




Source work

Published on 7 April 2009

By Christine Evans-Pughe



The UK's universities are leading the way in many aspects of ICT R&D. In the first of a three-part series, *E&T* surveys some of the most innovative projects.

Where would you expect to find the world's most influential computer scientists? A likely answer might be the United States of America, but a trip around the UK's university computer departments would change your mind. The labs in the UK are testbeds for some of the most innovative IT technologies that often start as funded initiatives, but in many cases turn into spin-off companies that go on to sell intellectual property to the bigger commercial players.

Take virtualisation – the concept used to turn server farms into flexible computing resources (or 'clouds') that can run multiple operating systems, and simultaneously shift work among servers. It is based largely on an open source standard, XEN, developed by researchers at Cambridge Computer Lab (see 'XGN and the Art of Server Maintenance, p57'). And ARM, the low-power computing engine for mobile phones, was originally the brainchild of Professor Steve Furber, who is now working on some of the most exciting ideas for next-generation computer architectures at the University of Manchester (see BIMPA Mentality, p58)

This series of three features in *E&T*, assesses the gamut of IT-related research now in progress in UK universities. This first piece focuses on a representative projects with the potential to feed through to real-world applications. We follow this in subsequent issues with reports on two contrasting university computing departments: Imperial College, London, renowned for its engineering focus; and Oxford, whose work on modelling and simulating the human heart is paving the way to a more quantitative approach to medicine.

Privacy in distributed systems

Bringing IT into public or private healthcare systems is a complex and expensive endeavour – as anyone involved with the NHS's 'Connecting for Health' programme will tell you. The goal of providing a central electronic record for patients and connecting all GPs with all hospitals and clinics requires the integration of a variety of different, distributed applications and the coordination of widely distributed operations as events occur, such as the referral of a patient from a GP to a clinic. All of this has to be done without compromising patient safety and privacy.

Dr Peter Pietzuch is leading a research project funded by the Engineering and Physical Sciences Research Council (EPSRC) at Imperial College called Smartflow, which aims to solve some of these issues with a new kind of middleware.

"Off-the-shelf middleware such as Java Messaging System (JMS) or IBM's Websphere works well in a regular business context because the middleware assumes a single administrative domain," Pietzuch explains, "but it isn't suited to the healthcare environment because there is a lot of autonomy of different hospitals, healthcare providers, and organisations within the NHS."

In this joint project between the NHS's clinical and biomedical computing unit, Imperial College, and the University of Cambridge, the researchers are building an extendable middleware layer around 'information flow control' – the idea being to ensure it is always possible to track how information flows from application to application and from organisation to organisation.

"We have ways of specifying what we consider are acceptable flows of information, and what is an illegal flow of information – for example, if you have a patient record, you can't release it to the general public without consent from that patient," Pietzuch adds. "If you build an application on top of this kind of middleware, the application will automatically satisfy the confidential integrity... Unlike using specific privacy solutions per application, by incorporating privacy policy in the middleware, it can be applied in a uniform way across different applications."

In this way, legacy applications can work with the new middleware by creating a 'wrapper' around them, says Pietzuch: "As long as you control the flow of data in and out of the legacy application, you can control what it can release to other applications."

Imperial and Cambridge are trying out their ideas with applications used by the Eastern Cancer Registration and Information Centre (ECRIC), an organisation that gathers cancer reports from different hospitals, and creates an archive for statisticians and researchers. The next stage is to deploy a prototype version of the middleware, to use as a case study.

Distributed and decentralised intelligent systems

Dealing with an environmental disaster or a terrorist incident involves rapidly shifting scenarios, where information changes constantly, and is often conflicting, making it hard to get an effective response and ensure the safety of emergency services personnel. The University of Southampton's electronics and computer science department is applying its expertise on autonomous agents to develop software programs that in such situations can interact robustly to maintain data and information systems.

