Whither Advanced Academic Networking?
<u>WITHER</u> Advanced Academic Networking?

[*Whither* == Interrogative: "Where shall we go now?" vs.
Wither == Verb: "To lose vitality, force, or freshness..."]

Joe St Sauver, Ph.D. (joe@oregon.uoregon.edu)

Merit Networking Summit 2014
September 23rd, 2014

I. Introduction

An obligatory disclaimer:

The opinions and perspective I'm going to share with you today represent my own point of view, and do not necessarily represent the opinion of any other person or organization.

One More Thing To Get Our Right Up Front...

- I'm **NOT** here to bash what the community has worked really hard to build.

- Everyone **HAS** worked really hard, and you've done a **great job.**

- The existing network at the national level (as well as our regional and campus networks) are, collectively, tangible evidence of what the community has been able to accomplish to-date.

- But having said that, we **MUST** talk about where we should be going **NEXT.**

- **We can't afford to just sit back and rest on our laurels and coast, and we can't afford to get distracted, either.**

Last Year and Today

- With that out of the way, let me thank Merit for the invitation to talk with you once again this year.

- Last year I talked about "Networking in These Crazy Days," describing some of the **security challenges** we faced, and urging you to **stay calm, get secure, and get involved.** (If you're interested, those slides are still available online at http://www.stsauver.com/merit-networking/)

- Many of you have taken those recommendations to heart and that's great. THANK YOU!

- Today, however, let's talk about the current state of **advanced academic networking, and its future prospects.**

Today's Questions

- Where's advanced academic networking "at" today? Where will it "go" in the future? Where ***should*** it be going?

- Is advanced academic networking thriving? Or is it searching for meaning and direction?

- How did we get where we are?

- What should be the role of Michigan's colleges and universities?

- Before we can talk about advanced academic networking, we need to begin by talking about **university research.**

II. Higher Education and Research

The Classification of Colleges and Universities

- There are literally thousands of American colleges and universities, each with its own unique and important role, and each striving to meet the needs of its particular community.

- Today, however, we're primarily interested in two (of the many) Carnegie categories:

  -- The 108 "<u>Research Universities (Very High Research Activity)</u>" [including, in Michigan, the *University of Michigan, Michigan State,* and *Wayne State University*, and in Oregon, the *University of Oregon* and *Oregon State University*], and

  -- The 99 "<u>Research Universities (High Research Activity)</u>" [including, in Michigan, *Michigan Tech* and *Western Michigan*, and in Oregon, *Portland State University*]

Research Universities Are The Core of the U.S. Advanced Networking Community

- While it goes without saying that Internet2 highly values ALL of its members (whether from higher ed or elsewhere, and whether research focused or not), **Internet2 was originally founded by research universities**, and **research universities** are, and should remain, **Internet2's "core constituency:"**

- 108 RU/VH schools (100% of the RU/VH category) are currently members of Internet2 (see http://www.internet2.edu/communities-groups/members/higher-education/level_1/all/all)

- 78 out of 99 RU/H schools (78% of the RU/H category) are also members of Internet2 (see http://www.internet2.edu/communities-groups/members/higher-education/level_2/all/all)

Association of American Universities (AAU)

- Other organizations also focus on **research universities**.

- For example, the AAU is an association of "**62 leading public and private RESEARCH UNIVERSITIES** in the United States and Canada." [emphasis added] In Michigan, AAU universities are the *University of Michigan* and *Michigan State*; in Oregon, there's the *University of Oregon.*

- AAU membership is by invitation, and their policy states that **"current members whose research and education profile falls significantly below that of other current members or below the criteria for admission of new members will be subject to further review and possible discontinuation of membership."**

Common Solutions Group (stonesoup.org)

- Another invitational group, CSG has just **29 members**, including both the University of Michigan and Michigan State. Stone Soup also has two consortial members, Educause and Internet2.

- CSG states, "**Intensively networked information technology is uniformly central to the work of major research universities**. [...] it is critical to the overall efficiency of **research universities** that they act collaboratively to influence commercial providers of information technology and, where the market fails to provide appropriate technology, work collaboratively to develop and disseminate common solutions to important IT challenges. Collaborative work of this sort requires open, sophisticated interaction within and across two groups, namely the technical staff of key **research universities** and the senior IT administrators responsible for strategic direction and resource allocation."

What Do Some Schools Themselves Say?

- "With expenditures in excess of $1 billion, **research is central** to U-M's mission and permeates all 190 schools and colleges..."

  https://www.umich.edu/research/

- "**Research, both basic and applied, is fundamental to the mission of the University** and is essential to Oregon's economic and civic vitality."

  http://research.uoregon.edu/

- While research is very important to higher education, **research at universities has been dramatically outpaced by research conducted elsewhere...**

U.S. R&D Expenditures by Sector Over Time: Higher Ed--Just a 13.9% Share as of 2012!

FIGURE 3. U.S. R&D, by performing and funding sectors: 1953–2012

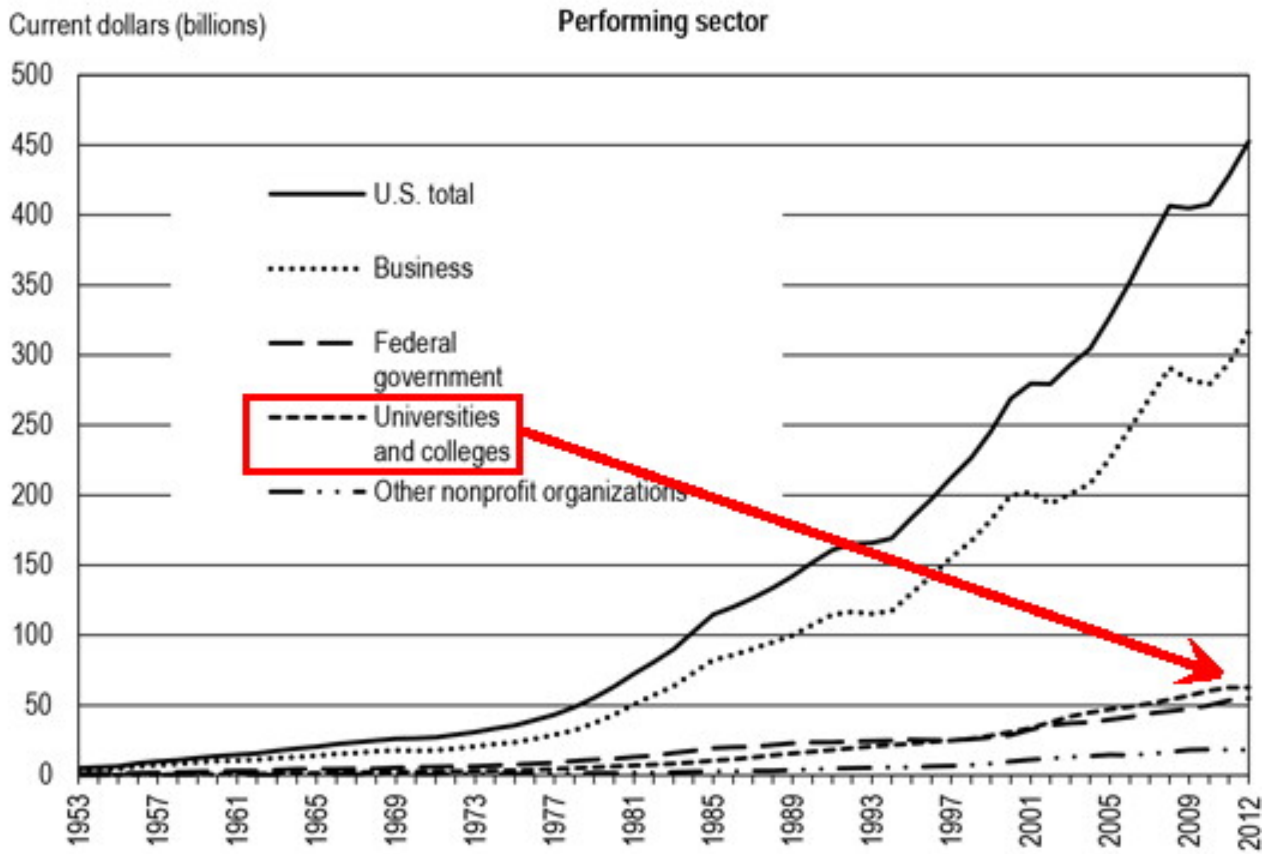

Source: http://www.nsf.gov/statistics/infbrief/nsf14307/

U.S. Federal Expenditures on R&D as a % of GDP: Investments Increased Dramatically From '53-'64

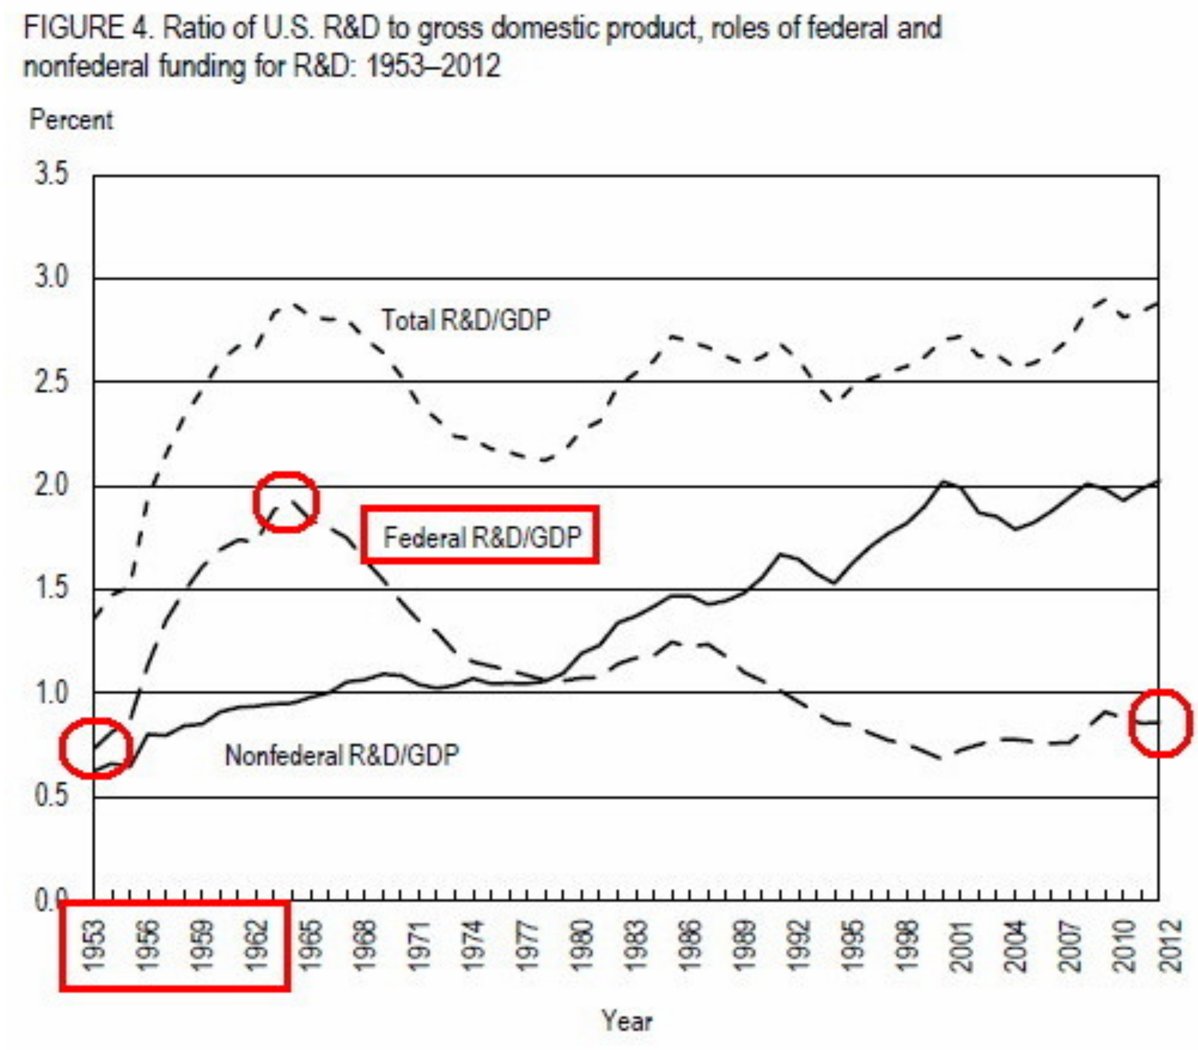

For context, the US spent 3.8% of its GDP in 2013 on military expenditures

Source: http://www.nsf.gov/statistics/infbrief/nsf14307/

Coming Back To R&D Spending...

- You've just seen how Federal investments in research and development soared from 1953-1964...

- That's now six decades ago.

- What the ***heck*** was going on ***then?*** Answer: the "Cold War."

III. The Cold War

*"**1957** – USSR launches **Sputnik** into space. In response, the USA creates the Advanced Research Projects Agency (ARPA) with the mission of becoming the leading force in science and new technologies.*

*"**1962** – **J.C.R. Licklider** of **MIT** proposes the concept of a **"Galactic Network."** For the first time ideas about a global network of computers are introduced. J.C.R. Licklider is later chosen to head **ARPA'**s research efforts.*

*"**1962** - **Paul Baran**, a member of the RAND Corporation, determines a way for the Air Force to control bombers and missiles in case of a nuclear event. His results call for a **decentralized network comprised of packet switches."***

"A History of the Internet: Internet Timeline,"
http://inthistory4u.blogspot.com/2012/06/internet-timeline.html

A Strange Time

- The 1950s were a strange time because of world events, but also because of the world's **leaders** during that period.

- **The success or failure of a nation (or any organization for that matter) is often the result of its leaders, their priorities, and their ability to execute and accomplish their missions.**

American Presidents of the 1950's

- There were just **two** men who lead the United States during the 1950s:
- **Harry S. Truman (D)**: Two term president, serving from April 1945 through **January 1953** (Truman assumed the presidency when FDR died). Made the decision to use America's new atomic weapons against Japan, ending World War II. Supported creation of the UN, but took a hard line against Soviet expansionism (aka, the "Truman Doctrine"). Rebuilt postwar Europe. **Created the DoD, Air Force, CIA and National Security Council.**
- **Dwight D. Eisenhower (R):** Two term President, serving from **January 1953 through January 1961.** Five Star Army General and former Army Chief of Staff under Truman. Supreme Allied Commander Europe from 1949-1952. First supreme commander of NATO. Created the Interstate and Defense Highway System. Originated the "Domino Theory." **President of Columbia Univ.**

What About The Soviet Leaders of the 1950's?

- **Joseph Stalin** lead the Soviet Union from the mid-1920s until his death in 1953. Stalin is thought to have been responsible for 15-20 million Soviet deaths. He also started the Soviet nuclear program. Many Americans were profoundly frightened of Stalin.
- He was followed by **Malenkov**. Few in the west know Malenkov. [It was during his tenure, in February of 1954, that **Crimea** was transferred from the Russian SFSR to the Ukrainian SSR]
- **Nikolai Bulganin** became Premier in February 1955, continuing through March of 1958. During the Suez Crisis of 1956, Bulganin **threated Britain, France and Israel with nuclear attack.**
- **Nikita Khrushchev** (Premier from March 1958-October 1964). Khrushchev was the Russian (after Stalin) who may truly have frightened the West the most, acting belligerently in the UN, fulminating the Berlin Crisis (1958-1961) and then causing the Cuban Missile Crisis in October 1962.

In 1953, Nuclear War Was Perceived As A Very Real Possibility

- In 1953, after both the United States and the Soviet Union successfully tested high yield thermonuclear weapons, *The Bulletin of the Atomic Scientists* set its "Doomsday Clock" to **two minutes to midnight,** the closest to midnight (e.g., doomsday) that its ever been.

