Viral Communications

Media Laboratory Research

Andrew Lippman
David P.Reed

DRAFT May 19th, 2003

Abstract

Communications are poised to become personal, embedded features of the world around us. New technologies allow us to make wired and wireless devices that are ad hoc, incrementally installed and populous almost without limit. They need no backbone or infrastructure in order to work — instead, they use neighbors to bootstrap both bit delivery and geolocation. This re-distributes ownership of communications from a vertically integrated provider to the end-user or end-device and segregates bit delivery from services. Communications can become something you *do* rather than something you *buy*. This new research program explores the enabling principles for these *viral communicators* and will demonstrate their fundamental ability to scale and automatically configure themselves through a diverse set of applications including live voice, secure transmission, low-power/high-availability signaling, and sensors with a sense of place. We will address this in economic and social cases that include telephony, media distribution, safety and emerging markets.

Introduction

Innovation often comes in waves when the social and economic environments synchronize around a technologically primed opportunity. This happened in the 1930s with the telephone, in the 1950s with the automobile and in the 1980s with the personal computer. The communication industry is facing a similar disruption. As in the past, vertically integrated giants tied to centralized or mainframe technologies and services are being eclipsed by newcomers with new ideas about individual ownership, incremental adoption and instant turnover. Technology enables the change by making local intelligence affordable; society transforms that power into something useful to them, and the potential for diffuse economic investment fuels new options. The ***Viral Communications Program*** at the **MIT Media Laboratory** addresses this opportunity; companies can succeed by using research to "see around corners" rather than plow a straight line.

The key idea is communications devices that work with no central backbone and scale almost without bound. They are based on reinterpreting the basic principles of wireless in the light of economically viable digital radios that can expand spectrum capacity even as they use it. This apparent contradiction is resolved by real-time RF processing that collaboratively distributes signals and reduces the power required at each node. As with the PC, this fundamental shift in architecture, moves communications intelligence from the core of a network to the ends, and builds upon a *viral architecture** that enables infinite growth and vastly reduced costs of innovation.

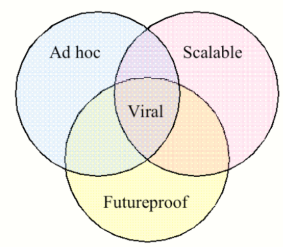

Applications range from wearable health monitors that find the nearest doctor, to smart parking meters that can download a movie to your car while billing you for parking. They can extend radio service into elevators, basements and stairwells, and can be both secure and multi-party. These emergent communicators need not require traditional "accounts" nor exist as centrally created services – they can be unlicensed, personalized, digital, renegade and "below the noise floor." *Because they can grow virally, they need not be provided by traditional communications companies.*

The program targets a wireless, copperless communications future – one that uses photons (both optical and RF) as its core networking technology. Instead of reserving wireless for special applications that require mobility, we view it as the norm. We anticipate a time when wires are the special case, needed primarily for power distribution. 802.11 points the way: it is the wrong technical solution to the right social and economic problem.

Our research frontier entails creation of infrastructure-free, scalable networks that conserve operating power, interoperate with existing systems, adapt to new radio techniques as they appear, and minimize the cost of functional evolution. This is feasible in the near term. We will demonstrate it through sample applications in domains as diverse as sensor nets and personal telephony.

---

* A Viral Communications Architecture is one where elements are independent, scalable and where each new element adds capacity to the system, so that it can be adopted incrementally from a small base and gains accelerating value with scale.

Intelligence at the Leaves

In the 1980's, the computing industry went through a major upheaval with the introduction of the personal computer. Until that time, computing had been controlled by mainframe providers and designed around a vertically integrated model that spanned the hardware, the software, the service and operation of computing functions. Intelligence was centralized, hardware was leased more often than it was purchased.

The PC shifted that intelligence to the end-users. CAD systems became shrink-wrap programs, printing migrated from copy centers to desktops, hardware became inexpensive and distributed. While the mainframe business did not disappear, its growth in comparison to the other segments of the industry dramatically dropped. Few mainframe manufacturers survived. The PC had two unstoppable advantages over the mainframe: the cost (or risk) of innovation was far lower, and the architecture permitted features that were simply not possible with a mainframe-attached terminal.

