Pawsey Supercomputing Centre
Annual Report 2017-18

A WORLD OF DIFFERENCE

Connecting Australia's HPC

CONTENTS

Providing outstanding leadership and knowledge to position Pawsey as a global player in science and innovation

Making Tomorrow Happen, Today

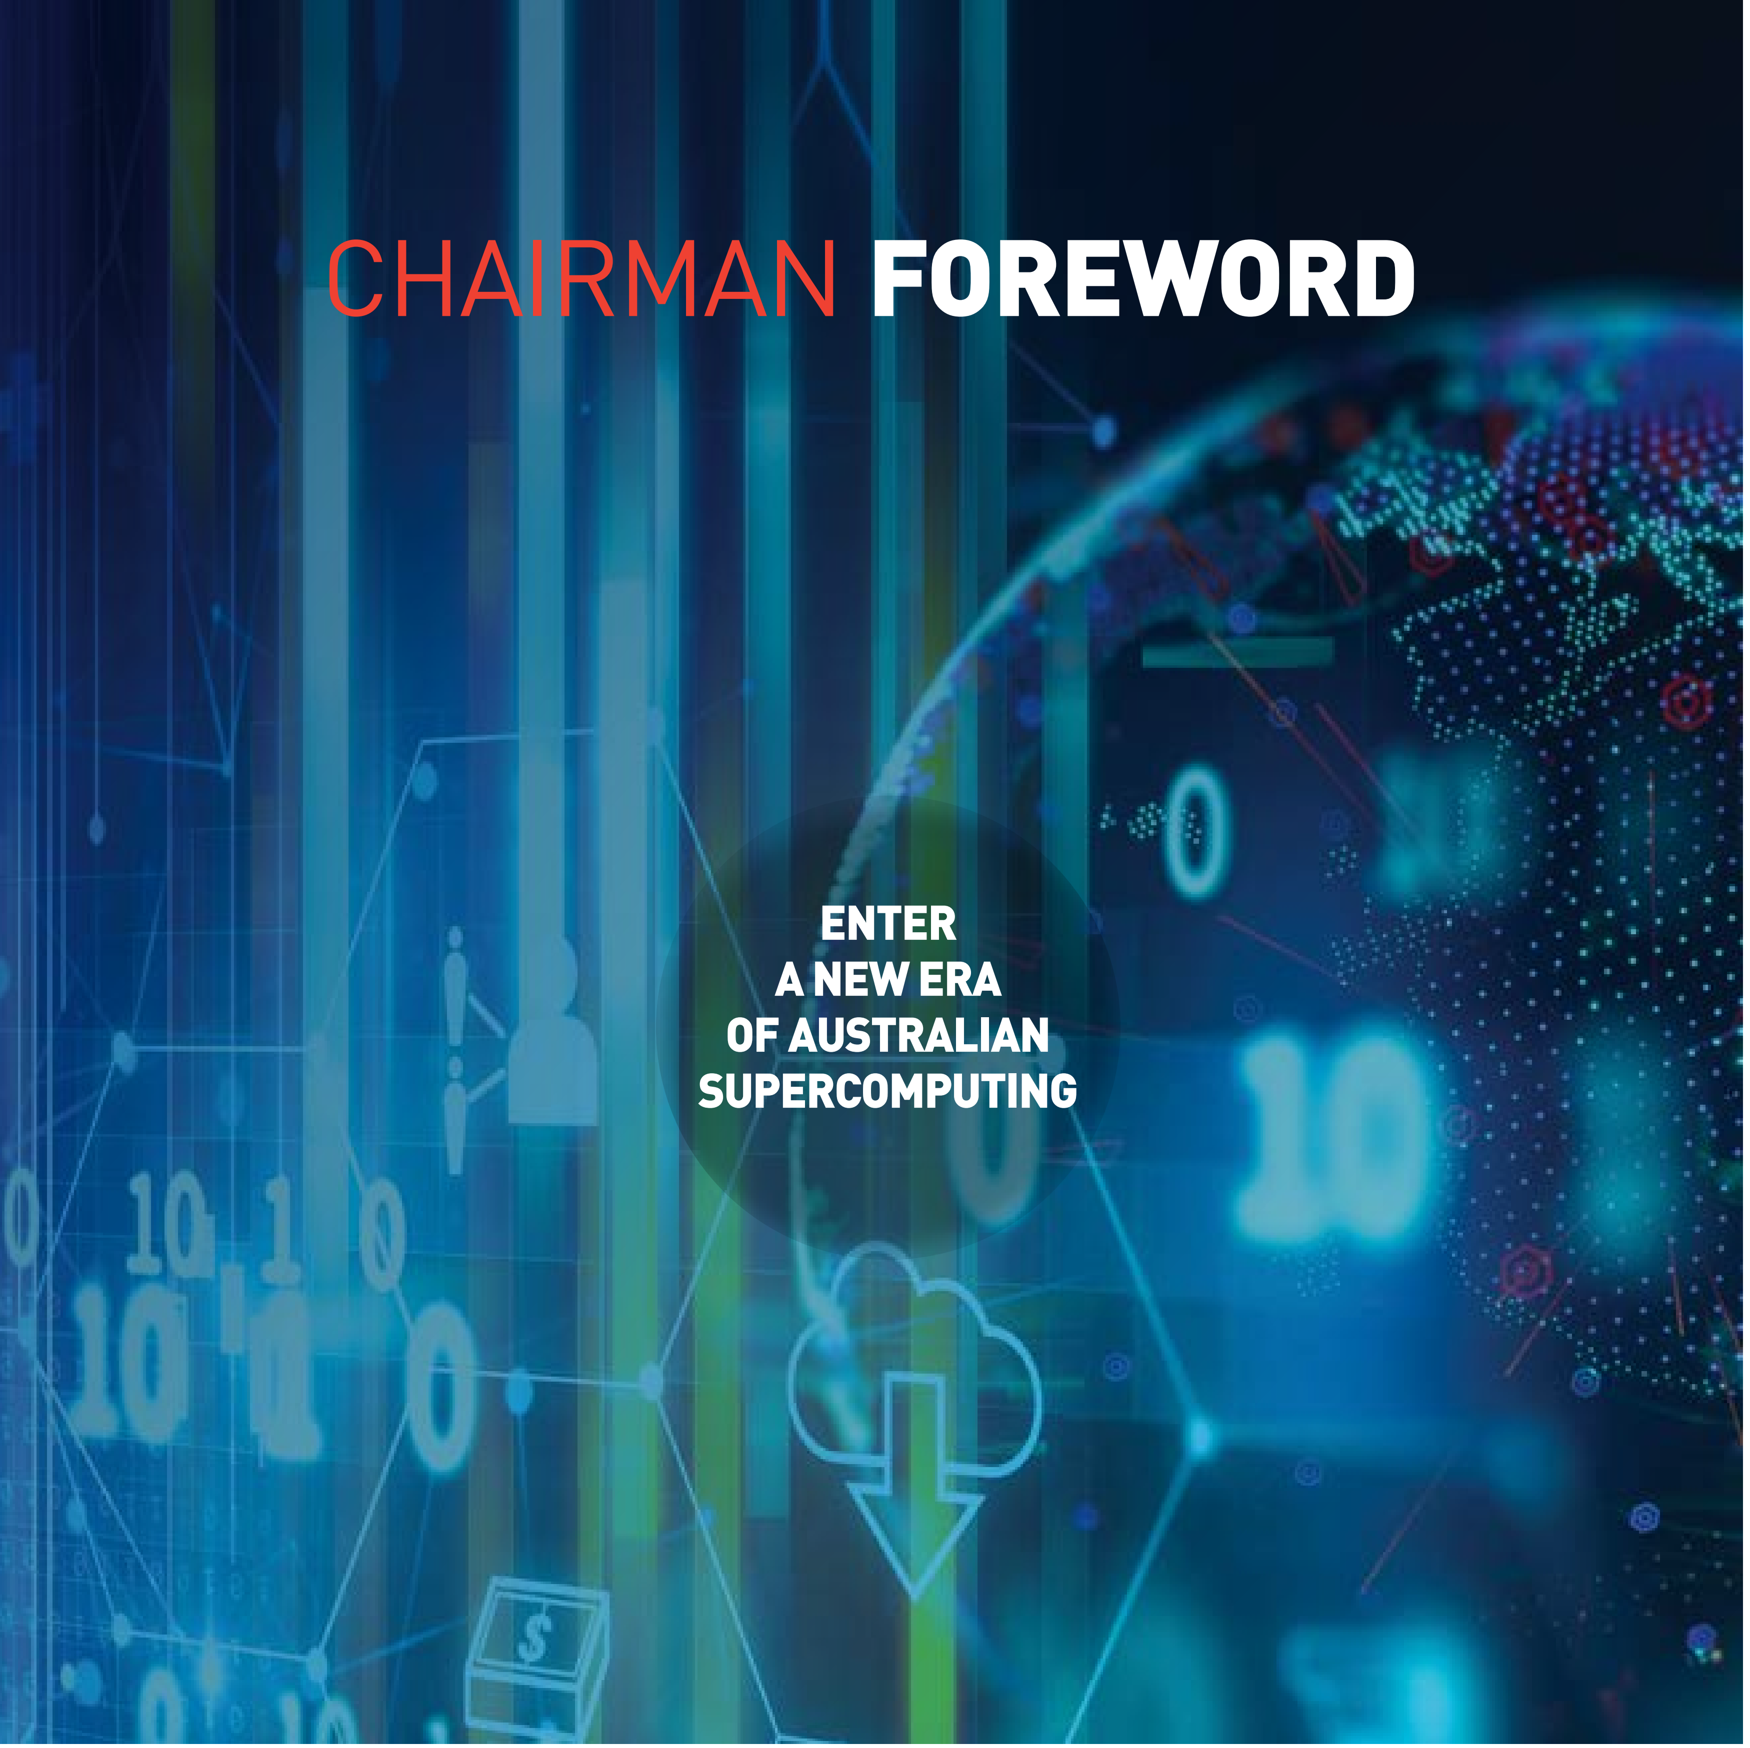

CHAIRMAN FOREWORD

ENTER A NEW ERA OF AUSTRALIAN SUPERCOMPUTING

WELCOME TO THE PAWSEY SUPERCOMPUTING CENTRE 2017–18 ANNUAL REPORT.

This year has seen major movements in policy, planning and placing the Centre firmly within the national research infrastructure and eResearch ecosystem, as we enter a new era of Australian supercomputing and data sciences.

It is an exciting time to be at Pawsey. **The Commonwealth Government's response to the 2016 National Research Infrastructure Roadmap** has recognised that Australia's two Tier 1 High Performance Computing (HPC) facilities – Pawsey and the National Computational Infrastructure (NCI) – need coordinated upgrades at regular intervals to keep pace with research needs and ensure Australia remains globally competitive. This commitment was recently demonstrated through the allocation of $70 million in funding towards the replacement of Pawsey infrastructure, as our flagship supercomputers Magnus and Galaxy approach the end of their operational lives.

Strengthening our partnership with **NCI** has been a focus throughout the year, and will continue to be into the future. Through sharing best practice and expertise, coordinating infrastructure upgrades, and working collaboratively in our operations, both across the nation and internationally, we are committing to a joint supercomputing agenda to support Australian innovation and development well into the future.

As we move into this new era, we must acknowledge the efforts made by our previous Executive Director, Neil Stringfellow, and Acting Executive Director, Ugo Varetto, for their work in establishing Pawsey and guiding it to maturity. This firm foundation is now a platform for future renewal, development, and ongoing success.

Our new Executive Director, Mark Stickells, is now driving our engagement with NCI and the international supercomputing community, and creating new opportunities for collaboration and partnership across a range of industry sectors. We welcome Mark to the Centre, and look forward to the results of his strategic focus moving forward.

Of course, business as usual has been all about enabling leading-edge and innovative science, accelerating research discoveries and solving even bigger scientific problems. Pawsey is continuing to underpin the activities of the radio astronomy community and the development of **the Square Kilometre Array (SKA)** precursors, including the doubling in size of the Murchison Widefield Array (MWA) earlier this year. At the same time we have been broadening our support of research across Earth science, medical science, engineering, chemistry, agriculture and more.

As HPC matures in Australia as a pivotal part of the research landscape, we will continue to broaden our horizons and partnerships to drive world-leading research that adds great value to Australia's economy and future.

I would like to recognise and thank my fellow **Board Directors** and the dedicated Pawsey staff for their effort and commitment as we continue our journey to remain at the forefront of this high-performance and nationally-critical digital domain.

John Langoulant AO
Chairman of the Board

This year has seen the ongoing development of the Pawsey Supercomputing Centre as a collaboration hub, helping more than 1,500 active researchers and nine Australian Research Centres of Excellence progress our state and national science priorities.

ACTING EXECUTIVE DIRECTOR'S REPORT

EXPANDING OUR COLLABORATIVE RELATIONSHIPS

Our collaboration efforts have actively extended further afield. We have made considerable progress in formalising our collaborative relationship with the NCI, to consolidate knowledge and best practice between Australia's two Tier 1 facilities. We have also developed formal relationships with equivalent international HPC centres in Singapore, Europe and the USA. Our international reach and reputation is growing.

ENGAGING WITH STAKEHOLDERS

While working closely with researchers, we have been exploring their future needs and objectives, to inform the configuration of the coming major upgrade of our HPC facilities. This engagement with stakeholders will continue through 2018 to ensure we develop improved facilities that continue to service the growing needs of Australia's research community.

EMERGING TECHNOLOGY ARCHITECTURES

To support both our research and upgrade planning, this year has seen the installation of the Advanced Technology Cluster, known as 'Athena', consisting of two emerging technology architectures. This facility allows Pawsey researchers to engage with cutting-edge technologies to accelerate their research outcomes, while providing a testbed to inform our future capital procurement.

SUPPORTING RESEARCHERS TO MORE EASILY FOCUS ON THEIR RESEARCH

We have also further developed the data analytics and visualisation capabilities at Pawsey, continuing to integrate them with our HPC facilities. This allows researchers to more easily focus on their research rather than computational frameworks, and translate big data into research insights.

UPSKILLING THE RESEARCH COMMUNITY

The focus is not only on infrastructure. Through the year we have continued to evaluate and improve our training strategies to reach, support and upskill our researchers and potential new users across Australia. This has seen an increase in the number of face-to-face training sessions offered, and the development of online training modules.

WORLD-CLASS SCIENCE OUTCOMES ACROSS A DIVERSITY OF FIELDS

Pawsey's combination of infrastructure, expertise and support has continued to underpin world-class science outcomes across a diversity of fields. This year our research partners simulated how a quantum computer behaves on a larger scale than ever before, and broke a world record in the process. They significantly improved our chances of detecting gravitational waves, in an ongoing contribution to the international collaboration that detected the first gravitational wave in 2015, and saw some of their team members awarded the 2017 Nobel Prize in physics for their discoveries. They identified optimal replenishment and abstraction sites on Perth's aquifers to maximise water draw while addressing environmental needs, began implementing cloud-based technologies to improve the monitoring of premature babies in neonatal units, and are developing artificial intelligence for targeted weed control in broadacre crops.

The breadth of research endeavour Pawsey supports can be seen throughout this annual report. I would like to acknowledge our staff, who go above and beyond to help researchers realise the opportunities HPC can deliver in furthering their research.

I have appreciated the opportunity to lead the Centre through this transition, and welcome Mark Stickells as the new Executive Director from July 2018.

Ugo Varetto
Acting Executive Director

Pawsey is all about getting things done that couldn't be done before, and solving the problems we have only just learned how to define.

NEW EXECUTIVE DIRECTOR'S REPORT

I feel privileged and excited to join Pawsey and I'd like to acknowledge the contributions of those who have gone before me, and thank the joint venture partners for their continued support of this wonderful facility.

Mark Stickells
Executive Director

IMPROVING UNDERSTANDING AND COMPETENCY IN DATA-LED SCIENCE

Pawsey is about people and innovative science – it is built on an amazing team with expertise across HPC and data science. We will continue to build on our strong reputation for training, user support, and end user engagement. This means continuing to invest in our people, our research users, and in the development of the next generation of data scientists and data-literate experts working across government, academia and industry. We support STEM, and increasingly STEAM, education through student engagement, career expos and internships, and intend to play an increasing role in improving understanding and competency in data-led science across multiple domains. We must grow and diversify the talent pool supporting the deployment of HPC for the benefit for all Australians.

CREATING AND MAINTAINING PARTNERSHIPS

Pawsey is built on strong partnerships. The nature of science and the increasing engagement with other research and infrastructure providers, government and industry, means that our partnerships in HPC will continue to grow. Our strong partnerships with the Western Australian universities and CSIRO, along with strengthening relationships with the NCI, national Centres of Excellence, and international HPC centres, will allow us to expand both our avenues of operation and our reach. Increased engagement will allow us to remain at the forefront of large-scale data science and HPC activities globally.

A FOCUS ON PERFORMANCE

Pawsey is synonymous with performance. With the support of the Commonwealth Government, we can continue to increase the power, speed and scale of our HPC infrastructure. Partnering that performance with our people, we can continue to meet the growing demands for data science and HPC across broadening domains to progress our national science, innovation and economic goals.

Performance relies on the right infrastructure platforms. With the expanding developments in cloud computing, visualisation, analytics and big data, we will continue to grow our investment in these areas, and the advanced software and applications our researchers require to access, manipulate, share, and draw insights from such large volumes of data.

WORLD CLASS PROJECTS SPANNING THE GLOBE

People, partners, performance and platforms enable projects that deliver results – Pawsey is all about getting things done that couldn't be done before, and solving the problems we have only just learned to define. We continue to support Australia's leading radio astronomy projects, as they push their world-class precursor telescopes towards the goal of the Square Kilometre Array. Equally, we hasten progress across many other research domains and industry sectors, as demonstrated by the ground-breaking science showcased in this report.

Pawsey has a tremendous foundation on which to build, and we are in a strong position to grow as a national facility with an international reputation. In this regard I would like to acknowledge the governance and support provided by the Board in developing Pawsey's strategy. Together we will deliver on Pawsey's potential to be an essential and critical enabler of our national science and innovation system.

HIGHLIGHTS

FUNDING

$70 million
TOWARDS A CAPITAL REFRESH

The Prime Minister and Minister for Science and Innovation visited the Centre to announce Commonwealth funding towards an upgrade of its supercomputing infrastructure.

$54 million
FOR OPERATIONS

Funding has been confirmed from the National Collaborative Research Infrastructure Strategy for operations from 2013 until mid-2023.

PLATFORMS

92 extra nodes

Added to Zeus, creating a mid-range cluster providing more than 20 million core hours per year to researchers.

12 NVidia GPUs

12 Graphics Processing Units nodes were added to the Nimbus cloud to accelerate Artificial Intelligence, HPC, and graphics applications, offering researchers up to 1.3 PetaFLOPS of deep learning performance.

PEOPLE

New Executive Director

Mark Stickells joined the Centre after the end of the reporting period. He is a strong advocate of science and innovation, with extensive experience in government and business engagement with research.

12 students

12 undergraduate students upskilled as part of Pawsey's Summer Interns program.

Over 600 people

Had HPC training in 27 different sessions across Australia.

4 Pawsey User Forums

Feedback received on over 52 topics.

COLLABORATIONS

2 institutions, one presence

NCI and Pawsey jointly represented Australia at Supercomputing 2017 in Denver, one of the world's largest international supercomputing conferences with 12,000 delegates.

400 members

400 members of the HPC community received the inaugural International HPC Best Practices Newsletter, initiated and coordinated by Pawsey staff.

7 NATIONAL AND INTERNATIONAL COLLABORATIONS WERE ESTABLISHED

University of Ontario
University of Technology Sydney
Western Australia Department of Health
National Supercomputing Centre Singapore
Partnership for Advanced Computing in Europe
National Energy Scientific Computing Center USA
Edinburgh Parallel Computing Centre UK

AWARDS

Best use of HPC in energy

2017 HPCWIRE Editors' Choice Award for best use of HPC in energy: Awarded to the Carnegie Wave Power Project undertaken by Pawsey and The University of Western Australia.

ABOUT **PAWSEY**

World Class

Pawsey Supercomputing Centre is a world-class facility that provides HPC resources to accelerate Australian scientific research for the benefit of all. The Centre, located in Perth, Western Australia, is home to the supercomputers Magnus, Galaxy and Zeus, and supports a steadily expanding network of users with supercomputing, data management, cloud, and visualisation services. Over 1,500 users benefited from Centre infrastructure in 2017–18, supported by over forty talented, highly-skilled centre staff.

