ANALOG DEVICES

AHEAD OF WHAT'S POSSIBLE™

New Ethernet Applications - Accumulated switch latency in industrial applications

January 24, 2018

JORDON WOODS

Director, Deterministic Ethernet Technology Group

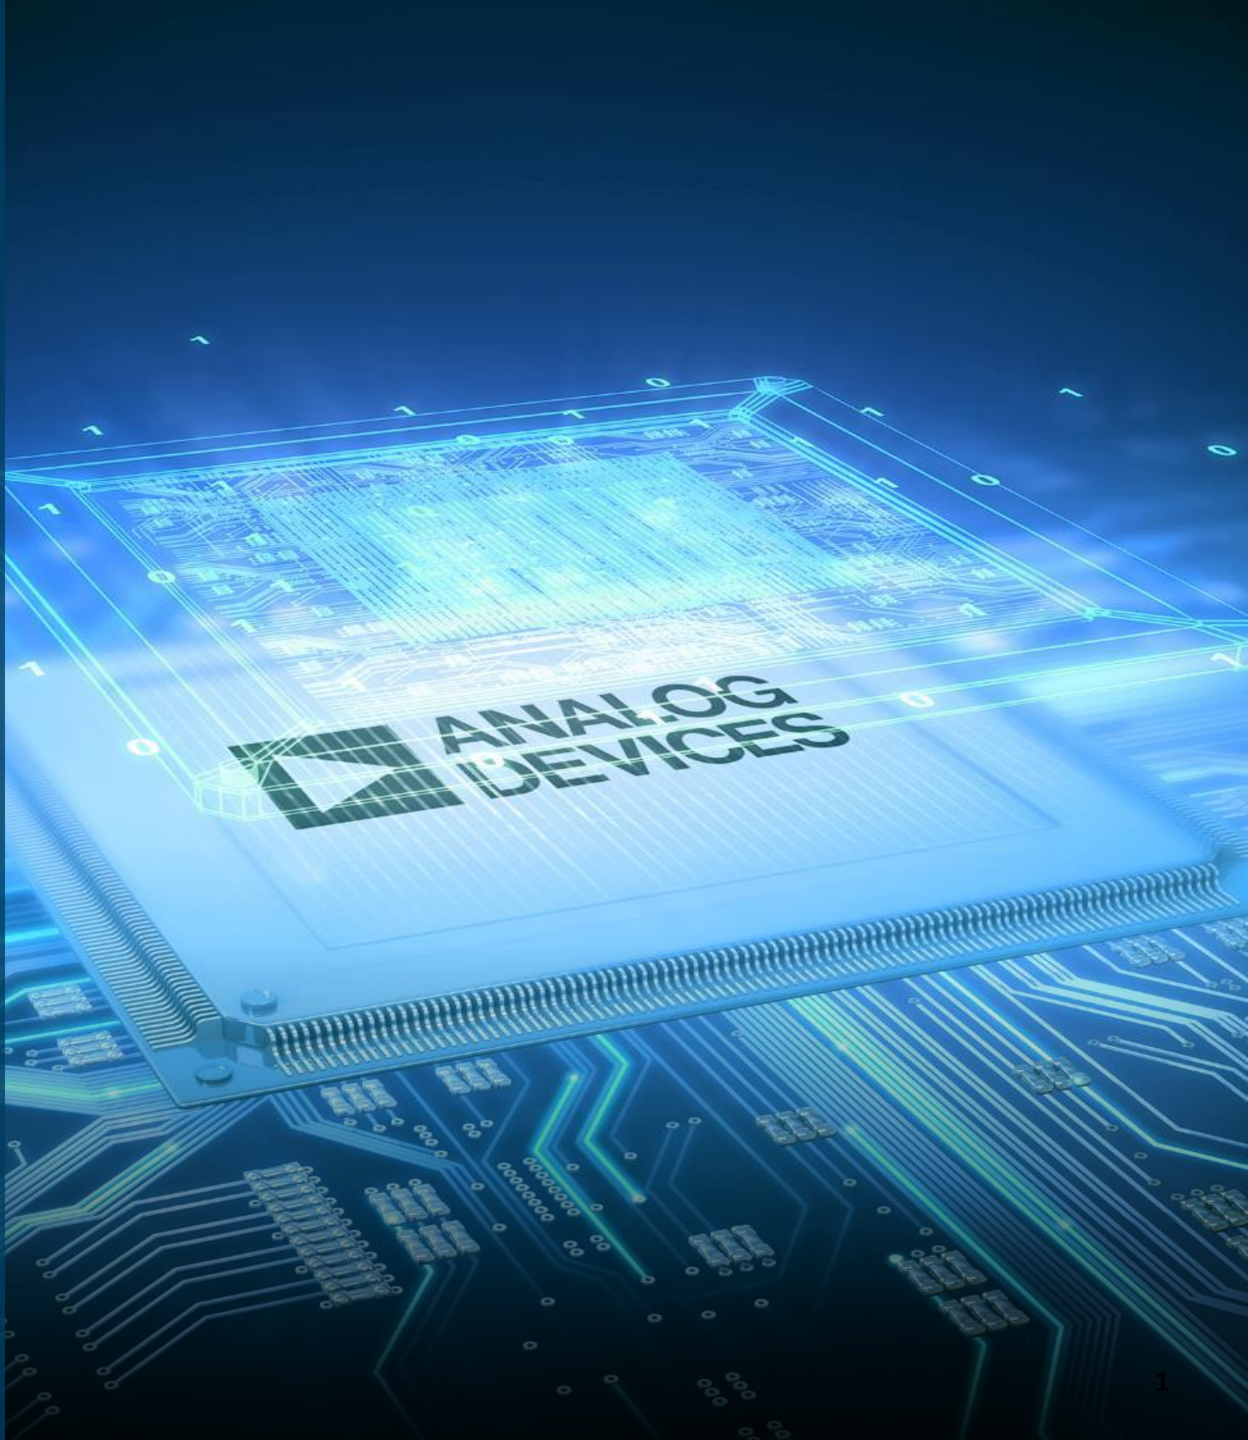

Synopsis of the Problem

- Industrial applications, such as machine control, are typically built in long line configurations. For these installations, to minimize wiring cost and complexity, typical installation uses “daisy chain” where each node has (2) external switched ports and an internal port that goes to the end-node.

