Per Hop Worst Case Class A Latency

Deggendorf
University

Christian Boiger
christian.boiger@fh-deggendorf.de
IEEE 802 Plenary Meeting
March 2011
Singapore

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 mentioned number of bytes are incl. preamble SOF and IFG
- The numbers in the graphics are µs
- @FE 7.2µs (90 bytes), 123.36µs (1542 bytes), 86.4µs (1080 bytes)
- @GE 1.92µs (240 bytes), 7.2µs (900 bytes), 12.36µs (1542 bytes)

Worst Case
Burst Size

Bursting Talker (Worst Case)

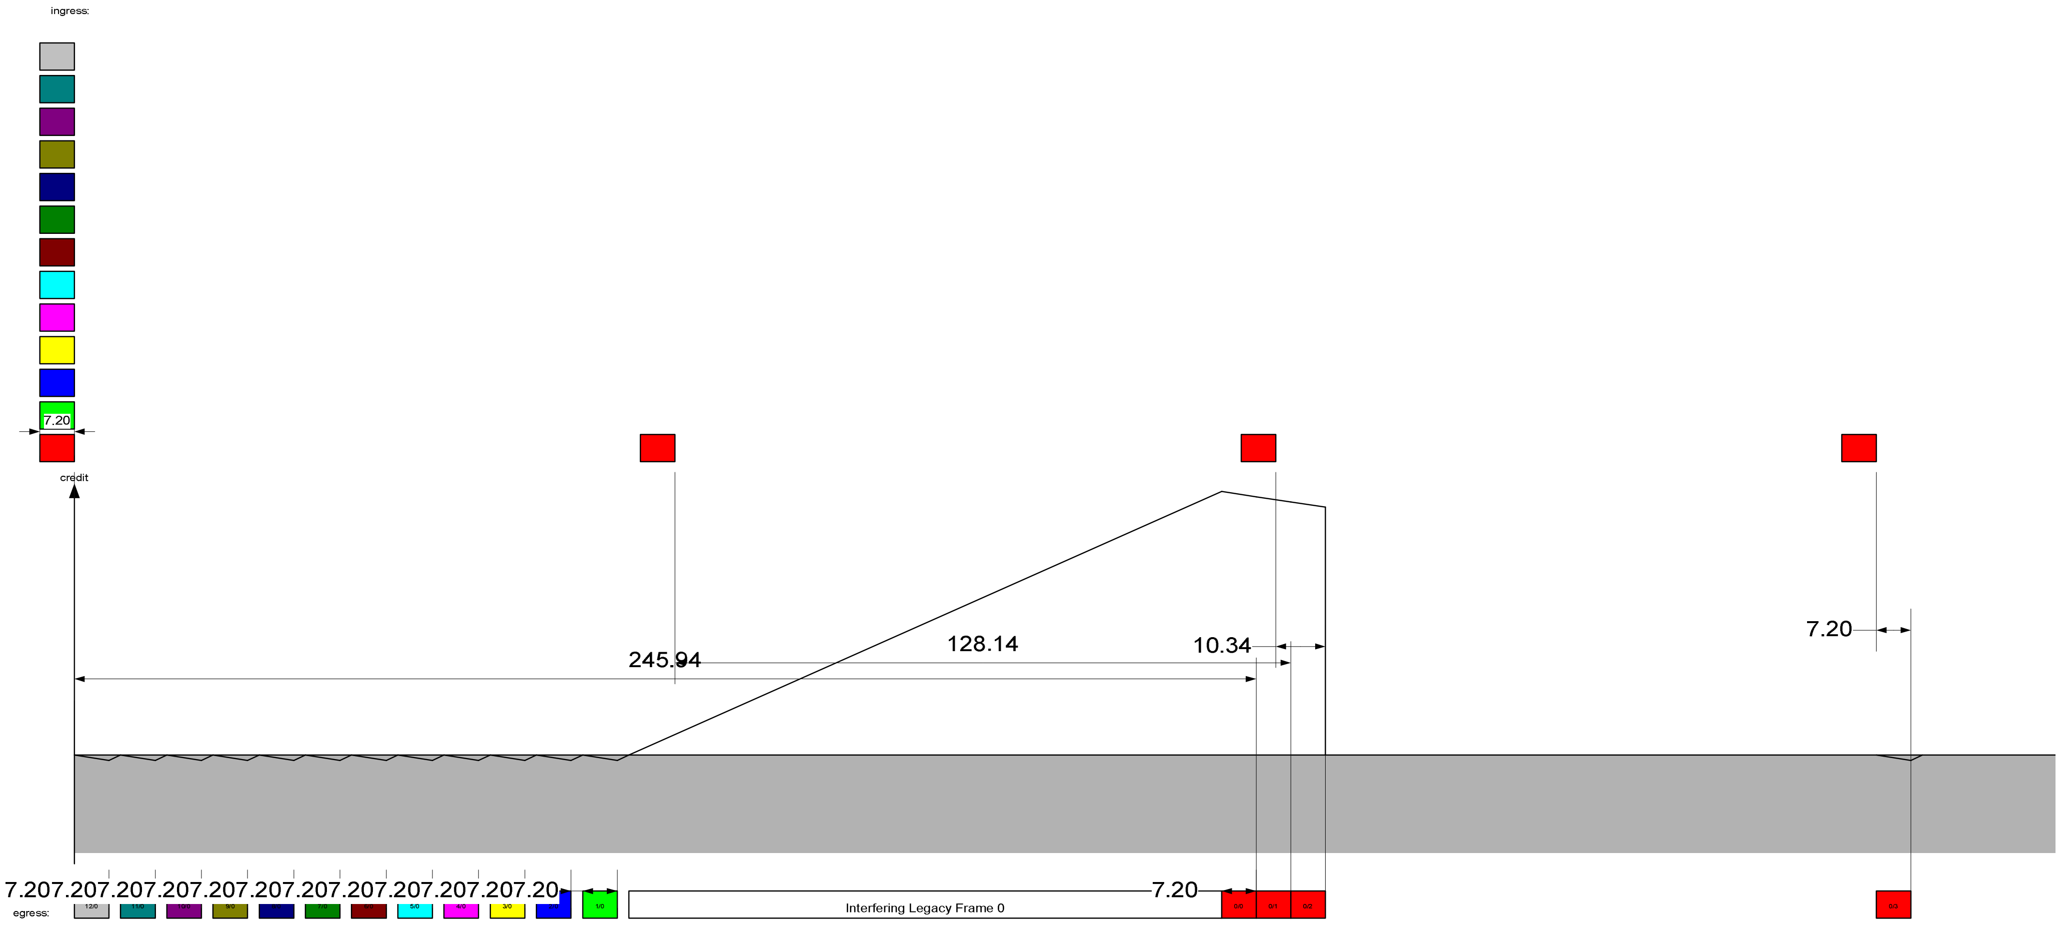

Hochschule
Deggendorf

Bridge After 13 Worst Case Bursting Talker

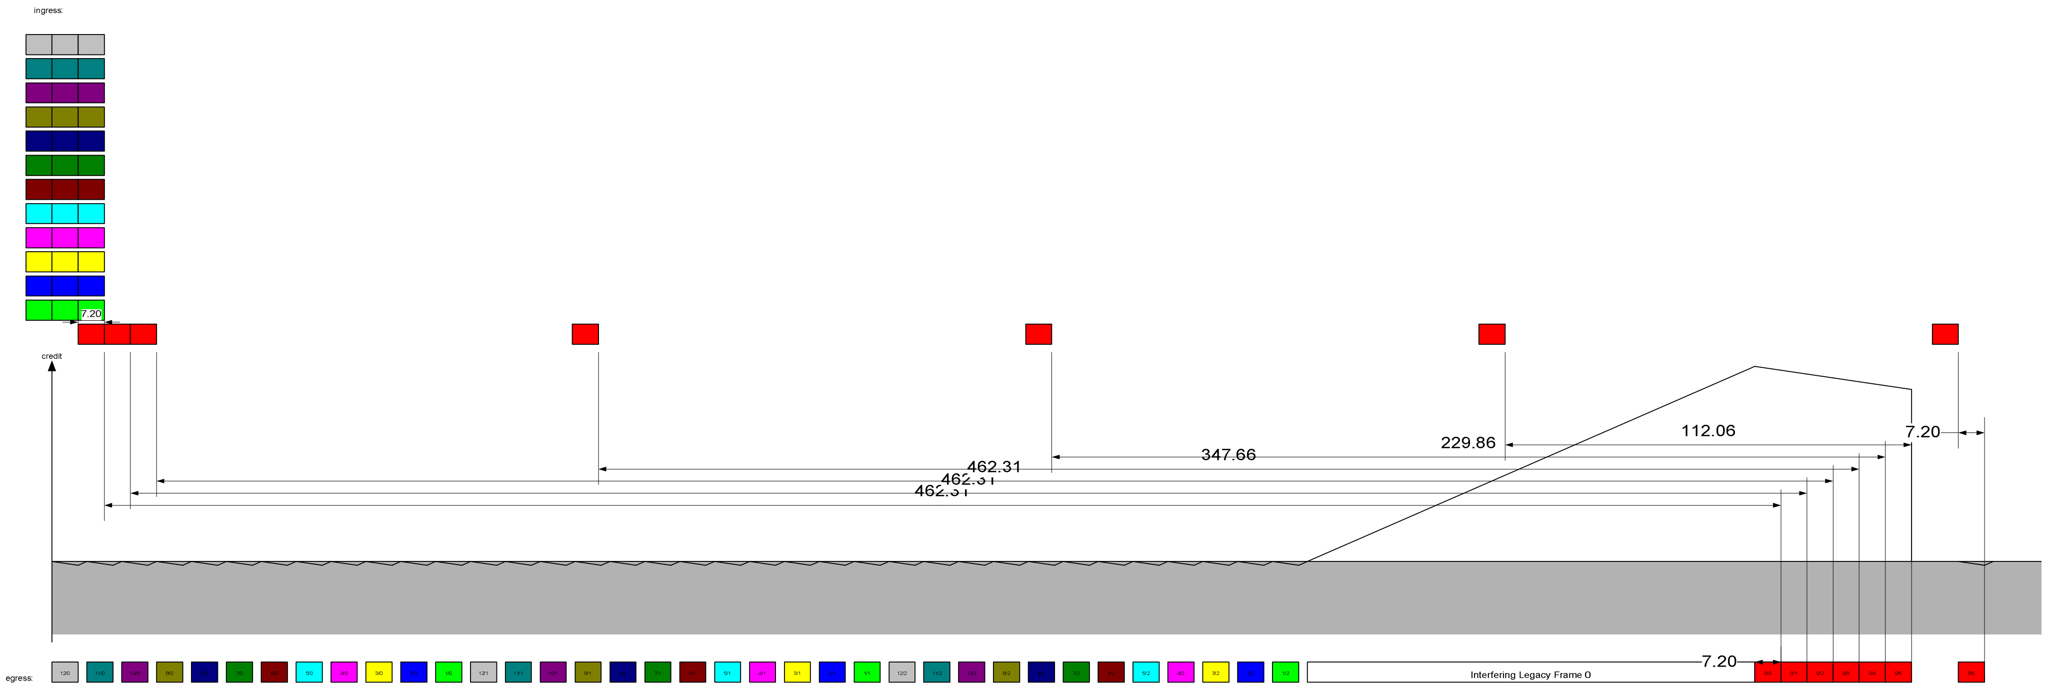

Results

- The burst sizes shown in my last presentations (and used in this presentation (following slides)) are not the worst case. If the interfering class A talkers only cause a delay to the first packet of the burst and then stop their transmission (or pause) bigger bursts are possible.

- The max size of the bursts correspond to the max latency of a frame at this port (e.g. if the max latency is 250µs (i.e. two measurement intervals), altogether three packets can get pushed together)

- As max latency depends on the interfering traffic, max burst is bigger when the interfering traffic contains bursts (higher latency)

Results

MaxBurst = min(floor(MaxCreditBurst/PacketSize), 1 + floor(MaxLatency / (t_interval – t_packet))

- MaxBurst = Maximum number of frames of the same stream sent as a burst (back to back)
- MaxCreditBurst = Maximum number of bytes transmitted as a burst after max credit considering the length of the stream frames.
- PacketSize = Size of stream frame plus preamble, SOF and IPG
- MaxLatency = Maximum per hop latency at the observed egress port
- t_interval = Measurement interval
- t_packet = time to transmit one stream frame plus preamble, SOF and IPG

Possibilities to Determine a Worst Case Class A Latency for Fast Ethernet Networks

What is the FE per Hop Class A Worst Case Latency without any further Limitations?

- Is there a general FE per hop class A worst case latency?
  - **No**
  - In an infinite big FE network the worst case traffic pattern which can egress of an infinite long line of bridges when 75 % of the bandwidth through this bridges is allocated but only one stream with 75%/13 of the bandwidth is transmitted, equals a stream which is using a certain amount of time (almost three measurement intervals) 100% of the bandwidth (100 Mbit/s) and after this time 75% of the bandwidth (75 Mbit/s)

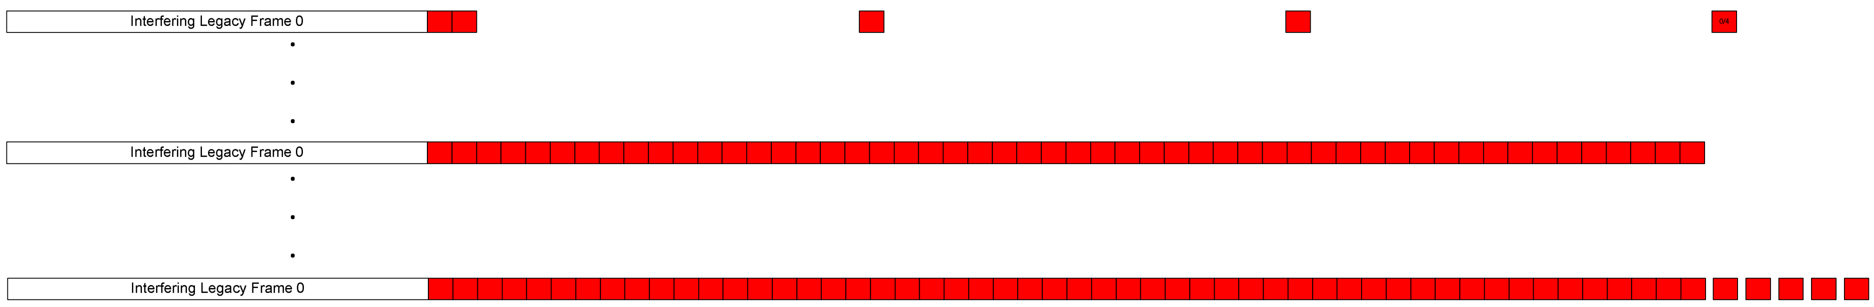

What is the Worst Case FE per Hop Class A Latency without any further Limitations?

- Is there a general FE per hop worst case class A latency?
  - In the worst case one egress port has to handle 13*75 Mbit/s (a certain time 100 Mbit/s)
  - In an infinite big FE network this traffic pattern causes a constant increasing per hop latency (assuming infinite big buffers)
  - In a finite FE network this traffic pattern can only occur a certain amount of time
  - As a result the class A per hop latency in a finite FE network has a limit but this limit depends on the topology of the network

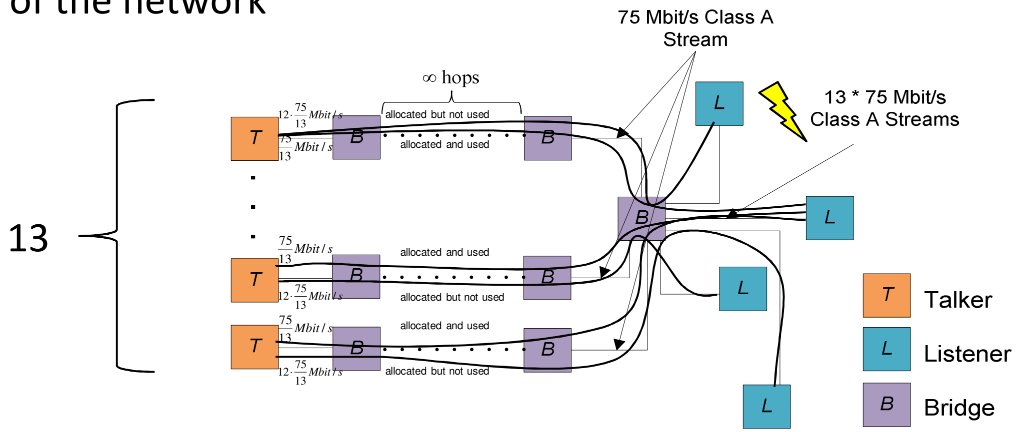

What Figure Should Be Reported by SRP to the Listener?

Is it realistic to report a smaller number with the “hope” that in the average this is correct?

- How should such a figure be determined?
- The averaging effect depends on the hop count
- A stream with a small hop count but a very high per hop latency would have a higher latency than reported by SRP

→ **It is very important that we can rely on the reported figures**

→ **As there is no general worst case latency figure it might be interesting to make the MIB variable portTCMaxLatency also writeable**

Possibilities to Archive a FE Worst Case per Hop Class A Latency

- Limit fan-in
- Decrease allowed max bandwidth (<75%)
- Limit max class A bandwidth and speed up shaper
- Limit network diameter (e.g. max diameter = 7 hops)

Other Ideas:

- No Plug and Play support for FE (only engineered networks)

Limit Fan-In

Fan-In = 13

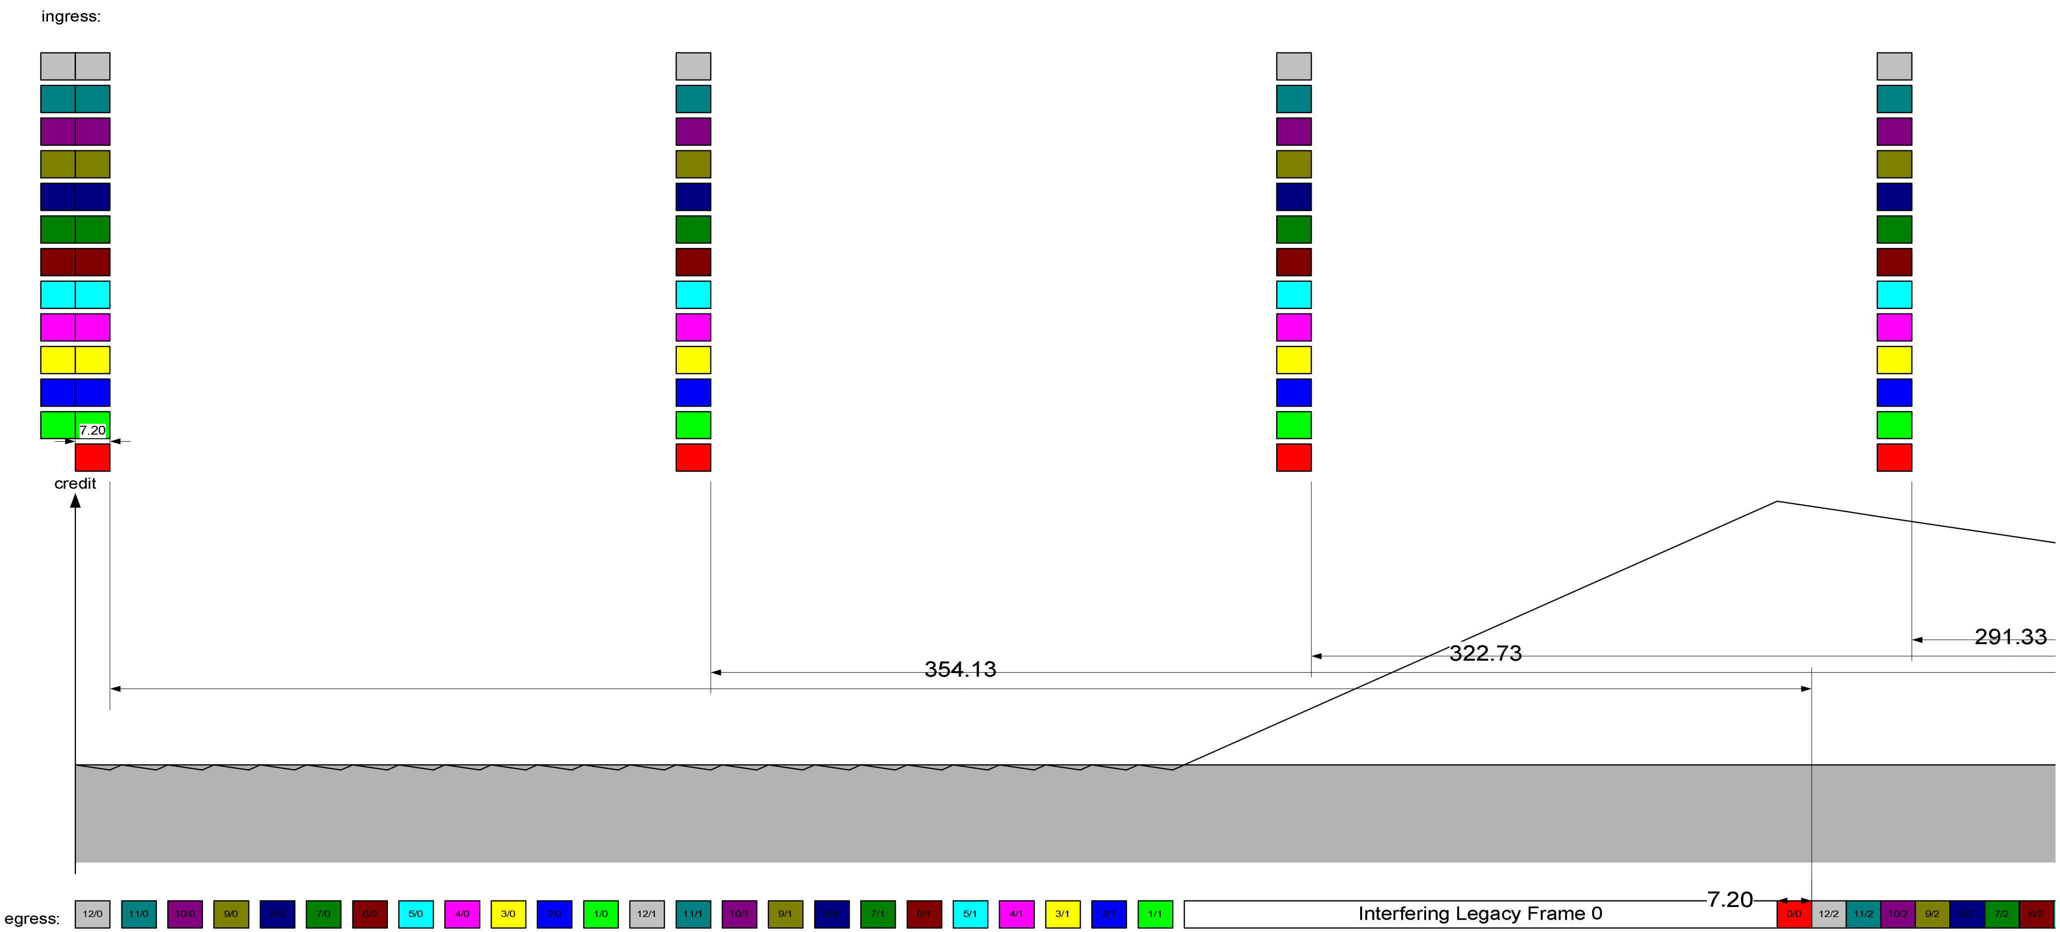
ingress:
7.20
credit
354.13
322.73
291.33
egress:
12/0
11/0
10/0
9/0
7/0
6/0
5/0
4/0
3/0
2/0
1/0
12/1
11/1
10/1
9/1
7/1
6/1
5/1
4/1
3/1
2/1
1/1
Interfering Legacy Frame 0
7.20
0/0
12/2
11/2
10/2
9/2
7/2
6/2

Fan-In = 6

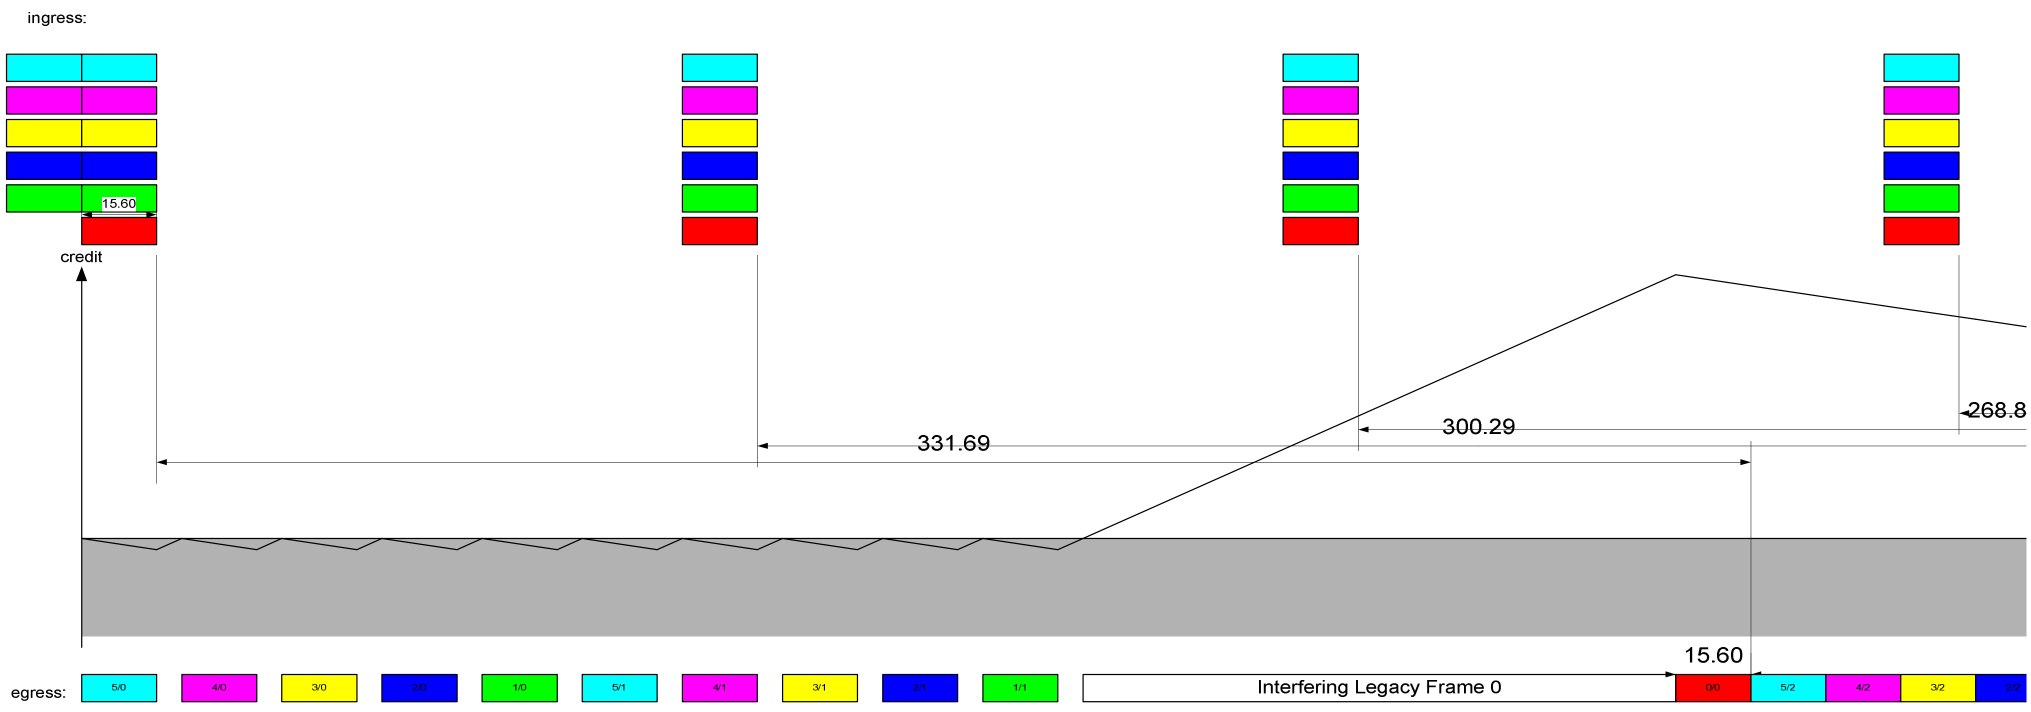
ingress:
15.60
credit
331.69
300.29
268.8
15.60
egress:
Interfering Legacy Frame 0

