On Standardization of Cut-Through Forwarding (CTF)

Johannes Specht

Preamble

- **This presentation collects technical work on cut-through forwarding (CTF).**
- **It is in support to:**
  - **Reach a common view in IEEE 802.1 amongst goals, needs, and operation of potential standardization activities around CTF.**
  - **Discuss potential standardization of CTF with other IEEE 802 working groups.**
- **Contributions by Johannes Specht to this presentation are individual contributions.**

Table of Contents

Proposal:
Focus subsequent technical work.

Proposal:
Reach a common point of view where and how to address these.

Proposal:
On contents, and distribution.

Where CTF Doesn’t Matter

Short overview, subsequently excluded

(Relatively) Slow Bridges

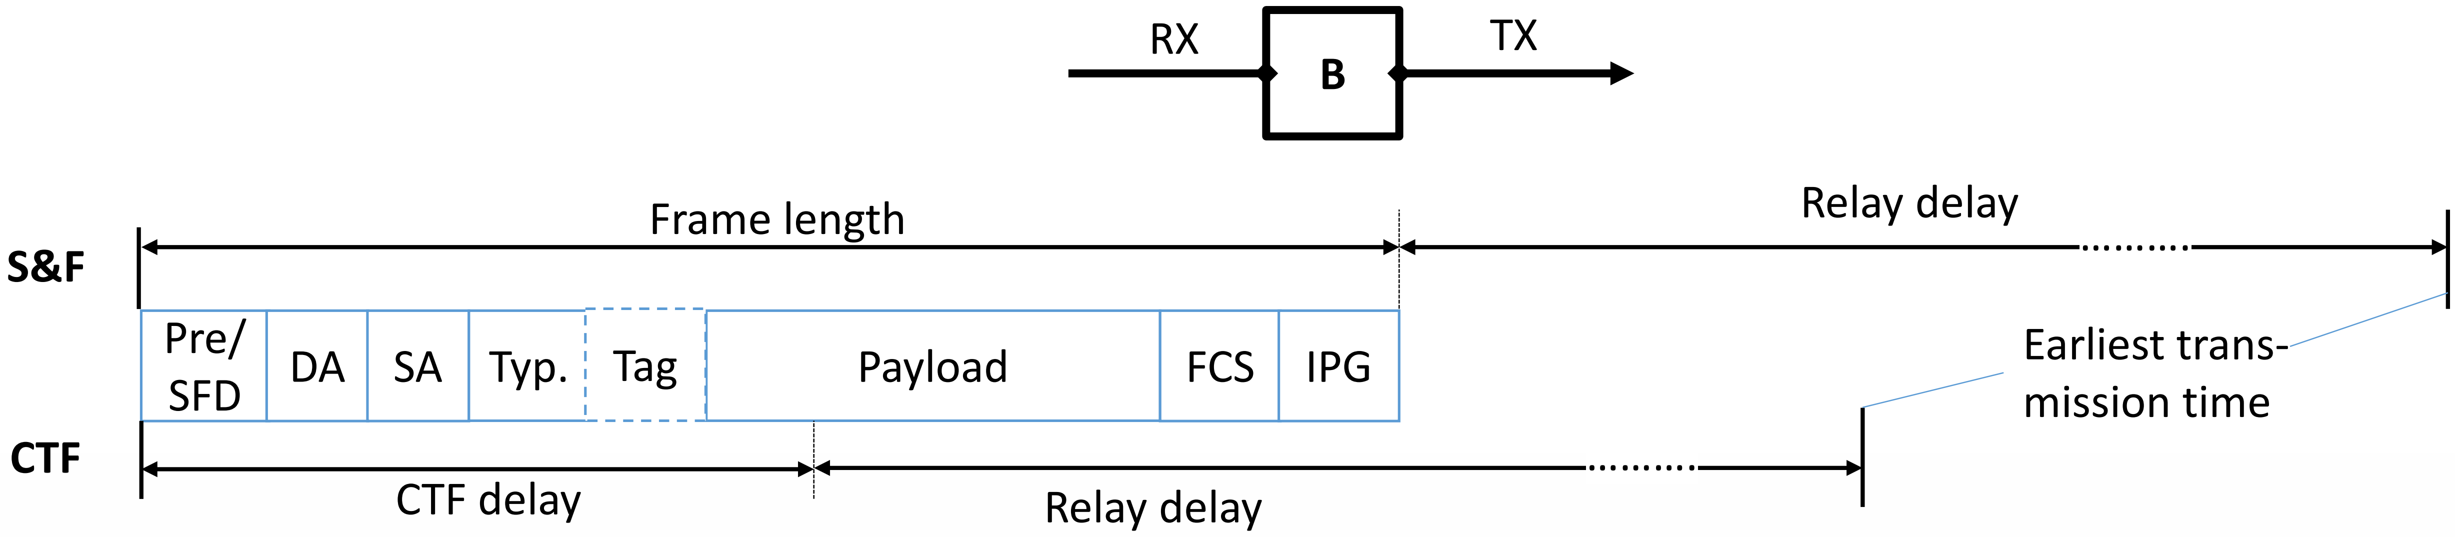

- Bridge-internal relay for both, S&F and CTF traffic, out-weights CTF delay savings
- Link speed relationship - For example:
  - At coast-to-coast links, propagation delays are several miliseconds.
  - At 100GBit/s, these delays can be magnitudes higher than the other relevant quantities (frame lengths, etc.)