At least part of the transformation of the industry was due to the vastly reduced cost (or risk) of innovation. Barriers such as high integration costs, centralized support and a brittle architecture don't exist in the PC. The cost of change is borne incrementally, as each user purchases a new package. Software can become a garage industry and a growth potential can be detected earlier. Consider the spreadsheet as an example. It was born in a Cambridge apartment and entered the market on a toy computer, the Apple II. Its adoption flowed from small businesses to large; to this date, there is still no spreadsheet integrated into a mainframe-based application, yet almost everyone uses them.

Even hardware changes are simplified. The open bus of the Apple II and the original IBM PC (1981) permitted a vibrant adapter industry. Industry groups and manufacturers maintained momentum through evolving standards such as the inclusion of a CD in 1989 and convergence around the Ethernet in the mid-1990's. These fostered growth and enabled competition.

A second relevant aspect of PCs versus mainframes is the fundamental architecture. By moving the intelligence to the desk, features not easily supported by a centralized structure become feasible. Dynamic graphics, sound, movies, high quality color and the web would not be possible using a terminal connecting to a computer center — the requisite bandwidth is not available. All depend on local access to local storage; they trade intelligence at the ends for bandwidth in the link, and enable high performance with readily available communications and installation costs.

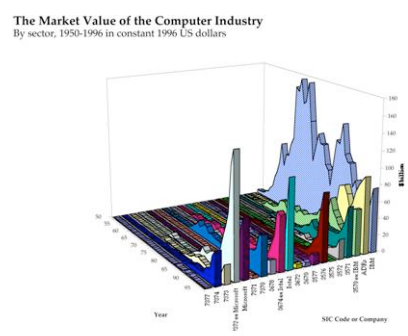

Many authors address the spurts of innovation that result from similar architectural shifts in other industries. Baldwin and Clark[1], for example, characterize this as modularity: in a modular system, designers can break through the limits of scaling a complex system. Fine[2] describes it as a periodic alternation between vertical integration and modular or horizontal organizations. He analyzes the forces that propel this shift and industry factors that dictate its speed. Christiansen[3] shows how vertically integrated businesses can be blind-sided by global changes while following a rational if shortsighted business plan.

While we dwell on the PC as an example, these authors attest to the universality of the concept: innovation and the potential for growth are coincident with major architectural shifts in the design of systems where the intelligence or adaptability is moved from the center to the ends. The argument applies to the automobile versus the train, the desktop copier versus the print center, and music industry.

The relevance to the current communications industry is painfully acute. In both broadcasting and telecommunications, the technology and regulatory environment reinforces centralized control and vertical integration. Innovation in radio systems is bottlenecked by spectrum allocation procedures predicated upon subdivision of a presumably scarce resource; wired systems are considered a natural monopoly. In both cases, the historically high cost of end-user equipment prohibited widespread diffusion of control or intelligence – television, for example, was optimized for a minimal-cost receiver; telephony places all control at the central exchange.

None of these premises is valid today. Computing, storage and communications access in mass markets devices cost less by the month, the spectrum need not be a scarce resource, and Internet demonstrates that a wired system need not be a monopoly to attract investment. Its end-to-end design principle proves that decentralization can scale, operate reliably, and permit unfettered competition and change. As with the personal computer, the enabling changes are (1) the reduced costs of the user-owned equipment, and (2) use of low cost programmability to move the value from static, fixed function designs to those whose functions are revisable and upgradeable over the useful lives of the physical device.

We argue that the impact of enabling architectural innovations is amplified when they are in synchrony with cultural change. This is what we mean by a *wave*. Such waves form when disparate industries act in synchrony rather than through deliberate ventures to overcome hurdles that none could broach individually. The simultaneity of action distributes the risk and the cost by allowing individual efforts to profit from the work of others. E-Bay, for example, resonates with modern culture but relies on the presence of consumer computing, a field of ISPs, graphic support and a body of owners. Its business plan could not finance that ensemble, but each segment for its own reasons could.