Decrease Max Class A Bandwidth

Max Class A Bandwidth = 75%

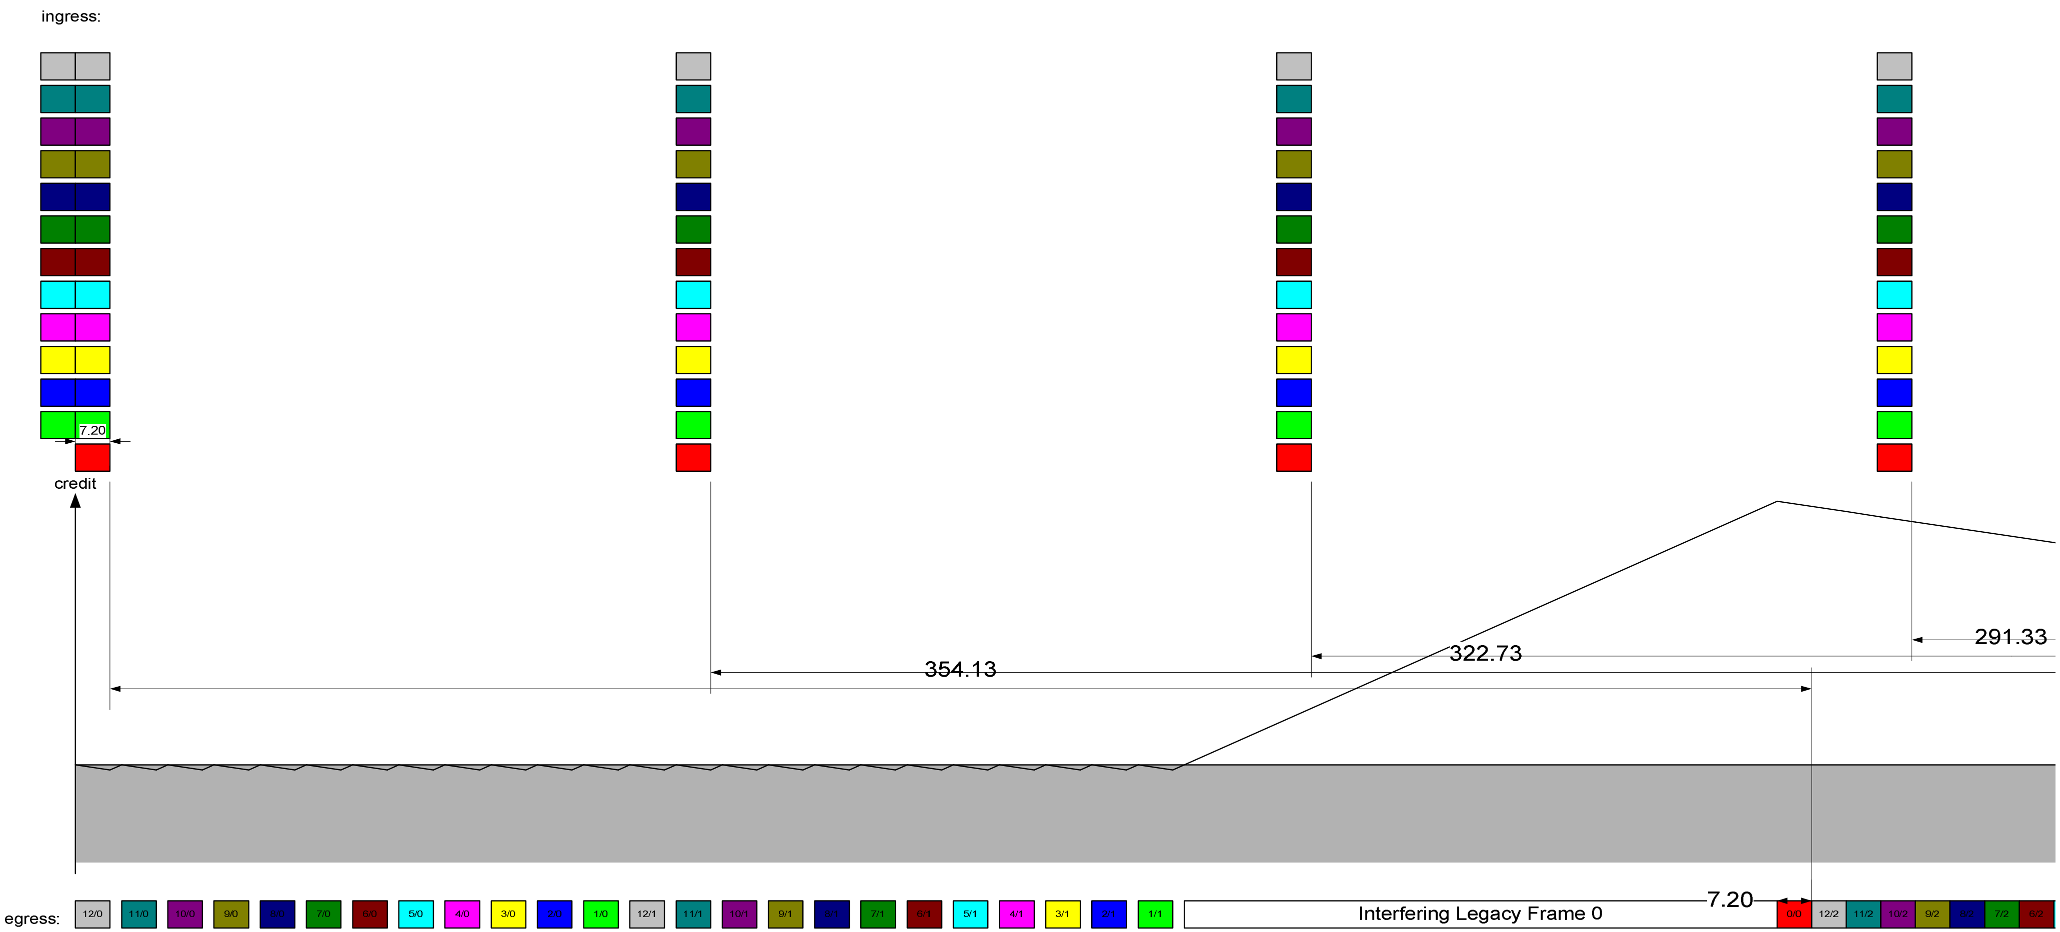

Max Class A Bandwidth = 46%

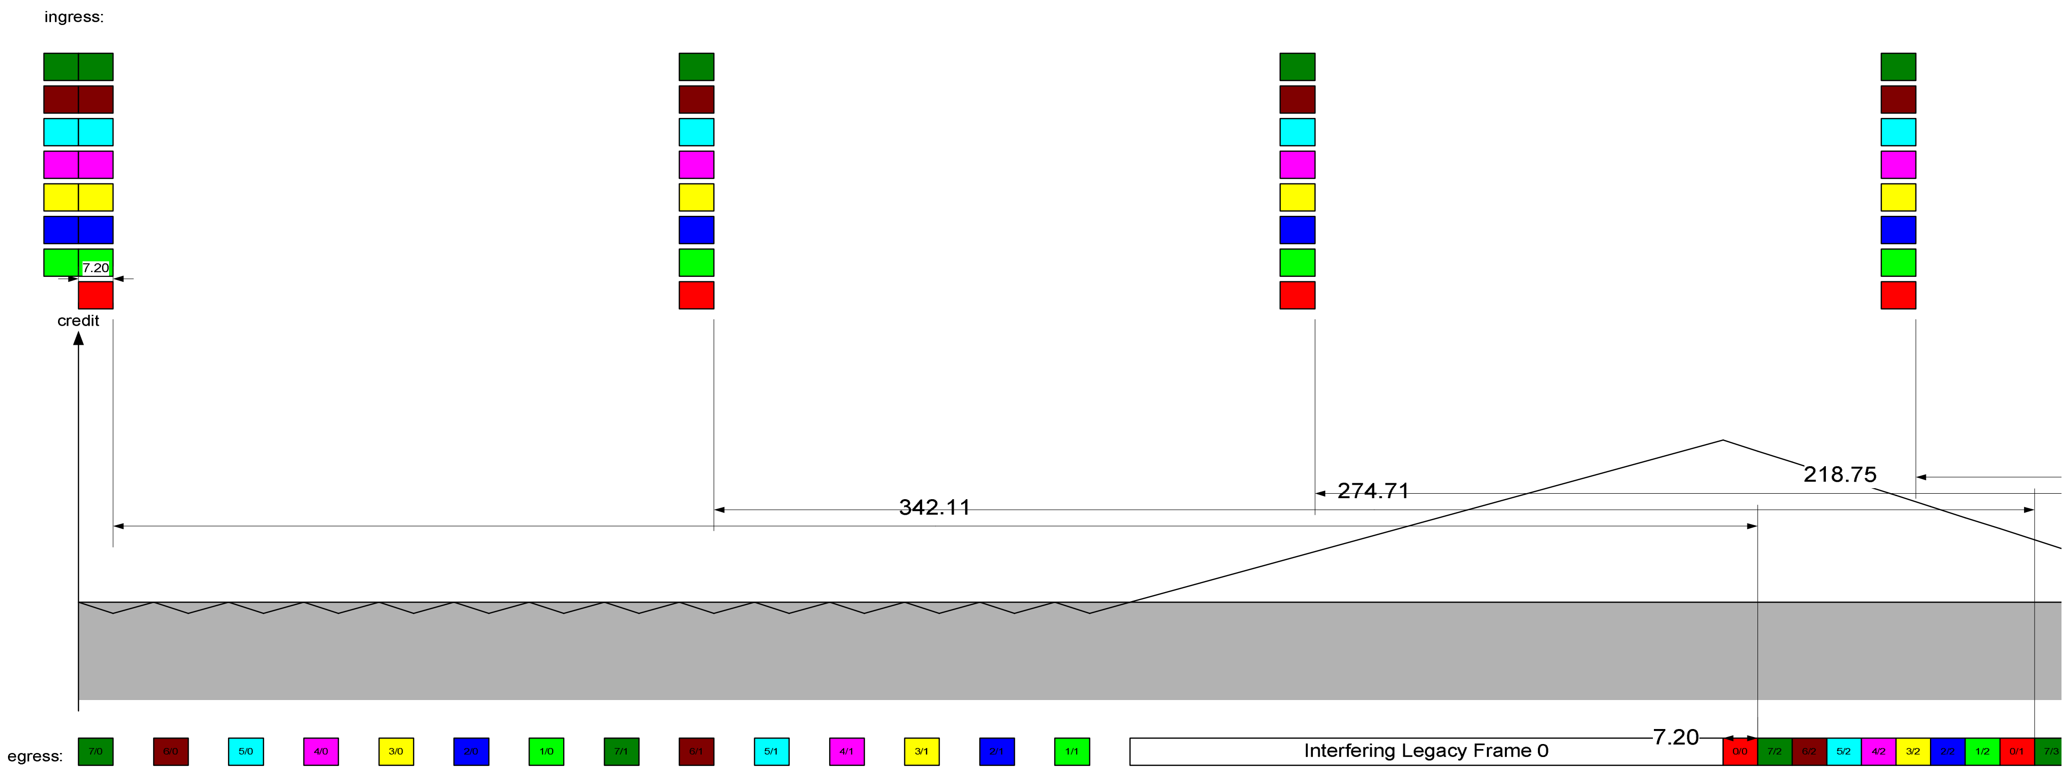

Results

- Limiting the fan-in or decreasing the max class A bandwidth does not provide an upper boundary to latency.
- The effects on latency are only very marginal.

Decrease Max Class A Bandwidth and Speed Up Shaper

Max Class A Bandwidth = 75%

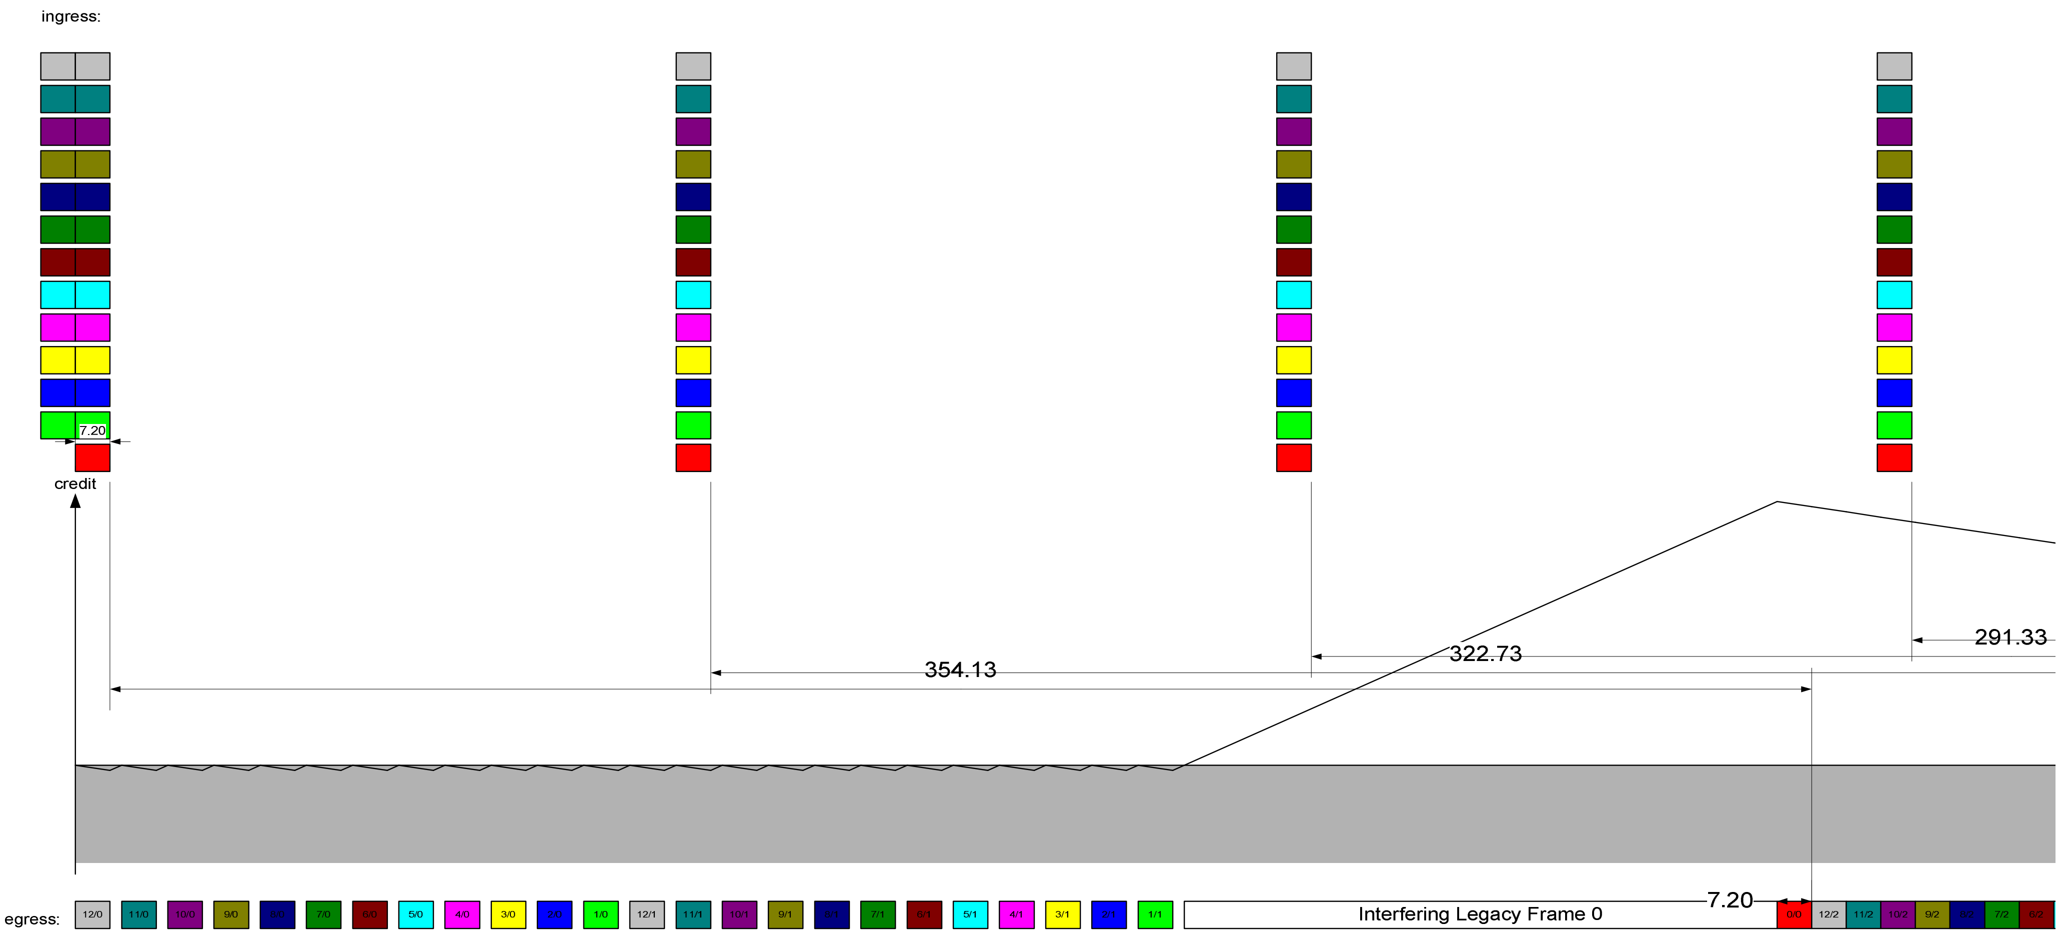

Hochschule
Deggendorf

Max Class A Bandwidth = 46% (allocated 75%)

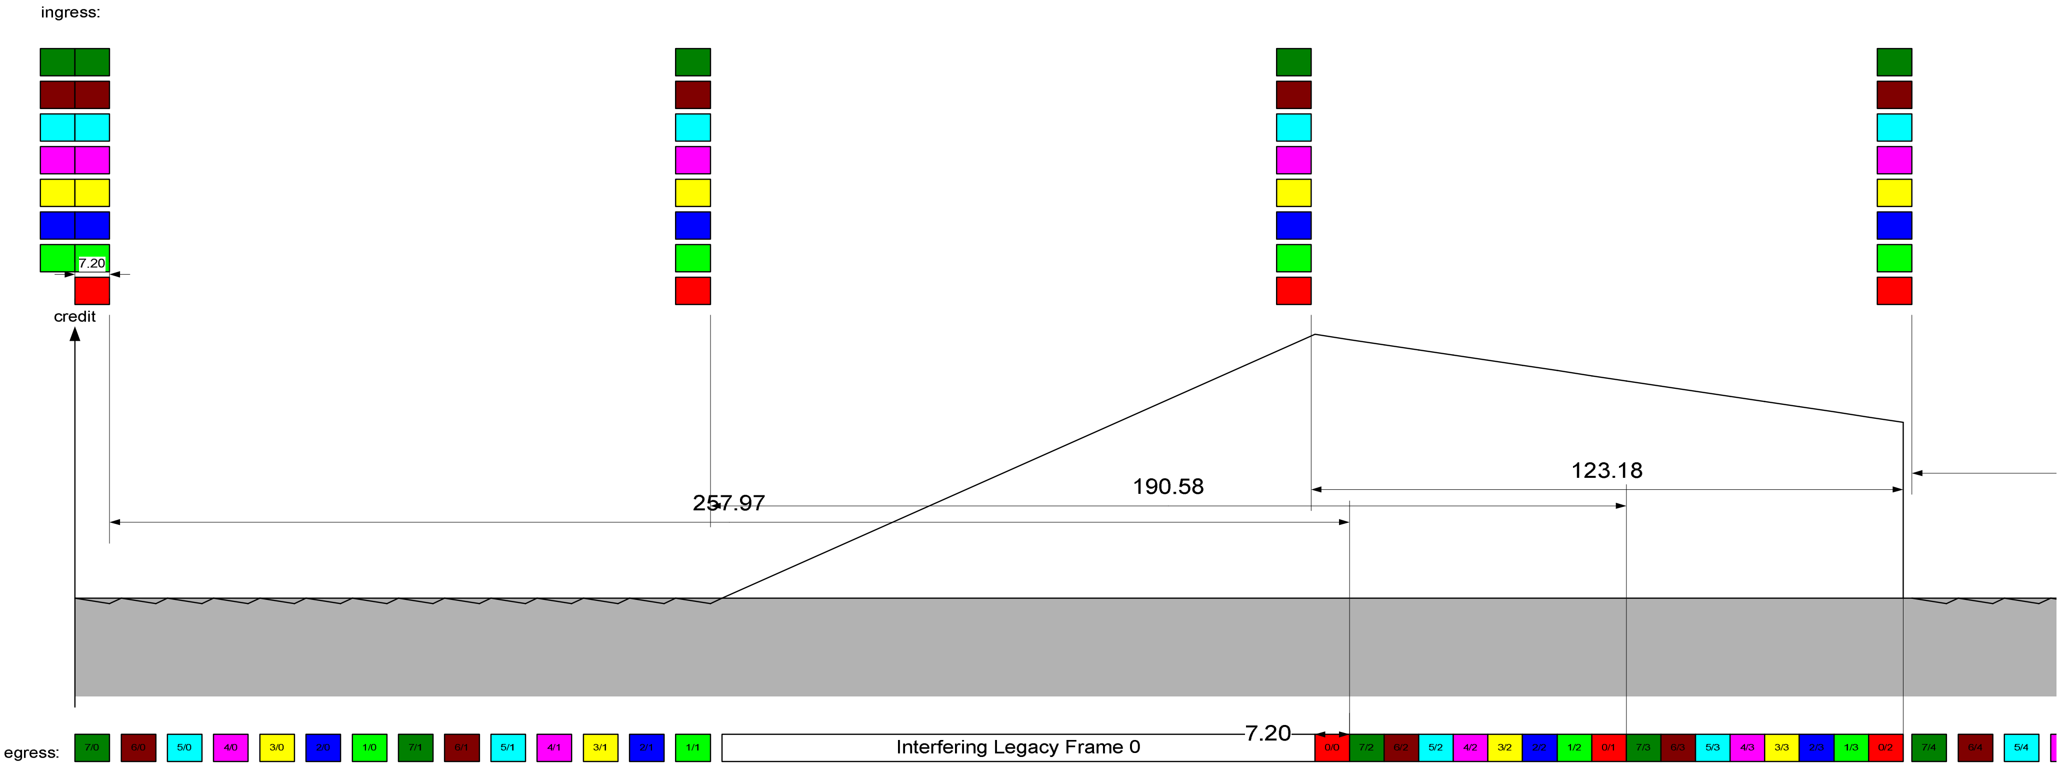

Hochschule
Deggendorf

Max Class A Bandwidth = 75%

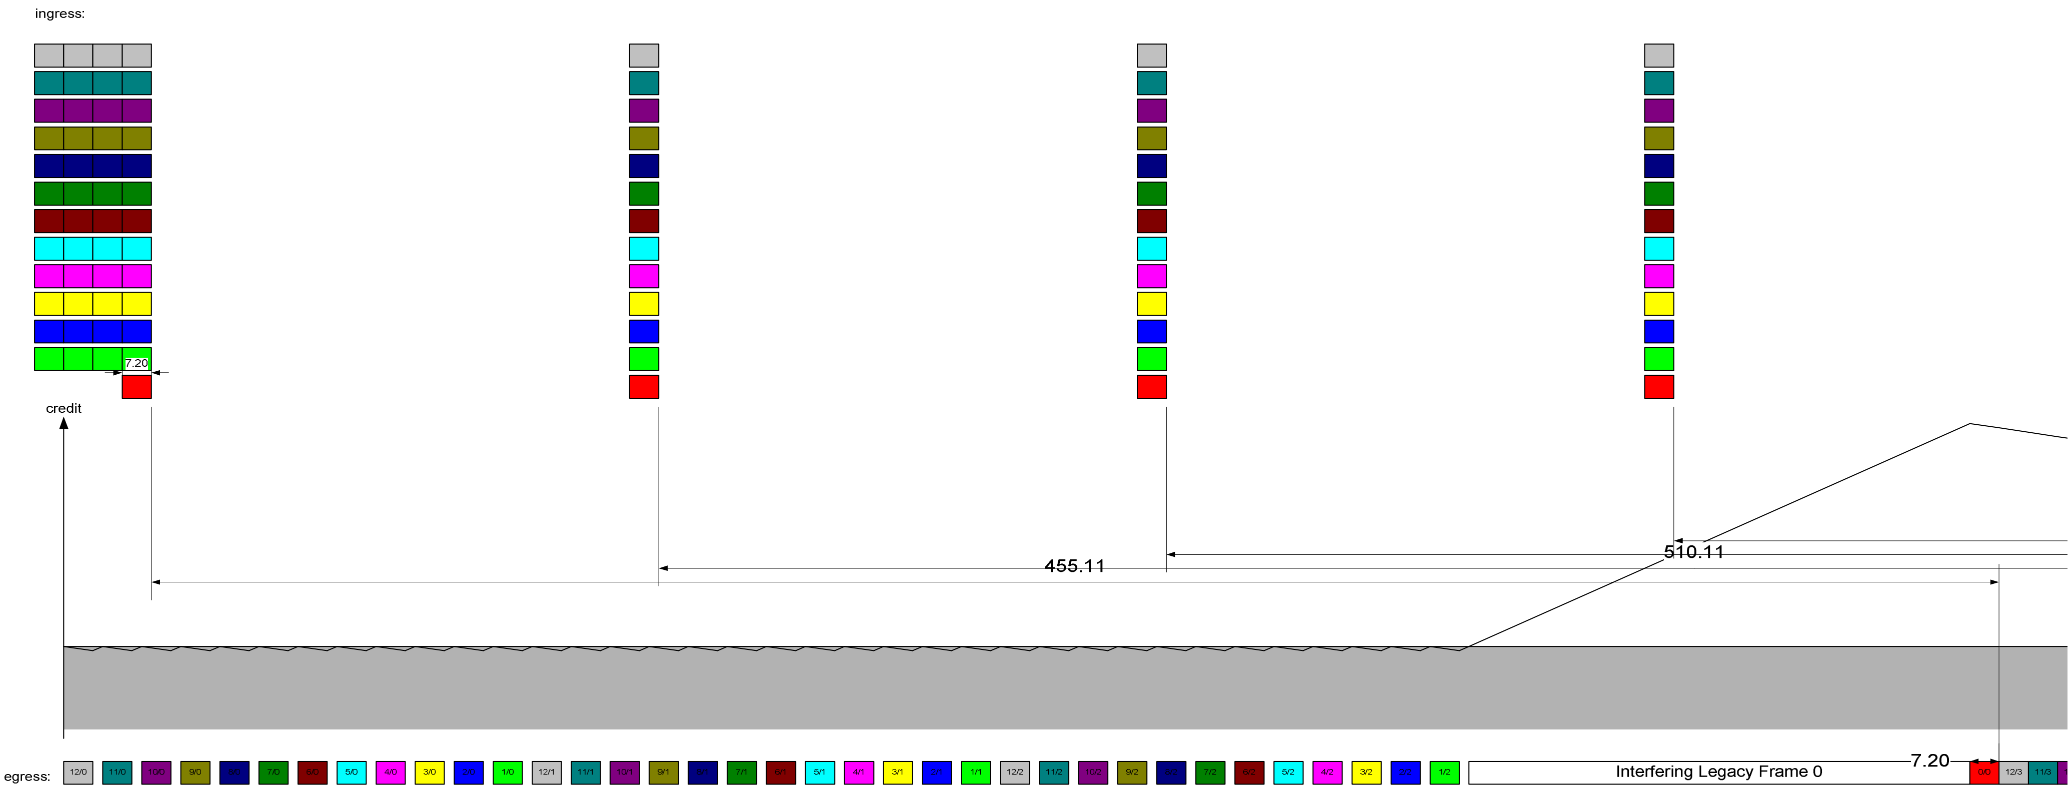

Max Class A Bandwidth = 46% (allocated 75%)

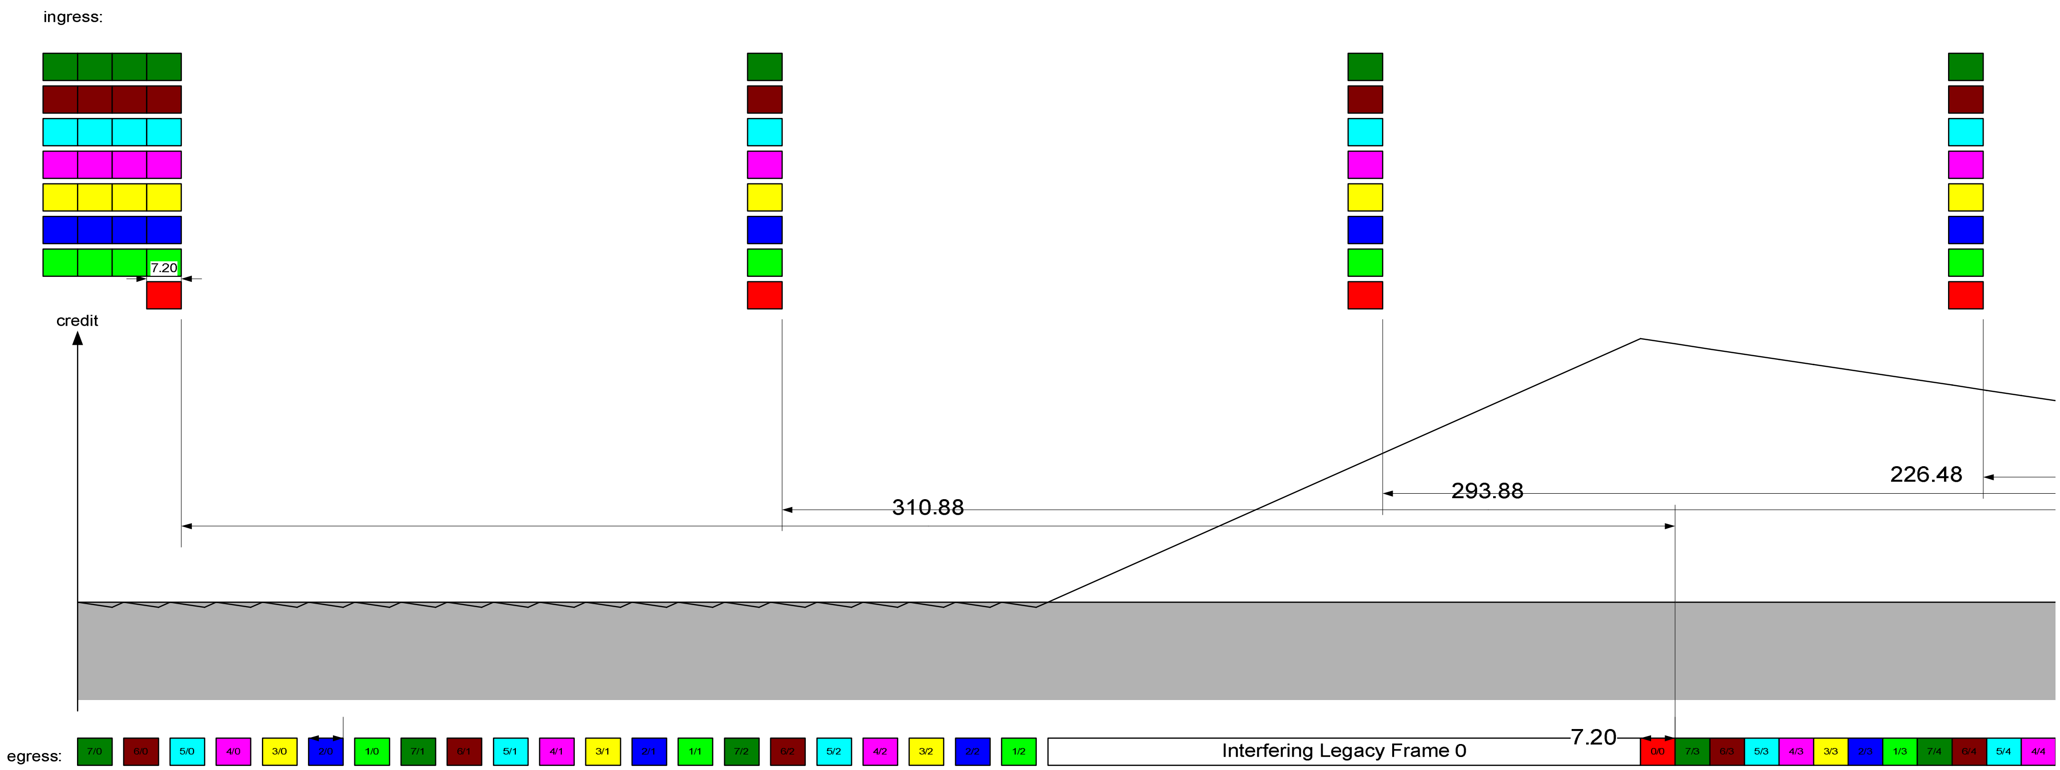

Results

- Speeding up the shaper does not lead to a topology independent per hop worst case latency but together with a bandwidth limitation it leads to a significantly lower latency.

