



## **Residential Ethernet: Time-of-day timer synchronization**

Maintained by David V James

July 2005


IEEE 802.3 RE Study Group  
San Francisco

1

This is a cumulative RE slide, with many slides created by DVJ.  
Credit is due to many others, whose reviews and comments drove this content.  
A few of those names, who are willing to support a majority of these concepts,  
are acknowledged above.

# Cascaded TOD synchronization

Wall-clock distribution model



July 2005

IEEE 802.3 RE Study Group  
San Francisco

2

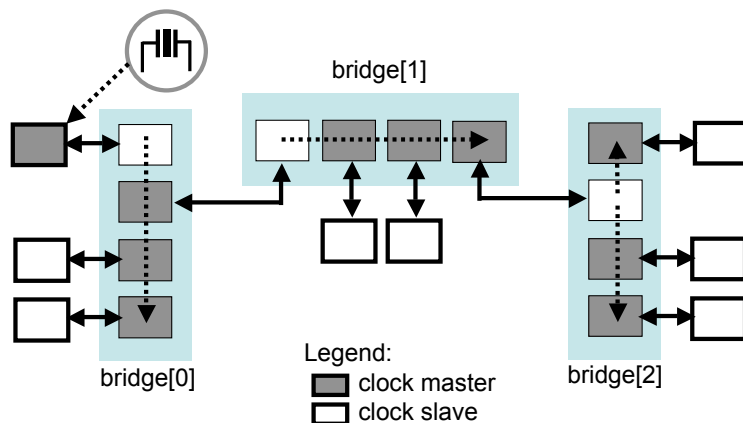
The clock master is the station whose time-of-day clock is the reference.  
From a logical perspective, the clock master broadcasts the current time-of-day to the attached stations.

From a physical perspective, a multicast time distribution is inaccurate:

- 1) There may be source transmission delays.
- 2) There may be bridge forwarding delays.

# Cascaded TOD synchronization

Cascaded adjacent-synchronization hierarchy



July 2005

IEEE 802.3 RE Study Group  
San Francisco


3

To simplify the problem, the specification defines:

- grand clock-master (grand master) selection
- clock-value synchronization (adjacent "ports" only)

## Adjacent-station synchronization

Offset value adjustments



July 2005

IEEE 802.3 RE Study Group  
San Francisco

4


In the next cycle, previously sampled values are communicated to one's neighbors.

In the case of stationB, the aTx and aRx values must be sent, since these were measured by stationA and are not known to stationB.

The value of bTx is (in concept) known to stationB and need not be transmitted. However, for simplicity, transmission of this value allows it to be more easily affiliated with the same-cycle indexed aRx value.

# Adjacent-station synchronization

## Differential rate adjustments



July 2005

IEEE 802.3 RE Study Group  
San Francisco

5


In the next cycle, previously sampled values are communicated to one's neighbors.

In the case of stationB, the aTx and aRx values must be sent, since these were measured by stationA and are not known to stationB.

The value of bTx is (in concept) known to stationB and need not be transmitted. However, for simplicity, transmission of this value allows it to be more easily affiliated with the same-cycle indexed aRx value.

# Adjacent-station synchronization

Snapshot value distribution



July 2005

IEEE 802.3 RE Study Group  
San Francisco


6

In the next cycle, previously sampled values are communicated to one's neighbors.

In the case of stationB, the aTx and aRx values must be sent, since these were measured by stationA and are not known to stationB.

The value of bTx is (in concept) known to stationB and need not be transmitted. However, for simplicity, transmission of this value allows it to be more easily affiliated with the same-cycle indexed aRx value.

# Minimalist HW design model



July 2005

IEEE 802.3 RE Study Group  
San Francisco

7

How can you be precise, when real-world PHYs typically have FIFOs?

Several strategies are possible:

FIFOs add ambiguity (existing hardware)


The MAC arranges for FIFOs to be nearly empty, at critical times

The PHY signals the actual clockSync arrival/departure times

What do you do when the client requires a distinct format?

SIMPLE: conversion from binary is not difficult.

# Uncompromised precision



July 2005

IEEE 802.3 RE Study Group  
San Francisco

8

How can you be precise, when real-world PHYs typically have FIFOs?

