Networking Technologies and Applications

Rolland Vida
BME TMIT

September 19, 2017

Aloha

- Hawaii – no telephone network at the end of the 70's
  - How to connect the computers on different islands to a network, containing the central computer in Honolulu?
- Solution: **ALOHANET** – low range radio
  - Norman Abramson, University of Hawaii
  - Each user terminal had a small radio device
    - Operated on two frequencies
      - One for the downstream, one for the upstream traffic
    - Downstream data broadcasted by the central computer, no problem
    - Contention on the upstream channel
      - If data reached correctly the central computer, it retransmitted it on the downstream channel
      - If the original sender did not receive it back, it was probably lost
        - » Retransmission required
  - If low upstream traffic, the solution is quite efficient
  - If higher traffic, the solution is unusable

Aloha

Advantages:

- Different size packets
- No need for synchronization
- Simple operation

If low upstream traffic, the solution is quite efficient

If higher traffic, the solution is unusable

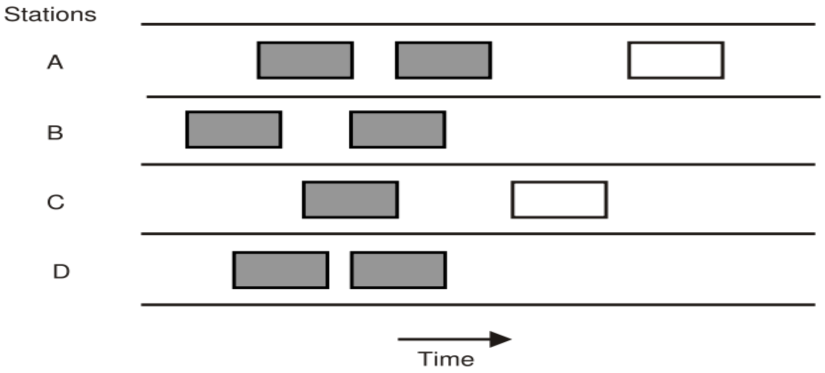

Slotted Aloha

- Time is divided into **slots**
  - Fixed length slots to transmit fixed size packets
- If a node wants to send, waits until a new slot begins
  - Need to **synchronize** the nodes
- If nobody else sends in the same slot, then the sending is sucessful
  - Otherwise collision, the packet is resent after waiting for a random amount of time

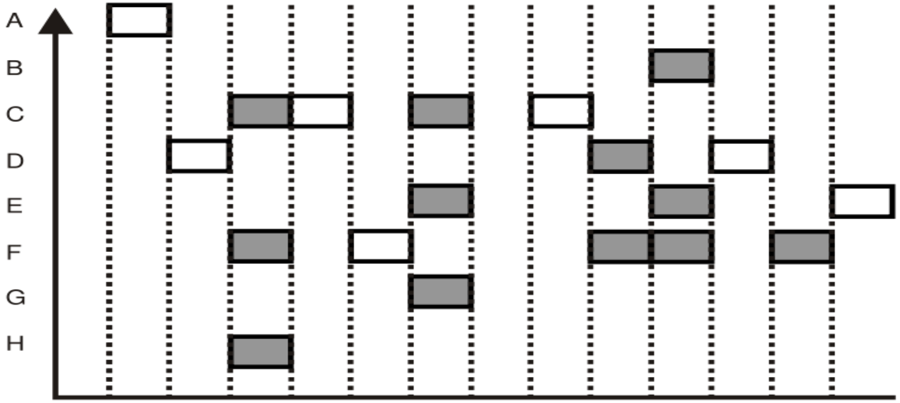

Slotted ALOHA protocol (shaded slots indicate collision)

Ethernet

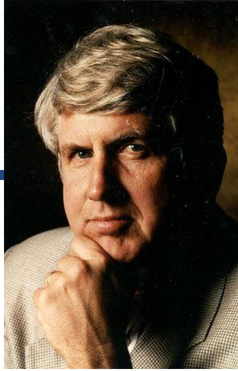

- Bob Metcalfe (MIT, Harvard) spends his holiday together with Abramson on Hawaii
  - Idea: let's do something similar, but for a wired network (1973)
  - First standard (DEC, Intel, Xerox) in 1982, IEEE 802.3 standard in 1983
- Many stations connected to the same cable segment
  - Everyone hears everyone, but without any central node
  - The Ethernet frames will have to have a destination address
    - Everyone hears it, but the frame will be processed only by the destination

