FUTURE SCIENCE

COMPUTER SCIENCE

Meeting the scale challenge

Australian Academy of Science

Australian Academy of Science 2013

GPO Box 783, Canberra, ACT 2601

This publication is also available online at: www.science.org.au

ISBN: 978-0-85847-342-3

Cover image: iStockphoto
Design and layout by WordUp! Websites and Graphic Design
Printed and bound by New Millennium Print

▶ CONTENTS

FOREWORD

Professor Suzanne Cory AC PresAA FRS
President
Australian Academy of Science

The Australian Academy of Science instigated a series of foresight studies in 2011 as part of its mission to promote scientific understanding and advance Australian science.

These studies offer insights into possible development pathways for different scientific fields. In doing so they provide guidance on likely research challenges that need to be overcome, and areas where productive opportunities for new scientific endeavour may exist. The future of computer science was chosen for this first report not only because it is an enabling science but also because its growing capacity to facilitated social interaction worldwide.

This document, Future Science – Computer Science, scopes future capacities and applications of computer science and, to the extent possible, identifies scientific challenges that are likely to drive the development of computer science over the next 10 to 20 years. High performance computing and the growing interface between humans and computers are particularly considered.

The Academy is very appreciative of the contributions made to this report by the members of the Expert Group. The Academy also thanks the Defence Science and Technology Organisation for providing financial support to prepare this document.

2

INTRODUCTION

In today's world, computing power, digital information and networking capacity are growing exponentially[1]. That growth has the potential to significantly improve the way individuals work, play and care for one another.

Computing power can help society to understand and overcome diseases, and to improve quality of life for the most vulnerable citizens. It can help society better manage the environment, by allowing problems to be predicted sooner and rectified more easily.

Computing power can also help society to accelerate the development of better energy sources and to more effectively use existing sources of energy. Computing makes it easier to extract the maximum value out of existing infrastructure, including improving our transportation.

Harnessing the power of computing can improve national productivity, significantly increase national wealth and create entirely new enterprises and categories of work. Planning for disasters and recovering from them are made faster and more effective thanks to computing power. It can also help societies improve the security of their physical and virtual infrastructure.

At a more personal level, computing power allows individuals to properly manage digital identities, improve trust in online interactions and develop and maintain personal relationships.

All of these benefits are possible because computing has become pervasive. There are now information and communications technologies (ICT) in almost everything, from mobile phones to cars, and from supercomputers to buildings. Economic sectors and scientific disciplines have been disrupted, revolutionised and reinvented as ICT has swept through them. This has transformed everything from the way people shop to the way they do science. ICT is a 'universal acid', eating through everything society does and leaving in its wake a transformed economic, social and scientific landscape.

The numbers that chart this change are staggering. There are hundreds of billions of computing devices, including those embedded in toasters, washing machines, cars, TVs and airplanes. There are more than a billion 'normal', non-embedded computers, such as PCs, laptops and iPads, and billions of mobile phones. By 2014, there are expected to be more than 2 billion such devices[2]. More than 2 billion people use the internet[3]. Wireless networks at gigabit/s speeds and wired networks at terabit/s speeds are now possible[4]. Annual global internet traffic is predicted to surpass one zettabyte (10^{21} bytes) by 2016[5]. The amount of data created and replicated in 2011 exceeded 1 zettabyte and is expected to reach 8 zettabytes by 2015[6].

This surge of information and processing capability is unprecedented and has created a significant

scale challenge for computing. Current computing methodologies are not designed to operate at such scales. These include programming technologies and software engineering techniques; data mining, analysis and other algorithmic methods; networking technologies and approaches to security and privacy; and human-computer interaction paradigms. Research in computing (and other areas) over the next 10 to 15 years will largely be driven by the problem of overcoming the myriad issues generated by this scale challenge.

Although Australia has a 'primary industries' economy, it is nonetheless vital that the nation operates at the forefront of this wave of opportunity. If it does not, significant parts of the economy -- for example health, services and education — will simply become uncompetitive. Similarly 'big data' is now at the heart of resource discovery and exploitation. And consumers will soon expect that the agricultural produce they consume will be tracked continuously, from paddock to plate, which requires a large-scale mix of sensors and data analysis. Being able to compute at scale is in the national interest.

The development of computer science has been closely related to the need to solve large, pressing problems. The initial sponsors of computer development were interested in the scientific applications of new technologies, and such developments have given rise to new possibilities. But to make the most of those developments there are a number of challenges and issues that will need to be overcome.

There have been numerous attempts to forecast the directions that computer science might take. As we have seen in the past, predictions on the future of computer science are rarely accurate and are of limited use, with the unexpected growth of smart phones and tablet computers being notable omissions from past predictions about present times.

1 Hilbert, M. & López, P. (2011) 'The World's Technological Capacity to Store, Communicate, and Compute Information'. Science, 332(6025):60-65

2 Gartner (2012) 'Gartner Says More than 1 Billion PCs In Use Worldwide and Headed to 2 Billion Units by 2014'. Available at: http://www.gartner.com/it/page.jsp?id=703807. Accessed 2 July 2012.

3 Internet World Stats (2012) 'Internet Growth Statistics'. Available at: http://www.internetworldstats.com/emarketing.htm. Accessed 2 July 2012.

4 Hickey, K. (2012) 'Terabit Transmission Shatters Internet Speed Record'. *Redmondmag.com,* 13/01/12. Available at: http://redmondmag.com/articles/2012/01/13/terabit-transmission-shatters-internet-speed-record.aspx. Accessed 2 July 2012.

5 CISCO (2012) 'The Zetabyte Era'. Available at: http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/VNI_Hyperconnectivity_WP.html. Accessed 02/07/2012

6 Ammirati, S. (2011) 'Infographic: Data Deluge - 8 Zettabytes of Data by 2015'. *Readwrite enterprise,* 17/11/11. Available at: http://www.readwriteweb.com/enterprise/2011/11/infographic-data-deluge---8-ze.php. Accessed 04/07/2012

The prudent approach adopted in this report is to identify areas where progress seems most likely to occur, based on known drivers of progress, e.g. demand and capacity, and the professional judgement of contributing authors.

Based on current trends, the future of computing is expected to be dominated by embedded and pervasive application drivers. Advances in algorithms, databases and machine learning methods are likely to be major drivers for new science in areas such as biological systems, infrastructure, and earth systems. Judgements about future capacities are problematic given the unknown nature of medium- to long-term discovery. The future of computing is unlikely to be shaped in the near term by exotic computation such as quantum or DNA computers. Computing power is considered likely to grow by 5 or 6 orders of magnitude over the next 20 years and it can be expected that algorithms will increase in power (e.g. speed, data handling capability) by up to 10 orders of magnitude.

To be of practical use, this document sketches a map for future computing research orientated around the scale challenge. It is structured into three major sections:

Grand Challenges looks at the opportunities that computing at scale will open up, and the current technical roadblocks to getting there, through the prism of several scenarios.

Research Opportunities outlines the areas where research is required to make computing at scale a reality, and the cross-disciplinary engagements required to tackle the grand challenges.

Education looks at the shifts in education needed to train a new generation of computing professionals proficient at computing at scale, and also at the impact that computing at scale could have on education.

3

A TRANS-DISCIPLINARY FUTURE

Information and communication technology is a polyglot discipline, having drawn from mathematics, the sciences (mostly physics, ecology and biology), engineering, the social sciences and design. A large amount of ICT research has been about applying results from other disciplines to computing, and about the discovery of 'good tricks'. For example, much work that has been undertaken on algorithms in areas such as machine learning, optimisation, or sorting and searching can be seen as a coupling of 'good trick' discovery with results from mathematics and statistics.

In another example, the field of human-computer interaction has focused on finding good ways to present information to, and get information from, people. Much of this has characteristics typical of a so-called 'wicked problem': one where the problem is clarified as iterative solutions are found, and solutions are 'satisficing' — a term that combines the ideas of satisfying and sufficing -- rather than correct, consistent or complete. In this case ICT has drawn from the social sciences and design.

Similarly, the need to build large software systems at the scale of millions of lines of code has led to a desire — if not always successful — to make ICT more like an engineering discipline.

Perhaps this tendency to appropriate from other fields is simply a reflection of the relative youth and immaturity of computer science. In any case, the scale challenge will force ICT to address the question of how to move beyond good tricks to a fully fledged engineering discipline. Much of the research agenda articulated in this publication is devoted to this. Engineering is about '-ilities' — reliability, dependability, repeatability, composability, scalability, maintainability — and also about standardisation, security, and coping with complexity. These are themes that will recur throughout the document.

ICT is distinguished from most other scientific and technological disciplines by the extent to which it is a universal *enabler*. It is rarely the whole solution to any problem, but it is becoming increasingly true that any solution to a real problem relies upon ICT. The evidence for this is clear — more than 99% of central processing units (CPUs) are *embedded*, that is, integrated in an essential manner into other devices and technologies. The use of ICT in industry reflects this universality: the largest productivity gains and opportunities for wealth creation arise from the use of ICT in diverse areas.

Consider, for example, the announcement in mid-2012 that particle physicists had detected a fundamental particle consistent with the long-hypothesised Higgs boson. The researchers involved in the search for the Higgs boson, credit open source software (Linux and ROOT) with a major role in the discovery[7].

The commoditisation of ICT infrastructure, especially the enormous cost reductions in data transmission and storage, suggests that the trends to ever-more embedded computing, larger-scale data generation and storage, and integration of ICT techniques into existing disciplines will continue. This will in turn have a profound impact on computer science as a discipline, as well as the way we approach research in computing.

The key point is that most problems cannot be solved by ICT alone. That implies that computer scientists need to be able to interoperate with other disciplines such as the social sciences and creative industries more effectively. To succeed, computer scientists need to embrace multiple ways of knowing[8] and doing. A good illustration of this is the work originated at Carnegie-Mellon University on *Computational Thinking*, which attempts to bridge interactions between disciplines by making algorithmic thinking universal. A second example is the growth of cross-disciplinary, domain-specific programming languages, intended to make it easier for programs to become shared artefacts that members of multi-disciplinary teams can all understand and shape.

Sophisticated ICT becoming *infrastructure*[9] highlights that the computing scale challenge is very much an engineering challenge. Coping effectively with scale means that much of computer science (and artificial intelligence) needs to turn to engineering disciplines, and to integrate with other disciplines such as neurobiology, sociology, systems science and geology.

This engagement also implies that the "use-inspired basic research"[10] model advocated by Donald Stokes becomes more important and natural for computer science research — the question is how else to focus the fundamental research in areas that will have the greatest enabling effect?

Finally, all of this suggests that traditional academic boundaries within computer science sub-disciplines — for example between theory and software engineering –– may no longer be appropriate, and that these require rethinking as part of the research journey foreshadowed in this document. Further, computer science needs to take on its enabler role more explicitly, reaching beyond its traditional boundaries and engaging more vigorously with other disciplines, so that future opportunities are not missed. Interesting challenges will now require a deep understanding of computer science along with sophisticated knowledge of other areas, such as biology, physics, engineering or social sciences.

7 See Ubunto Vibes (2012) 'Linux Played a Crucial Role in Discovery of 'Higgs boson''. Available at: http://www.ubuntuvibes.com/2012/07/linux-played-crucial-role-in-discovery.html. Accessed 08/07/2012.

8 See, for example, Vincineti, W. (1990) *What engineers know and how they know it*. London: The John Hopkins Press.

9 Infrastructure can be defined as something you only notice when it fails – see Bowker, C. (2008) *Memory Practices in the Sciences.* Cambridge, MA: MIT Press

10 Stokes, D. (1997) *Pasteur's Quadrant*. Brookings Institution Press.

4

GRAND CHALLENGES

This section outlines several grand challenges that are motivating the agenda of computer science research today and which could, if overcome, have a significant impact on Australia. This is not intended to be a complete set of possible challenges, but rather a group of illustrative examples.

4.1 PERVASIVE HEALTH DATA

The health sector has, for quite some time, collected information on hospital admissions, diagnosis and outcomes as well as clinical and service practices. This data spans a disconnected web of text, images, computerised records and registries in heterogeneous databases.

Making better use of this data requires new information-processing paradigms that can generate and test hypotheses, in addition to being agnostic and unbiased. The methods should be capable of managing very large and interconnected datasets, and of sourcing unexpected variances and latent outcomes. They should be able to handle complex data, of varying characteristics.

The healthcare sector needs computational models that can identify service risks and clinical risks, and discover latent patterns to inform policy. Models are also needed to identify the progression of disease over time, or complications, to inform strategies for intervention, and to model patient profiles and health plans for personalised use.

Computational models are also required to identify latent patterns that underlie best practice outcomes, to identify critical safety issues as well as service and clinical efficiencies, and to permit the formulation of new early-intervention methods to couple data from continuous monitoring.

Because poor coordination between the data and desired prediction can easily lead to 'data leakage' and flawed results, the algorithms needed in the healthcare sector should include machine learning models that can grow and re-parameterise as new data arrives. They should also deal with mixed data types without bias, generate and answer hypotheses, allow clinicians to interact and guide model growth, and incorporate new ways to evaluate algorithms in large datasets when ground truth is unavailable at large scales.

To deliver these outcomes, a multidisciplinary approach is required, with clinicians, health service managers and computer scientists working together.

4.2 GENOME SEQUENCING AND PERSONALISED MEDICINE

The term 'personalised medicine' was coined in 1999 by Robert Langreth and Michael Waldholz in a *Wall Street Journal* article to describe the development by pharmaceutical companies of "a cornucopia of personalized medicines that will produce huge profits into the next century".[11] Originally conceived of as arising from analysis of single-nucleotide polymorphisms (SNP) — DNA sequence variations arising from differences in single nucleotides -- the concept has since expanded to include all types of genomic, proteomic and metabolomic variation, together with any other kind of personalised biological information. The concept also includes all aspects of healthcare that can potentially be honed to the individual, including diagnosis, treatment, prevention and diet. While there are already indications that personalised medicine is becoming a reality, it is the application of ICT that will make this cost-effective on the massive scale needed to have real impact.

Just as social media data will present novel insights into the behaviour of individuals and groups, these biological data will permit novel insights into humans in health and disease, and fundamentally change the way we understand and manage health, disease and quality of life. This will also change the way that physicians and hospitals provide — and patients use — services, leading to changes in workflow, including data collection, storage and analysis.

The volume of data available for personalised medicine to become a reality is increasing at a staggering rate, driven by advances in sequencing technology and decreases in cost. This has led to the sequencing of hundreds of organisms, including many human pathogens and the human genome[12], and the subsequent development of the human haplotype map[13], which describes common patterns of human genetic variation.

