The Ring

The Newsletter of the University of Cambridge Computer Laboratory Graduate Association
Chairman: Professor Ian Leslie
President: Stephen Allott

LATEST TRENDS IN WIRELESS COMPUTING

A talk to the Cambridge Computer Lab Ring by Andy Hopper, Professor of Communications Engineering, Laboratory for Communications Engineering, University of Cambridge

Meeting report by Stephen Allott
Director of Development, Cambridge University Computer Lab

With widespread adoption of mobile computing and the vast sums spent on 3G licences, a large audience gathered to find out where wireless computing is heading. 802.11 wireless LANs have grown rapidly. Rollout of 3G services continues to be delayed. The telecoms and technology industry bust coupled with recession has further clouded the issue, making the rate of adoption of wireless data services even harder to predict. For these reasons, we invited Andy Hopper to talk about the latest trends in wireless computing. Andy first joined the Computer Laboratory in 1974 completing his PhD in 1978. He remained a member of the Computer Laboratory faculty until taking up his present chair in the Engineering Department in 1997. His distinguished academic career sees his list of publications run to 6 pages including work on the real Cambridge Ring, the data networking project conducted in the late 1970s. He has also had a successful business career. He has been a founding director of 11 companies and is an adviser to 2 separate venture capital firms. The Computer Lab has been able to welcome him back as a neighbour. The Laboratory for Communications Engineering, which he heads, moved into the 2nd floor of the William Gates Building last term.

Wireless communications systems can be classified according to speed, on one axis, and whether they are point to point, point to multi-point or mobile, on the other axis. While cellular is an example of a slow mobile technology, ultra wideband and fixed wireless access are examples of fast point to point technologies. Bluetooth and 802.11 are examples of medium speed point to multipoint technologies. Andy covered 3 topics in detail:

- Wireless LAN and PAN (802.11 and Bluetooth)
- Fixed wireless access
- Ultra wide band.

802.11 AND BLUETOOTH - WIRELESS LAN AND PAN

Bluetooth is coming along well. There are still some questions about interoperability between Bluetooth and 802.11. Bluetooth will soon be delivering 3 MBits/second and the price of silicon is falling. Soon we shall see a Bluetooth web server on a chip.

Wireless LAN has been a huge growth area. “Ethernet rules OK” in Andy’s words. Although the radios are problematic, faster versions are coming along soon. All the issues of Quality of Service, security and compression are being dealt with but in a chaotic way. Eventually however, 802.11 will “jam up” through over use. It also needs to be proactively managed.

Andy made a few comments on the slow mobile technologies. 3G will not make money. All is doom and gloom. 802.11 will take over in hot spots. The audience agreed that the future for 3G is poor. As far as satellite is concerned, Andy tested the data rate on a walk to the top of Greenland, proving that satellite really does work from anywhere.

He also commented on embedded radio systems. These fit below Bluetooth and are a very simple radio networking system with low but variable data rates. It’s a switch with a radio on it. There are 2 sorts of issues. Engineering issues and logical operational issues. An example of the latter is how it uses sleep mode, which is important to preserve power. The problem is that it cannot communicate or be woken up when in sleep mode. One solution is to use a mains powered base station to act as a relay when anything wakes up and sends a signal. This is an example of a fundamental issue of how to synchronise low power devices. The Zigbee alliance, on the West Coast of the US, has been formed to research these issues.

FIXED WIRELESS ACCESS

Although fibre can deliver huge bandwidth, fixed wireless access can deliver very large bandwidth indeed; such as 48Mbps at distances of up to 20 kms. This is much much more than DSL or cable modem. Andy described the Cambridge Broadband solution, Vectastar, in some detail. It is a complete systems design which is designed to interoperate. Some fixed wireless solutions have become discredited because of a lack of systems design. Essentially it is an ATM switch in the air. It uses fast error correction based on ARQ (like the old Cambridge Ring approach) rather than FEC.

A standard 4 sector base station can be deployed enabling subscribers to connect at high data rates and significant distances. Targeted at the business market, the system is too expensive for retail use.

Spectral efficiency is good because radio spectrum can be re-used. Opposite quadrants in the 4 sector approach can use the same frequency. Dynamic modulation is used to determine the data rate. A test packet is sent and the modulation level is then set to maximise data transfer rates given the quality of reception. The higher the base station above the ground or roof levels, the better. There is a pilot installation in Cambridge which has a reasonable performance up to 27 kms. This would be a good backhaul system for 802.11 based in local village hotspots. A 4 colour map, i.e. using 4 frequencies, can cover any size of area. Andy showed how one can daisy chain from one base station to another so that coverage can be extended into an area entirely using radio transmission for the backhaul as well as the access.