- A common application is motion control where fast loop times are required. 125 µs cycle rate is common. To support this, low latency for messages through the network is a high priority.

- Even Gigabit data rates are not sufficient to solve this problem. For instance, in a line topology of 64 hops, accumulated latency would exceed a 100 µs control loop even at Gigabit speeds.

- Theses systems often also have high EMC and there is a desire in some applications to support brown-field wiring. For these applications 100Mb/s rates are desired.

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.

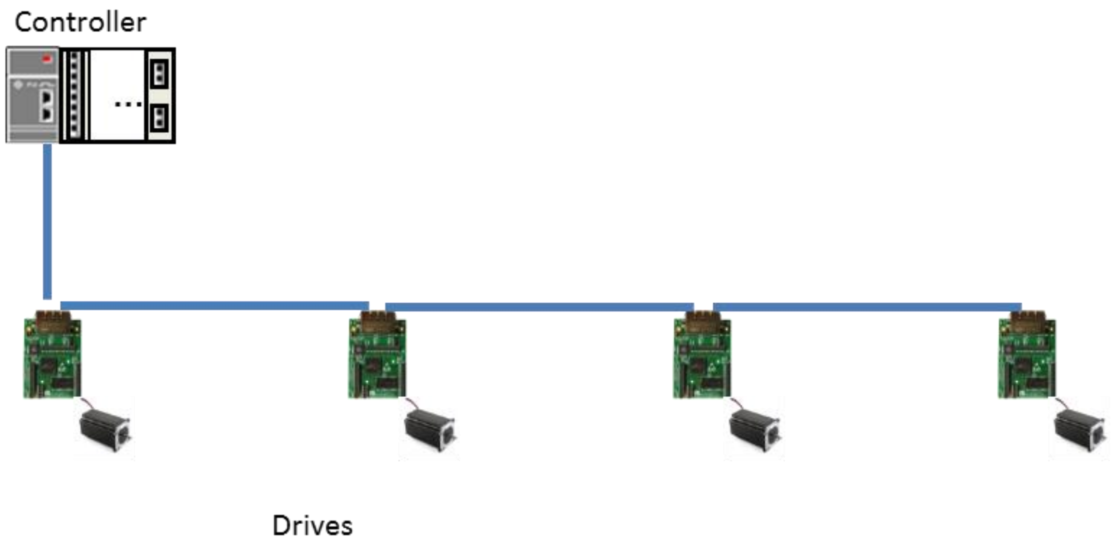

Why line topologies?

