THE IMPACT MADE BY UK Academic Computer Science Research

CPHC
The Council of Professors and Heads of Computing

UK Computing Research Committee

bcs Academy of Computing

Introduction

This report highlights the impact made by UK academic Computer Science Research within the UK and worldwide over the period 2008 - 2013.

It is based on the 280 impact case studies submitted to the 2014 Research Excellence Framework (REF) sub panel 11 Computer Science and Informatics by Eighty Seven institutions. Over 80% of the case studies had some form of economic impact, including spin-out businesses created by universities, software tools and techniques developed by research projects which have benefited the efficiency of both computing practitioners in large and small organisations, as well as standard security and communication protocols in daily use by millions of users. The annual revenue generated from those spinouts which included figures in the case studies, was in excess of £170 million and they had nearly 1900 employees. The additional sales revenues attributed to the academic research in industries such as aerospace, telecommunications, computing and energy was about £400 million. Some of the impact has been in the form of public policy, for example in terms of identifying security risks, informing healthcare decisions or public debate on ethical issues. There has been considerable social impact in terms of new healthcare procedures and treatments as well as aids for disabled or elderly people.

The following figure indicates the main types of impact in the submitted case studies listed in Appendix A.

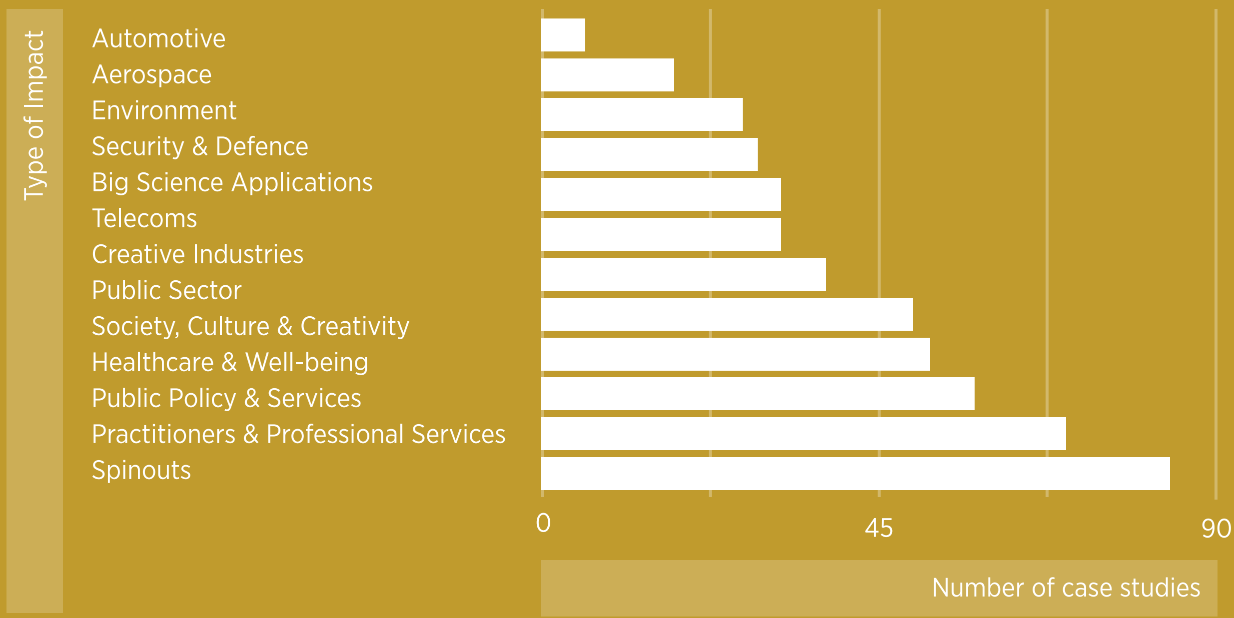

The report describes twenty of the submitted case studies and was jointly commissioned by UK Computing Research Committee (UKCRC), Council of Professors and Heads of Computing (CPHC) and British Computer Society (BCS) Academy.

The sub-panel assessors, which included eight people from industry and government appointed only to assess impact, recommended about fifty case studies as being potentially suitable for publicising UK academic Computer Science impact. These were not the fifty highest scoring case studies but were selected based on potential interest to the general public. An initial set of twenty case studies was selected from these to be written up in a form to make them more accessible to non-technical people. The selection criteria included ease of understanding of the technology underpinning the impact, potential interest by the public, examples from a wide range of different types of impact – both social and economic and showing that excellent impact can be generated from a range of universities with both large and small submissions to REF including post-92 universities, Russell Group and the other universities.

The 2014 REF was the first formal assessment of impact as part of the overall research assessment of UK academic institutions. The sub-panel assessors were very impressed by the extent to which UK academic research has had social and economic impact within the UK and often world-wide. The range of impact case studies included:

- Spin-out companies from universities, some of which had then been taken over by large international companies.
- Software tools and techniques either made available open-source or sometimes licensed to particular organisations with impact in automotive, aerospace, energy suppliers, media, gaming, healthcare, pharmaceutical, transport, retail as well as computing industry.
- Contributions to many different international standards e.g. telecommunication, web, compilers, security.
- Impact on government, healthcare and security policy as well as on public awareness about ethical and social issues.

The REF criteria stated that the research underpinning the impact must have taken place during the period 1 January 1993 to 31 December 2013 and be of a quality that is recognised internationally in terms of originality, significance and rigour (i.e. at least 2* quality, in terms of REF scoring), but the actual impact must have taken place during the period 1 January 2008 to 31 July 2013. The underpinning research described in case studies ranged from development of specific protocols, to formal methods used to reason about software design or to machine learning techniques. The underpinning research was often of the highest quality with publications in top conferences and journals.

The twenty case studies selected for this report were picked to reflect the range of those submitted and include spin out companies, software tools and techniques, commercialisation of open source software as well as a number of healthcare related applications and aids for people with disabilities. Some case studies indicate impact influencing public policy including issues relating to electronic payments, autonomous weapons systems and evaluation of health information systems.

The working group managing the report included Jon Crowcroft, David Duce, Ursula Martin, David Robertson and Morris Sloman. John Hill wrote the impact case study texts, in consultation with the relevant academics, and Naomi Atkinson was responsible for the design layout of the report.

Can a robot teach a child how to interact with others?

Meet KASPAR, the interactive playmate who is breaking down walls for children with autism.

1

Both parents and teachers agree: Eden has certainly opened up since she met her new friend.

It's 2011, and the four-year-old is growing up with autism. In the past, she would shy away from contact, and shrink when other girls held her hand.

"Affection always used to be on her terms", said her mother Claire. "If you didn't say, 'Can I have a hug?' and you just went and hugged her, then you might have got scratched or something. You always had to pre-warn her, whereas now it's much more spontaneous."

KASPAR

KASPAR is an unusual-looking friend, a simplified face without wrinkles or crinkles, tucked under a mop of wavy hair and a baseball cap. And it's true: KASPAR isn't like other children.

He's a robot, built by researchers at the University of Hertfordshire.

Some adults consider KASPAR a little creepy. He can blink, move his mouth, and gesture with his neck and arms, but his rubbery face lacks the intricate detail of a human face. It is precisely this simplicity that makes him so attractive to children with autism.

"Interacting with other people can be very challenging and frightening for autistic children", says Hertfordshire's Dr Kerstin Dautenhahn.

"The range of expressions can be very overwhelming and difficult to process. KASPAR has a very plain face, which is similar to a human face but simplified. By playing with KASPAR, children can spend time learning about human faces, and getting more used to human interaction."

KASPAR was born in 2005, a minimally-expressive robot prototype developed in just three months. However, the university's work stretches back to the late 1990s, when Dr Dautenhahn founded an umbrella project called Aurora to investigate how robotics could help with autism therapy. Her work was conducted alongside co-investigator Dr Ben Robins.

"When I started this work, I had many years of experience in social robotics. Autism started as a personal interest", she said. "When I did some research, I learned that many people with autism were much happier interacting with computers, as they found them safe and predictable compared to human beings.

"I started forming this idea. What if children were able to use a robot to practice social interaction and communication? Robots can be simplified, and can therefore be less stressful and easy to understand. We wanted to show them that social interaction didn't have to be frightening."

Once KASPAR was ready, the research group took him into schools, such as the TRACKS autism early years centre in Stevenage.

KASPAR is operated by a teacher or child using a remote control keypad. He is able to play a number of games that help children to explore human faces and interaction. In one game, KASPAR makes a gesture that the child has to mimic, such as putting his hands over his eyes or opening his arms wide.

Autistic children also struggle to determine how to touch others appropriately, often hurting other children by shoving them or pinching their nose. KASPAR reacts happily when he is touched in a friendly way, and shies away obviously when he is "hurt".

While KASPAR doesn't learn new behaviours himself, he can be adapted to play games that are tailored to a child's particular needs. Following suggestions from teachers, KASPAR was programmed to lead the children in song, with tunes like "If You're Happy and You Know It".

"You see them a little while later, loudly and enthusiastically singing the song and touching their nose. They're so happy", said Dr Dautenhahn.

"If you're a parent of a child with autism, one of the sad things is that it can be very difficult to play. If you can't play with your child, it can be very frustrating and upsetting. Can this robot help an autistic child to become closer to their family?"

The research group conducted long-term studies with 80 children between 2009 and the present day, and many parents reported that their children have noticeably softened after spending time with KASPAR.

The work done at the University of Hertfordshire also inspired former doctoral student Tamie Salter to set up her own company offering robotic devices for autistic children. Canada-based Que Innovations is currently trialing the QueBall, a ball-shaped robot which rolls around and reacts to the touch. It is designed for children on the lower-functioning end of the autism spectrum.

Other researchers have noted KASPAR's effect on children. At the University of Roehampton, Paul Dickerson's team undertakes conversation analysis to determine how people interact with each other. A study of autistic children who had interacted frequently with KASPAR uncovered some heartening results.

"If you analyse the interactions in detail, you can find instances in which children with autism display behaviour they're not meant to have", said Dr Dautenhahn.

"For example, children with autism are meant to be very bad with 'joint attention', the ability to have a common focus with someone on another object. But we have seen examples of children doing just that.

"Rather than focusing on what children with autism can't do - which is well documented - if you let them interact with a robot like KASPAR, you might see surprising things happening. We're trying to focus on the abilities of children, not the disabilities."

It has certainly proved beneficial for children such as Ronnie.

His mother Sheena said:

> "Before, he was always bashing the other children around. It's calmed him down a hell of a lot. I'd love to have KASPAR at home."

That dream may not be far off. In August 2012, a two-and-a-half year project began to re-design KASPAR for use by non-researchers, and robots with new hardware will be tested in schools and even houses. The aim is to enable some families to welcome KASPAR into their own homes.

Dr Dautenhahn said: “We’re still in the process of working out how this could happen, but we want to make KASPAR more widely available.

“Some people might be able to buy one, while others might rent one from certain suppliers. KASPAR may not be as expensive as some of those wonderful robots we see in the news, but we know that some families in need might not be able to afford him.

“We’d like to see as many children as possible benefit from KASPAR. We don’t want to see him become an expensive luxury object that only a few people can afford.”

Four of the UK's *five* biggest **dictionary publishers** now use a system built on technology developed at the University of Brighton to keep track of language and how it varies and evolves.

2

Language changes all the time, which makes the job of a dictionary particularly tricky.

Keeping track of new meanings and new words is a huge task for a human, which is why many major dictionary makers have switched to a computer system built on research undertaken at the University of Brighton.

The Sketch Engine is a system for people who want to get a data-driven idea of how words behave. It emerged from Brighton's research into computational lexicography, which has been a focus for the university since Roger Evans arrived in 1993. With EPSRC research funding, Evans brought in a research fellow called Adam Kilgarriff, who was the main instigator in creating a tool that has changed the way that dictionaries are compiled around the world.

"We started off thinking about how you build resources that you can use for computerised language processing systems", says Evans. "To do that, you need a lot of information about how words behave, so we started looking at dictionaries.

"What was going on there was more interesting and challenging than we expected. We ended up supporting the dictionary-making process rather than drawing from existing dictionaries. There was a definite progression from building computer systems to creating tools for production.

"We developed a tool that collects a lot of examples of language use and uses them to produce something that is effectively a draft dictionary of words and meanings that a lexicographer can work with."

To produce a dictionary, you need a large collection of language, which will tell you how a word is used, how often it appears and where. Dictionary-makers take advantage of a large repository of sentences and literature known as a Corpus, which contains millions of words. For example, the British National Corpus offers 100 million words drawn from literature and spoken conversation, providing a valuable idea of how contemporary language is used.

Previously, words would be picked out by poring through these corpora (the plural of 'corpus') and jotting them down on cards as they appear, then reviewing the context in which they appear. Definitions often would be based on these manually-collected examples, combined with the experience of the compiler.

Kilgarriff argued that the sense of a word needed to be based on its usage, and that the best way to approach the problem was to trawl through the entirety of the evidence and group word occurrences with similar meanings.

Adam Kilgarriff

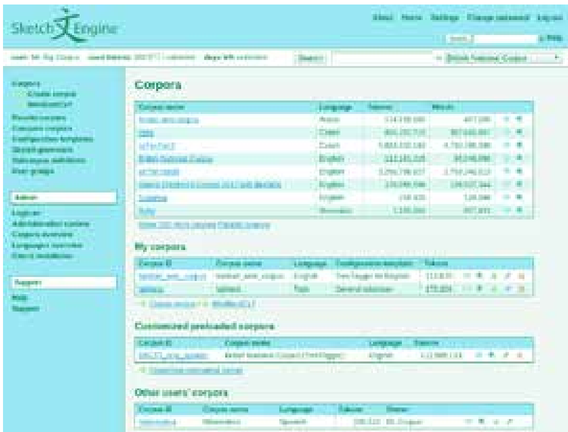

Corpus

"The thing I did that no one had done before, was to base the analysis of corpus data on grammar rather than just word-spotting", said Kilgarriff. "This took it into the realm of computational linguistics."

Working alongside research fellow David Tugwell with funding from EPSRC, Kilgarriff developed the idea of "word sketches", which present a word in the context of how it behaves grammatically and in conversation. This has changed the way in which lexicographers explore and analyse word usage.

goal *(noun)* **ukWaC freq = 168345** (107.5 per million)

object_of	58924	3.2	subject_of	25451	2.4	modifier	67879	1.6	modifies	11026	0.3
score	8390	11.28	score	903	8.59	ultimate	1911	9.27	scorer	389	9.39
achieve	9422	9.9	disallow	223	8.04	long-term	875	7.66	kick	634	8.86
concede	1421	9.39	concede	204	7.53	league	638	7.38	tally	129	7.9
accomplish	585	7.97	gape	76	6.5	winning	401	7.33	keeper	204	7.31
reach	1924	7.66	come	1316	5.44	primary	993	7.24	scramble	50	6.75
net	337	7.42	kick	76	5.44	second	2000	7.19	drought	78	6.65
pursue	648	7.41	rule	61	5.24	common	1529	7.17	difference	676	6.28
attain	400	7.35	orientate	34	5.06	strategic	645	7.1	cushion	53	6.26
grab	406	7.34	arrive	90	4.43	realistic	422	7.05	lead	267	6.24
set	2413	7.01	cap	20	4.38	achievable	290	6.97	setting	405	6.14
pull	501	6.88	beat	53	4.31	stated	259	6.8	kicker	25	6.04
disallow	190	6.67	direct	53	4.22	score	611	6.75	post	482	5.91

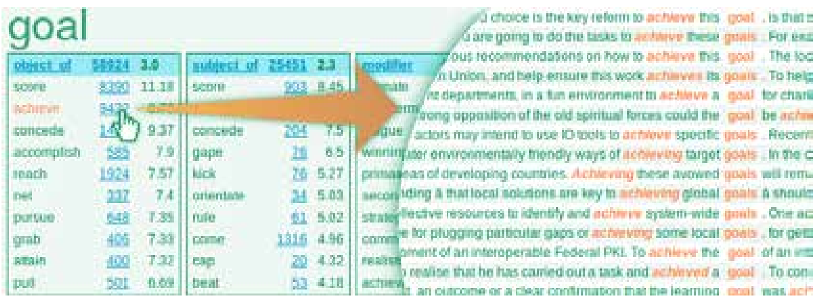

Kilgarriff says: "The first version of word sketching started as a small research project for Macmillan, who had been in discussions on how to improve their corpus use. It got very good reviews as it sped up their processing, made it more objective and helped them to include things that other dictionaries might have missed."

This research was at the heart of spin-out company Lexical Computing, which Kilgarriff formed in 2003. Its flagship technology - The Sketch Engine - enabled commercial lexicographers and dictionary makers to quickly analyse large amounts of evolving examples of language. It was a breakthrough that was particularly interesting to dictionary makers who were just getting to grips with digital publishing.

Macmillan had used the pre-Sketch-Engine word sketches, but Oxford University Press was the first publisher to adopt the Sketch Engine in its entirety in 2004. It has been using it ever since.

Oxford University Press say it has helped them to create a detailed profile of a word within seconds, hugely bolstering their in-house research. It has also been adopted by Collins and Cambridge University Press in the UK, as well as commercial publishers worldwide, such as Shogakukan in Japan, Le Robert in France, and National Language Institutes in Bulgaria, the Czech Republic, Estonia, Ireland, the Netherlands, Slovakia and Slovenia.

Lexical Computing's tool allows users to get information on between 30 million and 15 billion words in a variety of languages. That has potential for other organisations as well.

For example, the company sees huge potential in using the approach to help people get to grips with collocations in the English language.

"If your native language is not English, how would you know that we make mistakes but take care, have sex but make love?" says Kilgarriff.

Every year, Kilgarriff and his colleagues deliver an intensive, week-long training course called Lexicom, which attracts publishers, universities and government agencies. Organisations across the world have noted its usefulness for professional training and language teaching.

Artist David Levine and PhD student Alix Rule used the Sketch Engine to pick through thousands of emails from art-world email service e-flux, identifying which words and phrases came up more in "International Art English" than in everyday spoken English. They noticed a tendency towards longer sentences, and a significant upswing in uses of words and phrases such as "multitude", "the real world" and "the void".

It is also being used to analyse the language that children use, in collaboration with Oxford University Press and the BBC.

"There are lots of different companies using it in slightly different ways", says Kilgarriff. "Some hi-tech companies see it as an effective tool to help them find a good name for their business. Others want it for translation, language teaching or linguistics research.

"Dictionary companies are our biggest customers, but they represent less than half of company income."

The techniques and technology explored at Brighton have already changed the nature of lexicographical study and resulted in a successful company. So what's next?

"We've talked about computers that are capable of advanced language processing, but the way that we use language is hugely adaptive and context dependent", says Evans. "We can get computers to deal with one style and genre but that's not the way that people talk. We want them to be able to switch between styles in the way that people do, to answer one question in a technical way and then to talk about the weather or the football in the same conversation.

"If we can crack the single-style problem, which hasn't been done yet, we can create systems with the adaptability and context sensitivity to unlock a whole range of more natural language applications."

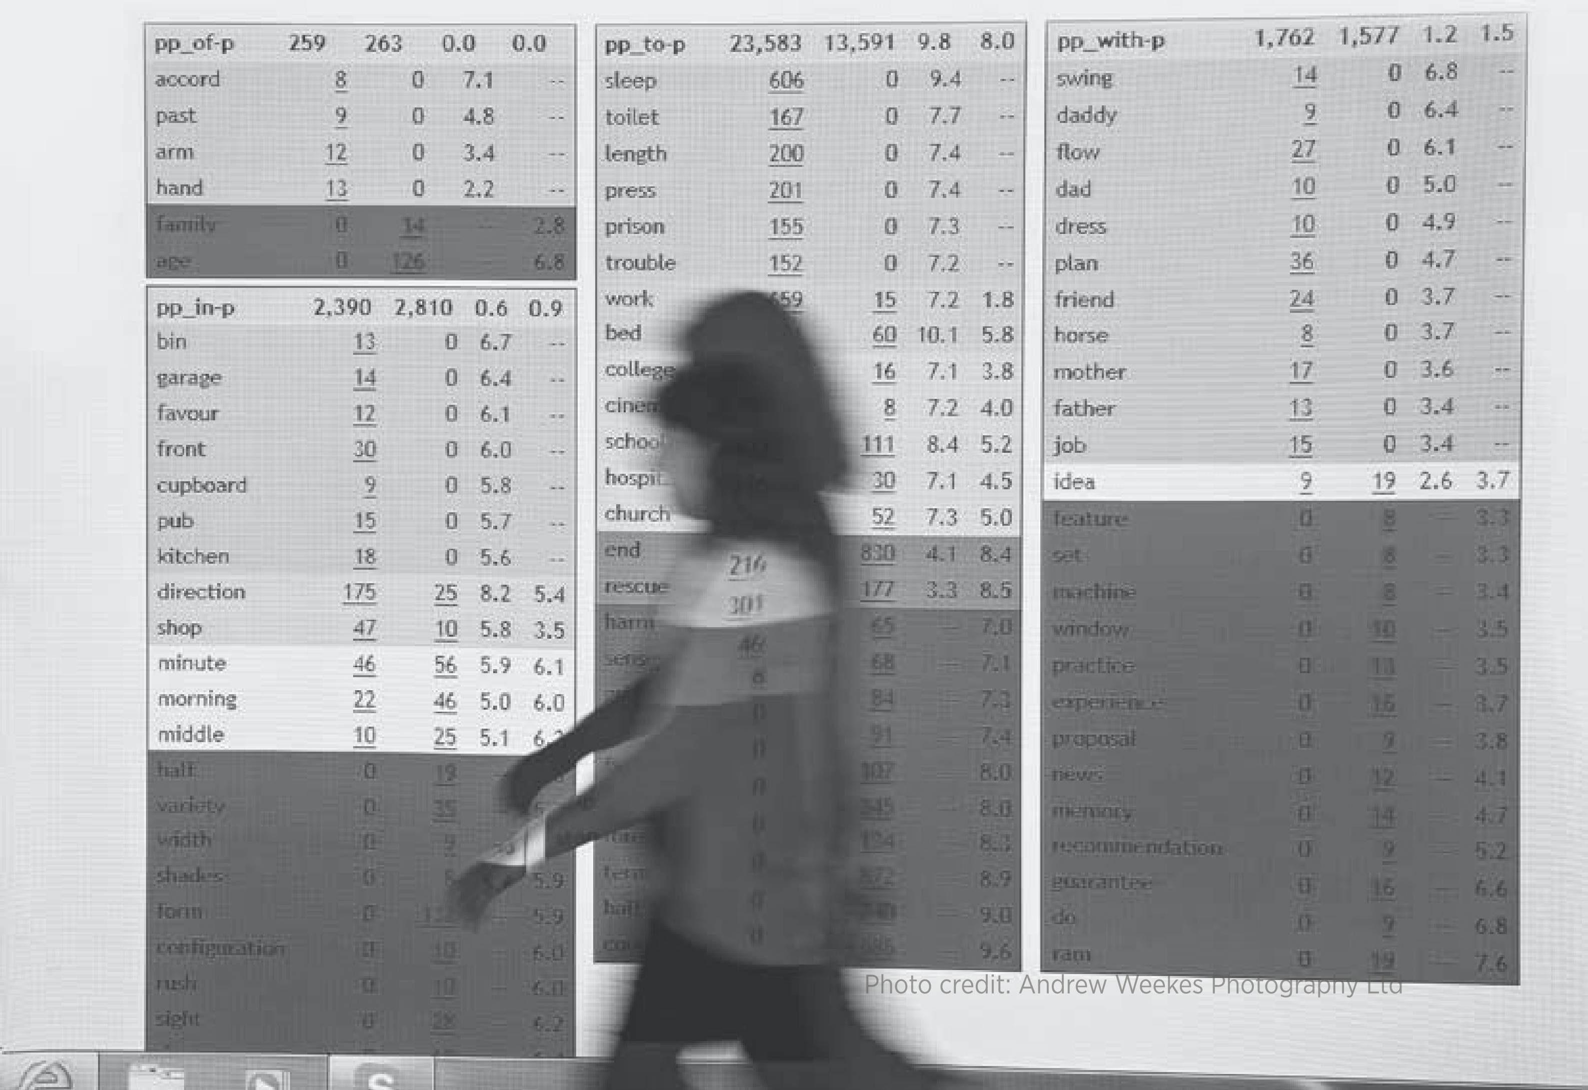

Photo credit: Andrew Weekes Photography Ltd

When you tell a vehicle to do something, timing is all important. That means you need on-board systems that can deal with a lot of processing and communications traffic efficiently and reliably without ever missing deadlines. Since the mid-nineties, the Department of Computer Science at the University of York has played a big role in making sure that happens.