An Example of Collaboration

The Centre is an unincorporated joint venture between CSIRO and the four public universities in Western Australia: Curtin University, Edith Cowan University, Murdoch University and The University of Western Australia. The Centre is also financially supported by both the Commonwealth Government and the Western Australian State Government. Originally established in 2009 within iVEC at the Australian Resources Research Centre in Perth, the Centre became part of the Commonwealth Government's Super Science Initiative to develop research infrastructure that would make a lasting contribution to Australian science. It was renamed the Pawsey Supercomputing Centre in 2014, in honour of the prominent Australian radio astronomer, Dr Joseph Pawsey.

One of the Word's Top 500 Computers

Pawsey Supercomputing Centre is now one of only two Tier 1 HPC facilities in Australia, the other being the NCI, based in Canberra. Both are on the **TOP500** list, which ranks the top 500 most powerful computers in the world. To maintain their positions (in a field where technological obsolescence is rapid) and ensure that Australia's HPC facilities continue to make internationally-competitive, computationally-intense research possible, the Commonwealth Government has committed funding to support a regular upgrade cycle enhancing capability across both Tier 1 facilities. This will be coordinated so Australia always has at least one facility operating at full capacity to keep pace with the increasing volume of complex computing demands across government, research and industry.

Real World Applications

Through supporting Australian researchers and their collaborators around the globe, the Pawsey Supercomputing Centre contributes to a wide range of data-intensive disciplines, including radio astronomy, geoscience, medical research, renewable energy, artificial intelligence, mining, food science, and agriculture. Research undertaken at Pawsey is rooted in the real world, from improving the health of all Australians, and safeguarding our food and water supplies,

to expanding our knowledge of the Universe. Centre infrastructure not only allows researchers to tackle previously-intractable challenges, but its data processing capability also allows researchers to work through data-intensive problems more quickly than ever before. For example, a medical researcher can use Pawsey facilities to process a year's worth of accumulated genetic data in five hours, accelerating the pace of research and hastening its path to impact.

Cutting-edge Science

Pawsey is playing a pivotal role in the cutting-edge science associated with the two SKA precursors, the **Murchison Widefield Array (MWA)** and the Australian Square Kilometre **Array Pathfinder (ASKAP).** Some of the finest minds in radio astronomy and computational science are working in partnership with Pawsey to process and analyse the massive amounts of data generated by these next-generation telescopes. These radio telescopes are making new discoveries about the history of our Universe, the future of our Solar System, and the vagaries of near-Earth solar weather, as well as laying the foundations for the SKA.

Upskilling Researchers, Students and Industry

Pawsey experts are drawn from the national and international supercomputing community and have a wealth of combined experience. They enable the big-thinkers in Australia to access the best computing services available to facilitate their research, allowing them to upscale their ambitions and consider new classes of problems. They facilitate access and use of HPC facilities through providing high quality training and support programs, and help upskill researchers, students and industry personnel via regular seminars and symposia covering data management, visualisation and supercomputing. These activities, along with internship opportunities, are helping develop the next generation of HPC and data scientists.

Driving Innovation

Pawsey Supercomputing Centre is an integral part of Australia's research landscape and is positioning Western Australia as a global hub for supercomputing expertise. The Centre is forging collaborative relationships with other supercomputing centres in the USA, Europe and Asia, to further develop the global impact of supercomputing. By driving innovation and accelerating discoveries across a multitude of different fields, the Centre is enabling Australian researchers and their collaborators to tackle the 'wicked' challenges facing our world.

Creating real world impacts through high-performance computing data science

Amazing Outcomes

CASE STUDY 1

Listening to the Ripples of the Universe

Gravitational waves are helping us delve deeper than ever before into the fabric of our Universe. Since the first gravitational wave was detected in 2015, a brave new field of astronomy has grown in response. The impact of this research is vast, plumbing the depths of black holes, checking the accuracy of Einstein's general theory of relativity, and helping us understand how our Universe fundamentally works.

Professor Linqing Wen and her team at the OzGrav ARC Centre of Excellence for Gravitational Wave Discovery at **The University of Western Australia** (UWA) are improving our chances of detecting these waves.

The team is focused on optimising detection methods for the Laser Interferometer Gravitational-Wave Observatory (LIGO) in the USA, as well as its European counterpart, Virgo. Professor Wen's team also contributed to the international effort which allowed LIGO to detect the first gravitational wave in 2015. LIGO scientists were awarded the 2017 Nobel Prize in physics for this discovery.

THE CHALLENGE

While our experience with gravity on Earth makes it the most familiar fundamental force, it is also the weakest. Radio, optical and many other types of telescopes can easily detect light and sound many galaxies over, but detecting gravitational waves is far more difficult.

Gravitational waves are so subtle that even our most advanced technology can only detect waves from the largest cosmic gravitational events. The first detected gravitational wave involved the merging of two black holes, each roughly 30 times more massive than our sun, allowing LIGO detectors to read a change in its detectors about 10,000 times smaller than a proton for just over 0.2 seconds.

The detection methods at the core of the gravitational wave searches must be accurate and responsive to react to these subtle changes in small time windows. This is where Professor Wen's work comes into play.

"These flashes from gravitational waves are very short and the electromagnetic light from them might only last for around two seconds. It's important for us to detect that blip quickly, then tell our optical and radio telescopes to point towards its origins so we can catch possibly extremely short-duration light from the gravitational wave event. Our software does this automatically and can send the trigger in under 30 seconds," said Professor Wen.

To perfect the detection method, Professor Wen's team needed to test it against large-scale simulated data as well as detector data from previous LIGO science runs.

THE SOLUTION

LIGO's third science run is due to start at the beginning of 2019 with new, advanced sensors capable of detecting fainter gravitational waves. In preparation for this run, Professor Wen's team are in midst of fine-tuning the search pipelines so the detection of gravitational wave signals run more smoothly and automatically. Optimizing and testing the robustness and detection efficiency of the searches requires an incredible amount of computing power. To accomplish these, Professor Wen's team used Pawsey Supercomputing Centre's Advanced Technology Cluster, Athena, as well as the Magnus supercomputer.

While Magnus is one of the most advanced public research supercomputers in the southern hemisphere, Athena was designed to provide researchers access to cutting edge technologies to propel their science. Professor Wen's team were able to take advantage of the state-of-the-art nodes available in Athena containing either many-core Xeon Phi processors or four NVIDIA "Pascal" Graphics Processing Units (GPUs).

"It would take us millions of Central Processing Unit (CPU) hours to test the program. We're trying to build a robust, low-latency program and that requires a lot of testing using simulations and data from previous science runs."

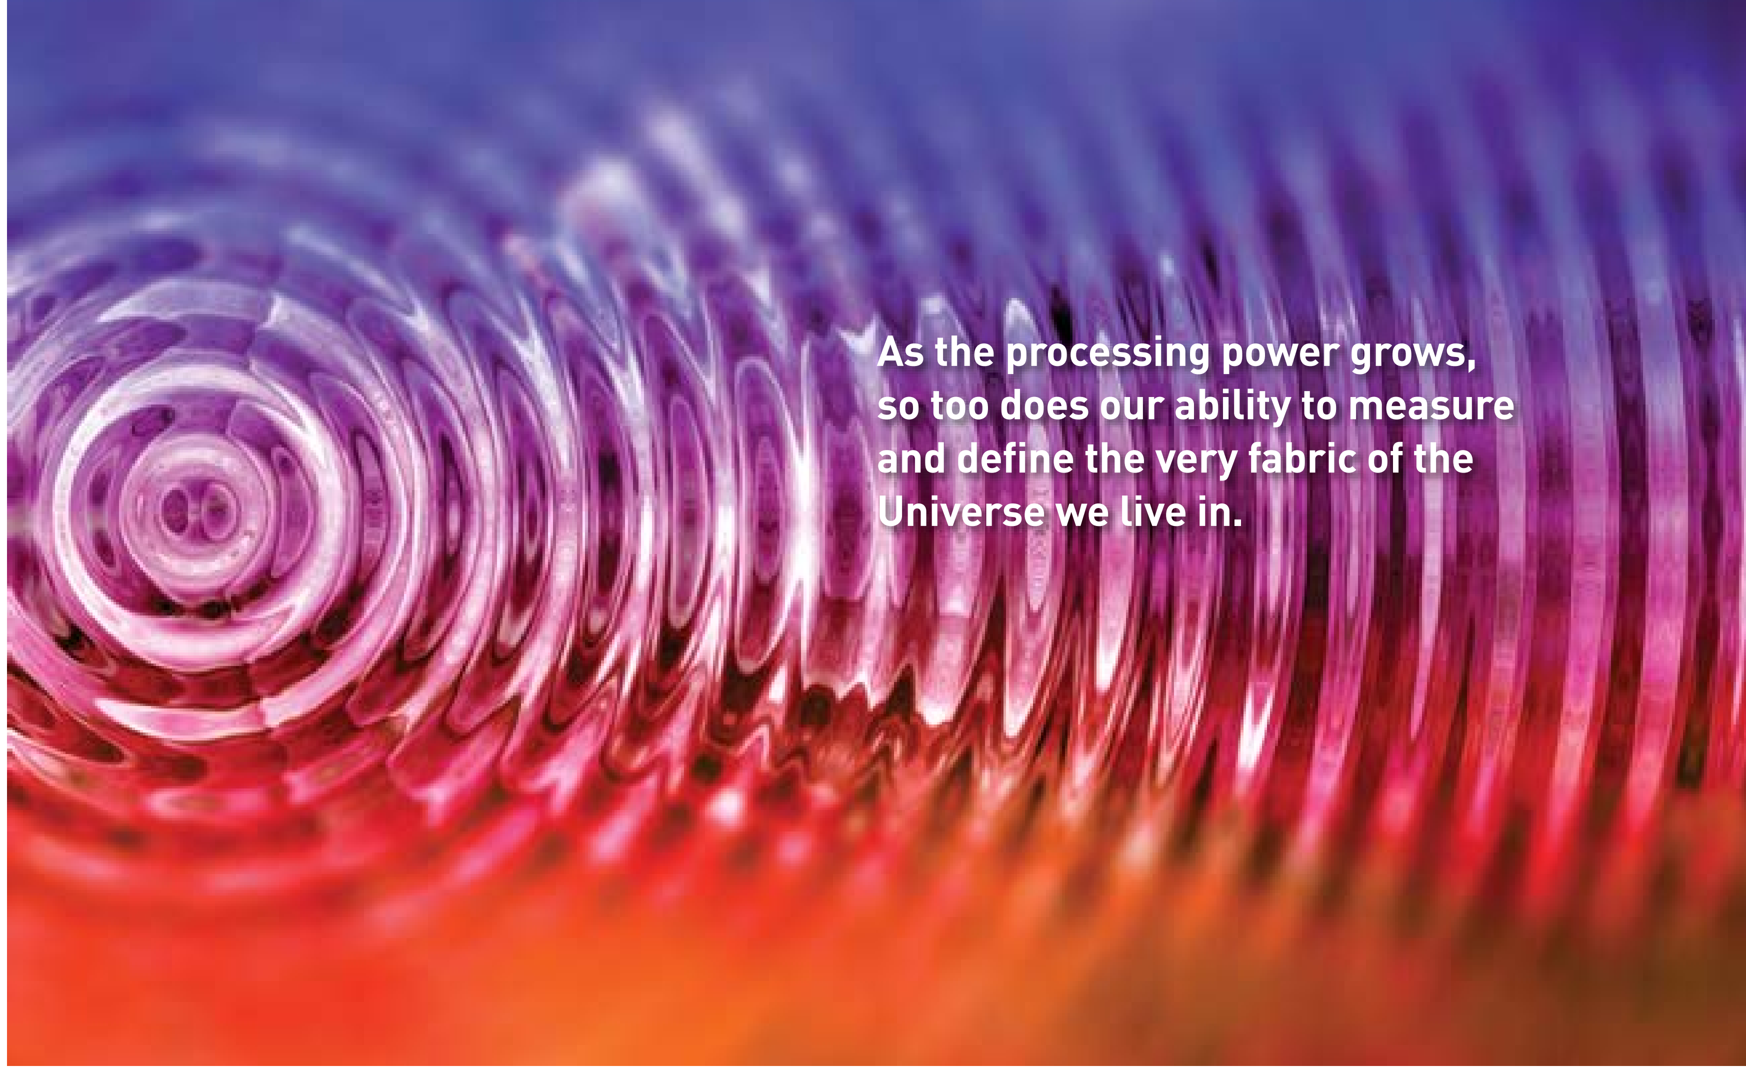

THE OUTCOME

Professor Wen's team are now preparing the newly improved detection method for incorporation into the upgraded LIGO systems. The team are in the final stages, checking over the program's data and sorting for errors. Soon, the work will become an important part of the search for more gravitational waves throughout the Universe; how and why they form.

"At the moment we have a full team checking the data around the clock to ensure we're not reading any false positives: that the program isn't accidentally reading gravitational waves when they're not there. We have regular meetings where we all get together to work as a team as the research requires expertise in many areas," said Professor Wen.

The demand for faster and better detection and localisation of gravitational wave signals requires ever more advanced supercomputing facilities like Pawsey. As the processing power grows, so too does our ability to measure and define the very fabric of the Universe we live in.

"We've been working with Pawsey for a while now and we hope to continue that relationship in the future. We have proprietary access to LIGO detector data and our work is very dependent on processing power. We're at the point now where we can collect so much data, but what we can do with it is dependent on how powerful our computers are."

Project leaders:	Professor Linqing Wen
Partner institution:	**The University of Western Australia**
System:	Magnus and Athena
Area of science:	Gravitational-wave Astronomy
Hours allocated:	1,000,000 (Athena) + 3,500,000 (Magnus)

CASE STUDY 2

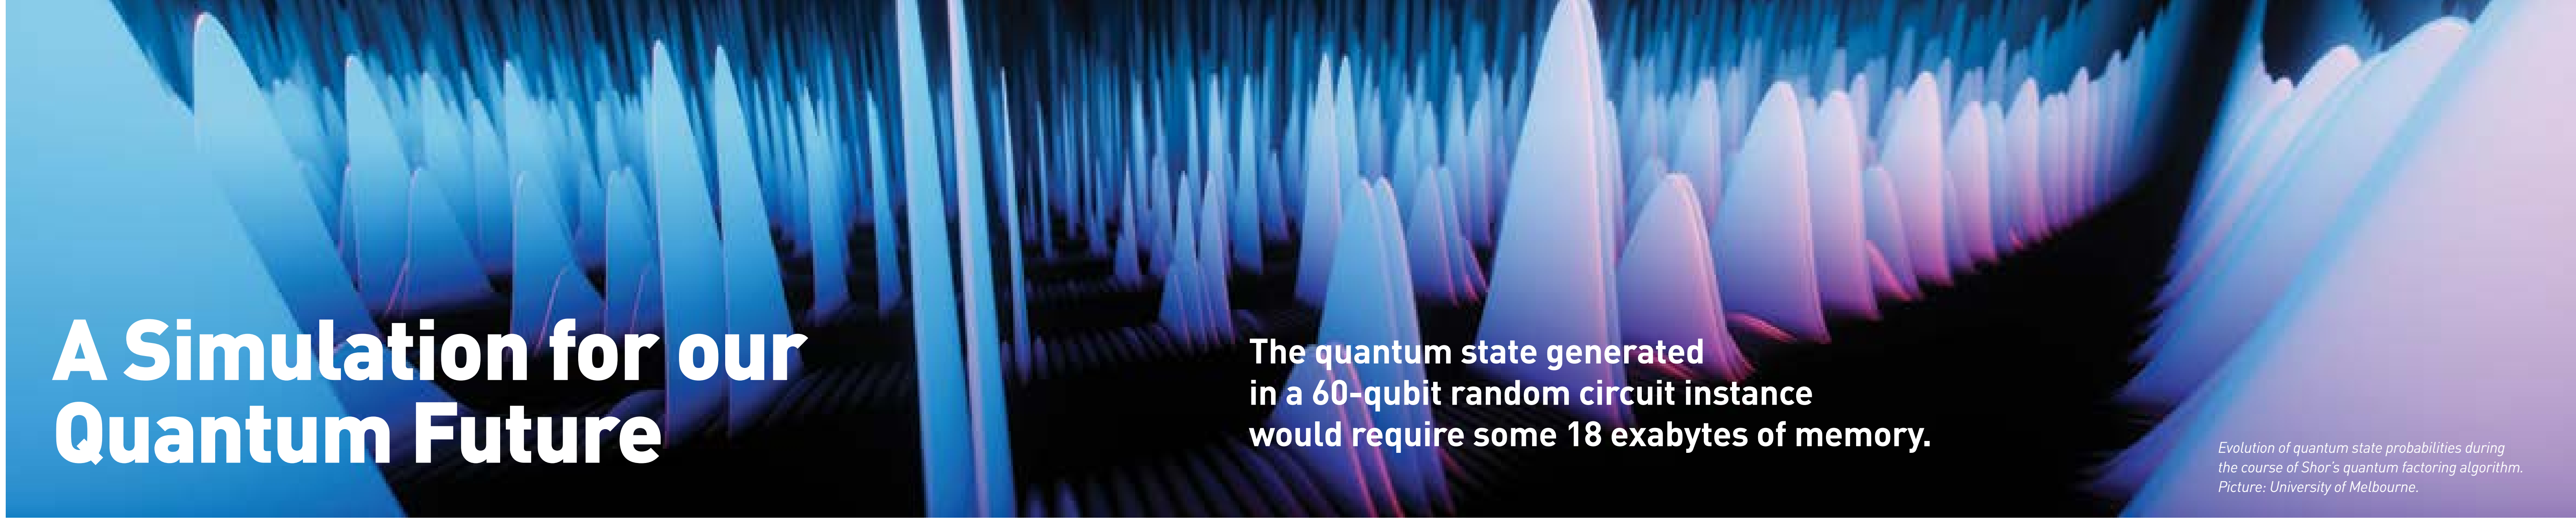

A Simulation for our Quantum Future

The quantum state generated in a 60-qubit random circuit instance would require some 18 exabytes of memory.

Evolution of quantum state probabilities during the course of Shor's quantum factoring algorithm. Picture: University of Melbourne.

Quantum computers are the revolutionary next step in the computing world. This technology harnesses the quantum behaviour of atoms to process information. Without the limitations of transistors, quantum computers have the potential to outstrip traditional computers in their problem-solving ability.

Professor Lloyd Hollenberg, Deputy Director of the ARC Centre of Excellence for Quantum Computation and Communication Technology at Melbourne University, is simulating how a quantum computer behaves with the help of Pawsey Supercomputing Centre and breaking a world record in the process.

THE CHALLENGE

Quantum computers stores information in qubits. Roughly equivalent to a traditional computer's bits, except with the potential to exist as both a binary 1 and 0 simultaneously. The number of qubits, their accuracy, and their ability to deal with complex algorithms are at the heart of quantum computer design.

Because a quantum computer can represent and process all possibilities of a binary string simultaneously, the longer the string, the larger the storage space a classical computer needs to match it, increasing exponentially with the length of the string. An accurate 50-qubit quantum computer generally requires several petabytes to simulate on a traditional computer.

"IBM and Google are producing systems around the 50-qubits or larger. The question is when they will get the errors down far enough for these machines to challenge conventional supercomputers," said Professor Hollenberg.

One way to test this accuracy is to simulate a quantum computer using a supercomputer. Using Pawsey Supercomputing Centre, Professor Hollenberg's team has created one of these simulations.

One common way to test these simulations is by having a supercomputer and quantum computer work on a random quantum circuit. If the quantum computer beats the supercomputer it has shown quantum supremacy.

Because quantum computer simulations require huge amounts of classical storage, the 50-qubit mark is considered the theoretical limit for simulating random circuits. However, Professor Hollenberg's team wanted to understand the limits of simulating useful quantum circuits. But such simulations also require significant supercomputing resources

THE SOLUTION

To go beyond the 50-qubit mark for a quantum algorithm, the team developed a sophisticated quantum computer simulator using Pawsey Supercomputing Centre.

"We didn't want to just simulate a random circuit because it doesn't do anything. It's an interesting problem, but we were more interested in what the maximum useful algorithm we can simulate is. The most iconic quantum algorithm is the factoring algorithm," said Professor Hollenberg.

This pragmatic approach, coupled with the team's sophisticated simulation framework specifically optimised for the quantum factoring algorithm, allowed them to test how a quantum computer might behave running the algorithm up to 60-qubits.

"To put that into perspective – to describe the quantum state generated in a 60-qubit random circuit instance would require some 18 exabytes of memory. It turns out the quantum algorithm, on the other hand, uses quantum states very efficiently."

Even with Professor Hollenberg's team focused on a resource-efficient quantum factoring algorithm, they still needed several terabytes of storage space and a large amount of processing power. To solve this problem, Professor Hollenberg's team were given a grant to carry out their simulation at Pawsey Supercomputing Centre.

Professor Hollenberg said it was a risky move to try and simulate a 60-qubit quantum computer in the finite space they had. The calculations, developed and run by MSc student Aidan Dang and postdoctoral fellow Dr Charles Hill, used up the 14 terabytes and most of the CPU allocation available to them, just nudging past the previous record of 56-qubits (albeit for the random circuit problem), set by IBM.

