

Communications Innovation Institute (CII)

Principal Points of Contact

Jon Crowcroft (Cambridge): jon.crowcroft@cl.cam.ac.uk

David Clark (MIT): ddc@lcs.mit.edu 617-253-6003

Project Description

The goal of the CII is to promote the progress and vitality of the communications industry as a whole. The structure of the CII unifies activities at MIT, Cambridge and UCL. Its vision has research and industrial partnership interwoven, with partners from a broad industrial spectrum. BT will be an anchor partner, and the diversity of programs will encourage others to join, as testified by the history of industry-supported research at MIT.

The CII plan for research is framed by a two-pronged approach, each leg of which motivates and nourishes the other. One leg characterizes and models the business structure and value chain dynamics of the industry; in our terminology it lays out roadmaps toward the future. The second leg addresses the enabling and disruptive technologies that can transform communications from a vertically integrated service that companies and consumers have typically *purchased* to a distributed one open to locally incremental innovation. In this way the CII can lay out a richer and more diverse future. Working closely with industrial partners, the overall program will develop the business structures, fundamental architectures and demonstrative applications that show the way.

The CII is a bold experiment both in terms of the vision behind the work, and in the manner of industrial cooperation by which we expect the program to develop. Most important, we are breaking new ground by building a large-scale, multi-faceted program that guides existing players while at the same time gives a seat at the table to companies who either had not considered communications a core competence or had been excluded from its deliberations in the past.

Key Participants

MIT: David Clark, Charles Fine, Frans Kaashoek, Andrew Lippman, David P Reed, John Wroclawski

Cambridge: Ross Anderson, Jon Crowcroft, Frank Kelly, Richard Penty, Richard Steinberg, Ian White

BT: Bob Briscoe

UCL: Tilman Borgers , Ingemar Cox, Mark Handley, Helen Margetts

Table of Contents

1. Institute Outline	3
2. Research Overview	5
2.1. Communications industry value chain	7
2.2. New architecture for tomorrow's networks	9
2.3. Economic & social processes and network design	11
2.4. Policy implications of new communications	13
3. Commercial Enterprise Component	15
4. Educative Component	16
5. Public Agency Component	17
6. Knowledge Exchange Component	18
7. Assessment and Study of KE	20
8. Intended Outcomes	22
9. Budget	24
10. Staffing & Logistics	26
11. Appendices	27

1. Institute Outline

The Communications Innovation Institute unifies activities at MIT and researchers at Cambridge, UCL and BT in a program to promote the progress of the entire communications industry. Key to our vision is that research and industrial partnership are interwoven, with participation of partners from across an industrial spectrum that comprises both traditional sector members as well as newcomers who have not had a voice in the industry but for whom its soundness and health is mandatory. The nature of the program is shaped by and predicated on industrial guidance, participation and responsive action. BT will be an anchor partner, and the diversity of programs will encourage others to join — the history of industry-supported research of the MIT principals attests to this. There are six components of industrial knowledge transfer woven through the work plan that are explicitly noted in Section 6.

The research in this project is framed by a two-pronged, interleaved approach. One branch characterizes and models the business structure and value chain dynamics of the industry, broadly conceived. This activity spans communications providers, suppliers, manufacturers, content developers and consumers and will be done with an international suite of participants that represent the interests and concerns of industry and economic segments. The notion of constructing a roadmap for such a diverse industry is daunting, and extending this to include dimensions of technology as well as use is unprecedented. We do not imagine that we will reach a single conclusion. We imagine that the ongoing process of discussion, refinement and disagreement will be the catalyst of progress and understanding. We expect that this evolving characterization of the rapidly changing communications industry will guide investment decisions by large-scale operating companies and will clarify the opportunities for entrepreneurial newcomers. This work will be closely coupled to the new Communications Futures Program agenda at MIT.

The second branch of our research addresses the enabling and disruptive technologies that can transform the communications sector, and provide our industrial partners a richer view of options for the future. The communications industry is facing a transition from a vertically integrated service that companies and consumers have typically *purchased* to a distributed one open to locally incremental innovation. Our research will help to define and articulate this future.

Heretofore, the technological economies and the regulatory environment have relegated communications to well-entrenched providers. Economic, regulatory, and technical decisions have been made based on speculation about services, users and uses, or with respect to historical and entrenched business practices. For example, spectrum has been allocated to communications providers, not customers, and telephone lines are the property of a regulated (near-) monopoly rather than the users. Today, end-to-end innovations, programmable digital technology and new ways of exploiting the continually increasing capacity of the RF and optical spectrum change this picture dramatically. For example, the recent lessons of WiFi have taught us that given an entrance ticket, distributed

participation can drive social and economic based innovations that legacy companies have abjured or overlooked. Wired and wireless systems can scale without limit and can operate independently of a backbone infrastructure. This lowers the risk barrier for new industrial entrants, enables decentralized services and applications, and permits all companies to integrate communications into their normal products and operations. We see those as ranging from personal expression to integrated business components.

The institute is an experiment in the vision behind the work, the approach to organizing the resources of the universities and in the manner of industrial cooperation by which we expect the program to develop. By reaching across technology, business architecture, regulatory dynamics and economics, CII gathers resources from the university that mirror the challenges facing the industry today. We are breaking new ground by building a large-scale, multi-faceted program that guides existing players while at the same time gives a seat at the table to companies who either had not considered communications a core competence or had been excluded from its deliberations in the past.

The ethos of the program is rooted in the Internet; it is our model. Its end-to-end design principle has permitted innovations and diffusion at rates not seen in staid industries and is a feature not to be disregarded lightly. On the other hand, its lack of central authority or planning potentially impedes the coordinated thrust sometimes needed to make large leaps such as broadband and quality of service provision. These are grist for the research, and they will validate the international and multi-part constituency.

The initial team features the strong presence of BT, as an anchor partner for the program. Through the interaction with the Communications Futures Program, newly launched at MIT in April, 2003, we gain the presence of corporate Media Lab sponsors: BT, HP, ICU (Korea), Learning Lab Denmark, Lego, Mastercard International, Motorola, Swatch AG, Telmex, and the United States Postal Service. We believe this initial set can help entice other usual and unusual partners to join. Industrial transfer features joint, cooperative projects built as cross-industry demonstrations and reports that can be understood by both members of the existing value chain as well as by companies that have no history in the debate. Since some aspects of the work relate to regulated enterprises, we expect representatives of national and international regulatory agencies to be active members. In short, we are the neutral turf on which one can invent new partnerships, where diverse companies are quite literally thrown together to see and suggest research opportunities, and where unbiased approaches to the value chain dynamics of the industry can be developed.

2. Research Overview

The computer industry has long been defined by a tradition of constant reinvention well stated by Alan Kay: “The best way to predict the future is to invent it.” The communications industry has been slow to embrace this sort of rampant innovation, because of its culture, regulatory history, and burdensome investment cycle. The so-called convergence of computing and communications is not a happy marriage for the communications partners, who see themselves saddled with sunk cost but driven by the innovative turmoil of the computer industry. Nonetheless, the communications industry must more and more deal with disruptive change both technically and socially – this is the motivation behind our program.

The research component of the CII is based around four interleaved areas, elaborated in sections 2.1-2.4:

- We will fund research that seeks to understand the value chain of the communications industry and to construct roadmaps into the future.
- We address new architectures for novel implications for both current communications industry members and newcomers.
- We will organize, build and test interdisciplinary and inter-industry approaches to heretofore intractable problems such as routing and quality of service that are based on economics options and social utility.
- We will formulate communications policies that are informed by technology and pursued in local and international fora.

Roadmapping, described in Section 2.1, is near the heart of the project. Our goal is to work with academic contributors with specialties all along the value chain, and to engage industrial partners drawn from a similarly broad base. This work will be centered in techniques and methods typically associated with schools of business and economics, but will be richly informed and refined by the broad mix of players in the project.

The second thread of our research addresses disruptive factors that are often overlooked by those who extrapolate current trends. This component is more technically centered, and is charged with bringing forth new architectures with novel implications for industry and for society. As Section 2.2 will detail, the core technologies of tomorrow’s communications infrastructure—wireless, fiber, digital signal processing—can be combined in novel ways that are more than technically intriguing and can redefine the commercial opportunities for our industrial partners.

These two research components thus complement each other. Our technical research, inventing alternative visions of the future, will impose a discipline on

the roadmapping activities and force them to contemplate the potential for disruptive shifts. At the same time, the roadmapping activities, centered in the tradition of business and economics, can demand of the technologists that they justify the real-world viability of their ideas, not just their “coolness”.

In section 2.3, we broaden the context. It is no longer sufficient for technical research to be academically novel. To have impact, the researchers need to take into account the economic, societal and policy implications of their work. If they can be persuaded to broaden their agenda in this way, it is often possible to use these new dimensions as part of the design. In other words, these exogenous factors such as economics become a part of the mechanism. Thus, for example, there is a small but growing trend looking at the economic factors impeding the deployment of better security, and economic tools for capacity allocation in networks. On the more theoretical front, game theory is becoming a tool of computer science research. Our existing industrial partners have indicated that they find this work of great importance.

The final component of the CII research plan, detailed in Section 2.4, is an agenda of technology/policy research. Most practitioners of policy-making have a background in economics, law or political science. As the pace of technology innovation increases, there is increasing need for policy-making to be grounded in current technology, both its constraints and its opportunities. By including a program of policy research in the CII, we can link this to our strong technology component, and thus cross-fertilize both policy and technology research, just as we will cross-fertilize the technology and the road-mapping.

All of these examples illustrate the overarching criteria we impose for CII research. It must be interdisciplinary, it must bear on real-world concerns, it must be innovative, and it must inform in some way our overall goal of aiding the health of the communications industry as it will become. We observe that research that takes into account both technical and social factors is hard to fund from traditional governmental funding agencies, both in the US and UK. Policy research is in general hard to fund, even though it is critical to the government’s own needs. Interdisciplinary work is risky for discipline-centered researchers, and they need the encouragement of a supportive community of like-minded individuals and suitable funding. CII can provide this, and has the potential to shift the center of gravity of communications research.

The most ambitious outcome of this research is that we build a neutral table for the development and realization of a vibrant, revitalized communications industry with new dimensions, new participants and new opportunities that combines ideas with the public, economic and cultural drivers.

