



A BRIEF HISTORY OF SUPERCOMPUTING



1 = 2000

iphone 6

CDC 6600

6000	MATHS speed (MFLOPS)	3
1400	CPU speed (MHZ)	10

What is a Supercomputer?

1



CDC 6600 Overview

3D rendering of an overview of CDC 6600 supercomputer, designed by Seymour Cray, with a figure as scale. Author FlyAkwa. Source commons.wikimedia.org

A supercomputer is basically an extremely powerful computer. It's a name given to the fastest, most powerful computers of the time. They're designed to work on problems too big and too complex for normal computers to deal with.

Having powerful processing capabilities is really what defines a supercomputer. ***The CDC 6600 was the world's first supercomputer; built by Seymour Cray, the "father of supercomputing" – today, it wouldn't compare to the processing power of our latest smart-phones.***

Supercomputers are akin to a cluster of normal computers working together. While a good PC might have one high-end processor, a supercomputer is likely to have over a thousand. These processors are designed to work in parallel to each other to solve extensive and complicated problems.

So what kind of work do supercomputers do? A common supercomputing task is to work with simulations. Working on models of the universe or different astral bodies, the interactions of atoms and molecules in physics, nanoparticle research for medicine, and fluid dynamics modelling for aquifers

The expense and rarity of supercomputers often means their time is in high demand. Not just anybody can walk into a supercomputing centre and use it. Official applications must be lodged for review.

For example, at Pawsey, projects should demonstrate:

- The research will make best use of Pawsey's supercomputing resources (a supercomputing justification).
- It will contribute to the State and National science and research priorities – such as radio astronomy, energy and resources.
- The research leader is employed at an Australian institution or an institution partnered with the Pawsey Supercomputing Centre

Supercomputing centres are often fully-booked and there is fierce competition between projects.

2 The Flop

The typical way to measure a supercomputer's power is by its FLOPS. 'FLOPS' stands for **F**loating **p**oint **O**perations **P**er **S**econd. You might not have heard of this term before because it specifically measures the power of computers.

On a PC you're probably used to measuring speed by clock speed (GHz). If you have a background in computer programming you might even be familiar with timing computers by instructions per second.

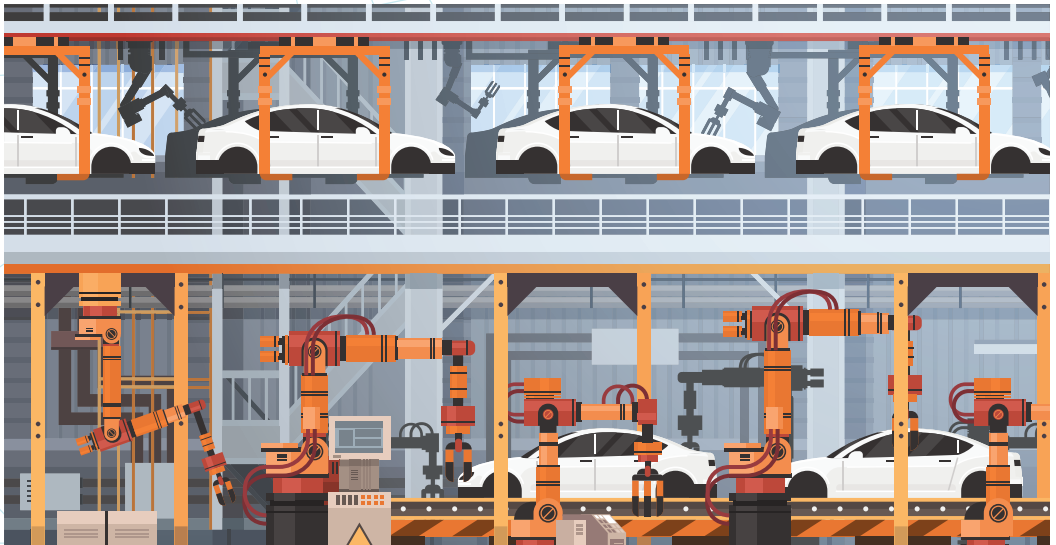
A PC is measured with these because of its ability to perform a lot of random tasks. It's designed to do things like run web browsers, watch videos, and play computer games. All of these things function using an input-output system. You might input a website address into your browser, which issues a command to the

computer to contact the website. Once the computer reaches the website, the site will input another command to the computer, giving you the output web page.

