THE AUSTRALIAN COMPUTER SOCIETY

PRESENTS

ANNUAL DENNIS MOORE ORATION DINNER

THE UNIVERSITY CLUB OF WESTERN AUSTRALIA
WEDNESDAY 16 OCTOBER 2019

Sensor Networks and Distributed Intelligence

PRESENTED BY

Associate Professor
Rachel Cardell-Oliver

WE WISH TO THANK OUR PLATINUM PARTNER SPONSORS

THE UNIVERSITY OF WESTERN AUSTRALIA

Curtin University

GOLD PARTNER SPONSOR

Murdoch UNIVERSITY

WE ALSO WISH TO THANK THE ORGANISING COMMITTEE

Dr Robert CROSS FACS CP
Dr David COOK FACS CP (CHAIR)
Mr Arnold WONG BSc FACS CP IP3P GAICD
Prof Terry WOODINGS FACS

SUPPORTED BY

MYSECURITY MEDIA

ANNUAL
DENNIS MOORE ORATION DINNER

THE UNIVERSITY CLUB OF WESTERN AUSTRALIA
WEDNESDAY 16 OCTOBER 2019
PRESENTED BY - ASSOCIATE PROFESSOR **RACHEL CARDELL-OLIVER**

EVENT PROCEEDINGS

Opening and welcome by **Arnold Wong**, ACS National Treasurer.
Opening address by Parliamentary Secretary, **Mr Chris Tallentire** MLA,
on behalf of **Hon Dave Kelly** MLA, Minister for Water; Fisheries; Forestry;
Innovation and ICT; Science, representing the Premier.

Entree Served

Introduction to the 1962 Prize and 1962 Medal
and 2019 award finalists by **Dr Bob Cross.**
Professor Dennis Moore to present finalist certificates, the 1962 Prize and the 1962 Medal.

Main Course Served

Introduction of the 2019 Orator by **Dr David Cook**
Distinguished Oration delivered by **Associate Professor Rachel Cardell-Oliver**.
Concluding comments by **Arnold Wong**

Dessert Served

ANNUAL DENNIS MOORE ORATION

Since 2012, to commemorate fifty years of digital computing in Western Australia, the WA Branch of the ACS has invited a distinguished scholar and researcher with a connection to WA to present a lecture on the leading edge of an important and emerging area of information and computer technology.

Prof Jingbo Wang with **Prof Dennis Moore** at the 2018 Oration

- In 2012, this was delivered by Professor Andrew Rohl on the subject of supercomputers.
- In 2013, the oration was delivered by Professor Ian Reid and examined the subject of computer vision.
- In 2014, it was delivered by Professor Craig Valli on the subject of cyber and network security.
- In 2015, it was delivered by Professor Svetha Venkatesh on the question of “Where is that Data Utopia?”
- In 2016, it was delivered by Dr Adrian Boeing who addressed the question “Autonomous Mining: What’s Next?”.
- In 2017, Professor Matthew Bellgard, gave the oration on “Embracing Digital Disruption to Advance Clinical Research”.
- In 2018, Professor Jingbo Wang presented the Oration “Quantum Computation”

Taaqif Peck receiving his certificate and prize from **Prof Dennis Moore**

1962 PRIZE

From a suggestion of Dennis Moore (and with his strong support), 2012 also saw the setting up of an annual prize for the best graduating student in ICT from a WA university. Although the primary criteria are based on academic performance, the candidates are also interviewed for their ability to promote their ideas in computing and contribution so far.

Previous winners of the 1962 Prize are:

2012	**Kevin Adnan** (Curtin University)
2013	**Laurence Da Luz** (ECU)
2014	**Anthony Long** (Curtin University)
2015	**Michael Martis** (UWA)
2016	**Dalibor Borkovic** (Murdoch University)
2017	**Mark Shelton** (UWA)
2018	**Taaqif Peck** (Murdoch University)

The 1962 Prize finalists who graduated in 2018 for consideration of the award in 2019 are:

Dishain de Silva (ECU)
Lauren Gee (UWA)
Amit Goutamas (Murdoch University)
Josh Lloyd (ECU)
Morgan Smolder (Murdoch University)
Isaac Ward (UWA)
Jarryd Wimbridge (ECU)

DENNIS MOORE
AM MA (CANTAB) FACS

A young Dennis Moore demonstrates the IBM 1620 at UWA in 1967.
Photo by Wayne McKenzie, WA Newspapers.

Dennis Moore was born in NSW in 1937. He was educated on scholarships at The King’s School, Parramatta where he was captain and dux of the school, and at Queens’ College Cambridge where he graduated in 1958 in mathematics.

After a period with commerce and industry in computing and operations research in NSW, he pioneered computing in Western Australia, installing the first computer at UWA in 1962. He introduced WA’s first computing qualification — the DipNAAC — at UWA. In 1965, he was responsible for the purchase and installation of the DEC PDP-6. This was the world’s first commercial installation of a time-shared computer and Australia’s first high precision graphics device.

He was foundation president of the WA Computer Society, which later merged with the Australian Computer Society, becoming the first WA Branch Chairman. He was Director of the Western Australian Regional Computing Centre in the sixties and seventies. This provided computing services to CSIRO and State Government Departments as well as the University.

He was executive director of Government Computing for WA from 1978 to 1984. During this period he promoted the development of inter-departmental systems and was closely associated with the development of the WA Land Information System and the WA Technology Park. This was followed by a two year stint managing a computer company in Malaysia, including a consultancy to the Sarawak Government.

