Tuesday 4th October, 2022
Swiss SKA Days 2022
Lugano, CH

# Co-design for the Square Kilometre Array

*Chris Broekema, Shan Mignot*
*on behalf of team SCOOP*

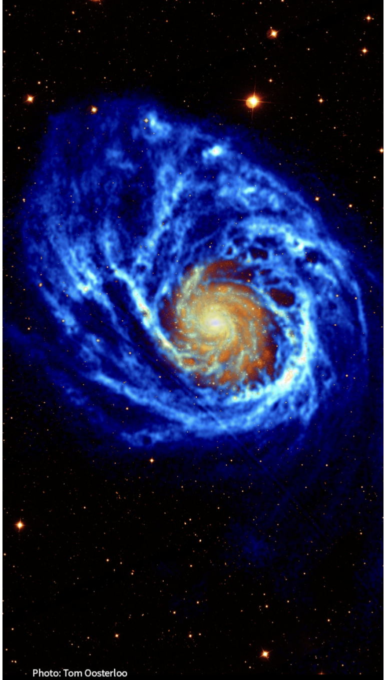


ASTRON
Netherlands Institute for Radio Astronomy

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

# Context

Radio astronomy relies critically on the availability of abundant and cheap compute resources to transform sampled sky data into scientifically viable data.

- Precursor instruments like LOFAR, NeNuFAR and MWA are already heavily limited by available resources
- Pathfinder instruments are generally even more so
- The SKA is expected to exceed all of these by at least one order of magnitude
- Each of the 2 SKA SDPs is expected to require[1]
  - $O(150)$ PFLOPS
  - $O(40)$ PB of fast buffer storage
  - to ingest $O(1)$ TB/s of data from the correlator/beamformer
  - to continuously export $O(100)$ Gbps of data products to the archive

[1] from SKA SDP CDR documentation and updated ICDs

# The problem

LOFAR was explicitly designed to leverage the inexorable and predictable increase of available compute technology over time.[2]
The demise of Dennard's scaling and Moore's law makes that we cannot make that same assumption for the SKA.
Although new generations of hardware will continue to be developed, HPC systems will

- become evermore complex
- likely increase in price due to more expensive development and production cost
- be far more difficult to use efficiently
- probably not able to run at full speed continuously
- require prodigious amounts of energy

Furthermore, these continuous one-time optimisations will inevitably end, sooner rather than later.[3]

[2] Bregman 1999, Bregman 2000, Bregman 2004
[3] See IEEE IRDS for example `https://irds.ieee.org/`

# Fortunately

Two main classes of IT system can be considered of similar scale as the SKA SDP.

1. general purpose supercomputers
2. cloud infrastructures

Both of these need to support a very diverse set of applications.

The SKA SDP in contrast can be specifically tailored to its specific tasks.

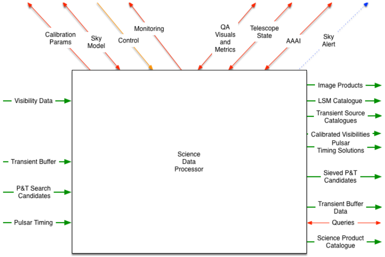


Figure: Data types handled by the SDP

# Enter the SCOOP Team

SKA CO-design and OPtimisation First envisioned just prior to the SKA system CDR: SKA Computing Hardware Risk Mitigation Plan[4]

To reduce cost and avoid idle compute hardware, SDP is procured as late as possible. This introduces its own problems:

- Prices for hardware are difficult to predict: identified as a financial risk (procurement, operations)
- Performance evolves
  - Inability to estimate performance of procured components in advance
  - difficult to predict number of components to achieve performance requirements
- Uncertainty regarding compute and storage needs (software under development)

Proposed mitigation: establish a team dedicated to track developments, and maintain expertise in performance optimisation for the state of the art compute hardware.

