Worst Case Latencies of Gen2 Proposals

Deggendorf
University

Christian Boiger
christian.boiger@fh-deggendorf.de
IEEE 802.1 Interim Meeting
May 2011
Santa Fe, NM

Latency Calculations

- The latency calculations in this presentation are last bit in – last bit out
- The MAC delays are not taken into account
- The switching delays are not taken into account
- The calculations are based on blocks representing a frame plus preamble, SFD and IPG (the shown latency includes also the last IPG, for the real latency figure this last IPG has to be reduced)
- The numbers in the graphics are µs
- @FE 7.2µs (90 bytes), 7.76µs (97 bytes), 11.84µs (148 bytes) 123.36µs (1542 bytes)
- @GE 1.184µs (148 bytes), 2.4µs (300 bytes), 7.2µs (900 bytes), 12.36µs (1542 bytes)

Fragmentation

- As the fragmentation method has not been defined, the following slides base on the following assumptions:
  - There is one highest traffic class
  - Traffic which would delay the transmission of this traffic class is fragmented
  - The fragmentation happens at the very latest after a 128 bytes interference (i.e. 148 bytes incl. preamble, SFD and IPG)
  - This worst case fragment of 128 bytes (148 bytes) can always occur (e.g. as a non fragmented frame)

Fast Ethernet Class A

+

128 Byte Fragmentation

Worst Case Talker Latency

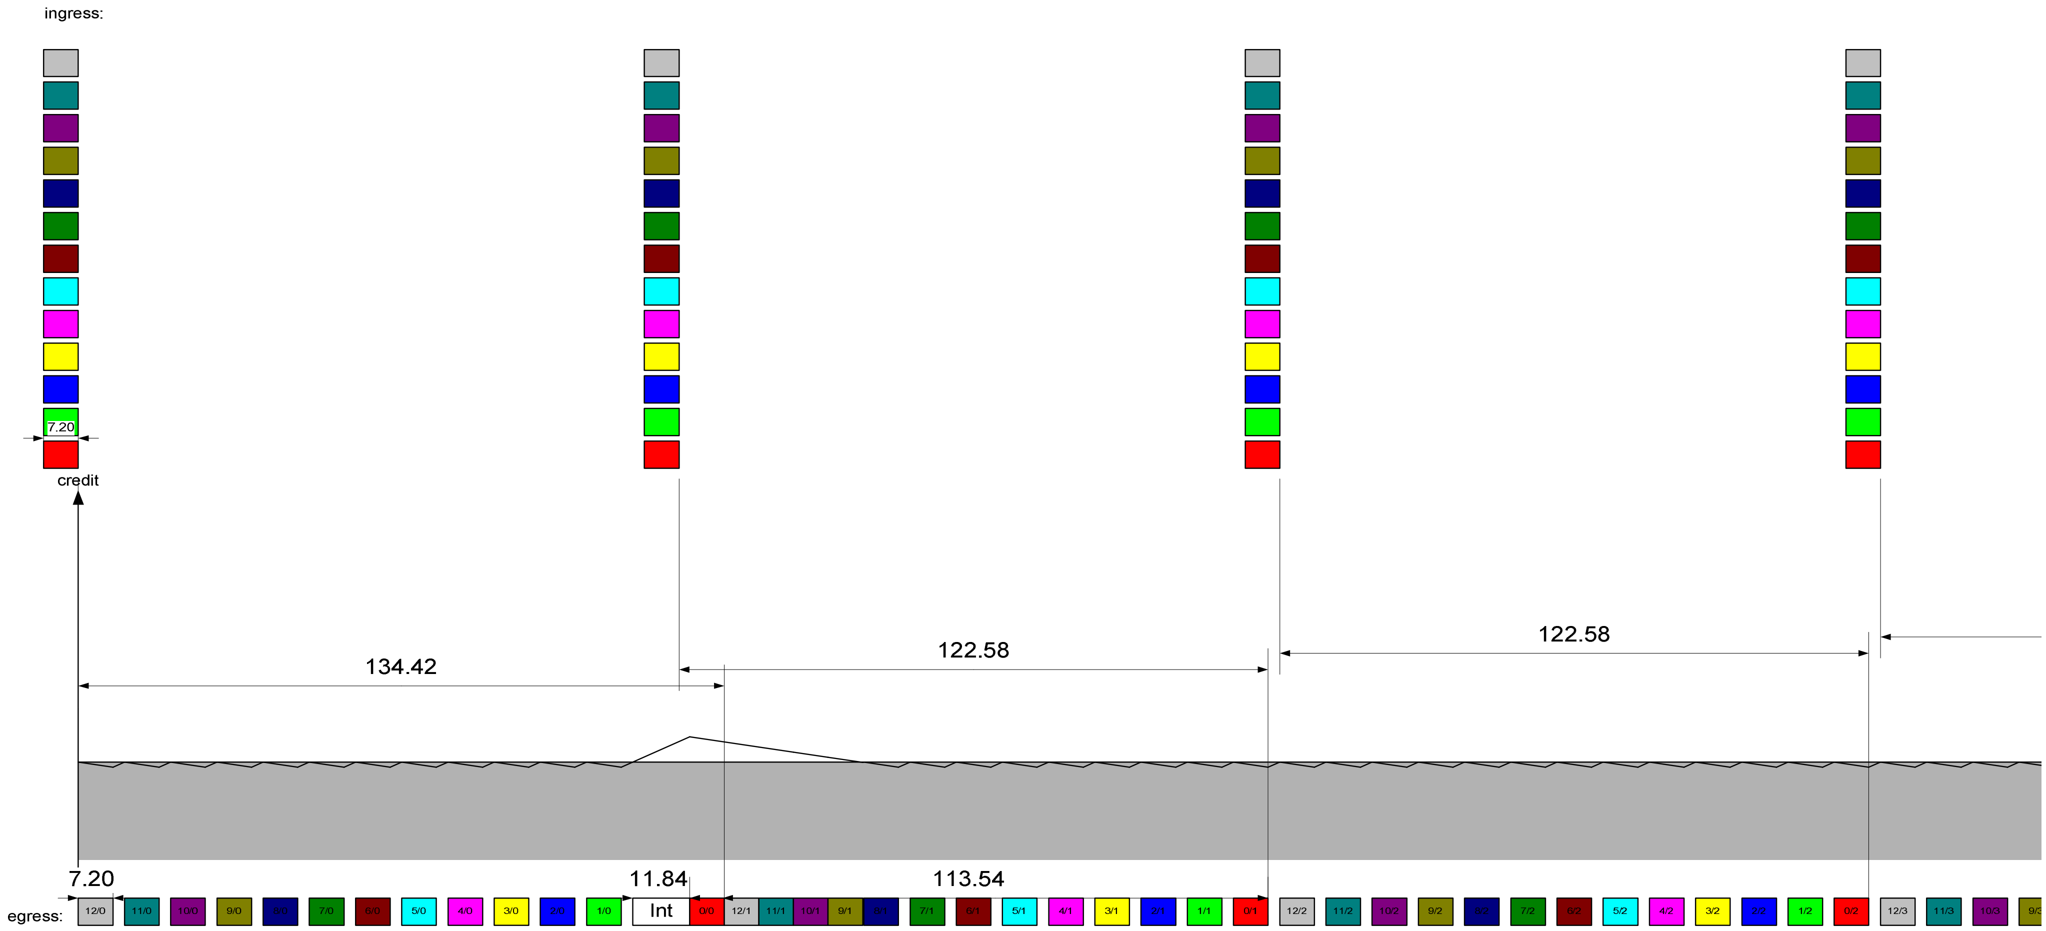

Creation of Burst

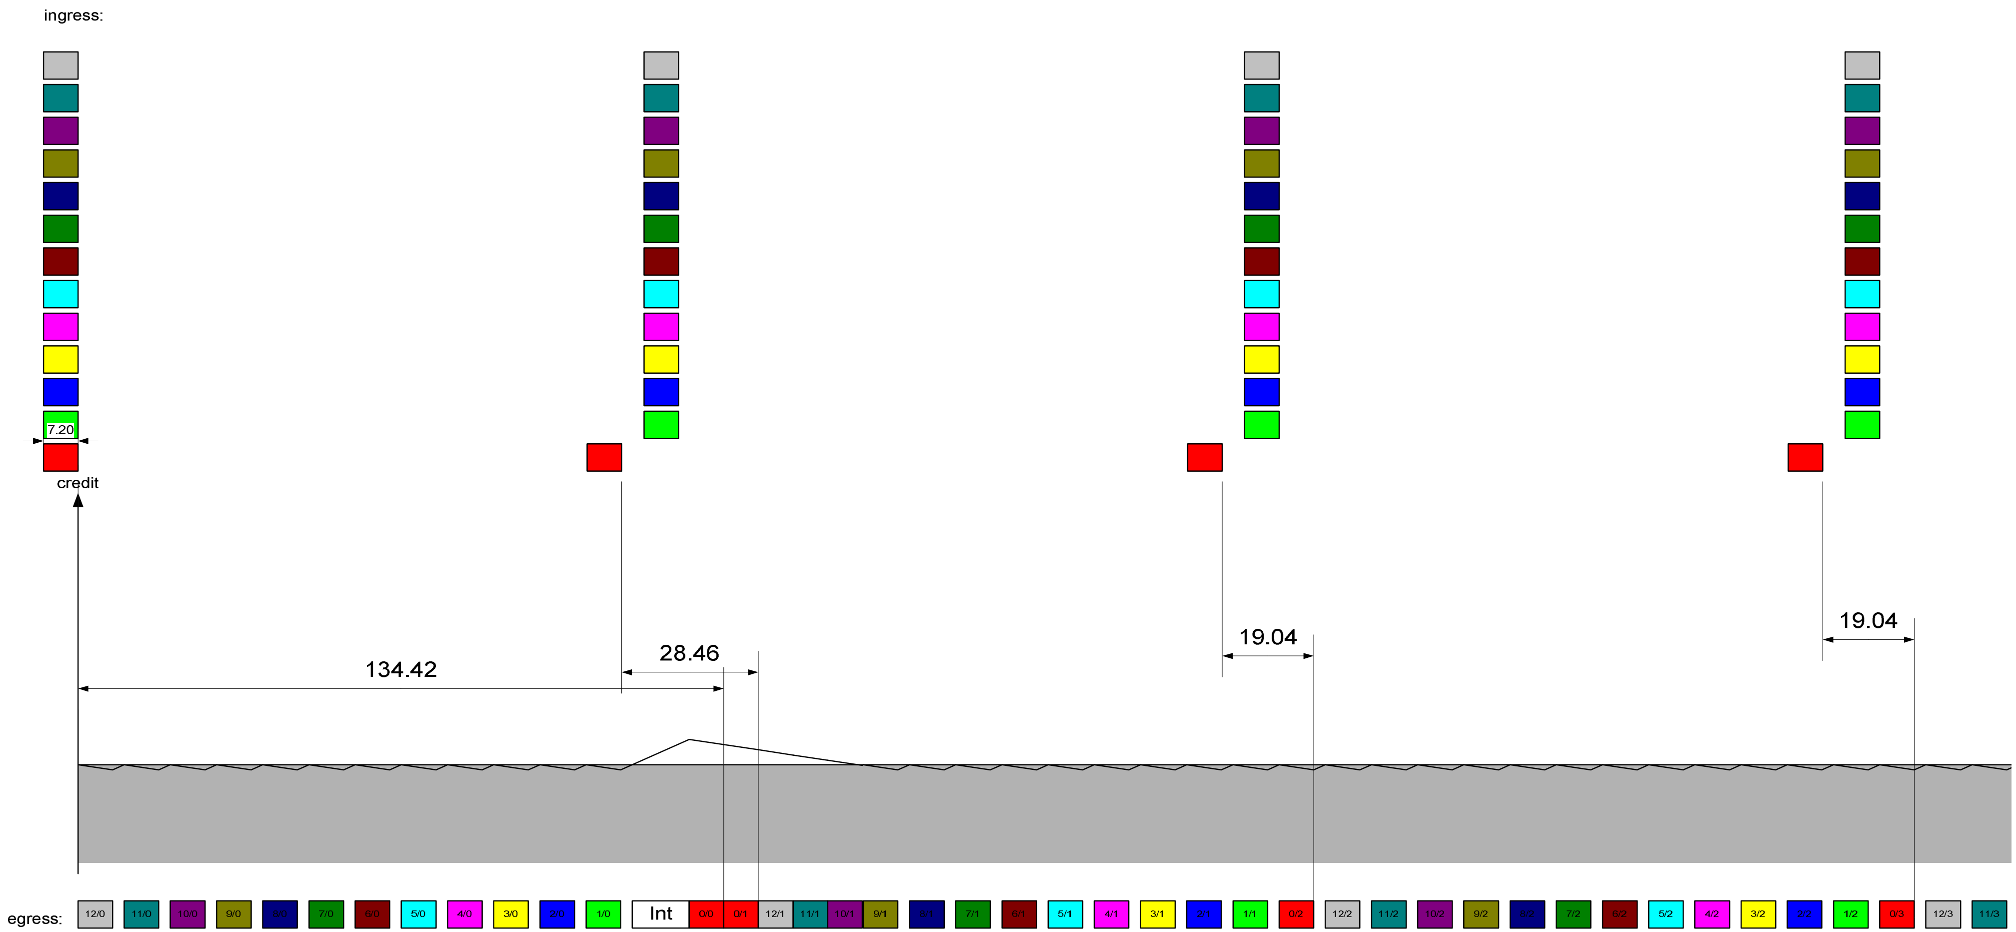

Consequences of a Bursting Ingresses

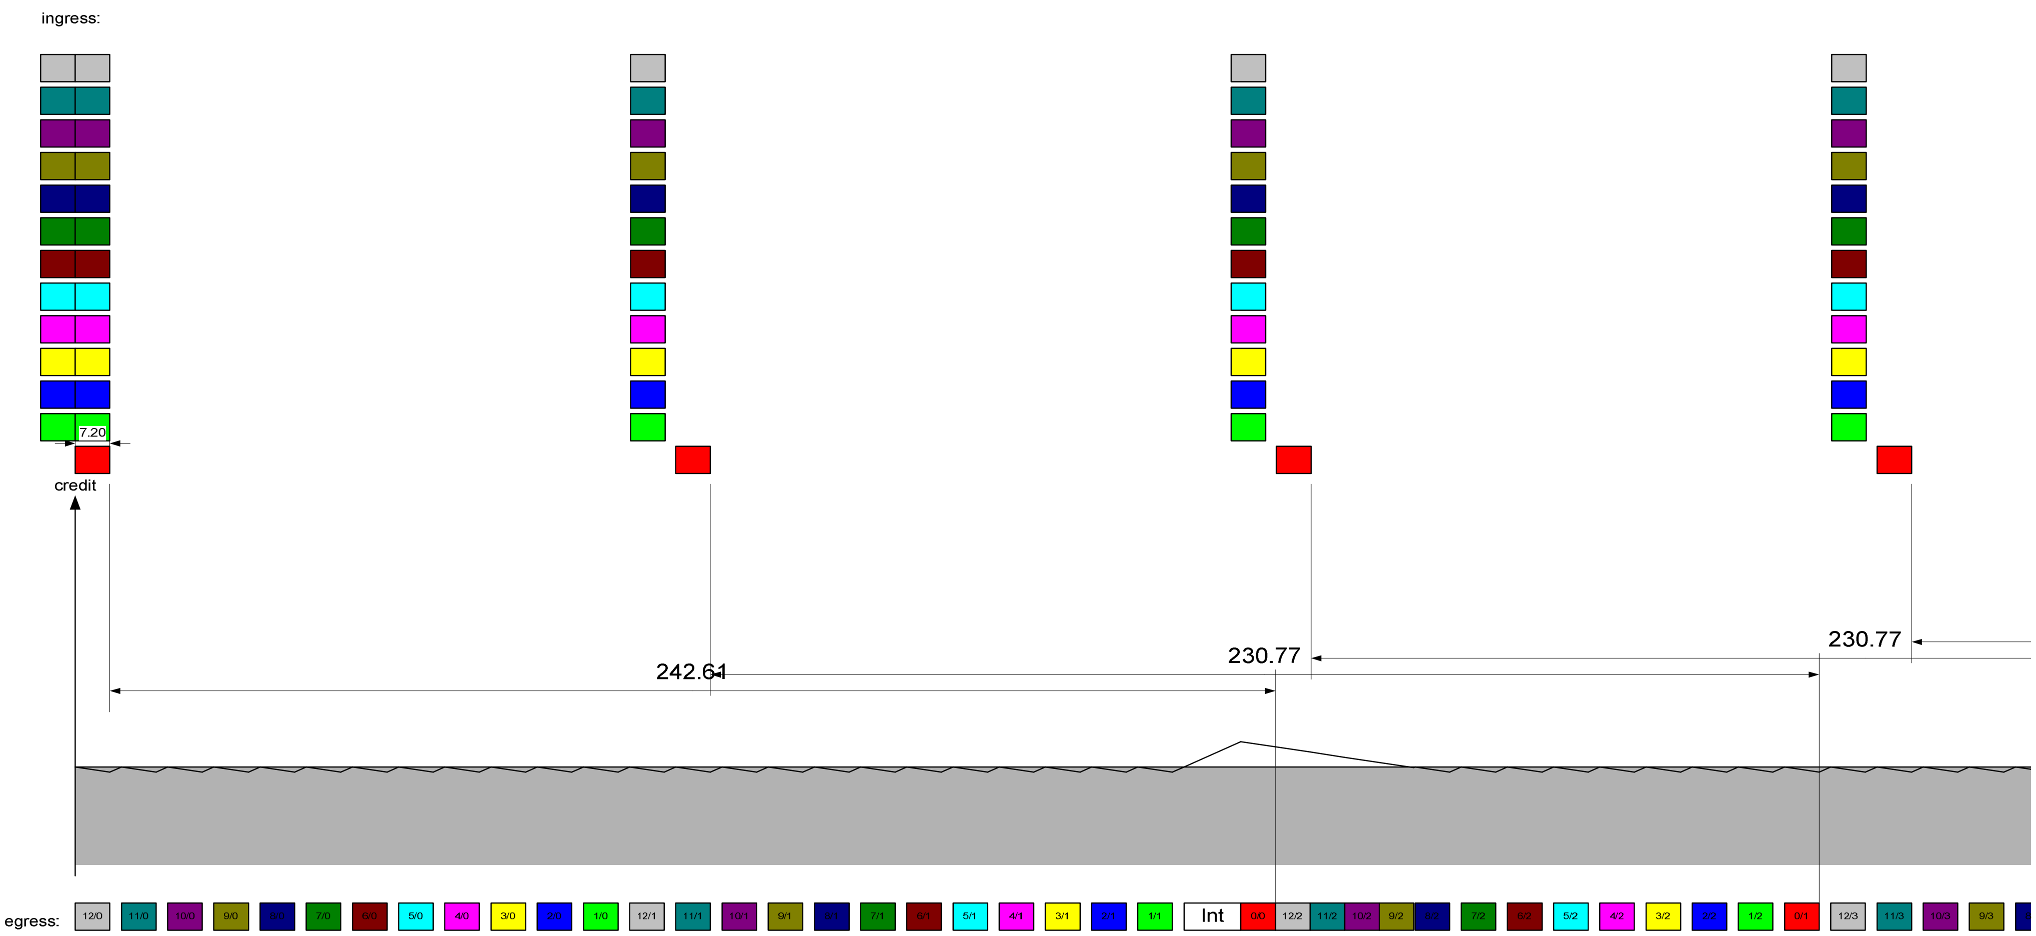

Consequences of Fragmentation @ High Bandwidth Utilization

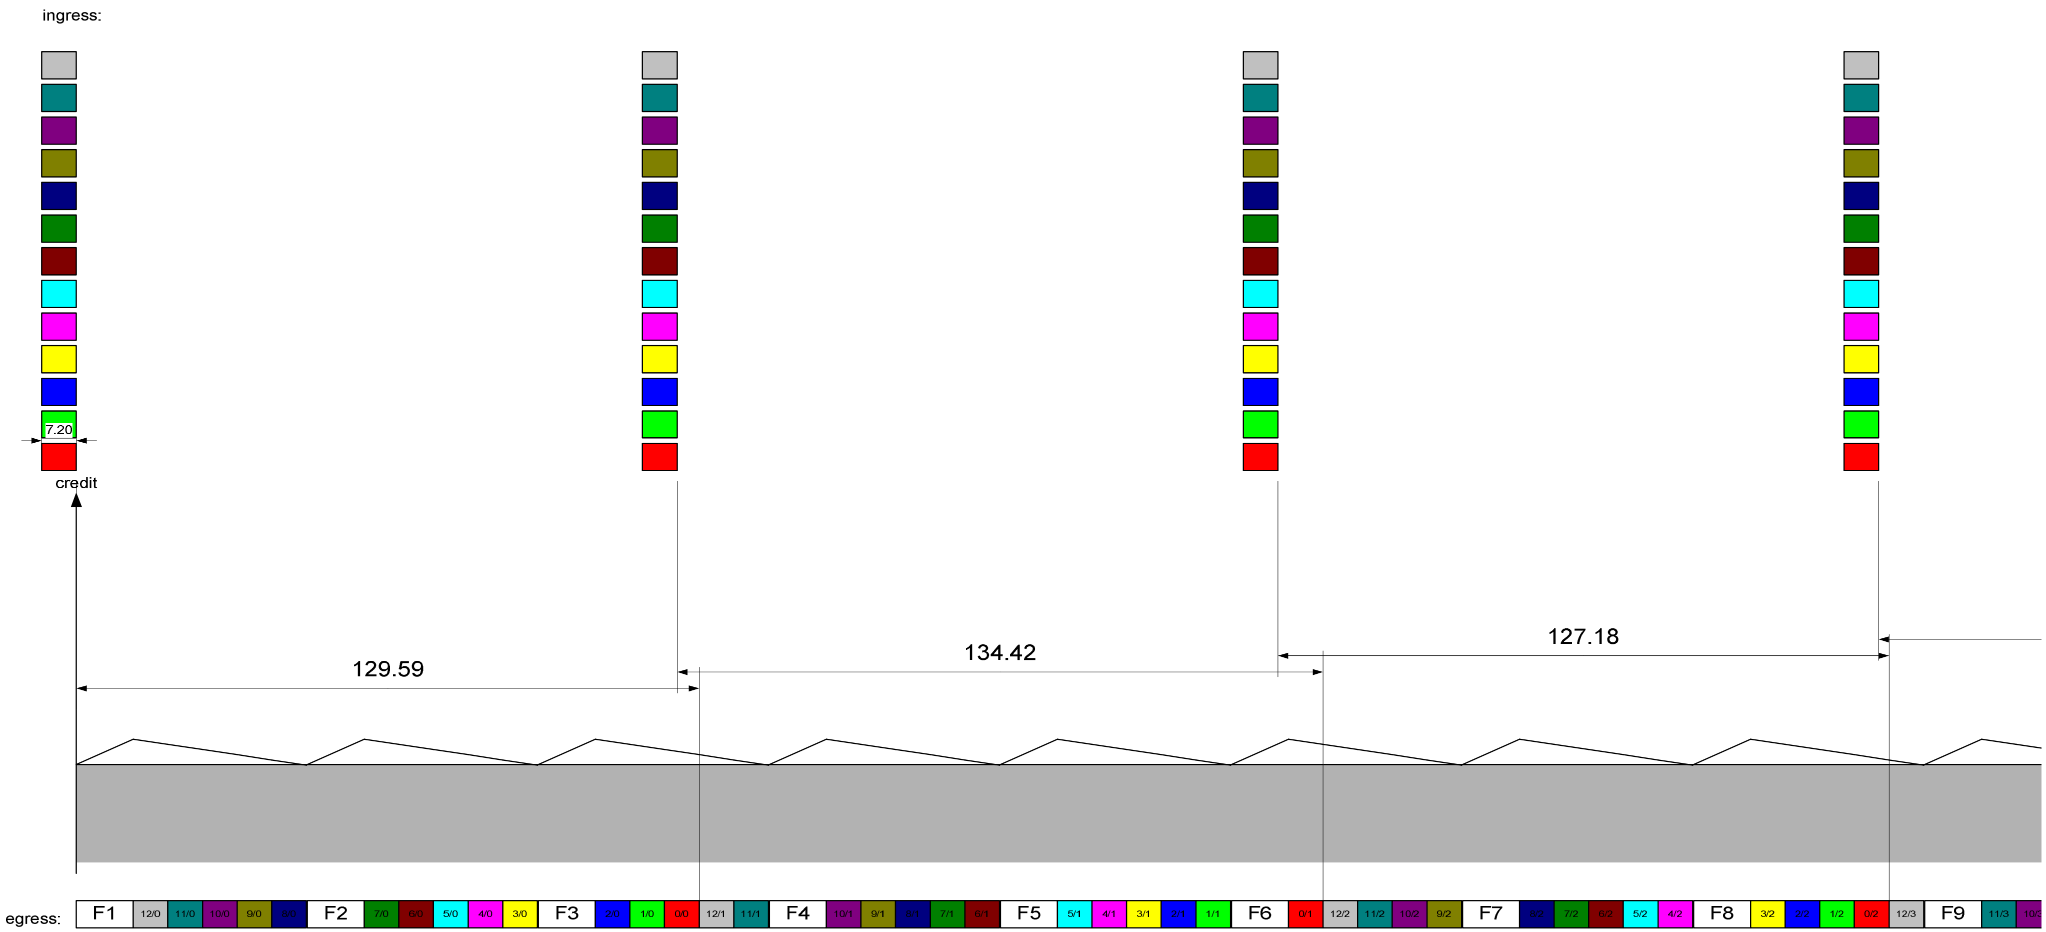

Results

- Fragmentation would decrease FE max class A latency significantly
- Bursts of AVB stream frames of the same stream are possible (consequences according to latency are the same as without fragmentation)
- The shaper produces, at a high bandwidth utilization with many small streams, many fragments (→ insufficient bandwidth use)
- But it might be interesting to define the fragmentation mechanism as an independent mechanism. The shaping mechanism which is used for the high priority traffic can be a time aware shaper, bursting shaper or AVB Gen1 shaper. It only should be guaranteed that the traffic is bandwidth limited.

