Communication in Avionics

Outline

- Basic Overview
- Communication architectures
  - Event Triggered
  - Time Triggered
- Communication architecture examples
- Case Study: How Data Communication Affects Scheduling

Communication

- A typical airplane includes over one hundred avionics components (computers) communicating with each other
    - ➡ each dedicated to a specific function

Generic Communication Architecture

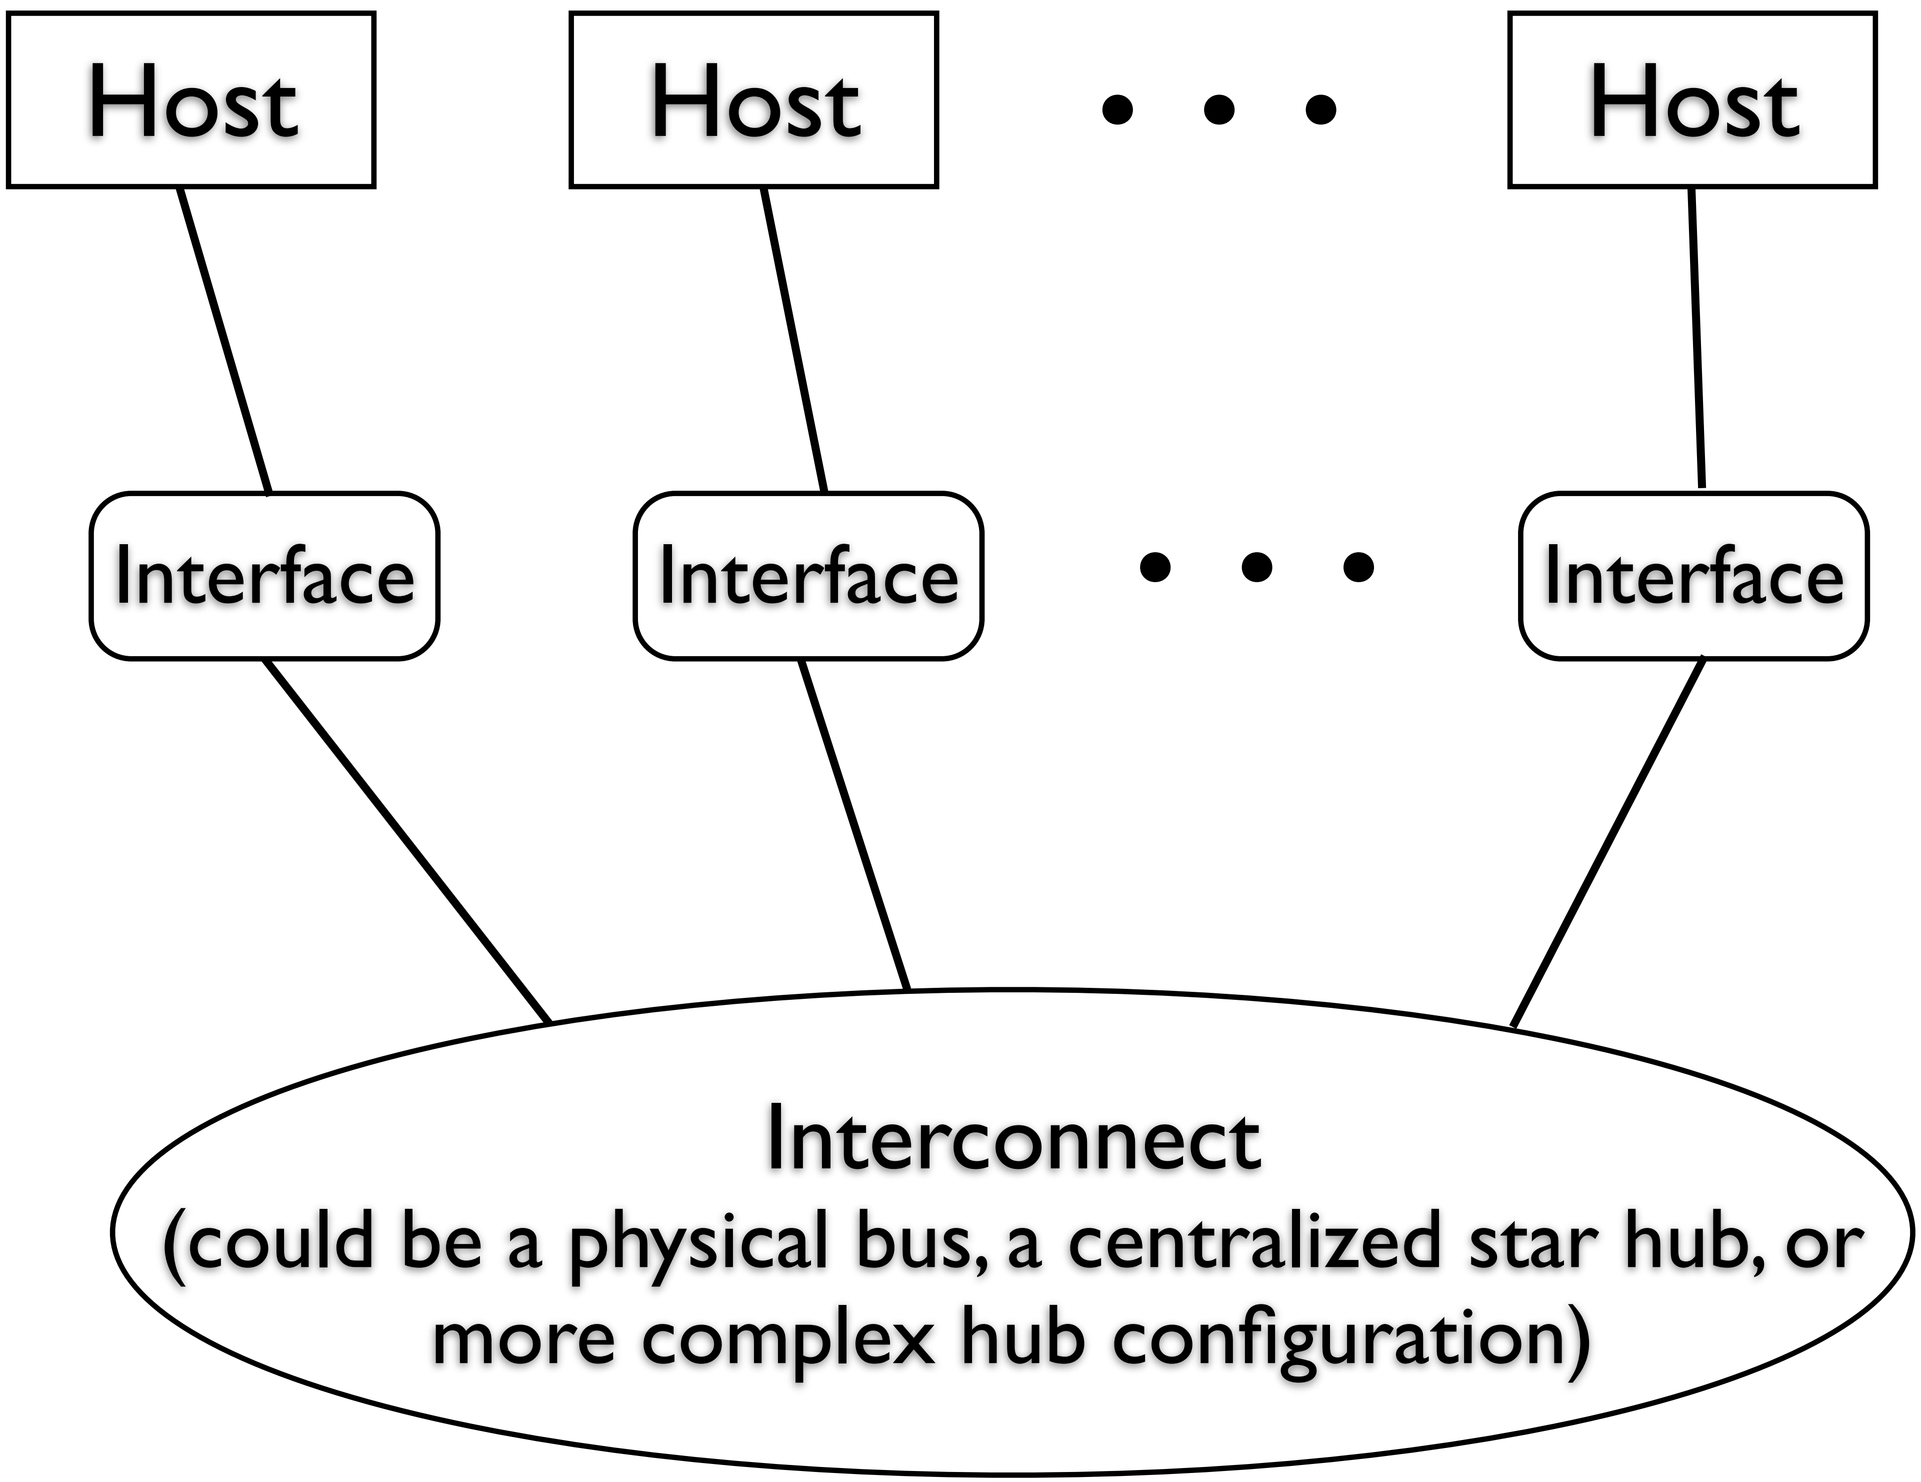

Outline

- Basic Overview
- Communication architectures
  - Event Triggered
  - Time Triggered
- Communication architecture examples
- Case Study: How Data Communication Affects Scheduling

Event-Triggered Communication

- Messages are generated based only on the need to transmit a new or changed piece of information
  - ➡ Based on an event that can occur at any time
  - ➡ E.g., "if the sensor reading changes, then broadcast its new value"

Event-Triggered Communication Protocols

- Only the sender has knowledge about the time point when a message has to be transmitted
  - ➡ Thus the error detection is based on a timeout at the sender waiting for an acknowledgement message from the receiver
  - ➡ In case of a message lost, k-fault-tolerance is achieved by retransmitting a lost message k times.
  - ➡ Need congestion control mechanisms to avoid congestion in case of coincident events

Scheduling in Event-Triggered Communication Systems

- Event-triggered systems are demand driven and thus require a dynamic scheduling strategy
- In general the problem of deciding whether a set of real-time tasks whose execution is constrained by some dependency relation can be correctly scheduled is NP-hard
- In practice most ET-systems resort to static priority scheduling
