Ethernet, AVB, and TSN introduction

LINKÖPING UNIVERSITY

Ethernet (CSMA/CD)

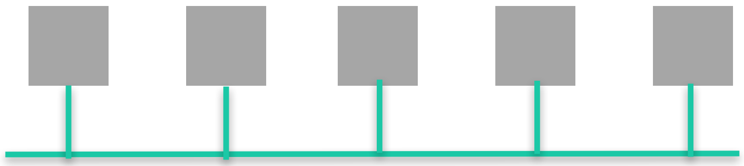

- Created in the mid 70s
- Ethernet bus: shared media access
- 46 to 1500 byte payloads; 10 Mbps to 10 Gbps
- Node transmits as soon as medium is free
- Collisions can occur during the interval of one slot after start of transmission (512 bits)
- Jamming signal (32 bits) is sent in case of collisions
- Retransmission after random amount of time
- Outdated
- Not appropriate for real-time systems

Switched Ethernet

- IEEE 802.1D and 802.1p. Now all in 802.1Q with Virtual LANs (VLANs) and priority queues
- Better but still gaps for real-time and safety-critical systems:
  - Priority inversions in FIFO queues
  - Interference through shared memory
  - Additional forwarding delay with address learning and flooding
  - Delays vary with switch technology
  - "Plug and Play" has been the mindset: lots of protocols that are not appropriate for real-time (long delays and nondeterminism)

Ethernet frame format

802.3 Ethernet packet and frame structure

Layer	Preamble	Start of frame delimiter	MAC destination	MAC source	802.1Q tag (optional)	Ethertype (Ethernet II) or length (IEEE 802.3)	Payload	Frame check sequence (32-bit CRC)	Interpacket gap
	7 octets	1 octet	6 octets	6 octets	(4 octets)	2 octets	46-1500 octets	4 octets	12 octets
Layer 2 Ethernet frame			← 64–1522 octets →						
Layer 1 Ethernet packet & IPG	← 72–1530 octets →								← 12 octets →

802.1Q tag format

16 bits	3 bits	1 bit	12 bits
TPID	TCI		
	PCP	DEI	VID

- Unicast frames (Direct addressing of the (only) receiver)
- Multicast frames (when there are multiple receivers)
  - Some reserved for Layer 2 protocols (e.g., Spanning Tree Protocol and Precision Time Protocol)
  - AVB and TSN streams are typically multicast
- Broadcast frames (FF:FF:FF:FF:FF:FF)
  - Flood frame on all ports

MAC address

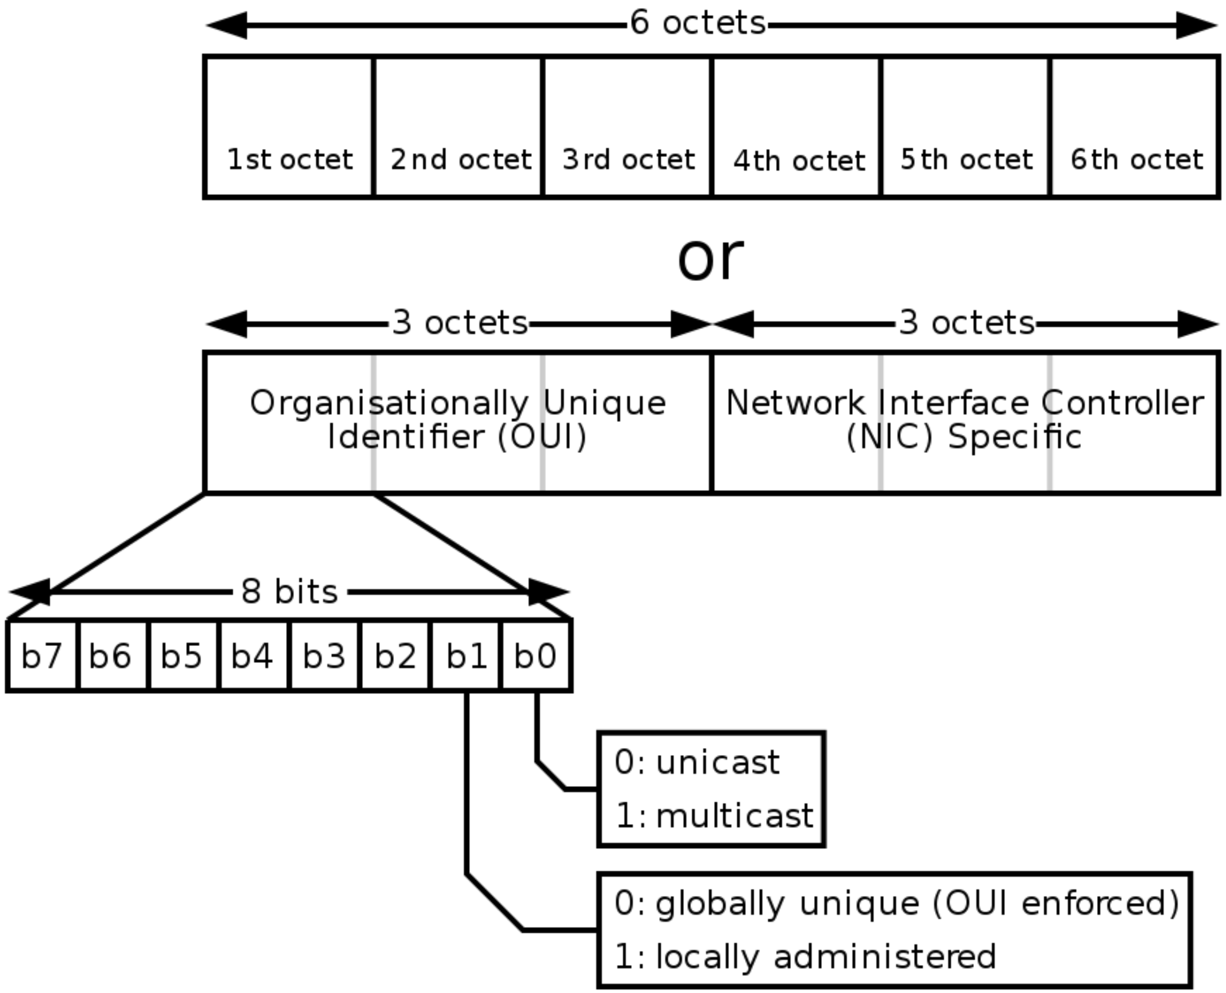
6 octets
1st octet
2nd octet
3rd octet
4th octet
5th octet
6th octet
or
3 octets
3 octets
Organisationally Unique Identifier (OUI)
Network Interface Controller (NIC) Specific
8 bits
b7
b6
b5
b4
b3
b2
b1
b0
0: unicast
1: multicast
0: globally unique (OUI enforced)
1: locally administered

