


# InspiredResearch

RESEARCH NEWS FROM THE OXFORD UNIVERSITY COMPUTING LABORATORY



## SINGLE-CHIP CLOUD COMPUTING: ADVENTURES WITH INTEL



### ROBOT RESCUE

OUCL team do battle at  
Robot World Cup - p12



### BADGER QUEST

Using sensor networks to  
track badgers in the wild - p3



### HEALTHIER HEARTS

Predicting the cardiac side  
effects of new medicines - p8



UNIVERSITY OF  
**OXFORD**

## Welcome

In*Spired* Research is the twice-yearly newsletter published by the Oxford University Computing Laboratory (OUCL). If you would like to learn more about anything in these pages please get in touch:  
Oxford University Computing Laboratory  
Wolfson Building  
Parks Road  
Oxford, OX1 3QD

editorial@comlab.ox.ac.uk

## Editorial Board

Bob Coecke, Ciara Dangerfield, Alastair Donaldson, Max Goulding Thorsten Hauler, Shoshannah Holdom, Muzammil Hussain, Sharon Lloyd, John Lugton, Thomas Lukasiewicz, Suzanna Marsh, James Osborne, Bogdan Panait, Stephen Pulman, Max Schaefer, Aadya Shukla.

## Contributors

Kelly Burrowes, Bernardo Cuenca Grau, Pamela Farries, Katherine Fletcher, Vojtech Forejt, Ian Horrocks, Aleks Kissinger, Markus Krötzsch, Marta Kwiatkowska, Alex Merry, Gary Mirams, Bill Roscoe, Julie Sheppard, Giorgos Stoilos, Thomas Wahl, Elizabeth Walsh, Helen Wilcox.

Where images within this publication have been used under a Creative Commons licence, credit is given alongside the photo.  
Printed on recycled paper.



## Start-up GroupSpaces.com secures \$1.3 million funding

**Oxford computer science graduate Andy Young, and David Langer (mathematics) have secured a \$1.3 million investment round for their venture GroupSpaces.com, a site for simple online management and administration of real-world clubs and groups.**

GroupSpaces arose from frustration at the multitude of different websites which clubs and societies at Oxford University were using to organise themselves online. The site has been adopted by most large clubs across the UK's top universities and is currently used to manage over 500,000 group memberships. To date, it has attracted over 100,000 beta testers in the UK student market and 70 blue-chip advertisers.

GroupSpaces aims to meet all the needs of groups such as sports clubs, hobby groups, associations and student societies. This integrated offering makes administration of

groups more effective, improves member communications and saves group organisers and officers time compared to more traditional methods.

GroupSpaces differentiates itself by only requiring group leaders to sign up, making it fast to adopt and get started for secretaries, treasurers and other officers of existing real-world groups. Group members can be imported from Excel and then receive emails, register for events and pay their membership dues, all without having to create a user account and password. GroupSpaces is free for groups with less than 250 members and for student groups of any size. Above 250 members, pricing starts at £5/month.

In April, the company was selected for WebMission as one of 20 promising UK web startups. Andy and David have also recently appeared in a Spectator article, tipping them to be two rising stars of the business world.

## Computer science careers fair announced

**The first dedicated 'Careers in Computing' event will take place on Thursday, 10 February 2011, giving companies offering graduate employment opportunities the chance to talk directly with Oxford students.**

The 'Careers in Computing' event is a joint initiative between the Careers Service and OUCL. It aims to connect employers to Oxford students looking for a career that involves the use of Computer Science.

The event will be open to both undergraduate and postgraduate students (taught and research) from all disciplines, and will give students the opportunity to discuss graduate opportunities with both large blue chip companies, and smaller local organisations. It will also give students the opportunity to gain a better understanding of what companies have to offer, and in turn what companies are looking for in their graduate intake.

It is anticipated that there will be a broad range of organisations from those dealing in computer systems and hardware, to software houses, consultancies, research institutes and professional services organisations from a variety of sectors.



The event will take the form of a networking fair which will run from 3pm–6.00pm in the Computing Laboratory, Parks Road, Oxford. There will also be a selection of talks covering topics such as internships, work experience, graduate traineeships, continuing professional development and membership of industry bodies.

All stand places for this year have been taken, but please contact Helen Hill of the University Careers Service via [fairs@careers.ox.ac.uk](mailto:fairs@careers.ox.ac.uk) to receive information about future fairs when available. To become a member of the OUCL Industry Relations Network, allowing your organisation to post vacancy and internship information to Oxford Computer Science students, you can sign up at [www.comlab.ox.ac.uk/industry/](http://www.comlab.ox.ac.uk/industry/)

# Boosting companies' R&D capacity with the Industrial CASE Programme

**The Industrial CASE programme is an excellent opportunity for companies who wish to kick-start or boost their research and development capability by collaborating with a university partner, such as Oxford.**

This programme enables a company to employ a highly qualified postgraduate trainee, and benefit from the university's research expertise and facilities. Each Industrial CASE studentship award made sees a post-graduate student carrying out a PhD – or a DPhil as they are known at Oxford – working on an industry research project for 3.5 years.

Oxford University's Computer Science department has collaborated with a number of industrial partners, including QinetiQ, IBM, Microsoft and Alcatel with several students funded in this way. Projects must be in the engineering or physical sciences (including computer

science) and are jointly supervised by the academic and industrial partners. The university and company together recruit a suitable, eligible candidate.

