PaCER: Accelerating Science on Setonix

Maciej Cytowski, Ann Backhaus, Joseph Schoonover

Pawsey Supercomputing Research Centre

An Australian National Tier-1 Facility

- Launched as Pawsey in 2014, UJV foundations since 2000
- Recent AU$70m capital refresh by Australian Government
- 60+ Staff employed via CSIRO, Australia's national science agency
- Home to:
  - 227+ research projects
  - 4,000+ researchers
  - 10,000+ training attendees
  - 1,000,000,000+ hours of research computing

CSIRO | Curtin University | ECU Edith Cowan University | Murdoch University | The University of Western Australia | NCRIS National Research Infrastructure for Australia — An Australian Government Initiative | Government of Western Australia

Setonix Supercomputer

- HPE Cray EX System named after the iconic Quokka
- 30x increase in compute power and more emphasis on accelerators with future-generation AMD EPYC™ CPUs and AMD Instinct™ GPUs

#1

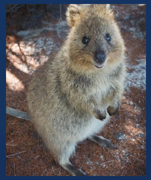

#4 The GREEN 500

#17 TOP 500

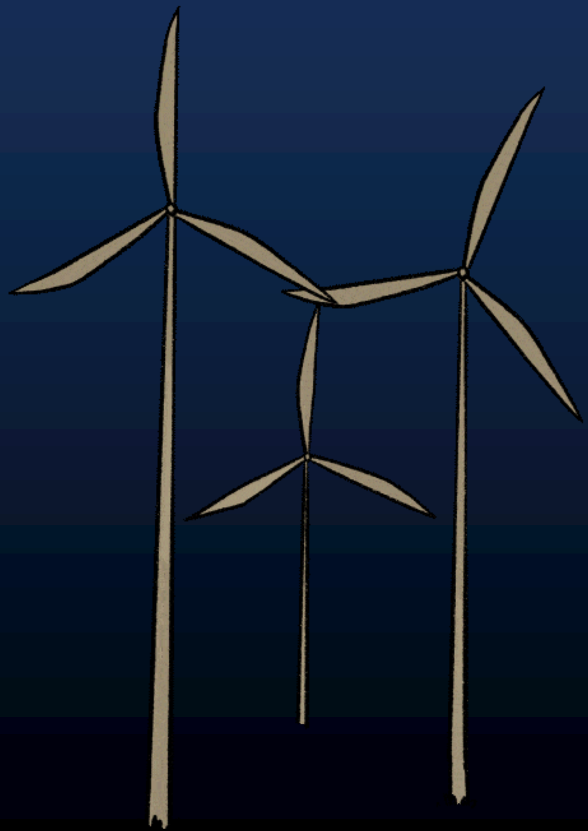

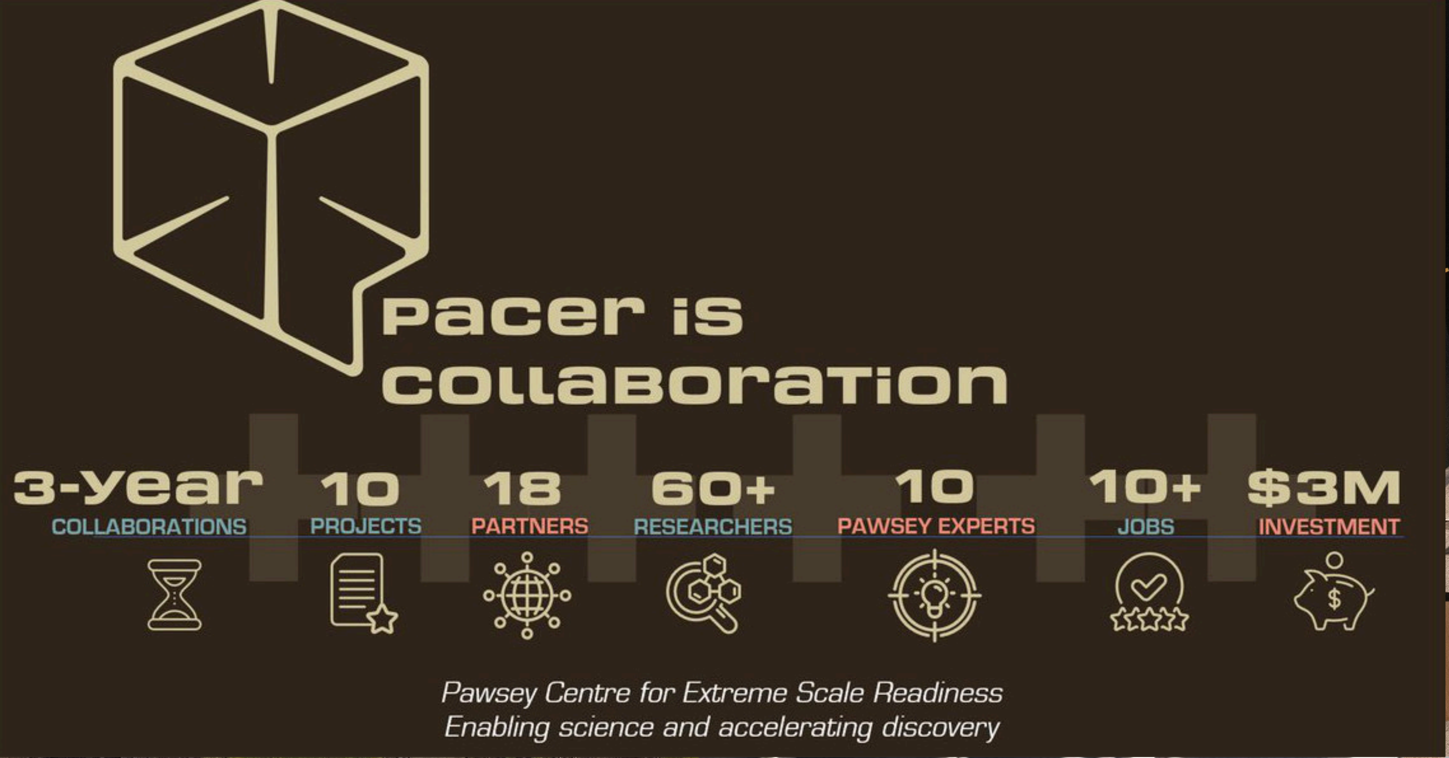
Pacer is collaboration
3-year
COLLABORATIONS
10
PROJECTS
18
PARTNERS
60+
RESEARCHERS
10
PAWSEY EXPERTS
10+
JOBS
$3M
INVESTMENT
Pawsey Centre for Extreme Scale Readiness
Enabling science and accelerating discovery



PaCER Projects

Visit https://pawsey.org.au/pacer for more details!

