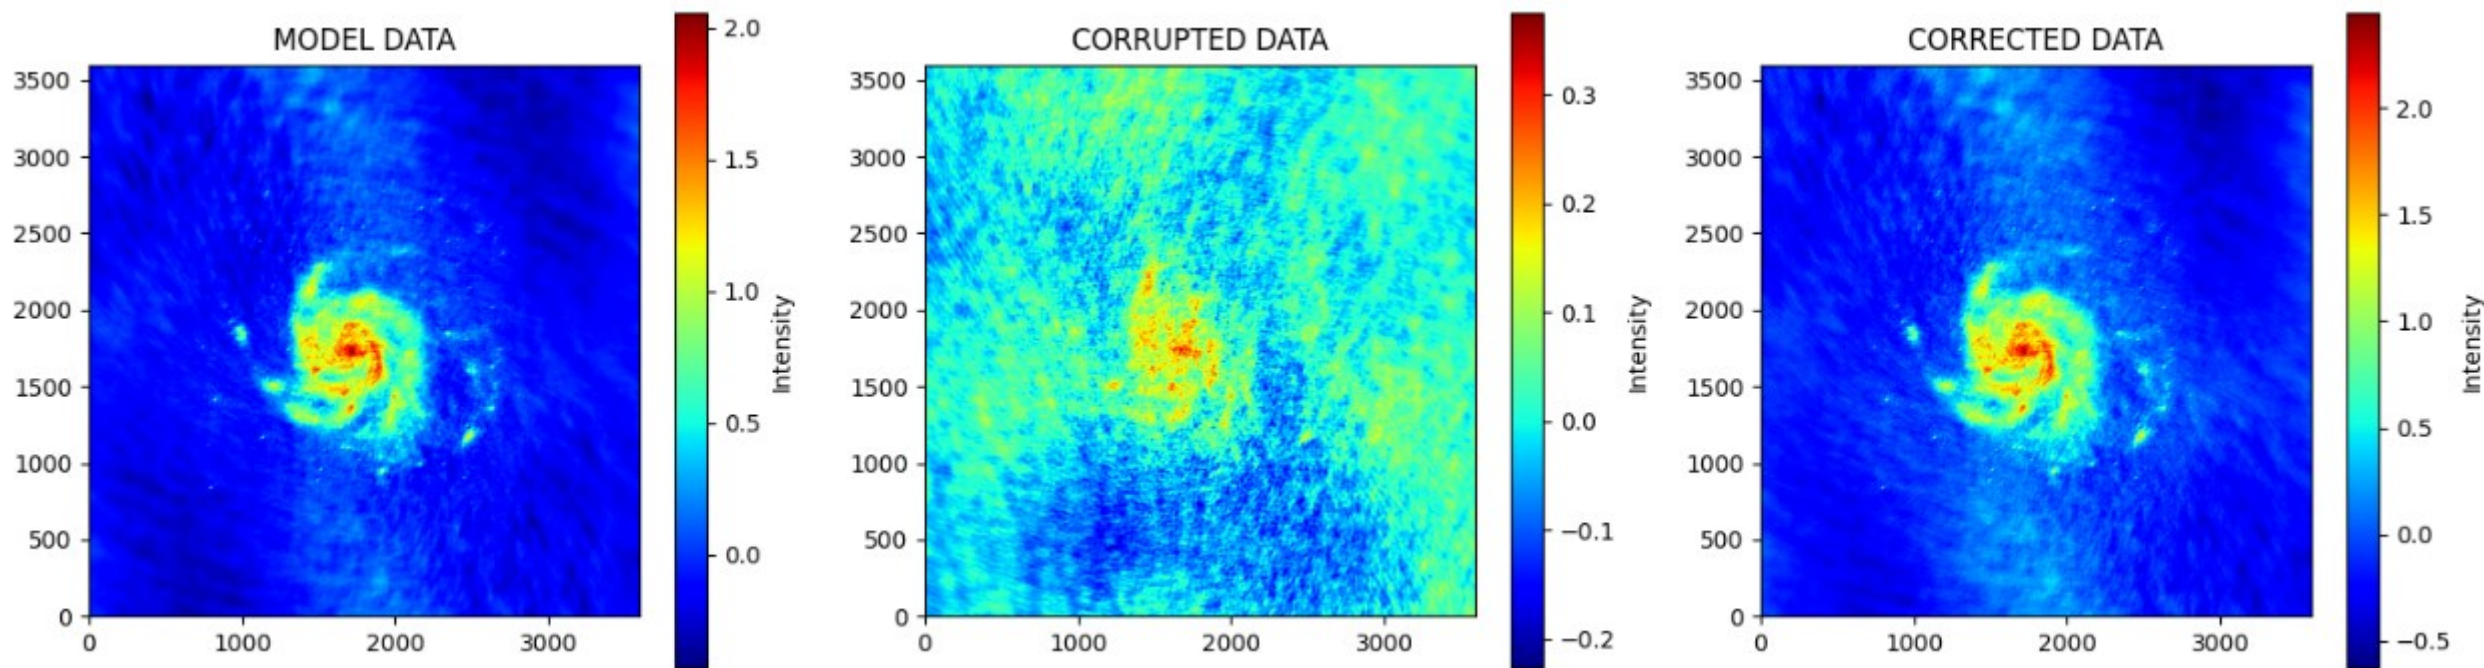
Credit: B. Besnard

Enabling technology: realistic observing scenario emulator