The company is required to provide additional funding to the student project, a minimum of a third of the EPSRC funding of £66,838 over 3.5 years (2010 figures) but the company can contribute more if they wish. From this the student receives an enhanced stipend payment.

Companies should also provide a placement at their premises (minimum of 3 months over the 3.5 years of an award), plus funding to cover the costs of travelling to and from the company and any accommodation or subsistence costs for the student while on placement.

Companies wanting to get involved need to define a student project and

contact the university department with a view to arranging collaboration with a particular member of academic staff who shares the company's research interests.

Studentships are allocated to large companies based on their participation in EPSRC-funded research. EPSRC appointed agents to help involve small- and medium-sized enterprises (SMEs) in the scheme: suitable SMEs receive a quota of Industrial CASE awards and the agents can assist these companies in making contact with appropriate university departments. Companies may also apply directly to EPSRC through an annual pool competition. Further information is available at [www.epsrc.ac.uk/funding/students/coll/icase](http://www.epsrc.ac.uk/funding/students/coll/icase)

Companies interested in working with OUCL via the Industrial CASE programme please contact Julie Sheppard at [sheppard@comlab.ox.ac.uk](mailto:sheppard@comlab.ox.ac.uk)

## WildSensing: tracking wild badgers

**WildSensing is a project to monitor wild badgers in their natural habitat in Wytham Woods, a Site of Special Scientific Interest in Oxfordshire.**

Current approaches using VHF beacons are highly labour intensive and involve night-time tracking. GPS tracking collars typically perform poorly under heavy tree cover. The WildSensing approach was to leverage the power of sensor network technology to provide continuous and long-term monitoring. WildSensing is a cross-disciplinary project undertaken by Niki Trigoni's Sensor Network Group, part of OUCL, along with the Cambridge Computer Laboratory and Oxford's Wildlife Conservation Research Unit. Over 40 animals were equipped with tracking collars which periodically

emit radio beacons. In order to rapidly deploy the devices, existing active radio-frequency identification (RFID) devices from WaveTrend were used. These transmissions were detected by a network of 30 sensor nodes placed at strategic locations, such as at the badgers' setts. In a one-year period, over 26 million transmissions were detected and logged.

The deployment experience has led to: the development of algorithms to deliver data using a hierarchical delay-tolerant approach; surprising results about the effect of rainfall on link behaviour; and hardware and firmware advances. A biologically inspired method of evolving code for distributed networks using an abstraction of gene regulation was also devised. The initial data has revealed some interesting badger behaviour which has previously been speculated upon, but never before been observed directly, such as dispersal of animals from one sett to another.

Another avenue of research that has been recently explored is localising badgers when they are underground. Radio signals are unable to penetrate soil, and thus conventional tracking modalities cannot be used. Instead, a novel method using low frequency magnetic fields was developed. As a

by-product of localising animals over time, the tunnel structure itself can be revealed. This new research area is likely to lead to some interesting discoveries about badger social structure underground, something which has been impossible to do until now.

Two papers have been accepted for the highly-regarded ACM Conference on Embedded Networked Sensor Systems (SenSys 2010). This is particularly impressive considering that only 25 papers in total were accepted from over 145 submissions.

## 2011 Industry Day announced

A date for your diary: the 2011 OUCL Industry Day will be held on Wednesday 1st June 2011, in Oxford.

This annual event is designed to give industry participants, through a series of talks and demonstrations, the opportunity to explore the exciting and groundbreaking computer science research carried out at Oxford, and find out about opportunities to work with the university. Further details, including how to book, will appear here in the coming months: [www.comlab.ox.ac.uk/industry](http://www.comlab.ox.ac.uk/industry)



Photo by: Andrea Abbott



## OUCL authors launch new books

**'The highest goal of programming language design [is] to enable good ideas to be elegantly expressed'. These are the words of Sir Charles Antony Richard**

**Hoare, or Tony to his colleagues and friends, a man who has inspired generations of researchers in computer science.**

Director of OUCL Bill Roscoe has co-edited a new book *Reflections on the Work of C.A.R. Hoare*. It explores Tony Hoare's ground-breaking contribution to the evolution of programming. The British computer scientist is perhaps best known for the development (at age 26) of Quicksort, one of the world's most widely used algorithms.

Hoare's other achievements include Hoare Logic for verifying program correctness; Communicating Sequential Processes (CSP) for describing the interactions of concurrent processes and the inspiration for the Occam programming language; and Unifying Theories which provide a common framework for the semantics of programming languages.

Tony Hoare's interest in computing was awakened when he studied philosophy at Oxford. He returned in 1977 to lead the Programming Research Group and later the whole Computing Laboratory before his retirement in 1999.

*Reflections on the Work of C.A.R. Hoare* presents a comprehensive review of Tony's achievements, with original contributions by more than 30 international leaders in the field. It contains papers from a network

of collaborators and disciples on contemporary topics to which Hoare has provided insight and direction. All book royalties are going to the Hoare Fund for the benefit of the laboratory.

