Proceedings of the

2017 Oxford Computer Science Conference

Department of Computer Science
Wolfson Building
Parks Road
Oxford, OX1 3QD

9th June 2017

Contents

Conference Chairs

Martin Dehnel-Wild, *General Chair*
Katriel Cohn-Gordon, *Conference co-Chair*
Kevin Milner, *Conference co-Chair*

With *wonderful* support from
Julie Sheppard, Sarah Retz, and Lyn Hambridge

Programme Committee

Prince Abudu, Vojtech Havlicek, Chad Heitzenrater, Dennis Jackson, Ahmet Kucuk, Ulrik Lyngs, Alina Petrova, Elizabeth Phillips, Arianna Schuler Scott, Matthew Smith, and Yuan Zhou.

Sponsor: Leonardo

Leonardo is an international leader in electronic and information technologies for defence systems, aerospace, data, infrastructures, land security and protection and sustainable 'smart' solutions. Based mainly in the UK and Italy, the company employs over 47,000 people, with overseas ventures in countries including USA, Saudi Arabia and Brazil. Our Luton site is a world-class Electronic Warfare Centre of Excellence and provides defence systems to detect, evade and counter a wide-range of current-day threats to military platforms. Our equipment protects a wide range of well-known UK platforms such as the Eurofighter Typhoon and the Apache Attack Helicopter.

Leonardo employ across a range of positions and, for our engineering roles, we are looking for enthusiastic applicants from Engineering, Physics, Mathematics and Computer Science backgrounds. For more information, please visit our website:
www.uk.leonardocompany.com/people-careers

1 Programme

Start	End	Session
9:00	9:05	*Welcome Address*, Martin Dehnel-Wild, General Chair
9:05	10:05	**Session 1: Algorithms** (Chair: Ulrik Lyngs) *Predicting Semantic Graphs with Neural Networks* Jan Buys *The Complexity of Counting Compactions and Surjective Homomorphisms* Jacob Focke, Leslie Ann Goldberg, and Stanislav Zivny *Flash Crash Contagion and Systemic Risk: An Agent-Based Model* James Paulin, Anisoara Calinescu, and Michael Wooldridge
10:10	10:30	**Session 2: Lightning Talks** (Chair: Ulrik Lyngs) *Evaluating Manual Intervention to Address the Challenges of Bug Finding in Symbolic Execution* John Galea, Daniel Neville, and Sean Heelan *Molecular Semantics* Luca Laurenti *Balancing cooperation and defection in a two-agent grid-world game: the Malmo Collaborative AI Challenge.* Adrià Garriga, Daniel Furelos Blanco and David Tena Cucala *Search for Quantum Advantage in Restricted Permutational Quantum Computing* Vojtech Havlicek
10:30	10:50	**Coffee Break** (Atrium)
10:50	12:10	**Session 3** (Chair: Chad Heitzenrater) *Exploring Weak Ciphers Usage in Business Aircraft Communications* Matthew Smith, Daniel Moser, Martin Strohmeier, Vincent Lenders, and Ivan Martinovic *Broken Hearted: A Novel Cross-Device Presentation Attack Against ECG Biometrics* Simon Eberz, Nicola Paoletti, Marc Roeschlin, Andrea Patané, Marta Kwiatkowska, and Ivan Martinovic *Model-driven-design of biological experiments: The case for the in silico lab* Daniel Nichol, Peter Jeavons, and Alexander Anderson *Novel Parallel Watershed for Faster Image Partitioning* Varduhi Yeghiazaryan and Irina Voiculescu
12:15	13:30	**Lunch and Poster Session** (Atrium)

Start	End	Session
12:15	13:30	**Lunch and Poster Session** (Atrium) Please see Section 3 for poster titles and abstracts
13:30	14:10	**Keynote:** *From Semantics of Computation to Physics and Back* **Professor Samson Abramsky FRS**
14:15	15:00	**Session 4** (Chair: Kevin Milner) *Composably secure time-frequency quantum key distribution* Nathan Walk *Data flow and signal processing for Airborne ESM* Harvey Alison (Leonardo)
15:00	15:30	**Coffee Break** (Atrium)
15:30	16:10	**Session 5: Quantum** (Chair: Vojetch Havlicek) *Representing encoded operations in quantum computation, with diagrams* Niel de Beaudrap *Completeness of qutrit ZX-calculus for stabilizer quantum computation* Quanlong Wang
16:15	16:40	**Session 6: Lightning Talks** (Chair: Elizabeth Phillips) *ICTs and Attention Management: Evaluating the Emerging Anti-Distraction Market* Ulrik Lyngs *Security Games with Multiple Uncoordinated Defenders* Jiarui Gan, Edith Elkind, and Michael Wooldridge *"Privacy is the boring bit": Perceptions and behaviour in the Internet-of-Things* Meredydd Williams, Jason Nurse, and Sadie Creese *Automatically Verifying Stateful Security Protocols in Tamarin* Nicholas Moore *Cyber cowboys and the wild west of insurance* Daniel Woods
16:45	17:00	**Poster Judging and Awards Presentation**
17:00		**Drinks Reception in Atrium**
19:00		**Drinks and Dinner at Rewley House**

