Computer Networking

MAC Addresses, Ethernet & Wi-Fi

Lecturers: Antonio Carzaniga
Silvia Santini

Assistants: Ali Fattaholmanan
Theodore Jepsen

USI Lugano, December 7, 2018

Changelog

- V1: December 7, 2018

Last time, on December 5, 2018…

Résumé

- The link layer
  - Move a datagram over a single communication link
  - Basics Services:
    - Framing
    - Link Access
    - Reliable delivery
    - Error correction and detection

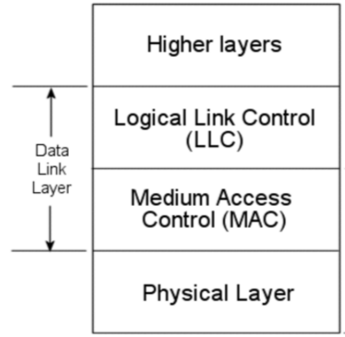

- Medium Access Control
  - Multiple nodes share a single broadcast link
  - Categories of protocols
    - Channel partitioning protocols
    - Taking-turns protocols
    - Random access protocols
  - Examples: ALOHA, Slotted ALOHA

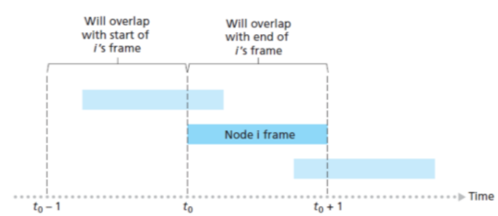

52

What about today?

- Link-layer addresses
- Ethernet (IEEE 802.3)
- Wi-Fi (IEEE 802.11)

Link-layer addresses

Image source: https://divansm.co/letter-to-santa-north-pole-address/letter-to-santa-north-pole-address-fresh-day-18-santa-s-letters/

Network adapters (aka: Network interfaces)

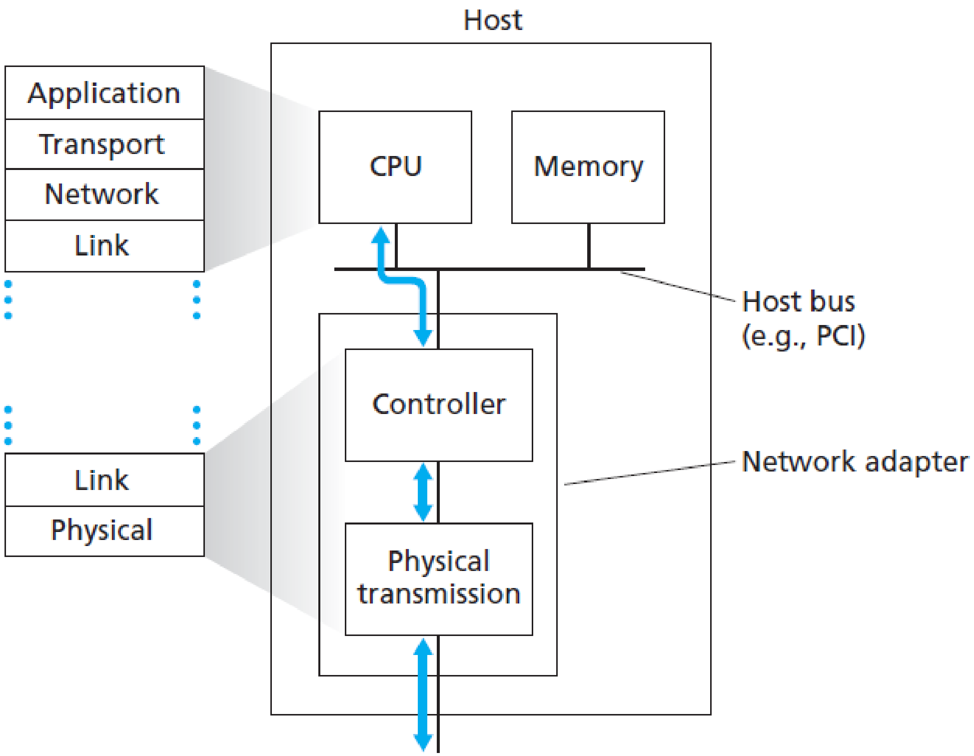

- A network adapter is a piece of hardware that connects a computer to a network
- Hosts often have multiple network adapters
  - Type *ipconfig /all* on a command window to see your computer's adapters

Image source: [Kurose 2013

Network adapters: Examples

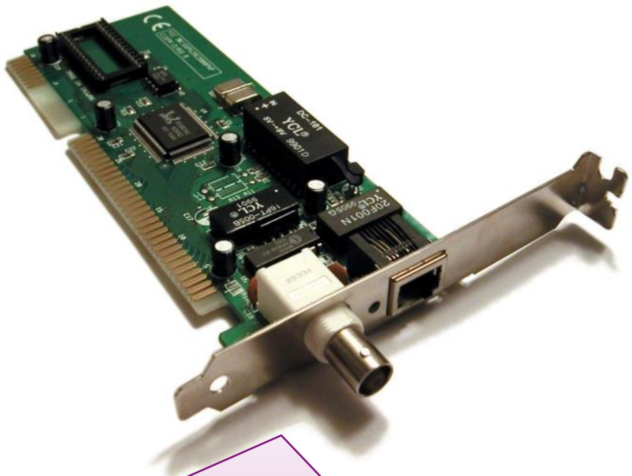

"A 1990s Ethernet network interface controller that connects to the motherboard via the now-obsolete ISA bus. This combination card features both a BNC connector (left) for use in (now obsolete) 10BASE2 networks and an 8P8C connector (right) for use in 10BASE-T networks."
https://en.wikipedia.org/wiki/Network_interface_controller

TL-WN851ND - WLAN PCI card 802.11n/g/b 300Mbps - TP-Link
https://tinyurl.com/yamo62z9

Network adapters: Addresses

- Each adapter has an own **link-layer address**
  - Usually burned into ROM

- Hosts with multiple adapters have thus multiple link-layer addresses

- A link-layer address is often referred to also as physical address, LAN address or, more commonly, **MAC address**

Format of a MAC address

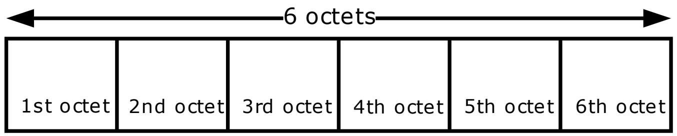

- There exist different MAC address formats, the one we consider here is the EUI-48, used in Ethernet and Wi-Fi
- 6 bytes, thus 2^{48} possible addresses
  - i.e., 281’474’976’710’656
  - i.e., 281* 10^{12} (trillions)

Image source: By Inductiveload, modified/corrected by Kju - SVG drawing based on PNG uploaded by User:Vtraveller.
This can be found on Wikipedia here: https://commons.wikimedia.org/w/index.php?curid=1852032

Do we have enough MAC addresses?

Internet of Things (IoT) connected devices installed base worldwide from 2015 to 2025 (in billions)

10^9

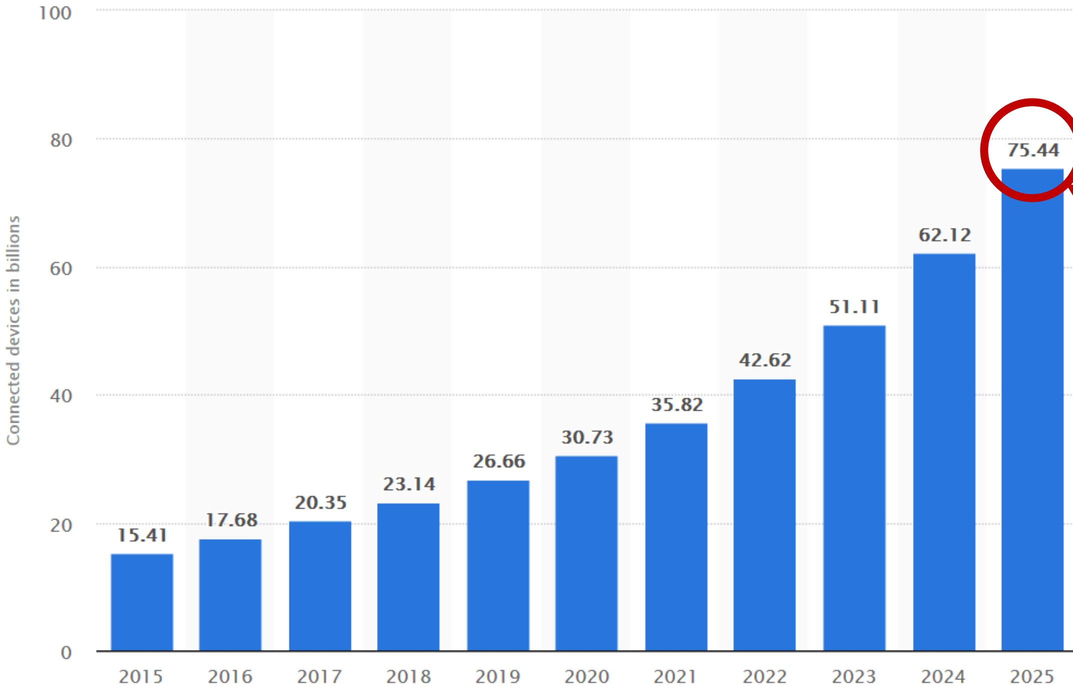

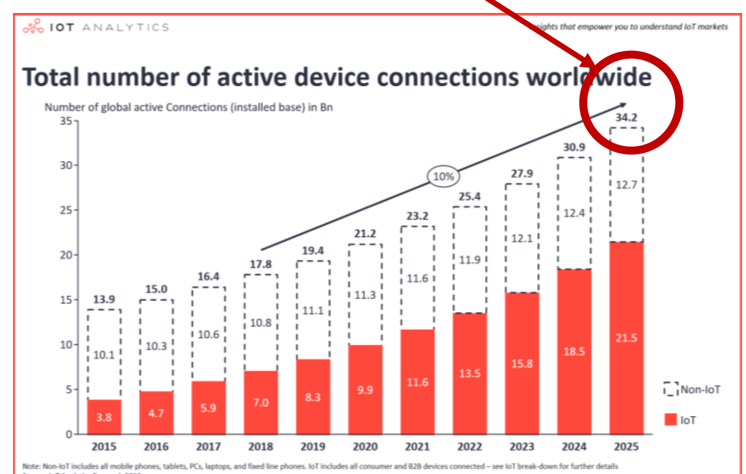

Image sources: https://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide/, https://iot-analytics.com/state-of-the-iot-update-q1-q2-2018-number-of-iot-devices-now-7b/

Format of a MAC address (continued)

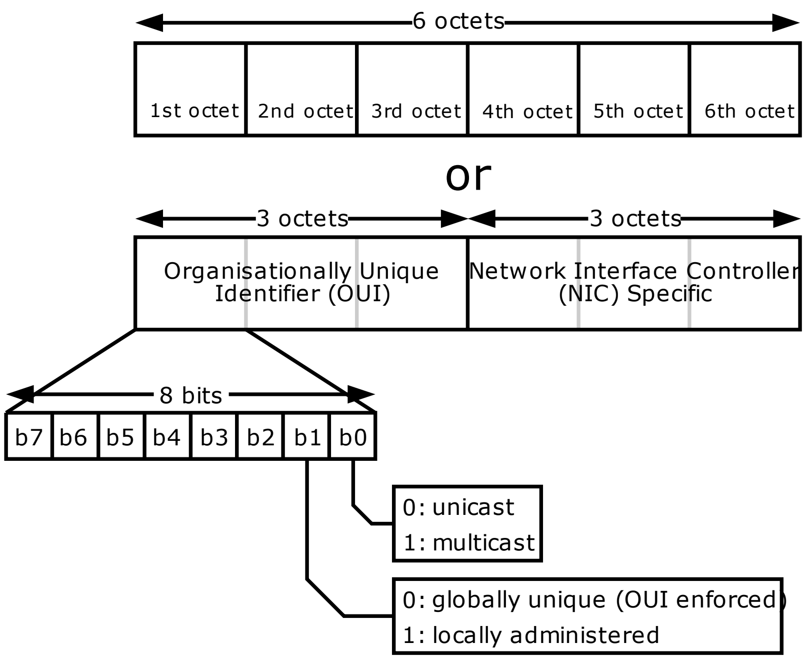

- OUI assigned by the IEEE
  - Globally unique
  - Can be purchased by vendors
- 24 bits for the OUI
- 24 bits for the NIC
  - Each OUI manages 2^{24} unique MAC addresses
- Example:
  - 44-85-00-11-3C-9E
  - 0100 0100 - 1000 01001 - 0000 0000 - 0001 0001 - 0011 1100 - 1001 1110

Image source: By Inductiveload, modified/corrected by Kju - SVG drawing based on PNG uploaded by User:Vtraveller.
This can be found on Wikipedia here: https://commons.wikimedia.org/w/index.php?curid=1852032

Are MAC addresses (really) unique?

