Demystifying Networking

Contents

Foreword

I thank Andy Oram for this foreword:

For decades, the most advanced communications technology most of us experienced was a telephone call. *Telecommunications, narrowly defined as one-to-one voice calls and dominated by monopoly telephone companies, was the gold standard for long-distance interaction.*

Today we have one-to-one and one-to-many exchanges of audio files, video files, text, and graphics over the Internet. Almost everything about the Internet is diametrically opposite from telecommunications: the Internet allows both immediate and deferred interactions (also known in the field as synchronous and asynchronous communications). It supports every format that can be digitized. It is exquisitely elastic, allowing side variations is speed and responsiveness. And it constantly sprouts new companies and applications.

But we still get all these communications through the old telecom network. Despite stunning technological advances, the business model for telecom is the same one imposed on us nearly a century ago by government. The carriers, be they phone companies or cable companies, live within a tightly controlled regulatory environment with synthetic rather than real competition.

Bob Frankston for several years has been laying out another vision that combines some unique proposals of his own with elements found in writings by other leading technologists in telecom and the Internet. Bob probes the implications of digitization, packet-switching, and smart transmission policies more deeply than most of us can go. He suggests that the nearly infinite adaptability of these technologies will provide solutions to our bandwidth problems, if they can be freed from atavistic business models. And by owning the wires and transmission facilities in our homes and neighborhoods, we can take back control of the Internet and create new markets.

Andy Oram, Editor, O'Reilly Media

Preface

"Give a man a fish, and you feed him for a day. Teach a man to fish, you feed him for life."

Contrast this with the philosophy of a provider – if you teach a man to fish you lose a customer.

Maintaining today's telecom model based on network providers costs the world economy a trillion dollars each year, and extracts an even higher cost in terms of limiting economic opportunity and growth. It's not just about the money but about jobs and our ability to take control and improve our own quality of life.

I was reminded of problems with the subscriber-centric focus at a recent meeting to discuss a new operating system for mobile devices. I had expected a discussion of the benefits to the user and developer. Instead it was framed in terms of how much traffic it would generate!

Fortunately we have a simple and generative choice: Funding infrastructure as a whole and creating opportunity without rent-taking.

We need to understand the simple but powerful ideas that lie behind what we call "The Internet" – how do we make connections and exchange bits of information?

This essay is not about "The Internet" as such. It is about what the Internet signifies and where the idea can take us

and why the idea of providers has run its course and now does harm to the world economy.

What is the Internet About?

Before we can have an informed discussion about the Internet and associated policy issues, we first need to understand what the Internet is, both as a technology and as a generative dynamic.

The Internet makes it very easy to exchange bits (or messages) between two points without worrying about how they get between the points.

Let's take the simplest example – a doorbell. You press a button at the front door and hear a bong sound inside the house. All the button does is close a contact so that electricity can flow causing the bell to make a sound. In some homes the front button might make bong sound and the back door button causes a different sound – bong-bong. It acts as a signal or message to the listener.

This is the basic idea behind digital technology. Devices can listen for a message on the wire – in this case a one (bong) or two (bongs).

Instead of having two wires we can share a single wire with one pulse for the front door and two for the back. We use an alphabet of 0's and 1's. Thus 01 might may cause the front door bell to ring and 10 for the backdoor.

The 01 or 10 acts as an identifier (or name) in same way writing "Joe" or "Joan" on an envelope says who the letter is for. If you live in the same house you can just leave the letter on a table and Joe or Joan will pick it up.

In colonial America you would extend this system by leaving a letter in the tavern because someone would notice it and pass it on. This quickly became formalized using the postal system but we know that you can also deliver messages other ways such as when you give a child their birthday invitations to give to their friends at school. The system can get more complicated when you send a letter "care of" someone else.

You can save money by using less expensive services that are less reliable or spend extra to track packages so you know if they don't get delivered.

Despite the various combinations, the basic idea is that you put an address on the envelope and the message gets there somehow. You may not know the details but you know there is no mystery. After all, it's just an envelope.

Internet packets work the same way. You put an address on a packet of bits just like you put a postal address on a letter. If they can be delivered locally, like within your house, they get picked up directly. Otherwise they get carried on a path that takes them to your destination. Instead of postal sorters we have packet routers.

The basic idea is really that simple.

Actually it's even simpler.

Or at least the idea is simple. The complexity we see today comes from trying to duck and dodge the existing telecommunications infrastructure and its myriad of twisting passages and focus on creating "billable events". We accept this system because it is self-perpetuating complexity but it is entirely unnecessary.

Once we remove the mystery from telecommunications and see that the complexity is unnecessary we'll be able to benefit from the powerful idea that underlies today's Internet.

