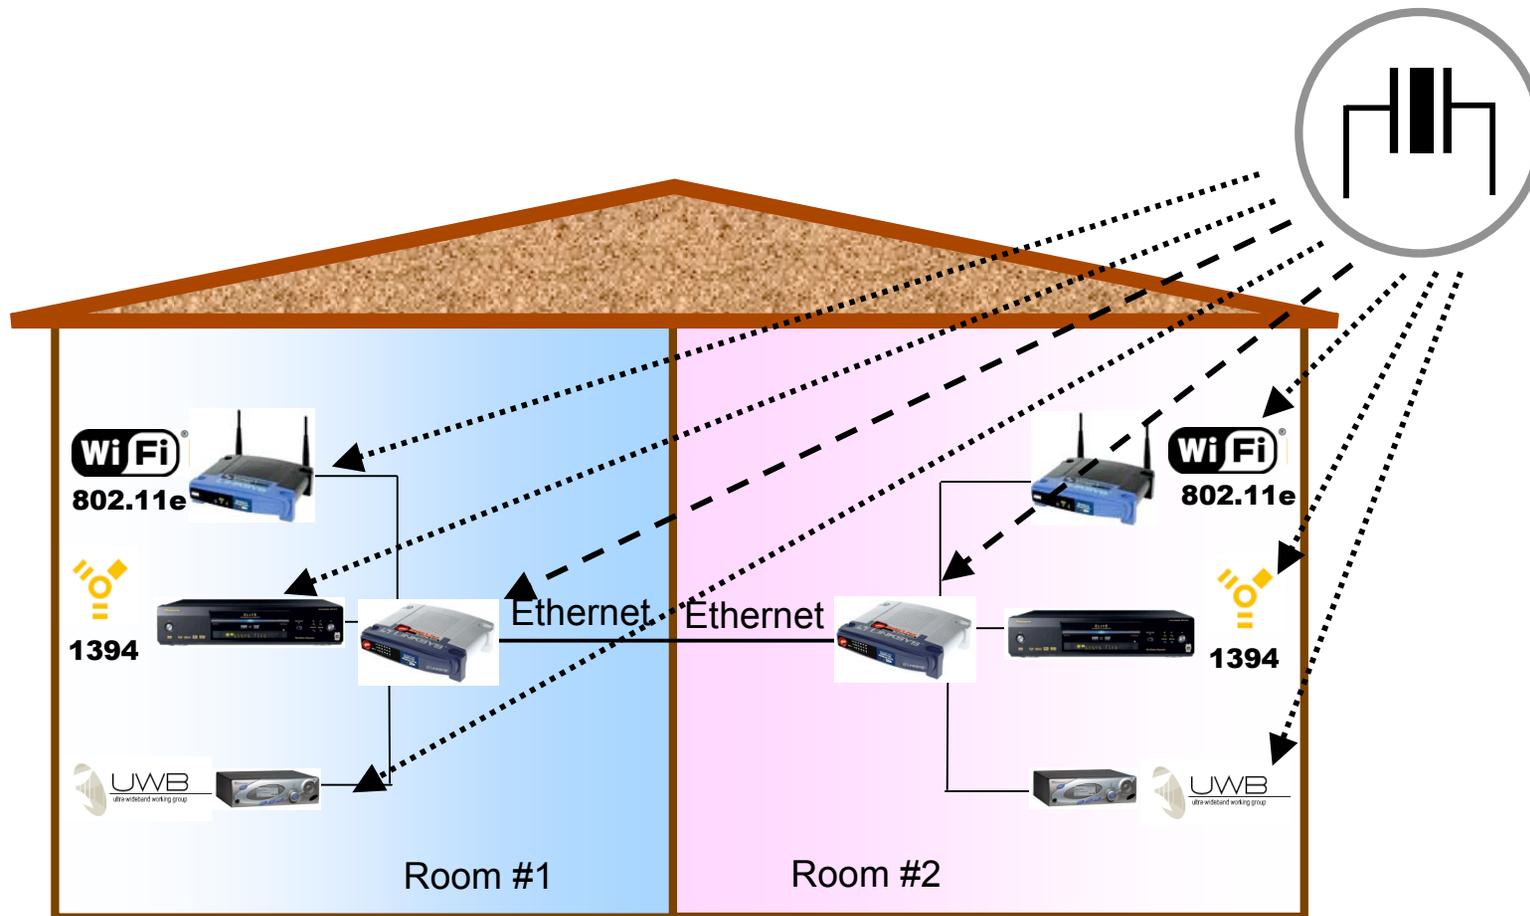

Clock synchronization

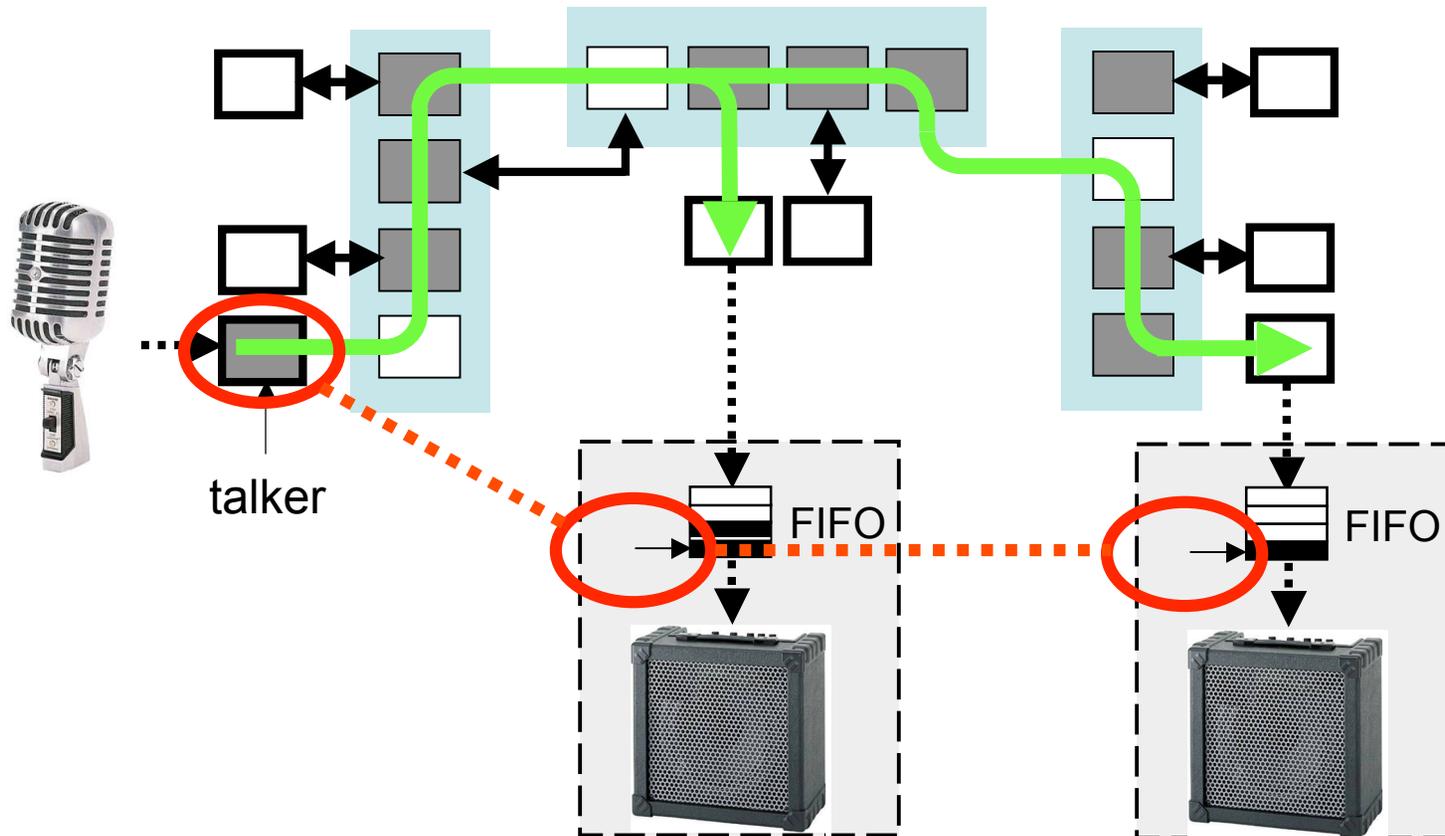
(a Residential Ethernet SG presentation)

