ETHERNOVIA

TRANSFORMING HOW CARS OF THE FUTURE ARE BUILT

IEEE802.1DG – TRAFFIC CATEGORIES | 2020-09-24

IEEE contribution

ETHERNOVIA
TRANSFORMING HOW CARS OF THE FUTURE ARE BUILT

IEC/IEEE60802(D1.2): 4.6 Industrial Traffic Types

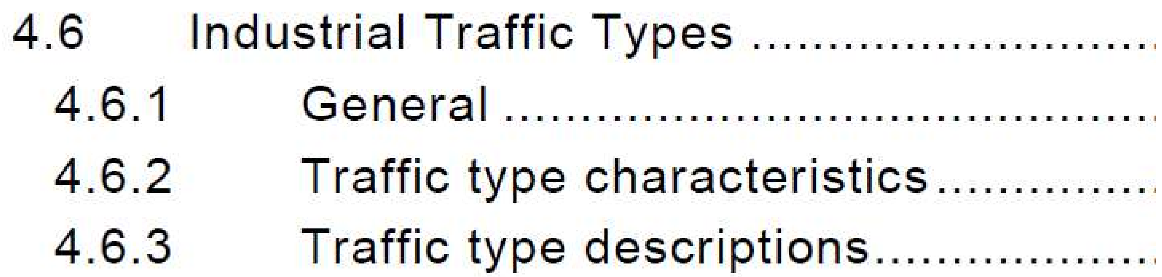

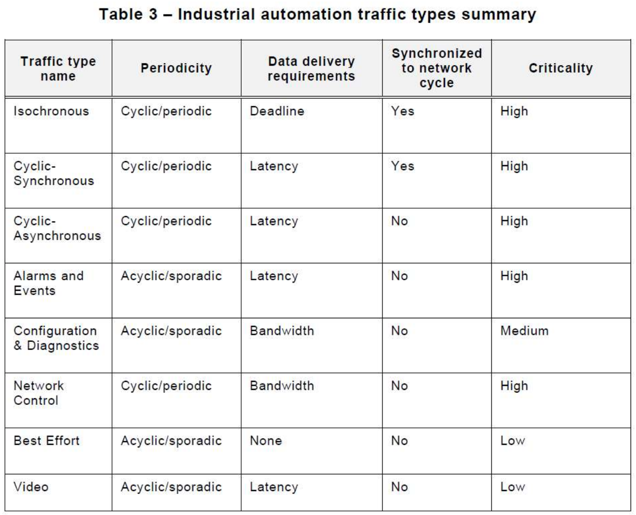

Table 3 – Industrial automation traffic types summary

Traffic type name	Periodicity	Data delivery requirements	Synchronized to network cycle	Criticality
Isochronous	Cyclic/periodic	Deadline	Yes	High
Cyclic-Synchronous	Cyclic/periodic	Latency	Yes	High
Cyclic-Asynchronous	Cyclic/periodic	Latency	No	High
Alarms and Events	Acyclic/sporadic	Latency	No	High
Configuration & Diagnostics	Acyclic/sporadic	Bandwidth	No	Medium
Network Control	Cyclic/periodic	Bandwidth	No	High
Best Effort	Acyclic/sporadic	None	No	Low
Video	Acyclic/sporadic	Latency	No	Low

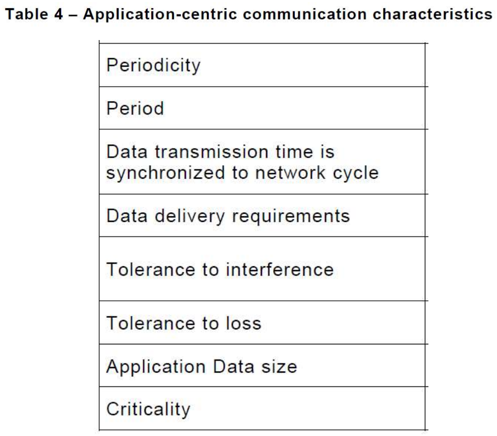

Table 4 – Application-centric communication characteristics

Periodicity
Period
Data transmission time is synchronized to network cycle
Data delivery requirements
Tolerance to interference
Tolerance to loss
Application Data size
Criticality

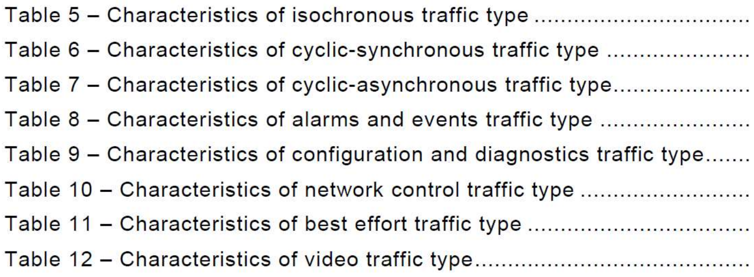

Automotive Traffic/Application Categories and Types

A personal view, based on experience, not deemed normative nor complete (yet)!

Audio	Video	Control	Sensor	Bulk
Entertainment Audio	Entertainment Video	Alarms	Small	Web-Download
Intercom	Human Assist (Parking Aide)	Control	List-Type	SOTA
Transient Noise Cancellation	Mirror Replacement	Events		OBD Flash-Update
Warning Chimes	Machine Vision			Off-Board "Sensor"
Emergency-Vehicle Detection				

Compare IEC/IEEE60802(D1.2) 4.6.1 General – Table 3.