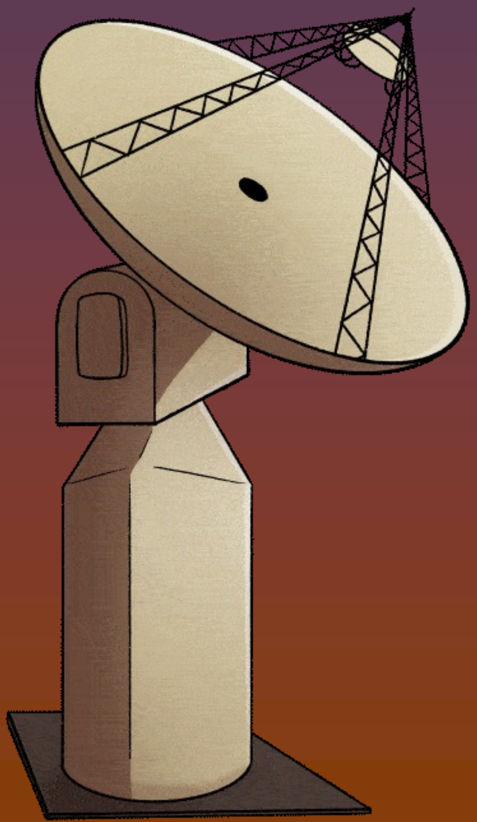

Project	Chief Investigator	Science Domain	Partners
EXESS	Dr Giuseppe Barca	quantum chemistry	ANU, Argonne, Ames, Monash, Flinders, Deakin
VISCOUS	Professor Debra Bernhardt	statistical mechanics, rheology	University of Queensland
MCCC	Professor Igor Bray	atomic and molecular physics	Curtin University
PIGI	Professor Melanie Johnston-Hollitt	radio astronomy	Curtin University, University of Toronto
EmPRiSM	Dr Waseem Kamleh	nuclear physics	University of Adelaide
MaPMoPS	Dr Christopher Leonardi	computational fluid dynamics, geoscience	University of Queensland
HiVIS	Dr Martin Meyer	radio astronomy, cosmology	ICRAR, UWA, CSIRO, SKA, ORNL, AusSRC
GTx	Professor Richard Sandberg, Professor Evatt Hawkes	computational fluid dynamics, turbulence	UMelb, UNSW, General Electric
EXA-GAMBIT	Professor Martin White	particle physics	UQ, Monash University, University of Adelaide
BLINK	Dr Marcin Sokolowski	radio astronomy	ICRAR, Curtin University, AusSRC

Collaboration model

- 2-3 year collaboration
- 0.2 Pawsey's FTE per project
- Co-funding for PhD/Postdoctoral positions (50/50)
- Early access to the system and testbeds
- Allocation for code development purposes
- Allocation for Grand Challenge Problems
- PCon – Pacer's Conference ('21,'22,'23)
- Trainings and GPU mentored sprints
- SC participation

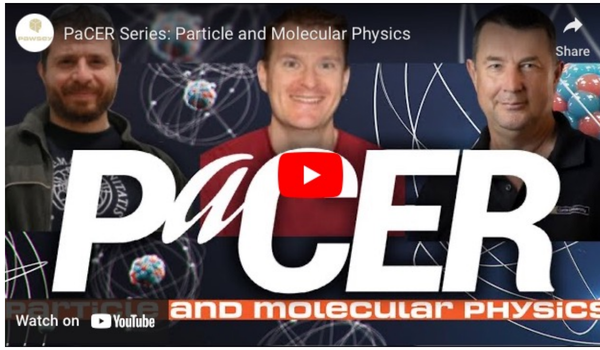

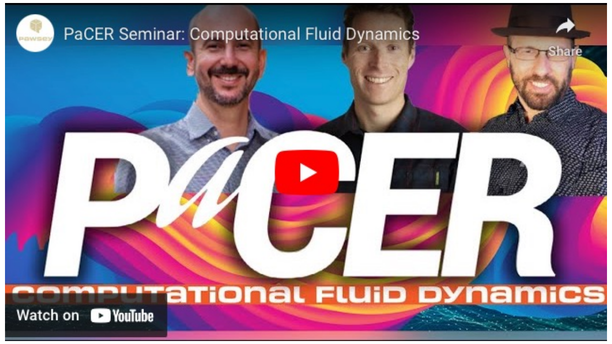

PaCon 2023

- 3-day events gathering all PaCER project technical staff, PhDs, postdocs
- 2021 – online, 2022 and 2023 in-person
- 40-60 participants
- Collaboration with HPE, AMD, Codee, Fluid Numerics
- Invited talks and guests: ORNL, LLNL, HPE Cray, ECP, LUMI
- 2023 event including mentored hackathon