2.1. **Communications industry value chain**

A core component of the CII research plan is value chain roadmapping. By this we mean the collaborative undertaking among members of the communications value chain to posit and explore multiple scenarios for various components of the industry, e.g., wireless personal communications, the home network, e-business services, last mile provision, etc. We think of this value chain as including users (consumer & corporate), service providers, appliance manufacturers, software, application, and content providers, network owners, network builders, equipment companies & system houses, component and device makers, materials and process equipment companies, and R&D organizations (including universities). Naturally, many of these members, in turn, have their own supply chains, so the value network is quite extensive.

We believe that such a roadmapping exercise is quite timely for the communications industry. History suggests that phenomenal productive growth can often result from an industrial environment featuring decentralized, free-market, “bare knuckles” competition (e.g., the personal computer industry in the 1980’s and 1990’s.) However, one can also observe industries and eras where some industry-wide coordination or leadership, driven either by government forces or powerful vertically-integrated firms, has supported rapid growth and wealth creation. During the middle part of the twentieth century in the United States one might have characterized the electric power industry (government coordination), the automotive industry (General Motors), and the telecommunications industry (AT&T) in this way.

In the communications industry today, many of the challenges to growth can only be addressed by resolving issues that cut across the value chain. Contentious issues, from digital rights to access competition policy, have constituencies from many components of the chain, as well as from consumer groups and regulatory/legislative bodies. In the absence of some degree of coordinated or collaborative processes to break logjams, the communications sector might remain moribund for quite a long time. Further absent the coordinating power of a large vertically-integrated industry leader (e.g., the old AT&T), we hypothesize that a roadmapping effort to provide some common ground (built on solid research) among members of the value chain and policy makers might provide some of the coordination that is missing in the industry today. (We think of the semiconductor technology roadmaps, collaboratively developed by members of Sematech and the Semiconductor Industry Association to aid the realization of Moore’s Law, as illustrating the value of this type of coordination, albeit in a much simpler value chain context with a much narrower roadmapping scope.)

We prefer to use “roadmap” primarily as a verb — the **process** of discussing and assessing future scenarios is a more fruitful undertaking than setting a target to write “*the roadmap*.” We will create multiple possible future roadmaps and assess features of each. The research tool for this will be systems dynamics modeling. Formally, we will build systems dynamics submodels for the dynamics of each of the following: technologies, business cycles, industry structure,

corporate strategy, public policy, capital markets, and consumer preferences. These submodels will be built from inputs from specific technology roadmapping efforts (see below) as well as data collection from industry and interviews with participants along the entire value chain. The resulting model will then be used for a range of “what if” exercises to explore the implications of policy options for the public sector as well as strategy and investment options for corporate and financial decision makers.

Some of these scenarios might appear sufficiently compelling as to trigger a set of companies across the value chain to pursue such a vision. Other assessments might convince participants that certain scenarios are unsustainable, resulting in better-focused research investments and business models. Additionally, some analyses and assessments might productively inform government policy.

Two of the technology roadmapping projects that will be pursued in depth are:

Roadmap for Smart Photonics:

The impact of cost-performance advances in photonics is not only raw speed, but also new freedom in higher layer design. Optimality will not be realised without KE between physical engineering, network control and socio-economic researchers. Understanding from building innovative, functional hardware already funded elsewhere, will lead to real advances in Internet provision. Deliverables will be road-mapping and technology transfer studies. This work will strongly interact with our research on wireless over optical and low cost WDM (Section 2.2), and that on new routing and control architectures (Section 2.3).

Roadmap for Wireless and Spectrum Usage:

The technical and commercial steps necessary to exploit recent results across disciplines showing that adding *co-operating* wireless nodes to a space can *increase* capacity and an incentive scheme can be devised to encourage co-operation. Although we have modulation & antenna design advances, problems of practicality (robustness, security, industrial incentives to deploy) mean that we need to convene the relevant groups, to develop a road-map with industry that limits exposure to risk balanced against support for services to the public. Initial deliverables will cover: i) resale of licenses; ii) ubiquity of WiFi; iii) multihop ad hoc radio. There is synergy here with our research plan in the areas of viral communication and 802.11ng (Section 2.2), and in the use of economic and social techniques for network control and the evolution of P2P (Section 2.3).

These technology-driven approaches will inform, and in turn be informed by, the overarching roadmapping exercise, ensuring the production of roadmaps which are technically, economically and publicly viable.

2.2. New architecture for tomorrow's networks

The second component of the CII research plan is the exploration of new and potentially radical and disruptive views of networking. These projects blend insights about technical capabilities and user needs to yield new proposals for network architecture: how technology parts are assembled to make a system. These projects are primarily centered in areas where we see rapid technical innovation: wireless and photonics. Our goal is to make the point that wireless and photonics are not best thought of as “better copper wires”, but as tools to build whole new sorts of networks. This research is designed to go hand in hand with the value chain and roadmapping work described in Section 2.1.

Viral Communications

We define a *Viral Communications Architecture* to be one where elements are independent, scalable and where each new element adds capacity to the system, so that a viral system can be adopted incrementally from a small base and gains accelerating value with scale. Examples include embedded devices whose functionality is significantly altered by communications capability, such as environmental sensors and actuators, inventory control systems, and monitoring systems. For these applications, a device is installed once and expected to continue to work even as new systems and devices are invented and deployed.

Personal communications are also amenable to viral techniques; e.g. one can envision a wireless telephone system where proximate phones talk directly to each other, interacting with a central antenna only for call setup. Such a system would localize channel re-use at the expense of a more complex end-user circuit; a reasonable option as digital processing reaches RF speeds. Such a telephone can operate in both point-to-point and broadcast modes, can carry private transmissions, and can propagate safety messages. The challenge here is to demonstrate that increased system utility outweighs the cost of complexity.

The novelty here is that we view bit distribution as a decentralized function that exploits intelligence and collaboration among the RF elements themselves. The collaboration exploits degrees of freedom that grow with the number of elements sharing a space and thus increases throughout as new elements are added. This view of the spectrum and its capacity is what has the potential to both disrupt and extend what the communications industry is, and who its members are.

802.11ng and Routing

802.11 (trade named Wi-Fi) is an ultra-low cost, easy to deploy wireless communications technology. Because it is low cost and operates in the “unlicensed” ISM bands, its deployment is easily financed by users themselves. The future of Wi-Fi architecture may support a “viral” scaling and evolution, in which user-owned equipment forms an infrastructure network cooperatively without central control, but stabilized by mechanisms arising from game theory, applied probability and the mathematics of network flow.

We propose experimenting with evolution of these IEEE 802.11 standards to include low level, built-in distributed control of the resources with integrated trust mechanisms. This will draw again on the combined expertise of economists, modelers and communications specialists at MIT and Cambridge University. By combining forces with our industrial partners (at BT, Intel, Motorola, and Nortel Networks), we expect to be able to generate significant industry leverage to prototype (and test) new chipsets that incorporate our design, and therefore evaluate real results that feed directly to both the standards committees (IEEE and related) and the regulators.

Wireless over Optical

The development of much more powerful mobile networks would open up new forms of commercial exploitation with advanced real-time internet mobile applications. We therefore propose building on recent results at Cambridge and UCL where we have demonstrated RF over in-building multimode fibre with very low distortion and SNR degradation. We will develop technologies which will allow much more flexible wireless networks to be deployed; for instance an architecture involving centralization of bandwidth in an enterprise central office with radio over fibre distribution of the RF signals to a network of antennas. This would consequently allow the dynamic redistribution of wireless LAN signal bandwidth to cope with non-uniform user bandwidth requirement distribution, or enable the simultaneous operation of several wireless standards over the same infrastructure, encouraging evolutionary advances. After successful initial prototypes, it is proposed that full deployment within a large scale WLAN network be carried out in Cambridge UK. This in turn would lead to opportunities for both companies and universities to carry out new higher level systems studies.

Low-cost WDM

We propose to develop and implement novel optical network topologies, primarily for local computer related networks but with potential scaling to the metropolitan area. In particular this work will study the use of multiple wavelengths and novel protocols (either fixed or rapidly reconfigurable) to enhance network performance. Effective utilization of the optical transport layer can be delivered through a distributed traffic control algorithm which makes use of the computing power and memory capacity at the edge of the network. The load control technique must deliver the richness service classes required for commercially successful data communications products, including QoS for some flows and rate adjustment on the timescales of Internet traffic dynamics for others.

New optical technologies such as wavelength division multiplexing can be expected to make great impact beyond the area of telecommunication transmission systems. One thrust is to extend the work to the new field of chip-to-chip networks, which is acknowledged to have considerable potential in the near future. The primary output of this research therefore will be the development of a new network technology for a range of computer and storage LANs, but it will also seek to determine optimum methods for reducing size so that the technology can also be applied to board-to-board bus and (later) chip-to-chip applications.

2.3. *Economic & social processes and network design.*

At certain points in the evolution of a network architecture there are points of structural flexibility, where pressure for change overcomes incumbent solutions and industry inertia. This change is triggered not by technology, but the larger economic and societal context in which these architectures exist. Today, there is a growing pressure for exploration of new network services: allocation of capacity and enhanced QoS, security, improvements in robustness and availability, and efficient placement of application-level functions within the network.

What marks most of these problems is that the barriers to progress are not purely technical, but commercial, legal, and societal. Research has provided potential technical solutions that have been ignored because they do not take into account this larger context. For this reason, the CII research component contains an explicit part that supports research that takes technical innovation and folds in these larger issues as co-equal factors in the design.

New Architecture for Routing and Addressing

The Internet hierarchy is still frozen in a topology determined by early 1980s technical decisions. A primary goal of this work is to design new inter-provider routing architectures that take account of current requirements. The purely technical side of the work is centered on the successor to BGP, but it has far broader ramifications. One area is in user choice of routes, which can foster competition and to allow the selection of services offered by one or another provider. The end user today desires choice (for QoS reasons, for rapid provider selection, or for resilience), as does the access network provider, choosing to select different long haul carriers for similar reasons.

Furthermore, users want an architecture that protects them from attack (denial of service, spam), with control over who can send to them – unlike today's transparent end-to-end architecture. Piecemeal solutions to denial-of-service are insufficient, invariably exposing alternative avenues for attack and blocking future evolution potential. An across the board set of architectural solutions, grounded in technical economic & legal realities, will create innovation opportunities in the security industry along with practical deployment options for service providers.