IEC/IEEE60802(D1.2): 4.6.2 Traffic type characteristics

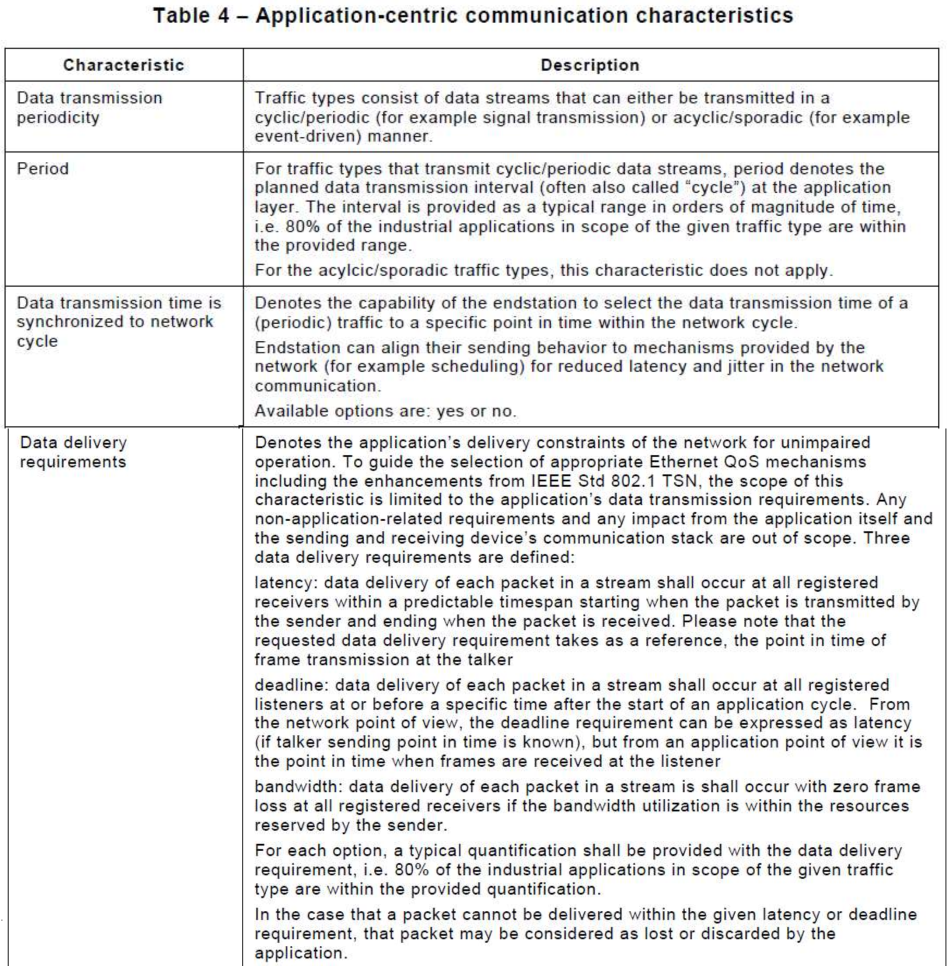

Table 4 – Application-centric communication characteristics

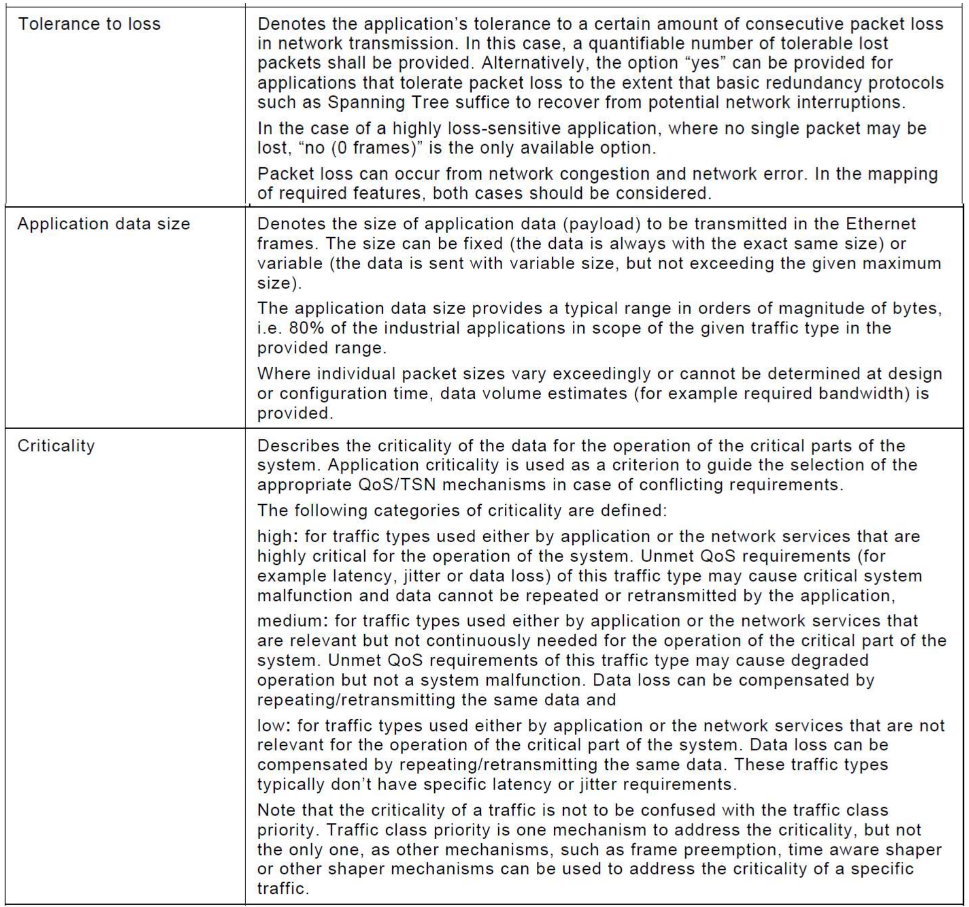

Characteristic	Description
Data transmission periodicity	Traffic types consist of data streams that can either be transmitted in a cyclic/periodic (for example signal transmission) or acyclic/sporadic (for example event-driven) manner.
Period	For traffic types that transmit cyclic/periodic data streams, period denotes the planned data transmission interval (often also called "cycle") at the application layer. The interval is provided as a typical range in orders of magnitude of time, i.e. 80% of the industrial applications in scope of the given traffic type are within the provided range. For the acylcic/sporadic traffic types, this characteristic does not apply.
Data transmission time is synchronized to network cycle	Denotes the capability of the endstation to select the data transmission time of a (periodic) traffic to a specific point in time within the network cycle. Endstation can align their sending behavior to mechanisms provided by the network (for example scheduling) for reduced latency and jitter in the network communication. Available options are: yes or no.
Data delivery requirements	Denotes the application's delivery constraints of the network for unimpaired operation. To guide the selection of appropriate Ethernet QoS mechanisms including the enhancements from IEEE Std 802.1 TSN, the scope of this characteristic is limited to the application's data transmission requirements. Any non-application-related requirements and any impact from the application itself and the sending and receiving device's communication stack are out of scope. Three data delivery requirements are defined: latency: data delivery of each packet in a stream shall occur at all registered receivers within a predictable timespan starting when the packet is transmitted by the sender and ending when the packet is received. Please note that the requested data delivery requirement takes as a reference, the point in time of frame transmission at the talker deadline: data delivery of each packet in a stream shall occur at all registered listeners at or before a specific time after the start of an application cycle. From the network point of view, the deadline requirement can be expressed as latency (if talker sending point in time is known), but from an application point of view it is the point in time when frames are received at the listener bandwidth: data delivery of each packet in a stream is shall occur with zero frame loss at all registered receivers if the bandwidth utilization is within the resources reserved by the sender. For each option, a typical quantification shall be provided with the data delivery requirement, i.e. 80% of the industrial applications in scope of the given traffic type are within the provided quantification. In the case that a packet cannot be delivered within the given latency or deadline requirement, that packet may be considered as lost or discarded by the application.