Sequencing costs, which originally followed a pattern similar to Moore's law, took a sudden nosedive, beginning in 2008 with the introduction of second-generation sequencers, and have since decreased in cost tenfold every 2-3 years. As a result, SNP genotyping of up to 1 million markers now costs a few hundred dollars, and an entire human genome can be sequenced for less than $10,000. Moreover, the capacity to sequence a human genome for $1000 now appears to be imminent, a price that is anticipated to be a tipping point in terms of the number of sequenced human genotypes worldwide. In 2012, tens of thousands of whole genomes are being sequenced across the world and this will grow exponentially. For medical applications, a variety of 'genomes' for a person will be sequenced at multiple time points. These genomes include DNA sequence, gene expression in multiple tissues, gene methylation and metabolites, leading to an integrative approach[14] to 'omics' profiling.

The availability of SNP and human genotype data, and the ability to analyse them, has already begun to transform health research. This has been led by increases in the computing and statistical power necessary to perform ever-more sophisticated analyses. But there remain significant computational and statistical challenges, particularly as the sophistication of the analyses increases to account for ever more nuanced associations with health and disease.

This is the challenge of the coming years: developing and using ever more sophisticated computational and statistical techniques to link personalised biological information with descriptions of health and disease—in order to have an increasingly sophisticated impact on clinical outcomes.

Genome sequencing will also become a routine tool for high-density data gathering in non-medical areas of research and application, such as plant and animal breeding, ecology and evolutionary biology. Bottlenecks for research and application will lie in asking the right question (or generating the right suite of hypotheses) and in analysing data accordingly. Limitations will be on the number of individuals (plants, animals) with phenotypic observations, because the cost of a whole genome sequence will be low, relative to the cost of a phenotype.

11 Langreth, R. and Waldholz, M. (1999) 'New era of personalized medicine: targeting drugs for each unique genetic profile'. *Oncologist*, 4(5): 426-7.

12 Consortium International Human Genome Sequencing (2001), "Initial sequencing and analysis of the human genome". *Nature*, 409(6822):860-921.

13 International HapMap Consortium (2003) 'The International HapMap Project'. *Nature*, 426(6968):739

14 Chen, R., Mias, G.I et al (2012) Personal Omics Profiling Reveals Dynamic Molecular and Medical Phenotypes. Cell, 148(6):1293-1307

4.3 THE AUGMENTED HUMAN

Overcoming the frailties of the human body is a constant endeavour of science, engineering and medicine, a theme in speculative writing, a matter of ongoing legal, ethical and moral debate and a challenge to social processes. The convergence of various computing-related technologies will lead to dramatic changes in the types of intervention that are possible. These will assist those who have disabilities and enhance the capabilities of the wider population.

Australian scientists have led breakthroughs in areas such as the cochlear implant[15] and they are taking a lead role in the development of bionic eyes[16]. Exoskeletons, robotic frames augmenting limb function, and neural-controlled prosthetic limbs are moving from the laboratory to real deployment. Pacemakers and other electronic or computer-controlled implants are commonplace. Sensor technology is advancing to ubiquity, with sensors literally woven into our clothes and implanted in our bodies, communicating via "body area" networks to mobile phones and on to the cloud. Google's Project Glass program[17] highlights the movement of augmented reality visual interfaces from prototype to mass use. Bell's MyLifeBits research project[18], automatically recording life activity, demonstrates the potential to augment personal memory.

Combining such technological breakthroughs creates enormous potential to augment and enhance human performance, to prolong quality of life, and increase human capability.

It will become commonplace for augmented reality interfaces to remind individuals of the name of the person in front of them, appointments they have coming up, salient features of buildings nearby, the cost of that car they are looking at, what the street looked like last time they visited, who they were with then, and what it would look like if the building they point to was not there.

It will be possible for health to be continually monitored and interventions scheduled as warning symptoms appear. Rooms will adapt to the presence of individuals in them, their preferences and their tasks. Common disabilities, such as paralysis and blindness, will largely be overcome. Prostheses may be a viable option to enhance "healthy" human performance for heavy manual occupations. It may become commonplace to provide enhanced quality of life to the elderly.

Achieving the augmented human is indeed a grand challenge. While computer science advances (noted in Table 1 in the Appendix) will underpin many achievements, they will require significant multidisciplinary collaboration with areas such as medicine, engineering, and psychology, to achieve the technical solutions. Success will also lie in collaboration with the fields of sociology, philosophy, political science and the legal community to cover legal, ethical and social dimensions.

15 Cochlear (2012) 'History'. Available at: http://www.cochlear.com/wps/wcm/connect/au/about/company-information/history/history. Accessed 1 September 2012

16 Bionic Vision Australia (2012) http://www.bionicvision.org.au/

17 Available at: https://plus.google.com/+projectglass/posts. Accessed 25 August 2012.

18 Microsoft Research (2012) 'MyLifeBits'. Available at: http://research.microsoft.com/en-us/projects/mylifebits/. Accessed 20 August 2012.

4.4 DISASTER MANAGEMENT

The impact of natural disasters has been increasing in recent decades and this trend is predicted to continue[19]. Natural disasters not only cause human tragedies; they also have significant, long-term economic consequences for entire regions, and governmental and commercial organisations.

Information and communication technologies have a fundamental role to play in mitigating the consequences of natural disasters. At a 2005 donors' conference for tsunami relief in New York City, a European ambassador said: "We don't need a donor's conference, we need a logistic conference"[20]. A US House of Representatives report[21], following the devastation caused by Hurricane Katrina in New Orleans, pointed out the need to go beyond situational awareness and to provide advanced decision support for identifying key decisions in these complex, fast-paced, and uncertain environments. Technology can play a key role in decreasing the number of people affected and the economic losses caused by disasters.

However, the computational challenges are daunting. Disaster management systems will need to collect, aggregate, synthesise, and communicate massive amounts of information from a wide variety of heterogeneous sources, including sensors and social media. They will require accurate predictions from complex simulators at unprecedented speeds.

These predictions will then feed optimisation algorithms that will compute in real time and continuously revise tactical, first-response, and recovery decisions. These algorithms will need to solve large-scale online stochastic optimisation problems over complex interdependent infrastructures, including the telecommunication, electrical power, and transportation networks. Their recommended decisions and their justifications will need to be communicated effectively to decision-makers and first responders using advanced 3D visualisation and "what if" analyses. Last but not least, these decisions will need to incorporate human response in emergency situations, a requirement that stresses the need for a close collaboration between physical and social sciences.

19 Department for International Development (2011) *Humanitarian Emergency Response Review*. London: DFID. March 2011. Available at: http://www.dfid.gov.uk/Documents/publications1/HERR.pdf

20 Managan, J., Lalwani, C. & Butcher, T. (2008) *Global Logistics and Supply Chain Management*. Chichester: John Wiley, cite Economist.com 'Global Agenda', 5 January 2005.

21 US House of Representatives (2006) *A Failure of Initiative*. Final report of the Select Bipartisan Committee to Investigate the Preparation for and Response to Hurricane Katrina. Available at: http://orise.orau.gov/csepp/documents/planning/evacuation-documents/federal-reports/mainreport.pdf

4.5 MONITORING AND MANAGING THE ENVIRONMENT

Meeting environmental challenges in a sustainable fashion is likely to be essential for maintaining the Australian quality of life -- assuming current predictions about climate change and other factors affecting the environment, including ocean change, continue to be validated. This critical need will require an extensive environmental sensor network for a massive space and time data repository, and coupled models defining behaviour. This system will be the basis for future planning and tactical day-to-day control of resources.

An integrated earth sensing, modelling, and action system implies both a massive physical size and time scale. It involves a sensor network and data service for limited resources such as water; highly accurate, continuously validated modelling of the environment; the ability to generate scenarios such as natural disasters; and the ability to undertake planning activities that on the one hand are grounded in accurate data, but on the other, allow broad consultation and engagement in the decision-making process, including engagement with social networks.

The information coming from sensor-based modelling ranges in size from sub-cubic millimetres for soil, through to an 8 million square kilometre continent. These models will affect political policy changes; resources management e.g., water, land, energy versus CO_2 production, agricultural productivity, mining; and any factors that signal natural disasters (see previous challenge 4.3).

In addition to a sensor network for natural resources and phenomena, man-made networks including highways, energy (electricity, coal, gas), and information flow are expected to be more heavily instrumented. For example, having information about the true state of Australia's moisture, temperature, and environmental health will improve weather predictability, aid agriculture production, predict disaster, and assist in handling disasters such as fires and floods.

The National Broadband Network (NBN) plays a critical role by providing an overall infrastructure for this comprehensive sensor network. For example, 9 million sensors would ideally provide a 1km² monitoring grid. As the NBN will only cover about half of the continent with fibre and wireless coverage, this would be coupled with satellite sampling to have a validated model for the entire country. Simply having sensors at every residential NBN access point provides a base of 8 million sensors that is virtually "free" and also provides imputed data about activity. Another network that is "free" to collect is dynamic sampling using smart communicating cars and mobile phones as sensors on moving vehicles such as trains and buses. By 2030, highways are likely to have sensors extensively embedded.

A near-term and achievable goal for this challenge is to have a complete model of all of the country's water resources that, coupled into a climate model, could answer any sort of "what if" queries relating to use, including agriculture, mining, city and regional planning, and energy generation.

The environmental challenge is interdisciplinary. Each of the networks, models and needs are in separate disciplines and industry silos e.g., agriculture, energy, mining, water management and regional planning.

4.6 MANAGING MASSIVE ENVIRONMENTAL DATA SETS

Australia is capturing and building a large number of massive data sets describing the natural environment. These range from continent-wide geological and geochemical data, acquired from a growing range of geophysical sensing campaigns, to terrestrial and marine environment data captured from a variety of sources, including Geoscience Australia, federal and state geological surveys, resource and energy companies, and the National Collaborative Research Infrastructure Strategy (NCRIS) programs in geoscience (AuScope), terrestrial ecosystems (TERN) and marine environments (IMOS).

These data sets are:

- Dynamic, being added to and growing at a continually high rate;
- Diverse, being increasingly acquired by different sensors at different times, places and scales;
- Require unique characterisation and discovery, because the location and composition of critical features in the data is sparse and often unknown *a priori*; and
- Massive in size, aggregating to at least hundreds of terabytes.

Managing and extracting value from these massive environmental data sets is a defining challenge in information technology. The challenge offers an opportunity to bring about a transformational change to the way society understands, manages and stewards its world.

Addressing this challenge will require major coordinated advances in three key areas:

1. **Information**: The assimilation, representation and management of massive real-time information. A central task is to develop extensible standards and abstractions to describe large-scale stochastic data in a manner that enables information to be efficiently fused or assimilated over time and space. Extensive work has been completed on much smaller scales with individual sensor types, including computer vision. A lot remains to be completed for much larger, more diverse sensor types spanning different spatial and temporal scales.

2. **Discovery**: Improvement of machine learning and feature discovery techniques to allow semi-supervised, unsupervised, and incremental machine learning and feature discovery in massive data. A defining characteristic of natural environmental data sets is that every feature of interest is different and hard to quantify *a priori*. Finding critical and relevant features goes well beyond a simple search and must instead address the core issue of discovery from data. Related work has been undertaken on a much smaller scale in areas such as dimensionality reduction and structure discovery.
3. **Computation**: Methods include data-centric fusion and machine learning algorithms that are amenable to distributed, cloud, and parallel computation. There is a major challenge in both scaling and applying current (and future) machine learning methods. Most critical to the discovery problem are the need to establish relations between physically disjoint data sets and the compromises that must be made to achieve this in distributed environments. Similarly, some methods, including dimensionality reduction, can be structured in a naturally exploitable, parallel and distributed manner.

Geology is an exemplar of this challenge. A wide range of diverse sensory information is collected from gravity and electromagnetics, through seismic and magnetotelluric imaging to direct drilling. Each of these sources of information allows inference to be made about the underlying geological structure and material properties including stratigraphy, density, porosity, chemical composition, thermal characteristics and stress fields. Together these data sources could allow geological inference to be undertaken at many different scales -- from individual mineralisation regions to continent-wide geological understanding.

At a continental scale, with the currently available data sets in Australia, the geological inference problem is hugely more challenging, in both size and complexity, than anything yet attempted in the machine learning field.

4.7 COMPETING EFFECTIVELY IN A LONG-TAIL WORLD

Marc Andreessen recently argued that "software is eating the world"[22]. To illustrate, consider Amazon and Borders. Amazon began as a discount online bookseller that thought its competitive advantage lay in avoiding retail shopfronts and selling books more cheaply. Its real advantage lay in its software systems and their huge repository of customer data. Choosing to expose all its software systems as retail services resulted in 'fortuitous opportunities' such as Amazon Web Services.

Today Amazon is a software company, a retailer, and the largest provider of cloud computing. Borders, by contrast, decided its competitive advantage was retail shop fronts, and used software for stock, inventory and payrolls, and even outsourced their online business to Amazon. Interestingly, Amazon is now a dominant player in the world economy, whereas Borders no longer exists within many markets and is a much smaller retailer than Amazon.

Other companies, such as FedEx, have undergone similar transformations. Pixar is a software company that now makes movies. Google has led a transformation of advertising and LinkedIn is having a similar impact upon recruitment. Walmart, a retail store chain, is now in large part a supply chain and logistics systems business.

22 Andreessen, M. (2011) 'Why Software Is Eating The World', *Wall Street Journal Online,* first published 20 August 2011. Available at: http://online.wsj.com/article/SB10001424053111903480904576512250915629460.html. Accessed 15/11/12

In each economic niche, participants form a 'long-tail' or power-law distribution, with a small number of large players, and a large number of smaller entities. Those on the tail must find a way to market their products widely, or settle for smaller returns. Thus, on iTunes and Amazon there are a small number of large-volume products, and many small-volume ones, but everything sells to someone. For example, when one aggregates all interest in, say, bagpipe music, these niches turn out to be profitable because Amazon and Apple allow products to reach a very wide audience.

Due to the size of Australia's local market, many Australian companies are going to be in the long tail. So they need to reach the widest possible audience, and operate as effectively as possible. Effectiveness comes from having and using the right data. Data can be mined and analysed to help performance, and to track activities across the entire value chain. This involves machine learning to identify trends and opportunities, and optimisation to improve efficiencies. The 'internet of things' -– sensor-based approaches to the automation of tracking objects and people — in turn, greatly increases the data pool that drives learning and economic optimisation.

