



NCI AUSTRALIA

PORTFOLIO OF SERVICES



NCI Australia
The Australian National University
143 Ward Road
Acton ACT 2601

T +61 2 6125 9800
E enquiries@nci.org.au
W nci.org.au

 @NCINews
 National Computational Infrastructure
 NCI Australia

© National Computational Infrastructure 2022

This work is copyright.
Apart from any use permitted under the *Copyright Act 1968*, all rights are reserved.

More information can be found on NCI's website at nci.org.au
To read this Portfolio of Services online, scan this code:

Produced by
NCI Australia
The Australian National University
143 Ward Road
Acton ACT 2601



Cover image: Turbulent particle motions from high-resolution simulations of star formation.
Image by James Beattie, The Australian National University

We acknowledge the Traditional Custodians of the ACT, the Ngunnawal people.
We acknowledge and respect their continuing culture and the contribution they make to the life of this city and this region.

CONTENTS

ABOUT NCI	// 1
WHAT WE DO	// 2
HIGH-PERFORMANCE COMPUTING – GADI	// 3
CLOUD COMPUTING – NIRIN	// 4
DATA STORAGE	// 5
DATA SERVICES & COLLECTION MANAGEMENT	// 6
CASE STUDIES – COVID-19 & BATTERY TECH	// 7-8
VIRTUAL RESEARCH ENVIRONMENTS	// 9
SCIENTIFIC VISUALISATIONS	// 10
USER SUPPORT	// 11
USER TRAINING	// 12
COMPUTATIONAL SCIENCE ENHANCEMENTS	// 13-14
CASE STUDIES – BUSHFIRES & GENOMICS	// 15-16
COLLABORATORS	// 17-18

CONTACT US

NCI Australia provides access to computing, data and virtual environments through schemes allowing access to researchers from many different institutional and discipline backgrounds.

Access to NCI facilities and services by commerce and industry is welcomed. Access is available on a Fee-for-Service basis. If your organisation is interested in getting access to NCI Australia capabilities, please submit a business enquiry at: enquiries@nci.org.au

Documentation about using and accessing supercomputer, cloud and data resources can be found in our online User Guides, visit nci.org.au.

You can find all of NCI's contact information by scanning this code:



ABOUT NCI

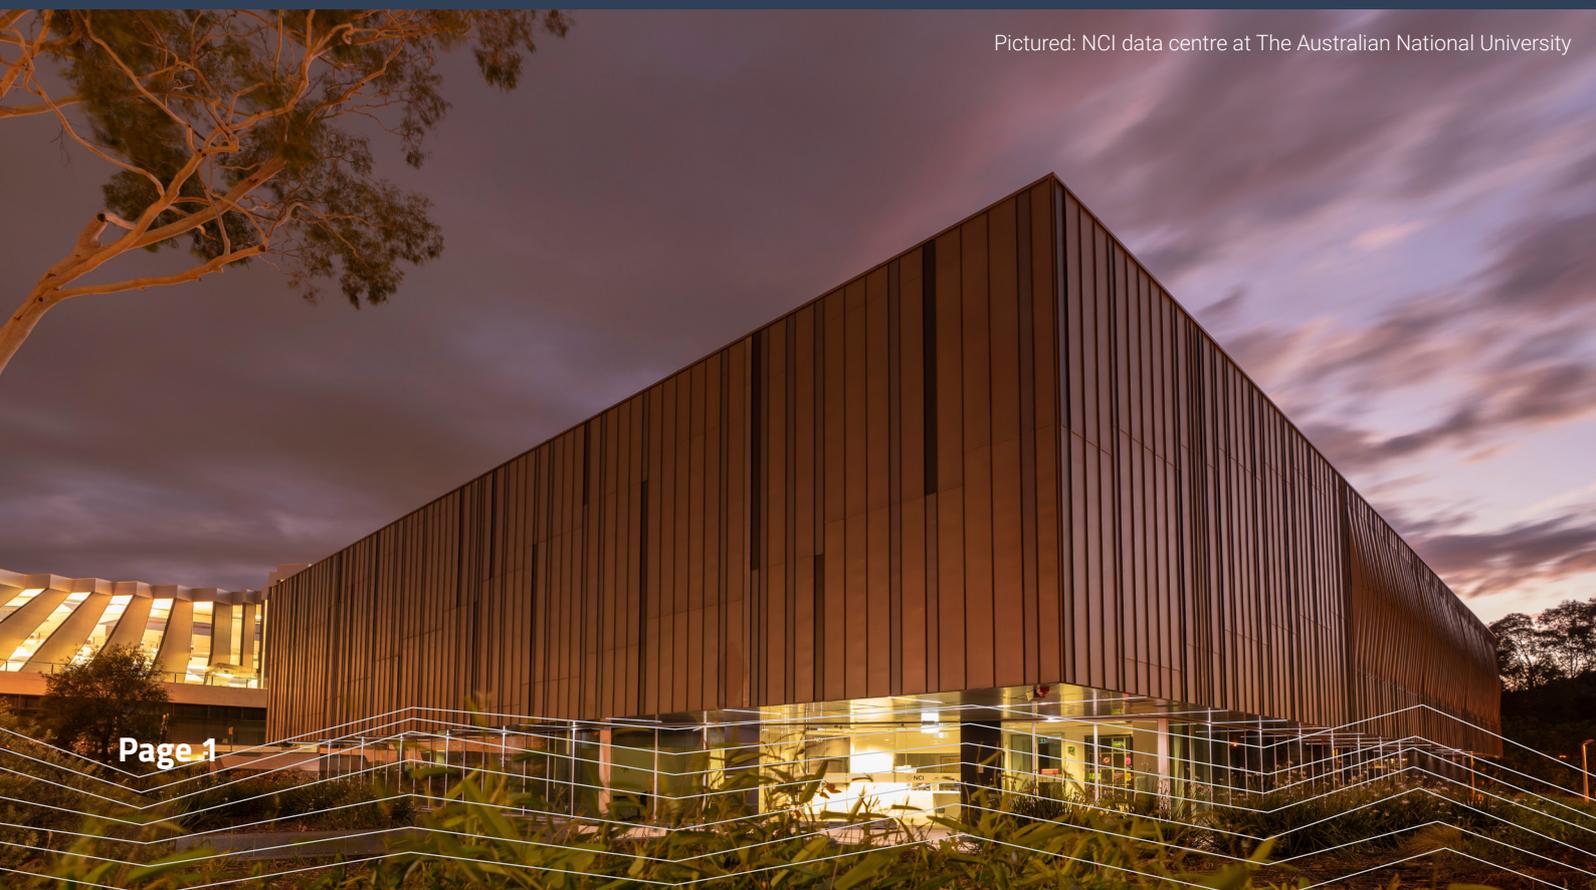
NCI is a collaborative hub at the centre of the Australian research ecosystem. We enable transformative science through big data and computing technologies, platforms and expertise. Through high-performance computing (HPC) and data (HPD) infrastructure, tightly coupled with our internationally renowned expertise, NCI provides essential services that support the needs of research and industry today and into the future.