Exchanging bits only seems mysterious because we can't really see that what is going on on the wires is not very different from what happens with letters.

Speed also makes a difference. In Victorian times there were multiple postal deliveries throughout the day like an early form of texting. Speed it up some more and we get highly interactive web pages.

The surprising discovery is that we don't need a reliable delivery. If some messages don't get through that's OK. You can resend the message or assume that the receiver is smart enough to handle a few missing messages.

The counter-intuitive idea that we can build reliable systems out of unreliable parts is at the heart of what makes the Internet possible. It shouldn't be that surprising. After all we might build a house out of boards which, individually, aren't very strong, but as a whole become structures that can last for centuries.

We don't need to rely on carriers for reliability – we can do it ourselves and we can do it better because we can build resilient systems. We can also choose innovative strategies such as filling in gaps instead of waiting a long time until every packet arrives intact and in order.

This is how streaming music works – if we lose a few packets we can make a very good guess to make up for what's missing. If we can't then we wait till the message gets resent.

This means we don't need a high priced carrier. You can carry your own message or hire people who use the common streets. (See Best Efforts for more on the concept).

You don't pay for each message when you use the streets. Cities don't run streets for a profit. They provide them as common infrastructure for the community to use.

The same goes for public transportation. The New York subway system started as a for-profit business but when that failed the city took them over so that the city could function.

Telecommunications vs. Networking

In the 1800's telegraph companies sprang up based on Samuel Morse's new technology – the telegraph. The business model was simple – if you wanted to send a message you would pay a telegraph operator to send the message over the wires owned by the telegraph company. The messages had to be sufficiently valuable to cover this cost of the dedicated wires.

The word "communications" is associated with carrying meaning. Thus the provider has to take care to assure that the meaning of the (valuable) message gets through. It has to be valuable to justify the effort.

This was a fine business model. After all why would you want to pay for an unreliable service? We don't need a provider just to communicate within our homes or with our neighbors. A telecom company provided value in enabling us to communicate a distance – the "tele" in telecommunications.

With the Internet we aren't asking the carrier to transport meaning over a distance. All we ask is that it makes a best effort to carry the bits. Bits themselves, like letters of the alphabet, don't have their own meaning. The meaning comes from what we do with it.

We can send "t-k-s 4 g-d t-m". The carrier doesn't know if we are saying "Thanks for a good time" or "thanks for the four good times" or maybe something entirely different.

We take responsibility for the meaning. If a few bits are lost we deal with it ourselves and decide if we care enough to ask the sender to try again or fill in the missing pieces ourselves. There is little value for the carrier to add.

Today we use the same technology to communicate between a computer and a nearby printer that we use to send email around the world. Yet the language of distance is implicit in "accessing the Internet" or "last mile". These phrases date from a time when crossing large distances was legitimately expensive and difficult. Conditions have since changed, to the extent that these assumptions no longer hold true. Unfortunately, they still frame our policies.

The telegraph business model of selling high priced services to cover high priced infrastructure – the wires and radios – works fine as long as there isn't too much competition to drive the price down. The telegraph business model of selling services has stayed the same even as the technology evolved. In fact we still use the term "message unit" for voice calls as if a phone call were a series of telegraph messages. Today we charge for SMS messages and phone minutes in the same way.

This model lasted because the money was in services. The FCC was created to maintain an orderly marketplace for service providers transporting messages. It was modeled after the ICC (Interstate Commerce Commission) which regulated the transportation system using trains and trucks.

If all we need is best efforts bit transport then the carriers can't make money by selling us services like telephony. We can do it ourselves as with Skype. And we can do it better without the arbitrary limits that the carriers imposed a century ago in order to maximize use of their scarce resources.

We should've had a wakeup call when the ATT business model was failing in the early 1980's. They accepted divestiture so that they could make money in the lucrative computer business. Or so they thought.

But it still seems that we need a carrier to manage the network in order to assure reliable delivery and to manage the network as a whole.

What if telecommunications providers were not necessary?

We've seen that we can use the wires in our homes and offices to exchange messages between computers by just placing packets on the wire and tagging them with the address of the destination.

We pay just once for our wires (and radios) and there is essentially no ongoing cost to using them within our homes and offices. The same is true when we share ownership of the wires in apartment house or a neighborhood or a city.

The basic idea is very simple. A router is like a street sign. Your packet has an address and follows the path towards the destination. Without being restricted to a single carrier's path we are free to find alternate routes when one path is congested.

The seeming complexity of today's Internet comes from two sources. One is that many of today's routing decisions are driven by the need to deal with the maze of billing paths created by the telecom business model – a source of

complexity that disappears once we aren't trying to contain bits in billable paths.