In the case of communications, the social landmark is the rapid adoption of unreliable, locally-based wireless networks (WiFi, or 802.11). While not the best technology for the job, the lesson is that local and locally-controlled communications is a similarly resonant chord in society as well as an efficient way to distribute bits.

A wave's growth is amplified even further when new forms of technology plasticity enable the technical architecture and social structures to co-evolve – when the design cycle for technology moves to the point of use from the point of technological conception and planning. The markers for this new wave are the emerging set of technologies loosely called "software radio" and "self organizing systems".

The sections that follow describe the characteristics of new architectures for communications, the potential opportunities they provide and some implications for the future of the industry. We show that high-speed, adaptable, digital devices allow us to revisit deeply entrenched tenets of communications such as the notion of a limited spectrum capacity and the inevitability of interference. Finally, we make some comments about the potential future of a communications industry where the bit delivery is divorced from the provision of services and where communications becomes something you do rather than buy.

Viral Architecture Requirements

We use the term *viral architecture* to mean a system that is adopted “virally” as that term has come to be used in the marketing industry. Viral adoption refers to a system architecture that can be adopted incrementally, and gains momentum as it scales. The growth behavior of such a system can be called *viral growth*. Though many systems are somewhat scalable, many are not viral, because they require a critical mass of adoption before any benefit is achieved. For example, a cellular telephone system is somewhat scalable, but until it reaches a certain scale of coverage, it provides no benefit to users (thus requiring a major upfront investment, and concomitant means to recoup that investment), and its scalability is limited by the inability to locate towers densely. In contrast, wireless Internet access using 802.11 architecture is more viral – an 802.11 access point eliminates the need to wire your home, and your laptop then can connect without effort in many other places as you take it with you

Each new element of a viral architecture must not deplete the capabilities of those that were there before -- to gain momentum, each new element must create more value from connecting into the system than from operating alone. That is, each adoption is a “win-win” decision – the existing elements gain a little more benefit from the new element, and each new element has a stronger value proposition for joining the system. Momentum results from this process, because a reluctant adopter will eventually be attracted to adopt when the scale reaches his cost/benefit tradeoff even when the architecture still has small reach. In the case of fax machines, this happened when enough of your contacts had or used fax machines that the case for owning one became compelling. A virtuous cycle results from a growing market cutting manufacturing costs, and increasing benefits to each new purchaser.

There are two primary design principles that lead to a viral architecture: *scalability and independence.* The first states that a viral system ought to be able to grow almost without bound, and the second requires that its elements operate autonomously, without connection to a central authority. In essence, one should be able to freely add elements and they should work without connection to a backbone. This works for automobiles as long as there are sufficient roads, and we will show that [1] that it can work for communications and [2] that the roadways are essentially infinite.

Secondary concerns are that a viral system be *future proof and adaptable*. In automobiles, the presence of either new cars or new routes does not obsolesce existing ones. A communications device should work indefinitely no matter what other communicators enter the environment and no matter how the underlying communications technology evolves.

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 it can be adopted incrementally from a small base and gains accelerating value with scale.

Examples include embedded devices whose underlying functionality is significantly altered by communications capability, such as remote controls, environmental sensors and actuators, inventory control systems, and monitoring systems. For these applications, the device may be installed once and be expected to continue to work even as new systems and devices are invented, installed or operating. A great many of these devices may need to be employed, as in the case of a building temperature system with 5000 nodes where the goal is to locate open windows or crowded rooms and alter the HVAC and lighting accordingly.

Often they have not been considered communicators at all: a dog collar that informs the owner where the dog is or what it is doing or an inventory system that directs someone to the right box or registers them, or a shipping crate that warns of expiration of its contents. The opportunity is immense: there are 10^3 cellular operators, 10^6 businesses, 10^9 portable phones and 10^{12} objects waiting to be connected[†].

More familiar communicators are also amenable to viral techniques. For example, 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, but there is nothing fundamental that precludes such a design. Such a telephone can be made to operate in both point-to-point and broadcast modes; it can carry private transmissions and propagate safety messages. The challenge here is to demonstrate that the increase in system capacity is worth more than the cost of additional complexity. Neither side of this equation is well researched (but circuit complexity is no longer showstopper it once was).