- Physical constraints make cabling for star topologies impractical
- The construction of the application naturally lends itself to point-to-point connectivity

A Simple Motion Control Model

- We'll focus on a part of the problem associated with network performance
- Ideally, we'd like all of the drives to transmit their output data simultaneously
- In this way the link between the controller and bridge is optimally utilized
    - Note: due to control loop timing constraints, this model assumes full-duplex operation.

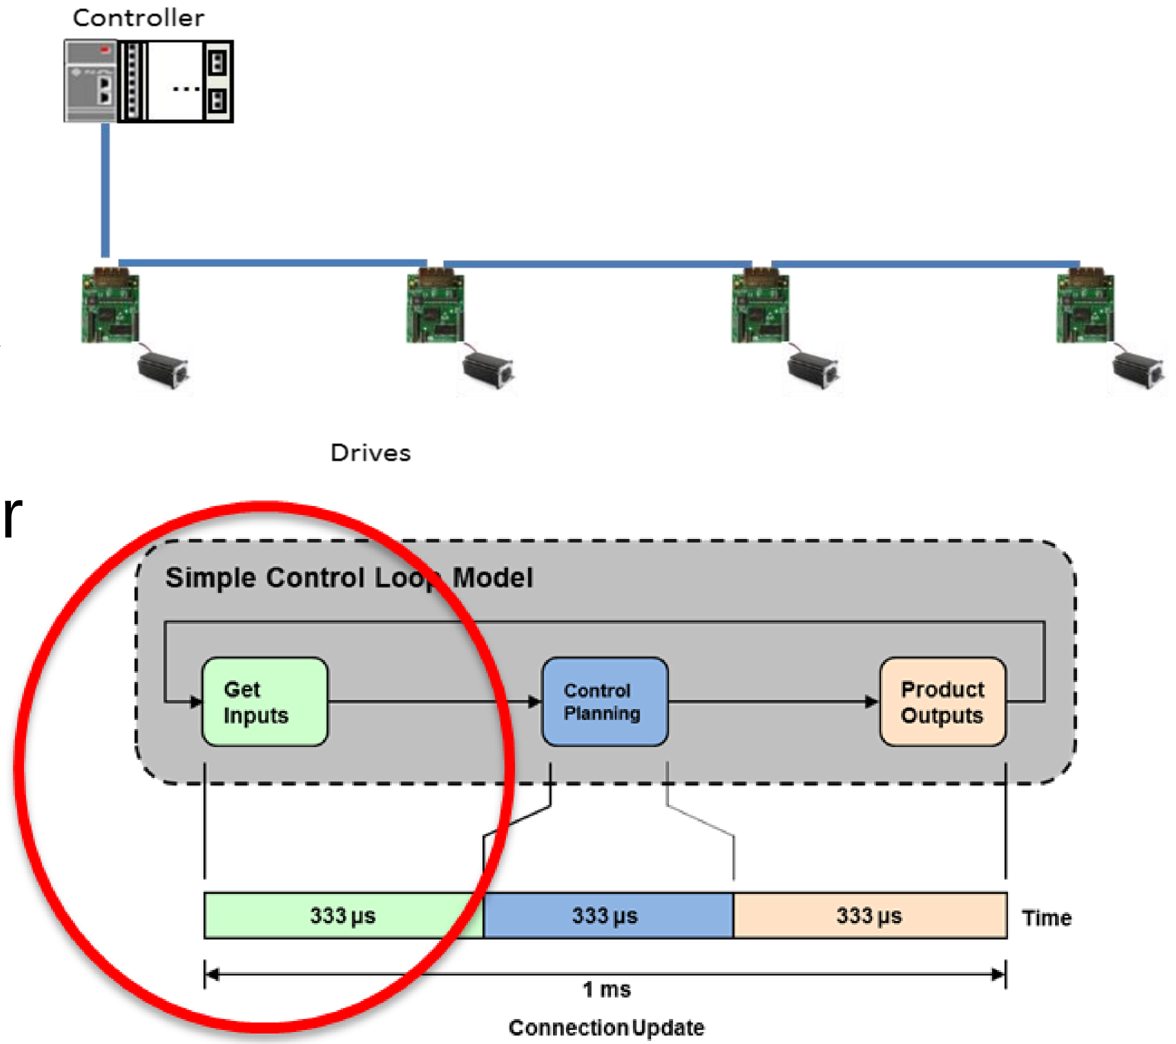

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

A Simple Motion Control Model

- *Max Axis = 1 + {1/3 * Connection Update Period – (Drive Transmission Delay + (m + 1) * Ethernet Transmission Time + m * Switch Latency + NIC Packet Processing Delay + Bus Interface Delay)}/NIC Packet Processing Delay*

