# CFP SPRING 2016 WORKSHOP PROCEEDINGS

April 13 & 21, 2016
Edited by Natalie Klym

## Infrastructure Matters, Dave Clark

The Internet has matured to the point where we can, and must, consider it as infrastructure, which means it has a social purpose. Google defines infrastructure as: "the basic physical and organizational structures and facilities (e.g., buildings, roads, and power supplies) needed for the operation of a society or enterprise." And like most infrastructures, it's something we take for granted most of the time. We don't notice it until it breaks, and consequently, it seems quite boring.

MIT loves innovation, and innovation is exciting! But infrastructure isn't, at least not in the same way. It can, however, be rather exciting when we realize it enables exciting opportunities. That means wonderful things like bringing economic opportunities to the developing world. It can also mean negative things, like censorship or government surveillance of citizens.

As infrastructure, the Internet is unique. For starters, it's much younger and more dynamic than other infrastructures like roads and water pipes. At my town meeting recently we discussed spending $16 million to replace the inflow water pipe from Nagog Pond in Acton. When we asked what was wrong with the pipe, they said, well, it was put in about 100 years ago and it's corroded. This is the first major upgrade in 100 years!

But more significantly, the Internet, here in the U.S., and in many parts of the world, is being built by the private sector. This raises two central questions: What is the incentive of the private sector to invest? And, will the private sector build the Internet that society wants?

The incentive problem is a long-standing concern of CFP. We published The Broadband Incentive Problem in 2005, which focused on the fixed Internet, and more recently, Mobile Broadband: Towards a Sustainable Ecosystem in 2014. Our key question was, what will create incentives to invest in the Internet, and what will it be used for?

Before getting too deep into the incentive discussion, I like to tell this story. In 1995, I wrote a report on the future of broadband for the academy of sciences. While briefing the telephone industry on this report, one executive said to me, "if we don't come to your party then you don't have a party. There's no business opportunity in selling broadband. The only way it will happen is if the FCC makes us."

First of all, he thought broadband was ISDN, but more importantly, it never occurred to him that the cable industry might come our party, and thus be a viable competitor. It was so NOT on his radar. Even when asked very explicitly about cable, he dismissed it. They would not allow us to brief them when the cable guys were in the room, so we had to have separate meetings. And when we met with cable, they said, let us tell you our 10-year plan. We have the capacity to sell Internet. Initially it won't work that well, but people will think that's the way it is. But we will get

better at generalized packet delivery. Meanwhile, the telephone industry won't be paying attention to what we're doing, and by the time we start selling voice and high-quality products, it will too late for them. The point is that the telephone guy was living in the past, while cable was looking into the future. Our job is to point out to our CFP members what it might be like to live in the future.

It's hard to make the private sector invest in infrastructure. If you're a monopolist, then you can make too much money and you end up getting regulated. On the other hand, if there's too much competition, you may not make enough. So that leaves us with the public infrastructure space. There are lots of rural and otherwise underdeveloped parts of the U.S. where the private sector doesn't think it can make money, and there are public infrastructure projects underway in some of those areas.

The current "plan" for the Internet is to stimulate investment by praying to the god of competition, but that's not the way to deal with an infrastructure that supports most, if not all, societal functions. What we see today is actually a mishmash of competition and regulation. The FCC for example regulates price, and piles build-out obligations onto mergers.

An important part of the problem is the fact that Internet access is a commodity; there's not much product differentiation. Combined with high sunk costs, the market outcome could be quite dismal. But perhaps there could be more product differentiation in broadband?

Platform theory is one way to frame the question of product differentiation. Platforms provide a set of technical capabilities upon which other businesses—I call them complementors—operate. Firms are either platform builders or platform users. Platforms can be built and controlled by a single firm, or multiple firms, and likewise used by a single firm, or multiple firms. In the case of multiple builders and users, the platform might be open or closed (walled garden).

The Internet is actually a series of platform layers, each of which provides an interface to services that are used by the layer above. But it is the *stability* and thus *predictability* of the lower layer that invites innovation at the layers above. In other words, an app developer does not want to sit on top of an unstable platform. The layers above thus need the ISPs to invest in their infrastructure to make it perform better in terms of speed and latency for example—not to innovate in a way that is disruptive to their operations. So this poses a dilemma: How to create product differentiation without destabilizing the complementors.

Another point I'll return to later is that not all the layers necessarily qualify as infrastructure. For example, is Facebook infrastructure? On the one hand, we think of it as an app we use for fun and entertainment vs one that serves a social purpose. On the other hand, more and more people use it as their primary form of personal communications. After Hurricane Katrina struck, Facebook was used to communicate with family and friends as well as a source of critical information for shelters and other types of assistance.

With regards to our second question, will the private sector build the Internet that society wants?, one answer is that the Internet will simply be whatever the private sector builds. I was told by a member of the private sector in 2000: "you are not allowed to have a vision for the future of the information infrastructure, because your vision may not align with my company's

vision." The Internet is therefore going to be whatever the most successful company builds. And that's the only answer the private sector will tolerate.