- Yes.

- Really?
- Yes, in theory.

- In practice we may find duplicate e-mail addresses, because we might have
  - Reutilization of old addresses by vendors
  - MAC address spoofing

Frame forwarding at the link layer

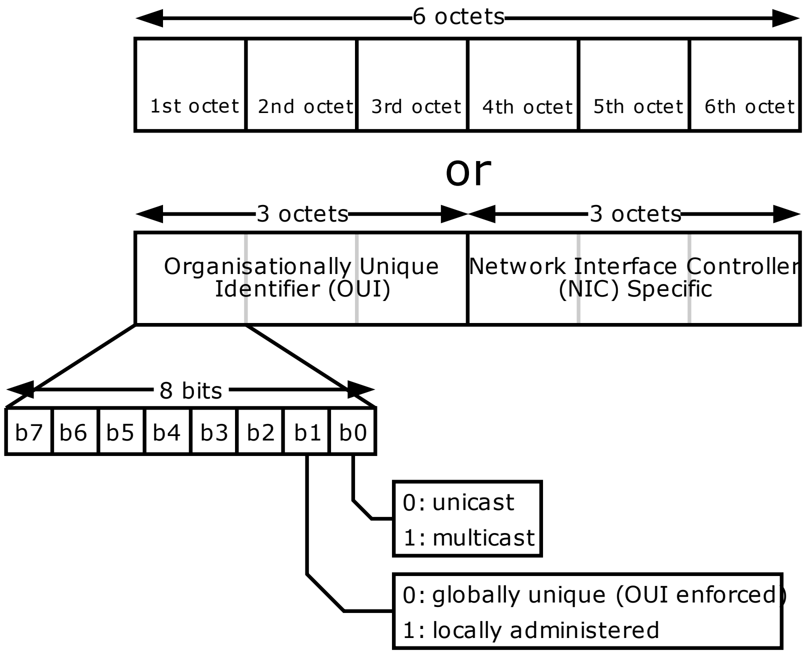

- Link-level frames contain the MAC address of the intended destination of the frame
- Router receives frame
  - Is the destination MAC address my MAC address?
    - No: Discard the frame
    - Yes: Pass the frame up to the network layer
- Broadcast address: FF-FF-FF-FF-FF-FF

Do we really need MAC addresses?

- Why not using IP addresses instead?
  1. Because then it would be difficult to support other network-layer protocols beyond IP
  2. Because then the address would have to be burn into the memory of the adapter each type it is re-assigned (e.g., at device startup or when the device moves into a new network)
  3. Because then each and every data-link frame would have to be passed to the network layer

- Different addresses at different layers allows he layers to be independent building blocks (that's why we have layers in the first place, after all…)

Why MAC addresses?

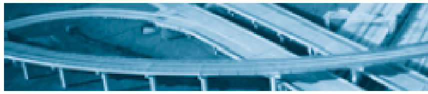

PRINCIPLES IN PRACTICE

KEEPING THE LAYERS INDEPENDENT

There are several reasons why hosts and router interfaces have MAC addresses in addition to network-layer addresses. First, LANs are designed for arbitrary network-layer protocols, not just for IP and the Internet. If adapters were assigned IP addresses rather than "neutral" MAC addresses, then adapters would not easily be able to support other network-layer protocols (for example, IPX or DECnet). Second, if adapters were to use network-layer addresses instead of MAC addresses, the network-layer address would have to be stored in the adapter RAM and reconfigured every time the adapter was moved (or powered up). Another option is to not use any addresses in the adapters and have each adapter pass the data (typically, an IP datagram) of each frame it receives up the protocol stack. The network layer could then check for a matching network-layer address. One problem with this option is that the host would be interrupted by every frame sent on the LAN, including by frames that were destined for other hosts on the same broadcast LAN. In summary, in order for the layers to be largely independent building blocks in a network architecture, different layers need to have their own addressing scheme. We have now seen three types of addresses: host names for the application layer, IP addresses for the network layer, and MAC addresses for the link layer.

The Address Resolution Protocol (ARP)

- IP actually the IP address of the next hop
  - The IPV4 header format

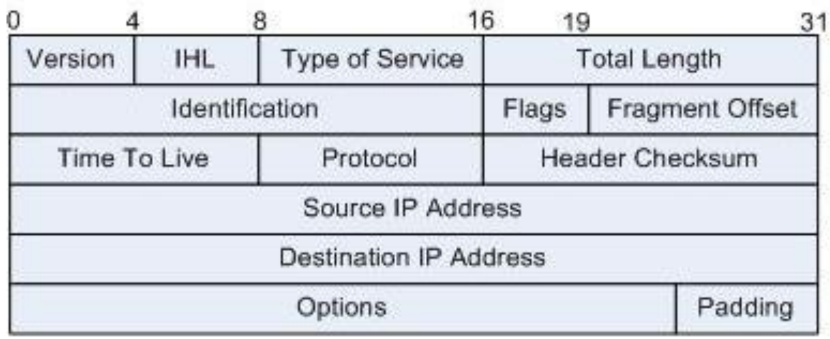

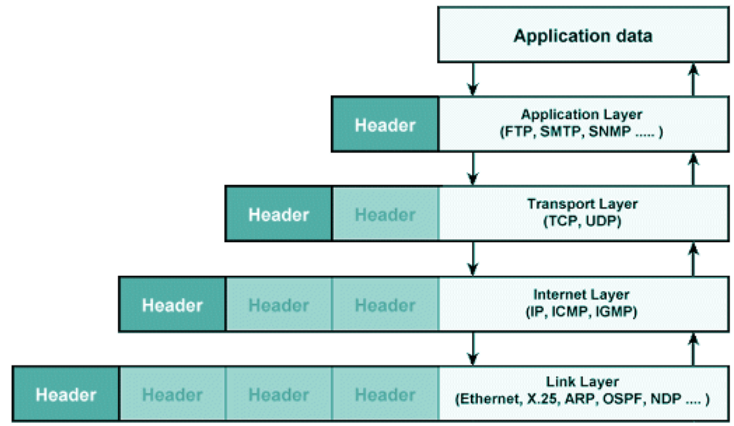

- The IP address must then be matched to a MAC address
  - *"The Address Resolution Protocol (ARP) allows a host to find the MAC address of a node […] when given the node's IP address."*
    https://technet.microsoft.com/en-us/library/cc961394.aspx

ARP: How it works

- ARP module in the sending host
  - Input: IP address of intended destination
  - Output: corresponding MAC address

- Important: ARP resolves addresses for hosts that are on the same subnet!
  - DNS resolves host names for hosts anywhere in the Internet

ARP: How it works

- ARP table with the IP <-> MAC mapping

IP Address	MAC Address	TTL
222.222.222.221	88-B2-2F-54-1A-0F	13:45:00
222.222.222.223	5C-66-AB-90-75-B1	13:52:00

- What if no entry in the table?
  - ARP query packet (broadcast)
  - ARP answer (unicast)

- BUT: What if destination is in another subnet?

ARP: How it works

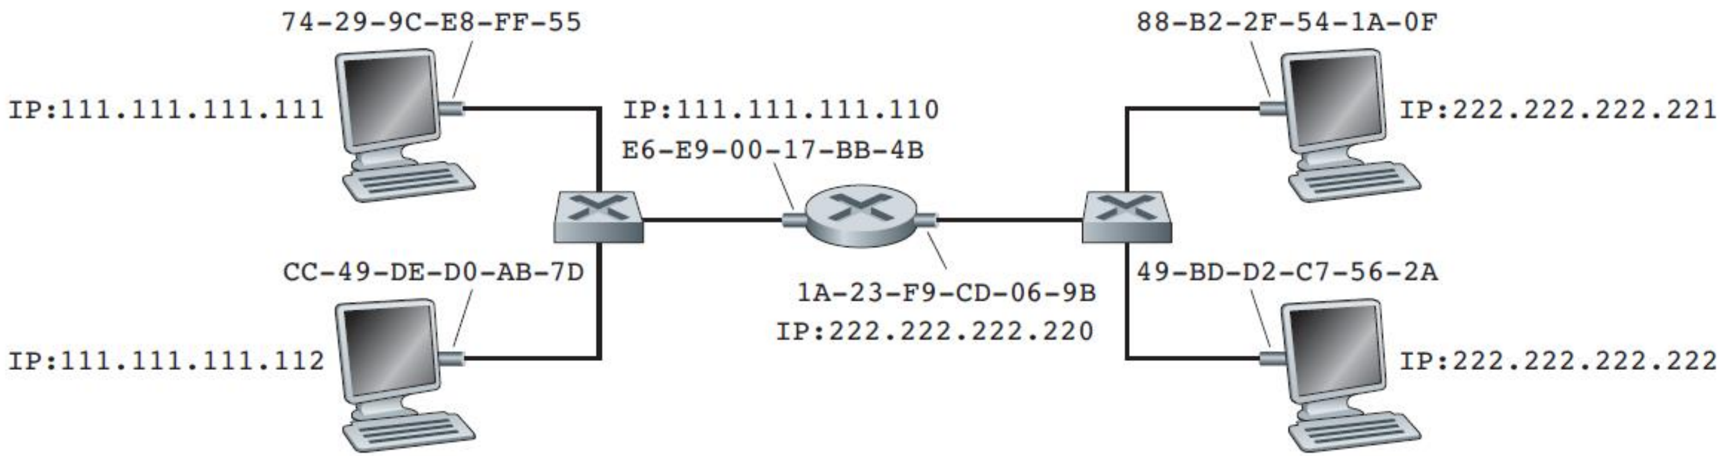

- If destination node is on another subnet, the frame must be passed to the router that interconnects the two subnets