Tolerance to loss	Denotes the application's tolerance to a certain amount of consecutive packet loss in network transmission. In this case, a quantifiable number of tolerable lost packets shall be provided. Alternatively, the option "yes" can be provided for applications that tolerate packet loss to the extent that basic redundancy protocols such as Spanning Tree suffice to recover from potential network interruptions. In the case of a highly loss-sensitive application, where no single packet may be lost, "no (0 frames)" is the only available option. Packet loss can occur from network congestion and network error. In the mapping of required features, both cases should be considered.
Application data size	Denotes the size of application data (payload) to be transmitted in the Ethernet frames. The size can be fixed (the data is always with the exact same size) or variable (the data is sent with variable size, but not exceeding the given maximum size). The application data size provides a typical range in orders of magnitude of bytes, i.e. 80% of the industrial applications in scope of the given traffic type in the provided range. Where individual packet sizes vary exceedingly or cannot be determined at design or configuration time, data volume estimates (for example required bandwidth) is provided.
Criticality	Describes the criticality of the data for the operation of the critical parts of the system. Application criticality is used as a criterion to guide the selection of the appropriate QoS/TSN mechanisms in case of conflicting requirements. The following categories of criticality are defined: high: for traffic types used either by application or the network services that are highly critical for the operation of the system. Unmet QoS requirements (for example latency, jitter or data loss) of this traffic type may cause critical system malfunction and data cannot be repeated or retransmitted by the application, medium: for traffic types used either by application or the network services that are relevant but not continuously needed for the operation of the critical part of the system. Unmet QoS requirements of this traffic type may cause degraded operation but not a system malfunction. Data loss can be compensated by repeating/retransmitting the same data and low: for traffic types used either by application or the network services that are not relevant for the operation of the critical part of the system. Data loss can be compensated by repeating/retransmitting the same data. These traffic types typically don't have specific latency or jitter requirements. Note that the criticality of a traffic is not to be confused with the traffic class priority. Traffic class priority is one mechanism to address the criticality, but not the only one, as other mechanisms, such as frame preemption, time aware shaper or other shaper mechanisms can be used to address the criticality of a specific traffic.

Network-Cycle vs. Application-Cycle

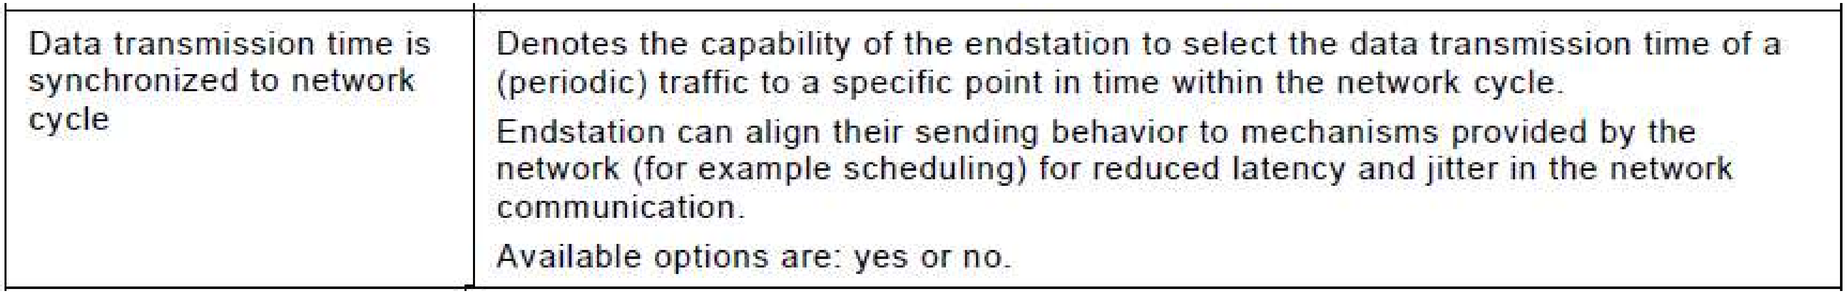

Data transmission time is synchronized to network cycle	Denotes the capability of the endstation to select the data transmission time of a (periodic) traffic to a specific point in time within the network cycle. Endstation can align their sending behavior to mechanisms provided by the network (for example scheduling) for reduced latency and jitter in the network communication. Available options are: yes or no.