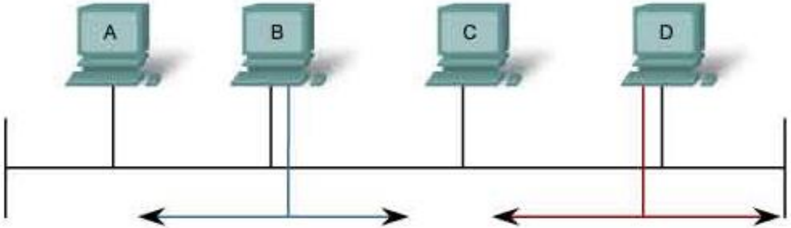

Ethernet = CSMA/CD

- Before transmission, hosts on the same Ethernet cable first listen to the channel (**CSMA – Carrier Sense Multiple Access**)
  - If busy, they wait for the transmission to end
  - If free, start sending
    - Not immediately, but after a **slot time**
      - If there is a signal on the channel, it leaves time for it to be received
    - Slot time = maximum round-trip time on the cable
      - For 10 Mb/s Ethernet it is 51,2 μs, for 100 Mb/s is 5,12 μs
- Two stations might think in parallel that the channel is free
  - Both start sending, a collision occurs
  - If collision, they detect it (**CD – Collision Detection**), and send a jam signal to ensure that others detect the collision as well

Ethernet = CSMA/CD

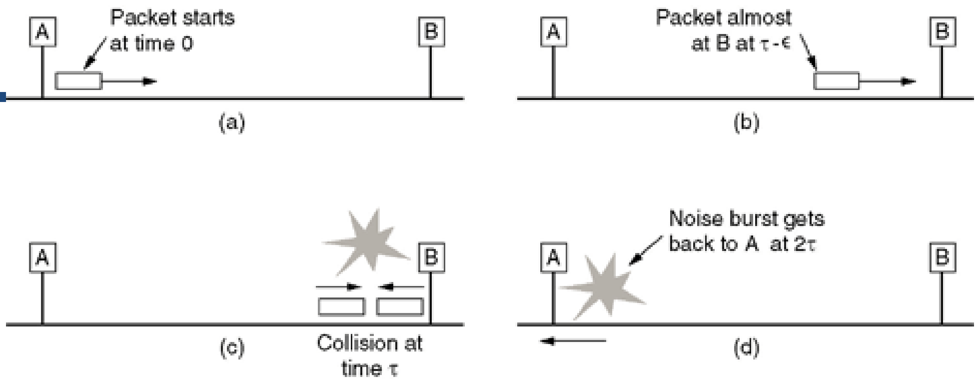

- Wait for a random time interval and retry afterwards
  - Set a timer to a random value from the [0, 1, .. , 2^{m-1}] x t_{av} interval, where t_{av} is the default waiting time (51,2 μs), m = min (10,n) and n is the number of collisions.
  - After each collision the maximum waiting time is doubled, until reaching an upper limit - **truncated binary exponential backoff**
- CSMA/CD not possible on Alohanet
  - Two users on the remote islands could not hear each other

