

# Co-Design Update

## Monitoring Benchmark Executions

Manuel Stutz  
27. January 2025


# Team SCOOP

## Challenges:

- Quality Assurance (QA) of output
- HPC requirements
- Real-time & batch processing
- HW: CPUs, Accelerators, Memories and Storage
- SW: benchmark, profiling and modelling tools

## Goal:

- Minimize cost/availability risk for SDP machines procurement
- Maximize the scientific value of SDP
- Inform procurement decisions



# Current PI Work



Use the AWS HPC cluster to deploy SDP pipeline modules in an HPC configuration



Use Benchmarkmon to provide automated benchmarking of the pipeline modules to the users



Evaluate the readiness of the different pipeline modules

# Current State



SEAC-Review identified some risks concerning maturity of software for SDP



Current PI: MVP for End-To-End-Processing of Visibilities to form Continuum Images at the End-Of-AA1-Scale



Concerns regarding Performance /  
Size of the SDP

Scaling Studies on the SKAO AWS Cluster will  
be conducted



Insufficient Performance = Less Science!

# Benchmarking Suite



**End-To-End Tests of  
Scientific Software**



**Measure Various  
Performance Metrics**

Runtime, Bandwidth, Mvis/s,  
Ops/J, ...



**Running on Different  
Hardware-Systems**



**Benchmarking-Code  
(«Glue») written in  
Python**



**Framework: ReFrame**


# Benchmarks: Reproducibility

What is a  
reproducible  
benchmark?

Which  
information do  
we need to  
know?

How do we  
ensure a  
benchmark is  
reproducible?

# Hardware



## CPU

- Architecture, Number of Cores, Threads per Core, Clock Speed, MIPS, Cache, ...



## Memory

- Size, Bandwidth, Number of Channels, ECC, Speed, ...




## Storage

- Mounts, Built-In vs Network, Size, Utilization, Bandwidth, IOPS, File System, ...




## Network Interfaces

- Interface Types, Bandwidth, Number of Interfaces, ...



## Accelerators

- Type, Location, Bandwidth, Capabilities, Clock Speed, ...



## Topology




## Cluster-Data

- Interconnect-Bandwidth, Round-Trip-Time, Topology (if available), ...

# Software



**Operating System**



**Environment**

- Environment Variables, Scheduler, ...



**Installed Tools**

- Loaded Modules, Activated Spack-Packages / - Environments, Easybuild, ...



**Python-Dependencies**

- If applicable: List of installed packages (pip, conda, poetry), versions, ...



# Execution



**Call Graph**



**Power  
Consumption**



**Tracing**



**Aggregation  
of Data**

# Conclusion



**SDP needs to show what's possible this PI!**



**Gathering everything at every benchmark is not feasible**

Still, much information can be gathered when need arises

Helps debugging issues with Hardware or Software (Environments, Dependencies, ...)

Provides an additional verification for the benchmarks

# Thank you!