IEC/IEEE60802(D1.2):
4.6.2 Traffic type characteristics
Table 4

Lessons Learned from FlexRay

- Software Cycle is difficult to synchronize to Network Cycle
  - (Fast) Parallel start-up without dependencies (clock, network, …)
  - Other application dependencies (sample rate, scan rate, revolutions, …)
  - Wall-Clock and "Media-Clock"/"Working Clock" required
  - Avoid re-sync after start-up for continuous operation
- "Bounded Latency" means maximum latency
  - TX-data delivered from SW to bus just after FR-cycle
  - RX-data delivered from bus to SW just after SW-cycle
- "Active-Star" drives Clock Requirements
- Software Cycles become Topology dependent
  - Delays depend on number of "Active-Stars"

A personal view, based on experience!

Discussion of Data Loss/Error

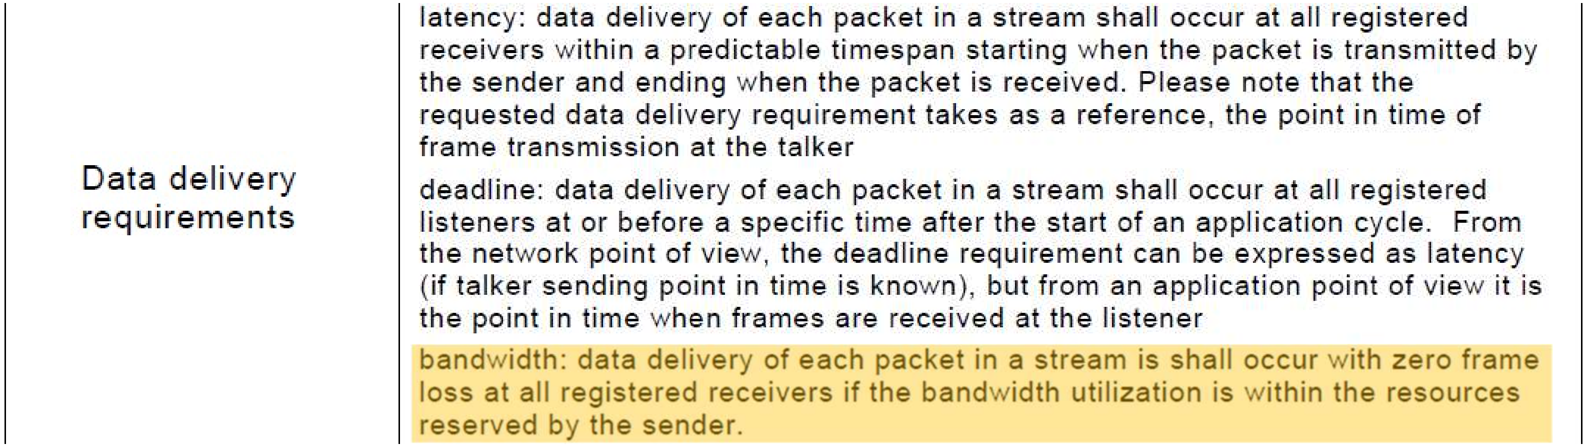

Data delivery requirements	latency: data delivery of each packet in a stream shall occur at all registered receivers within a predictable timespan starting when the packet is transmitted by the sender and ending when the packet is received. Please note that the requested data delivery requirement takes as a reference, the point in time of frame transmission at the talker deadline: data delivery of each packet in a stream shall occur at all registered listeners at or before a specific time after the start of an application cycle. From the network point of view, the deadline requirement can be expressed as latency (if talker sending point in time is known), but from an application point of view it is the point in time when frames are received at the listener bandwidth: data delivery of each packet in a stream is shall occur with zero frame loss at all registered receivers if the bandwidth utilization is within the resources reserved by the sender.

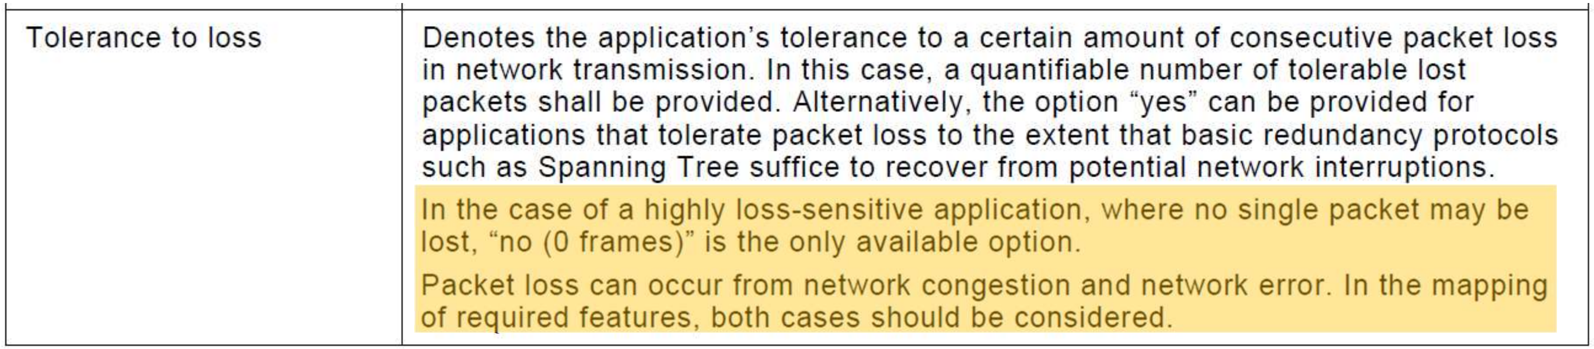

Tolerance to loss	Denotes the application's tolerance to a certain amount of consecutive packet loss in network transmission. In this case, a quantifiable number of tolerable lost packets shall be provided. Alternatively, the option "yes" can be provided for applications that tolerate packet loss to the extent that basic redundancy protocols such as Spanning Tree suffice to recover from potential network interruptions. In the case of a highly loss-sensitive application, where no single packet may be lost, "no (0 frames)" is the only available option. Packet loss can occur from network congestion and network error. In the mapping of required features, both cases should be considered.

