SWINBURNE

CENTRE FOR ASTROPHYSICS AND SUPERCOMPUTING

# Moving beyond the limits of desktop computing to improve research potential in astronomy with remote visualisation and ultra-high resolution displays

Bernard Meade

Presented in fulfillment of the requirements
of the degree of Doctor of Philosophy

June 2018

Faculty of Science, Engineering and Technology
Swinburne University

# Abstract

The desktop computer is becoming a bottleneck for many astronomers. The ubiquity of the desktop in the astronomy research workflow means its limitations are often overlooked, or simply accepted as part of the normal lifecycle of information technology infrastructure. As the demands of modern astronomical instruments and computer simulations continue to grow, the capability of the local desktop computer is being noticeably outpaced.

With regards to visualisation and analysis workflows, this disparity impedes astronomers in a number of ways. This research investigates solutions to two key problems. The first is the ability of astronomers to display Gigapixel astronomical imagery at native resolution on standard desktop displays. The second is the ability to store and process data that significantly exceeds the capacity and processing power of the desktop computer.

The use of an ultra-high resolution display space is one way to address the first of these challenges. Such a display maintains the pixel density of a standard display, but can achieve a scale that approaches the imaging resolution of modern astronomical instruments. This enables astronomers to find small but important features within the context of a much larger image more rapidly than with a standard desktop display. A Tiled Display Wall combines many smaller displays into a single display surface, making it possible to display high resolution images, or many related images or applications at the same time, without necessarily having to reduce the scale of the images to such a degree as to render some features invisible. Collaborations requiring high-resolution imagery are also aided by the use of Tiled Display Walls, with several astronomers able to be physically present in front of the display. This forms the motivation of the first of the three research questions answered in this thesis; *Are ultra-high resolution displays needed to help astronomers make sense of imagery containing billions of pixels or more?* A user study of astronomers performing astronomical feature searches with a Tiled Display Wall showed that the display technology resulted in improved user performance and user experience.

More important than simply presenting digital content on a physically large, pixel dense display, is understanding when such a display is useful, and where the greater benefits are when presenting content in this way. Combining several displays with different capabilities into a *display ecology*, where displays show content most appropriate to them, aids data analysis and communication. This concept motivated the second question addressed by this research; *Can using the right display for the right digital content improve research outcomes?* Applying a display ecology consisting of a mix of advanced displays and standard displays to an astronomical observation campaign, resulted in improved

communication and collaboration between the astronomers.

The second challenge is caused by the inherent limitations of the desktop computer. Hardware restrictions such as RAM, hard disk space and the number of processing cores can be simply too small to be used meaningfully to tackle modern astronomical problems. Yet the desktop is a crucial part of the astronomy workflow, allowing researchers to generate graphical interpretations of data, and prepare these data for publication.

Cloud computing offers a viable alternative to local processing and storage, overcoming the inherent limitations of the local desktop computer by combining the resources of many commodity computers in the form of virtual machines. These virtual machines can provide a virtual hosted desktop that performs the same functions of the desktop on a local computer, but with the added benefits of direct access to cloud storage and a far larger pool of computing resources. With the advent of GPU-enabled virtual hosted desktops, we consider how the remote desktop experience can match or even exceed a local desktop experience. Convincing the astronomy community that a cloud-based desktop could rival the performance and experience of a local desktop requires more than system specifications and benchmarks. This forms the basis of the final research question considered in this thesis; *Does a virtual desktop providing remote access to a data centre with astronomical datasets provide a viable alternative to a local desktop computer?* A study of astronomers performing typical astronomy tasks showed that a virtual hosted desktop can perform as well as a local desktop computer for many astronomy applications.

Desktop computing will remain a crucial part of the astronomers day-to-day research workflow for the foreseeable future. Fortunately, inherent physical restrictions such as the size and resolution of the desktop display, or the processing power and disk capacity of a local machine, can be overcome with the use of advanced displays, such as Tiled Display Walls, and virtual hosted desktops. Astronomers are now able to move beyond the desktop and harness the power of cloud computing and the visual scale of ultra-high resolution displays, to improve their research outcomes.

# Acknowledgements

I think if I had known at the beginning how hard it would be to get this point, I might not have started. A journey like this can sometimes feel quite solitary, but I have had a great deal of help along the way.

First and foremost, I would like to thank my principal supervisor, Christopher Fluke. Chris has been an incredible supervisor, supporting and guiding me every step of the way. He is patient and generous with his time, and also meticulous and determined that this thesis would be my best work. I couldn't have done this without him.

Thanks also to my co-supervisors, Steven Manos and Richard Sinnott. Steve helped me find my feet in the research world, particularly in the early days of my PhD, and Rich has provided his extensive experience whenever I asked.

In 2009, my close friend, Allyson Calderazzo, told me about a cool online astronomy course at Swinburne University of Technology. That started my journey into formal astronomy study. I thoroughly enjoyed the Masters of Astronomy course, but I met Chris Fluke around that time and he suggested I might enjoy doing research. Leon Sterling, then Director of E-Research at Swinburne, supported my application to join the Swinburne Centre for Astrophysics and Supercomputing as a research Masters student.

If I regret anything in this journey, it is that I couldn't spend more time with the wonderful people at Swinburne's Centre for Astrophysics and Supercomputing. Many of them helped me out by participating in my research as test subjects, and coming to my presentations, and always being enthusiastic.

In particular, I would like to thank Jeff Cooke and Igor Andreoni for letting me be a part of the *Deeper, Wider, Faster* project. It was the most enjoyable research activity I worked on during the whole PhD. I got to be involved in some really cool astronomy research and play a small but useful part.

Thanks also to Dany Vohl, a fellow student also supervised by Chris, who was always ready to provide assistance whenever I needed it, and I needed a lot.

I have completed this research while working full time, and I would like to thank the University of Melbourne for supporting me, both in terms of time and access to resources. In particular, over the last two years, Stephen Giugni and the team at Research Platform Services have been incredibly encouraging.

On a personal level, a PhD has a big impact on relationships and I would like to thank my friends, particularly those in the Albatross Club, for sticking by me despite all the times I couldn't go to "The Dan O'Connell", and other social outings.

My family have supported my academic journey for as long as I can remember. They have tolerated me, listened to my boring descriptions of tiled display walls and cloud computing, and still managed to be interested (or at least seem so). They have all been a wonderful support. My parents, Brian and Betty, have always encouraged me to strive to do my best, and have been a wonderful example. In particular, my father's dedication to the importance of research has been an inspiration during the more challenging periods of my PhD.

My penultimate thanks goes to my beautiful children, Amy and Jacob, who have changed my life in so many ways. They have heard me say, "...as soon as I finish the PhD..." more times than I can count. Thank you for your patience and support.

Finally, to my beautiful wife, Anita. It is impossible to describe how much your love and support has meant to me, in all aspects of my life. I told you when we first met all those years ago that I wanted to get a PhD one day, and now, thanks to you, I have achieved that goal. You are now, and always will be, my beautiful girl.

# Declaration

The work presented in this thesis has been carried out in the Centre for Astrophysics & Supercomputing at Swinburne University of Technology between 2012 and 2018. This thesis contains no material that has been accepted for the award of any other degree or diploma. To the best of my knowledge, this thesis contains no material previously published or written by another author, except where due reference is made in the text of the thesis. The content of the chapters listed below has appeared in refereed journals. All papers appear as published without alterations.

- Chapter 4 has been published as "Are Tiled Display Walls Needed for Astronomy?" in Publications of the Astronomical Society of Australia
- Chapter 5 has been published as "Collaborative Workspaces to Accelerate Discovery" in Publications of the Astronomical Society of Australia
- Chapter 6 Section 6.5 has been published as "Evaluating virtual hosted desktops for graphics-intensive astronomy" in Astronomy and Computing

My contribution to these papers was to devise and evaluate each of the studies and write the papers, accounting for 87 percent of the final manuscripts. My co-authors contributed by providing research program oversight and supervision, as well as assistance with technical editing and participation in the *Deeper, Wider, Faster* observation campaign, accounting for 13 percent of the final manuscripts. Co-authorship indication forms are included in Appendix B.

*Bernard Meade*
*Melbourne, Victoria, Australia*
*2018*

*Dedicated to my beautiful family, Anita, Amy and Jake, who kept me going.*

# Contents

# Acronyms

**AARNet** Australian Academic Research Network.

**ACID** Astronomical and physics Cloud Interactive Desktop.

**AR** Augmented Reality.

**ASKAP** Australian Square Kilometre Array Pathfinder.

**ASVO** All-Sky Virtual Observatory.

**ATCA** Australia Telescope Compact Array.

**AWS** Amazon Web Services.

**CANFAR** Canadian Advanced Network for Astronomy Research.

**CAVE** CAVE AudioVisual Experience Automatic Virtual Environment.

**CDA** Collaborative Data Analysis.

**CLaaS** Cluster as a Service.

**CTA** Cherenkov Telescope Array.

**CTV** Communication Through Visualisation.

**CVL** Characterisation Virtual Laboratory.

**DBaaS** Database as a Service.

**DECam** Dark Energy Survey Camera.

**EVL** Electronic Visualisation Laboratory.

**FRB** Fast Radio Burst.

**GPU** Graphics Processing Unit.

**GUIs** Graphical User Interfaces.

**HPC** High Performance Computing.

**IaaS** Infrastructure as a Service.

**LSST** Large Synoptic Survey Telescope.

**MASSIVE** Multimodal Australian ScienceS Image and Visualisation Environment.

**MR** Mixed Reality.

**NCI** National Computational Infrastructure.

**Nectar** National eResearch Collaboration Tools and Resources.

**NICT** National Institute of Information and Communications Technology.

**NIST** National Institute of Standards and Technology.

**PaaS** Platform as a Service.

**SAGE** Scalable Adaptive Graphics Environment.

**SAGE2** Scalable Amplified Group Environment.

**SDSS** Sloan Digital Sky Survey.

**SKA** Square Kilometre Array.

**Strudel** Scientific Remote Desktop Launcher.

**STScI** Space Telescope Science Institute.

**TACC** Texas Advanced Computing Center.

**TAO** Theoretical Astrophysics Observatory.

**TDW** Tiled Display Wall.

**UE** User Experience.

**UltraHD** Ultra High Definition.

**UP** User Performance.

**VDI** Virtual Desktop Infrastructure.

**VHD** Virtual Hosted Desktop.

**VL** Virtual Laboratory.

**VM** Virtual Machine.

**VNC** Virtual Network Computing.

**VO** Virtual Observatory.

**VR** Virtual Reality.

# List of Figures

# List of Tables

# 1 Introduction

*Astronomy, as nothing else can do, teaches men humility.*
—Arthur C. Clarke

## 1.1 Introduction

Since the advent of the personal desktop computer, the day-to-day workflow of the modern astronomer has been linked to the computation device on their desk. A great many technological advances have expanded the usefulness of the desktop and laptop computers, and as such they have become an integral part the astronomer's toolkit. At the same time, humanity has been able to extend its reach further out in space, further back in time, across the electromagnetic spectrum and even into the realm of gravitational waves, thanks to the technological advances in astronomical observing instruments. This is made possible by the symbiosis of computing and astronomical instruments. The volume of data captured by modern telescopes is so vast that it cannot be comprehended directly, but rather it is interpreted by algorithms operating over clusters of computer cores, before being inspected by astronomers and shared with collaborators around the world.

Yet a critical cog in the astronomy research machine is struggling to keep pace with the evolution of astronomy research. The local desktop computer, and more increasingly laptop computer, often does not have the storage capacity or processing power to deal with the volume of data now required for scientific discovery. The images produced by modern instruments are tens to hundreds or more times larger than the display screens on the astronomer's desk.

The aim of this research is to investigate the technologies that provide opportunities to overcome the bottleneck that is the desktop computer. The two broad areas of investigation are:

- the use of ultra-high resolution display screens that more closely approximate the growing image sizes compared to local computer displays; and
- the use of a remote computing co-located with astronomical datasets to overcome the storage and computing power limitations of local desktop computers.

The research challenges of information visualisation addressed in this thesis can be expressed in three questions:

1. Are ultra-high resolution displays needed to help astronomers make sense of imagery containing billions of pixels or more? See Section 3.2 and Chapter 4.
2. Can using the right display for the right digital content improve research outcomes? See Section 3.3 and Chapter 5.
3. Does a virtual desktop providing direct (remote) access to a data centre with astronomical datasets provide a viable alternative to a local desktop computer? See Section 3.4 and Chapter 6.

## 1.2 Background

Display technologies, and the means of driving these displays, are often overlooked by the astronomical community because computing technology has largely kept pace with the development of astronomical instruments. However, over the past two decades, this situation has changed, slowly at first, and more rapidly in recent years. Now astronomical instruments routinely produce far more data than a standard desktop computer can accommodate, either computationally or visually.

There are two significant consequences to this development. Firstly, the ability of astronomers to visually inspect enormous images, or a multitude of images, is hindered by the use of screens that are a fraction of the size of the content being displayed. Secondly, the ability to process big data requires the co-location of compute and storage resources, which limits the astronomer's ability to utilize standard desktop applications on a local computer.

### 1.2.1 Ultra-high Resolution Displays

While computer displays have been steadily increasing in resolution, the sizes of images captured by current and future astronomical instruments have increased far more rapidly. Interpretation of this information typically requires inspection by an astronomer, but the presentation technology is often inadequate for the task. For example, the high resolution Carina Nebula image available from the HST Gallery archive is over 200 $\times$ larger than a 1920 $\times$ 1080 display.

The ability of astronomers to find critical information within large format images can be augmented by digital processing techniques, however, it remains the responsibility of the astronomer to define the criteria for these processes. Events that have yet to be identified may therefore be overlooked by an automatic system that might be captured by the trained astronomer's eye, provided it is not lost in the scale of the image because it depends on the astronomer to serendipitously zoom on to the appropriate part of the image.

Comparing many images simultaneously can also be aided by presenting the images side by side, much as the typical person might do when comparing physical photographs. Rather than maintaining the photos in a stacked arrangement, it would be beneficial if the photos might be laid out on an available surface to be examined side by side. Likewise, spreading related images across a digital display space can aid sense-making.

Furthermore, parallel inspection by multiple astronomers collaborating on a task can be improved by employing a suitable display environment. Sharing a display environment allows researchers to physically share the space in front of the display without getting in each others' way should allow them to draw others' attention as required.

It is important to use the most appropriate display for a particular type of content, especially when there are multiple formats under consideration. Typically, users choose the display that will allow them to present all of their content reasonably well, without considering the potential of multiple displays of different sizes and capabilities. A *Display Ecology* typically combines several displays, with each focused on displaying the most suitable content.

### 1.2.2 Remote Desktop Computing

High Performance Computing (HPC) facilities have evolved from typically hosting monolithic vector-based supercomputers to employing clusters of commodity computers to achieve massive parallelism for a much lower, more scalable and flexible price point. Using a remote computing facility is a familiar experience for many astronomers that require additional computation power or storage, yet the standard desktop computer remains a key part of their process. More than providing a gateway to the remote facility, the desktop computer is also used for generating graphical interpretations of data and preparing research for publication and presentation.

A variety of graphical applications have been developed to allow researchers to understand, interpret and present data, however, many of these applications themselves require a graphical user interface to operate. Remote computing facilities often support these applications by streaming the graphical window to the local computer, but this experience is often less than optimal, and looked upon by astronomers as a last resort.

Recent advances in networking and Graphics Processing Unit (GPU)-enabled virtual computing able to support streamed applications presents an opportunity to reconsider the remote desktop not as a last resort option, but potentially a powerful alternative to an expensive local workstation.

### 1.2.3 The e-Research Landscape in Australia

In Australia, the research community has access to a federated research cloud, facilitated by the National eResearch Collaboration Tools and Resources (Nectar)[1]. Established in 2009 as part of the Australian Government Super Science initiative[2], the Nectar Research Cloud has several nodes (availability zones) distributed around the country, hosted by research and research support institutions. Access to the service is a mix of merit-based and institutionally provisioned resources. At present over 12,000 researchers utilise over 30,000 virtual CPUs[3]. Primarily operated as infrastructure as a service (where users can create and manage their own Virtual Machines) based on Openstack[4], the Nectar Research Cloud provides Australian researchers with a scalable, cost-effective (it is free at the point of service) and collaborative platform for conducting research.

Nectar also funded a number of Virtual Laboratory (VL)s, including the All-Sky Virtual Observatory (ASVO)[5], which supported the following research programs: Theoretical Astrophysics Observatory (TAO)[6], Skymapper, AAO Data Central, Murchison Widefield Array, and the CSIRO Australian Square Kilometre Array Pathfinder (ASKAP)[7] Science Data Archive. These programs rely on the Nectar Research Cloud to provide access to federated telescope data and computation resources.

Similar programs exist around the world. The Canadian Advanced Network for Astronomy Research (CANFAR)[8] provides services such as compute processing, visualisation and analytics on massive datasets captured by major Canadian and international observatories,

[1] `https://nectar.org.au/`
[2] `https://www.education.gov.au/super-science-initiative`
[3] correct as of 07/02/2018
[4] `https://www.openstack.org/`
[5] `http://www.asvo.org.au/`
[6] `https://tao.asvo.org.au/tao/`
[7] `https://www.atnf.csiro.au/projects/askap/index.html`
[8] `http://www.canfar.net/en/`

for the CANFAR partners, including the Canadian universities and the Canadian Space Agency. In China, the AstroCloud Virtual Observatory[9] provides Chinese researchers access to similar services.

## 1.3 Outline of thesis

This thesis is a primarily research papers that have been published in astronomy-related journals or presented at peer-reviewed conferences. Following this Introduction, Chapter 2 reviews the literature in the areas of big data, cloud and remote desktop computing, and Tiled Display Walls, where multiple computer displays are used as a single, large display. Chapter 3 discusses the specific research questions that have been explored and addressed in the thesis. The publications included in this thesis are focused around these research questions and are presented as Chapters 4, 5 and 6. Finally, concluding remarks and further research directions are presented in Chapter 7.

It is important to note that due to the nature of a thesis by publication, some repetition of content between the papers and the corpus of the thesis – particularly in the introduction of key concepts and methods – has occurred.

### 1.3.1 Literature Review

The review of existing literature initially focuses on the big data challenges confronting astronomy, and the approaches some researchers and research institutions have taken to address them. The various meanings attached to big data are all present in the context of astronomy. Some instruments produce data in the multi-terabyte scale and beyond on scales of hours or days. Often the data comes from a variety of sources.

Big data in a broader research context essentially means that data is beyond the capability of the local technology available to most researchers. This could mean the processor speed or hard disk capacity of their local desktop computer is exceeded, or the number of pixels in their attached display screen.

Understanding the nature of the challenges in an astronomy context are essential to ensure sensible investment in infrastructure. However, the traditional approach to HPC facilities does not always provide the necessary flexibility that these big data challenges require.

A more flexible approach has emerged in the form of cloud computing, where diverse computing resources are deployed in data centres that can be accessed remotely. The

[9] `http://astrocloud.china-vo.org/`

main shift however, is that resources are not assigned based on for example a position in a queue, but on more ad hoc basis. Physical infrastructure is virtualised and provisioned as a range of services, with the most common being VMs.

Flexibility is a key feature of cloud computing that makes it suitable to many astronomy applications. Astronomical research data requires considerable innovation in technology workflows to accommodate the data at the big data scale, but these innovations don't always align well with inflexible HPC systems. Because cloud computing is intentionally non-usage specific, it can be more easily modified to suit new workflows, and be easily restructured as new access and use models arise.

Cloud computing can also be used for more than purely computational workloads. As there is a significant dependence on graphical applications to interpret data, it is important that a cloud computing solution be able to accommodate this aspect of the astronomer's workflow as well. The ability to run graphical applications and stream the graphical data to the astronomer's local desktop computer has been possible for some time, however, it is only in the last few years that a cloud-based desktop solution was able to provide the same level of graphical processing power as a local computer with an internal graphics card.

This advance brings cloud-based computing to the point where it is important for astronomers to consider if the traditional local desktop or laptop computer is the best way to support their research.

Cloud computing only overcomes one of the bottlenecks caused by the local desktop computer. The second is the ability to display the ultra-high resolution imagery being produced by modern telescopes or through simulations. With images that exceed the resolution of a standard desktop or laptop display by many orders of magnitude, it is important that astronomers consider the risk associated with overlooking crucial data because an image is simply too big to be thoroughly reviewed on a display that is inherently too small. This applies to collaborative research where several astronomers need to discuss image data, but especially when these discoveries are highly time dependent.

Tiled Display Wall (TDW)s are a collection of commodity display screens that can be coordinated to provide a single display space with a resolution of all the screens combined. TDWs became popular around 2009, however, complications such as cost and ease-of-use resulted in slower uptake than originally expected. However, with the drop in price of many principal components, TDWs now have the potential to become a valuable addition to support astronomer's research workflow and to improve science outcomes.

### 1.3.2 User Evaluation Methodology

While astronomers are not strangers to new technology, as for all researchers, a new technology must prove its worth reasonable quickly or it will be discarded as a passing novelty. To assess a new technology purely on a whitepaper description is unlikely to be sufficient for most researchers. Desktop and laptop computers have been a part of the user experience for most astronomers for the past few decades, and it is would require a clear advantage for most astronomers to consider changing their workflow.

The scale of modern astronomy problems has been addressed for many astronomers by shifting computation and storage workloads to remote data centres and HPC facilities. This has proven very successful and is now considered standard practice. However, to convince astronomers to use a resource like a TDW or a desktop provisioned by a cloud service requires more than equipment specifications, performance benchmarks and cost-benefit analyses.

Users need to experience such services first-hand to fully appreciate how they might augment their existing workflow. To this end, the approach taken in this research was to observe astronomers using technology for real astronomy workloads.

Using scenarios outlined by Lam et al. (2012), which are explained in more detail in Chapter 3, studies were conducted to investigate the use of technology by astronomers. This provided evidence to support the hypotheses that TDWs can enhance certain astronomy visual inspection tasks and a desktop streamed from a VM in the cloud can sometimes be a better choice than a powerful local computer.

### 1.3.3 Ultra-high Resolution Displays

Around 2009, several research institutions in Australia and many more around the world, began building TDWs for their research communities. The purpose of these expensive (at the time) but impressive displays was to present the new generation of high-resolution imagery coming from several research sectors including astronomy. However, there was little specific evidence that TDWs actually resulted in improved research outcomes, and as is apparent from the literature review in Chapter 2, there are few papers that attribute a specific research outcome to the use of a TDW.

The main opportunity for TDWs was in the presentation of images with hundreds of millions to billions of pixels, something well beyond the capability of a standard display. Similarly, multiple high resolution images and digital content could be displayed simultaneously, capitalising on the surface display area.

However, other opportunities are also worth considering, such as operating linked applications, working with remote collaborators and working with remote facilities, such as telescopes or computation facilities. TDWs might also have a place in training and quality control for improving instrument workflows.

However, for a TDW to be useful, it does not have to be the critical cog that means a discovery is made or not made. In fact, it would almost be impossible to be sure that a discovery only occurred because a TDW was used, or a discovery missed because a TDW had not been used. More likely, a TDW might increase the chance of a discovery, or the rate at which a discovery might be made, or reduce the chance of missing a discovery.

However, up until 2014, it was largely anecdotal evidence and hypotheses that supported any given university investing in a TDW.

To address the lack of evidence that a TDW could improve science outcomes in astronomy, an investigation was devised to test astronomers performing the kinds of tasks for which TDWs were designed. The study compared TDWs and standard desktop computers in a head-to-head scenario, where the same tasks were completed in both environments. The performance of the participants (which included non-astronomers for comparison) showed that TDWs can indeed be useful in astronomy in a qualitative and measurable way.

Perhaps more importantly, the study showed that the user experience when using a TDW to search for astronomical features in a large image was better than for a standard desktop display. When this result was combined with the user performance improvements, the argument for a TDW became more compelling.

However, there are several reasons TDWs have not become more widespread and are not more tightly integrated into astronomy workflows. For example, in 2009, TDWs were expensive to build and operate, and were not as user friendly as they needed to be for widespread adoption. The screens themselves were also limited, particularly in resolution but perhaps more visibly in terms of the screen bezel, the plastic frame surrounding the display surface. Many users found the bezels off-putting, though there was little evidence that the bezels actually reduced user performance.

The most significant limitation for TDW over the last few years has been the dearth of research drivers arising not only in astronomy, but in the wider research community. More recently however, new areas of astronomy research have begun to grow quickly, especially in time-domain astronomy. In this context, the ability to process a large volume of visual information quickly and accurately have encouraged some projects to investigate the possibility of using TDWs to improve their visual inspection workflows.

Chapter 4 addresses the question: ***"Are ultra-high resolution displays needed***

***to help astronomers make sense of imagery containing billions of pixels or more?"***. Section 4.3 includes the paper published as Meade, B. F., Fluke, C. J., Manos, S., & Sinnott, R. O. 2014, Publications of the Astronomical Society of Australia, 31

### 1.3.4 Collaborative Workspaces to Accelerate Discovery

The key value of a TDW is that it can present certain types of information in a way that no other display can. It provides not only physical scale, but does so without sacrificing pixel density, and therefore detail. In this way, it provides an opportunity to see fine detail within a much larger image where understanding the detail requires the larger context. However, sometimes it is not necessary for these two capabilities to be used together. For example, where the feature resolution is not important, but the scale might be useful for the purposes of collaboration. In such a circumstance, it would not be appropriate to use a TDW as a generic display space, but rather use it as part of a wider display ecology.

The *Deeper, Wider, Faster* program is a contemporaneous multi-facility observing campaign designed to proactively capture fast transient astronomical events. An opportunity arose to support this endeavour, and observe the use of a TDW being used as part of a visual inspection workflow where time-critical image inspection was key to the success of the program.

The participating astronomers were unfamiliar with TDWs and agreed to participate in the study. The purpose was to determine if the user experience and user performance was improved in some measurable way due to the use of the TDW. As the study was being prepared, it became clear that some content intended for the TDW was better suited to another large-scale display. The curved projection display provided an immersive experience but with a far lower resolution than the TDW. However, the potential fast transient candidates were very small in size, and an automatic process was created to extract them as small thumbnail images. Inspecting tens to hundreds of these thumbnails was most easily accomplished by several people working in parallel and in close proximity, hence the scale of the curved projection display greatly improved the throughput of the inspection team.

The first of these observation runs that used the TDW and the curved projection display provided many lessons that were implemented for the second run. For the second run, a number of changes were made to the display ecology as it was now called, which again allowed the visual inspection workflow to evolve even further. Subsequent runs have seen the display ecology evolve in response to lessons learned from each preceding run, such that the more recent incarnation bore only a passing resemblance to the original configuration. However, this is the nature and value of a properly designed and refined display ecology.

Chapter 5 addresses the question: ***"Can using the right display for the right digital content improve research outcomes?"***. Section 5.2 has been published as Meade, B., Fluke, C., Cooke, J., et al. 2017, Publications of the Astronomical Society of Australia, 34

### 1.3.5 Using Cloud Computing to Support Virtual Hosted Desktops

In Australia, the Nectar Research Cloud has stimulated a number of new research activities that might have not had the opportunity to grow without it. Within astronomy, the ASVO has successfully supported several research programs. For example, the TAO is a VL to provide researchers with access to mock galactic survey data for the purposes of developing analysis pipelines and cosmic simulations. The tools provided by the VL allow astronomers to create simulated images of the sky and custom observer light cones without requiring programming knowledge (Bernyk et al., 2016).

#### Research Cloud Data Communities

Cloud computing is a relatively recent phenomenon within the research community, and the virtual labs have been critical in fostering adoption. They achieve this by providing access to datasets, tools and pipelines common to their research domain. However these funded projects cater to larger research communities and are therefore not always suitable or available for all research domains, so alternative approaches are needed.

A data community forms around significant datasets that have high value to many researchers. This activity is well supported by the Nectar Research Cloud. This is achieved by allowing researchers to collaborate irrespective of their location, by allowing the sharing of computation and data storage resources. This can be managed by the researchers themselves, through standardised, secure interfaces.

Perhaps the most important feature to support the development of the data communities is the ability for researchers to "fail fast", meaning VMs can be created and destroyed as required, without additional financial impact. This leads to more rapid development cycle of build, review, improve.

Finally, the Nectar Research Cloud is a scalable, flexible infrastructure that is designed around the needs of big data computing, with thousands of cores and petabytes of storage, all connected via a robust, high-bandwidth network and housed in enterprise-grade data centres.

More information about this study can be found in Chapter 6 Section 6.2, and published as Meade, B., Manos, S., Sinnott, R., et al. 2013, in THETA 2013, Hobart, Tasmania, ©2013 THETA: The Higher Education Technology Agenda

**Seeing the Big Picture: A Digital Desktop for Researchers**

Tapping into a resource like the Nectar Research Cloud, or any other cloud service, is relatively easy for many technically capable researchers, but not for all. Virtual Labs provide a valuable mechanism to utilise cloud resources, but this is not available to all. Command line interfaces do not suit everyone, especially those more comfortable with their familiar graphical desktop interface like Microsoft Windows, Mac OSX and Linux desktops such as KDE and Gnome. Such Graphical User Interfaces (GUIs) provide an intuitive way to interact with window-based applications that many researchers require for their workflows.

Like any other application, the desktop interface runs on the computing hardware, but in the case of the desktop, it manages the interface elements of the supported software. As such, the desktop itself can be run on virtual hardware and be operated remotely as a virtual desktop. When this desktop is provisioned within a cloud service, it is referred to as a Virtual Hosted Desktop (VHD).

As a VHD can provide a graphical interface to a remote computer, it is possible for such a desktop to drive the screens of a TDW. This would reduce the upgrade burden on the TDW as the computers used to drive the screens of the TDW only need to be powerful enough to display a stream of pixels from a remote server. With a suitable network, more powerful VMs could be linked to the TDW as and when required.

In 2015, few cloud services provided GPU-enabled VMs. An investigation was conducted using a commercial server solution that supported GPU-enabled VMs. The aim of this study was initially two-fold.

As a large TDW is not always readily available, the first challenge was to determine if an Ultra High Definition (UltraHD) (3180 × 2160 pixel) display had similar benefits as a TDW over a standard display, albeit lessened due to the reduction in pixels. To this end, the initial TDW study was repeated using the UltraHD display and confirmed the earlier results.

The second aim of the study was to determine if a GPU-enabled VM could drive an UltraHD display satisfactorily. A third aim also arose as it became clear that the research network in Australia needed to be capable of providing sufficient bandwidth with low latency and be robust, if GPU-enabled VMs were to be used to drive one or more

UltraHDs in a TDW configuration.

This study is presented in more detail in Chapter 6 Section 6.3, and published as Meade, B., Manos, S., Sinnott, R., et al. 2015, in THETA 2015, Gold Coast, Queensland, ©2015 THETA: The Higher Education Technology Agenda

**Evaluating Virtual Hosted Desktops for Graphics-intensive Astronomy**

Shortly after the study in 2015, several cloud services, including Amazon Web Services (AWS) and the Melbourne Node of the Nectar Research Cloud (hosted at the University of Melbourne, Australia) began offering GPU-enabled VMs. In 2016 it was possible to create a GPU-enabled VHD that operated with low enough latency as to appear almost indistinguishable from a local desktop. However, while some research had been published about the potential of using VHDs capable of handling 3D graphics, no research had specifically investigated the user experience of VHDs running graphically demanding astronomy software.

Based on system specifications and benchmarking to assess the impact of virtualisation, it seemed likely that a GPU-enabled VHD would closely match the performance of a local computer of a similar specification. However, the question remained about whether the user experience would be acceptable. Based on the fact that many astronomers are familiar with remote telescope operations being conducted over Virtual Network Computing (VNC) connections, and/or they are familiar with the need to use X11 Forwarding for graphical windows from HPC clusters, it was reasonable to expect most astronomers to consider a VHD as acceptable under certain circumstances, e.g. where an alternative was not available. The promise of the GPU-enabled VHD was that it need not be an option used as a last resort.

A study was undertaken to investigate the viability of a VHD supported by GPU acceleration as an alternative to a local desktop of similar specifications. The study recruited several a cohort of astronomers and had them complete astronomy-related tasks on both local desktops and GPU-enabled VHDs. The user performance was monitored throughout the tasks and aligned with the self-reported user experience to provide a clearer picture of the viability of the VHD.

Chapter 6 addresses the question: ***"Does a virtual desktop providing direct (remote) access to a data centre with astronomical datasets provide a viable alternative to a local desktop computer?"*** Section 6.5 has been published as Meade, B., & Fluke, C. 2018, Astronomy and Computing, (accepted April 4, 2018).

## 1.4 Summary

The desktop computer continues to be a key part of the astronomer's workflow, yet there are reasons to consider the line between the local and virtual more fluid than ever before. The limitations of the local computer and its display are possible to overcome, and in some cases, it is imperative to do so.

This research presents some of the technologies that overcome the limitations of a local computer. In particular, TDWs provide a mechanism for viewing ultra-high resolution astronomy imagery, which can be critical in some areas of astronomy, e.g. time-critical visual inspection workflows.

Finally, with significant increasing numbers of datasets hosted in data centres that are co-located with compute resources, it is now possible to shift the entire astronomy compute and visualisation workflow to the data centre as well, with the local desktop or laptop computer functioning solely as a window to the remote resources. Such explorations form the research questions that have driven this research.

# 2

# Literature review

The literature supporting this research can be broken into four parts. The first section looks at the broad challenges associated with *big data*, with a particular focus on its impact in astronomy, both now and in the near future.

Following on from big data, the second section focuses on the need for cloud computing and the state of the supporting technologies is in the astronomy context. This leads directly to the third section, which reviews the research into using cloud computing for visualisation tasks, and in particular, virtual desktop technologies and their use in supporting astronomical research.

The final section looks at the literature behind the development of TDWs. The underlying hardware to make TDWs possible has developed over several years. Tiling and synchronizing displays are a logical consequence to solving the challenges of immersive visualisation.

We note that the underlying network and infrastructure management technologies are outside the scope of this research.

## 2.1 Big data

The term big data refers not only to the amount of data being considered, but also to the importance placed upon the salient information contained therein. In fact, volume is only one of several features ascribed to big data. In the past, data served to explore testable hypotheses. However, in recent years, the data itself has become the starting point for developing a hypothesis.

This shift in perception of data collections was identified by Laney (2001). This was the first attempt to lay out the features of what would become known as big data. Big data is commonly considered in terms of 3 V's; Volume, Velocity and Variety. Volume was quite simply the amount of data, while Velocity refers to the speed with which the data is produced. The final V, Variety, refers to the fact that the data may no longer be coming from a single source, but rather multiple, often unstructured and hence heterogeneous sources.

Using these terms, “big” takes on different meanings depending on the context. In terms of volume, “big” means the number of bits representing the data exceeds some sort of measurement of capacity. This could mean something simple like it exceeds the local hard disk space, or something less obvious, such as exceeding practical inspection by traditional means.

As sensors and simulations increase in power and performance, the rate at which data is captured can quickly overwhelm typical means of processing. The Velocity of the data becomes “big” when extraordinary means of processing the data become necessary to understand it. This is the case for real-time data processing.

Data is now commonly linked to related but dissimilar datasets. This Variety of sources can yield important information not only about the phenomenon being studied, but it can also suggest other lines of inquiry. IBM among others has also suggested that a fourth V, Veracity needs to be considered[1]. This parameter considers the uncertainty in the captured data, trust in the source and value of the analysis (Ward & Barker, 2013). These authors consider several definitions, noting that magnitude always plays a significant part, however they settle on their own criteria of Size, Complexity and Technology. Size and Volume are somewhat interchangeable in this definition. Complexity includes structure and behavior, which incorporates Variety and Velocity, but also includes permutations of datasets. The Technology component is the key difference, embedding the idea that the processing of complex datasets is highly dependent on the technology being used.

Chen et al. (2014) adopt a less technical, and possibly more practical, definition of big data. They argue that the critical feature of big data is that it is untenable on traditional resources, e.g. it cannot be stored or processed on local computers. This intrinsically captures the 3 V’s and also the need to overcome technological limitations. Rather than use Veracity as a fourth V, the authors choose instead to use Value, which represents the benefit of capturing voluminous data with low-density yield. The authors also discuss the value of big data in terms of a data pipeline - generation, acquisition, storage and analysis. This is similar to the view presented by Zaslavsky et al. (2013) who refers to this value as the “semantic challenge” to find the “hidden gold”.

For some time it has been obvious for astronomers that the datasets captured by

[1] `http://www.ibmbigdatahub.com/infographic/four-vs-big-data`

modern telescopes or produced in modern simulations would present a huge challenge. Brunner et al. (2002) published a paper entitled "Massive datasets in Astronomy" which discussed the growing challenge of finding valuable information in enormous datasets. This led to the concept of Virtual Observatory (VO), first presented to the astronomy community in the Decadal Plan published by the National Research Council (U.S.) (2001). The principle underlying a VO was to provide access to accumulated data from telescopes to researchers, rather than requiring new observations of the same data. While the term big data had not yet been introduced to astronomy, Brunner et al. (2002) correctly predicted that the future of astronomy lay in data mining.

The shift toward data-intensive research was established by Jim Gray and published by Microsoft Research in a book called "The Fourth Paradigm" (Hey et al., 2009). This concept (after the first three paradigms of theory, observation and simulation) was further described in Bell et al. (2009), who also highlighted the decision to make the Sloan Digital Sky Survey (SDSS)[2] data available through VOs. While Bell et al. (2009) only touch on the coming "cloud", the authors identify the need to properly train researchers to successfully engage with the so-called Fourth Paradigm.

Processing massive datasets requires specialized techniques, often collectively called "data mining". Ball & Brunner (2010) offer a simple definition of data mining as the "...act of turning raw data from an observation into useful information". Taking this further, the authors describe how leveraging advanced technology is key to deriving human learning from an otherwise unassailable volume of data. The definition of data mining and its link to big data is further described in K.U & David (2014).

There are several useful examples of big data in science and more specifically in astronomy. Murphy et al. (2006) describe the use of VO tools to help cope with data from the Australia Telescope Compact Array (ATCA)[3]. The authors cite data transfer and storage along with processing and visualisation as key big data obstacles that can be overcome with VOs.

Brescia et al. (2010) describe a project called DAME (Data Mining and Exploration)[4], which was designed to allow astronomers to utilize standard scientific gateways to access and process massive datasets. These methodologies have given rise to Astroinformatics (Borne, 2008, 2010), which combines computer science and astronomy and knowledge discovery in databases (KDD; Brescia et al., 2012a,b; Ball, 2012a). Astroinformatics is the natural consequence of data-driven astronomy, where research is conducted using data

[2] `http://www.sdss.org/`
[3] `https://www.narrabri.atnf.csiro.au/`
[4] `http://dame2.na.astro.it/`

mining algorithms and massive datasets.

Cloud Computing is increasingly a key utility when dealing with big data. Wiley et al. (2011) employ the Hadoop[5] implementation of MapReduce (Dean & Ghemawat, 2004) to improve the image processing pipeline for SDSS image coaddition. This was essential in order to improve the signal-to-noise ratio for feature detection in image sets. A more detailed review of cloud computing appears in the following section of this chapter.

Feigelson & Babu (2012) provide an excellent overview of big data in astronomy. Their paper highlights the necessary shift in astronomy towards all-sky surveys, such as SDSS and the Large Synoptic Survey Telescope (LSST)[6] across a broad range of frequencies. Juric & Tyson (2012) provide more detail on the scale of the challenge, which they call "Petascale Optical Astronomy". This data is available in VOs to anyone to support their research. The use of individual resources to support individual researchers is diminishing. In fact, the Space Telescope Science Institute (STScI)[7] reports that more papers are now published based on existing archives than on new collections[8]. More specific information about the surveys themselves and the objects contained therein can be found in Mickaelian (2016).

Using this data remains a key challenge and Berriman & Groom (2011) argue that more needs to be done to improve the capabilities of the researchers probing these datasets. The authors identify the Software Carpentry[9] program as a way to combat poorly designed code that greatly hinders the use of massive datasets available through VOs. VOs can also provide built-in tools for handling big data, such as those described in De Pascale (2013).

While not explicitly dealing with big data, Goodman (2012) discusses the higher-dimensionality of new datasets and the need for astronomical phenomena to be considered in a greater context that may have been the case before. This links directly to the third V of the Gartner big data definition - Variety. Heterogeneity of data is increasingly becoming part of the big data domain.

Schlegel (2012) provides a rather intriguing take on the fluidic nature of the elements that are used to define big data. In this letter, he reminds the reader that big data is sometimes seen as a future state without the benefit of context. The author argues that while the data captured from LSST will be big by today's standards, it is not really appropriate to call it big data as technology development is currently keeping pace well with the predicted demand in this particular context. Hence it is worth reflecting on the

[5] `http://hadoop.apache.org/`
[6] `https://www.lsst.org/`
[7] `http://www.stsci.edu/portal/`
[8] `http://archive.stsci.edu/hst/bibliography/pubstat.html`
[9] `http://software-carpentry.org/`

earlier paper Ward & Barker (2013), where the technology component needed to be placed in the appropriate context, i.e. several years into the future.

Finally, summaries of the challenges of big data in astronomy, and the infrastructure and compute resources needed to solve them, can be found in Zhang & Zhao (2015) and Younge (2016).

There are many consequences of big data and considerable research has and continues to be conducted to find better, more efficient and more cost-effective ways of dealing with the data deluge. Of key importance is the efficiency of the underlying computing and storage infrastructures. At the petascale and beyond, investment in resources is significant and it is impossible for many research institutions to manage this alone. The traditional approach of buying monolithic computational resources optimised for a small number of important use cases risks leaving emerging research areas behind, and hence unable to establish a foothold in the research landscape. To address this challenge, universities are looking to more generalised computing options, that are able to accommodate a wider variety of methodologies and workflows. Of particular importance is the ability of such resources to be used for unanticipated requirements, either due to scale or the complexity of the environments. Large-scale, usage-agnostic resources that can support highly variable demand are possible through cloud computing services.

## 2.2 Cloud Computing

The origin of the term "Cloud Computing" is unclear, though the essential aspects of it seem to stem from the principles behind mainframe and Grid computing. In this model, the "cloud" component refers to a collection of resources provided to users without requiring them to have explicit knowledge of the underlying infrastructure. The generally accepted authoritative definition comes from the National Institute of Standards and Technology (NIST)[10] from the U.S. Department of Commerce.

NIST define cloud computing as follows:

> *Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model is composed of five essential characteristics, three service models, and four deployment models."* (Mell & Grance, 2011, p.2).

[10] `https://www.nist.gov/`

The authors then go on to define the characteristics, service and deployment models. The above definition is valuable as it provides a consistent basis for comparison of cloud computing deployments, particularly in the form of Infrastructure as a Service (IaaS), though many other services are worth considering such as databases, platforms, and clusters DBaaS, PaaS and CLaaS.

Infrastructure as a service is the provisioning of virtual servers, or VMs, where the computer in use is actually the result of software running on another computer, with the "hardware" components being emulated by software (Daniels, 2009). The virtualisation of computing resources has been around for many years (Buzen & Gagliardi, 1973; Goldberg, 1973) but has only recently become widespread due to the emergence of cloud computing. Because the virtual computer is the product of software, it is possible to "save" the VM in the same way one would save any computer file. Called an "image", this copy of the VMs can be transferred across the network, duplicated, archived, and re-instantiated when required.

A more recent version of this approach to virtualisation, sometimes called *lightweight virtualisation*, is *containerisation*. Rather than virtualising both the operating system and applications, as in the case of a VM, a *container* holds only the application, and runs in protected memory on an operating system with which the applications are compatible. Both approaches are useful in supporting scientific reproducibility (Chirigati et al., 2016; Emsley & De Roure, 2017). Details about the architectural and performance differences between VMs and containers can be found in Morabito et al. (2015). An example of containers in astronomy is *cyberhubs*[11], which provides interactive analysis of shared astronomy data using Jupyter[12], a python-based online notebook, and Docker[13], a leading container-hosting platform (Herwig et al., 2018). A detailed description of containers for the distribution of astronomy software can be found in Morris et al. (2017).

Integrating cloud computing into scientific workflows has been challenging because of the wide variety of demands and shifting usage patterns. Where business tends to have constant demand or highly predictable peaks and troughs, research demands greater flexibility. Traditional HPC is typically deployed as a monolithic static resource, where users' jobs are queued to access the resource. While this suits certain types of activity, it excludes a great many more. As demand for compute grows alongside big data, flexibility becomes key. Hoffa et al. (2008) were among the first to actively investigate cloud computing in the astronomy community. They compared VMs and virtual clusters with local compute

[11] `https://github.com/cyberlaboratories/cyberhubs`
[12] `http://jupyter.org/`
[13] `https://www.docker.com/`

resources at the University of Chicago. They found that while local resources could be sufficient in certain circumstances, they were not scalable. Virtual environments provided scalable solutions, but care is needed to manage storage, as this could easily undermine the performance.

Even before the birth of the Internet (or ARPANET as it was known in 1969), the idea that one day computing might be treated as a utility had already begun to form (Parkhill, 1966). In July 1969, Leonard Kleinrock, research scientist at the University of California, LA, identified the coming "computing utility" that would exist much like electricity or telephony (Kleinrock, 2005). It still took the Internet over 20 more years to enter mainstream. It is worth noting that the vision of utility computing is not a recent phenomenon.

Like any utility, cloud computing offers a way of accessing compute on demand, rather than purchasing in anticipation of need. This means that consumers only need to pay for what they use, and resources can be shared and distributed for optimal use. Buyya et al. (2008, 2009) discuss the potential for cloud computing as a compute utility. The progress from mainframes to cluster to grids to cloud can initially seem quite logical, yet Buyya et al. (2008) explain there is a difference between the first three and cloud computing. They argue that the virtualized nature of cloud computing is a significant deviation from the "bare-metal" approach of its predecessors. Buyya et al. (2008) focus heavily on the market drivers that make cloud computing a viable commercial activity divorced from the purpose of use. They compare several commercial providers and the markets they serve, and conclude that cloud computing is the best placed technology to provide the "fifth" utility (Buyya et al., 2009), after water, gas, electricity and telephony.

The principles of cloud computing and compute as a utility are well aligned. Armbrust et al. (2009) described cloud computing and its key elements. In this paper they begin by identifying three structural elements of importance. Firstly, that the resource *seems* infinite to the user. The second is that there is no requirement to purchase hardware ahead of use provided the resources purchased can cater for peak demand or be scaled down outside of peak usage. Lastly, that there is no commitment required of the user beyond the actual use, i.e. they do not necessarily have to be technically adept.

Armbrust et al. (2009, p.6) also argue that the reason cloud computing has become viable now is due to the convergence of technologies that make large-scale data centres possible. Commodity resources and the ubiquity of the Internet provide low cost expandability to meet and deliver demand as and where it is needed. Elasticity means resources are shared effectively and efficiently. Yet perhaps the most compelling aspect of cloud computing is the fact that it is use-agnostic. Resources are not committed to a single task. For example, hardware used for computer-generated imagery for a movie production house can be almost instantly repurposed to run financial simulations for banks and then reused again to run HPC jobs for a university.

Juve et al. (2009) compared the performance and cost of running scientific workflows in the cloud with traditional HPC solutions. They found that where the underlying resources and interconnect are similar, the performance is comparable. However, the cloud solution has the benefits of ad-hoc provisioning and scaling, through a "pay-as-you-go" business model, acknowledging the cost savings of co-location of data and compute resources to reduce ingress and egress fees. This work was applied in assessing the Pegasus Workflow Management System (WMS; Vöckler et al., 2011) to process Kepler data in the search for Earth-like exoplanets. Pegasus was used to distribute the workflow across several cloud services at the same time, showing improved scalability and flexibility that would otherwise be impossible.

Deelman et al. (2008) also investigated the viability of cloud computing to support an astronomical application, in this case, the *Montage* mosaic engine[14]. This application used data from sky surveys such as SDSS and 2MASS to produce mosaic images. At the time, cloud services like AWS were still relatively new and the techniques to best exploit them for research workloads were still being discovered. Deelman et al. (2008) found that because Montage depended more on throughput of data, and had relatively short computation processes for the meshing of images, it was well suited to the cloud. They concluded that for data-intensive applications, cloud computing could be very cost-effective. They also identified areas where improvements were likely to be made as the cloud matured, such as reliability and complexity challenges, along with concerns around privacy and security of data.

Using the Pegasus WMS, Montage performance was tested by Hoffa et al. (2008), who compared a single machine, a local cluster, a VM and a virtual cluster. Though the underlying hardware was not completely consistent across the environments, with the local machine having the lowest specifications, the results were still important. They found that the single machine performance was decidedly better than the other environments, yet the capability of the single machine was inherently limited, as the memory requirements and parallel processing of multiple mosaics grew beyond the capacity of the machine. The results for the other environments indicated that the interconnect between nodes played a significant role, something not plaguing the single or VM environments. The authors

[14] `http://montage.ipac.caltech.edu/`

concluded that the cloud environment could provide a suitable platform for clustered computing. This suggests that the virtualisation impact that was potentially limiting for some HPC workloads, was worth testing explicitly for particular applications.

Several published papers have tested the viability of existing clouds in scientific workflows. Berriman et al. (2010) published a study comparing Amazon EC2 with a local compute cluster for generating periodograms for light curves from NASA's Kepler Mission. The authors concluded that the Amazon service was better suited to processor and memory bound applications than I/O bound applications. They followed up with another paper Berriman et al. (2013) looking more generally at cloud computing in different scientific workflows, including the astronomy domain. In this second paper, they concluded again that data-intensive research, particularly astronomy but also other disciplines, would likely find cloud computing offered value for money. The authors highlighted the fact that the Amazon's storage costs make longer term use very expensive. As such they also considered the potential of a funded research cloud, where costs were not borne by the research department, but rather by the research institution or government.

Interestingly, the three years between the two Berriman et al. papers mentioned above, saw significant development in the extent, capability and price of cloud computing. For example, Ostermann et al. (2010) also conducted tests and benchmarks on the Amazon EC2 service in 2010, and while they recognized the potential of cloud computing, they argued that an order of magnitude performance improvement would be required for cloud computing to play a useful role in scientific computing. For example, the authors discuss the possibility of HPC as a Service, but argue that many researchers were unlikely to want to undertake the necessary performance tuning of the interconnect bottleneck to make it worthwhile.

Meanwhile, Rehr et al. (2010) also compared the performance of a local cluster with a virtualized cluster running on Amazon EC2. Their experiments showed that the EC2 cluster performance closely matched their local cluster. This result also concurs with Cohen et al. (2013) who looked in more depth at the challenges involved in migrating scientific computing to cloud computing. Focusing on high-energy physics and bioinformatics, these authors declared that they "found no fundamental impediment" to deploying HPC in a cloud computing context. Their research was undertaken as part of a project called RAPPORT (Robust Application PORting for HPC in the cloud). This position was also supported by Hiden et al. (2013), who presented three case studies for e-Science Central[15], namely in spectral visualisation, medical data analysis and chemistry. They

[15] `http://esciencecentral.org/`

argued that cloud computing could revolutionize scientific computing provided skills were properly developed.

In astronomy, Armstrong et al. (2010), Berriman & Groom (2011) and Ball (2012a) all identified the value and potential of cloud computing. Armstrong et al. (2010) focused on a Cloud Scheduler and a VM management system deployed by CANFAR. This paper discussed the system architecture for high throughput computing, pointing out that this was particularly useful for particle physics applications, where some application code ran on VMs at nearly the same speed as bare metal. Ball (2012a,b) provides a useful overview of cloud computing and CANFAR.

At the Australian Astronomical Observatory, Green et al. (2016) asked the question, “What will the future of cloud-based astronomical data processing look like?”, highlighting the importance of cloud-based parallel computing in addressing these challenges. The *Data Central* service was intended to be a long-term archive with the aim of automating the parallel processing of Python code in Jupyter notebooks[16]. This approach alleviates the need to download data to a local computer for processing.

Berriman & Groom (2011) also discuss cloud computing for both CANFAR and the Square Kilometre Array (SKA) as a potential solution for the “data tsunami”. The authors highlight the challenges of shifting to a cloud model and point out that self-taught programming is largely responsible for poorly performing code, no matter what environment. They identify that programming education such as the Software Carpentry[17] movement will play a big role in preparing future researchers to develop quality code that will perform in a cloud computing environment.

Extending previous work, Berriman et al. (2013) set about identifying workflows that might be suitable candidates for deployment on a cloud. The authors compare the aforementioned Montage, *Broadband*, an application for generating earthquake scenarios, and *Epigenome*, which maps DNA segments to a reference genome. Montage is I/O bound, Broadband is memory bound, and Epigenome is CPU bound. The comparison was run on a local compute cluster, on AWS, on FutureGrid and on Magellan, with several VM variations, and using different parallel filesystems. The long-term storage costs of AWS were identified as being expensive but in general, the cloud platforms provided viable options for CPU and memory bound applications. As expected, there was a difference in performance between dedicated high performance cluster hardware and the generic commodity hardware typically deployed in clouds, but the significant cost savings, scalability and flexibility of the cloud services were a clear indication of the benefit of cloud services.

[16] `http://jupyter.org/`
[17] `https://software-carpentry.org/`

More recent information about Montage can be found in Berriman & Good (2017).

Astrocloud[18] was launched in 2014 as part of the Chinese Virtual Observatory[19]. Cui et al. (2017) describe the service that provides access for both professional and amateur astronomers to data from five observatories in China. The service hosts over 15,000 datasets and has more than 17,000 registered users. Users are able to exploit compute and co-located data using VMs with pre-installed software and direct access to the required datasets. To make the process even easier, remote desktop capabilities have been added. The service is distributed across six major cites in China, reducing latency and improving service redundancy. Ease of use and data connectivity were identified as key drivers for deploying the service on the cloud.

The virtualisation of computational resources provides a valuable mechanism for sharing scientific workflows. VMs can be copied and shared with collaborators, pre-installed with a full application pipeline. These images can be published for subsequent use or review by other astronomers. However, the recent development of container technology, where the

One of the most influential papers that has opened up adoption of cloud computing across the globe was Vivek Kundra's "Federal Cloud Computing Strategy" when he was the U.S. Chief Information Officer in 2011 (Kundra, 2011). This paper presents the "Cloud First Policy" (Kundra, 2011, p.2) in which US government departments are required to adopt a policy of considering cloud computing as fit-for-purpose ahead of the traditional bare-metal approach. This does not constrain government departments from deploying physical resources, however it does require them to provide evidence to support such a decision, including why a cloud computing solution is not suitable. This policy shift has had the flow on effect of boosting confidence in cloud computing across the commercial world, which has seen dramatic uptake in the last few years.

Cloud computing is a rapidly evolving landscape, and identifying the correct solution can be a challenge for any organisation. Turilli et al. (2013) discuss the use of federated clouds in which a mix of public and private cloud computing can enable researchers to use in-house resources for high performance or highly sensitive processing, but allow access to public clouds like AWS EC2. This approach can also help limit unnecessary data transfers, and also provide the much-needed flexibility a research institution can often require.

A technical description of how a cloud computing service is configured, operated and maintained is beyond the scope of this review, however three useful papers providing this background are Smith & Nair (2005), Voorsluys et al. (2011) and Butt et al. (2012). A

[18] http://astrocloud.china-vo.org/
[19] http://www.china-vo.org/

general taxonomy of cloud infrastructure and survey of commercial clouds can be found in Rimal et al. (2010).

Cloud computing addresses several challenges for the modern astronomer, particularly in the collaborative and ad-hoc demand areas of astronomy research. However, the practice of using remote computing resources, even in astronomy, is still met with reluctance. In general, astronomers, like most people, still prefer to have a computation device at their desk to perform most immediate tasks. This includes basic tasks like office applications, but also more demanding research tasks like visualisation or small-scale data processing. Here, small-scale means anything that the local computer is *capable* of handling, but what constitutes capable is open to interpretation. Having a local computer processing for hours or days might be acceptable for one researcher but not another. Purchasing an external hard drive and downloading a multi-terabyte data set might work in some circumstances but not others.

Many astronomers have readily accepted the value of remote computation and storage, and are familiar with access methods where using the remote service is handled via a command line interface, such as via the SSH protocol[20], or a graphical windows allowing an application to be streamed to the local computer via SSH, as in the case of X11 Forwarding[21]. However, many applications don't have a command line interface, and the performance of graphical window-based applications can often be such that it is only used as a last resort.

A more advanced approach is the use of a remote desktop, where rather than individual application windows being streamed to the local computer, a suite of applications is presented as a streamed desktop in real-time, closely matching the look-and-feel of a local desktop interface. Several technologies exist to support this method of interacting with a remote compute and data service, and when correctly deployed and used over suitable network infrastructure, can even match the user experience of a local desktop. Using a remote desktop in the cloud and the challenges and solutions this implies, is the subject of the next section of the review.

## 2.3 Visualisation in the cloud

Visualisation is a key part of astronomy research, and there are many challenges to successfully using visualisation to improve research outcomes, especially in the big data era. Hassan & Fluke (2011) highlight the importance of visualisation-based knowledge dis-

[20]https://www.ssh.com/ssh/protocol/
[21]http://tldp.org/HOWTO/XDMCP-HOWTO/ssh.html

covery as "critical" for the foreseeable future of astronomy research. Of the six "Grand challenges" described in the paper, the two of most importance to this work are "*effective handling of large datasets*" and "*better workflow integration*". These two areas of interest can be addressed in combination by considering the use of cloud-based visualisation, where the workflow is presented on a remote desktop that is co-located with the (large) datasets for processing. Further extending this capability with GPU-enabled VHDs allows this approach to also begin to address the sixth "Grand Challenge" identified, "*encouragement for adoption of 3D scientific visualization techniques*".

Remote computing via command line interfaces has been used in astronomy for some time, and the forwarding of an X11 window for applications with a graphical user interface is a common method of accessing remote Unix and Linux computing resources. Microsoft Windows and Apple Mac OS have equivalent techniques. Forwarding a whole desktop interface became possible with the advent of suitable protocols, such as (VNC; Duato et al., 1997). This requires a VNC server to be running on the shared resource, which is then connected to via a client computer. This might be a standard computer with a client like TurboVNC[22] installed, or a thin-client (Nieh et al., 2000) that is designed to provide minimal local resources, and assumes that all computation will be streamed from the shared resource. Thin-clients are essentially a replacement for dumb terminals, where the device itself is not intended to provide a computation environment, but rather purely a means to connect to the remote computation environment. Keyboard and mouse controls for interactivity, and suitable graphics capabilities to display incoming content are all that is required.

Desktop virtualisation has been explored for over two decades, driven by the need to mitigate the cost and management of the burgeoning desktop computer fleets deployed in research workplaces. Understanding where this technology is best suited, and might replace the more commonly used local desktop, is challenging due to the rapidly changing landscape. The value proposition of cloud computing comes from balancing the use of lower-cost commodity hardware with the scale of deployment necessary to accommodate the research demand. The value of flexibility is lost if the resource is not used efficiently, noting that it is often difficult for researchers to invest the time needed to ensure this efficiency. In the case of virtual desktops provisioned in the cloud, often called Desktop-as-a-Service, Kishan et al. (2014) present a method to optimise resource allocation to improve efficiency, thereby improving the viability of providing cloud-based desktops.

Khalid et al. (2016) also provide a useful overview of the virtual desktop, detailing the

[22] `https://www.turbovnc.org/`

evolution of the technology and discussing a variety of models. Besides the cost advantages, the authors conclude that there are several additional benefits to virtual desktops, including improved security, reduced management and maintenance load, as well as increased flexibility. However, they also conclude that where performance is important, particularly with graphically demanding applications, virtual desktops are unlikely to be a good option. This perspective is not necessarily shared by everyone however.

Ravi et al. (2011) described the use of VMs that could access a GPU located on a host computer. Unlike the virtualisation of RAM and CPU, the GPU is not provisioned by software, but rather is made available in a "pass-through" mode, with access managed by the VM management software. This method allows the VM to route GPU calls, such as OpenGL or CUDA, through to a physical GPU for processing. In this way, GPU-based applications can function on the cloud with the same level of performance (albeit with the minor impact of the CPU virtualisation) as a physical computer. Each GPU would be occupied by a single VM, and as many modern GPU cards have several processors on them, multiple VMs could be provisioned with their own, dedicated GPU, for rendering or processing. This capability is available from several cloud providers, including AWS.

It is now possible to achieve GPU virtualisation in clouds, e.g. NVIDIA's vGPU[23]. Iserte et al. (2016) describes the process of slicing a single GPU into vGPUs that are assigned to VMs on the host. Rather than a single GPU per VM, the GPU can support multiple VMs, depending on the load from the user of the VM. For example, for light loads, such as basic windowing graphics, the GPU might provide 10 or more vGPUs. But for more demanding workloads, only one or two vGPUs might be made available. Because the hardware is virtualised, the number of vGPUs can be easily adjusted.

Taking this a step further, Hong et al. (2017) describes sharing the GPUs over the network. With suitable network bandwidth, it is possible to make a GPU resource from one host available on to a VM located on another host. This opens up many opportunities for dynamically provisioning and scaling GPU resources. AWS recently announced a new product line called *Elastic GPU*[24], that allows for GPUs to be attached to VMs as required, providing even more flexibility. With this model, a VM might be used for a variety of tasks that do not require a GPU, and therefore can operate on the cheaper CPU-optimised hardware. However, when GPU capability is required, a GPU can be attached and used with the additional costs incurred while the device is connected.

High-end graphics capabilities are important for understanding research data, yet the purchase of these specialist workstations is expensive, especially when demand is high.

[23] `https://www.nvidia.com/en-us/design-visualization/technologies/virtual-gpu/`
[24] `https://aws.amazon.com/ec2/elastic-gpus/`

This is exacerbated by low utilisation of resources, due to lack of portability. Miller & Pegah (2007) observed significant under-utilisation of high performance graphics workstations and investigated ways to improve this. Direct sharing of a workstation is problematic when more than one person needs to use the machine at a given time. Fortunately it is possible to provide coincident sharing using virtualisation. Aside from improved access, there are additional benefits including reduced deployment, power consumption, maintenance and enhanced security, especially when the resource is deployed in a data centre. In this model, the virtualised computer is presented to the remote user as a desktop, with all the same characteristics and capabilities of a local desktop computer.

An example of the use of cloud-based desktops in a scientific training exercise can be found in Berriman et al. (2012). Rather than require 160 astronomers install a suite of applications, a VM image was prepared on AWS with the applications pre-installed[25]. The astronomers were then able to launch their own VM instance based on that image and the applications and test data could be used immediately.

High performance and high throughput computing services usually require users to go through training to properly use the service. Transferring a workflow from a local computer to a HPC cluster is non-trivial, and thus represent a significant hurdle for some researchers. So much so, that some service providers opt to present their resources through user-friendly interfaces like the web or as remote desktops. For example, the Cherenkov Telescope Array (CTA)[26] project provides a *Science Gateway* for researchers to access the telescope data. Massimino et al. (2014) describe the Astronomical and physics Cloud Interactive Desktop (ACID)[27], which provides a VNC-User Interface to a remote desktop provisioned in their cloud environment. This offers a consistent user experience with applications and data automatically available to researchers.

Similarly, the Nectar Research Cloud supported several Virtual Laboratories, such as the Characterisation Virtual Laboratory (CVL)[28], which provides access to the Multimodal Australian ScienceS Image and Visualisation Environment (MASSIVE)[29]. Goscinski et al. (2015) discusses the access methods for MASSIVE, including the MASSIVE Desktop[30]. This method of access allows researchers to use the power of the MASSIVE compute resources through a remote desktop that has GPU pass-through capability. Users can request an interactive desktop session with particular features including the amount

[25] `http://nexsci.caltech.edu/workshop/2012/`
[26] `https://www.cta-observatory.org/`
[27] `http://acid.oact.inaf.it/ACID/Home_page.html`
[28] `https://www.massive.org.au/cvl`
[29] `https://www.massive.org.au/`
[30] `https://www.massive.org.au/userguide/getting-started/the-massive-desktop`

of RAM and number of CPUs, along with an estimate of how long they need the remote desktop for (called the wall-time in HPC parlance). These virtual desktops are automatically provisioned with GPU resources and provide the researchers with a desktop experience that parallels and often exceeds their local desktop. The underlying technology developed for this project, Scientific Remote Desktop Launcher (Strudel)[31] is in use at Australian national peak computing facilities, including the Pawsey Supercomputing Centre[32] in Western Australia, and the (National Computational Infrastructure (NCI); Druken et al., 2016) in Canberra.

The ability to move beyond the local desktop computer to a cloud-based desktop allows astronomers to take advantage of the flexibility and cost-effectiveness of cloud services. It also mitigates the need to transfer datasets that greatly exceed the capacity of their local computer, as the computational resources are typically located in the same data centre as the data itself. However, this does not solve the problem of visual inspection of images that have resolutions many orders of magnitude larger than the display screen attached to the local computer.

The data challenges posed by modern astronomical instruments have easily outstripped the resolution increases in physical computer displays. However, in much the same way that parallel computing allowed the rapid expansion of computation workrate using commodity computers, parallel visualisation allows the integration of discrete standard computer displays to function as a single, unified display space, called a Tiled Display Wall (TDW). The evolution of the TDW and related technologies is the subject of the final section of the review.

## 2.4 Tiled Display Walls

A key element to the success and value of a TDW is the way in which a user engages with the system. User behavior with TDWs has been explored over the past decade and generally has supported the expectation that within a limited scale, expanding the display environment for highly detailed content produces better engagement for users. The literature also establishes the concepts of physical and virtual navigation. Physical navigation is where the user employs their own body's physical capabilities to explore displayed content. This ranges from simple eye movement, to moving the head, to ambulation. Virtual navigation on the other hand requires the user to employ interface technology to achieve the same outcomes. The most common of these is the use of a computer mouse and/or

[31]`https://www.massive.org.au/userguide/cluster-instructions/strudel`
[32]`https://www.pawsey.org.au/`

keyboard to pan and zoom the displayed content.

TDWs have been around for several years, and have had some success in supporting research, education and engagement activities in a variety of disciplines. Several examples are discussed below.

### 2.4.1 Hardware

The idea that computers might be able to present data as an immersive and sensory experience was first put forward by Sutherland (1965) when he envisioned the "Ultimate Display". He extrapolated based on the emerging capabilities of displays to a future where computers could generate sensory stimuli from data such that ordinary objects could be replicated in both form and function. Sutherland believed such a display would allow researchers to step beyond our ordinary world and experience the extraordinary worlds such as the infinitesimal world of particle physics or the vastness of the cosmos.

This vision was captured in popular entertainment with the TV series "Star Trek: The Next Generation", where the "Holodeck" made its first appearance[33]. While the conceptual explanations from the show's writers were somewhat far-fetched, scientists began investigating how such a display might work. Sutherland had already been working on a display worn on the head, called a head-mounted display (Sutherland, 1968). The movement of the head could be tracked and the display image updated accordingly in real-time, giving the wearer the illusion of looking around a virtual environment. The early developments of the head-mounted displays can be found in Chung et al. (1989) and a summary of more recent virtual reality displays, including the Oculus Rift, in Boas (2013).

The second development path was the creation of the CAVE, which stands for Audio-Visual Experience Automatic Virtual Environment, a recursive acronym (Cruz-Neira et al., 1992), where the "C" stands for CAVE. This involved projecting the computer-generated environment onto each of the six walls surrounding the viewer. In this way, all visual information is supplied by the computer and coupled with a spatial audio system and head-tracking, thus providing a compelling sense of immersion. The authors describe maximum immersion as requiring a panoramic view to surround the user, head-tracking, and limiting the user's view of unrelated content, as well as the ability to move around in the environment. In essence, the CAVE offers the ability to "zoom" by walking closer or further from the display surface, and "panning", by turning the head and eyes.

However, due to the cost and spatial requirements, few full CAVEs have been built.

[33]Episode #1, "Encounter at Farpoint", 1987

Cheaper variations used fewer walls, such as the WEDGE displays, where only two or three walls were used (Boswell & Gardner, 2001; Gardner et al., 1999; Large et al., 2010). Alternative displays included large curved walls with multiple projectors combined to provide a seamless image using edge blending (Van Baar et al., 2003; Brown et al., 2005; van der Schaaf et al., 2007; Song et al., 2010). Edge blending compensates for the off-axis luminosity fall-off of the projection. Typically these displays also provide active stereo 3D, which required the projectors to have high enough frame rates with very little image cross-talk[34] to provide a useful stereo pair presentation (Konrad et al., 2000). These projectors were typically very expensive and despite providing a convincing sense of immersion, the expense associated with these installations meant very few were built.

The value of combining multiple displays to improve user performance had been identified in the mid-eighties (Woods, 1984). Increased screen real-estate allowed for considerably more data to be presented to the user however, the scalability of this approach on a single computer was limited to the rate of development of graphics cards. While the gaming industry has pushed the GPU development at an exponential rate, a more readily scalable alternative took advantage of the even faster expansion of network bandwidth. This resulted not only in the emergence of the cluster-based supercomputer (Hoffman & Traub, 1989), but also in the visualisation cluster (Wierse et al., 1993; Heirich & Moll, 1999).

Distributing imagery across several synchronized displays provided a cost-effective solution to increasing pixel count well beyond that of even the highest resolution monitors. Agana et al. (2010) compared TDWs with high-end immersive displays and found the user performance to be equivalent,

All TDWs use a similar principle for the coordination of the displays. Typically a dedicated head node manages the tile nodes supporting the screens that make up the display surface. Users may interact directly with the head node or with another interface device, but not directly with the tile machines. The tile machines may either display a pixel stream rendered on the head node or a remote cluster, or they may be directed by the head node to render content directly to their attached screens. The first two options require considerable network bandwidth as a stream must be generated for each tile node. The second option requires less bandwidth as only the rendering commands from the head node are sent, however the data to be rendered must be pre-staged on each machine for this to be useful.

Methodologies for building a TDW can be found in Hereld et al. (2000); Li et al.

[34]Cross-talk is the phenomenon where one side of a stereo image pair is seen by the wrong eye.

(2000); Kang (2007); van der Schaaf et al. (2007); Navrátil et al. (2009) and Scheidegger et al. (2012). Commercial options include SharpWall (Deshpande et al., 2009; Deshpande & Daly, 2010) and Samsung's MagicWall[35].

The OptIPuter project began in 2003 as a way of using optical data connections to connect research facilities (Smarr et al., 2003, 2005; Taesombut et al., 2006) and share large amounts of data over high-speed networks. A natural progression of sharing such data was to find a way of collaboratively visualizing it. Hence the OptIPortal[36] was born (DeFanti et al., 2009a; Almes et al., 2011). An OptIPortal is essentially a TDW with the additional capability of being connected to other TDWs via high-speed networks. Technically this feature is not unique to the OptIPortal, but is a defining aspect to this project, and lead to the idea of the OptIPlanet, a world-wide collaboration between capable research institutions (Smarr et al., 2009).

Around this time, another project called "WeSpace" was being developed at Harvard University, combining multiple high-resolution projectors with a multi-touch surface display(Wigdor et al., 2009). This is an example of an early display ecology, with the service coordinating a variety of display and interaction technologies to facilitate improved sense-making and greater collaboration. This project was developed and tested with astronomers from the university.

After an earlier divergence, TDWs and CAVEs began to converge again (DeFanti et al., 2010). Key drivers for this path was to reduce the physical footprint required by projection systems and to increase the resolution of the displays. One approach was the StarCAVE (DeFanti et al., 2009b), where all the projection screens were orientated toward the primary viewer. A subsequent version called NexCAVE used 9 Stereo capable LED displays to provide a forward-only immersive view, which also allowed for control from an iPhone or tablet device (Wedeen et al., 2014).

Early implementations of TDWs didn't use stereo-capable screens due to the expense of the devices and the issue of bezels disrupting the 3d effect (Grüninger & Krüger, 2013). However, as commodity prices for 3D LED TVs dropped dramatically and passive stereo became available, more TDWs began to employ these displays. The stereo effect could be maintained by keeping 3D objects behind the focal plane, where the bezels would not disrupt the view. This "French Window" effect was further enhanced by having the TDW management software employ a virtual occlusion of pixels as they "pass behind" the bezels, has been employed for 2D TDWs for some time (de Almeida et al., 2012).

[35] `http://www.samsung.com/au/business/solutions-services/smart-signage-solutions/smart-signage-solutions/magicinfo-videowall`

[36] `https://www.evl.uic.edu/entry.php?id=1547`

The convergence of CAVEs and TDWs can be seen in the CAVE2, which combines the best of both worlds (Febretti et al., 2013; Leigh et al., 2013). This immersive display environment is a 320-degree cylindrical wall standing 8-feet high. It is made up of 72 ultra-thin bezel stereo-capable displays. With combined bezel thickness of less than 8mm, and an optimal viewing distance of several feet, these bezels do not disrupt the stereo effect of objects presented in front of the focal plane, i.e. coming out of the screen toward the viewer.

Holding the record for the largest number of pixels in a display wall, the Reality Deck[37] is an immersive Gigapixel[38] display, with over 1.5 billion pixels. At 33 feet $\times$ 19 feet, the facility can host several people at once, and provides 20/20 visual acuity in a 360°visual environment. At this scale, the Reality Deck provides not only the display surface, but a fully immersive room, where the walls are part of the display surface (Papadopoulos et al., 2015).

Other challenges facing TDWs and their adoption into scientific workflows are steadily being addressed. Understanding the interaction with a TDW is key to their future development (Moreland, 2012; Rivera et al., 2013). Head tracking was investigated as early as 2007 (Wong et al., 2007) and more recently work has been done on including haptics (Lee et al., 2011), multi-touch control (Nishimura et al., 2012) and remote gaze tracking (Lee et al., 2013). A more extensive description of the underlying technology of TDWs can be found in Hagen (2011).

### 2.4.2 Software

Cluster management is often ad hoc for small clusters, while larger systems typical use a solution like ROCKS (Sacerdoti et al., 2004) to simplify the management of the machines in the cluster. Whatever solution is adopted, the underlying principle is the same: a head node coordinates the tile nodes by automatically launching software on the nodes and accessing data on a shared drive. Other considerations address the typical deployment and operation of a cluster.

In order to coordinate the display of data across the tiled screens, a variety of software has emerged. Li et al. (1997) described ParVox, designed to render volumes in parallel for distributed visualisation. Another project emerged called Xdmx (X Distributed MultiHead X)[39] (Faith & Martin, 2001), which allowed the presentation of X windows to be distributed across multiple screens. For basic content like text and static images,

[37] `http://labs.cs.sunysb.edu/labs/vislab/reality-deck-home/`
[38] A Gigapixel image contains over one billion picture elements, or pixels.
[39] `http://dmx.sourceforge.net/`

and for clusters of only a few machines, this process worked reasonably well. However, as the number of machines increased, the performance dropped off rapidly. Another impact to performance was the rendering of OpenGL content, however, the combination of Xdmx and Chromium[40], software designed to utilize local graphics cards and optimize the output streams, provided a way of extending the capability and size of a viable visualisation cluster (Humphreys et al., 2002). However, the scalability and management of the Xdmx/Chromium solution was limited and many in-house solutions evolved.

Another less well-known project called Vistrails[41] was also developed to assist in improving the graphics-processing pipeline (Callahan et al., 2006). While not the focus of the project, one of the components of Vistrails allowed the graphics output to be sent to a TDW. However, the more complex configuration required for Vistrails has meant there is relatively little adoption in Australia.

The COllaborative Visual Simulation Environment (COVISE[42]: Rantzau et al., 1996) is a scientific visualisation tool with a 3D Renderer called COVER (COVISE VR). This was further developed into OpenCOVER[43], which was able to be displayed on a TDW (Kopecki, 2011). It was also able to be used for video conferencing (Chu et al., 2008).

In 2009, the Texas Advanced Computing Center (TACC)[44] produced a white paper (Navrátil et al., 2009) detailing the development and motivation behind building the world's largest TDW of the time, Stallion[45], which is comprised of 328 megapixels across 80 LCD displays. Released in 2012, DisplayCluster[46] was developed to drive the display wall. According to Johnson et al. (2012), the DisplayCluster project aimed to address shortfalls in the other major display wall management suites of the time. New features not then available in other suites include touchless interaction using the Microsoft Kinect[47], a depth-sensing camera, and interaction control via mobile smart devices. Many of these features have since been incorporated into competing products.

Two projects of note that emerged around 2010 were Clustered GLX, or *CGLX* (Doerr & Kuester, 2011) from CALIT2 at UCSD; and *Scalable Adaptive Graphics Environment (SAGE)* (Renambot et al., 2004) from the Electronic Visualisation Laboratory (EVL) at the University of Chicago. In fact, early construction specifications for an OptIPortal

[40] `http://chromium.sourceforge.net/`
[41] `https://www.vistrails.org/index.php/Main_Page`
[42] `https://www.hlrs.de/en/covise/`
[43] `https://www.hlrs.de/en/solutions-services/service-portfolio/visualization/covise/opencover/`
[44] `https://www.tacc.utexas.edu/home`
[45] `https://www.tacc.utexas.edu/vislab/stallion`
[46] `https://www.tacc.utexas.edu/research-development/tacc-software/displaycluster`
[47] `https://developer.microsoft.com/en-us/windows/kinect`

required the use of CGLX on a ROCKS cluster (Sacerdoti et al., 2004).

CGLX allowed the management of the display nodes through an interface on the head node. It allowed the creation of "worlds", which defined screen groupings for the purpose of running display modules. In most cases, the "world" was defined as the entire available display space. Modules including "Imageblaster" allowed the display of multiple JPEG and PNG files as well a video stream receiver. Displayed elements could be positioned anywhere and at any scale within the "world" as defined by the head node. The video stream receiver could display a real-time multicast video feed, and was often used to display videoconference participants.