David V James JGG
Alexei Beliaev Gibson
George Claseman Micrel

House reference clock

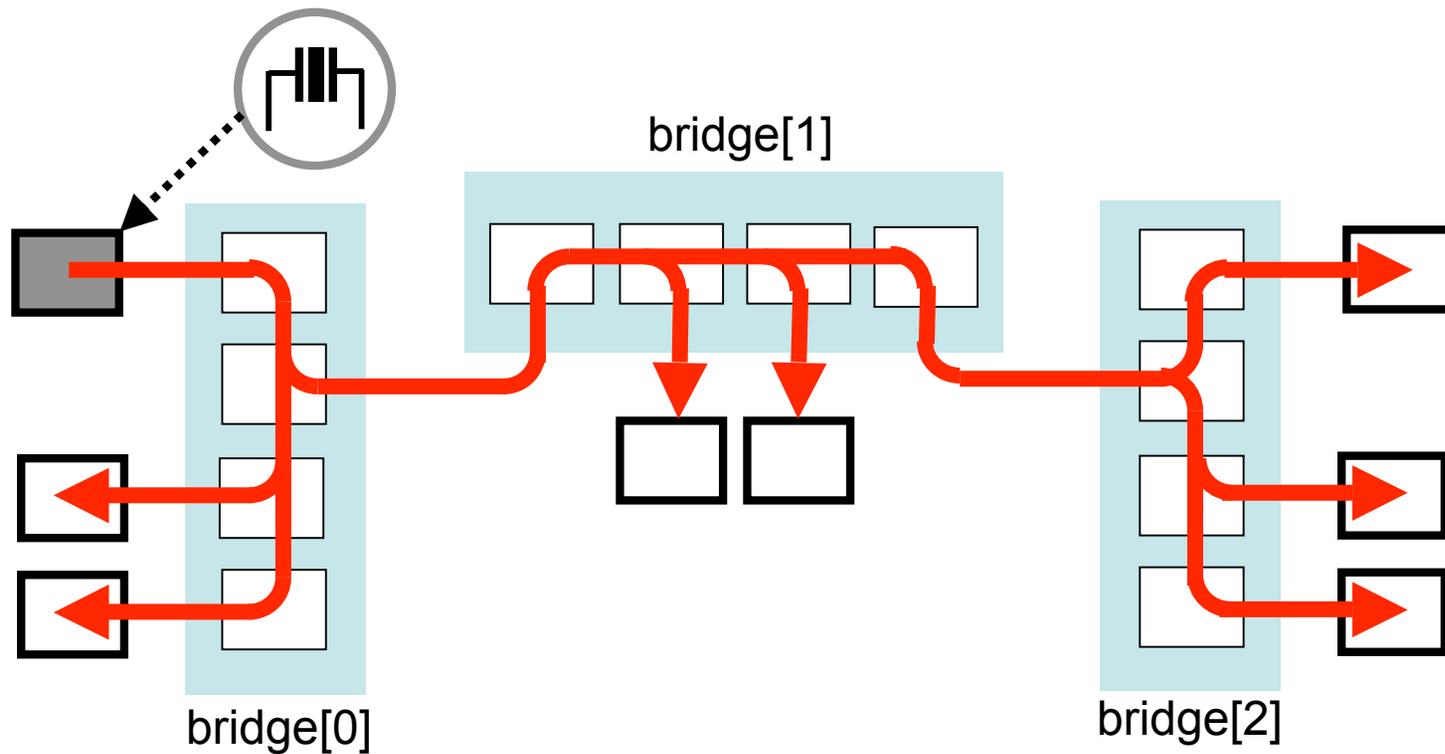


Precise time synchronization



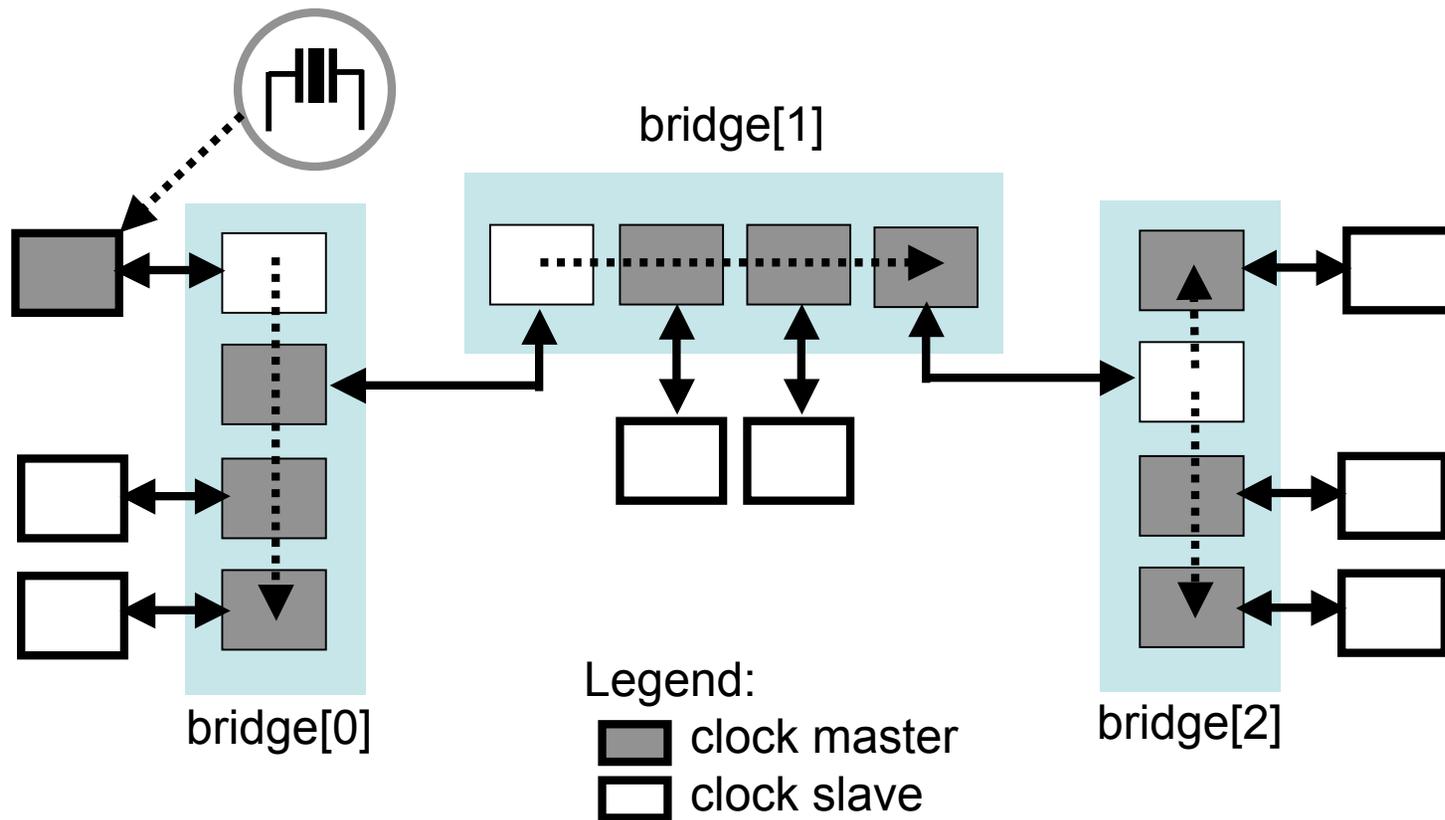
Cascaded TOD synchronization

Wall-clock distribution model



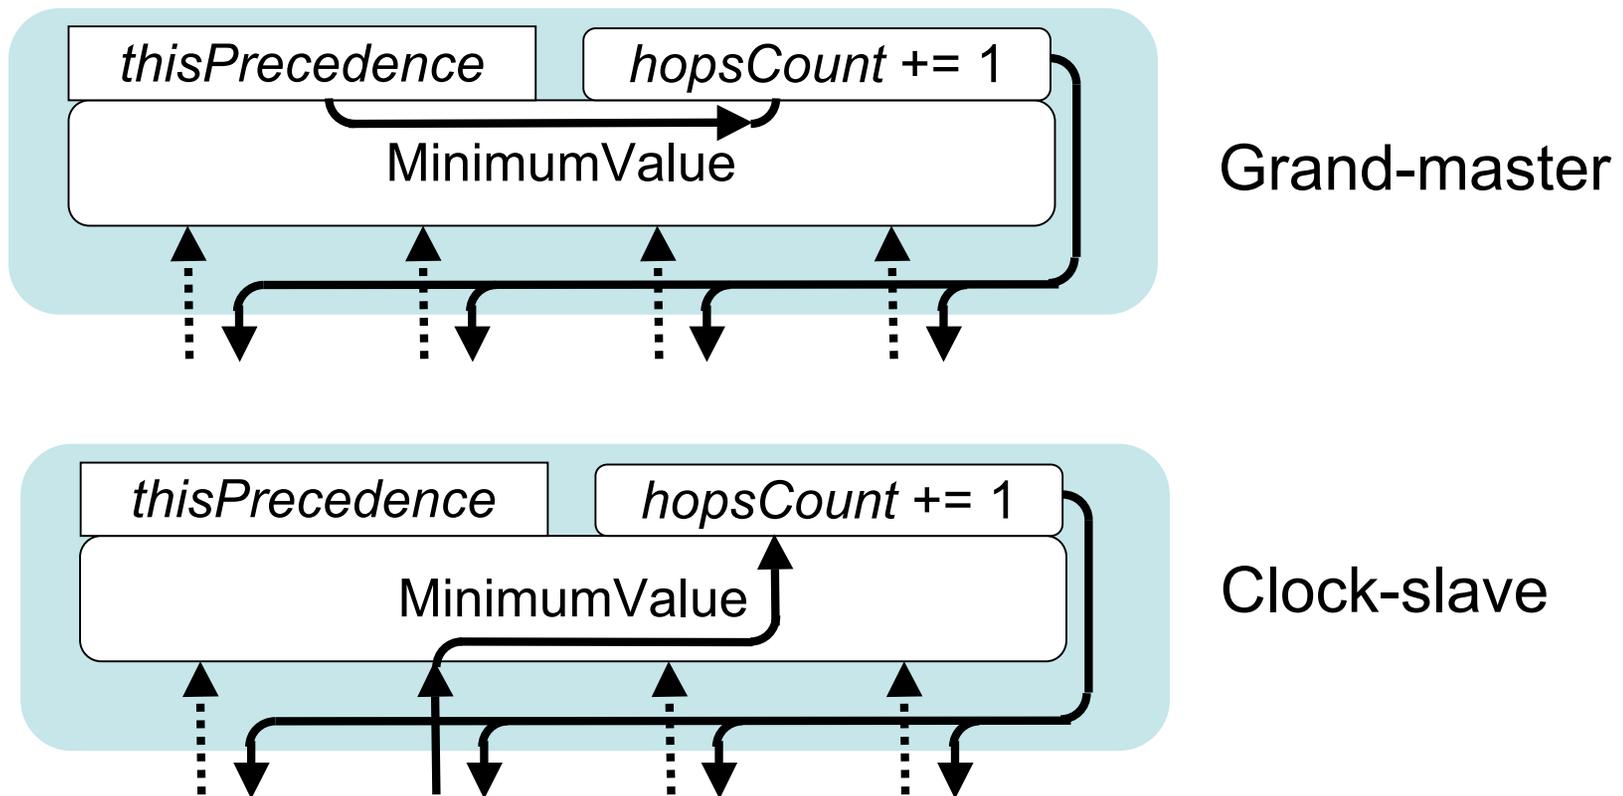
Cascaded TOD synchronization

Cascaded adjacent-synchronization hierarchy

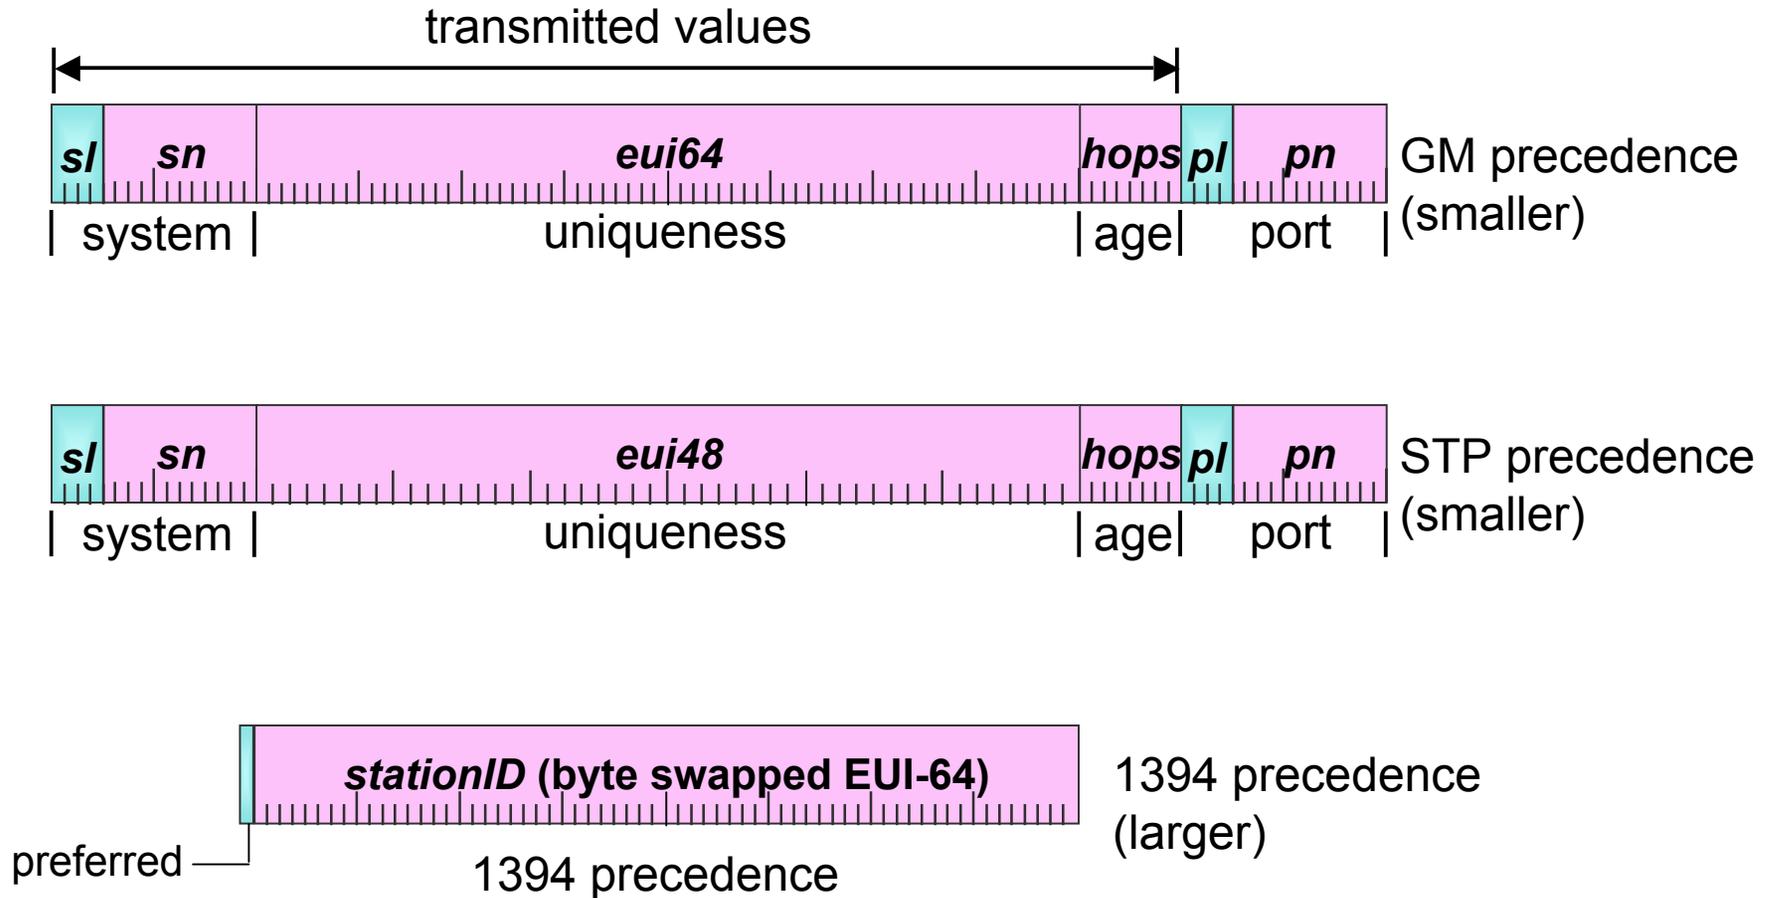


Grand master selection

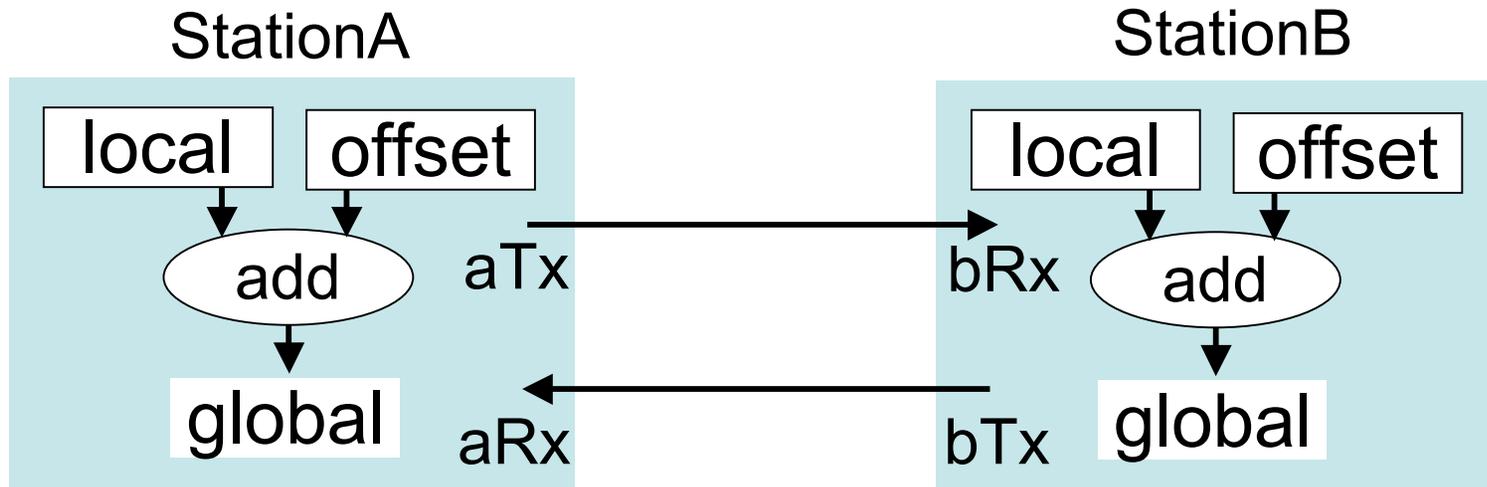
Grand-master selection protocol



Grand-master precedence

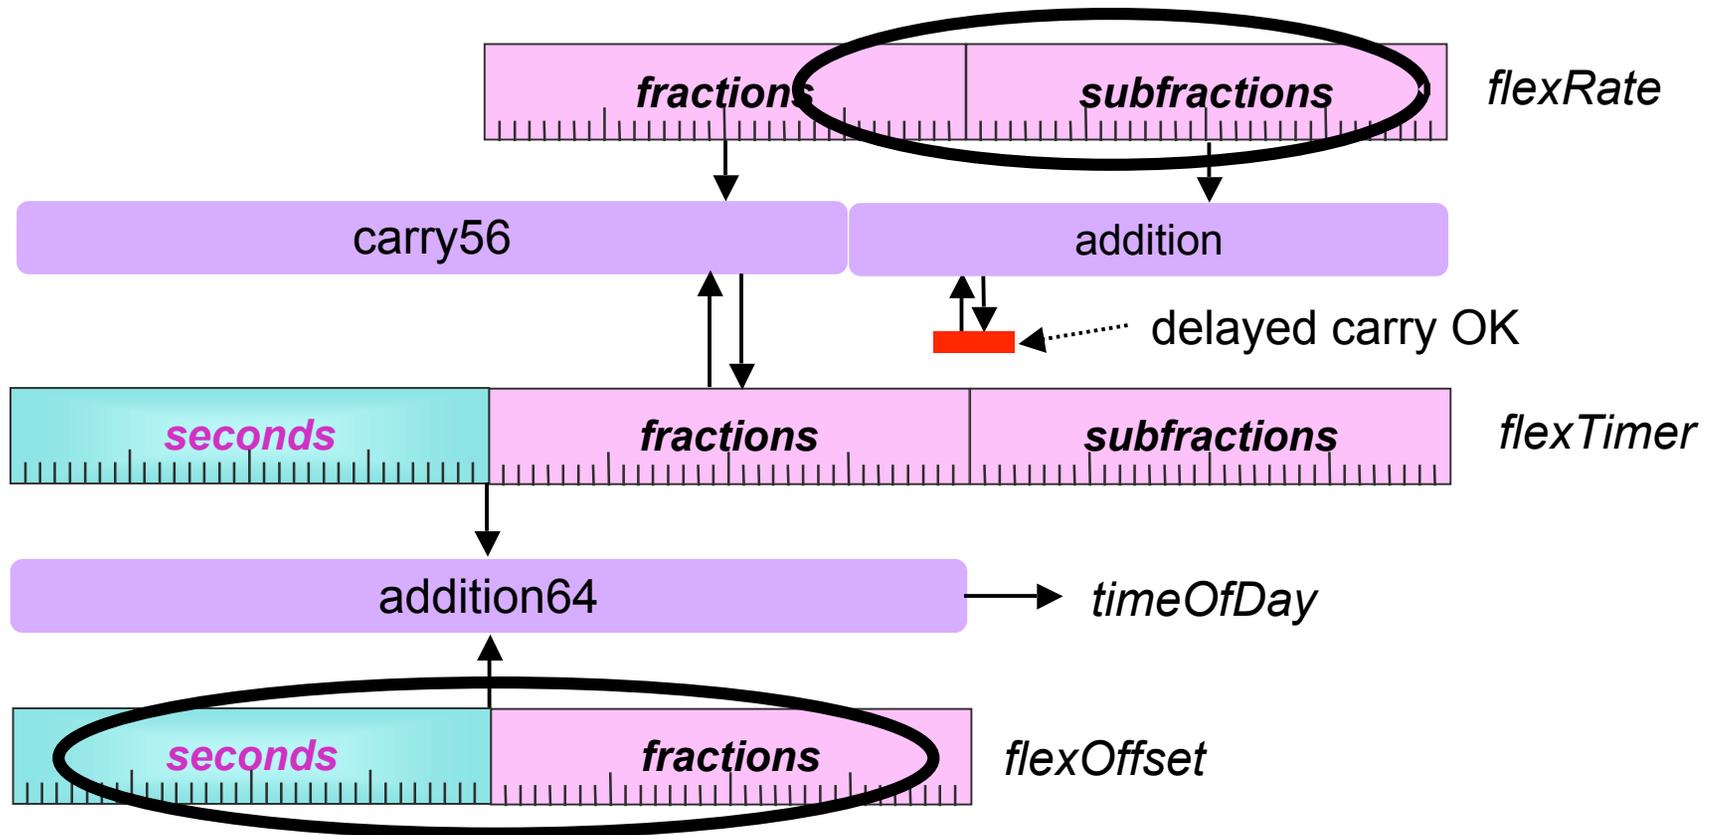


Basic snapshot assumptions



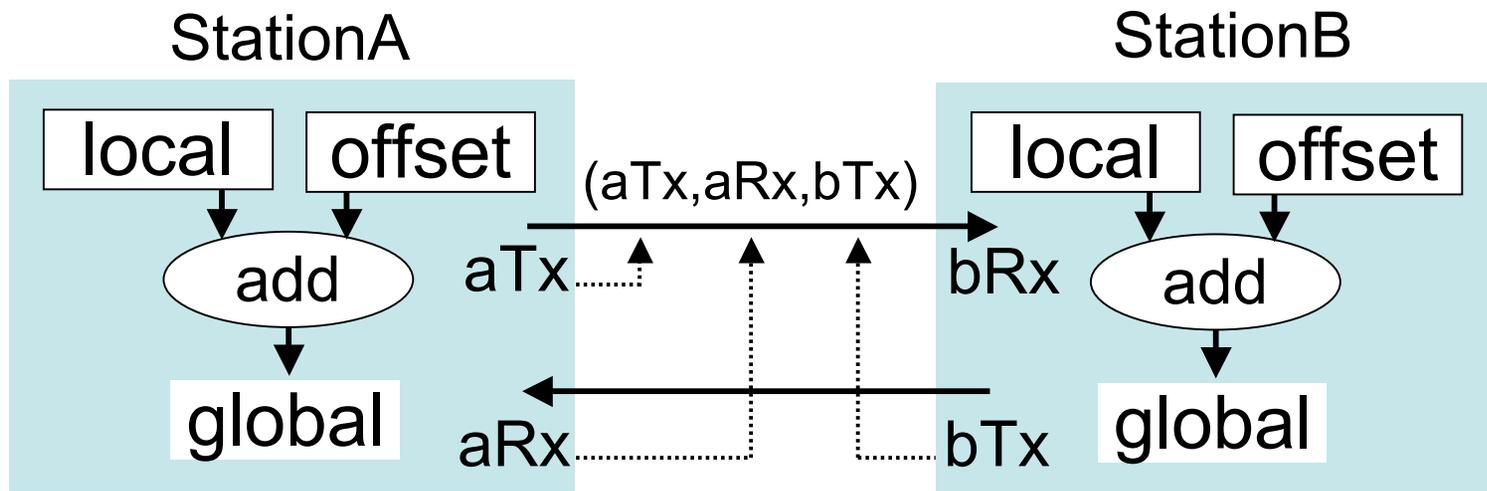
- Periodic distribution (10 ms)
 - Simple non-time-critical processing
 - Master/slave independence
- Pipelined computation
 - Enables SW-centric implementations

Adjustable *timeOfDay* timer



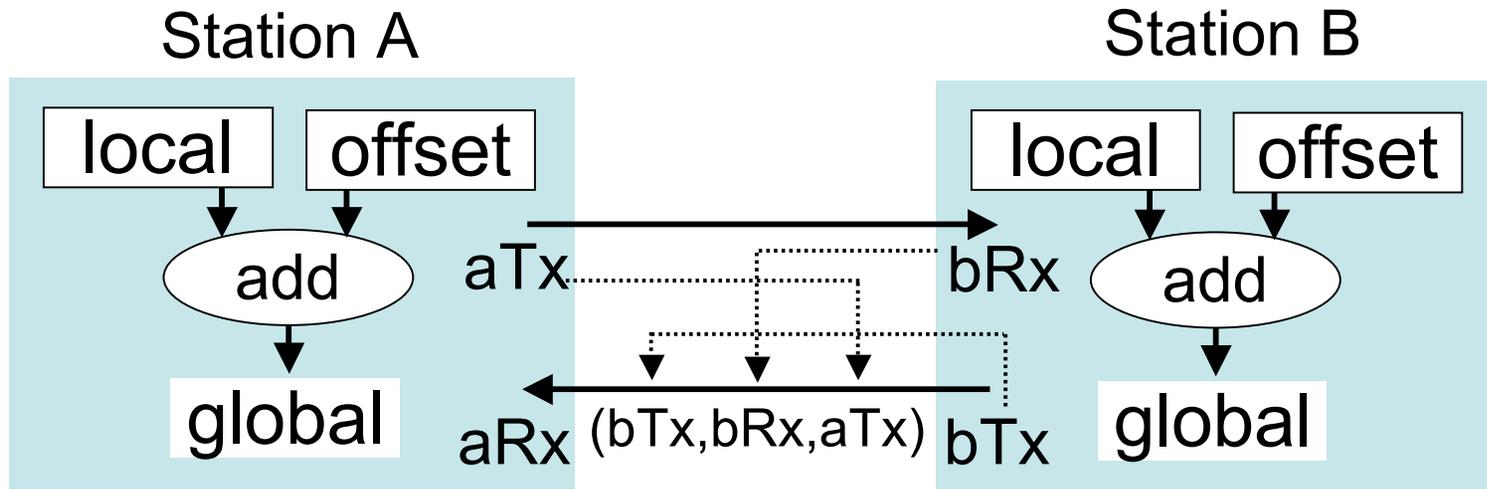
Adjacent-station synchronization

Snapshot value distribution
(information for stationB)



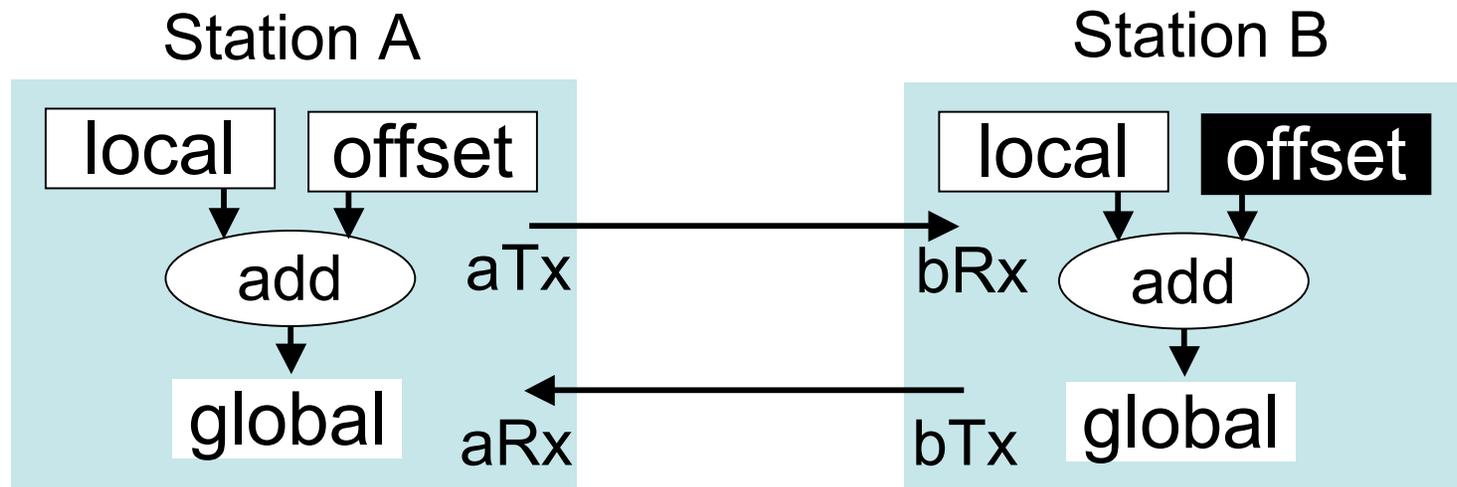
Adjacent-station synchronization

Snapshot value distribution
(information for stationA)



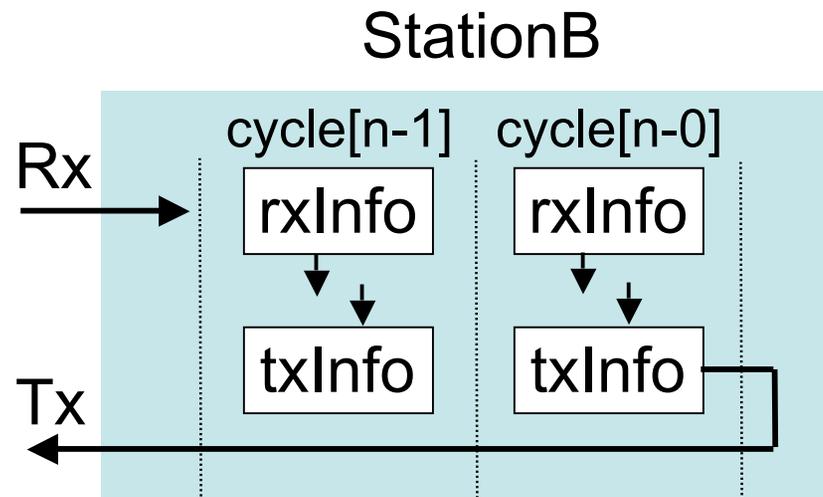
Adjacent-station synchronization

StationB offset adjustments

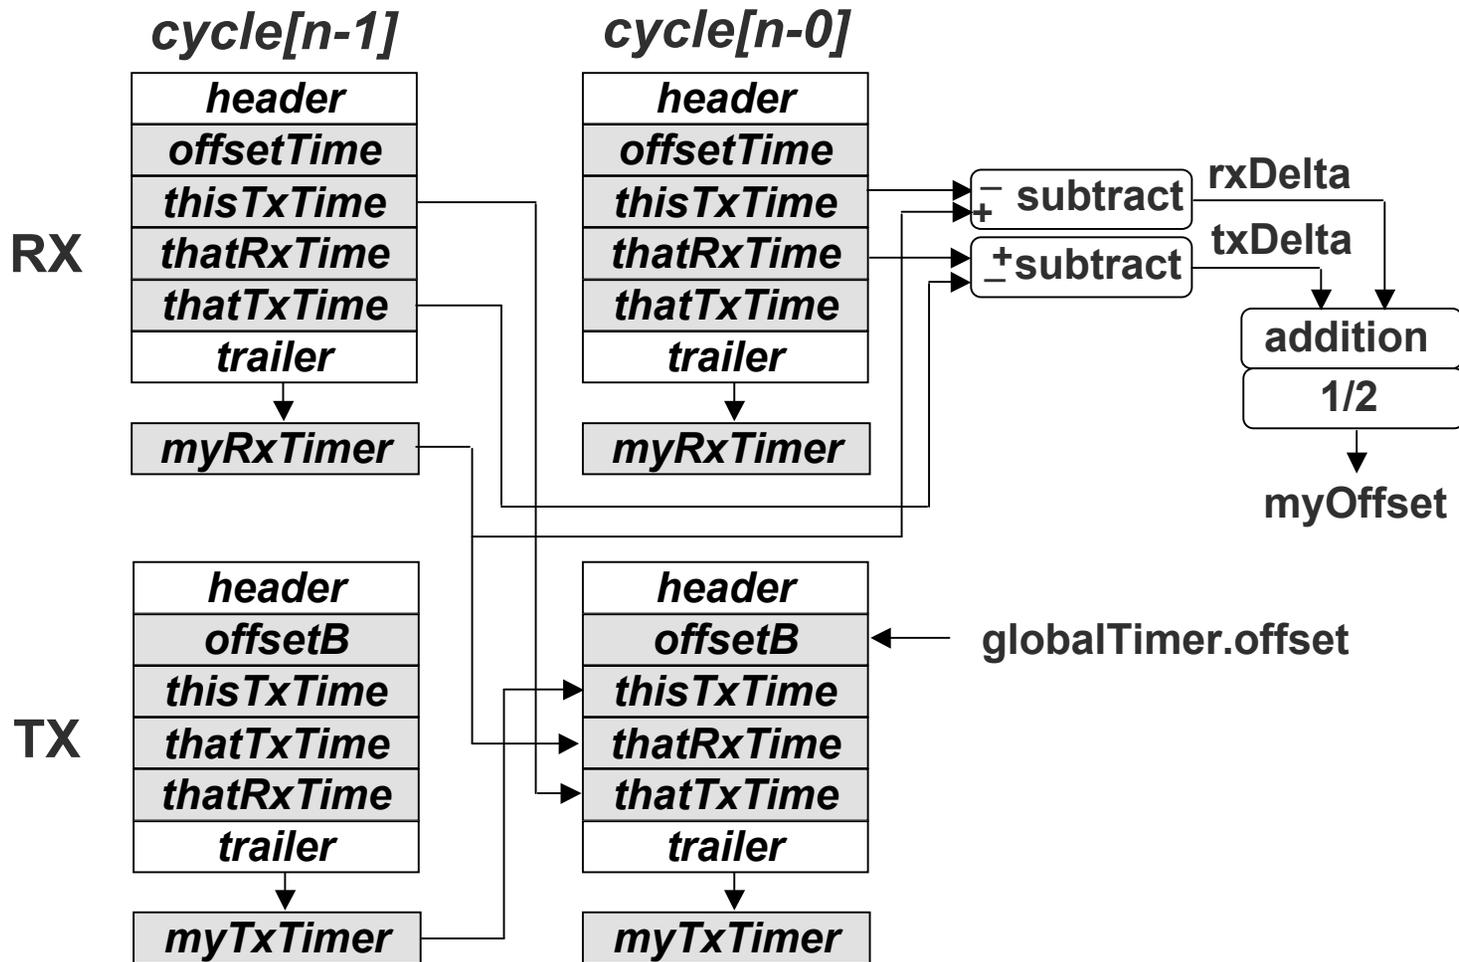


- $rxDelta = (bRx - aTx);$
- $txDelta = (bTx - aRx);$
- $clockDelta = (txDelta - rxDelta) / 2;$
- $cableDelay = (txDelta + rxDelta) / 2;$
- $offsetB = offsetA + clockDelta;$