Imagine you're in your car. You approach the traffic lights, and put your foot on the brake.

At this point, your brake lights go on and you slow to a stop. But there's a lot more going on than that.

Pressing the brake closes a switch, which is detected by an Electronic Control Unit (ECU). The ECU passes a message over a Controller Area Network to another ECU at the back of the car. The message is then decoded, and causes the brake light to turn on. All of this happens in a fraction of a second.

Now imagine similar events and responses happening hundreds of times a second throughout your car, controlling everything from gear changes to fuel injection and ignition timing. Oh, and there's a time limit for pretty much every one of them.

Over the last 25 years, under the guidance of Alan Burns and Andy Wellings, the Real-Time Systems Research Group at the University of York has been finding ways to make this system more efficient, more reliable and less memory-intensive. Its areas of research have sparked three separate start-up companies, and been used by automotive and aerospace companies worldwide.

Each strand of research represented a different challenge for the team, but helped to make vehicles work more efficiently and reliably as a result.

Let us begin with Controller Area Network (CAN) mentioned earlier. Before the 1990s, cars were fitted with point-to-point wiring, which was expensive, difficult to maintain, and left cars clogged with a heavy jumble of wires. For example, connections to a door in a high-end car required more than 50 wires.

As vehicles became more complex, CAN was devised to handle communication between Electronic Control Units (ECU) digitally, reducing the need for wires in that door from fifty to just four.

With CAN, messages are broadcast to all of the ECUs on the network; however, the number of ECUs keeps on rising. Whereas a mid-nineties car would have 5 to 10 ECUs, modern cars may have as many as 70. So now message traffic is an issue, and every message has an important deadline to meet.

In 1994, Ken Tindell, Alan Burns, and Andy Wellings introduced a different way to think about this potential log-jam. Using an approach known as schedulability analysis, they enabled designers to work out offline if all of the messages on the network could meet their deadlines. They did this by computing the longest possible time that each message could take to go from being queued by one

Alan Burns

Andy Wellings

ECU to being received by another.

After publishing their research, they were approached by Volvo Car Corporation. Two members of the group, Ken Tindell and Rob Davis, formed a start-up company Northern Real-Time Technologies to capitalise on the work.

Along with a Swedish company, they developed what became known as Volcano technology.

“Before this analysis was done, there wasn’t really any way to test how long a message could sit there in the queues”, says Rob Davis. “We provided a scientific method to analyse this offline and find the maximum time each message could take to be sent over the network.

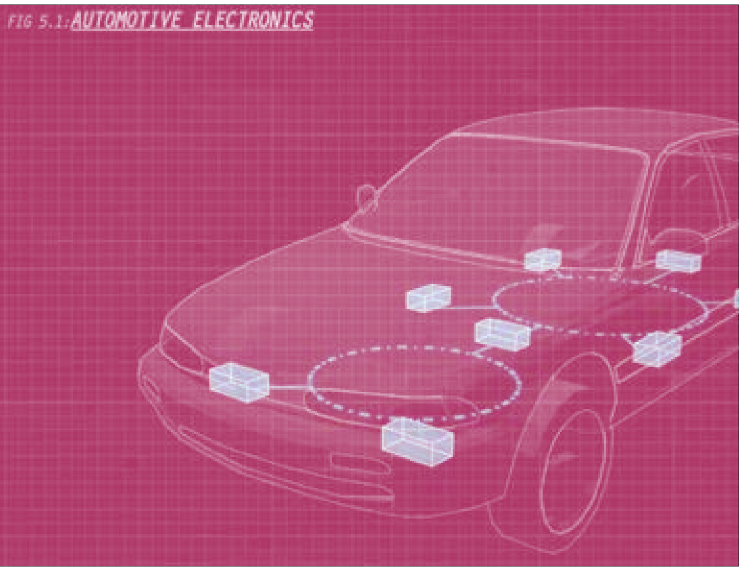

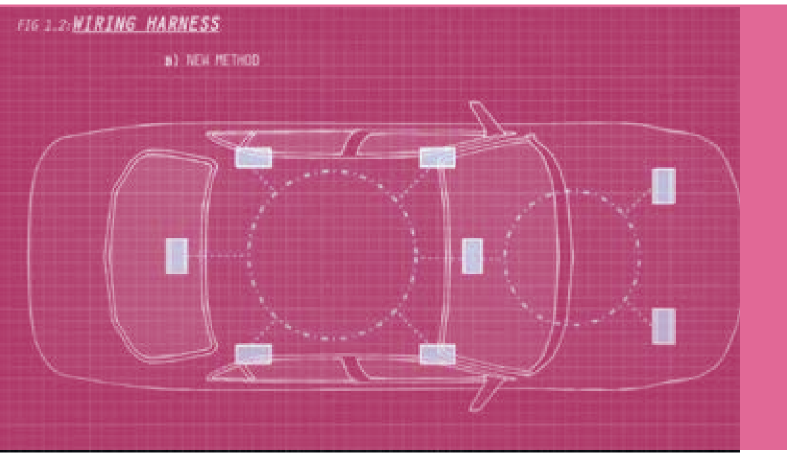

“You used to have a network that you could only run at 30% capacity before messages started missing their deadlines. With our approach, focused on using the correct priorities, we pushed that up to 70% or 80%.”

Volcano technology has been used in Volvo cars since 1997, and was added to the Ford Premier Automotive Group range after Ford bought Volvo in 1999. That meant it found its way into Jaguars, Land Rovers and Aston Martins as well. Today the Volcano technology is owned by Mentor Graphics.

Following their success with Volcano, Ken Tindell and Rob Davis set up another venture, Northern Real-Time Applications, focused on developing a real-time operating system for the low cost microprocessors used in ECUs. This work again capitalised on results from the Real-Time Systems Research Group, including research by Neil Audsley, Alan Burns and Andy Wellings on operating system design, processor scheduling, and response time analysis.

The two biggest pressures on the processors used in ECUs are processing load and memory use, while the biggest concern is cost.

“Memory on these processors is extremely limited”, says Rob Davis. “If you can use less memory, you can get individual microprocessors at slightly lower cost, and across a production run of a million or more cars that adds up to a big saving.”

“We started by realising that the functionality needed in a real-time operating system could be really quite basic. It needed to support the dispatching and running of tasks, and the sharing of data between them without that data being corrupted. We made sure that you didn’t have things going on in the operating system that didn’t need to be there. Effectively the operating system was tailored to the precise needs of each application”.

“The size of the code for the operating system is extremely small, about one and a half kilobytes, so small you can print it on a coaster.”

The company gained £1m in venture capital funding in 1998, and a further £9.2m in 2000, changing its name to LiveDevices a year later. Robert Bosch subsidiary ETAS bought the company in 2003, taking over the team of 20+ engineering staff.

"ETAS had their own operating system, and found it was not as efficient as the one we'd produced", says Rob Davis. "They could either start again from scratch or buy the company".

Since then ETAS have updated the operating system to the AUTOSAR standard and adapted it to work on more than 25 different microprocessors.

It has been used in about 1 billion ECUs by most of the world's leading car companies and their major suppliers. The total goes up by about 55 million a year. It's the world's smallest and fastest automotive operating system.

Research by the University of York's Real-Time Systems Research Group has reached further in recent years. Building on their industry knowledge, Guiem Bernat, Ian Broster, Antoine Colin and Rob Davis started a company called Rapita Systems in 2004 whose products give companies in the aerospace, automotive and space sectors more confidence that their systems will meet their timing requirements.

In these industries, if software takes longer than anticipated to execute, it can lead to issues with reliability or worse.

Using technology known as RapiTime, engineers are able to determine the longest time that each software component can take to run, called the Worst Case Execution Time, and budget accordingly. Previous approaches either tended to underestimate this Worst Case Execution Time, or required costly updates to the analysis tools for each new processor.

RapiTime was first evaluated on a system provided by Audi, and was used on BAE Systems' Hawk Advanced Jet Trainer project in 2006. The technology enabled BAE Systems to reduce the Worst Case Execution Time by 23%, by helping them identify the 1% of the code that was causing the bottleneck.

RapiTime technology has led to the creation of a number of high technology jobs in York, and is used in space, automotive and aerospace projects across the globe that would otherwise have to rely on manual measurement and analysis.

As one major aerospace supplier says: "Not only did we reduce our effort requirements for the testing, but we could use our results in ways that were infeasible before. It is now significantly faster for us to identify a timing issue, update the software to resolve the issue, test the updated software and verify that it's fixed."

As the automotive and aerospace industries develop new and more complex products, correct timing behaviour is crucial. A quarter-century of research and experience at the University of York has helped them be more confident that their systems will meet their deadlines.

"We've seen a huge revolution in electronics over the years, with significant improvements in reliability as well", says Alan Burns, current head of the Real-Time Systems Research Group, "It's great to know that our research has had a hand in it."

For more than 20 years, Professor Ross Anderson's research team at Cambridge University has been investigating vulnerabilities in payment systems. Its work has pushed businesses and organisations to re-consider their security, and raised awareness of the risks of certain transactions.

4

One day in Barcelona, a British sailor splashed out on the most expensive round of drinks of his life - without even knowing it.

It started off fairly innocuously, as he wandered into a bar and bought a round for 34 euros using his bank card. Shortly afterwards, 3,400 euros left his account. An hour later, the same thing happened again. And again, on the hour every hour for ten hours.

The sailor contacted his bank, which said its system was perfect and accused him of lying or being in league with the criminals. At this point, he got a lawyer, who sought some advice from an expert.

"I managed to provide the evidence that this incident was the result of malware in the terminal, using a particular protocol manipulation to which the card he used is endemically vulnerable", says Cambridge University's Professor Ross Anderson, a global expert in bank-related fraud. "It's something that the banks were trying to overlook. Eventually the bank gave this man his money back."

Ross Anderson

Professor Anderson is often sought out by fraud victims who need advice. Since 1992, his team at the Cambridge University Computer Laboratory has been investigating and testing the vulnerabilities of payment security. While other centres of expertise rarely publish their findings in the public domain, Professor Anderson's team is very willing to share with both the research community and the media.

"Once we've figured out what's gone wrong, we write a paper and send it off to various banking regulators, then they can then act as they see fit", says Professor Anderson.

"Banks are vast bureaucracies, with a dozen layers of management. The only way you can encourage a large slothful corporate to change its ways is to kick it in the balls."

Research carried out by the group has encouraged businesses to re-think their security and prodded authorities to improve their certification systems. In recent years, Anderson has worked directly with firms such as Google, Symbian and Samsung to make their platforms more secure.

The work has also led to the creation of a spin-out called Cronto, which provides authentication systems for online banking. Cronto was acquired by VASCO Data Security International for £17m in 2013.

Ross Anderson first began working at Cambridge University 23 years ago, after working as a bank security expert in Hong Kong. He initially discovered that the market for consultancy was "sewn up", but nevertheless found himself looking into the case of Andrew

Stone. Stone was eventually convicted of using information on discarded cash machine receipts to take a total of £750,000 from people's accounts.

Back then, receipts had the customer's full account number printed on them. Stone had been building access systems for cards with magnetic stripes, and had discovered that if you simply changed the account number stored on the stripe to another account, you could withdraw money using your own PIN. Andrew Stone would pick up the receipt that the customer had just thrown away and encode a card with the person's account number on it, which he could then use to withdraw money.

"The exploitable vulnerabilities at the beginning were mostly extreme stupidity", says Anderson.

Working as an expert consultant, he assisted barrister Alistair Kelman in a class action against 13 banks on behalf of 2,000 customers, asking for £2m back. Now banks only print the last four digits of the account number on the receipt.

"I got to sift through several dozen shelf-feet of witness statements, and I wrote a paper about it which appeared in 1993. It brought home to the research community how cryptographic systems fail in real life.

"This launched me on the path of being a thorn in the side of the British banking industry."

One of the research group's most famous discoveries came in 2010, when the team pointed out a notable flaw in the popular EMV smart-card system, known to many as "chip-and-pin". The system – which was introduced across the UK in 2004 – required users to input a four-digit PIN as well as having a chip on the card read by a point-of-sale terminal. The added PIN was hailed as a way to stop fraudsters from stealing cards and forging the signature, or cloning the magnetic strip on the card to use without the owner being aware of it.

Anderson's team observed that this process could be disrupted by criminals with as little as an undergraduate-level understanding of electronics. These findings were widely reported in the press, particularly by BBC current affairs show Newsnight.

When a customer taps their PIN into a terminal, the terminal sends a message to the card to confirm it is correct. If it gets a positive response, the transaction continues. However, the Cambridge researchers noted that it was possible to block this response by inserting a device known as a "wedge" between the card and the terminal. The device - which could be constructed using off-the-shelf materials - would then send the terminal a message saying the PIN was correct, regardless of the numbers that were typed in.

In a 2010 report called Chip-and-Pin Is Broken, the group said the vulnerability "exposes the need for further research to bridge the gap between the theoretical and practical security of bank payment systems. It also demonstrates the need for the next version of EMV to be engineered properly". This work led to changes to EMV in April 2012.

While the banks involved don't often welcome the attention, Professor Anderson states that all of his team's work is focused on pointing out issues so that they can be rectified in future. But he also believes that the Financial Conduct Authority needs to lean on the ombudsman to take complaints of bank-related fraud more seriously.

"It's good practice to shield your PIN at all times. But if you're an individual, there's really nothing more you can do about this. Fundamentally, this is a system maintained by the banks. Other than that, the failure is with regulation.

"Banks had been very careful with the truth when it came to the no-PIN chip-and-pin fraud. They said they didn't know about any cases in the UK, but were perfectly aware of cases in France where people had been prosecuted."

In past years, the group has unearthed a number of systemic failures in online consumer protection, including in the Payment Services Directive and in the Financial Ombudsman Service. Professor Anderson has been commissioned on three occasions by the US Federal Reserve to write papers for its Payment Systems Economics conferences, which are attended by industry leaders and policymakers from all over the world.

A product used by Barclays has also had a notable impact on security. Following a study of API security, Anderson's group suggested changes that would make the hardware modules used by banks less vulnerable. Security software firm Cryptomathic hired a member of the group - Dr Mike Bond - to become a security architect in 2006, and said it used the "great insights that arose from the Anderson group" to build a product called the Crypto Service Gateway. CSG has since been adopted by Barclays, where it improves performance and saves more than £1m a year on the development of new applications.

As well as shining light on some of the more significant vulnerabilities of transactions, the research has also supported the technology used by spin-out Cronto Limited. Team member Dr Stephen Murdoch is the chief security architect of the company, which is now helping to secure online banking in Chile, Switzerland and Germany. The system - which helps to shield consumers from malware - was rolled out to all of Commerzbank's 11 million retail customers in 2013.

Since 1993, customers around the world have become used to different methods of transferring money, from chip-and-pin to online payments and even contactless transactions. So how have the methods of fraud changed in that time?

"Things have become more industrialised", says Professor Anderson. "Around 20 years ago, card fraudsters would make their own equipment. Nowadays there are criminal workshops full of people with degrees in electronic engineering that manufacture this stuff, like card-cloning kits, and sell them into underground markets for several thousand pounds.

"The weaknesses we're noticing now are in many ways a reprise of 20 years ago. There are many gangs – particularly in places like Romania – that rig up ATMs with cameras, and use something called a 'Lebanese Loop' to physically capture the card you've put the ATM. If you go to an ATM and it eats your card, stand there and call 999 or call up the bank and make sure it's been cancelled.

"However, some people who do that have found the bad guys come back and take their card even if it's reported as stolen, as the banks allow it to be used in contactless transaction and still charge the customer. That's completely wrong."

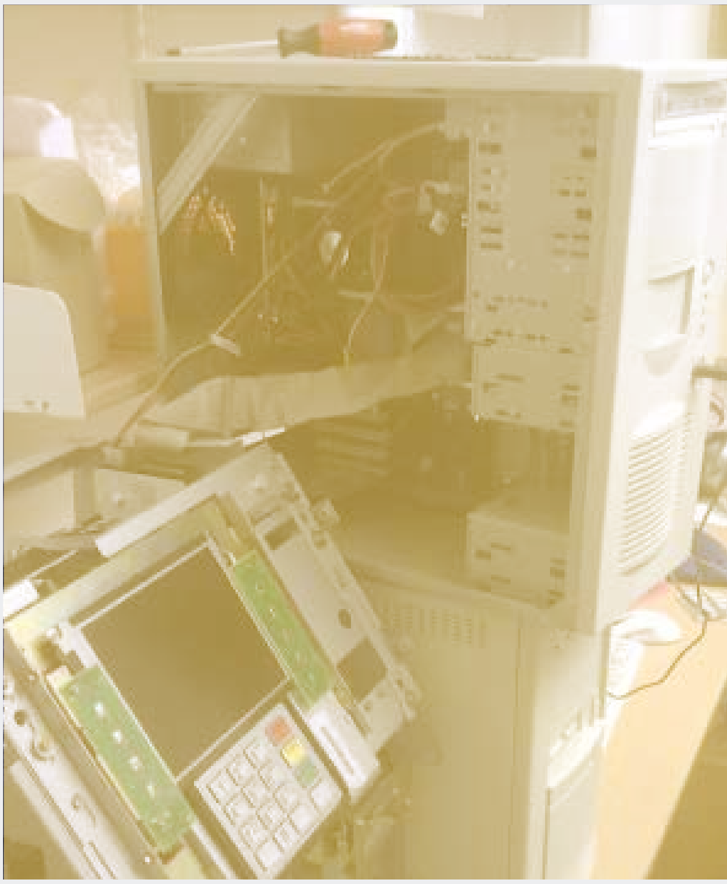

Example of extracting disk image from ATM

Crash testing vehicles was once a long-winded process involving frame-by-frame examination. Technology developed in Birmingham in the mid-nineties helped to streamline that process dramatically.

5

How long does it take to analyse a crash that only lasts a quarter of a second?

Thanks to technology developed by Birmingham City University, it's a process that's much shorter than it used to be.

Back in the mid-nineties -when it was known as the University of Central England – the university's Digital Media Technology research group started working on a new system for reviewing multimedia footage of crashes. It quickly revolutionised the way vehicle companies analysed crash data, and even prompted a new international standard.

The research has assisted manufacturers in creating safer vehicles, and enabled faster, more intuitive and more insightful analysis of crash test data.

"It proved very popular with the major manufacturers", said group leader Professor Cham Athwal. "And within a few years, everyone was bringing out copycat systems."

So why did it become so popular? To give you an idea of that, let's look at the amount of equipment that is generally required for the average crash test experiment.

Car companies know that thorough research into crash safety saves lives. Therefore they are keen to get as much information from their experiments as possible. As the test vehicle charges toward its impact, it is typically recorded by around 12 high speed cameras, 40 vehicle transducers, and 60 transducers attached to the dummy to register injury. The crash is recorded with a number of still photographs, while measurements are taken of the vehicle pre-and-post-crash. The investigation sparks a whole raft of paperwork, from explanations to graphs, tables, analysis and conclusions.

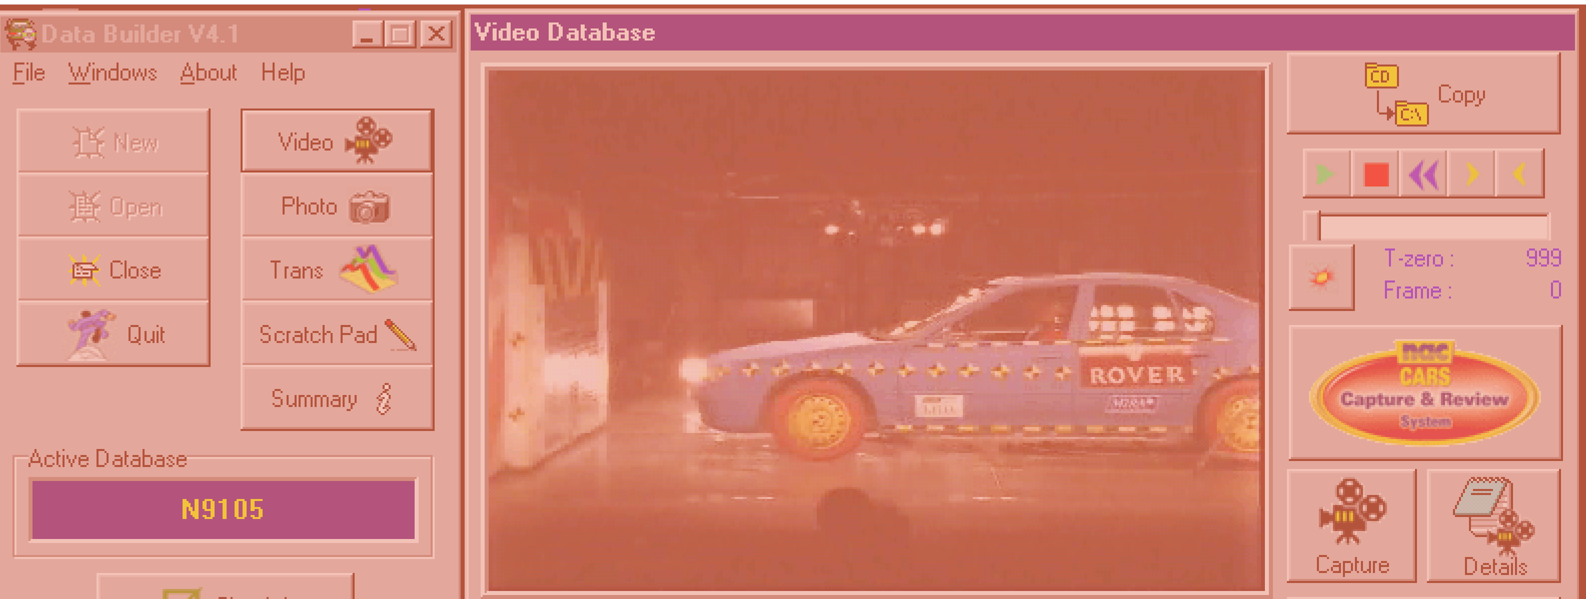

Now imagine that all of this footage is shot on 16mm celluloid film, at 1,000 frames a second, and has to be reviewed meticulously on a special projector, frame-by-frame, from different angles.

Exhausted yet?

"Before our system, people would have to spend days looking through rolls of film, picking through data and drawing up graphs", said Professor Athwal. "Now they can do the analysis in a day or half a day, and have the time to make better design modifications."

Working with the Motor Industry Research Association (MIRA), Professor Athwal's team researched and developed a system that could blend multimedia such as video footage, stills and sensor data from multiple sources, and mesh it into a single chronological record of the crash.

