


## 15-440 Distributed Systems

Lecture 2 – 15-441 in 2 Days


1



## Distributed Systems vs. Networks

- Low level (c/go)
- Run forever
- Support others
- Adversarial environment
- Distributed & concurrent
- Resources matter
- And have it implemented/run by vast numbers of different people with different goals/skills


2



## Keep an eye out for...

- Modularity, Layering, and Decomposition:
  - Techniques for dividing the work of building systems
  - Hiding the complexity of components from each other
  - Hiding implementation details to deal with heterogeneity
- Naming/lookup/routing
- Resource sharing and isolation
- Models and assumptions about the environment and components


3




## Today's Lecture

- Network links and LANs
- Layering and protocols
- Internet design

4




## Basic Building Block: Links



- Electrical questions
  - Voltage, frequency, ...
  - Wired or wireless?
- Link-layer issues: How to send data?
  - When to talk – can either side talk at once?
  - What to say – low-level format?

5




## Model of a communication channel

- Latency - how long does it take for the first bit to reach destination
- Capacity - how many bits/sec can we push through? (often termed "bandwidth")
- Jitter - how much variation in latency?
- Loss / Reliability - can the channel drop packets?
- Reordering


6

## Basic Building Block: Links

- ... But what if we want more hosts?



One wire




Wires for everybody!

- Scalability?!

7

## Multiplexing


- Need to share network resources



- How? Switched network
  - Party "A" gets resources sometimes
  - Party "B" gets them sometimes
- Interior nodes act as "Switches"
- What mechanisms to share resources?

8


## In the Old Days...Circuit Switching



9


## Packet Switching

- Source sends information as self-contained packets that have an address.
  - Source may have to break up single message in multiple
- Each packet travels independently to the destination host.
  - Switches use the address in the packet to determine how to forward the packets
  - Store and forward
- Analogy: a letter in surface mail.



10

## Packet Switching – Statistical Multiplexing




- Switches arbitrate between inputs
- Can send from *any* input that's ready
  - Links never idle when traffic to send
  - (Efficiency!)

11

## What if Network is Overloaded?

Problem: Network Overload




Solution: Buffering and Congestion Control

- Short bursts: buffer
- What if buffer overflows?
  - Packets dropped
  - Sender adjusts rate until load = resources → "congestion control"

12

## Example: Ethernet Packet

- Sending adapter encapsulates IP datagram (or other network layer protocol packet) in **Ethernet frame**



13

## Ethernet Frame Structure

- Each protocol layer needs to provide some hooks to upper layer protocols
  - Demultiplexing: identify which upper layer protocol packet belongs to
  - E.g., port numbers allow TCP/UDP to identify target application
  - Ethernet uses Type field
- **Type: 2 bytes**
  - Indicates the higher layer protocol, mostly IP but others may be supported such as Novell IPX and AppleTalk

14


## Ethernet Frame Structure (cont.)

- **Addresses: 6 bytes**
  - Each adapter is given a globally unique address at manufacturing time
    - Address space is allocated to manufacturers
      - 24 bits identify manufacturer
      - E.g., 0:0:15:\* → 3com adapter
    - Frame is received by all adapters on a LAN and dropped if address does not match
  - Special addresses
    - Broadcast – FF:FF:FF:FF:FF:FF is “everybody”
    - Range of addresses allocated to multicast
      - Adapter maintains list of multicast groups node is interested in

15


## Packet Switching

- Source sends information as self-contained packets that have an address.
  - Source may have to break up single message in multiple
- **Each packet travels independently to the destination host.**
  - Switches use the address in the packet to determine how to forward the packets
  - Store and forward
- Analogy: a letter in surface mail.