NCI supports the work of over 8,000 researchers from 35 universities, 5 national science agencies, a dozen NCRIS science infrastructure providers, 3 medical research institutes, and industry. Scientific research is highly dependent on the fusion of “big compute” and “big data” that NCI provides in areas such as weather and climate science, the earth sciences, earth observation, bioinformatics, and astronomy.



NCI is a collaborative hub at the centre of the Australian research ecosystem. We enable transformative science through big data and computing technologies, platforms and expertise.

Pictured: NCI data centre at The Australian National University



WHAT WE DO



NCI Australia is the country's leading high-performance data, storage and computing organisation, providing expert services to benefit all domains of science, government and industry.



The Gadi supercomputer, operated by NCI, can perform over nine quadrillion calculations per second.



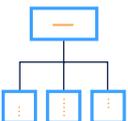
NCI operates the fastest filesystems in the Southern Hemisphere, and is home to more than 70 petabytes of critical research data.



The Nirin cloud supports interactive data analysis workflows with dedicated environments optimised for specific research communities.



NCI's data services allow users, data portals and external science cloud environments to access, interact with and extract value from our climate, weather, astronomy and genomics data collections.



NCI curates and optimises nationally and internationally significant datasets such as CMIP6 and ERA5, making them suitable for data-intensive science as well as broader access.



NCI's VizLab team generate images, videos, and VR experiences that extend the discovery process for scientists, and effectively communicate high-impact research to scientists and the general public.



NCI produces training resources and in-person training courses throughout the year to help develop the skills of the HPC and HPD user communities.



NCI provides User Support to assist users with administrative and technical issues around accessing and using systems and services at NCI.

HIGH-PERFORMANCE COMPUTING

GADI SUPERCOMPUTER

Australia's Tier-1 supercomputer, Gadi, is capable of over one billion hours of computing per year, across more than 100 supported software packages.

The name 'Gadi' means 'to search for' in the Ngunnawal language. It perfectly encompasses NCI's mission of scientific research and high-performance computing: to search for knowledge that can make the world a better place through enabling innovative world-class research.

- 9.26 PetaFLOPS (over nine quadrillion calculations per second)
- 185,256 CPU cores (24-core Intel Xeon 'Cascade Lake')
- 3,234 nodes total, including 3,074 CPU nodes and 160 GPU nodes
- 640 Nvidia V100 GPUs
- Linked with storage systems via 200Gb/s HDR InfiniBand by NVIDIA
- More than 100 scientific software packages.



Pictured: Gadi supercomputer, NCI

CLOUD COMPUTING

NIRIN CLOUD

NCI's newest cloud computing platform is called Nirin, meaning 'edge' in the Australian First Nations Wiradjuri language.

Nirin provides an Infrastructure as a Service platform which supports NCI and our stakeholders with a dynamic compute and storage environment which can scale up and down as needed. The Nirin Cloud provides the foundation for many of NCI's interactive data analysis, Virtual Laboratories and data publishing environments and is used extensively for processing and preparation of data used on the Gadi supercomputer.



DATA STORAGE DISK, TAPE & SSD

NCI operates multiple high-speed, low-latency global parallel filesystems to support the most data-intensive research workloads.

Image: NCI filesystems

We run the fastest filesystems in the Southern Hemisphere, linking high-performance computing (HPC) with high-performance data (HPD) via 100-gigabit network links. NCI's filesystems, catering to the needs of our research community, enable the next generation of computational tasks, including high-throughput computing (HTC).

NCI's filesystems contain around 70 Petabytes of data storage capacity, providing space for research data to be stored in five separate global Lustre filesystems, reaching a total aggregate IO performance of 450 GB/second.

NCI also runs a special IO-Intensive Platform, a dedicated filesystem using 576 2-Terabyte NVMe drives for cumulative performance around 960 Gigabytes per second.