This works well for a typical computer, but supercomputers are built to perform complex mathematical calculations on large numbers as quickly as possible. In order to do this efficiently, processors work in pipelines and in parallel.

It's a bit like building cars in a factory. The most efficient way to build cars is to have separate teams, each working in parallel on different parts. One team will build the engine, while another will build the frame, etc.

In a typical computer, one team will do all steps until the car's completion. In a supercomputer,

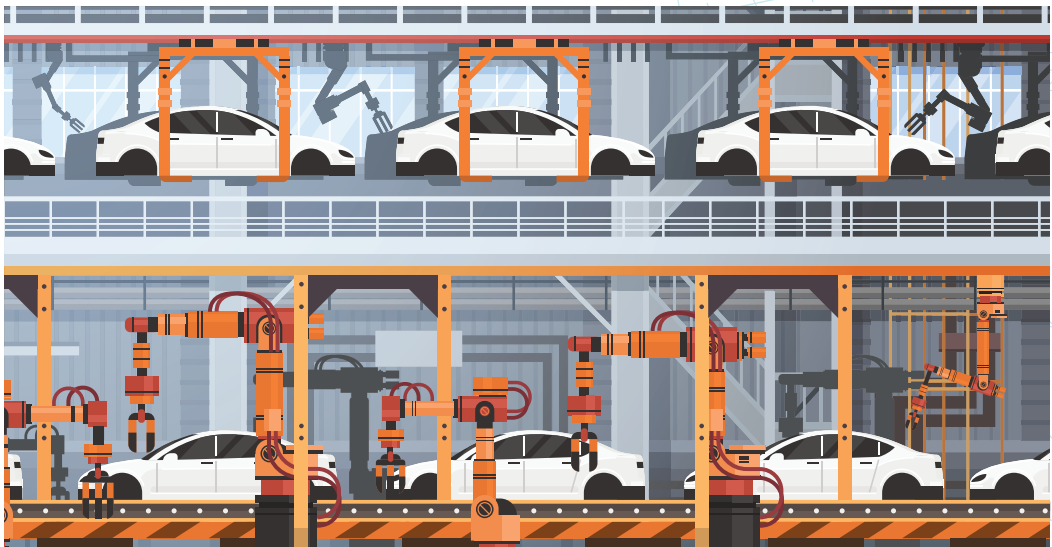


this works differently. As soon as the engine team has finished one engine, they begin another. While this is happening, another team might also be building engines, and another team is building frames. Each processor is a pipeline designed to work on a specific step. If it's designed well, this is much more efficient as every processor is always working at capacity.

So how many operations per second are we talking about here? The CDC 6600 was capable of three million FLOPS, using only a single processor.

In the modern era, the Magnus supercomputer is capable of 1 petaFLOPS (one quadrillion FLOPS, which is a one followed by 15 zeroes). The fastest supercomputer in the world, Summit in the USA, is capable of more than 200 petaFLOPS and has over 2 million processor cores!

The solutions to many of the biggest problems in science are limited by our ability to process data, so designing better supercomputers is a highly competitive and incredibly important field.



What Pawsey Does

Now that you know a little about supercomputers, here's how they are used at the Pawsey Supercomputing Centre

