

The Ring

THE JOURNAL OF THE CAMBRIDGE COMPUTER LAB RING

Issue XXVII — May 2011 — £20

Ring news <i>Ring AGM; Letter to the Editor</i>	2	The Cambridge Phenomenon	3
Who's who	4	Jeremy Bennett <i>Supplying industry with computer engineers</i>	6
Hall of Fame news	5	Sean Moran <i>Abduction and flying saucers</i>	8
		Great East London Software <i>It's a family business</i>	10
		BlinkPipe <i>David Gwilt's now-or-never moment</i>	12
		Jamie Shotton <i>Graduate story</i>	14
		Part Ib Group Projects <i>Jack Lawrence-Jones</i>	16

Ring news

Ring AGM 2011

This year's Ring AGM was held on March 7th.
Chairman's Report by **Professor Andy Hopper**.

It gives me great pleasure to welcome you once again to the Cambridge Ring AGM.

In terms of membership, the past financial year has been one of consolidation with membership remaining stable over the period. Indeed, we won't be able to truly measure our progress until 2012 when many members come to the end of their free trial periods and convert to fully paid up status. While converting those with easy access to events in Cambridge and London will likely be more straightforward, the challenge will be to keep those who live and work further afield. We hope that proposed changes to the Ring's Web site will do much to encourage these members to remain engaged and within the fold.

We have continued to provide an interesting and varied events calendar thanks to all the hard work put in by the Ring Council and London Ringlet committee. Alastair Gourlay's London Ringlet bars remain a stalwart of the London calendar, and I would like to thank him for making the events such a success.

Over the past year we have also held roundtable discussion events in both London and Cambridge, and I would like to thank Stephen Allott and Peter Cowley for organising two very well attended and well received events. I know that Peter Cowley has another Cambridge roundtable planned for May, and I hope that a many of you will come along.

The date of the London summer BBQ is always eagerly anticipated, and thanks again go to Lorenzo Wood and LBi for hosting another successful evening in 2010. We are always keen to hear from Ring members who would like to host an event, so please contact the Ring office if you are interested.

The 26th edition of The Ring newsletter came out in January. Its popularity lies in the fact that it allows members to stay connected to this valuable network no matter where they are based. So, if you would like to contribute to future editions, please get in touch.

Finally, I would like to thank two mainstays of the Ring's mentoring scheme — Peter Cowley and Stuart Newstead. Both have worked very hard over the last year mentoring a number of Ring members, and the Ring office receives nothing but expressions of gratitude for their advice and support. If you are interested in becoming a mentor and would like to better understand what it involves, please get in touch with the Ring office.

Letter to the Editor

This June marks the 40th anniversary of the graduation of the first Computer Science BA students in Cambridge's history. A single Part II course was announced late in the academic year 1969–70, and 27 natural scientists, engineers and mathematicians scrambled to register. It was a packed year — some would say that it only took a year to teach everything that was then known about the topic! We romped through numerical analysis, theories of computation, several then-current programming languages (including, of course, BCPL), the fledgling field of graphics, and the features of the Titan operating system. We had a 4K DEC PDP-15 to play with. In June 1971 we graduated and went our separate ways.

Now I have been gathering together as many as I can find, using this thing called the Internet, of course. I'm in touch with 15 of the 27 so far. We were pioneers in a small way, and it has been fascinating to find out what has happened to everyone. We've had a variety of careers, and about half of us have retired. Sadly, I know of one who has passed away. At least four (including me) are living overseas. Some of us are still in harness, and still programming; others have moved to other careers.

I hope to be able to put together some more details for the next edition of the magazine. Meanwhile, think of us this June, and maybe even raise a glass!

David Brooks (CC BA71)

The Cambridge Phenomenon

A richly illustrated book commemorating 50 years of the Cambridge Phenomenon, the technology cluster that has grown up around Cambridge, will be coming out towards the end of 2011. Not surprisingly, given the cluster's nickname of "Silicon Fen", the book will feature numerous stories about the people and companies that have come out of the Computer Lab.

Charles Cotton (Sinclair, Virata, Level 5 Networks/Solarflare) started the project, and has collected a veritable who's who of Cambridge entrepreneurs as his Advisory Board, including Hermann Hauser, Mike Lynch and Warren East. The book will feature a chronological view of the Cambridge Phenomenon since Cambridge Consultants was founded in 1960, and also take a sector approach, exploring how different technology sectors — particularly computing — have boosted the local economy.

Ever since Maurice Wilkes set the computer legacy in train with EDSAC in the 40s and 50s, the Computer Lab has driven a lot of the innovation, and provided the inspiration for many companies, around Cambridge. The presence of an Atlas 2 computer in Cambridge along with the expertise in the department, especially Charles Lang and his computer aided design (CAD) group, led the Labour Government to pick Cambridge as the site for a national CAD facility in the 1960s. CADCentre would eventually become a private company in 1983, list on the Stock Exchange in 1996 with a market capitalisation of £36 million, and change its name to Aveva in 2001. Headquartered in Maddingley and with some 750 employees around the world, 2010 revenue for Aveva was over £148 million.

Shape Data and Topexpress were two companies that came out of the Computer Lab in the 1970s. Founded by Charles and Jack Lang, respectively, Shape Data focused on CAD while Topexpress started out as a software consultancy. Both companies were acquired in the 1980s, while the Lang brothers went on to found or get involved in numerous other companies, including Three-Space Ltd, Electronic Share Information and Artimi.

The 1980s also saw two more Cambridge companies, Sinclair and Acorn, capturing the UK home computer market between them, with the ZX81, the ZX Spectrum and the BBC Micro. These machines, in turn, played their own role in inspiring teenagers and students to get involved in computing — Mike Lynch started learning program-

ming on his college room-mate's BBC Micro and went on to found Autonomy, the UK's largest software company; David Braben's first game, Elite, was written for the BBC Micro.

Sinclair and Acorn had all but disappeared by the end of the 1980s, but this wasn't necessarily a bad thing for computing companies in Cambridge — it led to the foundation of what began life as the Olivetti Research Laboratory and another string of start-ups, starting in 1993 with ATML, which changed business direction and became Virata. When ORL was finally closed in 2002, its disbanded staff turned their on-going projects into yet more companies, like RealVNC, Ubisense and Level 5 Networks/Solarflare.