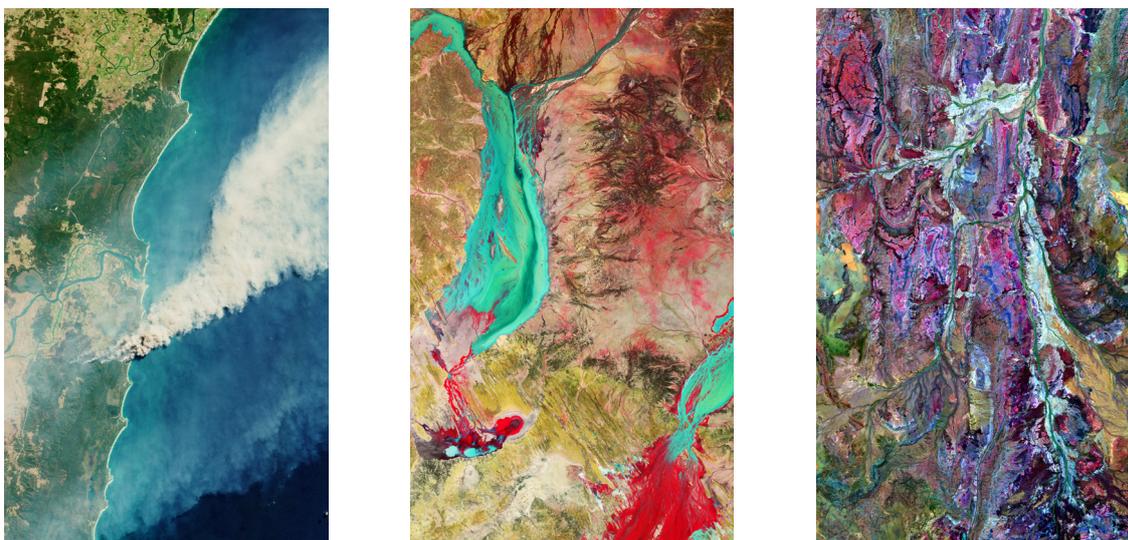
NCI also stores 50 Petabytes of archival project data in state of the art magnetic tape libraries. In total, NCI runs more than 15,000 hard drives from vendors including NetApp, DDN and HPE.



Image: NCI tape library

DATA SERVICES & COLLECTION MANAGEMENT

NCI increases the value and reach of research-ready datasets by providing high quality and high performance data-intensive services. As a trusted national data repository, NCI optimises some of the largest and most significant datasets requiring our high-performance capabilities. These cover climate and weather collections through CMIP6 and ERA 5, earth-observation through Landsat and Copernicus, and many more.



Images: Australian satellite imagery stored at NCI

We are constantly adding to our data collections' functionality to help users better access, analyse and share the data they need every day. A core aim is making sure that the data stored at NCI is safe, secure, accessible and readable. As a trusted repository for some of Australia's biggest datasets, we play a key role in the national data science community working on projects of environmental management, coastal erosion, bushfires and extreme weather.

All the significant national data collection at NCI, in total making up more than 10 petabytes of data, meet the international FAIR data standard: data should be Findable, Accessible, Interoperable and Reusable for the research community. Maintaining the datasets entrusted to us by the research community and the national science agencies for broad scientific use is a central facet of our data collections work. We are constantly working with our data providers to make sure that their data collections are up-to-date and accurate.

SUPPORTING COVID-19 RESEARCH

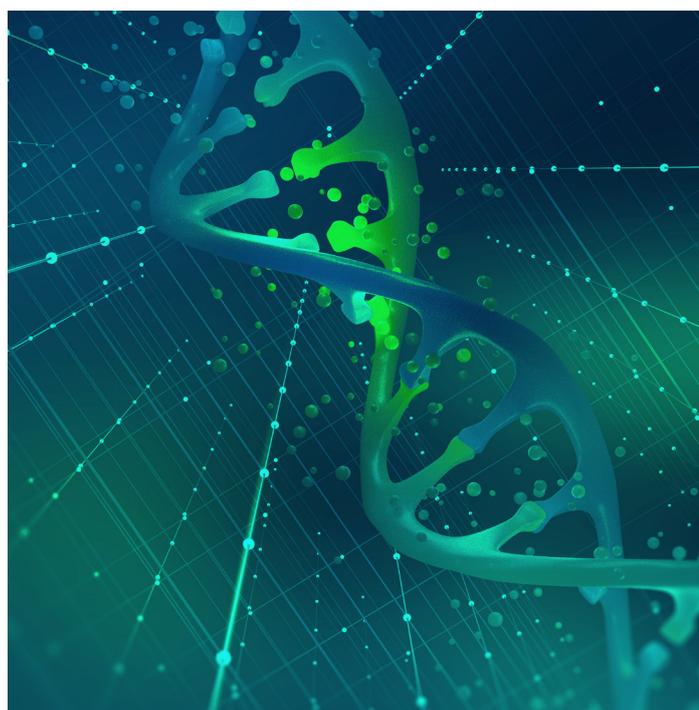
CASE STUDY



The work of computational researchers can sometimes take place years or decades before the products or technologies they help develop make their way into people's lives.

NCI provides the infrastructure that researchers need to respond to urgent national and global priorities as they arise. Streamlined access and the increased computational capacity of Gadi allow scientists to respond to issues impacting all Australians. Never has this been more apparent than with the plethora of COVID-19 research projects that NCI has been able to support. For NCI's users, high-performance computing and easily accessible resources made it possible for them to quickly focus their research efforts on this new direction.

High-resolution simulations of the molecular interactions on the surface of human cells require thousands of processors and millions of hours of computing time. Dashboards with up-to-date case numbers, maps of transmission locations and localised alerts require reliable and robust cloud environments to run around the clock.