Operators using this technology will first have to obtain radio spectrum from the regulator. An auction is planned for 3.4 GHz. Andy envisages local authorities being the operator and predicts that fixed wireless access supplying 802.11 will be the way of the future. Andy also pointed out that this system has taken 10 years to develop. John Porter, a computer lab grad, has been working on it.

ULTRA WIDE BAND

Ultra wide band goes back a long way ... to Guglielmo Marconi himself 100 years ago. Data rates over short ranges can be very high such as 500 Mbps over 5 metres. Signal strength decays rapidly. Current applications include linking digital cameras to laptops and so on. One can also use the technology to do precise location of devices at up to 50 metres. The US FCC has licensed UWB for in building applications whereas in the UK some vested interests are holding things back.

The Laboratory for Communications Engineering has been researching the use of UWB for location of objects and also the orientation of people. If one wears 2 active badges, one on each shoulder, an UWB system can determine the orientation of the person. One can also do environmental discovery and track people's movements.

SUMMARY

Wireless computing is fun yet hard. A systems approach is useful.

Some current research in machine learning

Sean B Holden, Lecturer in Machine Learning

My research for the last 13 years or so has been in the area of machine learning, and in particular the foundational issues in supervised machine learning.

What is supervised machine learning?

This is probably best explained by presenting a current application of interest to many of us: SPAM email filtering.

It is very easy for us to tell the difference between SPAM and non-SPAM email; however as with most tasks involving the reproduction of a cognitive ability that humans find easy, this particular competency proves very hard to replicate to sufficient accuracy in a software system. One possible approach to the production of filters is essentially rule-based: rules may be hand-crafted with the intention that they match SPAM, but do not match non-SPAM messages. For example a rule might specify that a message containing the words "diet" and "pill" in the same sentence has a high likelihood of being SPAM. This is labour-intensive and easily circumvented by any reasonably sophisticated spammer, and it also lacks the ability to adapt to the needs of specific users. For example, although messages containing "Viagra" will be SPAM to most people, they may well be critical communications to those working in the pharmaceutical industry. In a case such as this more detailed rules might go some way to solving the problem, but in general a more sophisticated approach is required.

A supervised learning-based solution to this problem would work roughly as follows. Users would initially mark unwanted messages as SPAM. This makes available a collection of messages marked as SPAM, and another collection known to be non-SPAM. Messages themselves will typically be processed so that only selected elements of their content are retained - these elements are typically referred to as *attributes*. Assume for the time being that each message of any type is reduced to a fixed number of attributes. We now have a collection of attribute vectors, each labelled according to whether it has the classification SPAM or non-SPAM. The task of a learning algorithm is to infer from this information, a *general rule* for mapping arbitrary attribute vectors to classifications. Thus, when a new email appears it can be transformed to a vector of attributes in the same manner as the earlier messages, and then classified using the general rule inferred by the learning algorithm. An initial rule can be obtained by training on a database of known SPAM and non-SPAM. (That is, messages which would be regarded as such by the majority of users.) However the ability of learning algorithms to take account of user-labelled data allows adaptation both to user preferences and to new SPAM to occur.

It should be mentioned that two further basic forms of machine learning also exist. In unsupervised learning we attempt to discern structure in un-labelled data, and in reinforcement learning we attempt to learn to act in the presence of rewards (and possibly punishments) as opposed to the scenario in which information is explicitly labelled.

Theoretical research in machine learning

Historically speaking, supervised learning has been studied for several decades, and arguably for much longer. Earlier in the twentieth century it would have been called statistical inference, and indeed much of the material developed around that time and based essentially on linear mathematical techniques remains as a critical foundational building block today. During the 1980s in particular, the popularity of *neural networks* was driven somewhat by their presentation as fundamentally new techniques, in some way removed from statistics and driven more by brain research. This is, however, in most cases an inaccurate viewpoint. A much fairer assessment is as follows: the emphasis on linear techniques in earlier statistics had much to do with the absence of sufficient computer power to explore anything else, and with the rise of cheap computer power in more recent years the scope of statistical research to explore more sophisticated non-linear techniques has resulted in major new developments. Neural networks of many kinds can now been seen as statistical techniques for the modelling of data, as can many other supervised learning techniques.

Supervised learning was introduced above using a specific example of its application; in fact it has been applied to a very wide variety of different areas, and in many industrial sectors. In practice it must generally be applied with due attention given to the specifics of the problem at hand. However it remains important to understand the process at a fundamental and application-independent level. Research in the underlying theory of supervised learning addresses questions such as:

- what is the computational complexity of the learning process; and,

- given a specific learning scenario, how many labelled examples might we expect to require in order to obtain a given classification performance in general?

