ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

Cut-Through Considerations and Impacts to Industrial Networks

May, 2017

JORDON WOODS

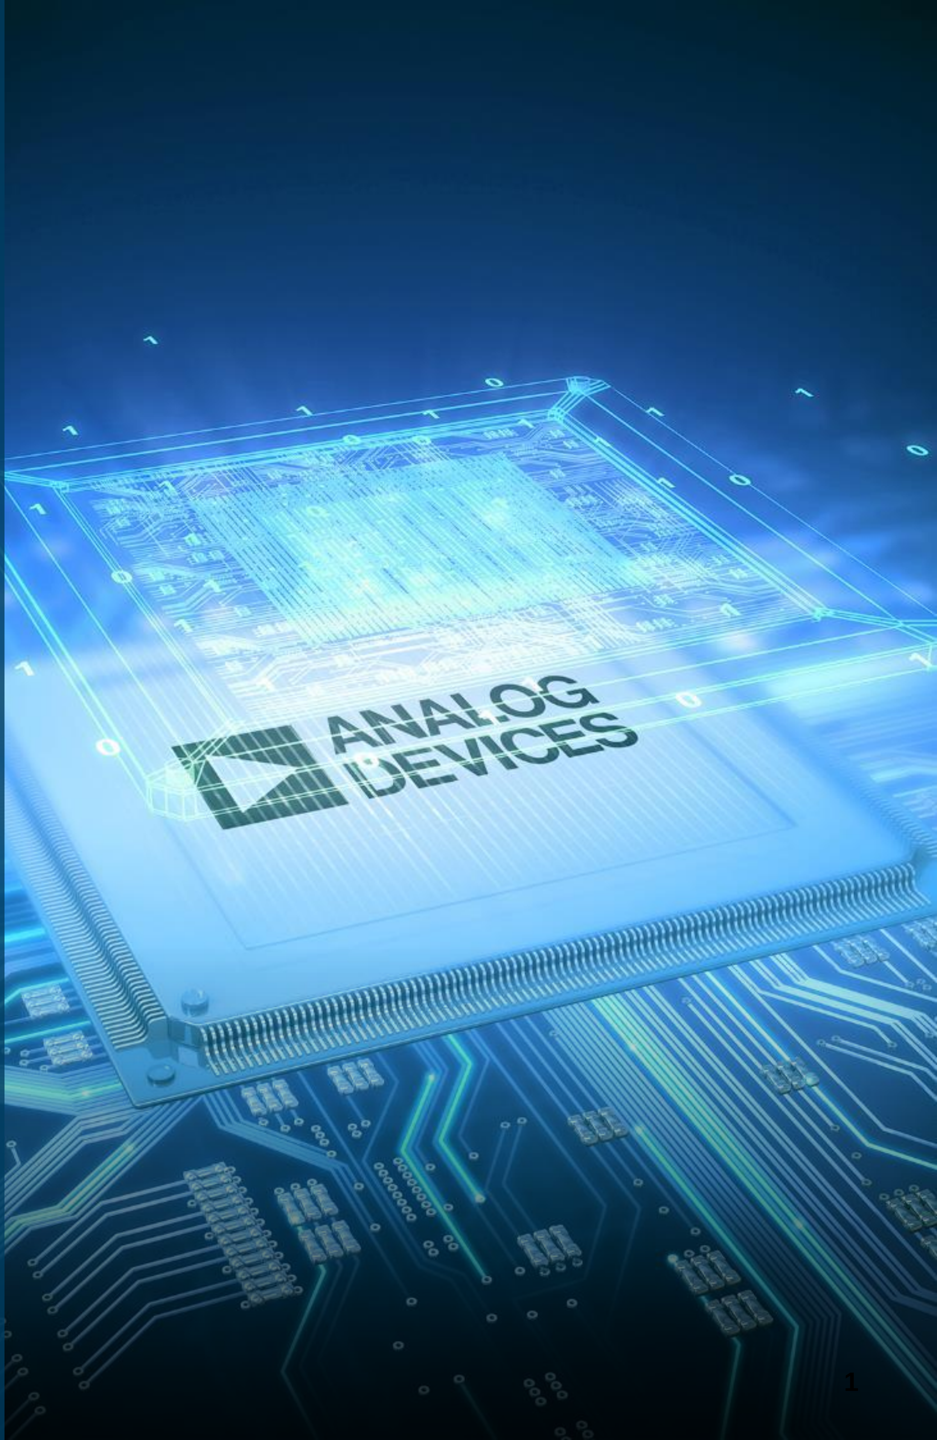

Background

- During the November, 2016 IEEE 802 Plenary. AVnu presented a liaison requesting guidance regarding the use of cut-through with IEEE802 technologies
  - http://www.ieee802.org/1/files/public/docs2016/liaison-woods-Avnurequest-1116-v00.pdf
- IEEE Responded with a request for contributions
  - http://www.ieee802.org/1/files/public/docs2016/liaison-response-avnu-1116-v01.pdf
  - Unfortunately, AVnu did not receive this request for contributions at the January IEEE 802.1 Interim meeting until the meeting was underway. Therefore, we were not prepared to contribute to the discussion.
- However, a contributions outlining some concerns regarding the use of cut-through technologies was made at that meeting. (Thank you, Pat Thaler).
  - http://ieee802.org/1/files/public/docs2017/new-tsn-thaler-cut-through-issues-0117-v01.pdf
- Avnu provided a response partially addressing some of the concerns and providing use cases for cut-through.
  - http://www.ieee802.org/1/files/public/docs2017/liaison-AVnuResponseCutthrgh-0313-v00.pdf
- The 802.1 WG agreed that the topic warranted further discussion and requested that the dialog be advanced via individual contributions.
- This contribution is intended to continue the dialog and hopefully provide context for the discussion.

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

Agenda

- A review of use cases
  - Control Applications (line topologies)
  - Preemption and cut-through
  - Redundancy (ring topologies)
- Cut-through implementations in industrial automation
