



International
Centre for
Radio
Astronomy
Research



CAASTRO
ARC CENTRE OF EXCELLENCE
FOR ALL-SKY ASTROPHYSICS

Future options for SKA-scale astronomical computing

Attila Popping

Collaborators: Richard Dodson (ICRAR), Kevin Vinsen (ICRAR), Chen Wu (ICRAR), Martin Meyer (ICRAR), Andreas Wicenec (ICRAR), Peter Quinn (ICRAR), Jacqueline van Gorkom (Columbia), Emmanuel Momjian (NRAO).

Second ngVLA Technical Workshop
December 8-9 2015
Socorro, New Mexico, USA



Curtin University



THE UNIVERSITY OF
WESTERN AUSTRALIA



Data is the new challenge




CHILES



x10

VLASS



x10

ngVLA

x100

Data is the new challenge




CHILES



x10

VCLASS



x10

ngVLA



x10




WALLABY/DINGO/ASKAP



x10

SKA

x100





- Small teams or individuals are replaced by large international teams
- Requirements are outgrowing the capabilities of a single institute.
- The capability of storing and processing data will become equally important as the observations or the telescope.
- Manual data reduction has to be replaced by automated workflows



- Small teams or individuals are replaced by large international teams
- Requirements are outgrowing the capabilities of a single institute.
- The capability of storing and processing data will become equally important as the observations or the telescope.
- Manual data reduction has to be replaced by automated workflows
- All this is challenging and daunting
- It will take time to develop





- Small teams or individuals are replaced by large international teams
- Requirements are outgrowing the capabilities of a single institute.
- The capability of storing and processing data will become equally important as the observations or the telescope.
- Manual data reduction has to be replaced by automated workflows
- All this is challenging and daunting
- It will take time to develop

CHILES is the perfect test particle!

COSMOS HI Large Extragalactic Survey

CHILES: $0 < z < 0.5$

single pointing in COSMOS
1000 hours integration
VLA B-configuration



	OLD	PILOT	NEW
Bandwidth (MHz)	6.25	240	480
Channels	31	16384	30720
Velocity resolution (km/s)	40	3.5	3.5
Instantaneous z coverage	$0 < z < 0.004$	$0 < z < 0.193$	$0 < z < 0.5$



178 hours observed and reduced in Fall 2013
 ~200 hours observed in Spring 2015
 ~230 hours allocated for Spring 2016

Imaging SKA-Scale data in three different computing environments

Richard Dodson^{a,*}, Kevin Vinsen^a, Chen Wu^a, Attila Popping^b, Martin Meyer^a, Andreas Wicence^a, Peter Quinn^a, Jacqueline van Gorkom^c, Emmanuel Momjian^d

^aInternational Centre for Radio Astronomy Research (ICRAR),
The University of Western Australia,
M468, 35 Stirling Highway, Crawley, Perth, WA 6009, Australia

^bAustralian Research Council, Centre of Excellence for All-sky Astrophysics (CAASTRO)

^cDepartment of Astronomy,

Columbia University,
Mail Code 5246, 550 West 120th Street,
New York, New York 10027, USA

^dNational Radio Astronomy Observatory,
1003 Lopezville Rd., P. O. Box 0,
Socorro, NM 87801, USA

Abstract

We present the results of our investigations into options for the computing platform for the imaging pipeline in the CHILES project, an ultra-deep HI pathfinder for the era of the Square Kilometre Array. CHILES pushes the current computing infrastructure to its limits and understanding how to deliver the images from this project is clarifying the Science Data Processing requirements for the SKA. We have tested three platforms: a moderately sized cluster, a massive High Performance Computing (HPC) system, and the Amazon Web Services (AWS) cloud computing platform. We have used well-established tools for data reduction and performance measurement to investigate the behaviour of these platforms for the complicated access patterns of real-life Radio Astronomy data reduction. All of these platforms have strengths and weaknesses and the system tools allow us to identify and evaluate them in a quantitative manner. With the insights from these tests we are able to complete the imaging pipeline processing on both the HPC platform and also on the cloud computing platform, which paves the way for meeting big data challenges in the era of SKA in the field of Radio Astronomy. We discuss the implications that all similar projects will have to consider, in both performance and costs, to make recommendations for the planning of Radio Astronomy imaging workflows.