Several strategies are possible:

- FIFOs add ambiguity (existing hardware)

- The MAC arranges for FIFOs to be nearly empty, at critical times

- The PHY signals the actual clockSync arrival/departure times

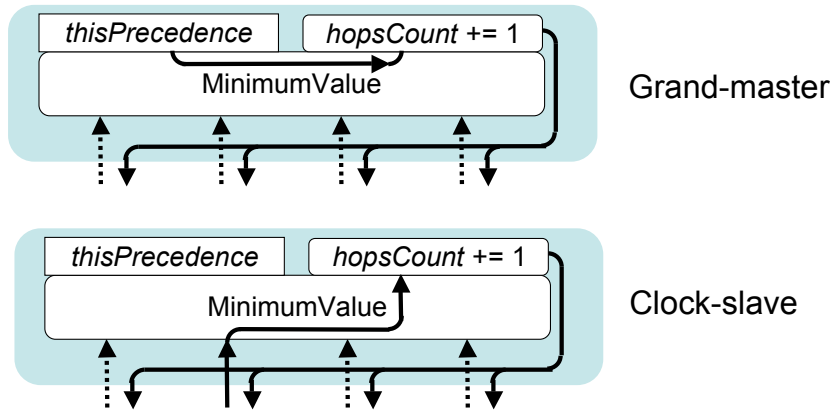
What do you do when the client requires a distinct format?

SIMPLE: conversion from binary is not difficult.

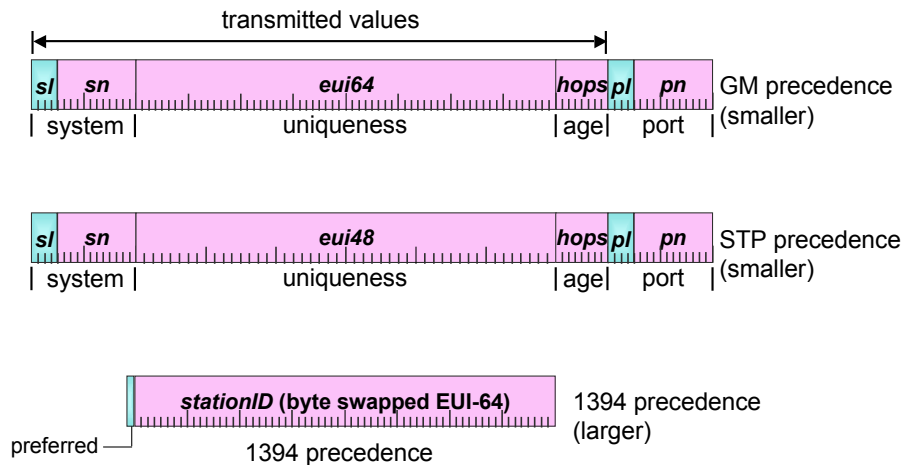


# Grand-master selection protocol

Grand-master precedence comparisons



# Grand-master precedence



July 2005

IEEE 802.3 RE Study Group  
San Francisco

10

The precedence numbers must be unique, so that only one clock master will be selected.

The existing spanning tree protocol selection criteria is sufficient, so another is not needed.

To communicate preferences, a *sp* (station priority) value is provided. This overrides the MAC address, allowing users to assert their preferences. This weighting can be accessed through the MIB.

For stations with equal *sp/systemID* values, the *stationId* becomes the tie breaker.

For stations with equal *systemTag/systemID* values, the *hops* becomes the tie breaker. This resolves grand-master in favor of smallest hops

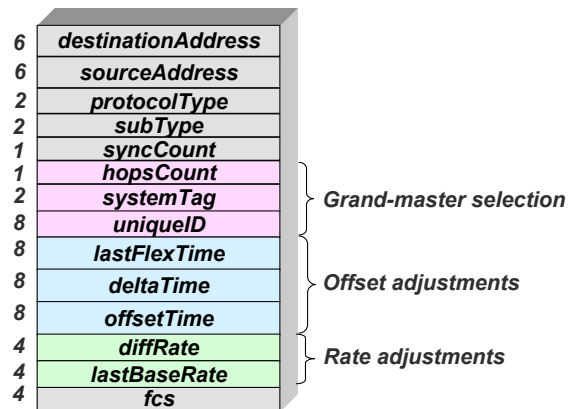
For ports on a bridge, the *portLevel* (*pl*) and *portNumber* (*pn*) are similarly used as tie breakers that select between available ports. This resolves same-hops messages on the basis of the arrival ports.

The lowest numerical value has the highest precedence.

Default weighting is the largest numerical value.

The setting of lower-valued weights is a higher level protocols and is beyond the scope of this standard.

# Frame format



July 2005

IEEE 802.3 RE Study Group  
San Francisco

11

What is contained within each frame?

Frames are sent from the listener to the endpoint talker:

The path to the talker need not be know to the listener.

Distinct coding allows the frame to be “hijacked” at bridge ports.