Example network

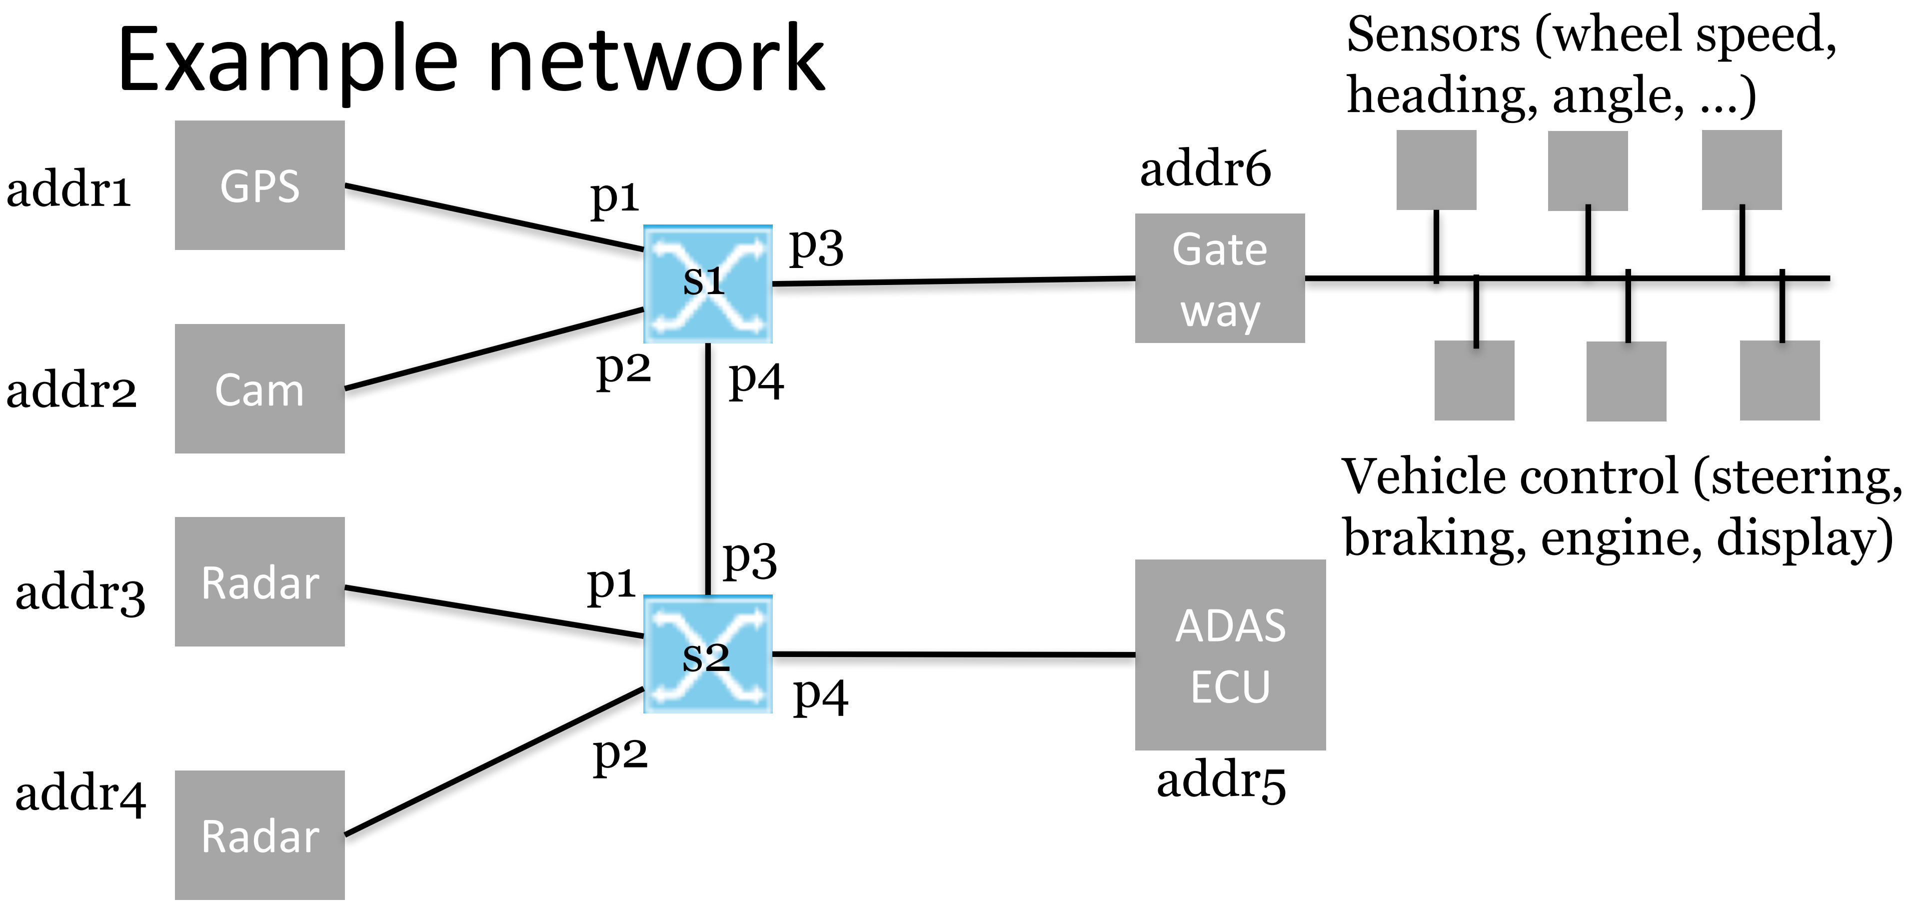

Example network

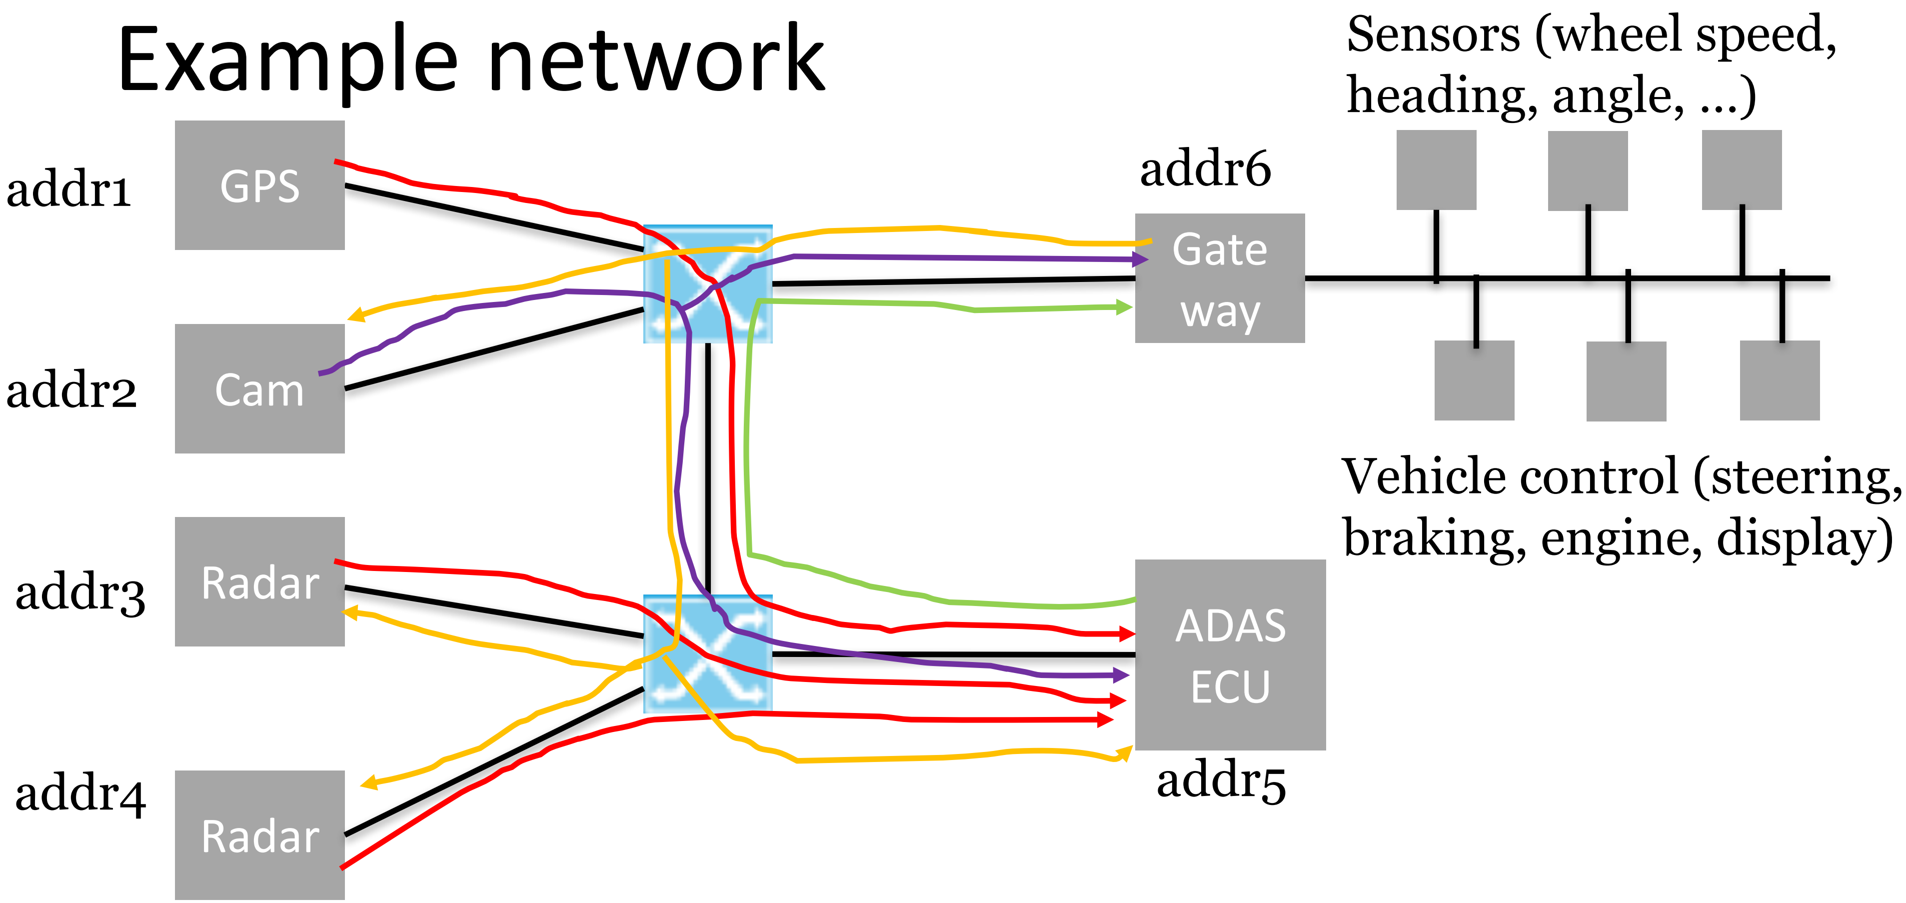

Streams:

- GPS and Radars to ADAS ECU
- Cam to ADAS ECU and Vehicle domain
- Vehicle sensors to Cam, Radars, and ADAS ECU
- ADAS ECU to Vehicle domain (actuation and corrected GPS)