Encapsulation

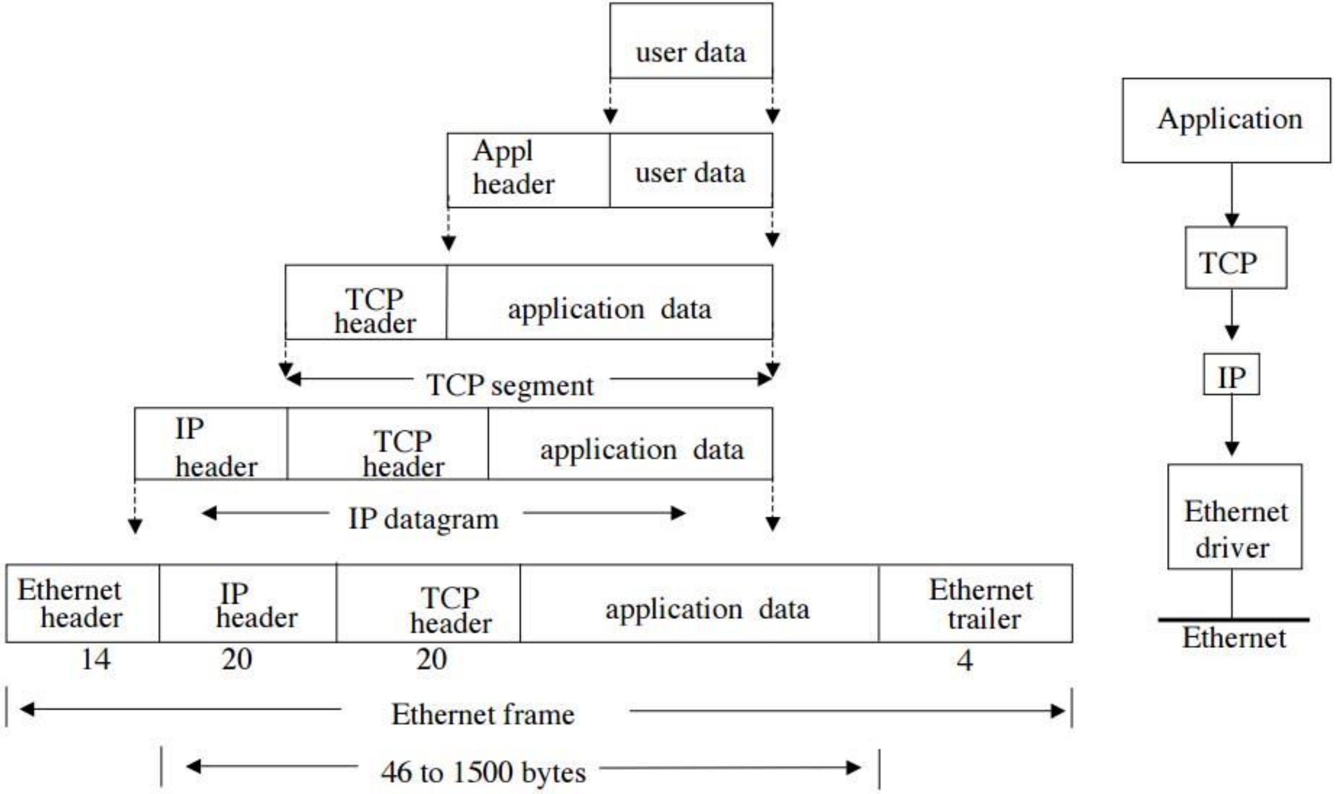

Ethernet Frame

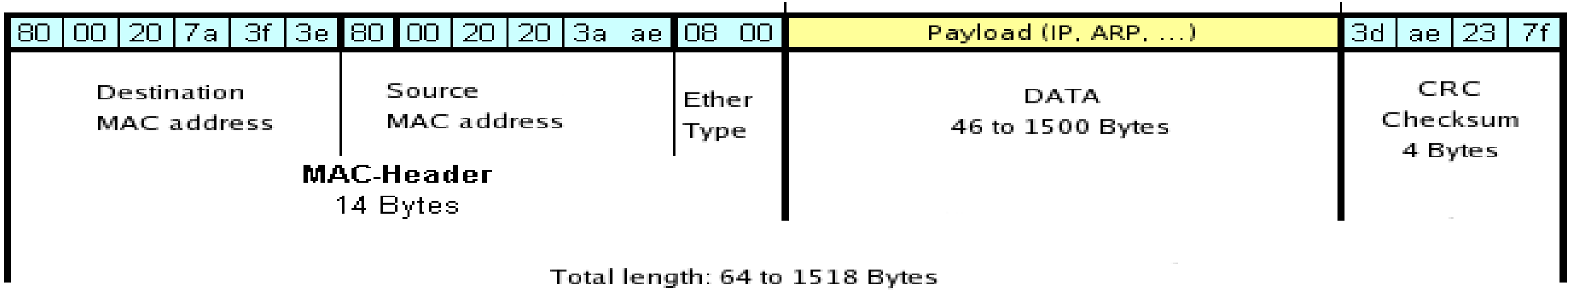

- The DATA field is at most **1500 bytes** (**MTU – Maximum Transmission Unit**)
  - If frame too large, it occupies the channel for long time
  - Higher possibility of an error, you will need to resend large frames
- Minimum length **46 bytes** (minimal frame size = 64 bytes = 512 bits)
  - If the frame too small, collision detection cannot be used
    - Transmission is terminated very fast, before the first bit reaching the end of the cable
    - 51,2 µs round trip time / 0,1 µs bit time = 512 bit
  - Even if there's a collision, the sender is not informed about it
  - Packets that are too small are filled with **padding data** (bits with no utility)

Carrier Extension

- If the speed of the network increases...
  - Either increase the minimum frame size...
  - Or decrease the maximum cable length
    - On a 2500 m cable, for a 1 Gb/s speed, the minimum frame size is 6400 bytes
    - If the minimum size is 640 bytes, the cable can be only 250 m long
  - Very annoying restrictions on a Gigabit network
    - Minimum frame size increased to 512 bytes

- **Carrier Extension**
  - The sender puts the useless bits after the CRC field
  - The receiver cuts it of, not included in the CRC
    - Still a serious waste of capacity
- **Frame Bursting**
  - During a single transmission, several consecutive frames transmitted
  - Increases efficiency considerably

First Ethernet versions

- 10Base5 – thick Ethernet
  - Coaxial cable, 10 Mb/s, 500 m long segments

- 10Base2 – thin Ethernet
  - Coaxial cable, 200 m segments

- 10Base-T
