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

Journey to the Centre of the Internet

Professor Brian E Carpenter
Department of Computer Science
The University of Auckland
September 2010

Some thanks

- Robert Amor, John Hosking and Bob Doran, for their part in getting me to Auckland.
- Dick Bellamy and the VC, for the administrative achievement of appointing me as a part-time Professor.
- The entire CS department, for making me welcome.

A little self indulgence

1964 Cambridge (physics)

1967 Manchester (computer science)

1971 CERN (control software)

1974 →

← 1976

MASSEY UNIVERSITY
TE KUNENGA KI PŪREHUROA

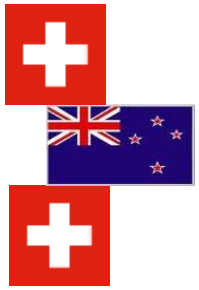

1985 CERN (networking)

1992 →

Internet Engineering Task Force

Internet Architecture Board
1994-2002

Internet Society
2000-2003

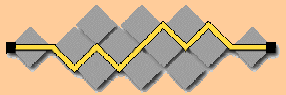

Chair, 2005-7

IETF

1997 IBM

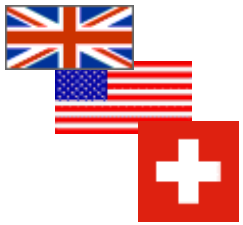

2007 THE UNIVERSITY OF AUCKLAND
NEW ZEALAND
Te Whare Wānanga o Tāmaki Makaurau

A little light cosmology

- By about -330, Aristotle was convinced by evidence that the Earth was the centre of the Universe.
- By about 1600, Copernicus, Galileo & Kepler were convinced by evidence that the Sun was the centre of the Universe.
- By 1915, Leavitt & Shapley were convinced by evidence that the Milky Way galaxy was the centre of the Universe.
- By 1929, Hubble, Slipher & Humason were convinced by evidence that the Universe was expanding in all directions around us (and the Milky Way was nothing special).
  - In fact, the Big Bang model suggests that *every* point is at the centre of the expanding Universe.
- How does this match the expanding Internet?

Topics

- Where did the Internet come from and how did it get here? How did it grow?
- Who controls it? Is that the centre?
- How are data routed? Is there a routing centre?
- Are there limits to growth?
- Aristotle or Hubble?

Warning: If you are hoping I will talk about the fluff on top of the Internet (things like Facebook, Google or spam) you will be disappointed.

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

...or maybe sooner

- Documents recently declassified by the UK and US signals intelligence services show that a data network for raw and decrypted intercepted signals was planned as early as 1946.
  - The plan was to link at least the UK, US, Oshawa (Ontario) and Melbourne (via Honolulu).
  - The UK-Canada-US link was apparently already in place by around June 1944.
  - One 1946 document refers to the Melbourne Sigint Centre as "not a purely national Centre. It is and will continue to be a joint U.K. - Australia - New Zealand organisation, manned by an integrated staff."

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.

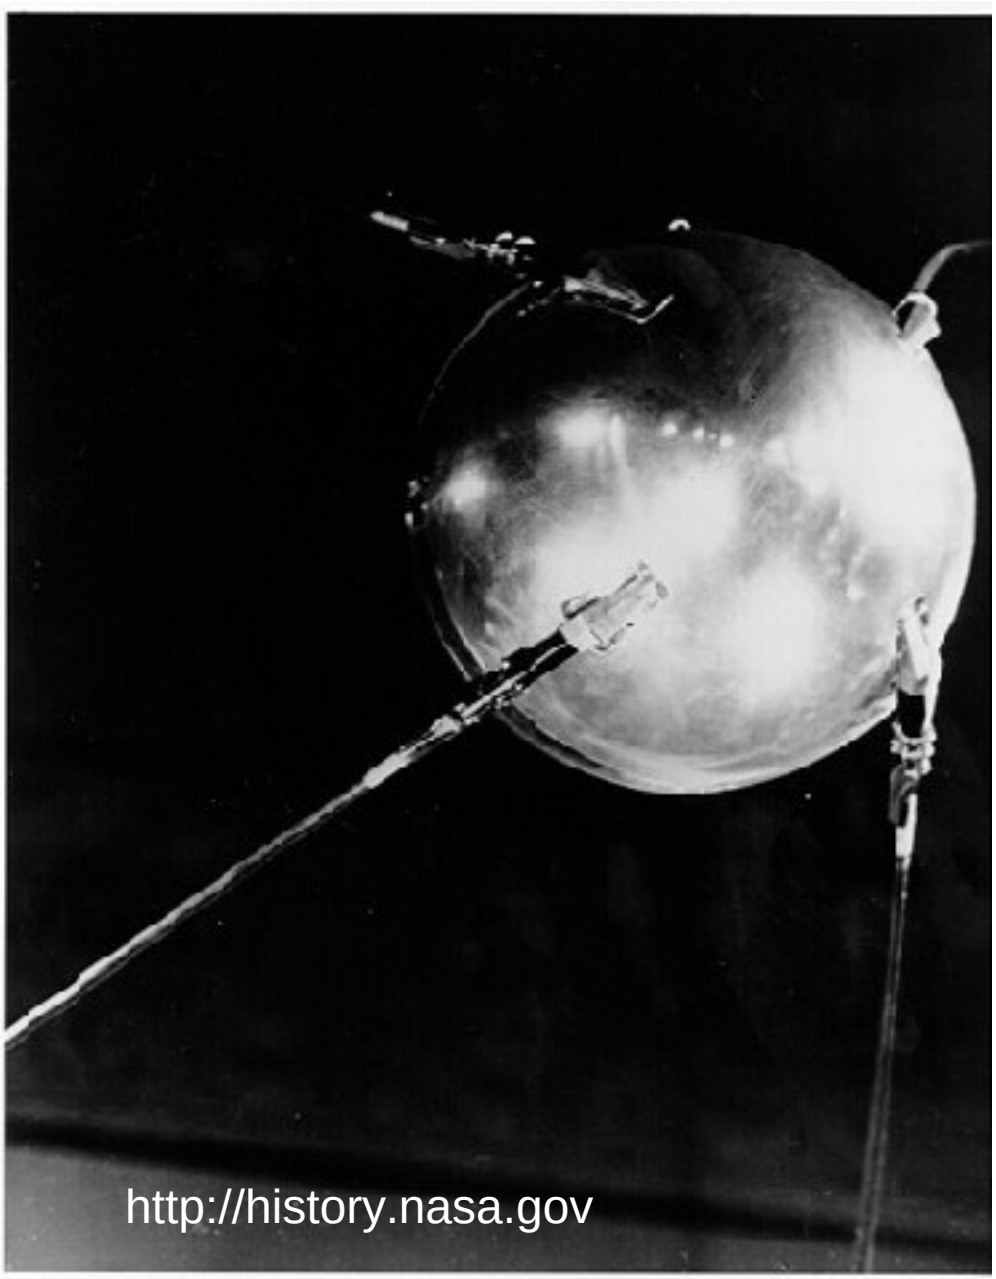

Helvetica font designed in 1957, too.

Another wake-up call

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

1962: Invention of packet switching

- Paul Baran at the RAND Corporation took on the task of designing networks to resist attacks (and other disasters). His ideas:

1. Redundancy and alternate paths
2. Split messages into small blocks (“packets”)
3. Adaptive routing (to route around failures)
4. Mixture of link types (e.g. cable and wireless)
5. Decouple the user's logical address from their physical address

Baran's 1962 topologies

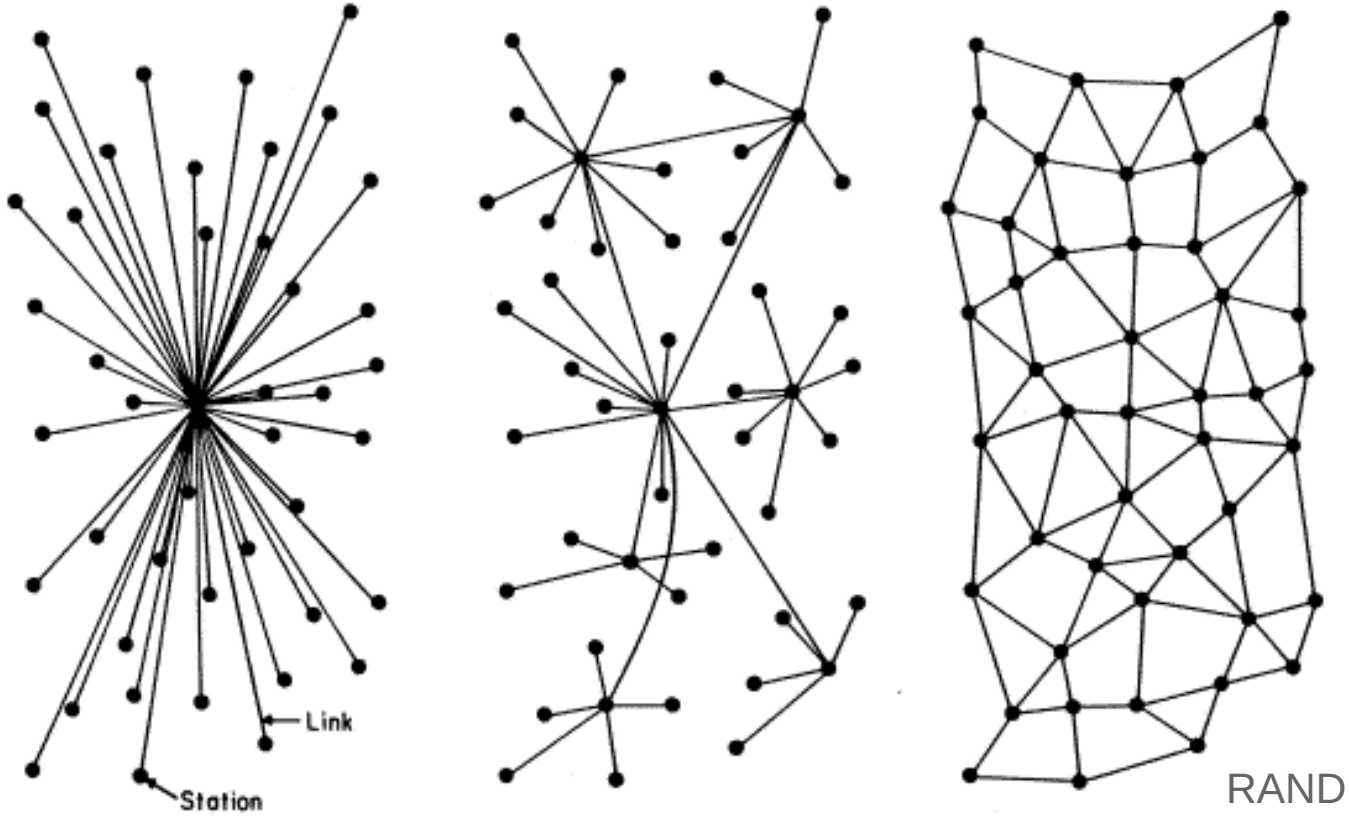

RAND Corporation

FIG. 1 – Centralized, Decentralized and Distributed Networks

The Internet today is a blend of these

Other origins of packet switching

- At least one WW II automatic cryptography system chopped messages into separate parts sent by separate routes.
- 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*)

*Massachusetts, not Manukau

The packet pioneers

http://www.ibiblio.org/pioneers

Baran

Photo by Louis Bachrach

Kleinrock

http://www.npl.co.uk

Davies

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)
- Nevertheless, the contract went to BBN, mainly to provide hardware & software for the *Interface Message Processors (IMPs)* at each ARPANET site.

http://ed-thelen.org/comp-hist

ARPANET in December 1969

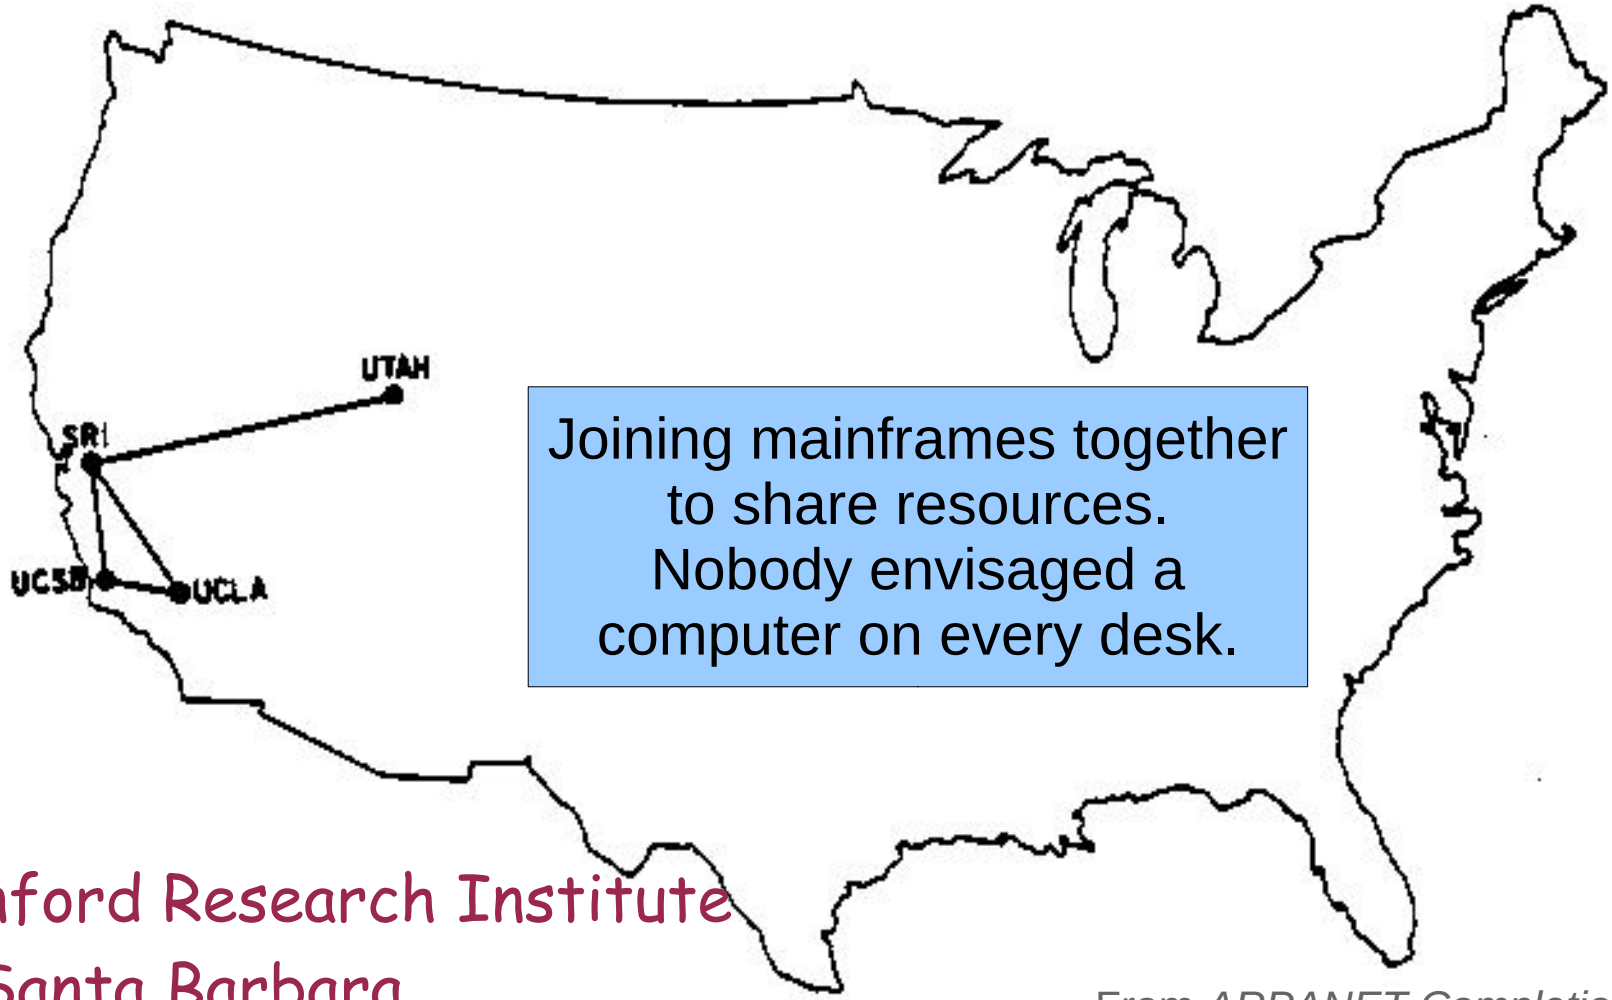

Joining mainframes together to share resources. Nobody envisaged a computer on every desk.