- High yield (multi-megaton) thermonuclear weapons were assumed to be the cornerstone of a "total war strategy," targeting major population centers (so-called **"counter-value targets"**).

Nuclear Delivery Technologies of the Early 1950s

- Fortunately, early thermonuclear weapons were **large and cumbersome,** which meant that the only practical way to deliver them against intercontinental targets was via **heavy bombers**.

- Bomber attacks would be **slow** and relatively easily **detected**.

- This lead authorities to develop complex **civil defense programs** for evacuating the populations of major American target cities in the **hours** between the time an attack was **detected**, and bombs could actually be **delivered**.

- This was also an era when nuclear-tipped surface-to-air **anti-aircraft missiles** were deployed across many parts of America to help defend major cities against incoming enemy bombers.

A Civil Defense Classic: "A Day Called X"

See also: http://www.atomictheater.com/civildefensefilms.htm

The Site From "A Day Called X", Portland (1957)

- 18,820 sq ft underground complex; 6.5 miles east of downtown Portland at 2960 SE 103rd Dr, Portland OR 97266
- Was to house 250 people for 2 weeks for "continuity of city government"
- 26 inch thick reinforced concrete roof and a 230 foot reinforced radio tower
- Cost $670,000 in 1956 dollars
- **First of its kind in the United States**
- Senator Wayne Morse derided these efforts as a pointless hoax.
- Portland subsequently quit the Civil Defense program in 1963.

http://kellybutteunderground.blogspot.com/

What About *Other* Hardened City Government Continuity of Government Centers? As of 1959, "The number [was] small but increasing" and included Detroit, as well as Portland...

Many of the cities and counties have established control centers. The number of adequately protected centers is small but increasing. Those centers with protective features include Portland, Oreg., Dade County, Fla., DuPage County, Ill., Detroit, Mich., Wichita, Kans., and Oakland, South Pasadena, and Glendale, Calif. Cities and counties are expected to use all the Federal matching funds allocated to this program in FY 1960 in excess of those used by the States.

Office of Civil and Defense Mobilization, Annual Report 1959, https://training.fema.gov/EMIWeb/edu/docs/HistoricalInterest/Office%20of%20Civil%20and%20Defense%20Mobilization%20-%201959%20-%20Annual%20Rep.pdf

Michigan Also Had Nike Air Defense Sites Which Were to Launch Surface-to-Air 20KT Warheads...

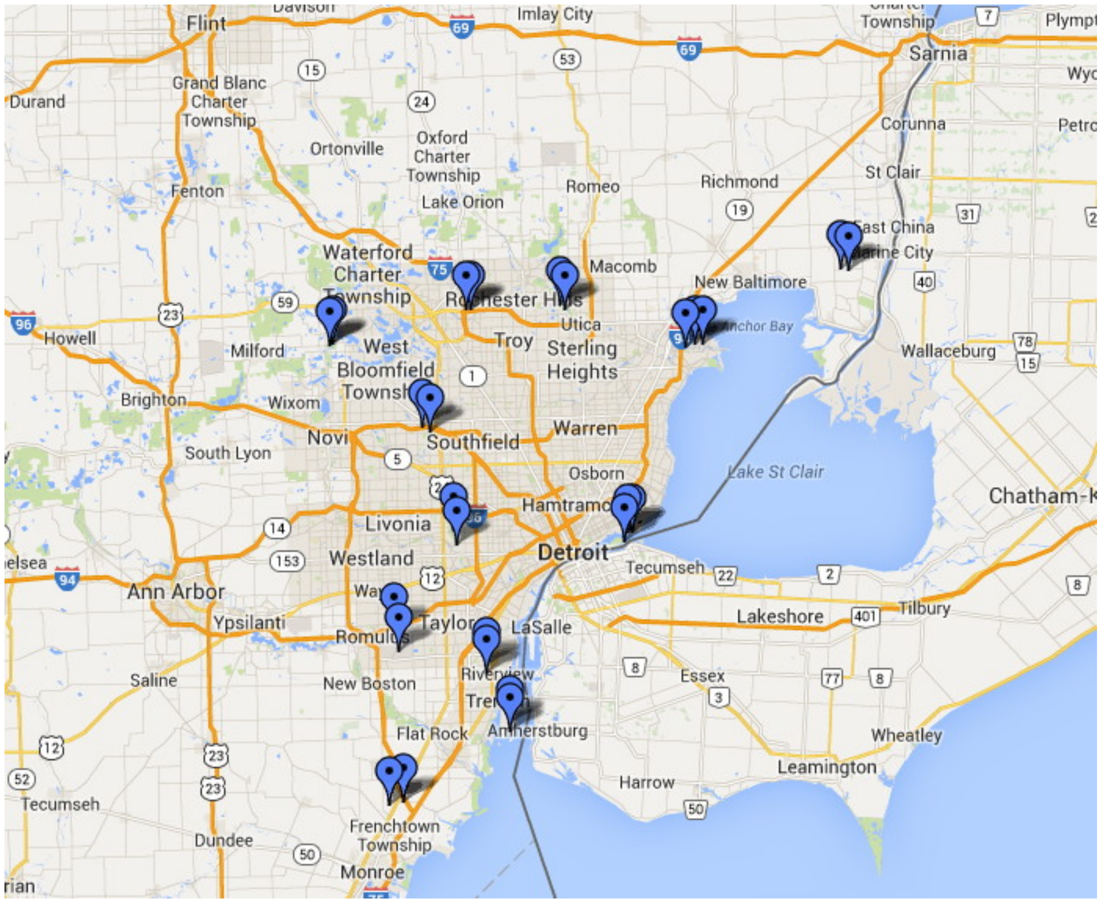

http://en.wikipedia.org/wiki/List_of_Nike_missile_sites

All VERY SCARY, But Mostly VERY SLOW

- As long as weapons were coming in via bombers, events moved in **human time frames**: the military had hours to predict aircraft courses/targets, and citizens had hours to potentially evacuate.

- That would change, if/when nuclear weapons were delivered by **missiles** rather than **bombers**.

- **Both sides** were working hard on missile technology in the cold war era, it's just that **America always assumed it would be the first to demonstrate expertise in ICBM-class technology.**

- The USSR's Sputnik Project demonstrated that that was a false assumption.

So What Was "Sputnik," Again?

- 57 years ago, Sputnik ("satellite" in Russian) was **the first artificial Earth satellite**, launched by the USSR from Baikonur Cosmodrome on Oct 4th, 1957. It orbited the earth for 92 days.
- It was a 23 inches metal sphere and had two pairs of external radio antennas. It weighed 184.3 pounds (112 pounds of that weight, 60.7%, was a power supply full of batteries that would last just 22 days)
- **Sputnik sent a piercing fast "beeping" sound** on 20.005 and 40.010 MHz. Want to hear it? You can, if you're curious: http://www.amsat.org/amsat/features/sounds/sputnk1b.wav
- NBC radio, in rebroadcasting Sputnik's beeping, said, "Listen now for the sound which forever more separates the old from the new" [quoted in "Red Moon Over the U.S.," *Time*, Oct 14th, 1957]

Sputnik 2

- If there had **only** been the first Sputnik, its impact might have been limited. However, **less than a month later**, on Nov 3rd, 1957, the USSR launched again.
- The 2nd Sputnik had a **1,121 pound payload,** included a live dog. The payload weight was what was important, however: that was the weight of a nuke.
- (BTW, poor Laika was doomed from the time she was launched because the USSR hadn't mastered re-entry as of Nov 1957)
- Anyhow, Sputnik 2 was a very big deal.
- President Eisenhower responded with a major radio and television address...

http://upload.wikimedia.org/wikipedia/commons/a/ab/Sputnik2_vsm.jpg

Eisenhower's Radio & TV Address of Nov 7th, 1957

- "My subject tonight is **Science in National Security.**"
- R&D for defense was already running over **$5 billion a year**
- Discussed the **B-52** jet bomber, and **nuclear subs & carriers**
- Mentioned that we have **more and better nukes** than the Soviets
- Discussed creation of an **early warning radar system**
- "According to my scientific friends [JES: remember, he was formerly president of Columbia University], one of our greatest, and most glaring, deficiencies is the failure of us in this country to give high enough priority to **scientific education** and to the place of science in our national life."
- "They believe that a second critical need is that of giving higher priority, both public and private, to **basic research**."
- Speech transcript linked from: www.eisenhower.archives.gov/research/online_documents/sputnik.html

Some Specific Outcomes, Both Good and Bad

- The **National Defense Education Act** was created, providing scholarships for those studying math and science at universities
- Funding for **basic research via the NSF dramatically increased**
- NASA and ARPA were both created in 1958
- The 'SAGE' Project was started (more expensive than the Manhattan Project that resulted in the first atomic bombs)
- And speaking of nuclear weapons, US and Soviet nuclear weapon counts started ramping up dramatically during this period. If you believe "more nukes=good," I suppose that's great, but if you believe "more nukes are bad," well there you go...

Stockpiles Ramp Up Sharply Beginning in 1957 (U.S. Finally Gets Back to 1957 Levels In 2007)

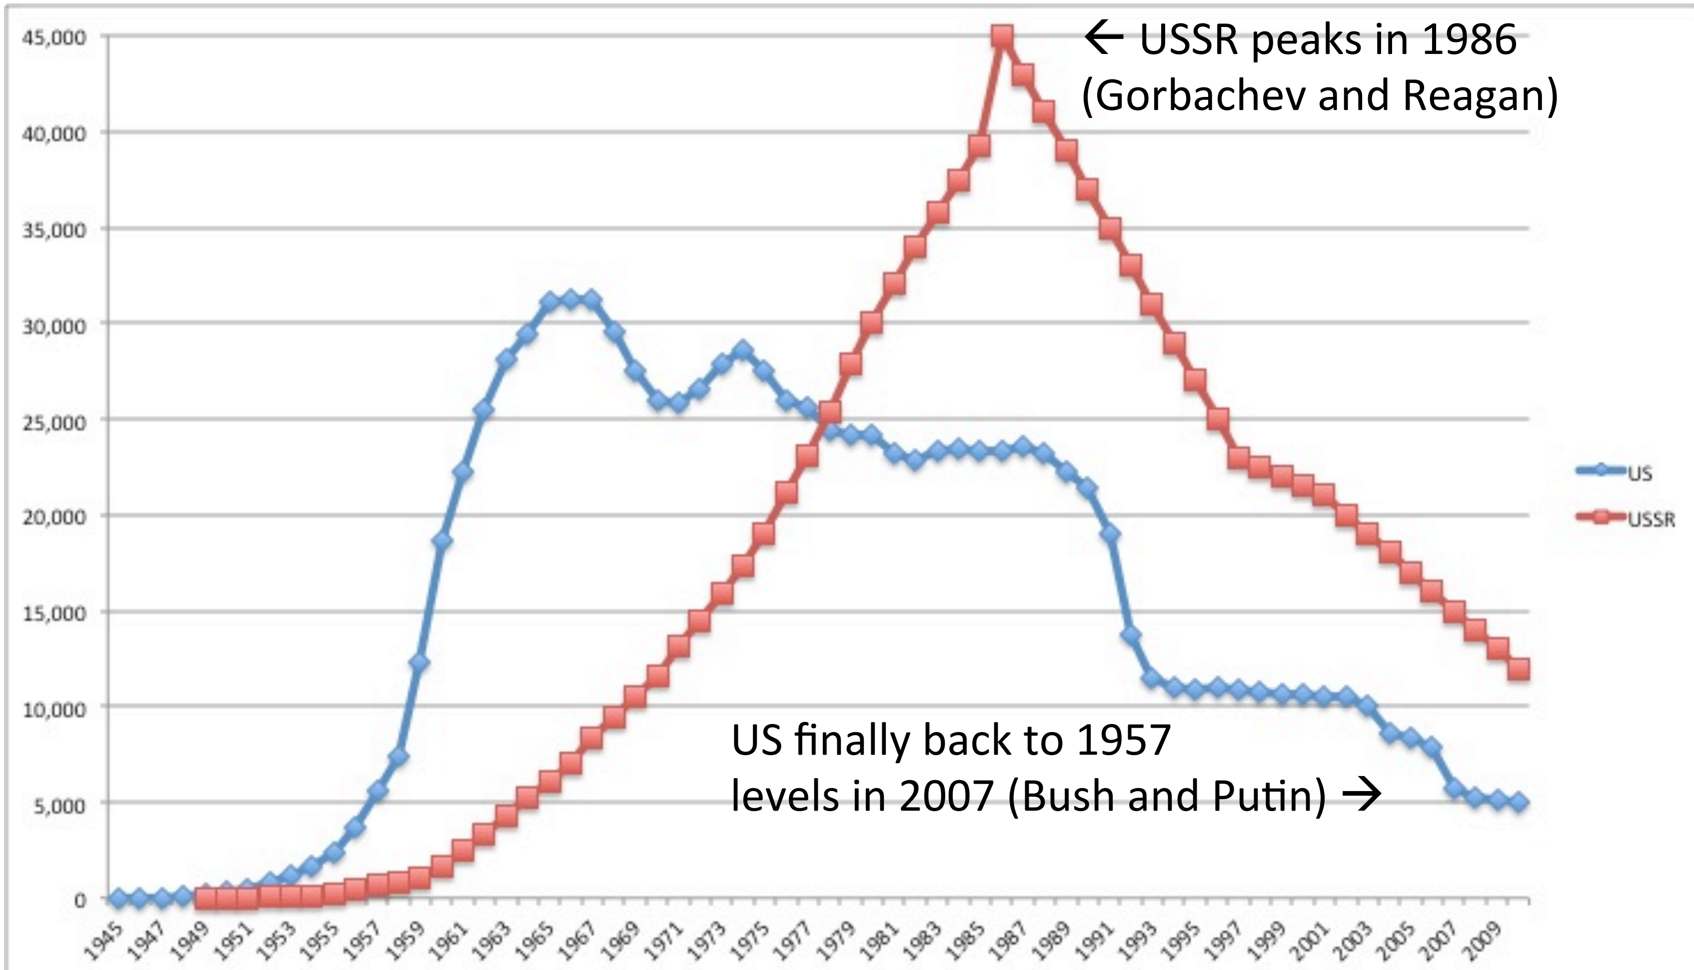

"Global Nuclear Weapon Inventories, 1945-2010," Bulletin of the Atomic Scientists, 4/2010.

The Key Project Many Never Heard Of -- SAGE

- "Semi-Automatic Ground Environment:" massive system for **US air defense, operating from 1959 through 1979**.
- Total SAGE costs are unclear, but it was believed to be **several times more expensive than the Manhattan Project** (the Manhattan Project cost ~$2 billion (1945 dollars); SAGE cost $8-12 billion (1964 dollars), equal to $61.5-92.2 billion today
- Noteworthy as one of the first wide area networks, built using Bell 101 (110 baud) modems (the first equipment to use ASCII).
- **24 of the largest computers ever built**, IBM AN/FSQ-7's, each with 60,000 tubes, drawing 3 megawatts, and weighing 250 tons.
- The work was lead by **MIT Lincoln Laboratory** however the **University of Michigan** had been a strong contender to do the R&D work (see "From Whirlwind to MITRE: The R&D Story of the SAGE Air Defense Computer" at pages 186, 216, etc.)

Michigan State Univ. During That Same Period

- Michigan State was lead by Pres. John Hannah from 1941-1969: http://www.archives.msu.edu/collections/presidents_hannah_j.php

- Quoting from that page "Hannah's government service included: [...] **Assistant Secretary of Defense, 1953–1954**; Chairman of the Commission on Civil Rights, 1957–1964; and, **Chairman of the United States Section of the Permanent Joint Board on Defense, 1954-1964.**"

- *PJBD???* "... formed in 1940 by the Ogdensburg Declaration, to create a body that could consider, in the broad sense, the security and defense of the northern half of the Western Hemisphere." blogs.ottawa.usembassy.gov/ambassador/index.php/tag/permanent-joint-board-on-defense/

Bottom Line on the Cold War Era

- The 1950s were a very scary time due to the threat of nuclear war, and the whole country was focused on nuclear defense (including both Michigan and Oregon).