Another module was "Videoblaster" which enabled the display of Quicktime, H.264 and MPEG movies. Again, several movies could be displayed simultaneously and scaled as required (Ponto et al., 2009). Ponto et al. (2011) also developed CGLXTouch, which allowed the connection of Touch screens to a CGLX-enabled TDW.

The "Tiffblaster" module could display a particular type of TIFF format called "Pyramidal TIFF" (Pitzalis et al., 2006). Images in this format are created from stacks of images, making use of TIFF pages feature. A tiled version of the full resolution image represents the bottom of the stack, and a smaller version sits above it, followed by a smaller version again, and so on. The value of such a format is that software like "Tiffblaster" need only load into memory the tiles from the layer that is most suitable for the size being requested for display. This provides a convenient way of quickly viewing images of tens or hundreds of megapixels, or even Gigapixels, as the image can be loaded and unloaded progressively. When using stored imagery that exceeds the scale of the TDW, additional methods for staging and enabling a high level of interactivity can also be required. Ponto et al. (2010) developed Giga-stack, a method specifically designed to address this challenge.

Several aspects of CGLX made it less than ideal for the operation of a TDW. Firstly, the creation of "worlds" in which only one module could be loaded, imposed a significant restriction on how the TDW was used, as content requiring different modules could not be displayed in the same "world", and if the content needed to be moved beyond the confines of the preconfigured display space of the module, the entire TDW had to be reconfigured. Also, the preparation of Gigapixel images into a pyramidal TIFF required significant processing, and was not practical to do in an ad hoc manner.

Furthermore, sharing content between remote sites often required staging of multi-gigabyte-scale content ahead of time, as real-time display connectivity was only viable for images in the tens of megabytes or smaller range. The local connectivity within the TDW cluster also required preloading the content on each of the display nodes for multi-gigabyte imagery, especially video or pyramidal TIFFs, as image tearing became obvious between node displays.

The virtues of CGLX were eventually outweighed by the challenges and most TDWs transitioned to SAGE by 2012, though some sites did continue to develop their own in-house solutions. One such solution was developed by Vohl et al. (2016) to display spectral volume cubes in the CAVE2 facility located at Monash Immersive Visualisation Platform (MIVP)[48].

Another astronomy-specific endeavour is described in Pietriga et al. (2016). Using a Java-based environment called the Zoomable Visual Transformation Machine (ZVTM; Pietriga et al., 2011) for clusters, the FITS-OW (*FITS on Wall*) application allowed astronomers to use a TDW to display astronomy applications like SAOImage DS9[49] and Gigapixel FITS images. The authors concluded that the use of TDWs could not only facilitate serendipitous discovery, but also enhance collaboration in the classification of astronomical objects, by virtue of the presentation of multiple sources of information.

The key benefit of SAGE was the unified display space. The entire TDW was available to display any content, and each element could be scaled, stacked and repositioned at will. This allowed for far simpler interaction control than with CGLX. However, because all content was stored and rendered at the head node, a high-speed local network was required to ensure content distribution was smooth. At 1Gbps, the University of Melbourne OptIPortal functioned well for static images and individual movies up to FullHD resolution, but multiple movies resulted in flooding the local network and reducing frame rates. No option existed to stage content locally.

Despite this, SAGE was readily adopted by most TDWs in Australia and many international sites. It supported more image formats than CGLX, though not the pyramidal TIFFs, which was perhaps the most valuable module of the CGLX suite. Besides the ease of use, SAGE also enabled connectivity between display walls. This was handled in two ways. The first was via the controller interface. The SAGE User Interface (SAGEUI) could connect to multiple TDWs at the same time, and libraries at each of the nodes were accessible by all connected SAGEUIs, of which there could be several. This allowed a researcher to display content on a local TDW, and push the same content to a remote TDW. As the SAGUI provided a preview of the remote TDW showing its layout, the local user could place their content wherever it was required on the remote TDW. This did not prevent the remote user from also moving or removing that content, or adding content of their own.

[48] `http://www.monash.edu/mivp`
[49] `http://ds9.si.edu/site/Home.html`

The SAGEUI is available as a standalone application that can easily be run cross-platform. This allows local collaborators sharing a TDW to each contribute content and arrange it as required. Each connection is displayed as a uniquely coloured mouse pointer accompanied by a username. In this way, several researchers could share content no matter where they were located.

Another benefit of SAGE is its extensibility. Several additional modules have been developed to extend the functionality, including an OpenGL wrapper module. This module allows any OpenGL application to have its OpenGL graphics rendering calls to be intercepted by the SAGE OpenGL wrapper (SAIL) and the contents duplicated and streamed to the TDW. Applications such as *Nasa Worldwind*[50] could be displayed on the TDW without modification. As the SAIL intercepts the OpenGL calls, the resolution being streamed to the TDW is exactly the same as that being streamed to the display. This dependence on screen resolution meant that while an OpenGL application could be scaled up visually on the TDW, the resolution was that of the smaller screen of the computer on which it was being displayed. The dependence of OpenGL and the graphics acceleration card was defined by the interaction between the screen and the card, however, using a virtual frame buffer allowed the graphics card to effectively be fooled into responding as though connected to a much large, higher resolution display. The pixels would then be rendered for the virtual display, up to the full capability of the graphics card, and streamed at high-resolution to the TDW.

Applications being presented on TDWs can also be remotely controlled through several mechanisms, but a specific solution was presented by Fujiwara et al. (2011) as an extension to SAGE. SAGE can also be combined with rendering pipelines to take advantage of clustered rendering, or renderfarms. Nam et al. (2009) describe a solution using Paraview on a cluster back-end to drive a visualisation on a SAGE display.

Modules such as web browsing and PDF viewing have been incorporated into the software. Kim et al. (2009) developed iTILE, a TDW management tool for displaying desktop applications on a TDW. Tada et al. (2011) developed a SAGE specific module to allow any XWindow application to be displayed in the SAGE display. See also Olsen et al. (2011). COVISE can also be integrated with SAGE (Shin et al., 2010).

Some software focused on specific user interactivity challenges. For example, when several sites are connected and wish to view a video stream simultaneously, network bridging and multicasting is required. The SAGE-Bridge module provided this functionality to SAGE TDWs (Renambot et al., 2009).

[50] `https://worldwind.arc.nasa.gov/`

In 2013, Vadiza[51] partnered with the University of Chicago to offer a commercialized version of SAGE. This partnership, with input from several developers, has led to a new release of the software called SAGE2 (Marrinan et al., 2014), with a change to the acronym - it is now called Scalable Amplified Group Environment to reflect the changing focus. The project has significant funding from the National Science Foundation[52]. Due to the shift in architectural design of SAGE2 to a web-based model, streaming data congestion is potentially a problem when several high definition streams are sent to a display wall simultaneously. (Kido et al., 2016) proposed a dynamic network routing, using Software Defined Networking (SDN)[53] to mitigate the bottleneck. This method improves the visualisation performance by making better use of available networks.

### 2.4.3 User Behaviour

Understanding how people interact with and use a TDW is key to ensuring they are developed and deployed appropriately, to ensure the most cost-effective outcomes and to enhance research activities. Not only is the display technology itself important, but so to is the space in which it is used, and the way people choose to interact with it. Peck et al. (2009) look at how people interact with displays of varying sizes where physical navigation is required, and discusses the importance of designing a TDW with this link in mind. They describe a link between perception and interaction in terms of the scale, and conclude that it is worthwhile designing installations to exploit this connection.

Ball & North (2005a,b) conducted some of the first experiments to investigate how people used a TDW. While their TDW consisted of only nine screens in a 3 × 3 configuration, they were able to confirm some basic advantages of a TDW - for example when all content needs to be seen at the same time, it is easier to match structural pairs, than when some of the matching components are off-screen. In their experiment, they used random grey dots on a black background with small red dot structures that needed to be paired.

Ball et al. (2007) also identified the importance of physical versus virtual navigation in a paper called "Move to Improve". This paper highlighted the fact that physical motion was an advantage when remembering which part of the image had already been inspected was important. Maintaining a sense of overall context was easier than with virtual navigation.

The use of multiple displays on a single computer has been available for some time, and studies have shown that not only do users prefer the additional screen real-estate,

[51] `http://vadiza.com/sage.php`
[52] `http://lava.manoa.hawaii.edu/sage2-the-scalable-adaptative-graphics-environment/`
[53] `https://www.opennetworking.org/sdn-definition/`

their actual performance also improves (Bi & Balakrishnan, 2009; Bi et al., 2014). It also improves their engagement, as evidenced by use in computer gaming (Lin et al., 2006).

Yost et al. (2007) introduces the concept of "visual acuity" in the context of TDWs, that is, the point at which pixel density and human perception cross. Modern displays can have pixel densities that exceed visual perception, such that the pixels themselves are not visible to the viewer, and the content appears seamless. Yost et al. (2007) determined that a display environment that met or exceeded visual acuity was able to produce increased efficiency and accuracy in the test subjects. Fluke & Barnes (2016) take this idea further, describing the conditions need for achieving the "Ultimate Display", where the user is no longer able to perceive the technology presenting the visual information, only the data itself.

However typical computer displays do not match the natural human method of sorting and consolidating information. Andrews et al. (2010) discusses how a TDW can be useful in digesting voluminous and/or heterogenous data. Much like a physical desk or whiteboard provides a large, flexible workspace, TDWs do not require content under consideration to be stacked as with windows on a standard desktop display.

Taking this further, Andrews et al. (2011) considered how TDWs allow researchers to work at "human scale", where the physical interaction with the displayed content required large, natural movements. This sort of interaction changes user behavior and enhances engagement and understanding.

TDWs are still unfamiliar to many researchers, and they are unlikely to be used properly without some guidance. Contrary to earlier findings suggesting physical navigation was preferable to virtual navigation (such as in Ball et al. (2007) and Liu et al. (2014)), Jakobsen & Hornbæk (2015) found that physical navigation does not necessarily improve performance when users are able to choose virtual navigation, especially when it is not necessary for the task at hand. The authors consider several reasons for the discrepancies, including the design of the experiment, the features of the displays and the rooms used, and the tasks required of the participants.

Furthermore, Liu et al. (2014) found that collaborative performance was not always improved with the use of a TDW when compared with participants using independent standard displays. However, they did find that as the complexity of tasks increased, the strategies employed by participants while using the TDW that had previously reduced performance, such as increased communication, resulted in improved performance over the standard displays.

This suggests that the design of the experiments to evaluate the usefulness of advanced displays needs to carefully consider the interaction techniques, directions and contexts that will influence participants in such a study. This is critical to understanding how these sorts of technologies will be used and valued in a functional research context.

Understanding the way humans interact with and absorb information from a TDW is key to using them effectively. Moreland (2012) discusses these and related challenges, effectively highlighting the fact that while the technology has advanced, insufficient research has been conducted to properly ascertain the research value of these devices, and thereby provide insight into their construction. Furthermore, evaluating applications, environments or users independently carries a risk of overlooking the complex interplay between these factors. According to Lam et al. (2012), many studies focus solely on the methods used to assess a particular feature of a problem, rather than taking a more holistic view of the scenario. The authors provide a framework of evaluation scenarios to aid researchers in developing suitable approaches to conduct their own investigations. This framework is discussed further in Chapter 3, where a clear alignment with parts of the framework and the following chapters of published papers is shown.

### 2.4.4 Examples

Perhaps the biggest challenge facing the adoption of TDWs is the number of suitable applications. While many disciplines produced large datasets that translated into images that significantly exceeded the display resolution, or where the numbers of images to be simultaneously displayed is necessary for comparison, integration into the researcher's workflow has not been smooth or simple enough for widespread adoption.

The literature supporting the use of TDWs in scientific endeavor is relatively thin, however a number of papers are available that highlight the successful application of TDWs.

One of the first discussions of TDWs being used in earnest can be found in Taesombut et al. (2006). This paper highlighted several challenges facing the Earth Sciences community, including the need for collaborative visualisation and communication in real time. The volume and scale of the image sets being captured required a display environment beyond what was typically available. The OptiPuter, the precursor to the Optiportal, was identified as providing the best solution to these problems. Other examples of TDWs being used in this field include Huffman et al. (2009); Tong & Zhao (2010) and Hsieh et al. (2011).

Many fields face the challenge of establishing effective collaboration environments, both for remote and local participants. The use of TDWs in a medical context is described in

Olsen et al. (2008) to address the needs of a radiology department communicating results with other hospital departments. Lau et al. (2010) employed TDWs to investigate the chemical structure of drugs and Son et al. (2010) did so for visualizing multiple high-resolution brain images.

TDWs have also been deployed to assist emergency services for monitoring and to aid in planning for emergency response. Sakuraba et al. (2013) describe the challenge of presenting disaster-related information rapidly and the importance of displaying content from multiple sources simultaneously. Hsieh et al. (2013) used a TDW to visualise tsunami simulations created with high-performance GPGPU[54] processing, showing water incursion impact enable better planning for future disasters.

The emerging field of Immersive Analytics, which brings together several visualisation and interaction technologies, is discussed in Sommer et al. (2017). MIVP operates the CAVE2 facility and the sensiLab[55], a resource to allow researchers to design and build custom devices for interactivity. The combination of technologies has facilitated several technically challenging research projects, including combining head mounted displays with the TDW. The use of Augmented Reality (AR) in combination with other display technologies is a nascent area of research (Nagao et al., 2016) that is likely to see rapid growth due to the interest in the general public around technologies like Microsoft HoloLens[56], Intel's Vaunt[57] and Magic Leap[58]. The combination of augmented and virtual reality environments with TDWs and other display technologies is called Hybrid Reality Environments, the subject of Febretti (2017).

The use of TDWs to support astronomy research is rare. The potential for this technology for astronomy was identified in 2006 by Fluke et al. (2006) and addresses some of the challenges presented in Goodman et al. (2011). Sims et al. (2010) discuss the use of TDWs to support the Mars Rover missions for NASA and Morikawa et al. (2010) mentions the use of TDWs to support the OneSpaceNet science cloud project at Japan's National Institute of Information and Communications Technology (NICT) studying Solar-Terrestrial Physics.

[54] General Purpose GPU
[55] https://sensilab.monash.edu/
[56] https://www.microsoft.com/en-au/hololens
[57] https://www.techradar.com/news/intel-vaunt
[58] https://www.magicleap.com/

## 2.5 Summary

Moving beyond the desktop is a challenge for the astronomy community, yet there is a clear need to do so. The big data challenge is far from addressed, however, rapid progress is being made in both technology and methods to ensure the value of capturing and/or producing so much data.

Operating at the petascale requires computation and storage resources to match, as well as the enabling technologies such as networking, security and infrastructure management. Solutions of this scale are expensive and few research institutions are so well funded as to be able to ignore the financial impact of conducting research in the petascale era. Cloud computing provides a mechanism to mitigate this cost, by ensuring resources can be easily redeployed to accommodate the changing demands of the research community. Cloud service providers such as Amazon Web Services, Google Cloud and Microsoft Azure, who primarily market to commercial organisations, operate at such scale that their purchasing power allows them to buy resources well beyond that of a typical university, and hence to accommodate unpredictable fluctuations in utilisation. However, commercial drivers are not always compatible with research drivers. This is where federated research cloud services like the Nectar Research Cloud is so valuable to the research communities they serve.

There is an opportunity to further exploit cloud computing. Shifting to a VHD is fast becoming a viable option to accommodate almost all of the computation requirements of the modern astronomer. With advances in streaming and GPU technology, along with network improvements within universities and to the homes and portable devices of researchers, it is likely that soon, very few will require a high-end graphics workstation to be deployed at their desk.

Being able to work with Gigapixel imagery requires a suitable display, and for most astronomers, the display attached to their local desktop or laptop computer is inadequate. However, the use of several commodity displays in a tiled configuration, presenting a unified display space, can increase the astronomer's display capability by orders of magnitude, without the commensurate cost.

Given the cost of building and operating modern telescopes, it is imperative that astronomers do their best to optimise their workflow. To this end, researchers must ensure that relevant content is being displayed on the most appropriate display technology, especially where this can be done cost-effectively.

Yet without the buy-in of the astronomy community, progress in these areas will be slow and face challenges in adoption. It is imperative that the viability of the cloud and capability display solutions be established with the end-users in mind. Benchmarking and cost-benefit analyses are insufficient alone and user experience testing is essential.

Determining the efficacy of a cloud-based desktop, and being confident that a TDW and a well considered display ecology can actually improve research outcomes in astronomy, are the subjects of this research.

# 3

# User Evaluation Methodology

## 3.1 Addressing the research questions

For most astronomers, their local desktop computer is a multi-function device that is used to perform everything from basic tasks like checking email and web browsing, to more complex tasks such as viewing astronomical imagery, collaborating with remote researchers and producing advanced visualisations. When demands placed on the local machine exceed its capacity, it becomes a bottleneck for research outcomes. This research investigates the three research questions posed in Chapter 1 and the options that are available to overcome this obstacle. This chapter serves to provide the background and approach to the specific investigations that led to the publications contained in Chapters 4, 5 and 6.

While it is clear that several technologies exist that are intended to address these question, there is a lack of evidence that astronomy researchers will be able to make appropriate use of the technology, thereby potentially limiting the usefulness of the technology. The focus of this research has been users and their experiences with technologies that promise to take them "beyond the desktop".

To best understand the value of display technologies, it is important to observe how a particular technology will be used by real people in an astronomy research context. Such observations eliminate the imprecise correlation between system specifications and benchmarks with real user experience as a means of determining suitability. This requires observing human subjects engaging with the technology to perform astronomy tasks. Lam et al. (2012) describes best practice in information visualisation, in seven scenarios. Of particular relevance to this research are the following scenarios:

- **Evaluating User Performance (UP):** Users are presented with a visualisation experience and challenged to undertake some set of tasks in a *controlled experiment*.

User performance is monitored by timing the completion of tasks, and *head-to-head* comparisons made between the environments under consideration.

- **Evaluating Communication Through Visualisation (CTV):** The objective of this approach is to determine if the presentation of information aids communication between users presented information both directly and as exposure to ambient information.

- **Evaluating Collaborative Data Analysis (CDA):** Following from CTV, this scenario considers the effect of information visualisation on the collaborative decision-making process. By sharing the data analysis experience, does the environment support better decision-making, and/or does joint decision-making improve the research outcomes?

- **Evaluating User Experience (UE)):** The experience of the user is the most critical element that is often overlooked when the viability of a technology is focused on the cost, benchmarking of results or feature sets. Understanding the experiences of users informs the design lifecycle, improving usability, while also improving user engagement and ultimately buy-in.

The above scenarios directly align with the development and execution of the published papers in this thesis.

## 3.2 Ultra-high resolution displays

In the literature discussed in Section 2.4 TDWs were often lauded as being highly useful in disciplines where visual content, in particular large imagery or multiple related images needed to be inspected. This is especially so in astronomy, where visual information has grown substantially. However, there has been little evidence to confirm this supposition. Therefore, a study was conceived to test the hypothesis that ultra-high resolution displays could improve research outcomes in astronomy. This was achieved by exploring two of the Lam scenarios described above, namely UP and UE.

For a TDW to become part of an astronomer's toolkit, there needs to be clear evidence that its use improves user performance. This might be achieved through the evaluation of several objective measures, such as:

- Improved (reduced) time taken to find critical feature(s) within an image;

- Improved success rate in object confirmation; and

- Improved parallel search performance for collaborative image inspection.

This can be understood by performing a head-to-head comparison between a typical desktop computer and a TDW, an approach also used by Liu et al. (2014) and Prouzeau et al. (2017). By presenting astronomers with Gigapixel astronomical imagery and timing their ability to find features in a controlled experiment, on both the standard desktop display and the TDW, it is possible to determine if a TDW can improve the astronomers' performance. In order to best understand the impact on an astronomer's workflow, a comparison was made with non-astronomers, to determine if the impact was commensurate with previous exposure to astronomical imagery or not.

Just as important as the improvements in user performance is evaluating the user experience with TDWs. If astronomers feel that a TDW can be of benefit to their research, or provides a more pleasant experience than performing the same task on a standard desktop computer, then the TDW has demonstrated value to the researchers' workflow. Such subjective measures are best obtained from the participants themselves, e.g. as solicited feedback through a survey or interview and/or unsolicited feedback by observing participants' mood and engagement with the task.

The resulting combination of objective performance measures with the evaluation of the subjective user experiences provides the basis for evaluating the efficacy of a TDW in the context of astronomy.

The results of this investigation were presented in the paper, "Are Tiled Display Walls Needed for Astronomy?", which appears as published in Chapter 4.

### 3.2.1 Overview of the Study

To address this question, an ultra-high resolution display was required. The OzIPortal facility located at the University of Melbourne, Australia, with a 98 Megapixel display area, was made available for this investigation. The study aimed to answer the following questions:

- Can an astronomer find increasingly difficult astronomical features within a large image?

- Can a TDW be used to improve collaborative inspection of a very large image?

The participants were presented with a set of large astronomical images and shown a feature found within each image, and given a limited time to locate the feature in the larger image. The feature search was conducted on the OzIPortal, with the image shown at native resolution (i.e. each pixel in the image was represented by a single pixel on the display wall), and also on a standard desktop computer display, where the participant could use a computer mouse to zoom and pan the image to find the feature.

### 3.2.2 Population/Sample

57 voluntary participants were recruited to the study, with a mix of astronomers (including astronomy students) and non-astronomers. All participants were over the age of 18 and in an approximate gender balance. All participants were fluent in English. Participants received no compensation for participation.

The choice to split the cohort between astronomers and non-astronomers was made to determine whether a performance improvement was seen relative to previous experience or not. In the case of astronomers, no distinction was made between fields or methods of study, as it was assumed that all astronomers participating would have more familiarity with astronomical imagery in general than the non-astronomers. This approach provided a control group for the use of the TDW for astronomy, where no astronomy experience was necessary. This allowed more accurate attribution of the outcomes of the study to the use of the TDW.

### 3.2.3 Location

The availability of a suitable TDW afforded limited choice, however, access was provided at one of the largest TDWs in Australia at the time, the OzIPortal, located at the University of Melbourne. The facility was located in a private room at the Parkville Campus. The facility was not being used for other research at the time of the study and it was possible to configure the space to the study requirements. The room layout is included in the published paper (p.9 Meade et al., 2014, see Chapter 4, Section 4.3).

### 3.2.4 Restrictions/Limiting Conditions

Because of the limited number of previous studies on the use of TDWs in astronomy, and the lack of research opportunities, it was necessary to construct a series of artificial activities that closely resembled the kinds of tasks astronomers might use such a resource for.

Once the room was configured for the study, it was impractical to change the setup to accommodate the personal preferences of the participants. Hence, taller participants experienced some difficulties using the lowest row of screens on the TDW, while the shortest participants found the topmost row difficult to inspect.

While it would have been ideal to explore longer-term exposure to the TDW, it was decided to limit the overall experience for the participants to around 30 minutes, including the post-investigation survey.

### 3.2.5 Procedures

Participants were asked to read the Informed Consent form before attending their session, and given time to ask questions and review the form before signing it at the start of the session. Once completed, the purpose of the study was described to the participant and the steps that they would be expected to follow. The form is available in Section A.1 of Appendix A.

The participants were alternately presented with the TDW or the standard desktop to begin their interaction. Each series of searches was limited to two minutes. A search target was shown on a large TV screen and laptop screen that was positioned where it was easily seen from a position in front of either the TDW or the standard desktop display. The participant would then point to the target in the source image when they had identified its position. In the case of the TDW, the participants would physically approach the wall (called *physical navigation*) to find the target, while for the standard desktop, the participant would employ a computer mouse to pan and zoom the image (called *virtual navigation*), again, pointing to the target when they were confident they had correctly identified it.

The targets were made increasingly difficult to find in the source image, as they were taken from increasingly smaller features of the source. Each image had twenty possible targets to find, with two at each size, progressively getting smaller. The number of targets found in the two minutes was recorded for each participant.

After the main investigation was completed, the participants were asked to complete a short survey comparing the standard desktop display with the TDW.

### 3.2.6 Materials

In order to ensure a suitable baseline for the cohort, an image was created using common English words[1] and presented at a range of sizes, from very large to very small. At the smallest sizes, the words were too small for the standard desktop display to render, forcing the participant to zoom in to see them. The same images were used in both environments,

[1]Words were taken from the list of common English words found at `Wordsearchdensitywordssource: http://www.anglik.net/english250.htm`

however the TDW allowed the smallest words to be rendered legible, though this required the participant to approach the wall to inspect more closely.

All search targets were presented at the same size, irrespective of their actual size in the source image. Targets were not rotated, appearing in the same orientation and colour as in the source image.

## 3.3 Workspace display ecologies

Having successfully investigated the usability of a TDW in the context under tightly controlled conditions, it was serendipitous that an opportunity arose to test the TDW described in Meade et al. (2014, see Chapter 4, Section 4.3) in a real astronomy observation campaign. The *Deeper, Wider, Faster*[2] campaign was designed to detect fast transient astronomical events in multiple wavelengths on many telescopes (Andreoni & Cooke, 2018). Having determined that the standard desktop and laptops used by the astronomers in the pilot phase of the campaign would be insufficient, the project leaders decided to employ the TDW located at the University of Melbourne, Australia. At that time, the TDW was housed in a room with additional display resources, including a large curved display wall with high-definition projection. Combined with several additional displays, a *display ecology* was formed, with different parts of the workflow presented on the most suitable display. Having the workflow entirely accommodated on the available displays meant that individual workstations were typically single-purpose, but visible by everyone in the room. This provided both the optimised visual presentation of information, as well as the casually available ambient information to support communication.

As the campaign ran several times, there was opportunity between runs to reconfigure and optimise the environment. While significant structural changes were made to the environment for each run, including moving much of the equipment to the Swinburne University of Technology Hawthorn campus (also in Melbourne, Australia), the objective of the display ecology continued to enhance the campaign.

During each observing run, data from several sources were displayed in the workspace simultaneously. Images captured by Dark Energy Survey Camera (DECam)[3] were processed to produce thumbnails of potential candidates and difference images of the full-resolution CCD images, as well as light curves of the candidates over time. Meanwhile, real-time communications with the remote telescope operators were maintained, with telescope status screens also being monitored, as well as connections to the HPC facility.

[2] http://www.dwfprogram.altervista.org/
[3] https://www.darkenergysurvey.org/the-des-project/instrument/

Rather than have all this information synthesised by one or two people, each station was monitored independently, with key information relayed to the principal investigators, allowing a CDA process to occur.

The co-location of the displays supported a rapid CDA process by enhancing the communication of information both within the workspace and to the remote telescopes operations sites. Not only was salient visualisation information readily communicable to the principal investigators, but also to other participants in the workspace (CTV). This made induction to the process quick and efficient, as new participants could learn the process by observing someone more experienced, until they were ready to contribute more actively. The ambient visualisations on the other displays in the workspace aided their understanding of the whole process, and therefore allowed them to more easily understand the importance of their contribution to the whole workflow. While the primary objective of the display ecology was to enhance collaborative decision-making, the overall display ecology also enhanced the communication among the *Deeper, Wider, Faster* team.

Participating in the *Deeper, Wider, Faster* campaign provided the opportunity to observe the suitability of the implemented display ecology to support both the CDA and the CTV scenarios, and hence the value of a properly designed display ecology.

The results of this investigation were presented in the paper, “Collaborative Workspaces to Accelerate Discovery”, which appears as published in Chapter 5. Details of the *Deeper, Wider, Faster* campaign can be found in Andreoni & Cooke (2018).

### 3.3.1 Overview of the Study

The researchers involved in the *Deeper, Wider, Faster* campaign had established the need to better handle the flow of images from DECam. Overwhelmed by the rate of images arriving, the team realised a collaborative, parallel inspection approach would greatly improve the throughput of image analysis. Also, because each image was $4096 \times 2160$ pixels, the standard display screens of the researchers’ desktops and laptops would require lots of virtual navigation, i.e. zooming and panning. Including a TDW in the workflow allowed the display of several full resolution images simultaneously. This meant that several astronomers could share the image inspection task, and thereby eliminate the need to pan and zoom the images.

However, the TDW was not the only piece of advanced display technology available. At the opposite end of the room was a large, curved display wall that used two overlapping $1920 \times 1200$ projectors. The output of another part of the workflow produced thousands of thumbnail size images, extracted from the DECam images, showing potential fast transient candidates. These candidates were projected together in groups of three, including the extracted target, a calibrated science image of the same spot in the full image, and a subtracted version. Tens to hundreds of these groups were projected simultaneously on the curved projection wall and inspected by another group of astronomers. Potential candidates were reviewed and discussed in the context of both displays.

Between these two groups, a third group of astronomers were facilitating the whole process. They coordinated the telescopes and the image processing pipeline, as well as reviewing astronomical catalogues, and ultimately making the decision to initiate a trigger to be sent to the standby telescopes to join the collective.

Before the observation run was conducted in the new workspace, the participating astronomers were asked to complete a short survey detailing how they hoped the new environment might improve on the previous run. In general, the group hoped that the addition of the TDW and the curved project wall would increase their ability to inspect the incoming images, thereby improving the throughput and ultimately allowing the principal investigator to make justified calls for fast transient event triggers to be sent.

During two observation runs, the astronomers were observed as they moved about and used the facility. The aim was to monitor whether the layout and configuration of the room and the display technologies augmented or detracted from the workflow. It was clear very early on that not only did the astronomers' performance improve, both in terms of throughput and confidence in the image inspection; they were also enjoying using the space itself. The ability to share discoveries almost instantly, while also being able to physically move about and see what others were working on, was appreciated by the participants. Training new volunteers was also made considerably easier as the new participants were easily able to shadow someone with more experience without getting in the way, until they were ready to go solo.

Observations of the positive impact that the display ecology had on the observing run was borne out in the follow up survey completed by the participating astronomers. While the space was not perfect, the participants felt that the display ecology did accelerate the workflow, improve communication and collaboration, and improve their enjoyment of the observation runs.

### 3.3.2 Cohort

The opportunity to participate in the *Deeper, Wider, Faster* campaign came up with relatively short notice, so there was not sufficient time to get ethics approval to conduct a human trial study during the observing run. The decision was made to include the participating astronomers as co-authors on the paper. This ensured that the survey results could be included and provided by the cohort of astronomers. Of the 16 authors, 14 participated in the survey, of which 3 were female and 9 were students. The two people who designed the survey did not contribute to the results.

### 3.3.3 Location

The two observing runs that were the subject of the Meade et al. (2017, see Chapter 5, Section 5.2) paper were both conducted in the Advanced Immersion Environment located at the University of Melbourne Parkville campus, Australia. The room was not used by any other project and was due for decommissioning, but was made available by the Property and Campus Services group for the duration of the observing campaign. This allowed the room to be freely reconfigured as required.

The room has since been decommissioned and repurposed, and the curved projection screen was dispose of. The room layouts are included in the published paper (p.6 and p.16 Meade et al., 2017, see Chapter 5, Section 5.2).

### 3.3.4 Restrictions/Limiting Conditions

No additional funding was made available to purchase equipment specifically for the display ecology, therefore the equipment was the best that was freely available at the time. The machines used to drive the TDW and the curved projection wall were functional, but not as powerful as desired. Network performance was limited by the 1Gbps network to the room. These limitations resulted in minor performance issues, such as slow loading of images on the display wall, but generally this did not impact on the result of the observing run.

Other technology limitations, such as the computer driving the curved display wall being limited to Windows, and the TDW being limited to certain file formats, forced some changes to the workflow to use them effectively, though not always efficiently. The file format limitation of the TDW was addressed in the second observing run by deconstructing the TDW into multiple smaller multi-screen workstations, thereby allowing each workstation to run independently and support the preferred image format without modification.

### 3.3.5 Procedures

Before the first observing run, the participating astronomers were invited to complete a short online survey. The purpose of the survey was to ascertain what the participants hoped the new display ecology might do to improve the campaign outcomes, and to establish information on their previous experiences with TDWs. A follow-up survey was sent immediately after the completion of this run, with more detailed questions focusing on their experiences during the observing run. This included reflection on the benefits of the display ecology, but also where they saw opportunities for improvement. Some of these suggestions were able to be implemented in the second observing run, such as the reconfiguration of the TDW into smaller workstations.

## 3.4 Remote desktops using the cloud

The use of cloud computing services presents both opportunities and challenges to the astronomy community. On the one hand, the promise of almost limitless resources that can be used and released as required is very attractive to the financially restricted researcher. On the other hand, understanding the best way to use a remote resource, and not get caught by the pitfalls of billing and security, might appear to be more hassle than its worth. Fortunately however, there is considerable research available to provide the guidance necessary to ensure the benefits of cloud computing.

This section is broken into three parts:

1. **Research Cloud Data Communities:** Cloud services designed to support research, called Research Clouds, are optimised to support a wide range of research disciplines. One of the key benefits to cloud computing is the ability to share resources and form communities around tools and datasets. In the case of the Nectar Research Cloud in Australia, considerable support has been given to developing *data communities*, that share resources designed to facilitate access to and processing of significant datasets.

2. **Seeing the Big Picture: A Digital Desktop for Researchers:** The inability to move beyond the local desktop has left many researchers caught in the upgrade cycle of their research institution, e.g. buying the most powerful computer they can afford at the start of each cycle in the hope that it will still be useful at the end of the cycle. However, when faced with the challenges of the big data era, where data is stored in a data centre and is too big to download to a local machine, the use of a remote desktop service becomes necessary. In fact, with appropriate infrastructure, these remote desktops can outperform many local computers.

3. **Evaluating VHDs for graphics-intensive astronomy:** Taking the previous step further, the use of commercial and research focused cloud services for provision-

ing VHDs is a reality that few researchers, let alone astronomers, have considered. However, with the improvements in network stability that most research institutions now provide as standard, and the broadband networks that reach into most homes, cloud-based desktop computing provides a cost-effective alternative to the traditional computer lifecycle.

### 3.4.1 Research Cloud Data Communities

The first of the papers in this section was intended to document the landscape of the emergent research cloud and the research communities that had begun to tap into it. The paper set about defining the basic concepts of big data and cloud computing, and then introduced the Nectar Research Cloud, the national federated cloud in Australia. Published shortly after the Nectar Research Cloud had launched, it highlighted the early success stories of the Research Cloud, and also described some of the challenges facing the early adopters. In particular, offering IaaS as the primary service meant that a significant learning curve existed for most researchers. While training was always a key strategy to address this learning curve, considerable effort was also put into developing the user communities. Sometimes these communities formed around specific tools, but more commonly around domains. Nectar funded several virtual laboratories, which provided a portal interface to tools and datasets for their communities.

However, beyond these funded virtual laboratories, many other research groups have established their own data communities, sharing resources and datasets, to enhance their research and collaborations. This paper aimed to provide guidance to those wishing to establish a community, and understand how best to engage with the emerging cloud technologies, and understand them in the context of more traditional options such as HPC.

The paper, "Research Cloud Data Communities", was presented at The Higher Education Technical Agenda (THETA) conference in Hobart, Australia, 2013, and appears as published in Chapter 6, Section 6.2.

### 3.4.2 Seeing the Big Picture: A Digital Desktop for Researchers

While TDWs have been shown to be a valuable asset in astronomy, they have significant drawbacks. In particular, they are orders of magnitude more expensive than high-end workstations, and they require a suitable space to house them. Smaller TDWs are increasingly available, with modern graphics cards capable of driving multiple standard sized displays at once. A cost effective compromise to consider is the commodity UltraHD display, sometimes called 4K, with a resolution of $3840 \times 2160$ pixels. With a variety of sizes available, these high-resolution screens are low-cost and provide viable intermediate options for researchers.

At this scale, an intriguing option becomes possible. At only four times the size of a $1920 \times 1080$ display, driving this display via a cloud-based desktop becomes a possibility. Provided sufficient network bandwidth is available, a VM configured with a suitable GPU can stream a 4K desktop to a local computer, provided that computer has sufficient power to drive the 4K display. Many modern desktop and laptop computers are able to support a 4K display at 60Hz. In 2015, GPU-enabled VMs were not available on the Melbourne Node of the Nectar Research Cloud, so a trial server courtesy of Dell Computers, with $2 \times$ Grid K2 graphics cards, courtesy of NVIDIA, was used to establish a 4K remote desktop. NVIDIA's Virtual Desktop Infrastructure (VDI)[4] solution was located in a data centre less than 1km away from the test site, with a 10Gbps network connection between buildings, down to a 1Gbps connection to the display computer from the building switch. The bandwidth and latency was such that the desktop experience closely matched a local desktop computer, though some frame delay was apparent.

While a campus network might be expected to provide reasonably stable connectivity, this is not necessarily the case for inter-institution networks, especially between states. The Australian Academic Research Network (AARNet)[5] provides the academic research network backbone across Australia. In conjunction with Pawsey Supercomputing Facility and NCI, supported by the Nectar Research Cloud, a test of the network performance was conducted. The aim of the test was to determine if suitable bandwidth and stability would allow a remote cloud service to provide a streamed desktop to a local computer with a 4K display, or a TDW. At the time of the test, GPU-enabled VMs were not available, so the test simply considered the performance with continuous file transfers, and the network load was monitored. The test confirmed that while the connection to NCI was stable and could provide the necessary connection conditions, the Pawsey connection was too erratic (at the time) to be reliable.

The results of this investigation were presented in the paper, "Seeing the Big Picture: A Digital Desktop for Researchers", which was presented at The Higher Education Technical Agenda (THETA) conference on the Gold Coast, Australia, 2015, and appears as published in Chapter 6, Section 6.3.

[4] `http://www.nvidia.com/object/vdi-unleashed.html`
[5] `https://www.aarnet.edu.au/`

### 3.4.3 Evaluating Virtual Hosted Desktops for Graphics-intensive Astronomy

While the use of cloud computing is growing in astronomy, the general perception in the community is that it is a purely compute and storage resource, more flexible but less powerful than a high performance cluster. Using a remote desktop, let alone a cloud-based remote desktop, was considered a necessary evil for many astronomers who relied on software or interfaces that demanded a graphical user interface. The aim of this study was to challenge this notion by allowing astronomers to experience an optimised VHD based on cloud infrastructure. Armed with objective experience, the participants would then be able to assess the viability of a VHD as a potential addition to, or replacement of, their existing desktop computing solution.

Following a similar design process as described above in Section 3.2, the VHD study was intended to explore the viability of a cloud-based desktop when compared to a typical local desktop for astronomy-related tasks. The study aimed to test the hypothesis that a cloud-based desktop could provide a user experience to closely match or exceed the performance of a local desktop, without diminishing the user's performance, when considering both objective and subjective measures. As with the TDW study, both the UP and UE scenarios described by Lam et al. (2012) were considered important to study.

Objective measures for a desktop in astronomy can be obtained by benchmarking software, but this does not always reveal sufficient information about the suitability of the underlying system. It is necessary to use the system in the context for which it is intended, namely supporting the required astronomy applications, but more importantly, it needs to be used by the researchers intending to use it. For example, a system that performs adequately from one astronomer's perspective might be grossly inadequate from another astronomer's perspective. One size does not necessarily fit all. When purchasing a desktop or laptop computer, the range of customisation is usually limited to a small catalogue. Also, once a purchase is made, it is often difficult to reverse that decision. Cloud computing offers an alternative that avoids the commitment of an outright purchase, and provides a far greater level of customisation and flexibility, but this is only useful if the cloud-based solution provides an acceptable user experience (UE). In this context, the cloud-based desktop might not match a local desktop, but the impact of customisation and flexibility, such as being able to choose more RAM or more CPUs as required, as well as the cost of service compared to outright purchase, might offset a lessened user experience. As before, evaluating the users' experience with a cloud based desktop is best achieved through user feedback.

The subjectivity of user experience can be combined with a head-to-head comparison of objective performance measures to provide a more complete picture of cloud-based desktops. Directly measuring UP provides a better objective measure of a system than simple benchmarks. It also allows for the variability of different users in a way that benchmarks cannot. Some objective measures considered include the following:

- Does the desktop environment (local or cloud) make a difference to a participant's ability to learn how to complete a task?
- How does the graphical performance (e.g. the frame rate during graphically intensive tasks) of the desktop environments compare?
- Does a cloud-based desktop provide a better value for money option?

This can be understood by performing a head-to-head comparison between local laptop computers and VHDs. By presenting astronomers with a set of simple astronomy tasks to be performed on a local laptop and on a VHD, we were able to provide a direct comparison of the environments. Using controlled experiments, where the participant provided real-time feedback and the system's graphical performance (frame rate) was simultaneously logged, the objective performance of the system could be aligned with their perceptions of the system performance (UE).

The combination of the participants' perceptions of performance and the reported system performance provides a further mechanism to challenge the suitability of basic benchmarks to determine the most appropriate desktop solution. The cost of purchase of a local laptop was also compared to the equivalent expenditure over a fixed period for two cloud-based solutions.

The results of this investigation have been accepted for publication (April 4, 2018) in the paper, "Evaluating VHDs for graphics-intensive astronomy", which appears as published in Chapter 6, Section 6.5.

### 3.4.4 Overview of the Study

This study involved a combination of user interviews and a practical component. For the practical component, the study presented a desktop experience on either a 2013 model Macbook Pro laptop or a 2017 model Macbook Pro laptop, with a VHD, provisioned on the Melbourne Node of the Nectar Research Cloud and AWS (Sydney data centre). The tasks presented were repeated as close to exactly the same as possible - though there were some minor interface differences.

Importantly, the investigation was designed to closely mimic the tasks an astronomer might undertake during a research activity. Rather than free exploration of the desktops, we wanted to focus the participants on completing work-like tasks. This would focus their reflection on the efficacy of the environment, rather than on any preconceptions they might have.

The interview was split into pre- and post-investigation:

- **Pre-investigation:** Participants were asked to respond to questions about their previous experience with cloud environments and remote desktops. The aim of this interview was to provide a profile of the cohort before they had experienced an optimised VHD, for comparison with the follow up interview.

- **Post-investigation:** Having completed the investigation, the participants were interviewed again, this time inviting them to compare the local desktop experience directly with the one or two cloud-based VHDs. They were also encouraged to consider the problems they experienced during the investigation, and to reflect on any shift in their perception after the experience.

During the task phase of the study, the participants were asked to complete two graphical-based astronomy tasks:

- **2D image alignment:** The first task was to perform a simple image alignment, so that the desktop experience could be compared between the local desktop and the VHD. Using standard 2D astronomical imaging software (SAOImage DS9[6] and IRAF[7]), the participants were timed aligning two FITS images with a known offset. No previous experience was expected or required, and the precise steps were provided to the participants. The tasks were repeated on a local laptop and either one or both of the cloud environments. The first half of the cohort were presented with the local laptop first each time, followed by the VHD in the Research Cloud or AWS, alternately. The second half of the cohort experienced a local desktop or a cloud desktop alternately. This was intended to help determine if the order of presentation made a difference to the outcome. Timing of each completed alignment was then compared for each participant and for each environment in the order they completed them. This would make it possible to determine if the impact of the environment was significant.

[6] `http://ds9.si.edu/site/Home.html`
[7] `http://iraf.noao.edu/`

- **3D spectral data cube manipulation:** The second practical task was the manipulation of a 3D spectral data cube using a 3D volume rendering software, Shwirl[8]. The task required the participant to load a 3D volume of a galaxy, and manipulate it with a variety of graphically intensive modifications applied. After each step, the participant reported a value between 0 (bad) and 10 (excellent), according to how they felt the desktop performed. At the same time, Shwirl logged the frame rate of the graphics card during the task, which could then be aligned with the participant performance perceptions later. As per the 2D task, the participants completed the 3D task on a local laptop and one or two cloud-based desktops.

As noted in Section 3.7 of Meade & Fluke (2018, see Chapter 6, Section 6.5), users participating in the study were asked for their subjective responses. The normalisation approach was introduced because it was apparent that in all cases there was a spread of values reported. Clearly, all users saw that some configurations were better than others. For this reason, normalisation between 0 and 10 was introduced, which allowed meaningfully comparisons of subjective responses across the cohort. However, this approach would not have been applicable if some of the cohort had reported no variation in their experience. Had this situation arisen, it would have been interpreted as an indication that benchmarking was completely irrelevant.

Combining the participants' reported experience during the tasks and their reflections (UE), and their ability to complete tasks (UP), through both timing (for the 2D task), and the logged frame rates (3D task), provides a useful evaluation scenario (Lam et al., 2012).

### 3.4.5 Population/Sample

The study was conducted in accordance with Swinburne University ethics requirements. Participants were solicited via email, announcements at departmental meetings, and through word of mouth.

The participants were recruited from the astronomy departments of Swinburne University of Technology, (Hawthorn, Australia), and the University of Melbourne, (Parkville, Australia). The participants were all astronomers, or astronomy students (including recently graduated students) and all were over the age of 18. Gender information was not captured. Astronomers from various areas of research were included, and not limited to observational astronomy. A total of twenty participants completed the study.

[8] `http://shwirl.readthedocs.io/en/latest/`

### 3.4.6 Location

The sessions with each participant were completed in private offices at Swinburne University Hawthorn campus, and in a private office at the University of Melbourne, Parkville campus. The VHDs were provisioned from the AWS data centre in Sydney, and the University of Melbourne data centre in Parkville, Melbourne.

### 3.4.7 Restrictions/Limiting Conditions

It was not possible to source a wider range of local desktop computers including high-end graphics workstations, nor investigate different cloud-based GPUs. At Swinburne University of Technology, the wireless network, *Eduroam* was used, while at the University of Melbourne, some of the participants used the 1Gbps wired network, while the others used the wireless network, *uniwireless*.

A key limitation for this study was the inability to conduct a longer trial, either with the participants performing a more extensive set of tasks, or repeating the tasks over weeks or months. While this would have been ideal, it would have been harder to recruit sufficient participants to take part in the study.

### 3.4.8 Procedures

Participants were scheduled for a 45 minute session. They were sent the Informed Consent form ahead of time to reduce the time required for this step during the session. The form is available in Section A.2 of Appendix A.

The session procedure was as follows:

1. The Informed Consent was signed by the participant and they were given a brief introduction to the study, expanding on the explanation contained in the Informed Consent form.

2. The network status was logged.

3. A brief interview was conducted to gauge previous experience with cloud and VHDs.

4. Participants were then taken step by step through the 2D image alignment task, without timing the process.

5. The first task was completed, using the local laptop directly, or using the local laptop to access a cloud-based VHD. This task was timed, and repeated on the other desktops.

6. The 3D spectral data cube manipulation task was then completed in each of the desktop environments. This task was not timed, but the GPU performance was logged during the task. The participants' perceptions of performance was logged for each step.

7. A final interview was conducted to capture the reflections of the participants.

8. The network status was logged.

9. Logs were saved with a filename that allowed them to be associated with a participant's task perception responses. These responses contained no participant identifiable information.

### 3.4.9 Materials

The FITS images used for the 2D image alignment task were sourced from the SDSS[9]. The 3D spectral data cube was NGC628 from The Hi Nearby Galaxy Survey data, sourced from http://www.mpia.de/THINGS/Data.html Participants at Swinburne University were provided with a Macbook Pro 2013 to complete the tasks. Participants at the University of Melbourne used a Macbook Pro 2017. All participants were given precise instructions (hard copy) with all the steps for each task clearly laid out.

[9] `http://www.sdss.org/`

# 4

# Ultra-high Resolution Displays

## 4.1 Overview

Astronomy has entered an era of unprecedented data capture, simulation and analysis, and as such it is necessary for the tools of the astronomer to evolve in step. Computing power has to scale to cope with the data deluge and provide the astronomer with connection to the information. However, the display technologies employed by the typical researcher have languished behind this evolution. TDWs provide researchers with a way of increasing the resolution available to examine their data, by parallelising the image across many displays. However, early attempts to engage researchers across all disciplines had been largely unsuccessful, due to expensive hardware, unstable software, and complex data preparation processes as typified by use of the OptIPortal. However over the last few years, these issues have been gradually addressed. Despite this, astronomy – one of the most visual of sciences – has largely failed to benefit from this approach to visualising data. The rest of this section argues that it is now time for astronomers and astrophysicists to seriously consider employing TDWs in their research workflow, to maximise the scientific returns of the data they work so hard to capture.

### 4.1.1 Ultra-high Resolution Images

Astronomy produces data at incredible rates, among the highest of any research discipline. New and proposed telescopes such as the Square Kilometer Array and the Large Synoptic Survey Telescope will produce datasets heading toward the exabyte scale (Quinn et al., 2014). Existing data analysis tools and methodologies, usually centred on the astronomer working directly on the data at the desktop, will be pushed to their limits. There will be an ever-increasing reliance on automated processes to identify objects of interest, e.g. Koribalski (2012) provides a summary of the so-called “source finding” problem in radio astronomy.

One of the solutions for viewing very high resolution images on standard display screens is the Hierarchical Progressive Survey (HIPS) scheme (Fernique et al., 2017). In this scheme, images are presented as tiles of increasingly finer resolution, and only those tiles necessary to fill the display size of the viewing screen are loaded. This facilitates multi-resolution panning and zooming, with lower resolution tiles displayed while higher resolution tiles are being loaded. This approach improves the viewing experience, but does little to aid the astronomer in finding targets of interest, or remembering which part of an image has already been searched. Figure 1 in Meade et al. (2014, see Section 4.3) shows an example of an image that exceeds both standard desktop displays and a 98 Megapixel TDW.

Advances in astronomy in the future will rely increasingly on the ability to discover important details hidden inside vast datasets. Examining gigabytes, petabytes or even exabytes of data requires moving beyond the small screens attached to desktop computers.

To tackle this requires the convergence and maturity of three technologies, which will help astronomers optimise their research workflows through enhanced capabilities for inspection of datasets:

1. Processing of data stored in remote archives in order to reduce the need to move enormous datasets around the globe;

2. Ultra-high resolution tiled displays, providing researchers with a means to display images at more meaningful sizes; and

3. Collaborative environments such as the OptIPortal network, to allow distributed research teams to work together simultaneously on the resulting imagery and analysis. This provides greater opportunity for collaborative research, which should lead to greater understanding of the data and hence scientific breakthroughs (Smarr et al., 2009; Sims et al., 2010; Yamaoka et al., 2011).

Astronomers deal with a variety of data types and formats, from observations across the electromagnetic spectrum and from computer simulations. While practical workflows have been developed to suit their processing and visualisation needs, these workflows do not always evolve at the same rate as the influx of new data. Figure 4.1 shows that the instrument image sizes now possible and those expected in the near future. As seen these are growing far quicker than the display technologies used to inspect them. In this section, we highlight several key areas in which a TDW could greatly enhance existing workflows.

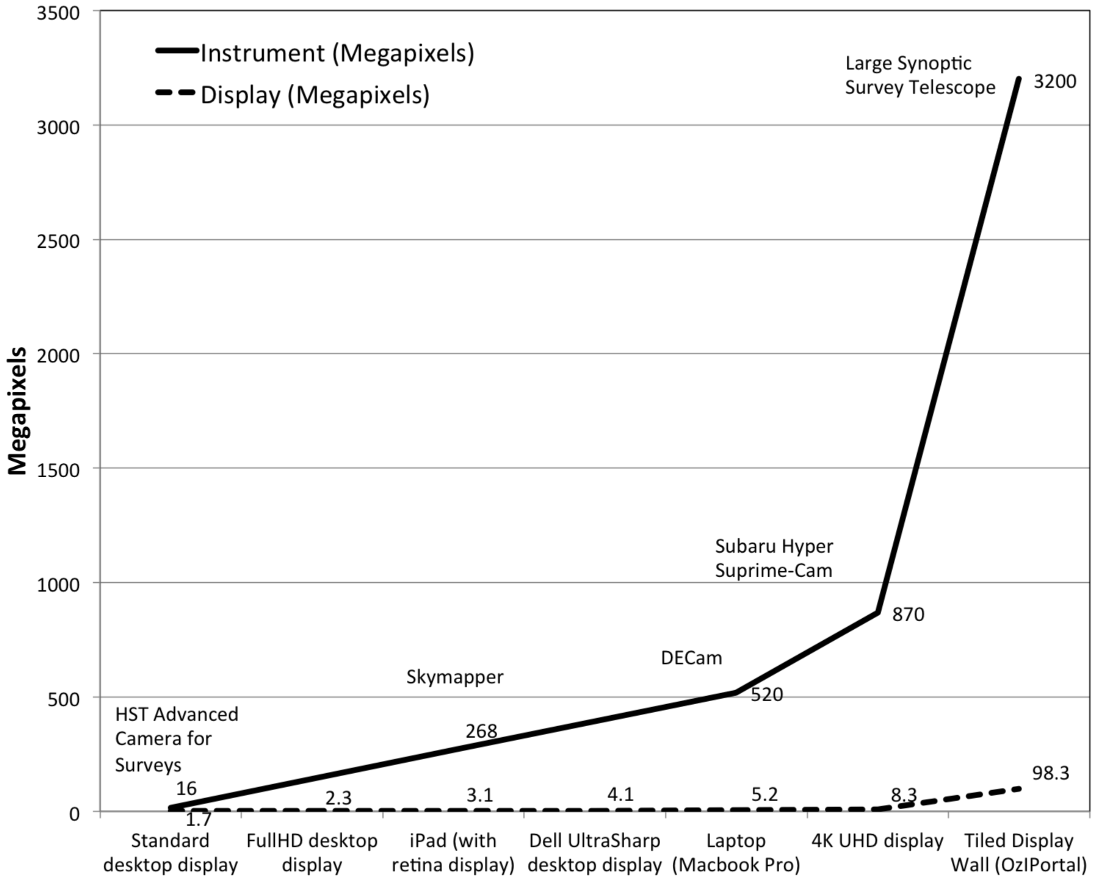


Figure 4.1 Growth in instrument image sizes far outpaces the growth in display technology.

### 4.1.2 Simultaneous Views of Multiple Images

Astronomers often need to examine subtle differences between several images at once. These images might be time-interval captures, the comparison of similarly structured galaxies, or a single dataset that has been imaged through several different filters (i.e. through different wavebands). The side-by-side comparison of many detailed images simultaneously can be difficult if they are all presented on the same desktop display, or scrubbed through in sequence. For example, a 100 Megapixel TDW could allow 24 individual 2K[1] images to be viewed simultaneously.

Activities such as transient surveys require the examination of multiple snapshots of particular regions of the sky taken over a period of time. Hundreds of images are used to determine candidates and while automation has improved the efficiency of this process, by-eye inspection remains critical to the success of these projects. Using a TDW to display batches of images simultaneously would allow an astronomer to utilise inherent pattern matching skills, gained over years of observation of the sky, to identify objects or regions of interest rapidly. See Chapter 5 for an example of an observation campaign supported by a TDW.

Using computers to perform image comparison often benefits from concurrent observations performed by trained astronomers, e.g to allow the refinement of the algorithms that will improve the reliability of computer detection. It is also important to note that computers will typically search for what they are directed to look for, so should new phenomena be captured, there is a significant chance that an automated detection process will overlook it.

### 4.1.3 Revisiting Archived Data

It is easy to only look forward when developing and employing new technologies, however, it may be that some discoveries are locked away in historic archived data. The value of a TDW would be well established if an observation of significance could be made using datasets that had previously only been examined using a standard computer display or was subjected to automated detection. It may be that some archived datasets containing astronomical treasures that have gone undiscovered simply require displaying on a suitable TDW in front of a new set of eyes.

While the principle is quite simple, it would likely be very difficult to effectively plan and execute such a project. The difficultly would lie in determining which datasets would

[1] 2K = 1920 × 1200 pixels

likely contain useful information and what size display would be useful to ensure that adequate coverage would be achievable. Should such datasets exist, they would probably consist of images or sets of images that closely matched the resolution of the TDW. Candidate data would also have been likely to have only had cursory consideration in previous examinations.

### 4.1.4 Linked Applications

Astronomers use many tools to interrogate their data, and quite often these datasets are high dimensional in that they go beyond simple 2D data (Goodman, 2012). However, it is not always convenient or even possible to meaningfully present all dimensions simultaneously. Also, such multi-dimensional data cannot necessarily be processed by a single application. Astronomers have developed many tools to solve certain problems and it is sometimes useful to coordinate these tools to process data from a single source, with each application utilising the dimensions relevant to its function. For example, when focusing on a particular object or region in the sky, one application might display the radio spectrum for the object, while another shows the X-Ray sources for the same patch of sky. Other windows might display optical images with a variety of filters, and others show the World Wide Telescope (WWT)[2] view of the region of sky. The result can be several windows, each displaying important information derived from a single dataset, or several related datasets. Using the WWT window, the astronomer can select an adjacent region of sky, and all the windows automatically update to reflect the new selection.

Live linking of these applications is partially possible already using Simple Application Messaging Protocol (SAMP; Taylor et al., 2015)-enabled applications, but this linking is likely to increase in the near future. Observing changes with several windows open on a standard desktop can be problematic, whereas the same presentation on a TDW would provide easy visual access to all applications simultaneously.

### 4.1.5 Remote Data Processing

With the increasing volumes of data being captured, it is no longer feasible to transfer this information to each independent site that wishes to inspect it. Instead, remote processing facilities, typically located in close proximity to the capturing instrument with high-speed interconnects can provide a high-definition video feed of the data being processed on a VM. The output from these VM operations can be rapidly distributed to numerous sites, while interaction with the VM itself can be shared with each site. In this way, the displayed

[2] `http://www.worldwidetelescope.org/`

information is limited in volume to the capability of the end display device, be that a TDW or a researcher's desktop. Anticipated expansions in network bandwith (Hancock, 2012) will allow further improvements in performance, such as increased frame rates and lower compression, but real-time interaction with petascale and larger datasets will be possible over any reasonable connection. Researchers will thus be able to access and process data in collaboration with colleagues using a variety of network conditions.

### 4.1.6 Collaboration and Training

Changes in technology and the ubiquity of network connectivity means researchers are not constrained by location when working in collaborations. However, dealing with such massive datasets can make remote collaboration difficult, as the data volume growth outstrips the increase in bandwidth. TDWs provide a mechanism to bring dislocated researchers one step closer together, by providing a shared display space to communicate while examining data. High-definition video-conferencing provides the necessary real-time audio and video feedback necessary for dynamic interaction between colleagues, while a unified display space allows them to brainstorm ideas and share information as if they were in the same room.

Sharing the virtual display space with software like SAGE2[3] allows collaborators to share their own local desktop, as well as content directly on the TDW. Pointers representing each site can be displayed, allowing remote participants to draw attention to features of interest in the display space. In such an environment, videoconferences can be streamed onto the TDW, though in practice it is usually more useful to position a remote video stream of a colleague to the side of the TDW, where one might expect a colleague to stand if they were in the same room.

Presenting digital content to more than one or two people usually benefits from having a larger display than is available on a standard desktop or laptop computer. Where several sources of visual information need to be understood in concert, and the links between them explained by an expert, a TDW can be a valuable training asset.

### 4.1.7 Quality Control

The rate at which data is captured or generated is already putting severe strains on the ability to conveniently store and transfer data. In the case of new telescopes such as the ASKAP, it will not be possible to retain the raw data for more than 12 hours. It will be necessary for astronomers to decide how best to process the data within a

[3] `http://sage2.sagecommons.org/`

given time window before the storage is cleared to make way for the next observation. Should a mistake be made at that time, the processed data would be useless and the data lost. Therefore, it is extremely important that astronomers are able to rapidly process and perform quality control on the processed data to ensure they have made full use of the captured data. Hassan et al. (2010) managed to optimise this process considerably, allowing data processed on a co-located GPGPU cluster to be distributed to remote sites in near-real time. This allowed several astronomers to observe the results and decide if the processing choices were in fact appropriate to the observation.

Taking this a step further, if collaborating researchers were able to quickly determine the success or failure of the processing, then it should be possible for an additional optimisation step to be included. For example, an alternative calibration model could lead to improvements in the signal-to-noise ratios of sources. The quicker the researchers can observe the results of data processing, the more time they will have to reprocess the data and improve the research outcome. Every additional reprocessing of the raw data can dramatically improve the results.

### 4.1.8 Time-Critical Tasks

When considering several of the scenarios described above, it becomes apparent that one aspect of examining data that may not have occurred to many astronomers, and indeed many other disciplines, is the time it takes to inspect a dataset with sufficient due diligence. One of the reasons many astronomers have not adopted TDWs in their workflow, is because they are actually able to achieve much of what they require with the equipment they already have. We are all familiar with the need to zoom in on high-resolution images in order to see the small details, and we often do so quite effectively even on displays with far lower resolution than the images being observed.

Aside from the question of whether this is always successful, it is also useful to consider how long such an inspection takes. It is easy to see how important rapid evaluation is when considering the limited window of access to data as described in the previous section. It is perhaps not as clear when considering the more typical workflow, when such limited access is not an issue. In this case, the value of a TDW is not necessarily whether or not a feature is discovered, but rather, how long it takes an observer to find it. If using a TDW can be shown to consistently reduce discovery time, then the value of the TDW can be found in the time saved.

In recognizing key features within an image, or identifying coincident phenomena across multiple images, the efficiency of the observation is greatly affected by the ability of the observer to retain the details in their own short-term memory. No doubt, zooming in and out several times aids the visualisation process, but this takes time, and for sufficiently large images, may prove impossible for some observers, e.g forgetting the image feature sought.

### 4.1.9 Outreach and Education

The visual scale of TDWs makes them an appealing tool for research promotion. Highly detailed imagery, or the presentation of multiple data sources simultaneously, provides a compelling experiences for audiences unfamiliar with such a display. Many people are familiar with the physical scale of the display through cinema experiences, however many TDWs have a resolution that match or in some cases exceed that of IMAX cinema[4]. A TDW is far more accessible, and content is more easily prepared than for an IMAX cinema. The Chicago-based Adler Planetarium[5] is an example of a TDW being used for outreach and education.

## 4.2 Obstacles to Adoption

The visual appeal of astronomy has meant that spectacular images from telescopes have often been used to demonstrate the visual impact a TDW can have. However, such demonstrations have hitherto failed to draw researchers into engaging with TDWs for displaying ultra-high resolution images as part of their mainstream research workflow. At this point, very few papers go beyond discussing the potential of TDW in astronomy, and none have been found where a specific research outcome has been achieved due to the use of a TDW. Some of the obstacles to adoption are outlined below.

### 4.2.1 Capital Expenditure

When the first TDWs became available, they were considerably expensive, costing many hundreds of thousands of dollars. However, with the continuing drop of the cost of commodity hardware, as well as the increasing functionality of components, the cost of an equivalent TDW today is a fraction of the cost of only ten years ago. The first TDW in Australia, 2008, an OptIPortal affectionately dubbed “OzIPortal” cost over half a million dollars, had 24 ultra-high resolution displays and required 12 high-end commodity computers to drive it. Today, ten years later, the 100 Megapixel TDW at Swinburne University

[4] `https://www.imax.com/`

[5] `https://www.necdisplay.com/success-stories/adler-planetarium/21`

of Technology is driven using 6 computers, and cost around a tenth of the price.

### 4.2.2 Ease of Use

As with most new technologies, TDWs suffered from the inevitable problems associated with immature systems. The software was unstable and the interfaces were unintuitive. Content required specific preparation and the complex process discouraged researchers who did not wish to spend time learning it. Since then developers have sought to address these issues and produced far more stable software able to handle most common file types. The underlying management of the hardware has also greatly improved, significantly improving the overall stability of the environment on the interconnected hardware.

### 4.2.3 Accessibility

One aspect of TDW that is not easily overcome is the simple fact that they are large and typically require a dedicated space, with appropriate power, network and air-conditioning needs. In most research institutions, space is at a premium and so suitable spaces for TDWs often compromise accessibility for general research use. Locating a TDW within a specific department or faculty can restrict use by other researchers.

### 4.2.4 Screen Limitations

TDWs are built by arranging many screens together to form a single, unified display space. However, when placed edge to edge, the individual screen bezels are combined with their neighbours' bezels to produce a double thickness bezel. This is often referred to as the "French window effect". Images displayed on the wall appear to be seen through a French window. This requires a user to decide if they want the bezels to occlude some of the data, making the image appear to sit behind the bezels, or to reveal all the data, making the image discontinuous at the bezels.

With larger bezels the effect is distracting, but can usually be accommodated after a short period, however this still puts users off. Recent developments have reduced the bezels in some screens to a combined thickness of only a few millimeters. Viewed from a comfortable distance, the bezels are easily overlooked in a TDW using these screens. A more expensive option is to use screens with essentially no bezel; however, such screens currently have a comparatively low pixel density and have therefore been largely restricted to large-scale advertising in airports and shopping centres.

Screen orientation also has the potential to impact on the effective use of the TDW. With all screens aligned to produce a flat display surface, there is no true "sweet spot" for an observer to stand. Screens directly in front of the observer are best, while screens to the left or right of this position will increasingly be viewed off axis. This is not a problem exclusive to TDWs as any large display will exhibit a similar effect, and modern digital displays allow for excellent off-axis viewing individually, in terms of intensity.

However, as the observer approaches the TDW, the effect of compression of the image at the periphery becomes more pronounced. One way to mitigate this effect is to curve the wall, angling the screens in toward the notional "sweet spot", where an observer can see all screens equally. In fact, this would also require not only the columns of displays to be rotated, but also the rows. In a regular grid of displays, this is impractical, though this approach has been used with some TDWs.

For most TDWs however, rotating the screens to place the observer at the focus reduces the usability of the TDW for multiple users. A gentle curve of the columns produces a pleasing visual effect whilst still allowing for excellent sharing of the environment by multiple users. This approach has the added benefit of increasing the stability of free-standing TDWs, but does require additional space and infrastructure to achieve.

### 4.2.5 Connectivity

Collaboration has become a key component to almost all research endeavors. The ubiquity of the Internet has made this possible. Researchers need to be able to share and discuss aspects of data in a timely manner. Quite often collaborators are not in the same (physical) location. A promise of the original OptIPortal was to connect the TDWs together so that researchers could work on the same data at the same time. Unfortunately, software and bandwidth limitations made connecting TDWs together difficult. More recently this obstacle has now been largely overcome.

### 4.2.6 Research Drivers

Research infrastructure requires sufficient demand from researchers to justify its continued existence, and numerous OptIPortals were built without research drivers to warrant them. Very few researchers produced ultra-high resolution images, and this was the primary target of the original promotion of TDWs.

Researchers have managed without TDWs for some time and generally accept that their current approach, while possibly limited, is adequate for the vast majority of their work. The occasions that are clearly requiring the significant infrastructure of a TDW are generally thought to be too few and far between to justify the investment in such technology. However, there are still find companies such as Toshiba, Sony and Panasonic producing higher and higher resolution screens, and professionals, academics and consumers are increasingly adopting the new technology.

Astronomy is already being impacted by the "data tsunami" (Berriman & Groom, 2011) and the next few years promise to see this problem expand dramatically. Huge investment has been made in the next generation of telescopes, cameras and detectors and with such a financial commitment, it seems prudent to make every effort to ensure the captured data can be adequately inspected, and to minimise the possibility of missed discoveries.

## 4.3 Are Tiled Display Walls Needed for Astronomy?

The following paper demonstrates the potential for TDWs for inspection of Gigapixel resolution images and the collaborative inspection of these images. It presents a use case where the image resolution exceeds a standard desktop display.

We assert that recent progress in technological capability and price of commodity hardware along with the stability of tiled display management software, has largely overcome many of the objections to the earlier incarnations of TDWs. In light of the heavy investment in astronomical equipment in recent times, it is necessary that appropriate effort also go into ensuring the full appreciation of the captured data is achieved. A TDW has become a cost-effective way of accomplishing this.

This Chapter comprises content published in the paper "Are Tiled Display Walls Needed for Astronomy?" by Meade, B., Fluke, C., Manos, S., & Sinnott, R. (2014). Publications of the Astronomical Society of Australia, 31. doi:10.1017/pasa.2014.29. Reprinted with permission of Publications of the Astronomical Society of Australia.

*Publications of the Astronomical Society of Australia* (PASA), Vol. 31, e033, 17 pages (2014).

doi:10.1017/pasa.2014.29

# Are Tiled Display Walls Needed for Astronomy?

Bernard F. Meade[1,2,3], Christopher J. Fluke[1], Steven Manos[2] and Richard O. Sinnott[2]

[1]Centre for Astrophysics and Supercomputing, Swinburne University of Technology, PO Box 218, Hawthorn, Australia, 3122
[2]The University of Melbourne, Parkville, Victoria, Australia, 3010
[3]Email: bmeade@unimelb.edu.au

(Received April 24, 2014; Accepted July 16, 2014)

**Abstract**

Clustering commodity displays into a Tiled Display Wall (TDW) provides a cost-effective way to create an extremely high resolution display, capable of approaching the image sizes now generated by modern astronomical instruments. Many research institutions have constructed TDWs on the basis that they will improve the scientific outcomes of astronomical imagery. We test this concept by presenting sample images to astronomers and non-astronomers using a standard desktop display (SDD) and a TDW. These samples include standard English words, wide field galaxy surveys and nebulae mosaics from the Hubble telescope. Our experiments show that TDWs provide a better environment than SDDs for searching for small targets in large images. They also show that astronomers tend to be better at searching images for targets than non-astronomers, both groups are generally better when employing physical navigation as opposed to virtual navigation, and that the combination of two non-astronomers using a TDW rivals the experience of a single astronomer. However, there is also a large distribution in aptitude amongst the participants and the nature of the content also plays a significant role in success.

Keywords: methods: data analysis – techniques: image processing

## 1 INTRODUCTION

Astronomy produces some of the largest volumes of scientific data. Future facilities such as the Large Synoptic Survey Telescope (Tyson 2002; Ivezic et al. 2008) and the Square Kilometer Array (SKA)[1] will produce final datasets heading toward, or even beyond, exabyte sizes.

In this ‘big data’ era of astronomy, existing data analysis tools and methodologies, where the astronomer works directly on the data at the desktop, will be pushed to their limits. There will be an ever-increasing reliance on automated processes to identify objects of interest. This includes the growing variety of generic approaches referred to as data mining (Ball & Brunner 2010; Brescia et al. 2012; Way et al. 2012), and discipline specific solutions such as automated source finders [e.g. see (Koribalski 2012) for a recent review of HI source finding strategies].

As valuable as automatic analyses of these enormous datasets are, astronomy still relies heavily on visual inspection. As the sensitivity of telescopes and detectors is improved, phenomena are increasingly being revealed at the boundary between the signal and the noise. In many cases, these phenomena are not even predicted, making automatic analysis meaningless. It is often a case of not knowing what you are looking for until you see it (Hassan & Fluke 2011).

Not only is the total volume of astronomy data increasing, but the size of individual images (and data cubes) is growing as well. For example, one of the highest resolution cameras currently available is the Dark Energy Camera (DECam), part of the Dark Energy Survey (DarkEnergySurveyWeb 2012). This camera uses an array of $62 \times 2048 \times 4096$ CCDs to form a 520 Megapixel image (Mohr et al. 2012). However, as Table 1 demonstrates, there is a growing divide between the resolution of images that can be recorded and the resolution of images that can be displayed on the desktop.

When exploring astronomical imagery, it is desirable to display an image at its native resolution, where there is a one-to-one correspondence between image and display pixels. A very large display with low resolution may reveal less information than a smaller display with a higher resolution. Such high-resolution images reveal more than just the detail of individual celestial objects. In fact, it is the combination of detail and context that make these images valuable: understanding the environment is critical to describing the

[1] http://www.skatelescope.org

**Table 1.** Comparison between the typical displays available to an astronomer, and the resolution of some of the current and proposed astronomical cameras. MPs = Megapixels.

| Capture device | Resolution | MPs | Reference |
|---|---|---|---|
| HST Advanced Camera for Surveys | $2 \times 2048 \times 4096$ | 16 | (ACSweb 2012) |
| Skymapper | $32 \times 2048 \times 4096$ | 268 | (Keller et al. 2010) |
| DECam | $62 \times 2048 \times 4096$ | 520 | (Mohr et al. 2012) |
| Subaru Hyper Suprime-Cam | $104 \times 2048 \times 4096$ | 870 | (HyperSuprimeCamWeb 2012) |
| **Display device** | **Resolution** | **MPs** | |
| Standard desktop display | $1680 \times 1050$ | 1.7 | |
| Full high-definition (FHD) desktop display | $1920 \times 1200$ | 2.3 | |
| iPad (with retina display) | $2048 \times 1536$ | 3.1 | |
| Dell UltraSharp desktop display | $2560 \times 1600$ | 4.1 | |
| Laptop (Macbook Pro) | $2880 \times 1800$ | 5.2 | |
| 4K ultra-high definition (UHD) display | $3840 \times 2160$ | 8.3 | |

phenomenon itself. When the image dimensions exceed the capabilities of a standard desktop display, then it is time to look to a non-standard display such as a tiled display wall.

### 1.1 Tiled display walls

A tiled display wall is an ultra-high-resolution display comprising a two-dimensional matrix of lower resolution display components, typically standard flat-screen monitors. While there are some slight differences in the way specific tiled display walls are assembled and configured, e.g. Hiperwalls (HiperwallWeb 2012), Powerwalls (LeedsPowerwallWeb 2009), or OptIPortals (see Appendix A.1), they still operate in a similar manner and hereafter are described simply as TDWs.

A key element of the design principle of TDWs is the use of commodity computers and displays. The computing power available in a standard desktop device with a typical graphics card is capable of driving stunning graphics across multiple displays at very high frame rates. Similarly, the expansion of capabilities of devices such as the emergence of multi-head graphics cards and additional expansion slots on motherboards, means that a single computer can now drive many displays. In fact, a modern computer containing a motherboard with three PCI-Express slots, each hosting a dual-head graphics card, with six Matrox TripleHead2Go[2]devices on each output, can drive 18 full high definition (FHD) displays.

While most TDWs are designed and built as flat screens, either free-standing or mounted on a wall, the use of individual display elements provides a great deal of flexibility in the geometrical configuration. The Mechdyne CAVE2 systems at the Electronics Visualisation Lab (EVL) (University of Chicago) (Febretti et al. 2013) and Monash University, Australia, wrap the TDW around the user, providing an extremely high-resolution immersive stereoscopic environment (74 Megaixels in 2D or 37 Megapixels in 3D). Two key advantages of using monitors over large-screen rear-projection, the usually approach for Cave Automatic Virtual Environments (CAVE; Cruz-Neira et al. 1992), is the increase in both the available pixels and the display brightness. A third advantage is the great reduction in physical footprint of the facility compared to the CAVE, which requires extra space outside of the walls to house the data projectors. The trade-off is a more complex computing and network back-end to drive $\sim 80$ individual panels, rather than the (maximum) six walls of a cubic CAVE. Additionally, there is the visible presence of screen bezels - the frame around each of the display elements.

While an ideal TDW would provide a seamless image, in reality the screen bezels introduce a windowing effect. Bezels can be distracting for certain types of content (e.g. office applications), whereas for other tasks they can actually provide a natural coordinate grid to aid in exploration (see Section 3.4). The display panels themselves continue to improve, including the appearance of screens with very thin bezels, such as the Christie Digital FHD551-X [3]with only 5.5mm combined bezel width.

We distinguish between resolution and pixel density when displaying images. For example, the first release of the Retina display for the 15 inch Apple Macbook Pro[4] had a resolution of $2880 \times 1800$, which greatly exceeds the typical resolution of a FHD home theatre display at $1920 \times 1080$. However, the home theatre's 2 megapixel display can extend over 150 inches (measured diagonally), while the Macbook Pro display crams its 5.1 megapixels into a 15 inch screen. The pixel densities of each configuration are at the extremes, with the Macbook Pro Retina providing a practically seamless image, while the FHD image projected to 200 inches would reveal the individual pixels quite clearly. At this time, there remains a significant price jump to move from FHD to the next off-the-shelf resolution of $2560 \times 1600$ pixels (e.g. Dell Ultrasharp). However, the recent emergence of commercially available 4K systems and the Retina displays from companies like Panasonic and Apple, will likely drive down the price of the $2560 \times 1600$ displays.

[2] http://www.matrox.com/graphics/en/products/gxm/

[3] http://www.christiedigital.com/en-us/digital-signage/products/lcd-flat-panels/pages/55-hd-lcd-flat-panel.aspx

[4] http://www.apple.com/au/macbook-pro/features-retina/,2013

A number of applications exist to simplify the management of a TDW. The two main contenders are Scalable Adaptive Graphics Environment (SAGE; SAGEWeb 2013) from the University of Chicago's EVL and the Cross-platform Cluster Graphics Library (CGLX; CGLXWeb 2012; Ponto et al. 2011) from UC San Diego's CALIT2, though several other solutions also exist. One of the principle benefits of SAGE is that it makes sharing content between TDWs easy (Fujiwara et al. 2011). More recently, Tada et al. (2011) have developed a visualization adaptor that extends the capability of SAGE to allow the display of any X-Window, which opens up the possibility of using almost any application on the TDW.

### 1.2 Background

Astronomical imagery is often seen on promotional material for TDWs, such as the SAGE gallery images[5]. Indeed, the very high resolution images captured by modern detectors (Table 1) do seem very well suited to the environment. TDWs have been used successfully as public outreach devices, such as the displays at the National Institute of Information and Communications Technology, Japan (Morikawa et al. 2010; NICTWeb 2012), and the Adler Planetarium, USA (Adlerweb 2012). However, there is a paucity of literature examining whether TDWs actually do improve understanding of any ultra-high resolution image in any scientific discipline.

