Introduction to Computer Networks

Spanning Tree and Ethernet

CS640

https://pages.cs.wisc.edu/~mgliu/CS640/F22/

Ming Liu
mgliu@cs.wisc.edu

Today

Last lecture

- How do frames traverse NICs/bridges?

Today

- How to avoid loops in the L2 switching?
- Why Ethernet dominates?

Announcements

- Lab2 due next Tuesday
- Quiz2 this Thursday

Question 2. [5 points] Hardware infrastructure

(a). [**1.5 points**] Suppose a NIC has 2 ports, where it can forward traffic between these two ports. A host only has 1 NIC. Can we build a 10 hosts network where all hosts are fully connected (i.e., there is a connection between any two hosts)? If so, please give an example.

(b). [**1.5 points**] The fat tree topology has been widely used to construct data center networks. Below is an example. How many communication paths are between NIC1 and NIC4?

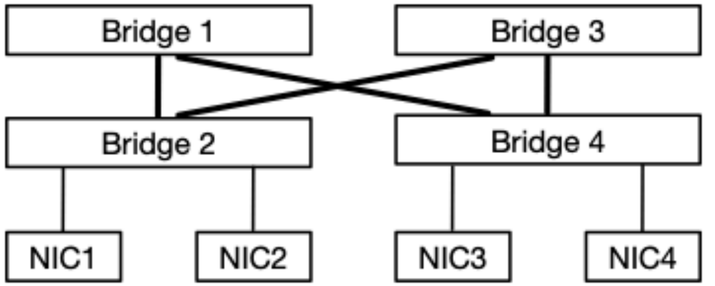

(c). [**2 points**] Given the above network topology, how many communication paths are between NIC1 and NIC2?

Q: How to avoid loops in the L2 switching?

Q: How to avoid loops in the L2 switching?

A: Spanning Tree

- Radia Perlman from the Digital Equipment Corporation
- IEEE 802.1 specification

Q1: How does the spanning tree algorithm work?

Q1: How does the spanning tree algorithm work?

A: Each switch chooses the ports to forward frames such that one reduces a topology graph to a tree. Three steps on a switch:

Q1: How does the spanning tree algorithm work?

A: Each switch chooses the ports to forward frames such that one reduces a topology graph to a tree. Three steps on a switch:

- #1: Select a root —> Use the switch ID
- #2: Decide the shortest path to the root —> Use the switch ID to break the tie
- #3: Configure the designated ports (and associated switches) for forwarding —> No blindly broadcast anymore

Q2: How does the spanning tree protocol work?

Networking Protocol

L2:

Protocol

An abstract object or module in a layered structure

- Vertical view: an interface to high-level protocols
- Horizontal view: a peer interface to a counterpart

Networking Protocol

Distributed logic that runs over one or more communication entities for certain functionalities

- #1: What states are maintained at each entity?
- #2: How do different entities coordinate?

Networking Protocol

Distributed logic that runs over one or more communication entities for certain functionalities

- #1: What states are maintained at each entity?
- #2: How do different entities coordinate?

Key principle: minimal states

- Anytime and anywhere connectivity -> highly scalable systems

Q2-1: What states are maintained at each entity?

Q2-1: What states are maintained at each entity?

A: A switch maintains the following per-switch/ per-port states:

- #1: Local switch ID
- #2: The switch ID of the root
- #3: The distance (i.e., the number of hops) to the root
- #4: Per-port action table

Q2-2: How do different entities coordinate?

Q2-2: How do different entities coordinate?

A: The configuration message, (Y, d, X)

- Y: the root switch ID in my view
- d: the distance to the root (i.e, the number of hops)
- X: my local switch ID

Spanning Tree Discussion

The resulting effect is blocking a port from forwarding

- A graph –> a tree

The protocol runs periodically

- When a switch recovers from a failure, it starts from the scratch
- The states of a switch are updated when the tree structure changes

Q: Why Ethernet dominates?

Q1: How to identify a frame from bit streams? => Framing
Q2: How to handle transmission errors? => Error handling
Q3: How do frames traverse NICs/bridges? => L2 switching
Q4: How to coordinate the NIC on two sides? => Flow control
Q5: How to orchestrate concurrent transmissions? => Access control

Q: Why Ethernet dominates?

Q: Why Ethernet dominates?

A: Easy-to-manage,

- #1: No central coordination

Q: Why Ethernet dominates?

A: Easy-to-manage, cheap

- #1: No central coordination
- #2: Coaxial cable, tap, and adapter

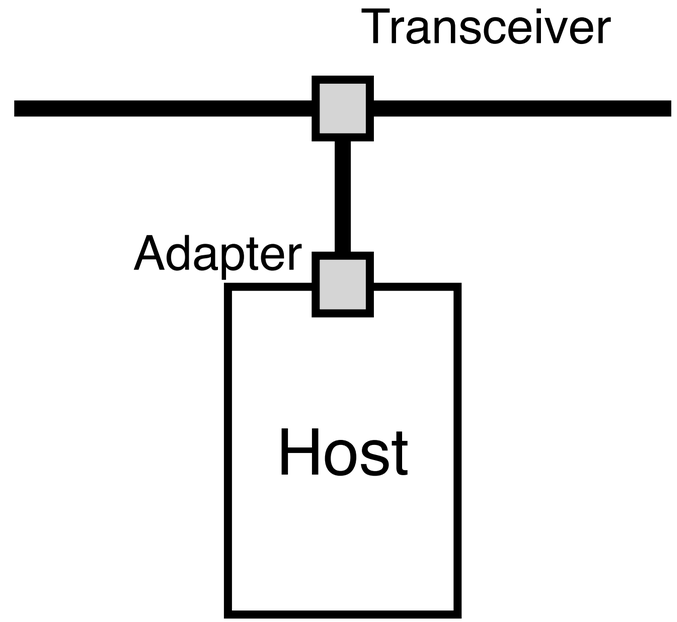

Q: Why Ethernet dominates?