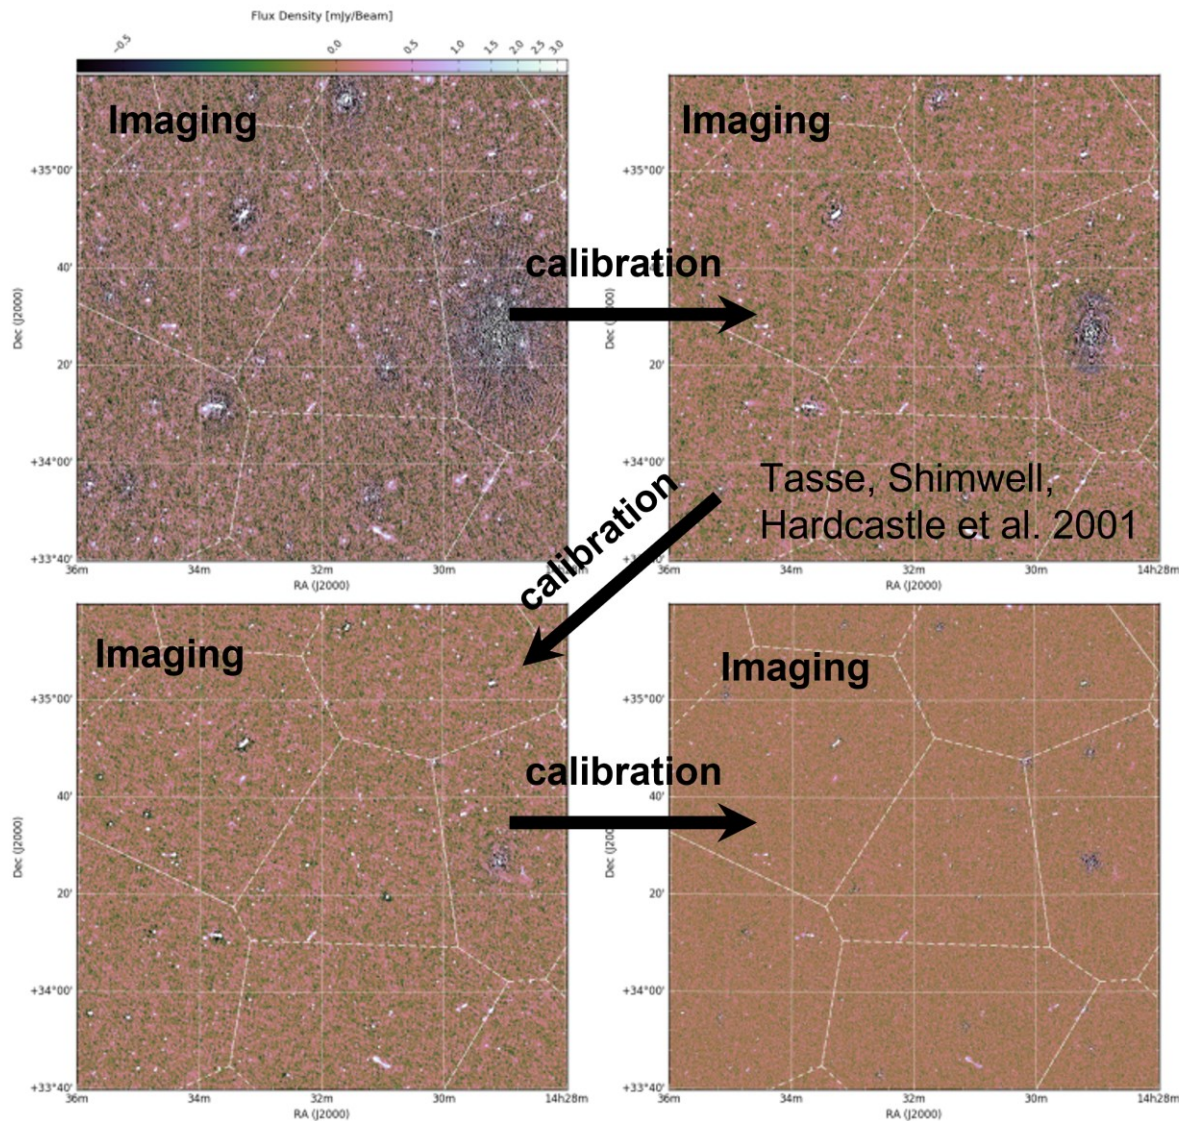
Enabling technology: realistic observing scenario emulator

