

NEW & IMPROVED!

# An event-based AVB synchronization architecture

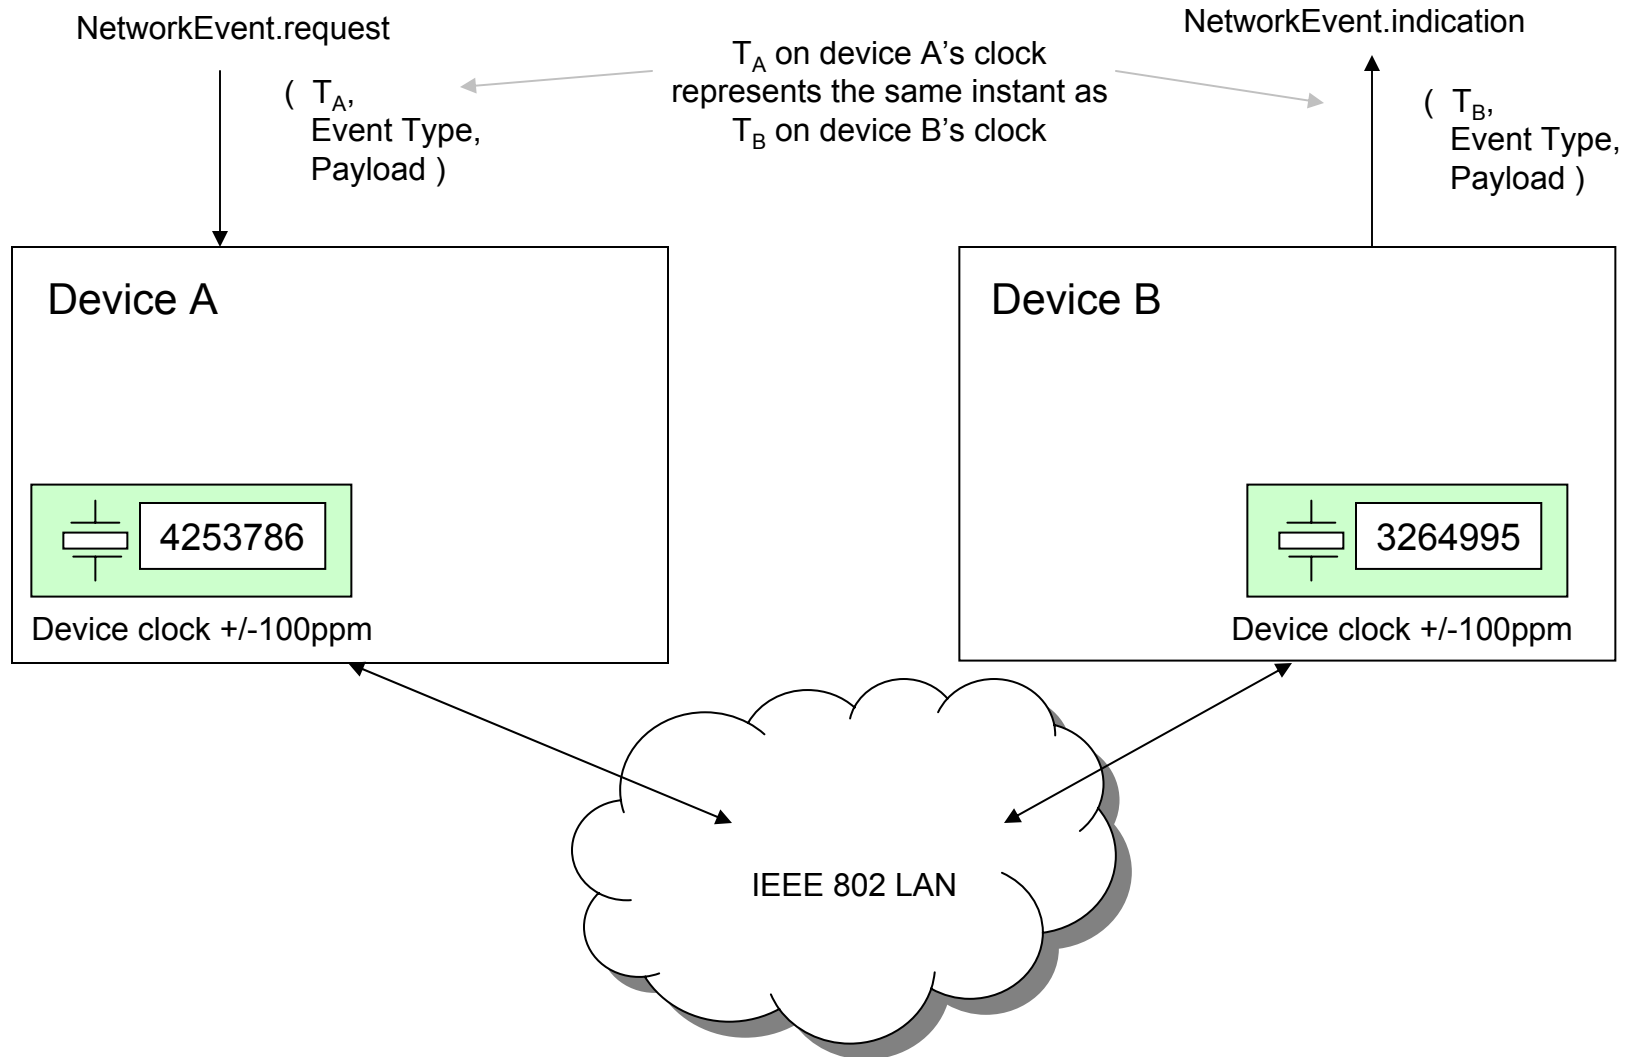
802.1AS Timing and Synchronization for Time-  
Sensitive Applications in Bridged  
Local Area Networks

Chuck Harrison  
10 February 2007

# Overview

- The proposed architecture separates a low-level, epoch-independent, **event-based synchronization service** as a sublayer below TAI time distribution
- The event-based synchronization service does *not* involve a Grandmaster clock but *does* rely on
  1. devices' ability to measure short durations of time
  2. the election of a “sync master” to issue periodic (but not particularly accurate) broadcast SYNC packets
- The proposal uses media-independent 802.1 interfaces to accommodate 802.3 and 802.11 PHYs without requiring BCs
- This proposal includes a P2P frequency and delay correction mechanism at the event sublayer, but the AVB/802.1AS performance goals can usually be met without using them
- A 1588-equivalent Ordinary Clock service is optionally co-located at any bridge and is synchronized by messages delivered by the event-based service
- This architecture permits continued operation (albeit with degraded performance) in a mixed network of both AVB and non-AVB bridges