He then undertook research in RAN DATA, an encryption company which he had helped establish, and was appointed foundation Head of School of Computing at Curtin University of Technology in 1987. From 1998 to 2002 he was Director of Academic Planning at Curtin. From 1995 to 1999 he was Chair of the State Government’s Information Policy Council.

Dennis Moore was elected a Fellow of the Australian Computer Society in 1970 and was made a Member of the Order of Australia for services to Information Technology in 1997. He retired in 2002 and was made an Honorary Life Member of the ACS in 2014.

In 2012 with Professor Andrew Rohl.

In 2013 with Professor Ian Reid.

In 2014 with Professor Craig Valli

In 2015 with Professor Svetha Venkatesh.

In 2016 with Dr Adrian Boeing

In 2017 with Professor Matthew Bellgard

1962 AWARDS ABSTRACTS 2019

1962 PRIZE FINALISTS / DIGITAL DISRUPTOR ENTRANT

Lauren Gee graduated with a Bachelor of Philosophy with First Class Honours in Computer Science and Software Engineering from the University of Western Australia. For her Honours thesis, she developed a smart traffic control strategy for autonomous vehicles based on vehicle to vehicle communications and virtual traffic lights.

Lauren has twice interned at Google and is currently reading for a PhD in the area of AI. She works as a lab demonstrator at UWA and volunteers for the Girl's Programming Network and the UWA Programming Competition Society.

Amit Goutamas graduated with a Bachelor of Science majoring in Computer Science and Business Information Systems from Murdoch University. With the maximum possible GPA, he was twice awarded a Vice Chancellor's Commendation and, on graduation, received the University Medal. He was elected President of the Murdoch Dubai Student Computer Society for his final year.

Amit has twice won first prize in 'bot' challenges and, in a final year team, built a proof of concept for a Workload Management System for academic staff. He now works for Omnicom Group.

Josh Lloyd graduated with Bachelor of Computer Science from Edith Cowan University. He had an early taste of project management when he coordinated a team and controlled schedules and deadlines for the project. In his final year, he had a leadership role as a member of a team that developed a prototype game for children at high risk of developing meningococcal.

Josh developed a weather app, initially on Android and also converted to iOS that is in use today. He has a part-time position as an Apple 'Genius', helping the public in an Apple store.

Dishain de Silva graduated with a Bachelor of Science majoring in Cyber Security from Edith Cowan University. He won Cyber prizes from Bankwest/CBA. During his final year, he worked on vault mobile applications for a client - the requirement was to discover unique properties of vaults including reverse engineering some of the programs.

Dishain was also involved in the launch of a bike-share project in the Joondalup area. While undertaking his studies, he also volunteered in outreach programs as a member of the Christian Union at the university.

Morgan Smolder graduated valedictorian from Murdoch University with a Bachelor of Science majoring in Computer Science. He twice won a Vice Chancellor's Commendation and, on graduation, received the University Medal. He also won prizes in the database, computer programming, systems analysis & design, games design and programming units.

Morgan held paid roles as Peer Assisted Study Officer and Peer Academic Coach within the university. He also undertook a major paid project developing a successful website while he was studying. Morgan is working as a contract developer while building a business as a games developer.

Isaac Ward graduated with a Bachelor of Science with First Class Honours in Computer Science and Software Engineering from the University of Western Australia. His Honours thesis was Deep Neural Networks for Indoor Robotic Navigation and involved the fields of AI, robotics programming, data collection and 3D graphics to achieve image localisation.

Isaac received a studentship to an AI project at the International Centre of Radio Astronomy Research, a scholarship funding a data analysis project in China, wrote a chapter for a computer vision book and collaborated on a paper submitted to the International Journal of Computer Vision. He is currently working at two jobs and considering opportunities for overseas higher studies.

Jarryd Wimbridge graduated with a Bachelor of Computer Science with First Class Honours from Edith Cowan University. For his Honours thesis, he used a coevolutionary algorithm to develop a machine learning approach to automating air combat maneuvers, working with the Defence Science & Technology Group at the Department of Defence. He has also worked with people with special needs.

In his final year of study, Jarryd was in the team that won a merit in the undergraduate awards at WAITTA, iAwards and APICTA and won the ACS Digital Disruptors award for Skills Transformation. He has now founded a small software company.

DIGITAL DISRUPTOR ENTRANT

Kira Molloy is studying a Bachelor of Advanced Science at Curtin University. As she hasn't graduated, she is ineligible to be considered for the 1962 Prize, however she is a worthy nomination for the ACS Digital Disruptors. During her first year of study, she decided that she wanted some experience. She approached the Harry Perkins Institute of Medical Research as she would like to use her time for a good cause. Her work gave researchers and pathologists a much faster and more user-friendly way of visualising Optical Coherence Tomography scans. The project required self-study of many topics that she hadn't yet covered in her course.

This year, Kira won the Peter Fillery Best Undergraduate Project award at the WAITTA INCITE awards.

1962 MEDAL

The 1962 Medal recognises the most outstanding Doctoral Research in the field of Information Technology and Computer Science. The 1962 Medal is awarded in 2019 for the first time. It is open to any Western Australian student for research in a specific field and is subject to the research being the completion of a Doctoral Thesis. The award recognises a research project over a sustained period of time resulting in an original thesis offering a significant new contribution to knowledge in the field of computer science.

In 2019, there are three finalists for the 1962 Medal. The finalists are:

- Dr Qiuhong Ke (UWA)
- Dr Matthew Peacock (ECU)
- Dr Max Ward (UWA)

1962 MEDAL FINALISTS

Dr Qiuhong Ke graduated with a PhD from the University of Western Australia. The title of her thesis was **"Deep Learning for Action Recognition and Prediction"**. The thesis presents new methods for skeleton-based action recognition: dividing the entire human skeleton into five body parts to learn deep part-based features; learning spatial-temporal features; and a new generation method. The research can infer a future action at its early stage and can identify a person across different cameras and temporal periods.