16

## Frame Forwarding




MAC Address	Port	Age
A21032C9A591	1	36
99A323C30842	2	01
8711C78900AA	2	15
201B2369011C	2	16
695519801198	3	11

- A machine with MAC Address lies in the direction of number port of the bridge
- For every packet, the bridge “looks up” the entry for the packets destination MAC address and forwards the packet on that port.
  - Other packets are broadcast – why?
- Timer is used to flush old entries

17

## Learning Bridges

- Manually filling in bridge tables?
  - Time consuming, error-prone
- Keep track of source address of packets arriving on every link, showing what segment hosts are on
  - Fill in the forwarding table based on this information



18


## Today's Lecture

- Network links and LANs
- Layering and protocols
- Internet design

19

## Internet

- An inter-net: a network of networks.
  - Networks are connected using routers that support communication in a hierarchical fashion
  - Often need other special devices at the boundaries for security, accounting, ...
- The Internet: the interconnected set of networks of the Internet Service Providers (ISPs)
  - About 17,000 different networks make up the Internet




20

## Challenges of an internet

- Heterogeneity
  - Address formats
  - Performance – bandwidth/latency
  - Packet size
  - Loss rate/pattern/handling
  - Routing
  - Diverse network technologies → satellite links, cellular links, carrier pigeons
  - In-order delivery

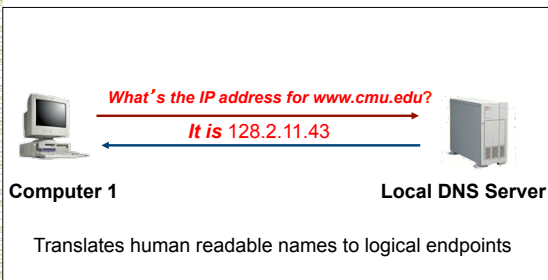
21

## How To Find Nodes?



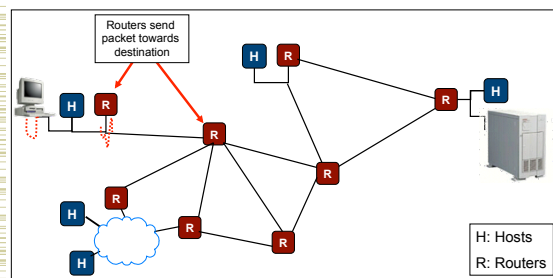
22

## Naming



23

## Routing



24

## Network Service Model



- What is the *service model*?
  - Ethernet/Internet: *best-effort* – packets can get lost, etc.
- What if you want more?
  - Performance guarantees (QoS)
  - Reliability
    - Corruption
    - Lost packets
  - Flow and congestion control
  - Fragmentation
  - In-order delivery
  - Etc...

25

## Failure models



- Fail-stop:
  - When something goes wrong, the process stops / crashes / etc.
- Fail-slow or fail-stutter:
  - Performance may vary on failures as well
- Byzantine:
  - Anything that can go wrong, will.
  - Including malicious entities taking over your computers and making them do whatever they want.
- These models are useful for proving things;
- The real world typically has a bit of everything.
  
- Deciding which model to use is important! 2

26

## Model Example: project 1



- Project 1: Build a bitcoin miner
- Server --- many clients
- Communication:
  - Send job
  - ACK job
  - do some work
  - send result to server
  - (repeat)
- IP communication model:
  - Messages may be lost, re-ordered, corrupted (we'll ignore corruption, mostly, except for some sanity checking)
- Fail-stop node model:
  - You don't need to worry about evil participants faking you out.

27

## Fancier Network Service Models



- What if network had reliable, in-order, mostly no-corruption, stream-oriented communication (i.e. TCP)
- Programmers don't have to implement these features in every application
- But note limitations: this can't turn a byzantine failure model into a fail-stop model...

28


## What if the Data gets Corrupted?



Problem: Data Corruption



Solution: Add a *checksum*



29

## What if the Data gets Lost?



Problem: Lost Data



Solution: Timeout and Retransmit



30

### What if the Data is Out of Order?

Problem: Out of Order

Solution: Add Sequence Numbers

31

### Networks [including end points] Implement Many Functions

- Link
- Multiplexing
- Routing
- Addressing/naming (locating peers)
- Reliability
- Flow control
- Fragmentation
- Etc....