Fan-in, Link Speed Transition

Large Bridge Fan-In of Asynchronous, Uncoordinated Transmissions

- Delay on a CTF frame X dominated by the sum of CTF frame lengths from all other Ports.

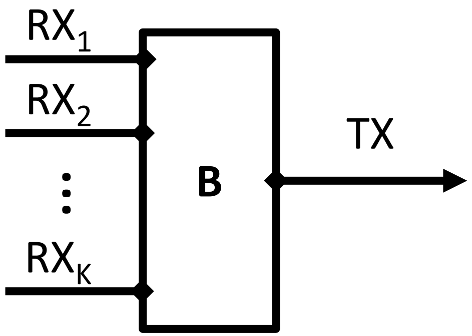

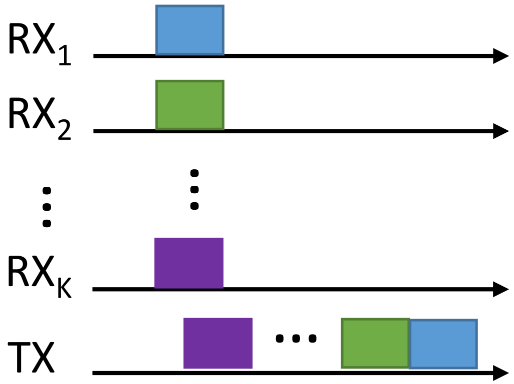

Link Speed Transitions

- <u>Slow-to-Fast:</u>
  Frame length known at end of frame reception, but required to avoid buffer underruns during transmission
  → no CTF possible, unless frame lengths already known, etc.
- <u>Fast-to-Slow:</u>
  Maximum delay savings out-weighted, though less complex to implement than Slow-to-Fast.

Where CTF Matters

Cases that benefit significantly from CTF. Can be extended, though it is suggested to provide technical proof.

General Assumptions

Identical properties for all CTF Streams

- Constant frame length
- Periodic, with period < max. E2E delay

Class based service

- CTF traffic in a dedicated traffic class, highest priority level
- S&F for all other lower priority traffic classes
- If preemption is used:
  - CTF traffic: preempting
  - S&F traffic: preemptible, 127B. Max. Fragment

Network

- Identical links speeds at all hops, typically low (e.g., 100 Mbit/s)
- Illustrated by Ring/Chain topology, indication given if this is not a requirement

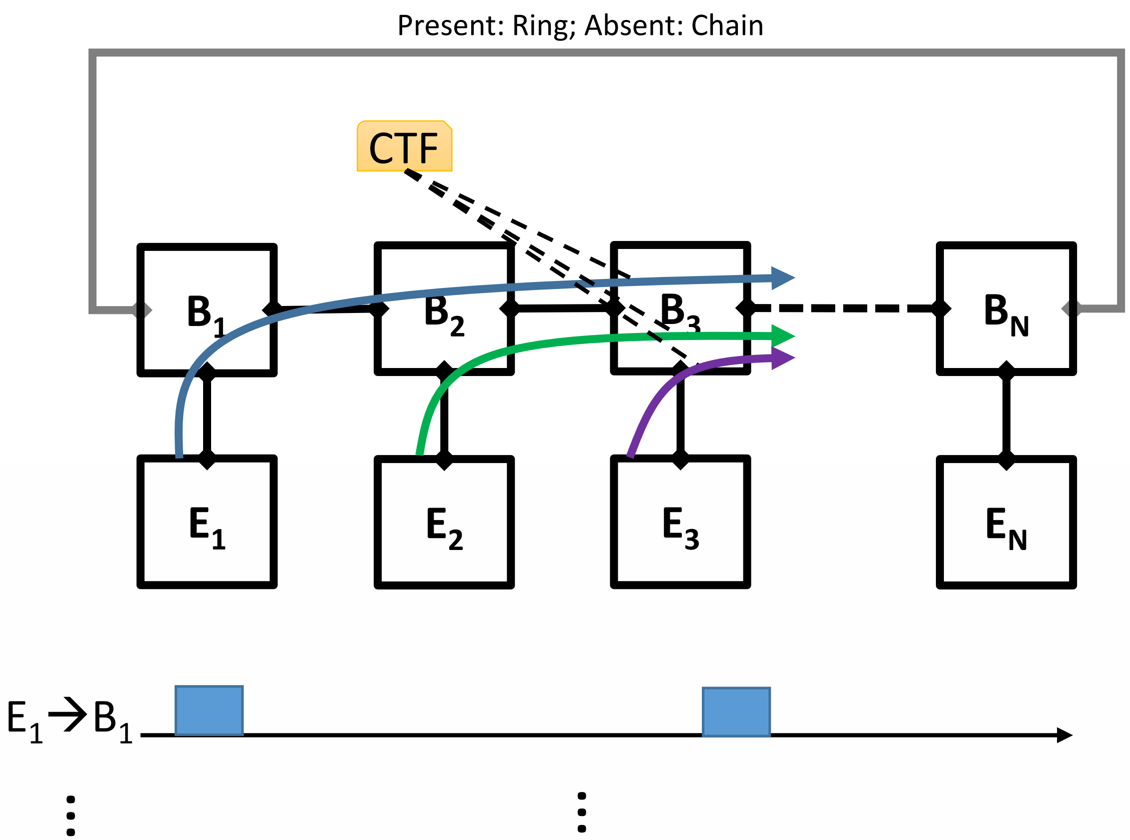

Bridges

- 64B. CTF delay (usually <a bit more than> 14 … <many> Bytes)
- 0B. additional Bridge-internal delay (Relay, etc.)