ALADDIN (Autonomous Learning Agents for Decentralised Data and Information Networks) is a multi-million pound, five-year project funded by the aerospace and defence group BAE Systems and EPSRC, with Southampton as the lead partner, and involving Imperial, Oxford and Bristol Universities.

"We've looked at two main scenarios: a city-wide urban response, where you might want to coordinate police, fire, and ambulance to respond to an incident or an evolving series of incidents; and sensors in the Solent, where we're looking at making predictions and extrapolations using data from sensors, which are sensing tide height, wind speed, and weather-related things," says Nick Jennings, professor of computer science at Southampton. "Weather follows a pattern, so what's further down the coast is – by and large – going to move up the coast in a few hours time."

By bringing together many different sources of data, some of which might be better than others, and making informed estimates about missing data, ALADDIN can develop predictive algorithms that tell how, for example, a fire might spread, based on patterns that have been seen before.

Jennings's team is applying a variety of techniques from machine learning, artificial intelligence, and Bayesian reasoning in this project. It is also using 'game theory', a theory of competition beloved of economists that is stated in terms of gains and losses among opposing players. "Game theory doesn't work very well on people, because they don't tend to behave predictably or rationally – but it does work well on software agents," Jennings explains. "We use incentives so that an agent gets rewarded for particular actions, which encourages them to behave in a particular way."

For instance, an incident commander may be tempted to overplay the severity of an incident in order to make sure that he or she receives an adequate number of ambulances or fire engines; but if everyone does this, allocation is not very efficient. The idea, then, is to put mechanisms in place that reward asking for the right amount of resource, and that punish requests for too much. This is tricky – not least because the systems involved are typically open, and anyone can join or add their piece of software or sensor into the network. Just like humans, software agents have to learn who is reliable and trustworthy – and who isn't.

The ALADDIN programme is over half-way through, scheduled to end in October 2010. The core intellectual property developed so far is a series of coordination algorithms, which pull different inputs together to coordinate a response. Already these coordination algorithms, and others for reasoning about uncertainty, are being tried out within BAE business units involved with logistics and coordinating supply chains. In the long term, the aim is that the technologies will benefit all the UK's emergency services.

Next-generation architectures

The desire for ever-improved computer performance has reached the point where multi-core, parallel processors are becoming the norm, but they bring in their wake new sets of hardware and software challenges. Associated closely with the 'many-core' trend are the effects of silicon scaling that mean moving data takes more time, and consumes more power, than number crunching. This, and a growing nervousness about the reliability of nano-scale transistors, is driving research into computer architectures in the UK.

At the University of Cambridge Computer Laboratory, Simon Moore leads the computer architecture group and is working on a concept he calls 'communication-centric design', which shifts the emphasis of computer system design away from computation, and towards communication between processors. In an EPSRC-funded project called C3D, working with colleagues Professor Alan Mycroft, Dr David Greaves and Dr Robert Mullins, Moore is looking to explore a number of questions. "For instance, when you have a large number of processors, to what extent can the processors do the scheduling or event handling in hardware? Also how, from an architectural point of view, can you design multiprocessor systems so that they help optimise the behaviour of 'parallel skeletons' – the higher-level building blocks used to design parallel applications?" Moore says. "Some of this applications analysis is used in traditional parallel processing – but we're trying to figure out how it all works in conjunction with chip multiprocessors."

Moore's Cambridge group is also looking at million-core power efficient processors within Steve Furber's latest EPSRC-funded project, BIMPA (Biologically Inspired Massively Parallel Architectures): Computing Beyond a Million Processors, which aims to deliver massively parallel machines with a million ARM processors, capable of modelling a billion spiking neurons – around 1 per cent of the human brain – in real time (see BIMPA Mentality, left).

Search engines

A thriving theme in IT is data sharing between organisations and data mining to find patterns in large datasets. This is an area in which the University of York's computer science department has been innovative, with the development of a 'search engine for signals', which it has been applying to pattern-matching across datasets terabytes in size using a GRID computer resource shared with the Universities of Leeds and Sheffield.

