Overview and Timing Performance of IEEE 802.1AS

Geoffrey M. Garner
Consultant

Michael D. Johas Teener
Broadcom Corporation

ISPCS
Ann Arbor'08

Outline

- Introduction
- Overview of IEEE 802.1AS
  - PTP profile
  - Synchronization
  - Best master selection
- Test configuration and hardware
- Test cases and results
  - Unfortunately, test results are not available as of the presentation of this paper. An amendment to this paper will be available in late September, 2008

Introduction – 1

- IEEE 802.1 Audio/Video Bridging (AVB) Task Group is developing four standards for transport of high-quality, time-sensitive audio/video (A/V) applications over IEEE 802 bridged local area networks
  - Precise network timing (IEEE 802.1AS)
  - Resource reservation (IEEE 802.1Qat)
  - Traffic shaping, queueing, forwarding (IEEE 802.1Qav)
  - AVB network requirements, i.e., parameters, configuration, etc. (IEEE 802.1BA)
- The current paper focuses on IEEE 802.1AS
  - Overview of the standard
  - Timing performance achieved with early implementations

Introduction – 2

- IEEE 802.1AS is based on IEEE 1588v2, and includes a PTP profile
  - Bridge acts as a boundary clock (but with peer-to-peer transparent clock formulation of synchronization)
    - Bridge participates in best master selection; this was a recent decision, driven by 3 reasons:
      - Fast reconfiguration to control phase transients when GM changes
      - Scalability (without best master selection at each bridge, larger timeout values needed for larger networks)
      - Data spanning tree determined by RSTP not necessarily optimal for synch
  - End station acts as ordinary clock

Introduction – 3

- Previously demonstrated via simulation that 802.1AS can meet the jitter, wander, and synchronization requirements for A/V applications (see [3], [4], [6], and [7] of paper)
- Subsequent test results reported at ISPCS '07 (see [8]) indicated ±500 ns synchronization could be achieved in 5 hop network with 1 Gbit/s links
- An amendment to this paper will report new test results
- As of the preparation of these slides, the latest draft of P802.1AS is D4.0 (August 26, 2008)
- Planned completion in 2009

PTP Profile Included in IEEE 802.1AS – 1

Profile Item	Specification
Best master clock algorithm (BMCA) option	Alternate BMCA (similar, but not identical, to 1588 clause 9)
Management mechanism	Still to be decided
Path delay mechanism	Peer delay mechanism
Range and default values of configurable attributes (likely 802.1AS will specify ranges; 802.1BA will specify precise values)	Precise values still to be decided. Ranges are: Sync interval: 0.01 – 1 s Announce interval: 1 – several s Pdelay interval: 0.1 – 1 s Announce receipt timeout: 3 announce intervals Sync receipt timeout: 3 sync intervals
Node types	Boundary clock (synchronization specified in manner similar to peer-to-peer transparent clock; BC and TC synchronization can be shown to be mathematically equivalent) Ordinary clock

PTP Profile Included in IEEE 802.1AS – 2

ISPCS Ann Arbor'08

Profile Item	Specification
Transport mechanism	Full-duplex IEEE 802.3 (may also model EPON as collection of full-duplex 802.3 links) Plan to have informative annex for coordinated shared network (CSN, e.g., MoCA), modeled as full-duplex 802.3 802.11 wireless is included, using facilities of 802.11v (not part of PTP profile)
Optional features	Bridges/end-station required to measure frequency offset to nearest neighbor (but not required to adjust frequency); frequency offset is accumulated and used to correct propagation time and compute synchronized time Standard organization TLV is defined Optional features of 1588 clauses 16 and 17 not used Annex K security protocol not used Annex L cumulative frequency scale factor not used (but cumulative frequency offset is accumulated)

Additional Network Assumptions – 1

- All bridges/end stations are “time-aware”, i.e., meet the requirements of 802.1AS
  - No ordinary bridges
  - Peer-delay mechanism used to detect non-802.1AS bridges
  - Except for peer delay, the 802.1AS protocol will not run on ports where a non-802.1AS bridge is detected
- Oscillator frequency of at least 25 MHz (40 ns granularity)
- ±100 ppm frequency accuracy
- Ethernet links are 100 Mbit/s or 1 Gbit/s