Gigabit Ethernet Class A

+

128 Byte Fragmentation

Worst Case Talker Latency

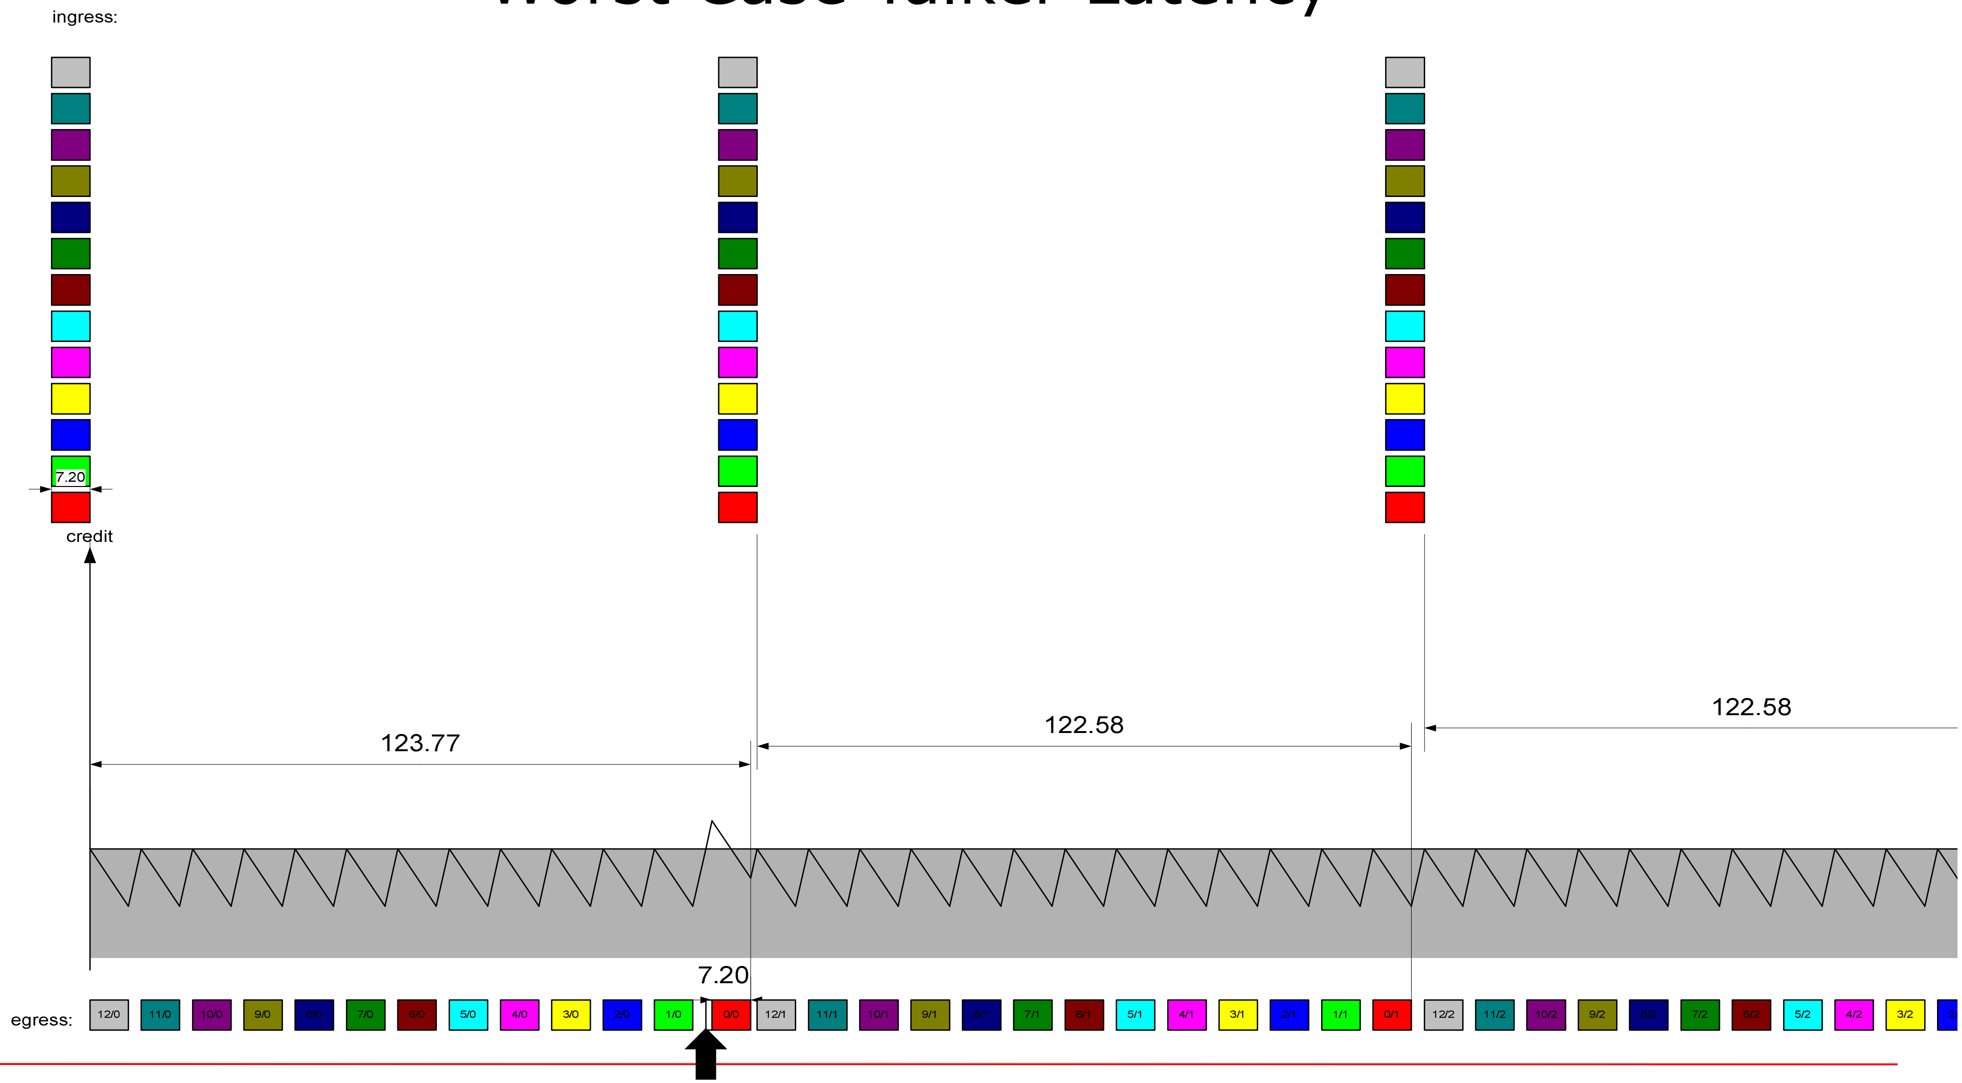

Creation of Burst in First Bridge

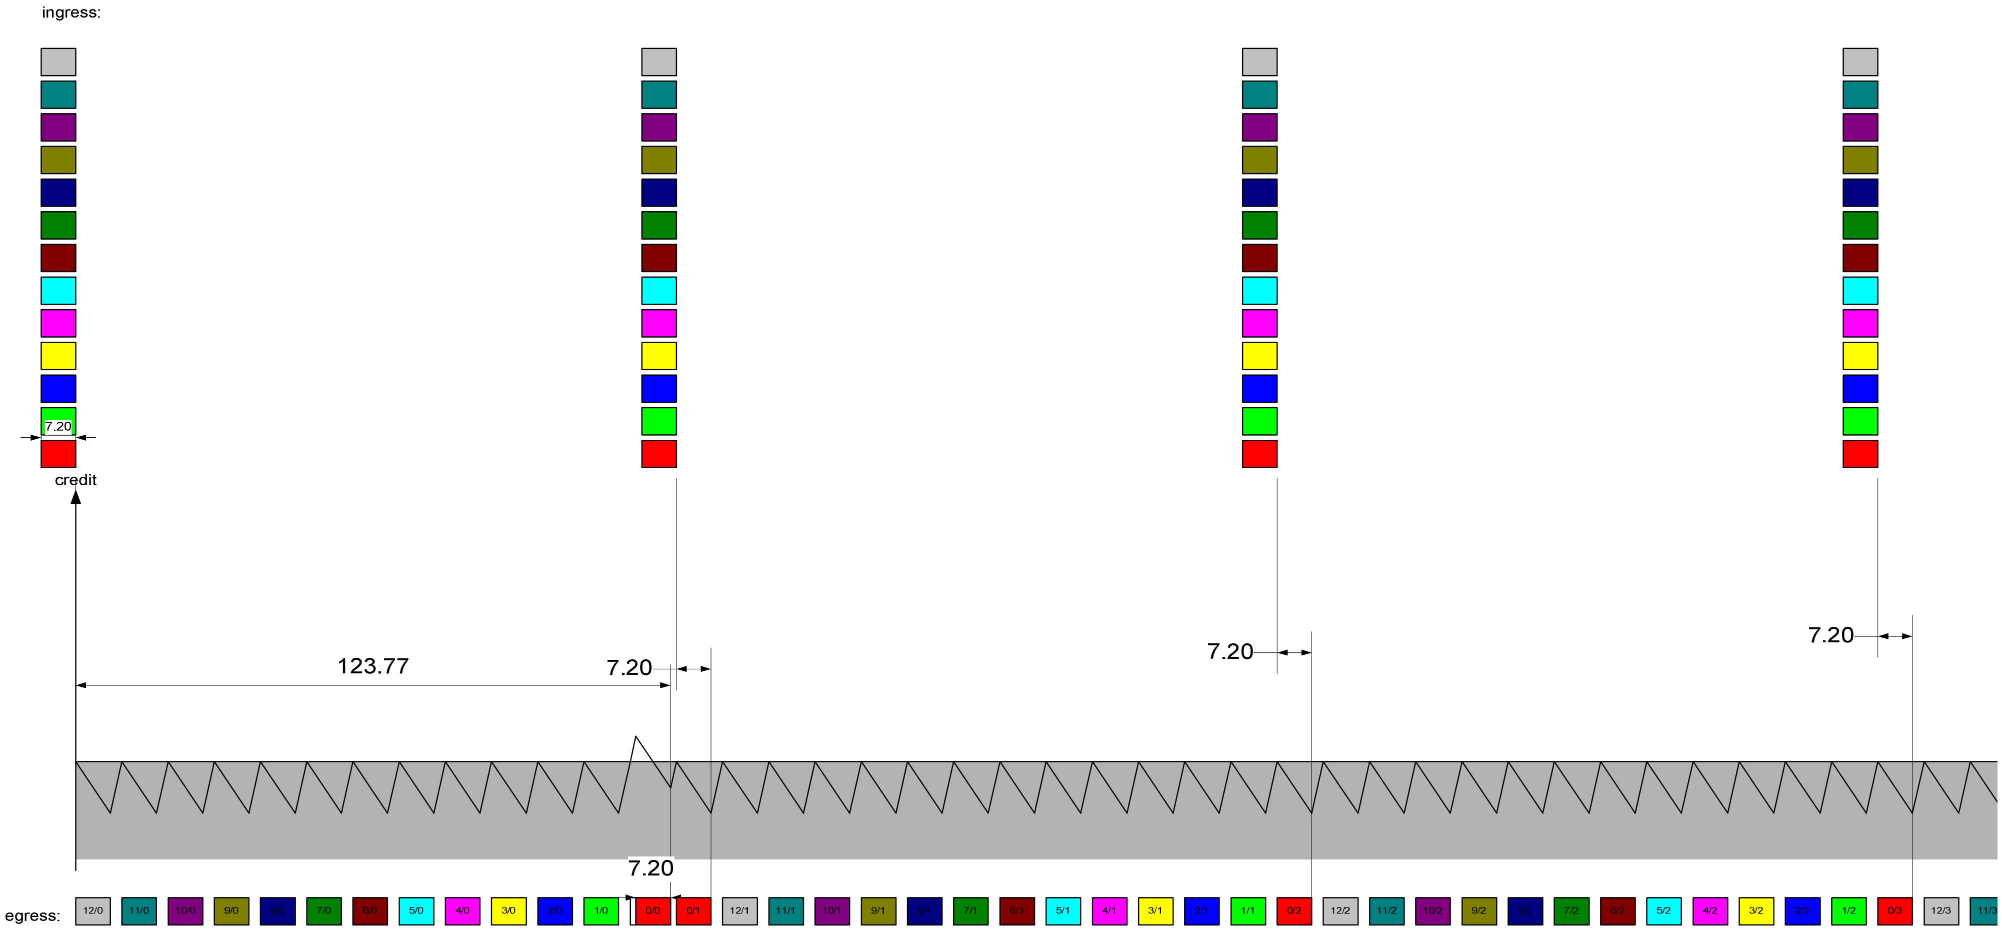

Consequences of a “Bursting” Ingresses

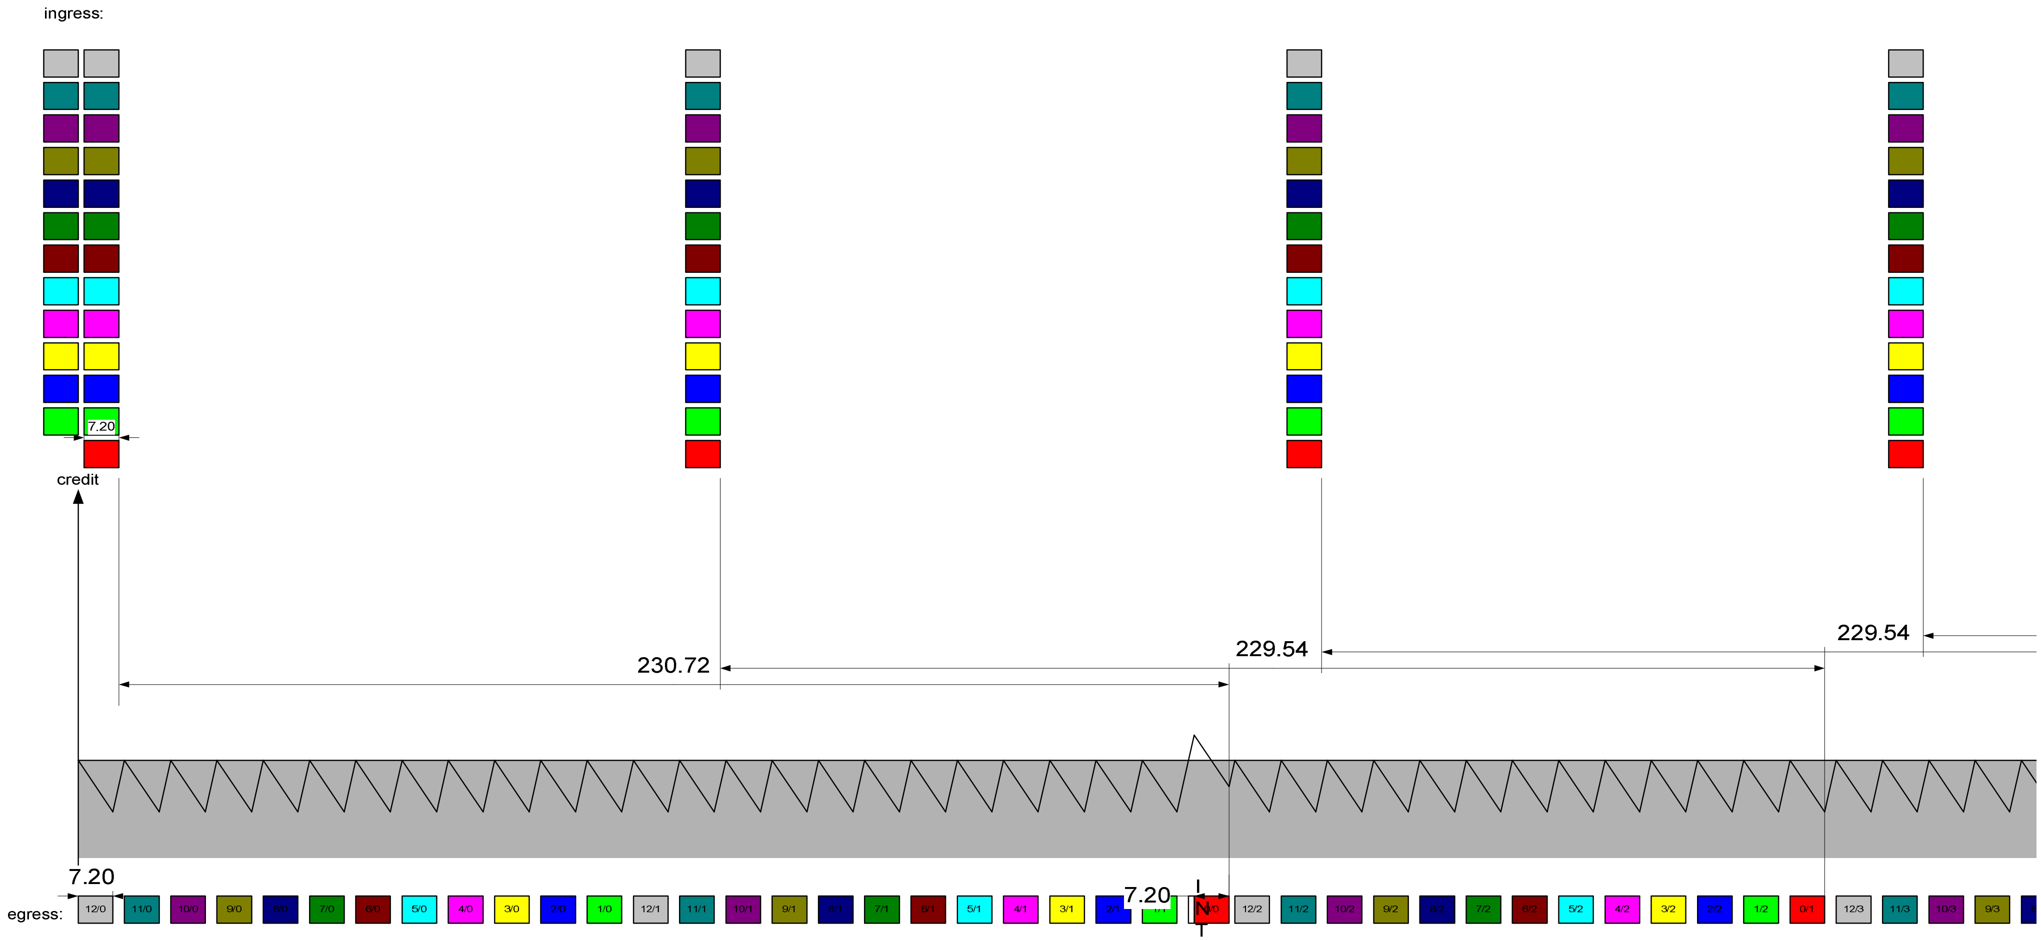

Results

- Fragmentation would decrease max GigE class A latency
- Bursts of AVB stream frames of the same stream are possible (back to back bursts are limited to small frames)
- The shaper produces, at a high bandwidth utilization with many small streams, many fragments (→ insufficient bandwidth use)
- The biggest (almost only) part of latency is a result of interfering frames of the same stream class (interfering legacy frames only cause an additional latency of 1.184µs)