Ethernet

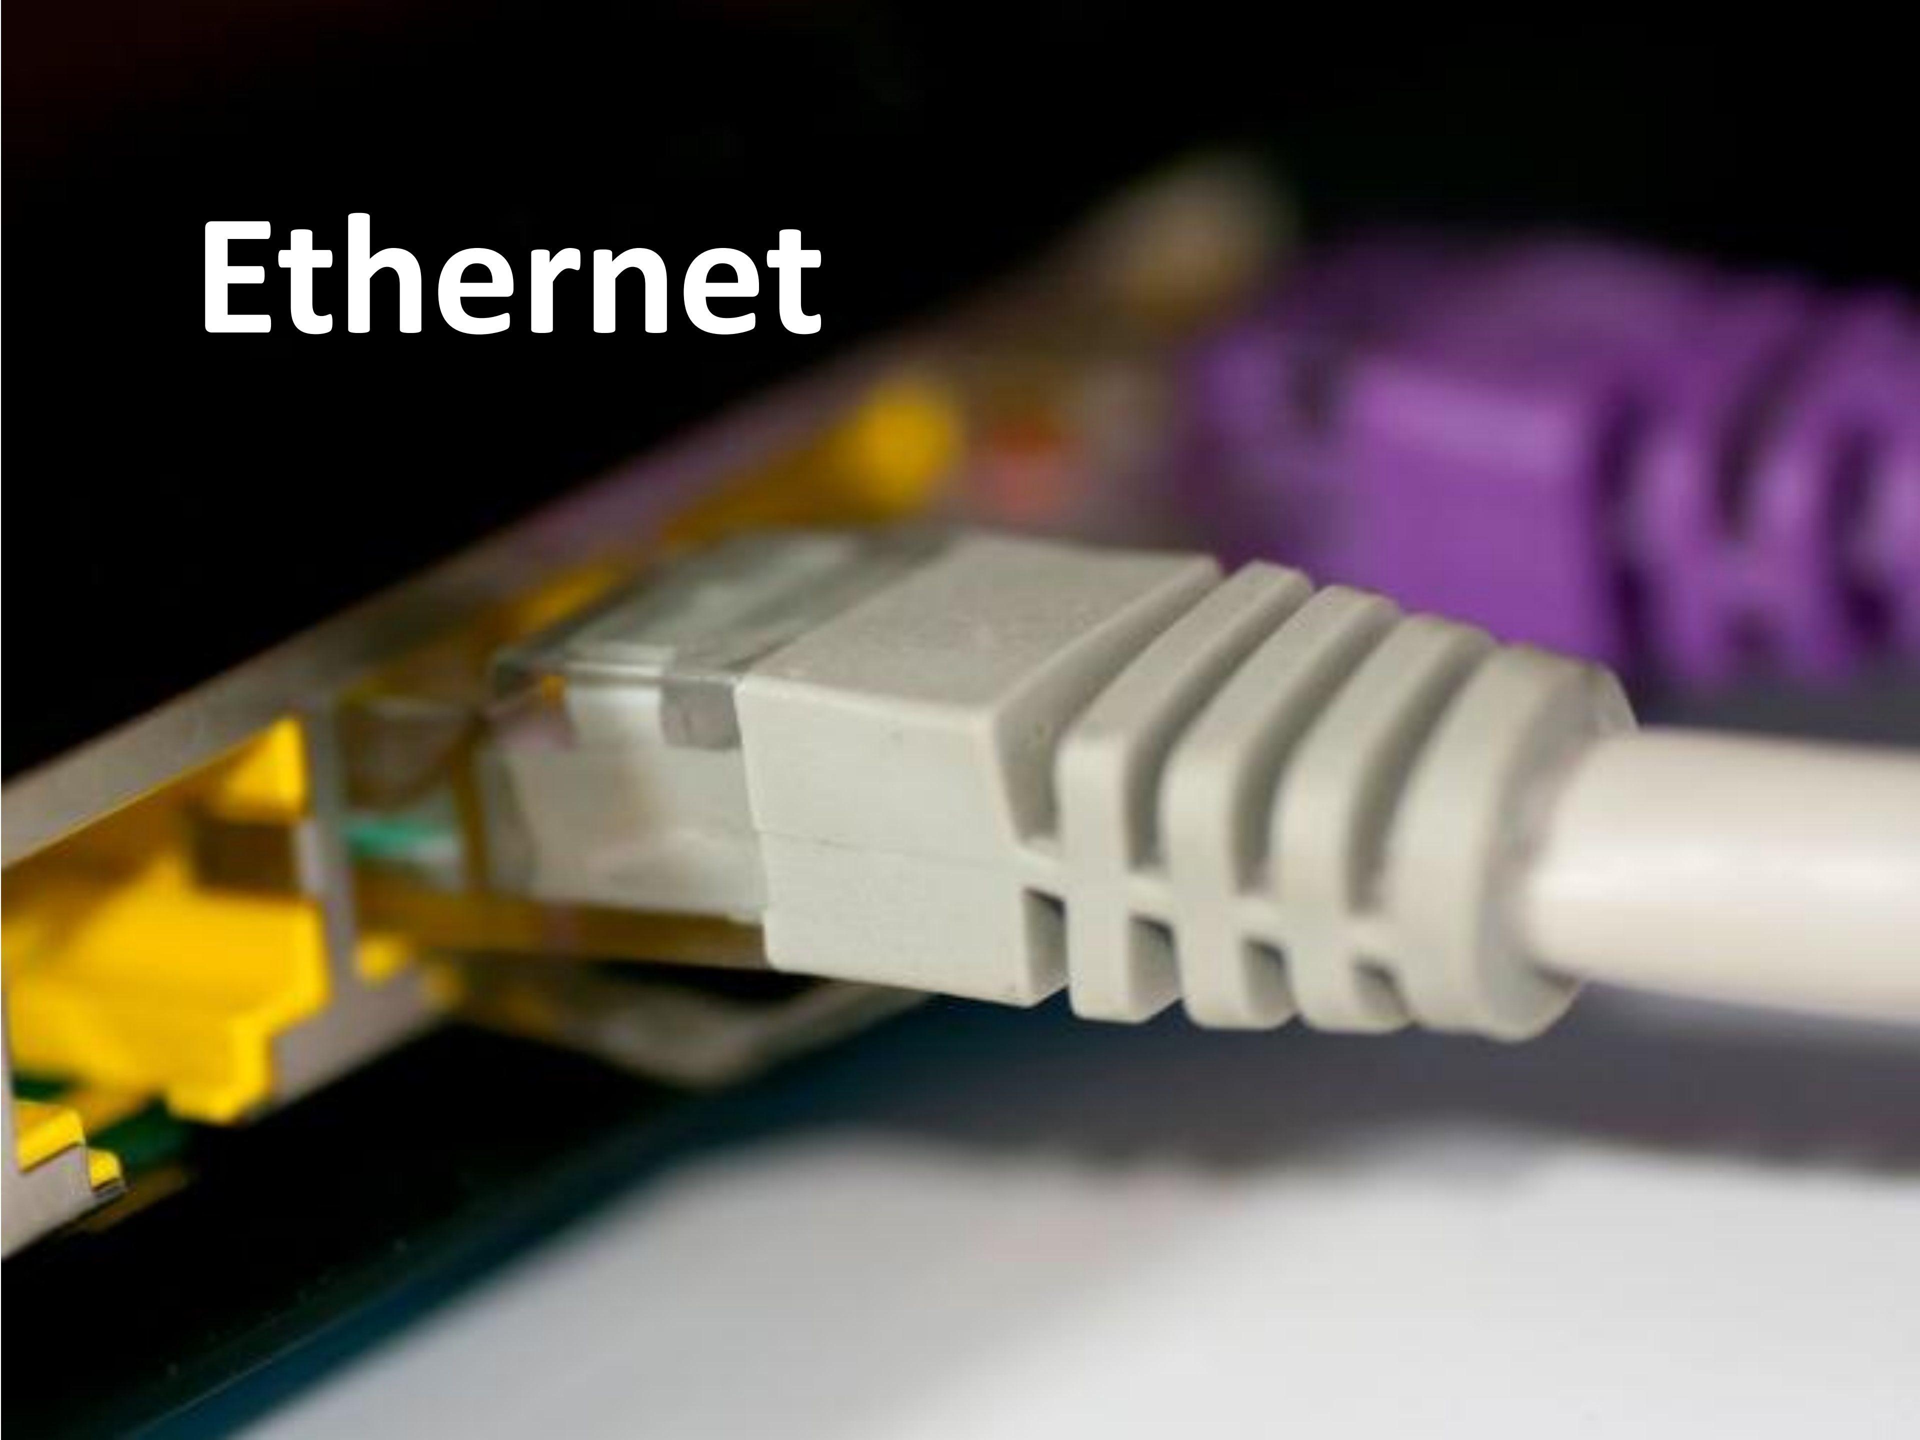

Bob Metcalfe and Ethernet

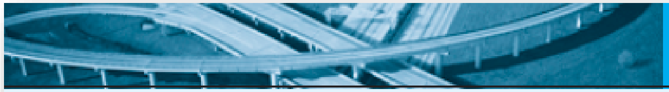

CASE HISTORY

BOB METCALFE AND ETHERNET

As a PhD student at Harvard University in the early 1970s, Bob Metcalfe worked on the ARPAnet at MIT. During his studies, he also became exposed to Abramson's work on ALOHA and random access protocols. After completing his PhD and just before beginning a job at Xerox Palo Alto Research Center (Xerox PARC), he visited Abramson and his University of Hawaii colleagues for three months, getting a first-hand look at ALOHAnet. At Xerox PARC, Metcalfe became exposed to Alto computers, which in many ways were the forerunners of the personal computers of the 1980s. Metcalfe saw the need to network these computers in an inexpensive manner. So armed with his knowledge about ARPAnet, ALOHAnet, and random access protocols, Metcalfe—along with colleague David Boggs—invented Ethernet.

Metcalfe and Boggs's original Ethernet ran at 2.94 Mbps and linked up to 256 hosts separated by up to one mile. Metcalfe and Boggs succeeded at getting most of the researchers at Xerox PARC to communicate through their Alto computers. Metcalfe then forged an alliance between Xerox, Digital, and Intel to establish Ethernet as a 10 Mbps Ethernet standard, ratified by the IEEE. Xerox did not show much interest in commercializing Ethernet. In 1979, Metcalfe formed his own company, 3Com, which developed and commercialized networking technology, including Ethernet technology. In particular, 3Com developed and marketed Ethernet cards in the early 1980s for the immensely popular IBM PCs. Metcalfe left 3Com in 1990, when it had 2,000 employees and $400 million in revenue.

Ethernet: The IEEE 802.3 standard

- Developed in the 70s at Xerox Palo Alto Research Center
  - Ethernet has its roots in the ALOHA network
  - 10 Mbps Ethernet standard published in 1978
  - First IEEE 802.3 standard (10BASE5) published in 1983

- 10BASE5 stands for:
  - 10 Mbit/s;
  - Baseband signaling (as opposed to broadband);
  - Maximum segment length of 500 m

Ethernet today

- *"Interestingly, modern Ethernet links are now largely point to point; that is, they connect one host to an Ethernet switch, or they interconnect switches. Hence, "multiple access" techniques are not used much in today's Ethernets. At the same time, wireless networks have become enormously popular, so the multiple access technologies that started in Aloha are today again mostly used in wireless networks such as 802.11 (Wi-Fi) networks."* [Peterson&Davie 2012]

- In the following, we focus on 10BASE5 as an example
  - "Historical" value
  - Multiple-access network
  - BUT: Newer (different) Ethernet standards are in use today

Ethernet (10BASE5): Physical properties

- In the Ethernet, a set of nodes send and receive frames over a shared link (bus topology)

- Ethernet 10BASE5 segment on a coaxial cable: max 500m
  - Hosts connect to an Ethernet segment by tapping into it
  - Minimal distance between taps: 2.5 m

Image source: [Elattar 2016]

(Digression) Bus topology

- A shared medium such as a cable is used to connect all hosts

- Advantages
  - Simple
  - Data sent by any host on the bus are received by all other hosts
- Disadvantages
  - Difficult to maintain
  - A damage to the bus causes network partitions

FREE EXTRA

Image source: [Elattar 2016]

Ethernet (10BASE5): Physical properties

- Any signal placed on a Ethernet segment by a host is broadcast over the entire segment (in both directions)
  - Repeaters and hubs forward the signal on all outgoing segments

- All hosts that compete for access to the same link are said to be in the same **collision domain**

Ethernet (10BASE5): Physical properties

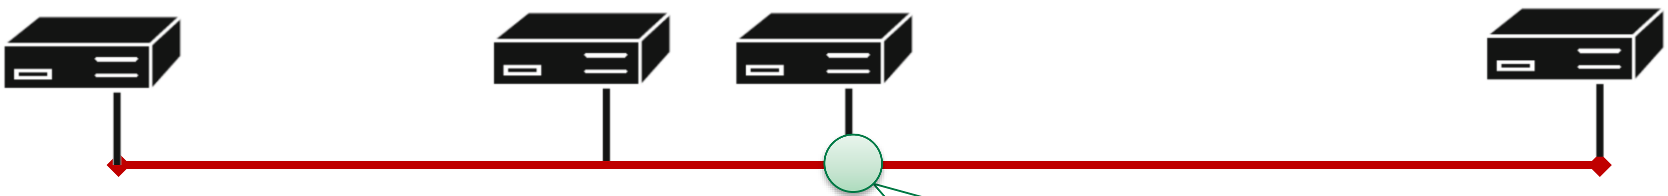

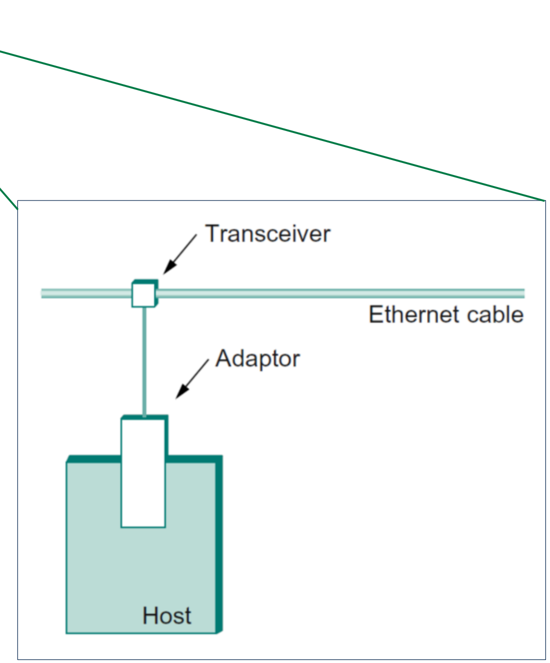

- Transceiver
  - Detects when line is idle and drives the signal when the host is transmitting
- Adaptor
  - Plugged into the host
  - Implements the actual protocol logic

Image sources: [Elattar 2016], [Patterson&Davie 2012]