NCI's resourcing in response to the COVID-19 pandemic extends beyond simply providing computing time. We also rapidly provisioned a cloud dashboard capability that has run for more than a year to provide graphical updates on the virus spread and transmission hotspots. CRISPER – the COVID-19 Real-time Information System for Preparedness and Epidemic Response – continues to share invaluable real-time case numbers, testing numbers and exposure site data that benefit public safety and lead to a more effective clinical response.

FUTURE BATTERY TECHNOLOGY

CASE STUDY



Every item and process in our lives is made up of technologies, discoveries and inventions that have been combined into a useful form. Every advance in creating a new material or method helps us gradually move society forward. In our electrified and energy-efficient future world, we will have high-capacity batteries that are lighter and safer than current ones, and energy-efficient computers made from silicon variants that have much lower energy consumption for each calculation. Professor Michelle Spencer from RMIT University and the Australian Research Council Centre of Excellence in Future Low- Energy Electronics Technologies (FLEET) is conducting research in these critical areas.

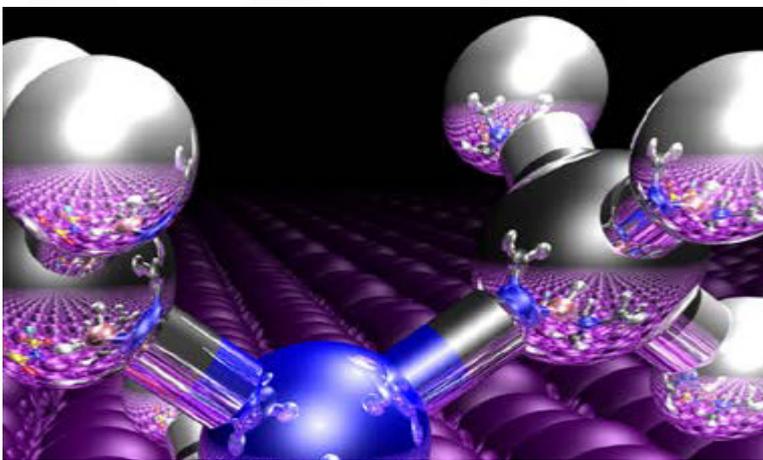


Image: Visualisation of Prof. Michelle Spencer's research at the molecular level

Professor Spencer uses the Gadi supercomputer to understand what is going on inside battery materials at the molecular level. The Gadi supercomputer allows the research team to run hundreds of simulations looking at complex, previously unresearched combinations of molecules.

Using specialised scientific modelling software running on thousands of processors at the same time, Professor Spencer can avoid wasting experimental time looking into unsuitable candidate materials. Furthermore, simulations allow the research team to learn about why certain behaviours occur in the materials of interest. This provides much greater insight and lets researchers everywhere build on and share their understanding of the fundamentals of the molecular interactions.

Efficiently searching for new materials that we can use to support a cleaner world requires the processing power and capabilities of supercomputers. NCI and Gadi give Australian researchers the performance they need to understand the fundamental characteristics of some of the materials that will soon power our planes, ships and computers.

VIRTUAL RESEARCH ENVIRONMENTS

Virtual research environments (VREs) are developed through coordinated engagement with research communities to maximise the use of NCI's compute and data infrastructure. NCI's national leadership in computational workflows on HPC systems and high-performance petascale reference data provides the only research platform in Australia capable of supporting intensive data analysis and simulation. While this is a collaborative effort with specific community organisations, one particular focus for our work is to integrate the community workflows and tools with our peak computing capabilities, as well as to interoperate with other communities at NCI.

- Climate Data-Enhanced Virtual Laboratory, a dedicated portal for accessing the CMIP6 modelling data, the Virtual Desktop Infrastructure, community user support and the ACCESS model control server.

- The Geoscience Data-Enhanced Virtual Laboratory (GeoDEVL) project is led by AuScope to provide a new class of services for Magnetotelluric (MT), seismic and geochemical samples.

- The Digital Earth Australia (DEA) project prepares vast volumes of Earth observation data of the continent and makes it available to governments and industry for easy use.

- NCI is the repository for the SE Asian set of the European Commission's Copernicus Earth observation data, collected by the European Space Agency and the European Organisation for the Exploitation of Meteorological Satellites' (EUMETSAT) Sentinel satellites.

- NCI has a long term involvement in the Australian All-Sky Virtual Observatory project led by Astronomy Australia.