I am particularly interested in questions such as the latter, which in recent years has benefited from deep and beautiful mathematical techniques ranging over fields as diverse as statistical physics, functional analysis, mathematical statistics, optimisation theory and beyond. In particular, this theoretical research has resulted in the development by a diverse group of researchers of what are at present perhaps the two most important techniques in applied supervised learning: boosting algorithms and the support vector machine. In recent years I have, in collaboration with my research student Matthew Trotter (based at University College London) and with GlaxoSmithKline, studied the application of support vector machines within the area of drug design. At present however my primary research interest is in the theoretical analysis of the behaviour of boosting algorithms.

Boosting algorithms allow us to construct a supervised learning system incrementally. Instead of obtaining a complete system in one step of training, a collection of simple systems is produced which are then combined to form a final classifier. Although this basic idea has been widely studied, boosting algorithms are unusual for two important reasons:

- they form the collection of classifiers in a highly novel manner; and,

- they appear to offer state-of-the art performance in practice in a wide variety of situations, *without overfitting*.

Overfitting is a ubiquitous problem in supervised learning whereby, essentially, the system rote-learns the training examples but the final rule produced performs badly in general. There are reasons for which boosting algorithms might be expected to suffer badly from this problem, but what makes them interesting from the point of view of a theorist is that *in general they do not* ! Several potential explanations have been proposed for this, but I feel the real answer remains elusive. My main research effort at present is to attempt to explain precisely why it is that boosting algorithms are effective methods that are resistant to overfitting in the manner observed.

Profile – Sophos Anti-Virus

In the third in the series of articles profiling companies founded by Computer Lab graduates, 'The Ring' was delighted to talk to Dr Jan Hruska, Chief Executive Officer of Sophos Anti-Virus, which he founded in 1987. Dr Hruska is a graduate of Downing College.

TR: Dr Hruska, to start with can you give me a brief description of Sophos in terms of history, product and geographical scope of your activities?

JH: Sophos was formed in 1985 by myself and my current business partner and co-CEO, Peter Lammer. We met at Oxford where we were both working on our doctorates in Medical Engineering. I was working on blood flow measurement and the creation of pulsating flow while Peter was working on post-operative pain control. This was our second attempt at establishing a business.
Two years before that (and to the horror of our supervisors), we formed Executive Computers Ltd, raised about £100,000 and designed a portable computer based on an 8086 processor. The funny thing was that when we started the design, we based it on the CMOS version of a Z80 and CP/M as the operating system. 8086 was not available in CMOS at the time, which we needed to minimise the power consumption. In the middle of the design, Harris Semiconductors announced the CMOS version of the 8086, so we completely turned the circuitry around, basing it on a 16-bit bus and MS-DOS instead. We got the design to a prototype stage (which we still have and which is still functional), complete with the plastic case design, and then failed to sell it to a third party to manufacture. We wound the company down in an orderly fashion, placed our tails between our legs and returned to finish our doctorates. I then spent a year as a conscript in the Yugoslav army (if you would like a feel of what that was like, read "The Mint" by T.E. Lawrence and multiply the discomfort factor by a factor of 10), while Peter finished his doctorate.

We formed Sophos in 1985 with one firm conviction: no more hardware. We started off by developing a number of software modules in C that we thought would be useful in constructing data security products, such as the Data Encryption Standard (DES), RSA encryption etc. This was the foundation of Sophos. When first viruses appeared in 1987 it seemed like a natural fit for Sophos as a data security company to develop an anti-virus product.

We did, and we are now world's 5th largest anti-virus software producer. This year we will probably become number 4. Like most software startups, we did not need a huge investment to start. We first raised about £40,000 in 1986 followed by £60,000 two years later. We have been fully profitable, cash-generative and cash-positive since 1991 and our bookings last financial year were around £48M.

From the beginning Peter and I were conscious that just growing turnover is not sufficient. A company has to grow profits at a similar or better rate. This was an unfashionable view during the .COM boom, but we stuck to our principles. Now profits are back in fashion, of course.

TR: What are the main problems with regard to viruses as far as their impact on companies?

JH: The main impact is the disruption they cause through effective denial of service. Virtually any company today needs its IT infrastructure to function and a virus attack disrupts that.

TR: Viruses are becoming more sophisticated. Which do you see as the most problematic viruses over the next year?

JH: The most problematic viruses are the ones that incorporate some form of mass-mailing capability. Everything is connected these days and if the virus is able to exploit that connectivity, it stands a much better chance of spreading around the globe effectively.