If, however, we believe that society should have a voice in defining the Internet, the question is, do we even know what we want? And who is "we"? Society? The government? The complementors? Last year, I complied a list of aspirations from various sources. The list was quite long (you can find the full list here) but I've included a few items here that give you an idea of what kind of aspirations I'm talking about.

- Available to all users, everywhere
- Low cost, but with constantly expanding capacity
- Open to all complementors and content (net neutrality)
- Trustworthy, secure, resilient
- Respectful of local values but a platform for a global world

These are all these wonderful goals until you realize: we don't know how to achieve them; they are inconsistent; and we can't see why the private sector would want to consider many of them. The interplay between what the private sector would like to do and what society's aspirations are is an interesting space.

This list leaves us with several questions. First of all, which of these aspirations relate to the parts of the Internet that we consider architecture. A related question is, exactly which parts of the Internet should be considered infrastructure? Recall that earlier I asked whether an application like Facebook should be considered infrastructure. And from there you need to ask what is "critical" infrastructure, which implies government involvement.

Second, which of these goals need to be applied in the same way, everywhere for the Internet to work? Back to my earlier point about commoditization, maybe there's more product diversity than we know about.

And finally, what are the policy tools that can "nudge" the Internet towards one particular aspiration or another?

## The Chart, Andy Lippman

I'm quite motivated by some of the words David used, like "infrastructure" and "nudged." I found a paper about the aspirations of 5G that envisions the future as "a networked society with unbounded access to information and sharing of data which is accessible everywhere, every time, for everyone." That sounds good, and everyone in this room would probably see that as an aspiration, but there's more to the sentence than that.

As Dave was saying, it's not just about the technology that connects us to everything and everybody, it's also about the society and the economy and policy decisions we make with respect to that technology. Someone is in the middle of that network, charging a lot to use it, and that has an impact on what grows around it and how. So, I thought, let's look at that process. Let's look at those different tracks together—society, policy, economy, and technology,

and put anything and everything that we can think of that relates to communications on that chart and look for the relationships; how do these different elements dance around the chart.

For example, a lot of people are thinking about things like the block chain, which I will talk more about later. It's a highly distributed way of dealing with information. But let's forget the technology for a moment and look at society. After 9/11, people put a lot of faith in centralized systems because they kept us safe. But Snowden's revelations in 2013 not only undermined our faith in centralized systems, they caused the world to think about and be more open to distributed systems. That's not to say that the world post-Snowden will persist either; the pendulum will swing back and forth and maybe settle into something in the middle. Another example of this dynamic between society and technology concerns MPEG. My group's contribution almost got me thrown out. We wanted to alter it to make it amenable to streaming; then Adobe came up with Flash, which led to an economic change in that it enabled new types of services in music and television. It's important to understand how you're being nudged.

And it's important to understand infrastructure. If you ignore infrastructure and communications, and think it's solved, or someone else will solve it, or build it for you, then you're the ox that is going to get gored because unlike public infrastructure, like roads and water, private companies are building the Internet. When they buy spectrum, they think they own it! And they can use it for whatever purpose they want, which may not align with the purpose you want. So I want you all to leave this room realizing that having such a diverse group of people in the room, having this discussion is of great value, regardless of what business you're in, or where on the chart you think you fit.

## DISCUSSION

Q: With regards to how the presence of a cable competitor for broadband drives competition, a study by Prof. Christopher Yoo at the University of Pennsylvania using deployment data from the EU and the US showed that in European countries that have cable deployed, telcos tend to upgrade to higher speeds using VDSL, G.fast and FTTP.

In Australia, the telephone company owns the cable company. The government thought it would be more efficient for deployment but forgot that cable and telco competition is what drives deployment. This means in Australia, the cable company doesn't offer HSD products bundled with video. It's an extreme example of gov't decisions micromanaging the industry structure.

The other point is how product differentiation in high tech industries drives competition – that means investment driving innovation which in turn drives product differentiation. For example, from my perspective, if there are restrictions regarding competition, e.g., ISPs not allowed to offer certain products like managed services inside the home by leveraging data to identify how bandwidth is being used or offer certain price models, which is a form of product differentiation, then that will lead to less competition.

DC: Turkey is another interesting case study. The government is trying to decide if it should privatize the cable industry, and they're having this exact same conversation. Will the private sector invest? Are there incentives? On the other hand won't competitive pressure be good for Turkish citizen?

Q: What are some of the technologies that have had potential for a while but haven't reached mainstream adoption?

DC: I mentioned VR earlier as one technology that still requires more bandwidth. The more general observation here is that infrastructure has always lagged behind the applications. This is a much longer conversation, but Steve is going to talk more about this after the break.

## The Superfast Internet, Steve Bauer

What are the implications of a one-gigabit broadband world? People are incorrectly assuming that there won't be anything new to do in 10 years that requires this capacity.