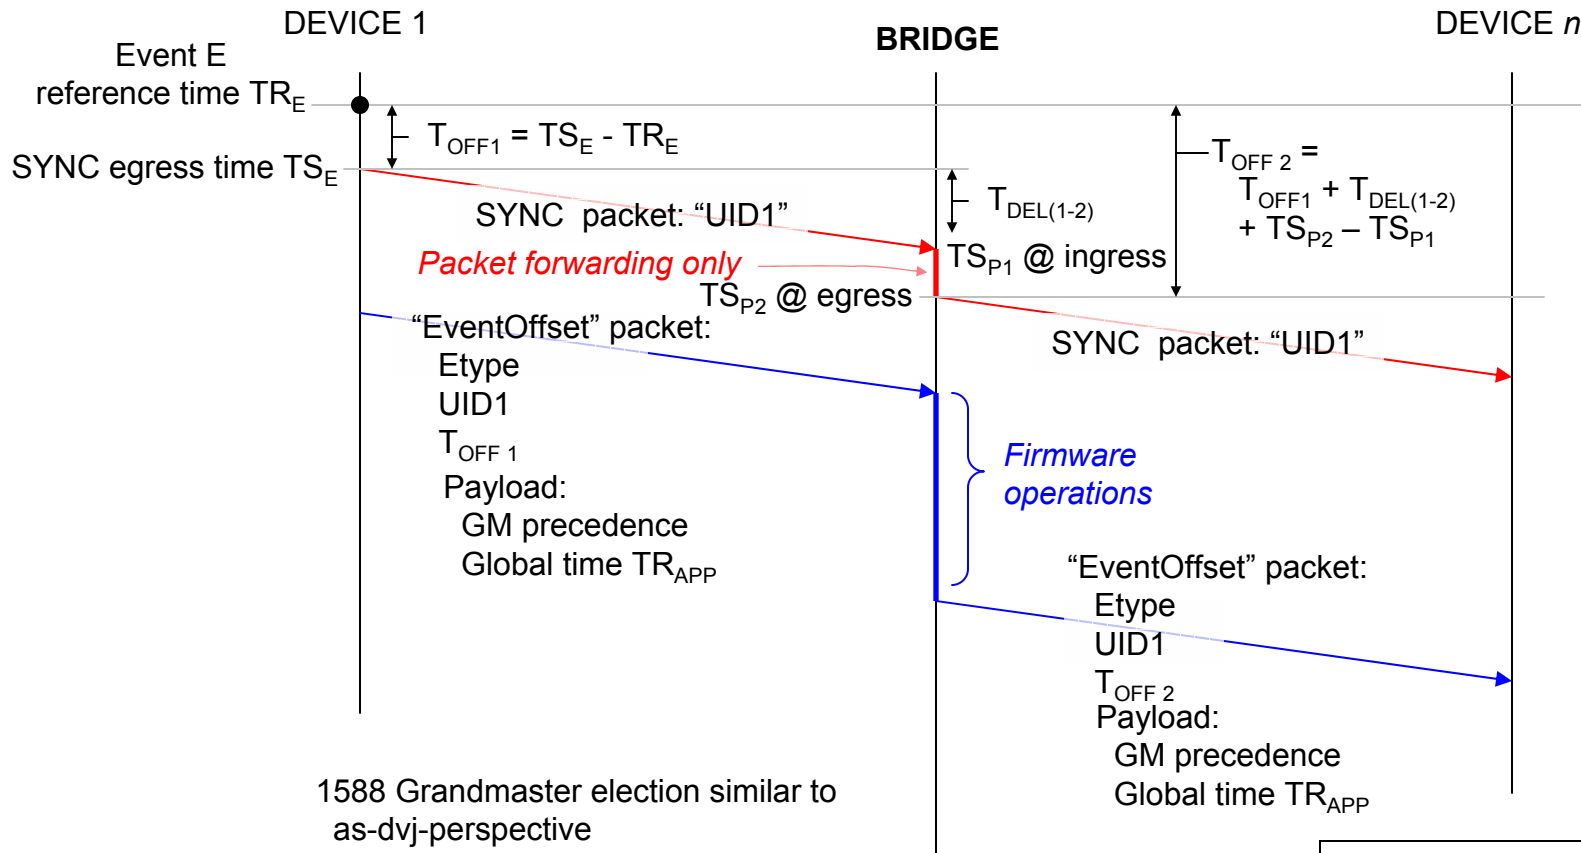
# Event Service



# Mechanism for Event-based synchronization service

1. A master SYNC generator floods network with periodic SYNC packets (exact timing of these packets is not important)
2. When we submit a NetEvent service request at device A, A waits for the next SYNC packet, then computes the *time offset*  $T_{\text{OFF}}$  between  $T_A$  and the SYNC packet's egress time. It transmits  $T_{\text{OFF}}$  'on the wire' as an EventOffset message
3. When the EventOffset frame arrives at AV bridge Z, the ingress port computes the corresponding time  $T_Z$  on its local timescale, and transforms the EventOffset frame into an EventTime frame that gets forwarded through the bridge's MAC Relay entity
4. Bridge Z's MAC Relay entity forwards the EventTime frame to a MAC port for transmission, where it is converted back to a new EventOffset frame and emitted from the PHY port
5. The same EventTime frame also causes a NetEvent service indication at the AV bridge's event access port containing  $T_Z$  and the original event's payload parameters.

# Master Clock distribution over event-based synchronization service



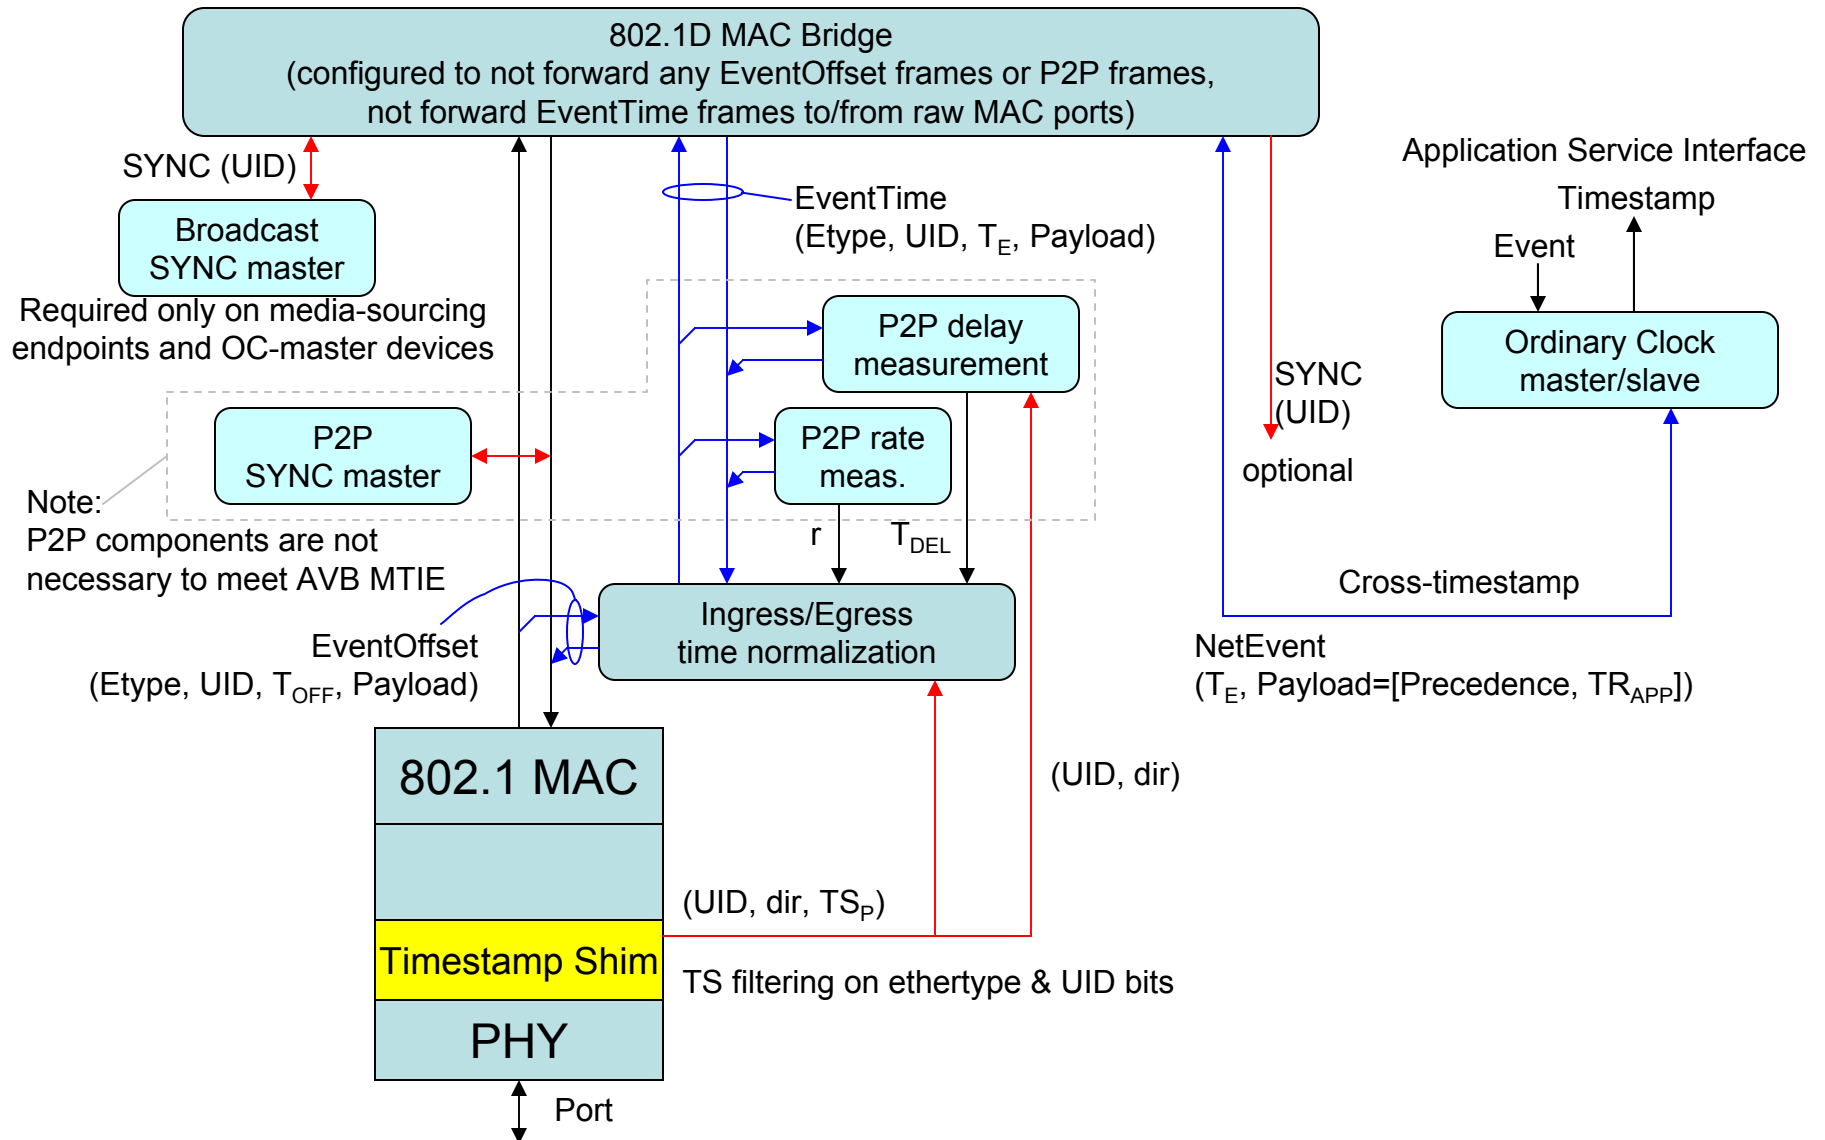
1588 Grandmaster election similar to as-dvj-perspective