4.8 SOCIAL MEDIA

Social media is "a group of Internet-based applications […] that allow the creation and exchange of user-generated content".[23] More generally, social media is the embodiment of various digital traces of large-scale human behaviour. Examples include blog posts, friendship entries on social network platforms, mobile phone calls with associated locations, swipe records for transit fare cards, and many more.

The volume of social media data is reaching unprecedented breadth, depth and scale, with emerging studies on 100,000 human movement trajectories[24] and 30 billion instant messages[25]. Such massive, dynamic and ever-growing data collection may help to solve difficult problems in both understanding human behaviour (i.e. social sciences) and building better computing systems (i.e. computer science).

The availability of large amounts of data and the ability to analyse them has transformed scientific research in biology and physics. Such a transformation on data-driven 'computational social science' is just happening now, as articulated by Lazer et al. in a visionary article published recently in *Science*[26] .

23 Kaplan, A.M & Haenlein, M. (2010) 'Users of the world, unite! The challenges and opportunities of Social Media'. *Business Horizons,* 53(1):59–68.

24 González, Marta C , César A Hidalgo, and Albert-László Barabási (2008) 'Understanding individual human mobility patterns'. *Nature,* 453(7196):779-782

25 Leskovec, J & Horvitz, E. (2008) 'Planetary-Scale Views on a Large Instant-Messaging Network'. *Proceedings of the 17th international conference on World Wide Web,* 915-924

26 Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabási, A-L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M., Jebara, T., King, G., Macy, M., Roy, D. & Van Alstyne, M. (2009) 'Computational Social Science'. *Science,* 323:721-723

A joint opportunity is now presented to social scientists and computer scientists, to effectively analyse and understand the massive longitudinal data sets of millions of people. These data will, like a magic looking glass, produce novel and potentially useful insights into the behaviour of individuals and groups, as exemplified by recent studies that used social media to explore the language of grief[27], and to understand the dynamic evolution of urban neighbourhoods[28].

Successfully tackling social media will produce paradigm shifts in several important areas such as better management of natural disaster response, with information disseminating through the most efficient route (see challenge 4.3); traffic patterns in a city will be better anticipated, with minimum road-incident delays; and data centres will be able to adapt to changes in usage patterns, improving execution efficiency and conserving energy.

27 Brubaker, J.R., Kivran-Swaine, F., Taber, L., & Hayes, G.R. (2012) 'Grief-Stricken in a Crowd: The Language of Bereavement and Distress'. *Sixth International AAAI Conference on Weblogs and Social Media.* 20/05/12.

28 Cranshaw, J., Schwartz, R., Hong, J. & Sadeh, N. (2012) 'The Livehoods Project: Utilizing Social Media to Understand the Dynamics of a City'. *Sixth International AAAI Conference on Weblogs and Social Media.* 20/05/12.

5 RESEARCH DIRECTIONS: TOWARDS COMPUTING AT SCALE

This section discusses each of the challenges raised by the issue of computing at scale. Note that while this list is wide-ranging and detailed it is not intended to be complete. Rather, these are illustrations identified as important, about which useful observations can be made.

Four broad thematic groups are presented in the following pages: big data, computation, interaction, and systems.

5.1 DOMAIN-SPECIFICITY

A common thread through all the grand challenges listed above is *domain-specificity*. In brief, it has been argued above that computer science needs to engage more closely with other disciplines. In order to bridge the gap it would be helpful to be able to program, analyse data, and perform other operations in a form and with a vocabulary accessible to those other disciplines. Computer scientists call tools that are adapted in this way *domain-specific*, because of this tailoring to a particular domain.

For example, storing and indexing massive and rapidly changing personalised biological information requires domain-specific data and storage paradigms that are aligned with the needs of biological disciplines. There are similarly specific requirements when dealing with genomes, mutations and single-nucleotide polymorphisms as basic data types, and for automating the generation of biological and health information.

Programming languages are another area where domain-specificity is an important challenge. Examples include domain-specific languages for geological application; and end-user programming languages and capabilities for research, clinical systems, and social media systems. The specific languages, patterns and tools for highly flexible business process engineering are further illustrations of domain-specificity.

In software engineering, challenges that are specific to individual domains include modelling and reasoning about the environment in terms of scenario planning; highly configurable enterprise systems and business process solutions for use by small enterprises; and customisation of applications and data to various clinical settings or disease systems.

Similarly, in the field of interaction, the healthcare domain requires seamless integration into clinical and service delivery practice.

These examples point to the gulf that exists between computing practitioners and researchers on the one hand, and the many communities of ICT users on the other. Computer scientists need to bridge this gap between computing and partner disciplines, so that practitioners and researchers in those disciplines can more effectively work with ICT professionals, as well as build and evolve ICT systems themselves.

This common need for domain-specificity requires a common response from the computing research community. A systemic response to these needs is required, cutting across all aspects of computing research. Otherwise the computing research community will run the risk of many hundreds of incompatible and idiosyncratic domain-specific tools emerging. This will make interactions between computing and other disciplines more difficult rather than easier.

Several common approaches are emerging that can form the foundation for this shift, including the concept of 'X-as-a-service' where developers abstract away from the details of a particular capability so that it can more easily be integrated into larger systems. Examples of this approach include the idea of machine learning as a service, or the shifting of large-scale data into the cloud, with an appropriate interface that allows users to treat the data and its storage/search/analysis capabilities as services.

Another approach that may help this shift is the use of domain-specific programming languages and libraries for existing languages, which introduce fundamental capabilities into the language to facilitate abstractions useful to specific domains which are not found in traditional programming languages. An example would be programming languages with genomes as basic data types, with associated semantics. This allows software development tools to reason about genomes, improving program performance as well as allowing some classes of software errors to be identified more easily.

Domain-specific engineering techniques — for example, particular algorithms or software patterns — are also part of the common response from the computing research community to the need for domain-specific ICT systems.

5.2 BIG DATA

Renowned computer scientists agree that nowadays "everything is data"[29]. "Big-data computing is perhaps the biggest innovation in computing in the last decade"[30]. "Big data" is a critically important future area of computer science[31 32].

Moore's law has enabled rapid growth in society's ability to process data. But the ability to move the data, store it and, crucially, to understand it, is not keeping up. New solutions are required for this challenge.

The data-centric view of the world leads to a new view of how science is done — summed up as "data-intensive scientific discovery" in Jim Gray's *Fourth Paradigm*[33].

The third paradigm -- simulation and computation -- has been the focus of much attention in computing recently. The crucial aspect of the fourth paradigm is the analysis. All the sciences[34 35 36 37 38 39] recognise the enabling role that data-centric computing will play in their future.

29 Cohen, J., Dolan, B., Dunlap, M., Hellerstein, J., & Welton, C. (2009). 'MAD Skills: New Analysis Practices for Big Data'. Available at: http://db.cs.berkeley.edu/jmh/talks/MADSkills-vldb09.pdf. Accessed 15/08/12.

30 Bryant, R. E., Katz, R. H., & Lazowska, E. D. (2008). 'Big-Data Computing: Creating revolutionary breakthroughs in commerce, science and society'. Computing Research Association. Available at: http://www.cra.org/ccc/docs/init/Big_Data.pdf. Accessed 23/08/12.

31 Bell, G., & Hey, T. (2009) 'Beyond the Data Deluge'. *Science,* 323(6 March):1297-1298

32 Bollier, D. (2010) *The Promise and Peril of Big Data.* Washington: The Aspen Institute.

33 Hey, T., Tansley, S. & Tolle, K. (2009) The Fourth Paradigm: Data-Intensive Scientific Discovery. Redmond, WA: Microsoft Research.

34 Nature Editorial (2008) 'Community cleverness required'. *Nature* 455(7209):1

35 Science Magazine (2011) 'Special Online Collection: Dealing with Data'. *Science,* 331(6018), collection available at: http://www.sciencemag.org/site/special/data/

36 Overpeck, J. T., Meehl, G. A., Bony, S., & Easterling, D. (2011) 'Climate Data Challenges in the 21st Century'. *Science,* 331(6018):700-702

37 Stein, L. (2010) 'The case for cloud computing in genome informatics'. *Genome Biology,* 11(207):1-7

38 King, G. (2011) 'Ensuring the Data-Rich Future of the Social Sciences'. *Science,* 331:719-721

39 Manyika, J., Chui, M., Brown, B., Bughin, J., Dobbs, R., Roxburgh, C., & Hung Byers, A. (2011). *Big data: The next frontier for innovation, competition, and productivity.* McKinsey Global Insitute.

Computer scientists need to develop new paradigms to address the needs of the data-rich society. In particular, computer scientists need to develop theory and algorithms that will lead to the production of new tools and techniques that are required to advance the state of the art in managing and making sense of data. They need to address major challenges with respect to data volume, large-scale knowledge discovery, data consistency and privacy and data trust, which are common across a large range of application domains. Big-data mining issues require distributed computing solutions that are robust, secure and trustworthy.

Big data can serve as a driver for much contemporary research in computer science. The subsections below focus on just three important areas where this applies:

1. Federated, distributed storage and computation (also known as cloud computing)
2. Machine learning
3. Security and privacy

5.2.1 FEDERATED, DISTRIBUTED STORAGE AND COMPUTATION (AKA CLOUD)

Federated, distributed storage and computation (Cloud) is the use of a network, usually external to the user, to provide access, storage and/or computational services. From a business perspective it allows costs to be shared, though ownership and control is also shared which presents additional limitations. Some of the key issues are:

- The cloud as home for complex simulations and data integrations.
- Storing and analysing (filter, correlate, learn) exascale data from trillions of devices in near-real-time.
- Managing millions of users and exabytes or more of data while supporting a high level of individualisation.
- Large-scale, distributed, parallel alignment with algorithmic techniques, work with physically distributed, heterogeneous and/or multiscale data sets.
- Storage and access to very large, distributed, heterogeneous data.
- Data acquisition — aggregation from trillions of diverse sources and devices in near real time, including approximate matching; data quality and cleansing.
- Automatically integrating community/vertical specific data types; and extracting metadata to publish into an environment to facilitate resource discovery and collaboration.
- Dynamically configurable exascale data storage and management systems that enable sub-second real-time queries across geographically distributed computation centres, and millions/billions of users.
- *In situ* processing to ensure scalability by moving computation to the data, eliminating the costs of data movement. This is especially challenging when analyses must span distributed data resources and be organised in computational pipelines or workflows.
- Integration of existing analytical and visualisation tools; enabling end users to better reuse and share appropriate domain-specific end user tools.
- End user programming – interfaces and languages that appeal to non-computer scientists for manipulating, analysing and querying exascale datasets.
- Creation/generation of domain-specific, individualised and customisable user interfaces; natural language processing (multiple data sets require NL or computational linguistics techniques, to make them understandable).
- Exascale testing techniques, operating on exascale data and geographically distributed computation units.
- Issues of ubiquity as related to privacy — techniques and systems that enable isolation and guarantee that privacy properties are respected according to different levels of privileges/requirements.
- New control centres that manage the future computer system (whatever characteristics it may have), tools to enable operators to better cope with failures that can propagate faster and wider, as well as exascale deployment and provisioning.
- Model generation/validation: new modelling techniques and formalisms to deal with exascale. How such things are defined, validated and made usable.
- From computation to action -- this is especially important in the big-data setting, but is more general. This is about how to take the results of computation and analysis, and integrate them into decision-making and operational activity. This includes 'complex event processing', as well as the 'near-real-time analytics' issue, but also has many business process and organisational structure issues.
- Context-awareness. For several of the applications, one needs to process in ways that differ according to context and purpose; a special case of this is personalisation, which is a huge research challenge.

5.2.2 MACHINE LEARNING

Machine learning (ML) is the science of making sense of data and thus a crucial component in dealing with the ever-increasing amounts of data available. However, at present it lacks many of the hallmarks of an engineering discipline. Some of the key problems are:

- How to standardise ML to make it easier for users? How to make ML composable? How to facilitate the interchange of models, methods and data between different ML solutions?
- Develop ML as a service.

- Human-computer interaction for ML: how to present the outputs of ML in a manner that facilitates their use? How to reason with uncertainty and represent it?
- Privacy, ethics and social aspects of ML: how to enable the benefits of ML on data associated with people, while protecting their privacy and not exploiting them?
- How to deal with heterogeneous data provenance? Data will become more valuable. How to do ML across data owned by various parties with complex constraints?
- Auditability and repeatability: how to deploy ML in mission critical systems that allows proper auditing of what the ML algorithms do?
- Software engineering for data-centric computing and the connection to ML. Develop new programming languages and frameworks that facilitate the deployment of ML. How to better integrate ML with embedded systems; web; mobile; database; enterprise etc.
- There is a compelling need for a more comprehensive theory of ML dealing with:
    - Different models of data (stochastic vs non-stochastic; streams; batch; online; active; continuous vs discrete; different notions of structured data; dealing with data heterogeneity)
    - A richer set of data imperfections (noise, missing data, etc.)

    - Reductions and relations between problems
    - Understanding the relationship between data and computational complexity.
- Develop a richer continuum of methods along the exploratory-confirmatory dimension.
- ML using different calculi of uncertainty.
- Decision-making based on ML. Develop systems theory for ML that facilitates its pervasive application.
- How to 'abstract' ML so that the user is unaware that it is taking place.

5.2.3 SECURITY AND PRIVACY

An important issue in the future of computer science networks and devices is their security.

5.2.3.1 MISTAKEN SECURITY ASSUMPTIONS

Virtually all applications of cryptography in today's networks and devices are built around algorithms and protocols whose security is dependent on one or more of the following assumptions:

a) An efficient mathematical attack will never be feasible; e.g., an algorithm for efficient factorisation does not and never will be developed;

b) Computing resources will be insufficiently powerful to reveal the plain text of a cipher text message by brute force or algorithmic attack;

c) Public key infrastructure is and always will be secure from attack;

d) Trusted third-party artefacts, services, design information, or sensitive records will never be compromised, or if they are, a compromise will have no security impact on reliant parties;

e) New technology, such as quantum computing, will never be developed to sufficient scale to support efficient operation of algorithms such as Shor's algorithm (an algorithm already developed that can efficiently break RSA, ECC and DH public key cryptography).

Most of these assumptions are now known to be false. There is, however, currently little evidence of action to mitigate these known risks.

Quantum key distribution provides theoretically secure communication that does not rely on the aforementioned security assumptions. In addition, quantum random number generation provides an ideal source of true random cryptographic key material for distribution by quantum key distribution systems.