2 Abstracts

Session 1: Algorithms

Predicting Semantic Graphs with Neural Networks
Author: Jan Buys
Abstract: An important goal in Natural Language Processing (NLP) is the automatic comprehension of unstructured text. Parsing sentences to machine interpretable semantic representations is an important step towards that goal. Yet there are many semantic relations which most statistical natural language parsers are not able to predict, due to assumptions to make inference tractable that lead to the widespread use of projective dependency trees. We propose a recurrent neural network (RNN) semantic parser that is able to predict labelled graphs based on Minimal Recursion Semantics (MRS), a linguistically expressive framework for compositional semantics. The stack-based parser builds the graph structure incrementally, as the RNN makes predictions conditioned on the state of the stack. Our parser is more accurate than baseline encoder-decoder RNNs with attention, and a GPU implementation makes it an order of magnitude faster than a high-precision grammar-based parser. The predicted MRS graphs are also more accurate than the upper bound on performance of Abstract Meaning Representation, a commonly used semantic graph representation. This research will enable the application of MRS graphs in end-user NLP systems.

The Complexity of Counting Compactions and Surjective Homomorphisms
Authors: Jacob Focke, Leslie Ann Goldberg, and Stanislav Zivny
Abstract: In Counting Complexity we investigate the difficulty of counting solutions to a given computational problem. #Hom(H) is the widely studied problem of counting homomorphisms from an input graph G to a fixed parameter graph H. This problem framework generalises important problems such as counting independent sets or proper k-colourings of a graph. A k-colouring of a graph G assigns a total of at most k colours to the vertices of G such that adjacent vertices are not of the same colour. In their seminal work, Dyer and Greenhill classify the complexity of #Hom(H) for all undirected graphs H. A natural question is asking for the number of k-colourings that actually use all of the k colours. The corresponding generalisation is counting surjective homomorphisms from input G to parameter graph H, which we denote by #SHom(H). In this work, we give a complete characterisation of the complexity of #SHom(H) and the related problem of counting compactions, i.e. homomorphisms from G to H that cover not only all vertices but also all (non-loop) edges of H. Interestingly, it turns out that the complexity of counting compactions does not coincide with that of #Hom(H).

Flash Crash Contagion and Systemic Risk: An Agent-Based Model
Authors: James Paulin, Anisoara Calinescu, and Michael Wooldridge
Abstract: Systemic risk refers to the risk of catastrophic failure of an entire system as distress propagates around the network of its interconnected components. Events such as 6th May 2010's *Flash Crash* demonstrate the vulnerability of global financial systems to coordinated distress and that such extreme phenomena act as a contagion channel for systemic risk, with a potential significance of trillions of dollars. Previous work has demonstrated the efficacy of network-theoretic concepts for characterising systemic risk, but the microscopic details of securities trading are often abstracted away. However, it is precisely when markets enter distressed modes that microscopic details can radically affect their dynamics. The present work seeks to integrate the macroscopic financial network approach to systemic risk with a realistically-calibrated microscopic behavioural model. The agent-based model we develop allows us to investigate flash crash contagion between assets held by multiple market participants with overlapping portfolios, with particular emphasis on the role of *position crowding* which occurs when many institutions hold similar positions. Our analysis aims to provide novel insights to better inform policy and decision-making by central banks and regulators.

Session 2: Lightning Talks

Evaluating Manual Intervention to Address the Challenges of Bug Finding in Symbolic Execution
Authors: John Galea, Daniel Neville, and Sean Heelan
Abstract: Whilst symbolic execution has shown its ability to discover security relevant flaws in software, it faces significant scalability challenges, such as dealing with the state space explosion problem. A collective intuition is believed, whereby the consideration of domain expert knowledge can help mitigate these limiting factors. However, little formal investigation of this approach has been carried out in previous work.