- Bursts are more likely with the speeded up shaper but in the worst case not bigger

Limit Network Diameter

Limit Network Diameter

- Limit = 7
- Resulting worst case topology:
  - The latency of the stream between T1 and L1 is observed
  - The talkers/talker networks above the dashed lines are only examples. In the worst case there are 12 of the shown type connected to the corresponding bridge.
  - The listeners for the talker T are not shown in the graphic they are always located one bridge after the talker/talker network enters the bridge. The talkers connected to the last bridge stream to the listener L1

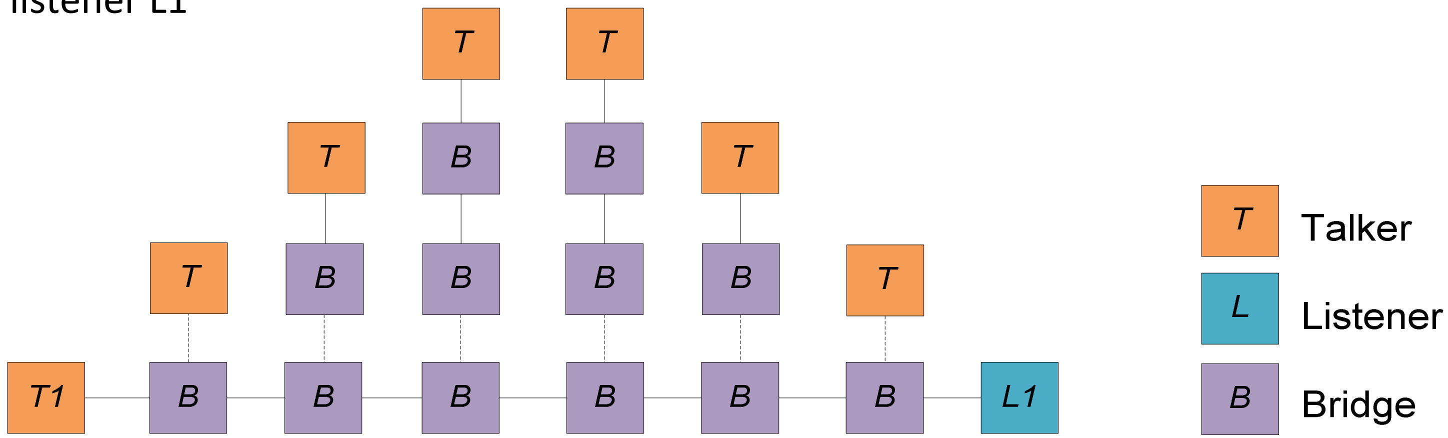

Example

Talker

talker: 245.94

total: 245.94

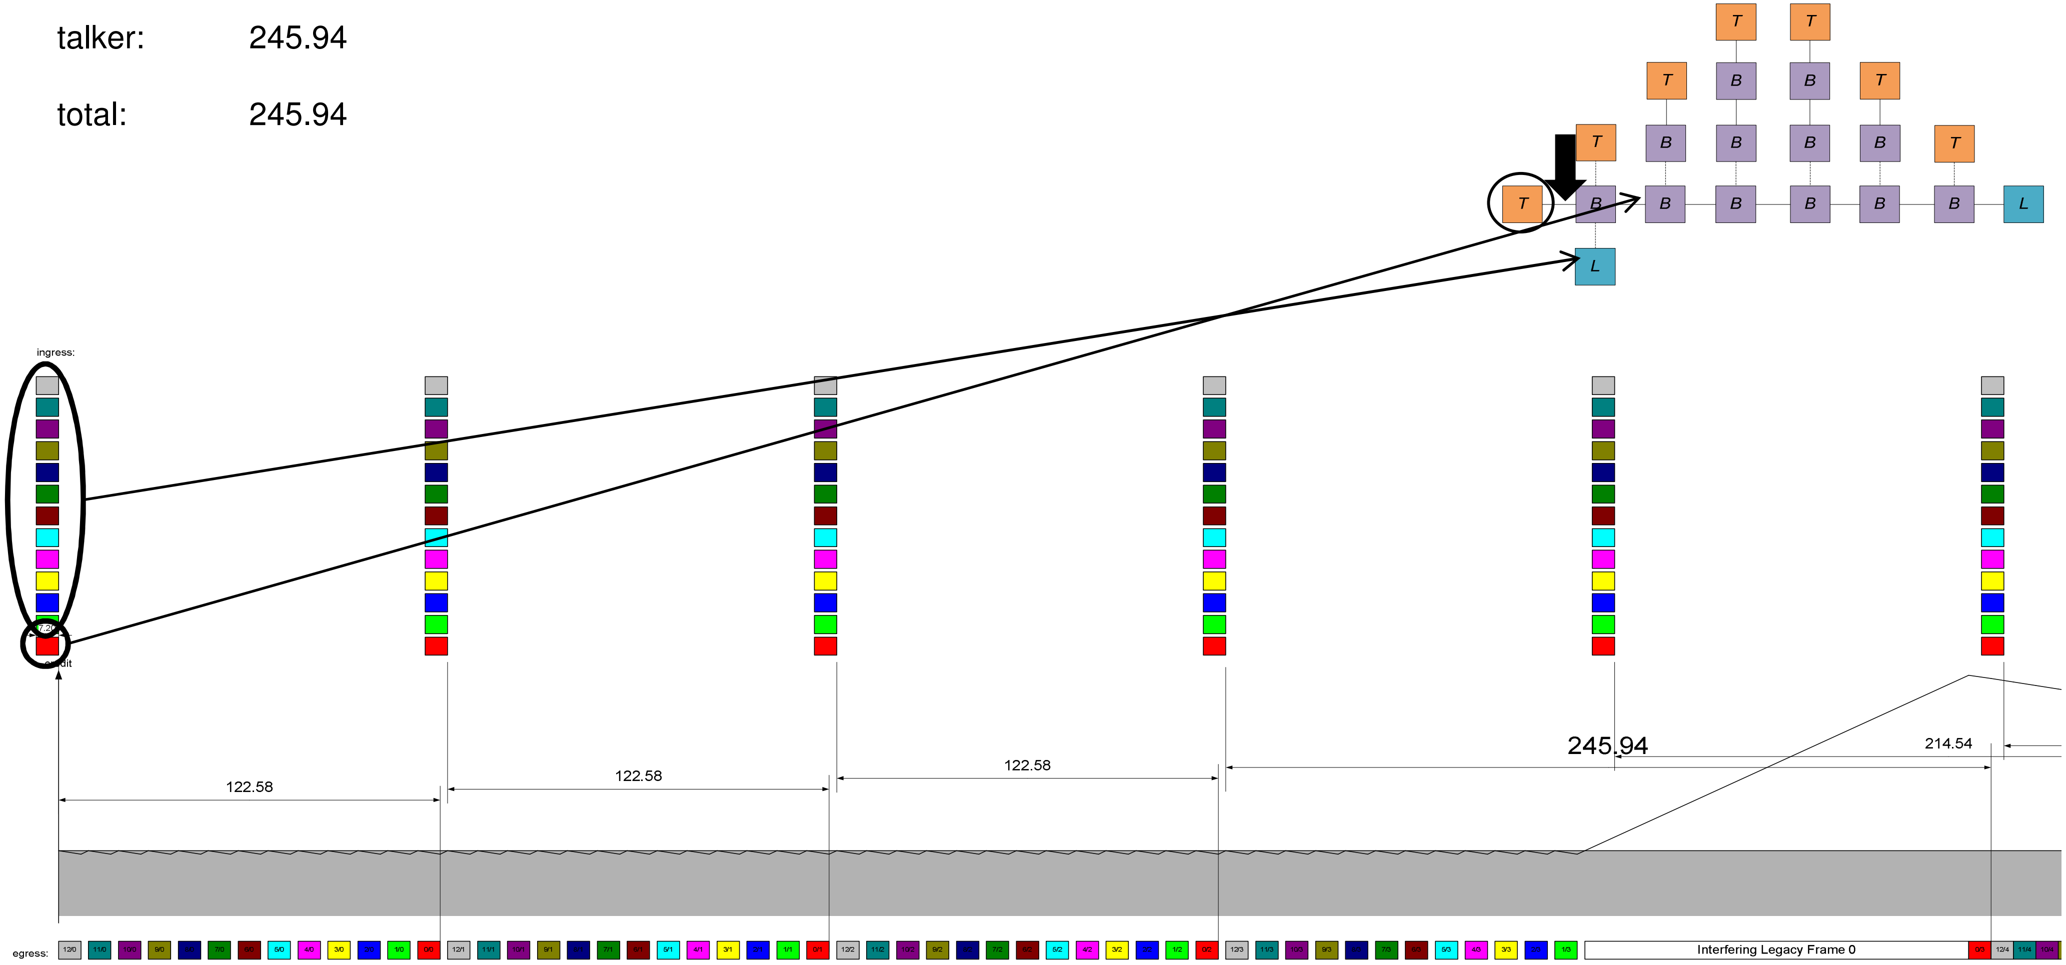

Interfering Talker

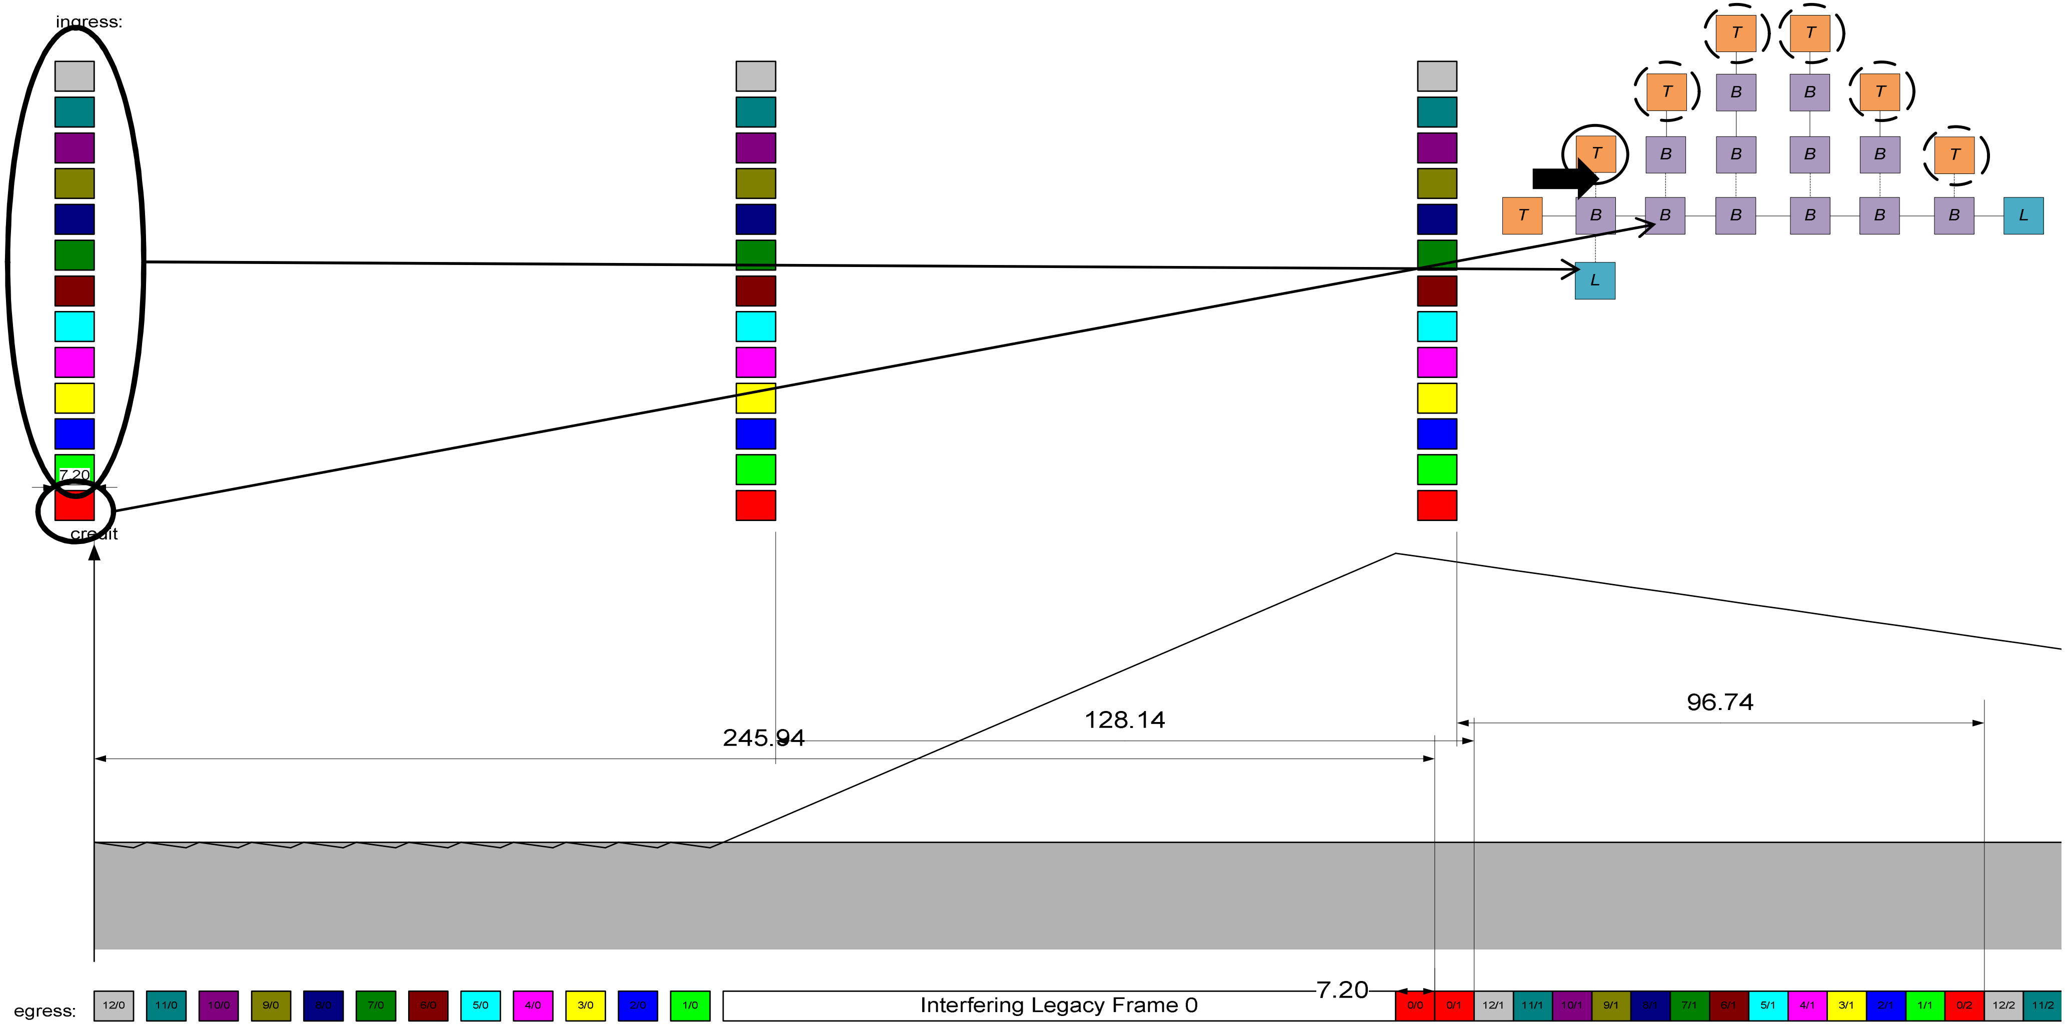

After First Bridge

talker: 245.94
1st bridge: 455.84

total: 701.78

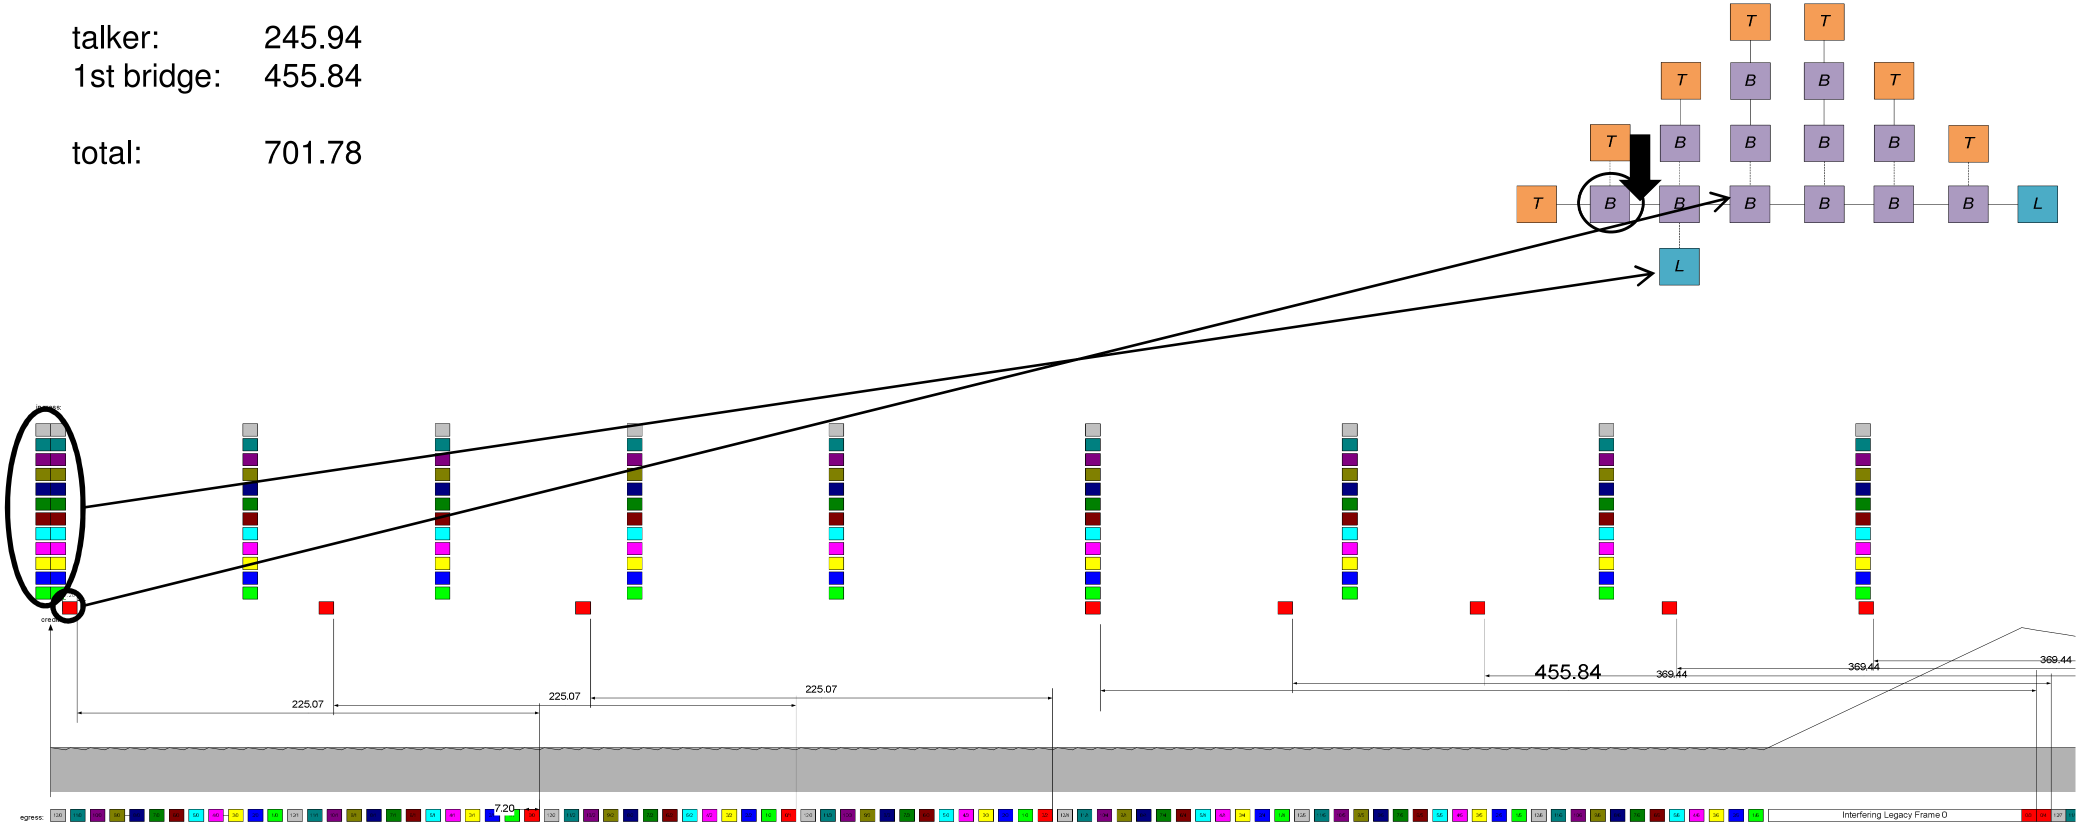

Hochschule
Deggendorf

