The Internet has come of age:

or has it?

Matthew Roughan
<matthew.roughan@adelaide.edu.au>

Discipline of Applied Mathematics
School of Mathematical Sciences
University of Adelaide

November 21, 2006

THE UNIVERSITY OF ADELAIDE AUSTRALIA

Douglas Adams on technology

1. everything that's already in the world when you're born is just normal;
2. anything that gets invented between then and before you turn thirty is incredibly exciting and creative and with any luck you can make a career out of it;
3. anything that gets invented after you're thirty is against the natural order of things and the beginning of the end of civilisation as we know it until it's been around for about ten years when it gradually turns out to be alright really.

- Give the start date of the Internet as 1997
  - the year my Mum heard of the Internet
- by Adams' argument, the Internet will come of age next year!
  - should we start planning the party?

Me?

THE UNIVERSITY OF ADELAIDE AUSTRALIA

So what am I doing here talking to you about the Internet?

- worked for AT&T in the US
  - arguably ran the largest single chunk of what we call the Internet
    http://www.att.com/news/2003/12/11-12627
- I spent my time trying to measure the Internet
  - traffic
  - performance
  - routing/topology

An outline

How do we decide if the Internet has come of age

- size?
- stability?
- maturity?

Oh, and somewhere in here, I'd better tell what the Internet is.

Does size matter?

THE UNIVERSITY OF ADELAIDE AUSTRALIA

Customer: I've been enjoying the Internet, and I wanted to show my Mum. Can I buy a copy of the Internet, please?

Sales: No, it won't fit on a disk.

Customer: Oh! How many disks will I need then?

Sales: about 200,000

Customer: (strange strangling sound)

Traffic volumes

- April 2003, in one day AT&T carried 1 petabyte of traffic a day in North America alone
- 1 petabyte
  - 10^{15} bytes
  - 2000 $\times$ 500 Gbyte disks
  - 180 $\times$ the US library of congress
  - estimated 2 petabytes as the contents of all U.S. academic research libraries

How big is it?

THE UNIVERSITY OF ADELAIDE AUSTRALIA

- Internet size estimate
  - 1 petabytes

  http://www.circuitcellar.com/library/priorityinterrupt/194.htm
- Google reported to have
  - 4 petabytes memory

  http://glinden.blogspot.com/2006/06/four-petabytes-in-memory.html
- Wayback Machine (archive of the WWW)
  - 2 petabytes

  http://www.archive.org/about/faqs.php

Dollars and sense

- Google+Yahoo!+eBay+Yahoo!Japan + Amazon.com
  - pre-2000: $2B
  - Nasdaq peak (2000), $178B
  - Nasdaq trough (2002), $32B
  - Nov.2006, $259B
- 8% of total US advertising online
  - estimate 13% in 5 years

Morgan-Stanley, Nov 2006

http://www.morganstanley.com/institutional/techresearch/webtwopto2006.html

Outline

THE UNIVERSITY OF ADELAIDE AUSTRALIA

How do we decide if the Internet has come of age

- size → clearly the Internet is big enough
- **stability?**
- maturity?

Early Internet Growth

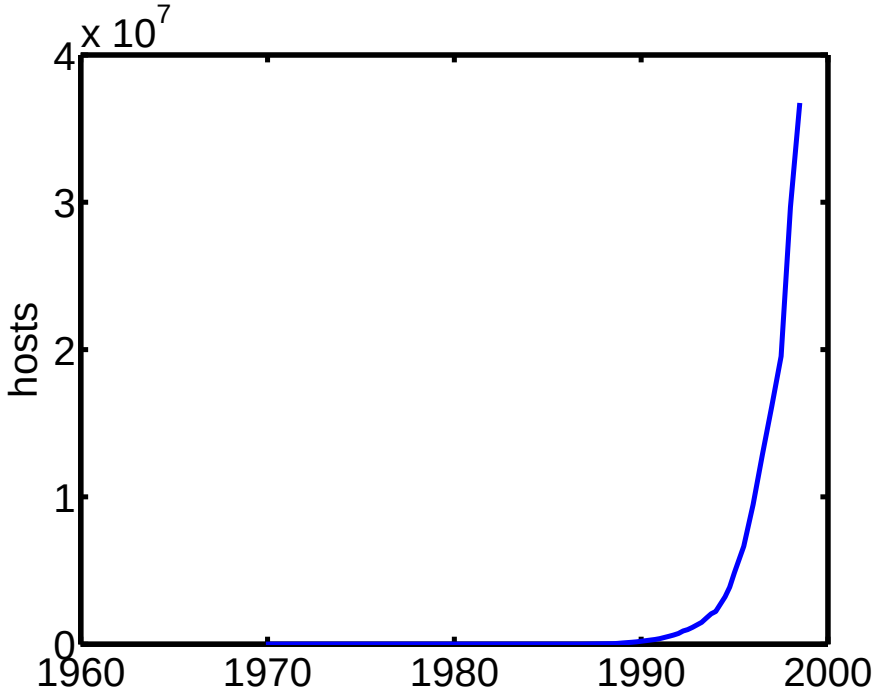

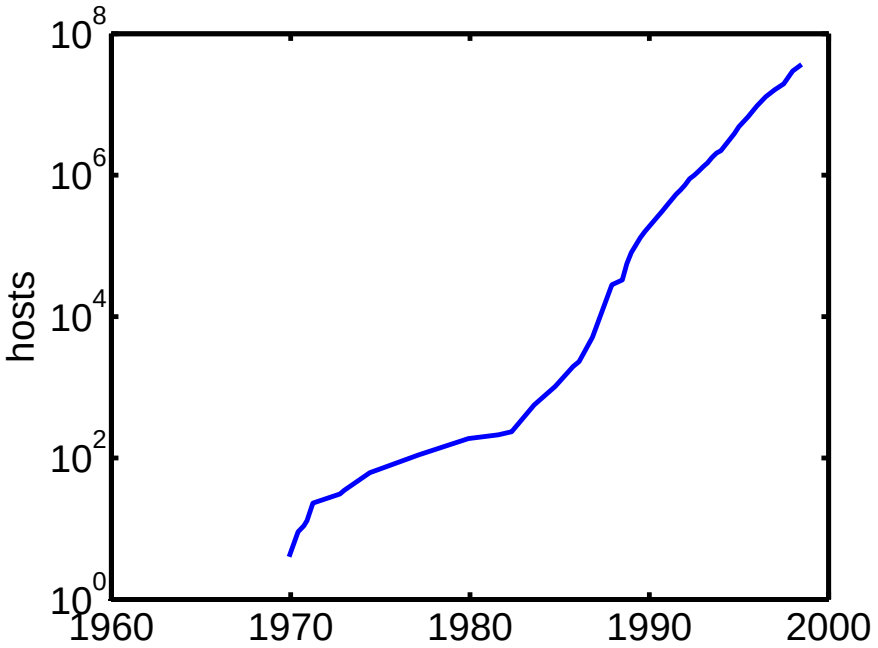

http://www.zakon.org/robert/internet/timeline/#1990s

Date (mm/yy)	hosts
08/1981	213
01/1992	727,000
01/1997	19,540,000

Internet Traffic Growth

Traffic roughly doubles every year [10].

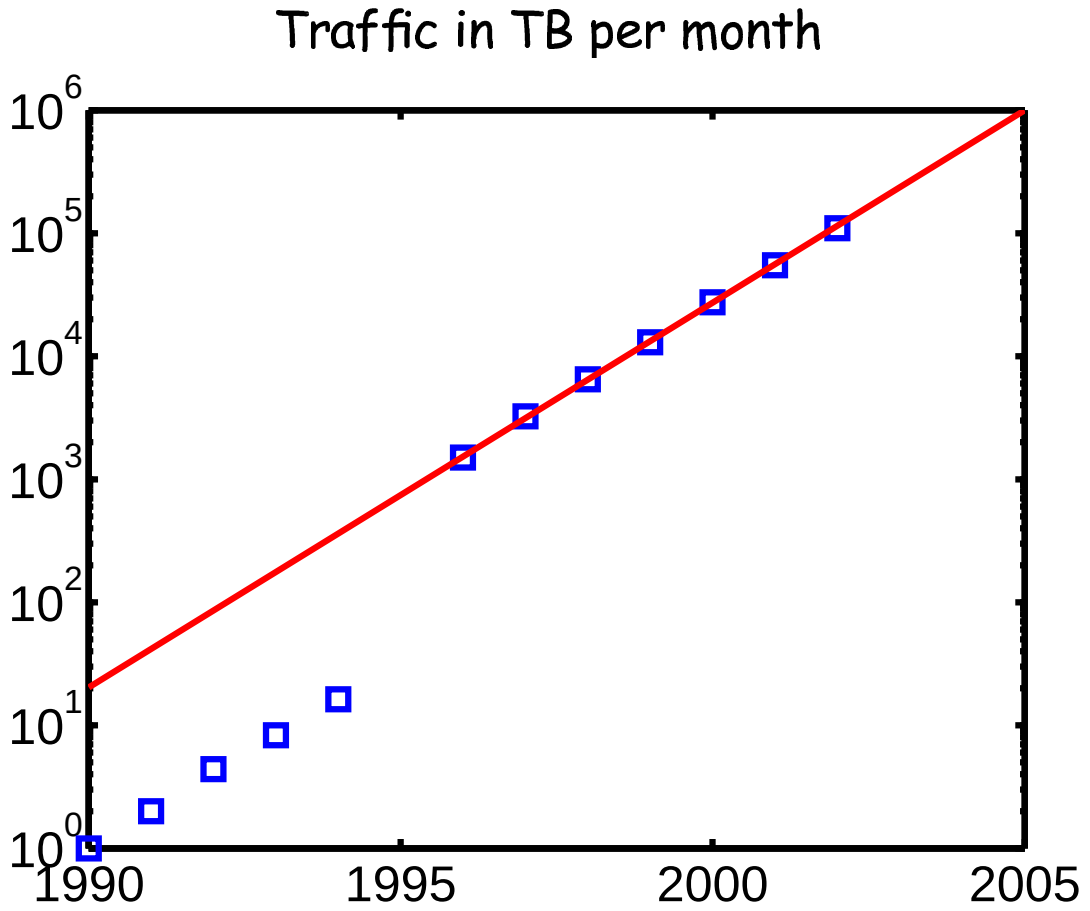

Combination of new users and higher bandwidth!

Drivers of growth

- Moore's Law
- Gilders Law (and Nielson's)
- Metcalfe's Law

Moore's Law

Moore's law: the speed of digital hardware increases by a factor of two every 18 months, or the number of transistors on a chip doubles, or the cost halves [4].

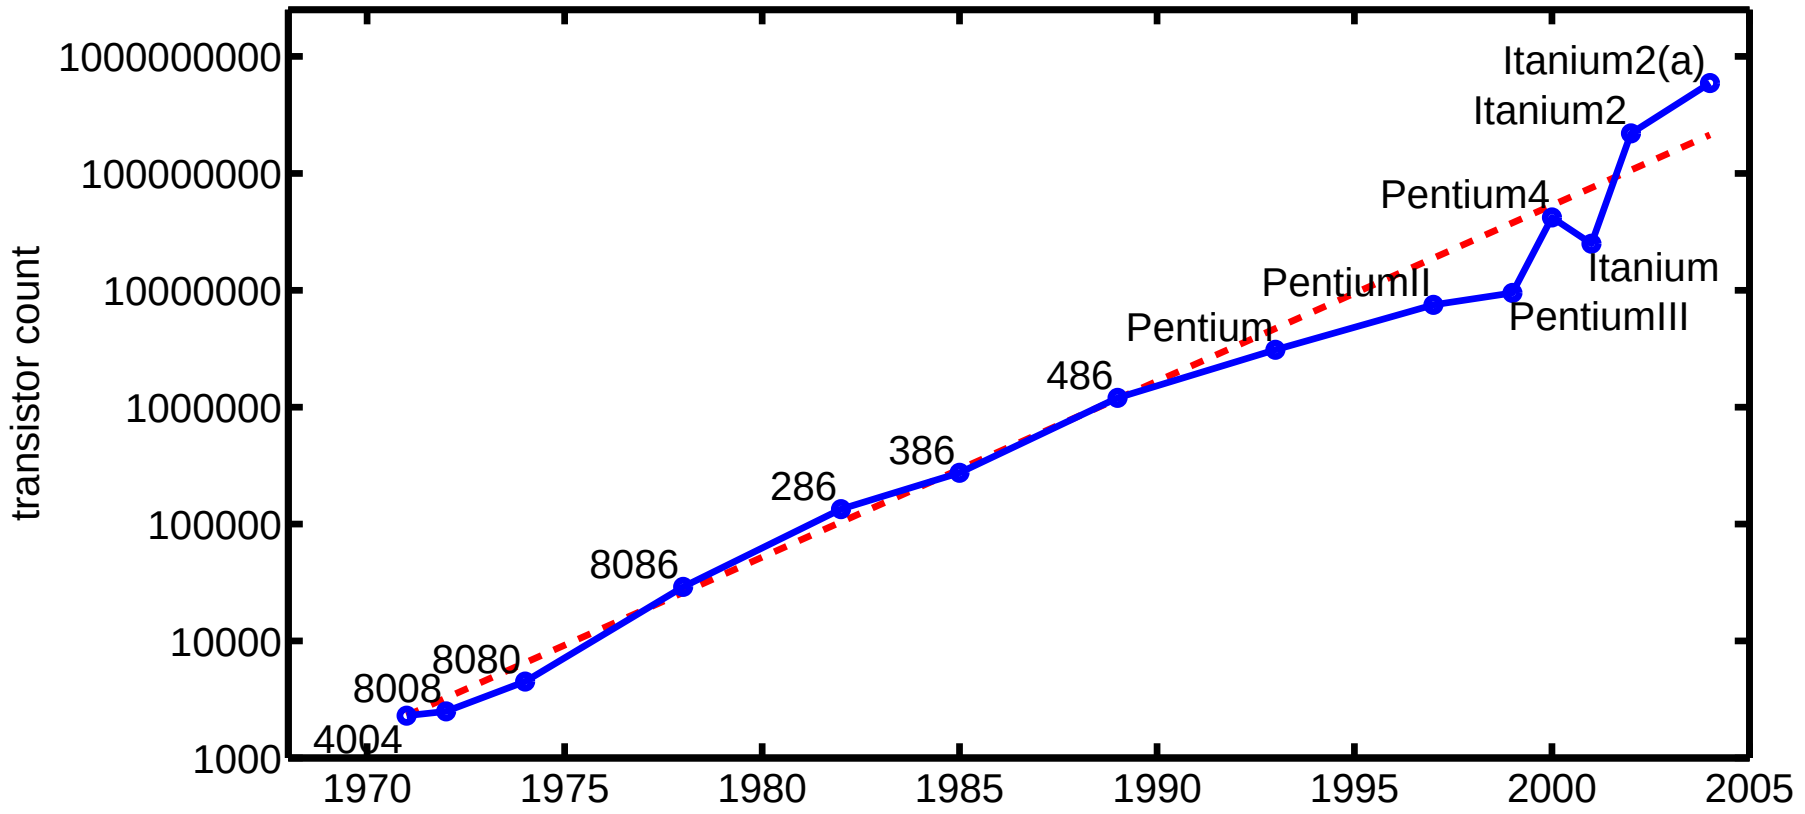

Actually looks more like a factor of 2 every 2 years.

Gilder's Law

Gilder's law: theoretical transmission capacity of a link increases by a factor of two every 12 months.