Scalable Radio

The gating function for wireless viral communications is the manner by which the systems scale. Ad hoc operation, adaptability, making them future-proof and interoperable are, to a great measure, by-products of a system that can grow nearly without bounds. Indeed, this is the technical distinction between the architectural change in computing versus communications: with PCs, each new one adds computing power, with communications, each unit is thought to divide.

Research in this area is a combination of networking, information theory, RF engineering, and physics of electromagnetic waves. In 1995, Tim Shepard demonstrated an ad hoc, edge-based scalable network architecture whose capacity increased as the number of user nodes increased.[4] This counterintuitive result contradicted the conventional wisdom that the capacity of radio architectures must be fixed and limited.

The scarcity of radio capacity has been used to argue that we must be careful to use radio only for the most important functions. Certainly the mere convenience of connecting devices without wires doesn't rank highly compared with public safety, the broadcast and land-mobile telephony infrastructures, and national defense.

But what if the conventional wisdom were wrong, and radio communications is not scarce? Can we make the roadway infinite?

Scaling Through Network Cooperation

Conventionally, wireless communications is limited by the assumption of limited spectrum bandwidth – the spectrum is thought to be a scarce but renewable resource. In the US, all of the

---

† Estimates are taken from personal statements within the communcations industry.

radio spectrum is fully allocated to services of various types: commercial, governmental, and educational. Only small segments are left to unlicensed or experimental use.

From the point of view of a lone radio receiver, capacity in bits per second is limited primarily by bandwidth of signals coming into its antenna. Signals whose frequencies overlap cannot be easily distinguished, hence the notion of interference. When the bulk of information services were broadcast services and radios where expensive, this view of spectrum made sense.

However, from the perspective of electromagnetic propagation in space, information capacity is unlimited. To see this, consider visible light or focused microwaves. There is no limit on the amount of light or microwave energy that can propagate through a given space – the photons do not interfere with each other, the air does not get overloaded or "go non-linear." With flashlights, lasers or focused antennas, the more receivers there are, the more communications there is. The limit is defined by the processes of detection, not the space itself. *The technology of the receivers and the computational architecture of the radio system, not the physics of the space, are the limits to communications capacity.*

There are a number of methods in common use that prove this point. Space division multiplexing, for example, as used in microwave and satellite links, re-uses the spectrum by restricting either the direction of emission or requiring directivity in the receiving antenna. License is granted to a *place*, not a merely a *band*. Experimental portable telephones use multiple antennas to control radiation. Most germane, phased array antennas rely on a collection of dispersed elements operating collaboratively to detect a signal that none individually could. This is the basis for radio-astronomy and communications with submarines (where the entire ionosphere is modulated.)

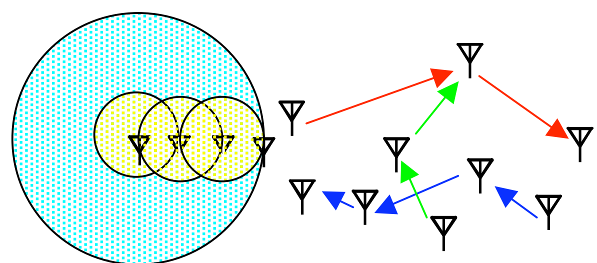

The essence of scalable wireless networks is *cooperation*, a generalization of the phased array. Shepard's packet repeater network design works because each node forwards packets of information on behalf of each other node. Since the power needed to reach an adjacent node is reduced by a factor equal or greater to the square of the distance, the total amount of energy used to carry a bit from source to destination is reduced. And since the energy radiates over a narrower region, the total amount of information that can be simultaneously traveling in the network increases as the nodes in the network get denser. In effect, each node is a "tower" for all of the nodes that are nearby; the "cells" are defined by who wishes to communicate with whom rather than the topography or zoning requirements of the place.

