

Organic networks

A Lippman and A Pentland

The topic of organic networks derives from the confluence of two distinct bodies of research that have been proceeding independently in the Media Lab for the past several years — ‘Viral Networks’, which focuses on the enabling technology underpinning end-to-end, grassroots communications systems, and ‘Influence Networks’, which encompasses ways that both first-world and third-world societies bend the technology of easy connectivity to suit their own economic, cultural and social interests. While the general method of research in the Media Laboratory (semi-autonomous groups following largely independent research tracks) implies that these two themes are somewhat segregated, their intersection carries implications and lessons in and of itself that are too strong to be ignored — hence this co-ordinated set of papers.

1. **Viral Networks — innovation follows architecture**

Viral communications derives directly from the end-to-end principle on which the Internet is based — the intelligence is in the end nodes, the network itself maintaining as little state as possible. As defined, the Internet is not optimised for any particular use or communications requirement, it merely forwards packets on a ‘best-effort’ basis. This principle kept the Internet open to innovation by reducing the architectural impact and cost or risk imposed on the development of any new application; applications could start small and propagate by pure popularity, no core change was needed, and no innovation had to be debugged well enough to ensure that it had no adverse impact network integrity.

One reason that this succeeded as a principle was that there were no economic stakeholders in the early Internet. Such a widely distributed communications system was unprecedented and its design was closely held by the engineers who used it. However, the same coupling of innovation with distribution had begun in the US telephone network with the advent of the FCC ‘Carterphone Decision’ in 1968. Until that ruling, no device other than one sold by AT&T (Western Electric) could be connected to the telephone network. Carter Electronics had wanted to bridge between their private mobile network and land-lines, and AT&T prohibited it.

When the prohibition was lifted, Western Electric lost the monopoly on terminal equipment, but the net amount of innovation in telephony was vastly increased. Facsimile, modem technology, answering machines, and private PBXs all flourished.

Similarly, there was ample economic force behind the mainframe computing industry by 1978, when the Apple II was introduced, but computing was a scarce resource and few envisioned its human-oriented or humane uses. By reducing the economic barrier to innovation and by adopting a more modular, flexible architecture, computing became accessible to entrepreneurs, small businesses, and ultimately consumers. If we use attendance at SIGGRAPH, the premier venue for innovators in graphics and interactive systems, as a measure of activity in the field, in 1974, there were 600 attendees; by 1985 it had grown tenfold, and in the 1990s 100-fold. Baldwin has plotted the scale of the entire industry and it reflects similar growth.

Other validations of the end-to-end principle abound. Innovation, while not divorced from economics, can be associated closely with architecture.

The theme of viral communications codifies this in a slightly different way, and explores how the end-to-end notion can be applied to wired and wireless networks. The challenge is most evident in the domain of radio, and this is where we focus the technical energy. In particular, a viral system is one that meets three criteria:

- it scales (almost) without bound,
- it can grow incrementally,
- each new element adds technical capacity to the overall system.

This is considered a radical concept in that convention proposes that a communications system has limited capacity which is divided among participating communicating elements — additional nodes subdivide available capacity rather than increase it. This lack of scalability is perhaps the most stringent restriction on system design, regulatory practice and economic development.

A fixed and limited resource results in the notion of ‘spectrum scarcity’, which directly implies allocation or auction of exclusive rights. Further, it dictates a particular design principle for communications systems. It is evident in mass media and telecommunications practice where the intelligence, power and architecture is centrally concentrated in broadcast power allowances, telecommunications towers, and interference regulations. The obvious distributed alternatives are in many ways creative failures — Citizen’s Band radio and FRS. As Yogi Berra says: ‘They are so crowded nobody goes there anymore.’

It is important to note that the issue is ill-posed. The spectrum itself is not scarce — there is no limit on the amount of electromagnetic energy it can carry. What has been limited is our ability to effectively distinguish the various signals that may occupy it. Division of radio spectrum by space, frequency, code or time are parameters of receiver design, not physical principles of electromagnetism. Indeed, each such manner of subdivision has evolved as the technology and techniques associated with it have matured and become realisable. Time division comes from telegraphy, frequency division derives from the linear electronic systems based on the Audion Tube of 1906, space division is a by-product of microwaves developed during World War II, and code division became realisable when digital processing reached speeds needed for real-time signal analysis.