Qiuhong has been involved with four patents, is a postdoctoral researcher with the Max Planck Institute for Informatics and will join the academic staff at the University of Melbourne.

Dr Matthew Peacock graduated with a PhD from Edith Cowan University. The title of his thesis was **"Anomaly Detection in BACnet/IP managed Building Automation Systems".** The research identified and evaluated a range of methods, particularly the Hidden Markov Models, which were accurate in classifying both known and unknown attacks in the evaluated BACnet/IP Building Automation Systems.

Matthew has had a book chapter published since completing his PhD and is a Technical Specialist working in the areas of machine learning, AI and protocol analysis at Sapien Cyber.

Dr Max Ward graduated with a PhD from the University of Western Australia. The title of his thesis was **"Algorithms for RNA Structure and Related Graphs".** The research showed that a focus on predicting Ribonucleic Acid (RNA) structures using advanced energy models could be achieved using polynomial time algorithms but ruled out further investigation in this direction. The thesis made contributions to advanced thermodynamic models, a linear time energy calculator for RNA free energy and an improvement to RNA structure prediction. The findings have been published in high impact journals. Max is a software engineer at Google in Australia.

Associate Professor
RACHEL CARDELL-OLIVER

Sensor Networks and Distributed Intelligence

Associate Professor
Rachel Cardell-Oliver

ABSTRACT

Distributed sensor networks comprise many small, independent, un-tethered computers that collaborate to acquire knowledge about their environment. Distributed sensor networks play a critical role in today's society where networked sensors, mobile phones, tablets, drones and micro-controllers are ubiquitous. Sensor networks are deployed in smart homes, health, agriculture, transport and logistics. This paper will review key research in the field of distributed sensor networks and its application to smart dust and smart cities.

1 Introduction

From the first civilisations humans have used measurement to understand and manage the world around us. Many breakthroughs in building, navigation, and commerce have been supported by the development of measurement instruments and techniques. For over 2 millennia most measurements were made manually, at a single point in space and at a single time. More recently innovations in digital sensors, low cost micro-controllers and radio transceivers have enabled us to measure the world at a high resolution of time and space. We can now build computer systems to measure, record and analyse complex spatio-temporal phenomena.

Distributed sensor networks is a field of Computer Science Systems research that aims to understand and build reliable, robust and cost-effective spatio-temporal measurement systems. The purpose of a distributed sensor network is to collect and analyse data. Distributed sensor networks are complex processing systems. They must be able to operate for long periods in hostile real-world environments with minimal support. In addition, in order to scale to large numbers of sensors, individual sensor nodes must be low cost, which means they will also have limited computational power and reliability.

What sorts of things can we measure with Distributed Sensor Networks? Everything! Smart homes, smart cities, precision agriculture, natural environments, transport and human activities, and even on or inside the human body. This paper will discuss key research areas for distributed sensor networks, and illustrate the development of the field including examples from West Australian research projects.

1.1 Key Research Ideas

What are the main ideas that have made the many advances in distributed sensor networks possible?

Theory: Distributed Agreement: In 1976 Edsger Dijkstra published A discipline of programming. In this book the task of developing computer programs was presented as a mathematical process for developing formally verifiable algorithms. This same (formal) approach was also used by distributed systems researchers. One of the basic modelling tools for distributed systems is finite state machines (FSMs). FSMs are used to describe the (distributed) state of the world and to model how that state can change over time as processes react with their environment or messages are exchanged.

In the 1980s two process models for distributed systems were developed. At the University of Oxford, Tony Hoare developed Communicating Sequential Processes (CSP). At the University of Edinburgh, Robin Milner (my PhD grandfather) developed the Calculus of Communicating Systems (CCS). These formalisms allowed researchers to explore the behaviour of distributed systems and problems of co-operation and agreement, livelock (getting stuck in a loop) and deadlock (getting stuck completely).

The problem of reaching agreement among remote processes is one of the most fundamental problems in distributed computing. In 1982 Lamport, Shostak and Pease published The Byzantine Generals Problem [13]. They considered a scenario in which distributed processors (generals) may malfunction or give conflicting information to different parts of a system. They proved that such a system can reach agreement using only oral messages if and only if more than two-thirds of the processors are “loyal generals”. So a single “traitor general” can confound two loyal generals. In 1985 the Impossibility of Distributed Consensus with One Faulty Process [6] placed an upper bound on what it is possible to achieve with distributed processes in an asynchronous environment. The Fischer, Lynch and Patterson theorem states that in an asynchronous network where messages may be delayed but not lost, there is no consensus algorithm that is guaranteed to terminate in every execution for all starting conditions, if at least one node may fail-stop. These results define the limits of what can be achieved by a network of communicating processes. In practice, distributed sensor networks have achieved many useful outcomes within this space.

Theory: Graphs and Set Cover: Another powerful way of thinking about distributed sensor networks is as a directed graph. The vertices of the graph are distributed sensor nodes. And the edges represent relationships between the nodes such as their ability to communicate or correlation in their sensor values. Given a directed graph representation, a common problem is to choose the smallest (or lowest cost) set of nodes that covers the whole network. This is called the set cover problem. The selected nodes can perform communication or sensing on behalf of the others thus saving energy for the whole network.

