THE UNIVERSITY OF AUCKLAND
NEW ZEALAND
Te Whare Wānanga o Tāmaki Makaurau

The Internet
Where did it come from?
Why did it succeed?

Dr Brian E Carpenter
Department of Computer Science
The University of Auckland
December 2008

Warning: scope of this talk

- I will be talking about the Internet's basic communications infrastructure.
- I won't talk about the fluff on top of the Internet.
  - The Web
  - File sharing
  - Spam
  - Skype
  - Social networks
  - IPTV
  - ....

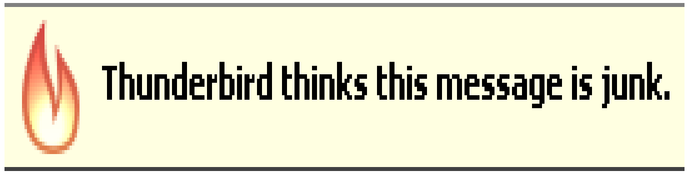

Protocol stack

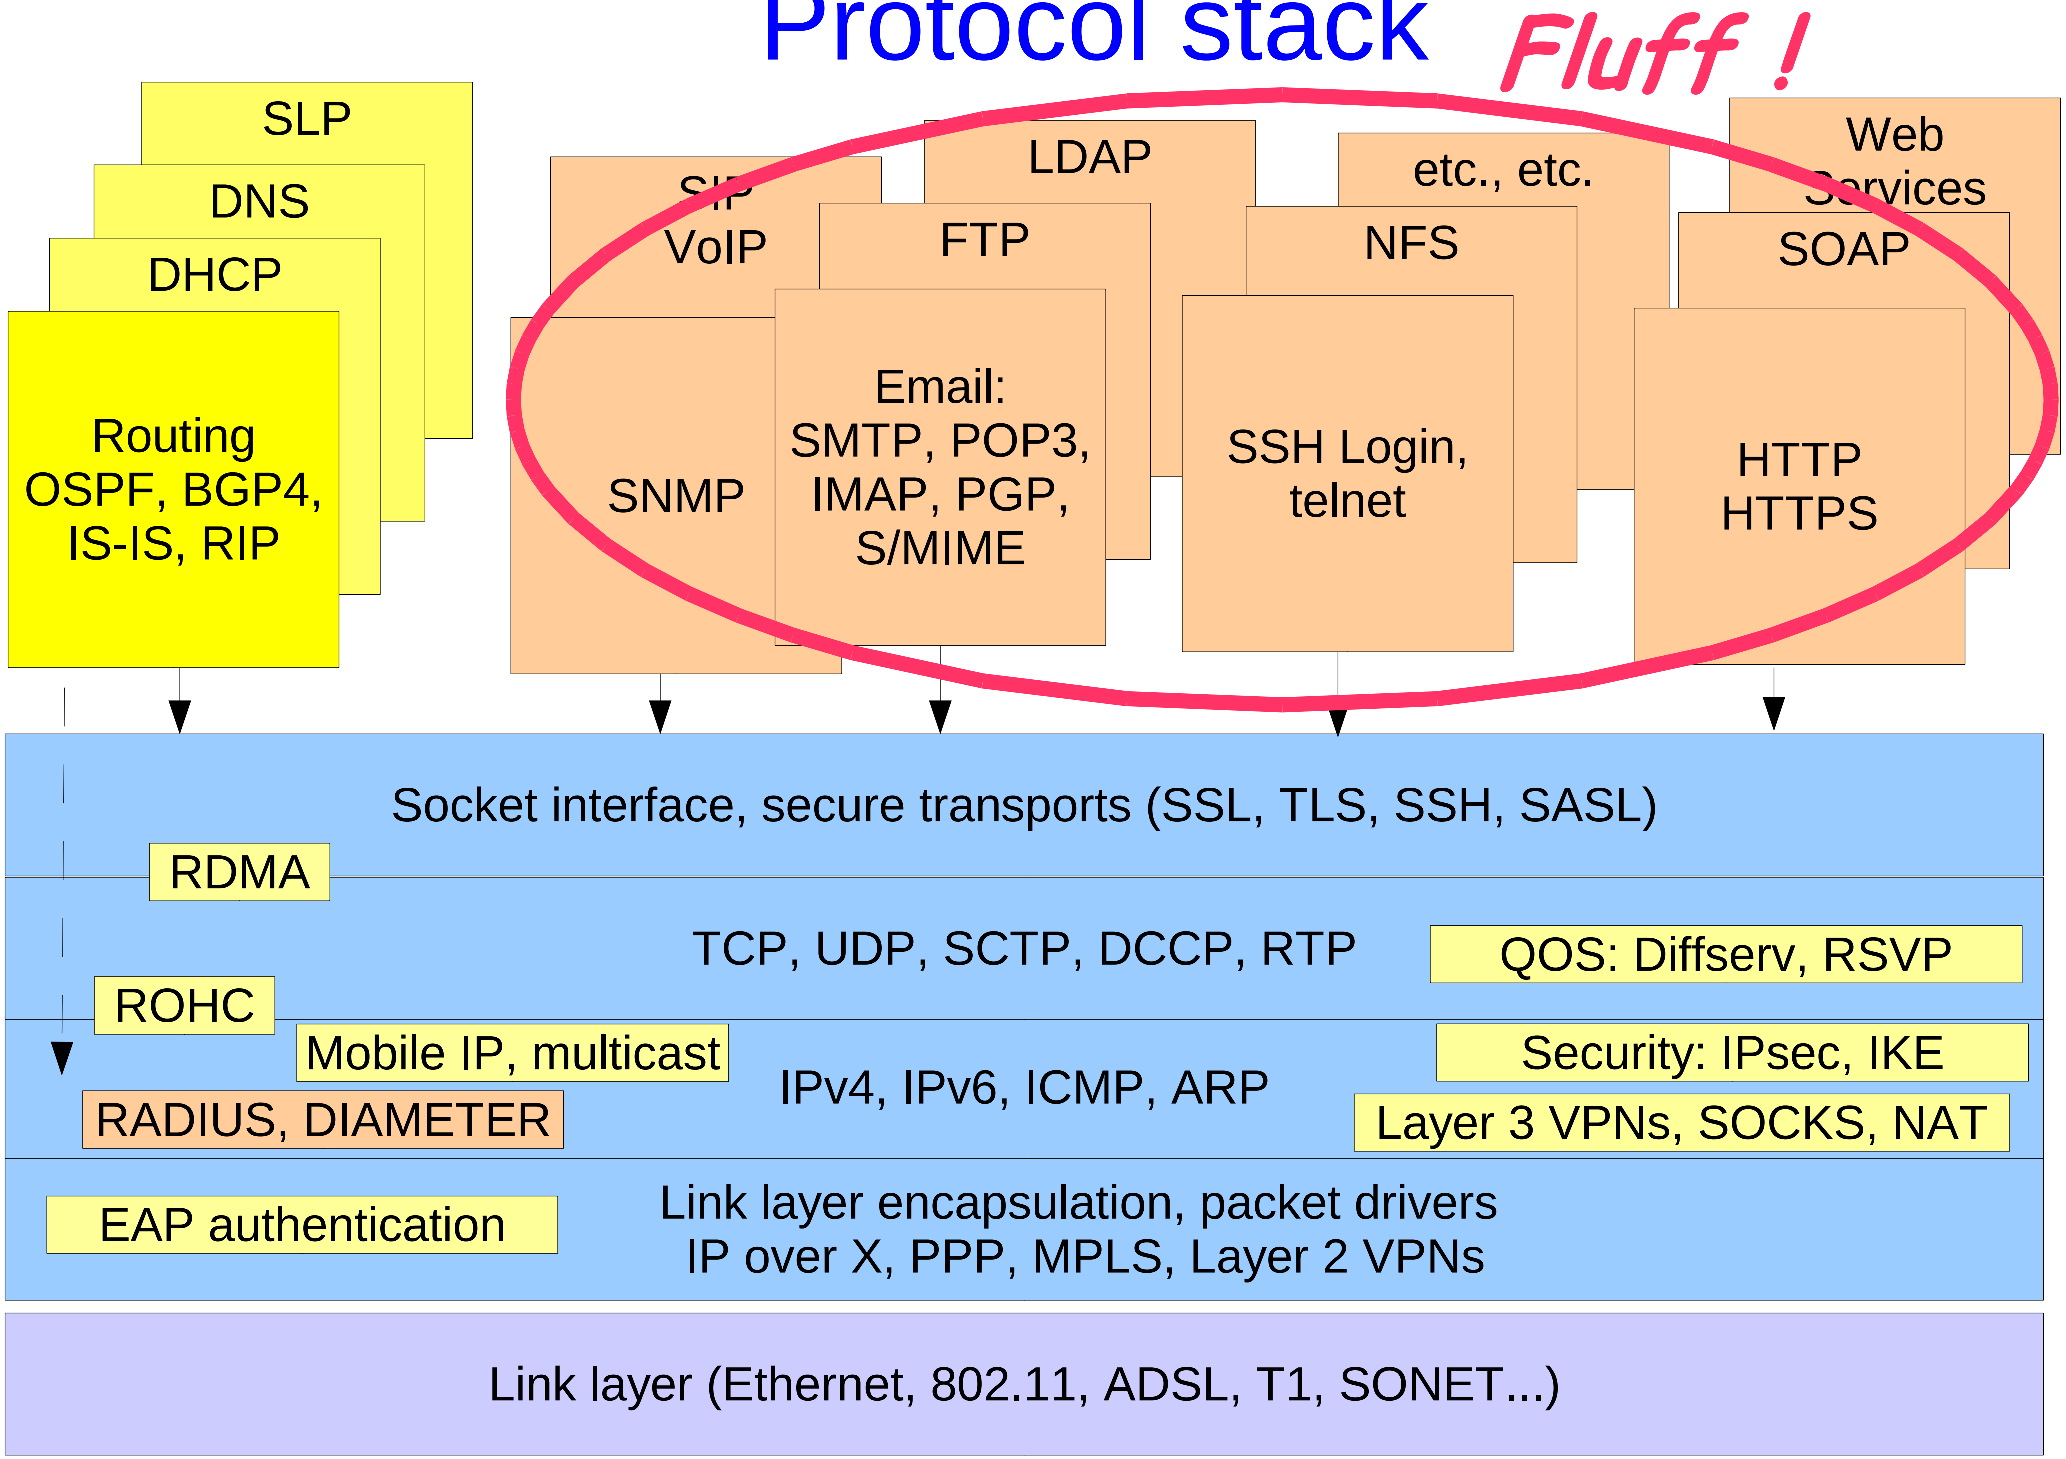

How many computers on the Internet?

(measured by published DNS names)

Data from http://www.isc.org/ and other sources

How many computers on the Internet?