Equation

MaxLatency(Talker)

$= t_{Device} + t_{(MaxPacketSize+IPG)} + (t_{AllStreams} - t_{(MaxStreamPacket+IPG)}) * transmissionRate/maxAllocatableBandwidth + t_{MaxStreamPacket}$

tDevice = the internal delay of the device (in slot times, i.e., increments of 512 bit times)

t(MaxPacketSize+IPG) = the transmission time for a maximum size interfering packet (maximum size interfering packet (156 octets)) plus IPG

tMaxStreamPacket = the transmission time for the maximum packet size of the stream that is being reserved

transmissionRate = transmission rate of the medium

maxAllocatableBandwidth = the maximum amount of Class A stream bandwidth the talker is able to allocate

tAllStreams = (maxAllocatableBandwidth * tInterval)/transmissionRate

tInterval = the Class A observation interval or 125 µs

The Bridge per Hop Worst Case Latency depends on the network topology

Fast Ethernet Bursting Shaper

+

128 Byte Fragmentation

Bursting Shaper Assumptions

- The following slides assume a bursting shaper
- The shaper is realized with the standard AVB Gen1 shaper
- The “base” of the shaper is not 0, i.e. the shaper starts with positive credit and also accumulates credit until it reaches this amount of credit
- The rest of the shaper mechanisms are not changed (no transmission with a negative credit, etc.)
- Max 25 Mbit/s allocatable bandwidth (4 streams, 125µs, 97 Bytes)

Bursting Shaper Without Non-Bursting-Class Interfering Frames

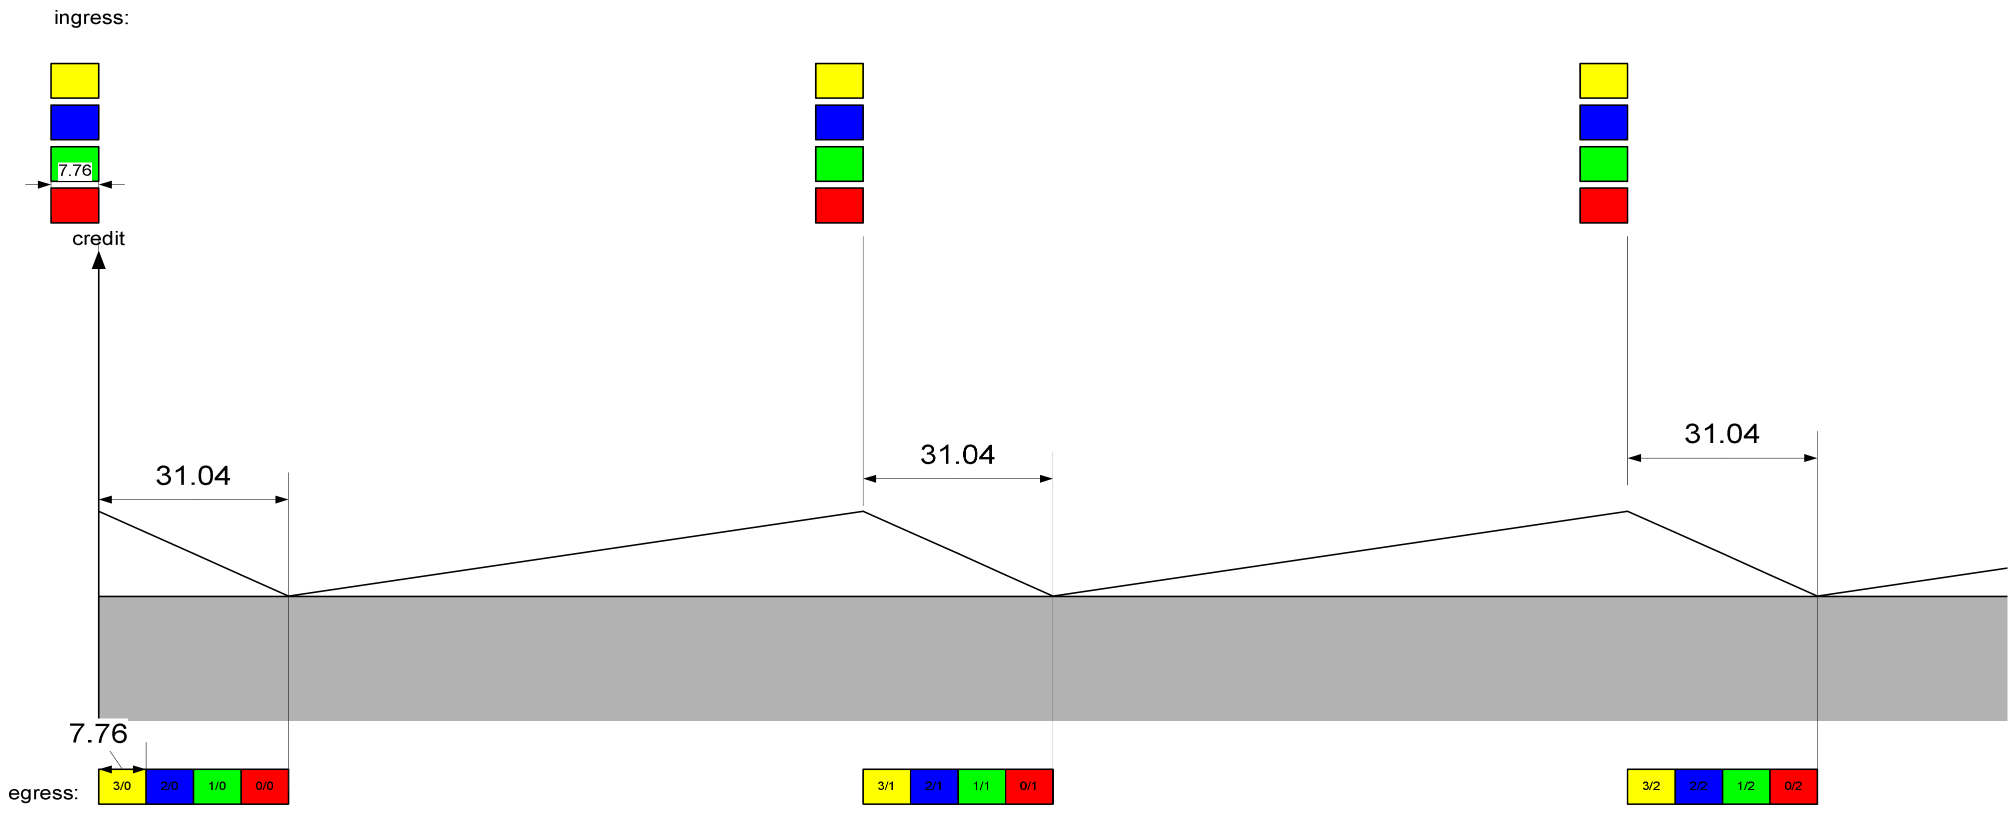

Worst Case Talker Latency (1)

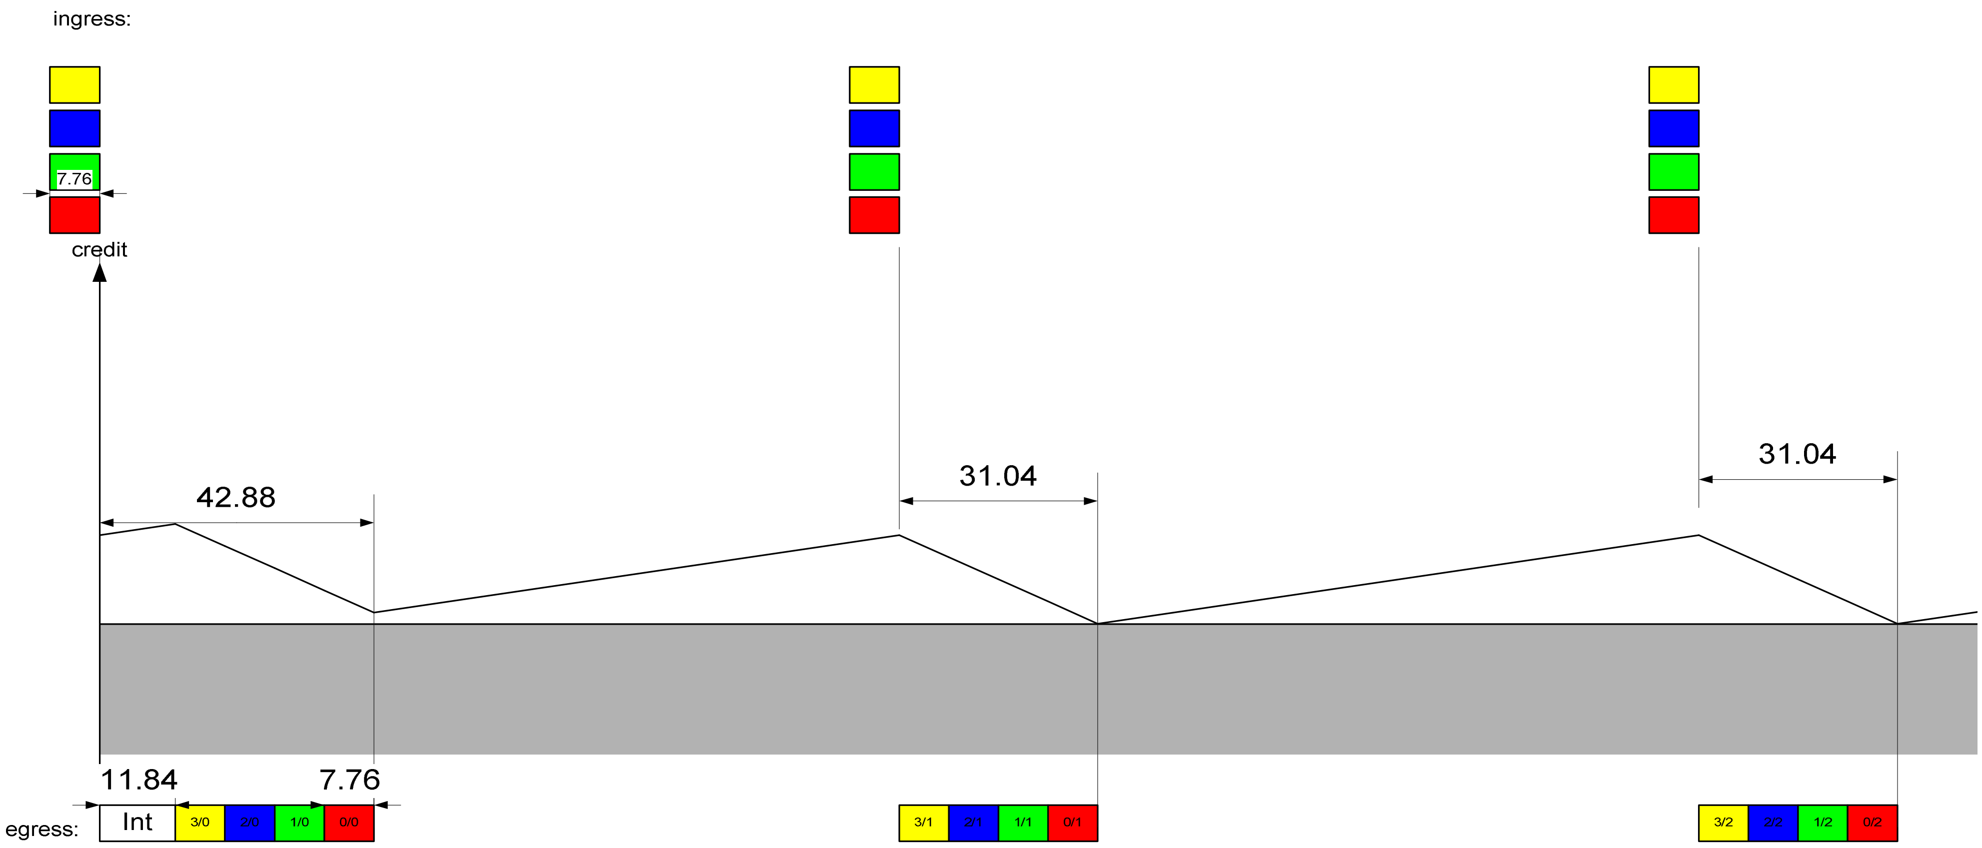

Worst Case Talker Latency (2)