- Many massive scientific research and development programs were getting underway, often motivated by defense concerns.

  These programs often leveraged academic expertise in partnership with industry. This period also saw noteworthy service by key academics in government.

- RAND also began research on the survivability of national command and control channels in the face of a nuclear first strike

IV: Mr. Paul Baran

- Born in 1926 in Grodno, Poland (now Belarus).
- Baran's family moved to the US in May, 1928.
- BSEE from Drexel (1949)
- 1949: tech at Eckert-Mauchly Computer Company, working on UNIVAC (the first commercial computer)
- MSE (Computers) UCLA (1959)
- At RAND from 1959 through 1968
- Wrote "On Distributed Communications"
- Testified before Congress on computer privacy in 1965 (first computer scientist to do so)
- Recommended the divestiture of ARPANET
- Founded Metricom in 1986, a wireless company that offered "Ricochet" wireless Internet service (even in Eugene, Oregon!)
- Died March 2011 in Palo Alto, at age 84
- www.ieeeghn.org/wiki/index.php/Paul_Baran

A Brilliant Cold War Era Scientist: Paul Baran

- It was in this apocalyptic cold war environment that Paul Baran, a scientist at RAND, worked on assured/survivable communications.

- Baran came up with many of the key concepts that were eventually incorporated into the modern Internet.

- Sadly, most people, including many long time Internet engineers, have never heard of Baran, in part because he was a very modest and self-effacing guy. I mention him here in the hope he will be remembered by at least those of you in this room today.

- Baran was also far ahead of his time. In particular, some of his ideas were hindered by being potentially economically disruptive.

"On Distributed Communications," August 1964

- "This Memorandum briefly reviews the distributed communications network concept and compares it to the hierarchical or more centralized systems. **The payoff in terms of survivability** for a distributed configuration [...] is demonstrated.
  "The requirements for a future **all-digital-data** distributed network which provides common user service for a wide range of users having different requirements is considered. The use of a **standard format message block** permits building relatively **simple switching mechanisms** using an adaptive store-and-forward routing policy to handle all forms of digital data [...]"

  Paul Baran, "On Distributed Communication," Volume 1, http://web.archive.org/web/20110330091634/http://www.rand.org/content/dam/rand/pubs/research_memoranda/2006/RM3420.pdf

The Logic of Positive Control

- In the 1950s, military command and control communications happened via AT&T Long Lines or high frequency ("skywave") radio. Both could end up disrupted in the event of nuclear war.
- Cut off from national authorities, nuclear base commanders might feel compelled to use their "best judgment" about whether or not to launch a counterstrike. This was NOT viewed as desirable.
- The goal for nuclear weapons has **always** been "positive control:"
  -- **ALWAYS** launch IF you receive a properly formatted and authenticated Emergency Action Message (EAM) from the National Command Authority, and
  -- Equally importantly, **NEVER** launch if you haven't.
- Launch instructions were (and are) apparently surprisingly succinct; see the discussion at http://mt-milcom.blogspot.com/p/what-is-emergency-action-message-or-eam.html

Baran's First Idea: Leverage AM Radio Stations

- Baran proposed a point-to-point ground wave communication network between AM broadcast stations for **"Minimum Essential Communications"** (e.g., nuclear launch orders).
- While traditional phone switching centers might have been destroyed and HF radio rendered unusable, many AM radio stations would have continued to work, and AM radio stations were already part of a national emergency warning system. They could easily have served as a path to pass along "go" codes.
- This plan was rejected by military authorities.
- Given the *fully authenticated* false alarm of February 20, 1971 (http://conelrad.blogspot.com/2010/09/code-word-hatefulness-great-ebs-scare.html) this might have been a *very* good decision. [for a subsequent failure, see also http://www.csmonitor.com/USA/2011/1109/Did-the-national-Emergency-Alert-System-mistakenly-play-Lady-Gaga]

Baran's Second Idea

- Baran then decided:

  "As I can't figure out what essential communications [may be] needed, let's take a different tack. I'll give those guys **so much damn bandwidth that they wouldn't know what in Hell to do with it all.** In other words, I viewed the challenge to be the design of a **secure network** able send signals over a network being cut up, and yet having the signals delivered with perfectly **reliability.** And, with **more capacity than anything built to date.**"

  See www.ieeeghn.org/wiki6/index.php/Oral-History:Paul_Baran

- That's what he proceeded to design. Like a good engineer, he also worked out cost estimates. They turned out to be surprisingly low.

Aside: Goring A Cash Cow Is Dangerous

- Baran's digital system was estimated to cost roughly $60 million in 1964 dollars.

- It could have rapidly replaced the traditional AT&T long distance telecom system in use by the American military, which cost American taxpayers $2 billion/year at the time.

- Excellent (but potentially financially disruptive) ideas can sometimes be smothered at birth if they rock the boat <u>too</u> much.

- Bottom line: Baran's packet switched network didn't get built (then).

Other Key Points to Note

- Military needs (assured communications to direct nuclear war fighting) were the original driver for Baran's work, as well as for early computing projects such as Project SAGE.
- Abundant **capacity** was important (even if exact capacity requirements were unclear); "when in doubt, overbuild."
- **Reliability** (through architectural redundancy) was key.
- **Business models** mattered when Baran's packet switched network ideas were being evaluated (a lesson he *rediscovered* decades later when building out Metricom's Ricochet service, a pioneering wireless network that serviced Detroit and Eugene in the mid-1990s)

V. "So Can We Finally Get To The ARPANET Era Now?"

No. Let's Go Backwards A Bit Further, Instead.

Far Before The Internet, People Still Wanted to Communicate, Share News, and Be Entertained

- This was manifest in many diverse projects and technologies:
  -- Postal mail
  -- Telegraphy
  -- Telephony
  -- Teleprinter/Teletype Networks
  -- Broadcast Radio

- All of these have helped set the stage for the Internet
- All have had a mix of academic, government and commercial involvement (the government has always been interested in improving communications, particularly for national defense requirements)
- Many have had ups and downs as business cycles go...

A Postal "Network," Falmouth MA to Savannah GA

1775.

The Continental Congress, by the act of July 26, 1775, provided that a Postmaster General should be appointed for the United Colonies, who should hold his office at Philadelphia, and should be allowed a salary of $1,000 per annum for himself, and $340 per annum for a secretary and comptroller, with power to appoint such and so many deputies as he deemed proper and necessary. The act continues:

"That a line of posts be appointed under the direction of the Postmaster General from Falmouth in New England to Savannah in Georgia, with so many cross posts as he shall think fit.

"That the allowance to deputies, in lieu of salary and all contingent expenses, shall be 20 per cent on the sums they collect and pay into the general post office annually, when the whole is under or not exceeding $1,000, and 10 per cent for all sums above $1,000 a year."

The Congress at once unanimously elected Benjamin Franklin as Postmaster General for one year.

Source: "Congressional Serial Set,"
http://books.google.com/books?id=4rMqAAAAYAAJ&lpg=PA102 (page 102)

A Few Notes About "The Postal Service"...

- Postal service in the 1700s bore little resemblance to postal service today.

- Sending a letter between two major American cities (such as New York and Philadelphia) might take weeks to go a distance of a little over a hundred miles.

- Letters were quite expensive. Ordinary people (who were often illiterate) might send just a single letter a year. (c.f., http://www.postalmuseum.si.edu/letterwriting/lw02.html)

- The postal system didn't even use standardized stamps until 1847.

The Postal Service Today: In a Business Death Spiral?

Executive Summary of Findings

The U.S. Postal Service will experience profound declines in its volumes of mail and its net income over the next decade under its current business model, presenting a grave threat to its viability. Massive structural changes are required to avoid this outcome.

We forecast U.S. postal volumes to decrease from 177B pieces in 2009 to around 150B pieces in 2020 under business-as-usual assumptions. Notably, volumes will *not* revisit the high-water-mark of 213B pieces in 2006 – on the contrary, the trajectory for the next 10 years is one of steady decline, which will not reverse even as the current recession abates. Expressing the decline in terms of pieces per delivery point highlights the challenge: we project pieces per household per day to fall from *four* pieces today to *three* in 2020 – driven by decreasing volumes delivered to an increasing number of addresses. We also project a rapid mix shift from very lucrative First-Class Mail to less-profitable Standard Mail. The volume decline and the mix shift, coupled with an increasing cost base, will cause profits to experience steep, unrelenting declines. Starting with the 2009 loss of $4B, we expect a steady string of increasing losses, culminating with an approximately $15B loss in 2020 (based on USPS and McKinsey cost forecasts).

These declining volumes are unlikely to reverse. First-Class Mail is succumbing to the online diversion of bills, invoices, statements, and payments. Senders are aggressively attacking the cost of paper transactions – both for sending mail and processing responses. More Consumers will move online when key barriers, like security concerns, are removed.

Source: "Projecting U.S. Mail Volumes to 2020," https://about.usps.com/future-postal-service/gcg-narrative.pdf

Telegraphy: The Morse System, 1837-1844

- The "Morse System" was patented in the United States by Samuel Morse in 1837. Information was sent as dits (dots) and dahs (dashes), also known as "Morse Code."

- In 1844, a message was sent by wire from the Capitol in Washington DC to Mt. Clare Depot in Baltimore,

  "WHAT HATH GOD WROUGHT"

  A profoundly insightful rhetorical question. :-)

- Morse code was sent by an operator using a key, and received messages were initially recorded via marks on paper tape (later operators found that they could transcribe messages via "ear")

1844: Early Wireline Telegraphy Had Close Ties to Academia... and to the U.S. Federal Government (Ezra Cornell Pulled Cable From DC to Baltimore)

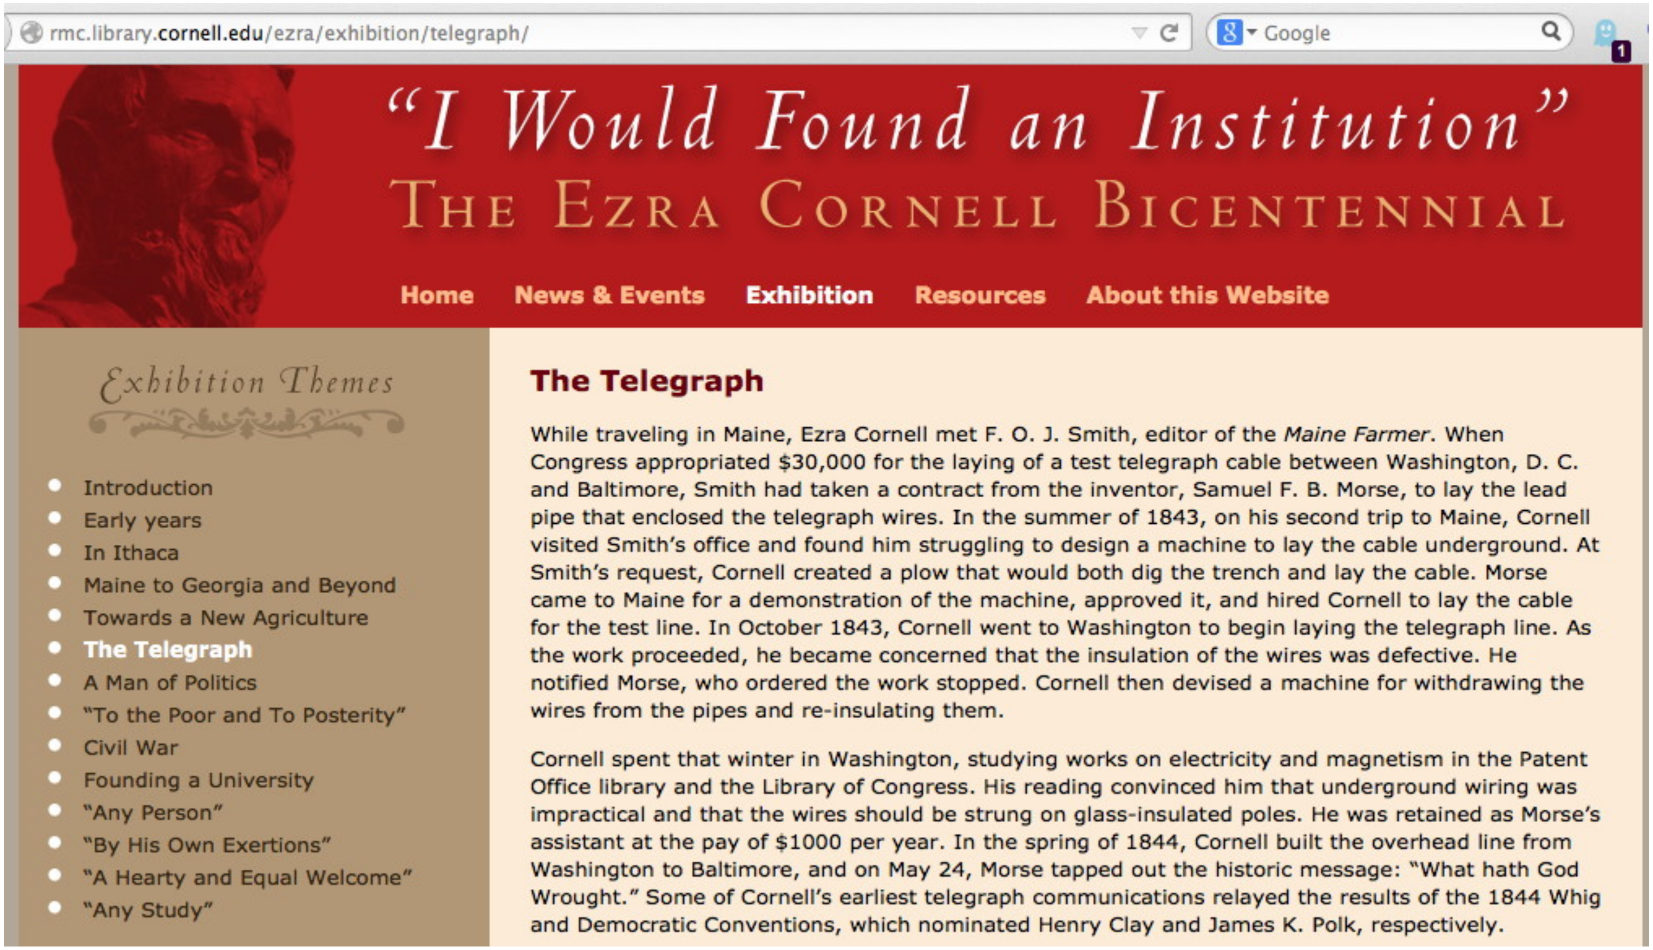

A Brief Diversion: The Pony Express

- During just two years (from 1860-1861) the "Pony Express" delivered ~35,000 letters (and managed to lose $200,000)

- **Why** did the Pony Express end in 1861?

  The **wireline telegraph** linked the coasts in 1861

  The telegraph also proved its value during the American Civil War (1861-1865)

Lincoln and The Telegraph at War

But, The Last U.S. Telegram Was Sent in 2006...

www.nytimes.com/2006/02/06/technology/06telegram.html?_r=0&pagewanted=print

Google

The New York Times
nytimes.com

February 6, 2006

Telegram Falls Silent Stop Era Ends Stop

By SHELLY FREIERMAN

Sometime on Friday, Jan. 27, Western Union, bowing to the ascendancy of modern technology like e-mail, sent its last telegram.

Western Union had its beginnings in 1851 in Rochester. Messages were transmitted by Morse code over wire, then hand-delivered by a courier. Ten years later, the company completed the first transcontinental telegraph line.

That drove the Pony Express, which had been operating for less than two years, out of business by offering customers delivery of a message across the country in less than a day (the average Pony Express delivery took 10 days). In the relatively recent era of e-mail and instant messaging, telegrams were usually delivered by overnight courier services.