SYNC packet is *not* generated from Grandmaster timescale

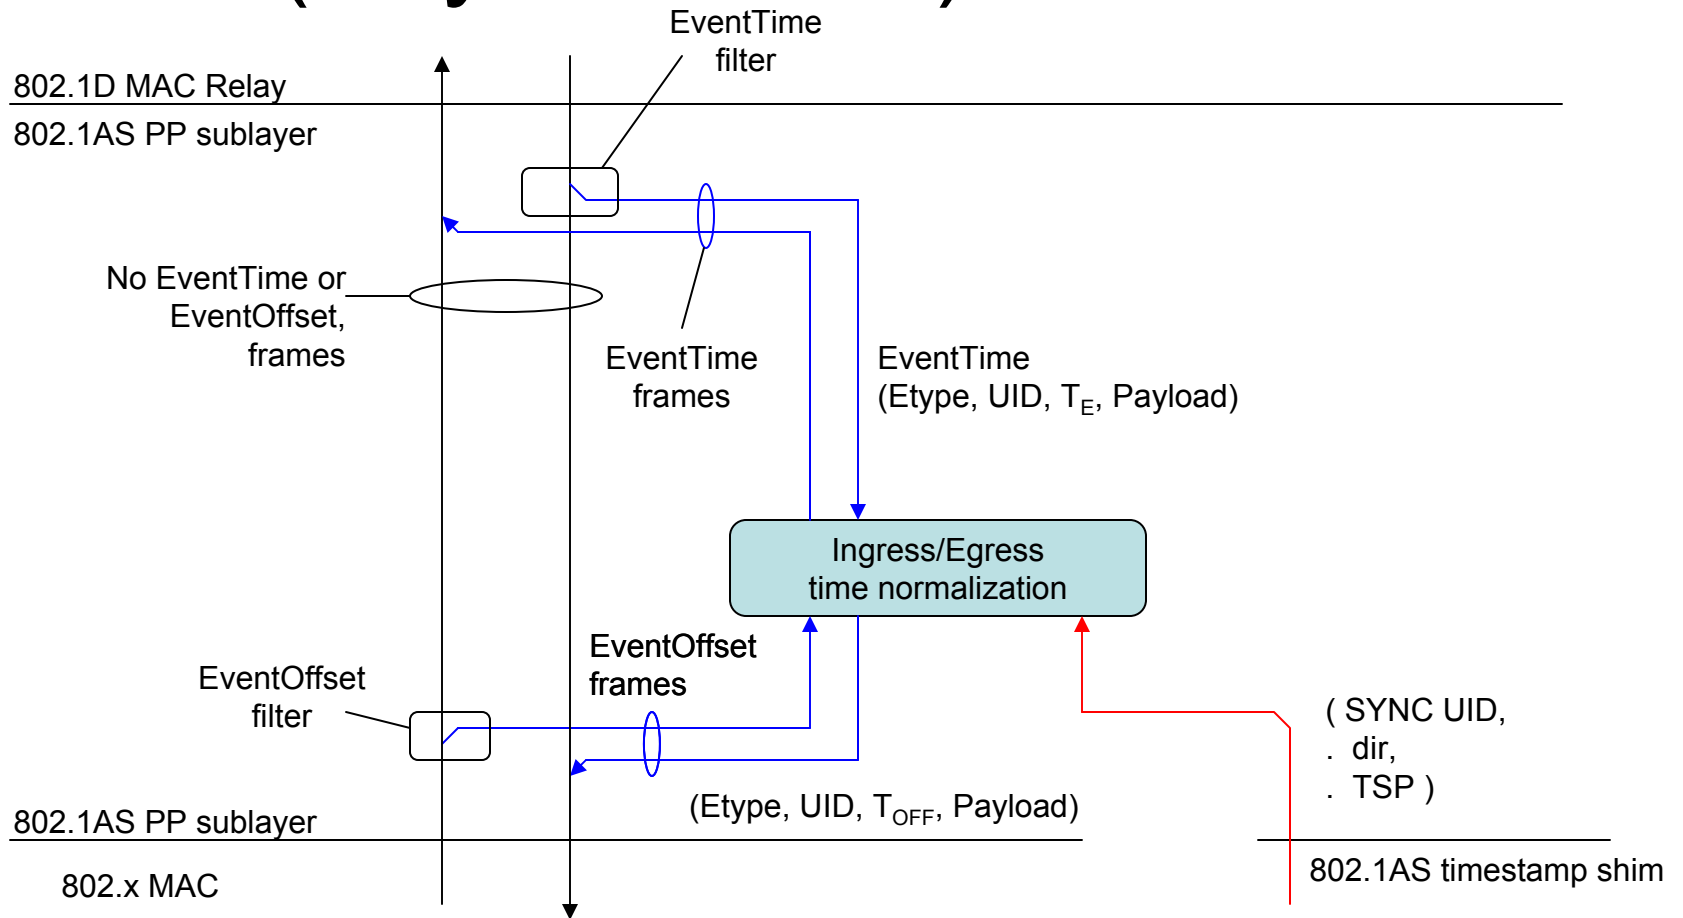
Cross-timestamp delivered to OC:  
 Local time  $T_E$   
 Global time  $TR_{APP}$