Asynchronous - Chains & Rings

Description

- Uncoordinated talker transmission times
- Preemption used
- Low fan-in (i.e., rings/chains)
- Medium…large CTF frames (!)

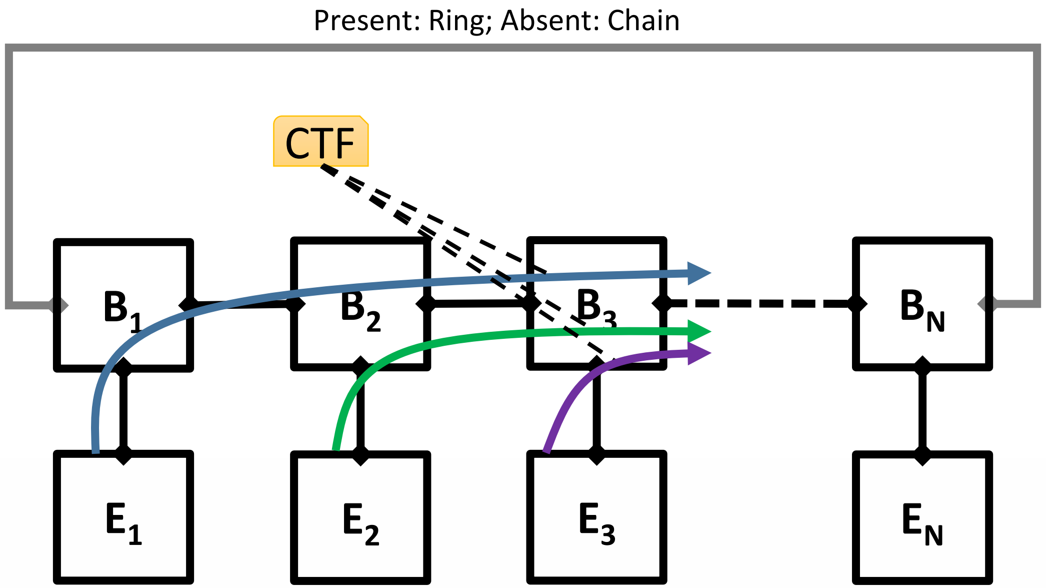

Reduction of Maximum Delay

128B CTF Frames → 17% lower
256B CTF Frames → 30% lower
1542B CTF Frames → 46% lower

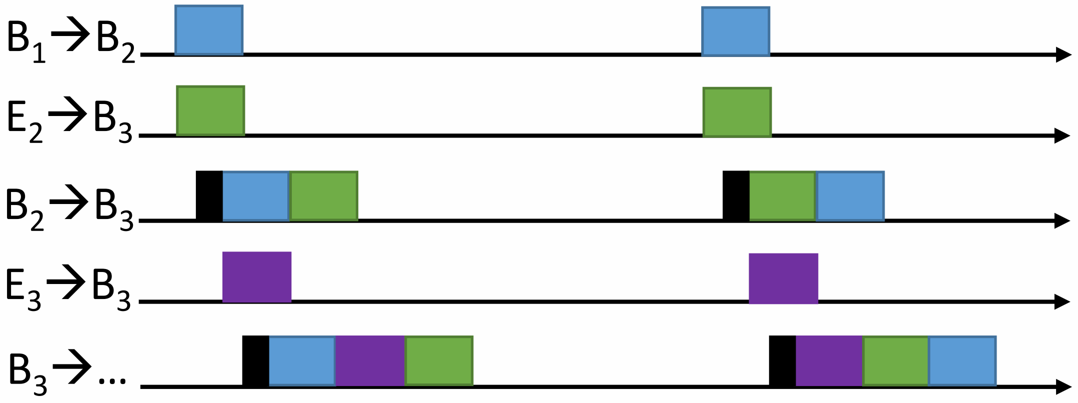

Global CTF Time Slice, Uncoordinated Talkers

Description

- “light-weight” time scheduling
  - CTF-traffic XOR S&F traffic at a time
  - Similarities with Cyclic Queuing and Forwarding

→No interference by S&F frames, no preemption required

- Low fan-in (chains/rings)