Ethernet (10BASE5): Physical properties

- Multiple Ethernet segments can be joined together by repeaters

- Max 4 repeaters between any pair of hosts
  - Total reach of only 2500 m

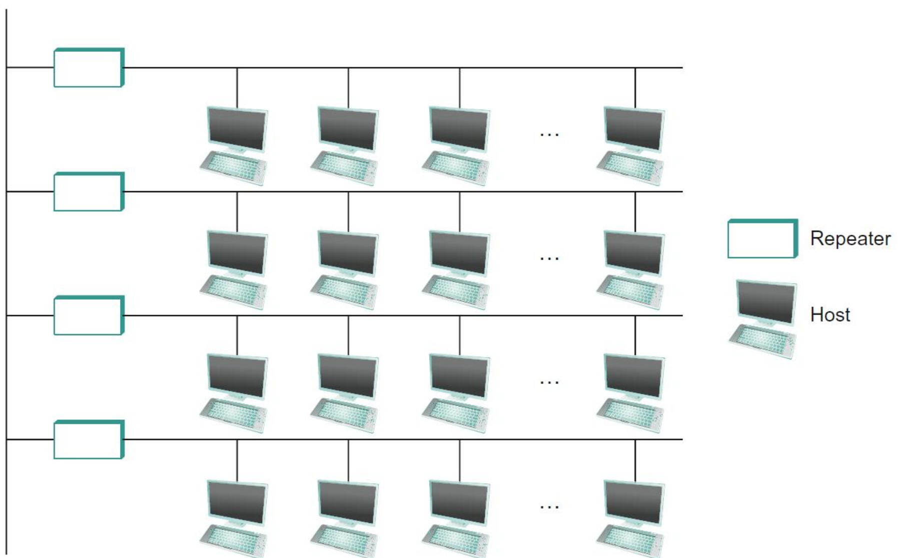

Image source: [Patterson&Davie 2012]

Ethernet (10BASE5): Physical properties

- A multiway repeater, also called a hub, repeats whatever it hears on one port out all its other ports

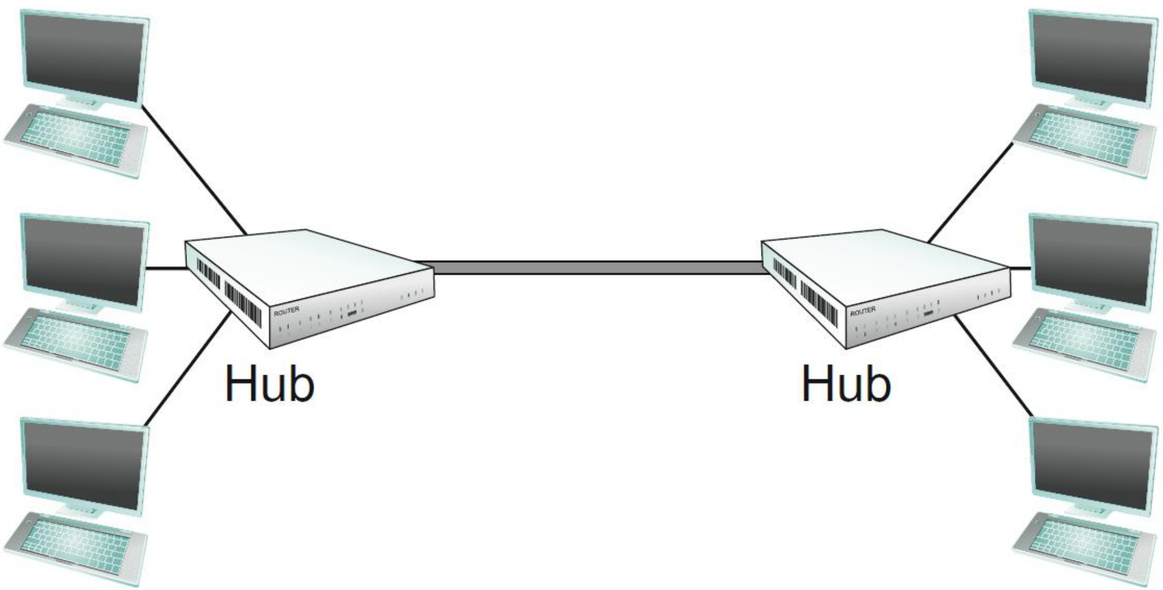

Image source: [Patterson&Davie 2012]

Ethernet (10BASE5): Packet and frame format

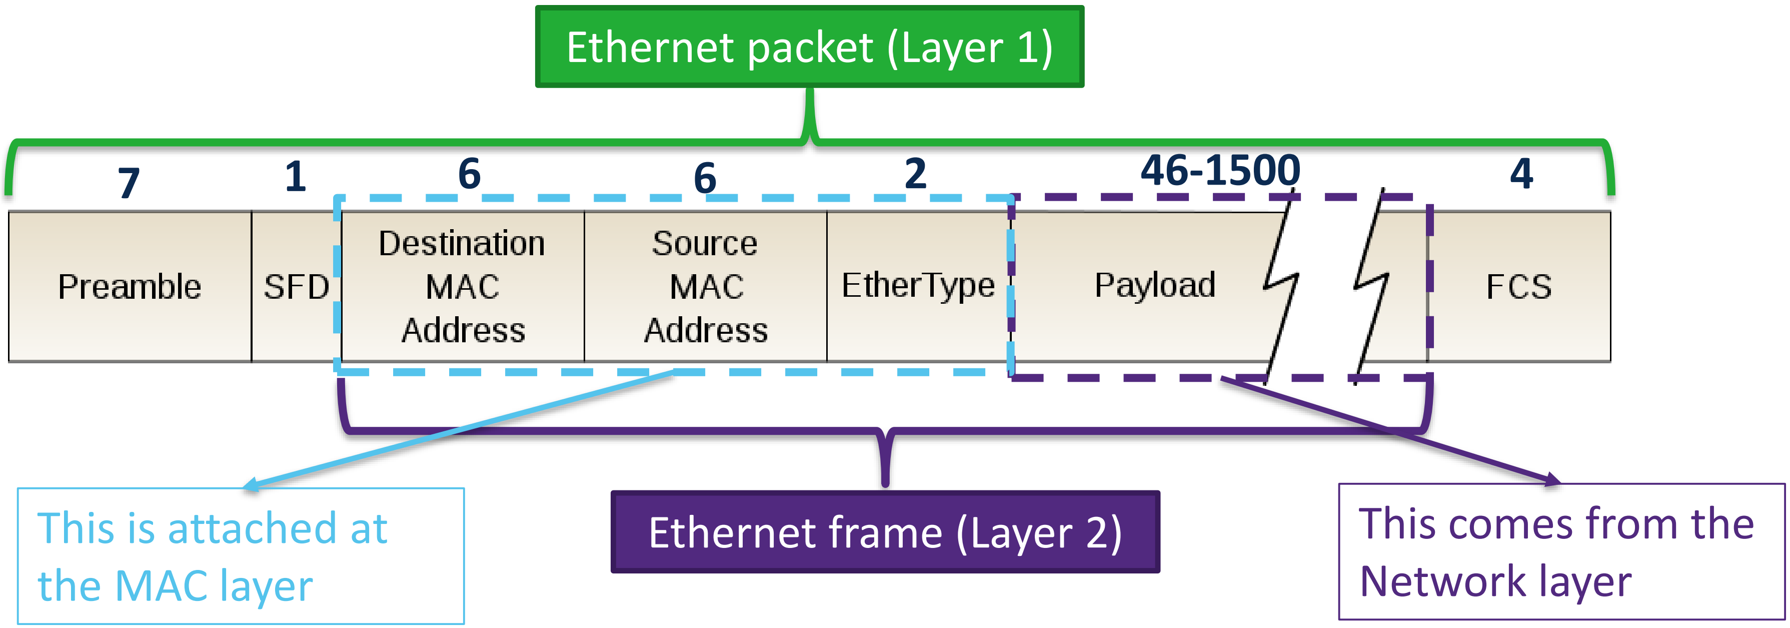

- Ethernet frame
  - 14-byte header: Source and Destination address and EtherType field
  - Size of the payload: min 46 bytes, max 1500 bytes
- The sending adaptor (physical layer) attaches the preamble, SFD and FCS before transmitting; the receiving adaptor removes them

Image source: https://upload.wikimedia.org/wikipedia/commons/4/42/Ethernet_frame.svg

Ethernet (10BASE5): Packet and frame format

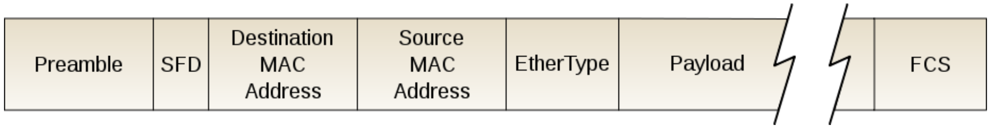

- Preamble: 7 bytes
  - Alternating 0s and 1s
  - 56 bits forming the sequence: 010101…….01
- SFD: 1 byte
  - Start of Frame Delimiter
  - Signals the end of the preamble and the start of the packet
  - Value of 171 (10101011), transmitted with least-significant bit first (0xD5)
- Destination (Source) MAC address: 6 bytes
  - Specifies the destination/source address of the communication
  - Technically, the address belongs to the adaptor, not the host
- Ethertype: 2 bytes
  - Identifies the higher-level protocol to which this frame should be delivered
- FCS: 4 bytes
  - Frame Check Sequence
  - Checksum to verify the integrity of the packet

Ethernet (10BASE5): Packet and frame format

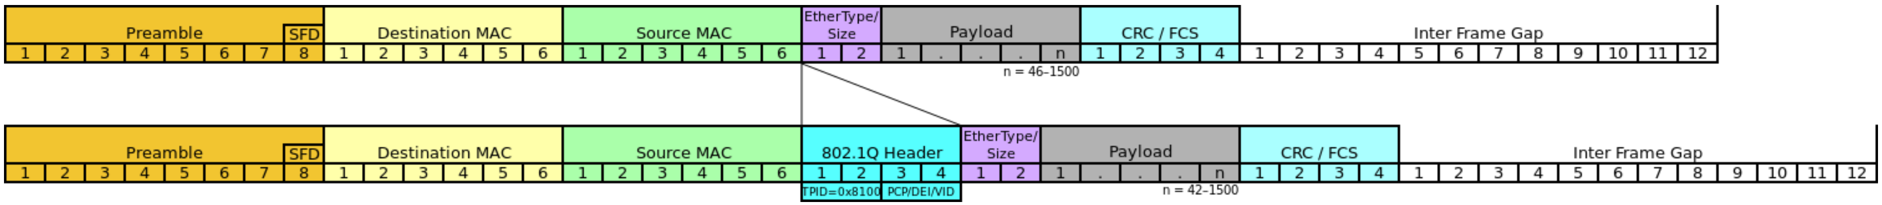

- A 802.1Q header might optionally be present
  - 4 bytes
  - First 2 bytes is the Tag Protocol Identifier (TPID)
  - TPID is a field set to a value of 0x8100 that identifies the frame as an IEEE 802.1Q-tagged frame
- Inter-frame gap
  - 12 bytes
  - Minimal pause between Ethernet frames

Image source: https://en.wikipedia.org/wiki/IEEE_802.1Q#/media/File:Ethernet_802.1Q_Insert.svg

Ethernet (10BASE5): Receiver algorithm

- An Ethernet adaptor receives all frames sent by any host and accepts:
  - frames addressed to its own address;
  - frames addressed to the broadcast address (all 1s);
  - frames addressed to a multicast address (start with 1), if it has been instructed to listen to that address;
  - all frames, if it has been placed in promiscuous mode.

- The adaptor passes to the host only the frames that it accepts

Ethernet (10BASE5): Transmitter algorithm

- The transceiver can detect whether the Ethernet link is idle or busy
  - Busy: if there is a signal on the link indicating another host is transmitting
- If the line is idle (i.e., free)
  - Transmit immediately
- If the line is busy
  - Wait for the line to go idle
  - Transmit immediately

Ethernet is a 1-persistent protocol (What is a p-persistent protocol?)

