Co-design activities for SKAO supercomputers

Shan Mignot for the SCOOP team

Journée SKA France 2022

Context

- Radio astronomy relies critically on computation
  - early digitization of signal: digital instrument
  - replication of data streams: commensal observations
  - data reduction & data analysis
- Precursor instruments like LOFAR, NeNuFAR and MWA are already heavily limited by available resources
- Pathfinder instruments are generally even more so
- SKA expected to exceed all of these by at least one order of magnitude

Telescope availability

- Telescope operational availability > 94%
  - Operational Capable: A telescope is operationally capable when it can perform astronomical observations. It is assumed that the telescope will be capable if more than 95% of its collecting area, signal processing and data reduction capabilities are available.
  - Asynchronous nature of SDP operation
    - SKAO-SDP_REQ-486: The SDP must achieve an average processing- over observation-time ratio of 1 over a period of 30 days.
    - SKAO-SDP_REQ-495: The SDP shall have an Inherent Availability (Ai) higher than or equal to 99.9%.
- unavailability of SDP due to saturation of resources not considered in requirements analysis

	MID-Telescope-Time %	LOW-Telescope Time %
Observing (OT)	88%	90%
Standby (ST)	0%	0%
Weather (WT)	3%	1%
Utility (UT) (External: power, cooling water and communications services)	4%	4%
Engineering / Maintenance (MT) (scheduled maintenance, off-line calibration, software updates or testing)	3%	3%
System Fault (DT)	2%	2%
Critical Repair Time (CRT)	1%	1%
Critical Support Delay (CSD)	1%	1%
	100%	100%

State		Definition
Operational	Observing	At least one sub-array is in the Calibrating or Observing mode, for the purpose of science observations.

from Jama specification

Feasible? Achievable?

- Critical design review assumptions
  - Dennard's scaling & Moore's law: predictable increase of performance over time
  - procure SDP as late as practicable
  - estimated max performance: ~125 PFlop/s ranks 6^{th} in Top500 today
- SKA computing hardware risk mitigation plan
  - financial risk vs. evolution of performance: capital cost of hardware
  - estimate uncertainties (software under development)
  - procurement strategy: collaborate with suppliers
  - power consumption: operational cost

Challenges (I)

- Data management
  - ingest data rates Mid & Low
  - streaming processing
    - close the telescope control loop
    - maintenance of calibration models
    - transients
    - live quality assurance for operators
  - buffer size & batch processing
    - introduce flexibility in scheduling
    - ~40 PB for 24h of storage
  - publish data products for archiving
    - ~300 PB/year/telescope

Mid

12 bytes/visibility
* C^2_{196} = 19110 baselines
* 4 polarizations
* 65535 channels
/ 0.14 s min integration
= 0.43 TB/s

Low

12 bytes/visibility
* C^2_{512} = 130816 baselines
* 4 polarizations
* 65536 channels
/ 0.9 s min integration
= 0.46 TB/s

numbers by P. Wortmann

Challenges (II)

- Energy budgets:
  - average: 1.3 MW Mid / 1.6 MW Low
  - peak: 2.0 MW Mid / 2.23 MW Low
  - Green500: Frontier, Lumy, Adastra achieve ~100 PFlops @ 2MW at maximum efficiency
  - diversity of tasks and operating conditions
- Staged deployment
- 50-year lifespan

SKA1-Mid SPC/ SOC Power Budget in Cape Town

Products	AA4 Long Term Average (>30min) [kW]	AA4 Peak Instantaneous (<5sec) [kW]	AA* Long Term Average (>30min) [kW]	AA* Peak Instantaneous (<5sec) [kW]
PDT4 - MID Digitisation	**230.8**	**323.4**	**230.8**	**323.4**
CSP.CBF	230.8	323.4	230.8	323.4
PDT6 - Network & Computing	**1641.7**	**2481.9**	**589.3**	**872.6**
SDP Hardware MID	1300.0	2000.0	325.0	500.0
PSS Hardware MID	296.0	414.0	222.0	310.5
PST Hardware MID	16.4	26.8	12.3	20.1
OMC Hardware MID	12.9	18.1	12.9	18.1
NSDN MID	5.6	7.8	6.6	9.2
CPF-SPC link MID	8.2	11.5	7.9	11.1
NMGR	2.6	3.6	2.6	3.6
Building losses and cooling	**374.5**	**561.1**	**164.0**	**239.2**
Commissioning Margin	**224.7**	**336.6**	**98.4**	**143.5**
Site Total	**2471.7**	**3702.9**	**1082.6**	**1578.8**