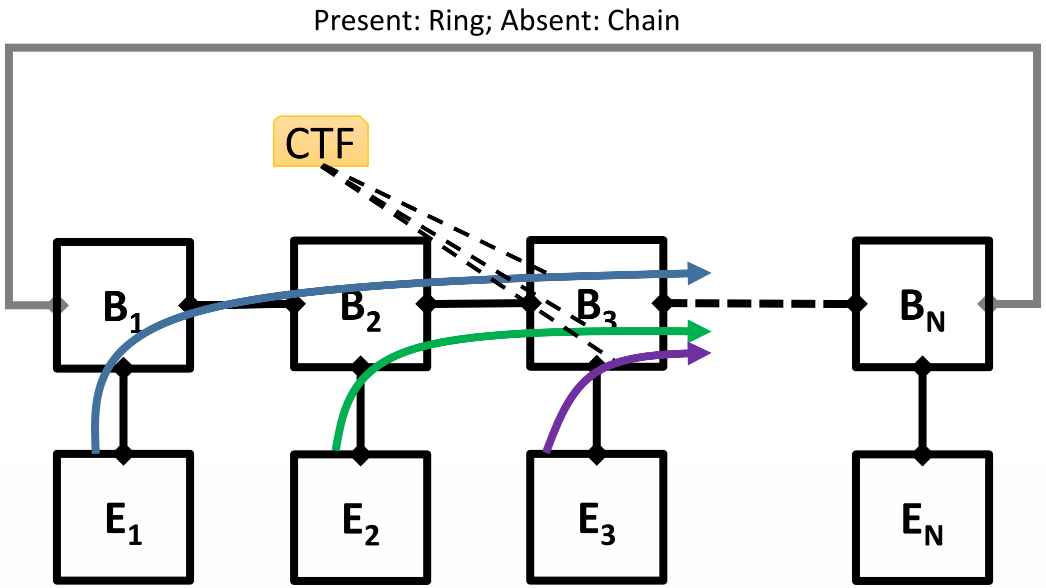

Reduction of Maximum Delay

128B CTF Frames	→ 25% lower
256B CTF Frames	→ 38% lower
1542B CTF Frames	→ 48% lower

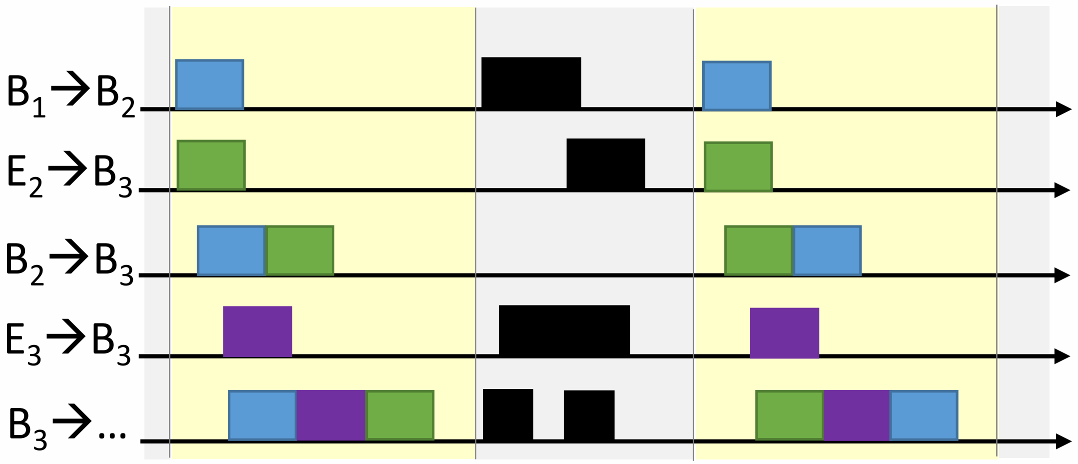

Coordinated Talkers

Description

- Coordinated Talker transmissions
  - Event-based/triggered by frame reception (e.g., responses aligned to initial PLC frames), or
  - Based on sync. time ("scheduled")

→No interference within CTF class - relaxes fan-in[1] & frame length limitations

- Preemption used

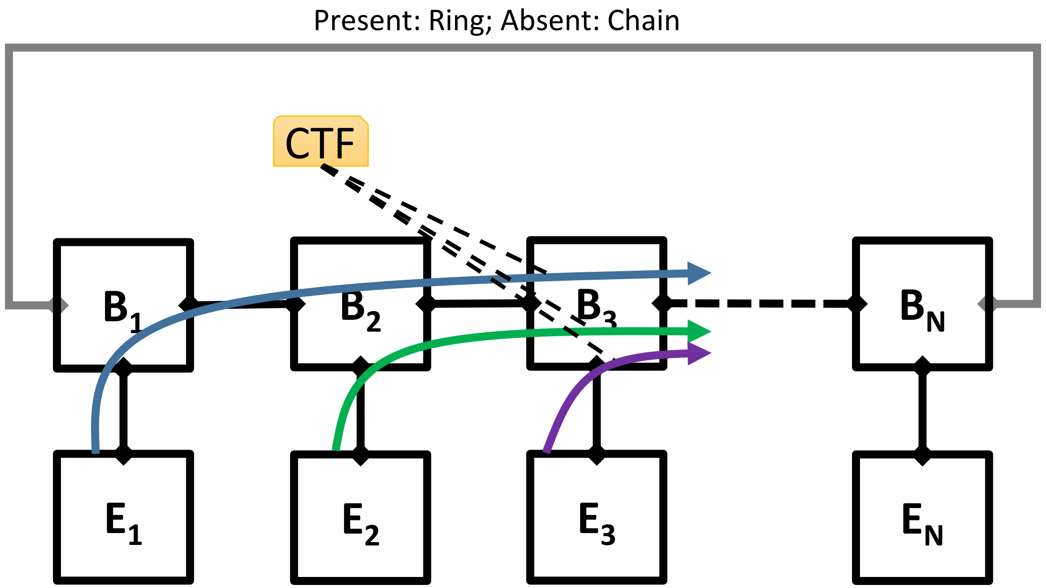

Reduction of Maximum Delay

128B CTF Frames	→ 25% lower
256B CTF Frames	→ 50% lower
1542B CTF Frames	→ 89% lower