The Square Kilometre Array

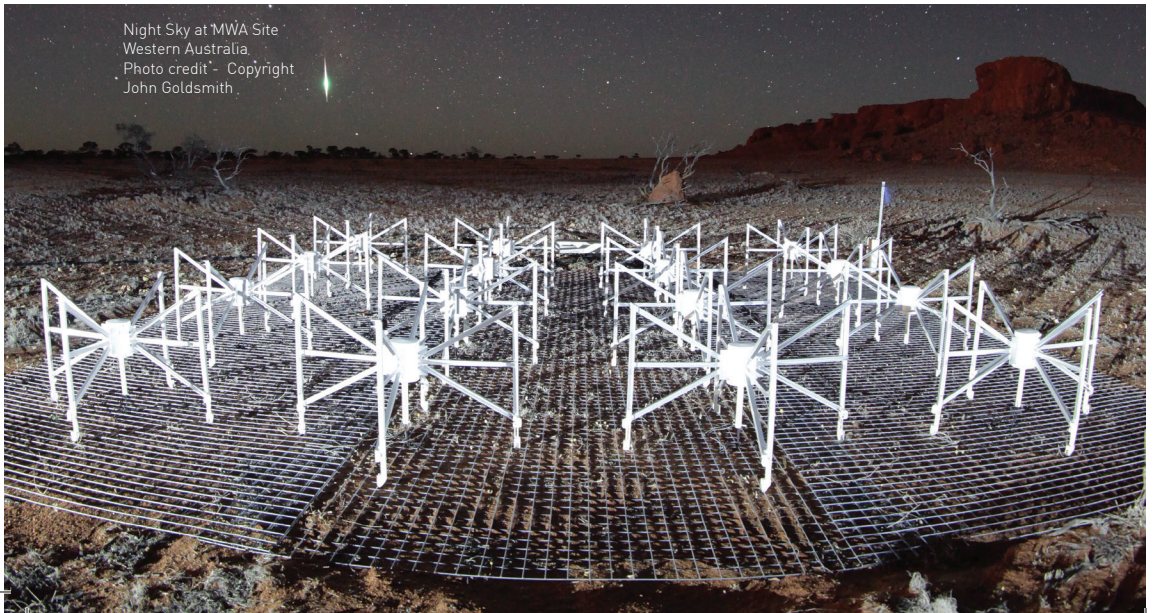
Once complete, the Square Kilometre Array (SKA) will be the largest radio telescope ever built. It's an international effort, built in two phases and features over 20 countries. Australia's role will be to house over 2,000 low-frequency antennas, designed to detect radio waves throughout the galaxy as they hit Earth. These radio waves often come from dying stars or powerful explosions in space.

Pawsey aims to process the scientific data from the SKA to produce images from the telescope. These images will be used by researchers to explore some of the most powerful and destructive events in our universe.

The data is so large it can only be stored and processed through the use of supercomputers. Pawsey will need to manage with terabytes worth of data each day, as the antennas and radio dishes observe space.

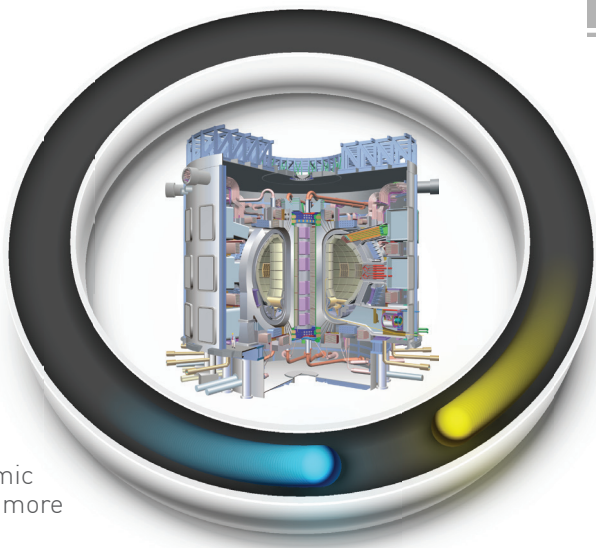
Even now, Pawsey is processing data from the Murchison Radio-astronomy Observatory (MRO), which houses precursors to the SKA. While it's a small-scale warm-up compared to the completed SKA, Pawsey and the MRO are producing science data for astronomy researchers.

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



Atomic Collisions

Let's go from the wide-open universe to the equally mysterious realm of atoms. Professor Igor Bray and his team at Curtin University are using Pawsey Supercomputers to map the collisions of atoms. Atoms and molecules colliding and interacting with each other makes up our very existence. As we map the atomic scale we are learning to build more precise technological tools.



One example is Professor Bray's research being used to help build the International Thermonuclear Experimental Reactor (ITER). The ITER will be the first fusion reactor in the world to produce a surplus of energy. Fusion energy is created by fusing together hydrogen atoms, instead of using uranium and other heavy elements. ITER will get its hydrogen from ocean water, producing vast amounts of energy.

