Astronomical Science, Astronomical Challenges @Pawsey Supercomputing Centre

Ugo Varetto, *Chief Technology Officer*

Pawsey Supercomputing Research Centre

Agenda

Pawsey	Who
Projects	What
Challenges	☹
Strategy	How
Future	?

The Pawsey Supercomputing Centre

pawsey

Pawsey Supercomputing Centre

Tier-1 Research Facility & Tier-0 Data Centre

Headquartered in Perth, Western Australia

Launched as Pawsey in 2014 with foundations back to 2000

Critical support for SKA infrastructure on Wajarri Yamatji country, 600km north of Perth

AU$70m capital refresh by Australian Government

50+ staff employed via CSIRO in UJV

1600+ users, ~200 projects

An emerging Indo-Pacific HPC Leadership Cluster

Pawsey Supercomputing Centre

200 Research Projects

42 Research Institutions

15+ Scientific Fields

194 Training Programs & Events

24/7 Data Ingest

Simulating atoms for green energy

Dr Yun Wang

Students work with big data, survey teen sleep

Dr Linda McIver, Australian Data Science Education Institute

Kicking goals with GPUs

Dr Sajib Mistry and Associate Professor Aneesh Krishna, Curtin University

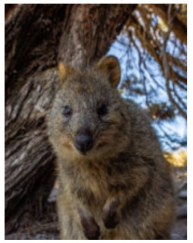

Mapping DNA to protect an iconic Australian species

Associate Professor Parwinder Kaur, Australian DNZ Zoo

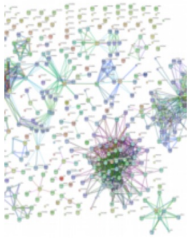

Charting how we feed a future world

Prof. David Edwards and Mr Philipp Bayer, University of Western Australia

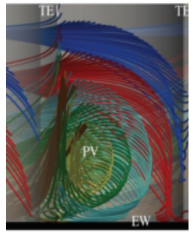

Propelling the environmental efficiency of jet engines

Prof. Richard Sandberg, Chair of Computational Mechanics, University of Melbourne

https://pawsey.org.au/case_studies

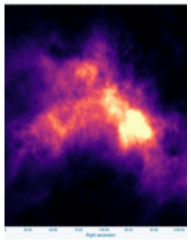

To boldly go where no supercomputer has gone before

Dr Chenoa Tremblay, Postdoctoral Fellow in Dark Magnetism. CSIRO

Tracking denizens of the deep to protect their future

Dr Ana Sequeira, University of Western Australia

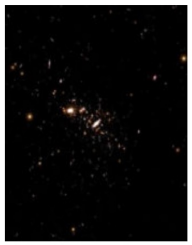

How do galaxies evolve, and is ours 'normal'?

Dr Claudia Lagos, ICRAR - UWA

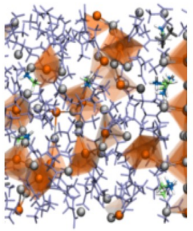

Capturing sunlight with supercomputing

Dr Widmer-Cooper, University of Sydney node of the ARC Centre of Excellence in Exciton

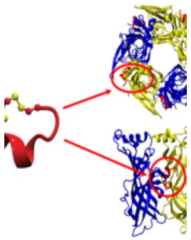

Turning toxins into treatments

Dr Andrew Hung, RMIT University

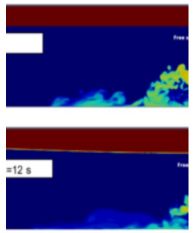

Safeguarding sea habitats from extra salinity

Dr Jason Antenucci, DHI Water and Environment

HPC @ Pawsey

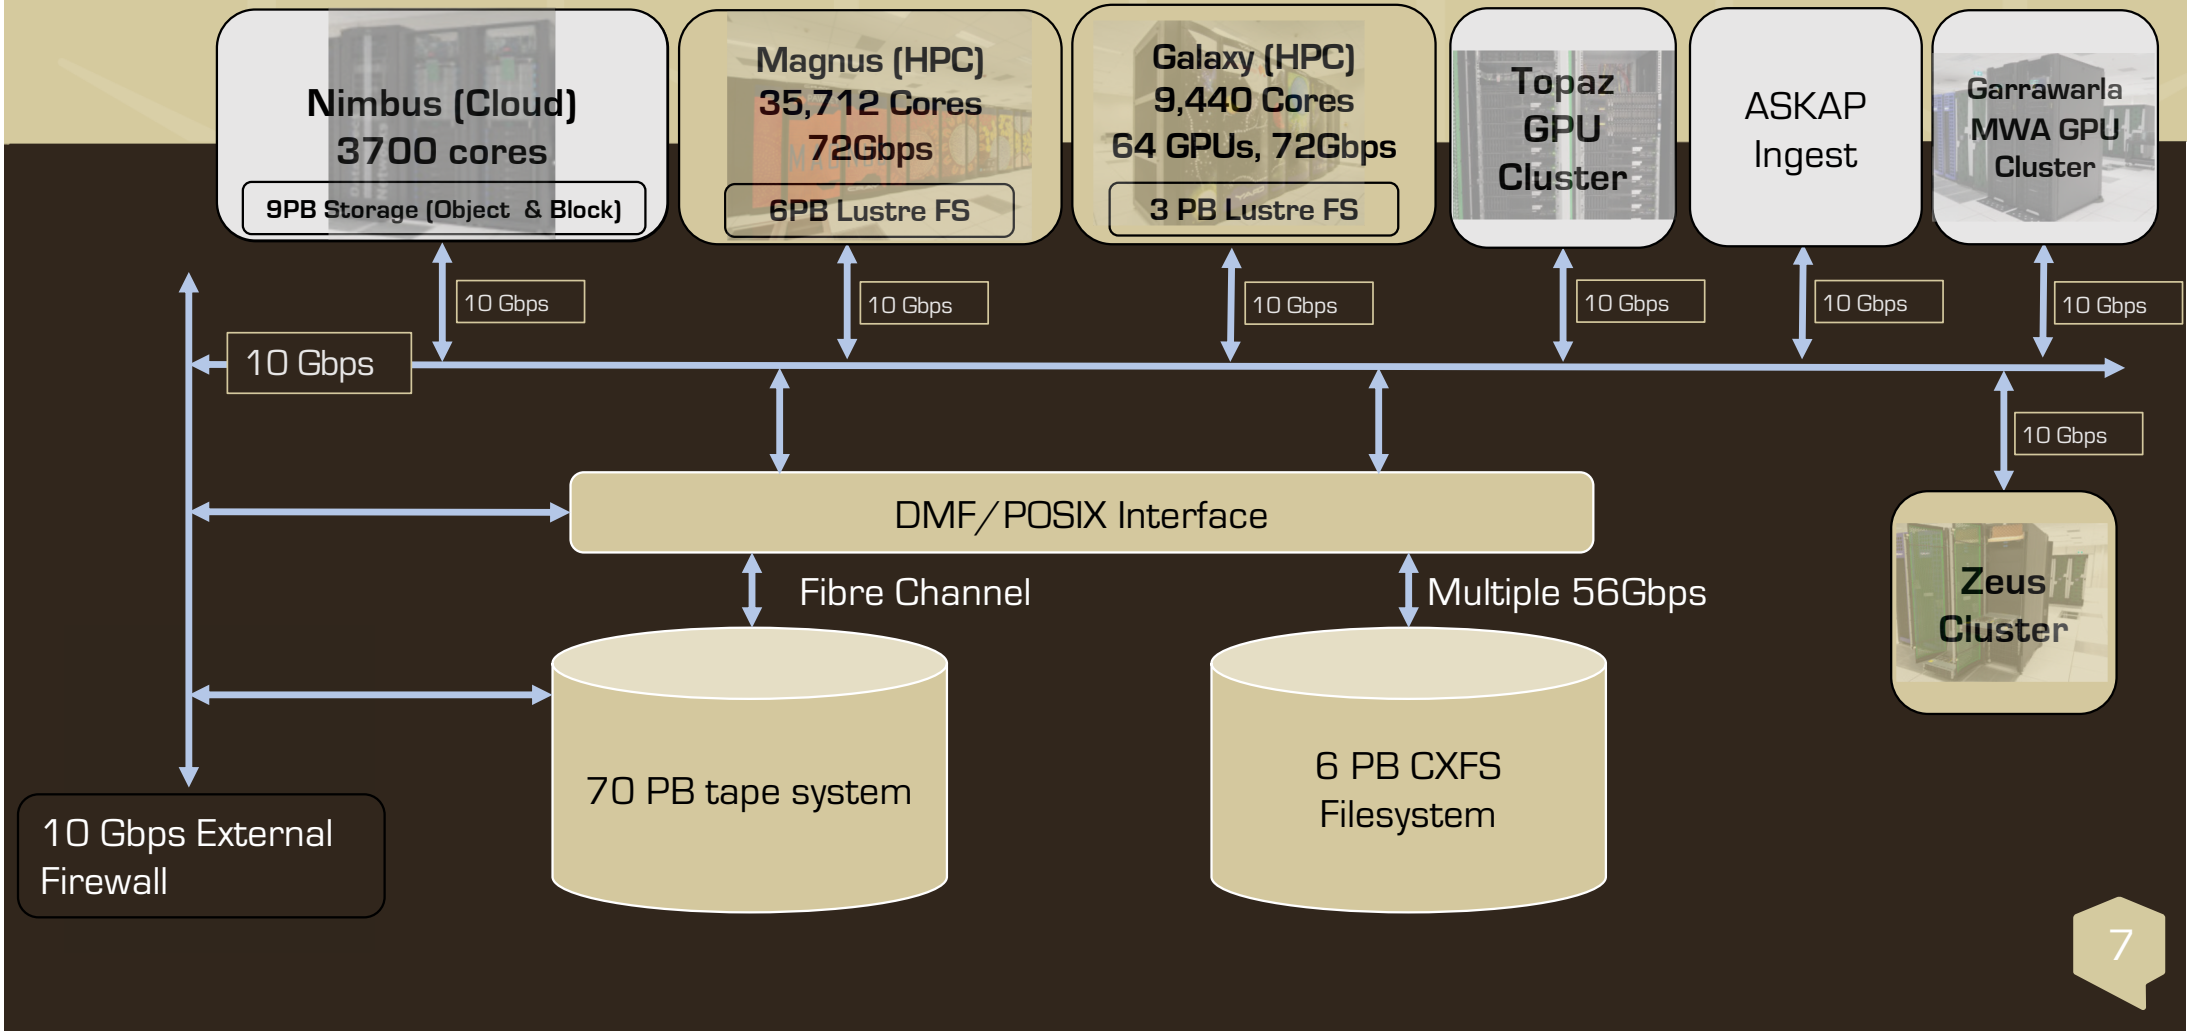
Nimbus (Cloud)
3700 cores
9PB Storage (Object & Block)
Magnus (HPC)
35,712 Cores
72Gbps
6PB Lustre FS
Galaxy (HPC)
9,440 Cores
64 GPUs, 72Gbps
3 PB Lustre FS
Topaz
GPU
Cluster
ASKAP
Ingest
Garrawarla
MWA GPU
Cluster
10 Gbps
10 Gbps
10 Gbps
10 Gbps
10 Gbps
10 Gbps
10 Gbps
10 Gbps
DMF/POSIX Interface
Fibre Channel
Multiple 56Gbps
Zeus
Cluster
70 PB tape system
6 PB CXFS
Filesystem
10 Gbps External
Firewall