SCIENTIFIC VISUALISATIONS

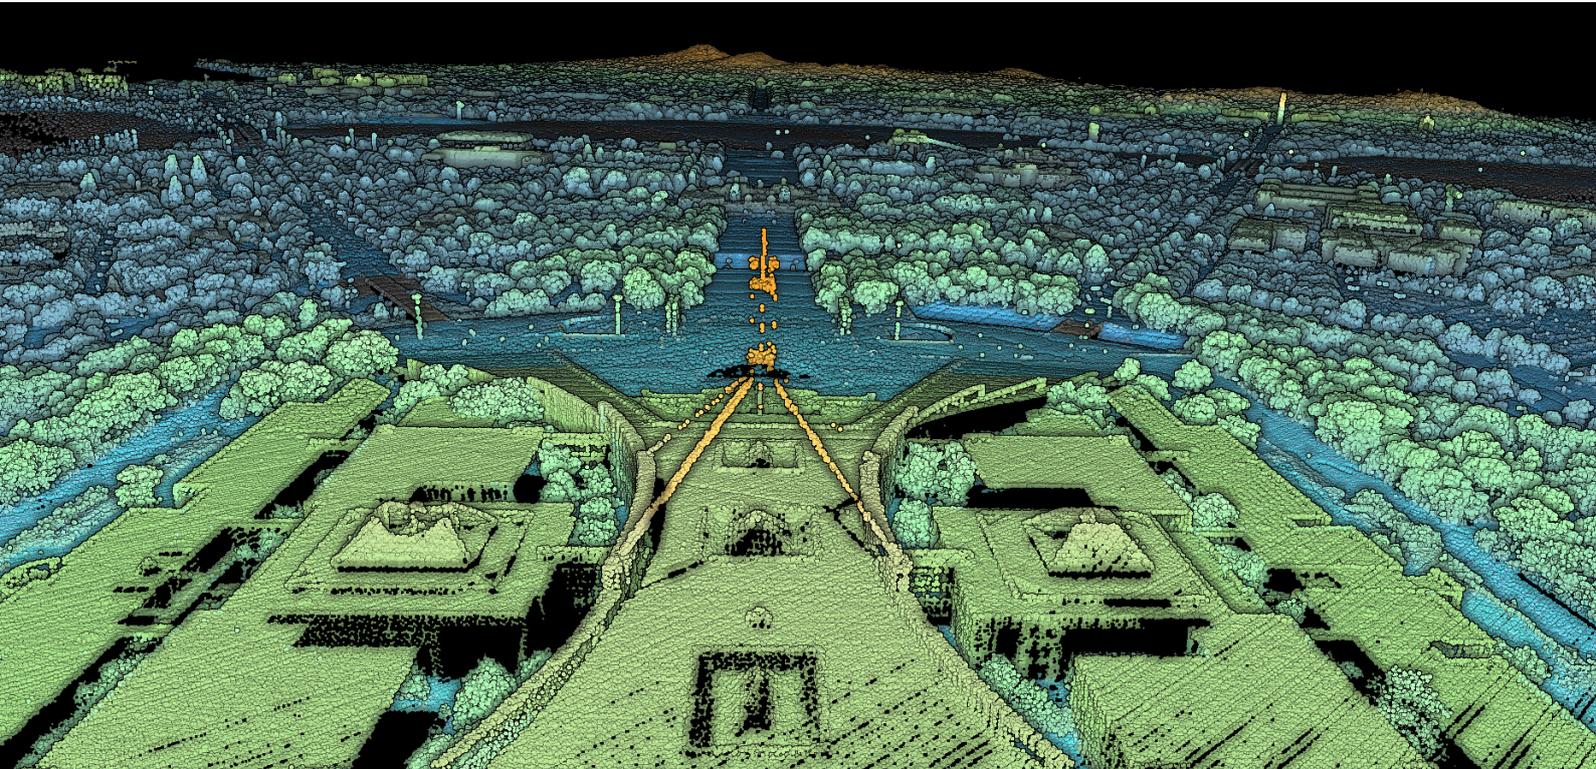


Image: Point cloud data of Canberra's Capital Hill, Ajay Limaye, VizLab NCI

NCI's scientific visualisations are a valuable tool that enables researchers to gain deeper insights into complex datasets and easily communicate their scientific results. Produced by the VizLab team, these visualisations aid researchers by helping them see their data in a new way, and more easily communicate aspects of the new insights with scientific peers and the general public.

The team works closely with both researchers and in-house computational and data experts to produce scientific visualisations for the research community. These have a particular focus on weather, climate, Earth Systems science, geophysics, and environmental science, and also cover areas such as astrophysics, materials science and medicine.

By leveraging the large data storage and computational capacity available at NCI, the VizLab team works with datasets hundreds of terabytes in size to create their visualisations. Using a combination of techniques, including video animation packages used in the movie industry and NCI-developed software, NCI's scientific visualisations are an important and striking element of some of the biggest computationally-intensive research projects in Australia.

USER SUPPORT

NCI's User Support Team provides expert advice, technical support and software installation services to the entire user community. From regular live Support Sessions for users with pressing questions, to dedicated support for priority disciplines, NCI's User Support raises the abilities and capabilities of our users.

NCI provides a helpdesk staffed by PhD qualified computational scientists to support users with administrative and technical issues to do with accessing and using systems and services at NCI.



The User Support function at NCI also provides in-depth documentation on the usage and technical details of each hardware and software platform that users can access. NCI's technical documentation repository is kept up to date and covers Gadi, Nirin, Data Collections, Training, Virtual Desktop services and more.

USER TRAINING

NCI continues to expand on the breadth and frequency of training for our cohort of Australian researchers. We recognise the need for Australian computational scientists to develop the skills necessary to make best use of HPC and HPD infrastructure.

Able assisted by the NCI User Support Team, NCI has been delivering a regular series of training webinars for beginners, with an emphasis on the fundamentals of using the Gadi supercomputer. In addition, the User Support Team has been holding virtual 'office hours' for NCI users each month. These sessions are a great opportunity for NCI users to get to know the Team, ask questions and highlight any challenges they have come across.



Pictured: Introduction to NCI training held at The Australian National University



COMPUTATIONAL SCIENCE ENHANCEMENTS

NCI's HPC Simulation, Scaling and Data Analysis Optimisation Team helps research groups understand the bottlenecks and inefficiencies in their code and plan for future growth in resolution or scale. The team can provide solutions for researchers to implement that save them time and computational resources and lead to better scientific outcomes. NCI's expertise in code optimisation starts with large climate and ocean models, and extends all the way to genomics, geophysics, fluid dynamics and astrophysics codes, as well as underlying tools beneficial to all disciplines.