Today, Cambridge is also home to Frontier Developments and Jagex, responsible for a number of hugely successful games, including Frontier: Elite II, Dog's Life, Runescape and War of Legends.

At the last count, graduates of the Computer Lab had founded 182 companies, with some still being within cycling distance of JJ Thomson Avenue.

Copies of the book can be pre-ordered via the publisher's Web site, <http://www.tmltd.com/shop/home/pld/150>.

If you've founded or work for a technology company in Cambridge, there's still time to tell the editor of the book your story. Kate Kirk can be contacted at kate.kirk@cambridgephenomenon.com

Who's who

Stephen Allott (T MA80) has been appointed Crown Commercial Representative at the Cabinet Office. His task will be to build a more strategic dialogue between HM Government and smaller suppliers.

David Aspinall (SE BA91) is a senior lecturer in the School of Informatics at the University of Edinburgh. David is a member of the Laboratory for Foundations of Computer Science and the Mathematical Reasoning Group.

Jonathan Ayres (R MA92) is Deputy CFO at C Hoare & Co.

Dan Barker (DOW BA) is CEO and creative director at Yearbook Machine, a company he co-founded in 2007. Yearbook Machine makes software to generate books automatically from a proprietary social network.

Judy Booth (CTH Dip84) works for Nominet UK where she is an analyst developer.

Youssef Bouguerra (PEM Dip98) is now a consultant at Rabobank International Brazil SA.

Richard Bradley (R BA05) works for Softwire where he is a software engineer.

Jiny Bradshaw (K BA96 PhD01) is group leader of Audio DSP Firmware at CSR plc.

Amir Chaudhry (DAR PhD) is a product manager at Red Gate Software where his roles are various and have included creating and running a seed accelerator programme (Springboard).

Tomas Cervenka (CHU BA09) works for VisualDNA in London where he is technical lead for Science and Technology.

Graham Cormode (K BA98) is a researcher at AT&T in New Jersey.

Peter Cowley (F MA77) has been appointed as Investment Director of Martlet, a corporate angel initiative formed by Marshall of Cambridge. Martlet will invest in several early-stage companies with high growth potential, as co-investees with other angels and seed funds. Investments will be made primarily in companies within Greater Cambridge and East Anglia, with sums of £25,000 to £100,000 being invested in each selected opportunity.

Robert Durkin (G BA07) is co-founder and CTO at FusePump.

Joe Farish (T MEng09) works for Barclays Capital as a software developer.

Robert Folkes (EM BA82) will continue as VP Commercial at Psymetrix following its acquisition by Alstom Grid in February. Robert said "I would like to thank Jan [Samols], Stephen [Allott] and the Ring — it was through Ring networking that I found this opportunity which enabled me to help build Psymetrix into an industry leader and execute a successful financial exit."

Martin Fulford (PET MA74) is now working as a senior software developer at Realworld Systems.

Rosemary Francis (N BA05 DAR PhD09) is CEO of Ellexus.

Dan Greenfield (TH PhD10) has co-founded Fonleap with fellow Lab grad, Alban Rustemi (ED Dip04 PhD09).

Thomas Haggett (SID BA07) has joined FreeAgent Central as a senior engineer working on a freelance/small business accounting Web application.

Andy Harter (F BA83 CC PhD90) has been awarded the coveted title of Businessman of the Year at the Cambridge News Business Excellence Awards.

Alex Howard (EM BA03) works for Detica where he is a principal consultant.

Laura James (CC MEng00 PhD05) is now at True Knowledge where she is Head of Knowledge.

Phebe Mann (HH BA01) has won the "Tomorrow's Leader" category at the UKRC Woman of Outstanding Achievement awards. She will receive her award at a special ceremony at the Royal Academy of Engineering. Phebe is a senior lecturer in Highway and Transportation Engineering at University of East London.

Russell Moore (K MA00) is a contact team leader at Great East London Software, a company he co-founded with his brother Philip.

The Ring, Issue XXVII, May 2011, price £20/€30/US\$40

Web: www.camring.ucam.org

E-mail: jan.samols@cl.cam.ac.uk

Tel: +44 1223 763585

Post: William Gates Building, Cambridge CB3 0FD

Published three times a year. Copy deadline for the September 2011 issue is August 5th 2011. All content is copyright ©The Cambridge Computer Lab Ring 2011 unless otherwise noted.

The Ring is the journal of the Computer Lab Ring, which is the graduate association of the University of Cambridge Computer Laboratory.

Governing council: Prof. Andy Hopper (TH78) (Chair); Stephen Allott (T80); John Brimacombe (T91); David Colver (CHR80); Peter Cowley (F77); Robert Folkes (EM82); Nigel Horne (CHU68); Lorenzo Wood (CHR93)

London Ringlet: Alastair Gourlay (SE02)

Hall of Fame news

Amir Nathoo (JN MEng02) has relocated to San Francisco. His company, WebMynd.com, has raised new funding from Founders Fund, Paul Graham, 500 Startups and other US-based angels. WebMynd.com makes software which all app developers need to integrate their apps into users' workflows. Developers can sign up for the beta at webmynd.com.

David Nissenbaum (JN BA09) is a technology analyst at Barclays Capital.

Valera de Paiva (LC PhD90) lives in California where she works on logical approaches to computation, especially using category theory.

Kim Powell (F BA06) has founded IT Business Solutions Ltd, providing business analysis for investment banks and financial institutions.

Andy Robinson (K BA76) played in a special production of Macbeth at the Globe Theatre. Andy can be found performing in all kind of gigs, recording sessions and theatre work in London's West End.

David Simner (JE BA07) works for Red Gate Software as a software engineer.

Pete Smith (Q BA05) works at Songkick, which he co-founded in 2007. Songkick works by indexing 137 different ticket vendors, venue Web sites and local newspapers to create the most comprehensive database of up-coming concerts happening around the world.

Thuckakorn Vachiramon (JN BA06) has co-founded boxmap.