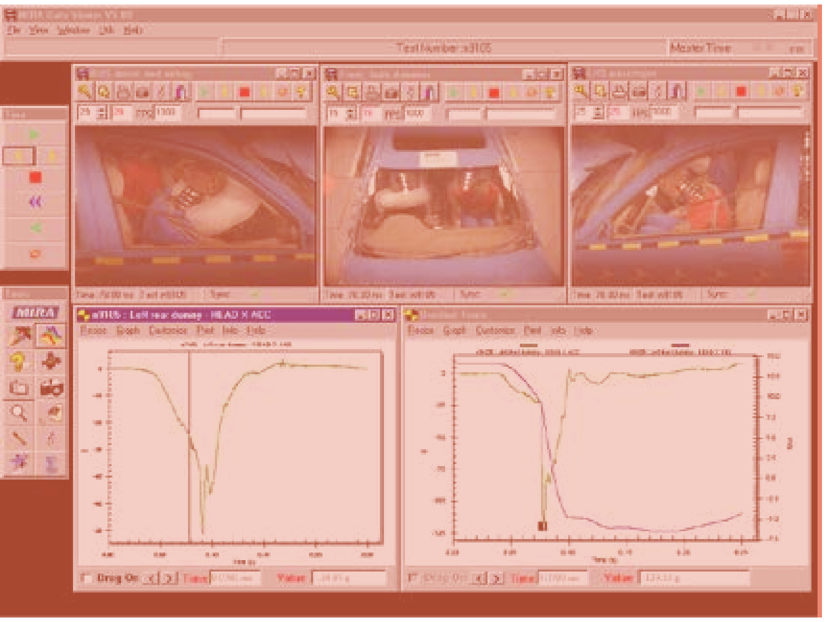

Investigators could pinpoint an instant in a crash, and see information from different angles and sources, synced to that exact split-second.

"In the past, people would have lots of computer files to check, with data coming from all the sensors, as well as lots of paper and rolls of film", said Professor Athwal. "We built a front-end system that allowed investigators to see all of the angles based on the specific time. You could then step through each of the videos and datasets and see what was happening at a certain point."

The system – which was made possible through funding from the Technology Strategy Board and EPSRC - reduced analysis time from several weeks to a matter of hours. It also allowed investigators to compare the test frame-by-frame against a previously-completed simulation, and assess whether the simulation was calibrated correctly or not.

"It allowed companies to make better design decisions more quickly, and dramatically shorten their design cycles", said Professor Athwal. "Plus, it helped to reduce injuries and prevent fatalities."

A pair of products were launched: DataBuilder structured and stored crash data, while DataViewer allowed companies to analyse it using multimedia. MIRA began using the technology in-house, and noticed significant improvements in efficiency and accuracy.

People sat up and paid attention. The prototype project won the Teaching Company Scheme's Best Technology Transfer Programme award in 1999. It emerged into the market in 2003, and before long major national crash test facilities were using the system, from European standard bearers such as Millbrook Proving Ground in Bedford, to international car giants such as Ford and General Motors.

It also had a lasting effect on how the industry worked. The team's research was incorporated into the international standard for crash data, and remains there to this day.

Considering the level of industry interest, a spin-out startup seemed inevitable. In the mid-2000s, research associate Jimmy Robinson set up Pixoft to work with companies and refine the research further. This fine-tuning resulted in the 2006 release of Vicasso, a market-leading product that is now used by prominent names such as Jaguar Land Rover and General Motors. It has also found markets outside the car industry, with clients including Duracell, Unilever and the US Army.

MIRA was also finding other uses for the system. It has been used for rail and plane safety, even down to studying how seats react to a collision. It has even been adapted to gauge how buildings and structures might respond to impact, from Westminster to Heathrow Terminal 5.

“The system itself is very flexible”, said Professor Athwal. “It was originally built with car crash testing in mind, but it has been useful for all sorts of engineering and scientific investigations where data and video are recorded together.

“We have had people use it to measure how bubbles form in biological systems, and to conduct spray analysis. One example we gave in our paper was analyzing astronomical data from solar flares, which involves looking at videos of solar flares and synchronising readings of radiation on the earth’s surface.”

Meanwhile the team’s work on digital media systems continues, including some projects that wouldn’t look out of place in a sci-fi blockbuster movie.

“At the moment, we’re looking at virtual TV systems that are interactive. We are researching the creation of scenes that show virtual objects that a person can grab and move around in real-time, like Iron Man in the movies.

“Imagine using something like this for consultations. A medical professional could show you a 3D projection of your knee, created from your latest CT scans. They could show you how it fits together and point out exactly what’s wrong with your knee by grabbing the model and manipulating it in their hands. The doctor would see it on the screen in front of him, while the patient at the other end of the Skype call sees it as if it was an actual object that was being held and moved about.

“That would certainly be a more intuitive solution.”

Back in the 1990s, Dr Heather Heathfield set out to understand the reasons ***why* so many new IT systems in healthcare failed.** She was part of a team that developed a systematic methodology for evaluating these systems, making them easier to use and maximising the benefits for patients and clinicians.

6

Switching to electronic health records can be extremely beneficial.

With electronic records, clinicians can provide accurate, up-to-date, and complete information about patients at the point of care. They can share electronic information with other health professionals. They can diagnose patients more effectively, reduce medical errors and provide safer care.

However, back in 2000, the NHS was still predominantly sticking with paper records, while other industries were gravitating towards IT systems. Why was this?

In short, systems were not designed to fit into clinicians' daily working practices, and involved spending large amounts of time typing in data. Consequently, they were shunned in favour of tried-and-trusted methods.

In 2001, the NHS Information Authority initiated The Electronic Record Development and Implementation Programme (ERDIP). The new programme funded a number of beacon sites to explore the issues associated with the creation of electronic health records.

Dr Heather Heathfield was part of a team commissioned to evaluate the programme. Dr Heathfield had designed electronic health records and decision support systems as part of her PhD work, and continued to explore this field both at IBM and later as a senior lecturer at Manchester Metropolitan University. She worked on the project with healthcare consultancy Secta and the Yorkshire Institute for Health Informatics at the University of Leeds.

"Having been involved in the design of a number of electronic health records and decision support systems, I was frustrated that despite the obvious benefits, these systems were not being used" she says. "I wanted to know what the barriers were and how we could overcome them."

Their evaluation helped them to discover a great deal about the barriers that stopped electronic health records from being implemented, and gave them some insight into how to systematically identify and address them.

Following several publications on the subject, the NHS Information Authority went one step further, and encouraged them to work on a practical toolkit, suitable for an age in which clinicians were being asked to take a larger role in evaluating and commissioning systems. The Project Review and Objective Evaluation toolkit - known as PROBE - promoted a way of thinking that placed evaluation at the start of the process, and encouraged organisations to be more rigorous with their processes.

Heather Heathfield

First released in 1996, PROBE was based on a six-step evaluation process, in which decision makers would be asked to agree on why an evaluation was needed, when, what and how to evaluate, analyse and report and finally to assess recommendations and decide on a course of action.

PROBE became a central part of NHS information strategy and policy. Since its revised publication in 2001, it has sparked the NHS to look at evaluation as a key strand of major IT projects.

Dr Heathfield says there are several "key principles of evaluation", including advance planning, the establishment of clearly defined aims and objectives, and the inclusion of formative and summative elements. Evaluation needs to be closely integrated into the life-cycle of the project, and evaluators need to be able to compare the results to what went before by collecting both quantitative and qualitative data throughout.

"An evaluation framework needs to address structure, processes and outcomes", says Dr Heathfield. "It should do this across five dimensions - strategy, operational, human, financial and technical."

In 2012, NHS Connecting for Health's head of patient and public partnerships said that Dr Heathfield's work "greatly improved understanding of ICT projects" by focusing analysis on the important questions, and providing more information on the tools and techniques available for effective evaluation.

The PROBE methodology is still widely used in the UK. For example, it was employed by the National Patient Safety Agency to evaluate systems designed to improve the safety of blood transfusions.

A total of five patients died after receiving incompatible blood during transfusions between 1996 and 2004. The issue contributed to the deaths of nine others, and caused major illness in 54 more. The National Patient Safety Agency used PROBE to evaluate systems that could improve the process.

Using PROBE helped the National Patient Safety Agency develop a detailed picture of the views of stakeholders, and encouraged decision makers to think about areas such as installation of the system, any changes that were required, the response of staff, and its effectiveness and reliability. The final report in 2011 observed that the system chosen as a result of this process reduced the time to resolve problems with transfusion from 123 hours to around 30 minutes

Dr Heathfield is now a principal and Head of Evaluation at OPM. Her work focusses on evaluating complex transformational initiatives in health and social care, with a particular emphasis on integrated care and new models of commissioning.

"My key aim at the time was to raise awareness and promote debate about the role and appropriate methods for evaluating IT systems in healthcare", Dr Heathfield says. "It certainly helped people design better evaluations. PROBE is still in used in teaching, by researchers and others involved in that process.

"I still use the principles set out in PROBE in all my evaluations. It's about being systematic and understanding the impact of an intervention or new model from a range of perspectives, including how it fits into local working practices."

Overview of the Evaluation Process

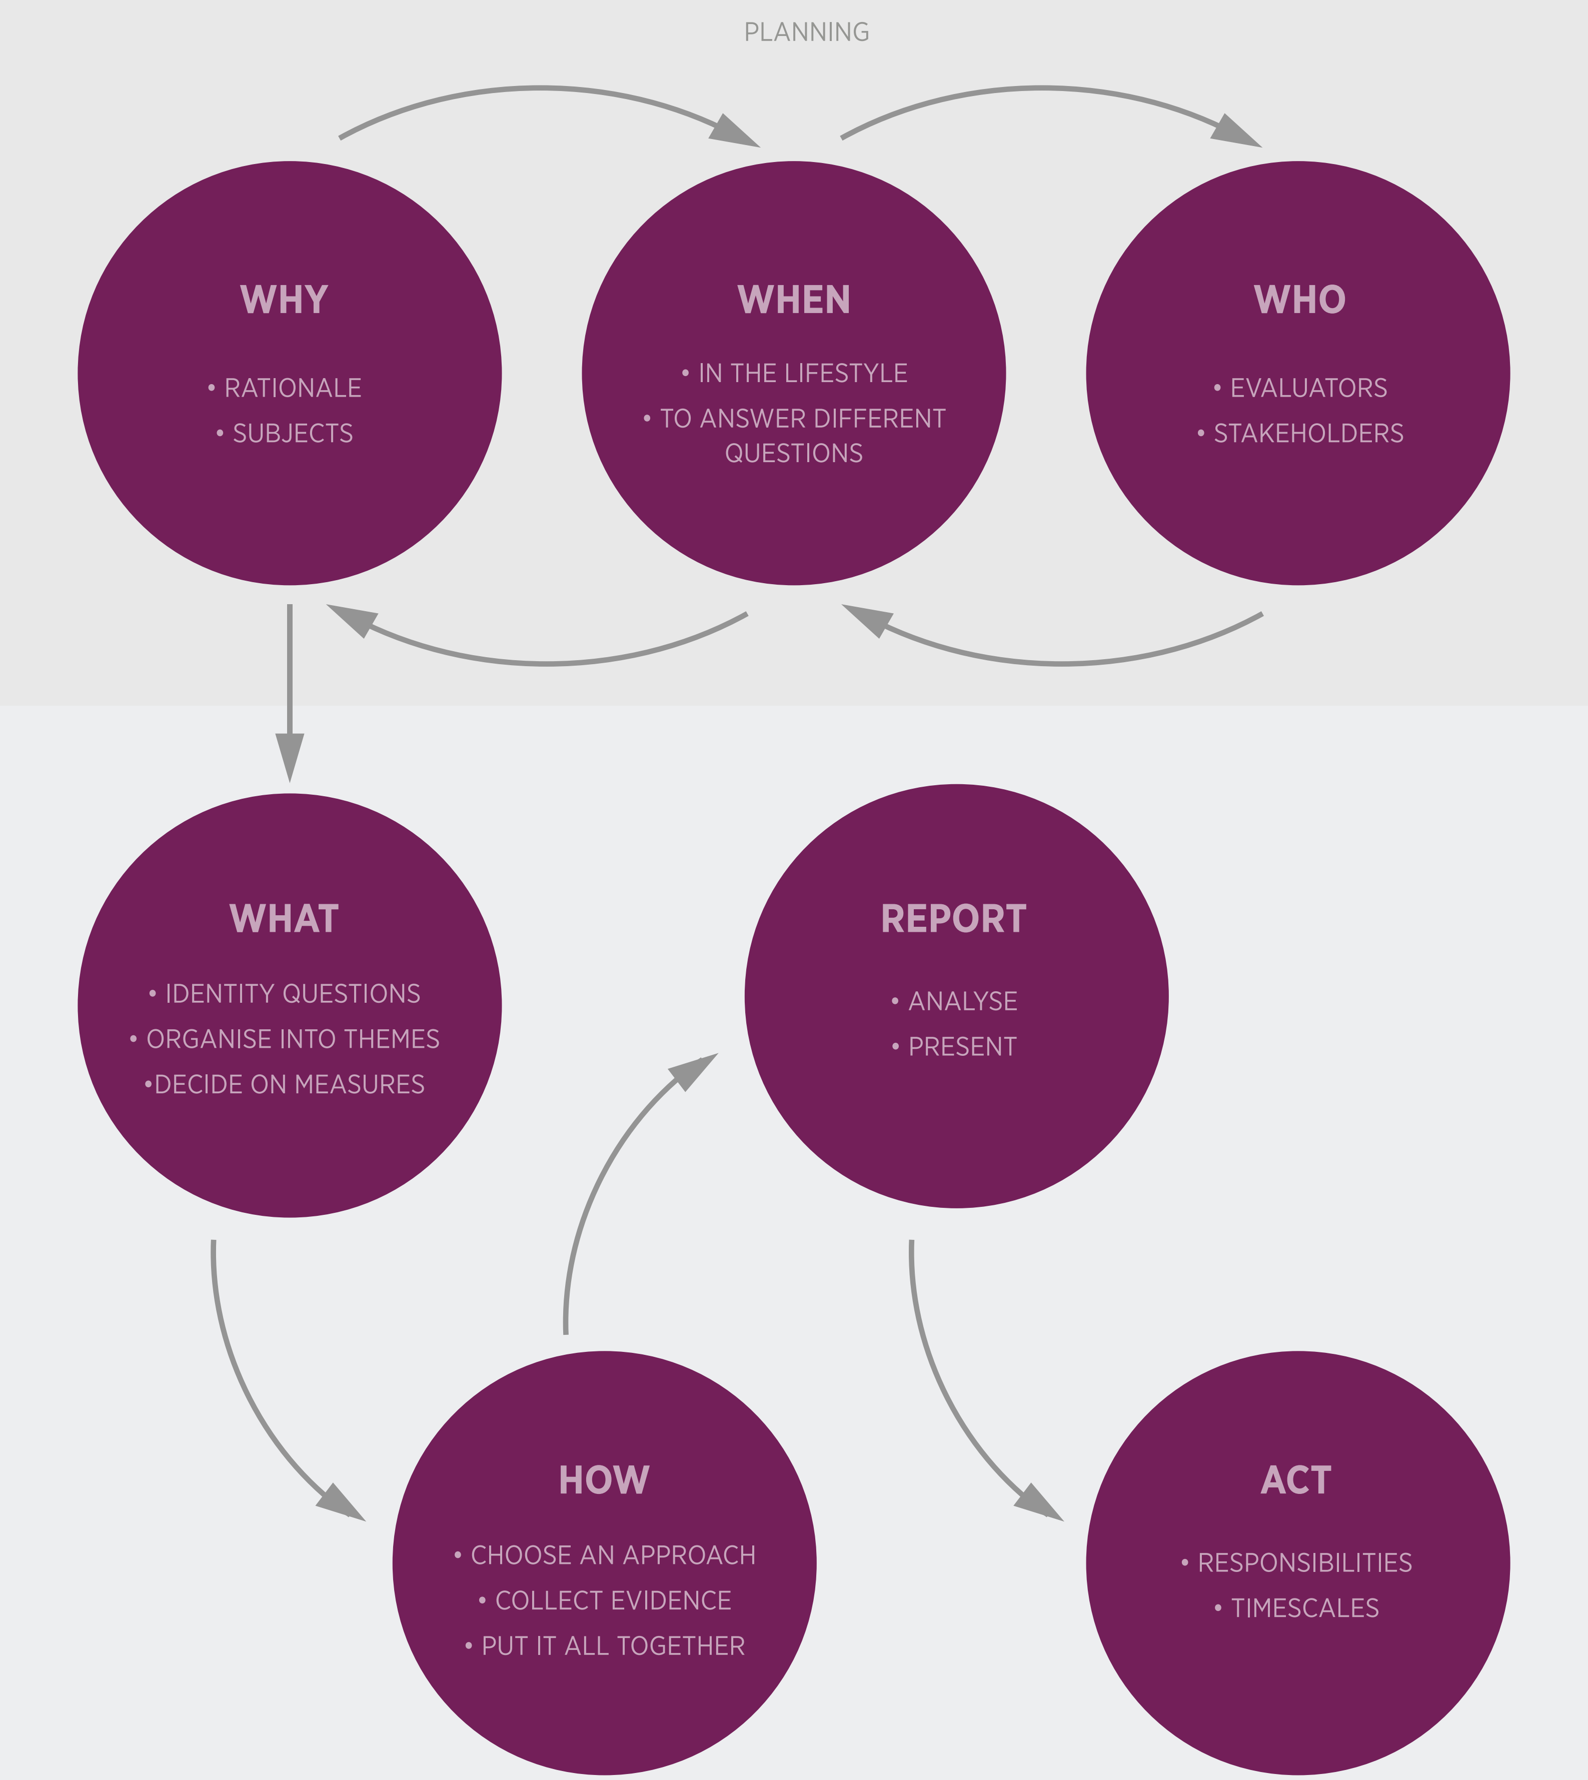

In the early nineties, University of Manchester Professor Madan Singh launched **an intelligent system that would help companies work out their pricing.** The resulting company has assisted more than 400 companies in 80 countries.

7

What is the right price for a product?

Dr Xiao-Jun Zeng

For many years, pricing wasn't exactly a scientific process. Businesses would weigh up the cost of parts and labour, the demand, and what the competition was doing.

Prices changed every day, and there wasn't really a very good way to reliably compute the best amount to charge, even when systems started being introduced in the mid-1980s.

This was particularly true in sectors such as retail and petroleum, where prices could vary across a number of different sites.

In 1993, Knowledge Support Systems Ltd began offering an approach which combined academic research into pricing with a computerised system which could analyse and react to the latest data.

Working with businesses in the retail and petroleum sector, KSS and its spin-off companies provided jobs for over 150 people, assisted more than 400 retailers in 80 countries, and produced turnover of £19.2m in 2012.

The origins of KSS came when Professor Madan Singh started working on pricing while he was a visiting scholar in France. His study continued for roughly ten years, by which time he had set up at the University of Manchester.

Dr Xiao-Jun Zeng was a PhD student at Manchester at the time, and became involved in the creation of Intelligent Pricing Decision Support Systems that could make reasoned decisions and track prices in the wider market.

"You need to combine knowledge and data", says Zeng, who is now a senior lecturer. "Once you have an accurate demand model, you try to maximise your profit based on demand. But at the same time you don't want to lose your market share, or affect your image.

"It's quite a complicated optimisation process, so you need a certain kind of algorithm. Our work was to create an accurate demand model from the data, and then to optimise it. Of course, you're not just optimising for short-term goals, but strategy goals as well."

Under Professor Singh's leadership, the team of Dr Jean-Christophe

Bennavail, Dr Zeng, and later Nathalie Cassaigne researched the necessary frameworks and algorithms to handle the task. But it also required a shift in culture for the industries in question.

"When we started, people didn't really collect all of the data required", says Dr Zeng. "They collected their own price data, but not competitor data as that was costly and they felt you only needed to survey next door or the next street. Clients started to realise they needed to do much more. Nowadays, you've got data providers who just collect the data and sell it to companies."

The products drew from two key areas of research. The first looked into frameworks that helped businesses make systematic decisions about prices, while the other used historical price-sale data to get an idea of the best demand models.

"You need to think about competitor price, and the cross-effect on other products. For example, if you do a promotion on Coca Cola, your Pepsi sales may be affected. I might say that if I have more customers, my price increases. But if do that, you will lose your market share. Thus you have to think about wider impact that isn't isolated to a single product, combine your short term and long term business goals, and you need to react quickly when costs change. It's too complicated for a human being to do that in an optimised way, but if you have a computing tool things get a lot easier."

Professor Singh of KSS continued as chief executive until 2001, when he announced he was stepping down to make way for a replacement with more management experience. He sadly passed away in May 2002.

Professor Singh's research remains a strong basis for KSS. In 2007, the company de-merged into KSS Retail and KSS Fuels. While the former brought science-based pricing to more than 30 companies from Kroger to Tesco and 7-Eleven, the other served more than 400 petroleum companies in a market where prices can change in a matter of hours. It did so using intelligent pricing systems known as PriceStrat and PriceNet respectively.

Controlled experiments with oil companies indicated that each fuel retail site could achieve profit improvements of between £4,000 and £10,000 per year, while convenience stores could boost profits by more than £50,000 a year on average.

Jeffrey Miller, the president of Miller Oil, has said that "KSS Fuels and PriceNet helped us revise our pricing process to be more accurate and timely in response to competition. PriceNet alerts us to stores needing attention so the pricing team can spend time where they add most value to the business, making price decisions".

In 2012, a 40% in KSS Fuel was sold for £7.2m, valuing the company as a whole at £18m. KSS Retail was sold to Dunnhumby Limited for £12.9m in 2010.

The research cooperation between the University and KSS Fuel and KSS Retail continues today. The most recent completed research includes the PhD project "Dynamic demand modelling and pricing decision support systems for petroleum" in 2014, jointly funded by EPSRC Case PhD studentship and KSS Fuel. One of the outcomes of this project is a "competitor prediction tool". KSS Fuel is currently in the process of applying for a US patent.

Further research efforts by Dr Zeng's team at the university include extending their research into demand modelling and pricing optimisation techniques relating to online retail pricing and real-time energy pricing.

> "The proportion of retail that takes place online gets bigger and bigger all the time, but there isn't really an optimisation tool to handle it."

"People's behaviour is different. They buy one or two items online, and there might be 100 to 200 different competitors. The effect might be different to a traditional grocery store, as it's difficult to collect competitor prices because you don't know who or where they are.

"You need to have dynamic optimised pricing so you can respond quickly and automatically when things aren't going your way. It's quite a challenge, but it's a topic I feel is very important and interesting."

Spot The Difference can be a relatively simple game. But what if the images themselves are deforming? Imperial College London spin-out IXICO has the ability to accurately spot subtle changes in the human body, even in three dimensions.

8

These days, it is reasonably commonplace for doctors to use scans to assess a person's health.

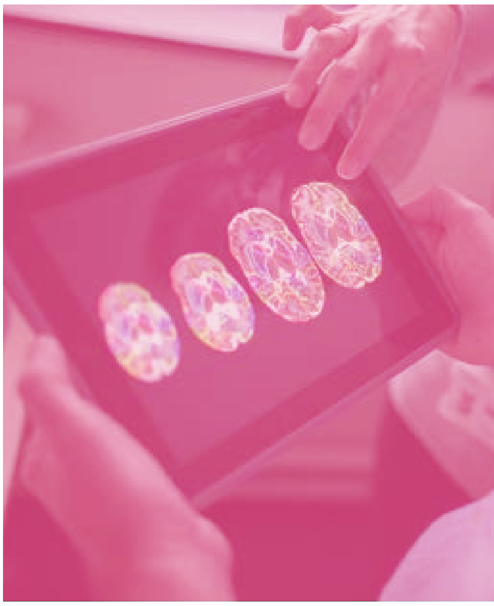

The clinician will review an image, looking for changes to previous scans or areas that may be anatomically abnormal. This is how they discover if a person has sustained injury or developed an illness that needs treatment.