# 802.1AS Device stack

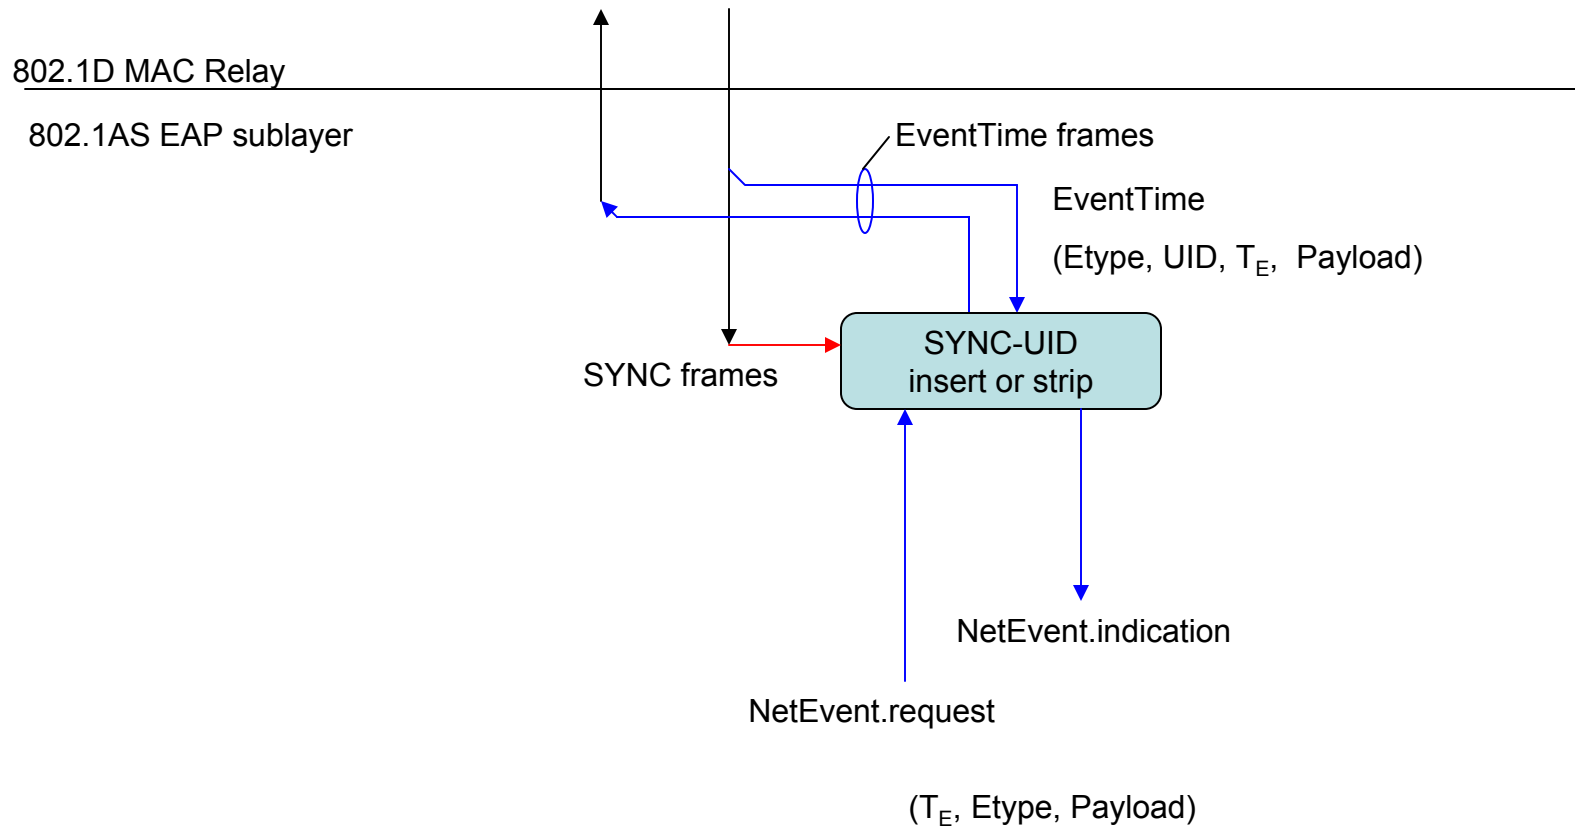


# 802.1AS ISS PP Sublayer (Physical Port) - basic

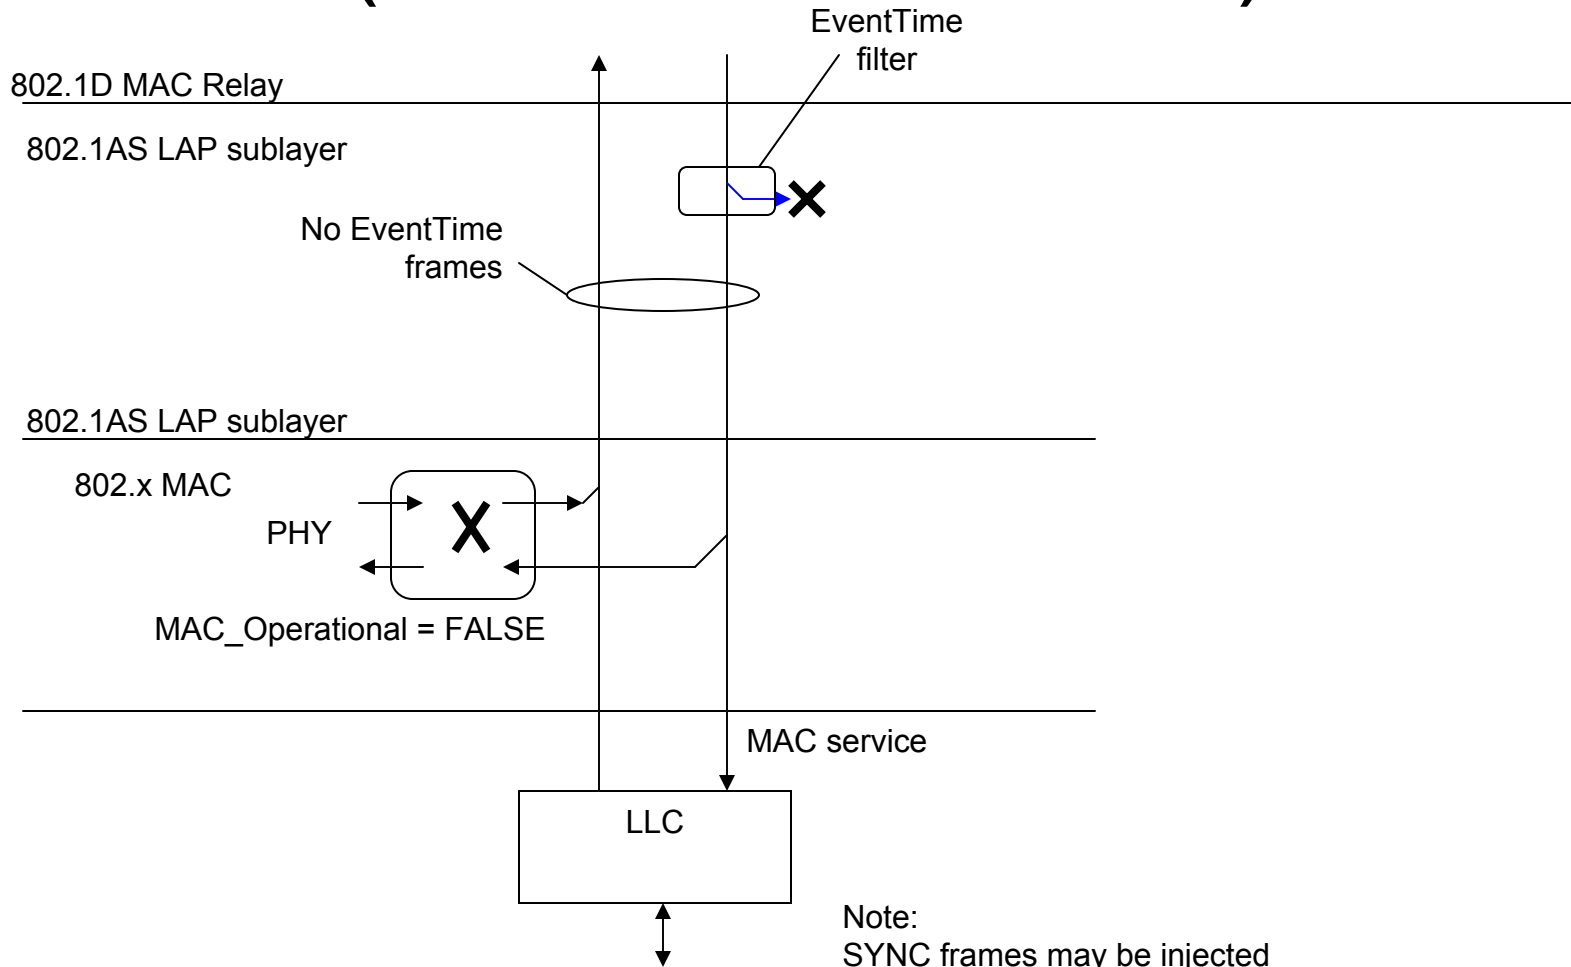




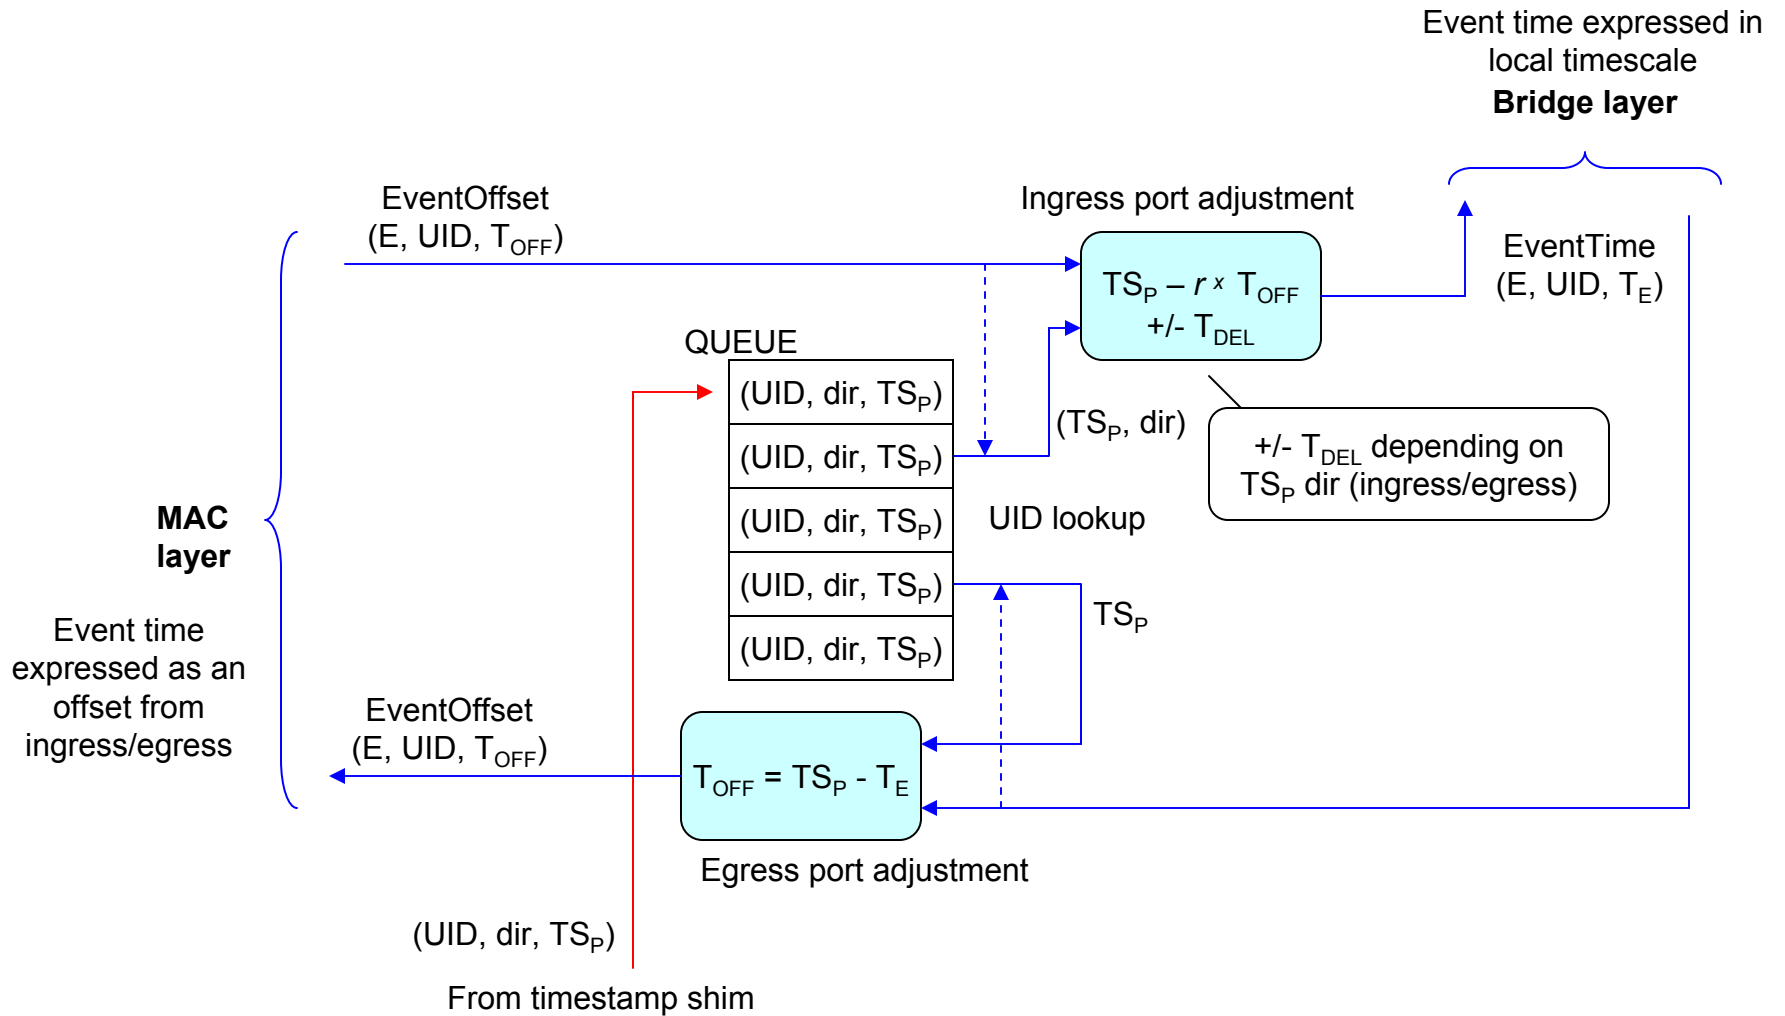
# 802.1AS ISS EAP Sublayer (Event Access Port)



# 802.1AS ISS LAP Sublayer (Local Access Port)



# Ingress/Egress time normalization



# Ordinary Clocks

- Ordinary Clocks are maintained in a 1588-inspired manner.
  1. The Grandmaster issues a broadcast NetEvent message, with each AVB bridge acting much like 1588 Transparent Clock.
  2. The NetEvent message contains a local time  $T_E$ , and its payload contains a Grandmaster time (and precedence).
  3. This results in a “cross-timestamp” delivered to each slave OC correlating the device’s local timescale to the Grandmaster’s timescale. (The cross-TS point may be in the recent past.)