Typical output (a.k.a. “dirty image”)



Credit: B. Besnard

Core composite: selfcal + imaging



Credit: C. Tasse

Selfcal + imaging: parallelization schemes

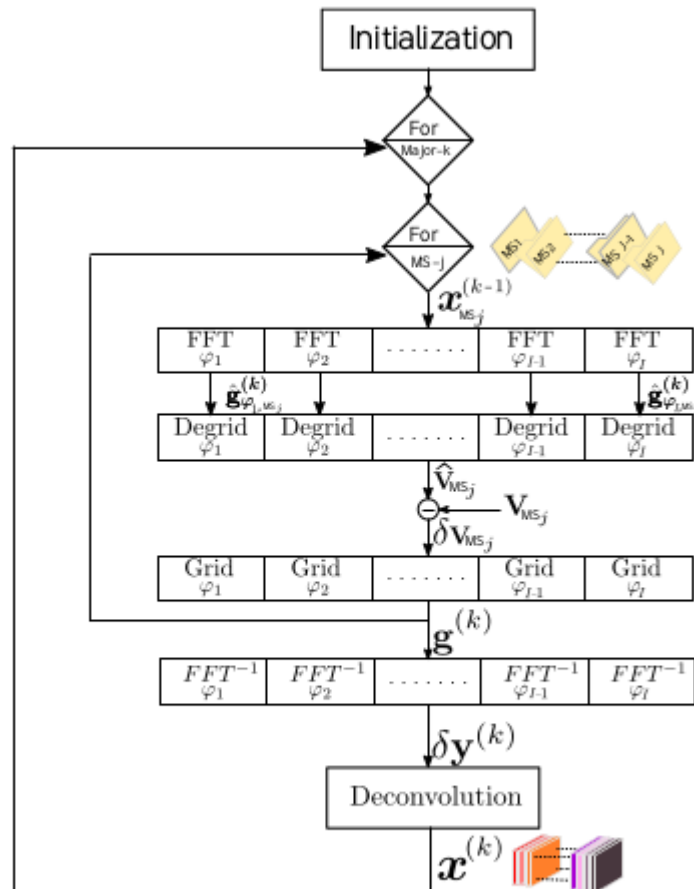
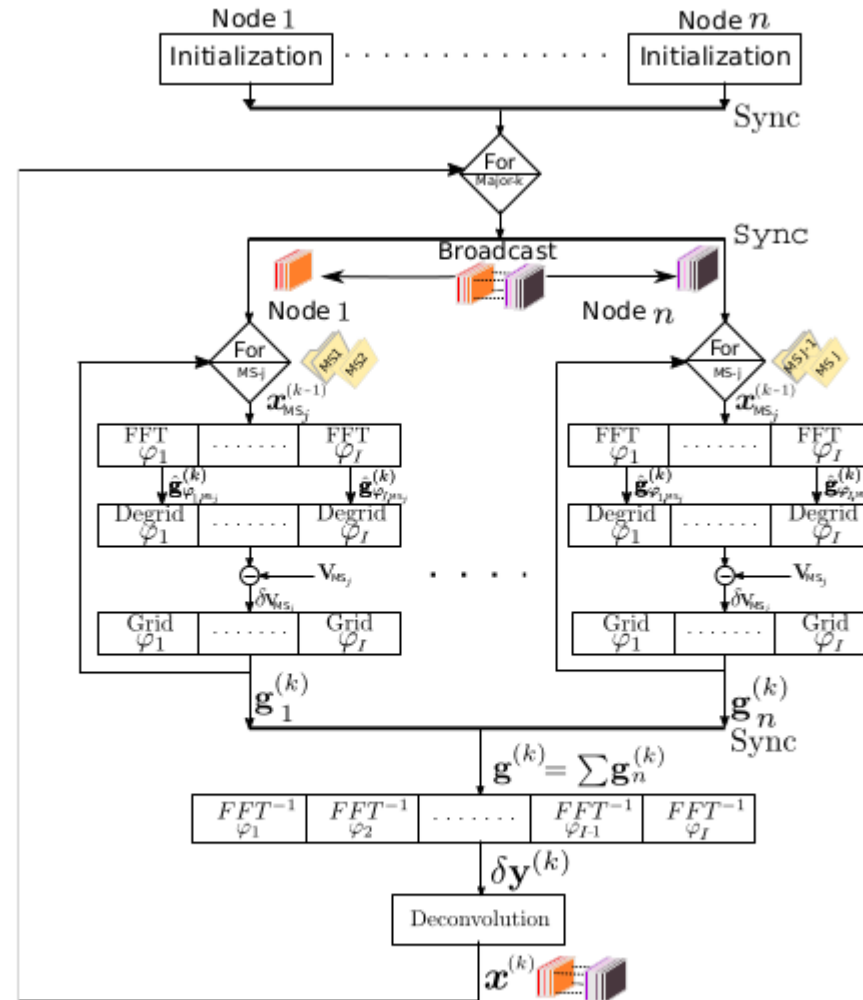