We present a novel corpus of over 130 bugs in real world software, including LibJPEG, TCPDump and LibTiff. We use the corpus to facilitate a thorough evaluation of KLEE, a symbolic execution engine considered to be state of the art, and classify the frequently occurring software patterns that are deemed problematic. In addition, a set of manual mitigations aimed at administering the underlying issues are assessed. Results show that, despite its inherent limitations, manual intervention is a feasible mechanism to increase both code coverage and the detection of bugs.

Molecular Semantics
Author: Luca Laurenti
Abstract: Molecular systems perform their computation in a noisy environment, and molecular interactions are inherently stochastic. However, molecular functions are surprisingly robust. Unfortunately, the mechanisms used by biological systems to perform their computation are still not very well understood. In this talk, I will first consider scenarios where the noise (stochastic fluctuations) needs to be reduced in order to enhance robustness. With a focus on gene expression, I will show that the basic mechanisms used by biological systems to deal with noise are equivalent to those currently used in electronic circuits. Then, I will consider scenarios where natural algorithms have evolved in order to explore stochasticity. I will consider our recent work and future challenges.

Balancing cooperation and defection in a two-agent grid-world game: the Malmo Collaborative AI Challenge.
Authors: Adrià Garriga, Daniel Furelos Blanco and David Tena Cucala
Abstract: The Malmo Collaborative AI challenge is a video-game-like rendition of the classic Stag Hunt game. In this case, two agents and a pig are inside a pen, which is divided in a square grid. The agents can cooperate to corner the pig, or exit the pen individually. We control one of the agents and, for us, collaborating to catch the pig has a much higher payoff than exiting the pen. However, the other agent may not be so inclined. The question then is how to determine whether the other agent will cooperate, and accordingly make movements in the grid to catch the pig or exit. One game instance is 10 episodic interactions with the same agent.

Traditional reinforcement learning (RL) without modifications is not well-suited to this problem: the other agent's action history plays a significant role in what the optimal action is, so the environment is not Markovian. On the other hand, the state space is not large. Our agent, which won the competition, employs a mixture of Bayesian inference to model the other agent's actions, and Monte-Carlo Tree Search to plan the optimal action. I will explain in more detail the algorithm of our agent, including tricks done to make inference data-efficient enough, and planning fast enough.

Search for Quantum Advantage in Restricted Permutational Quantum Computing
Author: Vojtech Havlicek
Abstract: As small scale quantum computers are becoming reality, it is important to provide evidence of quantum computing capabilities *provably* beyond classical. That is, accounting for limitations of the current hardware such as small computational register size (number of qubits) or limited set of operations, what is the simplest computational task for which the quantum device clearly outperforms *any* classical computer? It is for example known that Shor's algorithm for factoring provides computational advantage compared to the best *known* classical algorithm, although it is hardly possible to practically demonstrate in the near-term. The largest number factored classically has 768-bits, while the best implementation of Shor requires $2n+3$ qubits per n-bit input. To factor beyond classical, one hence needs a quantum computer with at least about $\sim$1550 qubits. The current state-of-the-art hardware has about ~20 qubits, making a quantum advantage demonstration by factoring unfeasible. It is natural to search for computational tasks allowing for simpler quantum advantage demonstration. I will outline the progress towards answering this and sketch out the quantum computational model of Permutational Quantum Computing, which was believed to achieve a substantial quantum advantage with moderate hardware requirements. I will conclude with necessary conditions on PQC problem instances to allow for quantum speedup.

Session 3

Exploring Weak Ciphers Usage in Business Aircraft Communications
Authors: Matthew Smith, Daniel Moser, Martin Strohmeier, Vincent Lenders, and Ivan Martinovic
Abstract: Aircraft Communications Addressing and Reporting System (ACARS) is a text-based avionic data link in wide use by commercial and private aircraft alike. Having been initially deployed the late 1970's, it is now used for many different purposes - some of which involve the transfer of potentially sensitive data such as position or destination. Due to the fact that no security exists by default, a number of post-hoc, proprietary solutions have been implemented. We examine one of these - a monoalphabetic substitution cipher in use by private business aircraft - by collecting ACARS data using commercial off-the-shelf components over a period of six months. This type of cipher offers no meaningful security and is trivially breakable by an attacker with limited resources, with the cipher sharing nine keys between all aircraft using it, even further weakening the security of the approach. We show that the vast majority of these aircraft attempt to protect their privacy by hiding their movements from flight tracking websites; many go on to transmit position or intention information under the 'protection' of this cipher thus breaching their privacy. Overall, we demonstrate that the cipher offers a false sense of security to those who are using it, since it undermines other efforts to protect privacy.

