PAWSEY
supercomputing centre

Pawsey Supercomputing Centre
Annual Report 2018-19

Growing Australia's Scientific Ambitions

by expanding our HPC infrastructure

TABLE OF CONTENTS

MAKING TOMORROW HAPPEN,TODAY

AMAZING OUTCOMES

VOICES OF SCIENCE

A NEW ERA OF RESEARCH

A WORLD OF DIFFERENCE

FINANCIALS


Making Tomorrow Happen, Today
A national research infrastructure and global player in science and innovation

CHAIRMAN FOREWORD

It has been a year of renewal, rearrangement and refocus as we continue our journey to remain at the forefront of HPC for Australia's benefit.

WELCOME TO THE PAWSEY SUPERCOMPUTING CENTRE 2018–19 ANNUAL REPORT

This year has seen our Executive Director, Mark Stickells, establish Pawsey's future course with a stronger commitment to engaging with our stakeholders and partners.

Our relationship with the National Computational Infrastructure (NCI), as Australia's two Tier 1 High Performance Computing (HPC) facilities, has strengthened significantly. We held our first Board-to-Board meeting in June, and are working to align our governance processes. We are also collaborating in our engagement with the Commonwealth Government, to ensure that Australia's national interests are served through the provision of globally-competitive HPC into the future. We are working together across a range of areas, ensuring that between the two facilities we support the numerous research fields requiring supercomputing across Australia.

With the NCI we are planning for future aligned and secure services. We are also working together to seed the development of skills in HPC across Australia, as our need for a HPC-skilled workforce will only grow in the future.

Engagement with our stakeholders and partners has also intensified as we have embarked on our $70 million Capital Refresh project. Exploring our users' current and future HPC requirements and identifying their key needs has been crucial, and their input has been central to designing the technical specifications for the infrastructure upgrades. Our stakeholders have also informed our planning for the procurement and installation process to minimise disruption to services as the upgrades occur. We are all excited by the enhancement to our facilities and capabilities that the project will provide, and look forward to extending our expanded services to all of our users in the future.

During this year our senior management team has settled around Mark, with more certainty about key appointments and organisational structure. The team is now entirely focused on improving relationships with Pawsey's external community moving forward. In particular, we're working hard to support the activities that are occurring in astronomy, particularly through the Square Kilometre Array (SKA) program. Highlights from those activities can be found throughout this report.

Moving forward, our connections with industry are slowly developing and we are using Mark's background and expertise to engage more effectively with industry. Building on our activities which can lead more directly to innovation and economic growth remains a priority.

I would like to recognise and thank the dedicated Pawsey staff for their effort and commitment over the year to providing Tier 1 HPC and accelerating Australian research across many domains. I would also like to thank the Board Directors for their focus and direction, and particularly acknowledge Dr Campbell Thomson from the University of Western Australia for the contributions he has made over several years, as he steps down to be succeeded by Dr Merran Smith.

It has been a year of renewal, rearrangement and refocus as we continue our journey to remain at the forefront of HPC for Australia's benefit.

John Langoulant AO
Chairman of the Board

EXECUTIVE DIRECTOR'S REPORT

It's been one year since I joined the Pawsey Supercomputing Centre, which makes it easy to reflect on the year's activities.

While exploring the community that exists around Pawsey – the partners, collaborators, research organisations and researchers – I've been working to build on and expand these connections. The nature of science and global research means collaborative partnerships in HPC remain vital to our future growth.

At the international level, Pawsey had a significant presence at the major supercomputing conferences SC18 in Dallas, US; and at ISC High Performance 2019 in Frankfurt, Germany. We also co-convened and presented at SCAsia 19 in Singapore.

Nationally, we are continuing to build a stronger relationship with the NCI, our partner Tier 1 HPC centre in Australia. Pawsey and NCI jointly exhibited at the three international conferences above, showcasing Australian HPC capability. The Centres also partnered on Data Science Week 2019 activities and held several joint training programs. Knowledge sharing is an important part of our relationship, and Pawsey has supported the NCI's recent capital refresh program with technical advice. The NCI in turn is now assisting with the initial stages of the Pawsey Capital Refresh.

Pawsey's Capital Refresh program has obviously been a major focus for the year, and we acknowledge the Commonwealth Government's support in investing in both the NCI and Pawsey technical infrastructure to ensure we can continue to meet the growing demands for HPC-powered science from data across Australia. The refresh is a multi-year effort, and this first year was largely about planning – establishing the governance and organisation for the process, engaging and consulting with users to understand their current and future requirements, and establishing the technical specifications for each part of the refresh.
In the coming year the major HPC systems will go to market, and the program is on track to be completed and fully operating in 2022.

Pawsey continues to support Australia's leading radio astronomy projects, which achieved some significant progress milestones this year. A priority for me was to visit the Office of the SKA Organisation in Manchester UK last October, prior to Australia becoming a signatory to an international treaty establishing the Square Kilometre Array Observatory, the intergovernmental organisation that will oversee the delivery and operation of the world's largest radio telescope.

Pawsey, with CSIRO and the International Centre for Radio Astronomy Research (ICRAR), has also contributed to the SKA's Science Data Processor consortium, which in May concluded the engineering design work for the supercomputers that will eventually process the enormous amounts of data produced by the SKA. One of these supercomputers will be located in Perth to process data from SKA-low, and Pawsey will continue to be involved in the planning for the transition to full operations in the coming decade.

Given Pawsey's historical and ongoing involvement with the SKA pathfinder telescopes, the Murchison Widefield Array (MWA) and the Australian Square Kilometre Array Pathfinder (ASKAP), it was a personal highlight of my year to have the opportunity to visit the Murchison Radio-astronomy Observatory last August, and see first-hand the amazing advances that are being made in radio astronomy on the path to the SKA.

Pawsey continues to support research in a multitude of other domains, including helping plan the data requirements of the new Australian Phenomics Centre, hosted by our partner Murdoch University. The Centre is also underpinning the Artemis Project, establishing a remote real-time health analytics platform, and appointed our inaugural Pawsey HPC research fellow Dr Jennifer Yeung in this field.

The Centre for International Economics was commissioned in November to undertake an economic impact assessment of the contribution of the Pawsey Supercomputing Centre. The report was released in March and demonstrates Pawsey's value in contributing to scientific knowledge and Australia's capabilities across many sectors. The case studies analysed in detail also demonstrated how economic value was created. A snapshot of these independent findings appears later in this report, along with other case studies that illustrate the breadth of research being undertaken at Pawsey.

Pawsey is an exciting place to work and its greatest assets are its people. I've enjoyed getting to know the wonderful experts in HPC and data science that I can now call colleagues. I'd like to thank the Pawsey team for their efforts in supporting the 1,744 researchers who undertook 200 projects at Pawsey over the last year. Their research benefits not just our own nation, but in many cases our global community. The community and networks in and around Pawsey are turning us into a truly global research facility, and I'm very proud to continue supporting their efforts.

Mr Mark Stickells
Executive Director

PEOPLE

50 per cent — *Increase in the number of women appointed*

Pawsey has doubled the number of women working at the Centre and almost half of the staff have a diverse cultural background, increasing gender and diversity balance and inclusion.

700 people — *focused on opportunities for women in data science*

More than 700 people participated in Data Science Week 2019 with over twenty events hosted in Perth and Canberra with the support of NCI and The Australasian eResearch Organisations (AeRO).

500 people — *were part of Pawsey's Open Day*

WA Minister for Water; Fisheries; Forestry; Innovation and ICT; Science, joined 500 members of the public at the Pawsey Open Day.

300 people — *met Magnus*

Staff and researchers from other supercomputing centres, government agencies and scientific facilities visited Pawsey, including delegations from Chevron Houston, the Committee for Perth, Liquified Natural Gas (LNG) China, and others.

more than **100** people — *inspired during Pawsey's careers nights*

Aiming to inspire future generations to pursue STEM, Pawsey's first Careers Night attracted 40 participants. A second session was hosted as part of Data Science Week with 68 participants.

3000 people — *were exposed to Pawsey services and scientific breakthroughs*

Pawsey staff engaged with over 3,000 people at conferences, tours and special events over the last year. At SC18 Dallas alone, Pawsey and the NCI engaged with more than 700 attendees to showcase Australia's HPC landscape and scientific achievements.

CAPABILITIES

200 projects — *using HPC, supporting 42 different institutions*

1,744 active users accessed Pawsey HPC services during the reporting period.

694 researchers — *trained in HPC*

On-line and face-to-face sessions were held across the country in over 23 different training sessions.

12 — *summer interns discovered HPC*

Pawsey's intensive 10-week summer internship program enabled 12 undergraduate students to develop skills in computational science to support their research.

13 projects — *accelerated via the Uptake program*

Projects from new and existing user groups accessed up to 45 days of Pawsey expertise each to accelerate their research outcomes.

70 — *bioinformaticians discovered supercomputing*

Pawsey hosted the first Pawsey-COMBINE Bioinformatics Symposium with participants from the bioinformatics community spanning academia, state departments and private organisations.

1000x speedup — *at the international GPU Hackathon*

Pawsey hosted its second GPU Hackathon jointly organised with Oak Ridge National Laboratories and NVIDIA where researchers and industry participants accelerated their work after a week of fast-track learning.

SYSTEMS

26 new projects — *activated on Magnus, the Supercomputer*

In January 2019, 26 new projects were succesful in their competitive application for allocation on Pawsey's supercomputer.

1st milestone — *towards Pawsey's growth was reached*

The governance mechanisms, terms of reference and membership of key User and Technical Reference Groups were established towards the $70 million upgrades to Pawsey's supercomputing infrastructure.

100 Gbps — *link live*

A 100 Gbps link between Pawsey and MWA at Curtin was completed and is now live.

WORKING TOGETHER

National Tier 1 Supercomputing Centres working together

Governance processes, future secure services, staff training and HPC skill development is being aligned between Australia's two Tier 1 HPC facilities.

Launch of the ADACS uptake program

Pawsey's Astronomy Data and Computing Services (ADACS) uptake program is embedding technical staff to support large-scale astronomy projects.

Strengthening engagement across Australasia

Pawsey partnered with the National Supercomputing Centre (NSCC) Singapore to co-organise SCAsia 19 in Singapore.

THE PAST

THE PATH TO A WORLD-CLASS FACILITY

Constructed in 2012, the Pawsey Supercomputing Centre, named in honour of the Australian radio astronomer, Dr Joseph Pawsey, is a federally funded purpose-built facility in Perth, Western Australia, It is a national HPC facility, delivering cutting edge computational support and services to key scientific areas such as radio astronomy, bioinformatics, energy and resources.

The Centre's compute capacity has grown by an order of magnitude since its inception, and it now provides world-class expertise and infrastructure in supercomputing, cloud services, and data-intensive analysis, storage and visualisation to over 1,500 researchers every year.

COLLABORATION

Pawsey is an unincorporated joint venture between CSIRO, Curtin University, Murdoch University, Edith Cowan University and The University of Western Australia. The Pawsey Supercomputing Centre is also funded by the State and Federal governments. One of two Tier 1 HPC facilities within Australia's national research infrastructure network, Pawsey partners with the NCI in Canberra to ensure researchers across Australia have access to world-class computing facilities.

ABOUT PAWSEY

THE PRESENT

REAL WORLD IMPACT

Through supporting Australian researchers and their collaborators around the globe, Pawsey enables high-impact research for the benefit of society across domains including energy and resources, engineering, bioinformatics, health sciences and agriculture. Pawsey-enabled research is building better engines, extracting energy from ocean waves, caring for premature babies, combating dementia, and safeguarding our food and water supplies.

WORLD-SCALE PROJECTS

Pawsey is the only HPC facility in the world to directly ingest and process data from world-leading scientific research instruments in real time. Pawsey supports the operation of two SKA precursors, the MWA and ASKAP, and has been processing and storing their data since they began operations. These radio telescopes are making discoveries about the nature of our Universe while contributing to the design and development of the SKA, the world's largest mega-science project. Pawsey is equally contributing to the design, development and testing of the data processing and storage components of the SKA.

STAYING AT THE FOREFRONT OF HPC

The Federal Government has committed funding to support a regular upgrade cycle enhancing capability across Australia's two Tier 1 HPC facilities ensuring there is always at least one facility operating at full capacity in a field where technical obsolescence is rapid. These upgrades will keep pace with the steadily increasing volume of complex computing demands arising across government, research and industry.

Pawsey is currently in the middle of a $70 million Capital Refresh project to secure its next generation of supercomputers, data and supporting infrastructure. This is keeping Pawsey at the forefront of global advances in supercomputing technology, and will provide significantly increased computing power and speed, expanded data storage capacity, and more efficient high-performance computing services for Pawsey users. The upgrade is being staged over time to ensure continuity of Pawsey services for the national research community. Installation and performance benchmarking of the new HPC infrastructure are anticipated to be complete by early 2022.

THE FUTURE

CUTTING-EDGE SCIENCE

Pawsey's petascale computing facilities will continue to provide essential tools for researchers to solve the problems of the future, faster. Through technology, expertise, and collaboration, the Centre upscales the ambitions, outcomes and impacts of Australian research.

DRIVING INNOVATION AND UPSKILLING AUSTRALIANS

Pawsey plays a pivotal role in Australia's research and innovation landscape, and through its outreach and international presence is positioning Western Australia as a global hub for supercomputing expertise. By facilitating access and use of HPC services through high quality training and support programs, the Centre is helping upskill researchers, students and industry personnel while developing the next generation of HPC and data scientists for Australia.

A CONNECTED FUTURE

The Centre is forging collaborative relationships with other supercomputing centres around the world to exchange knowledge and further develop the global impact of supercomputing. By driving innovation and accelerating discoveries across a multitude of different fields, Pawsey is enabling Australian researchers and their international collaborators to tackle the world's 'wicked' problems.


Enabling life-changing discoveries
Amazing Outcomes

PAWSEY IMPACT STUDY

The Pawsey Supercomputing Centre commissioned the Centre for International Economics to undertake an independent impact study of Pawsey's activities. The full Impact Report is available by request via https://pawsey.org.au/pawsey-impact-infographic/.

EVIDENCE-BASED IMPACT

Pawsey was recognised as making a major contribution to scientific knowledge and Australia's capabilities across many sectors. Several aspects of the Centre's operations were identified as generating economic value to Australia:

- The type of research undertaken using supercomputing is typically game changing — insights are more accurate, inevitably more timely, and many of the ideas that are tested could not otherwise be explored.
- Extremely large amounts of information can be stored at the Centre and accessed in real time, providing even greater research potential for current and future generations.
- The cost of research to users of the Pawsey Supercomputing Centre is substantially reduced, meaning research and development investments go further, and the process of research is more efficient and effective — identifying research endeavours with clear potential to generate value to Australia, and sharing learnings in a timely way to redirect other endeavours to more productive uses.

Three specific case studies from the 194 projects facilitated by the Centre in the past year were examined in detail to show the varied use of the Pawsey Supercomputing Centre to explore the biggest scientific questions, and quantify the economic and other impacts arising from Pawsey activities and the value this facility brings to Australia.

In the first year of Pawsey's operation as a node of the All Sky Virtual Observatory (ASVO), 202 individual users accessed approximately 640 TB of data.

CASE STUDY 1: IMPROVING GAS TURBINE EFFICIENCY

Research teams have undertaken multiple research projects at Pawsey to maximise gas turbine efficiency, accounting for 58.7 million core hours on Magnus between 2014–2018. Improvements have come about through determining the optimal design of each engine component to maximise performance, and new combustion modes have been investigated that have the potential to improve operational efficiency, reduce fuel usage and reduce greenhouse gas emissions.

The improvements researchers have developed for electricity generation and aviation gas turbines have resulted in:

- acceleration of the gas turbine research and design process by at least a factor of 10
- modelling of gas turbines enabling a one percent boost to engine efficiency and an equivalent reduction in the cost of air transportation and electricity generation
- creation of $11 million in net value to the Australian economy between 2014–2050, and environmental benefits from reduced carbon dioxide emissions worth half a million dollars over the same period.

CASE STUDY 2: THE MURCHISON WIDEFIELD ARRAY

The Murchison Widefield Array (MWA) is a powerful low frequency radio telescope which was the first operational pathfinder telescope as a precursor to the Square Kilometre Array (SKA). The Pawsey Supercomputing Centre is essential to MWA operations, as its observations are transferred to Perth and processed by Pawsey's Cray XC30 system Galaxy. Pawsey also stores the 30–35 TB of data produced daily, and is responsible for the global distribution of MWA data. In the first year of Pawsey's operation as a node of the All Sky Virtual Observatory (ASVO), 202 individual users accessed approximately 640 TB of data.

The impacts of the MWA include:

- enhancing our understanding of the Universe
- improving Australia's reputation and connections with the international research community, and contributing directly to Australia now co-hosting the SKA
- data processed and stored at Pawsey can be used to enhance the security and productivity of industry sectors impacted by solar weather, such as power supplies and communication networks
- providing an alternative to active radar monitoring for space debris, enabling industry innovations to better protect satellites from damaging impacts
- industrial applications of improved signal processing, data processing, data storage and computing knowledge
- upskilling researchers and enabling their research, with a value to Australia estimated at $1.4 million annually
- attracting and retaining world class research talent in Australia, worth $9.8 million in net value to 2050.

CASE STUDY 3: ALTERNATIVE ENERGY

Since 2014 the Carnegie Clean Energy hydrodynamic modelling team has been using the Magnus and Zeus systems at Pawsey to optimise the design of CETO, an experimental zero-emission technology that converts wave energy into electricity. The team has used over four million core hours to run computational fluid dynamic simulations to model the behaviour of CETO, producing more reliable estimates of power, loads and motions in moderate and severe wave conditions, reducing the need for expensive, higher risk tank testing and in-ocean trials of prototypes.

The research undertaken at Pawsey has enabled:

- six generations of CETO technology, nine patent families and over 140 patents/patents pending
- improved probability of addressing the economic and operational challenges needed to achieve commercialisation and adoption of the new technology by government and industry
- improvements in research efficiency estimated at between $9.6 million and $66.1 million in present value terms out to 2050 assuming eventual adoption is brought forward by one year.

In the following pages a wider selection of case studies showcases the range of amazing outcomes enabled by Pawsey-supported research.

CASE STUDY 1

Partner Institution:

CSIRO

Throwing Light on Dark Matter

Huge space projects have huge data processing requirements and Pawsey Supercomputing Centre is working with astronomers to meet their computing needs as they explore the cosmos.

The Australian Square Kilometre Array Pathfinder (ASKAP) became fully operational this year, with an array of 36 antennas mapping the night sky 30 square degrees at a time. Prior to its completion, astronomers performed early science with a subset of the antenna array. The Widefield ASKAP L-band All-Sky Blind Survey (WALLABY) is one of the ten Survey Science Projects established to take advantage of ASKAP's unique observing capabilities, and has just detected seven new galaxies.

More than 100 scientists from around the world are on board the Wallaby Survey. It is being led by Professor Lister Staveley-Smith from ICRAR.

THE CHALLENGE

Dr Karen Lee-Waddell, a data processing and imaging scientist on WALLABY, is measuring the hydrogen spectral line of a galaxy to determine its distance and mass. Galaxies spin faster than their mass would suggest, so by comparing the mass and spin of galaxies, WALLABY can calculate the amount of dark matter present. Dark matter is one of the greatest mysteries in science and measuring it allows astronomers to understand how it affects the universe.

"We don't fully understand dark matter. We assume galaxies are in these balloons of dark matter because of the extra mass, but we don't know the shape of that concentration. We don't really know how it affects things. We do know it has a gravitational pull, forming galaxies and stars," says Dr Lee-Waddell.

ASKAP's ability to map large sections of the sky is vital to this research, allowing astronomers to catalogue hundreds of thousands of new galaxies. But the array is huge and generates more than 200 terabytes of data per day. "The projected data rate of the full array is nine terabytes per hour," says Dr Lee-Waddell. "Currently, we're ingesting about half that. That's equivalent to streaming about 4,000 Netflix movies at once."

THE SOLUTION

The massive amount of data produced by ASKAP requires supercomputing power to process and store. Pawsey's Galaxy supercomputer performs batch processing for ASKAP, turning raw data into usable information for astronomers. The aim is to get as close as possible to 'real time' data processing, producing images as quickly as the raw data arrives.

"Right now, we're streaming four terabytes of raw data an hour, generally for 12 hours at a time. We have to do a lot of computations and processing to get usable data out of that. That's where Pawsey comes in. We take these terabytes of data and make insightful images of galaxies," says Dr Lee-Waddell.

Storage limitations mean data is refined down to five per cent of its initial size, concentrating the truly important data that astronomers can use. For WALLABY, it's galaxies that can be measured for their mass and other properties. However, the other data may still contain momentous discoveries, depending on what you're looking for. The current upgrade to Pawsey's supercomputing infrastructure will help astronomers process and store more data, helping fuel new discoveries.

THE OUTCOME

The early science recorded using only 12 of the 36 ASKAP antennas has already led to the detection of seven new galaxies, only scratching the surface of understanding dark matter and its distribution in the Universe. With ASKAP mapping more sky than ever before, many more galaxies are expected to be found, providing new insight into the physical processes occurring in these systems.

"With the full array, we will survey 75 per cent of the entire sky," confirms Dr Lee-Waddell. "WALLABY predicts the detection of well over half a million new galaxies. Understanding these different parts of the Universe helps us understand the fundamentals of everything around us."

WALLABY's work to measure dark matter will contribute to our knowledge of gravity, even as Australia harnesses this force to develop new industries and scientific discovery in space. With the help of Pawsey, Dr Lee-Waddell is processing terabytes of data to explore the Universe. CSIRO acknowledges the Wajarri Yamaji as the traditional owners of the Murchison Radio-astronomy Observatory site.

Understanding different parts of the Universe helps us understand the fundamentals of everything around us.