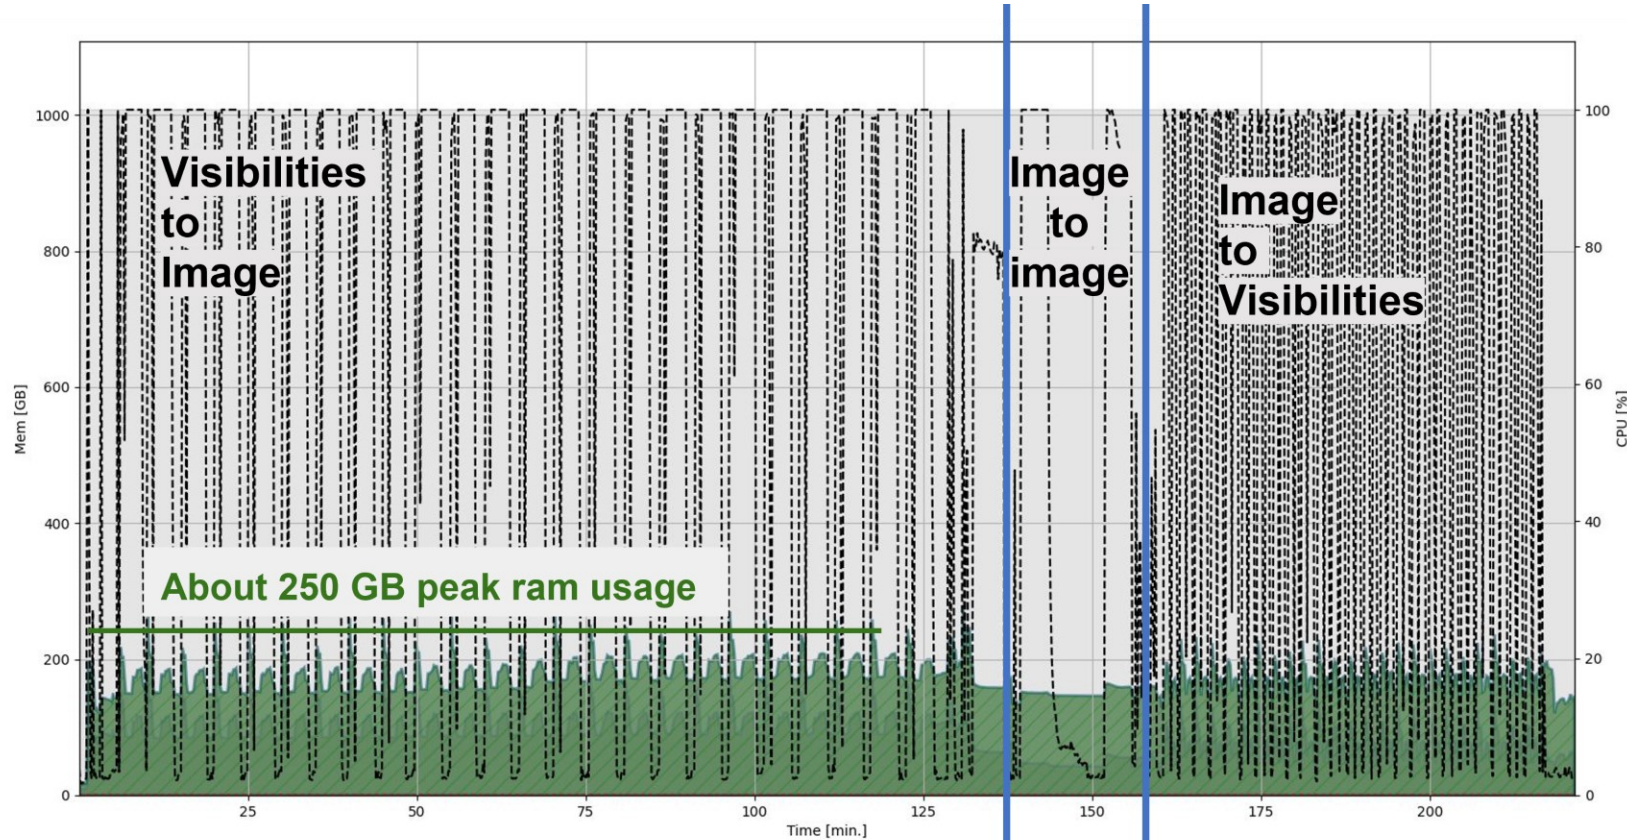
Figure 4.4: Multi-core parallelization



Credit: N. Monnier

Figure 4.7: Multi-core and multi-node parallelization based on independent MS files using distributed memory system.

Typical DDF-pipeline profile (single node, multi-core)



loop over MSlist:

- compute residual vis: degrid
- grid resid
- grid PSF

Minor cycle:
deconvolve given
- the dirty image- the
PSF image

loop over MSlist:

- compute residual vis: degrid
- grid resid

Credit: C. Tasse

Current efforts: optimizing the workflow

Consolidating distributed memory ports (selfcal + imaging)

- Validating multi-node deployment on large clusters (e.g. Jean Zay): [collab. ExaAToW, SCOOP-SKAO](#)