HPC @Pawsey – Allocations

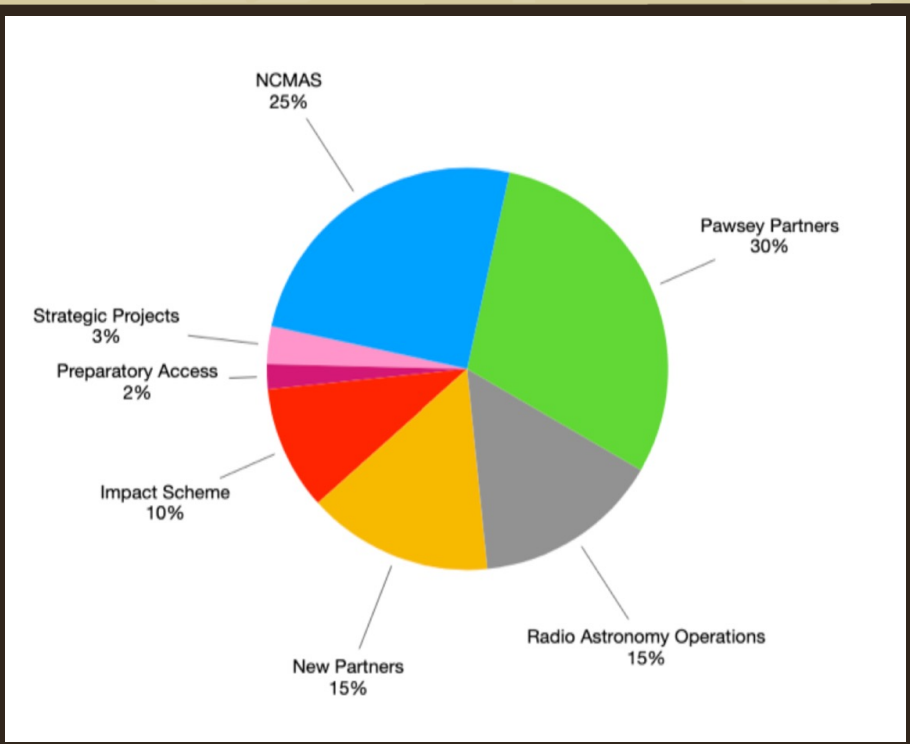

Allocations: ~413 million service units / year

Storage:

- 70 PB offline (mirrored)
- ~11 PB online

Radioastronomy operations:

- 134+ million service units / year
- ~80% long term storage (tape)

Projects

pawsey

Radioastronomy @Pawsey

SKA precursors

Data ingest

Pre/post-processing

Visualisation

Storage

Data lifecycle management

Data sharing (FAIR)