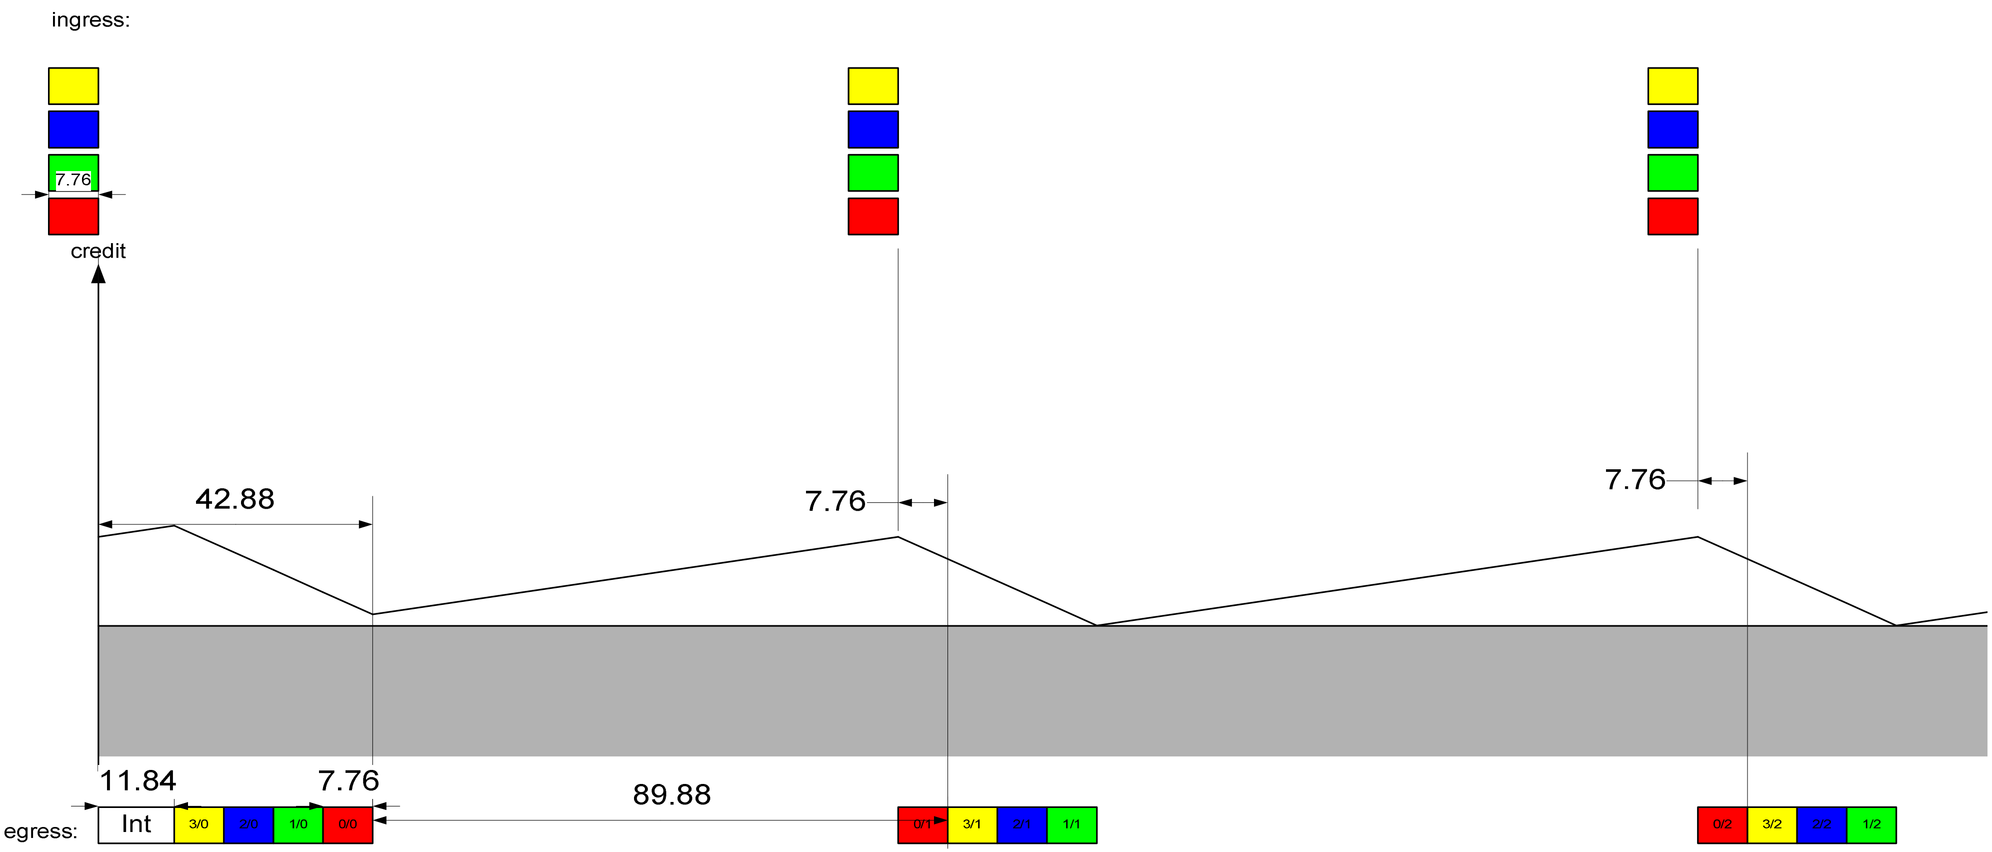

First Bridge After Talker

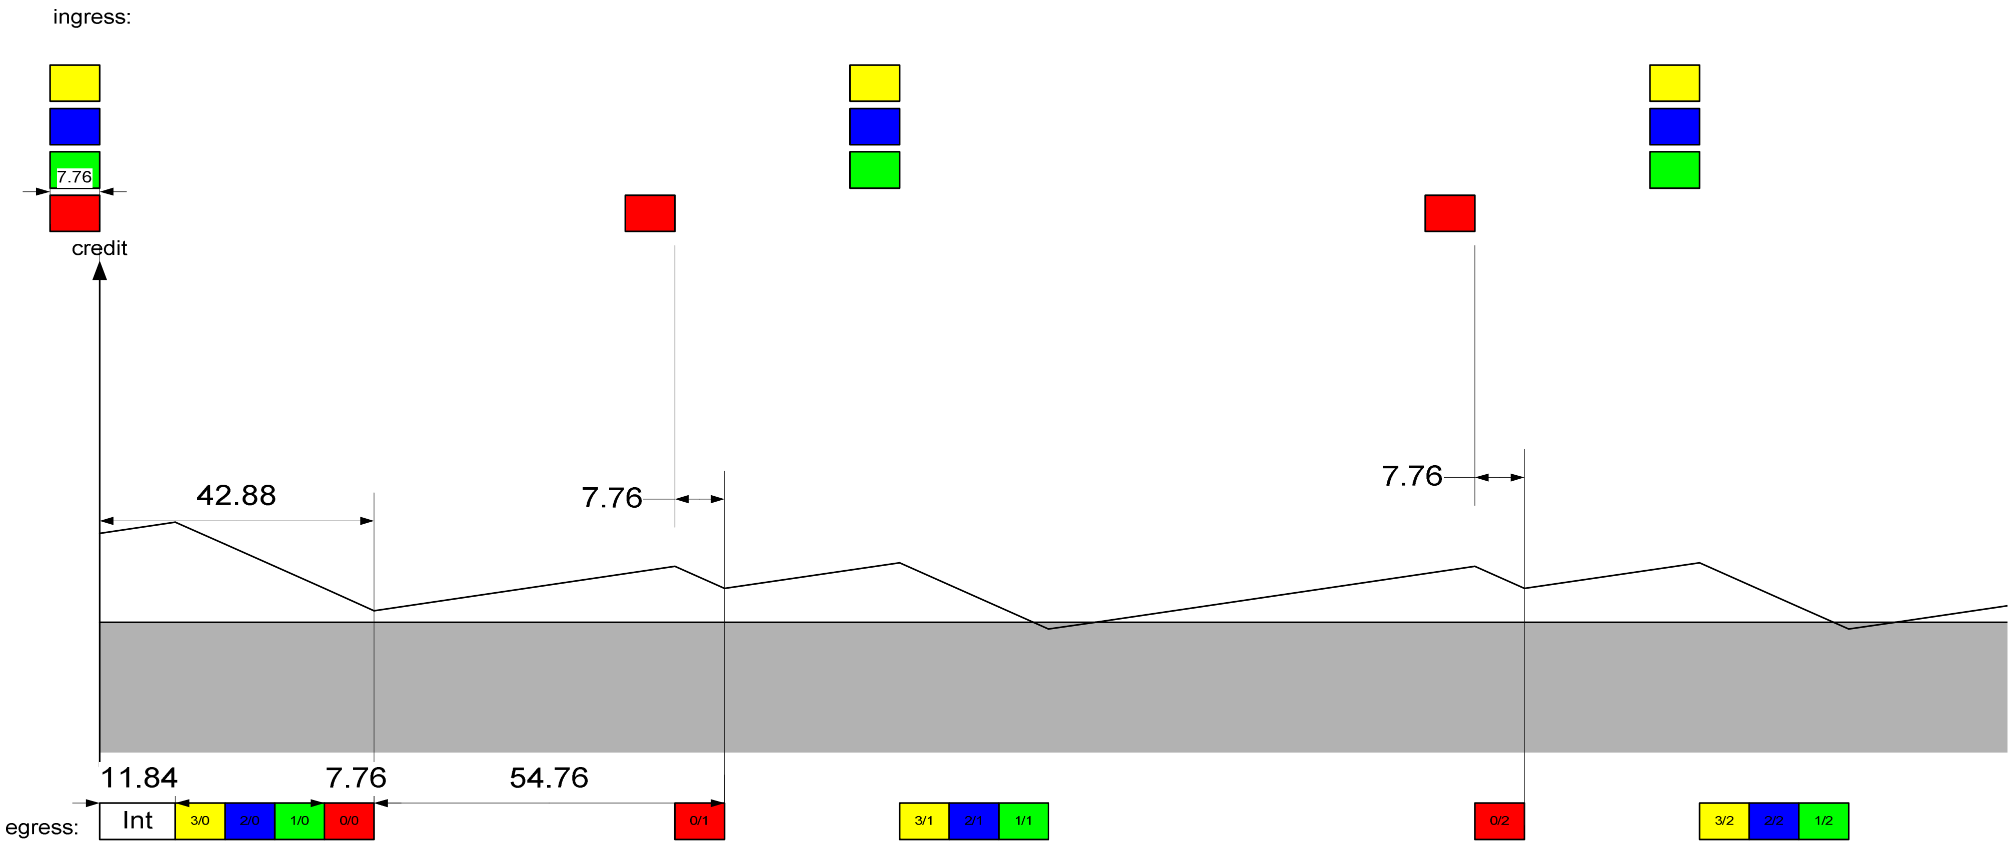

Second Bridge After Talker

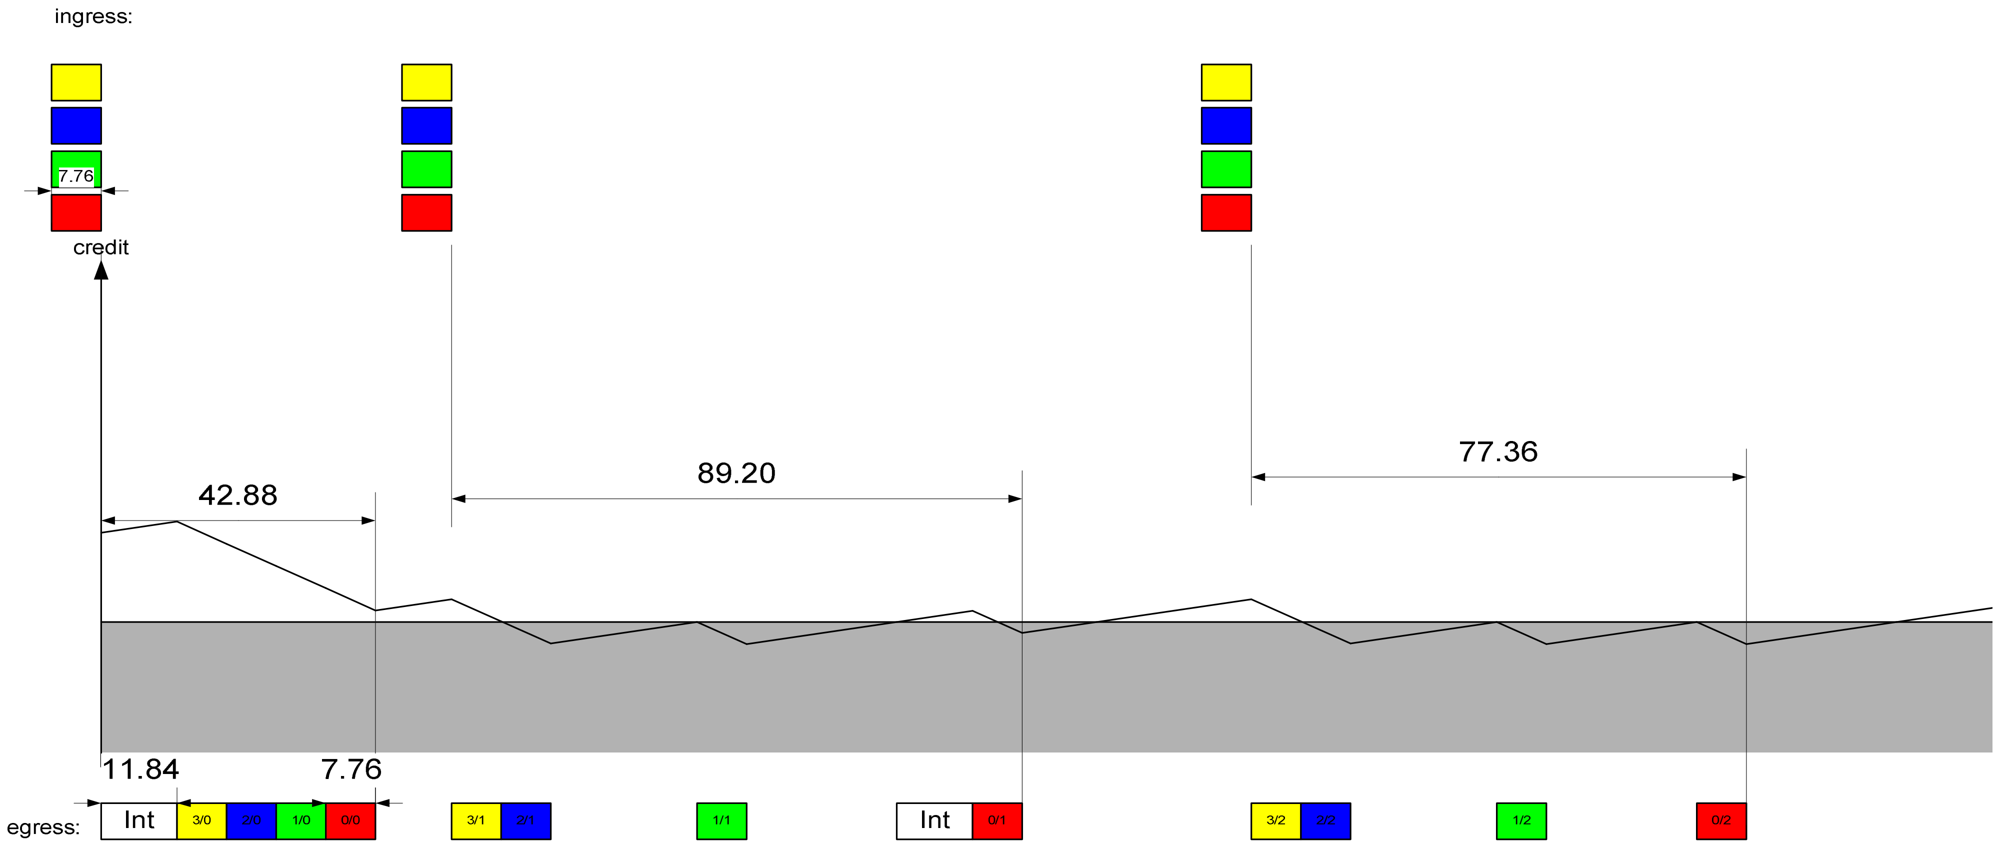

Results

- The biggest part of the latency is a result of interfering high priority frames (i.e. interference within the same class)
- Is it possible that frames of the same stream are getting pushed together to a burst
- The consequences of a burst or not equally spaced frames can only be determined when the shaping mechanism is defined (in the case of a shaper with a "positive base" the stream frames after a burst are getting shaped)
- Without any bursting ingress (of the same stream) the latency for a 97 byte stream incl. a tDevice of 5.12µs is **47.04µs** (assuming a max class AAA bandwidth of 25% and 128 byte fragmentation)