Krzysztof Wos (JN BA10) is an analyst at Goldman Sachs in London.

Acunu

Acunu, which provides software to address the new challenges of storing, processing and serving very large amounts of data at low cost, has secured £2.2m Series A financing to help bring its first product to market.

Acunu Storage Platform is in private beta with companies in areas including social media, advertising and cloud computing.

Acunu was one of 37 leading UK digital companies that were showcased as part of this year's UK Trade & Investment Mission to SXSW Interactive in Austin, Texas, in March 2011.

Jagex

Jagex, a leading independent games developer and publisher, has accepted a new investment from Insight Venture Partners, Spectrum Equity Investors and The Raine Group. The investment will enable Jagex to accelerate growth of its development and publishing capabilities in support of continued enhancement of the studio's portfolio of award-winning games.

Masabi

Masabi, the developer of mobile ticketing technology for the transport sector, has announced that the mobile rail ticket purchase application it developed for thetrainline.com now supports most high-end smartphones and everyday handsets.

Mobile users are able to search train times and purchase tickets using a credit or debit card, and then pick up their tickets at the station. The application will display tickets as secure barcodes on the phone's screen when rail operators start to support this in the coming months.

RealVNC

Hot on the heels of its success as Company of the Year at the Cambridge Ring Hall of Fame Awards 2011, RealVNC has been named Private Company of the Year in The Business Weekly East of England Business Awards.

TouchType

TouchType's SwiftKey has won best app at the Mobile Premier Awards.

The Swiftkey consumer product is one of the leading global Android apps, and has achieved over 600,000 downloads since its beta launch in July 2010. The paid version of SwiftKey remained one of the top five global apps on the Android Market for three months since launch in September and has achieved over 115,000 paid downloads.

Zeus

Zeus Technology, the only purely software-based application traffic management company, is helping BT, one of the world's leading providers of communications solutions and services, to deliver seamless on-line services across its premier Web sites.

Zeus now plays a key part in BT's on-line self-service strategy, acting as a traffic management layer, routing and managing customers' access to rich self-services, as well as personalising their end services through its automated on-line sales channel.

Supplying industry with engineers: where do we start?

Jeremy Bennett, Chief Executive of Embecosm Limited, has an answer.

Open up a smartphone and you will typically find it full of British technology. Almost certainly the main processor is from ARM plc in Cambridge. The Bluetooth chip is probably from CSR plc, also in Cambridge. There is a good chance the graphics subsystem is from Imagination Technologies in Kings Langley. Even where the chip is from overseas, it was very likely designed here.

Embedded engineering — that is computers and software that are embedded within devices such as mobile phones — is a British success story going back decades. As Simon Knowles, founder and CTO of Icera put it at last year's Future World Symposium (www.nmi.org.uk/events/nmi-international-conference-2010), “we based our engineering team in Bristol, not because we come from the UK, but because they are the only engineers in the world that can design our type of chip”.

Britain is not just good at this, it is the best in the world. If it can be done anywhere, it can be done here.

When I wrote an earlier version of this article, it was for teachers of children up to the age of 13. So why is this relevant to those teaching children at a young age?

The current leaders of this British expertise, the CTOs of these world-leading companies, grew up with the BBC Microcomputer. In the early 1980s, a brash young IT Minister, Kenneth Baker, a ground-breaking TV series and a high-performing and modestly-priced personal computer made technology fashionable both politically and socially. In 1983 there were more computers in the UK than in the USA, and teenage programmers gained national celebrity (and substantial fortunes) for their skills.

That wave of enthusiasm didn't come from nowhere. The BBC microcomputer drew on the skills of the Cambridge University Computer Laboratory, and their ground breaking research over the previous three decades. It led to a generation of children who chose to study computer science and electronic engineering at University, who became today's global industry leaders.

But recently we have run into a problem. We've stopped producing graduate engineers in electronics, software engineering and computer science. Companies like ST Microelectronics, with its huge UK research base, have been warning of the problem for some years. It seems all the more surprising, given that according to the government, software engineers can expect to make the most from gaining a degree of any subject — nearly a quarter of a million pounds over a career on average.

Up until now, it has been possible to replace lost team members with experts from overseas, who are only too happy to come to the UK to work and gain expertise. But the new government's decision to freeze immigration last year stopped that, and the problem was thrown into sharp relief. In the absence of any British candidates for jobs, major companies were forced to move their teams overseas to keep working: those jobs were then lost forever. The new regulations have somewhat eased that problem, but the underlying issue remains. There are not enough UK graduates to hire.

Talking to industry and universities, it becomes clear that the problem is not primarily in higher education. True, there are problems with some newer universities turning out IT graduates with skills that are of no use to the engineering community. But the problem is the lack of supply into university courses in the first place — even Cambridge University Computer Laboratory has trouble finding suitable applications.

The problem is in schools. All too often the pressure to deliver exam results means students are steered into soft “IT” courses, rather than the more rigorous “computer science” GCSE. Most teachers think these are the same subject — just one is a bit more difficult than the other. But IT GCSE grew out of the old secretarial skills courses — which is why it is primarily about learning word processor and spreadsheets — while Computer Science GCSE is designed as a science course.

In fact, for university courses, GCSE or ‘A’-level computing is far from necessary. But where schools have good students with maths and

science qualifications, they are rarely inspired to consider computer or electronic engineering as a degree subject, having been brought up to believe it is all about entering data into word processors.

Underlying this is the lack of trained engineers teaching in schools. A recent survey suggested that of 18,000 schools nationally, there were only 1,800 teachers who considered themselves qualified to teach IT. I have to confess that I had never met a school teacher with a degree in computing or electrical engineering until this January.

The result is poorly-taught courses, which inspire no one. The problem arises long before GCSE, which is why teachers of under-13's are so important. I see it personally — my son frequently comes back with homework of a screenshot of Microsoft PowerPoint or similar, asking him to label what the various buttons do. Yet this is a child who, like many of his friends has been writing programs in scratch (scratch.mit.edu) for several years. Thirty years ago, students his age would have been writing programs in BBC Basic.