Stanford Research Institute
UC Santa Barbara
UCLA
University of Utah

From *ARPANET Completion Report*, BBN, 1978

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)

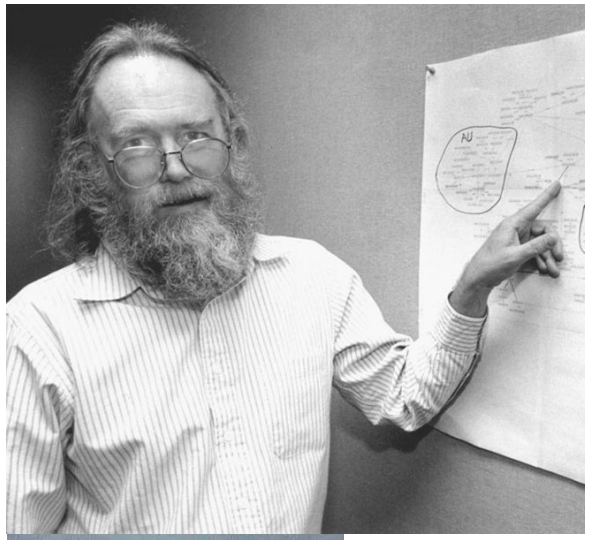

Postel

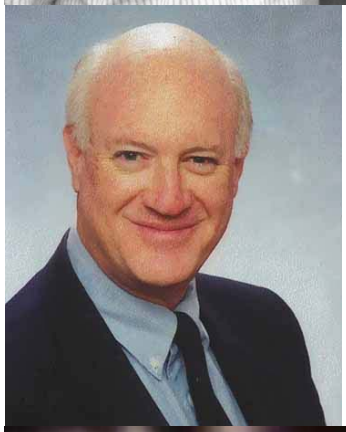

Crocker

Kahn

Cerf

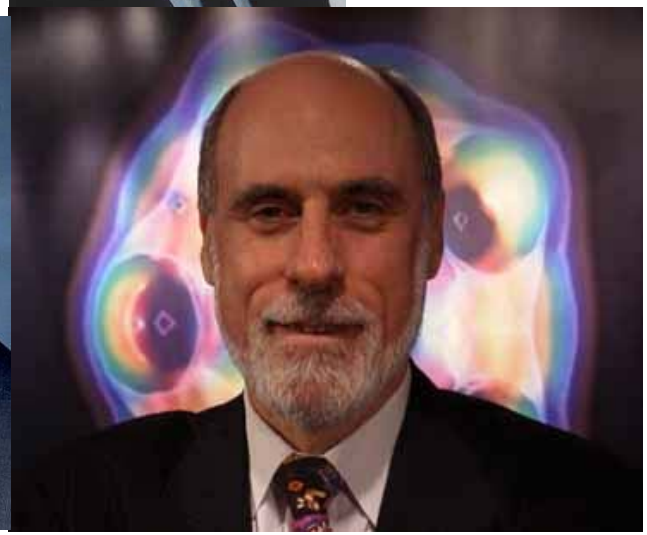

Photos: Irene Fertik, USC News Service, © 1994, USC; shinkuro.com; http://www.cnri.reston.va.us; Google

ARPANET in March 1977

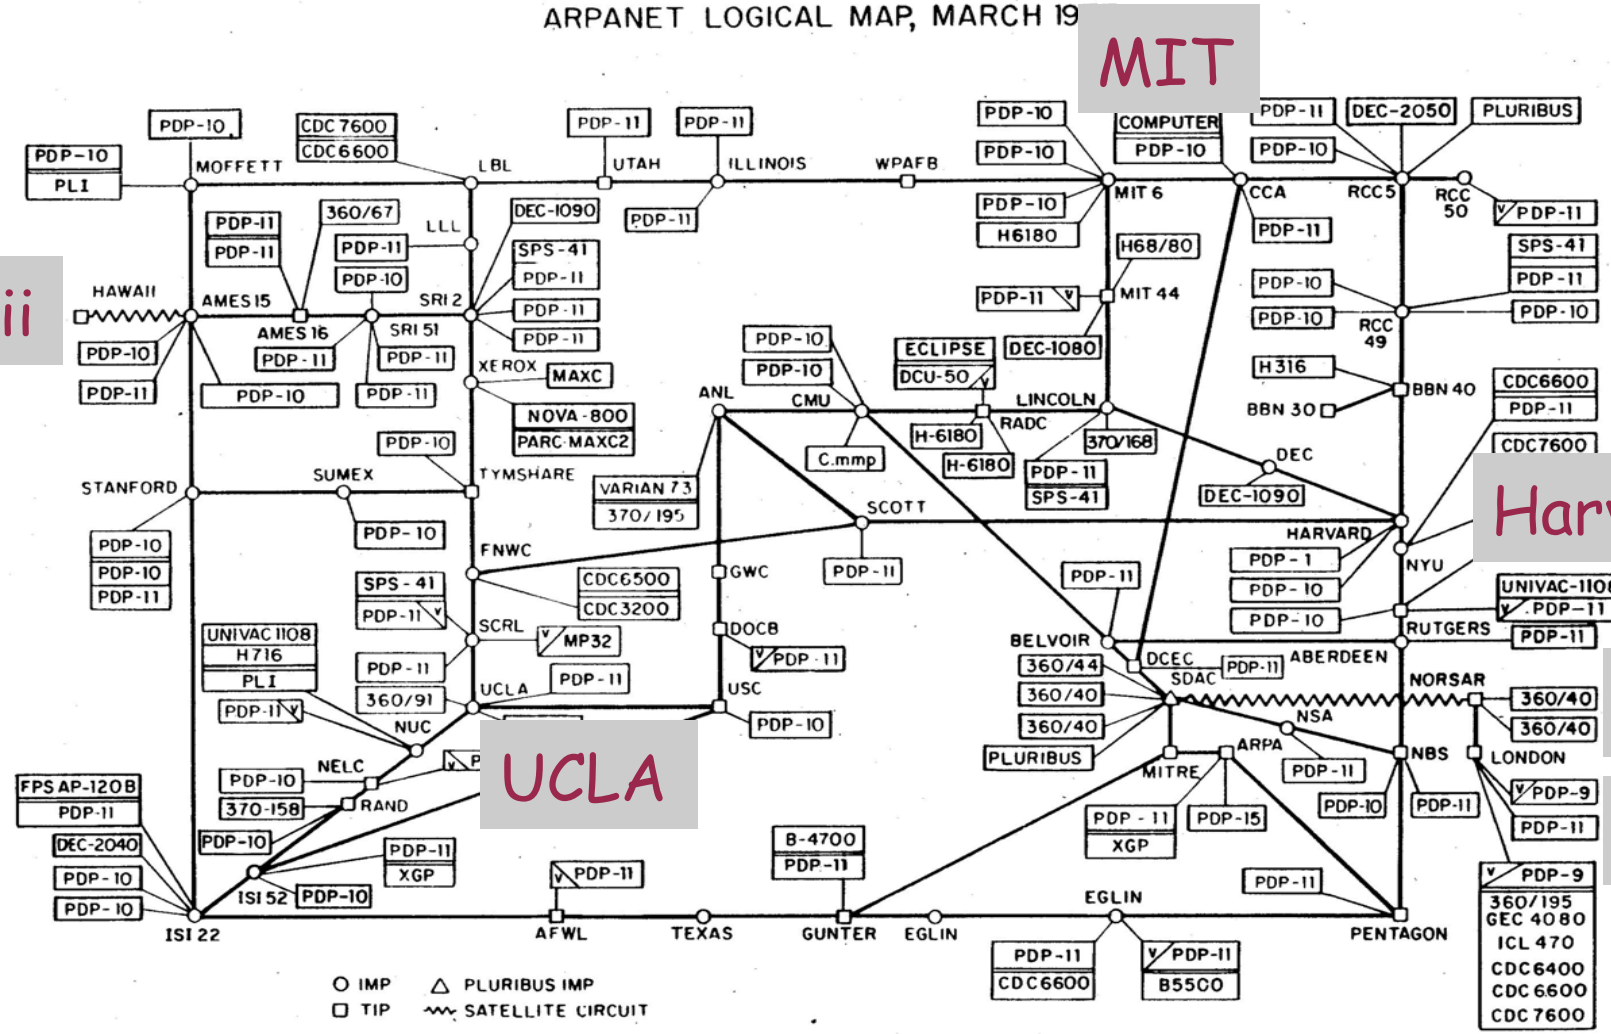

(PLEASE NOTE THAT WHILE THIS MAP SHOWS THE HOST POPULATION OF THE NETWORK ACCORDING TO THE BEST INFORMATION OBTAINABLE, NO CLAIM CAN BE MADE FOR ITS ACCURACY)

NAMES SHOWN ARE IMP NAMES, NOT (NECESSARILY) HOST NAMES

From *ARPANET Completion Report*, BBN, 1978

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.

How many computers on the ARPAnet?

