CSE 486/586: Distributed Systems

The Internet in Two Lectures (Part 1)

Ethan Blanton

Department of Computer Science and Engineering
University at Buffalo

Last Time…

- Attendance is required
- Academic Integrity is important
- Projects are individual
- A brief overview of distributed systems

Administrivia

Remember

- AI Quiz due Friday
- Programming Assignment 1 due next Monday

We’re using Piazza, be sure to join!

Introduction I

The next two lectures will be about the Internet.

Why?

- The design of the Internet influences distributed systems
- It’s also really good!
- Steve likes it, and so do I

This isn’t a networking course, just a refresher.

Introduction II

The Internet handles many distributed systems problems

- Unreliable communication
- Failure detection
- Agreement

(However, for reasons we'll get into, we have to solve some of them again!)

A Brief History…

- **1961:** Leonard Kleinrock writes the first paper on packet switching
- **1962:** J. C. R. Licklider of MIT wrote some memos about a “Galactic Network”
- **1962:** Licklider becomes head of DARPA Computer Research
  - Convinces his successors of the importance of networking
- **1965:** Lawrence G. Roberts & Thomas Merrill connect two computers, from MA to CA, over telephone lines
- **1966:** Roberts joins DARPA, develops a plan for the ARPANET
- **1969:** First two sites (UCLA and Stanford Research Institute) are connected to the ARPANET

1969

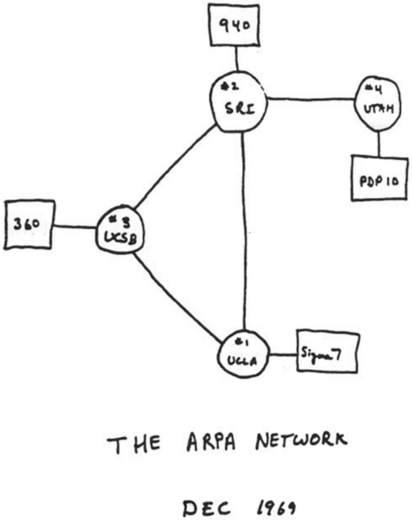
940
#2
SRI
#4
UTAH
PDP 10
360
#3
UCSB
#1
UCLA
Sigma7
THE ARPA NETWORK
DEC 1969