International Thermonuclear Experimental Reactor, the world's first fusion reactor to produce a surplus of power and relies on the interactions of hydrogen atoms

The research into atomic collisions has also been used to help create proton therapy, which aims to target and kill cancer cells by firing protons. This would mean patients could have their cancer treated without the side effects that come from chemotherapy.

Plotting how atoms act when they collide isn't simple work, however. Professor Bray's team have been working on the problem for more than 20 years. Much of that time involved the team using supercomputers at their capacity to simulate the atoms. As supercomputers became more advanced, the team were able to simulate more atoms and have even begun research into simple molecules. Professor Bray uses the Magnus supercomputer for the majority of his team's work.



Medical Technology

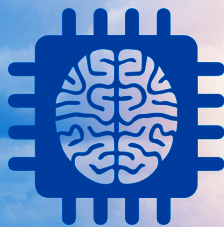
Another big project using Pawsey is the Artemis platform. Artemis is a medical program used in hospitals that stores a live stream of patients' health data. When monitoring a sick person in hospital, the machinery feeds doctors data as it happens. Blood pressure, pulse, and other vital signs are displayed on a screen moment-to-moment. This data is then deleted as the machine updates its readings.

The problem is, subtle patterns in the medical data that could trip alarm bells is lost because the data isn't stored or analysed. Late Onset Neonatal Sepsis is a type of infection newborn babies are vulnerable too. Often it's caught too late and the results can be deadly.

With the help of Artemis's data storage, doctors can analyse medical information over time and look for patterns in a newborn's health that can act as warning signs for infection.

The data storage and processing for this kind of work would be beyond many small and regional hospitals. That's where Pawsey comes in. The Nimbus cloud service is different to other supercomputers at Pawsey.

It's a digital cloud that can constantly stream and process data and is managed for very high uptime. This makes it the perfect hub for processing data for the Artemis. Even remote hospitals will need only monitoring technology, minimal onsite computing and a stable internet connection.



Artificial Intelligence

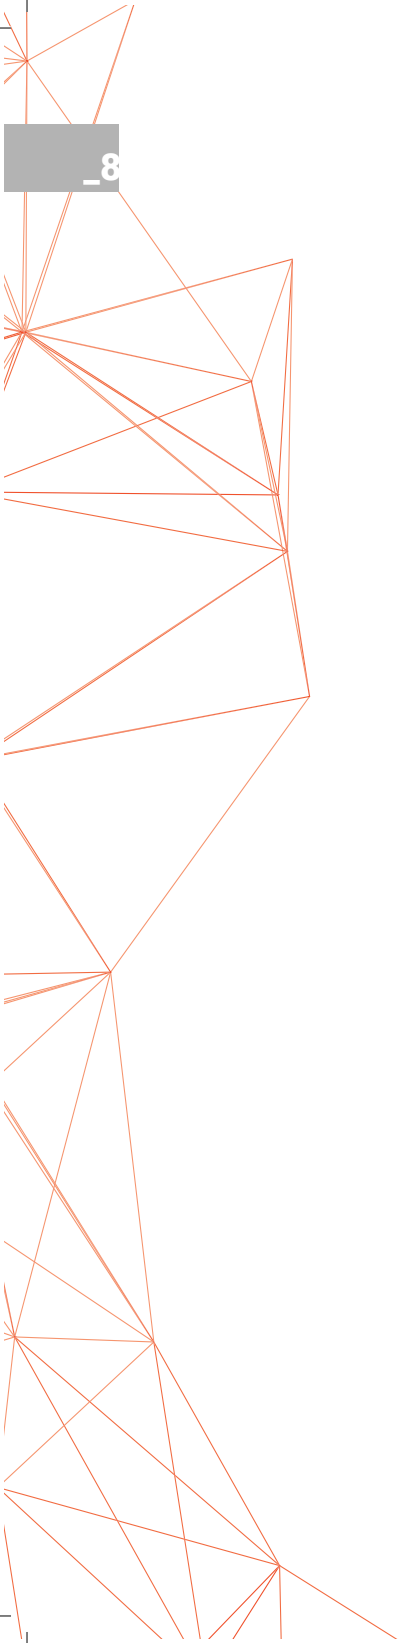
While getting rid of weeds on farms may not sound like Earth-shattering science, teaching a computer how to recognise a weed goes straight to some of the biggest challenges in machine learning in our age.