***Reflections on the Work of C.A.R. Hoare*. Edited by Cliff B Jones, AW Roscoe & Kenneth R Wood. ISBN: 978-1848829114**

### New Functional Algorithm Design Book

**Richard Bird's new book *Pearls of Functional Algorithm Design* was recently published.** This unique collection of 30 programming problems draws from various sources including games and puzzles, sorting, and problems in data compression. **ISBN: 978-0521513388**

## Oxford champions International Semantic Web Conference

**The Computing Laboratory is centrally involved in the organisation of this year's International Semantic Web Conference (ISWC), to be held at the Shanghai International Convention Centre, Shanghai, China in November. Professor Ian Horrocks is conference chair, and Dr Birte Glimm is proceedings chair.**

ISWC is the major international venue where the latest research results and technical innovations on all aspects of the Semantic Web are presented. ISWC brings together researchers, practitioners and users from universities, research labs and industry, and from a wide range

Photo by: Nicolas Rollier



of research areas, including artificial intelligence, databases, social networks, distributed computing, web engineering, information systems, natural language processing, soft computing, and human-computer interaction.

As well as a research track, the

conference also has a semantic-web-in-use track, targeted at deployed applications, and an industry track, which provides the business community with an opportunity to present innovative products that utilise or enrich the semantic web.

The ISWC programme also includes invited talks given by leading figures from both the academic and business world. Amongst this year's are Austin Haugen of Facebook, who will talk about Facebook's new Open Graph protocol, and Evan Sandhaus of the New York Times, who will explore the future of RDF and RDFa in the news industry, and who also promises to announce an exciting new milestone in The New York Times Company's linked data efforts. There will also be a talk from Jay Katzen, managing director of Elsevier's academic and government products for the science and technology division, in which he will announce a new Elsevier platform for the development of applications that leverage scientific literature to improve their search and discovery process. Further information: <http://iswc2010.semanticweb.org>

## Newton Fellowship extends researcher's Oxford stay

**A prestigious Newton Fellowship was recently awarded to Vojtech Forejt to undertake research at the Oxford University Computing Laboratory.**

This Fellowship is a joint initiative of the British Academy, The Royal Academy of Engineering, and the Royal Society. Its aim is to select the best postdoctoral researchers from outside the UK and offer them support for two years to pursue research at UK institutions, and afterwards provide additional 10-year support for maintaining connection with the UK. The success rate of the applications is about 7 per cent. Vojtech received his PhD from Masaryk University, Brno, Czech Republic, in 2009. He specialises in the theory of formal verification for probabilistic systems. Here at OUCL, he will be a member of the Quantitative Analysis and Verification group, where he will, along with other group members, work mainly on development and implementation of algorithms for real-time systems. Vojtech has been a visitor in the Laboratory since January 2010. The Newton Fellowship will allow him to extend his stay until the end of 2012.



# BCS recognises Oxford students

**Computer science graduate Anne-Marie Imafidon has reached the final of the British Computer Society (BCS) and *Computing UK*'s annual IT Industry Awards.**

Anne-Marie submitted a project she'd been working on as part of the requirements engineering element of her maths and computer science degree at Oxford, in order to reach the final ten in the Business Analyst of the Year category. The nominated project saw Anne-Marie tackling real-world computing problems for opticians Dollond & Aitchison.

Anne-Marie is 'very happy' to have been nominated, and commented 'I received some of the highest marks from my degree for this project, but I never thought it would get this level of recognition. It was a big surprise!' Anne-Marie added: 'It's really exciting and has fuelled my love for computer science. It just shows anyone can achieve anything if they give it a try.' The UK IT Industry Awards are a benchmark for excellence throughout the computer industry. The awards focus on the contribution of individuals, projects, organisations and technologies that have excelled in the use, development and deployment of IT in the past 12 months.

The Judging Day took place in October at the Madejski Stadium, Reading. The winners will be announced at a prizegiving ceremony, to be held in London in November.

Robin Booth, Publishing Director of *Computing* said: 'I would like to add my congratulations to the finalists. The UK IT Industry Awards are rigorously

judged and to be announced as a finalist is a great achievement.'

Anne-Marie who now works for Deutsche Bank in London has worked for a number of blue chip companies ranging from Hewlett-Packard to Lehman Brothers, and Dollond & Aitchison (Boots). She volunteers her weekends and holidays working with the Excellence in Education programme which supports inner city children in education and was recently named one of the top 100 Future Leaders in the UK.

## Computer scientist 'BCS IT Student of the Year Award' runner-up

**Oxford computer science student Thomas Gibson-Robinson was one of three people nominated for the Best Information Technology Student at the Science, Engineering & Technology Awards in 2010.**

The Science, Engineering & Technology Student of the Year Awards are well established as one of Britain's most important awards for science and engineering undergraduates. The awards are organised by the World Leadership Forum, with the award for IT student of the Year sponsored and judged by the BCS. Thomas, a recent graduate of the MA in Computer Science programme at Oxford (St Catherine's College) was one of two runners-up in the competition for his work on the automatic simulation of process algebras.

Peter Calvert, a student at Cambridge University, was revealed as the winner for his work on parallelisation of Java for graphics processors.

## Award for web ontology language research

**The Best Paper Award at the 12th European Conference on Logics in Artificial Intelligence (JELIA) in Helsinki was awarded to Markus Krötzsch, research assistant in the OUCL Information Systems Group.**

His work 'Efficient Inferencing for OWL EL' is concerned with the recently standardised W3C Web Ontology Language (OWL 2) that is used for modelling knowledge in science and industry. Prominent application fields of OWL 2 include the life sciences and the semantic web, and handling the huge volumes of information of such domains has become a major research challenge.

