

What Went Wrong?

How the Internet Stagnated

or

The 30 Years War

of the

Bellheads and the Bitheads

Science and Technology Innovators Lecture

University of Minnesota

John Day

Oct 2010

What Went Wrong?

- Network research and the Internet has stagnated.
 - This may seem an outrageous claim given its prominence and the daily headlines. But think carefully:
 - Is it quantity or quality?
 - Is it just much more of the same much cheaper or truly deep intellectual insights that are responsible?
 - Consider:


“A reviewer of an early draft of this report observed that this proposed framework – measure, develop theory, prototype new ideas – looks a lot like Research 101. . . . From the perspective of the outsiders, the insiders had not shown that they had managed to exercise the usual elements of a successful research program, so a back-to-basics message was fitting.”

Looking over the Fence at Networking,
Committee on Research Horizons in Networking,
National Research Council, 2001.

Everyone Knows the Story

- In the early 1960s, Paul Baran at RAND comes up with the idea of packet switching.
 - And independently by Donald Davies in the UK
- In the late 1960s, Larry Roberts heads a project at ARPA to build a packet switching network.
 - The first node is installed in late 1969 and quickly grows.
 - By 1973, the network has become an integral tool of daily work for a handful of researchers.
 - Developed by Operating System experts, not data comm.
- But networks were not new

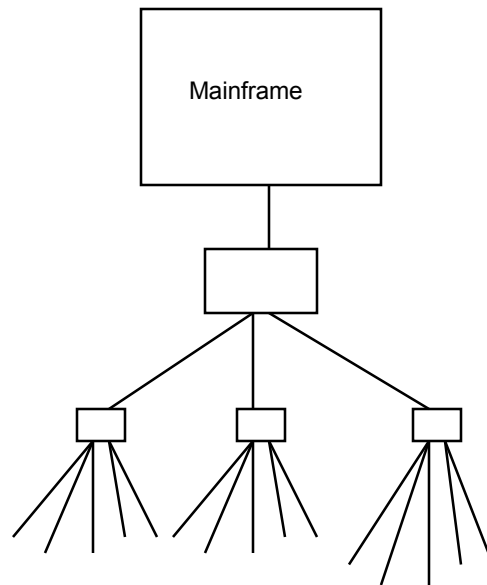
The Phone Company



- Beginning in the last quarter of the 19thC, telephony had always been based on a deterministic circuit model, which I call,
- “Beads on a String” which describes not only interfaces but who owns what (this is important).
- Remember at this time phone companies are monopolies
 - Outside the US, the phone company and the FCC were the same organization. They had immense power.

Mainframe Model: SNA

- Asymmetric/Deterministic/Hierarchical
- Everything is the “Mainframe” (and always will be) minimal functions at the Terminal. No question: everything but the wires belongs to IBM.
- The Mainframe and TPC co-exist nicely.
 - IBM has 85% of the computer market, very powerful.



But was Packet Switching a Major Breakthrough?


- Strange as it may seem, it depends.
 - During this period many things are age dependent.
- If your formative years had occurred prior to the mid-60s, your model of communication was defined by telephony.
 - Then this is revolutionary.
- If you are younger, your model is determined by computers.
 - Data is in buffers, How do you do communications: Pick up a buffer and send it.
 - What could be more obvious!

But then a Real Breakthrough: Datagrams

- Louis Pouzin had a bright idea: The network didn't have to be perfect.
 - This is completely unexpected.
 - Simple elegant. Dynamic, Non-deterministic.
 - Reliable communication with unreliable parts.
- CYCLADES is the first Connectionless network.
- While the idea of layering software had been part of the early implementation, this brings layers to networking as a key concept.
 - Baran had proposed a packet switching concept more like what became known as X.25 or ATM. There were no layers in Baran's work
 - and everything Baran did after this was connection oriented.
 - BBN followed Baran's lead in the ARPANet design.

Layers are Pretty well Fixed

Host or End System



- At least the lower layers. We don't know much about the upper layers. We only have 3 applications.