THE OUTCOME

Mr Dang's calculation successfully simulated how a 60-qubit quantum computer might solve the problem. It is an important find, because the factoring algorithm is the basis for almost all our modern encryption software.

"If you take DigiCert who produce Secure Sockets Layer (SSL) encryption keys, they estimate if you had a single core trying to factor a 2048-bit encryption key it would take a million times the age of the Universe. With a full-scale quantum computer, it might take days or months depending on the hardware."

Professor Hollenberg said this was an important space to watch. While quantum computers aren't here yet, their potential to solve the factoring algorithm means security systems need to become quantum-proof sooner rather than later.

"This is one of the problems for which a quantum computer will gain an exponential speed-up; for a range of other problems the speed-ups may be polynomial but still very significant. We need to watch this space."

In the meantime, simulating quantum computers allows a cost-effective way to train programmers and scientists in using the next generation of computer technology.

"There's a huge role now for education. We're now running our first subject in quantum computing, focusing on the computing aspects rather than the physics. It's of great interest to people of many backgrounds. We're teaching students from physics, maths, computer science, and engineering, using our specially developed 'quantum user interface' (QUI) software which allows students to gain hands-on experience with a quantum computer simulator. One day we hope to link this QUI to high performance computing (HPC) resources such as those in the Pawsey Centre," said Professor Hollenberg.

"There's great need to educate people on the logic of quantum information processing and programming so that we have people who can build the applications in various fields and work to make the Australian industry quantum ready."

While quantum computers still has some way to go before reaching mainstream use, Professor Hollenberg predicts quantum computers with hundreds or even thousands of qubits being utilised for specific problems in as little as five years.

With the future fast approaching, these simulations are vital for preparing for and understanding quantum computer technology. These simulations are only possible through partnering researchers like Professor Hollenberg with the resources of Pawsey Supercomputing Centre.

Project leaders:	Professor Lloyd Hollenberg
Partner institution:	University of Melbourne
System:	Magnus
Area of science:	Quantum Computing
Hours allocated:	37,444 core hours

CASE STUDY 3

When Atoms Collide

Our existence, as well as our ability to touch and interact with the Universe around us, is governed by the collisions of atoms and molecules.

The constant jostling and vibrating of these particles give us heat, light, and life. Professor Igor Bray and his team at Curtin University's Department of Physics and Astronomy are using supercomputers to chart movement at the atomic scale and unlock the amazing power possessed by these atoms.

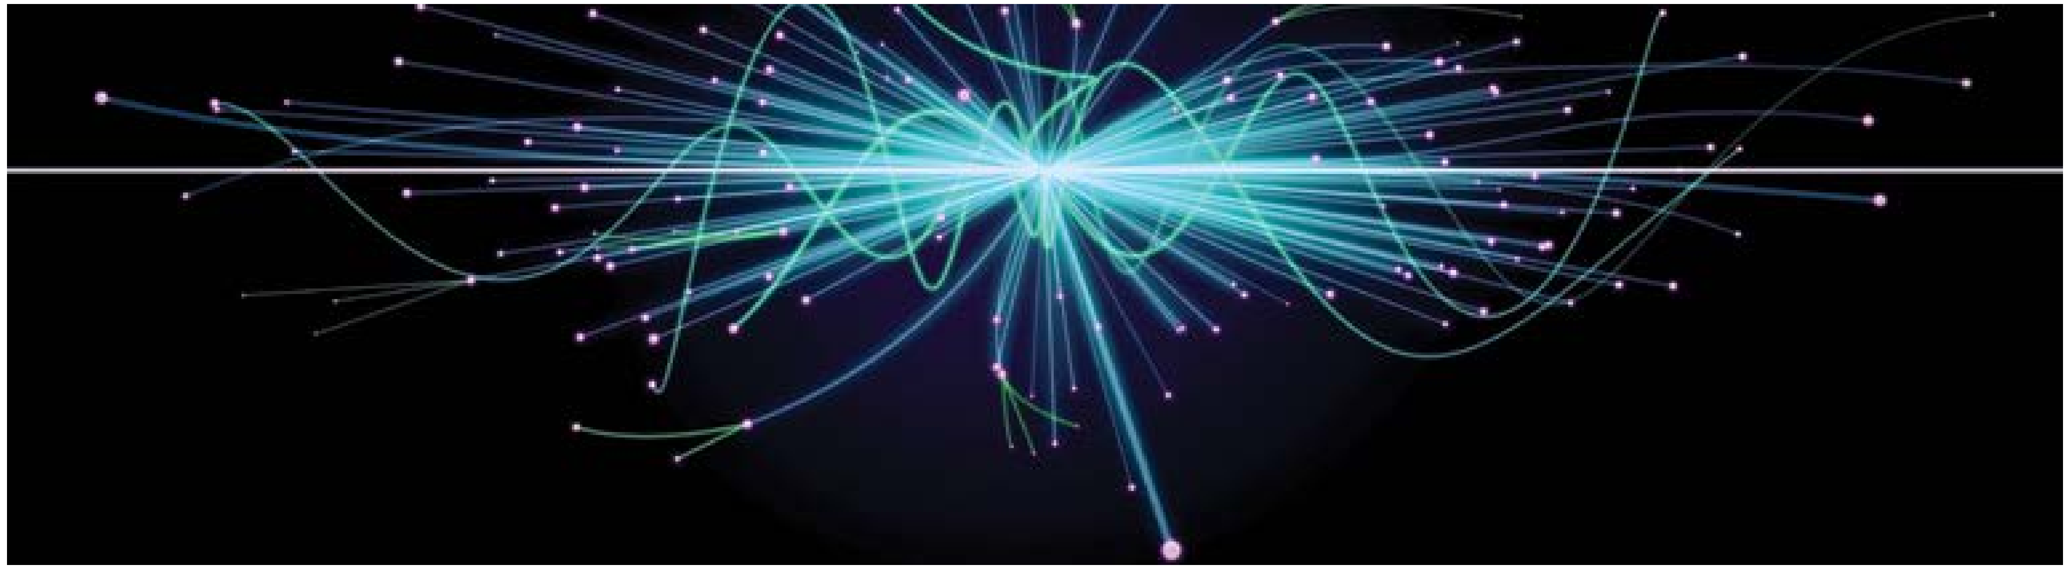

"With unlimited computing power, the sky's the limit."

THE CHALLENGE

Particle collisions happen on a scale that is both microscopic and complex. Many millions of particles will collide at each moment, transferring energy and sometimes emitting new particles.

It is only in the past few years that Professor Bray's team has had the computing power to deal with the more complex interactions of molecules.

"You can start off with an atom. It's a sphere, so the forces will be symmetrical on all sides. But then you move to molecules – even the simplest molecule, H_2 loses that symmetry," said Professor Bray.

"H_2 molecules orient in one direction. On top of that, they're now a molecule instead of an atom. So they'll also vibrate, which is another dimension. It becomes a much more complicated calculation."

While Professor Bray's team have set their eyes on larger molecules, the industries that use their work demand exact calculations. The new technologies exist in energy, medicine, and astronomy. So calculations must be surgically precise.

"It used to be, 20 years ago, astronomers would say if we could get data to within an order of magnitude or a factor of two that would be brilliant. That's not true anymore," said Professor Bray.

THE SOLUTION

To meet the demanding levels of computing power required for calculating particle collisions, Professor Bray's team have extensively used Pawsey Supercomputing Centre.

With the help of the Magnus supercomputer, one of the Southern Hemisphere's most powerful public research supercomputers, the team have been able to accurately calculate collisions.

"We've used Pawsey for everything. We've been utilising the Centre at its capacity throughout its existence. As the Centre became more powerful it opened up more possibilities. Now we've been doing molecular work over the past five years or so," said Professor Bray.

Professor Bray's team are still working on simple collisions, involving atoms or diatomic particles. The team are able to map more complex molecules, but the complexity of these collisions reduce result accuracy.

"With unlimited computing power the sky's the limit. We could do calculations with much larger molecules, like biological ones. But with what we have at the moment, it's important to decide how we're going to use those resources. The niche that we have chosen is to do whatever we do well or don't do it at all."

THE OUTCOME

The team's work has applications across many fields, from the lighting and television engineering industry, to some of the most ground-breaking science in the world.

Particle collision calculations made by Professor Bray's team are currently used in planning the International Thermonuclear Experimental Reactor. Much like a miniature version of the sun, it will be the world's first fusion reactor to produce a surplus of power and relies on the interactions of hydrogen atoms.

Professor Bray's work is also used in proton therapy, a new form of cancer therapy used to cure child brain cancer.

"You bombard the body with protons at the spot where the tumour is. To be able to do that the energy and direction needed for the protons to reach that spot must be known. So you can destroy the tumour without the bodily damage usually associated with X-ray radiation," said Professor Bray.

Most frequently, Professor Bray's work is used in astronomy. By analysing the unique signature of protons in space using satellites and observatories, particle collision calculations can tell astronomers what supernovae and solar winds are made of.

Despite the broad utility of Professor Bray's work, he insists that when research into particle collisions began many years ago, there was no specific outcome for the research in mind. It was simply curiosity about how the Universe worked at the atomic level.

"Is our research geared towards a specific application? Not really. That's not how we work. We work on solving problems. The fact they have applications is fantastic and it proves why we need to be working on these problems. But we don't know where these applications will be until we do the research," said Professor Bray.

"We know that certain problems deserve solving, even if we don't know what applications the solutions will have. It's also a feedback mechanism. When people know the solution they then wonder where it can be applied. That's why we work together with everybody."

Project leaders:	Professor Igor Bray
Partner institution:	Curtin University
System:	Magnus
Area of science:	Atomic Physics
Hours allocated:	15,000,000 core hours

CASE STUDY 4

GLEAM:
A Panoramic View of the Universe in Radio Colours

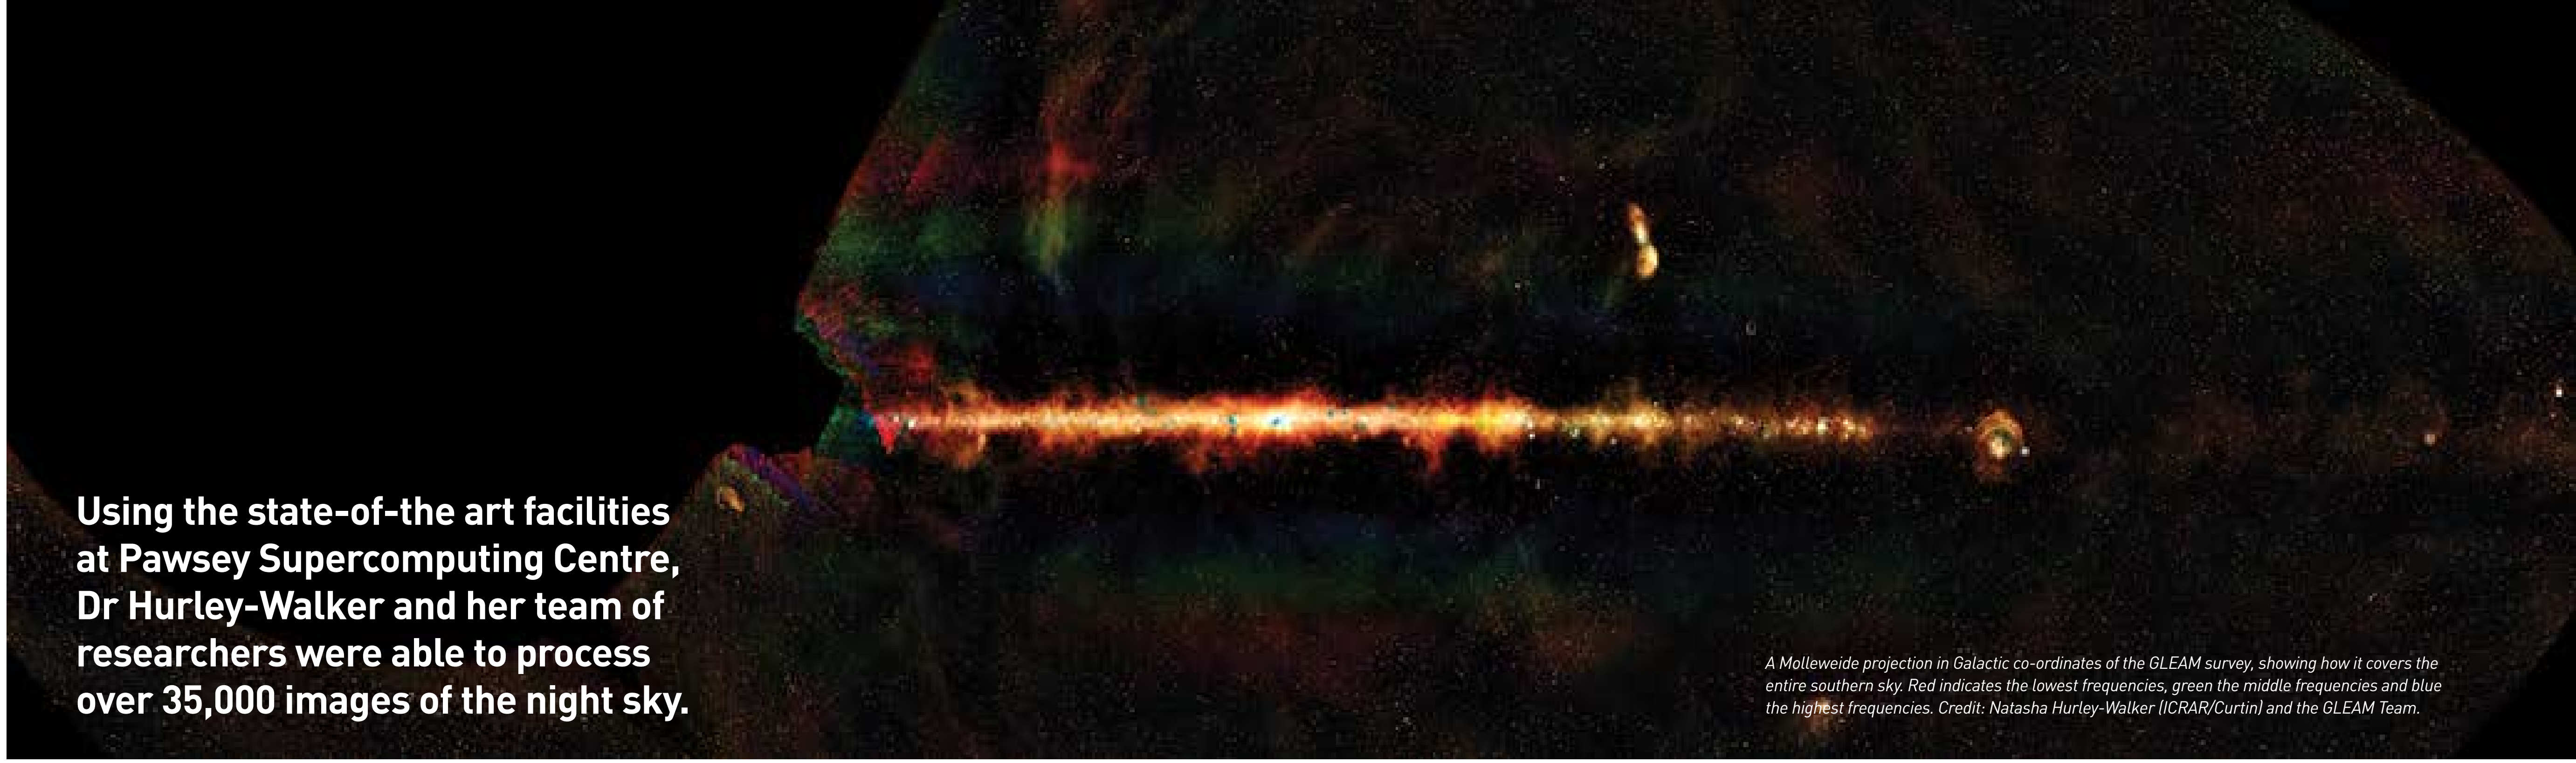

Using the state-of-the art facilities at Pawsey Supercomputing Centre, Dr Hurley-Walker and her team of researchers were able to process over 35,000 images of the night sky.

A Molleweide projection in Galactic co-ordinates of the GLEAM survey, showing how it covers the entire southern sky. Red indicates the lowest frequencies, green the middle frequencies and blue the highest frequencies. Credit: Natasha Hurley-Walker (ICRAR/Curtin) and the GLEAM Team.

In searching for cosmic events and the evolution of our Solar System, radio astronomers have traditionally surveyed only segments of the sky. Recently, however, the GaLactic and Extragalactic All-sky Murchison Widefield Array survey (GLEAM) project has collected and processed enough data to survey the entire southern sky: in 'radio colour.' GLEAM'S key researcher is Dr Natasha Hurley-Walker, a radio astronomer at Curtin University and the International Centre for Radio Astronomy Research (ICRAR).

Using the Murchison Widefield Array (MWA) telescope, and the world-class facilities at Pawsey Supercomputing Centre, the GLEAM survey has produced the first multi-coloured panoramic view of the Universe. Dr Hurley-Walker and her team have produced a catalogue of over 300,000 radio galaxies from the extensive sky survey which provides a valuable scientific resource for astronomers.

THE CHALLENGE

In the outback of Western Australia sits the Murchison Radio-astronomy Observatory (MRO). It is a radio-quiet zone where over 4,000 spider-like antennas make up the MWA radio telescope in 256 'tile' formations spread across ten square kilometres of land. The MWA, a precursor telescope for the Square Kilometre Array (SKA), collects low-frequency radio signals from the night sky. Signals are then correlated by researchers and sent through a fibre optic link to the Pawsey Supercomputing Centre.

The sheer number of the collective MWA antennas, and the correlations between them, compared to previous radio telescopes, means that vast amounts of astronomy data are generated. Data for the GLEAM survey was sent to the Pawsey Supercomputing Centre at the rate of about one Blue-Ray disc per minute, requiring immense data storage and processing. "Most past radio surveys were performed with dish-based interferometers with a low number of antennas, small fields-of-view, and narrow bandwidths," said Dr Hurley-Walker.

"The MWA is a massive aperture array with no moving parts, a wide bandwidth, and a very large field-of-view. The first step in transforming interferometric data into images is to perform a Fourier transform, which becomes very much more [computationally] expensive as you increase the field-of-view, the bandwidth and the number of antennas, so a supercomputer was very necessary to cope with this problem."

THE SOLUTION

By using Pawsey Supercomputing infrastructure, the GLEAM research team could store and process over 600 terabytes of radio astronomy data from the MWA telescope as it was piped down from the Western Australian outback.

"We used the supercomputers Raijin (at the National Computing Infrastructure, NCI) and Galaxy (Pawsey) to process the GLEAM survey. The way that the GLEAM observations were performed meant that we had around 7, 000 near-identical datasets to process," said Dr Hurley-Walker.

"This was perfectly suited for parallel supercomputing since each data set could be processed at the same time, subject to disk space limitations. Of course, we had to do this a few times, because each time we learned something new about the data! So all in all, we used a few million CPU hours to fully process the survey."

OUTCOME

Using the state-of-the-art facilities at Pawsey Supercomputing Centre, Dr Hurley-Walker and her team of researchers were able to process over 35,000 images of the night sky. With each MWA telescope capturing around one-tenth of the sky, Dr Hurley-Walker painstakingly combined the images to create a complete panoramic view of the southern sky in radio colour: a world-first.

"We created the images and performed source-finding to derive a catalogue of over 300,000 radio galaxies, each with 20 brightness measurements across the low-frequency band of the MWA," said Dr Hurley-Walker.

"The images and the catalogue are helping astronomers learn more about the environments around supermassive black holes, find colliding clusters of galaxies, and understand the cosmic ray electron budget of the Milky Way. The survey is also a perfect way of calibrating the low-frequency component of the SKA, due to be built on the same site as the MWA in coming years."

The GLEAM data set will soon be published in a series of papers that the radio astronomy community will be able to publicly access.

"It would not have been possible without the Pawsey Supercomputing Centre," said Dr Hurley-Walker.

Project leaders:	Dr Natasha Hurley-Walker
Partner institution:	ICRAR
System:	Galaxy
Area of science:	Radio Astronomy
Hours allocated:	41,877 core hours

CASE STUDY 5

Western Australia's Water Supply

Protecting Perth's Aquifers

One of Australia's most important – and sometimes most tumultuous – relationships with the natural environment is between the public and its water. In the south west of Western Australia in particular, delivering enough fresh water for a growing population is a pressing challenge when the conditions are drier than ever due to less annual rainfall as a result of climate change.