Forwarding tables

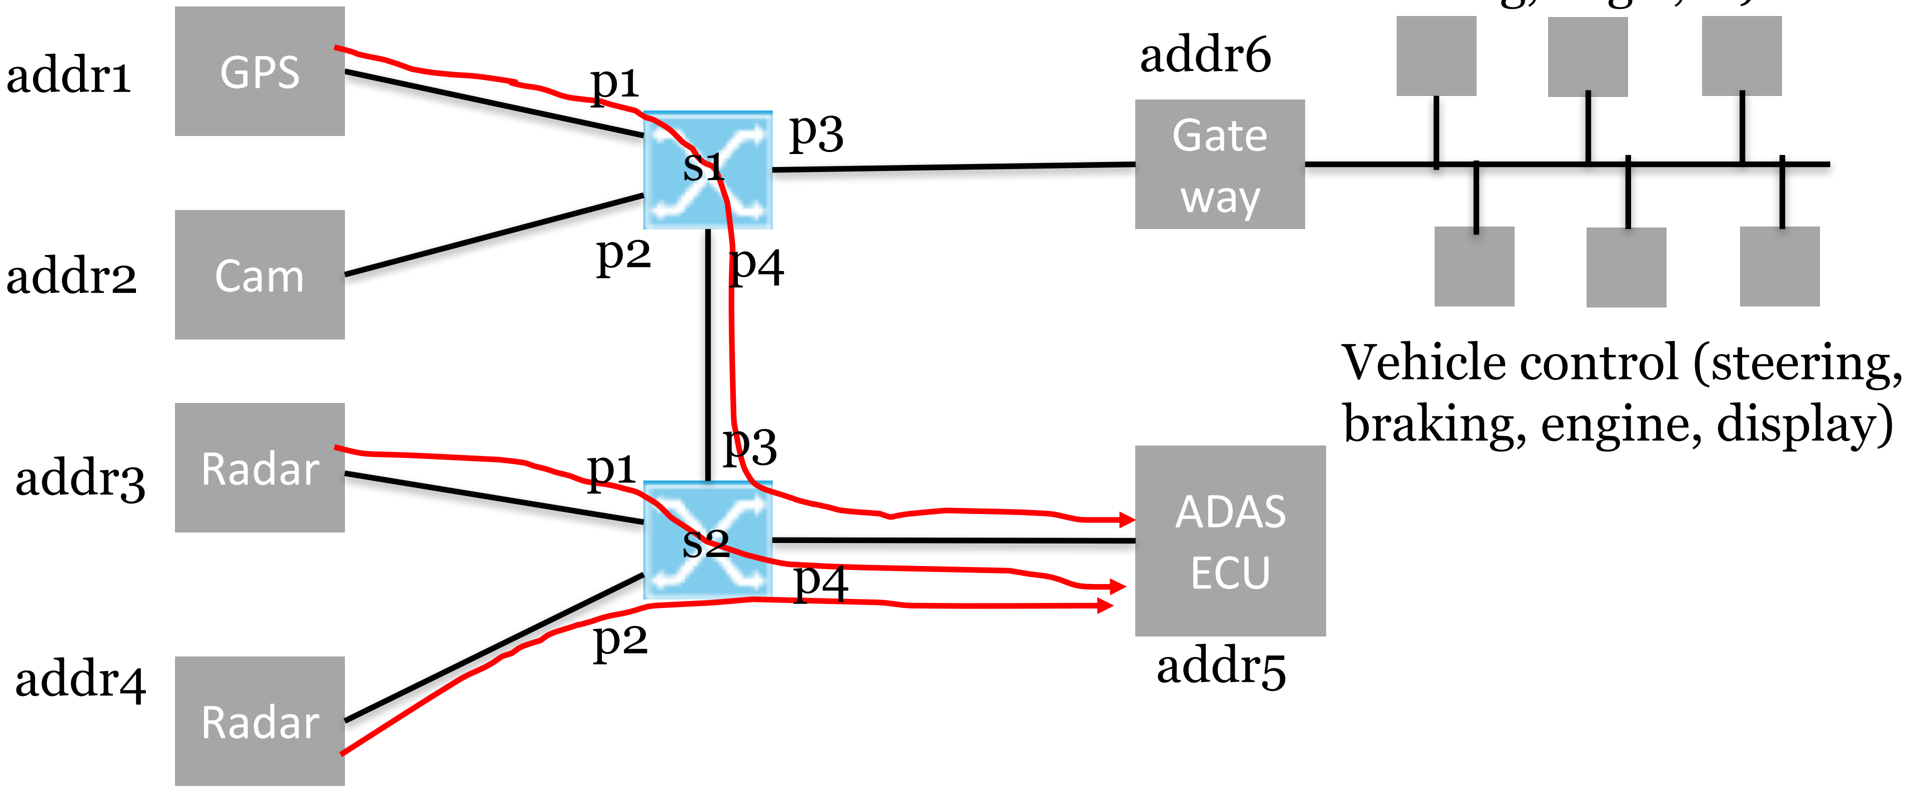

Switch s1 (DA -> port):
addr5 -> p4

Switch s2:
addr5 -> p4

Forwarding tables

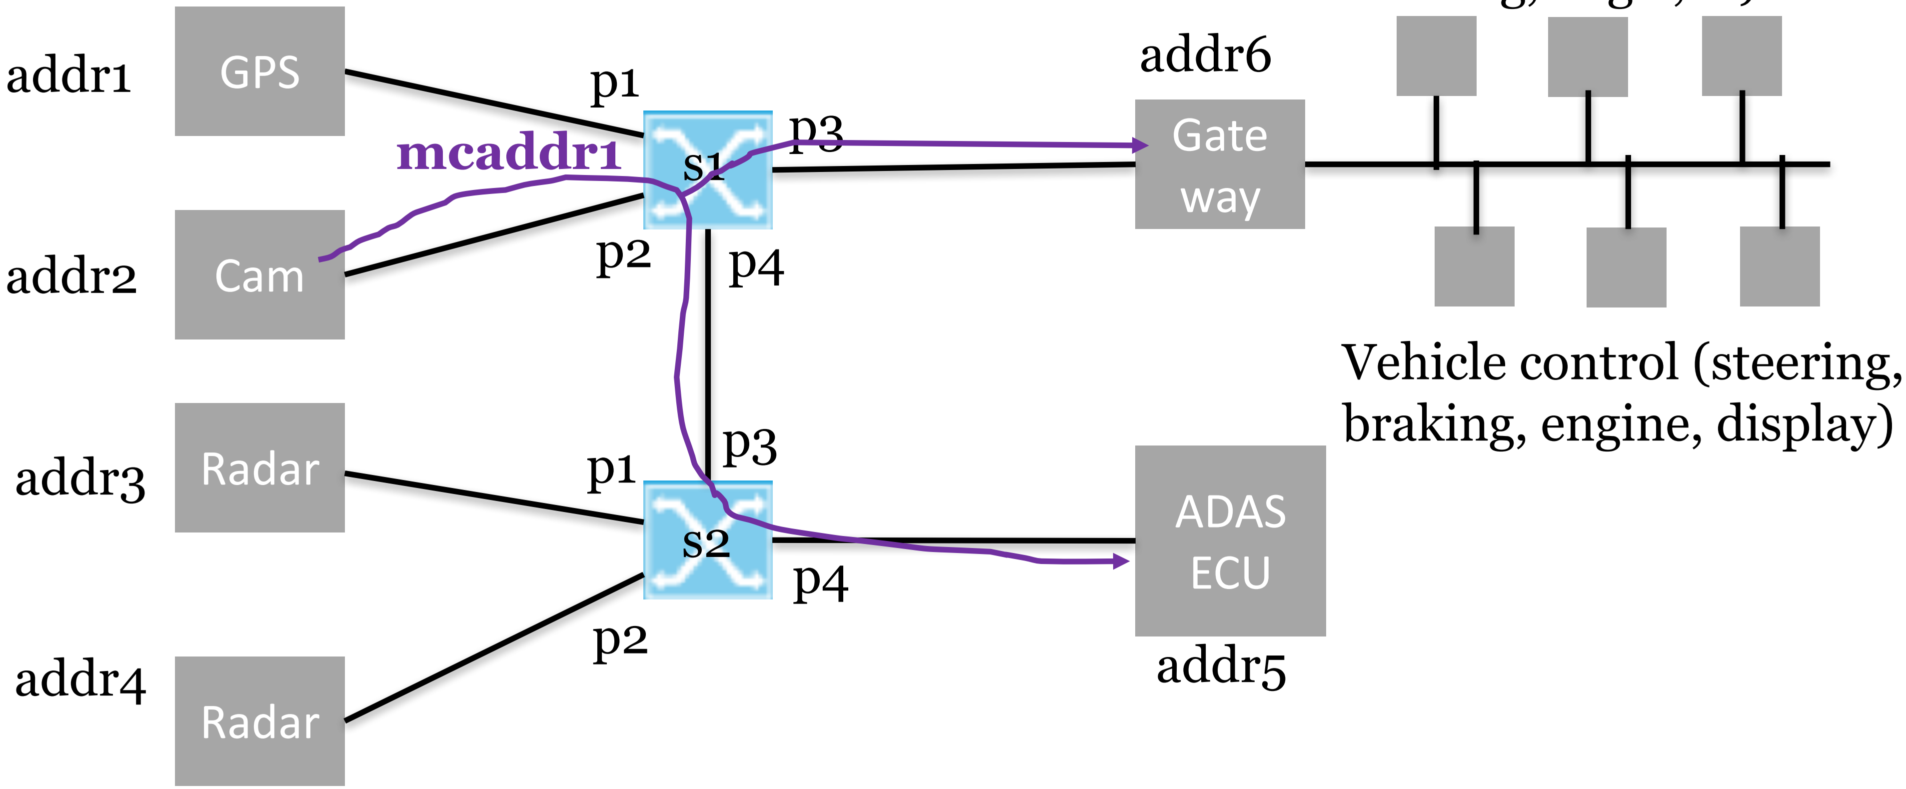

Switch s1 (DA -> port):
addr5 -> p4
mcaddr1 -> p3, p4

Switch s2:
addr5 -> p4
mcaddr1 -> p4

Forwarding tables

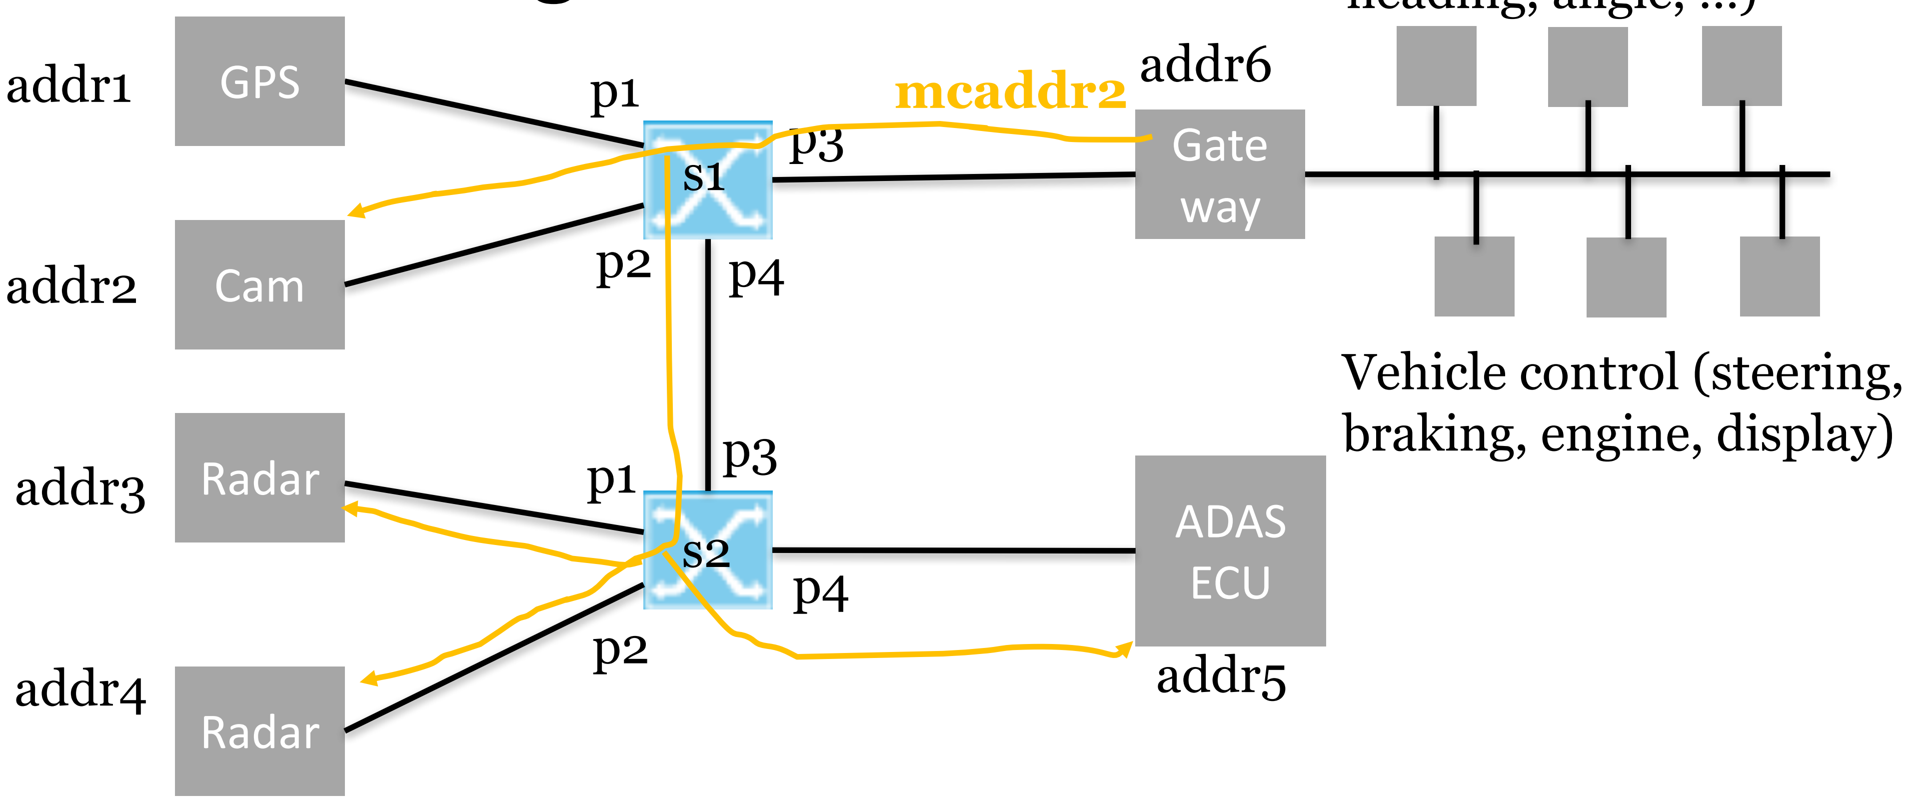

Switch s1 (DA -> port):
addr5 -> p4
mcaddr1 -> p3, p4
mcaddr2 -> p2, p4

Switch s2:
addr5 -> p4
mcaddr1 -> p4
mcaddr2 -> p1, p2, p4

Forwarding tables

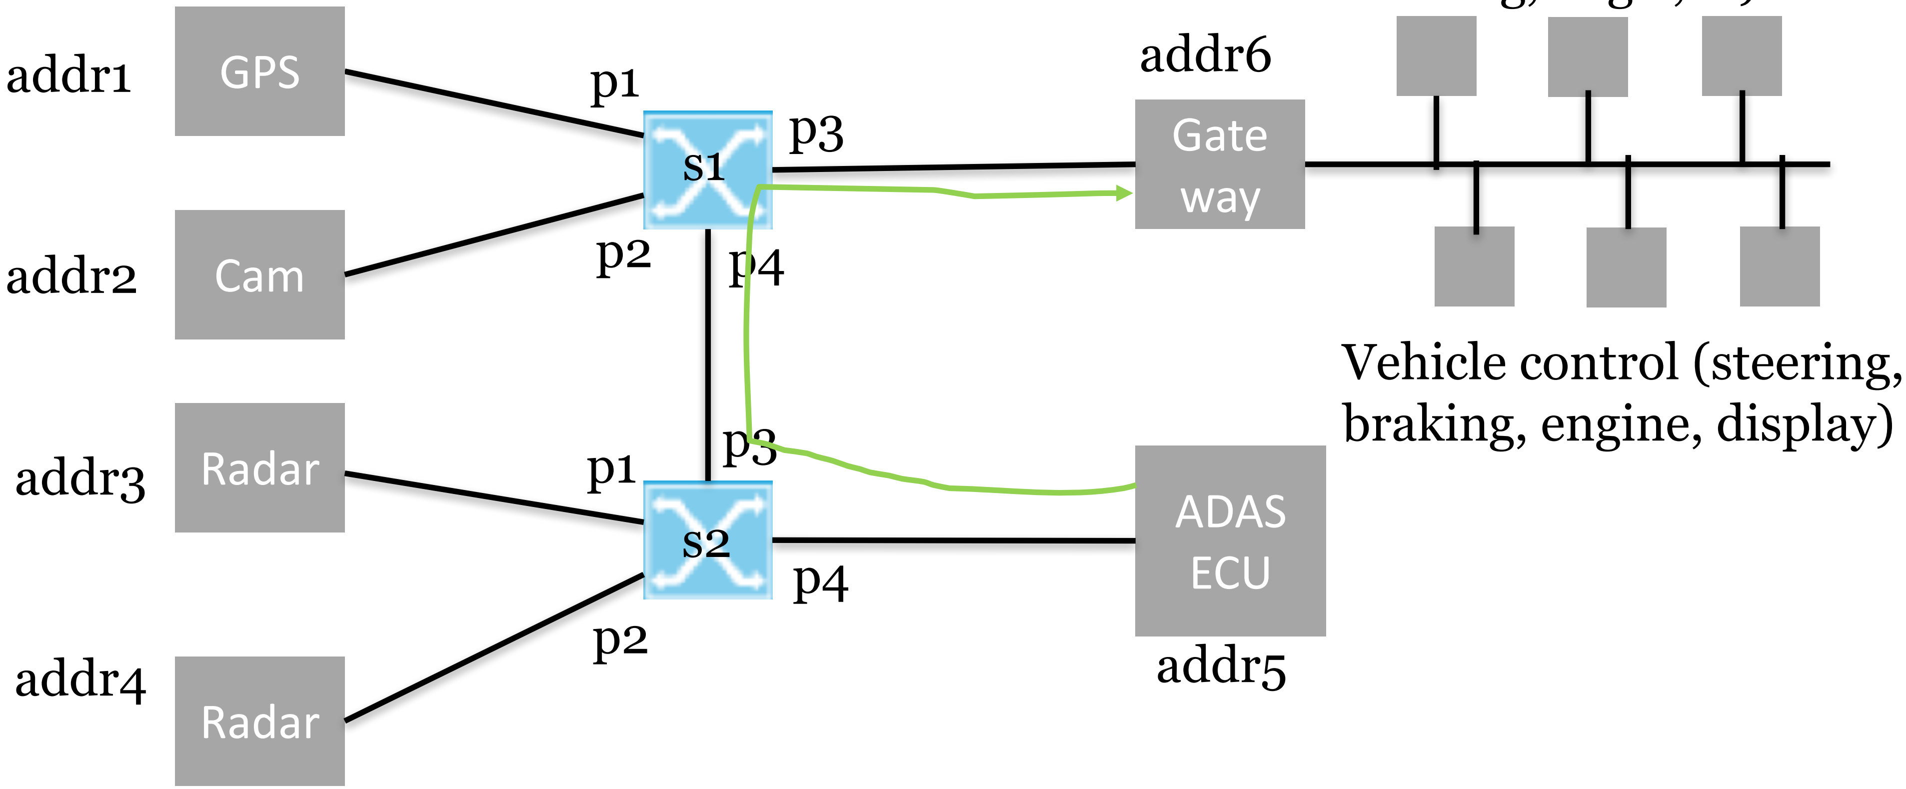

Switch s1 (DA -> port):
addr5 -> p4
mcaddr1 -> p3, p4
mcaddr2 -> p2, p4
addr6 -> p3

Switch s2:
addr5 -> p4
mcaddr1 -> p4
mcaddr2 -> p1, p2, p4
addr6 -> p3

Table can also be a mapping from DA and VLAN ID to ports

Configuration of forwarding tables

- For plug & play networks:
  - Flooding
  - Learning based on mapping ingress ports to source addresses (unicast)
  - Multiple stream registration protocol (multicast)
- Engineered networks (e.g., automotive, avionics)
  - Static configuration of forwarding tables
  - Protocols for learning are switched off (also helps to reduce risks for cybersecurity attacks)

Ethernet Switch

- Priority assigned to frame
- Priorities are mapped to queues
- Frame assigned to queue
- First in, first out (FIFO)