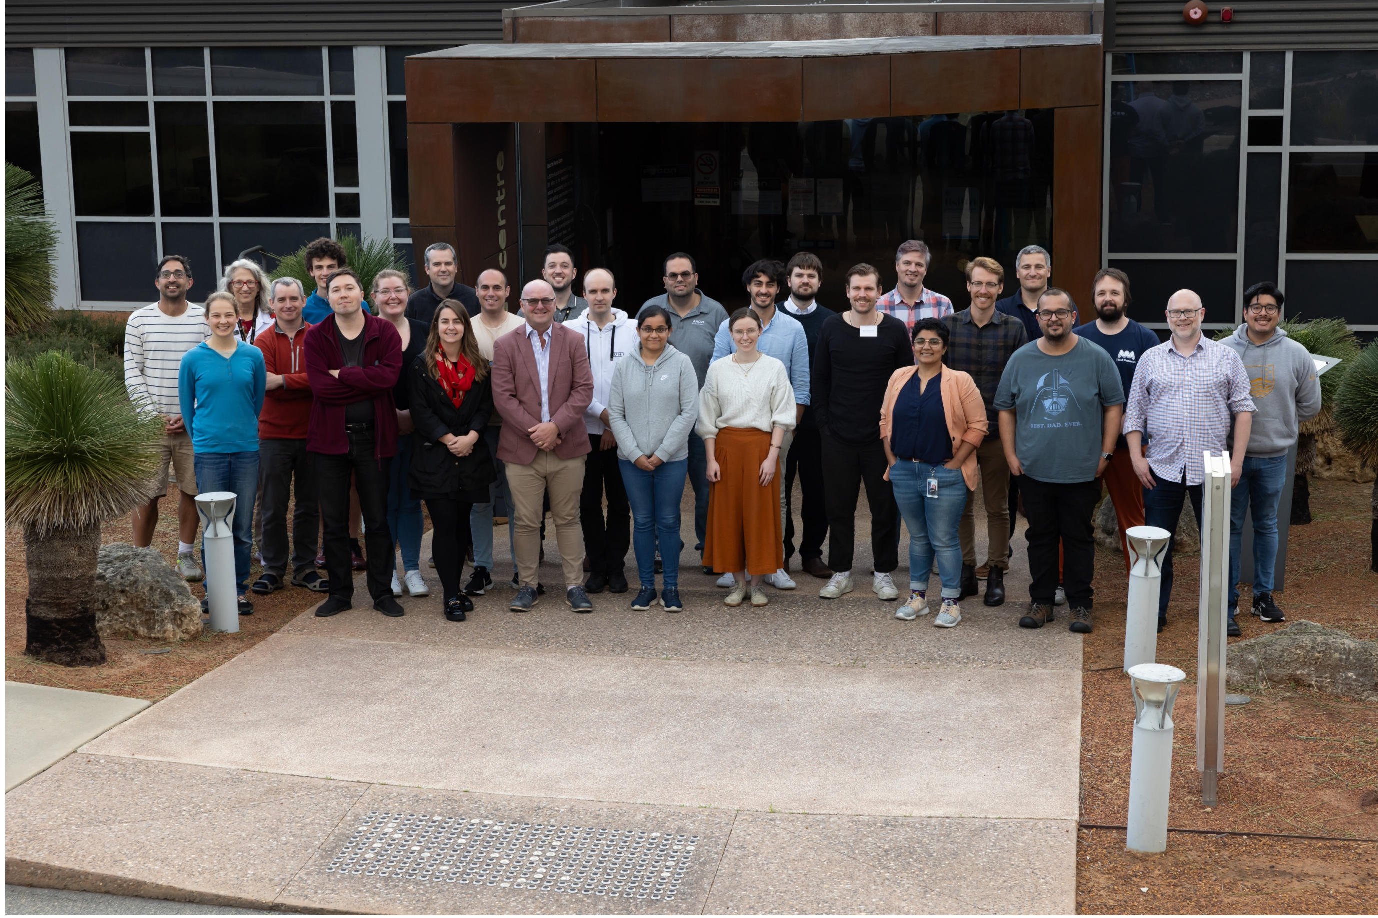

Code Change - Step 1 - 4

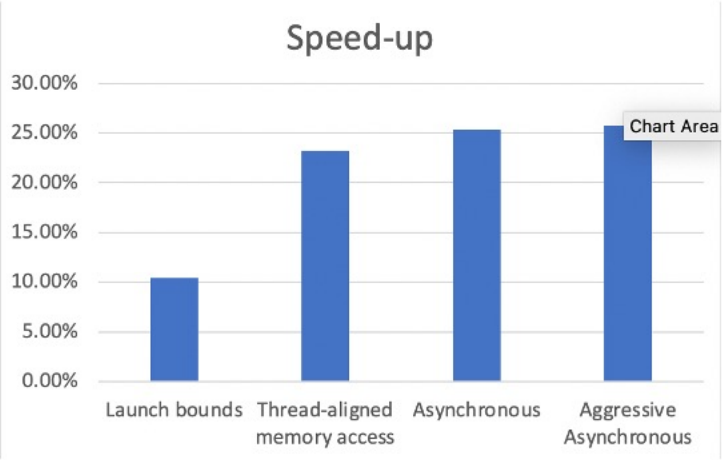

GPU Mentored Sprints

- Pawsey partnered with Fluid Numerics (https://www.fluidnumerics.com) to deliver GPU mentored sprints for 10 PaCER teams

Schedule

Time Period	Activity
-3 Week	Access confirmed to shared compute resources and software team's software. Initial team meeting. Establish goals/milestones for sprint.
-2 Week	Establish software build, testing, and develop (branching) protocols. Develop profile and call graph for existing application.
-1 Week	Draft & confirm sprint plan. Final meetings before the sprint begins.
Sprint Week	Daily Sync up meetings and active coding to execute the sprint plan
+1 Week	Develop retrospective : Incorporate sprint planning documentation with outcomes, hurdles, lessons learned, and next steps.

MCCC

Mentored Sprint Report

Main Challenges

- GPU porting and performance
- Portability and performance portability
- Scalability
- Data challenges: data reduction and in-memory processing
- New programming tools, paradigms, libraries: Kokkos, Julia, SLATE, Codee, AMD Profilers
- Directives: OpenMP vs OpenACC vs HIP
- GPU: reducing the register pressure
- Every project is different and moves at its own pace