In the capital city of Perth, where around 80 per cent of the State's 2.6 million population live, dams and rivers are all but redundant as water supplies. The answer to date has lain in the vast reserves of water our aquifers hold. But with climate change also impacting aquifer recharge, using these resources now presents its own difficulties for the State Government as it manages a reliance on over 400 gigalitres of groundwater a year to meet the city's public and private water supply needs.

With the help of Pawsey Supercomputing Centre and its most powerful computing system Magnus, Professor Brett Harris and his team from Curtin University's Department of Exploration Geophysics have helped create detailed 3 D models of Perth's groundwater aquifers by sending airborne electromagnetic survey measurements directly to the powerful supercomputer for near-real-time processing of data. These models are important components of innovative government moves to sustain these valuable water resources.

THE CHALLENGE

Aquifers are complicated hydrogeological structures that may span thousands of square kilometres. There are dynamic shallow aquifers that may be connected to streams and rivers, or play a role in keeping valued lakes and wetlands alive. There are also the older, deeper aquifers, where the groundwater may have remained untouched for tens of thousands of years.

Perth's largest groundwater source, the Gnangara system, is already one of the best understood systems in the world, but more extensive knowledge is required if the aquifer system is to be managed sustainably in the 21st century. By increasing the understanding of how the different parts of groundwater systems are connected, and where and how water moves through the systems, especially from the superficial aquifer into the deeper aquifers, the State Government can identify locations and rates for taking groundwater that have less impact on the environment and other water users.

The work also gives State Government the ability to identify the best locations where water can be put back – replenished or recharged – into the deeper aquifers, to optimise pumping for public scheme supply, and again, to improve environmental outcomes.

To support the State Government's plans for Water Corporation to put in new extraction and aquifer replenishment wells on the Gnangara system, an enormous and detailed 3 D picture from the surface to more than 2 km below the ground was needed.

To map Perth's complex multi-level aquifer systems, data spanning thousands of square kilometres had to be collected, processed, integrated and converted to information on rock type and groundwater chemistry.

"We had a total of about 100,000 readings, and a couple of hundred simulations had to be run on each reading," Professor Harris said.

So big, in fact, that the team claim a single CPU processing the data would have taken years. Instead, the team turned to the Magnus supercomputer, which reduced the time required to a matter of hours.

THE SOLUTION

With the sheer bulk of data being sent back-and-forth, the use of Magnus – one of the Southern Hemisphere's most powerful research supercomputers – was vital. It allowed Professor Harris and team to spot errors on-the-go, test new ideas and build their 3D map as quickly as possible.

"It was about using the computing speed and power of Magnus in order to do things that would be completely impossible otherwise," Professor Harris said.

Such a quick turn-around time for processing the electromagnetic survey data ensured the team could integrate data from multiple sources and accurately model the geology of the aquifers.

"Our role in this is building a framework, so when you run the model it will tell you what's going to happen in 50 or 100 years. It might take years to reverse any changes, so getting this big framework right is important," he said.

THE OUTCOME

The research has already helped the Department of Water and Environmental Regulation recommend replenishment and abstraction sites based on ability to deliver goals of both maximum redraw and optimum environmental benefits, supporting groundwater levels near lakes, wetlands and areas of seawater intrusion risk.

Senior Hydrogeologist for the Department of Water and Environmental Regulation, Dr Jon-Phillipe Pigois, is one of the scientists using this research.

"All of this technical information is going towards the Gnangara groundwater allocation plan. It was last released in 2009, so it's due for an upgrade," Dr Pigois said.

Dr Pigois notes that water replenishment can be highly beneficial if the water is put back into the best possible location for the health of Perth's underground aquifer systems, and this type of active management requires detailed knowledge of the entire groundwater system.

"Groundwater replenishment is the 'dam building' of the 21st century. It is accepted in areas like Perth as a better storage option for harvested and recycled water, and helps offset water losses from climate change, evaporation or water use.

"In the case of Gnangara, where the water accessed from the system is an essential component of Perth's public and private water supply, and would require billions of dollars to replace with alternative water supplies, the contributions of this science to getting water use management right and keeping aquifers sustainable, is returning major financial benefits to Western Australians.

The team is now writing code to automate processing and integration of data from many sources so that the electromagnetic readings taken can be instantly inverted to inform groundwater models and ultimately assist in water management decision making.

Project leaders:	Associate Professor Brett Harris
Partner institution:	Curtin University
System:	Magnus
Area of science:	Hydrogeophysics
Hours allocated:	8,500,000 (small part used for this project)

CASE STUDY 6

Artificial Intelligence for Targeted Weed Control

Weeds cost Australian grain growers $3.3 billion per year, and cause annual yield losses of 2.7 million tonnes.

Food security is one of the major global issues facing researchers and policymakers today, and crop infestation by weeds is one of the biggest threats to grain productivity. Weeds cost Australian grain growers $3.3 billion per year, and cause annual yield losses of 2.7 million tonnes. Currently, the most efficient method of weed control available is to use a blanket application of herbicides, but this wastes resources and cultivates herbicide-resistant weeds.

Dr Selam Ahderom and his colleagues at Edith Cowan University's Electron Science Research Institute (ESRI), alongside industry partners at the Grains Research and Development Corporation and Photonic Detection Systems, are developing technology that will allow grain growers to apply herbicides more selectively and efficiently. Using reflected spectral data from lasers and spatial information from cameras mounted on tractor booms, their technology will use artificial intelligence to discriminate between the crop and weeds such as ryegrass, wild radish, and wild oats, and will spray herbicide only where it is cost-effective to do so.

THE CHALLENGE

Weeds and grain crops may be easy for a trained human eye to discriminate between, but for a computer it is much more of a challenge. There are many variables that will affect the shape, colour and spectral properties of a plant, including its growth stage, the presence of disease, and water deficiency.

Furthermore, environmental factors like lighting, the presence of water droplets, and camera angle can change the plant morphology as measured by the sensors and camera. The ESRI team, led by Professor Kamal Alameh, needed to create a program that would be flexible enough to recognise weeds in various conditions with a high level of accuracy, while processing the data quickly enough to match the moving speed of a tractor, up to 25 kilometres per hour.

THE SOLUTION

The ESRI team designed a system using neural networks and deep learning to distinguish between crops and weeds, but needed a high-performance computing solution to realise their vision.

Just like a human brain, the neural networks require training, in order to calibrate the 'weights' of different filters, which the network uses like neurons in making the decision on whether an image is a weed or not.

"What each particular neuron does is not very important," said Dr Ahderom, "but when it has many different potential inputs, when it is trained by looking at the images in hundreds of thousands of different scenarios, a particular pattern emerges in the weight distribution in how each neuron should behave."

" If we are able to put this technology into farmers' hands, we could, in certain cases, reduce the amount of herbicide used in a field by up to 90 per cent."

The ESRI team used the Zeus cluster at Pawsey, allowing them to train the networks in different configurations simultaneously using multiple GPUs, each of which runs at more than ten times faster than a personal computer.

"Being able to see different models simultaneously, and being able to compare them, allows us to work very fast in evolving the system," Dr Ahderom said.

THE OUTCOME

Using the power of Pawsey's world-class facilities, the ESRI team have made significant progress towards their goal of creating a commercially viable variable-rate herbicide applicator. They are currently in the process of optimising their system, with the hope that they will be able to test a fully functional prototype in the field in 2018.

"If we are able to put this technology into farmers' hands, we could, in certain cases, reduce the amount of herbicide used in a field by up to 90 per cent," Dr Ahderom said. This would significantly reduce the estimated $113 per hectare that is spent on weed control by farmers each year.

Once the program has been optimised for wheat crops, the photonic weed control technology also has the potential to be adapted to other applications, including anywhere "an invasive weed is affecting either an industrial process or the natural ecosystem, whether in natural reserves, roadsides, or other areas," Dr Ahderom said.

Project leaders:	Dr Selam Ahderom
Partner institution:	Edith Cowan University
System:	Magnus and Zeus
Area of science:	Plant Protection, Neurocognitive Patterns and Neural Networks
Hours allocated:	50,000 core hours

CASE STUDY 7

Revealing Secrets from Ocean Currents

On 7 March, 2014, a commercial aircraft carrying 239 passengers vanished. Flight MH370 was scheduled to travel from Kuala Lumpur to Beijing, yet disappeared from civilian radar screens only one hour into the planned six-hour flight.

Despite exhaustive international searches within the South China Sea and Indian Ocean, investigators were unable to locate MH370.

In a new chapter of the search for MH370, investigators called on the expertise of oceanographers to pinpoint a crash site of the aircraft. Professor Charitha Pattiaratchi and Dr Sarath Wijeratne, oceanographers from UWA, used the Magnus supercomputer at Pawsey Supercomputing Centre to conduct oceanic drift modelling and predict the trajectory of aircraft debris from MH370.

Utilising the world class facilities at Pawsey, Professor Pattiaratchi focused on debris which washed up at Reunion Island, tracing it back to its original location through ocean currents. The pattern of debris dispersal in the ocean was used to assist search team investigators in locating a possible crash site.

"We were able to run 50,000 individual debris drift simulations at 25 locations over a two year period to identify the possible location of the crash site."

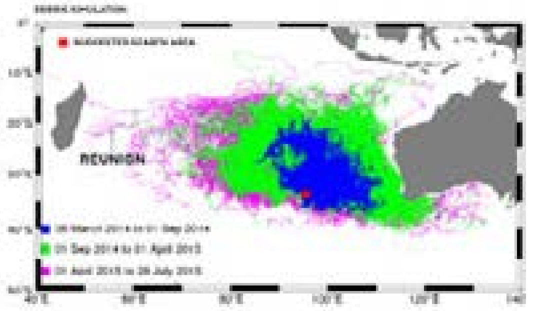

THE CHALLENGE

For reasons yet unknown, flight MH370 deviated from its original flight path. Experts believe the aircraft to have crashed into the Indian Ocean.

Initially, the search area for MH370 was defined by a series of seven satellite pings relayed from the plane's engines that told the now famous story of the plane going off-course in its final moments.

The Commonwealth Government, with assistance from Malaysia and China, were faced with a 120,000 square kilometre search area in the Indian Ocean near Western Australia. This search area, based upon the projected flight path provided by the satellite pings and some early oceanic debris modelling, was unfortunately found empty.

Now, UWA oceanographers have used a new model to run complex numerical simulations to search for the real crash site of MH370.

"We were able to run 50,000 individual debris drift simulations at 25 locations over a two year period to identify the possible location of the crash site," said Professor Pattiaratchi.

The calculations required far more than standard computational power, so to unravel this riddle, Professor Pattiaratchi and his team required the state-of-the-art infrastructure at Pawsey Supercomputing Centre. With it, UWA oceanographers could perform extensive dispersion modelling to investigate possible MH370 debris origins.

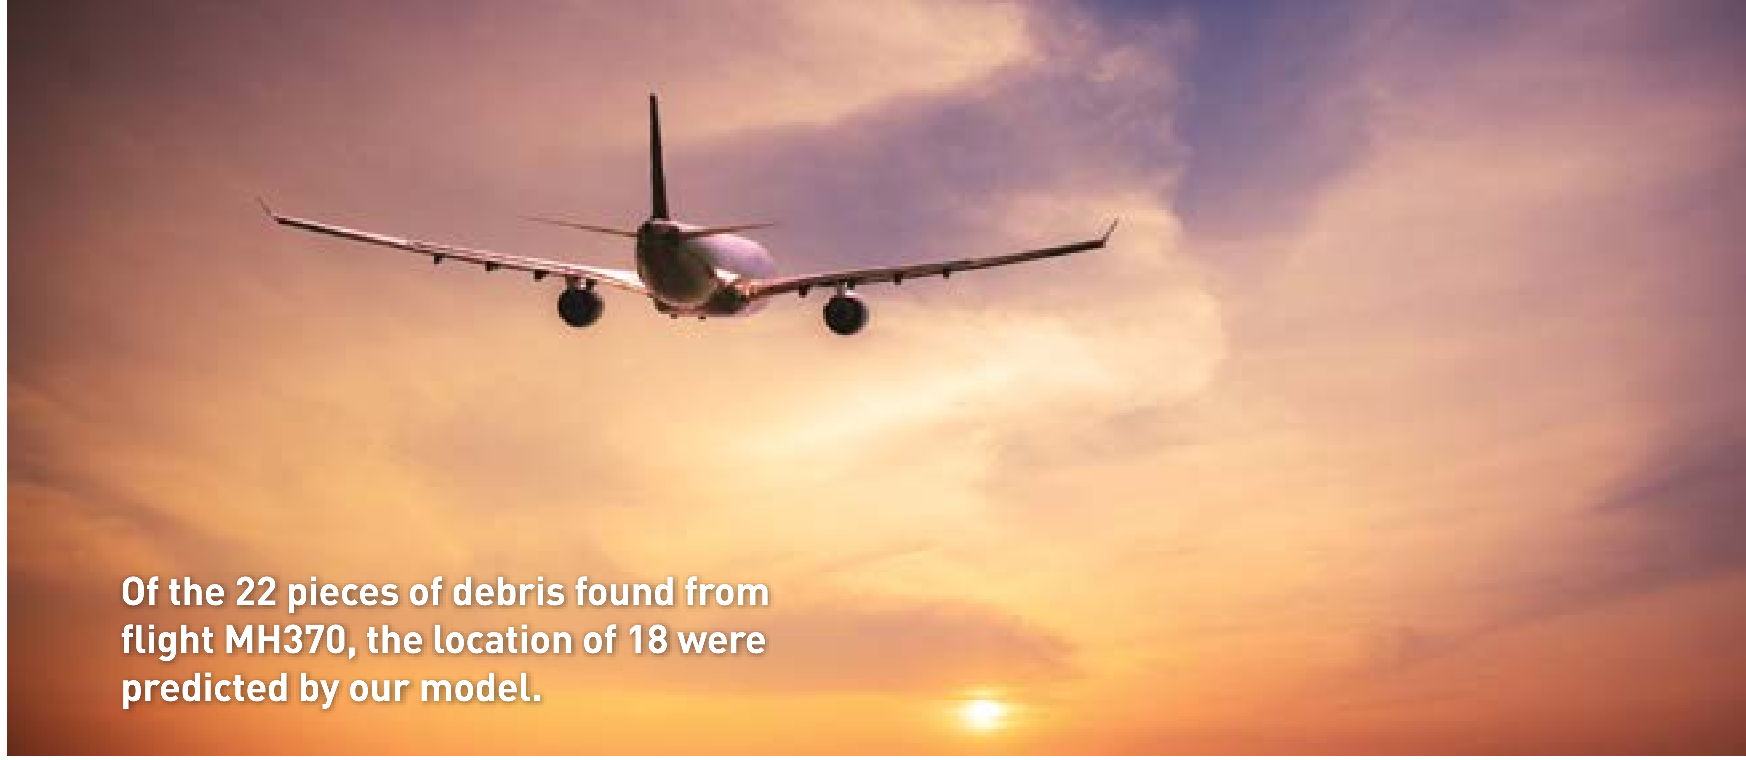

THE SOLUTION

With the discovery of a piece of a wing called a flaperon, Professor Pattiaratchi was able to calculate an approximate location of MH370 by predicting the debris path and the time that it took to reach Reunion Island.

Mapping the dispersion of flight crash debris requires complex mathematical equations to represent all the natural forces the debris encounters to affect its path. Light debris will float upon the surface and be subject to different forces than the heavier debris that sinks to the ocean floor.

"The drift models that simulate objects at the surface or at depth take into account the current field; the wind and wave field (if at the surface), as well as object-specific drift properties to advect the drifting object. In our models the 50,000 particles are used to propagate a range of debris behaviour statistically and deterministically," said Professor Pattiaratchi.

Whilst drift modelling can tell the story of airplane crashes, it can also be applied to a wide gamut of other situations where a subject's location is determined by ocean conditions. The dispersion model was run to include a range of different problems, to maximise the research benefit.

"The drift modelling conducted with the Magnus supercomputer was designed to include practical applications, not just for the search of MH370, but also for oil spill prediction, problems with seagrass wrack in Geographe Bay, turtle tracks, the exchange of larvae in marine populations, and marine debris pathways."

OUTCOME

The search for MH370 was suspended last year by authorities until credible new information could be found; but thanks to the work of Professor Pattiaratchi and others, the search began again in January 2018.

Though the crash site for MH370 has not yet been found, the debris modelling pointed to a new, significantly smaller, 25,000 square kilometre area – further north than previous predictions – which is currently being searched.

The research done with Pawsey has turned up promising leads. It was used to guide the search for MH370 coastal debris, to great success.

"Of the 22 pieces of debris found, the location of 18 were predicted by our model. Those not predicted were in Mauritius and Rodrigues Islands which may not be well represented in the oceanographic model."

Thanks to the help of Pawsey and UWA, the desperate search for answers to the MH370 mystery might soon come to an end for both the Australian Government and the loved ones of those lost passengers.

Project leaders:	Professor Charitha Pattiaratchi
Partner institution:	The University of Western Australia
System:	Magnus and Zeus
Area of science:	Physical Oceanography
Hours allocated:	100,000 core hours

CASE STUDY 8

BIG DATA ANALYTICS

Protecting Lives of Premature Babies

The Artemis Platform cloud service can connect to bedside monitors, allowing a live stream and analysis of around 1,256 data points per second, per patient, amounting to roughly 600MB per day.

When visiting a neonatal ward, you might find paper records of the heart rate, blood oxygen, and respiratory rate of premature and ill children, taken at hourly intervals. When Dr Carolyn McGregor AM of the University of Ontario walks into a neonatal ward, she sees the phenomenal potential for revolutionary advancements in healthcare.

The data recorded on paper represents a fraction of what monitors read. What is not recorded simply disappears, risking infant health and wellbeing; including the ability to predict infections such as Late Onset Neonatal Sepsis. Without these predictions, one in five low birth-weight babies can die from infection.

Dr McGregor and the team at the University of Ontario Institute of Technology have developed the Artemis platform; a system that supports the acquisition and storage of patients' physiological data streams and clinical information for real-time analytics, retrospective analysis and data mining. In collaboration with the Pawsey Supercomputing Centre, the University of Technology Sydney and the Western Australian Department of Health, Dr McGregor is bringing her work to the Australian community for the betterment of society.

THE CHALLENGE

In 2009, Artemis was deployed as a pilot research study to demonstrate how data collected by medical devices in hospital neonatal intensive care units (NICU) can be harnessed and used for new approaches to providing care to fragile premature and newborn infants.

"As a patient's condition changes, initially it can be very subtle," said Dr McGregor. "I realised that they didn't have a platform that could take in all of this data and help them to watch, and that we could learn from all of that data – things that we haven't learned before."

Artemis has supported research students in the NICUs of The Hospital for Sick Children in Toronto, Ontario; Women and Infants Hospital in Providence, Rhode Island; and The Children's Hospital of Fudan University in Shanghai, China. Now, the team behind Artemis are currently collaborating with the Department of Health and the Pawsey Supercomputing Centre to implement the system in Western Australia, enabling Dr McGregor to bring Artemis back to her homeland of Australia.

"Pawsey is a very large component in this work as we will be constantly streaming a lot of patient data from critical care monitors, analysing it in real-time and demonstrating the potential to provide the information back to the various healthcare workers in real-time," said Dr McGregor. "The unique nature of the computing facilities at Pawsey are very important for my work."

> **"Big data, stream computing and the internet of things (IoT) have the potential to bring about new game-changing approaches to health care."**

"Big data, stream computing and the Internet of Things have the potential to bring about new game-changing approaches to healthcare," said Dr McGregor. "We can't expect healthcare facilities in urban, let alone rural and remote, communities to try and provide these new forms of healthcare and support all the technology."

THE SOLUTION

Artemis doesn't require any local installations at hospitals, urban and remote, to access this level of healthcare, as it runs entirely on a cloud-based system; which is where the Centre's Nimbus cloud service comes in.

"Using cloud services and supercomputing, we can offer these new ways to analyse your current health state as a service," said Dr McGregor, "and we can easily scale the service to accommodate all those that we need to monitor."

As the research requires data streaming to be timely and accurate, it could not operate within the boundaries of typical supercomputers that experience down-times. With cloud computing services like Nimbus, the Centre can continue its focus on supporting research, and strengthen international collaborations through projects such as the Artemis Platform.

THE OUTCOME

Artemis has been ground-breaking in its monitoring of premature children in neonatal wards internationally. The Artemis Platform cloud service can connect to bedside monitors, allowing a live stream and analysis of around 1,256 data points per second, per patient, amounting to roughly 600 MB per day (as of 2012).

Artemis has significantly reduced false-positive frequencies in data collection, can diagnose different types of neonatal apnoea with 98 per cent accuracy, preventing babies from receiving too much oxygen and reducing the risk of permanent eye damage.

This research will improve healthcare services for some of the most vulnerable lives in Western Australia's hospitals, and that's only the beginning.