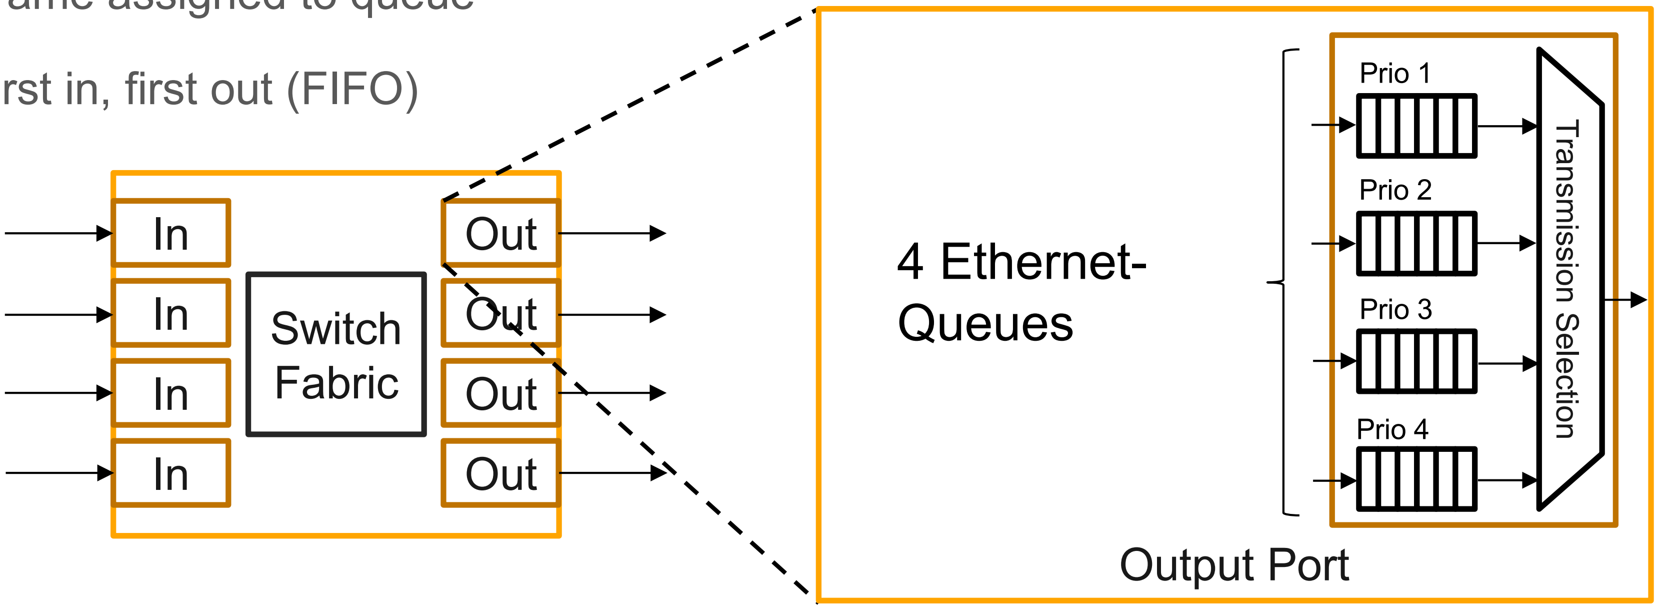

Classification of frames, thus different treatment is possible

Prioritization

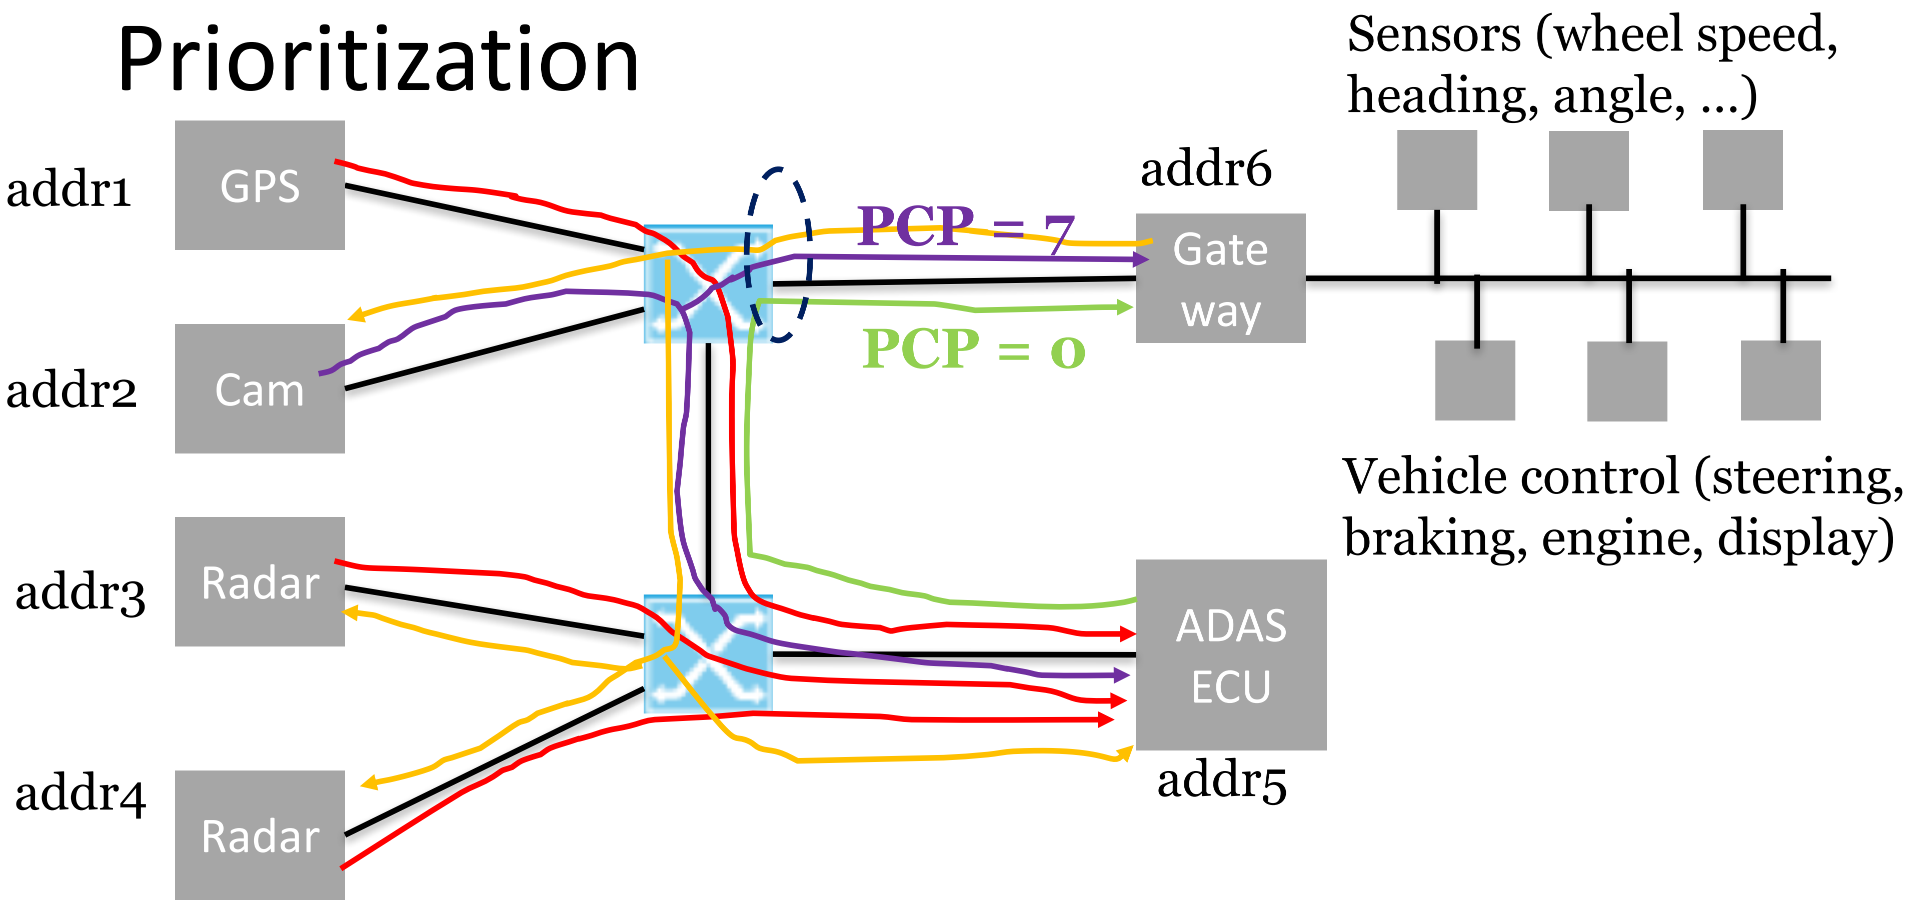

Streams:

- GPS and Radars to ADAS ECU
- Cam to ADAS ECU and Vehicle domain