The OWL 2 standard thus provides three 'profiles' – combinations of language features that are easier to implement efficiently. OWL EL is such a profile, but no complete algorithm for efficiently processing it had been published yet. The paper closes this gap by proposing new algorithms for OWL EL, and by showing how certain features of OWL EL affect the efficiency of these algorithms. This provides a foundation for improved implementations that achieve higher performance.

The biennial JELIA conference is one of the main conference series in the field, and the Best Paper Award includes an invitation to present at the International Joint Conferences on Artificial Intelligence (IJCAI) in Barcelona next year.


## Hillston gives Strachey Lecture

**This term, OUCL welcomed Jane Hillston (Professor of Quantitative Modelling and EPSRC Advanced Research Fellow in the School of Informatics at the University of Edinburgh) who was invited to give the Strachey Lecture.**

Her talk was entitled 'Making stochastic process algebras count: modelling collective dynamics.' The Strachey Lectures are a series of talks named after Christopher Strachey (1916–1975), the first Professor of Computation at the University of Oxford. He was the first leader of the Programming Research Group founded in 1965, which became part of Oxford University Computing

Laboratory. The first lecture of the series was given in 1995 by Dana Scott who had collaborated with him. The lectures are given once in each term by speakers who are highly regarded in their field of computing science – for example Don Knuth and Sir Tony Hoare.

The lectures attract distinguished speakers, and we welcome people from outside of the department who may also wish to attend. Tea and cakes are served beforehand. The events are free, and attendance is on a first-come-first-served basis. Forthcoming lectures will be publicised at: <http://web.comlab.ox.ac.uk/seminars/strachey>



## Collaborating with Intel: programming tools for the Single-Chip Cloud Computer

**Alastair Donaldson, an EPSRC Postdoctoral Research Fellow at OUCL, has been awarded a research grant with Intel, allowing him access to their novel Single-Chip Cloud Computer (SCC) technology.**

The project, in collaboration with Codeplay Software and Imperial College London, will involve designing programming tools for analysing and optimising SCC software.

### **The rise of multicore processors**

Complex computational problems, such as real-time medical imaging, financial analysis, particle simulation and high-definition video, demand faster computers. Until early last decade, the speed of processors increased

dramatically every year. This was due to frequency scaling, facilitated by Moore's law, which states that the number of transistors that can be inexpensively placed on an integrated circuit will double approximately every two years.

For many years, shrinking transistors allowed the clock speed of microprocessors to be increased, meaning that users requiring improved performance for an application could simply wait for the next generation of computer systems to be released.

Although transistor densities have continued to double every 18-24 months, further increases in processor frequency have been found to lead to prohibitively high levels of power consumption.

Instead of increasing the frequency of individual processor cores, manufacturers such as Intel and AMD have opted to deliver performance by using the additional transistors afforded by Moore's law to design processors consisting of multiple cores.

***“For application developers and users, this means that the “free lunch” is over: taking advantage of the performance offered by multicore processors requires serious programmer effort.*”**

The two main challenges are associated with multicore software development are programmability, and scalability. Designing an application to take advantage of multiple cores

often involves significant algorithm redesign. Correctly implementing parallel applications is also extremely challenging: the pitfalls of parallel programming are infamous. Given an application in parallel form, achieving high performance on a large number of processor cores can be difficult, due to the overhead of communication, reversal of sequential optimisations that are no longer applicable when parallelism is exploited, and limited memory bandwidth. On a standard multicore system, limited bandwidth means that many applications do not scale well beyond four cores, and very few applications scale well beyond eight cores.

## The Single-Chip Cloud Computer

Intel's Single-Chip Cloud Computer (SCC) is an experimental research microprocessor consisting of 48 processor cores. This is the largest number of Intel processor cores ever integrated on to a single chip. The aim of SCC is to provide a platform for research into methods for reducing the *scalability* problems outlined above.

The 48 SCC cores are divided into 24 tiles, with two cores per tile. The tiles are connected via an on-chip network, and cores communicate via message passing, for which hardware support is provided. The chip also provides facilities for software to dynamically manage power consumption by controlling both voltage and frequency.

**“ The name ‘Single-Chip Cloud Computer’ comes from the fact that the architecture resembles a cloud of computers, in miniature. ”**

SCC potentially promises high scalability by abandoning hardware cache coherence, seriously reducing the problem of a shared memory bottleneck.

## Collaboration with Oxford

Alastair Donaldson has been awarded a research grant from Intel, providing him with access to the SCC architecture. The project, which commenced in August, is in collaboration with Codeplay Software, an Edinburgh-based company specialising in compilers and development tools for multicore architectures, and the Software Performance Optimisation group at Imperial College London. The project's focus is on reducing the problems of

programmability, discussed above, in the context of the Single-Chip Cloud Computer.

SCC can be programmed at a low level, using *bareMetalC*. Essentially, this is ANSI C packaged with an application programming interface (API) exposing features of the SCC architecture. Inter-core communication is via RCCE, a message-passing API somewhat similar to message passing interface (MPI).

The low-level nature of this programming model opens a wide scope for programmer errors, leading to subtle concurrency-related bugs that are difficult to understand and fix. While detailed control over the low-level features of SCC will be essential for system-level programmers, such detail comes at the expense of portability. Regular application developers prefer to write portable programs at a higher level of abstraction, and would like to take advantage of the compute power of SCC without writing low-level message passing code to handle data-movement. The research project proposes to tackle these issues with two complementary tools - one to validate explicit low-level (*bareMetalC* and RCCE) code automatically, and the other to generate such low-level code from a higher-level programming language based on C++.