- backbone link capacity approx. doubling in 2 years
- access links likewise (Nielsen's Law)

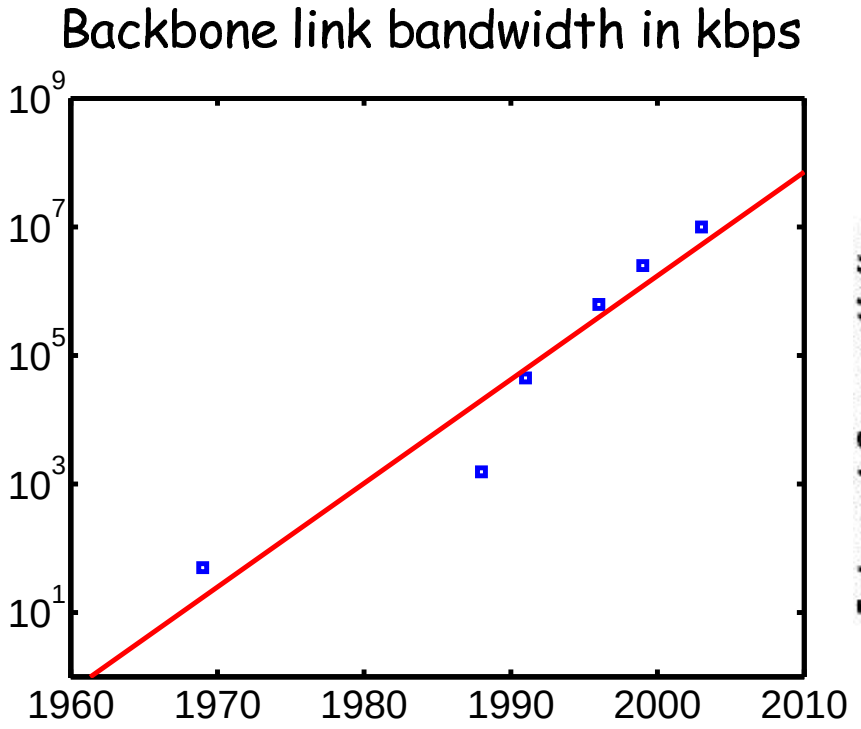

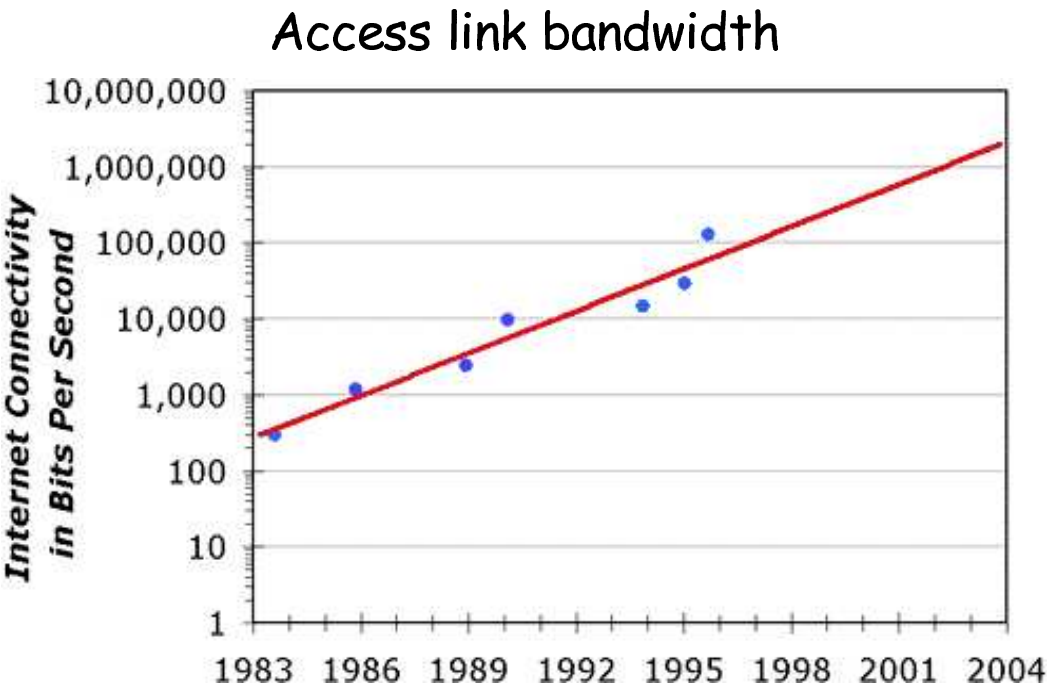

Metcalfe's law

- The value of a network goes as $O(N^2)$
- Metcalfe's law also drives the Internet
  - hence the failure of many "video-phone" trials
    - but success of most recent "camera phones"
- Actually probably more like $O(N \log N)$
  http://spectrum.ieee.org/jul06/4109

Exponential growth

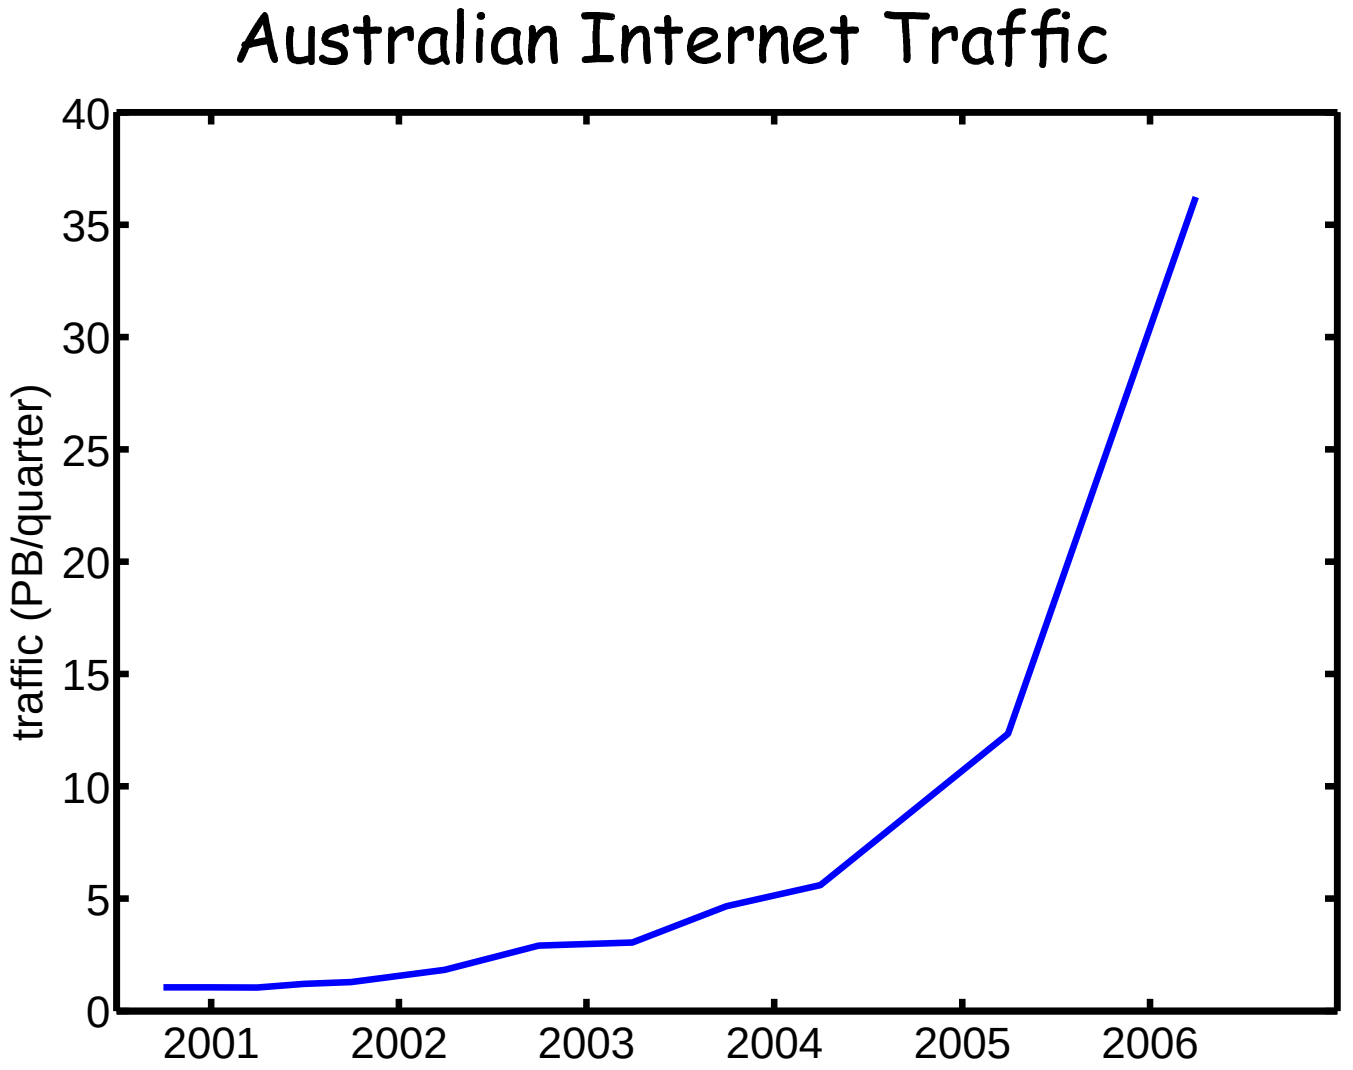

source www.abs.gov.au

Exponential Growth

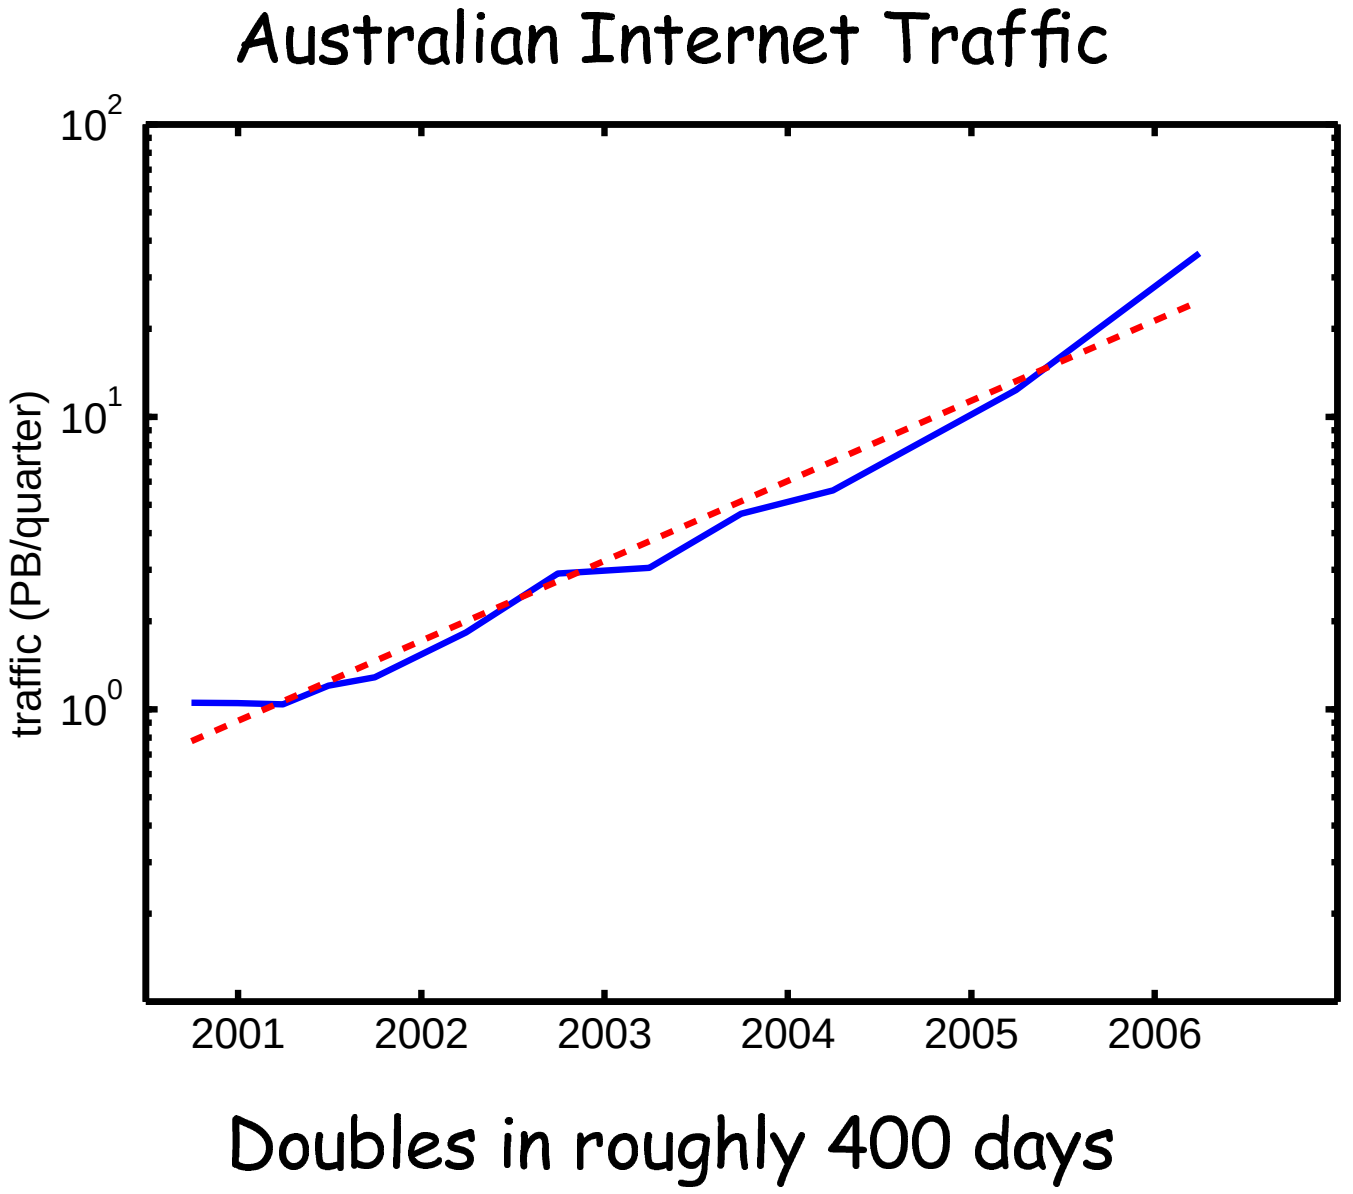

Will it ever stabilise?

Moore's law failure predictions have always failed

But

- the number of users is finite
- the amount of time they can spend on the web is finite
- so growth should at least slow to growth of access line speeds?
- maybe it will even drop back to linear growth?
  - most other technologies saturate the market at some point
- maybe it will still grow?
  - machine to machine traffic

Outline

THE UNIVERSITY OF ADELAIDE AUSTRALIA

How do we decide if the Internet has come of age

- size → clearly the Internet is big enough
- stability → clearly not!
- **maturity?**

Maturity

THE UNIVERSITY OF ADELAIDE AUSTRALIA

Lets have a brief history of networking

1. pre-industrial
2. 19th century
3. early 20th century
4. computer networks
5. early 21st century (now)

More detailed telephony timelines can be found at
http://www.telephonetribute.com/timeline.html
http://www2.fht-esslingen.de/telehistory/

Histories of computing and computer networks
http://en.wikipedia.org/wiki/Computing_timeline
http://www.isoc.org/internet/history/

Australian telecoms history
http://www.caslon.com.au/timeline.htm
http://www.anu.edu.au/people/Roger.Clarke/II/OzIHist.html

Pre-industrial

- Jungle drums
- Signal fires
  1184 BC, fall of Troy [1]
  1588 AD, Arrival of Spanish Armada
- Carrier pigeons
  700 BC, Olympic games
- Smoke signals
  150 AD, Romans
- Semaphore
  1791 AD, Chappe brothers
  later used by Napoleon

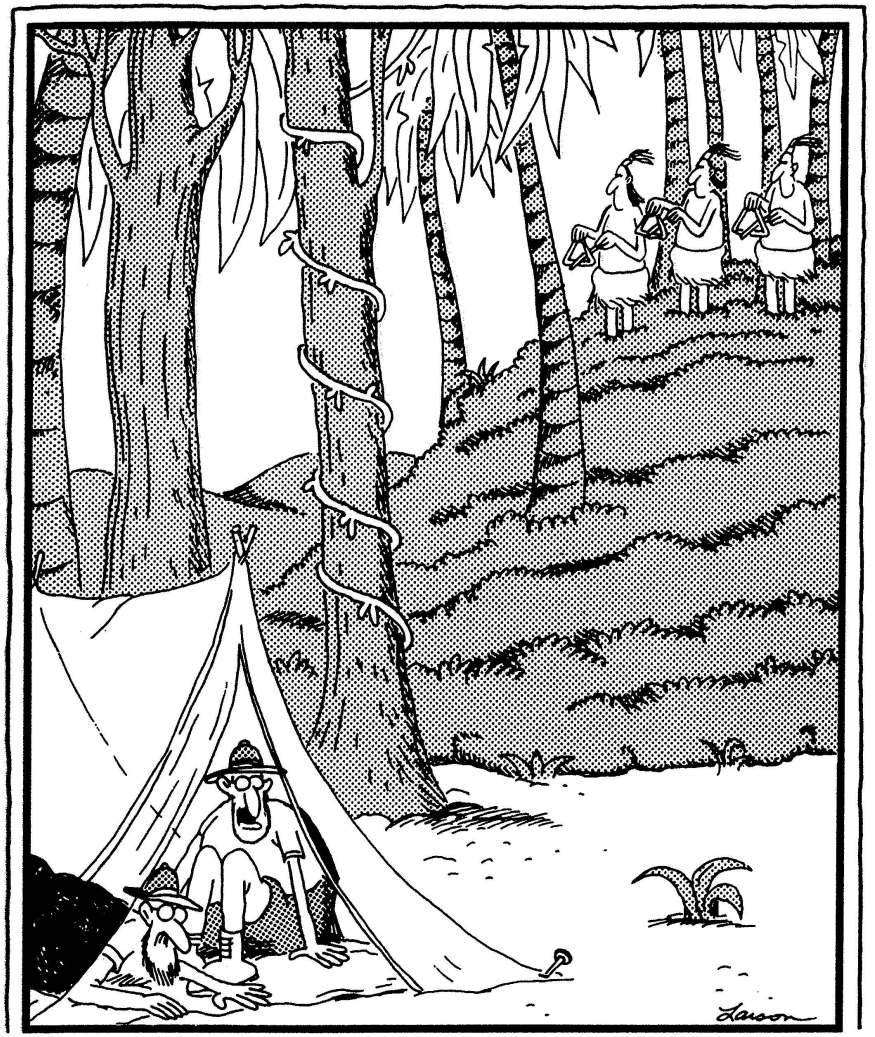

"Wait, Morrison! … It's OK—those are jungle *triangles*!"