A sampling of some of the last telegrams wired on the last day included birthday wishes and efforts by a few people, probably Western Union employees, to send the final message. The company had no information about the time of the very last message or the cities it bridged. At the height of business in 1929, more than 200 million telegrams were sent around the world. Just under 21,000 were sent last year.

What About Telephony?

http://en.wikipedia.org/wiki/Candlestick_telephone#mediaviewer/File:Genevieve-Clark-Bain.jpeg (circa 1915)

Telephony's First "Long Distance" (60 mile) Line: French Corral to French Lake, California (1877)

Source: http://www.noehill.com/nevada_county_california/images/telephone_plaque.jpg
See also: http://malakoffdigginsstatepark.org/?page_id=587

Coast-to-Coast? Less Than A Hundred Years Ago...

1915: First Transcontinental Telephone Call

Jan. 25, 1915 - If you collect U.S. postage stamps, you might have noticed one issued in February 1998 to commemorate AT&T's 1914 construction of the first transcontinental telephone line.

AT&T began building the nation's original long distance network in 1885. Starting from New York, the network reached Chicago in 1892. But, because an electrical signal weakens as it travels down a wire, that distance was close to the limit for a line built of thick copper. With the 1899 introduction of loading coils, which slow the rate at which a signal weakens, construction proceeded west. By 1911, the network stretched as far as Denver, but had reached the distance limit for loading coils.

In 1908, AT&T President Theodore Vail had made a transcontinental telephone line a major goal, even though he knew the technology to build one did not exist. The next year, Chief Engineer John J. Carty raised the stakes when he announced in San Francisco that AT&T would open a transcontinental line in time for the city's 1915 exposition to mark the completion of the Panama Canal.

But, without a scientific breakthrough, AT&T couldn't make good on that bet. To improve the company's odds, Carty not only hired physicist Dr. Harold Arnold to study the amplification of electrical signals, he also spread the word in the scientific and electrical-engineering community that AT&T would pay handsomely for an electrical amplifying device.

On June 27, 1914, the last pole was erected at Wendover, Utah, on the Nevada/Utah state line, and was topped with the American flag. Commercial service was started on Jan. 25, 1915. Inset: Stamp from the U.S. Post Office's "Celebrate the Century" series commemorates AT&T's achievement.

First Transoceanic Cables? Just 58 Years Ago!

1956: Transoceanic Telephone Cables

Today, international calling is routine, with a half-million transatlantic calls from the United States every day. But it took years of effort by AT&T scientists and engineers to make it so. They began to study deep-sea, long-distance submarine cable as an alternative to the telegraph and short wave radio in the early **1930s**.

Transoceanic telegraph cables have been around for almost 100 years, but the more delicate voice signals with their higher frequencies could not make it across the ocean without a periodic boost to the voice signal. Thus the challenge was to design amplifiers installable three miles below the ocean surface and capable of operating for many years without requiring a repair. Built around an electron tube especially developed for this purpose, the amplifiers or repeaters, encased in flexible multi-metal housings, were spaced at 4-mile intervals along the cable. The cable also had an outer sheath of armor wires to provide strength and protection against abrasion and an extra copper sheath to keep out marine worms, which were known to plague telegraph cables.

Laying each of the two cables that made up the systems (one for each direction of communication) took weeks and was carried out during two consecutive summers. The first day of commercial service, **September 26, 1956**, saw a 75 percent increase in transatlantic telephone traffic.

And Now Just This Year...

arstechnica.com/tech-policy/2014/01/att-plan-to-shut-off-public-switched-telephone-network-moves-ahead-at-fcc/

MAIN MENU MY STORIES: 25 FORUMS SUBSCRIBE JOBS

AT&T plan to shut off Public Switched Telephone Network moves ahead at FCC

FCC to greenlight IP phone trials, but won't completely deregulate AT&T.

by Jon Brodkin - Jan 30 2014, 1:00am PST

Share Tweet 159

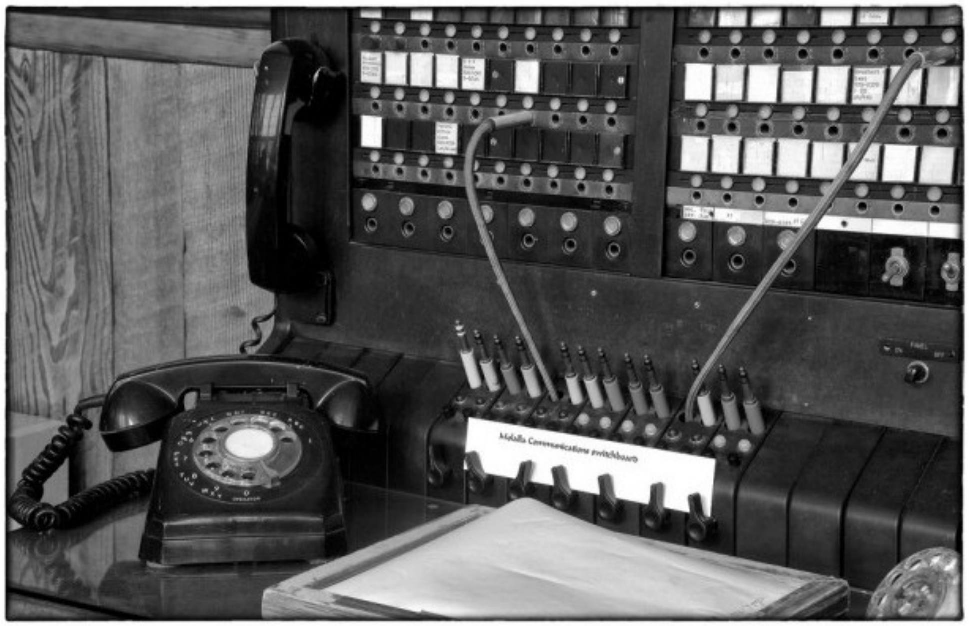

PhotoAtelier

On Thursday, the Federal Communications Commission is expected to take its first major step toward letting AT&T and other carriers replace the country's traditional phone system with one that works entirely over Internet Protocol networks.

What About <u>Teletype Networks?</u> UPI's National News Network

- Around the same time that telegraphy and telephony were becoming commercially successful, e.g., 1914, during WW I, national news syndicates began to use teletype networks to routinely transmit news around the country, see http://100years.upi.com/history_ms_1911.html

- National news stories (and later, even wire photos!) began to be transmitted nationally to subscribing newspapers.

- News syndicates weren't the only ones running **national scale networks** far before the modern Internet, either....

Aviation Weather Was Also Transmitted Via Teletype

- "Jul 1, 1928: The Commerce Department began using teletype machines to transmit aviation weather information. Among the first airport stations to receive teletypes were those at Hadley Field, N.J., Cleveland, Ohio, Chicago, Ill., and Concord, Calif. Those units were all connected with the central office at Washington, D.C., from which data were exchanged for all locations. **By Oct 1938, the teletype weather communications system had been extended to a total of 21,790 miles, covering all 48 states except Maine, New Hampshire, and South Dakota.**"

Source: https://www.faa.gov/about/media/b-chron.pdf

What About Wireless? Early Broadcast Radio

WWJ first signed on the air on August 20, 1920 under the call sign **8MK**, and was founded by *The Detroit News*; the mixed letter/number calls were assigned to the station by the United States Department of Commerce Bureau of Navigation, the government bureau responsible for radio regulation at the time. The 8 in the call sign referred to its location in the 8th Radio Inspection District, while the M in the call sign identified that the station operated under an amateur license.[3] It is not clear why the *Detroit News* applied for an amateur license instead of an experimental license. As an amateur station, it broadcast at 200 meters (the equivalent of 1500 AM).[4]

8MK was initially licensed to Michael DeLisle Lyons, a teenager, and radio pioneer. He assembled the station in the Detroit News Building but the Scripps family asked him to register the station in his name, because they were worried this new technology might only be a fad, and wanted to keep some distance. Later that year, Michael and his brother Frank, also assembled the first radio in a police car in Toledo, Ohio (with Ed Clark who started WJR, 760 AM, in Detroit). They captured a prowler using the radio, making national headlines. RCA got the contract to install radios in police cars across the country.

The Scripps family were also worried radio might replace newspapers if the medium caught on, so the family financially supported Michael. In fact, most early radio stations were built, for the same reason, by families who owned newspapers – out of concern that radio would put them out of business, on the basis that newspaper readers would find it more timelier to tune to listen to the headlines on radio at any given time than wait to read them in a daily newspaper the next day. Michael DeLisle Lyons was a descendant of Francois Bienvenu DeLisle, who served as Cadillac's lieutenant on the founding voyage of Detroit. Francois was also Detroit's first tavernkeeper.[5]

On October 13, 1921, the station was granted a limited commercial license and was assigned the call letters **WBL**. With the new license, the station began broadcasting at 360 meters (833 AM), with weather reports and other government reports broadcast at 485 meters (619 AM).[4][6]

Source: http://en.wikipedia.org/wiki/WWJ_%28AM%29#History

All Of Those Technologies Exhibit Commonalities

- They were all about **one-to-one messaging** (postal service, telegraphy, telephony), or **one-to-many messaging** (e.g., broadcast radio, news syndicate teletype networks, and weather teletype services, etc.)

- All technologies needed to operate over **substantial distances**

- **Faster** solutions rendered **slower** solutions **obsolete**

- **Cost was always critical**

- The **government** was almost always involved in one capacity or another, and quite often so was **academia,** as well as industry.

VI. *Now* We Come to The ARPANET

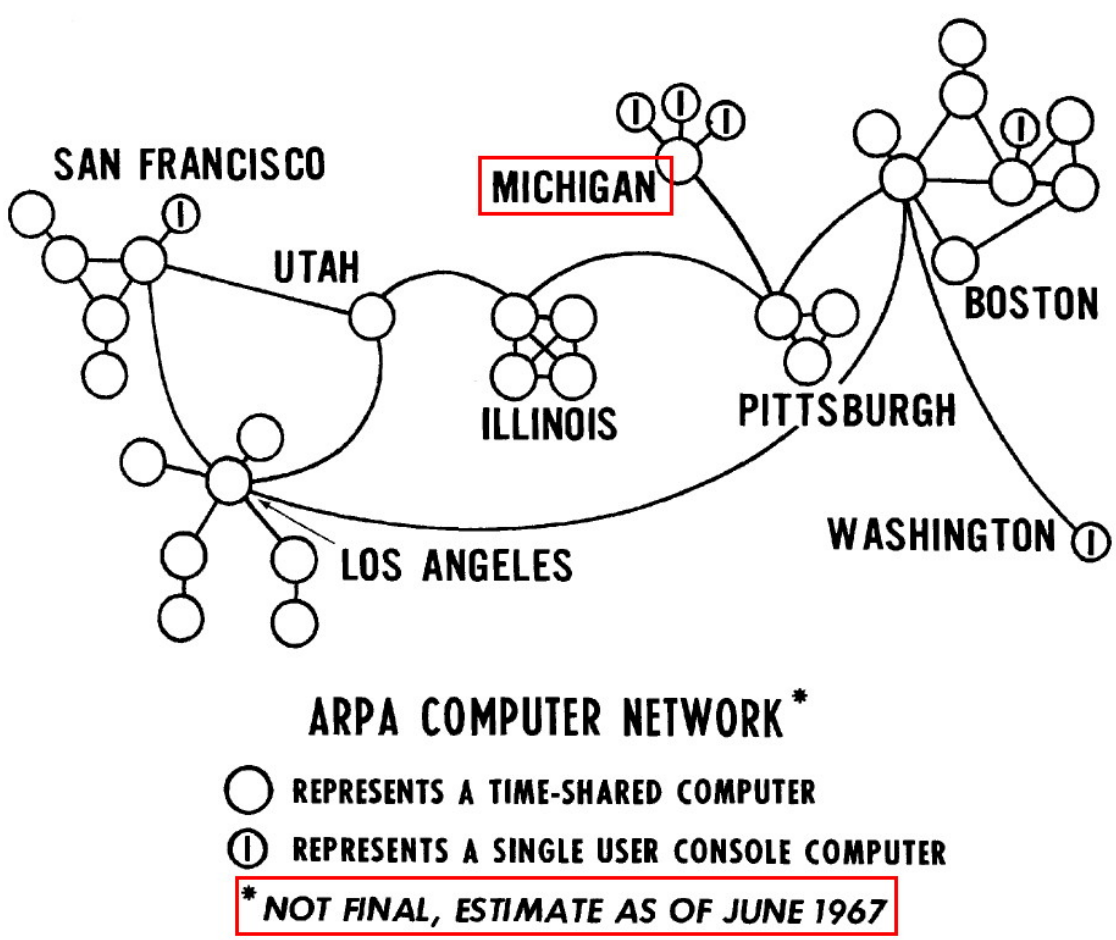

"Multiple computer networks and interconnection communications," Lawrence G. Roberts, http://dl.acm.org/citation.cfm?id=811680

The First ARPANET Meeting at the Pentagon, Oct 9-10, 1967 and UMich *Was* 1 of The Original 13

LIST OF ATTENDEES

G. Bell	Carnegie Institute of Technology	Attended October 10 only
A. Bhushan	Massachusetts Institute of Technology	
W. Clark	Washington University, St. Louis	Attended October 10 only
G. Culler	University of California, Santa Barbara	
L. Gallenson	System Development Corporation	
R. Kahn	Bolt, Beranek & Newman	
L. Kleinrock	University of California, Los Angeles	
M. Langtry	California Institute of Technology	Attended October 9 only
H. Magnuski	Bell Telephone Laboratories	BTL is not an ARPA contractor, only an interested observer
M. Pirtle	University of California, Berkeley	
L. Roberts	ARPA	
E. Shapiro	Stanford Research Institute	
R. Stotz	Massachusetts Institute of Technology	
T. Strollo	Bolt, Beranek & Newman	
B. Wessler	ARPA	
F. Westervelt	University of Michigan	

http://web.stanford.edu/dept/SUL/library/extra4/sloan/mousesite/Archive/Post68/ARPANETMeeting1167.html

Same Meeting: Routing and Baran's Work

- "3 Routing Procedures

  "3a It is anticipated that extremely dynamic traffic routing procedures will be employed, implemented by programs in each IMP. **In particular a version of the Baran (of RAND) hot potato method may he employed. The notion of the packet (an entity of 1000 bits maximum) was introduced, where a given message could be composed of many packets.** The routing mechanism would deal with the packet, thus packets of the same message may traverse different routes from source to destination. The problem now arises of packets of common message arriving at their common destination out of time sequence."

  http://web.stanford.edu/dept/SUL/library/extra4/sloan/mousesite/Archive/Post68/ARPANETMeeting1167.html

How Was The ARPANET *Supposed* To Be Used?

- **"Load Sharing"** ("... not a major consideration here")
- **"Message Service"** ("... **not** an important motivation for a network of scientific computers") [emphasis added]
- **"Data Sharing"** ("This type of use is particularly important to the military for **command and control**, information retrieval, logistics and war gaming applications.")
- **"Program Sharing"**
- **"Remote Service"** ("this will probably be the **most common** mode of operation until communication costs come down.")

See "Multiple computer networks and interconnection communications," Lawrence G. Roberts, http://dl.acm.org/citation.cfm?id=811680 [emphasis added]

Minimal Latency, NOT BANDWIDTH, Was Thought Key To The Original ARPANET Design...

Several network performance characteristics were specified. The average message delay for a short message to go from a source IMP to a destination IMP was to be less than one-half second for a fully loaded network. The probability of lost messages and message errors was to be very low. Interestingly, network capacity was considered third in order of importance and was defined to be the maximum bit rate that can be input at every node and still have the message delay remain less than one-half second; a 20 Kb network capacity was hoped for. A network model was presented.