The local listener address reduces local-talker resource allocations

This still works with >2 station links (802 wireless, 802.17, etc.)

The endpoint listener address must (of course) also be known.

The connection includes a talker “plug”, ala RSVP and 1394


The “required bandwidth” may need to include:

frameCount

byteCount

Since the frameCount may be associated with link-specific overhead (interframe gap, synchronization, etc.).

## Timing specifics...



(from IEEE 1588-2002, subclause D.1.1, page 127)

July 2005

IEEE 802.3 RE Study Group  
San Francisco

12

What (exactly) is the frame arrival/departure time?

Depends on the physical layer details.

Some are already specified, in IEEE 1588.

Parallel-bit transmission schemes may need clarifications.

1G CAT-5

10G (in general?)

## Basic requirements

- KISS (keep it simple, stupid)
  - Delayed snapshot processing
  - Periodic symmetric transmissions
  - Etc., etc.
- NTP (RFC-1305) and SNTP (RFC-2030)
  - Definition of the 64-bit time-of-day value
- For a detailed summary, see:
  - <http://dvjames.com/esync>
  - [dvjTimeSync2005Nov16.pdf](#) (or later revision)

While this application is a bit different from others, things can be leveraged. There is no point in being different, when concepts can be leveraged. However, the details are customized for the RE-specific environment.



**Backup slides for  
Residential Ethernet:  
Time-of-day timer synchronization**

Maintained by David V James


July 2005

IEEE 802.3 RE Study Group  
San Francisco

14

This is a cumulative RE slide, with many slides created by DVJ. Credit is due to many others, whose reviews and comments drove this content. A few of those names, who are willing to support a majority of these concepts, are acknowledged above.

## House reference clock



July 2005

IEEE 802.3 RE Study Group  
San Francisco

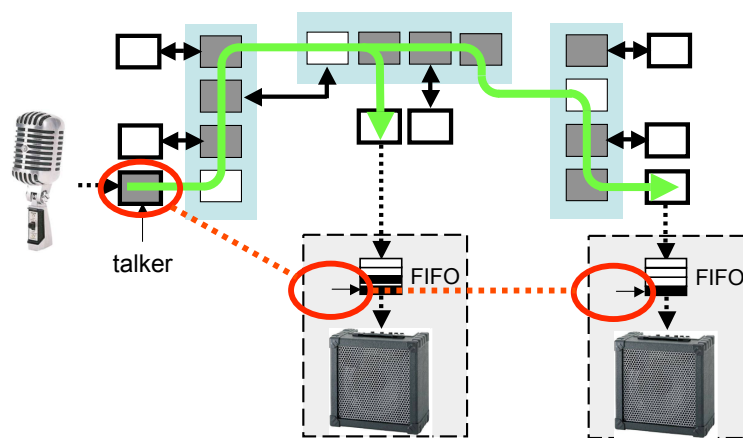
15

In support of synchronous transfers, all RE devices are assumed to have the same impression of time.

For this presentation, assume an 8kHz cycle time, although a decision on this value has not been finalized.

Requirement: 8kHz cycle frequencies are locked and the “same

## Precise time synchronization



July 2005

IEEE 802.3 RE Study Group  
San Francisco

16

Data arrives early; data is gated at its presentation time.  
(Each frame has a presentation time stamp.)

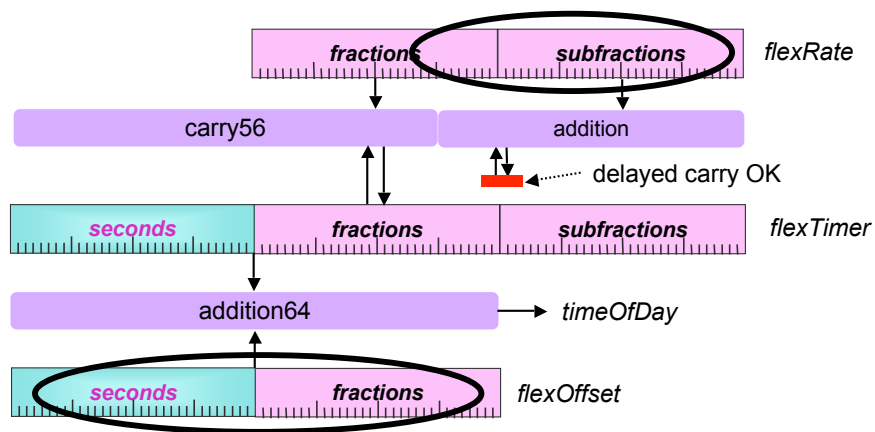
Bridge reclocking has a relatively modest clock-sync accuracy requirement, where microsecond deviations could be acceptable.