It doesn't have to be so. The solution may not lie within formal lessons, but in inspiring after-school clubs. In many schools near Cambridge, ARM engineers give up part of an afternoon each week to take MBED systems (mbed.org) into schools, to allow children to learn, hands-on, about modern electronics. Ten minutes to show children how to plug in the device to program it (using a standard PC), and then they are left to learn from themselves what they can do with the various displays, sensors and actuators, programming in C++. ST Microelectronics have their own equivalent Discovery Board (www.st.com/internet/evalboard/product/250636.jsp). None of this is expensive kit — the Discovery Board is around £10, while the MBED board is around £40. A more packaged (and expensive) alternative is the Lego Mindstorms (mindstorms.lego.com/enus/Default.aspx) robotics technology — something that all children take to very quickly.

These approaches show that it is possible to inspire young children to take an interest in engineering. It also shows that industry needs to help create the next generation. We cannot just blame teachers who do not have the skills necessary — industry can provide a way to supplement those skills. ARM shows how it can be done, and being a large corporation does so in a very planned and organized way. But this approach will work just as well with an informal relationship with a local computer or engineering company.

So how can we professional computer engineers help? The approach used by ARM in supporting after-school clubs is highly effective, but we need more of them and that needs more individuals and companies with the right skills to take part. We need to push this strategically at a national level, and also participate as individuals.

As well as running my own embedded software company, I spend two days each month working as the embedded systems champion for the Electronics, Sensors and Photonics Knowledge Transfer Network (ESP KTN). This is how I try to influence strategy at a national level.

The KTNs are funded by central government through the Technology Strategy Board, and are charged with improving the flow of knowledge between the academic and industrial worlds. We do not have a direct responsibility for education, but we are charged with using our resources to boost the competitiveness of UK industry. So that is why I encourage industry to reach out to schools to help boost the future supply of engineers.

The ESP KTN is free to join (ktn.innovateuk.org/web/espktn) and currently has around 7,000 members. We run a range of events, including seminars, industry brokerage and training courses, as well as running a Web site for the entire community. In the last few months we have worked with industry to put on a series of training courses on the MBED system, and in future plan to offer training courses on alternative systems, such as those from ST Microelectronics.

These low-cost courses have a broad remit of helping to train up current engineers to learn more about embedded systems. But they are also highly suited to technology leaders within the school system, who we encourage to attend by waiving the attendance fee.

The other way to reach out to schools as an individual is to become a STEM Ambassador. The STEM scheme is for professional scientists, technologists, engineers and mathematicians who wish to make their skills available to schools. This can be as simple as attending career days, or can be more substantial, providing teachers with specialist input for lessons, or helping to set up after-school clubs. The STEM scheme provides some basic training, a network to link schools up with ambassadors and insurance for individuals taking part. You can sign up via the STEMNet Web site (www.stemnet.org.uk).

Industry has identified the problem, and also shown one way it can be solved. In an ideal world we would have a BBC Micro for the 21st century. But in the meantime, opening children's eyes to the potential, by using technology such as MBED, and drawing on the skills of local engineers provides a way to inspire a future generation of British technologists.

An earlier version of this article was published in the ICT Broadsheet of the Society of Assistants Teaching in Preparatory Schools in January 2011. It was is licensed under the Creative Commons Attribution 2.0 UK: England & Wales License. To view a copy of this license, visit creativecommons.org/licenses/by/2.0/uk/

Abduction and flying saucers

Sean Moran asks whether you are taking seriously enough the steps necessary to protect the potential intellectual property that is taking life in the software you have under development.

What would the terms “abduction” and “flying saucer” immediately bring to mind? I doubt it would be the on-line search facilities at the European Patent Office. But perhaps that isn’t uppermost in your thoughts because you may not be taking seriously enough the steps necessary to protect the potential intellectual property that is taking life in the software you have under development. In two recent articles for *The Ring*, Nicholas Fox gave excellent and expert advice on what can and cannot be patented, and the impact on your business strategy that securing valuable assets in the form of rights of invention might have. But perhaps you should also be aware that others are already well ahead of you in the game. Before you sit down at the table with your pile of poker chips in front of you, remember there are some serious people with deep pockets who not only know the rules better than you but who are making the rules as the game proceeds.

Infringing a patent is trespassing on someone else’s idea. A patent is a private enforceable monopoly on an idea, granted in exchange for a disclosure to the public of how the invention works, and patent applications are held on public databases. An EPO database search on the term “flying saucer” will reveal that the former British Railways Board applied for a patent on a starship some time ago. Besides lacking essential component technology, such as room-temperature superconductors and a laser-ignited fusion reactor, the patent was granted, but has not been renewed and has lapsed. So do feel free to make your own contribution to this area of technology. Take care however if you are writing software that involves storage and retrieval. To the EPO, “abduction” has nothing to do with kidnapping, but is the kind of logical inference that generates a plausible hypothesis to explain a particular circumstance — in other words, “making a guess”.

Searching on the term “abduction” will lead you to a patent concerning the diagnosis of faults in printer or photocopier — data from test signals is used to have a guess at what the fault might be. If that isn’t your area of interest you might dismiss the patent as irrelevant, but you may be misjudging the breadth of coverage a patent can have. Does this patent cover the idea of abductive reasoning when applied to systems in general? From the patent wording it is very difficult to know.

Consider another patent, this time from the US Patent and Trade Marks Office search system, that describes a “system and method for data management according to the content of the data.” This continues with the description in the peculiar language of patent drafting: “The present invention enables data to be stored in one of a plurality of different storage options according to the at least one characteristic of the data, in which at least one characteristic is related to the content of the data. The present invention comprises a rule-based storage management mechanism for the processes of archiving and/or retrieving data. It should be noted that at least one storage option according to the present invention is optionally deletion and/or destruction of the data, such that the data may optionally be removed from storage media or may optionally not be stored initially on the storage media. Optionally and more preferably, the data is stored for a time interval according to the at least one characteristic of the data. Most preferably, the data is moved to a different type of storage option after an event occurs, for example the time interval has elapsed.”

So we may begin to discover that the software we write is unintentionally making use of constructs over which others have legal claim as private property.