1st Interfering Bridge

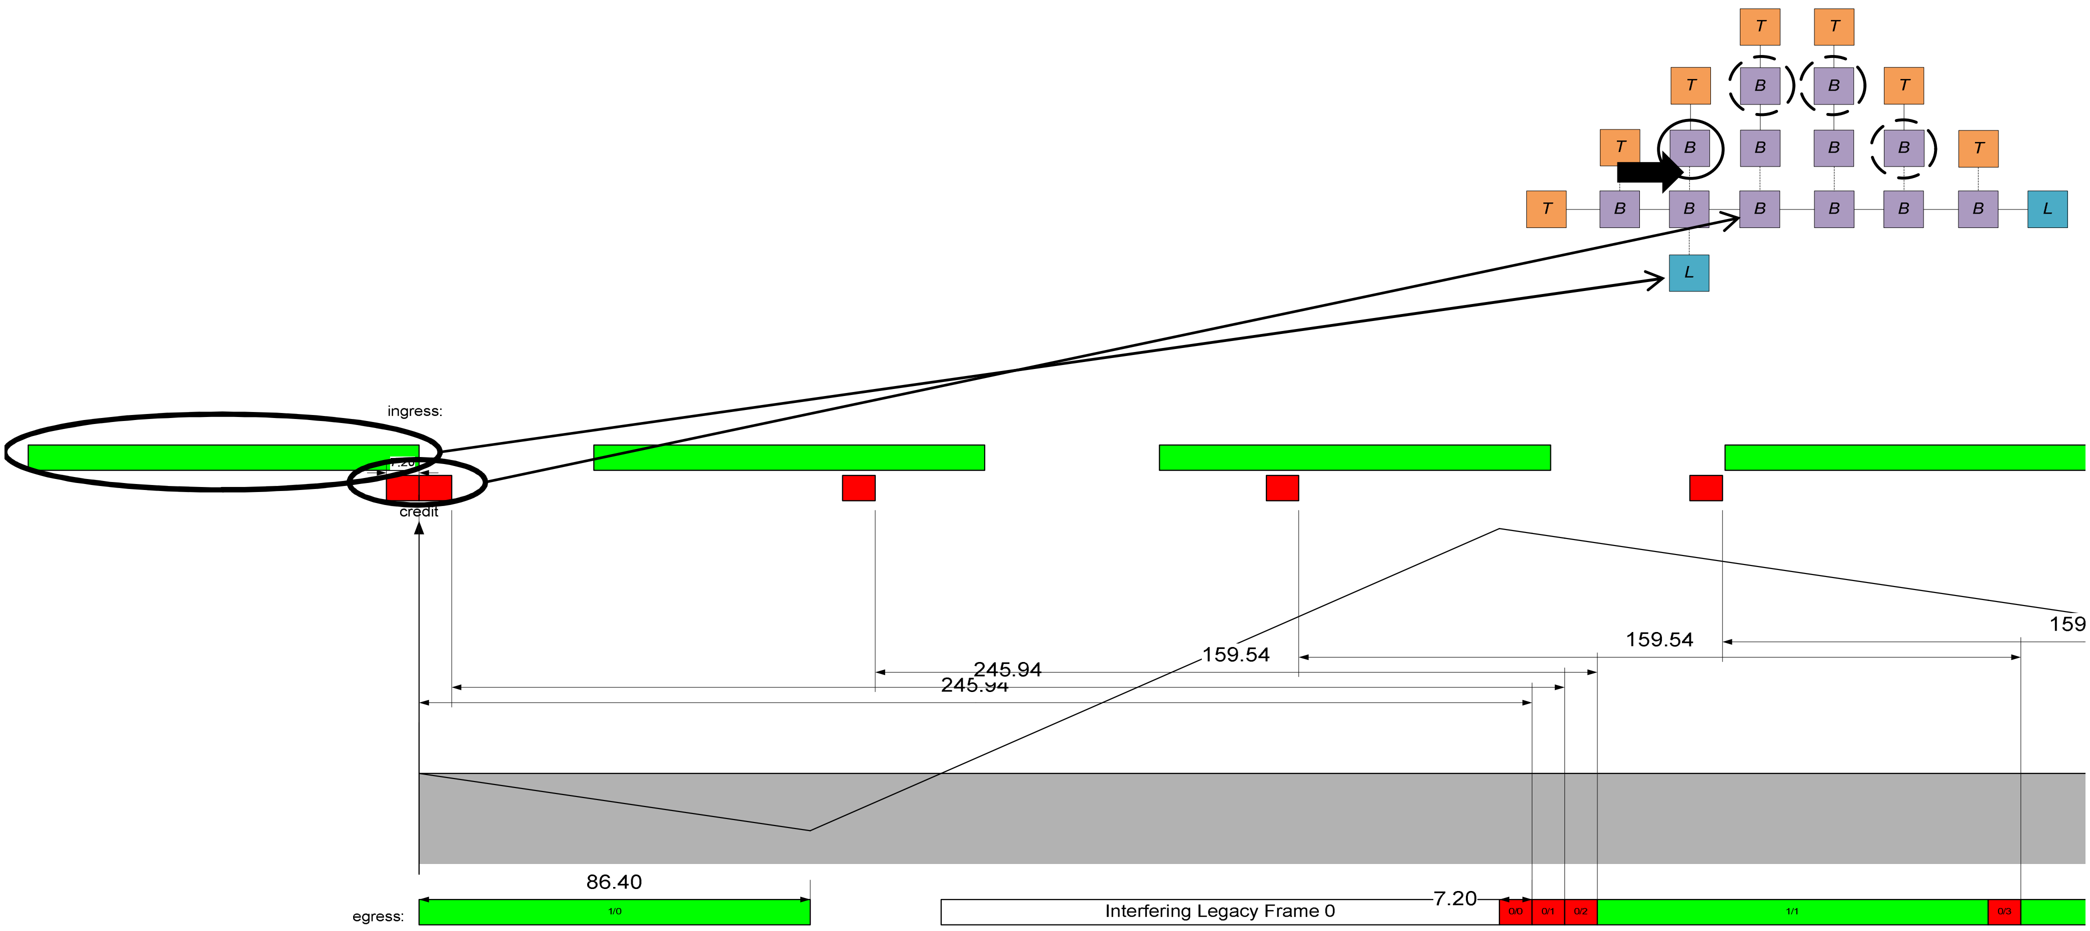

After Second Bridge

talker: 245.94µs
1st bridge: 455.84µs
2nd bridge: 542.37µs

total: 1244.15µs

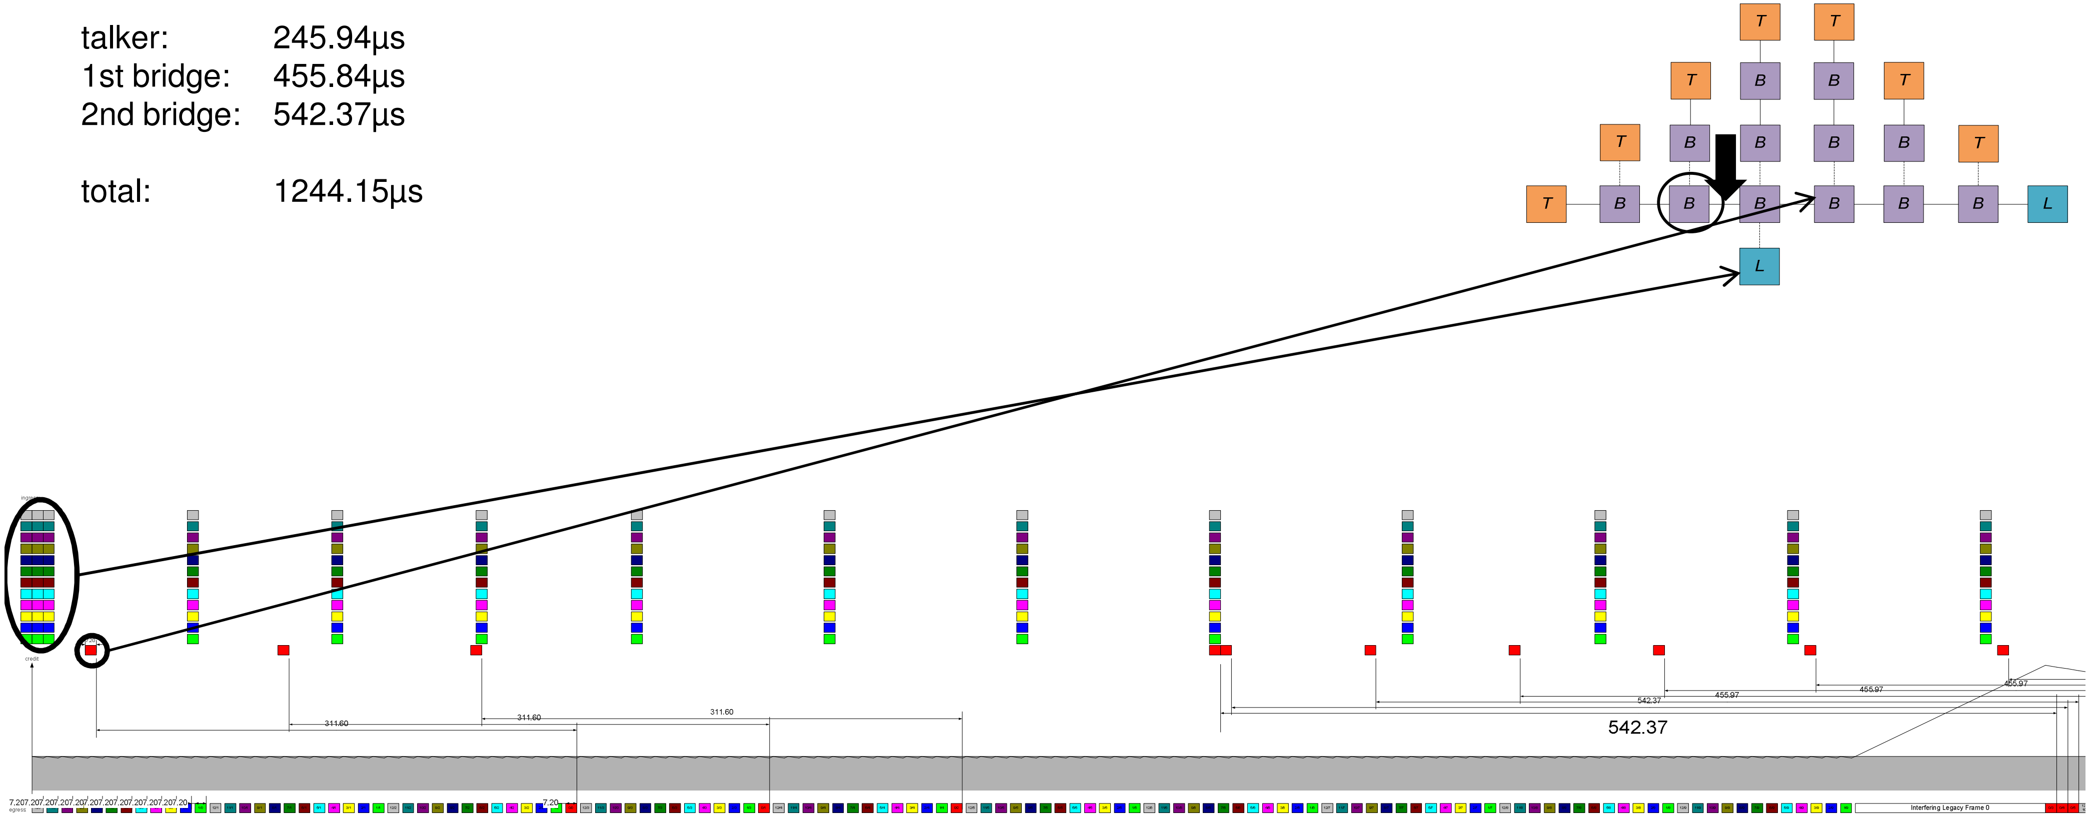

2nd Interfering Bridge

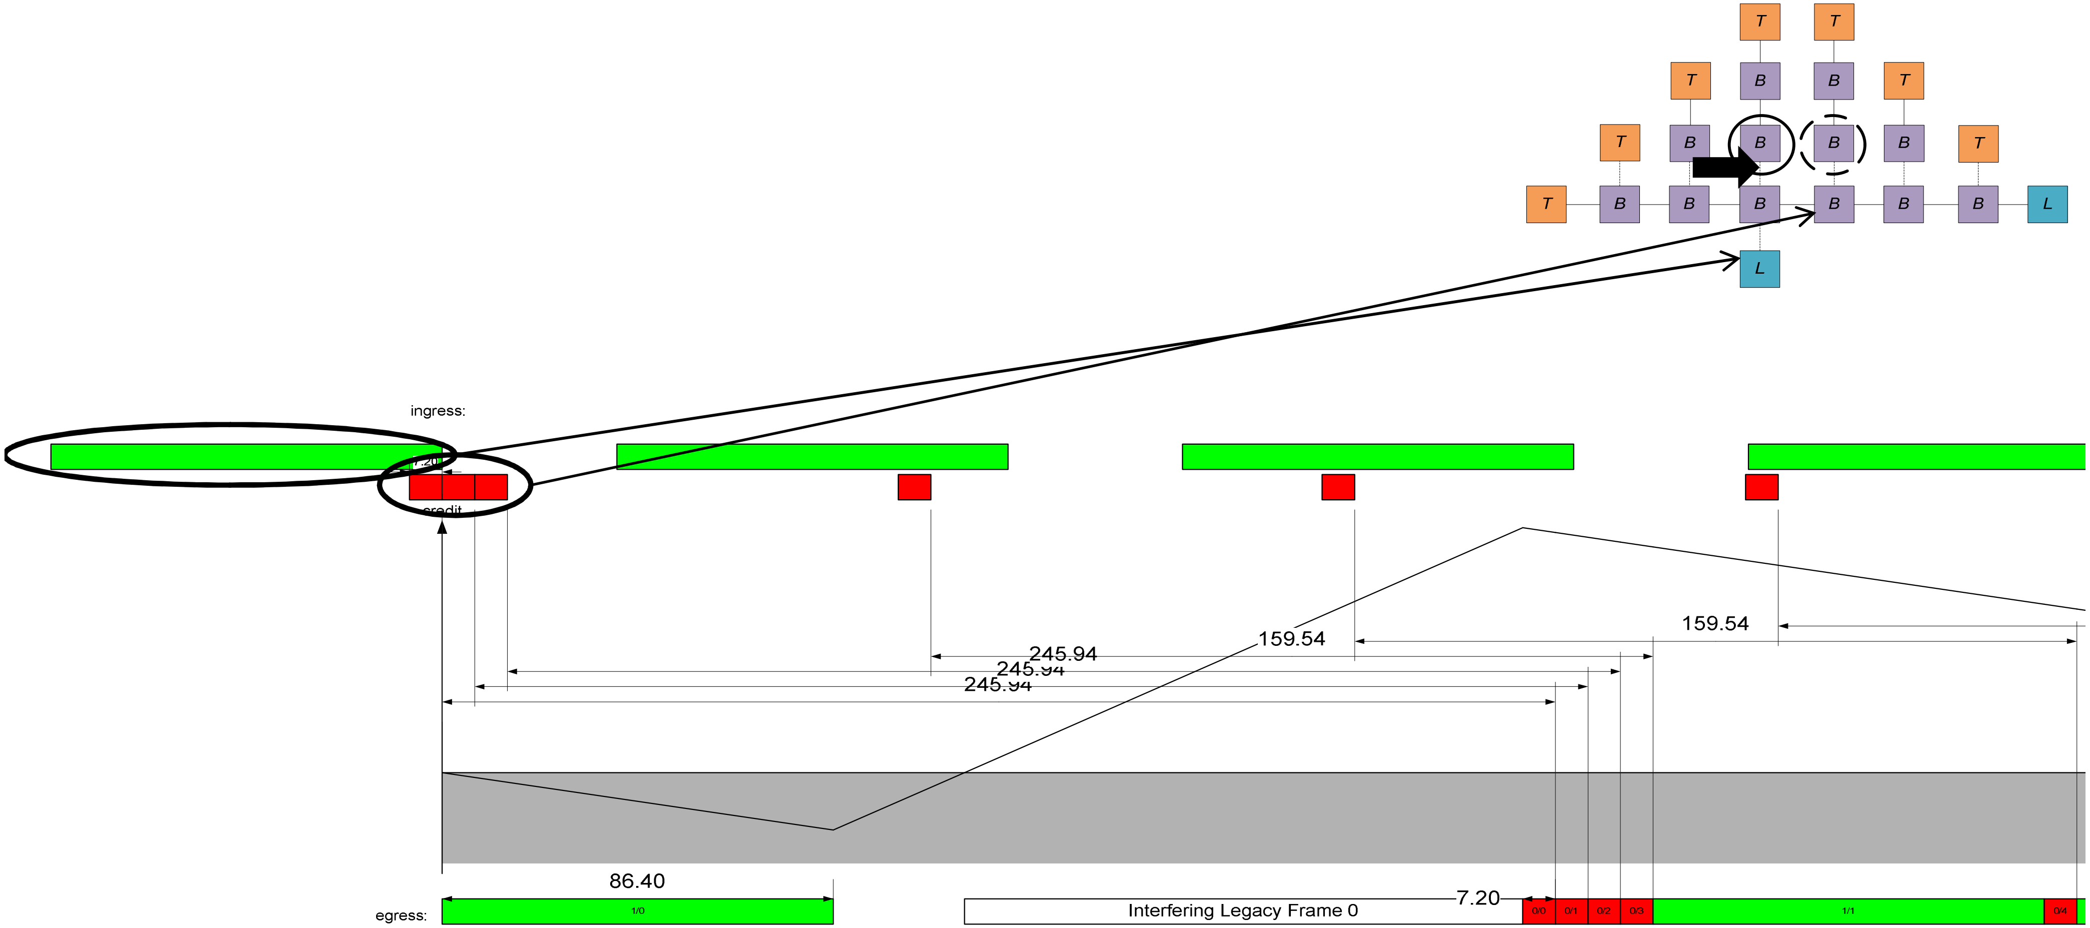

After Third Bridge

talker:	245.94µs
1st bridge:	455.84µs
2nd bridge:	542.37µs
3rd bridge:	628.92µs
total:	1873.07µs

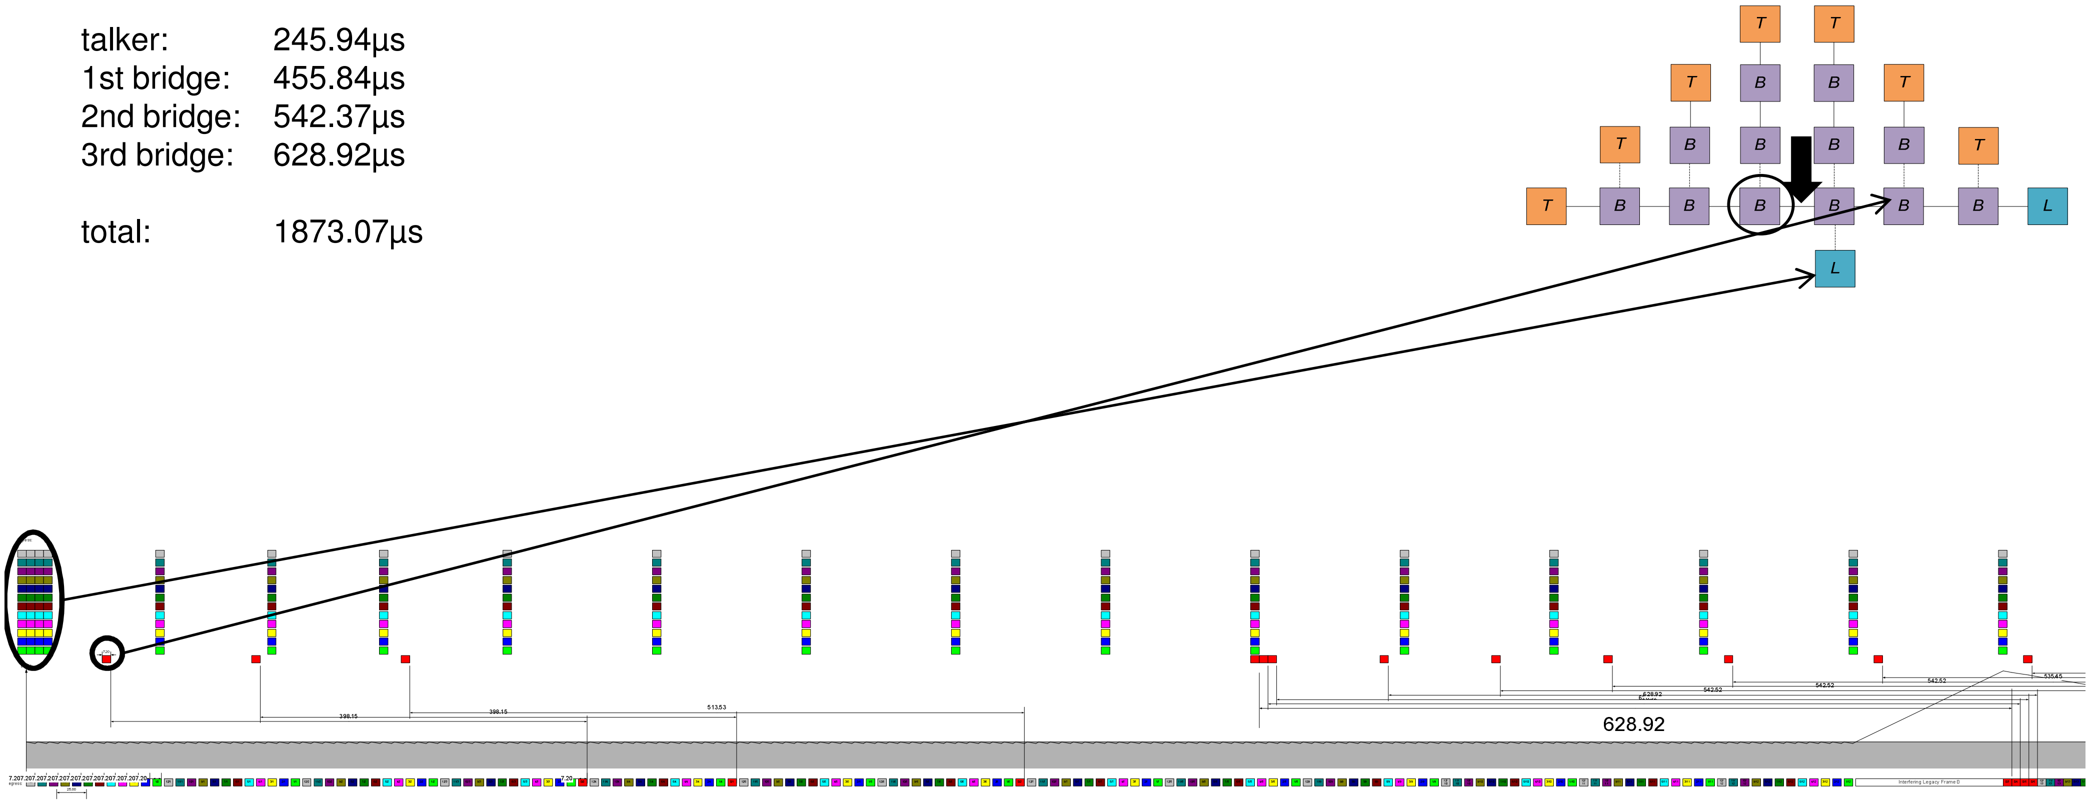

After Third Bridge (1)

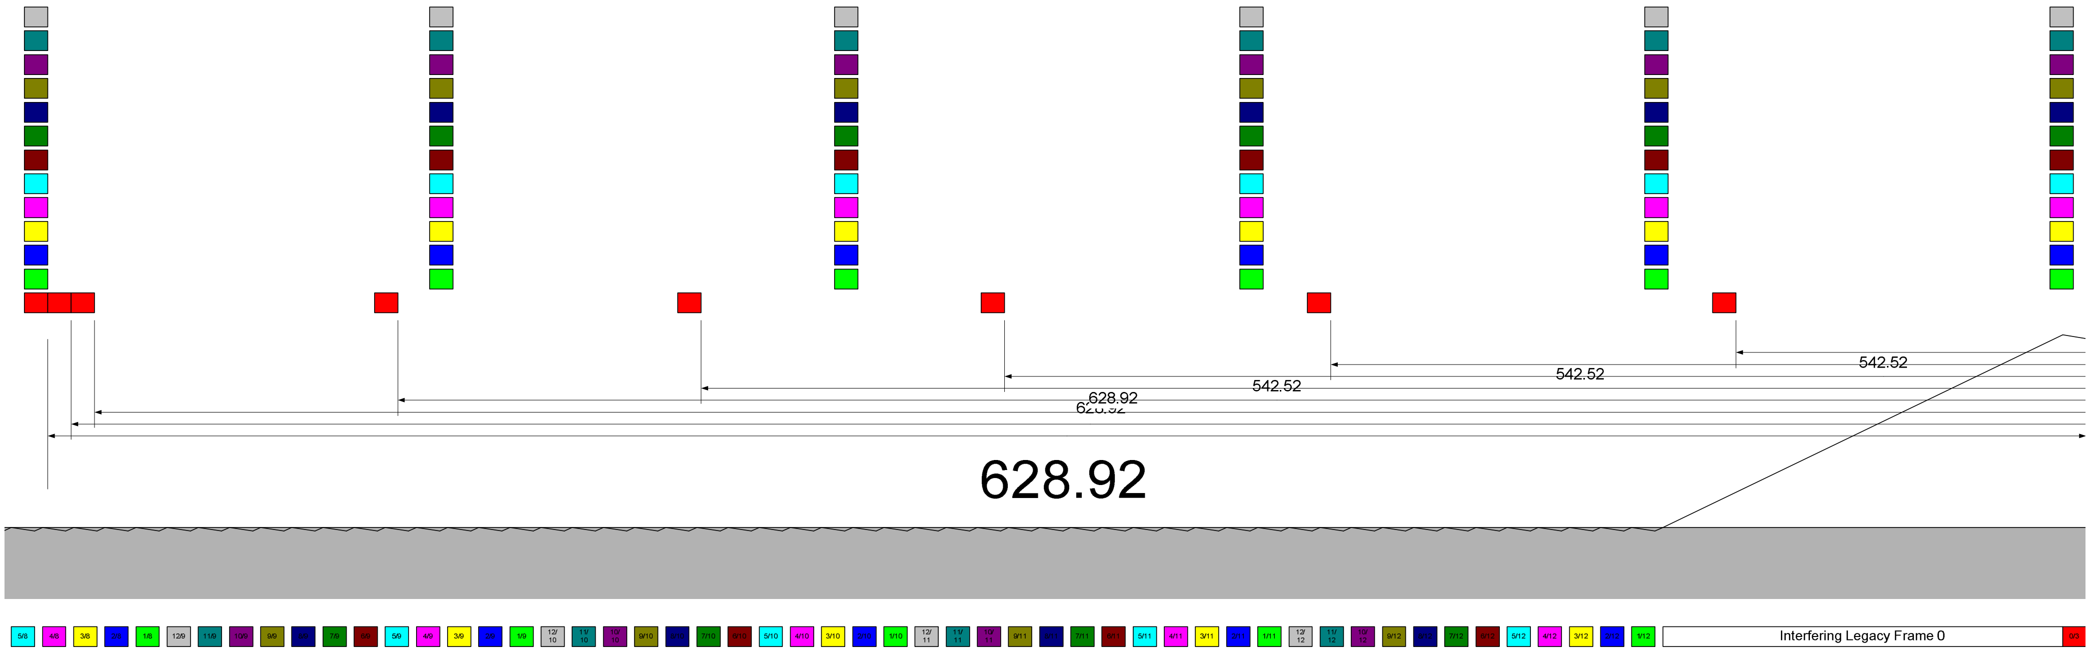

After Fourth Bridge

talker: 245.94µs
1st bridge: 455.84µs
2nd bridge: 542.37µs
3rd bridge: 628.92µs
4th bridge: 609.68µs

total: 2482.75µs

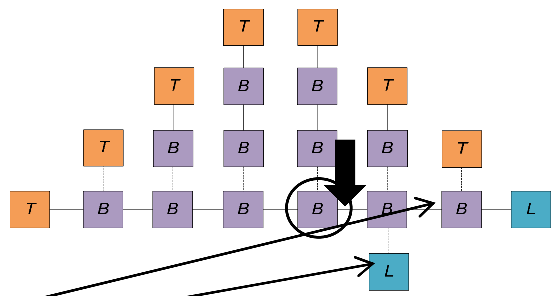

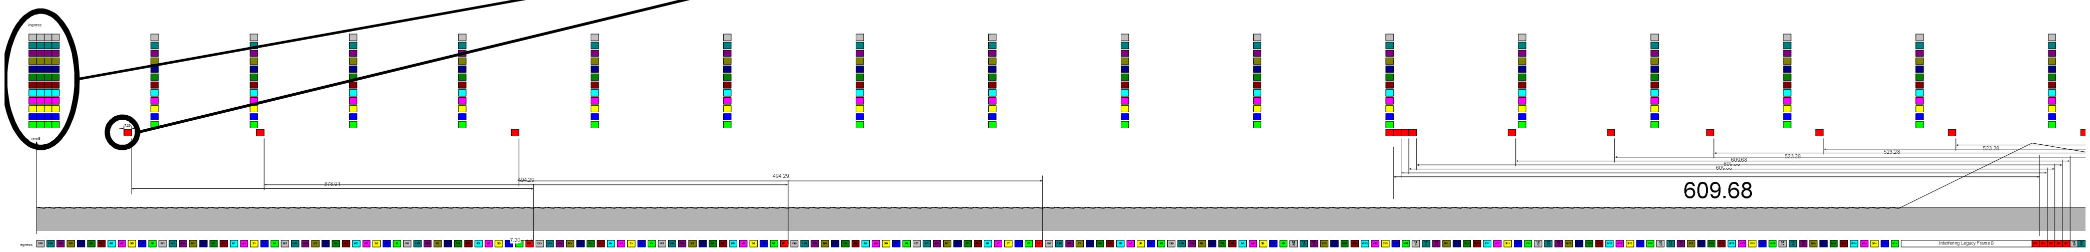

After Fourth Bridge (1)

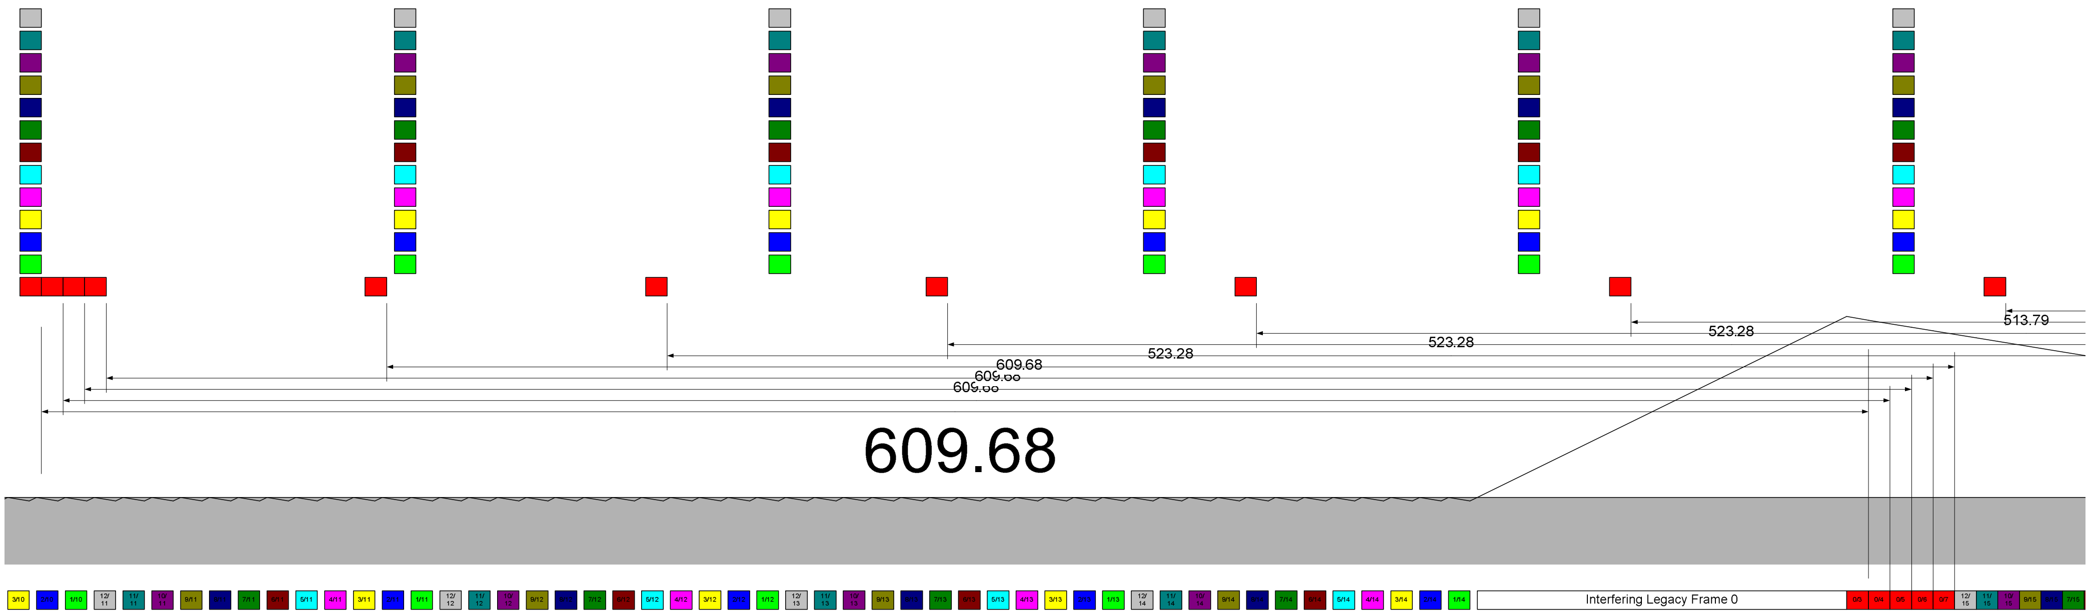

After Fifth Bridge

talker:	245.94µs
1st bridge:	455.84µs
2nd bridge:	542.37µs
3rd bridge:	628.92µs
4th bridge:	609.68µs
5th bridge:	484.69µs
total:	2967.44µs