SKA1-Low SPC/ SOC Power Budget in Perth

Products	AA4 Long Term Average (>30min) [kW]	Peak Instantaneous (<5sec) [kW]	AA* Long Term Average (>30min) [kW]	AA* Peak Instantaneous (<5sec) [kW]
PDT6 - Network & Computing	**1629.2**	**2270.9**	**429.2**	**598.4**
OMC Hardware LOW	12.9	18.1	12.9	18.1
SDP Hardware LOW	1600.0	2230.0	400.0	557.5
NSDN LOW	6.5	9.1	6.5	9.1
CSP-SDP LOW	7.2	10.1	7.2	10.1
NMGR	2.6	3.6	2.6	3.6
Building losses and cooling	**325.8**	**454.2**	**85.8**	**119.7**
Commissioning Margin	**195.5**	**272.5**	**51.5**	**71.8**
Site Total	**2150.6**	**2997.6**	**566.6**	**789.9**

from SKA-TEL-SKO-0000035-06

Co-design

- Maximize efficiency through co-design: tailor designs to needs
  - SDPs: some diversity but not versatile facilities
  - software still malleable (in development)
  - diverse HPC off-the-shelf hardware: platform design space
- Highly variable loads
  - early Array Assemblies & progressive deployment of SDP
  - commensal observations / channels: almost arbitrarily large loads
  - future extensions of arrays
- Science trade-off: schedule observations to adjust load

Co-design organisation

- In SKAO's SAFe framework: team SCOOP
  - Chris Broekema (Astron): product owner
  - Shan Mignot (OCA/Lagrange): Scrum master
  - Matthieu Carrère (OCA/DSI): developer (starting)
  - Stefano Corda (EPFL): developer
  - Xuezhou Lu (OP/LERMA): developer (ending)
  - Manuel Stutz (FNHW): developer
- Link with research communities
  - Alain Miniussi (OCA/DSI), Simon Prunet (OCA/Lagrange)
  - ANR Dark Era

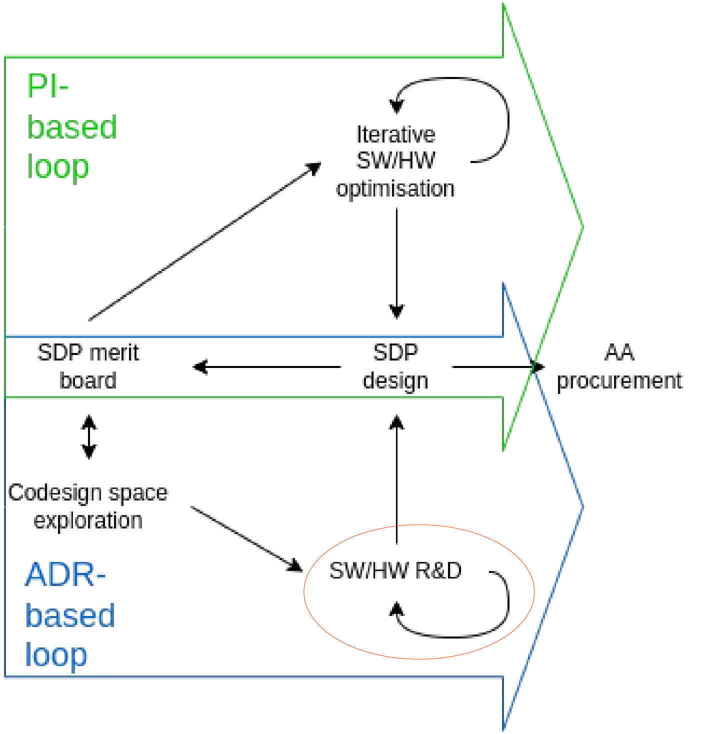

Co-design project goals

- Transverse co-design activity across the project (cross-ART)
- Collaborate on system-level issues to better integrate the SDP to telescopes’ designs
- Support procurement
  - collaborate with industry (sourcing)
  - develop benchmarking and acceptance testing code
- Interface with community
  - collaborate on defining applicable upstream R&D studies (within ECLAT)
  - leverage existing expertise in the community

Co-design technical goals

- Consolidate evaluation of needs through modeling and benchmarking
- Review and contribute to updating the SDP design
  - study trade-offs
  - evaluate alternative designs
- Strategy for adaptive management of pipelines & SDP resources
  - get more from installed system
  - smooth degradation instead of failure
- Long-term vision for managing SDP software & hardware

Co-design perspectives

- Co-design engineering change proposal
  - co-design in WBS: better define/organise co-design within SKAO
  - funding from SKAO (currently self-funded contributions)
  - invitation to tender: joint proposal from France & Switzerland
- Team
  - strengthen collaboration with OP/LERMA (recruitment of Aristide Doussot)
  - engineer CDD position available at Inria Lyon (Avalon team)
- R&D activities via ECLAT
  - défis Inria-Atos
  - environmental impact of SDP: collaborate with EcoInfo group
  - collaboration with LAB in Bordeaux