Explosive Growth!

- Well, *We* thought so! (within the community)
 - Everyone is building packet switch networks.
 - USING, the Users Interest Group, is formed with great ideas for making the ARPANet a true resource sharing network
 - Packet radio, packet voice experiments, distributed database access from workstation with plasma panel and touch screen, network software works, “distributed computation,” NLS with a mouse, IM’ing in 1972, etc.
- Research networks adopt the connectionless model
- Commercial networks, the connection model.
 - Layers? Levels of abstraction? Well, okay, but seems a little strange.

Regardless of Background

There is a Paradigm Shift Happening

(We all read Kuhn, right?)

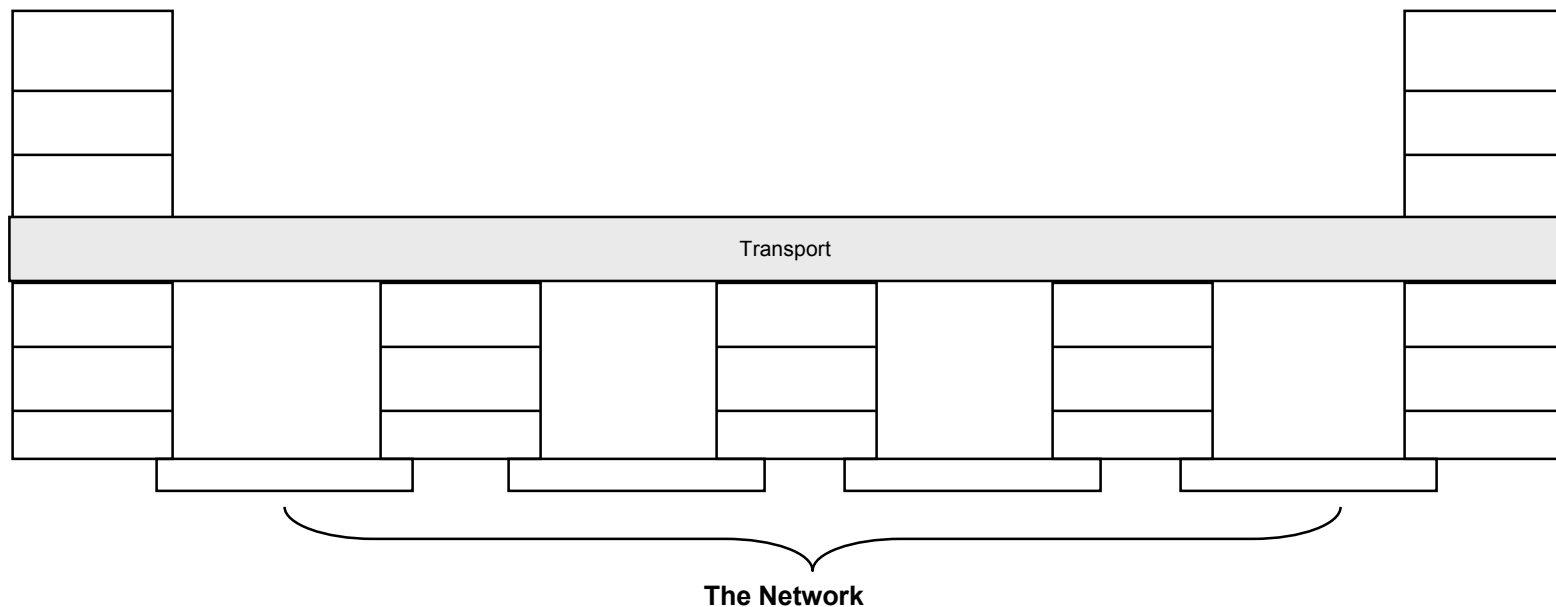
- Connection-oriented beads-on-a-string is being overthrown for connectionless layered model.
 - The new model easily explains many things
 - Coming up with a paradigm shift requires a different talent than required for “normal science.”
 - But is this “continental drift” or “plate tectonics?”
- As with any paradigm shift, we can expect a period of intellectual breakthroughs followed by a period of consolidation.
 - Each Requiring very Different Talents.
- Then as the new paradigm runs up against its limitations a new round of deep-thinking will be required to find the new paradigm
- As always not everyone gets with the program (never do).
 - Here, strong economic reasons not to and they are digging in for a fight!