Dr Karen Lee-Waddell
Project Scientist

36 antennas

7 newly detected galaxies

9 billion solar masses

78 million light-years away

Project Scientist: *Dr Karen Lee-Waddell* | **Systems**: *Galaxy* | **Areas of science:** *Astronomy*
Applications used: *ASKAPsoft*

CASE STUDY 2

Partner Institution:

Curtin University

Searching for Missing Matter

Up to 40 per cent of the matter in the Universe is 'missing'. Astrophysicists suspect it resides in the warm-hot intergalactic medium as filaments of a 'cosmic web' stretching between galaxy clusters and superclusters. Using the MWA and Pawsey, Mr Torrance Hodgson is developing indirect imaging techniques to try and map this cosmic web.

1,500,000
core hours allocated

7,500
observations to process

> 100 hours
observations to combine

Project Leader: *Mr. Torrance Hodgson*
Systems: *Magnus* | **Areas of science:** *Cosmology and Extragalactic Astronomy* | **Applications used**: *MWA tools*

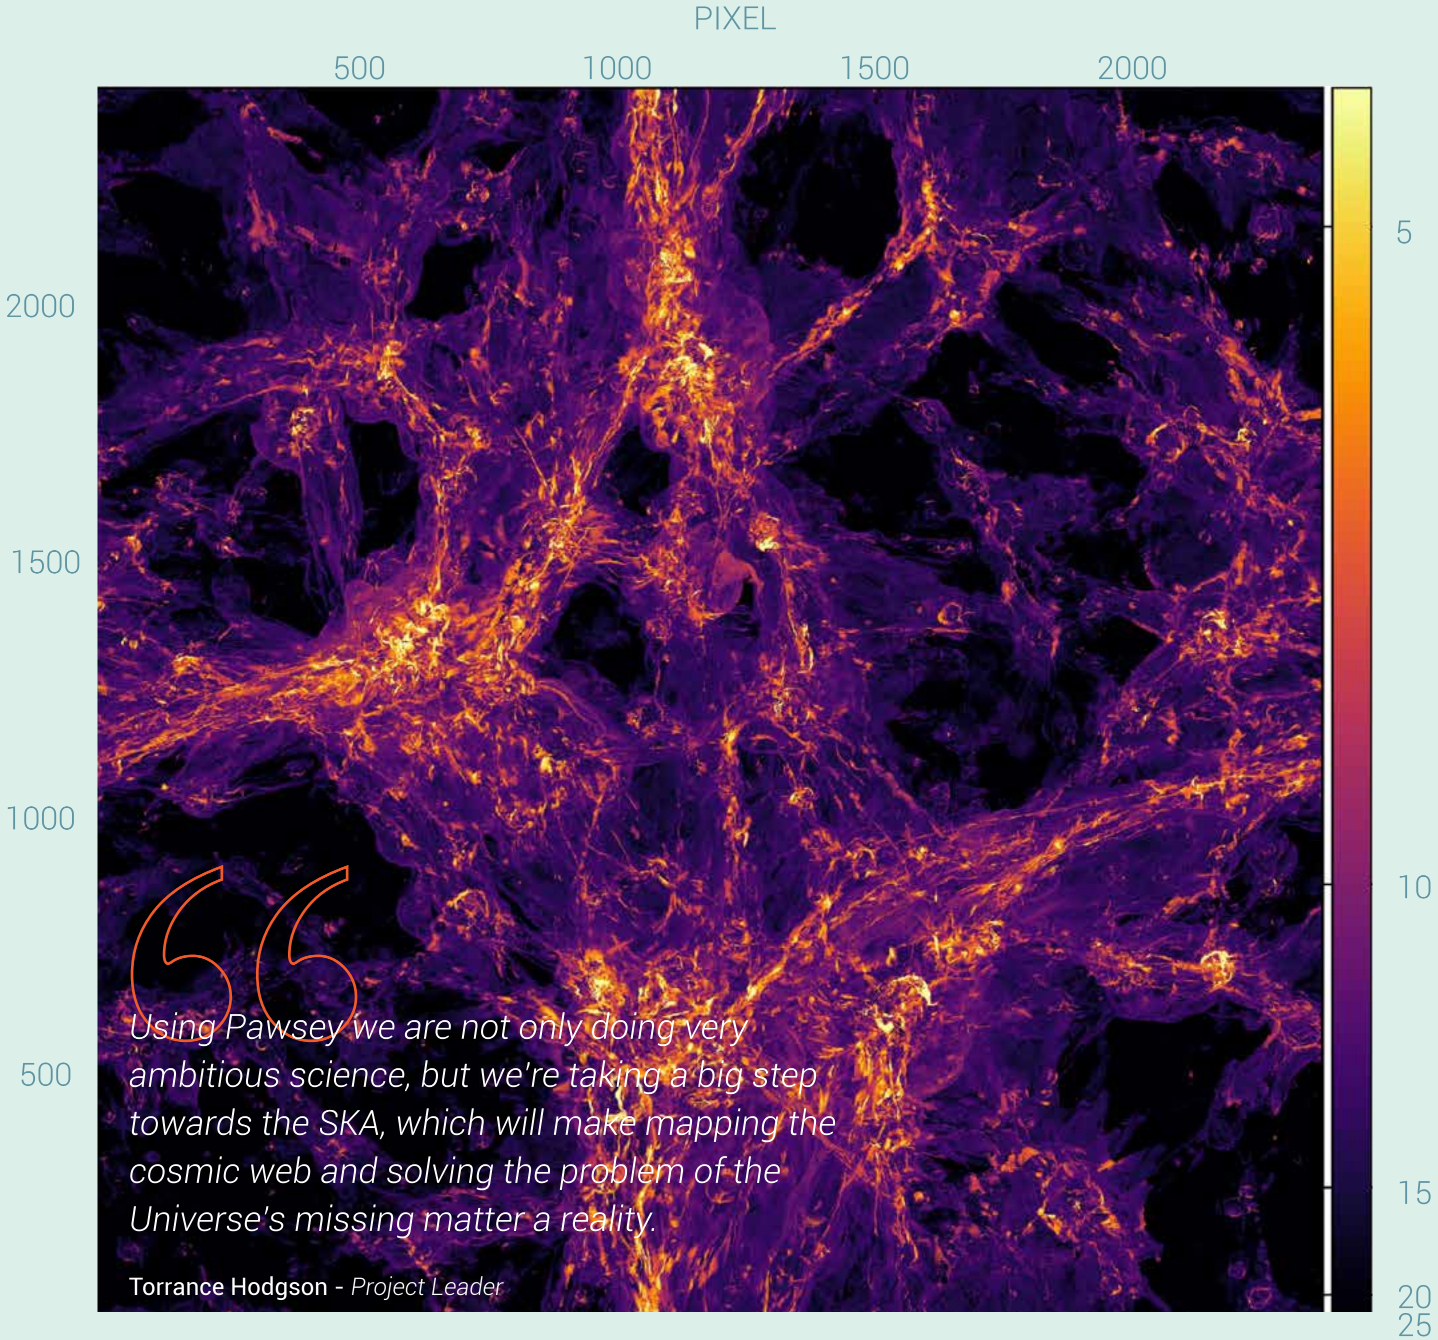

(Vazza, Ferrari, et al. 2015)

"Using Pawsey we are not only doing very ambitious science, but we're taking a big step towards the SKA, which will make mapping the cosmic web and solving the problem of the Universe's missing matter a reality.

Torrance Hodgson - *Project Leader*

THE CHALLENGE

Most matter in the Universe resides in galaxy clusters, which are visible to both optical and radio telescopes. The 'missing' matter is suspected to be in the warm-hot intergalactic medium, existing at a temperature that makes it effectively invisible to our telescopes. One indirect way to detect it is to look for radio emissions resulting from the electronic shock fronts that occur as this dispersed matter coalesces into larger filaments of the cosmic web.

This approach obviously requires looking deeply into the Universe across a large field of view for the very faintest signals. In all likelihood it will require the power of the SKA to make a positive detection of the cosmic web, but it hasn't been built yet.

THE SOLUTION

Mr. Torrance Hodgson, a PhD student at Curtin University, is working within the limitations of the MWA by compiling multiple datasets over time to increase its effective imaging power. He explains: "Phase I of the MWA was a three-kilometre diameter antenna array that had amazing sensitivity to large extended structures in the sky like the cosmic web, but with relatively low resolution. When the MWA was expanded to Phase II with a larger, six-kilometre diameter array, the resolution improved massively, but we lost the sensitivity to see those large-scale diffuse structures. We're trying to combine observations of the same part of the night sky taken four years apart with first the Phase I and then Phase II configurations, to combine their respective strengths and give a better chance of detecting the cosmic web."

The data sets are enormous, as a single two-minute observation is up to 15 gigabytes of data. Hodgson is attempting to combine about 50 hours of observations from Phase I with 50 hours of observations from Phase II. Pawsey supercomputing is required to not only store all of the data from the MWA, but also run the algorithms to integrate these two datasets.

"We're using a new wide-field imaging algorithm called image domain gridding," says Hodgson. "It has been developed by colleagues in the Netherlands to take advantage of GPUs so we can undertake massively parallel computation. We're one of the first research groups to use this new algorithm, as we need something scalable to combine all of these MWA observations. We're effectively testing and helping to de-bug the algorithm, which is designed to be able to scale up to SKA-levels of data management just by adding additional GPUs."

THE OUTCOME

Once the computation (requiring 3.5 million CPU hours to process) is complete, Hodgson will have created the deepest images ever with the MWA in search of the cosmic web. He's sanguine about the outcome. "We're unlikely to make a complete detection, as even the combined MWA data may not be enough," he admits. "But a statistical detection, or even a null result, will allow us to place an upper limit on the density of the cosmic web and the strength of its magnetic field. We'll know we have to look for something even fainter. That's still telling us more than we currently know about the properties of the cosmic web."

The developing and refining of new algorithms that can scale up to a problem this size, and equally to SKA-levels, is also a significant achievement. "Using Pawsey we are not only doing very ambitious science, but we're taking a big step towards the SKA, which will make mapping the cosmic web and solving the problem of the Universe's missing matter a reality."

CASE STUDY 3

Partner Institution: CSIRO

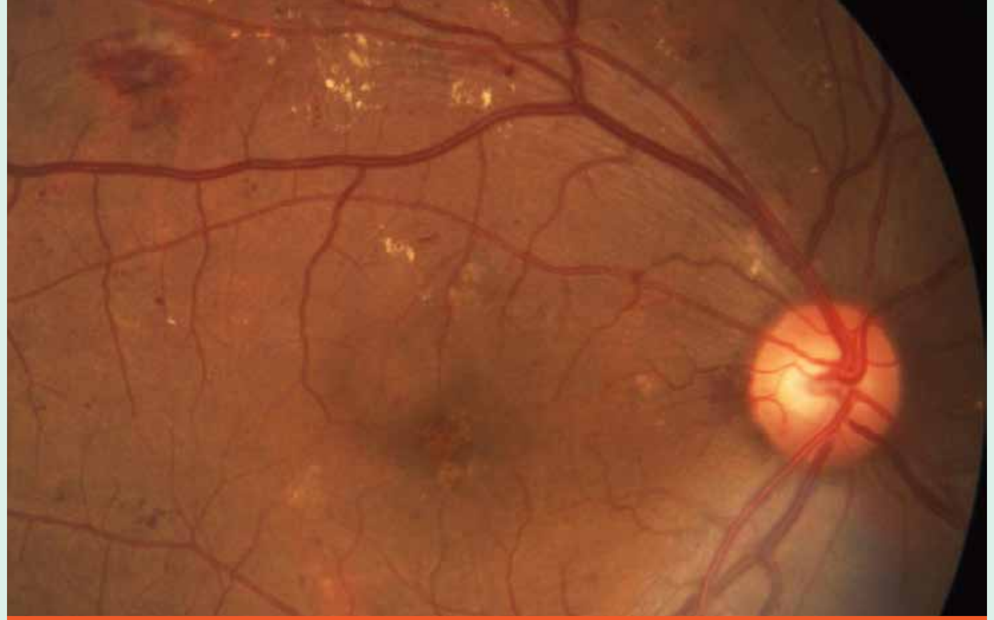

70,000 core hours allocated

291 patients tested in clinical trials

30 GP clinics providing service

Project Leaders: *Professor Yogesan Kanagasingam and Professor Xiao* | **Systems**: *Zeus*
Areas of science: *Medical data analytics* | **Applications used**: *Tensorflow, Keras, optical coherence tomography*

AI Learning Helps Save Sight

Diabetic retinopathy is a condition where blood vessels in the back of the eye begin to leak or bleed. As the condition worsens, the leakage damages a patient's sight. In the long term, untreated diabetic retinopathy can lead to irreversible blindness. It is the world's leading cause of vision loss, with anyone with diabetes at risk.

Diabetes affects over 1.2 million Australians, and one third show signs of diabetic retinopathy. If it is detected early it can be treated, and diabetics are encouraged to take yearly eye exams.

Professor Yogesan Kanagasingam is Director of the CSIRO Australian e-Health Research Centre in WA. With Pawsey Supercomputing Centre's Zeus system, Professor Kanagasingam, and his team, including Ms Maryam Mehdizadeh and Professor Xiao, developed an Artificial Intelligence (AI) program called Dr Grader to allow general practitioners (GPs) and primary care providers to detect diabetic retinopathy simply by taking a photo of a patient's eye.

> *Soon, this AI developed with Pawsey supercomputers may become a regular piece of equipment in every GP clinic, helping thousands of diabetes patients keep their sight.*
>
> Professor Yogesan Kanagasingam & Professor Xiao
> *Project Leaders*

THE CHALLENGE

Diabetic retinopathy cannot be diagnosed by a GP, but requires the expertise and equipment of an ophthalmologist. Professor Kanagasingam says many patients delay ophthalmologist visits due to long waiting lists and high cost, and put off these checks until their vision begins to deteriorate. Unfortunately, at this stage the diabetic retinopathy is already advanced and treatment becomes much more difficult.

To address this problem, Professor Kanagasingam and his team developed Dr Grader as a simpler way to detect diabetic retinopathy, allowing the check to be performed by GPs.

"There are a limited number of ophthalmologists, and they need to screen a large number of patients. But diabetics usually have to go to a GP as a first point of contact, so using AI there can make a big difference," says Professor Kanagasingam.

GPs will be able to take an image of a patient's eye, then Dr Grader can analyse the image to detect any signs of diabetic retinopathy. Having this detection system with GPs will decrease the load on ophthalmologists and encourage patients to get their eyes tested more regularly.

THE SOLUTION

Professor Kanagasingam and his team trained Dr Grader's AI using deep learning techniques, a form of machine learning for computer programs inspired by the human brain. The program is 'trained' to recognise diabetic retinopathy symptoms by giving it a large image data set of affected eyes to learn from.

Training the program requires a large amount of processing power, which is why Professor Kanagasingam needed to work with Pawsey Supercomputing Centre. "We can use the Pawsey systems to run these advanced learning techniques in a very short time. Timing is very important to process the images as each image is roughly 15-20 megabytes."

As well as recognising diabetic retinopathy, Dr Grader needs to be easy for busy GPs to use. Professor Kanagasingam has designed it to be fully integrated with GP medical databases, and include real-time image quality checking.

"We can't ask the patient to come back again if the image quality isn't high enough for Dr Grader, so we've included an AI quality control system which immediately checks the image and lets screeners know if they have to retake the photo. Now GPs can push the button, take the photo and everything is recorded, read and graded automatically."

THE OUTCOME

Dr Grader has already run a successful clinical trial at GP Superclinic @ Midland Railway Workshops in 2017 and 2018. There, Dr Grader successfully tested 291 patients for diabetic retinopathy. It is now being rolled out to 30 more GP clinics in Singapore and there are plans to implement Dr Grader at major hospitals in Western Australia.

"A major hospital may have 2,000–3,000 diabetes patients they scan annually. Ninety per cent of these patients are normal and test negatively with an ophthalmologist. With Dr Grader we can reduce strain on hospital resources. We're looking at setting up a server to process these hospital images now," says Professor Kanagasingam.

The Dr Grader technology has also been licenced by ophthalmology diagnostic imaging company, TeleMedC and is being rolled out in Singapore, Malaysia and South Asia. Its use is currently limited in Australia, but the technology is becoming cheaper and more portable. Soon, this AI developed with Pawsey supercomputers may become a regular piece of equipment in every GP clinic, helping thousands of diabetes patients keep their sight.

CASE STUDY 4

Partner Institution:
THE UNIVERSITY OF QUEENSLAND AUSTRALIA
CREATE CHANGE

Coal seam gas (CSG) is a valuable export commodity. It is also used domestically in a range of applications, from generating low-emission electricity and fuelling natural gas-powered buses, through to producing fertiliser and plastics. It's a vital part of the Australian economy. Pawsey is now helping researchers to develop more efficient ways to extract this natural resource.

Understanding Fluid Behaviour to Improve Gas Well Productivity

THE CHALLENGE

Dr Christopher Leonardi and Dr Mahshid Firouzi from The University of Queensland, along with graduate student Mr Travis Mitchell, are formulating advanced computational models to understand how gas and water flow in CSG wells. With these models, engineers can more accurately simulate the transport of hydrocarbons through pipeline networks, allowing them to design and build more efficient extraction systems.

"The fluid physics in a CSG well is an important boundary condition for the modelling of gas and water flow in the coal seam. It's also an important input to the downstream gathering network of gas pipelines. Understanding the physics is essential for predicting the flow regime and pressure distribution in the well," explains Dr Leonardi.

Coal seam gas is initially trapped by water pressure within the seam. This water must be pumped to the surface to liberate the gas. The characteristic ways that water and gas travel together in a pipeline are referred to as flow regimes. CSG wells exhibit a variety of regimes, from bubbles in a continuous fluid through to annular flow, where the gas travels through the centre of the pipe while the liquid water coats the pipe walls in a thin film. The dynamics and ultimately the impact of these flow regimes is what the team have modelled.

"If you don't understand the flow regime and pressure profile in the wellbore, then you'll not only have poor control over the individual well, but also over the upstream and downstream systems. You'll have a poorer understanding of how your production infrastructure is working as a system."

Modelling complex fluid dynamics is a challenging task that requires significant computing power. To further complicate matters, the team needed to simulate the flow under a wide variety of operating conditions, including the pipe's angle of inclination, the ratio of water to gas, and the well pressure. For this, they turned to the Pawsey Supercomputing Centre.

THE SOLUTION

To tackle these large-scale calculations, Mitchell used Pawsey's Magnus and Zeus supercomputers. By deploying multiple simulations at once, the team could drastically cut down on work time.

The team also took advantage of Zeus's post-processing software, which allowed them to visualise the results of their simulations at Pawsey. "One of the benefits of access to Zeus is the web-based visualisation tools," says Mr Mitchell. "We were able to post-process our results via the web using a simple laptop, without having to transport gigabytes of data across the country to local, high-powered machines".

The team's models have the ability to help other engineering fields that deal with liquid-gas flows, from oil and gas, to chemical reactors, to studying environmental flows.

Dr Christopher Leonardi
Project Leader

9	4,500,000	4	$5,000,000
team members	core hours allocated	industry partners	in funding received over the last five years

Project Leader: *Dr Christopher Leonardi* | **Systems**: *Magnus / Zeus* | **Areas of science:** *Petroleum and Reservoir Engineering*
Applications used: *Python, OpenMPI, cran/getopt, cran/optparse, cran/numbers, llaniewski/rtemplate, llaniewski/gvector, llaniewski/polyAlgebra, RInside*

THE OUTCOME

"Operating on large data sets is computationally demanding. If you tried to run it on a single processor, it could take months to do a simulation. With Pawsey we could bring that down to just a few hours, depending on the size and length of the simulation," says Dr Leonardi.

Not everyone who needs to simulate the extraction of CSG has access to a supercomputer. But that's where the team's research becomes really useful. Engineers can use Mr Mitchell's research to fine-tune simplified models developed for real-time computation on desktop hardware, improving their reliability and accuracy.

The team's models are also useful in other engineering fields that deal with liquid–gas flows, from oil and gas, to chemical reactors, to studying environmental flows.

"Because we're incorporating so much of the relevant physics into the model, it has a broad range of potential applications. Although our current focus is on bubble propagation, the same model can be quite easily manipulated for use in a range of other engineering disciplines," says Dr Leonardi.

However, the team's primary focus remains on trying to improve the efficiency, safety and fundamental understanding of the flow physics associated with CSG production, an industry that facilitates low-emissions energy and generates billions of dollars in export revenue. Using Pawsey, they are improving the productivity of a vital industry and shedding light on the fundamental behaviour of fluids.

CASE STUDY

5

Partner Institution:

UNSW AUSTRALIA

Climate modelling helps us predict and prepare for the future. Knowing how and when weather changes, helps protect our industries and cities. Professor Jason Evans of the Climate Change Research Centre at the University of New South Wales, and the ARC Centre of Excellence in Climate Extremes, is using the Pawsey Supercomputing Centre to create a more exact model of Australia's climate to better detail the changing climate in local areas.