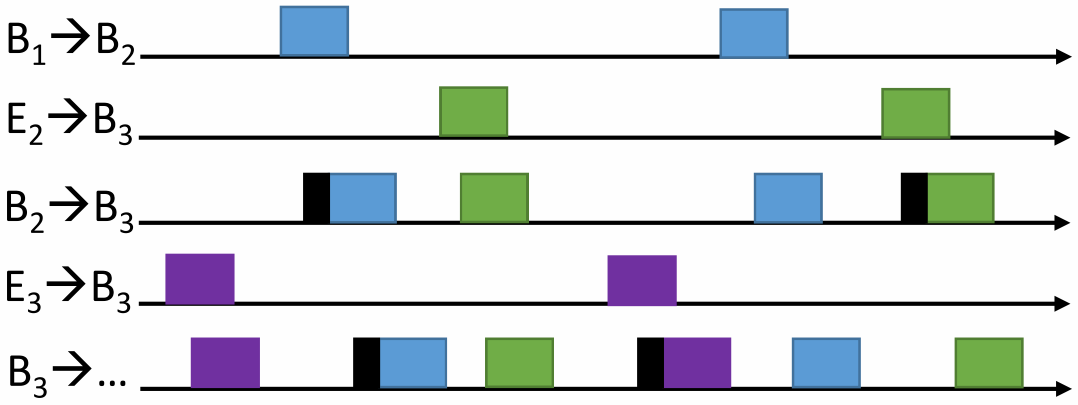

Note 1: Illustrated for Chain-/Ring Topologies, but not limited to these.

"Classic" Time-Division Multiplexing (TDM)

Description

- Global scheduling (every Port)
  - (CTF-frame A) XOR
    (CTF frame B OR S&F frame) at a time

→No interference by S&F frames, or by CTF frames from other reception Ports (i.e., relaxed fan-in limits)

→No preemption required

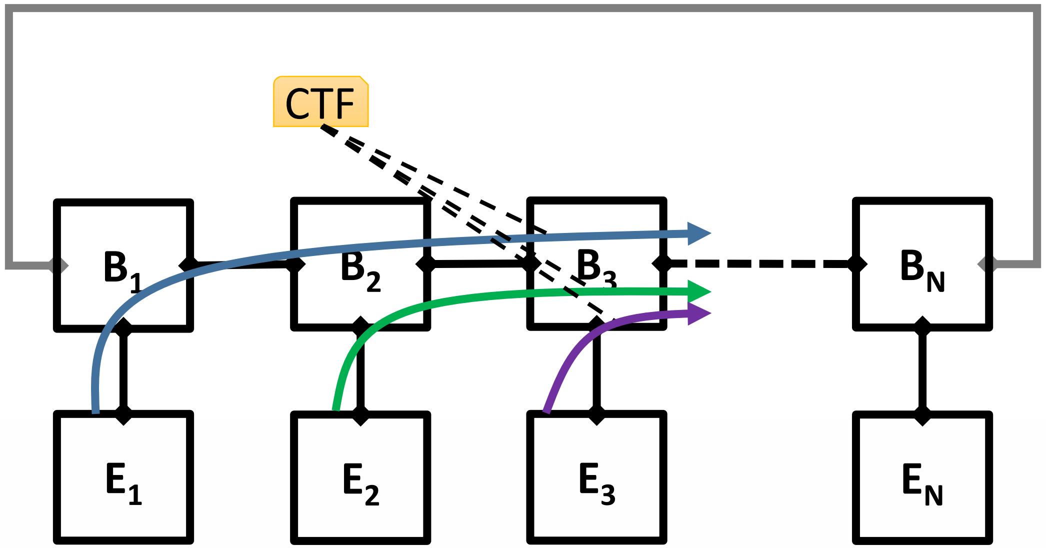

Reduction of Maximum Delay

128B CTF Frames	→ 50% lower
256B CTF Frames	→ 75% lower
1542B CTF Frames	→ 96% lower

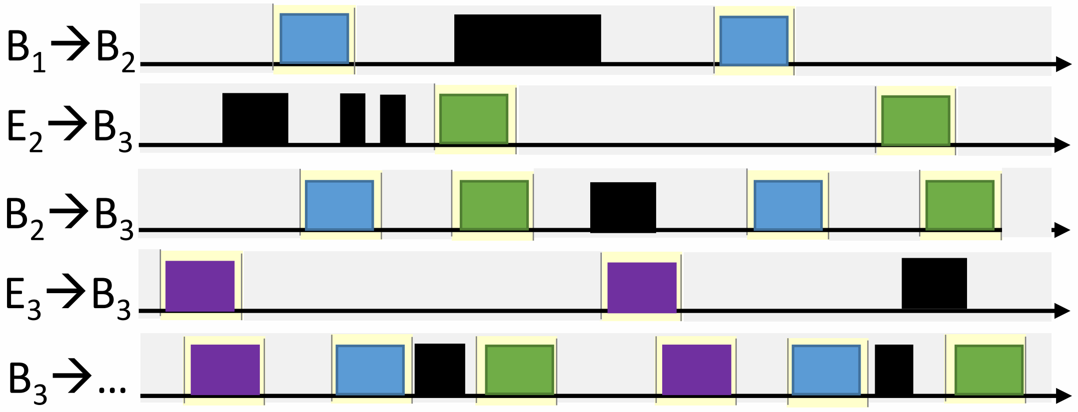

Note 1: Illustrated for Chain-/Ring Topologies, but not limited to these.