- Vehicle sensors to Cam, Radars, and ADAS ECU
- ADAS ECU to Vehicle domain (actuation and corrected GPS)

Different topologies

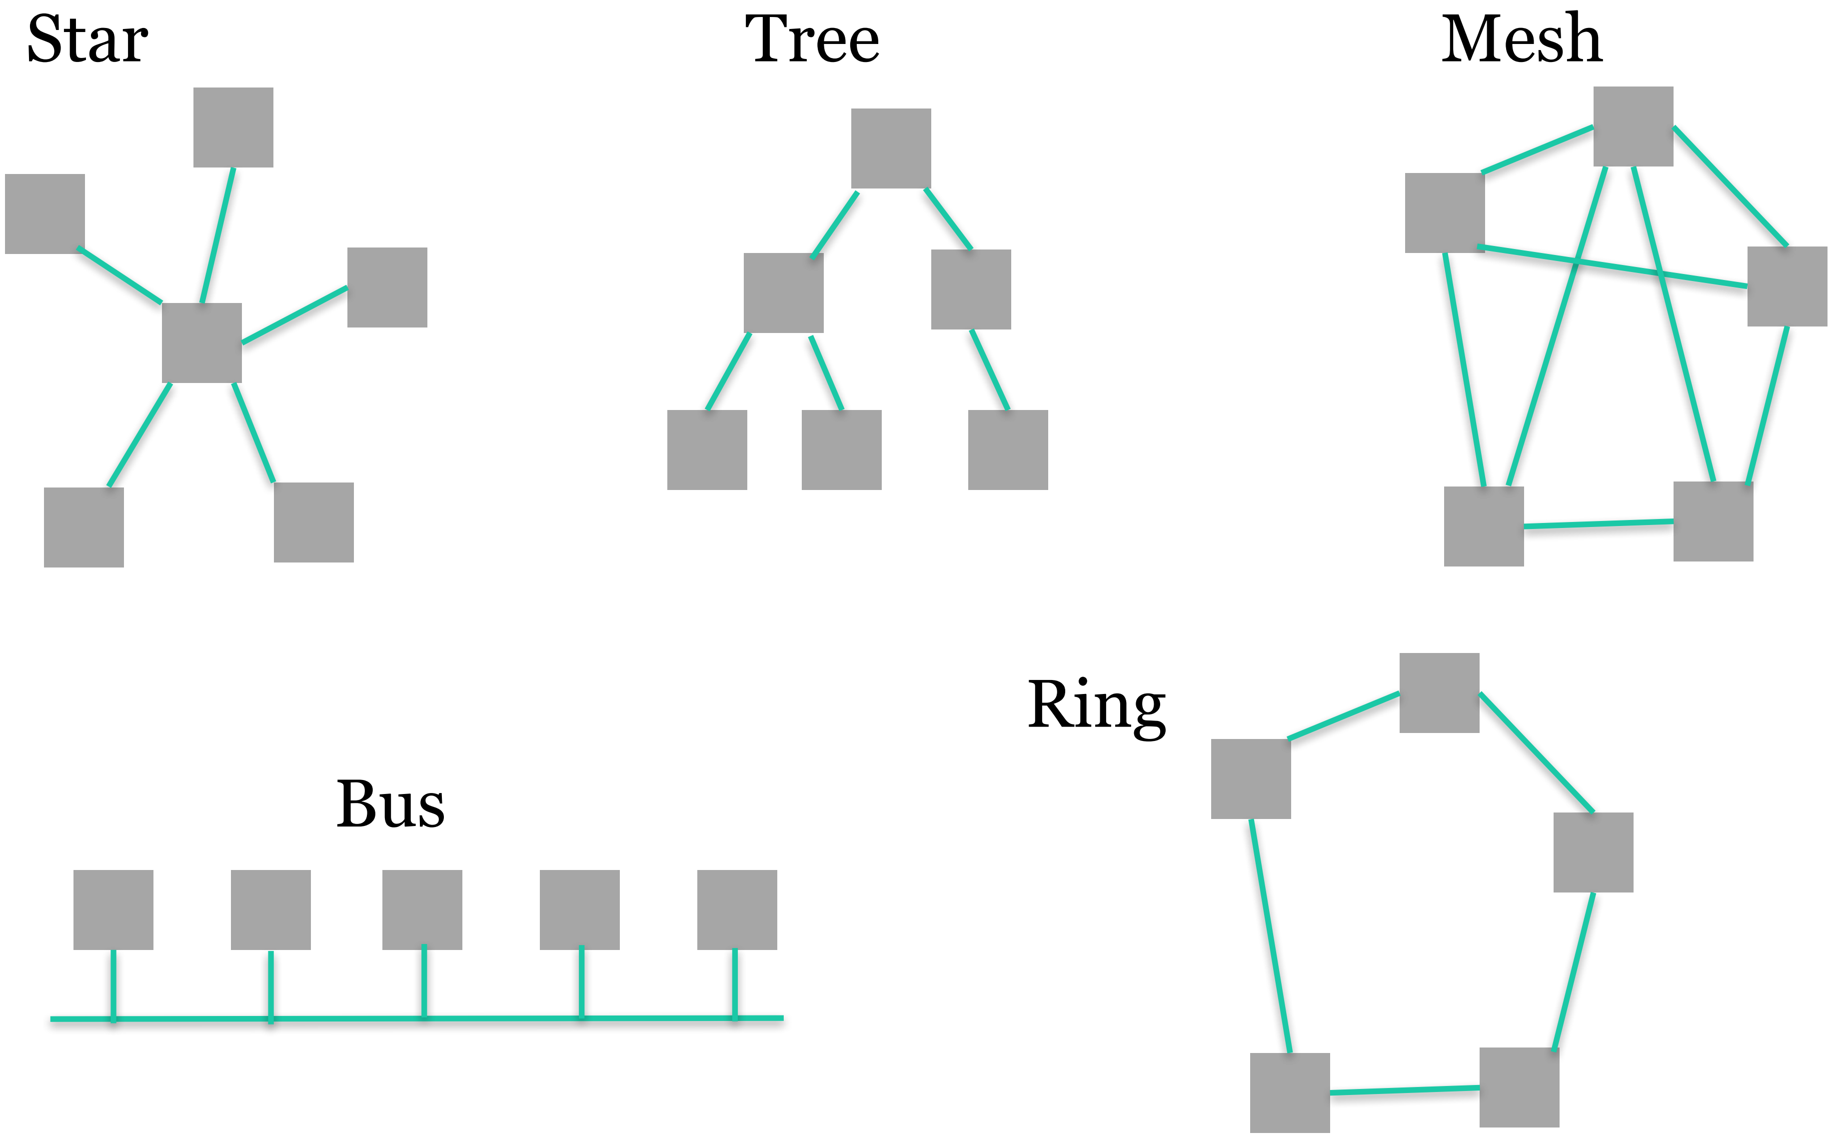

Star topology

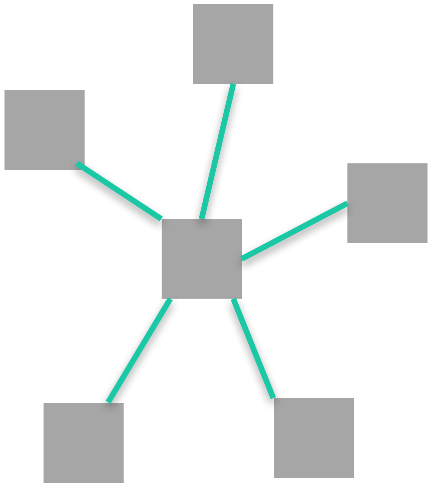

Some observations:

- One hop communication to all (low latency)
- Low cost
- Switch failure disconnects all communication
- Long wires, depending on physical location of endpoints

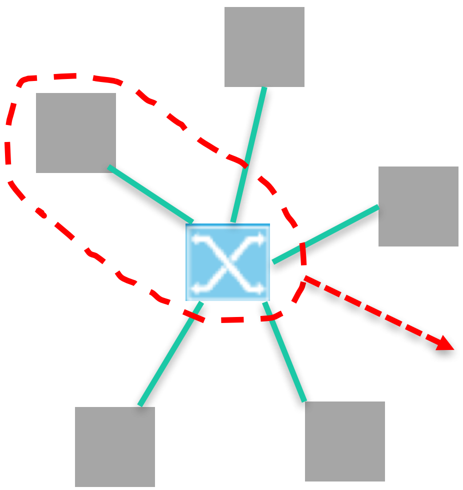

- Can be integrated in one module with MII, no PHY
- Integration is product and topology dependent

Tree topology

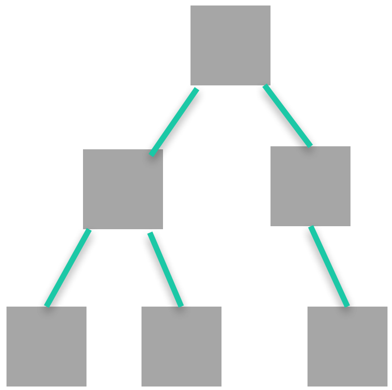

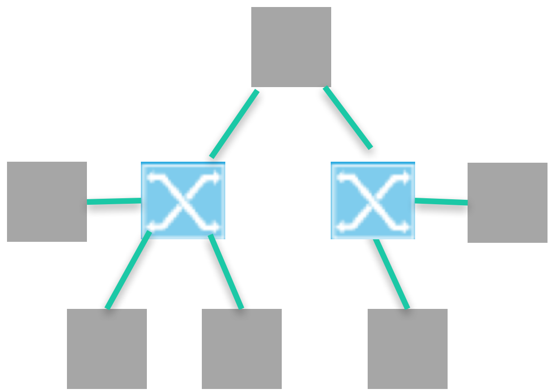

Some observations:

- Multiple hops between some nodes
- Switch failure disconnects several nodes
- Good for networks where communication mainly goes in one direction
- Segmented left and right side

Mesh topology

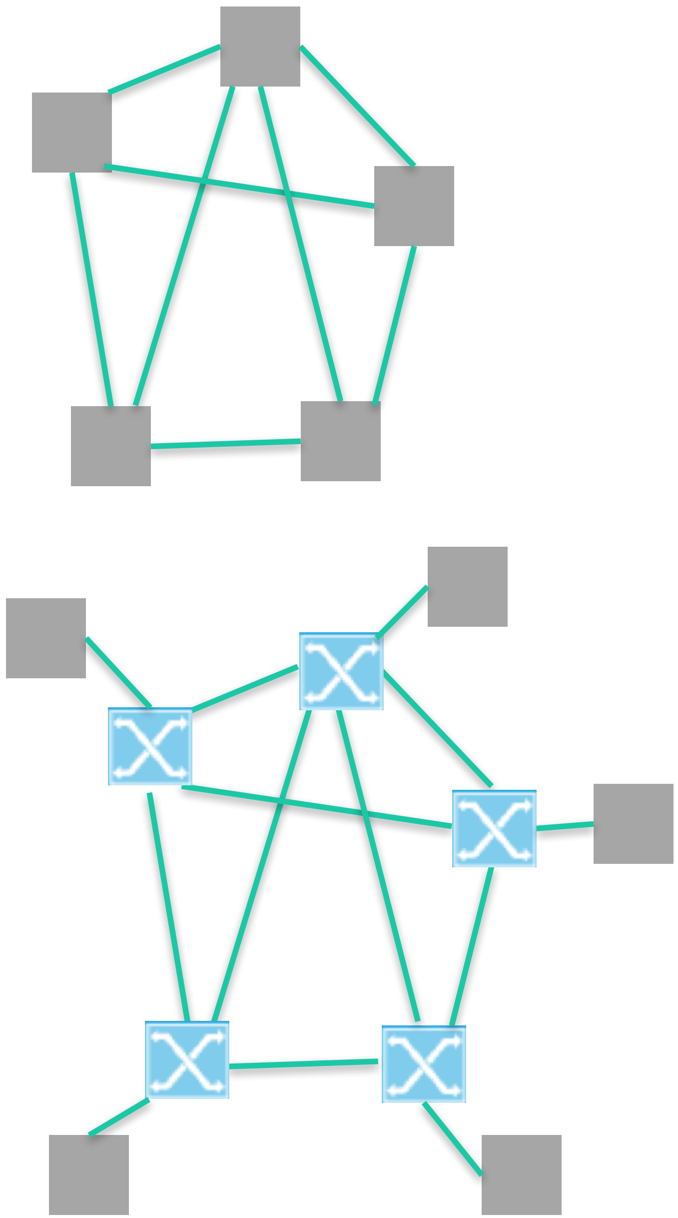

Some observations:

- Multiple hops between some nodes
- Reduced impact of one switch failure
- Multiple communication paths (easier to meet timing and redundancy needs)
- High cost due to many switches

Ring topology

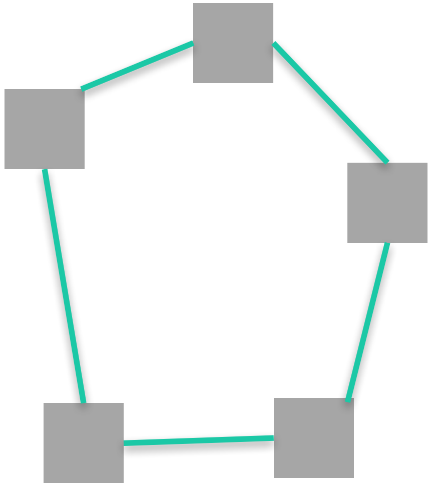

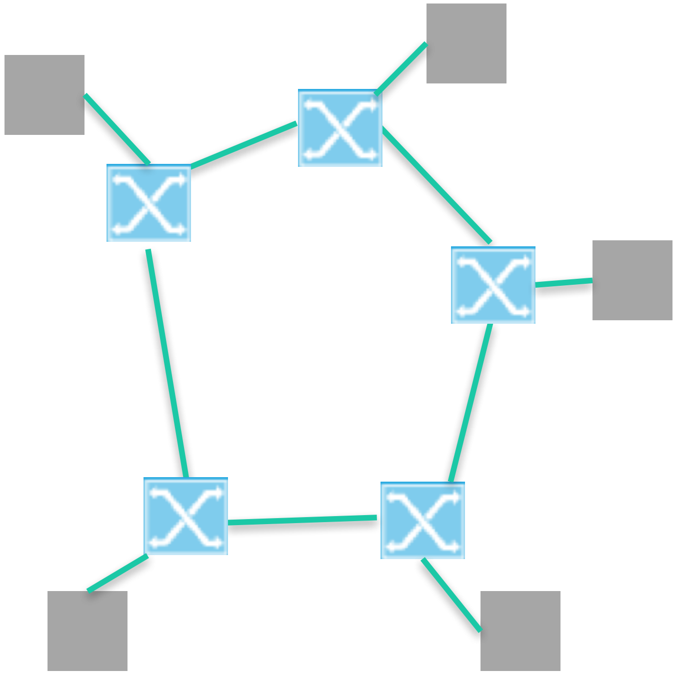

Some observations:

- Latency does not scale well with number of nodes
- Redundancy at lower cost than mesh solution
- Still many switches, but of lower port count
- Switch failure “only” disconnects one node

Bus topology

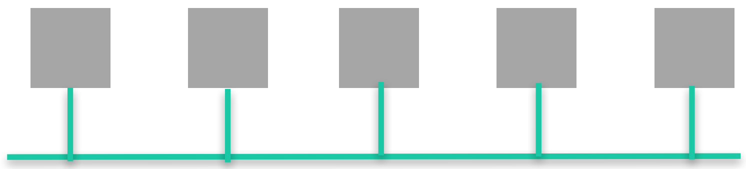

Some observations:

- Low cost (no switches)
- Single broadcast domain
- Shared media access
- Requires non-IEEE 802 mechanisms to handle (or avoid) collisions for real-time applications