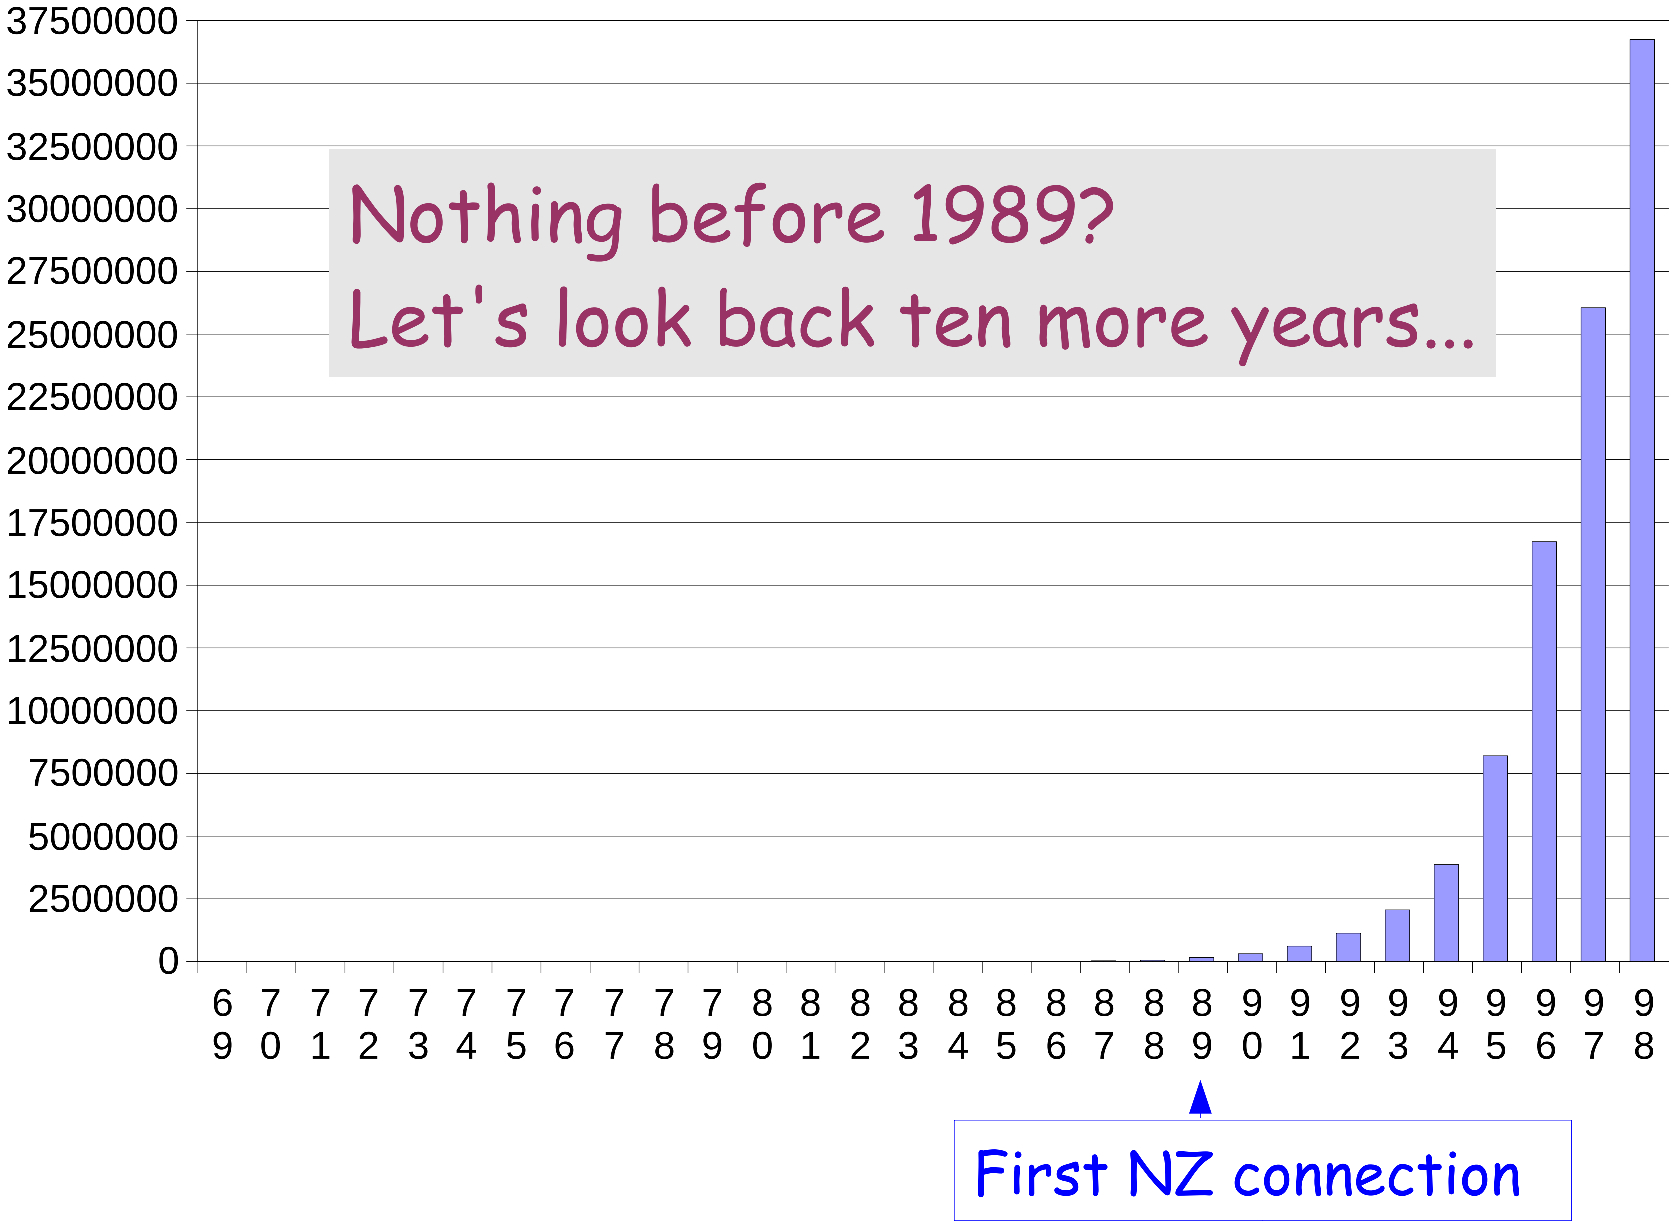

How many computers on the Internet?

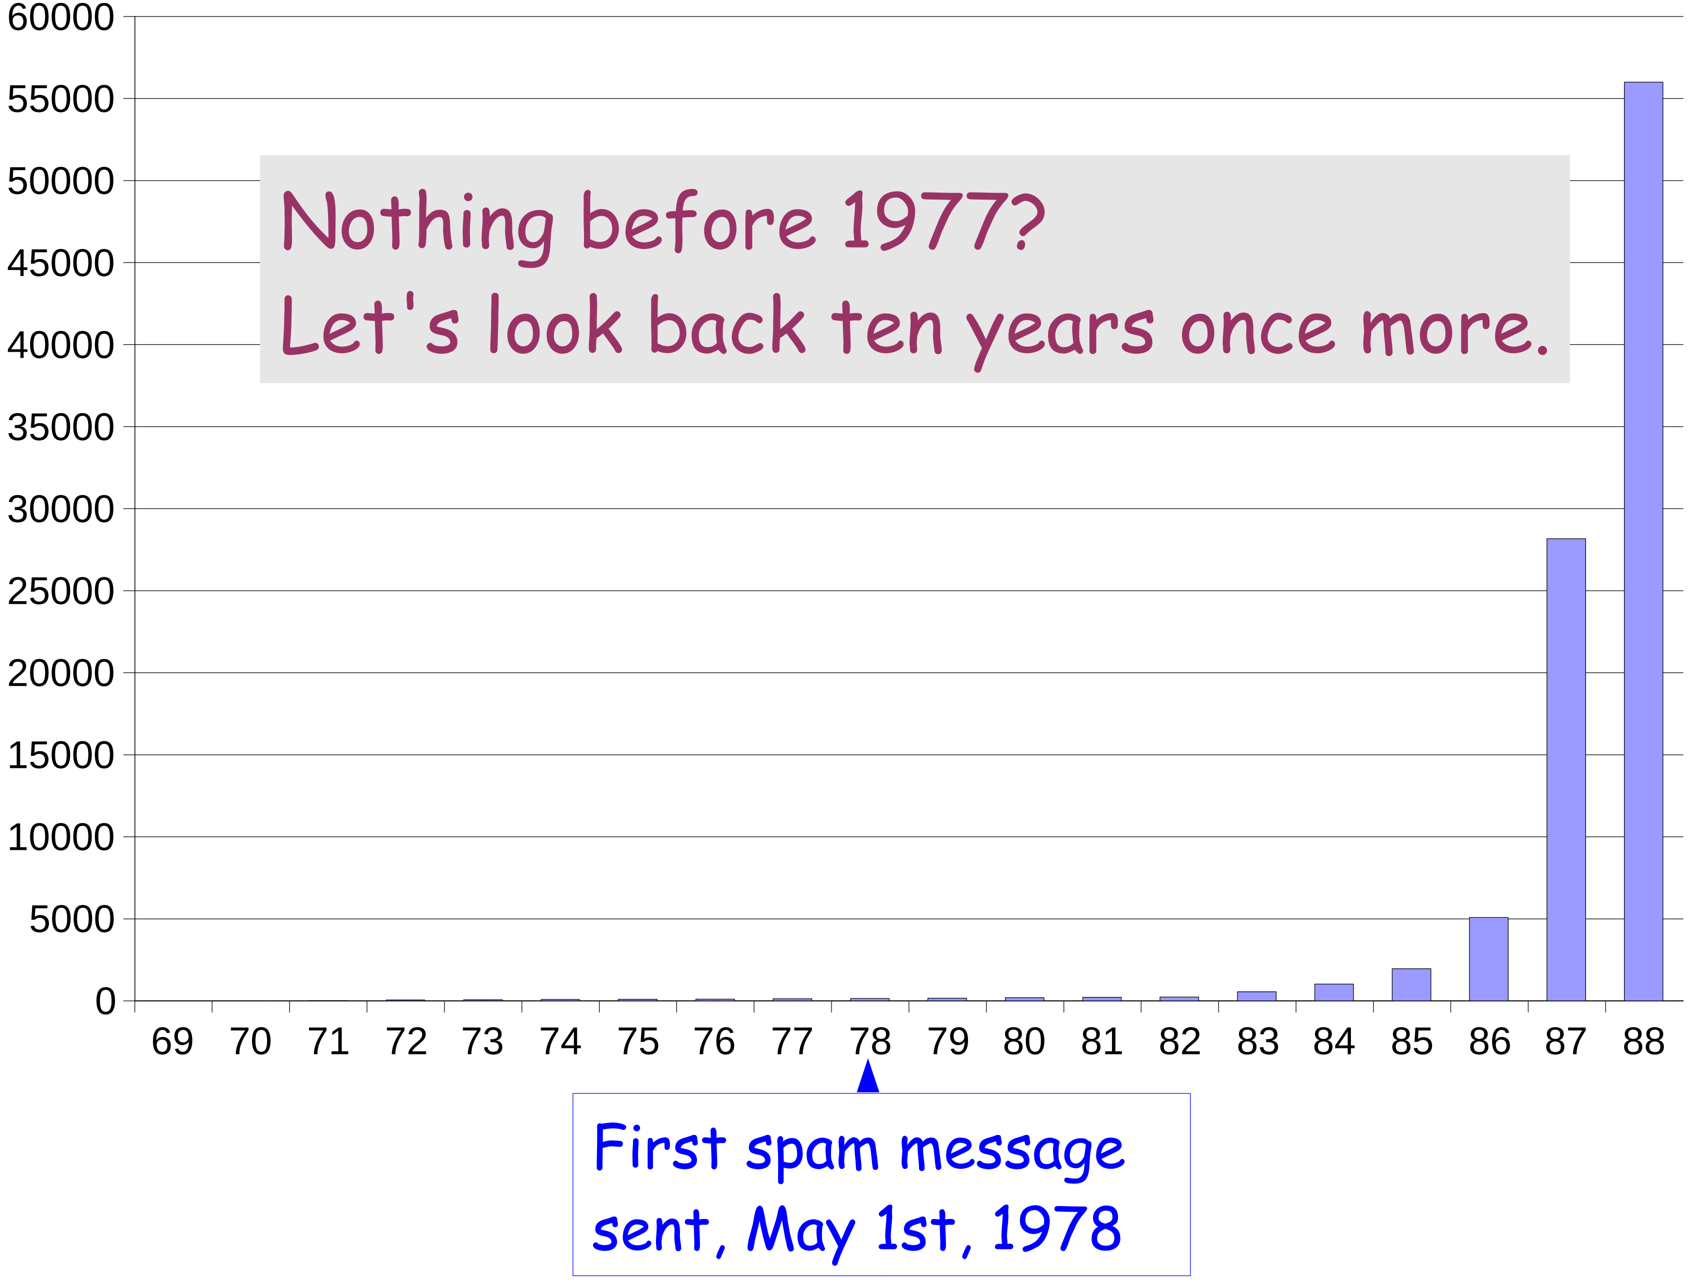

How many computers on the ARPAnet?