The other source is our failure to take advantage of the fact that the wires and radios tend to be fixed. Instead we have to maintain complex routing tables inside the network to keep track of the computers using the network. It's as if the post office tried to track people instead of addresses of homes and other buildings.

It's important to understand how much of today's complexity of networking comes from dealing with the very idea of having to have a network. It's this kind of circular reasoning that makes it so hard to understand the issues.

We have a regulatory edifice (which I call the Regulatorium) in which every element depends on every other one. This is why it so important to get back to the basics.

And it isn't very hard. When a telecommunication provider argues for the importance of QoS (Quality of Service) or a "control plane", you have to ask what problem they are trying to fix.

Even if we accept the idea that we need network operators, how can they make money if all they are each selling is "best efforts" transport of bits. Bits are like letters of the alphabet and not consumable like gas or electricity.

The carriers face the same problem as farmers with too much produce – the price must drop below cost. It's a "natural monopoly" in that there is one system. But it is not at all a monopoly in that there is no exclusion – anyone can add their own wires and radios to increase capacity. There is no need to keep bits contained in billable paths. So there is no monopoly in the sense of the old Ma Bell.

What is a "Network"?

If we don't need providers and can use any facilities available what do we mean by a "network"?

We talk about the "road system" as if it were a real thing. But it's really just a collection of local roads (and other paths) that we use to travel. We create the illusion of a system by posting signs saying "US 1" but there is no entity that owns the entire system; at least not in the US.

We do have some federal guidelines but they are interpreted locally and you can create your own "roads" by driving across fields or whatever paths the vehicle of your choice is capable of navigating. There are also regional facilities be they state highways or the interstate.

The road system is this mix and you are the ultimate navigator. We fund this hodge-podge with various ways but we are funding elements of a system and not trying to maximize the profit for the roads themselves.

Unlike railroads you don't have a single operator nor any entity that guarantees a reserved path.

It seems obvious that in order to carry on a phone conversation you need a guaranteed path so that the sounds get delivered at just the right time. That's what the PSTN (Public Switch Telephone Network) is supposed to do but that approach is being abandoned as all telephony is starting to go over IP.

Yet we still cling to the traditional view of the phone company being like a well-run railroad. Today computer systems can schedule a train and ensure the path is available so that it gets there on time barring unforeseen events. A road is very different because there is no central controller. You are likely to run into traffic jams

The telegraph system is like the railroads in that you have to have a clear wire to send a message. The phone system also has acted like a railroad – you couldn't place a call unless you had resources allocated from one end through every element of the path.

You don't need to do this for text messages since it doesn't matter if they get delayed. By tolerating delays you can share resources rather than having them dedicated. This might not have been obvious when we used modems to repurpose the phone network.

In the 1990's when the Web exploded onto the scene there was no way to provide sufficient dedicated capacity. The phone companies sounded the alarm and warned us that the modem users were going to bring down the phone network.

Of course that didn't happen. Instead we gave up on the idea of dedicating resources for each connection. Instead the packets competed for passage. They might get delayed or even lost but we could deal with that.

What did happen is that capacity increased dramatically because we didn't depend on the network for each new protocol and could add capacity using any combination of approaches.

We need to take care to understand the concept of "end-to-end". It's easy to confuse it with womb-to-tomb where a network provider takes responsibility for the entire path. The term is used in just the opposite sense to mean that only those at the end of the connection understand the meaning of the bits. Inside the network you don't know if the bit is part of a phone conversation or an email message.

We treated any available capacity as an opportunity and found a use for such as exchanging email because there wasn't enough capacity for voice. This allowed demand, in effect, to create supply, because we found value in what we could get but as we got more capacity we found more uses.

One example was repurposing a system designed for cable TV and using it to exchange bits. The TV technology is called "broadband".

Again and again we have to be careful of our use of the language. Language itself works because we can deal with ambiguity but when we use a word like "broadband" we fail to communicate because it sounds like the Internet is just another broadband TV Channel. And there is nothing to say otherwise because the entire array of telecom services seems to work just like they did before.

But it's an illusion, just an illusion. There is no network in the old sense; just operators pretending to provide network services.

Our metaphors fail us because the congested data highway now has millions of lanes and packets tend to make it through just fine. We don't really have cars, bits are not things.

The network owners are not contributing value to the applications. They are just transporting bits using their facilities. It doesn't really matter which facilities are being used.

We've seen that the electric grid is a natural monopoly in the technical meaning of the term – it doesn't make sense to have multiple grids to distribute electricity. So how does it make sense to have multiple systems for sharing bits?

It doesn't. If you look up at the poles festooned with "broadband" wires you are looking at profligate redundancy held in place by nothing more than bad metaphors.