For Networks, Standards are a Must

- As INWG is ramping up, it is discovered that the PTTs of the world are developing something called X.25 in CCITT.
 - It looks more like the early ARPANet, not our neat new idea.
 - We know its limitations, and they don't want a transport layer!
- An effort is made to get datagrams in X.25 and it succeeds in name. But the PTTs don't take these researchers too seriously.
- Furthermore, the CCITT is a treaty organization and only national representatives or phone companies can speak.
- The phone companies have this locked up.
 - Our introduction to electro-political engineering is a learning experience

Transport Relegates TPC to a Commodity Business


- Transport Seals Off the Lower Layers from Applications.
 - No such thing as Applications “*in the Network.*”
with very little possibility for value-add, and is neutral about who owns what.
- TPC counters that Transport Layers are unnecessary, *their* networks are reliable.



And they have their head in the sand, “Data will never exceed voice traffic”

IBM already had problems


- Computing and Memory Prices were headed South . . . Fast.
- Computing Power and Capacity were headed North . . . Fast.
- By the late 70's, it was clear that IBM's days as the dominant computer maker were numbered



And if that weren't enough.

IBM Found Itself in a Dead-End

**Cardinal Rule of Architecture:
You can always make a peer architecture hierarchical
But you can't go the other way.**



Had IBM made SNA a peer network and subset it for the 70's hierarchical market, the Internet would have been nothing but an interesting research project.

And if that were not enough

- The phone companies and IBM had carefully stayed out of each other's domain.
- Now the layered model created a complete overlap.
 - TWO 800 pound gorillas have been thrown in the room
 - and they are really pi . . . err very mad!
- And minicomputers were going to be important to exploiting the market.
 - Not IBM's favorite product line.

Warring Models? Indeed

- Intense negotiations during 1978 and early 79, finally lead to the CCITT agreeing to a joint project.
 - But the models are complete opposites!
 - O, we will figure something out.
- However, it isn't just the phone companies vs computer companies.
 - It is connections vs connectionless and Europe vs the US vs Japan
 - and Everyone against IBM.
- IBM counsils caution. A huge undertaking, we must move cautiously.
 - Some call this stonewalling. And since they chair most national committees they are in a good position to ensure “caution.”
- After INWG's recommendation of CYCLADES TS, DARPA bows out and tries to take the DoD with it.

We had Issues: I

- NCP wasn't going to scale, it needed to be replaced. But the datagram model indicated a direction.
 - CYCLADES pointed the way.
 - But the Internet had its own proposal.
- Had we gotten the basic structure correct?
 - This is all new. We haven't seen much of the problem space.
 - We knew we had a few kludges.
- What was the “upper layer” architecture?
 - 3 simple applications, but there is probably more here.
- Connectionless looks very promising but only used in the small, how would it scale?

We had Issues: II

- We needed application names and a directory.
 - What did it look like?
- We needed a solution to multihoming.
 - But being OS guys, the solution was clear. Addressing needed a hard look
- If addresses were location-dependent, what did that mean in a network?
 - This one is hard.

- Quite a List! Lots of material for researchers.

So What Happened?

- USING is shut down in early 74 after about 6 months.
 - Lots of good ideas, much excitement.
 - ARPA pulls the plug, fearing they are losing control.
 - For 20 years, no new applications in the Internet.
 - This will prove to be a critical juncture.
- Focus is now on implementing TCP.
 - Internet continues to grow. IP becomes a religion.
 - Researchers are intent on showing they can be good operators.
 - 2nd Generation effect, ensures tendency to stagnation.
 - Concentration on implementation, small incremental change, reverence for things past, NIH in spades, etc. This is artisan, not scientific behavior.
 - Moore's Law begins to be a factor.
- Nothing new is going to happen here.

