

A Software Development Environment for the CSP



Nicolás Erdödy

Open Parallel Ltd.

Computing for SKA Colloquium 2015

AUT University, Auckland, New Zealand
February 12, 2015

Agenda

- Introduction - Track record.
- Involved with the SKA since 2011.
- What is a SDE? What for?
- An Industrial-Scale Project.
- SKA is important for non-astronomers.
- Why are we doing this?
- What's next?

Past (from 2005)

- Sun Microsystems partners.
- T1 – T2 test bed.
- Work on RDMA, SDN, Snort, etc.
- First Sun Microsystems OpenSPARC Centre of Excellence of the world outside the US (w/University of Otago, Dunedin, NZ 2008).
- INTEL.
- ARM.
- Partnering with Catalyst since 2005 (largest OSS company in Australasia: 300+ devs).

SKA aligned prospects (2015)

- Software on Silicon - SPARC M7 (Oracle – Oracle Labs).
- Bare Metal work / OpenPower Foundation (IBM)
- OpenMP – C/C++ standards.
- Special projects collaboration (MIT).
- NVIDIA testing – Catalyst Cloud.
- NVIDIA – IBM (DoE)

Open Parallel **manages the
Software Development Environment
work package** (part of CSP)
and contributes to the
Common Software (part of SDP).

Open Parallel has been actively involved with the SKA since 2011.

- NZSKAIC member (2011-13).
- RWC2011 – SKA / Open Source workshop – Wellington.
- 1st Multicore World / CSIRO (2012).
- Officially selected by NZ Govt. for co-funding after international peer-review (only NZ Company), together with VUW & AUT (2012).
- Self-funded contribution since 2011.
- Formally part of CSP and SDP since 2013.
- Delivered v1.0 of Software DevPlan for CSP -also partially included in SDP's Preliminary Architecture (2014).

Goal:

To Build a
Common Solution
for a Massively Distributed
Software Development

As part of our contribution to the SKA, we have designed and delivered version 1 of a **proposed Software Development Plan for how to develop and deliver software and/or firmware** to achieve the software design goals established for the SKA project.

The plan includes activities supporting the framework for SKA-wide software and firmware development standardisation and co-ordination.

The Open Parallel Stack (TOPS)

“A formalised standard set of tools, processes and development infrastructure for projects at any scale”

(software bundle)

- To be built with the SKA in mind.
- Adaptable to any architecture.
- Evolves with customer's scaling.
- Bespoke R&D.
- Maybe a “SKA Distro”?
- Run on Catalyst Cloud -powered by OpenStack.

Developing and managing
large scale HPC applications
that typically involve development
across a wide variety
of hardware and software platforms
**is a complex and difficult
development challenge.**

Given the **enormous scope** of the SKA project, it is likely that a good deal of pioneering will be involved in coordinating such an ambitious project **across a variety of platforms and among widely distributed development teams.**

To inform the
Software Development Plan
we conducted a
preliminary survey
of the CSP and SDP teams
in 2014.

The primary objective of the survey was **to identify the widest range of development tools, languages and platforms** that are likely to be employed.

We have the expectation that the many parties involved in the SKA can **converge on a common set of tools and standards to deploy consistently**, thereby facilitating efficient coordination, communication, integration, testing and administration of the overall development process.

We received responses from:

- National Research Council (NRC) – Canada.
- Swinburne University of Technology - Australia.
- NZ SKA Alliance.
- SKA South Africa.
- ICRAR – Australia.
- ASTRON – Netherlands.
- INSPUR – China.
- Space Advisory Company – South Africa.
- CERN – Europe.

The SKA

- World largest, most powerful radio-telescope.
- Image resolution: 50x Hubble Space Telescope.
- Distributed between WA and Southern Africa.

- Also “The Largest Supercomputer of the World”.
- 1+ ExaFlop to process and reduce the massive amount of data generated by the sensors (that’s 10^{18}).
- All dishes together will produce ~360Tb/s.
- Extreme data throughput requirements (8x of today’s largest supercomputers).

In other words:

A huge supercomputer with
a radio-telescope “attached”.

The SKA's power will come from it being a **software-controlled and software-dependent telescope.**

This will allow the SKA's designers to continuously take advantage of advances in **computer power, algorithm design and data transport capability,** always keeping the SKA up to date.