More recently, researchers are taking a more physical approach to radio signalling. MIMO radio systems use multiple antennas and high-speed signal processing to separate a particular source from other signals and noise. It has been shown that one can either conserve energy in a transmission or reuse a given frequency band in the same area in proportion to the (smaller of) the number of antennas used at the transmitter and receiver. Reception is a matter of channel estimation rather than linear detection.

These ‘intelligent’, ‘collaborative’, or ‘cognitive’ radios can quite literally grow system capacity with an increasing number of nodes (although there is no firm consensus on the rate of growth) without violating the underlying mathematical theory of communications or electromagnetic physics. The difference between these newer systems, and those already in use, lies in where the complexity is placed — increasing processing in each radio (or wired node) similarly increases system capacity. The difference is that it is now potentially economic to do processing in the radio that was not even a dream 80, 50 or even 25 years ago, when the dominant spectrum uses (broadcast radio, broadcast television, and cellular telephony) were defined.

We elucidate this by analogy to light: the historical goal of spectrum regulation and radio design has been analogous to keeping a fixed and finite region (or radio band, etc) completely

‘dark’ except for one radiator, to protect it from interference. This is borne out by repeated measures of actual spectrum use that show that there is indeed almost no energy in most of the usable electromagnetic spectrum for most of the day (with some peaks in the cell-phone bands during rush hour). The potential of these techniques is that they break the model of scaling that has limited communications system design. A scalable system, where one can freely add elements, presents new opportunities and requires new approaches to regulation. The paradigm more closely approximates (mostly) non-exclusionary regulations used on the open sea than for real estate development — the sea scales without central monitoring as long as each user obeys reasonably simple rules of the road.

Nevertheless, there remain technical obstacles to a useful viral radio system. Most important, our knowledge of real channel characteristics is sparse — engineers often use fixed power-law functions of energy as a design parameter even when the system is to be used in a city or a building, where the channel changes continuously (but not arbitrarily) when radios or objects in the space move. Building a radio system that adapts in real time to channel parameters remains a challenge.

Similarly, there are unexplored issues of where the intelligence can effectively be placed. A multi-hop, wireless, *ad hoc* network can conserve power and it can scale, but there are delays imposed by each relay. An important question is whether we can eliminate this delay by processing the RF directly, using the derived physical parameters of propagation rather than by modelling the radio system as a mesh of ‘wireless wires’.

Finally, there are issues of security and robustness. It is reasonable to expect that a distributed system is inherently robust and can be secure (some of the founding goals of the Internet), yet technology allows modern forms of distributed attack that revive the debate. There are political as well as technical dimensions to the discussion.

We present these issues and some research progress in detail in the first paper, *Viral Communications*, by Lippman and Reed. Bletsas shows how these ideas can be applied to regions of the world that are industrialising and facing issues of initial deployment of communications almost *in vacuo*. We address some of the architectural issues of distributed networks for sensing in the paper by Paradiso et al on *Sensate Media* and by Bove and Mallett on *Collaborative Knowledge Building*.

Ultimately, if we can replicate the economic history of AT&T-post-Carterphone and the computing-post-mainframes, we can expect ‘grassroots’ and embedded communications systems that are released from the concentration of services and delivery in a small number of carriers or broadcasters. As with computing, innovation can come from the users as well as the providers, and almost anyone can be a user. This is half of the picture of ‘organic networks’.

2. Influence Networks — innovation follows society

Just as architecture defines innovation potential, the adoption of innovation is defined by society, not technology. This is the overriding premise of organic networks.

The historic approach to this issue cites the phonograph, the telephone, the automobile and early radio as examples: the first was invented as a recording medium for wills, the second as a hearing aid, the third as an ambulance, and the fourth as a point-to-point means for communicating with ships at sea. Yet each took violent turns as the technology became popularly accessible. Certainly there was individual vision behind the notion of publishing, records, and radio broadcasting, but it was the response of society as a whole that fleshed out the picture. In fact, the magic occurs when the technological development is 'in sync' with societal or cultural needs/desires. (The telegraph lay fallow until long-distance rail made it necessary for signalling and scheduling.)

More formally, an organic network is a viral system that gains widespread impact derived from local knowledge. In addition to being grassroots, it reflects the needs and interests of its local constituency and develops global, emergent behaviour. By global, we mean that it spreads in some universal manner either within a given society or quite literally, throughout the world.