This project will begin with a study of existing work, organize a workshop on the economics of routing, propose a strawman set of routing mechanisms, evaluate these with a second workshop, transfer this proposal into the Internet standards development process, and hold a final workshop on the broader issues of inter-provider business issues. The involvement of industry is critical.

Internet Congestion Control using Pricing

Debates over congestion pricing expose the tension at the heart of this Institute: that between a commoditised value chain open to distributed innovation and one where communications services are purchased from powerful vertically integrated operators. There is a need for resolution in two main areas: (i) analysis of local and global stability given the realities of feedback delay and of varying

application stability requirements, and (ii) relaxation of theoretical idealisations, instead allowing for the market realities. With regard to (i), the idea is to combine insights from control theory and economics to develop improvements to current schemes. With regard to (ii), as an example, on the demand side customers' willingness to insure against price instability reduces responsiveness from the ideal suggested by theory, while on the supply side, operators will wield their market power, holding back market information to prevent perfect competition.

Economics and Design of Third-Generation Peer-To-Peer Systems

Since mid-2000, there has been an explosion of interest in peer-to-peer systems – the business of building useful systems out of large numbers of intermittently connected machines, with a virtual infrastructure tailored to the application. Our goal is to build a next-generation peer-to-peer system and evaluate it by means of real applications, including a digital library system and an open news platform. The news platform will enable users, such as journalists and publishers, to post material (text, audio, or video) directly, and also to correct, syndicate and re-use other people's content to add value. Secondary applications such as e-clippings agencies will be easy applications to write; this will foster rapid innovation.

Building the platform will involve designing reputation systems that can work in highly distributed environments that are open to hostile attack. In the process, we will apply ideas from economics to shared control of distributed systems. Computer scientists think of this in purely instrumental terms such as Byzantine fault-tolerance and threshold signature mechanisms. Economists view the problem more broadly, in terms of the Arrow impossibility theorem and Sen's information broadening. Recent work on mechanism design provides a bridge between the two worlds. Peer-to-peer systems require broader control mechanisms than have been available through technology alone; we will explore a number of ideas about how economic mechanisms may be used to make them far more robust. We expect many of the mechanisms and ideas generated here will be applicable to the 802.11ng and viral communications areas (Section 2.2).

Participation in Internet-Mediated Interactions

A variety of government policies in many countries are aimed at promoting the use of the Internet across and between society, the economy and government, thereby increasing Internet-mediated transactions of all types. Implicit in these policies is the view that the take-up rates that would result without government intervention are too low. However, it is rarely articulated why government should play this role – and why these means of communication should be prioritised over others. Our aim is to develop theoretical models that allow us to assess the social optimality of take-up rates, and to take these models to data where possible. We will also look at the rationale for government target regimes aimed at increasing adoption rates. Principal-agent problems arise when central government agencies command public officials across government departments to introduce innovations, particularly where they are technologically based. Targets are one way of incentivizing officials – but they may not be the most effective method of overcoming principal-agent problems.

2.4. Policy implications of new communications

Wireless Spectrum Policy

Historically, communications has been regulated as a scarce resource whether in wires or in the air. As a result, it has generally been controlled as a monopoly and it has conflated information delivery with implicit services. Viral Communications stands in stark contrast by permitting incremental, open, scalable access to resources. By moving the intelligence to the end nodes of the system, this has the potential to separate delivery from service provision.

To date, systems that become available for user-financed deployment (e.g., Wifi, or 802.11) have used unlicensed spectrum, but they don't scale: they are viral in their *impact* but use a saturating technology: more users equals less bandwidth per node. Far more wireless capacity can be created by viral techniques, using collaborative and adaptive spectrum sharing than has been heretofore thought. This has been demonstrated in theory [c.f., Shepard, Gupta and Kumar, Gatspar and Vetterli] and our technical work will realize it as a compelling demonstration. However, its implications are not reflected in the current debates on spectrum policy. The simple fact that radio capacity can *increase* as more communications devices use it has ramifications on ownership, allocation, cost and power; the addition of devices whose operation does not rely upon a pre-existing infrastructure threatens the tenets of regulation that have existed since 1914.

We propose to carry out policy research on viral communications in a context of applications, demonstration, and experiments done with industry. The results of such research will have immediate value to Ofcomm, WARC and national regulatory agencies as well as the communications industry by enabling policies that open up wireless communications to innovative applications and new sources of investment. Indirectly, we expect such policies to lead to significant new sources of economic growth. We will explore these opportunities as an integral component of our research. The ability to use our three universities as neutral proving grounds and as points of origination for new policies is a key aspect of the program: we can thereby empower new entrants with a voice not normally heard in policy debates, and act within an impartial framework.

In parallel we plan to carry out an investigation of Government policy towards spectrum use. Specifically we would like to investigate: (i) do spectrum licenses need to be specific about the application for which they are used (like UMTS licenses), or does that hinder efficiency of spectrum allocation and innovation? (ii) does bidding in spectrum auctions confirm auction-theoretic predictions? (iii) does unregulated bilateral trading in spectrum licenses improve the efficiency of spectrum allocation, and does it allow innovators to acquire the spectrum rights that they need to implement their innovations? (iv) will contracts to compensate for interference lead to efficient spectrum use? In all areas we shall proceed by developing appropriate theoretical models, collecting evidence that might bear on the relevance of these models, and finally investigating policy implications.

Public Policy and Deployment of Broadband

There is much rhetoric about the “digital divide” and the inequitable deployment of broadband services. As a study from the US National Academies noted [cite], it is premature to conclude that there is a persistent problem here, and too soon to call for intervention. None the less, laws are being passed, programs such as rural subsidies are being created, and advocates have entered the field.

One of the most interesting activities is the direct involvement by un-served and underserved communities (often rural small towns, but not exclusively) to take direct action to deploy broadband. These actions include direct investment in infrastructure and operation of the system to provide services to the population. This project will track this phenomenon, pose and test hypotheses about why communities choose to act and when they are successful, and provide a summary of “best practice” to communities and policy makers. Past work in this area has equipped us to compare activities in the US and UK, and the MIT ITC has already acted as joint sponsor of a UK workshop on this topic.

Internet Piracy

There is currently much interest and concern regarding intellectual property rights and the Internet. The music and movie industries have acted as a catalyst for this interest yet, surprisingly, other IP-based industries have not been so alarmed. The content-industry's response to the Internet is varied and includes technological, legislative and legal strategies that have the potential to significantly affect the future of the Internet. However, it is unclear if the music industry's current malaise is a direct result of the Internet or coincidental.

For example, in 2001 USA Today (Monday, Feb. 25th 2002) reported a 10.3% drop in overall shipments of recorded music to retail stores. However, Soundscan measured only a 2.8% drop in direct sales (rather than shipped) for 2001. The latter number is more meaningful. In a recession, distributors will reduce the inventory leading to a larger decline in shipments versus direct sales. A 2.8% drop in CD sales in a recessionary year (2001) may be reasonable and unrelated to Internet music file sharing.

While the causality between music file sharing and reduced music sales has not been proven, there are very strong efforts to reduce internet piracy. Clearly, copyright holders deserve economic reward. However, it is imperative that the effect of the Internet on the music industry be better understood. And it is imperative that any solutions to Internet piracy maintain the health of both the content and communications industries.

This project seeks to provide reliable and impartial data from which a balanced solution to the piracy problem can be derived. To this end, we intend to (i) quantify and characterize the extent of music copied over the Internet, (ii) quantify the effect of music file sharing on the music industry, (iii) investigate the motivation for music file sharing, (iv) quantify the effect of current and proposed solutions to inhibit music copying and (v) investigate alternative solutions. This work will draw on separately funded parallel research being carried out at CU.

3. Commercial Enterprise Component

Key to the CII vision is a strong partnership with industry, not just for the dissemination of academic results, but to foster collaboration on the research. In particular, the roadmapping project will require strong industry participation.

All of the participating universities have a history of industrial partnership, and all have pre-existing collaborations of various scale. The founding partners already have wide-ranging contacts, including the members and associates of the MIT ITC, MIT Media Lab, FIPR, and others. Individually, we have contacts into a wide range of firms. Our plan is to fold these together into a larger whole of industry participation, and at the same time reach out to additional players, especially from parts of the value chain that are not usually represented in “communications” projects. The plan for how we merge these various pre-existing relationships is not yet concluded, and it will take some time to create the final construct, since the existing relationships were build using very different models of financial support, rights to intellectual property, and so on.

MIT has recently launched the Communications Futures Program, and as a result has had to contemplate some of these issues internally. Our view is that more than one sort of industrial participation must be accommodated. One is “top down”—a firm that is primarily interested in the overview that the roadmapping activity provides. Another is “bottom-up”—a firm with a specialty in one area that grows to understand why the bigger view helps their specific problem. A “top down” firm might join the CII directly; a “bottom-up” firm might first involve itself in a specific laboratory or research program, and then explore the CII from that vantage. Additionally, we anticipate that we will have to devise different means to reach out to big firms (which can commit both financial support and staff to engage us) and small firms (which must maintain a tight focus on their immediate goal, but which non the less benefit from exposure to our roadmapping.)

An important component of the CII plan for commercial enterprises is a relationship with CMI@Adastral, which will open up paths to UK start-ups and SMEs, and convene industry groups where required. **BT** helped create this Institute and will self-fund its participation in order to insure that its own plans for the future build on a broad industry vision. BT's desire for consensus led to the creation of its Future Communications Architecture research program, the main aim of which was to encourage transatlantic collaboration. BT is therefore fully committed to making wider academic and industrial collaboration a success.

In photonics, the academic groups have existing strong links with industries including Agilent, Bookham, Marconi, and Nortel. The Institute is novel in proposing that a Cambridge based company, **Scientific Generics**, play a key role in promoting new exploitation routes benefiting from its road-mapping exercises in the photonics & RF projects. Exploitation routes will therefore be managed jointly as a partnership between industry and academia, the Generics expertise being key in involving the widest range of UK companies.

4. Educative Component

UCL's School of Public Policy along with the Computer Science and Economics departments plans to create a new interdisciplinary **Masters level degree in the Internet Economy** for 2004/5. It will include courses in e-government, e-commerce and web-based innovation.

UCL's extensive postgraduate degree programme, includes the innovative MSc in Public Policy, the unique advanced MSc in Data Communications, Networks and Distributed Systems, and substantial outreach to industry through the **UCL@Adastral** interdisciplinary centre and the **BT MSc in Telecom Engineering**.