CTF-specific Issues and Mitigations

Overview of Issues

#1 Corrupted Frame Headers are discovered too late, i.e. after forwarding decisions have been made

- Frames forwarded to wrong transmission Ports or wrong Traffic Classes can circulate in topological rings
- Frames forwarded to wrong transmission Ports might be considered a security issue (i.e., contents readable in untrusted network segments)
- Frames forwarded to wrong transmission Ports can cause congestion (i.e., unplanned interfering traffic)
- Frames forwarded to wrong Traffic Classes can cause congestion (though transmission Ports might be correct)

#2 Oversized frames are discovered too late

- Unexpected large frames (e.g., jumbo frames) in a CTF Traffic Class can cause congestions for all other frames in this class (unplanned interfering traffic)

Mitigation: Classic TDM (by itself)

Description

- CTF frames “caught” in associated timeslots
  → Corrupted frames can be identified (i.e., out of timeslot), and dropped
- Extra scheduling constraints for coverage apply
- Already standardized in IEEE Std 802.1Q-2018:
  - Input Gates (PSFP)
  - Transmission Gates (EST)

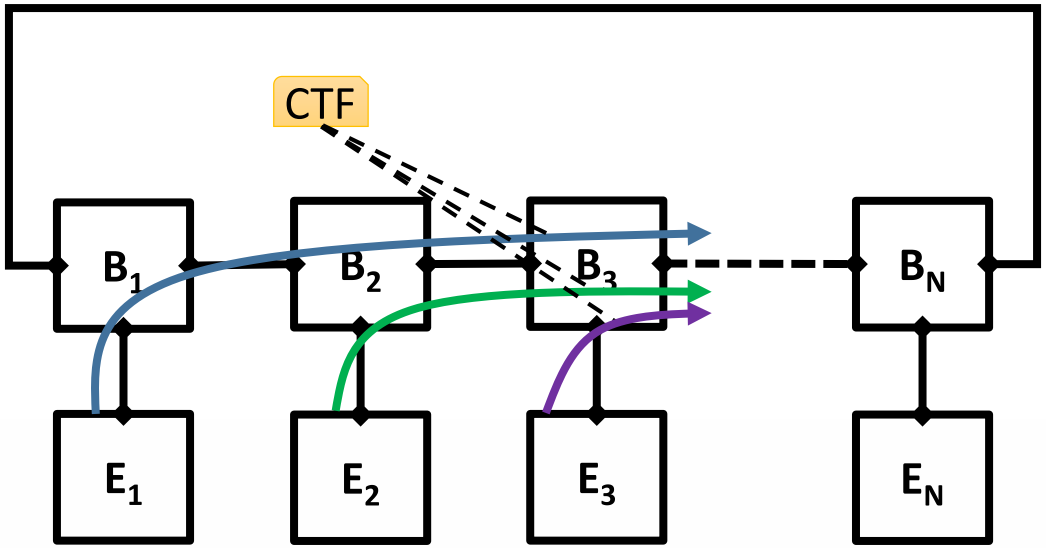

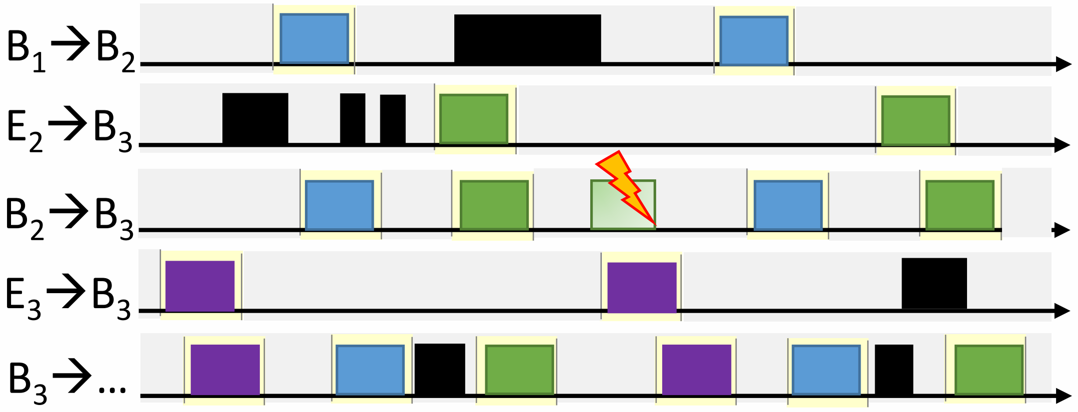

Mitigation: Full S&F Hops

Description

- General solution to handle corrupted headers.
- Moderate use (i.e., excessive use defeats the purpose) - examples:
  - At chain/ring entry points
    → Prevents colliding frame bursts from two reception Ports (congestion)
  - Once per ring
    → Prevents circulation – at most one round
  - From trusted to untrusted network segments
    → Prevents exposure of frame contents to untrusted network segments