TR: I believe that there were twice as many viruses in 2002 as in 2001. What can be done to stop this trend continuing?

JH: Viruses are like graffiti: the more effort you spend on cleaning them off the walls, the more they appear. Legislation is a step in the right direction and I am sure that fairly severe punishments given to virus authors so far are a deterrent to would-be virus authors. However, I do not see how the virus creation can be stopped in the long term. The internet allows virus authors to hide very effectively and unless they get careless, the chance of getting caught is close to zero.

Some people are calling for the freedom to write viruses as a manifestation of their freedom of speech. The University of Calgary, for example, is currently planning to offer a course in computer security which will, inter alia, teach students how to write viruses. When we (and the rest of the anti-virus industry) pointed out to them that that is not a good idea, the squeals invoking academic freedom to be able to teach whatever they wanted, had to be witnessed to be believed. It seems that they are going ahead with that project, regardless of the protests.

TR: Do you think education in "safe computing" is adequate? Do companies have sufficient computer security policies and follow safe computing practices?

JH: Safe computing relies on users observing rules. As soon as people are in the loop, the system will be fallible. Safe computing is necessary in the same way that safe driving is necessary, but accidents will, unfortunately, continue to happen for a variety of reasons. Other drivers may not be careful, cars may have defects etc.

TR: What are the key factors and challenges to Sophos's ongoing success?

JH: Surviving change is a major challenge. We are growing at some 40% per year, which in my opinion is about the maximum that a company can grow, while retaining stability. That kind of growth necessitates change, but people do not like change. We have to convince them that the change is good and since we mainly employ highly intelligent people, they do listen to and accept convincing arguments.

The other challenge is being able to continue to recruit sufficient numbers of top notch programmers. We are always on the lookout for new talent and like most expanding software houses, we have lots of projects and ideas, but an insufficient number of people to develop them.

TR: Where do you see the company within the next 3 years and what geographical area offers you the greatest potential?

JH: We have to make more inroads into two main IT markets in the world: the US and Japan, which together account for about 60% of the world anti-virus market of US$ 2000M. We have only some 3% of the market in each of these two countries and the potential for growth there is significant. Of course, our competitors for whom these countries are home turf, do not welcome our expansion there. But we do not mind a fight.

TR: Finally, what advice would you give to other graduates thinking of starting their own business?

JH: Think carefully before you do it. The rewards can be great, but the price of failure is also high. And if you decide in the meantime that you don't want to do it just now, join us! We are hiring.

Dr Jan Hruska will be giving a Hall of Fame Case Study in April 2004 – See Events Calendar.

Cambridge Computer Lab Ring Events calendar

If you have any suggestions about events you would like held or speakers you would like to hear, please contact jan.samols@cl.cam.ac.uk

October 21st 2003
Computing in the City: IT agility – how to support new products fast.
17:30 Royal Statistical Society, London

Speaker: Richard Muirhead, founder and CEO of Tideway

November 13th 2003
Careers talk
16:15 William Gates Building, Cambridge

Speakers: Dr Andrew Herbert , Managing Director Microsoft Research

Adam Twiss, CEO of Saviso Group and founder and former CEO of Zeus Technology

January 15th 2004
AGM
16:45 William Gates Building

Library House presents its Cambridge Cluster Report
17:15 William Gates Building

February 2004 (date to be confirmed)
Making the Numbers – Secrets of Software Sales
16:15 William Gates Building

Speaker: Stephen Allott, former president and CFO of Micromuse Inc

March 15th 2004
Lecture on 5 great Challenges in Computer Science followed by Annual Dinner at Churchill College

April 2004 (date to be confirmed)
Hall of Fame Case Study: Sophos anti-virus
16:15 William Gates Building

Speaker: Dr Jan Hruska, founder and Chief Executive Officer of Sophos

May 2004 (date to be confirmed)
Regulation of dominant positions
16:15 William Gates Building

Speaker: EU Commission

June 2004 (date to be confirmed)
London Garden Party

The talks are unticketed but capacity is limited so to reserve a place please email jan.samols@cl.cam.ac.uk or phone 01223 763585. If you would like to be emailed about upcoming events please let the office know.

How not to beat spam

By
Richard Clayton
Computer Security Researcher at the Lab and a trustee of the Foundation for Information Policy Research

(This article was first published in NewScientist on June 28th)

They have done it again. This month, it was finally proved that spammers trying to distribute advertisements for pornographic websites have been using "trojan programs", delivered by viruses to turn personal computers into servers capable of passing on 750,000 items a day. By using viruses, they have upped the ante once more in the fight against unsolicited and unwanted advertising, which now makes up about half of all email traffic. This confirms what has been apparent for some time: we will never beat spammers using technology alone. It is time to move on and tackle the problems in the courts.