Currently within the academic partners, a wide range of taught Masters courses are provided, and in addition specific modules are provided for Masters courses run elsewhere (for example in displays where Cambridge contributed to a UK wide course, and in optical data communications where Cambridge runs a module for a European Masters Course based in ENST Brest). This Centre will seek to enhance the knowledge base available to industry, not by developing new Masters courses (as we believe that there are sufficient already), but rather by making it more straightforward for colleagues in industry to take modules from a large range of Masters courses (both in technology and socio-economics), and indeed take some which are industry-based. This would allow those in industry to focus their activities on their own needs, the scheme allowing the award of a degree or certificate once a sufficient number of modules have been taken

The academic and industrial partners are all keen on visits and exchanges between each other, other universities and industry. Cambridge and UCL participants are particularly keen that Institute funding (whether initially by CMI or beyond that) will allow industry to have PhD internships every year, avoiding the "finish in 3 year" EPSRC funding rule. UCL already endorses a 2-way flow of academics and industry people - so they can second someone from academia to a start-up but also second someone from industry into the university(s) to add reality to a course. Institute funding is also ideal when a project needs interdisciplinary research, which would otherwise require rounding up the relevant experts across departments.

5. Public Agency Component

The successful achievement of the objectives of this institute in creating new, advanced businesses around the Cambridge-Ipswich corridor would justify the £2M support for UCL@Adastral from the East England Development Agency (EEDA).

- Standards:

The Institute will also have a strong role in International Standardisation fora, both educating those setting standards and directly proposing standards. It will engage with both technical (e.g. the Internet Engineering Task Force, the Global Grid Forum) and social/political standardisation bodies, whether public agencies, or self regulating industry bodies (e.g. the UK Internet Watch Foundation).

- Regulatory:

The FCC is already strongly engaged with MIT's ITC programme. We will welcome research commissioned by Ofcom or government departments into how public policy can become more innovative through the use of technology (see also the research components on "*New architecture for tomorrow's networks*" and "*Policy implications of new communications*").

Since the consortium can be expected to involve several new companies, it would be hoped that we would see the direct involvement of DTI, and EEDA and other development organizations in other areas linked to their partner companies. Special involvement will exist with the newly formed Centre for Photonic Integration at Adastral Park, where many of the necessary hardware components can be fabricated. Standards activities are strongly supported and hence the involvement of the IEE (Professor Seeds, being chair of the Photonics Network), IEEE and BCS will be strongly welcomed. It is additionally hoped that the CII would form close links with related UK initiatives such as the Grid and UKlight.

6. Knowledge Exchange Component

The constituent groups engaged in this program are motivated by the belief that communications evolution requires a broad, collaborative vision. To suit that end, we propose six novel ways by which the research in the program will be integrated into industrial methods both for large-scale industry members and entrepreneurial newcomers. We argue that this dissemination style is inseparable from the content of the technical program because of the very nature of the industry. Not only is communications a technology that impinges on basic freedoms of expression and is a basic staple of commerce, it is one whose history is intertwined with regulation and legislation. It is uniquely broad and pervasive; structures are slow to change, and innovation often occurs in spurts. Existing firms evolve slowly and entrepreneurs cannot do it alone.

Further, we endorse the notion that one measure of an invention is the speed with which it is appropriated from the inventors and remade in the image of the society that comes to own it. SMS is an exemplar -- it was invented for simple communications from the operator to the user but kids transformed it into a profitable, local service. Therefore, the lynchpin of the program is an innovative coalition of industrial partners from diverse industrial sectors as both the supporters and beneficiaries of the work. The team will be led by researchers at MIT, Cambridge, UCL and BT. We expect that within two years, there will be 30 other members drawn from the developed and emerging worlds, and including both small and large industries. These members are the experimental arena and the first distinctive element of the program. This partnership will be sustained through a traditional program of annual meetings, workshops, reports, exchange visits and so on, but also in the innovative ways discussed below.

The existing BT/CMI collaboration that we engage at the outset is the second experiment in knowledge diffusion. We have a history of working with BT that has already motivated their planning. With this program, we engage deeply into the operations of the company. We stress doing this through enabling opportunities and learning methodologies rather than toy applications. By mutual agreement, we will use BT as an experimental base. We will implement some of the ideas in BT systems and thus demonstrate how cooperation between academic theory and industrial realization can work.

Third, within the university we will construct a managed "horizontal overlay" on the traditional vertical structure laboratory-centered research. This layer in the KIC will deploy staff with the primary duty of translation of results from a form produced by academic researchers into a form that is directly accessible to industrial players, and will carry the responsibility of bringing forward the research results that feed into the roadmapping and planning process. It is our belief that high-level staff are required to insure that industry can extract value from academic research.

Fourth, we propose that we can release the "pinch points" that have impeded innovation in the past by establishing working groups mediated by unbiased

facilitators (or umpires) to foster communications between the partners. In the past, this facilitated element has either been absent or relegated to the courts. Consider that the very existence of VCRs in the US was threatened by content owners in the 1980's. In the best case they ever lost, the movie industry has benefited far more from the opportunity of the new market than it has sacrificed from its initial threat. Television is likewise threatened by personal digital recorders, the telephone industry is threatened by WiFi, and broadband is prevented by competing elements each demanding investment payback. A mediated link between disparate and potentially opposing interests can shortcut the process and speed innovative diffusion.

Fifth, we propose to learn from youth to drive discovery of startup opportunities. We recognize that *use* defines the success and utilization of communications innovations, but instead of inventing speculative uses *ab initio* in the lab or the sponsors, we exploit the resource of learning programs in the Media Laboratory to teach us how to connect basic ideas with social practice.

Sixth, we postulate that the interplay between big and small companies is far better than a strict focus on either. Like the smaller animals that piggyback on whales, the symbiosis between infrastructure providers and agile innovators is potentially stronger than either alone. Therefore, we do optimize the program for both working together. By contrast, most recent academic/industrial co-operations have focused on startups alone and have lacked the ability to propose entrepreneurial applications that can only work with cooperation of existing, large-scale players. In particular, if the only model of technology transfer is via venture supported startups, the faculty themselves may be the vector of transfer. But if one intends to reach out to larger firms in the field, a different approach (e.g. our managed horizontal overlay) is required.

Two-way understanding of context and problems, as well as two way exchange of research ideas and results between academia and industry does not 'just happen'. The MIT Internet Telecommunications Convergence program (ITC) has been a model of successful knowledge exchange between Internet architecture researchers, public policy academics & practitioners (e.g. the FCC), and industrial partners and associates. The **formula** involves directed funding and careful choice of an executive team also practicing the research. The ability of social sciences researchers to place themselves in a *believable* technical future envisioned by world-renowned Internet architects has created this unique flavor. Industrial partners are empowered to invest in ideas validated by such social, economic & technical assessment.

7. Assessment and Study of KE

Any consortium funded by industry that expects to survive must, of necessity, perform constant assessment to ensure that partners are finding value in the program. The “bottom line”, of course, is whether members sign up, refer other members, and renew. But one cannot wait for that signal; discovery that something is wrong comes much too late.

Much assessment is informal and personal—direct conversation with sponsors to determine what their attitudes are about the program. But this sort of anecdotal process needs to be complemented with more structured vehicles. MIT consortia use questionnaires to evaluate meetings, regular calls from staff to key contacts inside partner companies, advisory boards made up of key sponsors, and similar tools to assess their value to industry. CII will implement tools of these sorts for each of the parts of our KIC plan, and we welcome involvement of CMI in this process.

These comments are most relevant to the interactions between academia and larger companies—companies with the capacity to join in a sustaining partnership. Smaller companies need a different mode of interaction, and thus a different mode of assessment. The traditional model of small company interaction is the direct transfer of an idea into the commercial sector by venture investment and new company creation. Many labs track the number of “spin-offs” as a long-term measure of success, but the problem with this metric is that it is long-term. The process takes time, and the success of any one venture cannot be assessed for several years.

One novel activity that we have contemplated is the direct involvement of the venture capital community in the CII. What interests us is the question of why they choose *not* to fund proposals. Of course, many an idea is not funded because it is ill-formed, unrealistic, or has little apparent market demand. But anecdotes would suggest that many otherwise good ideas are rejected because there are external conditions in the market that throw up barriers to the innovation. In some cases, the barriers are intentional, but in many cases, barriers seem to reflect what we have called “pinch-points”—coordination problems that make it impossible for one sector to move independently. In a case like this, venture investment is blocked. A venture investor cannot reshape an industry; they can only invest in it as they find it.

Research and dialog centered in academia, because it is neutral, has a better chance of reshaping industry structure. Direct involvement of venture investors can help us identify these pinch-points, can help us assign priority to the creation of working groups to consider such problems, and can provide direct and immediate feedback as to whether our activities are being helpful. The measure of success is that venture investors tell us that we unblocked investment opportunities, independent of whether the underlying innovation came from inside the CII. By including the venture community in the CII, the scope of our impact is broadened, and the immediacy of the feedback is improved.

Summary: our “value proposition”:

The ***novel innovation for knowledge exchange*** we propose is value chain roadmapping. By this we mean the collaborative undertaking among members of the communications value chain to explore the joint technological, economic, and policy feasibility of certain value chain scenarios.

Our hypothesis is that a university is the optimal nexus for such a value chain roadmapping effort because the university can serve as a “neutral table” or “honest broker” for such interaction as well as contribute to, debate, audit, and refine the technological and business assumptions held by the various members of the industry. Having *three* universities with a variety of expertise, experience and culture means an broader scope of knowledge and an increased sphere of influence, but retains the neutrality we believe is required.

The ***measurable results*** we expect are that new ventures and technologies will be launched by participating industrial collaborators in the roadmapping process as a result of the collaborative effort undertaken.

We welcome the participation of CMI in the assessment process.

8. Intended Outcomes

Broadly speaking, the intended outcomes of CII are:

- To create joint industry/academic solutions to architectural problems that hold back innovations in communications: solutions that are socially and economically driven and feasible technically.
- To create the multi-faceted industrial consensus required to evolve the communications industry to a more open, responsive, vibrant and innovative one:
 - By giving voice to newcomers
 - By creating bottom-up and distributed structures that lower the thresholds of entry and innovation
 - By creating architectures that allow localized, small-scale, incremental change
- To explore through experience and experiment the social drivers of communications innovations, and to map those into industry imperatives.
- To learn through the experience of young inventors rather than speculation about "what people will want."