Ball & North (2005a) conducted some of the first experiments to compare individual computer displays with TDWs. They tested subjects on a single desktop display (17 inch, $1280 \times 1024$ pixels), a $2 \times 2$-display TDW ($2560 \times 2048$ pixels) and a $3 \times 3$-display TDW ($3840 \times 3072$ pixels). For a target search (red shapes on a black background with random grey dots), they found that participants performed far better when searching for small targets when they could see all the targets at once. The first part of the experiment involved the participants searching for a specific configuration of red dots, while the second part required the participants to match pairs of configurations. No statistically significant difference was observed for targets that could be easily seen on all display environments without needing to pan or zoom.

Ball & North (2005a) also found that the experience of virtual navigation, that is, using a mouse to zoom and pan on a single display, caused more frustration for the participants than the physical navigation required for the TDW. Here, physical navigation means observing parts of an image by physically moving the eyes, head or whole body to an optimum position. The more virtual navigation that was required, the greater the disorientation and agitation experienced by the participants. The authors suggest that being able to easily maintain context while searching for detail made the TDW a more acceptable experience. They also found that the physical construction of the TDW with the screen bezels dividing the image into segments aided the search process. We comment on this issue in Section 3.4.

Yost, Haciahmetoglu, & North (2007) studied the visual acuity of human perception with regards to high-resolution displays. In this context, visual acuity refers to the ability to perceive all displayed information: when an entire display surface is within the user's field of view, and the individual pixels remain discernible, the display is said to be within visual acuity. While increasing the pixel density is one way to exceed visual acuity, increasing both the pixel count and the display size is another. When a display exceeds visual acuity, there are always pixels that cannot be accurately perceived without physical navigation. However, increasing physical navigation does not negatively impact on the performance on most tasks, whereas virtual navigation of the same data or image does have a significant negative impact. This is contrary to the notion that there is no value in using a display that exceeds visual acuity. For practical reasons, it is harder to establish at what point this advantage disappears.

Andrews et al. (2011) focused their investigation on the human experience of using high-resolution displays. In this study, they defined large displays as being human-scale, that is, where the physical size of the display was of similar height and width to naturally occupy the natural field of view of an adult human. In this definition of a large, high-resolution display, this can be achieved through tiling of displays or using individual displays with greater pixel count and large physical size. The authors argue that the design of such displays as TDWs would greatly benefit from considering the physical nature of human-centric search techniques and creating displays that meet these needs.

Andrews et al. (2011) also considered the natural perceptions of users and how a display can affect these perceptions. There is a potential for TDWs to overwhelm the user with information, or induce physical fatigue due to the increased requirements of physical navigation of the environment. However other studies have shown that physical navigation can outperform virtual navigation for many tasks (Ball, North, & Bowman 2007), and that users quickly adjust to the information density shown on a large display (Andrews et al. 2010). The increased physical activity required for a large display has not shown significant increase in fatigue of subjects, though there is a possibility of some discomfort in the neck due to the increased turning of the head (Ball & North 2005b; Bi & Balakrishnan 2009).

Following on from these design questions, Bezerianos & Isenberg (2012) conducted experiments to determine how proximity to a TDW could affect the perceptions of the user, particularly focusing on angular distortion effects. The ability of participants to effectively estimate quantities such as angle, area and length of objects within an image were significantly affected when the angle of presentation was increased. Thus when a participant was very close to the TDW, their ability to estimate these quantities diminished as the distance of the object to the subject increased. Of particular interest are the results from the second experiment where some of the participants were required to remain in a fixed location, while others were allowed to move freely. The study found

[5] http://http://www.sagecommons.org/community/sage-walls/,2013


**Figure 1.** The 29566 × 14321 pixel Carina Nebula mosaic from Hubblesite.org, with OzIPortal (15360 × 6400), Dell Ultrasharp (2560 × 1600) and Standard Desktop Display (1680 × 1050) sizes overlaid.

that the static position yielded just as accurate results of the mobile position, but was less time consuming. Therefore, the authors' recommendation that users be encouraged to remain at a distance from the TDW where possible, or physically inspect objects positioned close to their position, ties in closely with the Ball & North (2005a) observation that users naturally avoid using virtual navigation unless they absolutely have to (Ball, North, & Bowman 2007).

Most of these studies focus on generalized examples of use of TDW, but the nature of the research disciplines also needs to be considered when investigating these displays. For example, the way an astronomer would use a TDW could have significant differences to the way an economist would use it. As Moreland (2012) argues, we already know how to build the displays, but we have little experience in considering domain-specific applications. The desirability of achieving a one-to-one correspondence between image and display pixels aside, it is far too simplistic to suggest that images of A x B pixels require displays of equivalent resolution. Instead, the need must be borne out of the research and the data, where the impact of virtual navigation impedes comprehension.

### 1.3 Overview

In this paper, we describe a series of experiments designed to investigate the assumption that TDWs are intrinsically beneficial in astronomical research. We focus our attention on targeted searches within high resolution images that exceed the available resolution of a standard desktop display. We consider the performance of both individuals and pairs of users at finding targets of decreasing size on either a standard desktop display or a TDW. The participants in the experiments included professional astronomers, experienced amateur astronomers and non-astronomers.

The TDW used in these experiments, the OzIPortal, was built by the School of Engineering at the University of Melbourne in 2008 and is now operated by the University's central IT department. The TDW comprises a $6 \times 4$ matrix of Dell Ultrasharp monitors ($2560 \times 1600$). With a total resolution of $15360 \times 6400$ pixels, it is capable of displaying 98.3 Megapixels. However, as Figure 1 shows, this is less than a third of the pixels available in images such as Hubble's Carina Nebula mosaic[6].

The OzIPortal initially used CGLX for the interface, but this was replaced with the somewhat more versatile SAGE software. We describe the history of the OzIPortal in Appendix A.

The remainder of this paper is set out as follows. In Section 2, we describe the OzIPortal experiments, including the image selections, participants and procedure. In Section 3, we show the experimental results. We look at the comparative performance of targeted searches using both standard desktop display and TDW environments. We consider the performance of the non-astronomer, astronomer and collaborative pair groups. We comment on key findings from

[6] http://imgsrc.hubblesite.org/hu/db/images/hs-2007-16-a-full_jpg.jpg

the post experiment survey and video observations. In Section 4 we discuss the implications of these results in the context of the potential use of TDWs in astronomy. We consider further experiments that are either extensions of the current work, or alternative aspects of using a TDW that might be beneficial to astronomers. Concluding remarks are made in Section 5.

## 2 THE OzIPortal EXPERIMENTS

In this section, we describe our experimental procedure to investigate the role TDWs might play in aiding knowledge discovery and comprehension of ultra-high resolution images (i.e. ~ 100 megapixels). These images provide researchers with an opportunity to seamlessly explore both context and detail at will. Yet on a standard desktop display (SDD), defined for our purposes as a 24 inch LCD with a resolution of 1680 × 1050 pixels,[7] a researcher must choose dynamically between context or detail, as both cannot be seen at once. In particular, we wanted to determine if there was indeed a definable performance improvement when using a TDW compared to a SDD, which corresponds to the popular expectation that big images need a big display to be seen "properly". The high-resolution images and target objects were chosen from three different categories: English words, galaxies and nebulae (see Section 2.2).

### 2.1 Participant selection

Participants were recruited from two different demographic categories: astronomers and non-astronomers. For the astronomers, participants included academics, postdoctoral researchers, research students and advanced amateurs. Within this group there was a mix of radio, optical and theoretical astronomers. Non-astronomer participants had a wide range of experience with astronomical imagery, ranging from none to a high level of familiarity. As such a secondary category of expert and non-expert was introduced, based on the participant's self-rated level of experience with astronomical imagery. Figure 2 shows that the non-astronomer cohort self-identified strongly with the low end of the experience scale, while the astronomer group is towards the high end. This self-rating reflects that, for example, a theoretical astronomer may not feel they have the same expertise as an optical astronomy who works constantly with images.

A total of 45 non-astronomers and 12 astronomers participated in a range of experiments. All participants had a reasonably high-level of familiarity with graphical user interfaces and the use of a mouse for panning and zooming, but few had any prior exposure to a TDW. We report here on the performance results of a subset of 30 participants, noting that:

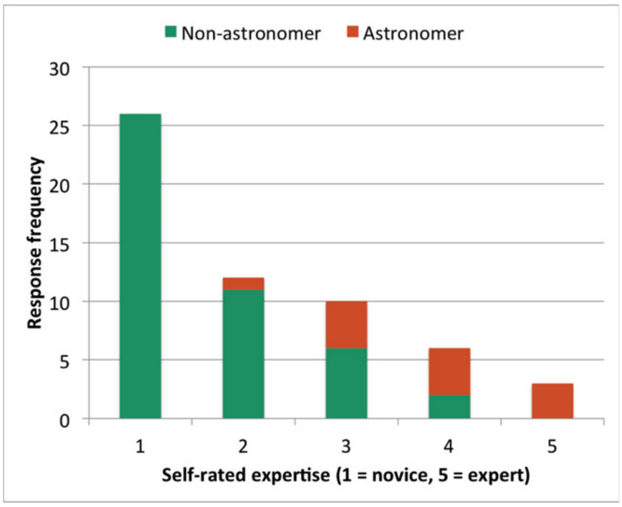


**Figure 2.** Survey results for self-rated level of expertise with astronomical imagery, for the astronomer and non-astronomer groups. The non-astronomer cohort (green) self-identified strongly with the low end of the experience scale, while the astronomer group (red) is towards the high end.

- The first five non-astronomer subjects participated in an experiment refinement phase and thus their performance results have been excluded from the results described below.
- 14 non-astronomers were presented with a slightly revised set of tasks to those described here. These additional tasks focused on a small target search and multiple image inspection. The small target search proved too difficult to complete in the SDD environment due to a "too-restrictive" time limit of two minutes, and too few participants were available to complete the multiple image inspection.

As all of these participants did complete the post-experiment survey, providing relevant comments on issues such as the suitability of the TDW for the target search task, we retained their survey responses for subsequent qualitative interpretation.

16 of the remaining non-astronomers completed a target search in pairs in order to investigate the process of collaborative inspection on SDDs and TDWs (see Section 2.4).

### 2.2 Image and target selection

In order to establish a common ground between the astronomer and non-astronomer groups, the first image was made up of black words on a white background at a resolution that precisely matched the TDW (15360 × 6400), as can be seen in Figure 3. At this resolution, all words were readable on the TDW without the need to zoom the image.

The words were taken from a list of the top 250 English words (AnglikWeb 2003) to ensure all participants were familiar with the targets. The words were rendered in Arial font and were sized in points of 1000, 300, 100, 30, and 10. Five

[7] The recommended size for centrally deployed computers at the University of Melbourne at the time of the experimental work

**Table 2.** Images used for the galaxy search and nebula search. Note that the Carina Nebula image was displayed at 50% of the native resolution for performance reasons

| Image ID | Field description | Resolution | Targets |
|---|---|---|---|
| Galaxy Set A | The Coma Cluster | 10816 × 7679 | Figure 4 |
| Galaxy Set B | CANDELS Ultra Deep Survey | 15516 × 8255 | Figure 5 |
| Nebula Set A | The Carina Nebula (NGC3772) | 29566 × 14321 @ 50% | Figure 6 |
| Nebula Set B | HST-Spitzer Composite of Galactic Center | 12203 × 4731 | Figure 7 |


**Figure 3.** The OzIPortal TDW with English word targets displayed at their native resolution. Arial font sizes used were 1000, 300, 100, 30, and 10 points. All words were visible on the TDW using physical navigation and no zooming.

targets were created at each size except the smallest where an extra five were added. For each size, an equivalent number of non-target words were added from the same list, to reduce the possibility of participants guessing based purely on size.

When viewed on the SDD, scaling the image to full screen reduced readability to words in 100 pt font or greater. For words 30 pt or 10 pt in size, zooming the image was necessary.

The astronomy targets were chosen to present a range of sizes similar to the word sizes described above, chosen from amongst the largest available on the HubbleSite gallery[8] - see Table 2 for details. For performance reasons, the Carina Nebula mosaic was shown at 50% of the native resolution.

From these images, the search targets were selected to roughly correspond to the physical sizes of the words, without following a strict sizing scale. The largest astronomy target was 2100 × 1730 pixels while the smallest target was 185 × 145 pixels. Images were not rotated, but were scaled to appear the same size on the search target presentation screen. Astronomical targets were chosen to reflect increasing difficulty. Due to the increased difficulty of the astronomical search compared to the word search, and the limited amount of time available for each participant to complete each task, the number of targets was restricted to 10 per image.

Targets selected from the galaxy images included structures around the galaxy. However, these targets exist on a black background and have no visible connectivity to the other objects in the image. Nebulae provide a fully connected structure with details visibly connected to the context. Figures 4 to 7 show each of the astronomical images and the targets. Additional galaxy and nebula images were used to introduce the environments but were not used during the experiment.

In order to eliminate any potential presentation bias, the image sets were shown alternating for the environments, so that half the participants saw the set A images on the SDD and set B images on the TDW, and vice-versa for the rest.

## 2.3 Procedure

Figure 8 shows the experimental set-up. The individual target objects were presented on a 40 inch LCD TV immediately adjacent to the TDW, as well as on a laptop sitting adjacent

[8] http://hubblesite.org/gallery/album/entire/hires/true/

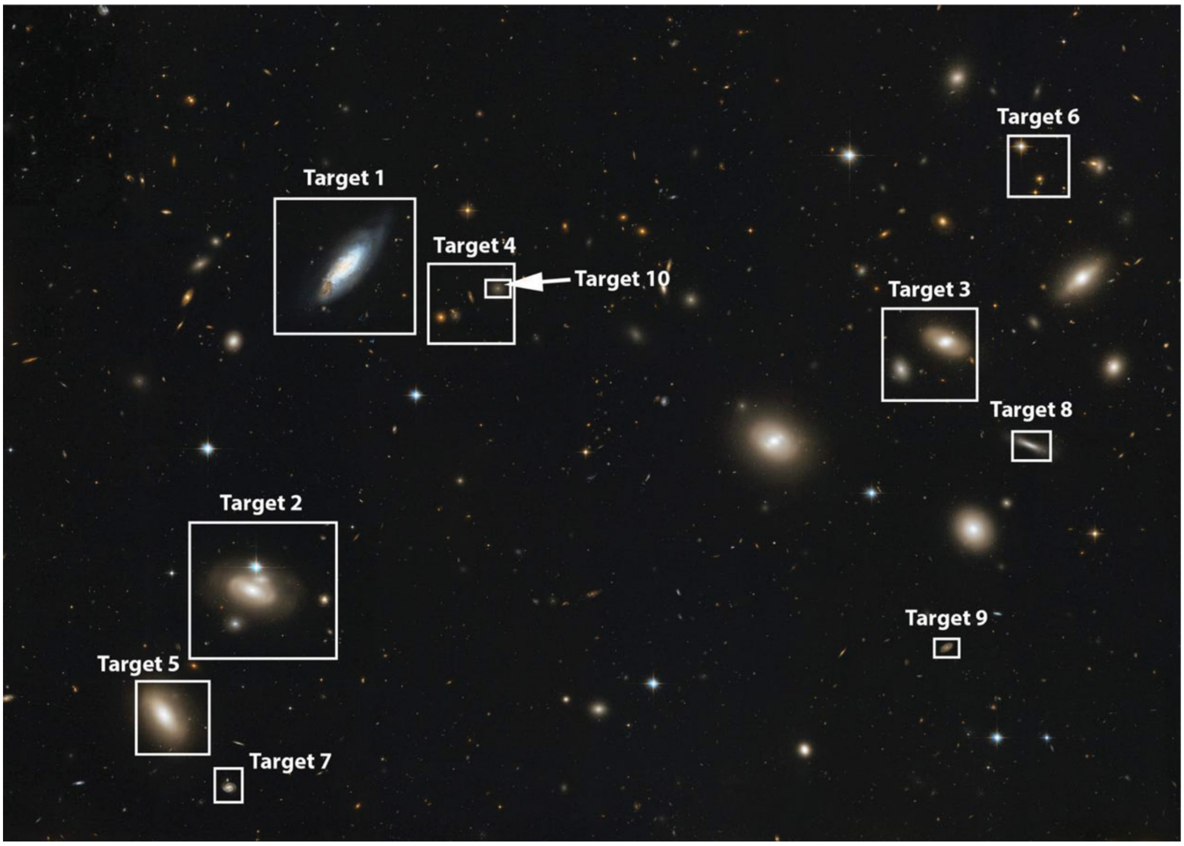


**Figure 4.** Galaxy Set A targets in the Coma Cluster (`http://hubblesite.org/newscenter/archive/releases/2008/24/image/a/`).

to the SDD. A standard Microsoft PowerPoint presentation was used to display the targets to the participant.

As the participant identified the target (or elected to pass on finding a given target), the presentation was advanced to the next target. Participants were given a total of 2 minutes to find as many of the targets as they could. Once the experiment had been completed on the SDD, the participant was then shown a new set of images and targets on the TDW, again given 2 minutes for each set. Several of the experiments were also filmed for later investigation as to how the displays were used.

Participants were introduced to the experiments, with a brief explanation of their purpose and a demonstration of how to use the two types of displays. The SDD was a familiar environment for all participants and very little introduction to the environment was required. In the case of the SDD, participants were advised that they would be able to find most of the large targets without using the mouse to pan and zoom, but would need to use virtual navigation for the very small targets. The mouse operation was already second nature, though most participants attempted to minimize the mouse use, preferring to lean closer to the screen.

Very few of the participants had ever seen a TDW before and so the experience was entirely new to them. Those that had encountered such a display before showed little if any advantage when engaged in the structured search experiment. The only significant advantage pre-exposure was that ability to “zoom” by physically approaching the TDW was already known.

In the initial experimental refinement phase, a test group of five non-astronomers was given the same introduction to the TDW as they were to the SDD. The result was that these participants all felt obliged to use the TDW in exactly the same way they had used the SDD, i.e. they sat well back from the screen to obtain the same field of view and used a mouse to zoom and pan rather than walk closer to the screen.

Due to the nature of the TDW software, zooming and panning resulted in some slight image tearing as the screen refresh was not always perfectly synchronized. Moreover, the zoom was not visually active with the image jumping between zoom levels rather than scaling dynamically, as participants were used to on their SDD.

While these issues can be mitigated with higher networking speeds, a simple alternative was found: the participants were told that standing and approaching the screen would more effectively function as zoom (i.e. physical navigation). This very simple training was included in the familiarization stage for the later participants. The use of physical navigation greatly improved the user satisfaction and performance with the TDW, and presented a more

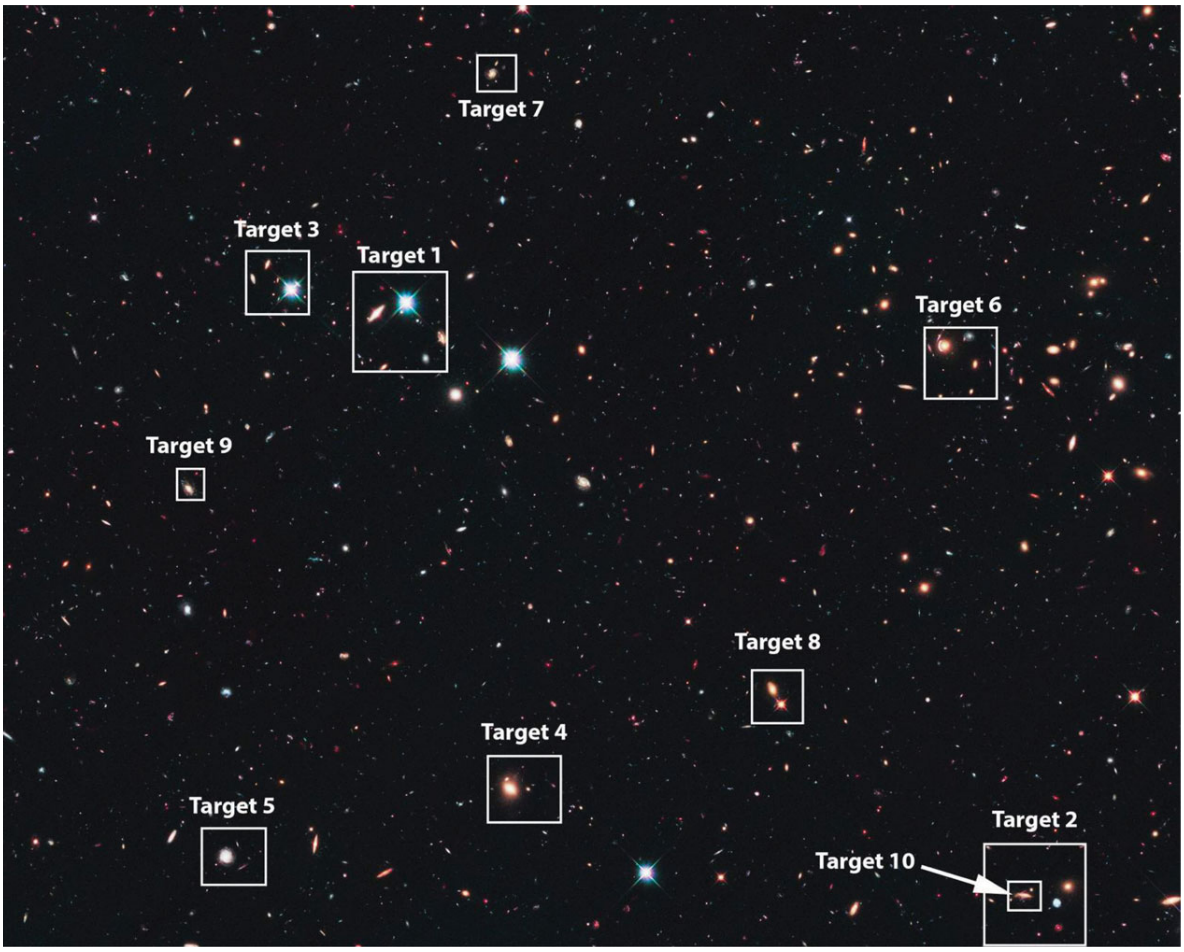


**Figure 5.** Galaxy Set B targets in the CANDELS Ultra Deep Survey (`http://hubblesite.org/gallery/album/entire/pr2013011b/hires/true/`).

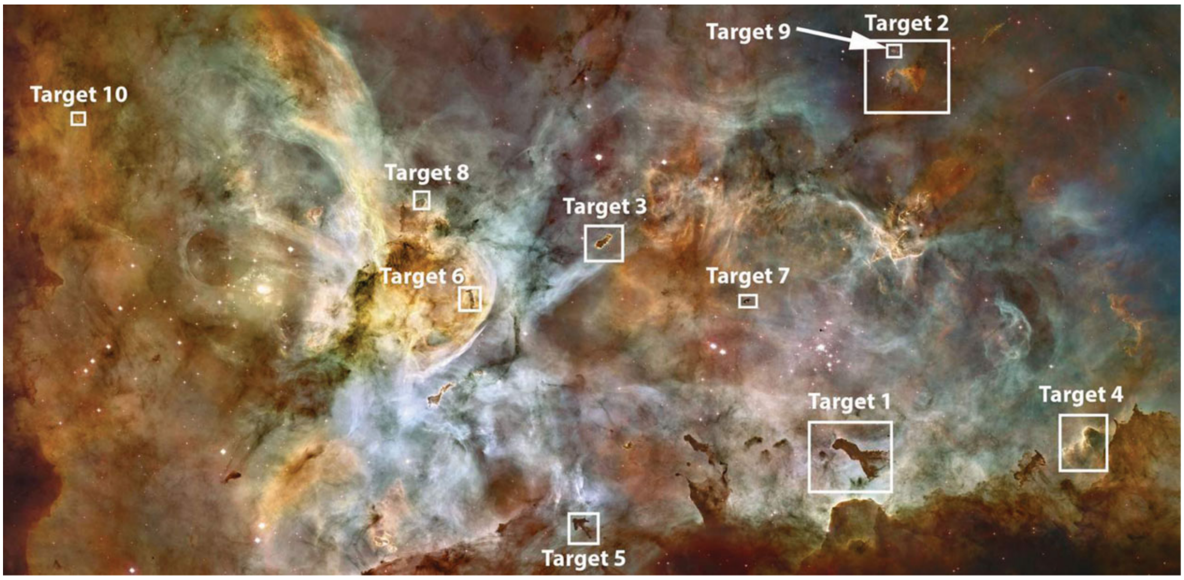


**Figure 6.** Nebula Set A targets in the Carina Nebula (`http://hubblesite.org/gallery/album/nebula/pr2007016a/hires/true/`).

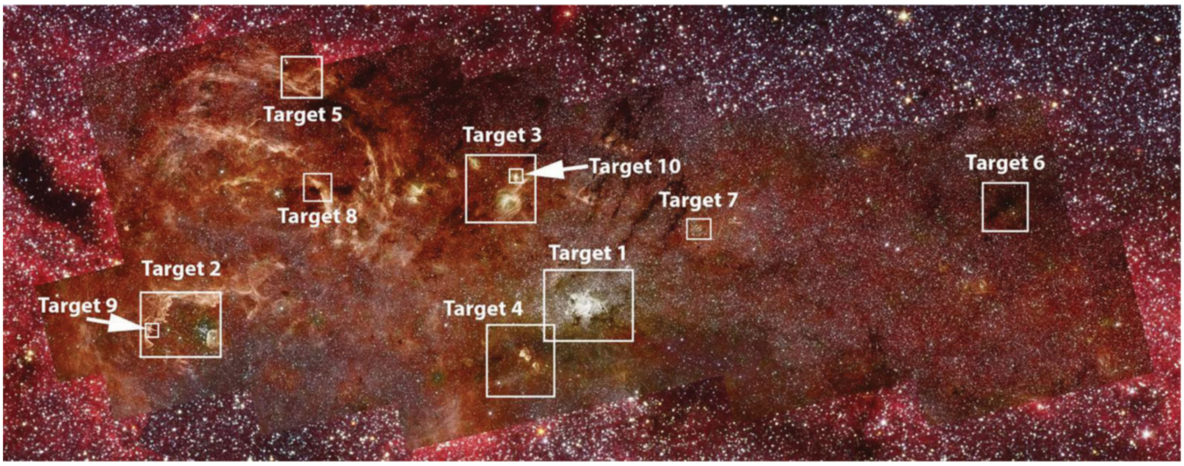


**Figure 7.** Nebula Set B targets in the HST-Spitzer Galactic Center composite (`http://hubblesite.org/newscenter/archive/releases/2009/02/image/d/`).

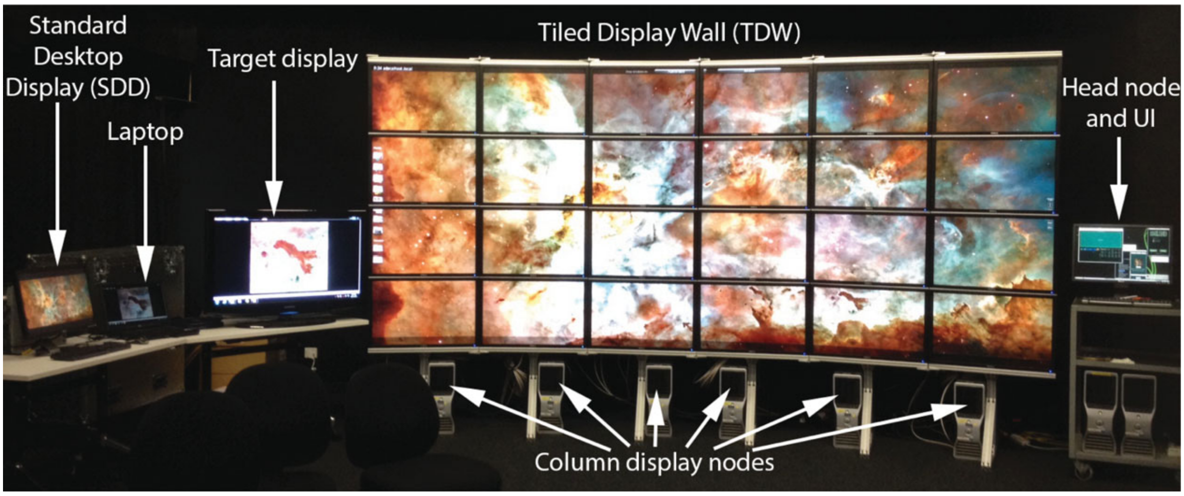


**Figure 8.** Experiment layout as described in Section 2.3. The individual target objects were presented on a 40 inch LCD TV (Target Display) immediately adjacent to the TDW, as well as on a laptop sitting adjacent to the SDD (Laptop). A standard Microsoft PowerPoint presentation was used to display the targets to the participant. The six columns of the OzIPortal are driven by six column display nodes, with master control under the SAGE environment from the Head node [Image: Carina nebula mosaic from http://www.hubblesite.org].

realistic assessment as to how the displays should be used in practice.

For the non-astronomer singles group, all participants began the target search using the SDD. For the astronomer group and the collaborative groups, we tested several participants with the TDW display first to see if there was any advantage to the order of exposure.

As Figure 10 shows, there is no significant difference in performance on the target searches regardless of the order of the environments used. The galaxy and nebula targets were alternated for the two environments in order to ensure that no advantage could be ascribed to a particular image/environment combination. After completing the experiment in both environments, the participants were then asked to complete a survey about their experience.

## 2.4 Collaborative pairs

16 non-astronomers were asked to complete the experiment in pairs. They received the same introduction as all the other participants, but no specific instruction was given to guide how they should share the task. They were required to determine the best way to operate between themselves as part of the task, and in all cases settled the matter of who would operate the interface (in the case of the SDD) or how they would split the search area (in the case of the TDW), with a very brief discussion. This process was occasionally completed during the introduction and so took no time during the task, however, in all cases it did not delay the image inspection process as the pairs began searching while discussing. No disadvantage was observed and therefore no adjustment

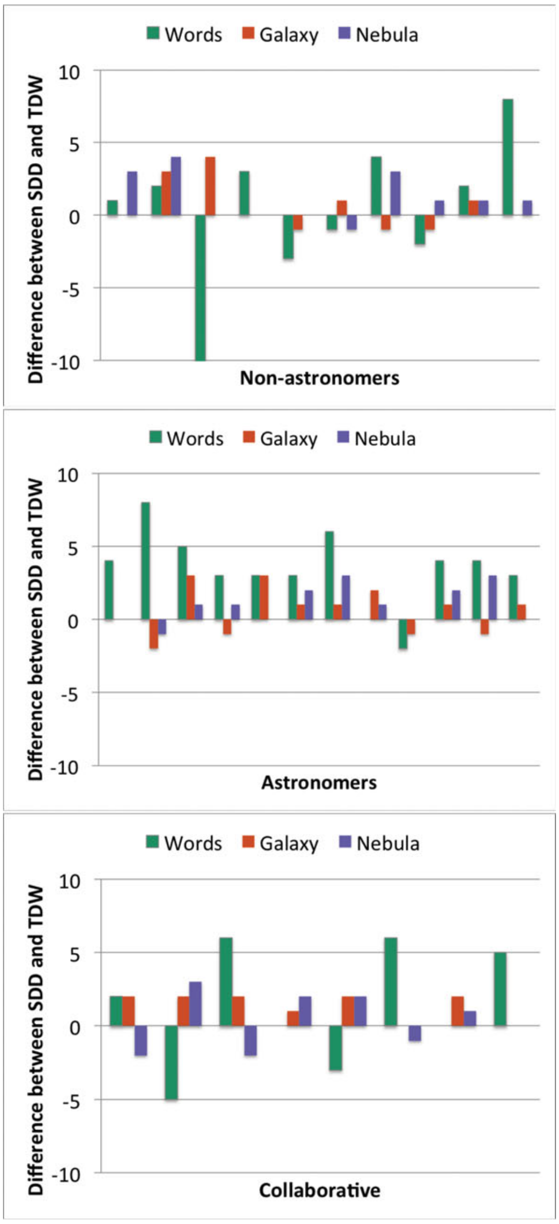


**Figure 9.** Comparison of the success rates of the TDW and SDD for 10 non-astronomers (top panel), 12 astronomers (middle panel), and eight collaborative pairs of non-astronomers (bottom panel). Positive numbers show success favoured the TDW, negative results favoured the SDD. The three types of experiments are the word search (green), galaxy search (red) and nebula search (purple).

has been made to the performance results. The only practical difference to the conditions of the experiment was that agreement was required from both participants for any ambiguous situation, for example, some targets could be mistaken for a similar looking non-target. This situation included when no target could be found, at which time both participants could agree to "pass".

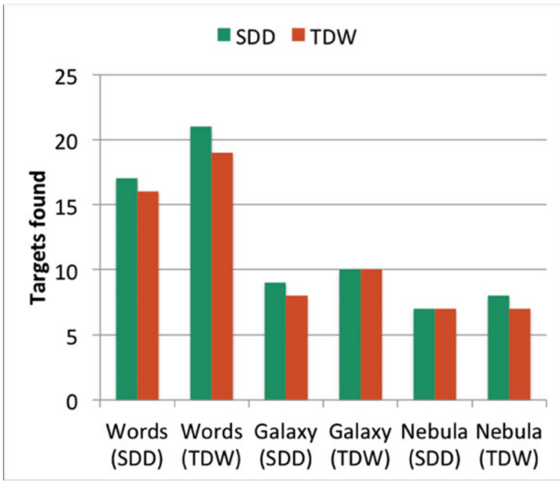


**Figure 10.** Results based on the presentation order of the display environments. Each pair of labelled bars indicates the image type and the first environment participants were exposed to (SDD or TDW in brackets). The green and red bars indicate the number of targets found using the SDD and TDW respectively. There is no strong dependence on which display technology that participants used first.

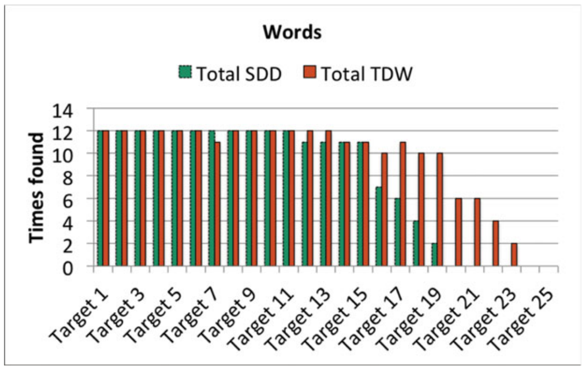


**Figure 11.** Success rate based on individual words. The green columns indicate the target words actually found using the SDD and the red columns are for the TDW. Only the results from the astronomer group were used, in order to align with results shown in Figure 12 as the image context was not recorded for the non-astronomer group.

## 3 RESULTS

The results of the experiments are presented in the next two sections: Section 3.1 shows the empirical results for the different image identification tasks and 3.2 contains our analysis of the post-experiment survey. We report also on the outcomes of the video observations (Section 3.3), which provide some valuable clues as to how to make better use of TDWs. In addition, we look at the specific feedback made with regards to the TDW bezels (Section 3.4).

### 3.1 Success rates

Figures 11 and 12 show the individual targets successfully identified. Figure 11 shows the first 15 target words were


doi:10.1017/pasa.2014.29

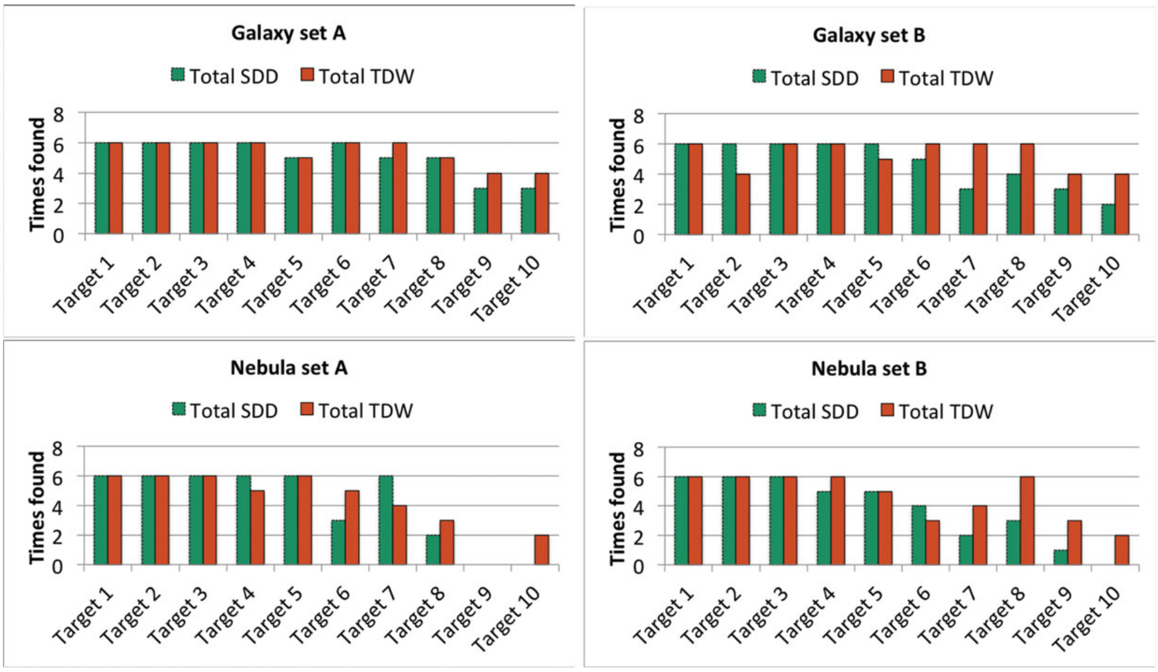


**Figure 12.** Success rates based on individual search targets with galaxy and nebula images. The green columns indicate the targets actually found using the SDD and the red columns are for the TDW. Only the results from the astronomer group were used as the image context was not recorded for the non-astronomer group.

found by most participants. This corresponds to point sizes of 1000, 300 and 100, which were easily readable in both environments. At point size 30, the words were no longer readable on the SDD and therefore virtual navigation was necessary, causing a performance decline. This can be seen by the rapid drop in the SDD success. Very few of the 10 point words were found in the SDD environment. However, the TDW success rate shows only a slight decline for the 30 point words and a steady decline for the 10 point words.

Figure 12 shows that the galaxy images in set A were fairly well matched for success in both environments, while set B showed a higher success rate for the TDW targets. Similarly for the nebula sets, with the exception of target 9 in the Nebula Set A, which was not found in either environment. As can be seen in Figure 6, target 9 was a subset of target 2, but not particularly more difficult than target 9 of Figure 7, also within target 2 for that set.

Table 3 shows the combined success rate for each group in each environment, where success refers to the number of targets identified during the test. These results indicate that generally performance on the TDW is slightly better than for the SDD for the same set of tasks, with the notable exception that self-rated experts actually performed slightly worse on the TDW for the Galaxy search. However there are other factors to consider. For example, the sample size is fairly small and the task is not necessarily indicative of typical astronomy work.

Table 3 also shows that the attempt to establish a consistent baseline between the cohorts was effective. In the word search on a SDD, where little experiential value could be ascribed, all groups achieved very similar results. Here the targets and the navigation method were familiar to all subjects. However, the non-astronomers did not experience an improvement in performance when searching for words on the TDW. This reason for this is uncertain, but could be because the astronomers' familiarity with exploring large images translated more easily to the TDW environment. Video observations reveal that astronomers tended to adopt methodical search strategies and were quicker to adapt their strategies to the TDW environment than non-astronomers. See Section 3.3 for more information on video observations.

A useful way to view these results is to consider the comparison of the results for the specific environments. Subtracting the results for the SDD from the TDW produces a simple comparison of the two environments, as seen in Figure 9 for the 10 non-astronomers (top panel), 12 astronomers (middle panel) and eight collaborative pairs of non-astronomers (bottom panel). The astronomer group test results are comparable to the non-astronomer cohort. The astronomers did demonstrate slightly better performance overall, particularly with the word and nebula search. However, the galaxy search results show that the astronomers tended to perform equally well on both the SDD and the TDW (c.f. Table 3). Observations supported by the video recordings show that the astronomers tended to have a more systematic approach to searching, and were less confused by targets split by the screen bezels (edges).

**Table 3.** Search success rates for non-astronomers (10 participants), astronomers (12 participants), non-astronomer collaborations (8 × 2 participants) and a combination of all three (10 + 12 + 8 = 30 sets of results). These results are based on the median values for the word, galaxy and nebula feature searches, with a quoted range of one standard deviation.

| | Group | $N_P$ | SDD | TDW |
|---|---|---|---|---|
| | Non-Astronomer | 10 | 66% ± 2.0 | 66% ± 4.4 |
| | Astronomer | 12 | 66% ± 1.8 | 80% ± 2.6 |
| **Word search** | Non-Expert (self-rated) | 15 | 68% ± 1.7 | 72% ± 4.3 |
| | Expert (self-rated) | 7 | 64% ± 2.3 | 76% ± 3.0 |
| | Collaboration | 8 groups | 66% ± 3.0 | 76% ± 5.4 |
| | Non-Astronomer | 10 | 60% ± 1.7 | 75% ± 2.5 |
| | Astronomer | 12 | 90% ± 1.9 | 90% ± 1.3 |
| **Galaxy search** | Non-Expert (self-rated) | 15 | 70% ± 1.8 | 80% ± 2.4 |
| | Expert (self-rated) | 7 | 90% ± 2.1 | 80% ± 1.3 |
| | Collaboration | 8 groups | 80% ± 0.9 | 100% ± 0.0 |
| | Non-Astronomer | 10 | 50% ± 2.0 | 60% ± 1.3 |
| | Astronomer | 12 | 65% ± 1.2 | 80% ± 1.6 |
| **Nebula search** | Non-Expert (self-rated) | 15 | 60% ± 1.9 | 60% ± 1.9 |
| | Expert (self-rated) | 7 | 70% ± 1.3 | 80% ± 1.3 |
| | Collaboration | 8 groups | 70% ± 1.2 | 80% ± 1.6 |
| | Non-Astronomer | 10 | 61% ± 3.3 | 67% ± 5.2 |
| | Astronomer | 12 | 71% ± 2.8 | 82% ± 3.3 |
| **Combined** | Non-Expert (self-rated) | 15 | 67% ± 3.1 | 71% ± 5.3 |
| | Expert (self-rated) | 7 | 71% ± 3.4 | 78% ± 3.5 |
| | Collaboration | 8 groups | 70% ± 3.4 | 82% ± 5.6 |

The TDW has often been cited as an ideal environment for research collaboration. The bottom panel of Figure 9 shows the results obtained by pairing two non-astronomers for the same task. Table 3 shows that non-astronomer collaborators match or exceed the performance of a single astronomer and show marked improvement of TDW over SDD.

Figure 9 shows an interesting anomaly with one participant finding the SDD to be far better for the word search than the TDW. In this case, the participant overlooked a 1000pt word and began to search among the smaller words. This highlights a potential trap with a TDW in that a large object presented on such a display may be too large to see, with participant approaching the TDW and effectively eliminating their chance of recognizing the word. This may be in part due to the way the brain recognizes words as a whole and therefore may not apply to astronomical targets.

## 3.2 Survey responses

The participants were asked to complete a short survey after the experiment, designed to gauge their experience using the two display environments. Participants were asked to rate their own experience with astronomical imagery, as can be seen in Figure 2.

Figure 13 show the survey results for the Ease of Use of the SDD and the TDW respectively. Results for both astronomers and non-astronomers indicate that the SDD is generally perceived as difficult for this kind of search while the TDW is generally perceived as easy to use for the same.

Participants were also asked to rate the suitability of the two environments to the tasks presented. Figure 14 show that the results for the SDD are skewed toward Unsuitable for the search tasks while the participants found the TDW was generally well suited.

## 3.3 Video observations

Several of the participants were also filmed to record the manner in which they used the display environments.

When using the TDW, several participants found themselves overlooking extremely simply targets, particularly in the word search, by assuming that the target they were seeking must be smaller than it actually was. However, generally the approach to searching was fairly uniform, with the subjects standing back to get an overview of the image, and then approaching promising regions of the image. In the case of words, the advantage to the TDW over the SDD was that even the very smallest of words could be clearly read when close, while the same target on the SDD was not even visible when zoomed to full extents.

However, the experiment was designed to make the first targets easy to find and subsequent targets were made progressively harder. This gave a distinct advantage to the SDD for the early targets, and therefore considerably more time was available for finding smaller targets. On the TDW, however, the larger targets were sometimes overlooked, occasionally due to the splitting of the target by the screen bezels, or due to the participants' assumption that the target must be smaller than it actually was.

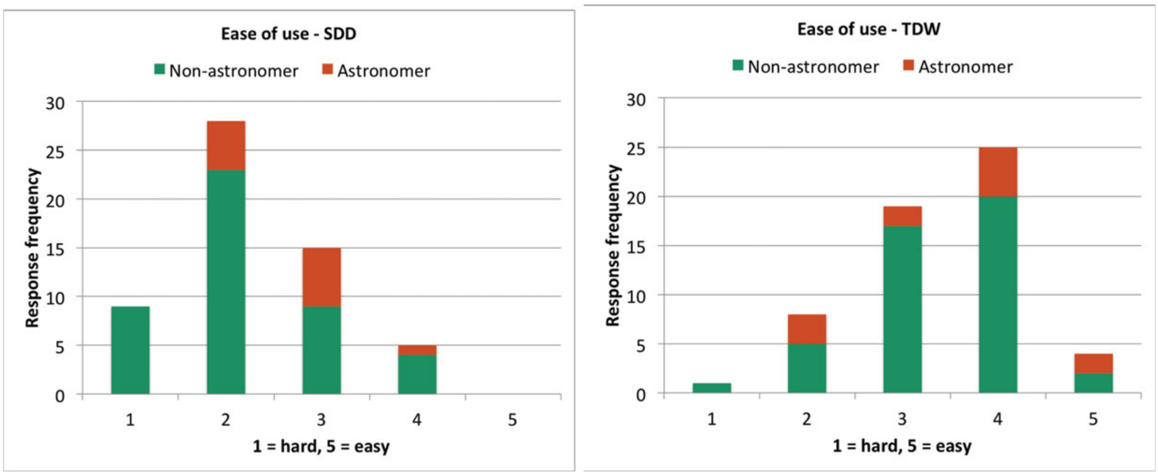


**Figure 13.** Ease of use of the display environments. Results were obtained from post-experiment surveys completed by all 57 participants.

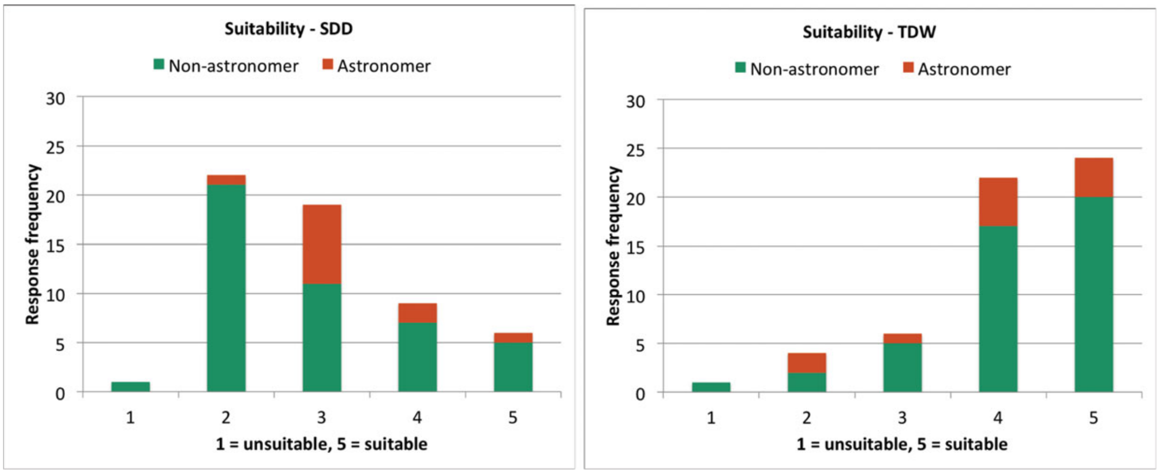


**Figure 14.** Suitability of the display environments for target searching in ultra-high resolution images. Results were obtained from post-experiment surveys completed by all 57 participants.

In general however, the video review shows participants were more methodical in searching for small targets on the TDW than on the SDD. It appears that the participants were more easily able to identify individual screens that they had already searched; compared to trying to remember which region they had searched of the image on the SDD. The screens of the TDW made for a simple segmentation that was easy to remember.

Furthermore, participants working with a partner found the TDW to be naturally separated into halves with each partner being responsible for their own half. Working together on the SDD, these participants found ceding control to someone else to be frustrating. Some chose to share the task of controlling the mouse, alternating between tasks, while others simply directed their partner by pointing in the direction they wished to explore. This resulted in some confusion, though some collaborations quickly settled into very effective teamwork.

On the TDW, splitting the screens into left and right did not always produce harmony. When one participant became convinced that the target was not on their side, they began to encroach on their partner's domain. For some, this resulted in an unspoken agreement to swap sides, while not so for others. However, the success of the collaborations between non-astronomers produced results that were generally better than a single astronomer (c.f. Table 3). Unfortunately there were not enough astronomers to test collaborative behaviours.

### 3.4 Feedback on bezels

Participants were also asked to comment on the structure of the TDW and how the "screen elements" impacted on the subjects' ability to complete the task. We defined screen elements as anything that interrupted the subjects' view of the search image. This included the SAGE icons and toolbars and the screen bezels. The results can be seen in Figure 15.

For the non-astronomer group, 19% found the screen elements were a continuous distraction while 25% found them initially distracting but quickly learned to ignore them. 27%

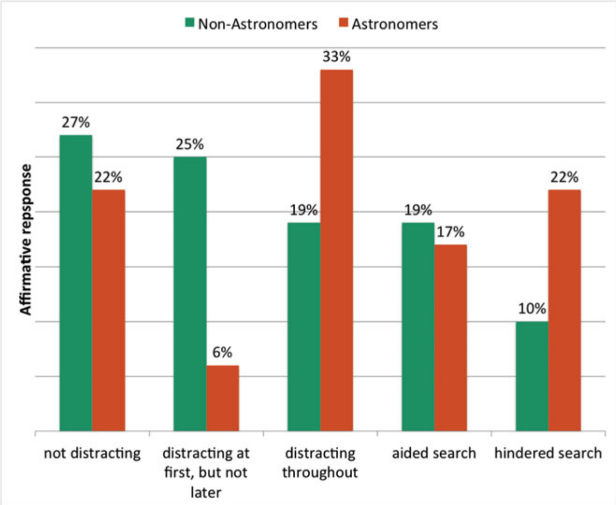


**Figure 15.** Impact of screen elements on the task based on post-experiment survey. The bars indicate the percentage of non-astronomers (green) and astronomers (red) who provided an affirmative response to questions regarding the level of distraction caused by screen elements, and their perceived impact on the search process.

found the screen elements did not distract them at all during the tasks. In fact, 19% of the non-astronomer respondents felt the bezels actually aided their search, compared to 10% who felt they made the task more difficult.

The astronomer group found the screen elements to negatively impact on their experience more so than the non-astronomer group. 33% found the screen elements distracting throughout the experiment, while 6% found them initially distracting, but not so later. However, much like the non-astronomer group, 22% found the screen elements did not distract them from searching for their targets and 17% found the bezels aided their search strategy. However, 22% found the screen elements hindered their efforts.

The general comments relating to the display environments and the screen bezels generally support the results described above. For the SDD, the comments suggest that the environment was well suited to examining very large images in a broad context, when the entire image could be seen. However, small details within these large images were much harder to find and the context was often lost, making strategic searches harder.

Contrasting this with the TDW, the expanse of the display itself sometimes made observing the whole image harder as the participant needed to be much further back to achieve the same field of view, or turning the head considerably more. However, the dynamic nature of physically approaching the screen and being able to see extremely fine detail within the large image made the search for small targets easier on the TDW. This combined with the physical break in the image due to the screen bezels provided subjects with an easier search methodology.

## 4 DISCUSSION

Previous experiments have shown that for searches of small targets within much larger images, a distinct advantage exists for a TDW (Ball & North 2005b; Ball, North, & Bowman 2007; Yost et al. 2007), when the case for physical navigation being preferable to virtual navigation is clear. However, when the target size varies, the advantage is less apparent, though the case for physical versus virtual navigation remains. This is because the larger targets can be found with relative ease in a large image, even when it is presented subsampled on a SDD.

Our study showed that participants typically attempted to gain an overview of the whole image to identify regions of interest. In the case of a large target, this was often more readily found on the SDD as it was quicker to obtain this overview and therefore ascertain the target. As we chose targets that would not easily be confused with other objects, this continued to be the case as the targets got smaller as the participants were able to recognize the approximate shape and zoom quickly only on that part of the image. However, as the targets became too small to even approximately identify, virtual navigation became essential and performance (i.e. success rate) declined rapidly.

While the results show a slight advantage to the TDW for the target searches, it was not as significant as expected. This is in part because the experiment deliberately spanned a range of difficulty, and thereby is inclusive of both the SDD and TDW advantages. However, results obtained from the post experiment survey indicate that participants decidedly preferred the TDW experience over the SDD, even if their performance results did not show a marked difference. Indeed, several participants believed themselves to have performed better with the TDW when they had in fact performed better with the SDD.

The novelty of the TDW environment cannot be ignored and the fact that observing very large images in such an immersive environment will have had some emotional impact on the way participants viewed the experiment. Also, the experiment was clearly investigating the perceived value of TDWs compared to the SDD, and may have skewed participants' perception in favour of the more novel technology. A cross-section of participants primarily sourced from universities may reflect that preference for new technology. However, the primary use of the TDW is in this sector and therefore the performance of such a cohort remains relevant.

The very fact that participants preferred the TDW environment is important even when it did not correspond to increased performance. This suggests that participants might be more inclined to persist with the TDW environment further than with the SDD, however this may be a result of the novelty factor.

No matter how much the novelty factor plays a part, we find that the experience of the participants in this study reflect the results of previous experiments that show performance improvement when virtual navigation can be avoided.

The results from this study indicate considerable opportunities for further work in general testing of TDWs and domain specific testing. This experiment used specifically constructed conditions to examine aspects of the display environments, however, these conditions aren't necessarily indicative of typical astronomical activities. Therefore, before astronomers are likely to include a TDW into their workflow, there needs to be evidence that TDW will actually improve efficiency and/or reduce errors or omissions.

Based on our understanding of how participants used the displays, and in response to individual comments, we identify several areas for future investigation:

- **Multiple image inspection.** Our experiments focused on one particular use case for TDWs - inspection of individual ultra-high resolution images. However, there is an alternative way to take advantage of the display pixels: instantaneous display of many individual lower-resolution images. Consider the simple case of classifying structures in images: a crucial skill in many fields of astronomy which remains difficult to fully automate. Using a TDW would allow astronomers to maintain a view of many classified structures, thereby assisting the evaluation of unclassified structures. This would provide an opportunity for refinement and reclassification, as previous decisions are still available for scrutiny.
- **Extended exposure.** The experiments in this study were designed to run within a 30-minute period to make it easier for participants to commit their time. The results from the survey indicate that most participants felt the TDW was easier to use than the SDD, as well as being more suitable to the task. This suggests that extending the exposure time in a variety of ways might yield more distinguishing results. For example, if the experiment were not time limited, how long would it take to find all targets? Alternatively, if the participants were to use the TDW each day for a period of time, for example, a week, would they improve their performance compare to a control group using SDDs?
- **Features retained over extended periods.** Following on from the previous item, it would be useful to learn whether a TDW aids in the recall of multi-scale features in ultra-high resolution images. Such an experiment would test a participant's ability to relocate features in an image that they had previous been able to find, or had been shown to them. This experiment would look at the difference between short and mid-term memory to see if the TDW exposure shows a difference compared to SDD exposure.
- **Collaborative exploration.** Our study looked at a very basic form of collaboration with two participants working together to share the task. However, there are several variations that would be worthy of further investigation. For example, rather than sharing a SDD, can the frustration caused by sharing control be alleviated by providing each participant with a SDD? It would be expected that more overlap of searched area would occur, but this might be mitigated if each participant could observe the other's display.

  Moreover, how exactly does communication between participants occur in this situation, either naturally or guided? Is physical proximity necessary, or is virtual proximing via teleconferencing facilities sufficient? Such a study has added relevance for the case where participants are working off physically remote but linked TDWs, as in the case of OptIPortals. Finally, increasing the number of participants beyond two, might establish a relationship between screen size and practical use with respect to the number of people observing that data.
- **Consumer 4K UHD displays.** With the recent availability of consumer-grade 4K UHD displays, it would be valuable to repeat the experiment in an environment that might represent an effective combination of the SDD and TDW. While not providing the number of pixels available on a TDW, the advantages of a SDD would be brought to bear, and might produce a cost-effective compromise. Depending on the screen size, this might also prove to be a viable collaborative environment as well as being suitable for an individual. While software like SAGE would work with a 4K display, a significant benefit of running a display from a single computer is that windowing environments can be configured easily and no inter-machine synchronisation is required. This means that all applications can be run without modification.

## 5 CONCLUSION

The amount of information captured by current and future astronomical instruments greatly outstrips the resolution of both current and on-the-horizon displays. TDWs provide a cost-effective method of achieving an order of magnitude increase in display resolution, thereby enhancing the presentation of astronomical data and potentially optimising the consumption of information. However, the notion that TDWs are essential when dealing with extremely large images is not so clear.

The results from this study indicate that TDWs provide a better platform for searching for discrete targets within large images than with a SDD. It also shows that astronomers perform somewhat better than non-astronomers at extracting information from extremely large images (likely due to their more systematic approach to searching), and that the collaborative combination of two non-astronomers using a TDW rivals the experience of an individual astronomer. However, the study also indicates that there is a great variety of aptitude of participants, suggesting that TDWs might greatly enhance the performance for some individuals, while providing little help to others. It also shows that the type of content also has a significant impact on the participants ability to identify targets.

This experiment has borne out the results of earlier research highlighting the benefits of physical versus virtual navigation, and the value of TDWs for searching for very small targets. However, in astronomy, as in many other disciplines, there is a great variation in the physical scale and the "visual connectivity" of the objects to be studied. Our experiments highlight the differences between looking at words, which are processed differently by the brain, compared to identifying isolated galaxies in wide field surveys or visually continuous nebula.

A TDW provides a very impressive environment to examine images and participants enjoyed the experience, which significantly influences their perception of the suitability of the TDW environment. While such value is difficult to quantify, it suggests that the availability of a TDW can be a useful addition to the astronomer's work flow - if only because using one is a more enjoyable task than being seated at the desktop. We are encouraged to believe that the TDW has now come-of-age for astronomy, particularly as a collaborative environment. Ultimately, the practicality of the wider up-take of TDWs for astronomy is contingent on increasing ease-of-use (e.g. through the SAGE environment), suitable interface options and simple training in the use of physical navigation.

## ACKNOWLEDGEMENTS

We thank the participants for contributing their time to this project. All experimental work was approved and conducted in accordance with the requirements of Swinburne Universitys Human Research Ethics Committee (SUHREC). We also thank Dr Simon Cropper and Dr David O'Conner (Psychology dept., University of Melbourne) for their advice on the experiment design, and Christoph Willing (University of Queensland) and Luc Renambot (University of Chicago) for their technical advice in the use of the SAGE environment. We also thank ITS Research Services at the University of Melbourne for the use of the OzIPortal for this experiment.

## REFERENCES

ACSWeb: ACS::The Advanced Camera for Surveys, http://acs.pha.jhu.edu/ (accessed 26 September 2012)
AdlerWeb: Shoot for the Moon Adler Planetarium, http://www.adlerplanetarium.org/experience/exhibitions/shootforthemoon/?searchterm=display%20wall (accessed 27 September 2012)
Andrews, C., Endert, A., & North, C. 2010, in proc. extended abstracts of ACM Conference on Human Factors in Computing Systems (hereafter referenced as (ACM Press)), 55
Andrews, C., Endert, A., Yost, B., & North, C. 2011, Inf. Vis., 10(4), 341
AnglikWeb: Anglik, http://www.anglik.net/english250.htm (accessed 18 October 2013)
Ball, N. M., & Brunner, R. J. 2010, IJMPD, 19(07), 1049
Ball, R., & North, C. 2005a, in (ACM Press), 1196
Ball, R., & North, C. 2005b, in Human-Computer Interaction - INTERACT 2005, ed. M. F. Costabile & F. Patern (Berlin, Heidelberg: Springer Berlin Heidelberg), 350
Ball, R., North, C., & Bowman, D. A. 2007, in (ACM Press), 191
Bezerianos, A., & Isenberg, P. 2012, IEEE Trans. Vis. Comput. Graph., 18(12)
Bi, X., & Balakrishnan, R. 2009, in (ACM Press), 1005
Brescia, M., Cavuoti, S., Djorgovski, G. S., Donalek, C., Longo, G., & Paolillo, M. 2012, in Astrostatistics and Data Mining, ed. L. M. Sarro, L. Eyer, W. OMullane, & J. De Ridder, 31
Calit2Web: Calit2?: Calit2 and University of Melbourne Initiate Australias Ultra-Resolution Global Collaboration Laboratory, http://www.calit2.net/newsroom/release.php?id=1219 (accessed 24 September 2012)
CGLXWeb: CGLX Project, http://vis.ucsd.edu/~cglx/ (accessed 26 September 2012)
Cruz-Neira, C., Sandin, D. J., DeFanti, T. A., Kenyon, R. V., & Hart, J. C. 1992, Commun. ACM, 35(6), 64
DarkEnergySurveyWeb: The Dark Energy Survey, http://www.darkenergysurvey.org/ (accessed 24 September 2012)
DeFanti, T. A., et al. 2009, Future Gener. Comput. Syst., 25(2), 114
DeFanti, T. A., et al. 2010, Cent. Eur. J. Eng., 1(1), 16
Febretti, A., et al. 2013, in Proc. SPIE 8649, ed. M. Dolinsky & I. E. McDowall, 86490386490312
Fujiwara, Y., Date, S., Ichikawa, K., & Takemura, H. 2011, J. Inf. Process. Syst., 7(4), 581
Hassan, A., & Fluke, C. J. 2011, Publ. Astron. Soc. Aust., 28(2), 150
HiperwallWeb: HIPerWall — hiperwall.calit2.uci.edu, http://hiperwall.calit2.uci.edu/ (accessed 26 September 2012)
HyperSuprimeCamWeb: Hyper Suprime-Cam, http://www.naoj.org/Projects/HSC/index.html (accessed 26 September 2012)
Ivezic, Z., et al. 2008, SerAJ, (176), 113
Keller, S. C., et al. 2007, PASA, 24(1), 1
Koribalski, B. S. 2012, PASA, 29(3), 359–70
LeedsPowerwallWeb: Leeds Virtual Reality Powerwall project wins a large extension - Faculty of Engineering - University of Leeds, http://www.engineering.leeds.ac.uk/faculty/news/2009/powerwall.shtml (accessed 26 September 2012)
Mohr, J. J., et al. 2012, ArXiv Prepr. ArXiv12073189
Moreland, K. 2012, in Proceedings of the IEEE Symposium on Large-Scale Data Anaysis and Visualization
Morikawa, Y., et al. 2010, AGU Fall Meet. Abstr., D5
NICTWeb: Universal Communication Research Institute — Information Services Platform Laboratory — NICT-National Institute of Information and Communications Technology, http://www.nict.go.jp/en/univ-com/isp/topics.html (accessed 27 September 2012)
OptiportalWeb: Optiportals, http://optiportal.org/index.php/Main_Page (accessed 11 March 2014)
OziPortalNewsWeb: Fast HD ultra-broadband link is powerful collaboration tool?: News?: The University of Melbourne, http://archive.uninews.unimelb.edu.au/news/4894/index.html (accessed 27 September 2012)
Ponto, K., Doerr, K., Wypych, T., Kooker, J., & Kuester, F. 2011, Future Gener. Comput. Syst., 27(6), 649
SAGEWeb: SAGE: SCALABLE ADAPTIVE GRAPHICS ENVIRONMENT, http://www.sagecommons.org/ (accessed 26 September 2012)

Sims, M. H., Dodson, K. E., & Edwards, L. J., 2010, LPI Contrib., 1538, 5614
Smarr, L., Brown, M., & de Laat, C. 2009, Future Gener. Comput. Syst., 25(2), 109
Tada, T., Shimojo, S., Ichikawa, K., Abe, H., & others 2011, in Granular Computing (GrC), 2011 IEEE International Conference on, 613
Taesombut, N., et al. 2006, Future Gener. Comput. Syst., 22(8), 955
Tyson, J. A. 2002, in Proceedings of the SPIE, ed J. A. Tyson & S. Wolff, 10
Way, M. J., Scargle, J. D., Ali, K. M., & Srivastava, A. N. 2012, Advances in Machine Learning and Data Mining for Astronomy (CRC Press)
Yamaoka, S., Manovich, L., Douglass, J., & Kuester, F. 2011, Computer, 44(12), 39
Yost, B., Haciahmetoglu, Y., & North, C. 2007, ACM CHI, 101

## A THE OzIPortal PROJECT

We now look at the OptIPortal project and a specific TDW, the OzIPortal, in more detail in order to understand some of the reasons why these devices have not already become standard elements of the research workflow across diverse scientific disciplines.

### A.1 The OptIPortal project

The OptIPortal project grew out of the Optiputer project, a US government funded project to connect high performance computing facilities together via optical networks, called Lambdas (DeFanti et al. 2009; Taesombut et al. 2006). With such powerful data processing and transfer resources in the background, an enhanced visualization capability was required. This was initially achieved using several projectors, with edge-blending techniques to compensate for luminosity fall-off between adjacent projections. The OptIPortal project took this methodology to the next level, adopting high-resolution off-the-shelf displays to create tiled surfaces. Software was developed to make the management of the TDW relatively transparent so that the users could focus on the research content (DeFanti et al. 2010, OptiportalWeb 2014).

Collaborative environments such as the OptIPortal network, were designed to allow distributed research teams to work together simultaneously on the resulting imagery and analysis. This would provide greater opportunity for collaborative research, leading to greater understanding of the data (Smarr, Brown, & de Laat 2009; Sims, Dodson, & Edwards 2010; Yamaoka et al. 2011).

While the project promised a simple, powerful and interconnected system, there were many problems with the early incarnations of OptIPortals, primarily due to immaturity of the associated software. To tackle this, workarounds were commonly employed, such as manual data processing steps. Whilst giving the appearance of success for visualization, this led to some misunderstandings as to what an OptIPortal actually is and its overall utility as part of research workflows.

### A.2 The OzIPortal experience

In 2008, The University of Melbourne launched the OzIPortal, a 98 megapixel TDW, with considerable fanfare. It was lauded as an amazing research tool: “In an Australian first, this next-generation platform set to revolutionize the way Australia interacts with the rest of the world allows real-time, interactive collaboration across the globe, combining high-definition video and audio with the sharing of ultra-resolution visualizations from a broad range of disciplines.” (Calit2Web 2012).

Despite the high level of interest generated by early demonstrations, the OzIPortal failed to attract a significant commitment from the research community.

The OzIPortal was configured using 24 $\times$ 2560 $\times$ 1600 LCD displays, initially in an 8 $\times$ 3 arrangement and later in 6 $\times$ 4. 12 slave nodes with dual-head graphics cards were used to drive two monitors each. Reducing the number of slave nodes improved the operation of the TDW without reducing the performance. Data was exported from the head node via NFS to each of the slave nodes. An additional machine was required to provide a real-time video stream that would allow any video signal captured via HDMI to be presented on the TDW. In this way, the OzIPortal was able to include low latency, high-definition video conferencing content alongside other stored content.

One of the initial drivers for the OzIPortal was to establish a dedicated gigabit link from The University of Melbourne through to CALIT2 in the United States (OziPortalNewsWeb 2012). This was successfully implemented on a layer 2 network via AARNet. The purpose of the network was to demonstrate the rapid transfer of massive datasets, to show how distant collaborators could work on the same datasets at the same time. In practice however, it was necessary to distribute the data in advance, as some of the data was too large to deliver in a timely manner, even over the dedicated link. Instead, the link was used primarily to stream uncompressed video directly to the TDW, instead of via a video conference codec. Finally, it took several technicians to satisfactorily operate the TDW, and a great deal of testing beforehand was required to minimize disruptions during events. Day-to-day operations could not be sustained with such human resource demands and as such provided a less than satisfactory experience for users.

Ultimately, what was required for the successful deployment of a TDW in the research workflow was a more stable system, that was easier to use, did not require high-level of support, automated preparation of content, and the ability to run applications specific to individual scientific disciplines. Our recent experiences with the OzIPortal, particularly through the use of the SAGE environment, is a positive step towards these outcomes.

## 4.4 Performance based on self-rated expertise and astronomy background

An additional comparison of performance based on self-rated expertise and on astronomy background was considered during this study, however the results were not included in Meade et al. (2014) paper. They are included here for completeness. During the interview phase of this study, participants rated their level of expertise in astronomical imagery from one to five, and to identify themselves as an astronomer or non-astronomer. 12 identified themselves as astronomers, while 45 identified themselves as non-astronomers. Of the 12 astronomers participating, 7 rated themselves as 4 or 5, while from the 45 non-astronomer participants, only 2 rated themselves as 4, and none rated themselves as 5. Of the 45 non-astronomers, 16 completed the tasks in collaboration with another participant, and 14 undertook another set of related tasks that proved inconclusive. Only the remaining 15 non-astronomer participants' results are considered below.

Figure 4.2 Panel A shows the median performance results for the word, galaxy and nebula feature searches for the 12 astronomers and 15 non-astronomers. Figure 4.2 Panel B shows the same performance results for the expert group and the non-expert group, where 4s and 5s as considered experts, and those with 3 or below are considered non-experts. In both panels there is no clear advantage for searching for astronomical features based on self-rated expertise, and a small advantage based on astronomy background.

Although we have claimed that there are differences in the two cohorts, we did not undertake a more formal statistical analysis. Normally this would be done with the one tailed T-test, which would allow us to determine if the difference seen in the performance of the astronomer cohort was in fact statistically significant compared to the non-astronomer group. Despite this, the insights gained about suitability of the TDW were examined and ultimately validated in situ as part of the *Deeper Wider Faster* observing campaign, described in Chapter 5.

## 4.5 Lessons learned

While TDWs might not have been the revolutionary research tool as they were touted, they do have a valuable place in a research workflow that relies on the manual inspection of imagery that exceeds the standard desktop display resolution by orders of magnitude. The efficacy of a TDW has been shown as a collaborative tool for simultaneous inspection of Gigapixel imagery, and for communication and training with such content.

This study tested the hypothesis that ultra-high resolution displays could improve

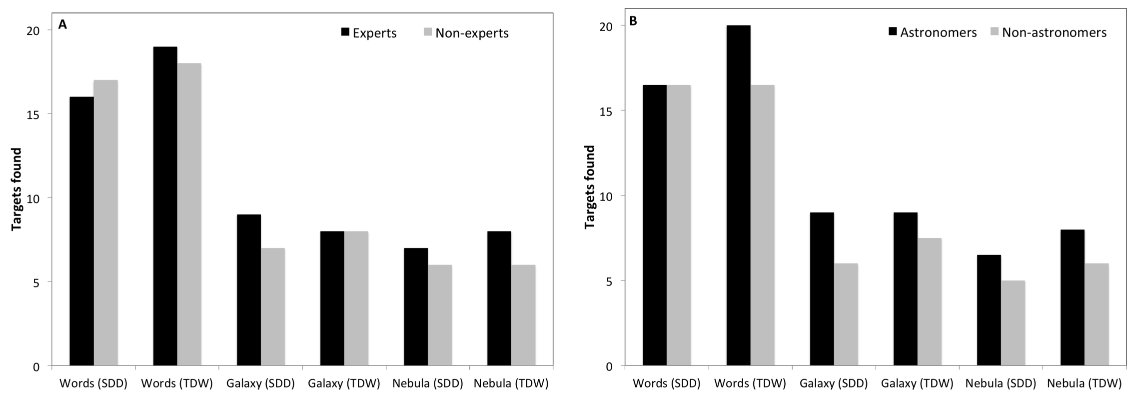


Figure 4.2 Participants' performance for identifying words, galaxies and nebula features, based on their self-rated level of astronomy expertise are shown in Panel A. Experts (those who rated their expertise as 4 or 5) did not perform better than those who considered themselves non-experts (self-rating 3 or below). Panel B shows the same results for 12 astronomers and 15 non-astronomers. A small advantage can be observed for astronomers in all but one category.

research outcomes in astronomy. This was achieved by exploring two of the Lam et al. (2012) scenarios, UP and UE, as described in Chapter 3.

This research has shown that a TDW can:

- reduce time taken to find critical feature(s) within an image, thereby improving user performance;
- improve success rate in object confirmation; and
- enhance parallel search performance for collaborative image inspection

Timing the ability of astronomers to find features in large-scale astronomical imagery under controlled experiment conditions, on both the standard desktop display and the TDW, showed that a TDW improved astronomers' performance when compared to a standard desktop. It was also shown that non-astronomers and astronomers both showed improved performance while using the TDW, and that two non-astronomers working in collaboration performed as well if not better than a single astronomer, thereby demonstrating the efficacy of a TDW for collaborative search tasks.

The use of virtual navigation on the TDW was not required for this study, however, many astronomical images exceed even the 98 Megapixel display wall used here. As can be seen in Figure 1 of Meade et al. (2014, see section 4.3), the Hubble image of the Carina nebula is more than four times the resolution of the TDW, and if displayed at full resolution, would require virtual navigation to fully explore. Additional planning would be needed to provide users with a suitable mechanism for panning and zooming the image, especially as traditional mouse and keyboard combination typically rely on a physical desktop surface to operate. Potential solutions such as multi-touch displays and mid-air pointing techniques are discussed in von Zadow et al. (2014), Nancel et al. (2015), and Jakobsen & Hornbæk (2016).

The study also found that for the tasks as described in Meade et al. (2014), participants found the experience of using a TDW to be more enjoyable than using a standard desktop display. While the novelty of the TDW probably contributed to this qualitative response, combined with the results of the quantitative results, it is clear that the TDW did provide the more suitable environment in this study.

# 5

# Workspace Display Ecologies

## 5.1 Overview

A rapidly growing area of research in astronomy is the field of fast transient events and in particular Fast Radio Burst (FRB), kilonovæ and gravitational wave events. The *Deeper, Wider, Faster* program aims to investigate these events by coordinating multiple telescopes and deep sky cameras to observe with multiple wavelengths, and as quickly as possible. The multi-facility campaign is proactive in that it targets a specific region of sky for several days catching as many events as possible in the observing period. The captured data is processed as quickly as possible in case a particularly interesting event is recorded, e.g. a supernova or Fast Radio Burst (FRB). If this occurs, additional facilities can be diverted from other tasks to participate to in follow-up observations such as collecting spectra. The events under consideration occur quickly, and hence it is imperative that sensible decisions about triggering additional resources to join the observation are made as soon as possible. To this end, the *Deeper, Wider, Faster* program has evolved an optimised workflow to capture, process and inspect the data as fast as possible.

Aside from the many techniques to improve data transfer and computation times, the inspection workflow itself has also been optimised. Lessons learned from the pilot observing runs led the principal investigators to seek support in the inspection part of the workflow. They identified that the inspection process was seriously limited by inadequate display technology, especially when dealing with multiple images that were several times larger than the screens being used to inspect them. This is particularly important when the objects of interest are only a few tens of pixels in diameter in a $4000 \times 2000$ pixel image and where many images needed to be inspected simultaneously. The Advanced Immersion Environment at the University of Melbourne hosted two advanced display technologies (in 2016) to enhance the inspection workflow of the observing campaigns.

The first of these displays was the OzIPortal, the 98 Megapixel TDW used by Meade et al. (2014, see Chapter 4, Section 4.3) to investigate the use of TDWs in astronomy. This TDW was used to display the difference images (created by subtracting the captured image from a reference image) at full resolution, for rapid inspection by volunteer astronomers.

The second display was a 6.9m × 2.2m curved projection screen, which used two 1920 × 1200 projectors with a 400 pixel overlap to create a seamless image across the display. This display showed the thumbnail images (usually less than 200 × 200 pixels) of potential candidates, along with their science and reference thumbnail counterparts.

These displays were used alongside several laptop and desktop computers to form a display ecology, where displays were chosen and optimised to support a particular function in the inspection workflow.

## 5.2 Collaborative Workspaces to Accelerate Discovery

The following paper documents the use of advanced display technologies to support the *Deeper, Wider, Faster* observing campaign. We found that matching suitable data content to appropriate display technologies enhances the efficacy and experience of the visual inspection workflow. Furthermore, the advanced displays used during the campaign supported both collaborative parallel searching, and the training of new astronomers in the potential candidate identification process.

The paper was published as "Collaborative Workspaces to Accelerate Discovery" by Meade, B., Fluke, C., Cooke, J., Andreoni, I., Pritchard, T., Curtin, C., . . . Reynolds, T. (n.d.). . Publications of the Astronomical Society of Australia, 34. doi:10.1017/pasa.2017.15. Reprinted with permission of Publications of the Astronomical Society of Australia.