IEC/IEEE60802(D1.2):
4.6.2 Traffic type characteristics
Table 4

Data Loss model (temporary loss only!)

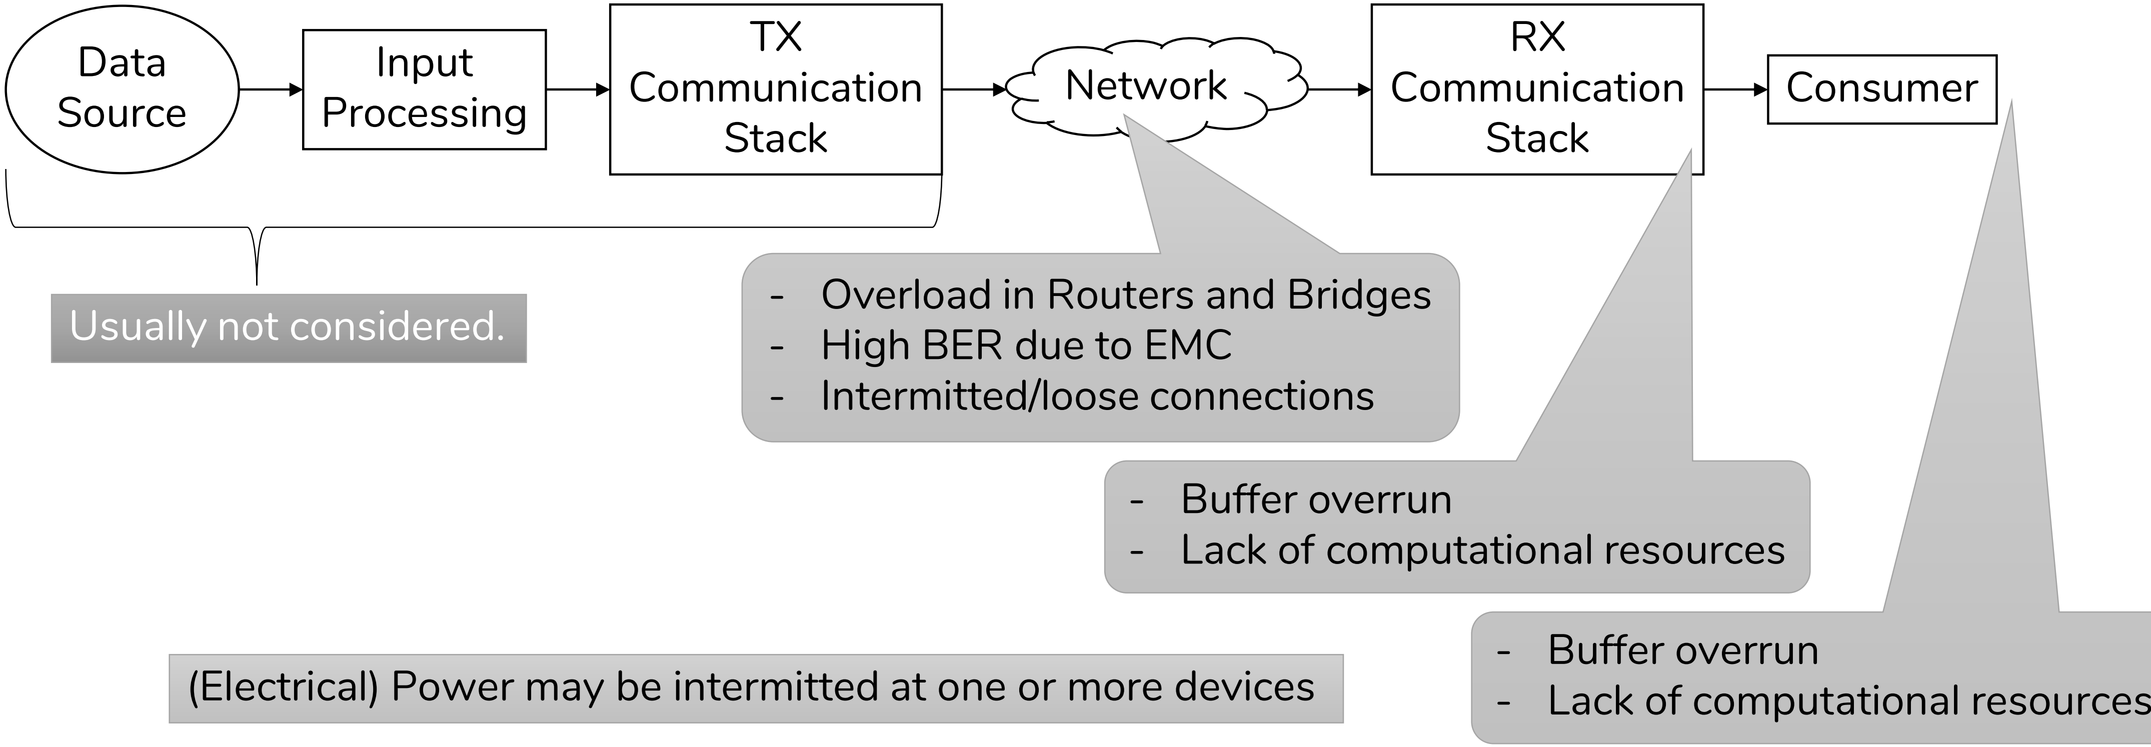

Data Source

Input Processing

TX Communication Stack

Network

RX Communication Stack

Consumer

Usually not considered.