Yet another example is the idea of a natural monopoly. It doesn't mean that we have a single network like in the days of ATT when the seeming inscrutability was abused by ATT. They convinced the government that only ATT (AKA Ma Bell) could avoid chaos and make the system work). They got the government to hand over control of the phone system to them with a guarantee of profit.

The cost plus model encourages the phone company to buy the most expensive gear and then defend the purchase in the name of "quality". But what does quality mean? Does it mean a phone that will last 100 years or is it a phone that does what people want. A good example of this is the concept of "call completion".

The FCC defines a "completed call" as one that merely rings. It's a perfect example of naïve indifference to the larger question of why we are using the phone. To a user (a word that makes us forget we are talking about people) a call is complete when you reach a person or, at least, leave message. Yet the phone companies didn't allow answering machines until the Supreme Court overruled them in the 1968 Carterfone decision.

This story is repeated again and again because it is at the very heart of the concept of telephony. In 1956 they lost the Hush-A-Phone decision. They tried to prohibit people from putting a box around their phone! That was the extent to which the providers went to preserve control and dictate how you were supposed to use *their* network.

If the phone companies lose control of telephony they have no reason to exist. Yet the heart of our policy, as embodied by the FCC, is the idea that there must be communications providers because, well, there must be.

The carriers lost control of their network because of good engineering. The simplicity of the simple red/green wire interface of the phone system made it too easy to see past their just-so stories.

The need to build services into the network is what doomed France's Minitel. In the early 1980's France gave everyone in the country a computer terminal – a very advanced idea. But the phone company defined how the system worked and billed based on the number you dialed. It was meant to be a proxy for the value of the service but real services aren't so easily pigeonholed.

We saw a similar problem when the US phone companies tried 900's as a way to bill for services. The inflexibility of the approach limited it to services like porn.

Today we see these problems reflected in the inflexibility of the cellular services which are more about carrier offerings than users.

Instead of addressing these problems we have managed competition – an oxymoron. The government sets rates and rules – rules that are increasingly bizarre and out of touch with reality as they attempt to keep the world as it was over a century ago. Even as all the phone traffic is

going over IP both inside and outside the phone companies we have a completely fictional system of rules that keep us living in the age of copper wires and rotary dials.

In fact the rotary dial itself is a good illustration of the problem. A rotary dial worked by opening and closing the circuit so that a switch could step through lines till the dial stopped. After a short pause it would send the next digit. As phone companies implemented their intelligent network they looked forward to getting a fee for completing each transaction.

The "problem" is that they also introduced the touch tone keypad using tones instead of breaking the circuit. This meant that the tones could go all the way to the other end of the call without requiring the phone company to be involved at all! They could not get a cut of the action!

The situation repeated itself in the 1990's when the providers (both cable and phone) companies first introduced their triple plays in which they would bring a fat pipe into the house and charge you for each service. In fact they tried to charge for each PC in the house but were frustrated by the NAT (Network Address Translator) which prevented them. They also failed to prevent the use of Webcams.

Yet the carriers have *Terms of Service* which tell you what you are allowed to do with the bits. In effect they are telling you what you are allowed to say. This is because the business model is based on charging you for the value you create outside the network and how you use the bits.

We see this again in cellular pricing where they attempt to limit your ability to relay bits via your cell phone. You have to pay extra to use your cell phone to do navigation if the application runs on a separate device sharing the connection (as in the example above). But if the application runs on your cell phone but just displays on the external screen you don't have to pay extra.

There is no difference yet the phone company insists on telling you how you are supposed to use the connection you paid for.

We see the same inconsistency with ATT's iPad service. If you pay for an unlimited connection everything works fine until you attempt to download updates. Then you get a message saying that you must use a separate Wi-Fi connection. But if you exchange the same bits but ATT can't anticipate that you will use 20 megabytes you can use the connection. How did ATT manage to worm its way into the iPad in order to impose such an arbitrary and annoying roadblock in our path?

This is the today's version of harming the network. The complexity of today's cellular network makes it easy to make self-serving engineering decisions seem necessary yet the inconsistency is still too obvious.

On one hand you are told that there is scarce capacity so the companies have to dole out bits very carefully while at the same time they are selling new video serves over the abundant 4G capacity. At least they should try to keep their story straight. Depending on confabulation is a poor basis for an business model.

How can you create value if you have carriers constantly second-guessing everything you do? It creates an environment that is hostile to innovation for no reason other than to tie us to the rails of an eighteenth century notion that bits are transported like farmers' wheat!

Yet none of this is necessary. We needn't hobble ourselves by trying so hard to preserve the limitations of the past.