It wasn't until Internet access went to the home, to consumers, that application innovation was unleashed. It was hard to predict bandwidth needs because until that time, it was used almost exclusively for research by academics and the army.

I believe there will be a huge wave of entrepreneurial and creative activity spawned by gigabit broadband. For example, all the video that is coming off of car DVRs combined with the push for more real-time capture of life in general will create opportunities for new services. The fitness market is another area of opportunity, e.g., videos will monitor people as they exercise at home for AI-based fitness training.

On a technical level, performance bottlenecks will move from access to Interconnection points, or maybe outside "infrastructure" altogether. What we are able to do, and at what speeds, has implications for how we do measurement, and for how measurement informs policy, e.g., what should basic/universal access be? What kinds of applications are "critical"? What are the policy, economic, and technical implications of these changes for your business?

### DISCUSSION

- Q: Is there an end in sight to the seemingly endless growth of bandwidth?
  - DC: Our predictions are never correct. The more important question is, are there other limitations besides technological? For example, battery power, or the cognitive abilities of humans. (However, there are technologies that are augmenting human cognition.)
  - SB: Even if it does stop at some point, it won't be before the one-gigabit transition.
- Q: Is the main driver of speed the arms race between providers? Bandwidth needs or demands? Technology progress?
  - BL: All of the above. Anyone who invests in new facilities will build to meet demand; bandwidth needs for new apps like video/RT capture of life will drive those needs; finally, with regards to technology; transmission, storage, and computing are substitutes. If the price of one of those goes to zero, then demand will shift in that direction. If someone says I can do it all with computing, someone else will say, I can do it with transmission. The question is, will we cycle around these different models?

- Q: This discussion is focused on fixed connections (home and business), but what about mobile: 4G, 5G?
    - SB: The same issues apply to wireless.
    - BL: 5G is forcing convergence between wired and wireless. But there's also an interesting pattern in the relationship between wired and wireless. A lot of the developing world is saying that they are leapfrogging directly to wireless, and the assumption is that wire line will never emerge. But that may be wrong. Instead we may see mobile first, followed by wire line to the home as demand for high bandwidth usage goes up. So it's a case of reverse evolution rather than leapfrogging
- Q: Qualcomm – why do I need a gigabit on my 5G smart phone – or do I? What kind of bandwidth do I need for a mobile device?
    - BL: You could say, I have a gig to my screen at home and mobile is just a thin client. But there are things you might want to do with your phone, and maybe not even in the digital space, but analog. Or, I might want to be running lots of different networks simultaneously and switching them.
    - SB: I told this story about the home vs work because it was hard to predict what was coming until the users were hooked up at home, and then those forces were unleashed.
- DC: Instead of asking, why do I need a gigabit to my smart phone? I think one should ask a different question: why should the capability of the radio wave be a gigabit? It may serve a lot of purposes other than the smart phone. It may very well be that it's the industry wanting to deliver video to my 4K TV. That may not seem realistic, they may not see a future of the industry they're tying to control being limited by access to smart phones.
    - Q: My question was very specifically about smart phones since my company makes money with those devices. And there's a debate going on right now about precisely that – what do you need on your device to sell it? There are a lot of people struggling with that question at our company.
    - We oscillate between the thin client/cloud model – this is the 4th time we've tried it—we call it cloud now – from the pov of the device, there's a trade off between how much app level state is in that device and how much is downloaded on demand, and if the network really isn't always there or fast enough that your transient downloads go faster than you notice, it might lead to a world where your device is just a display with a radio or two or three, with no storage. But that depends on the ubiquity and speed of wireless connectivity.
    - BL: Dina K is working on wideband sensors that she thinks can grab a gigabit/gigahertz for under a dollar. That means full bus speeds at short distances.
    - Q: Physics dictates that
- Andy: My answer to that question is also an alternate question. Let me explain. When you put the question like that (i.e., why do I need one gigabit to my smart phone), it's implicit that a smart phone is connected to an access network. The technical challenge is how much information flux we can get through a space in general, for example to beam 100megabits across my living room to my tv set, or trying to run a seismic experiment where I need a low level of latency among a lot of different things that are communicating directly. We may not need a gigabit to the phone to do the same things

I'm doing today, but we surely need to develop radio to get more than one gigabit/second through spaces in general, including access systems, local, systems, and longer distance ones. All will have to interoperate together.

- Q: At TI, we launched VOLTE and the bottleneck is access radio. If we have 100 users in one cell the throughput is reduced, so maybe this is the answer is. We have to augment the VOLTE technology for 3GPP to ensure QOS to the user to equal that of switched telephony
- NK: Re the high level question David asked, what do we want as a society, how does all this "performance" or capacity get distributed equally throughout society, or does it?
- This is the leading edge of what I think is going to happen for folks like you and I. But with regards to the other end of the spectrum, where people are just coming onto the network, I have aspirations in that area, but my focus is here because this is where the technological challenges are. Chintan has explored apps for the new sets of people getting onto the Internet.
- Karen: rather than looking at individual apps we need to understand categories based on resource requirements, what are the implications of having storage and computation in the network. We can't predict specific applications.
- But it doesn't keep us from trying. And Bill and I are interested in things like community clouds; it's different than what's provided by you're your access provider, it's computation and storage and communication that is somewhat local.