For more than a decade, NCI has been supporting the Consortium for Ocean-Sea Ice Modelling in Australia (COSIMA) with the development and optimisation of some of Australia's most important scientific codes: the various forms of the Australian Community Climate and Earth-Simulation Suite (ACCESS). From regional to global scales, and with a particular focus on Southern Ocean simulation, the ACCESS model plays a central role in national daily weather forecasts, seasonal planning and long-term climate projections.

NCI has helped COSIMA implement improvements to the ACCESS models that have led to higher resolutions, more accurate handling of the biogeochemistry of ocean circulation and more efficient operations. The large data collections produced by and used for running software such as the global ACCESS models on tens of thousands of processors at once also requires specialised development work from NCI's team. Processing, storing, labelling and making the data available for open-access sharing is an important part of the computational science enhancements that NCI offers researchers.



MITIGATING EXTREME BUSHFIRES

CASE STUDY



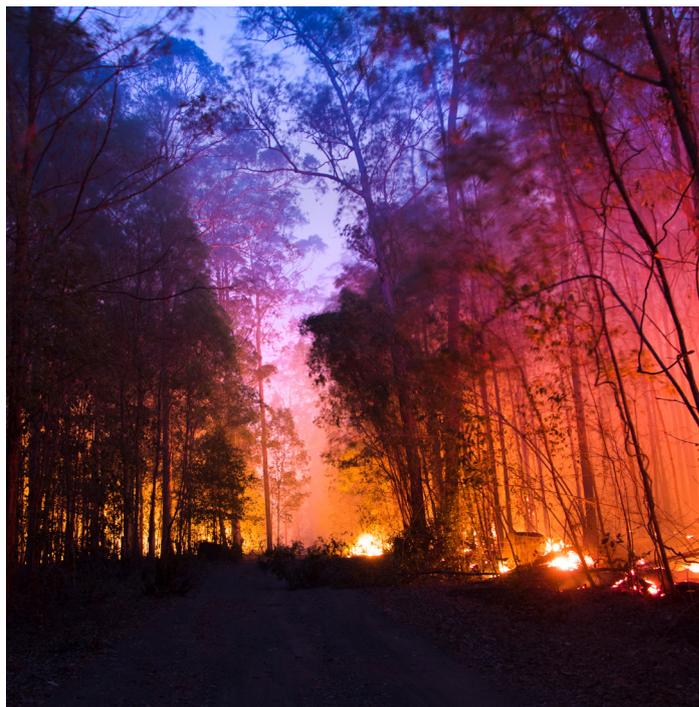
All fires start small, but some fires grow into large conflagrations that have the potential to cause massive destruction. The majority of damage caused by bushfires is from relatively short episodes as they escalate to 'extreme bushfires', involving widespread flaming and dense ember attacks. These fires are so fierce that they create their own miniature weather systems, such as fire-fuelled thunderstorms known as pyrocumulonimbus clouds.

By better understanding the forces and conditions that drive the transition of small fires into raging infernos, researchers hope to be able to better predict the outbreak of massive bushfires before they take place. Such research has significant implications for national security and prosperity, and has the potential to save lives and property. Professor Jason Sharples (UNSW Canberra) and his team are using sophisticated numerical models and observations to better understand how fires develop into extreme bushfires.

"Our research takes a deeper and more detailed look at the processes that drive bushfire propagation. These processes involve dynamic interactions between the fire and the surrounding environment, which comprises terrain, fuels and weather related factors", explains Professor Sharples.

Computational fluid dynamics (CFD) models have been used to understand interactions between the fire and the atmosphere at a finer scale.

"Based on this, we used the facilities at NCI to examine specific wind-terrain-fire configurations, and were ultimately able to demonstrate the existence of an unknown mode of fire propagation driven by pyrogenic vorticity", said Professor Sharples. The models that were used to produce these results are highly computationally intensive, and would not be possible to implement within project timeframes without the Gadi supercomputer.



VALUABLE GENOMIC DATASETS

CASE STUDY



Genetic medicine in coming years will benefit from the technical innovations that NCI and the Gadi supercomputer are now enabling. The Garvan Institute of Medical Research is refreshing its 4,000-strong database of genomes from healthy seniors – the Medical Genome Reference Bank (MGRB). Free from the markers of heart disease, cancer and diabetes, these healthy genomes form a robust comparison dataset for sequenced patient genomes that doctors might have.

Advances in the complex software controlling the hundreds of steps required have enabled reprocessing of all 4,012 MGRB genomes in one go. Reprocessing using the latest scientific understanding about human genomes allows researchers to have the most accurate data at hand for diagnosis and treatment. As genetic medicine becomes more central to our medical process, being able to compare patient test results with the rigorous baseline set out by the MGRB will be a key factor.



Reaching this number of processed genomes takes a lot of scientific effort. The sequencing data from thousands of genomes gets transferred from Garvan's Sydney laboratory down to NCI in Canberra, at which point the reprocessing begins. Thousands of tiny snippets of sequences get compared, lined up and combined into a long string making up the entire human genome.

From the snippets to the final sequence takes more than 40 different computational steps, all guided by the expertise of the bioinformaticians and programmers who built the software.

The final product is a treasure trove of valuable genomic data. The entire MGRB is securely shared with approved Australian and international genome researchers. The medical benefits to come from this modern dataset, built using the Gadi supercomputer's performance and filesystem speed, are only beginning to be realised.