1977

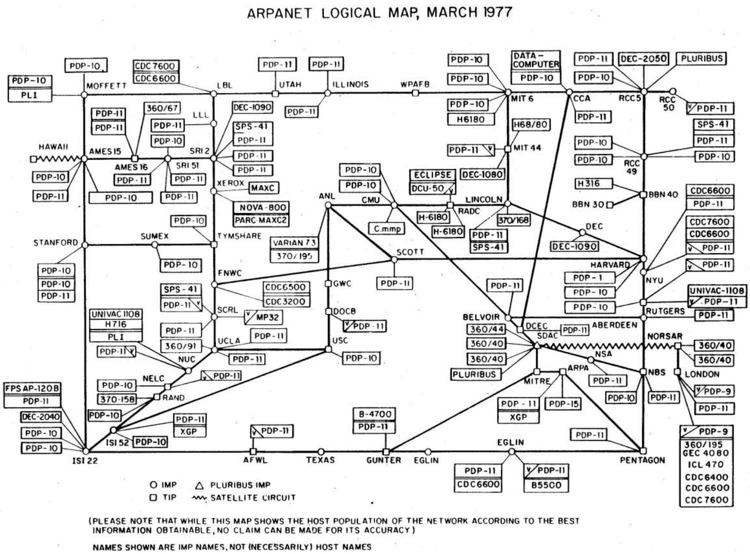
ARPANET LOGICAL MAP, MARCH 1977
PDP-10
PLI
PDP-10
MOFFETT
CDC 7600
CDC6600
LBL
PDP-11
UTAH
PDP-11
ILLINOIS
WPAFB
PDP-10
PDP-10
DATA-COMPUTER
PDP-10
MIT 6
PDP-11
PDP-10
DEC-2050
CCA
RCC5
PLURIBUS
RCC 50
PDP-11
PDP-11
PDP-11
360/67
LLL
DEC-1090
PDP-11
PDP-10
H6180
H68/80
PDP-11
SPS-41
PDP-11
PDP-10
PDP-11
SPS-41
PDP-11
PDP-11
PDP-10
SRI 2
PDP-11
MIT 44
PDP-10
PDP-10
HAWAII
AMES 15
AMES 16
SRI 51
PDP-11
RCC 49
PDP-10
PDP-11
PDP-11
PDP-11
PDP-10
PDP-11
XEROX
MAXC
PDP-10
PDP-10
ECLIPSE
DCU-50
DEC-1080
H316
BBN 40
CDC6600
PDP-11
NOVA-800
PARC-MAXC2
ANL
CMU
LINCOLN
BBN 30
PDP-10
RADC
H-6180
H-6180
370/168
C.mmp
CDC7600
CDC6600
PDP-11
DEC
STANFORD
SUMEX
TYMSHARE
VARIAN 73
370/195
PDP-11
SPS-41
DEC-1090
SCOTT
PDP-10
HARVARD
PDP-11
PDP-10
PDP-10
PDP-11
PDP-10
FNWC
GWC
PDP-11
PDP-1
PDP-10
NYU
PDP-11
SPS-41
CDC6500
CDC3200
UNIVAC-1108
PDP-11
PDP-11
SCRL
MP32
DOCB
BELVOIR
PDP-10
RUTGERS
PDP-11
UNIVAC 1108
H716
PLI
PDP-11
PDP-11
PDP-11
360/91
UCLA
USC
PDP-11
360/44
DCEC
PDP-11
ABERDEEN
NORSAR
SDAC
360/40
360/40
360/40
360/40
NSA
PDP-11
PDP-11
NUC
PDP-10
PLURIBUS
ARPA
NBS
LONDON
NELC
PDP-11
MITRE
PDP-11
FPS AP-120B
PDP-10
PDP-9
PDP-11
370-158
RAND
PDP-10
PDP-11
PDP-15
PDP-10
PDP-11
PDP-11
DEC-2040
PDP-10
PDP-11
XGP
B-4700
XGP
PDP-11
PDP-10
PDP-11
PDP-9
PDP-10
ISI 52
PDP-10
PDP-11
EGLIN
360/195
PDP-10
GEC 4080
ICL 470
ISI 22
AFWL
TEXAS
GUNTER
EGLIN
PENTAGON
CDC6400
CDC 6600
CDC 7600
PDP-11
CDC6600
PDP-11
B5500
○ IMP
□ TIP
△ PLURIBUS IMP
SATELLITE CIRCUIT
(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

2015

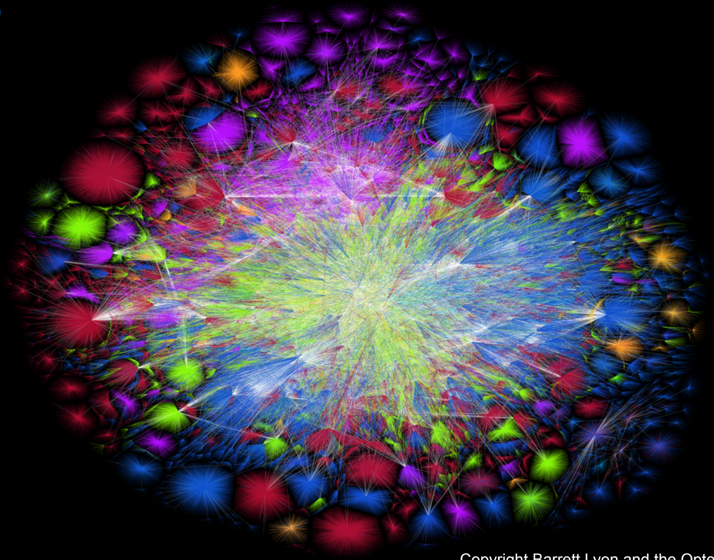

What Is the Internet?

It is a network of networks.

It was designed as standard protocols interconnecting different networks.

IP is the "narrow waist" of the Internet

There are many protocols above and below it.

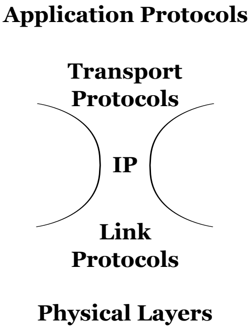

What *Is* a Protocol?

An agreement on how to communicate.

Every protocol consists of:

- Syntax: How things are "written down"
  - How is an integer represented?
  - Where is message length encoded?
  - *etc.*
- Semantics: What things mean
  - Who initiates communication?
  - How is lost information recovered?
  - *etc.*

In a good protocol, these things are precise.

Networks Before the Internet

Prior to the Internet, networks were walled gardens.

IBM machines on one network, Digital on another …

Different platforms and media used different protocols.