Comparing two images might seem straightforward enough, as long as you know where to look. However, organs don't always look the same. Aside from changes in position and perspective, issues such as tissue deformation and motion from breathing can even make two images of the same organ look very different.

This is why IXICO has proved so popular in the healthcare and pharmaceutical industry. Born out of the work of Imperial College London's Biomedical Image Analysis Group, the company licenses a system which can map one image onto another, even if it is distorted, giving clincians the opportunity to look for changes and possible symptoms.

The process has been used predominantly in the analysis of brain scans. It returns highly accurate results that are invaluable in diagnosis, and save a significant amount of money in clinical drug trials. It is even being trialled to discover how it might help spot the symptoms of Alzheimer's Disease at an earlier stage.

"Our technology might not lead to the discovery of a drug per se", says group founder Professor Daniel Rueckert. "But it does provide a very efficient way of testing whether it works or not."

Professor's Rueckert's background is in computer science, but he began to see how his skills could be applied to healthcare while completing a post-doctorate at a hospital. He joined Imperial in 1999, and founded the group.

The problem is this: if you are trying to compare two rigid 2D or 3D images, there are only a small number of ways an object could move. However, if the object is able to distort in some way, the possibilities multiply dramatically.

To solve this, Professor Rueckert's team adapted an approach used in computer graphics animation to address the problem, called B-spline free-form deformation.

"It is used to allow animated characters to move smoothly. We adopted a similar technique to animate a 3D volume, and match it to another by deforming it.

"It allows you to treat the image like an elastic material such as rubber. If you have a cube of rubber and pull on its corner, it will deform in quite a smooth way. You can pull on this cube at every point until you get the deformation you need. That way, we can match images that look very dissimilar."

This technique has allowed researchers to compute "imaging biomarkers", or measurements that are used to assess changes caused by a treatment or therapy.

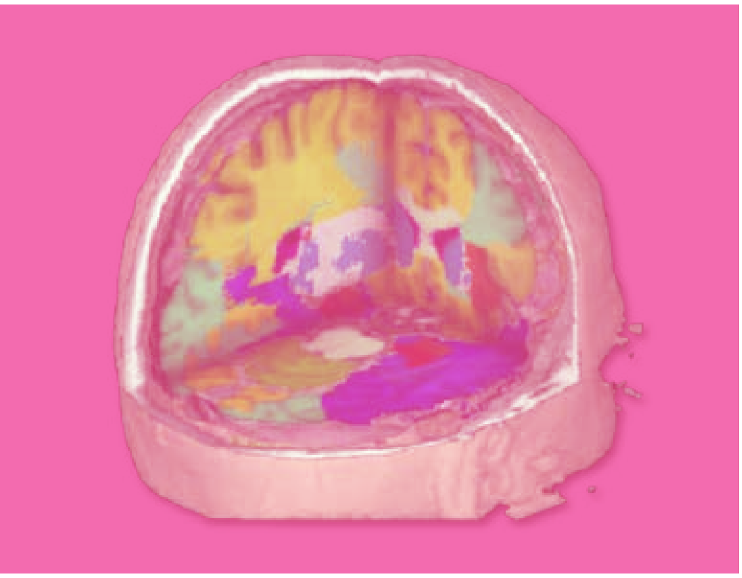

"We knew we could potentially solve a lot of problems, but we found it difficult to get the big companies like GE, Siemens and Philips to integrate it", says Professor Rueckert. "So we decided the best way to deliver these solutions was to set up IXICO."

IXICO was launched in late 2004. By 2006 the researchers had developed their work to the point where they could confidently begin applying it to various areas. One particularly useful application came in an area known as "image segmentation". Previously, experts would use an annotated "atlas" to identify which points belonged to a particular structure within the organ. The Biomedical Image Analysis team combines multiple images and a machine-learning element to build a more robust and accurate segmentation tool.

Their approach became a standard for medical image segmentation, particularly for cardiac, neurological and abdominal images.

IXICO focused its solution on two main areas: Healthcare diagnostics and clinical trials. The approach certainly made sense for the pharmaceutical industry, as clinical trials are by no means cheap. A Deutsche Bank Market Report in August 2012 estimated that the average cost of a new, approved drug had ballooned from $100m in 1979 to $1.9bn in 2011. IXICO's promise of more accurate results would allow companies to test fewer candidates, which was a big consideration.

"In commercial terms, the cost of a clinical trial is measured on how many patients you need to determine success. If you have inaccurate measurements you need a large number of subjects, because any individual measurement you make is not very reliable.

"If you say that every patient costs the company at least $30,000 per trial, accurate measurements can be a huge cost saving."

Interest from clients all over the world has helped IXICO grow from a five person company in 2008 to a 40-strong operation in 2013. It has secured more than £17m worth of business from global pharmaceutical companies such as GSK and Pfizer, and its revenues have trebled to £3.6m over the last three years.

It is also been used to analyse medical images from centres in North and Latin America, Europe, Asia and Australasia. It has been involved in more than 40 clinical trials, involving more than 10,000 subject visits.

IXICO has primarily focused on the brain so far. It has developed products for diagnostic use which are undergoing trials at new NHS brain health centres. There are also plans to look at the heart and the musculo-skeletal system.

Imaging biomarkers will almost certainly play a huge role in the future treatment of Alzheimer's

Disease. It is generally agreed that the disease is best treated in the pre-dementia phase, and the US Food and Drug Administration is currently reviewing an application to use MRI imaging to measure the volume of a segment of the brain called the hippocampus. This is a less invasive option than acquiring cerebrospinal fluid for analysis.

With these applications in the works, Professor Rueckert is keen to go further still.

“The next big step is to go beyond measuring things that we now believe are useful. Now we’re measuring the volume of a structure, such as the hippocampus. In future, we want to discover what we should measure in the images without the clinician telling us in advance.

“If you give us 1,000 images of some disease, we want to be able to find out what these patients have in common. We currently rely a lot on older textbooks about anatomy. A disease might have other important characteristics that are very subtle. We have a chance to find them using machine-learning.”

At the turn of the century, a team at UCL helped produce Britain's first major survey of body size. That research led to a 3D scanning process that has proved popular with retailers, and has medical and fitness applications as well.

9

Did you ever get the feeling that your clothes didn't quite fit like they should?

Using 3D human body measurement, researchers have been working on a way to make the latest styles just that little bit more comfortable.

"In the past, companies used to measure clothing sizes using fit models and tape measures", said UCL's Professor Philip Treleaven. "The fit models were based on the company's estimate of the size and shape of the average customer. They weren't particularly precise."

What did this mean for customers? For one, it meant that clothing might not be quite the right fit for most people, whether it was tight on the collar, loose on the shoulders or a little too snug around the waist.

At the turn of the millennium, Professor Treleaven's team started working on a solution, using 3D scanning. Throughout the mid-to-late nineties, they had developed the first size-extraction software for the Hamamatsu 3D scanner, which allowed them to draw linear measurements from a point cloud of data.

The team - and indeed the UK government - thought the technology would be prefect for cataloguing the size of the nation, and devising a size chart that would suit a modern population. A study was born, funded by a £2m Department of Trade and Industry grant and £2m pooled from 14 major UK retailers.

"That's where the national sizing survey started", said Professor Treleaven. "We measured a statistical sample of the population, and worked with clothing companies so that retailers would know exactly what an average size 12 person's body really was.

"The publicity was great. We were told that it was the most popular piece of work with the general public that they'd ever funded."

SizeUK was Britain's first national sizing survey. The UCL team worked with an American company called TC2 to develop the size-extraction software. Between 2001 and 2002, measurements were drawn from more than 11,000 people nationwide, selected from both genders, different ethnic and socio-economic groups, and seven age groups from 16 to 95.

They were scanned with a TC2 body scanner, which Treleaven describes as "like standing in a big telephone kiosk". The process would pick out key measurements such as chest, hips and waist. Fellow UCL professor Bernard Buxton developed an algorithm that would calculate average body sizes from the cloud of data created by the survey.

"The body scanner scans you like a flat-bed scanner, but in three dimensions", said Professor Treleaven. "It means you can extract measurements like size and shape very accurately."

The project didn't just involve measurements. It required a system of data management that could automate registration, store results and protect data that might be considered sensitive.

The idea caught on. Retailers such as Tesco, Arcadia, Next and Marks & Spencer participated in the process, and now all major UK retailers sell clothing that is based on the resulting size charts. Other countries came calling too: Professor Treleaven's team soon found themselves busy with SizeUSA, as well as surveys in Germany, France, Mexico, Sweden and Thailand. In Thailand, there has been a drop in overall clothing alterations of roughly 35%.

While it used to take up to five iterations to make a sample that was approved by the retailer, that number has now halved. That has decreased product development time by at least 25%.

"It took off around the world", said Treleaven. "Governments like it because it's popular science. It's not like explaining DNA coding or quantum computing. You could go into a pub and explain it, and everyone would get it."

Retail companies loved it too. A spin-out launched in 2000 called Bodymetrics develops body scanners for clothing stores. At Selfridges in London and Bloomingdales in west coast USA, customers can step into pods to discover which brand of jeans might suit them the best. Three years after it started being used in 2009, the pod accounted for 20% of premium jeans sales at Selfridges.

In 2008, the team set up another spin-out company called Sizemic to sell charts and mannequins based on this data. The mannequins are both physical and virtual, and are far more accurate than previous alternatives. Seismic has now grown into a 10-employee company with a turnover of £1m.

The results have proved useful beyond the clothing industry as well. Using data from the SizeUK, SizeUSA and SizeTHAILAND surveys, Professor Treleaven has been working with Professor Jonathan Wells at UCL's Institute of Child Health on studies into obesity.

Under the umbrella of another spinout called ShapeDynamics, they have developed a 3D system for GP surgeries that compares a patient's measurements to the size survey, plotting how they vary from the norm. The process could be particularly useful in detecting traits such as stomach shape, which could indicate a propensity for Type II diabetes. Since 2011, 240 patients have been scanned as part of trials in a large practice in the City of London.

Professor Treleaven believes the applications stretch to many different areas, from showing how people would look after they increase exercise or stop smoking to demonstrating the outcome of cosmetic surgery.

"For example, we think you could use these body scanners for sport", he said. "You can use it to track how people's bodies change during exercise, and you can also identify talent using kinanthropometry. A world-class swimmer has hands and feet like dinner plates and a body like a beanpole, and we can pick that out."

However, for the UCL team, the big change is coming as the technology becomes more accessible.

"The showstopper is going to be the arrival of a cheap and easy-to-use scanning system. When we started body-scanning, the scanners cost about $100,000, and could do about 10 measurements. We had to work with a scanner company to allow us to be able to improve that to 150, but now we can do an unlimited amount.

“Some parts of the body are still difficult. For example, if you’re scanning someone’s head, the hair deflects the system so you have to do it manually. But now the price of scanners has dropped to £7,000. If you can get that down to under £200, then people could have one in their bathroom. And then the possibilities are huge.”

Mr & Mrs Average

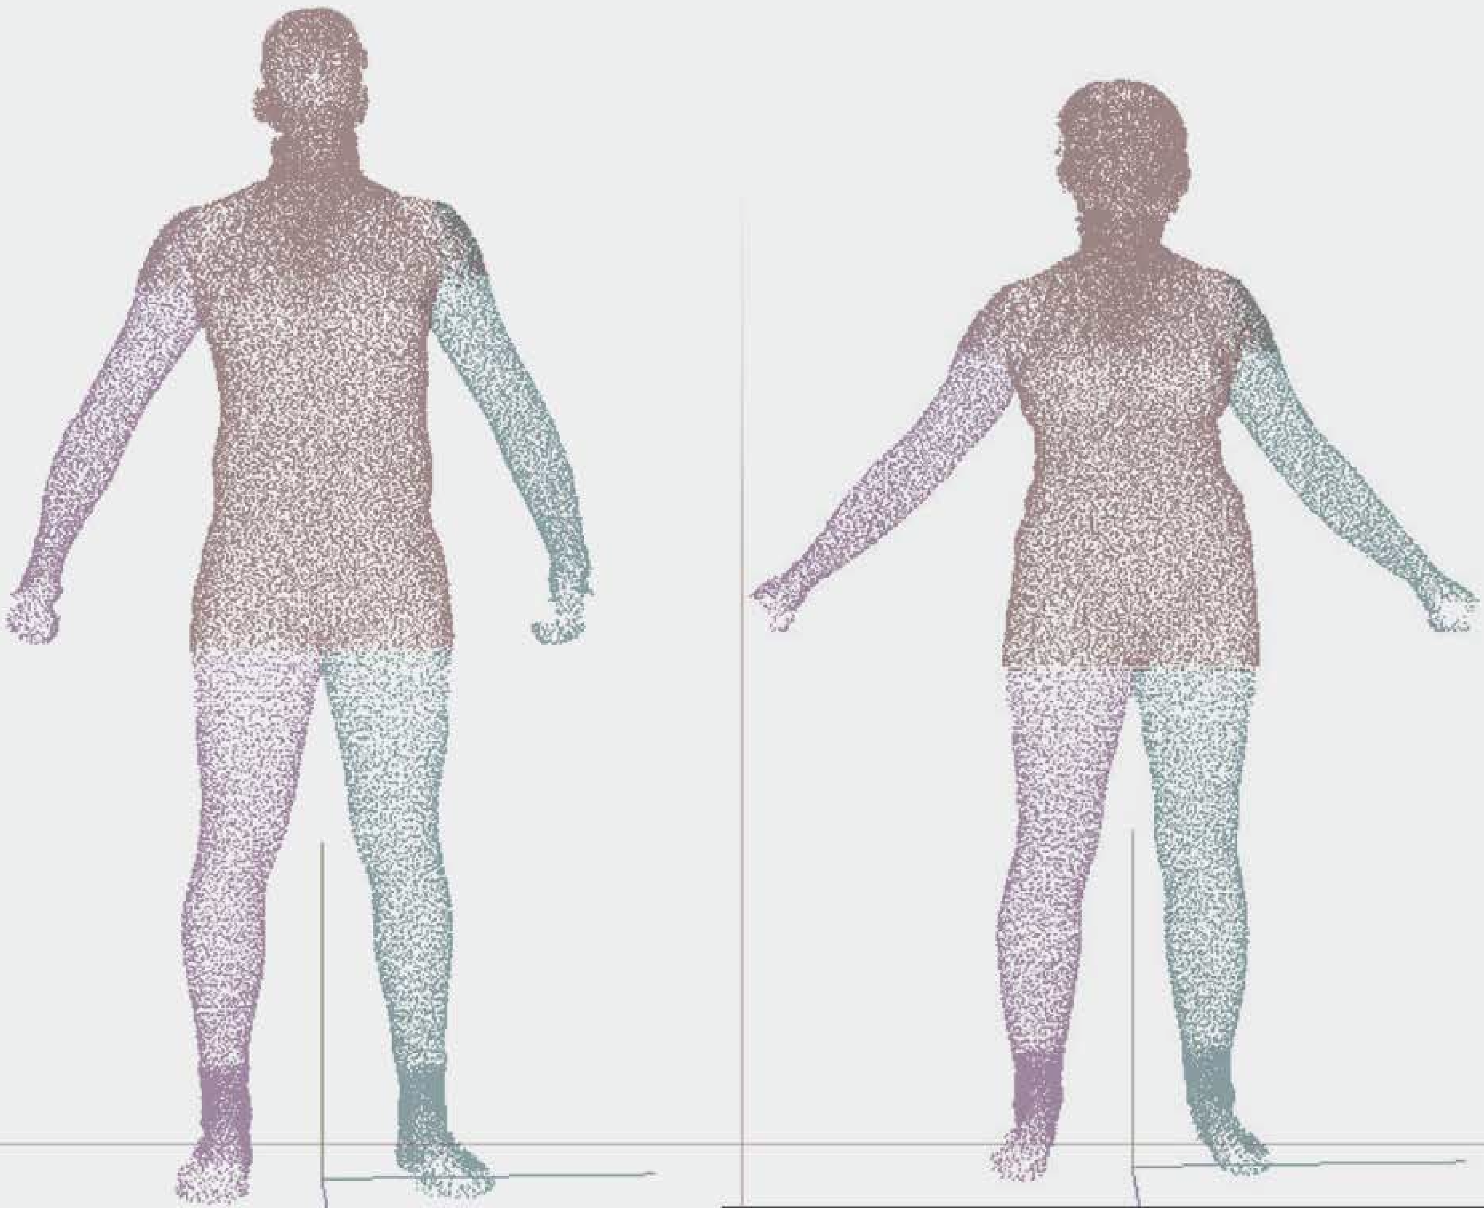

Translating speech into sign language isn't as simple as just replacing words with signs.

At the University of East Anglia, the Virtual Humans Group is blending AI, 3D animation and speech recognition to create a better way to get the message across.

10

Communication isn't just about words.

After all, "Good job" can be transformed from a compliment to an insult with a twist of the voice. Every day, we use our tone, our hands, our faces and our body language to add to what we say, and even to alter its meaning.

It's the same for sign language. So how do you teach a computer to read all that, and pass it on effectively?

"People imagine that sign language is all about the hands, but quite a lot comes from the face", says Professor John Glauert. "Your facial expression changes what you're talking about, and whether you're happy about something or not.

"Also, when people perform signs they make mouth movements that go with the words. Sometimes you change the meaning of a sign with a facial expression. To change the type of fish, you might sign fish, but mouth "salmon". The only difference between the sign variations is on the face, not the hands."

Since 1999, the Virtual Humans Group has been looking at ways to translate everyday communication into recognisable and nuanced sign language. They develop systems that can interpret speech and language, and animate a 3D character to make recognisable signs and gestures.

This requires a diverse mix of skills. Based at the University of East Anglia, the team has expertise in speech and language recognition, 3D character animation, AI and computational linguistics.

Professor Glauert says: "Around 50,000 people in Britain use British Sign Language as their first language. With a number like that, some might not be particularly motivated to produce tailored services for signing deaf people. But part of our work is to make it more cost-effective, so that people don't have the excuse not to do it.

"One of the things that's struck me during the time I've been doing this is the gap in people's understanding of the hearing-impaired community. There's a lot of misunderstanding, which leads to people not providing them with what they need."

It all started back in 1999, with an avatar called Simon the Signer. Simon translated text subtitles into animated sign language. It won two Royal Television Society Awards, but it was only a rough solution.

"Simon the Signer simply spotted important words and turned them into signs. So you're basically putting stuff out in the order it would be in English. However, sign language doesn't use the same order as the English language.

Professor John Glauert

"If you do it like that, you can certainly turn it into something that most signers can understand, but it's like turning a German or Spanish sentence into English without changing the word order. It doesn't look right."

The group originally started by looking solely at the linguistics side of the problem, but later created their own platform to animate the 3D character as well. The challenge was to balance two occasionally competing priorities: to make the signing movements as quick and fluid as possible, and to pass on the full meaning and nuance of the speech, often without any other visual means of communication.

By picking out words in order, Simon the Signer could translate quickly enough for the animation

Luna

to look fluid and natural. However, for the signing sequences to have the right structure, the system needed to know more about the sentence before translating.

Enter the TESSA project. TESSA was a speech recognition system designed to translate sentences and phrases into true British Sign Language by identifying a phrase as it was being said, and producing the corresponding signs with almost no delay.

Of course, it's a huge challenge to develop a system that can predict any sentence. There are so many variables. So TESSA was developed primarily for use in customer service situations, translating phrases spoken to customers by counter staff. In these situations, there are only a limited number of essential phrases that are likely to be said, so the system could spot them much quicker. In 2000, the technology was trialled by the Post Office in the UK.

Professor Glauert says: "The system has to recognise the whole sentence, but it can start making a pretty good guess midway through, and can come out almost straight away with an answer. It's not looking for every phrase. It starts, gets better information, and corrects itself."

The team also approached the issue of how to translate those meanings that weren't spoken. First arriving in 2000, the ViSiCAST system took a number of features of communication into account, such as eyebrow position, plural verbs, and the size and placement of gestures. In two more EU-funded projects called eSIGN and Dicta-Sign, the group has since fine-tuned algorithms that can deliver gestures that differ subtly in hand shape and location.

In order to achieve this, they used an established transcription system called the Hamburg Notation System, which tackled the sentences phonetically. This is converted into computer-readable XML, and then processed through a module that uses this information to manipulate the skeleton of an avatar character.

"The actual speech recognition element of our work is state-of-the-art, but not groundbreaking. The more challenging part is what we do with the animation. It's telling the system exactly how to move the bones of the fingers and arms to pass on a meaning.

"It's working on two parallel tracks: One communicates what the hand is doing and what the body is doing. The other tracks the face, eyebrows and eyes. For the mouth, we use an animation technique called Morphing, which carries a description of the mouth shape.

"You add a mesh with a skin and clothes over the top. So one of the great things about our system is that we can play the information back with different characters, from humans to robots, aliens and monkeys."

Their research has been applied in a number of different ways. IBM called their virtual signing technology "the most advanced and flexible system available", and integrated it into a real-time system called Say It Sign It. Thanks to a collaboration with Orange, that system was modified to work on mobile devices.

It has helped to translate pre-defined information into sign language, from train announcements to weather forecasts and warnings about avalanches. In conjunction with Action on Hearing Loss, the Visual Humans Group has built resources that have allowed others to create sign language content for websites, including Germany's Federal Ministry of Labour and Social Affairs and employment sites in the Netherlands.

There is also a valuable application in learning. Many children pick up words much more easily when there is a gesture associated with it, a process known as kinesthetic learning. With this in mind, a series of animated story DVDs have been released under the LinguaSign brand. They have been re-produced for English, Dutch, French and Portuguese. Following a 2013 trial of Key Stage 2 students at more than 50 UK primary schools, 62% of respondents confirmed that it improved a child's speaking and listening skills in a new language.

The avatar technology has already begun to supplement the interpreters we're used to seeing on TV. It has been showcased on Dutch programme Het Zandkasteel, and on online shows such as Wicked Kids. It has also been used by cultural heritage sites to help pass on stories from history using sign language.

Having almost mastered the hands, the team hopes to improve their work even further by improving their grasp of the expressive human face.

"It's not about dragging deaf people into the hearing world, but providing them with the sort of services we take for granted on their own terms", says Professor Glauert. "It's what they want, rather than what we think they should get."

A new approach to combatting disease has enabled Newcastle University spin-out *e-Therapeutics* to get a foothold in the pharma/biotech market, giving companies a new way to break down resistant bacteria.

11

Drug discovery is getting harder.

It is no secret that research and development spending by pharmaceutical companies has become less productive in recent decades. The Office of Health Economics estimates the average cost of developing a drug at £1.2bn. It can take around 10 to 15 years to reach the market. Many drugs fail clinical trials. And even then, there's concern that bacteria are becoming more effective at adapting to attacks.

In 2002, Malcolm Young and Peter Andras were among the co-founders of a university spin-out company called e-Therapeutics. Their perspective was that the industry needed to change the way it thought about combatting disease.

"The conventional way of doing it was to get a number of chemical compounds secreted by fungi or other bacteria, and see if they worked on a disease", says Peter Andras. "You'd look at molecules secreted by an organism that had antibiotic effects, and measure how effective they were.

"It was essentially a trial and error approach, but there weren't many other alternatives."