Murchinson Radio-astronomy Observatory

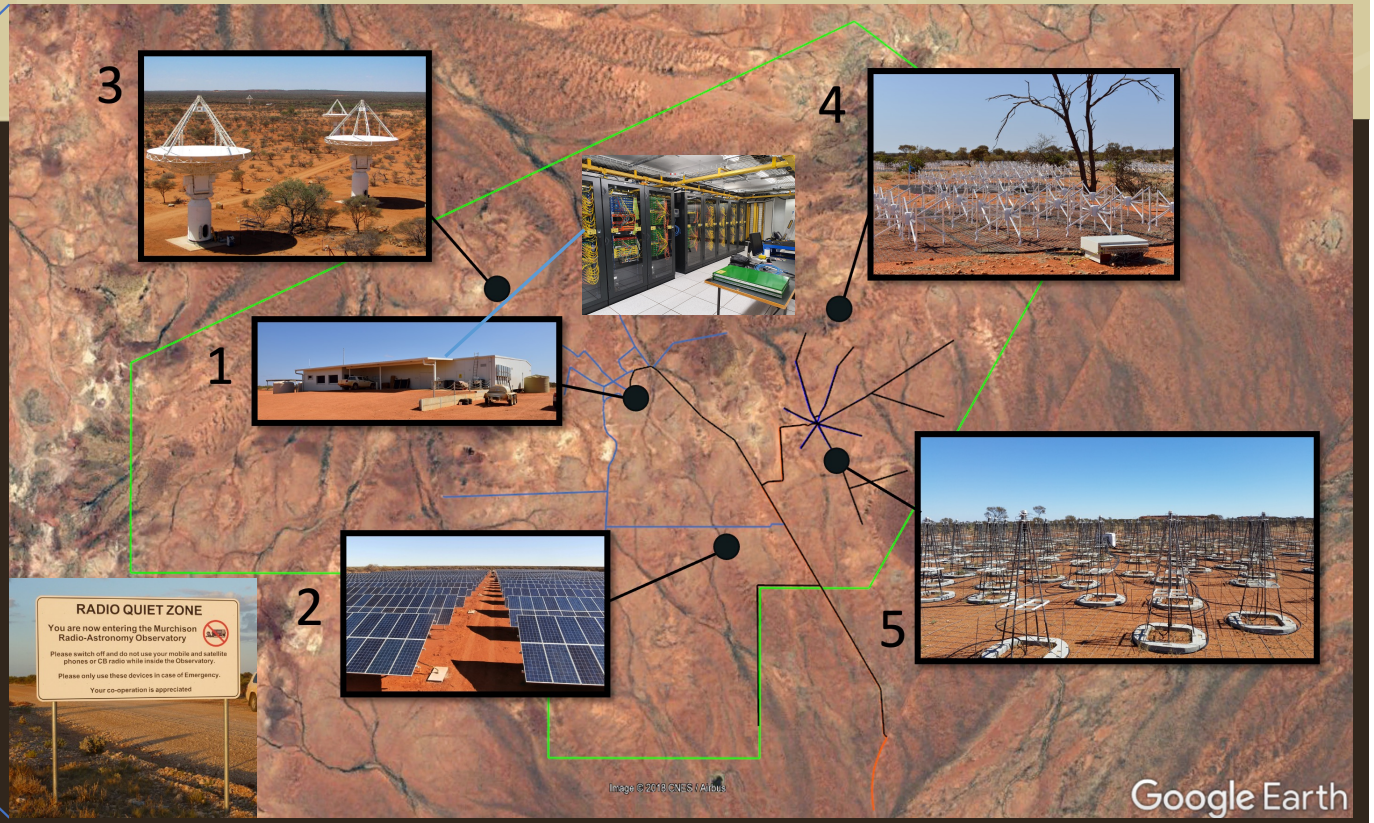

Murchinson Widefield Array

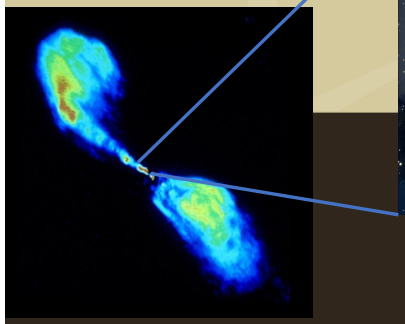

Radio “light”

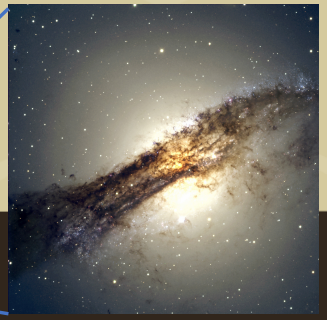

Visible light

70-300 MHz, ~m wavelength

Space Situational Awareness

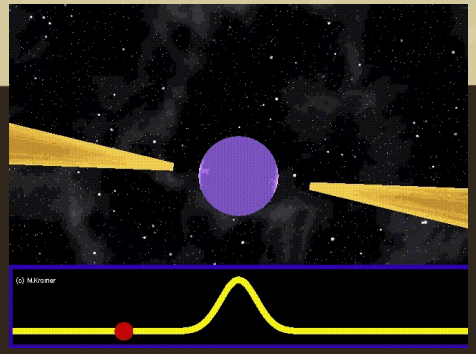

Fast radio bursts

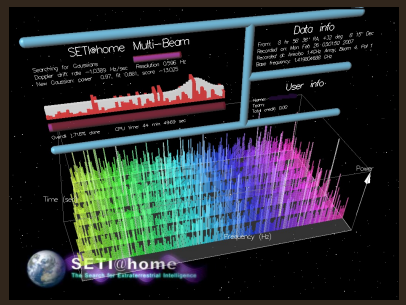

SETI

Epoch of Reionisation

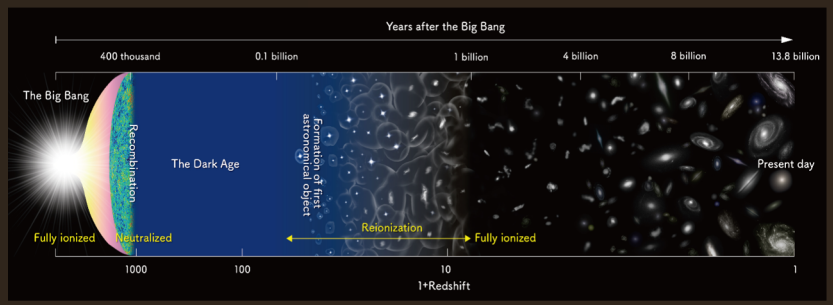

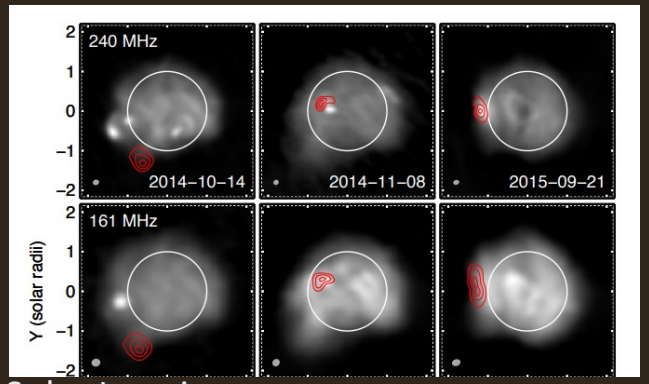

Solar Imaging

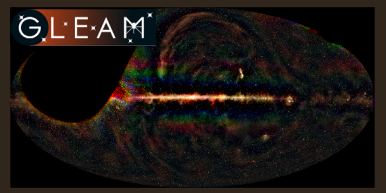

Galactic and Extragalactic All-Sky MWA Survey

Australian Square Kilometre Array Pathfinder

700 MHz – 1.8 GHz, ~cm wavelength

ASKAP EMU Evolutionary Map of the Universe

DINGO

WALLABY

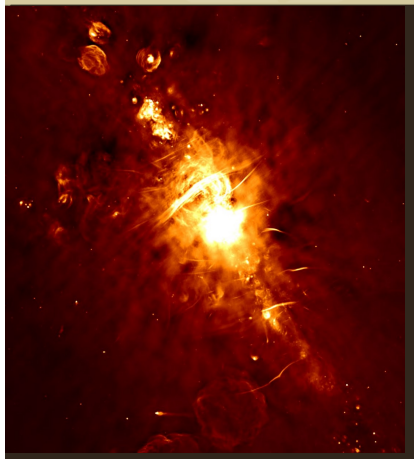

Centre of the Milky Way Galaxy, 28 antennas: *ASKAPSoft*/Wasim Raja.