    - ➡ No analytical guarantees about the peak load performance can be given

Example ET Communication System: Controller Area Network (CAN)

- A multi-master broadcast serial bus standard for connecting electronic control unit (ECU) (e.g., cruise control)
    - ➡ A master node issues the clock and addresses non-master nodes
    - ➡ Multi-master means that any number of master nodes can exist
- Each node is able to send/receive messages, but not simultaneously
- Nodes are not connected directly to the bus, but through a host processor and a CAN controller
    - ➡ The host processor decides what received messages mean and which messages it wants to transmit
    - ➡ The CAN controller stores received bits serially from the bus

Layered Structure of a CAN node

Application Layer

Object Layer

- Message Filtering
- Message and Status Handling

Transfer Layer

- Fault Confinement
- Error Detection and Signaling
- Message Validation
- Acknowledgement
- Arbitration
- Message Framing
- Transfer Rate and Timing

Physical Layer

- Signal Level and Bit Representation
- Transmission Medium

Data Transmission in CAN

- If the bus is free, any node may begin to transmit; If two or more nodes begin sending messages, only the message of the highest priority node remains and received by all nodes
- Priority is defined by the arbitration ID
- Arbitration free transmission
  - ➡ Lowest arbitration ID first

CAN Benefits

- Each ECU can have a single CAN interface rather than analog and digital inputs to every device
- Each device has a CAN controller chip and therefore intelligent
- All devices see all transmitted messages
    - ➡ Can decide if a message is relevant or it should be filtered