## 1. A Separate Internet for Video? Part 1 – Bill Lehr

This talk is really a thought experiment based on Cisco's prediction that video will be 80% of Internet traffic by 2019 (and most of that is entertainment video…).

We say that everything over IP is a good thing, but everything over IP need not mean everything on the Internet. The technical, business, and policy requirements of entertainment video may not mix well with those of the remaining 20% of applications, which are more about communications (for information gathering and control/decision-making).

From a technical perspective both sets of applications have very different QoS requirements, different security and reliability requirements, different privacy requirements, e.g., privacy in terms of what I watch is of a whole different order than problem privacy re my healthcare records. And it's not clear that you're going to get efficiencies from a single network.

From a business perspective, most of this video traffic is entertainment "content," where the economic concerns re differentiation and choice are key. With communication services, issues of interoperability are the problem. The concerns of entertainment and communications often overlap but different things motivate the technology requirements for each.

From a policy perspective, each regulatory sphere has policy issues that are domain-specific with little overlap. For example, broadcast regulation (e.g., domestic content requirements, free speech, domestic content, etc.) and telecom regulation (e.g., E911, wiretap) have nothing to do with each other, and bringing them into a single regulatory framework is a mess. And then we have the Internet, which historically, has not been heavily regulated, but as these legacy networks converge, how are we going to map this world onto the Internet?

A second question is do we want the Internet to be dominated by entertainment's concerns? Entertainment is about "leisure" and "attention," which is a very different world than things like smart grids, green homes, etc. When we explore why broadband is essential infrastructure, it's not clear that entertainment is "essential."

In entertainment economics, price discrimination is critically important; you want the ability to price differently based on the nature of the experience (movie theater, home viewing, real-time, delayed, high-res/low-res, etc.) to approximate first-degree price discrimination, but when you talk about the Internet, price discrimination is a big bugaboo because it violates Net Neutrality. The right way to pay for the infrastructure might not be to have a lot of price discrimination (e.g., a poll tax with subsidies for those who cannot afford might be best way).

If we did offer video on a separate Internet, how might we do it? The FCC has a defined a new separate service they call "broadband Internet access service" or "BIAS" that is to be regulated as Title II telecom service. Perhaps they could use a similar framework and define a "video-over-IP service" or "VIAS" as a Title II service to be offered to deliver video over the last-mile network. This could be regulated independently for interconnection, network management practices, and pricing. Given the hullabaloo about BIAS and the OIO, such an approach seems unlikely, but it's an interesting idea to consider.

From a technology perspective, it's not clear we would save anything in terms of capacity by separating the traffic. Thus, the video network could be managed to a high utilization rate; but the network for the other 20% of the traffic might need a sufficiently large peak capacity to accommodate low-latency requirements and the (potentially) very bursty nature of the other 20%.

Finally, it is worth noting that the entertainment industry is a highly competitive, global industry, and it's interesting to think about it in it's own right. At the same time, the whole thing is changing. We're in this transition now between new and old media, as an industry and as an experience. We need to ask what is entertainment? And that's our next talk.

## 2. A Separate Internet for Video? Part 2 – Natalie Klym

The argument for "a separate Internet for entertainment video" is usually made using examples of services that essentially comprise traditional television programming delivered over the Internet—whether it's a TV Everywhere implementation (e.g., Comcast Xfinity), HBONow, or Netflix, etc. But those services only represent part of the online video story.

The other part has to do with services that deliver new forms of visual media that we don't typically think of as television or even online video, but which nonetheless compete with these other services like Netflix and HBO Now and Hulu for users and advertising dollars. That's the part of the story that I've been exploring, and it's been remarkably challenging to get people to expand their definition of television, particularly in terms of WHAT we watch, regardless of the device and network it is attached to.

We've noted many times that "TELEVISION" is a loaded word. It's been around for over 50 years and implies a certain experience in terms of WHAT we watch as well as HOW, WHEN, AND WHERE. I therefore try to avoid using it and stick with "video," or invent new words like "televisual media" to break the association with past meanings. But, it's actually a pretty interesting word if you break it down to TELE and VISION and realize that television is basically about delivering video, or visual images, over telecommunication networks. And that basic capability can be organized into a variety of products and services that do different things and have different definitions, and operate on the basis of different business models. The form that emerged in the 1950s—in terms of content, devices and networks—was one of many possibilities.

