

CFP Bi-Annual Meeting: San Jose, CA

January 22-23, 2008

Meeting notes by Polychronis Ypodimatopoulos & Nadav Aharony

Tuesday, January 22, 2008

Keynote address

Mark Carroll, Sr. Director Video TG, NMAI, Cisco
P2P to Service Routing

P2P and Service Routing

Mark Carroll, Sr. Director Video TG, NMAI, Cisco

(NMAI = New markets Architecture and Introduction)

Appliance products / routers

Landscape today:

-Media rich applications and services

- Content providers looking to:
 - Concentrate core business
 - Monetize content
- P2P nets enable mass distribution of content to the end user
- SPs try to satisfy aggregators and add value

Access SP P2P Observation:

Issues

Containment

Collaborate

- Traditional methodology for peering
- Deploy
- Some do P2P caching
- BBC/BT trying to overcome P2P problem by means of P2P caching
- SPs need something to keep pace, need to offer value to the network, build applications timely

Service routing & Service Node:

- Replaces technology stretch
- Enables key new relationships and business models

New media optimized routing paradigm and thin layer implementation

IP infrastructure <-> provider apps, content provider, over the top apps

over-the-top players: built tools on top of platforms, add value by getting the key content providers on their side.

Towards a common platform: Service Routing: Cache, put/get, event, ...

- develop proximity, Delay-based balancing
- widget plugs-in without worrying about underlying caching, distribution, etc.
- This would allow SPs (service providers) deliver a platform over routers/switches

Key elements:

SG: service gateway

SN: service nodes (workflow, association with middle-layer)

FR: foreign resource

routers<->SGs<-> SN<->FR

each layer provides abstraction to the layers below it

Service routing protocol: Leveraging DHT technology:

- Host function (put/get functions)
- routing function
- forwarding function
- TCP/UDP/IP

Scaling services, Not servers:

Create non-linear relationships: you have users/apps in place, so now we can create bindings between users/applications

Potential provider models – capstone:

- Metro, ISP Backbone
- Self-peer, superPeer DHT (tier 1), Peer DHT (tier 2)

Service Networking – Recap:

Potential Topics collaboration:

- Searching in DHTs!
- Network search
- net security
- Lights out network management
- Presence and location layered on DHT
- Name/address/hierarchy
- Next generation media/apps/services
- Alternative business models

new models of collab. between industry, gov. and acad. - the possibilities are limitless

Q&A/Discussion

- Marie-Jose Montpetit:
- Q: can we collaborate on presence?
- A: yes, in the service routing protocol

- Q: Intel: the content provider is the consumer?
- A: they will become the same.

The P2P Elephant – What is P2P and why is it significance?

David Clark (MIT)
Chair/Introduction

- P2P is an opportunity. how do we turn into a research agenda? Cover the various aspects of P2P, get perspectives and redefine in more useful ways, talk about platform liquidity (be able to position what you're doing in a way that allows to float around the platform)
- traffic patterns, what are the traffic problems?
- Context aware platforms
- Case: television and P2P
- Open 700 MHz auction
- (generalizations of P2P)

David Reed (MIT) (Viral WG)
Beyond P2P: Liquidity and Virality

- P2P a year ago: why do we talk about P2P?
- If P2P is what people want, how do we monetize?
- Maybe P2P is what people really want.
- From one point of view, P2P is about resource allocation

- What are the apps going to be about?
- This is not about Tele-* terms

- Platform network Theory (Reed's Law):
- Not just about how bits flow, but about value that is created by interconnection structures, maybe on-demand

- Internet: we recognized that all different networks needed to be interconnected
- Create “excess” affiliation options that transcend boundaries

- Network invites Virality:
- 1: sudden emergence of a phenomenon
- Facebook: in 1 year, extreme speed
- 2: viral systems appear in the consumer market first, not in the enterprise market

- key things:
 - scale without bound
 - incremental growth

- Viral system: runs on a “host”, exploit its resources
- Contributory participation: fastest way to roll out a system, P2P is an example

- PlanetLab, Skype, value shared with users who then repay host
- Minimal cost, high “potential”

- Facilities as Albatross:
 - facilities are already commoditized
 - Amazon web services: cheap resources available to customers
 - very easy outsourcing to such service providers
 - But, users will not tolerate lockin to facilities

- The Net is now Function Liquid:
 - Where to store?
 - Where to compute?
 - How to communicate?
 - How to search?
 - How to authenticate?
 - How to charge?

- P2P provides an additional option.
- this liquid architecture is now flowing around

- Viral Growth and Evolution:
 - Biological evolution: competing for resources
 - Given enough scale, the virus becomes the host: Facebook was a virus but is now a host to new applications (eg. application “causes” has spread over Facebook, causes compete with each other)

- The Network Mirrors of “Us”:
 - the host of Facebook: the actual social network of humans, not the technology network.
 - Kwan Lee: create social network for quantitative analysts for Bank of America. More successful than a simple portal.

- Identity is relative: Networks of You:
 - Sociological def. of “identity”: relative to person, it is a set of networks of “You”
 - Content providers are only part of network.

- Awareness and Coordination:
 - We are social, cooperative spiders, building our own webs.
 - networks of not just switches and nodes, but of nodes that aware of each other (eg. Twitter web app.)

- The reason for cross-layer is not just optimization but awareness of other entities
- Polychronis: Presence in mobile networks
- Nadav: cognitive networking

- Integrating a community of people and devices:

- Attributes of Network:
 - mobility
 - awareness: ends sense and adapt
- Transport independent:
- Glue might limit need for transport layers, but this is not necessarily threatening to providers such as Cisco

Charlie Fine (MIT) (Value Chain Dynamics WG) Value Creation and Capture in P2P Systems

- Mark Carroll said: how to monetize this and that: The VCD group thinks about this.
- Where is value created/captured, how does this change over time?
- Cisco: we need to change the way we offer services to create/capture value
- Who are peers? What do they do?
- End user “peers”, Transport “peers”
- teenagers used to spend hours on the telephone, now they do that on Facebook
- consume information: free or buy information
- distribute information: my opinion or sell something
- Trans. peers: all for profit
- concerned with ROI, monetize their pipes
- What are platforms? What do they do?
- Aggregators/platforms/frameworks
- some for profit, some not
- they receive service from transport peers to provide a platform to end user peers.
- Multi-sided platforms:
 - ebay: buyers-sellers
 - google: content providers – advertisers – viewers

Roberto Saracco (Telecom Italia) The Telecom Italia NGN2

- big portion of investment on wireless (6.5bn euros total)
- pervasive fiber infrastructure feeding antennas
- upgrade: more bandwidth for even less money from the consumers
- Why? the main driver for this cost reduction: copper costs a lot, fiber does not
- Context – Service Platforms:
 - revolution is mainly due to opening up the network: provide platforms for end-users to better use the network. It's very difficult to get more money from the user. Instead monetize on relationship between end customers and other companies.
 - main services: voice, sms, ringtones

- we keep providing new services though. Out of the many thousands services to be deployed, only few will have ROI.
- New architecture/areas:
- Health care, etc. We must create an ecosystem for the services to be successful, not just provide a single service. Stakeholders: doctors, patients, infrastructure (hospitals etc)
- Infomobility platform:
- Eg, monitor a car, deploy services on the car, etc