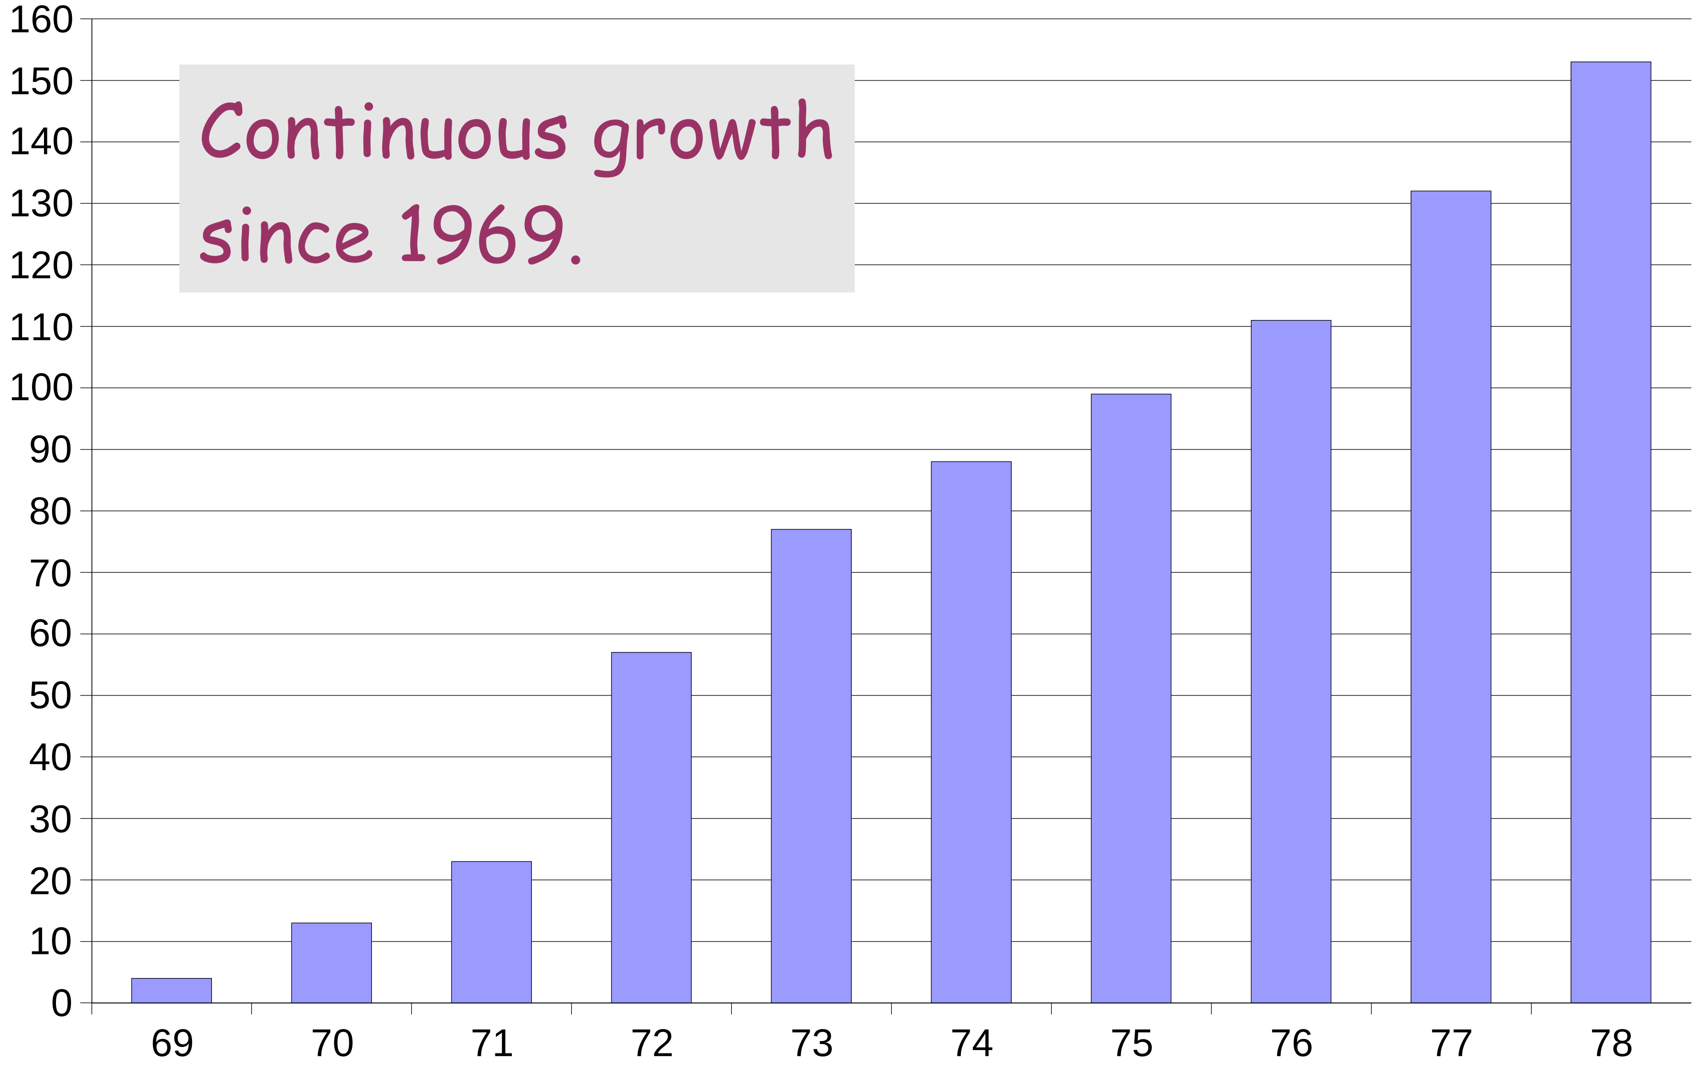

Kiwinet project (Vic/Massey)

Who knew about the Internet in 1973?

- *Nobody. The concept was invented in 1974.*

➔ ARPANET hosts in 1973: ~*75*

Who knew about the Internet in 1983?

- Vice-President George H.W. Bush? *Unlikely.*
- Bill Gates? *Possibly, as a vague memory from his time at Harvard.*
- University Presidents in the US? *Possibly, as a minor funding annoyance.*
- University Vice-Chancellors in NZ? Unlikely.

➔ Internet hosts in August 1983: *562*

Who knew about the Internet in 1993?

- Vice-President Al Gore? *Yes.*
- Bill Gates? *Yes, but there was no particular reason to support it in Windows 3.1.*
- The Editor of the Economist? *Possibly, but it had no particular business value.*

➔ Internet hosts in late 1993: *2,056,000*

Suddenly, rumours started that the Internet is important

- "Some of the top tier providers of the Internet have become very interested in business models. ... People are asking: What is the purpose of the Internet? What is the business model? ... If the predominate *[sic]* business model becomes one of support for mission critical Fortune 1000 business activities..."
  - *Cook Report on the Internet, September 1995*

Of course, not everybody "got it"

- *November 19, 1997:*
- An elderly couple from Portsmouth on the South Coast of England read ... that the BBC's Monitoring Unit at Caversham had a Web Site.
- So the couple got into their car and drove the 45 miles to Caversham.
- "Hello we have come to see the Website, we read ... where it says visit the BBC Website and I'm wondering if this is the right place and if you are open for a visit".

http://www.langston.com/Fun_People/1997/1997BIW.html

And of course, it didn't last...

SmartMoney

Search More Options

Quote Site Search

Ticker or Company N

advanced search

Select

Portfolio

Tools & Maps

Stocks & Options

Funds

ETFs

Personal Finance

Economy & Bonds

smSmallBiz

Login Special Offers Free Regis

Magazine Portfolios: Stock Portfolios: The One Telecom Stock to Buy Now

STOCK PORTFOLIOS

The One Telecom Stock to Buy Now

By James B. Stewart and Odette Galli | Published: October 2000

BOOKMARK RSS Send Comments

Add this column to your

Print Send

By James B. Stewart and Odette Galli
October 2000

IT SEEMS LIKE ONLY YESTERDAY that some of the hottest growth
new economy were the telecommunications carriers. With Internet tra
usage soaring, and with visions of one-stop shopping for voice, wirele
e investors' eyes, stocks in the sector were so

edicted that the summer of 2000 would prove a
ghtmare? **Verizon Communications** (VZ) wor
nalgam of Bell Atlantic and GTE reported weak
s plunged 12%. **WorldCom** (WCOM) had rock
of lower profits and revenue, and its stock also
elekom (DT) announced the $50 billion acquisi
s (VSTR), and its shares dropped 13%. **AT&T**
way. Slumping revenue from its business-servi
ninkable $30. On average, the shares of the se
own 28% year-to-date, lagging far behind the S

... except that you can't even see the bubble bursting on the growth curve.