In more detail, some selected milestones and outputs we expect to produce are:

Month 6:

- Initial roadmap for 802.11ng
- Software tools for Internet piracy traffic analysis
- Workshop with business and government on spectrum policy issues.
- Report on remit and constituency for smart photonics roadmap

Month 12:

- Workshop on the economics of route choice
- Smart photonics roadmapping workshop (with academic and industrial participation)
- Report on modeling and measuring the usage patterns of Internet-mediated interactions
- Workshop on 802.11ng ad hoc economics
- Communications Industry value chain & roadmap document (alpha)

Month 18:

- Design for centralised wireless-on-optical LAN
- Small scale demonstrator of next-generation P2P system.
- Realisation of electronic control for cheap WDM network node

- Demonstration of congestion pricing solution addressing the tension between infrastructure investment and innovation.

Month 24:

- Larger scale implementation of next generation P2P system.
- Report on spectrum policy implications; In-depth policy workshop with participants from business and government.
- Completion of work on principal-agent theory and target regimes; Seminar on government policy and intervention in take-up of internet-mediated transactions
- Final report on impact of Internet Piracy, and proposed solutions.
- Communications Industry value chain & roadmap document (beta)

Month 30:

- IRTF Routing Research Group Input - Informational RFC
- Smart Photonics Report on Network Strategies
- Implementation of dynamically reconfigurable centralised wireless-on-optical LAN network with access across specific sites in Cambridge.
- Completion of OPS cheap WDM network testbed

Month 36:

- Workshop on the economics of next generation routing & addressing
- Final roadmap document for 802.11ng
- Smart Photonics Final Roadmap Report
- Wiress-on-Optical: Network testing using live traffic.
- Extended form of next-generation P2P media platform
- Communications Industry value chain & roadmap document (gamma)

Due to space constraints, these represent a selected subset of the outputs we shall produce. In particular we envision a number of collaborative workshops and meetings within the CII to allow 'cross pollination' of ideas between researchers focusing on the four different components described in Section 2. Although we expect these collaborations to produce significant results in both the research components *per se* and in terms of knowledge exchange, it would presume to declare their precise outputs at this stage.

9. Budget by institution. First year and total 3 year costs.

PERSONNEL SALARIES	MIT YEAR 1	CU YEAR 1	UCL YEAR 1	MIT TOTAL	CU TOTAL	UCL TOTAL
Faculty/Senior staff	\$160,982	£30,000	£73,530	\$502,516	£90,000	£152,941
Post Doc/Tech staff (62% MIT OH, 70% CU OH)	\$72,558	£414,325	£145,883	\$226,496	£1,242,975	£437,649
Post Doc/Tech staff (10% MIT OH)	\$72,558			\$226,496		
Support staff	\$20,580	£9,496	£3,529	\$64,244	£28,488	£10,587
Research Assistants	\$164,916		£44,118	\$514,801		£132,354
Personnel allocation (No OH)	\$20,188			\$63,034		
Sub-Total	\$511,782	£453,821		\$1,597,587	£1,361,463	
Employee Benefits	\$80,784			\$267,405		
Vacation Accruals	\$19,476			\$60,086		
TOTAL SALARIES & BENEFITS	\$612,042	£771,496	£267,060	\$1,925,078	£2,314,487	£733,531
OTHER DIRECT COSTS						
Equipment (No OH)		£43,750			£141,250	
Foreign Travel	\$162,694	£103,750	£33,320	\$505,252	£311,250	£100,000
Materials & Services	\$10,000		£5,000	\$31,216		£9,000
Computer Resource Services	\$17,896			\$56,326		
Res. Asst. Tuition - No OH	\$72,324			\$289,984		
Materials & Services Alloc.	\$9,713			\$30,182		
TOTAL OTHER DIRECT COSTS	\$272,627	£147,500	£38,320	\$912,960	£452,500	£109,000
TOTAL DIRECT COSTS	\$884,669	£918,996	£305,380	\$2,838,038		£842,531
INDIRECT COSTS (F & A) MIT @ 62% , UC @ 70%	\$448,496	£317,675	£186,942	\$1,408,954	£953,024	£513,472
TOTAL COSTS	\$1,333,165 £833,228	£918,996	£492,322	\$4,246,992 £2,654,370	£2,766,987	£1,356,003
GRAND TOTALS	YEAR 1	£2,244,546		3 YEARS	£6,777,360	

Budget allocated by project components

PROJECT COMPONENT	Salaries			Approximate total costs		
	MIT TOTAL	CU TOTAL	UCL TOTAL	MIT TOTAL	CU TOTAL	UCL TOTAL
Administrative	\$235,110	£118,488		\$641,850	£240,531	
Research: Value Chain	\$452,994	£49,719		\$1,236,674	£100,930	
Research: New Architecture	\$366,759	£696,066	£205,000	\$1,001,252	£1,413,014	£309,550
Research: Economic and social	Note 1	£497,190	£257,000		£1,009,296	£388,070
Research: Policy	\$350,171		£257,000	\$955,967		£388,070
Commercial Enterprise	Note 2					
Educative			£152,941			£269,000
Public Agency	Note 2					
Knowledge Exchange	\$150,997		£10,000	\$412,222		
Assessment	Note 3					
TOTAL	\$1,556,031 £972,519	£1,361,463	£881,941	\$4,247,965 £2,654,978	£2,763,770	£1,354,690

Note 1: Funding for Kaashoek provided through another contract 2.73
 Funding for Clark included in Administrative 2.03
 1.51

Note 2: Support for these activities drawn from staff assigned to research and knowledge exchange

Note 3: Support for these activities drawn from staff assigned to administrative

Note 4: additional CU support for roadmapping drawn from staff assigned to research

10. Staffing & Logistics

This program will be guided and overseen by Jon Crowcroft from Cambridge and David Clark from MIT. Clark will allocate 50% of his time to this project, as the budget reflects. Crowcroft and Clark will be responsible for setting and review of research priorities, allocation of funds, responsiveness of the program to CMI objectives, and other matters that relate to overall success.

To carry out this responsibility, a management advisory board will be established, composed of a small set of research leaders from the academic institutions and major industrial partners. This board will offer advice on issues arising, and act to confirm the directions of the project.

A project of this magnitude will require a Director, who will take overall responsibility for management, KIC activities, assessment, and organization of the management advisory board. We have tentatively estimated that this will be a half-time task, but this may underestimate the demands. Reflecting the priority of the UK objectives for the CII, our current plan is to locate the Director in the UK.

At MIT, finance and personnel administration will be carried out using an existing laboratory or research center, such as the Laboratory for Computer Science or the Media Laboratory.

As the individual project descriptions indicate, the CII will make extensive use of staff, as well as students and faculty. Research will be carried out using a mixture of PhD level students and post-doc staff, the mix being determined opportunistically by the availability of suitable participants of each sort.

Staff within CII will have a joint responsibility to carry out their specific lines of research and to facilitate the other parts of the CII, in particular the KIC components.

11. Appendices

Letters of support are pending, and will be supplied as a supplement to this proposal.

A list of suggested reviewers will also be submitted as a supplement.

Brief biographies for key participants

Ross Anderson

Ross Anderson leads the security group at the Computer Laboratory, Cambridge University, where he is Reader in Security Engineering. He is a Fellow of both the Institution of Electrical Engineers and the Institute of Mathematics and its Applications. He is the author of the definitive textbook 'Security Engineering -- A Guide to Building Dependable Distributed Systems'.

One of his more influential papers was 'The Eternity Service' – a filestore distributed over the whole internet so as to make it highly resistant to censorship and sabotage. This provided the key idea for the development of peer-to-peer systems such as freenet, gnutella and mojonation.

He initiated the study of the economics of information security systems. Many security failures occur not so much because the designers used the wrong type of encryption, but because the people capable of protecting a system were not liable for the full costs of failure. An analysis of incentives is at least as important in system design as an analysis of the ciphers or protocols in use.

He also has well known publications on emission security, on techniques for removing copyright marks from digital media, on the tamper resistance of smartcards, and the robustness of cryptographic protocols. He was a coauthor of Serpent, which was a finalist in the competition to find an Advanced Encryption Standard. Other papers document the failure modes of a number of real world systems including automatic teller machines, prepayment electricity meters, and goods vehicle tachographs.

The unifying theme of all these projects is to provide robust control mechanisms for tomorrow's heterogeneous distributed systems. Other projects examine how to apply the lessons learned to new applications such as electronic publishing and medical information systems, and how such systems can uphold users' rights by supporting safety, privacy and consumer protection.

Ross has served on the programme committees of many of the security and cryptology conferences, three of which he founded – 'Economics and Information Security', 'Fast Software Encryption' and 'Information Hiding'. He also chairs the Foundation for Information Policy Research.

Tilman Börgers

1. Qualifications

- Habilitation, Universität Basel, 1993.
- PhD in Economics, London School of Economics, 1987.
- Diplom Volkswirt, Universität Köln, 1983.

2. Current Positions:

- Director, ESRC Centre for Economic Learning and Social Evolution (since 2002).
- Professor of Economics, University College London (since 1996).

3. Editorial Positions

- *Review of Economic Studies* (1992 – 1994: Assistant Editor, 1994 – 1995: Managing Editor, 1995 – 2001: Member of the Editorial Board).
- *Ricerche Economiche* (1996 - 2001: Associate Editor).
- *Mathematical Social Sciences* (1998 - 2003: Member of the Editorial Board).

4. Selected Publications

- Auction Theory for Auction Design, in: Maarten Janssen (editor), *Auctions Vs. Beauty Contests*, Cambridge: Cambridge University Press, forthcoming (with Eric van Damme).
- Awarding Telecom Licenses: The Recent European Experience, *Economic Policy* 36 (2003), 216-268 (with Christian Dustmann).
- Rationalizing the UMTS Spectrum Bids: The Case of the UK Auction, *Ifo Studien* 48(2002), 77 - 109 (with Christian Dustmann).
- Naïve Reinforcement Learning With Endogenous Aspirations, *International Economic Review* 41 (2000), 921-950 (with Rajiv Sarin).
- Is Internet Voting a Good Thing?, *Journal of Institutional and Theoretical Economics* 156 (2000), 531-547.
- Learning Through Reinforcement and Replicator Dynamics, *Journal of Economic Theory* 77 (1997), 1-14 (with Rajiv Sarin).
- Pure Strategy Dominance, *Econometrica* 61 (1993), 423 – 430.