Gary Larson, 1993

Pre-industrial

THE UNIVERSITY OF ADELAIDE AUSTRALIA

These had limitations

- Carrier pigeons: `1 short message per pigeon`
- Signal fires: `one bit per fire`
- Semaphore: `15 characters per minute.`

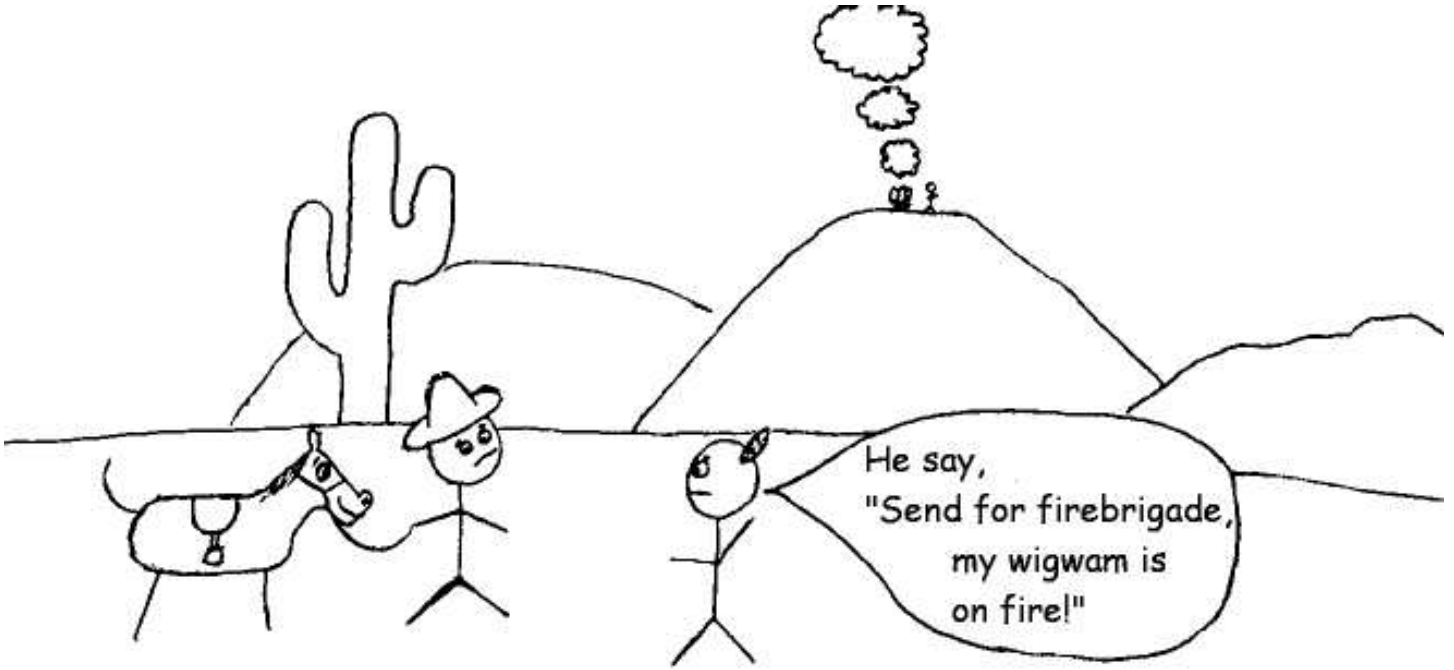

19th century

Electronic communication:

- telegraph
  - invented 1753
  - Morse code 1835
  - take off 1838
  - 1st transatlantic line 1866
- radio (Marconi, 1896)
- telephone
  - A.G. Bell
  - filed patent Feb. 14, 1876, 3 hours before Elisha Gray

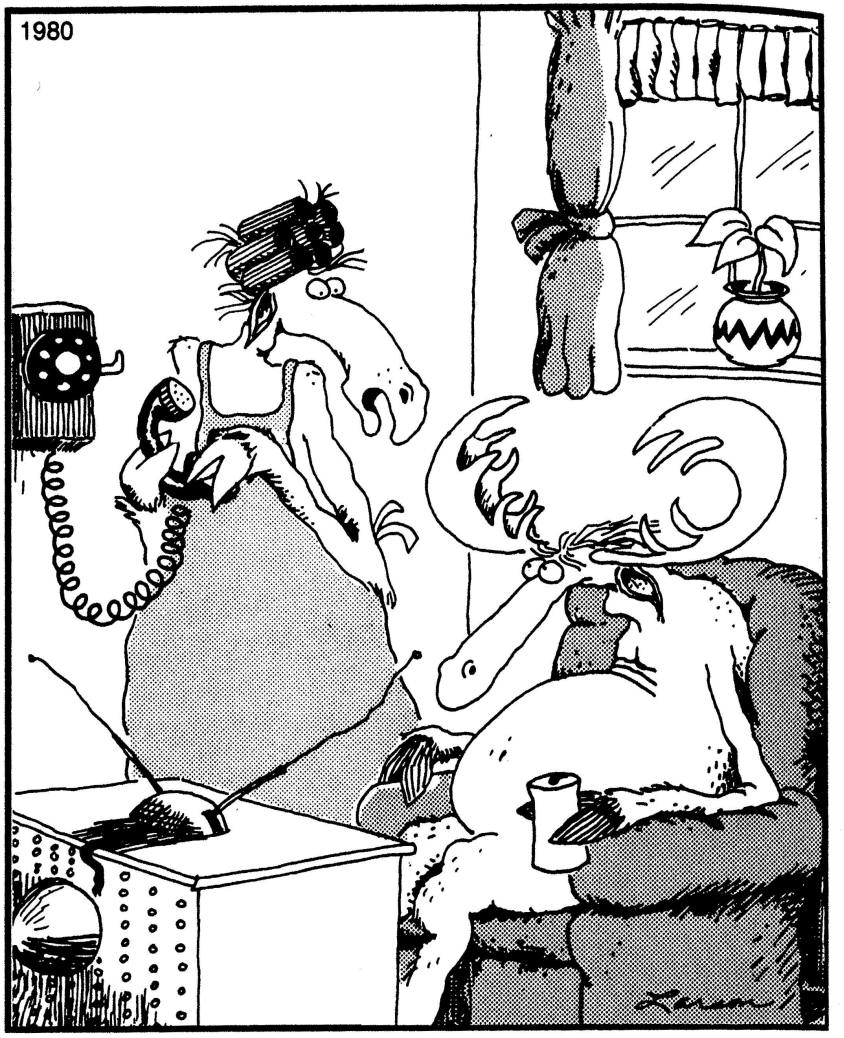

"It's the call of the wild."

Gary Larson, 1980

19th century

Simple telephone: connects two points with a wire

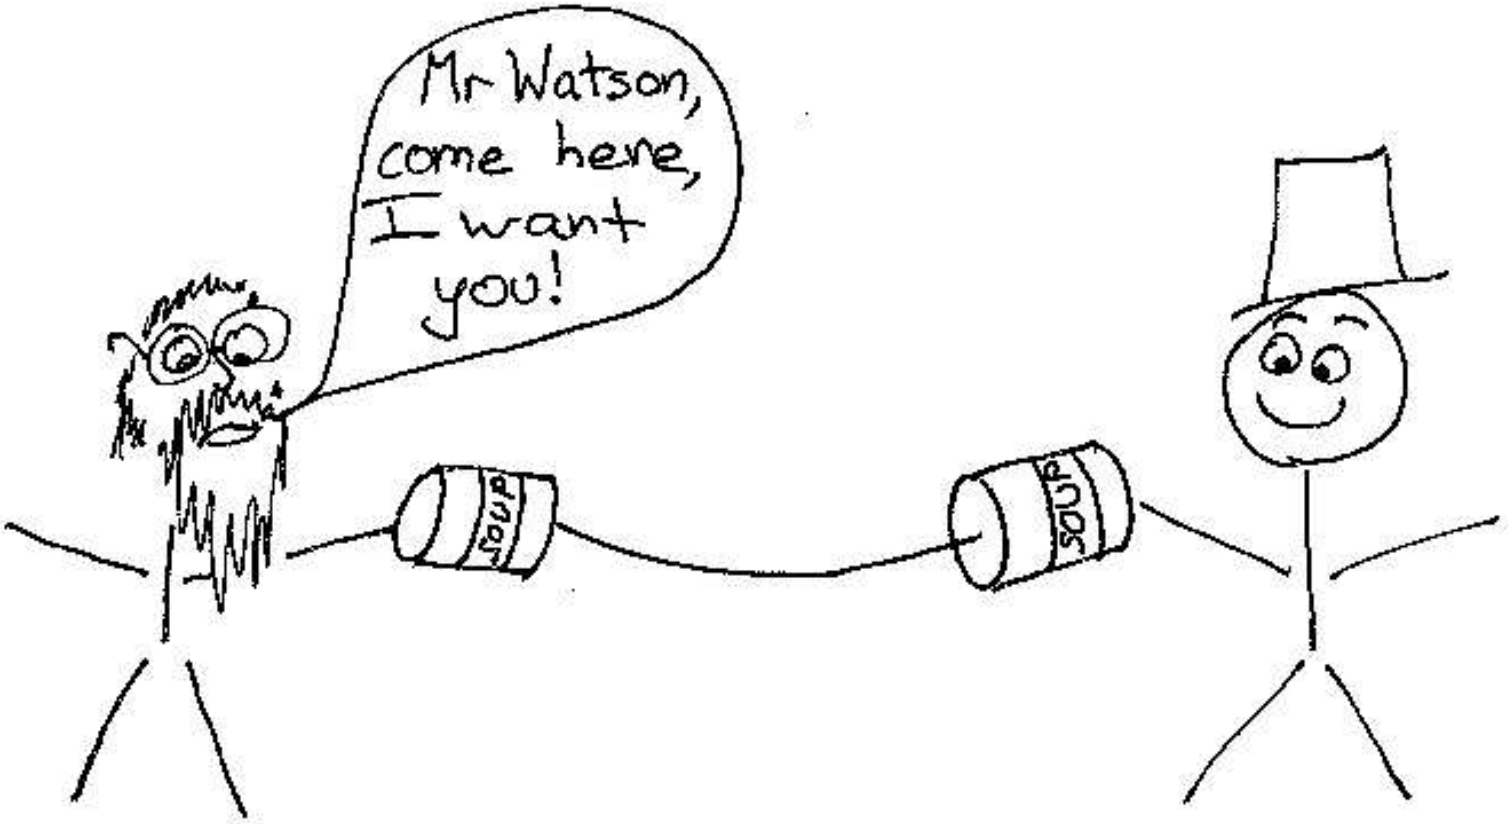

Reportedly, the first words over the telephone came when Bell spilled some acid on his pants, whereupon he call "Mr. Watson, come here, I want you!"