Peter Andras

Many traditional antibiotic molecules block the functionality of certain target molecules that play an important role in certain bacteria. The danger of this is that bacteria could adapt to familiar assaults, recruiting enzymes that break up the antibiotic molecules or render them ineffective through modification.

Instead, e-Therapeutics used algorithms to map the networks of interaction between proteins of bacteria and cells of their host organism. From there, it could identify targets; not just the network hubs and bottlenecks that drugs traditionally attacked, but other structurally-important elements that were previously ignored.
It also advocated combining drugs to hit more than one target simulaneously, making defence a much trickier task for the bacteria.

"If you attack unusual targets the chance of the bacteria working against them is much lower", says Andras. "If you can target two molecules rather than one, the bacteria needs to defend against two attacks at the same time."

This approach did not just open up new targets in the battle against complex diseases such as depression and cancer. It allowed drug companies to re-purpose drugs that may not have been effective on their own, by using them in combination with others.

Five years after being established, e-Therapeutics listed on the London Stock Exchange's Alternative Investment Market. By 2013, it was valued at £92.7m, and was the eighth-largest pharmaceutical/biotech company in the UK marketplace by market capitalisation.

According to the UK Higher Education Business Community and Interaction Survey, the average external investment in a university spin-out was around £700,000 per company. e-Therapeutics raised £18m in 2011, and a further £40m in 2013.

So how did it start?

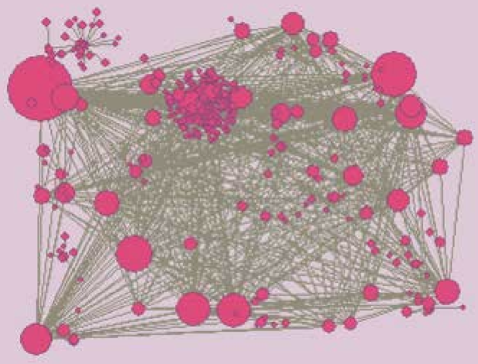

The earliest image representation of a protein interaction network

In 2002, Peter Andras joined Newcastle University's School of Computing Science as a Lecturer in Complex Systems. His expertise was in network analysis and its applications, and he worked with colleagues to develop algorithms that analysed how a bacterial protein interaction network was structured. The team then identified the right set of proteins to target, based on how much the protein representing nodes of the network contributed to the structural integrity of the network. For example, penicillin-binding proteins that are targeted by penicillin are structurally critical bottlenecks, while ribosomal proteins attacked by tetracycline are hubs in the protein interaction networks of bacteria.

Having developed an algorithm to pick out these targets, they put them in a workflow to help prioritise which drug candidates would be tested first.

Anil Wipat led the workflow development, working with e-Therapeutics as part of a Knowledge Transfer Partnership. The team also collaborated with the Faculty of Medical Sciences.

"Our most intensive research was from 2002 to 2007. We collaborated with biologists, and tried to develop the visualisation side as well.

"We published a few papers but mainly focused on patents. Our first patents were in 2007/8, and since that time we've added a number of patents in the UK, US and EU relating to this technology.

"The project managed to get a few extensions, and the company got some small amounts of funding from local investment funds. In 2007, the first major investment came in."

There are currently 20 e-Therapeutics employees located between sites in Oxfordshire and in Newcastle. Andras is not one of them. He retains a shareholding but began a new stream of research on network analysis in 2008. He is now at Keele University.

"My current work applies the same type of thinking to large-scale software systems. We look at automated tools that identify the right parts of the code to change in order to fix or modify the behavior of the software. We discover which methods or functions are most important, and which relate to certain requirements."

In the meantime, two drugs developed by e-Therapeutics are being evaluated in clinical trials. Anti-cancer drug candidate ETS2101 is in Phase 1 in the UK and US. The second is an anti-depressant called ETS601, which is beginning its Phase 2b trial shortly.

The contribution has not been forgotten. According to company CEO Malcolm Young, the work of Andras and Wipat at Newcastle University “has positioned e-Therapeutics as a world leader in the new science of network pharmacology drug discovery”.

P Aeruginosa

B Subtilis

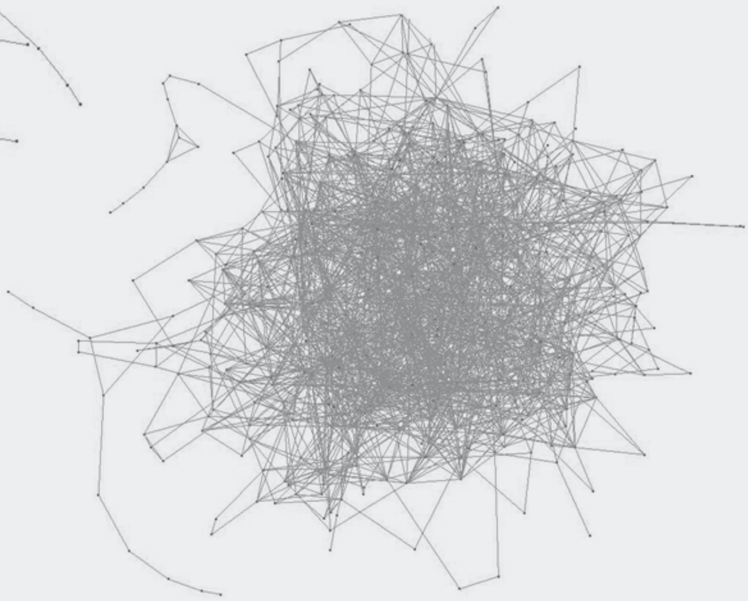

E Coli

CCTV systems are collecting information on people and places all the time. That's a lot of information. Researchers at Kingston University are working on intelligent systems that can help humans to monitor all these images, and they've even found their way into concert arenas and sports broadcasts.

12

While some say that CCTV “watches” us, in truth it’s sometimes just staring at a point.

If someone or something happens to move into that frame, it can be recorded and viewed. But what if you want something a little more intelligent and reactive?

In 2012, a prototype created by Ipsotek, a tech company in South West London, fought off Samsung and Panasonic to win CCTV System of the Year at the annual International Fire and Security Exhibition and Conference (IFSEC).

Its innovation? A system known as “Tag and Track”. Tag and Track allowed a user to identify an individual and follow them between multiple cameras by simply clicking on their image. The system could then search through its records to map out where that person had gone, or even make a prediction as to the route they might be taking. It was commercialised in conjunction with BAE Systems, one of the world’s leading defence organisations.

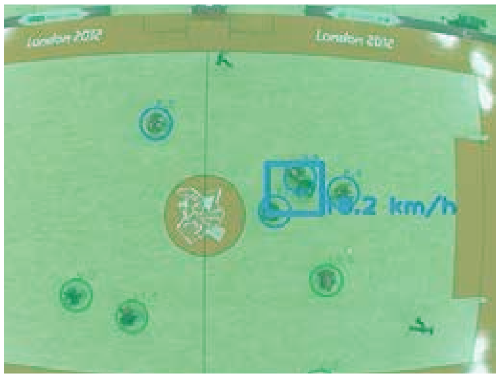

This is just one of the applications that has come out of Kingston’s Digital Imaging Research Centre (DIRC), a group that investigates how computer systems can be trained to track.

Did you make it to Led Zeppelin’s much-anticipated reunion show in 2007? The technology was deployed to help London’s O2 Arena record the size of the crowd filing into the venue. It also gave TV audiences watching the 2012 Paralympics a flavour of the speed and

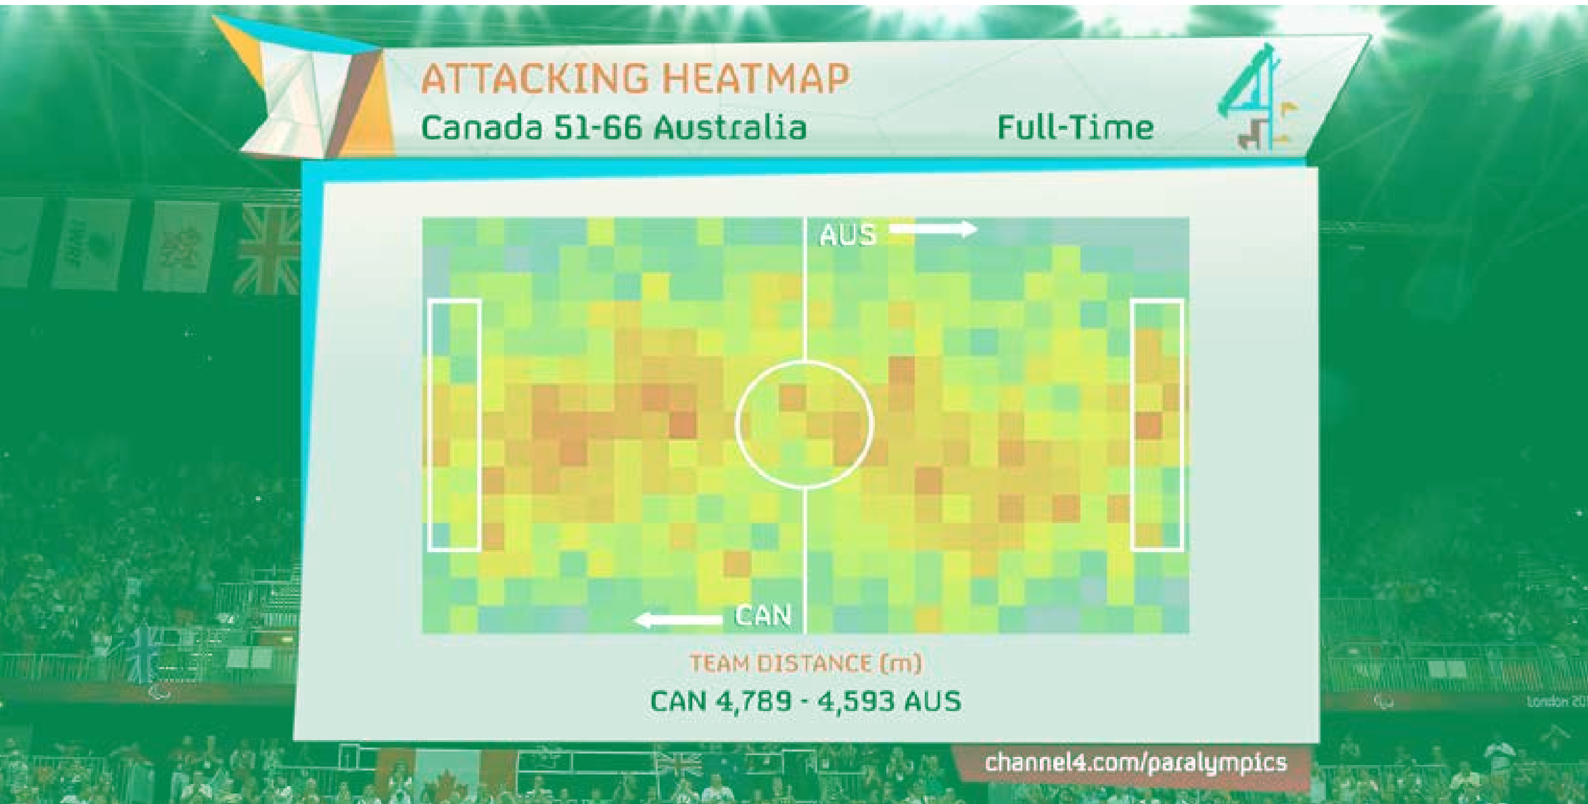

intensity of competition, measuring the speed at which wheelchair basketball and rugby players collided.

“The applications are in various fields, such as security, sports and commerce”, says lead computer scientist Dr James Orwell. “For example, you could be tracking a missing person, or modelling where people are going in a shop to help with layout design.

“This is a very fast moving area of research enabled by developments in modern computer technology. In the early ‘90s, it was very difficult to load even one whole image onto a computer. But in the last 20 years, we’ve seen an exponential surge in the power of computing devices.”

The Tag and Track prototype is one of a number of collaborations between DIRC and Ipsotek, whose customers include Transport for London, BP, BAA and the Australian government.

When the Digital Imaging Research Centre, at Kingston first formed in 1997, it was focused on how to discern interesting information from jumbled crowds captured on CCTV, especially across multiple cameras. In bustling spaces, it can be hard for humans to monitor a range of cameras on their own, so there is a need for systems to step in that can identify cues and threats with some level of intelligence.

Dr Orwell says: “Prof. Sergio Velastin did a lot of work on technology to alert transport providers to potentially-dangerous scenarios in their transport networks. If there’s a holdup in the flow of people and a couple of key points become blocked, it can quickly get to the stage where you need intervention.

“The research looked into systems that could identify anomalies, whether that was too much traffic or too little. For example, if there’s no one in a certain area, it might mean that a route leading to that area is blocked, or it can draw attention to an overcrowded part of the system.”

Kingston’s work on automatic video analytics was trailed in the field in underground stations in Paris, London and Rome. By teaching the system to identify flows of people and interpret what that might mean, operators were able to get useful information without needing to rely on the tracking of individual people.

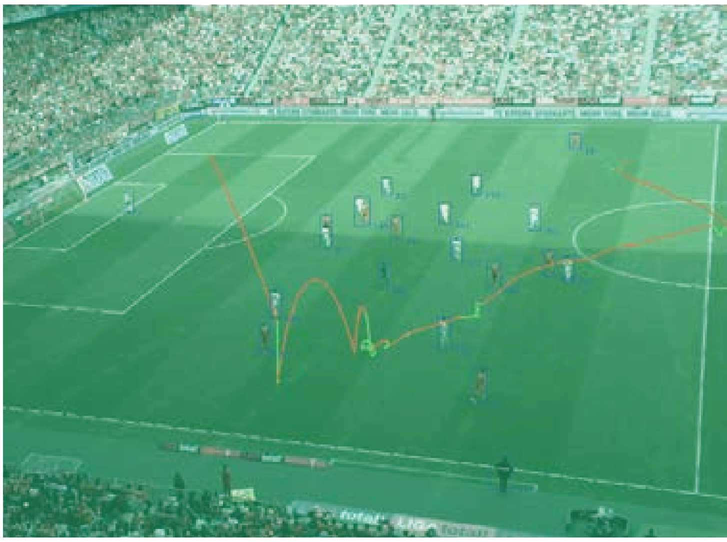

These systems formed the basis of the commercial products created by Ipsotek, which was co-founded by Sergio Velastin, a professor of Applied Computer Vision at the university.

While many systems searched for a single “trigger”, Ipsotek developed Scenario-Based Detection that would note when multiple events happened at the same time. So rather than flag up a vehicle every time it sat at an airport terminal for more than three minutes, it would do so when it confirmed that the vehicle was stationary, that it was in a particular area, and that it had been abandoned by its occupant.

In the meantime, work on individual detection was taking place. Working via a Knowledge Transfer Partnership with Ipsotek, researchers such as Professor Tim Ellis and Dr Dimitrios Makris shared their findings to help boost the capabilities of their technology in people tracking and counting. The project also benefitted from the hard work of Sateesh Pedagadi, a DIRC PhD student who also works at Ipsotek.

To track an individual, the system must be able to match a single object across multiple cameras, and differentiate them from passing people and background objects.

Dr Orwell says: “The difficult bit is determining what is caused by people moving around as opposed to shadows and windswept trees. This segmentation of ‘foreground’ from background is a core technology for many applications.

“Our approach uses optical tracking as opposed to RFID tracking, which would require tags. The advantage with optical is that it’s completely passive and non-invasive. But there are some real challenges, such as when crowds come together or it starts snowing.

“Our background models have now become much more sophisticated. Different climates can be learned, so a system can recognise a background is the same whether it’s sun or shade. It’s also using observations from different cameras to develop a clearer understanding of who’s where.”

This technology has been used in situations where crowds are particularly daunting. In 2010, Ipsotek deployed an “anti-backtrack” system which monitored the flow of passengers through airports including Edinburgh, Birmingham and Bristol. And it helped to count capacity crowds into the O2 Arena during the Led Zeppelin performance and the 2011 Prince concert. This footfall data was used to maintain safety, as well as calculate the amounts charged to promoters and advertisers for future events.

The intrusion detection technology developed for the London Eye demonstrated the system’s ability to monitor approaches on water as well as land.

Dr Orwell says: “This system was working in areas that included open water, as obviously the Eye is on the banks of the Thames. This showed that it was improving, and could cope with more complex moving environments as well as static land.”

Sport has also been an intriguing application of the technology. Initially, DIRC researchers set up cameras at football stadia to see if they could effectively track the positions of players. DIRC went on to work with Premier League football clubs to gather valuable information on performance, tactics and fitness.

In 2012, DIRC’s statistics were being beamed into living rooms around Britain during the Paralympic Games in London. It was commissioned by global sports business company deltatre to generate graphics that were seen during the coverage of 42 wheelchair rugby and wheelchair basketball matches on Channel 4.

Dr Orwell says: “We developed a system that was tracking individuals and collisions, and marking the speed at which that collision occurred. The production system allowed a graphic to be produced that could be overlaid on-screen by a director in a couple of seconds.“

DIRC has continued to develop its interest in sport, particularly in football. A project designed to accurately track all players on a football pitch began trials earlier this year.

As technology continues to improve, Kingston’s researchers are carrying on their work to develop systems that can accurately and promptly identify important events as they happen.

Dr Orwell says:

> “We’re talking about improving the effectiveness of algorithms and, ultimately, getting near or surpassing human performance.”

“If you’re trying to identify people across multiple cameras where you’ve got a constrained area and you’re walking from A to B, that works nicely. But if you have no idea where people are and when they’re coming through a door, it’s a much tougher nut to track. Looking at the future, the emergence of even higher resolution cameras can make detection even easier.”

The Internet isn't just about text and images anymore. As streaming video and audio become a big part of society's lives online, the world is benefiting from the standards first laid down in the mid-nineties.

13

Think about the many different ways in which you communicate. Now remove the ones that weren't available to you 20 years ago.

Developments in communications have emerged at a breakneck pace over the last decade or two. We can now have international video conversations on our desktops and mobiles. And that's not all: Our new TV's are connected and streaming video in real-time, while our music is played online rather than on CDs.

At the heart of this is a pair of international standards. The Real-Time Transport Protocol (RTP) helps deliver audio-visual data across a network. Meanwhile, the Session Description Protocol (SDP) describes the characteristics of the data and where it is headed. These standards were developed by the Internet Engineering Task Force (IETF), and are used by billions of devices worldwide.

The University of Glasgow's Dr Colin Perkins has played an integral role in improving these standards, and establishing a single standard framework for interactive real-time multimedia as opposed to several competing ones.

This work has improved the quality, speed and reliability of the media we digest online, paving the way for game-changers such as WebRTC, Apple's Facetime, and SIP phones.

"RTP used to be used by no one but a few researchers", says Dr Perkins.

"It was the protocol people settled on for interactive applications and multicast TV distribution. RTP was developed in the open, published in the open, and got a lot of traction from the research side. It's just spiralled since then. I guess we just happened to be in the right place at the right time."

Back in 1992, a demo took place in which audio was piped to 20 sites on three different continents, using tools built on top of a research network called DARTnet. By that time, PCs and workstations had sufficient power to grab and playback audio and video streams, while users could transmit real-time data to multiple people.

The tools that enabled that 1992 broadcast were known as MBone - or "multicast backbone" tools. RTP was one of the protocols used on the MBone.

While developing RTP to control the transportation of media, SDP was also created to fulfill a number of functions, such as telling sending and receiving systems what compression to use, and indicating the end-points of the media stream.

Dr Colin Perkins

Dr Colin Perkins was working on the development of Voice over IP tools that would make online telephony possible.

"I got involved with RTP in 1996", he says. "I was building tools on MBone, trying to make it work better. As a result I got involved in the standards and trying to make them more interoperable.

> "The internet in 1996 wasn't particularly widespread. It wasn't like today where people can access it on their home PCs."

Internet Explorer / Yahoo in 1996

Becoming the co-chair of the IETF's Audio/Video Transport Working Group in 1998, Perkins helped develop RTP into an accepted Internet standard used worldwide. Glasgow's research has helped advance the performance of multicast video, boosted congestion control and rapid synchronization, and it has assisted with standards that have improved the transport, stability and security of media online. Dr Perkins also edited a significant revision of the SDP standard while we was co-chairing the IETF's Multiparty Multimedia Session Control Working Group.

"The research was always aimed at solving a particular problem, such as making things more adaptive, avoiding network congestion if it gets popular, or understanding the behaviour of the network", says Dr Perkins.

"In terms of the standards work, it's very much about making sure things happen fairly, and are open to the greatest extent possible. There are a lot of companies with competing interests, trying to get their particular protocols into the standards, but having academics involved from the beginning was a help."

From 2007 to 2012, the University conducted tests on the limitations of networks to be used for Internet Television systems, demonstrating that it was possible to scale to the levels required. It also developed techniques to improve multimedia reliability and quality. Both projects were funded by Cisco.

Research funded by Ericsson involves developing a "circuit-breaker algorithm" which will allow high-quality video conferencing to be deployed widely and safely.

"There are probably three parts to this", says Dr Perkins. "The first is making the system safe to roll out. If we deploy very high-quality interactive video everywhere that's a tremendous load and we have to make sure the network can cope.

"Then there's delivering new types of application, and getting things into browsers with a much better Internet Protocol. Who knows what new peer-to-peer interactive video and data apps we'll end up with in future?

"The third thing is making it secure. The network is under constant surveillance and attack by those who would seek to eavesdrop on, disrupt, or subvert, phone calls and video conferencing, and we need to constantly review and improve the protocols to protect users."

Since 2011, Glasgow has been contributing to groups that are developing the WebRTC framework, to bring native video conferencing features to web browsers, without the need for Adobe Flash or Google Talk plugins. As well as helping to improve the quality and reliability of what we see and hear online, the presence of a non-corporate, neutral voice has enabled the standards to remain more fair and open than they would otherwise have been.

“The work we do in the standards community should come from a neutral point of view”, says Dr Perkins.

“Net Neutrality - for example - is a very difficult, multi-faceted problem. A lot of people are interested in network services; in prioritising and de-prioritising different types of traffic. There are good and bad points to that. If you’re using an internet-based telephone to ring 999, you obviously want priority. But for an Internet service provider to use the same system to give priority to their telephony or streaming service over a competing third-party service, might be considered anti-cooperative by some.

“In terms of the standard itself, it seems to be sustaining pretty well. It’s over the critical mass. That’s not to say that we couldn’t design something better, but would it be so much better that we should replace every smartphone in the world?”

Imagine you’ve got a very old system that is performing a very important function. **Do you risk updating it, or leave it alone?** Academic research has helped to create a company that enables the world’s largest businesses to save money, improve their processes, and reduce the risk of catastrophic system failures.

14

You may not realise it, but at some point in your day you may have interacted with a system that uses Assembler code – a low-level programming language specific to a particular computer architecture.

Assembler was widely adopted by large companies who needed to manage complex mainframes, from banks to government agencies. Of course, time passes and systems evolve, and many firms start to consider a change.

In the past, transforming reams of Assembler code into more efficient C code was a nightmare. Drawing on decades of research by Dr Martin Ward at Durham and De Montfort Universities, Software Migrations Ltd (SML) offered a less gruelling solution.

"There's a lot of assembler out there", says Dr Ward. "Every time you do a credit card transaction, it probably goes through an IBM mainframe running code in assembler. It's doing a lot of important work, but the assembler programs cannot be moved to new computer architectures without recoding.

"However, it's quite an esoteric language now. Students don't get taught it in schools, and it's more difficult to find people who can maintain your code.

"A lot of companies have these core algorithms that run their businesses, and nobody really knows how they work and what they're actually doing.