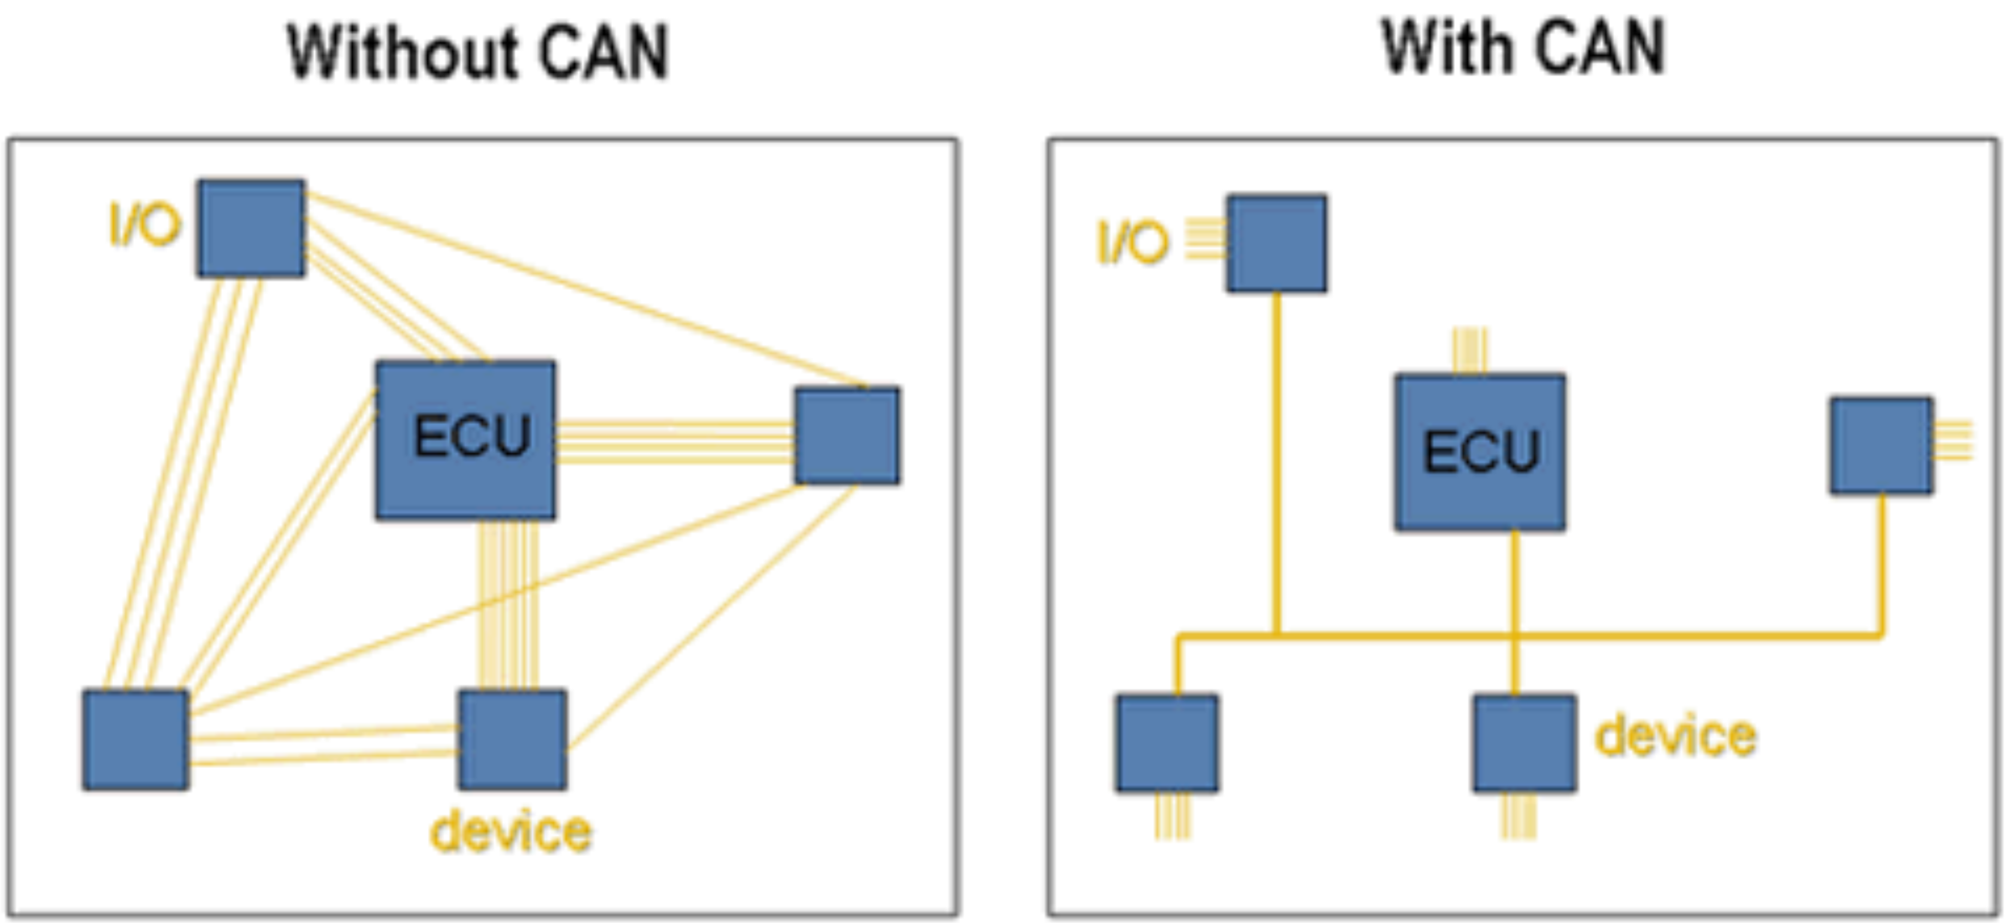

Time-Triggered Communication

- Occurs at specified times based on a globally agreed upon time base.
    - E.g., "if it's 20ms since the start of the frame, then read the sensor and broadcast its value"
- Scheduled with the passage of time and each node that is part of the network is given a finite amount of time, in which to transmit a message in each communication cycle.

Time-Triggered Communication

- Static preallocation of communication bandwidth in the form of a global schedule
  - ➡ each node knows the schedule and the time
  - ➡ time to send/receive messages
  - ➡ contention solved offline

Time-Triggered Communication Protocols

- The communication protocol in a TT-system is a reliable broadcast protocol
    - ➡ A protocol makes it possible for the receiver to detect the loss of any broadcast message
    - ➡ Such error detection is performed by the receiver of the information based on the global knowledge about the expected arrival time of each message

Scheduling in Time-Triggered Communication Systems

- Based on a set of static predetermined schedules
- Schedules have to consider all necessary task dependencies and provide an implicit synchronization at run time
- Cyclic executive schedule

Time-Triggered Communication

- Provide efficiency, determinism, and timing guarantees, but at the price of flexibility
  - ➡ Switch among several schedules optimized for different missions
  - ➡ Timesharing between time- and event-triggered operation

Example TT communication: Time-Triggered CAN

- Under the bitwise arbitration of CAN, the access may be delayed, if some other message is already in the process of transmission or if another message with higher priority also competes for the bus.
- The lower the priority of a message is, the higher the latency jitter for the media access
- Even the message with the highest priority may experience a small latency.

The Goals of TTCAN

- Reduce latency jitters
- Guarantee a deterministic communication pattern on the bus
- Use the physical bandwidth of a CAN network much more efficiently
- Developed by Bosch
- Standardized in ISO 11898

Extensions to CAN

- Guarantees the time triggered operation of CAN based on the reference message of a time master
- Fault tolerance of that functionality is established by redundant time masters (potential time masters)
- A globally synchronized time base is established
- A continuous drift correction among the CAN controllers is realized

TTCAN Operating Principles

- TTCAN is based on a deterministic temporal exchange, based on a time window of a pre-determined operating cycle
- All the messages traveling in the network between the CAN nodes are organized as elements of an X x Y matrix
  - ➡ The matrix defines the relationship between the time windows and the presence of messages in the network
- One of the nodes of the network is responsible for organizing the time division and allocation involved