As the traffic grows in a packet repeater network, there is a virtuous cycle: the total energy radiated by the network as a whole decreases, while the total capacity increases, and the peak transmission capacity that can be made available to each node increases, supporting the bursty and unpredictable communications typical of communications network.

But while the packet repeater architecture provides many benefits, it is not the full story. A deeper examination of multi-user information theory and the underlying physics of electromagnetic waves (or RF photons, to be completely precise) suggests that packet repeater networks are only one step toward realizing unlimited scalability.

For example, repeating at the packet level introduces latency or delay that grows with the density of nodes in the network. Work in our group has already shown that repeating at the symbol granularity, using selective RF-level repeaters, allows the same sort of scaling without attendant latency.[5]

The key insight for achieving scalability is that radio networks should not be modeled as a collection of isolated wire-like links among nodes, but instead as a collection of cooperating devices sharing a common medium that they collectively manipulate and sample to achieve communications.

As a cooperative system scales, full interoperability among any subset of participants becomes possible, because the radio layer by definition is fully connected. Easy interoperability creates new value, which typically scales in a way that creates increasing returns to scale[6], via scaling laws such as Metcalfe's Law[7] and Reed's Law[8]. In contrast, a system architecture based on subdivision of the underlying communications medium into distinct, non-cooperative channels must be made interoperable by adding gateway devices that use additional capacity and resources. The cost of deploying applications that provide increasing returns to scale is thus increased.

Independence

The second criterion for a viral system is independence of operation. While this is almost implicit by dint of the manner by which they scale, it is a requirement for decentralized growth. I.e., it is dictated by the applications. We require independence in order to extend communications to cases for which an account or registration process simply does not apply, where it is akin to filing a "drive" plan for a trip to the supermarket.

Wireless interconnections used with computers, such as 802.11b, is an exemplar of the trend. In particular, we refer to renegade 802.11 connections where individuals deliberately open their networks to outside use, and to non-standard extensions that permit multi-hop distribution.

But more important, "de-centrality" includes a range of autonomous, distributed devices as diverse as remote controls for home entertainment equipment, portable and wearable health monitoring equipment, things placed in the physical world such as security cameras, burglar alarms, environmental sensors, thermostats, and consumer equipment such as digital cameras, personal recorders, and so forth. It is unreasonable to expect that a radio-operated TV remote control would require an account with the local cellular operator or that one would need permission to download the photo from a digital camera.

Independence results from the use of a collaborative scheme for spectrum use. The principle is that each participant in a communication has the full capability of the network. Routing, relaying, and signal regeneration are embodied in each device. The notion is not radical and is implicit in ad hoc networks[9]. There, the general definition is that each node is a router. We likewise include routing in the devices, but we extend that notion to the radio layer itself and use the RF signal itself as a routing control parameter. Rather than complicate the design, this allows viral elements to literally sneak a signal around corners or through gaps in the local spectral flux.

Future proof: Anticipating Evolution

Viral devices must be future-proof. Upon installation, the purchaser will have the expectation that they should work for the expected lifetime of the device itself independent of any other change in the radio frequency environment or the specific manner by which it communicates. I.e., your television remote should not fail by the advent of a new generation of cordless phones or because you purchase a burglar alarm, garage door opener, or new refrigerator.

Because these devices are inherently digital, their operation is defined by software. They are adaptable by their nature: design decisions are effectively shifted to the point of use rather than the point of design conception. They can adapt to changes in the environment *at the time it occurs rather than at the time they are designed.* By this we mean that, in order to meet criteria described here, they can automatically change their mode of operation as circumstances demand. I.e., they are not like fixed radio equipment that must be designed once and operate indefinitely. Thus the users can assemble their own systems, rather than depending on manufacturers to anticipate their needs, marketing specific solutions to specific user markets.

Software based radio designs and software-controlled materials (nanotech and MEMS) allow radio to transcend the limitations of fixed physical components that have restricted radio architectures in the pre-digital era. In particular:

Parameters can be downloaded rather than designed-in;

Interference can be adapted to rather than regulated against.

Network intelligence can choose the best way to incorporate new radio innovations with existing networks and radios: the new radios adapt around the legacy systems, the new radio adopts the legacy protocols, or the legacy system systems upgrade themselves to work with the new technologies.