"Say you're selling holidays and you want to introduce new deals and discounts, that might involve updating your software. If you can't do that, your whole company gets held back. If you can migrate that code to a high-level language, it's understandable, maintainable and you can build all of your modern processes on top of it."

However, migrating code from Assembler isn't a picnic. A frequently-referenced quote from an Allen Eastwood paper in 1992 observes that software reengineering is "about as easy as reconstructing a pig from a sausage".

"When you're transforming it without maths behind you, it's very easy to put bugs in the system that don't show up until the end", says Dr Ward. "Even if you're as much as 99.99% correct, the chance of it still working at the end is only 30--40%. The only way to get

Dr Martin Ward

to 100% is to have a mathematical proof of correctness.

"In addition, it's a big project. For example, a company invested 40 man years of effort in re-writing 750,000 lines of Assembler code written by hand, and then had to give up without producing a useful result. We managed to migrate the system automatically in under 18 months with a team of four, and not all of them were working full-time."

Dr Ward's solution dates back to the 1980s, when he was working on a theory of "program transformation". This area of research focused on how to transform a program while still proving that the end product was semantically equivalent to the original, even if it took different steps.

"I was at Durham at the time, and we had a visit from a man from IBM", he says. "They had funding for blue-skies research, and we talked about taking the theory and applying it to Assembler. That turned out to be a lot harder than I thought as there were a lot of features to analyse."

By the time he left Durham for De Montfort in 1999, he had developed an initial version of the FermaT transformation engine, which forms the basis of SML's software and services. At that point, SML employed around six staff.

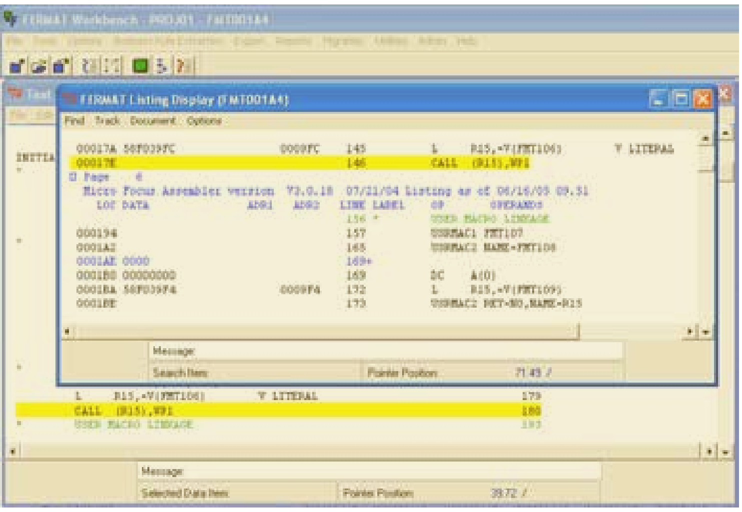

Dr Ward continued his research into the theory and application of "program slicing", which breaks down a program into chunks to check that each transformed element performs the same function as the original. For example, a "backwards slice" investigates how a variable came to a particular value, while a "forwards slice" looks at what would happen if the program were changed.

In 2001, SML was bought by Mike Dowd, who played a key role in commercialising this research. Between 2004 and 2006, Dowd invested more than £5m to help SML's research meet the requirements of industry, which included being able to successfully convert Assembler into Cobol, and later Java. The tools reached the marketplace in 2007.

SML has worked with a number of globally-recognised companies to address their "Assembler problem", and is working with the world's largest user of Assembler to transform more than 10 million lines of code.

However, it isn't just the transformation that proves challenging.

"The big problem is often the politics that bubbles away in big organisations", says Dr Ward. "Our biggest competitor is the temptation to do nothing, and to live with the existing system for another year until the CEO retires. In my view, that's really short-term thinking.

"Nowadays, PCs and workstations are very cheap and powerful, but if a system is written in IBM Assembler code it only runs on an IBM mainframe and that can be expensive to maintain. If you can take some of your processing off a mainframe and onto a PC, you can save hundreds of thousands of pounds."

"We've seen all kinds of horror stories involving Assembler. It can really hold you back when you're implementing enhancements or adding new functionality to a system. We even heard about two companies who were about to sign on the dotted line for a merger, but found that their two Assembler systems were so incompatible that the whole deal fell through."

SML remains the world leader in deciphering and migrating Assembler. The work has helped to fund three PhD studentships at De Montfort, where Dr Ward's team is looking to break even more ground.

"We're looking at how you take a procedural program and turn it into an object-oriented program", he says. "There's a lot of interest in Java now, and you can generate procedural Java, but the holy grail is be able to produce object-oriented Java as well, with pieces of code and pieces of data that operate together.

"I'd never have imagined all this when I was doing my PhD. I remember the first time we put a piece of Assembler through. It was about 200 to 300 lines long. The UNIX workstation chuntered away, and three days later it fell over in a heap trying to translate one module. These days it takes around seven seconds on average per module. That's down to the efficiency of the transformation process, and the quality of the machines running the software.

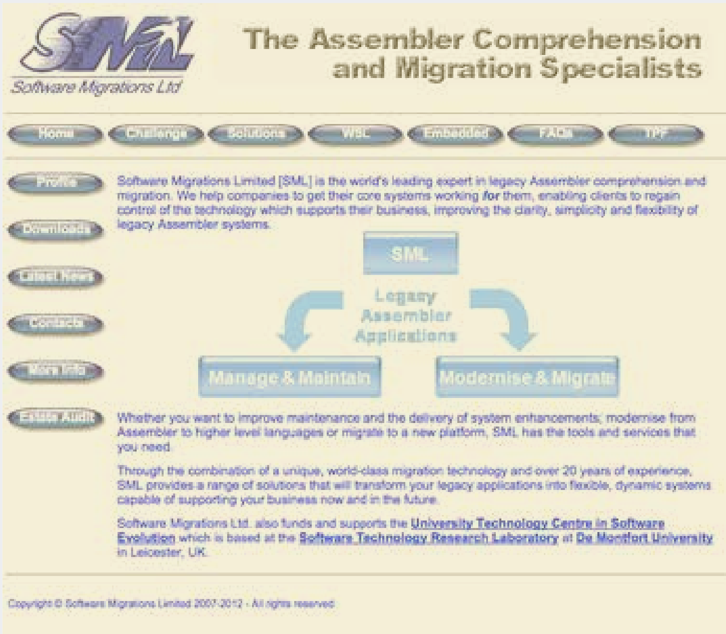

smltd.com

"That's what research is all about. You push the boundaries of knowledge, and sometimes you don't know what's going to have an impact and what isn't."

Are we ready to give robots the power to kill without human supervision?

Longstanding robotics expert Noel Sharkey thinks not, and he's taking the fight to the United Nations.

15

What should we do about killer robots?

This isn't a debate to be shelved until we reach some distant sci-fi future. As we delegate more and more tasks to robots, it's time to ask ourselves: Are we comfortable with giving them the power to decide when to kill?

In April 2015, a multilateral meeting was held at the United Nations in Geneva. Representatives from the Campaign to Stop Killer Robots hope it will be another step on a path towards a ban on "lethal autonomous weapons systems".

One of the leading campaigners pushing for this ban is Noel Sharkey, a Professor of Artificial Intelligence and Robotics at the University of Sheffield.

An experienced robotics researcher, Professor Sharkey is now focused on their ethical application in fields from health to education and the military. His expertise and campaigning have helped to shape the debate over the issue, and he has met with several military figures and governments to persuade them to take a stand against autonomous machines with the power to kill.

> "I have a real problem with autonomous weapons systems. I'm most concerned about the decision to kill being delegated to a machine, giving it the power to pull a trigger without human intervention."

Professor Sharkey has spent more than 15 years researching robot learning and behaviour. During this time, he has worked with colleagues on technical solutions for problems such as how to locate robots quickly and reliably, and how to help them avoid obstacles. He started researching the ethics behind the use of robots in 2005, and began speaking out against Autonomous Weapons Systems in 2007.

Professor Philip Alston - the former UN Special Rapporteur on Extrajudicial Executions - says that "because of the depth of his technical and scientific expertise, his work is an indispensable reference point for those working on these issues".

Professor Sharkey's concern lies with the development of systems that can identify a target and act on it themselves, without a human hand in the process. As part of his fight against this new development, he has appeared in print, radio, television and online over 1,000 times, in more than 50 countries.

Noel Sharkey

Matrix – with Dinka Dumicic (left)
Jody Williams (right)

He believes that these robots do not meet the standards for war set out in the Geneva and Hague conventions, as they will be unable to acceptably distinguish combatants from non-combatants, and the loss of life or damage will outweigh the military advantage gained.

He is also concerned that autonomous robots would be more vulnerable to malfunctions and cyber attacks, and that they would be more likely to strike illegitimate targets.

In a 2007 article in the Guardian newspaper, he said: “In reality, a robot could not pinpoint a weapon without pinpointing the person using it or even discriminate between weapons and non-weapons.

“I can imagine a little girl being zapped because she points her ice cream at a robot to share. Or a robot could be tricked into killing innocent civilians.”

Professor Sharkey has published several pieces in science and engineering journals, analysing whether robots have the cognitive capability to make lethal decisions. He has used his technical knowledge to explore questions of autonomy, and has argued that the failures that currently affect drones are likely to be even more pronounced with autonomous weapons systems. This work has been cited in academic publications more than 185 times.

His research led him to co-found the International Committee for Robot Arms Control in 2009. The ICRAC gathers experts in robot technology, ethics, international relations, arms control and human rights law to push for limits on the use of robot technology. Work with non-governmental organisations such as Human Rights Watch culminated in the launch of the Campaign to Stop Killer Robots in 2013. Human Rights Watch describes him as “one of the most effective spokespersons” in an international coalition of more than 40 NGOs, including Amnesty International, the International Peace Initiative and the Nobel Women’s Initiative.

Opponents to his line of argument have suggested that safeguards are preferable to a ban, and that the use of such machines could be restricted to non-human targets. However, Professor Sharkey counters that this approach is a “foot in the door” that can be opened further all too quickly.

Sharkey has already expressed concerns over the British unmanned supersonic super drone Taranis, which is pencilled in for introduction alongside manned combat aircraft after 2030.

“It might start off slowly, but if you look at aerial bombardment, people like Roosevelt initially tried to get treaties to stop that. Submarines were once completely against the rules of war as you were meant to show your colours, but they were so useful that eventually everyone adopted them.

“I’m convinced that if these things are in development, they will be used.”

ICRAC at the UN

As a result, he has been active in trying to influence policy-makers internationally. He has addressed senior military figures in 26 countries, and his research is referenced in officer training materials in the UK, US, France, and the Netherlands.

He has briefed UK parliamentarians on several occasions, including a cross-party briefing of around 20 MPs and peers in April 2013. He has also addressed groups in Germany in France, and was cited extensively in the 2013 European Parliament policy document on the use of unmanned robots in warfare.

Partly as a result of such efforts, the UN's Convention on Conventional Weapons agreed to meet last year to discuss such lethal weapons. Their second meeting this month is a continuation of that discussion.

"A lot of people in the military support us", says Professor Sharkey. "They don't want to see these weapons, as there's an accountability gap. Who is responsible if something goes wrong? Is it the commander, or the creator of the software?

"In fact, the UK is the only country that has taken a definitive stand on not banning them. The argument from our Ministry of Defence is that it has no intention of having autonomous weapons, but they will not support a moratorium or a ban."

If the April 2015 talks progress well, the next step in the process would be the establishment of a Governmental group of experts to discuss the best course of action.

"The process could last another four or five years. Maybe it could take longer, but we hope not. We'll have to wait and see.

"As an academic it's been fascinating. I wrote a lot of papers and gave a lot of talks and people were very interested, but I didn't know what the best course of action was. As soon as Human Rights Watch and other major Non-Governmental Organisations got involved, they knew exactly what to do and put me in touch with ambassadors and into diplomatic circles.

"We're now very much part of a large international civil society movement."

For more than 20 years, Aberystwyth University researchers have been investigating **teaching robots how to understand the world.** Their work has had a huge impact on helping global car companies to improve safety, efficiency and design.

16

Things go wrong.

It's a fact of life, and it's why testing matters. And when you've got something as complicated as a modern car, it's crucial to think carefully about how to minimise the effect when failures happen.

For many years, the process of checking through a car's electrical circuit was done by hand, by trained engineers. It could take weeks.

"They'd have to look at every wire in the system, and work out what it would do to the car if it broke", says Professor Mark Lee of Aberystwyth University. "They'd grade each one according to its seriousness. It was a very laborious process."

In the early nineties, a team at Aberystwyth University started looking at a more efficient option. The university's Advanced Reasoning Group was investigating how you might create systems that could reason about and automate the analysis of physical systems. In 1990, researchers started working with Jaguar to apply that reasoning to automotive design.

The result was a system that could do the analysis in hours rather than weeks. It also enabled companies to test their designs each time they improved them, boosting safety and efficiency.

Former Aberystwyth graduate Chris Price joined the university in 1987 after a spell in industry, and went on to play a leading role in a spin-out company called FirstEarth.

He says: "Cars seemed a natural application for this. Before this, you would use a diagnostic system in which you would write down all the possible problems, how serious they were, and what the solutions were. That's not exactly how engineers work.

"When something fails, an engineer has an underlying understanding of how cars work, and traces circuits to see what the problem might be. What this system is doing is looking at an issue and reasoning that there are 17 possible things it could be, and here are a few more tests that can be done to narrow it down."

The group's breakthrough involved applying an approach called Qualitative Reasoning to circuit design and testing. Instead of gathering a numerical reading such as a voltage from a wire in the circuit, the system is taught to apply a status label to it instead.


The FirstEarth team get into the Christmas spirit (Chris Price far right)

“When a repairman comes to mend your washing machine, they don’t go in and measure the voltages and currents”, says Professor Lee. “They look at the symptoms and figure out at a coarse level whether something is alive or dead.

“We started building circuit reasoning software that would do this. You pass around symbols that report whether a certain element is alive, dead or partially alive. In that way, you can very quickly analyse the state of a system as a whole.”

By 1993, the group was working on developing this idea with Jaguar and Ford, supported by funding from the EC and EPSRC. The system was originally called Flame, but their automotive partners weren’t exactly sold on the name.

“There were two guys called Steve Twitchett at Jaguar and Steve Leedham at Ford, whose job involved spending months analysing complex systems”, says Professor Price. “We decided to call the product AutoSteve, as it automated what Steve did, and gave him the time to go off and do other important things.”

The industrial version of AutoSteve was built by 1997 by a team that included Price, Neal Snooke and John Hunt, and Price set up FirstEarth in 1998 to market it to the wider industry. He took a year off from the university to get the company off the ground.

The software was awarded a Ford Europe Technology prize in 1999 for its impact on safety as well as cost efficiency. FirstEarth received some venture capital support in 2000 to allow it to develop its plans. Price returned to the university while maintaining an interest as Company Chairman.

Since the time taken to test a system effectively had been slashed, companies found themselves able to test earlier in development, and repeat the process more frequently.

“They found they could test the design before they built the car, and then again as they re-designed parts of it”, says Professor Lee. “They could also use it for diagnosis. It could also be applied to ‘sneak circuit faults’, where the circuit finds another route it shouldn’t be using when components are in a certain state. A computer can run through all that.

“At the time, it was a new approach. Qualitative circuit reasoning had not been applied to circuit analysis in this way.”

In 2003, the company was bought by Mentor Graphics, which has now installed AutoSteve with more than 200 automotive and aeronautic suppliers worldwide.

The benefits were clear. Ford estimated that the software had helped them to save nearly £14m a year. In one case, it enabled designers to cut around 35 pence from the cost of each automotive harness, which led to total savings of more than £1.7m over the cost of a whole production run.

FirstEarth employed 12 staff at its height, with nine trained by the university’s Advanced Reasoning Group. Six of those continued to work for Mentor after the buyout, while four are still key staff at Mentor’s Newbury research and development centre.

“At one time, it was the first or second-most popular software in their entire catalogue”, says Professor Lee. “It’s a great example of blue-skies research that found an application that had a real benefit.”

Under the terms of the buyout, Price and the team were restricted in how they could apply their expertise for three years. But in 2006, they began a project with BAE Systems, which involved helping to analyse safety cases for unmanned aerial vehicles in commercial airspace.

"We developed the process further to look at the modelling of other engineered systems", says Professor Price. "It was particularly useful, as you could build models of the system, and interrogate them."

Professor Lee's research now focuses on cognitive robotics, and how you help machines understand the world. Having applied the group's methods to hydraulics, Professor Price and Dr Snooke are now investigating using qualitative reasoning to reason about more complex systems, such as the effect of changes on how cities operate.

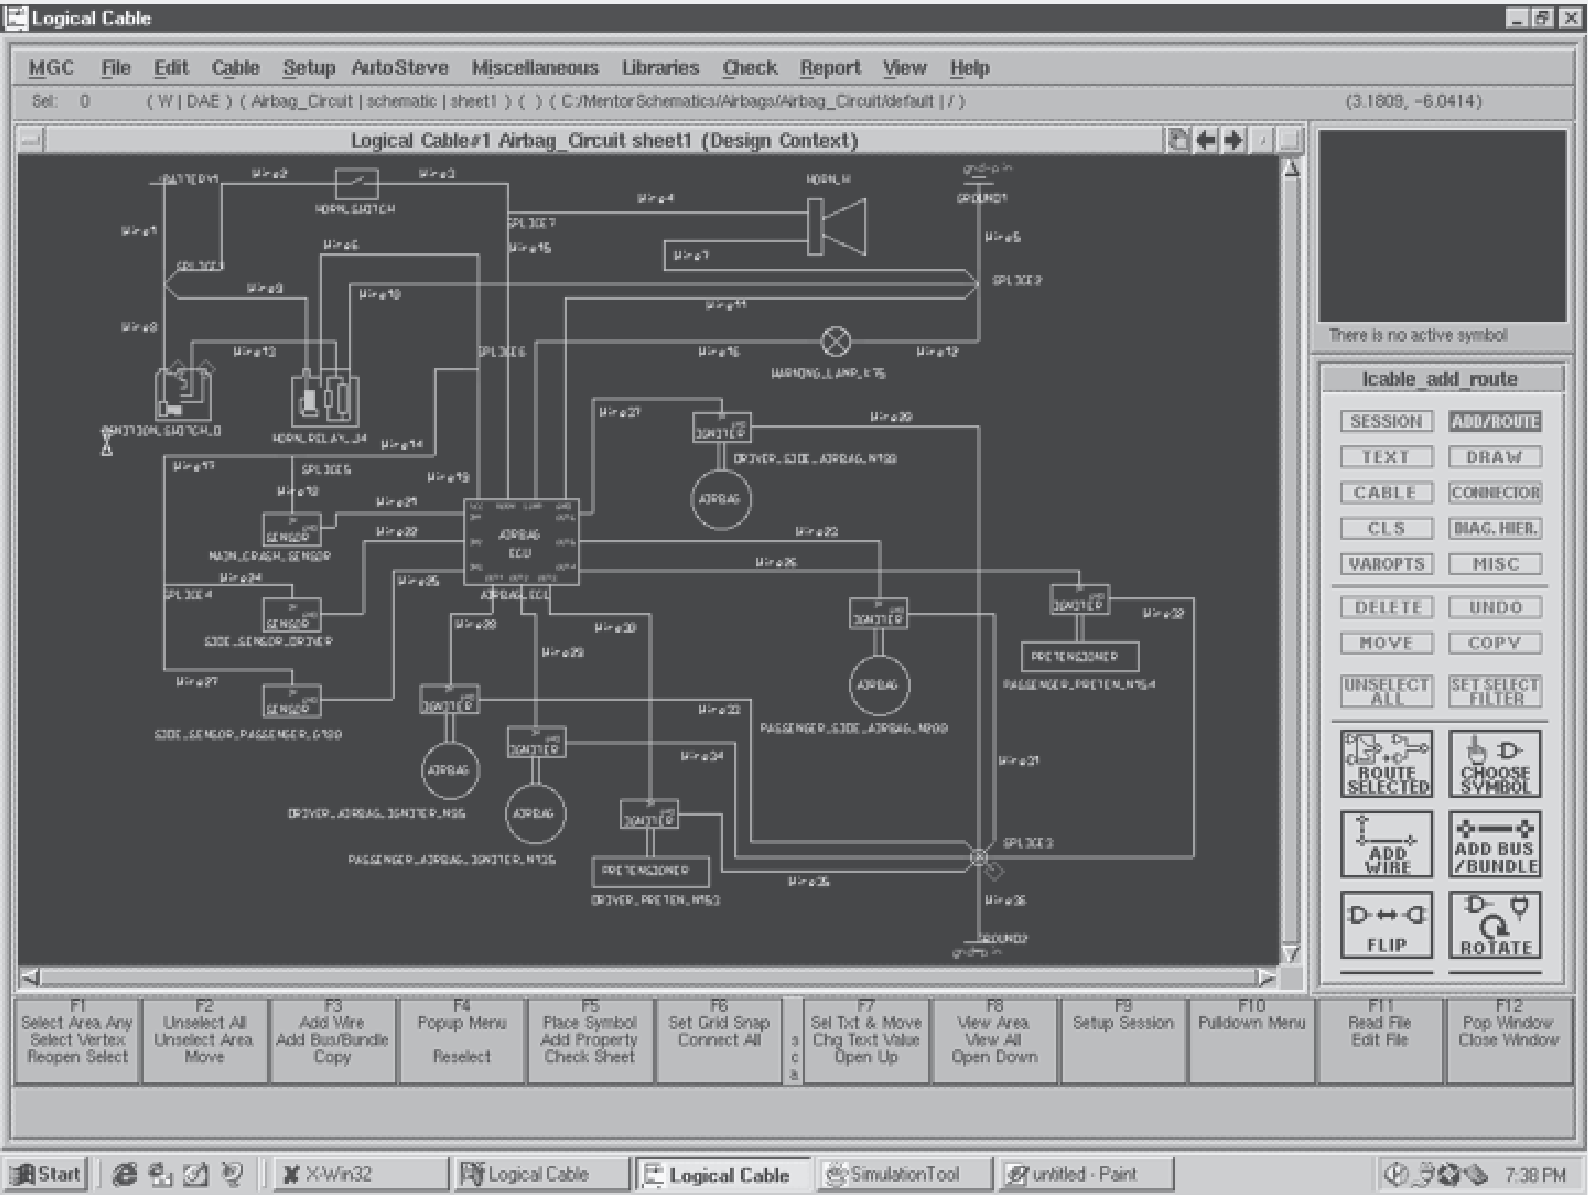

17

Putting staff in the right place at the right time is a big challenge, and it's even more difficult when you're responsible for a transport company. Researchers at the University of Leeds have spent decades perfecting solution methods that can handle the complexities of **crew scheduling**, and have created a multi-million pound international company as a result.

A successful business called Tracsis has spent the last few years helping bus and rail firms to construct their crew timetables effectively, saving them over £230 million between 2008 and 2013.

Its software is now being used by 70% of the train companies currently operating in the UK, as well as the UK's largest bus operator. And it all began at the University of Leeds.

The university's work on crew scheduling dates back to Anthony Wren from the 1960s onwards, but it was largely focused on bus scheduling until the 1990s. At this point, the UK Government began privatising the railways, and organisations such as the EPSRC funded Leeds researchers to tackle the more thorny issue of rail crew scheduling.

A major breakthrough in their work is PowerSolver, an algorithm which is able to crack scheduling puzzles that were previously considered too complicated and daunting.