Spammers are using ever more sophisticated techniques to get their ads into our inboxes. For years they used throwaway dial-up accounts and "open relay" mail servers that would accept a single incoming email and deliver it to hundreds of recipients. The ISPs have all but stamped this out, so the spammers have taken to hijacking insecure machines, especially those owned by broadband customers. Improving security here will take years. Software and internet experts are touting myriad technical solutions to the problem, but there are flaws in them all.

The technical solutions can be divided into four major categories: blacklisting, authentication, payment and filters. Blacklisting works by blocking email from locations that are being actively exploited by spammers. Unfortunately, the accuracy of blacklists varies considerably, and many ISPs consider it unacceptable to risk blocking legitimate traffic in this way. Nevertheless, blacklists have had some effect. The decision by many American ISPs to block everything from China or South Korea forced both countries to crack down on spammers and fix insecure systems, for fear their economies would suffer.

Authentication, whereby ISPs make you personally accountable for every piece of email sent using your cryptographic identity, and payment, whereby you are forced to pay for any spam you send, both suffer from the same problem. Innocent users who run insecure systems and have their identities stolen by spammers will simply not tolerate permanent disconnection or any significant monetary penalty. Economic solutions work only if you can guarantee that the spammers alone will pay, and they are unlikely to co-operate.

Filtering, the flavour of the month, involves identifying and discarding spam so that only genuine emails reach your inbox. But despite some short-term success, filtering is not a long-term solution to spam. An obvious problem is that it can lead to legitimate email being discarded. In February the UK Parliament's new filtering system blocked email relating to the Sexual Offences Bill and messages written in Welsh. Filters also need considerable personalisation: you or I may not wish to receive email about Viagra, but the people at pfizer.com certainly do. Even if filters are working well, they only prevent delivery to the end user, and this is just one part of the spam problem. The messages still travel through the network and incur bandwidth and processing costs, which someone has to pay for.

But the real show-stopper is that spam is rapidly evolving under Darwinian selective pressure to evade the filters, just as press advertising evolved to look like editorial (or art) and television commercials can now pass as miniature films. The filters will become ever fuzzier, discarding ever more genuine traffic, as mutations breach their defences, and they will be rendered useless when spam is indistinguishable from normal correspondence with friends, family or colleagues.

So if technology is not the answer – and the spammers have won every round so far – then we must turn to regulation. Essentially we have to make spam illegal everywhere. Sending bulk unsolicited email has been unlawful in Europe for years. Individual email addresses are classed as "personal data" (only role-based addresses such as sales@example.co.uk escape this definition), and there are strict limitations on processing personal data for "unfair" purposes – including sending unsolicited email. But anyone who looks into their mailbox will realise that most spam originates in the US, where anti-spamming laws are weak or non-existent.

Given the right powers, the courts could tackle the problem effectively because juries will be able to judge whether mail is unsolicited or unwarranted with more subtlety than any technology. We also need swingeing fines for companies that pay spammers to deliver their advertisements – precisely the economic lever that killed off Europe's pirate radio ships in the 1960s. Finally, anti-spam laws must be enacted planet-wide so that the spammers, who are far easier to trace than many suppose, can be prosecuted wherever they try to hide. The problem has been left for too long in the hands of programmers. It is time for diplomats and politicians to act.

ONLINE DIRECTORY

The online directory is a valuable resource for networking and staying in touch. If you have yet to update your personal profile and employment information please do so. If you have lost your username and password please contact the Ring office.

Who do you know?

Tim King *CC* BA76 PhD80

I took the Computer Science Tripos degree at the Computer Lab in 1976 – in those days it was a one year course and it was only the second year that it had been running. I continued at the Lab doing a Ph.D. – writing a relational database – under the supervision of Ken Moody, and I then stayed on at Cambridge for a further two years doing some post-doctoral work. So unlike many of you more recent graduates who simply took the three year Tripos and then left, I spent six years in Cambridge and met a lot of people who were connected with the Lab at that time. And a great number of those people have gone on to be movers and shakers in the computer industry.

Ever since then, I have used my Cambridge Lab connections extensively – the so-called "Cambridge Mafia" – in all aspects of my career. This networking has worked for me, and it also works for others with whom I network. I was offered my first job as a lecturer at Bath University by Prof. John Fitch, who had moved to Bath from the Lab a year or so earlier: he remembered me and invited me to apply for the post.

Later, when I joined a small start-up, I took with me an operating system based on work done at the Lab under the direction of Martin Richards in the late 1970s called Tripos. This eventually became the operating system for a home computer called the Amiga.