GASKAP

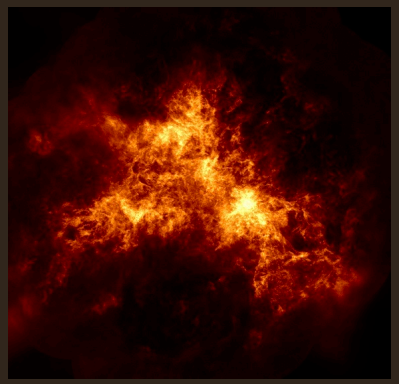

Atomic hydrogen gas in the Small Magellanic Cloud as imaged with CSIRO's Australian Square Kilometre Array Pathfinder (ASKAP). Credit: ANU and CSIRO

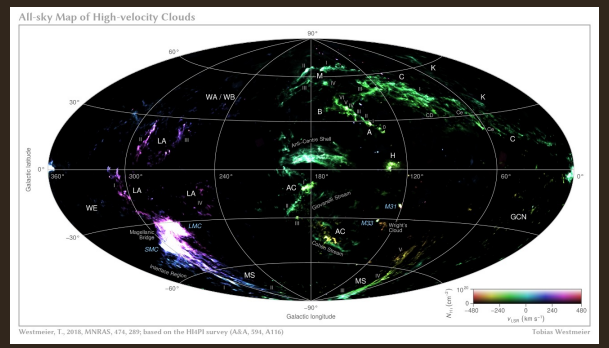

Data Processing Pipeline

1. Receive Signal
2. Amplify Signal
3. Digitise Signal
4. Combine Signals
5. Store Observation
6. Calibrate, post-process, analyse…
   → Science

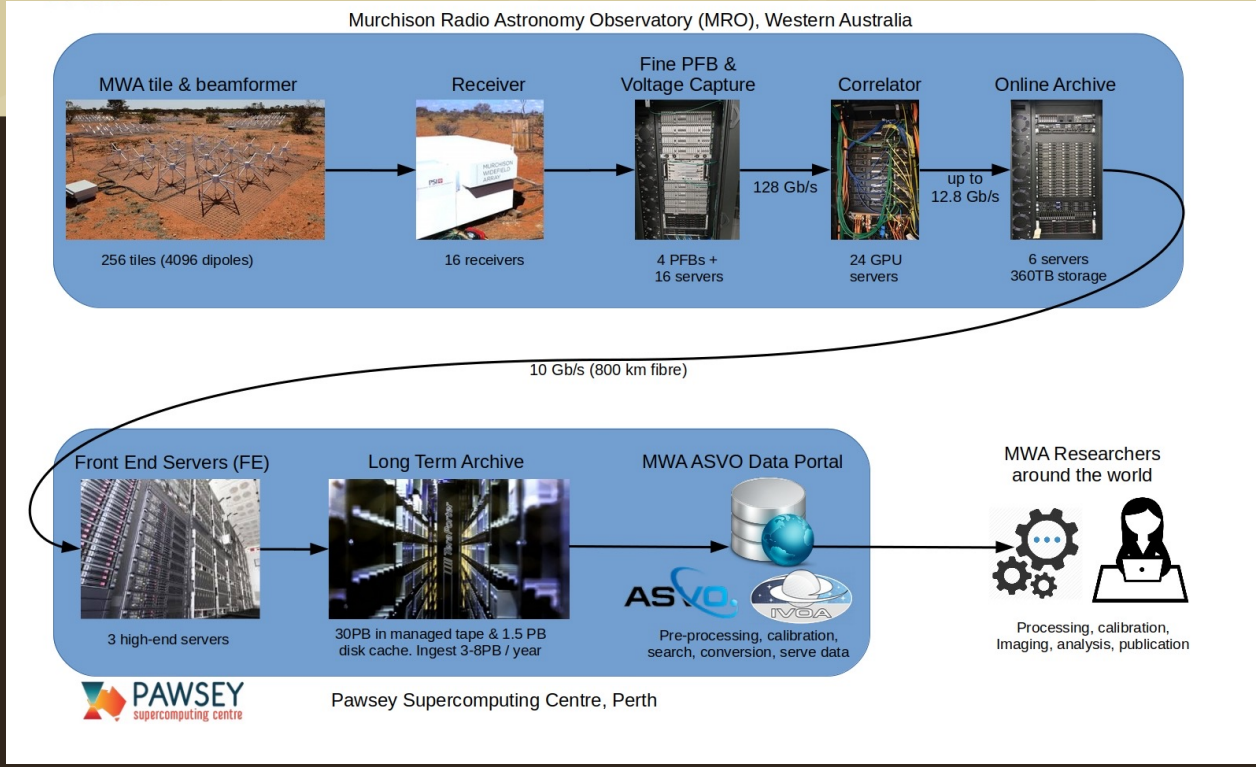

Credit: MWA

Post-processed observation (data cube, all frequencies, double precision): ~2TB

Data Processing Pipeline

Growing bandwidth (TB/s)
Growing data size (10^2 PB)
Large-scale processing workflows (denoising, deconvolution, calibration…)

Low-latency interconnect
Fast I/O
Fast compute performance
High-availability

Science-ready data

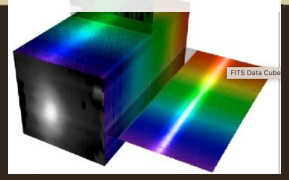

Credit: Gary Resnick, Scott Badenhorst
University of Capetown

Flexible post-processing workflows at scale

Chainer

ParaView

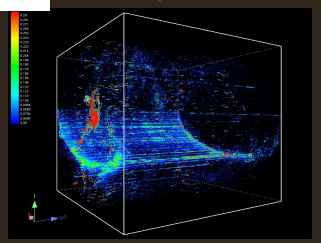

Credit: Russell Jurek (ATNF),
Amr Hassan (Swinburne University

Low-latency, high speed, high-bandwidth, software-configurable environment

Data Processing Pipeline

Future model:

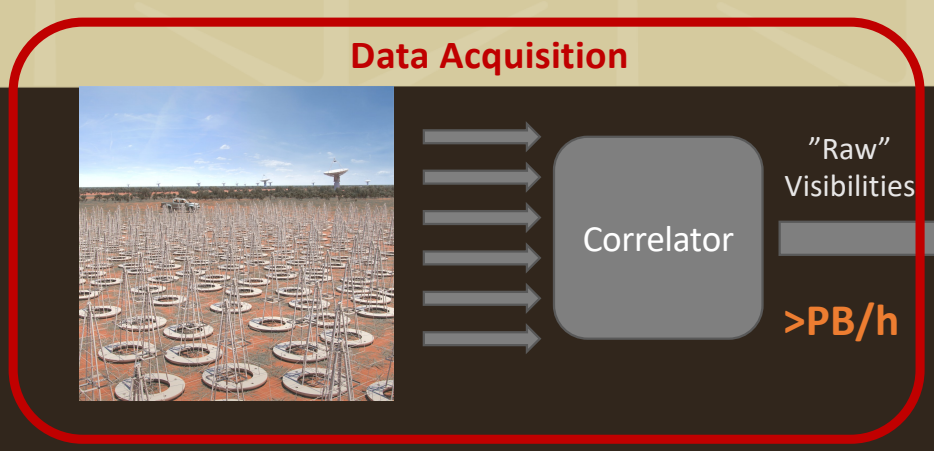

Credit: J.C. Guzman

Integrated HPC –Storage environment

All pre and post processing on HPC systems

Astronomy @Pawsey

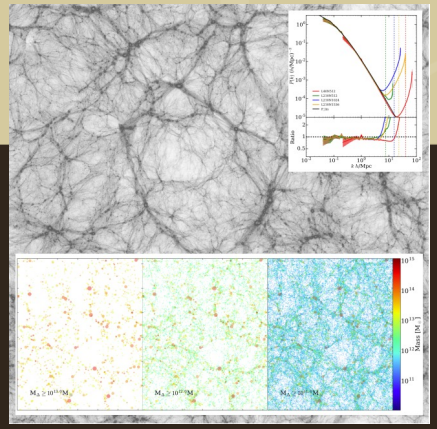

SURFS: Riding the waves with Synthetic UniveRses For Surveys
Pascal J. Elahi et al.