A prime example is IEEE802.11, or WiFi. The underlying technical standard is not viral (the radio itself does not scale nor does each transceiver add capacity to the system). However, in impact and as a practical matter, it is both viral and organic. For example, if we hypothesise that some proportion of people who purchase an 802.11 card for their laptop or palmtop also install an access point for their home, then they are indeed adding capacity to the system by dint of the hotspot they have thus created. Further, while too many nodes will ultimately saturate (or at least slow down) the throughput of any particular access point, as a practical matter, this is not noticeable by any single user doing standard access such as the Web or e-mail. The capacity is effectively infinite in that it is more than any one user needs.

An *ad hoc*, hotspot-based WiFi network, set up by the collection of individuals willing to share backbone connectivity, is built by an implicit social contract. Negroponte has likened this to the quaint notion of window flowerboxes: they are maintained both as a personal and as a communal good by the dweller. These kinds of socially constructed networks, whether technically viral or not, but at least somewhat so in practice, define an organic network.

Short messaging is a second example. It is a by-product of the digital PCS channel that was made available for cell operators to communicate with customers and distribute advertisements. If one accidentally leaves a GSM phone on during a flight across Europe, they will land with a plethora of messages announcing and offering connectivity from each carrier encountered *en route*. However, the vagaries of call pricing and the novelty of the user community has transformed it into a separate medium.

The coupled incentives of economic development and personal expression drive the Influence Networks' research agenda. Examples and principles of these networks are described in papers contributed by Schmandt and Lakshmiopathy, Pentland, Donath, Norton et al, Selker and Goler, Cavallo, and Resnick.

Pentland explores how we can sense, model, and potentially control the networks of human influence that drive technology adoption. This technology has already given us a better understanding of 'viral marketing' efforts, workgroup organisation, and knowledge management within organisations. Donath describes how people use public signals to help them navigate and profit from these influence networks. This work helps us to better understand the phenomenon of 'on-line identity' and the diffusion of innovation within on-line communities.

There are many well-documented problems with how organisations disseminate information and how decisions are made. These problems include the basic problem of security and privacy, but also problems in the dynamics of local groupings within the network — 'groupthink', risk aversion, and office politics. Several of the papers in this section describe communications systems that incorporate knowledge about human social behaviour in order to avoid these well-documented problems. The SAVE system addresses the problems of voting systems. SimPhony allows distributed groups to be more aware of the overall group dynamics, while AntiGroupWare and Second Messenger arrange communications to avoid some of the major problems in group decision making.

Most of these tools are aimed primarily at working professionals, but the same patterns of development are reaching far beyond this audience. Perhaps most unexpected is the extension of these principles to educational domains, as explored by Resnick and Cavallo. Education is perhaps an inherently viral system in that the intelligence is explicitly at the end nodes (the learners) and the system is designed to operate on their behalf. Yet it has been approached for the past one hundred years as a centralised, large-scale infrastructure problem, with all of the attendant benefits and problems thereby implied.

3. Final lessons

We present this portfolio as much as an initial exploration as a compilation of completed results. There have been many social and technical claims made about human and technical development since the consumerisation of computing and the emergence of the World Wide Web. These systems have been presented as democratising engines, as progenitors of economic development and as new political organisations rooted in information space rather than physical location. There is an element of truth in all of these visions as well as some well-intentioned dreaming.

Our goal in presenting this research arena is to illuminate the dimensions of the issue. We see an interconnection between technical development and social direction as a feedback loop that now engages a greater population than it has in the past.

The challenge now is for us to further explore the dual processes of innovation and viral networking, in order to better understand how to design communications and decision systems that are orders of magnitude better than those that exist today.

Organic networks



Andrew Lippman's work at the Media Lab has ranged from wearable computers to global digital television. Currently, he heads the Lab's Viral Communications program and co-directs MIT's interdisciplinary Communications Futures programme. He also directs the Digital Life consortium, which works to create a networked world where communication becomes fully embedded in our daily lives.

He has written both technical and lay articles about our digital future and given over 250 presentations on the future of information and its commercial and social impact. He received both his BS and MS in electrical engineering from MIT, and his PhD from the EPFL in Lausanne, Switzerland.



Toshiba Professor of Media Arts and Sciences Alex 'Sandy' Pentland heads the Media Lab's Human Dynamics group. One of the most-cited computer scientists in the world, his work encompasses areas such as wearable computing, communications technology for developing countries, human-machine interfaces, artificial intelligence, and machine perception. He is a co-founder of the Media Laboratory's Digital Nations consortium, the LINCOS project in Costa Rica, and of the Center for Future Health. Newsweek magazine named him one of

100 Americans most likely to shape the next century. He did his undergraduate work at the University of Michigan, and earned his PhD from MIT.