Additional Network Assumptions – 2

- 802.11 links are 100 Mbit/s (i.e., meet requirements of IEEE 802.11n)
- All time-aware systems are 2-step clocks
  - Always send Follow_Up and Pdelay_Resp_Follow_Up
- Bridges adjust time and frequency instantaneously, i.e., they do not do any PLL filtering
  - All filtering is done at end stations; this allows cost of filtering to be borne by applications
- 802.1AS network is single PTP domain (domain number 0)
- PTP timescale is used

802.1AS Architecture and Entities

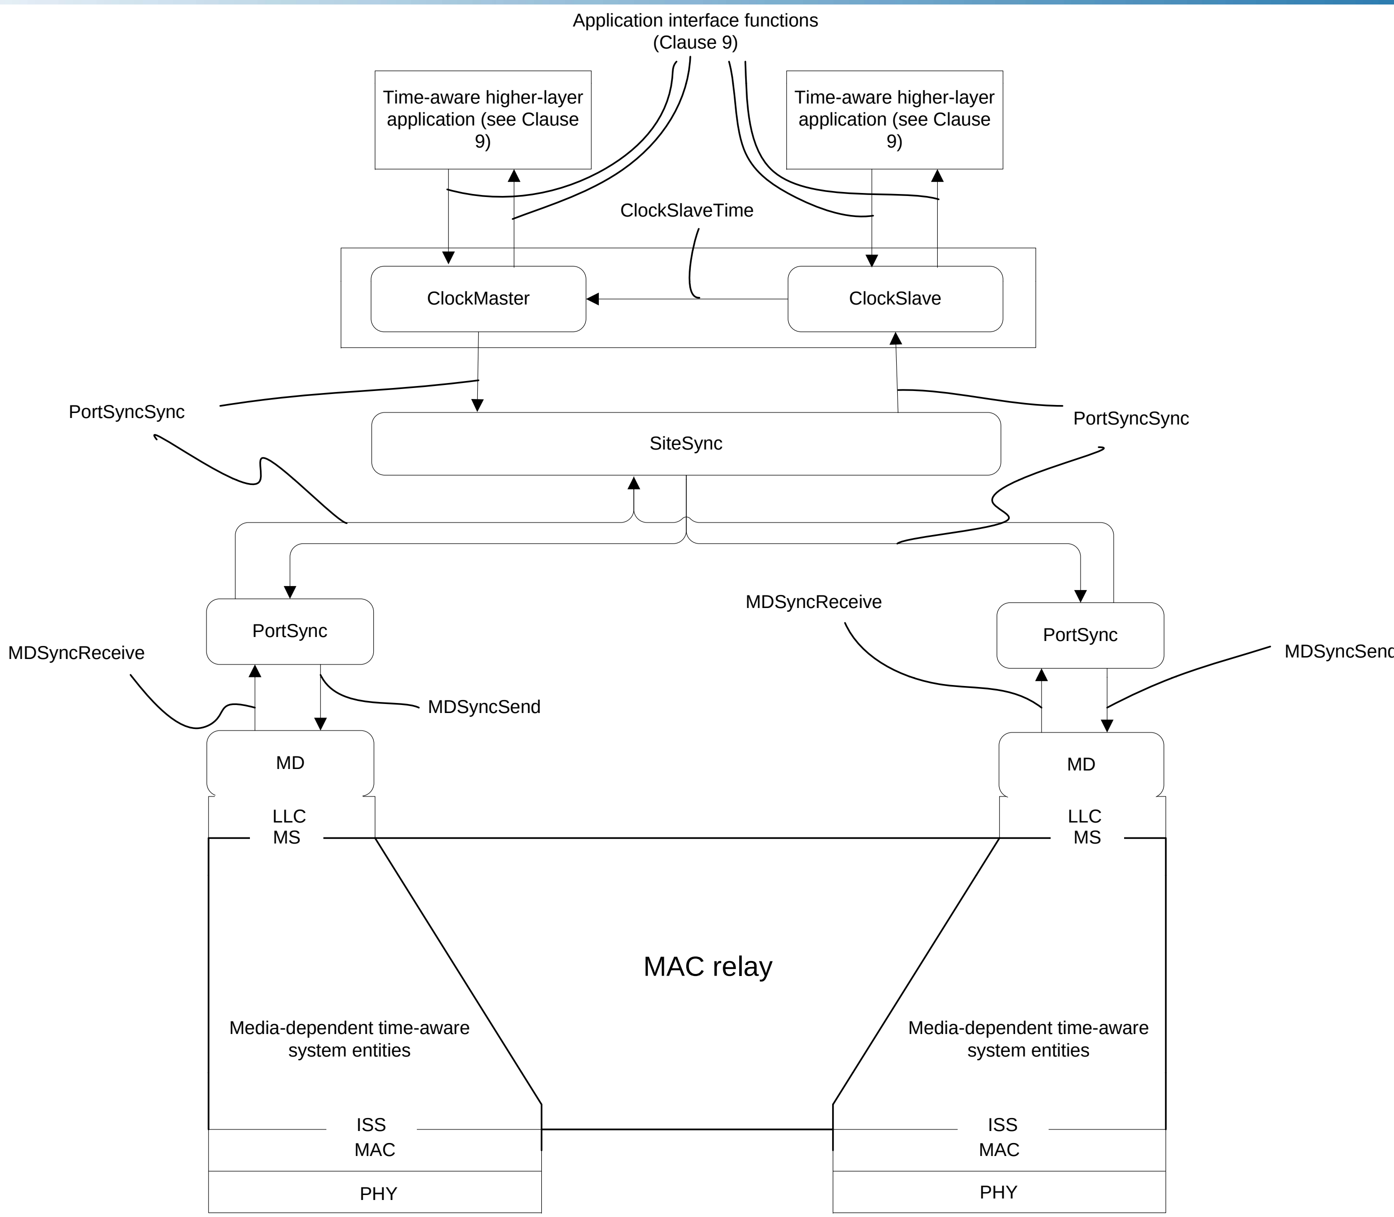

Synchronization in IEEE 802.1AS – 1

- Every IEEE 802.3 port of a time-aware system runs peer delay mechanism
  - Measure propagation delay as specified in 11.4 of IEEE 1588
    - Responder provides requestReceiptTimestamp and responseOriginTimestamp separately
  - Requestor uses successive responseOriginTimestamp values to measure frequency offset of responder relative to requestor
  - Frequency offset is used to correct propagation delay measurement (frequency offset multiplied by turnaround time)

Synchronization in IEEE 802.1AS – 2

- Frequency offset is accumulated in standard organization TLV (1588 clause 14)
  - TLV is attached to Follow_Up
  - Frequency offset is initialized to zero at grandmaster
  - Accumulation allows each time-aware system to know its frequency offset relative to grandmaster
- The advantage of accumulating the frequency offset relative to the grandmaster, rather than measuring it directly using Sync and Follow_Up, is that it can be determined on receipt of first Follow_Up after a change of grandmaster
  - This is because the nearest-neighbor offsets are measured constantly, on all links

Synchronization in IEEE 802.1AS – 3

- Each time-aware system sends Sync and Follow_Up on its master ports
- Normally, send Sync and Follow_Up as soon as possible after receiving Sync and Follow_Up on slave port
  - However, don't send until at least one-half sync interval has elapsed since last sync was sent, to prevent bunching of successive messages
  - Also, send Sync and Follow_Up after a sync interval has elapsed since sending of last Sync, even if Sync and Follow_Up have not been received

Synchronization in IEEE 802.1AS – 4

ISPCS Ann Arbor'08

- preciseOriginTimestamp and correctionField in transmitted Follow_Up
  - PreciseOriginTimestamp is copied from most recently received valid Follow_Up
  - Let T_r = syncEventIngress timestamp for most recently received valid Sync
  - Let T_s = syncEventEgress timestamp for Sync just transmitted (corresponding to Follow_Up)
  - Let C_r = correction field of most recently received valid Follow_Up
  - Let C_s = correction field of Follow_Up being transmitted