TTCAN Mechanisms

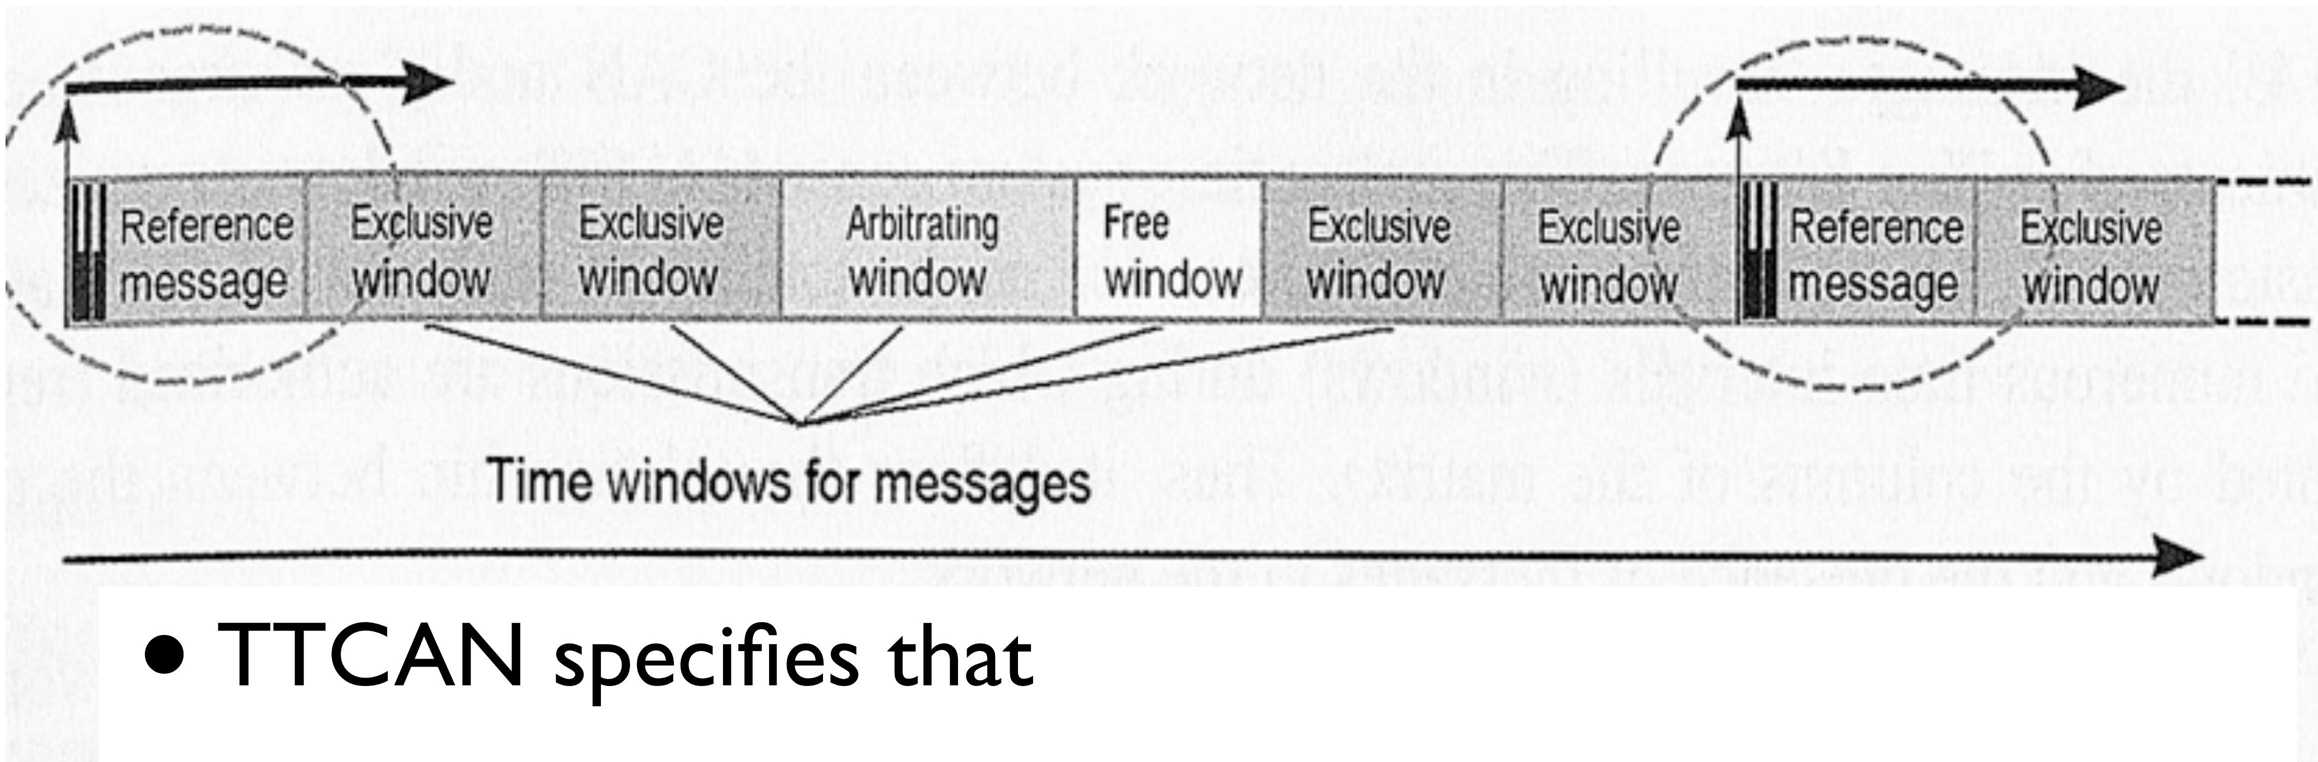

- TTCAN specifies that
  - periodic messages are included in exclusive time windows,
  - Occasional messages are included in arbitrating time windows
  - Free time windows are reserved as free spaces for all traffic
- The time sequence often has to be reconfigured according to external events, whether expected or unexpected

An Example: a 5-node TTCAN

- Suppose I am node 1, who is the controller of the system; I decide unilaterally to allocate the following time phases to the other four participants in the network.
- To do this, I start by sending them a minimum information required for the correct operation of the system, by means of a special generic message called a 'reference message', using a special identifier, indicating the following
- The time interval of the basic cycle will be 1 minute from now on and until further instructions. This is the time window sequence for each of you:
- Node 2, you do not have much to say, so you can talk from 00 to 05;
- I, Node 1, am known to be very talkative, so I shall speak from 05 to 20;
- Node 3, you will talk from 20 to 25;

An Example: a 5-node TTCAN

- All of you can talk from 25 to 30, subject to CAN arbitration;
- Node 2, you do not have much to say, so you can talk again from 30 to 35;
- Node 4, you will talk from 35 to 45;
- Node 2, you can talk again from 45 to 50;
- Node 5, you can talk from 50 to 55;
- From 55 to 60 it is free for all, everyone is welcome to talk - subject to CAN arbitration.
- So that you can synchronize your watches, I shall now tell you that it is 10.54 precisely.

TTCAN Mechanisms

- Each basic cycle starts with a reference message
- Given node 2 several chances to talk because, though it does not say much, it has to report information frequently
- The choice of time distribution is absolutely free and left to the discretion of the cycle manager
- For obvious reasons of synchronization and possible drift, the reference message must be sent periodically by the 'time master'

Comparison between TT and ET

- Event-triggered communication is generally good for aperiodic computations
  - E.g., interactions with the pilot generate aperiodic events, plane modes (air or ground)
- Time-triggered communication is good for periodic computations
  - E.g., fuel quantity updating, system data updating

Comparison between TT and ET

- Time-triggered communication enhances predictability by reducing latency jitters and provide higher dependability by making it easier to detect missed messages or illegal accesses to the bus
- However, event-triggered systems are more flexible to support configuration changes without a complete redesign and adapt faster to asynchronous events

Comparison between TT and ET

	ET	TT
Sporadic message	Yes	
Periodic Message		Yes
Flexibility	Yes	
Predictability		Yes

Outline

- Basic Overview
- Communication architectures
  - ➡ Event Triggered
  - ➡ Time Triggered
- Communication architecture examples
- Case Study: How Data Communication Affects Scheduling

Communication Architectures

- No one-size-fits-all architecture
  - ➡ Data throughput v.s. Latency guarantee
- Communication architectures may support different tiers of critical functionality
- Most systems can be divided into a hierarchy of functions from safety critical to mission critical and through a descending range to those rated as low criticality

MIL-STD-1553

- The aircraft internal time division command/response multiplex data bus, a military standard
- One of the first communication data bus standards
- First used in F-16 and Apache helicopter
- Define a dual-redundant communication bus that is used to interconnect nodes on a network
- In a dual redundant configuration, one bus is used to transmit data and the other is in hot backup status used only to send commands in the event of node failure