  4. Based on this cross-timestamp the OC is aligned to the Grandmaster (mechanism is implementation-specific)
- An OC “talks” *only* when it is Grandmaster
- Grandmaster election: an OC goes to slave mode when it gets a NetEvent from an OC with higher precedence; goes to master mode if it does not get any for a while

# Sync Master behavior

- Basic algorithm:
  - If you hear another SYNC master, be quiet
  - If you hear nothing, wait a random interval and start transmitting SYNC
- Broadcast SYNC master (~10ms)
  - sends broadcast SYNC packets to all ports of bridge
- P2P SYNC master (~100ms)
  - sends link-local SYNC messages to attached port

# P2P rate & delay protocol

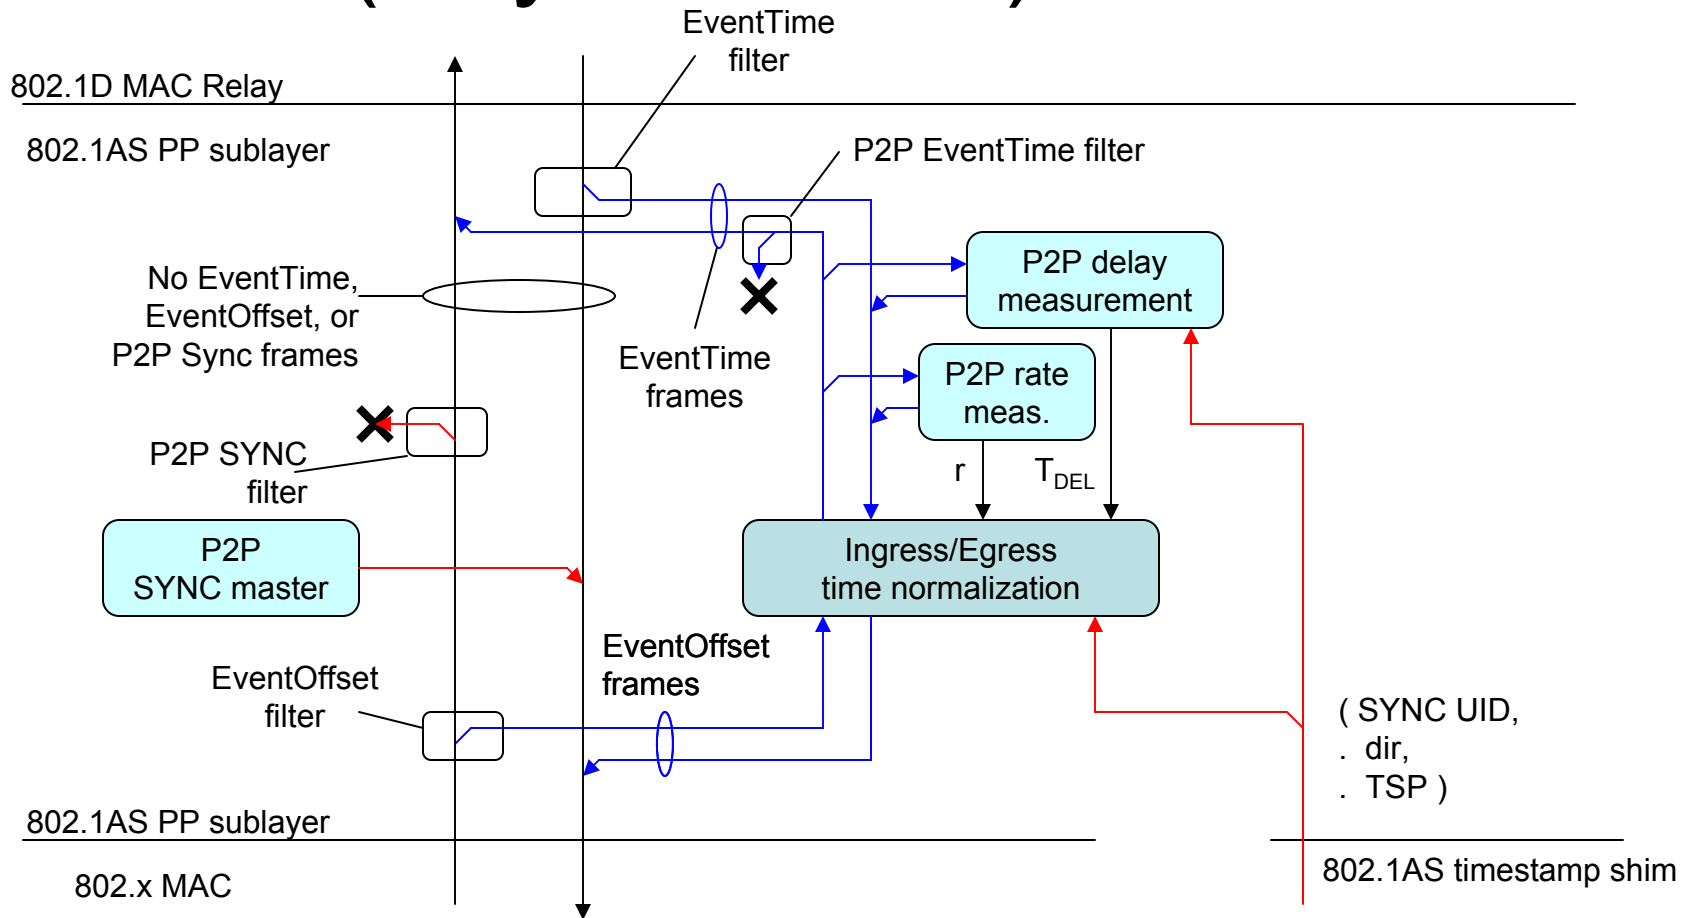
## 1. Rate correction (roughly equivalent to syntonization)