A: Easy-to-manage, cheap, simple multiplexing

- #1: No central coordination
- #2: Coaxial cable, tap, and adapter
- #3: Carrier Sense, Multiple Access with Collision Detection (CSMA/CD)

Ethernet Hardware in the Early Days

An Ethernet segment could support 500m originally

- No more than four repeaters could be positioned between two hosts ==> up to 2500m

Network repeater

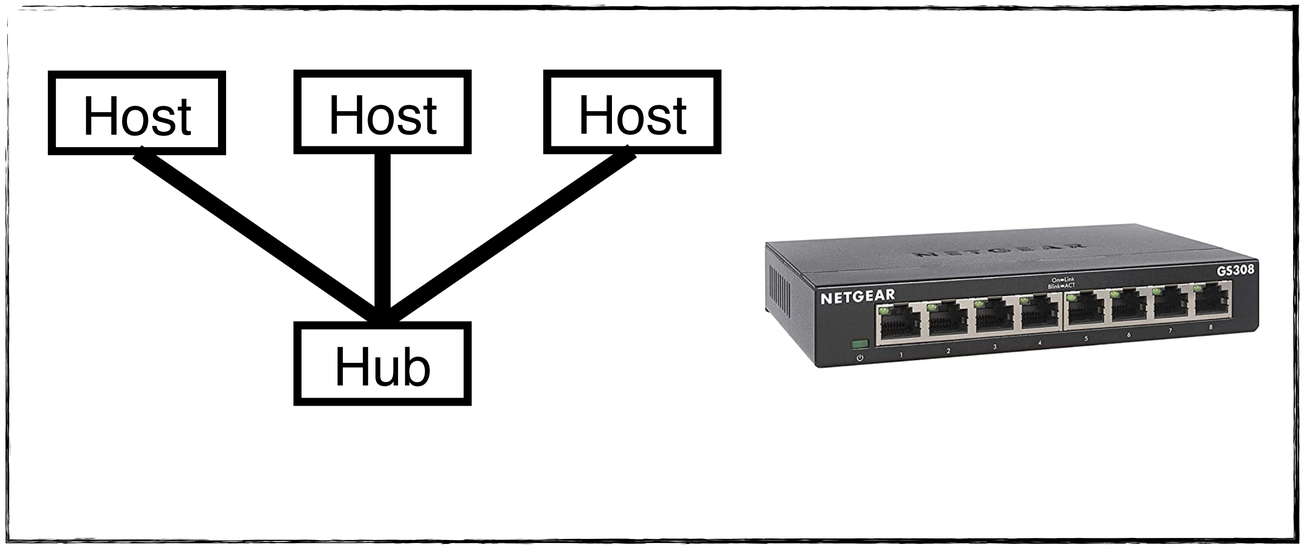

Network Hub

Key Idea: Access the Channel Randomly

#1: Carrier Sense (CS) — the transmitter senses the state of the communication carrier

- Idle or busy

#2: Multiple Access (MA) — the transmitter is allowed to send the frame with a probability p

- No coordination
- The transmitter makes the decision at the beginning of each time slot based on the result of CS

Key Idea: Access the Channel Randomly

#3: Collision Detection (CD) — When a collision happens, the transmitter aborts immediately, performs exponential back-off, and sends the frame again

- Postpone the transmission by an **internal T**
- The length of the **interval T** increases with **every collision, i.e., T = 2^(i -1) * X, where i is the number of retries**
- **X** was configured as **51.2us** originally, given the maximum 2500 Ethernet

Ethernet Discussion

CSMA/CD is mainly implemented at the sender

CSMA/CD enabled Ethernet segments are replaced by the switch-based Ethernet

- #1: Incremental deployment
- #2: Switching fabric cost drops significantly

CSMA/CD like mechanisms have been applied in the wireless network

Terminology

1. Host
2. NIC
3. Multi-port I/O bridge
4. Protocol
5. RTT
6. Packet
7. Header
8. Payload
9. BDP
10. Baud rate
11. Frame/Framing
12. Parity bit
13. Checksum
14. Ethernet
15. MAC
16. (L2) Switch
17. Broadcast

Principle

1. Layering
2. Minimal States

Technique

1. NRZ Encoding
2. NRZI Encoding
3. Manchester Encoding
4. 4B/5B Encoding
5. Byte Stuffing
6. Byte Counting
7. Bit Stuffing
8. 2-D Parity
9. CRC
10. MAC Learning
11. Store-and-Forward
12. Cut-through
13. Spanning Tree
14. CSMA/CD

Summary

Today's takeaways

#1: The spanning tree protocol reduces a topology graph to a tree by blocking certain forwarding ports

#2: Ethernet becomes popular due to the simple multiplexing support (i.e., CSMA/CD) as well as the manageability and cost benefits

Next lecture

- L2 flow control