What about TCP and DNS?

- There were many controversial aspects of TCP and at least 3 competitors: XNS seq pkt, CYCLADES TS, and delta-t.
 - While they had advantages, for the environment of the 70s, they weren't overwhelming.
 - And DoD was paying (and paying). TCP takes *8 years* of development and many band-aids before it is deployed.
 - Often touted during the 70s and 80s as the most tested protocol around. It was, in fact, only the most *used* protocol around. There was little real testing.
 - Almost every design decision in TCP is wrong for a large diverse Internet.
- DNS replaces the host file.
 - A directory? No, a synonym resolver. The minimum needed to replace the host file. Domain names are just macros for IP addresses.
 - Fails to introduce a level of indirection, crucial to a complete architecture.

Meanwhile Datagrams Get a Boost

But not without considerable angst

- LANs have been in the lab for nearly a decade. Ethernet is simple and inherently a datagram technology.
 - Digital, Intel, and Xerox (DIX) form IEEE 802, expecting to have a standard in 3 months.
 - They get a lesson in electro-political engineering.
 - 3 Years later they have much more.
- Ethernet, Token Bus and Token Ring and 3 forms of LLC.
 - To placate the stochastically averse and IBM
- But even so, it is popular and there is a lot of activity and product.
 - Some see it as a whole new paradigm for networking!
 - Only to be disappointed that it is just another data link protocol
 - This makes datagrams something that can't be ignored.

Can OSI Accomplish Anything?

- Not really. Everything it does is controversial.
 - Every word, every comma is argued over in every document.
- The PTTs don't want connectionless.
 - For the US, it is the only reason to be there.
- A major debate over whether a Transport Layer is needed.
 - PTTs, of course, say no. Computer companies say yes.
 - Europeans mistakenly hedge saying X.25 isn't completely unreliable
 - Result: 5 classes of transport.
 - Standard tactic: If you can't stop something, loading it up with options is almost as good.
- As part of the agreement for a joint effort, PTTs steal the Session Layer for two dead end applications.
 - This fatally impairs the ULA, which is otherwise quite good.

The Internet in the 80s

- Not at all. More concern with a stable operating platform than research.
 - Fast changing hardware is causing a lot of implementation work.
 - As far as cutting edge research goes, not much is happening. Implementing TCP, growth, but Moore's Law accommodates that, no new applications to generate new requirements.
 - Hey! We're having fun, what do you mean vision for the future?
- The 7 unanswered issues are largely forgotten. Newcomers see the Net as a complete and finished whole.
 - Not the half an architecture it is.
 - The Internet is the bible, kludges and all are now stone tablets.
 - All network research becomes Internet engineering.
 - The beads-on-a-string types retreat to ISDN and then ATM

Major Crisis in the Internet!

- In 1983, the Internet switches from the ARPANet NCP with strong internal congestion control to a pure connectionless network layer using IP.
- In 1986, the Internet suffers congestion collapse daily.
 - Why did it take 3 years? In early 80s, most hosts were connected to switches with ingress control (1822 or X.25). Through the 80s, hosts connected to LANs with no ingress control began to dominate.
 - As is often the case with converts, their zeal is greater than their understanding.
 - Others did. It was a hot topic in the late 70s
 - The Internet experts could not see how to put congestion control in the network layer without sacrificing connectionless, so it is put in TCP as a stop-gap. Even though, control theory says this is the worse place for it. Turning the Faber Marching Band down that dead end alley.