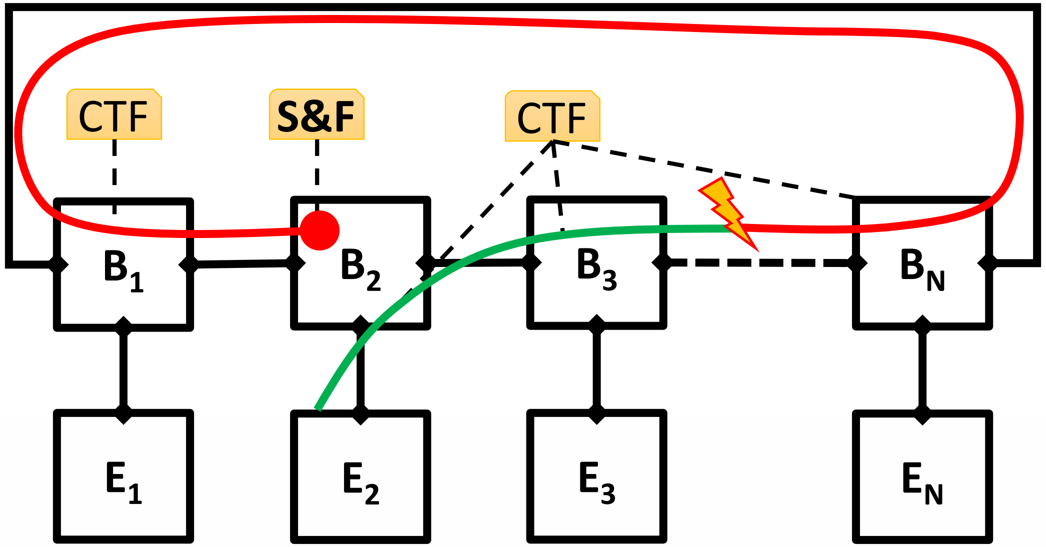

Mitigation: Distributed S&F Hop in Rings

Description

- Ring circulation: Full S&F hop may be split to different locations, as long as forwarding address sets are covered.
- Assumes header corruption leading to address set changes appears at most per round - more than once seems unlikely.

→Avoids S&F during fault-free operation, requires planning.

- Potential realizations for CTF/S&F decision:
  - FDB-based (though can be large)
  - IPV-based (smaller, split CTF class)

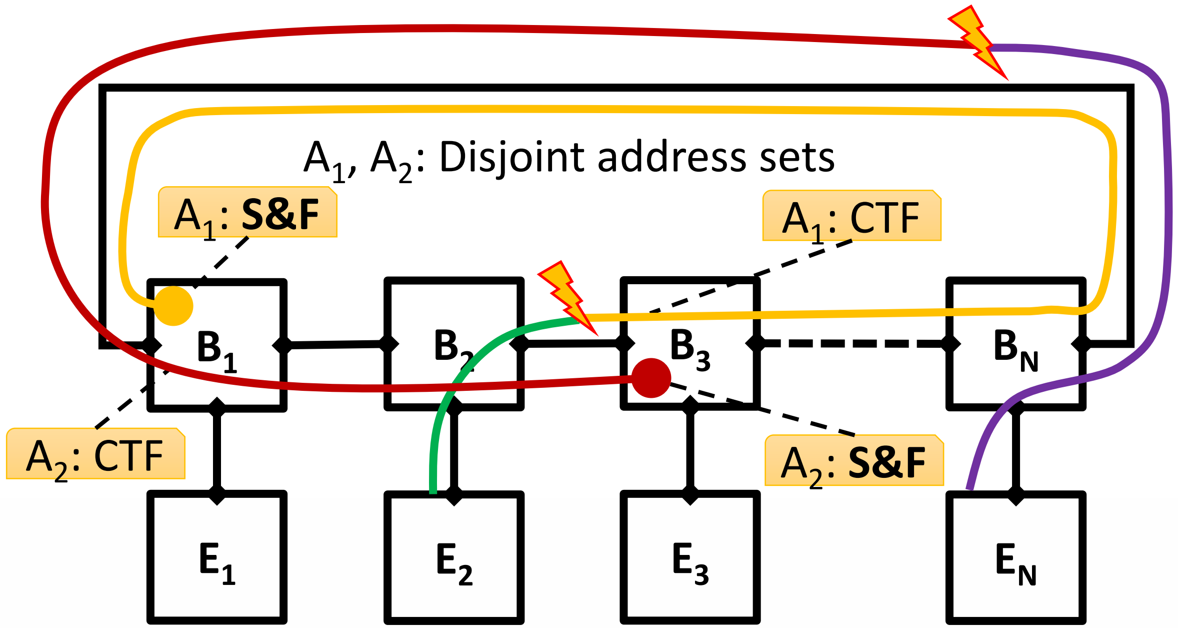

Mitigation: Tail Cutting and Invalidation

Description

- Corrupted CTF frame discovered at end of reception (e.g., FCS error), while already under transmission:
  - Stop ongoing transmission (i.e., “tail cutting”)
  - Invalidate (e.g., attach incorrect FCS)
- Prevents circulation, when rings are large, and affected frames are short enough
  → In conjunction with frame size limiters (IEEE Std 802.1Q-2018)