Dumb network design

One link between every pair who wish to speak

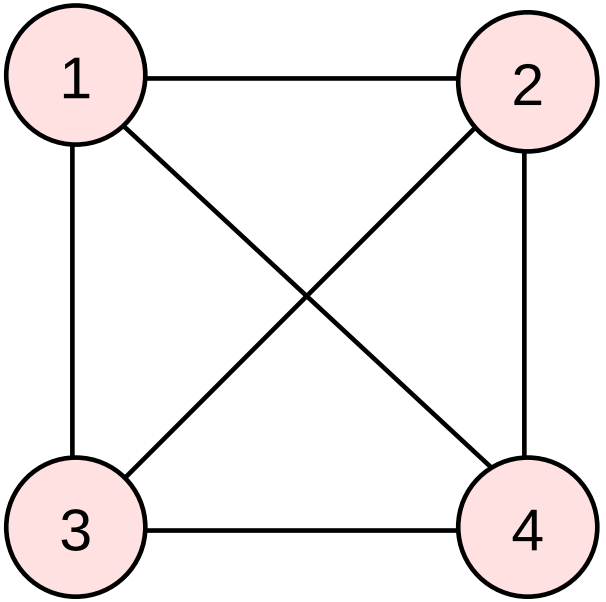

N=4
L=6

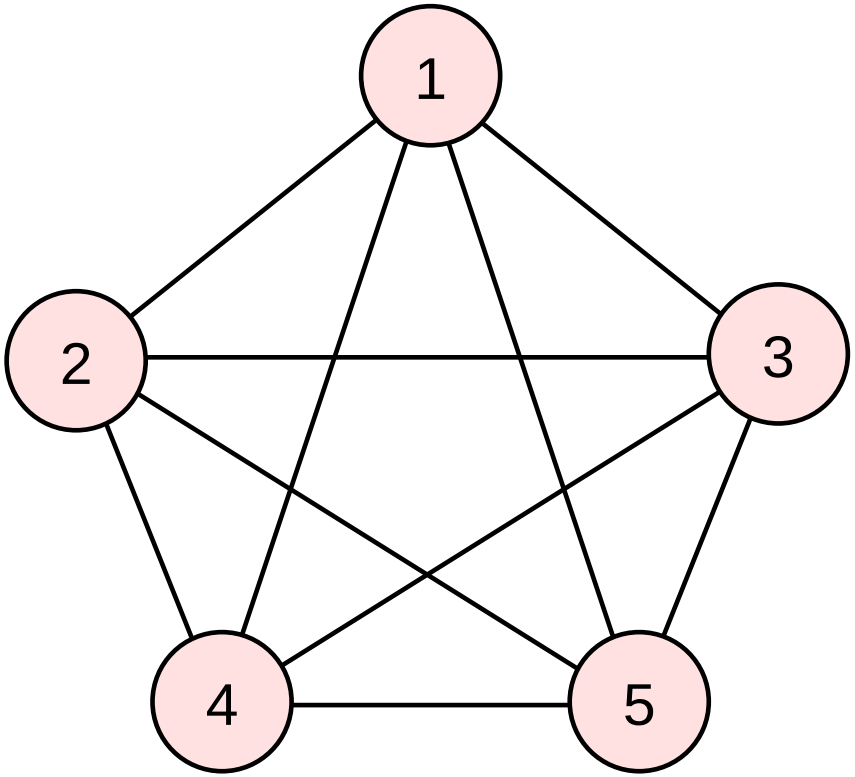

N=5
L=9

N nodes, then we have L=N(N-1)/2 links

Dumb network design

THE UNIVERSITY OF ADELAIDE AUSTRALIA

Pratt, Kansas

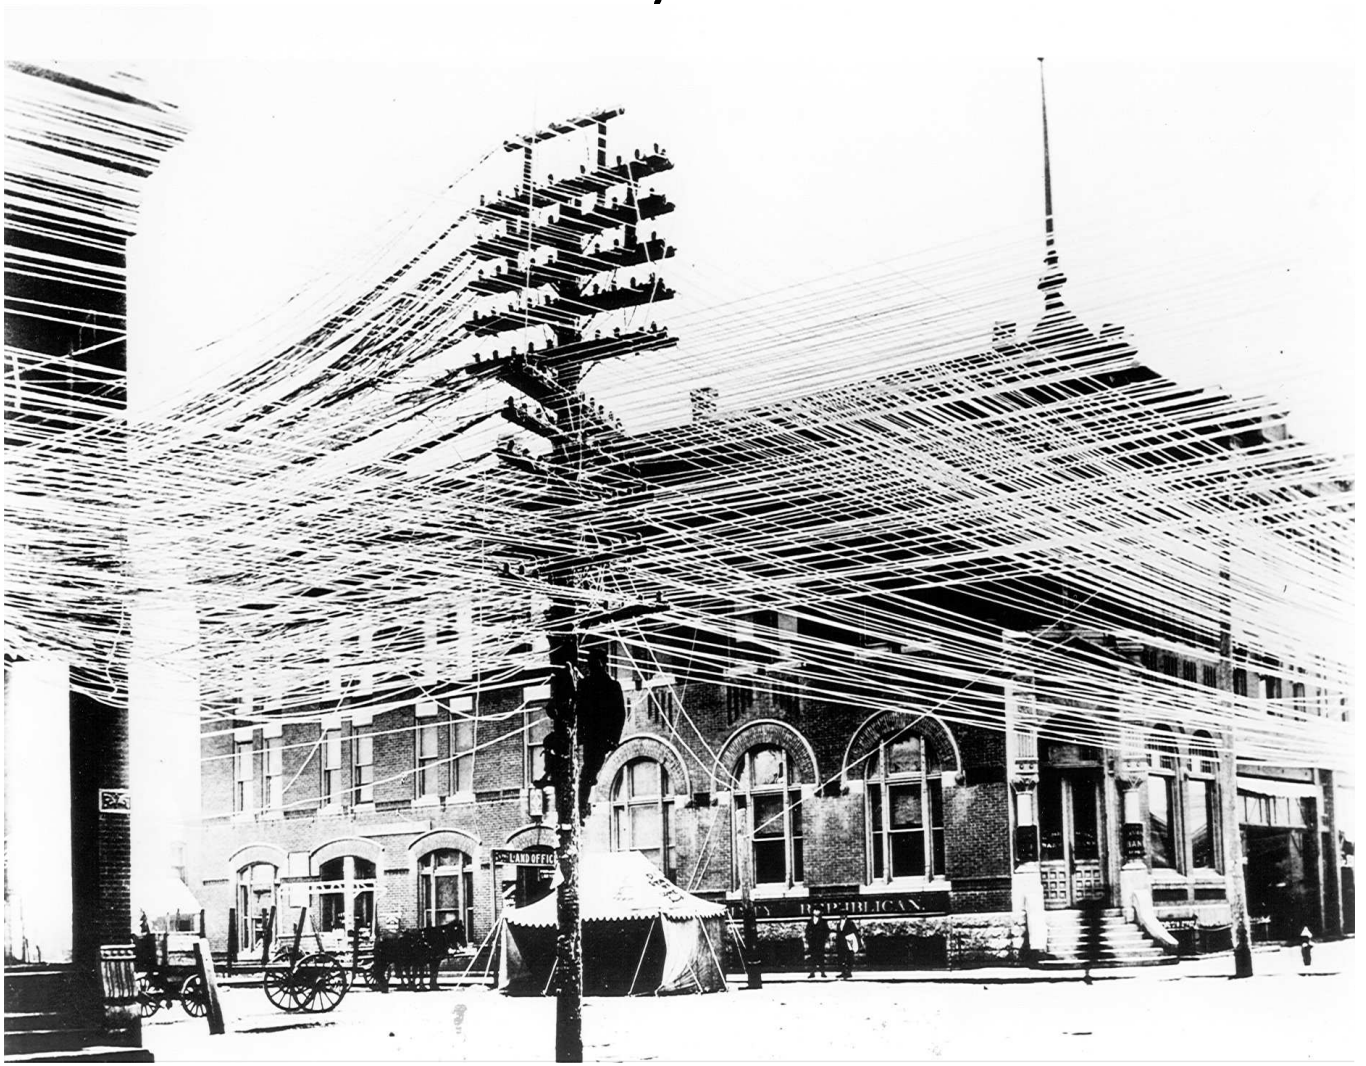

http://www.bellsystemmemorial.com/oldphotos_6.html

A switch

What if each person has one wire to a **switch**.

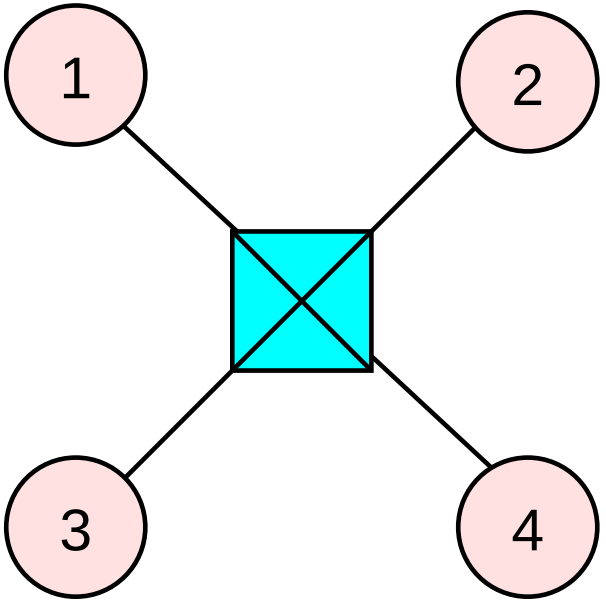

N=4
L=4

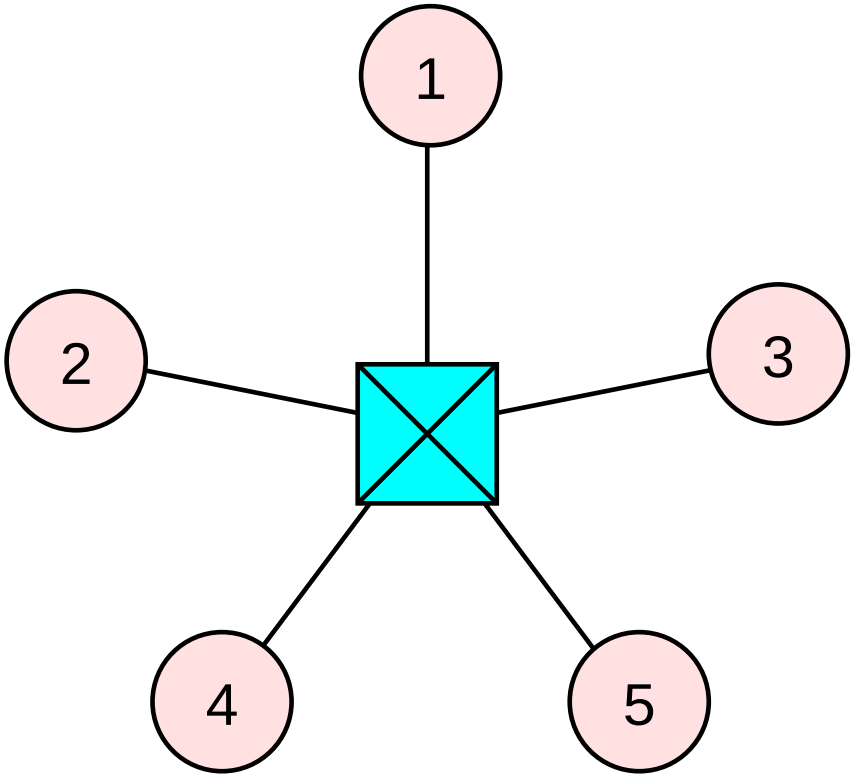

N=5
L=5

N nodes, then we have L=N links

Switchboards

THE UNIVERSITY OF ADELAIDE AUSTRALIA

So switches are great, but what is a switch?

http://www.bellsystemmemorial.com/oldphotos_6.html

Electromechanical switch

Almon Strowger was an undertaker in Kansas City in the late 1800's

- http://www.strowger.com/history.html
- Company 'Strowger Automatic Telephone Exchange' in October 1891
- step-by-step electromechanical switch [2]

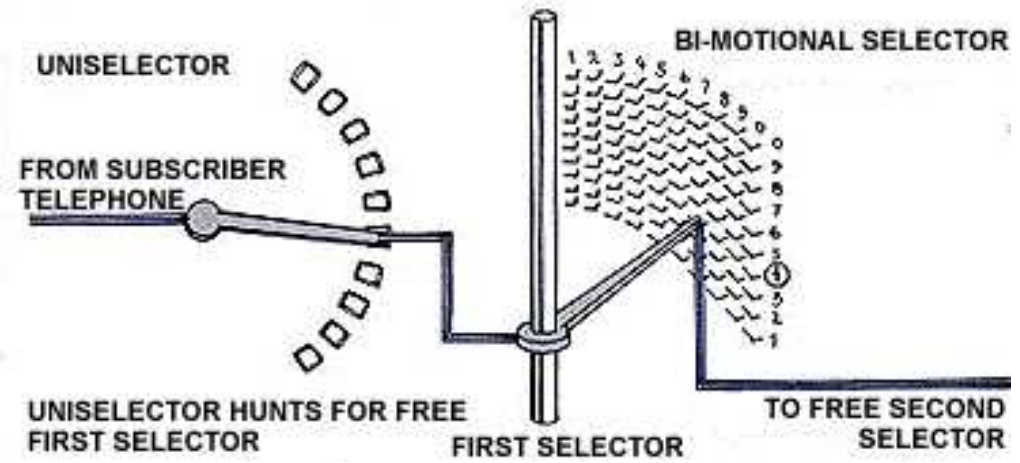

http://www.sigtel.com/tel_tech_sxs.html

Towards modern telephony

THE UNIVERSITY OF ADELAIDE AUSTRALIA

- switches get more complicated and sophisticated
  - electronic switch (instead of electromechanical)
  - 4ESS (like a building)
    http://www.att.com/history/nethistory/switching.html
- networks become hierachical
  - long distance versus local
- reliability and redundancy become important
  - alternate routing
- billing systems
  - harder than you think!

911

- 1889 (Britain) postal officials reprimand subscriber who used his phone to notify the firebrigade of a nearby fire. Directed to confine usage to his own affairs. [3]
- Great Britain (999), in 1937
- Australia (000), 1961
- investigated in US in 1958, legislation 1967
- 1st 911 call in 1968

911

- Skype has 113 million users (Sept. 2006)
  http://about.skype.com/2006/09/skype_introduces_video_calling.html
  - 54 million, Sept. 2005
- 7% of long-distance minutes (Morgan-Stanley)
- Skype's licence

  7.4 No Emergency Services. You expressly agree and understand that the Skype Software is not intended to support or carry emergency calls to any type of hospital, law enforcement agency, medical care unit or any other kind of Emergency Service. Skype, its Affiliates or Skype Staff are in no way liable for such emergency calls.

Computer networks

- First generation of electrical digital computers 1940s
- Second generation – late 1950s and early 1960s
  - transistor invented in 1947 (at AT&T)
  - direct networks: peripherals such as printers directly attached to computers
- Third generation, post-1964
  - integrated circuits
  - real computer networks start
- 1965, Moore's law discovered
  - computers get better and better ...

The "Internet"

THE UNIVERSITY OF ADELAIDE AUSTRALIA

What is the Internet?

- physical infrastructure
- protocols
- software
- services/applications
- standards
- architecture
- operational practices
- principles

All of the above, but lets think about its principles.

The Early Internet

Paul Baran, 1960s, envisioned a comm.s network that would survive a major enemy attack. The sketch shows three network topologies described in [6].

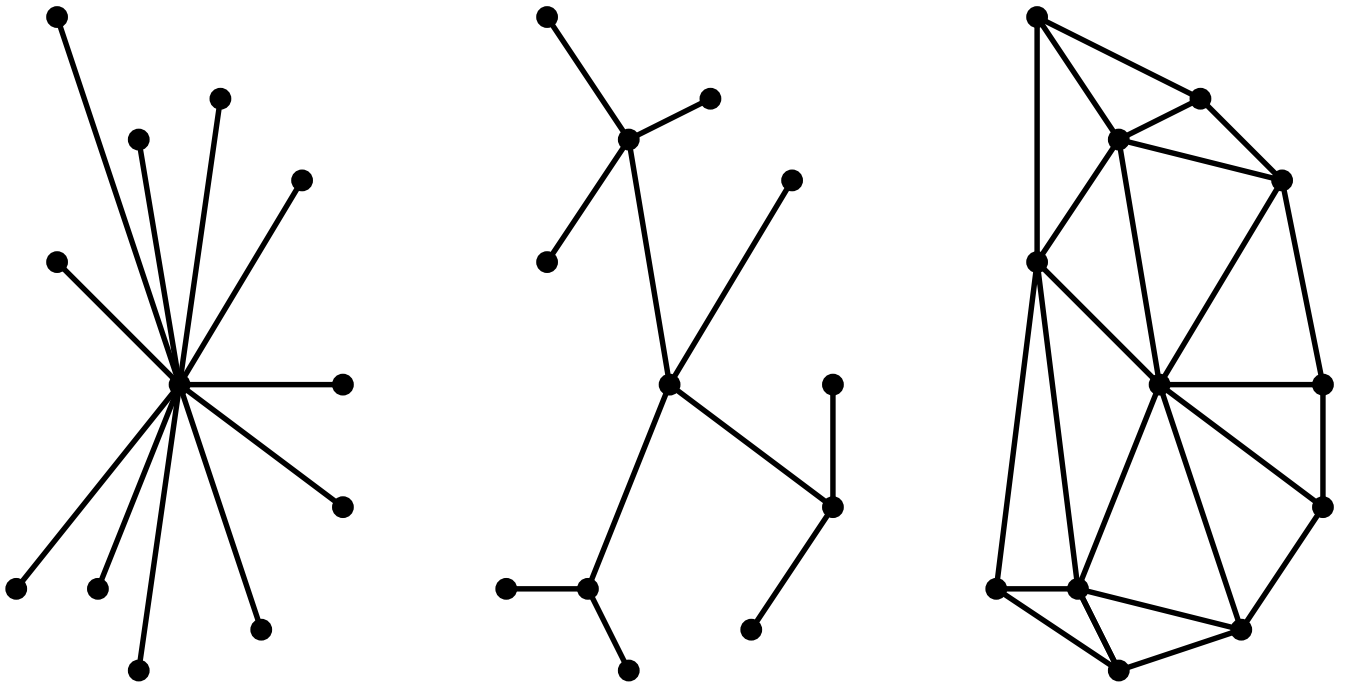

centralized **decentralized** **distributed**

Original available at
http://www.cybergeography.org/atlas/historical.html

The Early Internet

A rough sketch map of the possible topology of ARPANET by Larry Roberts. Drawn in the late 1960s as part of the planning for the network [7, p.50].