"While we are initially demonstrating the approach in neonatology, this will show the potential for holistic healthcare," said Dr McGregor. "In my other collaborations, we are also working now to demonstrate how these same tools can be used for mental health as well, and also the management of physical health and wellness in space on long-range missions."

"I am demonstrating how this can be used for differential diagnosis through the use of new cognitive computing technologies."

"I think in principle the intent for this project is to start with neonatal intensive care units, but technically you can deploy this kind of system to any hospital bed," said Dr McGregor.

"Then, you start talking about a much larger system, which might, in the future, be something that needs to be developed."

Project leaders:	Dr Carolyn McGregor
Partner institutions:	University of Ontario Institute of Technology, University of Technology, Sydney, Government of Western Australia, Department of Health
System:	Nimbus cloud
Area of science:	Medical Data Analytics
Applications used:	Artemis Platform

A total of 194 projects made use of Pawsey infrastructure and support during the 2017-18 reporting period.

These projects represent the value Pawsey adds to the research community, through hastening the pace and broadening the scope of their work, and allowing them to tackle previously-intractable problems. The knowledge gained through applying HPC to these research problems is helping keep Australian researchers globally competitive. The impact of these projects, across industries, economies and communities, will be felt for years to come.

In addition to the highlights below and the case studies presented in this report, our full list of projects is available at **www.pawsey.org.au**

Snapshots Across Science

SEARCHING FOR THE GENETIC BASIS OF DISEASE

Dr Alison Testa, Harry Perkins Institute

The team at the Harry Perkins Institute are sifting through human DNA to determine which genes are most likely to be involved in causing disease. Using publicly-available gene expression data, they are searching for variants in DNA sequences, and mapping their relationship with tissue-and cell-specific expression and disease phenotype. Patterns arising out of this combination of data will enable the team to predict which genes are most likely to be involved in the expression of various diseases.

HELPING CONSERVE AUSTRALIAN KOALAS

Dr Amanda Barbosa, Murdoch University

Using next-generation sequencing and integrated population modelling, scientists are for the first time determining the prevalence and genetic diversity of mixed infections in koala populations from different regions across Australia. By understanding the role of mixed infections in koala population dynamics, scientists and biologists will better understand their implications for management and conservation of koalas into the future.

MODELLING THE PHYSICAL BASIS OF SLEEP-RELATED BREATHING DISORDERS

Dr Julien Cisonni, Curtin University

Sleep-related breathing disorders are common and can significantly impact on wellbeing. Episodes of sleep apnoea are caused by periodic upper airway obstruction, and cause problems from sleep disruption to cardiovascular disease. This project is investigating the fluid–structure interactions between breathing airflow and the upper airway soft tissues. Modelling variations in breathing and upper airway architecture is providing a better understanding of the physical mechanisms involved in obstructive sleep apnoea and snoring, to better design and perform surgical methods to alter the upper airway.

MAPPING THE EVOLUTION OF GALAXIES

ASKAP, CSIRO

The Widefield ASKAP L-band Legacy All-sky Blind survey (WALLABY) is now running on the Australian SKA Pathfinder (ASKAP). WALLABY is an extragalactic neutral hydrogen survey, recording microwaves from the hydrogen line across three-quarters of the southern sky. Early observations are now complete, and will allow astronomers to improve their understanding of the processes involved in galaxy formation, evolution, and the role of mergers and galaxy interactions. WALLABY is providing the largest and most homogenous hydrogen line sample of galaxies yet made, and is an important pathfinder project for SKA science.

THE DEVELOPMENT OF NON-COMMUNICABLE DISEASE

Dr Carol Wang, UWA

The Western Australian Pregnancy Cohort (Raine) Study has been developing over 27 years, helping determine how events during pregnancy and childhood influence health later in life. Almost 3,000 families have provided a plethora of phenotype and genetic data to this research program. This program is focused on the developmental origins of health and disease, particularly cardiovascular disease, diabetes and obesity, and filling knowledge gaps related to human health and development. Raine Study research teams have produced over 60 publications in the last three years that relied on genetic data analysed using Magnus. These include research findings on genetic variants associated with fetal growth, body mass index, growth trajectories, immune function, asthma, respiratory health, mental health and vision.

ENERGY FROM THE OCEAN

Carnegie Clean Energy and Dr Hugh Wolgamont, UWA

The Wave Energy Research Centre is working to find cost-effective ways to generate wave energy on an industrial scale. Carnegie Clean Energy is working with them and Pawsey to tackle the nonlinear modelling and computational fluid dynamics analysis required to design an optimum energy capture device that can withstand the extreme responses and loads occurring in extreme wave conditions. Bombora Wave Power are also working with Curtin and Pawsey to prototype and optimise the performance of their wave energy converter.

ALL SKY VIRTUAL OBSERVATORY – MWA ASVO

Professor Melanie Johnston-Hollitt, Murchison Widefield Array Consortium

The All Sky Virtual Observatory (ASVO) for the MWA project is a project to make the data produced by the MWA project available to researchers worldwide. Only raw, uncalibrated visibility sets are currently available, but there are plans to enhance this service to include calibration of visibilities, to reducing the barriers that non-MWA radio astronomers face in accessing, processing and reducing MWA data.

Publications

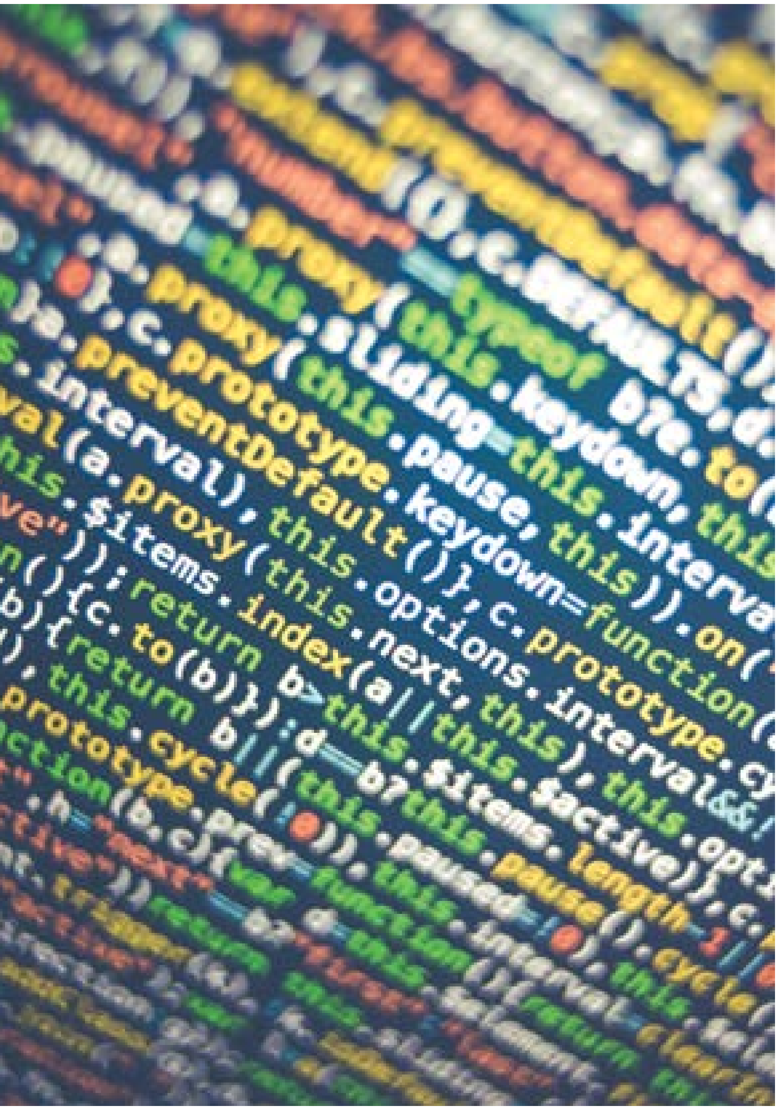

Research underpinned by Pawsey supercomputing and support services has been shared with the wider research community with 188 publications appearing in international journals and conference proceedings during the reporting period. These reflect the growing diversity of scientific research being accelerated through the application of HPC, as well as highlighting traditional areas of application and strength within Western Australia and our longer-term user base.

The radio astronomy, energy and resources base that Pawsey infrastructure significantly supports is well represented, with 42 papers published in astronomy and astrophysics publications, and 42 papers appearing in geoscience and geophysics publications. Physics research accounts for 20 titles, and bioinformatics is a recognised growth area at Pawsey with 15 publications this year.

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

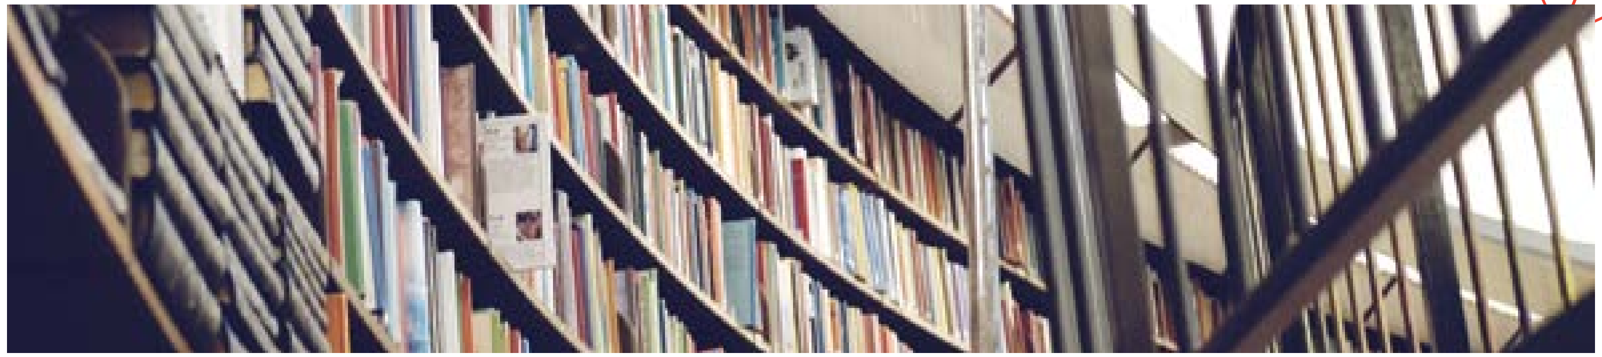

PUBLISHED WORKS ILLUSTRATING THE DIVERSITY AND IMPACT OF RESEARCH ARISING OUT OF THE CENTRE INCLUDE:

- **Radio Astronomy:** A large international collaboration established a robust method to detect and characterise the weakest solar emission features currently reported in the literature, using low radio frequency solar observations from the Murchison Widefield Array. **Their technique was published in the Astrophysical Journal,** established in 1895.
- **Geoscience:** A German and Australian team used small-angle x-ray scattering and molecular dynamics simulations to describe the forces that determine the colloidal stability of apolar nanoparticles, and explained the role of particle size and ligand shell structure, and why classical colloidal theory fails in these systems. **Their approach and insights were published in ACS Nano,** a leading forum for nanoscience and nanotechnology research.
- **Health Science:** An international team examined the possible causal effect of dairy consumption on systolic blood pressure and the risk of hypertension using instrumental variable analysis. **Their findings were published in BMJ,** one of the world's oldest general medical journals.
- **Bioinformatics:** A Czech and Australian team explored gene diversity across 18 wheat cultivars, identifying more than 36 million intervarietal single nucleotide polymorphisms across the pangenome. Their study provides insight into the genome diversity in elite wheat as a basis for genomics-based improvement of this crop, and **was published in The Plant Journal,** for the Society for Experimental Biology.
- **Bioinformatics:** UWA researchers reviewed the development of long-read sequencing and optical mapping, and the novel algorithms developed to more efficiently apply these techniques, and assessed their application in crop genomics for breeding improved crops. **Their assessment was published in Trends in Biotechnology**.
- **Chemistry:** German and Australian collaborators used *in situ* dynamic atomic force microscopy, supported by molecular dynamics and free energy calculations, to investigate the control of molecular self-assembly at solid-liquid interfaces. **Their findings were published in the flagship Journal of the American Chemical Society.**

Growing knowledge and expertise to advance research possibilities for the world

Voices of Science

IF YOU THINK OUR COMPUTERS ARE EXCITING?

WAIT UNTIL YOU MEET THE PEOPLE THAT BRING THEM TO LIFE.

OUR BOARD

Governance is exercised through the **Pawsey Supercomputing Centre Board of Management**, who provide leadership and oversight to the organisation. The Board's key responsibilities are to set the strategic direction of Pawsey, ensure the delivery of its objectives, and support staff to uphold Pawsey's core values. The Board is made up of representatives from Pawsey, our partner organisations – CSIRO, Curtin University, Edith Cowan University, Murdoch University, and The University of Western Australia, as well as independent members and an independent Chairperson. Having a Board, comprised of people with a vested interest in the future of Australia's scientific competitiveness, allows Pawsey to continue to strengthen its role as an enabler of Big Science outcomes and best practice for supercomputing.

OUR STAFF

The Pawsey team is made up of more than **45 staff members** who each have their own unique area of expertise, including specific scientific domain supercomputing, data, visualisation and operation. With diverse national and international backgrounds, the team is extremely experienced and knowledgeable in all aspects of high-performance computing (HPC) and data sciences.

Pawsey is continuously looking at ways to develop staff whilst expanding its team of world-renowned experts. During the reporting period, Pawsey welcomed more expertise in the Supercomputing, Data Services and Visualisation teams. The Centre also appointed Mark Stickells as Pawsey's new Executive Director.

Mr Stickells comes to Pawsey as an advocate of science and innovation and has extensive experience in the resources sector. He has played a major role in developing networks across Australia and overseas, to create opportunities for collaboration and partnership benefiting researchers, industry and ultimately the nation.

OUR RESEARCHERS

Our researchers have lasting and rewarding relationships with the Pawsey staff and Centre. Ongoing support and the mutual exchange of training and feedback create new possibilities for research in all areas of science.

Pawsey is currently broadening the Centre's capabilities to deliver services and expertise to increase both the number of scientific outcomes, and national uptake of its facilities. By demonstrating the capability of Pawsey supercomputers our team uncovers new possibilities for researchers in all aspects of science.

Our training programs support researchers in developing methods for incorporating the power of supercomputing into their existing areas of research and creating powerful results for the broader scientific community.

MERIT ALLOCATION COMMITTEES

Allocations to Pawsey and National schemes are executed by the **Merit Allocation Committees.**

The purpose of these committees is to allocate supercomputing resources and review project applications. All applications go through a rigorous process, assessing them against a set metric, with the input of a diverse committee. There are currently three allocation committees: Pawsey Partners, Energy and Resources, and the National Merit Allocation Scheme (NCMAS).

Each committee is broken into two main subsections: the Science Review Committee, which includes members from partners' institutions – for partner applications; or respected individuals who represent research across Australia – for general project applications. The second subsection is the Technical Review Panel, comprised of members of the supercomputing team familiar with the technical requirements and feasibility of each application.

In 2017, new Terms of Reference were agreed upon for the aforementioned Pawsey Partners and Energy and Resources Schemes. These have provided a clear framework for the chair of each Merit Allocation Committee to use to guide reviews, meetings and general committee standards.

In addition to this, two new committee members were introduced to the Energy and Resources Committee, bringing with them their expertise in plate tectonics and exploration geophysics.

The Merit Allocation Committees at Pawsey continue to provide a thorough review process and with annual and bi-annual position placements, the committees ensure the brightest minds and most appropriate experts are attracted to this essential role in Pawsey's operations.

Pawsey is About People as Much as Petabytes

John Langoulant AO, Chairman of the Pawsey Board

John is well known across government and business, having held a variety of prominent CEO positions for almost 20 years. He now consults to the education, banking and government sectors, and chairs a number of public, corporate and not-for-profit entities. He has been Pawsey's Chairperson for five years.

"I'll take on a Board because it's a challenge, because it takes me into new areas, and because it's likely to make a real contribution to diversify the economic base of Western Australia. Pawsey ticked all those boxes for me, and it's a very significant organisation in terms of the development of science and research across the State. Joining its Board was a very attractive proposition."

"I have an ability to work with people from diverse backgrounds, who in many instances have competing interests, forming Boards that through the right executive team can create and support good operating environments," explained John. "I'm a 'jack of all trades' really. I didn't bring any specific skills in HPC to the Pawsey Board, but over the last five years I've developed a real understanding of what is important in this space and what is not."

John has discovered that the Board of a supercomputing facility is a little different to others: "The level of technical oversight and guidance is higher than you might expect, especially with big developments, like supporting the SKA pathfinder projects. We've had to manage relationship issues arising out of technological roadblocks, especially during Pawsey's startup timeframe."

"Normally Boards only concern themselves with higher levels of policy and strategy, although we've also been fairly successful in that sphere. I'm proud of how we've managed all of the relationships within our unincorporated joint venture, especially as the partners are large organisations in their own rights. For example, bringing all of the staff out of the universities and unifying them under the umbrella of CSIRO was a challenge that caused a lot of uncertainty and caution at the time, but it was a smooth transition and materially simplified Pawsey's employment arrangements."

"I'm also proud of the action we took in engaging a new Chief Executive at a time where we saw a new direction and strategic focus for Pawsey's operations. Most recently, I'm particularly pleased that through our efforts, the Commonwealth has recognised the need to safeguard and strengthen Australia's Tier 1 supercomputing facilities within the national research infrastructure, and has acted to provide the funding base for that."

While world-class research and supercomputing is 'business as usual' for many people around Pawsey, John has a refreshing enthusiasm for its novelty in other arenas. "I love being able to say that I'm the Chair of a supercomputing facility! Given my background, people don't expect it, and they are intrigued. They always want to know what it involves, and if it's super-secret science stuff! It's great in that it gives me a chance to raise awareness across the various sectors that I work in that we do have a supercomputing facility right here in Kensington, doing amazing work."

John can see that Pawsey and the science communities around it are increasingly recognised as excellent, innovative and progressive. "Australia is seen to be stepping up as a science and research jurisdiction internationally, and we need to reach out to that global community to reinforce where Pawsey and our research collaborators sit in that landscape."

The challenge will be to consolidate capability and expertise moving forward. Just as astronomers are clustering in Western Australia associated with the SKA project, John is hopeful that more people skilled in HPC disciplines will graduate from Pawsey's partner institutions, and also be attracted from further afield, to work here in support of Australia's data-driven research and business interests. "People think the technology is difficult, but amassing and maintaining that human capital in this space will be the most important thing moving forward."

Professor Linqing Wen

Faculty of Engineering and Mathematical Sciences School of Physics, Mathematics and Computing The University of Western Australia (UWA)

Dr Linqing Wen is one of Australia's top astrophysicists. She is building and testing software to detect the tiny cosmic ripples generated by gravitational waves.

ABOUT PROFESSOR LINQING WEN

Professor Wen grew up in Fujian province, in the deep south of China. She studied condensed matter physics at the University of Science and Technology of China, before travelling to the USA and enrolling at Pennsylvania State University.

She moved to study astronomy for her doctorate at Massachusetts Institute of Technology. This is where Rainer Weiss, a member of the astrophysics division, with two of his colleagues would later go on to win the Nobel Prize for recording the first gravitational wave. It was this influence that led her to begin her research career in gravitational waves at California Institute of Technology.

She then moved to Germany to continue her career at the Max Planck Institute for Gravitational Physics, the largest hub for gravitational research in the world. She is now based at The University of Western Australia (UWA) and is a Chief Investigator in the **ARC Centre of Excellence for Gravitational Wave Discovery.**

WHAT DREW HER TO SCIENCE?

A high school maths teacher, by the name of Xun Li, first seriously inspired Professor Wen's interest in science. "He made everybody want to be a mathematician. He had just started teaching high school; we were his first batch of students. Now he's the deputy minister for education in the province. He's a legendary person."

"Part of the reason I became an astrophysicist is because where I grew up, we had a dark sky and I happened to be living on the top floor with a nice big balcony. I'd often spend nights watching the stars," Professor Wen said.

"When I was invited to a workshop in Perth, I realised this was a place I could stay. A place I could live. We took a tour to Gingin, and the sky was clear and free of light pollution, just like my childhood home. The tour of the night sky was fantastic – we were looking at Magellanic Clouds that you can only see in the Southern Hemisphere," she said.

Detecting the Smallest Signals with the Biggest Computers

"Not only that, the Square Kilometre Array (SKA) project was beginning, and they had the resources for astrophysics and gravitational waves – everything in science that I like and fits my training. As a bonus the nature was really good, and the climate was just like California!"

RESEARCH WITH SUPERCOMPUTERS

Modern stargazers use supercomputers, radio telescopes and advanced laser technology to probe the infinity of space and measure natural phenomena that we are only just beginning to understand. Since the first gravitational wave was detected in 2015, research in the field has plumbed the depths of black holes, checked the accuracy of Einstein's general theory of relativity, and helped us understand how the Universe fundamentally works.