*Publications of the Astronomical Society of Australia* (PASA), Vol. 34, e023, 18 pages (2017).

doi:10.1017/pasa.2017.15

# Collaborative Workspaces to Accelerate Discovery

Bernard Meade[1,2,8], Christopher Fluke[1], Jeff Cooke[1], Igor Andreoni[1,3,4], Tyler Pritchard[1], Christopher Curtin[1], Stephanie R. Bernard[5], Albany Asher[1], Katherine J. Mack[3,5,6], Michael T. Murphy[1], Dany Vohl[1], Alex Codoreanu[1,3], Srđan M. Kotuš[1], Fanuel Rumokoy[5], Chuck Horst[7] and Tristan Reynolds[5]

[1]Centre for Astrophysics and Supercomputing, Swinburne University of Technology, PO Box 218, Hawthorn 3122, Australia
[2]Research Platform Services (Doug McDonell Building), The University of Melbourne, Victoria 3010, Australia
[3]ARC Centre of Excellence for All-sky Astrophysics (CAASTRO), The University of Sydney, NSW 2006, Australia
[4]Australian Astronomical Observatory, PO Box 915, North Ryde, NSW 1670, Australia
[5]School of Physics (David Caro Building), The University of Melbourne, Victoria 3010, Australia
[6]ARC Centre of Excellence for Particle Physics at Terascale (CoEPP), School of Physics, The University of Melbourne, Victoria 3010, Australia
[7]Department of Astronomy, San Diego State University, San Diego, CA 92128-1221, USA
[8]Email: bmeade@unimelb.edu.au

(Received September 28, 2016; Accepted March 30, 2017)

**Abstract**

By applying a display ecology to the *Deeper, Wider, Faster* proactive, simultaneous telescope observing campaign, we have shown a dramatic reduction in the time taken to inspect DECam CCD images for potential transient candidates and to produce time-critical triggers to standby telescopes. We also show how facilitating rapid corroboration of potential candidates and the exclusion of non-candidates improves the accuracy of detection; and establish that a practical and enjoyable workspace can improve the experience of an otherwise taxing task for astronomers. We provide a critical road test of two advanced displays in a research context—a rare opportunity to demonstrate how they can be used rather than simply discuss how they might be used to accelerate discovery.

Keywords: techniques: miscellaneous

## 1 INTRODUCTION

Investment from research institutions and governments in new astronomical facilities, scientific instruments, and high-performance computing capabilities occurs at a great scale. The shared resources are typically heavily oversubscribed and time allocation on these instruments is extremely competitive. Simultaneously, new science, such as the search for fast transients, places an even bigger strain on the available resources as it requires several telescopes for coordinated observation, with additional telescopes to be on standby for immediate re-pointing if a significant event occurs. Therefore, it is imperative that all aspects of the scientific workflows engendered by this research infrastructure are scrutinised. While much effort is expended evaluating the processes for the observing, computation, and storage components of a workflow, only recently has attention been given to the operational workspace in which humans interact with the technological systems. For many years, the standard computer display served to present all manner of digital content, from text to graphs to images, with little consideration as to the appropriateness of the display to the content. In order to accelerate discovery, workspaces that facilitate collaboration and understanding in real time, both *in situ* and remotely, are fast becoming essential. A carefully considered *display ecology* (Huang, Mynatt, & Trimble 2006; Chung et al. 2015) that addresses specific visualisation tasks are a key component to achieving satisfactory scientific outcomes.

### 1.1. The role of the display

Computer displays, or monitors, have become such an integral component of the astronomer's scientific toolset, that it can be easy to overlook their significance or impact on the research workflow. It can be tempting to continue to use a display—even when its size or resolution begin to limit productivity—simply because it is available or on-hand rather than assessing the capabilities of an alternative solution.

It may be necessary to increase the physical size of the display in order to inspect very large datasets (e.g. thousands of DECam[1] 520 megapixel mosaics), where each image is orders of magnitude larger than the display it is being viewed

[1] http://www.ctio.noao.edu/noao/node/1033

on; achieve meaningful collaboration; and make discoveries that require the rapid or simultaneous validation additional astronomers and experts. This can be achieved through the purchase of a bigger—and thus usually more expensive—monitor, using a digital data projector, or by adding additional screens to the desktop. Occasionally, it requires a more radical re-thinking of what a display can be.

Some of the earliest work on alternative displays for astronomy was by Fomalont (1982) and Rots (1986). Norris (1994) examined the potential role for qualitative, comparative, and quantitative visualisation, and Fluke, Bourke, & O'Donovan (2006) presented options for collaborative environments: multi-projector tiled displays; digital domes; and large-scale collaborative, stereoscopic exploration of three-dimensional datasets, *viz* The Virtual Room.

### 1.2. Collaboration

Scaling up a display (in terms of physical size and the number of pixels) necessitates a move away from the desktop. There is an opening up of space around (or in front of) the display, encouraging researchers to stand up, move around, and share the workspace with their colleagues. These are key elements that can turn visualisation and inspection of data from a solitary pursuit into a collaborative one.

The value of collaboration in scientific discovery has been at the heart of endeavours such as the CAVE and the OptIPortal projects (Cruz-Neira et al. 1992; Smarr et al. 2003, 2005; DeFanti et al. 2009, 2010; Febretti et al. 2013). Placing multiple researchers in the same physical space and allowing them to interact with data together allows them to experience a shared engagement with the information. Contrast this with coincident engagement when they experience the data at separate locations at the same time, even while in contact via communication technologies.

Advanced collaborative workspaces with large, immersive display technologies have been in use around the world for over a decade, yet their impact on the research landscape has been relatively limited. While these facilities have been used as educational tools and high-impact demonstration environments (AdlerWeb 2007; SDOWeb 2015; QUTCubeWeb 2016), there is a dearth of published research that identifies dedicated collaborative display environments as a critical component in the workflow that has produced new scientific outcomes. Furthermore, beyond the advantages of the technology itself, the value of bringing astronomers together in a single space to collaborate in real time is considerable.

### 1.3. Tiled display walls

On the face of it, it seems likely advanced displays should lead to more rapid scientific discovery and would therefore be deemed essential. A specialised tool that improves engagement, by enhancing immersion or providing access to many more pixels, should afford greater insight to its users and scientific outcomes should follow. In reality, that has not been the case. But perhaps the issue is not that the displays themselves are not capable of achieving such goals, but that they have yet to find the right place in the research workflow.

Meade et al. (2014) consider some of these possibilities, while also testing the assumption that a display environment such as a tiled display wall (TDW) can actually improve research outcomes.

Clustering homogeneous computer displays to simulate a single continuous display canvas has been possible for some time. In this approach, the physical displays are placed in a fixed array and connected to one or more computers, often referred to as nodes or workers. These nodes are coordinated by a single head node, which typically does not take part in the display environment.

The content being displayed on any individual screen is synchronised by the head node to provide the user with the appearance of a single display. In this way, media content can be moved around the entire display almost seamlessly. High-resolution images (or movies) that greatly exceed the resolution of an individual display can then be shown at full resolution across several displays. The only interruption to the display space are the screen mountings, called bezels, at the edge of each screen.

The Meade et al. (2014) study concluded that in certain circumstances, and for certain people, a TDW will improve a user's ability to find small targets ($185 \times 145$ pixels) within a much larger image ($12\,000 \times 5\,812$ pixels). It also highlighted a tendency for individuals to prefer physical navigation, that is, the use of physical body movements such as eyes, head, and the whole body, to virtual navigation, which uses computing interfaces such as keyboard and mouse, when inspecting very large images. These findings were consistent with the outcomes of Ball and North (2005a, 2005b), Andrews, Endert, & North (2010), and Andrews et al. (2011).

### 1.4. The display ecology

By construction, Meade et al. (2014) used an artificial context that *resembled* a research activity from astronomy: visual inspection of an image looking for known types of objects in unknown locations. However, it also identified another context in which a TDW might be useful: parallel inspection of many images, as opposed to a single extremely large image. Yet, a TDW is not suitable for all types of data, and may enforce technological limitations of its own (such as the limited availability or functionality of software that can drive the display—see section 4.3).

Combining display technologies to form a display ecology (Huang et al. 2006; Chung et al. 2015) offers a best of all worlds approach. While each display can overcome a particular hurdle to understanding, they can also ignore or exacerbate others. Using the right display for the right content in concert improves a researcher's ability to draw on many sources to construct a more complete mental picture of the science at hand.

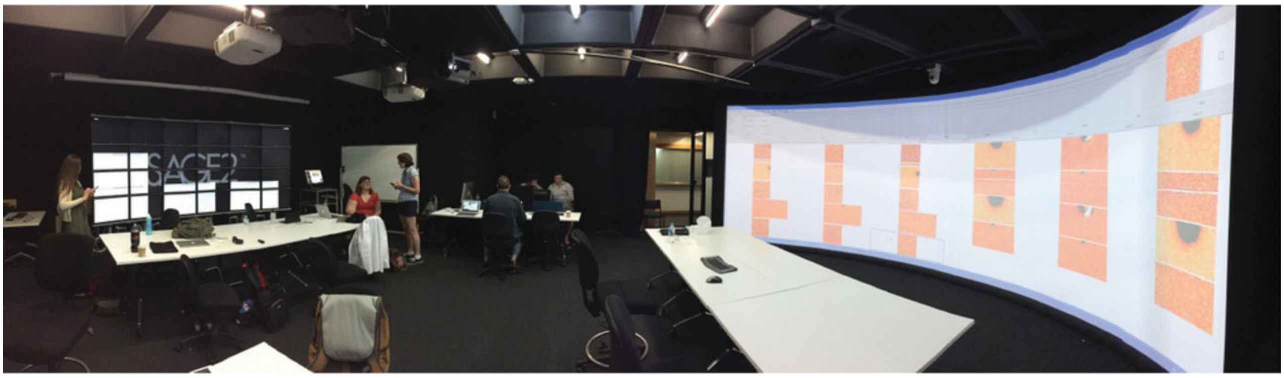

**Figure 1.** A panoramic view of the workspace used for O1, showing the TDW at the left of the image, the review, and control stations in the middle and the curved projection screen to the right. The whiteboard shown centre left was used to log potential candidates for review, as well as other important details including telescope on sky times.

### 1.5. Overview

In this paper, we present a case study based on our use of a collaborative workspace to support the *Deeper, Wider, Faster* initiative. *Deeper, Wider, Faster* is a coordinated, contemporaneous, multi-wavelength observing programme. It aims to make rapid, real-time identification of fast transients, i.e. those with a duration of milliseconds to hours, including Fast Radio Bursts (FRBs), Gamma-ray bursts, flare stars, kilonovæ, and supernova shock breakouts, using telescopes across the globe and in orbit.

A full description of the observing strategy, discovery pipeline, and detections from the first four campaigns are outside of the scope of this paper. Full details of the *Deeper, Wider, Faster* programme may be found in Cooke et al. (in preparation), Andreoni et al. (submitted), and future papers. We discuss only those aspects of observation, discovery, and analysis that informed our approach to understanding, adopting, and improving the display ecology.

The *Deeper, Wider, Faster* pilot programme (see Section 2) raised a number of issues relating to large-format image inspection and collaboration. While preparing for future campaigns, a TDW was identified as being a strong candidate to eliminate several key problems with the existing desktop-based workflow. This visualisation environment was augmented by the use of large-format curved display, with the two displays working in tandem.

In the remainder of this paper, we present the collaborative workspace used for the *Deeper, Wider, Faster* 2015 December 17–22 UT (operational run 1: O1) and the 2016 July 26–2016 August 7 UT (operational run 2: O2) campaigns. During these observing runs, up to 12 astronomers at a time shared the workspace.

For O1, our solution used two advanced displays, a 98 Megapixel TDW and a curved projection screen (6.9 m circular segment with 2.56 m radius, 2.2 m height), co-located in the Advanced Immersive Environment at the University of Melbourne—see Figure 1—along with several laptops and desktop computers.

One of the main objectives of the display ecology is to enable rapid identification (in minutes) of fast-evolving transient events to inform other telescopes of the location of the discoveries and to 'trigger' them to rapidly move to obtain spectroscopy or additional imaging of the objects before they fade.

During O1, several candidates were identified as potential spectroscopic trigger candidates, and a number of triggers were sent. For example, a live trigger was sent to Gemini-South for spectroscopic follow-up and did result in the successful acquisition of the spectrum of an extragalactic transient currently under investigation. Moreover, the process proved valuable in uncovering CCD artefacts, amplifier crosstalk, and other effects that can produce 'fake' fast transients. Also, the observing team were able to critically road test how advanced displays can be used—a rare alternative to previous discussions of how they might be used (Fomalont 1982; Rots 1986; Norris 1994; Fluke et al. 2006).

For O2, the processing workflow underwent several additional improvements based on our insights, and user feedback, from O1. Specifically, we rearranged the Advanced Immersive Environment at the University of Melbourne—see Figure 2, and integrated additional online tracking of candidates. The improved workflow, including advancements in the automatic candidate detection pipeline, resulted in three triggers sent to the Gemini-South Observatory and four triggers to the South African Large Telescope (SALT). In addition, round 570 spectra were obtained via the Australian Astronomical Telescope (AAT) with over 50 targets identified for follow-up with the Zadko Telescope (University of Western Australia) and Skymapper (Australian National University). In all, tens of thousands of candidate variable and transient objects were detected during this run.

Through a combination of pre-campaign questions, observations of usage patterns during the observing period, and post-campaign reflection, we

1. demonstrate that careful design of a collaborative workspace can greatly improve the rate at which CCD images can be inspected;

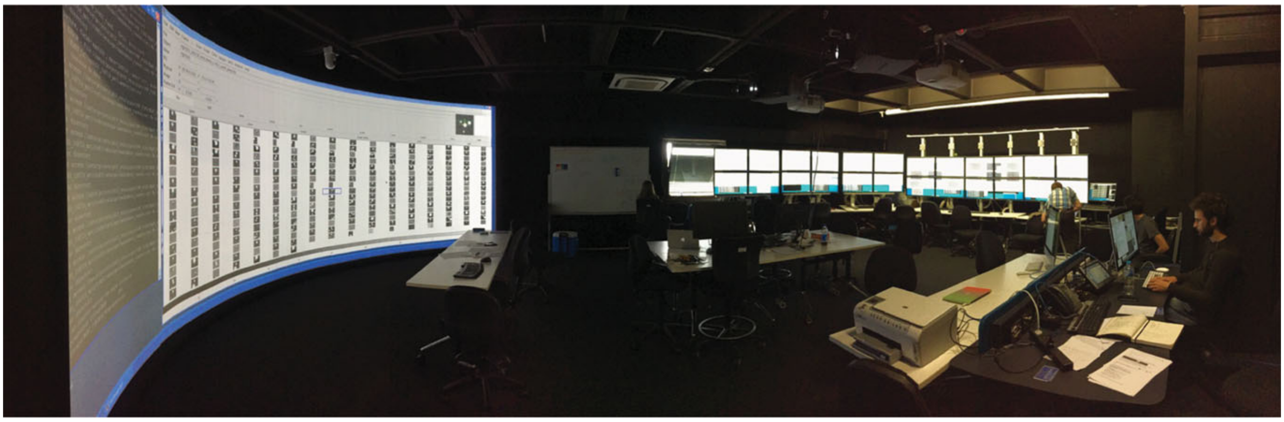

**Figure 2.** A panoramic view of the updated workspace for O2, showing the reconfigured TDW at the right of the image, the review, and control stations in the middle and the curved projection screen to the left.

2. show how facilitating rapid corroboration of potential candidates and the exclusion of non-candidates improves the accuracy of detection; and
3. establish that a practical and enjoyable workspace can improve the experience of an otherwise taxing task for astronomers.

The paper is set out as follows. In Section 2, we describe the pilot programme for the *Deeper, Wider, Faster* observing campaign, and identify the visualisation-based bottlenecks inherent in the original workflow. In Section 3, we discuss the setup for the 2015 December (O1) and in Section 4 we evaluate the collaborative workspace and the impact of the display technology on the workflow. The changes implemented for 2016 July/August (O2) are described in Section 5. In Section 6, we discuss planned improvements to both the process and the technological workflow in order to improve future scientific outcomes. Concluding remarks are made in Section 7.

## 2 DEEPER, WIDER, FASTER

*Deeper, Wider, Faster* primarily uses the Dark Energy Camera [DECam; Diehl 2012; Flaugher et al. 2012] on the Blanco 4-m telescope at the Cerro Tololo Inter-American Observatory (CTIO) in Chile to observe a region of the sky. These fields are simultaneously observed at radio wavelengths by the Parkes radio telescope and the upgraded Molonglo Observatory Synthesis Telescope (MOST) in Australia, the NASA *Swift* Space Telescope in low-Earth orbit, and occasionally other facilities such as the Very Large Array (VLA) in the US.

Should a suitable transient candidate be identified, such as a potential optical counterpart to an FRB, alerts were to be sent to additional telescopes, such as the Gemini-South Observatory, for targeted optical/infrared spectroscopic follow-up. In order to confirm an event, obtain its redshift, localise it, obtain its host galaxy properties, and understand its nature, spectroscopy needs to be acquired within minutes of the detection before the fast transient event fades, thus the urgency to process and identify candidate sources.

In the example of FRBs, optical and spectroscopic counterparts have yet to be identified and their behaviour at wavelengths other than radio are completely unknown, making it challenging to design and use a purely automatic detection pipeline to identify possible progenitors. For now, there is an important role for visual inspection of images and potential candidates at all stages of the workflow. This includes making judgements as to the likelihood that a potential candidate could be a counterpart, performing quality control tasks, or making serendipitous discoveries of as yet undetermined transient objects.

### 2.1. The pilot programme

The initial *Deeper, Wider, Faster* observing campaigns were held from 2015 January 13–16 UT (pilot run 1: P1) and 2015 February 27–28 UT (pilot run 2: P2). Figure 3 shows the flow of data from the DECam imager on the Blanco telescope at Cerro Tololo in the Chilean Andes to Swinburne University of Technology, Melbourne, Australia. While several other telescopes were involved, this paper focusses on the collaborative workspace used for visualisation and review of DECam images to discover transient sources.

For P1, several members of the observing team were situated at Cerro Tololo, to view and interact with the DECam images directly prior to transfer to Australia. The observers were able to use a six-panel tiled display (consisting of 27″ desktop LCD monitors @ 1 920 × 1 200 pixels on a standard desktop computer); however, individual CCD images could not be expanded across the full display. Instead, the screens were used to display six concurrent CCD images on individual monitors along with researchers' desktops and laptops. For P2, most of the team were located at Swinburne where all the analysis was performed using only desktop or laptop computers.

PASA, 34, e023 (2017)
doi:10.1017/pasa.2017.15

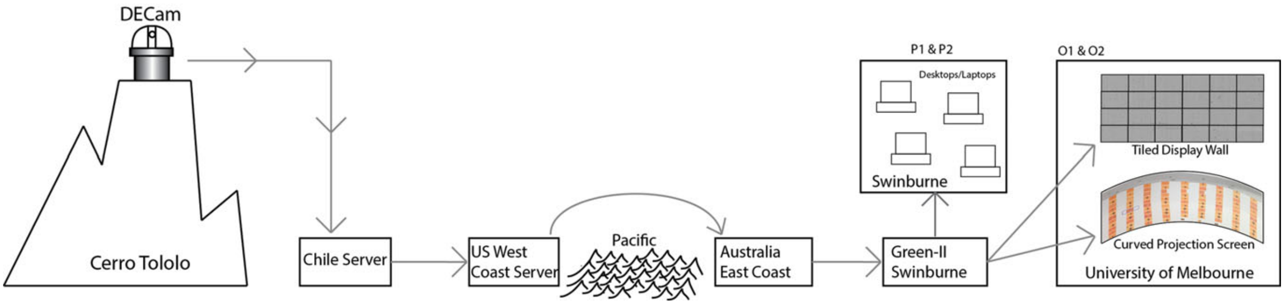


**Figure 3.** While data from many telescopes was collected, the focus of the data inspection optimisation for O1 and O2 was on the optical image data captured with the DECam imager in Chile that was then transferred to Green II supercomputer at Swinburne University for processing. In the pilot programmes, P1 and P2, the images were inspected on desktop and laptop computers in Chile and Swinburne University. In O1, after processing, the images were transferred to the University of Melbourne for inspection on the tiled display wall and on the curved projection screen (see Section 1). In O2, the images were inspected on tiled display wall reconfigured as six individual workstations (see Section 5), and on the curved projection screen.

The pilot runs P1 and P2 were designed as an opportunity to uncover and deal with obstacles that typically prevent real-time fast transient detection and study. This process identified the need for sophisticated visual inspection and in turn, the development of a supporting display ecology. In addition, P1 and P2 brought about the development of software to provide real-time data compression, processing, and analysis, and the software for real-time candidate identification.

The DECam CCD images were subtracted from a template to produce difference images. This provides the best opportunity to identify significant changes in the images since the template was captured that might indicate an event of interest. While perfect alignment is not possible and many artefacts remain after the subtraction, the combination of automatic catalogue lookup and the eyes of experienced astronomers are able to find the objects of interest. At 4 096 × 2 048 pixels, these were significantly larger than the resolution of the standard displays used (up to 2 560 × 1 440 pixels), therefore the inspection relied heavily on *virtual navigation*—zooming and panning—to search for potential candidates. This process was performed in parallel for each of the DECam CCDs, and occasionally with sections of the full DECam mosaic (60 CCD images).

The full images were examined (1) to understand the context and, equally, (2) to determine if the sources were CCD artefacts such as amplifier crosstalk (which requires full CCD inspection). As the physical scale of potential candidates is unknown, there is a risk of overlooking a feature of interest due to pixel subsampling. On a display that is considerably smaller than the image size, viewing at native resolution requires methodical scanning of the images, which is tedious. Zooming in and out on features of potential interest makes objects on the edge of detectability very difficult to find.

While the impact of pixel subsampling was not investigated in depth, a qualitative assessment was made by the authors by comparing the native resolution CCD images with a full screen version on the 2 560 × 1 600 displays. At this resolution, the CCD images are being displayed at less than 50% of their native resolution and the authors found the potential candidate detection to be far more difficult to perform.

The process of inspecting the individual CCD images in this manner was deemed to be a significant bottleneck in the selection of potential candidates for follow-up study. Even though the process was slow, it was still essential in determining potential candidates before a trigger to engage additional telescopes could be sent. The process was also greatly complicated when several CCD images needed to be compared. Finally, the lack of physical space in front of standard desktop displays limited collaboration, forcing the researchers to use multiple independent computers and displays, thus reducing the effectiveness of a parallel search.

With more observing runs planned, it was important to change the processes. The development of an automatic detection pipeline (described in Andreoni et al. submitted) and the use of advanced displays were expected to significantly improve productivity. With an emphasis on decreasing the time spent inspecting each CCD image, eliminating virtual navigation, enhancing collaboration, and integrating the automatic search more completely into the visual search, a new collaborative workspace was needed.

### 2.2. Workspace requirements

In preparing for O1 (2015 December 17–22 UT), there was a clear need to improve the visual inspection workflow. The five key requirements were as follows:

- Decrease the time to inspect a full CCD, or even the entire 60-CCD field of view of DECam: optimising the time taken to complete the visual inspection of difference images is critical for confirming suitable candidates for rapid follow-up, reducing the time from days or hours to minutes.
- Remove virtual navigation: by eliminating the need to pan and zoom images, inspectors can be more confident of complete coverage of an image and reduce the risk of overlooking potential candidates.
- Enhance collaboration: having inspectors working independently but immediately adjacent provides rapid corroboration of potential candidates, with minimal disruption to the inspection workflow.

- Integrate the automatic search more completely into the visual search: make better use of the astronomers' time to look at the most important things and to enhance the mutually supporting review processes.
- Completeness test: use the visual inspection as a means to provide a completeness test and training set for the software candidate selection.

Several constraints were imposed on the development of the display ecology, including

- no funding to secure hardware resources specifically to support the display ecology;
- no staff resources for developing a bespoke software solution to support the display ecology; and
- relocating the compute resources and inspection team to Chile to reduce the impact of the physical separation between the workflow and the capture device, DECam, would have been prohibitively expensive.

## 3 O1: THE 2015 DECEMBER CAMPAIGN

To address the shortcomings of the pilot programme (Section 2.1), a new workflow was developed for O1. This included the use of an automated candidate identification pipeline and an improved visualisation process to manually review the CCD images.

### 3.1. Automating candidate selection

The simultaneous multi-wavelength imaging strategy and real-time optical imaging analysis component of *Deeper, Wider, Faster* is as follows:

1. **Data collection and transfer:** The DECam electronics provides a 20 s readout time for the entire set of 62 CCDs. The *Deeper, Wider, Faster* programme chose to take a continuous stream of 20 s exposures to provide the fastest cadence to search for fast transients, while maximising survey depth and time on sky. Fields are simultaneously observed for 1–2 h by several observatories, with the time on field constrained by the coincident visibility by DECam in Chile and Parkes and Molonglo in Australia. As a result, around 100–200 DECam optical images are acquired per field. Image files are 1.2 GB (uncompressed), but total data increases by 3–4 times during processing. Images and processed files are stored on Swinburne University's Green II Supercomputer facility. While there are 62 science CCDs in the DECam array, only 59 were usable during O1 as two CCDs were non-functional and one had a damaged amplifier and could not be calibrated. These 59 CCD images are referred to as a *batch* for the remainder of the paper.
   JPEG2000 compression was performed at CTIO in order to compress the data sufficiently to expedite the transfer (Vohl et al. in preparation). For this purpose, we modified the KERLUMPH software (Vohl, Fluke, & Vernardos 2015) to convert files from the FITS format[2] into the JPEG2000 (ISO/IEC 15444) format (JPEG2000-part1. 2000). The level of file compression was determined on-the-fly to keep transfer time reasonable while maintaining sufficient information to achieve the science goals. Because several of the subsequent processing steps did not support JPEG2000, the images were converted back to FITS. Data transfer from CTIO to Swinburne University took between ~3 and ~15 min per batch of images during O1, and ~1 to ~5 min for O2. On reaching Swinburne, the images are uncompressed and processed (Andreoni et al. submitted).
2. **Initial processing:** Data were processed in stages using eight reserved nodes on the Green II supercomputer.

   a. Individual CCD images are calibrated using parts of the PhotPipe pipeline (Rest et al. 2005).
   b. The *Mary* pipeline (Andreoni et al. submitted) is used to coadd, align, and subtract the images, and to automatically search for transients. *Mary* identifies CCD artefacts and poor subtractions to reduce a sample of several thousand initial detections (across all CCDs) to a few tens of objects.
   c. Finally, *Mary* generates products for visualise inspection such as postage stamp images (varying between 80 px and 120 px per side) and region files identifying the nature of known variable and other sources from catalogues to assess potential candidates to follow-up.

3. **Visual inspection:** At the same time, as the *Mary* pipeline was extracting potential candidates, the full-frame difference images were also viewed in their entirety. While the preference was to maintain the FITS format, the requirements of the TDW necessitated converting the images to JPEG in order to display in the Scalable Amplified Group Environment (SAGE2)[3] environment. For O1, it was thought that this was more important to achieve than the use of FITS compatible software (see Section 3.3 for more detail), however for O2, with the more developed pipeline for eliminating unwanted candidates, the benefits of FITS was more important. This corresponded with requests from the O1 inspectors to reconfigure the TDW with only two rows for ease of use during O2. The use of FITS files enabled demarcating software identified candidates on the full CCDs, as well as known variable objects, known CCD crosstalk, and other information. This approach provided a visual confirmation of the efficacy of the *Mary* pipeline, as well as the opportunity to find targets possibly missed by the pipeline. Other problems with potential candidates that could fool the automatic system—but hopefully not a trained astronomer—include amplifier crosstalk and CCD defects (see Sections 3.5.2).

[2] http://fits.gsfc.nasa.gov
[3] http://sage2.sagecommons.org/


doi:10.1017/pasa.2017.15

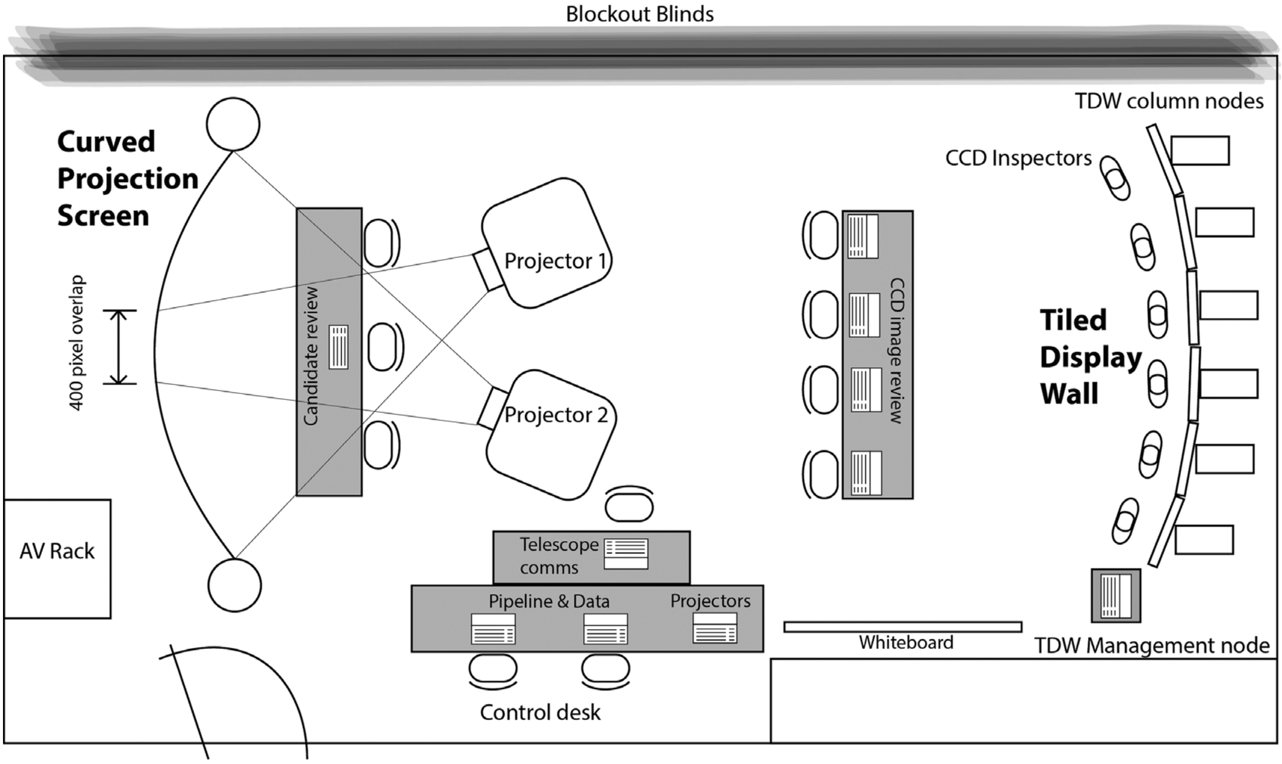


**Figure 4.** Floor plan of the Advanced Immersive Environment at the University of Melbourne for O1. The room configuration allowed the two principal activities, i.e. reviewing the software identified candidates on the curved screen and inspection of the CCD difference images on the TDW, to be conducted independently while supporting collaboration between these tasks. The control desk had an excellent view of both sides of the room, and team members here could easily respond to requests from either side.

### 3.2. Room configuration

The Advanced Immersive Environment at the University of Melbourne was chosen as the base of operations for O1. It offered access to a 98 Megapixel TDW, a large-area curved projection screen, and table-top work spaces for the team members to bring and use their own devices. Moreover, with around 100 $m^2$ of floor space, there was ample room for the team to move, work, and collaborate effectively.

During O1, the room was configured with two principal enhanced display technologies—see Figure 4 and Table 1. One end of the room was occupied by the TDW and was used to display the processed difference images in JPEG format. Each CCD image was 4 000 × 2 000 pixels and the TDW display area was 15 360 × 6 400 pixels. In order to optimise the 6 × 4 display configuration and ensure the images were shown at native resolution, the images were presented in a 3 × 3 matrix to allow clear space between the images, and reduce the need to use the uppermost region of the TDW. It was necessary to have the images appear across four screens with no pixels hiding behind the bezels, as shown in Figure 5.

At the opposite end of the room was the curved projection screen, which was used to show the postage stamps images of the software-detected candidates. Typically around one hundred of these candidates were displayed simultaneously across the 6.9 m × 2.2 m display using SAOImage DS9. Other applications could also be displayed alongside the DS9 window, such as IRAF, to assist in the evaluation of the candidates.

Other operations were positioned between these two displays to allow easy observation from the process facilitators.

### 3.3. Tiled display wall software

The SAGE2 software was used to manage the display windows being presented on the TDW. In this client server model, the SAGE2 Head Node acts as a HTML5 web server, with the 'clientID'tag being used to specify which window is to be streamed to the client. For example, with the head node running the 'node.js' based service, a Firefox window is launched on a tile display node, and directed to a particular URL for that frame. This environment was chosen as it made it possible to script the loading and display of the CCD images on the TDW, as well as log the time taken to do so. It also enabled easy review of the individual images at much larger scale for close scrutiny, as the images could be expanded across the entire display if required.

### 3.4. Workflow

As a batch of images arrived at the Green II Supercomputer, the *Mary* pipeline produced a collection of potential candidates for display on the curved projection screen. In parallel


doi:10.1017/pasa.2017.15

**Table 1.** Hardware specifications of the principal workstations and projectors used during O1.

| TDW SAGE2 | Specification | Display |
|---|---|---|
| Head node | Virtual server (NeCTAR Research Cloud)16 vCPUs, 64 GB RAM, 500 GB Volume Storage, 10 GB network, Ubuntu 14.04 LTS | No attached display |
| Management node | Dell T3400, Quad-core Intel, 2.4 GHz, 16 GB RAM, 500 GB HDD, 1 GB network, Ubuntu 14.04 LTS | 1 × 19 inch display (1 680 × 1 050) |
| Display nodes (×6) | Dell T3400, Quad-core Intel, 2.4 GHz, 16 GB RAM, 500 GB HDD, 1 GB network, Quadro FX570, 2 GB VRAM, Ubuntu 14.04 LTS | 4×Dell Ultrasharp 30 inch display (2 560 × 1 600) |
| Curved Screen | Specification | |
| Head node | Dual core Xeon 3.00 GHz, 3GB RAM, 500 GB HDD, 1 GB network, Windows XP SP4 | |
| Christie projectors | 2 × 1 920 × 1 200 (400 pixel overlap for image blending), fast phosphor, 120 Hz for active stereo (not used in this experiment) | |
| Control station | Specification | |
| Pipeline and data | iMac 27-inch (diagonal) LED-backlit display (2 560 × 1 440) | |

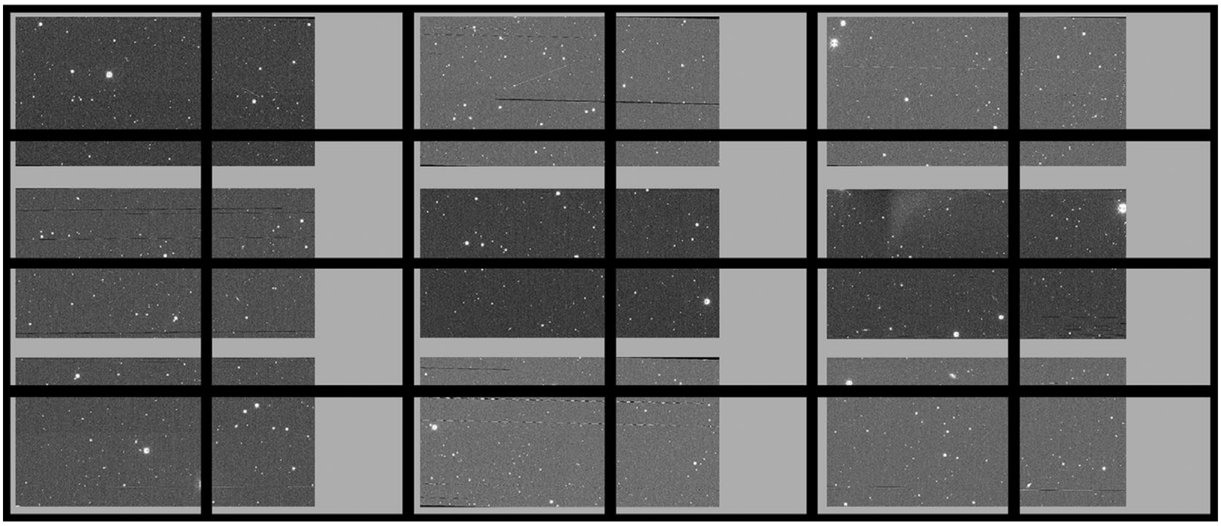

**Figure 5.** During O1, in order to avoid any image size reduction, the best image configuration for the TDW was 3 × 3. This provided clear separation between images but also meant that each image spread across four screens. The bezels did not obscure any image pixels.

to this process, *Mary* generates template-subtracted FITS files which are converted to JPEG. This conversion was necessary for O1 but abandoned for O2; see Section 4.3 for more detail. These were transferred directly to the head node of the TDW and visually searched for potential candidates (see Section 3.5.1).

Python scripts were used to present the images on the TDW running the SAGE2 interface. The display script automatically loaded and positioned nine images on the display initially and as the researchers completed the inspection of an image, it was replaced with a new image. After the initial images were loaded, an operator monitored the progress of the inspectors, and manually advanced the script to load the next image when it was clear the inspector had moved on to a new image. This process continued until the full set had been inspected. See Figure 6 (Top).

The images were presented in columns with an image placed at the bottom of each column and progressively moving up in rows. As each column was assigned to a researcher, the presentation order was intended to reduce the wait time for each researcher to start their inspection task. In this way, each researcher was presented an image in their assigned column before anyone else received a second image. The choice to start with the bottom row was decided by the people inspecting the images as preferable to loading top down.

As images were inspected, candidates of interest were flagged for follow-up, with approximate locations noted. Initially, interesting candidates were recorded using paper (as this was a natural reaction) but was then moved to the whiteboard as seen in Figure 1. These targets were then inspected on a standard laptop computer running SAOImage DS9, and on the curved screen, also running DS9. Where necessary, images could be recalled to the TDW for verification and comparison.

Several bash and python scripts were developed to expedite the workflow. These include such things as moving a completed batch of jpegs to a storage folder to make way for the next batch, or starting or stopping the TDW nodes.

**Figure 6.** (Top) The OzIPortal TDW with images displayed in 4 × 5 configuration during O1. Several configurations were tested but the 3 × 3 configuration was deemed most suitable. (Bottom) A large number of candidates, with science images and subtractions, shown as postage stamps, can be inspected at once by several researchers, and shared with anyone in the room. This was particularly useful in supporting novice inspectors.

### 3.5. Training the image inspectors

Several of the team members had never worked with a TDW or curved projection screen before, while others had extensive experience. Roles ranged from observing the use of the display technologies while working on their own tasks, to those who worked exclusively with the TDW and/or curved screen. See Figures 1 and 6 for examples of the displays in use during O1.

Introducing the candidate identification/rejection process required a short training session for the image inspectors. Meade et al. (2014) found that using a TDW was an unfamiliar experience for most people and without an introduction, it was unlikely to be particularly useful. However, a short explanation of physical navigation, i.e. physically moving your body to achieve the equivalent of panning and zooming, improved the experience and efficacy of using a TDW.

This orientation process was augmented for O1 by using sample images showing examples of potential candidates, as well as examples of system or processing errors, such as badly subtracted images and crosstalk. The collaborative environment meant volunteers could be trained 'on the fly', which

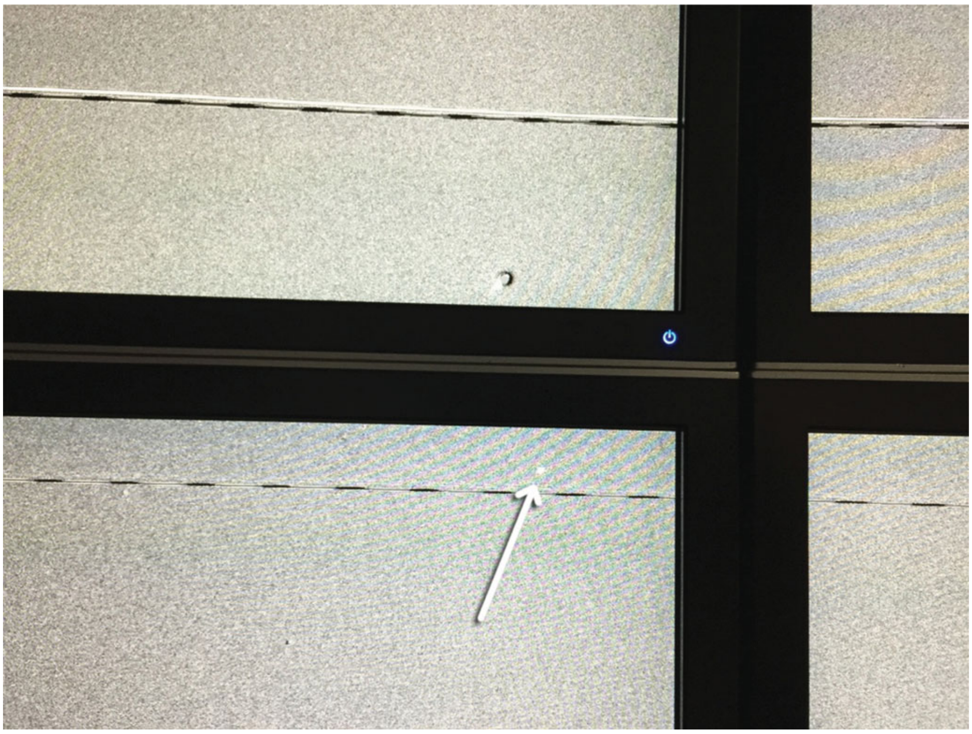

**Figure 7.** An example of a potential candidate on the TDW that meets all the necessary criteria for closer inspection and possible follow-up with other telescopes.

was very useful considering the dependence on volunteers with varying availability.

### 3.5.1. *Potential candidate*

Potential candidates are expected to appear as two-dimensional Gaussian point sources in the images and (roughly speaking) appear as small, round objects with soft edges and no black (negative flux) artefacts that could indicate poor subtractions of non-transients or CCD effects such as bad pixels or column subtractions. If a potential candidate met each of these conditions, they were usually corroborated by other researchers and then flagged for more detailed inspection, with approximate coordinates noted—see Figure 7.

### 3.5.2. *Amplifier crosstalk and CCD defects*

Each CCD has two amplifiers that can create artefacts when processed by the operating system electronics. When a source in the region of amplifier A saturates, it creates a crosstalk image in the region of amplifier B, equidistant from the line joining the amplifiers. Potential candidates that had a clear counterpart on the opposite site of the image could be eliminated from consideration, such as shown in Figure 8.

Occasionally what appears to be a potential candidate shows a negative partner observed at the same offset as other potential candidates within the image, as shown in Figure 9. While the precise nature of this effect is unknown, it is likely an artefact of the DECam CCDs and the fast data processing, and not celestial phenomenon. Fortunately, the display ecology helps easily identify the effect that could be missed by other conventional identification techniques.

### 3.5.3. *Time spent on tasks*

The workflow described above was used by the *Deeper, Wider, Faster* team over the six nights of O1. The team assembled from around 12:00 and prepared for on sky observations at 15:00 until 19:30. The direct measure of image loading time on the TDW was able to be tracked by a log generated by the script for displaying the targets.

The transfer of converted JPEG images from the Green II cluster at Swinburne University in Hawthorn to the SAGE2 head node located at the University of Melbourne's Queensberry Street Data Centre in Parkville, did not contribute significantly to the workflow and the transfer time was not tracked.

An image display control script was used to populate the TDW with images as quickly as possible, positioning the first nine images in a 3 × 3 matrix. As soon as an image was available, inspection started.

The aim of the control script was to ensure the participants always had a new image available when they were ready to move on. The initial loading time for the images was quite consistent, with an average time for the first nine images of 54.9 se. Each column had an image within 20 s, which includes additional scripted delays such as clearing the TDW


doi:10.1017/pasa.2017.15

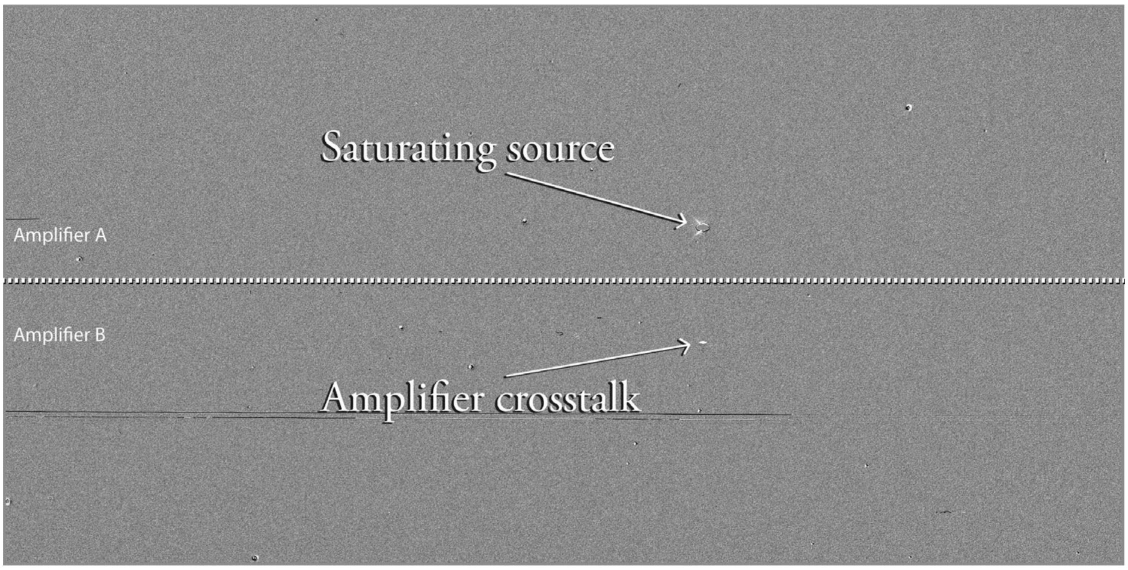


**Figure 8.** Each CCD has two amplifiers reading out each half of the image. Sometimes this will result in a crosstalk image of a saturated source from one amplifier to the other.

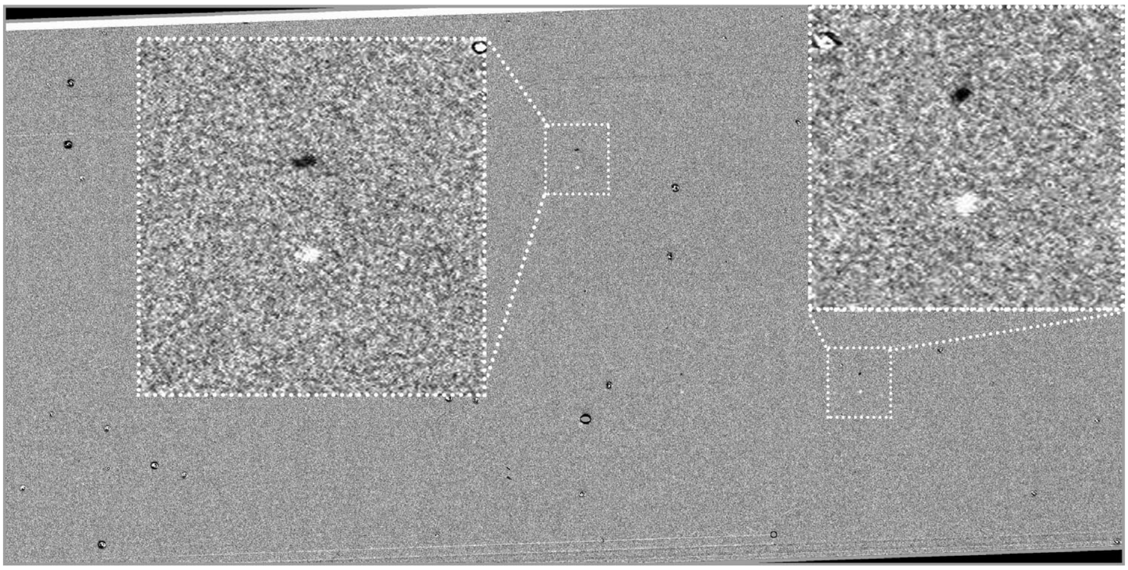

**Figure 9.** When several potential candidates show a negative partner offset by a regular amount, the potential candidate can be eliminated from consideration.

(2 s), loading and positioning images in each column (2 s to load and 2 s to place for each column, totaling 12 s). These delays were built in to the script to avoid race conditions, that is, where a compute process attempts to complete two or more tasks at the same time and fails, at the head node.

While the time to completely review a full batch of images was not formally recorded, Table 2 shows the duration of the TDW image review process taken from the first image loaded to the last image loaded for that day, and the number of images reviewed, as logged by the control script.

## 4 EVALUATION OF THE COLLABORATIVE WORKSPACE FOR O1

### 4.1. Expectations

Before O1 began, several members of the *Deeper, Wider, Faster* team reflected on the role that an alternative display ecology might have on overcoming the limitations of the pilot programme. The comments here refer only to the TDW, as the use of the curved screen for O1 had not been confirmed at the time. Four broad themes emerged.

PASA, 34, e023 (2017)
doi:10.1017/pasa.2017.15

**Table 2.** The image display control script was used to log the start and end times of image loading during O1. The shorter duration on the 23rd of December was due to problems with DECam that limited observing time.

| Date | Duration | Images |
|---|---|---|
| 2016-12-18 | testing | N/A |
| 2016-12-19 | 2 h 30 m | 424 |
| 2016-12-20 | 2 h 29 m | 303 |
| 2016-12-21 | 4 h 45 m | 494 |
| 2016-12-22 | 4 h 22 m | 267 |
| 2016-12-23 | 2 h 39 m | 341 |

**Throughput:** Utilising multiple astronomers to inspect the images in parallel should improve the throughput of the images in a set. By having the images automatically loaded and positioned on the TDW for the astronomers, there should be no wait time once the first image is available for inspection. This assumes that it is quicker to load new images than to inspect an image. Each astronomer can complete their images and should time permit, they can easily assist others.

**Rapid corroboration:** With astronomers inspecting images side-by-side, there is the potential for rapid corroboration of a suspected candidate. An astronomer can easily leave an image being inspected to assist a colleague nearby to determine the viability of a candidate. When complete, the astronomer can easily return to their own image. Because the researchers are in close physical proximity, this can happen very quickly. The short delay in inspecting a particular image should not make it difficult to return to the image and pick up where the astronomer left off.

**Native resolution:** A thorough inspection of each image is necessary as the potential candidates are likely to be represented by only a few pixels. Having the images shown at full resolution should reduce the possibility of overlooking potential candidates that might be missed due to subsampling caused by scaling, or when panning and zooming. This should also help rapidly identify artefacts and thereby reduce time spent on non-candidates.

**Workflow optimisation:** The use of an advanced display such as the TDW should improve the overall workflow and help design future workflows that are optimised for speed and accuracy. It should also help refine the pipeline in the identification of candidates for the future.

### 4.2. Impact of the tiled display wall

At the conclusion of O1, the *Deeper, Wider, Faster* team again reflected on their experiences, this time with both the TDW and the curved screen. While successfully meeting the expectations (Throughput, Rapid corroboration, Native resolution, Workflow optimisation), additional themes were identified.

**Candidate rejection:** Crosstalk artefacts are due to the dual amplifiers for each CCD. When this occurs, a potential candidate can be eliminated from further consideration because it is being generated by a non-candidate in the other amplifier. The observing strategy we adopted in O1 avoided performing dither patterns in order to maximise the continuity of sampling each part of the CCDs. As such, we uncovered the extent of this effect but, at the time, it was difficult to consistently anticipate crosstalk locations. Non-candidates that would have been otherwise discarded might appear as potential candidates in the reflected part of the image. The postage stamps themselves are not large enough to show evidence of this effect, however, it is quite easy to identify this phenomenon when looking at the whole image on the TDW.

**Quality control:** Other errors such as CCD defects, CCD processing problems from the real-time pipeline or telescope tracking or guidance problems are far more obvious on the full resolution images displayed on the TDW. When time is of the essence, rapid identification of faults is essential to avoid wasteful delays and prevent rapid-response telescopes triggers on non-celestial sources.

**Missed discovery:** As with any automated system, it needs specific criteria in order to make a selection. While this does not mean an entirely new phenomenon cannot be discovered, it does open the possibility of missing something that might catch the eye of a trained astronomer.

The sheer volume of data being collected from astronomical instruments these days mean it is essential to exploit automatic processes wherever possible, as typically there is simply too much information for human eyes to sift through in a meaningful time. The best option is the combination of automatic processes and manual inspection. As the automatic processes become more mature, they reduce the pressure on the manual processes, though it is hard to imagine if full discovery space can ever be fully automated. In the context of unbiased searches for fast or exotic transient events, the combination worked extremely well, with both the curved screen and TDW inspection processes being used to great effect to support each other.

**Throughput:** Images displayed on the TDW are able to be inspected far more quickly than was possible in the previous run of the experiment. Parallel inspection with several astronomers working on separate images significantly speeds up the process, with one observer estimating around 50% improvement in efficiency of detection confirmation or rejection of candidates.

**Native resolution:** The objects of interest are small, usually representing less than 0.08% of the image area. They are circular and have a soft edge i.e. a point source, thereby having a two-dimensional Gaussian-like shape, yet this is often lost when the image is subsampled, such as when viewed full screen on a display of lower resolution than the image. On such a screen, many more objects appear to have this profile until they are zoomed into, when they can be seen to be not circular, or have hard edges or other artefacts not apparent before. The TDW (and indeed any display capable of displaying images at native resolution) eliminates the need to zoom in, and so speeds up the rejection of non-candidates greatly.

The TDW encouraged whole body movement to scan images rather than just with eyes. This maintains the scale of an object in the context of the image which is difficult to match when panning and zooming on a standard desktop display.

**Human factors:** Another benefit of the TDW was the ability to recall images for the purpose of comparing epochs. To perform effective transient candidate detection, it is necessary to recall images from other epochs for comparison. During O1, the automatic process was not designed to efficiently crosscheck every candidate in previous images that the manual process was able to perform. This feature was added for O2, and complemented by the online logging tool (see Figure 11. Not only did this identify several interesting events worth following-up, it is also invaluable for maturing the automatic process for future observations. Giving objects 'running IDs' has since been employed to allow precisely this sort of temporal tracking for subsequent runs.

### 4.3. Problems with the tiled display wall

While aiding with the throughput, rapid corroboration, and overall experience, the TDW posed some logistical and operational challenges.

The physical height of the TDW made it difficult to see the upper regions of the images in the top row for some inspectors. The lowest regions of the bottom row also presented some difficulty as they required the inspector to bend or squat, which became uncomfortable after several hours of moving up and down through the images. This could be improved by reconfiguring the TDW into a $12 \times 2$ configuration, which would redeploy the top and bottom rows, providing a more comfortable fit with the average viewing height. This new configuration would allow additional columns of images, making it easier to include additional inspectors. The practicalities of changing the configuration made it too difficult to employ during O1 but was implemented for the 2016 July/August UT (O2) campaign.

As the TDW is necessarily made up of many smaller screens, it is impossible to avoid screen edges. While it is possible to purchase screens with negligible bezel (screen edges) size, these are very expensive. The screens used in this TDW have bezels of 20 mm, making a combined bezel width of 40 mm, and sometimes more due to slight gaps between the screens themselves. As image resolution exceeded the screen size, each image spanned four displays (see Figures 5 and 7), with a break in the image at the screen edge. Meade et al. (2014) showed the practical and psychological impact of the screen bezels on an observer is typically small, however inspectors reported it can be distracting. When potential targets that lay within a few pixels of the break in the image were encountered, the image could be shifted slightly to place the candidate in question in an unbroken region of the display. This however requires additional time, but fortunately happened only a few times, none of which resulted in a positive candidate selection. Reducing the physical size of the bezels would reduce this problem.

The depth of the bezels to the screen surface also meant that for the top row of screens, the outward protrusion of the bezel itself was obscuring pixels at the bottom edge of the screen, as seen by someone looking upward. Reducing the depth of the bezel and/or reconfiguring the tiles to reduce the need to look upward as much would reduce the possibility of missing candidates.

It was necessary to convert from FITS to JPEG in order to display on the TDW due to format restrictions of the SAGE2 software. This added an additional step to the workflow which, while relatively minor, became a tripping point on several occasions. Minor human-generated mistakes, such as beginning a transfer before the full set had been converted meant the process had to be repeated to pick up missed images. Transferring was initiated manually and on multiple occasions saw a set of images transferred too soon, overwriting a set of images during inspection. If the TDW could handle FITS images directly, possibly with alternative software, then the transfer step could more easily be automated and would reduce the potential for human error. Such an approach has been successfully tested by Pietriga et al. (2016) with the FITS-OW software, but that application is still in development. An alternative would be to configure the TDW in subgroups of $2 \times 2$ screens connected to a single computer, thereby allowing the direct use of SAOImage DS9 to display the FITS images.

The process of loading, resizing, and moving images was slower than expected due to race conditions at the SAGE2 web server. In order to avoid this, short delays of 1–2 s were built into the scripts to ensure a response from the web service. Once loaded, manual movement and scaling of the images was acceptable. These delays were added *in situ* to cope with problems as they occurred. While error trapping would have negated the need to incur delays on each load and move command, the time to develop such a solution was not deemed useful during O1, as the cumulative delays were only in the tens of seconds over a batch of images.

SAGE2 did not provide a convenient way to flag potential candidates and note their coordinates within an image. The number of promising candidates were relatively few and were relayed to the analysts by identifying the CCD and approximate location of the candidate either verbally or via notebooks and use of the whiteboard. While not ideal, this did not cause a major problem as it was typically done to verify a potential candidate within a current run, so the coordinates were reasonably well known. However, an online logging tool shared in real time has been developed for subsequent campaigns. This allows inspectors to use laptops and mobile devices to log potential candidates *in situ*, resulting in improved reporting consistency and tracking of review outcomes.

While the process of inspecting the images was entertaining and engaging, after closely inspecting several hundred images, the observers did become tired. This was due to the mental demands of being thorough and the physical requirement of standing in front of the TDW for several hours.

### 4.4. Impact of the curved projection screen

When viewing the postage stamps of potential candidates produced by the *Mary* pipeline, the curved projection screen provided a more suitable display surface than the TDW. The curved display has a resolution of 3 440 × 1 200 (due to the 400 pixel blending region) over a physical display surface of 6.9 m × 2.2 m.

The immersive nature of this display enhanced the experience for the researchers as several reported feeling more engaged with the information being presented. Driven by a single computer, the curved screen produces a more 'desktop' like experience that could be easily viewed by everyone at once, especially those across the room examining the full CCDs. There was a faster level of responsiveness when compared to operations performed through SAGE2. Without the physical presence of bezels on the TDW, the image blending of the two projectors provided an uninterrupted display area, and so all images remain unbroken.

Not only could many postage stamps of potential candidates be viewed simultaneously (100 images was typical), it is also possible to have other applications running alongside the SAOImage DS9 software. The X11 applications were forwarded from the Green II cluster. The bandwidth provided between sites (Swinburne University to the University of Melbourne) was adequate to operate the application with negligible lag. Other software such as IRAF and multiple terminal windows were also forwarded from Green II, and displayed alongside the candidates being presented by DS9. Having all the necessary information readily available and easily viewable by several people demonstrates the utility of the environment.

At 15.18 m$^2$, the actual surface area of the curved screen is much larger than the 6.17 m$^2$ of the TDW. This increased physical size made it easy for several people to collaborate on the same content at once. Moreover, the screen's curvature meant that content at the far edges was less horizontally compressed (from the central viewing point) than with a flat display of equivalent size.

### 4.5. Problems with the curved projection screen

As with the TDW, the curved projection screen posed several challenges with regards to its use, suitability, and display qualities.

When compared with the TDW, the curved projection screen has much lower resolution and contrast. The increase in area afforded by the curved screen was slightly counteracted by the lower pixel density: 15.9 pixels mm$^{-2}$ to 0.27 pixels mm$^{-2}$, respectively. Despite the lower resolution, the curved projection screen was ideal to display the postage stamps whose resolution is comparatively very low, thus, the pixels were resolved. Similarly, the text windows and graph displays, while not perfectly sharp, were quite adequate for the task.

While approaching the screen surface did result in shadows cast by the front projection system, this did not discourage the astronomers from getting very close to the screen to discuss objects of interest. There was a slight impact from the in-room lighting. For safety and general usability of the collaborative workspace, some lights were required to be on during the observation to facilitate people moving around the room. While the spill from the overhead lights was minimal, reducing it even further would have been desirable.

The lack of suitable drivers for the hardware used by the curved projection screen dictated the use of Microsoft Windows XP SP4. Rather than locating and installing Windows versions of the preferred Linux applications required, X11 forwarding was tested (with Putty[4] and XMing[5]) and found to perform very well. The applications were being forwarded from Green II at Swinburne, where the image data was stored and the cluster processing occurred.

As a single large display space, the curved screen functioned as a standard, albeit very large, desktop computer. A useful capability would be the ability to drive the display from another computer with a pre-configured environment more suited to the task, with drive paths and device drivers already installed. Also being able to have multiple people working independently on the same screen but in separate windows, with their own keyboard and mouse control, would greatly increase the versatility of the environment. We suggest to introduce an intermediate step to achieve this would be to use screen sharing, with researcher laptops being replicated on the curved screen.

## 5 O2: THE 2016 JULY/AUGUST CAMPAIGN

Operating from the 2016 July 26 until the 2016 August 7, with between 5.5 and 11 h each observing session, O2 adopted several improvements in workflow.

- **Display ecology**: The TDW was reconfigured, eliminating the top and bottom rows and spreading the screens out into six workstations with 2 × 2 screens, with each workstation computer operating largely independently, though sharing a file system with the other workstations. With 5 120 × 3 200 pixels, these workstations were able to display the 4 096 × 2 048 pixel images at native resolution, and provide sufficient screen real estate for the SAOImage DS9 toolbar. The addition of desks in front of the screens provided inspectors with a place to use their laptops to access the online spreadsheet, however, they could still stand if required. See Figures 2 and 10 for the updated layout.
  The new configuration eliminated the physical observation problems associated with the TDW for O1. With the top and bottom rows removed, the screen height was more consistently comfortable (though not customisable

[4] http://www.chiark.greenend.org.uk/~sgtatham/putty/
[5] https://sourceforge.net/projects/xming/


doi:10.1017/pasa.2017.15

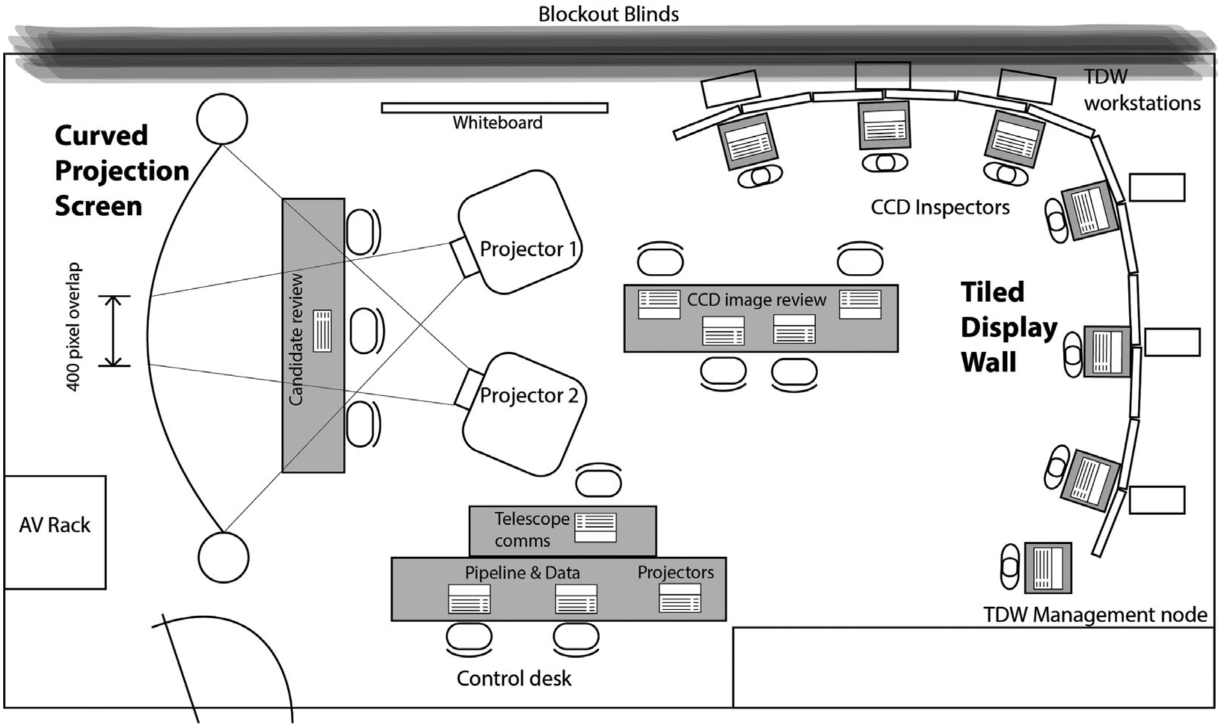


**Figure 10.** Updated layout of the Advanced Immersive Environment at the University of Melbourne. The curved screen for reviewing the *Mary* candidates remained unchanged from O1 to O2. The TDW was broken into six workstations with 2 × 2 tiled screens, and space for a users laptop. The central desk was also rotated to facilitate better movement between work areas.

to individuals). This also effectively removed the bezel depth occlusion problem described in Section 4.3, as the inspectors were easily able to reposition themselves to eliminate the issue.

- **Software:** Using DS9 allowed the inspectors to load FITS images directly from the TDW head node. Along with the FITS images, automatically generated DS9 region files were available for overlay on the images. These regions included the persistent Candidate ID numbers, making the process of identifying them within the full resolution image much simpler for the inspectors.
- **Event logging:** Using an online spreadsheet to log and track potential candidates, including their real-time light curves, the requirement for immediate inspection of the full CCD subtraction image was mitigated. Instead, individual or groups of images could be called up for review if they had already been identified by the *Mary* pipeline and logged in the spreadsheet, as seen in Figure 11.

The rest of the workflow remained relatively unchanged from O1. This preserved the collaborative and training benefits of the workspace from O1. While other optimisations to the detection pipeline were made, these did not affect the overall workflow significantly.

Improvements to the workflow from both the *Mary* pipeline and the display ecology resulted in significant outputs for O2. Several triggers were sent during the run to Gemini-South and SALT, with more than 50 targets identified for subsequent follow-up with Skymapper and the Zadko Telescope. Hundreds of candidates received spectra and follow-up imaging and tens of thousands of candidate variable and transient objects were detected.

As the principle objective of the display ecology had been established in the planning and execution during O1, no formal attempt was made to conduct an additional review for O2. Instead, the achievable recommendations from the inspectors after O1 were implemented and subsequent workflow improvements developed organically in response to the new configuration. The positive response from the inspection team was unanimous in supporting the need for the display ecology.

## 6 DISCUSSION

O1 operated on sky for approximately 4.5 h per d, with visual inspection continuing for up to an hour longer, over six consecutive days where the workflow was continuously refined. O2 operated for between 5.5 and 11 h, with additional time for visual inspection, over 13 consecutive days. As a mark of a successful endeavour, the focus shifted away from the workflow to the survey itself. Future refinements based on the experience acquired during each of these runs will greatly

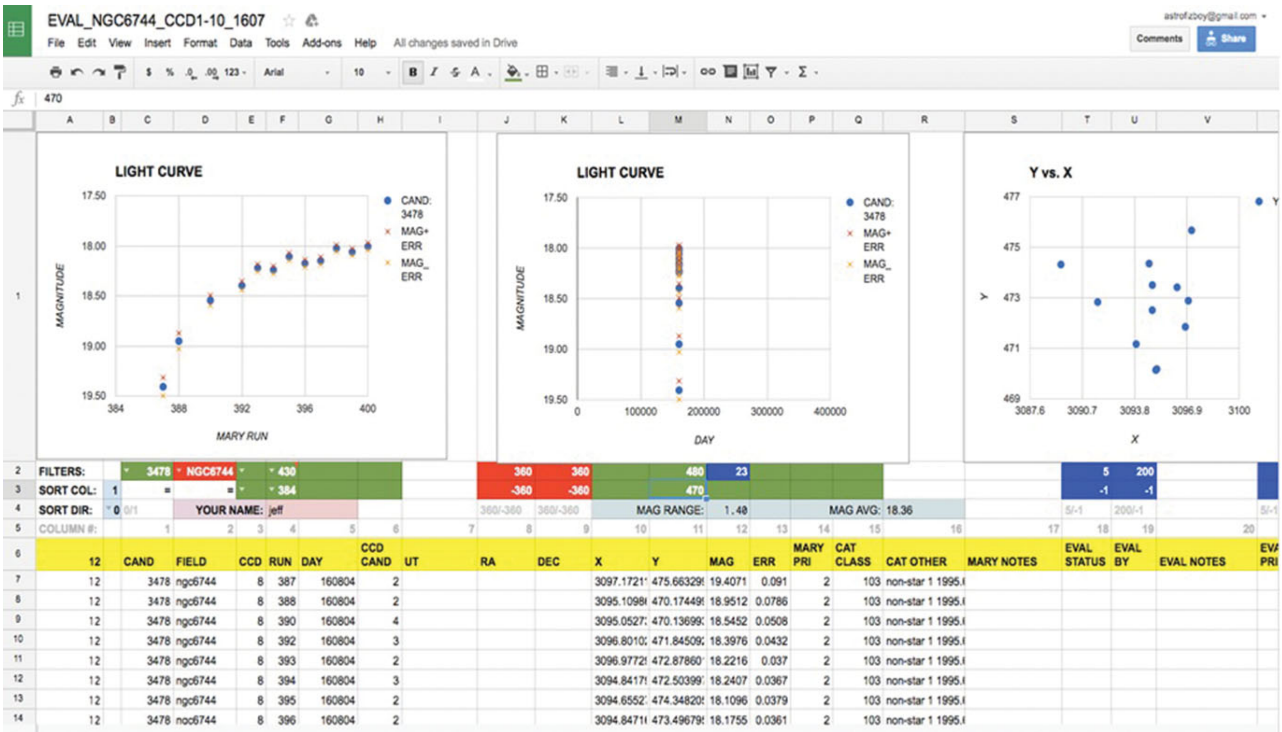


**Figure 11.** The *Deeper, Wider, Faster* online logging tool allowed the inspectors to track the light curves of the potential candidates, their postage stamp images, and candidate positions, magnitudes, and other information. From this tool, the inspectors could report targets of high priority to the principal reviewer for trigger consideration. However, the tool did not have the capability to show the full CCD images. This capability has been added in a later version of the tool.

improve the chances of successful real time, fast follow-up with additional telescopes.

The workflow adopted for O1 and O2 alleviated many of the frustrations associated with the pilot campaign. Establishing a functional display ecology with the ability to display all the relevant content and context simultaneously improved the confidence of the observers that they were getting all the necessary information to make the appropriate decision about candidates. The collaborative environment also improved the observers' experiences during the survey, to the point that reverting to the previous workflow could compromise the purpose of the survey.

The value proposition of advanced display technologies is not always clear. While an argument based on accelerating the time to reach a given scientific outcome is compelling, it is rarely enough to justify the expense on its own. However, it is becoming more relevant to respond rapidly to the influx of new data and science where the data needs to be analysed quickly, to ensure best-use of limited resources. In this work, we have examined improvements to a new programme aimed at detecting fast transients in real time requiring coordination of multiple observatories and astronomers and necessary rapid data analysis. In this context, it was imperative that the *Deeper, Wider, Faster* team was able to make rapid determinations of likely candidates to trigger multi-wavelength imaging and spectroscopic follow-up observations.

### 6.1. Potential improvements

The participating observers responded overwhelmingly positively to the combination of the TDW for O1, the independent workstations for O2, and the curved projection screen in both runs. However, there remain a number of opportunities for improving the display ecology through alternative choices of hardware and software.

A significant improvement would be to eliminate the bezels from the TDW in order to make it easier for the observers to see each entire, unbroken image at full resolution. This could be achieved by using ultra-thin bezel displays (though thin image breaks will still appear) or by using displays that match or exceed the resolution of the images being displayed.

Currently, the closest match to the DECam CCD image size is the 4K standard. Fully compliant 4K screens have a resolution of 4 096 × 2 160 pixels, which exceeds the resolution required to display the individual CCD images. However, these displays and projectors can be expensive. A more viable option would be the consumer version of 4K, more commonly called *ultra HD*. With a resolution of 3 840 × 2 160, these displays are not only more cost effective, but are also very close to the required resolution. In fact the images could be displayed at 96% of full resolution, which should not result in too much degradation due to a small amount of pixel subsampling.

The TDW offered a great deal of promise for the *Deeper, Wider, Faster* project. It addressed the need to be able to display multiple high-resolution images for a short time and then refresh these with new images at a fast pace. It allowed several researchers to search the images in parallel.

However, during O1, software limitations of the TDW were apparent that made it unsuitable for viewing some of the astronomical data needed for *Deeper, Wider, Faster*. In particular, the combination of SAOImage DS9[6] and IRAF[7] was critical to evaluating the software-detected potential candidates, but the TDW software did not provide an adequately performant mechanism to display this content.

There are several alternative software solutions for operating a TDW that were not explored during this campaign. SAGE2 was chosen for the TDW after earlier testing had shown it to be the most suitable for general applications. Meade et al. (2014) discussed solutions such as CGLX[8] and COVISE[9] and their relative shortcomings. Other solutions such as VisTrails[10] were not tested due to the time constraints of the campaign, but would be worth investigating in the future.

However, using a TDW as a fully integrated display was ultimately not the most appropriate use of the infrastructure, as the refinement of the workflow revealed. The improved display ecology for O2 highlighted the value of combining laptops with the new display configuration. Therefore, further investigations of alternative TDW software would have been fruitless.

Bertin, Pillay, & Marmo (2015) discusses alternative methods of dealing with the presentation of large astronomical imagery, which aims to solve the problem of performance of presenting extremely large, remotely stored astronomical images. In the context of *Deeper, Wider, Faster*, this approach might have rendered the transfer of the highly cadenced 4k images unnecessary for the purposes of review, provided a highly stable connection between Chile and Australia could be ensured. However, the *Mary* pipeline running on Swinburne's G2 cluster would still have required the transfer, and the CCD subtraction images were produced by this pipeline, making the transfer unavoidable. Still aspects of this approach are being considered for future runs.

### 6.2. Other applications

After using the environment extensively, several potential astronomical applications for the use of the advanced display environment were identified. These generally include any scenario where very small specific details contained within a very large context are critical to understanding the phenomenon being observed. Examples include the following:

- Comparing absorption features in different transitions in quasar absorption spectra.
- Large galaxy surveys looking for trends in shape and rotation curves.
- Viewing a large number of raw or reduced spectra from multi-object spectrographs to identify unusual objects, place preliminary redshifts, and run redshifting software.
- The TDW could help in the creation of training sets for machine learning software. Viewing thousands of images of real and non-real transient candidates in subtracted images to manually classify them for machine learning training sets would help produce more efficient automated software detections.

The successful use of a TDW as part of a collaborative workspace was consistent with the findings of Meade et al. (2014): physical movement of the eyes, head, and/or whole body was deemed preferable to using a keyboard and/or mouse to pan and zoom. There are several benefits to this approach as follows:

1. It is easier to remember areas of the image already searched.
2. It is easier to maintain a sense of scale of objects being considered as the image scale is consistent and persistent.
3. Physical navigation is often quicker than virtual navigation; improving the time to analyse data.
4. The activity is more stimulating than sitting and viewing in the one direction for prolonged periods; providing both physical relief and exercise.

## 7 CONCLUSION

Establishing a workflow that employs a suitable display ecology combining advanced displays with standard displays has proven essential in advancing the science outcomes of the *Deeper, Wider, Faster* campaign. The advantage of fast cadenced images can quickly become a disadvantage when manual inspection of the individual CCD images is required. The sheer volume of digital information makes it a challenging and cumbersome task for astronomers to achieve using standard, desktop-bound display technologies. We developed a suitable display ecology for postage stamp and CCD image review, and it is clear that without this approach, such a demanding workflow would have been cumbersome and unlikely to have resulted in two successful campaigns.

Dedicated advanced displays, such as a TDW or large-area projection screen, may only solve one part of the image inspection problem. For the *Deeper, Wider, Faster* programme, one display was more appropriate for parallel inspection of the multiple CCD images, while the other was more suited to displaying the numerous postage stamp candidates generated from the *Mary* automated source-detection pipeline. However, it was discovered that when used in conjunction with the online spreadsheet logging tool, independent work-

[6] http://ds9.si.edu/site/Home.html
[7] http://iraf.noao.edu/
[8] http://vis.ucsd.edu/~cglx/
[9] https://www.hlrs.de/en/covise/
[10] https://www.vistrails.org/index.php/Main_Page

stations with sufficient resolution for the CCD image review task was a better option. The use of standard laptops were well suited to interacting with the online spreadsheet. No display was well suited to all tasks and therefore only provided their maximum benefit when used in concert. The most appropriate devices are employed in an efficient manner to make all relevant information available in the most digestible, and actionable, form possible.

When it comes to processing vast amounts of data in useful timeframes, automation has allowed astronomy to advance well beyond human limitations. Despite this, it remains the purview of the astronomer to determine the nature and direction of these advances. Human inspection helps to train the software for better automated results and to place the detections in the larger context. Here, the eyes and experiences of astronomers remains a critical part of the discovery process. Employing the right technology to enhance this capability is every bit as important as deploying more advanced telescopes.