- Each partner issues a link-local SYNC packet, along with a link-local EventOffset message ( $E_{P2P}$ , UID,  $T_{OFF}$ ) at a periodic rate ( $\sim 100\text{msec}$ ). The  $E_{P2P}$  event is defined as the most recent 1-msec-granularity “tick” of the sender’s timescale. Receiver examines interval between adjacent  $E_{P2P}$  events, determines number of 1-msec ‘ticks’ (might not be exactly 100), and computes frequency ratio  $r$  between link partners’ timescales.
- The computed frequency ratio is filtered and passed to the ingress time normalization function as a correction factor for the incoming  $T_{OFF}$  field in EventOffset messages arriving on this port.

## 2. Delay correction

- Each partner sends a link-local EventOffset ( $E_{P2P}$ , UID,  $T_{OFF}$ ) message, referencing the UID of each incoming link-local SYNC packet. The  $E_{P2P}$  event is defined as above. Each partner now has two ways of computing the phase offset between 1-ms ‘ticks’ of his own clock and his partner’s: one based on the outgoing SYNC packet, as part of the rate correction scheme, and the other based on the incoming SYNC packet just described. While the phase offset itself is arbitrary, the difference between these two measurements of it is the residual uncorrected round-trip delay on the link.
- The residual round-trip delay is used to correct a filtered estimate of the actual delay. The estimated delay is divided by two and passed to the ingress time normalization function as a  $T_{DEL}$  correction to EventOffset messages arriving on this port.

# 802.1AS ISS PP Sublayer (Physical Port) - full



# A lot like 1588 with TCs, but...

- SYNC packets contain *only* an opaque UID and need *no* upper-layer processing
  - Capture timestamp *and* UID at MAC lower layer
  - Forward packet as standard 802.1D bridge
- SYNC packet is *not* queued waiting for the followup message; therefore standard 802.1D bridge logic is used and residence times are so short that synchronization is usually unnecessary in bridges
- SYNC master (loose ~10msec period) operates within the event-based sublayer and can be elected with a simple mechanism (no precedence issues)
- Grandmaster must wait until a SYNC issued by the sublayer comes by, then timestamp *that*.
- SYNC packets may pass either the same or opposite direction to other synchronization-related messages on a link
- Identical messages used for time propagation and for link delay measurement



# What about 802.11?

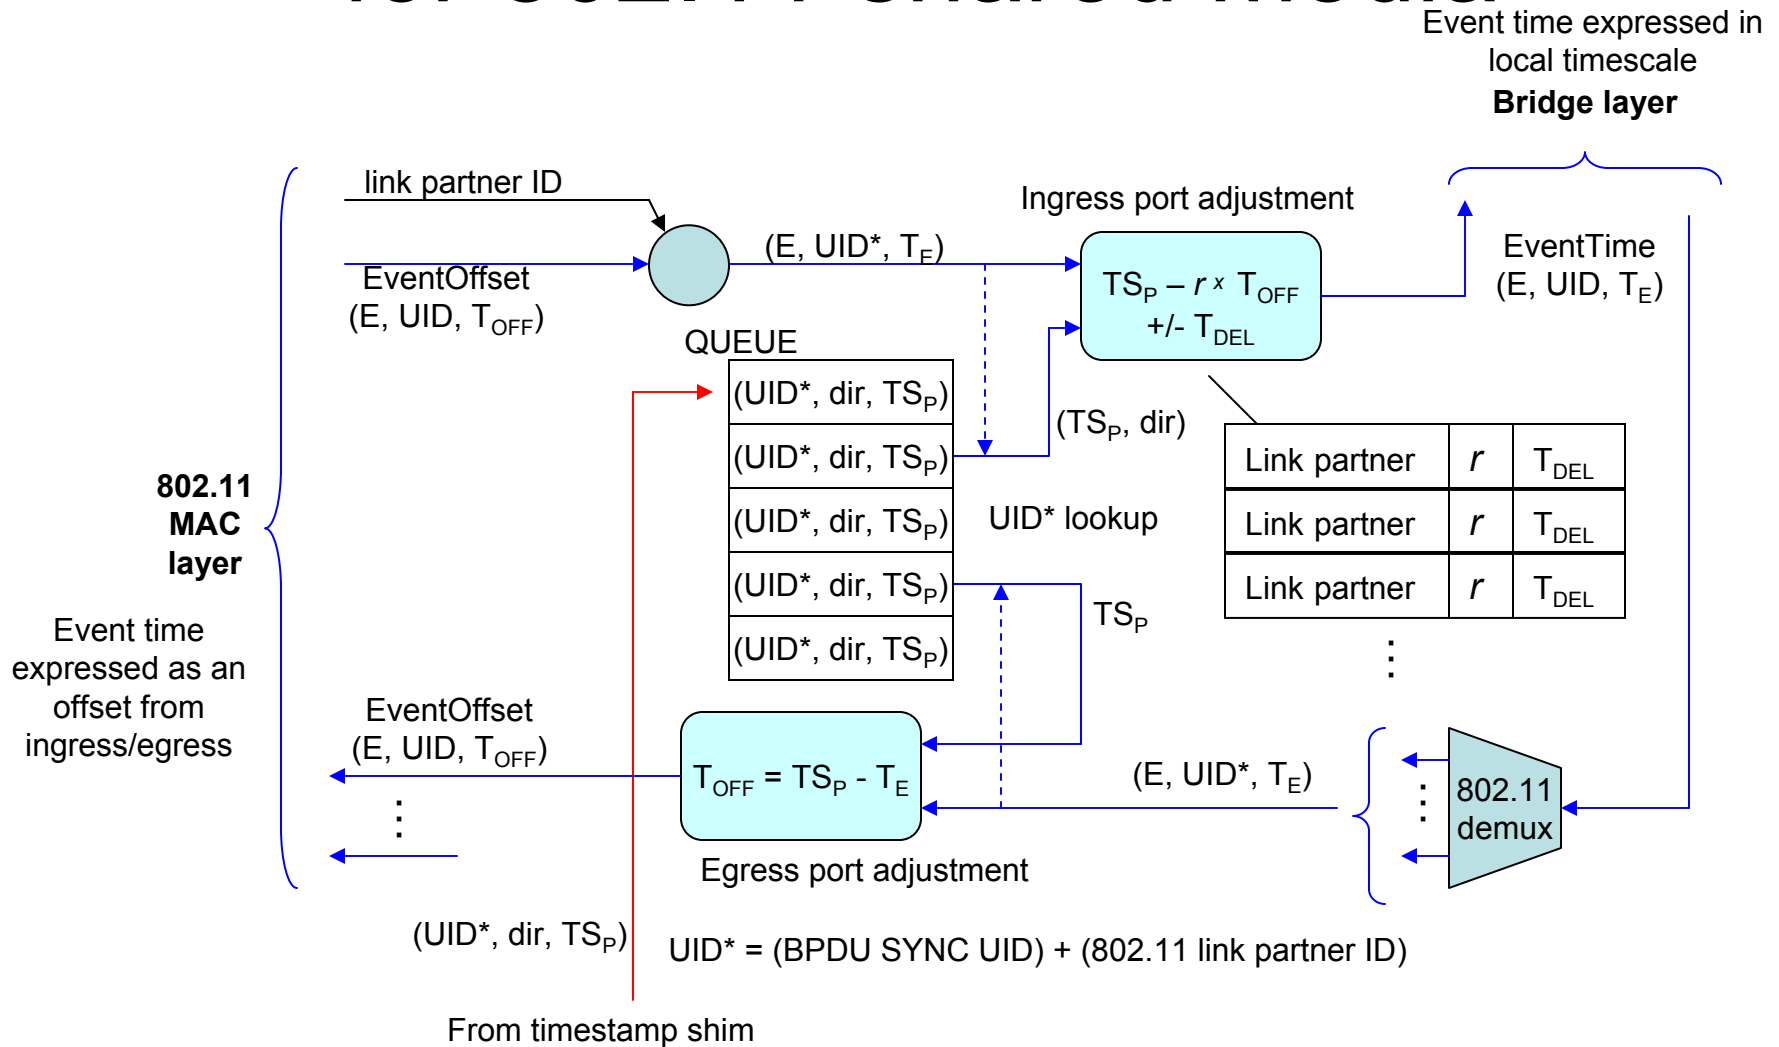
- Where SYNC propagates across a link, this protocol relies on tight correlation between egress timestamp at sender and ingress timestamp at receiver
- If 802.11 broadcast service is used, many STAs will miss SYNC packets, and will lose Events referencing those packets. This is not a disaster: the AppTime protocol for OCs is robust to moderate packet loss.
- Alternatively, if 802.11 P2P service is used (with ACK / retransmission), different STAs will experience different “residence times” at the same bridge port for the same SYNC UID.
- This can be accommodated by translating each broadcast SYNC BridgePDU and each EventTime BPDU into multiple unicast MacPDUs, modifying the lower layer so that SYNC UIDs are effectively extended by link-partner-ID bits (which are ultimately stripped at the receiving STA before forming BPDUs). SYNC UID timestamps propagate to ingress/egress normalization queue only after packet ACK.