## Techniques Formal analysis

Formal analysis techniques will be designed to automatically check correct usage of the RCCE message passing API by *bareMetalC* programs. The techniques will be based on symbolic model checking and abstract interpretation, and will detect bugs in SCC programs caused, for example, by races between messages, erroneously orchestrated messages that corrupt local memory, non-adherence to alignment requirements, and deadlock due to misuse of inter-core synchronisation primitives.

These techniques will be aimed at low-level programmers who require detailed control over the SCC hardware; the tool will boost the productivity of such programmers, aiding them in writing robust applications. This research will build on SCRATCH, a race analysis tool for the Cell Broadband Engine architecture, designed by Alastair Donaldson in collaboration with OUCL's Daniel Kroening and Philipp Rümmer.

## A high-level programming model

Developers wishing to write portable

code for a variety of architectures, *including* SCC, require higher-level abstractions than those provided by *bareMetalC* and RCCE. *Offload C++*, from Codeplay Software, is an automatic framework for offloading parts of C++ applications to run on multicore processors equipped with scratch-pad memories. *Offload C++* solves many of the code-duplication and data-movement problems associated with programming multicore architectures with multiple memory spaces. *Offload* has been successfully applied in the context of the Cell Broadband Engine architecture.

Codeplay will contribute to the project by developing a prototype version of *Offload C++*, dubbed *Offload-SCC*, targeting the SCC architecture. This will facilitate high-level, portable programming of SCC. In collaboration with Paul Kelly's group at Imperial College, *Offload-SCC* will be extended with data-movement optimisations geared towards maximising performance on SCC.

*Offload-SCC*, combined with data-movement optimisations, will contribute to the goal of performance portability for SCC, which has clear advantages for industrial practitioners. The success of *Offload-SCC* will be measured with respect to applications from high-performance computing, and computer games. The aim is to accelerate these applications on the SCC architecture without sacrificing portability. In addition, the formal analysis tools developed during the project will be used to analyse the *Offload-SCC* implementation.

## Industry website launched

Want to:

- Post job adverts for our students to see?
- Find out about new ways of working with us?
- Learn more about Computer Science at Oxford?

Sign your organisation up for free at OUCL's new industry website

[www.comlab.ox.ac.uk/industry](http://www.comlab.ox.ac.uk/industry)

## Improving the prediction of cardiac side effects

**More than a third of promising pharmaceutical compounds are abandoned during development due to the perceived risk of unwanted cardiac side effects. Gary Mirams from the University of Oxford's Cardiac Electrophysiology Group explains:**

One of the most dangerous arrhythmias, Torsades-de-Pointes (TdP) is linked to inhibition of the hERG potassium channel, which makes an important contribution to the electrical activity of the heart that coordinates its contraction. The earliest stage of screening for TdP safety currently consists of measuring the hERG channel block (rejecting compounds that bind and block this ion channel with high affinity). However, there are many drugs which are strong hERG inhibitors and are safe, and others which can cause TdP but are not strong hERG inhibitors.


Researchers from the Computational Biology Group, along with colleagues from the Department of Physiology, and a selection of external partners are working on the preDiCT project with GlaxoSmithKline (GSK) to improve the prediction of TdP risk in early

stage compound development, using mathematical models to integrate experimental data on the activity of various commonly measured ion channels. The researchers are using the experimental measurements provided by GSK on hERG, sodium and L-type calcium channel blocks, and comparing this with the clinical risk of TdP. Of the 31 drugs studied with known TdP risk, they have found that risk classification as 'safe' or 'dangerous' according to hERG inhibition alone misclassifies nine drugs, while the new methods correctly classified eight of these nine.

The suggested calculations can all be performed on a desktop computer in a small amount of time, and may improve the reliability of drug safety assessments beyond those based solely on a drug's hERG block effect.

The preDiCT team are in the early stages of a study of abandoned pharmaceutical compounds, where they will demonstrate what could be achieved given the information that the pharmaceutical companies had before deciding to stop development of the compound. They hope to predict what happened in the safety testing process, and see how early it would have been possible to predict that the compound should have been dropped.

preDiCT is one of several projects funded by the European Commission, as part of their support for the development of a virtual physiological human, coordinated by the Virtual Physiological Human Network of Excellence.



## EU award: validating changes and upgrades in networked software

**OUCL's Daniel Kroening has been awarded a new £426,000 EU grant for research into improving the reliability of networked software systems.**

The PINCETTE project is aimed at reducing the costs of failure detection and fault removal during software changes. The research targets the problem of analysing and validating complex systems upgrades. Under current practice, each change usually requires an expensive revalidation of the whole system, or is simply not checked thoroughly, thus potentially introducing new errors into the system design. State-of-the-art testing and formal verification tools are not optimised to validate system changes and upgrades, but instead focus on a single program version only.

PINCETTE will introduce pioneering technologies to enable systematic component substitutability checks, focusing on various classes of upgrades that are not solved by other approaches.

One of the main innovations of the PINCETTE approach is the integration of technologies that work at different levels or detail (abstraction), addressing different classes of problems and sharing a common solution framework. Another advantage of the approach for upgrade analysis is that models of the software system will be automatically extracted from either the source code (written in C or C++) or by monitoring system execution at run-time. Driven by these models, the PINCETTE paradigm will draw together and use the various

system upgrade checks throughout the system lifecycle from the design phase to the system deployment. It will therefore ensure strong coverage of safety and security issues of the evolving networked systems. The results of the analysis will be accompanied by quality metrics, allowing the user to measure the degree of confidence in these results, and to estimate the need for further analysis.