Experiments: Channel Measurements and Deployments: The theory of distributed agreement is very powerful for reasoning about distributed systems. But the real world that sensor networks inhabit turned out to be much messier than the theory allowed. An important development for distributed sensor networks was to improve our understanding of communication channels based on experimental measurements. Instead of modelling unreliable communication as a simple probability (e.g. a message has a 90% chance of getting through) researchers started measuring the behaviour of communication channels over long time intervals. We looked at properties such as auto-correlation: how likely are we to have a continuous block of lost messages or are these events independent? We looked at the type of corruption that occurs: are only few bits that damaged or whole packets? We looked at the ways in which reception varies in a spatial region.

Researchers have learned lessons from the many sensor networks that have been built and deployed in the past 30 years. We found that performance depends on many inter-dependent factors including environmental conditions, hardware and software, and communication protocols. Unfortunately as a com- munity we have been reluctant to publish results about failures and so we missed some opportunities to learn that way. But the current age of big data offers both interest and opportunity to publish more data from sensor network deployments.

COTS technologies: COTS stands for commercial off-the shelf technologies. Technology push has been a big driver for sensor networks research. Research and industry engineers have made many breakthroughs in the past 50 years that have resulted in near universal mobile phone network coverage, low cost computers such as the Raspberry Pi and Arduino, and software libraries that make those technologies accessible to all. Cloud services provide support to store data and share it via web pages. New software services make it easier to put together end to end systems for data collection, storage and analysis. The first generations of these systems were only accessible for experts in Computer Science. But later generations of hardware and software democratised sensor networks, so that now primary school students and hobbyists are also able to build their own applications. Flying drones and body sensors such as fit bits can be purchased and deployed by anyone. The science of these breakthroughs in sensing and communications are outside the scope of this paper. But the distributed systems science and technology that I will discuss has been enabled by that work.

People in the Loop: No systems approach is complete without considering how people and systems interact. People in the loop may be producers or consumers of sensor network data. The research methodologies for understanding people-in-the-loop come from disciplines such as Sociology, Psychology, Public Policy, Health and Law.

In this paper I will discuss two threads in the history of the development of distributed sensor networks, drawing on my own work and that of leading researchers. Distributed sensor networks is a wide field and so this is necessarily a personal and selective review.

2 Smart Dust

2.1 Beginnings

Smart Dust Proposal In 1998 Kris Pister of the University of California at Berkeley wrote a grant application. The proposal was for an idea he termed smart dust: millimetre scale sensing and communication platforms. He imagined distributed sensor networks comprising hundreds of thousands of dust motes and one or more interrogating transceivers. This was (and still is) a radical idea. He presented an application scenario for military applications in which dust motes would be scattered like seed from an unmanned plane or moving vehicle. These systems would gather intelligence for object detection or chemical attacks. This idea captured the imagination of the funders and the research community. Although 21 years later we are still some way from having thousands of senses scattered from an aircraft, the smart dust vision has led to many innovations.

Mica2 motes While Pister’s group was researching motes small enough to be scattered from an aircraft in their thousands, other research groups started thinking about how those motes might communicate and co-operate. Crossbow Technologies created the Mica2 processor and radio boards. These sold for around $200 each in the early 2000s. Berkley also developed the TinyOS operating system for programming the motes. It was also possible to buy sensing hats for connecting sensors to the Mica2. Although this system was expensive and actually didn’t work that well, Mica2s made it possible for research groups anywhere to buy a developer kit and start experimenting with the smart dust vision.

Pinjar soil moisture network Here at the University of WA we bought some Mica2 developer kits, wrote software and built field nodes. The first deployment was (appropriately for WA) in a Margaret River vineyard. David Glance was the investigator on this project and we deployed 12 motes to measure temperature and humidity in several vineyards. These motes collected data but they did not transmit their messages to a base station and so we had to visit the vineyards to download the data.

With another set of Mica2s and funding from the WA Water Corporation we built an application for measuring soil moisture. The goal was to understand how native Australian plants use water and the behaviour of WA sandy soil. This network was deployed at Pinjar just north of Perth. I spent a lot of time at Altronics and Bunnings and built the boxes and antennas as well as developing the software. Our Pinjar network, deployed in 2004, was quite innovative and the papers from this project are still being being cited today. A feature of our solution is its reactivity to the environment: when rain falls and soil moisture is changing rapidly, measurements are collected frequently whereas during dry periods, between rainfall, measurements are collected less often. Our field trial demonstrated the reactivity, robustness, and longevity of the network. It was able to collect about three months of data before the batteries ran out [3].

2.2 Wireless Communications

Aloha The Aloha protocol was developed in Hawaii in the 1950s. It is the simplest model for scheduling data transmissions over a shared channel: each node transmits data whenever it is ready. If the transmissions of two nodes overlap in time, then the messages will be corrupted, otherwise the message gets through. Although very simple, Aloha works well when the data rate for transmissions is low. It is still widely used today in many sensor network applications. However, as transmission frequency increases, there are more and more collisions and the