Q&A/Discussion

- Nokia Siemens: Virality and liquidity vs. keeping control over what's happening in the network (presented by Telecom Italia): how do these fit with each other?
- C. Fine: TI ecosystem of healthcare: in order for it to be successful, it must have some viral aspects
- Polychronis: TI's infrastructure can act as host to new viral services/apps
- DP Reed: Host is initially overly constrictive.
- Saracco: the value is moving from the infrastructure layer, you take that for granted. Wifi services are not really pervasive, we need passwords etc., so there still is a lot of work that still needs to be done.
- Clark: what were the control points to which TI wanted to get hold? Cisco tries to sell to every operator but is your architecture open to competing companies to sell their own products?
- Davie: you cannot tell a service provider to just buy something. instead we standardize. The new architecture demands new definition of SPs. they should compete at the service level.
- D. Clark: ISPs making new relationships: is this a space limited to ISPs?
- Reed: interconnection is important. at international level is it possible to route services to access local services?
- D. Clark: different view of what Akamai looks like: don't care about equipment.
- ??? : Cisco is open to standardization, but there is an efficiency in local optimization
- Whittaker: amazon is missing from this conversation. this is not a carrier/manufacture battle anymore
- ??? BT: obsession about standardization. Standards are not the success, it is scale. Eg. Apple, youtube
- they did not start by being open. We spend too much time on making standards and loose on time-to-market.
- Saracco: in order to make money we need both standards and time-to-market.
- D. Clark: truly scale does not need openness. Akamai is a 3rd party offering global scale. Operators though have to figure how to cooperate. Akamai did it with just a lot of computers. Operators need interfaces to allow cooperation, which also degrades product differentiation.
- Reed: akamai caught the wave before anybody else. We see the same with joost and skype. Operators have the infrastructure but they never move first! instead the let others make the first step towards new markets.

- Montpetit: peers need to be able to talk together. operators own the trust authority, not infrastructure providers. P2P needs standardization of interfaces between devices.
- Saracco: we missed skype, but are they making money? No. We are into making money. Internet over mobile is tiny in Italy. Monetize beyond connectivity.
- Fine: google brought a third party (advertisers) and made money. who can you bring in and create value to monetize on?

Platform Liquidity – Is P2P a transition to a “liquid” Internet?

David Reed (MIT) Introduction

- There will be 2 presentations in this session.
- PlanetLab – even though it is an academic project it could represent a model of networks of the future. It has an interesting business and control models. There might be a chance to commercialize it as well.
- We were also supposed to have a talk by Amazon, but they were not able to come. Instead I will discuss my thoughts in this area and review the companies currently working in this space and try to think where this can go in the future. Discussion will try to focus on what might be next steps for the operators and vendors in the room, and think about what business models could arise from this.

Jeff Sedayao (Intel) PlanetLab: A Platform for Planetary-Scale Services

- The talk is about what is PlanetLab, and what are planetary-scale services.
- PlanetLab is:
 - Technology
 - global network
 - Organization
 - Currently hosted by Princeton
 - Socially
 - Research infrastructure made available to the global community.
- 412 sites around the world (836 servers)
- Federated with PlanetLab Europe which is currently spinning off the main initiative.
- Mostly standard Linux. Number of global services.
- Many academic participants and industry participants.
- Content Distribution
 - 1993: in the beginning there was a single copy of the data.
 - 1998: end-to-end design breaks down; there were more than one point of serving the information.
 - Today – planetary scale service (like Akamai). Overlay concept.

What you need for planetary scale services: Pervasive, robust, adaptive to conditions in the network, scalable.

- You want multiple vantage points on the network.
- A little computation at many locations.
- Need to have management services that are appropriate for the task..

Solution: Share the platform. Everybody contributes a piece of the platform (minimum – 2 servers). Time sharing of the network among participants. This could be the model of future public computing utility.

Platform must provide: Isolation to protect services from one another. Market-based resource allocation

PlanetLab Services Are running:

Event processing, network mapping, distributed hash tables, content distribution, web casting.

Resource reservations:

CPU resources can be scarce during certain periods (e.g. before paper deadlines). There is a calendar service and resource reservation service. Increase priority for certain time periods. This seems to work well.

This is for CPU usage, not network usage so far. One cannot reserve network resources at the moment.

Planetlab Today:

- 836 machines, 412 sites. Principally universities, some enterprise. Globally distributed.
- Machines virtualized at the syscall level.

What we got right:

- Immediate impact within 18 months 25% of publications at top conferences were PlanetLab experiments.
- Became an “expectation” for validation of large system results.
- Incident response:
 - Early: very conservative.
 - Later: less restrictions
 - Education: learn from what caused the incidents.

We could have done better

- Community contributions to the infrastructure: infrastructure development remained centralized. Paying the price now..
- Support for long-running services:
 - Researchers aren’t motivated to keep services running for multiple years
 - Decreased the amount of service composition (can’t trust the dependent services will
- Admission Control

Open challenges:

- Community ownership of availability. Need to motivate decentralized management.
 - o Who keeps the nodes running?
 - o What happens when the nodes aren't running?
Added policies: if your nodes aren't up, you can't reserve time slices (to motivate users)
Showing public graphs of freeloaders – who uses the network while their own resources are down.
- Resource allocation aligned objectives
 - o Performance, innovation, stability
- Standardization
 - o Standard interface → platform stability
 - o Open architecture → improved innovation.
- Tech transfer
 - o How does one get the technology to commercialize?
 - o There are a number of private planetlabs out there.

Planetlab and Industry

- Global communications company
 - o Incubator for future internet infrastructure
 - o Emerging services become a part of the internet
 - o Intel tried to experiment with a private PlanetLab that connects regular computers (e.g. laptops) and allow users to create virtual machines on their peer's devices.
- Global computer vendor
 - o Platform for planetary scale services.
- Software company
 - o Test bed for next gen apps
- Fortune 500 Company
 - o Next generation opportunities for IT staff
 - o Leverage deployed PlanetLab services for CDN, object location, network health.

Q&A/Discussion

- David Reed: Princeton has done the core network management services well, its open source. But service composition still does not work well.
- ---??? -- do you find that there are coordination problem with management between different services? How do services do network monitoring?
- Jeff S: services tend to do this on their own (monitoring, management.) Princeton also provided some API for doing different things.
- ????: how do you convince people to install other people's experiments and software on their machines?
- How do you deal with heterogeneous systems, how do you get people to agree to install this
- How do you install your own software on the remote machine?
- Jeff S: as long as it works on Linux, you can compile and deploy it on the virtual machines. If you want to run something very different, like windows, you are out of luck. All servers run the same version of Linux.

- David Reed: the virtualization is in the user space and provides virtual resources (memory, disk space). Applications don't collide.
- Mark Carrol: This could lead to interesting models where many vendors can collaborate on such a distributed system. What Amazon and Google do today is closed and proprietary, but PlanetLab could be an interesting model for competitors to collaborate.
- Jeff S: Location of the servers really matters. There is only so much that a single company can do at its own locations.

Discussion led by David Reed (MIT) Will Clouds Reign? A Stormy Future for Pure-Play Communications

We've thought of communications as Shannon did – bits transferring over channels. But human communications are very different. Human communication is all about sharing, not being a post office of messages. We've forgotten that for humans, communication is a much richer space than just moving bits from place to place.

Agenda:

- Infrastructure is for applications
- Clouds: distributed computing (30 years late). Liquidity
- Mobility and Clouds.
- People and Clouds
- Business and Clouds

Infrastructure exists for application, not for itself. What applications are best served by network businesses and business models? These were good for conversations, broadcasts, remote access to distant resources, which are not people. What are the emerging applications?

- Locating relevant information
- Augmenting the Users' environment

When I move from my home to my neighbor's home I lose all of the environment I have at home (e.g. my Comcast subscription)

Tradition: connectivity platform.

- Model based on the fact that "Tele-" is the problem to solve.
- Resources are fixed – servers, people, content
- Transporting bits faster, farther, cheaper

Future is about hyper-mediation, not necessarily hyper-connectivity.