Broken Hearted: A Novel Cross-Device Presentation Attack Against ECG Biometrics
Authors: Simon Eberz, Nicola Paoletti, Marc Roeschlin, Andrea Patané, Marta Kwiatkowska, and Ivan Martinovic
Abstract: In this work we present a systematic presentation attack against ECG biometrics. We demonstrate the attack's effectiveness using the Nymi Band, a wrist band that uses electrocardiography (ECG) as a biometric to authenticate the wearer. The ECG signals are injected using an off-the-shelf audio player connected to the Nymi Band. In two sets of experiments we collect data from a total of 41 participants using a variety of ECG monitors, including a medical monitor, a smartphone-based mobile monitor and the Nymi Band itself.

We use the first dataset to understand the statistical differences in biometric features that arise from using different measurement devices and modes. Such differences are addressed through the automated derivation of so-called mapping functions, whose purpose is to transform ECG signals from any device in order to resemble the morphology of signals recorded with the Nymi Band.

Finally, we enroll users into the Nymi Band and test whether data from any of our sources can be used for a signal injection attack. Depending on the source device, we achieve a success rate of 43% when using raw data, and 62% after applying the mapping function. While we demonstrate the attack on the Nymi Band, we expect other ECG-based authentication systems to suffer from the same, fundamental weaknesses.

Model-driven-design of biological experiments: The case for the in silico lab
Authors: Daniel Nichol, Peter Jeavons, and Alexander Anderson
Abstract: In the biological sciences, mathematical modelling is often used as a tool to codify biological assumptions and generate hypotheses to be tested empirically. The results of these model-driven experiments are in turn used to re-examine the underlying assumptions and generate new testable hypotheses, completing a model/experiment cycle that is central to integrative mathematical bioscience. However, this approach can fail because of the biological complexity that underpins even the simplest experimental design. Indeed, erroneous conclusions can be drawn from biological observations as a result of incorrect intuition, resulting in the generation of new, but wrong, hypotheses. We demonstrate how explicit in silico simulation of biological experiments can serve as an experimental design tool to minimise these costly errors, and highlight the potential of this model-driven-design approach through examples drawn from our own work in experimental evolution. Through simulation of published in vitro evolution studies performed in both bacterial and cancer cell populations, we identify potentially incorrect conclusions regarding strategies to combat drug resistance. We then present preliminary results from our own evolution experiments designed using a model-driven paradigm. Finally, we conclude with a brief overview of model-driven experimental design in the mathematical biosciences and highlight some issues and their potential solutions.

Novel Parallel Watershed for Faster Image Partitioning
Authors: Varduhi Yeghiazaryan and Irina Voiculescu
Abstract: The watershed transform is an established image processing technique which represents the image as a topographic map and separates 'catchment basins', concavities around local minima, into different regions. There are popular sequential watershed algorithms, whereas parallel watershed algorithms gained interest only in recent years with the rise of general-purpose computing on graphics processing units.

We present a novel parallel watershed algorithm. While our approach shares similarities with previously suggested parallel watersheds, it stands out by addressing plateaux in a novel fast way with virtually no time overhead. We provide a description of the suggested algorithm along with a detailed discussion of the plateaux resolution. Our implementation of the algorithm relies on CUDA, a parallel computing platform and programming model by NVIDIA, and successfully transforms large medical images within seconds.

We report encouraging experimental results from having run our algorithm on real–life images and medical scans. In addition, we suggest using the repetitive application of our watershed for fast parallel hierarchical image partitioning mostly appropriate for large 3D medical data.

Session 4

Composably secure time-frequency quantum key distribution
Author: Nathan Walk
Abstract: Quantum key distribution (QKD) is the process of generating a common random key between two parties using a quantum communications protocol. The power of this method is that the security of the key distribution, and hence the subsequent communication via a one time pad, is established while making no assumptions about the technological capabilities or computational power of an eavesdropper. Instead, security stems primarily from the laws of quantum physics. On the theoretical side, one looks to derive a security proof that is composably secure such that the secret keys from multiple protocols can be safely combined. Practically, one searches for schemes that maximise both the raw clock-rate (the number of transmissions per second) and the number of secure bits per transmission to achieve the largest overall secret key rate at a given distance. I will present a protocol based upon the spectral properties of single photons, which performs well on both counts.