In fact, if you read the historical accounts of television technology, you realize that the original, pre-industrialized, concept of television was based on the notion of "seeing at a distance," or the visual extension of the telephone call. But, for very strategic reasons that are too detailed to get into here, television was industrialized as the visual extension of broadcast radio, which itself had emerged at that time according to a specific definition and business model (closed/centralized production of content, programming schedules, advertising based, etc.)

Since that time, we've seen many changes to television, e.g., new transmission technologies, the subscription and multi-channel based model, various interactive features, like the remote control, the EPG, time shifting and place shifting. Programming—what we watch and how it is produced—has evolved as well, in tandem with these shifts, but the overall model or experience did not really change in a radical or comprehensive way until about 10 years ago when YouTube arrived on the scene and allowed end users to upload content that they created (or ripped off) themselves. But YouTube was just the beginning.

Over the last few years, I've talked about services like Twitch, Vine, Snapchat, Instagram, Periscope, and Facebook as new distribution platforms for "television"; and Twitter, which has always been the back channel to live TV, just acquired the rights to distribute 10 NFL games this fall. Most if not all of these services are social media. So maybe the next conceptualization of television is as the visual extension of social media?

But, what is social media? A few years back, we talked about social networks like Facebook and Twitter as new forms of personal communications, replacing not just the phone call, but email and texting. But as they increasingly integrate video capabilities, they are becoming video distribution platforms for personal as well as professional content, including both traditional forms—like live football games over Twitter—but more significantly, the new forms that these platforms support, like Periscope's live-streaming, Vine's 6-second videos, Snap Chat's "stories" which compile a series of a user's daily Snaps to create a narrative. Following Twitter's announcement for NFL games, Snapchat announced a deal with NBC for streaming Olympic highlights. This is "secondary" or "ancillary" content but nonetheless, professional, licensed content. Another example is the Young Turk's "Final Judgment"—a news program that was created specifically for Facebook. The Young Turks have a YouTube channel, but Final Judgment was even more fine-tuned to what they saw were the unique characteristics of Facebook—shareability in particular (i.e., short, authentic/emotional, etc.)

For video publishers these platforms are attractive because they are mobile, engaging (built for sharing), and it's where the users are—which is what also makes them attractive to advertisers.

In effect, what we're seeing is the convergence of personal communications and publishing, i.e., content distribution services. This is all very new and there are lots of creative and business and other challenges associated with these models. (The recent tensions between Facebook and YouTube are reflective of some of the business challenges, as well as the fact that producers face the dilemma of wanting Facebook's reach but at the risk of losing control.)

Our current research challenges in this "social media" or "new media" space can be broken down into 5 areas:

TAXONOMIES: Taxonomies are the starting point for the other challenges. They organize the "media" space into categories of video services based on defining features, along a chosen set of dimensions. We have a basic taxonomy that distinguishes between "old" and "new" media, but breaking down services within the latter space is our current focus. Creating categories is challenging because new services and new features appear almost weekly. All content projects, including advertising, are experimental.

COMPETITIVE (AND COMPLEMENTARY) DYNAMICS: The "old" vs "new" media framework emerged out of our work on disruption outcomes, and served as the basis of understanding the competitive landscape, i.e., the competitive, as well as complementary dynamics between incumbent and entrant "television" services. (See The Ambiguity of Disruption: Discovering the Future of Video Content.)

DOMINANT DESIGN: Given the variety of new video experiences, one of our investigations concerns how the concept of dominant design (i.e., the emergence of dominant product features) applies to video. My colleagues at Sloan are looking for patterns regarding new features that suggest whether or if a dominant design has or will emerge, why or why not? These patterns should reflect whether competition is innovation-based or cost- or quality-based. They will also inform a discussion on whether we are moving to an era of multiple "television" service definitions and communication product (and thus industry) boundaries more broadly.

MONEY FLOWS: We're tracking money flows from traditional television space into the digital (in terms of subscriber and advertiser revenues).

NETWORK IMPLICATIONS: What are the (Internet) network architecture implications of the shift to online video distribution, particularly with regards to new/social media, for example: more user-generated (so could be coming from anywhere); more symmetric/interactive/real-time so less cacheable; new business models that make DRM and current payment models less clear or relevant.

To bring this back to the question regarding separate internets for different services or applications, I think the more general point is that, if we are going to look beyond convergence to separate Internets, we need to understand how services and apps are themselves evolving, and what their technical, economic, and policy considerations are. This case study is an example how to apply that kind of thinking.