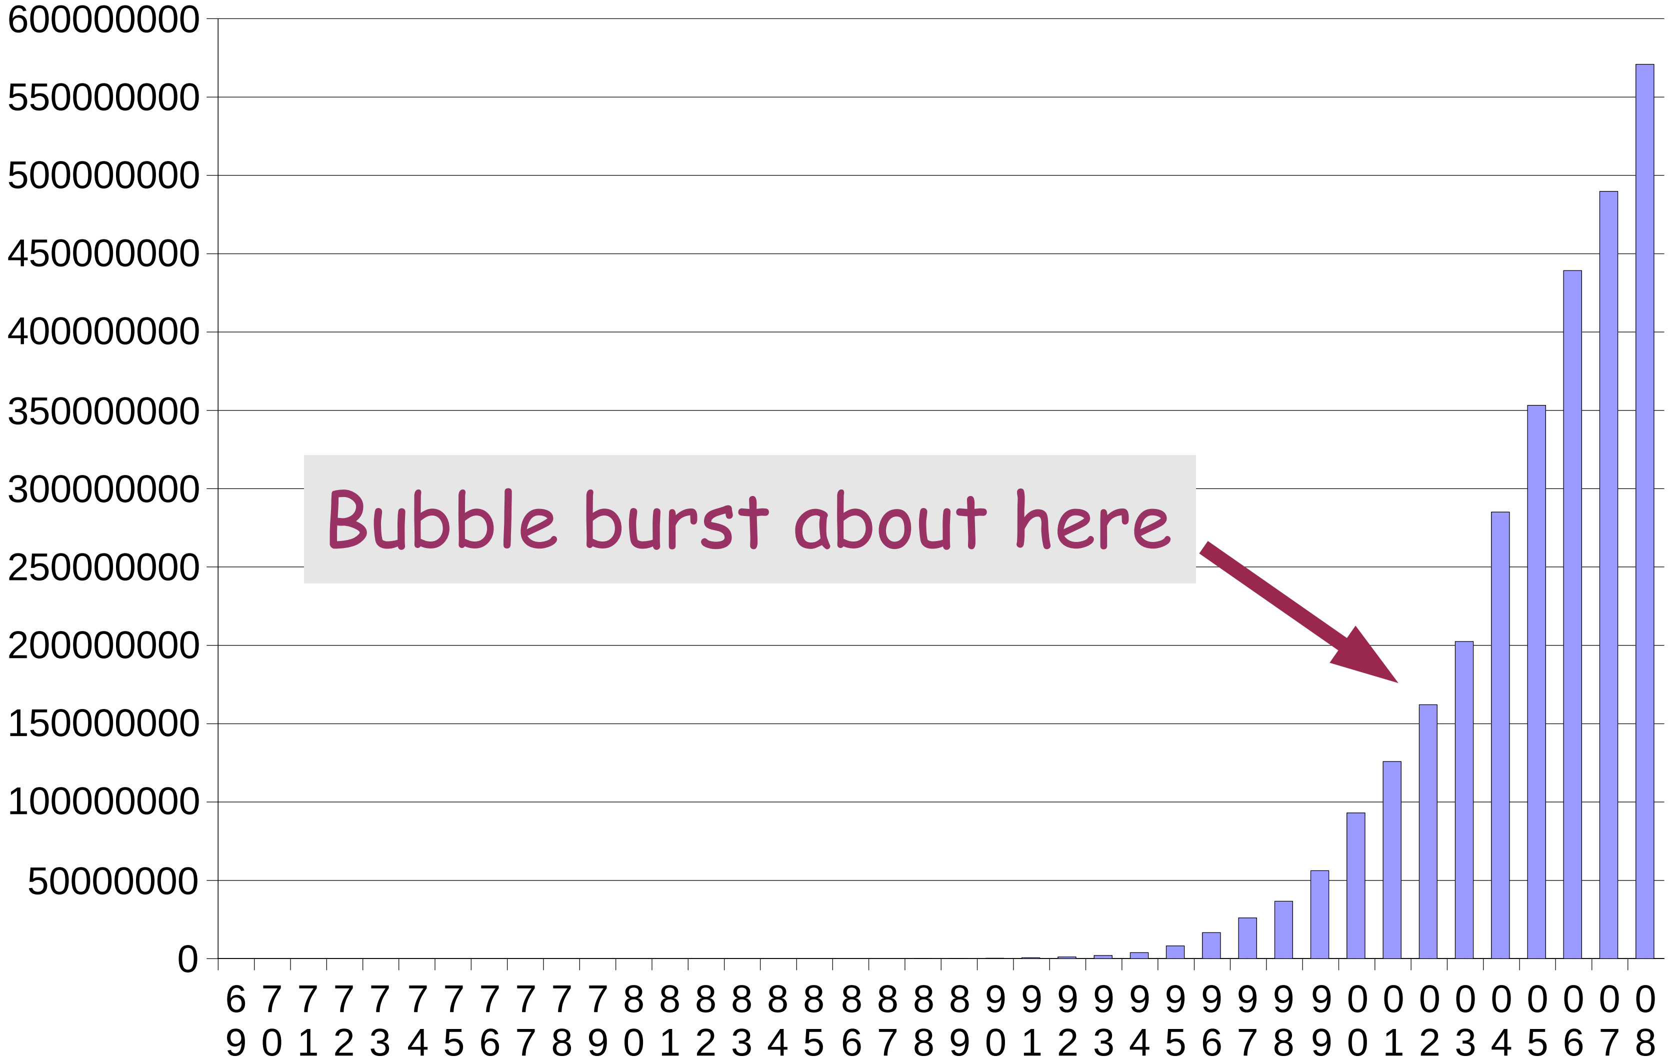

Data from http://www.isc.org/ and other sources

Subtle effects of the economy

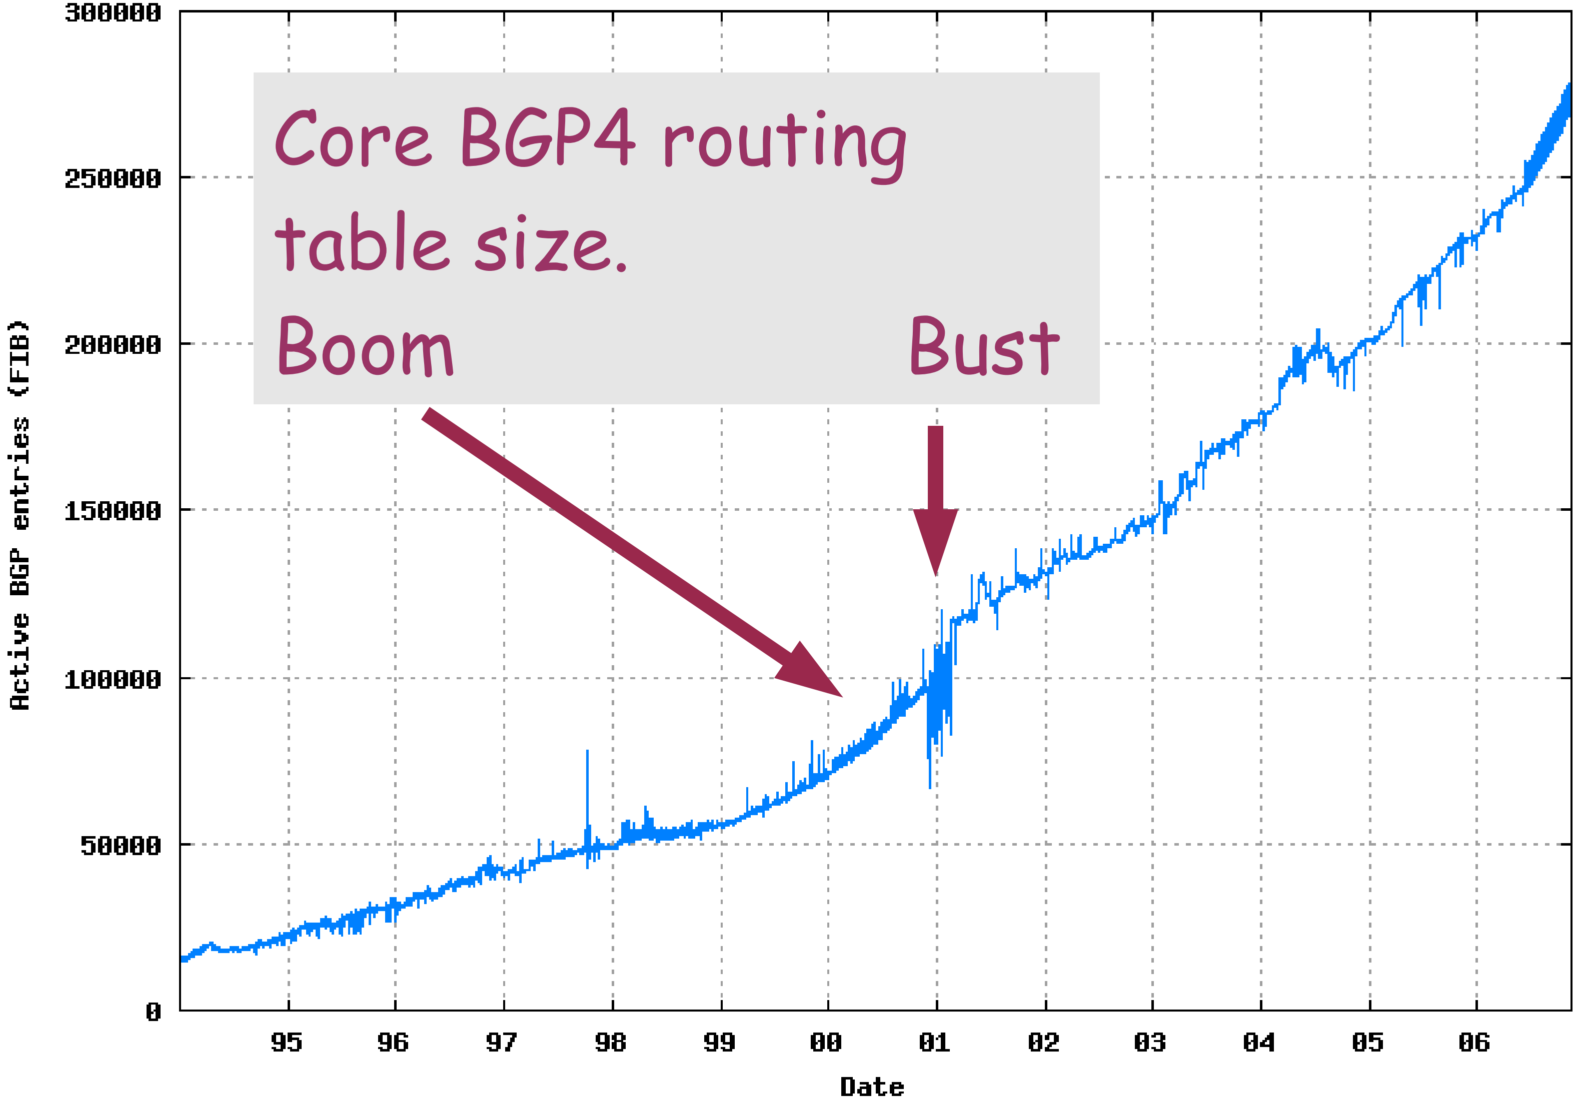

How many computers on the Internet?

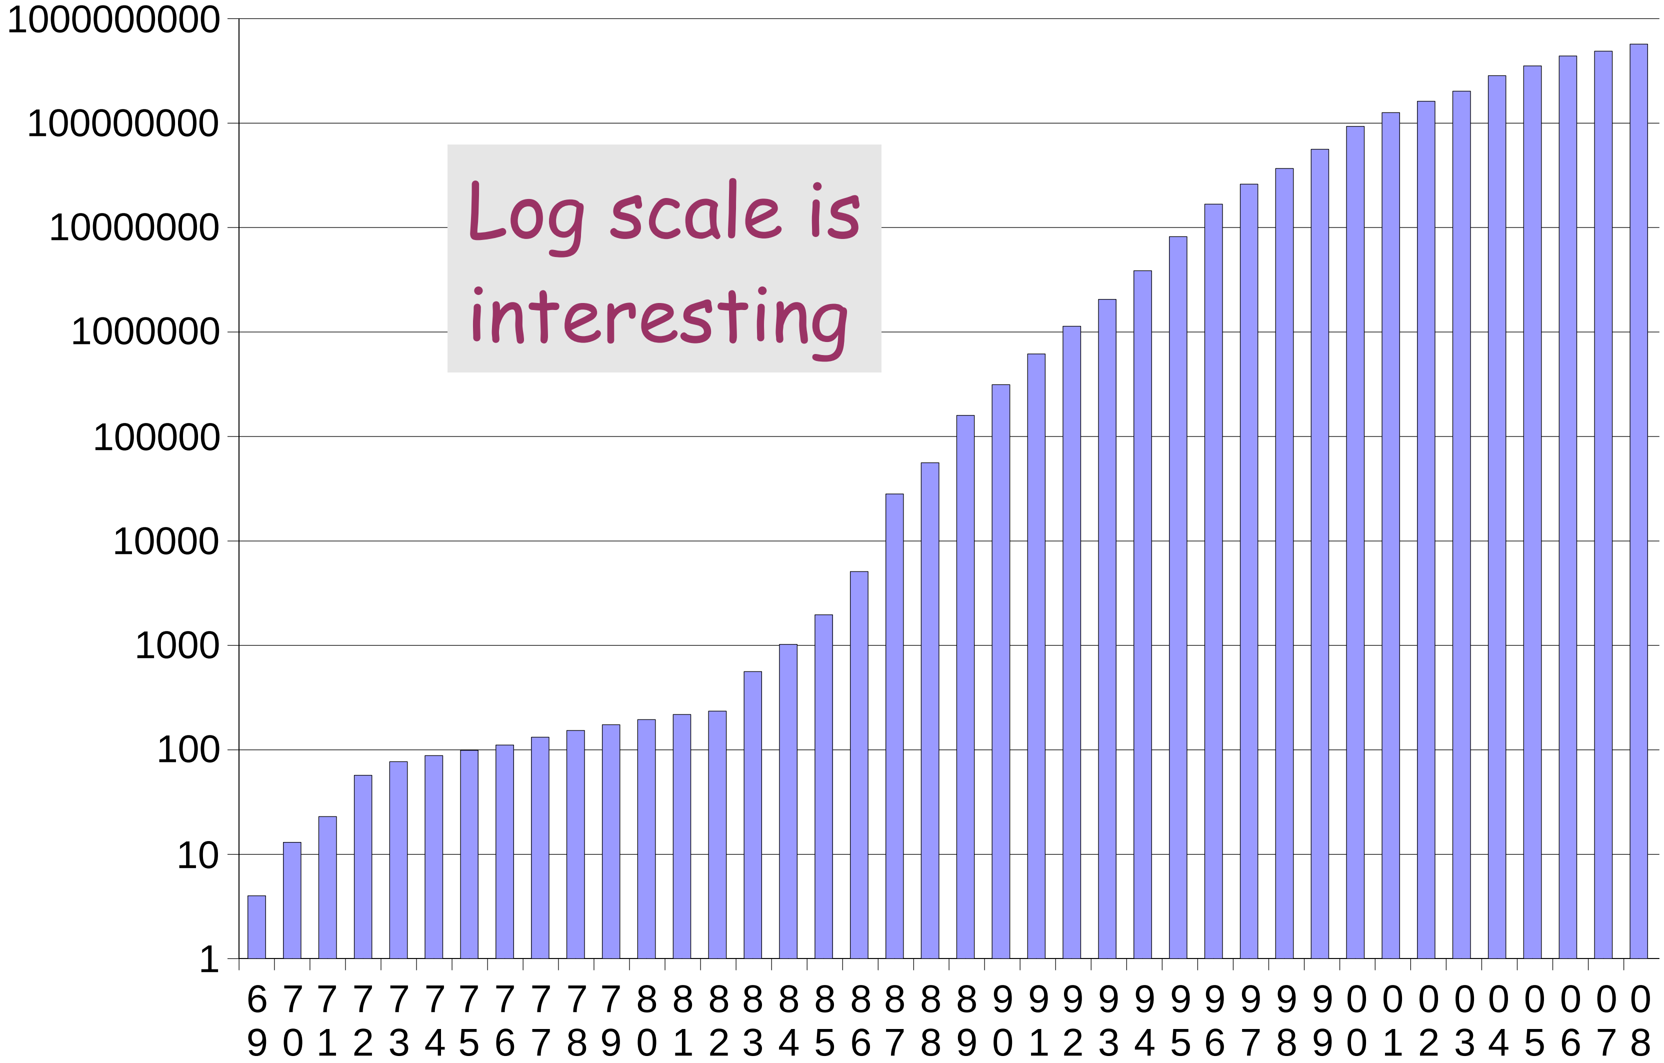

Data from http://www.isc.org/ and other sources

Who knew about the Internet in 2003?

- *Do I really need to answer that?*

➔ Internet hosts in late 2003: *180,000,000*

The underlying question

- Why did the Internet succeed?
  - sustained growth over almost 40 years
  - apparently recession-proof
  - transcends frontiers
  - swept aside "official" international standards
  - no sign of it stopping
- I want to attempt to answer that by reviewing history and highlighting underlying principles of engineering.

Stirrings

- The interconnection of computers was first envisaged in the 1950s, essentially for military or business purposes:
  - SAGE, the network for the early Cold War Distant Early Warning system
  - SABRE, the original airline reservation network
  - (surprisingly, the SWIFT banking network started as late as 1973)
- Mainly based around mainframe computers and pre-defined point-to-point modem links
  - vulnerable, inflexible, hand-crafted and expensive

A wakeup call

- СПУТНИК 1 (1957) led almost directly to the creation of the US *Defense Advanced Research Projects Agency (DARPA or ARPA)* in 1958.
- Budget and authority for basic technical research.