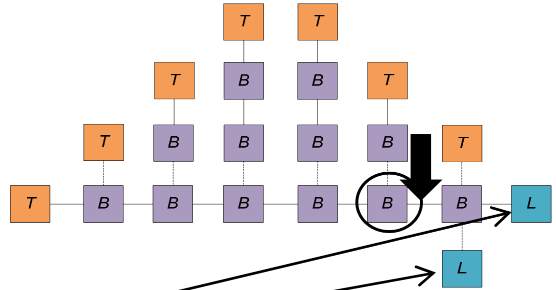

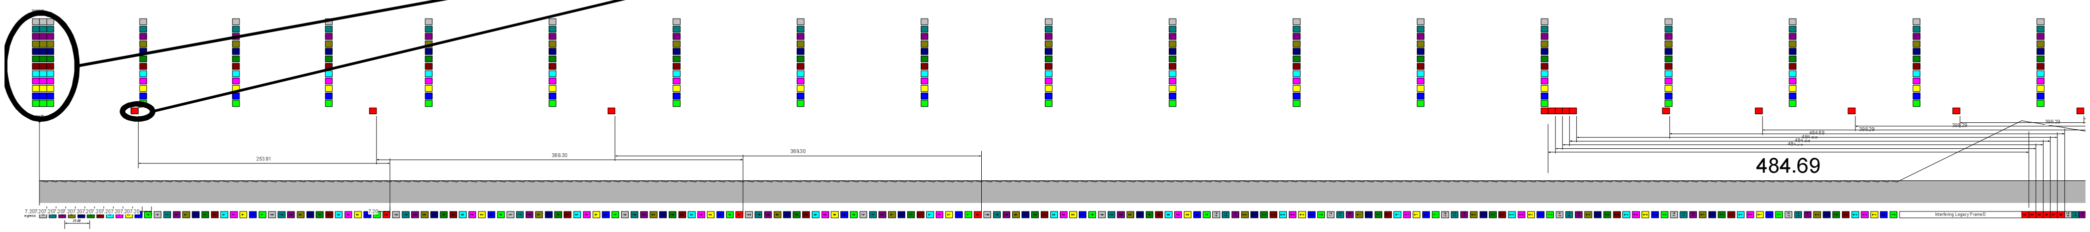

After Fifth Bridge (1)

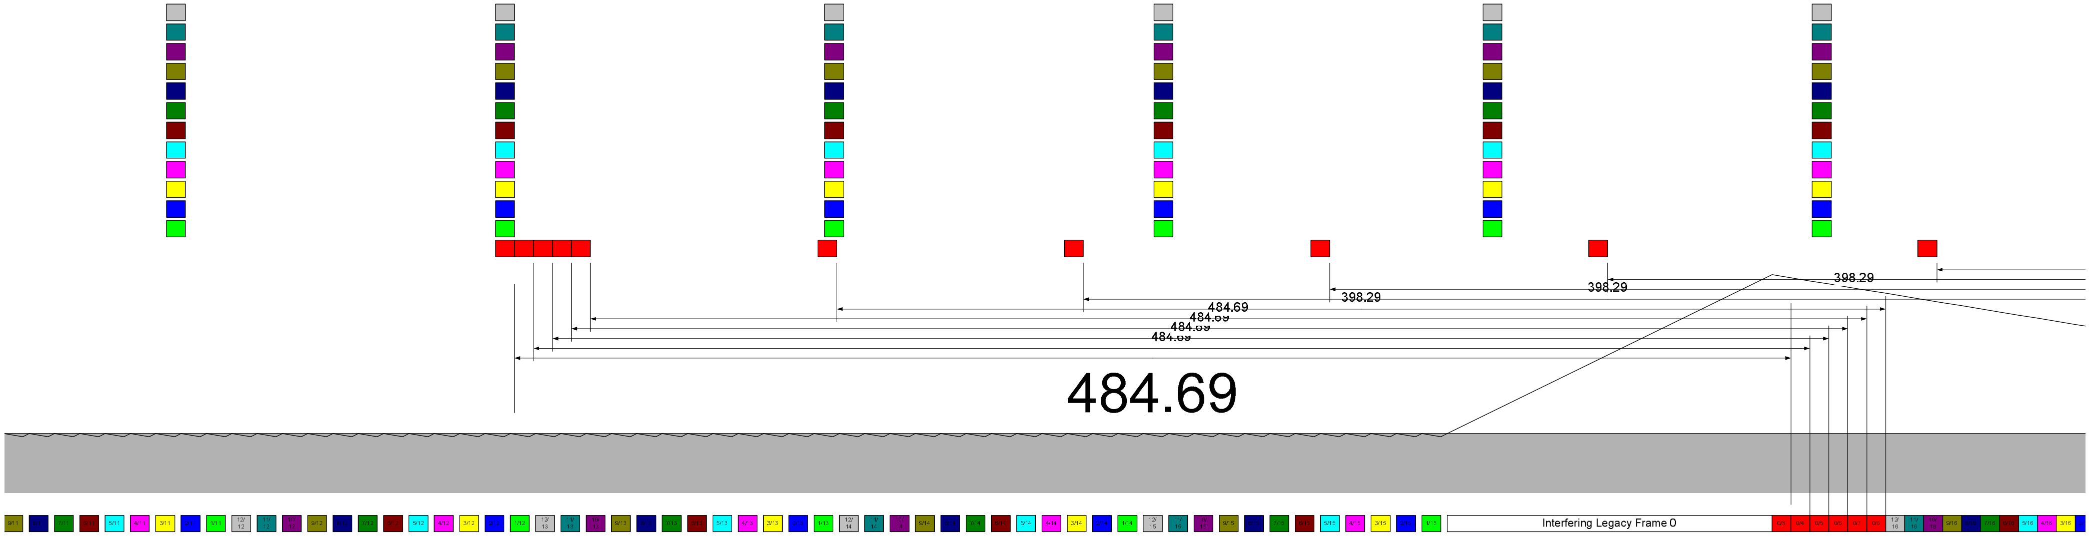

After Sixth Bridge

talker:	245.94µs
1st bridge:	455.84µs
2nd bridge:	542.37µs
3rd bridge:	628.92µs
4th bridge:	609.68µs
5th bridge:	484.69µs
6th bridge:	359.68µs (436.6µs)
total:	3327.12µs

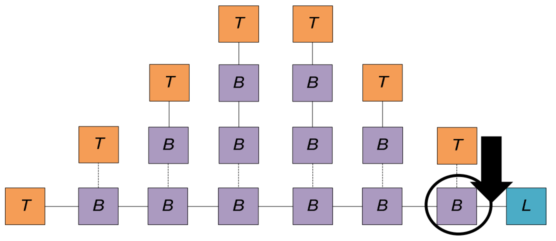

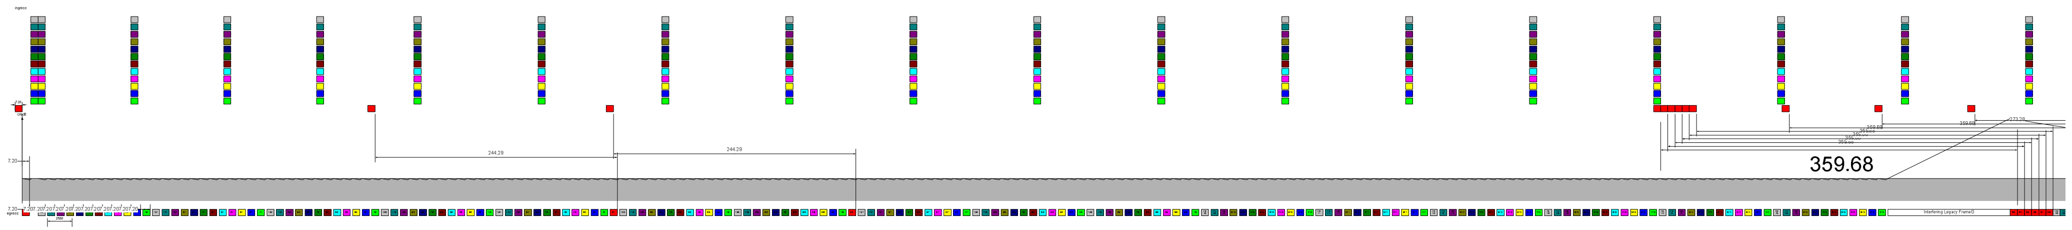

Results

- The example shows only “almost” the worst case. The growing gap between the observed frame and the frame before the observed one reduces latency. The result would be worse if the interfering burst would happen nearer to the delayed frame. At the last bridge this effect almost reduces the latency by 100µs. Additionally the bursts can be bigger (see first slides)

- Limiting the network diameter does not prevent from very high per hop latencies (>600µs).

⇒ Limiting the topology limits the latency, but without other additional limitations the worst case latency is still very high

→ **Limiting the diameter of a network is not feasible for industrial networks especially to such “small” numbers like 7 hops**

Limit Network Diameter

- Is it necessary to report the worst case at every hop?
  - The worst case will occur at the bridges 3 and 4 (counted from T1)
  - Reporting this worst case at all bridges is for the rest of the bridges to much
  - But it is possible that there is a talker connected to the bridge 3 and a listener to bridge 4 in this case there is no averaging effect

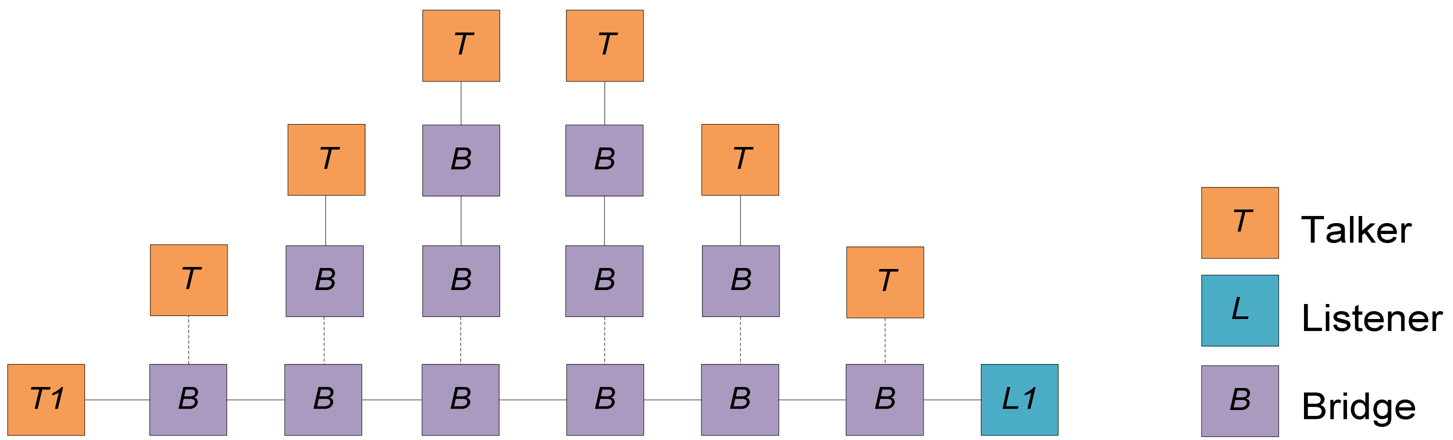

Max Class A Bandwidth = 46% + Speeded Up Shaper + Max Diameter = 7 Hops

Hochschule
Deggendorf

Talker

talker: 197.87 (245.94)

total: 197.87 (245.94)

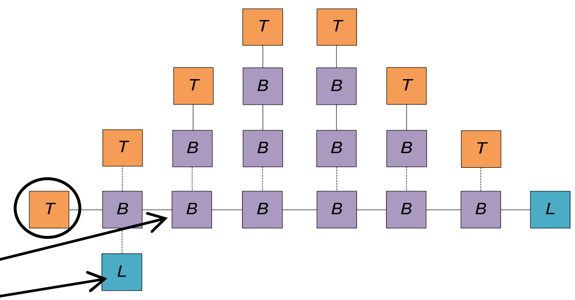

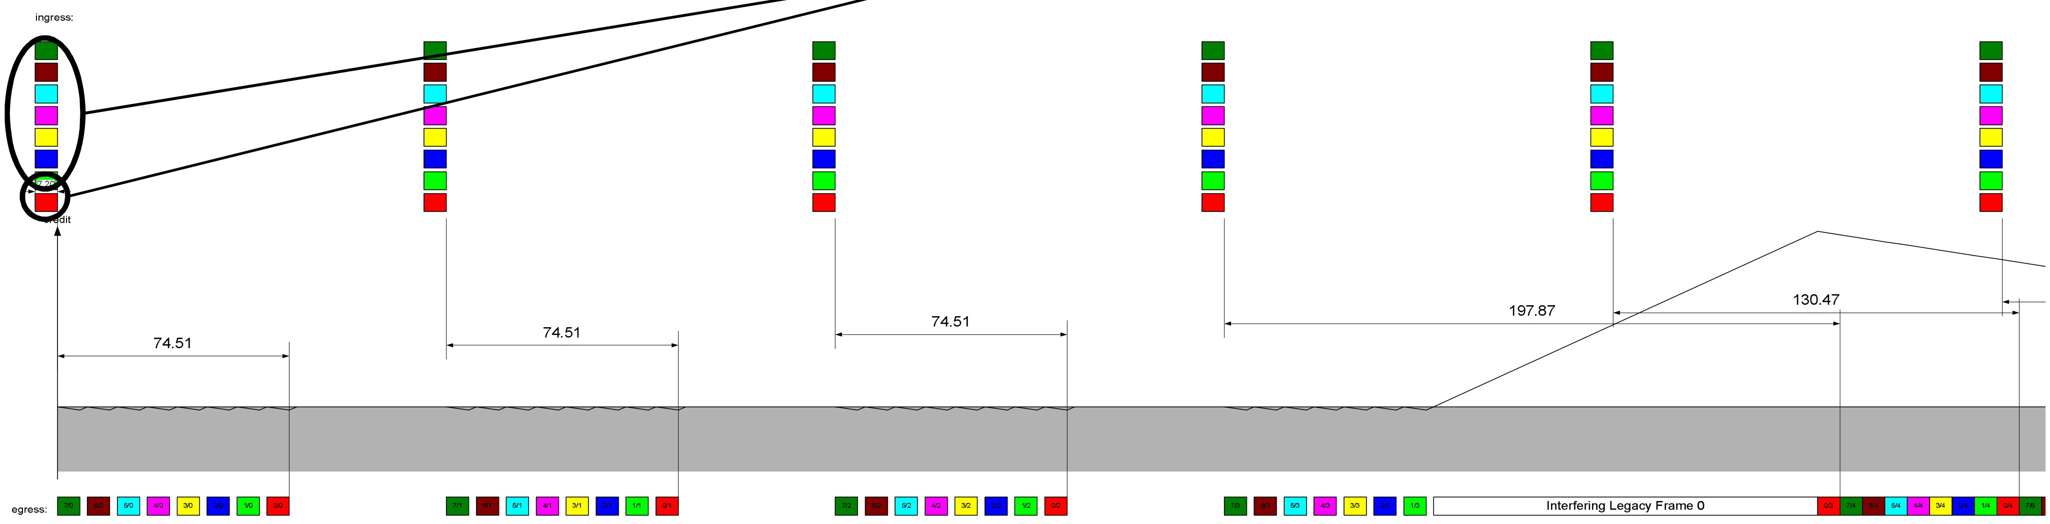

Interfering Talker

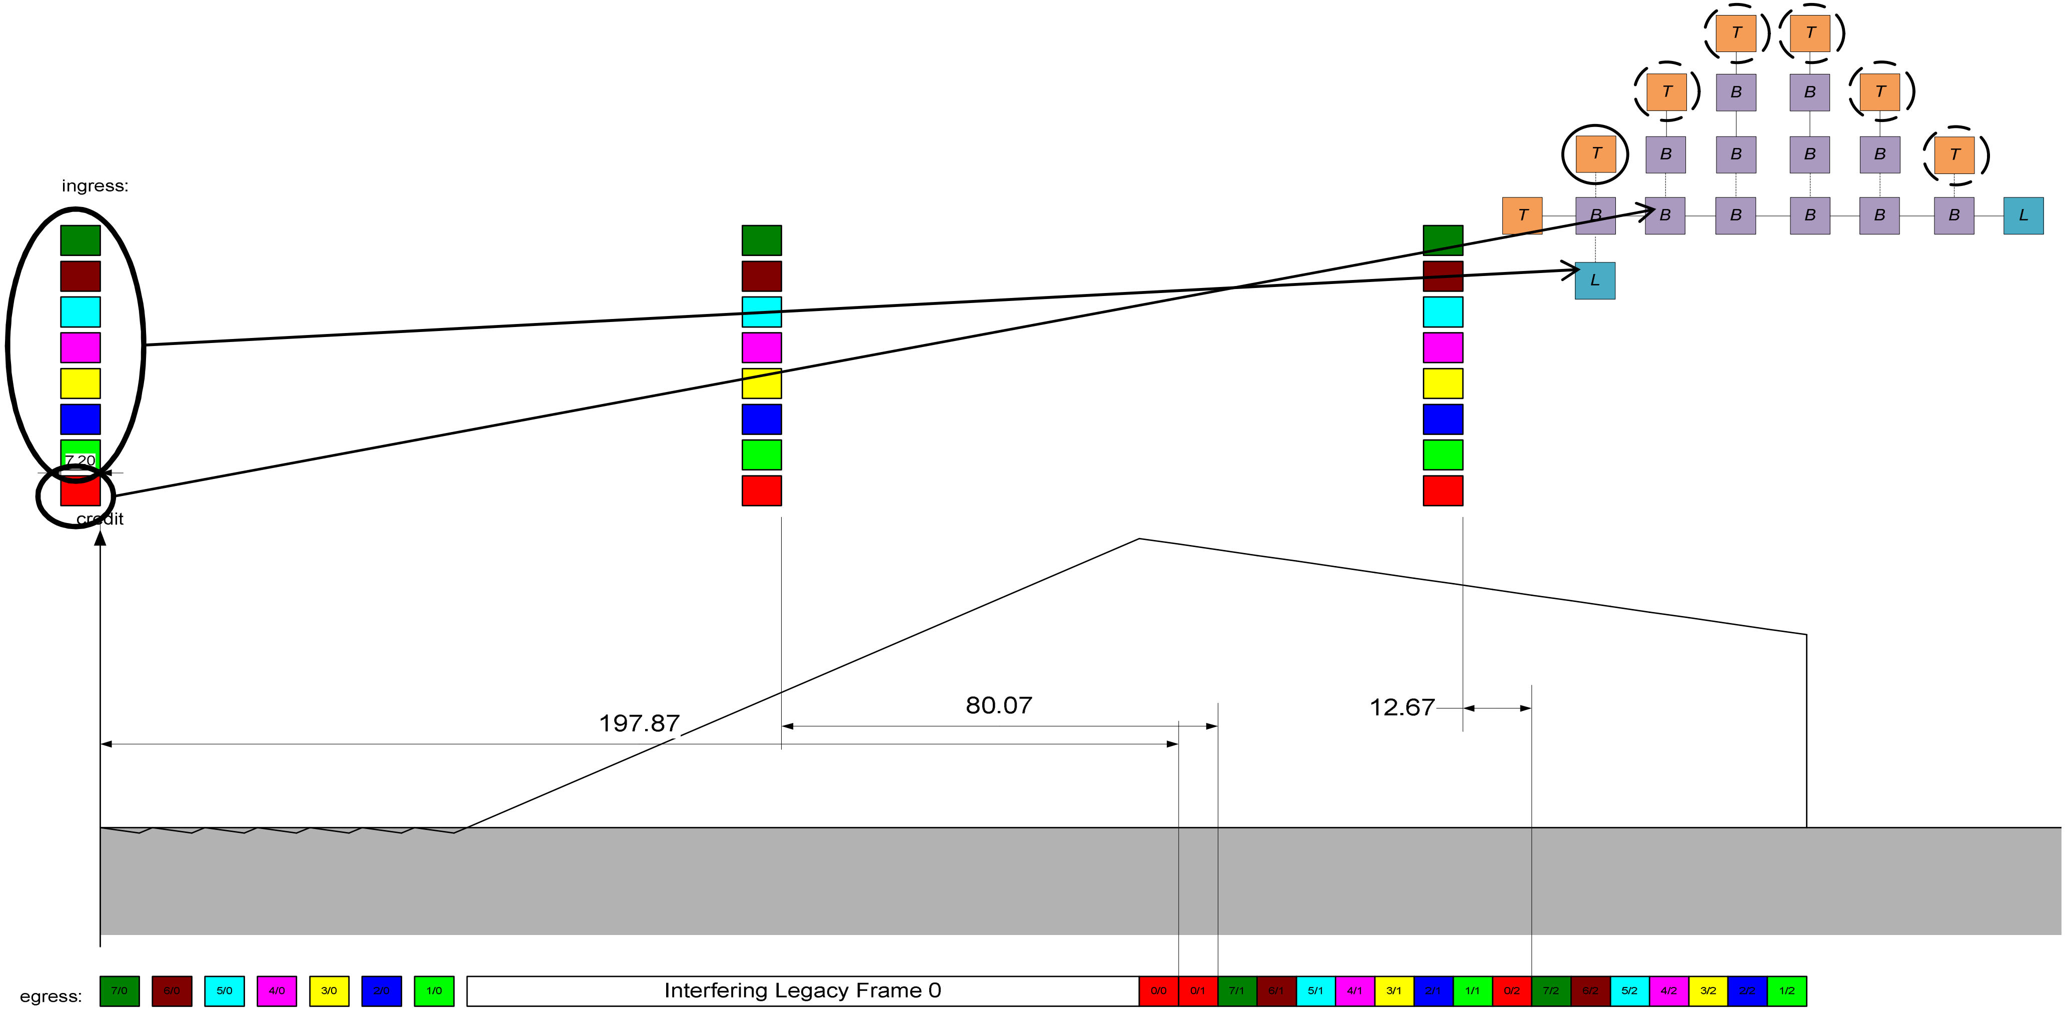

After First Bridge (1)

talker:	197.87µs (245.94µs)
1st bridge:	??? (455.84µs)
total:	??? (701.78µs)

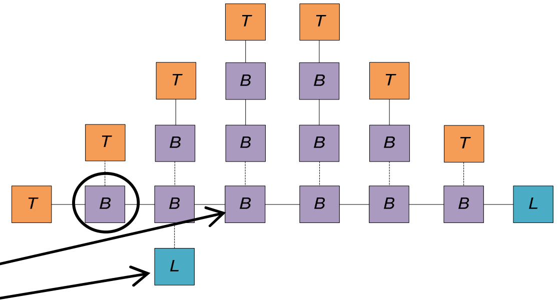

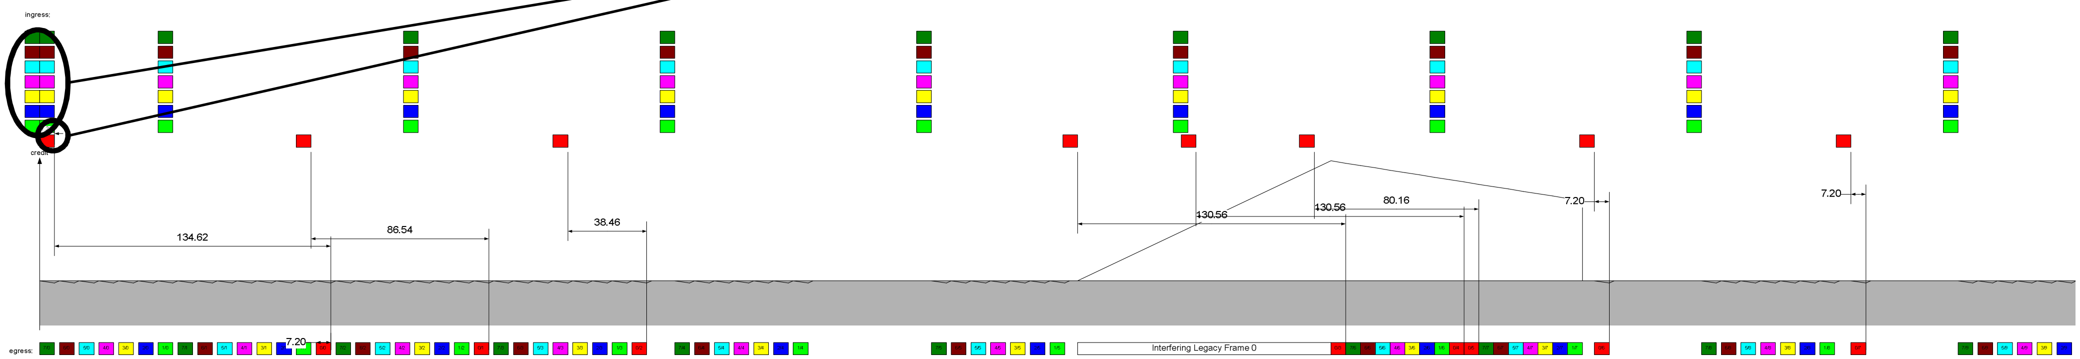

After First Bridge (2)

talker: 197.87µs (245.94µs)
1st bridge: 267.68µs (455.84µs)

total: 465.55µs (701.78µs)