5.2.3.2 QUANTUM KEY DISTRIBUTION

Quantum key distribution enables the distribution of cryptographic key material between two remote parties who are connected by an optical channel. This key material can be used as one-time pad key material, or can be managed as conventional key material for a variety of applications and devices, protecting information at rest or in transit.

The idea underlying quantum key distribution is that information is encoded onto quantum states of light, which are transmitted to a receiver that measures them. Security is guaranteed by the laws of physics, specifically the so-called no-cloning theorem and the Heisenberg Uncertainty Principle in quantum mechanics. This implies that an attacker, however powerful, cannot determine full knowledge of unknown quantum states.

There are two complementary approaches to quantum key distribution: discrete variable quantum key distribution, where information is encoded on single photons with measurements made by single photon detectors; and continuous variable quantum key distribution, where information is encoded on bright lasers and measurements made with homodyne photo detectors, which comprise two balanced photo detectors, a partially reflective mirror and a local oscillator laser. Both approaches have been proven to be theoretically secure, but the latter approach promises higher detection efficiencies, better compatibility with current telecommunications technologies and has the potential of achieving higher secret key rates.

5.2.3.3 QUANTUM RANDOM NUMBER GENERATOR

Nondeterministic random number generators are ideal sources for random cryptographic key material distributed using quantum key distribution systems.

Quantum random number generators derive random numbers from measurements conducted on specific quantum processes or quantum systems. The randomness of the outcomes of these measurements is of quantum origin as described by the laws of quantum physics. This is in contrast to hardware random number generators, which derive random numbers from physical processes.

Light is one possible source of quantum-random noise. If light is treated as consisting of particles, measuring which random path a single-photon will take when impinged on a partially reflective mirror is a source of quantum-random noise. However, detectors and sources of single photons are expensive and have limited bandwidths.

Alternatively, by treating light as waves, it is possible to use the quantum vacuum states of light measured using optical homodyne detectors as a source of quantum-random noise. In this case, owing to high-detection bandwidths, it is possible to generate true random numbers at higher rates, which can be used for cryptographic key material for quantum key distribution as well as other applications.

It seems likely that future developments in these areas of quantum key distribution and quantum random number generation will prove important in terms of the security of computer science networks and devices.

5.2.3.4 MANAGING DATA CURATION AND PROVENANCE

The emergence of large, multifactorial, federated data sets creates many privacy and trust issues. While some of these issues are clearly societal, legal and ethical approaches to dealing with them will most likely have technical implications. Challenges such as managing data curation and provenance, when to share data, how data can be anonymised before sharing (and when this is allowed to be done), and how to manage digital identities, are problematic enough at present, and likely to become much more so as we reach exascale, globally accessible data sets and beyond.

For example, can it be established that the data from a network of sensors was trustworthy and hasn't been tampered with or perturbed? How can the sensor networks that are relied on against large-scale cyber-attack, be protected? How can robust privacy-protection mechanisms be implemented in a world where hundreds or thousands of sensor networks could be collecting information on everyday lives, such as what is eaten, how many times the toilet is flushed, the speed at which cars are being driven and who people meet?

5.3 COMPUTATION

Facilitating computation and the nature of what can be computed are fundamental issues for computer science and engineering. In the scientific arena, the use of high performance computing to solve computationally intensive tasks is critical. This extends to the development of specialised programming languages that exploit established problem-solving techniques and programming abstractions to provide powerful computational paradigms. Using these languages allows the programmer to focus on the problem specification and leave the resolution to the programming language's computational mechanisms. From an engineering perspective, the discipline of design exploits these computational mechanisms in a systematic manner to develop large-scale solutions to complex problems.

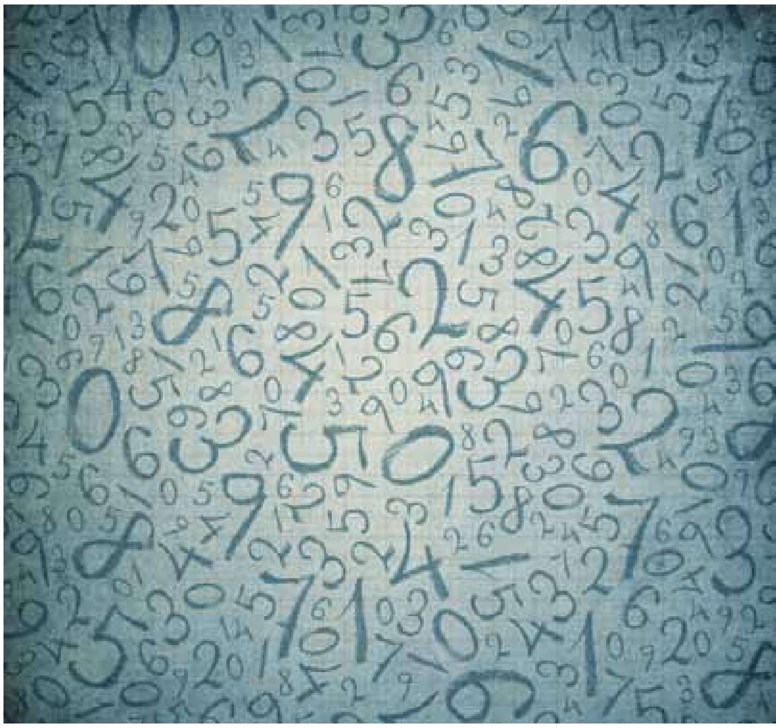

5.3.1 HIGH PERFORMANCE COMPUTING

Key challenges facing the high-performance computing (HPC) community include:

- The development of new high-performance computing (HPC) architectures and effective algorithms to apply data-intensive computing to exabyte- and zettabyte-sized data collections of the kinds described in the grand challenges section of this document.
- The development of new HPC architectures and energy-efficient algorithms that can cope with a 1000-fold scale-up of performance to achieve exascale[40] computations, while avoiding a linear scale-up of energy requirements.
- The effective application of hybrid (massively parallel, GPU, data-intensive, cloud) computing to an increasingly broad spectrum of disciplines.
- Means of parallelising computation and data applications in a wide range of research domains that are of interest to Australia.
- Reducing and greening the rapidly increasing amount of energy consumed by very large numbers of very small computing components.
- Finding ways for global research communities to address global research problems by using global HPC resources, computing data from collections dispersed globally over relatively narrow, global network pipes.

40 HPC computer speeds are measured in 'flops' – floating point operations per second. An exascale machine would be able to execute 10^{18} floating-point operations per second. Teraflop machines (10^{15} flops) are the current state of the art.

- Managing and ensuring the durability of data in a situation, where the capacity to produce data threatens to exceed by many times the capacity to store and manage.
- Development of domain-specific modelling, simulation, data analytics, image analysis and visualisation techniques for use by researchers in particular domains. An example is Galaxy, a toolset for genomics researchers.
- Training researchers from particular disciplines to develop the skills to create leading, domain-specific software tools. Skills in developing advanced research software tools are already in short supply, and these shortages will be accentuated to the point where they may hamper progress. Skills in the effective deployment of discipline-specific packaged HPC software will also be in strong demand and short supply.

5.3.2 PROGRAMMING PARADIGMS

New programming paradigms continue to be introduced, particularly suited to various tasks such as web programming, mathematical programming, etc. The dominant paradigms arguably remain:[41] procedural/imperative, object-oriented, declarative and functional.

Established domain-specific programming languages include Matlab (for matrix calculations, particularly in engineering applications), R (statistical applications) and Ruby (web applications). In addition, there are experimental languages that are starting to incorporate sophisticated problem-solving methods as part of the language, particularly in artificial intelligence. Instead of the programmer having to code search techniques, statistical and probabilistic reasoning algorithms or learning algorithms, these are provided as primitives in the programming language.

For example, at the NIPS 2008 Workshop on Probabilistic Programming[42], several probabilistic programming languages were discussed, including Church, DBLOG, PyBLOG, CP-Logic, Bach, Markov Logic, Dyna, FACTORIE and CSoft. Other experimental languages include constraint-programming languages, Answer Set Programming for optimisation problems, the Golog[43] family of languages for cognitive robotics and the Robot Learning Language.[44] These programming languages allow programmers to use problem-solving methods without having to program them from scratch or have them provided in a library. As a result, programming languages are becoming more powerful through the use of appropriate abstractions. A major challenge is to deal with the trade-off between the expressiveness of these languages and the complexity of the associated inference procedures.

The benefits of this approach are manifold. Enhancing programming languages with sophisticated in-built problem-solving methods increases the programmer's ability to deal with complex and rapidly evolving software. Many of the languages described above were developed to solve complex optimisation problems. From a software engineering standpoint these languages cater for modelling domain-specific aspects of problems and facilitate the assembly of complex decision support systems compositionally. They also make it easier to engineer reliable, scalable and adaptable software systems. Some of these languages are focused on providing massive-scale and massively parallel high performance computations. They also facilitate the development of rapidly evolving software through the abstractions they provide.

Challenges introduced by this approach include:

- The trade-off between expressiveness of a programming language and the computational complexity of the associated problem-solving procedure.
- The ability to control computation through heuristics.
- Computational complexity of algorithms/programs.
- How to take advantage of available resources (e.g., CPU, GPU, multi-core, etc.) without the user having to know the defaults.
- Formal verification of programs, especially in a seamless manner hidden to the programmer (e.g., static analysis, automatic verification).
- Guarantees of program correctness.
- Parallelism and concurrency.
- Defining the correct level of abstraction.
- Programming languages for lay people (i.e., non-programmers).
- Assurances for secure software.[45]

41 C.f. Van Roy, P. (2009) 'Programming Paradigms for Dummies: What Every Programmer Should Know', In G. Assayag & A.Gerzso (Eds) *New Computational Paradigms for Computer Music.* Delatour France: Ircam-Centre Pompidou.

42 See Probailistic Programming: universal languages and inference; systems; and applications. Available at: http://probabilistic-programming.org/wiki/NIPS*2008_Workshop. Accessed 10/9/12.

43 Reiter, R. (2001) *Knowledge in Action: Logical Foundations for Specifying and Implementing Synamical Systems.* MIT Press

44 This discussion is also important for the area of programming languages. See Pierce, B. (2009) 'Panel Summary notes on Grand Challenges in Programming Languages'. Available at http://plgrand.blogspot.com.au/. Accessed 30/7/12.

45 Thompson, K. (1984) 'Reflections On Trusting Trust'. *Communications of the ACM,* 27(8):761-763

5.3.3 SOFTWARE ENGINEERING

Central to many of the grand challenges facing computer science is the need to deal with size and complexity (in the organisational sense) of both data and the software used to manipulate it. Software engineering techniques are essential to the development of such complex high transaction systems, meaning that software engineering research is a key enabler of grand challenge outcomes.

There are numerous research challenges relevant to software engineering:

Modelling and abstraction: Model-driven and domain-specific modelling approaches are major software engineering techniques for dealing with scale. These techniques leverage common patterns of behaviour or structure into compact, composable abstractions, allowing much to be expressed simply, in terms closer to those of the problem domain. This has the potential to increase productivity and ease of understanding, and to provide more opportunity for end-user software development. The resulting models or abstraction compositions can then be used as the basis for automatic code generation or model transformation. An open issue is what abstractions and modelling approaches are suitable for grand challenge needs.

Software qualities: Dealing with the trade-offs between various software qualities is central to software engineering. The cross-cutting nature of qualities such as speed, reliability and security introduce significant complexity into software and the process of developing software. How to balance and meet such cross-cutting needs in an exascale computing environment is a major research challenge, particularly when these challenges must be addressed in real time without system downtime. For example, with respect to extensibility, whatever framework is produced for exascale data management should support a constantly evolving set of tools, analysis methods, datasets, and services, because in most areas, our understanding of our datasets is incomplete. Hence a hard-coded "toolbox" would fail to meet future user needs. Similarly, existing performance evaluation and techniques of high performance design and analysis must scale to match the new exascale computation challenge operating against extremely large datasets. Application areas commonly found in an internet-scale system (eg. Google, Facebook) will increasingly need to process more and more data while continually requiring sub-second response times that serve billions of people. (See Appendix for a list of qualities).

Automation: As the scale of both software and data increase, manual processes and techniques become impractical. For example, Google maintains a single code-base, which requires re-compilation several

times per minute. This would be unachievable without highly sophisticated and automatic continuous integration build processes, leveraging massive parallel computation capabilities. Other current examples of automation techniques include automatic fault identification mechanisms and traceability link recovery. Leveraging such developments to automate large parts of the software development lifecycle is a major research challenge.

Software engineering as a data science: Extending from the previous point, it is important to recognise that software engineering is itself a data science. Several trends illustrate this point:

- Mining of fault reports, documentation, and code for various salient features is becoming commonplace;
- Software pattern recognition techniques are being developed based on machine learning and data mining;
- Programmer productivity data is being mined for behavioural patterns and correlates.

Leveraging more systematic application of "big data" techniques within the software engineering lifecycle is a major research challenge.

Scaled techniques and tools: Tools and techniques appropriate to software engineering in a massively scaled world. These include mechanisms for dealing more effectively with millions of lines of code, geographically distributed computation units, long-lived systems, and modelling/testing based on huge data sets[46].

46 For example Google's entire code base is recompiled several times a minute, see GoogleTechTalks (2012) 'Building Software at Google Scale Tech Talk'. Video available at: http://www.youtube.com/watch?v=2qv3fcXW1mg. Accessed 18/8/12.

Heterogeneity: The heterogeneity of modern software environments, with multiple types of hardware, multiple flavours of software, disparate communication mechanisms and the need to integrate well with other systems (such as mechanical, electrical, biological and social), creates significant challenges. This implies the need for a broader systems-engineering approach that deals effectively with heterogeneity of computational systems and with multi-disciplinarity.

End-user software engineering: The need to place more power to develop software into the hands of end users, but with appropriate automated software engineering processes mitigating issues of quality and performance.

Software engineering as a human-centric discipline: Along with the emphasis on automation and scale it needs to be re-affirmed that software engineering is also strongly about people and their processes and that there are many important open research issues regarding the social aspects of software engineering. These will typically require multi-disciplinary approaches and much more involvement and partnership with social scientists than has been prevalent.

Engineering for evolution: Current approaches to software design and development result in software that is extremely brittle and difficult to maintain. We need to develop techniques that overcome these problems and result in software that is far more amenable to evolution over long time scales.