Remember that the FCC was created during the Great Depression when we didn't trust markets and that telecom is, in effect, a socialist experiment premised on the idea that regulators can regulate markets. It doesn't have the idea that the very nature of the market change in ways that are belie the presumptions underpinning the regulations. Nor is there the understanding that the conflict of interest inherent in the idea of telecommunications does not arise if we have a different marketplace configuration.

In fact, if networking as a service is only an accident of history and is not fundamental how can we justify continuing an experiment that has failed?

Networks vs. Networking

In the 1980's there was a "network" known as UUCP. I put the word in quotes because it wasn't a network in the sense of having a network provider. Instead it was a cooperative effort among computer systems. We could also say it was a cooperative (social) network among the system managers.

UUCP stood for Unix-Unix-CoPy for coping files. If you wanted to send a message to a "Joe@Mickey" your computer might dial a computer named "Pluto" and it would in turn dial "Mickey" where "Joe" had a mailbox. Remember that in those days long distance phone calls were expensive so you would want to avoid direct calls. It's a good example of being forced to take a more complicated approach to game the phone system.

What do we mean by "expensive"? We accept the idea that a product is expensive because the costs are high. But diamonds are expensive because De Beers is able to manipulate the marketplace. Thanks to a government-granted mo-

nopoly ATT had every incentive to use the most expensive gear even as digital technology greatly reduced the costs.

The idea of managed competition prevented the normal market from operating to reduce the prices and bring it more in line with costs. The problem is that with too many players there is too much supply and the price drops below cost.

Yet even today the carriers still charge us for phone calls even when we've already fully paid for a multimegabit pipe and the phone call doesn't really exist inside the network. With FiOS there is no copper phone line yet the price of a phone call has stayed the same.

We accept this because we believe there must be a network and there must be services.

The significance of UUCP is that it demonstrates how simple it is to do networking without a service provider. But the word "network" is so general that a social network is a network just like a phone network. This is why I try to be careful to focus on the activity of "networking" so as to emphasize that there is no provider charging for services that we create ourselves such as phone calls.

And once again we are lead astray by language. Even when we have computers that just place packets on a wire and listen for them we call it a network. And then we assume that we need a network operator.

But all we need are wires and radios. We can then use simple protocols to exchange bits locally and at a distance.

This is similar to the approach we take with local roads in each town forming a road system as you drive across towns. The interstates make the system work better but are not fundamental.

It costs money to lay down fibers but far less than roads. In fact we often put in fiber as we build roads because it's such a small part of the cost. Radios are very inexpensive but cell towers are expensive because of the complexities of synthesizing billable events and simply because carriers are rewarded for high expenses with high market caps (AKA high stock prices).

And we need people to keep the equipment working, again like roads. Again the cost is low. It's even lower once the facilities are not owned by high price network operators. We'll quickly see the kind of marketplace competition we've seen for home and office networking leading as we design equipment to make it easy to operate networks. If nothing else we'll avoid the frightful complexities of gaming today's systems of inter-carrier pricing.

The service-funding model requires scarcity in order to create value. This is why the cellular phone providers are complaining that too many people are using their services – it justifies their need to manage the network. If you have two phones next to each other they don't communicate unless there is a distant tower to create a billable event. It's a design decision chosen by business policies to justify the very same business policies. This is a recurring theme.

We have open access points everywhere in a city but they are locked down for the same reason. The carriers complain about the cost of backhauling cellular traffic to their networks. Yet they typically have huge capacity in the wires and fibers in many neighborhoods. They choose not to add inexpensive radios to the poles because it would make it obvious that wired and wireless bits are the same and that neither are scarce.

This scarcity story also applies to broadband. There are redundant wires on each poll but only about one percent is available as "Internet". And any one of the broadband wires (or fibers) could serve the entire community.

Why do we have a single electric distribution system yet separate wires for bits that are all the same and aren't even consumed? As the advertisements say, there is just one Internet. So why do we need a wire as if there were separate Comcast, Verizon, ATT, RCN, etc. networks? As if there were different a network for each purpose: safety, fire, policy, education, parks, homes, offices, dry cleaners, dogs ….

If we didn't try so hard to keep bits within billable paths we call networks, we would have abundant capacity with very simple protocols and no need be limited by the need to create billable events.

A New (kind of) Infrastructure

Infrastructure exists to serve the community. It may be the office equipment in a company or the streets in a city or a bridge across a river.

I chose the word "infrastructure" rather than "utility" to avoid the implication that we are paying for services rather than using what is available. We still pay people to keep the roads running and bits flowing but we're not paying for "consuming" the Internet.

The other important aspect of connectivity infrastructure is that it can be a common infrastructure for all kinds of services. If we don't need providers we don't need separate infrastructure for each purpose such as police radios, traffic light controls, educational networks or even turning on the street lights.