"A HIstory of the ARPANET: The First Decade," pdf page 64, describing the 1968/69 RFQ
http://www.dtic.mil/get-tr-doc/pdf?Location=U2&doc=GetTRDoc.pdf&AD=ADA115440

"The Largest Single Surprise of the ARPANET Program"

The largest single surprise of the ARPANET program has been the incredible popularity and success of network mail. There is little doubt that the techniques of network mail developed in connection with the ARPANET program are going to sweep the country and drastically change the techniques used for intercommunication in the public and private sectors. By hindsight, one can easily see the reasons for this success. The primary prior existing communications techniques (the U.S. postal service and the telephone) have certain serious deficiencies:

"A HIstory of the ARPANET: The First Decade," pdf page 152
http://www.dtic.mil/get-tr-doc/pdf?Location=U2&doc=GetTRDoc.pdf&AD=ADA115440

By July 1977, There Were Still Some Academic Sites But Gov/Mil Sites Were Now Numerous....

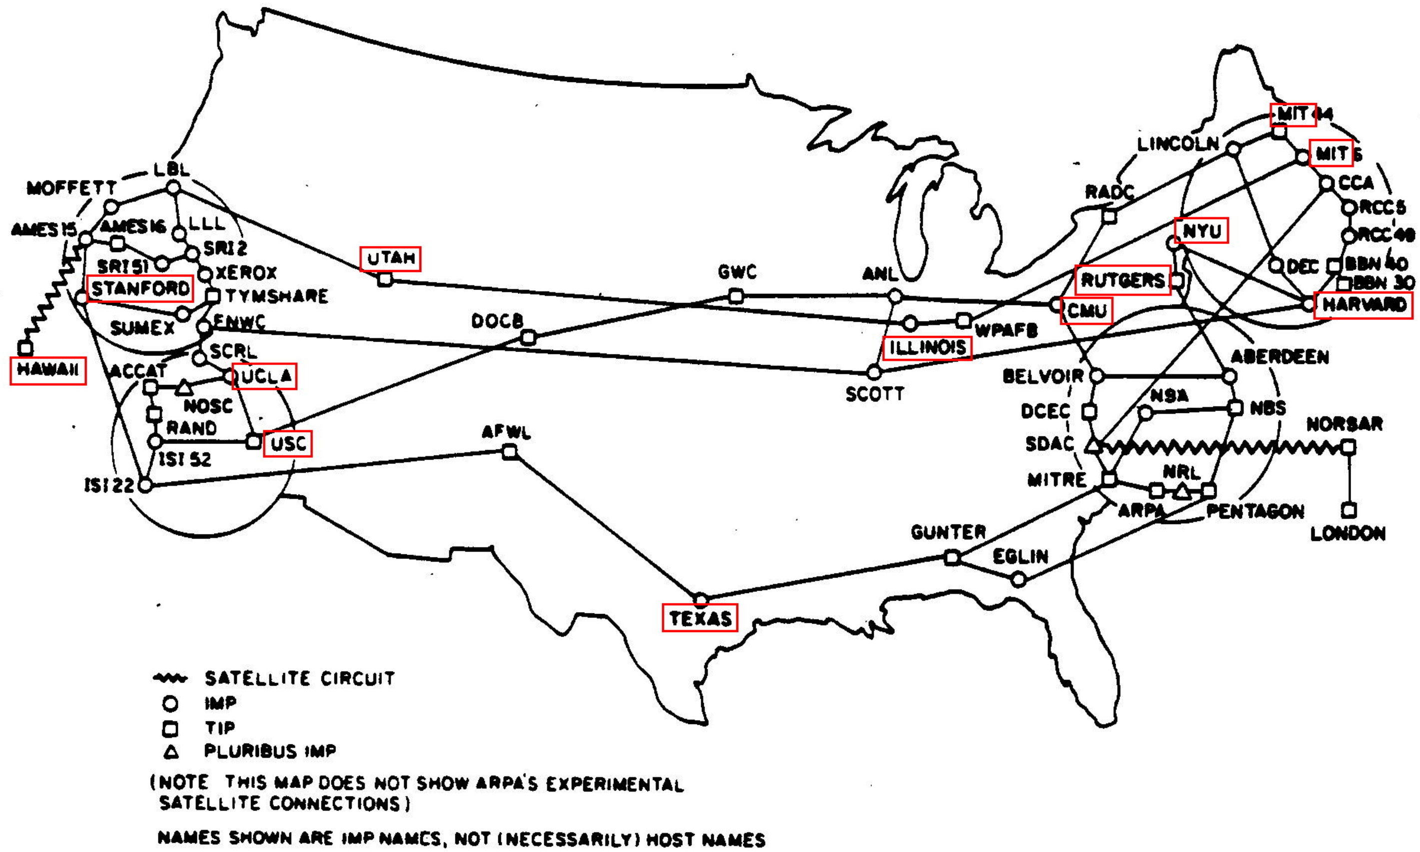

Source: http://som.csudh.edu/cis/lpress/history/arpamaps/
reportedly scanned from "ARPANET Completion Report," BBN, January 4th, 1978.

The Gov/Mil Guys Wanted A "Little Time Apart"

- "In the mid-1980s, NSF decided the time was right to try to link its regional university networks and its supercomputer centers together. This initial effort was called **NSFNET**.

  "By 1987, participation in the new NSFNET project grew so rapidly that NSF knew it had to **expand the capacity** of this new network. In November of that year, it awarded a grant to a consortium of IBM, MCI, and **a center at the University of Michigan called Merit** to create a network or networks -- or inter-net -- capable of carrying data at speeds up to **56 kilobits a second.** By July, 1987, this new system was up and running. **The modern Internet was born.**"

  http://www.nsf.gov/news/special_reports/nsf-net/textonly/80s.jsp [emphasis added]

The Six Node 56Kbps NSFNet Backbone Interconnected Supercomputer Sites

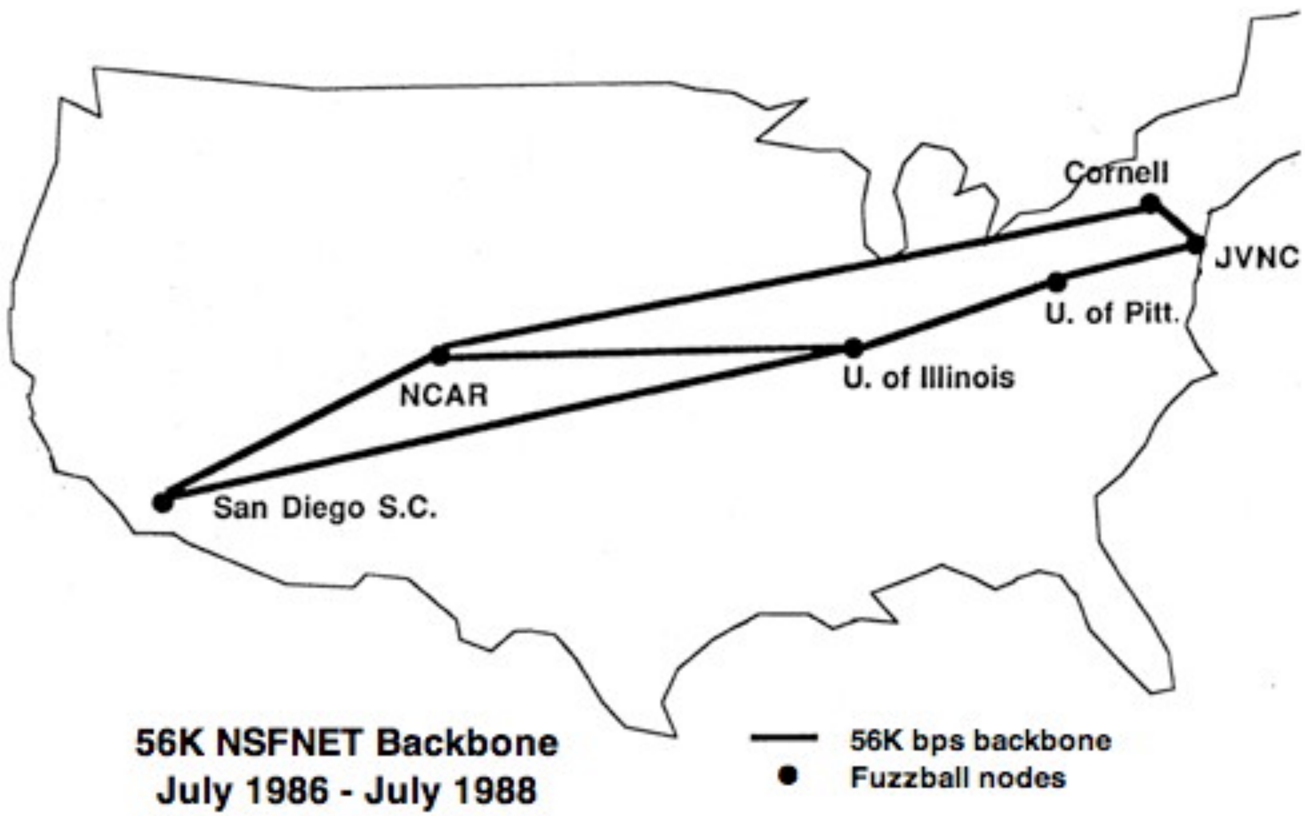

Source: http://en.wikipedia.org/wiki/NSFNET#mediaviewer/File:NSFNET-backbone-56K.png

VII. Going Faster

As the research and education network community evolved, there has been one relentless drumbeat: "Hey, go faster."

Good news? I think we're currently going as fast as we need to go.

Success!

But if so, now what?

The NSFNET at T1 (1.5Mbps) Speeds (~1990)

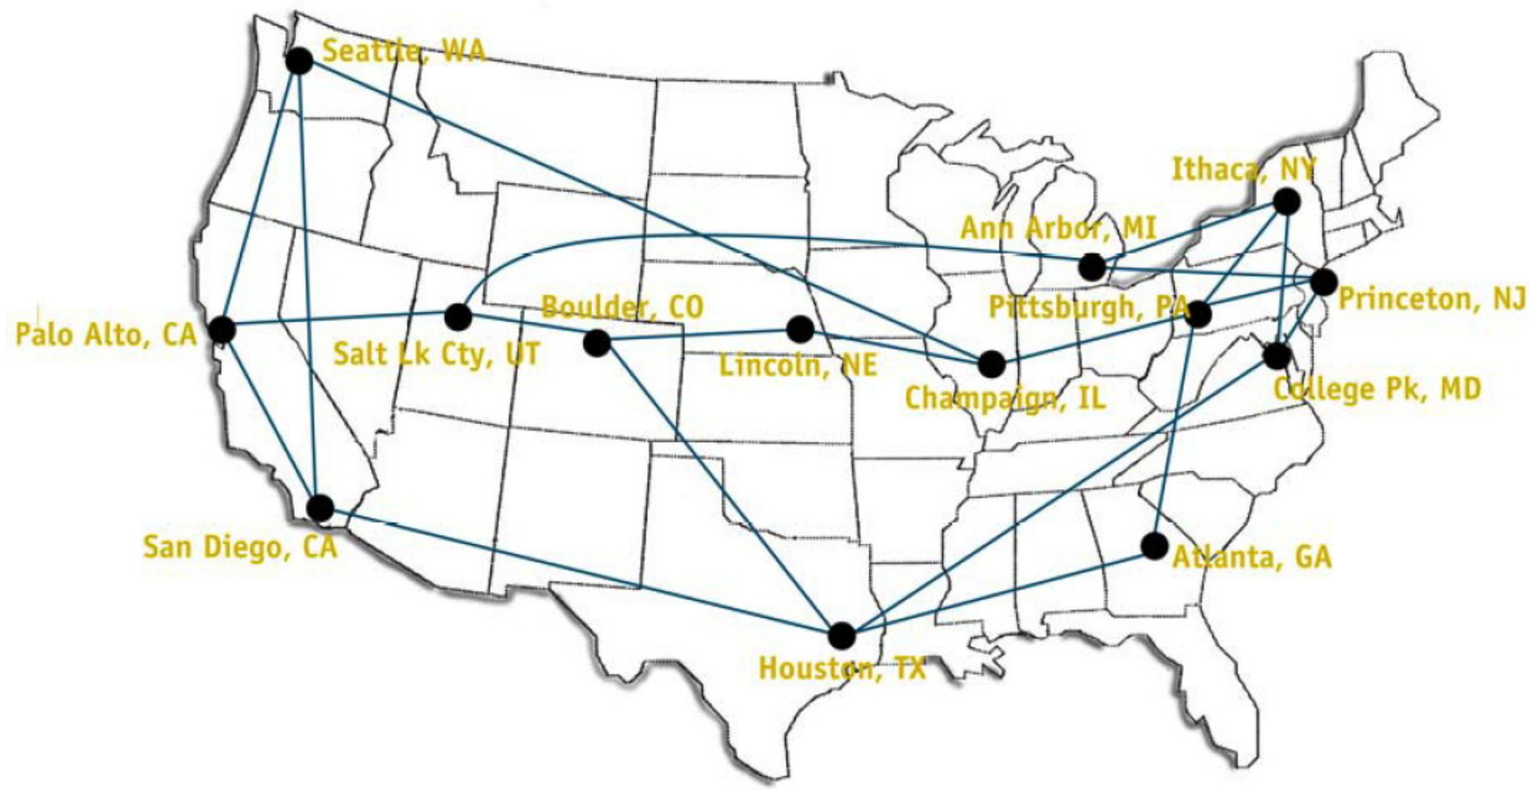

Figure 8: The Expanded T1 Network. An additional Atlanta node, deployed in 1990, brought to 14 the number of regional networks, midlevel networks, and supercomputer centers connected to the NSFNET at T1 speeds.

Source: https://web.archive.org/web/20111125110726/http://www.merit.edu/networkresearch/projecthistory/nsfnet/pdf/nsfnet_report.pdf

The NSFNET at <u>T3 (45Mbps)</u> Speeds (~1992)

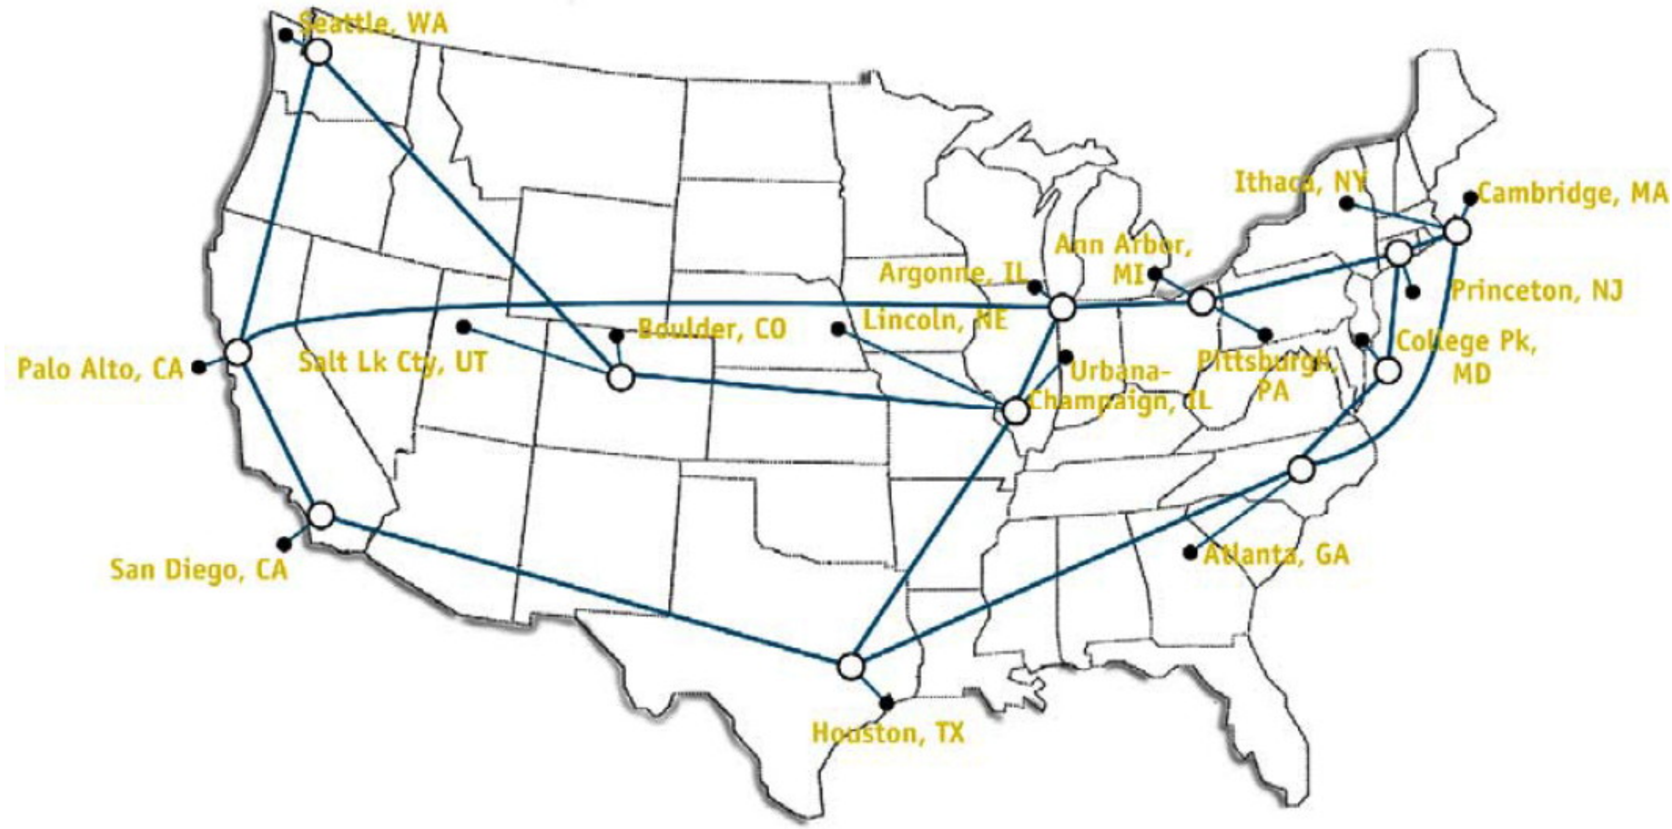

Figure 20: New T3 Backbone Service, 1992. Additional sites served by the T3 NSFNET included Cambridge MA (NEARNET) and Argonne IL (Argonne National Lab).