Network is usually a mix of different topologies, as well as, technologies

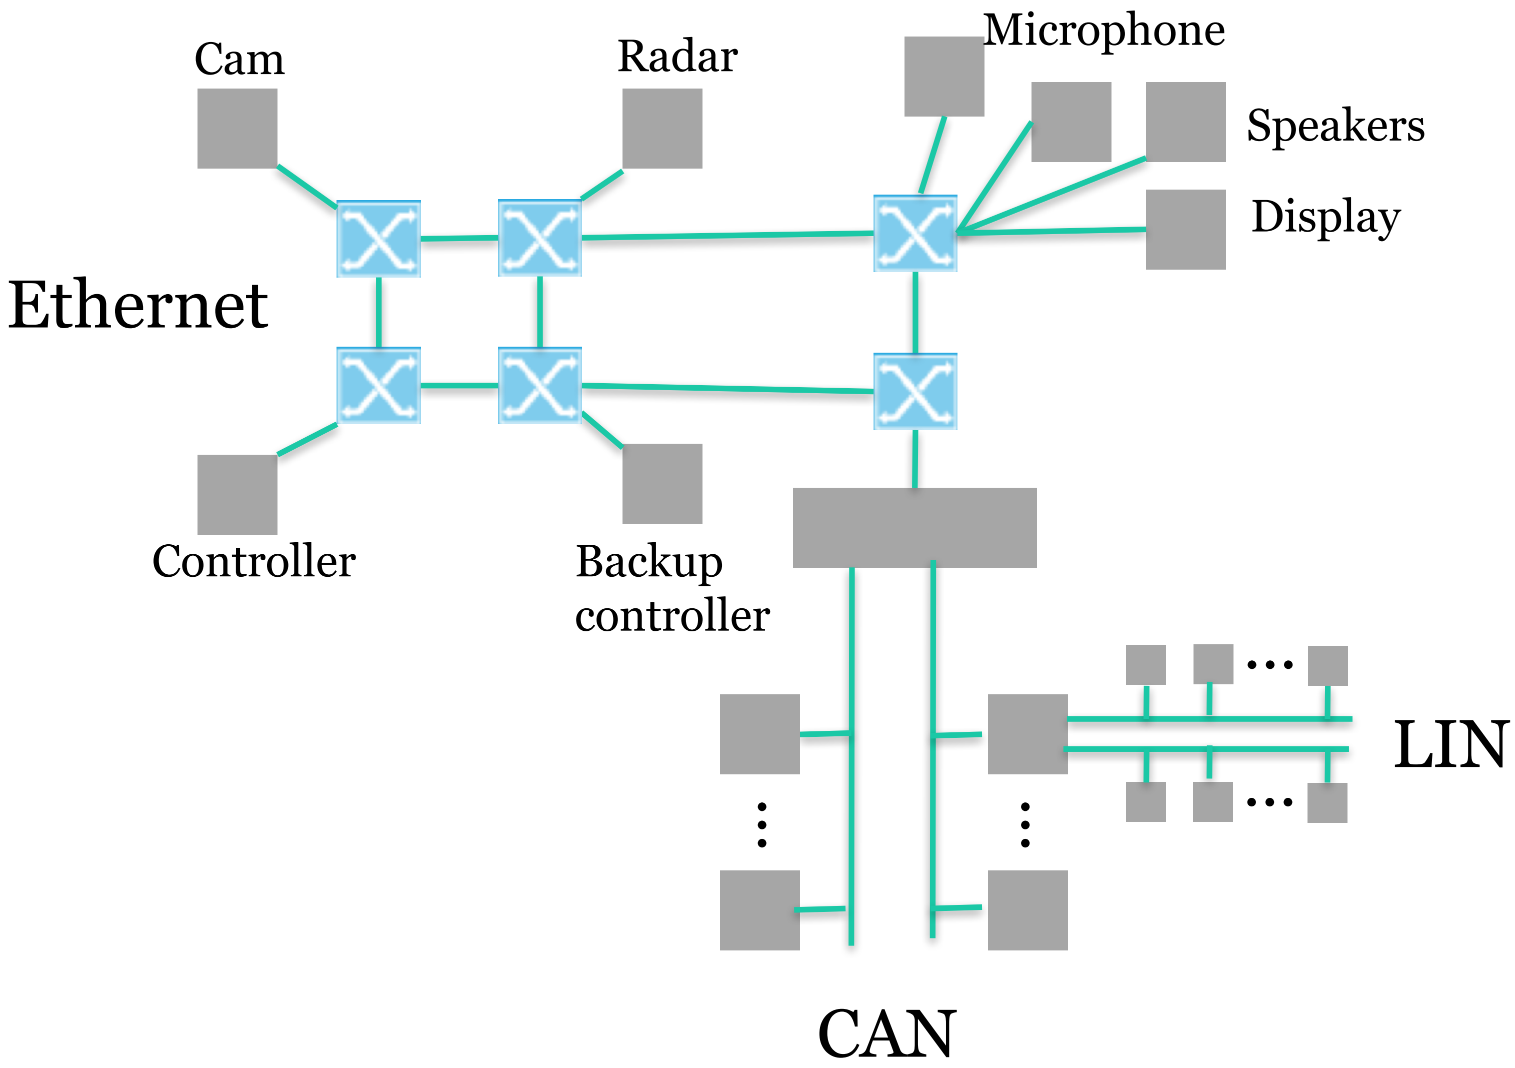

Ethernet PHYs

- Several standards have been created in IEEE 802.3
  - Full-duplex 100 Mbps and 1 Gbps Ethernet on single UTP (unshielded twisted pair) cable
  - Full-duplex 1+ Gbps is under development (e.g., we may end up with 2.5 Gbps and 10 Gbps)
  - 10 Mbps is under development too
    - "Multi-drop" (i.e., half-duplex, bus topology)
    - Some companies are looking at this as a replacement of CAN and FlexRay for the "All Ethernet car"

Adding bandwidth is not enough

- Support for real-time and safety-critical applications
- Several industries with similar communication requirements at OSI Layer 2 and above
  - Professional audio/video
  - Industrial automation
  - Automotive
  - Telecom
- Standardization development in IEEE 802.1
  - First, AVB. Now, TSN
  - Multi-hop, switched topologies make certain problems more complicated bus-based RT communication

AVB (Audio/Video Bridging)

- Main drivers:
  - synchronized audio and video applications on Ethernet
  - Plug and play
- IEEE 802.1AS: Plug and play Clock synchronization
- Amendments to IEEE 802.1Q:
  - Stream reservation protocol (admission control)
  - Credit-based traffic shaping (guaranteed bandwidth across priority levels; no starvation; zero congestion loss / no dropped packets)
- IEEE 802.1BA: Umbrella document

What was missing?

- Guaranteed lower latencies; searching for the lowest possible latency
- Error detection and isolation
- Redundant communication
- Redundant clock synchronization

New TSN standards

- TSN: Group renamed itself from “Audio/Video Bridging” to “Time Sensitive Networking”

- Guaranteed lower latencies; searching for the lowest possible latency
  - 802.1Qbv: A priori defined, pre-scheduled communication
  - 802.1Qbu: Frame preemption
  - 802.3br: Preemption required changes in MAC
  - P802.1Qcr: Asynchronous traffic shaping

New TSN standards

- Error detection and containment
  – 802.Qci: Monitoring of pre-defined “contract” between a data flow and the network
  – Meaning of “Contract” depends on the scheduling policy
- Redundant communication
  – 802.1CB: Frame replication and elimination
- Redundant clock synchronization
  – P802.1AS-REV: updates to the clock synchronization standard to support backup masters and multiple time domains

Security standards

- 802.1 has developed three main standards for secure Ethernet communication
- 802.1AE (“MACsec”): integrity and, optionally, privacy, is ensured by symmetric key crypto (128 and 256 bit keys are supported)
- 802.1X:
  - Port authentication
  - Key agreement protocol

Security standards

- 802.1AR (Secure Device Identity):
  - Device identity based on public key crypto (elliptic curves)
  - Public Key Infrastructure and certificates

- These standards, in particular MACsec and Secure Device Identity, require parallel crypto logic on chip
  - Challenging in terms of cost, power, and size
  - Automotive industry has started to look at how to secure Ethernet communication and which portions of the 802.1 security standards apply

TSN toolbox

Traffic shaping
- Credit based shaper .1Qav
- Time aware shaper .1Qbv
- Cyclic queuing and forwarding .1Qch
- Asynchronous Traffic Shaping .1Qcr

Resource reservation
- Stream reservation protocol .1Qat
- Stream Reservation Protocol Enhancements and Performance Improvements .1Qcc
- Path control and reservation .1Qca

Dependability
- Frame Replication and Elimination for Reliability .1CB
- Per-Stream Filtering and Policing .1Qci

Low latency
- Interspersing Express Traffic .3br
- Frame Preemption .1Qbu

Time sync
- Generalized Precision Time Protocol .1AS
- Enhancements for gPTP .1AS-Rev

TSN application areas

- Industrial control
- Avionics
- Automotive
- Telecom
- Aerospace
- Edge computing (a.k.a. “Fog” computing)
  - Industrial automation
  - Intelligent transport systems (road-side units)

Many non-standard solutions

- Ethernet POWERLINK (Master/Slave protocol; open source)
- EtherCAT (2-port switch in each node; cut-through forwarding
  - Beckhoff automation (IEC 61158)
- PROFINET (TDMA)
  - Siemens
- AFDX (static configuration, traffic shaping) (Airbus)
- TTEthernet by TTTech (TDMA at message level, and rate-constrained traffic class).

Summary

- Ethernet is a point-to-point network
- Direct addressing with unicast and multicast
- Switches forward frames towards the destination(s)
- Ports have multiple queues assigned to different priority levels
- Many different Ethernet topologies
- AVB and TSN standards added real-time capabilities
- There are also nonstandard solutions deployed