  - Performance and forwarding
  - Risk Mitigations
- Specifying Cut-through in IEEE802

Use Case 1 - Control Applications (line topologies)

- Control Applications (line topologies)
  - Utilization of line topologies is prevalent in motion applications utilizing embedded switch technology
  - There can be many hops along the line (64 hops or greater)
  - As indicated in the model, switch latency along these hops accumulates, eating into the time available for updates.
  - The schedule of drives can be individually adjusted to compensate for drive transmission delay and average switch latency (NOTE: Schedule does not necessarily refer to .1Qbv, scheduling may take place in the application).
  - However, the effects of these delays are cumulative. Each delay per hop consumes part of the time available during the cycle.
  - This is really a question of the accumulated latency per hop.

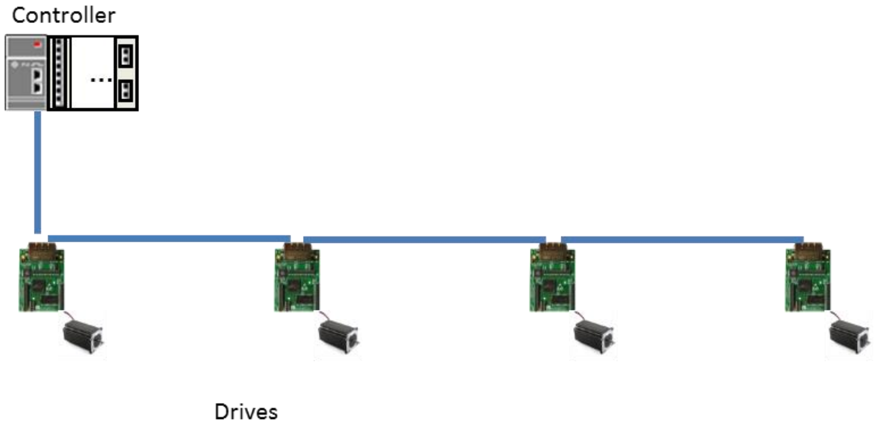

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

Use Case 2 - Preemption and cut-through

- Does cut-through only apply to .1Qbv?
  - No, preemption offers a means to limit the effect of interfering traffic on control traffic without the added complexity of scheduled traffic.
  - At the moment that an express frame preempts a best-effort frame, the conditions for cut-through apply, meaning that you know that the express frame can cut-through.
  - Properly engineered, line topology limits the effects of interfering traffic to a single hop (i.e. control traffic is transmitted in a burst) assuming preemption is enabled
  - With preemption, the effects of interfering traffic are minimal with respect to a 1 mS update cycle