5.3.4 OPTIMISATION

Optimisation is a process of making effective use of resources and information, often in the form of a mathematical algorithm, often implemented in a physical system. Optimisation systems are ubiquitous in our society: they run supply chains, the electricity grid, airlines, and steel manufacturing, to name only a few examples. Optimisation has significantly evolved in the past two decades, achieving substantial advances in algorithmic speed and functions. Yet the computational challenges generated by a variety of fundamental societal problems are pushing the boundaries of optimisation. It is now necessary to design:

- Robust and scalable, mixed, nonlinear optimisation tools that can provide decision support over complex infrastructures such as electrical power, gas, and transportation infrastructures;
- Complex simulations of natural phenomena that can produce high accuracy in real time;
- Online, large scale, decision-support algorithms that can support decision-making in uncertain environments;
- Mechanisms to provide incentives for independent agents to achieve global objectives in complex, combinatorial environments in which they also pursue their own goals;
- Optimisation algorithms conditioned on massive amounts of data;
- A new generation of optimisation algorithms to exploit large-scale parallel computing to maintain the significant computational progress achieved over the past two decades.

Modelling and simulation are fundamental techniques required to overcome many of our grand challenges. This points to the need for close integration between computing and relevant sister disciplines such as mathematics, statistics and relevant pieces of application domains, such as formulae developed in the domain.

Identified technical challenges include:

- Model generation/validation: new modelling techniques and formalisms to deal with exascale. How are such things defined, validated, and made usable?
- Continuous simulation models.
- The ability to correlate among properties, e.g., spatial properties.
- Handling uncertainty and ambiguity.
- The need to be able to operate at multiple scales.
- The ability to take data captured at incommensurate scales and still operate successfully (e.g., by interpolation).

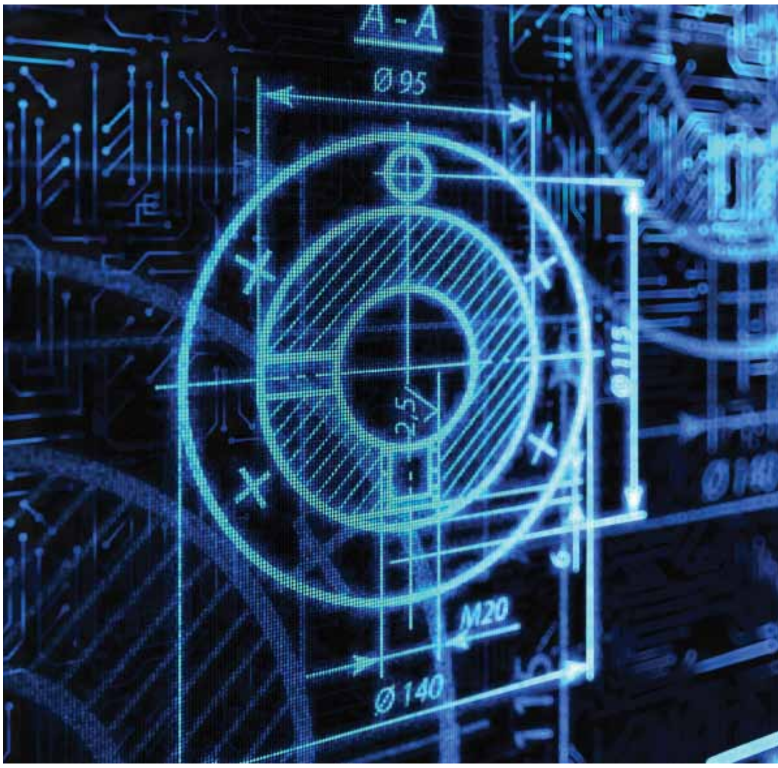

- The capacity to operate in real time.
- The ability to handle complex phenomena, including those from built and natural environments.
- The need to exploit latent massive parallelism.
- The capacity to model at 'real-world' scale. For example, simulation models for all of the separate physical domains including weather, ground water, power generation, habitation, etc.
- The ability to build effective models and simulate their effect on each other (e.g., weather and habitat).

5.4 INTERACTION

Working with massive-scale data and computation implies the need for new modalities for interaction: getting data into and out of the computation ecosystem, using computers to control massive families of sensors and actuators, and making what the computation ecosystem is doing (and the data it produces) comprehensible at a human scale.

5.4.1 HUMAN-COMPUTER INTERACTION

Traditional approaches to human-computer interaction will need to evolve to support computing at scale[47]. This evolution is likely to take many forms:

- *GUIs to gestures*: More fluid ways of interacting with computers, shifting away from mouse and keyboard to neural, gestural and multi-touch interfaces.
- *VDUs to immersed interactions*: Instead of viewing data through the 'porthole' of a traditional display, richer ways of viewing information are emerging, including 3D and immersive visualisations as well as virtual realities, augmented realities and time-series visualisations (animations).
- *Desk to world*: Increased use of ever more sophisticated mobile devices.
- *Shrink-wrapped to personalised*: The systems with which we interact will increasingly be tailored and customised – often automatically – to both our business and personal needs rather than everyone using the same, shrink-wrapped systems. Despite the increased personalisation, interoperability among users and systems will increase through development of appropriate standards.

47 Harper, R., Rodden, T., Rogers, Y. & Sellen, A. (eds) (2008) *Being Human: Human-Computer Interaction in the 21st Century*. Microsoft Research. Available at: http://research.microsoft.com/en-us/um/cambridge/projects/hci2020/downloads/BeingHuman_A3.pdf. Accessed 2/7/12.

- *Individual to collective*: Currently, collaboration and coordination capabilities must be 'bolted on' to computer systems that are inherently targeted at single users; we will see a shift to sharing, group awareness and coordination as a norm. This will allow both synchronous and asynchronous sharing and group work, and include 'stigmergic' capabilities to support group awareness.
- *Pixels to perception*: Data renderings tend to be unintelligent – essentially layers of pixels. Any 'smarts', such as annotations explaining the data, tend to be added by the developers. A shift to 'smart visualisations' will allow data-centric or evidence based explanations, resulting in interfaces that can make themselves intelligible to their users, and then reconfigure themselves appropriately as they are used.
- *Reactive to predictive*: Currently computer systems wait to respond to user actions; increasingly, however, they will understand the behaviour of their users and be able to respond predictively.
- *Ad-hoc to engineered*: All of the above require fundamental breakthroughs so that they can be treated as engineered services.

Each of these transformations represents a number of significant research challenges.

5.4.2 SENSOR TECHNOLOGIES

Sensors connect computing to the real world. Sensors are pervasive in embedded computing, from temperature sensors in an oven to the plethora of sensors now common in automobiles. The adoption of smartphones has put a highly sophisticated and powerful computing and sensing platform (an iPhone 4S has 10 integrated sensors) into the hands of hundreds of millions of people, with smartphone shipments expected to reach 1.7 billion by 2017.

The storage and processing of sophisticated sensor measurements now consumes enormous quantities of computing resources. From millions of video surveillance systems, through thousands of geological surveys, to astronomy's Square Kilometre Array, this data and computation is measured in petabytes per second.

Almost every industry has infused digital technology into their operation, including the roll-out of "smart" electricity grids and the use of precision GPS by farmers to optimise harvesting of crops. This explosion of digitisation is creating many opportunities and challenges. Estimates of the number of devices in the foreseeable future vary enormously, but they range from 50 to 100 billion devices through to more than a trillion. Whichever predictions turn out to be accurate, the volume of data will be huge.

While these sensors will obviously not all be connected to a single network, there will nevertheless be many scaling challenges related to the design and operation of networks connecting millions, or hundreds of millions, of sensors. Furthermore, the volume of data and the insights that people will want to make from this data, will require technology to aggregate, filter and correlate data from sensor networks in real time, at scales that are many orders of magnitude larger than exists today.

5.4.3 CONTROL SYSTEMS

Over the past 100 years, the control systems discipline has produced a substantial body of science that now forms the basis for a deep understanding of the way dynamic systems behave and how they can be affected to behave more desirably. Industrial processes, aeroplanes, automobiles, electrical power systems, irrigation systems, DVDs, container ships, unmanned aerial vehicles, industrial robots, ecological systems, biological systems and countless other commonplace examples are all cases of dynamic systems whose behaviour can now be regulated to some desired and optimal manner by control systems techniques.

The core problem for control systems is regulation of dynamic systems in the presence of uncertainty and unwanted disturbances. Control engineering has mastered this problem through fundamental advances in the understanding of adaptive feedback mechanisms and real-time modelling, estimation and optimisation. The problem studied by control scientists is the challenge of data-to-decision (or, equivalently, measurement-to-actuation) as it applies to dynamic systems. This is a core part of ICT and is set to play a leading role in addressing many of the great challenges discussed in this document.

The coming era, where millions of network-accessible sensors could be used to monitor various aspects of an extremely large distributed dynamic system, such as a river basin, provides exciting new research challenges for control engineering. One such challenge would be the real-time optimal management of giant, distributed dynamic systems, involving many millions of state variables. In such a case, measurement from various parts of the system and actuation points affecting various other parts of the system are communicated over data networks.

The central question of decentralised control is what measured information must be communicated to which particular actuation points in order to achieve robust, reliable, optimal performance of the entire system. This question remains unanswered. The integration of yet-to-be-developed advances in 'big data' science, network information theory and large scale decentralised dynamic modelling and optimisation appears to be required in order to address this important challenge.

Control engineering techniques will also play a critical role in handling the challenge of dealing with the vast amounts of data created by the massive sensor systems that will be implemented over the next few decades. The problem of deciding what data is needed in order to make a particular set of decisions to a specified level of optimality and confidence is essentially a control engineering question. The scheduling of sensors and associated adaptive measurement strategies in order to optimise a specified optimisation criterion is a simple example of this idea where the control actuation is the choice of sensor to be used.

Control engineering is concerned with the feedback interaction between dynamic systems in the presence of uncertainty and as such is well suited to addressing the challenge of developing human-actuated control systems that extend the capabilities of the human body, such as neural control of prostheses.

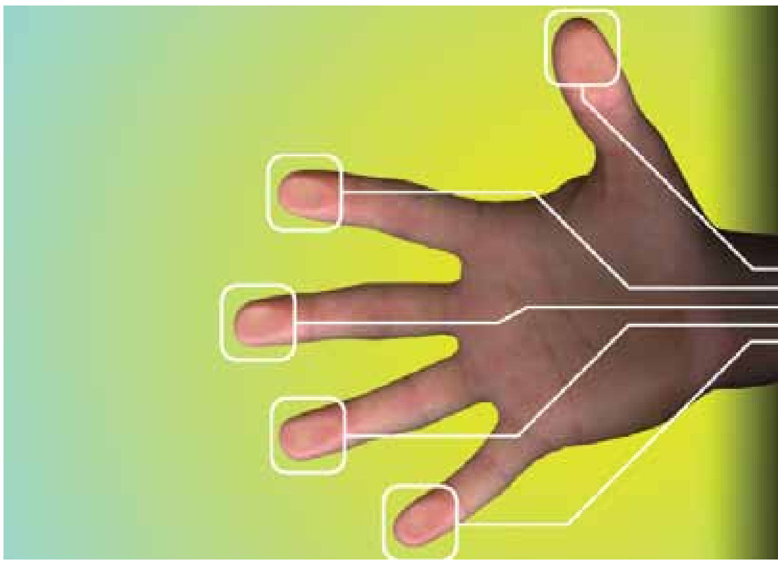

5.5 SYSTEMS

In computing, *systems software* is the conceptual layer that takes care of providing us with the illusion of the 'virtual machine' we use. Examples of the kinds of things this virtual machine provides include:

- Basic user interface capability: windows, mouse, cursors. Touch capabilities. The ability to drag, move, size, and gesture as part of the vocabulary of interface operations.
- Networking capability.
- The illusion that a computer is running multiple programs simultaneously.
- Access to the various devices (keyboards, cameras, mice, disk drives, scanners, printers, etc) connected directly to, or accessible to, the computer we are using.
- The ability to access resources that are not available locally, but give the impression that they are: remote data accessed from the cloud, running computations in the cloud and seeing their results, and so on.

The great success that has been achieved in building ever more complex and sophisticated modalities of computation and interaction via computing technologies is all due, in the end, to the advances that have been made in building (mostly) reliable systems software. Nonetheless, significant challenges remain, and Australia has demonstrated that it can make global contributions, for example making a major contribution to the development of wireless networking technologies.

48 Cisco's Visual Networking Index (VNI) on global mobile data traffic reported in February 2012 that the mobile broadband demand would increase by 18 times in the next five years where the usage doubled last year and is expected to double again in 2012. http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360_ns827_Networking_Solutions_White_Paper.html

5.5.1 NETWORKING

The ability to connect is paramount. All the data in the world is of no use if it cannot be accessed. Growth in data directly implies a necessary growth in networking capability.[48]

With a move to everything as a service and the "internet of things" (see 5.5.2 below), network availability and reliability become more essential. Networking needs to become a utility as reliable as reticulated electricity.

The technological challenges can be categorised under the following three headings:

5.5.1.1 THE CORE TECHNOLOGIES

The core networking technologies are developed by extensive human and technical ecosystems, mostly based on standardisation groups, which drive the evolution of fundamental core technologies. These ecosystems face multiple challenges, including the need for continually increasing bandwidth of wired and wireless networks.

Wireless networks are the greater challenge because of the finite natural resource of radio spectrum, the speed of uptake among users, and the fact that channels are intrinsically noisy and distance-sensitive.[49] These factors imply that it is essential to develop new methods for managing the radio spectrum, including dynamic allocation and cognitive radio. As wireless becomes interference-limited, dealing with the interference problem becomes a priority.

Over time, networks will become even more heterogeneous, especially as a solution to the wireless capacity problem[50 51]. This leads to many interoperability challenges.

49 Michael Kleeman, *Point of View: Wireless Point of Disconnect*, Global Information Industry Center, University of California, Dan Diego, October 2011.