The PINCETTE project is a Framework Program 7 STREP, made up of a consortium of partners working with the Oxford team on this important area. Partners include: IBM Israel, the Universities of Lugano in Switzerland and Milano-Bicocca in Italy, the Technical Research Center of Finland (VTT), ABB, and Israeli Aerospace Industries.



# Quantomatic: semi-automated graphical reasoning

Graphical theories are a powerful and widely applicable tool in virtually all branches of science and mathematics. They unite our innate ability to reason about visual and spacial constructions with the intuitive notion of a collection of items and relationships. Aleks Kissinger and Alex Merry from OUCL's Quantum Group explain:


We are particularly interested in a 'box-and-wire' style of diagram, as pictured. These diagrams use nodes to represent processes or machines, and wires to represent the connections between these machines (i.e. channels for data to flow). We can allow wires to 'hang free', or even be connected to themselves.'

Such diagrams could represent a computer network, a logical circuit, or even the evolution of a quantum state, the last being of particular interest to us. We can reason about these diagrams using graph rewrite theories, a collection of rules that specify the replacement

of one piece of a graph with another. We have formalised the operation of performing these kinds of rewrites, and have implemented automated graph rewriting in a tool called Quantomatic. Quantomatic is a collaborative effort between the Computing Laboratory and the School of Infomatics at the University of Edinburgh. It allows box-and-wire diagrams to be drawn and manipulated, and rewritten using a graphical theory.

This theory can be entered into the application and used to rewrite the graph in a consistent way, typically to reduce the graph to a simpler but equivalent form. The user can either exert complete control over the rewriting process by choosing specific rewrites to perform, or allow quantomatic to apply rules automatically until it is interrupted or runs out of rules to apply.

In a well-behaved theory (one that is confluent and terminating), this produces



a unique normal form: the simplest possible equivalent graph.

Having a tool with a sound mathematical underpinning to perform this work simplifies working with complex diagrams and makes their manipulation less error-prone. Although the visualisation of the graphs is currently tied to a particular type of diagram we use in the Quantum Group for reasoning about complementary observables and quantum entanglement, we intend to generalise it for easy application to other areas, such as boolean circuits and concurrency.

Quantomatic is an open source project, so is freely available at <http://dream.inf.ed.ac.uk/projects/quantomatic>

## And in other news...

- OUCL's Information Systems Group scored five accepted research papers, out of five submitted, at the 2010 International Conference on **Very Large Data Bases** (VLDB 2010) held in Singapore in September 2010.

- A large contingent from OUCL got papers accepted for presentation at the Annual International Conference of the **IEEE Engineering in Medicine and Biology Society** (IEEE/EMBC) in Buenos Aires in September.

- **Joel Ouaknine** has appeared as an invited speaker at a number of recent conferences including the 37th International Colloquium on Automata Languages and Programming (ICALP 2010), 26th Conference on the Mathematical Foundations of Programming Semantics (MFPS 2010), and Automatic Verification and Analysis of Complex System (AVACS 2010).

- Blanca Rodriguez gave a keynote lecture at the **International Congress on Electrophysiology** and at the Institut de Medecine Theorique in Lyon.

## Top theoretical computer science research

"Probabilistic Model Checking of Complex Biological Pathways" by OUCL's Marta Kwiatkowska and David Parker, recently received the Top Cited Paper award in the journal *Theoretical Computer Science*.

The paper was co-authored with Gethin Norman (formerly of OUCL, but now at the University of Glasgow), John Heath and Oksana Tymchysyn (the later formerly) of the University of Birmingham. Of the papers published in *Theoretical Computer Science* journal, this was the most heavily referenced in other academic research articles for the period 2005-2010.

Probabilistic model checking is a formal verification technique that has been

successfully applied to the analysis of systems from a broad range of domains, including security and communication protocols, distributed algorithms and power management. The award-winning paper illustrates the applicability of probabilistic model checking to a complex biological system: the Fibroblast Growth Factor (FGF) signalling pathway. FGFs are vital to processes including wound-healing and embryonic development.

The project uses the probabilistic model checker PRISM – a tool for formal modelling and analysis of systems which exhibit probabilistic behaviour. PRISM was created at the University of Birmingham, and since 2007 has been developed and maintained at Oxford.

## OOPSLA!

**Max Schaefer and Oege de Moor from the Programming Tools Group presented a paper on 'Specifying and Implementing Refactorings' at this year's OOPSLA conference in Reno, Nevada.**

This work, undertaken as part of the Aspect Refactoring Tools project, presents a new approach to describing software refactorings

that yields high-level specifications which are precise enough to cover real-world languages like Java. These specifications give rise to implementations that are compact, yet match existing implementations for features and outperform them in terms of correctness. OOPSLA is widely recognised as one of the premier conferences on object oriented programming.

## When a picture really does ‘paint’ words

OUCL computer scientists have won a Google Research Award to carry out preliminary research on a system that automatically generates a natural language description of a visual scene, as in the example pictured left.