The edge-based design of the Internet has already demonstrated the power of moving fixed application functions out of the network to the edge devices. The end-to-end argument used to organize the Internet protocols has allowed for major evolution in the application space. Such uses include the World-Wide Web, Napster, eBay, and Google. Asynchronously, the base technologies underneath the Internet have evolved from dedicated 50 kb/s telephone circuits and the original Ethernet to direct use of dark fiber networks and high-speed wireless LANs.

Evolving Radio: Getting there from here

Communications differs from PCs and automobiles and other examples of edge-oriented architectural innovations because communications has been regulated through history. This was absolutely true even with respect to print in Europe and only changed 200 years ago with the advent of the US Constitution. (One can argue that the Stamp Act played a large role in the American Revolution.) The full complexity of communications regulation is beyond the scope of this paper, however, we can suggest technically feasible means for wireless.

A wireless viral communications system is akin to a commons where each new cow adds grass — the space grows with membership. However, the two costs are (1) that each member need be a "good citizen" and use only enough power to pass a signal to the next transceiver, and (2) that each element bear the distributed cost of complexity to implement the relay function. The benefits are (1) ease of entry, and (2) reduced power consumption in devices. The regulatory issue hinges on whether this is a burden or an opportunity.

We advocate an "Open Spectrum", where large, usable swaths of spectrum are left unlicensed. Such a policy reflects the ready adaptability of digital, software radio in the sense that the actual signals, siting and radiation characteristics are left open to designers. In such a space, a relay-based network should out-compete other, more centrally organized architectures. We expect that viral communicators will demonstrate their utility in such circumstances primarily because their scalable and ad hoc nature allows demand to influence deployment directly.

Evolving the Communications Industry

In the end, viral communications transforms communication from something you *buy* to something you *do.* Independence of operation allows communications services to be separated from traditional service providers. At one level, this is a pure threat to the extent that the economics of the telecommunications industry business model depends on exclusivity of service provisioning. On the other hand, it reduces the bet. 3G telephony is betting that the next killer app is a cell phone with a camera in it. That may well be true, but it is a large bet for a small idea. SMS, on the other hand, was a by-product of system operation and proved to be a big winner. The lesson ought to be that low entry levels for new ideas can help more than it can harm. Secondarily, SMS also taught us that synchrony with the culture is more important than technology for some features — it was the populace at large that made it succeed, not the operators themselves.

Alternatively, low-cost real-time relays can be used to extend existing networks, by getting the signal into places where a tower can't reach such as elevators, basements, interiors. Relays can be placed literally anywhere and are potentially less environmentally evident. The inherent ability for viral systems to both broadcast and provide secure communications allows a new generation of emergency, safety and entertainment opportunities. Thus, the apparent threat of a new technology can also support existing businesses.

Our research program stresses other options. Specifically, we view the ability to freely add communicators as an opportunity to link up things in the world that are below the threshold of entry for a telecommunications provider but add to the quality of life. Their utility is borne out by the activities of home networking alliances[10] and research such as the Berkeley Motes[11]. Other labs operating under the aegis of extending WiFi as well as corporate ad hoc networking efforts (e.g., Ember) provide some validating experience. To the best of our knowledge, none of these programs directly address the scalability and cooperation issues presented here, and none migrate the network control to the RF layer.

The arguments presented above support a thesis that new growth in communications need not come through a centrally organized and controlled communications industry; indeed, it can be driven by the evolution of the end nodes, with applications emerging incrementally. Like computing, communications can become an embedded feature of any other device that has an electrical component.

Current Research (Partial List)

The research program is based on the following "mission statement", presented here in outline and represented throughout the Media Laboratory:

Photons for information; electrons for power:

- Basic Enabling (Viral) Technologies
  - Connectivity/Scalability
  - Power
  - Time and Place
  - Sensor Technologies
  - Representations (sound, image, bio)
- Focus on Kids, not businesses
- Design
- Industry and Regulatory Dynamics

The core of the program is the viral technologies described above and related elements that enable its application. We rely on youth to invent applications for us -- to a great extent, the measure of an invention is how rapidly it is wrestled from the hands of the inventors and re-made in the image of the users (witness SMS). Instead of doing this speculatively, we provide the basic technologies in constructive learning and play environments as quickly as possible. Design is essential, since communications is an element of lifestyle.