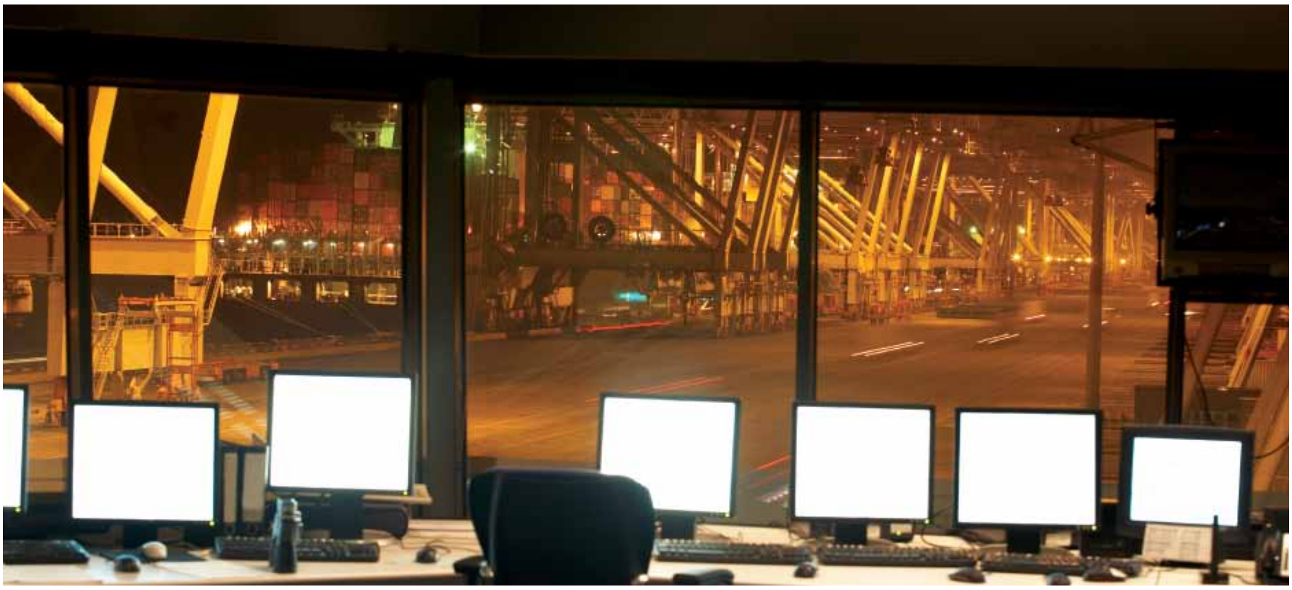

The heterogeneity and pervasiveness of networks, along with the demands for data, will mean that the old separation of data source and data sink will break down. Thus there are challenges in building distributed data centres *inside the network*.

New software engineering tools, languages and paradigms are also necessary in order to build reliable networked applications:

- Mobile wireless demands ever greater energy efficiency;
- Robustness requirements necessitate more powerful (transparent) proxies for both data and computation;
- Quantifiably reliable real-time guarantees are needed across heterogeneous networks.

5.5.1.2 MANAGEMENT

Management and understanding of network performance are important for maximising the operation of any sophisticated networked system. Some of the key challenges for this area include:

- Control and management of large-scale distributed systems.
- Managing privacy and security.
- Wireless spectrum resource allocation and control.
- Design for change ("adapting network architecture to meet future needs").
- Dealing with legacy systems -– networked infrastructure is enormously costly and ways need to be found to prolong its life even under scenarios of rapid growth and change.
- Application of advanced computational techniques such as machine learning to automate some aspects of network management.

5.5.1.3 INTERFACING

The human-computer interface (HCI) is currently undergoing a revolution, including some products that have deployed voice-recognition solutions. However, much is still to be achieved. In order to improve the use of HCI in networking we require better conceptual models (of both networks and interaction), better understanding of privacy and security from a user perspective, and the development of network-centric software engineering techniques, including languages, tools and formal methods.

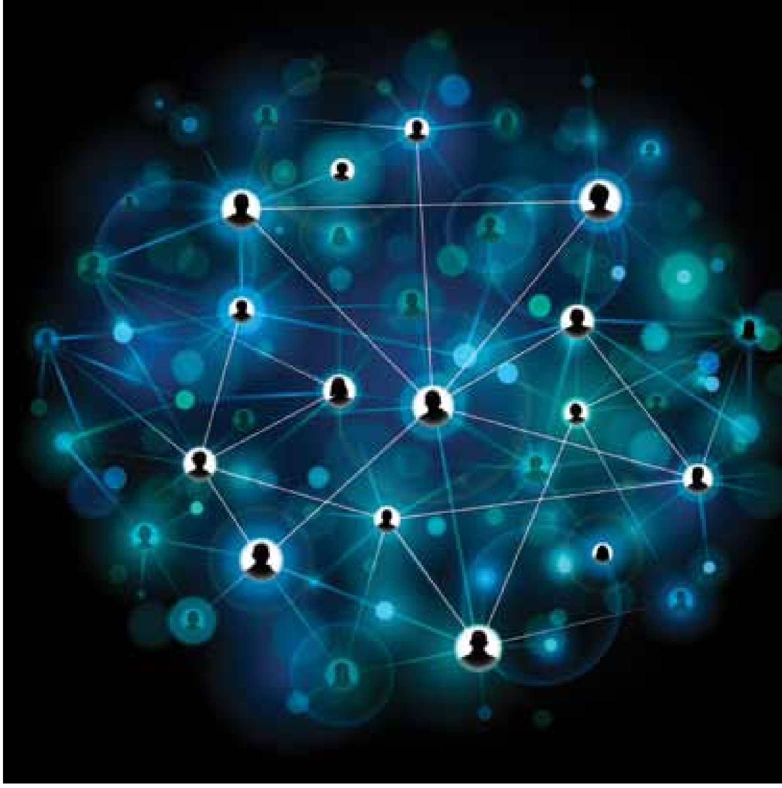

5.5.2 MIDDLEWARE AND OPERATING SYSTEMS

Traditionally, 'operating systems' have provided core system functionality, while 'middleware' has provided interoperability layers. However, these distinctions have become increasingly blurred, so in this document these areas are considered together.

Issues for systems software implied by the grand challenges outlined include:

- *Interoperability and connectivity across data sources*: Mechanisms that abstract the details of connecting to and working with data from many sources, so as to simplify the effort of developing highly scaled systems operating on massive data sets.
- *Adaptive distributed systems*: Mechanisms that simplify, abstract and hide the work required to get highly distributed computation units working together properly.
- *Middleware for highly flexible, rapidly evolving, open and service-based distributed enterprise systems*: Mechanisms and tools that make it possible to build, extend, maintain and evolve such systems.
- *Adaptors for easy integration into, and federation across, multiple independent cloud data and computation services*: See points on interoperability and adaptability above.

50 Landström, S., Furuskar, A., Johnsson, Falconetti, L. & Kroneatedt, F. (2011) 'Heterogeneous networks –increasing cellular capacity', *Ericsson Review.* Available at: http://www.ericsson.com/res/thecompany/docs/publications/ericsson_review/2011/heterogeneous_networks.pdf. Accessed 12/8/12.

51 See also Real Wireless Ltd (2012) T*echniques for increasing the capacity of wireless broadband networks: UK, 2012 – 2030*. Report prepared on behalf of Ofcom, March 2012. Available at: http://www.realwireless.biz/realwireless/wp-content/uploads/2011/11/Final-report-Capacity-Techniques-for-Wireless-Broadband-v1_15.zip. Accessed 12/8/12

The *internet of things* is a term for the internet with everything connected to it: smart meters, mobile phones, printers, computers, pacemakers, cars and so forth. This internet has literally trillions of devices connected to it, and significant systems work is required to make it possible for software developers to work with these data sources and sinks at an appropriately high level, that retains reliability. The internet of things is one of the major drivers of the scale challenge.

The 'internet of events' extends the concept of the internet of things, looking at event streams from devices rather than devices alone. Now, there are not only trillions of devices, but each of those is potentially the source of multiple event streams. The internet of events will mean that 'complex event processing' – a method of tracking, combining and analysing streams of data from multiple sources in order to infer more complex patterns -- becomes ubiquitous. Currently, support for complex event processing is embedded in enterprise resource planning systems or specialised software; creating common complex event processing capabilities as a systems software layer will allow these techniques to be used more generically. This is critical to being able to successfully scale event-based computations, such as the sensor-based scenarios in the grand challenges discussed.

The 'software qualities' in the Appendix are particularly important for systems software.

6

EDUCATION AND SKILLS

The emergence of computing has triggered a significant shift in the mix of skills that are required in a 21st century workforce. This section considers these shifts and likely implications.

6.1 BACKGROUND

At first glance, computing is like mathematics, physics or engineering: another advance in the development of 'scientific' thinking. Or, sometimes it is seen as simply a tool to support these disciplines (as well as business). From this perspective teaching computing should pose no special problems, or require any changes to the pedagogies and teaching approaches developed over centuries to train scientists, mathematicians and engineers. Unfortunately, this perception is incorrect, and any meaningful discussion of computing and its impact on the skills debate requires an understanding of why this is so. This understanding unfolds in three parts. First, an understanding of how problems are characterised is needed; then the different ways in which problems can be solved; and finally how these ideas are related to computing. With this machinery in hand, the impact on education can be considered.

6.1.1 PROBLEMS

It is useful to identify two broad categories of problems, 'wicked' and 'tame'.

Wicked problems are those that are difficult or impossible to solve because of incomplete, contradictory and changing requirements. The term was coined by Rittel and Weber[52], and characteristics include:

- Problems cannot be clearly defined until a solution has been found (each attempt to solve the problem changes the understanding of what the problem is itself).
- There are no stopping rules (the problem solver cannot know when the job is complete).
- There are not, in general, completeness or correctness characteristics for the problem.
- There are many stakeholders, each with their own perspective of the problem.
- Solutions are not 'true-or-false', but 'better-or-worse'.
- There may be many possible solutions, but they do not form an enumerable set.
- The problem solver has no right to be wrong (they are liable for the consequences of the actions they generate).

52 Rittel, H, and Webber, M. (1973) *Dilemmas in a General Theory of Planning.* Amsterdam: Policy Sciences, Vol. 4, Elsevier Scientific Publishing Company. [Reprinted in N. Cross (ed.), *Developments in Design Methodology* (1984) Chichester: J.Wiley & Sons

Tame problems, in contrast to wicked problems, can be characterised as follows:

- Problem and solution articulation are independent (the problem can be articulated without reference to any solution, and finding possible solutions doesn't change the problem).
- There are clear stopping rules (the problem solver knows when the job is complete).
- There are correctness and completeness criteria (the problem solver knows how well they have done).
- Solutions are therefore true-or-false.
- Solution spaces can be articulated.
- The problems and solutions can be articulated independent of any stakeholders.
- Solutions are 'independent' of the solver: due to the existence of stopping, correctness and completeness criteria, the solvers can prove the qualities of their solution and are therefore not liable for the solution (although they might be liable for the proof).

6.1.2 APPROACHES TO PROBLEM SOLVING

There are several approaches to understanding and solving problems, which here are referred to as thinking styles:

- *Analytic* thinking, also called scientific or reductionist thinking, refers to a style of problem solving in which larger problems are successively decomposed into independent, smaller ones that can then be solved or recursively decomposed. The solutions can then be composed to solve the overall problem. Analytic thinking is the dominant mode taught in western education systems.
- *Synthetic* thinking, also called design thinking, refers to a style of problem solving in which problems are solved bottom-up, by combining together ideas and solutions.
- *Computational* thinking refers to a style of problem solving by considering the solution to a problem as if it were a computer program, i.e. by articulating the data and algorithms necessary to provide a solution.
- *Systems* thinking is the process of understanding how things influence one another within a whole. Examples of systems thinking could include understanding how:
    - the various parts of an ecosystem depend on, or influence one another
    - people, structures, processes and technologies fit together to make an organisation
    - pieces of technology fit together to make, for example, a working Airbus A380; and how the aircraft fit into a larger air transportation system.

6.1.3 RELATING PROBLEM CHARACTERISTICS TO SOLUTION STYLES

Tame problems, and the associated methods for solving them, are at the heart of the scientific method: problems are understood and solved using a reductionist approach, usually focusing exclusively on the tame aspects of the problem, with wicked parts ignored or pushed to one side. Clearly this has been a hugely powerful and influential approach, and indeed one could argue that all of the advances of science and technology since the 17th century are due to the development of increasingly sophisticated and powerful methods and tools (such as mathematics) to deploy within the analytical/tame context.

As problems became steadily larger, the idea of systems thinking was developed to provide a superstructure: the problems one tackled were tame ones, solved using analytical approaches; and the larger-scale decomposition of problems and composition of solutions could be explained in a systems thinking context. The idea of tame/analytic problem solving, coupled with systems thinking, provide the foundation for the approaches used in engineering problem solving.

Independently, synthetic approaches to problem understanding and solving were developed. These took root in the creative/design disciplines, notably architecture; and also found a home in some of the social sciences, particularly sociology and ethnography. It turns out that synthetic approaches fit wicked problems far better than analytic approaches do: this is because, by their very nature, wicked problems do not lend themselves to mathematical reasoning or decomposition into independent sub-problems. Many of the problems designers face – for example in designing buildings – are wicked in nature: there's no one 'correct' building, the design process can potentially go on indefinitely, and so forth.

In the past few decades computing has become more ubiquitous, and it has become clear that a further way to think about problems is *computationally*: consider a problem or solution in terms of data flows and operations. This turns out to be useful regardless of whether the solution is implemented on computers or not. Initially those who knew computer programming techniques tended to use this kind of approach without thinking about it, but more recently it has become an approach in its own right[53].

As computing systems became a norm in organisations, they escaped from the drudgery of tame problems such as payroll calculations, stock management and account-keeping, and became deeply embedded in supporting all aspects of the behaviour of organisations. The important point here is that when one leaves behind the world of payrolls and general ledgers, and enters the world of business processes and systems, one also passes a threshold from tame problems to wicked ones. Business behaviours and problems are simply not reducible to formal models, reasoning, and so forth; nor can they be explained by the reductionist approach of analytic reasoning. Early attempts to deal with this through the 'normal' tame/analytic combination of problem characterisation and solving techniques were not successful[54], and a new combination of thinking styles emerged: systems thinking together with synthetic approaches to problem solving. The skills mix required by the emergence of social media and people-focused computer-based services have reinforced this trend. Today, the most sought-after business consultants are those that can combine all four thinking styles and are comfortable working with both wicked and tame problems.

Some general rules for the different problem types and thinking styles:

- Many problems have aspects of both wickedness and tameness. A key advantage of a broad portfolio of thinking styles, and the problem-solving skills that go with them, is that one can tackle a broad set of aspects of a large and complex problem, applying the relevant tools as required. Perhaps more importantly, one is not reduced to a thinking style of 'for a person with a hammer, everything is a nail', attempting to solve every problem from one perspective only.
- Analytic thinking and tame problems tend to go together
- Synthetic thinking and wicked problems tend to go together
- Computational thinking is a useful tool for all kinds of problems, often to clarify whether a problem is tame or wicked, and why.
- Computation is a critical implementation vehicle for all tame problems (and all the tameable parts of wicked ones)
- Systems thinking is a useful tool for understanding how the various components of a system should work together.