  - *(Where m = # of hops)*
  - *Drive Transmission Delay:* We'll assume all drives have outputs queued prior to transmission, so this is contribution is small with respect to other operands, effectively 0 *usec*
  - *Assume update packets are fairly small(124 bytes), so Ethernet Transmission Time is (124+20)*80ns/byte = 11.52 usec (at 100 Mbs)*
  - *Switch Latency = (interfering packet size+20)*80ns/byte*
  - *NIC Packet Processing Delay – There are techniques to ensure the network is the bottleneck (e.g. 2 cycle processing): 11.5 usec for 100 Mbs, 1.15 for Gigabit.*.
  - *Bus Interface Delay: has a lot to do with the overall system architecture. could go effectively to 0 (given good bus structure, DMA/ etc.). We'll assume 0 for this analysis.*

** Chaffee, Mark. "CIP Motion Implementation Considerations." Proc. of ODVA 2009 Conference & 13th Annual Meeting, Howey-in-the-Hills, Florida USA.*

Model at 100 Mbps

- Performance strongly influenced by interfering traffic and thus, limits the # of hops
- In practice, control systems will engineer the network to limit the size of interfering packets (In this case, 500 Bytes)

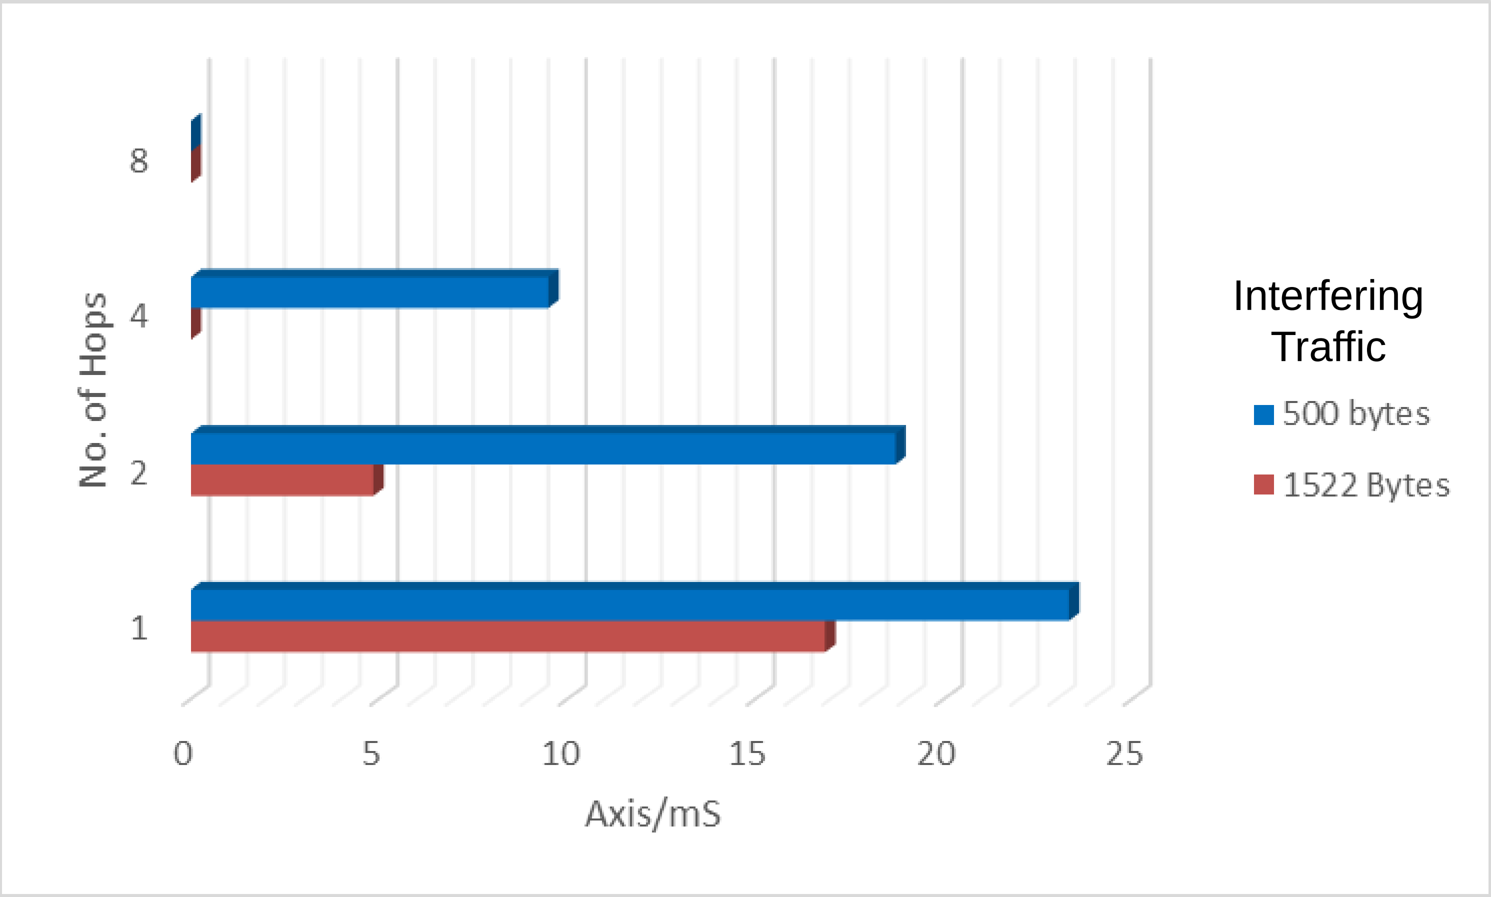

Model at 1 Gbps

- Gigabit line rates help, but still fall short of required performance in long lines

