

# Internet design philosophy

CSE 561, Winter 2021

Ratul Mahajan

# What we read

Two foundational papers that *extract* principles from Internet design

- [The Design Philosophy of DARPA Internet Protocols](#)

Clark, 1988

- [End-to-end Arguments in System Design](#)

Saltzer, Reed, and Clark, 1984

These principles were not even articulated, let alone be explicit design goal, when the Internet was being engineered

# Many applications, disparate requirements

File transfer: High throughput, reliability

YouTube: Min throughput, low jitter

Phone call: Low jitter, low latency

Zoom: Low jitter, minimum throughput, low latency

Web: Low latency

Your next idea

# How to design a network that satisfies them all?

## Two meta decisions

- Service model: Datagram or virtual circuits?
- What functionality to put in the network?
  - Reliable delivery
  - Delivery acknowledgement
  - Prevent duplication
  - Guarantee minimum throughput
  - Guarantee latency
  - FIFO
  - Encryption
  - Authentication
  - ...

# Two network service models

- Datagrams or packet switching

- Connectionless service
  - Like postal letters




- Virtual circuits or circuit switching

- Connection-oriented service
  - Like a telephone call




# Datagram model

- Packets contain a destination address; each router uses it to forward packets, maybe on different paths



# Datagram model (2)

- Each router has a forwarding table keyed by address
  - Gives next hop for each destination address; may change




# Virtual circuit model

- Three phases:
  1. Connection establishment, circuit is set up
    - Path is chosen, circuit information stored in routers
  2. Data transfer, circuit is used
    - Packets are forwarded along the path
  3. Connection teardown, circuit is deleted
    - Circuit information is removed from routers
- Just like a telephone circuit, but virtual in that no bandwidth need be reserved; statistical sharing of links


# Virtual circuits

- Packets contain a short label to identify the circuit
  - Labels don't have global meaning, only unique for a link



# Virtual circuits (2)

- Each router has a forwarding table keyed by circuit
  - Gives output line and next label to place on packet




# Datagrams vs virtual circuits

Issue	Datagrams	Virtual Circuits
Setup phase	Not needed	Required
Router state	Per destination	Per connection
Addresses	Packet carries full address	Packet carries short label
Forwarding	Per packet	Per circuit
Quality of service	Difficult to add	Easier to add
Masking failures	Easy	Hard

Goal1: Internet communication must continue despite loss of networks or gateways.

# Masking failures with datagrams



Routers have no connection state, so datagrams can be easily re-routed

- Possible but difficult with VCs

Hosts have connection state

- “Fate sharing”

Moot point when there are single-points of failure in the network

# What functionality to put in the network?

Internet: As little as possible

- Hard to think of a simpler network

Goal 2: The Internet must support multiple types of communication service.

Functionality in the network gets in the way of services that do not need it

Goal 3: The Internet architecture must accommodate a variety of networks.

Some networks may not be capable of providing expected functionality

End-to-end argument

Generally not possible for network to meet exact application requirements

These are different concerns

# The cost of in-network functionality

Consider reliability and common ways of providing it

1. Cache the packet and resend if it is not acknowledged
  - Needs extra memory and some compute in routers
2. Send multiple times
  - High network overhead
3. Error coding
  - Simple example:
    - Send  $A \oplus B$  in addition to sending A and B.
    - Recover A using B and  $A \oplus B$
  - Needs extra computation at the sender and memory at the receiver

# Were the other Internet goals met?

**Goal 4:** The Internet architecture must permit distributed management of its resources.

**Goal 5:** The Internet architecture must be cost effective.

**Goal 6:** The Internet architecture must permit host attachment with a low level of effort.

**Goal 7:** The resources used in the internet architecture must be accountable.

# Limitations of Internet's architecture

## Security

- DoS attacks – you can send an arbitrary amount, anonymously
- Phishing – identity spoofing
- Prefix hijacking

Hard to support highly demanding applications

Suboptimal efficiency and performance

## Privacy

# Discussion

Were those the right design goals for that time?

What should be the design goals now?

# End-to-End principle


# End-to-end Principle

- Broad networking principle
  - First implementation in French CYCLADES network (after ARPA) (1970)
  - Articulated in its most recognizable form by Saltzer, Reed, Clark (1981) [[paper](#)]
- Guidance on placing functionality such as reliability, security, etc.—in network or at endpoints (hosts)?
  - Argues for endpoint placement

# Multiple interpretations of the principle

- The network cannot be trusted. Do it yourself.
  - The network can suffer heavy damage
    - Nuclear attacks (but not DDoS attacks!)
  - Need end-to-end correctness anyway
- Diminishing returns from in-network functionality
  - Not everyone needs it
- Place functionality in the network only when necessary
  - E.g., for performance


# E2E Example: Error-correcting codes



IP:  
Host detects  
errors

802.11:  
Link detects errors


# E2E Example: ARQ



TCP:  
Host retransmits  
on failure

802.11:  
Link detects drops  
and retransmits


# E2E Example: In-order delivery



TCP:  
Host enforces in-order delivery

SS7:  
Network enforces in-order delivery

# E2E Example: Security



SSL:  
Host encrypts  
content

GSM:  
Network encrypts  
content

# End-to-End limitations


- Some functionality cannot be implemented at endpoints
  - NATs, DoS protection, ... the principle is silent on these
- Assumes a clear dividing line between network and endpoints
  - Reality of distributed applications (e.g., CDNs) is more complex
- No guidance on how much functionality can go in the network for performance

# What is the opposite of e2e argument?


1. Network with rich functionality that covers most requirements
  - E.g., phone network
  - Practical for a data network?
2. Network with multiple "lanes"?
  - CISC-like
  - Slight aside: Can applications be trusted?
3. Modular network
  - Applications mix-n-match what they need

# Your thoughts (1)

User-level stacks, kernel-bypass, abstractions ...



Monolithic applications – bundle everything



Modularity based on layers or libraries

# Your thoughts (2)

What about throughput?

Reliability has a throughput cost too

# Your thoughts (3)

Show me the data (for resource usage)

Assume that

- Error rate: 5%
- If fraction of applications that want reliability is 100%
  - With end-host recovery, 105 packets per 100 data packets
  - With network recovery, 105 packets ....
- If fraction of applications that want reliability is 10%
  - With end-host recovery, 100.5 packets per 100 data packets
  - With network recovery, 105 packets ...

Do you know the numbers for all networks and in the future?

# Summary

The Internet is but one possible design for a large internetwork

The design is optimized for

- Fault tolerance
- Diversity of applications and networks
- Keep network simple, push complexity to the ends

(Future classes will tinker with this in various ways.)

# Next class: Object lookup

Name to location mapping at Internet scale

- DNS – hierarchical names cs.washington.edu
- DHTs – flat names