## DISCUSSION

- Q: I'm surprised at how little you distinguish between mobile and fixed. I'm European and to me, everything except video entertainment is mobile. The fixed Internet has become a de facto back up network. Social media has moved completely onto mobile and if video does not work for mobile, it's irrelevant. In terms of "criticality" if fixed is down, we must move onto mobile. But if mobile is down, then we stop working.
    - NK: When I talked about new media, (social media) I didn't explicitly distinguish between fixed and mobile because I also assumed that it's primarily mobile. One of the reasons publishers and advertisers are attracted to social media as distribution platforms is because they are mobile.
    - BL: The access point is on a fixed network, so I need a lot of fixed infrastructure to support mobile. You want to make the relationship between the two seamless from the user's perspective, and if you want to make it seamless, then the better your fixed infrastructure, and the more coupled it is with mobile, the better.
    - SB: I'm surprised that you make such a distinction between fixed and mobile. My Google 5 phone, even though it's a mobile device, most of traffic it's sending and receiving is going over fixed networks.
    - BL: One advantage of separating the two has to do with the notion of critical infrastructure. If, for example, the FCC says everyone has to have a broadband connection, it's very hard to do that with a mobile connection if it has to be capable of delivering several channels of HDTV. But if the FCC doesn't want to fund video entertainment (i.e., does not consider it "critical") and all the video entertainment traffic is on a separate network, then it's more doable to do that over mobile.

- Q: Isn't it a bit dangerous to call video out as a separate content type and say it's going to have different QoS requirements? That video could change, or there could be metadata and other things. And also, you talk about video entertainment as "leisure," (not critical) who's defining leisure?
    - BL: There will still be a ton of video on the Internet because of things like video conferencing and some of the apps that Steve talked about. And there will still be large file transfer. With regards to your other point, people are thinking about mixing the capabilities of multiple networks.
    - Q: But it still feels like a step backward (from convergence)
    - BL: This is a thought experiment. Everyone bought into the convergence idea but we are exploring some of the limits of convergence as well as the paths to convergence.
    - DC – You can play this thought experiment at different levels. Bill said something very important – everything over IP is not the same as everything over the Internet – e.g., if you're a cable guy you can use DOCSIS to deliver a honking big flow of IP packets into the home, but conceptually, they would look like different services; they would have different regulatory regimes; they would have different requirements for resilience; they would have different manifestations throughout the network, but it could still be one pipe, one IP packet, one router. So one can play the game of asking this question at different levels. In other words, it doesn't have to be separate all the way down to the

physical level. But everything over IP is probably more the outcome than everything over the public Internet.

- Q: Is there anything analogous to Open Source that might disrupt this business, and get us out of this lock jam, this hegemony?
    - BL: People have looked into the physical un-bundling of the network (into components) and into services. Earlier work on pricing for last-mile infrastructure to potential multiple last-mile service providers found that it was hard to identify stable prices for a la carte that did not end up with need to price bundle to enable cost recovery of the underlying network (i.e., providers purchase capacity for triple play – video, voice, data). Otherwise, problem with operators defecting to the a la carte broadband service and using that as cheaper (and lower quality) wholesale service with which to deliver a competing (lower quality but cheaper for consumer) "triple play" alternative. Alternative is to look at the last-mile infrastructure as an investment in roads. Pay for it with a town 20-year bond and let folks use it by putting boxes on it. Do not try and recover cost of investment in service pricing. If you followed this "road model" then open source would be a logical way to go to make sure you had as many service/box options to choose from for use on the roads.
    - DC – one way to interpret Open Source in this context is a gigantic socially driven, collaborative activity vs the very centralized way we are currently building the Internet. A lot of people have dreamed of this approach because they're uncomfortable with the constraints that come along with the construction of the commercially driven infrastructure. They sometimes use the term—it's an oxymoron—the infrastructure-less Internet. I guess I'm pessimistic. It's really hard to imagine how that solution scales when you think of the questions of investment and incentives. And collective action in this case is far more complicated than the production of a piece of software. I don't find any stories that are realistic. And that makes me sad.

- Q: The inefficiencies of networks have to do with the way they have been sliced and diced in the past, but today we have ways to virtualize, to create software-based overlays that would resolve the bottlenecks and unleash the capacity.
    - BL: There's a lot of interest in that approach to architecture. It's not necessarily the right way to build all infrastructure, but even building parts of it that way may be enough to collapse the market, sort of what like PBXs did to CENTREX.
    - DC – I do think people are going to harvest the economies of integration where they can, but I look at things like the emergence of IP exchange (IPX)—a standard for the exchange of voice and xx traffic, which in many cases is being implemented using separate interconnect technologies—it may be the same network inside but there may be a different interconnect architecture, there may be places where it's integrated at the physical level, or separated, I think there's a space to act out this thought experiment in a more complicated way across the layers depending on the tradeoffs between costs and benefits of integration.

- Q: Who would pay for the separate Internet for television–the entertainment company? The content provider?