Better Climate Models to Protect Australia

Multi-model mean changes between the present day (1990-2009) and far future (2060-2079) in annual. a. Number of very heavy precipitation days (R20mm), b. contribution from very wet days (R95). c. maximum of consecutive dry days (CDD) and d. maximum consecutive wet days (CWD). Stippling indicates that the changes are 'significant agreeing' at the 5% level. White circles (left to right): Adelaide, Melbourne, Sydney, Brisbane

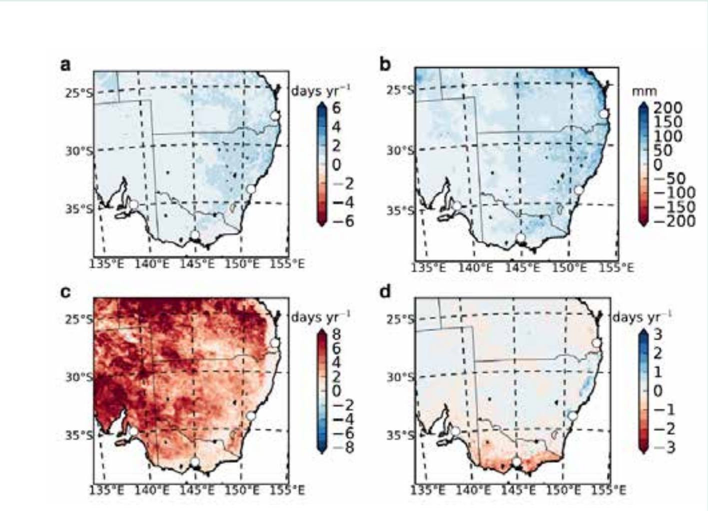

THE CHALLENGE

Professor Evans takes global models that predict how the climate will change in the future, then makes them more precise to map Australia and its surrounding waters in smaller sections, to understand exactly how local conditions will be affected as the climate changes.

"One of the biggest problems with trying to model climate is that there is this huge range of spatial scales that are really important to get right. From the micrometre scale of aerosol particles that act as nucleation sites for water droplets and ice crystals, to the kilometre scale of clouds, to the global scale of winds," says Professor Evans. As clouds and winds pass over, the weather they produce may change depending on the landscape. Water bodies, vegetation and terrain all play a role in affecting weather.

"At the moment we map in 50 km segments over the whole continent. We get down to 10 km in parts of Australia; mostly in the south east, as that's where most of the people live," says Professor Evans. "Mapping areas at this scale can show the differences in weather."

THE SOLUTION

To tackle climate data on this scale, Professor Evans needs supercomputers. The modelling required sending data across the country, to be processed at Pawsey and partially stored at the Climate Change Research Centre.

"These models need a lot of compute time to run and they produce a lot of data. There are challenges with data storage, data analysis and assessment. We were continuously transferring data back to UNSW from Pawsey. We're very appreciative of Pawsey's high-speed connection to Australia's Academic and Research Network."

Simulations are only as accurate as the data you gather for them. Variations between different climate models often occur because the Earth's climate is all interlinked, but only subsets of data get considered in each model. If you want to predict what the climate will be like in one Australian suburb, you really need to know what the climate is like everywhere else on the planet. That's an unrealistic amount of data.

To solve this problem, Professor Evans takes existing simplified climate models and then refines them with more comprehensive Australian data to give a more detailed view of local areas. "The processes that you need the model to capture occur on much smaller scales than we've been thinking about in the past. So, we're looking to go to higher resolution to better capture these climate extremes, and facilities like Pawsey are absolutely vital. They just can't be captured any other way."

THE OUTCOME

Part of Professor Evans' research considers how climate change affects Australian rainfall. Many industries need an accurate understanding of when and where rain will fall to plan their future. A change in rainfall amounts or seasons can wreak havoc on agricultural crops, extreme weather conditions can affect insurance, and the water cycle is vital for planning city growth.

"This research gets used in lots of different ways. We partner with water authorities who are interested in the security of water supply – what happens during dry spells and droughts. Other industries are much more interested in what's happening with very heavy rainfalls and flooding. Construction codes for street drainage need to be able to handle that one-in-a-hundred-year storm."

Professor Evans' research also aims to understand how extreme weather events can affect Australia's water cycle. Fires, extreme heat, rainfall, more intense tropical cyclones and droughts are all predicted to increase. "We're looking at what the models say about changes in heat waves and fire weather and compound events – multiple extremes occurring in quick succession or in nearby locations," explains Professor Evans. "Exploring how these events impact Australia will help us plan to minimise damage."

With the help of Pawsey Supercomputing Centre, Professor Evans will continue making more precise models of Australia's climate, to help us plan for the future and protect our industries, our cities and ourselves.

Project Leader: *Professor Jason Evans* | **Systems**: *Magnus / Zeus* | **Areas of science:** *Climate* | **Applications used**: *Weather research and forecasting model*

16	3,500,000	12	10 Km
team members	*core hours allocated*	*simulations*	*resolution models*

CASE STUDY 6

Partner Institution:

Murdoch University

7	606,600	5 km	$103,838
team members	core hours allocated	regional models	direct funding

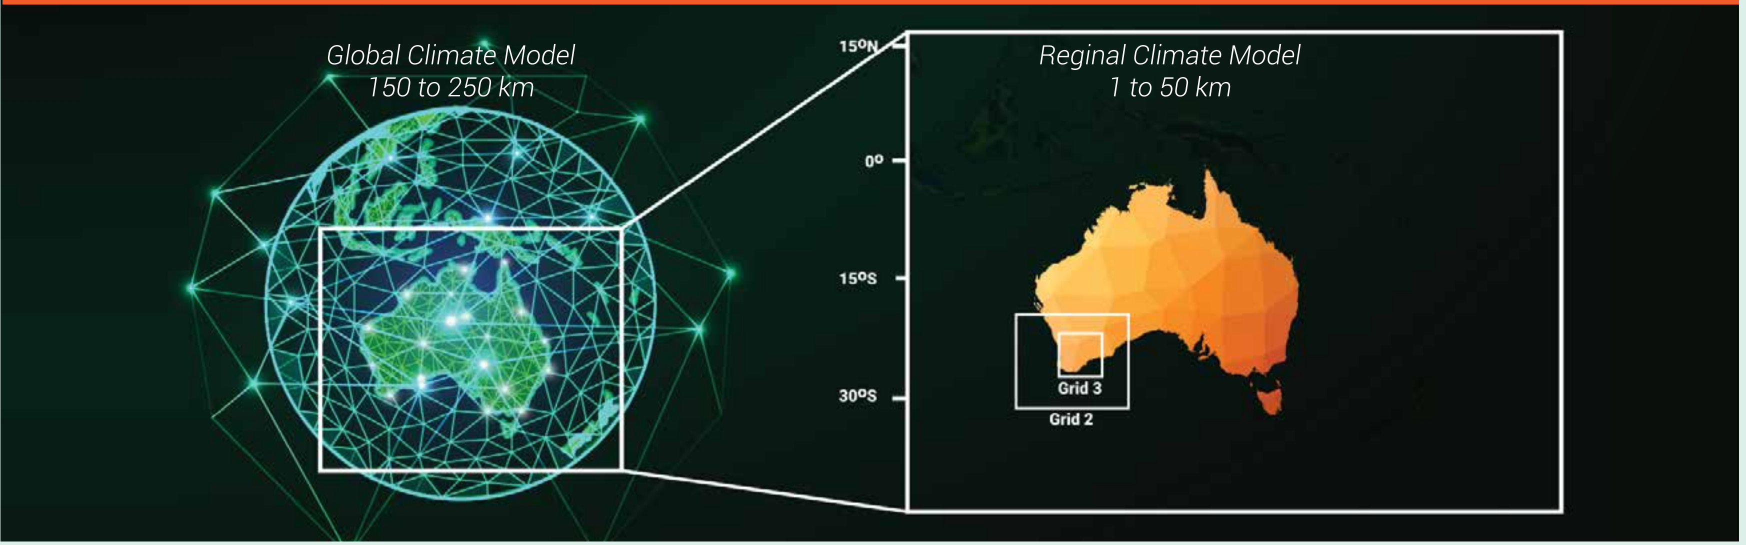

Project Leader: *Dr Jatin Kala* | **Systems**: *Magnus* | **Areas of science:** *Atmospheric Science* | **Applications used**: *Weather Research and Forecasting (WRF) model*

From Global to Local: Climate Change Predictions at the Suburban Scale

The south-west of Western Australia is a region of significant cropland farming and viticulture, worth $3,000-5,000 million annually. It is also a globally-recognised biodiversity hotspot and home to some of Australia's most iconic forests. Everything the south-west is notable for is sensitive to changes in rainfall and temperature.

Global climate models (GCMs) are used to simulate past, current and future climate across the entire globe, at a spatial resolution of 100–250 km. But to better understand how climate change impacts manifest at local scales, Dr Jatin Kala, an atmospheric scientist at Murdoch University, has been using Pawsey supercomputing and data storage to develop climate models at a 5 km resolution. These can be used to support local decision making in the south-west of Western Australia.

THE CHALLENGE

The south-west of Western Australia, a biodiversity hotspot and a region of significant agricultural production, is known to be vulnerable to climate change, with global models showing significant trends towards a warmer and drier climate across the entire south-west. However, Global Climate Models (GCMs) have a limited ability to simulate future climate changes at a more detailed local scale. To understand exactly how these changes will impact on the local environment, industries and communities in this region, and to plan appropriate response strategies, a more detailed regional climate projection is needed.

Dr Jatin Kala, an atmospheric scientist at Murdoch University explains: "The south-west of Western Australia is only represented by three or four points in the GCM, which is like saying the climate impacts on Perth and Busselton are exactly the same. But our climate and weather impacts are obviously more locally variable than that."

Regionally-relevant information is needed to assess how climate change will impact on local physical, ecological and socio-economic systems, and to inform policy and planning responses. The challenge is to create higher resolution projections of climate change, to inform adaption planning and impact assessment studies across agriculture, forestry, conservation and urban planning.

Kala's research team is using regional climate models, which can dynamically downscale GCMs to a resolution of 5 km, and incorporate the finer-scale influences of topography, land cover, and mesoscale weather systems like sea breezes and winter cold fronts. This requires the use of both high-performance computing and data storage resources as the models become exponentially more detailed and complex.

THE SOLUTION

Pawsey has been involved in climate modelling research since before it was even Pawsey. Kala's early research was carried out on an iVEC computer with a total of only 192 individual processors, and his climate simulations were always limited by the computational power and data storage available. But using Pawsey's current facilities, his regional climate projection models run using several hundred cores on Magnus (which has a total of 35,712 cores) and produce hundreds of terabytes of data.

"We need the supercomputing facilities to run the models, we need the data storage, and then to analyse all that data, we need to run even more code. This sort of research just isn't possible without Pawsey," says Dr Kala.

"We can now do higher resolution projections over longer periods of time. We can evaluate how well our models simulate detailed historical weather records across a region over the past 30 years, and then run them forwards in time with greater confidence. It is making our regional climate models much more useful."

THE OUTCOME

Dr Kala's research group is continuously refining and improving high-resolution climate projections for the south-west of Western Australia. These models are being used to explore climate change in the region, and specific issues like the impact on local viticulture or bushfire behaviour.

His datasets are currently available to other researchers on request, but together with other climate change researchers using Pawsey's partner institution the National Computational Infrastructure (NCI), Dr Kala is working towards making regional climate model outputs covering all of Australia more generally available to the community.

With modern supercomputing, providing state-of-the-art climate prediction modelling at a finer resolution for local impact assessment is becoming a reality.

Mean annual difference between historical (1970–1999) and future (2030–2059) simulations for the Diurnal Temperature Range (DTR), annual maximum temperature (TXX), annual minimum temperature (TNN), Frost Days (FD) and Summer Days (SU) indices. Differences for Frost Days (FD) and Summer Days (SU) are averaged over the growing season and stippling shows where the changes in the distribution of the indices is statistically significant at the 95 % confidence level.

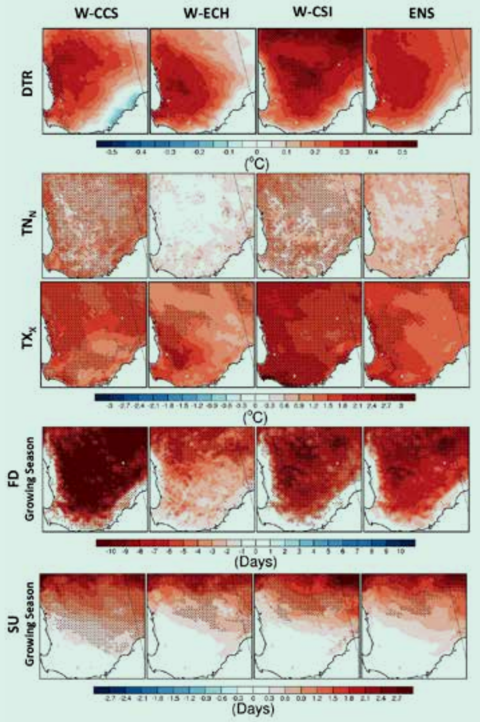

CASE STUDY 7

Partner Institution:

Curtin University

Pyrenophora tritici-repentis specific hairpin sequence element

10 chromosomal groups mapped | **8** new annotated genomes generated | **480,000** core hours allocated

Project Leaders: *Dr Paula Moolhuijzen and Dr Caroline Moffat* | **Systems**: *Magnus, Zeus, Nimbus* | **Areas of science:** *Bioinformatics* | **Applications used**: *BLAST, EvidenceModeller, KAKS Calculator*

Fighting Fungus: Using Genomics to Combat Wheat Diseases

Wheat is one of the world's most important staple grains, and meeting the projected demand for it in the future is a significant challenge. Fungal diseases that affect wheat grain yield are a big part of the problem. Research into the control and management of crop fungal pathogens is critical to safeguarding this food source into the future.

Dr Paula Moolhuijzen and Dr Caroline Moffat from the Centre for Crop and Disease Management (CCDM), a purpose built centre co-supported by Curtin University and the Grains Research and Development Corporation, are using the Pawsey's supercomputer, Magnus, to develop and analyse new fungal pathogen genome resources to better understand the underlying molecular mechanisms used by pathogens to infect and damage wheat.

> *... with Pawsey we could compare entire genomes, millions of combinations, that could run over 400 nodes on Magnus in under five hours.*
>
> Dr Paula Moolhuijzen
> *Project Leader*

THE CHALLENGE

Yellow spot, a global fungal wheat disease, is responsible for yield losses and control costs estimated at over $600 million annually in Australia alone. The fungus, Pyrenophora tritici-repentis, has evolved into different races around the world, making combating it extremely difficult.

Until recently, Australian researchers had limited genomic resources to study yellow spot. Only a single American reference genome of Pyrenophora tritici-repentis was available, which had been generated using older sequencing technology, and it was surprisingly different to the isolates found in Australia.

CCDM researcher's Drs Moolhuijzen and Moffat wanted to create new, high-quality reference genomes for yellow spot isolates across Australia and other agricultural areas to progress their research. "High-quality reference genomes are needed to identify key variations between different isolates and races of this pathogen. We now have sequencing technology available to create more complete genomes for more comprehensive comparisons. We also wanted to create genome resources relevant to the southern hemisphere and the wheat growing areas of Australia," says Dr Moffat.

Dr Moolhuijzen adds, "Sequencing individual isolates and comparing them lets us examine the genetic basis of this disease and pinpoint the evolution of different modes of action. This however, means comparing millions of pairwise alignments between different genes and features, and a high-throughput data processing environment is needed for this."

THE SOLUTION

Drs Moffat and Moolhuijzen used Pawsey's supercomputing facilities to comparatively analyse 12 different isolates from five different race groups and determined the sequence variability between them.

"On a single computer, an in-depth analysis could only be done on a subset of genomic data or a targeted region, because a job of this type for a whole genome would take too long to finish," explains Dr Moolhuijzen. "But with Pawsey we could compare entire genomes, millions of combinations, that could run over 400 nodes on Magnus in under five hours. If we didn't have Pawsey resources it would take us over a year. This has allowed us for the first time to run genome-wide in-depth analyses to look for significant patterns and variations."

THE OUTCOME

Drs Moffat and Moolhuijzen have produced a high-quality genomics resource relevant for the southern hemisphere, which has progressed the understanding of the different races of the yellow spot pathogen. "Through identifying differences between isolates, we have been able to identify genes under selective pressure," says Dr Moffat. "We've also found that 30 per cent of the gene content in these isolates is variable. That's a considerable amount of variation – that degree of chromosomal reshuffling can really drive the evolution of more damaging isolates of this pathogen. When previously we were working with one fragmented American reference genome, our understanding of the fungus was limited to that. But seeing this variability gives us much more to work with when trying to understand why some isolates infect wheat differently to others."

These genomic resources are now publicly available, allowing other researchers to identify exactly which genes and toxins are involved in infecting and damaging wheat, to help breeders develop more resistant wheat varieties to control yellow spot disease.

"Pawsey has become a fundamental research tool for us, to get the 'big picture' on this fungus," says Dr Moolhuijzen. "And it will only scale up. We'll be studying more strains from around the world, and are now also looking at the wheat side of the pathogen–wheat interaction. Given the wheat genome is six times the size of the human genome, Pawsey supercomputing is vital."

CASE STUDY 8

Partner Institution:

Landgate

Satellite Snapshots of our Changing State

Western Australia's government agencies working across agriculture, conservation, resource management, planning and development, and emergency services need to know how our State is changing over time. Seeing the 'big picture' and recognising changes in vegetation coverage and the impacts of development, salinity and climate change starts with satellite imagery.

The Land Monitor Project coordinates the production of vegetation and salinity maps from satellite data, and has tracked vegetation changes over time since 1988. Using the Pawsey Supercomputing Centre, Land Monitor data and products are becoming more readily available to the government agencies that rely on this information to safeguard our State.

THE CHALLENGE

Many activities and decisions rely on maps and an understanding of our environment. Having a single repository of this spatial information and understanding how it may change over time is invaluable to many of our government agencies and local industries. The Land Monitor project (https://landmonitor.landgate.wa.gov.au/home.php) is a long-term collaboration between Western Australian government agencies to provide this information and update it annually. Initially created through the National Dryland Salinity program, the cross-government consortium now includes:

- Western Australian Land Information Authority (Landgate)
- CSIRO/Data61
- Department of Biodiversity, Conservation and Attractions
- Department of Primary Industries and Regional Development
- Department of Planning, Lands and Heritage
- Department of Water and Environmental Regulation
- Department of Fire and Emergency Services
- Water Corporation.

Mr Dan Sandison from Landgate explains: "CSIRO takes the analysis-ready Landsat satellite imagery from Geoscience Australia and creates a mosaic of the entire area. It then uses algorithms to identify and map the vegetation cover. It does this every year, and then processes the data against the data sets from previous years to create products that show the trends over time. Landgate coordinates the program, and until now has distributed the new data set and map products to all of the consortium partners annually."

The core products are interactive maps of vegetation extent, change and trends, which can be used and developed further by individual government agencies to suit their particular needs.

Land Monitor originally only covered Western Australia's southwest agricultural region, but in 2018 expanded to deliver vegetation mapping and monitoring products across the entire State, increasing the area monitored from 18 million to over 252 million hectares per annum. As the Land Monitor program grows – with some products at 25 m spatial resolution now covering the entire State, and being added to every year – making sure everyone who needs access to this resource has it, becomes more difficult.

> *Each of the consortium partners can directly access not just the latest data but insights from the entire 30-year monitoring effort as they need it.*
>
> **Mr Dan Sandison**
> *Secretary of Land Monitor and Project leader*

20	8	252 M	25 m
TB data stored	*Government agencies*	*hectares per annum monitored*	*spatial resolution covering WA*

Project Leader: *Mr Dan Sandison* | **Systems**: *Pawsey Storage* | **Areas of science:** *Earth Sciences* | **Applications used**: *pshell*

Government agencies: *Landgate; CSIRO/Data 61; Department of Biodiversity, Conservations & Attractions; Department of Primary Industries & Regional Development; Department of Planning, Lands & Heritage; Department of Water & Environmental Regulation; Department of Fire & Emergency Services, Water Corporation.*

THE SOLUTION

Pawsey is now supporting the Land Monitor program by hosting its entire data repository, from the raw satellite data collected every year, to all of the vegetation and salinity mapping products developed and their trends over time. "Each of the consortium partners can directly access not just the latest data but insights from the entire 30-year monitoring effort as they need it," says Mr Sandison. "It reduces a lot of duplication, there are no issues with version control as the data set expands, everyone has access to the latest information in its historical context, and they can then use it as they need to."

THE OUTCOME

Pawsey's storage and access arrangements for Land Monitor will continue as the State's 'satellite snapshot' gets revisited every year. There is also significant scope for the mapping and monitoring program to grow.

Land Monitor has also been working with Geoscience Australia to develop the Copernicus Hub in Western Australia, which is hosted at Pawsey. This allows research on the European Union's Sentinel satellite series to map the state at 10 m pixel resolution with a more frequent orbit cycle.

Having access to this richer data also means more mapping products will be developed for the consortium partners. "With Pawsey providing such ease of access to the data, it will help us develop more detailed vegetation products, incorporating things like vegetation condition and research into new products, not just vegetation presence or absence over time. And all of our consortium partners will benefit from that new level of awareness of our State's health."

PROJECTS AND PUBLICATIONS

1,744 researchers | **200** projects | **306** publications

During the reporting period, 1,744 researchers undertook 177 projects using Pawsey supercomputing and support services. This research is regularly shared with the wider research community, with 306 publications arising from Pawsey research appearing in international journals and conference proceedings during the period. These reflect the growing diversity of scientific research being accelerated through the application of High Performance Computing (HPC), as well as highlighting traditional areas of application and strength within Western Australia and Pawsey's longer-term user base.

Radio astronomy and astrophysics research is a recognised strength and was well-represented with 68 publications. Engineering research accounted for 52 publications, and geosciences for 19 publications. Life sciences is a growing area of research at Pawsey, with 38 publications across plant science, environmental science, biology and health research. The nexus between chemistry, materials science and nanotechnology is also a very diverse area addressing applications from catalysis and energy research to quantum computing, and accounted for 101 publications.

The full list of publications can be found online at www.pawsey.org.au.

Examples of published works illustrating the diversity and impact of research arising out of the Centre include:

- **AGRICULTURE** – A large international collaboration has assembled the first pangenome of Brassica napus, the family of oilseed crops including rapeseed and canola. They have shown that homologous exchange is a major cause of gene presence/absence variation, with some of these variable genes predicted to be involved in important agronomic traits including flowering time and disease resistance. Their study highlighting the route to access this novel source of genetic diversity, which could be captured for the improvement of this important crop species, was published in the open access Plant Biotechnology Journal.