  - Twisted pair, star topology around a hub, 100 m segments

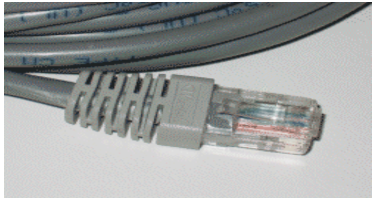

- 10Base-F
  - Optical fiber, 2km long segments

Fast Ethernet

- 100Base-TX
  - 2 twisted pairs
  - One for the upstream, one for the downstream, 100 Mb/s duplex speed
- 100Base-T4
  - 4 twisted pairs
  - One for the upstream, one for the downstream data, the other two can be used as needed
  - Maximum 100 m long segments
- 100Base-FX
  - Multi-mode fiber in both directions
  - 100 Mb/s duplex speed
  - Maximum 2 km between the hub and the stations

Gigabit Ethernet

- IEEE 802.3z (1998), 802.3ab (1999)
- Only point-to-point setups
  - No shared segments, as in traditional 10 Mb/s Ethernet
- Two operation modes:
  - Duplex – traffic in both directions in the same time
    - A central **switch** links the stations on the periphery
    - All the connections are buffered
      - Any station can send data at any time
      - No need to sense the channel, no contention
      - No need for CSMA/CD, not really Ethernet anymore
  - Half-duplex
    - Stations are connected to a simple hub
    - No buffering, collisions are possible

Gigabit Ethernet

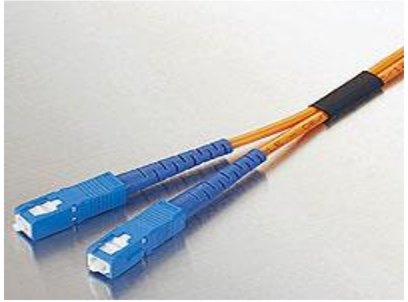

- Different versions
  - 1000Base-SX
    - Multi-mode fiber
    - Maximum 550 m long segments
  - 1000Base-LX
    - Single- or multi-mode fiber
    - Maximum 5000 m long segments
  - 1000Base-T
    - 4 pairs of Cat. 5 UTP cables
    - Maximum 100 m long segments
- IEEE 802.3ae – 10 Gb/s Ethernet (2002)
  - Only on optical cables
- IEEE 802.3ba – 40Gb/s and 100 Gb/s Ethernet (2010)
  - Lucent Technologies Bell Labs experimental results
  - Standard adopted in June 2010