- Web 2.0 – mashups, youtube
- Social networking – facebook, etc.
- MMORPG
- Online collaboration and simulation
 - Second life, coquet

- Augmented reality
- Mscapes, smart cars

Distributed Computing: Thirty years late:

- Location, Identity/relation, privacy/safety, resiliency, sensing/capture (cameras, temperature, other sensors), interaction/expression, computing, glue (comms)
- Infrastructure companies of today provide mostly the glue.

Architecture for hypermediated apps: composable functions

- Local computing and interaction (cell phones, laptops, other)
- Wireless glue , positioning, sensing
- Databases , filters, searching, recording
- Social glue and tracking
- Resilient resources configured on demand. Things are dynamic; we need to be able to dynamically assemble anything from an mini-computer to a super-computer, from flash drive worth of memory to a petabyte.
- Authentication, security
- Economic elements (billing, payments).

3 distinct clouds of resources:

- Peer to peer cloud (edge)
 - Skype and croquet
 - End users, personal computing
- Service-on-demand could (center)
 - Web 2.0, SL, World of Warcraft. The user does not need to run the WoW server, they just want the client. There is in reality a huge server cloud distributed in many locations.
 - Amazon, HP, IBM – “Cloud computing”
- Pervasive-context cloud (neighborhood, social context)
 - Social networking, augmented reality, navigation systems are also starting to provide this.
 - Google (search, OpenSocial, Android), Facebook. Less about search, more about creating this big context cloud.

Amazon web services:

- Key idea: dynamically scalable resources
- All the things one needs to develop a web based business. Help a small company grow and provide dynamic resources according to its needs.
- EC2 – servers configured on demand for pennies.
- S3 – reliable storage on demand
- SimpleDB – databases on demand
- FPS – payments on demand.

Google:

- Geospatial resources (maps, streets)

- OpenSocial – a platform framework for maintaining human-human relationships and identity
- Android – a personal platform that incorporates contextual framework
 - Going against a relatively open platform (S60)
- PlanetLab as a prototype:
 - Slices and slice management exemplify planet-scale coordination functions
 - “contribute to participate” business model
 - Avoids pricing by arbitrary cost accounting.

Why not create something similar for today’s service operators? Roaming agreements, etc. allow global deployment of services.

National boundaries traditionally block cross-country deployments of technologies (trying to protect local industry). Maybe that’s not the way to go, maybe giving the citizens more services is more important.

PlanetLab has its own challenges (e.g. porn taking over resources. Second Life is dealing with similar issues)

Mobility and Clouds

- A supercomputer in your hand
- A data warehouse on your belt
- Exchanging 1 terabyte in business cards – exchange a URL for the terabyte that’s sitting somewhere else.

People and Clouds:

- Your Life Bits are stored throughout the cloud
- Your social relationships are mediated through the cloud.

Businesses and Clouds:

- Enterprise provides additional context in each cloud
 - Roles/relationships
 - Data stores
 - Payment systems
- Inter-enterprise – clouds instantiate resources and value.

Communications value chain is redefined by this type of hypermediation.

- Comm. Platforms provide context
- Transactions, awareness, search, persistence, resilience become essential platform elements
- Value chain embodied in the cloud via mediated business arrangements.
- Resources become liquid.

Action items:

- Computing distribution networks in mobile spaces up for grabs (Akamai for mobile)

- Interoperability for mediated services across providers – a single company can't build world-wide infrastructure on its own.
- Peering required economic exchange framework

Q&A/Discussion

- Charlie Fine: I broke David's talk to 3 parts:
 - Thinking about Supply, Demand, what are the constraints we need to think about, what tests we need to pass before starting such an initiative.
 - Google seem like they are throwing many things on the wall and seeing what will stick.
 - ??? – unlimited outsourcing leads to a loss of control of one's destiny. If you make a business that relies on such services, if they decide to stop the service you are left with a need to adapt on your own. You put photo's on someone else's server and you may not know what could happen with it. What if the data is gone? What are they doing with it?
- David Reed: had to deal with such issues in the past when was involved with Lotus. Should we do our own O/S since we can't trust Microsoft? Sometimes you don't have a choice. When I put my photos in the clouds there is also a backup DVD in my closet, but I still rather access it through the cloud.
- Andy Lippman: one needs to find ways to store information in a way that's generic and transferrable between different service vendors. Not proprietary formats.
- David: Can we hold myspace to some society level requirement of what they are allowed to do or not?
- Karen Sollins: users may not understand the risks that they are facing (backup? Privacy?) they may not even be aware of these problems we listed above.
- Dirk (BT): If all is defined as an open policy, where does competition come in? How do you set up the economic ecosystem that won't kill competition?
- Would love a piece of software that made the user aware of the privacy and other concerns of users (e.g. over exposing private information in Facebook). This could be a potential business...
- David Clark: What is the sticky thing here, and what is the things that are not going to be centralized. Today Ebay holds your rating; facebook holds your social data.
- Creating the interconnection and negotiation systems to interface the different companies - that could be very complex. A deep research question is embedded in the interface and how much depends on the agent that works on your behalf:
- David Reed: my fantasy is that there will be bot-nets running around looking for evidence of our data and who is looking at it.
- ?? BT: there is need for legislation. For example, in California they made a law about real estate leasing. Online there are usage agreements and everybody just accepts them to get the service.
- Maybe we could have preset contracts that are mediated by an impartial party (e.g. the state)
- Perhaps we need some court cases that would happen before that.

- Dirk: there was a case where you can build a piece of software that scans agreements and looks for dangerous terms. There are ways to do this, but Nokia didn't turn this into a product.
- David Reed: human subjects committees try to protect people when dealing with research experiments.

Understanding Traffic Patterns – What's really going on with Internet traffic?

Bill Lehr (MIT) Chair/Introduction

- Why this session? There are many questions that require good empirical traffic data.
- Is there really a broadband investment problem? Are p2p and youtube traffic really that problematic? What resources are scarce?
- How do we decide what to observe? Timing? Frequency? How to observe? How to share it? Etc.

Andrew Odlyzko (University of Minnesota) Internet traffic growth trends

- Main points:
 - Internet traffic growth is slowing.
 - Hype is accelerating.
 - Even very biased hype is occasionally correct – trustworthy data collection is desirable.
 - There are huge sources of potential future traffic. The future traffic levels are a result of complex feedback loops.
- Internet growth hype: many people published prophecies of bandwidth becoming chronically scarce.
- Those rates given in the 90's and around 2000 were very exaggerated and did not come true when we extrapolate them to today. If you look at the numbers of a few petabits per second per person in the US you get that every person (300Milion) uses 10Mbps every second constantly (not realistic..)
- There were underestimates of the demand for high tech products (computers, mobile phones).
- When we look at actual results we see its less than predicted. Slide 6 presents some of the data.
- Australia is about a 1/3rd of use per capita compared to the US (US is about 2GB per month per capita). Korea and Japan are 6 times more than the US (~12-13GB per month per capita). The numbers don't seem to be directly related to the links of the country to the Internet. Japan has similar connections to Korea, for example, but its usage is similar to the US, 2GB per month per capita.