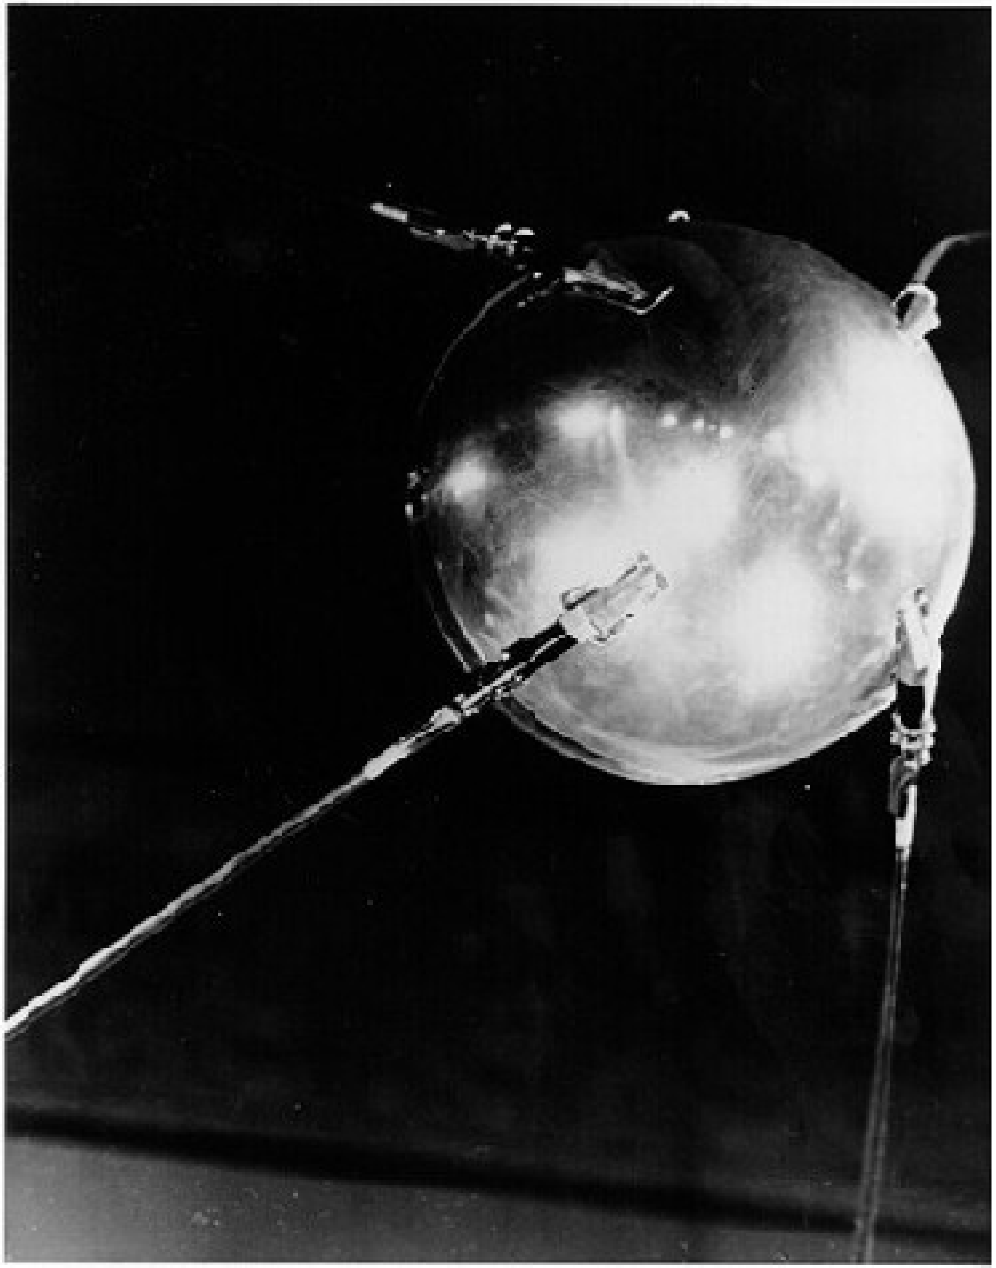

Helvetica font designed in 1957, too.

Another wake-up call

Microwave tower in Utah bombed in May 1961 by "American Republican Army." Widespread communications disruption.

http://long-lines.net

Fundamental concepts of "survivable networks" (1)

"I realized that the reliability of such a network could be far greater than the reliability of the elements that comprise that network.

...redundant units connected in parallel must all fail before the system would fail. This meant that it would be theoretically possible to build super reliable systems out of unreliable parts.

...extremely tough networks could be built at only moderate levels of redundancy."

Fundamental concepts of "survivable networks" (2)

"...The concept of a 'message block.' *(i.e. 'packet'.)*

The concept of adaptive routing.

The concept of decoupling the user's logical address from the physical address.

...the concept of building a network composed of a mixture of totally different types of media and data rates."

- Paul Baran, summarising in 1999 his work done in 1962, the year after the Utah attacks. All applies to the Internet today.

Baran's 1962 topologies

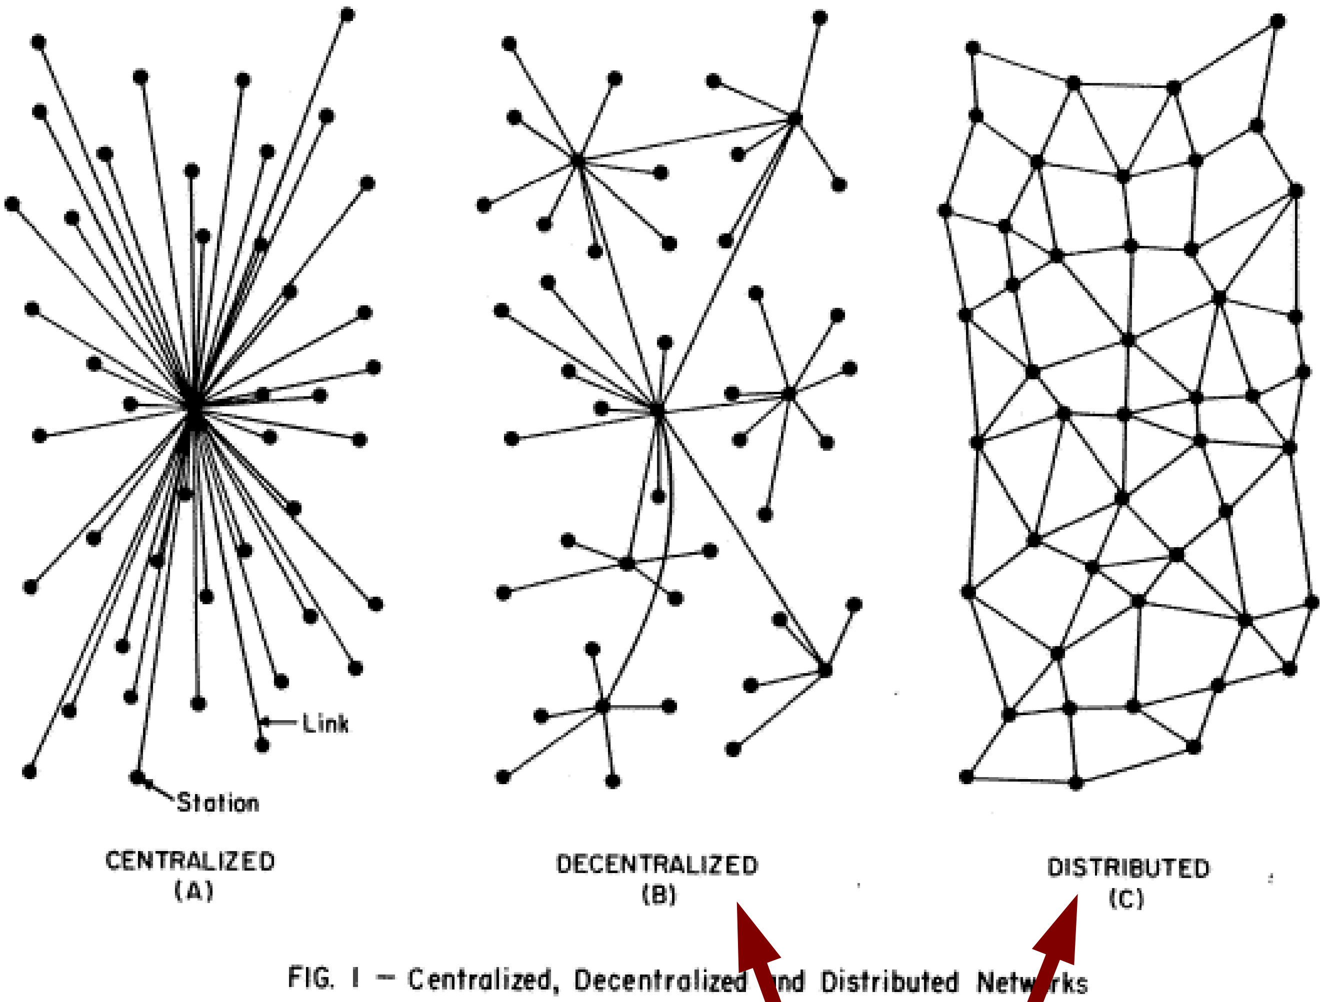

The Internet is a blend of these

Other origins of packet switching

- Formative work on queuing issues in message transmission systems (Leonard Kleinrock, MIT and UCLA)
- Independent invention of packet switching, but with emphasis on line-sharing aspects (Donald Davies, NPL, UK)
- Early experiments (NPL, MIT)

The packet pioneers

Baran

Kleinrock

Davies

Engineering principles known by the mid 1960's

- Mesh/star topology balancing redundancy and cost (Baran)
- Chop data streams into independent packets
  - resilience (Baran)
  - line-sharing (Davies)
  - mixture of media types and data rates (Baran, Davies)
- Adaptive routing (Baran)
- Logical addressing (Baran)
- Packet queuing disciplines (Kleinrock)

will mark important principles.

The ARPANET project

- Funded from 1968
- Potential bidder's reaction:
  "Frank... showed it to me. I couldn't imagine why anyone would want such a thing." (Severo Ornstein, of BBN, one of the bidders)
- Nevertheless, the contract went to BBN, mainly to provide hardware & software for the *Interface Message Processors (IMPs)* at each ARPANET site.

Some key people in 1969

- JCR Licklider (MIT and ARPA)
- Bob Taylor and Larry Roberts (ARPA)
- Kleinrock (UCLA)
  - Steve Crocker (UCLA)
  - Vint Cerf (UCLA/Stanford)
  - Jon Postel (UCLA)
- Dave Clark (MIT)
  - Bob Kahn (MIT/BBN/DARPA)