Hub

- Physical layer repeater device
  - Repeats the packet on bit level
  - An incoming packet is forwarded immediately on all the other interfaces
    - Everyone receives all the packets
- If many simultaneous transmissions - collision
  - The „collision domain” is not changed
- Usually a hierarchical, tree-like hub topology

Hub – advantages and drawbacks

- Each station can collide with any other station on the hub
  - Lowers the efficiency of the network
  - Lowers scalability
  - Anyone can see anyone's traffic
- Different Ethernet versions cannot be joined in the same network
  - If one 10Mbps station in the network, the entire network switches back to 10 Mbps operation mode

Hub

- Not efficient to build a large network using only hubs
  - One large collision domain

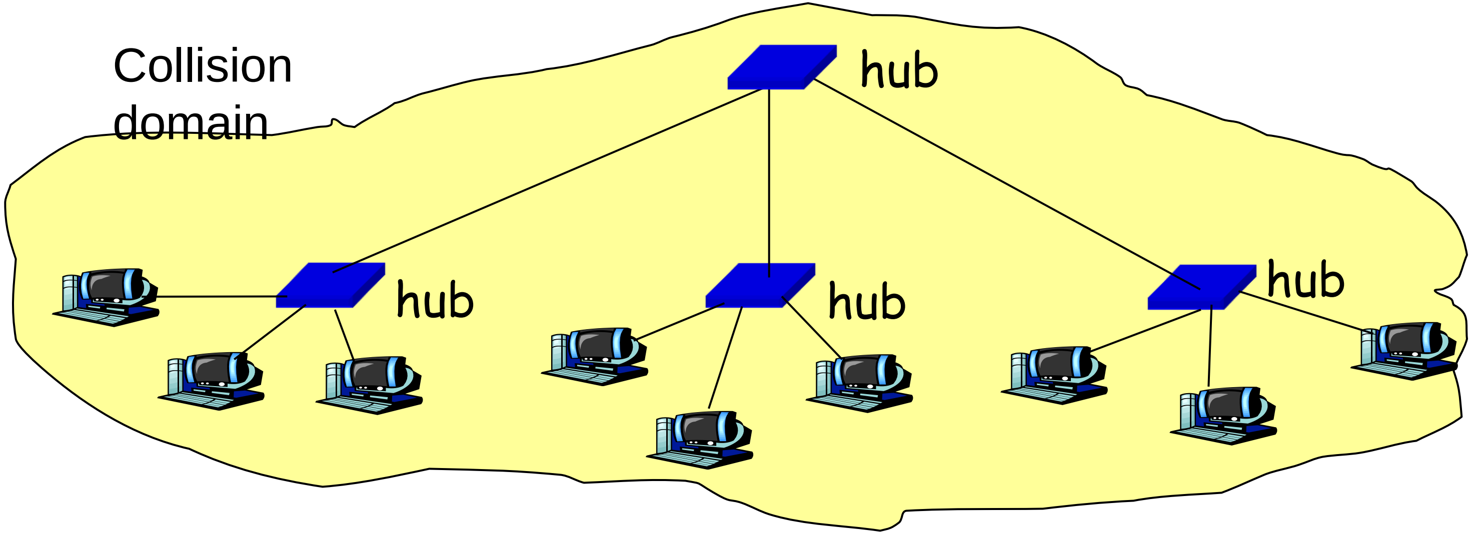

Switch (bridge)

- Link Layer device
  - Checks the MAC header, and forwards selectively
    - switch table: (MAC address, interface, timer)
      - Built from the received packets
      - If one address is unknown, the packet is forwarded to all the interfaces
  - Separates the collision domains
    - Buffers the packets
    - Forwards them only to the appropriate segments

Switch

- Advantages:
  - Higher scalability
  - More efficient, more secure
  - Buffering and switching tables makes the connection of different Ethernet versions possible inside the same network