- **CHEMISTRY** – Solar fixation of nitrogen to produce ammonia would make production of this essential agricultural and industrial feedstock chemical more sustainable and cost effective. A Chinese and Australian team used extensive first-principles calculations to design a metal-free single atom photocatalyst for the solar fixation of nitrogen. Their design and the highly stable catalyst's expected properties were published in the Journal of the American Chemical Society, a preeminent journal of chemistry founded in 1879.

- **ENGINEERING** – An American and Australian team have used experimentation and numerical simulation to study the dynamics of an optical axisymmetric scramjet – an air-breathing engine designed for hypersonic function – at Mach 4.5. They demonstrated supersonic flame holding and stable combustion under both mass and heat loading and described the scramjet's flow and flame behaviour. These design insights are aiming to improve combustion processes at altitude. The results were shared in the Proceedings of the Combustion Institute.

- **ENGINEERING** – The trailing edge of a high-pressure turbine blade is exposed to substantial heat loads and requires efficient cooling to operate. Researchers from the University of Melbourne together with two General Electric companies have applied machine learning to large-eddy simulation data to improve the nonlinear turbulence and heat flux modelling around these trailing edges and better predict cooling effectiveness for both the aviation and energy generation industries. The training and performance of the model was published in the Journal of Turbomachinery, produced by the American Society of Mechanical Engineers.

- **GEOSCIENCES** – Two CSIRO researchers have provided new insights into the genesis of magmatic nickel/copper/platinum group elements deposits which can inform the exploration targeting of these important mineral deposits. By using computational fluid dynamic simulations to model the multiple mineral deposition processes, kinetics was shown to play a major role, with the process operating at the shortest time scale dictating the location and type of ores formed. Their insights were published in the open access Geoscience Frontiers.

- **HEALTH** – A large international team has raised the possibility that genetic studies of hip shape may help in understanding the pathways involved in hip osteoarthritis and hip fracture. They have completed a genome-wide association study of hip shape and its relationship to the risk of both hip osteoarthritis and fracture, and identified eight single-nucleotide polymorphisms associated with variable hip shape. Their study was reported in the Journal of Bone and Mineral Research.

- **RADIO ASTRONOMY** – Our understanding of the population and behaviour of fast radio bursts originating beyond our galaxy is limited because until now only 34 transient radio bursts have been reported. A large Australian team conducted a well-controlled, wide-field radio survey for fast radio bursts and identified a further 20, including both the nearest and most energetic bursts detected so far. Their findings were published in Nature, the world's leading multidisciplinary science journal.

- **RENEWABLE ENERGY** – To design more effective fuel cell catalysts, a team from the University of Queensland have explored the effect of nitrogen-doping levels in metal heteroatom graphene catalysts for the hydrogen evolution reaction. They found that increased nitrogen doping makes the metal atoms more active towards hydrogen adsorption, providing a way to modify transition metals to improve their catalytic ability. Their results were published in the Journal of Chemical Information and Modelling, published by the American Chemical Society.

- **QUANTUM COMPUTING** – American and Australian collaborators used a scanning tunnelling microscope to investigate the influence of electronic valleys on exchange coupling interactions in a coupled donor-quantum-dot system, a basic building block proposed for quantum information processing. Their findings were published in Physical Review X, published by the American Physical Society to showcase research with the potential to have long-lasting and profound impact.

Meet the people accelerating scientific breakthroughs at Pawsey

Voices of Science

OUR PEOPLE

Dr Merran Smith Board Member

Synergies and emerging needs

Dr Merran Smith joined the Pawsey Board this year as the new UWA representative. She brings her background in science, health and economics to the Board, with strengths in strategic management and performance monitoring. She is also well versed in data management, as for more than 10 years she was in charge of the WA Department of Health's Information Centre. While with the Department, she established data linkage as a core service and served on a number of Australia's peak national health information committees. Since 2009 she has been the inaugural Chief Executive of Australia's Population Health Research Network.

Reflecting on Pawsey's current position in the data landscape, Merran is keen to help with the next steps: "I think Pawsey is doing an excellent job, particularly supporting some of the recognised science strengths in Western Australia such as astronomy, engineering and physics. But there are new opportunities with its expanding scale. It's very exciting now that Pawsey has additional funding and can plan for a broader role within the Australian research infrastructure ecosystem."

"My interest is in the expanded use and additional value that Pawsey will be able to provide in research areas across Australia, while continuing to support existing users. For example, until recently, health systems generally didn't require the compute and storage capacity of a supercomputing facility. That is changing rapidly with growth in demand related to genomics, other 'omics' and imaging technologies," says Merran. The amount of detailed clinical information is also growing, as more electronic health data is available across the country. "Australia is already an international leader with a wealth of linked health data, so I suspect that big-data-driven health research will be a significant growth area for Pawsey in the future."

"A challenging part of joining the Board was just getting my mind around what needs to be done over the next phase of Pawsey's growth. It's good to see that the collaboration with the National Computational Infrastructure (NCI) is strengthening, and through the National Collaborative Research Infrastructure Strategy (NCRIS) program we can coordinate and stage the infrastructure investment across Australia's two Tier 1 supercomputing facilities. I think Pawsey is already well-positioned within the broader national research infrastructure and that position will strengthen with the new investment."

Merran considers that Pawsey is very fortunate to have a highly experienced Chair of the Board (Mr John Langoulant AO), and other very well-credentialed Board members. "There is also a good working relationship between the Executive Director and the Board. This is particularly relevant in two key areas. One is the focus on the Capital Refresh. We all understand that it's a priority to get the tenders out to market, identify the preferred providers and organise the process of implementing the new, expanded systems. It sounds simple, but it's a complicated process, and we'll need to run both the current and new systems in tandem for some time, with the practical issues that causes. But the Board fully understands and is supporting the organisation in making this happen."

"There is also a focus on continuing to provide value for current users and stakeholders while the refresh is happening," Merran stresses. "The Board is working collaboratively with the organisation to ensure that service standards are maintained as this transition is occurring."

Merran is relishing the opportunity to make a difference to the reach and impact of High Performance Computing (HPC) across Australia. "As a Board we've got the opportunity to ensure that Pawsey's developing structures and processes continue to support our existing users and stakeholders, and can extend to support new activities. It is still early days for health but the demand could look very different in five years' time. Other fields are also starting to realise the potential in supercomputing. It's fantastic that this emerging need is happening at a time when Pawsey is expanding its capacity."

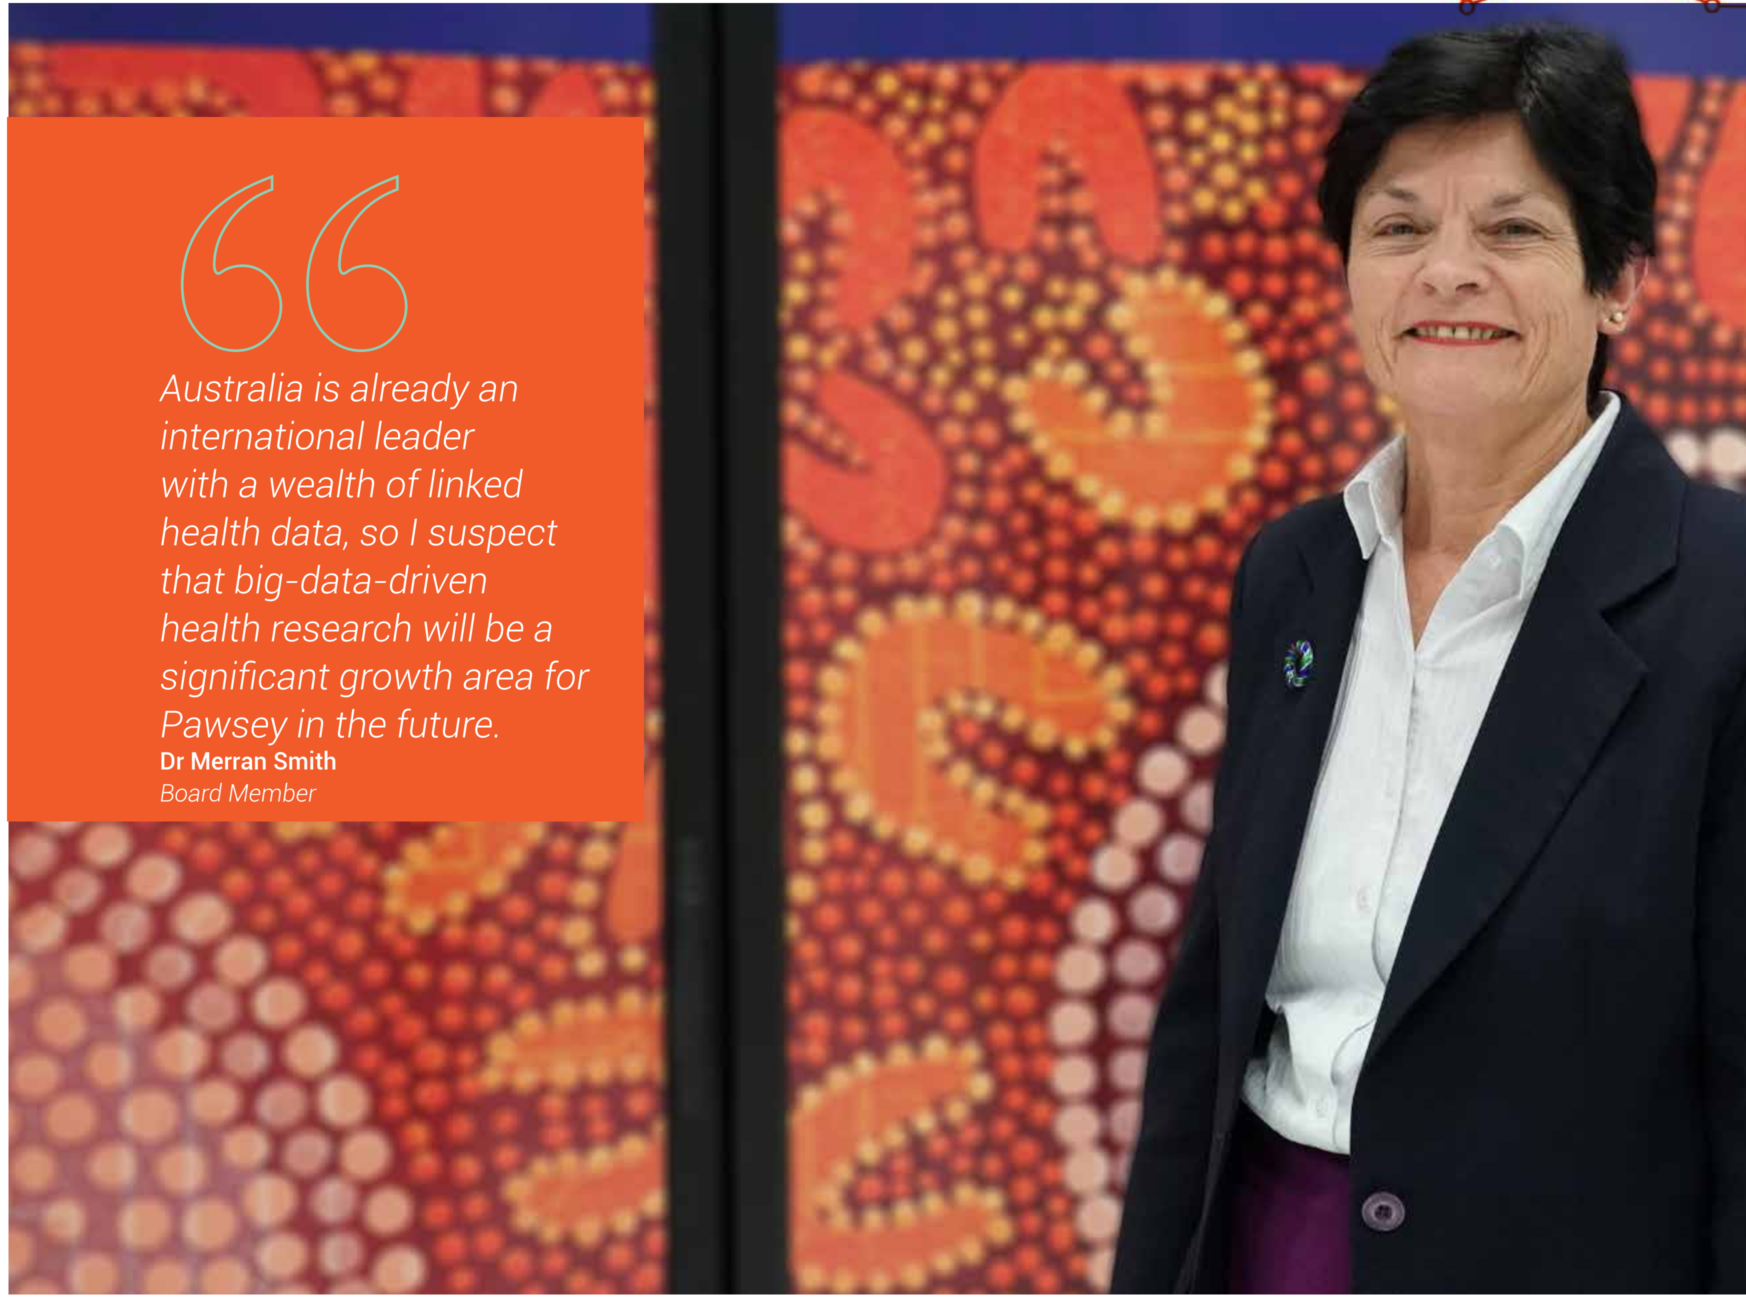

Australia is already an international leader with a wealth of linked health data, so I suspect that big-data-driven health research will be a significant growth area for Pawsey in the future.

Dr Merran Smith
Board Member

OUR PEOPLE

Professor Melanie Johnston-Hollitt
Professor, Curtin Institute of Radio Astronomy;
Director, Murchison Widefield Array
Area of Science: Astrophysics

Building big telescopes to unravel the structure of the Universe

Professor Melanie Johnston-Hollitt studies cosmic magnetism and galaxy clusters, the largest gravitationally-bound structures in the Universe. But to do this, she is also involved in developing and building the scientific instruments needed to study the Universe in such detail.

ABOUT PROFESSOR MELANIE JOHNSTON-HOLLITT

Melanie always intended to be an astrophysicist, but after starting out cleaning photomultiplier tubes to study cosmic rays at the Woomera Rocket Range, and then trying optical astronomy at the Anglo-Australian Telescope at Siding Springs Observatory, she settled on radio astronomy. She has since worked on the design and construction of several major international radio telescopes and is currently Director of the Murchison Widefield Array (MWA), the first operational pathfinder telescope for the Square Kilometre Array (SKA).

WHAT DREW HER TO SCIENCE?

"It sounds like a terrible cliché," Melanie admits, "but it started when I was about two, and my grandmother used to take me outside and show me the night sky, and tell me about the stars and the planets. I really loved it, so when my father suggested that I should follow that interest and become an astronomer, that's what I decided to do. I was an extraordinarily stubborn kid, and I held onto that idea right through school, aiming to eventually do a PhD in astrophysics."

RESEARCH WITH SUPERCOMPUTERS

Melanie has been constrained by computation her entire career. She explains: "During my PhD, I'd make observations at the telescope, but then have to go to CSIRO to process the data, as their computer cluster was bigger than anything we had at the university at the time."

"Eventually computing became more powerful, and there was a period about 10 years ago when you could produce and reduce astronomical data on your laptop. But now we've gone through a renaissance with radio astronomy, we've built a collection of new telescopes which produce absolutely prodigious amounts of data. Modern supercomputers are now absolutely necessary again to store, process and analyse it all."

The way of working is basically the same as it was before, but the scale improvements are literally astronomical. Putting it in perspective, Melanie reflects: "For my PhD it took me three years to collect my data and I ended up with just nine data points to draw my conclusions from. My students can now get that amount of data and that result in a day, just because the telescopes have got much better, and the processing has got so much faster. We're getting closer to the answers than we've ever been able to before."

REAL WORLD SOLUTIONS

While Melanie is trying to understand the history of how magnetism arose in the Universe, and how magnetism affects astronomical processes today, particularly in galaxy clusters, much of her research is working out how to scale up processing the data from these modern telescopes. Algorithm development and optimisation is critical to make the most effective use of modern instruments, and make next-generation telescopes like the SKA possible.

It's at the intersection of astronomy, physics, computing and mathematics that down-to-Earth applications occur.

"We're driven by curiosity, and curious about things that are really hard to understand, like galaxy evolution. But in the process of trying to answer our questions, we develop technology that is useful in so many other ways too. People who want to understand the physics of the Universe have come up with algorithms that led to our current Wi-Fi standards. Interferometry, the fundamental technique that modern radio astronomy is based on, is now used in medical imaging, and has been revolutionising diagnostic medicine and human healthcare in just the last 10–20 years. And through studying our ionosphere and working out how it distorts the radio signals we're trying to measure, radio astronomers have enabled accurate transmissions of signals from satellites to the ground, which underpins our Global Positioning Systems (GPS)."

OUR PEOPLE

Dr Tara Hopley
Research Scientist, Department of Biodiversity, Conservation and Attractions in Western Australia
Area of Science: Genomics

Dr Tara Hopley uses supercomputers to map patterns of genetic diversity in native plant species. Understanding the variability in plant populations is necessary to assess their conservation status, inform environmental impact assessment and rehabilitation work, and is a starting point for understanding any climate change adaptation.

ABOUT DR TARA HOPLEY

Tara studied environmental science with a focus on soils, but a seven-week botanical internship at the Australian National Herbarium put her attention on plants, which eventually led to a PhD in plant genetics. She now applies her skills to better understand and manage the diversity of plant life in Western Australia.

WHAT DREW HER TO SCIENCE?

"As a kid I spent a lot of time outside, hiking and camping with my family," Tara recollects. "I always felt at home out in the bush." Her subsequent career has been a natural progression of the things she enjoyed growing up. "I never really picked science, I just kept doing what I enjoyed at school, which ended up being more science, and then eventually environmental science at university."

"The plants and the diversity over here in Western Australia are so different to what I grew up with on the east coast. I love my field trips in Western Australia's south-west – the wildflowers are amazing."

Mapping Individuality: When do variations within a species become important?

RESEARCH WITH SUPERCOMPUTERS

Analysing and comparing the genomes of different individuals within a species of plant to map its natural variability across large areas of land requires a supercomputer. Tara explains: "Understanding the relationships among populations and species is essential for effectively identifying and managing species. To clarify the boundaries between species, sub-species and local varieties, or to identify the origins of invasive weeds, we need detailed genomic information on a range of individuals."

That detail may run to sequencing large parts of the nuclear genome or the entire genome of chloroplasts, generating hundreds of gigabytes of data. This data is analysed to provide information addressing a range of conservation and management issues such as identifying cryptic species (species that are genetically distinct and can't interbreed, but look identical), determining gene flow through pollen or seed dispersal, and understanding potential adaptation to climate.

Tara and the genetics group also use genomics to understand patterns of genetic diversity and gene flow across natural environments. "Comparing tens of thousands of single nucleotide polymorphisms in each individual is very compute intensive. If we're looking at the diversity of a species across a wide area, we might use up to 1,000 samples. So you really need a supercomputer."

A major focus of Tara's work has been analysing the genomes of species to identify areas that show adaptation to climate variables. This can inform management strategies such as assisted gene migration in forest management and climate-adjusted provenancing of seed in ecological restoration. "Identification of adaptation to warmer and drier conditions means we can move seed from these populations to other populations to prepare them for hotter conditions in the future," Tara explains.

REAL WORLD SOLUTIONS

The Department of Biodiversity, Conservation and Attractions is using this information to inform plant conservation and environmental restoration works.

"If we're sourcing seed to grow new populations, where should we be sourcing that seed from?" Tara asks. "Should we be mixing seed from different populations, are they too similar or too different – should we be keeping those populations distinct? Where should we source seeds with our future climate in mind? To get the best seed with adaptive potential for the future, instead of using local seed should we collect in areas that are historically warmer and drier so new populations can adapt to the future climate?"

Pawsey supercomputing and comparative genomics is helping to answer all of these questions, identifying the relative significance of local plant populations, the impact of their potential removal, and improving the outcomes of revegetation efforts.

OUR PEOPLE

Professor Jingbo Wang
Head of Physics Department,
The University of Western Australia
Area of Science: Quantum computing

Professor Jingbo Wang has spent her entire career solving the Schrödinger equation, trying to accurately describe the state of quantum-mechanical systems. She now simulates the smallest possible systems – single atoms, electrons and photons – in order to demonstrate the power of quantum computing.

ABOUT PROFESSOR JINGBO WANG

Jingbo grew up in a small Chinese city during the now-infamous Cultural Revolution, when over one billion people suffered extreme poverty and millions of children were without proper education. She was extremely fortunate to be given the opportunity to study physics at Sichuan University. She subsequently won a scholarship to complete a PhD in quantum physics at the University of Adelaide, and has been working in that field ever since.

Jingbo's research ambitions have seemingly scaled 'downwards' over time. Frustrated at the approximations needed to model the behaviour of condensed matter systems, she moved to much simpler systems that could be more precisely studied and controlled. From electron transport in crystalline material to diatomic molecular spectroscopy, to electron–atom scattering, and eventually to precision control of single atoms, electrons and photons, Professor Wang's research now has direct applications in quantum information and quantum computing.

RESEARCH WITH SUPERCOMPUTERS

Jingbo uses the Pawsey Supercomputing facilities to simulate the behaviour of prototype quantum computers. Where classical computers encode each bit of information as either a '0' or a '1', quantum computers hold information as qubits, which can encode both '0' and '1' at the same time.