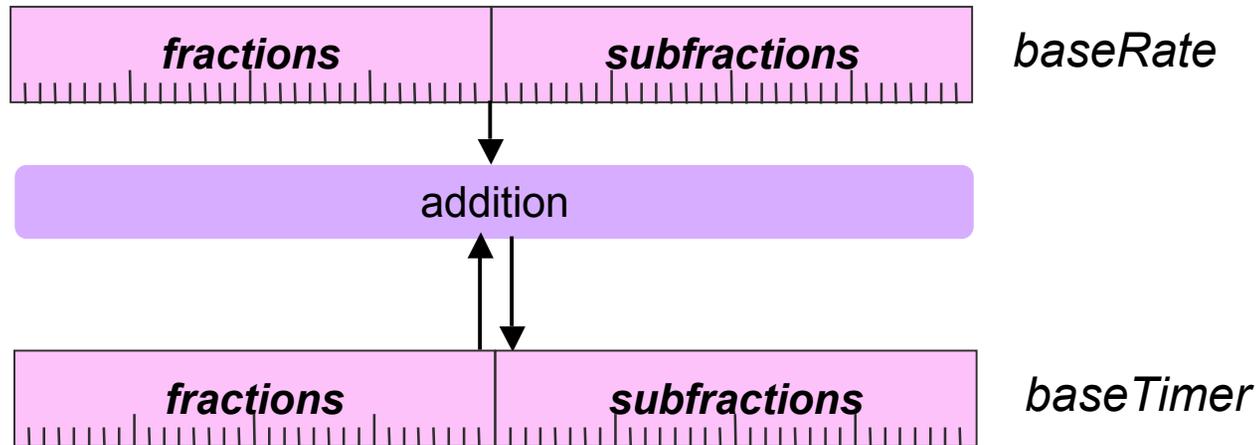
Clock slave details (1)



Clock-slave details (2)

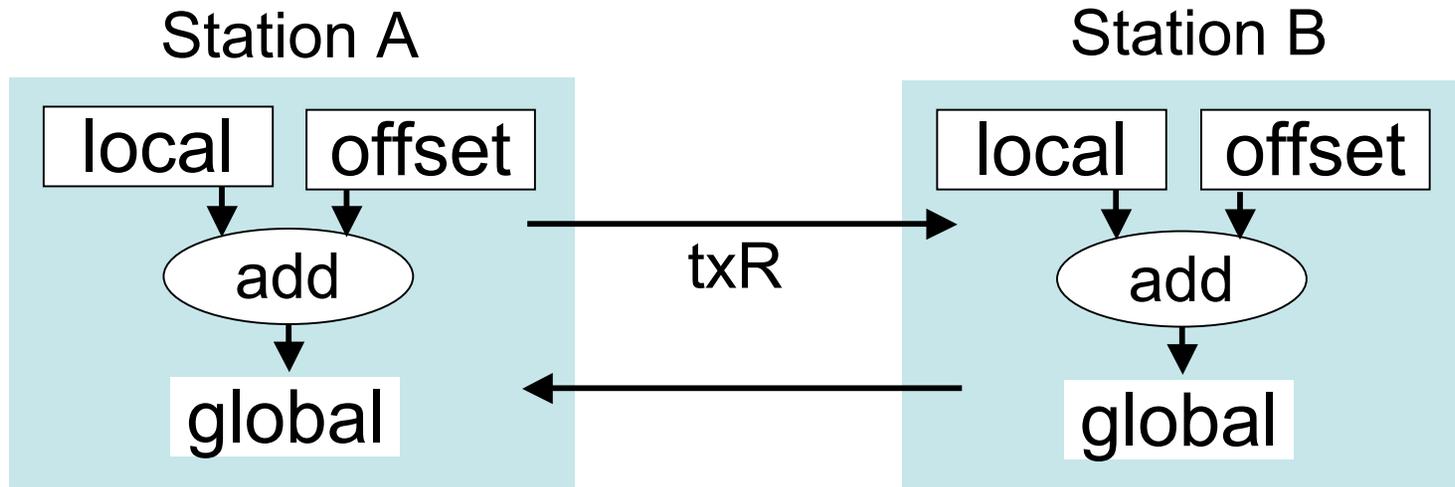


Rate-calibration timer



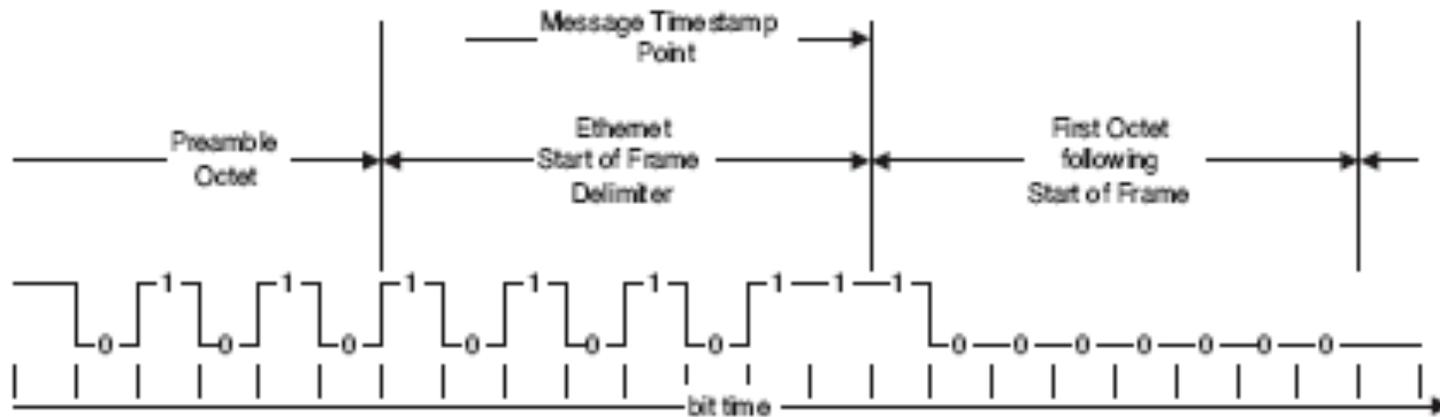
Adjacent-station synchronization

StationB rate adjustments



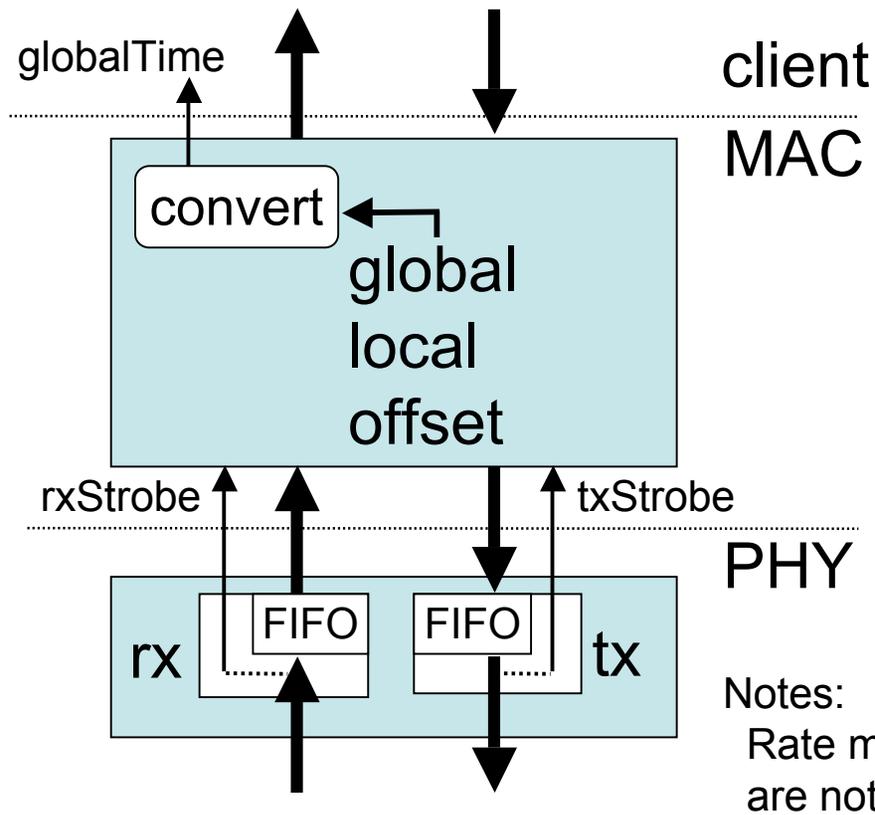
- $aDelta = (localA[n+1] - localA[n+0]);$
- $bDelta = (localB[n+1] - localB[n+0]);$
- $diffRate = (bDelta - aDelta) / aDelta;$

Timing specifics...



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

A viable design model



Residential Ethernet

(an unofficial cumulative slide set; 2005Mar11)

Maintained by David V James

Overview

(a Residential Ethernet SG presentation)

Categories of work

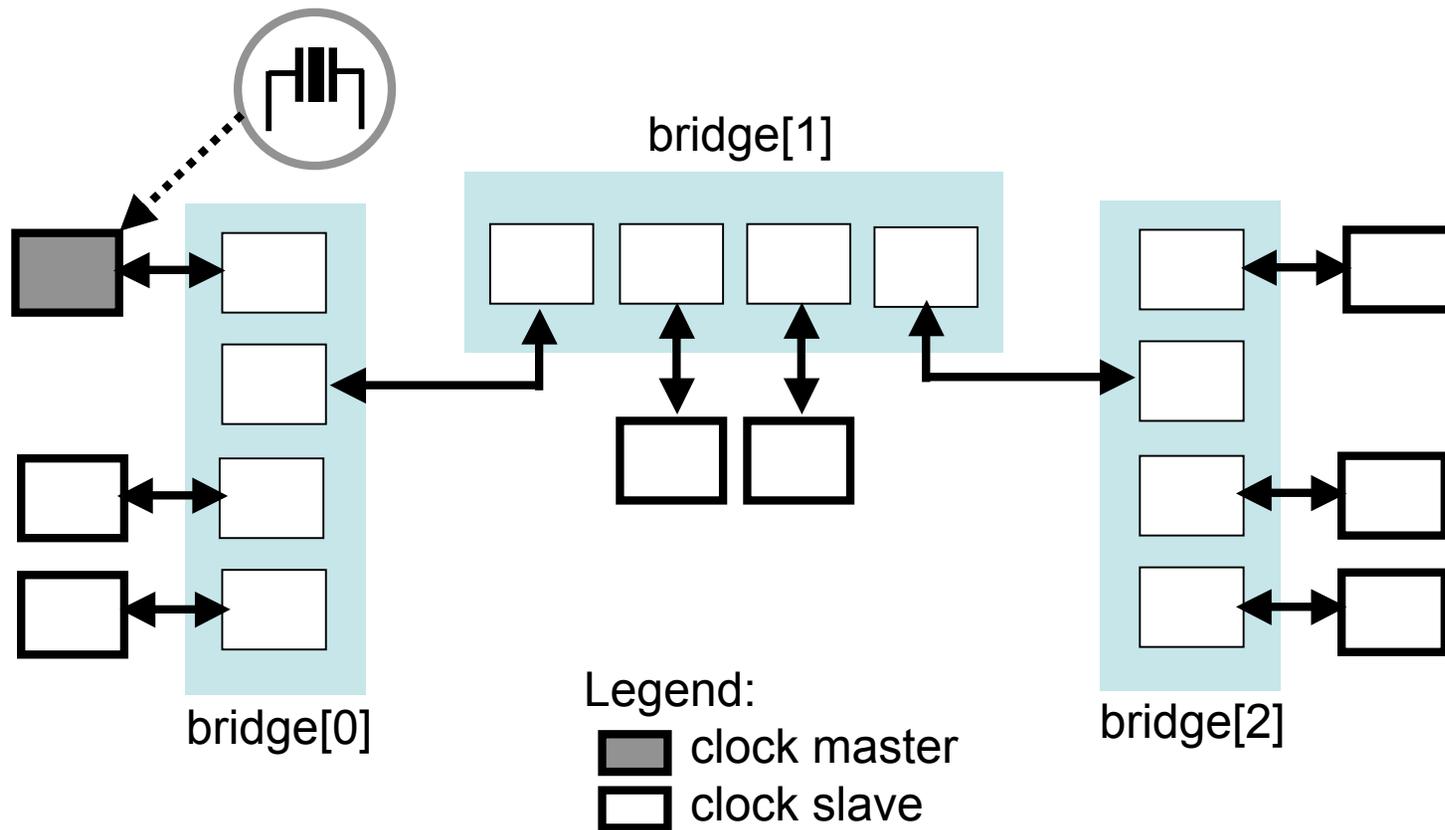
- Service discovery (out of scope)
 - Identify/control “talkers” and their available “plugs”
- Subscription (802.1 centric)
 - Establish conversation between talker and listener(s)
 - Reject unless: $linkBandwidth < linkCapacity$
- Clock synchronization
 - Synchronous reception, forwarding, and presentation
- Pacing
 - Talkers must not be well behaved
 - Bridges should “sustain” such behaviors
- Formats
 - Frame formats and content (stream IDs, time stamps)
 - Time aware service interfaces

Leveraged protocols

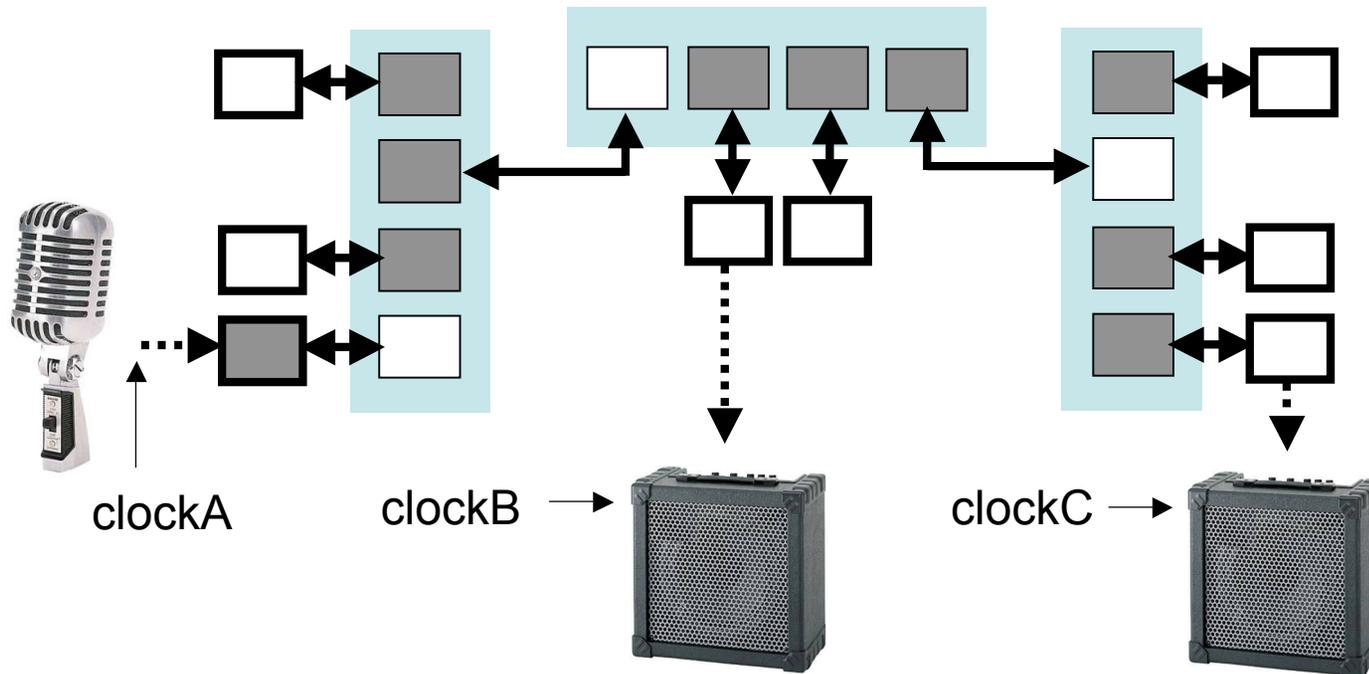
- Spanning tree protocol (STP)
 - Defines the grand-master precedence format
 - But, we use a distinct value and distribution protocol (The STP root and grand master could be distinct!)
- NTP (RFC-1305) and SNTP (RFC-2030)
 - Definition of the 64-bit time-of-day value
- IEEE 1588-2002
 - Techniques for delayed-sampling synchronization