"In running any public transport service, there are two main resources: Vehicles and manpower", says Dr Raymond Kwan, who has played a key role in the work for 25 years. "You have to schedule both resources to meet the timetables, or it doesn't work. It's complicated logistics. You have to devise a plan so that a train will always have someone to work on it, and the workforce will always be in the right place at the right time. And then - on top of that - it has to run like clockwork."

Before Tracsis software rolled out in 2004, most train companies organised their crew schedules manually. This was not a simple process. A working timetable needed to be effective without requiring an unsustainable increase in staff, had to meet limits on shift length, and accommodate the required break times for crew members.

Even state-of-the-art systems could only cope with the complexity of crews of 100 or less. Most companies needed to schedule more than 150, or sometimes over 300. The complexity increased further when the transport network became denser.

"It's an extremely difficult mathematical problem", says Dr Kwan. "We can formulate a mathematical method to solve the problem,

Raymond Kwan (middle) with post-doc Research Fellows Linus Lin and Eva Barrena but maths has limitations on how complex it can be.

"We wanted the most optimal results, but we also want to harness the maths in such as a way that we don't overburden it. We adopt a similar approach to taking a picture on a camera. It could be tens of millions of megapixels, but when you're sending it to your friends it's more compressed while retaining as much of the image quality as possible. To churn out a result, we might compress the problem to a smaller size, deciding what detail to remove. Then we pick out the important characteristics and do another run, and then another. After all of these cycles we end up with a good set of data points."

Bus and rail companies were trialling algorithms developed at Leeds as far back as 1994 to 1996. In 2004, Tracsis was founded to commercialise the software, and the company has grown into an international operation listed on the London Stock Exchange.

Early customers included ScotRail, First Bus and Translink, and others signed up following trials between 2004 and 2007. The company was floated on AIM in 2007. As of 2013, it employed 200 full-time equivalent staff in offices in Leeds, Derby, Tadcaster and Australia, and had market capitalisation of £46.7 million.

"We have worked with a lot of train companies over the last 20 years", says Dr Kwan. "With this software, we were able to show them that we could save them 2% to 12%, even compared to the most experienced schedulers."

"A big advantage is that by using a computer, the runtime is a lot shorter and can be done within a half hour, as opposed to a manual process that might take several weeks.

"Since the company started in 2004, the platform commercialisation side has become very mature, and there's a team of developers working on it. The company has expanded a lot. It started with just four people at the beginning.

"We tackle a problem where there is a high potential for savings. The return is high and we produce excellent results, which is why the research has survived."

Transport companies appreciate Tracsis software as it gives them the opportunity to be more efficient with their operation. For example, in December 2008, Virgin West Coast managed to increase their train services by 30% following modernisation despite only adding six crew members to their roster of 1,600.

It also caters for rescheduling at shorter notice. Tracsis software helped to schedule adjusted train services when the Rugby World Cup came to New Zealand, and it can also respond to the demands of large-scale sporting finals, engineering works or weather disruption.

It has been used by many rail contract bidders to prepare cost-effective and efficient bids.

Dr Kwan says: "Even though they sometimes have to do a bit of work to prepare the data and turn the results into a form they can use, they don't mind as the savings are so great."

Tracsis' reputation in the industry has led to their business being employed to deal with other performance issues, such as safety information management, data monitoring and passenger count surveys. Meanwhile, the team at Leeds has since received an ESPRC grant for a three-year project investigating scheduling for train vehicles instead of crew.

"That's probably even harder than scheduling crew", says Dr Kwan. "It's a very complex and interesting problem, and it's one we're working very hard to conquer."

Out in space, spacecraft monitor our planet and send intriguing whispers from others. But how do the various parts of a spacecraft talk to each other? While there were once as many methods as there were companies producing spacecraft, a standard developed by the University of Dundee is now used by space programs across the world.

18

Almost a hundred satellites, probes and other spacecraft are launched each year.

They are dispatched to monitor the Earth, explore other planets in our solar system, study the wider universe, and provide commercial communication services.

Most spacecraft are unmanned, and need a reliable and effective way to handle data, store it, process it and move it around the spacecraft from one unit to another. In the past, different organisations had different systems for achieving this, so communication and compatibility was a major headache.

Then came SpaceWire.

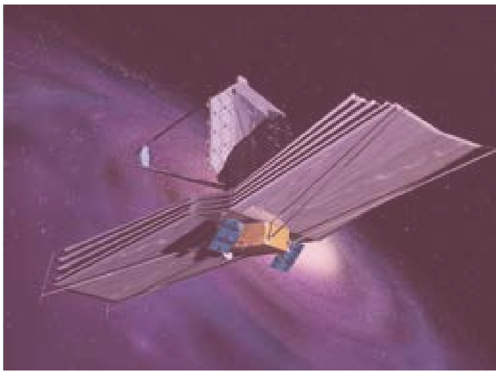

James Webb Space Telescope (JWST)
Photo credit: NASA

Photo credit: European Space Agency

SpaceWire is a standard for high-speed links and networks on spacecraft, connecting instruments, mass-memory and other on-board systems in a similar way to USB connecting peripheral devices to a computer.

This transformative standard is the result of the combined work of spacecraft engineers, the European Space Agency (ESA), and academics from the University of Dundee.

"SpaceWire is an on-board computer network, used for connecting electronic units", says Steve Parkes, who began work on spacecraft data network technology back in 1998. "It provides the nervous system of a spacecraft".

Sentinel – Photo credit: European Space Agency

"A typical spacecraft will have telescopes or other instruments looking down at the Earth, mapping the surface of another planet, or directed further out into space. SpaceWire enables spacecraft to gather data from these instruments to pass it efficiently to the on-board storage and processing systems, and for it to be sent to the radio downlink unit for transmission to Earth."

Dundee was a key player in the development of SpaceWire back in 2003, with funding from the ESA. It is now widely used in the industry, and is integrated in many current space missions run by the ESA, US National Aeronautics and Space Administration (NASA), Japanese Aerospace Exploration Agency (JAXA) and the Russian Federal Space Agency (Roscosmos).

Its popularity lies in the fact that it is simple, operates at high-speed with low power consumption, and has small implementation cost. It is also designed with compatibility in mind, which means units designed by one organisation can communicate with other units designed by another group or space agency. An instrument or other unit designed with a SpaceWire interface can be used on several missions across different space programs.

"SpaceWire is an open standard", says Parkes. "Anybody can use the technology to develop spacecraft on-board networks. At the University of Dundee, the main aim was to ensure that the standard was used as widely as possible. It's now being used by all of the world's space agents and most of the international aerospace industry. If someone in China makes an instrument to go on a spacecraft, they can fly it on an American spacecraft as there's a common interface."

The lack of an appropriate shared standard was a problem before SpaceWire. Many different communications links ended up being used on spacecraft interfaces, and these ended up adding to the mission cost and to the integration and test time. Attempts were made to use commercial technology like FireWire (IEEE 1394) and IEEE 1355 used in Science laboratories, but they were flawed in various ways.

Work had progressed towards a suitable technology but there were still a number of difficulties. At this point, the University of Dundee was given a contract by the ESA to work out some of the kinks in the nascent SpaceWire standard. Appropriate technology was designed, resulting in the standard being published in 2003.

“The key benefit of the technology we developed was that it was high-speed, could provide an arbitrary architecture, and used a small amount of gates”, says Parkes. “This sort of technology needs to be resilient against radiation. As soon as you put a spacecraft into orbit, it gets impacted by radiation and that can cause a chip to flip from a zero to a one and upset the whole working of a device.

“You combat this by having inherently reliable components and introducing redundancy. We developed the technology over time, and it was a case of understanding what the limitations were with these particular standards and improving on them.”

SpaceWire is now being used on well over £10bn worth of observation spacecraft, telecoms and GPS spacecraft, as well as disaster and environment monitoring satellites. The university’s work on SpaceWire also resulted in a successful spin-out company called STAR-Dundee.

SpaceWire has already been deployed in missions by the world’s leading space programs. It is used in ESA’s GAIA star mapping mission, which is currently in orbit gathering data to form a 3D model of the stars in our galaxy, the Milky Way. It is used onboard Europe’s Sentinel series of Earth Observation spacecraft, providing environmental scientists with crucial information about our planet, and will be part of the GOES-R spacecraft which will replace the USA’s current weather satellites. It is due to be integrated into the ESA’s Euclid space telescope, which will map the distribution of dark matter in the Universe

It is designed into the James Webb Space Telescope, which will be the biggest satellite ever launched other than the International Space Station. SpaceWire will also handle data on-board Europe’s first mission to Mercury when the BepiColombo launches in 2017 on a one-year data-gathering mission.

Parkes says: “We’ve been amazed at how many spacecraft SpaceWire is used on. We’ve already exceeded our ambitions by an enormous amount. We’ve now got 20 years of experience in SpaceWire technology, and we aim to make sure the next generation resolves any issues and makes everything fully backwards-compatible.”

Dundee is now contracted by ESA to explore higher-speed serial-link technology: SpaceFibre. A very early version of SpaceFibre specified by Dundee was used by NASA during testing for the Orion Launch Abort System back in 2009. SpaceFibre operates over fibre optic or electrical cables at data rates of 2.5 Gigabits/s using current radiation-tolerant technology, with high-speeds already planned and multi-laning reaching several tens of Gigabits/s.

“SpaceWire was vey much a partnership between academia, industry and the European Space Agency, with academia doing the research, the ESA providing funding, and industry doing initial chip design and being prepared to use it on crafts. It was everyone working together, and because of that we were able to see something with a real impact.

“We were lucky. It was the right thing at the right sort of time. But there was a lot of hard work too.”

Gaia– Photo credit: European Space Agency

We’re not short on data.
But how can we present it in such a way that we can learn from it? The team from the giCentre at City University London explores new ways of visualising large and complex datasets to help organisations make informed decisions.

19

Former Hewlett-Packard president Carly Fiorina once said that “the goal is to turn data into information, and information into insight”.

Humans have used visual aids such as charts and maps to understand and interpret information for generations. As the amount of data we can collect increases, how do we visualise it in a understandable way that helps us to take action?

Across the world, there is a growing community of researchers looking into how to approach the visualisation of this century’s more complex datasets. Amongst these groups is the giCentre, a team of academics at City University London that has developed software and techniques for interpreting this data in visual forms. The giCentre has spent several years working with organizations including Local Authorities, the Ministry of Defense, the British Library, Transport for London and energy firm E.ON.

“It’s all about shedding light on complexity with simplicity”, says Jason Dykes, professor of visualization at the centre. “We see visualisation as a really interesting solution to a broad range of challenges associated with making sense of large data sets.

Jason Dykes & Jo Wood
Image Credit: Mike Abrahams

An analysis of the workplaces of people using the London Cycle Scheme

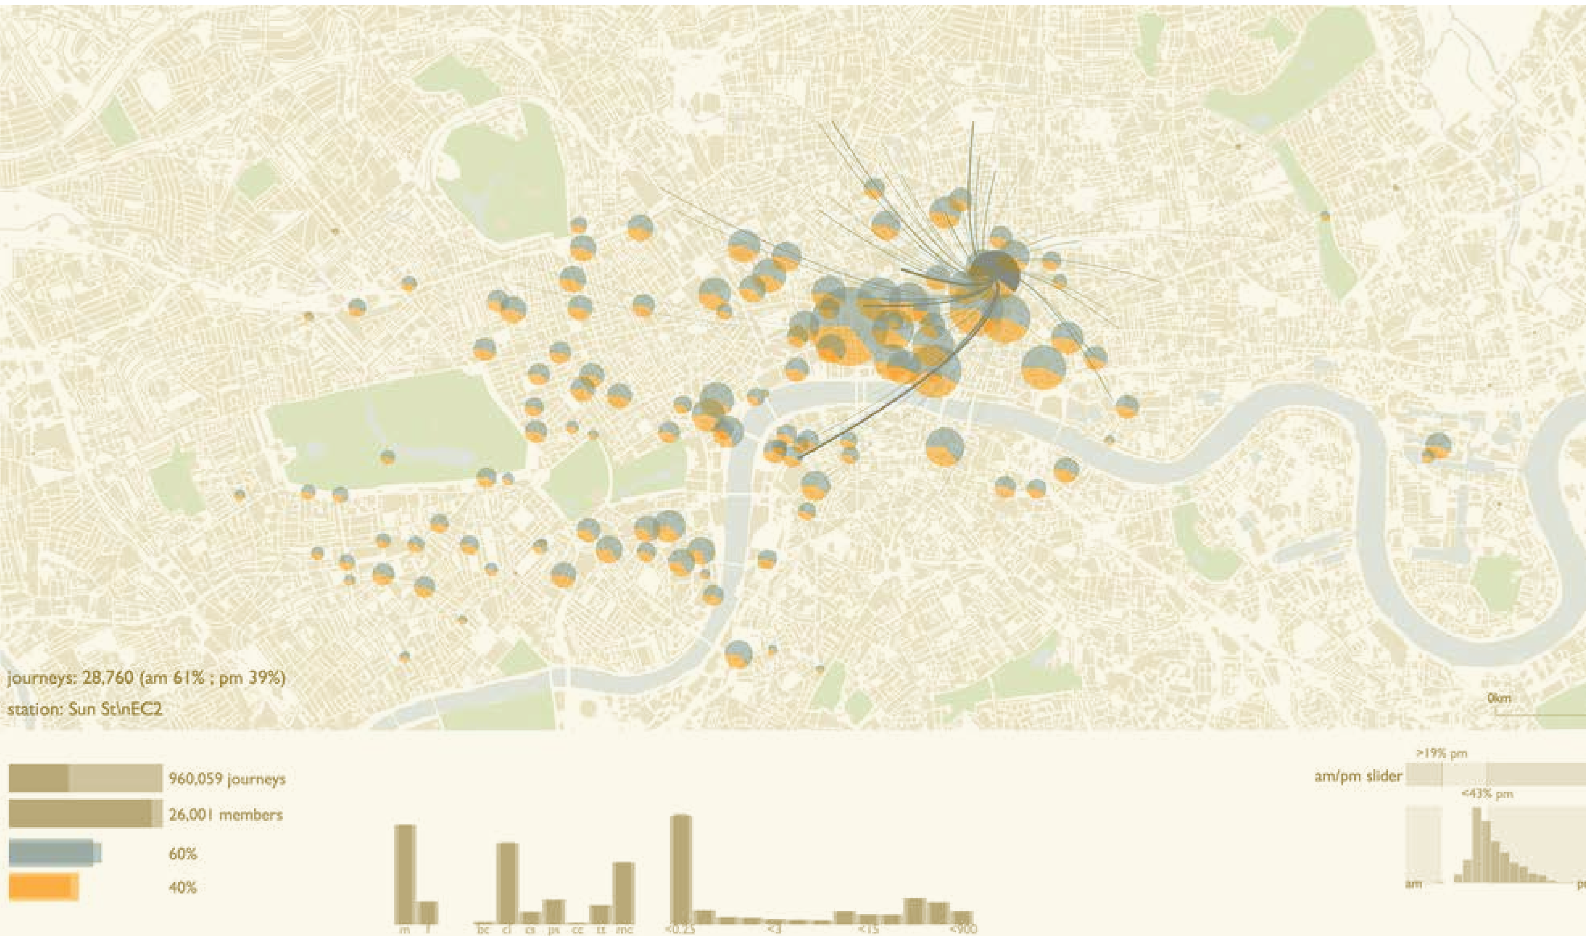

> We're trying to provide richer contextualised views of data, which gives more insight to decision makers.

There is a lot of discussion into how to improve processing and machine learning to meet the challenge of Big Data. Dykes believes that visualisation is often overlooked in that discussion, and that it should have a prominent role to play in the future Big Data research.

"We really believe that Big Data needs the visualisation discipline", he says. "It's important to recognise the value of good visual design as part of Big Data work, and to draw on the knowledge and technology that the visualisation community has developed."

A particularly fruitful area of the giCentre's research has been in using principles from cartography and information visualisation to present non-spatial information in interesting and informative new ways. Both Dykes and Professor of Visual Analytics Jo Wood are experienced in cartography and Geographical Information Science, and the group also features experts in visual analytics, geo-spatial analysis and applied data science.

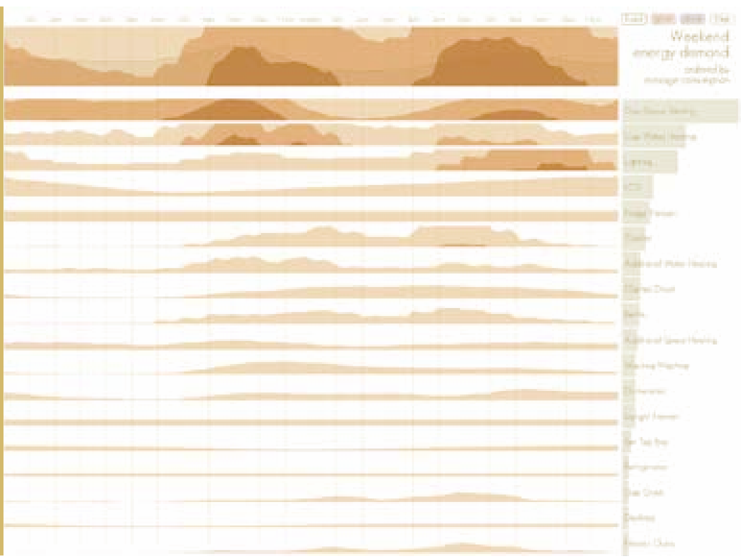

An interactive visualization of daily appliance usage designed as part of a project with E.ON

"We seek out interesting problems to solve", says Dykes. "We try to find combinations of people with needs and data with interesting characteristics, that stretch the current limits of visualisation."

"It's a process of showing people what is achievable, keeping your eyes open to see interesting possibilities, and then engaging with people to learn together so that they can make more of their data and we can expand our knowledge of visualisation and the application domains in which we use it."

In one case, the giCentre worked with Leicestershire Library Service to visualise data from the TALIS database, a detailed dataset containing weekly information on lending and customers at 54 libraries across the county. The process aimed to address questions on how performance varied at each library, where the key customers lived, and whether the library could predict usage based on a customer profiles.

Using data on 435,000 active library users collected over a two-year period, the giCentre structured and filtered the information, removing personal details.

"We get people's data visualised quite quickly using sketches, and then we iterate and flesh out our designs", says Dykes. "It's a way of enabling people to see what they have, and asking them how they want to delve further into it."

In this case, researchers adapted several visualisation methods to incorporate geospatial information. Using recognised techniques such as Spider Plots and Tree Maps, they preserved the geographical context which helped them to reveal subtle patterns in the data. These included the number of lapsed users in a certain area, lower numbers than expected in others, and whether some customers bypassed their own local library to visit another one further away.

The giCentre's research findings - such as the development of Spatial Tree Maps - usually emerge when the team has to take a creative approach to mine an organisation's data for the insight it needs.

How do organisations benefit from visualisations such as these? In one case, leading energy supplier E.ON had collected around 18 million data points from a 100-house Smart Home technology trial in Milton Keynes. Work with the giCentre and City's Centre for Creativity in Professional Practice resulted in four visualisation prototypes being developed. These showed daily peaks and troughs in energy demand, broke down consumption by appliance, pin-pointed trends or unusual instances, and allowed E.ON to gauge the impact of energy-saving technologies.

"You could spend months searching that amount of data for insights", said one E.ON data analyst. "But this just points you straight at it."

Researchers also worked with Transport For London's Cycle Delivery Planning team to raise their awareness of cycling behaviour, particularly where cyclists start and finish their journeys and how they behave. The project has influenced where the organisation places bikes for its cycle hire scheme.

"We had data on over 20 million journeys. If you try and plot that conventionally, you get a hairball", says Dykes. "We made judicious design calls, and highlighted the most important aspects of those journeys.

The study identified previously undiscovered patterns, such as how cyclists respond to station closures, the differences between locals and commuters, and the fact that female cyclists seem less likely to cross the Thames than their male counterparts.

"What TfL became interested in was trying to remove some of the noise in the data sets and look at the interesting structure underneath. We looked at where the workplaces were, as we wanted to remove the commuting patterns and see what happened under that. We ended up being able to see signals that were originally hidden in a mass of data."

The Defence Science and Technology Laboratory also worked with the giCentre to develop approaches for analysing the reports and other intelligence emerging out of areas of conflict. These are unstructured, uncertain and often contradictory, making a rich means of representation such as visualisation particularly appropriate.

Building on what they have already learned, the researchers continue to visualise new data sets in a variety of fields. They are working with police forces in the UK and Belgium to interpret crime reports, and are part of an Arts & Humanities Research Council project to develop software and new research methods to analyse scores, music audio and metadata in large music collections. The Digital Music Lab project also draws on the skills of Queen Mary University of London, University College London and the British Library.

"The process of enquiry does not always have an obvious end point, as the more you find out the more questions you want to ask", says Dykes. "People's expectations change as they interact with data, and some of the discoveries change how they think about what they see.

"We're trying to provide more detailed answers that show exactly where the patterns are and how they differ over time and geography. It gives just that little bit more to decision makers, and often results in organisations using data in new ways"

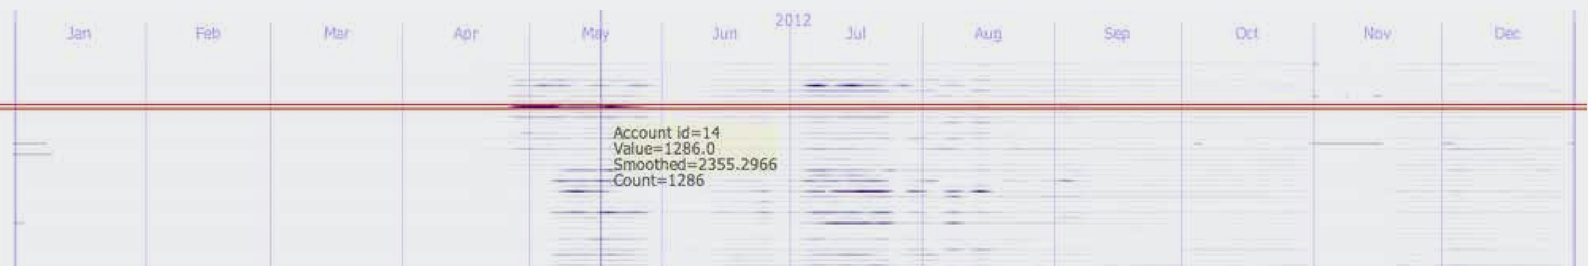

It takes more than just a technical brain to understand the Web. As it weaves ever tighter into our daily lives, the study of what it is and what it can be becomes ever more important.

20

“Why does the Web matter?”, Professor Sir Nigel Shadbolt asked a Royal Society audience in 2013. “Because it’s us.”

It may have been invented and shaped by British scientists, but the Web is no longer just a feat of engineering. It has been described as “the nervous system of the 21st Century” and is a huge part of nearly everything we do, from communicating amongst friends to sharing news and managing our finances.

So when Web pioneers such as Professor Dame Wendy Hall, Professor Sir Nigel Shadbolt and Professor Sir Tim Berners-Lee called for the creation of a new subject christened “Web Science” in 2006, they weren’t just introducing a computer science subject. They were inviting many different disciplines to investigate the many ways in which the Web shapes society, and society shapes the Web.

“We began to realise that you couldn’t just think about the Web from a computer science perspective. It is a socio-technical issue.”, says Dame Wendy. “The essence of the Web is that it evolves and grows because people put content on it. It’s crafted by millions of people.”

While it was CERN scientist Tim Berners-Lee that invented the Web itself in 1989, the University of Southampton has played a large part in of the story of its development. Just before the Web emerged, Wendy Hall’s team created Microcosm, a system that allowed users to browse collections of hypermedia information from text to sound and video. This was one of the first open hypermedia systems in the world.