When I came to start my own company in 1986, I looked for a partner and financial backer and immediately found one in the shape of Jack Lang, a well-known name around the Lab and someone who seems to know everyone in Cambridge. Jack's many introductions kept Perihelion alive with new ideas and new students, and as Perihelion joined the Computer Lab Supporters Club (which was very small in those days) I was able to keep up with other people who had either been at the Lab or who had some ongoing connection with it. Perihelion also employed a couple of ex-Lab people, but we were always at a disadvantage by not being in Cambridge or London, but in an obscure town in Somerset.

When I started the ISP UK Online in 1994 I turned for finance to probably the best-known face of Cambridge entrepreneurship, Hermann Hauser. I had met Hermann when he had a team of people wire-wrapping the prototype of the BBC Micro, and had kept in touch ever since. Not only did Hermann back my idea, he recommended my project to Olivetti who provided full finance within literally a couple of weeks.

Since selling UK Online in 1996 I have acted as an independent consultant, offering independent advice to Venture Capitalists on investments and other advice on software projects to organisations ranging from small start-ups to government departments. As such I don't advertise and I rely exclusively on my reputation and networking. I regularly get phone calls or email from people which start "You don't know me but I have this problem and you have been recommended to me…"

The mechanisms of networking are very ill defined, but it does centre on keeping your address book up-to-date. I find it very useful to send Christmas cards to people who might be important one day – for example I regularly exchange cards with the chairman of Misys (the UK's largest software company) because I helped him and Misys when I was a research student (they used Tripos at that point).

I always try and find out if someone else I am talking to in business has been to the Lab, and especially if he or she has done a Ph.D. at the Lab. There's a sort of unspoken camaraderie between us, and a type of recognition of each other's skills, that otherwise has to be delicately explored before relying on the opinion of the other.

I also keep in regular touch with the Lab, having helped suggest the idea of the Ring and even going so far as to lecture twice a year.

Finally, the education offered by the Lab always has been a badge of success. My favourite story concerns a time when I was sent to look at a potential investment in a company that was clearly on its last legs (and in fact with a very poor business proposition). After being given some flannel by the sales director I started to ask some searching questions from the CTO, who didn't know the answers. As I was being led out of the building I overheard the chairman of the company exclaiming to his colleagues "Of course the investigation didn't go well. They only sent a Cambridge F***** PhD!"

VISITS TO SCHOOLS

As part of a programme to encourage applications from the next generation of computer scientists, we are looking for volunteers to visit their local school or old school to talk about Computer Science at Cambridge. If you would like to get involved please contact the Ring office.

THE GAMES INDUSTRY: AT THE CUTTING EDGE OF TECHNOLOGY

Computer Games - a talk to the Cambridge Computer Lab Ring

Meeting Report by Stephen Allott
Director of Development, Cambridge University Computer Lab

Laboratory graduate Demis Hassabis (Queens Matr. 1994) addressed a crowded lecture theatre in the Computer Lab. The atmosphere was buzzing. Over 100 people attended including several faculty members and his former Director of Studies from Queens, Robin Walker. Demis is founder and CEO of Elixir Studios Ltd. He talked about the latest in games technology and his predictions for the next 4 or 5 years.

Whilst an undergraduate, he co-wrote Theme Park which is acknowledged as one of the most successful computer games of all time. He then joined Lionhead before setting up Elixir in 1998. He started by reminding the audience that games programming is hard and therefore games companies employ the best programmers around. The result is that games drive the leading edge of software in many areas, pushing people from a technical and a creative point of view at the same time.

The games industry has been changing over the past 10 years. Games now contain over 1m lines of code. Rigorous software engineering techniques are now required. Competition in the industry constantly drives innovation across many areas which he discussed in turn.

ARTIFICAL INTELLIGENCE

Strategic planning is important in team based games. Artificial team mates can now take orders but still lack the ability to think intelligently. Behavioural cloning is a technique where the game mimics the behaviour of human players. This is useful in multi-player games where someone loses a network connection or logs off for the night. The computer can continue to move that player's character following the style of the human player. Natural language processing is starting to be used although the voice recognition systems need a lot of training. Dialogue generation, which is not pre-programmed, needs to be developed. The current state of the art is still pre-programmed. A related topic is causal effects. Characters need to have different dialogue after something has happened. An example would be the death of a team mate. The character should be affected by the loss, not appear indifferent.