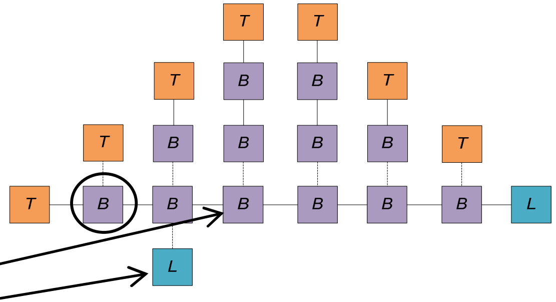

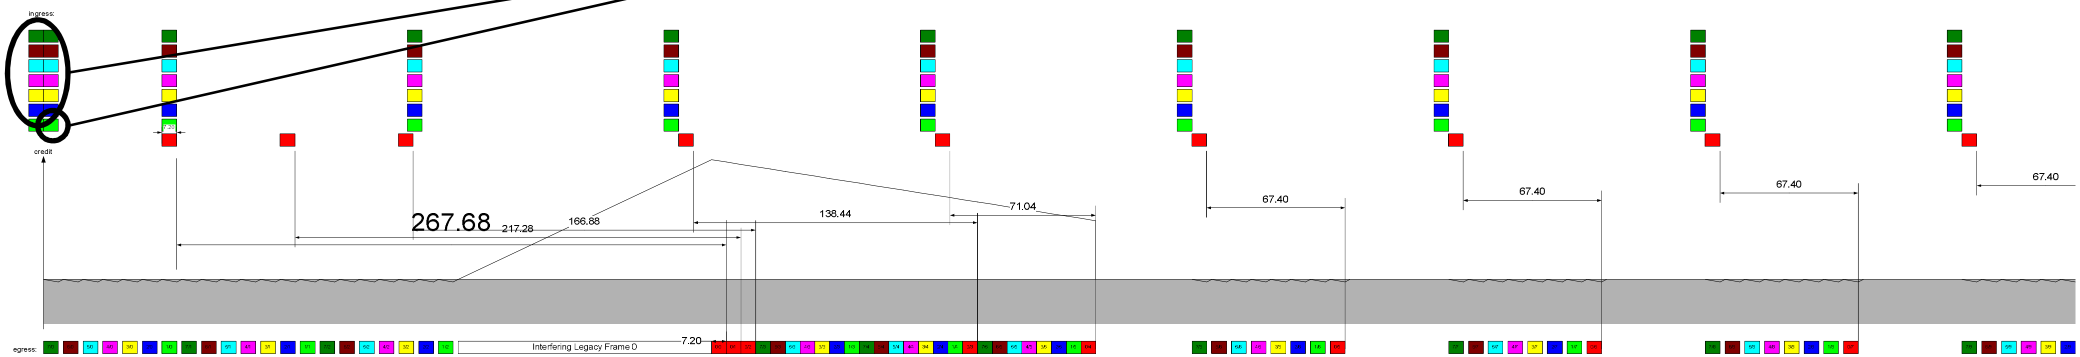

1st Interfering Bridge

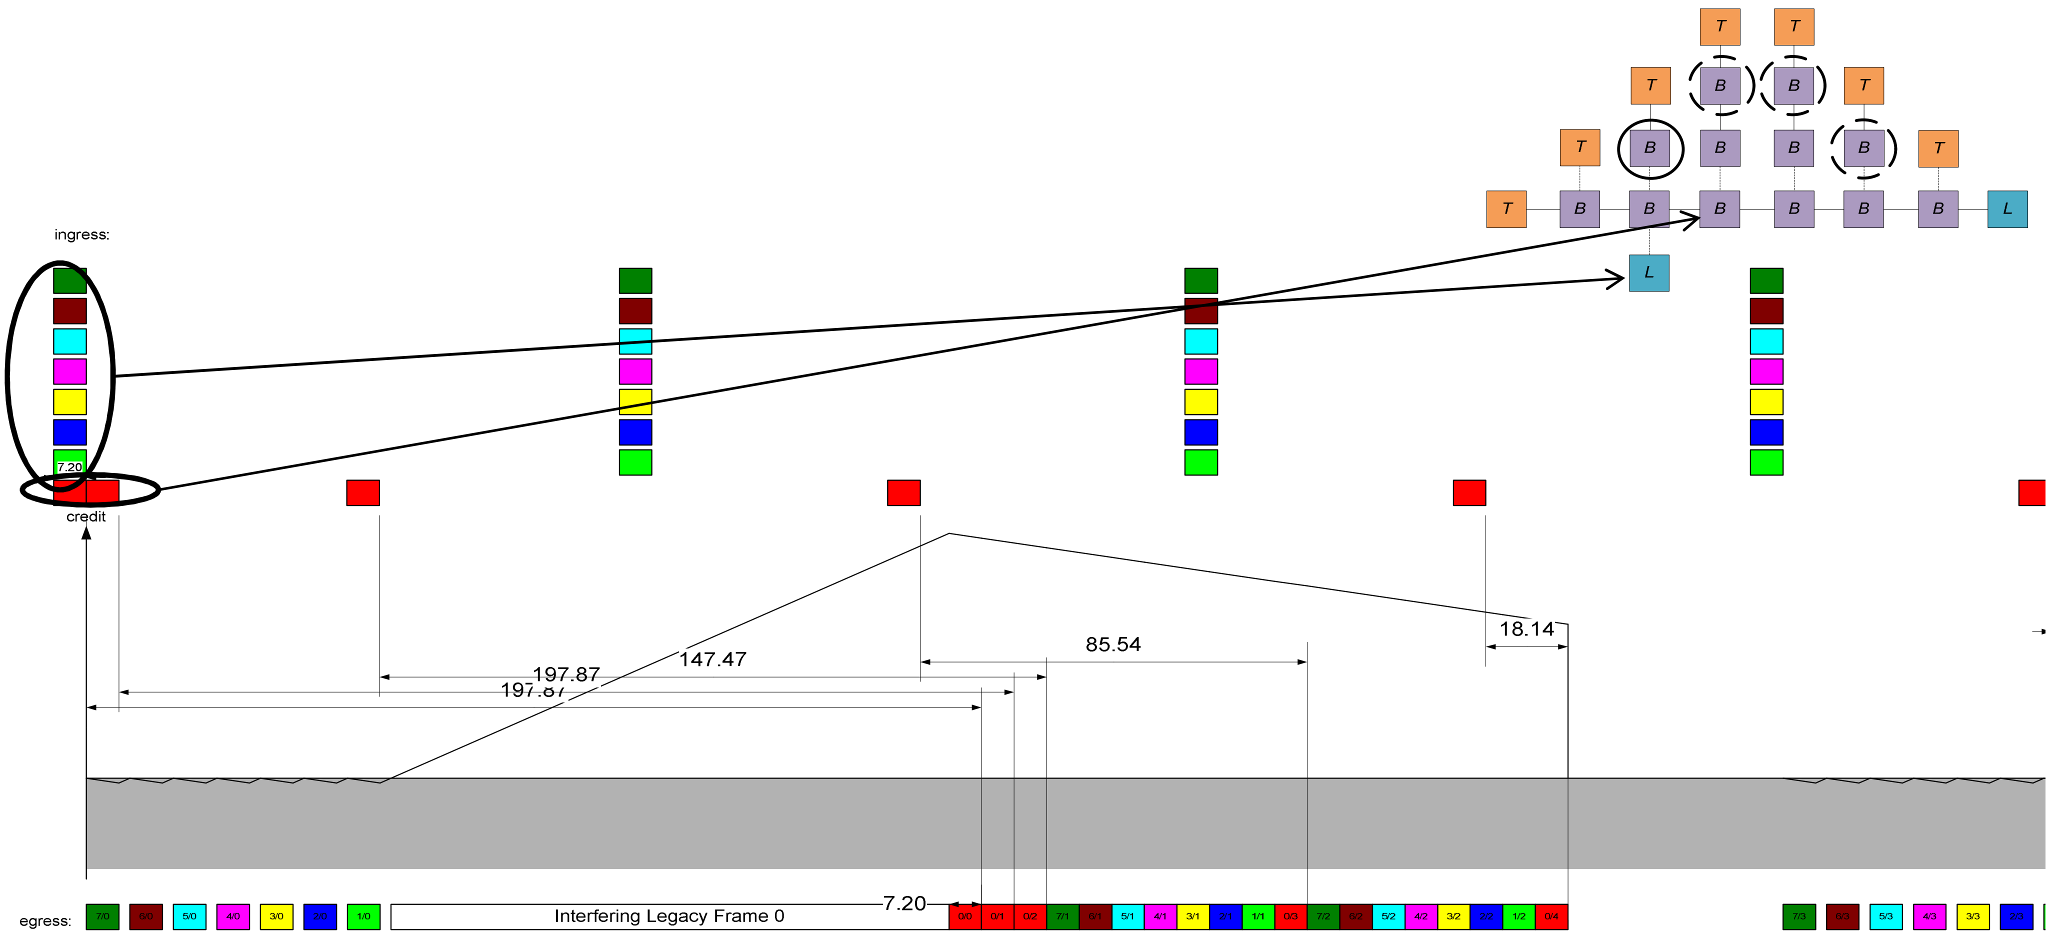

After Second Bridge

talker: 197.87µs (245.94µs)
1st bridge: 267.68µs (455.84µs)
2nd bridge: 337.50µs (542.37µs)

total: 803.05µs (1244.15µs)

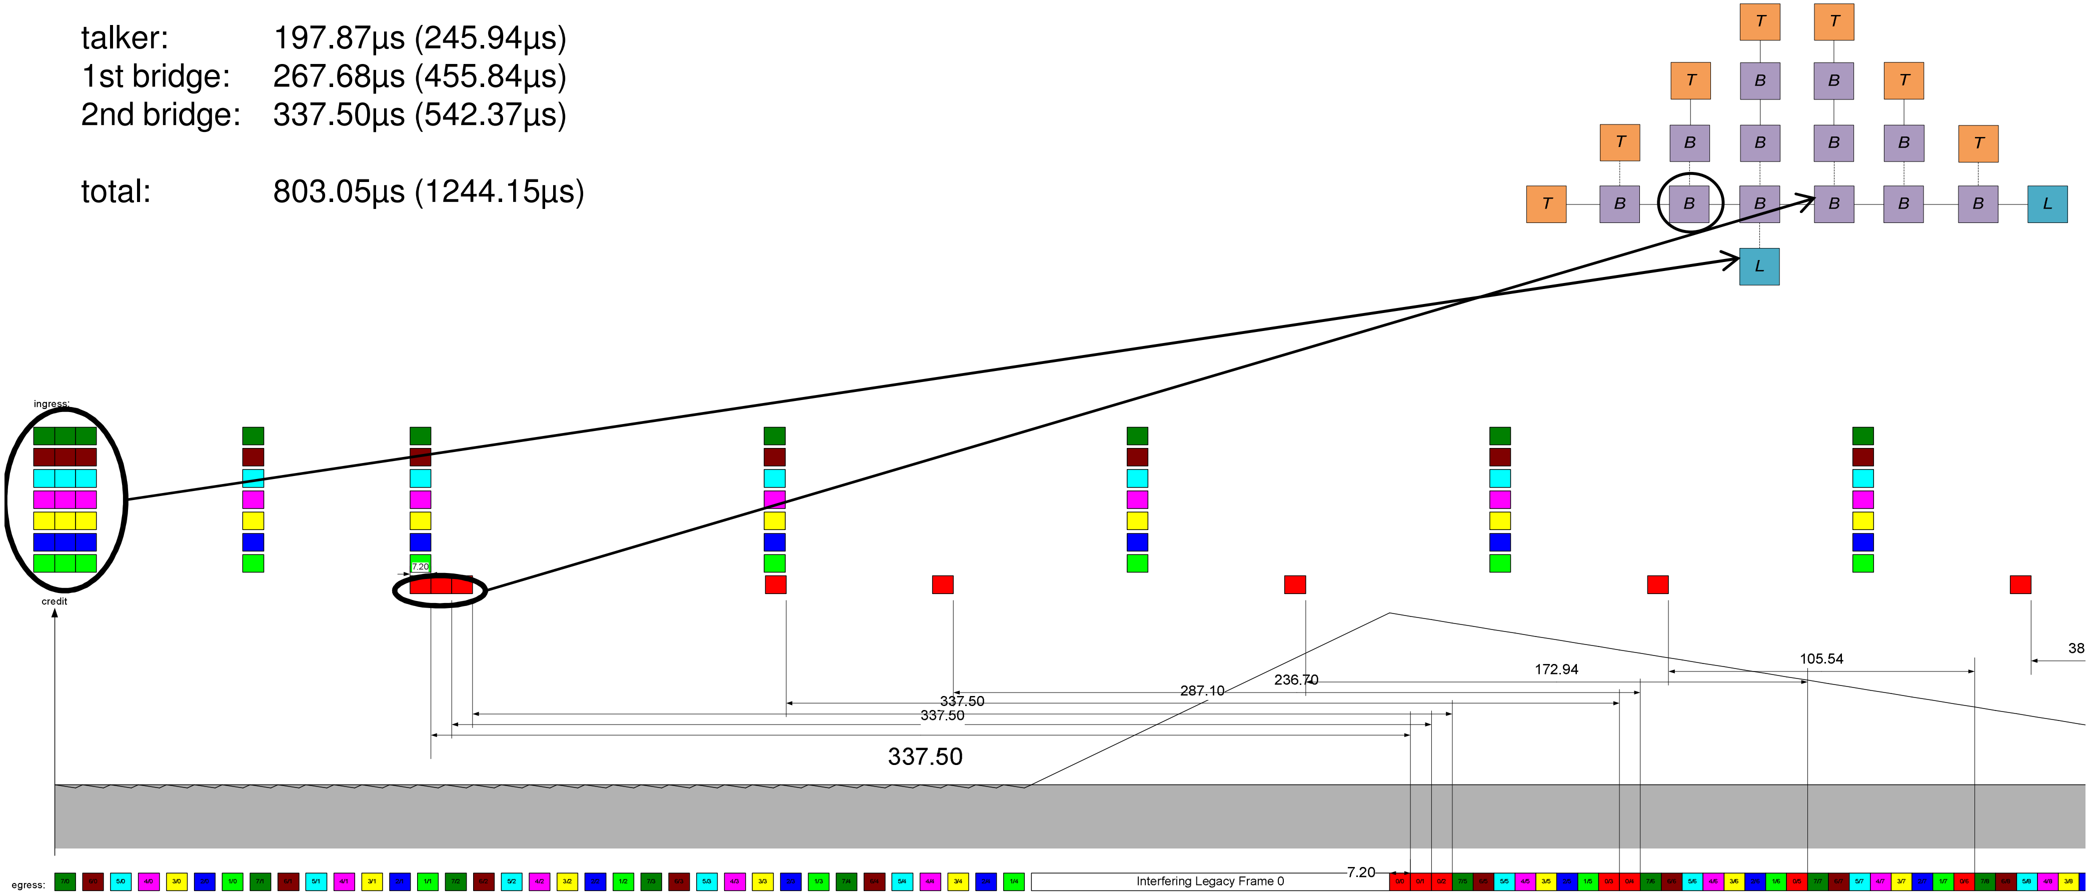

2nd Interfering Bridge

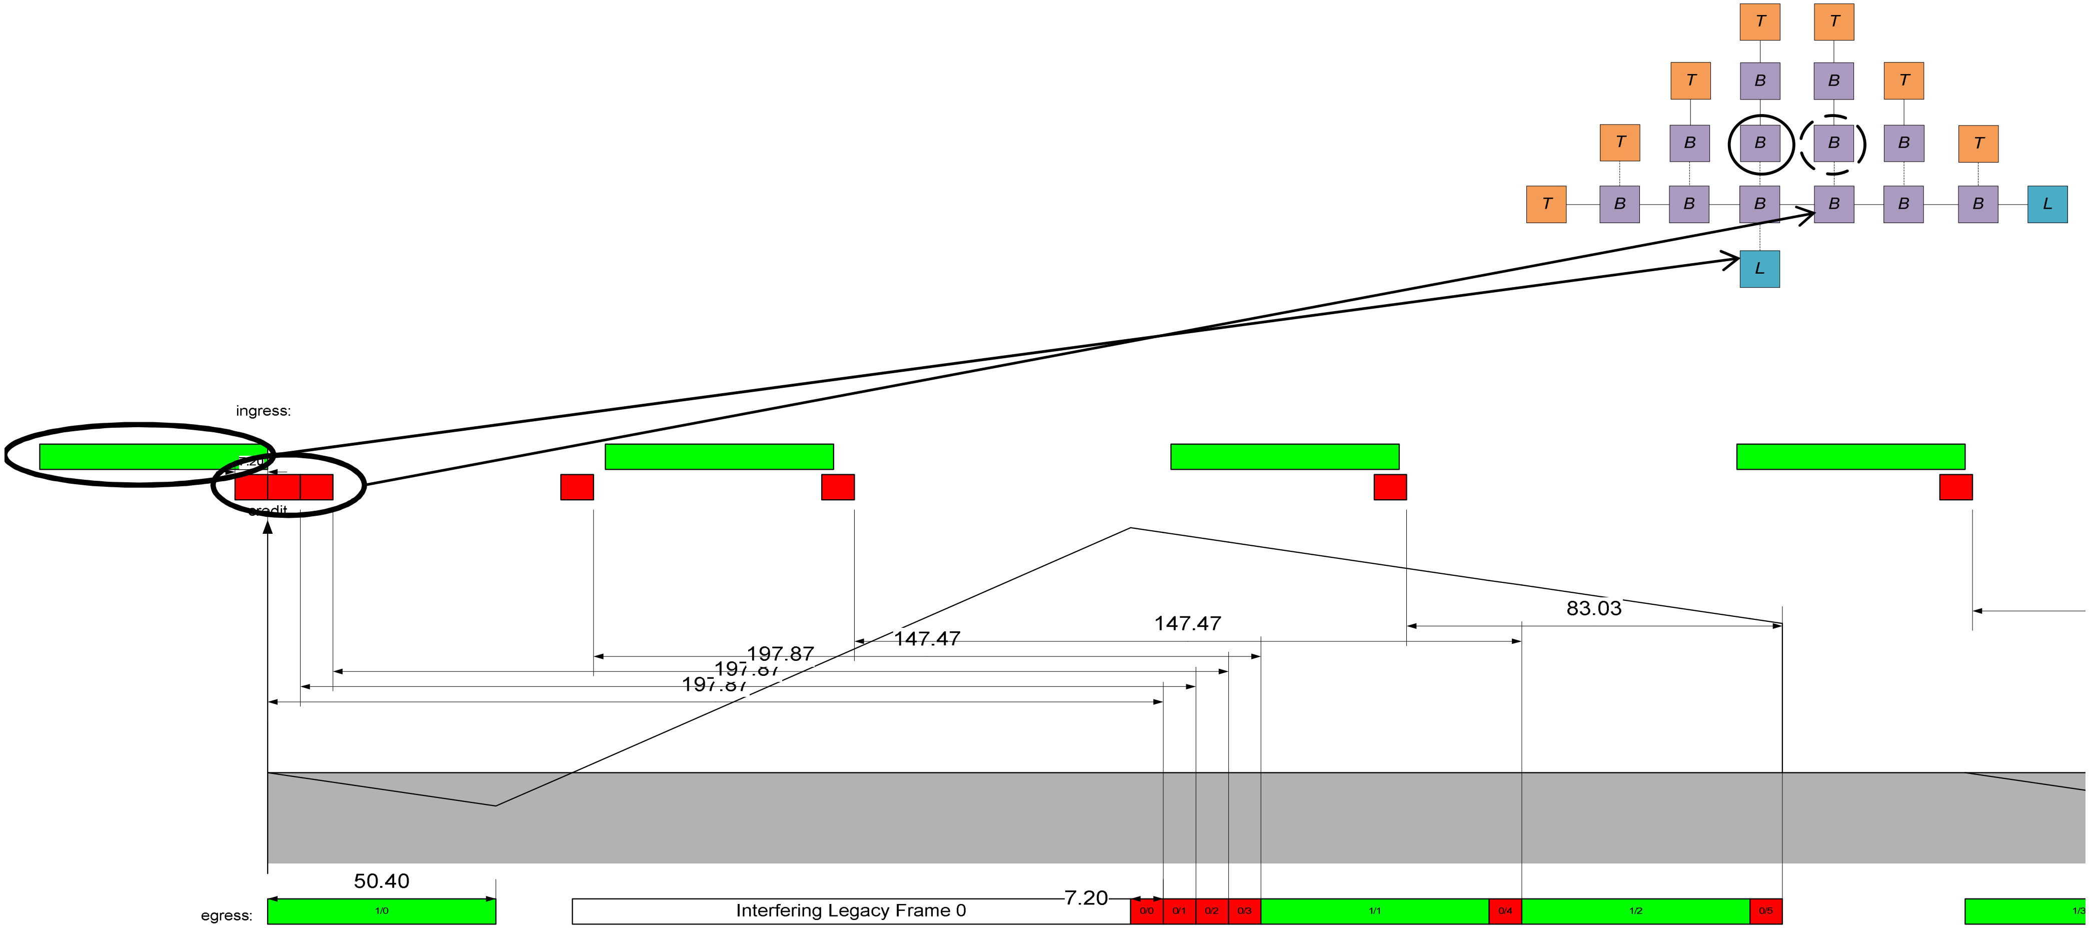

After Third Bridge

talker:	197.87µs (245.94µs)
1st bridge:	267.68µs (455.84µs)
2nd bridge:	337.50µs (542.37µs)
3rd bridge:	397.61µs (628.92µs)
total:	1200.66µs (1873.07µs)

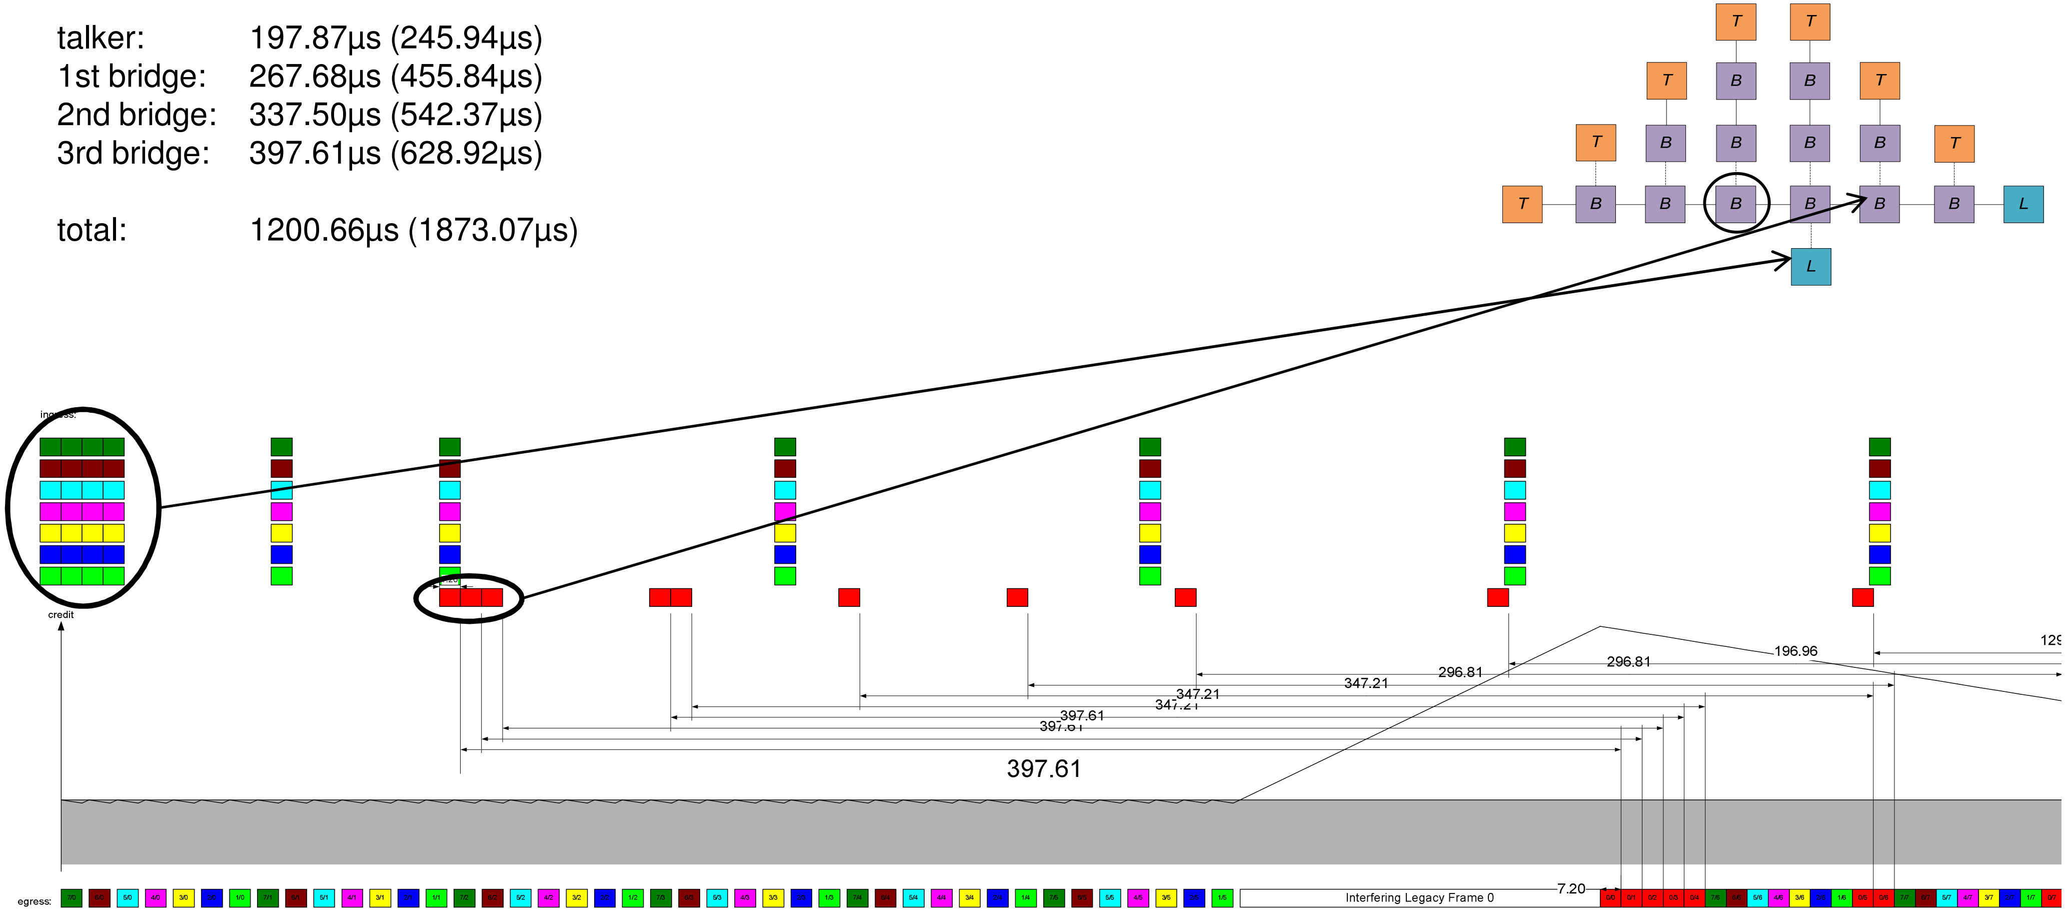

After Fourth Bridge

talker:	197.87µs (245.94µs)
1st bridge:	267.68µs (455.84µs)
2nd bridge:	337.50µs (542.37µs)
3rd bridge:	397.61µs (628.92µs)
4th bridge:	397.61µs (609.68µs)
total:	1598.27µs (2482.75µs)

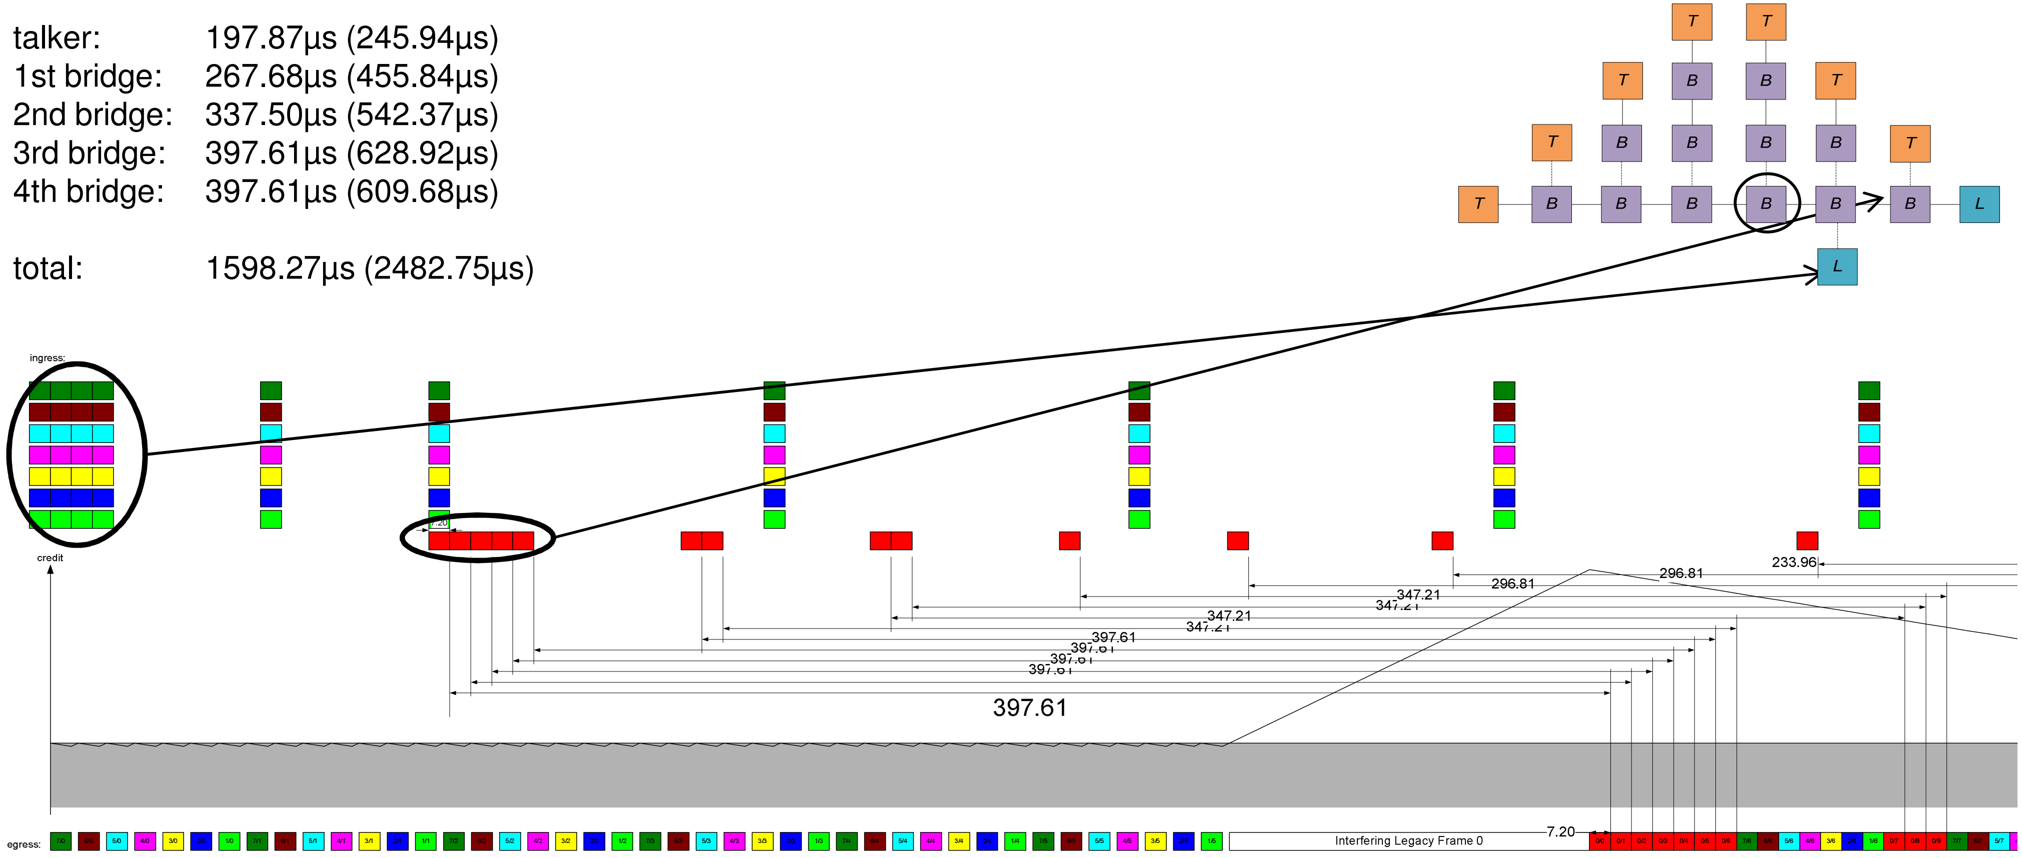

After Fifth Bridge

talker:	197.87µs (245.94µs)
1st bridge:	267.68µs (455.84µs)
2nd bridge:	337.50µs (542.37µs)
3rd bridge:	397.61µs (628.92µs)
4th bridge:	397.61µs (609.68µs)
5th bridge:	337.50µs (484.69µs)
total:	1935.77µs (2967.44µs)