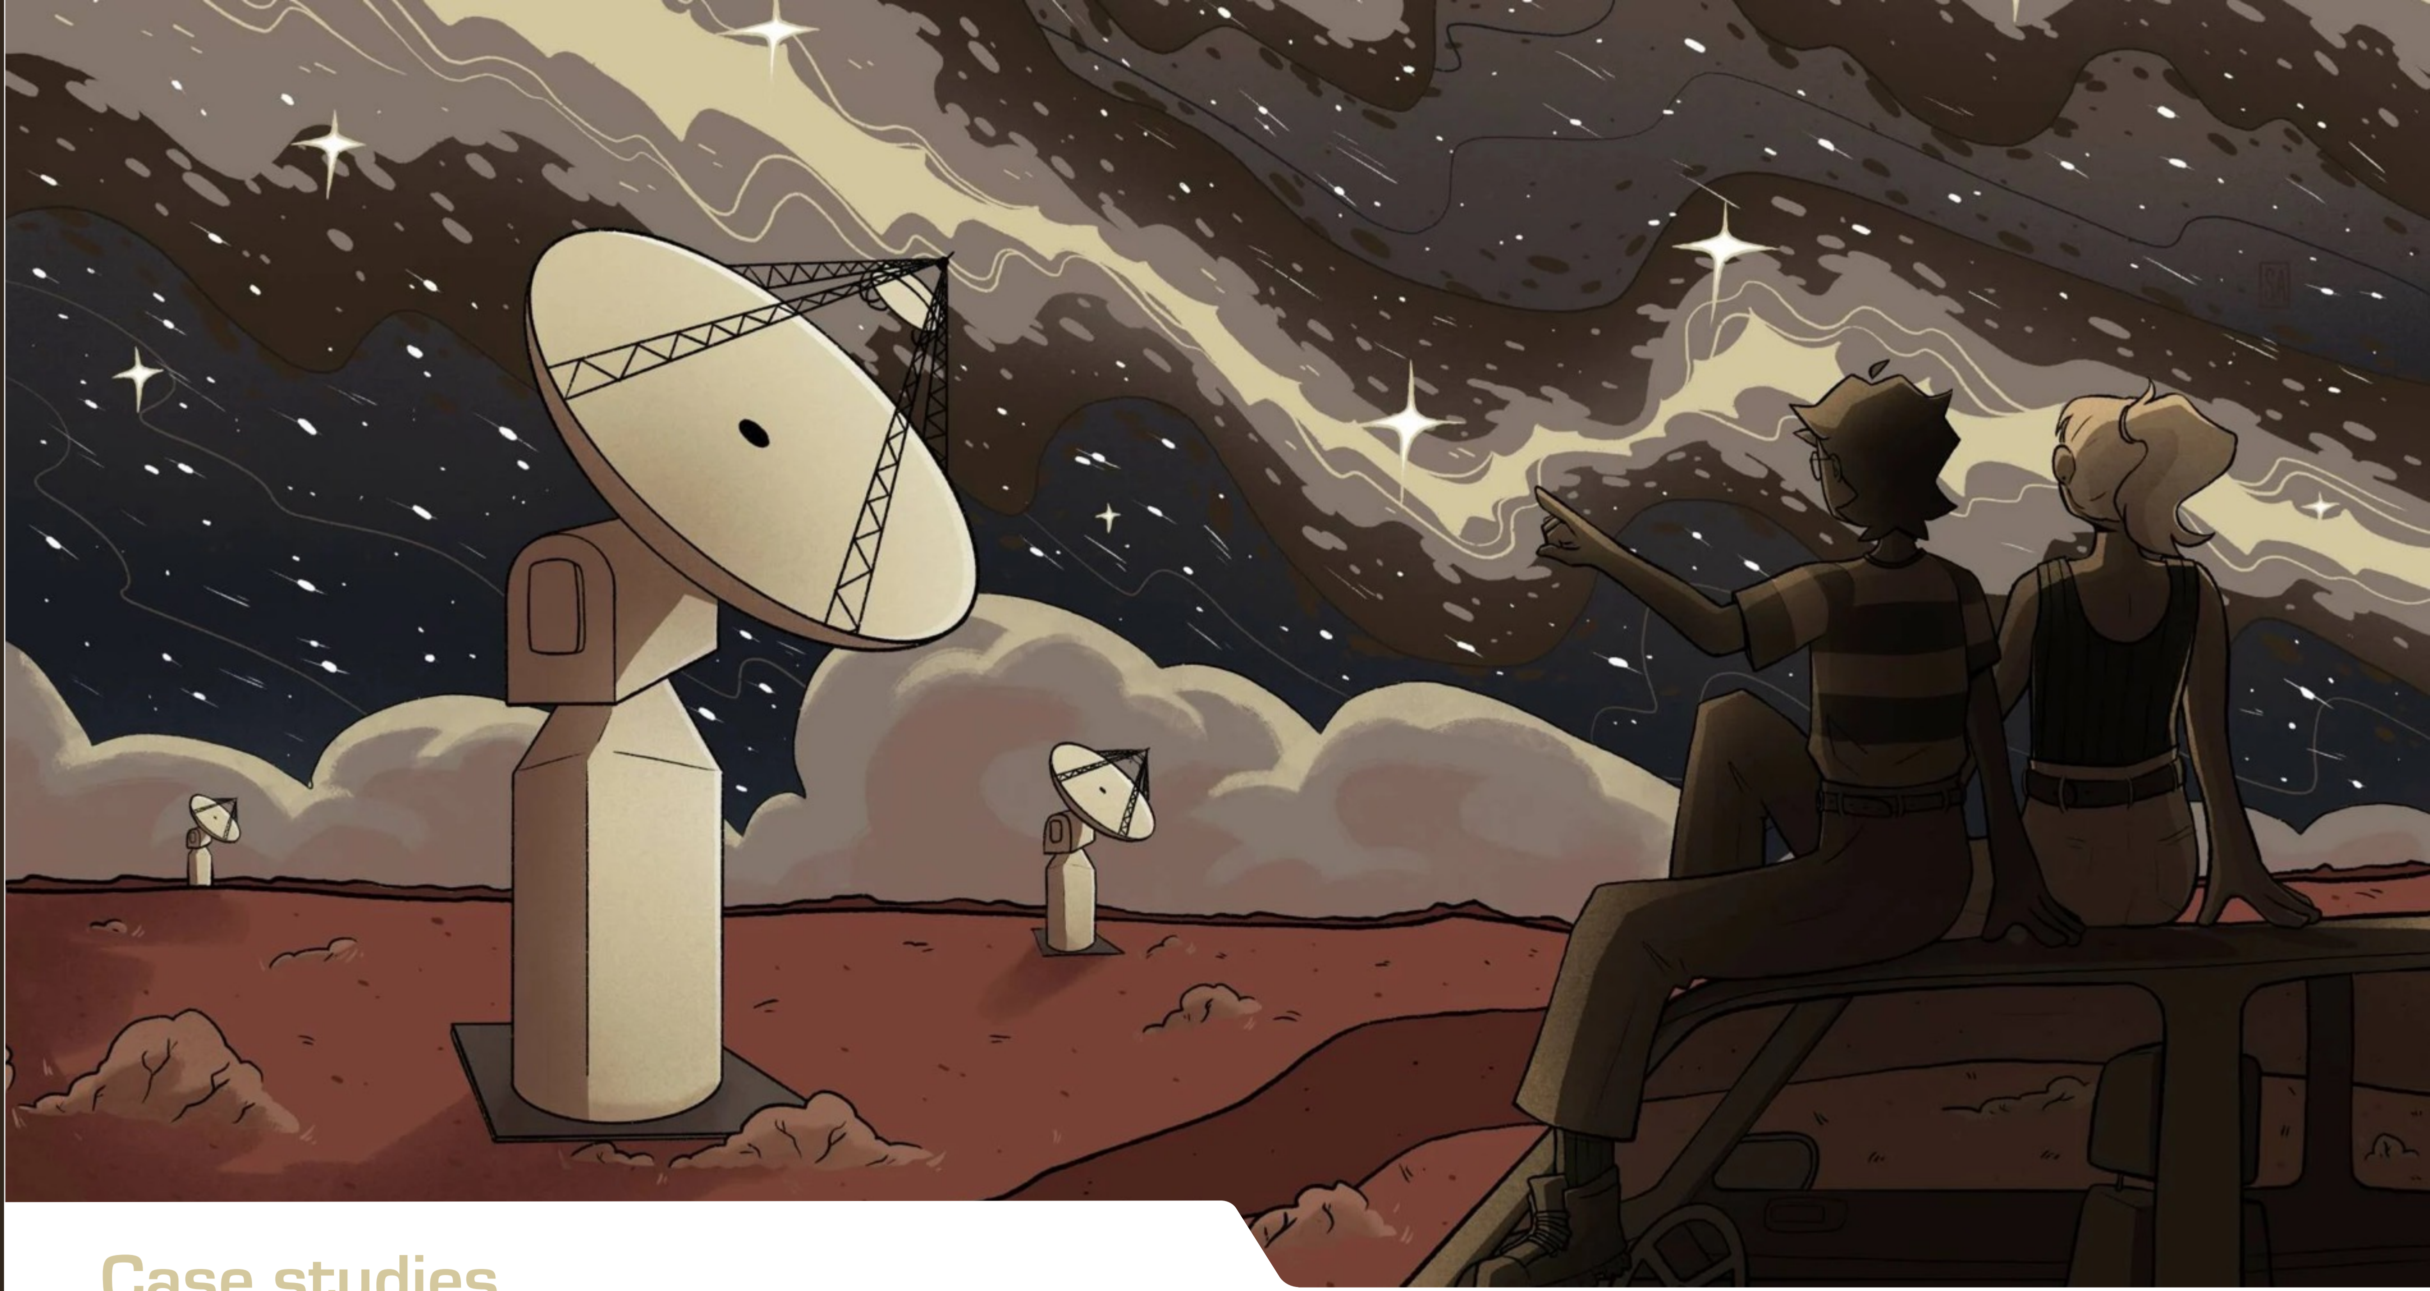

Case studies

Case study: GTx HIPSTAR

Prof. Richard Sandberg - University of Melbourne

- Running the Taylor-Green vortex case (triply-periodic) with 512^3 grid points takes about 0.1886 s