Animation straddles AI, Physics and Graphics. At the moment, every possible movement has to be pre-animated. Demis however foresees a time when programmers only do high level movements and AI animates the detail. Dynamic story creation is another leading edge area. An example would be providing a player with a real choice of corridor with different storylines depending on the path chosen. This is very hard to do. Facial expressions for characters linked to their emotional state should be automatic. The success of the SIMS is partly explained by the fact that they exhibit very human behaviour.

GRAPHICS

Turning to predictions for future development in graphics, in hardware, ATR graphics cards will get more powerful. In the Playstation 3, IBM/Sony cell chips control sub vector units in other machines. Lighting models are getting more important. Emphasis has moved on from the number of polygons to the light effects where Doom 3 is state of the art and was written by John Carmack, the greatest genius in graphics. Demis's new game, Republic, has multiple lighting sources.

The level of detail is the current technology frontier. As you get closer to an object, the level of detail increases. This has been an interesting challenge for Elixir in Republic. The standard technique would have been to switch through a variety of pre-set models as the scale of the model went up or down. In Republic, this is now achieved by a smooth flow without steps. Special effects such as hair, explosions and modelling of liquids are also important. Demis commented that it takes about 4 to 5 years for special effects to get from films to games.

PHYSICS

Rag doll movements used to be modelled with inverse kinematics. Contact points would be modelled from the ground up. Car dynamics are used in car games. Ballistics, fluids and cloth simulations are all used.

NETWORKING

Broadband allows much more complex games with multi-player architectures. Butterfly.net is a new service provided by IBM and Sony to games developers.

TOOLS AND ARCHITECTURE

C ++ is now totally prevalent and object orientation is important. Reusability of middleware is growing and it will get into graphics in 3 or 4 years time. Community tools are also growing.

INTERFACE DESIGN

Games are leading the way in interface design and providing lessons for more general systems. Nintendo games design is the best. Voice input is growing.

REPUBLIC - THE REVOLUTION

Republic is Demis's new game. Set in Eastern Europe in 1991 in a fictional country based on the Ukraine, the aim is to oust the President of the republic. The game will be available in July 2003.

TECHNOLOGY PREDICTIONS

Demis's predictions for games technology are:

1. Multi-player online games will grow
2. A new kind of storytelling will emerge
3. Games will grow and become even more ubiquitous
4. Games will rapidly become mainstream
5. Games will influence popular culture

6 We are at the dawn of an exciting era, the 4th era of games.

Pong was the 1st era, Super NES the 2nd, 3D the 3rd era and the 4th era will be about AI where games have believable characters. It's a good time to start working in the games industry. Demis is looking for people who have

- ¶ A passion for games
- ¶ Exceptional talent
- ¶ Enthusiasm and drive
- ¶ The personality to be a team player.

He recommends that people read Games TM magazine and join a UK developer with an R & D focus such as Elixir, Lionhead, Criterion (who have the best middleware in the world) or Sony. Games will appeal to a wider audience in due course. Elixir is leading the break away from shoot 'em up games. The UK is a world leader along with Japan and the US and a great place to practise computer science.

Iris recognition trials at Heathrow

Heathrow airport has concluded a successful iris recognition trial.

The algorithms for iris recognition were created and patented by the Computer Lab's John Daugman and they form the basis for all current iris recognition systems and products (see "Iris Recognition seeing International Deployments", The Ring September 2002)

Business Hall of Fame

Thank you to all those who contacted us to add a company to the list. They are as follows. The full list can be found on www.camring.ucam.org

Chris Charlton *Q* BA93 PhD99
Calum Grant *Q* BA96
SmartInfoSearch (f.2003)
Develops intelligent information retrieval technology which searches for text based on the meaning of words

Dominic Edmonds *K* BA87
Evertrack Ltd (f.1994)
Business process, technology management consultancy in the risk business sector: banking, insurance and betting and gaming

Andy Hopper *TH* PhD78
Ubisense Ltd (f.2003)
Location systems

John Kleeman *T* BA81
Questionmark (f.1988)
Questionmark software makes it easy for educators and trainers to write, administer and report on assessments, tests, exams and surveys using PCs, LANs, the Internet and intranets

Jack Lang *Em* Dip71
Artimi (f.2003)
Develops silicon solutions for ultra wide band technologies

Stewart Lang *Sid* Dip71 PhD
Micro Focus (f.1976)
Built COBOL compilers, and associated tools (interactive debugger, file managers, screen handlers etc) to run on IBM PCs and Unix. Company taken public in 1983 in the UK, and in 1984 in the US. Industry standard for COBOL

John Mannix *CC* BA94
Governor Technology Limited (f.2001)
ebusiness software consultancy