Cascaded TOD synchronization

Physical topology constraints

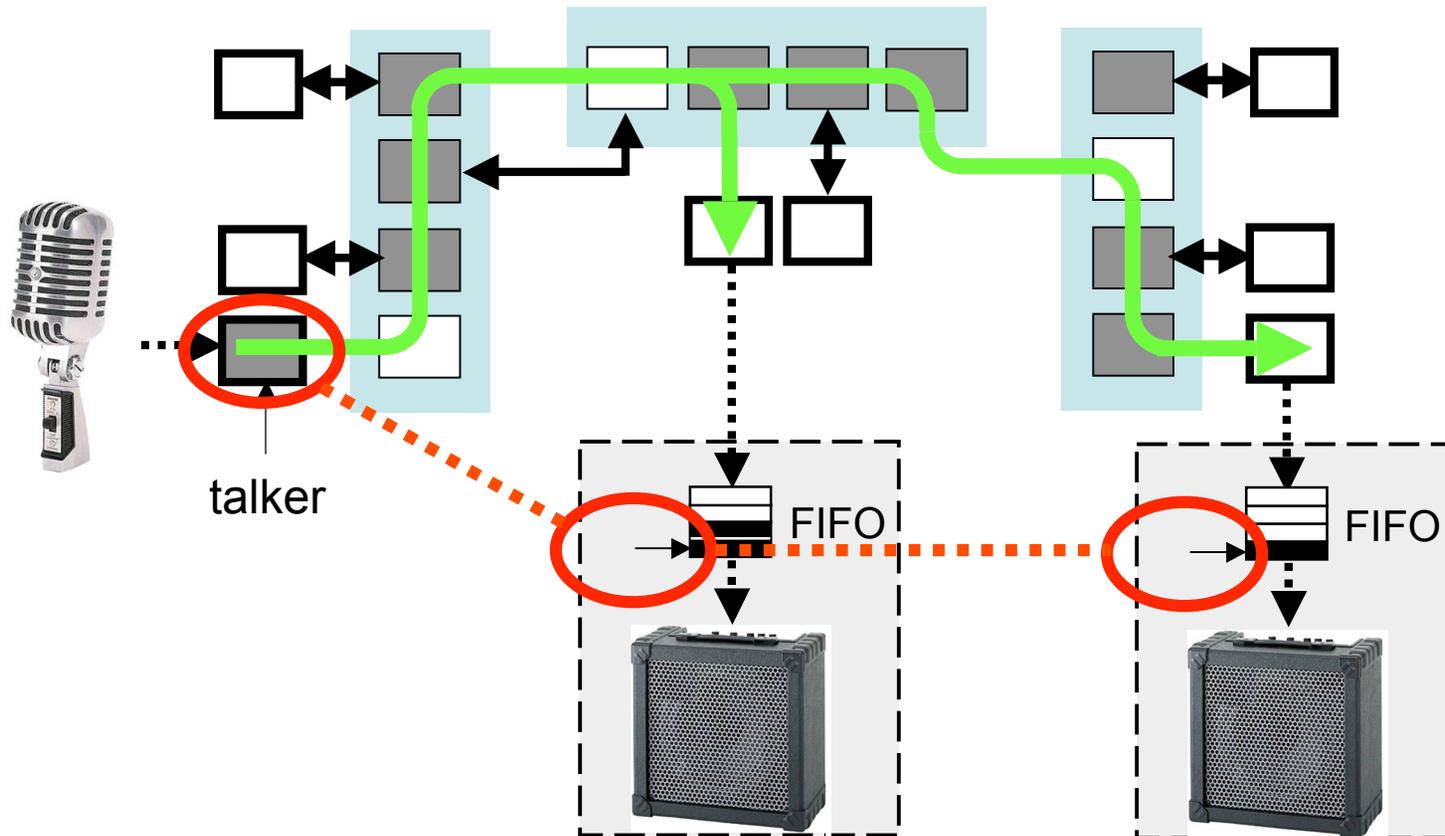


Synchronized reception/presentation

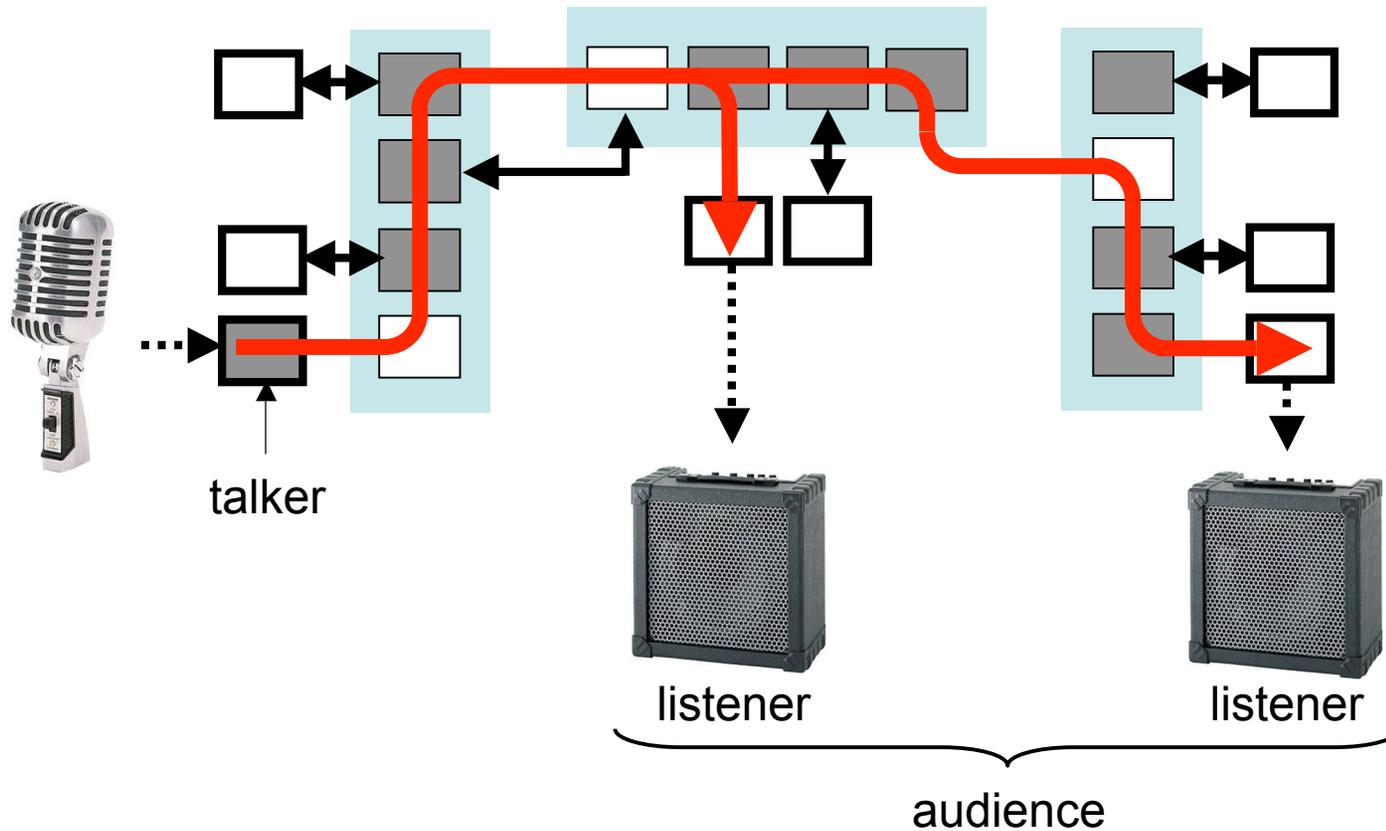


No long-term drift: clockA, clockB, clockC
Clock jitter: sub nanosecond (after PLL)

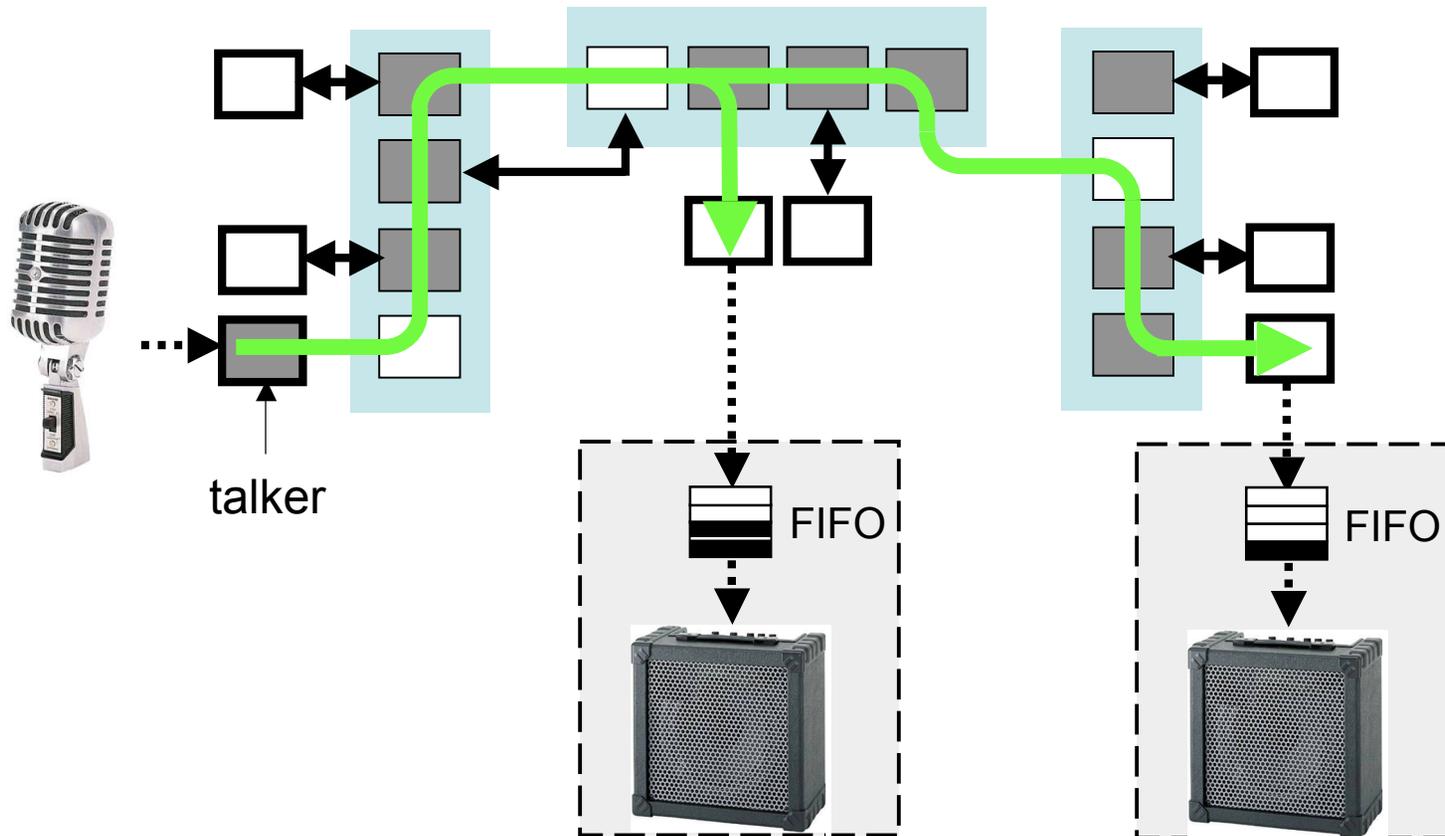
Precise time synchronization



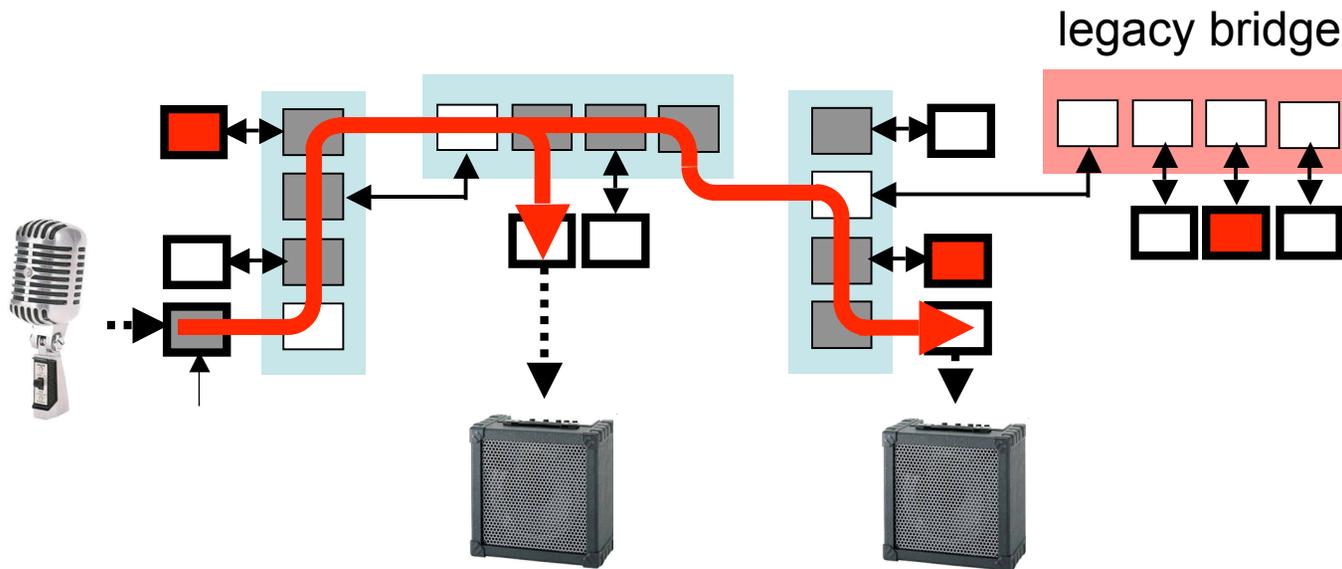
What is the application?



Subscription requirements



Ethernet compatibility (yes!)

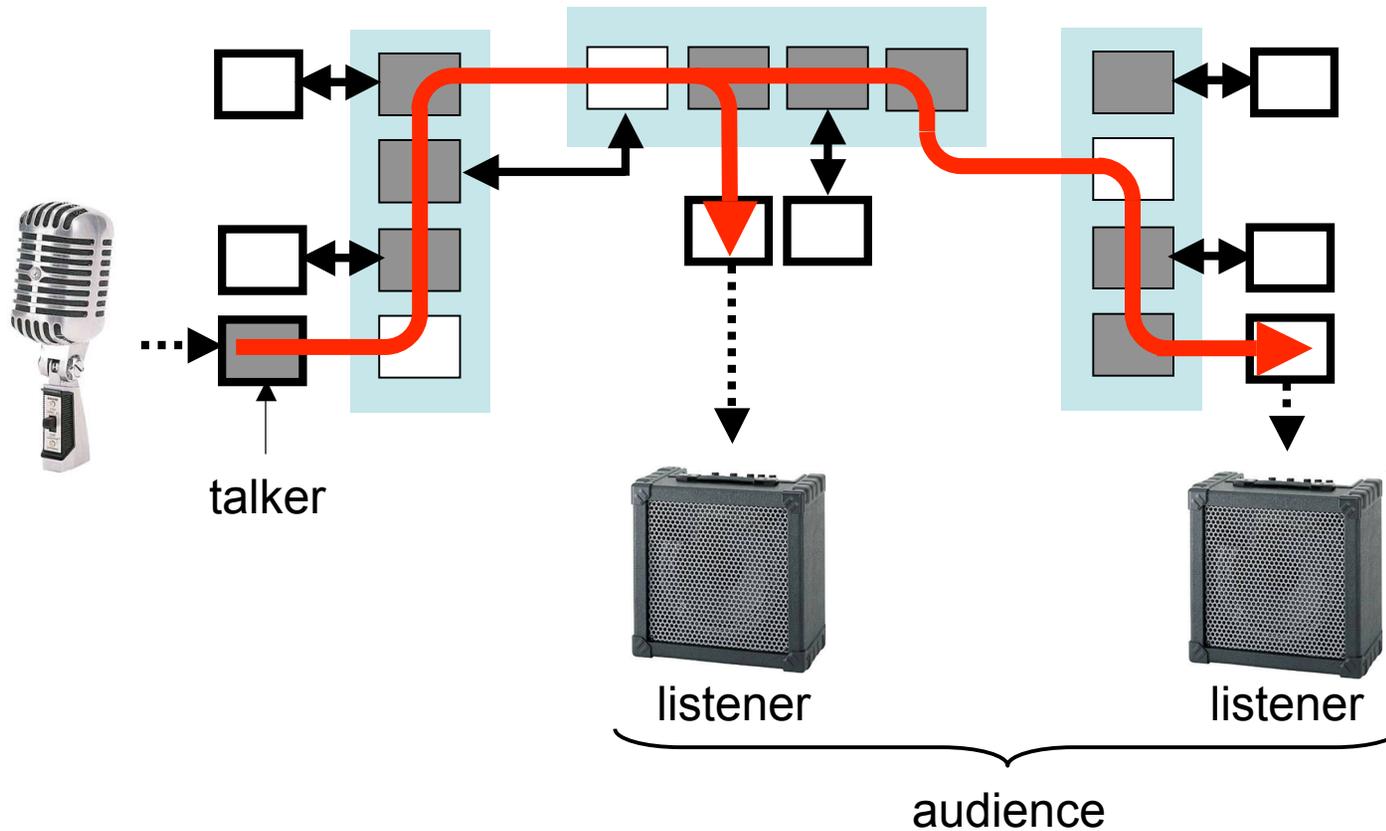


Legend:

□ □ □ □ legacy bridge

■ legacy endpoint

Vocabulary terms (1)



Vocabulary terms (2)

- audience
 - The set of listeners associated with a talker
- clock master
 - A bridge or end station that provides a link clock reference
- clock slave
 - A bridge or end station that tracks the clock reference
- grand clock master (grand master)
 - The clock master that provides the network time reference
- listener
 - A sink of a stream, such as a television or acoustic speaker
- path
 - A logical concatenation of links and bridges for a stream

Vocabulary terms (3)

- service discovery*
 - The process used to identify/control/configure talkers
- stream
 - An RE frame sequence with a distinct streamID
- subscription**
 - Establishing committed paths between talker and listener(s)
- talker
 - A source of a stream, such as a cable box or microphone

* A complete solution/vocabulary includes out-of-802 scope activities

** A complete solution/vocabulary includes out-of-802.3 scope activities

Residential Ethernet

(Problem components)

David V James JGG

Clock synchronization

(a Residential Ethernet SG presentation)

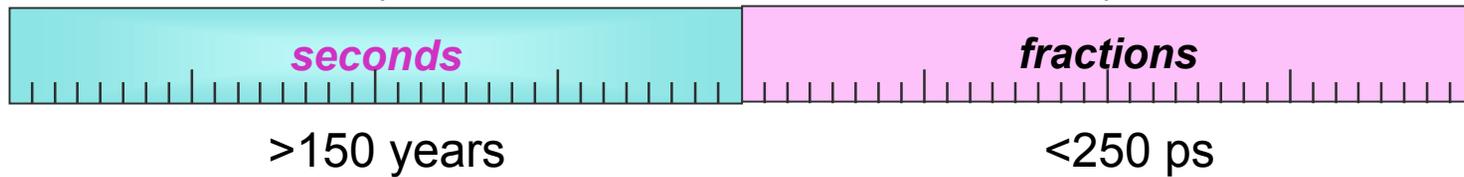
David V James JGG
Alexei Beliaev Gibson
George Claseman Micrel

Clock synchronization

What?

Time-of-day format options

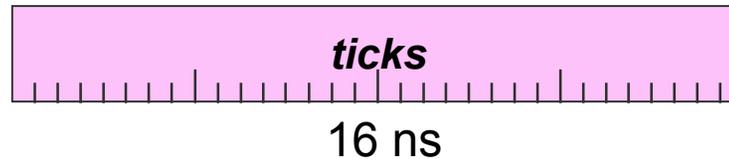
(NTP RFC-1305, SNTP RFC-2030)



(IEEE 1588)



(EPON)



OR
(...)

Format selection criteria

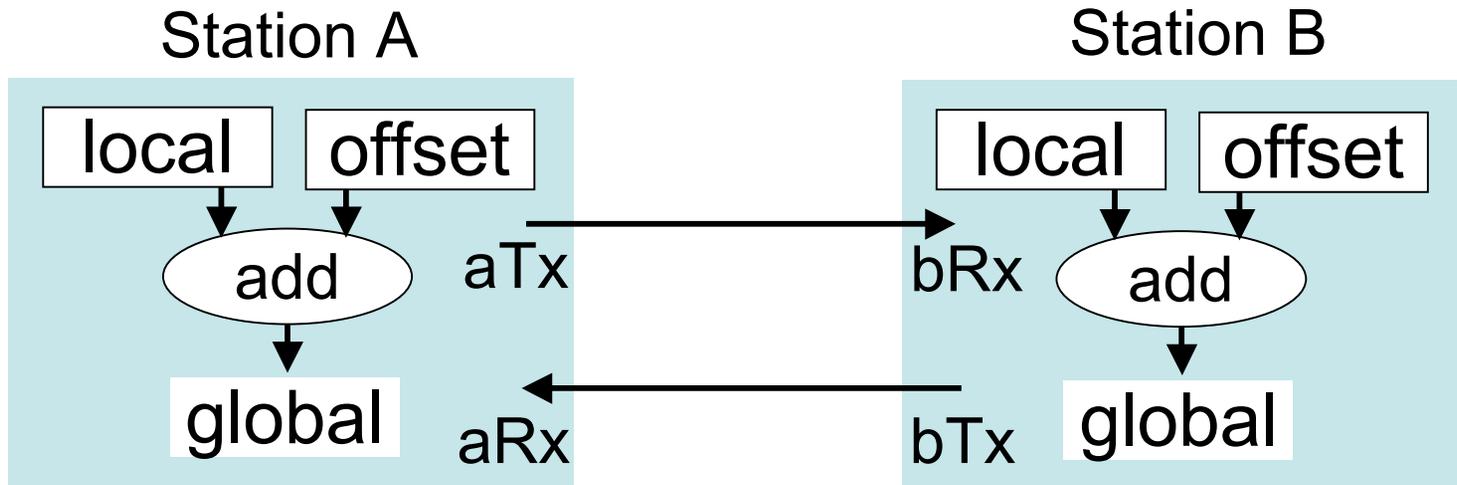
- Highest possible precision
 - Binary number (not BCD)
- Complete solution
 - 64-bit number
- Consistent with 1588, etc.
 - 32-bit seconds component
- Simple computations
 - 64-bit: seconds and fractions-of-second
- Client interface?
 - A logical interface (doesn't really matter)

Synchronized time-of-day clocks

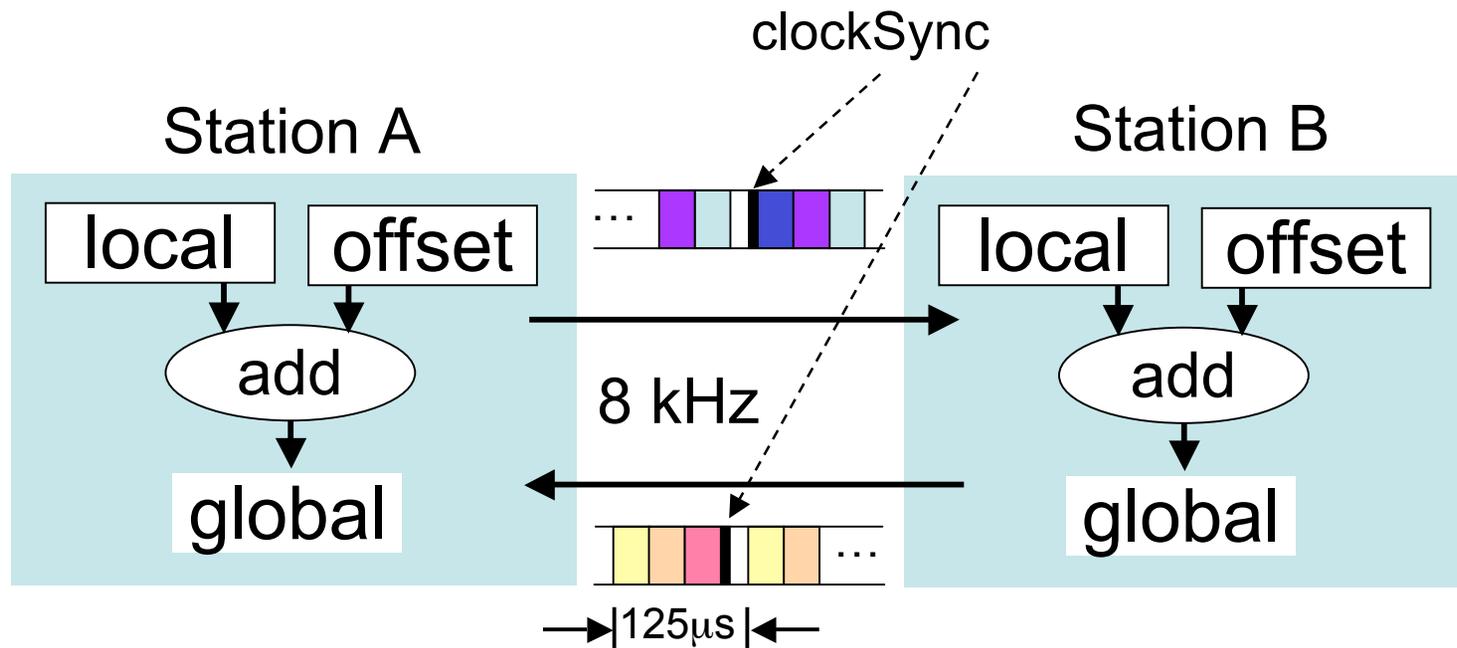
How?