"Three qubits can encode the first eight numbers, from zero to seven", Jingbo explains. "Adding a fourth qubit doubles its capacity to sixteen numbers. Each time you add just one qubit, you double the computing capacity of the system. That's why quantum computing has the potential to be so incredibly powerful. If you wanted to double Pawsey's computing capacity, you'd need to build a second Pawsey. But to double the capacity of a quantum computer, you only have to add one extra qubit."

The difficulty lies in that each qubit is a single atom, electron, or photon. Building the hardware systems that allow the control and manipulation of individual qubits is extremely challenging, and is still at the very cutting edge of research in only a few laboratories in the world. A different sort of computer logic is also required to manipulate the astronomically large number of calculations performed with multiple qubits. A set of quantum gate operations (not, Hadamard, phase and control-not) will replace the classical (not, or, and) logic gates, which operate in quantum parallel on a coherent linear superposition of states.

Successfully demonstrating the operation and power of quantum computing logic is very difficult when first-generation quantum computers do not yet exist. Jingbo points out: "The first quantum computers are going to be incredibly expensive and require extreme conditions to operate. To get there, we need to demonstrate that they can actually do something useful and provide the results we're after. So we're simulating these quantum logic computations for a few qubits using Pawsey, and demonstrating that they can be extrapolated and scaled up."

Programming the smallest particles to harness the power of quantum computing

REAL WORLD SOLUTIONS

Running quantum computation simulations on a classical supercomputer, Jingbo's team have been able to develop efficient quantum algorithms for machine learning, and to solve combinatorial optimisation problems such as working out the most efficient way to schedule jobs for multiple processors, or maximise the number of irregular objects that can be packed into a container.

They have also demonstrated how to extract local and global structures from complex networks, and decide which nodes are the most important – the equivalent of calculating Google rankings for every webpage on the internet. Another practical problem is that networks with similar topological structures can actually look quite different, and demonstrating the structural similarity is extremely difficult. Professor Wang's team have developed a quantum-inspired algorithm that can identify classes of networks with comparable structures, and provide a quantitative measure of network similarity.

"We used to run our simulations on desktop computers, but if we were looking at the behaviour of a network, we could only simulate very simple systems. That scale doesn't really characterise the properties of the network, and consequently we cannot expand the logic to larger systems. But using Pawsey we are able to demonstrate the algorithms, the software for quantum computing, does have practical application, and will be ridiculously powerful."

"We've really come full circle – we need a classical supercomputer just to simulate how we will eventually get a quantum computer to operate. Their fundamental elements are so simple – single photons or electrons – but showing that we can manipulate their interactions using quantum logic operations is taking all the computational capacity we currently have. We need to demonstrate the power and value of these systems, so quantum computing eventually becomes a reality."

OUR PEOPLE

Dr Raffaella Demichelis
ARC Future Fellow, School of Molecular and Life Sciences, Curtin University
Area of Science: Computational chemistry, geochemistry, materials science

Nature's Lego: Mastering the building blocks of the universe

Dr Raffaella Demichelis wants to understand not just how atoms interact, but why they interact the way they do, to better predict atomic interactions and then use this knowledge to eventually design specific materials. This understanding may eventually be the key to address many environmental, technological and medical challenges in the future.

ABOUT DR RAFFAELLA DEMICHELIS

Raffaella studied chemistry and then discovered the power of computing to explore the theoretical aspect of chemistry at the University of Torino in Italy. After completing her PhD, she was keen to experience living in another country, and took up a postdoctoral position within the Computational Materials and Minerals group at Curtin University to further develop and apply theoretical and computational methods to problems in chemistry. She was subsequently awarded an Early Career Research Fellowship at Curtin, and then an ARC Future Fellowship, extending the 'living abroad' experiment for almost nine years, and starting a family in Australia in the process.

WHAT DREW HER TO SCIENCE?

"I was always interested in science, and got into chemistry quite early. I remember being in grade seven, learning what water molecules look like and how they rearrange and structure themselves during melting, freezing, evaporation and condensation," she recollects. "That was probably the start of my interest in why molecules are the shapes they are, how they fit together, and how that affects their behaviour."

Chasing this fundamental understanding of chemistry led to a career in theoretical and computational chemistry, although Raffaella admits: "When I started at university I wasn't interested in computers at all, I didn't even know what computational chemistry was. But I realised very quickly that theoretical chemistry basically doesn't exist without computation. The scale of the calculations required makes it impossible to do any other way."

RESEARCH WITH SUPERCOMPUTERS

Trying to model chemical interactions at the atomic and electronic scale requires solving many complex equations that describe how atoms interact. As the atoms move around, these equations need to be solved over and over again.

"It's not just atom A interacting with atom B either", Raffaella explains. "If you're looking at a real-world example like how a molecule interacts with a surface, you need to simulate a few thousand atoms to generate a realistic surface, and reactions often occur in water, so you need to simulate all the water molecules near the surface and around your molecule of interest as well. And then you need to model all those interactions for microseconds to identify the structural arrangements that are the most stable and lowest energy, which tells you chemical information about the system."

It's millions and millions of computer operations – there's an awful lot of atomic movement between thousands of atoms in a couple of microseconds. "You can't run this on your laptop. It takes days even with 200 Central Processing Units (CPUs)."

REAL WORLD SOLUTIONS

Raffaella is using these techniques to work out how to mimic some of the amazingly precise reactions that happen in nature, and harness those processes to address human needs. One strand of her research is studying crystal growth, and how it is affected and controlled by biological environments.

"If we can understand how nature shapes shells, bones and teeth, how biological molecules affect the shape of the growing carbonate, phosphate and silicate minerals, we could design and grow materials with specific shapes, sizes and functions," she explains. "That level of control is beyond the current state-of-the-art in material science and design."

Her research not only focuses on controlling mineral shape, but also mineral reactivity. Another major project is studying how mineral surfaces react with carbon dioxide, to emulate naturally-occurring reactions that convert carbon dioxide to more 'useful' organic molecules.

"If we can accurately model and understand these natural processes, we can design chemical systems to harvest carbon dioxide and produce fuel. It's a potential path to producing clean energy while also reducing our greenhouse gas emissions. But to get there we need to develop a clear understanding of the chemical reactions at a molecular level."

OUR PEOPLE

Professor Julio Soria
Director of the Laboratory for Turbulence Research in Aerospace and Combustion (LTRAC);
Director of Research, Department of Mechanical and Aerospace Engineering, Monash University
Area of Science: Aerodynamics and fluid mechanics

Overcoming drag

Professor Julio Soria studies turbulence, the key to reducing drag on vehicles, improving the efficiency of wind turbines, and efficiently transporting fluids in pipes.

ABOUT PROFESSOR JULIO SORIA

Julio is a researcher who attacks problems both theoretically and experimentally. He started his career as an experimentalist, building the equipment and instrumentation needed to study heat transfer in a transitional boundary layer subjected to surface vibrations. At the same time, he was studying numerical methods, because as he explains: "The idea of doing Direct Numerical Simulations (DNS) of turbulence was only just coming into being in the 1980s. I could see that it would be a fantastic way to get data we couldn't access experimentally, if we had the computing power to do it."

After completing his PhD, Julio was awarded a CSIRO postdoctoral fellowship which allowed him to go to Stanford University and NASA Ames Research Centre, where he learned everything he could about the direct numerical simulation of turbulent flows from the pioneers in the field, who also had access to more computing power than was available in Australia at the time.

Since returning to Australia, he has continued to explore and manipulate turbulent flows, both with newer laser-based optical experimental methods and more powerful DNSs.

WHAT DREW HIM TO SCIENCE?

"I loved mathematics," Julio recalls. "Mathematical analysis was my major interest when I finished school, but I couldn't see a career path for mathematicians at the time. I decided to study mechanical engineering because I figured I would still be exposed to mathematics, but I might actually get a job at the end of it!"

Julio had never considered a career in research until one of his university professors put the idea into his head as he was finishing his Honours degree. So, he embarked on his PhD research project, and has never looked back.

"I just love research, I love discovery, I love the analysis of the data and connecting the dots. There are beautiful challenges every day."

RESEARCH WITH SUPERCOMPUTERS

Julio's research is focused on understanding the structure of turbulent flows, particularly the turbulent boundary layer that is in contact with a solid surface – like air against an aeroplane wing, water against a ship hull, or oil against a pipe wall. The turbulent layer accounts for the majority of the drag which is why we need engines and pumps to move things around, burning fuel and emitting carbon dioxide in the process.

The equations that describe the flow of fluids have been known for over 100 years, but turbulence is still today an immature science that lacks fundamental understanding of its complex multi-scale dynamics. "Knowing the equations isn't enough," says Julio. "They're non-linear partial differential equations, with very limited analytical solutions. It is almost impossible to predict flow behaviour from equations without fundamental understanding and correct interpretation of real-world data."

Turbulent flows need to be resolved on multiple scales – from the big eddies which contain most of the kinetic energy, to the smallest where kinetic energy is converted to internal energy and for all intents and purposes, is lost. And not all of this data can be measured in experiments.

"Experimental methods can't resolve all of the scales we need to study in turbulent flow," says Julio. "With DNS we can resolve all the scales of turbulence, and see how they interact and change over time. We can also generate data for quantities that are virtually impossible to measure, such as pressure fluctuations inside the flow."

Practical, real-world systems have a very large range of scales contained within the turbulent flow, so the computing challenge becomes the limitation. Julio admits: "We've used every supercomputer that's been available to us, because these DNSs take of the order of 50 million CPU hours to do. It's because it doesn't give you a 'solution'; the DNS provides the spatial details of the entire turbulent flow over time, so we can then study the statistics of the structures within it."

REAL WORLD SOLUTIONS

Using Pawsey, Julio is starting to see the devil in the detail. "The simulations give us the complicated dynamics of these turbulence structures, how they're created, how they interact, and how they die and get dissipated. No experimental method can give us this data, so now we can ask more relevant questions about the dynamics, and explore how to manipulate these structures using different control strategies to reduce drag or increase lift."

"Around half of the energy used worldwide to move fluids through pipes and propel cars, trains, ships and planes is dissipated by wall-bounded turbulence. It's all drag," says Julio. "So you can see why we want to understand and control it better. Even getting a one per cent reduction in fuel use overcoming drag equates to saving $390 million and 900 million kilograms of carbon dioxide across Australia's transport system, each year."

OUR PEOPLE

Stacy Tyson

Project Manager for Pawsey Capital Refresh

Delivering growth

Although new to Pawsey in 2019, Stacy Tyson has the project management experience to see Pawsey's major system upgrades through to a successful conclusion and fully expanded operation.

With 15 years of project management experience in the mining industry, particularly managing major upgrades to mining railway signalling and control systems, Stacy has already seen almost every problem that can potentially derail a big project. "I've worked with international project teams developing the hardware and software solutions required for safe and efficient rail services for massive mining operations, so I'm using that experience to anticipate and negotiate any issues and steer Pawsey's Capital Refresh to full delivery in 2022," she says.

Stacy meets regularly with Pawsey users through both the Pawsey Capital Refresh User Reference Group and the Technical Reference Group. The User Reference Group is developing and endorsing the various current and future user requirements for each of Pawsey's systems, whereas the Technical Reference Group is endorsing the hardware and software specifications for each component of the upgrade to be able to meet these requirements.

"As a project manager, it's my job to bring a diverse group of people together to deliver a system that meets both the user and technical requirements," Stacy explains. "A lot of each day is spent communicating with the Pawsey team and our users, monitoring progress, and then reporting that progress to all of the stakeholders. I'm basically responsible for working with the Pawsey team to deliver the Capital Refresh on time and on budget."

A key part of Stacy's planning is managing the change from existing systems to the new systems while minimising any impact on current Pawsey users. This requires staging the installation works and allowing time to migrate users onto new systems, while ensuring the new systems can simultaneously meet the needs of existing users, and emerging future users.

"The Capital Refresh is a number of interconnected procurements that together will provide significantly increased computing power

and speed, expanded data storage capacity, and more efficient HPC services to Pawsey's users. In the first year of the project we've already progressed the procurement of additional long-term storage, and upgraded high-speed storage for radio astronomy to service the Australian Square Kilometre Array Pathfinder (ASKAP) and the Murchison Widefield Array (MWA). We've also developed the specifications for the upgrades of both the Centre's cloud compute services and general use high-speed storage."

The team is now developing the specifications for an additional compute cluster for radio astronomy, and has developed and endorsed the user requirements for the general-purpose supercomputer, in advance of finalising the technical requirements.

Coming from an engineering background, Stacy is particularly enthusiastic about the scientific developments that are being advanced using Pawsey facilities. "It's great being a part of the team that is upgrading the Centre's services, which will then facilitate even more exciting breakthroughs. With new users in areas like bioinformatics, and Pawsey's involvement in the SKA Project, the future for this Tier 1 HPC facility and the researchers we support is looking really bright."

With data processing and storage needs anticipated to grow exponentially in the future, Stacy expects that this refresh is only one (big) step on that upwards trajectory. "HPC evolves so quickly, it is hard to imagine the technology that will be involved in 10 years – it will be well beyond anything available now. Pawsey will probably be in the middle of another upgrade by then!"

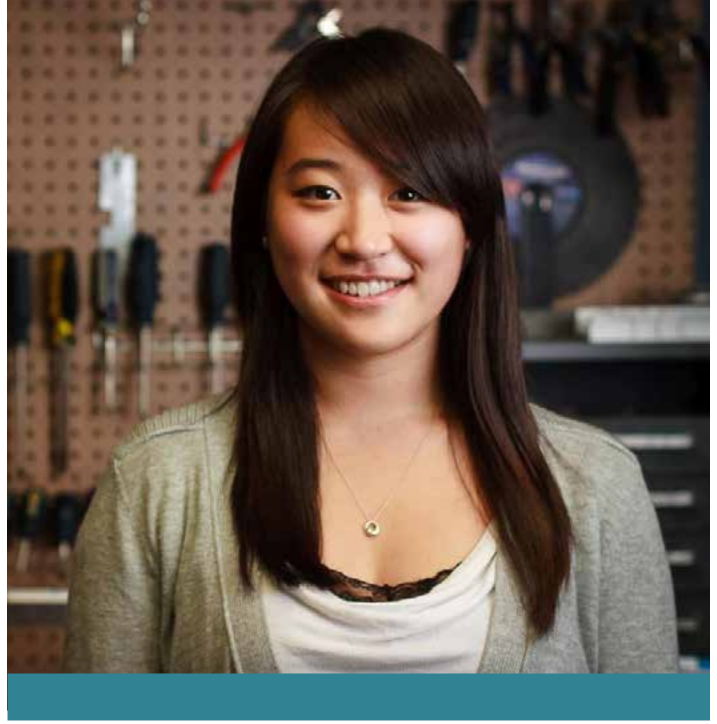

Jennifer Yeung

Pawsey HPC Research Fellow

Protecting babies and astronauts

Being a Pawsey user has led to trailblazing of a different kind for researcher Jennifer Yeung.

With a degree in aerospace engineering and a professional master's in aerospace design management, Jennifer worked in Canada's aerospace industry for five years and taught herself programming to use the data generated in aircraft designs to improve aircraft certification processes. Her study and work led her into big data analytics and her Master of Computer Science research that combines her technical skills in space and computer science with her passion for making positive change directly for people.

"My research at Pawsey involves an online health analytics platform called Artemis. Artemis provisions physiological data collected from bedside monitors in neonatal intensive care units with clinical algorithms to detect subtle changes in a patient's condition," Jennifer explains. "For remote hospitals where patient transport to a tertiary unit may be required to treat unique conditions but this may not be readily available, data can be streamed from the hospital through Pawsey cloud computing service, Nimbus, to be seen by urban specialists. This enables real time clinical decision support at the hospital without transporting the patient right away."

"Artemis can enhance the quality of patient care by allowing more timely diagnoses and faster recognition of deterioration in a patient's condition. This can improve health care and increase survival rates for our most vulnerable patients," she says.

In addition to monitoring premature babies on Earth, Jennifer's research is also extending the Artemis platform to similarly monitor astronaut health in long duration spaceflight.

Jennifer originally worked on Artemis through the Queen Elizabeth II Diamond Jubilee Research Scholarship at the Department of Health Western Australia as a Policy Officer. As a result of her term there, she was invited to become the inaugural Pawsey HPC Research Fellow in March 2019, a fellowship aimed at diversifying the technical professionals in the HPC industry and bringing different perspectives to HPC applications and solutions.

"I'm humbled by the opportunity to work with healthcare professionals, government officials, and supercomputing experts to implement innovations in healthcare that will make a positive difference for patients," Jennifer admits. "At the same time, working amongst the research and innovation opportunities here is incredible to witness – there are so many relevant issues that are actively being addressed and advanced by supercomputing."

From teaching herself to code as an aerospace engineer in Canada, Jennifer is now making a difference for patients in healthcare in Australia. "I'm incredibly proud of this feat, but it's quite the plot twist if you told me I'd get to work at a supercomputing centre, five years ago."

Working at Pawsey inspires Jennifer as it exposes her to projects that impact real-world problems that directly affect people and the planet. In turn, as the Pawsey HPC Research Fellow, she is a passionate advocate for diversity in Science, Technology, Engineering and Mathematics (STEM). "Whenever I get a chance to speak at events, conferences and seminars, I encourage students, staff and technical professionals to explore their passions, diversify their interests and explore cross-disciplinary careers. I'm also realising as a technical specialist that I have a responsibility to share the benefits of these innovations more widely, for example by being transparent about how user data is collected, how user privacy is protected, and ultimately how HPC and big data analytics can be regulated safely to benefit our day-to-day lives."

OUR PEOPLE

Dr Maciej Cytowski
Supercomputing Application Specialist

Upscaling research ambitions

A mathematician with a PhD in computational science, Dr. Maciej Cytowski has been working in supercomputing for 15 years. His expertise is focused on the optimisation and development of research applications on large scale (massively parallel) and accelerated HPC systems.

"It means that I spend a lot of my time working with Pawsey users to optimise and scale their applications, and port them to new architectures like Graphics Processing Units (GPUs)," explains Maciej. "Collaborating with them in this way lets them upscale their research, and really impacts on their research outcomes." His normal day sees him addressing any problems that have arisen for users, before jumping into programming, benchmarking and optimising research applications.

"We've got a great group of data, supercomputing and visualisation specialists here. Working with the researchers, we can really help them achieve things that were previously not possible. People notice the tools we have – supercomputing, visualisation walls, large scale storage – but the most important thing is how we help researchers take advantage of those tools to progress science outcomes."

Maciej particularly enjoys the diversity that comes with his daily work. "It's great to work across the various research fields that are represented by Pawsey users, and learn about their specific computational challenges. It's very rewarding to help them solve their problems, expand their research, and see the broader impact that has on the Australian research community."

Depending on the challenges that researchers bring to him, Maciej regularly collaborates with other specialists within Pawsey to optimise research applications. He admits: "Although Pawsey has a defined management structure, teamwork and collaboration tends to come ahead of management! I just talk to and collaborate with everyone, particularly the Cloud, Data, Visualisations and Marketing teams."

With his expertise in benchmarking and evaluating emerging computational architectures, Maciej is also heavily involved in Pawsey's Capital Refresh Project. "Right now is the most interesting time there has ever been (in my career) to work in HPC. There are so many CPU and GPU architectures being developed and in the market, and Artificial Intelligence (AI) and Cloud solutions are now also converging with HPC. It is transforming how we will use HPC resources in the future. As a result of this refresh, Pawsey will be leading the way in supporting Australian researchers to achieve completely new scales with their computationally-intensive research."

"We've grown from being a 'regional' HPC centre to a national resource, and prominent on the international HPC stage with strong collaborations with computing centres in Asia, Europe and the US. With the continuing support of our partner institutions and the State and Federal governments, and most importantly with the ongoing collaborations with our research users, we are consolidating our position as a truly world-class HPC node locally and internationally."

Dr Chris Harris
Senior Supercomputing Specialist

From the beginnings of GPU to HPC for all

Dr. Chris Harris has had a passion for computers as long as he can remember. He cites an early example: "For my Year Three presentation in primary school, I took apart a Compucorp 625 Mark II computer to show to the class, complete with Zilog Z80 8-bit processor, 64 KB of RAM, dual 5.25" disk drives, 80x16 character display, and built-in dot matrix printer. Soon after that I started learning to program in BASIC on an Amiga 1000."

That interest persisted to a PhD focusing on the use of GPUs for signal and image processing, and his research played a role in pioneering their use for radio astronomy signal correlation. From Year Three, his developing skill set was tailor-made for Pawsey, and he joined the Supercomputing Technology and Application Programming Team in 2012.

Since then, Chris has worked across many facets of Pawsey. "I started out providing helpdesk support for Pawsey researchers – helping them install software, set up workflows, and use the systems at scale. Over time, I've also coordinated the merit allocation processes for our systems, led our uptake projects, and been involved in outreach, training and user feedback activities."

"I really enjoy the fact that there is no 'normal working day' at Pawsey. There's such a variety of work, depending on the projects currently underway. I could be working on procurement documents, technical reviews of merit applications, or calls for uptake projects. There's always technical discussions with colleagues, and travel to provide face-to-face training and hold user forum discussions."

Chris is always looking ahead as HPC develops, exemplified by his work on technology demonstrator projects. "I loved working on our Advanced Technology Cluster 'Athena' – collecting the user requirements, coordinating the technical review before procuring the system, running the early adopter process and reporting on the appropriateness of the technology."

Projects like Athena have been critical in allowing Pawsey users to try out emerging technologies in preparation for the current Capital Refresh. Equally, the user forums that Chris coordinates are opportunities for both Pawsey users and staff to discuss Pawsey services and infrastructure, allowing staff to better understand user needs, and users to better understand Pawsey services and infrastructure. The resulting regular feedback is used to improve many aspects of Pawsey operations, and inform new developments like the Capital Refresh.