- Benchmarking I/O (see next slides) : [funded by ExaDoST](#)
- New algorithmic approaches (see next slides): [funded by ExaDoST](#)

New parallelization strategies: [collab. ODISSEE, SARAQ, Breakthrough Listen](#)

- GPU acceleration for gridding / degridding
- GPU acceleration for deconvolution
- Pipelining data simulation & processing
- More realistic simulation scenarios

Efficient workflow management: [collab. KerData \(E2CLab\) PEPR Cloud + BSC](#)

Current efforts: benchmarking I/O

Goals :

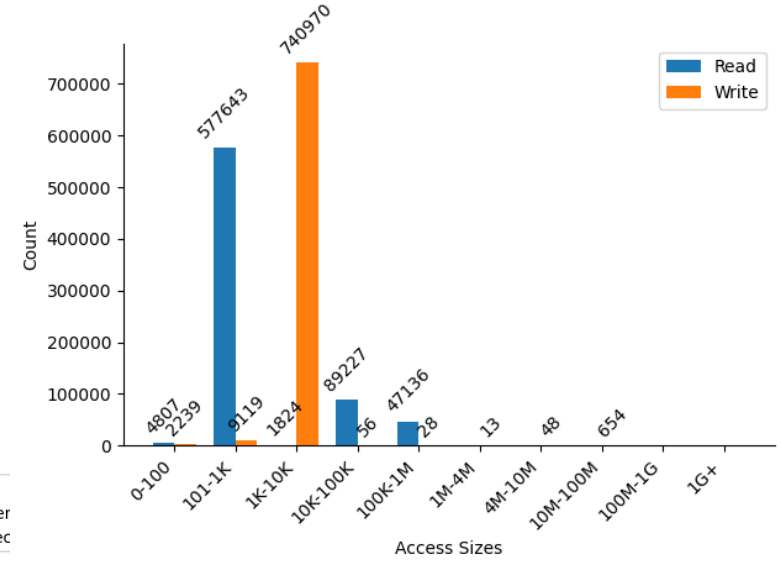

- Build a global understanding of performance and scalability in the DDFacet/KillIMS pipeline