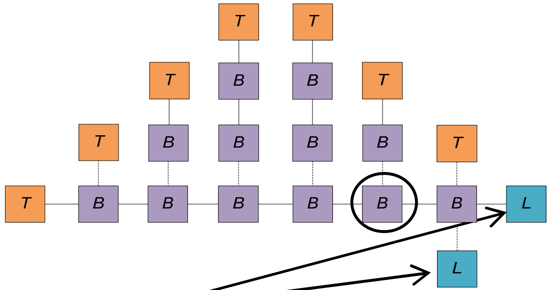

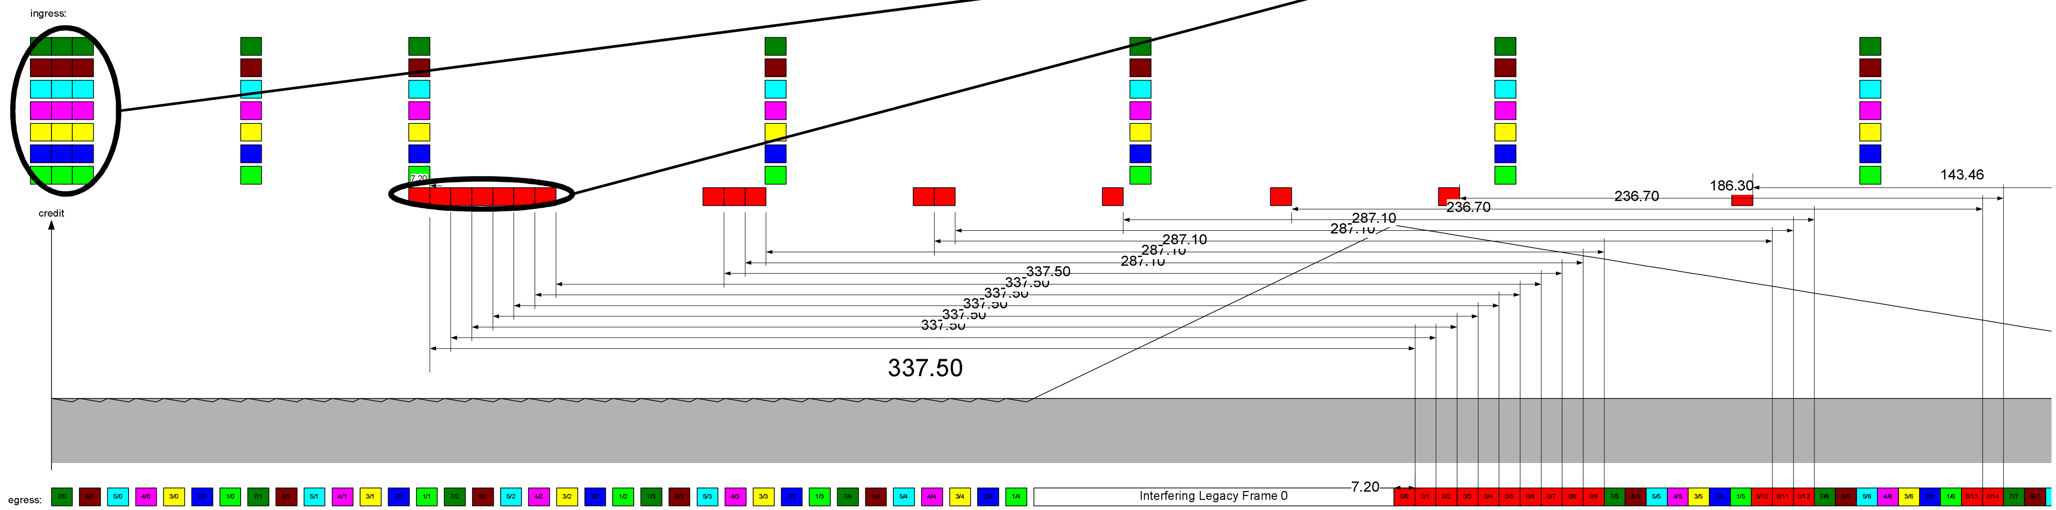

Hochschule
Deggendorf

After Sixth Bridge

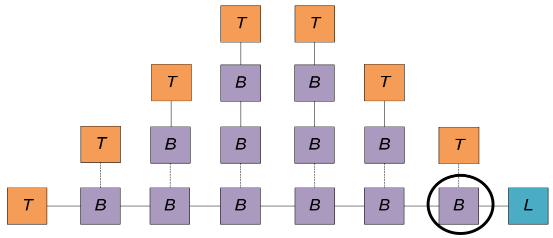

talker:	197.87µs (245.94µs)
1st bridge:	267.68µs (455.84µs)
2nd bridge:	337.50µs (542.37µs)
3rd bridge:	397.61µs (628.92µs)
4th bridge:	397.61µs (609.68µs)
5th bridge:	337.50µs (484.69µs)
6th bridge:	267.68µs (359.68µs)
total:	2203.45µs (3327.12)

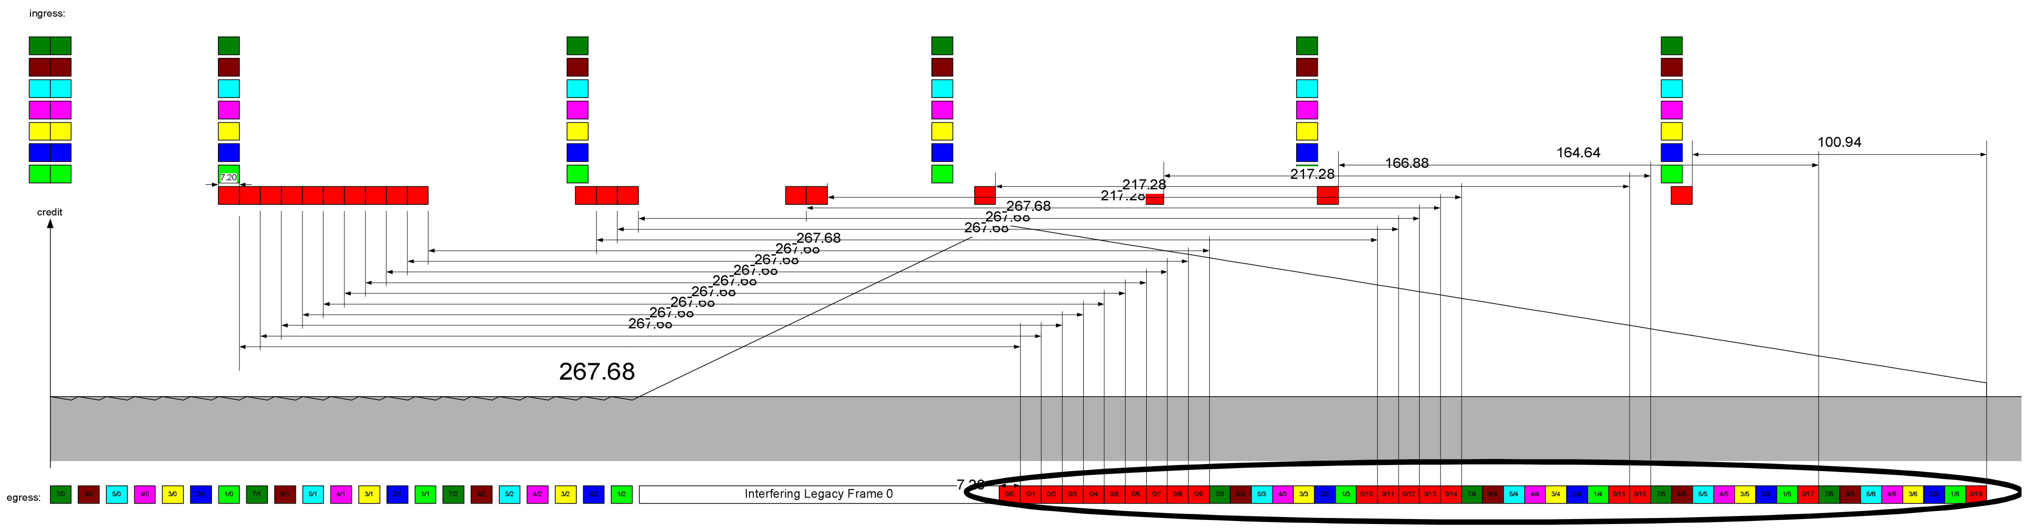

Bursts are more likely than with the original shaper

Results

- The speeded up shaper makes bursts more likely.

- It seems that the shaper allows the creation of bigger bursts in a network with the same topology compared to the current shaper. But this is not correct. It is possible that a talker with the current shaper generates a burst of three frames. Also bridges can generate a bigger burst.

- Limiting the network diameter does not prevent from very high per hop latencies (>600µs).

⇒ Limiting the topology limits the latency, but without other additional limitations the worst case latency is still very high

→ **Limiting the diameter of a network is not feasible for industrial networks especially to such "small" numbers like 7 hops**