Richard Mason *Q* BA84
Ionysys Technology Corporation (f.1996)
Provided e-commerce consulting services, both advisory and implementation-related. Based in Vancouver, Canada. Sold to Pivotal Corporation in 2000

Gerald Ratzer *Jn* Dip64
Media Dynamics (f.1968)
Computer Consultancy

Computer Laboratory News

Student Awards

The following students have been nominated for the Science, Engineering & Technology Student of the Year Competition:

Alexandre Mathy (Trinity)
Object Tracking in Video Scenes

Philipp Michel (Churchill)
Support Vector Machines in Automated Emotion Classification

James Murphy (Jesus)
Visual Simulation of Smoke

Dominic Wee (King's)
An Implementation of a-Prolong

Prize winners

Part II Data Connection Prizes were awarded to:

Rick Griffiths (Trinity)
Alexandre Mathy (Trinity)
Philipp Michel (Churchill)
James Murphy (Jesus)
Dominik Wee (King's)
Andrew West (Churchill)

AT&T Prize for Excellence in Computer Science was awarded to:

Dominik Wee (King's)

Part IB Industrial Supporters Prize was awarded to:

Viktor Vafeiadis (Selwyn)

The Data Connection Prizes for achievements in Part IA were awarded to:

Kai Arne Krueger (Churchill)
Richard Low (Trinity)
Jonathan Mark (St John's)

Twenty six candidates were successful in the Diploma, the following with distinction:

Gavin Dupoy (St Edmund's)
William Greenleaf (Trinity)
Ben Nicholson (Queens')
Ka Man Wong (Fitzwilliam)

The AT&T Prize for top Diploma Student was awarded to:

William Greenleaf (Trinity)

Moving towards a 4 year course?

While the current ComSci course has been pretty stable for 10 years, the subject has developed in breadth and depth and a number of external factors have changed which make it appropriate to reconsider arrangements for teaching Computer Science at Cambridge. Increased collaboration with Mathematics, Engineering and other disciplines, the need to find sufficient time for practical work and accrediting authorities expecting a four-year course all contribute to discussions to offer a fourth year. While discussions are in their early stages, they are likely to see further development over the coming year.

The Lab presents the inaugural Diploma ceremony

Dr Ken Moody, Coordinator of the Diploma course, welcomed students, family and friends to the William Gates Building for the diploma ceremony, the first time it has been held. After a few words from the Head of Department, Professor Ian Leslie, Dr Moody was delighted to introduce Professor Sir Maurice Wilkes, the first head of the Lab and the architect of the first proper computer, to present the certificates. After the ceremony, everyone retired to the Street for refreshments.

Surgeries for Entrepreneurs

Do you and your venture need expert advice on particular business matters?

Following the success of the pilot scheme, Cambridge Entrepreneurship Centre (CEC) will be hosting a series of ***Surgeries*** to help University of Cambridge start-ups, entrepreneurs and innovators deal with, and hopefully resolve, issues and concerns that need to be addressed.

Aims of Surgeries

- To provide no-cost face-to-face expert advice
- To run these ***Surgeries*** as part of a concerted support process to all eligible University start-ups, entrepreneurs and innovators
- To regularly provide assistance in the main categories of help that start-ups need.
- To enable Advisers to participate in the exciting Entrepreneurship Centre culture of the University and to give something 'back' to the system
- To provide a frequent forum for all University start-ups at the University's first dedicated facility for tenant and non-tenant ventures
- *Not* an opportunity for Advisers to sell or promote their wares!

Subject areas covered

Accountancy and Tax, Banking, Early stage funding and infrastructure support, General Law, Marketing & strategy, Patents and other forms of IPR and finally Venture Capital Funding.

For more information, please email business.creation@cec.cam.ac.uk or phone 01223 763 751.

Graduates in the News

We would welcome news of any appointments, distinctions gained or honours and awards made to graduates of the Laboratory. Please contact the Cambridge Computer Lab Ring office.

Chris Charlton, *Q* BA93 PhD99, and **Calum Grant**, *Q* BA96, have received a DTI SMART award for their recently founded company SmartInfoSearch. (The SMART scheme, which helps individuals and small and medium-sized businesses to research and develop technologically innovative products and processes is currently being replaced by the Research and Development Grant Programme)

Dr Andrew Herbert, *JN* PhD75, was appointed Managing Director of Microsoft Research in March 2003. He succeeded the founding director, Professor Roger Needham.

Prof Andy Hopper, *TH* PhD78, was awarded The Royal Academy of Engineering Silver Medal on June 4th this year. The Academy's Silver Medals, instigated in 1995, are awarded annually to engineers aged 50 or under who have made outstanding contributions to British engineering. Only four awards may be made each year.