OSI in the 80s

- The Internet is restricted to only DoD contractors, hence to the vast majority of business, networking is going to be OSI.
- In 1986, there is a big demo. This sparks yet another battle for who sets the commercial direction.
 - NIST, Implementor's Workshops, GOSIP, MAP/TOP, COS, DEC, IBM, PTTs, etc. Everyone has their own agenda.
- Development continues.
 - By 83 realized that the upper 3 layers are a single layer. Application structure is understood, a full addressing architecture, link-state routing. But for the most part, the documents are so dense, no one outside OSI is aware of it.

Separation bred Suspicion and Contempt

- Beginning in 1979, the separation begins.
- The Internet is not even on the screen of the business world. The connectionless advocates in the business world see it as an ally and are keeping a foot in both camps.
- The Internet sees itself as the threatened underdog and all of OSI as the enemy. We are right, OSI is wrong.
 - OSI is all politics, a bunch of phone company types.
 - It wasn't
 - If OSI did it, good bad or indifferent, it must be a bad idea.
 - Not let us show you how to get it right.
 - Groupthink has set in.

Groupthink

According to Smolin:

1. Tremendous self-confidence,
2. An unusually monolithic community,
3. A sense of identification with the group,
4. A strong sense of the boundary between the group and other experts.
5. A disregard for and disinterest in the ideas, opinions, and work of experts
6. A tendency to interpret evidence optimistically,
to believe results are true because they are “widely believed,”
7. A lack of appreciation for the extent to which a research program ought to involve risk.

In the The Trouble with Physics.

Quoted in Patterns in Network Architecture

Then the Nineties

- OSI is specified and implemented but no network.
 - And internal dissension is even worse. OSI self-destructs.
 - Lesson of OSI: Don't invite the Legacy to the new game.
- Suddenly a lot is happening all at once.
 - Another Crisis: Router table size and Address Exhaustion
 - Entrenchment reaches a peak. Anything but CLNP; IPv6 is answer
 - Major effort by politicians, leads Internet to NSFnet and the public.
 - Maybe Al Gore did invent the Internet, opening the door for,
 - The web breaking on the Internet, the Internet boom is on.
 - Confirming that we were right all along.
 - Ignoring that this is the same argument that DOS is a great OS.
 - Students are taught the Internet, not the principles. Universities do vo-tech
 - It is an unfinished demo living on Moore's Law and 30 years of band-aids.
 - Given the nature of software as long as no one did anything really dumb (and they were far from that), the Internet would work.
 - Errr, . . . The 7 questions? What 7 questions? still unanswered.

Craftsman and Seers

- Smolin also points out the increased professionalization of the sciences, selecting “master craftsmen,” but not theoretical “seers.”
 - Master craftsmen are those who were really good at math and science and solve problems quickly. We call them super-coders.
 - Incremental change and running code selects for craftsman.
 - Seers are also good at math and science but are more interested in getting to the bottom of things, what is really going on.
- Both are necessary:
 - Master craftsmen are best when exploiting a paradigm, but when you get to the edge and need a new paradigm, you need a different kind of head. This is when you need the seers.
- The emphasis in the Internet has been on master craftsmen to the total exclusion of seers.

By 2000, There is a Bad Feeling all is not well.

- Signs the Internet architecture is running out of steam.
 - Internet research is stagnate. Back to our NRC quote:

“A reviewer of an early draft of this report observed that this proposed framework – measure, develop theory, prototype new ideas – looks a lot like Research 101. . . . From the perspective of the outsiders, the insiders had not shown that they had managed to exercise the usual elements of a successful research program, so a back-to-basics message was fitting.”
 - Must be sobering to be told you don’t know how to do research.
 - Something new is needed, but everyone has been taught for the last 25 years everything is perfect. Unanswered questions? What unanswered questions!
 - This is a success!
 - There is nothing wrong with TCP and below. The problem must be elsewhere.
 - Reaction of master craftsman to finding a new paradigm:
 - Start building things! GENI?
 - Not start thinking hard.

So What Went Wrong?