Once we understand digital systems we can do much more. A home fire alarm can now simply contact the fire station by sending a message, a rich message. A rich message can provide information such as what chemicals are sensed at what temperature in what room. Of course privacy is a concern so home owners should be able to decide what information they will want to make available.

The ability to connect devices becomes even more important when we have powerful devices that we can use for any purpose. The computer power now locked within a smart phone becomes far more valuable when we can treat it as a general purpose device rather than just a telephone.

You want a doorbell with video? Just take a $200 smart phone and stick it to the wall. That's it. You can use the ambient connectivity to connect it to the apartments. But that's just the starting point. You might also support one time pass codes and even video check-in for packages or plumbers.

Need a medical monitor with video? Same approach: Just connect additional sensors if needed.

You can also have other form factors such as a wrist mounted device with a small screen for convenient viewing and, perhaps, the ability to monitor blood flow with a sensor.

Carriers

We have phone companies, cable companies, Internet Service Providers and various other forms of service providers or carriers. The details and business models vary greatly. Phone companies make money by creating services within the network. The cable companies started out as community antenna companies bringing in distant television signals.

Much of the focus in this essay has been on phone companies because of their historic role. But ultimately they all share the common property of owning the means of communications which they fund by selling services.

Cable companies today are less dependent upon the network itself and are further reducing their dependency. Comcast is buying NBC Universal and Time-Warner spun out their cable division.

A number of the traditional phone companies have looked to cellular telephony as a new opportunity for revenue. The seeming complexity has afforded them some degree of protection from commoditization but it's only a temporary respite.

They have also tried to become "cable" companies themselves. In fact DSL was originally created to carry TV programming but it took Verizon well over a decade to become a "cable" company with a multibillion dollar investment in new fiber (FiOS).

"TV-Anywhere" is a technology that allows cable companies to offer their content over IP. They are no longer tied to their own wires. This is a new dynamic because they would be subject to reciprocal discrimination were they to impose byte counts or other limitations. (Of course they could collude by making private arrangements).

What happens when they start marketing their services to users who aren't on their fiber? They will be able to charge less (or have a higher profit) once they are freed of the burden of maintaining a private infrastructure.

The carriers understand that their business model must change and that's a source of optimism. They are reluctant to give up the advantages that come from owning the network but increasingly they have a stake in a post-network world.

Perhaps Verizon realizes this because it has stopped deploying new fiber for FiOS.

The problem is that the very structure of the business makes it difficult to change. We saw this when ATT tried and failed to become a computer company. We should encourage their effort to adapt to change but we can't depend on it and must therefore force the issue.

It's a lot to ask a company or, more to the point, people in a company, to accept the idea that their business model doesn't make sense. The belief that a smarter network is necessary feeds a need and is consistent with the notion that network operators are necessary. We see this in government policy as well. It's not just the misguided broadband initiative; it's also research efforts like GENI which seek solutions in the network. We should be focusing on research on how to use networking instead.

Collateral Damage vs. Discovery

I have a spare cell phone – an Android G1. I wanted to use it as a development platform but found that I couldn't use it without having an unlocked SIM card. This is an example of collateral damage. Not all phones have this problem but the fact that some do is troubling.

If the purpose is simply to prevent unauthorized calls then it would be simple to prevent calls as if there were no service available. Instead I can't use the phone at all. This isn't problem for the cell companies because what other purpose does the phone serve in their world?

These architectural couplings are expedient and good engineering within the given design point but they are very bad public policy.

As I wrote in "Beyond Limits", one way of enabling generative hypergrowth (the kind we associated with Moore's Law") is to decouple markets so that neither is dependent upon the other.

We see this in the "end-to-end" principle of the Internet that stats that services can be created outside the network at the end points without depending on services within the network. But if you own the network then you will build a reliable transport because there is no reason not to and common sense argues that it is necessary. It also gives you a way to add value.

If you don't control the network you don't have the temptation to build smarts into the network. This is why I refer to an end-to-end constraint. Because of this constraint we discovered that we could do voice over an unreliable path.

But we didn't really discover that until the network had enough capacity to make it work. Until then we had to find other uses of connectivity. I go into more detail in my essay on "Purpose vs. Discovery".

In terms of public policy the telecom silos are an example of the price we pay for expedient solutions within silos. We can create generative opportunity by unlocking the value latent in the silos.

This is not just a problem with carriers. We see the same problem with cloud computing as service providers create protocols tethered to the cloud (AKA, their servicers and services). Not only does it avoid the hard work associated with true peer protocols but it also helps make customers dependent upon the services.