per timestep on one GPU node compared to 4.8 s on a CPU node
- The code is >**25**× faster and >**6**× more energy efficient on a GPU node

Fixed total problem size:

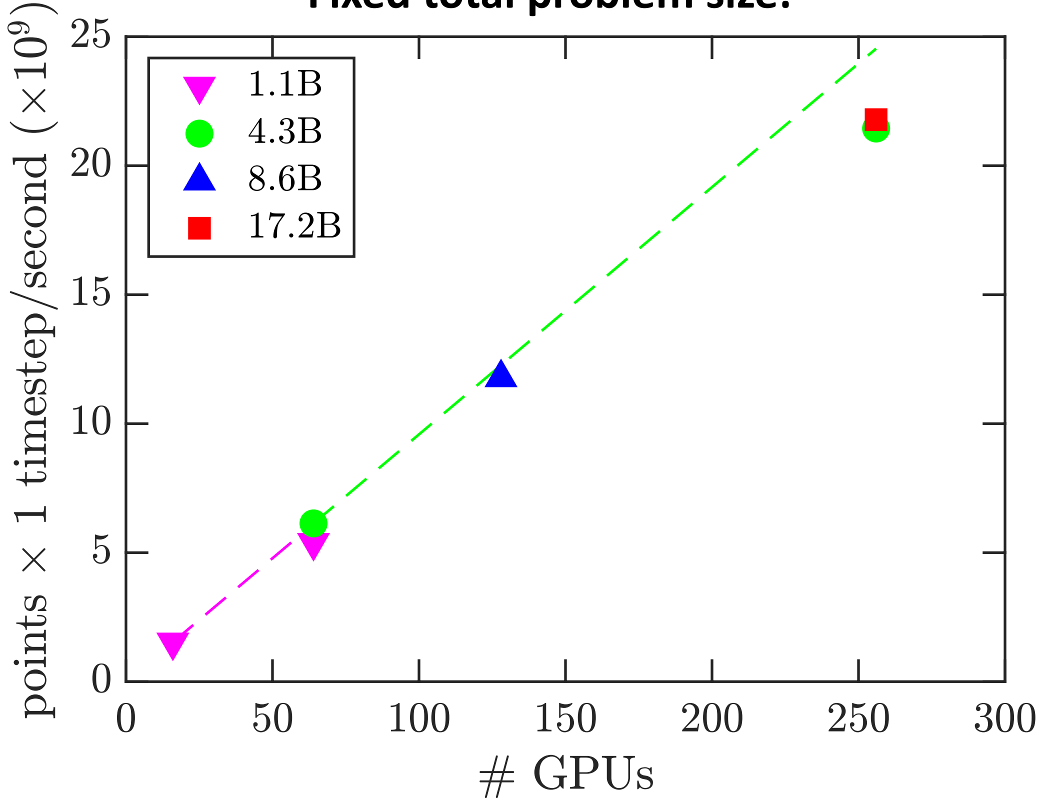

Fixed # grid points/GPU:

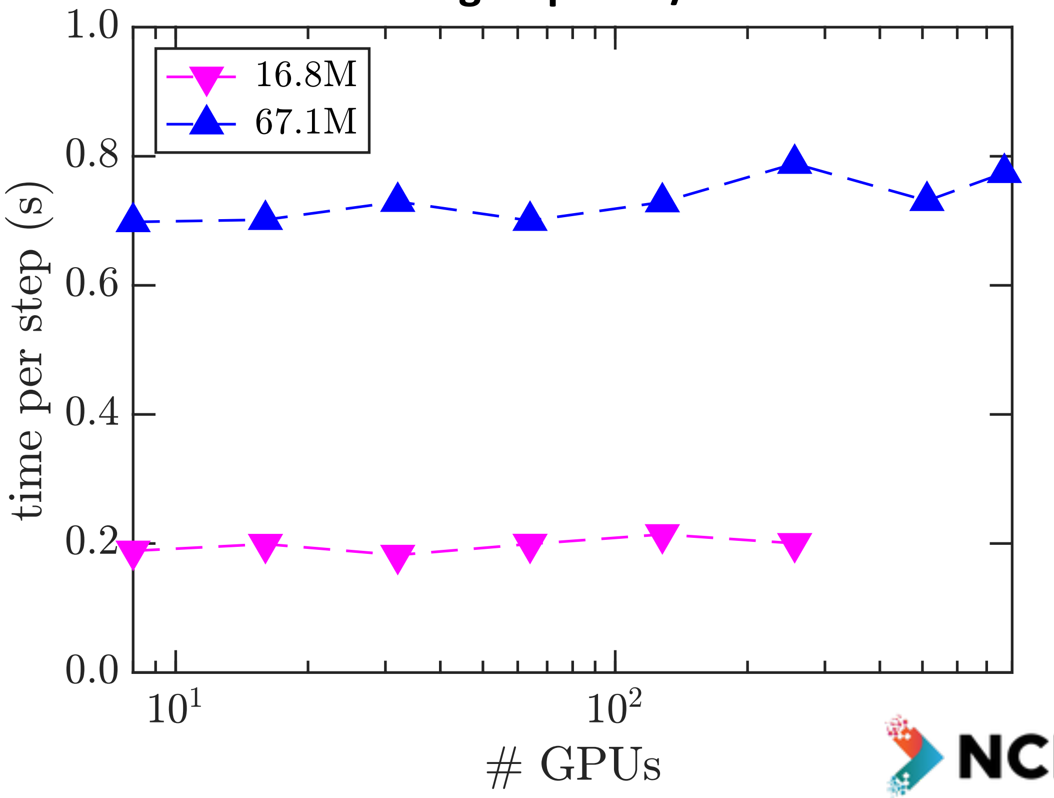

NCMAS

Case study: EmPRISM
Dr Waseem Kamleh - University of Adelaide

Emergent Phenomena Revealed In Subatomic Matter (QCD)

- How does the **foundation of matter** evolve to regimes relevant to the interior of neutron stars and the early universe?
- Can the **centre-vortex structure of QCD ground-state fields** account for the confinement of quarks and the dynamical generation of mass?
- How does the **spectrum of hadron resonances** emerge, and can a precision understanding guide the search for new physics?

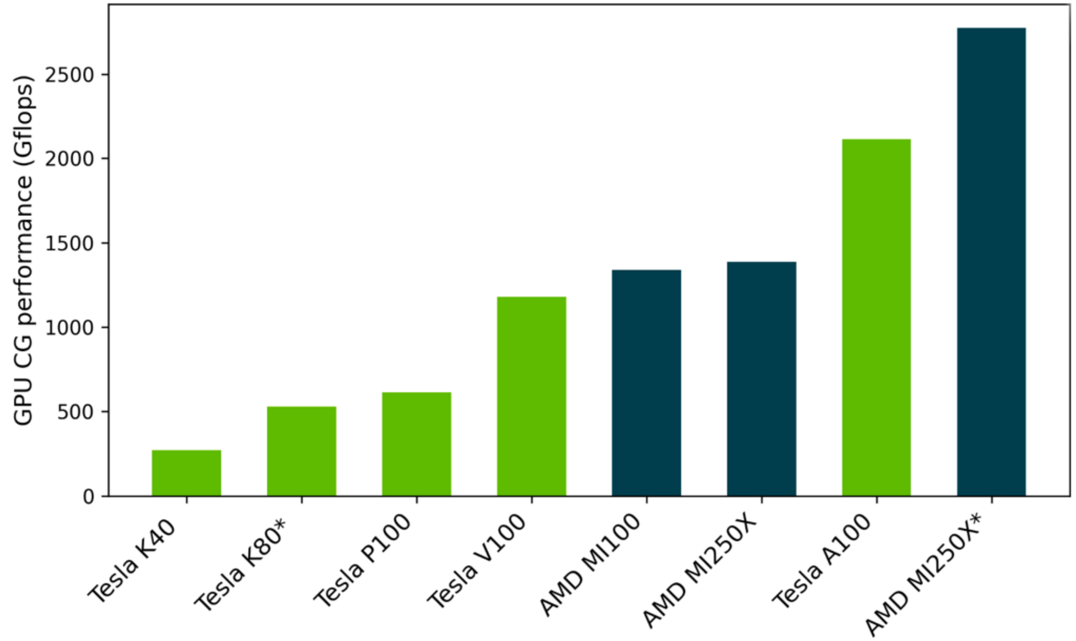

- In-house code optimised for multiple architectures
- Fortran ’08, multi-GPU
- Achieved platform portability from CUDA to HIP
- Explored different data layouts and code factoring
- Performance portability with minor code changes
- GPU acceleration critical to addres science challenges

NCMAS

Case study: EXESS

Dr Giuseppe Barca - Australian National University

The EXtreme-scale Electronic Structure System, predicting the chemistry of nanomaterial interfaces

- framework for a continuing community-wide effort to develop next-generation molecular modelling capabilities that support a broad spectrum of chemistry, biology and material science research on computing systems ranging from petascale to exascale and beyond