Industry and Regulatory Dynamics are the focus of the Communications Futures Program at MIT. This is a joint effort of the Media Laboratory, Sloan School, and Lab for Computer Science. Charles Fine, David Clark and the authors of this paper jointly direct it.

Infinitely Scalable Radio (A. Lippman, D. Reed)

We will design and engineer prototype radio systems where additional elements add spectrum capacity rather than share it. These are low-cost, low-overhead nodes that relay information in real time. These systems permit operation without an infrastructure, conserve energy and ultimately segregate delivery of bits from the provision of services. They are a combination of the heretofore distinct areas of ad hoc networking, diversity radio, and multi-path, bit-by-bit routing.

Intelligent RF (Elevator Radio) (A. Lippman, D. Reed)

We propose the construction of a collaborative wireless system that uses a combination of broadcast and multi-hop techniques to provide signaling to places that traditionally are isolated, such as elevators, basements, shadowed and weak-signal environments. This can be implemented as a safety system that gets a message to all members no matter where they are, or as an emergency alert system that warns of traffic jams and accidents.

Diffused Location (A. Lippman, A. Bletsas)

We develop techniques for automatic, decentralized computing the position and time of adhoc nodes in a network. These work both inside a building (unlike GPS) and with automatic corrections. The goal is to be able to pepper a region with devices that are sensors and relays without having to map their position in advance. Each learns it place and time by interaction with neighbors.

Expectation Propagation for Adaptive Wireless Signal Detection (R. Picard)

There exist many wireless communication standards, such as GSM, iDEN, Edge, and CDMA, etc. They differ from each other in many technical aspects. But almost all of these digital communications systems contain an algorithm for demodulation, to recover the original modulated information bits. We have developed a new algorithm, which could be applied to these real-world communication systems with some modifications to fit a specific standard. Specifically, we use a hybrid dynamic system to jointly model both the transmitted symbols and the communication channel coefficients. Then, we develop a new algorithm based on expectation propagation.

In our simulations, the performance of the new detectors is very close to the genie bound, which achieves the best estimation result by given additional observations, and much better than the classical differential detector.

Finally, we would stress that our algorithm could also be applied to other digital communication problems, such as joint decoding and channel estimation, and multi-user detection.

Opportunistic Storage (A. Lippman)

Opportunistic storage is a group of small devices that communicate with a short range wireless network. Anyone can come in and store data on this network, and it gets automatically and randomly split into pieces among the individual nodes. When the memory of one node fills up, it sends the data elsewhere. More nodes add capacity. Later we can collect the data and re-assemble it. The devices organize themselves and collaborate to reconstruct the data.

Gossiping Sensors (V. M. Bove)

We are building smart audio and visual sensors that autonomously collaborate and only talk to us when they collectively figure out an answer to a problem. A group of microphones can follow a single, moving conversation at a crowded party, a set of cameras can detect items borrowed from your office or pull out a face from a crowd. This is basic research on media representations in high-density networked environments.

Conversation Finder (C. Schmandt, S. Marti)

Small personal speech detection hardware with forms a self-organizing wireless network, with each node independently determining with whom its wearer is conversing. Conversation detection is based on how turns of speech either overlap or collide in normal conversation; repeated collisions identify independent separate conversations. The goal of this work is to provide context to control the level of intrusiveness of incoming communication activity from a remote party.

ListenIn (C. Schmandt, G. Vallejo)

Small IP-based audio sensors of several types are distributed around a living area; when detected acoustic energy changes they transmit audio samples to an on-site server. The server attempts to classify the audio, including detection of a speech signal. The server then conveys similar representative sounds to a remote monitoring location, or, in the case of speech, the actual signal modified to prohibit intelligibility. These are limited to short sounds, and are meant to convey an ambient auditory awareness of the home to a remote family member.