MIL-STD-1553

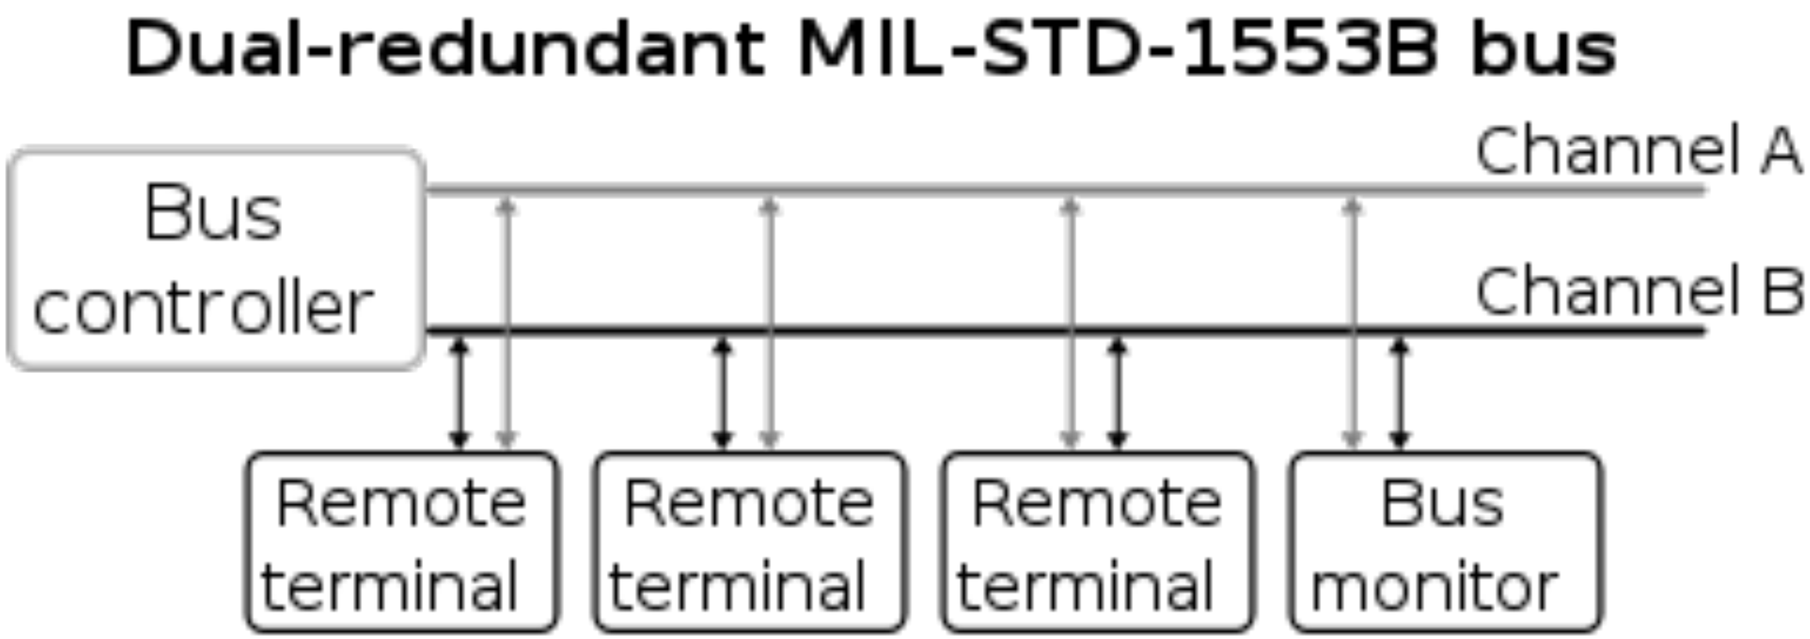

- A dual redundant bus - transmission media
- A bus controller - supervise the bus access
- A bus monitor - monitor and record bus transactions
- Multiple remote terminals - provide an interface between the data bus and an attached subsystem, or a bridge between a bus and another bus
    - e.g., interfaces that switch on the headlights and the landing lights in an aircraft

MIL-STD-1553

- The communication bandwidth is 10Mb/s
- Time division multiple access (TDMA) allows communication between the interconnected nodes

TDMA

- Time division multiple access, a channel access method for shared medium networks
  - ➡ share the same frequency channel by dividing the signal into different time slots
  - ➡ each user transmits using his own time slot
  - ➡ A type of time-division multiplexing