Once we start looking for examples they are everywhere. Look at the Set Top Box vs. the TV vs. other sources and the lack of protocols to make the process of coordinating the components anything less than very awkward.

This news keeps coming.

As I was editing this article my son called to tell me that Comcast's Internet is down! We joke about the Internet filling up or breaking but that can happen if you use a brittle architecture. This particular problem appears to be DNS-related – an unnecessary dependency upon a central authority. What is important is that this problem lasted for hours. People often say that we need providers because cities can't even fix potholes. But this is worse than not fixing a pothole. We should expect some accountability and explanation.

Remember that many cities do operate complex infrastructure like electric distribution systems.

Breaking news:

Worse, Comcast is now holding their customers hostage according to Level 3!

Connectivity is vital infrastructure and must not be in the hands of companies that treat it as just another TV channel. Cities do know how to deal with vital infrastructure. They may not have the skills on staff but they do know to hire suppliers who can be responsible. There might still be failures but the outcry will be strong enough to demand immediate change.

A mayor must answer to users each election but a cable company's responsibility is to its shareholders who have very different incentives.

Creating (Generative) Opportunity

I argue that the business model of telecom has failed. It might have made sense in the 19th century but why are we still paying for services we can do better ourselves? How is it that we pay for a phone call even when we have a 10 or 100 Mbps (Megabit per second) connection?

China depends on its Internet connections to the world. Considering how small the cost is relative to the value, it would make sense for China to invest in a mesh of cables across the Pacific to assure capacity and resilience. Instead it is dependent upon private providers whose business needs depend on maximizing the value of their investment by limiting capacity.

We also see this same lack of capacity in South Africa where efforts to extend education to townships run into limitations from an artificial scarcity of "Internet" because even local connectivity is subject to limitations lest the students use too many bits over the expensive (AKA thin) connection to the rest of the world.

With an infrastructure approach we can have the ability to communicate with zero marginal cost. Imagine the innovation unlocked when the new computing devices can assume connectivity.

Today the carriers must capture the value created using the network to fund their infrastructure. Applications that don't meet their price aren't considered worth supporting. Remember that price is not the same as cost.

Email: A Precedent

We have an important precedent in electronic mail. In the 1980's the international standards bodies agreed on a standard for email called X.400. It embodied all of the protocols of the telecommunications industry mashed up with the post office so you could provide seemingly necessary services like reliable delivery and confirmation.

There were two problems. The big one is that it imposed all the assumptions and limitations of the hierarchical world of telecommunications and government post systems (AKA, the PTT or Postal, Telegraph and Telephone systems) on electronic communications.

It failed because too many people understood email and protocols like SMTP (Simple Mail Transport Protocol) were too easy to implement. It only took a weekend to implement SMTP and another weekend to improve it. Changes to X.400 took a decade.

Sure, email didn't have the authoritative guarantees of the X.400 hierarchy but those assurances didn't really guarantee anything. You still had to know how to trust. Just because someone tells you to do something you don't just do it. You still have to verify and use your own judgment.

If you want an authority to vouch for the contents you can implement your own system for exchanging messages and there are many choices.

Finally in the 1990's the US Government relented and X.400 was no longer mandatory.

Perhaps the biggest difference was that email became common before it was captured.

Getting There

It would be nice if we could just make all of the wires and radios available as basic infrastructure like we do the roads and paths so that we can do our own networking without being beholden to a provider. Yet we will have to live in a world where telecom providers still have exclusive control of much of the wired and wireless infrastructure. We still have an FCC still exercising prior restraint on speech as if electronic speech wasn't protected by the First Amendment.

In 1934 protecting the providers was the best way to assure the availability of a communications infrastructure. It was just a continuation of the approach used for the railroads. Even if it were the right approach then it is certainly not so now.

Railroads have been subject to common carriers rules which prevent them from discriminating between various types of cargo. Discrimination by network providers is more subtle and insidious in that they can provide self-serving arguments for how to operate the network favoring applications they argue are more important.

This is why the idea of an open Internet is so important as a way to limit the carriers ability to meddle, even if with the best intentions, that make it difficult to innovate outside the network by doing our own networking.

Today a community that wants to pay for its own infrastructure finds itself thwarted rather than supported. We see an example of the problem with Utopia in Utah. Utopia was modeled on the high system as basic infrastructure. But instead it must be funded in the provider model as another backend provider.

Yet even with these impediments the idea of funding the wires and radios as common infrastructure is compelling. Just as we pay once for the wires and radios in our home, offices and campuses we can fund the facilities in our towns and cities. We can then have direct connections between these communities without having a provider in the middle of all conversations.

The question is how hard we are going to try to preserve the past and how willing we are change and embrace a rapid transition to infrastructure.