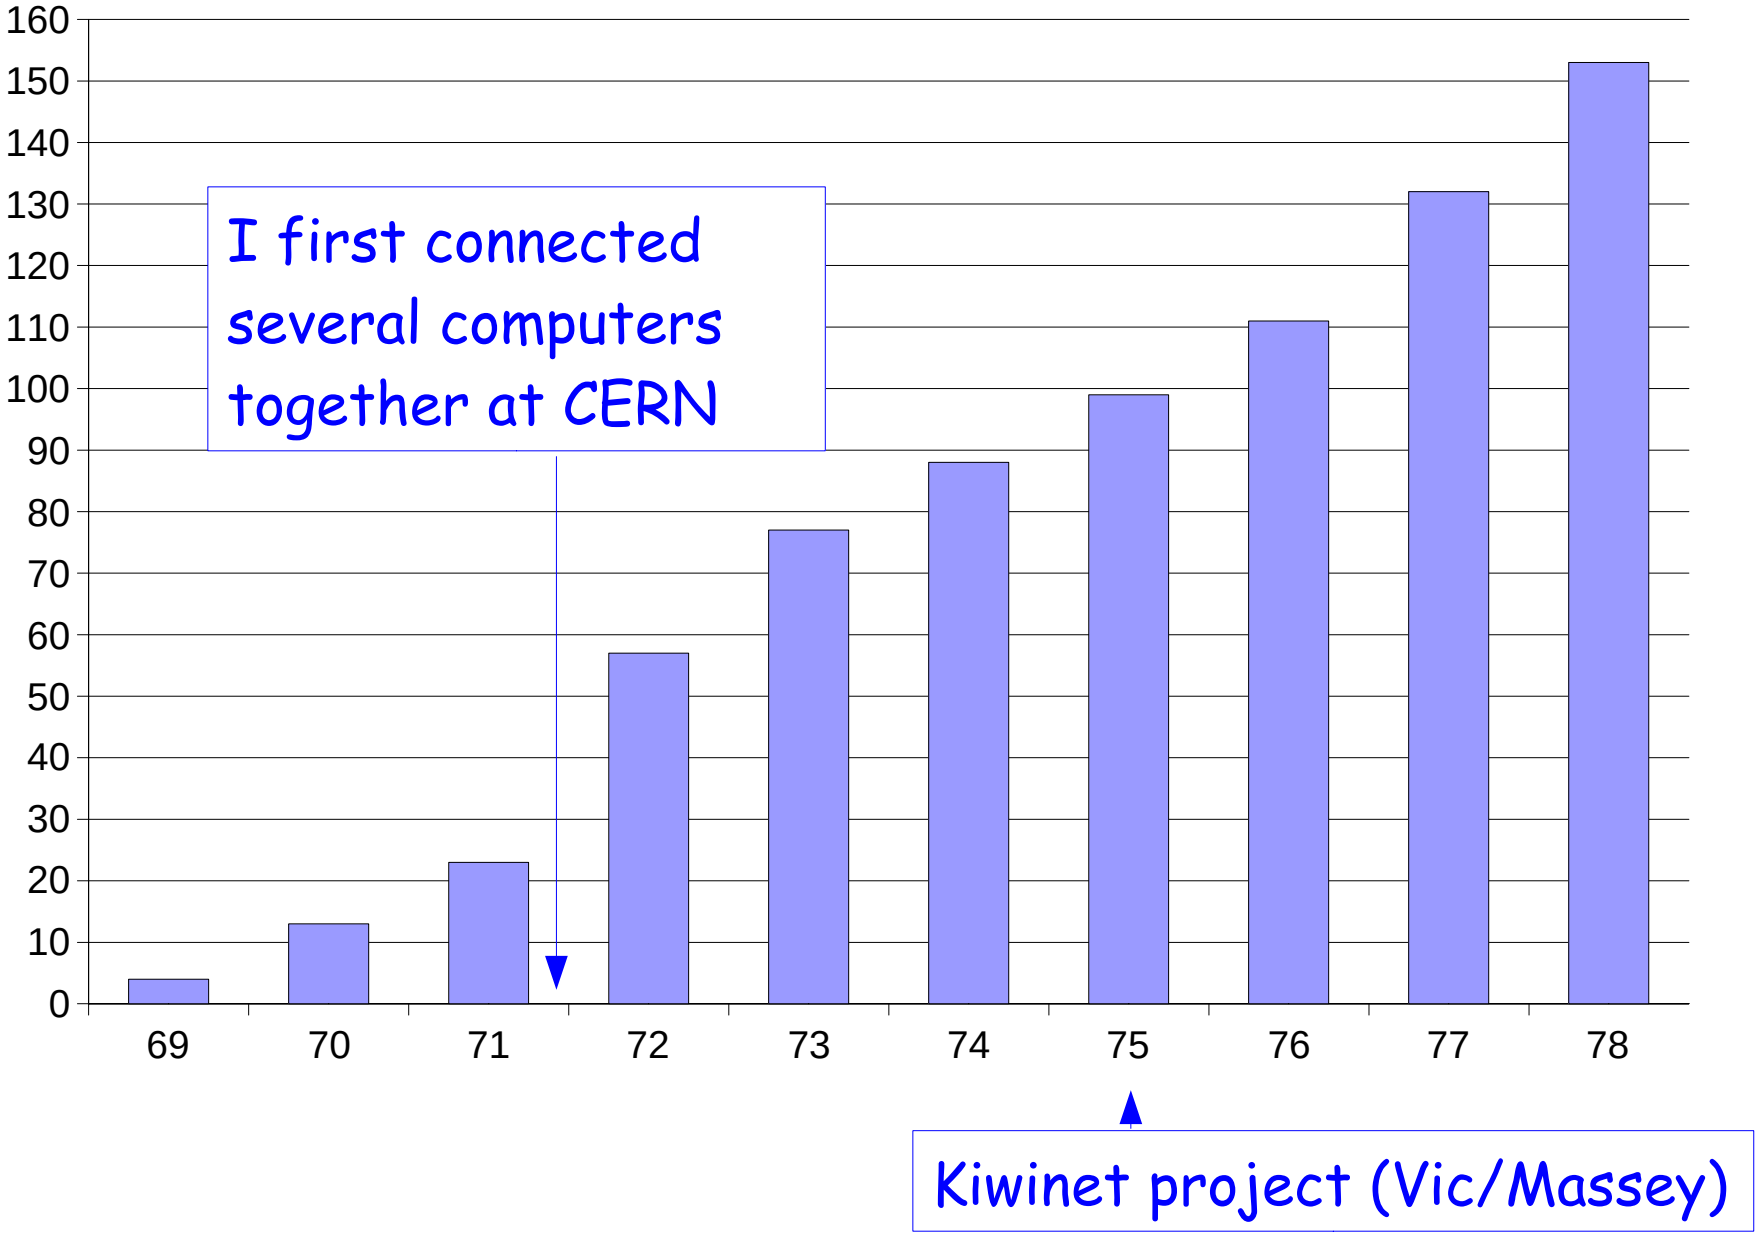

Using the network I built in 1971

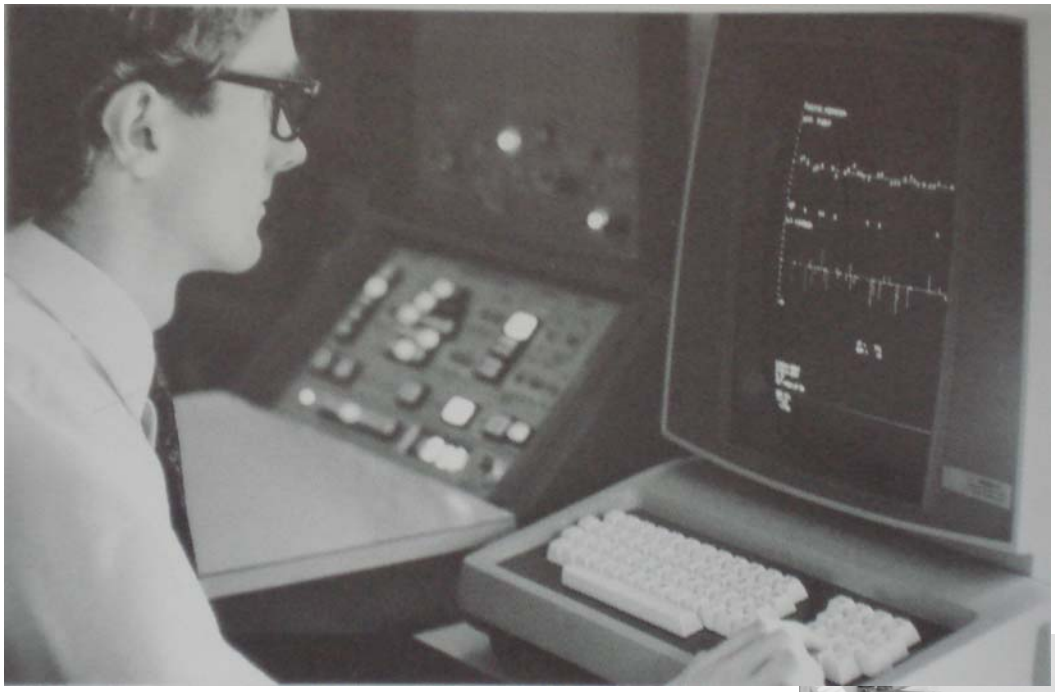

Imlac PDS-1
graphics computer

Photos CERN

Injection of 4
proton beams into
the CERN Proton
Synchrotron

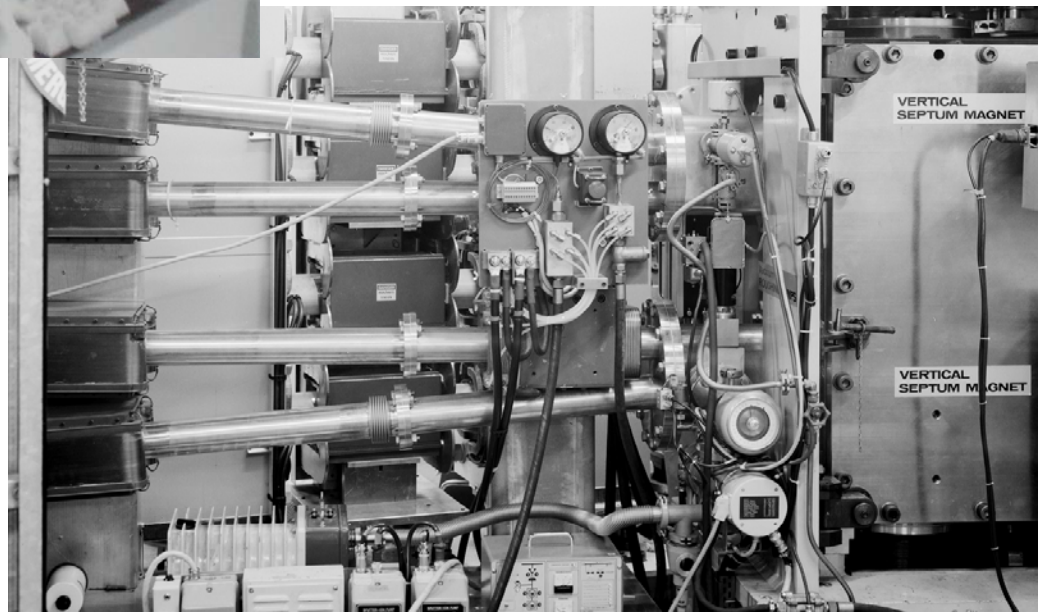

How many computers on the Internet?

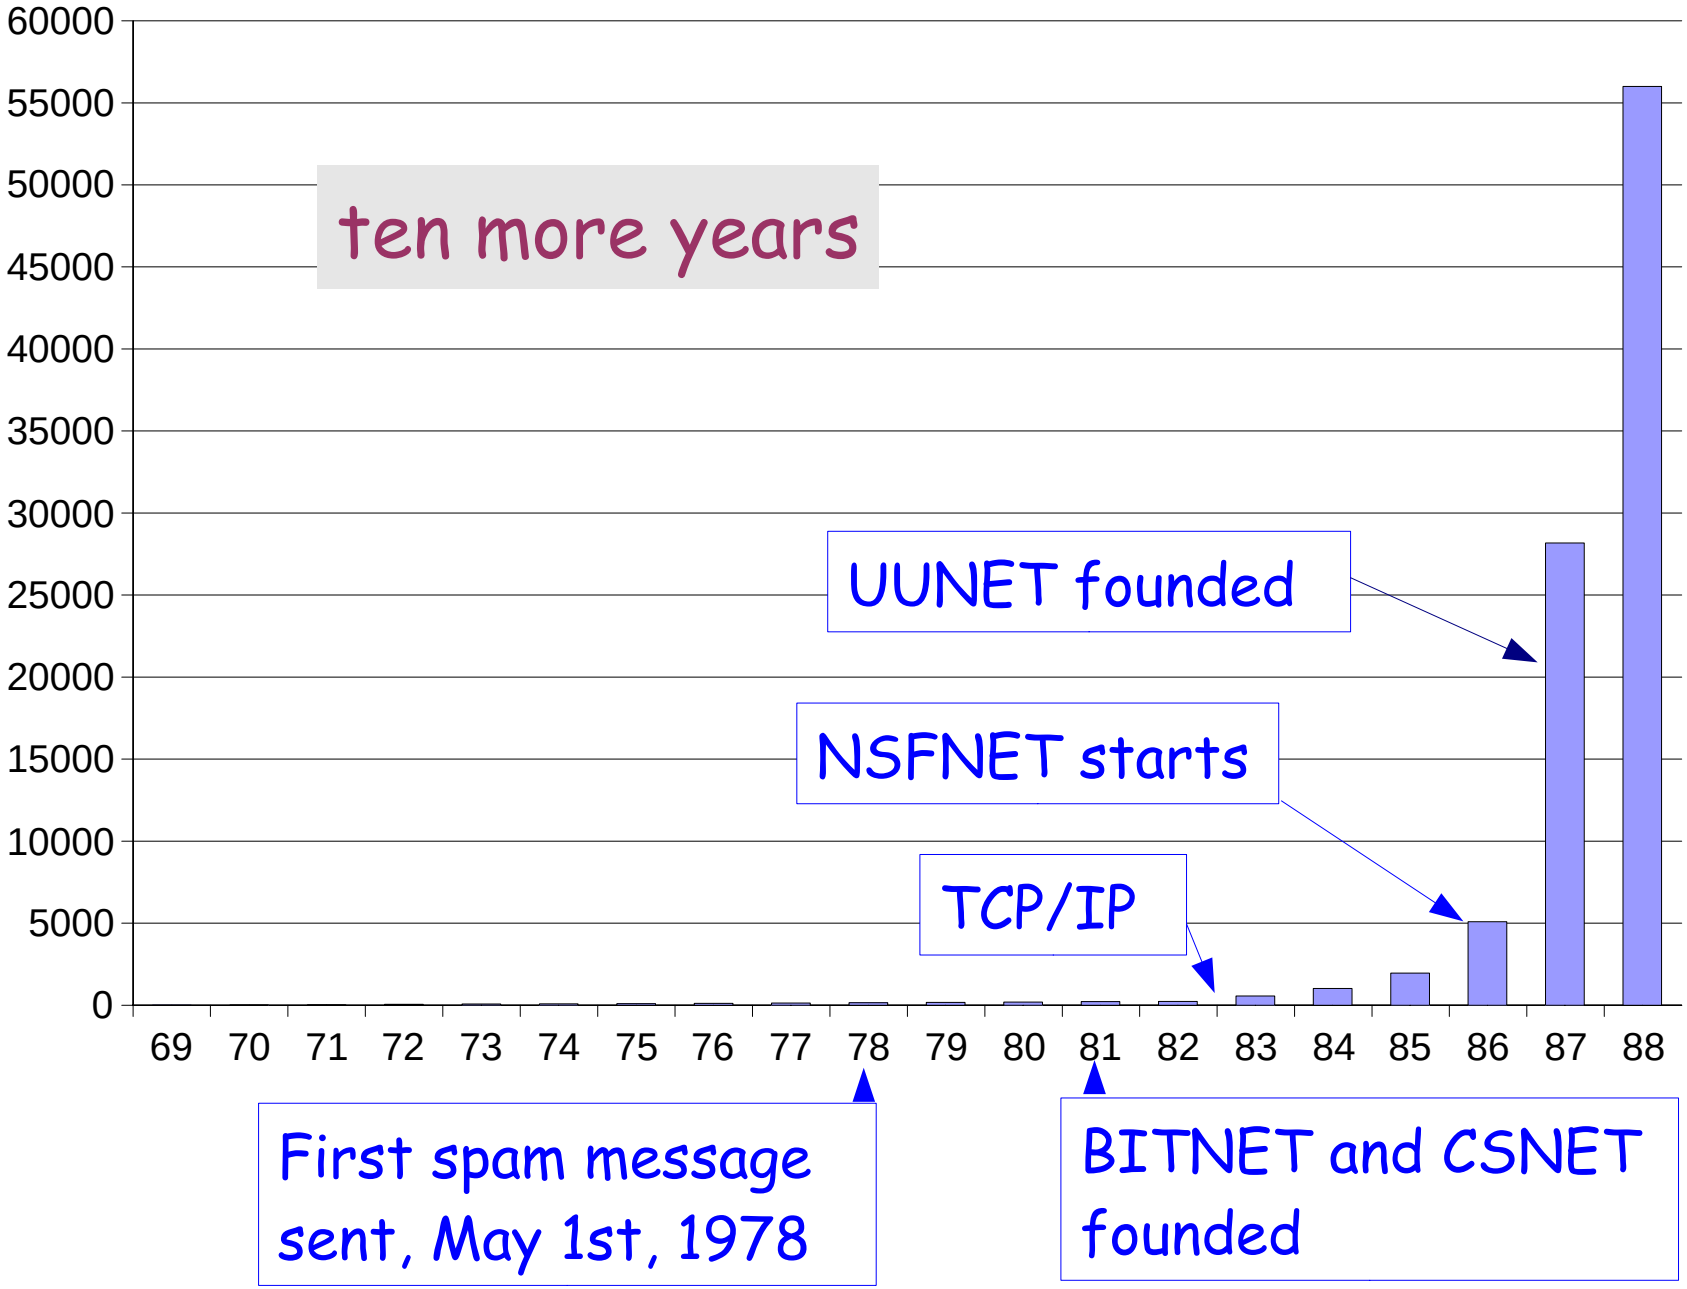

When did the ARPANET become the Internet?