Adjacent-station synchronization

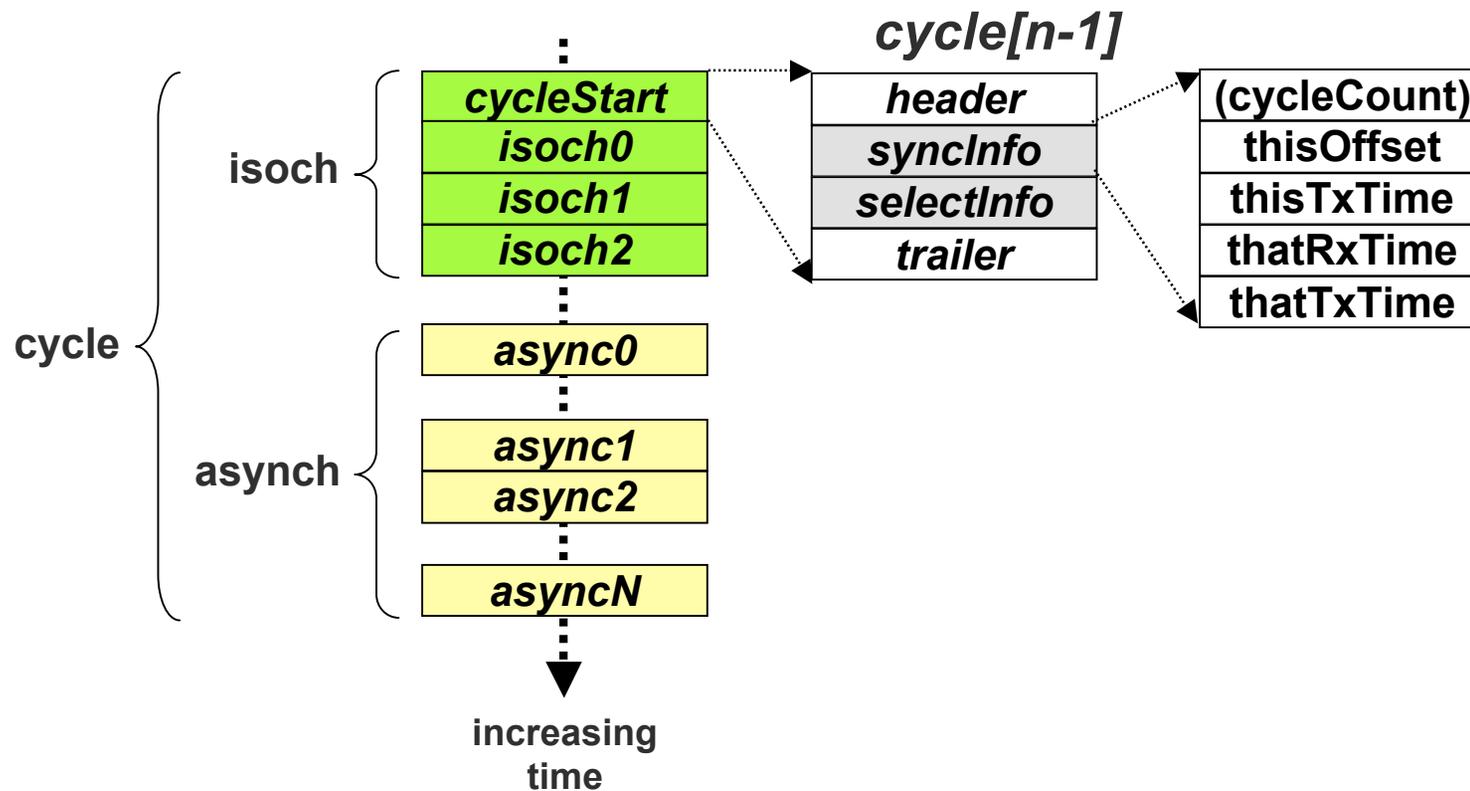
Timing snapshots



Adjacent station synchronization



clockSync frame format...



Why synchronous clock-syncs?

- Simple
 - Uses existing isochronous transmission state machine
 - Like IEEE 1588, requires only frame-sent snapshots
- Timely
 - Quasi-periodic transmissions
- Responsive
 - Sampling-to-adjustment delays are minimized
- Efficient
 - *sequenceNumber* is the *cycleCount*
- Consistent
 - Just another isochronous frame...

Conversion example: 1394

(NTP RFC-1305, SNTP RFC-2030)



Notes:

Two 32-bit additions for b:

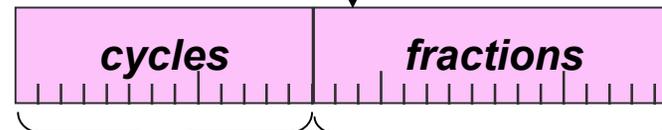
$$b = ((a \ll 7) - (a \ll 2) + a) \gg 7;$$

One 16-bit addition for d:

$$d = ((c \ll 2) + c) \gg 6;$$

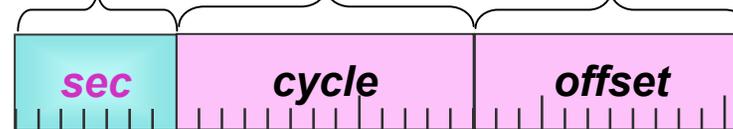
a ↓

$$b = (a * 125) \gg 7$$



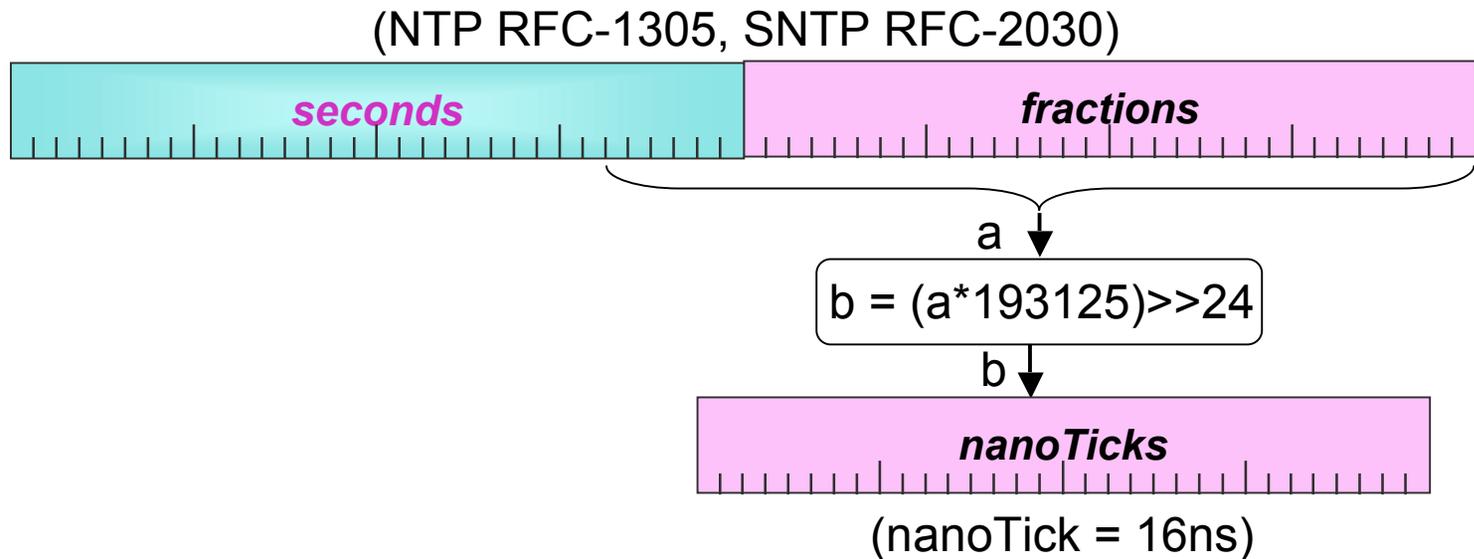
c ↓

$$d = (c * 3) \gg 6$$



(IEEE 1394 CYCLE_TIME)

Conversion example: EPON

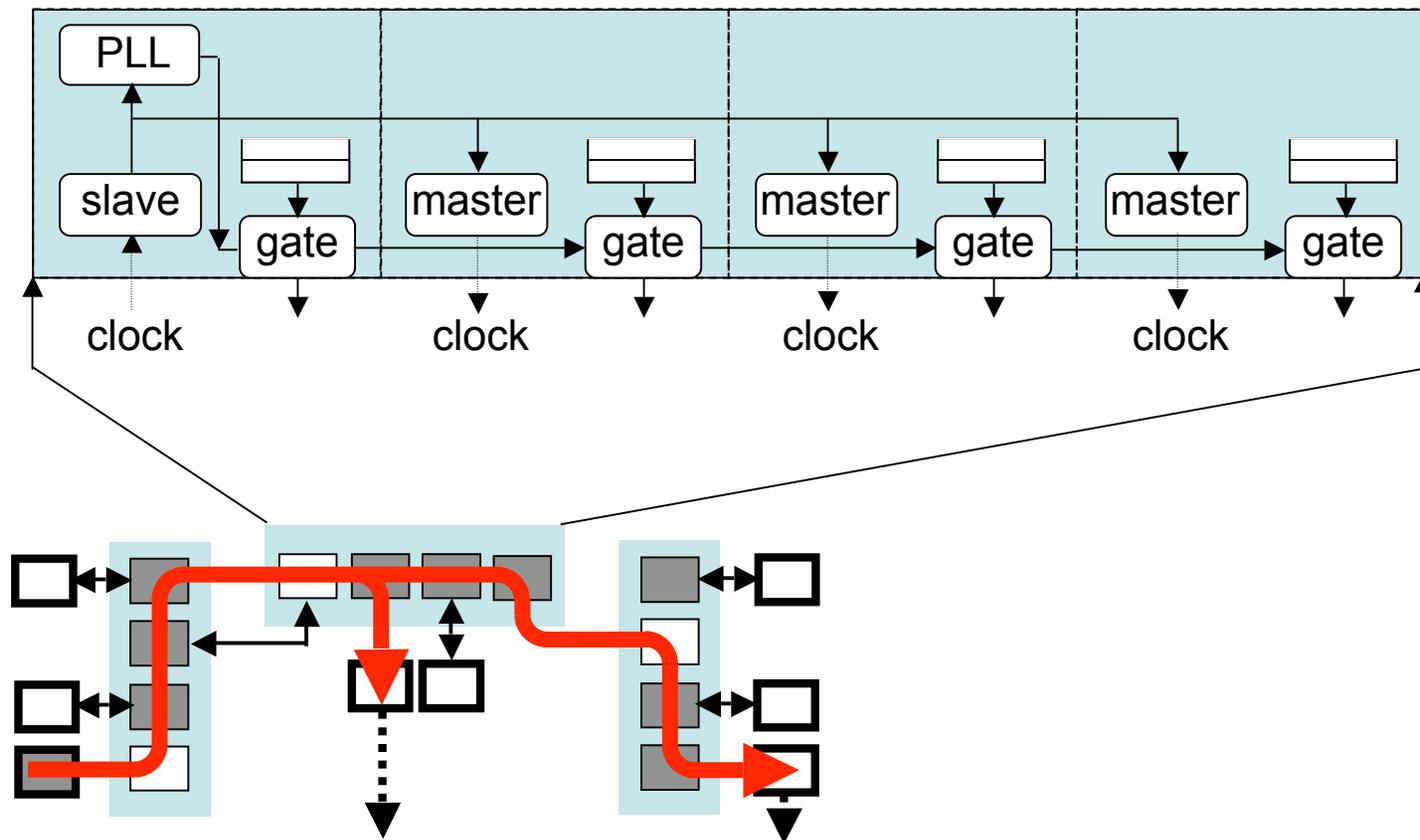


Notes:

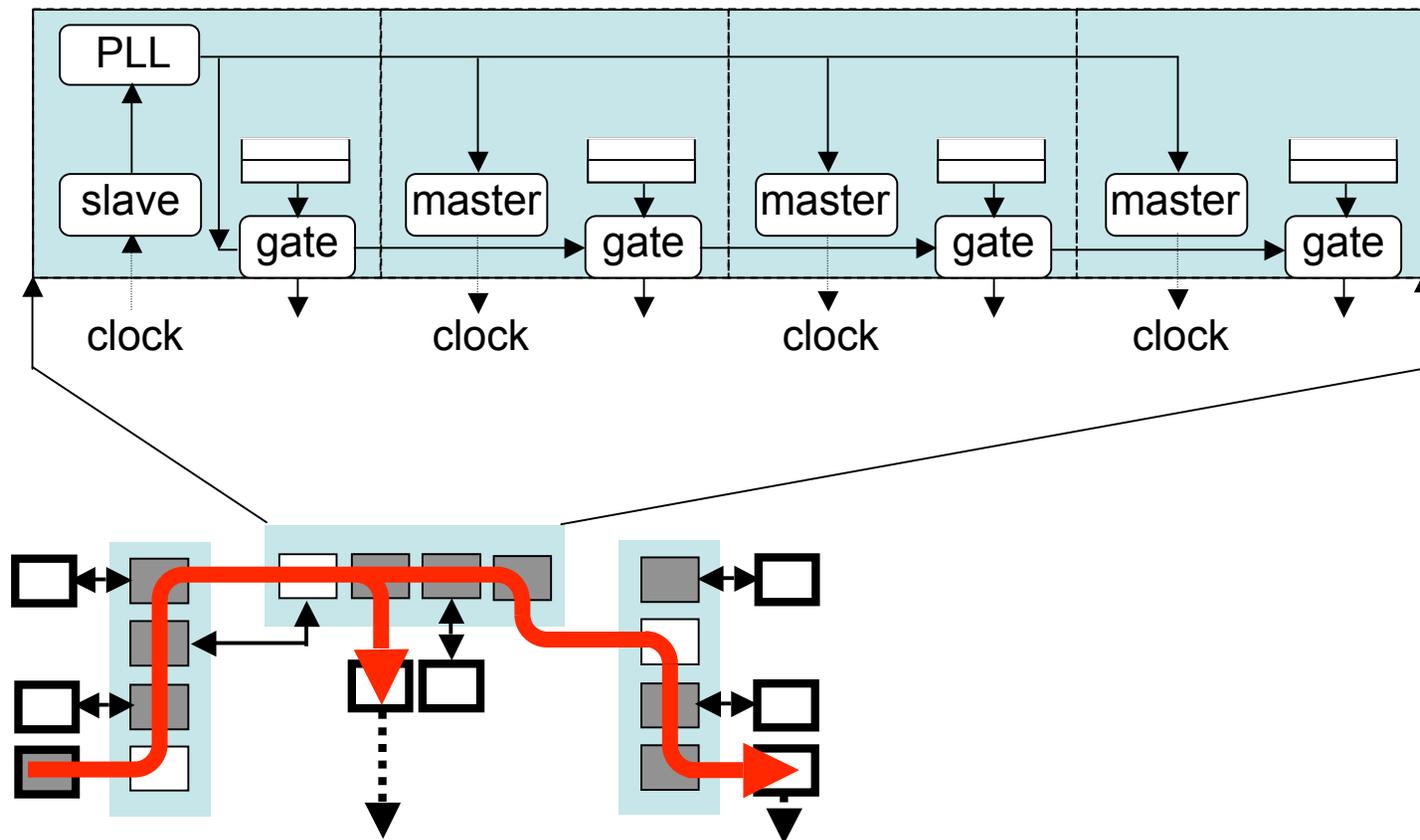
Thus, eight 36-bit additions compute a very precise b value:

$b = ((a \ll 17) + (a \ll 16) - (a \ll 12) + (a \ll 9) + (a \ll 6) + (a \ll 5) + (a \ll 2) + a) \gg 24;$

Passby PLLs (proposal 1)



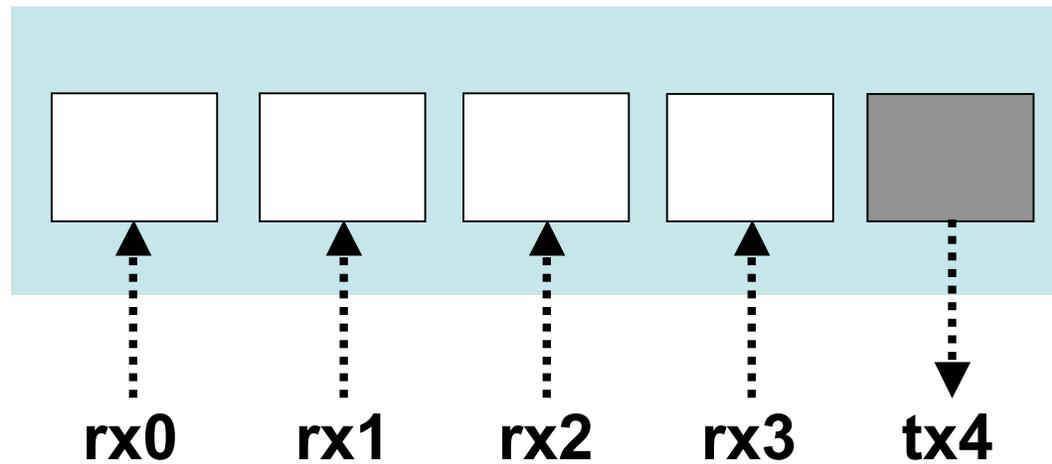
Passthrough PLLs (option 2)



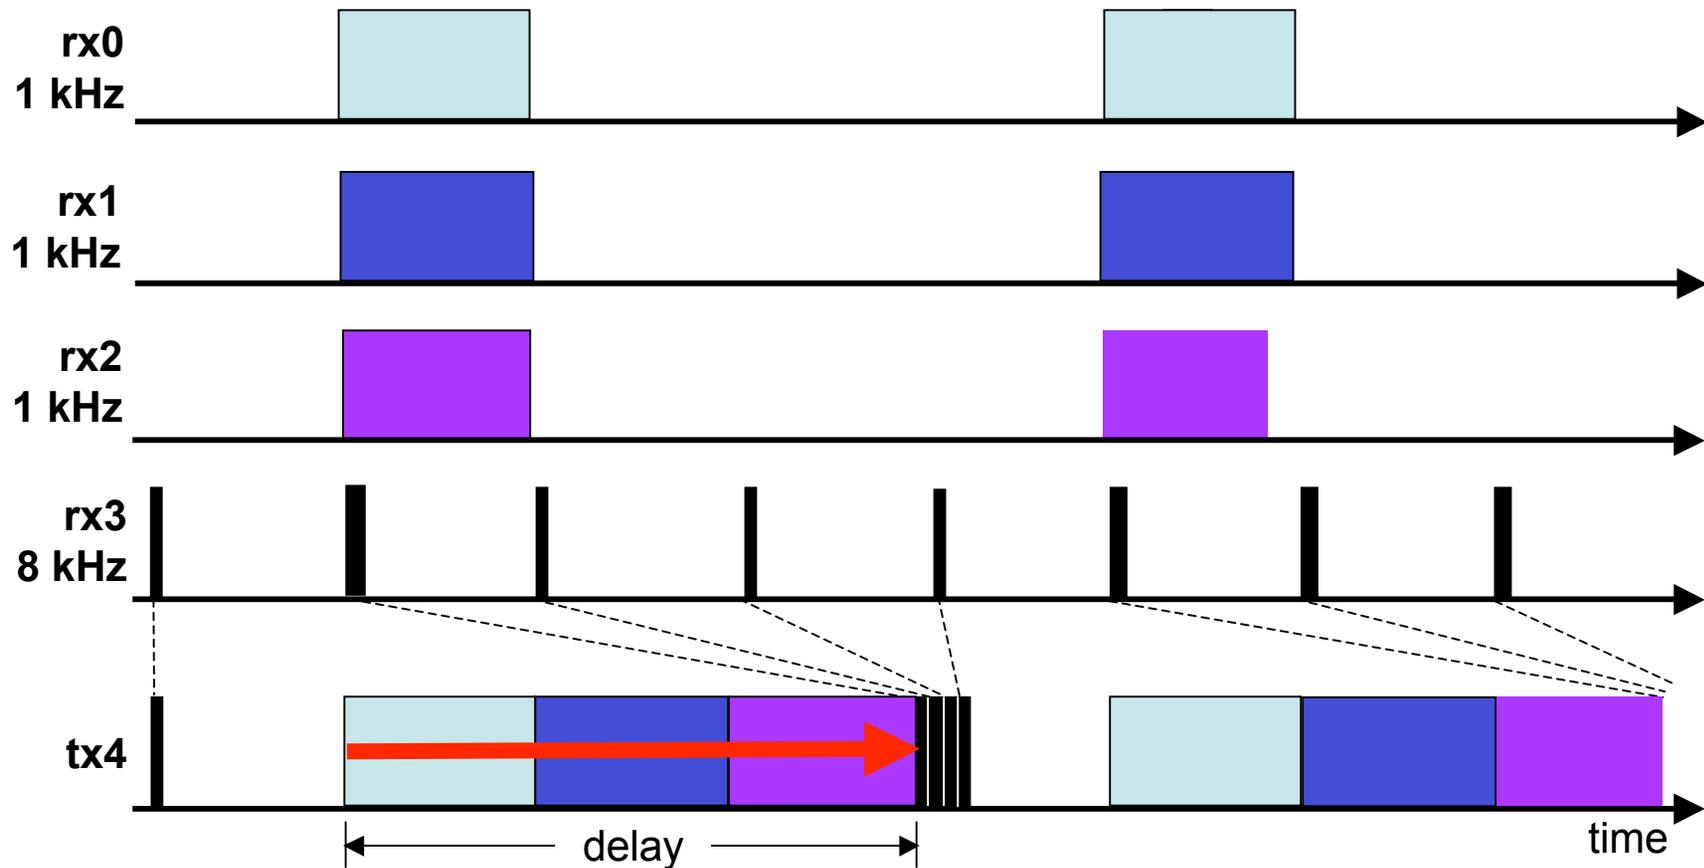
Synchronized time-of-day clocks

Why?