The Right to Communicate

The United States has been fortunate in having the First Amendment which gives us the right to unfettered speech. It means ideas get a chance to be heard without having to first prove they are good ideas. In fact the ideas that have given us the most have been disruptive or bad in that that they upset the status quo.

The Internet has given our free speech far greater reach. We can explore ideas because we are curious without having to justify them to investors. Individuals can effect real change.

Yet the very idea of the government licensing radios violates the First Amendment because we understand how to build a communications infrastructure without creating the artificial scarcity that justifies such intervention. Using spectrum to manage wireless communication assures scarcity.

If you use a single frequency then you have to make sure no one uses the same frequency thus you have to exclude others. It's as if you had to make sure that no two people wore shorts of the same or even similar colors in a room because you used only the shirt color to distinguish between people. The problem is not that we have a shortage

of spectrum (colors), it is that we are using in an extremely foolish way. The fifty year old idea of auctioning off portions of the spectrum might provide a minor improvement in usage but doesn't address the fundamental flaws in the basic approach.

The tragedy is the entire edifice of the FCC is premised on this scarcity and the mission has crept. The presumption of scarcity is used to justify managed competition with providers in the middle of the path setting a price hurdle on speech. And without common carrier protection we are back to the days of the robber barons charging us based on what we say not just the cost of carrying the bits.

We have a right to communicate. If we fund infrastructure instead of charging for services we can realize that right.

There: What we can achieve

Where is there? Let's go back to the doorbell. You can just put it in your pocket and assume it will ring when someone presses the button at your house. Of course you don't really carry a doorbell with you – you would instead have an app on your personal device (or smartphone) to alert you.

What is important is that it "just works". You simply decide the policy of whether you want to hear the doorbell when you're not at home based on your needs and not any technical hurdles.

We haven't just let you take your doorbell with you, we've created opportunity. You can now start innovating and add capabilities such as a video feed from your door. Of course the "doorbell" might be a medical monitor or an alarm or any connected device.

It's a cliché to say we're limited only by our imagination. But today we can create solutions by taking an idea and sharing it using connectivity and some software.

Once we can see past telecom to the radical simplicity of the Internet the question is not whether we can move beyond telecom but what took us so long?

Summary

Ultimately the Internet is about the idea that we can focus on our applications without being concerned about the complexities between the end points.

This is the end-to-end principle, not to be confused with womb-to-tomb control of the path.

The key insight is that we can create meaning outside the network without having to require more than *best efforts* from an intermediary. In fact making best efforts a constraint has created a dynamic in which demand created supply and innovation has allowed us to take advantage of opportunities rather than limit ourselves to what is provided.

This frees us from having a rent-taker in the path but can only be fully realized when we are no longer funding the infrastructure by paying for provided services and, in fact, free ourselves from the concept of a network as a service.

Instead we should fund physical infrastructure – pay for copper and glass for wires and stop prior restraint on the use of wireless devices by taking advantage of our modern understanding of physics and signaling.

Today's Internet protocols were meant as prototypes and work very well to connect computers but they don't support long term stable relationships. Once we can assume unfettered connectivity we'll see more focus on the relationships between applications (and the people that use them).

An application like connecting a button at a door to the doorbell is complicated because we don't have stable names for the end points and easy ways to maintain the relationships. Various applications like Skype have their own approaches. In the future we'll see more focus on relationships because merely exchanging bits will no longer be the difficult problem it is today.

Further Reading

I can only touch upon many of the topics in this essay. You can read more details in other essays that I have written.

My Writings

It's difficult to get past our implicit assumptions about the purpose of networks as I explain in "Purpose vs. Discovery" and "On NOT Baking in Special Services". Government research is still focused on finding answers in the network itself (Project GENI) rather than in understanding how to do networking. Concepts like "security" and "trust" are not to be found in the Procrustean plumbing of the network.

"Our Copper, Fiber and Radios" is my attempt to take a positive direction. Instead of criticizing the carriers I observe that what we need to do is take advantage of the physical facilities all around us. We also need to think in terms of the neighborhood around us – The First Square Mile.

Being able to presume Ambient Connectivity is a start. In "Maker Disconnect" I start to look at other enabling tech-

nologies though I note we won't be able to fully explore the ideas until we can presume connectivity.

In 1997 I gave a talk about how to embrace these principles in consumer electronics.'

Finally there's my chapter on hypergrowth (or what I would now call generative economics): Beyond Limits.

Others

I may add to this list over time until I create a stand-alone bibliography.

"Turnpikes and Toll Roads in Nineteenth-Century America" by Daniel B Klein and John Majewski. The transition from private pikes to today's public roads is fascinating and offers lessons us today in how to create public infrastructure.