- Sources of data: Public disclosures, Company updates about backbone and wireless data traffic. Sometimes gets it on confidential bases. Some information comes from carriers under an NDA.
- Information is public on their website. <http://www.dtc.umn.edu/mints>
- They have a program that collects such information out of databases, they reverse engineer the feeds and graphs and run various studies over the data.
- ESnet: longest available run of reliable traffic statistics (us department of energy funded network). 17 years of traffic. That data shows compound annual growth of 85%.
- Hong Kong – government sponsored data collection. Per capita traffic intensity in HK is about 6x the US level.
- Huge potential sources of additional internet traffic:
 - Storage
 - World wide digital storage capacity at 2006 - 185,000 Petabytes
 - Year end 2006 worldwide internet traffic – about 2,500 PB per month.
 - Broadcast TV
 - Year-end 2006 – internet traffic per capita – 2GB/month
 - Year end 2006 US tv consumption per capita: ~40GB/month
- Revenue per MB:
 - SMS - \$1000.
 - Cellular calls - &1
 - Wire line voice – 0.1
 - Residential internet – 0.01
 - Backbone internet traffic - 0.0001
- Volume is not value, but is indicator of ecosystem health and growth.
- Would love to share data and collaborate, especially if someone has more links to traffic data to make the results more accurate.

Q: Data about number of subscribers.

A: Most of the growth is per subscriber, not necessarily new subscribers. Not much visibility on China.

**Eve Schooler (Intel) & Jaideep Chandrashekar (Intel)
Inside the Forbidden City: A Look at End-Host Traffic Inside a Modern Enterprise**

Had access to large set of enterprise data.

Motivation: rarity of end-host enterprise data.

- Internet VS enterprise traffic

- Most enterprise data was from universities
- Edge vs. in-the-middle monitoring
- Holistic analyses of the end-host
 - o Wired vs. wireless, internal vs. external, managed vs. unmanaged.
- Debunk the myth of “normal” traffic
 - o Impact of context on the norm
 - o Quantity variable in the norm
- Quest for simulator realism.
- Network security and manageability:
 - o Can we use such knowledge to design better anomaly detectors?
 - o Collaborative detectors
 - o Management traffic impact
 - Trying to learn how distributed or p2p control protocols look like. Also, how does it look like on wireless?
 - o Comms optimizations
 - NIC i/o architecture
 - Energy efficiency

How and what to collect?

- App-level wrapper: two end-host monitors:
 - o Windump – free packet trace and capture software
 - o UserMon (Kenny) – sample machine level variables periodically to track user activity (restart if it dies)
- Traffic captured from all interfaces w/pcap
 - o Wire, wireless, VPN intercepted pre-encapsulation
 - o Confect packet traces to flows post facto
- File naming convention “anonymizes” user ID
 - o Eg. userHash_ipaddress_location_vpn_timestamp
 - o Distinguish address from user (when re-used).
Wanted to correlate the user but also the address that the user has. Address does not necessarily equal the user.
- The proxy problem:
 - o Wanted packet headers → instead, collected 150B prefix
 - o Post process: find actual dest in payload, rewrite header, and discard rest.
- Periodically uploaded to remote file store
 - o Opportunistic connection, max storage used, max timer, max attempts
- Challenges:
 - Privacy
 - o Require user consent, therefore use must pull the app.
 - o
 - Overhead Concerns
 - o Disk storage, network delays they might encounter when uploading traces, CPU, memory.
 - Population Diversity:
 - o Cross-geo call for participation

- Cross-organizational – R&D, business units, Network operations
- Sample size
 - Asked ~4000 individuals, got several 100 volunteers
 - Asked via organizational mailing lists, newsletters, web pages, at quarterly all-hands meetings, security oversight committee.
- Enterprise configuration
 - More machine heterogeneity than anticipated
 - Myriad of firewall, NAT, VPN issues thwarted uploads.

Deployment ramp up – slide 7.

At the end ~400 users, ~5 weeks, ~400 GB of data.

Early findings:

- Traffic locality – most traffic never leaves the enterprise. Be careful where the statistics are collected.
- Only ~12% went outside the enterprise.

User Locales:

- A month in the life of a laptop
 - You can tell a lot just from the interfaces used by the machine. (Slide 12)
- Diversity across environments –
 - Being inside or outside of the network is very different as far as behavior is concerned. Just when looking at durations of connectivity.
 - There is diversity between users as well.
 - When doing network or capacity planning these patterns need to be understood.
 - Threshold on behavioral anomaly detectors should also factor in these differences.
- Performance of anomaly detectors:
 - Individual vs. averaged thresholds
 - Standalone vs. collaborative
- Conclusions:
 - Synchronize launch across all organizations, synchronize PR in all media. → more complete dataset.
 - Appeal to individuals' sense of good will ... or entice with cash!
 - - participation can help save the world
 - Participation can help one's own research/causes
 - Lobby diverse teams to scrutinize resultant dataset – diverse findings strengthen case that data collection should be repeated regularly by ID instead of a small research team
 - No shortage of interesting Enterprise phenomena.
- Legal Concerns:
 - Data collection: CFP is an “informal” contract
 - Required Opt0in participation
 - A-priori, stated intended future use, by whom, and scope of usage
 - Aimed to preserve privacy of user identity and packet content
 - Known risks:

- resource usage – storage, time, b/w, CPU, memory
- Temporary collection of packet content
- Data sharing – maintain contract, protect community
 - Potential risk to the organization that users can:
 - Reverse engineer network topology, network policies
 - Gain entrepreneurial advantage
 - Reduce risk

Future work:

Separate management from user traffic in analysis

- Delve more into volume, packet-level & user level metrics
- Quantify impact of context-aware local detectors?
- Launch another trial?
- Is there an end-host knowledge plane? If so, who can/should have access to it?

Q&A/Discussion

- David Clark: in the consumer market there is a difference between heavy/regular/light users. There is a question whether user characteristics persist over time or does this circulate among all users?
- Do you have a notion about how this is in the enterprise world?
- Eve: this information not available now, this is one of the goals for the upcoming evaluations.
- Q: ??
- Eve: also planning work on volume of traffic in the future work.
- Kevin Brooks - Motorola: interested in a tool like “Kenny” as part of a p2p study
- Eve: the same sort of activity Kenny is performing could be used in such experiments and that’s what could be used in such experiments.
- KB: were machines slower?
- Jaideep: Not noticeably
- Sollins: did they measure how much their measurement affected the measurement?
- Lehr: what would be good follow up to the experiment is to follow a single user around for a much longer time
- David Reed: one of my theories is that we don’t do enough measurements; this challenge of giving people incentive to participate may be the core of the idea. Any idea what would get people to sign a long time contract?
- Dirk: we also looked at the concept of participatory sensing. Especially hard when a mobile device is involved since there are other issues like battery life.
- Nadav Aharony: Was there any monitoring of the state of the device, whether it was being charged or not, as part of the logging?

- Eve: not in this experiment.
- Dirk – in our experiment they did monitor the charger status and only upload results when charge is on.

Dirk Trossen (BT) **Collective delusions behind how capacity gets shared**

- Collective delusions behind how capacity gets shared
- Freedom: to limit the freedom of others, or not?
- Two incompatible partial world views
- TCP: per data flow, instantaneous
- Operators: per user, over time
- This talk shows status on our attempts to unveil multiple
- Base example:
 - unattended apps: bittorrent, etc
 - attended apps: web-browsing, etc
 - unattended get disproportional share of traffic over attended (from 20x to 500x more)
- Realistic numbers:
 - 45 TCPs per torrent and 40/torrent
- Most users hardly benefit from bottleneck upgrade:
 - expect about 300k more by increasing the bottleneck link, but only get 60k in the end.
- Higher investment risk:
 - ISP requires everyone to pay for 300k more
- ... but we still see enough investment:
 - because of:
 - subsidies
 - weak competition
 - throttling
- trend towards bulk enforcement:
 - as access rates increase
 - operator choices
- so the internet way was wrong (and operators right)?
- neither side is right
- networks making choices for users:

- we have at “usage” rather than “users”
- choices:
 - 1) bulk: configure software
 - 2) selective: using deep packet inspection (DPI)
- DPI is not fair to apps that use the medium fairly
- There are better solutions than fighting:
 - slow down heavy users
 - allow light users to go faster
- BT's solutions
 - tactical: adaptive mechanism to assign different priorities to different applications
 - strategic: re-feedback, evolution of weighted TCP
- D.Clark - Nadav: Bandwidth allocation based on per-user vs. time over which bandwidth allocation is measured
- D. Clark: ECN-TCP flag shows whether flow needs to be policed or not
- Montpetit: we should encourage better design in network applications
- D. Clark: but apps are inherently not built to be friendly to the ISP
- Montpetit: can we give incentives to be “nice” to the network?
- Clark: Deep down Bob believes in the dynamic allocation of the medium
- Nadav: the ISPs provide as much bandwidth as they can
- Montpetit: Instead of putting penalties all the time it may be better to provide incentives instead.
- Andrew Odlyzko: differentiation between provision for wired/wireless allocations, provide low transaction latency
- Intel: is it possible to provide positive feedback instead of negative to the users (participation of all stakeholders to the benefits of using the medium)
Reed: you need to classes of traffic (attended vs. unattended)

Wednesday, January 23, 2008

Context-aware platforms – How is our notion of communication changing?