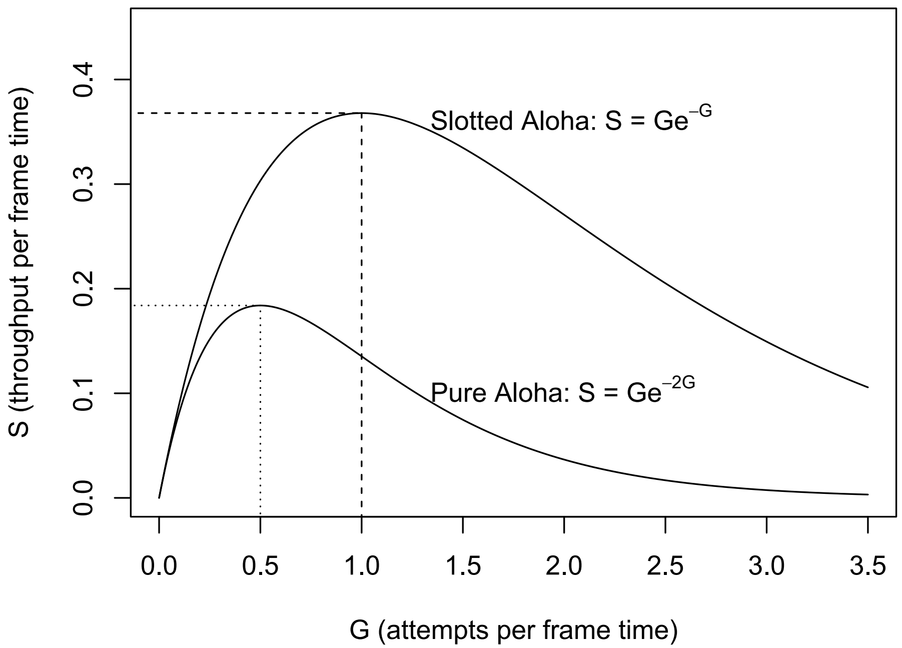

overall throughput of the system falls away sharply. The best utilisation of the channel that can be achieved with pure Aloha is ~18%.

FlexiTP: There are many protocols for organising networks of nodes to share their wireless medium more efficiently than Aloha. For example, listening before your transmit, backing off when the channel is busy, and transmitting only on fixed time boundaries [18]. The FlexiTP protocol [14], for example, uses timed slots for transmission, but allows nodes to build, modify or extend their communication schedules as their requirements for data reporting change.

Constructive Interference: An interesting recent development for sharing wireless channels turns the conventional view of wireless interference upside down [5]. Channel interference is made to be an advantage rather than a problem. Flooding in a distributed sensor network is a protocol for sharing a message with every node in the network. One originating node sends a message, and then each node that receives the message forwards it. Eventually the information should be propagated to the whole network. However, there are several disadvantages to this approach. Many transmissions will be corrupted because of wireless interference and many redundant transmissions will occur. Even with these redundancies, the flooding protocol does not guarantee that the message is received by all nodes. The Glossy protocol [5] published in 2011 harnesses interference. Glossy makes simultaneous transmissions of the same message so that messages interfere constructively. It can achieve flooding reliability of 99.99%, approaches the theoretical lower bound for latency (that is message delay) and provides network-wide time synchonization for free.

2.3 From Data to Information

Compression and Suppression Since we know that the capacity of shared channels is low, an intelligent approach for network protocols is to limit node transmissions wherever possible. That is, we need to take a do-more-with-less approach. One way of doing more with less is for sensor nodes to use temporal compression. The idea dates back to Claude Shannon's work on information theory:

IF Y CN RD THS ...

Much of the data communicated by distributed sensor networks contains redundancies. Our algorithm for Data-Aware Resource-Aware Lossless Compression [2] first investigates a sample of the data that node is collecting to identify where there are redundancies. Many sensor streams are smooth, with small differences between one reading and the next. The most common differences might be 0, 0.5, -0.5 etc. So a stream of sensor readings can be encoded using a base value and a list of the following differences. Knowing the first value we can reconstruct all the others using the differences. A special code, called a Huffman Code, can encode the most common symbols with a few bits. Longer symbols are used for rare symbols. In our experiments, this approach was able to reduce the size of a whole day's data to a few characters that could be sent in a single SMS message [2].

Co-operation: Many sensor network systems deployed today have multiple independent sensors communicating with a base station. For example, the Argo oceans observatory (www.argo.ucsd.edu) has over 3000 ocean diver sensors that independently submerge, take measurements, surface and report their data. But in many applications it is desirable or necessary for sensor nodes to co-operate. We now consider how to optimise the performance of distributed sensor network systems using co-operation amongst the nodes. An important reason for co- operating is that nodes share a wireless medium for communication. So they should co-operate in order to communicate and avoid the Aloha limit. Another important reason for co-operation is that nodes can save energy by co-operating on the task of data collection and interpretation.

Virtual Sensing: In the 1990s there were several approaches to reduce sensor network traffic by suppressing some nodes. Sensor networks have both temporal redundancies and spatial redundancies. In most of this work a sensor network is modelled as a directed graph in which the vertices are nodes and two nodes are connected by an edge if they are able to exchange messages or if one node can represent the other (that is, their data is correlated). In this setting, problems are expressed as variants of the set cover problem.

The 2009 paper Using mice to catch elephants showed how a small sample of sensors can be used to capture large scale events [7]. At the same time, many algorithms were developed based on dominating sets. The problem is to cover a 2D space with the fewest sensors where each sensors' space is a 2D disk [17]. If we find that two sensor spaces overlap, then one node can stop transmitting, so saving energy, while the other transmits. Sensors can take it in turns to sense and transmit in order to balance the energy use of all sensors. But one problem with this approach is that in many cases the disk model does not capture the true dependencies between sensors.

Our system, BuildSense [4], is a general framework for managing sensor sampling. The motivation was a project monitoring two rammed earth houses in Kalgoorlie. We had over 100 sensors per house and the cost of installation, maintenance and management was high. So we proposed a sense-learn-predict approach. First there is a training phase in which temporary sensors are deployed throughout a building (or other environment). Then we train prediction functions so that one sensor's readings can be predicted by another. Accurate predictor pairs are not always close by and may even be measuring different phenomena. After training, we remove the temporary sensors, leaving a subset of permanent sensors each of which predicts one, two or more virtual sensors. Overall, we were able to reduce the cost of a building sensor network by 60 to 80% by replacing physical sensors with virtual ones while still maintaining accuracy of 1.0 degree and fault-tolerance of at least 2 predictors per virtual sensor.