Chris points out that it's all fundamentally about enabling research. "I really enjoy working with researchers to help them get the most out of Pawsey's infrastructure. Coordinating the merit allocation calls, I've developed a real appreciation for the breadth and depth of computational research that we support across hundreds of projects and users from so many different scientific disciplines. It's so rewarding to see the transformative effect our services can have on our users' research. We've taken workflows that take months to process on workstations and moved them to supercomputing systems where they complete in hours. This increases the researchers' productivity by drastically reducing the time to result, and often also improves research quality, by allowing researchers to iterate and improve on those results in a reasonable timeframe."

"I think HPC is becoming a mainstream necessity as the scale of data and associated computation grows across all areas of scientific research. It will be an ongoing process for us to become more accessible for a wider range of end users."

Growing to upscale our scientific ambitions, outcomes and impact

A New Era of Research

PREPARING FOR GROWTH

Pawsey Capital Refresh

To ensure the Pawsey Centre can continue to provide internationally-competitive supercomputing and data services, a $70 million capital refresh project funded by the Federal Government is underway to secure the next generation of supercomputers, data and supporting infrastructure for Australia's research communities.

This upgrade will allow Pawsey to keep pace with global advances in supercomputing technology, and provide significantly increased computing power and speed, expanded data storage capacity, and more efficient high-performance computing services for Pawsey users.

In this first year of the project, the Pawsey Board has approved governance and sub-committee frameworks for the capital refresh to provide oversight, accountability, and to direct the expenditure of funds. A project manager has been appointed to oversee the program, and the terms of reference and membership of four Reference Groups – User, Technical, Contract and Finance, and Change Management – have been established.

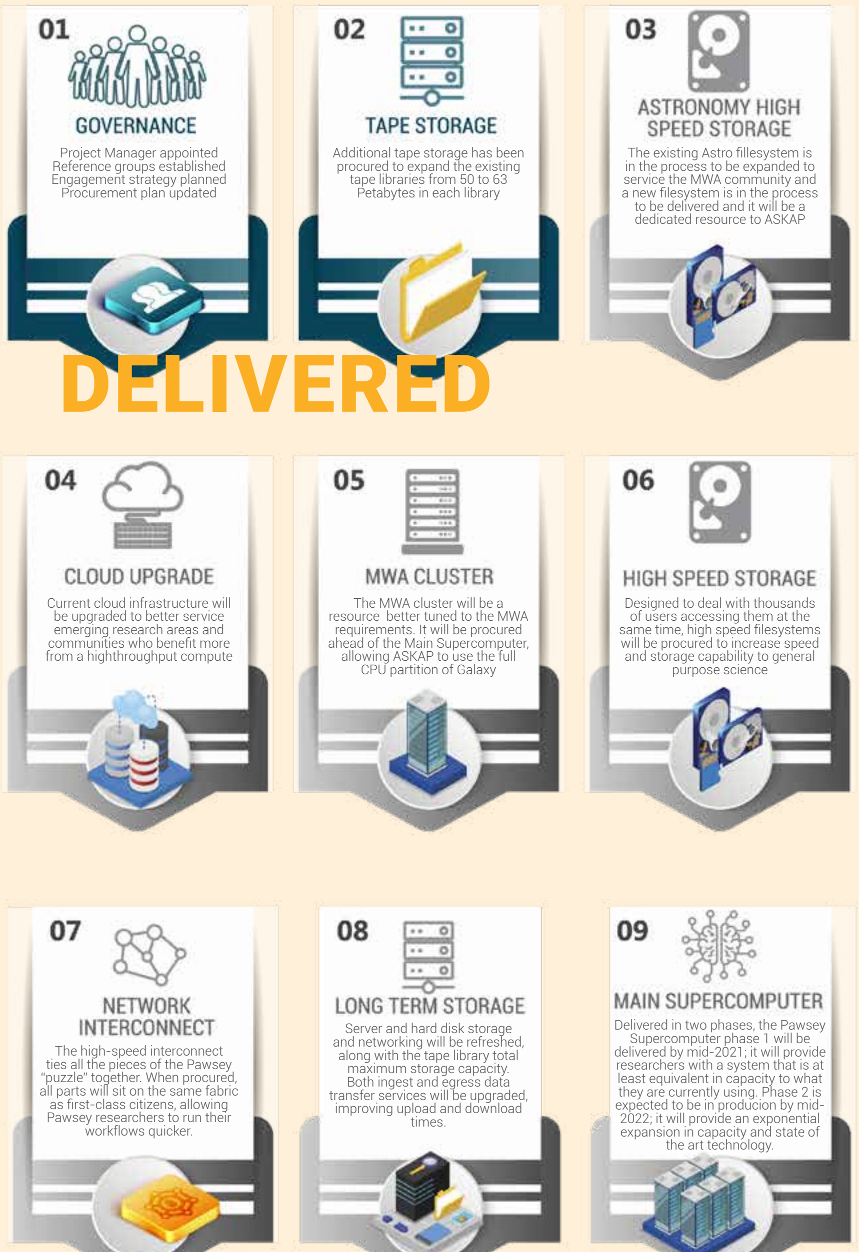

Through these reference groups and regular user forums, Pawsey's researcher and user community and the broader Australian research community are informing the upgrade process, to deliver supercomputing systems that meet their functional requirements and future needs.

A communications strategy for the capital refresh has also been established, with program communications including regular updates to stakeholders as a key requirement of the strategy. These updates will be via news stories published both in the media and on the Pawsey website, direct emails to users, quarterly podcasts and continued engagement through user forums.

Four procurements are complete or underway, sourcing high-speed storage, long-term storage and ancillary systems for a total of $8.5 million.

The first procurement completed was additional tape storage to expand the existing tape libraries. User requirements for the overall data management have been endorsed by the User Reference Group and Technical Requirements are being developed for submission to the Technical Reference Group in Q4 2019.

The next procurements to be delivered are the procurement of an expansion to the existing Astronomy Filesystem and a new Buffer Filesystem. These procurements will triple the temporary storage available to the Murchison Widefield Array (MWA) and the Australian Square Kilometre Array Pathfinder (ASKAP) as well as significantly increase the temporary storage performance and reliability. Dedicating the Astronomy Filesystem to MWA and the Buffer Filesystem to ASKAP will minimise the input/output performance impacts of both groups on each other, as each group will have their own completely separate filesystem. These procurements are well progressed with final acceptance testing occurring in August 2019.

A further filesystem upgrade is being progressed at present with the procurement of a replacement for the existing Scratch filesystem. Delivery of the new Scratch filesystem to Pawsey users will occur in Q1 2020.

Also progressing at present is the procurement of a new Cloud Compute System. Delivery of the new Cloud Compute to Pawsey users is to occur in Q4 2019. The development of the Statement of Requirements for the Pawsey Supercomputing System (the main procurement of the refresh) is well developed and is expected to go to market by the end of Q1 2020.

Key procurements to be considered over the coming months include the MWA Compute Cluster and the Pawsey Supercomputing System, the main procurement for the upgrade.

Until recently, MWA and ASKAP have both been able to share the Galaxy supercomputer. However, the now increased demands of processing from the ASKAP precursor telescope, and increased demand from MWA means that Galaxy is no longer able to support both teams within the remaining timeframe of the overall Pawsey Refresh. The procurement of the MWA Compute Cluster will provide MWA access to the latest generation technologies tailored for their workflows and provide Pawsey a platform to build new services that will be key to the success of the Pawsey Supercomputing System.

The procurement of a Compute Cluster dedicated to MWA is well progressed with delivery of the system to occur in Q1 2020.

The Pawsey Supercomputing System will replace the functionality of the Magnus, Galaxy and Zeus supercomputers in a single system. This will result in one large-scale, balanced High Performance Computing (HPC) system to support the broad scope of research users at Pawsey and the operational requirements of both ASKAP and the MWA.

The upgrade is being staged over time to ensure continuity of Pawsey services for the national research community, ASKAP and MWA. Installation and acceptance testing of phase 1 of the supercomputer is anticipated to be complete in Q2 2021. Phase 2 will be delivered following user migration from Magnus and Galaxy to the new system.

With the upgrade, Pawsey will allow Australia's researchers to remain globally competitive, upscale their research ambitions and accelerate scientific discovery, and continue to support the growing demand for supercomputer capacity from ASKAP and MWA.

SYSTEMS AND SERVICES

Allocation in numbers

Over 500 million
hours were requested

256 million
core hours were allocated

77 projects
were granted a NCMAS allocation

61 projects
accessed Magnus via Partner Allocation

37 projects
received an allocation via Energy and Resources scheme

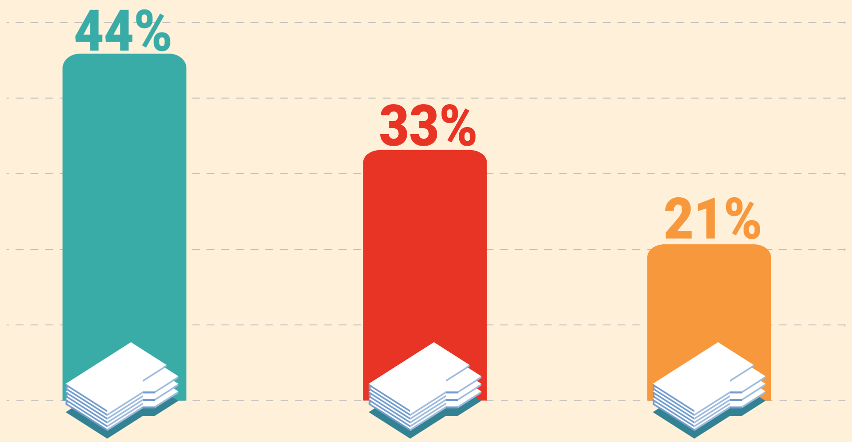

ACCESS TO PAWSEY SUPERCOMPUTERS

Pawsey provides supercomputing access to researchers predominantly via competitive merit allocation schemes. Through these merit allocation schemes, the Centre strives to:

- maximise the overall research impact of Pawsey supercomputing resources
- promote scientific advantage in priority domains such as radio astronomy and geosciences
- provide leading-edge supercomputing resources for researchers in Pawsey partner institutions
- enable wider adoption of, and benefit from supercomputing across Australia.

The Galaxy supercomputer is only available for radio astronomy-focused operations. Specifically, it is used to support ASKAP and MWA operations, which are two of the Square Kilometre Array (SKA) precursor projects currently under way in the north-west of Western Australia. For ASKAP, Galaxy acts as a real-time computer and is a part of the telescope, allowing direct processing of data delivered to the Pawsey Centre from the Murchison Radio Observatory.

RESOURCE ALLOCATIONS

Over 500 million core hours were requested on the supercomputer Magnus during the reporting period, almost twice as many hours as were available. Two hundred and fifty-six million core hours were allocated.

The competitive merit allocation process in 2018 continued to underpin Pawsey's commitment to high-end science. As part of the 2019 allocation released in January 2019, 177 projects received an allocation on Pawsey supercomputers and 26 new projects were activated on Magnus.

Supercomputer Allocation Scheme

THE NATIONAL COMPUTATIONAL MERIT ALLOCATION SCHEME (NCMAS)

Twenty-five per cent of the supercomputer time available at Pawsey is allocated through this scheme, along with compute-time allocations from other Australian HPC centres including the National Computing Infrastructure (NCI), the University of Queensland, and Monash University. The scheme operates annual allocation calls, which are open to the Australian research community and which provide substantial amounts of compute time for meritorious, computational-research projects.

THE PAWSEY ENERGY AND RESOURCES MERIT ALLOCATION SCHEME

Fifteen percent of Pawsey's total supercomputing resource is allocated through this scheme. Open to the Australian energy and resources research community, it provides significant amounts of compute time for meritorious research projects in these domains. Energy examples include generation, storage and distribution, while resources examples include exploration, minerals extraction and processing, groundwater, and waste management. This scheme operates annual calls, with an out-of-session application process for newly eligible project leaders.

THE PAWSEY PARTNER MERIT ALLOCATION SCHEME

Thirty percent of the whole supercomputer time available at the Centre is allocated through this scheme. Calls are open to researchers in Pawsey Partner institutions - CSIRO, Curtin University, Edith Cowan University, Murdoch University and The University of Western Australia - and provide significant amounts of compute time for meritorious computational research projects. This scheme operates annual calls, with an out-of-session application process for newly eligible project leaders.

THE PAWSEY DIRECTOR'S ALLOCATION SCHEME

Five per cent of the compute time available may be allocated to short-term projects requiring modest amounts of compute time. The scheme is open to all Australian researchers and is intended to support researchers with less experience of supercomputer-scale computational research (start-up grants) and to support speculative investigations designed to assess the potential of supercomputing for a particular application. This scheme is also used for commercial access.

Galaxy supercomputer (25 percent of Pawsey's total HPC resources) is fully allocated to support ASKAP and MWA operations.

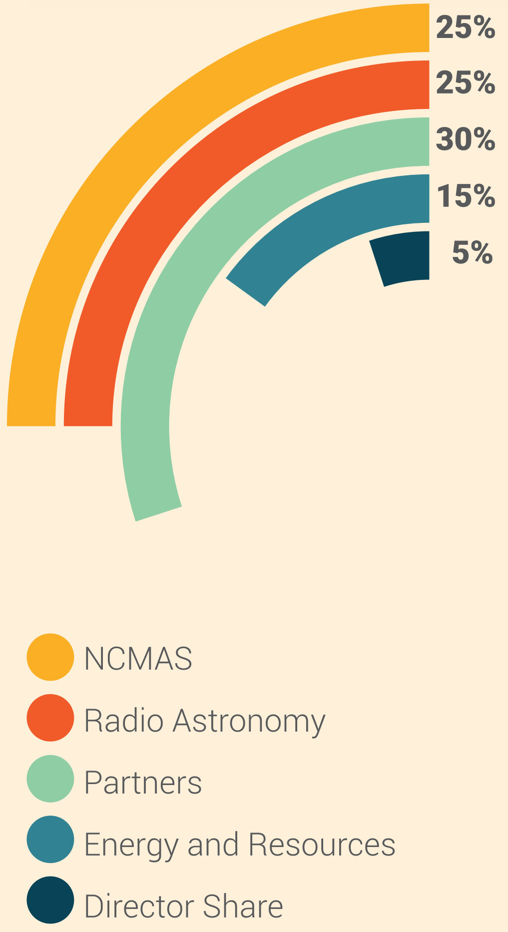

A National Facility

The Pawsey Supercomputing Centre supports researchers at universities across the nation. As part of the NCMAS, the two national Tier 1 centres, Pawsey and NCI, with Monash University and the University of Queensland contribute a number of core hours from their systems to service the national research community via meritorious allocations. Pawsey provides 100 million core hours to NCMAS for allocation to projects every year.

During the reporting period the top three institutions benefiting from those resources were Monash University, followed by the University of Melbourne and the University of New South Wales.

National access is also provided via the Energy and Resources Merit Allocation Scheme and the Pawsey Director's Allocation Scheme.

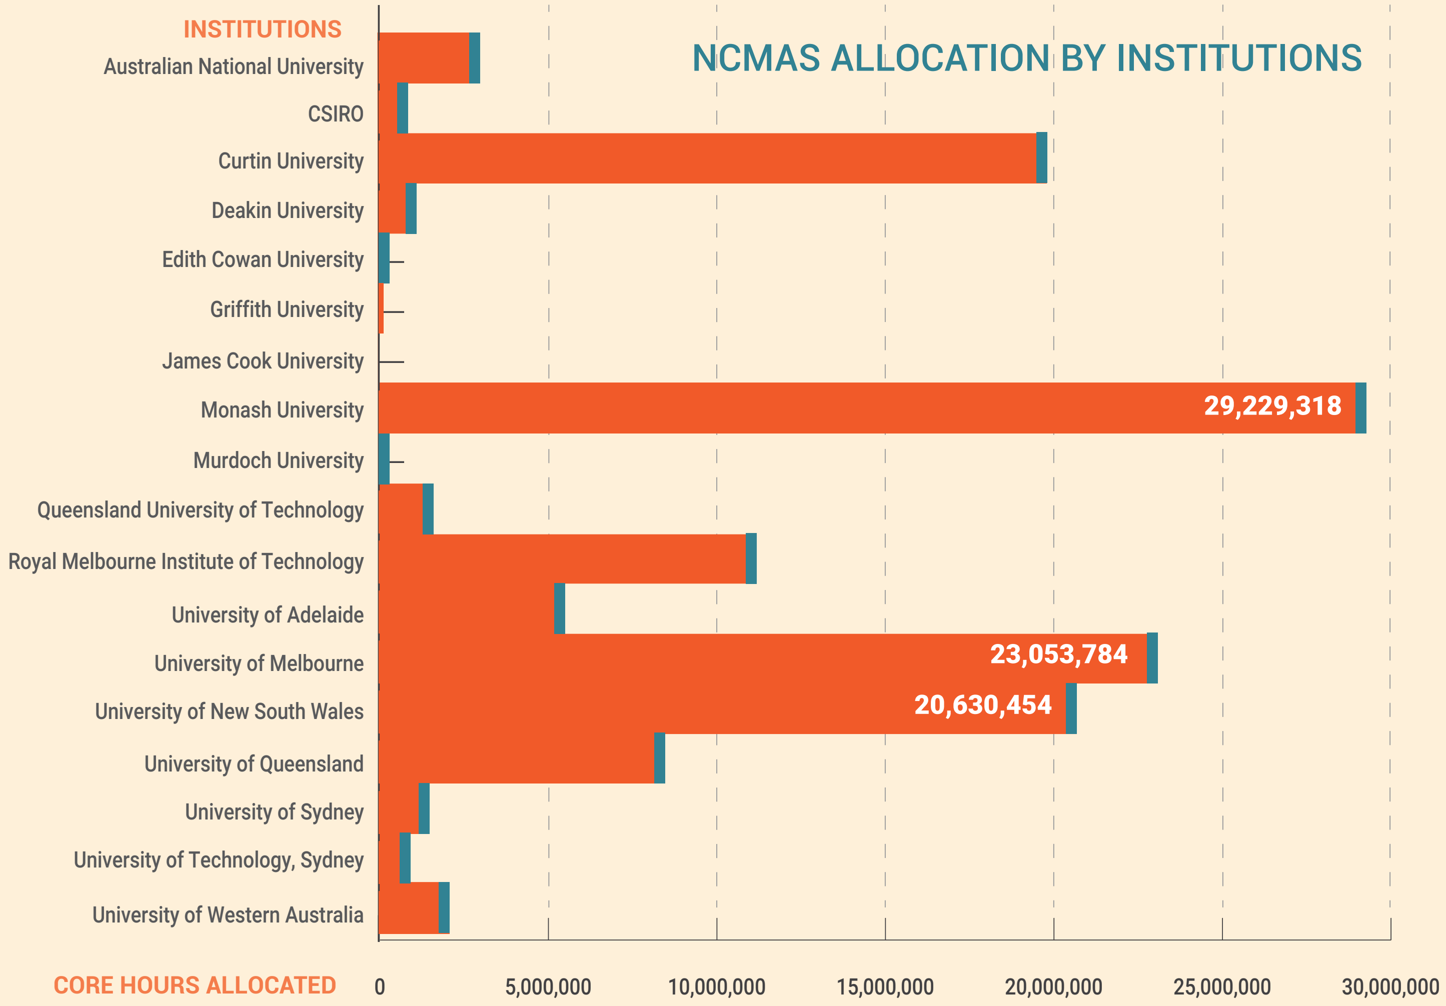

Supporting Australian Science Priorities

The Centre supports the National and State Science Priorities. Among the 177 projects that Pawsey supercomputers supported this year, Magnus saw a massive uptake from engineering, accounting for 38.7 percent of its use, followed by 21.9 per cent for chemical science projects.

Since Galaxy is Pawsey's real-time radio astronomy supercomputer and it is physically connected to the telescopes in the Murchison its usage was 100 per cent by radio astronomy.

As the first step toward the procurement of new infrastructure, the radio astronomy community required an increase in capacity. To ensure the speed and capacity required by astronomers will be available to carry out vital research projects and enable them to continue their search for the unknown more efficiently, the expansion of the Astronomy filesystem for the MWA and the purchase of a Buffer filesystem for ASKAP are underway.

Successful vendors to expand the Astronomy filesystem and procure a new filesystem named Buffer have now been selected, with delivery and installation scheduled for next financial year. When these are available it is expected that each project will benefit with three times as much storage and three times as much performance.

Zeus, Pawsey's Graphics Processing Unit (GPU) stepping stone, was predominantly used to support medical and health sciences with over 20 million core hours (over 45 per cent) used for that purpose. Engineering projects also accounted for 37 per cent of the total usage on Zeus, with over 16 million core hours.

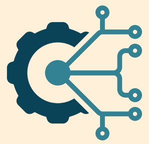

engineering
117,222,604
core hours used in Magnus

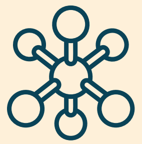

chemistry
66,299,446
core hours used in Magnus

radio astronomy
51,094,342
core hours used in Galaxy

health
20,543,909
core hours used in Zeus

SUPERCOMPUTER SPECIFICATIONS

GALAXY

- Peak performance in excess of 200 TeraFLOPS
- Real-time system for the Square Kilometre Array (SKA) pathfinders and radio astronomy projects
- Cray XC30 system with over 9,000 Intel processor cores
- Over 30 TB of memory
- 64 NVIDIA Kepler K20X GPUs.

MAGNUS

- World-class supercomputer in excess of 1 PetaFLOPs
- Cray XC40 featuring the 72 Gb/s Cray Aries interconnect
- Over 35,000 Intel 'Haswell' processor cores
- 3 PB of scratch file storage
- Over 95 TB of memory.