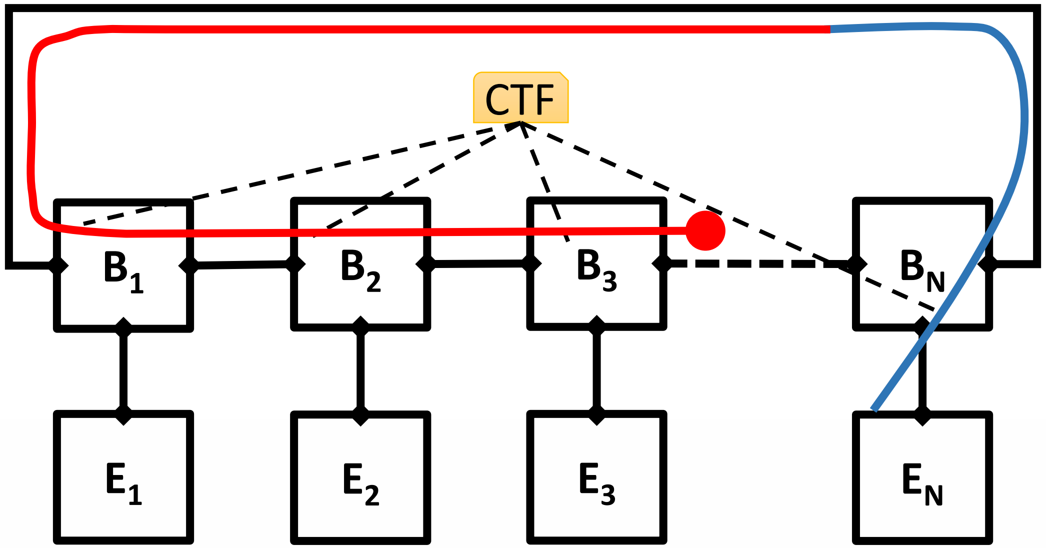

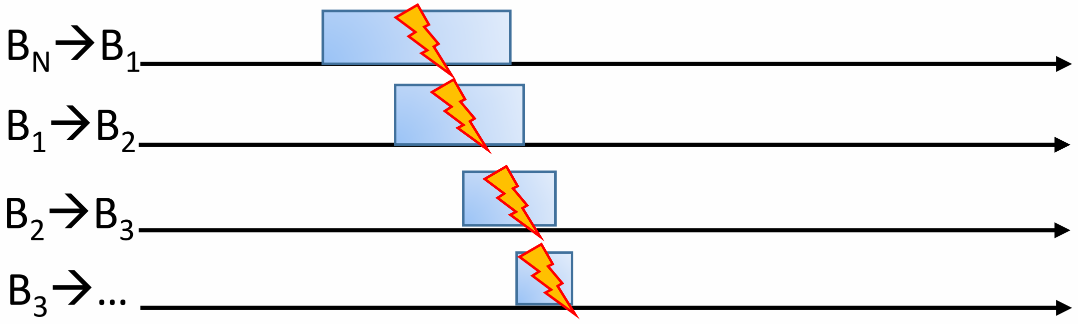

Further Thoughts on Mitigations

Definition of Goals Required

- Example:
  A frame shall not pass the same transmission Port more than once due to frame header corruption at one hop.

- Enhances technical discussion
- Identify appropriate/effective mitigation methods/combinations thereof for particular scenarios

Network Aspects

- Mitigating CTF issues is network dependent – specific Bridges mechanisms may or may not be sufficient in a given network.
- Example: Tail cutting in 3 hop rings appears insufficient, but may be ok in a 30 hop ring.

What to Standardize, and Where

IEEE 802.1 Parts (1)

Dedicated standard – not an Amendment to IEEE 802.1Q

- Not limited to IEEE 802.1Q
  - Other IEEE 802.1 Standards (e.g., IEEE 802.1CB, IEEE 802.1AS)
  - End station aspects
  - Network aspects
- Limit to “CTF capable” functions and protocols (e.g., from the Forwarding Process):
  - Only contained IEEE 802.1 Stds functions and protocols work in presence of CTF
  - No statements on (im)proper operation of functions and protocols beyond the listed ones (i.e., these are “out”, no need to address compatibility with all IEEE 802.1 protocols)
  - Start with the important protocols and functions, then extend
- Readability
  - Avoids cluttering, keeps functional additions and changes organized
  - IEEE 802.1Q alone has more than 2000 pages …

IEEE 802.1 Parts (2)

CTF standard contents (not a Standards document clause structure)

- **Existing Bridge functions/protocols** (e.g., pieces from the 802.1Q-2018 forwarding process)**, and modifications** (if needed)
- **New Bridge functions** (e.g., tail cutting/CRC invalidation)
- **Management Interfaces**, Counters, etc.
- **Where CTF Matters -** and where not?
- **Header corruption mitigation**
- **Device conformance**
- **Network conformance** - IEEE conformant networks using CTF should apply mitigations to handle CTF-specific issues!
- **MAC Service Interface (???)**, frame reception

Service Interface for Frame Reception

Involve IEEE 802.3

- Reach agreement in IEEE WG 802.1 on CTF in general
- Formulate requirements, for example:
  - Early reception event: Header fields/tags for Bridge forwarding decisions
  - Late reception event: FCS Status
  - Continuous: Transient Frame Length and payload
- Request a solution from IEEE WG 802.3

Summary

1. Where CTF matters and where not on a technical level
2. CTF and appropriate mitigations, definition of goals may help
3. IEEE 802.1 Standard, not an amendment
4. Work towards a common view in 802.1 for subsequent discussion with other groups (e.g., 802.3)

Thank you for your Attention!

Questions, Opinions, Ideas?

Johannes Specht
Dipl.-Inform. (FH)

johannes.specht.standards@gmail.com
T +49 (0)170 718-4422