Nadav Aharony (MIT)
Learning and Using Social Context in Communication Networks: Social Area Networking and the Comm.unity Platform

- Have the network more informed of our social context = better performance, better social interactions
- Why social awareness?
- People act differently in different social contexts
- This is a communication protocol – high level, loose

- Communication devices use hard-coded protocols and behaviors – every device treated equally regardless of owner
- Different economic incentives depending on who you're communicating with

- IMPLEMENTATION DOMAIN: Face to Face networking
- Find a domain where a problem is not solved
- Face to face networking – there is no good enough paradigm for digital exchange, still have to email something half way around the world – wireless devices still used to access the bigger network – it's like using FedEx to hand you a pencil
- MJM – Bluetooth is used already to do this
- NA – but not by most people, hard-coded apps. Our devices should be open in this space, open for developers
- The other motivation is to use geo proximity to boost other interactions, e.g., neighbor's blog about the neighborhood – locality important aspect of context
- Jeff – using the FedEx analogy – there's a big difference between FedEx cost and Internet cost structure
- NA -- Using someone else's resources, you have to know someone's email address re local sensing
- KS – the other point is that if you imagine this expanding – scaling problems hit you if everyone sends their stuff over the Internet
- KS -- Technology needs to match the circumstances
- DT – if everyone relied on Bluetooth, it would be simple –
- KS – even if we're using diff technologies, can we still find a way to communicate when in close proximity without using the big Internet?
- DT – most people are using various access networks to access the Internet, but still can't discover themselves – solving the problem of discovering everyone's devices, that's not trivial –
- Xx – your FedEx analogy is not achieving its goal – it's too vague
- Nadav – the demo allows you to share files with people around you, belonging to different social groups. Not a straight forward exchange
- Adding more modules as we go along, add more types of radios, e.g., Bluetooth, IR
- Rules for communication, e.g., using IR vs Bluetooth
- How do we learn social context?
- Reality mining project – 200 phones given to people for 2 months, dumped usage data, could rebuild the social network structure of the people, the MIT hierarchy, could tell who friends from you hung out with, what kind of day you're having ... with very high accuracy
- USER INTERACTION
- We're experimenting with different interfaces
- Distance is social distance from you – circles of trust – this has network implications – e.g., will not be bothered by strangers... as the system learns, increased interactions might enter closer systems, the user can drag people in an out of circles manually
- Network configuration will adapt to user groups
- The extended self – all of my devices – will have the closest trust relationship with me
- SOCIAL DASHBOARD –

- Re interaction design – very important part of the system – bad design could harm social interactions
- Try to make the platform as generic as possible
- Written in Python, Flash
- GENERALIZING SOCIAL APPLICATIONS ON THE GO
- Design concerns – online presence & social network merging with real world social networks

Polychronis Ypodimatopoulos (MIT) Parallel Internets and Ultra-Local Economies

- The project is about organizing presence, social interaction people in the same geo area
- Quick look at Facebook – most messages sent to about 4 people, I don't see them everyday, most live in foreign countries
- Paradox about Facebook – maximizes at about 9pm and then drops during the weekend – does the social interaction btwn students starts at the dorm and end at the weekend? Doesn't make sense
- Diff between messages sent between summer break and winter
- Hints that Facebook is more of a communication platform than one that establishes and maintains social interactions
- Social interaction based on physical proximity but Facebook type platforms offer a virtual location
- By focusing on virtual, miss out on hidden,
- E.g., laptop users in a café can't discover each other
- Two strangers at the airport take a taxi to the same place, but don't know, take separate taxis
- Our approach is to build a mesh network of devices in close proximity and have those entities share profiles
- If we put this info together useful, can generate useful results
- Making two basic assumptions – user carries wifi in every device and regularly updates profile with current needs, moods, etc.
- Could have multiple profiles depending on context, e.g., at work, on the street, in the supermarket
- Updating the profile isn't something users will easily do
- Solution – cerebro – have a set of nodes in a common physical area – cerebro allows users to establish a network – user is in the center, different levels of resolution... helps in scalability
- Given that users share profiles – serve queries from other users about our own profile
- Always a path to discover and communicate with people and
- This is a mesh network of nodes in the same area – main gain is to interconnect mesh networks into a large network, by means of Internet, can add more info that we can query and therefore more value in being part of any of those networks
- So the system allows us to detect presence of nodes, but also a gateway to access hundreds of more profiles
- What are the biz models that can emerge from a scenario

- SMS \$1000/mb – if our mobile phone could discover other nodes through a mesh network, we could forward all this data – through the mesh network rather than through the operator. The operator is crucial here,
- What other opps are there for the operator in this case?
- Reed – suppose a taxi carries an ad like this vs what they normally carry –
- Prevent my device from advertising my presence to certain people
- Discussion

Bernardo Huberman (Hewlett Packard Labs) Social Dynamics in the Age of the Web

- With the advent of the internet we can study things that have been hard to do before
- We have much more precision now, we can generalize much more easily today because we have so many measurements.
- Analyzed 4.2 users from 500 schools on Facebook. To discover who is communicating with whom and when.
- Empirical analysis of real world examples lead to the discovery of social dynamics on a massive scale, including consensus formation, collective categorization, and temporal patterns.
- There are interesting patterns throughout the day or the week.
- In the age of the Internet, the scarcity, the bottleneck, is in our attention.
- The economics of attention
 - How do we present the most salient items while taking into account the visual real estate available on a given device
 - If a user can only attend to a finite number of items in a given time interval, what should be presented up front.
- In information rich environments, the scarce resource is attention. Its very valuable, ephemeral, and hard to explain.
- Getting the attention of a group
 - Broadcasting
 - Virality – rumors, communications
- Did work on how to get the attention of a group in a viral fashion. Got data from Amazon, looked at 15M recommendation data.
 - Viral marketing – the network dimension. Looking when users actually purchase based on peer recommendations.
- Attention – the temporal dimension
 - We share with others what captures our attention
 - Analyzed digg.com
 - Looked at 1 million users. Got a fit of attention to a log-normal distribution.
- “The decay of novelty”: Story half-life - ~69 minutes.
- Currently doing a study of YouTube.
- People influence each other. An immense number of opinions about things are formed by accessing the web. Can we get a fix on how opinions form over time.
 - Public opinions: an expression of the collective intelligence
 - How do opinions form and evolve
- A paradox – why do people bother to post opinions? Why do people vote?

- Looking at voting patterns in different situations and websites.
- Studying the influence of one vote/review on the end result or on others. Trying to measure the influence in analytical terms.