ZEUS

- Manufactured by HPE/SGI (USA)
- 92 nodes with 28 'Broadwell' Intel Xeon cores with a 100 Gb/s Omni-Path interconnect and 128 GB of memory per compute node
- 80 nodes with 64 Intel Xeon Phi 7210 cores with a 100 Gb/s Omni-Path interconnect and 96GB of memory per compute node
- 11 nodes with four NVIDIA Tesla P100 SXM2 GPUs with a 100 Gb/s InfiniBand interconnect and 256 GB of memory per compute node
- Six nodes with 16 'Broadwell' Intel Xeon cores and 1 TB of memory per compute node
- SUSE Linux (SLES) 12 operating system.

Increase in GPUs Demand and Other Technology Updates

2019 GPU Hackathon

GPUs offer exceptionally high memory bandwidth which accelerates or increases efficiencies of data processing using parallel codes. GPU acceleration paired with thousands of cores is enabling scientists and researchers to achieve performance increases of up to ten times in their computations.

During this period, Pawsey decommissioned **Athena**, the Advance Technology Cluster, as a separate system and its nodes were re-deployed to Zeus to reduce the amount of operational debt required to look after two systems. One of the outcomes from the projects running on Athena was the increasing usefulness of **GPU resources**.

Zeus, the HPE/SGI Linux cluster that supports pre- and post-processing of data, throughput workflows and simulations was originally purchased with 'Kepler' class NVIDIA GPUs for general purpose computing and Quadro CPUs for remote visualisation. With the arrival of Athena in 2017 the 'Pascal' class NVDIA GPUs demonstrated how much GPU computing power has increased in the last five years.

During the reporting period, a procurement for a new commodity Linux cluster started, to support the increased demand for GPUs, especially in the field of machine learning. The new system will replace the original 42 GPU nodes in Zeus which were purchased as part of the initial Pawsey project. The new cluster will be made up of 42 nodes, 22 with dual NVIDIA V100 GPUs for computational work, machine learning workflows and data analytics, and 20 state-of-the-art Quadro NVIDIA GPUs for **remote visualisation** with the deployment of FastX.

Zythos, the large memory SGI UV2000, was retired. With six high memory nodes part of the Zeus cluster, Zythos was no longer required.

Support for **container technology** was set up on all Pawsey systems, i.e. Magnus, Galaxy, Zeus, and Nimbus. This is enabling researchers to improve the reproducibility and portability of their workflows, allowing them to develop them on local workstations, and then to seamlessly deploy them on Pawsey computing infrastructure. The largest positive impact is currently being made in the life and health sciences, as well as in deep-learning intensive research projects.

DATA STORAGE INFRASTRUCTURE

As part of the data services provided at the Centre, Pawsey hosts a series of significant data research collections from national universities and government departments across Australia. In addition to housing in excess of 40 Petabytes of data storage resources, Pawsey's research collection service provides data sharing facilities to support scientific collaboration among researchers nationally and internationally.

During the financial year, 64 projects were granted over 13 TB of long-term storage. Although most projects are radio astronomy related, Pawsey also supports another 17 organisations, including Landgate, Telethon Kids Institute, Western Australian Department of Transport, and Department of Health.

NIMBUS CLOUD

Nimbus, Pawsey's cloud service, has seen significantly increased uptake during the reporting period. In the previous financial year, 89 per cent of the cores available in Nimbus were allocated. In this financial year, 6,801 cores were allocated on the service. This is above the normal service limit of the existing system with all GPU nodes allocated. Over 60 different projects benefited from Pawsey's cloud services during the reporting period.

The cloud has been used as a stepping stone for researchers who find their computing capacity is not enough but have not identified a need for supercomputers or don't know how to use them. In some cases, Nimbus has been used by researchers as a resource additional to Magnus. More recently, Nimbus uptake from bioinformaticians has increased as the GPU technology included as part of the system has been demonstrated as extremely valuable to running their workflows.

Nimbus is about to be expanded as part of the Pawsey Capital Refresh project, with the technical specifications been completed during the reporting period.

Ceph, a flexible, reliable and easy to manage object-based storage for HPC is currently being tested with the support of CERN, the European Organisation for Nuclear Research, as a possible host for HPC workloads.

VISUALISATION

Pawsey's data visualisation services help researchers gather real-time insights from their data and allow the data to be more easily interpreted by translating it into more meaningful or recognisable forms.

A new web-based remote visualisation service went into production during the reporting period. The team is using the large-scale client-server ParaView setup for remote visualisation. There are also plans to provide novel display technologies by upgrading current systems.

A new visualisation resources booking system integrated with high-end Windows workstations was also launched to improve accessibility for users. During the reporting period 5,928 hours of visualisation resources were booked through the new system.

NETWORK

Pawsey's network provides high-speed connectivity both within Pawsey itself, and with Pawsey's research partners in Western Australia through peerings with CSIRO, the University of Western Australia and Curtin University. Pawsey also peers with AARNet, providing access to the NCI and other national research institutions, and international research and education networks. Pawsey has multiple high-speed connections to the Murchison Radio-astronomy Observatory, allowing the ingest, processing and storage of huge volumes of radio astronomy data.

During the reporting period the Centre added a direct 100 Gbps link to the Curtin Institute of Radio Astronomy (CIRA), and some multi-rate equipment to support the expansion of a variety of compute and storage infrastructure which occurred during the year.

OTHER SYSTEM SPECIFICATIONS

TAPE LIBRARIES

- Up to 100 PB of storage
- Two duplicate libraries for added resilience
- Connected at up to 40 GB/s.

NIMBUS CLOUD

- 3,000 cores and 16 TB of RAM across 46 compute nodes
- 300 TB (usable) of volume storage
- 6 x HPE SX40 nodes, each node has 2 x NVIDIA V100 GPUs.

VISUALISATION WORKSTATIONS

- Two workstations, Intel(R) Xeon(R) CPU E5-2697 V2 @ 2.70 GHz
- Two processors
- 256 GB RAM
- Two GPU components; Quadro K5000 and Tesla K20c
- 512 GB SSD.

NETWORK

50 devices supporting in excess of 6,000 ports

- ~1,400 non-blocking 10 Gbps ports
- ~70 non-blocking 100 Gbps ports
- The remainder are 1 Gbps ports, mainly used for management and monitoring purposes, and staff access.

Extending the reach of Australia's HPC expertise and its scientific endeavours

A World of Difference

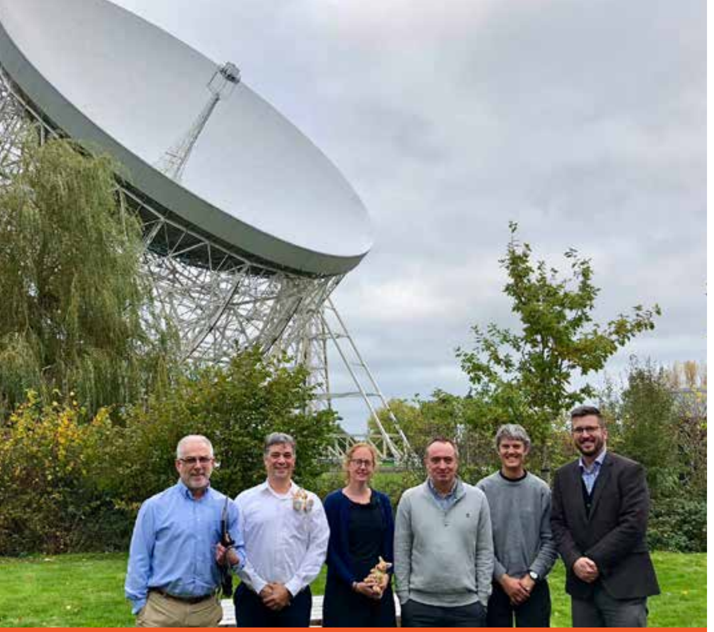

Connecting the dots

The Pawsey Supercomputing Centre is unique in the Australian scientific landscape as an unincorporated joint venture that brings together the Commonwealth Government, Western Australian Government, university partners - Curtin University, Edith Cowan University, Murdoch University and The University of Western Australia - CSIRO and collaborating organisations in a consortium that has been steadily producing outcomes for more than fifteen years.

By bringing researchers, industry, supercomputers and technical support together, the Centre provides access to infrastructure, expertise and opportunities at a scale previously unachievable. Collaboration at the national and international level is crucial to maximise Pawsey's utility within the national research infrastructure.

INTERNATIONAL COLLABORATIONS

Pawsey's HPC future is hyper-connected. The Centre collaborates at scale across geographic boundaries to tackle issues that affect people worldwide. The physical ability for Pawsey to be able to do so is being reinforced by Project Indigo, the new high-speed optical fibre nearing completion between Singapore, Jakarta, Perth and Sydney. The subsea cable system includes consortium members from Singapore, Indonesia and Australia, and will strengthen links between Australia and south-east Asia, providing high-speed broadband with low latency and enhanced reliability.

This world-class data connectivity will provide critical infrastructure and greater flexibility as the Asian zone of HPC activity continues to develop, supporting Pawsey's growth in collaborative education, research, development and the digital economy.

Pawsey is building on the Memorandum of Understanding signed in 2018 with the National Computing Centre (NSCC) in Singapore to collaborate in the fields of supercomputing, networking, data analytics, scientific software applications and visualisation. Staff visits and discussions were held in both Perth and Singapore, the Centre co-organised SC Asia19 in Singapore and Mr Mark Stickells was invited to present at the conference. As part of SC Asia19, Pawsey supported two Australian researchers from the fields of quantum computing and astronomy to showcase how their projects have benefited from High Performance Computing (HPC) resources. Pawsey is taking a leading role in supporting HPC in Asia and strengthening engagement across the region including proposing conferences and workshops, and joint access proposals.

In 2018, Pawsey was the first Centre outside of Europe to sign a Memorandum of Understanding with the Partnership for Advanced Computing in Europe (PRACE) to promote research and development in computational science and engineering, and share activities around access schemes, training and exchange schemes. There are ongoing discussions with PRACE regarding how best to drive the collaboration forward. Pawsey hosted delegates from Europe during Western Australia's European Business Week in 2019 and showcased collaborative opportunities between Europe and Western Australia/Australia.

Pawsey staff also visited CERN, the European Organisation for Nuclear Research, and started work on Pawsey–CERN collaboration plans. After attending the CERN OpenLab meeting and meeting with senior CERN staff, collaboration around training and internships was discussed. Plans were also developed to build collaborations around advanced storage technology - dynamic file system deployment with Ceph - with Pawsey and the CERN staff.

Pawsey continues to participate in international large-scale data projects, such as the earth observation Copernicus program run through the European Space Agency. With Landgate and Geoscience Australia, Pawsey is providing access to Copernicus data for Australia.

In the UK, Mark Stickells' visit to the Edinburgh Parallel Computing Centre (EPCC) and Hartree Centre provided excellent insights into security protocols, industry engagement and scaling their workforce to maximise impact. Discussions were also held with the Jodrell Bank Discovery Centre, highlighting future data processing and service opportunities to support the Square Kilometre Array (SKA).

In North America, Pawsey's 2018 Letter of Intent with the National Energy Research Scientific Computing Centre (NERSC) in the US is allowing the Centre to further develop its collaborative links. Pawsey also hosted a senior technical delegation from Chevron Houston who engaged Pawsey's data innovation and services team, and toured the facilities while exploring collaboration opportunities.

Pawsey is also involved in a collaboration with the WA Department of Health, University of Ontario Institute of Technology and University of Technology Sydney to provide hospital-based informatics. As a result of this collaboration, the Centre hosted a PhD student from Canada as Pawsey's inaugural HPC research fellow, who worked on the Artemis project for six months while gaining insight into the HPC landscape in Australia and helping move the project one step closer to implementation.

Mark Gray and Ugo Varetto are assisting Compute Ontario, serving on the Industry Engagement Committee, an advisory committee to its Board.

Pawsey representatives visited The New Zealand eScience Infrastructure (NeSI) and the National Institute of Water and Atmospheric Research (NIWA) while at eResearch NZ 2019. Opportunities for collaboration were identified around training and skills development and OpenStack infrastructure deployment for virtual service management.

These connections and partnerships underpin Pawsey's efforts to deliver better, faster and more innovative scientific outcomes to benefit not just the nation, but the global community.

Left: *SKA team welcomed Mark Stickells during his visit to their headquarters at Jodrell Bank Observatory, in the UK.*
Right*: Singapore's Finance Minister Heng Swee Keat visiting Pawsey's booth at Supercomputing Asia (SCA) 2019.*

Pawsey Board visited NCI and hosted a joint board session

NATIONAL

Pawsey continues to work closely with the National Computational Infrastructure (NCI) to increase coordination in the delivery of services to researchers across Australia through national HPC and data infrastructure. A current project is scoping a common data access interface and shared cloud environments, with provision for shared authorisation and authentication. More broadly, Pawsey is working with the Australian Access Federation (AAF) and the NCI to scope ways in which authentication and authorisation can be harmonised in peak national facilities.

Other initiatives include sharing staff and training resources to enhance knowledge and skills across both teams and the researchers using the facilities, and developing a nationally strategic petascale dataset for improved data management of national collections. The NCI and Pawsey are also working closely to replicate and back-up data between the two facilities.

Pawsey continues to provide data infrastructure and remains the largest provider of research data collections in Australia through its work with the Australian Research Data Commons (ARDC). In collaboration with ARDC, AAF, AARNet and the NCI as the 'Digital Data eResearch Platforms (DDeRP) group', the Centre is developing a framework of how to best facilitate service provision for researchers.

Pawsey is also:

- Working closely, alongside NCI, with members of the Integrated Marine Observing System (IMOS) community to support numerical ocean modelling
- Part of the consortium of organisations involved in developing the Murdoch University-led Australian National Phenome Centre (ANPC), and will be providing data services and infrastructure to support healthcare research
- Providing infrastructure support on BioPlatforms Australia's 'Australian BioCommons Pathfinder Project' which looks to enhance the ability of Australian research to understand the molecular basis of life across all research domains.

Pawsey and the SKA

Pawsey has continued to support Australia's international commitments to the SKA project through the ongoing ingest, processing, storage and provision of data for the radio astronomy community. Activities have related in particular to supporting the SKA precursor telescopes; the Murchison Widefield Array (MWA) and the Australian Square Kilometre Array Pathfinder (ASKAP), and towards delivering the SKA.

Pawsey, together with the International Centre for Radio Astronomy Research (ICRAR), CSIRO and SKA Australia were involved in the five-year design of the Science Data Processor (SDP) that will process the data collected by the SKA. The design was finalised in May, with SDP functions to be split across two supercomputers: one in Perth for SKA low frequency, and one in Cape Town, South Africa for SKA mid frequency.

The combined power of the SDP is expected to produce the world's fastest supercomputer, with a total compute power of around 250 PFlops.

Pawsey in particular developed the local infrastructure interface specification between the SDP and the hosting data centres, as well as porting and testing automated scaling test frameworks on the Centre's supercomputers.

Throughout the year, Pawsey compute, storage and network infrastructure provided to the SKA precursors has exceeded 90 per cent availability. As part of the Capital Refresh project, the long-term tape storage upgrade was completed to increase storage capacity for radio astronomy projects and other data collections.

Pawsey continues to support the development of the MWA All Sky Virtual Observatory (ASVO) with a dedicated staff member to support the MWA/AVSO portal and astronomy-related projects.

The Centre has continued to grow the development of expertise in the Australian astronomy community through its continued involvement in the Astronomy Data and Compute Services (ADACS) project, funded by Astronomy Australia Limited. In collaboration with Curtin University and Swinburne University of Technology, Pawsey provides astronomy data analysis services and training to astronomy researchers nationally, growing the level of expertise of researchers in this domain to be able to effectively use Pawsey infrastructure and services.

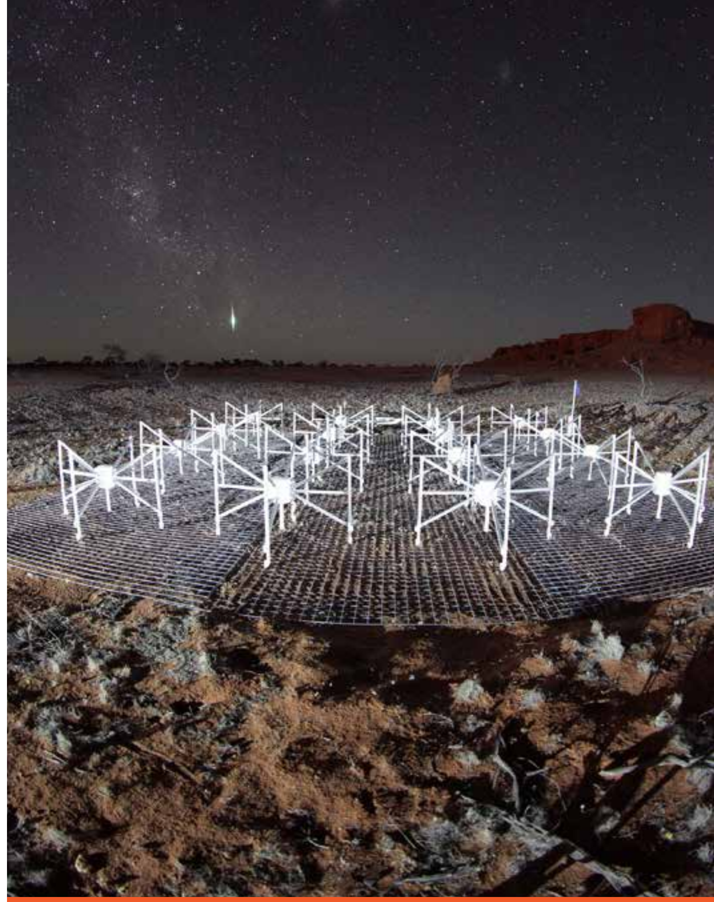

Night Sky at MWA Site Western Australia - Photo credit - Copyright John Goldsmith

Growing expertise

The Pawsey Supercomputing Centre is working to build a critical mass of advanced computing knowledge in the research community. Its unique range of programs continues to grow the expertise of the current Pawsey community, as well as the next generation of supercomputing specialists, ensuring the creation of a skilled workforce across Australia.

Pawsey Summer Internship poster presentation jointly hosted with the first 2019 Pawsey Friday

RESEARCHER DEVELOPMENT

Pawsey focuses significant effort on growing and diversifying the talent pool supporting the deployment of HPC in Australia, investing in the development of the next generation of data scientists, data-literate subject experts, and data-literate 'enablers' who can bring industry, researchers and HPC experts together.

The Centre's focus on on-site and online training continues to build, so the Pawsey research community can effectively use and maximise their benefit from Pawsey resources. Pawsey's program of training has been developed to support users across all levels of experience, and courses are provided nationally to researchers from academia, industry and government. During the reporting period Pawsey staff trained almost 700 people. Training occurred in Perth, Adelaide, Melbourne, and Sydney and has included:

- introductory and intermediate courses on supercomputing
- advanced sessions on MPI, OpenMP and GPU programming
- introductory courses on Unix and pshell
- Nimbus research cloud training
- containerising workflows
- remote visualisation courses.

To supplement these training sessions, the Centre continued to grow its open repository of technical reference materials covering different aspects of HPC systems usage, parallel programming techniques, and cloud and data resources, available via the Pawsey Confluence webpage. The Centre has also launched a collection of training and researcher highlight videos through its YouTube channel.

694
researchers were trained

195
people participated on webinars

30
people joined the GPU Hackathon

After last year's success, the number of webinars has grown in frequency and attendance. Four online webinars were held during the reporting period. Topics included Containers in HPC, GPUs and Python. These training sessions represented almost 30 per cent of total training participation for the period.

During the second edition of the GPU Hackathon, researchers and industry participants accelerated their work by nearly 1000 times after a week of fast-track learning process that helps them optimise their code using GPUs. This was achieved with the support of Pawsey supercomputing experts and internationally renowned mentors.

The GPU hackathon is a 5-day coding event in which teams of developers port their applications to run on GPUs, or optimize their applications. This Hackathon is jointly organised with Oak Ridge National Laboratories and NVIDIA.

The event featured 5 teams and more than 30 participants representing academia, industry, and government. Teams included Let's Get Physical from Curtin University, HiPSTAR EXPRESS from the University of Melbourne, Team ASKAP from the CSIRO, Multi Party Time! from Data61 and Perth IT company Manufacturing Intelligence.

Information sessions and presentations are also regularly held at Pawsey by visiting researchers and computational experts. During the period, presenters from NERSC, University of Trier, University of Warsaw, and the High-Performance Computing Center Stuttgart were featured together with researchers from partner universities.

Twelve students participated in Pawsey's Summer Internship program, spending 10 weeks delving deeper into their scientific areas using HPC, data analytics and visualisation. Their projects covered condition monitoring of assets in the resources industry, particle simulation code, atomic and molecular simulation, next-generation sequencing, eDNA analysis, image processing analysis, malaria, dengue and zika disease control, renewable wave energy, machine learning, genetic signatures of complex diseases and quantum combinatorial optimisation.

PAWSEY TEAM DEVELOPMENT

To support the retention of Pawsey staff and to enable them in providing the best supercomputing, cloud compute, visualisation and data management service and expertise to the research community, Pawsey has made a concerted effort to ensure staff have access to training and education to hone existing skills and develop new skills. To this end, an Education and Training Manager has been appointed to develop a training strategy and to manage Pawsey's overall education and training program. Internal training and skill building undertaken during the year has included high-performance storage systems training for HPC administrators, software carpentry training for uptake staff, and technical training around specialised team and individual needs. Discussions are also underway with NCI to develop a program of joint training activities.

Teams across Pawsey are proactive in identifying and investigating emerging technologies for HPC and big data, with vendor presentations and information sessions being held throughout the year. Where emerging technologies are adopted, training is developed and presented to Pawsey staff, researchers and the wider community to ensure that users understand the new technology.