Focus areas :

- I/O behavior and data movement
- Memory & CPU utilization

I/O-oriented tools :

- Darshan : I/O statistics & visualizations
- DFTracer : I/O profiling & Syscall-level tracing

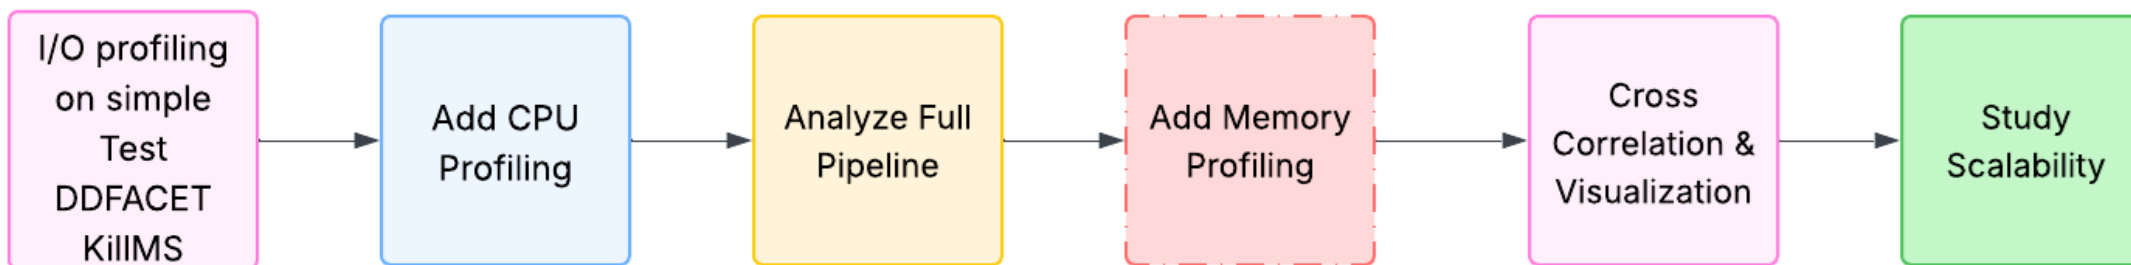


Credit: I. Becher

Current efforts: benchmarking I/O

Next steps :

- I/O analyses of full pipeline
- Integrate memory and CPU traces for full pipeline insight
- Pipeline scalability study