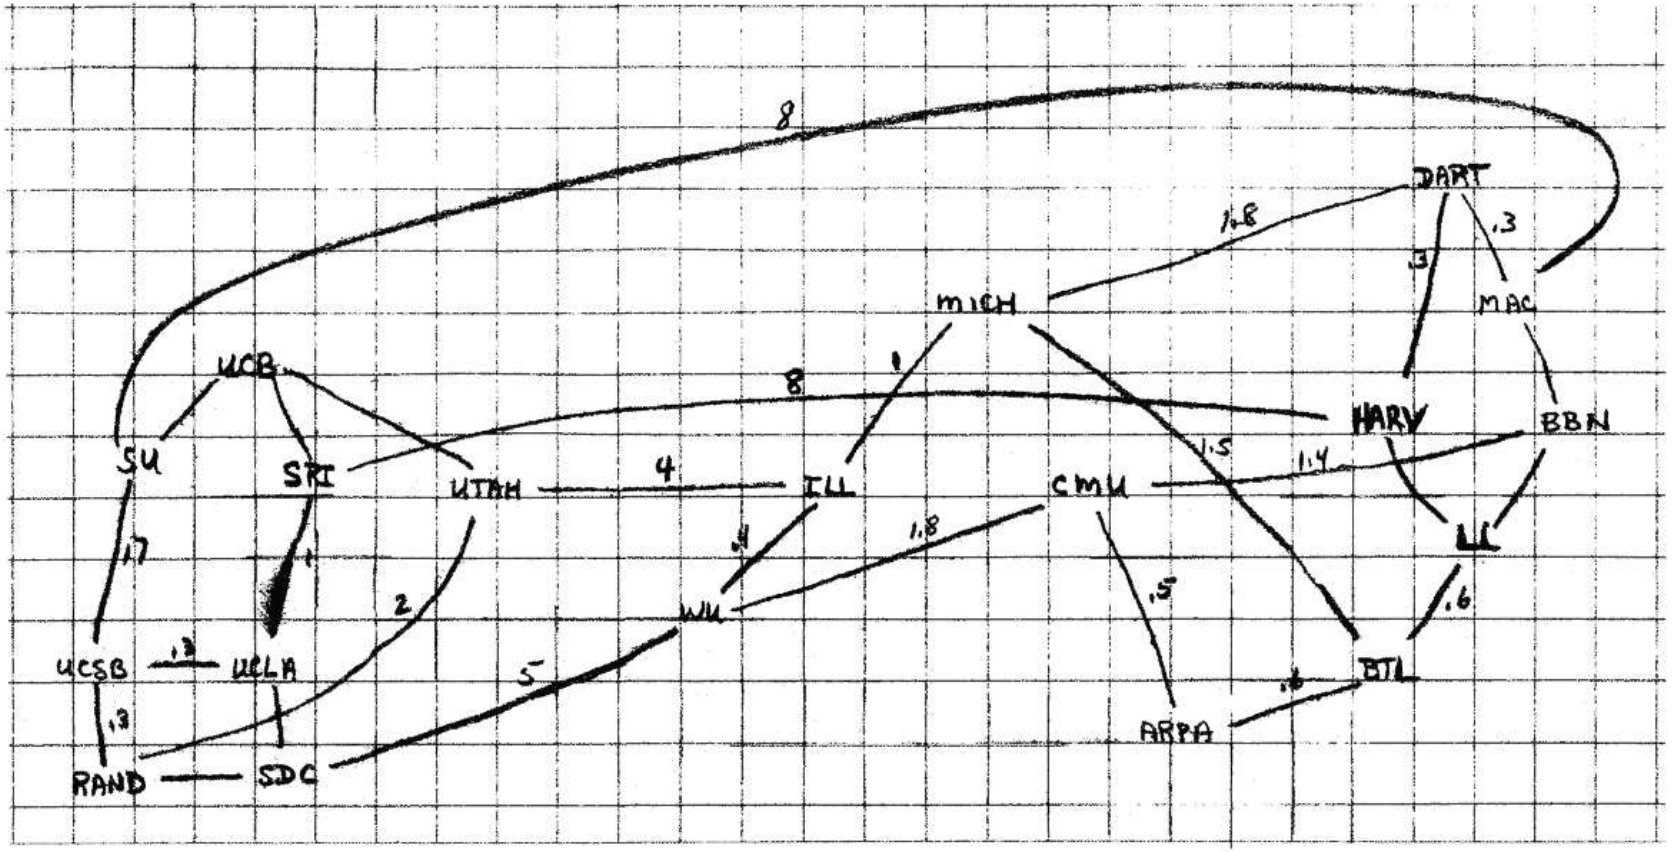

http://www.cybergeography.org/atlas/historical.html

The Early Internet

The first node on ARPANET at University California Los Angeles (UCLA) on the 2nd of September 1969 [8].

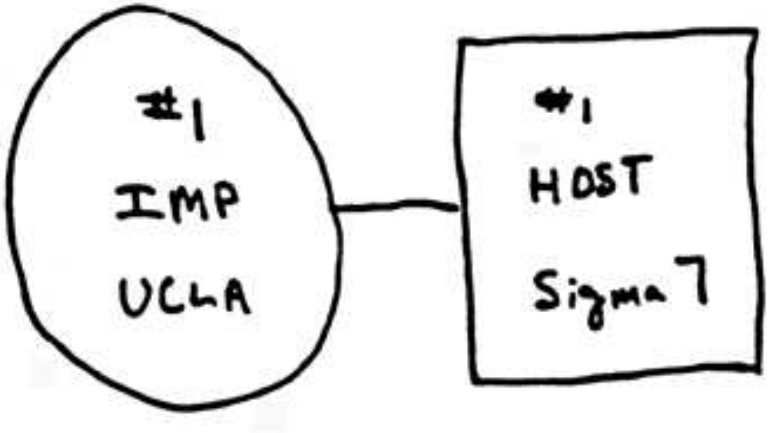

IMP = Interface Message Processor
what we would call a router
TIP = Terminal IMP
IMP to which terminals can directly connect
Host = computer (which provides services)

Available at
http://www.cybergeography.org/atlas/historical.html

The Early Internet

Dec 1969 "ARPA NETWORK". 4 nodes: Uni. of California Los Angeles (UCLA), Uni. of California Santa Barbara (UCSB), Uni. of Utah and the Stanford Research Institute (SRI) [8].

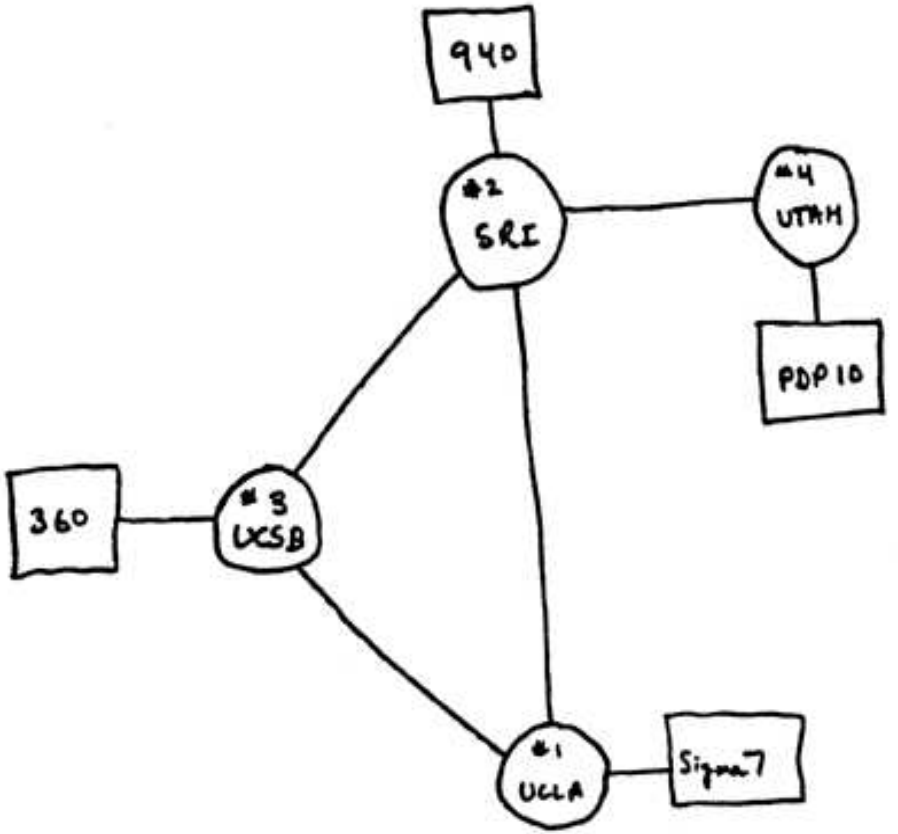

Available at
http://www.cybergeography.org/atlas/historical.html

The Early Internet

The map above shows the logical topology of ARPANET in April 1971. (computers connect direct to IMPs) [8].

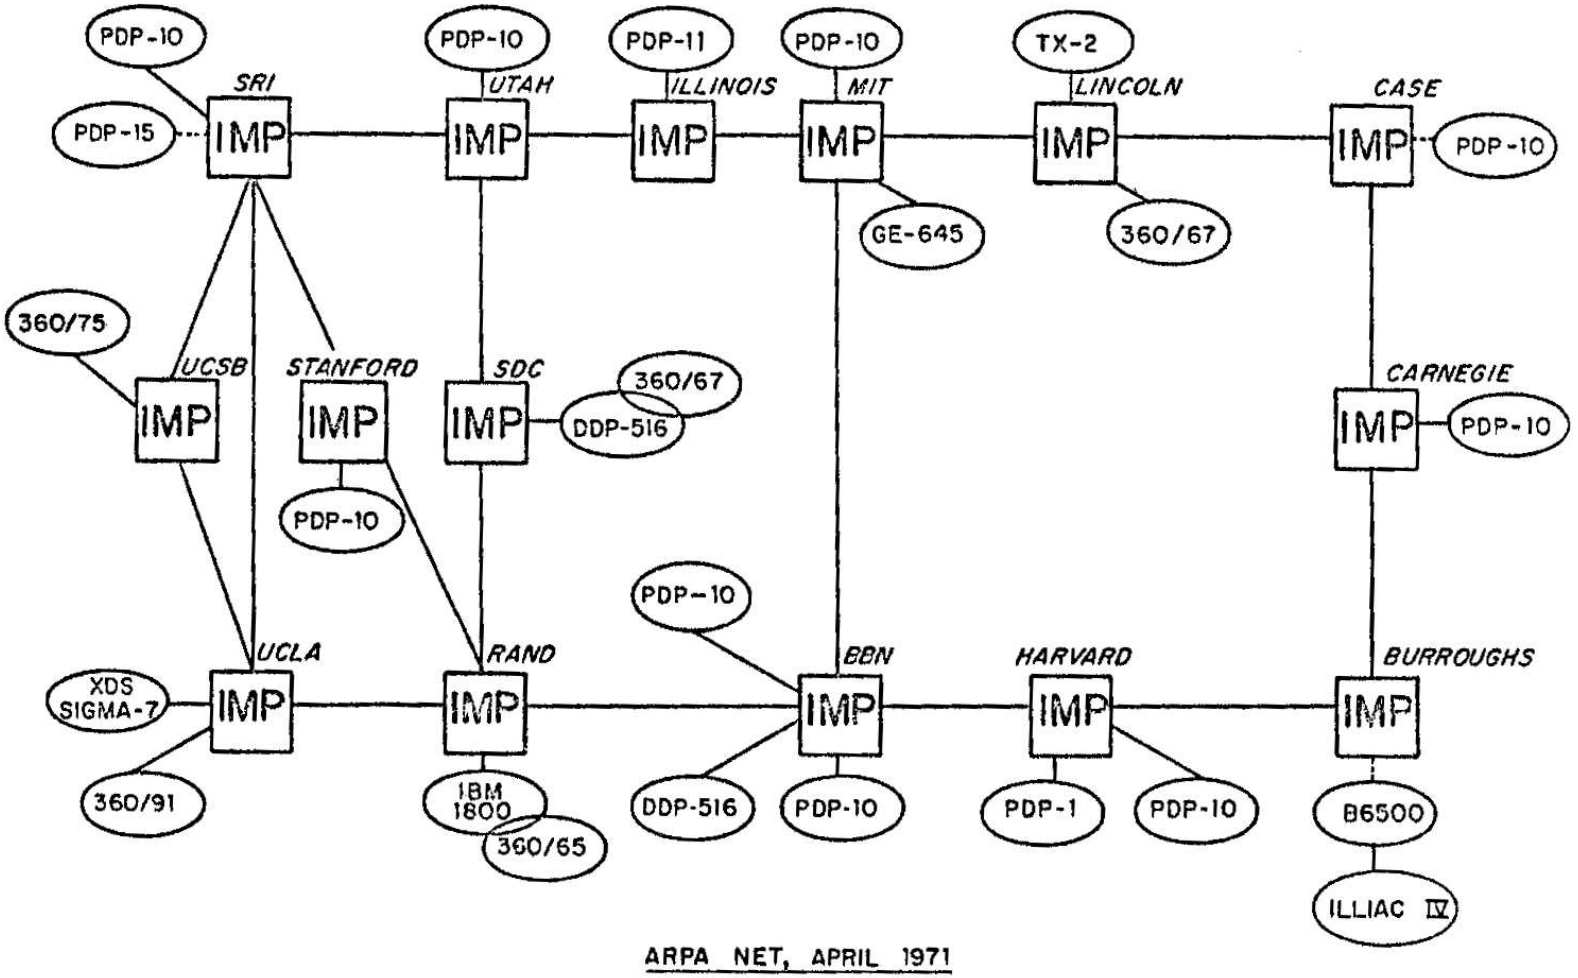

The Early Internet

ARPANET grew rapidly as more sites are connected [8].

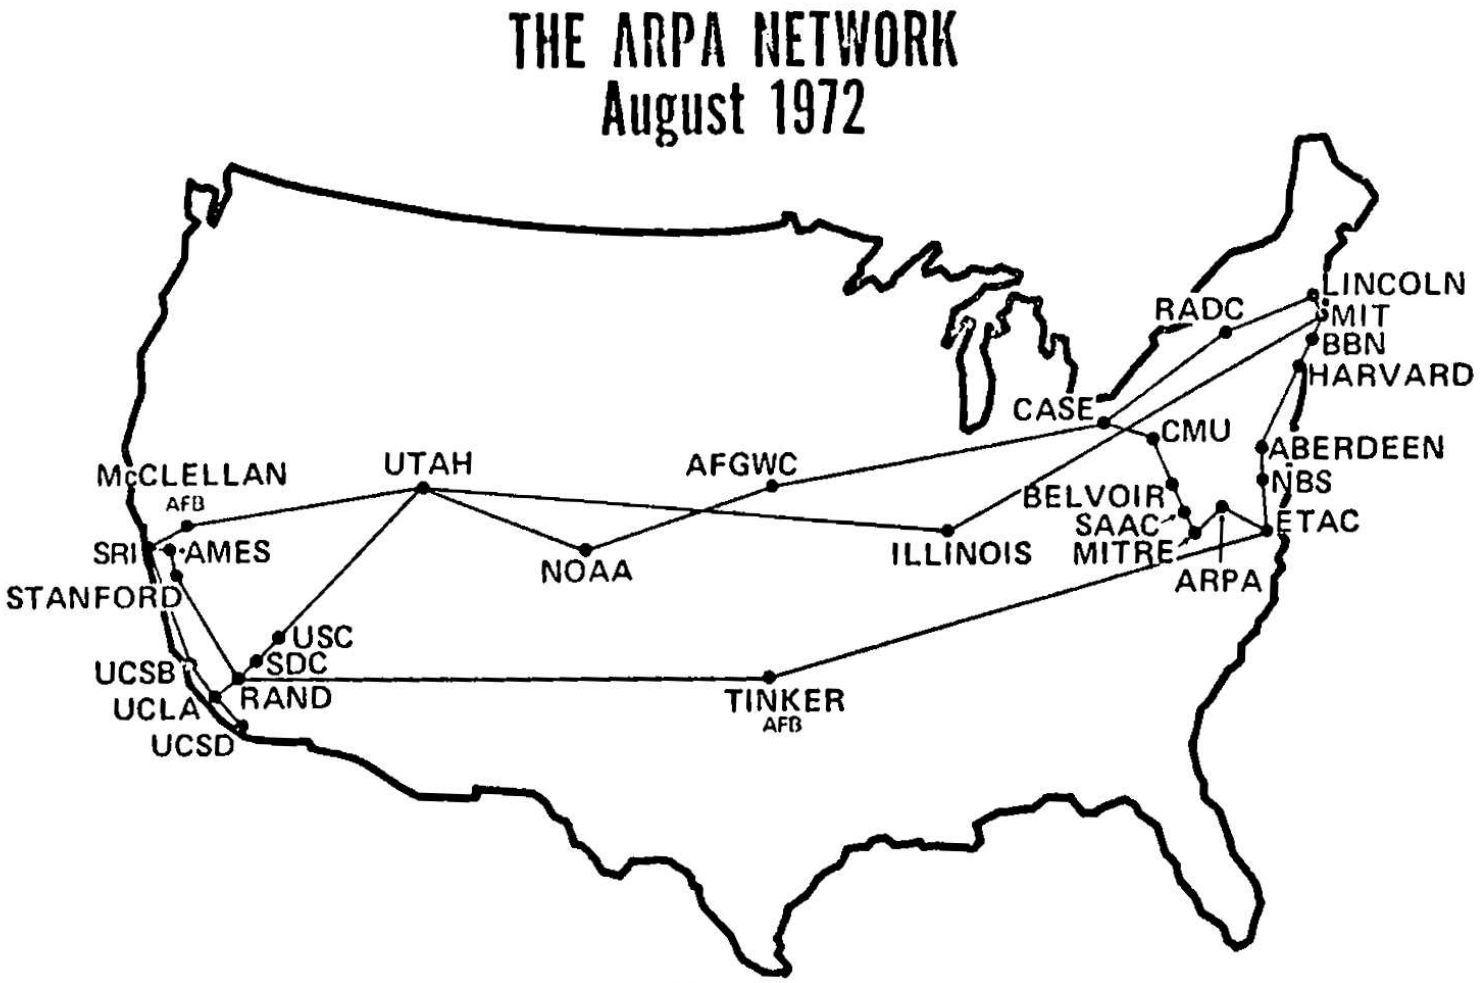

The Early Internet

ARPANET grew rapidly as more sites are connected [8].