Q&A/Discussion

- Huberman: In these kinds of studies we are not looking at the motivations of the participants. It's like studying traffic – we don't know the reasons for people's driving motivations. But If we look at the number of cars we can make inferences of how traffic behaves as a whole and how people behave as a group.
- Brooks: you don't know the cultural background of the people you are studying. Have you seen patterns related to culture?
- Bernardo Huberman: we have not done these studies yet.
- Karen Sollins: we need to see how technology affects the social interaction
- Huberman: This whole issue of how to configure websites and how to present information to users is educated directly by this research. For example we have an algorithm to reorganize links in a site based on the results of such studies.
- We have a tool that we use within HP labs which is called "Watercooler" Many people in HP blog, the tool collects all the RSS feeds of these blogs and uses this knowledge to display the information in appealing ways.

Karen Sollins (MIT) Privacy and Security in Context-Aware Platforms

At MIT we can look at trajectory paths of technologies and try to analyze/predict where it is going. I come at it from a security/privacy perspective.

Overview: pulling on several threads:

- Evolving communications paradigms.
- Evolving social model
- Evolving security challenge

Communications: E2E. How communications evolved:

- Point to point – letter/email, telephones, TCP connections
- Broadcast/multicast (print media, radio/tv, ip multicast)
- From source to destination, less about content

Overtime a change in perspective was that the other endpoint was not another person but a server or another type of remote invocation. It could be a cluster or server or a P2P system.

Intermediated communications: Loosing End to End. Things happen inside the network (forwarding, firewall, caching, rendezvous (e.g. for multimedia conferencing). Beginning to break direct, real time communication.

It's now about the content. WWW and URL. Time and Space separation – not a question of where and when, but more about the "what".

Pub/sub – information based. Specification of what something is, and consumers specify what they are interested at or subscribe.

Key Components:

- Information
- Set of publishers
- Set of subscribers
- Attributes: how to publish or subscribe
- Policies
- Trust model

This model can probably be simplified to achieve any of the other models, subsumes them.

The evolving social model: trust and security. Trust can be based on confidence in unmodified delivery. Confidence in correctness of the source.

And along came: forgeable email addresses, IP addresses are forgeable, worms, viruses, business opportunities, enterprise and other organizational controls.

Its not all bad stuff, but there are competing objectives.

Tussle: competing concerns.

- Why do we care?
 - Sharing
 - Cooperation
 - Exposed contention
- Can we design for it?
- Is it monolithic?
- Where are the control points?

Security challenges: trust model.

- Not universal: regional, topical → context
- Not binary or pairwise: scalable, commutative → value based
- Not immutable → evaluable, assignable.

We need to develop models of trust that are much more adaptive.

Current agenda of the security and privacy working group: exploring the nature of identity in an information based communications paradigm, as a basis for examining the nature and capabilities required for trust and security.

David Clark: I would like to review the structure of the working groups in CFP. In the handouts there is a list of the working groups and contact information. Value Chain, Viral Communications, Security and Privacy, which are represented here. Spectrum policy and Interconnection policy – they still do not have an industry chair, I and Bill Lehr are working in this space. With the launch of the new website we are creating a new category of activity were we can keep the CFP informed about activities that we are doing but have not crated a formal working group with regular teleconferences. Another category is going to be about longer term research dealing with Internet architecture, and we'll use this new mechanism to report on how its going. If there is more interest from the industry we can create more formal working groups. There is also an Optical technologies work group.

The Future of TV – What's the role of P2P?

Natalie Klym (MIT) Introduction

Trends disrupting the traditional TV industry:

- Enhancement & extension of traditional platforms
- new delivery platforms
- new end-user devices
- new content types

More complex and interconnected value-chains in the video landscape

2-high level categories for services:

- Carrier services
- online services

Services complement and compete with each other

Live TV still dominates... how and when will this change?

As video services evolve, the TV experience is being redefined:

- time shifting
- screen shifting
- place shifting
- Sharing/redistributing
- Enter social networking

Marie José Montpetit (Motorola) Community Networking: Getting P2P out of Prison

What is P2P?

P2P is a “loaded” word:

- Bluetooth devices do file transfers
- DLNA and UpnP
- Seti
- etc...

Why “Community Based”?

Just re-interpret P2P:

- Move the network functionality to the edge
- Exchange deploying more expensive uplink bandwidth for shared community resources
- Less people pay but they pay more
- Use hybrids: servers offer connectivity to outside services (eg PVR), while serving the community
- Defines multisided applications without the center

What is a community?

- “sneaker net”
- Home net

- WPAN, WLAN, WWAN
- Neighbors and family
- Global social network

Need to get away from the “lord of the boxes” approach

- no need to have the device that “rules them all”. can we provide instead a community of devices into our house? I think Yes.
- Harness the power of home devices
- Focus on Autonomics, Social Networking, Communities of interest
- Scalability of the architecture

Peering at Different Layers:

Physical Layer: big challenge in IPTV is changing channels. Out of 500 channels about 20 are ever watched at the same time. So reduce number of streams by peering at the neighborhood level.

network layer:

- advanced policy management
- discovery “2.0”: not just what they are, who they are and how to access them (heterogeneous peers, capability discovery)
- Local routing and load balancing
- Scalability: propagation of discovery information
- QoS

service layer:

- use social network as virtual operator?
- define common middleware across platforms
- feature

Business aspects:

what would make operators like Comcast adopt a community approach?

- uplink glut
- sticky-ness of their customers
- new provisioning and customer management
- QoS/ policy management

New method to distribute content?

Business Aspects:

Use of game theory to evaluate the impact of new entrants on incumbent's business

Sample result: Long primary performance acquisition delay allows the incumbent to catch up on entrant's ancillary performance

Conclusion:

- Evolve the network offering to the community
 - offer some new paradigms for the management of these communities
- P2P maybe key to mobile internet taking up.

Q&A/Discussion

- Reed: is there a huge amount of unused uplink information capacity?
- Marie: Yes. But operators have decided themselves to throttle it because of the cost of CMTS processing for this traffic.

Henry Holtzman (MIT) Community PVR

Community PVR

CE tend to be more complicated, poorly connected and blind to the users' needs

CE2.0 initiative:
Make it radically simple!

PC vs. CE platform
No solution for browser, GUI in CE devices (PUI is another issue)
PUI: too many inconsistent controls

Laws of Simplicity:

1. reduce
2. organize
3. time
4. learn
5. context
6. emotion
7. trust
8. ...

iPOD example: extreme simplicity won over price, features, storage

CE 2.0 = simplicity + Web 2.0 + CE + intelligence

Exemplars

CE 2.0 initiative roadmap

CE 2.0 initiative members

Subtopics:

- simplicity
- awareness
- security and privacy
- connectivity and interoperability

examples:

- Tivo remote

- structure (product + content delivery channel)
- information

Awareness:

- better citizen
- contributor
- discoverer
- learner
- self-preserving

Security and privacy subgroup

Connectivity and interoperability

Prototyping

Television meets Facebook: social networking via Consumer Electronics:
Connect CE with soc. networking applications

Application overview:

application on Facebook

we get data from the web, information from within facebook about people's interests, what they watch and their ratings

DVR and UI functionality:

Evaluation:

in a month or so

- did it work as specified? enjoyed using it? etc...

Q&A/Discussion

- ????: have you developed metrics for engineering the complexity?
- Henry: good idea, but don't have metrics for simplicity yet.
- ????: robustness requires complexity
- Henry: complexity is under the hood, not necessarily exposed to the user

- ????: disability access: ties with simplicity?
- Henry: we have not explicitly dealt with that.

- Kevin Brooks: there are some metrics for simplicity, but not symptoms. Mainly done through usability tests.

- ????: authentication for a shared device such as PVR?
- Henry: this is hard area and there is no single solution right now.