# Additional material

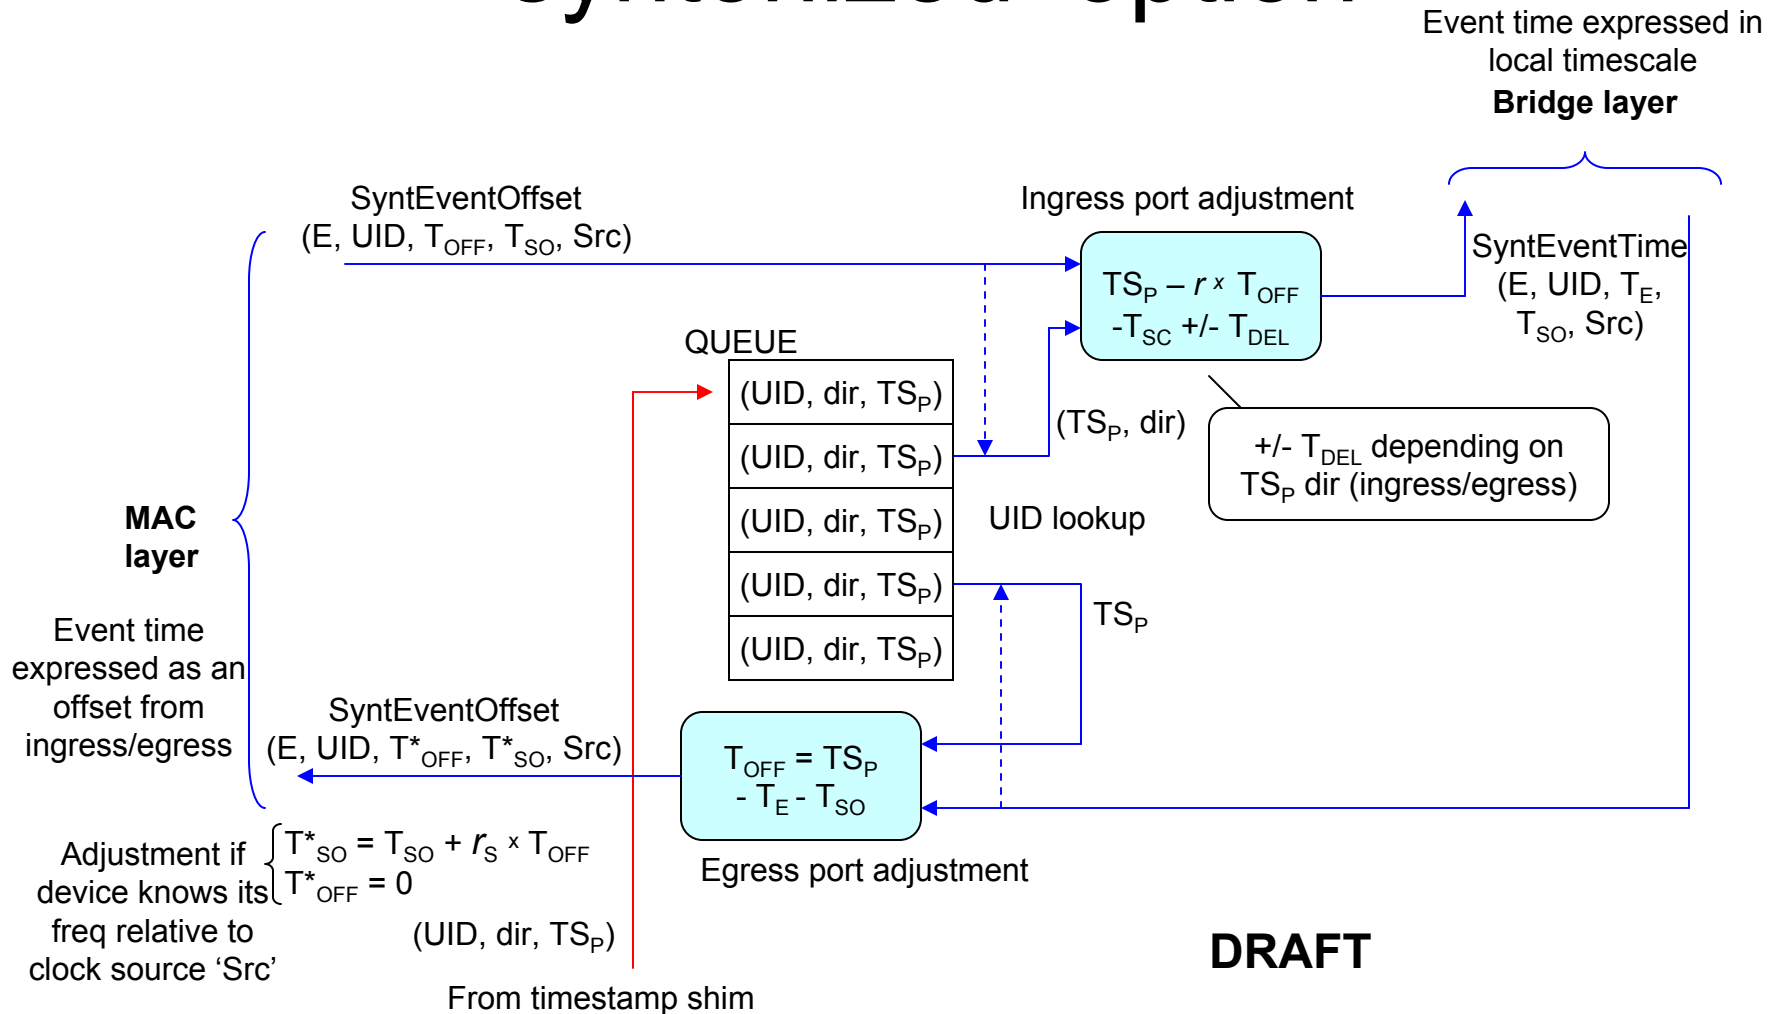
## Incompletely developed ideas

- Potential 802.11 implementation, see slide
- Syntonized offset: assume some part of  $T_{\text{OFF}}$  is known “precisely” (ref to a Grandmaster, say) and the rest is accumulated through non-syntonized bridges. By identifying the “syntonized” component of EventOffset, residence time error accumulation can be minimized.
- Independent MAC timebases: MACs for different media (or on different blades) may operate from different crystals. If the “timestamp shim” sublayer accepts a hardware event input, these independent MAC timescales may be correlated and made invisible to higher sublayers.
- Link Aggregation: this protocol should work with 802.3 clause 43 link aggregation. The timestamp shim report should include a link identifier so that distinct  $T_{\text{DEL}}$  values can be maintained.
- Broadcast, multicast, unicast, link-local, see slide
- MACsec: encryption/authentication on everything except SYNC packets is transparent. In principle this works on SYNC packets, too, but it requires two portholes into the stack (below and above MACsec), so it may in practice be more complex. With common MACsec authentication suites, underlying packet is unmodified so we could identify SYNC ethertype and UID inside the MACsec wrapper at a low-layer porthole.