Source-data and presentation-data clocking requirements are more severe.

- 1) Frequency drift is unacceptable, since dropped/replicated values are audible.
- 2) Presentation time jitter is sub nanosecond, based on slew rates and D/A accuracies.



# Adjustable *timeOfDay* timer



July 2005

IEEE 802.3 RE Study Group  
San Francisco

17

The rate can be adjusted, by updates to flexRate.

The offset can be adjusted, by changes to flexOffset.

How is a fractions-of-second timer implemented?

Clocks are typically multiples of something else!

The clock can tick at a natural rate, and the tick-size need-not be "one".

For example, consider a clock that is updated with a 16ns clock.


The update value is 68.719476736

Since the LSB is insignificant, the carry can be delayed

The more significant 12-bit rate field is a constant

The less significant bits provide precise rate adjustment, for precise tracking

# Time-of-day format options



July 2005

IEEE 802.3 RE Study Group  
San Francisco

18

What is synchronized?

Preferably a binary seconds and fractions-of-seconds value.

Doesn't overflow within our lifetimes.

Time resolution induced errors are insignificant.

Easily added and subtracted.

Readily converted to other formats...

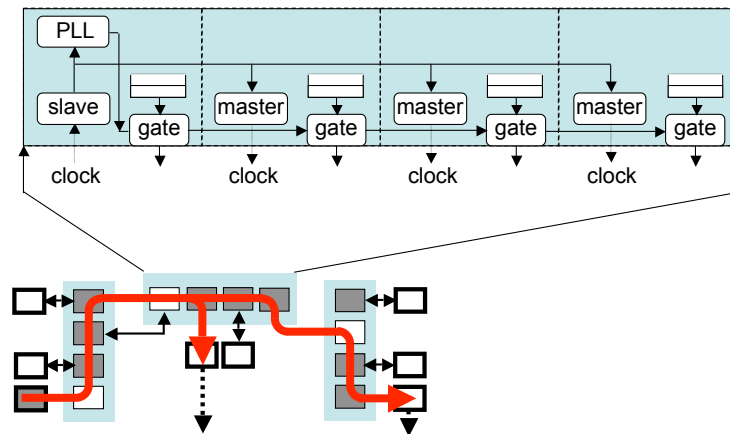
Other formats are also possible:

because we have 10 fingers

a multiple of the cycle frequency

....

## Passby PLLs (proposal 1)



July 2005

IEEE 802.3 RE Study Group  
San Francisco

19

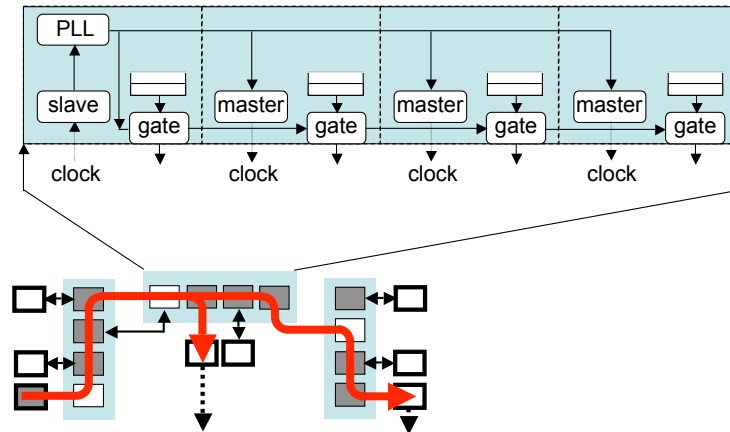
There are two proposals for PLL functionality within bridges.

The considerations revolve around the effects of cascaded PLLs.

The passby option could use a PLL for gating frame transmissions  
(although even this may be unnecessary, studies are TBD)

The assumption is that additional PLLs would not help, but would actually increase the effective clock-jitter noise at the output of endpoint PLLs.

## Passthrough PLLs (option 2)



July 2005

IEEE 802.3 RE Study Group  
San Francisco

20

The passthrough option would use a PLL for gating frame transmissions.

This (or another PLL) would also be used to “clean up” the clock reference.

If this approach is taken, the properties of intermediate PLLs must be specified.

There is a possible concern (particularly for high-fidelity audio) that the bridges' PLLs might do more harm than good.

Both options should be considered.

Experts, calculations, and simulations are encouraged, so that the impacts/values are better understood before either one is selected.