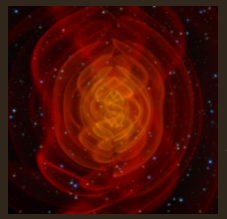

Real-Time Gravitational Wave Search
Prof. David Blair & Prof. Linqing Wen, University of Western Australia

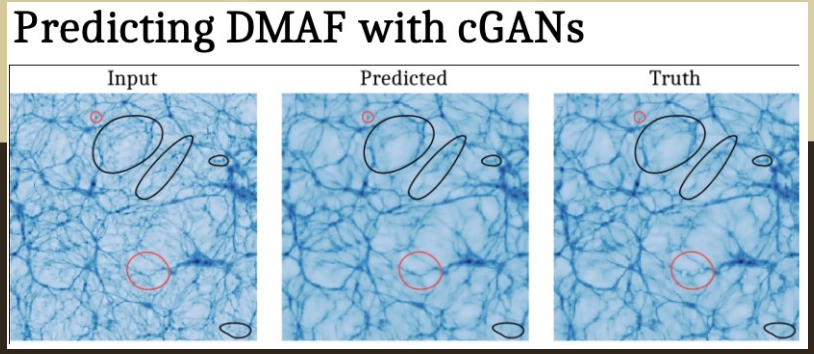

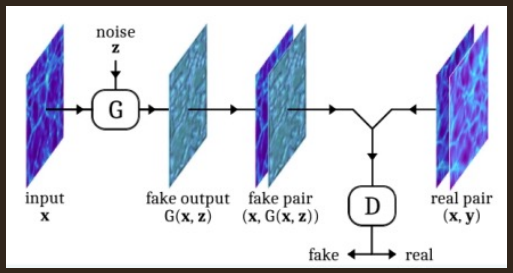

A Novel Scheme for Dark Matter Annihilation Feedback in Cosmological Simulations
Floarian List, Geraint F.Lewis, Nikolas Iwanus, **Pascal Elahi**, Ishaan Bhat

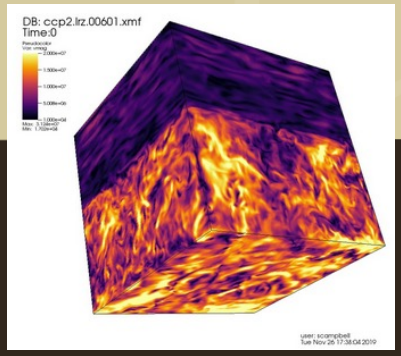

Convective-Reactive Nuclear Burning and Turbulence Boundaries in Stars
Simon Campbell - Monash University

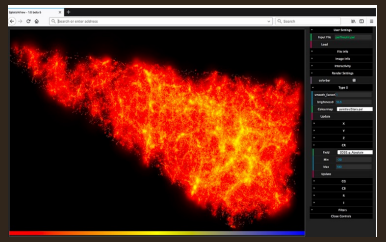

Real-time Web-based Remote Interaction with Active HPC Applications
Tim Dykes, Ugo Varetto et al.

Challenges

pawsey

Challenges – Science

Auto-tuning Parameter Choices in HPC Applications using Bayesian Optimization Harshitha Menon , Abhinav Bhatele , Todd Gamblin. Lawrence Livermore National Laboratory

P. Düben; U. Modigliani; A. Geer; S. Siemen; F. Pappenberger; P. Bauer; A. Brown; M. Palkovic; B. Raoult; N. Wedi; V. Baousis Machine learning at ECMWF: A roadmap for the next 10 years

Kashinath K *et al*. 2021 Physics-informed machine learning: case studies for weather and climate modelling. *Phil. Trans. R. Soc. A* **379**: 20200093. https://doi.org/10.1098/rsta.2020.0093

Current:

- Modelling and Simulation
- Visualisation
- Data Management
- High Throughput Computing
- Machine Learning
- Real-time Data Processing

Future:

- High Resolution Simulations + ML
- In-situ Visualisation
- (FAIR), Federated Data Management
- Higher Throughput Computing
- Large Scale Machine Learning
- Low latency, Large Scale, (24x7) Real-time Data Processing

Challenges – Ecosystem

Real-time data ingest

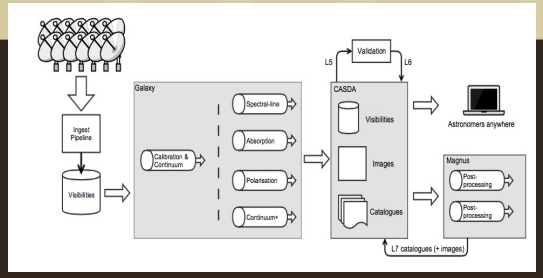

Credit: JC Guzman (ASKAP)

Simulation

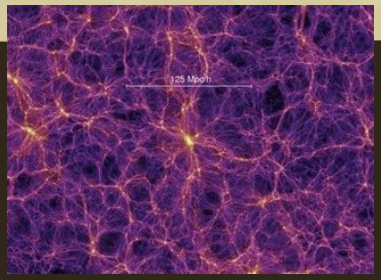

Credit: Springel et al. (2005)

Visualisation

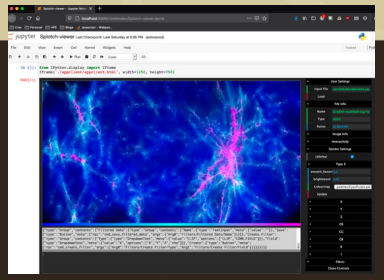

Credit: Tim Dykes (HPE), Ugo Varetto

Machine Learning

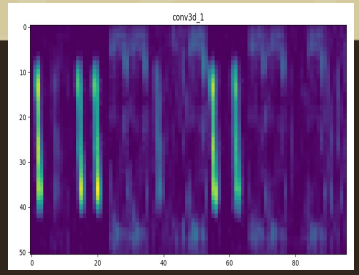

Credit: Observatoire de Paris

Dissemination

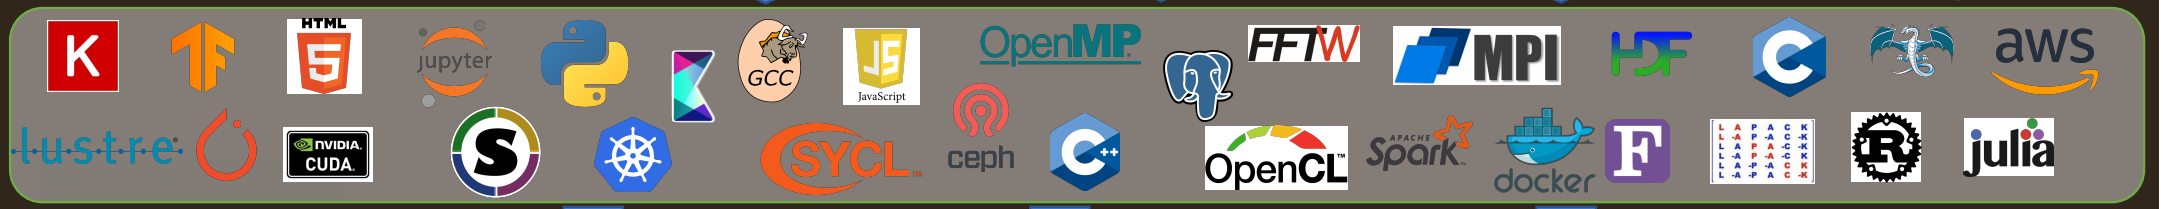

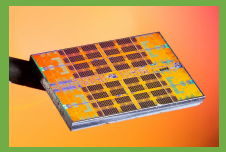

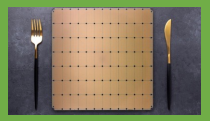

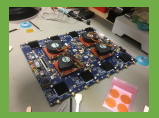

Challenges – Operations

Increase Utilisation

Concurrent Heterogeneous workloads
Variable Resource Demand
Complex Workflows
Flexible Allocation Schemes
QoS