Ethernet (10BASE5): Transmitter algorithm

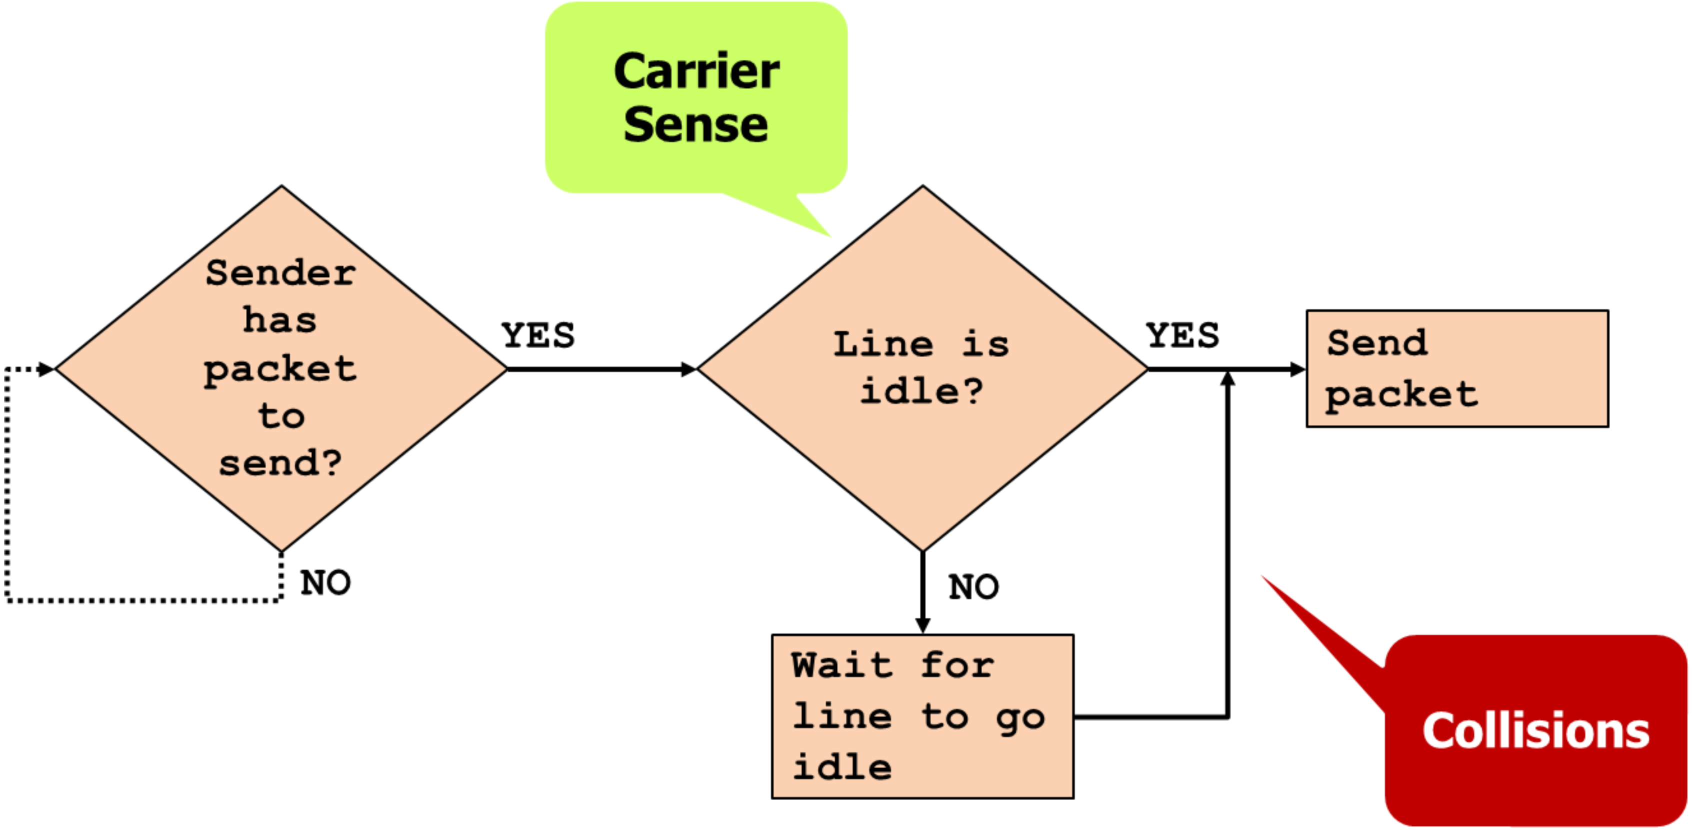

Ethernet (10BASE5): Transmitter algorithm

- Collision
  - Occurs if two (or more) hosts transmit at the same time
  - Both hosts found the line idle or were both waiting for a busy line to become idle

- Collision detection
  - A sender senses a signal from another host while its own transmission is in progress

Image source: http://murraygrp.com/wp-content/uploads/2012/06/when-do-i-drop-collision-coverage-300x225.jpg

Ethernet (10BASE5): Transmitter algorithm

- If a collision is detected
  - Transmission of a 32-bits long "jamming sequence"
  - Stop transmission
  - Backoff
  - Retransmit

- How can be guaranteed that a collision is <u>always</u> detected? What is the worst-case scenario?

Ethernet (10BASE5): Collision detection/Worst-case scenario

- Host A and B lie at the opposite ends of an Ethernet segment
  - A sends a frame at time t
  - A's frame arrives at B at time t+Δ
  - B begins transmitting at time t+Δ and collides with A's frame
  - B's runt (32-bit sequence) frame arrives at A at time t+2Δ
- Host A will detect the collision at time t + 2Δ

Ethernet (10BASE5): Collision detection/Worst-case scenario

- In the worst-case scenario, host A must transmit for 2Δ to be sure it detects all possible collisions
- Round-trip delay: 51.2 μs (empirical)
  - Maximal length of an Ethernet segment is 2500 m
  - 10 Mbps Ethernet
- Minimum packet length?
  - 512 bits (or 64 bytes)

Ethernet (10BASE5): Transmitter algorithm

- After a collision has been detected, the transmitter waits for a time $\Delta_{backoff}$ before attempting again to transmit
  - $\Delta_{backoff}$ is chosen at random in the set $\mathrm{D} = \{k \cdot (2\Delta), k = 0 \ldots (2^n - 1)\}$
  - n = min(c,10), where c is the number of collisions experienced
  - $2\Delta = 51.2$ µs

- Example
  - 1st retransmission attempt at: either 0 or 51.2 µs
  - 2nd retransmission attempt at: either 0, 51.2 µs, or 102.4 µs
  - 3rd retransmission attempt at: either 0, 51.2 µs, 102.4 µs, 153.6 µs
  - …

- Constraints
  - Number of retransmissions attempts limited to 16
  - Length of backoff window limited by capping n at 10

Ethernet (10BASE5): Transmitter algorithm

- Main techniques: CSMA, CD, Exponential Backoff

- CSMA (Carrier sensing multiple access)
  - Multiple stations can listen to the link and detect when it is in use or idle

- CD (collision detection)
  - If two or more stations are transmitting on the link simultaneously, they will detect the collision of their signals

- Exponential backoff
  - Wait before retransmitting after a collision is detected

Ethernet (10BASE5): The good and the bad

- The good
  - An Ethernet is extremely easy to maintain and administer
  - An Ethernet is cheap (only cables and adaptors needed)

- The bad
  - Best performance under light load conditions
  - Under heavy loads (over 30% of capacity) many collisions occur
  - Thus: In practice, only few host in a Ethernet (max 200) and max length of a segment much shorter than 2500m

Switched Ethernet

- Today, most Ethernets are so-called switched Ethernets
- Hosts are connected to a switch and do not share the medium with other hosts

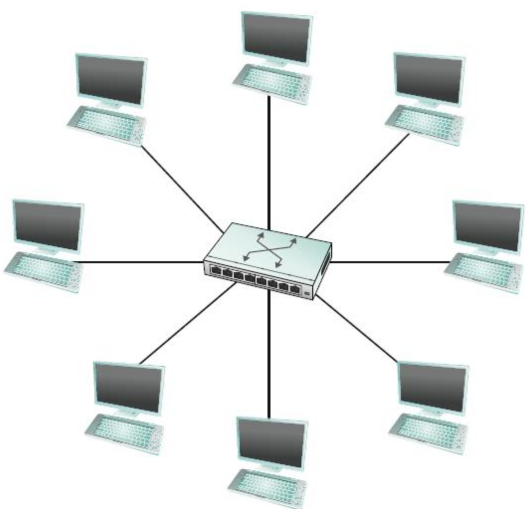

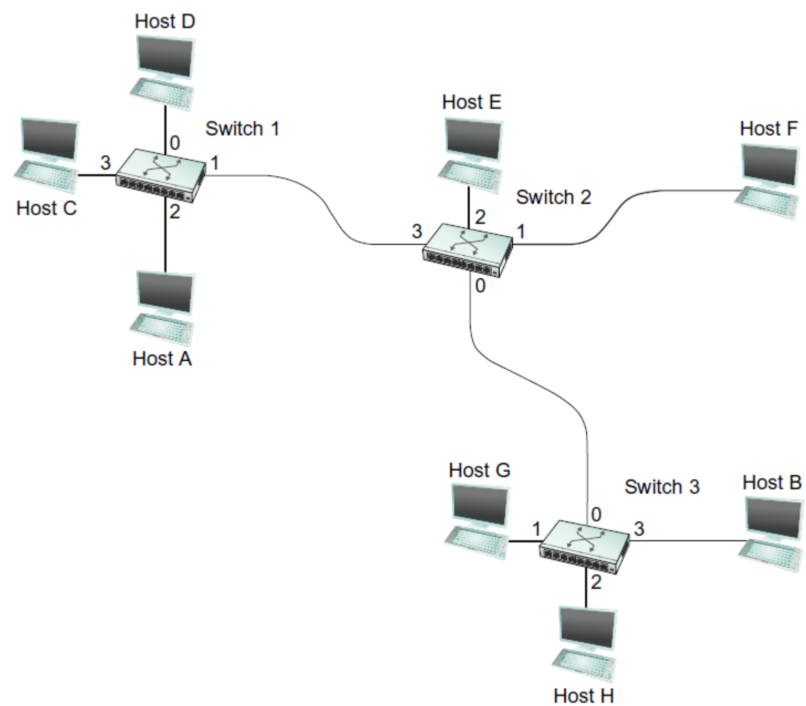

Images source: [Patterson&Davie 2012]

The Ethernet family of communication protocols (1)

Ethernet Standard	Year	Description
Experimental Ethernet	1973	2.94 Mbit/s (367 kB/s) over coaxial cable (coax) bus
Ethernet II (DIX v2.0)	1982	10 Mbit/s (1.25 MB/s) over thick coax. Frames have a Type field. This frame format is used on all forms of Ethernet by protocols in the Internet protocol suite.
IEEE 802.3 standard	1983	10BASE5 10 Mbit/s (1.25 MB/s) over thick coax. Same as Ethernet II (above) except Type field is replaced by Length, and an 802.2 LLC header follows the 802.3 header. Based on the CSMA/CD Process.
802.3a	1985	10BASE2 10 Mbit/s (1.25 MB/s) over thin Coax (a.k.a. thinnet or cheapernet)
802.3b	1985	10BROAD36
802.3c	1985	10 Mbit/s (1.25 MB/s) repeater specs
802.3e	1987	1BASE5 or StarLAN
802.3d	1987	Fiber-optic inter-repeater link
802.3i	1990	10BASE-T 10 Mbit/s (1.25 MB/s) over twisted pair
802.3j	1993	10BASE-F 10 Mbit/s (1.25 MB/s) over Fiber-Optic
802.3u	1995	100BASE-TX, 100BASE-T4, 100BASE-FX Fast Ethernet at 100 Mbit/s (12.5 MB/s) w/autonegotiation
802.3x	1997	Full Duplex and flow control; also incorporates DIX framing, so there's no longer a DIX/802.3 split

Source: https://en.wikipedia.org/wiki/IEEE_802.3

The Ethernet family of communication protocols (2)

Ethernet Standard	Year	Description
802.3z	1998	1000BASE-X Gbit/s Ethernet over Fiber-Optic at 1 Gbit/s (125 MB/s)
802.3y	1998	100BASE-T2 100 Mbit/s (12.5 MB/s) over low quality twisted pair
802.3-1998	1998	A revision of base standard incorporating the above amendments and errata
802.3ac	1998	Max frame size extended to 1522 bytes (to allow "Q-tag") The Q-tag includes 802.1Q VLAN information and 802.1p priority information.
802.3ab	1999	1000BASE-T Gbit/s Ethernet over twisted pair at 1 Gbit/s (125 MB/s)
802.3ad	2000	Link aggregation for parallel links, since moved to IEEE 802.1AX
802.3ae	2002	10 Gigabit Ethernet over fiber; 10GBASE-SR, 10GBASE-LR, 10GBASE-ER, 10GBASE-SW, 10GBASE-LW, 10GBASE-EW
802.3-2002	2002	A revision of base standard incorporating the three prior amendments and errata
802.3af	2003	Power over Ethernet (15.4 W)
802.3ak	2004	10GBASE-CX4 10 Gbit/s (1,250 MB/s) Ethernet over twinaxial cables
802.3ah	2004	Ethernet in the First Mile
802.3-2005	2005	A revision of base standard incorporating the four prior amendments and errata.
802.3aq	2006	10GBASE-LRM 10 Gbit/s (1,250 MB/s) Ethernet over multimode fiber