OUCL’s Stephen Pulman, Phil Blunsom, and Nigel Crook aim to explore an existing analogy between machine translation and object description, but to use contemporary structured translation models, particularly synchronous tree-substitution grammars. This research combines state-of-the-art computer vision with natural language processing. Images will be segmented into regions that correspond to objects in the visual scene. Relationships between the segments will then be identified based on spatial prepositions (e.g. ‘by’, ‘above’, ‘under’) and possibly comparative adjectives (e.g. ‘brighter’, ‘further’), or some simple verbs (‘sitting’, ‘running’), forming a graph structure.

A probabilistic model using synchronous tree-substitution grammars will be used to generate a natural language sentence that describes the original image by relating the image graph to a natural language dependency tree.

They envisage the process of generating a description of an image as requiring in the following three phases (i) the relationships between the objects in the segmented




image are represented as a labelled graph, then (ii) a complete spanning tree is formed from the graph, from which subtrees will be extracted corresponding to components of the image. Finally, (iii) a synchronous tree substitution grammar will map subtrees to dependency trees, and this will be intersected with a language model to generate a sentence.



Photo by: Darrell Godlman

‘A man on a bench under a tree’

## Integrating database technologies and description logics for expressive ontology-based data access

The Information Systems Group at OUCL has recently been awarded a £683,000 EPSRC grant to fund the ExODA project: research into the integration of ontology and database technologies for ontology-based data access.

Sources of semi-structured, overlapping, and semantically-related data on the web are currently proliferating at a phenomenal rate, which has created a demand for more powerful and flexible information systems. This new generation of information systems will need to integrate incomplete and semi-structured information from heterogeneous sources, employ rich and flexible schemas, and answer queries by taking into account both knowledge and data.

Ontology-based data access has recently been proposed as an architectural principle for such systems. The main idea is to develop a unified view of the data by describing the relevant domain in an ontology, which then provides the vocabulary used to ask queries. The information system can

use ontological statements, such as the concept hierarchy, to derive new facts and thus enrich query answers with implicit knowledge. Current ontology-based systems, however, suffer from two main problems. Firstly, the modelling capabilities of ontology languages are often insufficient for real-world applications. They are usually capable of describing only basic relationships with relatively simple powers of reasoning and deduction. Furthermore, ontology languages typically employ the open-world assumption: that the universe of discourse is not limited to the data items explicitly mentioned in the database. However, when answering queries over large amounts of data, the closed-world assumption is often more appropriate. Secondly, the existing ontology-based information systems typically do not scale well to the size of information sets available in the real world.

Database theory and practice can provide partial solutions to these problems. In databases, complex domains can be described and queried using dependencies and the closed-world

assumption, which has allowed practical relational database management systems to scale to very large data sets. However, dependencies can also be used similarly to ontologies to derive implicit knowledge, and they often cannot model arbitrarily large structures and thus do not cover all practical modelling use cases. Furthermore, generalising the query-answering techniques used in practical relational database management systems to the case where information-deriving dependencies must be taken into account is still an open problem.

The ExODA research into the next generation of ontology-based information systems is thus based on a synthesis and an extension of ontology and database systems and techniques, providing data handling capabilities similar to current relational database management systems, but with schemas that are rich, flexible, and tightly integrated with the data. The work is being led by OUCL’s Michael Benedikt, Georg Gottlob, Ian Horrocks, Thomas Lukasiewicz, and Boris Motik.

# Grant for logic-based methods for ontology mapping

OUCL researcher **Bernardo Cuenca Grau** has been awarded a new EPSRC grant.

Bio-medical ontologies such as SNOMED CT (a component of the NHS's patient record service) are becoming core platforms for accessing, gathering and sharing medical knowledge and data.

To exchange or migrate data between ontology-based applications, it is crucial to establish correspondences (or mappings) between their ontologies. Creating such mappings manually is often unfeasible due to the size and complexity of modern ontologies. Therefore, the problem of automatically generating mappings between ontologies has been investigated extensively in recent years. Despite their already mature state, bio-medical

ontologies still pose serious challenges to existing techniques.

The grant's main research hypothesis is based on the observation that existing techniques for ontology mapping often disregard the logic-based semantics of the input ontologies. As a result, they fail to take advantage of the available semantics, and of the highly effective reasoning services for modern ontology languages. The goal is to rethink the foundations underlying the current state-of-the-art thinking in the field by incorporating reasoning-based techniques in each of the steps of the ontology mapping process.

The expected results of the research will be directly relevant to users in the bio-medical domain, where knowledge and data integration is a matter of major concern.


# AI award

**Giorgos Stoilos, Bernardo Cuenca Grau, and Ian Horrocks** have won one of the two outstanding paper awards at the 24th AAI Conference on Artificial Intelligence for their paper 'How Incomplete is your Semantic Web Reasoner?'

The Conference is one of the two most prestigious international conferences in the field of artificial intelligence (AI).

The work presents a formal framework and algorithms for the evaluation of incomplete reasoning systems. Although they do not always compute all answers to queries, incomplete systems are widely used in semantic web applications since they are highly scalable in practice (a recent example being BBC's world cup website).

The techniques presented in the paper represent a promising step towards a deeper understanding of the trade-off between completeness and performance of semantic web reasoning systems.



# 2010: Year of the lung

**Kelly Burrowes is embarking on two new projects involving computational modelling of the respiratory system, that are to be funded through the European Commission's Seventh Framework Programme (FP7).**

Both projects fall within the Virtual Physiological Human call (5.3) which aims to fund patient-specific computational modelling and simulation of organs or systems targeting specific clinical needs such as prediction of diseases, early diagnosis, disease quantification, surgery planning, treatment and training.