Gigabit Ethernet

Speeded Up Shaper + Max Class A Bandwidth = 20%

Talker

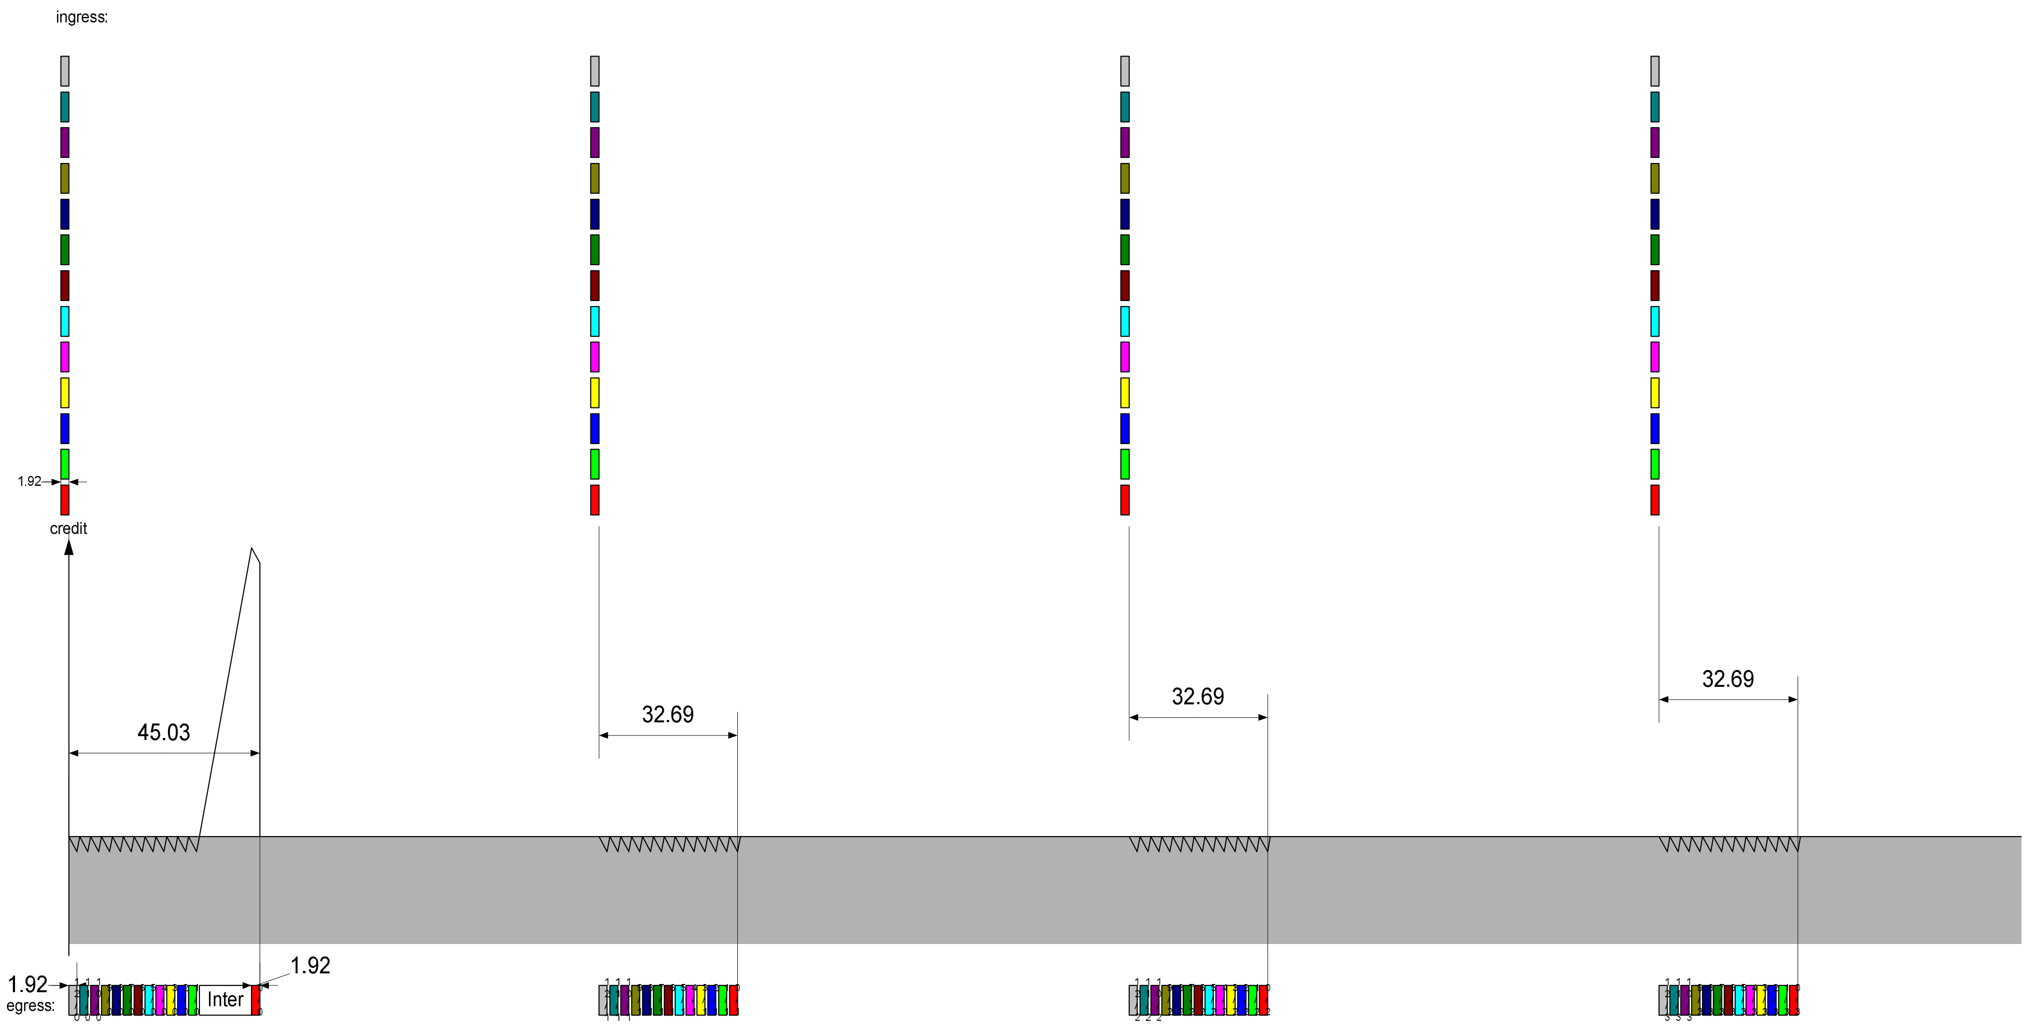

Creating a Burst – First Bridge After Talker

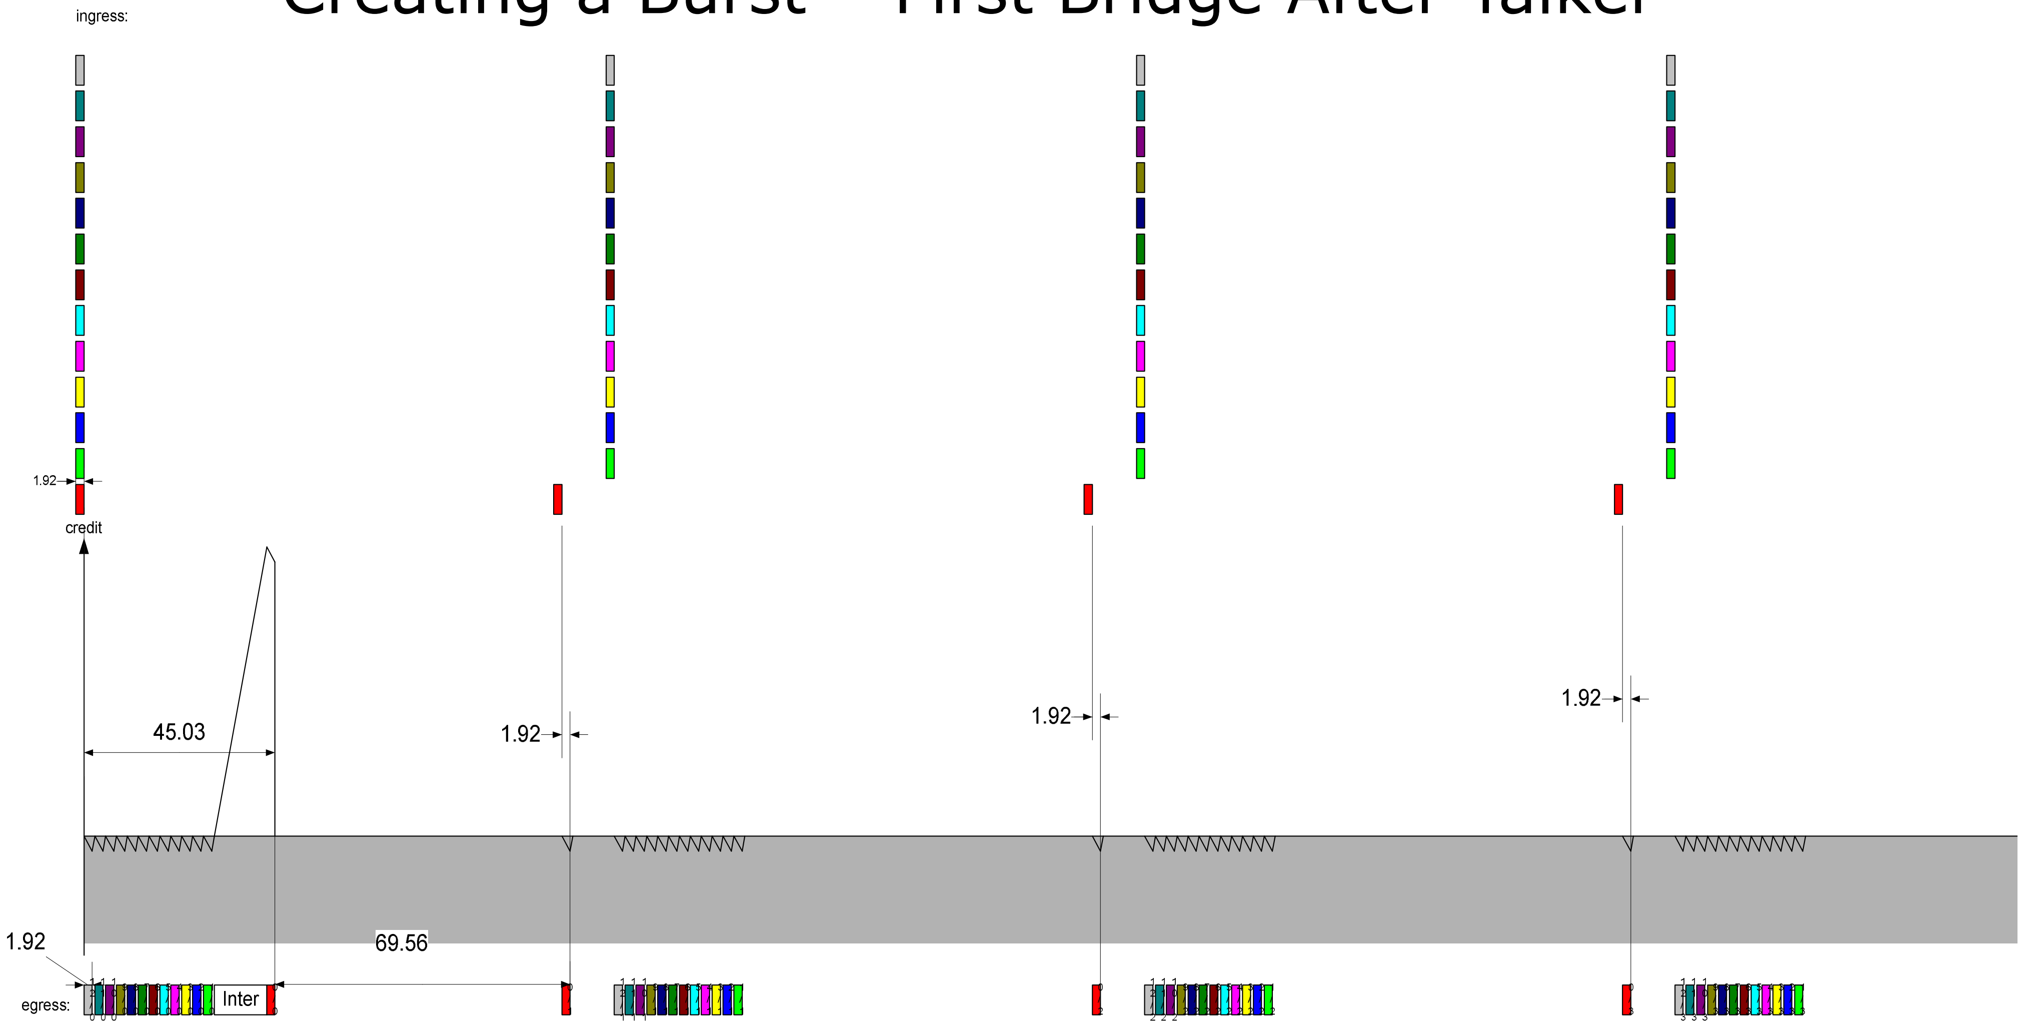

Creating a Burst – Second Bridge After Talker

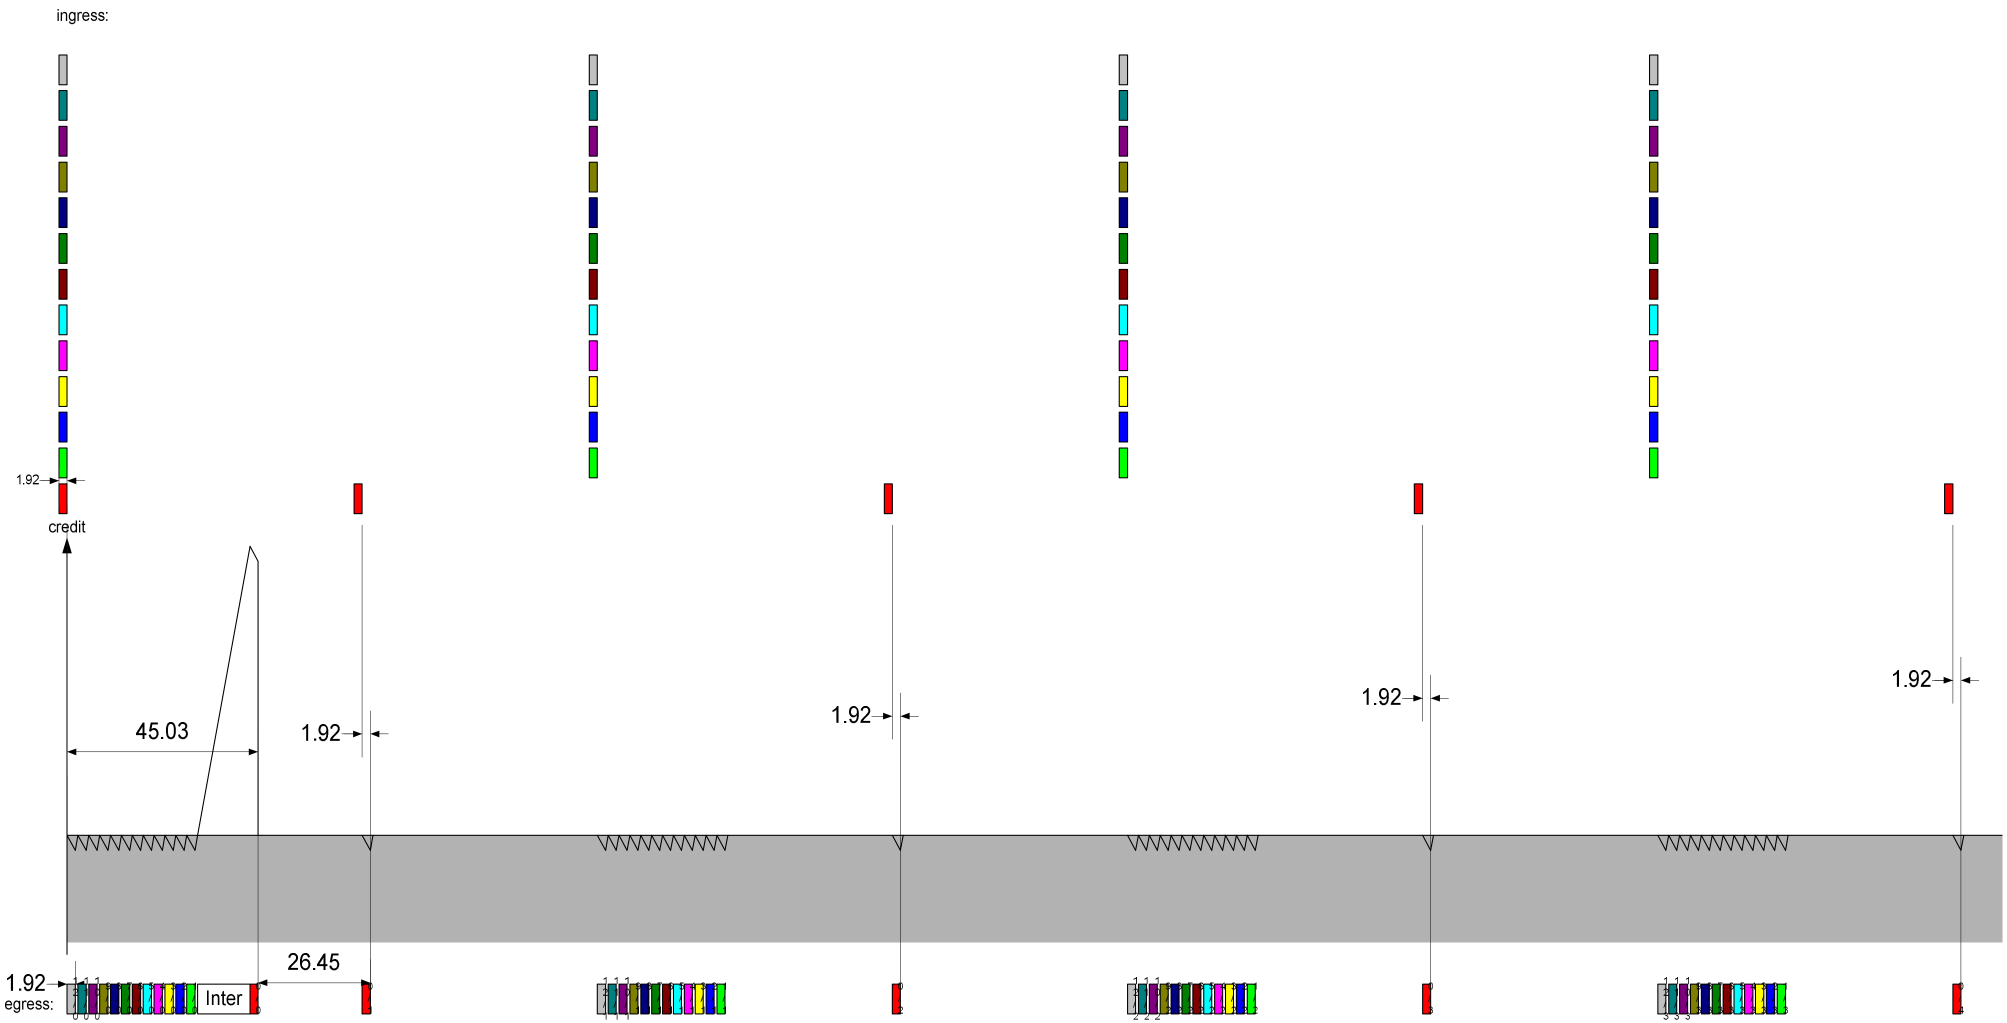

Creating a Burst – Third Bridge After Talker

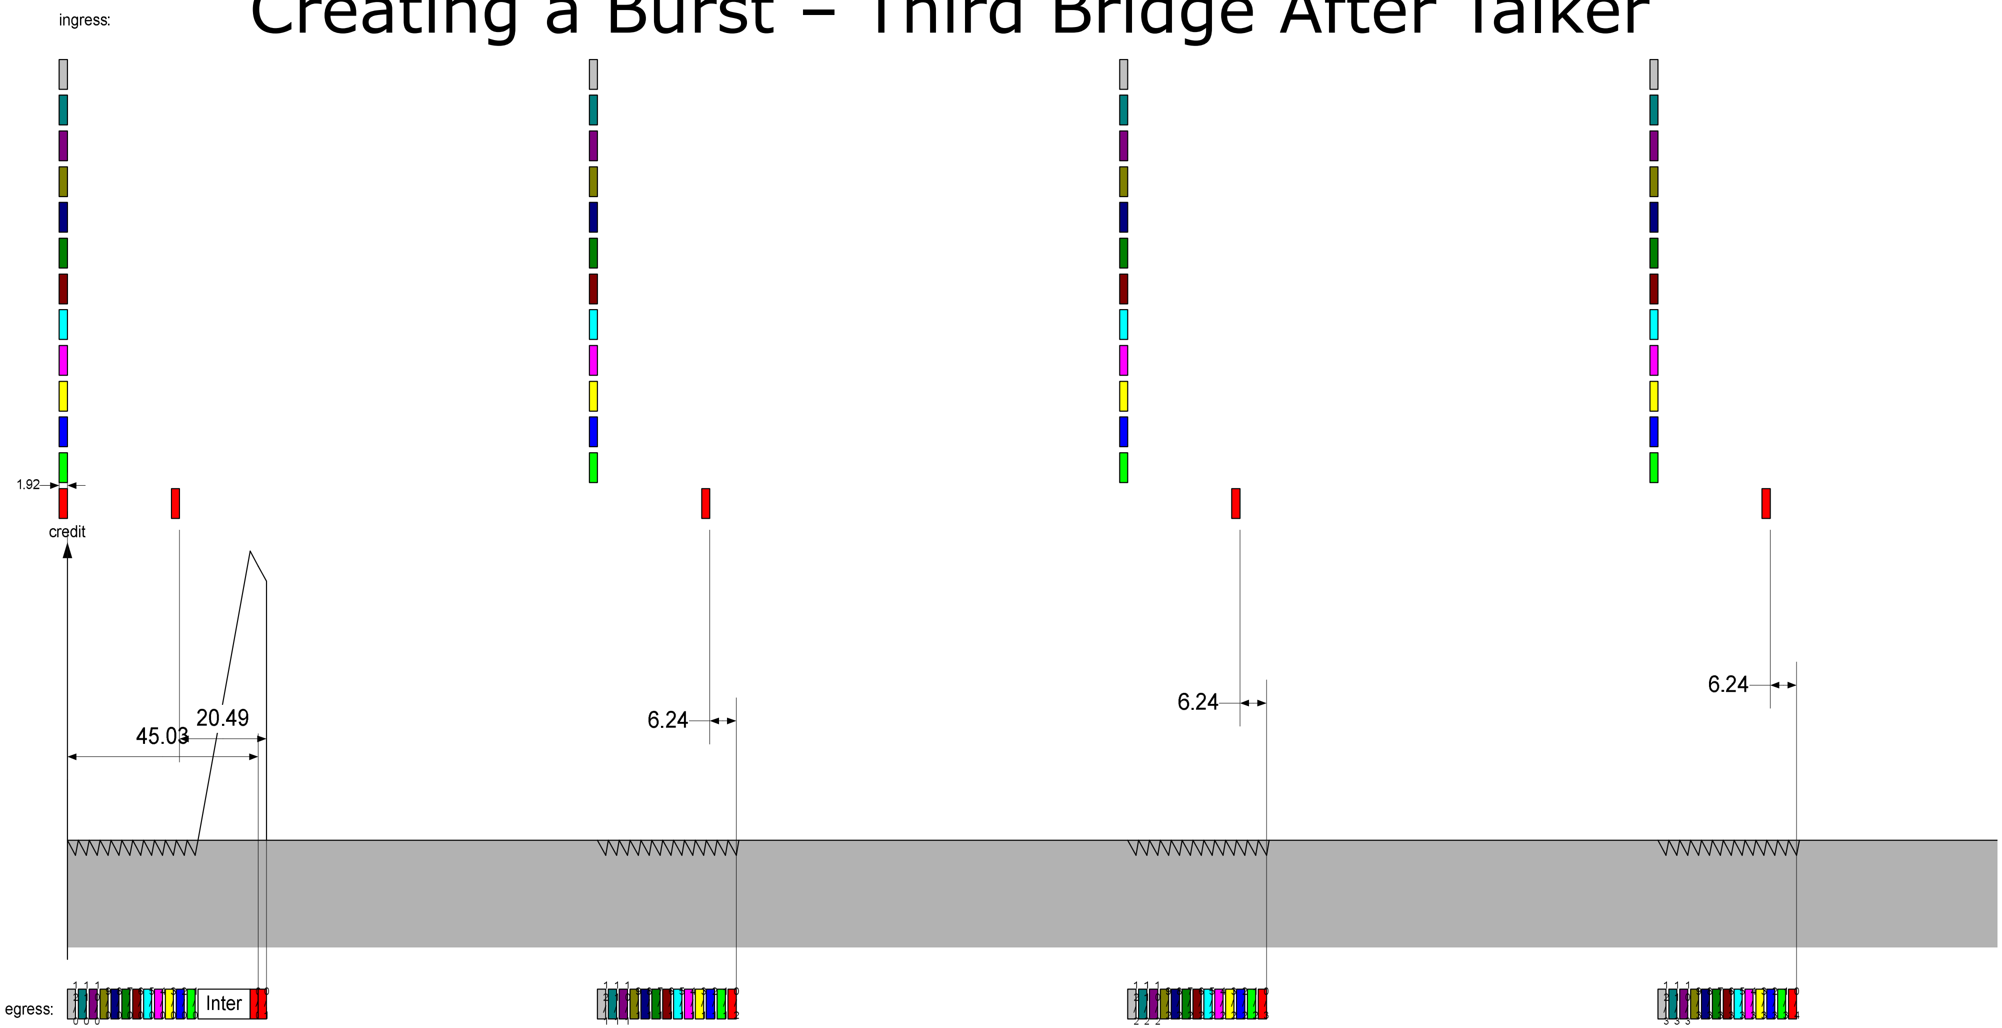

Consequences of Bursts – Burst of Two

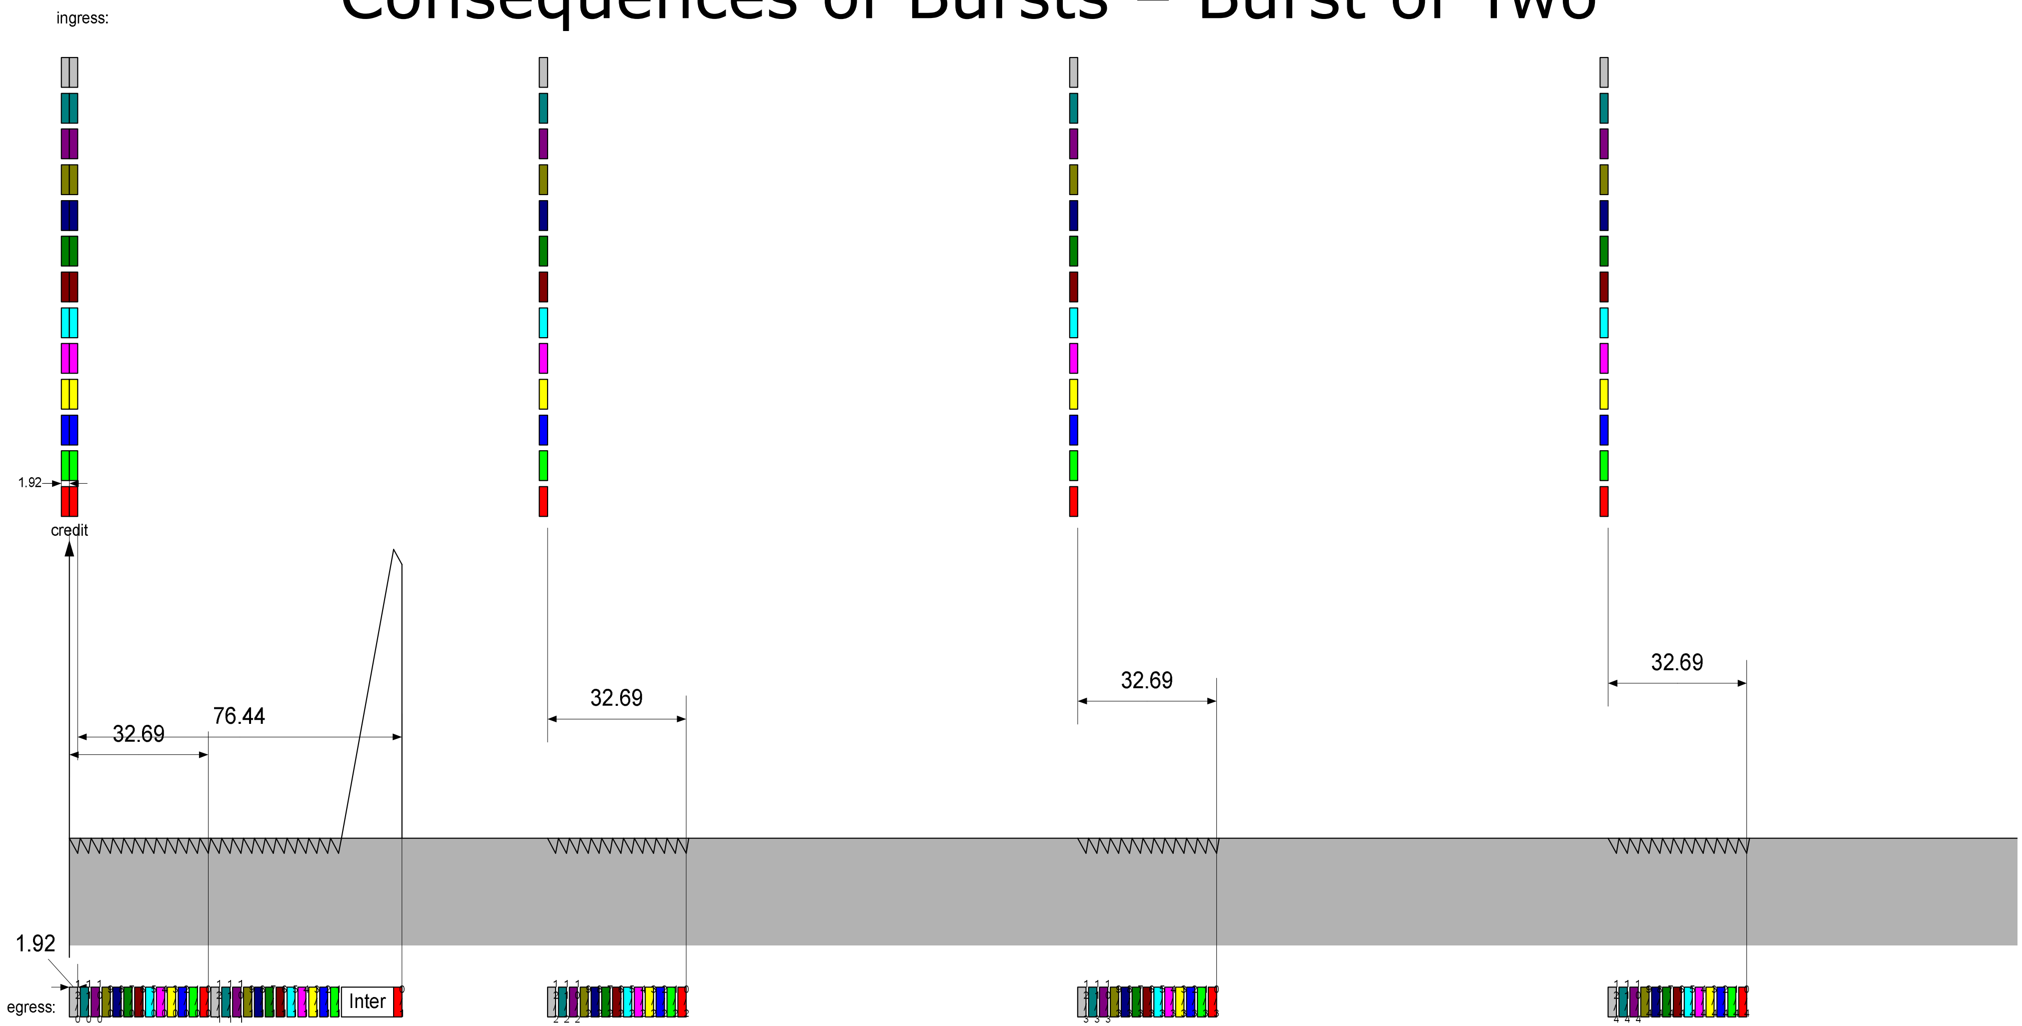

Creating a Burst – 10th Bridge After Talker

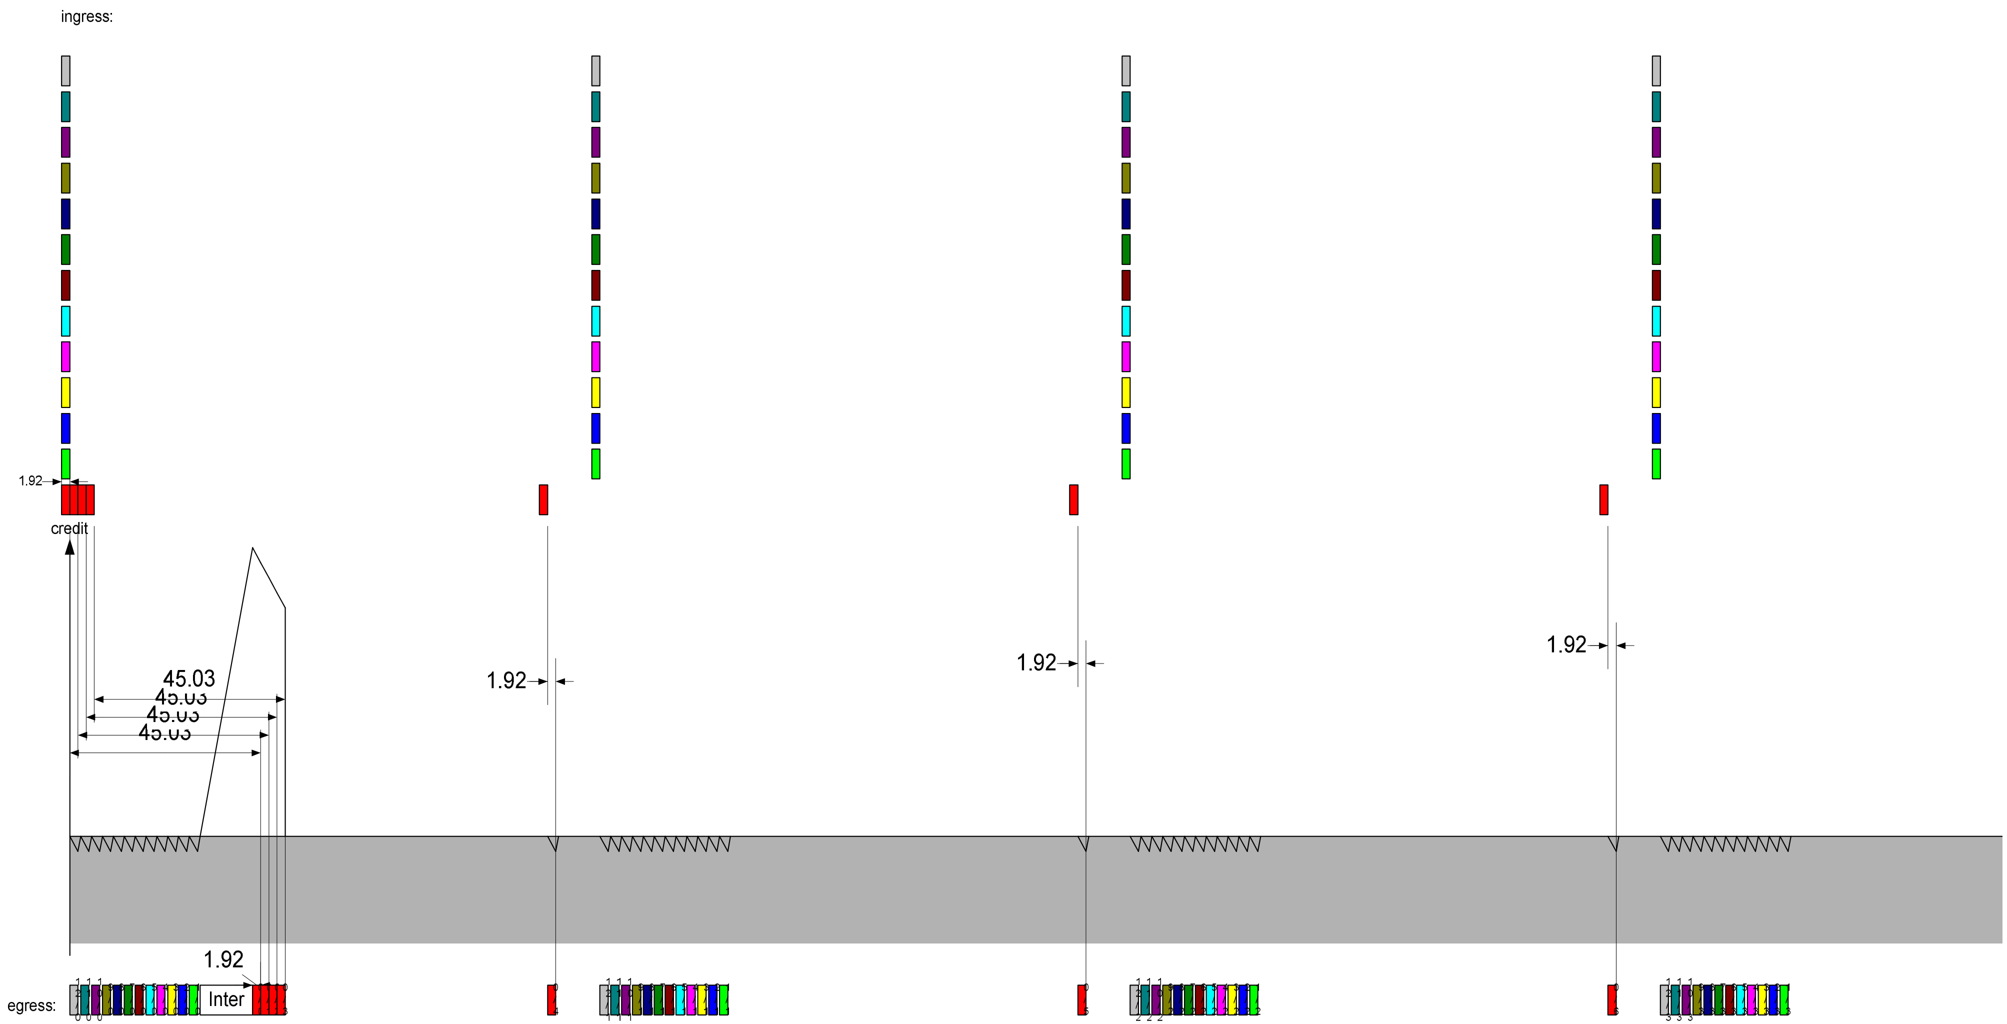

Consequences of Bursts – Burst of Four

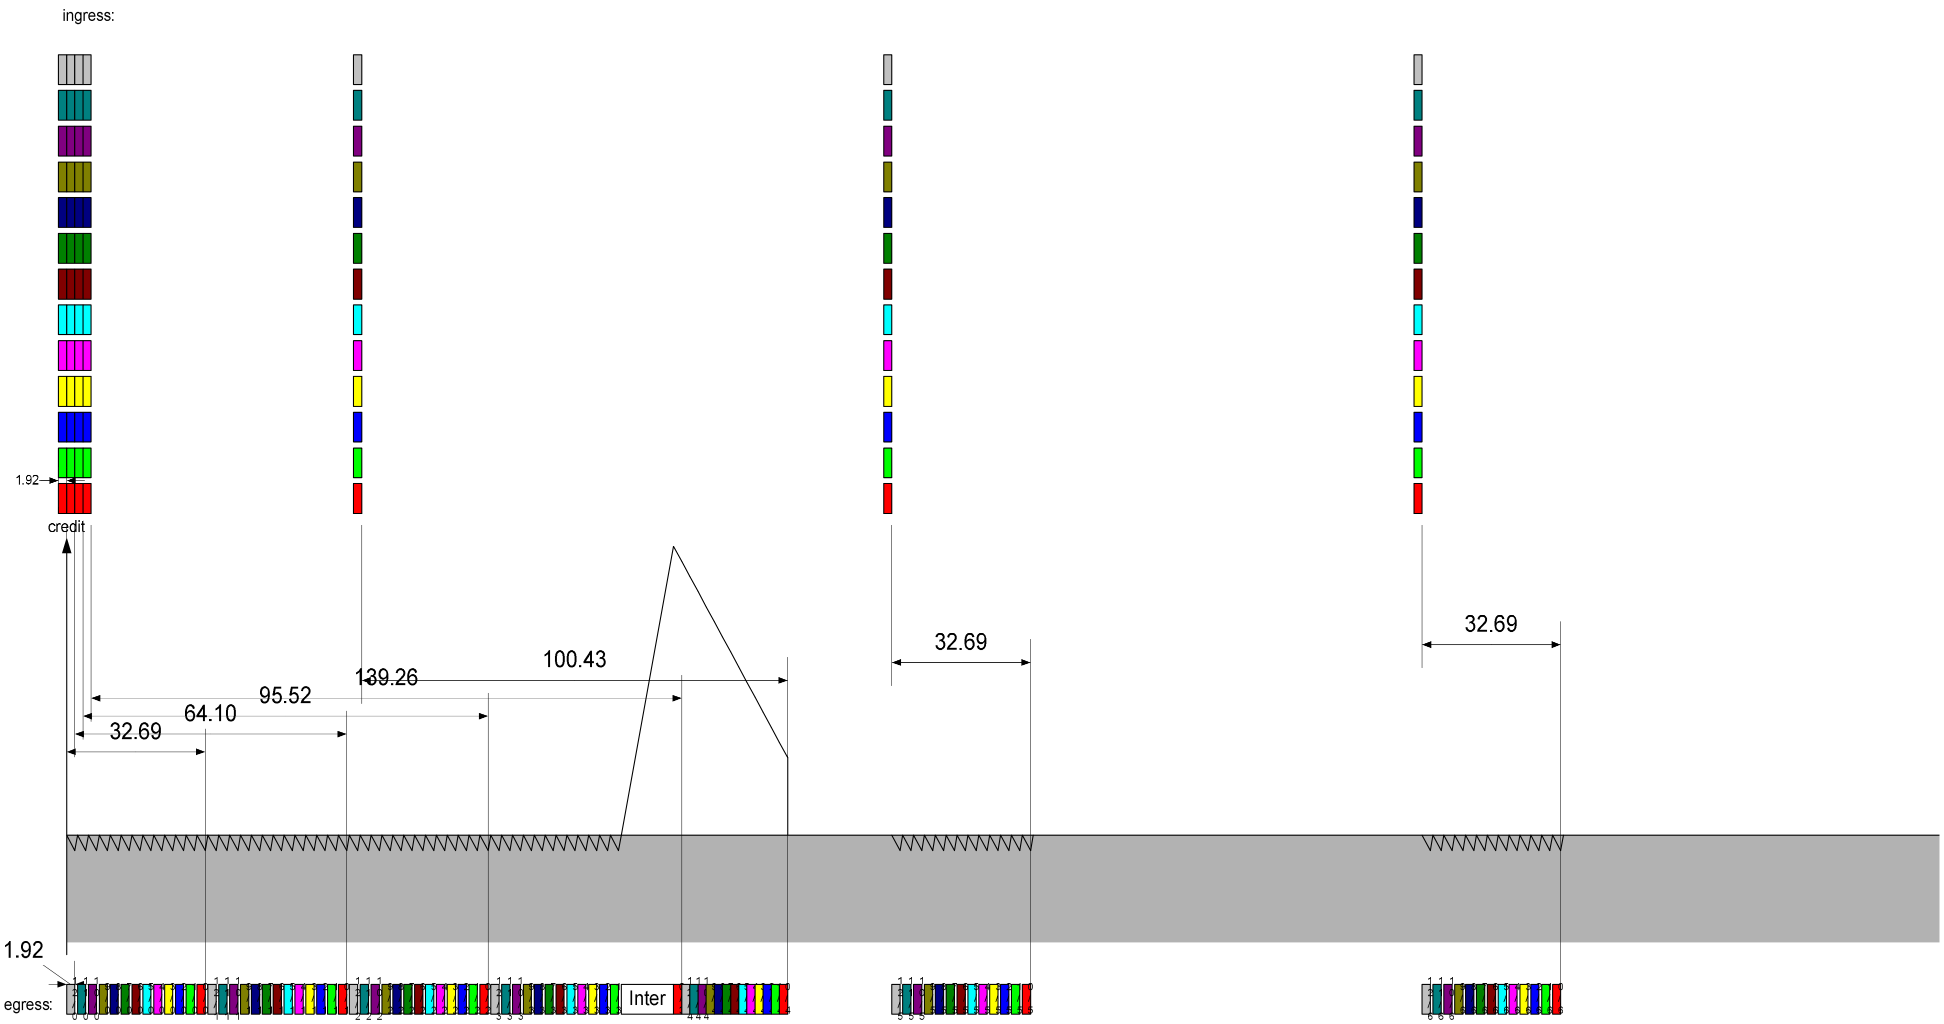

Results

- The speeded up shaper does not provide a topology independent per hop worst case latency.
- It takes several hops to push packets together.
  → Smaller bursts in big topologies
- Effects of bursts are limited to a short period of time (a few measurement intervals).
- Much interfering traffic is necessary to generate a burst in GigE (similar to current shaper).
- With the current shaper it is possible to push more than 2 packets together at one hop (see first slides of this presentation). With this shaper this is much more unlikely.
- If the allowed bandwidth is further reduced the latency gets even better.
  → configurable bandwidth boundary might be interesting
- The worst case latency is significantly smaller than in GigE networks with the current shaper.

Thank You