Source: https://en.wikipedia.org/wiki/IEEE_802.3

The Ethernet family of communication protocols (3)

Ethernet Standard	Year	Description
802.3an	2006	10GBASE-T 10 Gbit/s (1,250 MB/s) Ethernet over unshielded twisted pair (UTP)
802.3as	2006	Frame expansion
802.3au	2006	Isolation requirements for Power over Ethernet (802.3-2005/Cor 1)
802.3ap	2007	Backplane Ethernet (1 and 10 Gbit/s (125 and 1,250 MB/s) over printed circuit boards)
802.3aw	2007	Fixed an equation in the publication of 10GBASE-T (released as 802.3-2005/Cor 2)
802.3-2008	2008	A revision of base standard incorporating the 802.3an/ap/aq/as amendments, two corrigenda and errata. Link aggregation was moved to 802.1AX.
802.3av	2009	10 Gbit/s EPON
802.3-2008/Cor 1	2009	Increase Pause Reaction Delay timings which are insufficient for 10 Gbit/s (workgroup name was 802.3bb)
802.3bc	2009	Move and update Ethernet related TLVs (type, length, values), previously specified in Annex F of IEEE 802.1AB (LLDP) to 802.3.
802.3at	2009	Power over Ethernet enhancements (25.5 W)
802.3ba	2010	40 Gbit/s and 100 Gbit/s Ethernet. 40 Gbit/s over 1 m backplane, 10 m Cu cable assembly (4×25 Gbit or 10×10 Gbit lanes) and 100 m of MMF and 100 Gbit/s up to 10 m of Cu cable assembly, 100 m of MMF or 40 km of SMF respectively
802.3az	2010	Energy-efficient Ethernet

Source: https://en.wikipedia.org/wiki/IEEE_802.3

The Ethernet family of communication protocols (4)

Ethernet Standard	Year	Description
802.3bd	2010	Priority-based Flow Control. An amendment by the IEEE 802.1 Data Center Bridging Task Group (802.1Qbb) to develop an amendment to IEEE Std 802.3 to add a MAC Control Frame to support IEEE 802.1Qbb Priority-based Flow Control.
802.3.1	2011	MIB definitions for Ethernet. It consolidates the Ethernet related MIBs present in Annex 30A&B, various IETF RFCs, and 802.1AB annex F into one master document with a machine readable extract. (workgroup name was P802.3be)
802.3bg	2011	Provide a 40 Gbit/s PMD which is optically compatible with existing carrier SMF 40 Gbit/s client interfaces (OTU3/STM-256/OC-768/40G POS).
802.3bf	2011	Provide an accurate indication of the transmission and reception initiation times of certain packets as required to support IEEE P802.1AS.
802.3-2012	2012	A revision of base standard incorporating the 802.3at/av/az/ba/bc/bd/bf/bg amendments, a corrigenda and errata.
802.3bk	2013	This amendment to IEEE Std 802.3 defines the physical layer specifications and management parameters for EPON operation on point-to-multipoint passive optical networks supporting extended power budget classes of PX30, PX40, PRX40, and PR40 PMDs.
802.3bm	2015	100G/40G Ethernet for optical fiber
802.3-2015	2015	802.3bx – a new consolidated revision of the 802.3 standard including amendments 802.3bk/bj/bm
802.3bw	2015[4]	100BASE-T1 – 100 Mbit/s Ethernet over a single twisted pair for automotive applications
P802.3ar	Cancelled	Congestion management (withdrawn)

Source: https://en.wikipedia.org/wiki/IEEE_802.3

The Ethernet family of communication protocols (5)

Ethernet Standard	Year	Description
802.3bj	Jun-14	Define a 4-lane 100 Gbit/s backplane PHY for operation over links consistent with copper traces on “improved FR-4” (as defined by IEEE P802.3ap or better materials to be defined by the Task Force) with lengths up to at least 1 m and a 4-lane 100 Gbit/s PHY for operation over links consistent with copper twinaxial cables with lengths up to at least 5 m.
802.3bp	June 2016[2]	1000BASE-T1 – Gigabit Ethernet over a single twisted pair, automotive & industrial environments
802.3bq	June 2016[3]	25G/40GBASE-T for 4-pair balanced twisted-pair cabling with 2 connectors over 30 m distances
802.3by	June 2016[5]	Optical fiber, twinax and backplane 25 Gigabit Ethernet[6]
802.3bz	Sept 2016[7]	2.5GBASE-T and 5GBASE-T – 2.5 Gigabit and 5 Gigabit Ethernet over Cat-5/Cat-6 twisted pair
802.3cc	~2017	25 Gb/s over Single Mode Fiber
802.3bs	~2017	400 Gbit/s Ethernet over optical fiber using multiple 25G/50G lanes
802.3bt	~2017	Power over Ethernet enhancements up to 100 W using all 4 pairs balanced twisted-pair cabling, lower standby power and specific enhancements to support IoT applications (e.g. Lighting, sensors, building automation).
802.3cd	~2018	Media Access Control Parameters for 50 Gb/s and Physical Layers and Management Parameters for 50 Gb/s, 100 Gb/s, and 200 Gb/s Operation
802.3ca	~2019	100G-EPON – 25 Gb/s, 50 Gb/s, and 100 Gb/s over Ethernet Passive Optical Networks

Wi-Fi

What is “Wi-Fi”?

- The Wi-Fi alliance is a *“global non-profit association with the goal of driving the best user experience with a new wireless networking technology – regardless of brand”* http://www.wi-fi.org/who-we-are

- Member companies include
  - Sponsors
    - Apple, Microsoft, Cisco, Dell, LG, Nokia, Qualcomm, Motorola, Samsung, Sony, Texas Instrument, Intel, Broadcom, Huawei, Comcast
  - Regular members
    - Google, HTC, D-Link, Fon, Fujitsu, Canon, Ericsson, TeliaSonera, Fastweb, …

Wi-Fi Alliance Vision & Mission

- Wi-Fi Alliance Vision
  - Connecting everyone and everything, everywhere

- Wi-Fi Alliance Mission
  - Foster highly-effective collaboration among stakeholders
  - Deliver excellent connectivity experiences through interoperability
  - Embrace technology innovation
  - Promote the adoption of our technologies worldwide
  - Advocate for fair worldwide spectrum rules
  - Lead, develop and embrace industry-agreed standards

Wi-Fi technology

- Wi-Fi is often (improperly) used as a synonymous of **WLAN** (Wireless Local Area Network)

- Wi-Fi devices are devices based on any of the IEEE 802.11 radio standards

- IEEE 802.11 extends the common wired Ethernet local network into the wireless domain
  - http://grouper.ieee.org/groups/802/11/

IEEE 802.11 (WLAN)

- IEEE 802.11 is a media access control and physical layer specification

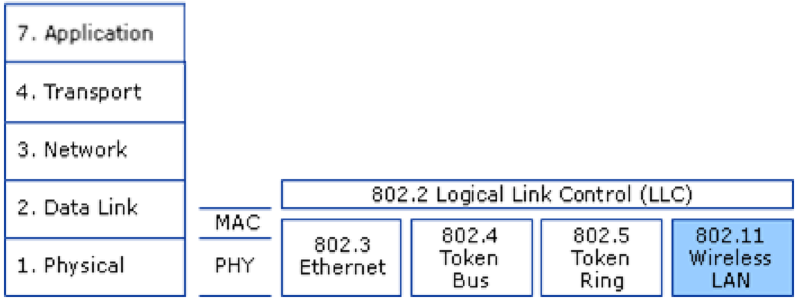

- Designed for use in a limited geographical area
  - Home, office buildings, campuses, etc.

53

Ethernet (IEEE 802.3) vs Wi-Fi (IEEE 802.11)

- Ethernet

- Wi-Fi

IEEE 802.11: Logical architecture

- Main components
  - Station (STA)
  - Wireless access point (AP)
  - Independent basic service set (IBSS)
  - Basic service set (BSS)
  - Distribution system (DS)
  - Extended service set (ESS)

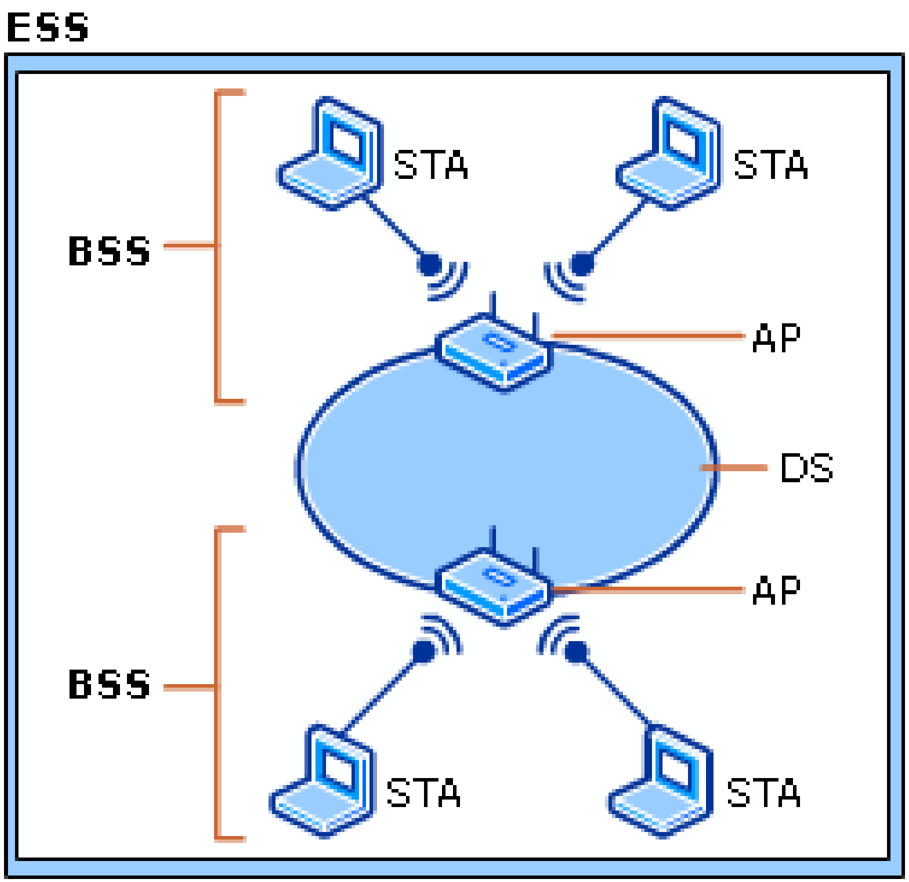

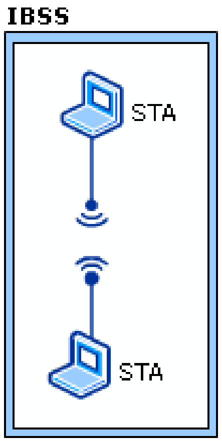

Image source: [Microsoft 802.11]

IEEE 802.11: Operating modes

- Infrastructure mode

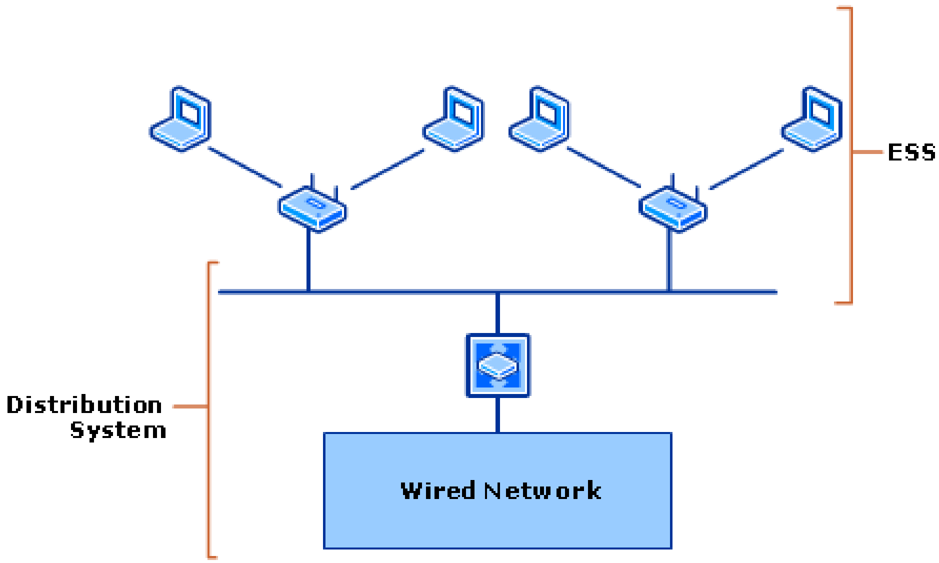

- Ad hoc mode (peer-to-peer)