While a human might easily pick out a weed in a row of crops, it's incredibly difficult to teach a robot to do the same. The robot needs to recognise what a weed looks like. However, even plants of the same species can vary in appearance. Compound that with lighting problems, camera angles, and the many other problems that come from using imaging in an open setting, and it becomes almost impossible to get a machine to reliably pick out weeds.

While machines are faster than humans at processing simple, repetitive tasks, they lack even the most basic ability to think or adapt. Dr Selam Adherom at Edith Cowan University is using machine learning to create a robot capable of recognising and killing weeds.

To do this Dr Adherom's team trains neural networks. Neural networks are computer programs that learn to identify patterns similar to a human brain. These networks need to be trained on large data sets. That's where Pawsey's supercomputers come in. Pawsey's Zeus cluster allows Dr Adherom's team to train networks using different data sets simultaneously. Soon, the team are hoping to reveal a machine capable of targeting weeds and spraying them with herbicide. This would allow farmers to use about 90% less herbicide.





_8

1965 The Birth of Supercomputing

9

The CDC 6600, the world's first supercomputer

The CDC 6600 is considered the first supercomputer. It was designed by Seymour Cray – often considered the father of supercomputing – while he was working at the Control Data Corporation (CDC). It's considered the first supercomputer because, at the time of its release, it was up to ten times faster than the world's fastest supercomputer, the IBM 7030 stretch.

The CDC 6600 was roughly the size of four filing cabinets, which makes its performance even more incredible – considering the IBM 7030 was around 600 square metres, or roughly the size of an average house.

The CDC was a relatively small company and the release of its supercomputer completely disrupted the computing industry. At the time, IBM CEO Thomas Watson Jr. released an internal memo saying:

"Last week, Control Data... announced the 6600 system. I understand that in the laboratory developing the system there are only 34 people including the janitor. Of these, 14 are engineers and 4 are programmers. Contrasting this modest effort with our vast development activities, I fail to understand why we have lost our industry leadership position by letting someone else offer the world's most powerful computer."

CDC 6600 Overview
3D rendering of an overview of CDC 6600 supercomputer, designed by Seymour Cray, with a figure as scale. Author FlyAkwa. Source commons.wikimedia.org



_10

1972

The ILLIAC IV, the most infamous computer

ILLIAC IV computer. In this SIMD parallel computing machine, each board has a fixed program that it would farm out to an array of Burroughs machines. It was the "cutting edge" in 1966. Credits: Steve Jurvetson from Menlo Park, USA Flickr

The ILLIAC IV was an important disaster. Only one model was ever built. It cost four times as much as initial estimates, and was finished years late. That's the disaster part, but the important? Well, the ILLIAC IV was the first computer to be built with parallel architecture. This means it had multiple processors working together, it was seen as a way to get around the technological limitations of processing power at the time.

It was suppose to have up to 256 processors. However, it was only partially completed, leaving it with 64 processors. By the time it reached this

partial completion other supercomputers – like the CRAY-1 – were already outstripping its performance.

Despite its tremendous failure as a supercomputer in its own right, the ILLIAC IV's clever use of multiple processors inspired many later models of supercomputers. Eventually, as processors became cheaper to manufacture, this became the dominant form of supercomputer construction.

Nowadays all supercomputers are built with multiple processors – often thousands of them.