- Overload in Routers and Bridges
- High BER due to EMC
- Intermitted/loose connections

- Buffer overrun
- Lack of computational resources

- Buffer overrun
- Lack of computational resources

(Electrical) Power may be intermitted at one or more devices

A personal view, based on experience!

Default answer: Re-Transmit if Lost

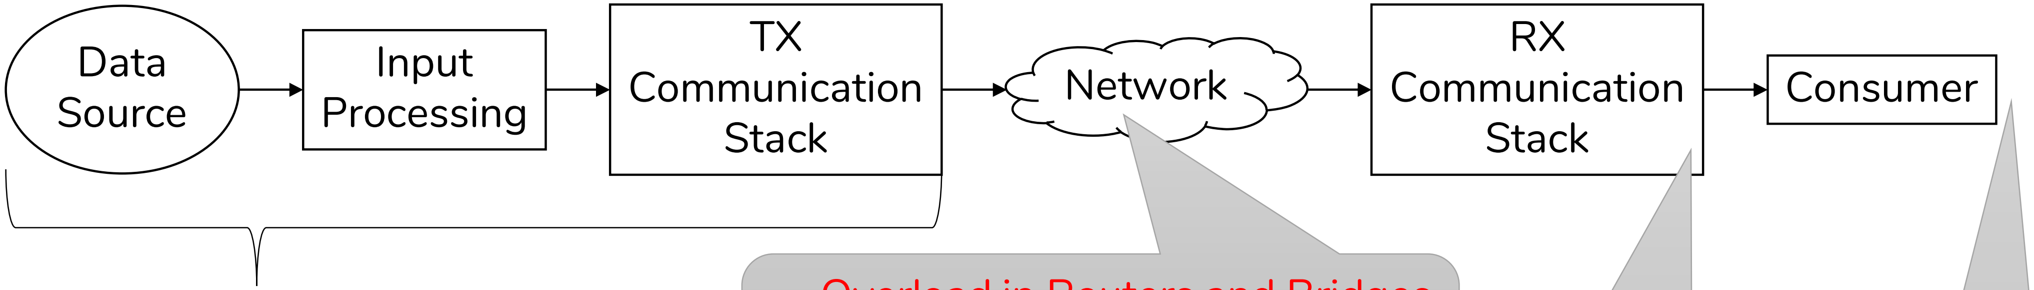

Usually not considered.

- Overload in Routers and Bridges
- High BER due to EMC
- Intermitted/loose connections

- Buffer overrun
- Lack of computational resources

- Buffer overrun
- Lack of computational resources

A personal view, based on experience!

Fast/Frequent Re-Transmits only make these worse!

Message Delivery Policies – Part 1

- Ensure in order delivery
  - Classical TCP use-case.
  - It is important that all data be delivered to the Consumer in the correct (byte-)order.
  - How long the total transmission takes is secondary for the functionality.
  - TCP congestion control algorithm variety shows dependency on loss model.
  - Receiving application may still not consume the data.
- Request/Response
  - Re-transmission only if Response is not received within application timeout.
  - Similar to TCP, but timeout implemented at application level for short (single frame) messages.
  - Receiving application has consumed the data if Response is sent.
- Fire and Forget
  - Send transient event data only once.
  - May be used for user observed functions:
    - Volume control
    - CD eject
    - Reading lamp
  - User will initiate re-try if function does not execute.
  - Impaired quality perception if loss is frequent.

Message Delivery Policies – Part 2

- Last is best
  - Classical automotive cyclic-transmission use-case.
  - Updated information is transmitted at a certain period that is often shorter or at least similar to a potential need for retransmission detection (more BW, easier test).
  - The period is chosen to be short enough to tolerate the loss of at least a single slot.
  - For safety critical information the alive-counter may be more important than the application data.
- Safeguard single timely delivery
  - Delivery of transient event must be assured within minimum time.
  - Different solutions depending on loss-model:
    - Send message burst (often 3x) – if BER is assumed main reason for loss.
    - Repeat at regular period (last is best) – if temporal overload (buffer or computation) is assumed main reason for loss.
    - Request/Response, where Request is repeated quickly if no Response is received (similar to message burst) – assumed loss model?
    - Redundant data path?

Fail-Safe vs. Fail-Operational

A System relies on an Input for executing its Mission.

- Fail-Safe
  - After an Initial Error to the Input, the System fails, but assumes some Final Safe State, that will not cause further harm, but it can no longer perform its Mission.
  - A Secondary Error is not considered.
- Fail-Operational
  - After an Initial Error to the Input, the System has some Alternate Input enabling it to continue its Mission for a Limited Time.
  - After some Time or Secondary Error the System may
    - fail or
    - go into a Final Safe State.
  - A Ternary Error is (usually) not considered.

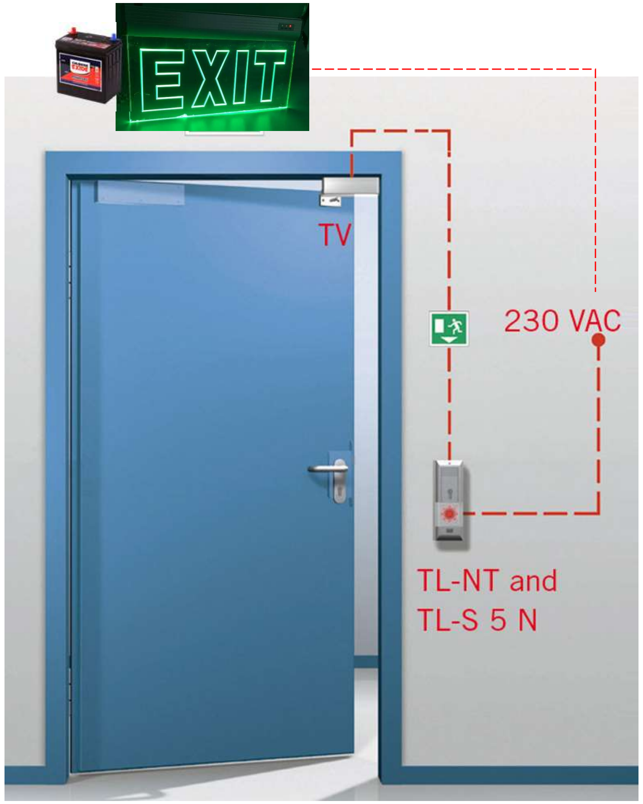

A personal view, based on experience, not deemed normative nor complete (yet)!