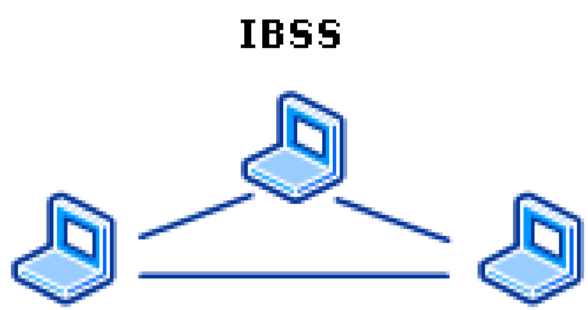

"In both operating modes, a Service Set Identifier (SSID), also known as the wireless network name, identifies the wireless network. The SSID is a name configured on the wireless AP (for infrastructure mode) or an initial wireless client (for ad hoc mode) that identifies the wireless network. The SSID is periodically advertised by the wireless AP or the initial wireless client using a special 802.11 MAC management frame known as a beacon frame." [Microsoft 802.11]

Image source: [Microsoft 802.11]

Collisions in a wireless environment

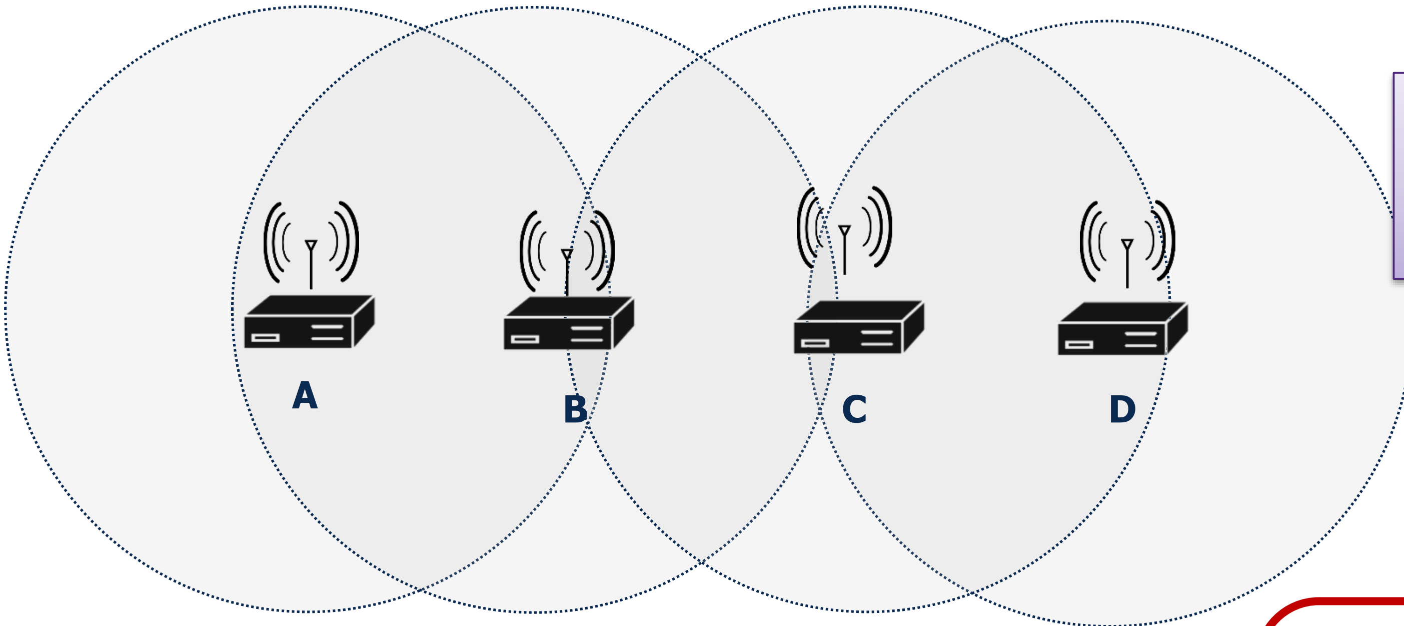

B can communicate with A and C but not with D

C can communicate with B and D but not with A

- Communications from A and C to B can collide
  - But A cannot hear C and C cannot hear A!
- Node C overhears when B communicates with A
  - But this should not hamper C to communicate with D

Hidden terminal problem

Exposed terminal problem

These problems are addressed by the Carrier Sense Multiple Access with Collision Avoidance (CSMA / CA) protocol

IEEE 802.11: Collision avoidance

- At the MAC layer: CSMA/CA protocol
  - Carrier Sense Multiple Access with Collision Avoidance

- At the physical (PHY) layer: Use the wireless medium efficiently
  - Frequency-Hopping Spread Spectrum (FHSS)
  - Direct-Sequence Spread Spectrum (DSSS)
  - Orthogonal Frequency-Division Multiplexing (OFDM)

Basic technique: CSMA / CA

- Carrier Sense Multiple Access with Collision Avoidance
  - CSMA / CA

- Carrier Sense
  - Before sending a packet, the transmitter checks if the wireless channel is free or busy
  - If the channel is free, the transmitter starts a RTS/CTS handshake
  - If the channel is busy, the transmitter backs off (random exponential back off similar to that of the Ethernet)

59

RTS/CTS Handshake

Let's assume B is the access point

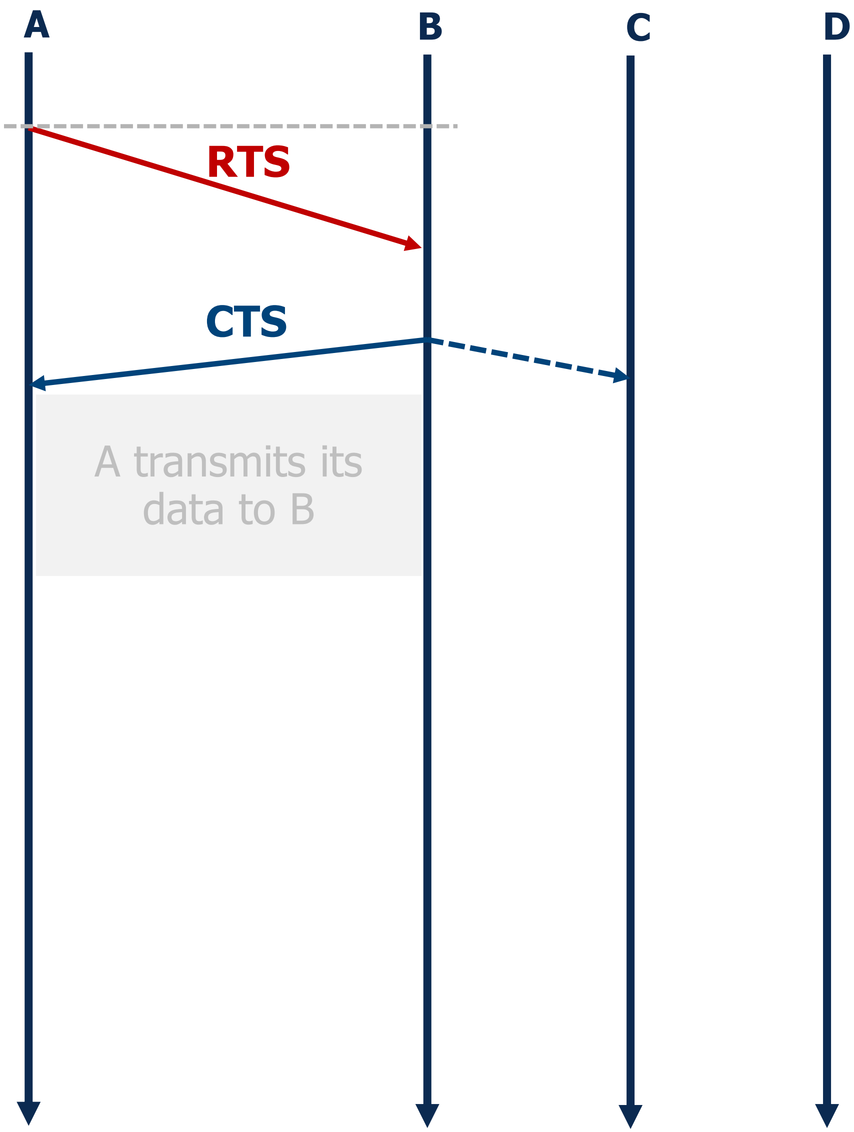

Request To Send (RTS)
Clear To Send (CTS)

- RTS and CTS are broadcasted
- RTS
  - Includes a field that indicates the length of the packet to transmit (i.e., for how long the sender wants to hold the medium)
- CTS
  - Echoes the length field back to sender
  - Each node receiving the RTS or CTS frames knows that A is going to send and for how long

RTS/CTS Handshake with ACK

Let's assume B is the access point

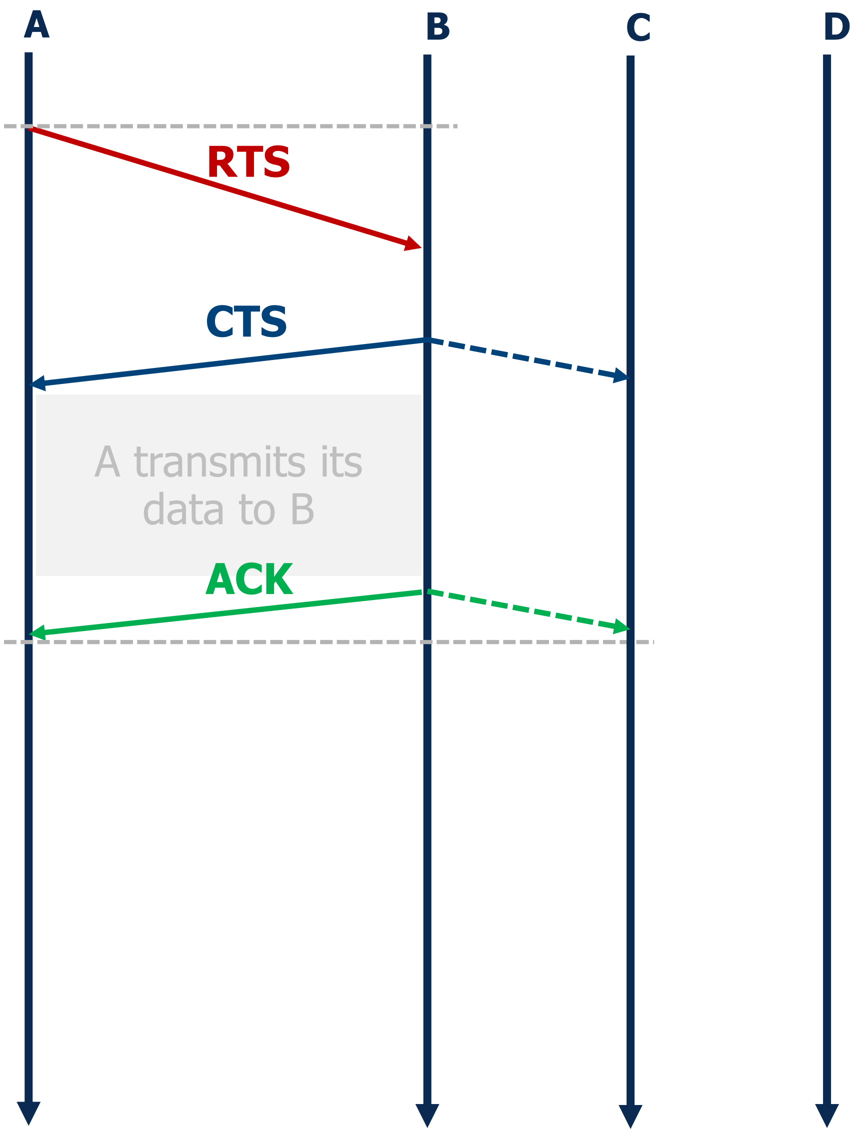

Request To Send (RTS)
Clear To Send (CTS)
Acknowledgment of successful reception (ACK)