Artificial intelligence expert Nigel Shadbolt arrived at Southampton in 2000. His research helped to develop the vision of the “Semantic Web”, in which information on the Web is converted into a form that allows machines to interpret and connect it in different ways.

With others, Hall, Shadbolt and Berners-Lee produced a paper published in Science, which called for an interdisciplinary study of the Web, its impact and its development.

“Fundamentally our ambition was to promote the notion that the emergence of this global eco-system was an object of study in its own right”, says Sir Nigel. “There were constituent existing disciplines, but they needed to be brought together to provide a

Professor Dame Wendy Hall

Professor Sir Nigel Shadbolt

systems-level view. We wanted a way of looking at this system that did justice to its technical, social and dynamic aspects".

In 2006, MIT and Southampton combined to form the Web Science Research Initiative, which aimed to probe the technical, social and scientific questions behind the growth of the web. In 2009, WSRI was transformed into the Web Science Trust, a charity that facilitates the development of Web Science research and education globally and coordinates WSTNet, a network of 15 Web Science research labs around the world. The trust also coordinates the Web Observatory, a decentralized repository of data about all aspects of Web and out interaction with it. The Web Observatory is designed to advance our understanding of the Web and how it supports economic and social prosperity.

In 2009, the University of Southampton was awarded £6m in EPSRC funding to create a Doctoral Training Centre for Web Science, partnering with organisations such as the BBC, NHS, IBM and Dow Jones to furnish graduates with the skills to understand and develop the Web. These graduates have gone on to roles private and public sector organisations such as Price Waterhouse Coopers and the UK Civil Service. The Web Science Trust's curriculum development programme has helped to develop the content of Web Science degrees from the UK to South Korea.

"There are many paradoxes to the Internet", Dame Wendy adds. "There's open versus closed, surveillance versus privacy, and freedom of speech versus closed networks".

"All the things that happen in 'real-life' are happening on the Internet at scale in a way that we don't fully understand, and it's happening very fast. Web Science is about looking at what we can learn, what we can predict, and how we can protect without shutting something down.

"We learned how to create a more civilised society over time, and that took centuries to evolve. The Internet is evolving in a very short space of time. So how do we make sure it's good for humanity and enables people to share, disseminate and create new knowledge, while at the same time making it as safe a space as possible?"

The Web and Internet Science group's work at Southampton has also helped to shape World Wide Web Consortium standards in areas such as rich linking and the PROV language relating to provenance. It was instrumental in the organisation of the world's first International Web Science Conference in 2009. Two years later, the Association for Computing Machinery officially recognised Web Science as a research discipline.

Former Google CEO Eric Schmidt has stressed that web science "represents a pretty big next step in the evolution of information". He adds that it is a discipline that is "likely to have a lot of influence on the next generation of researchers, scientists and entrepreneurs". Google's support has been echoed by firms such as BT, IBM, InfoSys, NESTA, the Web Foundation and Switch.

Many public and private organisations have heeded the group's advice on how to present data for an evolving web. Sir Nigel and Sir Tim led the UK effort to publish as much government data as possible as open data which resulted in data.gov.uk and the founding of the Open Data Institute. Open data is an active area of research, development and deployment in Web Science. The principles for public data publication suggested by the University of Southampton have been adopted by the UK Government, which publishes data in such a way that computers can connect data that isn't explicitly linked. Examples include the National Archives' legislation.gov.uk and DCLG's opendatacommunities.org. The London Gazette is also published using linked data. This trend toward open linked data pioneered by the UK is now an international trend. In the United States, the National Climate Assessment report to Congress makes its sources publicly available, referencing its author team and more than 550 technical inputs from papers to datasets and graphs.

Southampton's semantic web research birthed a spin-out called Garlik, which protects consumers from identity theft and fraud. It had over half a million users in December 2011, and has since been acquired by global firm Experian.

"The Web originated from a collegiate community of scientists and engineers", says Sir Nigel. "As it becomes more mainstream, assumptions of trust can't hold. You have fewer assurances that the content you see is what it claims to be. That's going to be even more important as more aspects of life get moved into the digital realm."

Research into Web Science is playing a key role in influencing how we see the Web, and prompting discussion on what we want it to be in future. As well as exploring issues such as copyright protection and net neutrality, the group hopes to develop predictive models that would give organisations a better idea of how best to approach and connect with users online.

"You can't predict that one raindrop will fall in a certain place, but you can predict that a large amount of rain will fall in one area", adds Dame Wendy. "We've become better at predicting the weather over the years, as we have decades worth of data.

"The more we understand about human behaviour in these systems, the better the chance that we can help build the evidence base and robust models to more accurately predict possible futures for the Web. It's helping us look after the health of the digital planet."

The Web Science Centre for Doctoral Training aims to give graduates the skills to understand and develop the web

Appendix

Impact Case Studies submitted to Computer Science and Informatics sub panel 11

	University	Case study title
1	Aberdeen	Data2Text: natural language summaries of complex data sets
2		Public Understanding of Artificial Intelligence
3		Enabling Semantic Reasoning for Linked Data
4	Aberystwyth	Automated Design Analysis and Generation of Diagnostics
5		Planetary Exploration Based Camera Technology for Precision Agriculture Applications
6		The Robot Scientist: automating science
7	Aston	Impact of Machine-Learning based Visual Analytics
8		Improving the Management of Uncertainty on the Web: UncertML
9		The GRiST computer decision support system for assessing and managing risks associated with mental-health problems.
10	Bangor	Enhancing of National Grid Stability via optimisation of Dinorwig Hydro Power Station
11		Minimally invasive procedural training for clinicians using virtual patients
12	Bath	Collaboration research influencing defence and security policy, strategy and stabilisation operations
13		Commercialised advances in computer algebra
14		Virtual warehousing and market intelligence in online book retailing
15	Bedfordshire	Cyberstalking countermeasures adopted by Government bodies
16		Enabling high-quality, low power mobile broadband services
17		Integrated computer hardware and software to support design capability of Energy supply industry and chip manufacturers
18		Powering Electronic Devices from Ambient Radio Signals
19	Birkbeck College	Informatics support for the management and integration of large-scale life sciences data
20		Intelligent Constructionist Tools for Learning and Teaching
21		Participatory Cyber-physical Computing: applications in Monitoring Biodiversity, Healthcare, Urban Development, Transportation and Art
22	Birmingham	Automated detection of faults and undesirable scenarios in very large telecom systems
23		Fighting the Malicious Web
24		Fixing and verifying authorisation protocols of an industry-standard security chip
25		Providing Accessibility to Scientific Documents to Visually Impaired Readers via Mathematical Formula Recognition
26		SIAscopy for rapid non-invasive in-vivo quantification and assessment of skin histology in dermatology and cosmetics
27	Birmingham City	Computational Models to Enhance Design Processes for Engineering Jet Engines at Rolls-Royce
28		Vehicle Crash Test Analysis Systems
29	Brighton	Using Computational Lexicography for Dictionary Production with the Sketch Engine
30		Using Diagrams at Nokia to Protect Privacy

	University	Case study title
31	Bristol	CRYPTO: Elliptic curve research impacts on international standards and products.
32		ProVision: wireless video transmission
33		SLAM: computer-vision-based simultaneous localisation and mapping technology for applications in personal localisation and robotics.
34		VAB: Visual Animal Biometrics to automate animal recognition
35		XMOS: spin-out fabless semiconductor manufacturer
36	Brunel	Algorithms to support faster, cheaper and more accurate diagnosis of specific medical conditions
37		High Performance Simulation techniques to reduce industrial production and logistics costs through better management
38		Improving Data Models in Operational IT Systems
39		Software Engineering Impact through Fault Analyses
40	Cambridge	Vulnerabilities in Electronic Payments
41		Iris-based People Recognition
42		Virtual Network Computing (VNC) remote access technology
43		Security Economics: influence on public policy
44		Ubisense: location sensing
45		Xen virtual machine monitor used by Cloud Services
46	Cardiff	Enabling the Catalogue of Life to index the world's species
47		Realising the potential of 3D scanners through reverse engineering and digital shape reconstruction
48	City University	Design diversity for safety and reliability in software-based systems
49		Making sense of complex data through innovations in visualisation
50		Making the results returned by search engines more relevant
51		Reducing risk in critical computer-based systems by using assurance cases
52	Coventry	Digital Environment Home Energy Management System (DEHEMS)
53		Serious Games Innovation and Business Engagement

	University	Case study title
54	De Montfort	Core Underpinning of Fuzzy Logic
55		Programme transformation
56		Using Research to Change Policy
57	Derby	Knowledge Transfer of Innovative Cloud Computing Technologies
58		Sustainable Cloud Computing
59	Dundee	Augmentative and Alternate Communication (AAC)
60		SpaceWire computer network for use on board spacecraft
61	Durham	From Formal Methods to Software Migration
62		Stereoscopic Imaging
63	East Anglia	Avatars for Visual Communication
64		Improved Insurance Products for the Multinational Insurance Industry
65		Imsense Ltd: The Pursuit of Perfect Photographs
66	East London	Improving mobile service engineering in the Italian Telecommunication Industry
67		Using secure software systems engineering to improve business processes and information systems
68	Edge Hill	Simulation software to improve technology enhanced learning of modern computer architecture.
69		Improving quality assurance of a large software model through relative debugging.
70	Edinburgh	Actual Analytics Ltd: automated processing of video data to reduce the use of laboratory animals in scientific research
71		Automatic detection of defects in multi-threaded enterprise Java codebases
72		Clinical and commercial applications of text-to-speech synthesis technologies
73		Milepost open source GCC compiler and machine learning based compiler technology
74		Shaping the XML technologies used to manage the world's data
75		Speech Graphics Ltd: Audio-driven Animation
76		The EnCore Microprocessor and the ArcSim Simulator
77		The Moses Machine Translation Toolkit
78		The Natural Language Toolkit (NLTK)

	University	Case study title
79		The Tegola Wireless Community Broad-band Project
80	Edinburgh Napier	Enabling greater citizen participation in governance: e-petitioning
81		Standards for Taxonomic Classification of Biodiversity Data
82	Essex	Designing Virtual Worlds for games
83		Optical Switching for High Performance Networks
84		Type-2 Fuzzy Logic: Managing uncer-tainty and imprecision in telecoms and finance
85		UltraSoC: Commercialisation of a novel debug support architecture for mul-ti-processor systems on a chip
86	Exeter	Cost Effective Design of City-Wide Water Distribution Infrastructure
87		Ensuring air traffic safety: optimising short term conflict alert systems
88	Glasgow	Developing algorithms to optimise paired kidney donation in the UK
89		International Standards which have fuelled the rapid global development of telecommunications technology
90		Kelvin Connect – spin-out providing advanced mobile data capture systems for police officers and healthcare professionals
91		Shaping Policy, Legislation and Regu-lation in European Air Traffic Manage-ment
92		Transforming computing science education to confront global industry skills gap
93	Glasgow Cale-donian	Technology to Optimise Movement for Health and Wellbeing
94		User Authentication Methodologies for Secure and Competitive Business
95	Glyndwr	EASYLINE+: Low Cost Advanced White Goods for a Longer Independent Life of Elderly People
96		Software Integration and Visualisation for Complex Electrical Motor Design Programming, Simulation and Model-ling
97	Goldsmiths' College	Multimedia: the impact of Con-tent-based Multi-media algorithms
98		SpendInsight: spend analysis software
99		Winkball: connecting people through video
100	Greenwich	Computational modelling, simulation and visualisation research supporting fire safety engineering

	University	Case study title
101		Computer methods for assessing relia-bility of complex structures
102	Heriot-Watt	Bringing Computer Science, Program-ming and Computational Thinking into the Classroom
103		Enhanced reservoir management in the oil/gas sector via new algorithms for large-scale optimization
104		3D texture capture system enables 'vir-tual' production of IKEA catalogues
105		Pioneering Web Portals for Health Information
106	Hertfordshire	Digital and Accessible Information: Accessibility for All
107		Industrial applications of Automatic Differentiation and advanced methods in compilation technology
108		Robot-assisted Play for Therapy in Children with Autism
109		SBML: the Systems Biology Markup Language
110	Huddersfield	Development of a Next-Generation Stu-dent Response System for Academia and Industry
111		Engineering Knowledge for Autono-mous and Intelligent Systems
112	Hull	HiP-HOPS: A novel method and tool for dependability analysis and optimisation of systems
113		Vertual and VERT (Virtual Environment for Radiotherapy Training)
114	Imperial Col-lege London	Machine Learning for Agrisciences (Syngenta)
115		Reconfigurable Computing for High Performance Applications
116		Applications of Computational Opti-mization under Uncertainty in Decision Support
117		IXICO : Quantitative Image Analysis – Novel Biomarkers for Clinical Trials and Diagnostics
118		Knowledge Management Technology for Pharmaceutical and Healthcare Industries (InforSense)
119		Body Sensor Networks for Healthcare and Sports
120	Keele	Automated object recognition and focussing for Medical Applications
121		Classification within forensic datasets
122	Kent	Communicating Process Architectures: the Future for Systems

	University	Case study title
123		Greenfoot: Transforming the way programming is taught
124		Improved Functional Programming Practice through Refactoring
125		PERMIS – A modular authorisation infrastructure
126	King's College London	Contributions to Industrial Mobile Telecommunications Standards
127		Creation of Wireless Technology Spin-out Companies
128		Data provenance standardisation
129		International software modelling standards
130		Portable Electronic Red Palm Weevil (Rynchophorus Ferrugineus Olivier) Larvae Detector
131		Search Based Software Testing as a Practical Tool
132	Kingston	Economic benefits from sales of people-tracking and crowd-monitoring technology
133		Social and economic benefits from development of sports tracking technology
134	Lancaster	Efficient video for wireless broadband standards
135		Impact of QoS research on the global TETRA radio standard
136		Investigative toolkit and associated techniques to support online child protection based on digital persona analysis
137		Research on natural language processing leading to improved language tests and dictionaries for millions of language learners
138	Leeds	Commercial software for modelling and visualising manufacturing variation in automotive products
139		Human motion analysis research enables global retail analytics industry
140		Scheduling research leads to optimised cost efficient public transport – the Tracsis spin-out
141	Leeds Metropolitan	Green IT: Data Centre Management
142		Improving the online promotion of UK contemporary art on Axisweb.org
143	Leicester	Enhancing the effectiveness of educational games and learning tools
144		Software companies stay innovative and win business in fast-moving and competitive market
145		Underpinning successful UK-based innovation in security alarm systems
146	Lincoln	Giving Medicine a Better Image with Wafer-scale CMOS Imagers
147		User-trainable visual anomaly detection for quality inspection tasks in the food industry
148	Liverpool	Agent-Oriented Software Engineering: The Gaia Methodology
149		In My Shoes: A Software Tool for Professionals assisting Children and Vulnerable Adults
150		National Gas Demand Forecasting
151	Liverpool Hope	MAGIC2VIP: Making Accessible Graphic Information Context to Visually Impaired People
152		Segmentation and Watermarking of Peripheral Blood Smear Images
153	Liverpool John Moores	Securing Networked Systems
154		Fire Prevention and Community Safety
155	London Metropolitan	An Innovative Intelligent Keyboard Design for disabled community.
156		Development of Inclusive Participative Media and games for audiences with special educational needs.
157	Loughborough	Commercialisation of novel knowledge-based computer tools for process plant design check and hazard identification
158		Enhanced products and services through low-cost wireless solutions
159		Image Enhancement Pipelines that Allow Digital Images to be Viewed as Seen by Human Eyes
160	Manchester	Active Shape and Appearance Models for image interpretation, in industrial inspection, medical image analysis, and face tracking/recognition
161		Dynamic binary translation for virtualisation enables applications compiled for one architecture to run on another
162		ICARUS – Interactive Construction of 3D Models from Digital Images
163		Intelligent Pricing Decision Support Systems
164		OWL – an Ontology Language Standard with Sound Logical Underpinning
165	Manchester Metropolitan	Methodology for Evaluating Health Information Technology Systems
166		Synthetic Biology and Citizen Science
167	Middlesex	Algorithms for Bio-imaging

	University	Case study title
168		Applied Ethics – embedding ethics in the design of technology projects
169		Complex User Interfaces: Design of Products and Services
170		ICT Inclusion and Design for All
171		Model Checking Multi-Agent Systems
172		Visual Analytics - Interactive Visual Search and Query Environment
173	Newcastle upon Tyne	Expansion of the middleware software market
174		Improved Processes for the Development of Dependable Systems
175		Novel computational approaches to discover medicines
176		Worldwide adoption of asynchronous circuits and improved business process modelling
177	Northumbria at Newcastle	Cyber Security: Situational Awareness and Infrastructure Protection research changing policy and practice
178		Database migration and data conversion for improved, consistent and integrated address database for the Government of Gibraltar
179	Nottingham	Aptia Solutions - Maximising Use of High Value Materials
180		Broadcasting Thrill for Television, Advertising and Public Engagement
181		Improving Take-Off Time Predication at Heathrow Airport
182		Transforming Theatre, Games and Television
183		Transforming Transport Planning Operations at 3T Logistics
184	Nottingham Trent	Using interaction technologies to help people tackle the effects of stroke and other impairments
185	Nottingham Trent	Wireless and mobile computing for sustainable urban mobility and social inclusion
186	Open University	Empowering people through technologically enhanced senses
187		Enabling exploration of hidden, contextual knowledge within large collections of documents
188		Increasing society's capacity to tackle complex, socio-technical dilemmas
189		Supporting effective live, visual, virtual collaboration
190	Oxford	Automated Software Design and Verification

	University	Case study title
191		Risk Mitigation in the Commercial Design of Microprocessors
192		Securing Data with Database Firewall
193		Semmle: a powerful query language for analysing large data sources
194		Boinc – Open Source tools for Volunteer Computing
195		Validation of Embedded Systems with Bit-Accurate Floating Point
196		Automated Verification and Validation for Defence, Aerospace and Automated Embedded Software
197		HermiT: Reasoning Systems to support ontology systems
198	Oxford Brookes	Reducing fraud in internet registry companies to understand typosquatting and improve abuse detection
199		Sony's Wonderbook: theoretical mathematics contributes to enriching the gaming experience
200	Plymouth	Neural Machine Learning System for MOD Defence Contracts
201		Robotics Applications in Health, Education and Entertainment
202	Portsmouth	Clinical Outcome Modelling Saves Lives
203		Improved Mobility and Quality of Life for Children with Disabilities
204	Queen Mary	Automatic memory safety verification for critical software
205		Causal Bayesian reasoning in critical decision making
206		Cognitive Science Research Group and Chatterbox
207		Semantic Web for Music (Centre for Digital Music)
208	Queen's Belfast	Affordable Confocal Microscopes Enabled by Novel Digital Imaging Techniques
209		Applications of Novel Speech and Audio-Visual Processing Research
210		The commercial impact of scheduling and optimisation on university space planning and utilisation
211	Robert Gordon	AmbieSense: Ambient, context-aware and mobile applications
212		Data-driven Decision Support
213		Optimised Retrieval for Reusing Insurance Underwriting Cases
214	Royal Holloway	Analysis of IT Security Techniques for International Standardisation

	University	Case study title
215		Cryptographic Analysis and Improvement of Transport Layer Security (TLS)
216		VC Learning Theory and Support Vector Machines
217	Salford	Collaborative Visualisation: how people communicate through and around shared simulation
218		The Pattern Recognition and Image Analysis (PRImA) Research Laboratory: Digitising Europe's printed cultural heritage
219	Sheffield	ASR: Commercial, societal and cultural benefits of new advanced Speech Recognition Technology
220		GATE: General Architecture for Text Engineering
221		OAK: Harnessing the power of information for situation awareness and organisational intelligence
222		Shaping international policy and stimulating international public debate on Autonomous Weapons Systems (Ethics)
223	South Wales	Linking Archaeological Data - enabling semantic infrastructure in the digital archaeology domain
224		Mobile applications and technologies making economic impact
225	Southampton	Applications of agent technology
226		Intelligent Energy Management
227		Leading the open data revolution
228		Open Access and Digital Archiving
229		Web Science: Designing a Pro-Human World Wide Web
230	St Andrews	CloudSoft: Distributed Mediation
231		Extending Open Virtual Worlds for Cultural Heritage and Education
232		Migrating IT Infrastructure to the Cloud
233	Staffordshire	Adaptive Video Analytics Software
234		Computer Based Methods for Diagnosing and Predicting River Health
235	Stirling	Models of the control of Koi Herpes Virus
236		Sentic Computing: sentiment analysis in natural language processing
237	Strathclyde	Improved user experience of the European Digital Library through user-centred evaluations
238		Improved video surveillance and customer relations management through efficient data representation.
239		Widening public access to judicial decision making through information systems.
240	Sunderland	Development of Smart Planning Tools for BT and Network Optimisation
241		Sunderland Software City: Developing the Software Sector in the North East
242	Surrey	Adaptive Information Systems
243		Financial Fraud Detection
244		Introducing a secure electronic voting system to the State of Victoria, Australia
245	Sussex	Augmented digital representations of cultural heritage enabling interactive virtual museums
246		Automatic grammatical analysis enabling advanced text processing in commercial applications
247	Swansea	Empowering rural digital communities in the developing world and the UK
248		Improving Processes and Policies in the UK Railway Industry
249		Safer Human-Computer Interaction for Healthcare
250		Document Engineering of Clinical Guidelines
251		Interactive Storytelling Technologies
252	Ulster	A new mobile-based reminding product for connected health
253		A new product for creating annotated data sets within smart environments
254		A new range of outdoor clothing for the active ageing based on wearable technologies
255		Modelling phases of care for acute stroke patients
256		New software products for programming wireless sensor networks
257	University College London	3D body scanning in clothing manufacturing and retail, and healthcare
258		A clinical management service for stroke prevention
259		Camino diffusion MRI toolkit: microstructure imaging and connectivity mapping to avoid cognitive deficits after neurosurgery
260		Enhanced photo and special effects processing for professional and amateur photographers
261		Human-centred security in government and commercial applications

	University	Case study title
262		Improving prostate cancer diagnosis and care using computer simulation and medical image registration
263		SIP/SDP as an enabler of real-time internet communication
264		xlinkit for fast, cheap, reliable banking with automated verification of over-the counter derivatives trading
265	Warwick	High-performance-computing and parallel software development tools
266		Impact of image analysis research on fingerprint biometrics and multimedia forensics
267		Impact of performance modelling research on IT capacity planning and national nuclear security
268	West London	Model-driven software methods that support knowledge-based, process-driven (mobile) service oriented architectures
269		Institutionalising HCI in Asia: focusing on India and China
270	West of England	Facilitating System Evolution during Design and Implementation: CRISTAL
271		More efficient and effective requirements engineering in the aerospace industry using the newly developed ontology-driven methodology, OntoREM
272	West of Scotland	Centre for Enabling Technology Generates Economic Impact on a Worldwide Scale
273		Economic, Societal and Policy Impact of ICTE Research Centre
274	Westminster	Balanced Model Truncation (BMT) and its Applications in DSP System Modelling and Computational Complexity Reduction
275		Component-based Highly Productive Methodology for Software Development in Grid and Cloud Computing
276		Gateway technologies for high-performance computing in business, industry and science
277	York	Guaranteed Performance on Controller Area Network (CAN)
278		Real-Time Operating Systems (RTA-OSEK & RTA-OS)
279		The Goal Structuring Notation
280		RapiTime: Worst-Case Execution Time technology

Report written by John Hill – jpj.hill@gmail.com
Design by Naomi Atkinson – hello@naomiatkinsondesign.com