Source: https://web.archive.org/web/20111125110726/http://www.merit.edu/networkresearch/projecthistory/nsfnet/pdf/nsfnet_report.pdf

Remember the vBNS? May 20^{th}, 1997 (17 Years Ago)

- "Twenty-four Internet2 members were among the awardees today as Vice President Gore announced $12.3 million in grants to 35 research institutions across the United States. The National Science Foundation awards will allow them to connect to the **very high speed Backbone Network Service (vBNS)** and to communicate with other Internet2 members at **speeds 10 to 100 times greater than is possible through today's Internet. These grants bring to 44 the number of Internet2 institutions connected to the vBNS."**

  https://web.archive.org/web/19970607044127/
  http://www.internet2.edu/html/20_may_1997_nsf_awards.html

Just Two Years Later, 1999: Internet2 @ 2.4Gbps

www.zdnet.com/news/real-bandwidth-internet2-goes-live/101671

Real bandwidth: Internet2 goes live

Summary: *Researchers get first crack at the bigger bandwidth but will the benefits trickle down?*

By Maria Seminerio February 24, 1999 -- 00:00 GMT (16:00 PST)

Get the ZDNet Insights newsletter now

Comments 0 Votes 0 more +

The $500 million high-speed Internet2 network linking 37 U.S. universities, which went live Wednesday, will help researchers collaborate on projects such as real-time surgeries. Part of a planned demonstration today in Washington will include a segment where an offsite surgeon aides colleagues in a live operation, according to the University Corporation for Advanced Internet Development, sponsor of the Indiana University-based project.

High bandwidth

Internet2, also known as the Abilene Project, is a boon for the universities, which will benefit from the network's 2.4 gigabit-per-second speed said Michael McRobbie, vice president for information technology at Indiana University, in a statement. When the five-year project is completed in 2003, 140 universities will be connected to the network, which runs some 45,000 times faster than 56K modems, he said.

"Faculty and researchers now have access to unsurpassed networking capabilities for teaching and research," McRobbie said.

Consumer boon

The project is also expected to spur next-generation achievements in networking that could ultimately make the consumer Internet experience much richer, said Douglas Van Houweling, president and CEO of UCAID, in a statement.

"Just as the research networks of a decade ago produced technologies that have transformed the way we all work, learn and live today, Abilene will help develop the technology we all will use tomorrow," Van Houweling said. The 10,000-mile backbone runs on equipment donated by Qwest Communications, Cisco Systems, and Nortel Networks. The Abilene Project Network Operations Center is housed at Indiana University.

So What Was Internet2 Trying To Do, Again?

https://web.archive.org/web/20000301065128/http://www.ucaid.edu/abilene/html/techinfo.html

About Abilene ♦ Technical Info ♦ Network Status ♦ Connecting ♦ News & Highlights

SEARCH

Abilene Technical Information

architecture | status

Abilene is designed to support the Internet2 project by providing an interconnect backbone, connecting Internet2 gigaPoPs and universities at high speed, with advanced Internet2 functionality, and cost-effectively.

Key goals include:

- **Low latency**: Packets should experience little queueing delay and minimal propagation delay,
- **High reliability**: Failures should be few, and repair should be rapid, and
- **Advanced functionality**: Internet2 functionality, including multicast, quality of service, and measurements, should be supported.

Components of this area describe:

- **Network architecture:** How do the Core Architecture and the Access Architecture combine to provide Abilene functionality?
- **Advanced functionality:** How does Abilene support advanced Internet2 functionality?

Thinking About Those Objectives...

- **Low latency:** check -- just like the original ARPANET objective.
- **High speed:** check -- just like Baran's game plan (if in doubt, give 'em lots of bandwidth)
- **Highly reliable:** check -- another Baran objective.
- **Cost effective:** check -- another Baran theme.

The Shape of Internet2 in August 2000

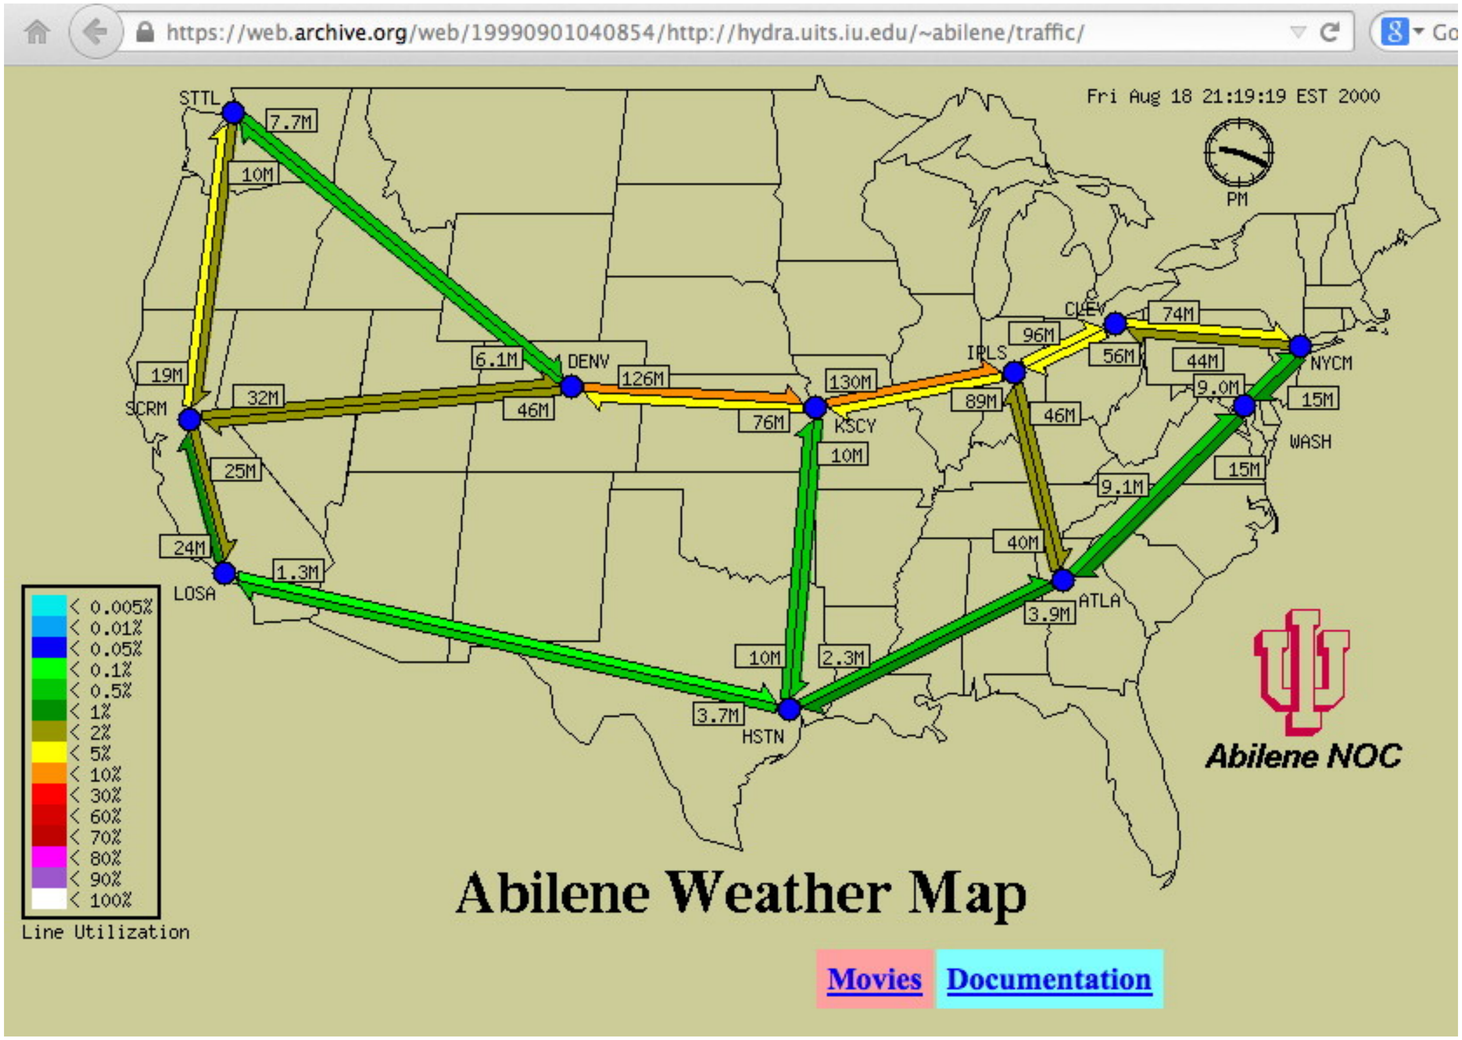

I2-NEWS: Abilene Network Upgrade to 10 Gbps Complete

- *To*: <i2-news@internet2.edu>
- *Subject*: I2-NEWS: Abilene Network Upgrade to 10 Gbps Complete
- *From*: "Michelle Pollak" <mpollak@internet2.edu>
- *Date*: Wed, 4 Feb 2004 10:54:41 -0500
- *Delivered-to*: i2-news@internet2.edu
- *Importance*: Normal
- *Organization*: Internet2
- *Reply-to*: mpollak@internet2.edu
- *Sender*: owner-i2-news@internet2.edu

```
ABILENE NETWORK UPGRADE TO 10 GBPS COMPLETE
New OC-192 circuits provide Abilene participants leading-edge networking
capability

WASHINGTON, D.C. - February 4, 2004 - Abilene, the most advanced research
and education network in the United States, today announced the completion
of its upgrade from 2.5 Gigabits per second (Gbps) to 10 Gbps.  The
Internet2(R) backbone network upgrade quadruples the capacity to more than
15,000 times faster than a typical home broadband connection.  Abilene
partners, Indiana University, Juniper Networks and Qwest Communications,
provided the equipment and services to successfully implement the network
upgrade.

The Abilene upgrade gives researchers, students and professors at more than
200 Internet2 member institutions a more robust network, on which to conduct
research; supports both Internet Protocol version 4 (IPv4) and native IPv6
networks; and increases overall network performance.

"With the upgrade now complete, Abilene continues to provide the Internet2
membership with a vehicle for advancement and development through a
leading-edge, national network environment," said Internet2 president and
CEO, Douglas Van Houweling.  "Many thanks go to our partners who made the
upgrade a smooth process with their ongoing support."

New routers were installed across the 13,000 route-mile national backbone,
and all but one OC-48c circuit was replaced with an OC-192c circuit.  In the
process, several router locations were changed but all remain collocated in
Qwest points of presence (PoPs).
```

https://web.archive.org/web/20050317090529/
http://archives.internet2.edu/guest/archives/I2-NEWS/log200402/msg00001.html

www2.ntia.doc.gov/grantee/university-corporation-for-advanced-internet-development

University Corporation for Advanced Internet Development

United States Unified Community Anchor Network (US UCAN)

As part of a longstanding project to connect essential community anchor institutions across the country, and facilitate closer collaboration and long-term benefits for education, research, healthcare, public safety, and government services, the University Corporation for Advanced Internet Development (UCAID) proposes a comprehensive 50-state network benefitting approximately 121,000 community anchors. The project proposes a large-scale, public-private partnership to interconnect more than 30 existing research and education networks, creating a dedicated 100-200 Gbps nationwide fiber backbone with 3.2 terabits per second (TBps) total capacity that would enable advanced networking features such as IPv6 and video multicasting. The project plans to connect community anchors across all disciplines into virtual communities with shared goals and objectives, including colleges, universities, libraries, major veterans and other health care facilities, and public safety entities, with additional benefits to tribes, vulnerable populations, and government entities.

Total Award: $62,540,162

Project Fact Sheet
Recipient Project Website

BTOP In Action

In March 2011, the University Corporation for Advanced Internet Development (UCAID), also known as Internet2, began upgrading its advanced middle-mile backbone network in support of the United States Unified Community Anchor Network (U.S. UCAN) project. Extending across 50 states, this upgraded network will enable high-speed broadband connectivity for up to 121,000 additional community anchor institutions. The project is a large-scale, public-private partnership that will interconnect more than 30 existing research and education networks, creating a dedicated fiber-optic backbone that will enable advanced broadband capabilities such as video multicasting, telemedicine, distance learning, and other life-changing Internet-based applications. UCAID's BTOP-funded fiber-optic network will also offer more than 8.8 TBps of capacity.

As of July 2011, UCAID has upgraded and activated more than 4,828 miles of its proposed 16,312 mile fiber network. In addition, UCAID anticipates the first coast-to-coast links, from New York City to Sunnyvale, Calif., to be installed and operational this summer with the entire network expected to be completed by early 2013.

In addition to the network infrastructure, UCAID also initiated outreach and awareness-building activities to demonstrate how advanced broadband applications and services could be deployed throughout the United States. In April 2011, UCAID hosted an annual two-day conference meeting in Arlington, VA, where more than 65 members attended to receive a status update on the project's construction efforts and learned about the new services the organization will provide through the advanced network.