- BL: In some ways it's the model we have today: Comcast has a bunch of infrastructure in place and if they didn't get any revenue from TV services then they wouldn't have any programming costs either and they would be a different company--maybe smaller, maybe even more profitable. So, unless you bundle these services together you can't pay for the Internet you want. Some people complained about telephone poles costing money, others said, thank god they weren't free because they paid for the Internet. Other people are saying that television is paying for the Internet, but the question is, is it paying for the RIGHT Internet? So, this is a thought experiment. If we fundamentally change the way people use these services at the entertainment level, we might end up with something more community based, e.g., TV becomes about what happens in my neighborhood, or it's big files that are cachaeable with some addition of live real time broadcast. They are radically different traffic scenarios and if you put the alongside different architectures, that would be an interesting experiment.
- DC: Natalie talked about the fact that one thread of the research is money flows. One answer is to chase the money back. There are only 3 sources of money – consumers, advertising, and the public sector. I mentioned earlier that the public sector is doing build-outs in certain areas. There's only so much advertising money to go around, and the number of people going after the same pool of advertising money is going to increase, and dilute that pool.
- Q: Isn't the advertising money actually coming from consumers?
- DC: well, it's a fraction of it.

## 3. How Private is Your Privacy? Karen Sollins

To a certain degree, my work has been living in the backwater for a while but the thunder is coming! In the more public arena, we've had cases like Snowden and Apple vs the FBI. In civil society, although this started a long, long time ago, the tension between the rights to privacy and the expectations or requirements for surveillance is growing rapidly. And in the industry arena, companies like Facebook and Google are providing options for consumers that they claim will enhance their privacy, like all kinds of confusing settings (not sure that's the right way). And there are more and more services claiming to give you privacy upfront, like Duck-duck-go, Bitcoin, etc. Whether it's at the app level or lower-level communication services—it's all about data; how it's collected, how it's used.

We are looking at the problem from the perspective of what mechanisms are possible; what technologies are coming along, what are their capabilities. We've divided up the space into the data stakeholders and the data life-cycle.

The stakeholders range from the subjects of the data, the data collectors, curators, analysts, platform providers, all the way down to the policy enforcers and auditors. The data lifecycle looks at how does that data come into existence, what kinds of policies are applied to it, what kind of enforcement models do we have, down to data destruction.

Our focus is on the tradeoff between risk and trust: how do we measure and evaluate those dimensions. There are several technologies that have metrics built into them, but the problem is

that we don't really know what these values actually mean. The challenge therefore is to take all these metrics and make them meaningful to the stakeholders.

Trust is a key question in this space. There will always be risk, so the question is how much do I trust that those risks are less threatening, i.e., how to achieve a balance between trust and risk?

**DISCUSSION**

- Q: The challenge for me is: I want privacy, but I also want the convenience of relevant info being pushed on me, e.g. Google tracks me all the time, but it tells me if my flight is delayed so it's worth the trade offs.
    - KS: yes – one of the hard questions is, we need to talk about how private is our information, individuals want to make that choice – I'll take more risks by exposing my data because I trust the service, or because at least I know what I'm getting into, but if someone else says, your privacy rating is 2.3 – what does that mean? Is that useful?
- Q: I'm a recent bitcoin convert, and for the first time, I'm trusting the math of cryptography. In my perfect world, I would have all my data cryptographically secured and pushed on a public block chain and I can temporarily give data access to individually people, I can revoke it, as I switch services--I'm not locked into a service.
    - KS: you would be comfortable saying you trust the math, but others want the soft squishy feel of "Google is my friend" – different people will have different needs with regards to how to build that relationship
    - Q: If you consider Lessig's four forces – architecture, markets, laws, and norms – all of those are far squishier than we thought, so even though we say it's just math, every talk I've seen on cryptography in the last couple of years is preceded by "to the degree that we understand number theory, we believe that…" That never happened before. We have a better sense now of how squirrely code is. Some people trust in code, some trust in markets, etc. There's no one perfect hook you can hang your trust on.
    - DC: Amateurs think you break the algorithms; professionals steal the keys – there's a deep truth in that.

## 4. Block chain: A New Infrastructure?

When you talk about Bitcoin and the block chain you often find an utter lack of knowledge or an utter lack of interest. It reminds me of talking about the Internet in 1995. The audience would always ask, what does the Internet mean to me? For example, Heineken said, I can't sell beer over the Internet, why is it important to me? I didn't have a compelling answer besides "It's the ante. If you're not on the Internet, you're not in the game."

Today, I'm going to tell you why I think you should be interested in Bitcoin. I'm constantly learning, and I change my mind every week as to whether it's something that's only good for illegal activity or fundamentally revolutionary in how we do legal things, or both.

First, there's bitcoin the currency; it's a monetary system that is unpegged. That means that there may be an exchange rate for Bitcoin, but Bitcoin itself is not tied to any sovereign-based

currency. That has not been tried before at scale. Bitcoin is an independent currency, and that's what's novel about it. And that's what ensures the integrity of the block chain. If it were tied to dollars for example, then someone could cheat and cash out. But they can't cheat because that would devalue what it is they're stealing. There are about 6,000 people maintaining the integrity of the ledger.

Bitcoin solves 2 problems associated with digital currency. The first problem is the "double spending" inherent in anything digital. There's always another exact copy, unlike something physical like library books. I can solve that problem by registering the transaction, i.e., writing it down into a ledger. The second problem is that you need to be sure the money wasn't counterfeited.