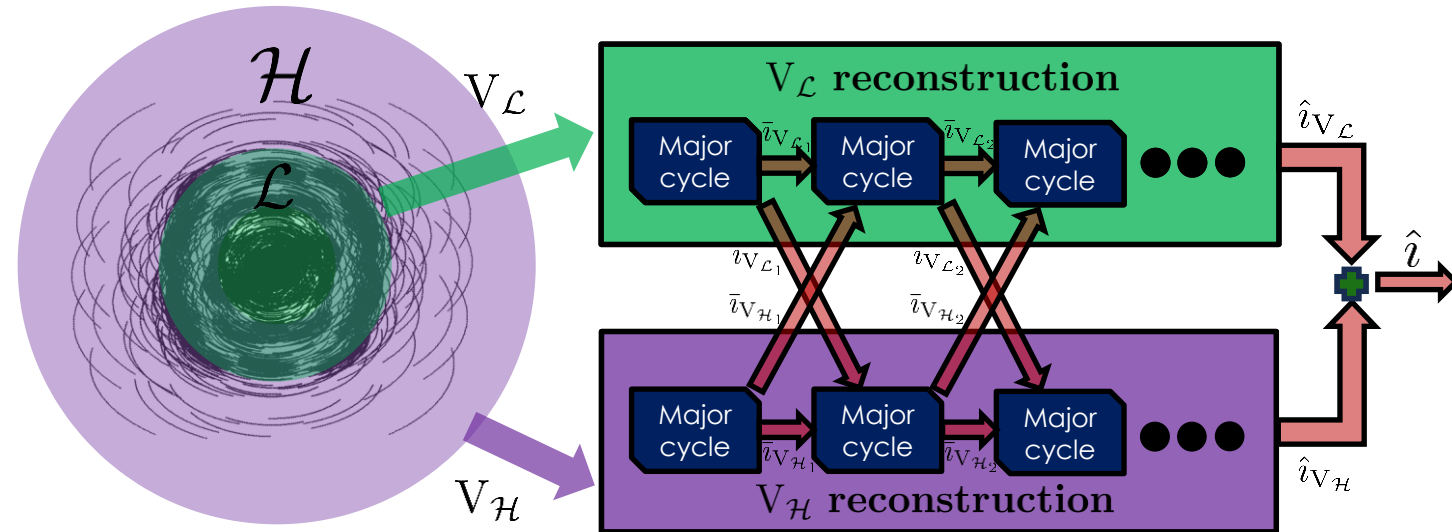


Credit: I. Becher

Current efforts: new algorithmic approaches

Spatial frequency decentralized imaging framework

- Each node reconstructs full-resolution images from partial spatial-resolution visibilities
- Operates within traditional major-minor cycle paradigm
- Existing pipelines can be easily molded into framework
- Largely agnostic of internal algorithms.
- Should work in tandem with time/em-frequency parallelization

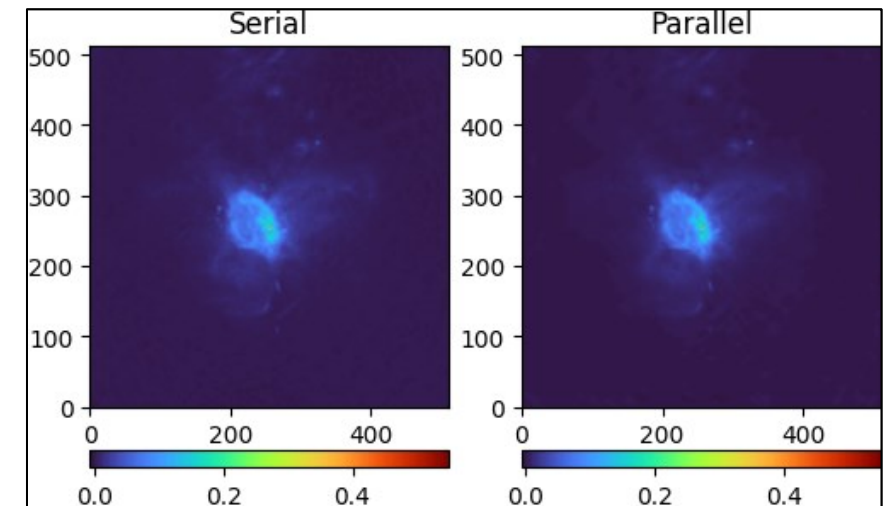




1. [2]Wang, S., et al. "Multi-step reconstruction of radio-interferometric images." *Astronomy & Astrophysics* 692 (2024): A61
 2. [2]Wang, S., et al. "A Decentralized Framework for Radio-interferometric Image Reconstruction." *The Astronomical Journal* 169.6 (2025): 289.

Current efforts: new algorithmic approaches

Initial results on MeerKAT dataset

- 5 major cycles
- Run on Jean-Zay cpu_p1 nodes
- Achieved $\sim 4.12x$ speed up for 5 partitions
- Scaling becomes worse with larger number of partitions, achieving only a $9.02x$ speedup for 15
- Non-perfect load balancing and data duplication most likely causes



Example reconstructions

Credit: S. Wang

Current efforts: new algorithmic approaches

Future works

- Improve load balancing
 - Currently under testing
- More realistic datasets
 - Longer observation times and more channels
 - More realistic image sizes
- Evaluation metrics for convergence

Perfectly load balanced visibilities





Image source: Heywood et al.[1]



Credit: S. Wang



RÉPUBLIQUE
FRANÇAISE

*Liberté
Égalité
Fraternité*



Inria



PROGRAMME
DE RECHERCHE
NUMÉRIQUE
POUR L'EXASCALE

Retrouvez toutes nos actualités

 NumPEX