- Concept of a network of networks (originally called 'catenet', then inter-net) (Louis Pouzin, 1974)
- Splitting the network software into two layers (Cerf and Kahn, late 1970s)
  - Transport (end to end data streams, congestion management: TCP)
  - Internet Protocol (packets and routing)

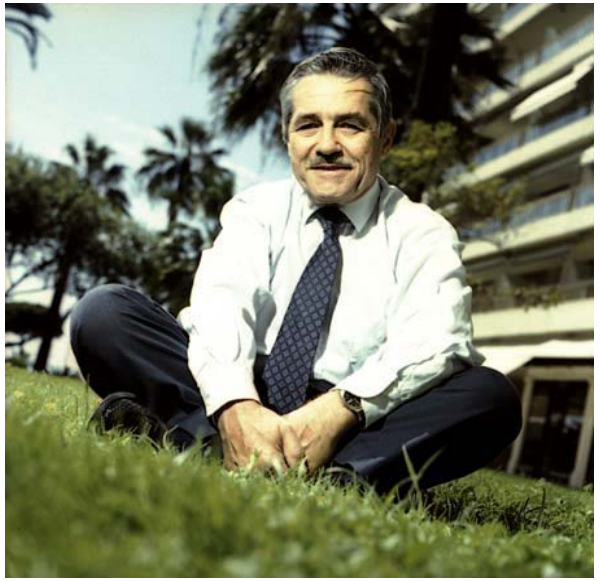

http://olivier.roller.free.fr

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

Pouzin's 1974 drawing

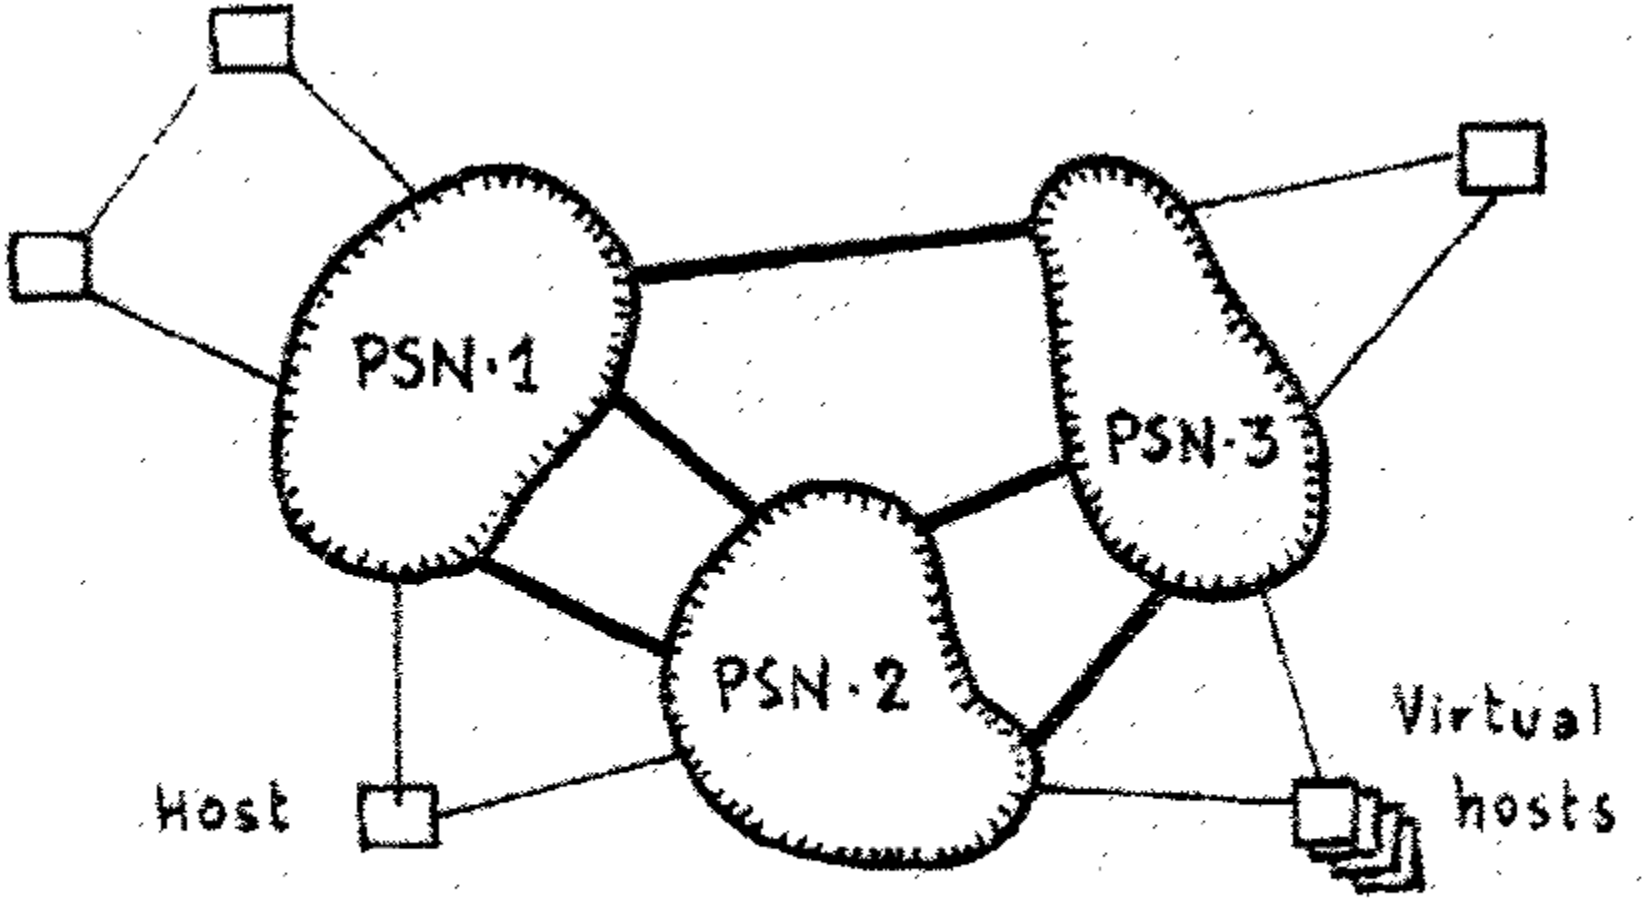

Fig. 1 - Network interconnection

L. Pouzin, 7th Hawaii International Conference on System Sciences, 1974

How many computers on the Internet?

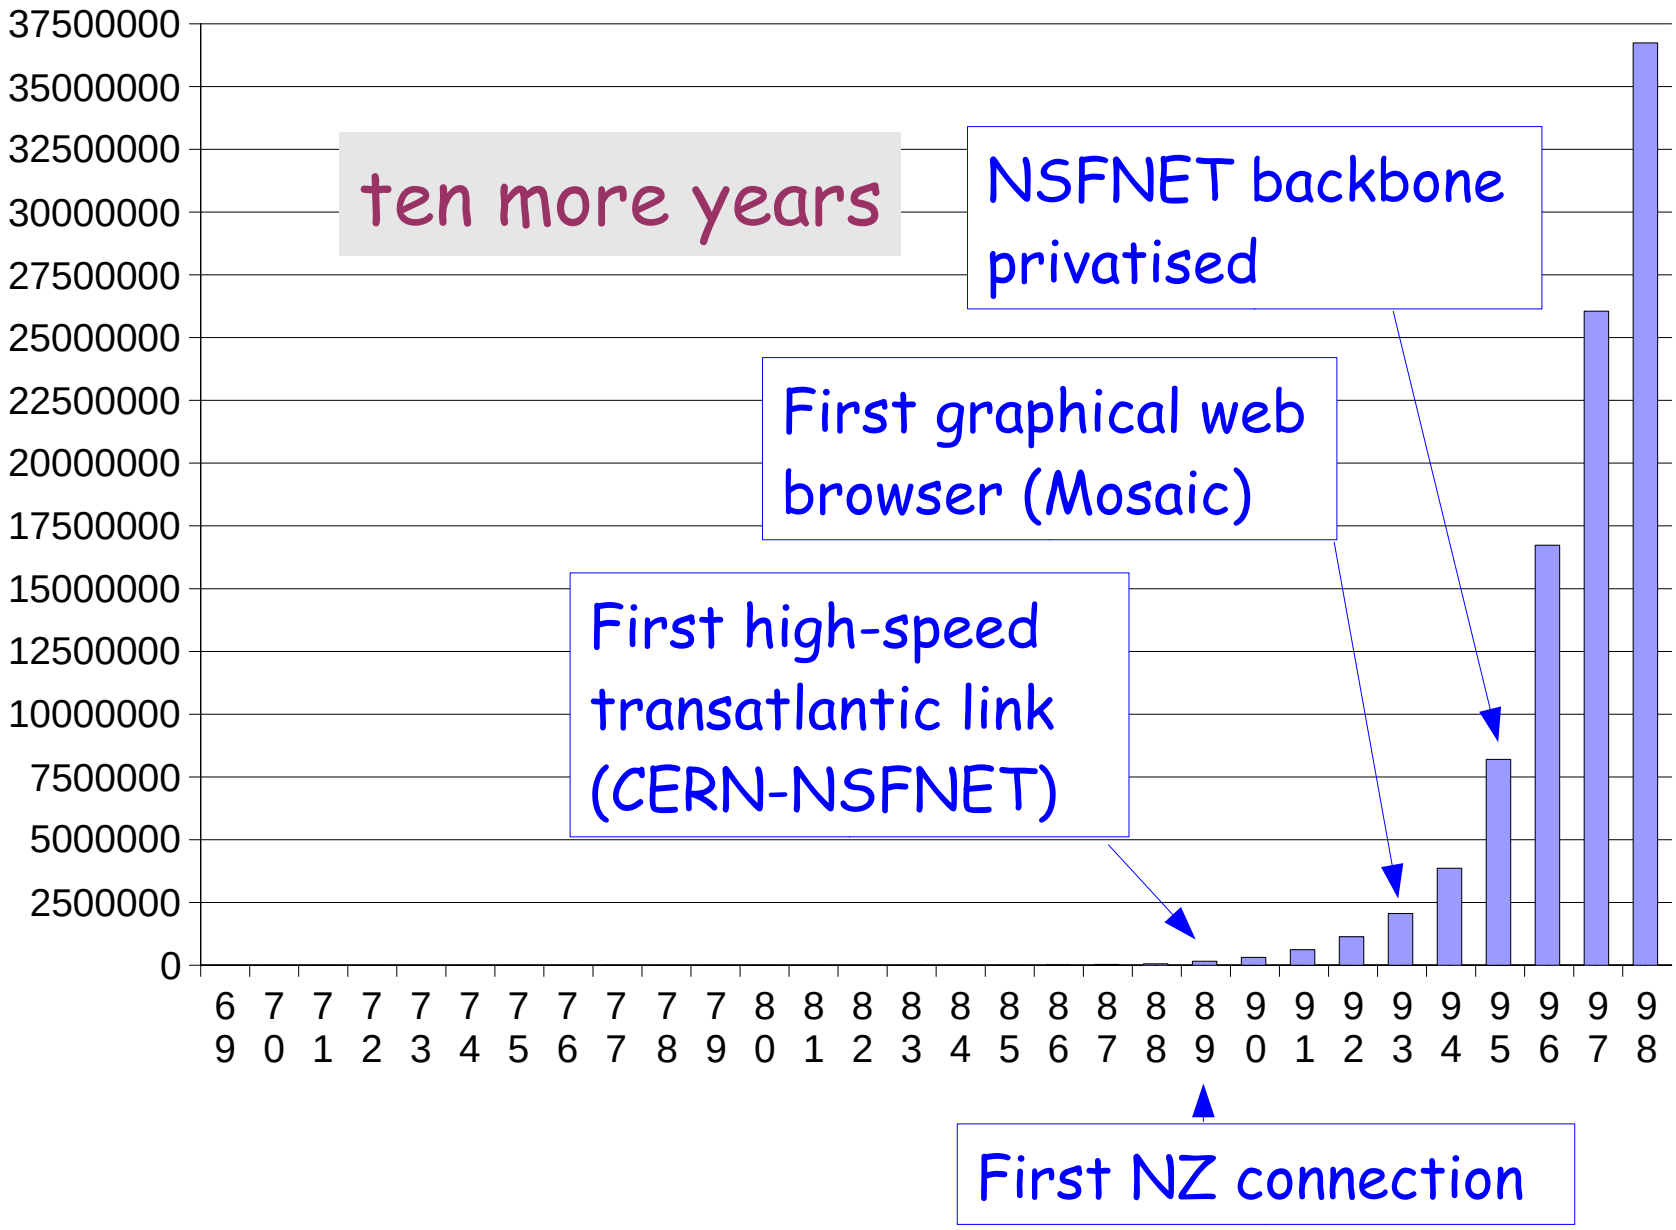

How many computers on the Internet?

(measured by published DNS names)

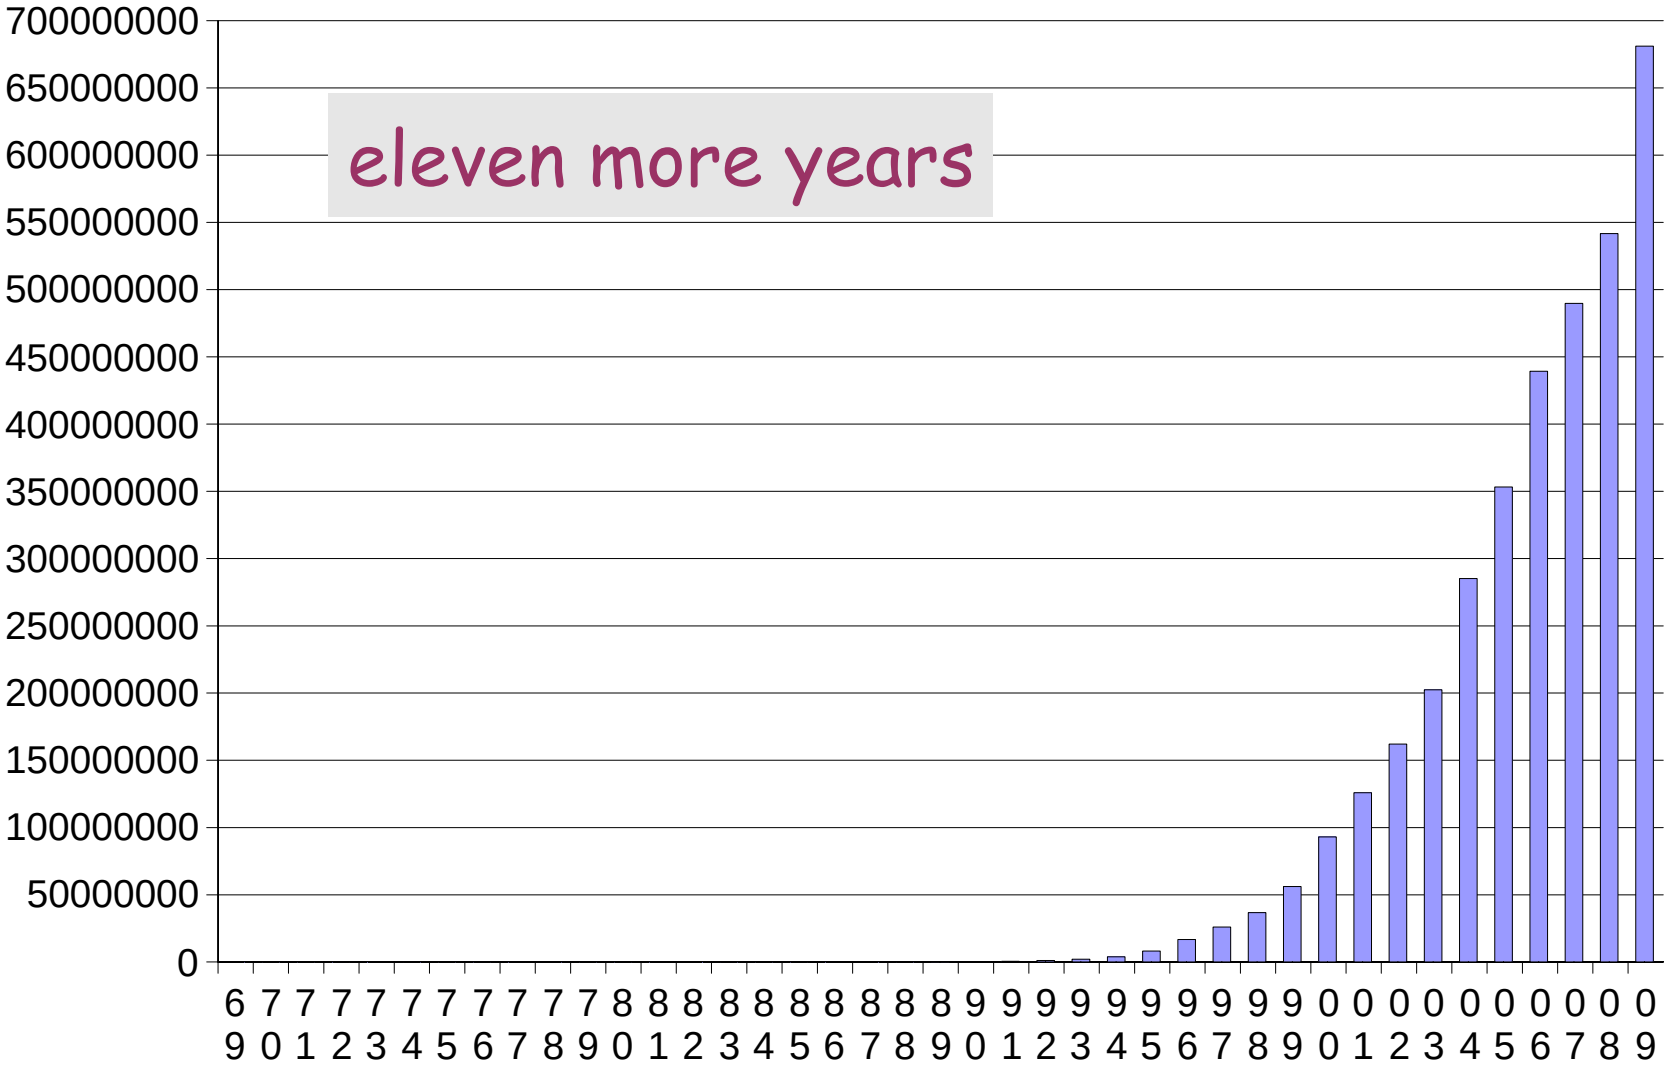

Data from http://www.isc.org/ and other sources

How many computers on the Internet?