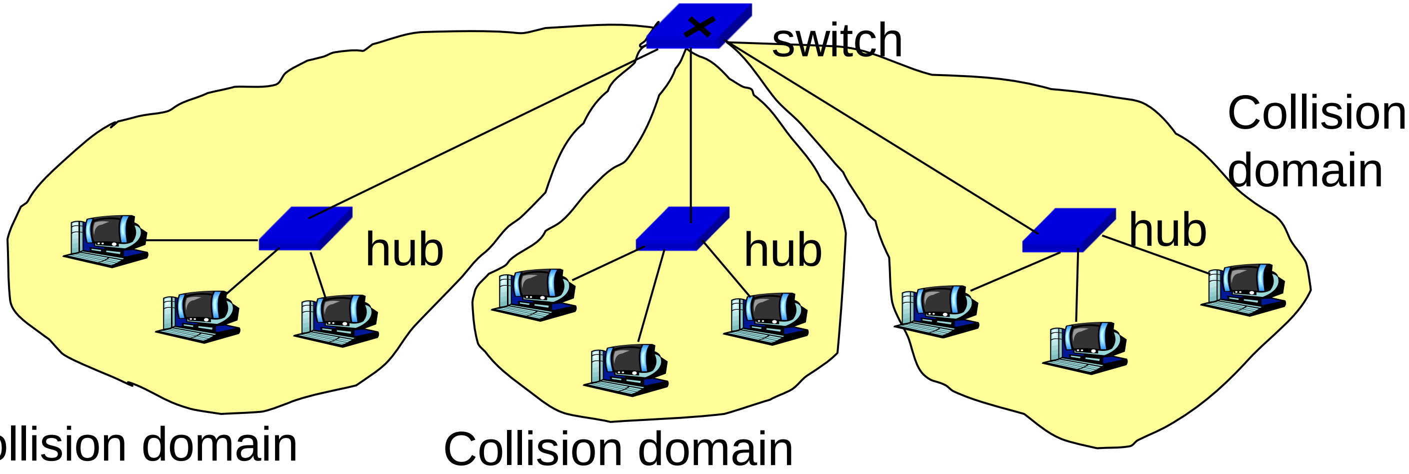

Corporate network

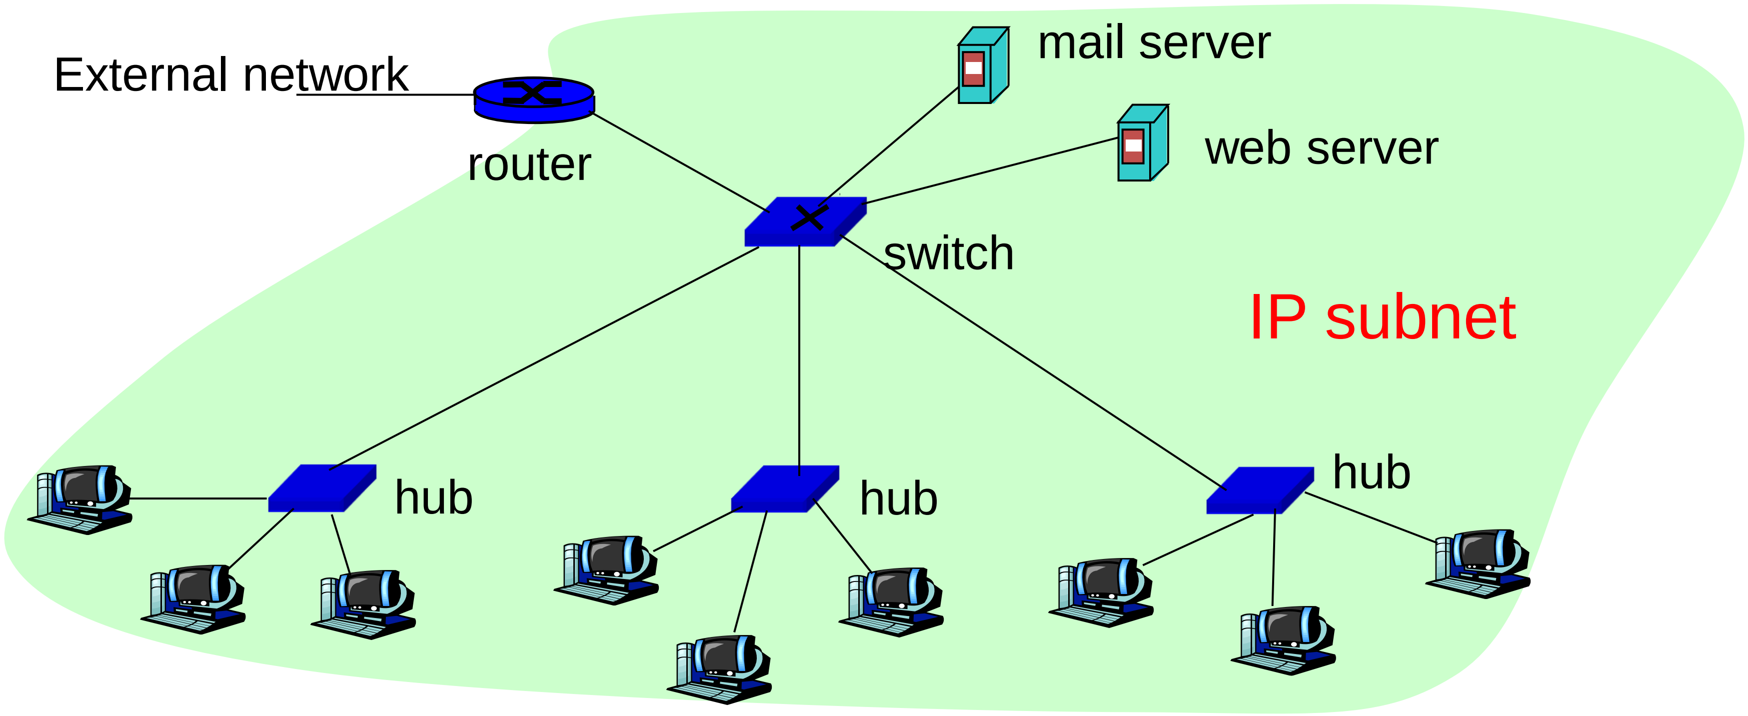

Switch (bridge) vs. router

- Intelligent store-and-forward devices
- Router
  - In the network layer (L3), based on IP addresses
  - Stores routing tables, uses routing protocols
- Switch
  - In the data link layer (L2), based on MAC addresses
  - Stores switching tables, uses address learning algorithms