Possible Design Goals for Redundancy

During operation:

- No Redundancy: Fail safe – after loss immediate transition to a local safe state
- Extended wearout: Ignore initial failure, second failure will loose system functionality (likely only for quality driven data)
- Fail gracefully: Redundant data after initial failure used to mitigate transition to a system safe state (choose time interval to limit chance of second failure)
- Lip home: Continue mission for extended period, maybe with reduced performance, but no reduction of safety level (likely requires tripple redundancy)

Availability: Does a failure (or false positive) at start-up prevent start of mission?

Also consider Error in Data!

Also consider Loss of Data Source!

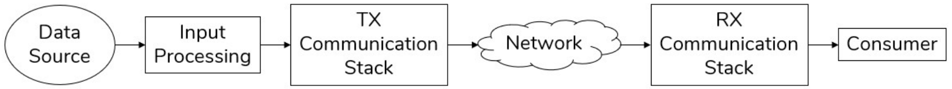

Categorizing Audio traffic

- Entertainment Audio
  - 1.5kByte frame(s) (multi channel)
  - Period: #us-1ms
  - Latency: #ms (AVnu-CBS 1.33ms) – lip-sync!
  - Policing: drop
  - Loss: quality issue
  - Preemtion: may be preempted
- Intercom
  - 100Byte frames (single channel)
  - Period: #us
  - Latency: 100us (AVnu-CBS 125us)
  - Policing: drop
  - Loss: quality issue
  - Preemtion: may be preempted
- Transient Noise Cancellation
  - 100Byte frames (single channel)
  - Period: #us
  - Latency: 50us
  - Policing: drop
  - Loss: quality issue
  - Preemtion: may preempt other flows
- Warning Chimes
  - 100Byte frames (single channel)
  - Period: #us
  - Latency: #ms (AVnu-CBS 1.33ms)
  - Policing: no
  - Loss: warning, visual alternate (redundancy!)
  - Preemtion: ??
- Emergency Vehicle Detection
  - 100Byte frames (single channel)
  - Period: #us
  - Latency: #ms (AVnu-CBS 1.33ms)
  - Policing: no
  - Loss: warning, abort/not-start automation mission
  - Preemtion: ??

Categorizing Video traffic

- Entertainment Video
  - 1.5kByte frames
  - Period: #us
  - Latency: #ms (AVnu-CBS 1.33ms) – lip-sync!
  - Policing: drop
  - Loss: quality issue
  - Preemtion: may be preempted
- Human Assist
  - 1.5kByte frames
  - Period: 10-30ms
  - Latency: #ms (AVnu-CBS 1.33ms)
  - Policing: drop
  - Loss: black screen – do not display old or false picture
  - Preemtion: may be preempted
- Mirror Replacement
  - 1.5kByte frames
  - Period: 10-30ms
  - Latency: #ms??
  - Policing: no
  - Loss: warning, black screen – do not display old or false picture
  - Preemtion: ??
- Machine Vision
  - 1.5kByte frames
  - Period: 10-30ms
  - Latency: #ms??
  - Policing: no
  - Loss: abort/not-start automation mission
  - Preemtion: ??

Categorizing Control traffic

- Alarms
  - 64Byte frame
  - Period: 5-50ms
  - Latency: lowest possible
  - Policing: no
  - Loss: warning, abort/not-start automation mission
  - Preemtion: may preempt other flows
- Control
  - 64Byte frame
  - Period: 5-50ms
  - Latency: #ms??
  - Policing: no
  - Loss: warning, abort/not-start automation mission
  - Preemtion: may preempt other flows
- Events
  - 64Byte frame
  - non periodic
  - Latency: #ms??
  - Policing: drop
  - Loss: quality issue
  - Preemtion: neither

Categorizing Sensor traffic

- Small Sensor
  - 64Byte Frames
  - Period: 5-50ms
  - Latency: #ms??
  - Policing: no
  - Loss: warning, abort/not-start automation mission
  - Preemtion: ??

- List-type Sensor
  - 1.5kByte Frames
  - Period: 30-100ms
  - Latency: #ms??
  - Policing: drop
  - Loss: warning, abort/not-start automation mission
  - Preemtion: ??

Categorizing Bulk traffic

- Web-Download
  - 1.5kByte Frames
  - non periodic
  - ⇨ Latency: sec.
  - Policing: drop
  - Loss: quality issue
  - Preemtion: may be preempted
- SOTA
  - 1.5kByte Frames
  - non periodic
  - ⇨ Latency: sec./min.
  - Policing: drop
  - Loss: quality issue
  - Preemtion: may be preempted
- OBD Flash-Update
  - 1.5kByte Frames
  - non periodic
  - ⇨ Latency: #ms?? (special vehicle state?)
  - Policing: no
  - Loss: quality issue
  - Preemtion: n/a
- Off-Board “Sensor”
  - 1.5kByte Frames
  - non periodic
  - Latency: sec.
  - Policing: drop
  - ⇨ Loss: abort/not-start automation mission?
  - Preemtion: may be preempted