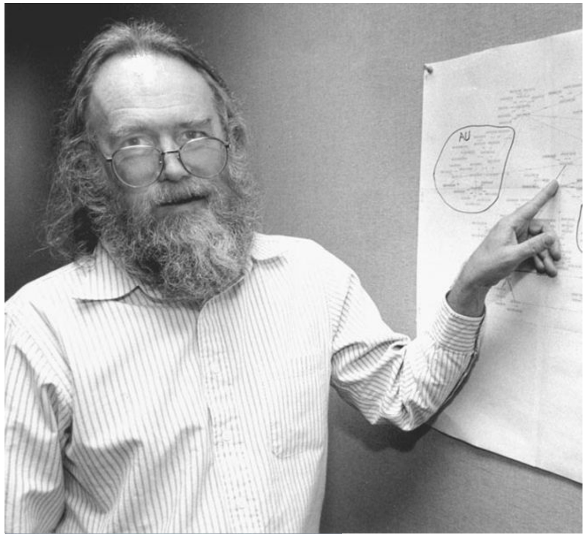

Postel

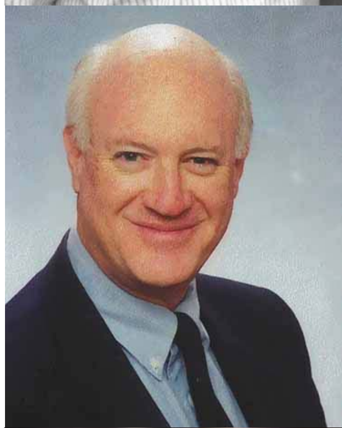

Crocker

Cerf

Kahn

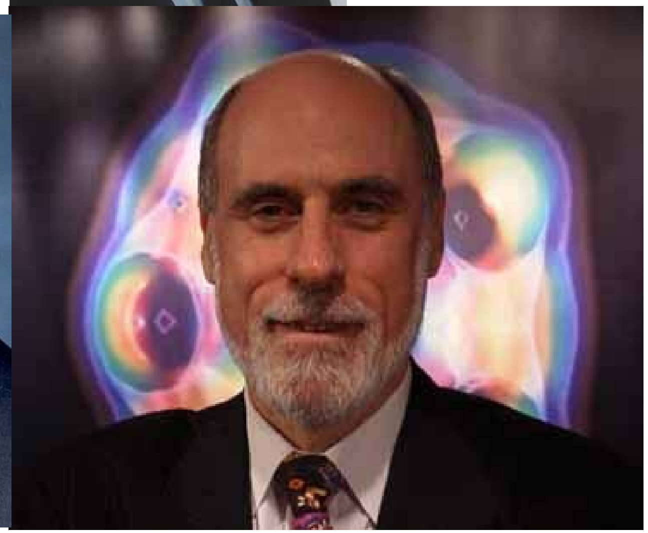

ARPANET in December 1969

Stanford Research Institute
UC Santa Barbara
UCLA
University of Utah

ARPANET in March 1977

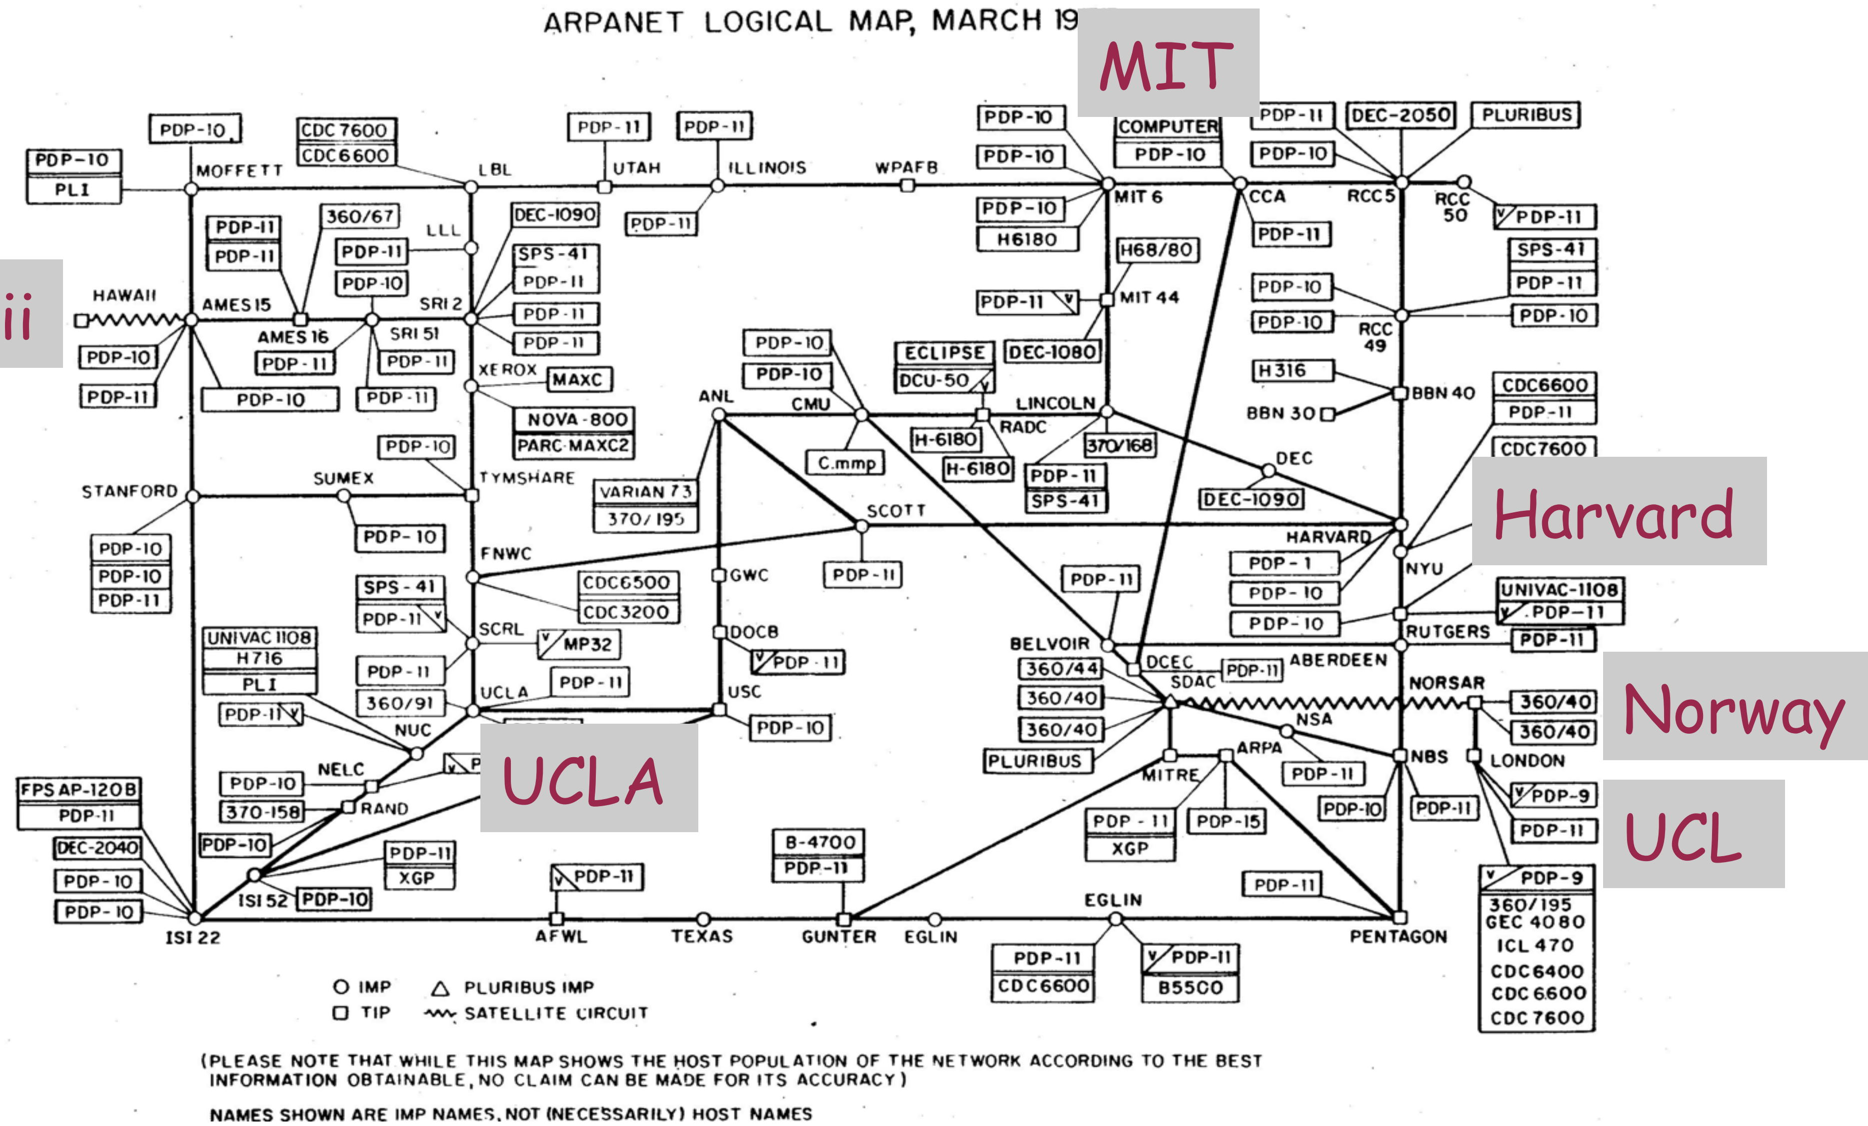

This was not a trivial testbed

- In fact it was a sustained, major, cooperative effort involving both *development* of radically new technology and *active use* of that technology.
  - The beginning of a long history of combining *research about networking* and *networking for research users*.
  - The beginning of a long history of pragmatic engineering informed by operational problems and user feedback.
  - The beginning of a long history of cooperative engineering aimed at the common good.

Problems in initial ARPANET design

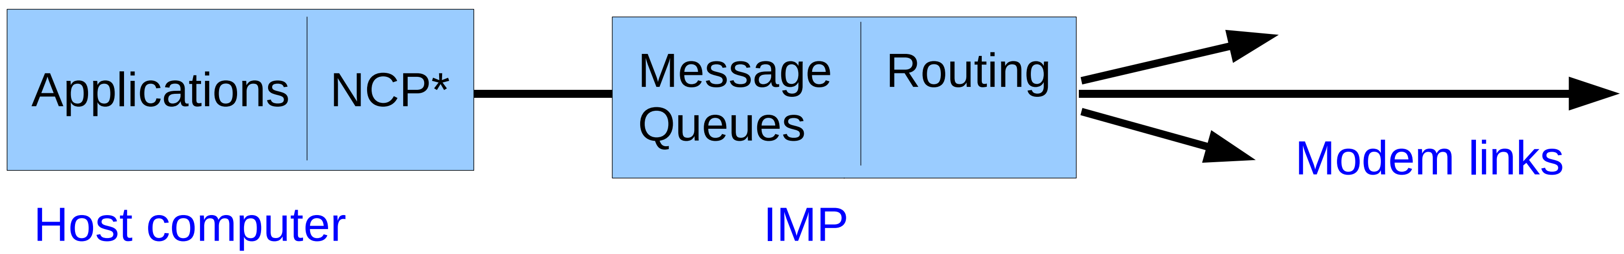

- No distinction between raw packets and application messages - NCP was monolithic
- Acknowledgments and retransmission handled by IMPs - no flexibility for host software
- IMPs couldn't prevent congestion overload
- Routing collapse was a possibility
  - actually occurred on October 27, 1980 - the last time the whole network was unusable.

*Network Control Program

Engineering progress in the 1970's

- Concept of a network of networks (originally called 'catenet', then inter-net) (Louis Pouzin)
- Splitting the NCP function into two layers (Cerf and Kahn)
  - Transport (end to end data streams, flow control, retransmission: TCP)
  - Internet Protocol (packets and routing)
- Resilient routing protocols

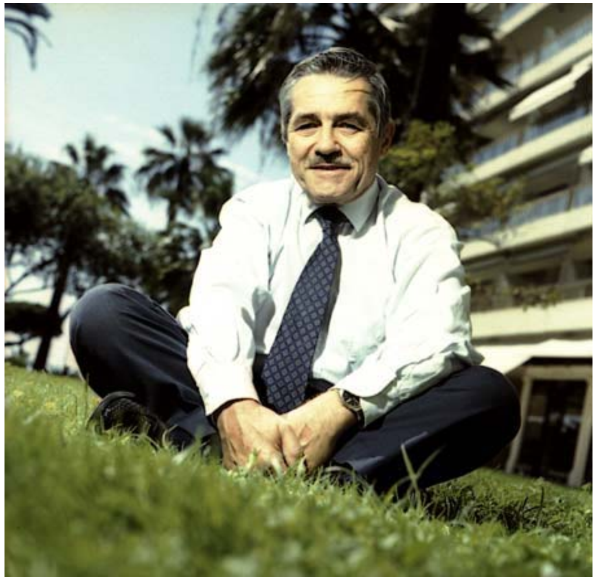

➔ The network switched to TCP/IP on January 1, 1983, and thereby became the Internet.

Pouzin's 1974 drawing

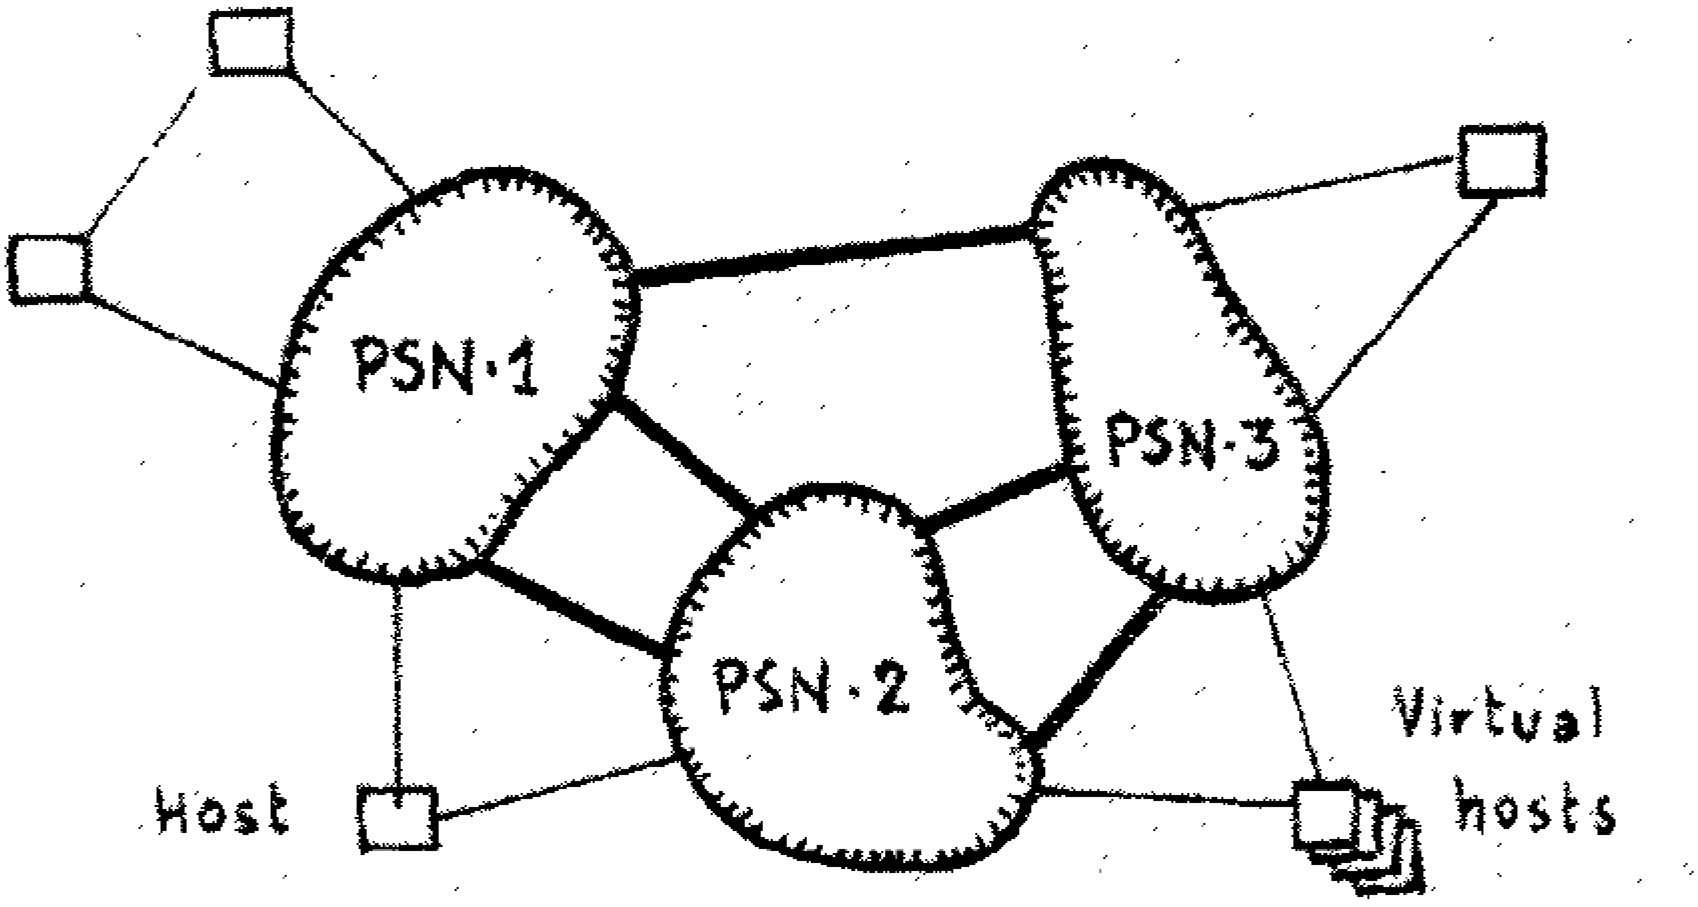

Fig. 1 - Network interconnection

Engineering principles known by the late 1970's

- Network-of-networks model (Pouzin)
- Layered protocol model (many people)
- End to end flow control in TCP (Cerf, Kahn)
- Resilient routing algorithms (many people, depends on graph theory)
- Pragmatic engineering informed by operational problems and user feedback.
- Cooperative engineering for the common good.