Ecological Thinking (Mitchel Resnick)

Success in a networked society will require not just new skills and new knowledge, but new ways of thinking. Rather than seeing the world as a clockwork mechanism, people will need to think in more ecological terms, recognizing the importance of adaptation and improvisation, and understanding how patterns can arise from many simple, local interactions. We are developing new technologies and activities to help people develop as "E-Thinkers."

System Play (Oren Zuckerman and Mitchel Resnick)

We are developing a new generation of digital toys and games that promote system-thinking intuition through hands-on activity. System Blocks, for example, are a set of computationally enhanced blocks that enable children (and adults!) to explore systems concepts such as feedback, causal loops, interdependencies, and side effects. Each block has a pre-defined behavior that represents a conceptual building block in a dynamic system. When children snap the blocks together, data is transferred from one block to another, forming a simple system simulation.

Solving the Last 25km Problem (A. Pentland)

Many regions in the world are embarking on infrastructures that do not use the presence of wired and cellular telephony as the model. Bhutan, for example is moving from pre-cable television to national Wifi; China may follow suit. In this area, costs and ubiquity are the main issue. We will develop analytical and engineering tools for connecting communities that couple people and devices, connect to backbone infrastructures as well as operate independently of them.

Developing Viral Communications (A. Pentland)

Just as the cell phone industry is more in developing countries than in developed countries, so, too, new types of communications infrastructure are seeing their first light of day outside the American and European markets. Driven primarily by low capital cost of entry and scalability, urged on by declarations by Kofi Annan and bets by forward-looking IT ministries, WiFi based viral communications infrastructures are already appearing on a small to medium scale. We have developed two types of viral communication systems, one extremely low cost and based on mobile transmitters, the other customizing the 802.xx family, and are now conducting field tests both for quality of service and for profitability of the associated business model.

---

[1] Baldwin, C. and Clark, K., *Design Rules: The Power of Modularity*, MIT Press 2000.

[2] Fine, Charles, *Clockspeed : Winning Industry Control in the Age of Temporary Advantage*, Perseus Publishing, 1999.

[3] Christensen, C., *The Innovator's Dilemma*, Harvard Business School Press, 1998.

[4] Shepard, T., Decentralized Channel Management in Scalable Multihop Spread-Spectrum Packet Radio Networks, MIT Laboratory for Computer Science Technical Report MIT/LCS/TR-670, July 1995.

[5] A. Bletsas and A. Lippman, "Efficient Collaborative (Viral) Communication in OFDM Based WLANs", IEEE/ITS International Symposium on Advanced Radio Technologies (ISART 2003),□ Institute of Standards and Technology, Boulder Colorado, March 4-7, 2003.

[6] J. Brian Arthur, *Increasing Returns and Path Dependence in the Economy*, Ann Arbor, University of Michigan Press, 1994

[7] Metcalfe's Law says that the value of the network that comes from switching (connecting one element to another) grows as the square of the number of elements. It is named after Robert Metcalfe, inventor of the Ethernet.

[8] Reed's Law says that the value of the network that comes from supporting the formation and sharing of information among persistent groups (group forming networks) grows exponentially in the number of elements. It is named after David P. Reed, one of the authors of this document..

[9] An ad-hoc network is a network built by its users out of a collection of communications elements that can be arranged to work together. In Latin, *ad hoc* literally means "for this," further meaning "for this purpose only," and thus usually temporary. We apply this term to networks in which new devices can be quickly added to the network in a way that allows them to use all the resources already present in the network.

[10] E.g., The Internet Home Alliance (www.internethomealliance.com) mission is "to accelerate the development of the market for home technologies that require a broadband or persistent connection to the Internet." The HomePNA (www.homepna.com) is organized to promote home phoneline networking. The Wireless Communications Alliance (www.wca.org) is organized to share information on a variety of wireless communications options. The WiFi Alliance (www.weca.net) focuses on the 802.11 WLAN technologies.

[11] Berkeley Motes are tiny computer-based sensor nodes that can be deployed easily to sense and communicate. They are being developed at UC Berkeley by the WEBS project.