53 See, for example, the Centre for Computational Thinking at Carnegie-Mellon University. Available at: http://www.cs.cmu.edu/~CompThink/. Accessed 15/12/12.

54 For a detailed discussion of why this is so, see for example Suchman, L. (1987) Plans and Situation Actions : *The Problem of Human-Machine Communication.* New York: Cambridge University Press.

7

RECOMMENDATIONS

7.1 RECOMMENDATION ONE – TALENT DEVELOPMENT

Talent Development: government, schools and universities should, in collaboration with industry:

1. Increase the quality, diversity and size of the ICT talent pool by adapting for Australia and then broadly adopting the recommendations of the Royal Society's report Reboot or Restart for schools (whereby a distinction is made between generic ICT skills, and computing science)[55].
2. Ensure the development of computational and synthetic thinking skills in both secondary and tertiary education.
3. Re-energise the tertiary ICT education system and reduce attrition by refocusing curricula, changing pedagogy, suggesting that all ICT graduates engage with at least one other discipline in reasonable depth (at least to the level of a minor). Align with international best practice such as the ACM 2013 draft curriculum[56], and in particular adopt the principles there, especially
   i. Computer science curricula should be designed to provide students with the flexibility to work across many disciplines
   ii. Computer science curricula should be designed to prepare graduates for a variety of professions, attracting the full range of talent to the field.
   iii. Computer science curricula should be designed to prepare graduates to succeed in a rapidly changing field.
   iv. [The curriculum] should provide the greatest flexibility in organizing topics into courses and curricula.
   v. It is naturally tempting to associate each Knowledge Area with a course. We explicitly discourage this practice in general.
4. Adapt the curriculum to more strongly focus on the key areas described in the report.
5. In particular, ensure that the emerging data sciences are appropriately represented in the undergraduate curriculum.

55 The Royal Society (2012) *Shutdown or Restart*. Available at http://royalsociety.org/uploadedFiles/Royal_Society_Content/education/policy/computing-in-schools/2012-01-12-Computing-in-Schools.pdf

56 See the Joint Task Force on Computing Curricula Association for Computing Machinery IEEE Computer Society (2013) *Computer Science Curricula 2013*. Available at: http://ai.stanford.edu/users/sahami/CS2013/ironman-draft/cs2013-ironman-v1.0.pdf

7.2 RECOMMENDATION TWO – SUPPORT THE GROWTH OF COMPUTING SCIENCE

Continue to support the growth of computing science in Australia by:

1. Continuing to push for open data/open standards etc. This facilitates adoption and transformation of industries.
2. Continually investing in national computational infrastructure.
3. Providing ongoing support for a broad range of computing science research initiatives.
4. Encouraging multinational corporations that invest heavily in computing to establish bases in Australia.

7.3 RECOMMENDATION THREE – AN ONGOING ROLE FOR LEARNED ACADEMIES AND PEAK INSTITUTIONS

Leverage the capacities of Learned Academies and peak institutions to:

1. Champion the ongoing role of ICT as a transformative enabler.
2. Facilitate an ongoing dialogue between ICT researchers and their colleagues across academia and with government and industry.

7.4 RECOMMENDATION FOUR – STRENGTHEN INTERNATIONAL COLLABORATION

Research institutions should be supported to strengthen their international collaborations and focus additional research effort into the research areas highlighted in this report, namely:

1. Algorithms
2. Machine learning
3. Optimisation
4. Programming languages (including domain specific languages) and programming paradigms
5. Service orientation
6. Security and privacy
7. Software engineering
8. Interaction
9. Systems, especially big data infrastructure and embedded systems

This should be done in the context of embedded computing science, exemplified by the grand challenges in this report.

7.5 RECOMMENDATION FIVE – RECOGNIZE THE IMPERATIVE FOR TRANSFORMATION THAT FOLLOWS FROM THE EMERGING DIGITAL ECONOMY

Industry should recognize the imperative for transformation that follows from the emerging digital economy, and seek to ensure that they have the appropriate capacity for resilience and change. This will necessarily include:

1. Recognising the ways in which their business models and organisational culture can change.
2. Partnering with Australian research institutions on collaborative research projects in their particular domain.
3. Encouraging a wider role for computing scientists throughout organisations, particularly those qualified to facilitate the adoption of state-of-the-art technologies.
4. Encouraging technical staff to improve their value to their company by undertaking further computing science education.
5. Appointing senior staff with a computing science background as even businesses that are not primarily about computing are increasingly dependent on ICT for their future success.

8

APPENDIX

8.1 RELATION TO RESEARCH THEMES AND AREAS

	Table 1. Research Directions (Section 5) in relation to Grand Challenges (Section 4)								
	Relevance to Challenge								
	Research Area	Pervasive Health Data	Genome Sequencing and Personalised Medicine	The Augmented Human	Disaster Management	Monitoring and Managing the Environment	Managing massive environmental data sets	Competing Effectively in Today's Market	Social Media
Big Data	• Federated, Distributed Storage and Computation (Cloud) • Machine Learning • Security and privacy	• Developing new indexing, search and retrieval mechanisms for tera- or exa-scale multimedia data. • Latent feature discovery, factor understanding and application to clinical and service practice • Disease and complication progression models and appropriate predictive systems • Discovery of service and clinical factors associated with best practice • Privacy preserving mechanisms for real time risk assessment systems	• Mechanisms for storing and indexing massive and rapidly-changing personalised biological information • Dealing with genomes, mutations, SNPs, etc. as basic data types • Mechanisms for mining biological and health data repositories as a service, so these capabilities become widely available • Properly protecting user privacy • Securing massive, highly replicated and rapidly changing biological and health data.	• Manage needs around sensitive personal health, financial and social data • Large-scale data matching including evidence-based health monitoring and diagnosis using machine learning techniques • Efficient access to, and storage of information • Securely retain and serve information resources for a global community	• Developing algorithms to collect, aggregate and synthesize massive amount of real-time data on complex infrastructures • Learning complex models of human behaviours in life-critical situations	• Time scales range from sec/min/hour/day through historical records • Physical scales range from sub-millimetre sampling of water and soil samples, to cities and regions • Appropriate data curation and provenance management	• Developing new storage paradigms, massively parallel algorithms, indexing mechanisms for tera- or exa-scale data • Appropriate data curation and provenance management • Stochastic data assimilation, feature discovery, characterisation • Security and privacy (eg data anonymisation and embargoes)	• Machine learning as a service, to allow widespread uptake of trend and opportunity recognition • Federated data storage to increase uptake of advanced ICT in SMEs and evolution of the digital economy • Appropriate regimes for privacy, security and trust.	• Mechanisms for storing and indexing massive and rapidly-changing social media repositories • Dealing with graphs as basic data types • Mechanisms for mining these repositories as a service, so these capabilities become widely available • Properly protecting user privacy • Securing massive, highly replicated and rapidly changing social media data • Better tools for dealing with privacy • Parallel searches of massive social media data stores
Computation	• High Performance Computing • Programming Paradigms • Software Engineering • Optimisation, modelling and simulation	• Seamless integration into clinical and service delivery practice • Software engineering techniques for privacy preserving, real time systems for risk alerts • Optimisation techniques that incorporate group and context information	• Parallel searches and analyses of massive biological and health data stores • End-user programming languages and capabilities for research and clinical systems • Sequencing technologies • Wearable and implantable sensor, delivery or computing capabilities • Automating biological and health information generation • Customization of applications and data to various clinical settings or disease systems	• Engineering for reliability, scalability, adaptability • Architectural issues such as power/computation trade-offs; optimal locality of computation; accuracy and speed trade-offs • Semantic analysis and reasoning over heterogeneous data sources and repositories, especially social media data • Large-scale data-based optimization of resources. • Population-scale computations	• Developing optimisation and simulation algorithms for reasoning over complex infrastructures • Developing novel, scalable optimization techniques for dynamic decision- making under uncertainty • Developing real-time, incremental simulation algorithms for complex infrastructures and natural phenomena • Developing domain-specific languages and software architecture for assembling complex decision support systems compositionally • Exploiting massive parallelism in optimization and simulation algorithms	• Computing power to process weather, ocean, and surface conditions for validation with sensors • Simulation models for all of the separate physical domains including weather, ground water, power generation, habitation, etc. and their effects on each other • Software engineering techniques for massive-scale, massively-parallel, long-lived systems • Modeling and Reasoning about the environment in terms of scenario planning	• Computing power to process weather, ocean, and surface conditions for validation with sensors • Simulation models for all of the separate physical domains including weather, ground water, power generation, habitation, etc. and their effects on each other • Software engineering techniques for massive-scale, massively-parallel, long-lived systems • Modeling and Reasoning about the environment in terms of scenario planning	• Highly configurable enterprise systems and business process solutions for use by small enterprises • Software ecosystems that encourage rapidly evolving software • Appropriate engineering approaches • Languages, patterns and tools for highly flexible business process engineering	• End-user programming languages and capabilities for social media systems such as Facebook (for example to build apps and games)

	Table 1. Research Directions (Section 5) in relation to Grand Challenges (Section 4)								
	Relevance to Challenge								
	Research Area	Pervasive Health Data	Genome Sequencing and Personalised Medicine	The Augmented Human	Disaster Management	Monitoring and Managing the Environment	Managing massive environmental data sets	Competing Effectively in Today's Market	Social Media
Interaction	• Human-Computer Interaction • Sensor Technologies • Control Systems	• Making sense of huge data sets and models, through visualization and summarization • Human computer interfaces for appropriate clinical and service delivery • Incorporation of new pervasive sensor data into existing datasets for early intervention	• New HCI paradigms for clinical setting/physician and tech interaction	• Expanding the meaning of HCI from glass screen to true neural integration • Controlling prosthetic devices and systems of sensors in adverse environments with utmost reliability • Analysing visual sensor data through computer vision technologies	• Developing optimization algorithms for strategic and dynamic sensor placement; • Developing 3D visualization and simulation for complex decision-support processes	• Sensor management and synchronization, and validation with historical records • Making sense of huge data sets that come from changes in time and space. • Social networks will come into play in order for citizens to be able to more effectively participate in the trade-offs	• Sensor modelling • Sensor management and synchronization • Making sense of huge data sets and models, through visualization, immersion, summarization and drill-down • Human behaviour prediction	• User interface and visualization techniques integrated with machine learning and optimization systems, so that business users can derive maximum benefit from these services • User interface and visualization techniques to build, understand and evolve business processes. • Effective coordination and collaboration technologies for large-scale, rapidly-changing and globally-distributed organisations	• New HCI paradigms for social media and social interaction (noting that social media is human-to-human communication mediated by computers, rather than tradition HCI • Building 'social cues' such as stigmergic coordination capabilities into social media tools and apps. • Personalisation of applications and data
Systems	• Networking • Middleware and operating Systems	• Reliable interoperability and connectivity across data sources	• Massive, mobile, heterogeneous networking for biological data collection and healthcare applications and their carrier devices (such as smartphones, tablet) • Body area networking with very low power requirements • Middleware to support physician monitoring and feedback.	• Advances in large-scale heterogeneous mobile networking • Adaptive distributed systems		• Numerous independent networks will be fused, including fixed function for water and traffic, auto and mobile phones, and satellites • Sensor scale will require new networks for coupling and managing the shear number • Reliable interoperability and connectivity • Verifiable behaviour • Seamless integration among the many coupled system	• Meshes for sensor connectivity (scale, dynamic structure) • Teraspeed wireless networking to cope with data volumes • Reliable interoperability and connectivity • Verifiable behaviours • Seamless integration	• Middleware for highly flexible, rapidly evolving, open and service-based distributed enterprise systems • Adaptors for easy integration into, and federation across, multiple independent cloud data and computation services	• Massive, mobile, heterogeneous networking for social media applications and their carrier devices (such as smartphones, tablet) • Wearable and implantable computing capabilities • Body area networking with very lower power requirements • Social middleware to support social cues such as stigmergic coordination • Automating social media generation – the internet of things and internet of events

8.2 RESEARCH REQUIREMENTS FROM GRAND CHALLENGES

Domains		Research requirements from Grand Challenges
Big data	Federated, distributed storage	• Developing new storage paradigms, massively parallel algorithms, indexing, search and retrieval mechanisms for tera- or exa-scale data.
		• Developing algorithms to collect, aggregate and synthesise massive amounts of real-time data on complex infrastructures and at multiple physical and temporal scales
		• Federated data storage
		• Data curation and provenance management
		• Parallel searches of massive social media data stores
	Machine Learning	• Stochastic data assimilation, feature discovery, characterisation, factor understanding
		• Reification of machine learning capabilities to specific domains, for example biological, health, geological
		• Machine learning as a service, so that capabilities become widely adopted
		• Latent feature discovery, factor understanding and application domains
		• Large-scale data matching including evidence-based, domain-specific machine learning techniques, for example health monitoring and diagnosis
	Security and Privacy	• Domain-appropriate security and privacy (e.g. data anonymisation and embargoes, real-time risk assessment under privacy constraints)
		• Privacy preserving mechanisms for real time risk assessment systems
		• Securely retain and serve information resources for a global community
		• Securing massive, highly replicated and rapidly changing data sets, e.g. biological and health data.
		• Disease and complication progression models and appropriate predictive systems
		• Discovery of service and clinical factors associated with best practice
		• Learning complex models of human behaviours in life-critical situations
Computation	High performance computing	• Computing power to process huge geological data sets using massively parallel algorithms
		• Population-scale computations
		• Computing power to process weather, ocean, and surface conditions for validation with sensors
	Programming paradigms	• Software ecosystems that encourage rapidly evolving software
	Software engineering	• Software engineering techniques for massive-scale, massively-parallel, long-lived systems
		• Modelling and Reasoning about systems behaviour.
		• Software engineering techniques for privacy preserving, real time systems for risk alerts
		• Engineering for reliability, scalability, adaptability
		• Architectural issues such as power/computation trade-offs; optimal locality of computation; accuracy and speed trade-offs
		• Semantic analysis and reasoning over heterogeneous data sources and repositories, especially social media data
		• Software engineering techniques for massive-scale, massively-parallel, long-lived systems
		• Developing domain-specific languages and software architecture for assembling complex decision support systems compositionally
	Optimisation	• Optimisation techniques that incorporate group and context information
		• Developing optimisation and simulation algorithms for reasoning over complex infrastructures.
		• Developing novel, scalable optimization techniques for dynamic decision-making under uncertainty.
		• Large-scale data-based optimization of resources
	Simulation	• Simulation models for geological data: continuous, correlations among spatial properties, handle uncertainty and ambiguity, operate at multiple scales
		• Developing real-time, incremental simulation algorithms for complex infrastructures and natural phenomena
		• Exploiting massive parallelism in optimization and simulation algorithms.
		• Simulation models for all of the separate physical domains including weather, ground water, power generation, habitation, etc. and their effects on each other