+

Reduce Cost

Energy Efficiency
Simplify System Management
Easy integration and expansion

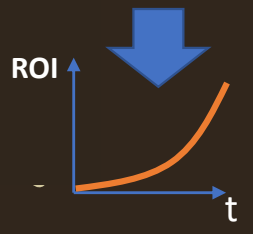

Strategy

pawsey

Operations

Open Standards
Interactive
Elastic
24/7
Expand vs Replace

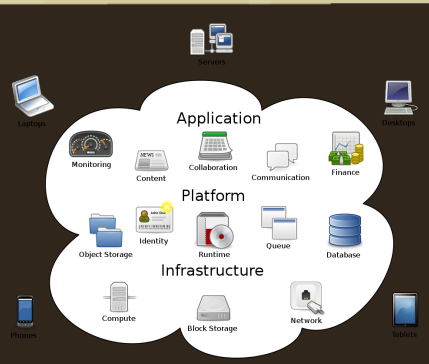

Low-latency (RDMA)
Fast I/O
High Performance

All images:

Setonix would like to say hello

Australia's new 50 petaflop system

200000+ AMD Milan (CPU) Cores

750+ AMD MI-Next GPUS (128GB/GPU)

548+ TB system memory

Near-node NvME storage

15 PB ClusterStor Lustre filesystem with 2.7PB SSD

(90 PB (~60 PB with 8+3 erasure coding) Ceph storage/S3)

Integrated CPU-GPU-NIC (IF)

Software configurable HPC environment

Open Standards

Quality Of Service

Congestion Control

HPC @Pawsey

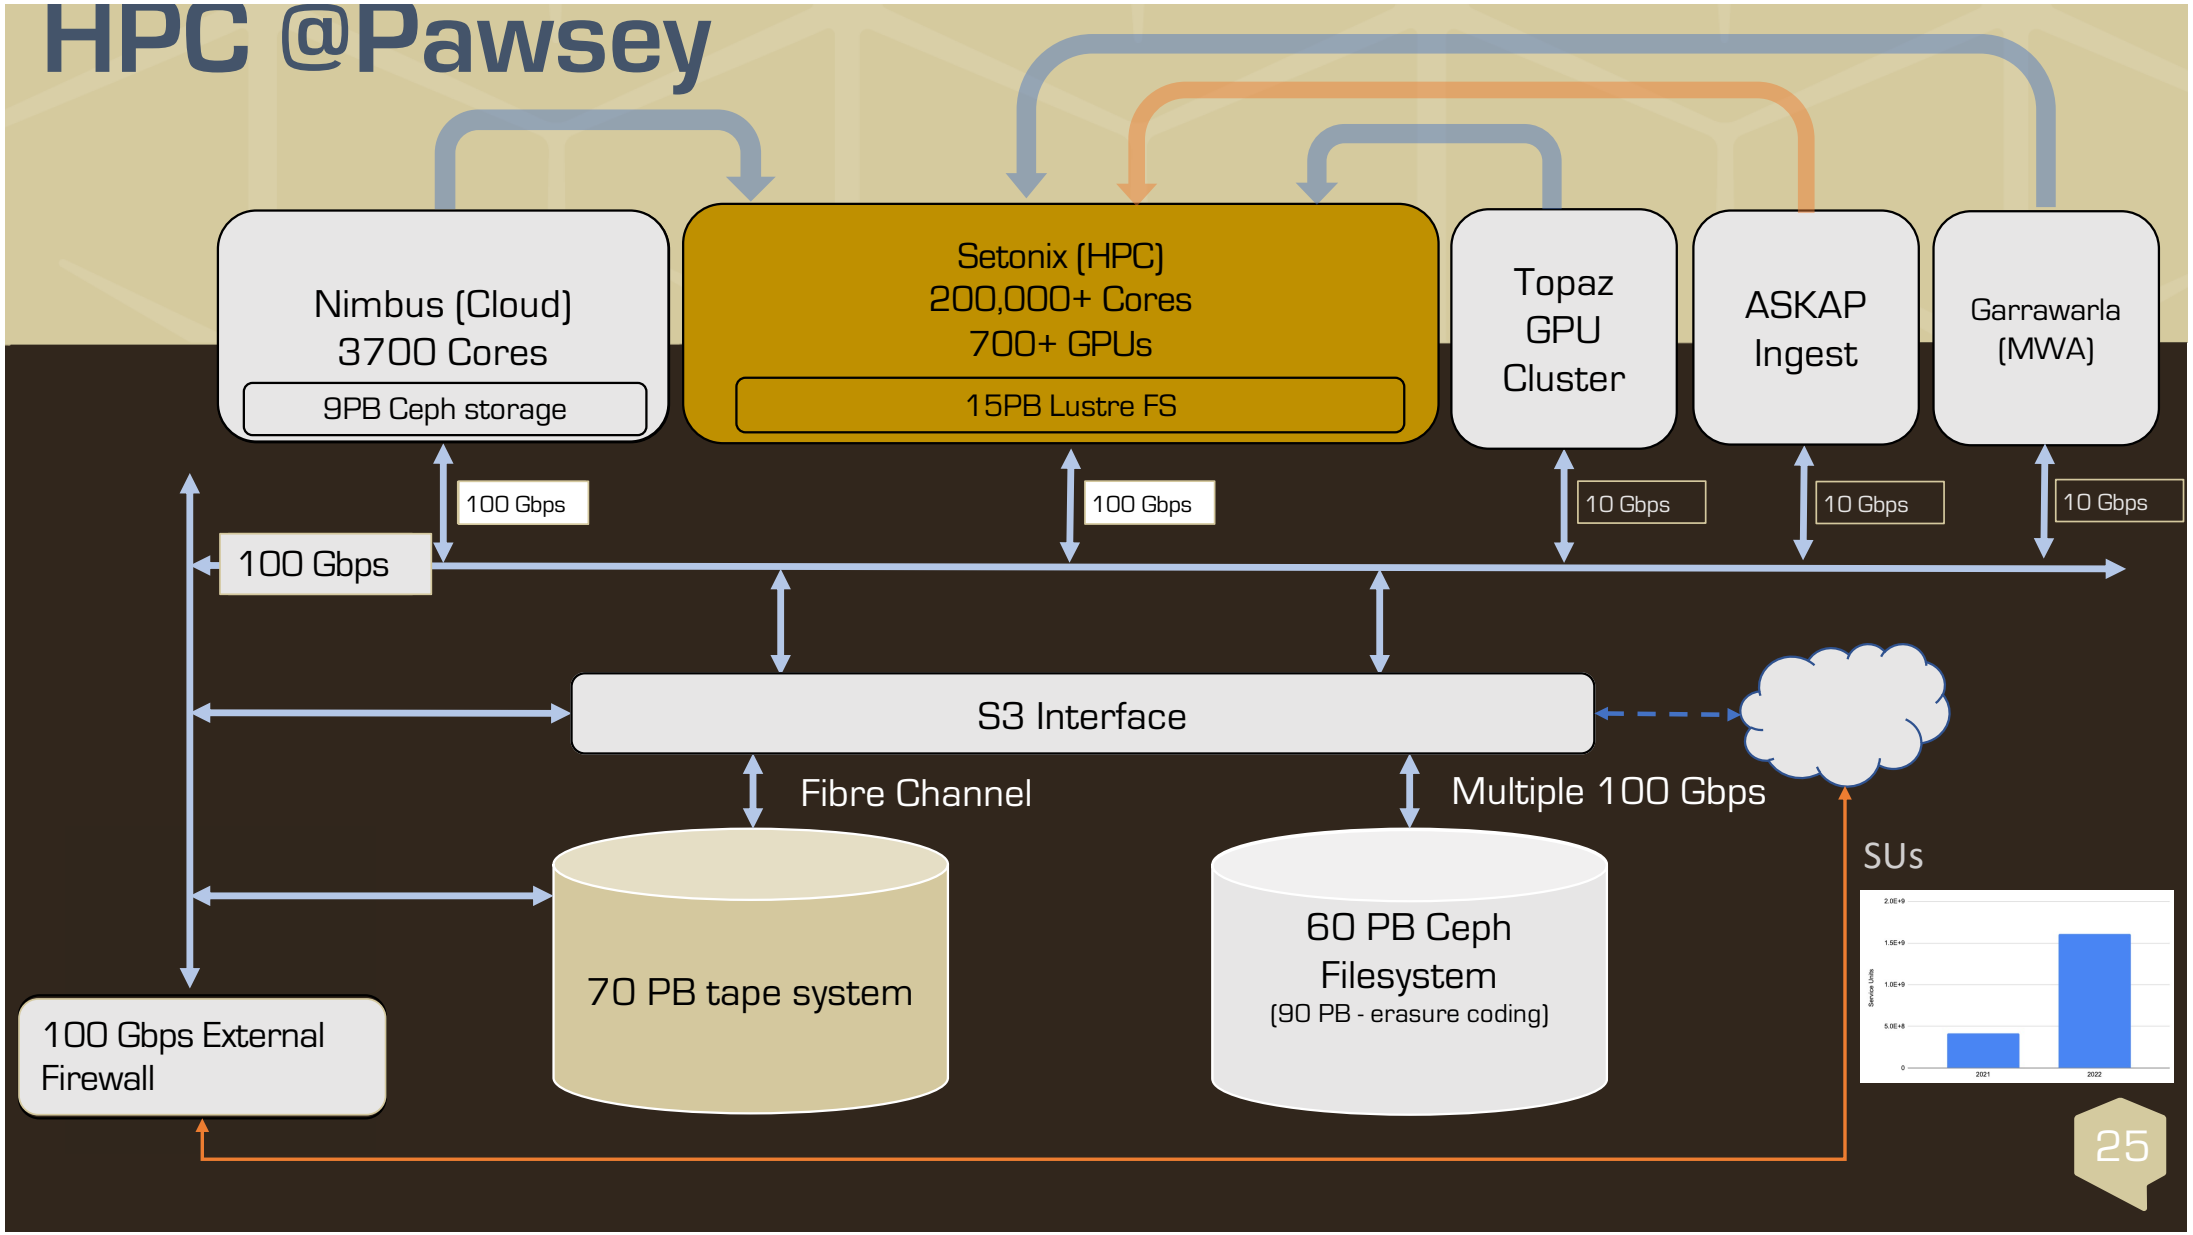
Nimbus (Cloud)
3700 Cores
9PB Ceph storage
Setonix (HPC)
200,000+ Cores
700+ GPUs
15PB Lustre FS
Topaz
GPU
Cluster
ASKAP
Ingest
Garrawarla
(MWA)
100 Gbps
100 Gbps
10 Gbps
10 Gbps
10 Gbps
100 Gbps
S3 Interface
Fibre Channel
Multiple 100 Gbps
70 PB tape system
60 PB Ceph
Filesystem
(90 PB - erasure coding)
100 Gbps External
Firewall
SUs