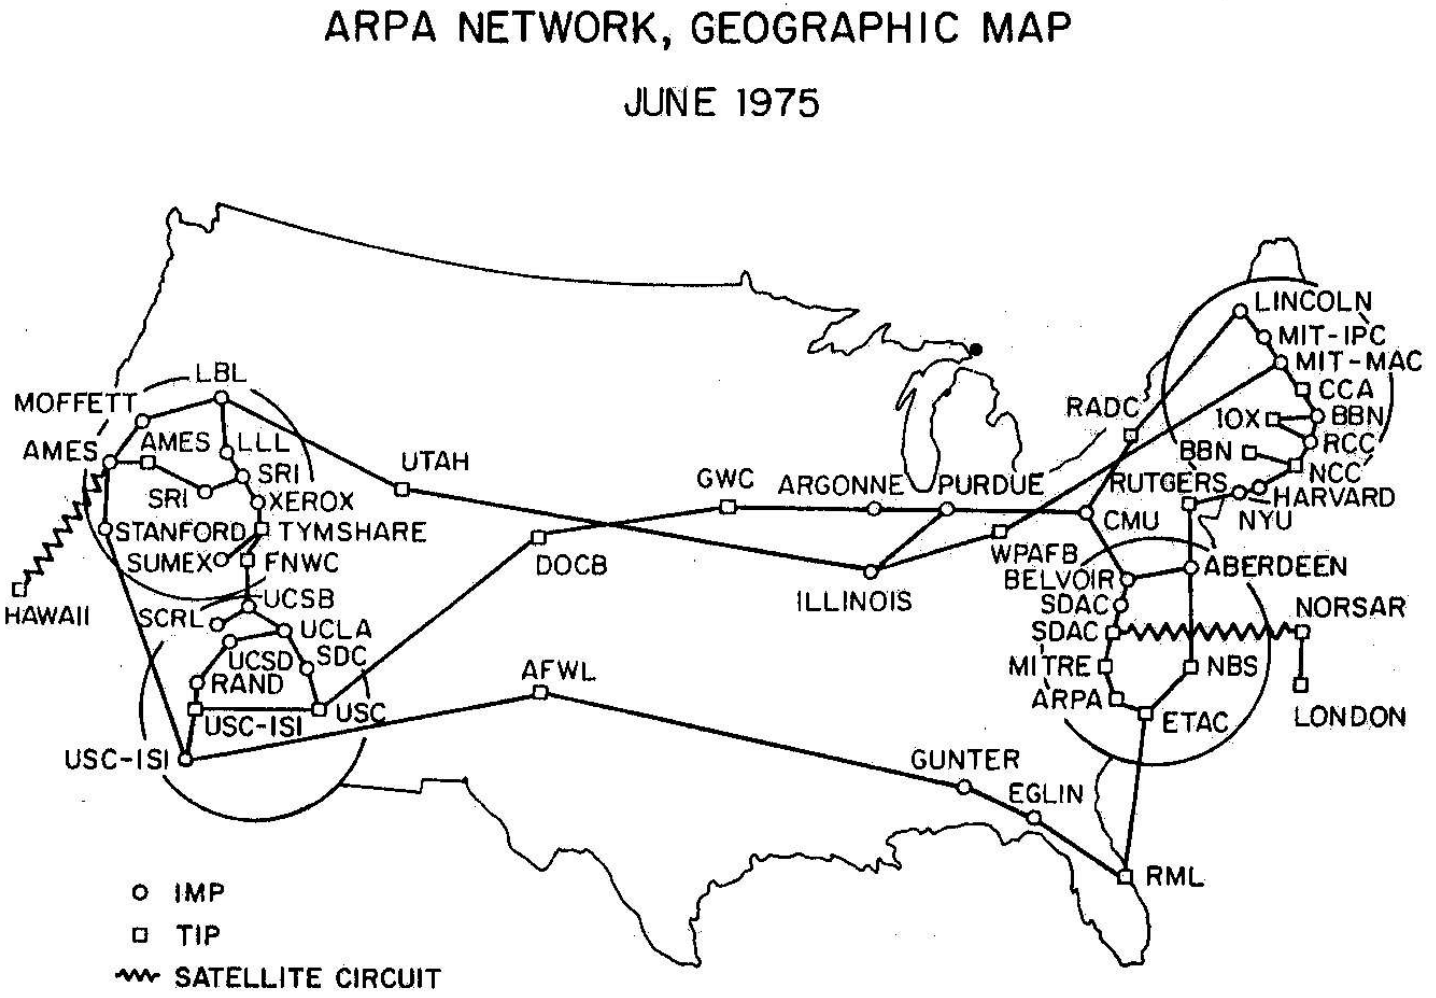

The Early Internet

ARPANET grew rapidly as more sites are connected [8].

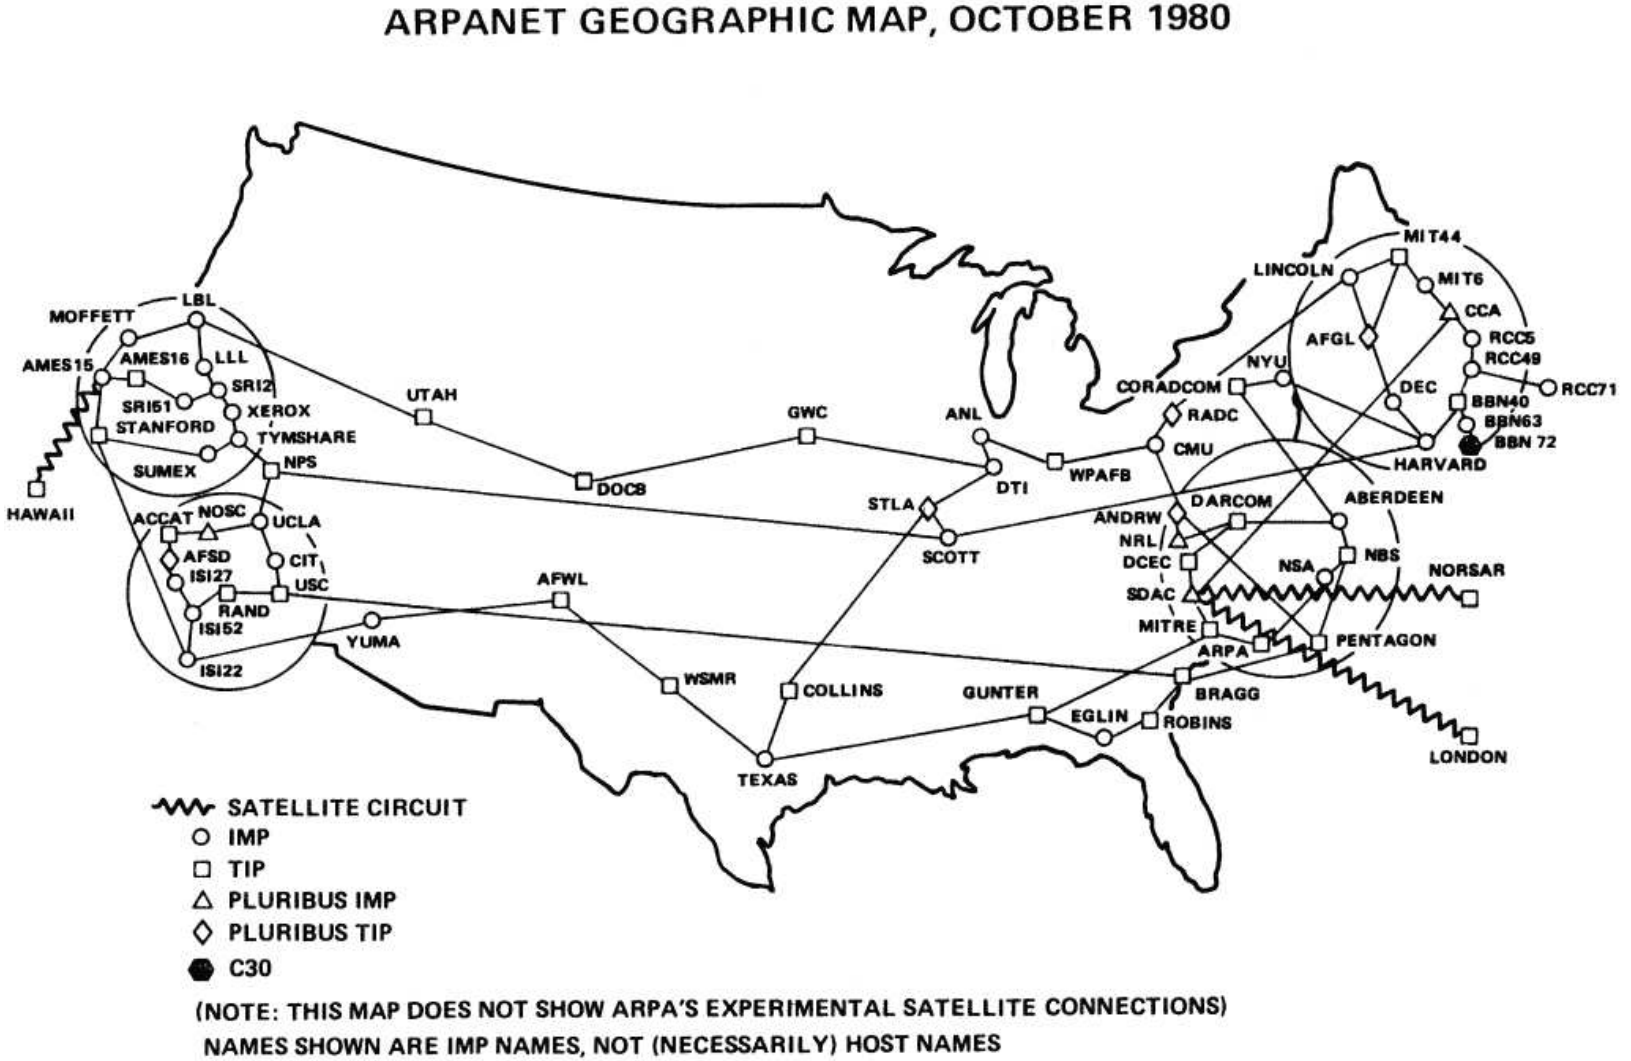

http://www.cybergeography.org/atlas/historical.html

The Internet: the 80's

- new developments
  - Personal Computers (PCs)
    $\Rightarrow$ lots more computers to network
  - Ethernet (1973, Robert Metcalfe) creates LANs
- the Internet
  - TCP/IP provides a way to hook up the LANs and PC over wide areas (standard in 1980)
  - scale gets bigger
    - numbers increase
    - becomes international
  - partitioning
    - ARPANET splits into MILNET and ARPANET in early 80's, followed by further additions

The Internet: the 80's

THE UNIVERSITY OF ADELAIDE AUSTRALIA

ARPANET/MILNET [81].

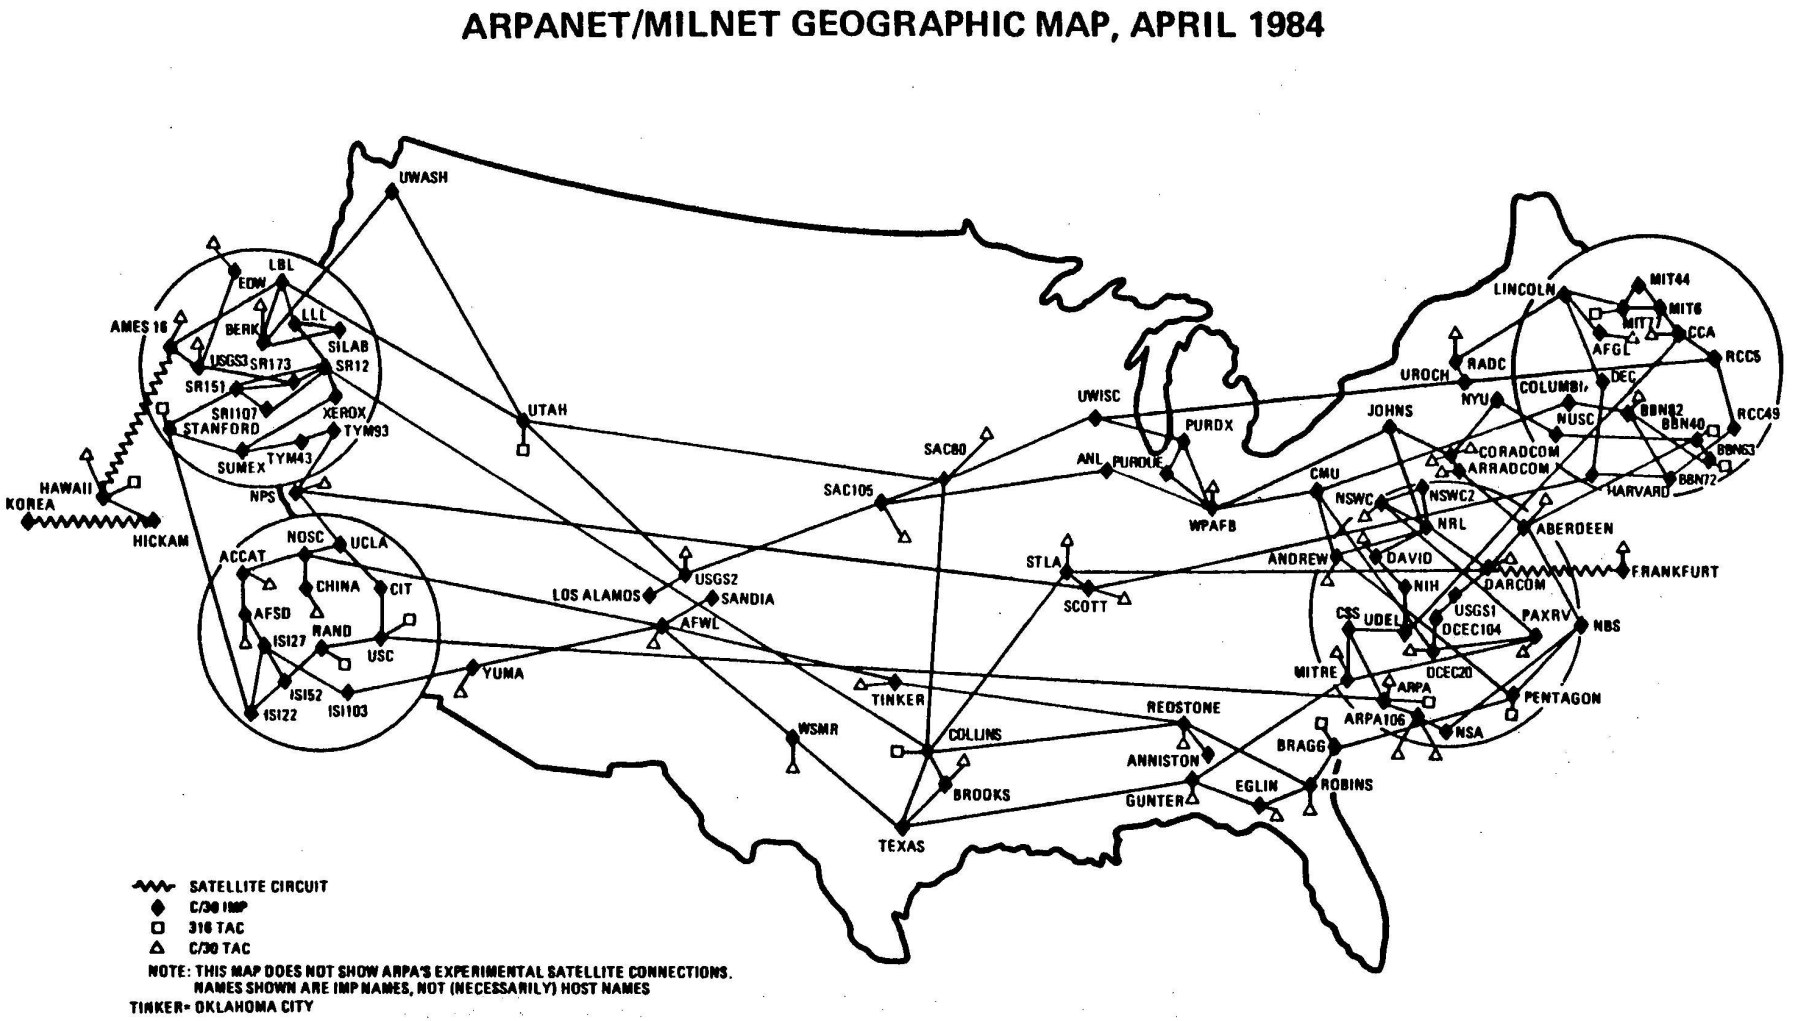

The Internet: the 90's

- ARPANET decommissioned 1990
  - NSF Backbone connects many other networks
    - Australia connected in 1990 [9]

```
It was the first, and being first, was best,
but now we lay it down to rest.
Now pause with me a moment, shed some tears.
For auld lange syne, for love, for years and years
of faithful service, duty done, I weep.
Lay down thy packet, now, O friend, and sleep.
```