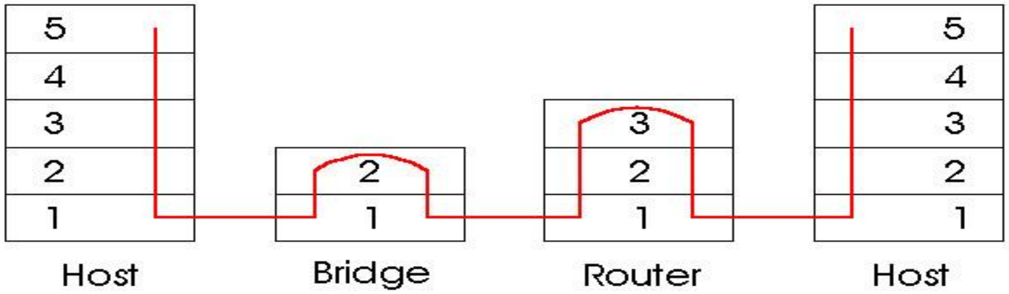

STP

- Spanning Tree Protocol
  - Part of the IEEE 802.1D standard
    - Radia Perlman (MIT, DEC)
  - Loop-free trees on a bridged LAN
    - No TTL in Ethernet (Time To Live)
      - In case of a loop, packets travel indefinitely in the network
    - Need for redundancy
      - In case of an error, there should be an alternative path

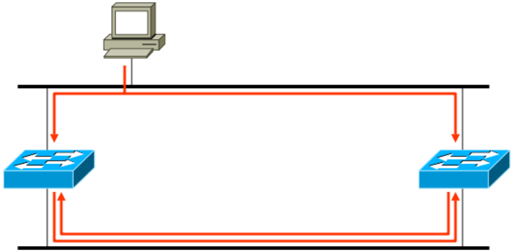

Example topology