Gigabit Ethernet Bursting Shaper

+

128 Byte Fragmentation

Bursting Shaper Without "Low Priority" Interfering Frames

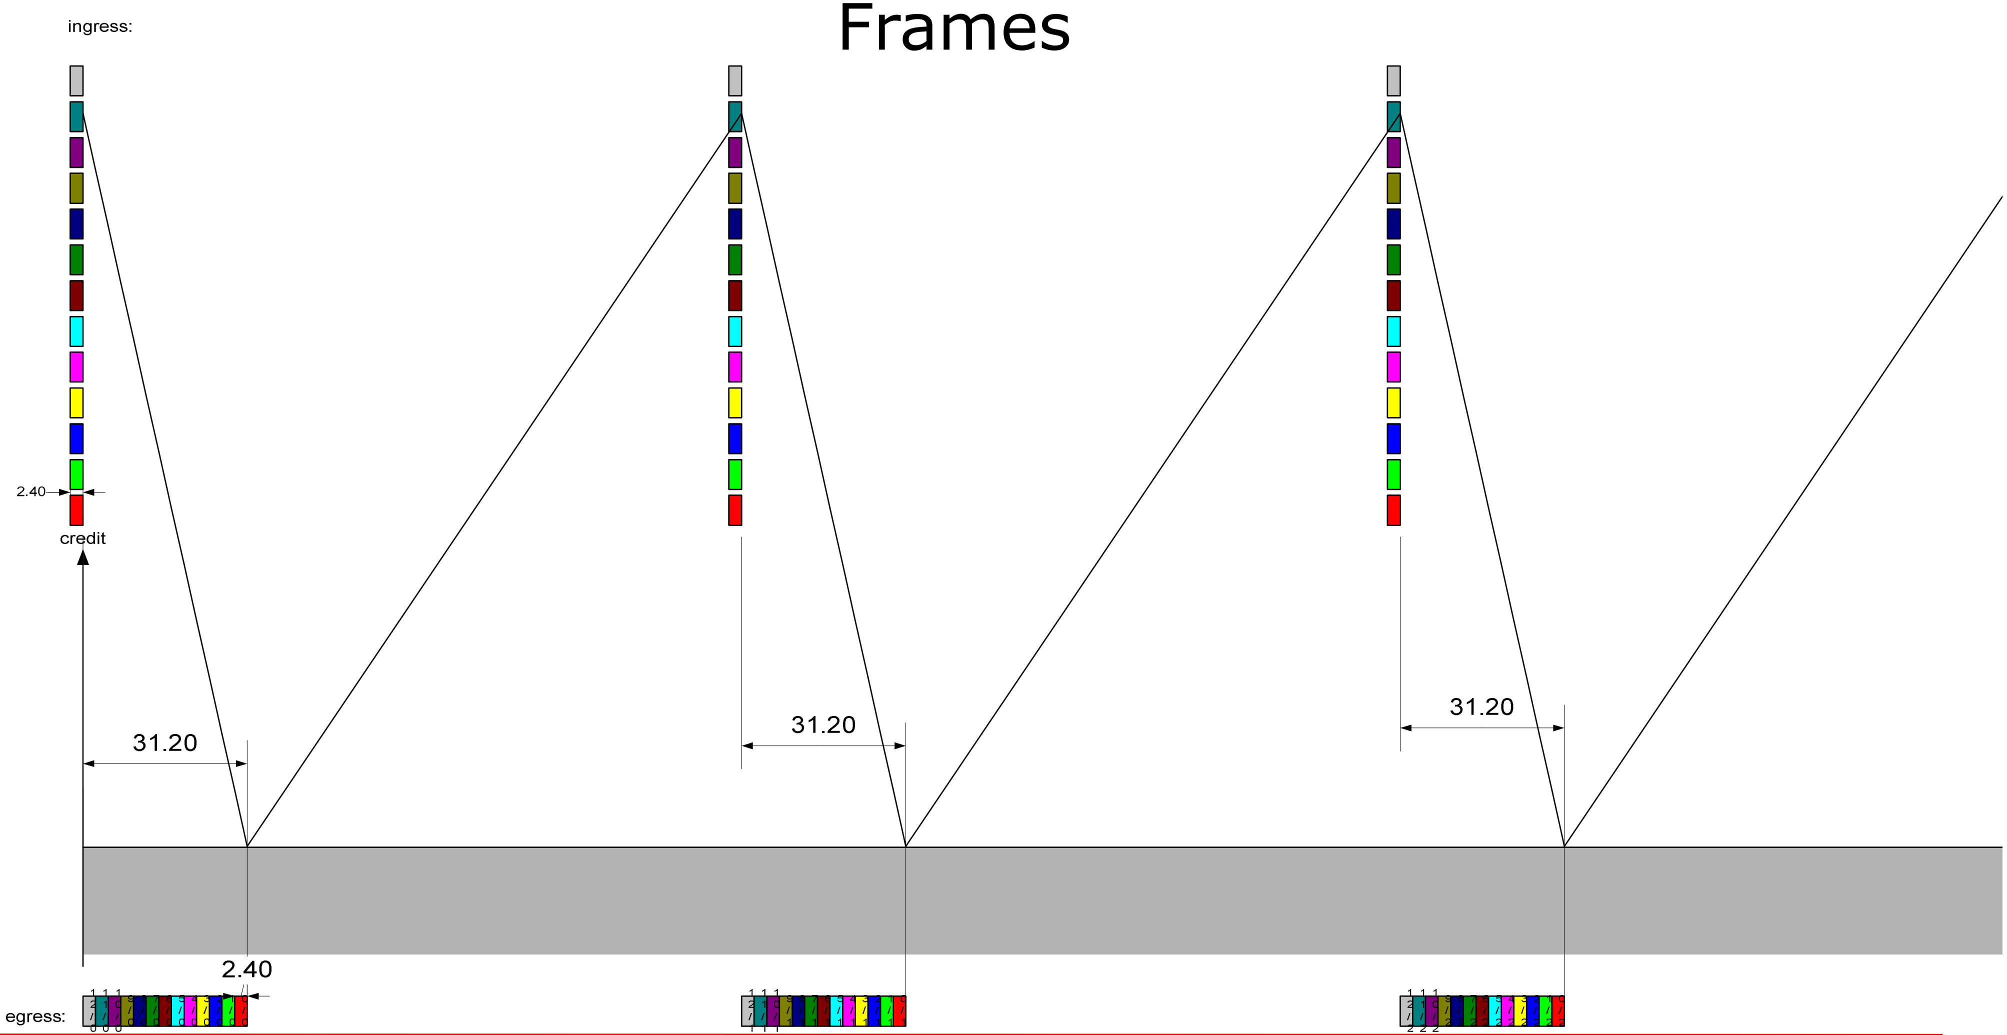

Worst Case Talker Latency

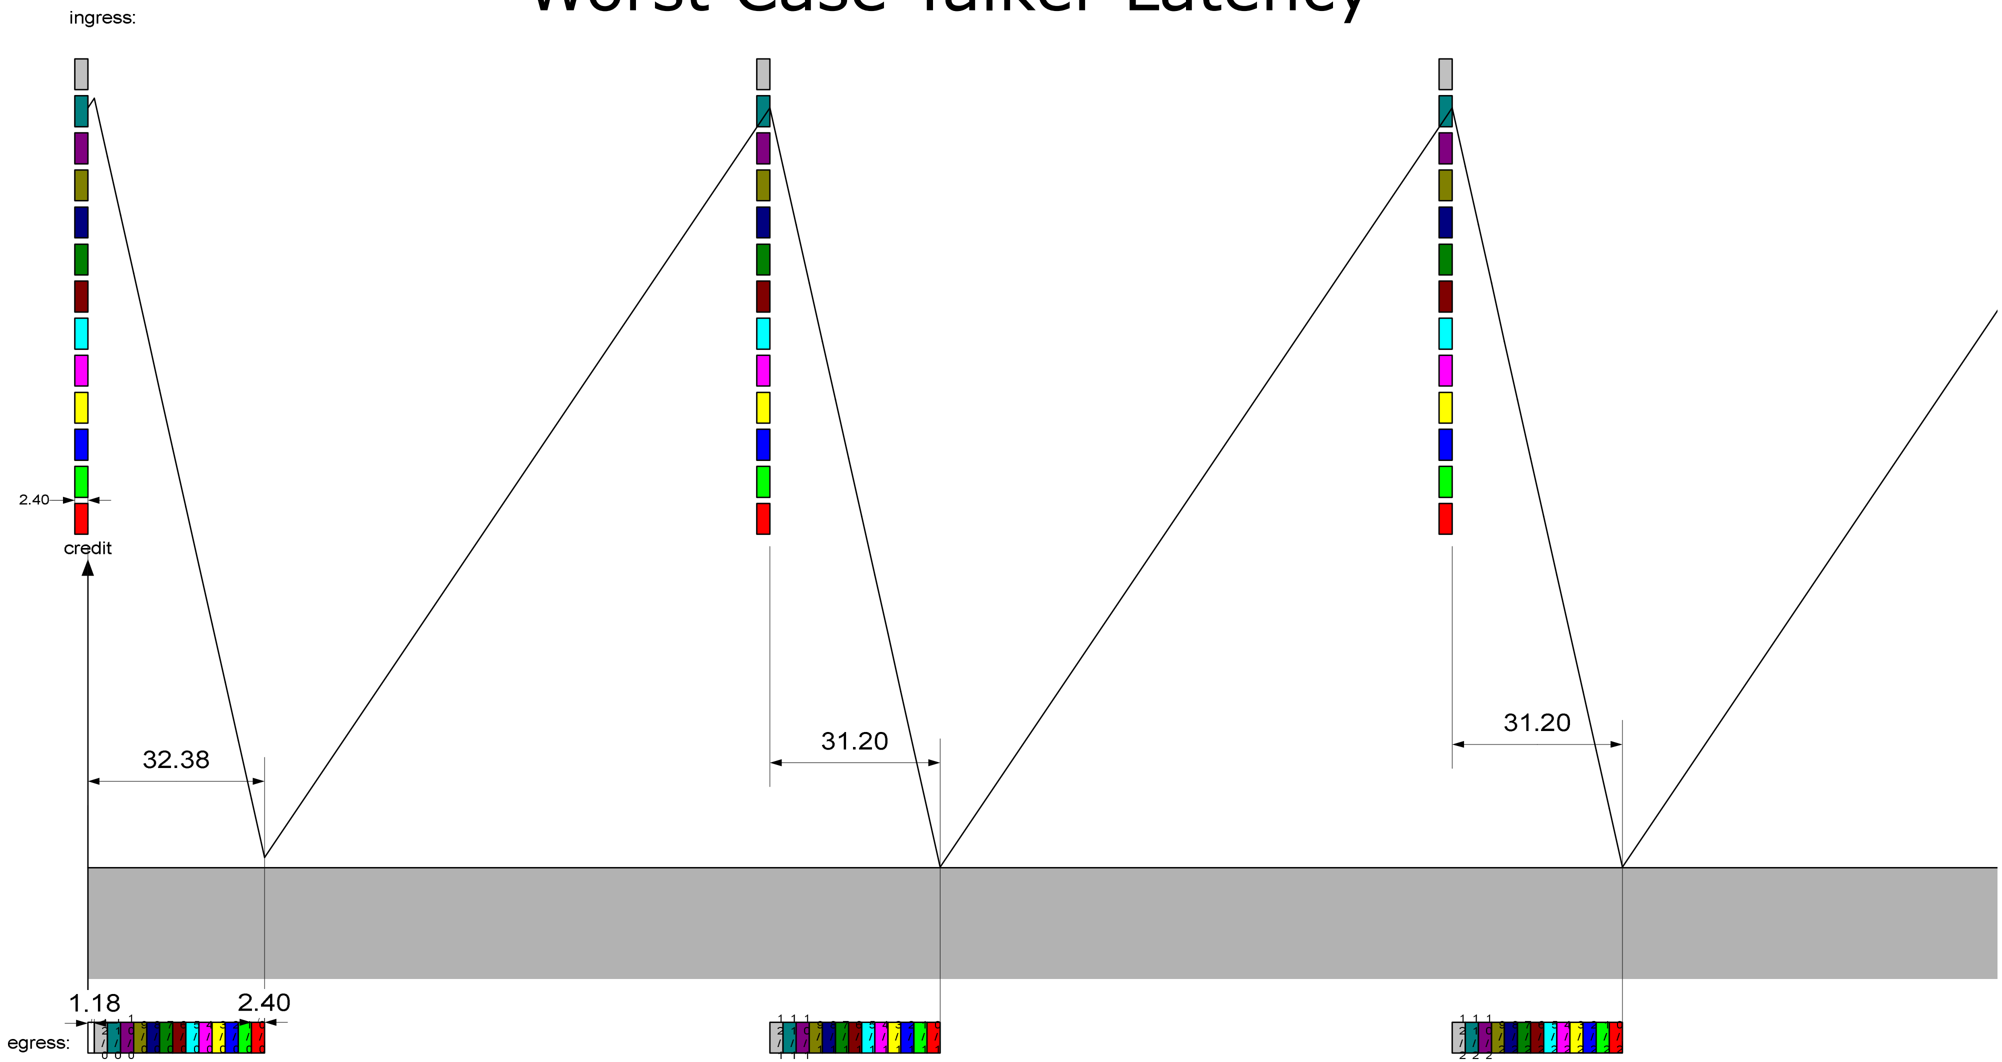

Pushing Stream Frames Together

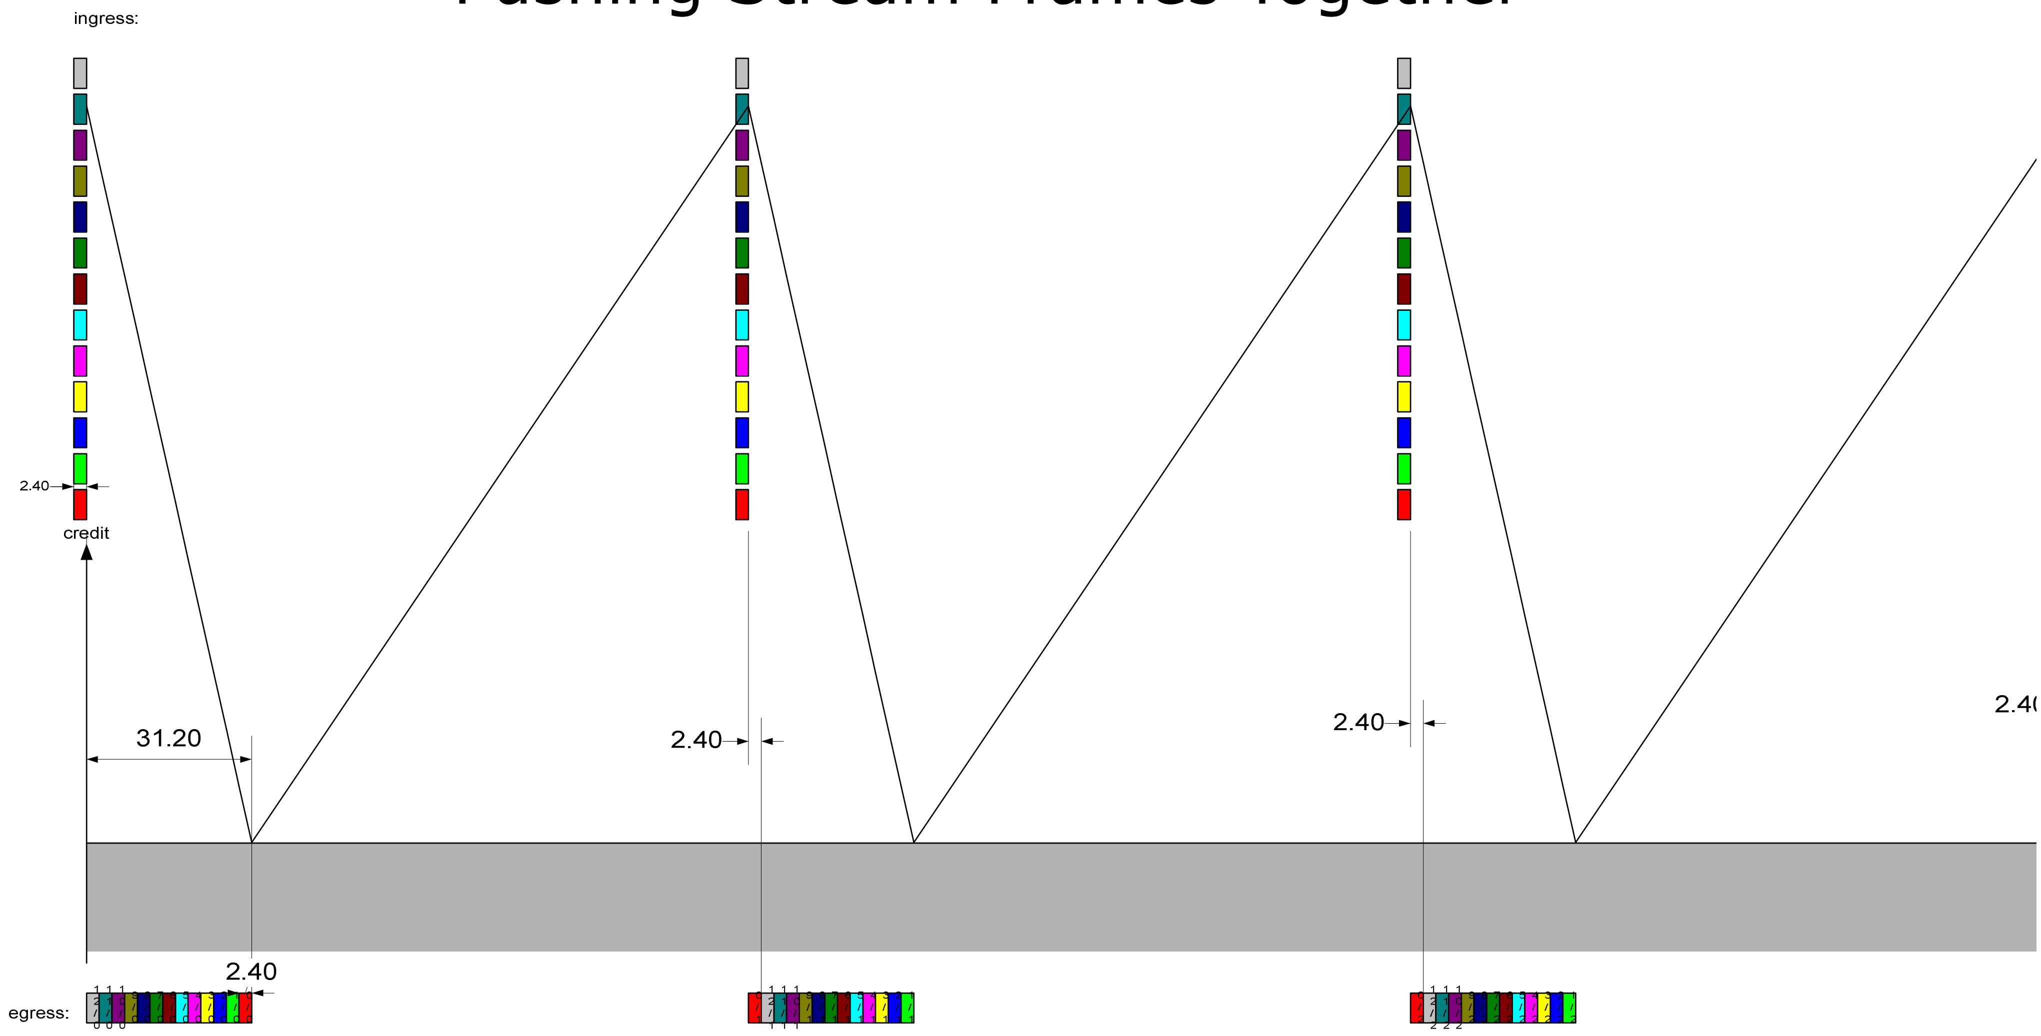

First Bridge After Talker

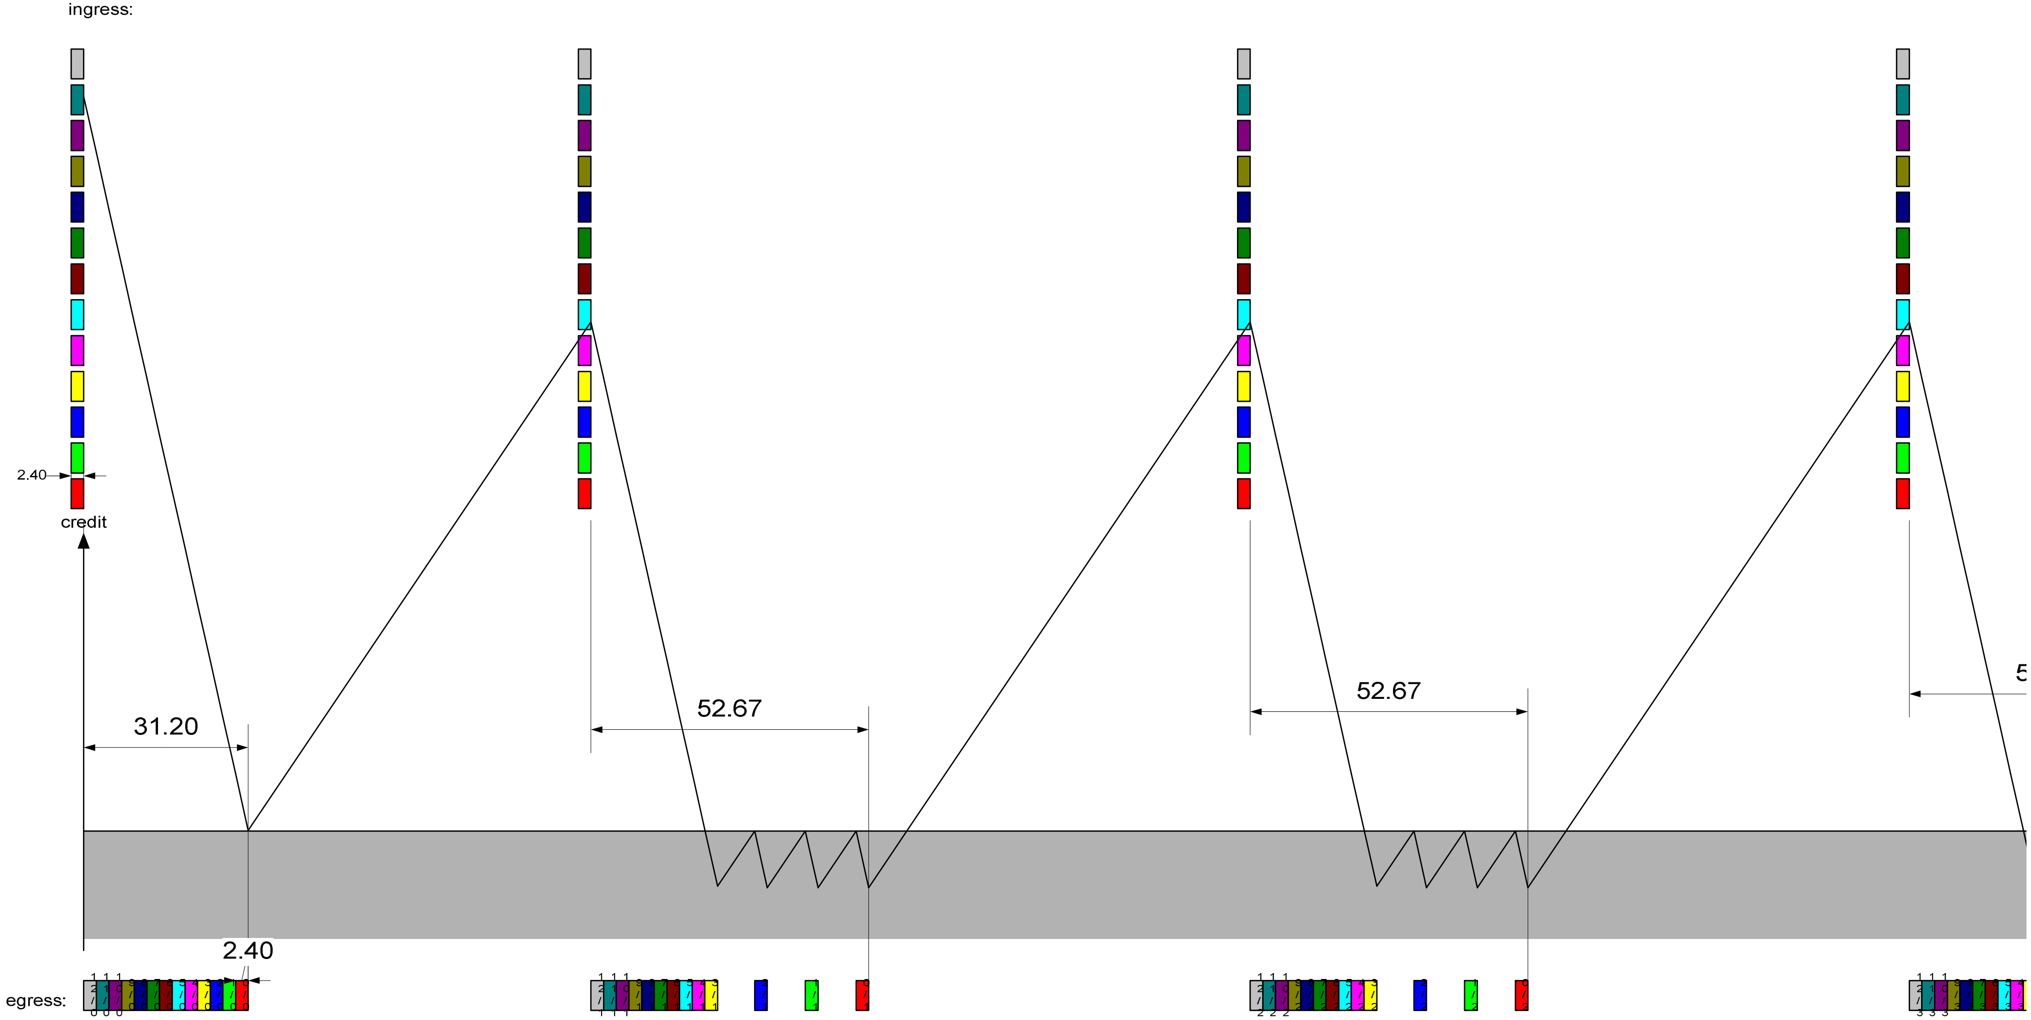

Equation

MaxLatency(Talker)

$= t_{Device} + t_{(MaxPacketSize+IPG)} + (t_{AllStreams} - t_{(MaxStreamPacket+IPG)}) + t_{MaxStreamPacket}$

tDevice = the internal delay of the device (in slot times, i.e., increments of 512 bit times)

t(MaxPacketSize+IPG) = the transmission time for a maximum size interfering packet (maximum size interfering packet (156 octets)) plus IPG

tMaxStreamPacket = the transmission time for the maximum packet size of the stream that is being reserved

transmissionRate = transmission rate of the medium

maxAllocatableBandwidth = the maximum amount of Class A stream bandwidth the talker is able to allocate