- Reed: it's remarkable how CE community has failed to address simplicity in terms of local proximity which could contribute a lot towards achieving simplicity at communities of devices.

The 700 MHz Auction – What role does wireless play in a liquid world?

Bill Lehr (MIT) Introduction

- Why the 700MHz auction?
- It's 62MHz of prime spectrum that's being freed up due to the transition to digital television.
- People tend to think of dollars/MHz-population.
- When you look at mobile companies, numbers of more than a dollar or even 4 dollars are common. The most recent big auction that happened in the US was only ~\$0.70 /MHz-POP.
- Starts January 24th.
- Potential implications for Telecom/Internet value chain?
 - National broadband competitor
 - Ubiquitous coverage (“digital divide”)
 - New (more open?) platform (“Net Neutrality”)
 - Google? Frontline?
 - Convergence
 - Satellite
 - Cable
 - 3G/4G services: AT&T, VZ (roaming)
 - Public Safety national broadband network
 - Emergency response coverage
 - Interoperability
 - Govt/commercial sharing

“D Block”

- National broadband network to interoperate with and support public safety/emergency responsiveness
- 10MHz with reserve price \$1.33B
- Frontline shuttered doors on Jan 11, 2008

In the policy space: FCC Commissioner Michael Copps: “Wireless needs to be as open as the Internet...” and some other statements about opening up the wireless space.

Wireless research questions of interest:

- What is the value of spectrum?
 - Elasticity of demand?
 - Value of flexibility?
 - Economic impact of BB, wireless
 - Issues: band planning

- When will the spectrum regime no longer determine industry structure?

Building The ecosystem: making wireless resources liquid

- Observation: need to access intelligence and resources at the edge
- Challenges: unbundling the wireless edge
- Spectrum management reform
- Frequency agile devices
- Secondary markets
 - Liquidity!
 - Transaction costs

IEEE DySPAN – IEEE conference where those issues will be discussed.

Dale Hatfield (University of Colorado) Applying New Technologies to Old Spectrum Management Problems

Comment on Bill's presentation: I think that if you get the D band you can negotiate with the public safety people to get additional spectrum from them when you are not using it.

Quick review:

- Major steps in spectrum management (Things that the FCC does): Allocation, Service Rules, Assignment and enforcement, primary and secondary status (secondary user can't harm the usage of the primary user).
- Agencies responsible for spectrum management
 - Federal Government, and FCC that covers non federal use
- Traditionally there is centralized "command and control" approach to management. It's very hard to do from a centralized point. Not really compatible with the rapid technological changes in this industry.
- Pressure on resources:
 - More users, more uses, greater bandwidth
 - In the past how did we handle this?
 - Reallocation – take from someone, give to someone else. But prices have become too high to do that nowadays
 - Move higher in frequency
 - Increased sharing – people don't want to share, afraid from interference.
 - Improved technology
 - Focus of this talk is on increased sharing of the spectrum. There are 2 types of sharing: cooperative by the 2 parties, and involuntary, or even mandatory, as enforced by the government.

Constraints and criticism of the traditional approach:

- Excessively rigid

- Stifles technical and service innovation
- Lacks incentives for efficient use of the resource
- Creates barriers to sharing
- Erects barriers to other beneficial agreements and transactions. Companies couldn't do deals because of limiting regulations.
- Invites rent seeking behavior.

Rigidities in the allocation, allotment and assignment of spectrum: spectrum is not really congested. More dynamic/decentralized approaches to managing the resources were often hampered by technology and equipment limitations. Under utilization due to administrative scarcity and equipment constraints has been verified by recent spectrum measurements.

Challenges:

- Reducing rigidities in the current system
- Taking advantage of advanced technologies
- Vision of a more flexible future

Competing approaches/Models to solve the problem

- Property rights /market incentive model
- "Commons" or "unlicensed access" or "license exempt" model
- Command and control/engineering model

Advances in wireless technologies:

- Technological advances for reducing past constraints
 - Software defined radios, cognitive radios, policy based radios
- "Intelligence at the Edge" and its implications

Tracing one thread of a reform initiative:

I worked on two inter-related proposals directed at increased sharing through market forces:

- Modifying/adopting rules to allow market transactions – a "secondary market" in spectrum, esp. lease transactions.
- Modifying adopting the equipment authorization rules to facilitate the regulatory approval of SDRs/CRs/PBRs

This allows marketplace forces to reduce scarcity through increased sharing.

Sharing is voluntary and potentially cooperative.

That was 10 years ago when I went back to the commission.

Status of the initiative today:

- It's now legal to lease spectrum. There is even a limited market for these transactions.
- The FCC has modified its rules to allow approval of equipment that has characteristics of SDR/CR/PBR
- There have been challenges and issues with involuntary sharing, but those are different issues.

- For cooperative sharing there are no longer regulatory barriers so such transactions but very few such transactions have actually occurred.

Why hasn't this worked? Possible explanations:

- Remaining transaction costs - the cost of legal negotiations may not be worth the revenue.
- Convergence – fear of creating new competitor.
- Insufficient number of sellers/leasers. Sellers prefer to keep it out of the hand of potential competitors.
- Exclusion of government spectrum. Its hard to find out who in the government has spectrum and there's no motivation for the government employees to lease their spectrum.
- Fear of permanent reallocation – sellers fear the leaser will go and ask the FCC to give them permanent allocation, since property righthood is not clear.
- Economies of Scale/Scope and first mover advantages make it difficult for any new entrant.
- Immature technology.

Final Thoughts: SDR/CR/PBR have potential to solve many of the problems associated with the spectrum, but incentives are still not strong enough.

Q&A/Discussion

- Andy Lippman: So why do an auction now?
- Dale: It might be the best way to still do it. How else would we decide? How do we decide who will get it – for example MIT or Stanford? Both are good universities. Other option – Flip a coin? Auction is a reasonable solution to do it.
- David Reed: The thing that puzzles me about these processes is that its frustrating that people want to get things done but at the end it doesn't happen. Since I am terribly focused not on metro or large area communication but on close vicinity communication, it seems to be that it's a simpler problem to solve in that space.
- Dale: Isn't 3650 (?) a little bit of that? I want to create an environment with incentives for voluntary sharing.

Andy Lippman (MIT) Respondent

In the past we had a program about television. It wasn't an attractive area for graduate students and researchers to work on. So broadcasters got together and set up an advanced center for television research. Computers at the time were also reaching the speed that they can start processing the TV signals. We thought about processing, encryption, hi definitions, etc.

The same thing is happening now with the intelligent radios. For a long time it was a boring area. Marconi already solved that problem decades ago. Radio was pretty dead, we knew what it was. Then the technology started catching up. Computers were becoming strong enough where they could generate the RF signal itself, not just the digital picture.

In those days I visited Sarnoff labs in Princeton, they had a special filter attached to the ceiling with huge capacitors as big as beer kegs to amplify a full power TV signal. Nowadays we can probably do all those things we could never dream about in those days. Technology and economics made it possible and now a whole lot of kids are interested in radio again.

If we are at the threshold of this great fundamental change, then why are we giving away the ground underneath it (with the auction)? This might be the academic perspective. But we need to think away if we want to give away our assets on the verge of such a change. I feel we didn't get the balance of that right.

Now I'm done with this response so I'll talk about some other things. I am on sabbatical now, in Nortel. I get to watch a lot of things. Listening to thoughts about AT&T coming back together and the IPHONE, things seem bad in wireless to me. People were drawing lines leading to HTDV in people's pockets. Nobody was believing these – not the presenters and not the audience. Then you look at the Iphone and realize those projections are never going to pass. At that time the Iphone was the most significant change to cellular telephony in a long time. It's the first device that is intractably hackable, what happened on the phone was under the control of the owner and not the carrier and there was no way to stop it. Second significant thing about the Iphone is who the customer is. When you buy and Iphone you are a customer of apple. You buy it at the Apple store, not from AT&T. You look to them for revisions, modifications, or upgrades. Your relationship is with them, not the carrier.