COLLABORATORS



Supported By



Major Collaborators



Australian
National
University



Collaborators



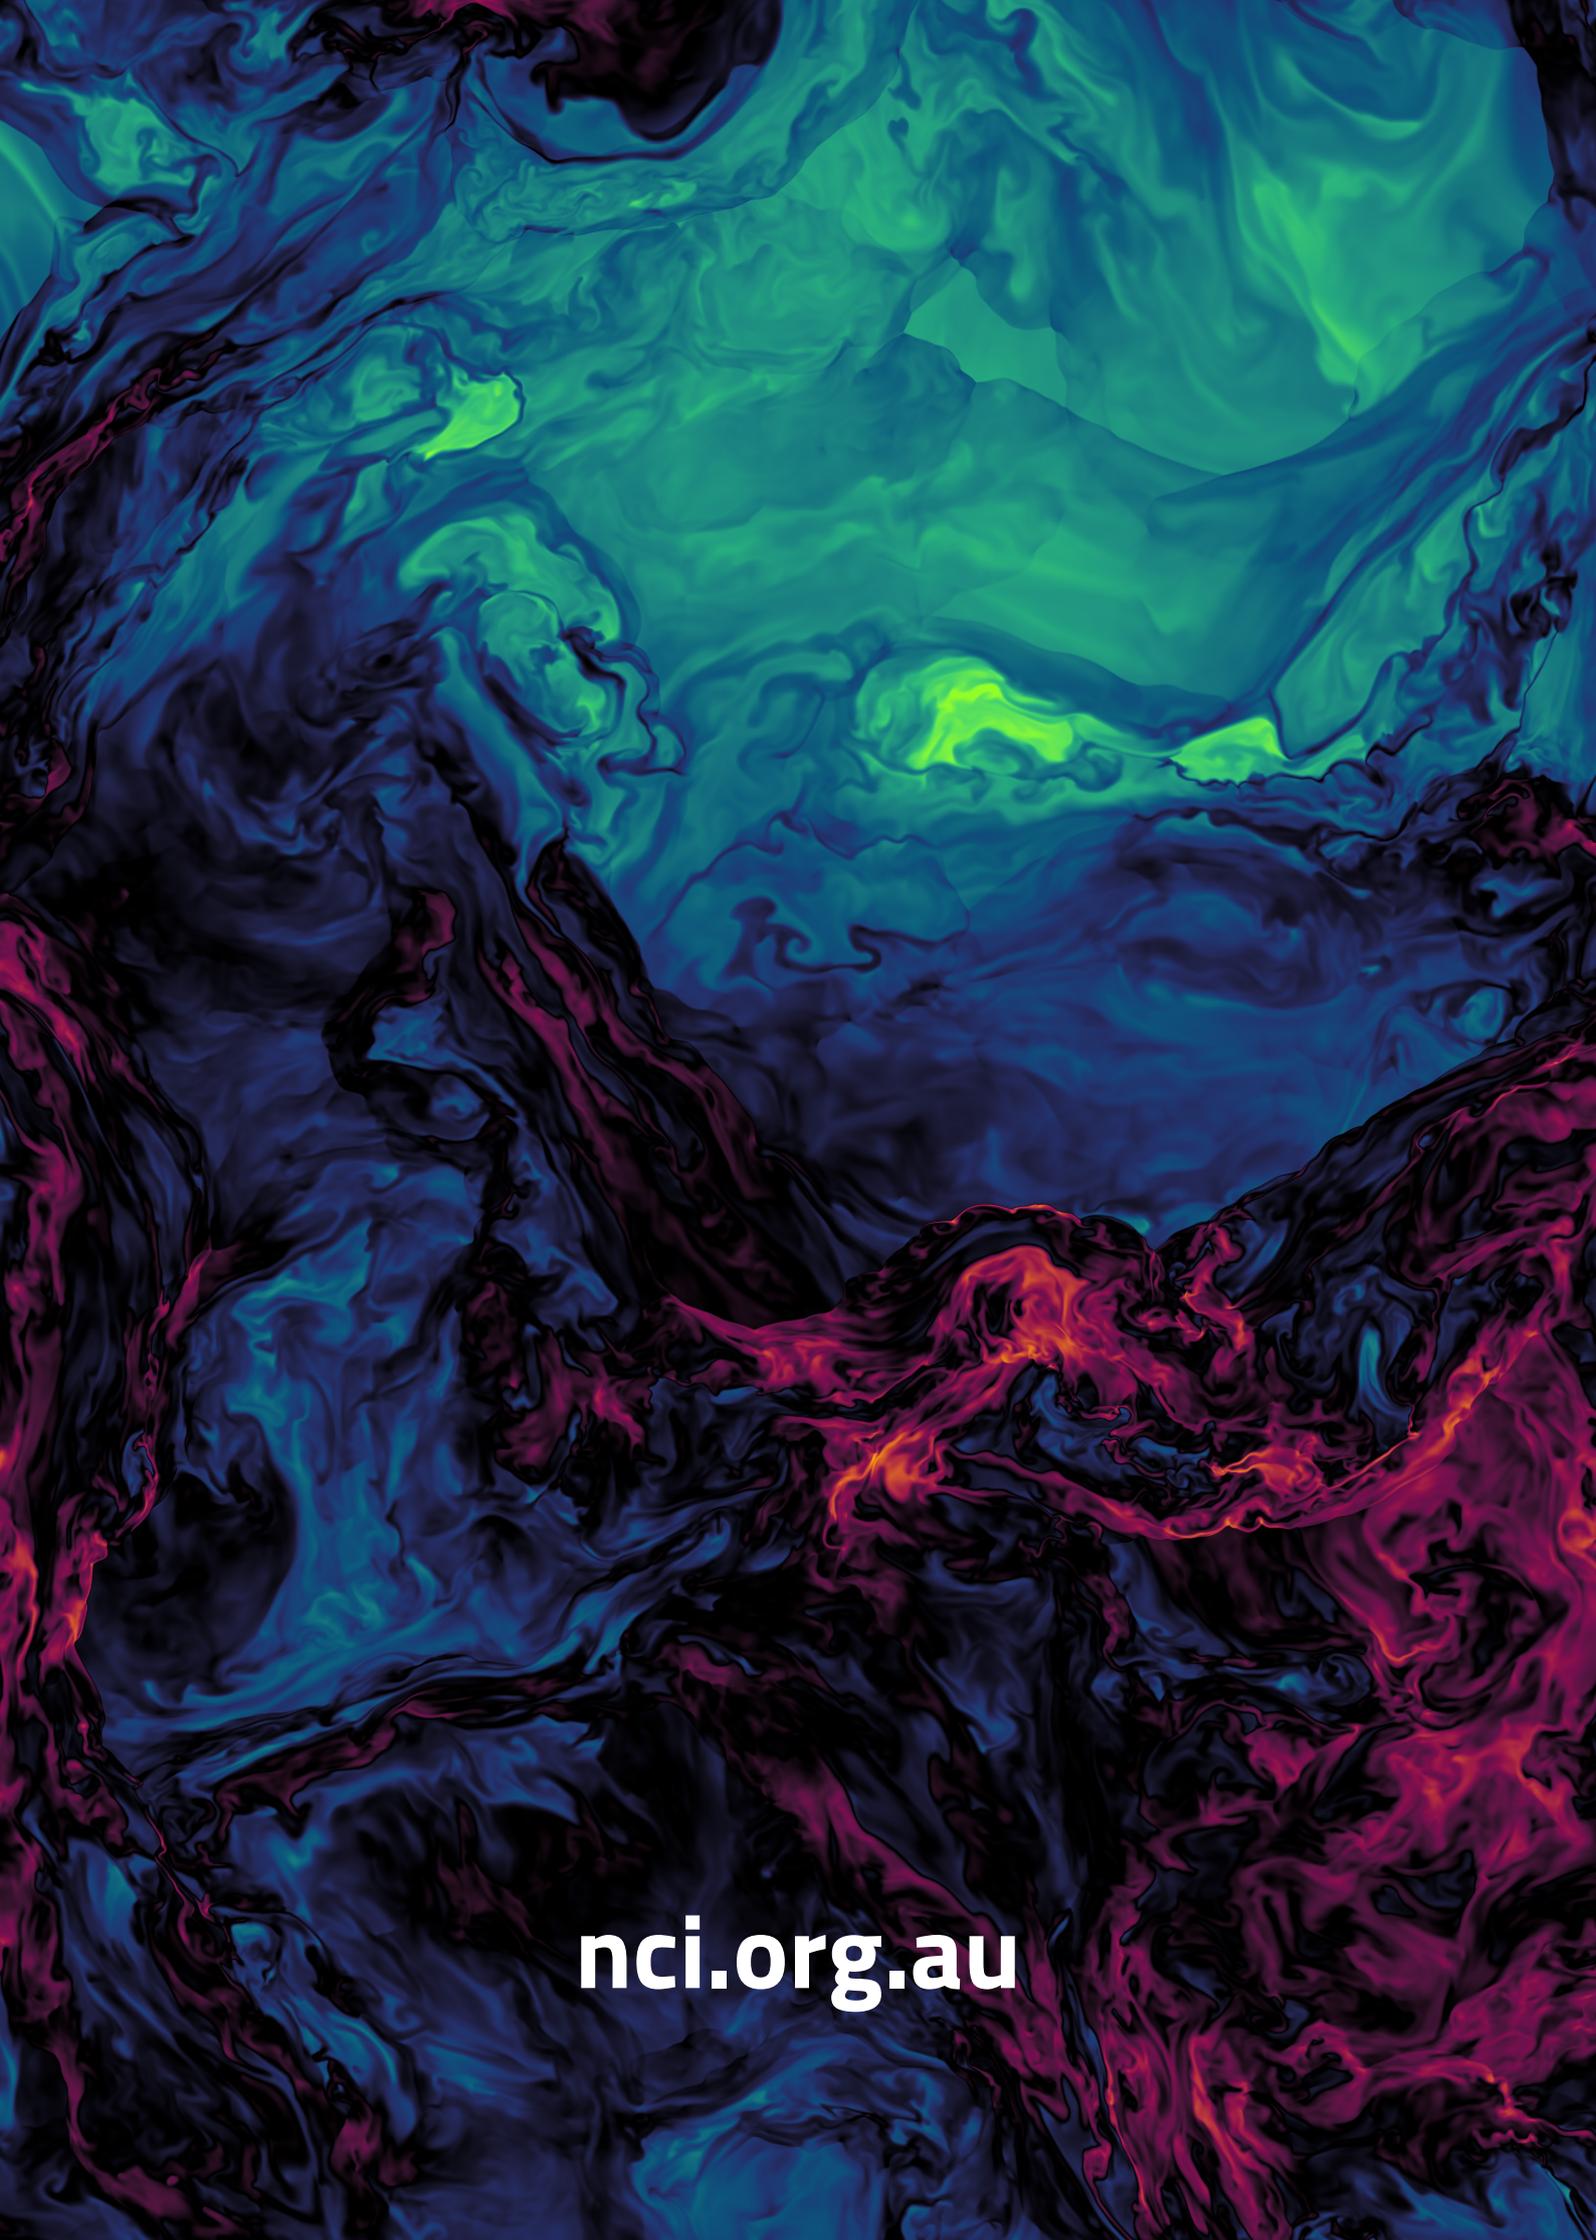
COLLABORATORS

Other Contracts



Merit Flagships





nci.org.au