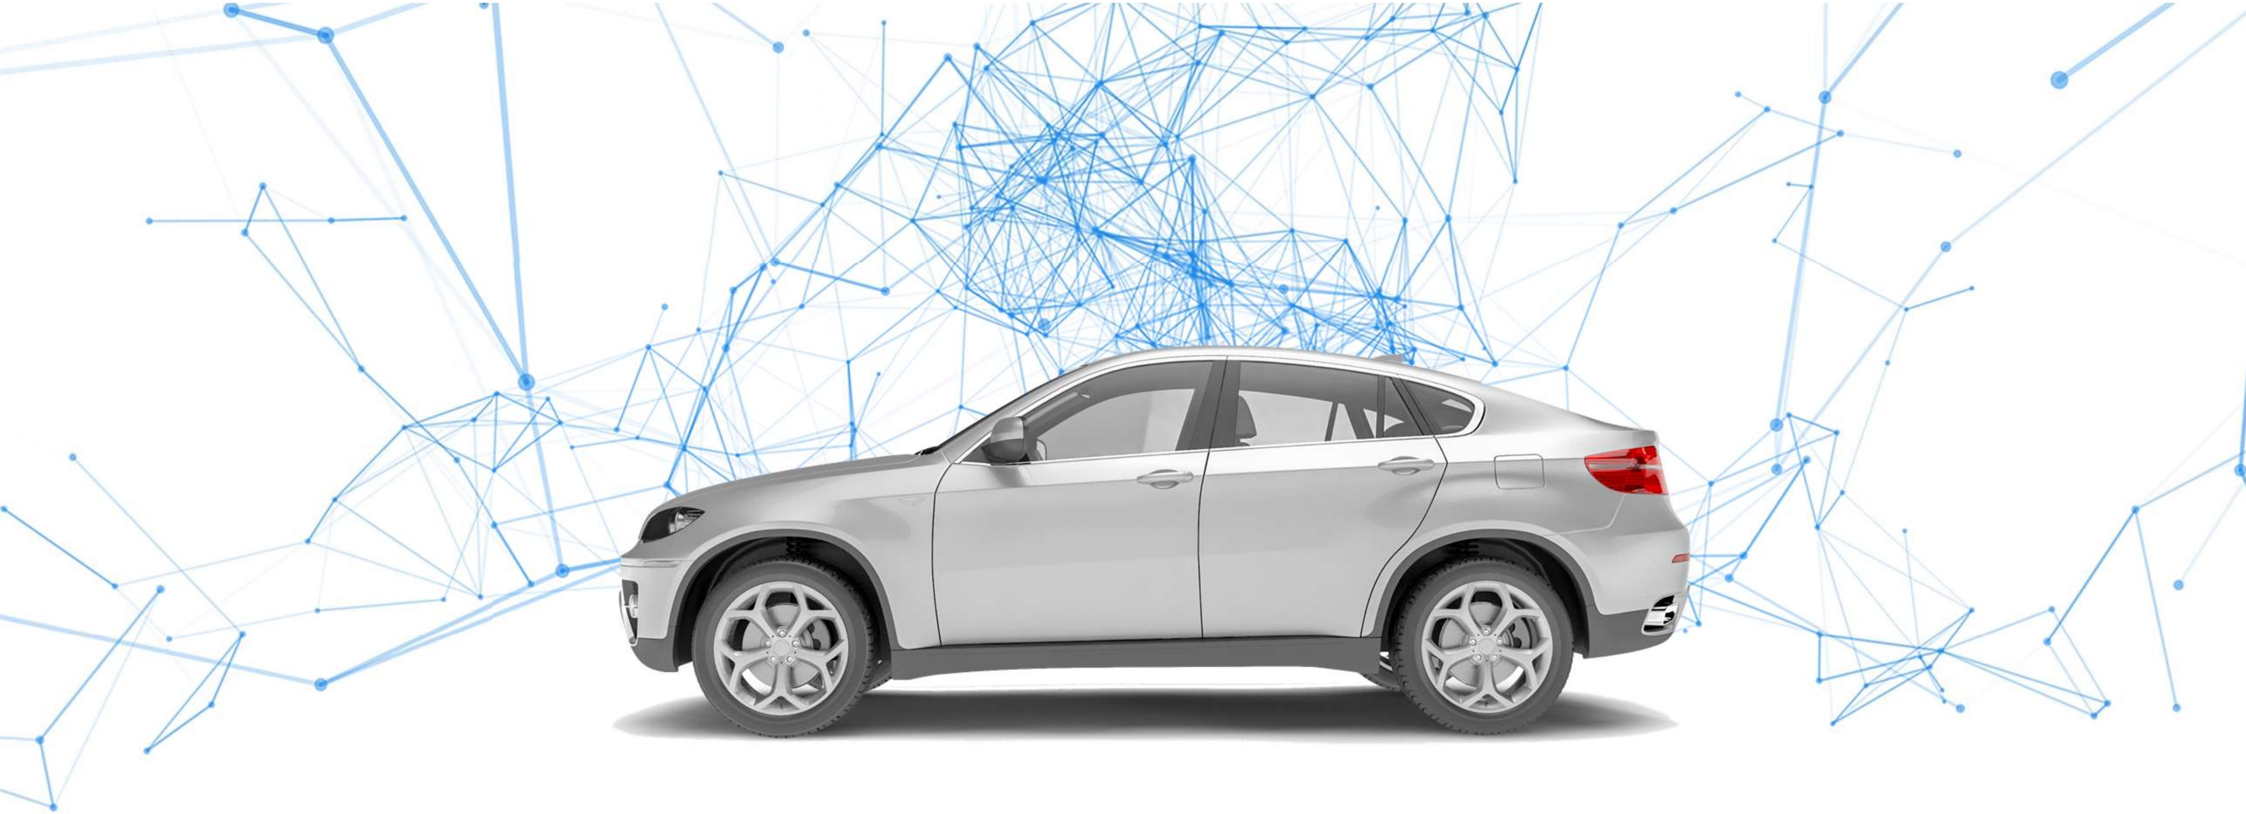

THANK YOU

ETHERNOVIA | max.turner@ethernovia.com