## ACKNOWLEDGEMENTS

The authors thank Nino Colella, Carlo Sgro, and the Learning Environments team (University of Melbourne) for their support in using the curved projection screen, Ken Hodgson and Tony Mazzei (University of Melbourne) for assistance with building access and security, and Luc Renambot (University of Chicago) and Dr Ian Peake (RMIT) for their technical advice in the use of the SAGE2 environment. We also thank Dr Steven Manos (Director, Research Platform Services at the University of Melbourne) for the use of the OzIPortal TDW and the curved projection screen for the observing campaign. Research support to IA is provided by the Australian Astronomical Observatory (AAO). This research was supported by use of the NeCTAR Research Cloud and by the Melbourne Node at the University of Melbourne. The NeCTAR Research Cloud is a collaborative Australian research platform supported by the National Collaborative Research Infrastructure Strategy.

## REFERENCES

AdlerWeb: NEC Case study 2007, http://www.necdisplay.com/case-study/adler-planetarium/21 (accessed 20 May 2016)
Andrews, C., Endert, A., & North, C. 2010, in Proc. of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10) (New York: ACM Press), 55
Andrews, C., Endert, A., Yost, B., & North, C. 2011, Information Visualization, 10, 341
Ball, R., & North, C. 2005a, in Proc. of the 2005 Conference on Human Factors in Computing Systems (CHI '05) (New York: ACM Press), 1196
Ball, R., & North, C. 2005b, in Human-Computer Interaction-INTERACT 2005, Lecture Notes in Computer Science, Vol. 3585, eds. M.F. Costabile & F. Paternò (Berlin, Heidelberg: Springer), 350
Bertin, E., Pillay, R., & Marmo, C. 2015, A&C, 10, 43
Chung, H., North, C., Joshi, S., & Chen, J. 2015, in IEEE Conference on Visual Analytics Science and Technology (VAST) (Chicago: IEEE), 33
Cruz-Neira, C., Sandin, D. J., DeFanti, T. A., Kenyon, R. V., & Hart, J. C. 1992, Communications of the ACM, 35, 64
DeFanti, T. A., et al. 2009, Future Generation Computer Systems, 25, 114
DeFanti, T. A., et al. 2010, CEJE, 1, 16
Diehl, T. 2012, PhPro, 37, 1332
Febretti, A., et al. 2013, in Proc. SPIE 8649, The Engineering Reality of Virtual Reality, ed. M. Dolinsky & I. E. McDowall (Bellingham: SPIE), 864903
Flaugher, B. L., et al. 2012, in Proc. SPIE 8446, Ground-based and Airborne Instrumentation for Astronomy IV (Bellingham: SPIE), 844611
Fluke, C. J., Bourke, P. D., & O'Donovan, D. 2006, PASA, 23, 12
Fomalont, E. B. 1982, Synthesis Mapping, NRAO Workshop Proceedings 5, eds. A. R. Thompson & L. R. D'Addario (Green Bank: NRAO), lecture 11
Huang, E. M., Mynatt, E. D., & Trimble, J. P. 2006, in Pervasive Computing, eds. K. P. Fishkin, B. Schiele, P. Nixon, & A. Quigley (Berlin, Heidelberg: Springer), 321
ISO/IEC 15444-1:2000 2000, Information technology - JPEG 2000 image coding system - Part 1: Core coding system
Meade, B. F., Fluke, C. J., Manos, S., & Sinnott, R. O. 2014, PASA, 31, 29
Norris, R. P. 1994, in ASP Conf. Ser., Astronomical Data Analysis Software and Systems III, Vol. 61, ed. D. R. Crabtree, R. J. Hanisch, & J. Barnes (San Francisco: ASP), 51
Pietriga, E., et al. 2016, in SPIE Newsroom, 99130W, doi:10.1117/2.1201605.006505
QUT - The Cube 2016, http://www.thecube.qut.edu.au/ (accessed 19 March 2016)
Rest, A., et al. 2005, ApJ, 634, 1103
Rots, A. 1986, Synthesis Imaging, eds. Perley, R. A., Schwab, F. R., & Bridle, A. H. (Green Bank: NRAO), 231
SDOWeb 2015, Video from Solar Dynamics Observatory wows museum visitors | Smithsonian Insider, http://insider.si.edu/2015/04/video-from-solar-dynamics-observatory-wows-museum-visitors/ (accessed 19 March 2016)
Smarr, L., Ford, J., Papadopoulos, P., Fainma, S., DeFanti, T., Brown, M., & Leigh, J. 2005, in Optical Fiber Communication Conference (Washington, DC: Optical Society of America), OWG7
Smarr, L. L., Chien, A. A., DeFanti, T., Leigh, J., & Papadopoulos, P. M. 2003, Communications of the ACM, 46, 58
Vohl, D., Fluke, C. J., & Vernardos, G. 2015, A&C, 12, 200

## 5.3 Lessons learned

The *Deeper, Wider, Faster* program stands as an example of the need to consider a display ecology in the planning of a complex observing program. After the pilot run of the project, it was clear to the principal investigators that access to astronomical instruments and computation was not the most serious bottleneck they faced. The most critical and time-sensitive decisions required to determine follow-up targets were almost impossible to make due to the overwhelming amount of information that needed to be understood in a short time period. The investigators realised that trying to inspect all forms of content on standard displays was not viable. The nature of the workflow required different pathways for different forms of information, and they determined that much of this activity could be processed in parallel, either by computation resources or by human observers.

Parallel observation of images improved the performance of the observers, but the limited screen size of standard displays made virtual navigation essential for the image inspection. Moving to a TDW allowed the observers to not only see the images at full resolution – thereby eliminating the need for virtual navigation – but also to easily collaborate search regions and targets.

Because the principal investigators were forced to consider a display ecology that would support the workflows, and have benefited from the experience, it has become part of the continuous improvement activity of the program.

After the observing runs described in Meade et al. (2017), the Advanced Immersion Environment at the University of Melbourne was shut down due to building renovations. The TDW was dismantled and transferred to Swinburne University of Technology, Hawthorn campus, however the curved projection display was discarded due to the impracticalities of moving it. Improvements in automatic candidate detection shifted the inspection focus to potential candidate thumbnails, reducing the urgency of the full image inspection and hence the dependency on the TDW. To better fit the changing workflow, the TDW has undergone several reconfigurations, and has since been replaced by a 100 Megapixel TDW with considerably more graphical processing power that cost a tenth of the price. This new TDW at Swinburne University of Technology forms a key component of an expanded display ecology. Located within the Remote Telescope Operations Centre, with partial funding from the Eric Ormond Baker Charitable Fund, the facility was commissioned in 2018 to support the *Deeper, Wider, Faster* program over the next three years (see Figure 5.1).

The lessons learned during the *Deeper, Wider, Faster* observing runs have supported the development of the automatic detection pipeline used (Andreoni & Cooke, 2018). De-

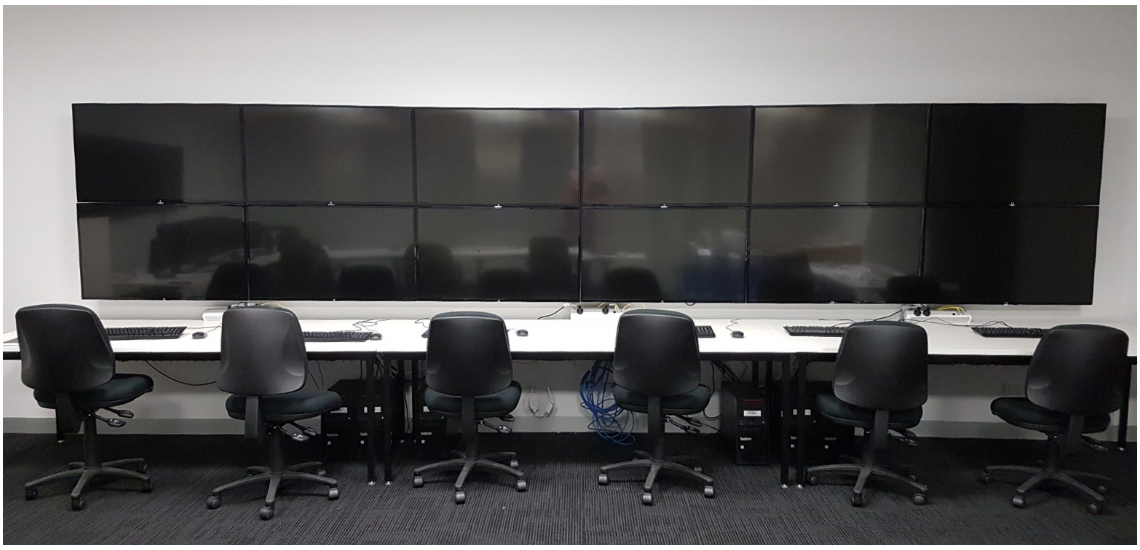

Figure 5.1 This custom display ecology was purpose-built at Swinburne University of Technology to support the *Deeper, Wider, Faster* program.

spite advances in automation, especially with revolutionary techniques like Deep Learning[1] (Chen & Lin, 2014; Schafer, 2017), there will continue to be a role for visual inspection by astronomers for the foreseeable future.

Of key importance is the need to apply these learnings from the *Deeper, Wider, Faster* observing runs to emerging astronomical endeavours such as the SKA and its pathfinder projects. It will be necessary for SKA scientists to work similarly with visualisation experts to prepare and validate workflows and solutions that will scale effectively and maximise the scientific return of the data captured by such facilities.

[1] `http://deeplearning.net/`

# 6

# Using Cloud Computing to Support Virtual Hosted Desktops

## 6.1 Overview

The big data challenge confronting astronomy has forced a shift in the workflow of the modern astronomer. Data is often gathered from remote instruments that can be many orders of magnitude larger than a desktop computer can actually store and process. Simulations are often performed on remote HPC clusters, harnessing the power of hundreds and even thousands of processing cores, and again, producing an overwhelming amount of data. Yet many of the software applications that exist to aid understanding and presenting astronomical data are bound to a graphical desktop interface. However, this desktop is no longer bound to a single physical computer. It is now possible to employ cloud services to provision both computation and storage solutions, but it is increasingly desirable and possible to also host virtual desktops.

This Chapter investigates the opportunity in overcoming the limitations of the physical desktop by moving to a cloud-based virtual desktop. Section 6.2 discusses the use of the Nectar Research Cloud in Australia to foster data communities that form around high-value datasets, hosted in data centres around the country, along with highly flexible Infrastructure as a Service offering. Section 6.3 bridges the gap between TDWs as discussed in Meade et al. (2014, see Chapter 4, Section 4.3) and the emergence of GPU-enabled VHDs. It also considers the fundamental dependence on the network infrastructure in Australia, in particular the network connecting the national peak computing facilities including the Pawsey Supercomputing Facility in Western Australia, and the National Computation Infrastructure in the Australia Capital Territory. Finally, Section 6.5 takes a closer look at the viability of VHDs for astronomy, by conducting a human study [using the UP and UE approaches described in Lam et al. (2012)] with practicing astronomers and astronomy students.

## 6.2 Research Cloud Data Communities

Since it's launch in 2012, the Nectar Research Cloud in Australia has grown to over 12,000 users and nearly 30,000 virtual processing cores[1]. The strong uptake in the research community reflects the growing need for a flexible approach to research computing and data storage. In particular, the uptake in the astronomy community was positive, in part because the astronomy community was already supported by several institutional and national HPC facilities, and used to the concept of remote computing. A case in point is the TAO, which operates as part of the ASVO, a virtual laboratory funded by the Nectar program to support several astronomy programs in Australia. The aim of Meade et al. (2013) was to identify the state of the technology and it relevance to astronomy computing challenges. The paper focused on the formation of *data communities* around particular high value datasets and collections, such as might be captured by astronomical instruments like the Australian Square Kilometer Array Pathfinder ASKAP[2] or SkyMapper[3], and the potential to enhance collaborative research through the sharing of computation and data resources.

This section comprises content published in the paper "Research Cloud Data Communities" by Meade, B., Fluke, C., Sinnott, R., Manos, S., Sinnott, R., van der Knijff, D., & Tseng, A. (2013, May). Paper presented at the THETA 2013 Conference, Hobart, Australia., 34., reproduced with permission from CAUDIT (Council of Australian University Directors of Information Technology).

[1] According to `https://status.rc.nectar.org.au/growth/infrastructure/`, accessed on the 18/03/2018

[2] `https://www.atnf.csiro.au/projects/askap/index.html`

[3] `http://skymapper.anu.edu.au/`

# Research Cloud Data Communities

Bernard Meade[1,2,*], Steven Manos[2], Richard Sinnott[2], Christopher Fluke[1], Dirk van der Knijff[2], Andy Tseng[2]
[1]Swinburne University of Technology
[2] The University of Melbourne
*Corresponding author email: bmeade@unimelb.edu.au

Big Data, big science, the data deluge, these are topics we are hearing about more and more in our research pursuits. Then, through media hype, comes cloud computing, the saviour that is going to resolve our Big Data issues. However, it is difficult to pinpoint exactly what researchers can actually do with data and with clouds, how they get to exactly solve their Big Data problems, and how they get help in using these relatively new tools and infrastructure.

Since the beginning of 2012, the NeCTAR Research Cloud has been running at the University of Melbourne, attracting over 1,650 users from around the country. This has not only provided an unprecedented opportunity for researchers to employ clouds in their research, but it has also given us an opportunity to clearly understand how researchers can more easily solve their Big Data problems. The cloud is now used daily, from running web servers and blog sites, through to hosting virtual laboratories that can automatically create hundreds of servers depending on research demand. Of course, it has also helped us understand that infrastructure isn't everything. There are many other skillsets needed to help researchers from the multitude of disciplines use the cloud effectively.

How can we solve Big Data problems on cloud infrastructure? One of the key aspects are *communities based on research platforms*: Research is built on collaboration, connection and community, and researchers employ platforms daily, whether as bio-imaging platforms, computational platforms or cloud platforms (like DropBox).

There are some important features which enabled this to work.. Firstly, the borders to collaboration are eased, allowing communities to access infrastructure that can be instantly built to be completely open, through to completely closed, all managed securely through (nationally) standardised interfaces. Secondly, it is free and easy to build servers and infrastructure, but it is also cheap to fail, allowing for experimentation not only at a code-level, but at a server or infrastructure level as well. Thirdly, this (virtual) infrastructure can be shared with collaborators, moving the practice of collaboration from sharing papers and code to sharing servers, pre-configured and ready to go. And finally, the underlying infrastructure is built with Big Data in mind, co-located with major data storage infrastructure and high-performance computers, and interconnected with high-speed networks nationally to research instruments.

The research cloud is fundamentally new in that it easily allows communities of researchers, often connected by common geography (research precincts), discipline or long-term established collaborations, to build open, collaborative platforms. These open, sharable, and repeatable platforms encourage coordinated use and development, evolving to common community-oriented methods for Big Data access and data manipulation.

In this paper we discuss in detail critical ingredients in successfully establishing these communities, as well as some outcomes as a result of these communities and their collaboration enabling platforms. We consider astronomy as an exemplar of a research field that has already looked to the cloud as a solution to the ensuing data tsunami.

**Keywords:** Big Data, cloud computing, virtual infrastructure, virtual machines, platforms, communities, discipline-specific support

**Index Terms:** Big Data, The cloudscape

## Introduction

The research landscape is changing rapidly. More and more, we are being confronted by the "Big Data" revolution. Yet research methodologies are sometimes slow to change and it can seem an almost insurmountable challenge to draw meaningful research from the "data deluge". The timely arrival of cloud computing has been held up as a way for researchers to engage with this new data paradigm, providing a simple, efficient way to adopt Big Data into research activities. But the promise and the reality are often separated by a skills chasm.

The NeCTAR (National eResearch Collaboration Tools and Research – www.nectar.org.au) Research Cloud (RC) was launched in February 2012, with the lead node hosted at the University of Melbourne (NeCTARWeb 2012). By 2014, seven more nodes are expected to come online around the country. Over 1,650 researchers have begun using the RC to underpin their research, with several research groups hosting virtual laboratories directly tackling Big Data problems. From web servers and blog sites, through to ad hoc cluster computing, the RC is in active use across Australia.

Each of these research activities helps us understand better how to use the cloud computing infrastructure to address Big Data challenges. It is clear that the two most significant elements are the combination of community and research platforms. Research is not conducted in isolation, but in collaboration. Connection and collaboration technologies are essential elements in both forming and supporting such communities.

The RC gives researchers an opportunity to change the way they engage with Big Data. It is a new way to work and no doubt this will be challenging for many. But the potential benefits of forming communities with Big Data at the core, connected through research precincts or via disciplines, are enormous. Collaborative platforms that are sharable and repeatable, can be open or tightly secured encourage coordinated use and development, fostering community-orientated methods.

In this paper we discuss in detail the specifics of establishing these communities, as well as some of their research outcomes derived from use of collaboration enabling platforms. We start with a general background to the concepts of Big Data and cloud computing, followed by a discussion of the NeCTAR Research Cloud specifically, focusing on those aspects that can benefit data communities, as well as addressing some of the potential risks. Next, we look at Communities and introduce the idea of Virtual Laboratories, highlighting some of the current projects already running on the RC. Using astronomy applications as an example we then discuss cloud computing platforms, followed by a discussion of the relationship between cloud computing and HPC. We also consider the challenges and potential of cloud computing in terms of data management and provenance, as well as the need for effective integration into an institution's IT ecosystem. Finally we discuss what we might expect the research landscape to look like a few years from now.

## Background

New research instruments, sensor networks and computer simulations are producing data at an unprecedented rate. Scientific disciplines, such as astronomy, have been dealing with Big Data challenges for several years. However, the value of Big Data is now being recognised across many more "non-traditional" fields, e.g. the humanities and social sciences.

## What is Big Data?

Big Data means different things to different people, but the generally accepted concept is that the accumulated data exceeds the capacity of typical or traditional processing means. See Table 1 for some examples of Big Data. The size of data collections stored in services such as Research Data Storage Infrastructure (RDSI) will most likely follow a power distribution, where there are a few very large collections (1PB+) such as the LHC, EBI, etc., more biomedical and imaging DBs on the scale of 100's of TB, and then 1000's of smaller - but equally important datasets - such as survey results - in the order of GB's or MB's. This can mean the data volume exceeds the capacity of local databases, or even local hard-drives, or it may mean the data is accumulating too fast for a desktop computer to process. It can also mean the data required is sourced from a variety of repositories, and is heterogeneous in nature. In all cases, Big Data means local storage and manipulation is impractical at best, impossible at worst.

| Resource | Data volume |
|---|---|
| Large Hadron Collider (LHC) | 1TB/second, 13PB in 2010 |
| Human Genome (e.g. European Bioinformatics Institute (EBI)) | 100GB/personal human genome, 30,000 human genomes processed in 2011 |
| Research Data Storage Infrastructure (RDSI) | Expected to exceed 100PB |

Table 1. Examples of Big Data [source: (Brumfiel 2008; "Data, Data Everywhere | The Economist" 2010; "Another Node Announced for Research 'big Data' Project - Research Data Storage Infrastructure - The University of Queensland, Australia" 2012)]

The best use of these expanding networks is to provide access to remote data stores for researchers. To paraphrase a saying, if the data won't come to the computation, then the computation must go to the data. Indeed, using remote computing with services such as VNC (Virtual Network Computing), researchers are provided with an interface to a virtual desktop that operates very much like the one on the local computer. With the explosion of mobile devices such as smart phones and tablets, the performance of the virtual interface is every bit as good on an iPad as it is on the very latest desktop computer, provided sufficient network bandwidth is available.

## What is cloud computing?

Cloud computing offers a way to obtain computing resources on demand, rather than having to commit to potentially unnecessary hardware. It allows an economy of scale to the service provider, and provides consumers with a cost effective way of harnessing the required computing power. For example, by purchasing an amount of computing resource or storage from a cloud provider, a user can ensure that they only pay for what they use, as opposed to a computer under a desk that is paid for whether it is being used or not. The US Federal Government created its "Cloud First" policy to ensure departments investigated the potential of cloud services before investing in IT (Kundra 2011).

Cloud computing is also very attractive for a web service, particularly when the server experiences sporadic loads. For example, a web resource that experiences low usage by students during semester might come under significant strain during exam time. Rather than pay for a high-end computer that can handle the maximum load, and have it sit almost idle for most of the year, a cloud computing hosted virtual web server can exist as a small server costing very little until the demand exceeds a certain level, when additional servers are automatically brought online to cope, instantly balancing the load. This expansion on demand is known as cloud-bursting or elastic cloud as typified by Amazon's Elastic Compute Cloud offerings ("Amazon Elastic Compute Cloud (Amazon EC2), cloud computing Servers" 2013).

Cloud computing also provides an opportunity to test configurations without risk. Launching an instance of a VM typically happens in a matter of minutes and can be terminated just as quickly, making it "cheap to fail". The image can be cloned and modified, launched several or even hundreds of times. It can even provide an ad hoc expansion to an HPC cluster. Images can be used like templates, preconfigured and shared like documents, with links to databases and application already installed, ready to go.

Cloud computing is typically built on big infrastructure, and is therefore ready to handle Big Data. The high-speed interconnects provide excellent access to data stored either adjacent to the compute resource, or via multi gigabit links to other parts of the country, or even the world.

## NeCTAR Research Cloud

The NeCTAR Research Cloud was launched in February 2012 and in its first fifteen months of operation has seen over 1,650 research individuals and more than 110 projects sign on. Berriman et. al. (2010; 2013) provide an excellent summary of cloud computing in scientific workflows when comparing commercial clouds such as Amazon and institutional HPC facilities. However, the NeCTAR RC blazes a new trail for research communities. Rather than weighing the cost benefits of internal resources versus commercial cloud providers, it aims to weigh the value of research opportunities and outcomes against the cost of purchasing and supporting institutional facilities.

There are many directly measurable benefits of cloud computing, and these become even more obvious in the NeCTAR RC context. Initial outlay of capital, operational costs of maintaining space for equipment, power, cooling are easily measured. However, the most significant benefits stem from the fact that hitherto impractical research activities become viable. Many researchers confronted by Big Data are finding new ways to engage with their data, and ultimately produce valuable new research.

### Community benefits of the Research Cloud

There are many benefits to using the RC as opposed to deploying your own infrastructure. Understanding the value of communities around Big Data is key to successfully utilising the RC to extend and enhance collaborations.

1. *Borders to collaboration are eased.* Communities need to be able to share resources, and research collaborations are often national if not global in nature. Fast and efficient sharing of resources, either as infrastructure or information, is essential to the success of these teams. Having the ability to create "instant" computing resources as required, and having full control over the access to that resource, allows researchers to work together no matter where they are in the world, in a secure environment, and to make their work available to a global audience as necessary.
2. *Free and easy.* The NeCTAR RC is free for Australian researchers, allowing them to build virtual servers and infrastructure as required to facilitate their research. This has the benefit of allowing for experimentation, with servers able to be launched and terminated with ease and without penalty. Moving beyond code testing, researchers can now test servers and services in ways that were simply impractical, impossible or simply too expensive before.
3. *Sharing infrastructure.* Perhaps the most exciting aspect of virtual infrastructure is that it can be shared between/across collaborations. In the past, sharing code, systems and results between remote collaborators, writing papers together over long distances has been non-trivial. To develop/integrate code from multiple sources often required researchers to be physically co-located. With RC, virtual servers can be connected to from anywhere in the world by multiple people concurrently. What's more, the actual virtual machine (VM) itself can be shared, cloned and archived. Others can extend the research activity by launching a copy of a preconfigured VM, running simulations or data interrogations with their own parameters. And this can happen in a matter of mere minutes (Hiden et al. 2013).
4. *Big Data Infrastructure.* Today's data centres are built with the capacity to handle Big Data. Physical machines are packed closely together with extremely fast interconnects between them. These racks of machines are in turn connected to high-speed Internet backbones, giving the very best speeds available to other facilities. This greatly exceeds the capability of a typical desktop computer. For many researchers working with Big Data, the proximity of the data to the processing facility is a necessity.

### Risks of the Research Cloud

As with any new technology, there can be significant risks associated with early or insufficiently planned adoption. Cloud computing in general and RC specifically is not a panacea to Big Data difficulties. It is important that institutions and researchers consider their own application before employing the RC for their research (Canon 2011).

1. *Ethical considerations.* Many datasets have strict use controls that limit the way data can be distributed. In some cases, this may preclude storing or transferring the data via public networks.
2. *Security management*. Like any server operating online, there is an onus on the operator to ensure the system cannot be easily compromised and exploited. For many researchers, this will mean employing a system administrator to maintain their servers. The lack of financial barrier to entry may tempt cash strapped researchers without sufficient experience to try to manage their own server, which may result in their systems being compromised. There is also the chance that data stored online might be compromised if the hosts security prevention measures are overcome. In recent years, even high-security services such as those used at financial institutions, have been shown to be not immune to breaches, so it is reasonable to expect that successful attacks will happen for services running on the RC, either through unpatched exploits in the system or inadequate security measures on the VMs themselves.
3. *Network dependence.* While many researchers are already dependent on the presence of a robust network, for cloud computing it is imperative. Large institutions such as the University of Melbourne have high-bandwidth and high-quality network services. However, it is essential that researchers consider the stability of their own environment before committing themselves and their research to the RC. Fortunately, most Universities and research institutions around Australia have excellent network infrastructure, and connectivity to the wider community via broadband networks like the NBN (National Broadband Network) ensures that the reliability and bandwidth of networks will only get better.
4. *Sustainability and technical capabilities.* It is hard to predict the impact of some of the challenges relating to the long-term sustainability of cloud services. At this time, Government funding for the NeCTAR project is uncertain beyond 2014, and the potential for the service to be fully funded by research institutions independently is by no means certain. Sustainability also relies on the persistence of technical capabilities of those creating, operating and maintaining VMs and Virtual Laboratories (VLs). There is a risk that without adequate documentation, once systems are put in place, the processes for establishing new or improved services could be lost.

## Communities

Research communities are the backbone of research. The communities can form around disciplines, institutions, and even methodologies. Communities provide support and form the basis of the peer-review system. The 'dude who knows about computers' is often your PhD student or a postdoc.

In the era of Big Data, communities can also form around datasets and data collection resources and methods. Because the value of the data goes beyond the initial collection motivation, further research based on a dataset or collection of sets is brought about by community awareness. This potential for reuse of data for entirely new research is a key ingredient to justifying expenditure on high-end resources, rather than myriad low-end resources.

## Virtual Laboratories

The NeCTAR RC is aiding the formation of data communities with the VL concept. A VL, also known as a remote laboratory, is an online resource that provides remote access to data collection and analysis tools, and/or data archives. A VL will typically allow resources to be used in very much the same way as if they were stored locally, however, the potential for collaboration is greatly enhanced. Access to the VL is no longer determined by proximity to the computation or the data collection equipment. Processing the data is equally simplified. Table 2 shows some of the current RC Virtual Laboratories.

| Virtual Laboratory | Purpose |
|---|---|
| Virtual Geophysics Laboratory | Scientific workflow portal for Geophysicists |
| Virtual Genomics Laboratory | “Sequence-oriented” genome-related molecular bio-sciences |
| Marine Virtual Laboratory | Marine and ocean-climate science |
| The All Sky Virtual Observatory | “Hardware, tools and services to bring together data from radio telescopes, optical telescopes and supercomputers, covering all parts of the southern sky, under a Virtual Observatory” |
| Climate and Weather Science Laboratory | “Support an intrinsically complex Earth-System Simulator that allows scientists to simulate and analyze climate and weather phenomena.” |
| Humanities Networked Infrastructure (HuNI) | Unlocking and uniting Australia's cultural data |
| Characterisation Virtual Laboratory | Research environments for exploring inner space |

Table 2. NeCTAR Research Cloud Virtual Laboratories [source: (NeCTARWeb 2012)

## Cloud computing platforms for astronomy

As an example of the way RC can support scientific communities, we look to a field where Big Data is already a reality: astronomy. For astronomers, the challenge of coping with new telescopes such as the Square Kilometer Array (SKA) is a real and present concern. While network bandwidths are increasing, astronomers are loath to forego their traditional approach to interrogating data. However, in the next few years, even with significant expansion of bandwidth, the networks will be overwhelmed by the appropriately nicknamed “data tsunami” (Berriman & Groom 2011). Table 3 shows some examples for Big Data in astronomy.

| Resource | Data Volumes |
|---|---|
| Sloan Digital Sky Survey | 357 million unique objects, 15.7TB FITS images, 26.8TB Other data objects, 18TB catalogs |
| Large Synoptic Survey Telescope | Will capture 20TBs/night, 60PBs over ten years |
| Australian Square Kilometer Array Pathfinder | 72Tb/second raw data stream, enough to fill 120 million Blu-Ray discs/day |
| Square Kilometer Array | ~1EB/day (2x global daily internet traffic, 100x Large Hadron Collider data collection) |

Table 3. Examples of Big Data in Astronomy [source: (“SDSS Data Release 7” 2013), (“LSST Data Management | LSST” 2013), (“CSIRO Launches the ASKAP Telescope – and a New Chapter for Radio Astronomy Begins” 2012), (“Amazing Facts - SKA Telescope” 2013)]

Like many disciplines, researchers in astronomy have been confronting the problem of working with datasets that are simply too large to transfer. The Big Data challenge is currently met by remotely processing data using collocated HPC facilities, such as the International Virtual Observatory Alliance (IVOA) (“International Virtual Observatory Alliance” 2013). However, HPC resources are not always a viable solution for many researchers. For one thing, there is a significant learning curve in developing suitable code to run efficiently on such systems.

A model adopted by international facilities like the CyberSKA in Canada ("CyberSKA: Authorized Application Tokens" 2012; Willis 2011) or "OneSpaceNet" from the National Institute of Information and Communications Technology, Japan (NICT) (Morikawa et al. 2010) is that of a portal to a remote processing facility. Similarly in the case of the RC, the portal is created in virtual machine hosted in the Cloud. This portal already has links into both HPC and data storage facilities, often with the two connected with very high-speed interconnects. The user can submit requests via the portal to the HPC system, often with preconfigured widgets, which in turn draws on the data from the connected store, either adjacent to the facility or from wherever it is located on the globe. Only the results of the processing are sent to the researcher (see Figure 1). This methodology has also been adopted by the Canadian Astronomy Data Centre in the form of the Canadian Advanced Network for Astronomical Research (CANFAR) (Ball 2012).

The collaborative potential of this approach is for several researchers to work together to determine the parameters of the request, with the results distributed to each researcher simultaneously. In the case of astronomy, these results could be ultra-high resolution images automatically displayed on remote tiled display walls. Being able to observe the images and discuss the results in real-time, would allow the researchers to refine the parameters and resubmit their query. For astronomers interested in real-time quality control of terabyte-scale radio astronomy data before the raw data gets erased, this may allow for essential refinement of parameters and result in a significantly better scientific outcome.

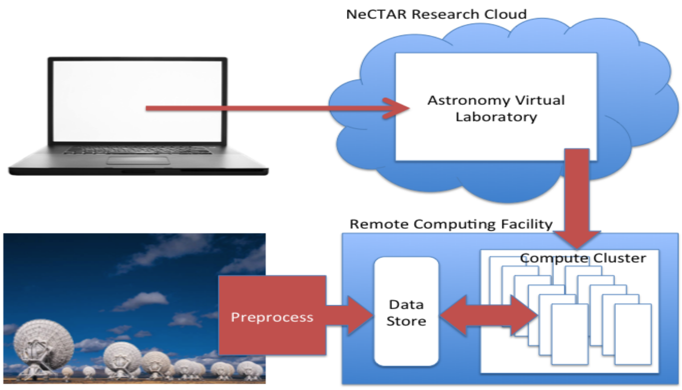


Figure 1: Model for an astronomy virtual laboratory

Recent results from the Space Telescope Science Institute (STScI) shows that more papers are being produced using archived data than from new data ("HST Publication Statistics" 2013; Berriman & Groom 2011). This means the value of the stored data has tipped from validation of research to maximising the scientific return of captured data. These massive datasets can therefore become the core of a research community. The reuse of data increases the potential of research instruments and aids in the justification of expenditure.

Forming communities around data and data-generating instruments, such as telescopes and HPC clusters, is easily facilitated using the Research Cloud. For example, a research group investigating a particular data set,can produce a VM with all their code and links to the dataset in place. This VM can be stored along with data for both provenance and sharing. Another group wishing to extend the original research could clone the VM and conduct new investigations, furthering the original research.

Archiving and provenance as ends in themselves are also better served using VMs that can be backed up and transferred at will. As technology advances, out-dated equipment is typically decommissioned, sometimes to the detriment of being able to reproduce the original environment of the research. With a VM, the virtual environment in which experiments and data were created can be persisted, however this also has challenges that must be overcome, e.g. for how long should they be stored?

## Research Cloud and HPC

HPC can be seen as the forerunner to cloud computing. Rather than utilising local desktop computation resources, HPC allowed users to take advantage of available compute cycles on a massive remote resource. cloud computing achieves a similar outcome. Both HPC systems and cloud computing are based on clusters of computers interconnected by some high-speed network, often managed by a dedicated additional (head) node.

Cloud computing and HPC differ in that HPC systems are predominantly task based whereas cloud computing is more often characterized as Infrastructure as a service (IaaS). On HPC systems, users submit tasks to a queuing system, which then allocates resources to the task as they become available. User tasks all run in the same software environment. cloud computing on the other hand allows the users to develop VMs with their chosen software environment, which they then submit to an allocation system that allocates them the resources they need.

The major differences are that on HPC systems, users are guaranteed exclusive access to the allocated resources for a limited time and sharing is accomplished by having tasks wait on a queue until resources become available, while in the Cloud resources are shared by being oversubscribed, but VMs are allowed to be persistent. This leads to the two systems having different best use situations. HPC, as the name implies, is most suited to well defined and bounded computational problems, whilst Cloud is most suited to ongoing continuous loads. Cloud systems also have the capability to add VMs in a dynamic fashion to cope with varying demand in a way that HPC systems find difficult, and this makes them suited to many collaborative activities where demand is hard to predict (Cohen et al. 2013; Suresh, Ezhilchelvan, and Watson 2013).

## Data management and provenance

As research outcomes becoming more varied and versatile, data management becomes a crucial component of research when dealing with massive datasets. It is essential for research institutions to establish relevant policies and services in order to address these ever-increasing Big Data challenges (Turilli et al. 2013).

Reliability is particularly fundamental when it comes to managing high volumes of research data. Transferring the research data to a trusted cloud environment, that has been set up specifically to accommodate researcher's needs, dramatically reduces the risks of their valuable data being lost or stolen, at the same time lowering the time and resources needed compared to managing data stored in different locations.

For instance, in 2011, the Higher Education Funding Council for England ("Higher Education Funding Council for England (HEFCE)" 2013) has announced the availability of the Universities Modernisation Fund (UMF) to assist UK universities and colleges to take the advantage of the new cloud computing technology to provide more efficient cloud-based services that can be utilised and shared by all research communities.

Three key areas were identified in the UMF initiative:

- *Infrastructure as a Service* (IaaS) offers access to virtual servers, data storage and high-performance computation;
- *Platform as a Service* (PaaS) provides virtual tools for researchers to develop and host individual customised applications; and
- *Software as a Service* (SaaS) enables the users to publish their applications online for easy public access.

Another critical element of cloud-based data management is the data provenance. Data provenance is important because not only does it identify the source or origin of the data, but also ensures its integrity and quality as well. The Open Provenance Model ("The Open Provenance Model" 2013), for example, is a community-driven model providing guidelines on how to allow provenance information to be exchanged between systems which in turn enables developers to build and share tools that operate on the same agreed provenance model.

Cloud computing enables data to be stored and accessed from the very same shared, remote environment as software and computation power. It empowers researchers with a greater control in what they could do with their research data better than they could have imagined which leads to a more productive research experience.

## Cloud as part of an effective institutional IT ecosystem

Sustainable research communities need a good base to be built upon. To tune this base of services to meet the needs of academics is often seen as too challenging. This is understandable as the Research IT environment is quite complex. The customers come from diverse disciplines, each with their own tools, data formats, experience levels and expectations of quality and price ('but DropBox is free?'). Users are geographically dispersed, academics consume collaborations, not services, yet we provide services. Innovation is occurring at breakneck speed elsewhere on the internet, injecting free and easy to use services direct to academics. So, what is the role of the Research Cloud and the institution more broadly in that environment? It is to complement the evolving continuum of services that are provided by local, departmental, faculty, state and national levels, as well as the myriad of other service providers.

However, the final hurdle often remains the incompatibility of the traditional IT helpdesk with researchers. The problem here is that "The very first assumption about an IT helpdesk is that the researcher will know that IT can help them with their problem." The mapping of research problem to IT problem is often the biggest hurdle. This is where growing communities is imperative. They can enable researchers to identify their IT problem more clearly and in context of their discipline, and thereby begin a course of action to solve their problem.

To meaningfully support data communities, IT services need to be made up of a few things to be effective:

1. Good communications, helping researchers understand the benefits in a way that are adapted to discipline-specific audiences and skill levels
2. Community & connection & trainer knowledge
3. Flexible underlying (technical) services that give users *full control* – Academics are very self-sufficient, so enabling them to take ownership and control of their services is key (e.g NeCTAR dashboard).

## Discussion

It seems inevitable that cloud computing will become standard practice, even to the point of overtaking the typical desktop computer. Laptops, tablets and even mobile phones now provide our typical access to the network resources and this will only increase, probably to the point of rendering a local, "anchored to the desk" PC redundant. Our work activities are also shifting to cloud platforms, such as online email, web browsing, journal access and office suites such as Google Docs or MS Office 365. We are already using many cloud platforms and in the future, the seamless integration of these environments will possible (Fransham et al. 2010; Armstrong et al. 2010).

In the next few years, e-Research will have evolved to simply being Research. Researchers will expect a high-bandwidth, "always there" network with simple and efficient access via devices they carry on their person. The data collected for their research will be entirely managed in datacenters across the globe and will be accessible by others in their research community, and beyond. They will also have access to data collected by others, with little difference in procedure between newly collected data and archived material.

When a researcher needs data to drive their research or to support their hypotheses, they will be able to access relevant Big Data stores almost instantly. Where archived data lacks appropriate information, researchers will be able to collect new data from remote facilities, contributing to these online datasets. Research students will be able to complete their research degrees using nothing but archived data. Research communities will collectively decide on the use of limited access facilities such as telescopes, capturing datasets that will satisfy the largest number of research activities. All collected data will also be available to citizen scientists, who in turn will be able to work with research communities to aid the research endeavours.

Big Data and cloud computing will underpin the majority of research activities in the next few years. Whether as primary methods of supporting new research or as supplement, both Big Data and cloud computing will become so ingrained in research methodology and computing in general, that like the "e" in e-Research, they will simply merge into the term, "Research".

## Conclusion

Big Data and cloud computing have already begun to change the research landscape. Researchers have begun to embrace both in an effort to continue to produce cutting edge research. Big facilities like the Pathfinder projects for the Square Kilometre Array and the Large Hadron Collider produce Big Data, but Big Data can also come from sensor networks and crowd-sourced repositories. The volume of data being captured often provides a resource well beyond the original purpose, and it heralds a new way of thinking for many researchers. New skills are needed and this is where communities and the associated platforms are critical to success.

Over the next few years, cloud computing services like NeCTAR RC will prove key to the development of research data communities. With six nodes online by mid-2014, NeCTAR RC will represent a crucial computation resource for a wide variety of projects. Virtual Laboratories from numerous disciplines will exist, with dozens of communities forming around these resources. Communities will develop platforms that will be able to cross disciplines, and make the use of Big Data a natural extension of research activity.

The next few years will provide an opportunity to observe and understand how cloud computing and Big Data changes how researchers work. The combination of community and research platforms will enable far greater collaboration and in turn, better research outcomes. The reuse of platforms and Big Data datasets will be made possible by the ability of cloud computing proliferate customized VMs throughout a research community.

This future will not be without challenges of its own. It is imperative the due diligence be paid to issues such as security and skills development, as well as improving the stability of the underpinning technology. As more research finds its way into the cloud, frailties of the system will be exposed, and will need to be addressed decisively.

While these risks exist and need to be attended, the potential benefits are enormous. The simple fact that Big Data offers such a rich opportunity for research, and is reusable in ways beyond the original purpose, justifies the effort to capture and retain this scale of information. Research communities that form in precincts, around disciplines or even around Big Data, can create collaborative platforms that are shareable and repeatable. Adept users can manage their VMs fully, creating open systems for the wider community to use, or highly secured systems to protect valuable or sensitive data.

The future of cloud computing is all but assured, growing with the same inexorability as the Internet itself has over the last decade. Provided we understand this growth and the opportunities it presents, it can only serve to enrich research as we know it.

## References


Amazing Facts - SKA Telescope. (2013). Retrieved January 17th, 2013 from http://www.skatelescope.org/media-outreach/fun-stuff/facts-figures/.

Amazon Elastic Compute Cloud (Amazon EC2), cloud computing Servers. (2013). Retrieved January 17. http://aws.amazon.com/ec2/.

Another Node Announced for Research 'big Data' Project - Research Data Storage Infrastructure - The University of Queensland, Australia. (2012).Retrieved from http://www.rdsi.uq.edu.au/nodes-announced.

Armstrong, P., A. Agarwal, A. Bishop, A. Charbonneau, R. Desmarais, K. Fransham, N. Hill, I. Gable, S. Gaudet, and S. Goliath. (2010). Cloud Scheduler: a Resource Manager for Distributed Compute Clouds. arXiv Preprint arXiv:1007.0050. Retrieved from http://arxiv.org/abs/1007.0050.

Ball, N. M. (2012). Astroinformatics, cloud computing, and New Science at the Canadian Astronomy Data Centre. In *American Astronomical Society Meeting Abstracts* 219.219:145.11. American Astronomical Society Meeting Abstracts.

Berriman, Bruce, and Steven L. Groom. (201). How Will Astronomy Archives Survive the Data Tsunami? In *Communications of the ACM* 54 (12) (December 1): 52. doi:10.1145/2043174.2043190.

Berriman, G. B., E. Deelman, G. Juve, M. Rynge, and J. S. Vöckler. 2013. The Application of cloud computing to Scientific Workflows: a Study of Cost and Performance. In *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 371(1983). Retrieved from http://rsta.royalsocietypublishing.org/content/371/1983/20120066.short.

Berriman, G. B., G. Juve, E. Deelman, M. Regelson, and P. Plavchan. (2010). The Application of cloud computing to Astronomy: A Study of Cost and Performance. In *e-Science Workshops*, *2010 Sixth IEEE International Conference, 1–7.* Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5693133.

Brumfiel, Geoff. (2008). LHC by the Numbers. In *Nature*, September 9. doi:10.1038/news.2008.1085. Retrieved from http://www.nature.com/doifinder/10.1038/news.2008.1085.

Canon, S. (2011, June). Debunking some common misconceptions of science in the cloud. In Raicu, I Beckman, P and Foster, I (Chairs), *ScienceCloud2011*, 2nd *Workshop on Scientific Cloud Computing*. Workshop co-located with the ACM HPDC 2011 (High Performance Distributed Computing), San Jose, California, June 8th 2011. Retrieved from http://datasys.cs.iit.edu/events/ScienceCloud2011/

Cohen, J., I. Filippis, M. Woodbridge, D. Bauer, N. C. Hong, M. Jackson, S. Butcher, D. Colling, J. Darlington, and B. Fuchs. (2013). RAPPORT: Running Scientific High-performance Computing Applications on the Cloud. In *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 371 (1983). Retrieved from http://rsta.royalsocietypublishing.org/content/371/1983/20120073.short.

CSIRO Launches the ASKAP Telescope – and a New Chapter for Radio Astronomy Begins. (2012).Retrieved from http://theconversation.edu.au/csiro-launches-the-askap-telescope-and-a-new-chapter-for-radio-astronomy-begins-9991.

CyberSKA: Authorized Application Tokens. (2012) Retrieved February 10, 2013 from. http://www.cyberska.org/pg/oauth/catalogue?offset=0.

Data, Data Everywhere (2010) In| *The Economis*t. Retrieved from http://www.economist.com/node/15557443.

Fransham, K., A. Agarwal, P. Armstrong, A. Bishop, A. Charbonneau, R. Desmarais, N. Hill, et al. (2010). Research Computing in a Distributed Cloud Environment. In *Journal of Physics: Conference Series* 256 (1): 012003.

Hiden, H., S. Woodman, P. Watson, and J. Cala. 2013. Developing Cloud Applications Using the e-Science Central Platform. In *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 371 (1983). Retrieved from http://rsta.royalsocietypublishing.org/content/371/1983/20120085.short.

Higher Education Funding Council for England (HEFCE). (2013). Retrieved January 17, 2013 from http://www.hefce.ac.uk/.

HST Publication Statistics. (2013). Retrieved January 17, 2013 from. http://archive.stsci.edu/hst/bibliography/pubstat.html.

International Virtual Observatory Alliance. (2013). Retrieved January 17, 2013 from http://www.ivoa.net/.

Kundra, V. (2011). Federal cloud computing Strategy. White House,[Chief Information Officers Council] Retrieved from http://www.theresearchpedia.com/sites/default/files/Federal%20Cloud%20Computing%20Strategy.pdf.

LSST Data Management | LSST. (2013). Retrieved January 17, 2013 from. http://www.lsst.org/lsst/science/concept_data.

Morikawa, Y., K. T. Murata, S. Watari, H. Kato, K. Yamamoto, S. Inoue, K. Tsubouchi, et al. (2010). A Science Cloud: OneSpaceNet. In *AGU Fall Meeting Abstracts*. December: D5.

NeCTARWeb. (2012). Home | NeCTAR. Retrieved from http://www.nectar.org.au/home.

SDSS Data Release 7. (2013). Retrieved January 17, 2013 from. http://www.sdss.org/dr7/.

Suresh, Visalakshmi, Paul Ezhilchelvan, and Paul Watson. (2013). Scalable and Responsive Event Processing in the Cloud. In *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*. 371. doi:10.1098/rsta.2012.0095. Retrieved from http://rsta.royalsocietypublishing.org/content/371/1983/20120095.abstract.

The Open Provenance Model. (2013). Retrieved January 17, 2013 from. http://openprovenance.org/.

Turilli, Matteo, David Wallom, Chris Williams, Steve Gough, Neal Curran, Richard Tarrant, Dan Bretherton, et al. (2013). Flexible Services for the Support of Research. In *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Science.* 371.

doi:10.1098/rsta.2012.0067. Retrieved from http://rsta.royalsocietypublishing.org/content/371/1983/20120067.abstract.

Willis, A.G. (2011). The Canadian CyberSKA Project. Presented at the 19th Annual Meeting of Astronomy and Astrophysics, May 24, Aveiro, Portugal.



UPDATE 23/06/2018: On page 6 we incorrectly stated the role of the IVOA. The role of the IVOA is to develop standards for interoperability to allow astronomers to discover distributed datasets through common interfaces.

## 6.3 Seeing the Big Picture: A Digital Desktop for Researchers

To move astronomers "beyond the desktop" requires an environment that emulates the user experience that astronomers have come to expect and have developed their workflows around. However, GPU-enabled cloud-based desktops were not widely available in 2015, and this technology only existed in other virtualisation platforms. For example, the VMWare Horizon[4] product provided VDI, which was essentially another term for VHD. Combining VMWare Horizon with Citrix XenApp and XenDesktop[5] on a powerful Dell server containing Virtual Desktop Infrastructure (VDI) graphics cards provided a VHD experience that was almost indistinguishable from an equivalently configured desktop workstation.

The investigation compared the use of a single Ultra HD display[6] to a standard desktop display using the same conditions as described in Meade et al. (2014, see Chapter 4, Section 4.3). It also investigated the use of a VHD provisioned on a VMWare server, linking the use of TDWs in astronomy to cloud-based desktops.

Finally, the paper considered the networking capacity and reliability to sustain VHDs provisioned from a remote peak facility, where high-value datasets were hosted. The paper identified that if the network is suitable and the VHD appropriately resourced, then there was no longer a need to transfer the datasets to a local computer for processing.

This section comprises content published in the paper "Seeing the Big Picture: A Digital Desktop for Researchers" by Meade, B., Fluke, C., Sinnott, R., Manos, S., Killeen, N., Mignone, P., & Wang, M. (2015, May). Paper presented at the THETA 2015 Conference, Gold Coast, Australia., 34., reproduced with permission from CAUDIT (Council of Australian University Directors of Information Technology).

[4] `https://www.vmware.com/au/products/horizon.html`
[5] `https://www.citrix.com.au/products/xenapp-xendesktop/`
[6] Ultra HD = 3840 × 2160 pixels

# Seeing the Big Picture: A Digital Desktop for Researchers

**Bernard Meade[1,2], Christopher Fluke[1], Richard Sinnott[2], Steven Manos[2], Neil Killeen[2], Paul Mignone[2], and Michael Wang[3]**
[1]Swinburne University of Technology; [2] The University of Melbourne; [3] NVIDIA Corporation

The rapid increase in size of experimental and simulation data requires researchers to rethink the way they interact with data to discover new knowledge. One of the many challenges of big data is how to support visual inspection of very large datasets. With sophisticated software, extremely large datasets can be reduced to more understandable graphical summaries. However, these data reduction methods can make it difficult to observe unexpected phenomena at the limit of detectability. In the case of very-high resolution images or image collections, it is beneficial to include a manual inspection stage to support and verify automatic detection algorithms. Tiled Display Walls (TDW) provide a valuable aid for such a process, but because of costs and physical size, have been overlooked by many researchers as a viable option. The recent availability of commodity UltraHD screens offers a cost-effective alternative. For desktop-based activities that draw data from several sources, having a display that allows all these items to be displayed simultaneously improves cognitive performance (Ball & North 2005).

A second consideration is how to use TDWs or UltraHD screens effectively for remote collaboration. While networks have become increasingly robust and reliable, the bandwidth is not expanding at the same rate as data collection technologies. Local storage often represents a potential single point of failure and traditional local backup methods are no longer as cost effective as online options. Also, many datasets used by modern researchers exceed the storage capacity of local systems.

Our particular application area of interest is astronomy: where high-resolution images, vastly exceeding the resolution of standard displays, are generated at a rapid pace from new observational facilities. In this paper we discuss the research underpinning the use of TDWs in astronomical research. We consider UltraHD displays as intermediate options between standard desktop displays and TDWs, and discuss the practicalities of using such displays to enhance the typical desktop environment. Finally we test the capabilities of the Australian Academic Research Network (AARNet) in terms of very large file transfers. Transfer tests have shown that for files from one gigabyte to one terabyte, the network scales up approximately linearly, particularly for some parts of the country, such as Canberra to Melbourne, but less so for other places, such as Western Australia to Melbourne. This allows us to put limits on the image size, and interaction speed, for remote collaborative inspection of high-resolution images.

## Introduction

Technological advances in High Performance Computing infrastructures, such as distributed architectures, graphics processing units and cloud computing, have led to a dramatic increase in the volume of data available for scientific purposes. The challenge now is to determine how best to meaningfully interpret these enormous datasets to enhance and advance knowledge discovery.

Visualization is often the key ingredient to understanding data. Presenting information in graphical summaries can help reduce an overwhelming volume to its essence, and provide improved insight into the phenomena being studied. One of the most challenging aspects of this reduction process is to ensure that no salient information is discarded or compressed beyond detection. The reality facing the modern, data-rich researcher is that the number of data points to display vastly exceeds the number of screen pixels available.

While parallel computing has given rise to parallel visualization [e.g. ParaView; (Ahrens et al. 2005)] display technologies have not necessarily kept pace with computer power. As Table 1 shows, typical desktop or portable devices are only able to display images ranging from 2-5 Megapixels in size. Rather than being restricted to a single display, a reasonable alternative is to spread the data across multiple displays. We refer to such a solution as a tiled display wall (TDW).

Table 1
*Screen resolutions of standard displays, UltraHD displays and tiled display walls.*

| **Display** | **Resolution** | **Image Size (Megapixels)** |
|---|---|---|
| Standard desktop display | 1680 x 1050 | 1.7 |
| iPad (with retina display) | 2048 x 1536 | 3.1 |
| Dell UltraSharp desktop display | 2560 x 1600 | 4.1 |
| Macbook Pro | 2880 x 1800 | 5.2 |
| 4K UltraHD display | 3840 x 2160 | 8.3 |
| OzIPortal | 15360 x 6400 | 98.3 |

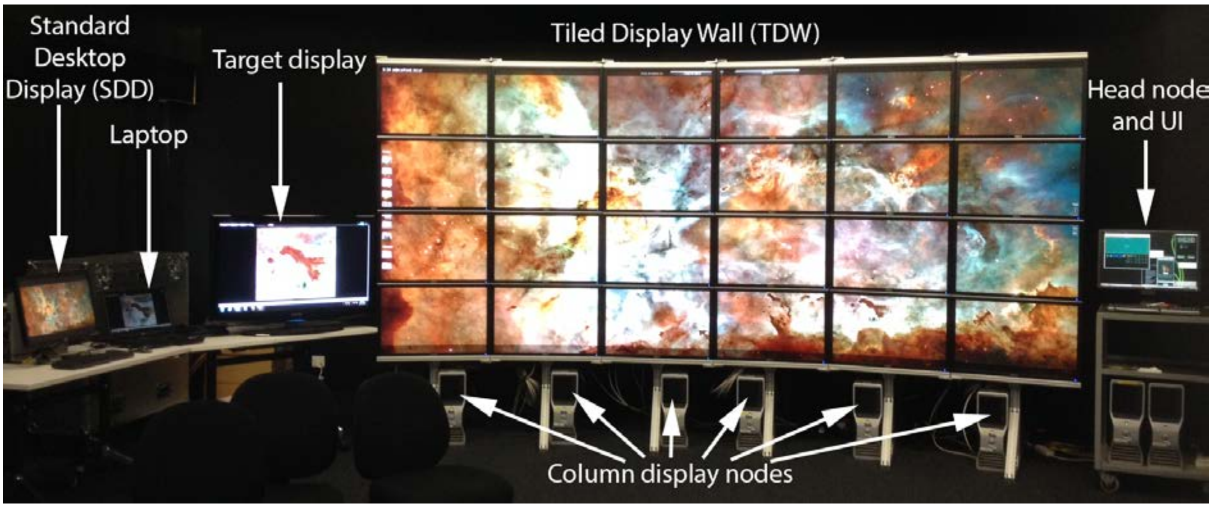


*Figure 1*. Configuration of the OzIPortal tiled display wall used by Meade et al. (2014). The six columns and four rows of Dell UltraSharp displays (2560 x 1600 pixels each) combine to produce a 98.3 Megapixel image (15360 x 6400 pixels). Note the presence of bezels between pairs of displays, which can distract from the visual content. The displayed image is the Carina nebula (image source: http://hubblesite.org/gallery/album/nebula/pr2007016a/hires/true/).

A typical TDW (see Figure 1) comprises a matrix of commodity displays. User interaction is via a head node, which coordinates communication with the individual compute nodes that generate pixel data for a column of displays. First appearing more than a decade ago through initiatives such as the OptIPortal project (DeFanti et al. 2009a), TDWs were expected to become the display of choice for high-resolution data sets. Unfortunately, the uptake of TDW at research institutions has been limited in practice due to their perceived complexity, cost and space requirements.

In the last few years, a new option has appeared: the low-cost, consumer UltraHD display (typical resolution of 3840 x 2160 pixels). While it is possible to use UltraHD displays in a TDW configuration, more value is likely to be obtained by simply making use of the display as a stand-alone, high-resolution desktop display. Costing a few per cent of the price of a TDW, it becomes a far more attractive option for researchers who need more screen real estate. The UltraHD display has the benefit of replacing the multi-display desktop that is increasingly common among researchers, allowing simultaneous heterogeneous content display. Several studies (Czerwinski et al. 2003; Robertson et al. 2005; Ni et al. 2006) have shown performance improvements with typical office-like applications when display windows can be spread out, much the way a traditional desk with paper worked.

TDWs and UltraHD displays (collectively *large-format displays)* can also enhance collaboration between remote colleagues. Many applications now exist to provide effective real-time collaboration, due largely to improved bandwidth and reliability of underlying networks (e.g. Skype, Google Hangouts, Zoom, EVO and others). Using larger displays allows visual presence of a remote colleague via videoconference, as well as shared document workspace and other communication technologies, such as shared desktops and digital whiteboards. This approach is teaching researchers to engage with content that is not directly attached to their own local computer.

As data volumes become increasingly difficult to transfer, the separation of researchers from their data is causing new challenges. One solution is to transfer *some* of data in order for it to be displayed and inspected. A recent development is the Virtual Display Infrastructure (VDI) concept: the delivery of a user's computing desktop or windowed application from a remote service. In this case, a virtual desktop is created on a virtual machine, often hosted in a cloud. Pixels that would ordinarily be directed to an attached display are instead streamed, in real-time, to the remote user's computer/display. With suitable bandwidth and graphics processing capabilities of the VDI host, the experience for the user can be smooth enough that no distinction can be made between a local and remote interface.

In this work, we road-test the UltraHD display. We compare its suitability for large-format image inspection with the investigation of TDWs in Meade et al. (2014). We describe and demonstrate how VDI can be used to deliver content to an UltraHD display as a requirement for remote collaboration. Ultimately, the ability to make use of remote collaboration and VDI depends on the underlying bandwidth. We examine whether the existing national research infrastructure in Australia – the Australian Academic Research Network (AARNet) and the National e-Research Collaboration, Tools and Resources (NeCTAR) Research Cloud – has the capabilities required for remote display of UltraHD video streams in real-time.

We find that the research network is robust between well-established centres such as National Computational Infrastructure in Canberra and the University of Melbourne, but is considerably less stable between the Pawsey Supercomputing Centre in Western Australia and the University of Melbourne. We consider the implications for researchers in a bandwidth-limited environment.

## Background

Visual representations have played an important role in helping researchers engage with their data, recognizing important elements and trends that lead to new knowledge (Fluke et al. 2006). Astronomy is a scientific discipline where this is particularly true. Astronomy has traditionally been a visual science, both in terms of the way that it involves the collection of images, and in the role that visual inspection of data has played in identifying anomalies, image-based artefacts and for knowledge discovery.

Existing and next generation cameras and detectors will take astronomy deeper into the realm of "big data" (see examples in Table 2). Working at the exascale (Quinn et al. 2015), the Square Kilometer Array will produce approximately 1 exabyte of data per day, about 10 times the global internet traffic (SKAFactsWeb). As the quantity, resolution and rate of astronomical images grows, astronomers will increasingly rely on fully automated calibration and analysis pipelines. It is expected that data mining techniques will play a significant role in many new astronomical discoveries (Ball & Brunner 2010).

While a data-mining algorithm can make a discovery, it still requires an astronomer to explain it. Invariably, this requires an understanding of an individual object of interest and its environment. For example, the evolutionary history of an individual galaxy is strongly affected by the density of material that surrounds it – galaxies living in isolation experience very different lives to those at the centre of a gravitationally bound cluster of many hundreds of galaxies [e.g. (Peng et al. 2010)]. As such, manual inspection of extremely high-resolution images or image sets remains an important step for verification of the automated processes (i.e. quality control) and for discovering unexpected phenomena at the limits of instrument sensitivity.

Table 2
*Typical image sizes for existing and proposed astronomical cameras.*

| **Facility** | **Image Size (Megapixels)** | **Reference** |
|---|---|---|
| HST Advanced Camera for Surveys | 16 | (ACSWeb 2005) |
| Dark Energy Camera | 520 | (Mohr et al. 2012) |
| Subaru Hyper Suprime-Cam | 870 | (HyperSuprimeCamWeb 2011) |

As Tables 1 and 2 demonstrate [see also Table 1 of Meade et al. (2014)], there is a clear mismatch between the resolution of the data and that of the typical desktop-based display resources used by the majority of astronomers. There is considerable support for the value of multiple displays, and more and more computer users are working in a multiple display environment. However, there remains an opportunity to make more use of non-standard display environments that could greatly enhance visual inspection, collaboration, communication and training (Fluke et al. 2006). Of particular relevance are tiled display walls and UltraHD displays.

## Tiled Display Walls

TDW comprise a collection of individual, standard desktop displays linked to compute nodes that are coordinated by a head node. The displays are arranged in a tiled configuration, either as a flat or nearly flat wall, or in an arrangement that surrounds the user, such as the CAVE2 (Febretti et al. 2013) or the StarCAVE (DeFanti et al. 2009b). The combined resolution of the individual displays can produce display environments of more than 300 megapixels (OptipresenceWeb 2009), although 50-100 Megapixels is more common.

Management software such as Scalable Adaptive Graphics Environment (SAGE; (SAGEWeb 2012) provides an efficient, easy-to-use method for displaying a variety of multimedia content. SAGE also makes connecting remote walls possible, enabling improved collaboration and control of remote displays. TDWs need not be the same physical arrangement or scale to be connected. Moreover, a user at a TDW with many screens can easily and effectively collaborate with a colleague using a single screen.

While the use of commodity displays and compute nodes drives the dollar per pixel cost down, large TDWs can cost well over AU$100,000, and require considerable cluster management expertise to run the underlying infrastructure. Additionally, they are not suited for installation in a typical office area, requiring dedicated spaces to house them. Both limited physical location options and the high upfront capital costs have resulted in a relatively small number of TDWs being rolled out to Australian universities. Another challenge for TDWs is to get researchers to incorporate them into their research activity, which is often impractical if access is limited.

## UltraHD displays

A cost effective intermediate step is to use UltraHD screens. These displays can be purchased off-the-shelf for less than AU$1000, though better quality devices are closer to AU$3000. With a resolution of 3840 x 2160 pixels, or four times the resolution of Full HD (1920 x 1080), a 55” display provides approximately the same pixel density as a high quality standard display, and the visual display space of four 24” desktop monitors. This also eliminates the distraction of bezels when multiple monitors are placed in a tiled configuration.

As a single display running on a single machine, no additional software is required to coordinate across machines. However, a suitable graphics card, such as a modern professional card (e.g. NVIDIA Quadro K2200), is required to drive the Ultra HD display. This often means the built-in graphics capabilities of mini, all-in-one and laptop computers is insufficient. Fortunately, any recent desktop computer capable of housing a standard PCIe graphics card should be sufficient to provide excellent performance (i.e. up to UltraHD resolution @ 60 fps).

### Large-format displays in astronomy

Meade et al. (2014) undertook the first detailed study on the use of large-format displays in astronomy through a series of "visual source-finding" experiments. In these experiments, participants were presented with images created to match the resolution of the University of Melbourne's 98 Megapixel OzIPortal TDW: 15360x6400 pixels (a matrix of 24 screens with 2560x1600 pixels per screen – see Figure 1). These images included extremely high-resolution astronomical images of galaxy clusters and gaseous nebulae taken from the Hubble Space Telescope, as well as a word field comprising well-known English words.

The purpose of these experiments was to discover if participants could find objects of varying size (larger being easier) within a set time period. Here, performance referred to a participant's ability to find more and/or smaller images. Comparisons were made between performance on the TDW and a standard desktop display (with 1680 x 1050 pixel resolution). The 57 participants included both astronomer and non-astronomer groups. Overall, both groups showed better performance when using the TDW as opposed to the standard desktop display, with the astronomer group generally performing better than the non-astronomers. Of interest is the result that when working in pairs, non-astronomers performed as well as individual astronomers. The aim was to test the notion that TDWs provide better understanding for extremely large astronomy images or image sets.

The Meade et al. (2014) results were consistent with earlier studies (Ball & North 2005; Ball et al. 2007; Bi & Balakrishnan 2009; Bezerianos & Isenberg 2012; Andrews et al. 2010) suggesting physical navigation of a large image (where the participants had to move their head and/or whole body) was more effective than virtual navigation (where a mouse was used to pan and zoom the image). When searching for small objects, more effective search strategies were produced when the context of the overall image could be maintained. This was especially true when visual inspection was conducted as a collaborative exercise – but it is not always possible to get two astronomers in the same room in front of the same TDW.

Based on the approach taken in Meade et al. (2014), but without the benefit of the original study participants, we repeated the image search experiment using a LG UXD7000 UltraHD display, connected to a Dell Precision T3400 with a NVIDIA Quadro K2200 graphics card. The number of participants (n = 4 and included authors of this current work) was too small to produce a statistically significant comparison, but the results were consistent with the expectation that the UltraHD display would function better than the standard desktop display but not as well as the TDW. Observing the manner of interaction of the participants with the UltraHD display was also consistent with the Meade et al. (2014) observations that virtual navigation was only used when physical navigation techniques had been exhausted. Therefore, the increased display space while maintaining suitable pixel density improved the performance of the participants, at least within the range of display sizes tested (viz. standard 24inch display, UltraHD display, 98 megapixel TDW).

### Virtual Desktop Infrastructure

An increase in resolution at the desktop for an UltraHD display requires a corresponding improvement in the graphics capabilities of the host computer. But an over-powered desktop machine may not be the solution if the full display capabilities (viz. pixels and frame rates) are not required all of the time. An emerging option is the Virtual Desktop Infrastructure (VDI) (Miller & Pegah 2007). VDI formally denotes a completely isolated virtual desktop environment, i.e. one desktop per user. The VDI desktop referred to in this paper is actually delivered using a Remote Desktop Services (RDS) model, through Citrix XenApp. The reason for this is that, at the time of writing, only Citrix XenApp was able to deliver GPU-accelerated applications and desktops capable of being scaled to the native UltraHD resolution on the client display. In order to most closely replicate the 1:1 nature of true VDI, only a single user session existed on the host VM during all remote desktop testing.

For applications that support it, the addition of a recent release graphics processing unit (GPU) card to the host server of a virtual machine provides the necessary power to generate a real-time compressed video stream at sufficient frame rate and resolution to provide an excellent desktop experience. Indeed, when combined with additional CPU cores, the power of the virtual machine can be considerably greater than a typical desktop configuration. For example, we were able to stream a Windows desktop at FullHD to a remote client consuming between ~5 and ~50 Mbps of bandwidth using a server with the following configuration:

- Dell PowerEdge R720 server, with:
    - 2x Intel(R) Xeon(R) CPU E5-2670 @ 2.60GHz (total 16 physical cores, 32 logical);
    - 192GB ECC RAM (32GB allocated to the host VM);
    - 1 TB VMFS5 storage (80 GB dedicated to the host VM); and
    - 2x NVIDIA GRID K2 with 4GB GDDR5 RAM (1 GPU dedicated to the host VM)

On gigabit Ethernet, up to 60fps with imperceptible compression is possible. Applications and desktops are streamed using XenApp 7.6, installed in a Windows Server 2008 R2 host VM. The hypervisor is VMWare vSphere 5.5. Mouse and keyboard events are transmitted to and from the virtual desktop, providing a seamless desktop experience. Furthermore, when streamed to a machine capable of driving an UltraHD display, the resolution is automatically increased accordingly and continues to run at acceptable frame rates, though some frame delay is apparent. Performance of the machine decoding the video stream also affects the display frame rate. Using VDI in this way requires a robust network to maintain the connection, but the payoff is impressive.

By combining the above server configuration with a suitably capable receiving client, we suggest that an acceptable user experience with either a FullHD or UltraHD desktop is possible over a 1Gbps network. There are several benefits of this approach:

1. Most universities opt for a regular life cycle for desktop computer infrastructure, usually around 3 to 5 years. In this model, computers are often purchased that exceed requirements in the first year, are ideal for the second year, and noticeably underpowered in the last years of the life cycle.
2. Additional processing power can be allocated as required. As e-mail, web browsing and office applications typically need relatively little processing power and memory on modern processors, fewer resources need to be consumed. Additional cores and memory can be made available when higher-than-normal processing capabilities are needed.
3. High-performance systems, such as multi-core, large memory and GPU-accelerated virtual machines can be created on demand, and shared between users.
4. Virtual machines can be left in a “powered on” state, so that they can continue processing even while the user is no longer connected.
5. Connections can be made from anywhere a suitable network is available. In many cases, a researcher’s personal home network is considerably slower than that provided by their research institution. As the virtual machine would be hosted in a highly connected data centre, the network performance does not diminish, even when operated from home. This is because the user is only receiving a display stream via their home network.

The ability to use VDI effectively for remote collaboration with large-format displays depends on the existence of a suitable remote processing facility to host and serve virtual machines, and sufficient bandwidth between the local and remote facilities. In the next section, we describe the current state of both of these capabilities in Australia.

## Infrastructure for Remote Collaboration

Australia's national research computing infrastructure comprises (amongst other things):

- High performance computing facilities, including the National Computational Infrastructure (NCI) in the Australian Capital Territory and the Pawsey Supercomputing Centre (Pawsey) in Western Australia;
- The internet backbone provided through the Australian Academic Research Network (AARNet); and
- A growing research-focused Cloud computing capability, most notably offered through the National e-Research Collaboration, Tools and Resources (NeCTAR) program.

For the majority of Australian Higher Education researchers transferring data beyond their home institution, bandwidth is provided by the AARNet. AARNet currently provides a 40 Gbps backbone to most Australian universities, and up to 100 Gbps in some places. Locally, many universities support connectivity at 10 Gbps throughout their campuses, although it usually falls to 1 Gbps to the desktop. At most institutions, this wired network is also supported by 802.11n wifi networks, providing up to 308 Mbps for wireless devices. The reliability of network infrastructure at a university is typically around 95% or better, inclusive of both planned and unplanned outages. Such an environment has allowed research to become a largely "online" activity, with much of the data and reference materials sourced remotely.

The NeCTAR program was established in 2011 as part of the Federal Government's Super Science Initiative. The NeCTAR Research Cloud first came online in February 2012, with the lead node established at the University of Melbourne. Since then, additional nodes have been added at Monash University, University of Tasmania, Australian National University, Queensland University of Technology, e-Research South Australia, Intersect and at Pawsey. The principal capability of the Research Cloud nodes is to provide the Australian research community with free and easy access to computational resources in the form of virtual machines (VM) offering Infrastructure-as-a-Service (IaaS) capability. VMs can be created and terminated with ease, and can be used for a multitude of research purposes.

Individual researchers working at Australian universities can be authenticated via the Australian Access Federation (AAF) and gain access to the Research Cloud dashboard. Trial resources for 30 days consisting of 2 compute cores, 8 GB of RAM and 70 GB of disk space are automatically allocated, however users can request additional merit-based, long-term resources for research projects through NeCTAR. Larger coordinated research efforts have also obtained funding and resources to establish Virtual Laboratories (VL). These typically provide the discipline specific research community with appropriate tools for the given domains. The Research Cloud VMs contribute to the underlying infrastructure for these applications. Some examples of highly successful VLs include the Genomics Virtual Laboratory (GVLWeb), Characterisation Virtual Laboratory (CVL) and the All-Sky Virtual Observatory (ASVOWeb).

Currently the Research Cloud has over 20,000 cores, with more than 17,000 cores in use. Nearly 5,000 users have registered with the Research Cloud, launching over 6,000 VM instances. The Research Cloud uses OpenStack to provide the service and the National Endpoint Status (NES) reports an availability of 99.854% for the component supporting the VMs (called Nova) over the most recent 6-month period (22$^{nd}$ of August, 2014 – 22$^{nd}$ of February, 2015). The overall performance of all components, including monitoring, storage and access security is 99.751% for the same period.

## Testing the Nation's Research Bandwidth

Having established the suitability of TDWs and UltraHD displays as advanced display infrastructures for knowledge discovery, we now turn our attention to the bandwidth required to transfer high-resolution images at reasonable frame-rates. To achieve this, we performed a simple, yet instructive experiment: measuring the time taken to transfer files of known size.

The purpose of the tests was to look at how the network bandwidth and stability would affect a typical researcher trying to retrieve large datasets from remote repositories for visualization on a local high-resolution display. In an ideal situation with a direct connection between sites, a 1 GB file would travel at the maximum speed of 1 Gbps link to the desktop (since it is limited by the 1Gbs connection to the desktop machine). Adding in hops between routers and switches, and allowing for other network traffic, introduces delays and instabilities. The longer the transfer time, the greater impact of both systemic and transient effects. Therefore, transferring a single gigabyte might show acceptable transfer times, yet not provide a reliable indication of scalability or consistency of the network. As the datasets get bigger, the network stability plays a much larger role than bandwidth.

With an emphasis on high-resolution astronomical image collections, two major facilities hosting petabyte-scale astronomy data were chosen to connect to the University of Melbourne:

- NCI hosts the the Skymapper (Keller et al. 2007) dataset through the All-Sky Virtual Observatory project. This facility also houses the ANU node of the NeCTAR Research Cloud.
- Pawsey is the repository for the Australian Square Kilometer Array Pathfinder project. Compute time was provided on the Galaxy supercomputer.

These two sites also provided an opportunity to compare the well-established East coast AARNet connection with the relatively new high-speed link to Western Australia.

As we were interested in exploring the instantaneous bandwidth for transfers between NCI, Pawsey and Melbourne, we generated a set of 10 fiducial files as strings of random numbers using the Linux *dd* command:

*dd if=/dev/urandom of=file-1GB.txt bs=1048576 count=1024*

Here, the input parameter *bs* sets the block-size; equal to 1048576 bytes in this example. The minimum file size was 1 GB, and the maximum file size was 1000 GB, as each site had only 2 TB of storage available.

For simplicity, we performed the transfers with the *scp* protocol. While faster options such as GridFTP and Aspera do exist, they are not typically available to researchers. A short bash script initiated the transfers and the timing of the transfer results subsequently logged.

### Transfer speeds

Figure 2 shows the transfer rates of files from 1 GB to 1000GB using the *scp* transfer protocol between a VM hosted on the NCI node of the NeCTAR Research Cloud and a Linux desktop computer on a 1 Gbps wired connection, connected to a 10 Gbps building switch and the 40 Gbps AARNet border router. The data centre hosting the VM at NCI in Canberra is connected via 40 Gbps border router and 10 Gbps switch to the server running the VM. Each fiducial file was transferred three times (consecutively), and the median transfer time was recorded. The transfer experiment was run continuously over a period of one week, to account for daily variations in the network. As the figure shows, the network between the sites is very stable for all transfers. However, variations are more obvious on the large file transfers as the extended duration makes them subject to greater instabilities. Overall, it appears that the NCI to University of Melbourne network is stable for any file size up to one terabyte.

Running the same experiment from Pawsey to the same desktop at the University of Melbourne showed the connecting networks are much less stable – see Figure 3. It was necessary to restart the experiment several times due to unexpected outages that caused the automatic transfers to fail. The results shown are for a single run rather than a full week, due to several transfer failures during the experiment. Even ignoring outlying data points, it is clear that the network between the Pawsey Centre and the University of Melbourne is not currently as stable as between NCI and the University of Melbourne.

There are several possible reasons to account for the increased instability between Pawsey and the University of Melbourne:

1. The network is relatively new and may not be optimally configured yet;
2. The Galaxy server, located in the Pawsey Centre, has experienced several system errors in recent months, which have contributed to the failed transfer attempts, and may also have contributed to delays even for the successful transfers;
3. The timing of the experiment may have occurred during an atypical period of network instability.

Figure 4 shows the results of transfers of a 1GB file with a five-minute sleep between transfers over the course of a week, from NCI to University of Melbourne and Pawsey Centre to University of Melbourne.

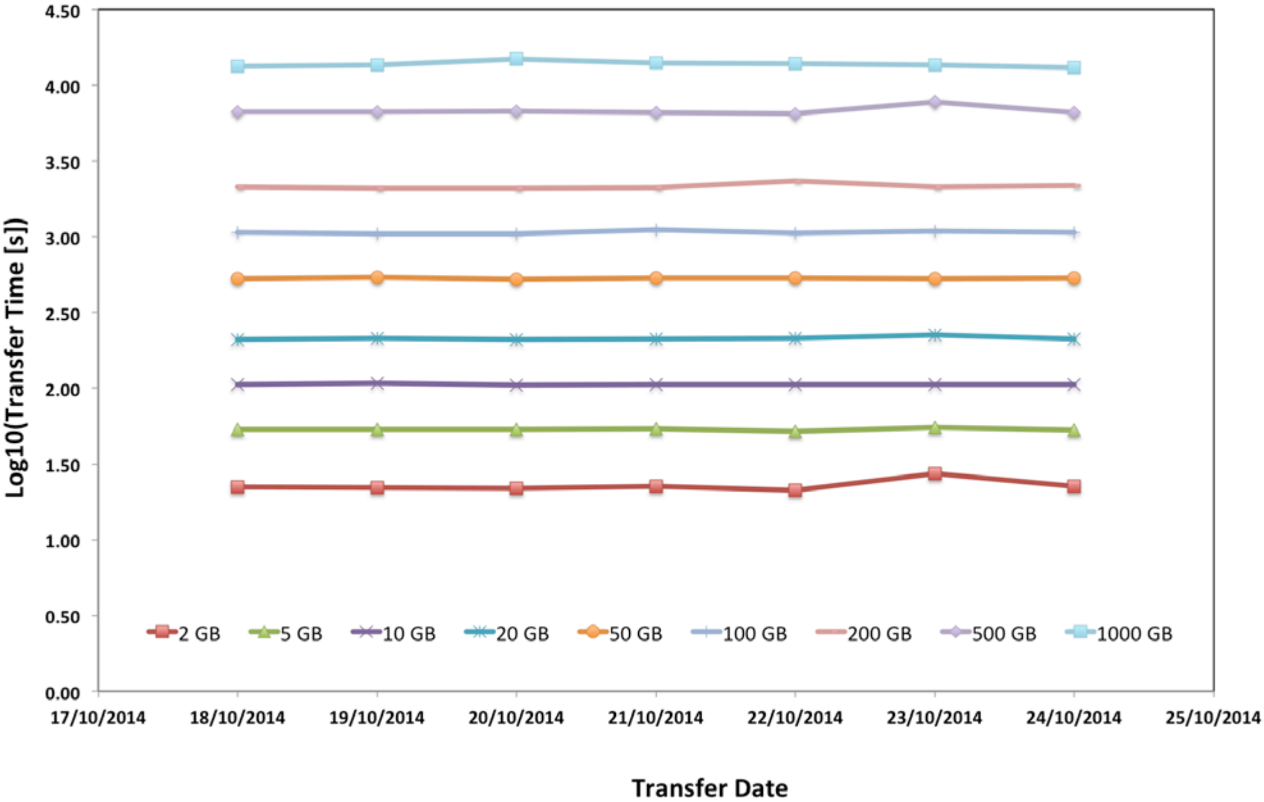


*Figure 2.* Daily data transfer speeds between the NCI and the University of Melbourne. There is little variability in the transfer times during the week for each of the file sizes.