To improve detection methods for gravitational waves, they need to be tested against large-scale simulated data as well as detector data from previous science runs. This requires an incredible amount of computing power, provided by the Pawsey Supercomputing Centre's Advanced Technology Cluster, Athena, as well as the Magnus supercomputer.

AREA OF SCIENCE

Astrophysics and Astronomy
Gravitational Wave Physics
Signal Processing
Algorithm Development

REAL WORLD SOLUTIONS

Professor Wen's career has run parallel to the construction and opening of the **Laser Interferometer Gravitational-Wave Observatory** (LIGO), now the most important tool for gravitational physics. Two LIGO detectors detected the first gravitational wave – caused by the merger of two black holes, each roughly 30 times the mass of our sun.

Professor Wen and her team are designing and testing automation software to make it easier to detect gravitational waves, improving capability in both their rapid detection and early warning. She is fine-tuning search pipelines for two gravitational observatories, both LIGO in the US and Virgo in Italy. The improved detectors are painstakingly accurate and responsive: at their most sensitive state, the LIGO detectors are accurate enough to detect a change in space 10,000 times smaller than a proton, across the four-kilometre observatory. Through Wen's work, they can also respond fast enough to tell optical and radio telescopes to look in the right part of the sky to catch any electromagnetic radiation from a gravitational wave event.

Professor Richard Sandberg

Chair of Computational Mechanics
Department of Mechanical Engineering, University of Melbourne

Professor Richard Sandberg is saving the planet and our wallets through his research in the efficiency of turbines. He is examining the effects turbulence has on aircraft engine efficiency, helping to create greener and cleaner methods of air travel.

In Australia alone, approximately 8.2 billion litres of aviation fuel are consumed per year. A single return flight from Australia to Europe produces over 4.5 tonnes of carbon emissions. As such, even a one per cent reduction on these figures could save billions of dollars and tonnes of carbon emissions each year.

ABOUT RICHARD SANDBERG

Professor Sandberg is a world-renowned academic in the space of 'high-fidelity computational fluid dynamics' – this combines high-order accurate numerical methods with high-performance computing. Through this work he is helping the world better understand turbulent flows, what they are, how they affect engine efficiency and noise generation.

It was his passion for aircraft engines that brought Professor Richard Sandberg to Australia. He relocated from the UK with his family three years ago to take up the position of Chair of Computational Mechanics in the Department of Mechanical Engineering at the University of Melbourne.

Prior to moving to Australia, he was Professor of Fluid Dynamics and Aeroacoustics in the Aerodynamics and Flight Mechanics research group at the University of Southampton. Professor Sandberg also headed the UK Turbulence Consortium, a group of UK institutions.

WHAT DREW HIM TO SCIENCE?

Richard's childhood ambition was to be a pilot. However, this changed the first time he saw an aircraft at close quarters. "Once I got up close to an aircraft, I realised I was more interested in how the engine worked than flying it, and over time, aircraft engines became my passion."

This made the decision to pursue aerospace engineering at the University of Stuttgart easy and from there, he went on to complete a PhD in Aerospace Engineering in the USA.

Now as a teacher his passion for science extends to the opportunity to shape the thinking processes of the next generation of bright young engineers and scientists in Australia who are the researchers of the future.

"I don't teach students what to think, but how to think, which means I can encourage students to better understand how they can solve problems with the high performance computers."

Aircraft Engines Propel Professor Sandberg

RESEARCH WITH SUPERCOMPUTERS

Professor Sandberg wants to help the world better understand turbulent flows and how they affect efficiency and noise generation in an industrial setting. His secret weapon to achieve this has been the use of some of the world's fastest computers. His research into the flow of gases in aircraft engine components, such as low and high-pressure turbines, have enabled him to advance the understanding of turbine technology that would have been almost impossible to achieve in the past.

"We rely on supercomputers because turbulence is a complicated problem. I'm focused on solving equations that are very non-linear and complex that we have not been able to find general solutions for to date. However constantly evolving technology means that we can now use high-fidelity simulation approaches to conduct world-leading research for industrially relevant problems - and I have no doubt that in another decade, we will be able to study even more involved problems."

AREA OF SCIENCE

Turbulent Flow
Turbo-machinery
Computational Fluid Dynamics
Numerical Modelling

Professor Sandberg has developed his own computer program for use on supercomputers to gain better understanding of turbulence. He is able to use around 100 million computer hours per year which is equivalent to using a desktop or laptop computer for about 3,000 years.

He believes that Australia's continual investment in supercomputers is paramount to ensure that it keeps pace with a range of other leading scientific communities around the world.

REAL WORLD SOLUTIONS

Accurately predicting and modelling the gas flows in turbine engines has the potential to enable engineers in the commercial and public domains to design the next generation of green engines, making plane travel and power generation more affordable and more efficient, and leading to fewer emissions and less noise.

Professor Igor Bray

Head of Physics and Astronomy, Curtin University

Professor Igor Bray is a world expert on atomic collisions. His passion for studying atoms is allowing him to solve real world problems across the globe.

Building a Career with the Universe's Building Blocks

ABOUT PROFESSOR IGOR BRAY

With a career spanning roughly 30 years, his work features in some of the most frontier-expanding science in the world. Analysing mighty supernovae, treatment of childhood brain cancer, and the engineering behind fusion power all require his team's work on atomic and molecular collisions.

After graduating from the University of Adelaide in 1986 with the PhD paper 'Gravitational Lens Effect of Galaxies and Black Holes', he began his research career at Flinders University of South Australia, before moving to Perth, Western Australia.

He now teaches and researches as Head of Theoretical Physics and Head of Department of Physics and Astronomy at Curtin University.

WHAT DREW HIM TO SCIENCE?

"Maths and science were innate, I just loved problem solving – I still do. I'm a compulsive problem solver. For me, if something breaks it's an exciting moment because I wonder how I'm going to fix it. It's a mindset."

"I love the symmetries of mathematics and the application to real-world problems. I realised while I was at university that, as much as I love mathematics, I really wanted to apply it to real things. The transition from pure abstract to something practical and real is very interesting," he said.

An impassioned and engaging speaker, as well as a highly-regarded scientist, Professor Bray brings a 'real world' approach to all areas of his work. He often perforates his scientific findings with near-philosophical comments on the scientific method's application to greater life problems, and isn't afraid to bring out his sense of humour – including the odd Monty Python reference here and there.

RESEARCH WITH SUPERCOMPUTERS

Professor Bray's research uses supercomputers to simulate and test complex particle interactions. His team has been using the Pawsey Supercomputing Centre facilities since it's foundation in 2000 employing Pawsey's systems to test their theories with different targets and atomic projectiles.

"We've used Pawsey for everything. We've been utilising the Centre at its capacity throughout its existence. As the centre became more powerful it opened up more possibilities. Now we've been doing molecular work over the past five years or so," said Professor Bray.

"Our careers are critically reliant on supercomputer technologies."

AREA OF SCIENCE

Physics and Astronomy

REAL WORLD SOLUTIONS

While Professor Bray's work spans multiple areas of science and projects, some of his most recent work has contributed to:

Fusion Energy

His work features in **the International Thermonuclear Experimental Reactor**, set to be the world's first power-supplying fusion reactor. Once completed, it will use sea water as fuel, providing a clean and readily-available energy source.

Treating Cancer

Professor Bray's team has also developed the most technologically advanced forms of cancer treatment used to treat childhood brain cancer. It involves firing protons to destroy tumours, without causing damage to the body associated with traditional X-ray radiation treatments.

Particle Mapping

His team developed the convergent close-coupling computational approach to atomic and molecular collisions. It is the most accurate computational theory to date for particle collisions and involves mapping particles as they collide with each other on an atomic level.

To meet the demanding levels of computing power required for calculating particle collisions, Professor Bray's team have extensively used Pawsey Supercomputing Centre. With the help of the Magnus supercomputer, one of the Southern Hemisphere's most powerful public research supercomputers, the team have been able to accurately calculate collisions.

THE KEY TO SUCCESS

There are many doors open to scientists outside of academia. Reflecting on his own life he sees his career as successful because of two main traits: excellence and team work.

"It begins with research excellence. In academia, research excellence opens doors. ...The ability to help and work with others complements the credibility associated with research excellence."

"...It's a missed opportunity to compartmentalise the scientists to just their science. Their skillset applies much more broadly to build a better society. Science is extraordinarily creative.
It's extraordinarily emotional. I love what I do..."

Ashley Chew
Supercomputing Systems Administrator

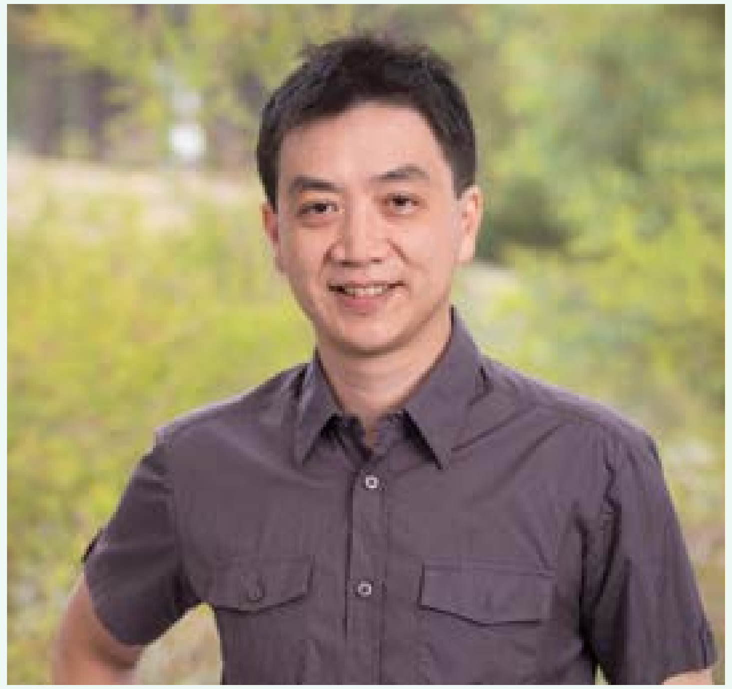

From the Beginning...

Ashley has more than 18 years' experience in the supercomputing industry, beginning as a system administrator/programmer at The University of Western Australia. He joined iVEC in 2011, when he was drafted into commissioning and looking after the Fornax supercomputer at UWA, as part of the original 'Pawsey Project'. Fornax was a precursor cluster that preceded the commissioning of Magnus, Galaxy and Zeus, and the development of the Pawsey Supercomputing Centre as we know it today.

Ashley is now a supercomputing system administrator at Pawsey, where he uses his years of Linux and Unix experience to look after the high-performance computing (HPC) facilities.

"Working at Pawsey means working in teams – the number of different skills required means people are inherently reliant on others to achieve common goals," explained Ashley. "It's amazing to work with the latest technology, but the buzz comes from seeing the researchers trying to solve problems using these HPC systems that just weren't possible to approach before."

"It's really rewarding to get to know our users, meet face-to-face, and see outcomes of research that are beneficial for everyone. It's unfortunate that as our user base has expanded, many of them are offsite and the interaction is necessarily more remote. I've always enjoyed having an open-door policy, but it's hard for many of our users to just drop in. The door is still open though!"

Looking after the users is obviously Ashley's first priority. His average day always starts with addressing any user questions or issues that have arisen. Only then does he move to project work aimed at enhancing the supercomputing systems he manages, to make them even more usable for the researchers.

His career highlights centre around providing a first-class computing facility for Australia. "I'm very proud to have been a part of the original Pawsey Project, specifically Pawsey 1B aka Fornax, because we showcased that Western Australia really could deliver a world-class HPC system. And Fornax became the precursor to the main systems in place at Pawsey now. More recently I'm also proud to have been a part of the Australia-wide Nectar Cloud program, helping to provide flexible, scalable computing power and related infrastructure online to researchers across Australia. Users can store, access and run their data any time – from their desktop."

"I think the technology will eventually converge, where HPC, cloud and data will be provided by one platform, and where resources can be allocated dynamically based on user demand. Given that we consider three-year-old systems 'old' here at Pawsey, I can guarantee that the Pawsey Supercomputing Centre will be nothing like it is today within ten years!"

> "It's amazing to work with the latest technology, but the buzz comes from seeing the researchers trying to solve problems using these HPC systems that just weren't possible to approach before."

VOICES OF SCIENCE – STAFF

Jenni Harrison
Director of Strategic Projects and Engagement

Both Sides of the Fence

Jenni understands what it is like to be a researcher accessing digital support services, having spent several years since her PhD in research, both within academia and industry. It informed her subsequent work applying novel digital technology at NHS Education for Scotland, and as the eResearch Policy Lead at the Ministry of Research Science and Technology in New Zealand.

Her considerable expertise on the delivery of collaborative Information and Communication Technology projects and working with stakeholders in research, management and government has been put to work at Pawsey over the last seven years, originally as Head of Data, but more recently as the Director of Strategic Projects and Engagement.

"I work at a strategic level across all areas of Pawsey, and interact with a wide range of people externally, which makes my role both challenging and really interesting," Jenni revealed. "I help ensure that Pawsey supports researchers as effectively as possible. I'm not quite comparing supercomputing to trains, but it's like managing any public utility – when it works well, most people are completely unaware of its presence, but they're still dependent on its service to be able to function effectively."

Depending on priorities and projects, Jenni can be found anywhere across Australia, ensuring strategic projects progress smoothly.

"As a team, it's great to feel like we're making a difference, supporting others to make great scientific achievements. One of my roles is to co-direct **the Astronomy Data and Computing Services project.** We're working to provide astronomy-focused training, support and expertise to allow astronomers to optimise their data and computing infrastructure for ease of use. The aim is to allow the next generation of researchers to spend more time on astronomy challenges, and less time trying to figure out how to use the facilities."

When looking forward, it is easy to take for granted how much has been achieved already, and how relatively recent some things have been. Jenni remembered: "About six years ago I asked my team if we could bring a 'Big Data Week' to Perth, to raise awareness of the opportunities with large scale data in research. At the time there were no organisations in the Southern Hemisphere participating in this global initiative. Thanks to the tireless efforts of several teams within Pawsey,

'Data Science Week' **(www.datascienceweek.org)** is now an annual collaborative week-long festival of events, all related to data science."

Things are still changing: Jenni suspects that connectivity with other HPC centres and consortia is likely to grow in significance, along with running software from cloud-based environments, while deep learning, artificial intelligence and quantum computing are likely to have significant longer-term impacts. Jenni is clear about her role in all of this moving forward: "I'm working to make sure Pawsey is well placed to support Australian researchers working as part of both national and international research projects, and trying to make sure the individuals I support have as seamless (and happy) an experience as possible when working with us."

> "As a team, it's great to feel like we're making a difference, supporting others to make great scientific achievements."

VOICES OF SCIENCE – STAFF

Mark Gray
Cloud Leader

Clouds:
Both Atmospheric and Electronic

Mark has a history with clouds. Earlier in his career he worked at NASA Goddard Space Flight Centre as a research software engineer, working in global cloud analysis. After returning to Australia he was still focused on processing system development for large-scale satellite data, assisting the Integrated Marine Observing System and the Terrestrial Ecosystem Resource Network and doing research in shallow water ecosystem analysis.

He joined Pawsey (then iVEC) in 2012. "At Pawsey I could apply my research development and operations skills, as well as my experience with high volume datasets and advanced data techniques," he explained. He is now the Cloud Lead at Pawsey, managing the research cloud service Nimbus and the team of system administrators who operate it. Mark also works directly with researchers and manages project work as needed, particularly projects of key strategic importance to the Centre.

"The best thing about Pawsey is getting to work with a diverse group of researchers tackling critically important research across all science sectors," he enthused. "Our researchers and collaboration partners are achieving science that just can't be done anywhere else outside the Tier 1 HPC facilities in Australia. We're helping them take their research up to that next level. Seeing our researchers undertake truly world-class research and achieve real world impact is the most rewarding aspect of my job."

"For example, recently I had the opportunity to help the Artemis team get started in Western Australia. Artemis is a Neonatal Intensive Care Unit bed monitoring system, with a cloud-based analytics system. We're establishing a prototype in Western Australian hospitals, using Pawsey resources as the processing engine. It's a new way for the hospitals, the Department of Health and Pawsey to collaborate. The interaction is all about listening to each other – when the science, motivation and Pawsey services intersect, we have the opportunity to do great work."

Mark's work also includes leading, developing and delivering training materials at Pawsey. "Often researchers come to us because they've reached roadblocks in their research that can only be solved computationally, and it's the first time they've had to deal with this. It's quite normal for teams to come to us with little or no background in the use of HPC infrastructure," Mark explained.

"**Our training programs** ensure Pawsey users are able to make full use of our facilities, and gain an appreciation of how they can extend or improve their workflows at Pawsey. Our training is free and covers a range of teaching tools – online, in-person, self-study – depending on user need. Our users are a diverse group of people with a wide range of needs and experience, so we provide diversity in both infrastructure and approach so that we can meet the researchers where they are."

The effort goes beyond training to ongoing research support, as Pawsey staff help users bridge the knowledge gap to adopt scalable computational approaches to tricky problems.

Mark believes: "HPC is in a period of dramatic change, moving towards providing a diversity of resources that can match an increasingly diverse range of research needs on much shorter timescales." So the current diversity in training and support approaches is likely to be a feature of Pawsey operations for a long time to come.

"Our training programs ensure Pawsey users are able to make full use of our facilities, and gain an appreciation of how they can extend or improve their workflows at Pawsey."

William Davey
Systems Administrator

From Software to Hardware...

William started his career in software development, developing and maintaining both back end and front end applications at iiNet. His first exposure to the Pawsey Supercomputing Centre came soon after, when he became a computer science student at Murdoch University. The USA-based **Student Cluster Competition** is designed to introduce the next generation of students to the HPC community, and Pawsey took the first Australian student team, including William, to the competition.

William demonstrated outstanding skills as part of that ground-breaking team, and was favourably remembered when he applied for a position at the Centre in 2016. He is now a system administrator within the supercomputing team.

"I joined Pawsey as I wanted to try something new," William elaborated. "I'd had plenty of experience working with hardware, but not in a professional capacity, and not at the scale we have here at Pawsey."

His role involves close communication with a number of other teams at Pawsey, as many of the systems are interconnected and rely on one another. Working on the supercomputing user portal means he also acts on experience other Pawsey teams have gained regarding user interaction and management. William also works directly with research groups, and investigates any problems they experience using Pawsey systems, "to better the way Pawsey can do things. It's all about working together to achieve our common goals".

William successfully co-developed and released a user and project management web application – Origin – and has worked on the Puppet project to evolve the way systems and services are managed at Pawsey.

"I don't have a 'normal' working day," admitted William. "Most days start with checking the status of our systems and services. If there are no major issues then it's usually onto project work that can involve research, documentation, configuration and programming, or any number of physical tasks on the systems themselves."

"It's really fulfilling to see the amazing research projects that are facilitated by the resources and services that the Centre provides. A lot of the mathematics and computational work behind the research projects is over my head, but it's great that we hold 'Pawsey Fridays', where we can hear firsthand from the researchers themselves what is being achieved."

"Researchers just want to focus on their particular domains, so with respect to HPC, they'll usually take the path of least resistance to get their work done. With the current upsurge in cloud computing services as their cost and ease of use make them increasingly viable options, we need to learn from these companies and continue to make our HPC facilities more accessible. I want Pawsey to be an innovator in the way our systems are run and how we support our users."

"It's really fulfilling to see the amazing research projects that are facilitated by the resources and services that the Centre provides."

Providing access to new technologies and techniques to stay at the cutting-edge of scientific research

Tools for Change

PAWSEY HOUSES A CLUSTER OF WORLD-CLASS CAPABILITIES **INCLUDING HIGH-PERFORMANCE COMPUTERS, BIG DATA SERVICES, CLOUD AND VISUALISATION RESOURCES AND AN EVER-GROWING POOL OF APPLICATIONS AND WORKFLOWS. THIS INFRASTRUCTURE IS MANAGED BY A TEAM OF STAFF WHO ARE EXPERTS IN THEIR FIELD.**

PROVIDING CUTTING-EDGE RESOURCES FOR RESEARCHERS ACROSS MANY SCIENTIFIC DOMAINS IS A CONSTANTLY EVOLVING GOAL. PAWSEY INFRASTRUCTURE, ALTHOUGH STATE-OF-THE-ART AT THE TIME OF INSTALLATION AND COMMISSIONING, IS CONSTANTLY REFINED, DEVELOPED AND EXTENDED TO KEEP PACE WITH INCREASING USER DEMAND AND AMBITION.