5. Research Grants: Multiple research grants from the ESRC.

6. PhD students: Supervised seven PhD students (three ongoing).

7. Consulting: Auction advice for government and private clients.

Bob Briscoe

Bob Briscoe joined BT in 1980 and now leads the Edge Lab, one of the Research Labs of BTexact Technologies. In the late-1980s he managed the transition to IP of many of BT's R&D networks and systems. In the mid-1990s he represented BT on the HTTP working group of the IETF and in the ANSA distributed systems research consortium, which led to the creation of the OMG and CORBA. In 2000 he initiated and was technical director of the Market Managed Multi-service Internet (M3I) consortium, a successful European collaborative project investigating the feasibility and user acceptability of controlling Internet quality on fast time-scales through pricing. His published research, standards contributions and patent filings are in the fields of loosely coupled distributed systems, scalable network charging and security solutions (esp. multicast), managing fixed and wireless network loading using pricing and on the structure of communications markets. He is also studying part-time for a PhD at University College London.

Selected publications

- Bob Briscoe, Vasilios Darlagiannis, Oliver Heckman, Huw Oliver, Vasilios Siris, David Songhurst, and Burkhard Stiller. A market managed multi-service internet (M3I). *Computer Communications*, 26(4):404-414, February 2003.
- Vasilios A. Siris, Bob Briscoe, and Dave Songhurst. Economic models for resource control in wireless networks. In *Proc. 13th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC 2002)*, September 2002. IEEE.
- Bob Briscoe (Ed.). M3I architecture. Deliverable 2, M3I Eu Vth Framework Project IST-1999-11429, URL: <http://www.m3i.org/>, July 2000.
- Mike Rizzo, Bob Briscoe, Jérôme Tassel, and Kostas Damianakis. A dynamic pricing framework to support a scalable, usage-based charging model for packet-switched networks. In *Proc. Int'l W'kshp on Active Networks (IWAN'99)*, volume 1653, February 1999. Springer LNCS.
- Bob Briscoe, Mike Rizzo, Jérôme Tassel, and Konstantinos Damianakis. Lightweight, end to end, usage-based charging for packet networks. In *Proc. IEEE Openarch 2000*, pages 77-87, March 2000.
- Bob Briscoe. The direction of value flow in multi-service connectionless networks. In *Proc. International Conference on Telecommunications and E-Commerce (ICTEC'99)*, October 1999.

David Clark

David Clark is a Senior Research Scientist at the MIT Laboratory for Computer Science, where he has worked since receiving his Ph.D. there in 1973. Since the mid 70s, Dr. Clark has been leading the development of the Internet; from 1981-1989 he acted as Chief Protocol Architect in this development, and chaired the Internet Activities Board. More recent projects include extensions to the Internet to support real-time traffic, explicit allocation of service, pricing and related economic issues, and policy issues surrounding the Internet, such as local loop deployment. He has also worked on computer and communications security.

In addition to his appointment in the Laboratory for Computer Science, Dr. Clark oversees the Internet and Telecomms Convergence Consortium at MIT. This consortium examines the broader context of the Internet—economics, societal impact and policy. The goal of this interdisciplinary consortium is to shape technological innovation and business planning by articulating this larger context for the Internet.

Dr. Clark is chairman of the Computer Science and Telecommunications Board of the National Research Council, and has contributed to a number of studies on the societal and policy impact of computer communications. He is a member of the National Academy of Engineering and a Fellow of the ACM and the IEEE.

Recent and significant publications:

David D. Clark, Craig Partridge, and J. Christopher Ramming, "A Knowledge Plane for the Internet", submitted for publication.

Clark, D., J Wroclawski, K. Sollins, R. Braden. "Tussle in Cyberspace: Defining Tomorrow's Internet", Proceedings of the ACM SigComm 2002 Conference, Pittsburg, PA August, 2002, Computer Communications Review, vol. 32, num. 4, Oct 2002 <http://www.acm.org/sigs/sigcomm/sigcomm2002/papers/tussle.pdf>

Marjory S. Blumenthal and David D. Clark, "Rethinking the design of the Internet: The end to end arguments vs. the brave new world". *ACM Transactions on Internet Technology*. August 2001, Vol. 1, No. 1

S. Shenker, D. Clark, D. Estrin, S. Herzog. Pricing in computer networks: Reshaping the research agenda. *Telecommunications Policy*, Vol 20(3). 1996

D. Clark. Implications of Local Loop Technology for Future Industry Structure. In S. Gillett and I. Vogelsang, (Eds.) *Competition, Regulation, and Convergence: Current Trends in Telecommunications Policy Research*. Lawrence Erlbaum Associates, 1999 pp.283-296

Clark, D., and Tennenhouse, D., "Architectural Considerations for a New Generation of Protocols", *SigComm Symposium*, ACM, September 1990.

Clark, D., "The Design Philosophy of the DARPA Internet Protocols", *SIGCOMM Symposium*, ACM, August 1988, pp. 106-114.

Saltzer, J., Reed, D., and Clark, D.D., "End-to-End Arguments in System Design", *ACM Transactions on Computer Systems*, Vol. 2, No. 4, November 1984, pp. 277-288.

Ingemar J. Cox

Ingemar J. Cox received his B.Sc. from University College London and Ph.D. from Oxford University. He was a member of the Technical Staff at AT&T Bell Labs at Murray Hill from 1984 until 1989 where his research interests were focused on mobile robots. In 1989 he joined NEC Research Institute in Princeton, NJ as a senior research scientist in the computer science division. At NEC, his research shifted to problems in computer vision and he was responsible for creating the computer vision group at NECI. He has worked on problems to do with stereo and motion correspondence and multimedia issues of image database retrieval and watermarking. In 1999, he was awarded the IEEE Signal Processing Society Best Paper Award (Image and Multidimensional Signal Processing Area) for a paper he co-authored on digital watermarking. From 1997-1999, he served as Chief Technical Officer of Signafy, Inc, a subsidiary of NEC responsible for the commercialization of watermarking. Between 1996 and 1999, he led the design of NEC's watermarking proposal to the Copy Protection Technical Working Group of the DVD Consortium and later collaborated with IBM in developing the technology behind the joint "Galaxy" proposal supported by Hitachi, IBM, NEC, Pioneer and Sony.

In 1999, he returned to NEC Research Institute as a Research Fellow. He is currently Professor and Chair of Telecommunications in the Departments of Electronic Engineering and Computer Science at University College London.

He is a senior member of the IEEE and on the editorial board of the Int. Journal of Autonomous Robots and Pattern Analysis and Applications Journal. He is co-author of a book entitled *Digital Watermarking* and the co-editor of two books, *Autonomous Robots Vehicles* and *Partitioning Data Sets: With Applications to Psychology, Computer Vision and Target Tracking*. He is inventor/co-inventor on 22 patents and author/co-author of over 100 papers.

Jon Crowcroft

Jon Crowcroft is the Marconi Professor of Networked Systems in the Computer Laboratory, of the University of Cambridge; prior to that he was professor of networked systems at UCL in the Computer Science Department.

He is a Fellow of the ACM, a Fellow of the British Computer Society and a Fellow of the IEE and a Fellow of the Royal Academy of Engineering, as well as a senior member of the IEEE. He was a member of the IAB, and was general chair for the ACM SIGCOMM 95-99. He is on the editorial team for COMNET, and on the program committee for ACM SIGCOMM and IEEE Infocomm. He has been editor of IEEE/ACM Joint Transactions on Networks, editor of IEEE Networks magazine, and served on IEE, IEEE and ACM Electronic Publication Committees. He has published 5 books, the latest on the Linux TCP/IP Implementation (Wiley 2001).

His industrial experience includes work for Bloomsbury Computer Consortium and a sabbatical at Hewlett Packard Research Labs Bristol. He has also been on the technical advisory board (TAB) for startups such as Ensim, Orchestream, Bandwiz, Nexthop, Interprovider & Endace and consulted for Reuters, BBC, Nortel, Cisco & Oftel. He currently serves on the TAB for Microsoft Research.

Home Page: <http://www.cl.cam.ac.uk/users/jac22/>

Recent Publications:

- *Preferential Treatment of Acknowledgement Packets in a Differential Services Network*, Dina Papagiannaki, Patrick Thiran, Jon Crowcroft, Christophe Diot, Proc of IEEE IWQoS 2001, Germany.
- *Lower than Best Effort: a Design and Implementation*, Ken Carlberg, Panos Gevros, Jon Crowcroft, in ACM SIGCOMM First Workshop in Data Communications in Latin America and the Caribbean, 2001.
- *Congestion Control Mechanisms and the Best Effort Service Model*, Panos Gevros, Jon Crowcroft, Peter Kirstein, Saleem Bhatti, in the IEEE Network Special Issue on the Control of Best Effort Traffic, May/June '01.
- *Router Level Filtering for Receiver Interest Delivery*, Manuel Oliveira, Jon Crowcroft, Christophe Diot, in Proc. of NGC2000, ACM Press.
- *'An Architecture for Application Layer Routing'*, Ghosh, Fry & Crowcroft, Yasuda, H. (Ed), Active Networks, LNCS 1942, Springer, pp 71-86. ISBN 3-540-41179-8 Springer-Verlag
- *Herding Cats: Modelling the Internet*, Jon Crowcroft, Royal Society talk, Dec 1999. pub in August 2000 issue of the Philosophical Transactions of the Royal Society, vol 358, number 1773
- *Watercasting: Distributed Watermarking of Multicast Multimedia*. Ian Brown, Jon Crowcroft, Colin Perlins, appeared in NGC1999. Also an International and US Patent

Charles H. Fine

Charles H Fine is the Chrysler LFM Professor of Management, MIT Sloan School of Management, 77 Massachusetts Avenue, Room E53-393, Cambridge, MA 02139; **tel:** (617) 253-3632, **fax:** 617-253-6720, **email:** charley@mit.edu. He holds an M.S. in Operations Research and a PhD in Business Administration, both from Stanford (1981 and 1983 respectively).

Books Published:

- *The U.S. Automobile Manufacturing Industry*, U.S. Department of Commerce, Office of Technology Policy, (joint with R. St. Clair), 1996.
- *Clockspeed: Winning Industry Control in the Age of Temporary Advantage*, Perseus Books, 1998.