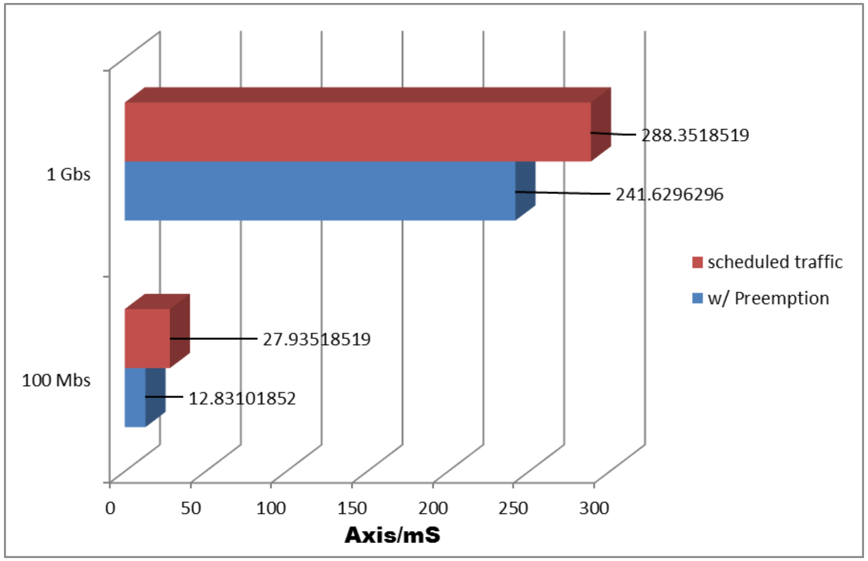

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

Use Case 3 - Redundancy (ring topologies)

- Typical topology for redundancy in industrial networks is a ring:
  - inherently different packet times on the network along the different routes
  - depending on the setup, packet time on the two paths has extreme deviation
  - depending on allowed reception window of redundancy mechanisms, ring size is limited
  - Additionally ring size is limited by tolerable jitter of RT-application
- With 100 Mbit/s, 300 byte packet and 100µs packet deviation:
  - store and forward: max. difference in path is 4 hops
  - cut-through: max. difference in path is up to 38 hops
- With 1GBit/s, 300 byte packet and 100µs packet deviation:
  - store and forward: max. difference in path is 34 hops
  - Cut-through: max. difference in path is up to 141 hops

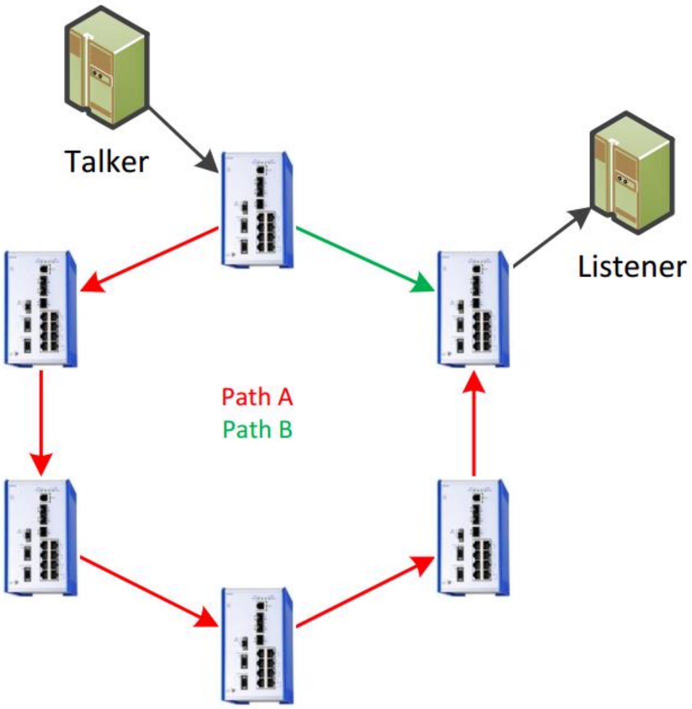

Cut-through performance

- In Industrial use cases, there are two basic approaches to the timing of cut-through:
  1) Ensure that a minimum of 64-bytes have been received before starting transmission of a frame to avoid propagation of runt frames.
  2) Receive the minimum number of bytes necessary to make a forwarding decision.
- For most industrial protocols, the avoidance of runt frames is not a major consideration. Waiting for a minimum sized frame avoids only that single class of potential errors.

Cut-through performance

- This leads us to the number of bytes necessary to make a forwarding decision. There are various answers depending on the forwarding process. Some of the most common (though certainly not all approaches) are shown below:
  a) Destination address only
  b) Destination address and VLAN Tag (if present)
  c) Destination address, EtherType and a protocol-specific field (assumes no VLAN header)
- In addition to delays incurred for the received bytes, there is also the receipt/transmission of the preamble, any table lookup time, and queuing delays.