1976₁₁

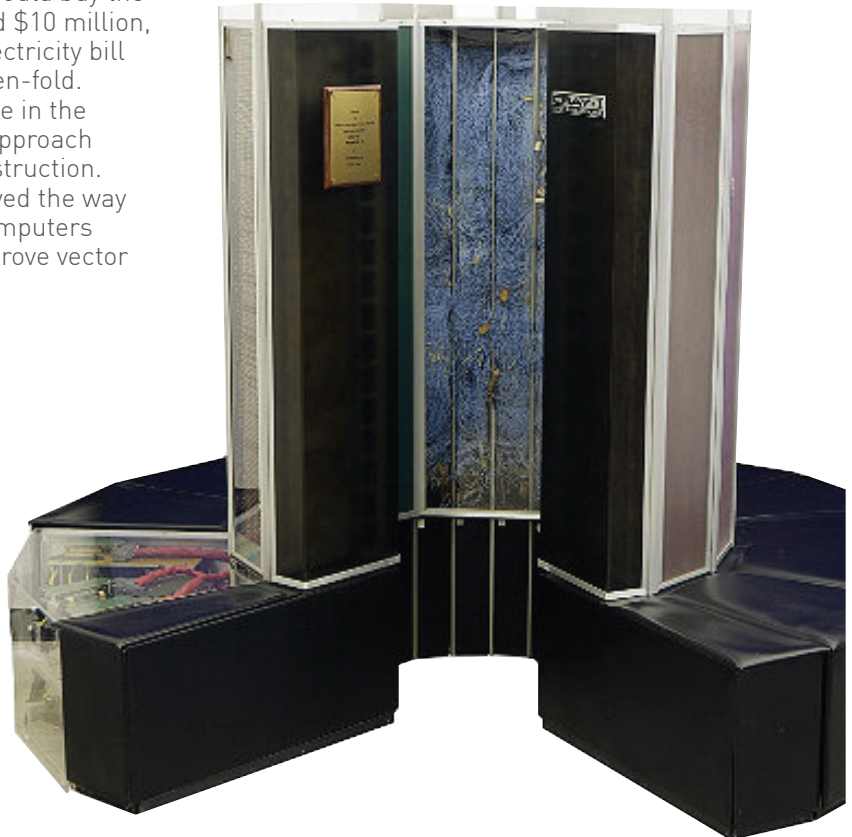
The Cray-1 and Vector Programming

By the time the ILLIAC IV was commercially available it had been beaten by the Cray-1. After developing the CDC 6600, Seymour Cray left the company and founded his own. Cray had quit after becoming tired of management looking for cost-cutting measures that would lead to an inferior product.

At the time, you could buy the Cray-1 for around \$10 million, although your electricity bill would increase ten-fold. Cray didn't believe in the multiprocessor approach to computer construction. Instead, he believed the way to make supercomputers faster was to improve vector programming.

The idea behind vector programming was to slim down the input/output wait times that occur in calculations by trying to rid itself of loops. Rather than loop an operation, running different variables each time, the supercomputer searched and found all variables, while simultaneously preparing the calculation.

Cray-1A; Cray Research, Inc. (CRI)
Credit:
Gift of Lawrence
Livermore
National
Laboratory
Photographer
© Mark Richards



_12

1987 The Birth of Australia's Supercomputers

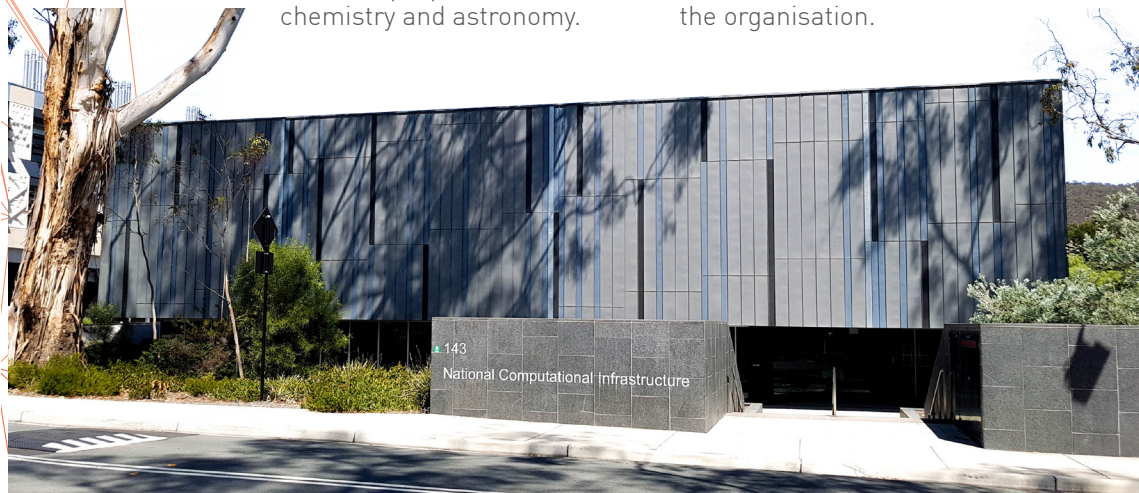


The Australian National University (ANU) broke ground in 1987 by buying a Fujitsu FACOM VP 100. At the time, the CSIRO and Bureau of Meteorology already had supercomputers in Australia, but the ANU's supercomputing facility became the foundation for Australia's supercomputers. Other universities began to build their own supercomputers for research purposes – often for chemistry and astronomy.