Protocols seldom crossed physical media boundaries: coaxial cable, twinaxial cable, twisted pair, *etc.*

Internet Design Goals

To “develop an effective technology for multiplexed utilization of existing interconnected networks.”

At the time it was assumed that there would be other sorts of networks [than the ARPANET] to interconnect, although the local area network had not yet emerged.

1. Be robust to failure of networks or gateways
2. Support multiple types of service
3. Accommodate a variety of networks
4. Permit distributed management
5. Be cost effective
6. Permit easy host connection
7. Be resource accountable

From David Clark’s 1988 SIGCOMM paper (required reading) [1]

Internet Design Solution

Layering — the network layer and up

Existing networks underneath the network layer, and different types of service above.

Each layer can:

- rely on layers below, and
- provide services to layers above.

This provides:

- Encapsulation (layer details are hidden)
- Polymorphism (layer implementations can be replaced)

Layers

We usually recognize five layers:

- Application (HTTP, SMTP, IMAP, XMPP)
- Transport (TCP, UDP)
- Network (IP — v4 or v6)
- Link (Ethernet, Wi-Fi, DOCSIS, LTE)
- Physical (coaxial cable, twisted pair, RF)

This is the standard CS approach to problems: add a layer of indirection!

The OSI Model recognizes 7, but it's a bad retcon[1].

[1] Wikipedia: Retroactive continuity, [or retcon], is a literary device in which established facts [...] are adjusted, ignored, or contradicted by a subsequently published work[.]

But Which Layers?

Where and how to draw the layer boundaries is important.

Obviously on top of the link layer …

Packet switching was selected, with the “bottom” layer handling:

- Addressing
- Routing
- Fragmentation and Framing

This is the network layer, IP (the Internet Protocol!) [2].

Everything else goes on top!

The Network Layer

This division provides a “least common denominator.”

Addressing papers over the disparate local network address formats and semantics.

Routing handles walled gardens that don’t know how to communicate.

Fragmentation and framing solves the problem of link layers with wildly different packet sizes — or no packets at all!

Remaining Problems

We said the Internet solves some distributed systems problems.

But …the network layer didn't. What does?

Remaining Problems

We said the Internet solves some distributed systems problems.

But …the network layer didn't. What does?

Another layer: the transport layer.

Why?

What Layers, Part II

By giving each layer minimal functionality, we:

- reduce complexity
- increase flexibility
- serve multiple purposes

The Network Layer, Part II

Why “network” layer?

Where does it do its work?

(Who handles routing? Who do addresses belong to?)

The Network Layer, Part II

Why “network” layer?

Where does it do its work?

(Who handles routing? Who do addresses belong to?)

…the network! But what does that mean?

We divide duties between:

- the network: routers, gateways, *etc.*
- endpoints: communicating entities

Why Route in the Network?

Remember this?

1. Be robust to failure of networks or gateways
2. Support multiple types of service
4. Permit distributed management

Routing at the endpoints complicates these design goals.

What Layers, Part III

A big question:
Where *should* problems be solved?

A related question:
What sorts of problems *do* we solve?

Should the answers to these questions affect layering?

What Layers, Part III

A big question:
Where *should* problems be solved?

A related question:
What sorts of problems *do* we solve?

Should the answers to these questions affect layering?

It turns out they do.

What Layers, Part IV

What about the other layers? Do they belong to the endpoints?

What Layers, Part IV

What about the other layers? Do they belong to the endpoints?

Find out Next Time!

Summary

- A protocol is an agreement on how to communicate.
  - Syntax & Semantics
- The Internet is, primarily, a network of networks.
- Its durability and longevity are due to good layering.
  - Routing and addressing are handled in the network.

Next time:

- Reliability and Survivability
- Transport protocols

References I

Required Readings

[1] David D. Clark. "The Design Philosophy of the DARPA Internet Protocols". In: *Computer Communication Review* 18.4 (Aug. 1988), pp. 106–114. URL: http://ccr.sigcomm.org/archive/1995/jan95/ccr-9501-clark.pdf.

Optional Readings

[2] Jon Postel. *Internet Protocol*. RFC 791. Sept. 1981. 45 pp. URL: https://www.rfc-editor.org/rfc/rfc791.txt.

Acknowledgements

These slides are based on slides from Steve Ko, used with permission.

Those slides contain the following attribution:

- These slides contain material developed and copyrighted by
  - Indranil Gupta ant UIUC
  - Mike Freedman and Jen Rexford at Princeton