TDMA

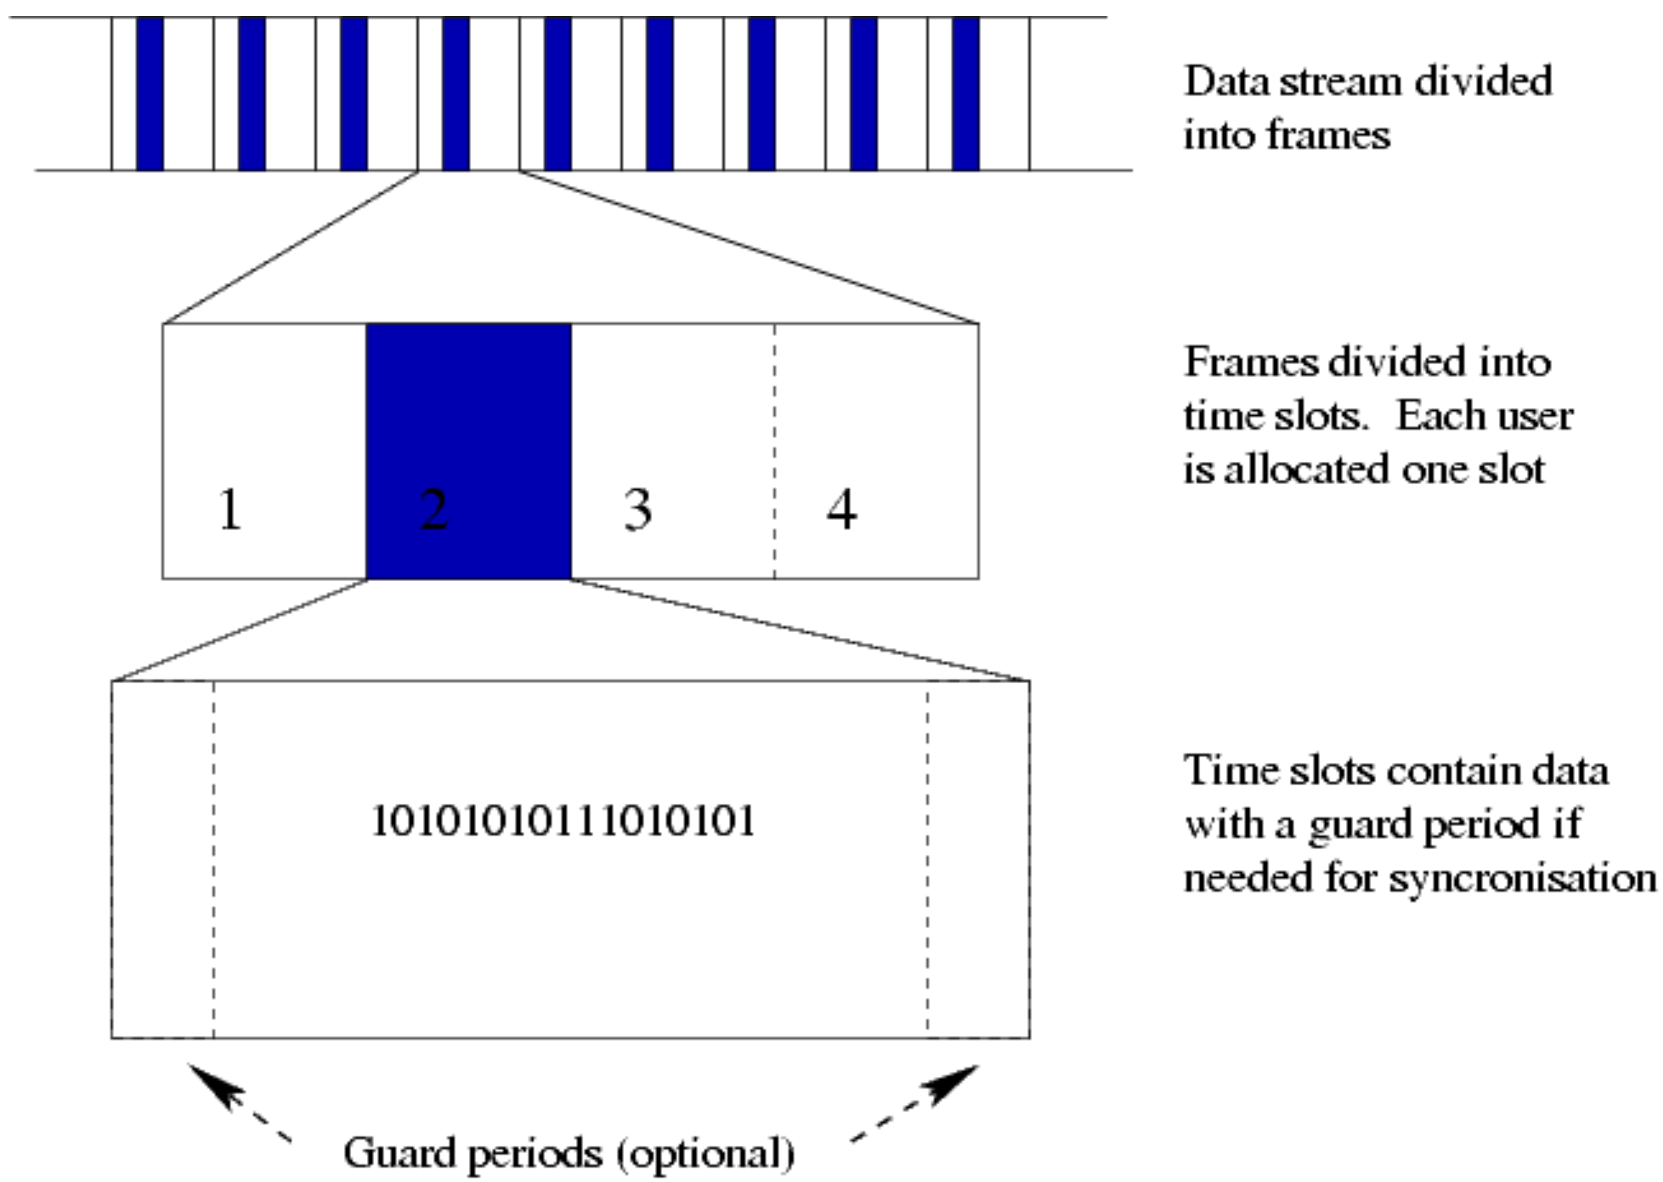

TDMA frame structure showing a data stream divided into frames and those frames divided into time slots

MIL-STD-1553

- Extremely reliable, widely used in military and space applications
- However, unable to transmit larger amounts of data at near real-time rates
- High cost due to the niche market targeted by suppliers

SAFEbus

- Honeywell implementation of ARINC 659 (an avionics data bus)
- Part of Boeing 777 avionics architecture

SAFEbus

- A quad redundant bus
  - ➡ two data lines and one clock line comprise each line
  - ➡ four data nodes and four IO nodes
  - ➡ Full duplication of bus interface units is provided at each node
  - ➡ Has a transmission rate of 60Mb/s

SAFEbus

- Data transmission is time-triggered and is governed by a message schedule
- Synchronized timing of messages delivered is maintained using a global clock
- Clocks are synchronized via periodic pulses on the dedicated clock line

SAFEbus

- Extremely high level of reliability and redundancy
- Good for safety critical functions in aircrafts
- Fault hypothesis - it is guaranteed to tolerate one arbitrary fault, but may tolerate multiple faults
- Very expensive - the hardware is redundant as a pair of pairs at all levels

Time-Triggered Communication Architecture

- Developed at the University of Vienna
  - ➡ using a time-triggered communication protocol called TTP/C
- Designed for hard real-time communication requirement
- Used in the Airbus A380 cabin pressure control system

Time-Triggered Communication Protocol

- A fault-tolerant TTP providing important services such as autonomous message transport based on a schedule with known delay and bounded jitter
- Current implementation uses a global clock to establish a time base, membership services to inform all nodes of the health status of the other nodes
- Masterless - allow communication to continue between the remaining nodes when other nodes fail

Time-Triggered Communication Protocol

- Designed to be physical layer independent
- Currently support communication at 5Mb/s over RS-485 (a flexible communication protocol), 25Mb/s over the ethernet physical layer, and 1Gb/s using gigabit ethernet
- Fault hypothesis guarantees that the communication system can tolerate any single fault in any architecture component, and multiple faults depending on the application

IEEE 1394b

- A communication architecture that has generated much interest in aerospace applications
    - ➡ used in F-22, F-35, and NASA's space shuttle
- Communication is specified over twisted, shielded and unshielded, pairs as well as plastic and glass optical fiber
- Support data rate from 100Mb/s up to 3.2Gb/s

IEEE 1394b

- The communication protocol used is characterized by an isochronous transmission phase and an asynchronous transmission phase
  - ➡ Isochronous transmission refers to broadcast transmissions to one node or many nodes on the network without error correction or retransmission
  - ➡ Asynchronous transfers are targeted to a specific address on the network and are acknowledged by the recipient, allowing error checking and the retransmission of message
- No timing guarantees

Space Wire

- Developed in Europe for use in satellites and spacecraft
- Used in NASA's Swift spacecraft

Space Wire

- The transmission physical layer is shielded twisted pair and point-to-point
- A large network of devices can be created using cascades of hubs or switches that route messages from one node to another
- Max data transfer rate is 400Mb/s
- Event-triggered
- No timing guarantees

Ethernet 10/100 Base-T

- One of the most widely used communications architectures
- Also, used on the international space station
- Provide data transmission rate of 10Mb/s and 100Mb/s over unshielded twisted pair
- Hubs connected by twisted pair facilitate "star topology"