- ARPA killing USING removed the only driver for change.
 - With Moore's Law, growth alone was not enough to cause it.
- Research was pushed on the world stage too soon: bunker mentality
- By mid-70s, second generation effect was in place.
 - Either old guys protecting legacy, awe of the old guys, IT mentality
- Groupthink. So much emphasis on practice there was no theory to see where it was going and no one really trained to do it.
- Textbooks and courses are vo-tech: What is out there, not why.
- Reticence to break with the past.
 - Okay for old guard to criticize but no one else.
- The success of VC (short ROI) lead to emulation by research funding.
 - Again an emphasis on technique (solutions) not theory (answers).

What Now?

- ARPA funds NewArch for 2 years. Comes up dry.
 - Nothing, Nada
- NSF creates FIND and GENI. (Mostly same people)
 - Promises of breaking new ground! Grand new ideas!
 - For awhile, “clean slate design” is the buzzword
 - But with whiteboards, the technology seems to have been lost.
 - The results are no better.
 - Emphasis on group effort ensures no one is going to break new ground.
 - How to fit favorite projects into some project plan. Bassackward!
 - Primary purpose is to maintain revenue stream for their institutions.
 - Now, rationalization has set in (also known as seeking political cover).
 - “The Internet has always progressed by evolution, by incremental change. There isn’t *that* much wrong. A few fixes here and there. All will be well.”
 - Clearly a craftsman approach is not good for finding new paradigms.

We Interrupt this Slide for Breaking News

- In the Fall of 2006, we notice that multihoming is causing router table size to increase dramatically.
 - A problem we have been ignoring since 1972!
- Projections are that Moore's Law won't bail us out.
 - (and the projections are probably low)
 - The problem is worse with IPv6
 - Already seeing other scaling problems with v6
 - And there are route calculation issues that have not been looked at
- And time is short. What to do?
 - What were these guys doing? We could have fixed this problem 25 years ago.
 - We tried to fix it 15 years ago and failed because of poor education.
 - Cisco is proposing a solution, but they have missed the point and it doesn't look like it will scale.

We return you now to your regularly scheduled presentation (in progress)

What Now?

- ARPA funds NewArch for 2 years. Comes up dry.
 - Nothing, Nada
- NSF creates FIND and GENI. (Mostly same people)
 - Promises of breaking new ground! Grand new ideas!
 - For awhile, “clean slate design” is the buzzword
 - But with whiteboards, the technology seems to have been lost.
 - The results are no better.
 - Emphasis on group effort ensures no one is going to break new ground.
 - How to fit favorite projects into some project plan. Bassackward!
 - Primary purpose is to maintain revenue stream for their institutions.
 - Now, rationalization has set in (also known as seeking political cover).
 - “The Internet has always progressed by evolution, by incremental change. There isn’t *that* much wrong. A few fixes here and there. All will be well.”
 - Clearly a craftsman approach is not good for finding new paradigms.

Dead Ends Everywhere

- IETF decisions are based more on “tradition” than “science.”
 - A Craft Guild, not an Engineering Discipline
 - New patches constantly accrete, never replace and simplify.
 - Complexity and interactions among them have become dangerous.
 - The IETF has become the new bellheads.
- But the Internet is too big to change!
 - They have been saying this since 1980.
 - Just indicates the lack of imagination
 - Assumes the Internet is near the end of its growth.
 - We haven’t started.
 - Said the same thing about the phone system in 1975.

The Problem is They Have Been Asking the Wrong Question

(still don't know how to do research)

- They have been asking,
 - What do We Build?
- They should have been asking,
 - What don't We *Understand*?
- Time to go back to doing Science
 - Maximizing the invariances and minimizing the discontinuities
 - This is what we have been doing for over a decade and it has led to a new game. A simpler, more powerful model.

What New Game?

- Exploring PNA
 - The presentation I didn't give today.
 - A much simpler repeating architecture
 - Securable, scalable
 - A Framework of principles for finding principles
- Patterns in Network Architecture
 - Prentice Hall
 - www.pouzinsociety.org
 - csr.bu.edu/rina
 - Talk to day@bu.edu