Software Environment and Tools

- Software development plan is required for how software and/or firmware will be developed and delivered in accordance to a design specification
- Determines suitable software development environment/tools, coding standards, templates, revision control methods, etc
- Determines software libraries/tools that will be common across the SKA
- Open Parallel leads in CSP and SDP, involves Catalyst and other NZ expertise
- Spinoffs for local software development industry and tool vendors, head-start developing required software

SDE / CSP– Software DevPlan (v1.0)

Table of Contents

(excerpts)

- Development Processes.
- Development Infrastructure & Environment.
- Software Modelling.
- Development Tools.
- Coding Standards.
- Testing.
- Security.
- Runtime Environments.
- Common Libraries.
- Governance Processes.
- Change Control Management.
- Risk Management.

Team – Contributors to v1.0

- Technical Advisor, Open Parallel (DirectX media platform, massively parallel distributed computing solutions).
- Technical Advisor, Open Parallel (25+ years in scalability and distributed systems).
- Solutions Architect, Catalyst.
- Cloud Engineer, CERN.
- Director, NZ SKA Alliance.
- Assistant Director, CSIRO.
- Former SKAO Software & Computing Domain Specialist.

Reviewed by: - SKAO's SDP & TM Systems Engineer.

- CSP PM.

The right Network = The right Team

Multicore World – a global network



Audience at II Multicore miniconference (LCA2011) in Brisbane, Australia 25 January 2011. In the front row Vint Cerf, Dirk Hohndel and Linus Torvalds listening to Paul McKenney's keynote



4th Multicore World

& workshop

“HPC in the Cloud: the SKA example”

Wellington, 17-18-19 February 2015



**The exascale computational requirements
for the SKA are beyond
the capabilities of existent technologies:**

They are needed to enable
the SKA to scan the skies
thousands of times faster than ever before,
producing vast amounts of valuable data,
at rates in to Tb/s
(100x the global internet traffic).

Some of the talks & panels next week
at **Multicore World 2015 & HPC for SKA:**

- **“Revolutionary Multicore Computing for Exascale”** (UI)
- “Emerging Trends in GPU Computing” (NVIDIA-U of Malaga)
- **“HPC Supporting the Cloud Supporting HPC”** (UI - CREST)
- “The hidden costs of the Parallel World” (Codecentric AG)
- **“SKA imaging Compute Requirements: the whats -and the Watts”** (DH – DIA, fmr SKAO)
- “Accelerating discovery via cloud services” (Argonne-U of Chicago)
- **“Computer-Aided Discovery in Astronomy”** (MIT)
- “Supercomputing in the Cloud” (Nyriad)
- **“HPC Spinoffs from the SKA”** (AUT – NZ SKA Alliance)
- “Is Exascale Computing the Panacea we've all been waiting for?” (Panel)

What can the SKA do for NZ?

- NZ Government initially underestimated significance of New Zealand's global role and technical/scientific/business contribution to the SKA.
- NZ has a very active participation in the SKA -particularly in computing, it isn't “just a follower”.
- A large involvement in the pre-construction phase (currently happening -that needs to keep happening) will ensure a strong NZ participation in the tender for construction (i.e. associated with Canada).
- Juste Retour principle would apply.
- **Broader benefits “Beyond Astronomy”**: create a new IT industry for NZ, technology spillovers, STEM subjects, applications to primary industries, etc.

What's Next?

- Process feedback from PDR into CSP work package.
- Contributing with Catalyst to the SDP.
- Possible Demonstrator Platform (?) -powered by OpenStack, running pilot cases in the Catalyst Cloud (can we test it with your data?).
- Keep defining the space for The Open Parallel Stack: “A formalised standard set of tools, process & development infrastructure for projects at any scale”.
- Showcase (something) at Multicore World 2016.
- Full SDE design (2015-16) & prototype -towards CDR.
- Applications to industry-scale problems in Australia-NZ and globally & keep working towards NZ becoming a global hub / CoE in multicore & parallel computing.

Open Parallel

Software for Multicore and Parallel Computing



MULTICORE WORLD 2015

catalyst 

open source technologists

Committed with the CSP!



“A Software Development Environment for the CSP”

Auckland, New Zealand.
12 February 2015.

Thank you!

Nicolás Erdödy

Founder, Open Parallel.

nicolas.erdody@openparallel.com