Data flow and signal processing for Airborne ESM
Authors: Harvey Alison (Leonardo)
Abstract: Electronic Warfare is defined as "action involving the use of electromagnetic (EM) energy to determine, exploit, reduce or prevent hostile use of the electromagnetic spectrum (EMS), and action which retains friendly use of the electromagnetic spectrum." This presentation introduces ESM (electronic support measures), a particular application challenge within EW, and presents an overview of a particular technical solution. Relevance to computer science and algorithm development is highlighted, and a view presented of future challenges for research in this rewarding domain.

Session 5: Quantum

Representing encoded operations in quantum computation, with diagrams
Author: Niel de Beaudrap
Abstract: Quantum computing is shifting from an abstract computational model towards deployment as a technology. However, to be scalable and effective, quantum devices will need to protect data against errors. A leading choice of error correction — particularly for modular or distributed architectures — is the surface code with logical two-qubit operations realised via "lattice surgery". These operations consist of "merges" and "splits", which do not fit easily in the usual description of quantum transformations by unitary transformations. This raises the question of how best to reason about lattice surgery, for instance to efficiently manage resources in a quantum computer. We demonstrate that lattice surgery is well represented using the "ZX calculus" — a diagrammatic notation for tensor networks in terms of interacting Frobenius algebras, using labelled graphs. The "split" operations can be represented by nodes with in-degree 1 and out-degree 2, and the "merge" operations can be represented in terms of nodes with in-degree 2 and out-degree 1, in an axiomatic system for rewriting labelled graphs. This provides us with a starting point for developing more nuanced compilation tools for quantum computers, using abstract notation to describe the management of resources motivated by practical choices in quantum architectures.

Completeness of qutrit ZX-calculus for stabilizer quantum computation
Author: Quanlong Wang
Abstract: Recently metaplectic non-Abelian anyons were shown to be naturally connected to ternary (qutrit) logic in contrast to binary logic in topological quantum computing, based on which there comes the Metaplectic Topological Quantum Computer platform. Comparing to higher dimensional based quantum computers, qutrit-based computers are space-optimal in certain sense. On the other hand, the current theoretical tools for qutrit-based quantum computing are dominated by quantum circuits. However, the quantum circuit notation has a major disadvantage: It is not easy to transform one circuit diagram into another, since complicated diagrammatical equations will be involved. By contrast, the ZX-calculus which is founded on the framework of symmetric monoidal categories, has intuitive and simple rewriting rules represented by diagrams for quantum computing. So far the ZX-calculus has been successfully applied to (topological) measurement-based quantum computing and quantum error correction within stabilizer formalism. Here we show that a qutrit version of ZX-calculus is complete for pure qutrit stabilizer quantum mechanics, which means any equality that can be derived using matrices in the stabilizer system can also be derived diagrammatically.

Session 6: Lightning Talks

ICTs and Attention Management: Evaluating the Emerging Anti-Distraction Market
Author: Ulrik Lyngs
Abstract: Information communication technologies like laptops and smartphones have enabled users to do anything anywhere anytime. However, with this ability comes a challenge to manage attention and avoid perpetual distraction, with which most users struggle. As a developer response, a growing market has emerged for anti-distraction tools that adjust user interfaces in ways intended to help users regulate their behaviour. At a high level, these tools take one of three routes: remove distracting functionality from the user's behavioural options (e.g. *Freedom*, *SelfControl*, *News Feed Eradictor*), track user behaviour over time and visualise app usage to the user (e.g. *RescueTime*, *meTime*), or directly punish or reward undesired user behaviour (e.g. *Forest*, *Write or Die*). Some of these tools now have millions of users, but no research exists on whether they actually change user behaviour, which cognitive frameworks explain their effects, and what design lessons can be learned from them. My DPhil project will map the design space of anti-distraction tools, evaluate how currently available tools influence user behaviour and perception, and prototype and test a new tool which allows users to set custom latencies on app and website loading times.

Security Games with Multiple Uncoordinated Defenders
Authors: Jiarui Gan, Edith Elkind, and Michael Wooldridge
Abstract: Stackelberg security games have received much attention in recent years. While most existing work focused on single-defender games, there do have some real-world scenarios involving more than one defenders (e.g., counter-piracy actions over international waters). This motivates us to consider security games with multiple defenders, and particularly, uncoordinated defenders each of whom optimizes their own utility function. We analyse the existence of exact and approximate equilibria under different settings, and answer the associated computational problems such as the complexity of deciding the existence of equilibria when the existence is not guaranteed.