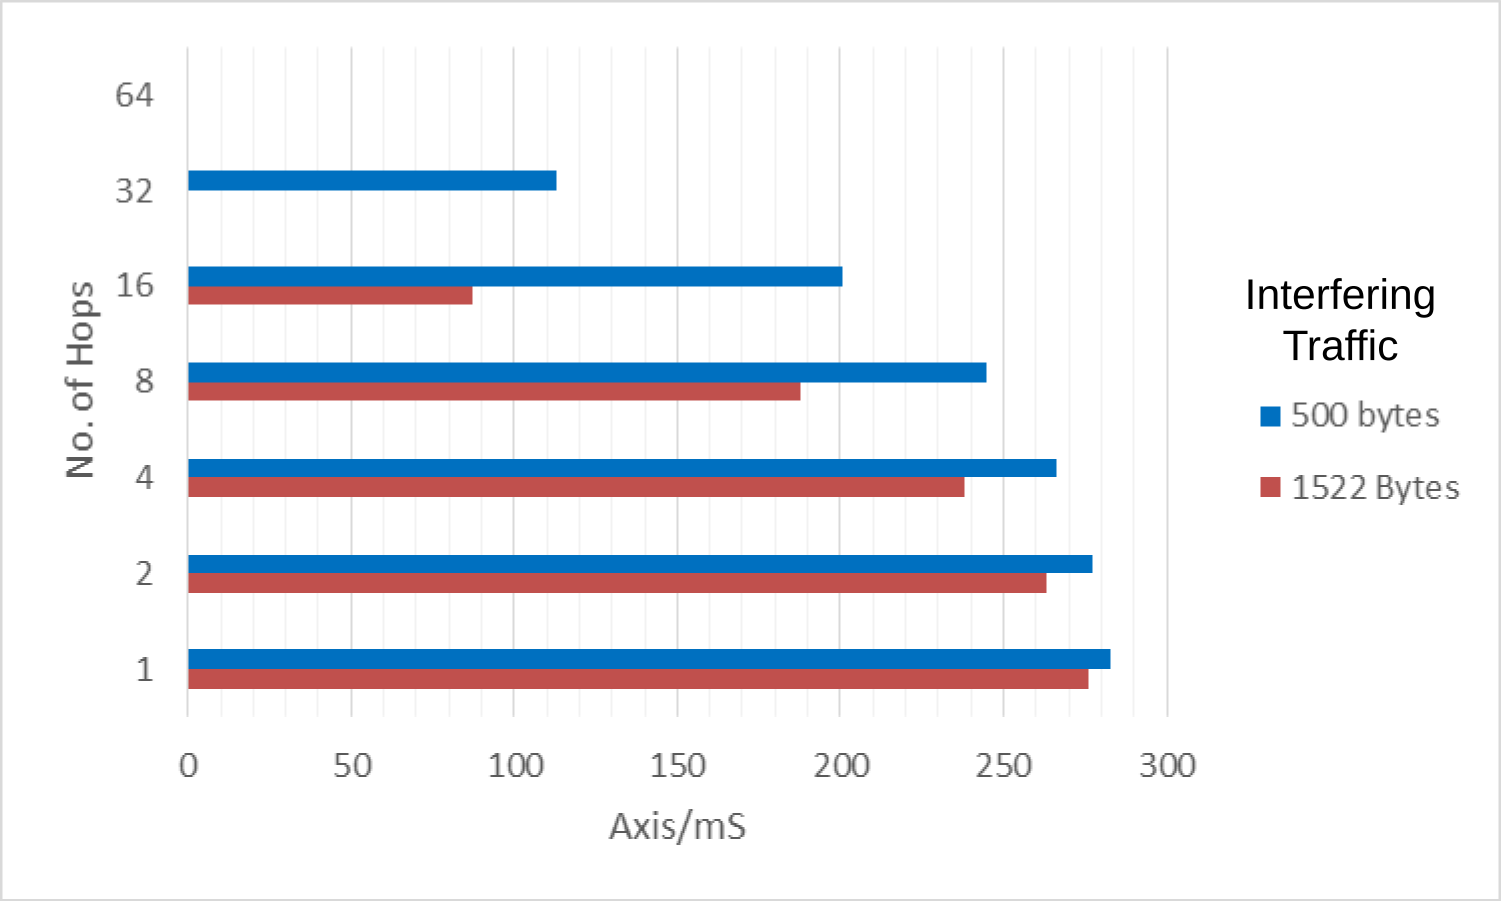

Gigabit Performance at Faster Cycle Times

- In fact, there are many precision motion applications which require even greater performance
- Gigabit line rates fall short of satisfying these requirements