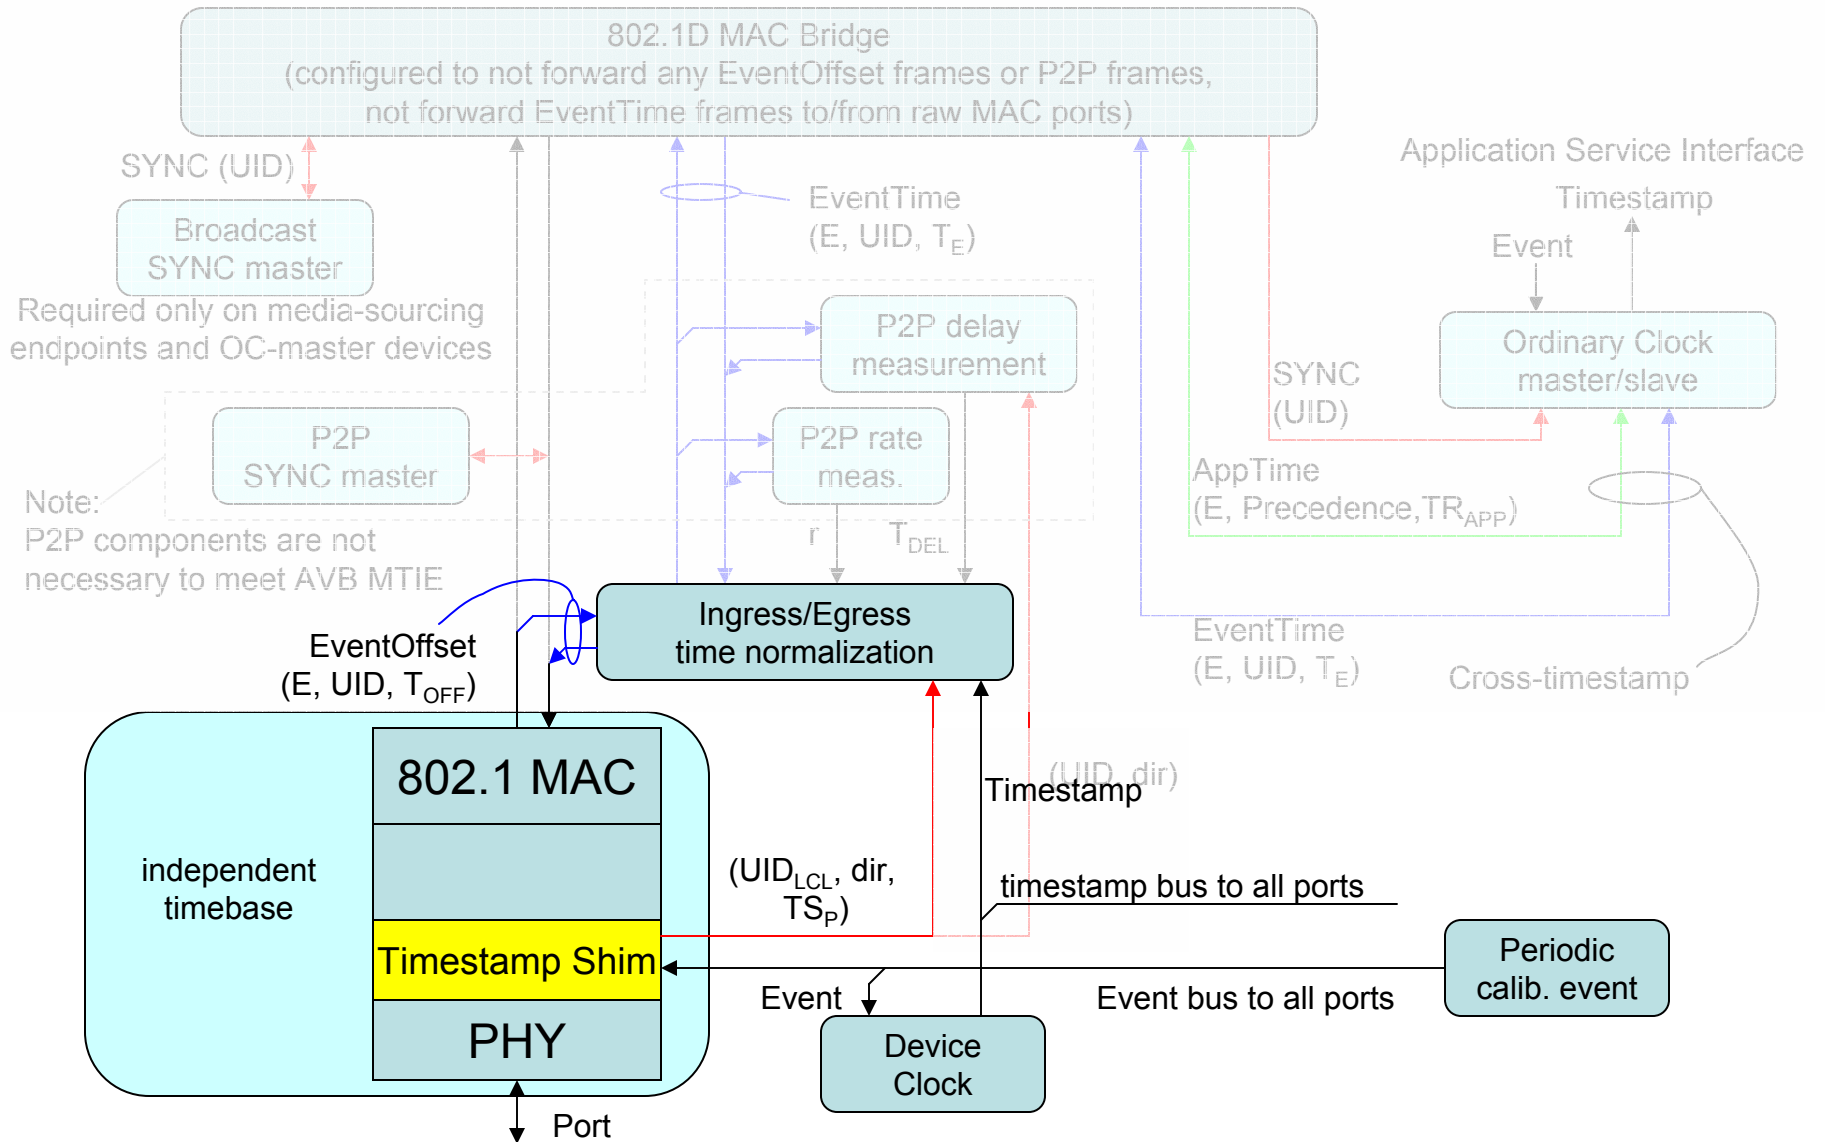
# Ingress/Egress time normalization for 802.11 shared media



# Ingress/Egress time normalization 'syntonized' option



# Independent MAC timebases



# Broadcast, multicast, unicast, etc

- This architecture is based primarily on a 802.1D MAC bridge model, so many 802.1 addressing options should work (in addition to broadcast) for Event and AppTime messages. Group multicast and VLAN should work.
- The basic protocol (without P2P rate & delay measurement) will propagate through non-AVB bridges with reduced accuracy. If P2P packets are sent on a link-local MAC address, as proposed, the protocol cannot traverse non-AVB bridges. Alternatively, P2P SYNC and EventOffset may be sent with ordinary MAC addresses, in which case they will traverse non-AVB bridges and the intervening bridges will be modeled as link delay. This provides graceful degradation in mixed networks.