"Privacy is the boring bit": Perceptions and behaviour in the Internet-of-Things
Authors: Meredydd Williams, Jason Nurse, and Sadie Creese
Abstract: In opinion polls and surveys, the public claim to value their privacy. However, individuals often disregard the principle when using social networks and novel gadgets. This contributes to an opinion-action disparity labelled the 'Privacy Paradox'. The Internet-of-Things (IoT) frequently places privacy at risk, whether through poor authentication or ubiquitous surveillance. However, despite research interest in the area, the Paradox has been underexplored. In addressing this, we first survey the privacy opinions of the general public (N = 170). Reflecting on both IoT and less-novel products, we find users knowingly purchase risky devices. With these gadgets rated both less usable and less familiar, we assert that the IoT constrains private behaviour. To explore this hypothesis, we perform contextualised interviews with 40 participants. In these 20-minute dialogues, owners discuss both their opinions and actions with a personal device. We find the Paradox is significantly more prevalent in the IoT, justified by short-term necessity and a lack of awareness. We highlight the qualitative comments of ordinary users, who criticise product design and social norms. We finish by proposing solutions aligned with participant recommendations. These include IoT-focused awareness campaigns and opt-out privacy settings.

Automatically Verifying Stateful Security Protocols in Tamarin
Author: Nicholas Moore
Abstract: Security protocols are increasingly moving to using more and more state, with the aim of achieving stronger security properties. Unfortunately, the current 'state of the art' in protocol verification tools do not always handle state well - often causing them to fail to terminate, require additional expertise to model the protocol, or even preventing the protocol from being modelled at all.

In this talk I'll be going over some of my work to improve the protocol verification tool Tamarin, as well as looking at other attempts to do so.

Cyber cowboys and the wild west of insurance
Author: Daniel Woods
Abstract: At the turn of the century, Bruce Schneier described a world in which the "computer security industry will be run by the insurance industry". Such a vision was motivated by an analogy to property insurance where sprinkler systems are installed because insurance policies demand it. Fifteen years on, the cyber insurance market has been described as the "wild wild west" of the insurance industry, not least because of the continuously shifting threat landscape. The talk explores how the reality of the market sits between the vision of centralised control outlined by Schneier and comparisons to the lawless American frontier. The talks reflects on a series of interviews with cyber insurance professionals working in London. By taking a high level view of the problem, we hope other researchers will see links to their own research.

3 Posters

Hearing Attacks in Sonified Network Data
Authors: Louise Axon, Sadie Creese, and Michael Goldsmith
Abstract: Sonification is a promising technique for signalling network attacks and network-security information, which could complement the range of security-monitoring tools currently used in SOCs. Prior work in sonification for network monitoring has not assessed the effectiveness of the technique for enabling users to detect network attacks accurately and efficiently. We investigate the concept that sonification can represent network datasets such that a range of network attacks can be detected and identified by humans. In this work, we present a parameter-mapping, musical sonification system for network packet captures. We report the results of a user experiment to assess the utility of the sonification for signalling attacks.

Available Labelled Flexible Realistic Data-generator
Authors: Alastair Janse van Rensburg and Louise Axon
Abstract: We present a method for the generation of full synthetic network-attack datasets for the evaluation of network-monitoring tools. Measuring attack-detection performance using network datasets is important for the improvement and evaluation of developed intrusion-detection tools. There is in general a lack of appropriate public datasets for assessing network-monitoring systems. Real-world datasets are often unlabeled or highly anonymised, and the network-attack datasets available are limited. Synthetic dataset generation is an approach to solving this problem: current solutions enable the injection of attacks into existing network datasets. We address the requirement for a tool that enables network-security researchers to generate complete synthetic network-attack datasets for tool evaluation. Our contribution is ALFRED, a tool to generate labeled, realistic network-attack datasets. Default settings are available based on existing research, and can also be specified by users using a graphical user interface. We intend this tool as a ready source of synthetic datasets for use by researchers developing network-monitoring tools.