Vinton Cerf, 1989

- commercial Internet services evolve
  - 1995 NSFNET terminated (replaced by vBNS)
  - effectively fully privatised Internet
  - links through exchange points

The Internet: the 90's

Australia's network 1991

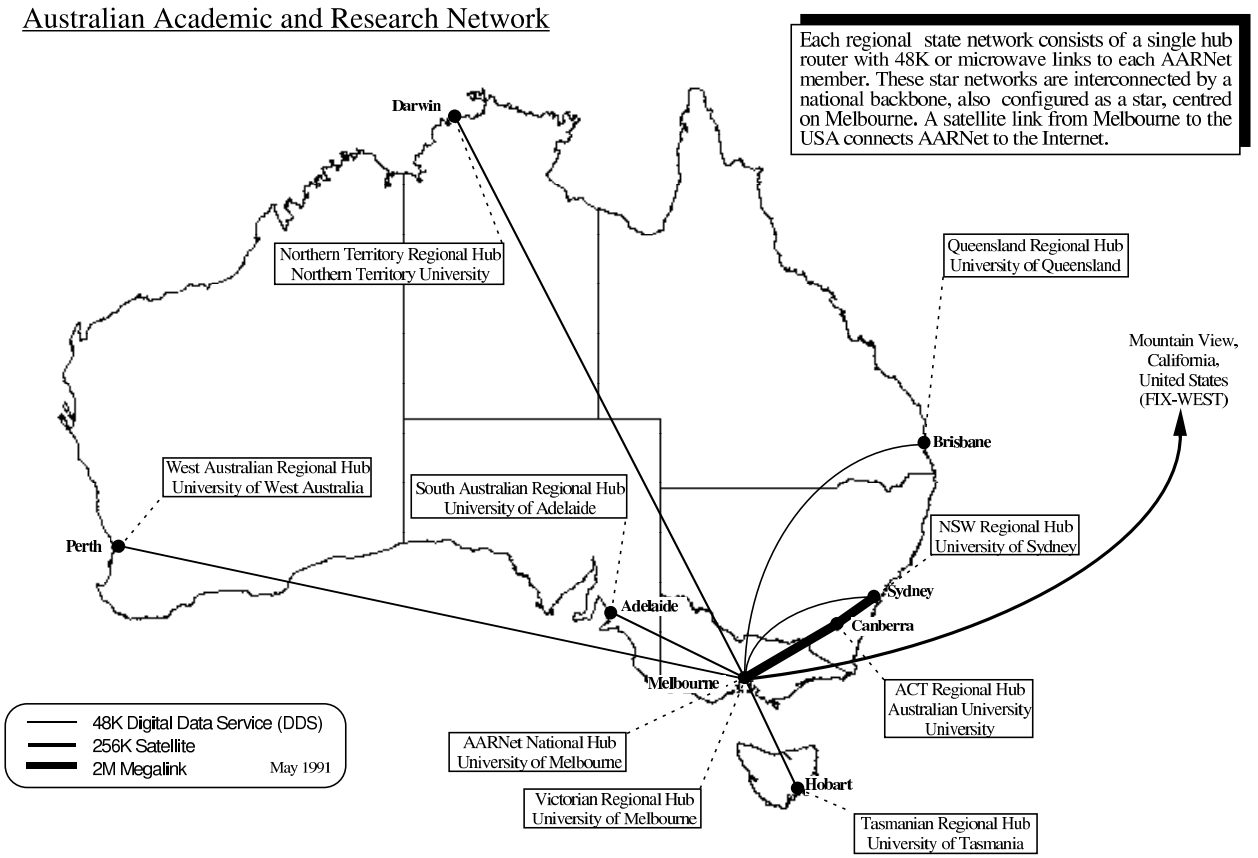

http://www.ucs.ed.ac.uk/fmd/unix/edftp/pub/maps/

New network
http://www.aarnet.edu.au/engineering/aarnet3/

The Internet: the 90's

http://www.w3.org/History.html

- 1990: World Wide Web
  Tim Berners-Lee created HyperText Markup Language, or HTML. Along with URL (Uniform Resource Locators), and HTTP (HyperText Transfer Protocol), created the web. Based on earlier work at CERN (1980).
- 1993: Mosaic (Marc Andreesen, NCSA)
  Mosaic became the first popular web browser. It was not only easy to use to access the World Wide Web, but it was also extremely easy to download and install!
- Killer app => the Internet takes off in a big way

These days

The Internet is broken into around 20,000 Autonomous Systems (ASs)

- AS is a separately managed network
- within an AS may use different routing, technology, management, ...
- may be a LAN, WAN, or combination
- example ASs:
  - ISP (Internet Service Provider)
  - Campus network
  - Enterprise network
  - Hosting center
- interior details of an AS are not exposed