32

### What is Layering?

- Modular approach to network functionality
- Example:

33

### What is Layering?

Modular approach to network functionality

34

### Layering Characteristics

- Each layer relies on services from layer below and exports services to layer above
- Interface defines interaction with peer on other hosts
- Hides implementation - layers can change without disturbing other layers (black box)

35

### What are Protocols?

- An agreement between parties on how communication should take place
- Module in layered structure
- Protocols define:
  - Interface to higher layers (API)
  - Interface to peer (syntax & semantics)
    - Actions taken on receipt of a messages
    - Format and order of messages
    - Error handling, termination, ordering of requests, etc.
- Example: Buying airline ticket

36

## IP Layering

- Relatively simple

Host    Bridge/Switch    Router/Gateway    Host

37

## The Internet Protocol Suite

The waist facilitates interoperability

38

## Layer Encapsulation

User A    User B

39

## Multiplexing and Demultiplexing

- There may be multiple implementations of each layer.
  - How does the receiver know what version of a layer to use?
- Each header includes a demultiplexing field that is used to identify the next layer.
  - Filled in by the sender
  - Used by the receiver
- Multiplexing occurs at multiple layers. E.g., IP, TCP, ...

V/HL	TOS	Length
ID	Flags/Offset	
TTL	Prot.	H. Checksum
Source IP address		
Destination IP address		
Options..		

40

## Protocol Demultiplexing

- Multiple choices at each layer


41

## Today's Lecture

- Network links and LANs
- Layering and protocols
- Internet design

42


## Goals [Clark88]



- 0 Connect existing networks**  
initially ARPANET and ARPA packet radio network
- 1. Survivability**  
ensure communication service even in the presence of network and router failures
- 2. Support multiple types of services**
3. Must accommodate a variety of networks
4. Allow distributed management
5. Allow host attachment with a low level of effort
6. Be cost effective
7. Allow resource accountability

43


## Goal 0: Connecting Networks



- How to interconnect various network technologies
  - ARPANET, X.25 networks, LANs, satellite networks, packet networks, serial links...
- Many differences between networks
  - Address formats
  - Performance – bandwidth/latency
  - Packet size
  - Loss rate/pattern/handling
  - Routing

44


## Gateway Alternatives



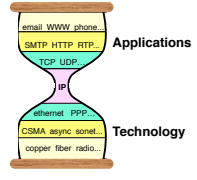
- Translation
  - Difficulty in dealing with different features supported by networks
  - Scales poorly with number of network types ( $N^2$  conversions)
- Standardization
  - “IP over everything” (**Design Principle 1**)
  - Minimal assumptions about network
  - Hourglass design

45

## IP Hourglass




- Need to interconnect many existing networks
- Hide underlying technology from applications
- Decisions:
  - Network provides minimal functionality
  - “Narrow waist”



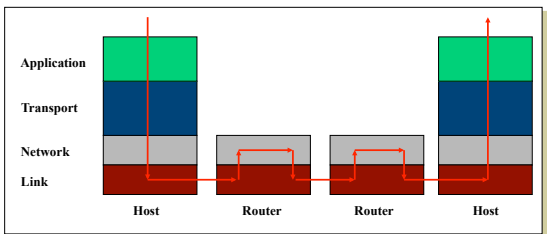
**Tradeoff: No assumptions, no guarantees.**

46

## IP Layering (Principle 2)




- Relatively simple
- Sometimes taken too far



47

## Goal 1: Survivability



- If network is disrupted and reconfigured...
  - Communicating entities should not care!
  - No higher-level state reconfiguration
- How to achieve such reliability?
  - Where can communication state be stored?

	Network	Host
Failure handing	Replication	“Fate sharing”
Net Engineering	Tough	Simple
Switches	Maintain state	Stateless
Host trust	Less	More

48