Last Updated: October 18, 2011

Reports and Documents

Reports to NTIA:
Third Quarter, 2010

BTW, Nice Summary From Chris From 06/05/12

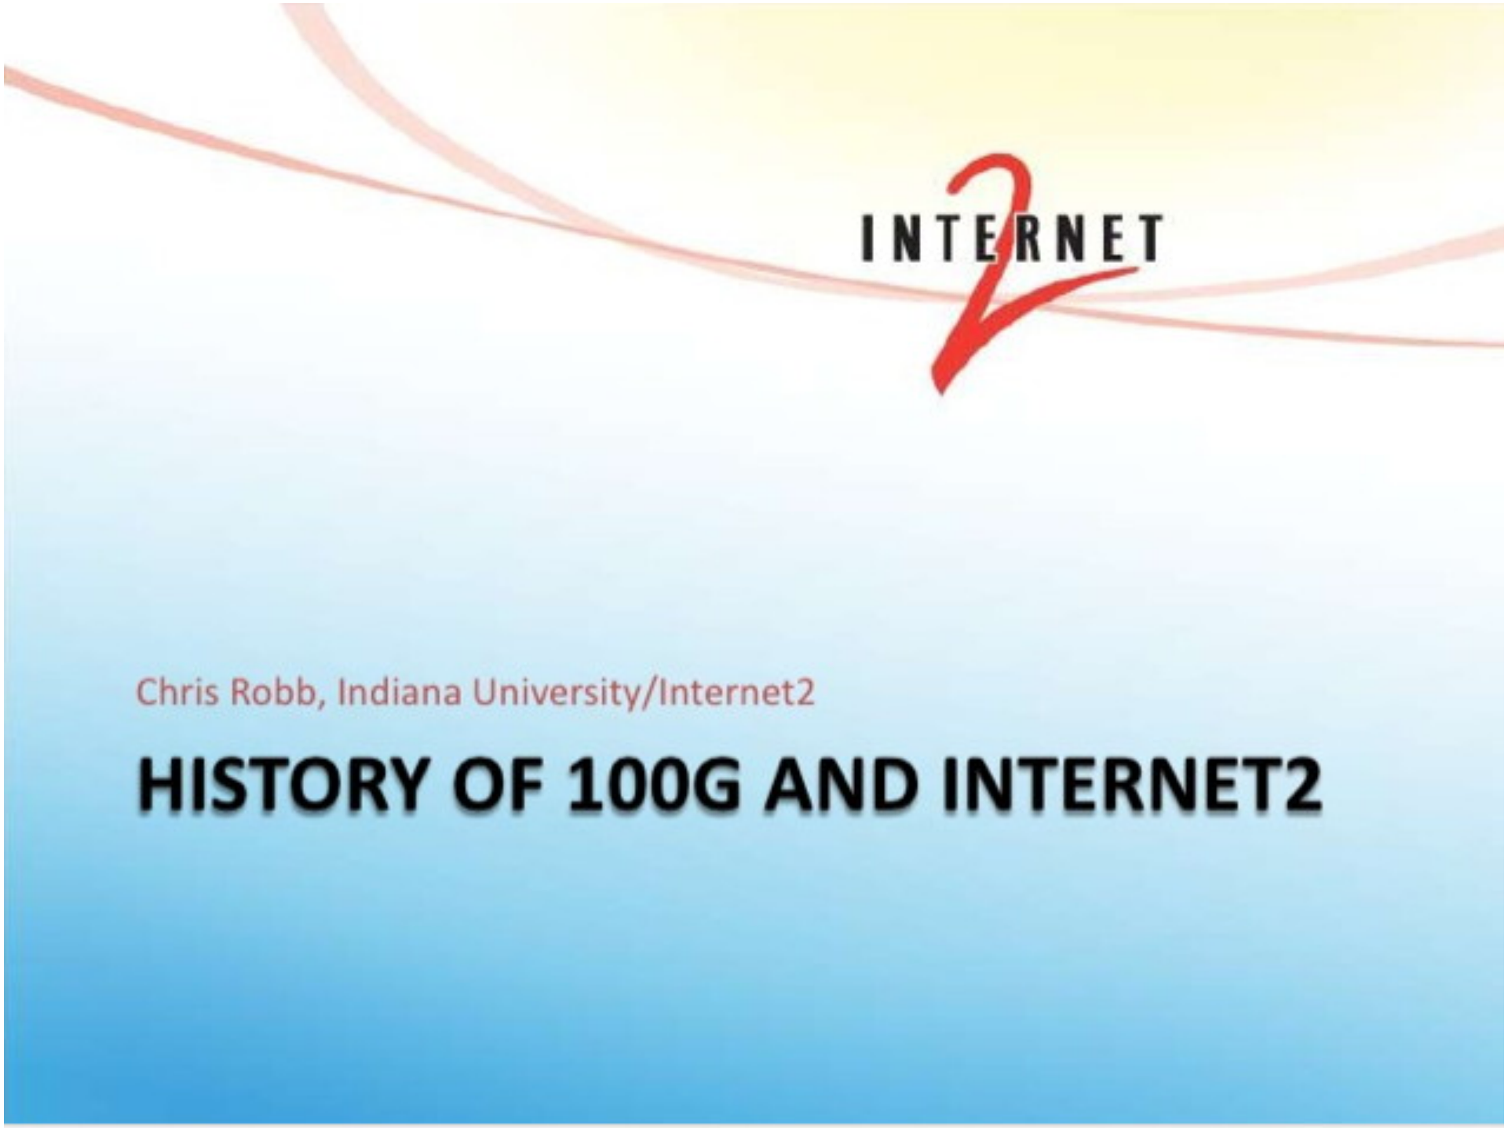

Internet2 *Aggregate* Traffic Over The Last Year

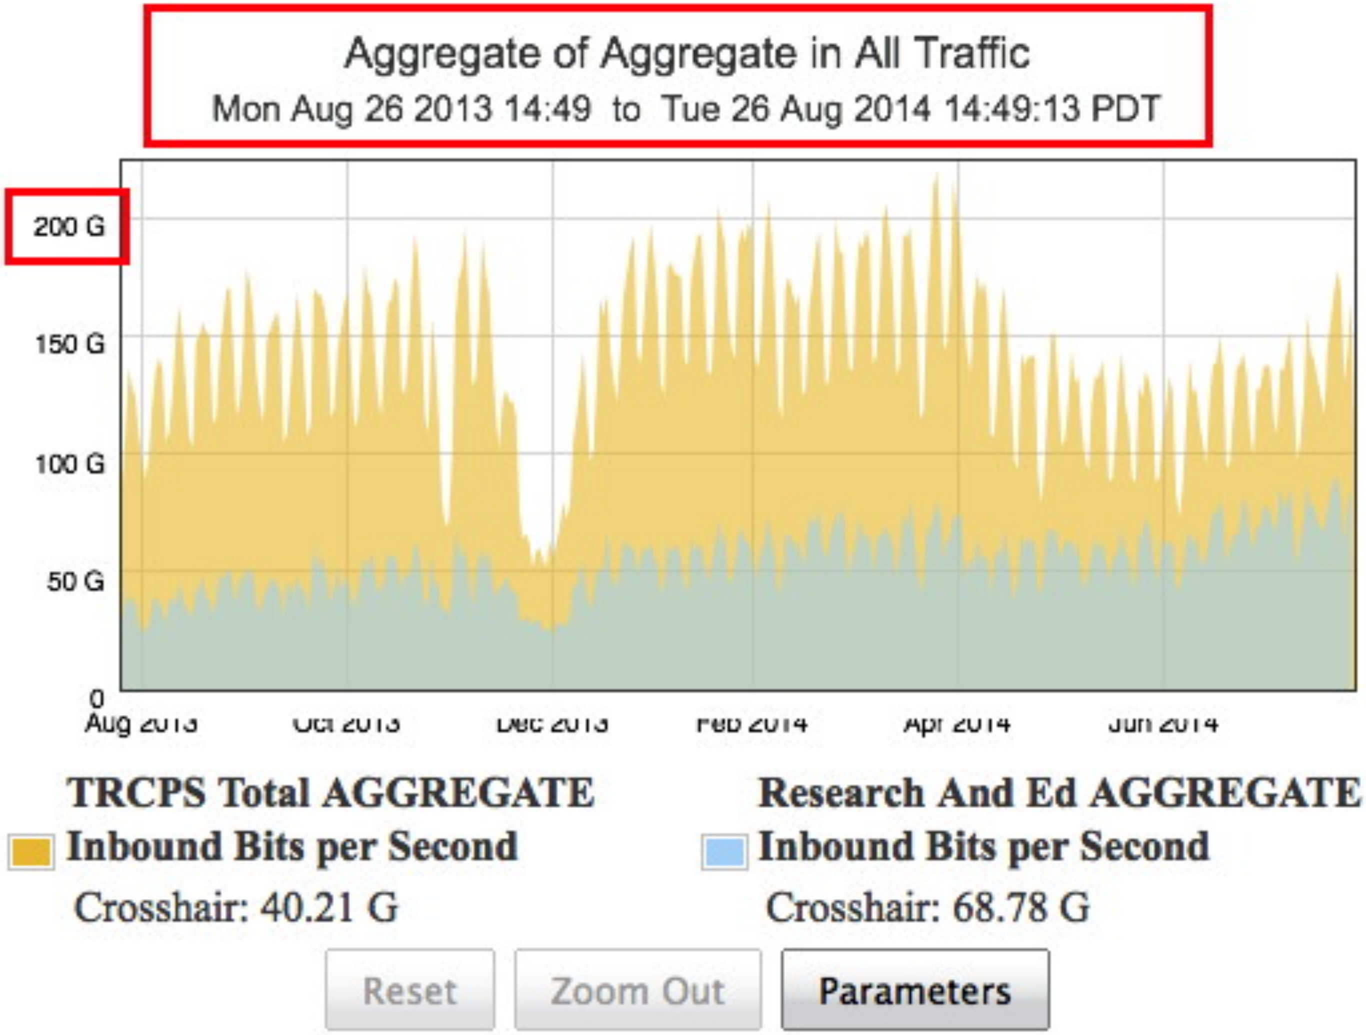

Source: http://noc.net.internet2.edu/

Thinking About The Preceding Graph A Little...

- Interpreting that graph requires a little guidance. It shows the SUM of ALL traffic into to the Internet2 network. The TOTAL input traffic shown there peaked at a little over 200 Gbps. That represents the aggregate of traffic coming from ALL regional optional networks. Individual connectors and **individual backbone links will** NOT normally see traffic at that level. Individual 100 Gbps links will have plenty of headroom, even when striving to avoid any possibility of congestion/packet loss.

- That that graph is the superposition of two datasets: research and education traffic (greyish) and Transit Rail/CPS traffic (goldish).

- Commodity Internet peering traffic increasingly dominates community R&E traffic.

Last Month...

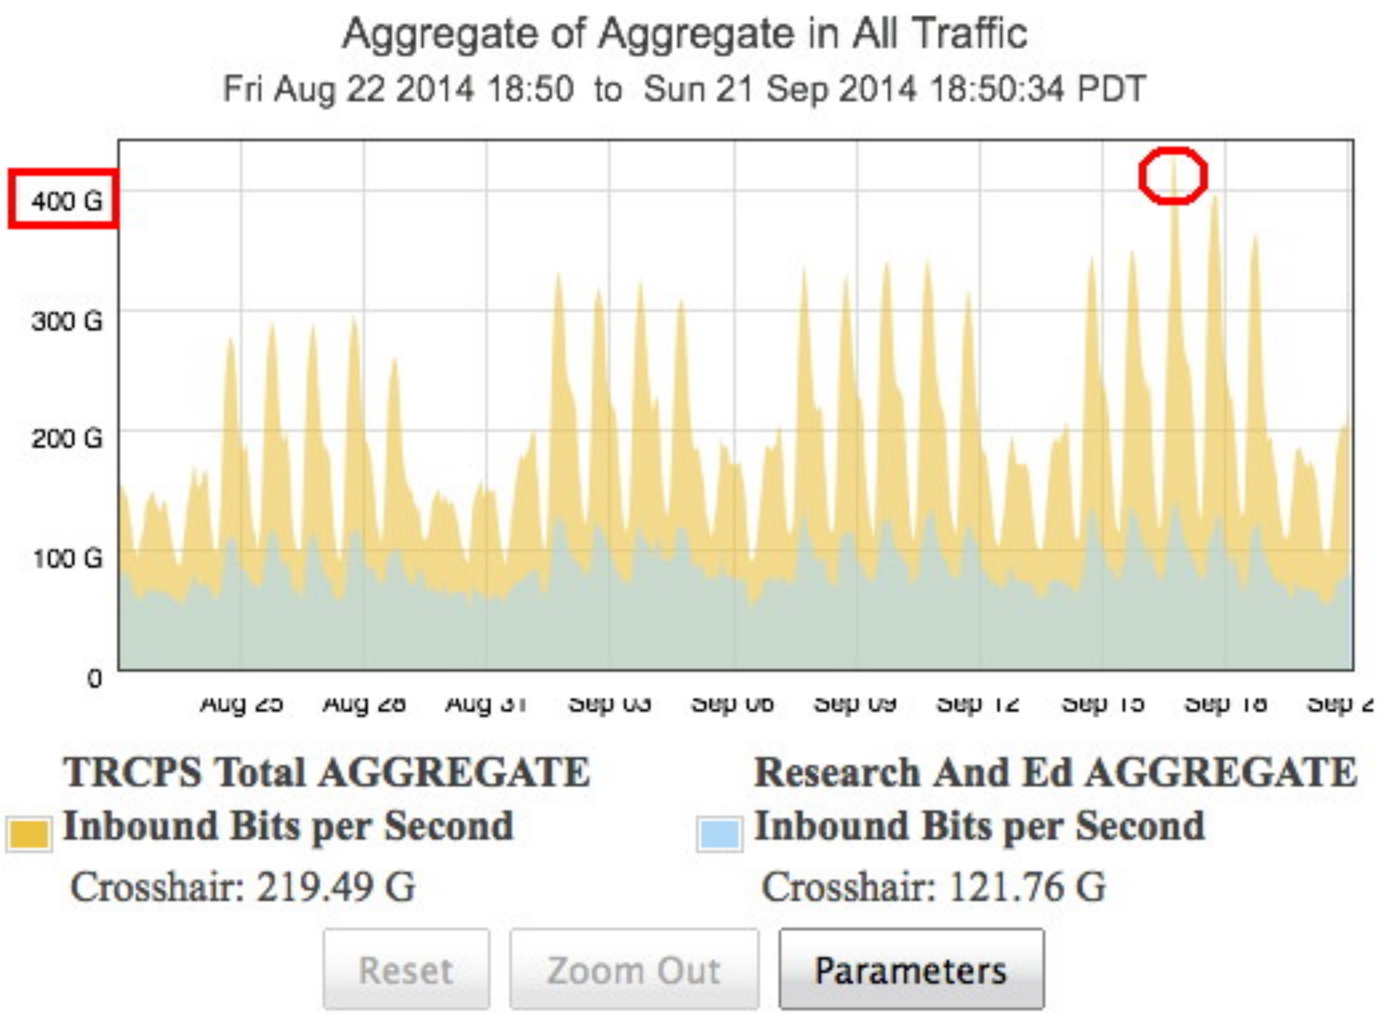

Aggregate of Aggregate in All Traffic
Fri Aug 22 2014 18:50 to Sun 21 Sep 2014 18:50:34 PDT
400 G
300 G
200 G
100 G
0
TRCPS Total AGGREGATE
Inbound Bits per Second
Crosshair: 219.49 G
Research And Ed AGGREGATE
Inbound Bits per Second
Crosshair: 121.76 G
Reset
Zoom Out
Parameters

Are We Changing What We Do?

- Should our goal be to meet the community's networking needs, whatever those needs may be?

- **If most of the community's load turns out to now be commodity peering traffic, is that our new reason for being? Are we okay with that?**

- And if so, is our pricing right? What does 100Gbps peering currently cost, anyhow?

AMS-IX As A Benchmark

- If you happen to be a network engineer in Amsterdam, and you want 100Gbps peering,AMS-IX sells a 100GE peering port for EUR 5,000/month (~$6,600/month) plus VAT (see https://ams-ix.net/services-pricing/pricing). Doing the math:
  -- $6,660/month*12 months=$79,200/year
  -- Going the other direction, $6,600/100,000Mbps= $0.066 per Mbps/month. **That's pretty cheap.**

- For comparison, from Internet2:
  -- A 100GE Advanced Layer 3 port costs a higher education member $200,000/year, or $16,666.67/month (>2.5X AMS-IX)
  -- A 100GE Advanced Layer 2 port from Internet2 costs a higher ed member $165,000/year, or $13,750/month (>2X AMS-IX)

- Is there anything cheaper still?