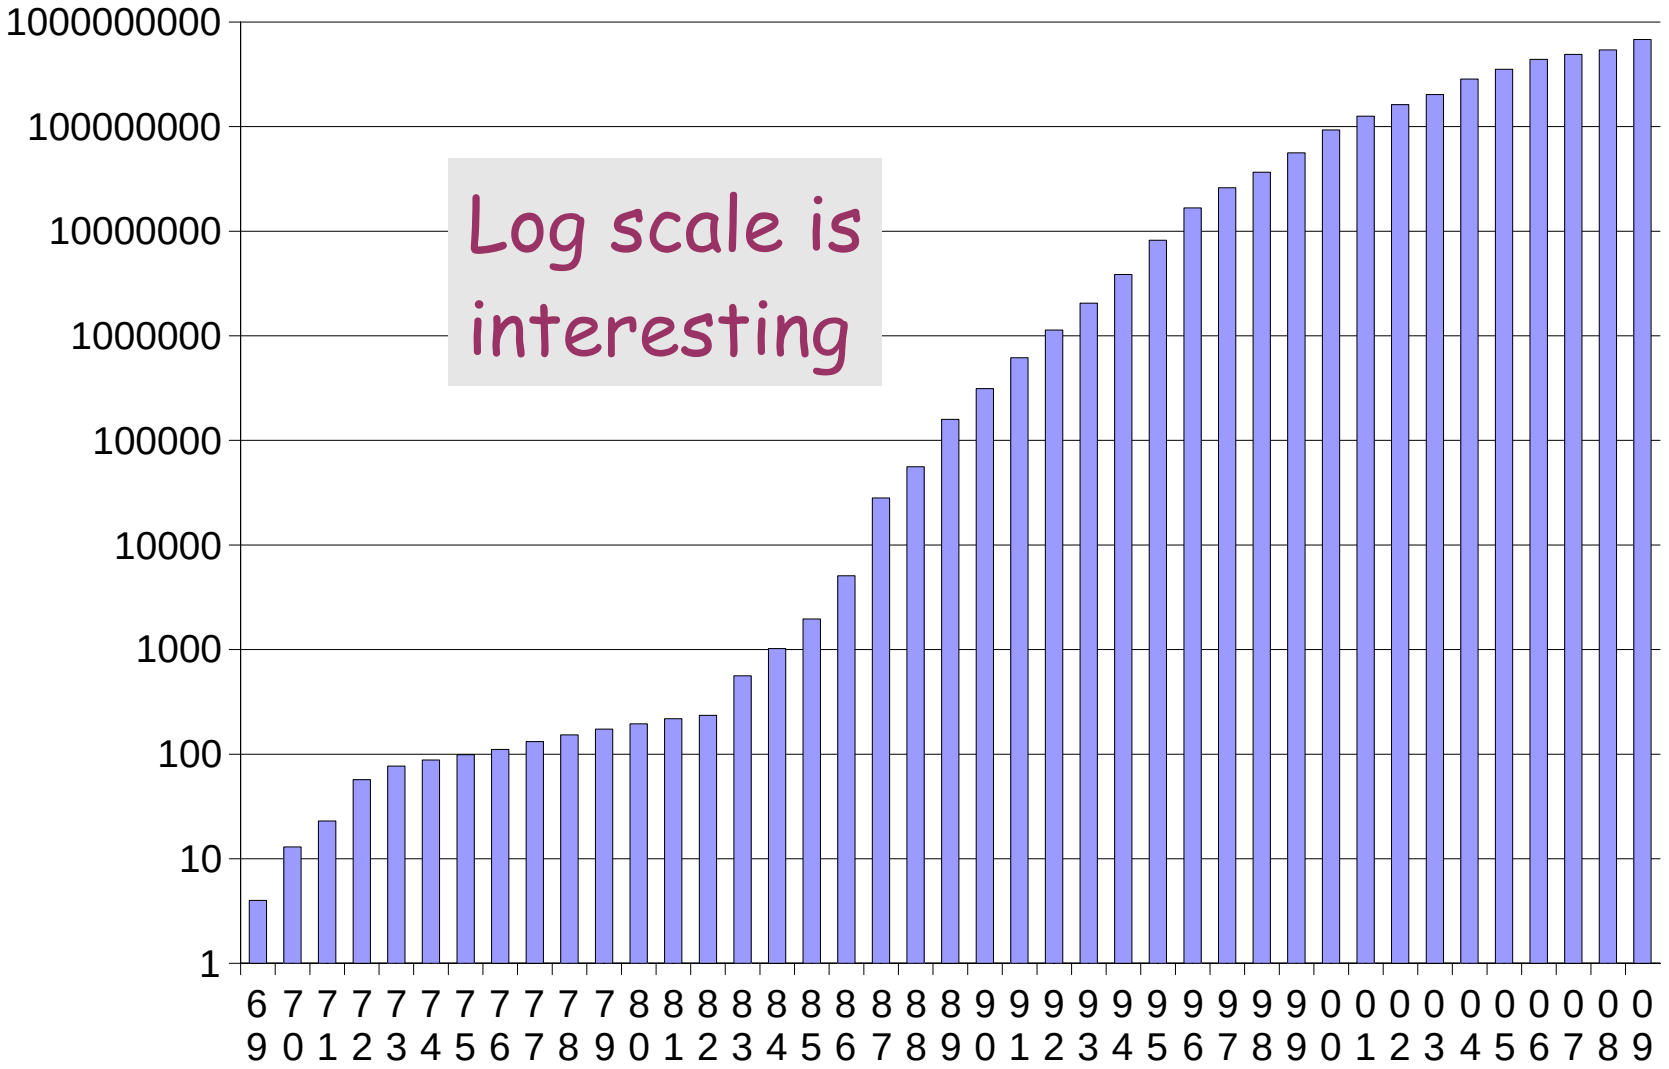

Data from http://www.isc.org/ and other sources

Key moments

- 1962 - packet switching invented
- 1974 - internet concept invented
- 1983 - TCP/IP, BSD Unix, Ethernet, workstations, routers, first intercontinental networkshop
- 1986 - NSFNET (and international peers)
- 1989 - High speed international links (for researchers)
- 1993 - Mosaic released (the “killer app” at last)
- 1994 - Policy-based routing (BGP4)
- 1995 - NSFNET stops, competitive ISPs established, Windows supports TCP/IP, commercial/private use explodes.

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

SmartMoney

Search More Options
Quote Site Search
Ticker or Company N
advanced search

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

shed: October 2000

end Comments
dd this column to your News Alerts

of the hc
iers. Wit
pping for
in the se

f 2000 w
unicatio
d GTE re
m (WCO
ue, and
the $50 b
dropped
om its bu
, the sha
gging far

Tribune de Genève

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

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

A list of who might control the Internet

- Governments
- The UN or the ITU (International Telecommunication Union)
- ICANN (Internet Corporation for Assigned Names and Numbers)
- Regional Internet registries
- Name registries and registrars
- The standards community (IETF, W3C,...)
- NGOs (InternetNZ, ISOC, Internet Governance Forum...)
- Equipment vendors (Cisco, Huawei, Juniper,....)
- Computer & software vendors (IBM, Microsoft, Apple,...)
- The open source community (GNU/Linux, Apache, Mozilla,...)
- Internet Service Providers (AT&T, Qwest, Sprint,...)
- Enterprise customers
- Consumers

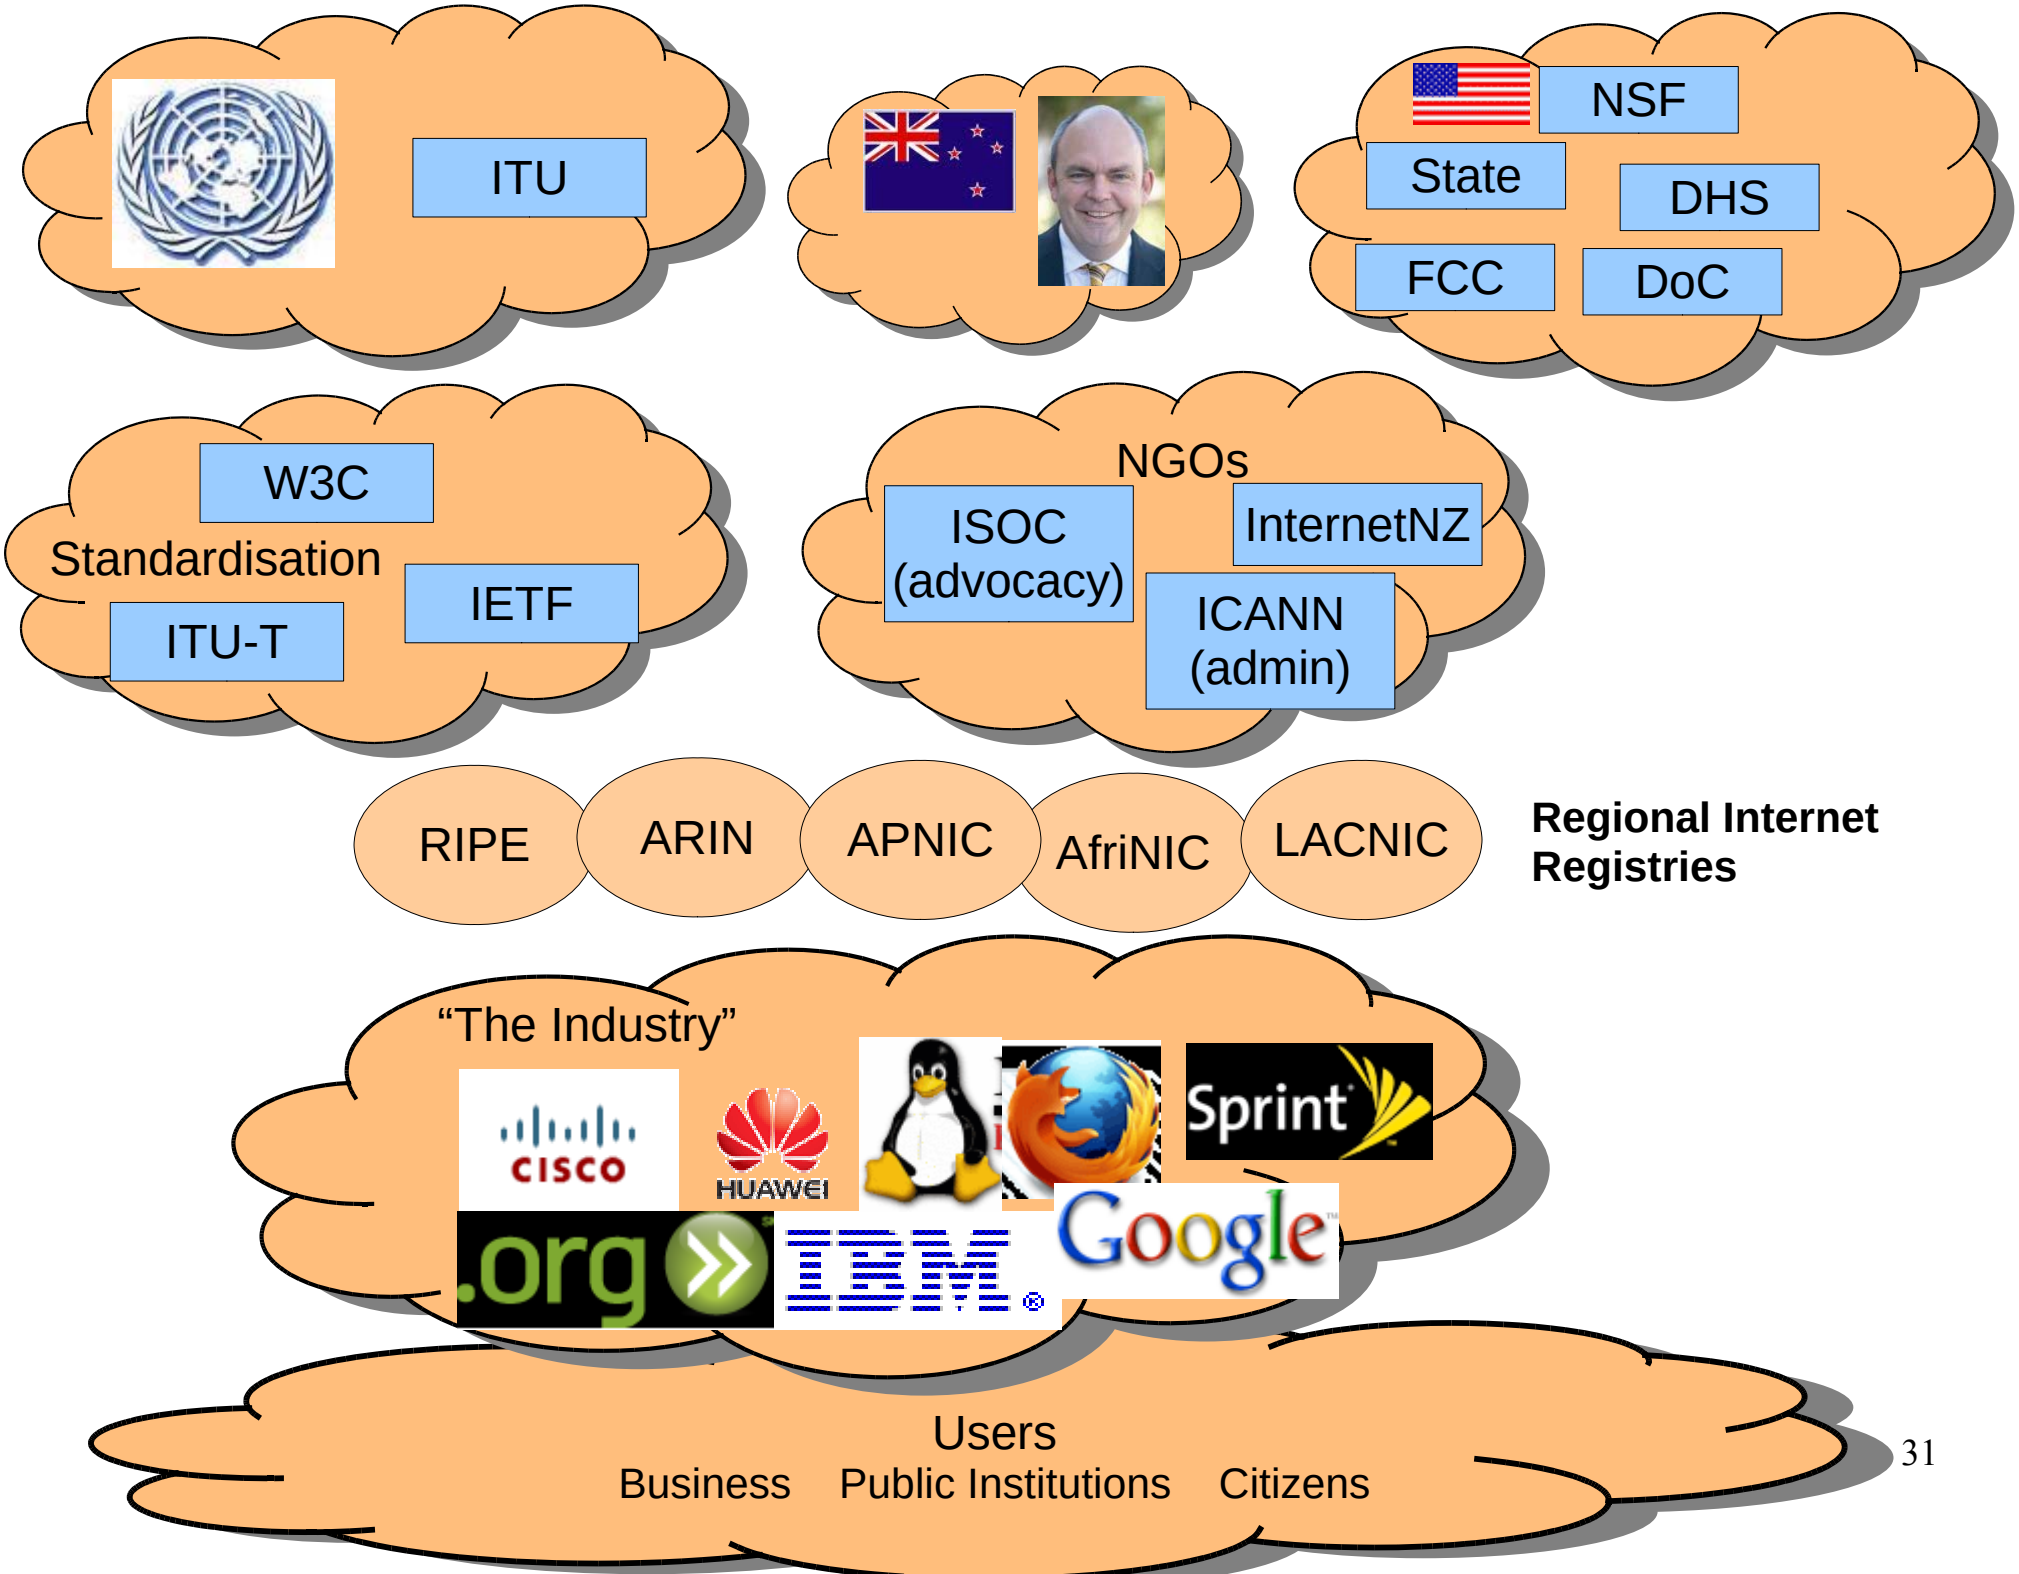
ITU
NSF
State
DHS
FCC
DoC
W3C
Standardisation
IETF
ITU-T
NGOs
ISOC
(advocacy)
InternetNZ
ICANN
(admin)
RIPE
ARIN
APNIC
AfriNIC
LACNIC
Regional Internet Registries
"The Industry"
CISCO
HUAWEI
Sprint
.org
IBM
Google
Users
Business
Public Institutions
Citizens