tAllStreams = (maxAllocatableBandwidth * tInterval)/transmissionRate

tInterval = the Class A observation interval or 125 µs

The Bridge per Hop Worst Case Latency depends on the network topology

Fast Ethernet Time Aware Shaper

Time Aware Shaper – 1 Slot For All Stream

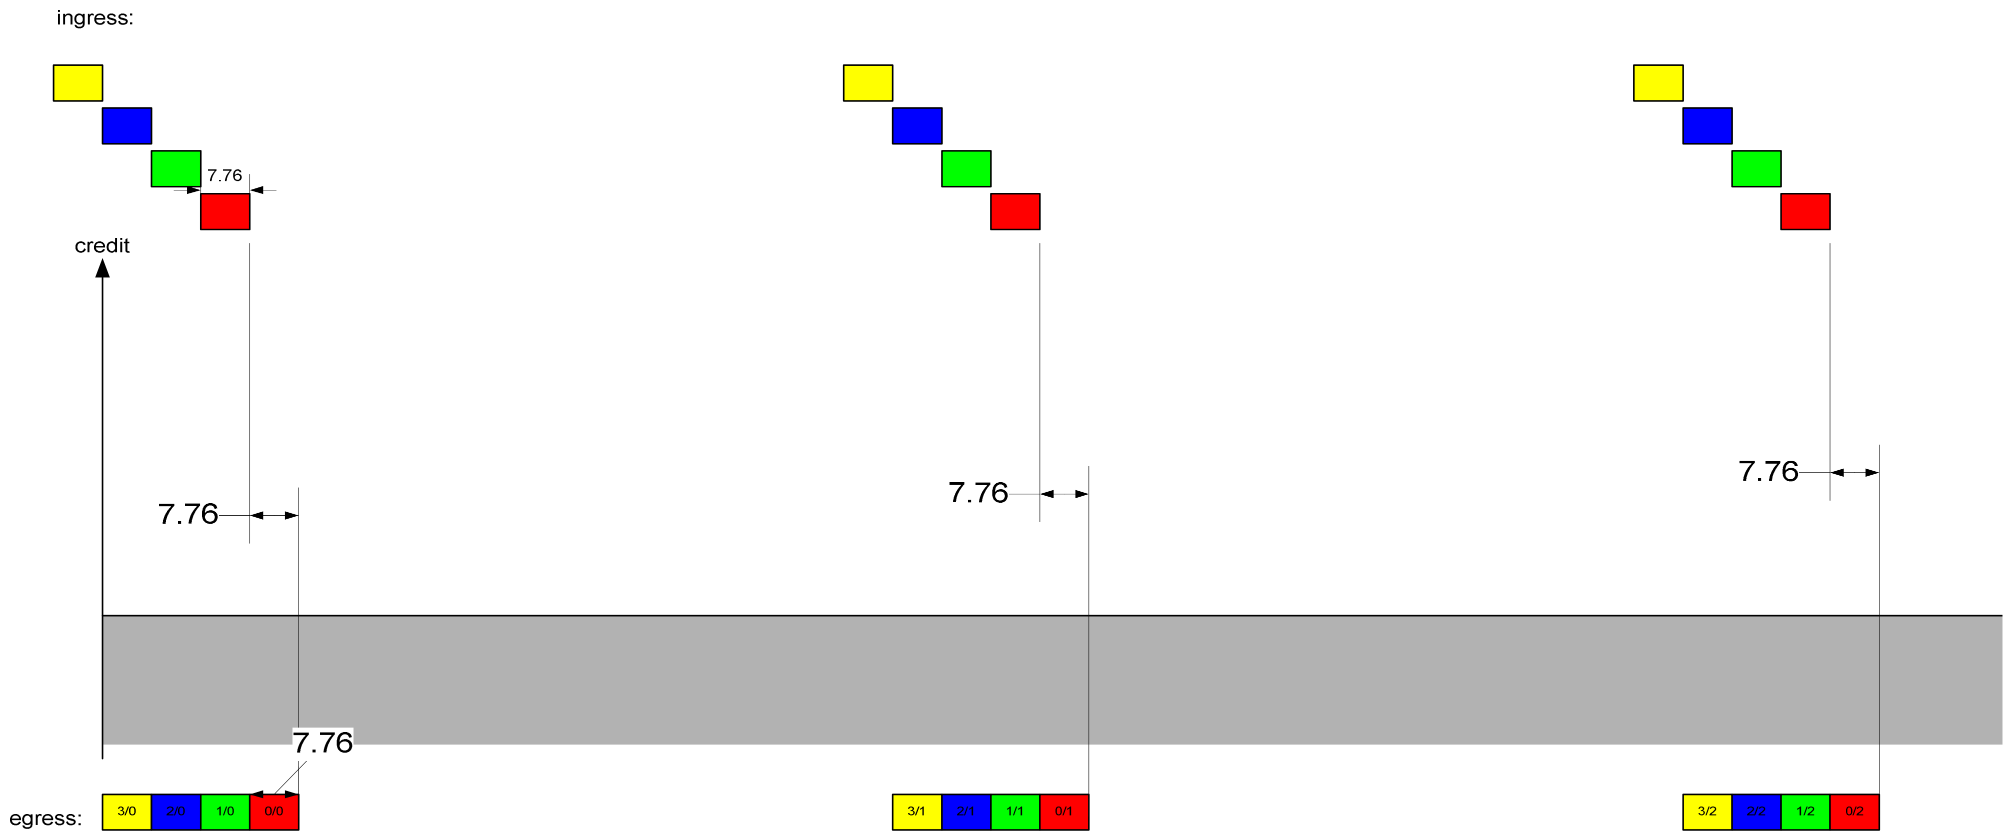

Time Aware Shaper – 1 Slot Per Stream

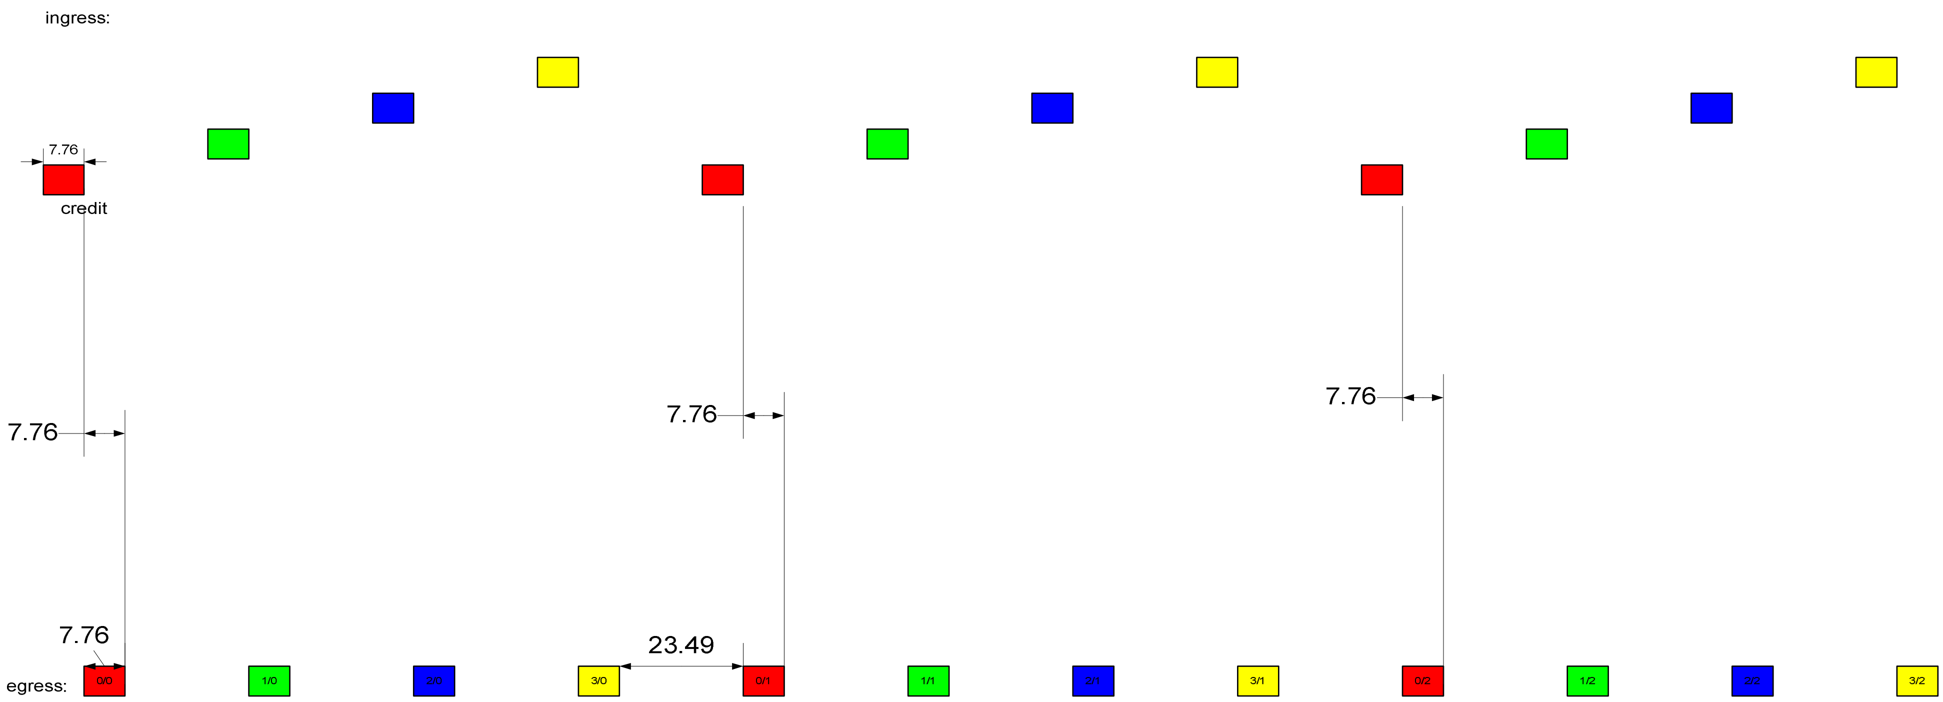

Gigabit Ethernet Time Aware Shaper

Hochschule
Deggendorf

Talker/Bridge Latency

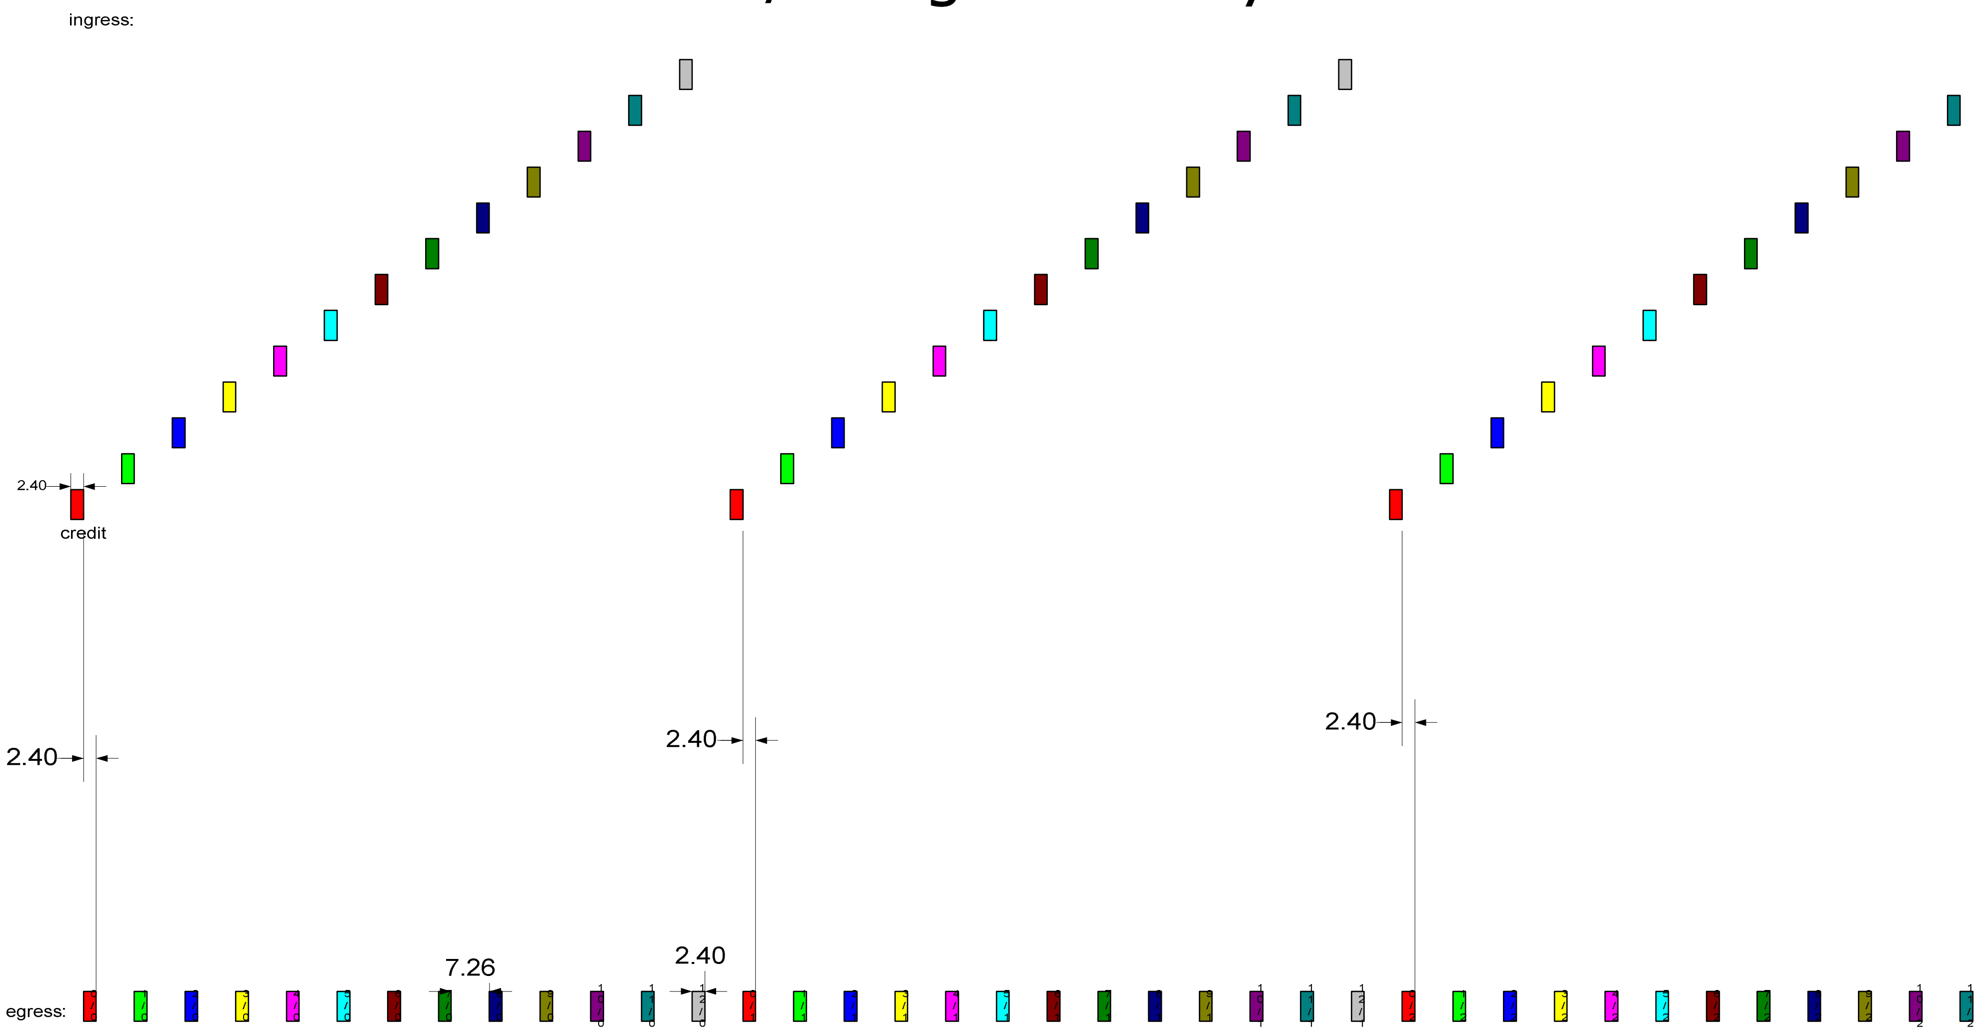

Results

- The latency only consists of the store and forward delay and the internal delay of the device
- The space between the high priority transmissions has to be big enough to transmit a min size fragment
- Bursts of frames of the same stream are not possible as the frames are equally spaced during the whole transmission
- The one slot for all streams approch decreases the available bandwidth for streams and legacy traffic significantly
- The per hop bridge latency does **NOT** depend on the network topology

Equation

$$MaxLatency(Talker) = t_{Device} + t_{MaxStreamPacket}$$

tDevice = the internal delay of the device (in slot times, i.e., increments of 512 bit times)
tMaxStreamPacket = the transmission time for the maximum packet size of the stream that is being reserved

$$MaxLatency(Bridge) = t_{Device} + t_{MaxStreamPacket}$$

tDevice = the internal delay of the device (in slot times, i.e., increments of 512 bit times)
tMaxStreamPacket = the transmission time for the maximum packet size of the stream that is being reserved

Further Latency Effects (1)

In some cases it is possible that the gaps between the high priority streams are not big enough to transmit a min size fragment:

- One can delay the high priority frames as long as it is necessary to transmit a min size packet in between of the
- One can transmit the high priority frame immediately, but in this case the bandwidth between the two frames is lost (possible as long as there are enough slots for legacy traffic)

Further Latency Effects (2)

In a bridge the slots for the transmission of high priority streams may be blocked by another high priority stream:

- The high priority stream may be delayed or
- The reservation of the stream fails

Both effects may add additional latency. But the worst case is still easy computable e.g. with SRP Gen2

Thank You