- RTS, CTS, and ACK are all broadcasted
- RTS
  - Includes a field that indicates the length of the packet to transmit (i.e., for how long the sender wants to hold the medium)
- CTS
  - Echoes the length field back to sender
  - Each node receiving the RTS or CTS frames knows that A is going to send and for how long
- ACK
  - All nodes must wait for this ACK before (re-)trying to transmit
  - "Virtual carrier sense“ [Ye 2004]

Network Allocation Vector (NAV)

RTS / CTS Handshake: RTS collisions

- RTS packets from two or more nodes may collide
  - "Implicit" detection in absence of CTS packet (after a timeout)

- A RTS request may hamper the reception of a CTS packet
  - A new RTS would then be accepted
  - Solution: CTS packets longer than RTS packets (this way, collisions are detected)

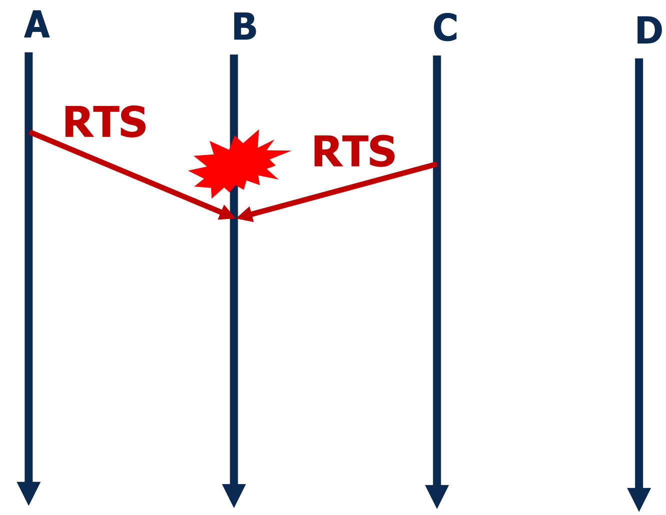

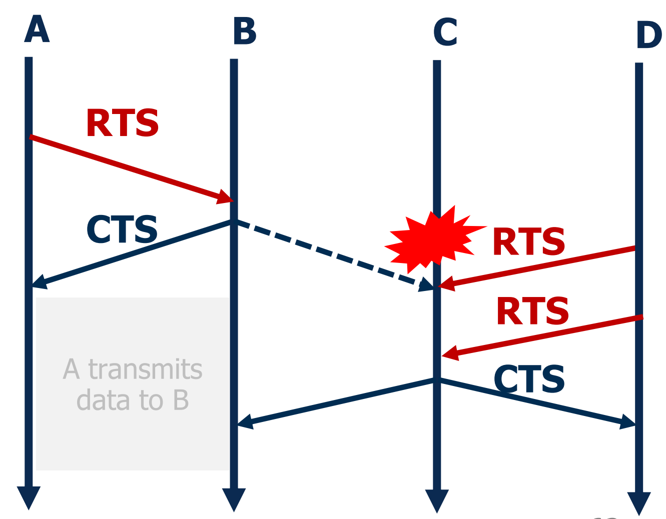

62

How does a station knows to whom to send the RTS?

- APs advertise their presence using **beacons**
- Stations request to connect to one of the available APs
  - If more access points available: connect to the one that has the highest RSSI (Received Signal Strength Indicator)


Conclusion

Résumé (December 7, 2018)

- Link-layer addresses
  - EUI-48 format
  - ARP protocol

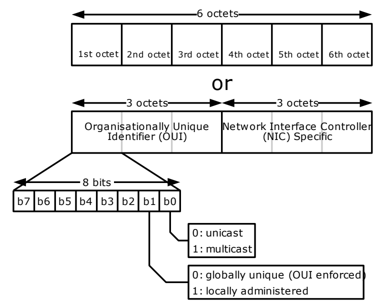

- Ethernet (IEEE 802.3)
  - 10BASE5 standard (1983)
  - CSMA/CD
    - Worst-case scenario analysis

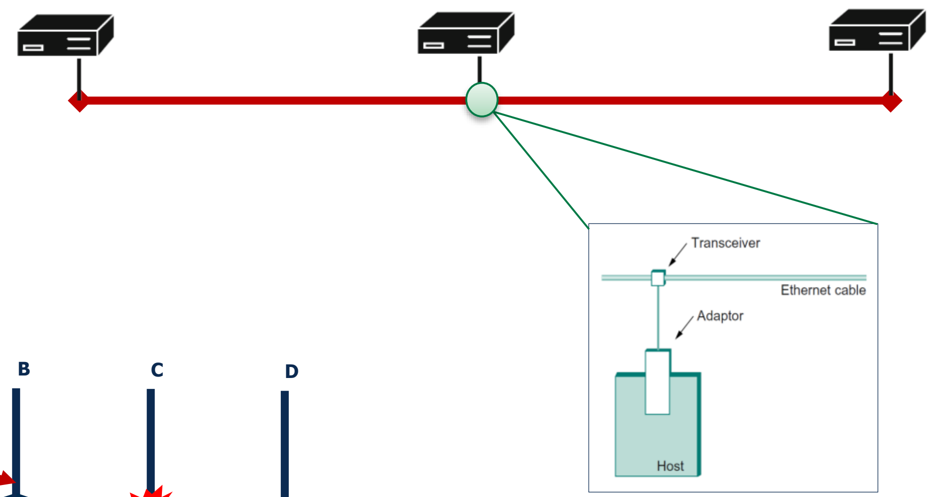

- Wi-Fi (IEEE 802.11)
  - CSMA/CA
    - RTS/CTS handshake

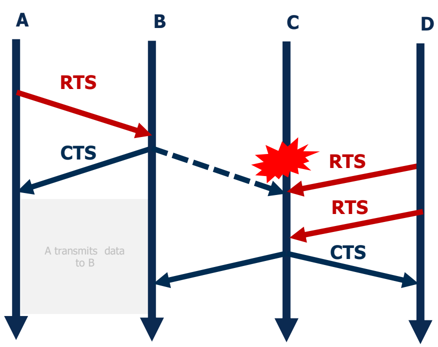

Required readings

[Kurose 2013] James F. Kurose and Keith W. Ross. Computer Networking: A Top-Down Approach. Pearson, 6th Edition 2013. **[Section 5.4 (excl. 5.4.3 and 5.4.4)]**

Additional references

[Microsoft 802.11] Microsoft TechNet: How 802.11 Wireless Works. Available at: https://technet.microsoft.com/en-us/library/cc757419(v=ws.10).aspx
(Excluded "802.11 PHY Sublayer" and subsequent paragraphs)

Acknowledgments

- The following sources or authors have directly or indirectly contributed (through their ideas, papers, presentations, and more) to the realization of these slides:
  - …
  - and others that might have been omitted unintentionally.

Comic of the day

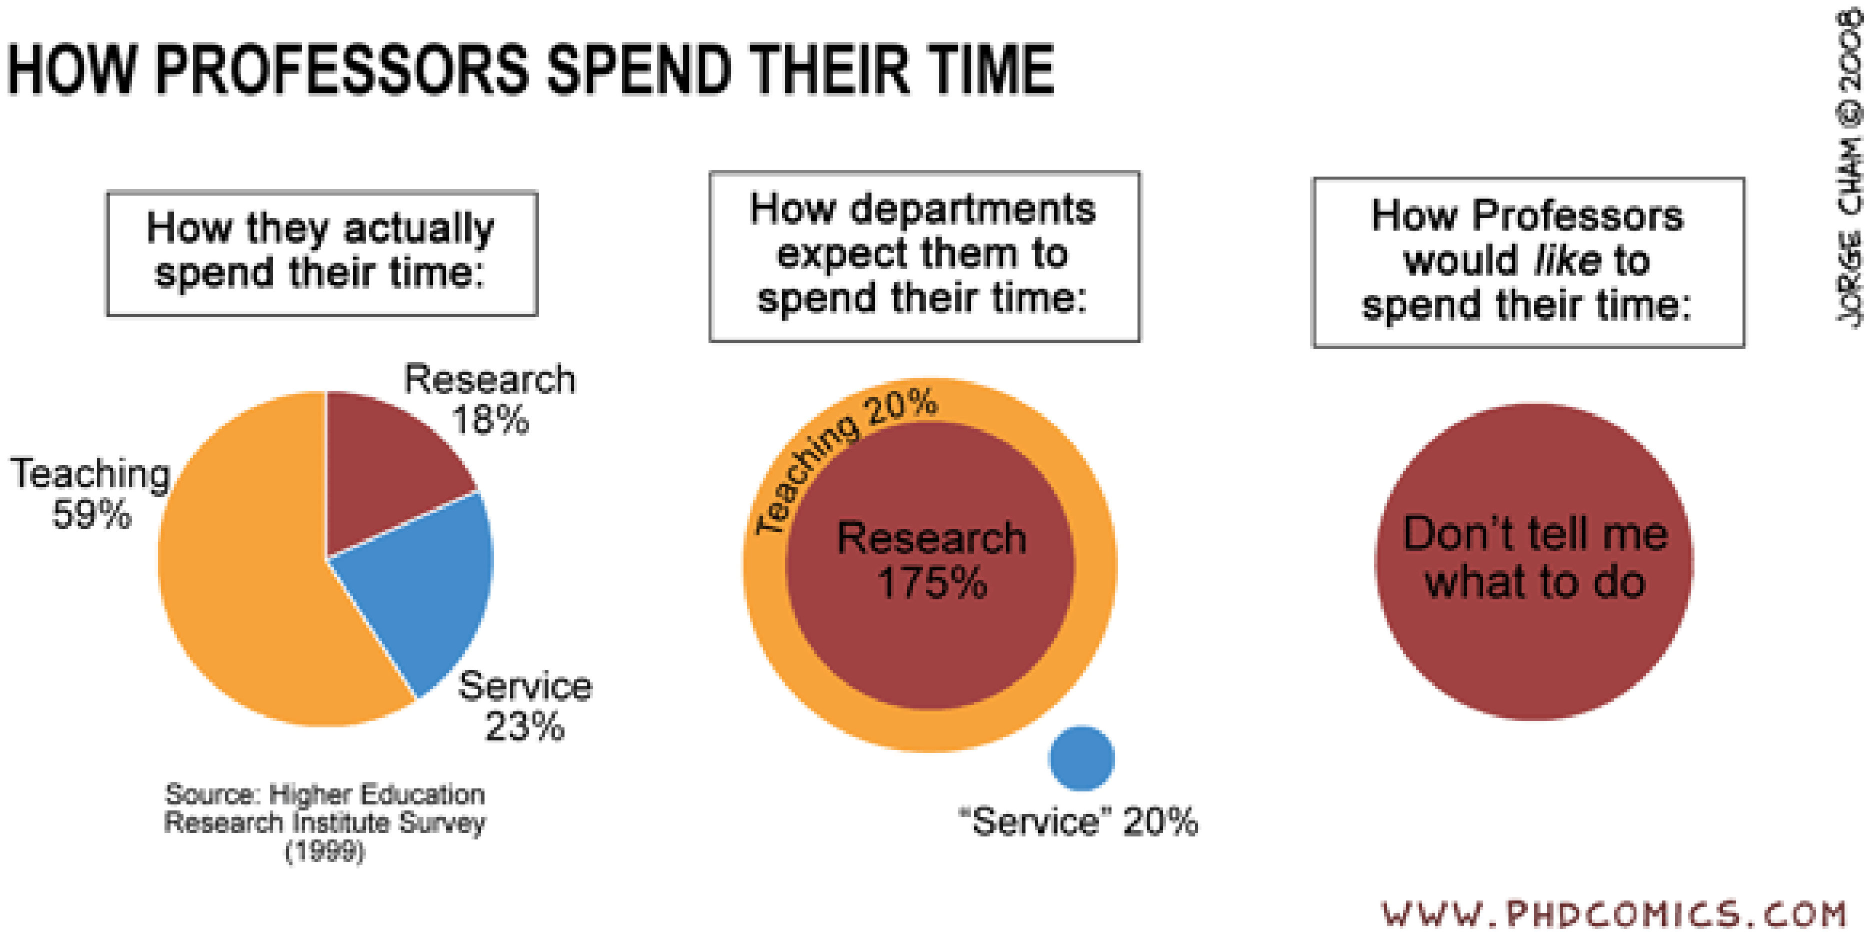

Source: http://phdcomics.com/comics/archive.php?comicid=1060