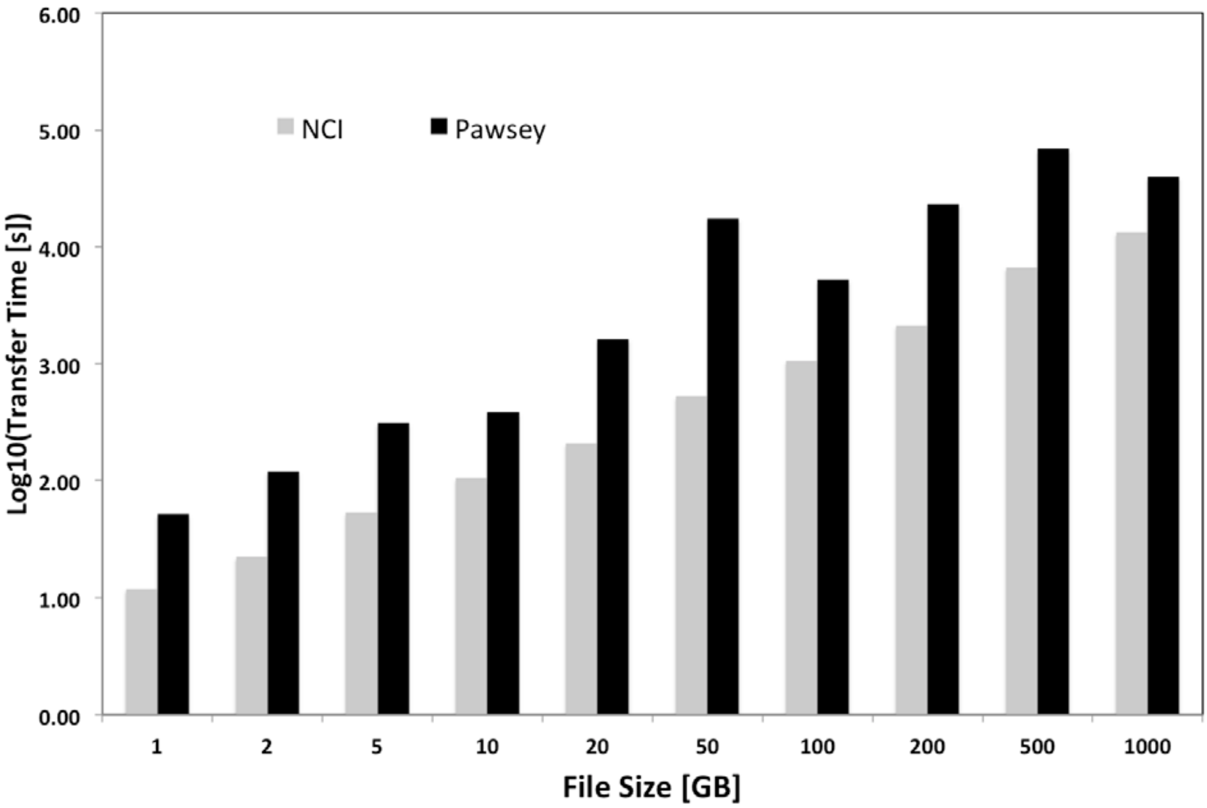


*Figure 3.* Data transfer times to the University of Melbourne from NCI (grey) and Pawsey (black). In all cases, transfer time from Pawsey is longer, and shows more variability.

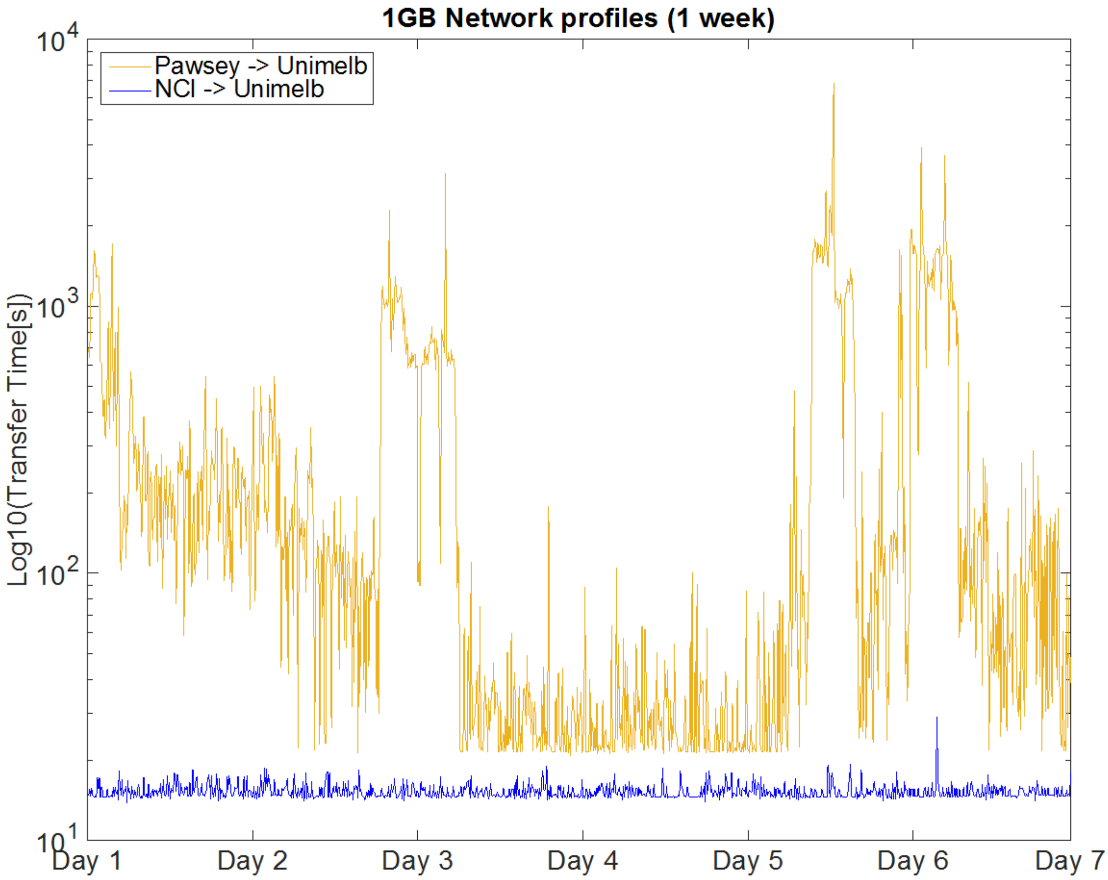


*Figure 4.* Network profile for 1 GB transfers to the University of Melbourne from NCI and Pawsey.

## Conclusion

Combined, the network transfer results provide a compelling case for remote processing, rather than the transfer of large volumes of data to local computers. With a highly stable network such as between National Computational Infrastructure and the University of Melbourne, the growing volume of datasets quickly overwhelms bandwidth and the capabilities of a local computer, both in hard disk storage and RAM and CPU capacity. The situation becomes much more challenging when dealing with additional network instabilities such as those connecting the Pawsey Supercomputing Centre to the rest of Australia.

The stability of the network becomes critical when a fully interactive graphical user interface is being used. The NCI to University of Melbourne link is clearly capable of providing a sustained link of suitable bandwidth for VDI up to UltraHD resolution. However, the link to Pawsey is unlikely to be able to sustain a useable VDI link, and in some cases would be unable to maintain a simple X11 forward. However, this is likely to be a short to medium term issue as the demand for stable network increases and problems are overcome.

Cloud computing has become an essential tool for big data research, and is changing the way research is conducted in many disciplines. Combining computational resources in a shared pool provides far greater performance and economies of scale. But it is also essential to provide solid underlying networks to ensure a high quality user experience, especially in the case of VDI and window forwarding. Fortunately considerable effort is being invested in improving networks across Australia and to international connections. The reality of a fully remote desktop streamed to a local thin client for all researchers is very near.

## Acknowledgments

This work was supported by resources provided by the Pawsey Supercomputing Centre with funding from the Australian Government and the Government of Western Australia. This research was undertaken on the NCI National Facility in Canberra, Australia, which is supported by the Australian Commonwealth Government. We also acknowledge use of computing resources from the NeCTAR Research Cloud http://www.nectar.org.au NeCTAR is an Australian Government project conducted as part of the Super Science initiative and financed by the Education Investment Fund and National Collaborative Research Infrastructure Strategy (NCRIS). We would also like to acknowledge the assistance of Hans Bauer, Dell Computers, for providing the resources for the VDI testing, and Marko Alat (University of Melbourne) for network support.

## References

ACSWeb (2005). *The Advanced Camera for Surveys (ACS)*. Retrieved from http://acs.pha.jhu.edu/

Ahrens, J., Geveci, B., & Law, C. (2005). *The Visualization Handbook, ParaView: An End-User Tool for Large Data Visualization*, ed. CD Hansen & CR Johnson (p. 717). Burlington, MA: Elsevier.

Andrews, C., Endert, A., & North, C. (2010, April). Space to think: large high-resolution displays for sensemaking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 55-64). New York: ACM.

ASVOWeb (2015). *The All-Sky Observatory | NeCTAR*. Retrieved from https://www.nectar.org.au/all-sky-observatory.

Ball, N. M., & Brunner, R. J. (2010). Data mining and machine learning in astronomy. *International Journal of Modern Physics D*, *19*(07), 1049-1106.

Ball, R., & North, C. (2005, April). Effects of tiled high-resolution display on basic visualization and navigation tasks. In *CHI'05 extended abstracts on Human factors in computing systems* (pp. 1196-1199). New York: ACM.

Ball, R., North, C., & Bowman, D. A. (2007, April). Move to improve: promoting physical navigation to increase user performance with large displays. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 191-200). New York: ACM.

Bezerianos, A., & Isenberg, P. (2012). Perception of visual variables on tiled wall-sized displays for information visualization applications. *Visualization and Computer Graphics, IEEE Transactions on*, *18*(12), 2516-2525

Bi, X., & Balakrishnan, R. (2009, April). Comparing usage of a large high-resolution display to single or dual desktop displays for daily work. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1005-1014). New York: ACM.

Czerwinski, M., Smith, G., Regan, T., Meyers, B., Robertson, G., & Starkweather, G. (2003). Toward characterizing the productivity benefits of very large displays. In *Proceedings of INTERACT* (Vol. 3, pp. 9-16).

DeFanti, T. A., Leigh, J., Renambot, L., Jeong, B., Verlo, A., Long, L., ... & Smarr, L. (2009). The OptIPortal, a scalable visualization, storage, and computing interface device for the OptiPuter. *Future Generation Computer Systems*, *25*(2), 114-123.

DeFanti, T. A., Dawe, G., Sandin, D. J., Schulze, J. P., Otto, P., Girado, J., ... & Rao, R. (2009). The StarCAVE, a third-generation CAVE and virtual reality OptIPortal. *Future Generation Computer Systems*, *25*(2), 169-178.

Febretti, A., Nishimoto, A., Thigpen, T., Talandis, J., Long, L., Pirtle, J. D., ... & Leigh, J. (2013, March). CAVE2: a hybrid reality environment for immersive simulation and information analysis. In *IS&T/SPIE Electronic Imaging* (pp. 864903-864903). Bellingham, WA: International Society for Optics and Photonics.

Fluke, C. J., Bourke, P. D., & O'Donovan, D. (2006). Future Directions in Astronomy Visualization. *Publications of the Astronomical Society of Australia*, *23*(1), 12-24.

GVLWeb (2015). *Genomics Virtual Laboratory | NeCTAR*. Retrieved from https://www.nectar.org.au/genomics-virtual-laboratory-0

HyperSuprimeCamWeb (2011). *Hyper Suprime-Cam*. Retrieved from http://www.naoj.org/Projects/HSC/index.html

Keller, S. C., Schmidt, B. P., Bessell, M. S., Conroy, P. G., Francis, P., Granlund, A., ... & Waterson, M. F. (2007). The SkyMapper telescope and the southern sky survey. *Publications of the Astronomical Society of Australia*, *24*(1), 1-12.

Meade, B. F., Fluke, C. J., Manos, S., & Sinnott, R. O. (2014). Are tiled display walls needed for astronomy?. *Publications of the Astronomical Society of Australia*, *31*, e033.

Miller, K., & Pegah, M. (2007, October). Virtualization: virtually at the desktop. In *Proceedings of the 35th annual ACM SIGUCCS fall conference* (pp. 255-260). New York: ACM.

Mohr, J. J., Armstrong, R., Bertin, E., Daues, G. E., Desai, S., Gower, M., ... & Yanny, B. (2012). The Dark Energy Survey Data Processing and Calibration System. *arXiv preprint arXiv:1207.3189*.

Ni, T., Schmidt, G. S., Staadt, O. G., Livingston, M. A., Ball, R., & May, R. (2006, March). A survey of large high-resolution display technologies, techniques, and applications. In *Virtual Reality Conference, 2006* (pp. 223-236). IEEE.

OptipresenceWeb (2009). *Research Projects: OptIPresence Tele-Immersion Testbed – Gravity*. Retrieved from http://vis.ucsd.edu/mediawiki/index.php/Research_Projects:_OptIPresence_Tele-Immersion_Testbed

Peng, Y. J., Lilly, S. J., Kovač, K., Bolzonella, M., Pozzetti, L., Renzini, A., ... & Cassata, P. (2010). Mass and Environment as Drivers of Galaxy Evolution in SDSS and zCOSMOS and the Origin of the Schechter Function. *The Astrophysical Journal*, *721*(1), 193–221.

Quinn, P., Axelrod, T., Bird, I., Dodson, R., Szalay, A., & Wicenec, A. (2015). Delivering SKA Science. *arXiv preprint arXiv:1501.05367*.

Robertson, G., Czerwinski, M., Baudisch, P., Meyers, B., Robbins, D., Smith, G., & Tan, D. (2005). The large-display user experience. *Computer Graphics and Applications, IEEE*, *25*(4), 44-51.

SAGEWeb (2012). SAGE: *Scalable Adaptive Graphics Environment*. Retrieved from http://www.sagecommons.org/

SKAFactsWeb (2015). *Amazing facts - SKA Telescope*. Retrieved from https://www.skatelescope.org/amazingfacts/

## Appendix A

Traceroute results for NCI to University of Melbourne Desktop:

traceroute to 128.250.7.99 (128.250.7.99), 30 hops max, 60 byte packets

1 ncihpchub-vlan-256.nci.org.au (130.56.248.4) 0.364 ms 0.314 ms 0.290 ms

2 182.255.121.17 (182.255.121.17) 0.309 ms 0.285 ms 0.293 ms

3 et-5-3-0.pe1.crlt.vic.aarnet.net.au (113.197.15.22) 7.744 ms 7.726 ms 7.703 ms

4 ae9.bb1.b.mel.aarnet.net.au (113.197.15.97) 7.862 ms 7.832 ms 7.813 ms

5 tengigabitethernet2-1.er2.unimelb.cpe.aarnet.net.au (202.158.200.99) 7.791 ms 7.772 ms 7.844 ms

6 gw1.er2.unimelb.cpe.aarnet.net.au (202.158.206.162) 20.547 ms * 11.707 ms

7 * * *

8 * * *

9 * * *

10 * * *

11 128.250.7.66 (128.250.7.66) 8.236 ms 8.360 ms 8.338 ms

12 128.250.7.99 (128.250.7.99) 8.281 ms * 8.239 ms

Traceroute results for Pawsey Centre to University of Melbourne Desktop:

traceroute to 128.250.7.99 (128.250.7.99), 30 hops max, 60 byte packets

1 146.118.80.1 (146.118.80.1) 0.346 ms 0.281 ms 0.281 ms

2 * * *

3 146.118.1.89 (146.118.1.89) 1.005 ms 0.919 ms 0.887 ms

4 ivec-bdr1-te1-4.ivec.org (202.8.32.33) 1.620 ms 1.452 ms 2.574 ms

5 wa-bdr1-te4-4.gw.csiro.au (130.116.129.73) 1.685 ms 1.394 ms 1.408 ms

6 tengigabitethernet2-2.er2.csiro.cpe.aarnet.net.au (202.158.198.233) 1.467 ms 1.391 ms 1.411 ms

7 ge-4-0-0.bb1.b.per.aarnet.net.au (202.158.198.49) 1.442 ms 1.416 ms 1.428 ms

8 ge-6-0-0.bb1.a.per.aarnet.net.au (202.158.194.1) 1.794 ms 1.778 ms 1.715 ms

9 ge-4-0-0.bb1.a.adl.aarnet.net.au (202.158.194.8) 28.008 ms 27.884 ms 27.854 ms

10 so-0-1-0.bb1.a.mel.aarnet.net.au (202.158.194.18) 36.885 ms 36.948 ms 36.911 ms

11 xe-0-0-0.er1.unimelb.cpe.aarnet.net.au (202.158.210.26) 36.919 ms 36.917 ms 36.959 ms

12 gw1.er1.unimelb.cpe.aarnet.net.au (202.158.200.250) 37.015 ms 37.135 ms 37.072 ms

13 * * *

14 * * *

15 * * *

16 128.250.7.66 (128.250.7.66) 37.645 ms 37.434 ms 38.935 ms

17 128.250.7.99 (128.250.7.99) 37.903 ms 37.319 ms 37.330 ms

**Corresponding author:** Bernard Meade, bmeade@unimelb.edu.au

**Please cite as:** Meade, B., Fluke, C., Sinnott, R., Manos, S., Killeen, N., Mignone, P., & Wang, M. (2015, May). *Seeing the Big Picture: A Digital Desktop for Researchers*. Paper presented at the THETA 2015 Conference, Gold Coast, Australia.

Note: All published papers are refereed, having undergone a single-blind peer-review process.

## 6.4 UltraHD screens and Tiled Display Walls

VHDs can also be considered for more graphically challenging applications, such as driving UltraHD (3840 × 2160 pixels) displays. Higher resolution and multiple displays are becoming more common as part of the researcher's workstation, and with suitably powerful GPUs in the cloud, along with robust networks, VHDs can be a cost-effective way to exploit these displays. Meade et al. (2015) found that the network bandwidth required to drive a UltraHD using a VHD (using an NVIDIA Grid K2 GPU), was less than 50Mbps. With the cost of some UltraHD displays now below $500, they are becoming viable desktop display options.

TDWs are a powerful way to provide a productive and engaging way to work with big data (Meade et al., 2014, see Chapter 4, Section 4.3), however, they are typically quite expensive. While the cost of the displays themselves has come down considerably, they still require considerable compute power to run. However, as these facilities are typically a shared resource, the workstations are often idle when not being used for the TDW.

Many miniPCs are already capable of driving an UltraHD display at 60Hz, but they lack the horsepower to function as a serious workstation. However, costing only a few hundred dollars and small enough to be glued to the back of a display screen, these devices could be used to provide a TDW for a fraction of the cost of most TDW facilities. For example, for $729 ($229 miniPC and $500 UltraHD) per unit, the 98 megapixel TDW described in Meade et al. (2014, see Chapter 4, Section 4.3) could be constructed for under AUD$9,000.

## 6.5 Evaluating Virtual Hosted Desktops for Graphics-intensive Astronomy

While the use of VHDs has grown in recent years, there is little research focusing on the comparison between a VHD and a standard desktop on a local computer. Anecdotal evidence suggests that when it comes to computation, astronomers are willing to use remote desktops when absolutely necessary, but not as a principal desktop interface. However, as more datasets residing in data centres become impractical or impossible to download to a local machine, the use of expensive workstation computers dedicated to a single user needs to be reviewed.

This paper challenges the commonly held belief that a VHD cannot provide the same user experience as a local installed desktop. More importantly, using a VHD will negatively impact a researcher's work output, either in terms of quality or rate. There are certain scenarios where astronomers are willing to accept these impacts as an unavoidable consequence of undertaking research, such as accessing remote telescopes via VNC, or forwarding X11 Windows applications from HPC clusters, especially where no alternative is made available.

However, following on from the experiences documented in Meade et al. (2015), it was clear that a GPU-enabled cloud-based desktop could compete with a standard desktop computer interface. Working with AWS and the Melbourne Node of the Nectar Research Cloud hosted by the University of Melbourne, a study was conceived to investigate at what point a VHD became an acceptable alternative to an astronomer.

This section comprises content accepted for publication (April 4, 2018) in the paper "Evaluating Virtual Hosted Desktops for Graphics-intensive Astronomy" by Meade, B., & Fluke, C. J. . Astronomy and Computing, reproduced with permission of Astronomy and Computing (Elsevier).

Available online at www.sciencedirect.com

ScienceDirect

Astronomy and Computing 00 (2018) 1–20

# Evaluating virtual hosted desktops for graphics-intensive astronomy

Bernard F. Meade[1,2*], Christopher J. Fluke[1,3]

[1]Centre for Astrophysics & Supercomputing,
Swinburne University of Technology, PO Box 218,
Hawthorn, Victoria, 3122, Australia

[2]Infrastructure Services (Doug McDonell Building),
The University of Melbourne, Victoria 3010, Australia

[3]Advanced Visualisation Laboratory,
Digital Research and Innovation Capability Platform,
Swinburne University of Technology, PO Box 218,
Hawthorn, 3122, Australia

## Abstract

Visualisation of data is critical to understanding astronomical phenomena. Today, many instruments produce datasets that are too big to be downloaded to a local computer, yet many of the visualisation tools used by astronomers are deployed only on desktop computers. Cloud computing is increasingly used to provide a computation and simulation platform in astronomy, but it also offers great potential as a visualisation platform. Virtual hosted desktops, with graphics processing unit (GPU) acceleration, allow interactive, graphics-intensive desktop applications to operate co-located with astronomy datasets stored in remote data centres. By combining benchmarking and user experience testing, with a cohort of 20 astronomers, we investigate the viability of replacing physical desktop computers with virtual hosted desktops. In our work, we compare two Apple MacBook computers (one old and one new, representing hardware and opposite ends of the useful lifetime) with two virtual hosted desktops: one commercial (Amazon Web Services) and one in a private research cloud (the Australian Nectar Research Cloud). For two-dimensional image-based tasks and graphics-intensive three-dimensional operations – typical of astronomy visualisation workflows – we found that benchmarks do not necessarily provide the best indication of performance. When compared to typical laptop computers, virtual hosted desktops can provide a better user experience, even with lower performing graphics cards. We also found that virtual hosted desktops are equally simple to use, provide greater flexibility in choice of configuration, and may actually be a more cost-effective option for typical usage profiles.

*Keywords:*
methods: miscellaneous, cloud computing, graphical user interfaces

## 1. Introduction

Astronomy, as with many other scientific disciplines, is now in the petabyte-data era (Brunner et al., 2001; Borne, 2009; Juric & Tyson, 2012). This growth in the total volume of data is due, in part, to the improvements in resolution that modern instruments and detectors are able to access and record. Alongside this is the increased computational power available for numerical simulations.

Visualisation is a crucial component of knowledge discovery. As both the size and complexity of astronomical data sets continue to grow, the existing paradigm of the astronomer visualising and analysing data at the desktop is being pushed to the limit. The high computational and graphics-intensive requirements for many research workflows now exceed the processing, storage, and memory capabilities available with standard desktop-based solutions (Berriman & Groom, 2011; Hassan & Fluke, 2011).

A compelling option is to move all of the processing requirements away from the desktop to a dedicated remote data centre or into the cloud. Here, on-demand computational resources can be co-located with the data such that computation and analysis can be performed at an appropriate scale.

Choosing the right mix of dedicated compute resources that suit the needs of all users is complex. The availability of cloud services allows for flexibility and experimentation with config-

*Corresponding author
*Email addresses:* bmeade@unimelb.edu.au (Bernard F. Meade[1,2]), cfluke@astro.edu.au (Christopher J. Fluke[1,3])

urations that is not always possible with a fixed-purpose data centre. Cloud computing abstracts the hardware aspects of computing away from the user. This takes away the burden of managing hardware, and allows the user to consume the service like a utility such as electricity or network connectivity. However, there is much that is still unknown, and untested, regarding the suitability, choice of hardware, cost effectiveness, and user experiences afforded by commercial and research clouds for supporting astronomical workflows.

To this end, the Square Kilometer Array (SKA) Telescope organisation[1] and Amazon Web Services (AWS)[2] jointly announced the formation of the "Astrocompute in the Cloud" (Astrocompute, 2015) program in April 2015. This program was proposed as a way to explore potential roles for AWS infrastructure to be used for current astronomy projects, and in the future for SKA-related research and operations. This included opportunities to improve research outcomes through the application of additional on-demand compute power, storage and other capabilities.

### 1.1. The need for a virtual desktop

Allocation and scheduling of computing resources through a prioritised batch queue is the preferred approach for most high-performance tasks. For workflows that are computationally limited, any reduction in the overall processing time is beneficial. The overhead in waiting for a workflow to be executed is amortised by the reduction in wall-time once the job starts. However, many astronomy applications – especially data visualisation tasks – require an interactive, on-demand desktop window interface to operate. Such an option is not always compatible with queued access to remote compute resources.

The paradigm of "moving the computation to the data" applies to both traditional computational tasks for analysis and knowledge discovery, and in the use of Virtual Hosted Desktops (VHDs; see Miller & Pegah (2007)). Here, the astronomer's virtual workspace resides entirely in the cloud, and is unlocked from the reliance on the processing capabilities of a physical desktop. A low powered local computer is only required as a gateway between user inputs (e.g. keyboard and mouse interaction) and streaming of images back to the display. Input requires minimal bandwidth; response speed is limited by the remote processing time and the overhead in returning image data to the display device, which scales linearly with the number of pixels.

VHDs can be provisioned on standard workstation or server computers, with the desktop environment presented as a Graphical User Interface (GUI) application management suite. Alternatively the service can be installed on a virtual machine hosted by a cloud provider.

VHD capability has existed for some time, e.g. via X11 window forwarding, where graphics are remotely rendered and streamed via a connection protocol such as SSH, or through Virtual Network Computing (VNC; Duato et al., 1997). These approaches are usually reserved for circumstances where the performance of the environment itself is not critical. In this way, VHDs have often been used as a last resort due to their bandwidth and latency issues. With the rapid improvements in modern networking, it is now possible to employ a VHD in a manner that is almost indistinguishable to the local desktop.

For graphics-intensive work, the result of using a VHD has not always been satisfactory. This, too, has changed with the advent of graphics processing unit (GPU) acceleration of remote desktops. Indeed, GPU manufacturers such as NVIDIA are creating graphics cards specifically for operation in cloud infrastructures (e.g. NVIDIA Grid K1 and M10)[3]. It is timely, therefore, to explore whether a virtual hosted desktop is a functional replacement for a local computer in astronomy.

### 1.2. Overview

In this work, we investigate the suitability of VHDs for performing visualisation tasks from the domain of astronomy. Combining benchmarks with user experience testing (Lam et al., 2012) through the involvement of a cohort of astronomers, we compare software performance and user experiences between local computers (two generations of Apple Mac laptops) and two VHDs, provisioned by AWS and the Australian National eResearch Collaboration Tools and Resources (Nectar)[4] Research Cloud (NRC). These options cover three potential choices for upgrading a computing environment for use in graphics-intense workflows: buy a new local computer; purchase time through a commercial cloud; or, if the option is available, utilise a national research cloud infrastructure.

One of our motivations is to provide astronomers with the knowledge to make more informed decisions when it comes to investing in either a new physical desktop or a VHD. Choosing an alternative infrastructure requires a consideration of operational factors, suitability, user experience, and financial matters. While subject to change without notice, we compare pricing models (at the time of writing) for both physical hardware and cloud services.

This remainder of this paper is laid out as follows. A review of the previous work done in this research area and an introduction to VHDs is presented in Section 2. In Section 3, we benchmark two laptops and two VHDs, and describe the user experiences in these environments for 2D and 3D astronomy tasks. In Section 4 we discuss the results of the user experience testing, and compare the costs associated with VHDs and laptops. Concluding remarks are made in Section 5.

## 2. Cloud computing and virtual hosted desktops

### 2.1. Cloud computing in astronomy

Cloud computing allows clustered commodity computers to be provisioned in the form of Product-as-a-Service, where the product being consumed could be a database (DBaaS), a development platform (PaaS), or most commonly, infrastructure

[1] `https://skatelescope.org/`
[2] `https://aws.amazon.com/`
[3] `https://www.nvidia.com/en-us/gpu-cloud/`
[4] `https://nectar.org.au/research-cloud/`

(IaaS). Most commercial cloud services provide a mix of these options. On allocation of the resource, the virtual instances can be used to perform a variety of tasks, ranging from scientific computation to running web servers.

Cloud computing for scientific workflows has been investigated by several groups over the last decade or so (e.g. Deelman et al., 2008; Hoffa et al., 2008; Juve et al., 2009). One of the first investigations into the use of cloud specifically in astronomy workflows was published by Berriman et al. (2010) and was extended with more detailed benchmarking in Vöckler et al. (2011). These papers showed that commercial clouds, in these cases AWS, could be used cost effectively to provide substantial ad-hoc computation resources.

Ball (2013) described data mining with the machine learning platform, SkyTree[5], running on CANFAR[6], the cloud computing platform for the Canadian Astronomy Data Centre[7]. This research established that cloud computing was a viable option for certain types of computation in astronomy, in this case, the data mining of a 13 billion object catalogue.

Beyond on-demand computation, cloud computing can provide a suitable platform for visual tools. The Montage Image Mosaic Toolkit, as mentioned in Deelman et al. (2008) and Hoffa et al. (2008), was used on AWS to create a Galactic Plane atlas, which combined data from the 2MASS, GLIMPSE, MIPSGAL, MSX amd WISE sky surveys (Berriman et al., 2013, 2016; Berriman & Good, 2017). These studies found that cloud infrastructure provided increased flexibility in resource provisioning, reducing initial financial outlay and costs overall. However, an increased understanding of service models and cloud resource management was required to effectively use cloud services. For applications such as Montage, with short job runtimes, the cloud approach provided good compute resource utilization, while longer, more computationally intensive jobs were less cost-effective. There is also the risk of resource availability and network connectivity introducing unexpected and indeterminate delays.

Cost-benefit analyses have been conducted in relation to the use of cloud with the SKA pathfinders, such as LOFAR[8] and CHILES[9]. Sabater et al. (2017) used cloud infrastructure to run the LOFAR calibration pipeline, finding the flexibility and ad hoc availability of the cloud provided a better option than traditional on-premise HPC services. Dodson et al. (2016) conducted direct comparisons of the CHILES imaging pipeline using a local cluster, a National Peak cluster (Magnus at the Pawsey Supercomputing Centre[10], Western Australia), and cloud infrastructure from AWS. For both the LOFAR and CHILES projects, the cloud platforms were found to be highly competitive across most measures, where costs such as operations are offset against capital expenditure on a local cluster.

A more general discussion of the taxonomy of cloud service providers can be found in Rimal et al. (2010). Further discussion of cloud, high performance computing and big data, including the impact of virtualization can be found in the PhD Thesis of Younge (2016).

### 2.2. Virtual hosted desktops in astronomy

The provision of a VHD is achieved by connecting a local computer to a remote server, which appears on the local computer as a desktop, with all the pre-installed applications ready to use. A thin client (Nieh et al., 2000; Deboosere et al., 2012) is so called because it requires the local client computer to perform very little computationally, while the power to drive the application comes from the server the client is connected to. A common method of connection is via VNC which uses the Remote Desktop Protocol (RDP; Khalid et al., 2016). Many astronomers are familiar with telescope operations being managed using VNC (e.g. Caton & Hawkins, 2009). More general detail about Desktops-as-a-Service can be found in Bipinchandra et al. (2014) and Khalid et al. (2016).

An early example of the use of cloud-based desktops in astronomy is detailed in Berriman et al. (2012). During the 2012 Carl Sagan Workshops[11] hosted by NExScI[12], Berriman and his team successfully used AWS' Elastic Compute Cluster (EC2) service to provision VHDs for use in training 160 astronomers to reduce and analyze Kepler light curves. Rather than have the participants install and configure the raft of applications required for the workshop, a pre-installed suite was available to connect to via VNC.

Some national peak facilities provision Desktops-as-a-Service as a means to access computation and storage services. ACID (Astronomical and physics Cloud Interactive Desktop) is a suite of desktop applications for astronomers and physicists provided by the Cherenkov Telescope Array (CTA)[13] Science Gateway (Italy) for the research community (Massimino et al., 2014). This service provides a VNC User Interface which is accessible through a web browser. This approach eliminates the need for a VNC client installation on the local device.

Many astronomy applications require three-dimensional (3D) graphics acceleration. Services like MASSIVE[14] augment cloud services with GPUs to support these applications (Goscinski et al., 2015), usually via a "pass through" model, where each GPU is used to support a single virtual machine (Ravi et al., 2011). This allows a virtual machine to direct OpenGL calls to the physical GPU on the host, allowing the virtual machine to run GPU-dependent applications, with very little impact from virtualization.

It is now possible, though not necessarily widely available, for cloud services to virtualize the GPUs in the hosts. For example, a GPU existing in one node can be shared with several of the virtual machines running on that host. Like the core utilization, the virtual GPU is managed to respond to GPU requests coming from each virtual machine (Iserte et al., 2016). This

[5] http://www.skytree.net/
[6] http://www.canfar.net/en/
[7] http://www.cadc-ccda.hia-iha.nrc-cnrc.gc.ca/en/
[8] http://www.lofar.org/
[9] http://chiles.astro.columbia.edu/
[10] https://www.pawsey.org.au/
[11] http://nexsci.caltech.edu/workshop/2012/
[12] http://nexsci.caltech.edu/
[13] https://www.cta-observatory.org/
[14] Multi-modal Australian ScienceS Imaging and Visualization Environment https://www.massive.org.au/

approach can be extended to provide network access to virtual GPUs, allowing virtual machines on other hosts to access GPU resources. This is often used to support general purpose GPU computation for numerical calculations, but can also support GPU integration to VHDs (Hong et al., 2017).

### 2.3. Comparing cloud and physical desktops

A physical computer, in the form of a desktop workstation or laptop computer, is an essential tool for modern astronomy. It is important to ensure that any replacement, such as a VHD, is capable of providing a better value proposition. A dependence on computers means that many astronomers have a high-level of technical computing competence, and are often quite particular when it comes to choosing IT hardware. When purchasing a computer, factors such as number of computer cores, clock speed, RAM and GPU capabilities are all important in making a decision.

The same is true for choosing a virtual machine, whether it is to be used as a VHD or not. However, choosing a virtual machine from a cloud provider typically offers far more potential for customization, unless the researcher is willing and able to build a physical computer from parts. More importantly, making a mistake is far easier to correct with a virtual machine, as the chosen options can be discarded, and a new specification built in its place.

If a virtual hosted desktop is to be considered a viable alternative to a local desktop for graphics-intensive workflows, it must meet certain criteria:

1. It must be a simple process to use the environment. If astronomers find it difficult or impractical to use, or requires significant education to learn how to use, then it is unlikely to be adopted by the community.
2. It must be a smooth experience. Even if the process to use the environment is simple, it must be able to run astronomical and related software smoothly. Low-responsiveness to keyboard and mouse movements, and delays in running applications and loading data will result in a frustrating experience for astronomers.
3. It must be demonstrably cost-effective. Many researchers do not wish to change from an environment that they are familiar with, but if a solution can be shown to be cost effective, or show significant benefits in other tangible ways, they may be more willing to explore it.
4. It must be powerful enough to do the tasks required. Modern astronomy workloads are increasingly demanding, either computationally, data intensively, or visually. The ability to choose a fit-for-purpose compute capability for a specific task is one of the main attractions of cloud computing. If the selected environment is under-powered, a more powerful option can be selected next time. If the environment is more powerful than required, a lower-powered option can be selected. Flexibility means the right resource is available when required, and can be relinquished when it is not.
5. It must be available when required. Astronomers, like most researchers, are turning to portable devices for research and other work. The portability of modern devices allows for work to be conducted in non-traditional settings. This has led to an expectation that research can be conducted anywhere and at anytime. However, a cloud based service has an increased dependency on network connectivity, which may not always be present or sufficient.

In the next section, we describe the performance of VHDs using both quantitative and qualitative measures, by comparing system specifications and benchmarks result, as well as participants' reactions to the environments.

## 3. User experiences with virtual hosted desktops

To date, the majority of the discussion on the usefulness of virtual hosted desktops is based on anecdotal evidence and supposition. To remedy this, we recruited 20 astronomers to participate in user experience testing. This is an established approach to testing the suitability of an environment, application or interface that depends on the effective use by human operators. User experience testing methods include informal evaluation, observations of how software or systems are used – in the field or in controlled settings – and questionnaires [see the taxonomy of evaluation methods in Lam et al. (2012) in the context of information visualisation]. While mindful of limitations based on the number, and experience levels, of participants, user experience methodologies recognise there is value in obtaining immediate, subjective responses from even a small cohort of users. User experience testing approaches can be used in conjunction with more objective measures, such as bench-marking data, to improve the value outcomes when purchasing computational resources (Bevan, 2009; Rampersad et al., 2017).

During the user experience tests, we record two main types of data:

1. Quantitative. By timing how long participants take to complete a task or monitoring frame rates during a task, we gain insight as to which classes of tasks are suited to each of the desktop environments.
2. Qualitative. By asking astronomers to perform typical visualisation tasks with current astronomy software, we are able to use the reflections of the participants to gauge the human experience factors of VHDs.

### 3.1. Properties of local and virtual desktops

Consider two common purchasing scenarios faced by astronomers:

1. I am a new staff member or postgraduate research student. What is the best standard option I can access? In this case, a new device is anticipated.
2. My computer is old, can I get a replacement? In this case, the device in question may be out of warranty and considered suitable for replacement under standard university IT renewal schemes.

Table 1: Technical specifications for the four computing environments used for benchmarks and user experience testing.

| Environment | Operating system | CPU/Cores | RAM | GPU |
|---|---|---|---|---|
| MB13 | MacOS 10.10.5 | Intel i7-3740QM CPU(2.70GHz) x 4 | 16GB | NVIDIA GeForce GT 650M (1024MB) |
| MB17 | MacOS 10.10.11 | Intel i7-6700HQ CPU(2.60GHz) x 4 | 16GB | AMD Radeon Pro 450 (2048MB) |
| NRC VM (mel.gpu-k1.large) | Ubuntu 14.04 | Intel Haswell CPU(2.30GHz) x 4 | 16GB | NVIDIA K1 GPU (256MB) |
| AWS VM (g2.2xlarge) | Ubuntu 14.04 | Intel Xeon CPU(2.60GHz) x 8 | 15GB | NVIDIA GRID GK104 GPU (256MB) |

For our user experience testing, we selected four different computing options that addressed such purchasing scenarios.

As many research institutions use a 3 to 4 year renewal cycle for their computer fleet, a MacBook Pro 2013 (MB13) and a MacBook Pro 2017 (MB17) were used to represent the two lifecycle edge cases of a physical desktop environment. They were chosen because they represent the mid-to-high-end options for researchers purchasing a new laptop in 2013 and 2017. Both MacBooks were equipped with graphics accelerators and have a screen resolution of 2880 × 1800 pixels.

For comparison with these physical desktops, we chose to investigate VHDs offered by one commercial cloud – Amazon Web Services because it is available worldwide, has competitive pricing structures, and a large selection of infrastructure options – and one national research cloud service – the Nectar Research Cloud, which is a private option, available to all Australian researchers.

*3.1.1. Amazon Web Services*

Public cloud providers like AWS offer a fully online service where anyone with a credit card can sign up and start a virtual machine in a matter of minutes.

AWS provides a wide range of pay-for-use computation and data resources accessible over the Internet, along with a number of managed services. The IaaS mode is the most common use, where users request a virtual machine with specific characteristics, such as number of compute cores, amount of RAM, amount of attached storage, and an operating system. A virtual machine is launched on AWS infrastructure in one of several data centres located around the world, and made available to the user via the Internet.

The user connects to the virtual machine via certain protocols, most commonly SSH, and can install any required software. If a windowing (i.e. desktop) environment is installed, the user can also configure the virtual machine to allow VNC connections. With a local desktop VNC client, the user can then connect to the virtual machine desktop and operate the virtual machine as if it were their local desktop, with local keyboard and mouse activity being passed through to the virtual machine.

A full discussion of the pricing models of AWS is beyond the scope of this paper, however, two options are relevant:

1. *On-demand pricing* is the simplest to use and plan for, but is also the most expensive. Using this model, a fixed price for a virtual machine configuration is known and agreed to before the virtual machine is created. Importantly, the On-demand price guarantees the availability of the virtual machine while it is being used. This provides surety and clarity when planning the actual usage of the resource.
2. *Spot pricing* allows AWS to sell the unused capacity in its data centres at a far more attractive rate than the On-demand price. However, as demand on that resource increases, the Spot price will rise. Once the Spot price exceeds the user's bid price, the user's virtual machine will be stopped. The user will need to increase their bid to allow the virtual machine to be restarted.

In this work, we only use AWS On-demand instances, as interactive visualisation workflows require guaranteed and continuous access to the VHD.

*3.1.2. Nectar Research Cloud*

Many research institutions or federations offer private research clouds specifically for their research communities. In Australia and New Zealand, Nectar established the largest private Research Cloud in the Southern Hemisphere. Private research clouds are not as big as public cloud providers like AWS, but they are generally more suited to the demands of research.

Operating in a similar way to AWS, the Nectar Research Cloud is an Australian Federal Government initiative, which commenced operation in 2012. It is designed to support the computation and storage needs of the Australian research community using a federated private research cloud (Meade et al., 2013). The NRC offers over 32,000 cores, distributed between nine physical nodes, and has supported in excess of 10,000 users from the Australian research community.

Access to the NRC is either through a research merit application or via a host institution's private infrastructure. Researchers are not directly charged for their use of the NRC under the merit scheme, and institutions determine their own access model for their private resources. NRC primarily provides IaaS to the Australian research community, though new services continue to come online as the service matures.

The Melbourne Node of the NRC, hosted at the University of Melbourne, offers limited GPU capability for VHDs. A mix of NVIDIA K1s and M10s is provided (though the M10s were not available during our user study), with the GPU-enabled hosts operating in a "pass-through" configuration to the virtual machine.

*3.1.3. Technical specifications*

Table 1 summarises the system specifications for the MB13, the MB17, the NRC virtual machine and the AWS virtual machine.

The cloud virtual machines were chosen based on availability. At the time of the investigation, the AWS virtual machine g2.2xlarge was the cheapest of the available options. The more expensive g2.8xlarge provided considerably more CPU computation power than either of the local computers, and so was not used. The NRC virtual machine chosen was the most closely matched to the AWS virtual machine instance from the available flavours. During the study, the virtual machine environments were run in full screen mode to match the graphical rendering load of the local laptop screen.

Due to licensing requirements from Apple Computer PTY LTD, the Macintosh operating system, MacOS X, is not able to be used in a cloud environment. While the Microsoft Windows license does allow for use in cloud environments, it does not support some of the applications needed for the investigation, so the selected operating system for the VHDs was Ubuntu Linux 14.04 LTS (Trusty Tahr). The MB13 was used at Swinburne University of Technology and the MB17 was used at the University of Melbourne. TurboVNC[15] was used to connect to the VHD. When operated in full screen mode, TurboVNC automatically adjusts screen resolution according to the size of the attached display, which is $2880 \times 1800$ for both the MB13 and MB17.

A network with sufficient bandwidth and stability is critical to providing a persistent connection to the virtual machine supporting the VHD. Both the University of Melbourne and Swinburne University of Technology have substantial network infrastructure, both wired and wireless. At Swinburne University of Technology, the network used was *Eduroam*[16], a federated wireless research network with peering institutions all around Australia. Six of the University of Melbourne participants used the University's wired 1Gbps network, while the remainder used the University's wireless solution, *uniwireless*. The network performance was measured before and after the tasks using the Speedtest website[17], to ensure the network was stable throughout.

To minimize the impact of network latency during the user experience testing, the AWS data centre located in Sydney Australia, was chosen as this is the closest option to the University of Melbourne and Swinburne University of Technology.

Results from these networks tests are shown in Appendix A.

### 3.2. Benchmarking the environments

Computing products are released with technical specifications, which are usually considered objective measures of a product's performance in certain conditions. However, computational environments are complex, and individual components might not be operating in ideal conditions, resulting in less than optimal performance. To accurately determine the true performance of a complete system requires the performance measurement to be conducted under the conditions of intended use.

The most common way to compare systems' performance is to run benchmarks. Benchmarks are only useful if they test a system's capacity in conditions that fully expose the system's limits. While many benchmark options are available, not all are suited to the situation being investigated here. The Unigine Valley and Heaven benchmarks[18] were chosen for heir cross-platform availability, free accessibility and use of highly demanding 3D graphics computation. The Heaven benchmark tests vertex and texture operations as well as lighting effects, while Valley, released later, expands these to include atmospheric effects and performance optimizations. Figure 1 shows the results for each of the four computing environments, with the scores being calculated based on CPU and GPU performance, and maximum, minimum and median frame rates. These results are intended for comparison between the four systems in this investigation, rather than an independent objective measure against any system.

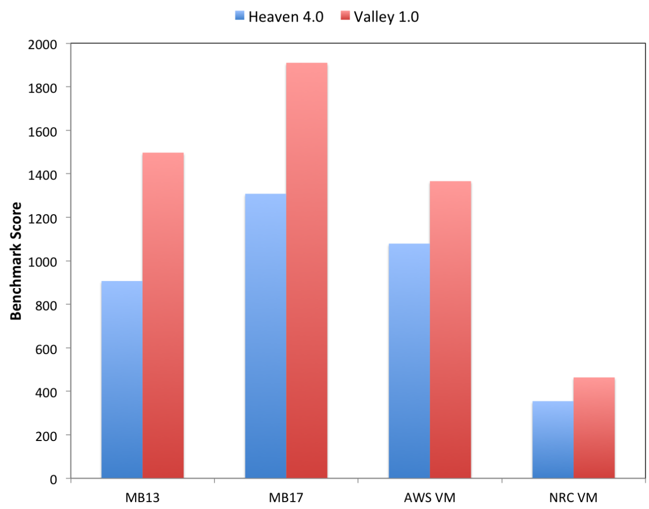


Figure 1: Scores for the Unigine Heaven 4.0 and Valley 1.0 Basic benchmarks obtained for each of the four computing environments. The Valley benchmark includes improvements on the older Heaven benchmark. The score is calculated from a combination of the maximum, minimum and median frame rates, as well as CPU and GPU performance. This graph indicates the MB17 is the most performant system, while the NRC VHD is the least performant.

The local computers performed very well for the Unigine Valley and Heaven benchmarks. The AWS virtual machines performed slightly better than the MB13, but the NRC virtual machine was easily the lowest. This partially aligns with the expected performance based on the currency of the GPU, with the MB13 being the oldest, the AWS GRID GK104 being next, and the MB17 being the newest. The NRC K1 is approximately the same age as the MB13, but is not as performant. In each case, the predicted performance does align with the benchmark results.

Benchmarks provide a valuable method of comparing systems out of context when in context comparison is not available. For the most part, this is an accepted approach to determining the potential suitability of a system for a task that is similar to the benchmarking application. However, this is less useful if the nature of the intended task is uncertain, or exposes an unanticipated demand on the system.

[15] https://www.turbovnc.org/
[16] https://www.eduroam.edu.au/
[17] http://www.speedtest.net
[18] https://benchmark.unigine.com/

Furthermore, a value choice might be informed by benchmark results, but factors such as user experience must also be taken into consideration. For example, there is little value in buying the best graphics card if users are unable to distinguish the difference when compared with a lower performing card.

Having completed benchmarking, we investigated whether user experiences with the physical and virtual hosted desktops were consistent with a benchmark-only approach.

### 3.3. Participants

Because the focus in this investigation is the astronomers' experience, the participant cohort was limited to astronomers, either current academic staff or postgraduate research students (including recently graduated students). Ten participants were recruited from the astronomy department at Swinburne University of Technology and a further ten from the University of Melbourne. Three participants were academic staff and 17 were postgraduate students.

No previous experience of the software or techniques involved was required to complete the tasks. Where previous experience with the software was identified, these participants were encouraged to adhere to the instructions provided, even if they differed from their usual practices.

Each participant was asked to complete a brief interview before the hands-on component of the user investigation was undertaken. This survey was designed to understand the cohort's collective experience of cloud computing in general and VHDs in particular.

From this survey, 60% of the cohort could provide a reasonable definition of cloud computing, with the remaining 40% either unsure or confused it with online storage. 80% were able to provide a reasonable definition of a VHD, though only 20% had used VHD in their research, compared to 60% who had used the cloud (including cloud storage) for their research.

17 of the participants had used X11 Forwarding or VNC (typically for telescope operations) previously, but their experiences ranged from "terrible" to "fantastic". All of the participants said they had experienced limitations when using their local computer for their research, with problems including lack of power, memory and storage.

### 3.4. User experiences

Typical visualisation tasks that astronomers might encounter in a data analysis workflow were presented using the local and virtual desktops. Due to the limited time available with each participant - around 35 to 45 minutes - it was not possible to explore a wide range of applications, or to delve too deeply into the applications chosen. We chose one task that was not overly demanding of the GPU for computation, which simulated basic desktop operations, and one where graphics card performance would be paramount.

In a 2D environment, where no GPU acceleration is required, the task was a simple image alignment. The focus on a specific task, rather than free exploration of the environment, provided a more objective method to evaluate whether or not the environment itself impacted on the completion of the task positively or negatively, as opposed to whether the environment was enjoyable or not. Completing the tasks required the use of application windows, menus, keyboard commands and a mouse. The participants were asked to inspect two FITS images (Abolfathi et al., 2017) with DS9[19], and align them with IRAF[20]. The participants were guided through each step of the task, and then asked to repeat the same task three times, once via a local desktop and either once or twice in a cloud environment.

Many astronomical applications require advanced 3D graphics. While it is a relatively simple matter to monitor a system's performance metrics, this does not necessarily coincide with the perceptions of an astronomer. To determine if such a correlation exists, a GPU-accelerated 3D application called Shwirl (Vohl, 2017) was used to monitor how the environment performed under increasing load, while the participants' perceptions were also recorded. Shwirl uses graphics shaders operating on the GPU to perform interactive, real-time volume rendering of 3D spectral data cubes. For this task, a spectral data cube[21] was loaded into Shwirl and adjustments were made to the volume rendering, mimicking steps in a workflow that might occur in visualisation and analysis of a spectral data cube. After each change was performed, the participant was asked to provide their perception of the performance of the system at that time.

See Appendix B for a more detailed description of steps in the user testing procedure.

### 3.5. Setup

Each participant was presented with a laptop computer with a standard mouse (to avoid possible issues with the use of a trackpad). The starting environment was already loaded, with half the cohort seeing a local desktop first, and the rest seeing a cloud desktop first. The cloud desktops were used in full-screen mode and minimised when not required. Switching between desktops was performed by one of the investigators.

In all cases, to avoid having to train the participants in how to create a virtual machine or use the cloud, the environments were setup in advance, and each was preconfigured in such a way as to avoid the need for the participant to "learn" their way around. Shortcuts were placed on the desktops for the applications, and terminal windows were already running. Participants were presented with a standard desktop environment that closely resembled what they are most likely already used to.

For the NRC virtual machine, a volume mounted disk access issue caused significant delays to the loading of DS9, IRAF and Shwirl, but only for the first time they were run. Subsequent executions did not exhibit the problem. This is a known issue with the GPU nodes of the Melbourne Node of the NRC. To reduce the impact, the NRC virtual machine environment was prepared in advance by doing a first run of each application before the start of the user testing. Hence, the participant experienced the cached version of the application, which closely matched the other environments. As this technical issue does

[19] `http://ds9.si.edu/site/Home.html`

[20] `http://iraf.noao.edu/`

[21] NGC628 from The HI Nearby Galaxy Survey (THINGS; Walter et al., 2008), data from `http://www.mpia.de/THINGS/Data.html`

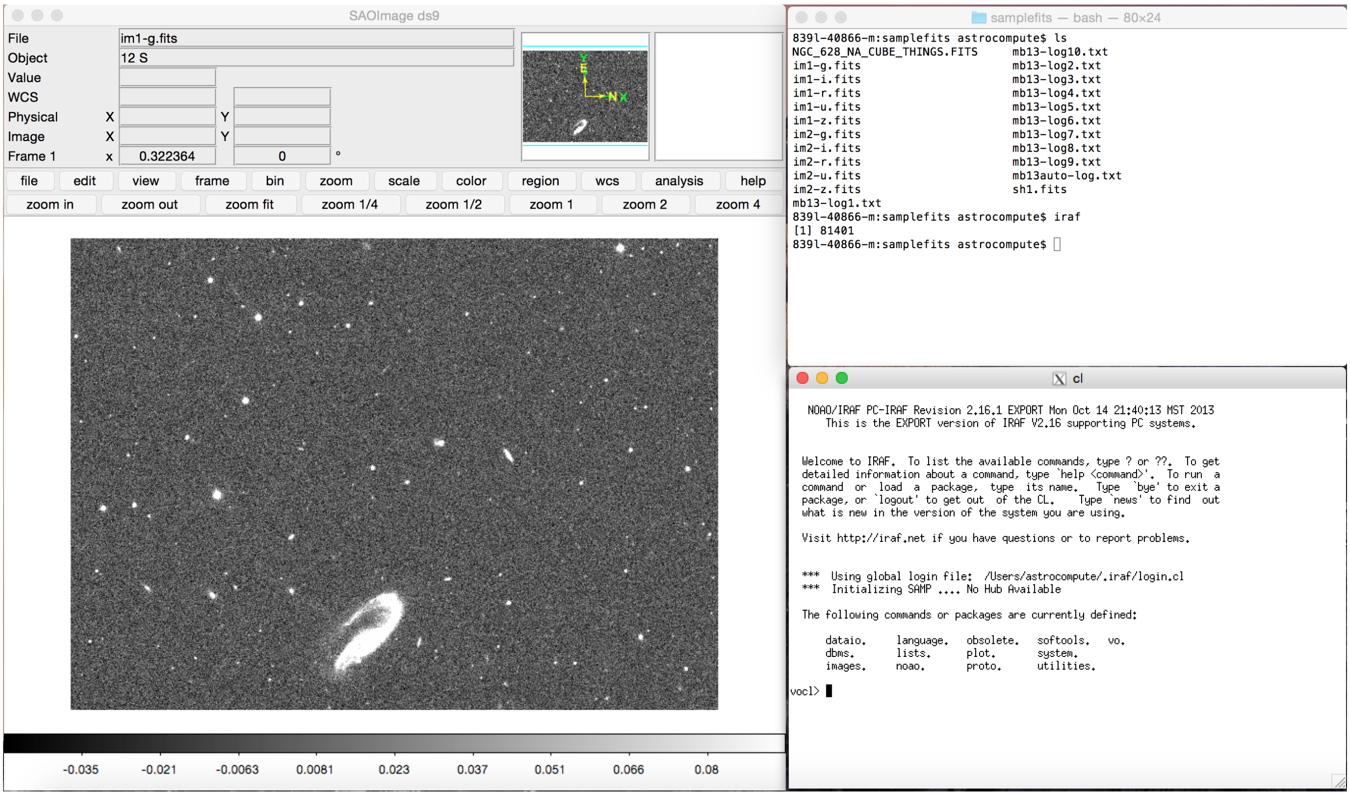


Figure 2: The layout of the display for the 2D image alignment task. The terminal window (upper right) was running before the participant began the investigation. SAOImage DS9 (left) was launched from a desktop icon and IRAF (lower right) was started from the command line.

not impact all NRC virtual machines or applications, it was decided this was the fairest way to compare the environments.

Every effort was made to ensure that the experience of the local and virtual hosted desktops presented to the participants was the same, and that the same set of tasks was completed. Unfortunately, due to time restrictions imposed by working with volunteers, three of the 20 participants were only able to complete one local and one VHD version of the Shwirl task. Additionally, a log file was not recorded for one participant while completing the Shwirl task on the AWS virtual machine.

### 3.6. The 2D image alignment task

The 2D image alignment activity was undertaken first. The participant was provided with a sheet of paper with explicit instructions on how to proceed, including the precise IRAF commands and offsets needed to align the images. The process was timed to ensure the participant remained focused on completing the task, and to provide a means of comparing the change in performance with subsequent runs.

To ensure all participants were familiar with the instructions, the operator guided the participant through every step without the timer running. Once the participant had completed this initial pass, they were then instructed that they would be timed for the next three passes. After some of the participants made minor mistakes that resulted in noticeable delays, it became clear that it was necessary for the operator to intervene if it was apparent that an error was about to occur, such as missing a step, that would result in a significant loss of time. This was

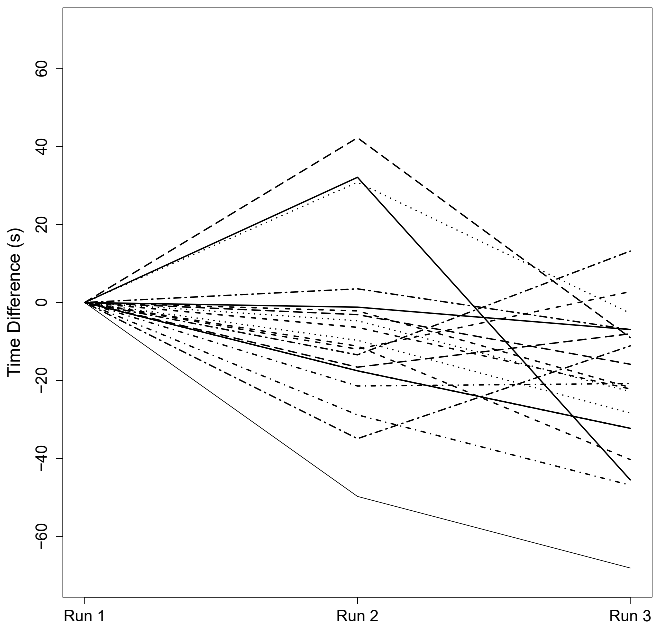


Figure 3: This graph shows the timing change from the first timed run for each participant. All but three participants showed an overall improvement, a few experienced minor problems, which caused delays. However, these delays were caused by minor mistakes made by the participants, and not the environments.

Table 2: Shwirl provides a number of modules that allow the adjustment of the loaded volume to improve visual understanding.

| Module | Purpose | Impact on GPU |
|---|---|---|
| Camera and transforms | Adjust the field of view and scale the loaded volume in X, Y and/or Z direction | Scaling the volume in the Z direction significantly increases the computation load on the GPU |
| Colour | Apply a colour scheme to the loaded volume | No additional calculations are required by the GPU, so impact is negligible |
| Filter | Filter the volume to eliminate data above and below set values | No additional calculations are required by the GPU, so impact is negligible |
| Smooth | Apply smoothing to the volume data by averaging between neighbouring points | As smoothness is increased, the range of neighbouring points increases and requires increasing number of calculations by the GPU |

deemed acceptable because our purpose was not to evaluate the participants' ability to learn a task. See Figure 2 for the screen layout during this activity.

The participants were timed by one of the investigators while completing the task, but were encouraged to simply "follow the instructions" without worrying about the timing. By repeating the task, participants generally improved their time, which suggests that they were learning while working. Such learning effects are expected in task-based user studies. Varying the order in which participants were presented with a local or cloud desktop was necessary to determine if the environment itself impacted on this learning process. Timing the process encouraged the participants to focus on performing the task itself rather than whether the environment lived up to their expectations.

For the 2D image alignment task, the cloud and local environments performed equally well. Figure 3 shows the individual time changes based on each participants' initial time, with most participants improving their performance each time. In each case where a participant saw a decrease in performance (corresponding to an increase in time taken to complete the task), a simple error such as clicking the wrong button or misreading an instruction, was identified as the cause of the problem. In only one case did a participant attribute a loss of performance to the environment, and in that case, it was the lack of familiarity with the local desktop operating system (MacOS X) that was identified as being the problem. One other participant encountered a minor issue where they accidentally switched out of the virtual machine client. In that instance, because it was the first time trial, and no steps had been completed, the task and timer were restarted. The order of presentation of the environments did not appear to have any impact on the performance or the participants.

Based on user testing, we conclude that both the local and virtual hosted desktops were equally suited for tasks of this nature, especially before significant demand was placed on the graphics systems. Participants reported very positive experiences with the VHDs, and many expressed surprise at the level of performance and ease of use for the cloud environments.

### 3.7. The 3D spectral cube rotation task

The 3D spectral cube rotation activity was not timed by hand because the Shwirl software generates a log of the frame rate once per second, linked to the corresponding states of the Shwirl options. This provided a means to compare the system performance with the participants' perception scores. See Figure 4 for the screen layout during this activity.

As the participant completed each section, they gave the environment a score out of 10, where 10 was identified as being "as good as they could possibly want" and 0 being "something they would never willingly use again". As this measure is very subjective, the responses were normalised for each participant, such that the highest score they gave for all three environments became a 10, and the lowest score they gave became a 0. This provided a direct comparison of the participants' perception of their experiences of the two or three environments encountered. These results were then compared with the system performance for each stage of the 3D tasks.

The participants' responses were recorded for each step and environment, along with a Shwirl output log for each. This way, each log could be associated with the corresponding perception scores. Table 2 shows the modules used during the task and how they are used to test the GPU in the environment. More detailed descriptions of the modules can be found in Vohl (2017).

Using Shwirl's auto-rotate feature, the same steps to be used by the participants were applied to create a baseline for the investigation. Figure 5 shows the median frame rate for each environment for each stage of this task. This graph shows that the GPUs in the local laptops initially performed much better than the GPUs in the cloud environments. However, as the load was increased on the GPU, such as in the Z-Scale and the smoothing steps, the performance of the local GPUs dropped more than the cloud environments. Once smoothing was applied at increasing levels, the difference in the performance became almost negligible, with the AWS virtual machine decreasing at a slower rate than the others.

#### 3.7.1. User perceptions of computing environments

Participants naturally have different approaches to assessing an environment, so it was necessary to provide some guidance to ensure some similarity in the assessment. The very first time a participant was asked to provide a ranking for an environment, they typically opted for numbers around 7 or 8, as these reflect the generally positive experience, without having over committed. This approach gave the participant room to give subsequent environments a higher score if they performed better, or lower if they performed worse.

The first two tasks were intended to give the participant an opportunity to calibrate their perception. The first task provided

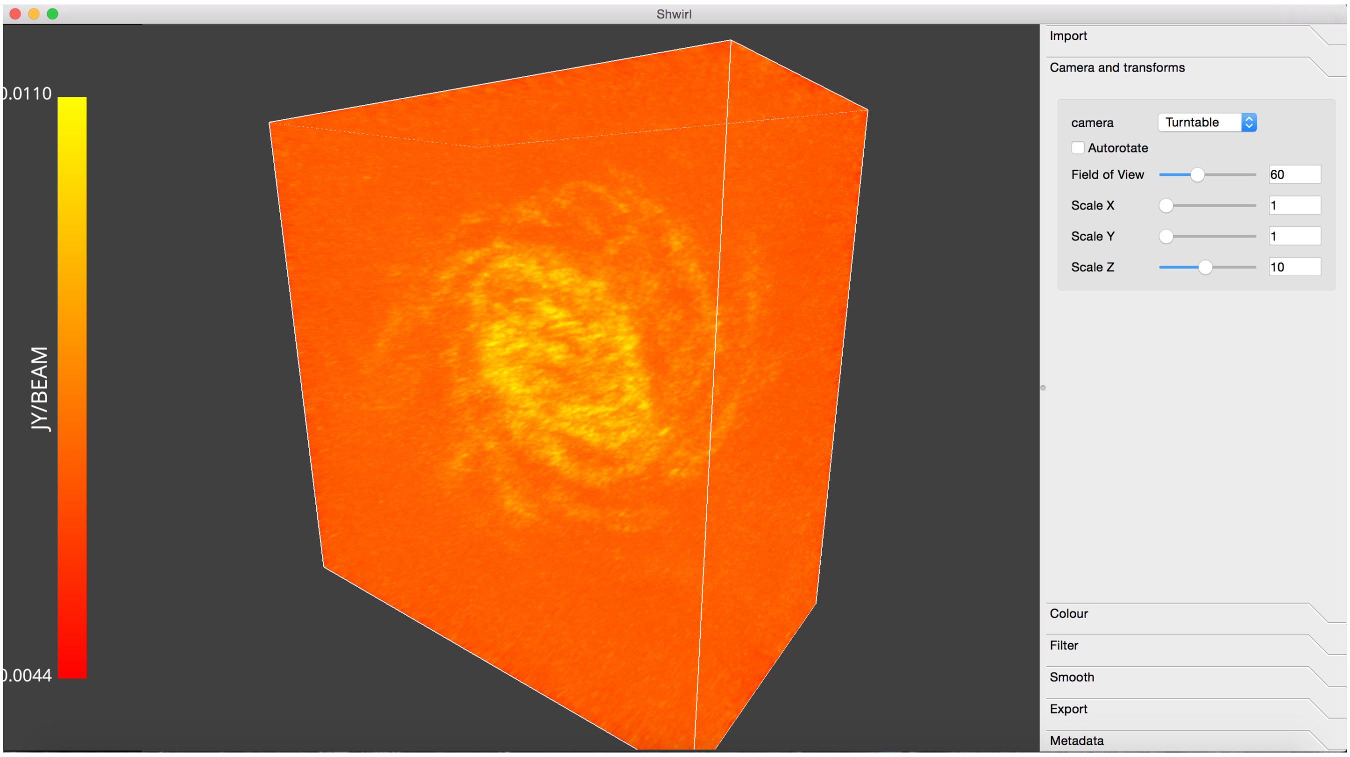


Figure 4: The Shwirl application was launched from the command line. The tabs on the right side of the application are the Shwirl modules (see Table 2), and correspond to the tasks for each step of the 3D activity. This image shows the spectral cube with a scaling of 10 in the Z direction.

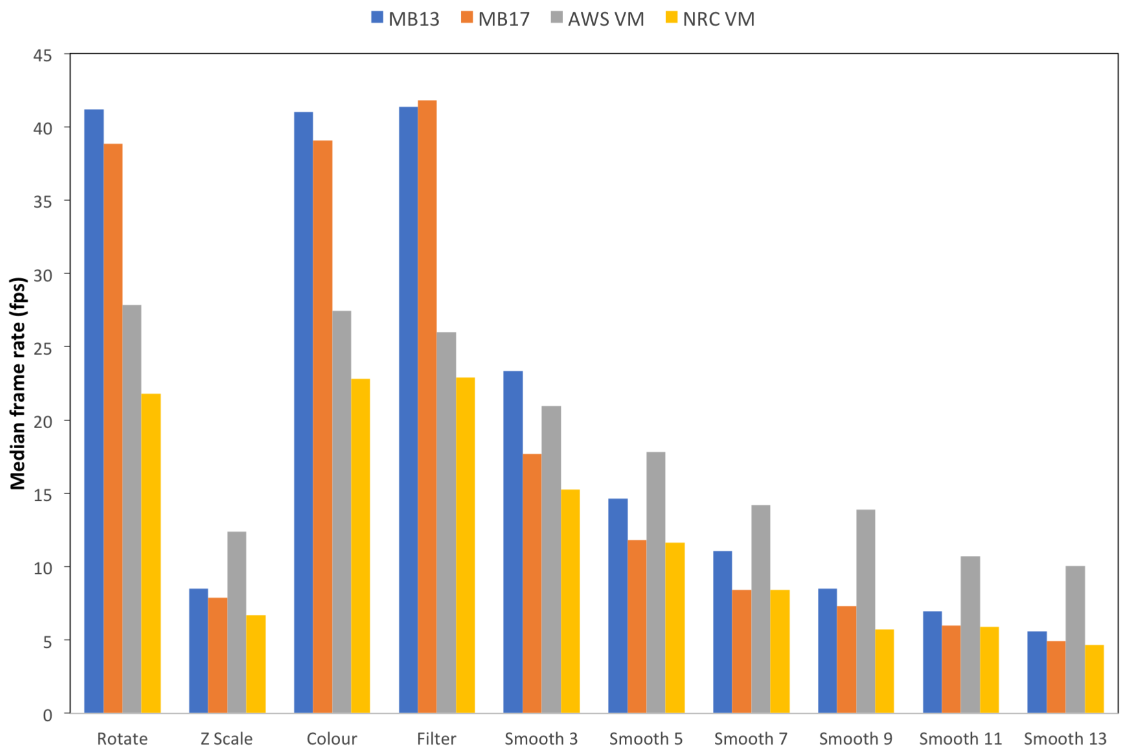


Figure 5: Using the auto rotate feature in Shwirl, each of the steps in the 3D component of the task were completed to establish a benchmark for the platforms. The Rotate step shows the environments without additional computation load, while the Z-Scale shows an increased load. These steps were included to give the participants the opportunity to calibrate their perception. The steps are included here as they correspond to the steps undertaken by the participants.

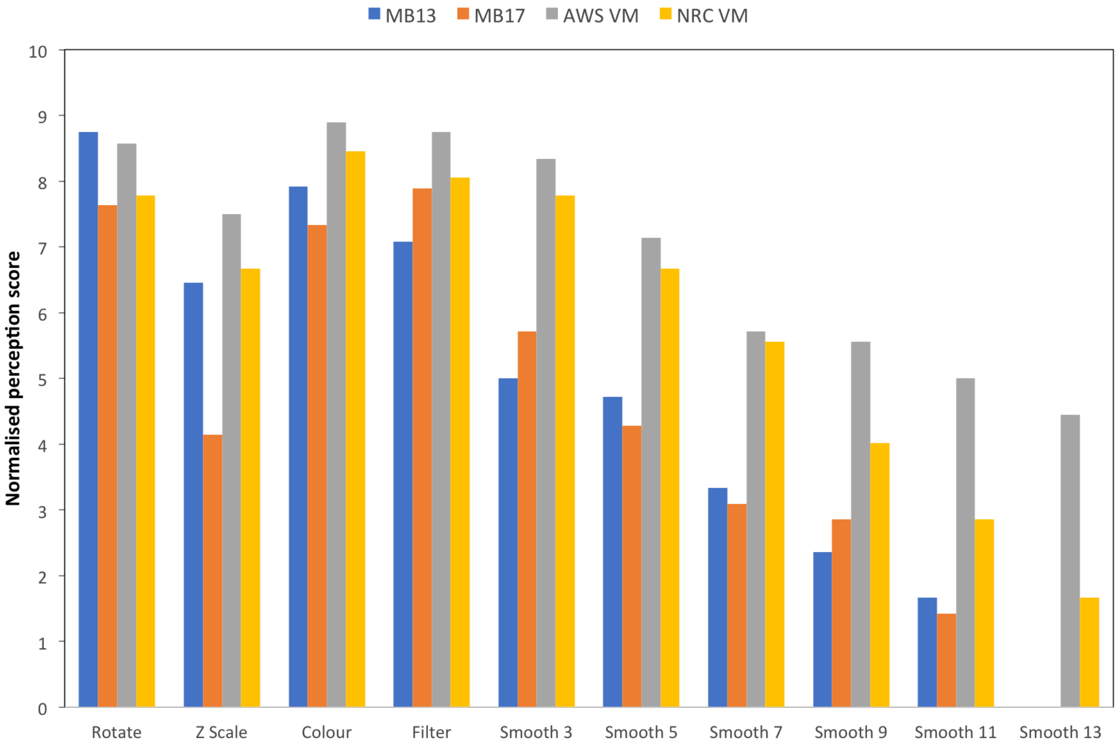


Figure 6: After each step of the 3D tasks, the participant scored the environment between 0 and 10. The normalised median perception scores are shown for each environment and for each stage of the 3D component of the investigation. The overall trend follows the one seen in 5. Each of the participants had access to both the cloud environments and only one local laptop, so the above scores are each drawn from 10 × MB13, 10 × MB17, 18 × AWS virtual machine, and 18 × NRC virtual machine results.

an unaffected interaction with the loaded volume, while the second task placed considerable load on the GPU. Despite this calibration step, it was still necessary to normalize the responses.

Figure 6 shows the corresponding stages of the 3D spectral cube rotation as median perception scores for each environment. The overall trend follows the one seen in Figure 5:

1. a relatively high initial score;
2. a drop for the Z-Scale step;
3. a return to the previous levels for the Colour and Filter steps;
4. a steady decline for the smoothing steps.

It is interesting that despite the high frame rates for the local laptops in the Rotate, Colour and Filter steps, the perception scores for all four environments are very close. This suggests that the usefulness of an environment is acceptable above a certain frame rate, and that additional performance does not necessarily correspond to a better experience.

### 3.7.2. *Performance*

Manipulation of a large spectral cube places a significant load on a GPU. Most laptop GPUs are designed to provide an optimal GUI experience, and are not designed for these types of workloads.

Increasing the Z-Scale factor to 10 had a huge impact on the GPU performance, with all environments dropping significantly. This resulted in an expected drop in perception score, though the MB17 showed a much bigger drop than the others. When the Z-Scale was returned to 1, and a colour filter was applied, the frame rates returned to the previous levels, as did the perception scores, with a slight improvement for the cloud environments over the local environments.

Applying the Colour or Filter option does little to affect the GPU performance and this is reflected in frame rates as seen in Figure 5. As expected, the perception scores also remained relatively unchanged, still with a slight favoring of the cloud environments.

As before, when a greater load is placed on the GPU, in this case applying a Smoothing algorithm, the local computer GPU frame rate drops markedly. Oddly, MB17 dropped more than the MB13, which suggests that the GPU in the MB13 is better suited to this sort of workload than the one in the MB17 (see Section 3.1 for more details). Yet at smoothing value of 3, the cloud virtual machines' perception scores are now markedly higher than the local computers.

As the smoothing value is steadily increased to 13, we see a steady drop in all perception scores, though a clear separation between the local laptops and the cloud environments is apparent. By smoothing value of 5, the AWS virtual machine outperforms the other environments and continues to do so for the rest of the task. For the NRC virtual machine, even though its frame rate for the higher smoothing values are almost the same as the local computers, it maintains a higher perception score until the end of the task.

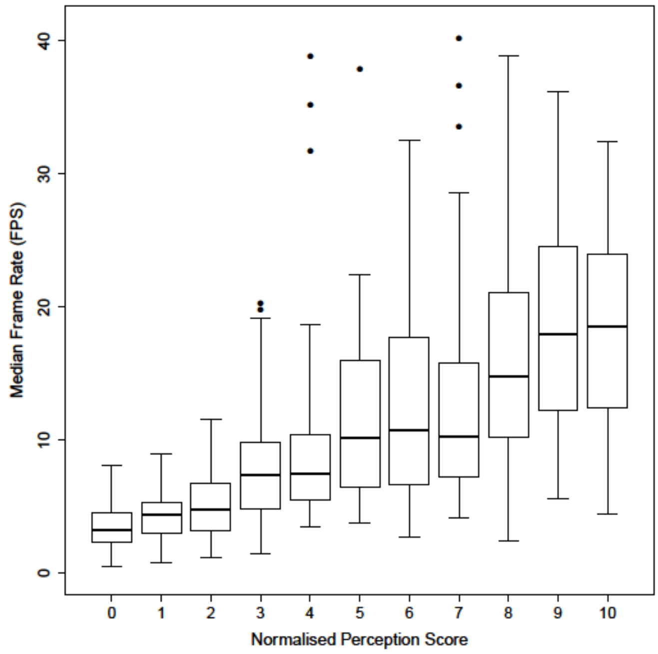


Figure 7: The normalised median perception scores versus the median frame rates shows that in general, participants' increasing perception scores corresponded with the increasing frame rates. However, there is also considerable variation, and sometimes high frame rates correspond to low scores, and low frame rates can correspond to high scores. This reaffirms the idea that frame rates are not the only factor participants consider when evaluating an environment. For each Score, the whiskers extend to the data value that is no more than twice the interquartile range from the box.

### 3.7.3. *Perception versus frame rate*

Figure 7 shows that while there is a correlation between the normalised perception score and the median frame rate, a great deal of variation is still present. Figures 8 A, B, C and D show that the variation of frame rate on the local environments is considerably greater than the cloud environments, and likely is the most significant factor for determining the perception score. The bigger the drop in the frame rate experienced, the lower the perception score, regardless of the highest frame rate experience. This is why low perception scores are observed even for the highest of frame rates. We also see some surprisingly high perception scores for low frame rates, suggesting that for some participants, the frame rate itself did not determine their experience of the environment for a particular task. That is, a low frame rate can still provide a satisfactory experience for some researchers, depending on their expectations and needs.

## 3.8. *Post-task user reflections*

After completing the hands-on component of the investigation, the participants were asked to reflect on their experience with either the 2D or the 3D task. They were asked again to rate the environments they had used from 0 to 10, but this time they were focusing on performance, ease of use, and suitability for the tasks conducted. Four of the participants chose to focus on the 2D experience, while 16 chose the 3D experience. This split might have occurred as the 3D experience was more readily recalled because it more recent, or it may have been chosen because it was perceived as more enjoyable. Since the purpose of the post-investigation survey was to focus the participant on the functional purpose of the environments, and to be able to meaningful comment on the viability of the VHDs as a replacement for a local desktop computer, this split does not affect the overall outcome. The survey results are shown in Figure 9.