We may discover that the software we write is unintentionally making use of constructs over which others have legal claim

Now I have a personal rule-based storage strategy — for example, I keep last year’s tax return in my desk, but my collection of Byte Magazine from the 1980s is in the garage. I even throw some things away after a time. Do I infringe this patent? That could be a matter for the lawyers and the courts to decide. In fact each decision of the courts changes the law. The current legislation in the UK placed computer software in the category of “excluded matter”. Nicholas Fox acknowledged this in his first article, but lumped computer programs with “business methods” and “ways of playing a game”. Had he revealed the other categories of excluded subject matter he could have restored the

gravity of the exclusion of computer programs — laws of nature are excluded from patent, as are mathematical methods, mental processes and aesthetic creations.

Now while computer programs clearly have little to do with natural laws, to many people they are certainly aesthetic creations, and provably they can be identified with mathematical algorithms.

Without going into the Theory of Computability and Universal Turing Machines and other arcana, it is evident that there is great difficulty in seeing any dividing line between the abstraction that lies behind a computer program and the abstraction that constitutes a mathematical construction. Both capture behaviour. In principle, as excluded subject matter, “computer programs as such” should never even be considered for the further tests of inventiveness and industrial application essential for consideration for patent protection. In practice the two words “as such” have provided a handhold for the lawyers, and after the recent Symbian case to which Fox made reference, it seems that having a linking loader that eliminates a single branch instruction is sufficient “technical effect” to allow the courts to side-step the legislative prohibition. So we may begin to discover that the software we write is unintentionally making use of constructs over which others have legal claim as private property.

The problem is compounded because the courts can decide after the fact just what scope a patent claim may cover. The Unisys patent on a disk controller that did lossless compression of data using the Lempel-Ziv-Welch algorithm was enforced by Unisys against anyone of substance who made use of the GIF image file format, a popular method for Web graphics and also in medical imaging, that used the same algorithm. A proprietary technique embodied in a piece of hardware can suddenly spread its influence into software through the doctrine of equivalence.

All this raises a number of questions for anyone undertaking software development. You have already received sensible advice about taking steps to protect your innovation. But do you use the international databases of Patents to see where you might be treading on other peoples’ property? Or to collect ideas — the “teaching” that patent disclosure is supposed to deliver to the world at large in return for monopoly protection — perhaps to develop alternatives that step around the scope of the patent, or to identify potential partners or licensing deals where that isn’t possible? Or do you make use of the Open Software Community in the hope that by doing so it may reduce the risk of ruinous litigation?

Has your company developed a policy for risk management in this area, or have you experience of patent disputes that you might be willing (and at liberty) to share? Is the risk of a lawsuit against you insurable? Would you fight, or fold? The patent system is alleged to encourage innovation. Is that your experience or do you find the game intimidating?

In the nineteenth century Charles Babbage successfully defended a friend against a patent action by finding the one surviving example of an obsolete lace-making machine that cut the ground away from the lawsuit by proving the “invention” was not novel at all, but this gets difficult to do with the abstract ideas embodied in pieces of software.

Do you have powerful friends if you wander into a patent minefield?

Sean Moran (T67) has an interest in property rights in computer software from a development career spanning four decades with companies such as CADCentre, Leica Microsystems, Cambridge Control and several others. He has an ambivalent position in respect of patents — being a named inventor on a couple of patents but concerned at the encroachment of private ownership on the cultural common pool of software techniques and the history of the science. He would very much like to hear your views pro, ante or indifferent — please e-mail t.s.moran@open.ac.uk

Great East London Software

It's a family business.

Starting a company with your twin could put a strain on your personal relationship. However **Philip Moore** (JN MA00), with his twin brother **Russell** (K MA00), looks to achieve boom without the bust-up!

TR: Philip, tell me about Great East London. What does it do and how did you come up with your business idea?

PM: Great East London is a bespoke software company serving the financial sector, including investment banks and funds. We focus on the creative phases of the software life cycle, combining rigorous business and technical analysis with industrial-strength implementation.

We maintain a community of the brightest and most talented people in the City, and this allows us to put together teams that can deliver very high quality results across a range of business lines and technologies.

The idea really came from meeting and working with remarkable individuals at a range of organisations. We met a lot of exceptionally smart people, but they were mostly contractors, mostly working alone. We wanted to bring them together, to create a place where they could work on great projects, develop their interests, and meet similar people. Great East London was the outcome — and as a result our teams can deliver projects that would be beyond the reach of other firms.

TR: What were you and Russell doing before this?

PM: We were both IT contractors, working freelance for a number of different investment banks. This taught us a great deal about the industry and allowed us to make a lot of contacts, but it also showed us the limitations of working alone, without being part of a bigger team — especially one which actively supports your growth and success.

TR: What are the biggest differences between working with relatives and working with strangers?

PM: I think the main thing is the common ground you already share. When I am speaking with my brother, I can often get a message across in half as many words as it would usually take, and I never have to repeat myself. This works because we have so much shared context.

Trust is also important. If you have a family member who you can trust implicitly, then that is a powerful thing. When you know people's motives are pure, you can expend less energy worrying, and get on with creating great things.

Then there are some extra responsibilities that come from working with relatives: after all, you have a duty not to ruin Christmas! So you have to keep things in perspective, and respect the relationship.

TR: Did you first operate your business from home? If so, what were the challenges and benefits from this strategy?

PM: Yes. Like a lot of companies, we found this was a good way to reduce our initial costs. In the early days it can also be useful to have your work available at all times, but this is something you have to shift away from as you grow.

One difficulty with working from home is where to meet clients and colleagues. I developed an encyclopaedic knowledge of London's best coffee shops, which still comes in handy today.

TR: How did you and Russell decide to delegate responsibilities?

PM: Although we are twins, we have quite different personalities, and so the responsibilities mostly fell quite naturally to one or the other. From the outset, I mainly worked on finance and operations, with Russ focusing on the branding and communication. I am not sure I got the better deal, but the result has been effective.

TR: What mistakes did you make early on?