Privacy-Preserving Targeted Advertising for Mobile Devices
Authors: Yang Liu and Andrew Simpson
Abstract: Targeted Mobile Advertising (TMA) has emerged as a significant driver of the Internet economy. It enables organisations to tailor advertisements to specific consumers by analysing the personal information collected from consumers' mobile devices. Although TMA offers great benefits to advertisers, the privacy concerns associated with it may reduce the advertising effectiveness. It follows that there is a need for an advertisement selection mechanism that can support the existing TMA business model in a manner that takes into account consumers' privacy concerns. The research described in this poster session explores the delicate balance between privacy and utility in this emerging area and presents such an ad selection mechanism that has the potential to provide benefits to both consumers and advertisers.

Closed-loop quantitative verification of rate-adaptive pacemakers
Authors: Nicola Paoletti, Andrea Patane, and Marta Kwiatkowska
Abstract: Rate-adaptive pacemakers are cardiac devices able to automatically adjust the pacing rate in patients with chronotropic incompetence, i.e. whose heart is unable to provide an adequate rate at increasing activity levels. Rate-adaptation parameters depend on many patient-specific factors, and effective personalisation of such treatments can only be achieved through extensive exercise testing, which is normally intolerable for a cardiac patient. In this work, we introduce a data-driven and model-based approach for the automated verification of rate-adaptive pacemakers and analysis of personalised and population-wide treatments. We develop a dual-sensor pacemaker model based on accelerometer and QT interval metabolic sensors. Our approach enables personalisation through the estimation of heart model parameters from patient data, and closed-loop analysis through the online generation of model-based physiological signals. To capture the probabilistic and non-linear dynamics of the heart, we define a probabilistic extension of timed I/O automata with data, and employ statistical model checking for quantitative verification of rate modulation. We evaluate our pacemaker design on three subjects and a pool of virtual patients, demonstrating the potential of our approach to provide quantitative insights into the closed-loop behaviour of the device.

Applications of Reinforcement Learning to Medical Image Segmentation
Authors: Edoardo Pirovano and Irina Voiculescu
Abstract: This project investigates the application of reinforcement learning techniques to the segmentation of medical images. In particular, we present a novel approach that is based on learning how to grow a selection with regions obtained from an image partition forest (IPF) based on various attributes of the regions. Our algorithm is almost automatic, although we discuss why it cannot be classed as entirely automatic.

We then proceed to quantitatively evaluate this method against two datasets of manually segmented gold standards of femurs in MRI scans of knees. In the first dataset, we achieve an average DSC of 0.97±0.01 on unseen images after training. This is similar to the performance of current state-of-the-art algorithms.

The second dataset will present more challenges as the scans in it are lower resolution, less uniform and have less contrast between the bone and surrounding tissue. In this dataset, our algorithm performs worse achieving an average DSC of 0.89±0.07 which, while reasonable, is inferior to that achieved by modern algorithms. Nonetheless, we discuss some advantages this method offers over others and present ways in which it could be improved in future.

Automated Experiment Design for Efficient Verification of Parametric Markov Decision Processes
Authors: Elizabeth Polgreen, Viraj Brian Wijesuriya, Sofie Haesaert, and Alessandro Abate
Abstract: We present a new method for statistical verification of quantitative properties over a partially unknown system with actions, utilising a parameterised model (in this work, a parametric Markov decision process) and data collected from experiments performed on the underlying system. We obtain the confidence that the underlying system satisfies a given property, and show that the method attains data efficiency and is robust to the amount of data available. These characteristics are achieved by firstly exploiting parameter synthesis to establish a feasible set of parameters for which the underlying system will satisfy the property; secondly, by actively synthesising experiments to increase amount of information in the collected data that is relevant to the property; and finally propagating this information over the model parameters, obtaining a confidence that reflects our belief whether or not the system parameters lie in the feasible set, thereby solving the verification problem. We implement our method and evaluate its efficiency, robustness and ability to handle unbounded-time properties over partial probabilistic models with actions and nondeterminism.

4 Keynote

This year the committee is thrilled to be hosting Professor Samson Abramsky FRS, who will provide the 2017 OxCSC Keynote address.

The keynote is entitled “From Semantics of Computation to Physics and Back”.

Biography: Samson Abramsky is Christopher Strachey Professor of Computing and a Fellow of Wolfson College, Oxford University. Previously he held chairs at the Imperial College of Science, Technology and Medicine, and at the University of Edinburgh.

He holds MA degrees from Cambridge and Oxford, and a PhD from the University of London. He is a Fellow of the Royal Society (2004), a Fellow of the Royal Society of Edinburgh (2000), a Member of Academia Europaea (1993), and a Fellow of the ACM (2014). His paper “Domain theory in Logical Form” won the LiCS Test-of-Time award (a 20-year retrospective) for 1987. The award was presented at LiCS 2007. He was the Clifford Lecturer at Tulane University in 2008. He was awarded the BCS Lovelace Medal in 2013. He received the Alonzo Church Award for Outstanding Contributions to Logic and Computation in 2017.