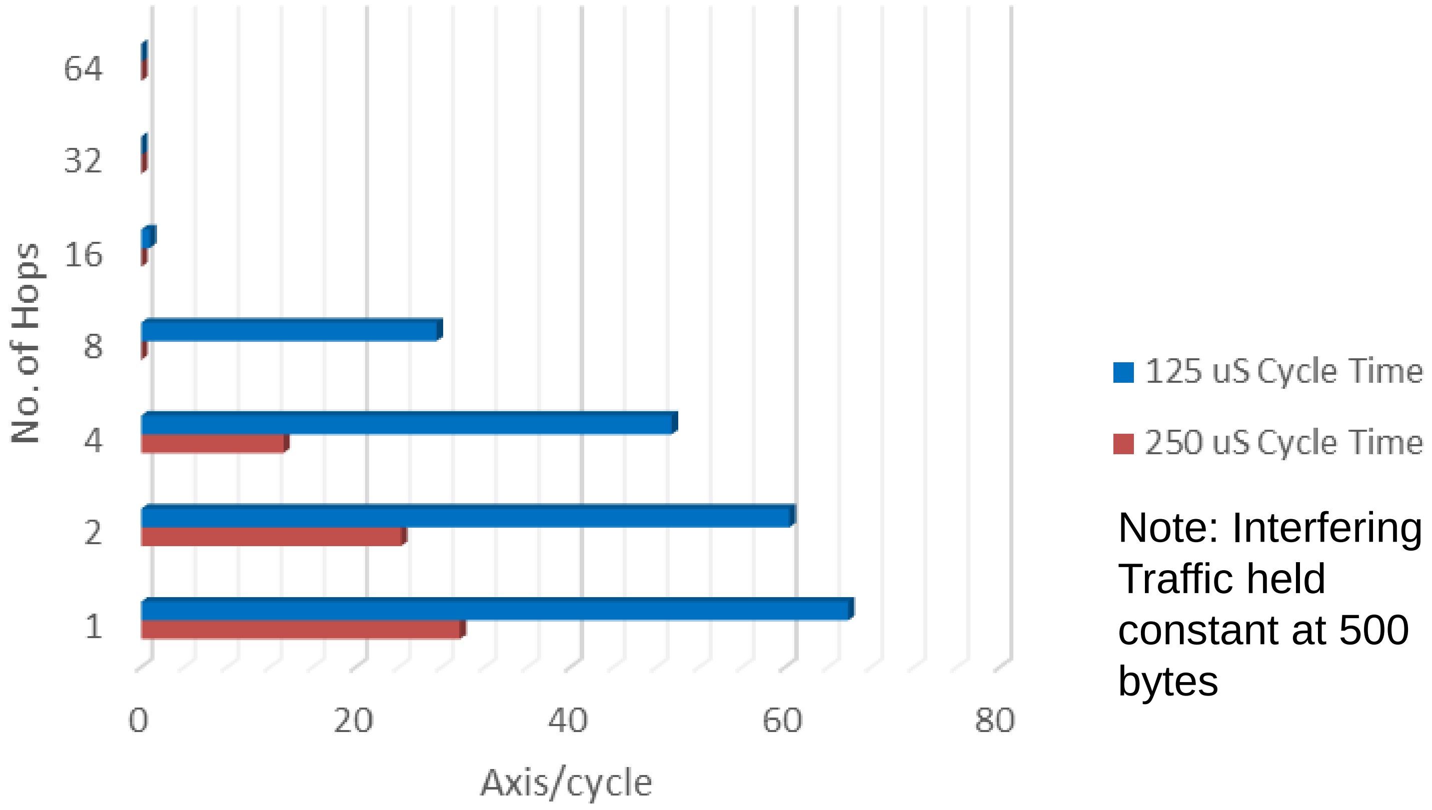

Note: Interfering Traffic held constant at 500 bytes

Use Case 2 - Redundancy (ring topologies)

- Typical topology for redundancy in industrial networks is a ring:
  - Inherently different packet latency on the network along the different routes
  - Depending on the setup, packet latency on the two paths can have extreme deviation
  - Depending on the allowed reception window of redundancy mechanisms, ring size is limited
  - For instance, for a 300 byte packet and 100 us packet deviation:
    - At 100 Mbit/s: the max. tolerable difference in the path is consumed in 4 hops
    - At 1 Gbit/s: the max. tolerable difference in the path is consumed in 34 hops