Synchronization in IEEE 802.1AS – 5

- Let y_a = frequency offset of current time-aware system relative to grandmaster (accumulated in Follow_Up TLV)
- Then
  - $C_s = C_r + (T_s – T_r)*(1 + y_a)$
- This result is very similar to what is done for a transparent clock, with $(T_s – T_r)*(1 + y_a)$ = residence time corrected for frequency offset
- The sum of C_s and the preciseOriginTimestamp is equal to the synchronized time at the instant the Sync message is timestamped
  - A BC would put this sum (except for any sub-ns portion) in the preciseOriginTimestamp

Synchronization in IEEE 802.1AS – 6

- Then, the only difference in the manner in which a BC and TC transport synchronized time is in how the time value is distributed between the timestamp and correction fields
  - The two are mathematically equivalent
- A BC and TC do differ in other ways, e.g.,
  - A BC invokes best master selection; a TC does not
  - A TC ordinarily sends Sync as soon after receiving Sync as possible; a BC does not
    - But note that it has been shown that phase accumulation in a chain of BCs is reduced if this time is made smaller and, ideally, zero

Synchronization State Machines

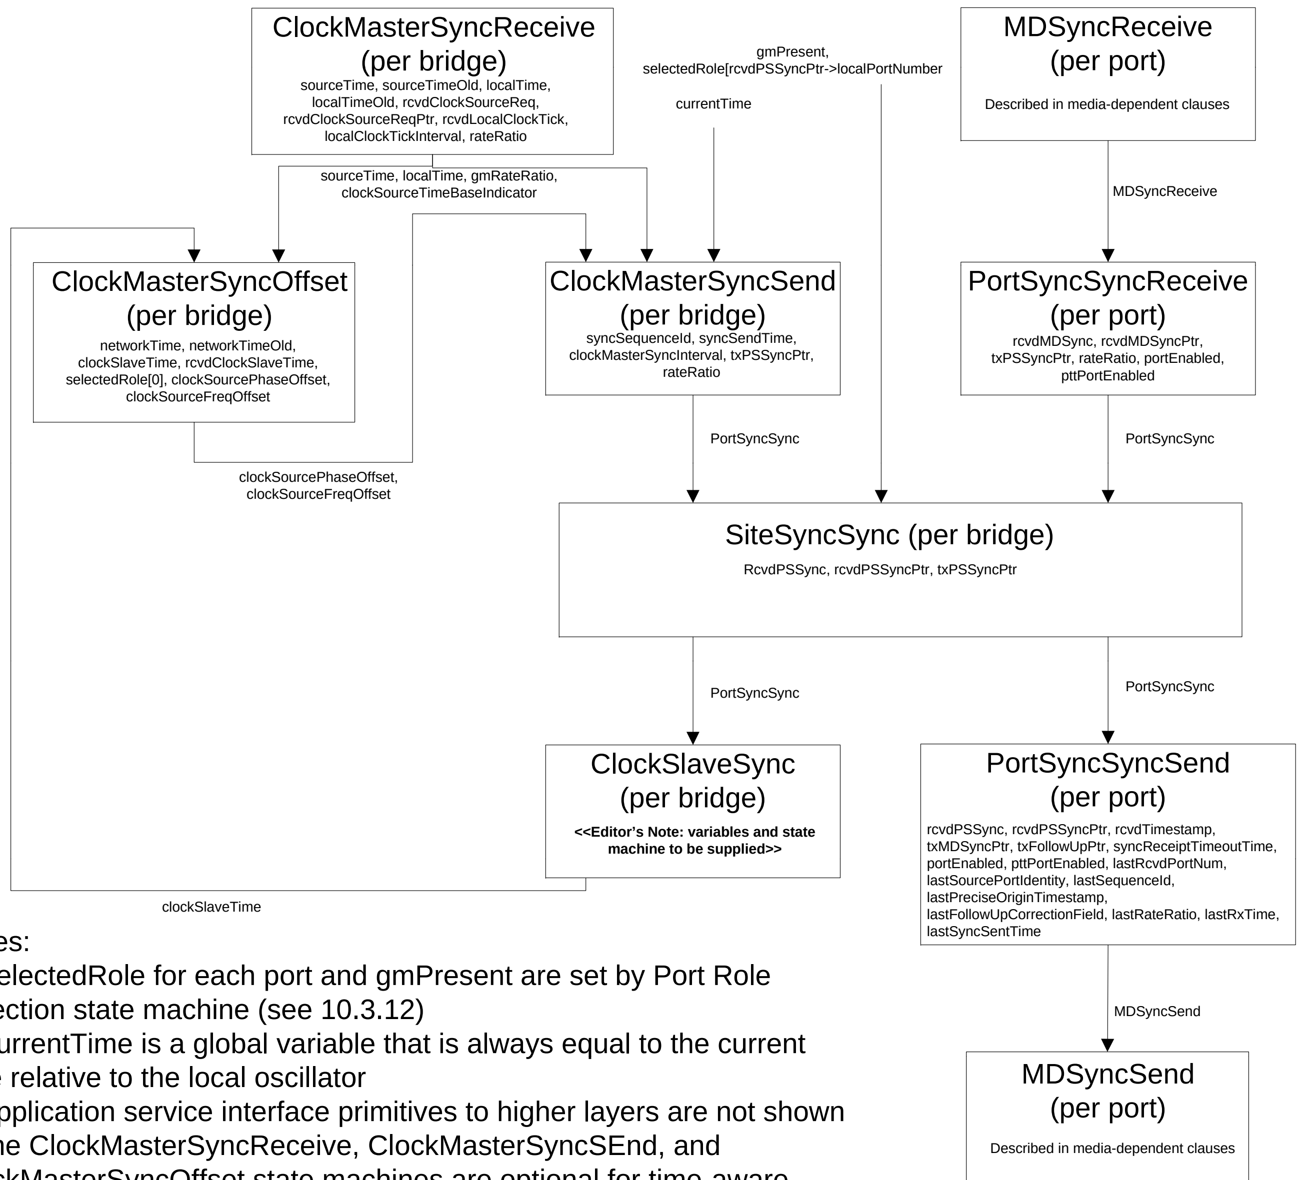

Notes:

a) selectedRole for each port and gmPresent are set by Port Role Selection state machine (see 10.3.12)

b) currentTime is a global variable that is always equal to the current time relative to the local oscillator

c) application service interface primitives to higher layers are not shown

d) the ClockMasterSyncReceive, ClockMasterSyncSEnd, and ClockMasterSyncOffset state machines are optional for time-aware systems that are not grandmaster-capable.

Best Master Selection in IEEE 802.1AS – 1

- IEEE 802.1AS uses a mechanism that is very similar to the default mechanism; there are 3 main differences
  - No qualification of Announce messages, and therefore no consideration of foreign masters
    - BMCA runs on receipt of an Announce message on any port attached to another time-aware system
    - This was done to speed up reconfiguration when the grandmaster changes
  - The pre-master state is eliminated; a port that is determined to be a master port immediately goes to the master state
  - The uncalibrated state is eliminated, because PLL filtering is not done in bridges