PM: I think if I had my time again, I would get more organised early on. At the time it seems like overkill to think about processes and admin, because times are exciting and you need quick decisions. Later though, when you have a growing firm, you need these things working smoothly, and by then you have even less time to do anything about it! So maybe I could have saved myself some all-nighters if I had known this earlier.

TR: What key characteristics of family companies give them a competitive advantage?

PM: I think that family companies often have a ready-made culture, or ethos, which can go a long way to preventing board-room divisions, especially in the start-up phase, when the company is young and undergoing rapid change. That is not to say there will not be vigorous discussions — but I think the will to find a united way forward is stronger.

TR: What hurdles or pitfalls can de-rail family firms?

PM: Well, I hope never to find out first hand. Clearly there are boundaries which need to be considered when working with your relatives: everyone argues and falls out at times, but you need to make sure these disagreements stay under control and do not affect the family. It is harder to walk away if things go wrong, and that is a risk you need to understand from the outset. You have to be sure that your family friendships will survive come what may.

TR: What challenges do family-owned companies face in attracting and keeping professional managers?

PM: Provided the company is professionally run there should not be any problems. In the end, it does not matter if the directors are colleagues, friends or relatives: their role is to guide the company and influence the culture so that talented people choose to stay with them.

TR: Are there management practices, commonly seen elsewhere, that are given less priority in family firms?

PM: With the exception of political infighting, I would hope not.

TR: Has your experience in running the business been different from what you expected?

PM: Well, I think this is true for every entrepreneur out there. It has been extremely hard work, but also vastly more satisfying than I could have imagined. We have achieved some amazing things for our clients, and created some brilliant opportunities for our associates — the road has not always been smooth, but it has always been exciting.

TR: Do you manage to have business-free family gatherings?

PM: Yes, we do. It's very important to be able to unwind and relax together. There are times when even the most understanding family members do not want to discuss BPM and Enterprise Java!

*Find Great East London Software at
<http://www.greateastlondon.com>*

BlinkPipe

When Alex Nancekievill was looking for a partner for his start-up, **David Gwilt** realised the now-or-never moment had come.

TR: What is BlinkPipe doing?

DG: We're developing meeting-room video conferencing that's as natural as a handshake, as reliable as the telephone and as easy to install as a toaster.

TR: How did the idea for your business come about?

AN: I was CIO at ARM here in Cambridge, and our IT group was getting a lot of demand for video conferencing from all corners of the company. It seemed lots of people could see the potential advantages for improving communication and reducing travel. At the same time it was clear that the communication landscape was a bit of a jumbled mess of technologies that aren't particularly intuitive and don't play that well together. The result left users properly confused. Some meetings really only wanted to share slides; others wanted to be able to see the person at the far end; yet others wanted to be able to get round a virtual whiteboard. Some times it would be two sites communicating, sometimes three or four. Often a few meeting participants would be working from home, or dialling in from an airport lounge. Whatever their other requirements, no one wanted to have to prepare the technology in advance of the meeting, and no one wanted to attend training courses to learn how to use it.

No one wanted to have to prepare the technology in advance of the meeting, and no one wanted to attend training courses to learn how to use it.

We looked at all the technologies we had on the shelf, and surveyed as many alternatives as we could. The result was pretty clear. The best hope for making communication intuitive and straightforward was to implement a single vendor solution across our phone system, video conferencing, slide sharing and e-whiteboards. Unfortunately that would involve replacing a lot of perfectly good equipment. We'd also

have to make a significant investment in our network infrastructure. Even then, video calling to other companies would require some preparation and wouldn't always be possible. We were looking at a massive project for IT and a major capital spend, and it was still only going to get us part of the way there.

It was about this time that I first heard about Y Combinator and I read this from Paul Graham's blog (<http://ycombinator.com/ideas.html>)

"So if you're working for a big company and you want to strike out on your own, here's a recipe for an idea. Start this sentence: "We'd pay a lot if someone would just build a ..."
Whatever you say next is probably a good product idea."

Bingo! If someone could take all the hassle out of remote meetings we would pay good money for that. What was needed was a system that didn't require any preparation for the user and didn't make you decide in advance which technology you were going to use. No more meeting rooms with three or four different communication devices in them. What our staff wanted was to walk into the meeting room and start the meeting. No technical hassles getting started, no interrupts half way through because someone else in the building is watching iPlayer, and absolutely no training courses for the equipment. However this was achieved, it needed to be done within the existing telephone and network infrastructure because these are major capital costs. If it could be made trivial to install that would be even better (one less headache for IT!).

TR: What prompted you to start your own business after many years of comparative safety at ARM?

DG: I was sharing an office with Alex when he handed his notice in, and was intrigued to find out more. Alex was looking for a partner, and having looked into the technical and commercial opportunity I decided it was worth a go. I had the right background, and had worked with Alex before. I've always had an entrepreneurial leaning, from doing Young Enterprise at school to being a pub owner for a few years

in the Noughties. For me it was a now-or-never moment, having been with ARM since college. I don't see any real downsides, so we're just going to give it our best shot.

TR: Where do you see the big drive for using video coming from?

DG: For me, I don't see any big demand-side events that say "now we have to use video". Sure, there is the increasing cost of fuel and a global drive towards more environmental sustainability in business, but that is a slow burn in my view. It is the supply side of the equation that I feel is missing a trick: classic "crossing the chasm" stuff. If the market provides a product that's as easy to use as a telephone, and at a price that means it's a local manager decision rather than a board decision, then market can take off. I believe we have that product.

TR: Have you been able to find funding?

DG: We were introduced to Simon [Galbraith] and Neil [Davidson] of Red Gate Software by a member of the ARM board. Red Gate had experience of corporate venturing through its Springboard programme. Once Alex and I had met with both Simon and Neil, things moved very quickly: they liked our technology and team, and could see that we were serious — both having handed in our notices by this time. Taylor Vinters had worked with Red Gate before, and were able to provide a fast and affordable service to get us through the legals. Much sooner than we'd expected we had the funds we needed. We count ourselves very fortunate to have such great investors. Simon and Neil have been there and done it — they really get what's it like to be starting a business from scratch.

TR: Are you able to share pricing with us at this stage?