The first project, **Airway Disease PR**edicting **O**utcomes through Patient Specific Computational **M**odelling (**AirPROM**), aims to develop a multi-scale integrated patient specific model of the airways within the lung in order to investigate the complex multi-scale

processes occurring in patients with asthma and chronic obstructive pulmonary disease (COPD). This consortium consists of 34 partners establishing an extensive database of knowledge spanning from the omics level to the whole organ.

Synergy-COPD, a smaller project consisting of eight partners across Europe, will develop a simulation environment and a decision-support system aiming at enabling deployment of systems medicine. More specifically, this project will integrate models of oxygen transport throughout the body from the respiratory system to skeletal muscles and ultimately down to the sub-cellular level. The focus will be to highlight underlying mechanisms of COPD phenotypes associated with poor prognosis. Both projects are due to start on 1 February 2011.

# And in other news...

- James Worrell and Stephan Kreutzer recently co-organised **GAMES 2010** – the Annual Workshop of the ESF Networking Programme on Games for Design and Verification, in Oxford. Topics covered include the mathematical and algorithmic analysis of finite and infinite games, the interplay of games with automata theory and logic, and applications of games, automata, and logic for the design and verification of computing systems.

- Bill Morton, former Professor of Numerical Analysis and Director of the Computing Laboratory receives the **De Morgan Medal** the London Mathematical Society's premier award.

- OUCL's Georg Gottlob has been elected a **Fellow of the Royal Society** for his fundamental contributions to both artificial intelligence and database systems.

## OUCL competes at robot World Cup

Scientists from the Oxford University Computing Laboratory took their 'rescue robot' to RoboCup 2010 – the robot World Cup – in Singapore, and achieved fourth place in the simulated rescue competition.

Earthquakes, tsunamis, explosions: disasters like these make natural and urban environments both inaccessible and dangerous for human search and rescue teams. That's when it's time for robots to come to our rescue. Rescue robots can be sent into areas that are too dangerous for humans, or physically impossible to access because of their size, or the lack of oxygen. In an emergency, accurate information on the number and location of victims, and the obstacles that lie between them and the rescue workers, is crucial. Robots have the promise of being able to bring back this information for human search and rescue teams, without putting more human lives at risk.


RoboCup allows teams of computer scientists and artificial intelligence experts from around the globe to demonstrate the latest in intelligent robotics, as their robots battle it out.

The RoboCup-Rescue competition puts virtual robots through their paces in a simulated disaster area, for example a collapsed building.


The robots explore the environment – effectively a giant obstacle course – mapping it as they go, and relaying back vital information to their human handlers. Robots are scored on their mobility, on how much of what's around them they are able to sense, how they work together and how many disaster victims they are able to find successfully.

RoboCup-Rescue therefore sees competitors tackling some of the grand challenges in artificial intelligence: how to get robots to navigate, to sense objects around them (such as victims, and hazards), how to relay this information back to humans, and how to get teams of robots to work effectively with each other.

Oxford University computer scientists – who joined forces with the University of Amsterdam – were the only British team competing in the virtual robot category in this year's RoboCup-Rescue competition.



Doctoral student Jie Ma, with robot



Dr Stephen Cameron, Reader in Computer Science at Oxford University Computing Laboratory, who led the *Amsterdam Oxford Joint Rescue Forces* team said: 'These simulated robots, acting within a virtual environment, are used to push the boundaries of robot intelligence, without the constraints of what is physically possible today. The science behind the virtual robots competing today will form the basis of the next generation of physical robots.'

The competition was launched after the Great Hanshin earthquake, which hit Japan in 1995, killing more than six thousand people. RoboCup-Rescue brings together many smart young minds from around the globe to solve computer science problems that will ultimately save lives in future disaster situations.

The Oxford team are next year planning on working with University of Wales, Newport researchers to take a team of physical robots – similar to those pictured – to the next RoboCup competition in Istanbul, in July 2011. If you are interested in sponsoring them please contact: [Stephen.Cameron@comlab.ox.ac.uk](mailto:Stephen.Cameron@comlab.ox.ac.uk).

## Oxford Prof 'No. 1' says Microsoft Academic Search

Professor Ian Horrocks has been ranked as the number one author in the world wide web research domain by Microsoft Academic Search.

Microsoft Academic Search is a public search engine for academic papers and literature in the field of computer science, which is developed by Microsoft Research Asia. It covers a multitude of scientific papers, conferences, journals, and authors, mainly in computer science and related subjects, and is constantly growing and evolving.

Uniquely, in Microsoft Academic Search, objects in the search results are sorted based on two factors: their relevance to the query and their global importance. The relevance score of an object is computed by its attributes; and the importance score of an object is calculated by its

relationships with other objects. At the time of writing Ian is listed as top author in the World Wide Web field, with an impressive 302 publications and 9479 citations. Ian Horrocks's research focuses on knowledge representation, particularly ontology languages, description logics, and optimised reasoning algorithms. His work on tableau reasoning for very expressive description logics has formed the basis of most description logic reasoning systems in use today, including Racer, FaCT++ and Pellet.

His achievements also include being jointly responsible for development of the OIL and DAML+OIL ontology languages, and he played a central role in the development of the Web Ontology Language OWL.