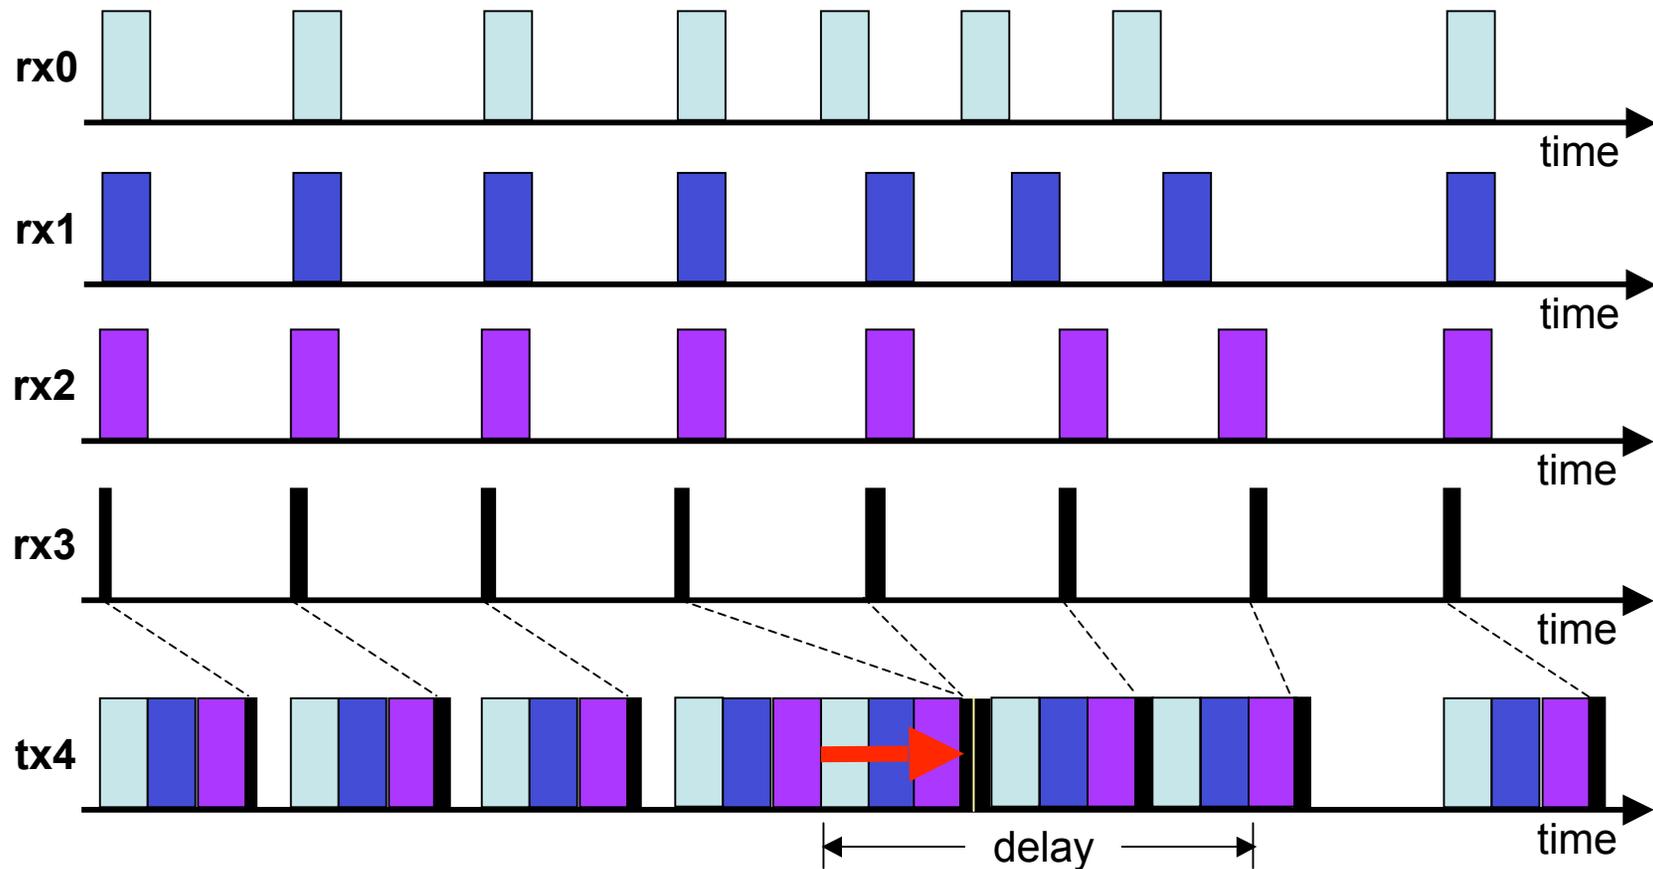
Consider possible congestion...



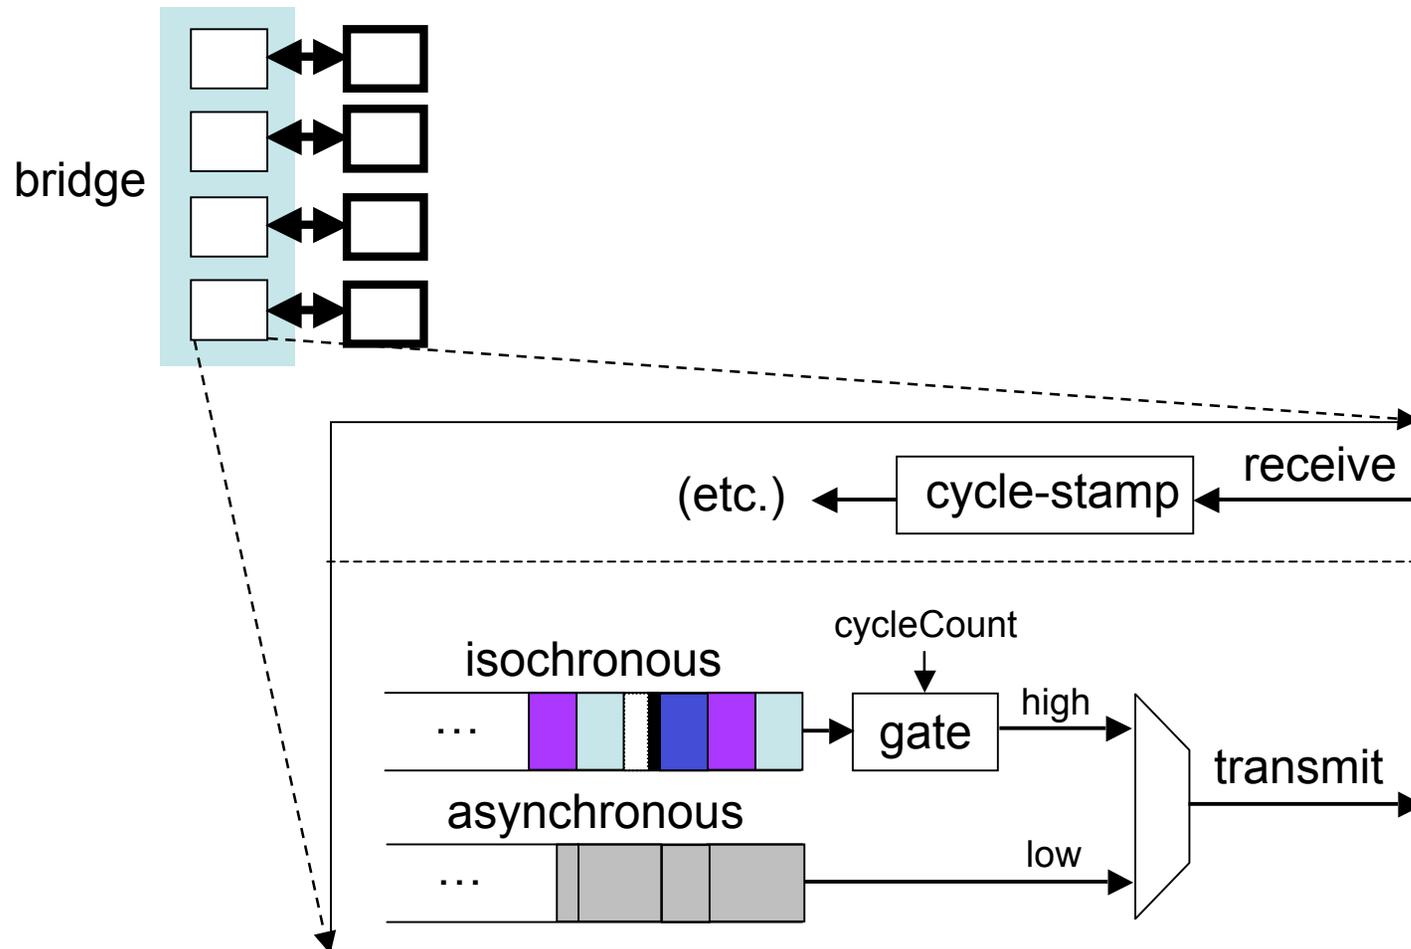
Bursting causes jitter



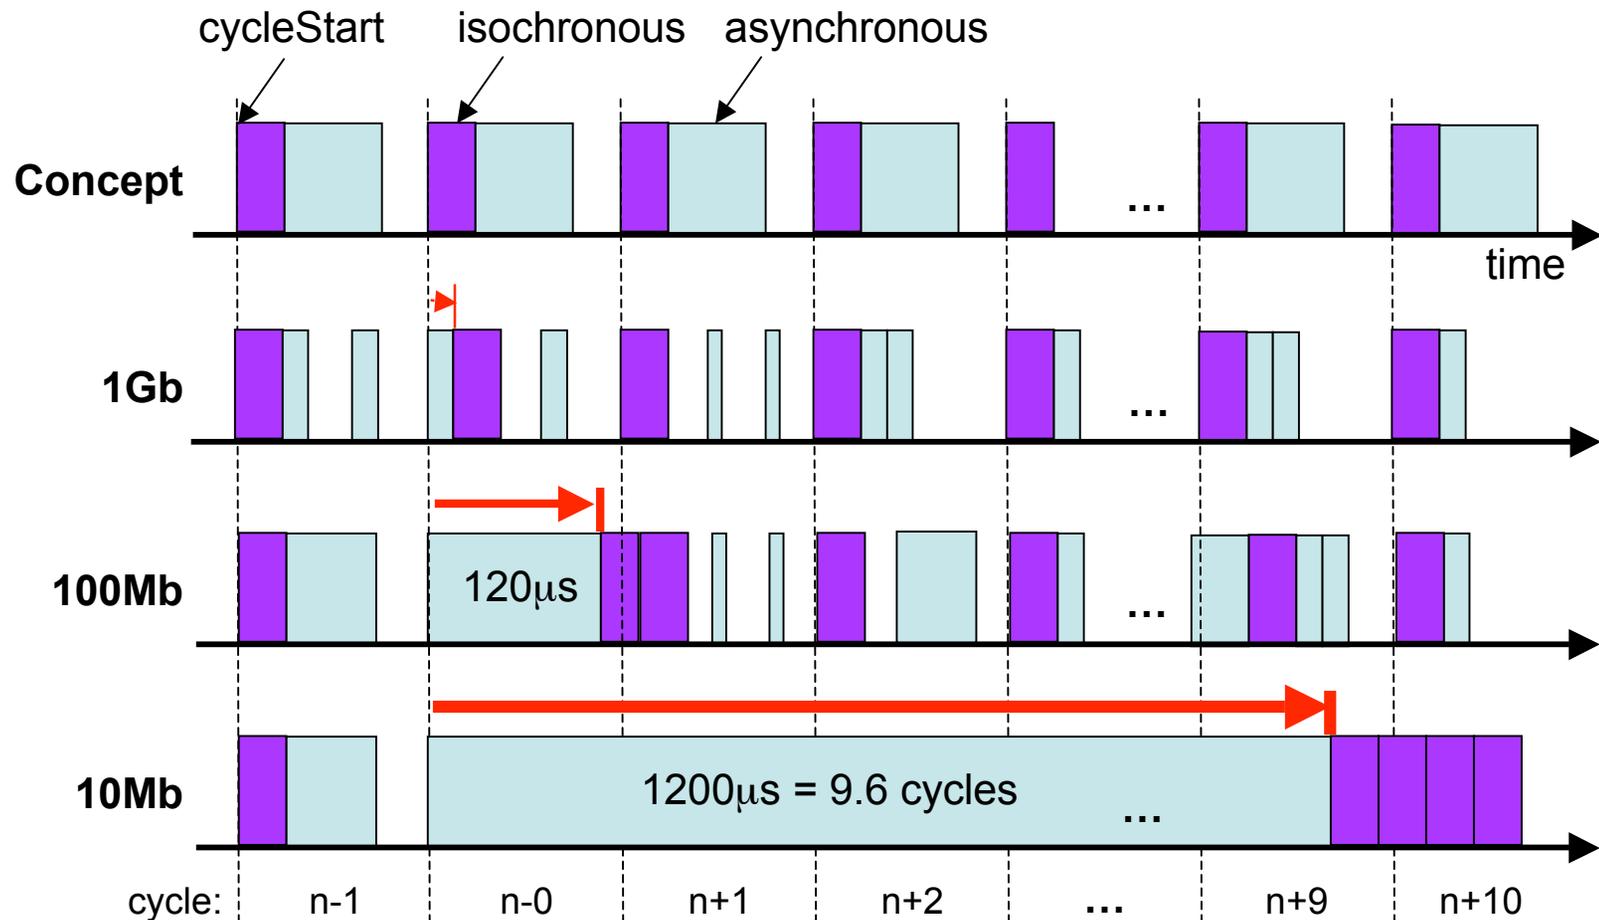
Bunching causes jitter



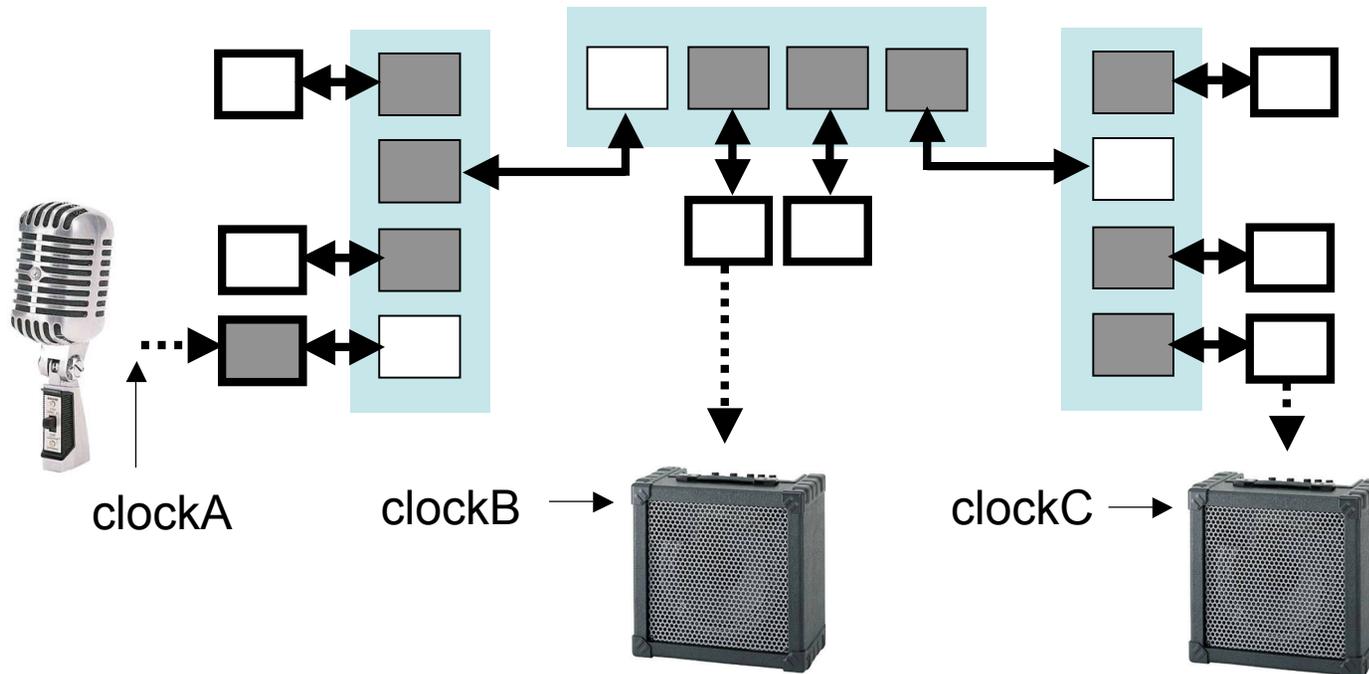
Bridge re-clocking contains jitter



Frame transmission timings



Synchronized reception/presentation



No long-term drift: clockA, clockB, clockC
Clock jitter: sub nanosecond (after PLL)

In summary

- Time-of-day synchronization (house clock)
 - Global synchronization is required
 - Implemented as cascaded adjacent synchronizations
- Time synchronization formats
 - Binary time is accurate with simple add/subtract
 - Clock-master voting: 48+ or 64+ selection priorities
- Time-of-day applications
 - Synchronous reception and presentation, within applications
 - Synchronous re-clocking within bridges
- Time-of-day distribution
 - Pipelined sampling for highest accuracies
 - Cable delays can be derived, based on the same information

Synchronized time-of-day clocks

Questions?

Synchronized time-of-day clocks

Backup slides...

Residential Ethernet

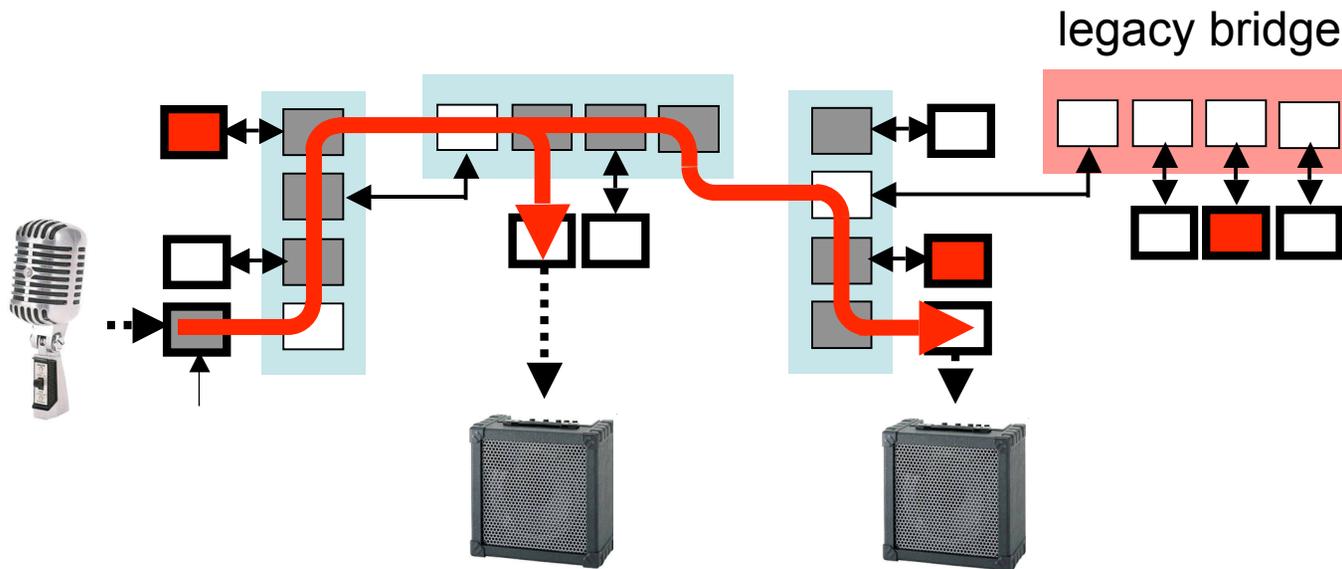
(an unofficial cumulative slide set; 2005Mar11)

Maintained by David V James

Categories of work

- Service discovery (out of scope)
 - Identify/control “talkers” and their available “plugs”
- Subscription (802.1 centric)
 - Establish conversation between talker and listener(s)
 - Reject unless: $linkBandwidth < linkCapacity$
- Clock synchronization
 - Synchronous reception, forwarding, and presentation
- Prioritized queues
 - Talkers and 100Mb bridge ports must be gated
- Formats
 - Frame formats and content (stream IDs, time stamps)
 - Time aware service interfaces

Ethernet compatibility (yes!)