AN: It's a little early to be concrete with pricing. However, I can say that in the meeting-room video market, we are expecting to be significantly more accessible than the likes of today's incumbents, Polycom and Cisco.

TR: When will the product be available?

DG: We both have enough experience of product development to know that our first attempt will not be perfect, so rather than shut ourselves away and come out with a big "tah-dah...oops", we are planning to run limited field trials with a small number of prototypes in the Cambridge area in the second half of 2011. We want to learn fast, and iterate quickly towards the go-to-market product. On that note, if any of your readers would like to join our trial — helping a local company to bring meeting-room video to the masses — please contact either Alex or myself through our Web site at www.blinkpipe.com

Hall of Fame Winners 2011

Company of the Year

RealVNC

Founded in 2002 to develop remote access software, RealVNC now has 100,000 customers globally and is licensed on 20m desktops. It has 65 staff and sales of £5m as well as making over 20% net profit margin. It has grown organically without external investment.

RealVNC has recently signed an OEM deal with Intel to embed VNC technology into Intel's latest chipsets. It has some 50 OEMs in total and has become a truly global brand. VNC for Apple and Android devices is popular.

Product of the Year

Cambridge Broadband's VectaStar Gigabit Point-to-Multipoint Microwave Backhaul Platform

75% of mobile network traffic is now data. Cambridge Broadband's platform is now being widely deployed in mobile carrier networks.

Publication of the Year

Robert Watson and Jonathan Anderson for "Capsicum: Practical Capabilities for UNIX"

This research in operating system and application security was performed in collaboration with Ben Laurie and Kris Kennaway at Google. The paper also won the best student paper award at the 2010 USENIX Security Symposium.

Jamie Shotton

Jamie Shotton graduated from the Computer Lab in 2002 and went on to do a PhD at the Department of Engineering from 2003-2007. Jamie now works for Microsoft at its Cambridge research laboratory, where he has been intimately involved in the development of Kinect for Xbox 360 — of which eight million sold in the first two months alone. Jamie tells his behind-the-scenes story about Kinect.

I joined the Machine Learning & Perception group at Microsoft Research Cambridge (MSRC) in June 2008 as a post-doc to continue my research in computer vision. Little did I know how quickly I would get pulled into the frenzy of development around Kinect, or how practical my blue-skies PhD research would turn out to be.

As my PhD topic, I had focused on automatic visual object recognition: teaching computers how to recognise different categories of object (such as cars, sheep, or trees) in photographs. I had taken a machine learning approach. First, you build up a varied training set of images where you label each pixel by hand with a colour, according to which object category it belongs to. So, for example, you label all ‘cow’ pixels in blue, and all ‘tree’ pixels in green, using a simple painting application. Second, you give this training data to a machine learning algorithm that does some number crunching to mine patterns of image appearance that correlate with the presence or absence of the various object categories. The training procedure results in a model that efficiently encodes these correlations, and hopefully generalises to new unseen data. Finally, you then show the model a new image that it has never seen before, and it works out to which object category each pixel belongs.

A couple of months into my post-doc at MSRC, I got a call out of the blue from the Xbox product incubation group who, having come across some of my earlier publications, wanted to discuss an “important, top-secret scenario” with me. They described their goal — real-time robust human body tracking — and how it could be used for playing computer games. Now, this had been a dream of science fiction for many years, and still is a hugely active topic in the computer vision community. Indeed, several of my fellow PhD students, including Dr Bjorn Stenger, had had this as their PhD topic. But it was always seen as being “five years away” from being commercially viable, certainly at a consumer price point. So of course I was rather sceptical anything could come of this, especially given Xbox’s ambitious plan to launch in late 2010.

But then they mentioned the new depth-sensing camera hardware they were busy developing. I had seen depth cameras before but only at very low resolution (about 10x10 pixels). The new Kinect camera worked at 320x240 pixels and 30 frames per second, and the depth accuracy really got me excited — you could even make out the nose and eyes on your face. Having depth information really helps for human pose estimation, as it removes a few big problems. You no longer have to worry about what is in the background since it is just further away. The colour and texture of clothing, skin and hair are no longer an issue. The size of the person is known, as the depth camera is calibrated in metres. Further, since the depth camera is active — shining out its own structured dot pattern of infra-red light into the room — the camera can work with the lights turned off.

But even with depth cameras, it’s not all plain sailing. There is still the whole gamut of human body shapes and sizes and, worse, people can get themselves into an incredible variety of poses (body positions). Just think about how many positions you can put your right arm in, then multiply that by the number of positions for your left arm, your right leg, and so on, and you rapidly end up in the hundreds of billions.

The Xbox group also came to us with a prototype human tracking algorithm they had developed. It worked by assuming it knew where you were and how fast you were moving at time t . It would then estimate where you were likely to be at time $t + 1$, and then refine this prediction by repeatedly making small adjustments while comparing a computer graphics model of the human body at the prediction to the actual camera depth image. The results of this system were incredibly impressive: it could smoothly track your movements in real time, but it had three significant limitations. First, you had to stand in a particular “T” pose so it could lock on to you initially. Second, if you moved too unpredictably, it would lose track, and as soon as that happened all bets were off until you returned to the T pose. In practice this might typically happen every five or ten seconds. Third, it only worked well if you had a similar body size and shape as the programmer who had originally designed it. These limitations were all show stoppers for a possible product.

And so our brief back at MSRC was to overcome these problems somehow. I sat down with colleagues Dr Andrew Fitzgibbon and Prof. Andrew Blake and we brainstormed about how we might solve the problem. A first observation was that when you look at a photo of a person, you can tell where their limbs are even though the person is not moving. If we could remove the temporal dependency, we would remove the need for the initial T pose, and be able to recover if we lost track. Another thought was that to cope with the variations in human size and shape we should use machine learning, rather than try to somehow directly program for all possibilities by hand: instead, we would encode these possibilities in the training data.