Keywords: methods: data analysis, Parallel Architectures: Multicore architectures, Distributed architectures: Cloud computing, CHILES

*Corresponding Author

Email addresses: richard.dodson@icrar.org (Richard Dodson), kevin.vinsen@icrar.org (Kevin Vinsen), chen.wu@icrar.org (Chen Wu), attila.popping@icrar.org (Attila Popping),

Preprint submitted to Astronomy and Computing

November 3, 2015



Dodson et al. 2015

Imaging SKA-Scale data in three different computing environments

arXiv: 1511.0040

- 1) Local Cluster (Pleiades)
- 2) Super Computer (Magnus)
- 3) Cloud computing (AWS)



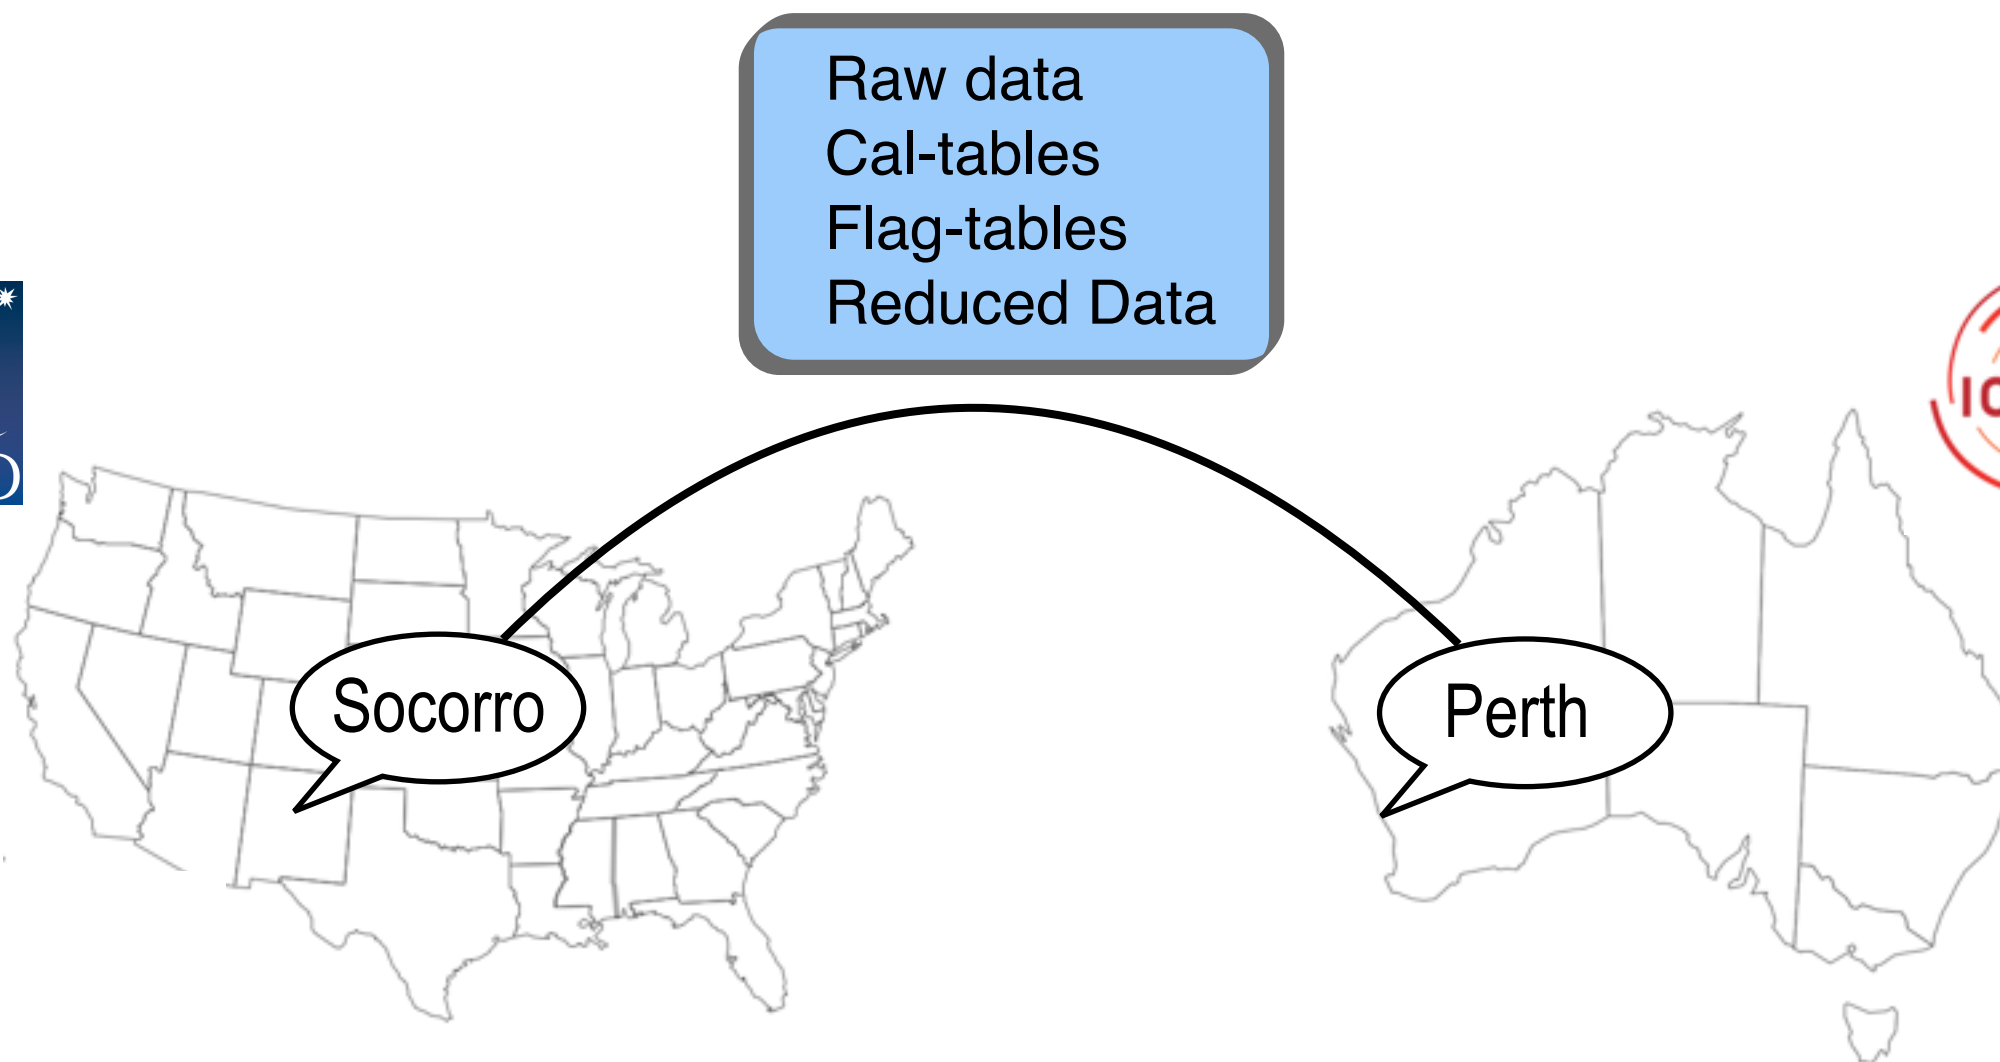
CHILES workflow



Observing
Calibration
Flagging
Quality Control





CHILES workflow




Observing
Calibration
Flagging
Quality Control

Backup
Combination
Imaging











Data reduction (imaging)