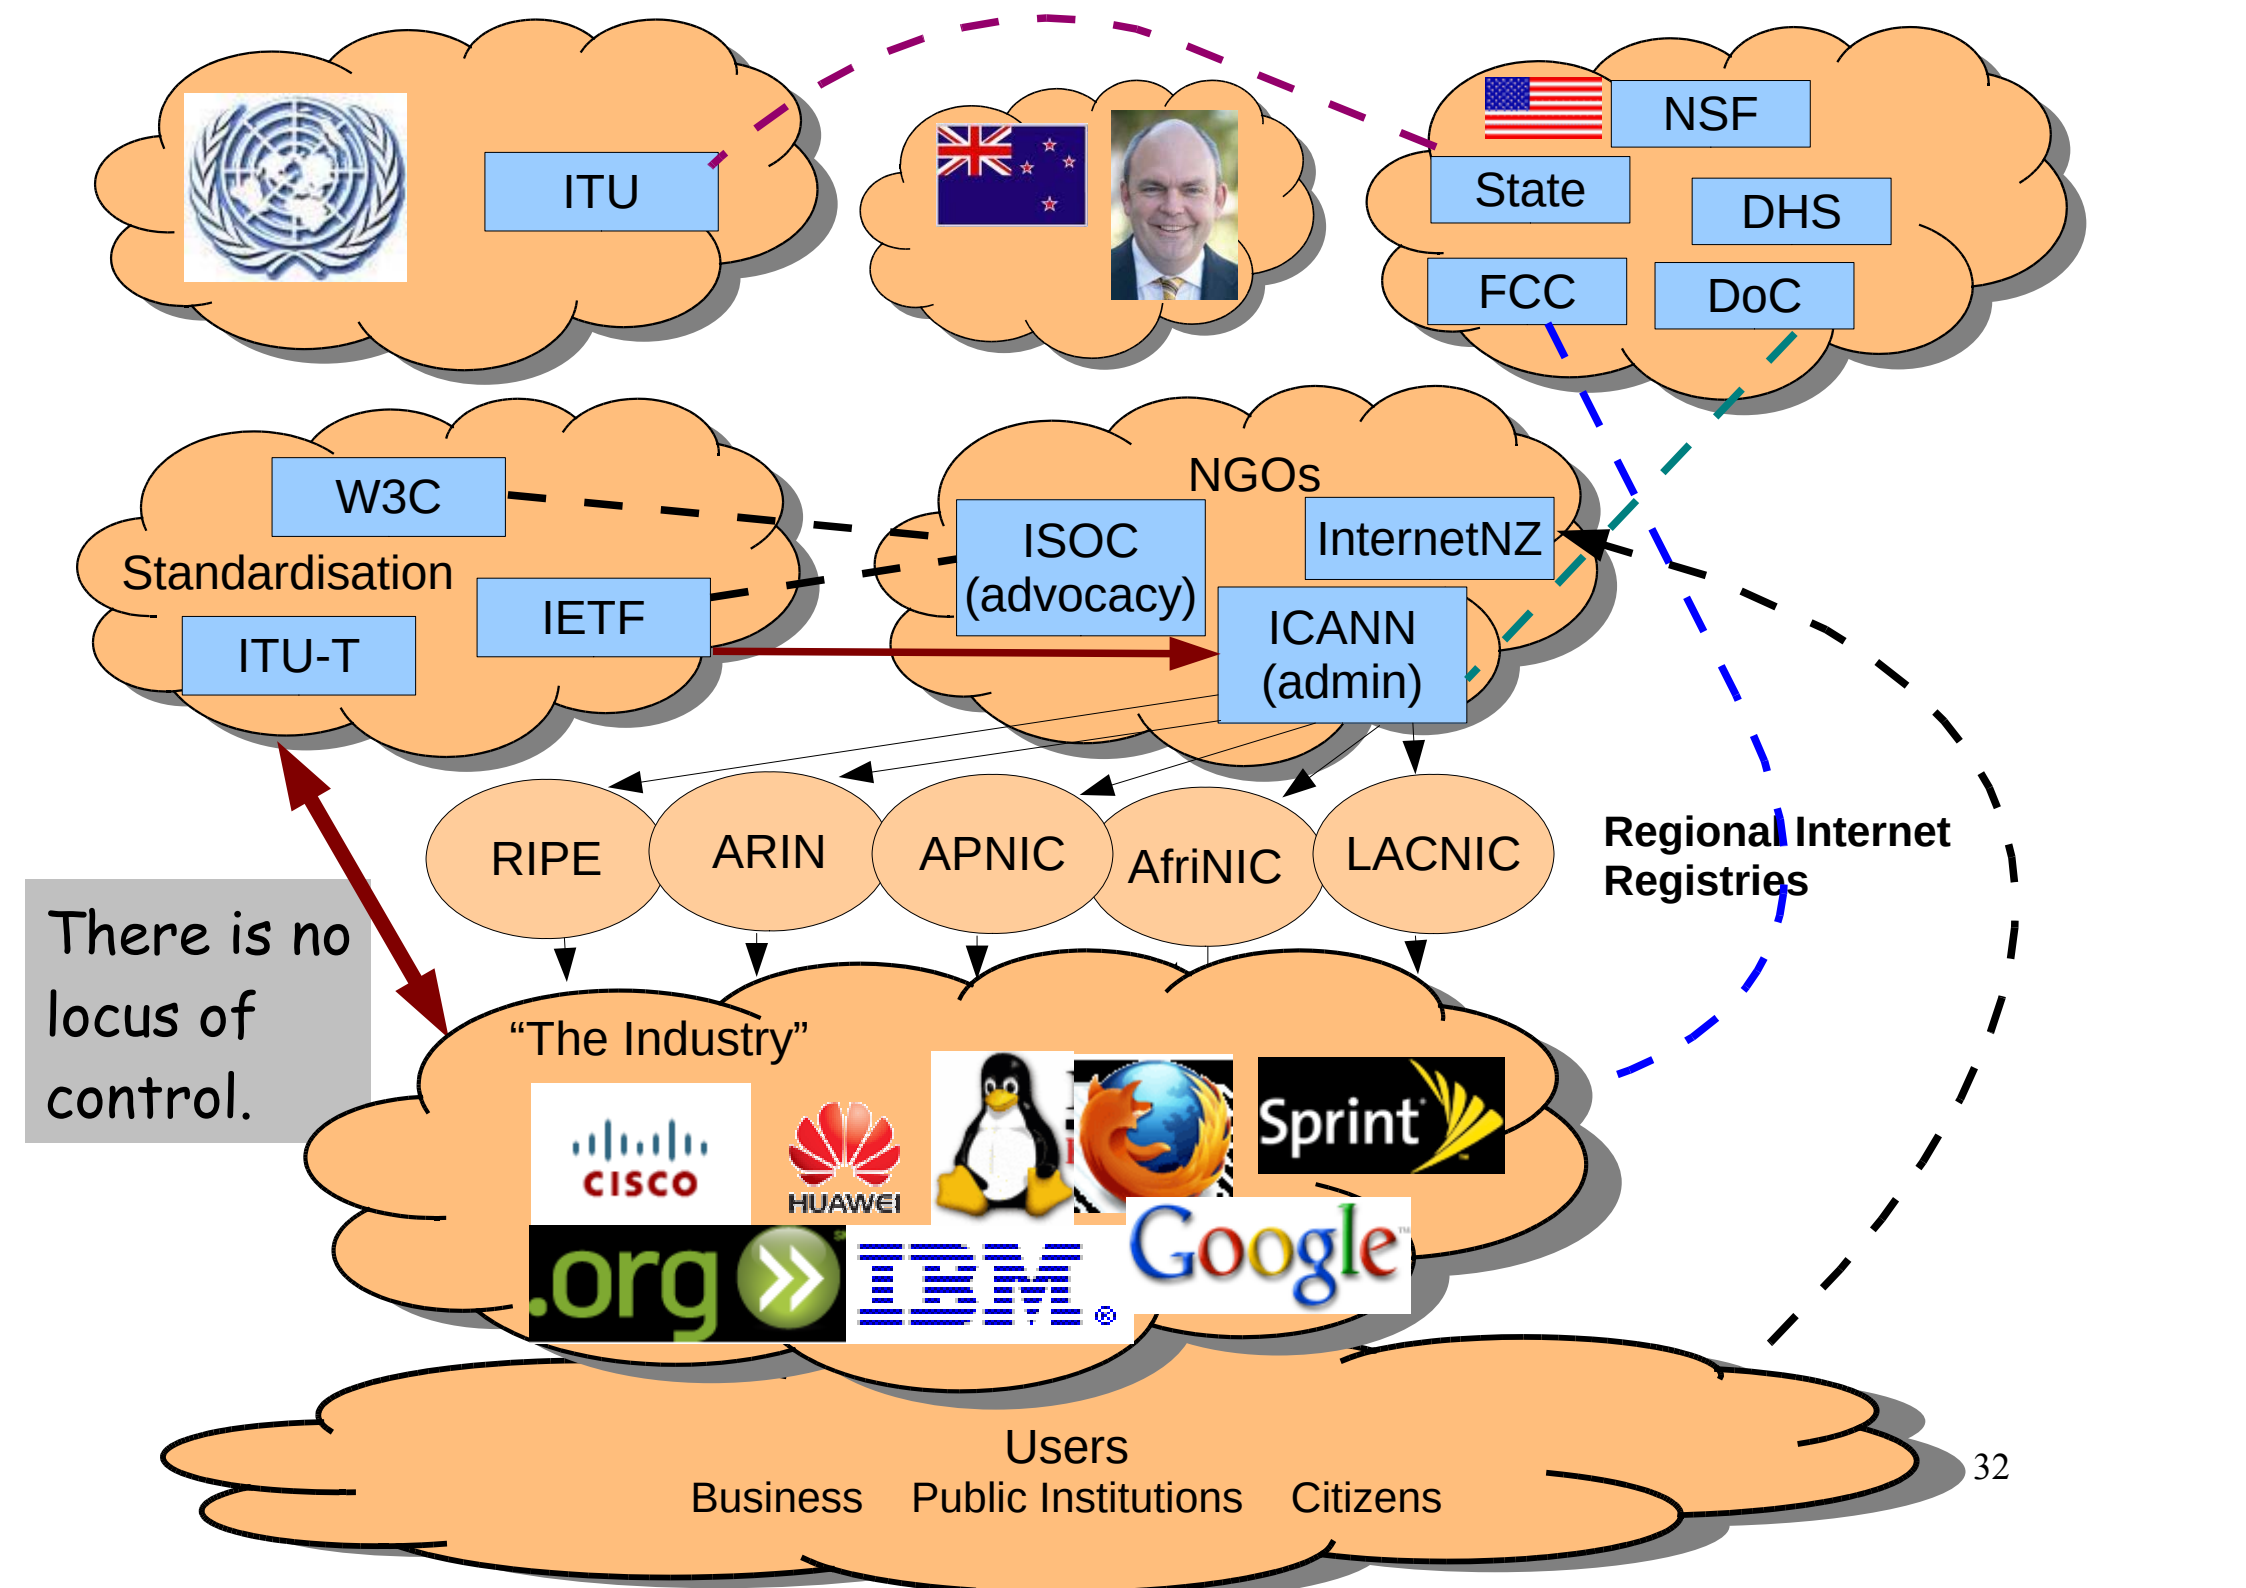
ITU
NSF
State
DHS
FCC
DoC
W3C
Standardisation
ITU-T
IETF
NGOs
ISOC
(advocacy)
InternetNZ
ICANN
(admin)
RIPE
ARIN
APNIC
AfriNIC
LACNIC
Regional Internet
Registries
There is no
locus of
control.
"The Industry"
CISCO
HUAWEI
Sprint
.org
IBM
Google
Users
Business
Public Institutions
Citizens