a network of networks

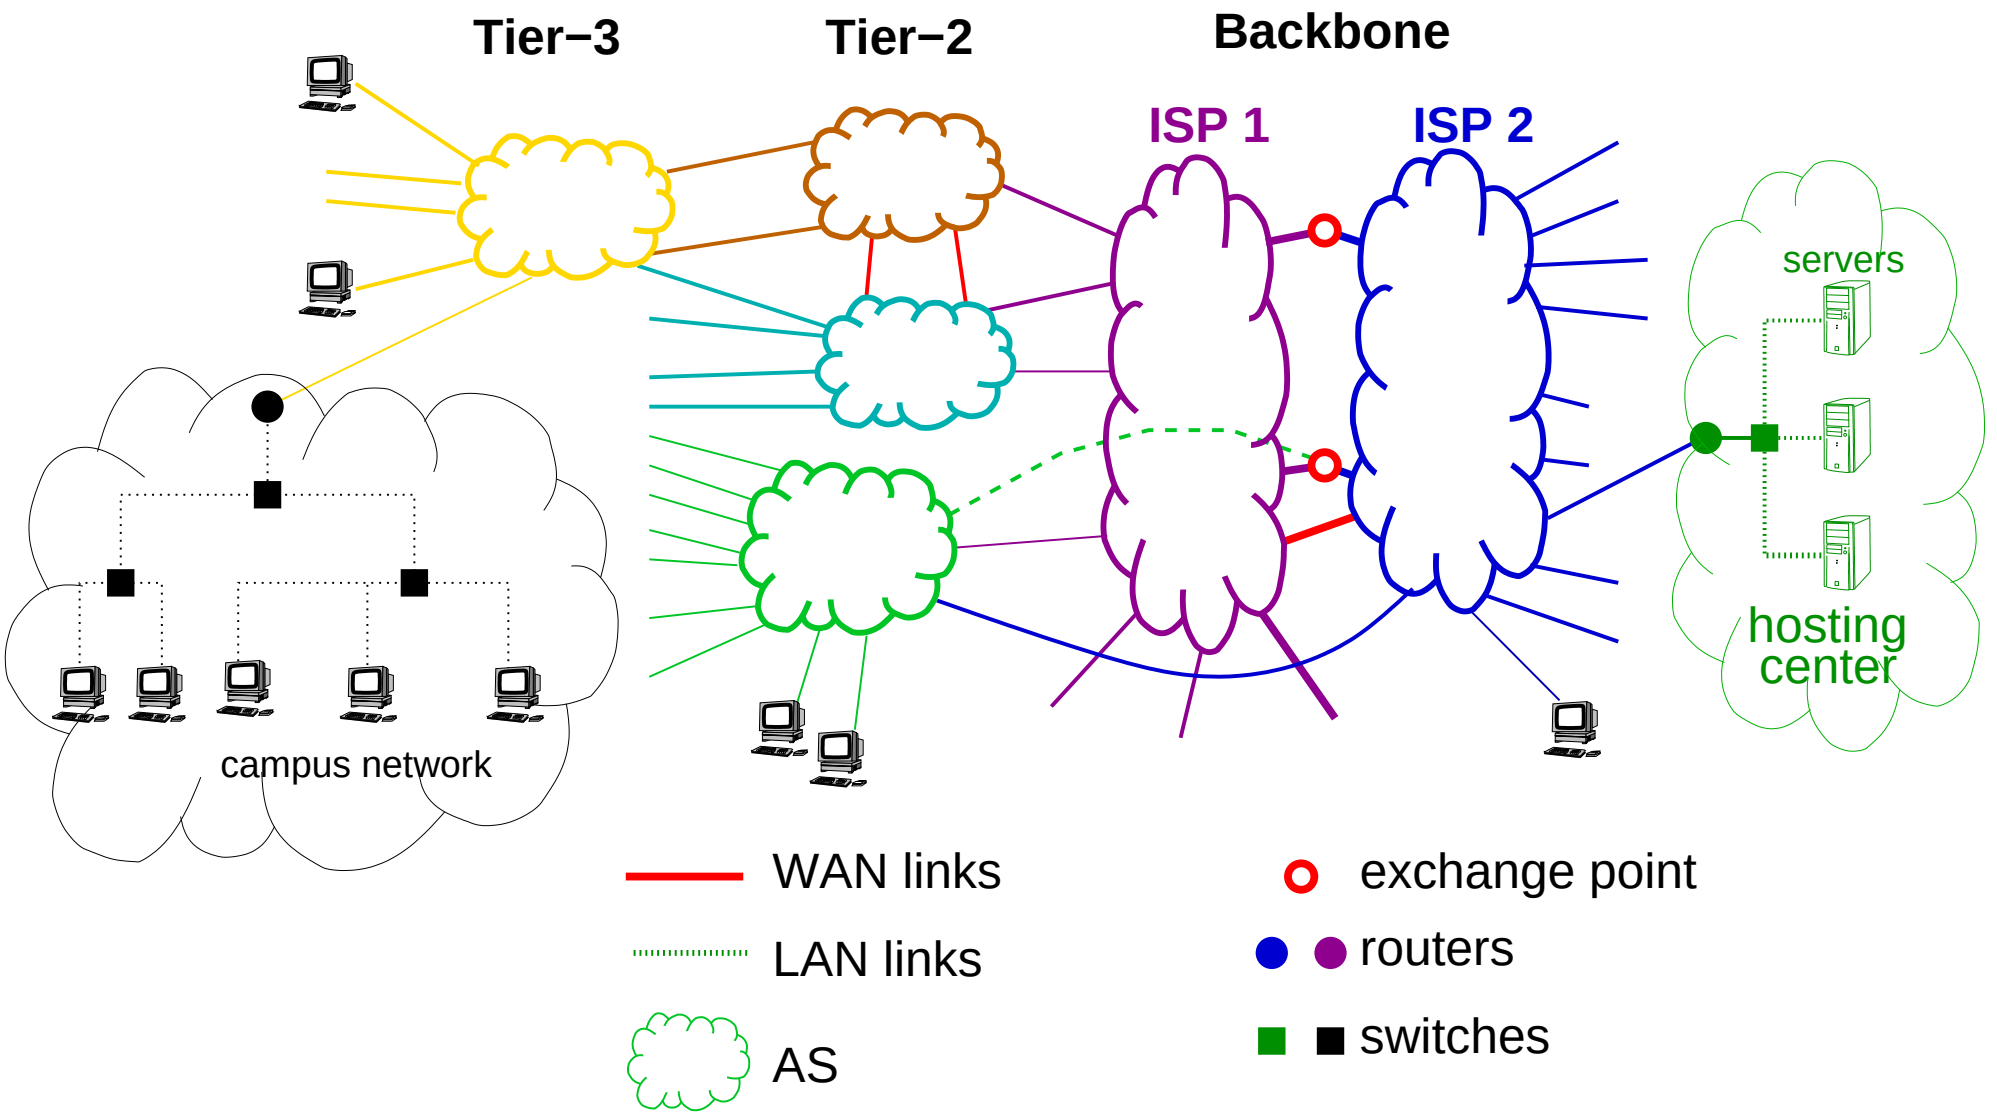

a network of networks

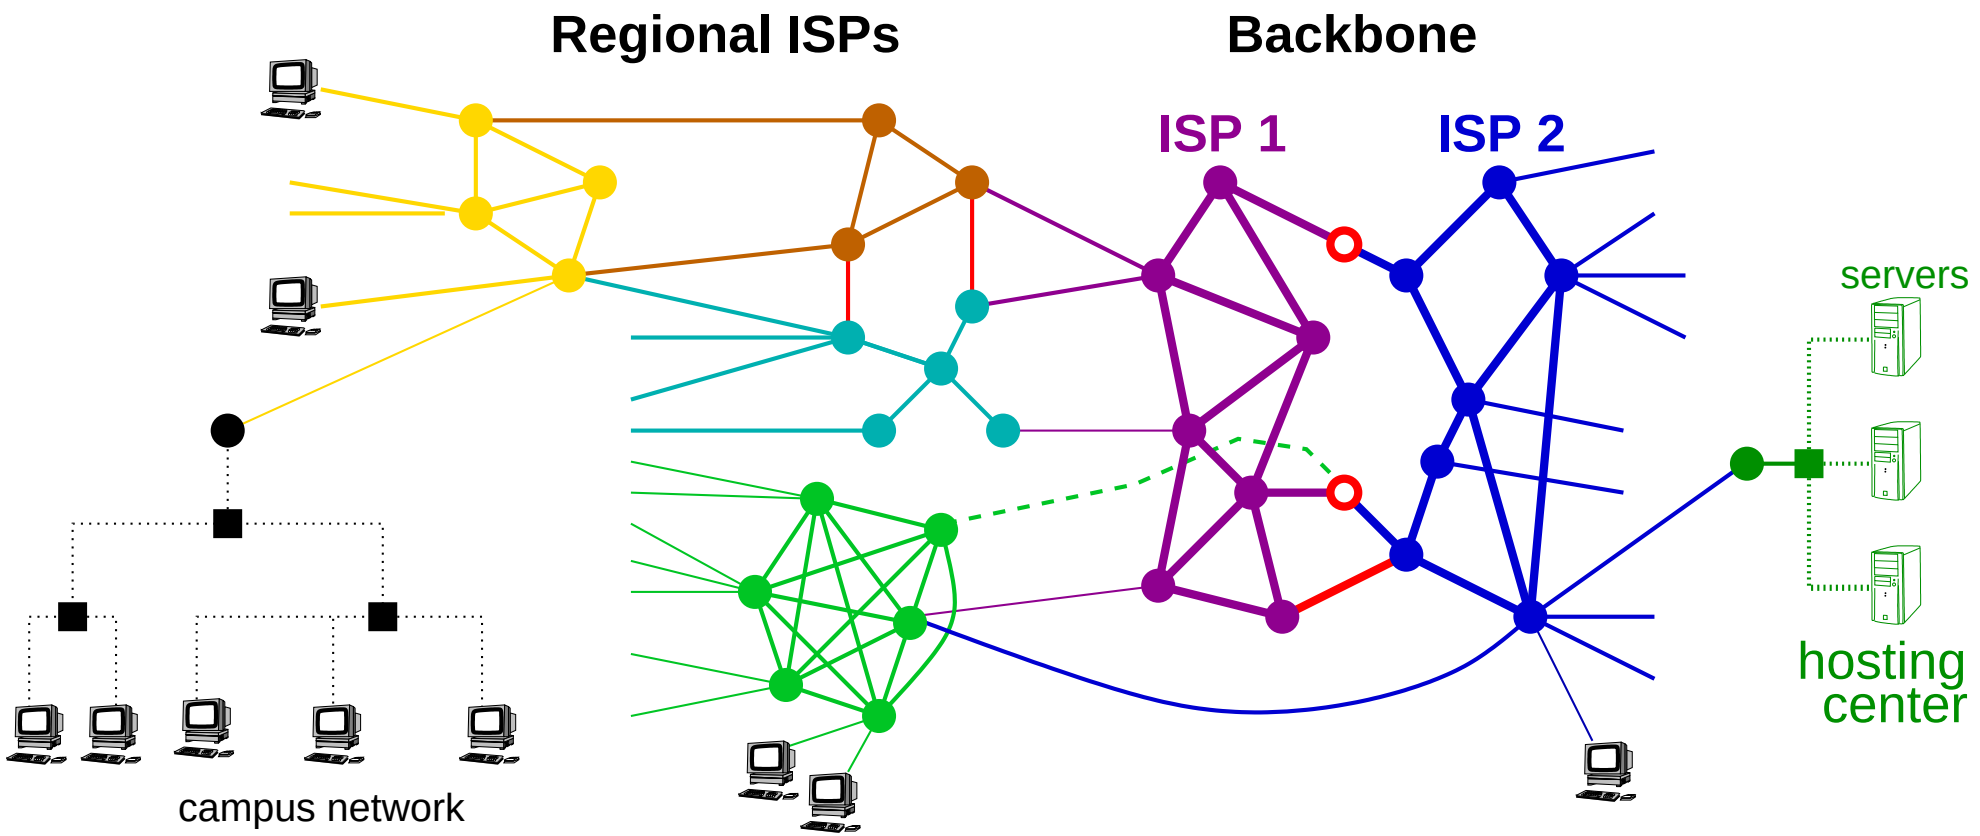

Number of subnets

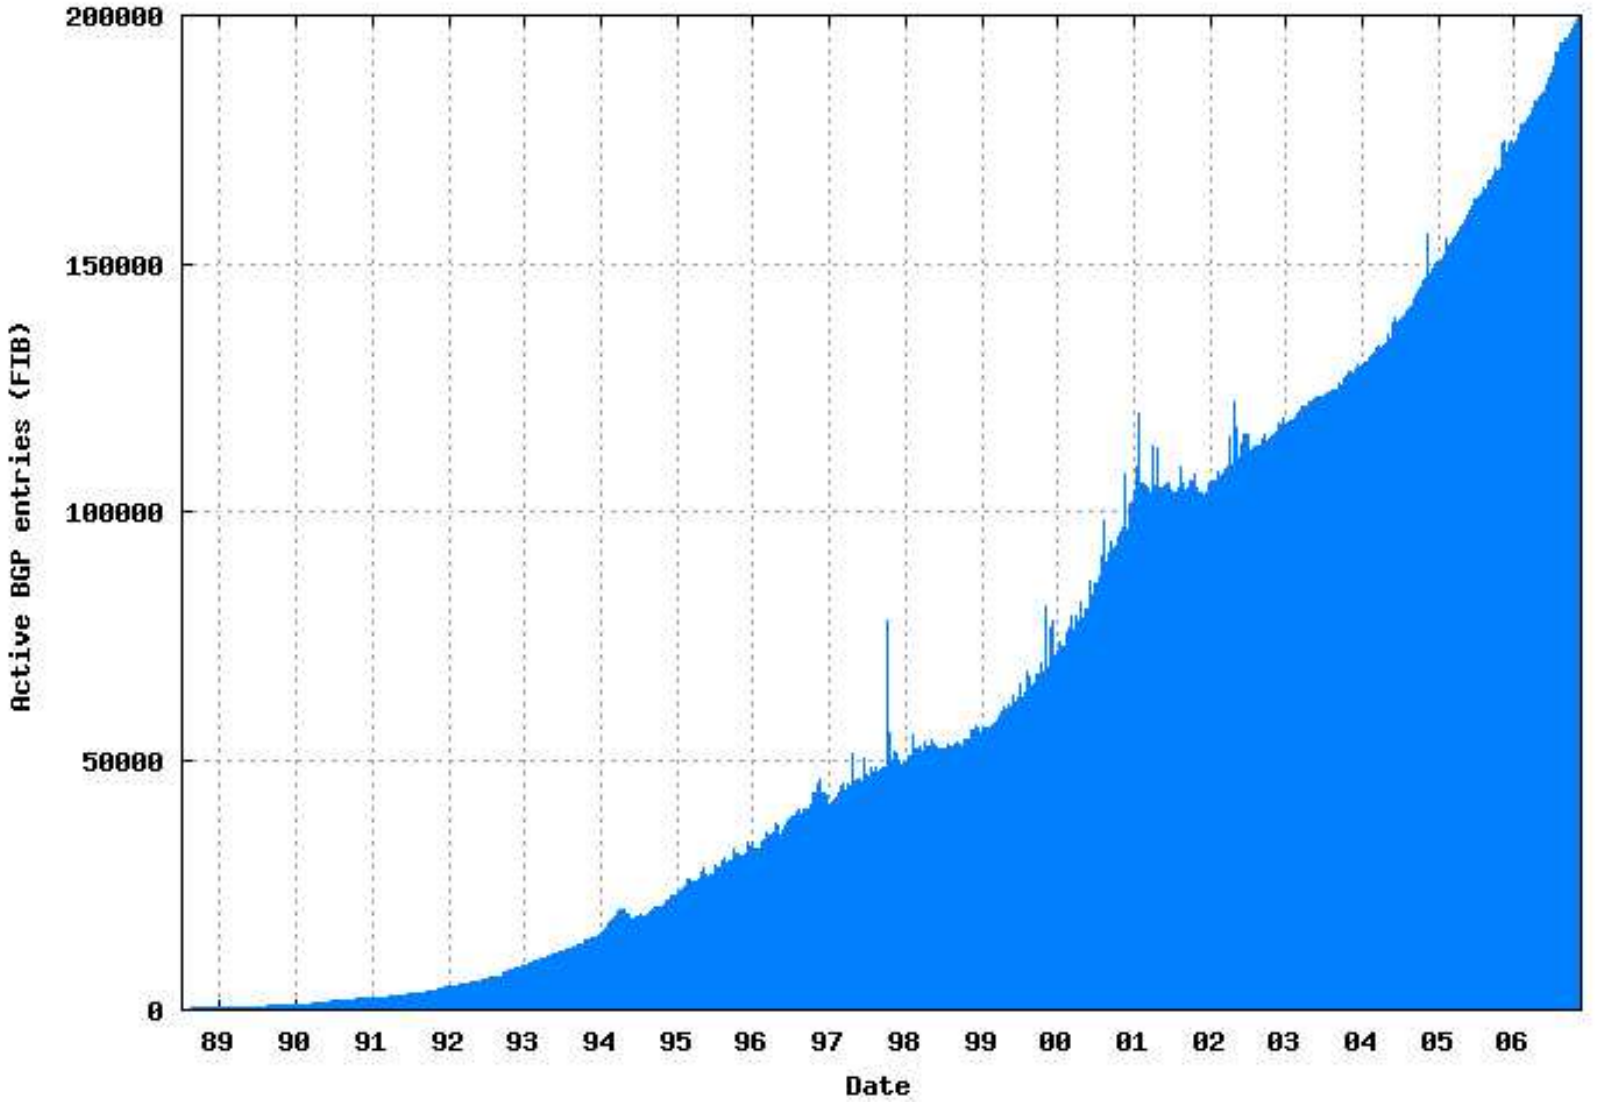

http://www.cidr-report.org/

Internet characteristics

THE UNIVERSITY OF ADELAIDE AUSTRALIA

The Internet

- highly dynamic (exponential growth)
- highly heterogeneous (20,000 Autonomous Systems, different technologies, and applications)
- 100's of thousands of subnets
- many millions of hosts
- (largely) decentralised
- (largely) stateless
- (largely) self-configuring

In many ways it resembles some kind of living thing!

The Great Big Jellyfish

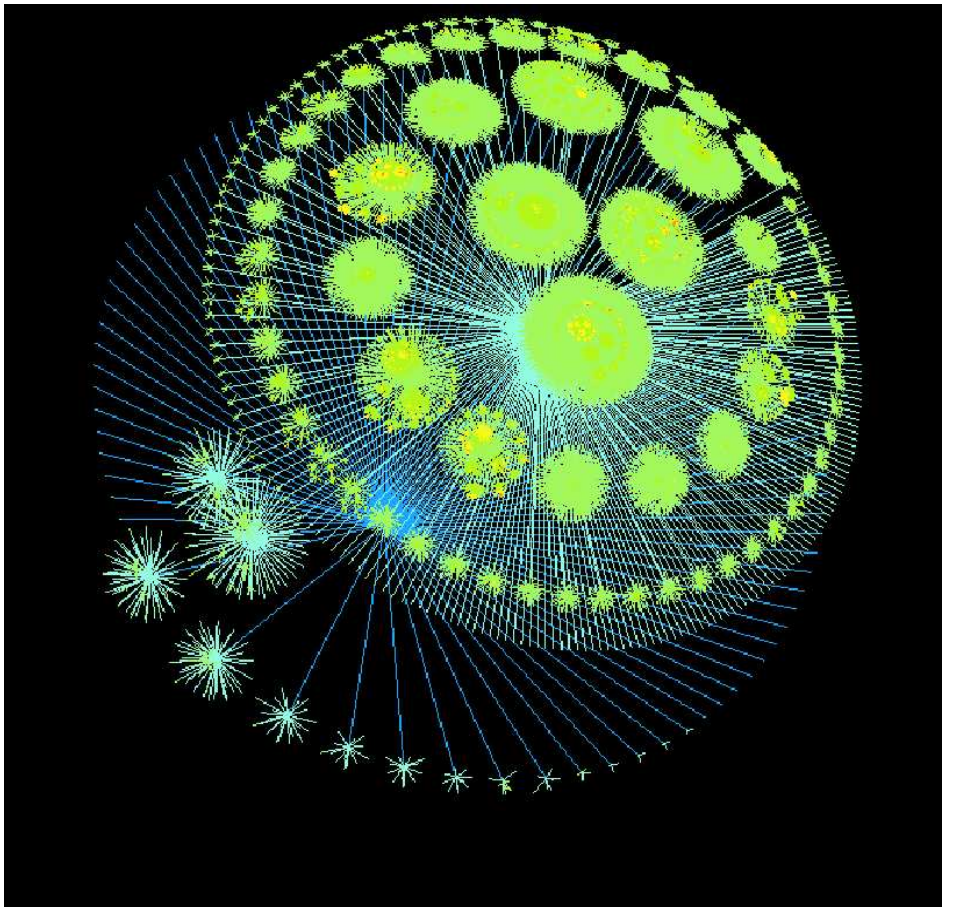

Created by Ashley Flavel using Walrus.

Distributed control

- Distributed control is the key to making this work
  - its one of the fundamental principles of the Internet's design
- But its not all roses and jellyfish
  - how do you design/optimize something like this?
    - you can't (at least not easily)
      - tragedy of the commons
      - multiple objectives of different ASs
      - unwillingness to co-operate
    - no guarantees of convergence when you try!

Internet Design Principles

THE UNIVERSITY OF ADELAIDE AUSTRALIA

- **distributed control:** as compared to centralized, or decentralized [14].
- **packet switching not circuit switching:** Don't reserve bandwidth for a connection.
- **layered model:** with a thin waist.
- **robustness principle:** Be liberal in what you accept, and conservative in what you send [11, 12].
- **end-to-end principle:** Smart terminals, dumb network [13, 14].
- **general issues:** simplicity, modularity, performance [12].
- **deployment issues:** scale, incremental deployment, heterogeneity [12].

The Early Internet

THE UNIVERSITY OF ADELAIDE AUSTRALIA

Kleinrock's insight [5]

- computer traffic is bursty (it comes in spurts)
- more efficient to transmit packets of data on-demand than to reserve circuits between computers
  - setting up a circuit takes time (high latency)
  - keeping up a circuit set up is inefficient
    - not used most of the time
  - all you want to do is send one little chunk of data
    - example: typing – one character at a time
    - even a whole email is quite small
  - alternative: send data as packets

Packets vs circuits

THE UNIVERSITY OF ADELAIDE AUSTRALIA

- Bell-heads vs Net-heads
  - Bell-heads from old Bell system (AT&T included)
  - Net-heads: new generation, who grew up on the ARPANET/Internet
- Bell-heads believe you need a dedicated circuit
  - like a phone line (but higher speed)
  - said the ARPANET would never work
- Net-heads think circuits are a waste of time
  - poor use of resources when traffic is bursty [15].
  - invented the ARPANET/Internet
- this is a theological debate

packets are not all good

Some things IP packets don't do well

- billing
  - circuits are easy to bill
  - packets are not
    - most Internet charging is flat rate
    - Australia is an exception
  - does this really matter? [17].
- QoS (Quality of Service)
  - e.g. network not design to provide voice quality
  - maybe it can be fixed? Lots of research.
- security (crypto doesn't fix DoS, Worms, ...)
- washing dishes

MPLS

THE UNIVERSITY OF ADELAIDE AUSTRALIA

Multi-Protocol Label Switching

- MPLS creates virtual circuits between end-points
  - connections are not between end-users though
  - allows **multiplexing** of traffic inside a connection
  - multiplexed traffic is less bursty
- MPLS seems to be changing the underlying infrastructure of the Internet.

History's Lesson

- 2 steps forward, 1 step back
  - Strowger
    - bigger, faster networks
    - lost operators' local information
  - distributed control
    - robustness
    - cost of optimality
  - packets
    - cheap, efficient, high-rate data
    - poor billing and QoS
  - wireless telephony
    - mobility
    - loss in quality

History's Lesson

- step back is often perceived as bad
- has it always been that bad?
  - Strowger
    - aren't phone directories better than operators?
  - distributed control
    - optimality doesn't matter as much as robustness
  - packets
    - flat rate billing isn't such a bad thing
  - wireless telephony
    - quality wasn't as important as once thought
- really, there are tradeoffs

Fundamental Changes

THE UNIVERSITY OF ADELAIDE AUSTRALIA

Rather too many to go into today

- Wireless Internet
- Network convergence
- Net neutrality
- IPv6
- real Worldwide governance
- TCP landspeed
- Knowledge plane
- Asia-Pacific – the new engine of growth

Has the Internet come of age?

THE UNIVERSITY OF ADELAIDE AUSTRALIA

- size → clearly the Internet is big enough
- stability → clearly not!
- maturity →
  - we're still arguing about some basic issues
  - but there are some good signs along the way

Has the Internet come of age?

THE UNIVERSITY OF ADELAIDE AUSTRALIA

- size → clearly the Internet is big enough
- stability → clearly not!
- maturity →
  - we're still arguing about some basic issues
  - but there are some good signs along the way
- perhaps best charactered as a "teenager"
  - going through a bit of a growth spurt
  - rebellious
    - its going through its "packets" phase
  - spends a bit too much time in the bedroom
  - likes loud music a bit too much

Conclusion

THE UNIVERSITY OF ADELAIDE AUSTRALIA

Has the Internet come of age?

- Douglas Adam's may be off by a couple of years
- When will I say its come of age?
  - I still have LPs, but when I'm giving a class and want to talk about them to students, I have to show them one to explain
  - When I have to do that for the Internet, then whatever it has become will have come of age.
- Finally: will people please stop talking about Internet enabled fridges.

References

[1] Aeschylus, Agamemnon. http://classics.mit.edu/Aeschylus/agamemnon.html, 458 B.C.E.

[2] A. B. Strowger, "Automatic telephone exchange." United States Patent Office, patent no. 447,918, March 10th, 1891.

[3] A. M. Odlyzko, "Pricing and architecture of the internet: Historical perspectives from telecommunications and transportation," in TPRC, 2004.

[4] G. E. Moore, "Cramming more components into integrated circuits," Electronics, vol. 38, April 1965.

[5] L.Kleinrock, "Information flow in large communication networks." RLE Quarterly Progress Report, July 1961.

[6] P. Baran, "On distributed communications: 1. introduction to distributed communications network." RAND Memorandum, August 1964.

[7] K. Hafner and M. Lyon, Where Wizards Stay Up Late: The Origins of the Internet. Touchstone, 1996.

[8] V. Cerf and B. Kahn, "Selected ARPANET maps," Computer Communications Review (CCR), vol. 20, pp. 81–110, 1990.

[9] J. Abbate, Inventing the Internet. MIT Press, 1999.

[10] A. M. Odlyzko, "Internet traffic growth: Sources and implications," in Optical Transmission Systems and Equipment for WDM Networking II, Proc. SPIE, vol. 5247, pp. 1–15, 2003.

[11] R. Braden, "Requirements for Internet hosts – application and support." IETF, Request for Comments: 1123, 1989.

[12] B. Carpenter, "Architectural Principles of the Internet." IETF, Request for Comments: 1958, 1996.

[13] J. H. Saltzer, D. P. Reed, and D. D. Clark, "End-to-end arguments in system design," in Second International Conference on Distributed Computing Systems, pp. 509–512, April 1981.

[14] D. Clarke, "The Design Principles of the DARPA Internet Protocols," Computer Communication Review, vol. 25, January 1995.

[15] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," IEEE/ACM Transactions on Networking, vol. 2, pp. 1–15, Feb 1994.

[16] J. Postel, L. Kleinrock, V. Cerf, and B. Boehm, "Act One - The Poems." IETF, Request for Comments: 1121, 1989.

[17] A. M. Odlyzko, "The history of communications and its implications for the Internet." http://www.dtc.umn.edu/~odlyzko/doc/networks.html.