York's approach was developed largely within an EPSRC-funded project called DAME (distributed aircraft maintenance environment) between Rolls-Royce, DS&S, and the universities of Sheffield, Leeds, Oxford, and York, in which the partners showed how the GRID and Web services could ease the design and development of systems for diagnosis and maintenance applications, which combine geographically distributed resources and data.

"In an aircraft engine, a signal event you're interested in might be a bird strike or just a squeaky bearing," explains Professor Jim Austin, whose team developed the pattern-matching tool called SDE (Signal Data Explorer). "You may not have seen this particular noise before, so you can search a database to find a similar event and see how someone else has fixed it before." The DAME project began by collecting and analysing vibration data from aircraft engines for diagnostics and prognostics and was then broadened to include simulation and modelling of engine systems, and after-market support.

Signal Data Explorer is being used in two other projects. One is EPSRC-funded Grid-computing project CARMEN (code analysis, repository, and modelling for e-neuroscience), involving 11 UK universities, with the objective to create a virtual laboratory in which data on neuronal activity can be shared, stored, manipulated, and modelled. The other is Freeflow, a multi-partner venture jointly funded by EPSRC, Department of Transport, and the Technology Strategy Board to develop intelligent transport systems including

signal-matching to optimise traffic flow.

“If you can take distributed data from the traffic, from inductive sensors in the road for example,” Austin says, “you can recognise patterns within those signals, and then use that information to advise various traffic controls to ensure, say, that the buses run on time.”

The partners are now developing a demonstrator to manage bus timings on the Hull road in York. Next stop is Hyde Park Corner in London. SDE is being commercialised by Cybula Ltd, Austin’s spin-out company, where it is looking at applications with Rolls-Royce and also in areas such as oil, gas and rail industries.

Computing for the future of the planet

IT has a central role in society, whether it is helping to run global enterprises efficiently or enabling social networking. While computerisation has been gathering pace, it is only comparatively recently that the energy and resources it consumes have been rigorously scrutinised.

The University of Cambridge Computer Laboratory is looking at these, and related issues, within its ‘Computing for the Future of the Planet’ framework. As Dr Andrew Rice – who is developing this project with Professor Andy Hopper – points out, research in this field could make huge contributions to society.

For instance, while there is growing awareness about power consumed by data centres, we tend to overlook the energy used to manufacture the microchips running the show, which is an order of magnitude more (by weight) than most other manufactured goods. Conversely, computing can have a positive environmental impact, such as using sensor data to optimise transport systems or improving modelling techniques to make better predictions about climate change, or by providing digital alternatives to physical activities. Take iTunes, for example, Rice says: “We don’t download music because of carbon trading incentives, but because it’s more convenient than buying CDs. If we could arrange other digital alternatives that have a lower environmental impact, then we’re onto a winner.”

Hopper and Rice manage a panoply of projects, ranging from trying to minimise the power consumed by spinning disks in data centre servers, to how computing might exist more sympathetically in the environment. For example, can we modulate power consumption so it works with the ups and downs of the UK’s National Grid electricity supply network? “So, when everyone puts the kettle on at the end of a TV show, all data centre computers power-save for two minutes, and reduce the variability of the supply, like a virtual battery,” says Rice. “If you can reduce power consumption on demand, it might also exist better with renewable energy. Maybe we can build our datacentres next to wind farms or solar power generation schemes and move our compute jobs around the world, chasing energy where it’s available, and using it where it’s spare.”

Rice refers here to container-isation technologies that can be used to move computer tasks when they are still running.

Ideas being explored are long-term explains Rice, looking perhaps 10 or 20 years ahead, and thinking about benefits to the world outside of computing.

Further information:

www.doc.ic.ac.uk/~peter/smartflow.html

www.ecs.soton.ac.uk

www.cl.cam.ac.uk

www.cs.york.ac.uk

www.cs.manchester.ac.uk

Comments

[All comments](#)

You need to be registered with the IET to leave a comment. Please log in or register as a new user.