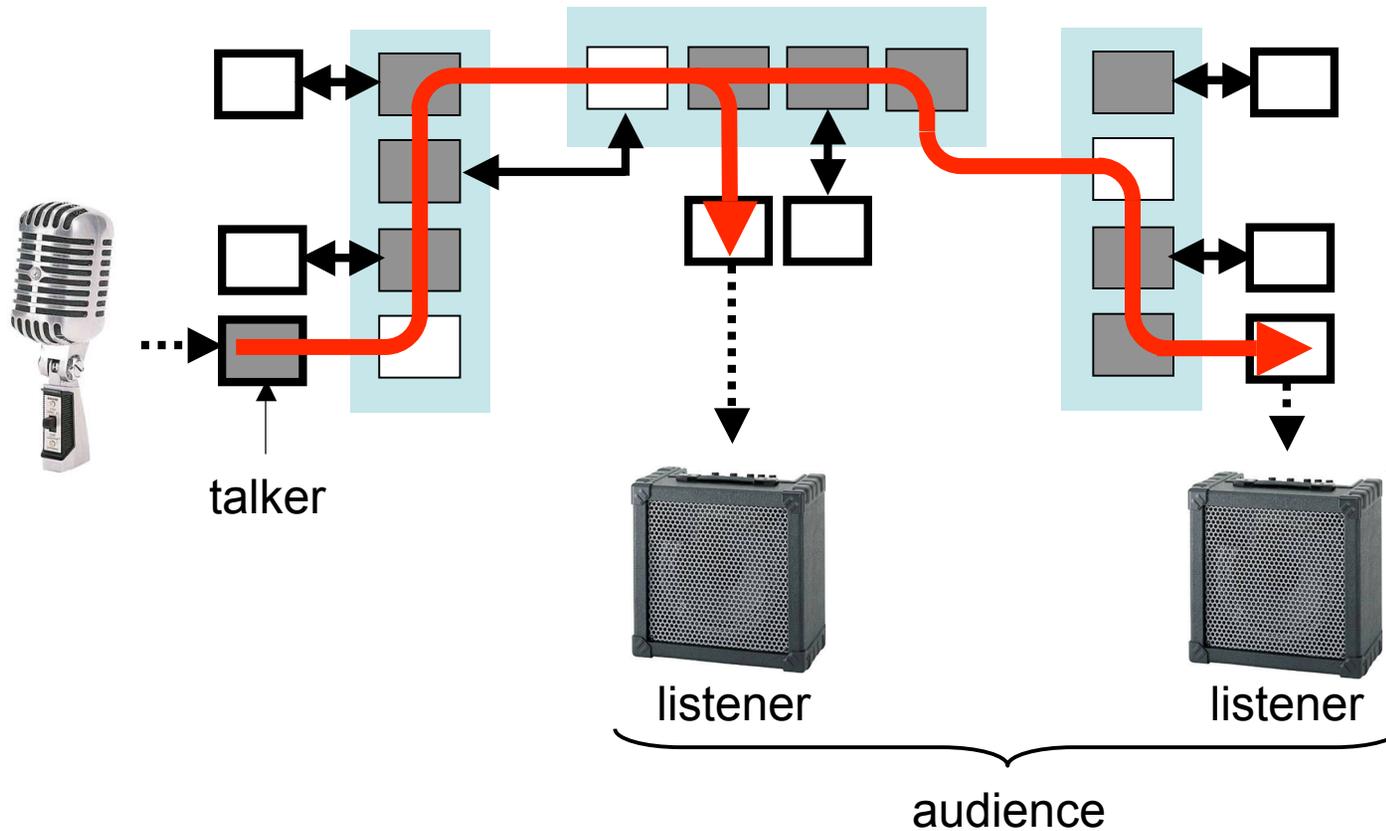


Legend:

□ □ □ □ legacy bridge

■ legacy endpoint

Vocabulary terms (1)



Vocabulary terms (2)

- audience
 - The set of listeners associated with a talker
- clock master
 - A bridge or end station that provides a link clock reference
- clock slave
 - A bridge or end station that tracks the clock reference
- grand clock master (grand master)
 - The clock master that provides the network time reference
- listener
 - A sink of a stream, such as a television or acoustic speaker
- path
 - A logical concatenation of links and bridges for a stream

Vocabulary terms (3)

- service discovery*
 - The process used to identify/control/configure talkers
- stream
 - An RE frame sequence with a distinct streamID
- subscription**
 - Establishing committed paths between talker and listener(s)
- talker
 - A source of a stream, such as a cable box or microphone

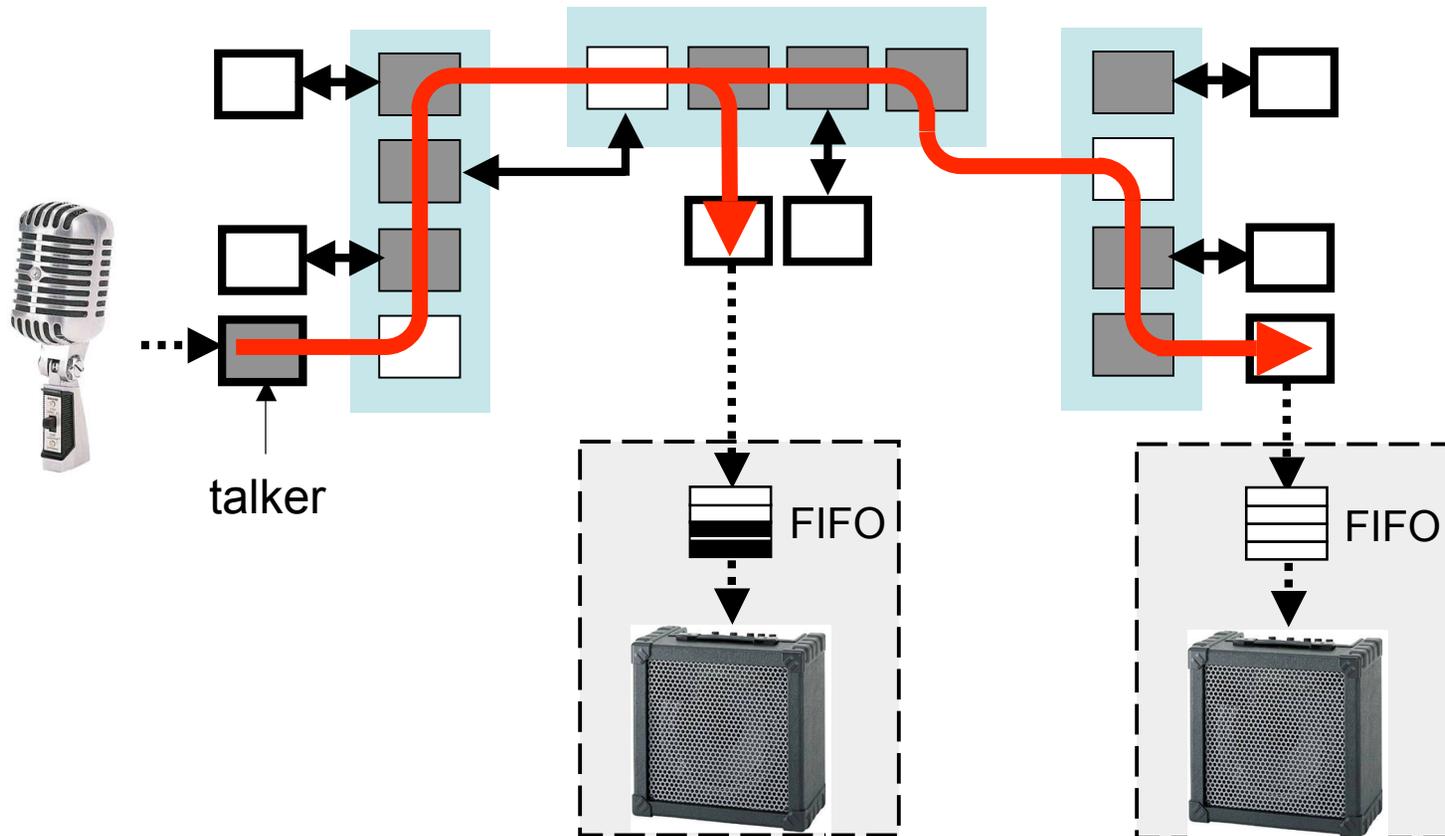
* A complete solution/vocabulary includes out-of-802 scope activities

** A complete solution/vocabulary includes out-of-802.3 scope activities

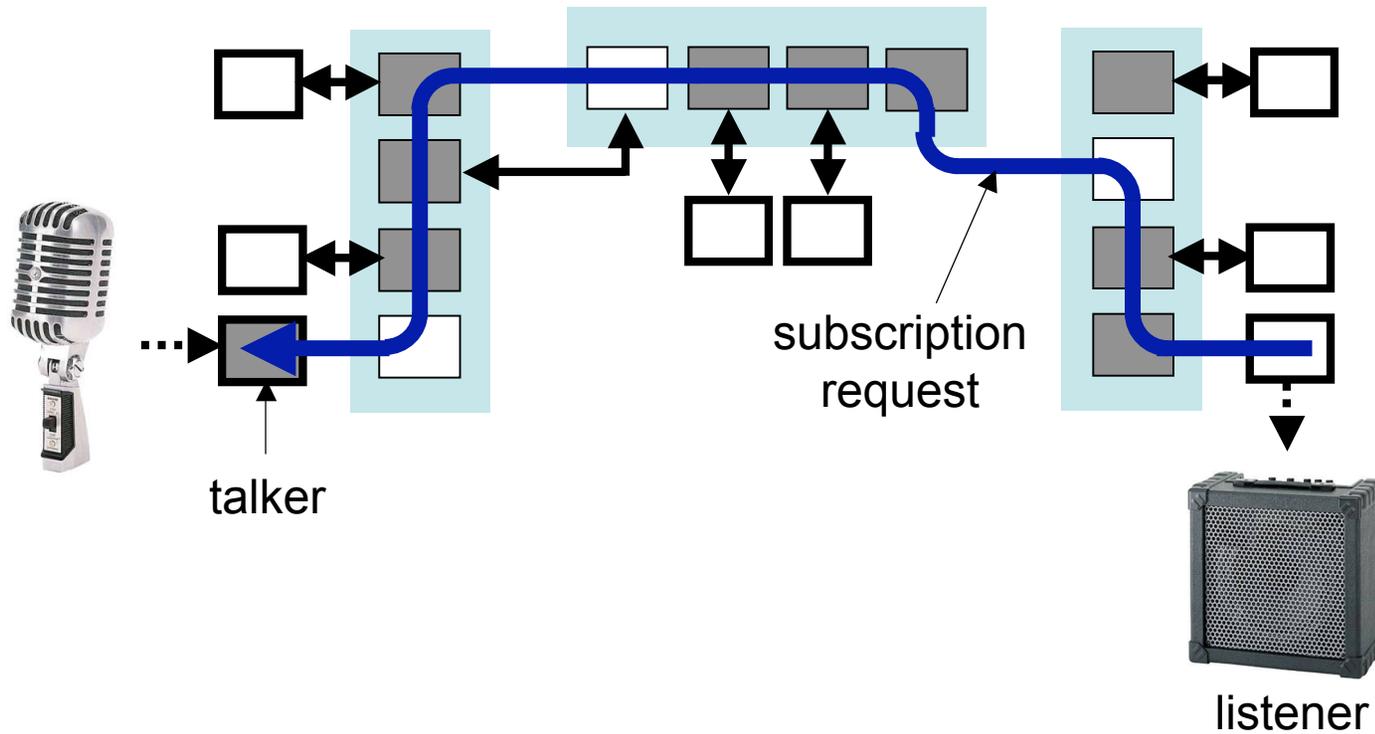
Subscription

(some possibilities)