Now take this and look at the 700MHz auction. For the first time we have an auction which won't go to the usual suspects. For the first time the winner may not use it in the traditional way that we already know of.

There is nothing more fragile than a business plan that tells you what will happen 3 years from now. And its especially true today with the huge amount of spectrum and devices not going the traditional way. Now we generalize that.

In 2000 technology was the disruptor. In 2008 the disruptor is speed:

- Clock rate of society.
- Diffusion of ideas and collaboration
- Business plans go stale quickly

Even Verizon as the token for the tightly knit carrier, even they are making noises about opening up their network to programmable devices. That happened within 2 months since the Iphone and Google phone came out.

Youtube was founded in 2005. Within the 1st 9 months the TV world was changed. Facebook is 3 years old. Second Life started in 2003.

Spectral examples: Identity.

Some things that we are working on in the Viral Comm research group. We are slowly developing a program we call “living the future”, to give technology to the kids at MIT. We’ve been started to do the fundamental experiments towards that wireless future.

Identity: we are working on identity, but on a small piece of it. The component of identity that’s physically embodied in the things we carry around with us, like RFIDs. Today we have no control over what info is exchanged.

We want to move to an active identity, negotiated by devices. Challenge, communicate, cooperate.

The problem with digital and software radios today – power. Battery life doesn’t last long enough or becomes illegal to take on a plane. Job 2: interference.

But it’s easier to work in the RFID frequency range. This is why we work in that range right now. We also do work on reading multiple RFIDs at the same time. Can we build a radio that can take apart signals on the same frequency? The answer is yes, if you have the smarts to take it apart. In order to do this, you need a radio that runs twice as fast, as long as the 2 radios don’t occupy exactly the same degrees of freedom.

Using such technologies we can do cooperative sharing of the spectrum.

- Halifax/San Francisco cvs Katrina / VA Tech:
- E911, security services. I think we are doing these wrong.
- The basis is context: who can help, who should know, mashing up the real world.
- In 1906 there was a huge earthquake in SF. The army was nearby; they arranged supplies and help quickly.
- In Halifax in 1917 a munitions ship blew up on the harbor, middle of WWI. 3100 people died. Because the military was there order was restored quickly.
- In Katrina, everyone had communications, but when the comm. Systems went down the population became part of the problem and not part of the solution. Ambulances were routed with mapquest, but routes were flooded. Nobody could mark that the streets are blocked. Calling 911 is not the same in an apartment in NY. It’s been diverted to be thought of as a carrier based system than a social safety net. If you have a heart attack in a mall today and call 911, the owner of the mall won’t know that a cavalry is on the way to his mall. Also nobody would know there is a defibrillator 40 feet away from you.
- 911 could be used as an example of using communication to fix such problems. I don’t think these things are so hard.
- The message I got from the 2 days of this meeting. Instead of communications in the future being silos, communications is now a horizontal assembly of function. I call those context, and this E911 is also based on context.
- I don’t like the word “services” since service is something you get charged on. These are building blocks to help us form the applications that do catch hold like facebook and myspace and work in the silos. What’s interesting about those functions and those horizontal sub-services is that they are occurring at the semantic level and not just in the bit level. Queries, using human words of civilians.

David Clark (MIT)
Closing remarks/Wrap

- P2P is a signal of a major change, a reminder that function is flowing (Liquid Platforms, Reed).
- It is also a new way to introduce new services and apps -- there is a lot of computing power at the edge, enabling users to download software and give them a new capability, e.g., video, music apps, with increasing rapidity (Speed, mentioned by A. Lippman)
- But P2P is a very fragile phenomenon. If Time Warner creates new price tiers for usage, users will turn off P2P apps – in this way, pricing models could have a profound, quick impact on P2P
- That’s looking at P2P from the negative side, as a bandwidth hog. But if people are going to download music over P2P, there’s less load on the backhaul if they download it from their neighbor.
- As a Robinhood phenomenon, it has blinded many people to the positive side, preventing legit players from getting into the business
- Positive side – P2P can be a robust dissemination tool if we decouple who owns, and controls the content from how its disseminated, then you can build more robust dissemination tools. i.e., separate the identity of the info from how it’s sent
- This is going to happen, the community is pushing for that – that’s one way for P2P to grow up
- It’s important to distinguish between peer production of the content vs peer networking – structure of the underlying system that it’s delivering it. The design of the supporting system will be up for grabs over and over again, but peer production is here to stay.
- “One of the nice things about the telephone is that we don’t need Hollywood to script our phone calls.”
- Identifying sources of high quality content to jam through their services is wrong...
- We have to think about win/win models that enable ISPs to get some of the joint value. P2P is not just a caching mechanism, it’s an architecture that allows many players to participate. The key is to make it legitimate, to let the legitimate players join in the ecology.
- To the extent that it’s a flexible arch, we can’t predetermine how money will be made. We have heard a bit about the roadmap and the fundamentals including the collective production of the experience, rich computing, and the separation of how content is signed, protected, etc. and delivered
- What’s up for grabs is how to use pricing at the edge to limit what the end user does at the edge
- “Where’s the money?” is an important question. Peer production for example is valuable but there are not a lot of opportunities to make money
- Why are there centralized sources of peer production, e.g., Facebook? Advertising, capture of aggregated and personal id – ownership of the context – that’s where the value is – will they sell it? Can we pry it out of them?
- Or can you make money delivering this stuff... might be the best thing to do, it’s debatable

- I'm interested in questions that relate to the architectures Cisco is putting forward – how do those facilitate new relationships between actors in the marketplace so new opportunities emerge, e.g., something as mechanical as ad placement, demographics
- LTF – central to that is the idea of you as a participant, or as a device, have an ID in that system and protect it and how to deal with context shaped computing that addresses people's privacy concerns.
- Which actors have an advantage?
- Charlie – talked about service aggregators – roles for people who facilitate the construction of these services ...what's the relationship between the facilitator and the facilities provider? The Value Chain group can answer those questions.
- Interconnection can help – not in terms of wires but in terms of new ways for how the actors can interconnect to construct new services

- Questions to the group:
- This meeting had a somewhat coherent theme, although we played a bit fast and loose, but we did not have WG reporting sessions... do people like the format of this meeting?
- Trossen -- I liked that for the first time in a long time there are common strands between the WGs. Each group deals with the theme in different terms but closely related concepts – very promising – common area – economic and architectural aspects of this ecosystem
- I like the idea that the interconnection group is interested in looking at a collection of players rather than wires – that WG could drive the framework of CFP
- re the format – I did like having a theme, it pushed us forward, but I don't like abandoning the WG sessions entirely – keep a slot, maybe smaller than in the past, but still have working group reports
- Polychronis: Every preso mentioned SN and Facebook with a few exceptions. Maybe there should be a horizontal WG on social networking to cover common aspects for other “vertical” WGs.
- Tony – I agree with Dirk – I like the idea of modified inclusion of WG meetings.
- For the web site – how much interactive capability in terms of blogs?
- Dirk – can the members create or just respond?
- David – want a cap to invite members to write op eds – web site will evolve
- Dirk – comment re presos – don't email the presos – had to delete
- MJM -- A lot of us had similar presos, maybe before we have these meetings there's a pre-meeting – e.g., call last week to coordinate – we could have gone deeper...
- Clark – as faculty we can have slides ready in advance, but execution is difficult
- Dirk – request for having a good synopsis that connects the presentations