At this point, I will largely stop citing individuals. The work has been communal for many years.

Side note: software engineering in the 1970's

- Modular programming was a new buzzword.
- Layered architectures were considered leading edge and probably inefficient.
- Languages above assembler were considered hopelessly inefficient for real time work.
  - C was a lab project.
  - Strong typing and structured programming were "in", but object-oriented programming was a weird idea from Norway.
- The Internet's early expansion was in that context; actually the TCP/IP layering was very innovative.

Progress in the 1980's

- Free TCP/IP in BSD4.2 Unix (1983).
  - led to dominance of the C language in Internet software.
- Emergence of affordable routers, workstations, and PCs with TCP/IP support (1983-1985).
- TCP congestion control (1986+).
- Creation of the Domain Name System (1983-87).
- Consolidation of technical collaboration in the Internet Engineering Task Force (IETF) (1986).
- Rapid growth of TCP/IP throughout academia, with national research & education networks (NRENs) emerging in most developed countries.
  - extensive international collaboration

Side note: the protocol wars

- The 1980's also saw the protocol wars between proprietary solutions (DECnet, SNA,...), the "official" standards (OSI), and TCP/IP.
- TCP/IP was free, open, easy to use, and efficient.
- Thus it triumphed, firstly over OSI (despite government support) and secondly over the proprietary solutions.
- The war was all over by 1990, except for a few final skirmishes.

Killer applications by 1990

- email
- news groups (uunet merged with the Internet)
- remote login (telnet)
- remote file access (ftp, afs, nfs)
- information retrieval (wais, archie, gopher, www*)

➔ the scientific community showed a bottomless appetite for these applications, especially driven by Big Science with its enormous datasets.

*yes, that's the text-only line-mode world wide web, developed by the high energy physics community.

Engineering principles known by the late 1980's

- Names (identifiers), addresses (locations), and routes are different and should not be confused.
  - Unfortunately, we *did* partly confuse identifiers and locations.
- The end-to-end principle:
  - Each packet travels independently.
  - The end systems should not assume any function inside the network *except* the best-effort delivery of packets.
  - Functions such as error detection & correction, congestion control, retransmission and security should be provided solely by the end systems.

➔ This builds directly on Baran's original concept of robustness via redundancy

Sobering truths known by ~ 1990

- Scaling up the IP addressing system is a very hard problem, but 32-bit IP addresses will run out one day.
- Scaling up the wide-area routing system is a very hard problem, but the routers have a hard time keeping up with growth.
- There are bad people using the Internet. Security is a very hard problem.
- Somebody is going to invent a really compelling application sometime, and these problems will get worse as a result.

Example: the addressing problem

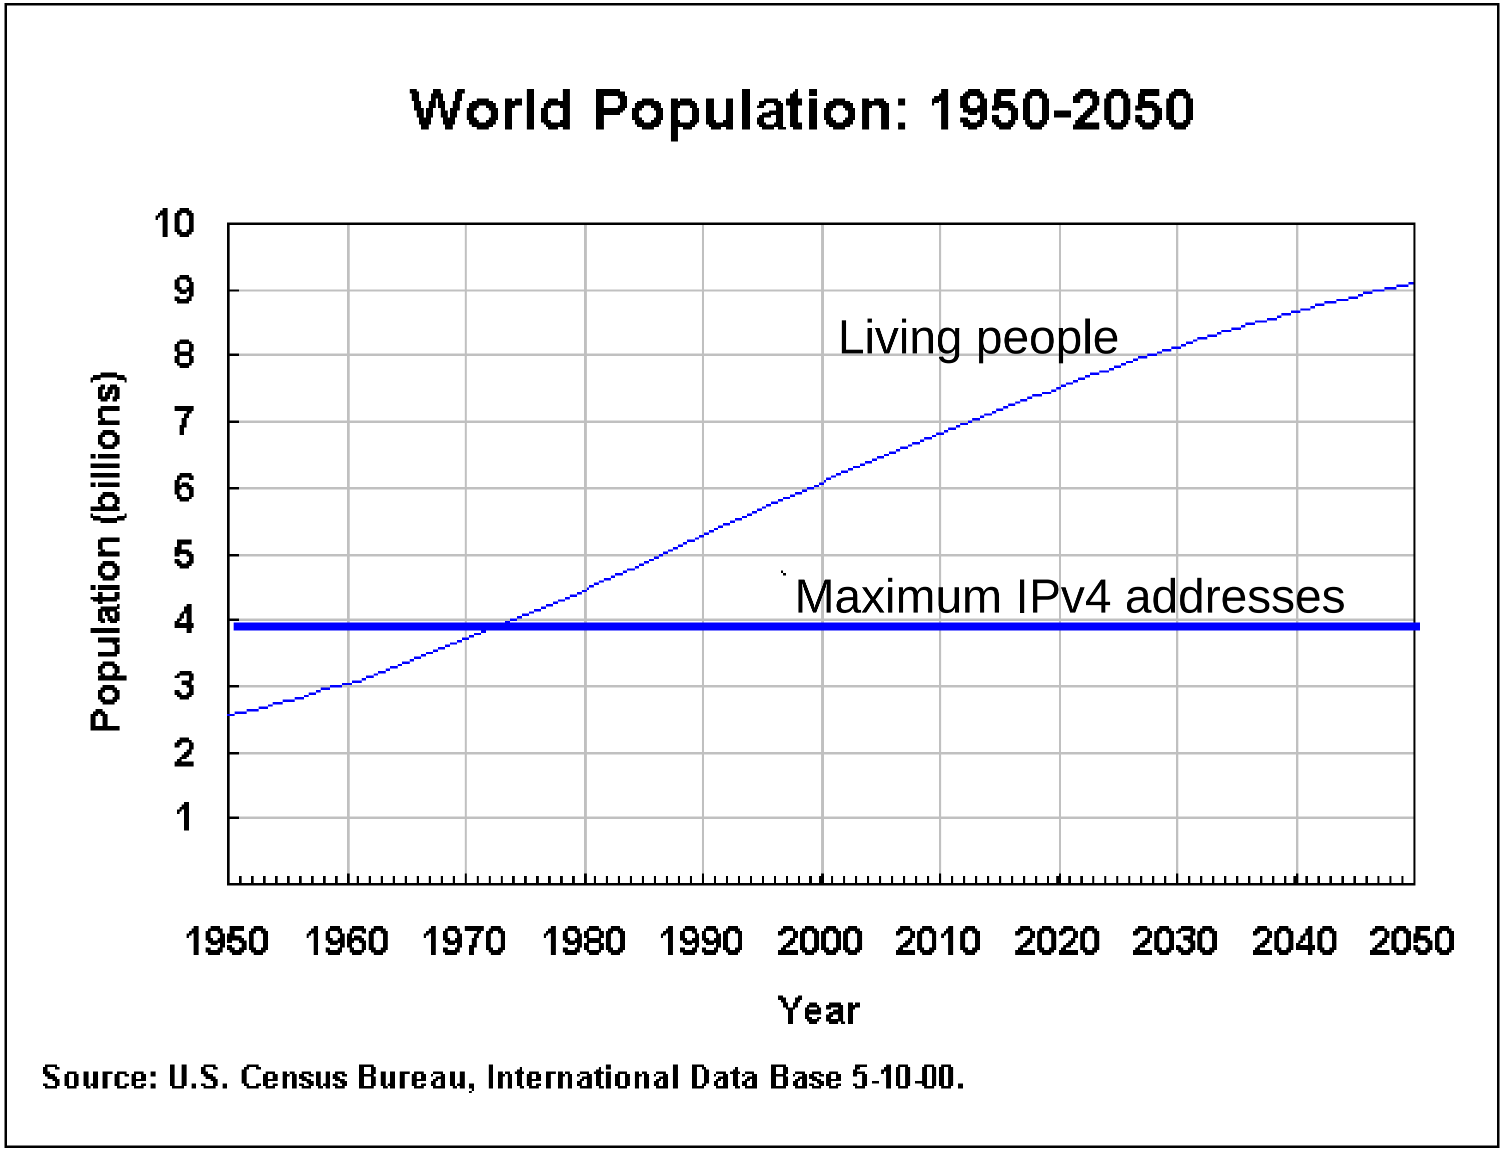

Obviously, having fewer addresses than people is silly

The perfect storm of 1993-95

- **The Mosaic web browser was released in 1993**
  - spread like wildfire
  - Netscape released in 1994
  - Internet Explorer released in 1995
- **The US Government privatised NSFnet in 1995**
  - Telecom liberalisation also began to take real effect
  - Internet service provision became a competitive sport
- **PCs got cheap; Windows 95 came with TCP/IP**
- **Companies from IBM to Enron declared the Internet to be The Next Big Thing.**

We know what happened next

- The Internet was transformed from a tool for research into a gold rush...

1999

...until the price of tulip bulbs collapsed.

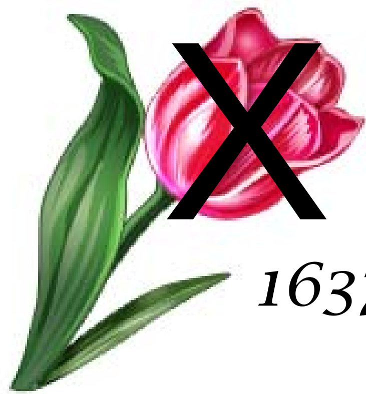

1637

Lessons from the perfect storm

- The underlying robustness principles of the Internet *actually work* during a perfect storm.
  - The Internet never crashed during its most spectacular growth period.
  - US telephone system was unusable from ~10 a.m. on September 11, 2001. The Internet just worked.
  - The Internet easily survived the various telco industry collapses when the bubble burst.
- But... the hard problems (addressing, routing, security) just got harder.
  - and as you'll recall, underlying growth continued.

Engineering changes since 1990

- Significant improvement in security protocols
  - but still no basic defence against denial of service
- Significant evolutionary improvements in routing
  - especially for inter-ISP mechanisms
- Improved congestion management
  - Traffic statistics matter
  - Queueing theory matters
- Surprisingly little basic change, except...
  - Widespread deployment of Network Address Translation boxes
  - Widespread deployment of security firewalls

Summary of engineering lessons

- Mesh/star topology for redundancy and cost
- Chop data streams into independent packets
- Adaptive and resilient routing protocols
- Logical addressing across a network of networks; names, addresses, and routes are different
- Layered protocol model
- End to end principle
- Pragmatic engineering with operational feedback
- Cooperative engineering for the common good

➔ These robustness principles *actually work* during a perfect storm. We'll forget them at our peril.