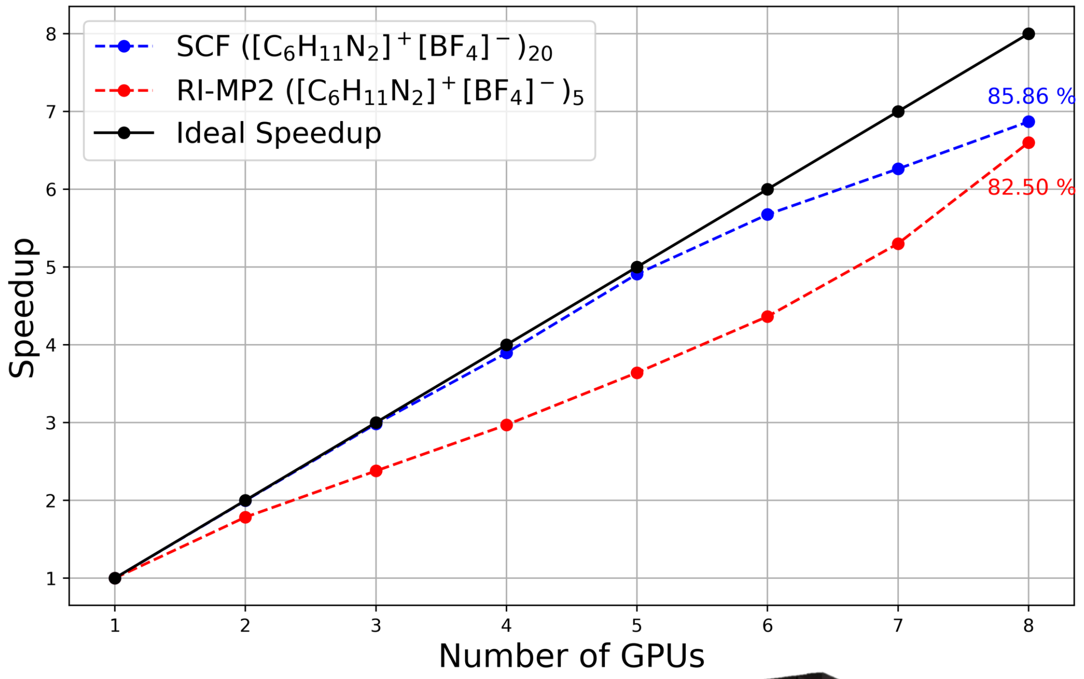

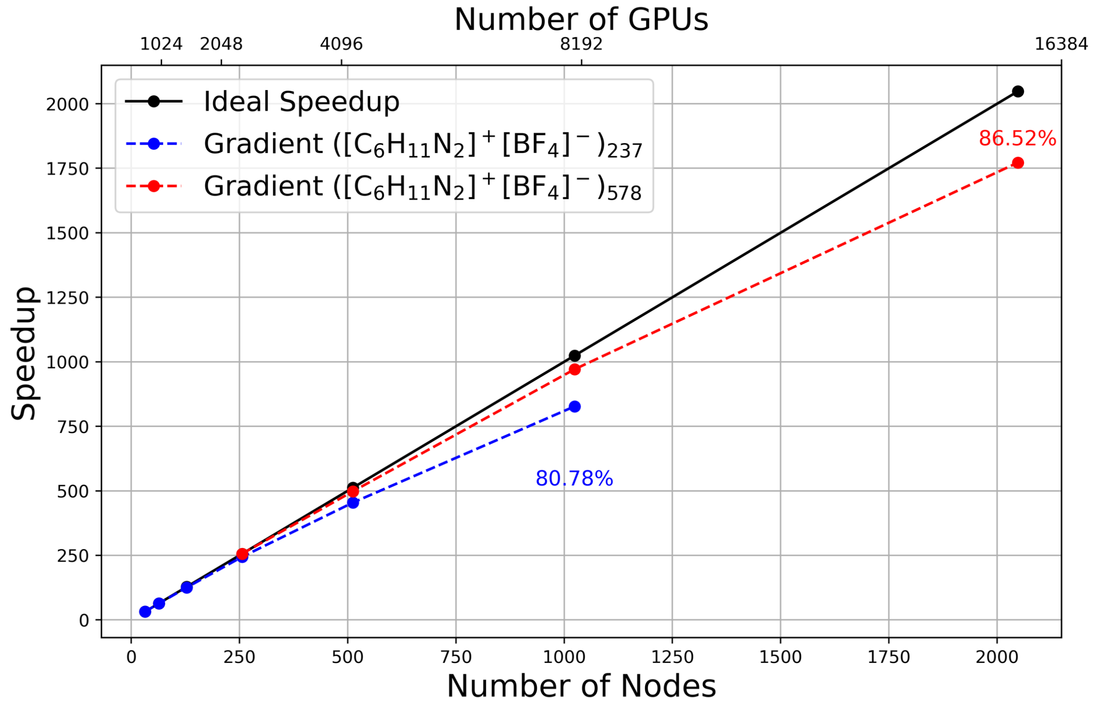

Case study: HiVIS

Dr Martin Meyer – University of Western Australia

- Project aims to add much needed IO functionality to data products produced by the ASKAP radio telescope and other radio telescopes

- The project has potential to move SKA-low from EB-scale project to more manageable 100 PB scale project

- Adaptive I/O System Version 2 (ADIOS2) library and new data reduction technologies with Multi-Grid Adaptive **Reduction of Data** (MGARD)

- Data processing performance improvement and ~10x reduction without major impact on science results

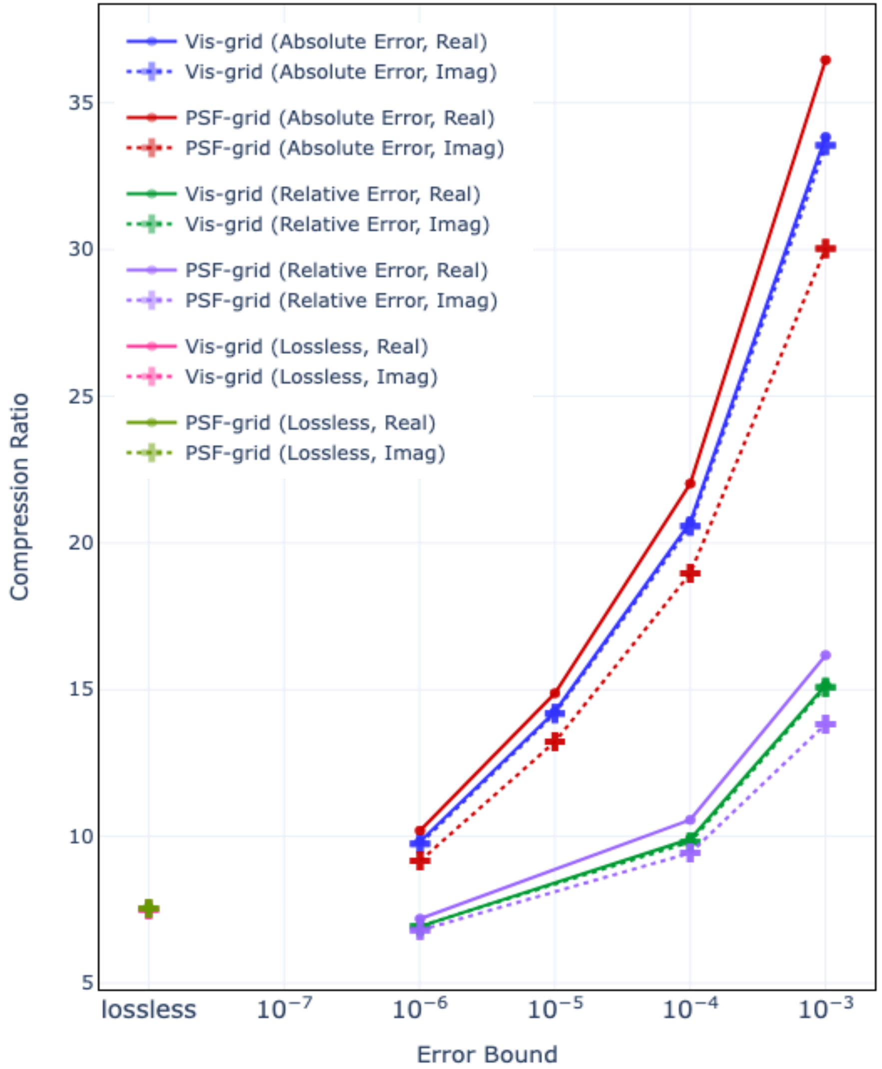

CUG24 presentation “Optimising the Processing and Storage of Radio Astronomy Data” by Alex Williamson