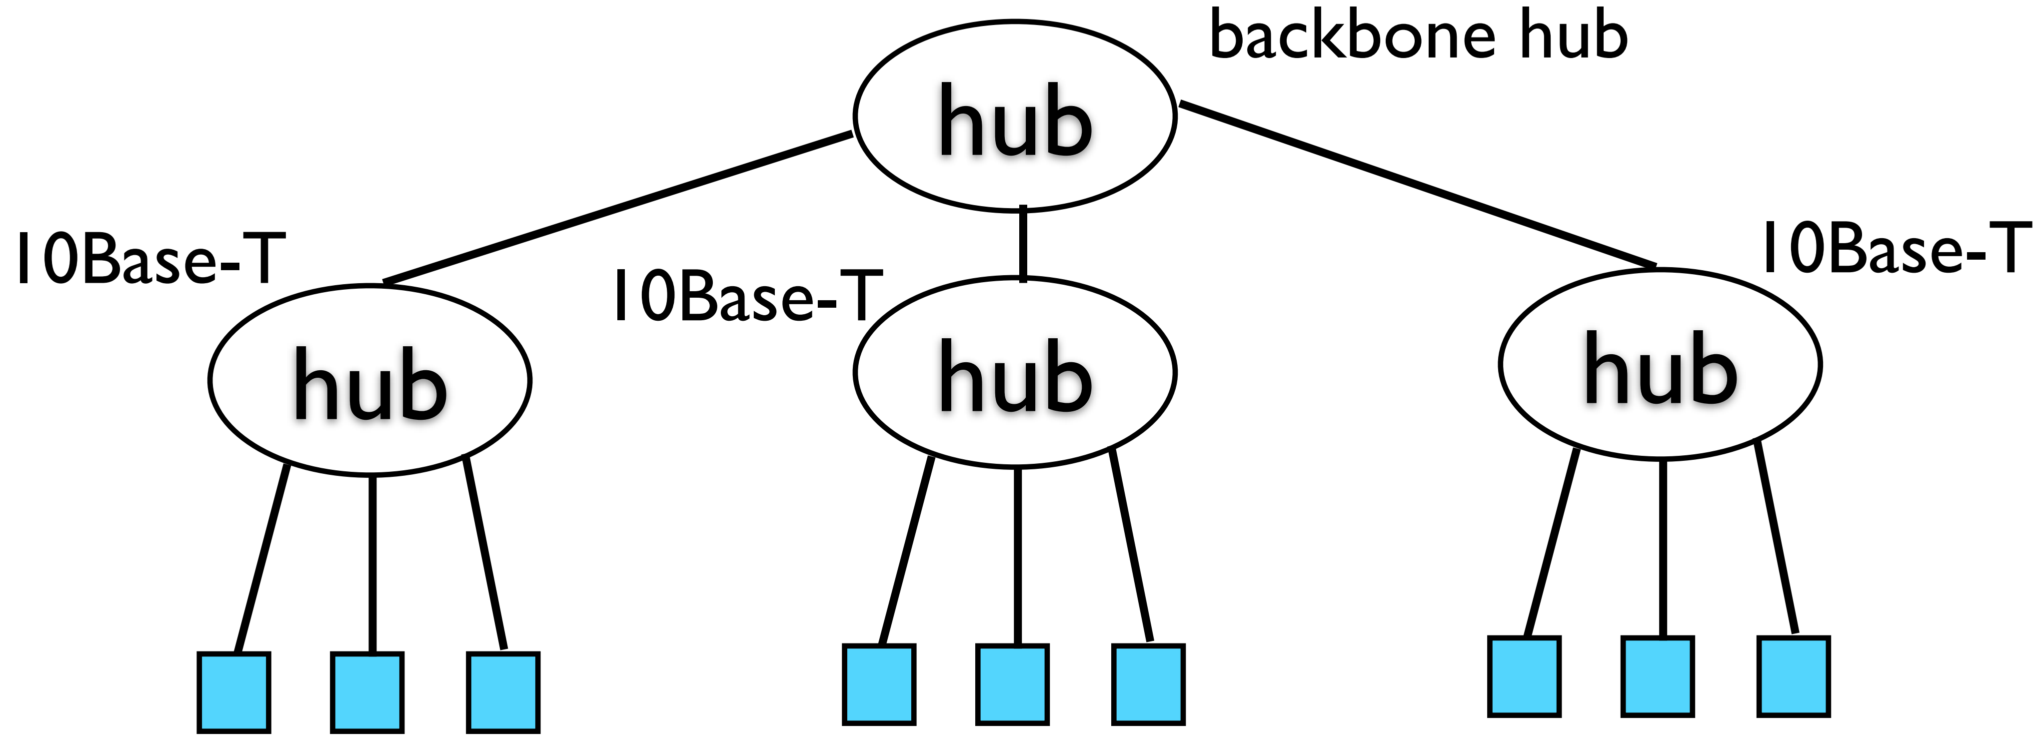

Ethernet 10/100 Base-T

- Can operate in half-duplex mode (i.e., communication in both directions, but only one direction at a time)
  - ➡ all nodes share the same cables
  - ➡ CSMA/CD governs the way computers share the channel
- Can also operate in full-duplex mode (i.e., allows communication in both directions to happen simultaneously)
  - ➡ nodes can communicate over dedicated cabling with one other device

CSMA/CD

- Carrier sense multiple access with collision detection
- Data is transmitted in the form of packets
- Sense channel prior to actual packet transmission
- Transmit packet only if channel is sensed idle; else, defer the transmission until channel becomes idle
- After packet transmission is started, the node monitors its own transmission to see if the packet has experienced a collision
    - ➡ Collisions can be spotted since they are generally higher in signal amplitude than normal signals
- If the packet is observed to be undergoing a collision, the transmission is aborted and the packet is retransmitted after a random interval of time

Ethernet 10/100 Base-T

- The protocol used to send messages affects the reliability of the transmission
- User datagram protocol (UDP)
  - ➡ Unreliable connectionless protocol with no guarantee that the data will reach its destination
  - ➡ Provide barebones service with very little overhead
- Transmission control protocol (TCP)
  - Add significant overhead compared to UDP
  - Provide a reliable connection

Avionics Full-Duplex Switched Ethernet (AFDX)

- A trademark of Airbus
  - ➡ Used in the A380 plane
- A standard that defines the electrical and protocol specifications for the exchange of data between avionic subsystems using Ethernet 10/100 Base-T for the communication architecture
- Derived from commercial databus standards (Ethernet 10/100 Base-T, Ethernet medium access control addressing, IP, and UPD)
  - ➡ Add deterministic communication and redundancy management
  - ➡ Provide secure and reliable communications of critical and noncritical data

Avionics Full-Duplex Switched Ethernet (AFDX)

- Deterministic communications
  - ➡ Achieved by defining communication virtual links between end systems with specified bandwidth, bounded latency, and frame size
  - ➡ Messages are sent with a sequence number that is used on the receiving end to verify that the sequence numbers are in order - integrity checking

Fibre Channel

- A high-performance data transport connection technology supporting transmission via copper wires or fiber optic cables over long distances
  - Originally developed for storage applications, also selected for use in military aircraft avionics
    - ➡ Notably used in F/A-18

Fibre Channel

- A full-duplex communication architecture that supports a variety of topologies such as point-to-point, arbitrated loop, and switched fabric