DURING THE REPORTING PERIOD, NEW TECHNOLOGIES WERE TESTED AND IMPLEMENTED. UPGRADES WERE MADE TO EXISTING SUPERCOMPUTING SYSTEMS, NEW SOFTWARE WAS DEPLOYED AND CODES WERE DEVELOPED TO IMPROVE EFFICIENCY, AND REMOTE AND FLEXIBLE ACCESS SERVICES WERE EXPANDED.

Systems and Services

MAGNUS

One of the most powerful public research supercomputers in the Southern Hemisphere

GALAXY

Real-time computing for Australian radio astronomy

Both **Magnus** and **Galaxy** were upgraded to the latest Cray Operating Systems, to extend their useful lifespan while increasing their stability and reliability. This upgrade ensures Pawsey researchers have access to the latest generation of compilers and development tools.

Magnus, Pawsey's flagship system, is now due for a capital upgrade to increase its capacity and efficiency, to maintain Australia's international competitiveness. Part of the **$70 million upgrade** to the Centre infrastructure will improve high performance computing and data science capabilities in line with the ever-increasing expectation and demand from our researchers. With funding secured, planning for this major development is now underway. In combination with **the upgrade to NCI's Raijin supercomputer**, Australia's Tier1 HPC facilities will continue to provide the technological edge researchers and industry require to stay at the forefront of the innovation curve.

MACHINE LEARNING AT SCALE:
Bringing Technology to Life

Machine learning is a statistical approach giving computers the ability to 'learn' from data and better predict outcomes. Without being explicitly programmed, computer algorithms develop and improve their predictive response through exposure to both training and validation data sets.

Pawsey staff have been working with Cray to develop machine learning at scale. The novel approach is working to split the problem into multiple nodes, with the aim to provide researchers with future systems that can fulfil all of their large-scale machine learning requirements.

Based on this work, Pawsey and Cray are developing and publishing guidelines and best practices for scaling machine learning workloads. Staff are now working with Curtin University to implement these machine learning workflows using real world data on HPC systems.

Following the procurement of the Centre's newest system Athena, and in response to a significant demand from researchers, machine learning support at Pawsey is steadily increasing.

SPECIFICATIONS

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

GALAXY
- Peak performance in excess of 200 TeraFLOPS
- Real-time system for the 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

200%
Magnus was fully allocated through 2017-18. Its allocation application was over-subscribed by 200 per cent

270
MILLION CORE HOURS AWARDED

77
NCMAS APPLICATIONS
(National Computational Merit Allocation Scheme)

62
PARTNER APPLICATIONS

34
ENERGY AND RESOURCE APPLICATIONS

ATHENA

Cutting-edge technologies to keep researchers ahead

During the reporting period the Centre's new advanced technology cluster 'Athena' was made available. It was procured to test new and novel technologies, and to ensure Pawsey researchers are well equipped to engage in emerging computational techniques such as deep machine learning.

In 2017 the first group of 14 early adopters were provided access to Athena, and embarked on their projects with the assistance of a dedicated Pawsey expert. The results of their interactions with the new architecture has helped inform the capital upgrade for Magnus.

Athena has now been incorporated into the upgraded mid-range cluster, Zeus, and the old nodes are dedicated visualisation or Graphics Processing Units (GPU) partitions, allowing researchers faster access to resources.

CODE OPTIMISATION:
Increasing Efficiency

A new profiling tool from ARM called 'Forge' is now being used to highlight bottlenecks in code, allowing researcher's code to be refined and improve its efficiency before processing. On-site training and a hands-on workshop on Forge were run for both Pawsey staff and research partners as part of the International GPU Hackathon hosted by Pawsey with the support of the USA Department of Energy's Oak Ridge National Laboratory and Nvidia Corporation.

Identifying areas where software is most active allows these areas of code to be targeted and re-assessed, maximising speed increases with minimum effort. By focusing on bottlenecks in code, significant reductions in time-to-solution can be achieved.

NUMBER OF CORES AVAILABLE IN ATHENA: 5,428.

ALLOCATION: 75% ALLOCATED VS 90% USED.

SPECIFICATIONS

- 80 nodes with 64-core Intel Xeon Phi 7210 processors with a 100 Gb/s OmniPath interconnect
- 11 nodes with four NVIDIA Tesla P100 GPUs with a 100 Gb/s InfiniBand

ZEUS

Pawsey's HPC stepping stone

In 2018, Zeus was upgraded to become Pawsey's mid-range cluster. It is now used by research communities such as bioinformatics, which have large compute requirements but cannot take advantage of the Cray's high-speed network. Zeus now provides approximately 20 million more core hours each year to researchers.

NUMBER OF CORES AVAILABLE IN ZEUS: 2,576.

ALLOCATION: 25% ALLOCATED VS 87% USED.

SPECIFICATIONS

- 92 nodes with 28-core Intel "Broadwell" Xeon processors
- 128 GB RAM per node
- 100 Gb/s OmniPath interconnect

NIMBUS-CLOUD

An integrated data-intensive infrastructure

The cloud service 'Nimbus' was launched in mid-2017 as a free service for national researchers requiring flexible access to compute resources. In 2018, Pawsey invested in GPU expansion. GPUs were added to provide a new set of functionalities to researchers, such as the capability to support machine learning and increase usability when carrying out development work. Twelve NVidia GPU nodes were added offering researchers up to 1.3 PetaFLOPS of deep learning performance.

Pawsey staff are now developing the use of containers, allowing researchers to easily package software developed on a local workstation in a container, to then be deployed on computing infrastructure.

DEEP MACHINE LEARNING IN THE CLOUD

A new product from IBM called 'Conductor' has been evaluated to be used with existing analytics and filesystems. The introduction of Conductor will support deep (machine) learning and the Jupyter notebook; an open-source web application that allows the creation and sharing of documents containing live code, equations, visualisations and narrative text.

VCORES AVAILABLE: 4752

VCORES ALLOCATED: 4261

SPECIFICATIONS

- 3000 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

WORLD CLASS DATA STORAGE

Projects using the Pawsey Supercomputing Centre's resources often produce vast amounts of data. Not only does this data need to be stored, but scientific research standards for computing place increasing demands on researchers to curate and share their data, as well as to ensure the value, accuracy and longevity of their data sets. Pawsey staff help researchers to maximise the potential of their vast data sets through a number of support services while also providing storage for researchers across Australia.

PROACTIVE SUPPORT TO DATA GROWTH

In an ongoing commitment to the MWA and CSIRO ASKAP Science Data Archive, Pawsey has expanded its disk capacity, doubling the file-system size available to the projects through the Hierarchical Storage Management system. The radio astronomy projects are expected to generate four times the current data level over the next 12 months, making the technical upgrade critical to Pawsey being able to effectively support the project.

In further actions to support the MWA, the Data Team have prepared for the use of new technology, the Intel Skylake CPU, to handle large data sets.

DATA RESILIENCE BACK UP

As part of a wider collaboration with NCI, Pawsey is working toward an offsite data resilience back up. The result will be a full offsite back up of all Pawsey researchers' data. Currently the project is in the early stages of planning and development and will continue to be executed in 2018-19.

SPECIFICATIONS

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

VISUALISATION

Data visualisation is the process of applying advanced algorithms and computer graphics to represent data in other ways. This allows researchers to better interpret their raw data by translating it into more meaningful or recognisable forms.

Applying visualisation techniques to complicated datasets can require specialist hardware, such as high-end graphics cards for handling large datasets in real time, novel display technologies to fully exploit the human visual system, and user interface devices to facilitate the interaction. Pawsey has increased the capabilities of visualisation services with multiple updates to the infrastructure and the introduction of new software to handle these requirements.

As part of a commitment to make high quality resources more readily available to researchers, a new web-based remote visualisation service has been trialled. This approach allows visualisation of large data sets generated from supercomputers without moving data across networks. The associated high-end Windows workstations are also configured for remote access, enabling users to access this resource from their offices. These upgrades are leading to exciting new research and development using virtual reality technology to visualise scientific data.

ORIGIN: HANDING CONTROL TO RESEARCHERS

In 2018 Pawsey's web-based portal 'Origin' was deployed, allowing researchers to view their allocations across both computing systems and storage, and manage researcher access to existing projects. This has enabled smoother activation of projects on Magnus, Galaxy and Zeus and reduced traffic through the helpdesk.

The Building

The Centre is a purpose-built facility housing supercomputers, data storage and associated infrastructure. It is located in Kensington, Western Australia, on CSIRO-owned land adjacent to the Australian Resources Research Centre, approximately six kilometres from Perth's central business district.

The Centre incorporates **a number of specific design features to minimise its environmental impact,** and specifically employs best practice technologies to reduce energy usage. As a purpose-built structure, the Centre's design anticipates the expanding power, cooling and physical footprint requirements of the next generation of supercomputers. Scalable cooling and electrical services enable flexible supercomputer expansion within the 1,000 square metre technical floor space.

Pawsey has a high speed data connection to **AARNet**, Australia's National Research and Education Network (NREN) (and via AARNet to other NRENs linking the Centre to researchers and facilities world-wide) as well as high speed direct connections to CSIRO, The University of Western Australia, Curtin University, and the Murchison Radio-astronomy Observatory (MRO) some 800 km north of Perth. During the reporting period the Phase 2 expansion of the MWA was undertaken, with test data streamed to Curtin University at 100 gigabits per second. A further 100 gigabits per second link has been established between Curtin and Pawsey over existing optical fibre. This link will become operational in the coming year, supporting the operation of MWA Phase 2.

Pawsey uses a unique **groundwater cooling system** for the supercomputers, with no net consumption of water. Part of the CSIRO Geothermal Project, the groundwater cooling system pumps water from the aquifer below through an above-ground heat exchanger to cool the supercomputers, and then reinjects it back into the aquifer. Its environmental impact has been assessed as being close to zero, and there have been no adverse temperature or water level variations recorded in the monitoring bores. An electromagnetic device was installed in the system during the year, minimising water fouling issues in the heat exchanger and allowing the system to operate more efficiently.

During the reporting period, Pawsey operations have saved 4,808,582 litres of water through this system, compared to conventional cooling. More than 25 million litres have been saved since 1 January 2014.

The building's shaded facade contains a Photo Voltaic (PV) system and an extensive PV array occupies most of the roof. This installation generates 140 kW of electricity onsite, which acts to offset the electrical and CO_2 footprint of the Centre, and makes the groundwater cooling system zero emission.

In combination, Pawsey's building technologies place its 'green' credentials at position 154 on the **Green500 list,** a ranking of the top 500 supercomputers in the world by energy efficiency.

A World of Difference

Developing global partnerships to promote Western Australia as a leader of science and innovation

A Model for Collaboration and Partnership in Australia

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.

Pawsey is a clear example that the whole is greater than the sum of the parts. By bringing researchers, industry, supercomputers and technical support together, its community has access to infrastructure, expertise and opportunity at a scale unachievable for any individual institution or business.

National Research Infrastructure

The Pawsey Supercomputing Centre and its sister institution, the National Computational Infrastructure (NCI) are now recognised in **the National Research Infrastructure Roadmap** as being pivotal to the ongoing strength of our national research capability. While the Commonwealth Government has acted on this recognition to secure appropriate funding for high-performance computing (HPC) infrastructure upgrades moving forward, Pawsey is also working closely with the NCI to develop a coordinated approach to the provision, operation and access to HPC facilities nation-wide. Strengthening its partnership with NCI to work collaboratively in their operations both nationally and internationally moving forward has been a major focus during the reporting period.

Planet-Spanning Collaborative Projects

Pawsey is playing a key role in one of the largest and most ambitious scientific endeavours in history, **the Square Kilometre Array (SKA) project**. This international project will build a next-generation radio telescope to help scientists answer fundamental questions about our Universe.

The SKA is a joint effort between institutions from more than 20 countries, with the physical components of the SKA to be co-hosted in southern Africa and Australia. Pawsey underpins Australia's big data and supercomputing component of the SKA by supporting the computational and data storage requirements of Australia's SKA precursor telescopes, **the Murchison Widefield Array (MWA)** and **the Australian Square Kilometre Array Pathfinder (ASKAP)**. The Centre receives huge amounts of data from both of these telescopes, situated in the Murchison Radio-astronomy Observatory (MRO), to a dedicated supercomputer, Galaxy. Processing and data storage for these telescopes are significant activities for the Centre.

It is predicted that the SKA will eventually generate in excess of one exabyte of raw data (five times the global internet traffic) per day, and drives Pawsey staff and its other SKA collaborators towards continuously improving on state-of-the-art supercomputing and data analytics to eventually manage the requirements of true megadata.

Milestones supporting the SKA project in 2017–18 included completing service responsibility agreements to further improve collaboration, cooperation and communication between Australian participants, delivering additional storage space for the MWA, and supporting MWA's All Sky Virtual Observatory team to make MWA data more readily accessible to the wider scientific community.

Working with Government for the Public Good

Pawsey's research collaborations span countries, discipline boundaries, and government agencies, particularly in the area of public health. Whereas previously access to the Centre facilities was provided to State Government through the Director's Share Allocation, from 2017 state government departments could access services direct through Pawsey's Partner Allocation Scheme.

As an example, Artemis is a ground-breaking and lifesaving platform allowing monitoring of premature children in neonatal wards internationally. This is a collaborative project between the University of Ontario, Swinburne University, Western Australia Department of Health and Pawsey. The project enables the Artemis Platform cloud service to connect to bedside monitors, allowing a live stream and analysis of over 1,000 data points per second per patient, enabling real-time interpretation and response.

Other departments currently benefiting from Pawsey infrastructure include the Western Australian Department of Biodiversity, Conservation and Attractions, Department of Health, Department of Primary Industries, and Department of Water and Environmental Regulation.

Industry Collaboration

To accelerate innovation and development, Pawsey fosters relationships between researchers and industry, allowing industry collaborators to access supercomputing facilities, technical expertise, training and data storage. This year saw another industry collaboration take place through the Director Share allocation. **INTECSEA** WorleyParsons Group is an offshore engineering consultancy aiming to provide a tool to assess the integrity and safety of pipelines laid on the bottom of the oceans. They are now working with Pawsey to run computational fluid dynamic simulations of flows around subsea pipelines, to support subsea asset management.

In another example helping to develop real-world wave energy applications, Pawsey is working with both **Bombora Wave Power** and **Carnegie Clean Energy** to accelerate progress towards zero emission, high volume electricity. Several university-based researchers are also supporting industry-partnered projects through their allocations. The Centre intends to further develop its industry collaboration, preferably in association with independent research partners, in the coming years.

Building Engagement Through Training

Over 600 people attended training sessions delivered by Pawsey across Australia. In combination with the innovative activities undertaken this year to deliver training and information to a wider audience, Pawsey has significantly increased its reach and engagement.

The Pawsey Supercomputing Centre focuses on developing training programs in order 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

Through the reporting period, Pawsey has delivered a mix of on-site and online training, roadshows, user forums and hackathon events. Training modules have included basic computer science training for non-experienced users, introductory and intermediate courses on supercomputing and cloud systems, parallel programming courses, GPUs hackathons and customised training for specific research groups.

The success of these programs has been driven by an ever-increasing range of industries now using supercomputing and data facilities to support their research. Based on this increased demand, Pawsey now conducts training, in different Australian capital cities, nearly every month.

To further support the community of researchers, and increase accessibility of supercomputing both nationally and internationally, Pawsey has developed an **open repository of materials**. This repository covers different aspects of HPC systems usage, parallel programming techniques and cloud and data resources.

The graphs (Figures 1 and 2 next page) shows the significant growth in attendance at five core training modules of the National Training Program over the past three years. Pawsey believes this trend will continue with the further development of online training offerings.

CONNECTING ONLINE

Pawsey is continuing to explore ways to make training more accessible to researchers.

This year the Centre developed a new online course format designed especially for advanced computing and data infrastructure topics. This new format of training has proven to be very successful.

To increase the reach of Pawsey's training resources, this year has also seen the introduction of online training, via webinars and video tutorials. The Centre's first two webinars, organised by the Supercomputing Team, attracted an overwhelming number of registrations.

Previously, advanced training sessions were part of the face-to-face National Training Program with training modules covering MPI and Open MP with serial code optimisations offered during the last day of the training. The low number of registrations (Figure 2 next page), low turnout to these training sessions, and the challenge to address the range of knowledge requirements prompted the development of more specific and targeted online training.

Participation in the new Pawsey webinar sessions increased over the reporting period, as did the interaction and communication demonstrated during the sessions. Researchers were encouraged to use chat room messaging tools and online team collaboration tools. Both channels greatly enhanced engagement with the Centre and between users.

A suite of training and researcher highlight videos were developed and will be made available on **Pawsey's Youtube channel** during the next financial year.

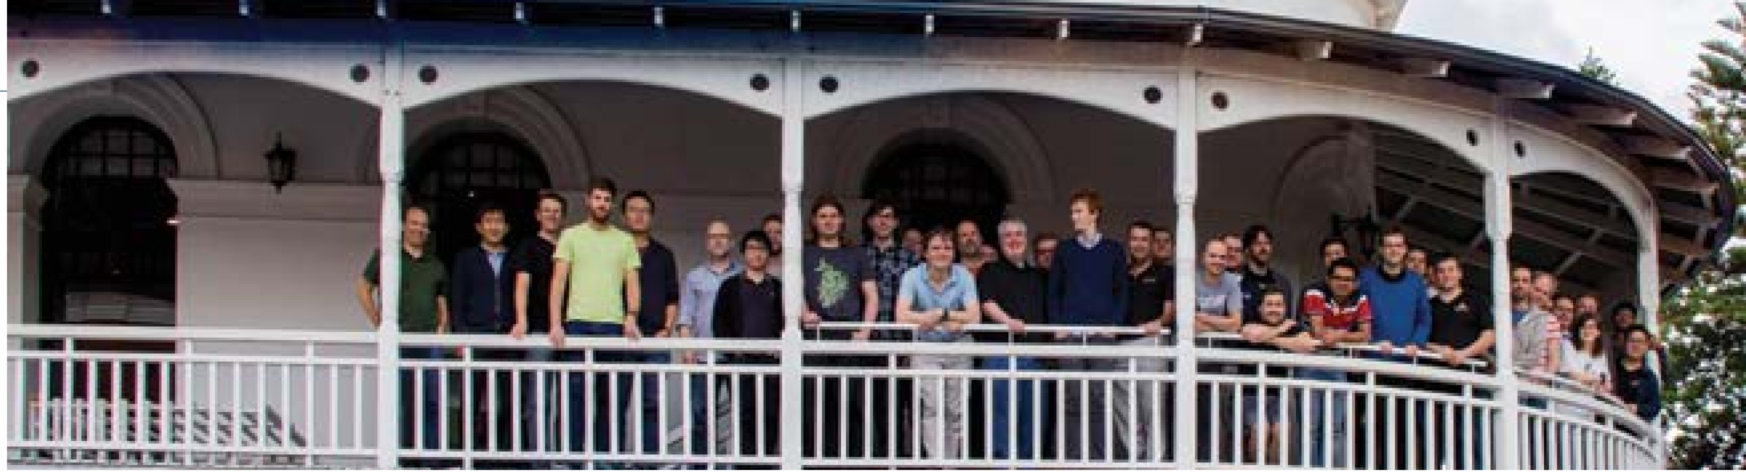

GPU Hackathon

TRAINING HIGHLIGHTS

GPU HACKATHON

Pawsey hosted the first **GPU Hackathon** in Australia. General-purpose GPUs are very appealing to researchers, offering exceptionally high memory bandwidth and performance for a wide range of applications, resulting in increased efficiencies and decreased data processing time.

In the past, Pawsey researchers have been more familiar with Central Processing Units (CPUs) than GPUs. The event was developed as a method to reduce this gap in knowledge. 'Hacking' GPUs and accepting the challenge of using and programming these accelerators means in the future researchers and programmers can optimise their code to solve bigger problems.

With a focus on collaboration and teamwork, 14 mentors worked with six teams of researchers and developers from around the world (Australia, USA, Asia and Europe) over a five-day period to solve difficult challenges. The skills and information gathered during this intensive program, as well as insight into possibilities for the future were taken back to their respective institutions.

SUMMER INTERNSHIP

The Summer Internship Program enables undergraduate students to develop skills in computational science through an immersive 10-week experience in advanced computing research projects. Interns are supervised by leading researchers and undertake challenging research projects utilising the Pawsey Supercomputing Centre.

With 14 places available this year, interest from both students and researchers has steadily grown with the 2017-18 Internships attracting submissions from 49 researchers and 70 students. This year was the first time applications were open for projects and students from interstate, with one project from University of Queensland successfully completing the program.