Multi-TB Drives Shipped Via FedEx?

- If you're just transferring bulk data, and you don't need interactivity, "Never underestimate the bandwidth of a station wagon full of tapes hurtling down the highway," as Andrew S. Tanenbaum famously said – or now, the bandwidth of 6TB drives via an overnight shipper such as FedEx.
- Any average researcher can now buy a brand new 6 terabyte USB external drive for $300 and promptly load it at up to 5Gbps via a USB 3.0 port, no special network "hoop jumping" skill required.
- If they need to share that drive with a colleague on the other coast, they can get it there via FedEx Overnight for less than a hundred bucks, or if they're not pushed for time, they can get it there in two days for half that price. That's pretty cost effective.
- I'm not sure I want to engage in that economic (or throughput!) "race."

So How's Internet2 Different From the Regular Commodity Internet, Again?

- At one point, our key differentiator was speed. We were faster than what you could get from the commodity Internet. But as you can see from the AMS-IX pricing, they're doing 100GE these days, too.

- Being fast, in and by itself, is just no longer enough to set us apart from the "regular Internet."

- That's okay, we also strived to deliver advanced protocols.

IP Multicast, IPv6, QoS...

- Historically the community has also put heavy emphasis on our **support for advanced protocols**, such as IP multicast, IPv6, or QoS.
- We now know that the world has looked at **IP multicast** and shrugged its shoulders: IP multicast is dead, Jim. It was a technically elegant protocol, and one that many of us loved, but Netflix, Hulu and Youtube have won. It's time to stop flogging it.
- **IPv6** is doing somewhat better, but uptake in commercial ISPs (such as Comcast) now challenges uptake in higher education. If a person's grandparents can get IPv6 from their cable Internet provider, does that really still qualify as something "exotic" or "cool" that we can still brag about? I don't think so.
- And when it comes to QoS...

QoS, Another Fallen Advanced Network Service

- *"Between May 1998 and October 2001, Internet2 worked to specify and deploy the QBone Premium Service (QPS) [QBone], an interdomain virtual leased-line IP service built on diff-serv [RFC2475] forwarding primitives and hereafter referred to simply as "Premium". Despite considerable effort and success with proof-of-concept demonstrations, this effort yielded no operational deployments and has been suspended indefinitely. [...]*
- *"The costs of Premium are too high relative to the perceived benefits. Moreover, even if successfully deployed, Premium fundamentally changes the Internet architecture, running contrary to the end-to-end design principle, and threatening the future scalability and flexibility of the Internet."*

"Why Premium IP Service Has Not Deployed (and Probably Never Will)," https://web.archive.org/web/20081203141945/http://qos.internet2.edu/wg/documents-informational/20020503-premium-problems-non-architectural.txt

VIII. So What SHOULD We Be Working On These Days, Eh?

The Community Really Needs To Do Some Planning

- If you're going to get where you need to go when it comes to advanced academic networking, **you need a plan.**

- Internet2 went through a strategic planning exercise once before, in 2008, see https://wiki.internet2.edu/confluence/display/I2SP/2008+INTERNET2+STRATEGIC+PLANNING , but that plan's now woefully out of date.

- Without a current plan, how will we all be able to get where we need to go?

- Answer: we won't. We'll end up lost, perhaps with priorities that have subtly shifted to areas that are important to some, but which aren't consistent with a primary focus on RESEARCH.

We Need To Re-Engage With Academic Researchers, And With Campus <u>Technical</u> Networking Leaders

- Of late, there has been much effort to better connect with university Chief Information Officers (CIOs), and that's great. Our CIOs lead production IT in our schools, and they write the checks that we keep our network infrastructure operating. They need to be "on-board." However, CIOs tend to worry primarily about "CIO stuff."
- As a result, there's a risk of misalignment: what CIOs want or thinks their school needs may not be what faculty researchers want or need, or what technical networking leaders see as critical steps in preparation for meeting the community's network needs.
- We need renewed national dialog with our research faculty members (and senior academic staff, such as our Provosts and VPs for Research), AND with our technical networking staff.

Let's Never Forget Why We're Here: Research!

- At one time, research computing needs dominated campus network planning. Hosts connected directly to the network, rather than living behind firewalls (or other throughput-limiting and protocol-constraining middleboxes). ***Sensitive administrative hosts were the exception, and they were handled by putting *them* in a specially protected network enclave.***
- **These days, things have gotten "all turned around." These days, administrative computing security requirements dictate the default network posture for everyone,** including most campus researchers, and that means that most people are behind a firewall by default. It is only the rare uber-high performance research host that is allowed to be "daring," and live outside the campus firewall in a specially plumbed "Science DMZ." This is crazy – we're letting the administrative <u>**tail**</u> wag the research <u>**dog**</u>. **Research needs should drive university network architectures.**

The Rising Tide That Lifts All Boats vs. Specialized Solutions For "Unusual" Needs

- Treating high performance networking as a "special case" scenario has many implications, but one of the most important ones is that the effort (and money!) that a site puts into delivering high performance network doesn't help everyone on campus -- it only helps the few researchers who may be lucky enough to have hosts living in the campus "Science DMZ."
- While I know that every campus has a few research stars who have the budgets (and clout!) to get whatever they want, I guess I'm enough of a populist that I'd like to see network investments help all members of the university community, if only because I have a hard time guessing in advance who will turn out to be the NEXT research star, and building capacity accordingly.
- Fix the general case, don't micro-optimize just the "exceptions."

Network Confidentiality with IPsec

- Recent revelations about pervasive network monitoring by Edward Snowden have made it clear that we can no longer assume that international (or even domestic) network traffic will be immune from interception (if we ever could)
- Moreover, trust in even application layer crypto has been severely undercut by discovery of implementation flaw-after-implementation flaw in widely deployed cryptographic libraries. (wasn't "Heartbleed" just a ton of fun?)
- We need a layered cryptographic approach for redundancy, and to regain practical network confidentiality.
- Specifically, we've never seen substantial deployment of IPsec (see relevant RFCs at http://en.wikipedia.org/wiki/IPsec). **The time has come for that to change. It is time for the community to work on getting IPsec broadly deployed.**

Success in Deploying Opportunistic Encryption

- As we think about deploying IPsec as a community, I want to highlight the fact that we've actually seen commercial success when it comes to deploying another type of encryption, namely opportunistic encryption of email traffic from mail server to mail server, at least here in the US.
  See https://www.google.com/transparencyreport/saferemail/

- **Is your campus mail server protecting MTA-to-MTA flows?**

The Challenge of Traffic Analysis

- Even if application traffic is fully encrypted, simply knowing that two particular parties are communicating can tell a analyst a lot. For example, if a person living under a repressive regime is seen sending an encrypted email to a human rights organization, even if there's no way to tell *what's* being said in that email, the mere fact that there's any communication may still send up a red flag.
- Similarly, knowing WHEN someone's communicating, or how extensively they're communication, or how a sequence of communications occurred can all tell a trained analyst a lot, even if the body of the communications is totally encrypted.
- This is another area where we need to work as a community -- we need to ensure that we have trustworthy solutions that will effectively resist traffic analytic approaches. This is a potentially a very difficult goal to try to satisfy.

Surprisingly, Paul Baran Anticipated This, Too...

- "The proposed network is a <u>universal high-secrecy system</u>, made up of a hierarchy of <u>less-secure sub-systems</u>. It is proposed that the network intentionally treat all inputs as if they are classified, in order to raise the intercept price to the enemy to a value so high that interception would not be worthy his effort. Of course, that extra layer of <u>conventional</u> cryptography would be maintained for use in those extremely sensitive cases where the proposed approach might seem risky.

  "Thus, fullest advantage is taken of the mechanism within the proposed system that takes a channel or a message and chops it into small pieces (like a fruit salad), transmitting it on a series of message blocks, each using a different path. Additionally, much unclassified material is purposely transmitted cryptographically, and perhaps even a light dose of obsolete traffic is mixed in. Given a big enough bowl, it becomes very difficult to separate the garbage from the salad." [emphasis in original]

  "On Distributed Communications: IX. Security, Secrecy and Tamper-Free Considerations," http://web.archive.org/web/20111004164807/http://www.rand.org/content/dam/rand/pubs/research_memoranda/2006/RM3765.pdf

Securing Wide Area Routing

- Another area that has seen woefully little progress to-date has been the area of securing BGP.

- This is an area that the FCC Communications Security, Interoperability and Reliability Council has recently been concerned about, and one that we should also be concerned about -- and working to address as a community -- too.

- You can read a one page outline of the issue here:

  http://pages.uoregon.edu/joe/csric4-wg6/bgp-routing-security-outline.pdf

The 10,000 Mile Problem (International Networking)

- If you're fortunate enough to be a researcher in the United States or Canada, or the European Union, or one of the advanced countries in the Far East such as Japan or South Korea, you've typically got **pretty good connectivity.** Give thanks for your good fortune!

- When it comes to much of the rest world, connectivity is still all-too-often expensive and limited.

- **Should we (those of us who are fortunate to have so much), voluntarily help our less fortunate colleagues with their hugely expensive transoceanic connections? The United States and the EU have already done some of that sort of thing, but pragmatically, have we done enough?**

Why Should Developed Nations (Like the US) Pay To Improve Connectivity to Developing Nations?

- Millions of new Internet users are coming online from developing nations every month. Many of those users are using insecure systems that aren't patched up-to-date, and as a result we're seeing spam and other unwanted traffic from them.
- *Why* aren't those systems secure?
  - Users in some of these regions may be using older systems that run old and no-longer-supported software, or they may be using pirated software.
  - Most times, however, users may have poor connectivity, so it takes too long to download updates over the network.
- The FCC worries about trivial levels of infected hosts in U.S. service provider networks, but that's crazy when there are countries where botnets are endemic – and those systems are targeting the United States. Focus on the low hanging/worst problems first!

% of Infected Internet Users, Selected Countries

#1	Cypress	15.73%
#2	Dominica	6.36%
#3	Gabon	6.08%
#11	Iran	2.23%
#12	Vietnam	2.11%
#40	Russia	0.87%
#63	India	0.58%
#98	Brazil	0.38%
#132	China	0.20%
#138	**US**	**0.16%**
#141	UK	0.16%
#146	Japan	0.15%
#154	Canada	0.12%
#171	Netherlands	0.08%

cbl.abuseat.org/countrypercapita.html

So where's the REAL problem with insecure (botted) systems? Is it the US? Or is it elsewhere?

Should we obsess about fixing tiny residual issues at home, or should we get our priorities straight and focus on helping countries overseas that are totally overrun?

Finally, The "Five-to-Ten Millimeter" Problem

- Everyone tells me "Joe! The future is all about mobile devices."
- Today's smart phones -- things like the new iPhone 6, or the Galaxy S5 -- sure are terrific, aren't they? Everybody probably loves them -- except, I guess, me. I have **multiple problems with mobile devices**, but let me just talk about one of them, what we might call the **"five-to- ten millimeter problem."**
- That's the typical thickness of the human skull bone, the obvious physical barrier between mobile devices and the brain. As a result:
  - I still get output from smart phones via my worn-out eyes and ears, and
  - I still (try) to enter text by typing with my too-large fingers or by using my voice ("Siri, why are smart phones still so damn primitive?")
- I don't want a 1920's interface solution for my 2014 devices!

Connecting to Your Personal Communications/ Entertainment System in The 1920s... and Today

Right image source: http://upload.wikimedia.org/wikipedia/commons/e/e3/Shane_Morris_September_2013_%28cropped%29.jpg

A Medical Example of Man-Machine Coupling

- 360 million people worldwide have disabling hearing loss (that amounts to over 5% of the world's population)
- Notwithstanding the substantial cost (averaging $60,000 per ear), at least 300,000 people worldwide (0.083% of those who might be helped) now have a cochlear implant.
- Even production of simple hearing aids is just at less than 10% of the level needed to meet worldwide demand. [Thus many who are deaf or hard of hearing rely on sign language, send text msgs, etc.]

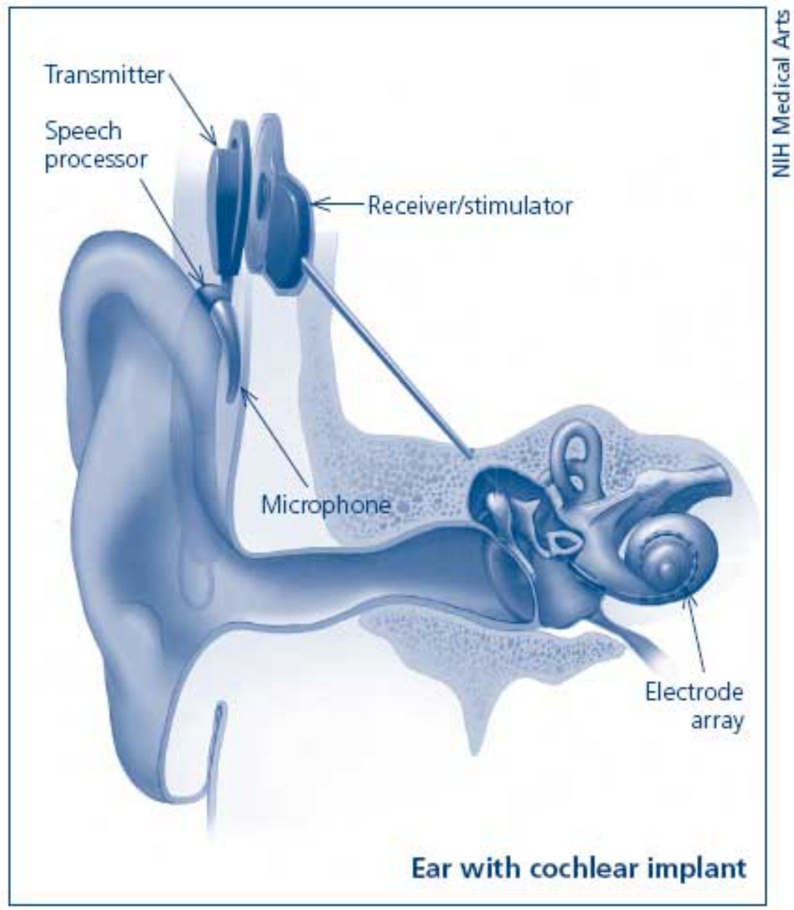

Thinking About Man-Machine Interconnections

- First of all, let me be clear that I'm **not** suggesting that the hearing get cochlearized! THAT's **not** my point. I merely raise that as one example of a current effort to directly couple users & processors.
- Many issues remain, including the fact that we don't know how to do link systems and brains for all relevant senses on a routine basis. The insides of our own heads? Still largely *terra incognita.*
- Costs are still prohibitive, and will likely continue to be so.
- There are potentially profound security issues (I don't know if I want to have a direct connection into my head to get hacked).
- Typical mobile device life cycles run around two years. Unless upgrades can be constrained purely to external components, I don't think you want invasive surgery on a two year schedule. :-;
- And yet, shouldn't we be pushing the frontier somewhere? If not here, where?

One Last Historical Note...

- Remember Licklider, mentioned on slide 16 of this talk? I like to think that he'd particularly like to see this topic finally getting a little of the attention it deserves, too... He wrote:

- "At present, however, there are no man-computer symbioses. [...] The hope is that, in not too many years, human brains and computing machines will be coupled together very tightly, and that the resulting partnership will think as no human brain has ever thought and process data in a way not approached by the information-handling machines we know today."
"Man-Computer Symbiosis," J. C. R. Licklider, March 1960
http://groups.csail.mit.edu/medg/people/psz/Licklider.html

- Let's work on making his vision -- or your vision -- come true.

Thanks For The Chance To Talk Today

- Are there any questions?