A few words about cooperative engineering

- The Internet Engineering Task Force performs cooperative protocol design.
- Regional and local bodies perform cooperative operational coordination
  - e.g. APRICOT meetings at the Asia-Pacific level
    NZNOG at the New Zealand level
- Registries also work on a socially cooperative (although technically hierarchical) model, e.g.
  - IANA at global level
    - APNIC at Asia-Pacific level
      - InternetNZ/NZRS at national level

Cooperative engineers

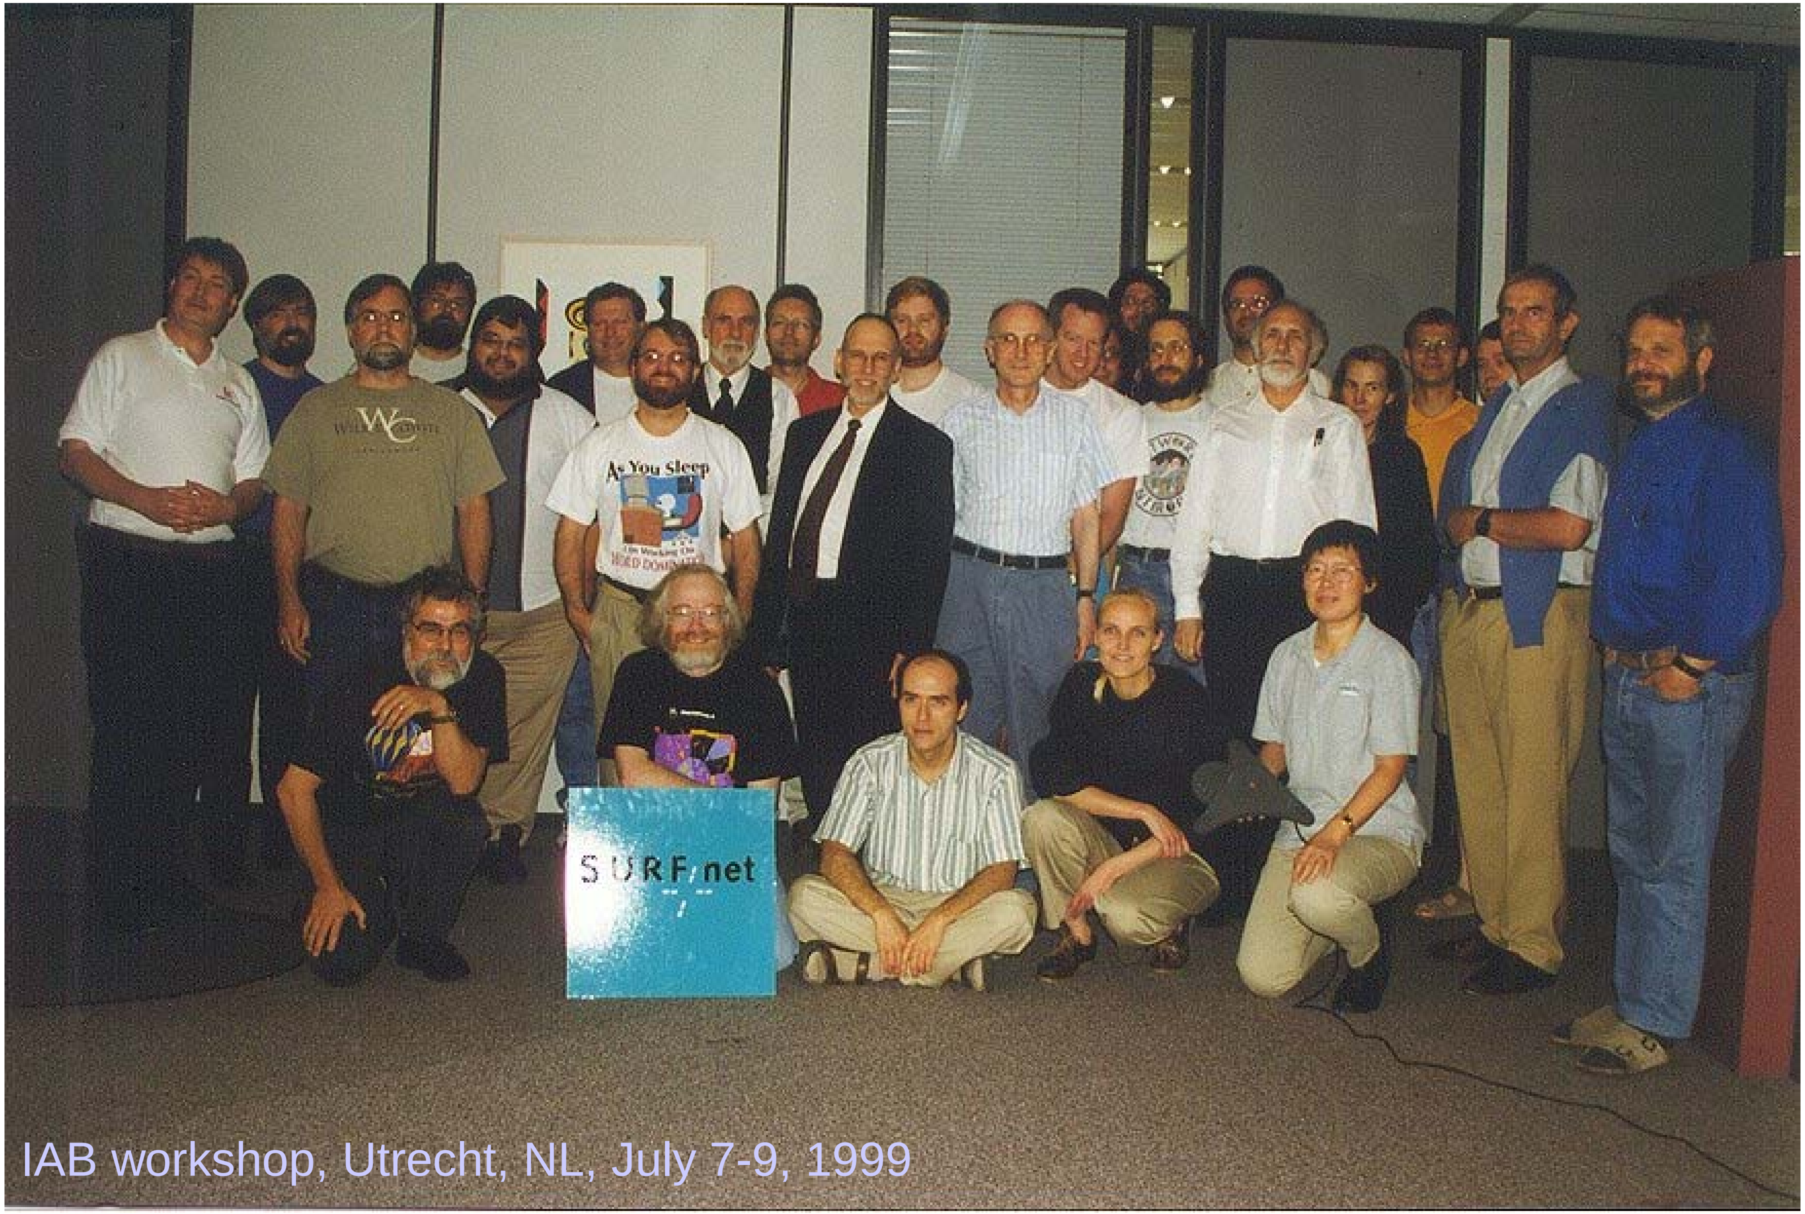

IAB workshop, Utrecht, NL, July 7-9, 1999

Engineering challenges for the future

- Address shortage
  - we *must* deploy IP version 6
- Loss of network transparency due to address translation and firewalls
  - slow deployment of complete end-system security
- Basic difficulty in avoiding *unwanted traffic*
- Concerns about wide area routing
  - we'll need to support 10,000,000,000 nodes
- Mobility
- Demand for increased bandwidth and quality
  - especially to support audio and video services

Where is the Internet going?

- Who knows? Nobody knew in 1969 where it would be today.
- I believe the engineering challenges will all be met in due time.
- As long as vested interests don't block innovation, with hundreds of millions of people connected to one network, there is no thinkable limit to what might be invented.

WATCH THIS SPACE!

Fixing the address shortage

- Today we use IP version 4 (32 bit addresses)
- Next step is IP version 6 (128 bit addresses)
  - Trillions of addresses possible, no more shortage
  - But old computers and software simply can't understand IPv6
  - Therefore, deploying IPv6 smoothly is far from easy
    - Must coexist indefinitely with IPv4
  - Really needs to happen over the next 5 years

The transparency problem

- Tight supply of IP addresses has caused widespread use of private address space in enterprise and domestic networks
  - Address translation (private<>public) breaks up the logical addressing of the Internet
  - With luck, IPv6 will fix this
- Insecure end-systems have caused widespread use of firewalls at the public/private boundary
  - Firewalls block innovative applications as well as unwanted traffic
  - Not obvious how to clear this blockage

Unwanted traffic

- Unsolicited or objectionable commercial email,
- Fraudulent email,
- Objectionable web sites,
- Bogus web sites,
- Floods of garbage ("denial of service attacks"),

➔ are all extremely hard to stop, since the network is only doing its job (delivering packets).

➔ *Expect this battle between good and evil to continue indefinitely; it's about people, not technology.*

Who's (not) to blame

- Note that susceptibility to viruses and worms is a problem of the end systems and their operating systems.
  - Blaming the network is like blaming the postal service for delivering a letter bomb. It doesn't solve the problem.

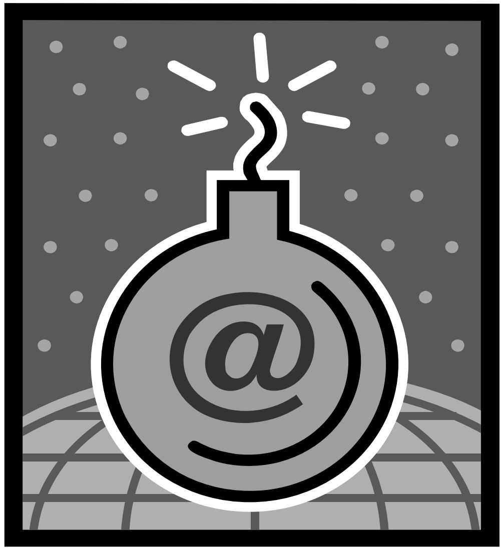

Wide area routing

- This is the "rocket science" of the Internet
  - Mechanism designed for the mainly academic network in 1988-92 now supports 550 million nodes and 45 thousand autonomous routing systems.
  - Apart from just routing trillions of packets a day, it must allow ISPs to manage traffic flow patterns and must allow for customers to "mix and match" ISPs
  - At the scale and speed required some years from now, this is still a research problem.

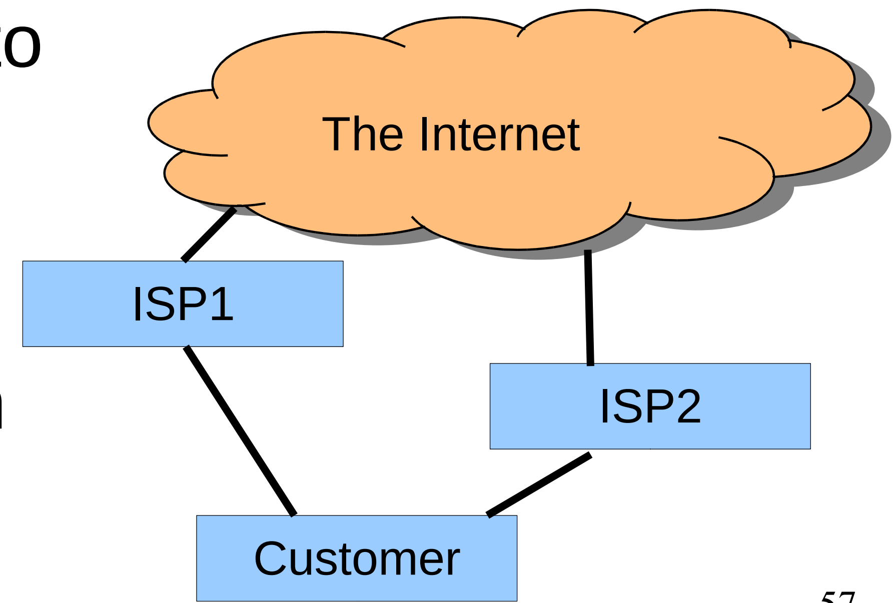

Mobility

- Mobile phones only have to solve the problem of roaming for single lengthy connections.
- Mobile Internet devices (and complete mobile networks) have to solve a harder problem
  - quickly (re)attaching to the Internet routing system
  - handling many packet streams in parallel
- Mobile ad hoc networks are even more interesting
  - allow systems to assemble themselves into a network on the fly and (re)attach to the Internet
- Moving from research to standards development

Bandwidth and quality of service

- Apparently, you can't stop progress. Higher bandwidth technology always seems to be under development.
  - Deployment is limited by business acumen and economics
- Up to now, quality of service has been provided by bandwidth management.
  - Relatively simple QoS technology is available to provide differentiated quality of service (e.g. for telephone traffic vs browsing)

More information

- Where wizards stay up late, K. Hafner & M. Lyon, Simon&Schuster, 1996.
- Casting the net, P.H. Salus, Addison-Wesley, 1995.
- Weaving the web, T. Berners-Lee, Harper Collins, 1999.
- How the web was born, J.Gillies & R. Cailliau, OUP, 2000.
- Exploring the Internet, C.Malamud, Prentice-Hall, 1992.
- Netscape time, J.Clark, St Martins Press, 1999.
- Burn rate, M.Wolff, Simon&Schuster, 1998.
- Why the Internet only just works, M.Handley, 2006, http://www.cs.ucl.ac.uk/staff/M.Handley
- http://www.isoc.org
- http://www.ietf.org