Cut-through performance

- To calculate the cut-through delay for a) through c) above we'll use the following formula
  - ***Switch Delay = (P+Nb) * Tb + Lu + Q***
- Where:
  - ***Switch Delay*** is the time from receipt of SFD on the ingress port to the transmission of SFD on the egress port
  - ***P*** = number of bytes in the preamble
  - ***Nb*** = number of data bytes in the frame necessary to make the forwarding decision
  - ***Lu*** = look-up/processing time to compute forwarding destination
  - ***Q*** = internal queueing times (including MAC traversal, memory delays, etc.)
  - ***Tb*** = Time necessary to transmit a byte (e.g. 80 nsec for 100 Mbit, 8 nsec for 1Gbit)
- So, on a high-performance cut-through switch you can have numbers something like:
  - ***Lu*** = 160 nsec (this process is simplified on a two-port switch)
  - ***Q*** = 320 nsec

Cut-through performance

- Now for the cases above with an 8-byte preamble at 100 Mbit you'd see
  a) ***Switch Delay = (8 + 6) * 80 + 160 + 320 = 1.6 usec***
  b) ***Switch Delay = (8 + 18) * 80 + 160 + 320 = 2.56 usec***
  c) ***Switch Delay*** = same as B (Ethertype and 16-bit protocol-specific field same delay as 32-bit VLAN header)
- If we apply the same values to a Gbit interface we see
  a) ***Switch Delay = (8 + 6) * 8 + 160 + 320 = 592 nsec***
  b) and c) ***Switch Delay = (8 + 18) * 8 + 160 + 320 = 688 nsec***

Cut-through performance

► If we take a minimum small frame (64-bytes) and a large frame (1500 bytes) and have each traverse a 64-node line network we see the difference in latency. We'll use case b) above and ignore PHY and cable delays for this computation.

► **100 Mbit Cut-through Network**, 2.56 usec switching delay per hop, 64 hops -> latency = 64 * 2.56 = **163.84 usec** switching delay, for both the small frame and the large frame

► For a store-and forward approach assume frame time (e.g. 64-bytes + 8-byte preamble) and 480 nsec queueing and switching delay.

► **100 Mbit store-and-forward Network** -> ((64 + 8) * 80 + 480) * 64 nodes = **399.36 usec switching delay** for the small frame

► **100 Mbit store-and-forward Network** -> ((1500 + 8) * 80 + 480) * 64 nodes = **7.75168 msec switching delay** for the large frame

Risk Mitigations in Industrial Use cases

- Risk: Little benefit in case of different link speeds at bridge ports
  - Industrial networks are usually:
    - heavily engineered
    - using the same link-speed (at least within network segments)
    - a line or a ring topology
  - **Exactly** the situation where cut-through will offer its benefits

Risk Mitigations in Industrial Use cases

- Risk: Bit errors in headers can change fields including address, VLAN, and priority fields leading to incorrect forwarding:
- Mitigations in Industrial Use Cases
  - Mitigated by FCS: the receiving node will still detect a bad FCS
  - Mitigated by application:
    - Many industrial protocols are connection-based meaning received packets without the correct connection ID (or equivalent) are dropped
    - Applications are typically tolerant of 2-3 missed updates

Risk Mitigations in Industrial Use cases

- Risk: Bit errors in headers can change fields including address, VLAN, and priority fields leading to incorrect forwarding:
- Mitigations in Industrial Use Cases
  - Mitigated by topology: line topologies minimize the opportunity for misrouted traffic to compromise the network
  - In ring topologies, zombie frames are prevented from infinitely circulating by:
    - A ring “master” which blocks traffic on one of its ring ports, effectively establishing a line topology
    - Special HW or SW specifically designed to detect and eliminate zombie frames (HSR)
    - Special timing (i.e. the frame is only allow to forward during a particular schedule period).

Risk Mitigations in Industrial Use cases

- Risk: Bit errors in headers can change fields including address, VLAN, and priority fields leading to Higher congestion risk and violation of delay guarantees:
- Mitigations in Industrial Use Cases
  - Mitigated by design: traffic in control segments is typically constrained and control traffic packets are small.
  - Mitigated by application:
    - Applications are typically tolerant of 2-3 missed updates

Risk Mitigations in Industrial Use cases

- Risk: Bit errors in headers can change fields including address, VLAN, and priority fields leading to security concerns (i.e. Packet payload may become visible on links where it shouldn't be seen):
- Mitigations in Industrial Use Cases
  - Mitigated by topology: line topologies offer little opportunity for misrouted traffic to compromise the network
  - Mitigated by application:
    - Confidentiality is not a primary concern in these use cases
    - The larger problem is authentication which is not adequately addressed in this market. Typically, these applications cannot tolerate the hop-to-hop latency introduced by MacSec or similar authentication schemes.

Specifying Cut-through in IEEE802

- What is specified?
  - Behavior (when do we make the forwarding decision)?
  - Management (controlling cut-through and reporting performance)?
- How to we specify and limit the impact to IEEE802 standards?
  - A single “special” cut-through traffic class
  - No queueing required
  - Potentially a form of “express-traffic” path through the bridge

ANALOG DEVICES

AHEAD OF WHAT'S POSSIBLE™

THANK YOU