Year	2015	2016	2017
Number of projects submitted	28	33	49
Number of students who applied	51	46	70
Number of places allocated	12	16	14

Figure 1

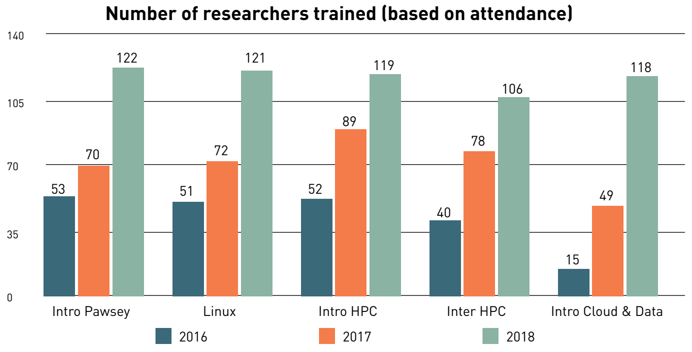

Figure 2

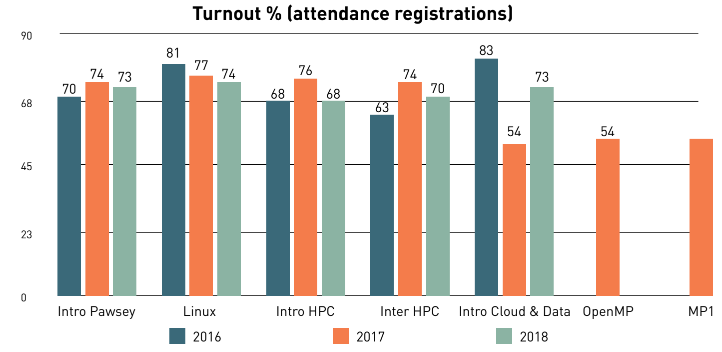

Our unique range of programs continues to grow the expertise of the current Pawsey community, as well as the next generation of supercomputing specialists, ensuring the creation of a skilled workforce across Australia.

NATIONAL TRAINING: CAPITALISING OPPORTUNITIES

Several opportunities were taken, as part of major events hosted for researchers in Australia, to showcase researchers' impact and also to upskill current and potential researchers to use Pawsey facilities.

ANITA WORKSHOP

The aim of the **Australian National Institute for Theoretical Astrophysics** (ANITA) is to ensure that Australian theoretical astrophysics is represented at a national and international level. Pawsey joined the event and showcased their involvement in astronomy, one of the Pawsey focus areas and a state science priority, with ANITA theory workshop and science school hosted in Perth in 2018.

C3DIS MELBOURNE

Pawsey exhibited at the first **Computational and Data Intensive Science** (C3DIS) Conference of 2017 in Melbourne. Pawsey staff hosted a workshop with users showcasing how its projects were taken to the next level with the help of supercomputers, and conducted an 'Introduction to Pawsey' session.

At the 2018 C3DIS conference, Pawsey hosted its first joint training session with NCI. Pawsey also conducted a workshop on building applications for the scientific community.

These conferences attract scientists, researchers, and computing, data and information management specialists from publicly-funded Australian research organisations. They provided important opportunities for Pawsey to showcase their expertise and capabilities, and reinforce the significance of Pawsey as a world-class supercomputing centre.

RESBAZ

The Research Bazaar (ResBAZ) is a worldwide festival promoting digital literacy emerging at the centre of research. In addition to playing an integral role in facilitating this project, this year staff from Pawsey provided a one-day training session around PySpark, Nimbus research cloud, serial optimisation and remote visualisation.

eRESEARCH AUSTRALASIA

eResearch Australasia is a conference providing an opportunity to engage, connect, and share ideas and exemplars concerning new information-centric research capabilities. Attendees learn how information and communication technologies help researchers to collaborate, collect, manage, share, process, analyse, store, find, understand and re-use information.

This year, Pawsey exhibited in partnership with NCI at the conference. Staff from Pawsey also provided an update on the latest technologies available for researchers across Australia, and presented Pawsey's role as part of the Astronomy Data and Computing Services (ADACS).

RNA-SEQUENCING WORKSHOP

Pawsey hosted a **RNA-Seq workshop** organised by the Australian Bioinformatics and Computational Biology Society and COMBINE. The event targeted the bioinformatics community with more than 25 attendees from across the country. Pawsey is deeply invested in supporting this relatively new area of research and is currently developing facilities to further its capacity to service the bioinformatics community.

OPENSTACK MEETUP GROUP

Over the reporting period, Pawsey staff identified and engaged with the OpenStack community in Perth. OpenStack is a cloud operating system controlling large pools of compute, storage, and networking resources throughout a datacentre. A dashboard gives administrators control, whilst also empowering their users to provision resources through a web interface.

A special session was hosted at the Centre, where a Pawsey representative explained how Pawsey is using OpenStack to integrate research and computing services whilst providing opportunities for researchers to investigate novel approaches to solving problems.

Engagement

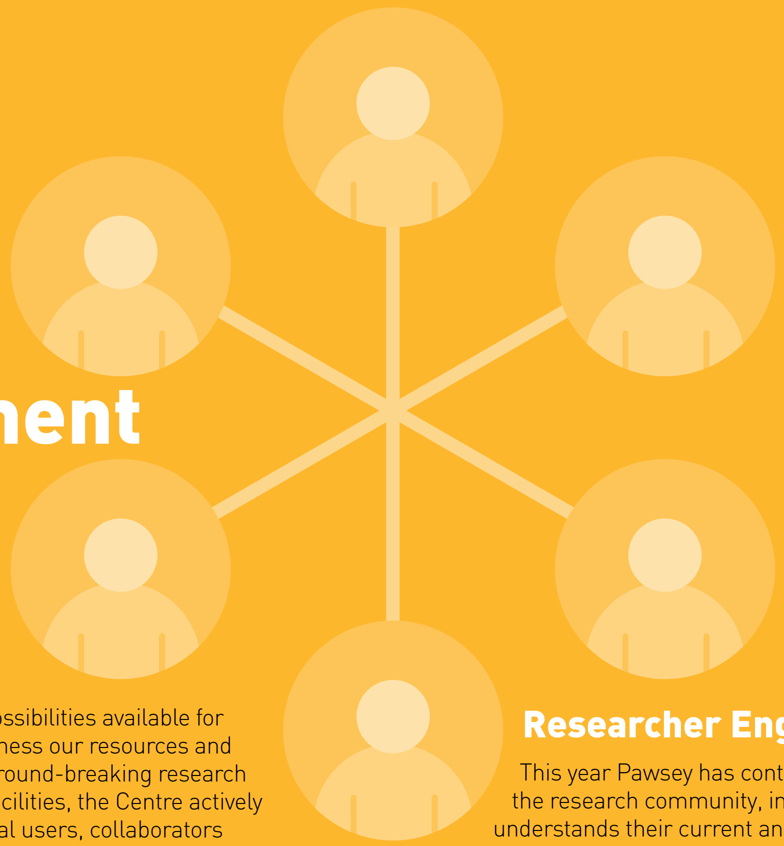

To raise awareness about the possibilities available for researchers and industry to harness our resources and expertise, and to highlight the ground-breaking research being produced using Pawsey facilities, the Centre actively engages with numerous potential users, collaborators and partners. Activities have targeted a range of different stakeholders, with a diverse focus spanning the research, education, industrial and HPC communities, both nationally and internationally.

Through tours, conferences, researchers' presentations, roadshows and user forums, Pawsey aims to connect more closely with researchers, the science community, government and industry, as well as developing its international relationships and standing in the international supercomputing and data sciences community.

Consolidating knowledge and best practices between the Pawsey and the NCI, the two facilities that embody Australia's Tier 1 HPC capability, has been a particular focus this year. Progress has been made towards formalising a collaborative relationship, and our unified presence was apparent from joint booths at eResearch Australasia 2017 and the Supercomputing Conference, SC-17, held in the USA. Pawsey and NCI staff have also embarked on a number of collaborative projects during the reporting period.

Moving into the new financial year with a new Executive Director and planning for an infrastructure refresh, Pawsey will be working to strengthen relationships with industry and drive collaborations that can translate into innovation and economic growth.

Researcher Engagement

This year Pawsey has continued to strengthen links with the research community, in particular to ensure that it understands their current and future HPC, data analytics and storage needs before starting the procurement process for the coming infrastructure upgrade. These linkages also serve to improve the Centre's current researcher experience, and create awareness of Pawsey services with potential new users, prior to the opening of requests for allocations at the end of the calendar year.

IMPROVING USERS EXPERIENCE: PAWSEY USER FORUM

Pawsey hosted the new Pawsey User Forums in four different states across Australia, engaging directly with over 30 researchers. The forums provided researchers with an opportunity to discuss their experiences and give invaluable feedback to the Centre. All feedback has been made publicly available through the Pawsey Portal and has also been circulated to the researchers using Pawsey services.

SHOWCASING IMPACT: PAWSEY FRIDAYS

Every two months, Pawsey hosts Pawsey Fridays at the Centre as an opportunity to demonstrate how industry partners, universities and government departments have taken advantage of the data, supercomputing and visualisation facilities and staff expertise available to enable significant scientific outcomes. Five events were held during the reporting period, with 140 people learning about researchers' impact while networking with Pawsey staff.

Credit: Sahlan Hayes, Official Photographer to Prime Minister Turnbull

LOOKING FOR THE NEXT 'BIG SCIENCE' PROJECT: PAWSEY ROADSHOWS

Pawsey staff has been visiting universities across Australia to raise awareness of the services available to Australian researchers, while showcasing current experiences from users and the impact Pawsey infrastructure has had in their projects. As part of this strategy, Pawsey was present during Research Week at each of the Western Australian universities, and also ran a session as part of C3DIS in Melbourne. During the reporting period more than 100 attendees joined Pawsey staff and researchers willing to find out more about how to access our services.

At the C3DIS conference in May 2018, Pawsey's presence extended well beyond the roadshow. It also hosted a conference exhibition booth and the team ran several pre-conference workshops. One of these was a joint session run with the NCI: an Introduction to High-Performance Computing. This workshop included an overview of Pawsey's use of machine learning and containerisation technology to make the use of HPC and parallel systems even more accessible for researchers across a range of technologies.

The Centre also hosted a booth as part of the previous C3DIS conference in July 2017, with an introduction to Pawsey services and joint panel discussion with NCI again part of the program.

A Roadshow was also hosted during eResearch Australasia (eRA). The event was opened to eRA attendees as well as the research community in general.

Government and Industry Engagement

Australia's Prime Minister Malcom Turnbull visited the Centre, joined by the Honourable Michaela Cash, Minister for Minister for Jobs and Innovation, and met with Pawsey staff for the announcement of the Commonwealth Government investment to secure the next generation of supercomputers.

Pawsey also hosted Senator the Honourable Arthur Sinodinos, in his role as Minister for Industry, Innovation and Science. Minister Sinodinos was welcomed by a group of Australian researchers, industry representatives and national supercomputing experts, who showcased the importance of the two HPC Tier 1 facilities to ensure Australia remains globally competitive.

The State government is also a strong supporter of the activities underway at the Centre. The Minister for Water, Fisheries, Forestry, Innovation and ICT, Science, the Honourable David Kelly, joined Pawsey for the launch of Data Science Week, Perth's festival of data. He also returned later in the year to learn how detailed 3D models of the Perth's groundwater aquifers created by Curtin University's researchers using Pawsey infrastructure has informed the Department of Water and Environmental Regulation's water management strategies.

Pawsey then took centre stage at the 2017 Department of Mines and Petroleum Open Day poster session. The event was full of opportunities for the resources sector and government to exchange updates on developments, gather stakeholder feedback and identify future trends. Opened by the Honourable Bill Johnston, Minister for Mines and Petroleum, the day attracted over one hundred key industry players and representatives who visited the Pawsey installation to explore how government and industry could benefit from the Centre's HPC infrastructure and support.

Aiming to further increase awareness of Pawsey services across the resources sector and industry, the Centre also hosted a booth as part of the Australasia Exploration Geoscience Conference held in Sydney. During the conference Pawsey gained valuable insight into the exploration industry and how we can enable more geoscientists to achieve more with their work.

Data Science Week is the biggest collaborative event Pawsey coordinates every year. This year 13 organisations ran 18 events across five days, celebrating data science and exploring its impact in Western Australia. The event was launched by Professor Peter Klinken, Chief Scientist of Western Australia and an avid Pawsey enthusiast.

Global Engagement

Pawsey's global engagement efforts has seen the development of formal associations with several international supercomputing and data centres. These relationships allow Pawsey to consolidate its international reputation and will provide the basis for exchanges on best practice, the development of joint training activities, and will allow staff to share information on operational procedures, governance and sustainable funding models.

COLLABORATIONS

Pawsey staff have been actively working to widen the national and international reach of the Centre. As a result, Pawsey staff have signed a Memorandum of Understanding with both the **Partnership for Advanced Computing in Europe** (PRACE) and **the National Supercomputing Centre in Singapore** (NSCC).

A letter of intent has been also signed with **National Energy Research Scientific Computing Centre** (NERSC) to facilitate and encourage the exchange of information relating to practices around the management of computational resources and technologies.

A series of seminars took place during the reporting period as a result of a collaboration with the **Edinburgh Parallel Computing Centre** (EPCC) UK. Presentations entitled 'Novel Approaches to HPC User Engagement' and 'ISO Certification' were delivered by EPCC staff online for Pawsey stakeholders and staff in Perth. In addition, a presentation entitled 'Using containers in HPC' was presented online by Pawsey staff for the EPCC community.

Based on these collaborations, Pawsey staff have since developed a biannual international best practices newsletter targeting supercomputing centres and researcher facilities around the globe. The newsletter is a collaboration between Pawsey, PRACE, NERSC and the Texas Advanced Computing Center (TACC). The newsletter content is topical and built through collaboration with other international HPC centres. **The first issue,** focused on training and outreach best practice, was published one week prior to the Supercomputing Conference, SC2017, and reached more than 400 interested parties. Articles were contributed by the founding collaboration partners.

The **second edition,** focused on allocation schemes, was issued before ISC18, an annual international conference for high-performance computing, networking and storage. Articles were contributed by NCI, NERSC and Pawsey. This project will continue into the next reporting year.

CONFERENCES

Building its international profile, the Centre partnered with the community-led **HPC Advisory Council (HPCAC)** to organise the first Australian HPCAC Conference, in Perth. The two-day event saw key HPC players from Australia and around the world present to more than 70 attendees drawn from the local, national and international HPC community.

To continue growing Australia's reputation as a leader in supercomputing and data sciences and to showcase the extraordinary outcomes researchers using Pawsey facilities are achieving, a joint booth was hosted by Pawsey and NCI at Denver's Supercomputing Conference, **SC2017**, the largest and more prestigious international conference for HPC. SC attracts scientists and engineers, software developers, policymakers, corporate managers, chief information officers and HPC administrators from universities, industry, and government agencies. SC has grown to become truly an international conference, and SC2017 featured more than 350 exhibitors from industry and research communities, attracted 12,000 delegates, and at least 117 countries were represented.

During the exhibition, Pawsey received the **HPCwire Reader's and Editor's Choice Award** in recognition of the Centre's outstanding performance. Pawsey took home the award for Best Use of High Performance Computing in Energy, for the Centre's collaboration on the Carnegie Wave Power project, alongside The University of Western Australia. Pawsey was also nominated in the field of life sciences, for the Centre's collaboration with Dr Laura Boykin and her Cassava Warriors team in their fight against insect and virus outbreaks on cassava crops.

Frankfurt then became another stage to showcase Australia's Big Science, this time at **ISC18**, an annual international conference for high-performance computing, networking and storage. Pawsey engaged with over 200 visitors on the exhibition floor of the conference, and also participated in the 'Birds of a Feather' discussion forum on HPC training.

Pawsey staff have presented at a number of international conferences both in Australia and overseas in the reporting period. They include:

- presentations at the HPCAC conference in Perth in August 2017.
- presentation at eResearch Australasia 2017 on the ADACS consortium.
- keynote presentation at the European HPC Summit Week/ PRACEdays18 conference at the end of May. This conference brought together experts from academia and industry to present their advancements in HPC.

Users Across the Globe

43	EUROPE
25	AMERICA
17	ASIA
5	NEW ZEALAND
1	AFRICA

Financials

THE PAWSEY SUPERCOMPUTING CENTRE RECEIVED

$70 MILLION IN FEDERAL FUNDING

FOR THE REPLACEMENT OF INFRASTRUCTURE.

"It is an exciting time to be at the Pawsey Supercomputing Centre"

– Ugo Varetto, Pawsey Acting Executive Director.

Funding to Accelerate Science and Innovation

The Pawsey Supercomputing Centre received $70 million in Commonwealth funding for the replacement of infrastructure, allowing the centre, already a world class facility, to further increase its capacity to service the national and international research community.

"It is a reflection of the government's understanding of the value that the Pawsey Supercomputing Centre delivers to the Australian scientific landscape.

[We are] accelerating innovation and increasing opportunities for engagement between Australian researchers and their peers internationally," said John Langoulant AO, Chair of the Pawsey Supercomputing Centre.

The funds are being allocated towards a capital refresh of Pawsey's infrastructure, in particular Magnus and Galaxy, as both systems are nearing the end of their operational lives.

The Centre's expert staff have been liaising with Australia's researchers to identify their needs and guide the configuration of the new systems.

Magnus, a Cray XC40 supercomputer, was the most powerful public research supercomputer in the Southern Hemisphere at the time of its installation. Pawsey aims to return to that level of international recognition of its facilities with the upgrade.

Galaxy is currently dedicated to the operational requirements of Australia's SKA precursor radio telescopes: the Australian Square Kilometre Array Pathfinder (ASKAP) and the Murchison Widefield Array (MWA). The increased capacity of the upgraded system will support Pawsey's ongoing and essential contribution to these projects.

This investment enables Pawsey to continue to drive innovation and accelerate discoveries in medical science, nano technology, engineering, geoscience, marine science, chemistry, food, agriculture and more.

Financial Report 2017–2018

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

The Commonwealth Government provided the Centre the initial capital investment for the construction of the Centre's building and high performance computing (HPC) and data infrastructure. The Commonwealth continues to support the ongoing operation of the Pawsey Centre through the National Collaborative Research Infrastructure Strategy (NCRIS) program.

The Western Australian Government, along with the unincorporated joint venture partners, provide the other portion of the operational funding for the running and maintenance of the Pawsey Supercomputing Centre. Pawsey, through CSIRO, employs professionals and experts on HPC to run the infrastructure and provide exceptional services supporting data storage, visualisation, user training and consulting.

Table 1: Financial Report for FY17-18

Pawsey Supercomputing Centre Statement of Income and Expenditure
01 July 2017 to 30 June 2018

	FY17-18 Actuals $'000
Revenue	
External Project Funding	10,908
Internal Joint Venture Partner Subscriptions and Recoverables	3,510
Total Revenue	**14,417**
Expenditure	
Asset Procurement (Capital Expenditure)	2,466
Labour Staffing	6,192
Machine Maintenance Contracts	2,642
Operating Costs	
Utility Charges	1,521
Other Operating Expenses	1,964
Total Expenditure	**14,785**
Surplus / (Deficit)	**(368)**

42%	NCRIS Funding (2017-19)
29%	WA Government
24%	Joint Venture Partner Subscriptions and Recoverables
3%	RDS, NCRIS, University of Queensland
2%	AAL Contract NCRIS

Revenue %

Figure 1: Sources of Revenue

The funding model for Pawsey aims to reflect the proportionate usage of machine on projects with approximately 46 per cent of operational funds provided by the joint venture partners and the Western Australia Government whilst 54 per cent of operational funding is provided by the Commonwealth through its NCRIS and NCRIS related program.

In FY17-18, capital expenditure on asset procurement represents only 17 per cent of its total expenditure. However, with the new grant of the $70 million capital refresh received in April 2018 from the Commonwealth, Pawsey started planning a significant capital investment for the next four years to replace its ageing infrastructure and to maintain its world-class status.

The major cost for the Centre is labour, representing the staffing costs for technical experts and allied support services, which is 42 per cent of its expenditure.

Machine maintenance comes third at 18 per cent and it is followed by other operating expenses at 13 per cent and utility costs at 10 per cent, this includes electricity and water usages. The Centre receives electricity credits from the solar panels installed at the building.

42%	Labour Staffing
18%	Machine Maintenance Contracts
17%	Asset Procurement (Capital Expenditure)
13%	Other Operating Expenses
10%	Utility Charges

Expenditure %

Figure 2: Expenditure Breakdown

Acknowledgements

The Pawsey Supercomputing Centre is supported by $90 million funding as part of the Australian Government's measures to support national research infrastructure under the National Collaborative Research Infrastructure Strategy 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.

GOVERNMENT OF WESTERN AUSTRALIA

NCRIS National Research Infrastructure for Australia An Australian Government Initiative

CSIRO Curtin University ECU EDITH COWAN UNIVERSITY AUSTRALIA Murdoch UNIVERSITY THE UNIVERSITY OF WESTERN AUSTRALIA

pawsey.org.au

CONTENTS