Maciej Cytowski, Pawsey Supercomputing Applications Specialist presenting on Containers in HPC from Pawsey/NCI joint booth at SC18 Dallas

Minister for Water; Fisheries; Forestry; Innovation and ICT; Science, Hon. Minister Kelly started his day and National Science Week celebrations at the Pawsey Open Day. He teamed up with attendees on the beanbag-sorting competition, a game to help understand the concept of parallel computing.

Engagement and Outreach

To encourage greater awareness of the supercomputing, large-scale data and scientific visualisation capabilities available within Western Australia, and highlight the possibilities available for researchers and industry to harness Pawsey resources and expertise, the Centre actively engages with potential users, collaborators and partners.

Each year Pawsey attends and/or presents at a number of international events to increase the exposure of the work Pawsey is involved in, as well as promoting supercomputing, large-scale data and scientific visualisation capability in Australia more generally.

- At SC18, Dallas, Pawsey's joint booth with NCI attracted at least 700 visitors, while Pawsey staff presented a workshop on HPC education and training and contributed to three 'Birds of a Feathers' (BoFs).
- Pawsey co-organised SCAsia 2019, Singapore. Focusing on 'HPC Futures – Hyperscalers, Exa, AI, Quantum and Beyond', the team engaged with conference participants as an exhibitor, and presented at conference streams, workshops and keynotes to showcase Australia's scientific outcomes enabled by the power of supercomputing.
- At eResearch Australasia, Pawsey welcomed visitors to the Centre's joint booth with NCI, and contributed to two panel sessions, one focused on eResearch democratisation, and the other on engagement and culture. Pawsey also participated in a BoF discussion session on advancing HPC and data collaborations in Australasia, where an Australasia collaborative group initiative was announced.
- Pawsey staff also presented at CS3, Italy; ISC19, Germany; and in Australia at the 3rd International Whitefly Symposium, 21st Australasian Fluid Mechanics Conference, the Accelerated Computing for Innovation Conference, C3DIS 2019, and ResBaz Sydney 2019.

Various events and activities were organised and attended to connect Pawsey more closely with researchers, the science community, government and industry.

- Pawsey's second bi-annual Open Day was held during National Science Week on 11 August 2018, attracting over 500 attendees including Minister for Water Fisheries; Forestry; Innovation and ICT; Science, Hon. Dave Kelly. Visitors toured the Centre and attended information sessions and presentations by Pawsey researchers and industry representatives. Following the Open Day, the virtual tour of Pawsey was launched jointly with the Murchison Radio Observatory tour (https://tour.pawsey.org.au/).
- Tours of Pawsey facilities are held throughout the year, with over 300 visitors representing staff and researchers from other supercomputing centres, government agencies and scientific facilities visiting the Centre during the reporting period.
- Pawsey continued gathering and sharing Best Practices via newsletters with collaborator contributions from the EPCC, NERSC, NeSI and NCI. Researcher outcomes enabled by the Centre have also been widely shared through case studies and researcher profiles promoted via Pawsey's broader communications channels.
- Over the year, six Pawsey Fridays were held, spanning from Rottnest swim ocean forecasts and tropical cyclones predictions all the way to teaching schoolchildren about quantum mechanics and general relativity. Around 200 people participated in these sessions.

Russell Potapinski - General Manager and Head of Intelligent and Autonomous Systems at Woodside Energy during his presetation at HPC AI advisory Coucil Conference in Perth

- Pawsey staff travel the nation showcasing how HPC is impacting researcher outcomes. Across this period, Pawsey visited Melbourne, Sydney and Brisbane to engage with the research community more broadly, making them aware of the services Pawsey has to offer and how to access them.
- During the reporting period, face-to-face User Forums were held in Sydney, Adelaide, Brisbane and Perth. User Forums provide researchers with the opportunity to discuss their experiences and give invaluable feedback to the Centre. This has also been an opportunity to understand their future needs and workflows to inform the procurement. For the first time, attendees had the opportunity to join a forum remotely via WebEx and based on positive feedback, this experience will be repeated for future forums. The feedback and sharing resulted in over 60 issues being addressed in total and shared with the wider research community. Pawsey's end of year user satisfaction survey recorded that 81 percent of users had their expectations met or exceeded.
- Quarterly newsletters reached over 2,000 people throughout the year, including the International Best Practices newsletter focusing on 'User management' made with contributions from EPCC (Edinburgh), NERSC (US) and NCI. A second International Best Practices newsletter, focusing on 'Data-centric HPC', was made with contributions from the NCI and New Zealand eScience Infrastructure (NeSI). In addition to the newsletters and news stories released, project case studies and research profiles were produced and promoted via Pawsey's broader communications channels.
- Pawsey's new website was launched on 31 October 2018 to showcase and more easily navigate the modern Pawsey. Case studies and news articles are regularly added to the website highlighting the work of the researchers and other stakeholders who use the Centre's infrastructure and services. A comparison of FY17–18 vs FY18–19 shows a boost in Pawsey's website audience. Users increased almost 11 per cent over the latter period with 20,948 new users, compared to 18,914 in the former.

Pawsey hosted several events across a broad range of research disciplines and stakeholder groups.

- On 28–29 August 2018, Pawsey hosted the second international HPC–AI Advisory Council Conference in Fremantle. The conference focused on exploring major and emerging domains and insights of the industry, with a specific focus this year on artificial intelligence, data sciences, HPC, deep and machine learning, and visualisation.
- On 6–7 September 2018, Pawsey hosted the inaugural joint Pawsey-COMBINE Bioinformatics Symposium. The two-day symposium focused on aligning the bioinformatics researchers of today with the researchers of tomorrow and how their work can be supported by HPC. The event attracted 70 people from diverse backgrounds including academia, CSIRO, state departments and private organisations.
- Pawsey organised Data Science Week which for the first time was a national event, run in partnership with NCI and the Australasian eResearch Organisations (AeRO) Collaborative Conference on Computational and Data Intensive Science (C3DIS). The week had a focus on 'Women in Data Science', and included over 20 presentations and training sessions held in Perth and Canberra with over 700 participants. Pawsey also hosted the launch of the Western Australian Data Science Innovation Hub, and chaired the Cray User Group.
- STEM in schools played a special role in this period's engagement activities, with Pawsey providing support for 14 schools to take part in the Western Australian Information Technology and Telecommunications Alliance (WAITTA) Game Changer awards, a program that brings together students who are interested in STEM and innovation across Western Australia.
- Pawsey hosted its first two careers nights which attracted Year 10 students and their parents from schools across Perth. The Centre also ran a masterclass in 'Supercomputing and Data Science' targeted at school teachers, as part of its commitment to building STEM capabilities in the state and the nation.

IN TOTAL, PAWSEY STAFF ENGAGED WITH OVER 3,000 PEOPLE AT CONFERENCES, TOURS AND SPECIAL EVENTS IN THE LAST YEAR.

DIGITAL ENGAGEMENT GROWTH - NEW WEBSITE

Page views were boosted over 14 per cent during the period - increasing from 79,095 to 90,447

DIGITAL ENGAGEMENT GROWTH - YOUTUBE CHANNEL

2017

Impressions	Views ↓	Average view duration	Watch time (minutes)
–	643 100.0%	3:00	1,937 100.0%

2018

Impressions	Views ↓ ⚠	Average view duration	Watch time (minutes) ⚠
19,301	2,223 100.0%	2:17	5,096 100.0%

In a comparison with the last two financial years, engagement has nearly quadrupled: FY17–18 had 643 views for the year while FY18–19 amassed 2,233. The channel's watch time increased from 1,900 minutes in 2017 to over 5,000 minutes in 2018.

Users and collaborations across the globe

AMERICA

NERSC

HPC-AI ADVISORY COUNCIL
NETWORK OF EXPERTISE

EUROPE

PRACE

epcc

CERN openlab

ASIA

SG NSCC SINGAPORE National Supercomputing Centre

Researchers

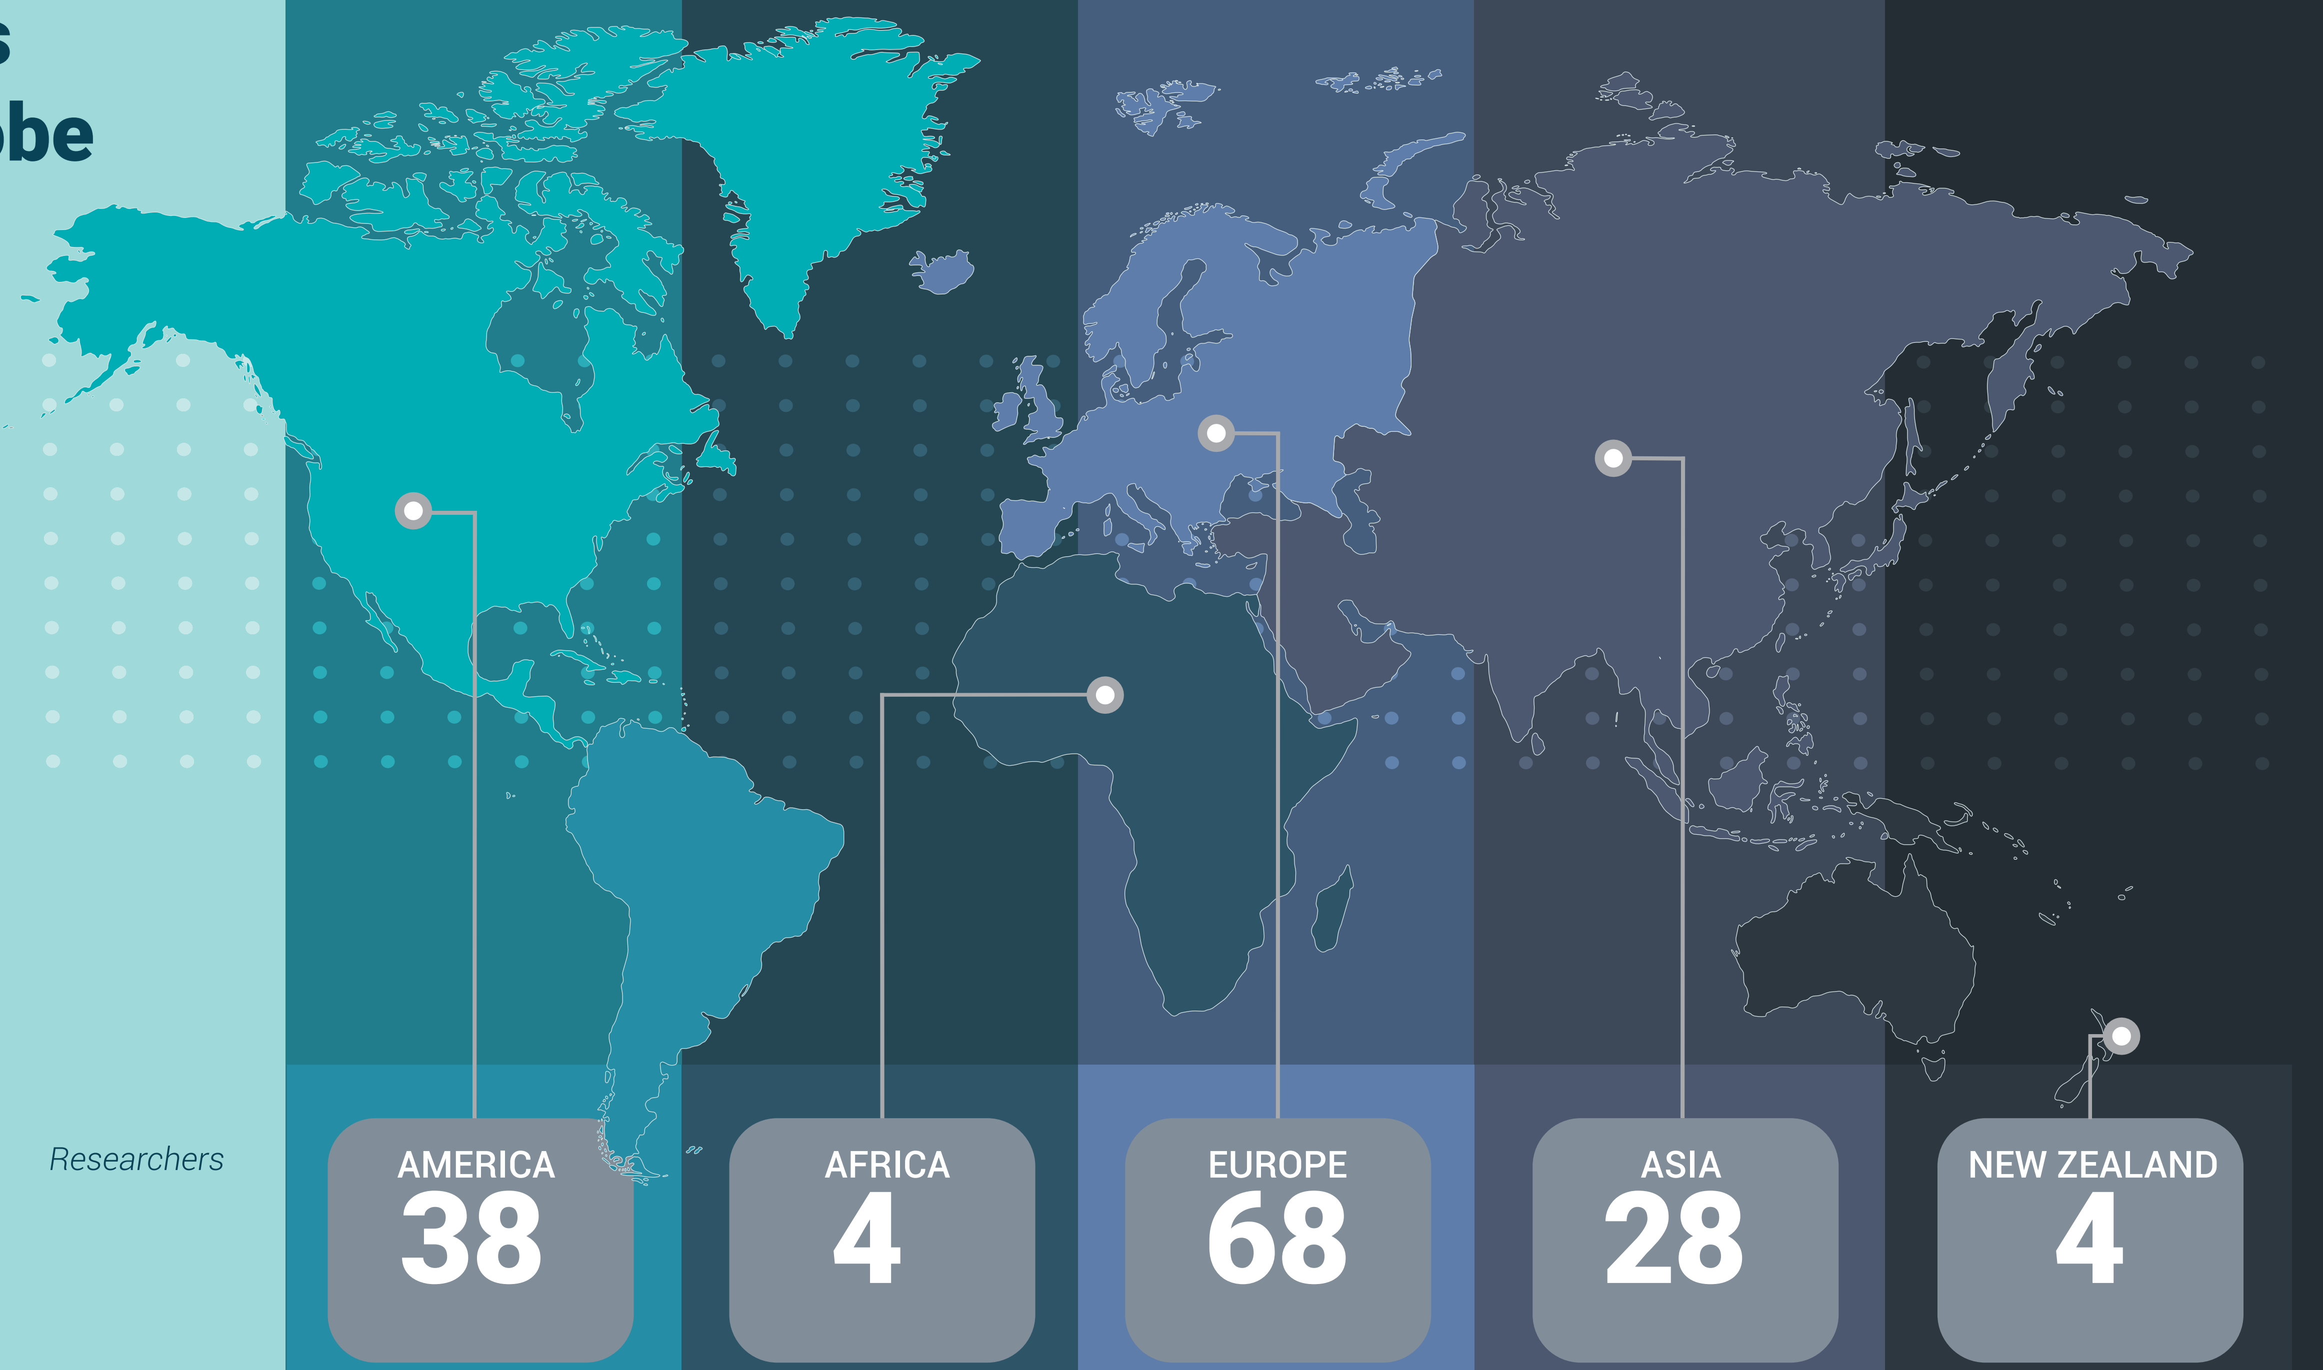

Financials

The Pawsey Supercomputing Centre is an unincorporated joint venture between Commonwealth Scientific and Industrial Research Organisation (CSIRO), Curtin University of Technology, Murdoch University, Edith Cowan University and The University of Western Australia. CSIRO as the lead agent, holds and manages Pawsey's assets and finances. Pawsey Centre is required to adhere to CSIRO's reporting, budgeting and auditing framework and requirements.

The Commonwealth Government provided the Centre with the initial capital investment for the construction of the Centre's building and High Performance Computing infrastructure. The Commonwealth Government continues to support the infrastructure maintenance and operation of Pawsey through the National Collaborative Research Infrastructure Strategy (NCRIS) program run by the Department of Education and Training.

The Western Australian Government along with the unincorporated joint venture partners, provided the other portion of the operational funding for the running and maintenance of the Centre.

Pawsey, through CSIRO, employs professionals and High Performance Computing experts to run operations of the facility and provide exceptional services on supercomputing access, cloud, data storage and visualisation, training and consulting. As a Tier-1 High Performance Computing facility in Australia, Pawsey's function is to accelerate scientific research for the benefit of the nation and the continuance of these funding support are essential for its existence.

Pawsey Supercomputing Centre
Statement of Income and Expenditure
01 July 2018 to 30 June 2019

	$'000
Revenue	
External Project Funding	11,412
Internal Joint Venture Partner Subscriptions and Recoverable	3,405
Total Revenue	**14,817**
Expenditure	
Asset Procurement (Capital Expenditure)	2,031
Labour staffing	6,835
Machine Maintenance Contracts	3,042
Operating costs	
Utility charges	1,580
Other Operating expenses	2,371
Total Expenditure	**15,859**
Surplus / (Deficit)	**(1,041)**

*Note: Deficit relates to carried forward activities funded in 17/18 and completed this period.

NCRIS Grant 41%
WA Government FAA 29%
Joint Venture Partner Subscription and Recoverable 23%
Subcontract Projects (ADACS-AAL & ARDC-Monash University) 5%
Other Revenue: Pawsey Services & Pawsey Capability Refresh (2%) 2%

Revenue %

Figure 1: Sources of Revenue

FINANCIAL REPORT 2018-19

The funding model for Pawsey aims to reflect the proportionate usage of machines on research projects with 23% of funds provided by the joint venture partners and 29% from the Western Australian State Government through the Financial Assistance Agreement. The Commonwealth Government provided the majority of other funding (41%) through the NCRIS program.

In FY18/19, the major cost driver for the Centre was labour, representing the staffing costs for technical experts and allied support services which is 43% of its total expenditure. Machine maintenance comes at 19% followed by other operating expenses at 15%. Utility costs are 10%, which includes electricity and water usages. The Centre receives electricity credits from the solar panels installed at the building. Asset procurement costs were 13% which represents the Capital Expenditure funded through NCRIS program for minor upgrades to the existing infrastructure.

The Pawsey Capability Refresh grant of $70 million from the Commonwealth Government funds the major upgrade to the supercomputing facility which kicked off during the first quarter of FY1819. This is financially managed and reported by Curtin University, a joint venture partner of Pawsey Supercomputing Centre. A strong governance process over the procurement has been implemented to involve all members of the unincorporated joint venture.

43% Labour staffing
19% Machine Maintenance Contracts
15% Other Operating expenses
13% Asset Procurement (Capital Expenditure)
10% Utility charges

Expenditure %

Figure 2: Expenditure Breakdown

Acknowledgment

The Pawsey Supercomputing Centre is supported by the Australian Government through a $70 million grant made under the Industry Research and Development Act and administered by the Department of Industry, Innovation and Science. Pawsey is also supported by the Australian Government under the National Collaborative Research Infrastructure Strategy (NCRIS) and related programs through the Department of Education. The Centre would also like to acknowledge the support provided by the Western Australian Government and its Partner organisations.

NCRIS National Research Infrastructure for Australia
An Australian Government Initiative

GOVERNMENT OF WESTERN AUSTRALIA

CSIRO Curtin University ECU Edith Cowan University Australia Murdoch University The University of Western Australia

pawsey.org.au