During my studies, I had interacted with Dr Stenger whose research uses a technique called “chamfer matching” to match a whole image of the body against the training set of body images. By finding the closest match (the “nearest neighbour”) you can then transfer the known 3D human pose from the training image to the test image. We tried this technique out, and had some success getting a coarse human pose out without using any temporal information. The problem was, however, that to get the level of detail we needed would have required so many ‘exemplar’ training images to cover all possible body shapes and sizes that the matching process could not run in real time on the limited processing hardware we had available.

So we went back to the whiteboard. What was now clear was that we had to divide up the body into parts and somehow match each part independently to avoid the combinatorial problems with matching a whole pose at once. I hit on the idea of revisiting my PhD work on object recognition, but this time instead of object categories, we were going to use body parts such as left hand or right ankle. We designed a pattern of 31 different body parts and then trained an efficient decision-tree classifier to predict the probability that a given pixel belongs to each part of the body. If you can accurately predict these part probabilities from a single depth image, regardless of body shape, size, or pose, then you get 3D proposals for the locations of many body joints at extremely low computational cost.

This turned out to be the winning formula, but it still needed a lot of engineering to scale up to the level of accuracy we needed. The larger and more varied we could make the training set, the better it was likely to perform in your living room. So we turned to Hollywood, who have for many years been building advanced computer graphics models of the human body for their movies. We recorded hours of footage at a traditional motion capture studio of several actors doing various moves that could be useful for gaming: dancing, running, fighting, driving, etc. This “mo-cap” data was then retargeted to different human shapes and sizes and used to automatically animate computer graphics. We ended up with a vast training set of millions of synthetically gener-

ated depth images. Moreover, the graphics algorithm could easily render the corresponding body part images we needed for training as a texture map.

The final piece of the puzzle was how to deal with these millions of training images. My previous work on recognition in photographs had taken a day or two to train from only a few hundred images, and using this approach directly on millions of images would have taken weeks if not months, prohibitive on our tight schedule. We enlisted the help of our colleagues at Microsoft Research in Silicon Valley who had been developing an engine called Dryad for efficient and reliable distributed computation. Together, we built a distributed training algorithm that divided up the millions of training images into smaller batches and trained off each batch in parallel on a networked cluster of computers. Using a cluster of about 100 powerful machines, we were able to bring the training time down to under a day.

All the pieces were in place now, and we worked with the Xbox team to put everything together. Our body part recognition algorithm gives fast and accurate proposals about the 3D locations of several body joints which are then taken and processed by the Xbox group’s tracking algorithm to stitch the skeleton together (another fantastic effort of engineering!) and ensure a smooth, seamless, multi-player experience. This skeletal tracking, together with other new technologies such as voice recognition, give game designers the platform on which to build the magical experiences you get with games such as Kinect Sports and Dance Central.

But of course, gaming is just the beginning, and I foresee this technology fuelling rapid advances in augmented reality and tele-presence, personalised shopping, and healthcare, to name just a few. We are even looking at how touch-free interaction could find its way into the operating theatre so that the surgeon can quickly navigate the patient’s data without risk of contamination from a mouse or keyboard.”

If you are interested in finding out more, please contact Dr Jamie Shotton by email: jamie@shotton.org.

Part IB Group Projects

Jack Lawrence-Jones reports.



Early on a fine Saturday morning last term I found myself, amidst the seething February mist, defending my stumps from salvos of precision balls expertly delivered by my off spin opponent. But what, I hear the astute reader cry, has this to do with *The Ring*? Could it be the inherent parallels between cricket and the tech industry — drive, foresight and a keen eye for opportunities? Or was it the sprawl of laptops and radio receivers surrounding the boundary, recording data being transmitted from the MeLock (©) devices around our necks?

Part of the IB Computer Science syllabus is a 6-week mandatory coursework to be completed in the Lent term. The year is split up into around 10 groups of 6, and each group is assigned to a particular project often devised by a participating Ring member. The projects are chosen to present a challenge to students and are often fairly quirky — previous years have had them building robotic chess boards, Twitter data mining systems and even “Real guitar hero”.

This year we were given the choice of a diverse array of projects. My client was Peter Cowley (see *The Ring*, issue XV). My task was to produce an automated cricket commentary system using just Peter’s hardware — a simple proximity-sensing system originally designed to lock workstations when the user moved a certain distance away.

Although students are allowed to express their preference for a project, the teams are put together entirely at the discretion of faculty staff, with the intention that you end up working with peers you don’t necessarily know very well — no supervision partners or college friends. This simulates a business environment, of which the reader probably knows more than I. I thoroughly enjoyed the experience. My teammates were all very friendly, easy to get on with and assets to the project. Although there were one or two minor disputes along the way, resolution was always just around the corner — I think this was all part of learning how to function in such conditions, and we were quickly having efficient yet fun meetings and managing to set the right tone when working together.

We were given a timeline, with client meetings pre-scheduled every two weeks in order that we maintained a sensible development rate, as it would be all too easy during our first big independent coding project to either neglect our other studies and complete it all in the first few weeks, or conversely leave too much until the last moment. Client meetings comprised of keeping our clients in the loop, ensuring their objectives were achieved, and also enabled teams to ask for help and make additional hardware or software requirements (both in our case — it turned out that our system required the sensors’ firmware to be slightly modified).

Finally Presentation Day arrived. I was tasked with presenting our product in a five-minute speech in front of our entire year, various professors and participating clients and other guest Ring members. I found this a nerve-wracking but valuable experience, and subsequently our team was lucky enough to be awarded the IBM Technical Prize for the most technically accomplished project, as voted by the audience. It was really nice to receive some formal recognition, and the cash prize was a bonus!

For Ring members, it’s worth getting involved in the Group Project scheme for many reasons — not only recruiting future interns, but also a publicity opportunity for participating companies.

What did I take away from the exercise? A new-found interest in cricket, which I certainly didn’t expect to be a product of my academic studies (the World Cup’s on in the background — Malinga bowls... howwwzaat!). It was also a gentle easing into the world of team software development, and has definitely inclined me towards the practical side of Computer Science. In retrospect I liked that we were left to work a lot out for ourselves as it meant we learnt first hand about functioning in a team. Also there was a networking aspect — a few of my peers managed to glean internships from impressed clients. All in all, an invaluable experience.