- **Performance:** 17 participants had previous experience with X11 Forwarding or VNC, and they generally found its performance adequate for the tasks they required it for. However, after having completed the user tests, these participants found the cloud environments matched or outperformed both the local environments and their previous experiences. While the best frame rate was recorded for the MB17, it also received the lowest score for performance. Figures 8A and B, reveal greater variation in the median frame rates than the cloud environments shown in C and D, which influences participants' perceptions and hence their overall impression of the environments.
- **Ease of use:** While MB17 received the lowest Performance score in the post-investigation survey, it received the highest score for Ease of use. While not specifically recorded, most participants were quite comfortable with the MacBook Pro laptops and scored them positively. Overall, the four environments evaluated in the investigation were considered easy to use, which suggests that once familiarity is gained in a cloud environment, there is little difference between operating a virtual machine and a physical computer.
- **Suitability:** This score was intended to give the participants the opportunity to summarize their experience. As Figure 9 shows, the participants generally found that while the local environments were quite acceptable, they were exceeded by the cloud environments. Interestingly, despite the AWS virtual machine showing a clear lead in the perception scores for the 3D tasks, it was seen as being just as suitable as the NRC virtual machine.

During the post-investigation interview, the participants were asked to reflect on the problems they experienced with the different environments they used. Figure 10 shows the frequency of these issues, which are categorized into four themes:

- **GPU:** The GPU performance was considered the main problem for the completion of the task
- **System:** The system experienced a momentary freeze, crashed, was unfamiliar to the participant, slow to load the application, or slow to load data
- **Video:** Most commonly experienced as video tearing, where a mismatch in the GPU rendering frequency and the screen display refresh caused a momentary splitting of the image. Also where the colours displayed were not as expected

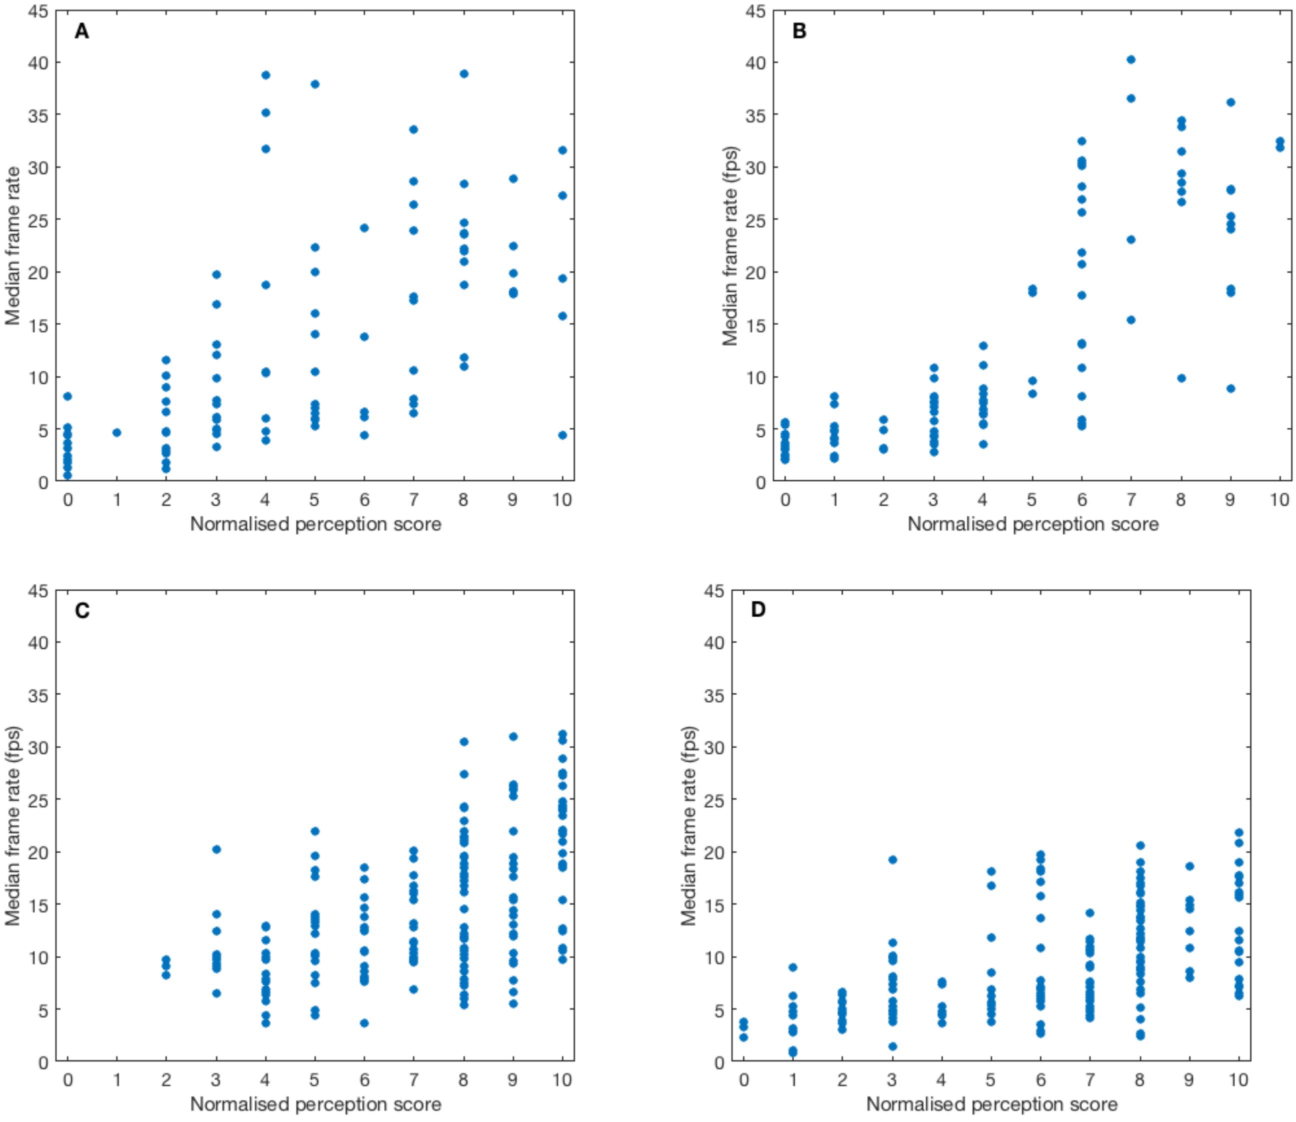


Figure 8: The normalised median perception scores versus the median frame rates for the individual environments shows greater variation in the local laptops (A) MB13 and (B) MB17 than the cloud environments (C) AWS virtual machine, and (D) NRC virtual machine. This suggests that the variability in the frame rate is linked to the user experience.

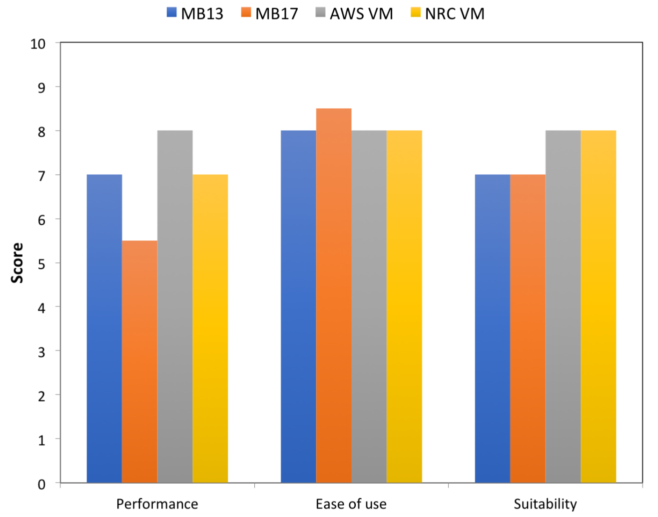


Figure 9: Having completed the tasks, the participants were asked to rate (out of 10) the NRC and AWS virtual machines and the laptop they used for Performance, Ease of use, and Suitability. The values shown are the median response scores for the whole cohort, with 10 responses for MB13 and MB17, and 18 responses for AWS virtual machine and NRC virtual machine. The performance of the MB17 was rated lowest, despite showing the best performance using a standard benchmark (see Figure 1). Overall, the cloud environments were seen as being equally performant, easy to use, and suitable for the tasks undertaken.

- **Latency:** the only noticeable demonstration of this issue was when the DS9 application window was moved, but tracked slightly slower across the screen than the mouse

Notably, the MB13 received the most complaints about the GPU performance, however it was the only one not to receive comments about video tearing. Video tearing was most apparent on the MB17, while the AWS virtual machine was the only environment where the latency was an issue.

## 4. Discussion

To be a viable replacement or adjunct to a desktop machine, a VHD needs to support research workflows in a seamless way. The user experience – compared to a physical desktop – must be maintained or even exceeded. Astronomers are more likely to accept and adopt the use of a fully remote desktop if the interaction speeds (e.g. from movement of mouse on the desktop to movement of the mouse cursor on screen) are no worse than those currently achieved on a local desktop.

The key outcomes of our user experience testing demonstrate the following:

- VHDs are as easy to use as a standard desktop;
- A correctly resourced and configured VHD provides a suitable environment to run typical astronomy software;
- A correctly resourced and configured VHD can provide a better user experience than a local laptop computer; and
- VHDs can provide a viable desktop alternative for astronomers.

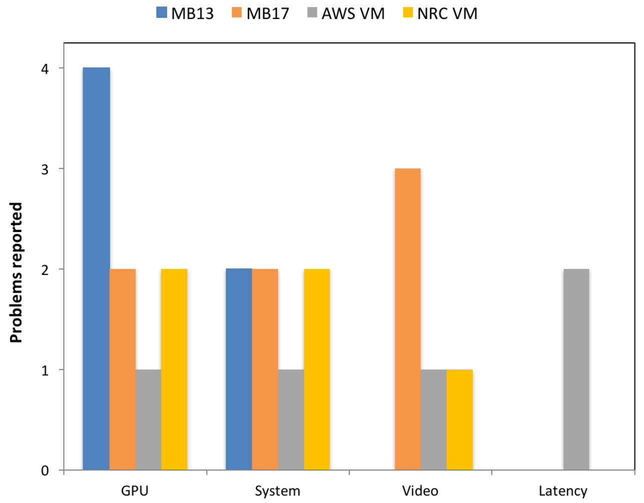


Figure 10: During the post investigation survey, participants identified issues with the environments. These can be categorized into four themes, relating to the GPU, the System, the Video display and Latency. The MB13 GPU was the most complained about, followed by the video tearing on the MB17. Latency was only observed for the AWS VM. In general, participants experienced fewer problems with the cloud environments.

However, the tasks did not touch upon some other elements of VHDs that are part of the cloud experience. For example, creating, configuring and maintaining a virtual machine is a non-trivial task and requires a reasonable level of technical skill. Factoring this aspect into the experience might have had a negative impact on the impressions of the participants, and therefore it might be argued that this should have been included. Yet cloud specific impacts need to be measured against local computer impacts.

It is also important to consider whether pre-built cloud images can mitigate much of the challenge associated with the creating, configuring and maintaining a virtual machine. Further, a managed service can eliminate the need for an astronomer to manage the VHD entirely. Because of the complexity surrounding the establishment of a VHD in comparison to a local computer, the investigation was deliberately limited to focus purely on the operational use of the environments.

While all participants agreed that there was considerable potential in the use of VHDs in their workplace, some skepticism remained. The idea of committing to a cloud-based service continues to be a source of concern for many participants, and reflects the wider community attitudes. However, there are many elements to consider when choosing a suitable environment.

### *4.1. Pricing comparison*

The initial outlay for a computer of reasonable power is something that many researchers and research departments take for granted as they prepare their annual expenditure projections. For a typical PhD research student, an estimate of the likely computation power they will require is made at the start of their research journey. However, it may take several months

Table 3: This table shows the basic pricing for each environment, and the life expectancy of the hardware purchased. The AWS On-demand and Spot pricing are included separately. An ASUS i7 laptop is included for comparison, though one was not used during the study. VM = Virtual Machine. UoM = University of Melbourne. Prices are shown in Australian dollars, noting that there are also country-specific variations for hardware purchases.

| Environment | Approximate Price | Life expectancy | Comment |
|---|---|---|---|
| MB13 | \$3,000 (2013) | 4 years | Best option available in UoM standard catalogue |
| MB17 | \$4,500 (2017) | 4 years | Best option available in UoM standard catalogue |
| NRC VM | \$200 per core per year (2017) | New hardware added each year | Limited GPU options |
| AWS VM (On-demand) | \$0.90 per hour (2017) | New hardware added each year | Additional costs for storage and data ingress/egress |
| AWS VM (Spot) | \$0.34 per hour (as at 2017-11-15) | New hardware added each year | VMs are stopped at bid threshold |
| ASUS i7 | \$2,100 (2018) | 4 years | Comparable system to the MB17, from the Amazon online catalogue (Australia) |

before they reach a point of actually needing that power. Until that point, they are financially over-committed on the purchase. Having started to make reasonable use of the computer, they may well find that soon the demands of their research outgrows their local machine, and they need to move to remote high-performance computing facilities. Once again, they are over-committed on the purchase of the computer.

To determine how powerful a computer should be to meet a researcher's need, the most common approach is to estimate the peak workload expected, and then buy a machine that most closely achieves that requirement within the available budget envelope. However, unless tested directly with the workload intended, simple benchmarks and specification comparisons may not provide an adequate indication of the machine's suitability. Combined with the fact that the machine may only be used at that peak for a fraction of its lifetime, a financial over-commitment is likely.

Alternatively, a more modest purchase paired with a suitable virtual machine from a cloud provider could provide a far more financially responsible option. Furthermore, an out of warranty machine may still be able to perform sufficiently well as to provide connectivity to a virtual hosted desktop for several additional years. A local computing device and a suitable network to access the virtual machine is still required, but the local device does not need more power than is required for that function.

A service like the Nectar Research Cloud is not directly charged to the researchers for research supported by national funding, and several participating institutions also provide additional resources internally to their own researchers. Commercial service providers like AWS are fee-for-service and are an excellent way to explore options without committing significant financial resources. See Table 3 for a pricing table.

More extensive testing with a larger range of laptop and desktop computers, as well as other virtual machine variants and locations, would have been ideal. However limitations on participant time and numbers, as well as the lack of financial resources to purchase additional equipment beyond the laptops included in the study, made this impossible. Furthermore, as Lam et al. (2012) indicates, the immediate subjective feedback from user experience is more important than the range of technical specifications in this case. While a Linux laptop was not able to be purchased for the study, a model with similar specifications to the MB17 (an ASUS i7-7700HQ[22] with a more performant graphics card than the MB17) is included in Table 3 for the purpose of comparison. The model ships with Microsoft Windows 10 but could be reinstalled with the same Linux Ubuntu operating system and applications used in this study.

The pricing comparisons in Table 4 show the costs of laptops and the cloud environments under different usage models. Hardware purchased, either as laptops MB13 or MB17, or as part of the NRC provides, provide the best value for money for the maximum use scenario, where the computer is used at capacity for its useful life. However, other than high performance compute clusters, this is an unlikely usage pattern. The high use scenario, where the device is used 6 hours per day, 5 days per week, 42 weeks per year for four years, shows a close alignment between the MB13 and the NRC VM, and again with the MB17 and the AWS VM. However, based on the experiences reported by the participants, the cloud environments would still provide the better value for money. Finally, the low use scenario consisting of 2 hours per day, 3 days per week, 32 weeks per year for four years, shows a clear price advantage for the AWS VM, costing less than \$700. As above, the ASUS i7 is included in Table 4 for reference.

The above discussion focuses on the On-demand pricing for the AWS VM, and further cost reduction can be made by using Spot pricing, though additional care needs to be taken to monitor the resource availability. It is also important to note that the cloud pricing for both NRC VM and AWS VM include all costs, while the laptop costs include only the initial purchase price.

### 4.2. Risks

When considering the viability of a VHD as a replacement or augmentation of an existing desktop computer, it is important

[22] Specification: 16GB DDR4 RAM, 256GB SSD drive, 15.6inch Ultra HD display, NVIDIA GTX1050 (4G RAM), Microsoft Windows 10. Purchase price of ~AUD\$2,100.00 was correct at 18/03/2018

Table 4: Direct pricing comparisons can be difficult due to the difference in usage and service models. This table shows the price comparison for a maximum use model (i.e. 24 hours per day, 7 days per week, 52 weeks per year, 4 years), high use (6 hours per day, 5 days per week, 42 weeks per year, 4 years), and low use (2 hours per day, 3 days per week, 32 weeks per year, 4 years). * indicates the fixed value. Non-starred values are calculated based on the fixed value and the maximum, high and low usage configurations. NRC virtual machine priced based on estimated $200 per core per year. An ASUS i7 laptop is included for comparison, though one was not used during the study. Prices are shown in Australia dollars.

| Environment | Approximate Price | $ per hour (maximum) | $ per hour (high) | $ per hour (low) |
|---|---|---|---|---|
| MB13 | $3,000* | $0.09 | $0.60 | $3.91 |
| MB17 | $4,500* | $0.13 | $0.90 | $5.86 |
| NRC VM | $3,200* | $0.09 | $0.64 | $4.17 |
| AWS VM (maximum, On-demand) | $31,379.71 | $0.90* | $0.90* | $0.90* |
| AWS VM (high, On-demand) | $4,525.92 | $0.90* | $0.90* | $0.90* |
| AWS VM (low, On-demand) | $689.66 | $0.90* | $0.90* | $0.90* |
| AWS VM (maximum, Spot) | $11723.71 | $0.34* | $0.34* | $0.34* |
| AWS VM (high, Spot) | $1,690.92 | $0.34* | $0.34* | $0.34* |
| AWS VM (low, Spot) | $257.66 | $0.34* | $0.34* | $0.34* |
| ASUS i7 | $2,100.00* | $0.06 | $0.42 | $2.73 |

to consider the risks. As part of the post-experience evaluation, participants reported on their concerns. Possible solutions are available to each of the risks using either direct mitigation or user education.

#### *4.2.1. Network availability*

The concern raised most often by the participants when considering the viability of a VHD was the availability of a suitable network. The network needs to provide sufficient bandwidth and low latency, but must also remain consistent. This concern reflects the changing way that researchers work, with portable devices and working in a variety of conditions with a variety of networks now common. Most research institutions now recognize the importance of providing a highly stable network.

However, while campus networks might be stable, the local café or airport wifi may not be. Other networks such as ADSL or cable connections in the home, or 4G mobile networks, are more subject to congestion than campus networks, and therefore may provide a suitable network at some times, and not others.

While network connectivity is important when using VHDs, an unexpected disconnection is disruptive to workflow, VHDs are quite tolerant of this sort of event. When a client suddenly loses connectivity, the VNC connection is dropped, but the running VHD session does not terminate. Instead, it remains active and available for when the connection to the client can be re-established.

#### *4.2.2. Hidden costs*

Using a cloud service like AWS requires careful planning to ensure costs do not blow out. While the On-demand pricing captures the essential rate for the virtual machine, it does not include things like data ingress and egress, or additional storage capacity. While there are several tools for monitoring active resources, it is quite easy to overlook certain components, which can result in a significant bill.

For a service like the Nectar Research Cloud, the hidden costs are more significant. For the individual researcher, the service can appear to be completely free, as they never see a bill, nor are they necessarily required to answer to their department for their usage costs. However, institutions like Nectar, the University of Melbourne and Swinburne University of Technology invest capital and operational funds to provide the service, and are answerable for the efficacy of that expenditure. Without a mechanism to remind researchers of the value of the service they consume, the claimed resources are often left idle. For many such cloud services, the resource utilization can be as low as 10%, even as the resource allocation approaches 100%.

#### *4.2.3. Security*

As more research is conducted online, the security of online resources is of ever increasing importance to researchers and their institutions. Most researchers are familiar with online storage, but many are not aware of the potential security implications. This includes the possibility of hacking, ethical considerations, and the unwitting loss of intellectual property rights.

Hacking is one of the highest concerns, where unauthorized access to research data can occur. In the case of a VHD, there is an additional risk that the VHD itself may be managed by a non-expert, who may not be as proactive in securing a system as they need to be. A breach might also occur with a properly secured system, simply because of a bug in the connection service, or one of the many other system services, that can be exploited by a person with nefarious intent. However, training and diligence are very effective tools for minimizing the possibility of this happening.

In some cases, though unlikely in the field of astronomy, the data being processed might have ethical requirements that prevent it being stored in a public cloud, or prevent it being moved out of a geographical location. This is a common issue for medical and biological science data. In this case, the use of private research clouds is usually an acceptable approach.

Finally, the storing of data on a public cloud may subject the data to intellectual property laws that relate to the location of the data centre rather than the origin of the data. Protections applicable in one country are not necessarily available in another country. It is important that researchers are aware of where their data will be stored, and what laws have jurisdiction over their data.

### 4.3. Convenience

The convenience of using a laptop or desktop computer, without needing to connect to a VHD, was cited by one participant as being their main reason for not accepting a VHD as a viable alternative. However, as common as this view might be, it does not take account of the risks associated with the local computer itself. For example, a local computer can fail for many reasons, such as loss of power, breakage, or theft. While these issues also prevent access to a VHD from that device, they do not prevent access from another available device. In fact, reducing the dependence on a specific local device reduces the risk of mid-to-long term access to resources. Theft from a properly secured data centre is rare, and while hardware failure is quite common, cluster-based computing, and the services which operate on top, are typically very tolerant of such failures due to their frequency.

Interestingly, the criteria defined in Section 2.3 are measured in comparison with the operation of a local computer, but such a local computer need not meet these criteria to be considered acceptable. One participant stated that even though they recognized that the cloud environments performed better, they would still choose a physical desktop computer over a cloud-based desktop solution, even if it didn't perform as well, because it could be used in places where the network might not be adequate.

## 5. Conclusions

As we move closer to the Square Kilometre Array's (SKA) exabyte-data era, it will be increasingly impractical to visualise data products on a local desktop with a standard display device (Hassan & Fluke, 2011). Not only will the data vastly exceed the local storage, memory and computation capabilities of desktop computers, but the image resolution and graphics capabilities required will force astronomers to consider large-format displays and specialised graphics processing hardware.

Today, many instruments already produce datasets that are too big to be downloaded to a local computer, so it stays in the data centre. The compute power also sits in the data centre, close to the data. Astronomers are happy to use remote computing as part of their workflow, but prefer to continue use desktop applications for some tasks, such as visualisation. However, this typically means transferring data to the local computer. As this becomes harder to do, having windowed applications forwarded to the local desktop is one solution, but the performance is limited.

Cloud services with GPU hardware allow GPU-enabled applications to be used in a VHD environment on a remote computation platform, with just the screen updates transmitted to the local computer.

By combining benchmarks with user experience testing, we found that VHDs provide a viable, cost-effective desktop alternative for typical astronomy applications, particularly for graphics-intensive tasks.

While benchmarking may approximate the intended workload, only direct testing with astronomy workflows operating under load will really determine the adequacy of a compute resource. Standard benchmarking applications, Unigine Heaven and Valley, indicated that a new laptop computer was more performant than a cloud-based VHD, but user experience testing revealed that for some tasks, a VHD can provide a better solution.

Through our combined use of benchmarking and user experience testing we compared two laptops, one four years old and one new, with two cloud-based desktop environments (AWS and the Nectar Research Cloud). We have shown that:

1. For the 2D and 3D tasks, the environments were equally simple to use. During the investigation, it was apparent that the cloud environments operated as hoped for the tasks presented, easily matching the local computers. As many astronomers use standard applications, and these applications function the same way across the available environments, the only differences are the environments themselves.
2. The cloud environments provided an equally smooth experience for the required tasks as the local computers. Figure 10 shows that other than minor latency caused by the distance to a data centre (in the case of AWS virtual machine located in Sydney, Australia), the cloud environments matched or bettered the local computers. In fact, one of the likely reasons for the positive responses to the cloud environments was due to the more consistent performance shown by the cloud environments when compared to the more erratic performance (in terms of frame rate) shown by the local computers as the workload increased. High frame rates do not necessarily correspond to a better user experience, and for many participants, the more consistent the frame rate, the better the experience, depending on their expectations and needs.
3. Table 4 shows that other than the maximum use case, where an environment is used continuously for four years, the cloud environments provide competitive alternatives when compared to purchasing mid-to-high end laptops and workstations. They also offer a degree of flexibility that increases their financial suitability.
4. Being powerful enough for the task required does not only mean achieving the maximum, but also achieving a suitable minimum. A local computer chosen for its suitability in a high compute demand scenario, will be greatly overpowered for standard office operations, such as checking email or editing LaTeX documents. This investigation has shown that cloud environments may be better suited to the tasks, as presented, than the local computers.
5. High availability is critical to the business practices of cloud providers, and research network infrastructure is just as critical to research institutions. Networking performance during the tasks showed that a reasonably stable network with moderate bandwidth is sufficient to use a VHD. This type of networking is available at most re-

search institutions in Australia, where our user experience testing was conducted, and much of the world.

## Acknowledgements

We thank the participants for contributing their time to this project. All experimental work was approved and conducted in accordance with the requirements of Swinburne Universitys Human Research Ethics Committee (SUHREC). We also thank Anita, Amy and Jacob Meade for their assistance in preparation and data entry of the participant results, Dr Dany Vohl (Swinburne University of Technology) for his support with Shwirl, and Dr Glenn Kacprzak (Swinburne University of Technology) for sample image data. We also thank Dr Stephen Giugni and Dr Steven Manos (University of Melbourne) for supporting this research, and Terry Brennan (University of Melbourne) for assisting with the costing of the Nectar virtual machines, Dr David Perry and Dylan McCullough (University of Melbourne) for providing the build script for the cloud virtual machines. This research was supported by use of the Nectar Research Cloud and by the Melbourne Node at the University of Melbourne. The Nectar Research Cloud is a collaborative Australian research platform supported by the National Collaborative Research Infrastructure Strategy. This research was also supported by Adrian While and Craig Lawton through the "AWS Cloud Credits for Research" program from Amazon Web Services.

Funding for the Sloan Digital Sky Survey IV has been provided by the Alfred P. Sloan Foundation, the U.S. Department of Energy Office of Science, and the Participating Institutions. SDSS-IV acknowledges support and resources from the Center for High-Performance Computing at the University of Utah. The SDSS web site is www.sdss.org.

SDSS-IV is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS Collaboration including the Brazilian Participation Group, the Carnegie Institution for Science, Carnegie Mellon University, the Chilean Participation Group, the French Participation Group, Harvard-Smithsonian Center for Astrophysics, Instituto de Astrofísica de Canarias, The Johns Hopkins University, Kavli Institute for the Physics and Mathematics of the Universe (IPMU) / University of Tokyo, Lawrence Berkeley National Laboratory, Leibniz Institut für Astrophysik Potsdam (AIP), Max-Planck-Institut für Astronomie (MPIA Heidelberg), Max-Planck-Institut für Astrophysik (MPA Garching), Max-Planck-Institut für Extraterrestrische Physik (MPE), National Astronomical Observatories of China, New Mexico State University, New York University, University of Notre Dame, Observatário Nacional / MCTI, The Ohio State University, Pennsylvania State University, Shanghai Astronomical Observatory, United Kingdom Participation Group, Universidad Nacional Autónoma de México, University of Arizona, University of Colorado Boulder, University of Oxford, University of Portsmouth, University of Utah, University of Virginia, University of Washington, University of Wisconsin, Vanderbilt University, and Yale University.

## References

Abolfathi, B., Aguado, D.S., Aguilar, G., Prieto, C.A., Almeida, A., et al., 2017. The Fourteenth Data Release of the Sloan Digital Sky Survey: First Spectroscopic Data from the extended Baryon Oscillation Spectroscopic Survey and from the second phase of the Apache Point Observatory Galactic Evolution Experiment. arXiv:1707.09322 [astro-ph] ArXiv: 1707.09322.

Astrocompute, 2015. Seeing stars through the Cloud - SKA Telescope. AstrocomputeWeb.

Ball, N.M., 2013. CANFAR+ Skytree: A Cloud Computing and Data Mining System for Astronomy. arXiv preprint arXiv:1312.3996 .

Berriman, B., Groom, S.L., 2011. How will astronomy archives survive the data tsunami? Communications of the ACM 54, 52.

Berriman, G.B., Brinkworth, C., Gelino, D., Wittman, D.K., Deelman, E., et al., 2012. A Tale Of 160 Scientists, Three Applications, A Workshop and A Cloud. arXiv preprint arXiv:1211.4055 .

Berriman, G.B., Deelman, E., Juve, G., Rynge, M., Vöckler, J.S., 2013. The application of cloud computing to scientific workflows: a study of cost and performance. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 371.

Berriman, G.B., Good, J.C., 2017. The Application of the Montage Image Mosaic Engine to the Visualization of Astronomical Images. Publications of the Astronomical Society of the Pacific 129, 058006.

Berriman, G.B., Good, J.C., Rusholme, B., Robitaille, T., 2016. The Next Generation of the Montage Image Mosaic Toolkit. arXiv preprint arXiv:1608.02649 .

Berriman, G.B., Juve, G., Deelman, E., Regelson, M., Plavchan, P., 2010. The application of cloud computing to astronomy: A study of cost and performance, in: e-Science Workshops, 2010 Sixth IEEE International Conference on, pp. 1–7.

Bevan, N., 2009. Usability, in: LIU, L., ÖZSU, M.T. (Eds.), Encyclopedia of Database Systems. Springer US, Boston, MA, pp. 3247–3251. DOI: 10.1007/978-0-387-39940-9_441.

Bipinchandra, G.K., Aluvalu, P., Singh, D., Shanker, A., 2014. Intelligent Resource Allocation Technique For Desktop-as-a-Service in Cloud Environment. arXiv preprint arXiv:1404.7494 .

Borne, K., 2009. Scientific data mining in astronomy. arXiv preprint arXiv:0911.0505 .

Brunner, R., Djorgovski, S., Prince, T., Szalay, A., 2001. Massive datasets in astronomy. Arxiv preprint astro-ph/0106481 .

Caton, D.B., Hawkins, L., 2009. Remote Observing: Equipment, Methods and Experiences at the Dark Sky Observatory, in: American Astronomical Society Meeting Abstracts #213, p. 427.

Deboosere, L., Vankeirsbilck, B., Simoens, P., De Turck, F., Dhoedt, B., et al., 2012. Cloud-Based Desktop Services for Thin Clients. IEEE Internet Computing 16, 60–67.

Deelman, E., Singh, G., Livny, M., Berriman, B., Good, J., 2008. The cost of doing science on the cloud: the montage example, in: Proceedings of the 2008 ACM/IEEE conference on Supercomputing, IEEE Press. p. 50.

Dodson, R., Vinsen, K., Wu, C., Popping, A., Meyer, M., et al., 2016. Imaging SKA-scale data in three different computing environments. Astronomy and Computing 14, 8–22.

Duato, J., Yalamanchili, S., Ni, L.M., 1997. Interconnection networks: an engineering approach. IEEE Computer Society Press, Los Alamitos, Calif.

Goscinski, W.J., Paterson, D., Hines, C., McIntosh, P., Thompson, D., et al., 2015. MASSIVE: an HPC Collaboration to Underpin Synchrotron Science .

Hassan, A., Fluke, C., 2011. Scientific Visualization in Astronomy: Towards the Petascale Astronomy Era. Publications of the Astronomical Society of Australia 28, 150–170.

Hoffa, C., Mehta, G., Freeman, T., Deelman, E., Keahey, K., et al., 2008. On the use of cloud computing for scientific workflows, in: eScience, 2008. eScience'08. IEEE Fourth International Conference on, IEEE. pp. 640–645.

Hong, C.H., Spence, I., Nikolopoulos, D.S., 2017. GPU Virtualization and Scheduling Methods: A Comprehensive Survey. ACM Computing Surveys 50, 1–37.

Iserte, S., Clemente-Castello, F., Castello, A., Mayo, R., Quintana-Orti, E., 2016. CLOSER 2016: proceedings of the 6th International Conference on Cloud Computing and Services Science: Rome, Italy, April 23-25, 2016. SCITEPRESS - Science and Technology Publications, Lda, Setbal, Portugal.

Juric, M., Tyson, T., 2012. LSST Data Management: Entering the Era of Petascale Optical Astronomy. Proceedings of the International Astronomical Union 10, 675–676.
Juve, G., Deelman, E., Vahi, K., Mehta, G., Berriman, B., et al., 2009. Scientific workflow applications on Amazon EC2, in: E-Science Workshops, 2009 5th IEEE International Conference on, IEEE. pp. 59–66.
Khalid, F., Shoaib, U., Sarfraz, M.S., Shabbir, A., Shaheed, S.M., et al., 2016. Desktop Virtualization: An Art to Manage and Maintain affordable PC infrastructure. International Journal of Computer Science and Information Security 14, 187.
Lam, H., Bertini, E., Isenberg, P., Plaisant, C., Carpendale, S., 2012. Empirical studies in information visualization: Seven scenarios. IEEE Transactions on Visualization and Computer Graphics 18, 1520–1536.
Massimino, P., Costa, A., Becciani, U., Vuerli, C., Bandieramonte, M., et al., 2014. ACID Astronomical and Physics Cloud Interactive Desktop: A Prototype of VUI for CTA Science Gateway, in: Manset, N., Forshay, P. (Eds.), Astronomical Data Analysis Software and Systems XXIII, p. 293.
Meade, B., Manos, S., Sinnott, R., Fluke, C., Knijff, D.v.d., et al., 2013. Research Cloud Data Communities, in: THETA 2013, Hobart, Tasmania. Copyright 2013 THETA: The Higher Education Technology Agenda.
Miller, K., Pegah, M., 2007. Virtualization: virtually at the desktop, ACM Press. pp. 255–260.
Nieh, J., Yang, S.J., Novik, N., 2000. A comparison of thin-client computing architectures. Technical Report. Technical Report CUCS-022-00, Department of Computer Science, Columbia University.
Rampersad, L., Blyth, S., Elson, E., Kuttel, M.M., 2017. Improving the usability of scientific software with participatory design: a new interface design for radio astronomy visualisation software, ACM Press. pp. 1–9.
Ravi, V.T., Becchi, M., Agrawal, G., Chakradhar, S., 2011. Supporting GPU sharing in cloud environments with a transparent runtime consolidation framework, in: Proceedings of the 20th international symposium on High performance distributed computing, ACM. pp. 217–228.
Rimal, B.P., Choi, E., Lumb, I., 2010. A Taxonomy, Survey, and Issues of Cloud Computing Ecosystems, in: Antonopoulos, N., Gillam, L. (Eds.), Cloud Computing. Springer London, London, pp. 21–46. DOI: 10.1007/978-1-84996-241-4_2.
Sabater, J., Sánchez-Expósito, S., Best, P., Garrido, J., Verdes-Montenegro, L., et al., 2017. Calibration of LOFAR data on the cloud. Astronomy and computing 19, 75–89.
Vöckler, J.S., Juve, G., Deelman, E., Rynge, M., Berriman, B., 2011. Experiences using cloud computing for a scientific workflow application, in: Proceedings of the 2nd international workshop on Scientific cloud computing, ACM. pp. 15–24.
Vohl, D., 2017. Shwirl: Meaningful coloring of spectral cube data with volume rendering. Published: Astrophysics Source Code Library.
Walter, F., Brinks, E., de Blok, W.J.G., Bigiel, F., Kennicutt, Jr., R.C., et al., 2008. THINGS: The H I Nearby Galaxy Survey. 136, 2563–2647. 0810.2125.
Younge, A.J., 2016. Architectural principles and experimentation of distributed high performance virtual clusters. Ph.D. thesis. Indiana University.

## Appendix

### A. Network

The network performance for connections to the VHDs, while varied, delivers sufficient bandwidth and is robust enough to provide a user experience that matches the local desktop.

Network Speedtests shown in Figures A1 and A2 were conducted between July and September, 2017.

Despite the longer ping times for participants 7 and 10, and the larger variation in times for participants 17 and 19, the network did not affect the performance of the VHDs during those tasks.

Figure A3 shows the network latency from the study locations to the virtual machines at the University of Melbourne data centre and the AWS Sydney data centre. Included are the

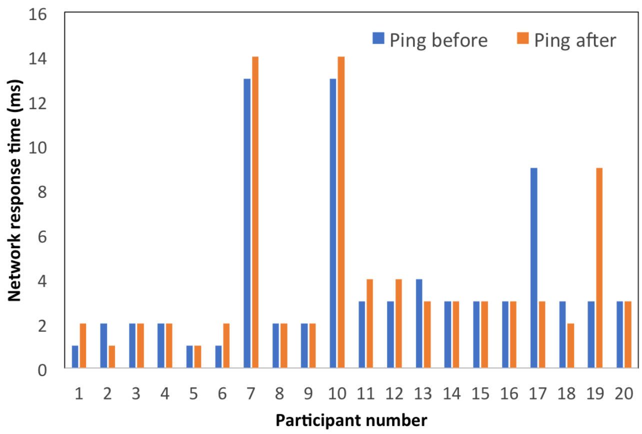


Figure A1: The network response time was measured before and after the tasks. The figure shows that the start and end state of the network environment did not vary greatly. Despite the variation of ping response times, participants did not report any latency impact during their use of the VHDs.

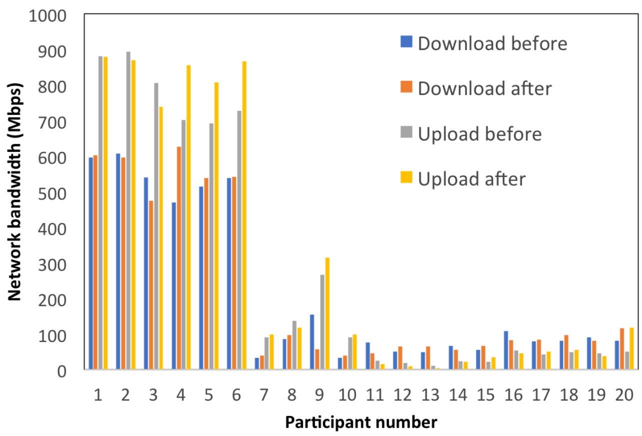


Figure A2: The network bandwidth for download and upload was measured before and after the tasks. Although they varied significantly due to the use of different connection types, no apparent impact on the performance of the cloud environments was identified by the participants.

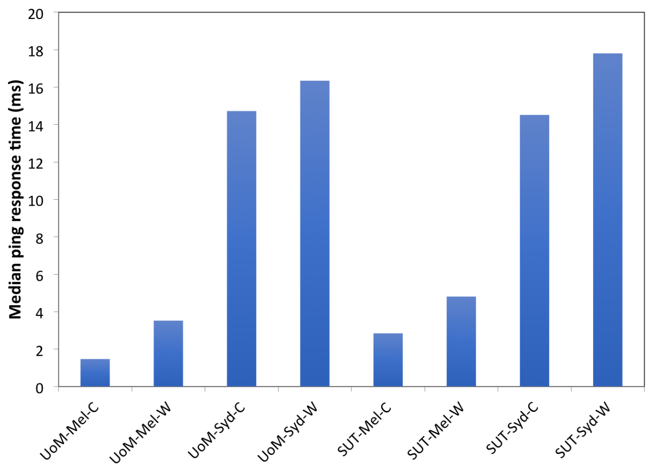


Figure A3: Network latency ping test results from study locations at Swinburne University of Technology (SUT) and the University of Melbourne (UoM) to the UoM data centre in Melbourne (Mel), Australia, and AWS data centre located in Sydney (Syd), Australia. C = laptop connected to network using a cable. W = laptop connected over wireless network. Ping tests were conducted in March, 2018.

ping response times over a cable connected network and the wireless networks available at each testing site.

### B. *User experience testing: procedure*

The tasks completed during the user experience tests were intended to engage the participants in a manner that closely resembled typical astronomy use. The most common experience is using windowed applications with a local mouse and keyboard, with information displayed on a local computer screen. The design of the 2D phase of the investigation met that aim, by challenging the participant to use common astronomy software to complete a common astronomy task. Using a clearly defined activity allowed the participants to focus their reflections on the usability of the environment in question. This was preferable to having the participants freely explore the interface, which might have yielded more 'operating system'-centric evaluations.

The purpose of the 3D phase of the investigation was to test the capacity of the environments under graphically intensive load. Many astronomy tasks involve the use of 3D models and volume, so it is imperative to understand how VHDs handle these workloads, and if they are comparable to local GPU computation.

Each user testing session took 35-45 minutes depending on the participant. This consisted of:

1. Introduction and pre-investigation interview (5 minutes)
2. 2D image alignment activity (10 minutes)
3. 3D spectral cube manipulation activity (15 minutes)
4. Post-investigation interview (5 minutes)

The 2D image alignment procedure was as follows:

1. Start DS9
2. Load first image
3. Load second image in new frame buffer
4. Observe the offset using the Blink mode
5. Start IRAF
6. Using the `imshift` command, create a shifted image
7. Load the shifted image into the second frame buffer in DS9
8. Observe the lack of offset using the Blink mode

The 3D spectral cube manipulation procedure was as follows:

1. Start Shwirl and load the spectral cube
2. Switch to the Camera and Adjustment tab and set the Z-Scale to '10'. Rotate the volume using the mouse, and then reset the Z-Scale to '1'.
3. Switch to the Colour tab and choose a new colour scheme for the volume. Rotate the volume using the mouse and then reset the colour scheme (Note: some participants did not reset the colour scheme, but no performance difference was observed in the results).
4. Switch to the Filter tab and adjust the High and Low filters. Rotate the volume and then reset the values.
5. Switch to the Smoothing tab, and set the Smoothing value to 3. Rotate the volume.
6. Repeat the previous step with smoothing values of 5, 7, 9, 11 and 13.

## 6.6 Lessons learned

Cloud computing provides the modern astronomer with more flexibility than is possible with monolithic computation infrastructure, or high-end localised workstations. Because data is overwhelming individual compute systems, a clustered compute approach is rapidly becoming the only option for data-intensive astronomy. However, while many of these visualisation applications are bound to a graphical interface, it is now possible to use these applications in a cloud environment, with GPU accelerators available to handle complex 3D graphics demands.

The flexibility of cloud services provides researchers with a way of "failing fast", where choosing the wrong computer is solved by simply deleting the instance and starting again. Configurations that are potentially useful to others can be cloned and distributed, and instances can be archived and made available with associated data to foster scientific reproducibility.

Recent advances in cloud-based graphical computing make cloud even more attractive. For example, NVIDIA announced a new product called GeForce Now[7], where a GPU-enabled cloud-based computer can be rented from NVIDIA, pre-configured to match a high-end gaming computer. While this service is aimed at the gaming community, it can also be an important resource for the astronomy researcher with significant graphics demands.

Understanding the true costs associated with commercial cloud services such as AWS is critical to successfully evaluate their viability for mid to long-term use. The AWS Elastic Block Storage (EBS) product is the standard volume storage option used in combination with EC2 instances, such as those used in the study presented in Meade & Fluke (2018, see Section 6.5). Pricing at different data centres varies for different products and use cases. For example, EBS costs AUD$0.16 per GB-month at both Northern California and Sydney data centres, while Elastic File Storage, a more flexible product, costs AUD$0.45 and AUD$0.49 per GB-month respectively. Figure 6.1 shows the monthly charge for a 100GB, 1TB and 10TB EBS volume, for each of the Northern California and Sydney data centres. Volumes of 100TB and greater are not considered here as they are typically in the purview of HPC facilities. Data ingress is free of charge.

AWS provide a Simple Monthly Calculator[8] to help users determine monthly costs. Table 6.1 was generated using this calculator and provides indicative pricing of operating a VM on AWS over a four year period. These costs include the g2.2xlarge VM that was

[7] `https://www.nvidia.com/en-us/geforce/products/geforce-now/mac-pc/`

[8] `https://calculator.s3.amazonaws.com/index.html`

used in Meade & Fluke (2018), an EBS volume and a data egress that matches the size of the volume each month. As Table 6.1 shows, costs can increase rapidly and quickly exceed the cost of an outright purchase of a computer. Herein lies the challenge of choosing between purchasing and renting a computing solution. On one hand, the value of an outright purchase can only be determined by reflecting on the usage of the device over its lifetime. On the other hand, while the value of the rented solution is known at the start, the anticipated usage requires some guesswork to estimate the load over the same lifecycle. Similarly, the purchase price of a computer is fixed at the time of purchase, while the rented solution provides the ability to adjust usage patterns to manage ongoing costs.

Commercial companies like AWS have very different drivers to most research institutions. Renting resources is well suited to business processes supported by cash flows, which is typical of many businesses. Funding for research often comes in the form of grants, which require very clear upfront costing, and often don't allow for significant adjustments as a research project progresses. Therefore, the application of a commercial model to research can be problematic. AWS often provide incentives to research institutions, typically in the form of credits for compute and storage, and data egress charge waivers, providing valuable opportunities for researchers to explore cloud computing solutions. However, without long-term commitment to these cost-reduction measures to support ongoing research, it will be difficult for researchers to commit to a commercial solution for any but the shortest of research projects.

As more research institutions move to cloud-based generic IT services like Microsoft's Office 365[9], and Google's G-Suite[10], research workloads will likely follow suit. AWS has a wide range of in-house products, but a number have companies have emerged with value-add services like Teradici's PC-over-IP[11]. Sharing powerful computation resources, rather than deploying discrete computing capability for each researcher, can improve value for money for research institutions and ultimately allow more funding to be invested in the research itself.

[9] https://www.office.com/
[10] https://gsuite.google.com.au/
[11] https://www.teradici.com/

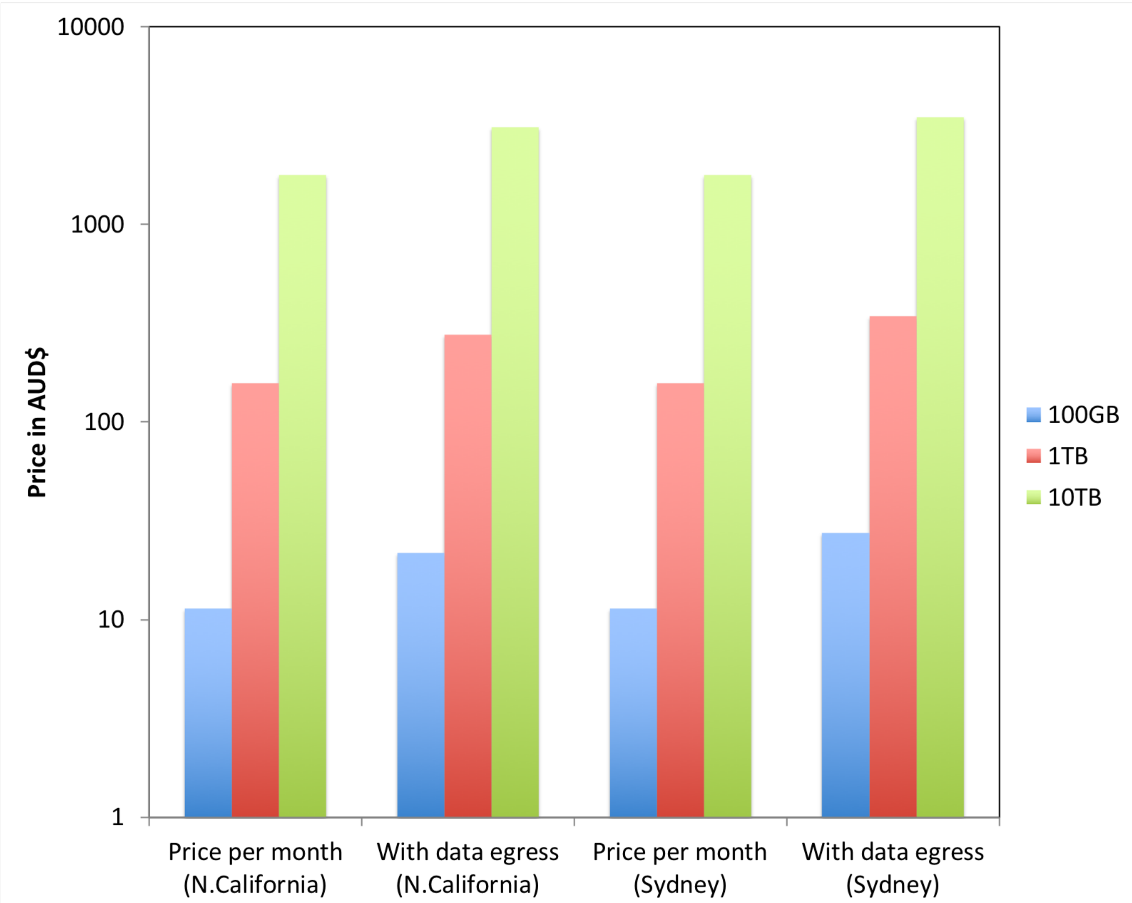


Figure 6.1 Shown are the base prices (in AUD$) of AWS Elastic Block Storage per GB-month, and the price combined with the egress charges for a single download of the full capacity in a month.

Table 6.1 This table shows the cost of operating a VM on AWS over four years. The three usage loads are maximum use (i.e. 24 hours per day, 7 days per week, 52 weeks per year, 4 years), high use (6 hours per day, 5 days per week, 42 weeks per year, 4 years), and low use (2 hours per day, 3 days per week, 32 weeks per year, 4 years), corresponding to those described in Table 4 of Meade & Fluke (2018). The costs include the g2.2xlarge VM, and assume a single full capacity download each month. The default setting was used for all other options. The AWS Simple Monthly Calculator was used to generate the prices, based on the Sydney data centre. Prices are shown in Australian dollars based on conversion rate as of 22/06/2018.

| | | **Month** | **Year 1** | **Year 2** | **Year 3** | **Year 4** |
|---|---|---|---|---|---|---|
| **Maximum use** | **100GB** | $901 | $10816 | $21632 | $32447 | $43263 |
| | **1TB** | $1231 | $14769 | $29537 | $44306 | $59074 |
| | **10TB** | $4829 | $57944 | $115889 | $173833 | $231777 |
| **High use** | **100GB** | $152 | $1827 | $3655 | $5482 | $7309 |
| | **1TB** | $468 | $5618 | $11236 | $16854 | $22473 |
| | **10TB** | $3990 | $47879 | $95758 | $143636 | $191515 |
| **Low use** | **100GB** | $46 | $547 | $1094 | $1641 | $2188 |
| | **1TB** | $361 | $4338 | $8676 | $13014 | $17351 |
| | **10TB** | $3873 | $46471 | $92941 | $139412 | $185882 |

# 7 Conclusions

## 7.1 Overview

The desktop computer is an easily overlooked part of the typical astronomy research workflow, yet it and the related technologies that support it have played a crucial role in helping our understanding of the universe at an unprecedented rate. However, because it is so ubiquitous, the inherent limitations often go unacknowledged.

To keep pace with the evolution of astronomy research, the way we use a desktop computer also needs to evolve. The display attached to a laptop or desktop computer can be augmented with UltraHD displays, or connected to TDWs, providing a more natural way to engage with astronomical imagery and data. The device itself can provide access to flexible and scalable compute resources that more closely match current actual demand, i.e. the local computer need not be so powerful as to accomplish all or most of the research tasks, but only so powerful to be able to access a VHD.

This research has investigated two particular approaches to resolve the over-reliance on local computers for conducting research. First, it is clear that using a display that is far smaller than the content needed to be displayed, be that Gigapixel imagery or heterogeneous content from images, movies and applications, results in a reduction in the time taken to assimilate the information when the information can be spread out and viewed on a TDW. More importantly, using the right display for the right content can greatly improve sense making and decision making in time-critical research tasks.

Second, just as centralised compute and storage has made it possible to store and process vast quantities of information from astronomical instruments and simulations, centralised cloud services now make it possible to move the functions of the local desktop into the cloud, where a VHD capable of performing the same functions can be deployed. The advantages of this approach include co-location of the desktop applications with the data on which they will operate. Security, scalability, scientific reproducibility and collaboration can also be significantly enhanced.

## 7.2 User Evaluation Methodology

Understanding the value of a new technology or methodology requires understanding how the user experience and performance will be affected. A positive user experience means that researchers are more likely to accept changes to a more familiar workflow. Similarly, improvements to a workflow that result in increased user performance provide a compelling reason to change. The user evaluation scenarios described by Lam et al. (2012) were used to devise empirical studies to investigate the technologies and methodologies reviewed in this research. The combination of user experiences and tracking user of performances showed that for certain astronomy visual inspection tasks, both user performance and the user experience were improved when using a TDW compared to a standard desktop display.

Communication through visualisation and collaborative data analysis – two more user evaluation scenarios described by Lam et al. (2012) – were used to evaluate the efficacy of a display ecology to support the *Deeper, Wider, Faster* fast transient observing campaign. This study not only confirmed the value of the advanced displays used in the observing runs, but also the important role a display ecology plays in improving the visual inspection workflow, including the refinement of automatic detection algorithms.

The advanced displays, in the form of a 98 Megapixel TDW and an immersive high definition projection display, supported the astronomers' need to work collaboratively with multiple image presentation formats, as well as their ability to communicate the salient information to their peers, and ultimately to the principal investigators. Most importantly, even though a purely objective measure of evaluating the efficacy of the display ecology was impossible to capture, the participating astronomers were convinced that the use of the display ecology significantly improved their ability to perform their tasks.

The demands of modern astronomy have forced the moving of digital workflows to data centres hosting HPC facilities with high performance networking and storage attached. Increasingly, these data centres also host cloud computing services, which are also tightly coupled with other resources within the data centre. This provides astronomers with more computation options than even before, where it is easier than ever to accommodate complex workflows that do not sit neatly in the typical batch job submission model of HPC. Cloud computing provides a flexible platform for trial and error experimentation with computation pipelines. Moreover, with the availability of GPU-enabled VHDs, cloud computing can be used to support almost the entire end-to-end digital workflow of most astronomers.

The key to the successes of these technologies lies in the user experiences. To this end, this research has focused on understanding how the user experience is changed when standard astronomy tasks are performed in such environments.

## 7.3 Ultra-high Resolution Displays

The rate of increase in the size and resolution of individual computer displays is considerably slower than the equivalent increase in the capability of telescopes. Despite the promise of machine learning and deep learning, it is expected that for some time yet, the role of the astronomer in inspecting critical visual data will remain, albeit augmented by advances in automatic feature detection. An astronomer is trained to understand the features of an image in the context of other, non-visual data, such as parallel research and emerging hypotheses, a feat still beyond the reach of neural networks. Therefore it will be important to furnish the astronomer with appropriate visual tools such as TDWs to improve science outcomes.

Further value of TDWs have been identified in this research. For example, a TDW can provide a useful tool for collaboration, where several astronomers can share a single display space containing Gigapixel imagery or multiple sources of related data. While two or three astronomers can reasonably share a standard desktop display, this becomes increasingly problematic as numbers increase. In many cases, this is solved using projectors, where the physical scale of the display is increased, but often not the resolution. A TDW can provide the same physical display scale, often at a comparable price point, but with far higher resolution. There is also the added benefit that a TDW does not suffer from shadows cast on the display surface by the participants.

Perhaps the most important consequence of this research is understanding the importance of considering the right display for the right content when confronted with a critical visually-based workflow. Designing a display ecology around the visual demands of an inspection workflow can greatly improve the rate and confidence of decision making under time-constrained conditions. This is accomplished by making critical information easily observable, and by improving the experience of the participants involved.

### 7.3.1 Tiled Display Walls

The last decade saw a rush to build TDWs in Australian research institutions, however, the enthusiasm was short-lived. Despite the dropping price of the base components, TDWs failed to attract many researchers, most likely due to their complexity to operate. The principal problem was that most of the TDWs were built on the idea that "If you build it, they will come."[1] However, without properly evaluating user engagement in terms of the scenarios described in Lam et al. (2012) in domains such as astronomy, it is not surprising that widespread success did not follow. Perhaps the major problem TDWs suffered was purely that of meeting the marketing hype, which suggested that TDWs would enable discoveries that would otherwise be impossible. While this may be true, no such discovery in astronomy has been attributed to the use of a TDW to date.

In fact, the key value of a TDW lies in the margins of achieving discoveries. By making it easier to observe small features in large images, a TDW might allow an astronomer to find a feature of interest more quickly than they would have on a standard desktop display. Just as importantly, by displaying content at native or close to native resolution, a TDW might reduce the chance of missing a discovery, or recognising an artificial artefact from an astronomical phenomenon because of the context within the larger image, as happened during the *Deeper, Wider, Faster* observing campaign (Meade et al., 2017, see Chapter 5, Section 5.2).

This research has shown that users can perform certain visual inspection tasks, still a critical part of many astronomers' workflow, better with a TDW than with a standard desktop display. It has also shown that users prefer conducting these tasks with a TDW, where they can employ physical navigation to search for objects, rather than rely on virtual navigation.

Meade et al. (2014, see Chapter 4, Section 4.3) addressed the lack of published evidence that a TDW could improve scientific outcomes in astronomy. This was accomplished by comparing a 98 Megapixel TDW with a standard desktop display in head-to-head astronomical tasks. This study showed that TDWs not only improved user performance, but provided a more satisfactory user experience.

Many of the issues that plagued the first TDWs have been addressed to varying degrees. Firstly, the price of building a TDW has plunged, e.g. a 100 Megapixel display wall recently built at Swinburne University of Technology cost around 10% of the 98 Megapixel OzIPortal, built at the University of Melbourne almost ten years ago. Secondly, the software needd to operate a TDW, in particular SAGE2, is far more user-friendly than previous

[1] From the movie "Field of Dreams", 1989

incarnations. Screen elements such as bezel thickness and resolution have improved such that current UltraHD screens are now cheaper to buy and run, physically lighter, and have much thinner bezels, significantly improving the tiled appearance.

However, the research drivers for TDWs are still emerging slowly. While campaigns such as *Deeper, Wider, Faster* are becoming more common, it seems other research areas have not yet felt the need to adopt TDWs as a part of their digital workflow. Perhaps as the demand for processing large volumes of visual information quickly and accurately becomes more critical to more workflows, TDWs might see more uptake in the future, especially in the area of communication and collaborative research.

### 7.3.2 Display Ecologies

Desktop displays, UltraHD displays, projectors and TDWs are all valuable tools in the right context. However, most researchers try to use a single display, typically their local desktop computer, to accommodate their full visual workflow. When data becomes big in terms of volume, velocity or variety, that is, when the local computer cannot handle the data ingest volume, or rate, or the data comes from different sources in different forms. In this case it becomes important to consider the display ecology needed to understand it. For example, a TDW provides a far larger physical display surface than a standard desktop display, while maintaining the pixel density of the standard desktop display. A projector can produce a very large image, but the resolution is the same as the standard desktop display. A standard desktop display can display any content, but is physically small, making it ideal for a single user displaying low resolution content, such as small images and text. While any one of these can display any digital content, they are not equally suited to display all content. This can become crucial when time becomes a factor in making discoveries.

The fast transient search program, *Deeper, Wider, Faster*, provided an opportunity to explore the use of a display ecology, where two advanced displays were used as part of a visual inspection workflow, and where time-critical image inspection was key to the success of the program.

Meade et al. (2017, see Chapter 5, Section 5.2) showed that the *Deeper, Wider, Faster* astronomers found the display ecology provided improved their ability to complete the visual inspection component of the workflow more quickly and with more confidence than they could have achieved using standard desktop displays only. Furthermore, the display ecology supported the training of participants in the visual inspection workflow, bringing them up to speed more quickly than using a standard desktop display. But the most valuable use of the display ecology was the ability for many astronomers, up to 20 in come cases, to work collaboratively to analyse the incoming data quickly and efficiently, such that the principal investigators were able to reach key decisions quickly and with a high degree of confidence.

It is the nature of a display ecology to evolve with the changing needs of the program it supports, and this occurred several times during the *Deeper, Wider, Faster* observing runs. Lessons learned from each run informed the configuration for the next run, and so on. Changes to the automatic detection pipeline changed the focus of parts of the visual inspection workflow, which also informed the display ecology. Personal experience from the participants was also used to enhance subsequent designs.

## 7.4 Virtual Hosted Desktops

Astronomers have been ready adopters of remote computing technology, as it has become essential in many research workflows. However, there remains a reluctance to “move beyond the desktop” and utilise a VHD as an alternative to a local desktop. Some of this reluctance comes from unsatisfactory past experiences, and some from an assumption that the performance (usually in terms of latency) of a remote desktop will make it a poor facsimile of a local desktop. However, this research has shown that using a GPU-enabled VHD can provide a user experience that can match and in some cases exceed that of local desktops.

The assumption made by many researchers that virtualisation will drastically reduce the experience of a VHD is not borne out in studies conducted during this research. In fact, the apparent lack of impact due to virtualisation came as a surprise to several participants. Using a suitably provisioned VHD to perform astronomy tasks not only provides a satisfactory experience, the flexible nature of cloud-based resources can provide significant cost benefits as well.

It is important to note that VHDs are themselves no better or worse than the infrastructure on which they run. That is, when matched with a local desktop for systems specifications, e.g. processor speed, RAM type and amount, etc., they perform as well as the local machine. The critical consideration is that the difference is undetectable for most practical purposes. This is not to say that issues such as network latency cannot be exposed on a VHD with some effort, but more that for the majority of common tasks an astronomer might perform, no difference will be noticed. This makes a VHD a viable alternative to a local desktop computer.

However, there are several ways that a VHD might be a better option than a local computer:

- Purchase of a local computer requires an upfront commitment of financial resources, with little opportunity to correct over- or under-spending. A VHD is effectively "rented" from the cloud service provider, and therefore can be relinquished as required, and is typically operated in a "pay-as-you-go" model. For federated cloud services such as the Nectar Research Cloud, the researcher is not charged directly, and yet still has the opportunity to "fail fast", creating and destroying VMs with ease.
- The pressure of making the right choice in purchasing a local computer can result in the over estimation of system specifications and benchmarks. As seen in Chapter 6 Section 6.5, benchmarking a system can provide misleading results, and only by testing in context, i.e. with the intended workload, can a true understanding of a systems capabilities and limitations be understood. This limitation is just as true for a VHD, but the consequences are less severe, as an inappropriately resourced VM can be reinstantiated on more appropriate virtual hardware.
- Resources can be more easily shared in a virtualised environment than with a physical desktop computer. For example, a high-end graphics workstation purchased for a research group can only have a single researcher using it at a time. It would also require the researchers to be physically present at the machine to use it. When physical access is not possible, such as out-of-hours, the machine might go for extended periods of little or no utilisation. In contrast, a VHD can be accessed by several people at the same time, and the underlying server can often host several virtual desktops on the same hardware simultaneously, allowing several researchers to share the resource at the same time. Like a physical desktop, a VM providing the VHD might be under utilised when researchers are not using it, but being able to access the resource from any networked location increases the opportunity to use it. The resources of an under-used VHD can also be redeployed easily to other functions, such as supporting a compute cluster.

### 7.4.1 The e-Research Landscape in Australia

The research community in Australia has access to Nectar, a national federated research cloud establish by the Australian Federal Government in 2009. This multi-node dedicated research cloud provides tens of thousands of virtual cores and supports over 12,000 researchers across the country. Many of the nodes also provide additional cores to their local research communities as well. The highly successful Virtual Laboratory program resulted in new research tools and collaborations for many research domains, with the Australian astronomy community supported by the world-class ASVO, which in turn provides access to services such as the TAO, Skymapper and ASKAP data collections.

### 7.4.2 Research Cloud Data Communities

Shared access to high-value datasets, along with shared tools and workflows to exploit the data, are key drivers for data communities. The Nectar Research Cloud has fostered the formation of research data communities in Australia, by establishing virtual laboratories such as the ASVO, and supporting nationally meritorious research undertaken by Australian researchers.

Meade et al. (2013) describes the basis for the formation of the communities on the Nectar Research Cloud. The key elements identified include:

- Researchers can "fail fast", meaning they can quickly establish a digital workflow using the cloud infrastructure, which can be rapidly evolved in a *create, refine, terminate* cycle of improvement.
- The Nectar Research Cloud is built on big data infrastructure, meaning VMs are hosted on infrastructure in the same data centre, connected over the same high speed network, as the high-capacity storage service, providing petabytes of storage that can be accessed directly from the VMs. The data centres are enterprise grade, and provide a high level of security and reliability.
- VMs can be created, cloned, deleted, archived and shared with collaborators around the world with ease. Access to resources for meritorious projects are free at the point of service, meaning researchers are free to explore options as they see fit, without having to be concerned about the financial impact.

### 7.4.3 Seeing the Big Picture: A Digital Desktop for Researchers

One of the main barriers to adoption of cloud-based services is the lack of experience within the research community in using remote services. While virtual laboratories provide portals to interface with their resources, this is not available to most research projects. Most researchers are quite familiar with a graphical interface to their applications. Desktops like Microsoft Windows and Mac OSX are among the most common, with Linux variants KDE and Gnome popular with Linux users. VHDs are available on the Nectar Research Cloud, though their use requires additional steps to make them work.

With the exception of Mac OSX, all standard desktops available on a physical computer can be provisioned as a VHD. The principal difference between a VHD and a local computer desktop is the method of connecting. A VHD requires a local computer to run client software to connect to the remote service, and then to pass keyboard and mouse commands through to the VHD, which in turns sends screen updates to the local machine. With a suitable network and local machine, the user experience of a VHD can be equivalent to a local desktop.

Taking this a step further, Meade et al. (2015) investigated the potential for a VHD to drive a local UltraHD display. It was first identified that an UltraHD display, while considerably smaller than some TDWs, such as the 98 Megapixel OzIPortal, the increase in screen real estate over a standard desktop display showed similar, albeit smaller, improvements when used to repeat the investigation conducted in Meade et al. (2014, see Chapter 4, Section 4.3). At the time, the Nectar Research Cloud and other cloud services like AWS, did not provide suitable GPU accelerated VHDs, so a VMWare server hosted in the same data centre was used instead. A VHD with GPU acceleration was provisioned on the server and used to drive an UltraHD display located nearby over the University of Melbourne network. The study confirmed the viability of using a GPU accelerated VHD over this network, with the screen resolution of 3840 × 2160 operating at 60 frames per second requiring less than 50 Megabits per second to provide a seamless experience.

The study also investigated the state of the Australian research network (AARNet) connecting research institutions to the national peak computing facilities, Pawsey Supercomputing Facility in Western Australia, and NCI in Canberra. At the time, the more mature connection between the University of Melbourne and NCI showed a highly stable network capable of supporting several VHD streams at UltraHD resolution, while the less mature connection to Pawsey Supercomputing Facility would not have been capable of sustaining a VHD of that size. However, the issues plaguing the Pawsey connection at that time have largely been addressed, and it is likely the network is more stable now. This suggests it should be possible to drive a TDW made up of UltraHD displays using a cluster of GPU-enabled VHDs from a national peak facility.

### 7.4.4 Evaluating Virtual Hosted Desktops for Graphics-intensive Astronomy

Even as the Meade et al. (2015) was being conducted, both the Nectar Research Cloud and AWS had already announced plans to support GPU-enabled VMs, specifically for the purpose of supporting VHDs. It still took over a year to get a GPU-enabled desktop to work with either cloud providers' service, but eventually the opportunity to explore GPU-enabled VHDs became possible.

Meade & Fluke (2018, see Chapter 6, Section 6.5) found that a VHD using GPU pass-through technology (where a single physical GPU was made available exclusively to a VM) exhibited such low latency as to appear almost indistinguishable from interacting with a local desktop. More importantly, this remained true for graphical 3D applications such as volume renderers.

Comparing 3D graphics benchmarks for local laptop computers with GPU-enabled VHDs revealed that while the frame rates for VHDs were not as high as the local computers, they still performed well. Using a cohort of astronomers, the study showed that the user experience of VHDs was equally acceptable to the participants as the local desktops for basic astronomy tasks, such as simple image alignment. As expected, the participants all found that as the graphical demand of the 3D application used was increased, the user experience decreased. An unexpected outcome however, was that the user experience for the VHDs was noticeably better as the load increased than for the local laptops. Comparing system performance with the reported user ratings confirmed the user experience, i.e. the load on the GPU increased, the frame rate of the GPU dropped more rapidly on the local laptops than on the VHDs.

The study identified a clear risk of over valuing benchmarks when determining how suitable a computer is for some tasks. It also dispelled the assumption that virtualisation would dramatically impact the user experience, by showing that for many research tasks, most astronomers would be unlikely to even notice a difference. When this result is combined with a cost benefit analysis of purchasing a powerful graphics workstation outright versus renting a similarly capable cloud based VM, it is clear that a VHD should be considered as a highly competitive option.

This led to the conclusion that a GPU-enabled VHD could indeed become a viable alternative to a local desktop computer for many astronomy research workflows.

## 7.5 Future Research Directions

Throughout this research, several questions arose that were beyond the scope of the individual investigations presented herein. These questions are still worthy of further investigation and are presented here as future directions of research within each of the subject areas listed below.

### 7.5.1 Tiled Display Walls

TDWs were initially intended to be used for large-scale visualisation, far grander than a standard desktop display. However, as many researchers use multiple screens attached to a local computer, or choose larger format displays such as UltraHD screens, the line between a TDW and a multi-screen, or large format screen workstation is rapidly blurring. Other technologies to enhance the ability to consume visual data have also emerged to change the potential use of TDWs.

Due to the time pressures on the participants, it was impractical to observe the TDW use over longer periods of time than 30-45 minutes. However, it would be useful to understand how a TDW might be used if it was deliberately included within the research workflow of a research group over a period on weeks or months. If accessibility and usability were no longer an impediment to use, would the research group adopt the TDW going forward if the option were available? Would it be possible to objectively determine if the group's performance was improved as a result of their experience with the TDW?

#### Virtual Reality

Virtual Reality (VR) headsets were a technology that suffered a similarly disappointing uptake when they were first proposed as the next big research tool. However, in 2012, the Oculus Rift[2] was launched on crowd fund-raising site, Kickstarter[3]. Launched by Facebook in 2016, the Oculus Rift had already generated considerable interest in virtual reality. The main focus for VR was in its potential for immersive gaming, but there were many potential research applications as well. Of particular relevance to this research is the fact that a VR headset can essential virtualise a display screen, in much the way a VHD virtualises the local computer desktop.

It is worth noting that despite the drop in the price of the commodity components of a TDW over the last decade, a VR headset is still a considerably cheaper option.

In the context of a TDW, the scale of visualisation and resolution afforded by clustering high-resolution displays together can be achieved through virtual means by way of a VR headset. For example, with a VR headset, a user's entire view is occupied by the display within the headset, but the movement of the user's head and body effectively disguises the presence of these small displays, essentially convincing the wearer that the view before them is a form of "reality". As the resolution of the headset displays improves, and the field of view widens, this effect will become even more compelling.

[2] `https://www.oculus.com/`

[3] `https://www.kickstarter.com/`

A VR headset therefore, can simulate the scale and resolution of a TDW of effectively any size and resolution. Physical navigation is translated into virtual navigation by the VR environment, making the process transparent to the wearer. Collaboration and training can be achieved within the environment with even more ease than for a TDW, as the physical presence of the participants is not required, nor do participants impede one another, and the environment can be shared from anywhere in the world.

However, despite the improvements to VR headsets over the last few years, driven by the rapid uptake in the wider community, it is unlikely VR headsets will replace standard desktop displays in the short- to mid-term, which suggests that large collaborative displays, either in the form of projection, large UltraHD displays, or TDWs, will likely continue to be useful for some time.

#### Augmented Reality and Mixed Reality

AR displays entered the public consciousness with the announcement of Google Glass in 2013. Despite the project being shelved by Google in 2015 for unspecified reasons, the technology did show considerable promise, and remains part of Google's AR plans. The principal difference between AR and Mixed Reality (MR) is the way digital content is displayed in relation to the real world. In both environments, the "reality" component comes from the real world, either passively through transparent displays, or actively via attached camera(s). In the case of augmented reality, the digital content is overlaid on the real world content, often displaying related information, such as GPS coordinates, but the content is not directly anchored to the real world. With MR, the digital content is connected to the real world. For example, with Microsoft Hololens, digital content such as a 3D model can appear to be sitting on a real world table. This effect is reinforced when the wearer approaches the table, or moves around it, because the MR system updates the headset display to adjust the appearance of the 3D object to match the user's movements, making the 3D object appear to be a part of the physical world.

AR and MR can also be achieved without the use of head mounted displays, and many popular smart phone and tablet applications exploit built in motion sensors to achieve the same effect.

There is an opportunity to investigate the combination of AR/MR with TDWs, where features displayed on the TDW image can be tagged and displayed as a digital overlay within a MR headset. In the case of the study conducted in Meade et al. (2014, see Chapter 4, Section 4.3), an astronomer could tag a feature of interest, which could be picked up by another astronomer, either at a later time and/or another location, making the collaborative searching more effective.

#### The Ultimate Display

Following on from Sutherland (1965), Fluke & Barnes (2016) have posited the requirements for an "Ultimate Display", that immerses a researcher in a distraction free environment in which the only visual input experienced are the data themselves. While such a display environment is relatively easy to conceive, there remain technological hurdles to overcome before it is a reality.

However, even if such technology is possible, the question remains as to how likely it will be widely adopted. For researchers to embrace the opportunities such advanced displays afford, they will need to appreciate the advantages in their day-to-day research. As the costs of displays fall, in both the purchase price and ongoing costs, it is conceivable that the physical wall space within a standard office might be overtaken by displays. With an expanded display environment, researchers would be able to anchor digital content to their display surface. This would allow them to capitilise on their spatial memory for rapidly retrieving content. It would also allow them to personalise the display environment to suit their preferred style. It may be that this final point is the key to widespread adoption of TDWs as part of a researcher's workflow.

### 7.5.2 Display Ecologies

Designing display ecologies to better align with visual inspection workflows could benefit from incorporating VR/AR/MR headsets. For example, given the collaborative nature of the *Deeper, Wider, Faster* observing campaign, VR might not fit easily into the dynamic environment with people moving around in the space. In this case there is an opportunity to use MR to link automatic detection pipelines to the headset of an inspector looking at potential targets. Critical information about the target being viewed, such as light curves and details of the position from a catalogue being displayed in real-time, could decrease the time taken to identify targets of interest. Communication between inspectors can also be enhanced, as objects of interest can be tagged and easily shared with others, even in very large images.

### 7.5.3 Virtual Hosted Desktops

Investigating the use of VHDs to support astronomy research workflows has resulted in several questions about their use.

Firstly, if VHDs, by virtue of being located in a highly connected data centre, can be used to access data that is larger than the capacity of a local disk drive, what applications might take advantage of this? Can a VHD be linked directly to a HPC cluster, allowing direct job submission, e.g. to a HPC batch queue?

Secondly, as participant time during the Meade & Fluke (2018, see Chapter 6, Section 6.5) study was limited to 30-40 minutes, it would be valuable to observe participants' usage over longer periods of time, such as week or months. Extended use might expose potential limitations such as network latency that were not apparent in the controlled study. This usage would also give a more useful indication as to typical usage patterns, making cost comparisons with outright purchase more meaningful.

Also, there are a number of competing approaches to desktop and application virtualisation, such as containerisation and application streaming (see Section 2.2), which might provide better options for different circumstances. These technologies can be valuable in mitigating the complexities of creating and maintaining astronomy software, potentially making cloud adoption more appealing to astronomers. They can also add to the reproducibility of scientific outcomes by allowing researchers to capture and publish the operating environment with links to online datasets used to produce their results. Examples of the use of containers in astronomy can be found in Morris et al. (2017).

## 7.6 Summary

The use of local desktop and laptop computers will remain an essential supporting technology in the astronomy research environment for some time. However, it is now possible to overcome some of the limitations associated with the local desktop. TDWs provide a platform to display billions of pixels worth of digital content in the form of Gigapixel images, high definition movies, animations, web pages and text, or any mix of these. They can be displayed at a scale that makes inspection, especially for collaborations, easy and enjoyable. More importantly, where time-critical decisions need to be made, a TDW as part of a wider display ecology supporting a visual inspection workflow can improve the rate and confidence of decision-making.

The limitations of the desktop computer include not only the number of processing cores, RAM and hard disk space, but also the ability to change these features as required. While some components can be replaced, others cannot. Cloud computing alleviates this concern by allowing resources to be repurposed as required, expanding and contracting as demand shifts. This is possible from the most basic single core VM, to a cluster comprised of thousands of cores.

Shifting the function of a local desktop to the cloud in the form of a VHD can provide all the capabilities of a local desktop, but with the added benefits of flexibility, optimised financial commitment, and direct access to larger datasets.

The astronomy community should no longer be limited by a local computer. It is now possible to move beyond the desktop and embrace the power of cloud computing and the visual scale of ultra-high resolution displays.