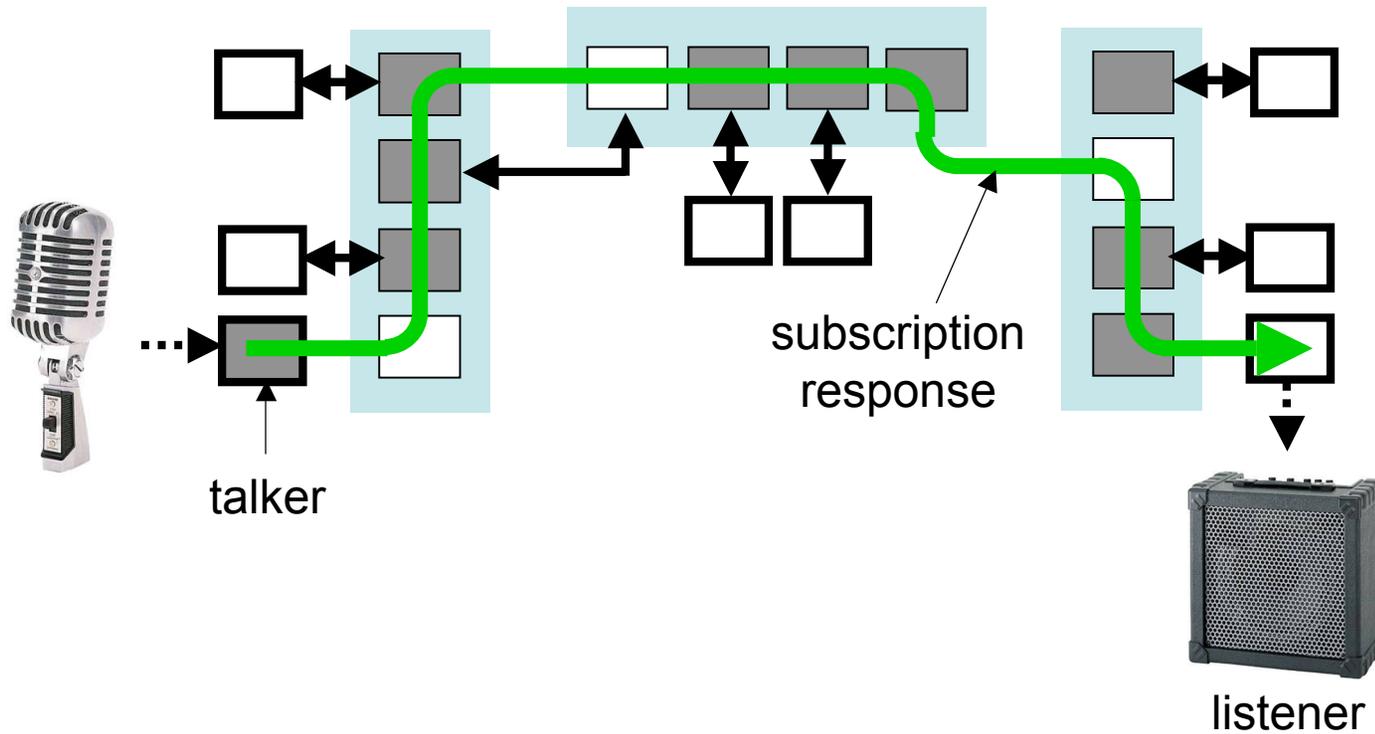
Delay-matching listener FIFOs



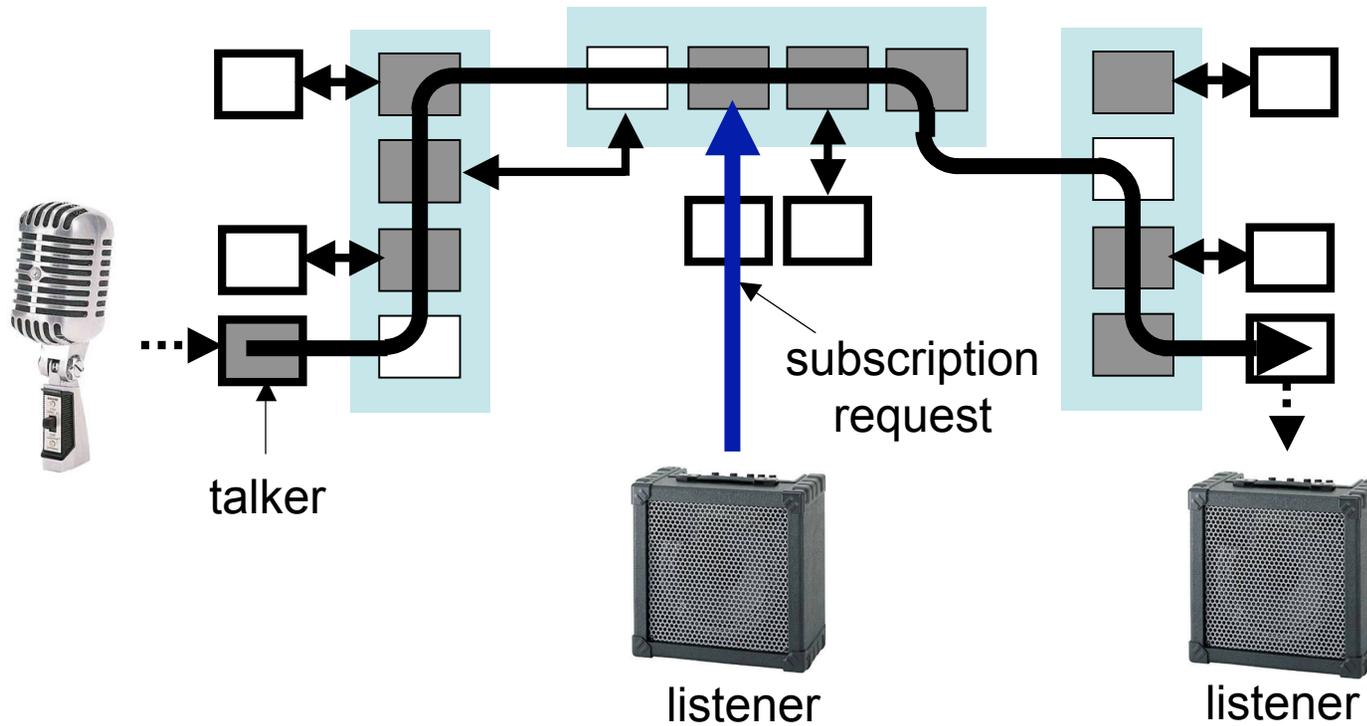
First access request



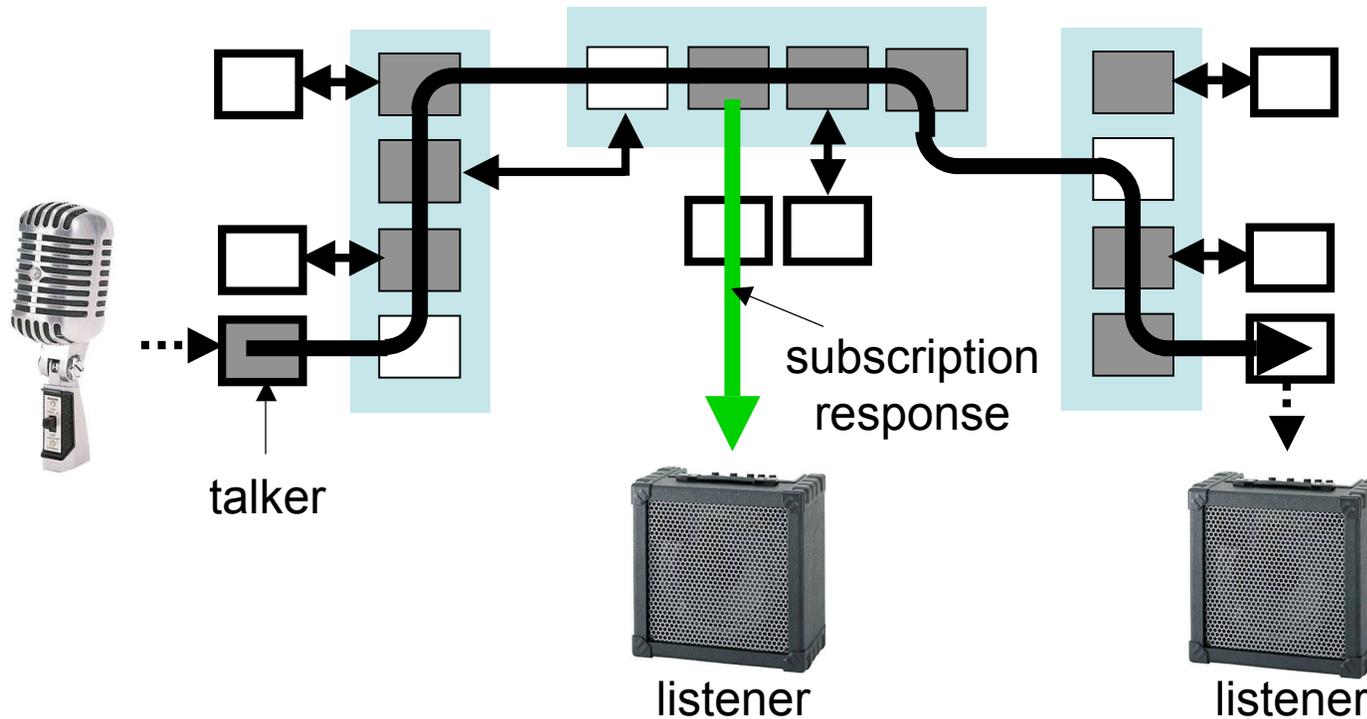
First access response



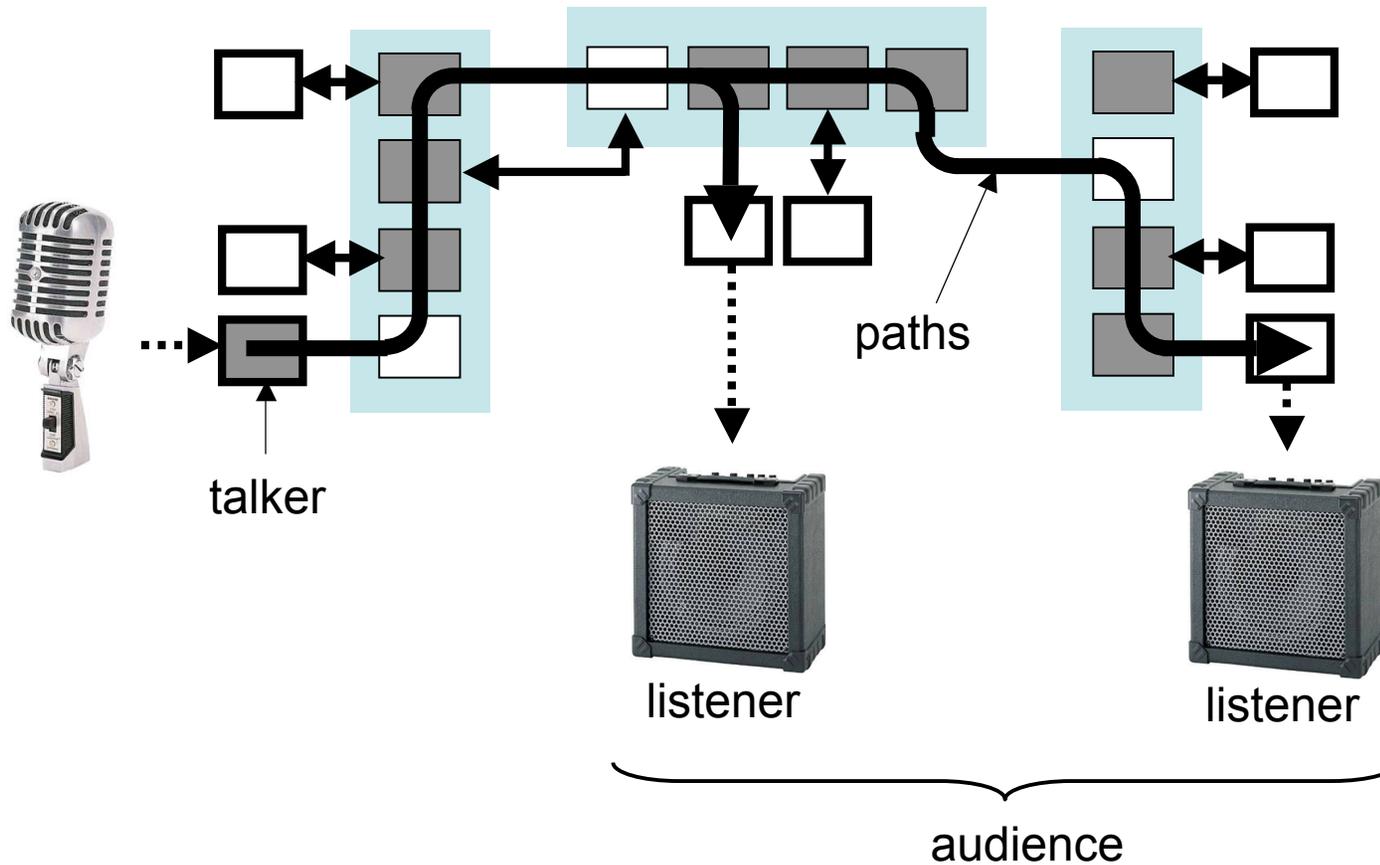
Second access request



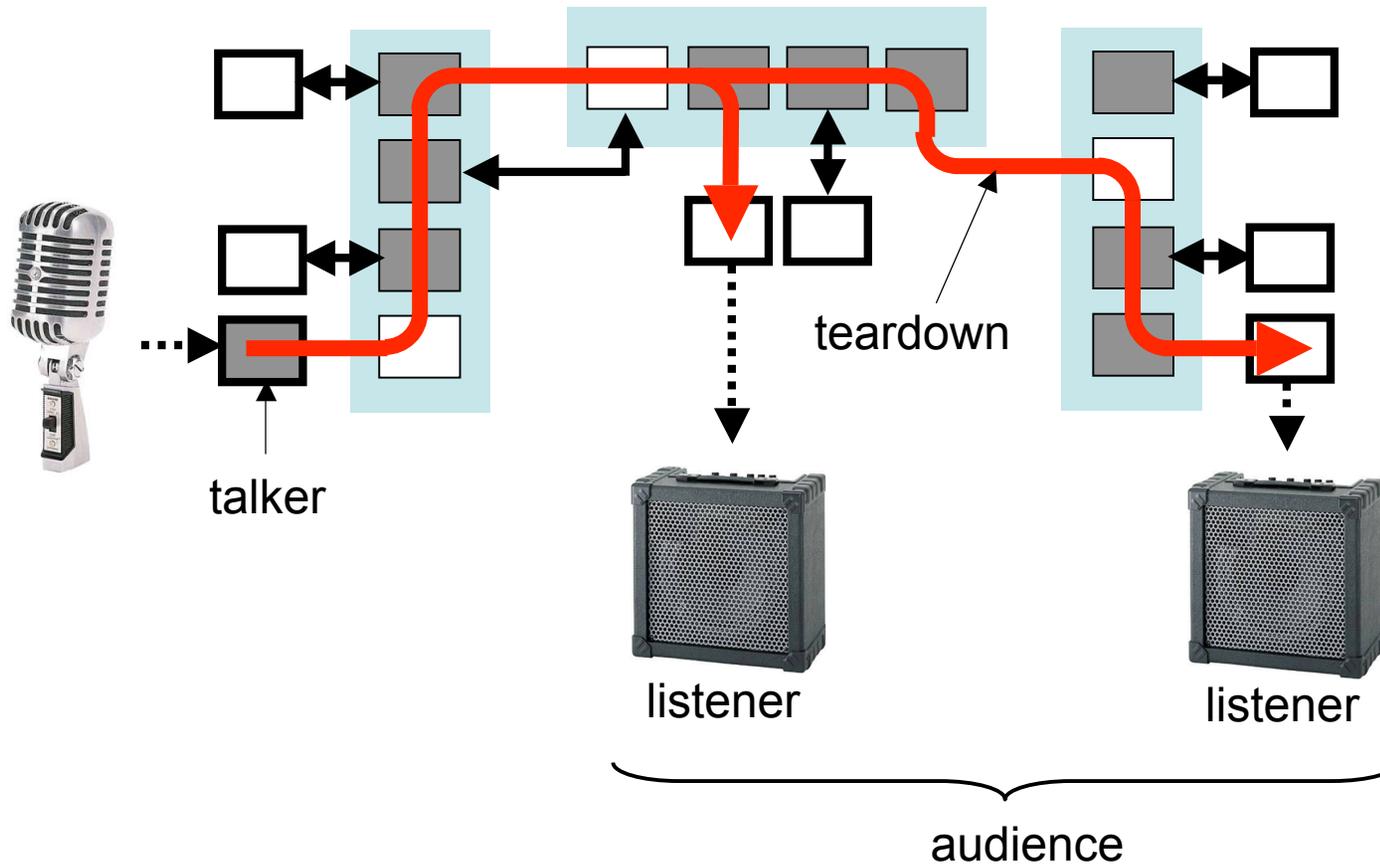
Second access response



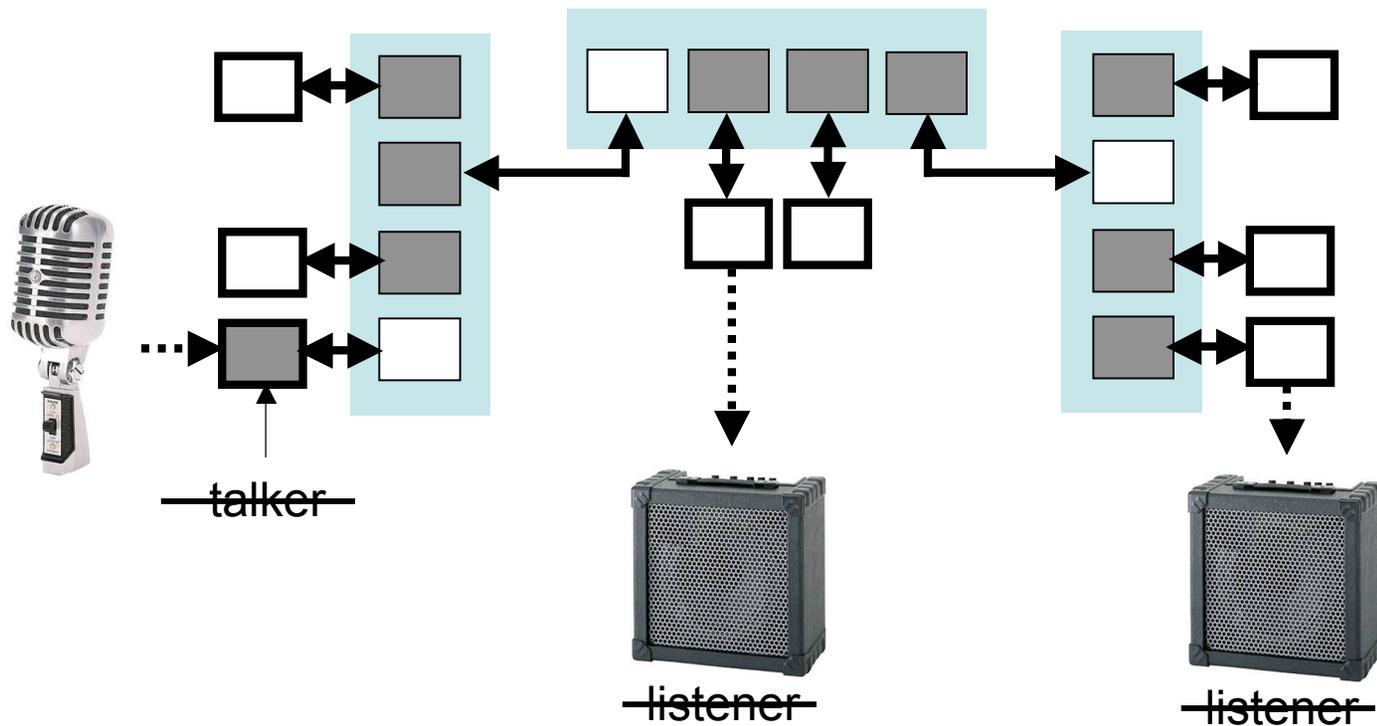
Established paths



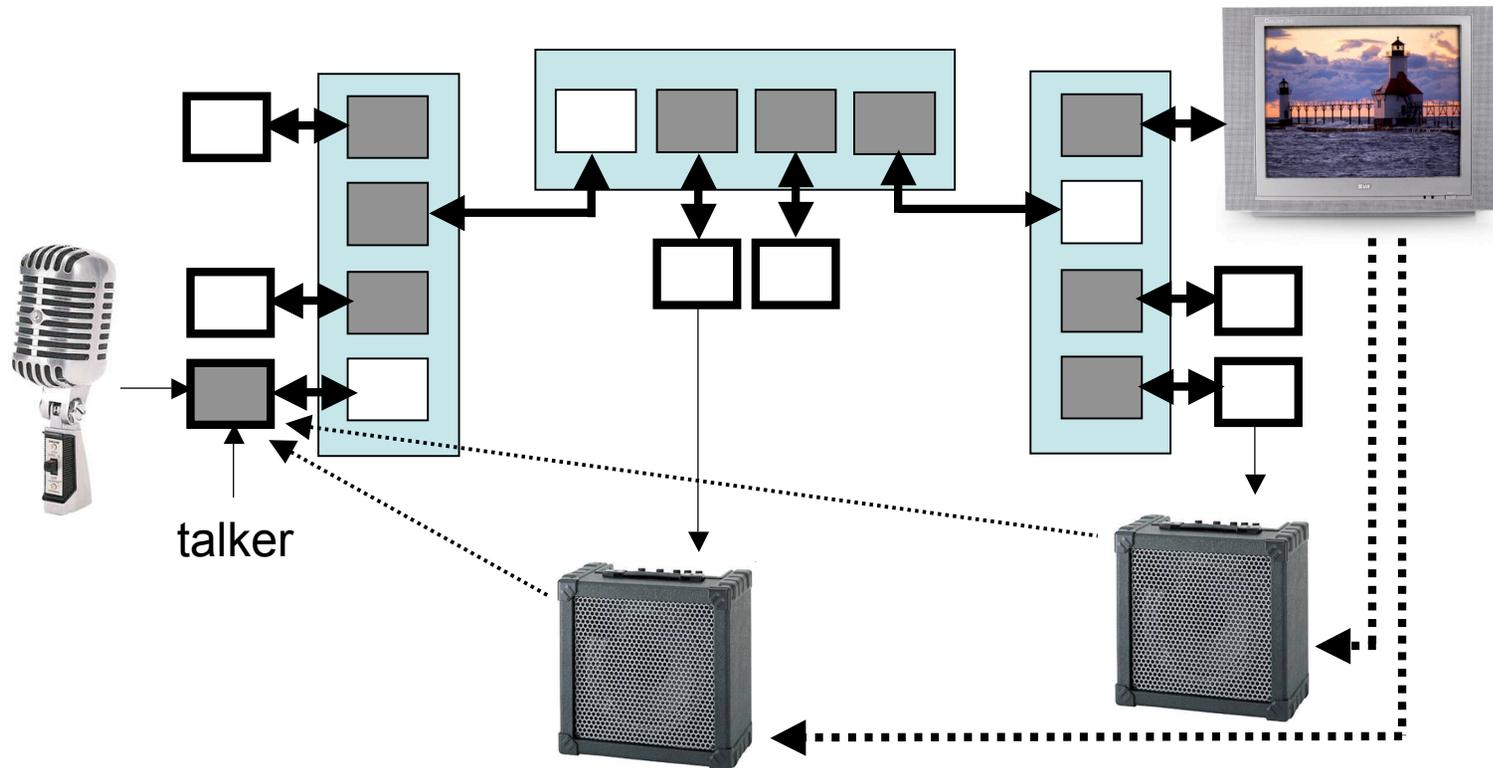
Teardown: talker → audience



Teardown completed



Third-party activations



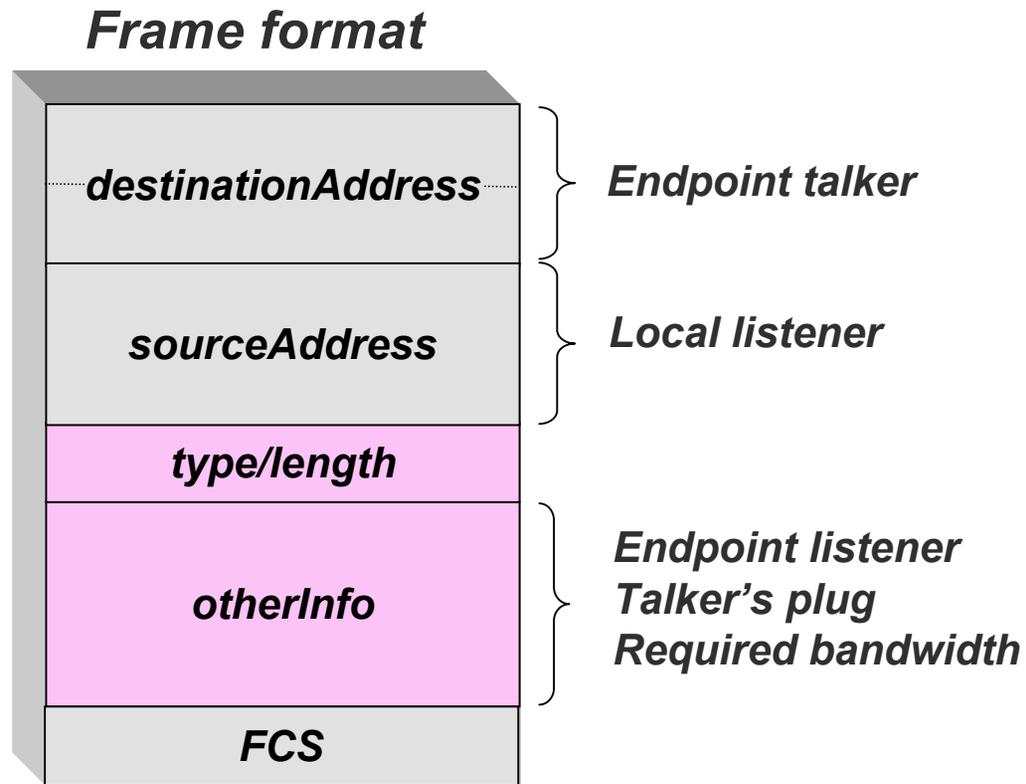
Legend:

- ◀..... Third-party activates the listener, provides talker's streamID
- ◀..... Listener subscribes to the talker's streamID

Listener-initiated heartbeats (RSVP)

- RSVP based timeout strategy
 - Periodic local listener confirmations
 - Confirmations are really just subscription requests
 - Tolerate single-frame losses
 - Tolerates configuration changes
- Talker “leaves” if no requests observed
- Listener “leaves” if no responses returns
- Talker “responses”
 - Is the flow of isochronous traffic sufficient?

Subscription requests



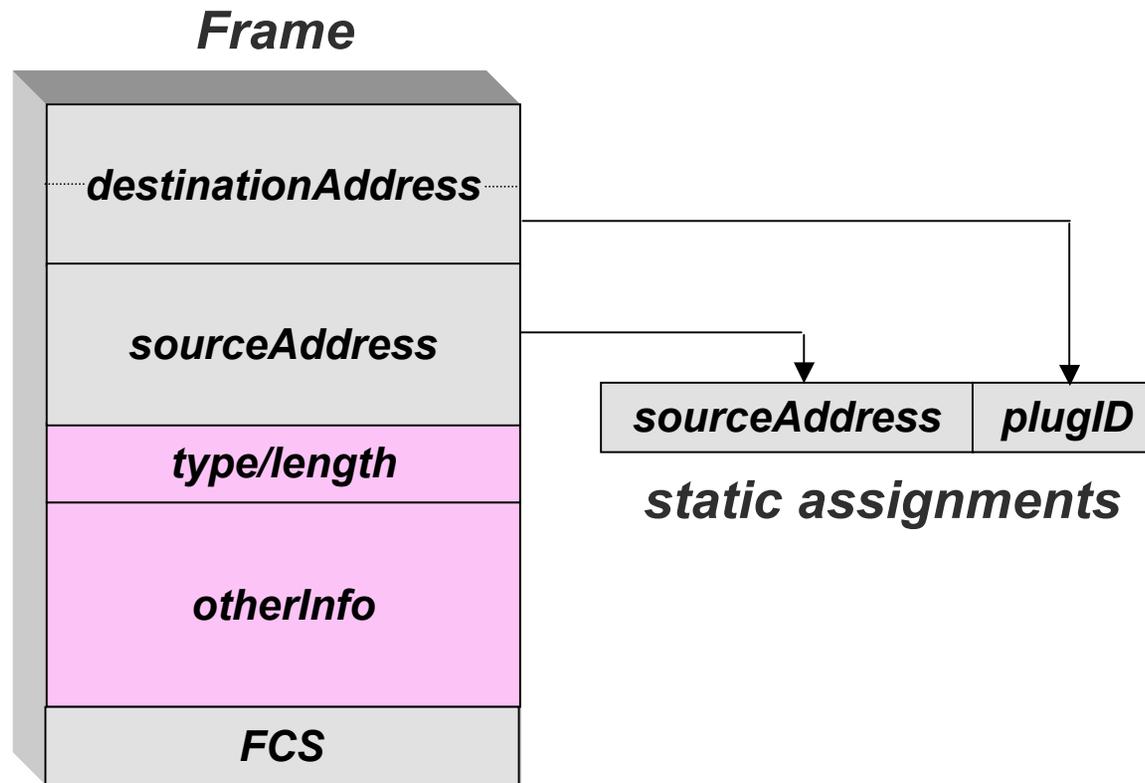
Stream addressing?

- A New Resource ReSerVation Protocol (RSVP)

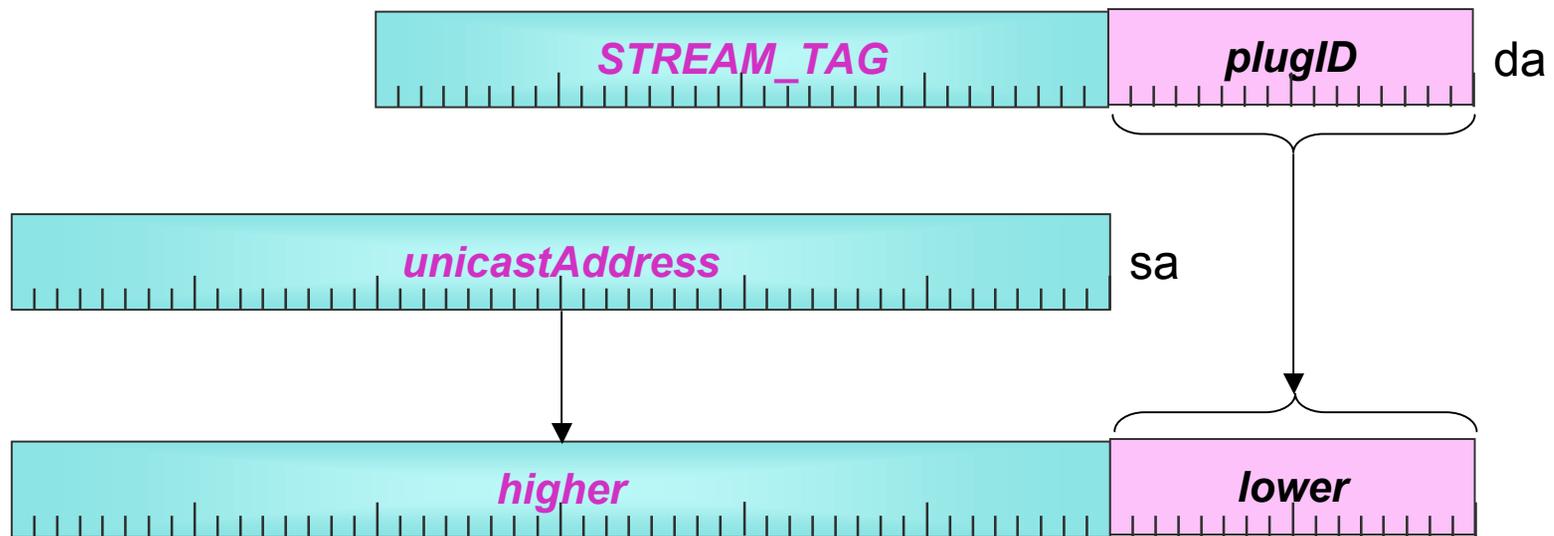
(Again, in the Internet context a data source can be specified by the source host address plus source port number. We only refer to the source host address here.)

Stream addressing?

(a possibility)



StreamID details



Are bridge changes necessary?

- End-station throttling (assuming highest priority)
 - With 1Gb/s switches, this may be sufficient
- Source device spreads transmissions evenly
 - Bunching may be tolerable within the home (needs study)
 - Must ensure that nothing else uses the highest priority
- But, some access control changes needed anyway
 - Bridges are naturally encountered along the path
 - Central topology database is thus unnecessary

Summary

- We need an RSVP-like lower level protocol
 - Restricted to 1-to-N traffic
 - Restricted to “no-filter” streams
- The IP admission control alternative
 - Out of scope
 - Layering violations (non-IP synchronous traffic?)
 - Not generally supported by residential bridges
 - Multiple components (and synchronized use) required:
 - A multicast address server?
 - IGMP/snooping for multicast setup?
 - RSVP/snooping for bandwidth negotiation?

Clock synchronization

(a Residential Ethernet SG presentation)

Leveraged protocols

- Spanning tree protocol (STP)
 - Defines the grand-master precedence format
 - But, we use a distinct value and distribution protocol (The STP root and grand master could be distinct!)
- NTP (RFC-1305) and SNTP (RFC-2030)
 - Definition of the 64-bit time-of-day value
- IEEE 1588-2002
 - Techniques for delayed-sampling synchronization

Overview

- What?
 - The clock slaves time-of-day tracks the grand master
 - No requirement for slaves to be clock-synchronous
- How?
 - Periodic exchanges of small messages
- Why?
 - Bridges: synchronized 125us cycles
 - Applications: accurate presentation times