Engagement

Training and Education

Joint R&D: Embed staff in research groups

Collaboration

Cross cutting: Learn from other domains

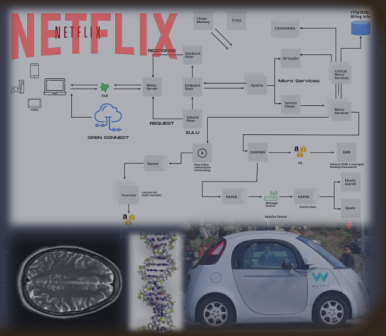

Engagement – PaCER

PaCER: *Pawsey Centre for Extreme Scale Readiness*

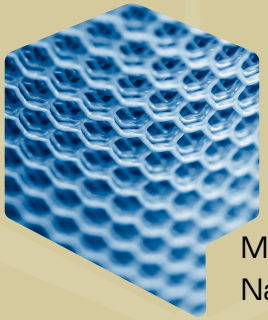

Material Science,
Nanomaterials

Engineering

Molecular Physics

Virology

Ann Backhaus
Education Manager

Radio Astronomy

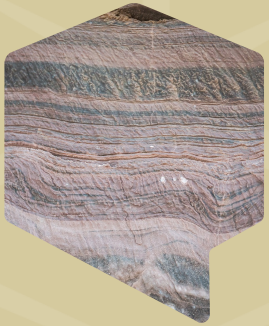

Geoscience

Particle Physics

Maciej Cytowski
Head of Scientific Services

pawsey

PaCER Projects

Principal Investigator/s	Institution/s	Project Title
Richard Sandberg Evatt Hawkes	University of Melbourne, UNSW, GE Global Research	Towards exascale simulations for efficient, low-emissions gas turbines
Giuseppe Barca	ANU, Argonne National Labs, Ames National Labs, Monash, Flinders, Deakin	The Extreme-Scale Electronic Structure System (EXESS): Predicting the Chemistry of Nanomaterial Interfaces
Martin Meyer	ICRAR, UWA, CSIRO, SKA, Oak Ridge National Labs, AusSRC	Delivery of a next-generation data storage approach to unlock deep SKA and pathfinder observations
Christopher Leonardi	University of Queensland	Massively Parallel Models of Particle Suspensions
Melanie Johnston-Hollitt	Curtin University, University of Toronto	Parallel Interferometric GPU Imaging
Igor Bray	Curtin University	Calculation of collisions with molecular targets using the convergent close-coupling method
Marcin Sokolowski	ICRAR, Curtin University, AusSRC	BLINK and you'll miss it: blazingly-fast all-sky radio astronomy pipelines
Waseem Kamleh	University of Adelaide	Emergent Phenomena Revealed in Subatomic Matter
Debra Bernhardt	University of Queensland	Towards a molecular level understanding of flow-induced physical and chemical reactions
Pat Scott	University of Queensland, Monash University, University of Adelaide	Searching for New Particles From the Attoscale to the Exascale with GAMBIT
Robert Edwards	Flinders University	Computational exploration of the virosphere by deep data mining the sequence read archive
Juan Guzman	CSIRO	Exa-scaling the Science Data Processing (SDP) pipelines to support the operations of the ASKAP Full Survey Science Program and the future SKA