Selected Articles (authored and co-authored):

- "Optimal Investment in Product-Flexible Manufacturing Capacity," *Management Science*, 1990.
- "Interdisciplinary Product Design Education," *IEEE Transactions on Engineering Mgmt*, 1990.
- "Manufacturing Equipment Development and Sourcing," *Toray Business Review*, Nov 1993.
- "Technology Supply Chains and Business Cyclicalilty," *Toray Business Review*, Nov 1994.
- "An Empirical Study of Manufacturing Flexibility in PCB Assembly" *Operations Research*, 1996.
- "Time vs. Market Orientation in Product Concept Development," *Management Science*, April 1997
- "Upstream Volatility in the Supply Chain," *Production and Operations Management*, Fall 2000.
- "Is the Make-Buy Decision Process a Core Competence?" in Moreno Muffatto and Kulwant Pawar (eds.), *Logistics in the Information Age*, Servizi Grafici Editoriali, Padova, Italy, 1999.
- "Biography of a Killer Technology: Optoelectronics Drives Industrial Growth with the Speed of Light" (available at <http://web.mit.edu/ctpid/www/people/Fine.html>)
- "Business Cycles and Productivity in Capital Equipment Supply Chains," in S. Tayur, M. Magazine and R. Ganeshan (eds.) *Quantitative Models for Supply Chain Management*, Boston: Kluwer Press, 1998.
- "Rapid Response Capability in Value-Chain Design," *Sloan Management Review*, Winter 2002.

Mark Handley

Mark Handley will be Professor of Networked Systems in the Department of Computer Science at University College London from July 2003. Prior to this he held the position of Senior Scientist at the International Computer Science Institute in Berkeley, California.

He has been an important participant in the area of Internet standards. The principle Internet standards organization is the Internet Engineering Task Force (IETF). Mark is a member of the Internet Architecture Board (IAB), which is the committee of the IETF that oversees the Internet standards process, including technical oversight of the architecture of the protocols being developed and the publication of the standards themselves. Prof Handley was chair of the IETF working group on Multimedia Session Control, and was chair of the research group on Reliable Multicast Transport. He is also a member of both the IETF's Transport Area Directorate and Routing Area Directorate.

He has served on the programme committee for many conferences and workshops, including ACM Sigcomm, IEEE International Conference on Network Protocols, ACM Future Directions in Network Architecture, and the ACM Networked Group Communication workshop. His research spans a broad range, from multimedia systems and applications, through network congestion control, to routing protocols and network architecture.

Home page: <http://www.icir.org/mjh> or <http://www.cs.ucl.ac.uk/staff/M.Handley>

Selected Publications:

- M. Handley, O. Hodson, E. Kohler, "XORP: An Open Platform for Network Research", ACM Computer Communications Review, Jan 2003, ACM, New York, NY, 2003.
- B. Braden, T. Faber, M. Handley, "From Protocol Stack to Protocol Heap - Role-Based Architecture", ACM Computer Communications Review, Jan 2003, ACM, New York, NY, 2003.
- M. Handley, S. Floyd, J. Padhye, J. Widmer, "TCP Friendly Rate Control (TFRC): Protocol Specification", IETF RFC 3448, January 2003.
- M. Handley, C. Kreibich, V. Paxson, "Network Intrusion Detection: Evasion, Traffic Normalization, and End-to-End Protocol Semantics", Proc. USENIX Security Symposium 2001, Washington DC, August 2001.
- R. Rejaie, M. Handley, D. Estrin, "Quality Adaptation for Congestion Controlled Playback Video over the Internet", Proc. ACM SIGCOMM '99, Cambridge, MA, USA, September 1999
- M. Handley, H. Schulzrinne, E. Schooler, J. Rosenberg, "SIP: Session Initiation Protocol", March 1999, IETF RFC 2543.

Frank Kelly

Frank Kelly is Professor of the Mathematics of Systems in the University of Cambridge.

His main research interests are in random processes, networks and optimization, and especially in applications to the design and control of communication networks. He and colleagues at Cambridge developed, jointly with British Telecom, the routing scheme implemented in BT's main digital telephone network. His current research is directed at understanding methods of self-regulation of the Internet.

Frank Kelly has been awarded the Guy Medal in Silver of the Royal Statistical Society, the Lanchester Prize of INFORMS, and the Naylor Prize of the London Mathematical Society. He is a Fellow of the Royal Society. He has chaired the Advisory Board of NRIC (the Mathematics Enrichment Project), and the Management Committee of the Isaac Newton Institute for Mathematical Sciences. He currently serves on the Scientific Council of EURANDOM and the Conseil Scientifique of France Telecom.

Andrew Lippman

Andrew Lippman is a Senior Research Scientist at MIT and director of the Digital Life Research Program. His research group is the Media and Networks group at the MIT Media Laboratory, and he recently began the Viral Communications Program, in concert with David Reed. In the past, his research has focused on the enabling technologies of human communications, beginning with early research on video displays, interaction systems, and ultimately visual communications. He was one of the first members of the MPEG group, and he developed scalable video techniques for low-rate and high-definition television. He has taught at MIT for 25 years in positions ranging from professor to his current position, and his courses included digital video and the technology and policy of media. He received his BS and MS degrees from MIT and a PhD from EPFL (Lausanne.)

Home Page: www.media.mit.edu/~lip

Selected Publications:

"Efficient Collaborative (Viral) Communication in OFDM Based WLANs," Bletsas, A. and Lippman, A. Proceedings of the International Symposium on Advanced Radio Technologies," March 2003.

"A Multicast Transport Protocol for Cooperative Media Distribution," Vyzovitis, D. and Lippman, A. Extended Abstract, Proc. 10th IEEE International Conference on Network Protocols (ICNP 2002), Paris, FR, November 2002

"Evaluation of Kalman Filtering for Network Time Keeping, " Bletsas, A. and Lippman, A. MIT Media Laboratory, November 2002.

"Cooperative Multicast for Scalable Real-Time Media Distribution," Vyzovitis, D. and Lippman, A. MIT Media Laboratory, September 2002.

"Statistical Models of Video Structure for Content Analysis and Characterization," Vasconcelos, N., and Lippman, A. IEEE Transactions on Image Processing, vol. 9, no. 1, January 2000.

"The Media Bank: A Fault Tolerant Heterogeneous Distributed Pull Based Architecture for the Delivery of Object Oriented Multimedia Across Packet Switched Networks," Kermode, R., Holtzman, H., Lippman, A., Proceedings of the 7th International Workshop on Packet Video, Brisbane, Australia, March 1996.

"Digital Television: A Perspective", Netravali, A., and Lippman, A.B., Proceedings of the IEEE, Vol. 83, No. 6, June 1995.

"Scalable Open-Architecture Television", Bove, V.M., Lippman A.B., SMPTE Journal, January 1992.

"Feature Sets for Interactive Images", Lippman, A.B., CACM, Vol. 34, No. 4, April 1991.

Helen Margetts

Education: PhD in Government, LSE, 1996; MSc Politics, LSE, 1990; BSc Mathematics, University of Bristol, 1983.

Positions held: Professor of Political Science and Director, School of Public Policy at University College London (from 2001, Reader from 1999), Senior Lecturer in Politics, Department of Politics and Sociology, Birkbeck College (from 1994) and Course Director, MSc Public Policy and Management; previously Research Officer in Department of Government, LSE (from 1992); Senior Systems Analyst, Amoco (from 1987); Computer Programmer/Analyst, Rank Xerox (from 1984).

Policy Experience: Consultant to: the National Audit Office (1992-2003); OECD (1996-9); the Government Office for London (1997-8); the Independent Commission on the Voting System (the 'Jenkins' Commission) (1998); Lewisham Council (1999).

Current Major Research Council Grants: *Public-private Partnerships and Policy Transfer in Central Government ICT Systems* (funded by the UK Economic and Social Research Council, 2001-2003)

Recent Relevant Publications

- (with Patrick Dunleavy) *Cultural Barriers to e-government*, academic article for the report Better Public Services Through e-government (London: National Audit Office, 2002, HC 704-III).
- (with Patrick Dunleavy) *Government on the Web II* (London: National Audit Office, 2002) HC764 –(see www.governmentontheweb.org).
- (with Keith Dowding, James Hughes) *Challenges to Democracy: Ideas, Involvement and Institutions* (London: Macmillan, 2001)
- (with Patrick Dunleavy, Stuart Weir and Trevor Smith) *Voices of the People* (London, Politicos, 2001).
- 'Political Participation and Protest' in P.Dunleavy, A. Gamble, I. Holliday and G.Peele (eds.) *Developments in British Politics 6* (London: Macmillan, 2000) pp. 185-202.
- (with Patrick Dunleavy) *Government on the Web* (London: National Audit Office, 1999). –(see www.governmentontheweb.org).
- (with Klaus Goetz) The Solitary Centre: The Core Executive in Central and Eastern Europe, *Governance*, vol 12 no 4, 1999, pp.425-454
- *Information Technology in Government: Britain and America* (London: Routledge, 1999) pp. 1-208.
- 'Computerising the Tools of Government' I. Snellen and W de Donk (eds. 1998), *Public Administration in an Information Age* (Amsterdam: IOS Press) pp. 441-457.
- 'The National Performance Review: A New Humanist Public Management' in A. Massey (ed.) *Globalization and Marketization of Government Services* (London: Macmillan, 1997) pp 47-70.

Richard V Penty

Professor Richard V. Penty studied for his doctoral research in the Cambridge University Engineering Department on optical fibre devices for signal processing applications, receiving his Ph.D. in 1989. He was then a SERC Information Technology Fellow, again at Cambridge, working on all-optical non-linearities in waveguide devices. In 1990 he was appointed to a Lectureship in the School of Physics at the University of Bath and in January 1996 took up a Lectureship in the Department of Electrical and Electronic Engineering, University of Bristol, subsequently being promoted to a Professorship in 1998. He joined Cambridge University Engineering Department in October 2001, becoming Professor of Photonics in October 2002.

His research interests include high-speed optical communications systems, wavelength conversion and WDM networks, optical amplifiers, optical non-linearities for switching applications and high power semiconductor lasers. He has been an author of in excess of 250 refereed journal and conference papers, including over 25 invited papers.