Device-Free Sensing Human-sensing is typically done with devices such as smartphones or wearables that are attached to a person. But these approaches have high costs of installation and maintenance which limits their application in some areas. Environmental monitoring applications, such as our rammed earth houses, also impose limits on the placement of sensors. One solution to these problems is device-free or non-intrusive sensing.

Dina Katabi's group at MIT design systems for sensing people [1]. Human bodies are around 60% water and wireless signals are attenuated by water. Their idea is to use wireless devices for non-intrusive detection of small human movements for health applications. Machine learning algorithms extract human signals from the readings of the wireless devices. This approach can be used to recognise signals as small as heart beats and sleep states without having to instrument the patient with intrusive contact sensors and it is able to sense people through walls and occlusions.

Cameras are an obvious choice for contactless sensing. As we learned from Ian Reid's Dennis Moore Oration vision researchers have made great progress in understanding images. However, camera images are too intrusive for many human sensing applications because they invade privacy. To address this problem, we developed a contactless occupancy sensor using a low-pixel thermal array [19, 20]. This system preserves the privacy of the occupants, but still delivers reasonably accurate estimates of the number of people in the space. For environmental applications, we are further developing the idea of non-intrusive sensing. For example, PhD student Omar Anwar is investigating whether non-intrusive sensing can be used to classify the health of honey-bees in their hives.

3 Smart Cities

The second part of this paper considers distributed sensor networks research from the perspective of the smart city. So far, we have considered how to gather data effectively and reliably. This section considers how to ensure that data is useful and some of the consequences when the sensors are people.

Ever since people have been living in cities, they have striven to improve city living, from Ancient Babylon's hanging gardens to the elegant boulevards of Haussman's Paris. Norman Bel Geddes' 1939 Futurama 20-year vision was characterised by automated highways and the American streamlined style. China's Shenzhen is a brand new liveable city, one of the fastest growing cities in the world with a permanent population of 13 million people working in its high tech industries.

Many cities are striving to become sustainable cities, water sensitive cities, liveable cities and smart cities. Perth has projects in each of these areas. But as Ben Green argues in his 2019 book The Smart Enough City taking an exclusively technical view of urban life (he calls this tech goggles) will lead to cities that appear smart but are rife with injustice and inequality:

> ***We shape technology by embedding values in its design and developing it to achieve particular outcomes ... We must be critical about the values embedded in these tools and who gets to choose them. Many technologies are designed to remedy social issues by enhancing efficiency . . . but that approach does not make them value- neutral [9].***

Unmanned flying drones are an interesting example of this phenomena. There are many compelling public-good applications for drones. For example drones could deliver medicines and supplies in poor countries where roads are often un-passable for much of the year. But there is currently little social licence or community support for large scale use of drones. This may be because of the perceived or actual possibilities for misuse: the theft of a drone or its cargo, use for bombing or spying, and potentially dangerous cluttering of airspace.

3.1 Actionable Insights

Distributed sensor networks gather data from distributed sensors and that data is analysed to provide actionable insights. The Co-operative Research Centre (CRC) for Water Sensitive Cities was launched in 2012 with an ambitious goal:

> ***The Water Sensitive City requires the transformation of urban water systems from a focus on water supply and wastewater disposal (the "taps and toilets" water utilities) to more complex, flexible systems. The output of the CRC will guide capital investments of more than $100 Billion by the Australian water sector and more than $550 Billion of private sector investment in urban development over the next 15 years.***

My work with the CRC has been on smart metering, a technology for real- time measurement and reporting of household water demand. The traditional method of measuring water use is quarterly readings made by a human meter reader who visits each property. Smart meters give water utilities continuous, real-time information about demand for a whole population. Smart metering solves many occupational health and safety (OHS) problems associated with the traditional meter reading approach. However, apart from the OHS bene- fits, early adopter water utilities did not really know what else to do with their new hourly per-household water demand data. Our work investigated several avenues. For example, empowering users by interpreting the different types of use and the contribution of each type to (maybe high) water bills. Or empowering utilities by identifying different types of customers and their patterns of demand.

The impact of our smart metering research was significant, but not in the way I expected. One of the data mining algorithms we developed was for automatically identifying regular, but non-periodic activities such as garden irrigation. This analysis supported queries such as: what is the total volume of irrigation demand? what is the spatial and temporal distribution of irrigation demand? and how does irrigation command compare with other demands? We identified and tested ways that this information could be used for day to day operations by irrigation inspectors or for behaviour change campaigns. However the most impactful output of this work turned out to be this visualisation. It shows demand in Litres per household per day separated into discretionary (likely irrigation) and non-discretionary (indoor) demand.

What we had identified was that discretionary use of water (by irrigation) was (1) a higher than expected proportion of all use; (2) concentrated in a peak around 6am; and (3) how much it depends on the season of the year. These results are important because water distribution infrastructure is usually over- engineered to allow for peaks in demand. Therefore if you can reduce peak demand then millions of dollars can be saved by not upgrading or building new infrastructure. Steve Atkinson of the Water Corporation took my visualisation and pasted it up in the office with a big arrow labelled can we shift this peak?. The impact was to change how people

think about demand. Since then, the Water Corporation has invested in data science as part of its core business by making a number of permanent data scientist appointments. They are building this type of thinking into the whole organisation's planning for the future. More widely, partly through the CRC's research, water businesses now look to ways in which distributed sensor networks can help in sweating the assets [21].