Australia National Computation Infrastructure (NCI), Canberra
Currently one of the Australia's peak research supercomputing facilities

It wasn't until the late-90s that the Australian government intervened and organised a national supercomputing framework for Australian universities, tying all the supercomputing centres together in partnership with the Australian Partnership of Advanced Computing (APAC).

After 2007, APAC became the National Computational Infrastructure (NCI), which still exists today. This involved several million dollars of funding to make Australia's supercomputers internationally competitive. The NCI also wanted to devote the majority of its time to a merit program, to ensure important public research was the main focus of the organisation.



1993-2001

13

Birth of the and Australia's First Ranking



If you want to find some of the fastest supercomputers in the world, the Top500 is a good place to look. It ranks supercomputers by applying the LINPACK Benchmarks to them – a set of linear equations designed to test a supercomputer's FLOPS.

It wasn't until 2001 that Australia managed to rank in the Top500 with its Compaq/HP Alphaserwer SC system, reaching number 31. The new supercomputer was built for the APAC's national facility and managed to break into teraflop territory. Not a bad start for our country, but Japan was about to revolutionise processing power just a year later.

2002

The Earth Simulator

In 2002, Japan revealed the Earth Simulator. Designed to predict movements of the Earth's crust and solve environmental problems, it was the most powerful supercomputer in the world. Its designers had wanted to create a virtual Earth to test environmental simulations on.

It was capable of a monstrous 32 teraFLOPS of performance power. Compare that to its closest competitor, the US ASCI White system, which was running just over seven teraFLOPS. The Earth System had 640 shared memory nodes and eight vector processors. It remains the last world-class



supercomputer to use the classical vector processing design approach. It stayed on the Top500 for two years until the IBM Blue Gene appeared.

Earth Simulator in
Japan JAMSTEC
2007.
Credits:
Manatee_tw

2004

IBM Blue Gene

The IBM Blue Gene was originally designed to help biologists simulate protein folding and gene development, but it soon became such a popular supercomputer its uses expanded well beyond this.

Supercomputers were getting larger and larger, and the power required to keep them running was skyrocketing with each new model. The developers of Blue Gene realised that if the world continued to build supercomputers this way, it

would soon require a football field just to house them, and take the power cost of a mid-sized town.

Blue Gene was designed to use 131,000 processors underclocked to reduce power and heat generation. Each dual processor was designed to be small enough to fit 32 of them on a single microchip. This massively reduced the size of the Blue Gene compared to its competitors. It had the processing power of around half a petaflop.

2011

The K Computer

A man look at the K computer
Photo: Nicolas Datiche/Hans Lucas

The K Computer was a design named after the Japanese word 'kei', meaning ten quadrillion. It was the first supercomputer to reach over 10 petaFLOPS – 10 petaFLOPS is roughly 10 quadrillion FLOPS, hence its name – though its average closer to nine.



When it hit number one on the Top500 list it was more powerful than the following five supercomputers combined. While energy-efficient for its size, the K Computer was a monster. It had 864 cabinets and needed a room roughly half the size of a football oval, using over 1,000km of cable.

2012

Raijin Australia's Highest-Ranking Supercomputer

Named after the Japanese thunder god, when Raijin debuted in 2012 it ranked 24 in the Top500. Given its name, it's probably not surprising that Raijin was designed to work in the field of climate science. It had a peak performance of 1.1 petaFLOPS, which was needed to deal with the demanding atmosphere and ocean simulations Raijin was expected to perform.

It remains the highest-ever ranked Australian supercomputer, with over 57,000 cores. With a peak performance of around 1.1 petaFLOPS. Thanks to an agility upgrade in 2015, which jumped up its processing power, Raijin is still in use at the NCI. True to its name you can recognise it by the steam cloud that rises from the NCI building, a side effect of the constant cooling needed to keep the machine running.