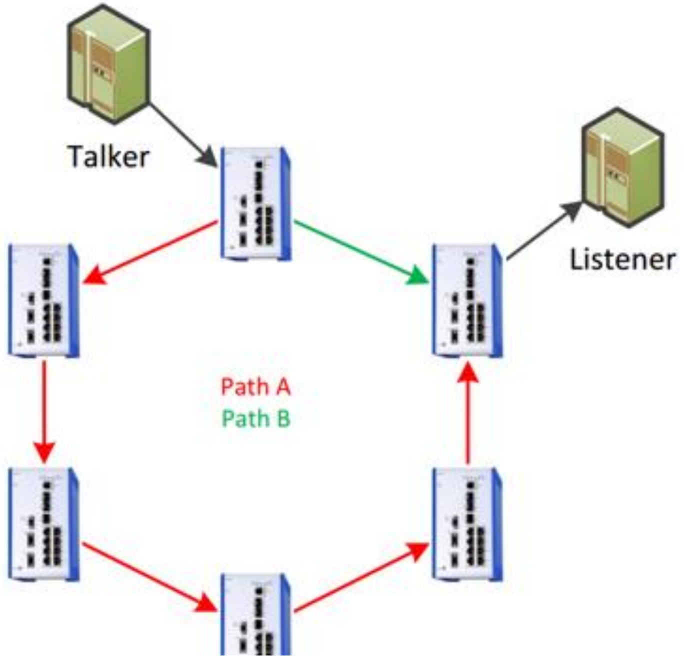

Looking to the future

- The promise of Industrie 4.0 has led to a desire for increased performance and network convergence amongst automation companies.
- This trend, in turn, has led great interest in TSN technologies amongst automation companies.
- The iIOT promises to be significant market and an early adopter of TSN technologies.

Looking to the future

- Today, industrial networks solve these problems using a variety of techniques.
  - In many cases, these techniques violate IEEE802.3 and IEEE802.1 standards.
- Therefore it seems prudent that IEEE802.3 and IEEE802.1 consider these use cases to ensure an approach consistent with IEEE standards.

AHEAD OF WHAT'S POSSIBLE™

THANK YOU