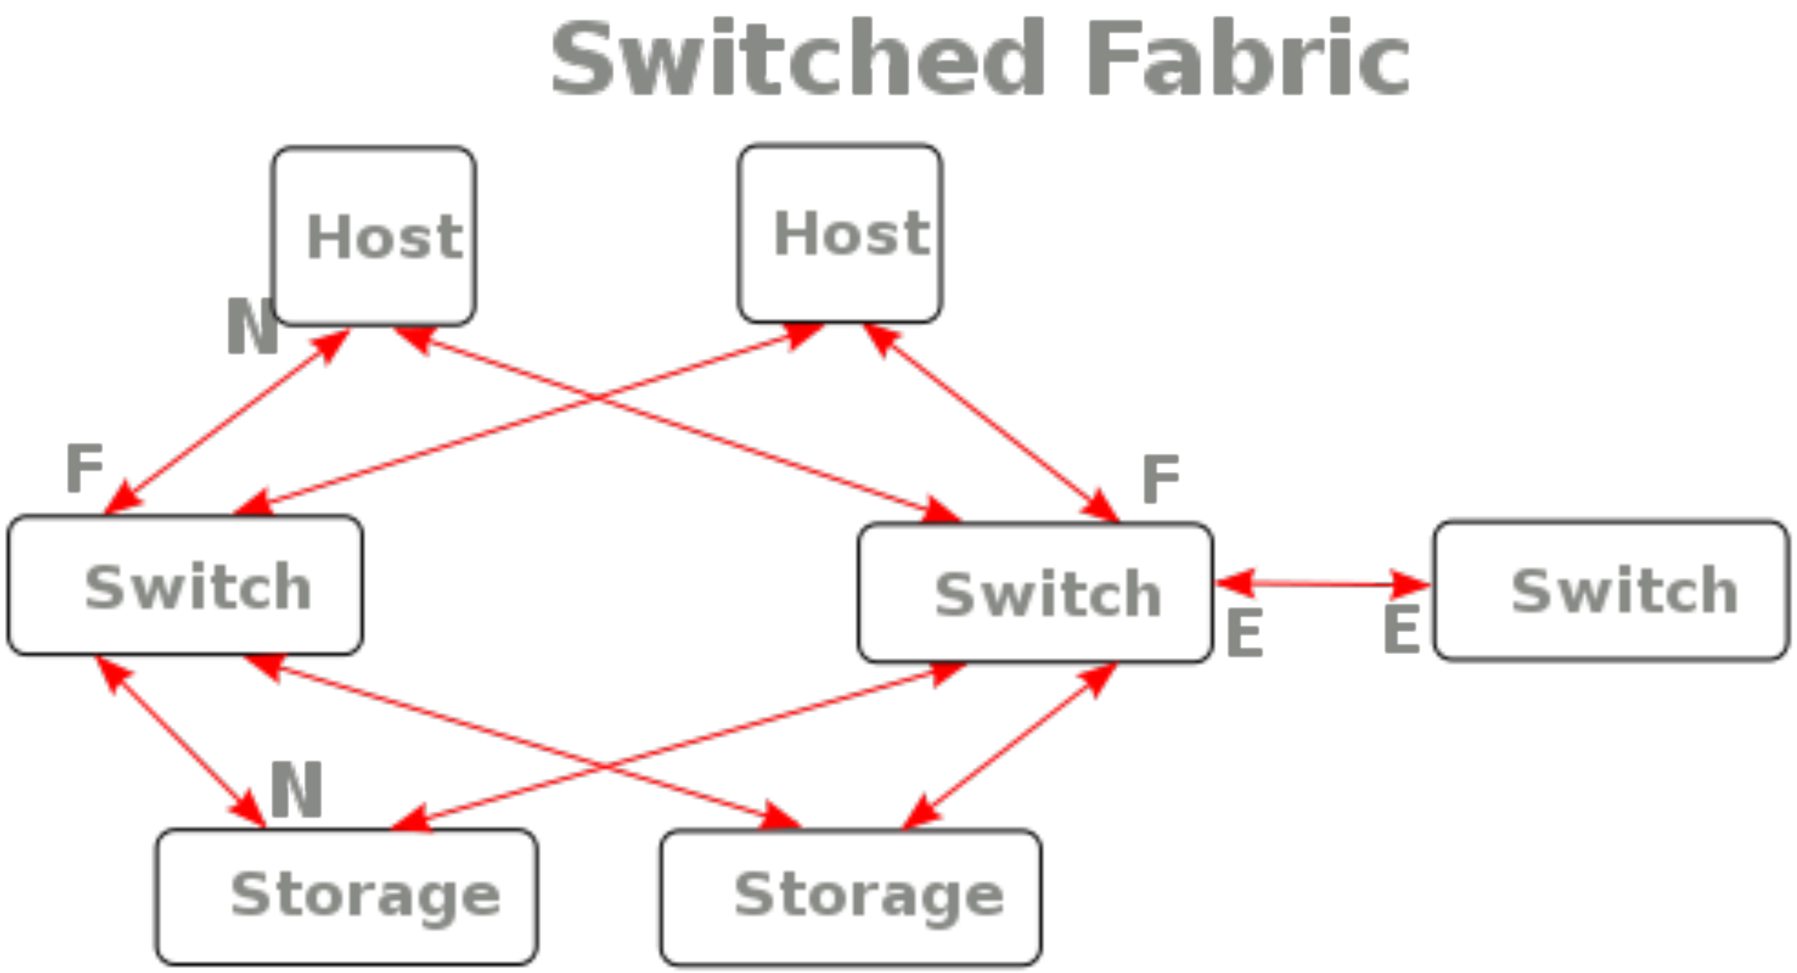

Fibre Channel

- Support several classes of transmission
  - ➡ Class 1 - provide a dedicated connection with acknowledgement, guaranteeing delivery and message sequence
  - ➡ Class 2 - connectionless and may provide message out of order, delivery confirmation is provided
  - ➡ Class 3 - connectionless and unconfirmed
- Extremely fast, but not deterministic

Gigabit Ethernet

- A combination of the Ethernet 10/100 and the Fibre Channel physical layer standards
- Support both copper wire and fiber optic transmission media
- Very fast transmission rate over long distances
- Other than the differences in the physical layer, it operates the same as Ethernet 10/100 Base-T

Communication Architecture Selection Rationale - A Case Study

- Select the best-fit architecture to support a hard real-time distributed control system for safety critical systems in a spacecraft

TTP/C wins!

- TTP/C is designed specifically for safety critical, hard real-time distributed control
    - ➡ Provide the guaranteed latency
    - ➡ The use of a predefined message schedule with a fault-tolerant global clock provides known and exactly predictable communication bus loading and message sequencing
    - ➡ Masterless - failure of some nodes does not prevent synchronized communication from continuing between the remaining nodes
    - ➡ Physical layer independent - higher speed transmission can be obtained by moving to an appropriate physical layer
    - ➡ Cost-effective

Outline

- Basic Overview
- Communication architectures
    - ➡ Event Triggered
    - ➡ Time Triggered
- Communication architecture examples
- Case Study: How Data Communication Affects Scheduling

ARINC 659 Scheduling Constraints brought by Data Communication

- ARINC 659 is an international standard which specifies the communication protocol between the different systems on board
- The ARINC 659 protocol is a communication protocol with a rate of 100 Kbits per second
- An ARINC 659 bus is a set of channels, each carrying data packets
- The ARINC 659 interface system is in charge of acquisition of the data packets received on several parallel input channels
- SAFEbus is a Honeywell implementation of the ARINC 659

ARINC 659 Scheduling Constraints brought by Data Communication

- Boeing 777 integrated airplane information management system (AIMS) is built based upon ARINC 659
  - ➡ high resource utilization
  - ➡ strict partitioning and provable performance via pre-scheduling of processing and communication resources
  - ➡ But, has many theoretical and practical problems brought by data communications

Dataflow

- Communication requirements are defined in terms of input and output messages
  - ➡ Come from or go to either other processes, external communication ports, or both
- 5900 unique data items in the system
- Each data item is transferred across the SAFEbus at its respective rate
  - ➡ Resulting in 17000 total messages per frame
- The scheduling protocol must be efficient enough to schedule such huge dataflows

Data Consistency

- Scheduling must maintain data consistency
- A process cannot read data from nor write data to a buffer while it is being transferred across the SAFEbus
  - ➡ Data transfer must follow the expiration of the deadline of the producing process
  - ➡ Must be transferred before the release of a consuming process

Data Size

- Each message has a specified length corresponding to the number of words to transfer
- The schedule must provide appropriate time to transfer a message of the size across the SAFEbus

Data Rate

- Data is scheduled to be transferred at the minimum necessary rate to meet uniqueness and latency requirements
- If a high rate consumer specifies a tight latency on external input, then the data needs to be scheduled for transfer in a rate sufficient to meet the latency requirements

Master / Shadow

- Data may be specified Master / Shadow, meaning that there exist multiple processes which can generate that data
  - ➡ These processes must be scheduled so they simultaneously satisfy all the relevant data transfer constraints of sharing the same time on the bus

Synchronization

- Sufficient synchronization messages will be allocated on the SAFEbus to maintain appropriate timing and recovery characteristics
- The number of synchronization messages should be minimized in order to maintain time accuracy and minimize consumption of bus bandwidth

Sources

- http://en.wikipedia.org/wiki/MIL-STD-1553
- http://www.ttagroup.org/ttp/articles.htm
- http://en.wikipedia.org/wiki/Avionics#Communications
- http://en.wikipedia.org/wiki/Time_division_multiple_access
- http://en.wikipedia.org/wiki/Controller_area_network
- http://en.wikipedia.org/wiki/Avionics_Full-Duplex_Switched_Ethernet
- D.A. Gwaltney and J. M. Briscoe, Comparison of Communication Architectures for Spacecraft Modular Avionics Systems
- T. Carpenter, K. Driscoll, K. Hoyme and J. Carciofini, ARINC 659 Scheduling: Problem Definition
- http://www.can-cia.org/index.php?id=166
- http://en.wikipedia.org/wiki/Time-Triggered_Protocol
- F. Cloute etc., Hardware/software co-design of an avionics communication protocol interface system: an industrial case study
- http://en.wikipedia.org/wiki/Carrier_Sense_Multiple_Access_With_Collision_Detection
- H. Kopetz, A Comparison of Can and TTP
- http://www.google.com/imgres?imgurl=http://www.idea-consulting.eu/images/Avionik.jpg&imgrefurl=http://www.idea-consulting.eu/de/Service/