[4] SKA-TEL-SKO-0001083-01

# SCOOP - SKA system-level goals[6]

For the SKA system: optimise the hardware/software match to reduce procurement risks and maximise the value of the investment

- Lead a transverse co-design activity across the DP, Services and OMC ARTs
- Collaborate on system-level issues to better integrate the SDP to telescopes' designs
- Review and contribute to updating the SDP design
- Enumerate the value and cost of various s/w and h/w choices[5]
  - value can be performance, accuracy, programmability, correctness, quality of implementation, etc
  - cost can be capital, operational, energy consumed, developer time, etc
- Enable well-reasoned procurement decision to be made
- Support procurement with hardware benchmarking and acceptance testing code

[5] https://arxiv.org/abs/1806.06606

[6] https://confluence.skatelescope.org/display/SE/Co-design+purpose+and+organisation



ASTRON

# SCOOP - SKA team-level goals

For the SKA project, in particular for the Data Processing software development teams:

- Enable seamless performance regression testing
- Provide optimisation and hardware expertise and experience
- Support development with performance regression testing capabilities
- Make developers aware of energy consumed by code with power measurement framework
- Mitigate the budgetary/performance imbalance risk via
  - repeated assessment of the match between the SW/HW design and the needs
  - modelling performance and operation of the SDP
  - devising strategies for reducing operations' sensitivity to SDP resources

# SCOOP - Outward-facing goals

Involve industrial partners:

- Provide a set of benchmarks suitable for component and system evaluation
- Increase collaboration/communication
- foster detailed understanding of the SKA requirements with industry
- Inform on the progress of the project
- Be informed on technological advances made by industry

Involve academic partners:

- Collaborate on defining R&D studies applicable to co-design for the SKA
- Leverage existing expertise in the community

The SCOOP team is the conduit between industry and academia on one hand, and the SKA project on the other.

# Current status

The team has been operational since the start of 2020 (PI6). First together with the networking team in the DP ART, since PI12 in the Services ART, since PI15 as a separate team.

Currently self-funded by

- Observatoire de la Côte d'Azur (FR)
- EPFL (CH)
- Observatoire de Paris (FR)
- FHNW (CH)
- ASTRON (NL)

Funding from the SKA construction budget is underway, subject to an ECP which will lead to an invitation to tender for which a joint French-Swiss proposal is expected. Will follow the just-retour principle (FR and CH).

# Some early results

- SDP benchmarking framework[7] based on ReFrame[8]
- AMD-optimized Image-Domain Gridder (IDG) (Stay tuned for Stefano's talk next)
- Contribution to SKAO's Pulsar Search architecture review[9]
- IDG microbenchmark[10]
- Added AMD energy monitoring to existing power sensor framework[11]
- Initial attempt at simulating execution using SST[12](work in progress)
- Explored performance on NVIDIA DGX-1 server
- Identified and investigated performance anomalies in IDG gridding/degridding code

[7] https://developer.skao.int/projects/ska-sdp-benchmark-tests
[8] https://github.com/reframe-hpc/reframe
[9] https://confluence.skatelescope.org/display/SE/Architecture+Reviews%3A+PSS
[10] https://gitlab.com/ska-telescope/sdp/ska-sdp-idg-bench
[11] https://gitlab.com/astron-misc/libpowersensor
[12] https://sst-simulator.org/

# Roadmap of future work

Priorities will be defined on a per-PI basis in close collaboration with business owners, architects and developers. Resources will depend on funding.

- contribute to SDP architecture review (Oct 2022)
- revisit hardware risk mitigation plan
- re-invigorate industry engagement
- continue characterisation of various accelerators
- add additional computational kernels and pipelines to benchmark suite
- expand focus to include high-performance buffer
- start collecting performance and energy consumption library
- work on a renewed cost model for SDP
- establish a high-level software model for (parts of) SDP
- revisit the SDP h/w architecture
  - compute node model (host+accelerator, number of accelerators)
  - networks (IB, OMP, BXI, Ethernet)
  - high-performance buffer (all flash, size, interconnectivity)



ASTRON

# Perspectives

- Scope and organisation of co-design still being discussed with SKAO (ITT should help clarify things)
  - SKAO's intent is to fund co-design through the savings it will allow
  - SKAO-funded resources will be limited (manpower, hardware)
  - External collaborations are valuable
  - See for instance the SKAO/CERN/Prace/Geantt
  - DICES proposal for Horizon Europe
  - Open for more engagement
- The SCOOP team
  - Chris Broekema (ASTRON, product owner)
  - Shan Mignot (OCA, scrum master)
  - Stefano Corda (EPFL, developer)
  - Xuezhou Lu (Observatoire de Paris, developer)
  - Manuel Stutz (FHNW, developer)



ASTRON