3.2 Public Good

For a large class of distributed sensor networks the sensors are or are carried by people. People-based sensors lead to two important considerations for sensor networks applications: people are mobile and people must be treated respectfully. The WA Planning and Transport Research Centre (PATREC) is a multi- disciplinary centre for applied research in transport and land use planning. My research with PATREC analyses sensor data from public transport smart card tickets with the goal of making sense of how cities are growing and the ways that people are using them. In this smart city network the sensors are people using their smart cards to travel on the public transport network using Perth's SmartRider system. Researchers need to consider two perspectives. First, the technical problems of how to analyse trip records to understand how the whole system is being used, and second how to do this without compromising individuals' privacy. Smart card data is extremely sensitive because it contains personal identifiable information. Simply replacing names with numbers is not sufficient to protect privacy. For example, Victoria released the MyKi trip records of Melbourne commuters as part of a data science competition. Recent news re- ported a University of Melbourne analysis that highlighted the privacy breaches of publishing this supposedly de-identified data [8]. Earlier studies have found that more than 90% of people can be identified by just 4 data points of where they have been and when they were there [9]. PATREC's policies protect the data we use from that type of breach, but the recent Victorian news is a timely reminder of the risks of research using individual's travel records.

One problem we researched is where people go and what they do there. In order to do this in the safe way we aggregated the raw data so that instead of reporting on individuals we analyse places. Instead of looking at the trips we looked at the gaps between trips because that gives you an indication of the activity that might be taking place. Analysing these stays showed the spatial regions that attract people and the mixture of activities that people do when they arrive. Our hubs and stays model is now part of an online planning support system for Perth, RailSmart, for modelling and optimisation of employment creation and public transport use. It is one of the tools used for planning around WA MetroNet's northern line railway extension [16].

4 Discussion

4.1 Achievements

This paper has reviewed research in distributed sensor networks using two narratives: smart dust and smart cities. We have seen how the vision of smart dust led to the development of low-cost and easy-to-use hardware and software that has enabled distributed sensor networks to be widely deployed in many application areas. Present day visions of smart cities build on these developments. But smart city applications also highlight the complexities of socio-technical systems at scale, and their important ethical, policy and public good implications.

Major achievements of distributed sensor networks research include understanding the real-world behaviour of sensor network systems; designing software, protocols and hardware that make them easy to design and deploy; and data science advances for collecting the right data and asking the right questions.

4.2 Future Research

Distributed sensor networks research has now left the hype cycle of the early 2000s. Some of the early ambitions have been realised and there has been a huge transfer of research ideas into mainstream products. However, there are still many challenging research problems to be addressed, particularly around scalable systems. I will conclude with three areas of research that I think will be important in the coming years.

Social licence: One of the biggest challenges we face is the responsible design and use of distributed intelligent systems. Research is needed on questions such as: Who owns the data being collected? How is privacy protected? and How can we make distributed sensor networks sufficiently secure? UWA's Technology and the Public Interest group aims to tackle difficult inter-

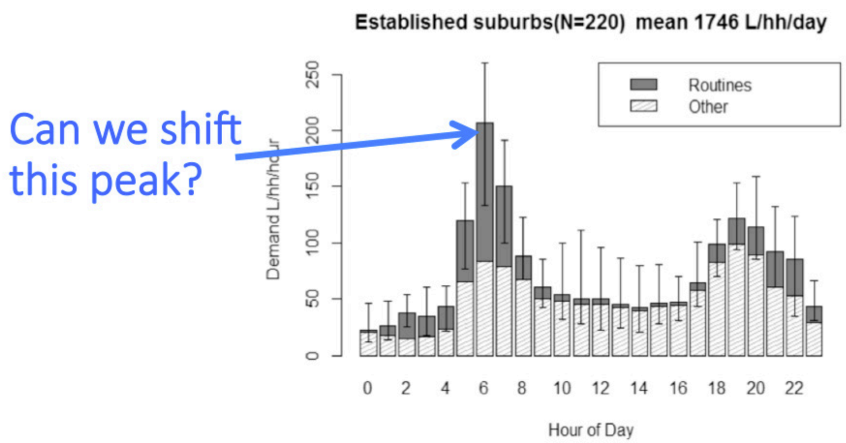

disciplinary research problems in this space and to provide subject-matter expertise on policy and adoption of new technologies including distributed sensor networks. Over 60 researchers from disciplines including law, computer science, urban design, environmental science, sociology and philosophy are involved in this initiative.

Novel sensing and communication modalities: As the demand for mobile internet grows the wireless spectrum is getting crowded. New communication mediums are being investigated.

For example, visible light can be used to transmit information using human-imperceptible flickering of the light from a single LED [12]. The idea requires a small microchip to be included in every potential illumination device. And there are 14 billion light bulbs in the world that could be harnessed for this purpose. On the sensor side, contactless sensors are important for monitoring humans or applications in dangerous or sensitive locations.

Energy Harvesting: Using batteries to power distributed sensor nodes has been a major barrier for scaling this technology. The maintenance costs of changing 10 batteries a year may be acceptable but for 1000s of sensors it is not feasible. One solution to this problem is to harvest energy from the environment [15]. But this energy harvesting scenario requires new programming models that are able to adapt to a scarce, highly variable, intermittent power supply [10]. As Hester and Sorber put it in a recent new directions paper, the Future of Sensing is Batteryless, Intermittent, and Awesome [11].

References

[1] Fadel Adib, Hongzi Mao, Zachary Kabelac, Dina Katabi, and Robert C. Miller. Smart Homes that Monitor Breathing and Heart Rate. Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15, pages 837–846, 2015.