Domains		Research requirements from Grand Challenges
Interaction	**Human-computer interaction**	• Making sense of huge data sets and models, through visualization, immersion, summarization and drill-down • Human behaviour prediction • Human computer interfaces for appropriate clinical and service delivery • Expanding the meaning of HCI from glass screen to true neural integration • New HCI paradigms for clinical setting/physician and tech interaction • Developing 3D visualisation and simulation for complex decision-support processes. • Making sense of huge data sets that come from changes in time and space. • User interface and visualisation techniques integrated with machine learning and optimisation systems, so that business users can derive maximum benefit from these services. • User interface and visualisation techniques to build, understand and evolve business processes. • Improved automation and tracking technologies that can be easily adopted by small enterprises. • Effective coordination and collaboration technologies for large-scale, rapidly-changing and globally-distributed organisations. • New HCI paradigms for social media and social interaction (noting that social media is human-to-human communication mediated by computers, rather than tradition HCI) • Social networks will come into play in order for citizens to be able to more effectively participate in the trade-offs • Building 'social cues' such as stigmergic coordination capabilities into social media tools and apps. • Personalisation of applications and data.
	Sensor technologies	• Sensor modelling • Sensor management and synchronisation • Analysing capability-specific sensor data through appropriate technologies (for example, using computer vision algorithms to analyse visual sensor data). • Developing techniques for strategic and dynamic sensor placement; • Sensor management and synchronisation, and validation with historical records • Incorporation of new pervasive sensor data into existing datasets for early intervention
	Control systems	• Controlling prosthetic devices and systems of sensors in adverse environments with utmost reliability • Wearable and implantable sensor, delivery or computing capabilities
Systems	**Networking**	• Meshes for sensor connectivity (scale, dynamic structure) • Teraspeed wireless networking to cope with data volumes • Reliable interoperability and connectivity • Advances in large-scale heterogeneous mobile networking • Massive, mobile, heterogeneous networking for biological data collection and healthcare applications and their carrier devices (such as smartphones, tablet) • Massive, mobile, heterogeneous networking for social media applications and their carrier devices (such as smartphones, tablet) • Body area networking with very lower power requirements • Numerous independent networks will be fused, including fixed function for water and traffic, auto and mobile phones, and satellites • Sensor scale will require new networks for coupling and managing the sheer number • The internet of things and internet of events
	Middleware and operating systems	• Seamless integration • Reliable interoperability and connectivity across data sources • Adaptive distributed systems • Middleware for highly flexible, rapidly evolving, open and service-based distributed enterprise systems • Adaptors for easy integration into, and federation across, multiple independent cloud data and computation services • Seamless integration among the many coupled system • Reliable interoperability and connectivity • Verifiable behaviour • Middleware to support physician monitoring and feedback. • Social middleware to support social cues such as stigmergic coordination.

8.3 'ILLITIES'

There is a family of non-functional properties required by many, if not all, large-scale computing systems, listed below. These remain research challenges for the computing and information sciences communities.

- *Composability*. The ability to compose together elements of a system so that they can be easily integrated and reused.
- *Reliability*. The ability to articulate (and ideally guarantee) under what circumstances the component will operate.
- *Standardisation*. Shifting from bespoke systems and components to components that have standardised interfaces and well-defined behaviours. Also standardising the entities into which software is integrated, so that both software and hardware can be easily reused.
- *Auditability.* The ability to record and reason about the actions that a component has performed.
- *Repeatability*. Presented with the same data and context, a component will behave in a predictable fashion.
- *Complexity*. The ability to build systems that are highly complex. This is more than just very big — complex systems exhibit properties such as emergent behaviour and self-organisation. Successful engineering and maintenance of very large scale, very long-lived systems, requires that advantage can be taken of these properties.
- *Heterogeneity*. A system has parts of many different kinds or types.
- *Verifiable behaviour.* The ability to reason about, and prove certain properties of a system.
- *Dynamic*. The system is subject to constant change, and can cope with this without violating any of the other properties listed here (or at least can identify when these violations are occurring).
- *Real time*. The system can sense and respond to stimuli (e.g., incoming data streams) in real time. The size of data streams can be immense, for example the proposed Square Kilometre Array is expected to produce over an exabyte of data each day.
- *Curation*. Mechanisms to curate very large scale data over very long timeframes.
- *Managing Service Quality*. Any ICT system should be able to operate at various service quality levels and cope with changes (for example, degradation in network speeds).
- *Modelling*. The ability to model and reason about systems behaviours.
- *Simulation*. The ability to simulate systems behaviours.
- *Incremental update*. With very large data sets it is infeasible to repeat computations each time the set changes. Ideally computations can proceed incrementally, updating themselves in response to data changes in time proportional to the size of the change rather than the size of the data set.
- *Multi-scale operation*. The data needed to use for modelling and simulation will be at multiple scales, and this needs to be accounted for seamlessly in computations and the components built.

8.4 RELEVANT REPORTS OF INTEREST

Altman, E. & Weisong, S. (Eds) *Network and Parallel Computing: 8th IFIP International Conference, NPC 2011, Changsha, China, October 21-23, 2011, Proceedings (Lecture Notes in ... Computer Science and General Issues)*. Heidelberg: Springer.

Andreessen, M. (2011) Why Software Is Eating the World. *The Wall Street Journal*. http://online.wsj.com/article/SB10001424053111903480904576512250915629460.html

Britton, J.W., Sawyer, B.C., Keith, A.C., Joseph Wang, C.-C., Freericks, J.K, Uys, H., Biercuk, M.J. & Bollinger, J.J. (2012) 'Engineered two-dimensional Ising interactions in a trapped-ionquantum simulator with hundreds of spins'. *Nature*, 484:489-492.

Cerone, A. & Pihlajasaari, P. (Eds) (2011) *Theoretical Aspects of Computing -- ICTAC 2011: 8th International Colloquium, Johannesburg, South Africa, August 31 -- September 2, 2011, Proceedings computer science and general issues)*. Heidelberg: Springer.

Chief Scientist of Australia (2012) *Health of Australian Science Report*. http://www.chiefscientist.gov.au/wp-content/uploads/Report-for-web.pdf

Committee on Information Technology Literacy, National Research Council (1999) *Being Fluent with Information Technology*. Washington, D.C.: The National Academies Press.

Gladwell, M. (2011) 'Creation Myth Xerox PARC, Apple, and the truth about innovation'. *New Yorker*. 16/05/2012. http://www.newyorker.com/reporting/2011/05/16/110516fa_fact_gladwell#ixzz236hxVMEE

Lampson, B. (2003) 'Computing Meets the Physical World'. *The Bridge*, 33(1):4-7.

Loth, S., Baumann, S., Lutz, C.P., Eigler, D.M. & Heinrich, A.J. (2012) 'Bistability in Atomic-Scale Antiferromagnets'. *Science*, 335:196-199.

National Research Council (2003) *Innovation in Information Technology*. Washington, D.C.: The National Academies Press.

Nielsen, M.A. (2011) *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge: Cambridge University Press.

Royal Society (2012) *Shut down or restart? The way forward for computing in UK schools*. http://royalsociety.org/uploadedFiles/Royal_Society_Content/education/policy/computing-in-schools/2012-01-12-Computing-in-Schools.pdf

U.S. Department of Education, Office of Educational Technology (2012) *Enhancing Teaching and Learning Through Educational Data Mining and Learning Analytics*.

8.5 PREVIOUS COMPUTER SCIENCE PREDICTIVE REPORTS

Amblee, R.S. (2011) *The Art of Looking into the Future: The Five Principles of Technological Evolution*. Gloture Books.

Australian Academy of Technological Sciences and Engineering (2010) *Cloud computing: opportunities and challenges for Australia*. http://www.atse.org.au/resource-centre/func-startdown/263/

Brennan, J.E. (2010) *The Future of Analytics: the trends, the implications, and who will lead*. Song Bird Hill Media.

Committee for Advancing Software-Intensive Systems Producibility; National Research Council (2010) *Critical Code: Software Producibility for Defense*. Washington, D.C.: The National Academies Press.

Committee for the Workshops on Computational Thinking; National Research Council (2010) *Report of a Workshop on The Scope and Nature of Computational Thinking*. Washington, D.C.: The National Academies Press.

Committee on Assessing the Impacts of Changes in the Information Technology R & D Ecosystem: Retaining Leadership in an Increasingly Global Environment; National Research Council (2009) *Assessing the Impacts of Changes in the Information Technology R&D Ecosystem: Retaining Leadership in an Increasingly Global Environment*. Washington, D.C.: The National Academies Press.

Committee on Modeling, Simulation, and Games; Standing Committee on Technology Insight--Gauge, Evaluate, and Review; National Research Council (2010) *The Rise of Games and High Performance Computing for Modeling and Simulation*. Washington, D.C.: The National Academies Press.

Committee on the Future of Supercomputing, National Research Council (2003) *The Future of Supercomputing: An Interim Report*. Washington, D.C.: The National Academies Press.

Committee to Assess the Scope and Direction of Computer Science and Technology, National Research Council (1992) *Computing the Future: A Broader Agenda for Computer Science and Engineering*. Washington, D.C.: The National Academies Press.

Council of Canadian Academies (2010) *Catalyzing Canada's Digital Economy*. http://www.scienceadvice.ca/en/assessments/other/digital-economy.aspx

D'Mello, S., Graesser, A., Schuller, B. & Martin, J-C. (Eds) (2011) *Affective Computing and Intelligent Interaction: Fourth International Conference, ACII 2011, Memphis, TN, USA, October 9-12, 2011, Proceedings, Part I Vision, Pattern Recognition, and Graphics*. Heidelberg: Springer.

Ford, M. (2009) *The Lights in the Tunnel: Automation, Accelerating Technology and the Economy of the Future*. Acculant Publishing.

Fuller, S.H. & Millett, L.I. (2011) *The Future of Computing Performance: Game Over or Next Level?*. Washington, D.C.: The National Academies Press.

Fuller, S.H. & Millett, L.I. (Eds) *The Future of Computing Performance: Game Over or Next Level?*. Washington, D.C.: The National Academies Press.

Graham, S.L., Snir, M. & Patterson, C.A. (Eds) (2004) *Getting Up to Speed: The Future of Supercomputing*. Washington, D.C.: The National Academies Press.

Harper, R., Rodden, T., Rogers, Y. & Sellen, A. (Eds) 2008, *Being Human: Human-Computer Interaction in the year* 2020. http://research.microsoft.com/en-us/um/cambridge/projects/hci2020/downloads/beinghuman_a3.pdf

Hoare, T. (2003) The Verifying Compiler: A Grand Challenge for Computing Research, *LNCS*, 2622:262-272.

Hopper, A. & Rice, A. (2008) 'Computing for the future of the planet'. *Philosophical Transactions of the Royal Society*, 366:3685-3697

IT Industry Innovation Council (2011) *Cloud Computing – Opportunities and Challenges*. http://www.innovation.gov.au/Industry/InformationandCommunicationsTechnologies/ITIIC/Documents/CloudComputingOpportunitiesandChallenges.pdf

Lazere, C. (2010) *Natural Computing: DNA, Quantum Bits, and the Future of Smart Machines*. New York: W.W. Norton & Co.

Lazowska, E. (2008) 'Envisioning the Future of Computing Research'. *Communication of the ACM*, 51(8):28-30.

Lohr, S. (2006) 'Computing, 2016: What Won't Be Possible?'. *New York Times*. 31/10/2006. http://www.nytimes.com/2006/10/31/science/31essa.html?ex=1319950800&en=74e41f37896acb9d&ei=5090&partner=rssuserland&emc=rss

Nagel, W.E. (2012) *High Performance Computing in Science and Engineering '11: Transactions of the High Performance Computing Center, Stuttgart (HLRS) 2011*. Heidelberg: Springer.

National Academies of Science (1995) *Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure*. Washington, D.C.: The National Academies Press.

National Academy of Engineering (2008) *Introduction to the Grand Challenges for Engineering*. http://www.engineeringchallenges.org/cms/8996/9221.aspx

Pato, J.N. & Millett, L.I. (Eds) (2010) *Biometric Recognition: Challenges and Opportunities*. Washington, D.C.: The National Academies Press.

President's Council of Advisors on Science and Technology (2010) *Designing A Digital Future: Federally Funded Research And Development In Networking And Information Technology - Report To The US President And Congress*. http://www.whitehouse.gov/sites/default/files/microsites/ostp/pcast-nitrd-report-2010.pdf

Schmidt, E. & Cohen, J. (2010) 'The Digital Disruption'. *Foreign Affairs*, Nov/Dec. http://www.foreignaffairs.com/articles/66781/eric-schmidt-and-jared-cohen/the-digital-disruption

Schwartz, P., Taylor, C. and Koselka, R. (2006) 'The future of computing. Quantum leap. Brain prosthetics. Telepathy. Punctual flights. A futurist's vision of where quantum computers will take us'. *Fortune magazine*. http://money.cnn.com/2006/07/26/magazines/fortune/futureoftech_quantum.fortune/index.htm

Thornton, J., Blake, A., Combes, F. & Palmer, T. (Eds) (2010) *Computational frontiers audio recordings*. http://royalsociety.org/events-computational-frontiers-audio-recordings/

US Department of Energy (2010) *Future Computing Trends 2009*. www.cs.berkeley.edu/~demmel/.../lecture21_Future_jwdhds10.ppt

FUTURE SCIENCE COMPUTER SCIENCE EXPERT GROUP

Dr Gordon Bell, Microsoft

Professor Hugh Durrant-Whyte FAA (Chair), NICTA

Professor Robin Evans FAA, University of Melbourne

Professor John Hosking, Australian National University

Professor Simon Kaplan, NICTA

Professor Ramamohanarao Kotagiri FAA, Unviersity of Melbourne

Associate Professor Maurice Pagnucco, University of New South Wales

Professor Pascal Van Hentenryck, NICTA and University of Melbourne

Professor Svetha Venkatesh, Deakin University

Mr Glenn Wightwick, IBM Research

Professor Robert Williamson FAA, Australian National University and NICTA