Wide area routing (BGP4 protocol)

- This is the “rocket science” of the Internet
  - Mechanism designed for the mainly academic network in 1988-94 now supports 680 million nodes and 35,000 autonomous routing systems.
  - Routing table has grown to about 330,000 entries
  - 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

Subtle effects of the economy

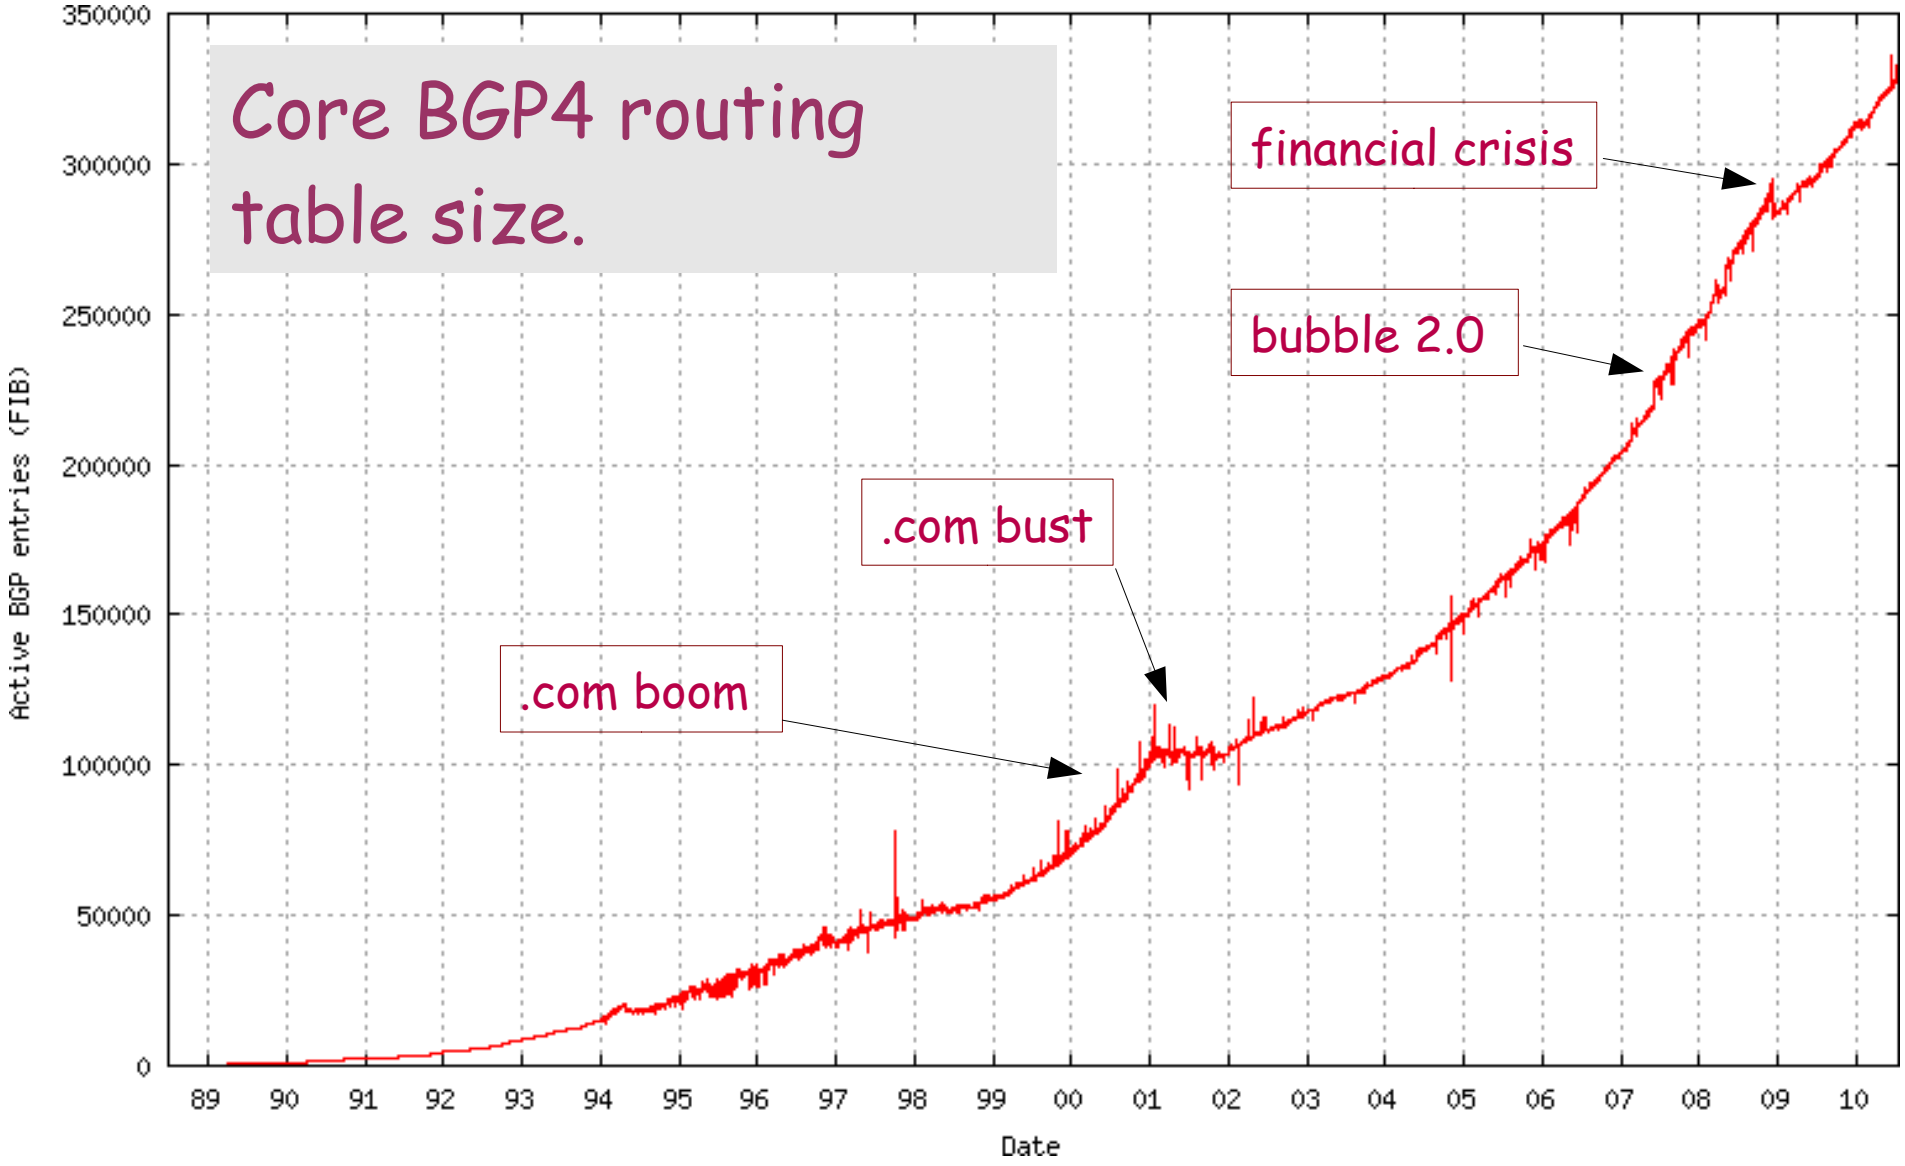

http://www.potaroo.net/

Abstract view of the BGP4 system

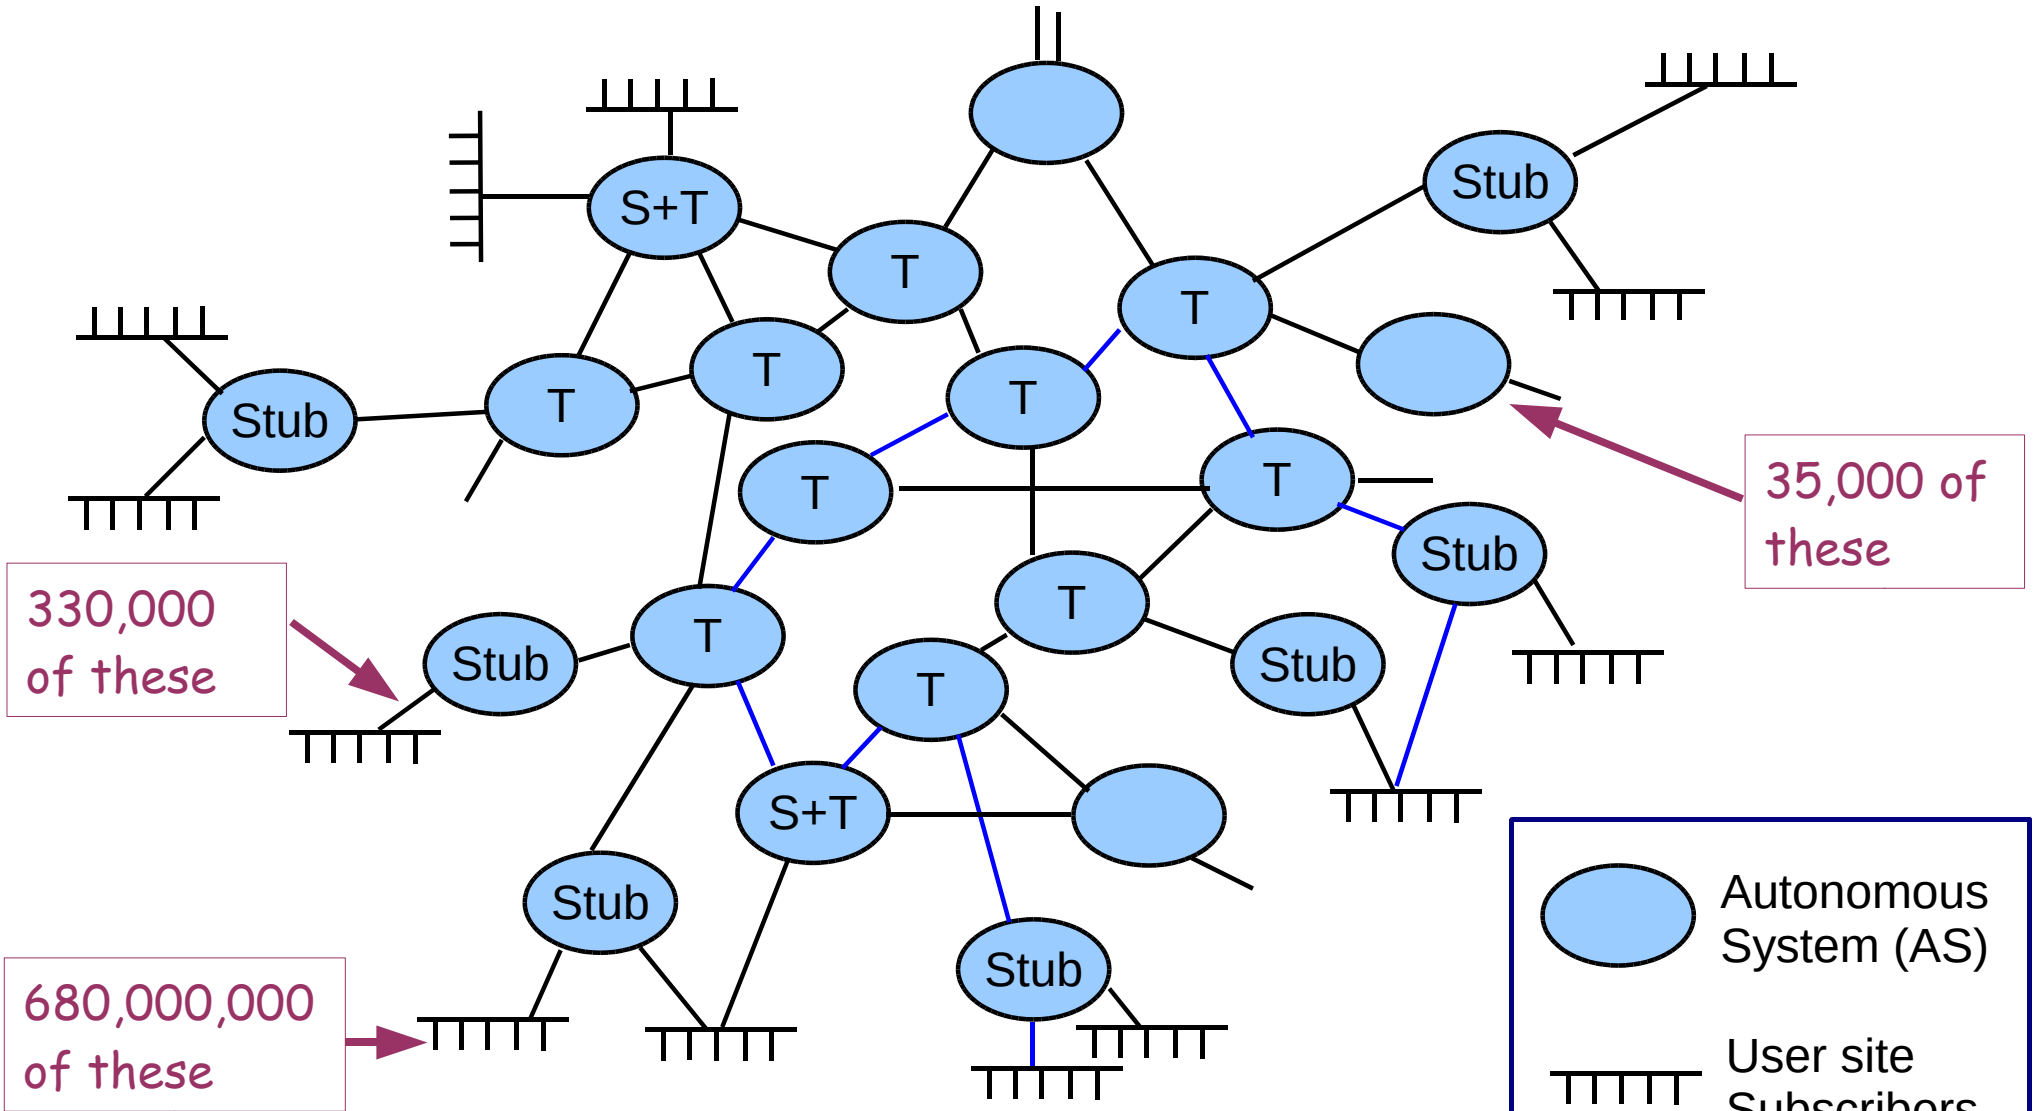

T = Transit AS; S+T = Stub + Transit AS

Autonomous System count vs “Computers on the Internet” count (1994-2009)