BLINK: Breakthrough Low-latency Imaging with Next-generation Kernels

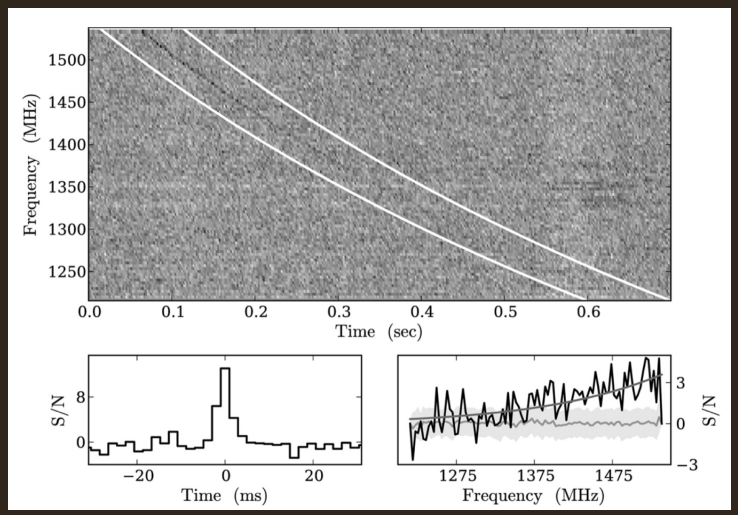

Challenges and Goals

Develop data processing pipelines and imaging algorithms for detection of Fast Radio Bursts (FRBs) by Murchison Widefield Array, the foremost precursor telescope for the low-frequency component of the Square Kilometer Array (SKA)

Achievements so far

Detailed profiling of the existing pipeline on GPU-based system identified the need of:

- Removing unnecessary intermediate I/O,
- Developing scalable GPU imaging algorithms on AMD GPUs with **HIP**.

Ultimate goal: Generating 100 high-res images of the radio sky per each second of data observed, as fast as possible (real-time)

ICRAR / Curtin University

Marcin Sokolowski, Ramesh Bhat, Danny Price, Sammy McSweeney, Susmita Sett

Pawsey Supercomputing Research Centre

Cristian Di Pietrantonio

Future

pawsey

Square Kilometre Array

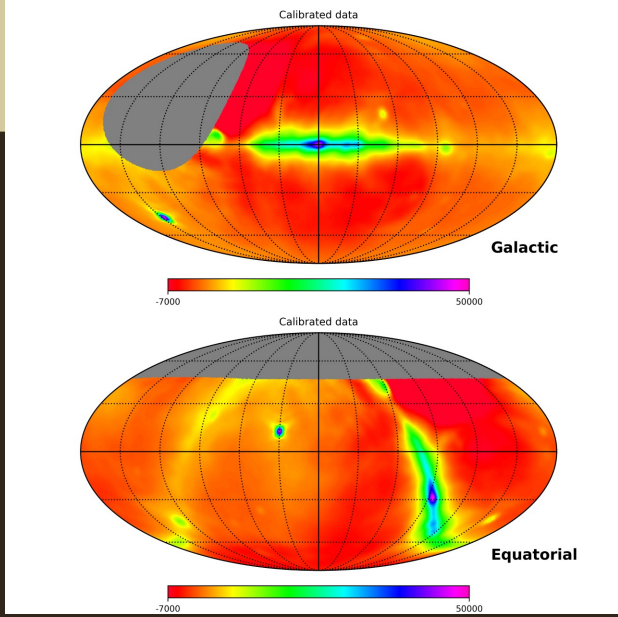

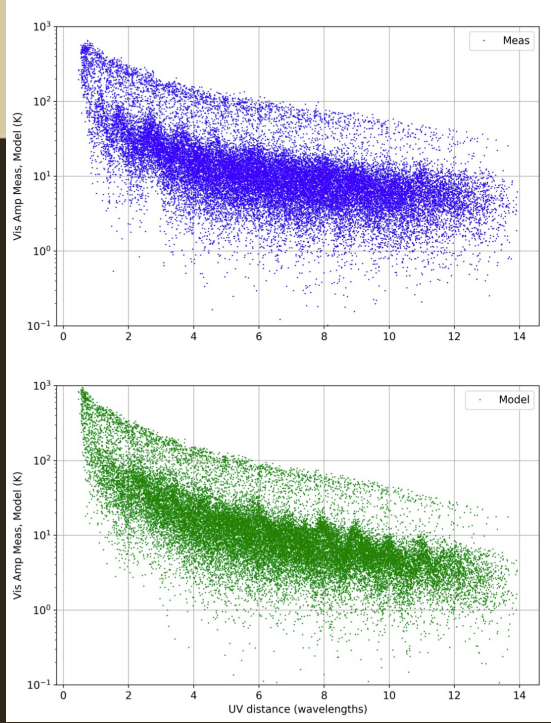

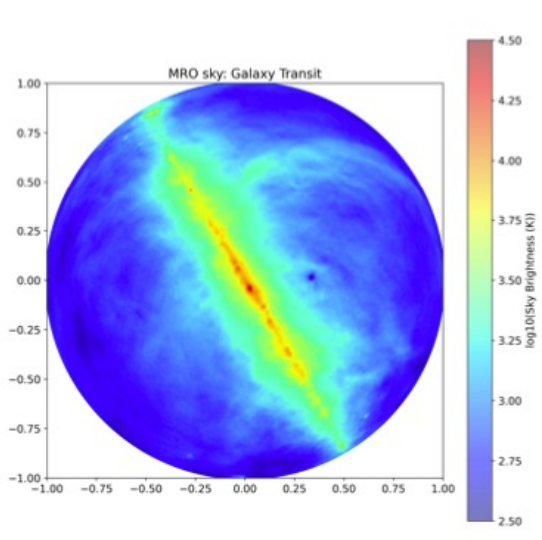

Images generated from first SKA-low station @MRO
Credit: Ravi Subrahmanyan (CSIRO)

Correlator at Pawsey → > 1Pb/s bandwidth

Post-processed observation (data cube, all frequencies, double precision): ~1300TB (x650 increase)!

Quantum

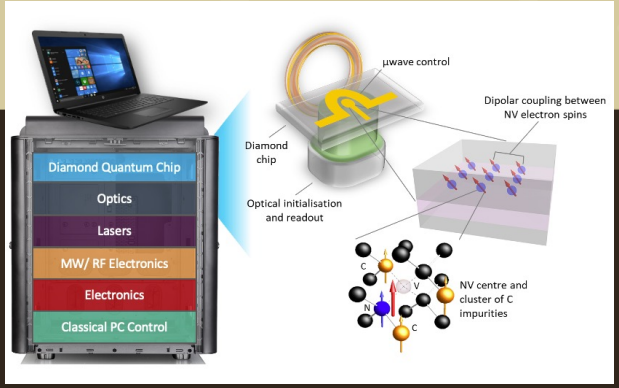
µwave control
Dipolar coupling between NV electron spins
Diamond chip
Optical initialisation and readout
Diamond Quantum Chip
Optics
Lasers
MW/ RF Electronics
Electronics
Classical PC Control
NV centre and cluster of C impurities


QUANTUM BRILLIANCE

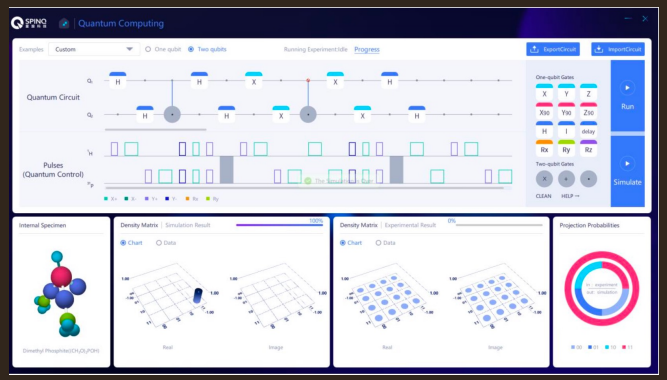
Quantum Computing
Quantum Circuit
Pulses (Quantum Control)
Run
Simulate


pawsey
Thank you!
Aperture Array Verification System 1