He has played a leading role in the development of game semantics, and its applications to the semantics of programming languages. Other notable contributions include his work on domain theory in logical form, the lazy lambda calculus, strictness analysis, concurrency theory, interaction categories, and geometry of interaction. More recently, he has been working on high-level methods for quantum computation and information. He introduced categorical quantum mechanics with Bob Coecke. He introduced the sheaf-theoretic approach to contextuality and non-locality with Adam Brandenburger, and has contributed extensively to developing a structural theory of contextuality and its applications.

5 Conference Dinner

This year's conference dinner will be hosted at Rewley House. A drinks reception will be held starting at 19:00, with dinner commencing at 19:30.

Rewley House is on Wellington Square, and is accessible via a 6 minute walk, across St Giles and Little Clarendon Street. Directions to Rewley House can be found below.

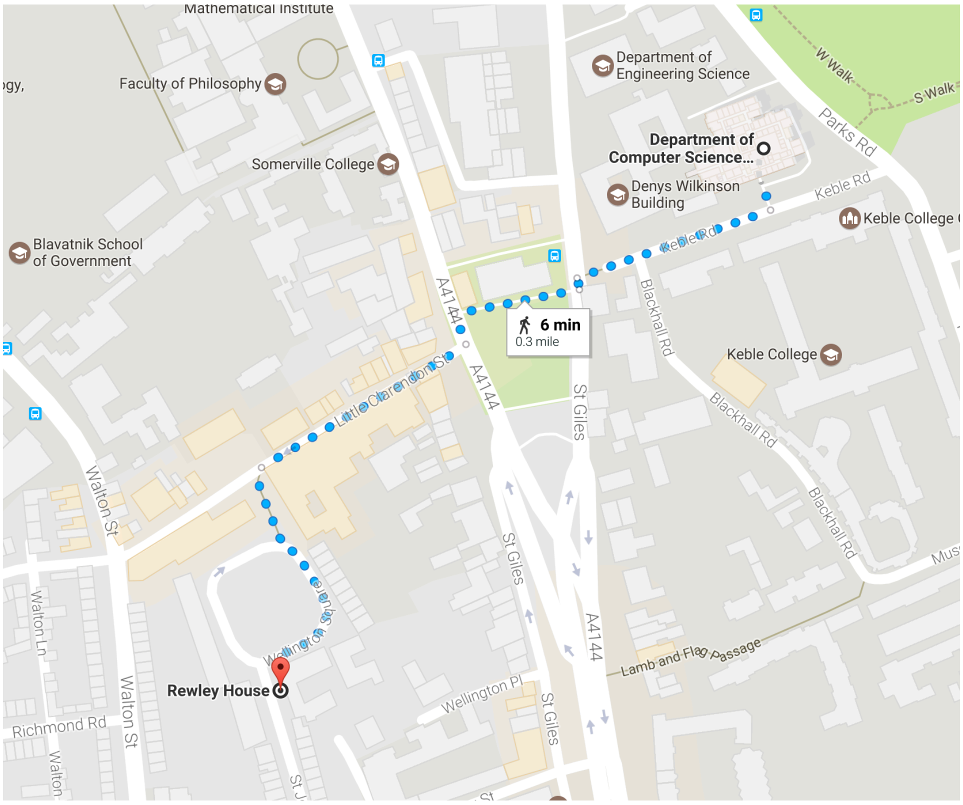

6 Awards

The following awards will be selected from the participants of OxCSC by a panel of judges. Each award recipient will receive a certificate and £150.

Best Abstract

Best Presentation

Best Poster

The judges for this year's awards are:

- **Professor Samson Abramsky** FRS, *Christopher Strachey Professor of Computing* and 2017 OxCSC Keynote Speaker
- **Professor David Kay**, *Director of Graduate Studies in the Department of Computer Science, University of Oxford*
- **Dr. Harvey Alison**, *EW Capability Manager, Leonardo MW Ltd.*
- **Martin Dehnel-Wild**, *General Chair, OxCSC 2017*
- **Kevin Milner**, *Co-Chair, OxCSC 2017 and CoGS President*

These awards were made possible by kind sponsorship from Leonardo MW Ltd.

LEONARDO