[2] Rachel Cardell-Oliver, Stefan B¨ottcher, and Christof Hu¨bner. Data-Aware, Resource-Aware, Lossless Compression for Sensor Networks. In EWSN'13 Proceedings of the 10th European conference on Wireless Sensor Net- works, volume 7772 LNCS, pages 83–98. Springer-Verlag Berlin, Ghent, Belgium, 2013.

[3] Rachel Cardell-Oliver, Mark Kranz, Keith Smettem, and Kevin Mayer. A Reactive Soil Moisture Sensor Network: Design and Field Evaluation. International Journal of Distributed Sensor Networks, 1(2):149–162, Mar 2005.

[4] Rachel Cardell-Oliver and Chayan Sarkar. BuildSense: Accurate, Cost- aware, Fault-tolerant Monitoring with Minimal Sensor Infrastructure. ACM Transactions on Sensor Networks, 15(3):36:1–36:23, Aug 2019.

[5] Federico Ferrari, Marco Zimmerling, Lothar Thiele, and Olga Saukh. Effi- cient network flooding and time synchronization with Glossy. Proceedings of the 10th ACM/IEEE International Conference on Information Process- ing in Sensor Networks, pages 73–84, 2011.

[6] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. Journal of the ACM, 32(2):374–382, 1985.

[7] Sorabh Gandhi, Subhash Suri, and Emo Welzl. Catching elephants with mice. ACM Transactions on Sensor Networks, 6(1):1–27, 2010.

[8] Mary Gearin. 'Shocking' MyKi privacy breach for millions of users in data release, 2019. ABC News 2019/08/15

[9] Ben Green. The Smart Enough City. MIT Press, 2019.

[10] Josiah Hester and Jacob Sorber. Flicker: Rapid Prototyping for the Bat- teryless Internet-of-Things. ACM SenSys, pages 1–13, 2017.

[11] Josiah Hester and Jacob Sorber. New Directions: The Future of Sensing is Batteryless,

Intermittent, and Awesome. SenSys 2017 - Proceedings of the 15th ACM Conference on Embedded Networked Sensor Systems, 2017- Janua:1–6, 2017.

[12] Lennart Klaver and Marco Zuniga. Shine: A step towards distributed multi-hop visible light communication. Proceedings - 2015 IEEE 12th In- ternational Conference on Mobile Ad Hoc and Sensor Systems, MASS 2015, pages 235–243, 2015.

[13] Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine Gen- erals Problem. ACM Transactions on Programming Languages and Sys- tems, 4(3):382–401, 1982.

[14] W.L. Lee, A. Datta, and R. Cardell-Oliver. FlexiTP: A Flexible-Schedule- Based TDMA Protocol for Fault-Tolerant and Energy-Efficient Wireless Sensor Networks. IEEE Transactions on Parallel and Distributed Systems, 19(6):851–864, Jun 2008.

[15] Robert Margolies, Peter Kinget, Ioannis Kymissis, Gil Zussman, Maria Gorlatova, John Sarik, Gerald Stanje, Jianxun Zhu, Paul Miller, Marcin Szczodrak, Baradwaj Vigraham, and Luca Carloni. Energy-Harvesting Ac- tive Networked Tags (EnHANTs). ACM Transactions on Sensor Networks, 11(4):1–27, Nov 2015.

[16] PATREC. RailSmart Planning Support System railsmart.patrec.org, 2019.

[17] B. Pazand, A. Datta, and R. Cardell-Oliver. A review of redundancy elimination protocols for wireless sensor networks, LNCS Volume 6018. Springer- Verlag, Berlin, 2010.

[18] Lawrence G. Roberts and Lawrence G. ALOHA packet system with and without slots and capture. ACM SIGCOMM Computer Communication Review, 5(2):28–42, Apr 1975.

[19] A. Tyndall, R. Cardell-Oliver, and A. Keating. Occupancy Estimation Using a Low-Pixel Count Thermal Imager. IEEE Sensors Journal, 16(10), 2016.

[20] Liam Walmsley-Eyre and Rachel Cardell-Oliver. Hierarchical Classifica- tion of Low Resolution Thermal Images for Occupancy Estimation. In IEEE SenseApp 2017 IEEE 42nd Conference on Local Computer Networks Workshops (LCN Workshops), pages 9–17, Singapore, Oct 2017. IEEE.

[21] Zhiguo Yuan, Gustaf Olsson, Rachel Cardell-Oliver, Kim van Schagen, An- gela Marchi, Ana Deletic, Christian Urich, Wolfgang Rauch, Yanchen Liu, and Guangming Jiang. Sweating the assets – the role of instrumentation, control and automation in urban water systems. Water Research, 2019.

Speaker Biography

Associate Professor ***Rachel Cardell-Oliver*** studies distributed sensor networks, designing systems that integrate data measurement using environmental sensors; data collection with wireless communication systems; and data analysis using data mining techniques. Working with multi-disciplinary teams, she has researched environmental challenges such as understanding public transport use, reducing household water consumption, measuring water use by native Australian plants, and the performance of rammed earth for sustainable buildings in outback Australia. Rachel was introduced to distributed systems during her Honours and Masters degrees in Computer Science at UWA, leading to a PhD on formal methods for distributed systems at the University of Cambridge. She has worked at the University of Essex in the UK and the University of Western Australia, where she is now Head of the Department of Computer Science and Software Engineering.


acs
REIMAGINATION'19
THOUGHT LEADERS' SUMMIT
Amazing Humans, Epic Tech
4 December | MCEC, Melbourne