# Bibliography


Agana, A., Davalath, M., McNamara, A., & Parke, F. 2010, The effect of tiled display on performance in multi-screen immersive virtual environments, in 2010 IEEE Virtual Reality Conference (VR) (IEEE), 249–250

Almes, G. T., Jent, D., & Stewart, C. A. 2011, Campus Bridging: Data and Networking Issues Workshop Report, Tech. rep., Indiana University

Andreoni, I., & Cooke, J. 2018, The Deeper Wider Faster program: chasing the fastest bursts in the Universe, in IAU Symposium 339

Andrews, C., Endert, A., & North, C. 2010, Space to think: large high-resolution displays for sensemaking, in Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '10 (ACM Press), 55

Andrews, C., Endert, A., Yost, B., & North, C. 2011, Information visualization on large, high-resolution displays: Issues, challenges, and opportunities, Information Visualization, 10, 341

Armbrust, M., Fox, O., Griffith, R., et al. 2009, Above the clouds: a Berkeley view of cloud computing, Technical Report No. UCB/EECS-2009-28, University of California, Berkeley

Armstrong, P., Agarwal, A., Bishop, A., et al. 2010, Cloud Scheduler: a resource manager for distributed compute clouds, arXiv preprint arXiv:1007.0050

Ball, N. M. 2012a, Astroinformatics, Cloud Computing, and New Science at the Canadian Astronomy Data Centre, in American Astronomical Society Meeting Abstracts, Vol. 219, American Astronomical Society Meeting Abstracts #219, #145.11

Ball, N. M. 2012b, Focus Demo: CANFAR+Skytree: A Cloud Computing and Data Mining System for Astronomy, in Astronomical Data Analysis Software & Systems (ADASS) XXII, ASP Conference Proceedings

Ball, N. M., & Brunner, R. J. 2010, Data Mining and Machine Learning in Astronomy, International Journal of Modern Physics D, 19, 1049

Ball, R., & North, C. 2005a, Analysis of user behavior on high-resolution tiled displays, in Human-Computer Interaction-INTERACT 2005 (Springer), 350

Ball, R., & North, C. 2005b, Effects of tiled high-resolution display on basic visualization and navigation tasks, in Proceedings of the 2005 Conference on Human Factors in Computing Systems (Portland, Oregon, USA: ACM Press), 1196

Ball, R., North, C., & Bowman, D. A. 2007, Move to improve: promoting physical navigation to increase user performance with large displays, in Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '07 (ACM Press), 191

Bell, G., Hey, T., & Szalay, A. 2009, Beyond the data deluge, Science, 323, 1297

Bernyk, M., Croton, D. J., Tonini, C., et al. 2016, The theoretical astrophysical observatory: cloud-based mock galaxy catalogs, The Astrophysical Journal Supplement Series, 223, 9

Berriman, B., & Groom, S. L. 2011, How will astronomy archives survive the data tsunami?, Communications of the ACM, 54, 52

Berriman, G. B., Brinkworth, C., Gelino, D., et al. 2012, A Tale Of 160 Scientists, Three Applications, A Workshop and A Cloud, in Astronomical Data Analysis Software and Systems XXII

Berriman, G. B., Deelman, E., Juve, G., Rynge, M., & Vöckler, J. S. 2013, The application of cloud computing to scientific workflows: a study of cost and performance, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 371

Berriman, G. B., & Good, J. C. 2017, The Application of the Montage Image Mosaic Engine to the Visualization of Astronomical Images, Publications of the Astronomical Society of the Pacific, 129, 058006

Berriman, G. B., Juve, G., Deelman, E., Regelson, M., & Plavchan, P. 2010, The application of cloud computing to astronomy: A study of cost and performance, in e-Science Workshops, 2010 Sixth IEEE International Conference on, 1–7

Bi, X., & Balakrishnan, R. 2009, Comparing usage of a large high-resolution display to single or dual desktop displays for daily work, in Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '09 (ACM Press), 1005

Bi, X. J., Bae, S. H., & Balakrishnan, R. 2014, WallTop: Managing Overflowing Windows on a Large Display, Human-Computer Interaction, 29, 153

Boas, Y. A. G. V. 2013, Overview of Virtual Reality Technologies, in Interactive Multimedia Conference 2013, accessed from `http://www.dmarlett.com/s/yavb1g12_25879847_finalpaper.pdf`

Borne, K. 2008, Scientific Data Mining in Astronomy, in Next Generation of Data Mining, ed. H. Kargupta, J. Han, P. Yu, R. Motwani, & V. Kumar, Vol. 20084157 (Chapman and Hall/CRC)

Borne, K. D. 2010, Astroinformatics: data-oriented astronomy research and education, Earth Science Informatics, 3, 5

Boswell, R., & Gardner, H. 2001, The Wedge Virtual Reality theatre, in Apple University Consortium Academic and Developers Conference

Brescia, M., Cavuoti, S., Djorgovski, G. S., et al. 2012a, Extracting Knowledge from Massive Astronomical Data Sets, in Astrostatistics and Data Mining, ed. L. M. Sarro, L. Eyer, W. O'Mullane, & J. De Ridder, 31

Brescia, M., Longo, G., Castellani, M., et al. 2012b, DAME: A Distributed Web Based Framework for Knowledge Discovery in Databases, Memorie della Societa Astronomica Italiana Supplementi, 19, 324

Brescia, M., Longo, G., & Pasian, F. 2010, Mining knowledge in astrophysical massive data sets, Nuclear Instruments and Methods in Physics Research A, 623, 845

Brown, M., Majumder, A., & Yang, R. 2005, Camera-based calibration techniques for seamless multiprojector displays, IEEE Transactions on Visualization and Computer Graphics, 11, 193

Brunner, R. J., Djorgovski, S. G., Prince, T. A., & Szalay, A. S. 2002, Massive datasets in astronomy, in Handbook of massive data sets (Springer), 931–979

Butt, S., Lagar-Cavilla, H. A., Srivastava, A., & Ganapathy, V. 2012, Self-service cloud computing, in Proceedings of the 2012 ACM conference on Computer and communications security (ACM), 253–264

Buyya, R., Yeo, C. S., & Venugopal, S. 2008, Market-oriented cloud computing: Vision, hype, and reality for delivering it services as computing utilities, in High Performance Computing and Communications, 2008. HPCC'08. 10th IEEE International Conference on (Ieee), 5–13

Buyya, R., Yeo, C. S., Venugopal, S., Broberg, J., & Brandic, I. 2009, Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility, Future Generation Computer Systems, 25, 599

Buzen, J. P., & Gagliardi, U. O. 1973, The evolution of virtual machine architecture, in Proceedings of the June 4-8, 1973, national computer conference and exposition (ACM), 291–299

Callahan, S. P., Freire, J., Santos, E., et al. 2006, VisTrails: visualization meets data management, in Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data, SIGMOD '06 (ACM Press), 745

Chen, M., Mao, S., & Liu, Y. 2014, Big Data: A Survey, Mobile Networks and Applications, 19, 171

Chen, X.-W., & Lin, X. 2014, Big Data Deep Learning: Challenges and Perspectives, IEEE Access, 2, 514

Chirigati, F., Capone, R., Rampin, R., Freire, J., & Shasha, D. 2016, A collaborative approach to computational reproducibility, Information Systems, 59, 95

Chu, R., Tenedorio, D., Schulze, J. P., et al. 2008, Optimized rendering for a three-dimensional videoconferencing system, in eScience, 2008. IEEE Fourth International Conference on (IEEE), 540–546

Chung, J. C., Harris, M. R., Brooks, F. P., et al. 1989, Exploring virtual worlds with head-mounted displays, in OE/LASE'89, 15-20 Jan., Los Angeles. CA (International Society for Optics and Photonics), 42–52

Cohen, J., Filippis, I., Woodbridge, M., et al. 2013, RAPPORT: running scientific high-performance computing applications on the cloud, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 371

Cruz-Neira, C., Sandin, D. J., DeFanti, T. A., Kenyon, R. V., & Hart, J. C. 1992, The CAVE: audio visual experience automatic virtual environment, Communications of the ACM, 35, 64

Cui, C.-Z., He, B.-L., Li, C.-H., et al. 2017, AstroCloud: A Distributed Cloud Computing and Application Platform for Astronomy, in Proceedings of the 3rd International Conference on Wireless Communication and Sensor Networks (WCSN 2016), Advances in Computer Science Research (Atlantis Press)

Daniels, J. 2009, Server virtualization architecture and implementation, Crossroads, 16, 8

de Almeida, R. A., Pillias, C., Pietriga, E., & Cubaud, P. 2012, Looking behind bezels: french windows for wall displays, in Proceedings of the International Working Conference on Advanced Visual Interfaces, AVI '12 (ACM Press), 124

De Pascale, M. 2013, Automatic analysis of big data sample in astronomy., PhD Workshop Proc

Dean, J., & Ghemawat, S. 2004, MapReduce: simplified data processing on large clusters, Communications of the ACM, 51, 107

Deelman, E., Singh, G., Livny, M., Berriman, B., & Good, J. 2008, The cost of doing science on the cloud: the montage example, in Proceedings of the 2008 ACM/IEEE conference on Supercomputing (IEEE Press), 50

DeFanti, T., Leigh, J., Renambot, L., et al. 2009a, The OptIPortal, a scalable visualization, storage, and computing interface device for the OptiPuter, Future Generation Computer Systems, 25, 114

DeFanti, T., Dawe, G., Sandin, D., et al. 2009b, The StarCAVE, a third-generation CAVE and virtual reality OptIPortal, Future Generation Computer Systems, 25, 169

DeFanti, T. A., Acevedo, D., Ainsworth, R. A., et al. 2010, The future of the CAVE, Central European Journal of Engineering, 1, 16

Deshpande, S., & Daly, S. 2010, 80.1: Technologies and Applications for Large Sized High Resolution Tiled Display System, SID Symposium Digest of Technical Papers, 41, 1188

Deshpande, S., Yuan, C., Daly, S., & Sezan, I. 2009, A Large Ultra High Resolution Tiled Display System: Architecture, Technologies, Applications, and Tools, in Proc. of 16th Int'l Display Workshops, Miyazaki, Japan

Doerr, K., & Kuester, F. 2011, CGLX: A Scalable, High-Performance Visualization Framework for Networked Display Environments, IEEE Transactions on Visualization and Computer Graphics, 17, 320

Druken, K. A., Trenham, C. E., Steer, A., et al. 2016, Lowering the Barriers to Using Data: Enabling Desktop-based HPD Science through Virtual Environments and Web Data Services, AGU Fall Meeting Abstracts

Duato, J., Yalamanchili, S., & Ni, L. M. 1997, Interconnection networks: an engineering approach (Los Alamitos, Calif: IEEE Computer Society Press)

Emsley, I., & De Roure, D. 2017, A framework for the preservation of a Docker container, in 12th International Digital Curation Conference

Faith, R. E., & Martin, K. E. 2001, Xdmx: Distributed, multi-head X

Febretti, A. 2017, Multiview Immersion in Hybrid Reality Environments, PhD Thesis, Argonne National Laboratory

Febretti, A., Nishimoto, A., Thigpen, T., et al. 2013, CAVE2: a hybrid reality environment for immersive simulation and information analysis, in Proc. SPIE 8649, ed. M. Dolinsky & I. E. McDowall, 864903

Feigelson, E. D., & Babu, G. J. 2012, Big data in astronomy, Significance, 9, 22

Fernique, P., Allen, M., Boch, T., et al. 2017, HiPSHierarchical Progressive Survey Version 1.0, IVOA Recommendation 19 May 2017

Fluke, C. J., & Barnes, D. G. 2016, The Ultimate Display, in Proceedings of Astronomical Data Analysis Software and Systems XXV (Sydney, NSW, Australia: ASP Conf. Series)

Fluke, C. J., Bourke, P. D., & O'Donovan, D. 2006, Future Directions in Astronomy Visualization, Publications of the Astronomical Society of Australia, 23, 12

Fujiwara, Y., Date, S., Ichikawa, K., & Takemura, H. 2011, A Multi-Application Controller for SAGE-enabled Tiled Display Wall in Wide-area Distributed Computing Environments, Journal of Information Processing Systems, 7, 581

Gardner, H. J., Boswell, R. W., & Whitehouse, D. 1999, The WEDGE Emmersive Projection Theatre

Goldberg, R. P. 1973, Architecture of virtual machines, in Proceedings of the workshop on virtual computer systems (ACM), 74–112

Goodman, A. A. 2012, Principles of High-Dimensional Data Visualization in Astronomy, Astronomische Nachrichten, 333, 505

Goodman, A. A., Udomprasert, P. S., Kent, B., Sathiapal, H., & Smareglia, R. 2011, Astronomy Visualization for Education and Outreach, in Astronomical Data Analysis Software and Systems, Vol. 442 (Astronomical Society of the Pacific), 659–662

Goscinski, W. J., Paterson, D., Hines, C., et al. 2015, MASSIVE: an HPC Collaboration to Underpin Synchrotron Science, in Proceedings of ICALEPCS2015, Melbourne, Australia

Green, A. W., Mannering, E., Harischandra, L., et al. 2016, What will the future of cloud-based astronomical data processing look like?, Proceedings of the International Astronomical Union, 12, 27

Grüninger, J., & Krüger, J. 2013, The impact of display bezels on stereoscopic vision for tiled displays, in Proceedings of the 19th ACM Symposium on Virtual Reality Software and Technology, VRST '13 (ACM Press), 241

Hagen, T. M. S. 2011, Interactive Visualization on High-Resolution Tiled Display Walls with Network Accessible Compute-and Display-Resources, PhD thesis, University of Tromsø

Hancock, C. 2012, New Broadband Frontiers, RADCOMMS2012, Melbourne

Hassan, A., & Fluke, C. 2011, Scientific Visualization in Astronomy: Towards the Petascale Astronomy Era, Publications of the Astronomical Society of Australia, 28, 150

Hassan, A., Fluke, C., & Barnes, D. 2010, GPU-Based Volume Rendering of Noisy Multi-Spectral Astronomical Data, in Astronomical Data Analysis Software and Systems XIX, Vol. 434, Sapporo, Japan

Heirich, A., & Moll, L. 1999, Scalable distributed visualization using off-the-shelf components, in Proceedings of the 1999 IEEE symposium on Parallel visualization and graphics (IEEE Computer Society), 55–59

Hereld, M., Judson, I., & Stevens, R. 2000, Introduction to building projection-based tiled display systems, IEEE Computer Graphics and Applications, 20, 22

Herwig, F., Andrassy, R., Annau, N., et al. 2018, Cyberhubs: Virtual Research Environments for Astronomy, ApJ Supplement Special Issue on Data (accepted)

Hey, T., Tansley, S., & Tolle, K. M. 2009, Jim Gray on eScience: a transformed scientific method., in The Fourth Paradigm, ed. T. Hey, S. Tansley, & K. M. Tolle (Microsoft Research)

Hiden, H., Woodman, S., Watson, P., & Cala, J. 2013, Developing cloud applications using the e-Science Central platform, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 371

Hoffa, C., Mehta, G., Freeman, T., et al. 2008, On the use of cloud computing for scientific workflows, in eScience, 2008. IEEE Fourth International Conference on (IEEE), 640–645

Hoffman, A. R., & Traub, J. F. 1989, Supercomputers: Directions in Technology and Applications (National Academy Press)

Hong, C.-H., Spence, I., & Nikolopoulos, D. S. 2017, GPU Virtualization and Scheduling Methods: A Comprehensive Survey, ACM Computing Surveys, 50, 1

Hsieh, T.-J., Chang, Y.-L., & Huang, B. 2011, Visual Analytics of Terrestrial Lidar Data for Cliff Erosion Assessment on Large Displays, in Satellite Data Compression, Communications, and Processing Vii, ed. B. Huang, A. J. Plaza, & C. Thiebaut, Vol. 8157 (Bellingham: Spie-Int Soc Optical Engineering)

Hsieh, T. J., Liang, W. Y., Chang, Y. L., Satria, M. T., & Huang, B. M. 2013, Parallel tsunami simulation and visualization on tiled display wall using OpenGL Shading Language, Journal of the Chinese Institute of Engineers, 36, 202

Huffman, J. N., Forsberg, A. S., Head, J. W., Dickson, J. L., & Fassett, C. I. 2009, Testing Geoscience Data Visualization Systems for Geological Mapping and Training, in Lunar and Planetary Inst. Technical Report, Vol. 40, Lunar and Planetary Institute Science Conference Abstracts, 2086

Humphreys, G., Houston, M., Ng, R., et al. 2002, Chromium: a stream-processing framework for interactive rendering on clusters, in ACM Transactions on Graphics (TOG), Vol. 21 (ACM), 693–702

Iserte, S., Clemente-Castello, F., Castello, A., Mayo, R., & Quintana-Orti, E. 2016, CLOSER 2016: proceedings of the 6th International Conference on Cloud Computing and Services Science: Rome, Italy, April 23-25, 2016, ed. J. Cardoso, D. Ferguson, V. Méndez Muñoz, & M. Helfert (Setúbal, Portugal: SCITEPRESS - Science and Technology Publications, Lda)

Jakobsen, M. R., & Hornbæk, K. 2015, Is Moving Improving?: Some Effects of Locomotion in Wall-Display Interaction, in Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI '15 (New York, NY, USA: ACM), 4169–4178

Jakobsen, M. R., & Hornbæk, K. 2016, Negotiating for Space?: Collaborative Work Using a Wall Display with Mouse and Touch Input, in Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, CHI '16 (New York, NY, USA: ACM), 2050–2061

Johnson, G. P., Abram, G. D., Westing, B., Navrátil, P., & Gaither, K. 2012, DisplayCluster: An Interactive Visualization Environment for Tiled Displays, in Cluster Computing (CLUSTER), 2012 IEEE International Conference on (IEEE), 239–247

Juric, M., & Tyson, T. 2012, LSST Data Management: Entering the Era of Petascale Optical Astronomy, Proceedings of the International Astronomical Union, 10, 675

Juve, G., Deelman, E., Vahi, K., et al. 2009, Scientific workflow applications on Amazon EC2, in E-Science Workshops, 2009 5th IEEE International Conference on (IEEE), 59–66

Kang, Y. B. 2007, A scalable, high-resolution tiled display system, IMECS 2007: International Multiconference of Engineers and Computer Scientists, Vols I and II, 2009

Khalid, F., Shoaib, U., Sarfraz, M. S., et al. 2016, Desktop Virtualization: An Art to Manage and Maintain affordable PC infrastructure, International Journal of Computer Science and Information Security, 14, 187

Kido, Y., Ichikawa, K., Date, S., et al. 2016, SAGE-based Tiled Display Wall enhanced with dynamic routing functionality triggered by user interaction, Future Generation Computer Systems, 56, 303

Kim, S., Kim, M., Cho, Y., & Park, K. S. 2009, iTILE Framework for constructing interactive tiled display applications, Grapp 2009: Proceedings of the Fourth International Conference on Computer Graphics Theory and Applications, 367

Kishan, G., Aluvalu, R., & Shanker Singh, A. 2014, Intelligent Resource Allocation Technique for Desktop-as-a-Service in Cloud Environment, International Journal of Computer Applications, 96, 43

Kleinrock, L. 2005, A Vision for the Internet, ST Journal for Research, 2, 4

Konrad, J., Lacotte, B., & Dubois, E. 2000, Cancellation of image crosstalk in time-sequential displays of stereoscopic video, IEEE Transactions on Image Processing, 9, 897

Kopecki, A. 2011, Collaborative Integration of Classic Applications in Virtual Reality Environments, Journal of Systemics, Cybernetics, and Informatics, 9, 24

Koribalski, B. S. 2012, Overview on Spectral Line Source Finding and Visualisation, Publications of the Astronomical Society of Australia, 29, 359

K.U, J., & David, J. M. 2014, Issues, Challenges and Solutions : Big Data Mining, in Sixth International Conference on Networks & Communications (NeTCoM 2014) (Academy & Industry Research Collaboration Center (AIRCC)), 131–140

Kundra, V. 2011, Federal cloud computing strategy (White House,[Chief Information Officers Council])

Lam, H., Bertini, E., Isenberg, P., Plaisant, C., & Carpendale, S. 2012, Empirical studies in information visualization: Seven scenarios, IEEE transactions on visualization and computer graphics, 18, 1520

Laney, D. 2001, 3-D Data Management: Controlling Data Volume, Velocity and Variety, Application Delivery Strategies by META Group Inc., 949

Large, M. J., Large, T., & Travis, A. R. L. 2010, Parallel Optics in Waveguide Displays: A Flat Panel Autostereoscopic Display, Journal of Display Technology, 6, 431

Lau, C. D., Levesque, M. J., Date, S., Chien, S., & Haga, J. H. 2010, Visualization of virtual screening results on tiled display walls (TDW), Faseb Journal, 24

Lee, H. C., Lee, W. O., Cho, C. W., et al. 2013, Remote Gaze Tracking System on a Large Display, Sensors, 13, 13439

Lee, Y. J., Zitnick, C. L., & Cohen, M. F. 2011, ShadowDraw, ACM Transactions on Graphics, 30, 1

Leigh, J., Johnson, A., Renambot, L., et al. 2013, Scalable Resolution Display Walls, Proceedings of the IEEE, 101, 115

Li, K., Chen, H., Chen, Y., et al. 2000, Building and using a scalable display wall system, IEEE Computer Graphics and Applications, 20, 29

Li, P., Whitman, S., Mendoza, R., & Tsiao, J. 1997, ParVox-a parallel splatting volume rendering system for distributed visualization, in Parallel Rendering, 1997. PRS 97. Proceedings. IEEE Symposium on (IEEE Comput. Soc. Press), 7–14,

Lin, T., Hu, W. H., Imamiya, A., & Omata, M. 2006, Large display size enhances user experience in 3D games, Smart Graphics, Proceedings, 4073, 257

Liu, C., Chapuis, O., Beaudouin-Lafon, M., Lecolinet, E., & Mackay, W. E. 2014, Effects of display size and navigation type on a classification task (ACM Press), 4147–4156

Marrinan, T., Aurisano, J., Nishimoto, A., et al. 2014, SAGE2: A new approach for data intensive collaboration using Scalable Resolution Shared Displays, in Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom), 2014 International Conference on (IEEE), 177–186

Massimino, P., Costa, A., Becciani, U., et al. 2014, ACID Astronomical and Physics Cloud Interactive Desktop: A Prototype of VUI for CTA Science Gateway, in Astronomical Society of the Pacific Conference Series, Vol. 485, Astronomical Data Analysis Software and Systems XXIII, ed. N. Manset & P. Forshay, 293

Meade, B., & Fluke, C. 2018, Evaluating virtual hosted desktops for graphics-intensive astronomy, Astronomy and Computing, 23, 124

Meade, B., Manos, S., Sinnott, R., et al. 2013, Research Cloud Data Communities, in THETA 2013, Hobart, Tasmania, copyright 2013 THETA: The Higher Education Technology Agenda

Meade, B., Manos, S., Sinnott, R., et al. 2015, Seeing the Big Picture: A Digital Desktop for Researchers, in THETA 2015, Gold Coast, Queensland, copyright 2015 THETA: The Higher Education Technology Agenda

Meade, B., Fluke, C., Cooke, J., et al. 2017, Collaborative Workspaces to Accelerate Discovery, Publications of the Astronomical Society of Australia, 34

Meade, B. F., Fluke, C. J., Manos, S., & Sinnott, R. O. 2014, Are Tiled Display Walls Needed for Astronomy?, Publications of the Astronomical Society of Australia, 31

Mell, P., & Grance, T. 2011, The NIST definition of cloud computing, Special Publication 800-145 (U.S. Department of Commerce)

Mickaelian, A. M. 2016, Astronomical Surveys and Big Data, Open Astronomy, 25

Miller, K., & Pegah, M. 2007, Virtualization: virtually at the desktop, in Proceedings of the 35th Annual ACM SIGUCCS Fall Conference, SIGUCCS '07 (ACM Press), 255–260

Morabito, R., Kjällman, J., & Komu, M. 2015, Hypervisors vs. lightweight virtualization: a performance comparison, in Cloud Engineering (IC2E), 2015 IEEE International Conference on (IEEE), 386–393

Moreland, K. 2012, Redirecting Research in Large-Format Displays for Visualization, in Proceedings of the IEEE Symposium on Large-Scale Data Anaysis and Visualization

Morikawa, Y., Murata, K. T., Watari, S., et al. 2010, A Science Cloud: OneSpaceNet, AGU Fall Meeting Abstracts, D5

Morris, D., Voutsinas, S., Hambly, N. C., & Mann, R. G. 2017, Use of Docker for deployment and testing of astronomy software, Astronomy and Computing, 20, 105

Murphy, T., Lamb, P., Owen, C., & Marquarding, M. 2006, Data Storage, Processing, and Visualization for the Australia Telescope Compact Array, Publications of the Astronomical Society of Australia, 23, 25

Nagao, K., Ye, Y., Wang, C., Fujishiro, I., & Ma, K.-L. 2016, Enabling interactive scientific data visualization and analysis with see-through hmds and a large tiled display, in Immersive Analytics (IA), 2016 Workshop on (IEEE), 1–6

Nam, S., Jeong, B., Renambot, L., et al. 2009, Remote visualization of large scale data for ultra-high resolution display environments, in Proceedings of the 2009 Workshop on Ultrascale Visualization, UltraVis '09 (ACM Press), 42–44

Nancel, M., Pietriga, E., Chapuis, O., & Beaudouin-Lafon, M. 2015, Mid-Air Pointing on Ultra-Walls, ACM Transactions on Computer-Human Interaction, 22, 1

National Research Council (U.S.). 2001, Astronomy and astrophysics in the new millennium (Washington, D.C: National Academy Press)

Navrátil, P. A., Westing, B., Johnson, G. P., et al. 2009, A practical guide to large tiled displays, in Advances in Visual Computing (Springer), 970–981

Nieh, J., Yang, S. J., & Novik, N. 2000, A comparison of thin-client computing architectures, Tech. rep., Technical Report CUCS-022-00, Department of Computer Science, Columbia University

Nishimura, J., Sakamoto, N., & Koyamada, K. 2012, Tiled Display Visualization System with Multi-touch Control, in Proceedings in Information and Communications Technology, Vol. 4, Advanced Methods, Techniques, and Applications in Modeling and Simulation, ed. J. H. Kim, K. Lee, S. Tanaka, & S. H. Park (New York: Springer), 492–497

Olsen, B. I., Dhakal, S. B., Eldevik, O. P., Hasvold, P., & Hartvigsen, G. 2008, A Large, High Resolution Tiled Display for Medical Use: Experiences from Prototyping of a Radiology Scenario, Ehealth Beyond the Horizon - Get It There, 136, 535

Olsen, B. I., Laeng, B., Kristiansen, K.-A., & Hartvigsen, G. 2011, Spatial Tasks on a Large, High-Resolution, Tiled Display: A Male Inferiority in Performance with a Mental Rotation Task, in Engineering Psychology and Cognitive Ergonomics, ed. D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, & D. Harris, Vol. 6781 (Berlin, Heidelberg: Springer Berlin Heidelberg), 63–71

Ostermann, S., Iosup, A., Yigitbasi, N., et al. 2010, A performance analysis of EC2 cloud computing services for scientific computing, in Cloud computing (Springer), 115–131

Papadopoulos, C., Petkov, K., Kaufman, A. E., & Mueller, K. 2015, The Reality Deck–an Immersive Gigapixel Display, IEEE Computer Graphics and Applications, 35, 33

Parkhill, D. 1966, The Challenge of the computer utility (Addison-Wesley publishing)

Peck, S. M., North, C., & Bowman, D. 2009, A multiscale interaction technique for large, high-resolution displays, in 3D User Interfaces, 2009. 3DUI 2009. IEEE Symposium on (IEEE), 31–38

Pietriga, E., Huot, S., Nancel, M., & Primet, R. 2011, Rapid development of user interfaces on cluster-driven wall displays with jBricks, in Proceedings of the 3rd ACM SIGCHI symposium on Engineering interactive computing systems (ACM), 185–190

Pietriga, E., del Campo, F., Ibsen, A., et al. 2016, Exploratory visualization of astronomical data on ultra-high-resolution wall displays, in procspie, Vol. 9913, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 99130W

Pitzalis, D., Pillay, R., & Lahanier, C. 2006, A New Concept in High Resolution Internet Image Browsing., in Proceedings of the 10th International Conference on Electronic Publishing, 291–298

Ponto, K., Doerr, K., & Kuester, F. 2010, Giga-stack: A method for visualizing giga-pixel layered imagery on massively tiled displays, Future Generation Computer Systems, 26, 693

Ponto, K., Doerr, K., Wypych, T., Kooker, J., & Kuester, F. 2011, CGLXTouch: A multiuser multi-touch approach for ultra-high-resolution collaborative workspaces, Future Generation Computer Systems, 27, 649

Ponto, K., Wypych, T., Doerr, K., et al. 2009, VideoBlaster: a distributed, low-network bandwidth method for multimedia playback on tiled display systems, 2009 11th IEEE International Symposium on Multimedia (Ism 2009), 201

Prouzeau, A., Bezerianos, A., & Chapuis, O. 2017, Trade-offs Between a Vertical Shared Display and Two Desktops in a Collaborative Path-Finding Task, Proceedings of Graphics Interface 2017, Edmonton, Alberta, Canada, 16-19 May 2017, 214 - 219

Quinn, P., Axelrod, T., Bird, I., et al. 2014, Delivering SKA Science, in Proceedings of Science

Rantzau, D., Lang, U., Lang, R., et al. 1996, Collaborative and Interactive Visualization in a Distributed High Performance Software Environment, in High Performance Computing for Computer Graphics and Visualisation, ed. M. Chen, P. Townsend, & J. A. Vince (London: Springer London), 207–216

Ravi, V. T., Becchi, M., Agrawal, G., & Chakradhar, S. 2011, Supporting GPU sharing in cloud environments with a transparent runtime consolidation framework, in Proceedings of the 20th international symposium on High performance distributed computing (ACM), 217–228

Rehr, J. J., Vila, F. D., Gardner, J. P., Svec, L., & Prange, M. 2010, Scientific computing in the cloud, Computing in science & Engineering, 12, 34

Renambot, L., Jeong, B., Hur, H., Johnson, A., & Leigh, J. 2009, Enabling high resolution collaborative visualization in display rich virtual organizations, Future Generation Computer Systems, 25, 161

Renambot, L., Rao, A., Singh, R., et al. 2004, Sage: the scalable adaptive graphics environment, in Proceedings of WACE, Vol. 9 (Citeseer), 2004–09

Rimal, B. P., Choi, E., & Lumb, I. 2010, A Taxonomy, Survey, and Issues of Cloud Computing Ecosystems, in Cloud Computing, ed. N. Antonopoulos & L. Gillam (London: Springer London), 21–46

Rivera, L. R., Viveros, A. M., & Vergara, S. C. 2013, User Interface Features For Tiled Display Environments, 2013 10th International Conference on Electrical Engineering, Computing Science and Automatic Control (Cce), 296

Sacerdoti, F., Chandra, S., & Bhatia, K. 2004, Grid systems deployment & management using Rocks, in Cluster Computing, 2004 IEEE International Conference on (IEEE), 337–345

Sakuraba, A., Noda, S., Ishida, T., Ebara, Y., & Shibata, Y. 2013, Tiled Display Environment to realize GIS based Disaster Information System, 2013 16th International Conference on Network-Based Information Systems (Nbis 2013), 311

Schafer, C. 2017, The Potential of Deep Learning with Astronomical Data, in American Astronomical Society Meeting Abstracts, Vol. 230, American Astronomical Society Meeting Abstracts #230, 104.02

Scheidegger, L., Vo, H. T., Krüger, J., Silva, C. T., & Comba, J. L. D. 2012, Parallel large data visualization with display walls, in Visualization and Data Analysis 2012, 82940C

Schlegel, D. 2012, LSST is Not "Big Data", letter submitted to the NSF Astronomy Portfolio Review

Shin, S., Kang, Y., Lee, H., et al. 2010, Development of v-DMU Based on e-Science Using COVISE and SAGE, in Future Application and Middleware Technology on e-Science, ed. O.-H. Byeon, J. H. Kwon, T. Dunning, K. W. Cho, & A. Savoy-Navarro (Boston, MA: Springer US), 21–30

Sims, M. H., Dodson, K. E., & Edwards, L. J. 2010, Hyperwall Use as a Tool for Collaboration, LPI Contributions, 1538, 5614

Smarr, L., Brown, M., & de Laat, C. 2009, Special section: OptIPlanet — The OptIPuter global collaboratory, Future Generation Computer Systems, 25, 109

Smarr, L., Ford, J., Papadopoulos, P., et al. 2005, The OptIPuter, quartzite and starlight projects: A campus to global-scale testbed for optical technologies enabling LambdaGrid computing, in Optical Fiber Communication Conference (Optical Society of America), OWG7

Smarr, L. L., Chien, A. A., DeFanti, T., Leigh, J., & Papadopoulos, P. M. 2003, The OptIPuter, Communications of the ACM, 46, 58

Smith, J., & Nair, R. 2005, The architecture of virtual machines, Computer, 38, 32

Sommer, B., Barnes, D. G., Boyd, S., et al. 2017, 3D-stereoscopic immersive analytics projects at Monash University and University of Konstanz, Electronic Imaging, 2017, 179

Son, S., Hong, J., Bae, C., Jun, S. C., & Kim, J. 2010, Interactive Scientific Visualization of High-resolution Brain Imagery Over Networked Tiled Display, Future Application and Middleware Technology on E-Science, 125

Song, Z., Gong, G., Huang, Z., Han, L., & Ding, Y. 2010, A new edge blending paradigm for multi-projector tiled display wall, in Computer Application and System Modeling (ICCASM), 2010 International Conference on (IEEE), V5–349–V5–352

Sutherland, I. E. 1965, The ultimate display, Multimedia: From Wagner to virtual reality

Sutherland, I. E. 1968, A head-mounted three dimensional display, in Proceedings of the December 9-11, 1968, Fall Joint Computer Conference, Part I, AFIPS '68 (Fall, part I) (ACM Press), 757

Tada, T., Shimojo, S., Ichikawa, K., Abe, H., et al. 2011, A visualization adapter for SAGE-enabled tiled display wall, in Granular Computing (GrC), 2011 IEEE International Conference on, 613–618

Taesombut, N., Wu, X., Chien, A., et al. 2006, Collaborative data visualization for earth sciences with the OptIPuter, Future Generation Computer Systems, 22, 955

Taylor, M. B., Boch, T., & Taylor, J. 2015, SAMP, the Simple Application Messaging Protocol: Letting applications talk to each other, Astronomy and Computing, 11, 81

Tong, R., & Zhao, D. 2010, Interactive Real-Time Visualization for Large-Scale Spatial Datasets Based on OptIputer, in 2010 International Conference on Computational Intelligence and Software Engineering (IEEE), 1–4

Turilli, M., Wallom, D., Williams, C., et al. 2013, Flexible services for the support of research, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 371

Van Baar, J., Willwacher, T., Rao, S., & Raskar, R. 2003, Seamless multi-projector display on curved screens, in Proceedings of the workshop on Virtual environments 2003 (ACM), 281–286

van der Schaaf, T., Germans, D., Bal, H. E., & Koutek, M. 2007, Lessons learned from building and calibrating the ICWall, a stereo tiled display, Computer Animation and Virtual Worlds, 18, 193

Vöckler, J., Juve, G., Deelman, E., Rynge, M., & Berriman, B. 2011, Experiences using cloud computing for a scientific workflow application, in Proceedings of the 2nd international workshop on Scientific cloud computing (ACM), 15–24

Vohl, D., Fluke, C. J., Hassan, A. H., & Barnes, D. G. 2016, An interactive, comparative and quantitative 3D visualization system for large-scale spectral-cube surveys using CAVE2, in Proceedings of Astronomical Data Analysis Software and Systems XXV, ASP Conf. Series, Sydney, NSW, Australia

von Zadow, U., Büschel, W., Langner, R., & Dachselt, R. 2014, SleeD: Using a Sleeve Display to Interact with Touch-sensitive Display Walls, in Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces, ITS '14 (New York, NY, USA: ACM), 129–138

Voorsluys, W., Broberg, J., & Buyya, R. 2011, Introduction to cloud computing, Cloud computing: Principles and paradigms, 3

Ward, J. S., & Barker, A. 2013, Undefined by data: a survey of big data definitions, arXiv preprint arXiv:1309.5821

Wedeen, M., Darling, J. M., Cagampan, C., et al. 2014, Controlling the CAVE: A User Study Exploring Handheld Controllers in a Virtual Reality Environment, in CHI 2014 (Toronto, Canada: ACM Press)

Wierse, A., Lang, U., & Rühle, R. 1993, Architectures of distributed visualization systems and their enhancements, in Eurographics workshop on visualization in scientific computing, Abingdon

Wigdor, D., Jiang, H., Forlines, C., Borkin, M., & Shen, C. 2009, WeSpace: the design development and deployment of a walk-up and share multi-surface visual collaboration system, in Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (ACM), 1237–1246

Wiley, K., Connolly, A., Gardner, J., et al. 2011, Astronomy in the cloud: using mapreduce for image co-addition, Astronomy, 123, 366

Wong, C. O., Kyoung, D., & Jung, K. 2007, Adaptive context aware attentive interaction in large tiled display, Universal Access in Human-Computer Interaction: Ambient Interaction, Pt 2, Proceedings, 4555, 1016

Woods, D. D. 1984, Visual momentum: a concept to improve the cognitive coupling of person and computer, International Journal of Man-Machine Studies, 21, 229

Yamaoka, S., Manovich, L., Douglass, J., & Kuester, F. 2011, Cultural Analytics in Large-Scale Visualization Environments, Computer, 44, 39

Yost, B., Haciahmetoglu, Y., & North, C. 2007, Beyond visual acuity: the perceptual scalability of information visualizations for large displays, ACM CHI, 101

Younge, A. J. 2016, Architectural principles and experimentation of distributed high performance virtual clusters, PhD thesis, Indiana University

Zaslavsky, A., Perera, C., & Georgakopoulos, D. 2013, Sensing as a service and big data, in Proceedings of the International Conference on Advances in Cloud Computing (ACC), Bangalore, India

Zhang, Y., & Zhao, Y. 2015, Astronomy in the big data era, Data Science Journal, 14, 11

# A
# Ethics approval and Informed Consent forms

## A.1 Ethics approval and Informed Consent documentation for "Are Tiled Display Walls Needed for Astronomy?"

Rec 7/4/14



SWINBURNE UNIVERSITY OF TECHNOLOGY

**SWINBURNE RESEARCH**

# Human Research Final Report Form

## SECTION A: PROJECT DETAILS

HREC Project No: **2013/098** **Report Date: 7-4-2014**

Project Title: **Astronomy data visualisation experiments: studying the use of display technologies**

Chief Investigator: **Assoc. Prof. Christopher Fluke**

Co-Investigators:

Student Investigators: **Bernard Meade**

Currently Approved from **17-05-2013** to **28-02-2014**

If you have recently submitted to Swinburne Research a progress / final report, please check this box ☐ and give the date of submission . dd / mm / yyyy

If you are now submitting the report, please indicate if this report should be treated ☐ or Final Report ## ☒

**If there are any corrections or further details needed in the details above, please clarify below.**

NB If there are personnel changes not previously notified, please check C2 as applicable

## SECTION B: PROGRESS SUMMARY

### B1 Status of the Project (as at the date of this report)

Please check [*double-click*] one or more of the following:

- ☐ Data collection/access yet to commence [Explain below then go to B3 then Sect D]
- ☒ Data collection/access commenced/continuing [Explain below then go to B2 →]
- ☒ Data collection/access completed [Explain below then go to Sect C →]##
- ☐ Project Abandoned before data collection/access [Explain below then go to Sect D]
- ☐ Project abandoned during or after data collection/access [Explain below then go to Sect C →]##
- ☐ Other

Please provide a brief explanation as to the project status indicated, including any delays.

All experimental data has now been collected and analysis is well underway. We are preparing to submit the first set of key findings for publication in the journal Publications of the Astronomical Society of Australia within the next few weeks.

[## *Please note that completion of data collection or project abandonment effectively means this completed form equates to a final report. Unless otherwise advised (eg, due to a participant complaint, a research conduct inquiry or other significant issue arising), no further action will be needed beyond the submission of this report.*]

**B2 Participants**

How many participants have been recruited to date? 57

**B3 Given any delays or other factor** (as per B1 above), **do you need an extension of ethics clearance without any other modification to the project?**

☒ No ☐ YES. If YES, please give the new date for end of data collection/access:

dd / mm / yyyy

NB information given below re changes may mean a simple extension of ethics clearance cannot be given and a separate ethics clearance application process for modifications needs to be followed.

## SECTION C: CONDUCT OF PROJECT

**C1 Compliance with the approved protocol**

Has the project been conducted in line with the approved protocol, including standard and any special conditions of ethics clearance?

☒ Yes ☐ NO

If NO, please provide a brief explanation. *Please also note information to be given at C2.*

**C2 Project / Protocol Modifications / Additions**

Please check [*double-click*] one or more of the following if there are any:

- ☐ Changes/additions to project investigators (including students) and personnel accessing identifiable info
- ☐ Changes/additions re project personnel deriving a personal benefit from the research
- ☐ Changes/additions to the research protocol (title, aims, procedures, measures, sampling, etc)
- ☐ Changes/additions to the consent instruments/arrangements (including re personnel)
- ☐ Changes/additions to the recruitment material/methods
- ☐ Changes/additions to participant sampling or numbers
- ☐ Changes/additions to project resourcing (financing or otherwise)
- ☐ Other changes/additions

Have any of the above changes been put for ethical review?

☐ Yes ☐ NO

Whether Yes or No, please briefly explain the situation, including any separate submission(s) for ethics clearance for the modification(s) indicated. NB This form cannot be used for modification requests.**

No changes made.

** Information on applying for ethics clearance for any modification(s) can be found at: http://www.research.swinburne.edu.au/ethics/human/monitoringReportingChanges/

## C3 Project incidents

Have there been any incidents that affected the conduct of the project or which have impacted adversely on participants and/or the researchers?

☒ No ☐ YES

If YES, please provide a brief explanation including with respect to any reporting to Swinburne or other authority:

**Level of impact of incident:**
Would any of the incidents related above be considered serious or worse?
[Serious adverse events (SAEs) include, eg, harm or distress to individuals or groups, loss of significant or sensitive data, breach of confidentiality.]

☐ No ☐ YES

If YES, please provide a brief explanation including with respect to any reporting to Swinburne or other authority, where appropriate attaching a copy of the report(s):

## C4 Issues or experiences of ethical significance?

Have there been any issues or experiences which have been or remain of ethical significance, especially as regards the ethical conduct of the project and/or project outcomes, including any actual or potential conflicts of interest not identified previously or formal complaints received/processed?

☒ No ☐ YES

If yes, please briefly explain.

## C5 Project Outcomes (as at the date of this report)

Please check all of the following:

| | Yes | No | n/a |
|---|---|---|---|
| Compensatory payments made or prizes awarded and records kept | ☐ | ☐ | ☒ |
| Student thesis/theses submitted for examination | ☐ | ☒ | ☐ |
| Results have been published or presented | ☒ | ☐ | ☐ |
| Results are to be published or presented | ☒ | ☐ | ☐ |
| A lay summary of the project outcomes is given below^ | ☒ | ☐ | ☐ |
| Project outcomes have been made available to participants | ☐ | ☒ | ☐ |
| Project outcomes are to be made available to participants | ☐ | ☒ | ☐ |
| Other | ☐ | ☐ | ☒ |

If Yes or No, please provide a brief explanation as to the items checked as appropriate:

- Bernard Meade expects to submit his Masters thesis incorporating experimental results in the second half of 2014.
- Results have been presented at the following conferences by C.Fluke: Astroinformatics 2013 (Sydney); 7th international PHISCC workshop (Dwingeloo; the Netherlands in March 2014).
- Key results from the experiments will be submitted to the journal Publications of the Astronomical Society of Australia (submission excepted prior to end of April 2014).
- Should any participant ask us for information on the outcome of the experiments, we will provide a copy of the journal article (as appropriate).

^Brief lay summary of project outcomes (not more than ¼ page):

Astronomy has entered the "big data" era, with data collection now vastly exceeding the desktop computing environments of all astronomers. Indeed, it may no longer be feasible for many standard data analysis and visualisation tasks to be undertaken at the desktop. In our data visualisation experiments, we examined the role that non-standard display technologies might play. In particular, we compared Tiled Display Walls comprising a matrix of off-the-shelf LCD monitors with a Standard Desktop Display.

Our experimental participants were presented with a set of target identification tasks, including searching for words, galaxies and nebula from very high-resolution images that vastly exceed the resolution of a standard desktop display. Our experiments included astronomer and non-astronomers – both groups were equally unfamiliar with the use of Tiled Display Walls.

We found that Tiled Display Walls do provide some improvements when used to find regions of interest within an image that significantly exceeds the resolution of a Standard Desktop Display. We found that most participants preferred using the Tiled Display Wall for the tasks, and performed around 10% better overall in that environment. We also found that non-astronomers working in collaborative pairs performed as well if not better when using a Tiled Display Wall than a sole astronomer on the same task.

**C6 Study Materials/Documents**

Please check one or more of the following:

☒ Project documents/material securely stored for the minimum period

☐ Project material to be made available for future research/other researchers. If so, in what form?

Briefly explain what storage or archiving has occurred, including the location(s) and length of secure storage as well as intended secure data disposal arrangements:

Storage of materials/documents has been undertaken in compliance with our approved protocol.

All video recordings are stored on a password-projected disk drive, directly accessible only by Mr Meade. At the completion of the project, recorded footage that was not used as part of any publication, or for the Masters thesis work of Mr Meade, will be deleted. Recorded footage that was used will be copied to an DVD/external hard-drive, that will be stored in a locked filing cabinet in A/Prof Fluke's office for the time period required by journal or other research publisher, after which time it will be deleted/destroyed.

Once final analysis of survey forms is completed, the hard-copy forms will be stored in a locked filing cabinet in the office of A/Prof Fluke for a relevant time period following publication of research results, after which time they will be securely disposed of.

Are research material retention and disposal arrangements in line with what was outlined in the approved project protocol? ☒ Yes ☐ NO

If NO, please explain why.

**C7 Project Audits**

Please check one or more of the following:

☒ Project self-audit(s) have been conducted during or at conclusion of project

[For a self-audit tool, see http://www.research.swinburne.edu.au/ethics/human/monitoringReportingChanges/]

☐ Swinburne audit(s) have been conducted during or following completion of the project

☐ External audit(s) have been conducted during or following completion of the project

Please provide a brief explanation as to any audits conducted:

C.Fluke and B.Meade have followed procedures in the self-audit.

## SECTION D: DECLARATION BY CHIEF INVESTIGATOR/SUPERVISOR

**DECLARATION BY CHIEF INVESTIGATOR(s)/STUDENT SUPERVISOR(s)**

I declare that the above report accurately reflects the outcome or progress of the project to date

I acknowledge that an internal Swinburne or external audit may be conducted on the conduct of the project and as regards secure data retention/disposal.

Signature & Date: ...Christopher Fluke............................ 7/4/14.....

Name of Signatory & Position: ...Associate Professor Christopher Fluke..........................

**Student Investigator(s)** (where possible)

I agree with the above declaration signed by the Chief Investigator/Supervisor

Signature & Date: ...................................................... 7/04/14 ......................................

Name of Student:......Bernard Meade..........................

***For official use only:***

| | | |
|---|---|---|
| Progress Reports Received | ☐ entered | ☐Not applicable |
| Final Report Received | ☐ entered | ☐Not applicable |

Research Ethics Office Action Taken/Notes:

To: A/Prof C Fluke Mr B Meade FICT

Dear Christopher and Bernard,

**SUHREC 2013/098 Astronomy data visualisation experiments: studying the use of display technologies**
A/Prof C Fluke Mr Bernard Meade FICT
Approved duration: 17/05/2013 To 17/10/2013 [Adjusted]

I refer to the ethical review of the above project protocol undertaken on behalf of Swinburne's Human Research Ethics Committee (SUHREC) by SUHREC Subcommittee (SHESC2) at a meeting held on 10th May 2013. Your response to the review as e-mailed on 16th May 2013 was reviewed.

I am pleased to advise that, as submitted to date, the project may proceed in line with standard on-going ethics clearance conditions here outlined.

- All human research activity undertaken under Swinburne auspices must conform to Swinburne and external regulatory standards, including the current *National Statement on Ethical Conduct in Human Research* and with respect to secure data use, retention and disposal.

- The named Swinburne Chief Investigator/Supervisor remains responsible for any personnel appointed to or associated with the project being made aware of ethics clearance conditions, including research and consent procedures or instruments approved. Any change in chief investigator/supervisor requires timely notification and SUHREC endorsement.

- The above project has been approved as submitted for ethical review by or on behalf of SUHREC. Amendments to approved procedures or instruments ordinarily require prior ethical appraisal/ clearance. SUHREC must be notified immediately or as soon as possible thereafter of (a) any serious or unexpected adverse effects on participants and any redress measures; (b) proposed changes in protocols; and (c) unforeseen events which might affect continued ethical acceptability of the project.

- At a minimum, an annual report on the progress of the project is required as well as at the conclusion (or abandonment) of the project.

- A duly authorised external or internal audit of the project may be undertaken at any time.

Please contact the Research Ethics Office if you have any queries about on-going ethics clearance. The SUHREC project number should be quoted in communication. Chief Investigators/Supervisors and Student Researchers should retain a copy of this email as part of project record-keeping.

Best wishes for project.

Yours sincerely,

Ann Gaeth

Dr Ann Gaeth
Administration Officer (Research Ethics)
Swinburne Research (H68)
Swinburne University of Technology
P O Box 218
HAWTHORN VIC 3122
☎ +61 3 9214 8356
🖷 +61 3 9214 5267

Figure A.1 Ethics approval confirmation for human study as published in Meade et al. (2014), Chapter 4 Section 3.2.

To: A/Prof C Fluke Mr B Meade FICT

Dear Christopher and Bernard,

**SUHREC 2013/098 Astronomy data visualisation experiments: studying the use of display technologies**
A/Prof C Fluke Mr Bernard Meade FICT
Approved duration: 17/05/2013 To 17/10/2013 [Extended to 28/02/2014]

I refer to your request to modify the approved protocol for the above project as per your email of 9 October 2013. The request, to extend the project duration and to increase the participant numbers to 100, was put to a SHESC2 delegate for consideration.

I am pleased to advise that, as modified to date, the project may continue in line with on-going ethics clearance conditions previously communicated and reprinted below.

Please contact the Research Ethics Office if you have any queries about on-going ethics clearance, citing the SUHREC project number. Please retain a copy of this email as part of project record-keeping.

As before, best wishes for the continuing project.

Yours sincerely

Ann Gaeth
Ph +61 3 9214 8356

Figure A.2 Ethics confirmation of extension for human study as published in Meade et al. (2014), Chapter 4 Section 3.2.

# Consent Information Statement

**Swinburne University of Technology**

**Project Title:** Astronomy data visualisation experiments

**Principle Investigators:** A/Prof Christopher Fluke (Swinburne University of Technology) and Mr Bernard Meade (Student investigator, Swinburne University of Technology)

## About the Project

Many scientific disciplines are now entering the petabyte-data era, where the quantities of data stored vastly exceed the capacity of standard desktop-based research workflows. This growth in the total volume of scientific data is due, in part, to the increases in image resolution that modern instruments and detectors are able to record. However, the use of display technologies capable of showing high-resolution images (e.g. typically greater than 10 million pixels per image) is not widespread.

In this project, we are investigating the impact of large-format, tiled displays on common visualisation and anomaly identification tasks from the domain of astronomy. By combining many smaller displays, large-format tiled displays provide many more screen pixels than a standard desktop display or high-definition projected image.

While there appear to be benefits from being able to view high-resolution images using such displays, there up-take in astronomy over the last decade has been low. In part, this is due to the higher cost of installing and running such a system.

We invite you to participate in a series of Astronomy data visualisation experiments, with which we will critically assess the suitability of large-format, tiled displays compared with lower-cost, standard desktop displays. Our goal is to obtain an understanding of the role that different display technologies play in supporting visualisation and anomaly identification problems in astronomy, including quantitative measurements of the time taken to complete tasks and qualitative examination of how the displays are used in practice. The outcomes will inform future research on appropriate software and hardware solutions to maximise the potential of each display type for big-data tasks.

## Project and researcher interests

This project is being undertaken wholly to satisfy Mr Bernard Meade's academic qualification for the degree of Master of Science. Access to the display technologies is provided by the Information Technology Services (ITS) Research group at the University of Melbourne.

## What the Project will involve

As a participant in the Astronomy data visualisation experiments, you will use up to three different display technologies (standard desktop monitor, high resolution data projection, OptIPortal large-format tiled display) to complete a set of image identification tasks. In each experiment, we will ask you to identify specific patterns within images with resolutions up to 16,000 x 9,000 pixels. The time taken to complete each task will be recorded for later analysis.

Centre for Astrophysics and Supercomputing
Swinburne University of Technology

Mail H30, PO Box 218
Hawthorn Campus
John Street Hawthorn
Victoria 3122 Australia

Telephone +61 3 9214 5569
Facsimile +61 3 9214 8797
http://astronomy.swin.edu.au/

Each experiment will be filmed, providing insight into how individuals or small groups cooperate to use each type of display to complete the assigned tasks.

At the completion of the experiments, we will ask you to complete a short survey form to record your impressions on the suitability and ease of use of the display technologies for each of the experiments.

**Participant rights and interests**

**(i) Risks**

When filming participants completing the Astronomy data visualisation experiments, every effort will be made to avoid identification of individuals. However, we are aware of the potential for an individual to feel embarrassment while being filmed. Your informed consent (see below) regarding the use of filmed material can be withdrawn at any stage, and any relevant video recordings of the session will be deleted immediately. Analysis of the experimental outcomes (timing and survey results) can still occur without the filmed content.

**(ii) Benefits**

Participation in this project will contribute to knowledge on the use and suitability of different display technologies to the image-based problems that are arising in the petabyte-scale data era. Participants will also have an opportunity to experience, first hand, astronomical imagery on a large-format, tiled display.

**(iii) Free Consent/Withdrawal from Participation**

Participation in the project is voluntary. Participants have the right to withdraw participation, data or material contributed at any stage without question or explanation. Your consent to participate in the project will be indicated by completion of the signed Informed Consent document.

**(iv) Privacy & Confidentiality**

Signed consent forms will be stored separately to any data collected and will be accessible only to the Principle Investigators. Anonymous, hardcopy survey forms will be stored separately from other physical materials, accessible only to the Principle Investigators. All electronically recorded material will be stored on password protective hard-drives, accessible only to the Principle Investigators

**(v) Research output**

The data collected during this project will be analysed and interpreted, and prepared for publication and presentation in Mr Bernard Meade's Masters thesis. Publications will be prepared for relevant astronomy and/or computing journals, and presentations of this work will be made at relevant astronomy and/or computing conferences.

Participants wishing to receive copies of any publications are invited to provide their contact details to the Principle Investigators, and this information will be stored separately to all data collected during the project.

SWINBURNE UNIVERSITY OF TECHNOLOGY

Centre for Astrophysics and Supercomputing
Swinburne University of Technology

Mail H30, PO Box 218
Hawthorn Campus
John Street Hawthorn
Victoria 3122 Australia

Telephone +61 3 9214 5569
Facsimile +61 3 9214 8797
http://astronomy.swin.edu.au/

SWINBURNE UNIVERSITY OF TECHNOLOGY

**Further information about the project**

**If you would like** further information about the project, please do not hesitate to contact:

A/Prof Christopher Fluke
Centre for Astrophysics & Supercomputing
Swinburne University of Technology
Mail H30
PO Box 218
Hawthorn VIC 3122

Telephone: (03) 9214 5828
E-mail: cfluke@swin.edu.au

**Concerns/complaints about the project**

This project has been approved by or on behalf of Swinburne's Human Research Ethics Committee (SUHREC) in line with the *National Statement on Ethical Conduct in Human Research*. If you have any concerns or complaints about the conduct of this project, you can contact:

Research Ethics Officer
Swinburne Research
Swinburne University of Technology
Mail H68
PO Box 218
HAWTHORN VIC 3122

Telephone: (03) 9214 5218
E-mail: resethics@swin.edu.au

Centre for Astrophysics and Supercomputing
Swinburne University of Technology

Mail H30, PO Box 218
Hawthorn Campus
John Street Hawthorn
Victoria 3122 Australia

Telephone +61 3 9214 5569
Facsimile +61 3 9214 8797
http://astronomy.swin.edu.au/

# Informed Consent

SWINBURNE UNIVERSITY OF TECHNOLOGY

1. I consent to participate in the project named above. I have been provided a copy of the project Consent Information Statement to which this consent form relates and any questions I have asked have been answered to my satisfaction.

2. ***In relation to this project, please circle your response to the following:***
   - I agree to allow my participation in the Astronomy data visualisation experiments to be filmed or otherwise recorded by electronic device **Yes** **No**
   - I agree to complete questionnaires asking me about experiences relating to the completion of the Astronomy data visualisation experiments **Yes** **No**

3. I acknowledge that:

   (a) my participation is voluntary and that I am free to withdraw from the project at any time without explanation;

   (b) the Swinburne project is for the purpose of research and not for profit;

   (c) any identifiable information about me which is gathered in the course of and as the result of my participating in this project will be (i) collected and retained for the purpose of this project and (ii) accessed and analysed by the researcher(s) for the purpose of conducting this project;

   (d) my anonymity is preserved and I will not be identified in publications or otherwise without my express written consent.

By signing this document I agree to participate in this project.

**Name of Participant:**

……………………………………………………………………………………………………

**Signature & Date:** ………………………………………………………………

**Centre for Astrophysics and Supercomputing**
**Swinburne University of Technology**

Mail H30, PO Box 218
Hawthorn Campus
John Street Hawthorn
Victoria 3122 Australia

Telephone +61 3 9214 5569
Facsimile +61 3 9214 8797
http://astronomy.swin.edu.au/

## A.2 Ethics approval and Informed Consent documentation for "Evaluating Virtual Hosted Desktops for Graphics-intensive Astronomy"

SWINBURNE UNIVERSITY OF TECHNOLOGY

## Consent Information Statement

**Swinburne University of Technology**

**Project Title:** The Astronomer's Virtual Hosted Desktop

**Principle Investigators:** A/Prof Christopher Fluke (Swinburne University of Technology) and Mr Bernard Meade (Student investigator, Swinburne University of Technology)

### About the Project

Many scientific disciplines are now entering the petabyte-data era, where the quantities of data stored vastly exceed the capacity of standard desktop-based research workflows. It is no longer viable to transfer many of these datasets to a local computer for processing, and so the computation needs to be performed by computational resources co-located with the data in a data centre. However, there remain a number of software tools that require a level of interactivity and/or graphical display that requires a computing desktop window interface to operate.

Virtual Hosted Desktops function in the same way as a standard desktop computer, with the notable exception that the computer driving the windows environment is located remotely from the user, typically in a data centre. The user is able to interact with the virtual hosted desktop using their local desktop as a display device. In this way, the hosted desktop may be far more powerful than the local computer being used to display the desktop.

Amazon Web Services (AWS) is a commercial, public cloud service that provides compute and storage on demand. With appropriate configuration of a virtual machine, a virtual hosted desktop can be effectively rented from AWS by the hour.

The NeCTAR Research Cloud is a Federally funded cloud service designed to support the computational requirements of Australian researchers. It consists of eight geographically distributed nodes, managed by a central core service operation hosted at the University of Melbourne. Computational resources, typically in the form of virtual machines, are provisioned ad hoc and on demand to the research community based on research merit.

This project aims to compare the use of virtual hosted desktops available through the NeCTAR Research Cloud and AWS with a typical local desktop configuration. We are investigating the suitability, usability and cost-effectiveness of these approaches in order to understand the strengths and weaknesses of these alternatives for graphics intensive, visual-based analysis tasks.

### Project and researcher interests

This project is being undertaken wholly to satisfy Mr Bernard Meade's academic qualification for the degree of Doctor of Philosophy. Access to the NeCTAR Research Cloud has been determined by a merit allocation application. Access to the AWS cloud services was awarded through the AWS Cloud Credits for Research scheme.

### What the Project will involve

As a participant in the Astronomer's Virtual Hosted Desktop experiments, you will be asked to use three astronomy-themed applications with various datasets. In each

Centre for Astrophysics and Supercomputing
Swinburne University of Technology

Mail H29, PO Box 218
Hawthorn Campus
John Street Hawthorn
Victoria 3122 Australia

Telephone +61 3 9214 5569
Facsimile +61 3 9214 8797
http://astronomy.swin.edu.au/

experiment, we will ask you to complete tasks using these applications on both a local laptop and a cloud virtual hosted desktop. The time taken to complete each task, along with hardware performance characteristics, will be recorded for later analysis.

The experimental process will commence with the completion of this form, followed by a short interview and introduction to the purpose of the research and how this experiment will be used to compare the technological environments. You will then participate in the main activity of the experiment, which will take around 30 minutes.

Prior to attempting each task, you will be shown how to use the software and taken through each step to complete the task. You will then perform the same task on different sized datasets, with the time to complete each task being recorded.

At the completion of the experiments, we will ask you to complete a short interview to record your impressions on the suitability and ease of use of the local and cloud based desktops for each of the experiments.

**Participant rights and interests**

**(i) Risks**

While there are no perceived risks in the operation of this experiment, your informed consent (see below) can be withdrawn at any stage, and any relevant data recorded during the session will be deleted immediately.

**(ii) Benefits**

Participation in this project will contribute to knowledge on the use and suitability of a virtual hosted desktop solution in the petabyte-scale data era. Participants will also have an opportunity to experience, first hand cloud based desktop computing with GPU acceleration via AWS and the NeCTAR Resarch Cloud.

**(iii) Free Consent/Withdrawal from Participation**

Participation in the project is voluntary. Participants have the right to withdraw participation, data or material contributed at any stage without question or explanation. Your consent to participate in the project will be indicated by completion of the signed Informed Consent document.

**(iv) Privacy & Confidentiality**

Signed consent forms will be stored separately to any data collected and will be accessible only to the Principle Investigators. Anonymous, hardcopy interview forms will be stored separately from the signed consent forms, accessible only to the Principle Investigators. All electronically recorded material, such as computer logs, will be printed and the digital files purged, with the printed copies accessible only to the Principle Investigators.

**(v) Research output**

The data collected during this project will be analysed, interpreted, and prepared for publication and presentation in Mr Bernard Meade's PhD thesis. Publications will be prepared for relevant astronomy and/or computing journals, and presentations of this work will be made at relevant astronomy and/or computing conferences.

Participants wishing to receive copies of any publications are invited to provide their contact details to the Principle Investigators, and this information will be stored separately to all data collected during the project.

SWINBURNE UNIVERSITY OF TECHNOLOGY

Centre for Astrophysics and Supercomputing
Swinburne University of Technology

Mail H29, PO Box 218
Hawthorn Campus
John Street Hawthorn
Victoria 3122 Australia

Telephone +61 3 9214 5569
Facsimile +61 3 9214 8797
http://astronomy.swin.edu.au/

**Further information about the project**

**If you would like** further information about the project, please do not hesitate to contact:

A/Prof Christopher Fluke
Centre for Astrophysics & Supercomputing
Swinburne University of Technology
Mail H29
PO Box 218
Hawthorn VIC 3122

Telephone: (03) 9214 5828
E-mail: cfluke@swin.edu.au

**Concerns/complaints about the project**

This project has been approved by or on behalf of Swinburne's Human Research Ethics Committee (SUHREC) in line with the *National Statement on Ethical Conduct in Human Research*. If you have any concerns or complaints about the conduct of this project, you can contact:

Research Ethics Officer
Swinburne Research
Swinburne University of Technology
Mail H68
PO Box 218
HAWTHORN VIC 3122

Telephone: (03) 9214 5218
E-mail: resethics@swin.edu.au

# Informed Consent

1. I consent to participate in the project named above. I have been provided a copy of the project Consent Information Statement to which this consent form relates and any questions I have asked have been answered to my satisfaction.

2. ***In relation to this project, please circle your response to the following:***
   - I agree to allow my participation in the experiment to be recorded **Yes** **No**
   - I agree to complete an interview about my experiences relating to the completion of the Astronomer's Virtual Hosted Desktop experiments **Yes** **No**

3. I acknowledge that:

   (a) my participation is voluntary and that I am free to withdraw from the project at any time without explanation;

   (b) the Swinburne project is for the purpose of research and not for profit;

   (c) any identifiable information about me which is gathered in the course of and as the result of my participating in this project will be (i) collected and retained for the purpose of this project and (ii) accessed and analysed by the researcher(s) for the purpose of conducting this project;

   (d) my anonymity is preserved and I will not be identified in publications or otherwise without my express written consent.

By signing this document I agree to participate in this project.

**Name of Participant:**

……………………………………………………………………………………………………

**Signature & Date:** ……………………………………………………………

**Centre for Astrophysics and Supercomputing**
**Swinburne University of Technology**

Mail H29, PO Box 218
Hawthorn Campus
John Street Hawthorn
Victoria 3122 Australia

Telephone +61 3 9214 5569
Facsimile +61 3 9214 8797
http://astronomy.swin.edu.au/

To: A/Prof Chris Fluke, CAS

Dear Chris,

**SHR Project 2017/108 - The Astronomer's Virtual Hosted Desktop**
A/Prof Chris Fluke, Bernard Meade (Student) – CAS
Approved duration: 03-06-2017 to 02-12-2017 [adjusted]

I refer to the ethical review of the above project by a Subcommittee (SHESC3) of Swinburne's Human Research Ethics Committee (SUHREC). Your responses to the review as emailed on 02 June 2017 were put to the Subcommittee delegate for consideration.

I am pleased to advise that, as submitted to date, ethics clearance has been given for the above project to proceed in line with standard on-going ethics clearance conditions outlined below.

- The approved duration is **03 June 2017** to **02 December 2017** unless an extension is subsequently approved.
- All human research activity undertaken under Swinburne auspices must conform to Swinburne and external regulatory standards, including the *National Statement on Ethical Conduct in Human Research* and with respect to secure data use, retention and disposal.
- The named Swinburne Chief Investigator/Supervisor remains responsible for any personnel appointed to or associated with the project being made aware of ethics clearance conditions, including research and consent procedures or instruments approved. Any change in chief investigator/supervisor, and addition or removal of other personnel/students from the project, requires timely notification and SUHREC endorsement.
- The above project has been approved as submitted for ethical review by or on behalf of SUHREC. Amendments to approved procedures or instruments ordinarily require prior ethical appraisal/clearance. SUHREC must be notified immediately or as soon as possible thereafter of (a) any serious or unexpected adverse effects on participants and any redress measures; (b) proposed changes in protocols; and (c) unforeseen events which might affect continued ethical acceptability of the project.
- At a minimum, an annual report on the progress of the project is required as well as at the conclusion (or abandonment) of the project. Information on project monitoring and variations/additions, self-audits and progress reports can be found on the Research Intranet pages.
- A duly authorised external or internal audit of the project may be undertaken at any time.

Please contact the Research Ethics Office if you have any queries about on-going ethics clearance, citing the Swinburne project number. A copy of this email should be retained as part of project record-keeping.

Best wishes for the project.

Yours sincerely,
Astrid Nordmann
(for Sally Fried, SHESC3 Secretary)

**Dr Astrid Nordmann | Research Ethics Coordinator**
Swinburne Research| Swinburne University of Technology
Ph +61 3 9214 3845| anordmann@swin.edu.au
Level 1, Swinburne Place South
24 Wakefield St, Hawthorn VIC 3122, Australia
www.swinburne.edu.au

Figure A.3 Ethics approval confirmation for human study as published in Meade & Fluke (2018), Chapter 6 Section 6.5.

On 1/11/2017, 16:14, "resethics@swin.edu.au" <resethics@swin.edu.au> wrote:

Dear Christopher,

Re: Final Report for the project 2017/108

'The Astronomer's Virtual Hosted Desktop' (Report Date: 01-11-2017)

The Final report for the above project has been processed and satisfies the reporting requirements set under the terms of ethics clearance.

Thank you for your attention to this matter.

Regards
Research Ethics Team

Swinburne Research (H68)
Swinburne University of Technology
PO Box 218
HAWTHORN VIC 3122
Tel: 03 9214 3845
Fax: 03 9214 5267
Email: resethics@swin.edu.au

Figure A.4 Ethics confirmation of final report for human study as published in Meade & Fluke (2018), Chapter 6 Section 6.5.

# B
# Co-authorship indication

## B.1 Co-authorship indication for "Are Tiled Display Walls Needed for Astronomy?"

SWIN BUR NE — SWINBURNE UNIVERSITY OF TECHNOLOGY

**Swinburne Research**

# Authorship Indication Form

## For PhD (including associated papers) candidates

### NOTE

This Authorship Indication form is a statement detailing the percentage of the contribution of each author in each associated 'paper'. This form must be signed by each co-author and the Principal Coordinating Supervisor. This form must be added to the publication of your final thesis as an appendix. Please fill out a separate form for each associated paper to be included in your thesis.

### DECLARATION

We hereby declare our contribution to the publication of the 'paper' entitled:

Are tiled display walls needed for astronomy?

**First Author**

Name: Bernard F. Meade Signature:

Percentage of contribution: 85 % Date: 09/03/2018

Brief description of contribution to the 'paper' and your central responsibilities/role on project:

Devised and evaluated the TDW experiment, wrote the paper

**Second Author**

Name: Christopher J. Fluke Signature: Christopher Fluke

Percentage of contribution: 10 % Date: 12/03/2018

Brief description of your contribution to the 'paper':

Research program oversight and supervision, assistance with technical editing

**Third Author**

Name: Steven Manos Signature:

Percentage of contribution: 3 % Date: 22/3/2018

Brief description of your contribution to the 'paper':

Research supervision, assistance with technical editing

**Fourth Author**

Name: Richard O. Sinnott Signature: Richard

Percentage of contribution: 2 % Date: 12/03/2018

Brief description of your contribution to the 'paper':

Research supervision, assistance with technical editing

Principal Coordinating Supervisor: Name: Christopher Fluke Signature: Christopher Fluke

Date: 12/03/2018

In the case of more than four authors please attach another sheet with the names, signatures and contribution of the authors.

## B.2 Co-authorship indication for "Collaborative Workspaces to Accelerate Discovery"

**Swinburne Research**

# Authorship Indication Form

## For PhD (including associated papers) candidates

## NOTE

This Authorship Indication form is a statement detailing the percentage of the contribution of each author in each associated 'paper'. This form must be signed by each co-author and the Principal Coordinating Supervisor. This form must be added to the publication of your final thesis as an appendix. Please fill out a separate form for each associated paper to be included in your thesis.

## DECLARATION

We hereby declare our contribution to the publication of the 'paper' entitled:

Collaborative Workspaces to Accelerate Discovery

**First Author**

Name: Bernard Meade    Signature:

Percentage of contribution: 85 %    Date: 09 / 03 / 2018

Brief description of contribution to the 'paper' and your central responsibilities/role on project:

Devised and evaluated the experiment, wrote the paper

**Second Author**

Name: Christopher Fluke    Signature: Christopher Fluke

Percentage of contribution: 6 %    Date: 12/03 /2018

Brief description of your contribution to the 'paper':

Research program oversight and supervision, assistance with technical editing

**Third Author**

Name: Jeff Cooke    Signature:

Percentage of contribution: 4 %    Date: 09/ 03 /2018

Brief description of your contribution to the 'paper':

Principle Investigator for *Deeper, Wider, Faster,* assistance with technical editing

**Fourth Author**

Name: Igor Andreoni    Signature: Igor Andreoni

Percentage of contribution: 4 %    Date: 09/ 03 /2018

Brief description of your contribution to the 'paper':

Assistance with scientific description of *Deeper, Wider, Faster,* assistance with technical editing

Principal Coordinating Supervisor: Name: Christopher Fluke    Signature: Christopher Fluke

Date: 12/ 03 / 2018

In the case of more than four authors please attach another sheet with the names, signatures and contribution of the authors.

## B.3 Co-authorship indication for “Research Cloud Data Communities”

SWINBURNE UNIVERSITY OF TECHNOLOGY

**Swinburne Research**

# Authorship Indication Form
## For PhD (including associated papers) candidates

### NOTE

This Authorship Indication form is a statement detailing the percentage of the contribution of each author in each associated 'paper'. This form must be signed by each co-author and the Principal Coordinating Supervisor. This form must be added to the publication of your final thesis as an appendix. Please fill out a separate form for each associated paper to be included in your thesis.

### DECLARATION

We hereby declare our contribution to the publication of the 'paper' entitled:

Research Cloud Data Communities

**First Author**

Name: Bernard F. Meade Signature:

Percentage of contribution: 75 % Date: 09/03/2018

Brief description of contribution to the 'paper' and your central responsibilities/role on project:

Devised and evaluated the experiment, wrote the paper

**Second Author**

Name: Christopher J. Fluke Signature: Christopher Fluke

Percentage of contribution: 5 % Date: 12/03/2018

Brief description of your contribution to the 'paper':

Research program oversight and supervision, assistance with technical editing

**Third Author**

Name: Steven Manos Signature:

Percentage of contribution: 3 % Date: 22/3/2018

Brief description of your contribution to the 'paper':

Research supervision, assistance with technical editing

**Fourth Author**

Name: Richard O. Sinnott Signature: Richard Sinnott

Percentage of contribution: 3 % Date: 12/03/2018

Brief description of your contribution to the 'paper':

Research supervision, assistance with technical editing

Principal Coordinating Supervisor: Name: Christopher Fluke Signature: Christopher Fluke

Date: 12/03/2018

In the case of more than four authors please attach another sheet with the names, signatures and contribution of the authors.

**Additional Author**

**Name:** ANDY TSENG **Signature:** Andy Tseng

**Percentage of contribution:** 7% **Date:** 12/ 03/ 2018

**Brief description of your contribution to the 'paper':** wrote the data section of the paper

**Additional Author**

**Name:** DIRK VAN DER KNIJFF **Signature:**

**Percentage of contribution:** 7% **Date:** 13 / 3 / 2018

**Brief description of your contribution to the 'paper':** wrote the HPC section of the paper

## B.4 Co-authorship indication for "Seeing the Big Picture: A Digital Desktop for Researchers"

SWINBURNE UNIVERSITY OF TECHNOLOGY

**Swinburne Research**

# Authorship Indication Form
## For PhD (including associated papers) candidates

### NOTE

This Authorship Indication form is a statement detailing the percentage of the contribution of each author in each associated 'paper'. This form must be signed by each co-author and the Principal Coordinating Supervisor. This form must be added to the publication of your final thesis as an appendix. Please fill out a separate form for each associated paper to be included in your thesis.

### DECLARATION

We hereby declare our contribution to the publication of the 'paper' entitled:

Seeing the Big Picture: A Digital Desktop for Researchers

**First Author**

Name: Bernard F. Meade Signature:

Percentage of contribution: 80 % Date: 09/03/2018

Brief description of contribution to the 'paper' and your central responsibilities/role on project:

Devised and evaluated the experiment, wrote the paper

**Second Author**

Name: Christopher J. Fluke Signature: Christopher Fluke

Percentage of contribution: 6 % Date: 12/03/2018

Brief description of your contribution to the 'paper':

Research program oversight and supervision, assistance with technical editing

**Third Author**

Name: Steven Manos Signature:

Percentage of contribution: 3 % Date: 22/3/2018

Brief description of your contribution to the 'paper':

Research supervision, assistance with technical editing

**Fourth Author**

Name: Richard O. Sinnott Signature: Richard Sinnott

Percentage of contribution: 3 % Date: 12/03/2018

Brief description of your contribution to the 'paper':

Research supervision, assistance with technical editing

Principal Coordinating Supervisor: Name: Christopher Fluke Signature: Christopher Fluke

Date: 12/03/2018

In the case of more than four authors please attach another sheet with the names, signatures and contribution of the authors.

**Additional Author**

**Name:** NEIL KILLEEN **Signature:** Neil Killeen

**Percentage of contribution:** 2% **Date:** 12/3/2018

**Brief description of your contribution to the 'paper':** assistance with technical editing

**Additional Author**

**Name:** Michael Wang **Signature:**

**Percentage of contribution:** 2% **Date:** **14 / 3/ 2018**

**Brief description of your contribution to the 'paper':** assistance with technical editing

**Additional Author**

**Name:** PAUL MIGNONE **Signature:**

**Percentage of contribution:** 2% **Date:** 14/03/18

**Brief description of your contribution to the 'paper':** assistance with technical editing

## B.5 Co-authorship indication for "Evaluating Virtual Hosted Desktops for Graphics-intensive Astronomy"

SWINBURNE UNIVERSITY OF TECHNOLOGY

Swinburne Research

# Authorship Indication Form

## For PhD (including associated papers) candidates

### NOTE

This Authorship Indication form is a statement detailing the percentage of the contribution of each author in each associated 'paper'. This form must be signed by each co-author and the Principal Coordinating Supervisor. This form must be added to the publication of your final thesis as an appendix. Please fill out a separate form for each associated paper to be included in your thesis.

### DECLARATION

We hereby declare our contribution to the publication of the 'paper' entitled:

Evaluating virtual hosted desktops for graphics-intensive astronomy

**First Author**

Name: Bernard Meade Signature:

Percentage of contribution: 90 % Date: 09 / 03 / 2018

Brief description of contribution to the 'paper' and your central responsibilities/role on project:

Devised and evaluated the experiment, wrote the paper

**Second Author**

Name: Christopher Fluke Signature: Christopher Fluke

Percentage of contribution: 10 % Date: 26/03 /2018

Brief description of your contribution to the 'paper':

Research program oversight and supervision, assistance with technical editing

**Third Author**

Name: ____ Signature: ____

Percentage of contribution: ____% Date: __ / __ / ____

Brief description of your contribution to the 'paper':

**Fourth Author**

Name: ____ Signature: ____

Percentage of contribution: ____% Date: __ / __ / ____

Brief description of your contribution to the 'paper':

Principal Coordinating Supervisor: Name: Christopher Fluke Signature: Christopher Fluke

Date: 26/03 /2018

In the case of more than four authors please attach another sheet with the names, signatures and contribution of the authors.