Richard Steinberg

Richard Steinberg is Lecturer in Operations Management in the University of Cambridge. He received his B.A. (Mathematics) from Reed College, his M.Math. and Ph.D. (Combinatorics and Optimization) from the University of Waterloo, and his M.B.A. from the University of Chicago. His main research interests are in telecommunications economics, computer networking, and combinatorial auctions. He has also worked in the marketing/manufacturing interface, cost and revenue allocation, and transportation networks.

His current research is directed at understanding pricing issues for the Internet, including, pricing for congestion control, ISP pricing to end users for profit maximisation, and how the owner of a dedicated communication network should price bandwidth to businesses that host Internet applications.

Dr Steinberg has held full-time appointments at the University of Chicago and Columbia University, and visiting appointments at MIT in the Operations Research Center, CORE (Center for Operations Research and Econometrics), and the London School of Economics. He has also worked at Salomon Brothers, AT&T Bell Laboratories, and GTE Laboratories. He is Senior Editor on *Manufacturing & Service Operations Management*, and also serves on the editorial boards of the *International Journal of Production Economics* and *IIE Transactions*.

Ian White

Ian White is currently the Van Eck Professor of Engineering at Cambridge University and a Fellow of Jesus College, Cambridge. Between the years of 1996 – 2001, Ian was Professor of Optical Communications, University of Bristol, and latterly Head of the Department of Electrical and Electronic Engineering. Previously he also held the position of Professor of Physics at the University of Bath. His current research interests have been in the area of high speed communication systems, local area networks using optical links, RF over fibre, nonlinear optics and photonic components.

Ian has written over 400 publications, and holds 20 patents. Ian won the Ambrose Fleming Premium in 1984 and the Blumlein-Browne-Willans award from the Institution of Electrical Engineers in 1993, for journal papers. Outside of Cambridge University, Ian is a Member of the IEE Photonics Professional Network Committee and a Member of the Strategic Assessment Team of EPSRC for Computing & IT. In addition, Ian is Editor of Optical and Quantum Electronics and an Honorary Editor of Electronics Letters, a Fellow of the Institution of Electrical Engineers and Senior Member of the Institution of Electrical and Electronic Engineers. Conference Committee Membership includes ECOC, Globecomm, ECIO, Photonics West and CLEO-Europe.

Selected Publications:

- "Mode-hop-free, singlemode operation of 2D lattice distributed reflector laser under 2.5Gbit/s modulation", AB Massara, LJ Sargent, PJ Heard, RV Penty and IH White, *Electron. Lett.*, Vol.36, No.2, pp.141-142, 2000
- "Demonstration of high-speed optical packet routing using vertical coupler crosspoint space switch array", S Yu, M Owen, R Varrazza, RV Penty and IH White, *Electron. Lett.*, Vol.36, No.6, pp.556-558, 2000
- "1.3 micron quantum well InGaAsP, AlGaInAs and InGaAsN laser material gain: a theoretical study", JCL Yong, JM Rorison and IH White, *Journal of Quantum Electronics*, Vol. 38, pp.1553-1564, 2002
- "0.6 Tbit/s/km multimode fibre feasibility experiment using 40 channel DWDM over quadrature-subcarrier transmission", P Kourtessis, T Quinlan, E Rochat, SD Walker, M Webster, IH White, RV Penty, and MC Parker, *Electronics Letters*, Volume 38, pp. 813-815, 2002
- "Wavelength switching components for future photonic networks", IH White, RV Penty, M Webster, YJ Chai, A Wonfor and S Shahkooh, *IEEE Communications Magazine*, Vol. 40, pp. 74-81, 2002 (invited paper)
- "Compact integrated silica wavelength filters", CN Morgan, S Yu, RV Penty, and IH White, *IEEE Photonics Technology Letters*, Vol. 14, pp. 1303-1305, 2002

BIBLIOGRAPHY

The following are a selection of articles and books which are related to the research component of this proposal.

"Optimal Investment in Product-Flexible Manufacturing Capacity," *Management Science*, Vol. 36, No. 4, 1990, 449-466, (joint with R. Freund).

- "An Empirical Study of Manufacturing Flexibility in Printed Circuit Board Assembly" (joint w/ F. Suarez and M. Cusumano), *Operations Research*, Vol. 44, No. 1, pp. 223-240.
- "Time vs. Market Orientation in Product Concept Development: Empirically-based Theory Generation" *Management Science*, April 1997. (joint with Gary Burchill)
- "Biography of a Killer Technology: Optoelectronics Drives Industrial Growth with the Speed of Light" with L.K. Kimerling, published in 1997 by the Optoelectronics Industry Development Association, 2010 Mass Ave, NW, Suite 200, Wash. DC 20036-1023. (Also available at <http://web.mit.edu/ctpid/www/people/Fine.html>)
- "Business Cycles and Productivity in Capital Equipment Supply Chains," in S. Tayur, M. Magazine and R. Ganeshan (eds.) *Quantitative Models for Supply Chain Management*, Boston: Kluwer Press, 1998. (joint with Edward Anderson).
- "Is the Make-Buy Decision Process a Core Competence?" in Moreno Muffatto and Kulwant Pawar (eds.), *Logistics in the Information Age*, Servizi Grafici Editoriali, Padova, Italy, 1999, pp. 31-63. (joint w/ D. Whitney).
- "Upstream Volatility in the Supply Chain: The Machine Tool Industry as a Case Study," *Production and Operations Management*, Vol. 9, No. 3, Fall 2000, pp. 239-261. (joint w/ E. Anderson, G. Parker).
- "Rapid Response Capability in Value-Chain Design," *Sloan Management Review*, Vol. 43, No. 2, Winter 2002, pp. 69-75. (joint w/ R.Vardan, R. Pethick, J. El Hout).
- A network flow model for mixtures of file transfers and streaming traffic. Peter Key, Laurent Massoulié, Alan Bain and Frank Kelly International Teletraffic Congress 18, Berlin, 2003.
- Modelling incentives for collaboration in mobile ad hoc networks. Jon Crowcroft, Richard Gibbens, Frank Kelly and Sven Östring WiOpt'03 Workshop.
- Fairness and stability of end-to-end congestion control. Frank Kelly *European Journal of Control* 9 (2003) 149-165.

- A contract and balancing mechanism for sharing capacity in a communication network. Edward Anderson, Frank Kelly and Richard Steinberg.
- Mathematical modelling of the Internet. In "Mathematics Unlimited - 2001 and Beyond" (Editors B. Engquist and W. Schmid). Springer-Verlag, Berlin, 2001. 685-702.
- Distributed admission control. F.P. Kelly, P.B. Key and S. Zachary IEEE Journal on Selected Areas in Communications 18 (2000) 2617-2628.
- Tilman Börgers and Eric van Damme, Auction Theory for Auction Design, in: Maarten Janssen (editor), *Auctions Vs. Beauty Contests*, Cambridge: Cambridge University Press, forthcoming.
- Tilman Börgers and Christian Dustmann, Awarding Telecom Licenses: The Recent European Experience, *Economic Policy* 36 (2003), 216-268.
- Tilman Börgers and Christian Dustmann, Rationalizing the UMTS Spectrum Bids: The Case of the UK Auction, *Ifo Studien* 48(2002), 77 - 109.
- Tilman Börgers and Christian Dustmann, The British UMTS Auction: A Response to Klemperer and Schmidt, *Ifo Studien* 48 (2002), 121 - 122.
- Margetts, H. and Dunleavy, P. *Cultural Barriers to e-government*, academic article for the report Better Public Services Through e-government (London: National Audit Office, 2002, HC 704-III) - download at www.governmentontheweb.org.
- Dunleavy, P. and Margetts, H. *Government on the Web II* (London: National Audit Office, 2002) HC764 - download at www.governmentontheweb.org.
- Dunleavy, P. and Margetts, H. *Government on the Web* (London: National Audit Office, 1999) download at www.governmentontheweb.org.
- "Cruel, Mean or Lavish?: Economic Analysis, Price Discrimination and Digital Intellectual Property", James Boyle, *Vanderbilt Law Review*, 53,6,2000
- "DRM {and or v's} the Law", Pamela Samuelson, *Communications of the ACM*, 46, 4, 2003
- "The darknet and the future of content distribution", P. Biddle, P. England, M. Peinado and B. Willman, *Proc ACM Workshop on DRM*, 2002
- Pricing Internet Service. R. Steinberg. In: A.K. Chakravarty and J. Eliashberg, eds., *Managing Business Interfaces: Marketing, Engineering, and Manufacturing Perspectives*, Dordrecht: Kluwer Academic Publishers/ISQM, 2003.
- Congestion Pricing and User Adaptation. A. Ganesh, K. Lavens, and R. Steinberg. *IEEE INFOCOM* (2001), 959-965.

- Internet Service Classes under Competition'. R. Gibbens, R. Mason, and R. Steinberg. *IEEE Journal on Selected Areas in Communications* 18, 12 (2000), 2490 -2498.
- Ross J. Anderson, The Eternity Service, Pagocrypt 96, 1996.
- Ross J. Anderson, Vaclav Matyas, and Fabien A. P. Petitcolas, The Eternal Resource Locator: An Alternative Means of Establishing Trust on the World Wide Web, 3rd USENIX Workshop on Electronic Commerce, 1998, Boston.
- JianXin Yan, Stephen D. Early, Ross J. Anderson, The XenoService - A Distributed Defeat for Distributed Denial of Service, Information Survivability Workshop, Oct 2000, Boston.
- Richard Clayton, George Danezis, Markus G. Kuhn: Real World Patterns of Failure in Anonymity Systems, IHW 2001, Pittsburgh, April 25-17, 2001.
- Frank Dabek, M. Frans Kaashoek, David Karger, Robert Morris, and Ion Stoica, Wide-area cooperative storage with CFS, ACM SOSP 2001, Banff, October 2001.
- Ion Stoica, Robert Morris, David Karger, M. Frans Kaashoek, and Hari Balakrishnan, Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications, ACM SIGCOMM 2001, San Deigo, CA, August 2001
- Emil Sit and Robert Morris, Security Considerations for Peer-to-Peer Lookup Systems, the 1st International Workshop on Peer-to-Peer Systems, MIT Cambridge, March 2002.
- Ross J. Anderson, Why Information Security is Hard -- An Economic Perspective, in Proceedings of the Seventeenth Computer Security Applications Conference (IEEE Computer Society Press, 2001), pp 358--365; also given as a distinguished lecture at the Symposium on Operating Systems Principles, Banff, October 2001