Pawsey's Magnus Supercomputer Makes Records

2014

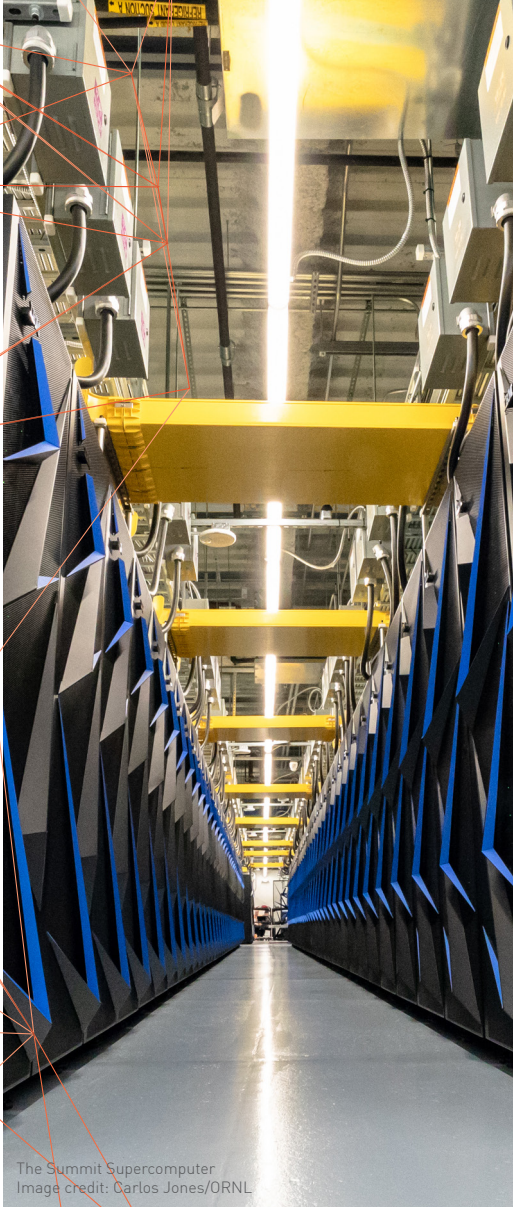
Magnus, Pawsey
Supercomputing
Centre

In 2014, Magnus, the supercomputer was built. Magnus, which is Latin for 'great', ranked at #41 on the Top500 and was the most powerful publicly used research computer in the Southern Hemisphere at the time.

With over 35,000 cores, Magnus is the culmination of Pawsey's petascale research goals. It is fast and in high demand across a number of fields, including geoscience, astrophysics, chemistry, and biology.



The China/US Supercomputer Race



The Summit Supercomputer
Image credit: Carlos Jones/ORNL

At the start of 2013, The Tiahne-2 was unveiled two years ahead of schedule. Tiahne means Milky Way, and the supercomputer has its own minor galaxy of processor cores, standing at over 3 million cores in 16,000 nodes. It beat the previous chart-topper, the US supercomputer Titan, by almost doubling its processing power. Tiahne-2 is capable of over 30 petaFLOPS. The Tiahne-2 maintained its number one spot for three years.

The US and China were in fierce competition to pass the 100 petaflop mark for supercomputers. The Tiahne-2 was due to be upgraded, potentially breaking this barrier.

China then built a new supercomputer, the Sunway TaihuLight. The TaihuLight is capable of a massive 93 petaflops of processing power and over 10 million cores. The TaihuLight took over from the Tiahne-2, maintaining the number one spot in the Top500 until June 2018. The US finally took the spot back with the \$200 million Summit computer, a supercomputer capable of over 122 petaflops, powered by over 2 million cores.

Unsurprisingly, the race between China and the US to build the fastest supercomputer is still very much underway.

Resources

Want to leave your mark on the world, but don't know where to start? Here are some resources to help guide you into the supercomputing world:

CoderDojo WA

<https://coderdojowa.org.au/>

The Hour of Code

<https://hourofcode.com/au>

Code Monster

<http://www.crunchzilla.com/code-monster>

Code Camps in WA

<https://codecamp.com.au/coming-soon/>

Code Clubs

<https://codeclubau.org/find-a-club>

Scratch

<https://scratch.mit.edu/>



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.



Curtin University



THE UNIVERSITY OF
WESTERN
AUSTRALIA

www.pawsey.org.au