- Exclusively done using CASA
- A lot of decisions we made were to get around ~~limitations~~ features of CASA
- We trialled 3 ways of splitting out the data
 - SPLIT
 - CVEL
 - MSTRANSFORM
- Tried to keep a common code base in GitHub
 - PBS - Pleiades
 - SLURM - Magnus
 - Python/Boto - AWS

<p>Conventional Cluster (pleiades) 5 nodes each node has 2x Intel Xeon X5650 2.66GHz CPUs (6 cores / 12 HTs)</p>	<p>Enough computing power, however it would take weeks</p>
<p>Super computer (MAGNUS) Cray XC40 - 24 cores per node 2.6GHz Intel Xeon E5-2690V3 64GB per Node 35,712 cores available 3PB of storage #58 in the world</p>	
<p>AWS Whatever we wanted r3.xlarge 16 cores 122GB Ram</p>	



	On demand	Spot Price
r3.4xlarge	\$1.68	\$0.20
r3.2xlarge	\$0.840	\$0.09
m3.xlarge	\$0.392	\$0.04
m3.medium	\$0.098	\$0.01




Works!
costs so far : ~\$2000

Spot Instance Pricing History

Product : **Linux/UNIX** ▾ Instance type: **r3.4xlarge** ▾ Date range : **1 week** ▾ Availability zone: **All zones** ▾




We have now combined 42 observing runs (~ 20 Tb)
and imaged $2048 \times 2048 \times 31,000$ pixels (~ 500 Gb),
covering the redshift range $z=0 \sim 0.5$



1020 - 1040 MHz

We have now combined 42 observing runs (~ 20 Tb)
and imaged $2048 \times 2048 \times 31,000$ pixels (~ 500 Gb),
covering the redshift range $z=0 \sim 0.5$



1020 - 1040 MHz

Operation	Platform	Peak Memory	I/O Throughput	CPU Usage	I/O Characteristics
SPLIT	AWS (EBS)	420MB	<10MB/s	40%	Sequential read/write dominate
	Magnus	545MB	40 ~ 100 MB/s	100%	
	Pleiades	390MB	60 ~ 100 MB/s	100%	
INVERT	AWS (SSD)	60GB	70 ~ 500MB/s	400%	Random writes and sequential reads dominate
	Magnus	30GB	50 ~ 400MB/s	100%	
	Pleiades	35GB	50 ~ 400MB/s	400%	

	AWS	Magnus (HPC)	Pleiades
Completion Time	96hr	110hr	1,060hr (est)
Capital Costs	AUD\$0	AUD\$12,000,000	AUD\$50,000
Operational Costs	AUD\$2,000	AUD\$3,240 (free)	-
Control	Root	Limited	Root
Usability	Complex	Good	Good

CHILES - The return of the Jalapeño. (SKA-AWS project)

- Using more AWS working with NRAO
 - NRAO to upload Measurement Sets to AWS Oregon
 - Code run at the Oregon data centre from Perth
 - Scaling Groups will fire off the instances to split the code
 - Code rather than “a person” looking at the spot prices
- Using SKA-SDP Prototype code
 - Docker
 - DROPS
 - CasaPy
- JPEG2000 Visualisation

- Department Cluster
 - Not very satisfactory
- HPC
 - Very fast
 - You can't have root access
 - Installing new software is done by the admin's
 - In WA it is effectively free
- Cloud
 - You can do what you like (a good and a bad thing)
 - EBS volumes are slow
 - Directed attached SSDs are very quick
 - You pay for what you use... And if you forget to turn it off you are still paying