Case study: MCCC

Prof. Igor Bray – Curtin University

- MCCC project is studying collisions between tiny particles which helps us understand everything from fusion energy to cancer treatment
- Implementation of Convergent Close-Coupling method for studying ion-atom collisions in Fortran, MPI and OpenMP
- Simulating such collisions is very time consuming
- Successfully ported to run on GPUs (Graphics Processing Units)
- Portability across AMD and NVIDIA GPUs
- Software deployed on other centres in Australia as well as through ACCESS in US

Portable GPU implementation of the WP-CCC ion-atom collisions code, I. B. Abdurakhmanov, N. W. Antonio, M. Cytowski, A. S. Kadyrov, https://doi.org/10.48550/arXiv.2403.04252 (to be presented at ISC24 with the The HPC on Heterogeneous Hardware H3 Workshop)

Case study: BLINK

Dr Marcin Sokołowski – Curtin University

Goal: enable real-time processing of MWA data to detect Fast Radio Bursts (FRB) candidates

Initially:

- several legacy codes forming the pipeline
- every step reading and writing data to disk
- different data formats, converters in use

Currently:

- 3/6 steps implemented as GPU kernels (portable across NVIDIA and AMD GPUs)
- all data kept in-memory
- extreme speedup of selected components of the pipeline
- IOPs friendly

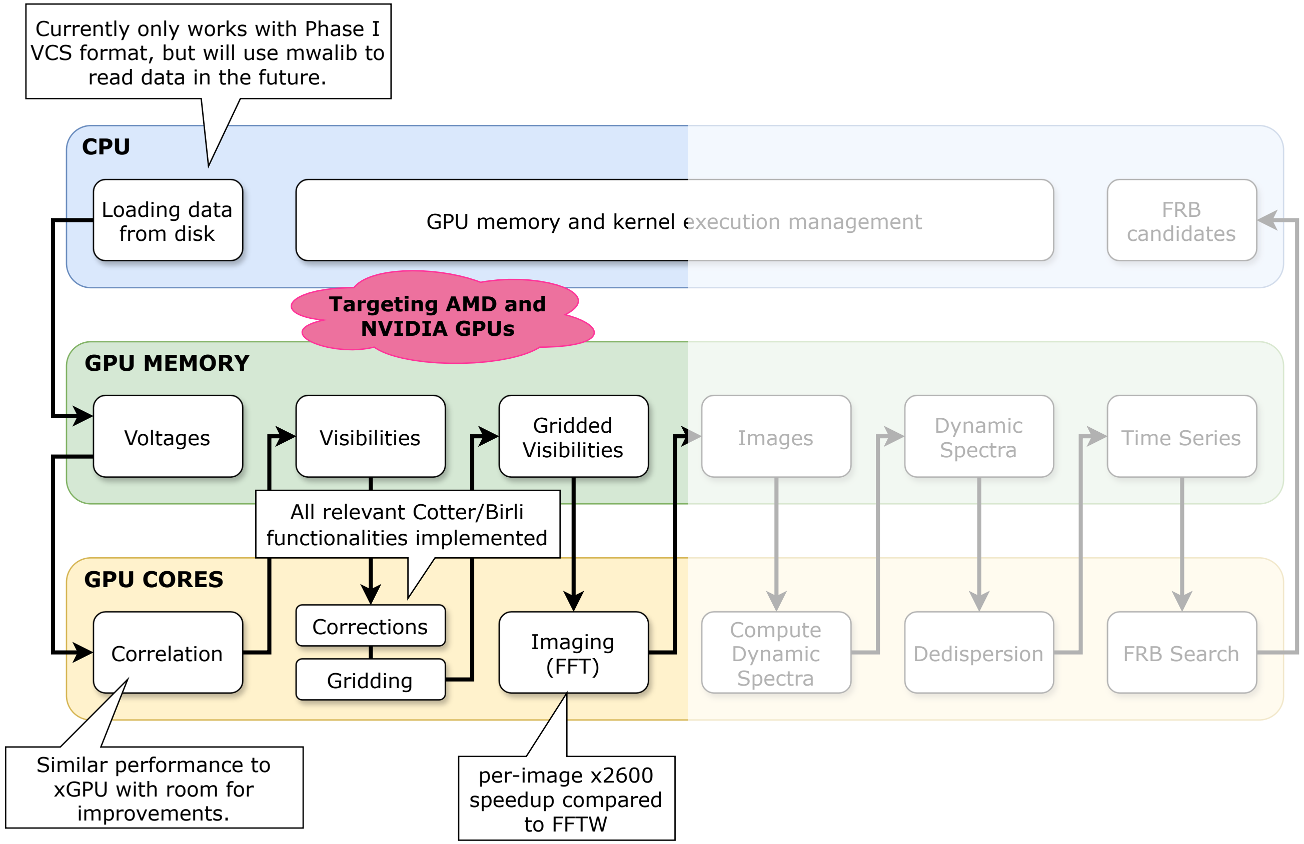

"*High-Time Resolution GPU Imager for FRB searches at low radio frequencies*", M. Sokolowski, G. Aniruddha, C. Di Pietrantonio, C. Harris, D. C. Price, S. McSweeney, R. B. Wayth, and N. D. R. Bhat, to be published in Publications of the Astronomical Society of Australia

Conclusions

- We collaborated with the biggest computational and data science projects in Australia to optimize their workflows on Exascale-class supercomputer Setonix
- We accelerated science, improved energy efficiency and portability of codes
- We enabled research previously not possible
- We built cases and requirements for next upgrades
- We created unique communities, collaboration and careers
- Software readiness programs should always be part of new system deployments
- Collaboration with vendors is crucial

Any questions?

Thank you