Bitcoin solves both those problems with the block chain. The block chain is essentially a bunch of people around the world emulating a central server acting as that ledger. It's a distributed version of that ledger. The 2 problems I mentioned above inform the design of the block chain, i.e., it records the value of every transaction all the way back to its origins so you know that coin was not wrongfully invented and it records every transaction so there's no double spending.

The other thing to note is that you can put a script in a bitcoin transaction, i.e., you can specify conditions like, you can't have the money until Tuesday. So that's the money part of it.

But we're primarily interested in the generalized value of the block chain. That's the more important concept. The block chain is a mechanism for dealing with things in a distributed way. It's a write-only, indelible, public ledger that records everything.

In the world of music for example, you have the performing rights organizations like ASCAP and BMI that register when people play recorded music in public places, and in some statistical proportion pay the authors of that music. Few authors of music actually understand how that works or if it's accurate, they just get periodic checks in the mail. But you could register their music on the block chain and automate that payment and you have a whole new world of automated international contracts.

If you imagine a global, public block chain, you could store fundamental ID information, and you could always refer to it without worrying about the integrity or security or reliability of a third party. Another fundamental point is that you can store the signature of a document on the public block chain. That separates the provenance of a document from its authenticity. For example, I can't give you my medical record today – you have to get it from Mass General because you're not going to trust that I didn't change my cholesterol levels—but if I store the code for that document on a public registry you don't have to care where it came from. The legitimacy of the document is ensured by its code.

These are the kinds of opportunities that you have with block chain. Are they real? Will they work? A lot of them begin with the assumption that a block chain exists, and it is maintained. And if you assume a block chain then there are lots of wonderful opportunities. It's a load of fun to imagine some of these. E.g., it facilitates a number of criminal contracts, e.g., assassinations. It's hard for criminals to operate in an environment of trust. But it works in good places too.

Bitcoin is just one use for the block chain, but it's proof that such a system can work. There's $6 billion in the Bitcoin system right now, and about $175 million changes hands every day. There are other currency systems like Etherium. DNS is an obvious case where a block chain is applicable. There are banking applications—the banking industry does not want to be Ubered. The banking system is full of incompatible legacy systems. If I told you how checks were cleared, you'd invest in PayPal tomorrow.

I don't care which one wins because it's the metaphor that's important. It's not about Bitcoin per se, it's about the fact that we've had a system that that's lasted 7 years that hasn't broken. It's woken up the field of distributed computing. That will have an impact on how we do things, and it may have security advantages because it there's no single point of attack. And that might be enough to dissuade attackers. Centralized systems like the U.S. government are under attack all the time, and they don't always win.

DISCUSSION

Q: You sound a bit skeptical about Bitcoin.

AL: Bitcoin has serious governance problems. In every other system like this (I'm thinking of the Internet) there was a governance system in place before it became immensely valuable. There's a bit of tulip fear because there's nothing maintaining it besides its user community. But, it's the public blockchain, so you have to root for it! I'll devote myself to Bitcoin if I can find a way to make it solve the problems of people who don't have money, and can't afford the banking and credit systems that are in place, and that's what some of my students are doing. So I'm skeptical but also hopeful.

Q: How does Bitcoin affect the money supply in the U.S?

AL: Bitcoin doesn't qualify as money. There are two things you need to qualify as money. One is that it has to be commonly and easily used in trade. You can buy stuff with Bitcoin, but it hasn't crossed that line yet. Right now the U.S. administration views it as a commodity, meaning it's traded but not regarded as a full part of the money supply. But it'll get there… something will get there. Digital currency is too good an idea not to be.

Q: But at some point, it will be counted as part of the money supply.

AL: I suppose that's true, but… there used to be these measures, but it's so synthetic now. That I don't even though if those are used. We don't track the money supply the way we used to.

BL: But it's still a trivial amount of money right now, relative to the total supply. Part of our interest in this beyond what it tells us about the limits of centralized management is the fact that you can value it -- demonstrates that it's not decentralized because the value comes from people going to exchange it.

DC: I want to share a more trivial piece of information as we wrap up. Andy talked about the difficulty of criminals communicating and trusting and that made me think of a book: The Codes of the Underworld: How Criminals Communicate. It's a charming book. The author studied the mechanisms for criminals to establish trust. He made analogy between how Italian academics communicate and how Italian criminals communicate.

AL: One more comment that brings us back to Dave's court. The anonymous guy (or group) who wrote the paper on Bitcoin was not Dave Clark, nor Vint Cerf, nor Bob Khan. I don't think he knew enough about networks, but he knew enough about consensus and he knew enough about computing power. But he didn't know about DOS attacks, latency attacks, computational attacks, network partitions, overwhelming TCP ports, etc. There's a lot of room left to figure out how to build an infrastructure and a network for this system.