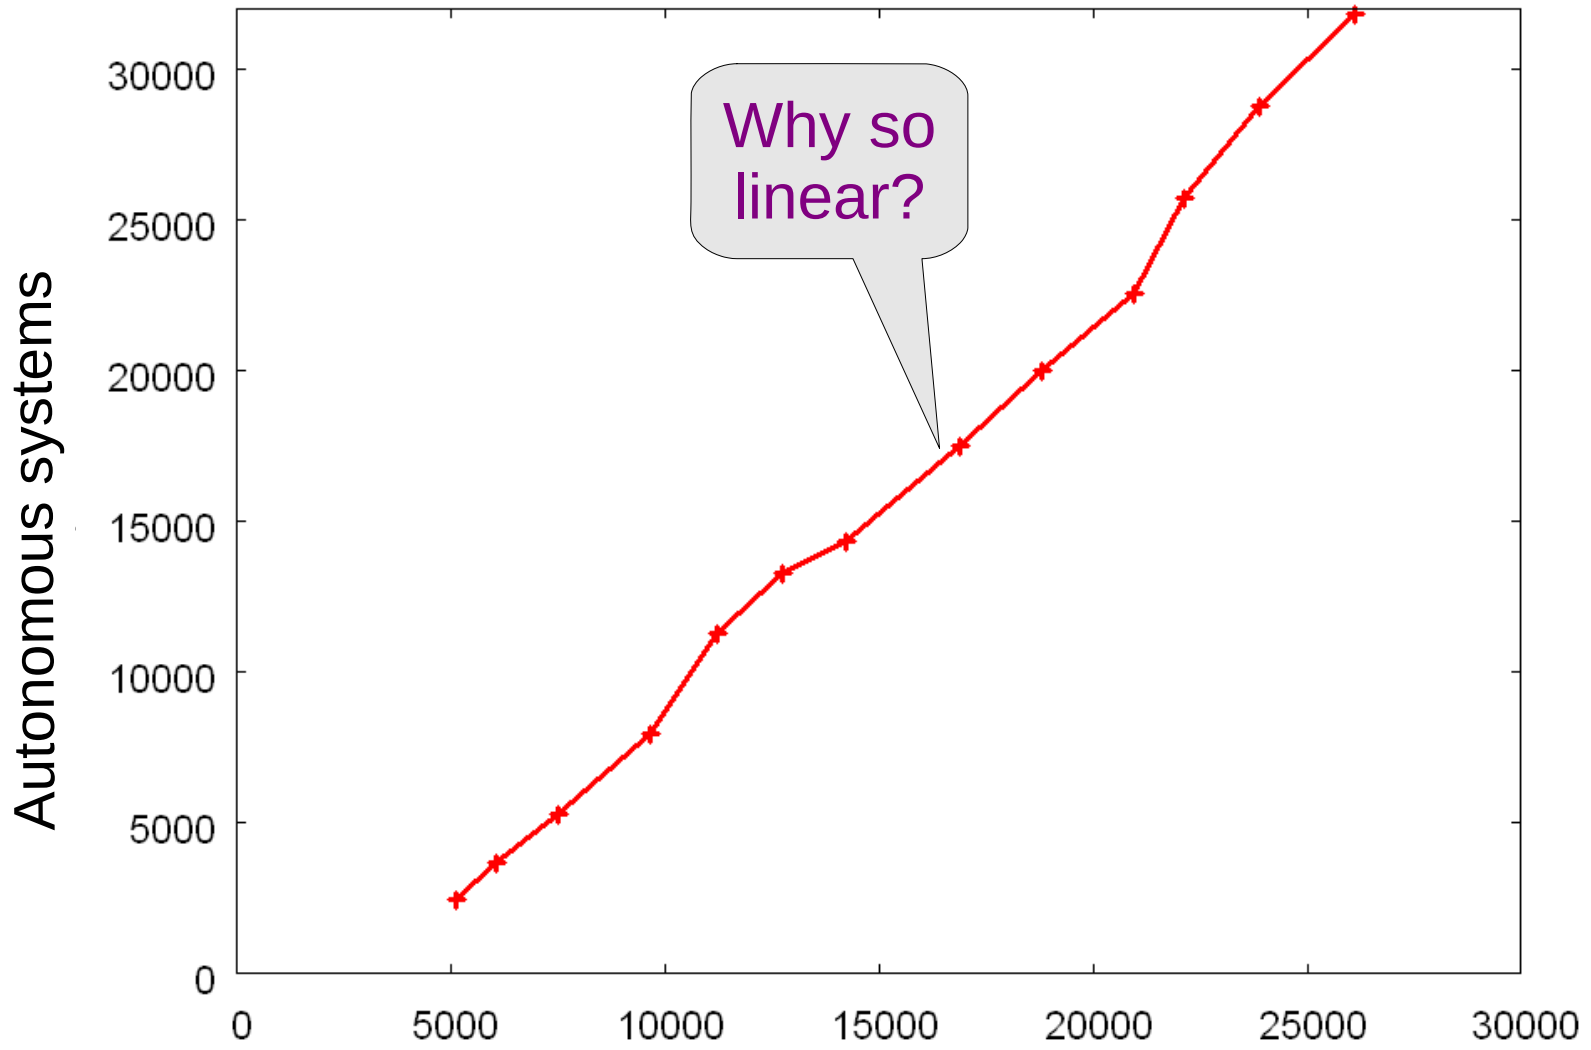

Why a square law?

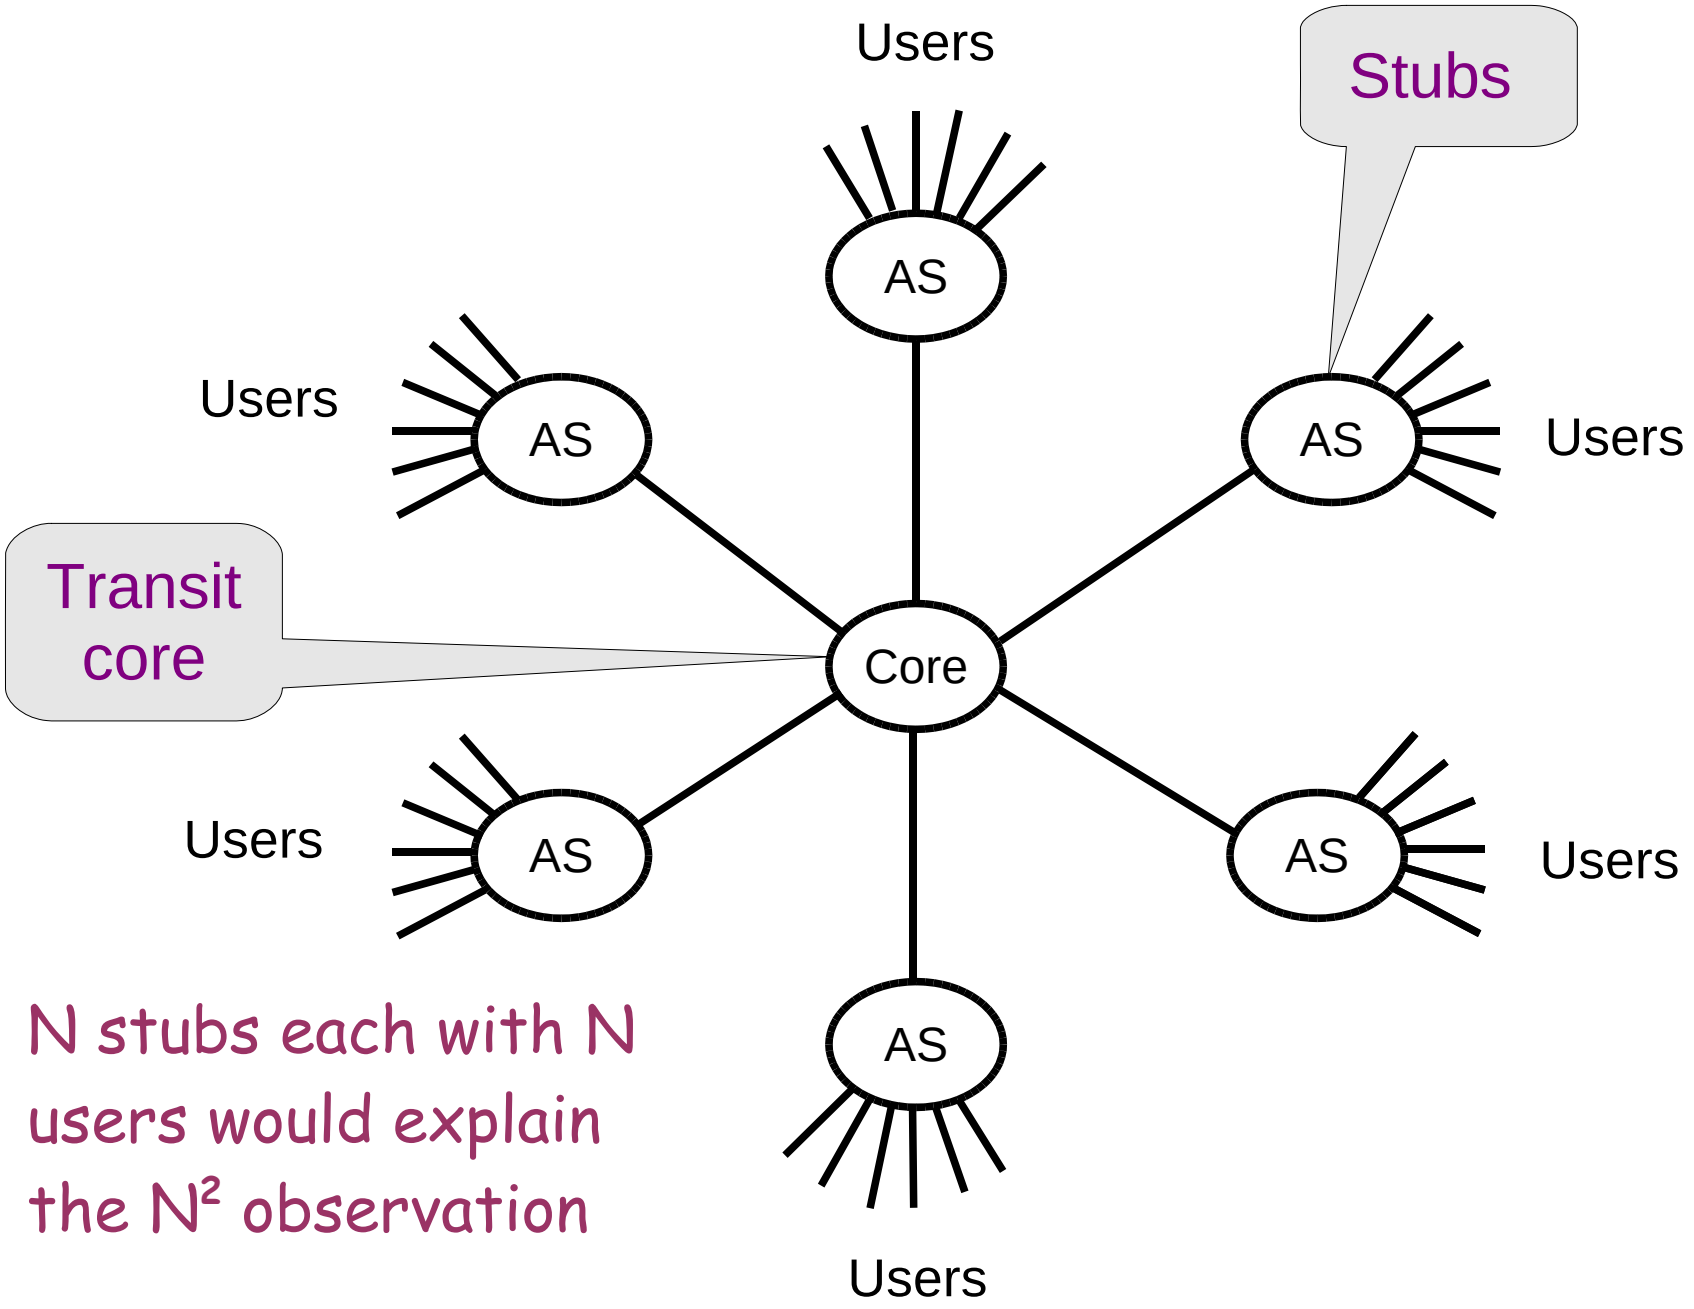

N stubs each with N users would explain the N^2 observation

Did we find the centre?

- Unfortunately not. Even if we regard the 14% of autonomous systems that provide transit to others as the core, that core is a mesh of about 5000 separate networks.
- Even at intercontinental level, there are competing transit providers.
- Also note that the “balance of power” between the major ISPs and the major content providers is constantly changing, so the topology of the network is in constant flux.

Limits to growth?

- There’s no control centre, no routing centre, and apparently endless growth. Where is that going to take us?
  1. Network capacity
  2. Energy
  3. Routing system
  4. Address space

Network capacity

- That's bits per second
  - (or megabytes per second)
  - Sometimes misnamed as "bandwidth"
- There's plenty; optical fibres are cheap.
  - ~$^{NZ}100 per capita buys a new NZ-US cable
  - But installing fibre to the home is expensive (k$ per dwelling)
- No practical limit except economics

Energy consumption

- A worry.
  - Broadband network estimated to consume 0.4% of all electricity (in developed countries). Will probably grow to 1 or 2% [Baliga et al, 2009].
- A bigger issue is the content servers.
  - Google investing directly in wind power
  - Numerous instances of data centres placed next door to power stations to reduce power loss in transmission
- Carbon credits for the Internet?
  - A telemeeting consumes less energy than a face to face meeting involving air travel
  - Email, downloads, on-line tax forms, etc., offset snail mail

Routing system “explosion”

- Since at least 1990, the routing community has been concerned that BGP growth might exceed technological limits.
- So far, Moore’s law has prevailed.
- But will that still be possible if the routing table increases by another one or two orders of magnitude?
- This is still an important research topic.
  - No imminent worries, but perhaps there is a technical limit somewhere in the future.

Address space

- The late 1970's design of the IP protocol allows for about 4 billion addresses (32 bits)
  - It seemed a lot at the time (after all, there were only 4 billion people alive then).
  - But the central registry will run out in June 2011, and ISPs expect to start running out in ~2015.
- It really isn't hard to grasp the conclusion - we have to start using larger addresses (128 bits), which means installing a new version of IP (known as IPv6)
  - That sounds easy, but it ultimately affects every computer in the world.
  - If you're running any kind of IT service, today is the day to start planning for this upgrade.

Internet cosmology

- We observed growth over 40 years from a small ball of energy (the ARPANET pioneers).
- We didn’t find the centre, and there’s a lot of anisotropy.
- We certainly didn’t find anyone in charge.
- We don’t know of any definite limits to growth.
- The Internet is surely Aristotelian by being rooted on Earth, but it definitely feels as though its Big Bang isn’t over yet.