Best Master Selection in IEEE 802.1AS – 2

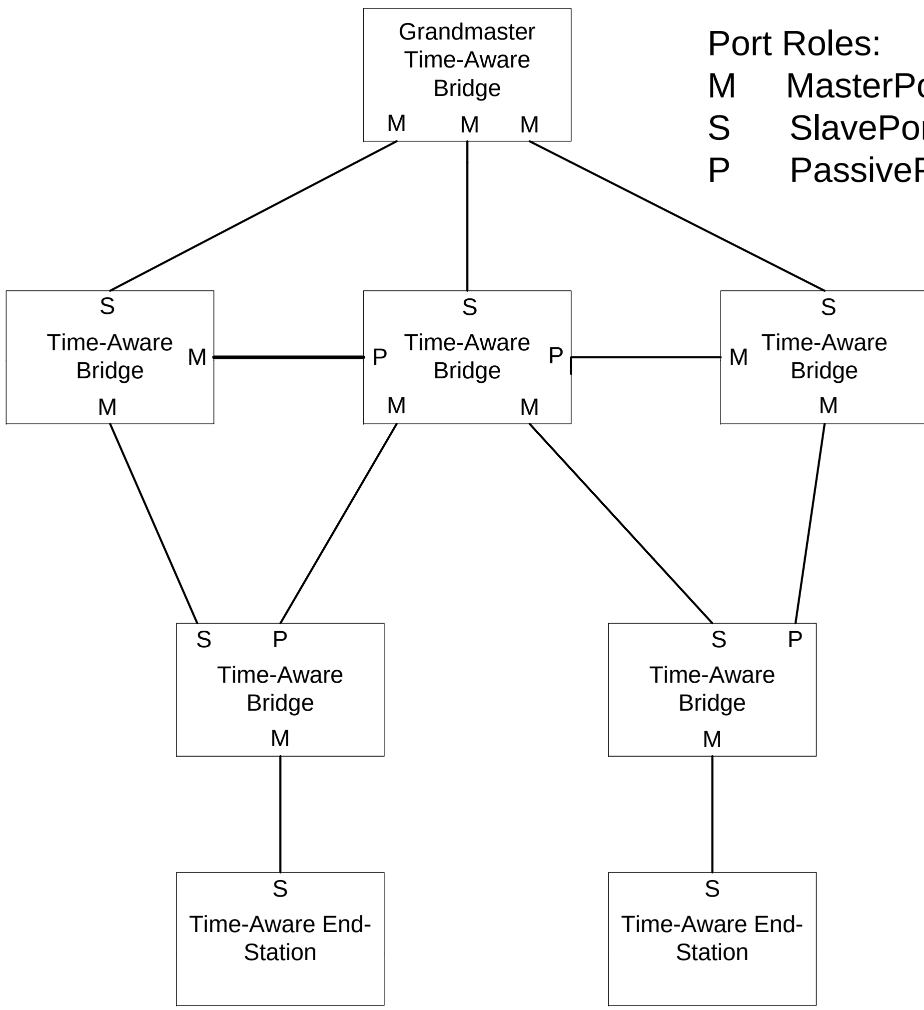

Port Roles:

M MasterPort

S SlavePort

P PassivePort

Best Master Selection in IEEE 802.1AS – 3

- The BMCA is expressed using a subset of the Rapid Spanning Tree (RSTP) protocol formalism of IEEE 802.1D and IEEE 802.1Q
- This formulation is mathematically equivalent to the dataset comparison and state decision algorithms of IEEE 1588
  - Aspects of RSTP pertaining to updating the forwarding data base of a bridge are not needed for BMCA
- BMCA creates a spanning tree, with the GM at the root (unless no time-aware system is GM-capable (see below))
  - May or may not be the same as the data spanning tree created by RSTP

Best Master Selection in IEEE 802.1AS – 4

- The attributes priority1, clockClass, clockAccuracy, offsetScaledLogVariance, priority2, and clockIdentity are concatenated as unsigned integers into an overall attribute systemIdentity
- Part 1 of the dataset comparison algorithm is expressed as a comparison of systemIdentity attributes (smaller is better)
- The most significant bit of priority1 is used to indicate whether a time-aware system is GM-capable (0 if so, 1 if not)
  - If no system is GM-capable, Sync is not sent by any system

Best Master Selection in IEEE 802.1AS – 5

- A spanning tree priority vector is defined, using the root systemIdentity, rootPathCost (number of hops from the root, i.e., 1588 stepsRemoved), sourcePortIdentity, and portNumber of receiving port
- Following IEEE 802.1D, 6 different, but related, priority vectors are defined
- These priority vectors are set and compared in 4 interacting state machines
  - The machines also set the ports to Master, Slave, or Passive
  - The operation of these state machines is equivalent to to the dataset comparison and state decision algorithms (see [14] for an example)

Best Master Selection State Machines

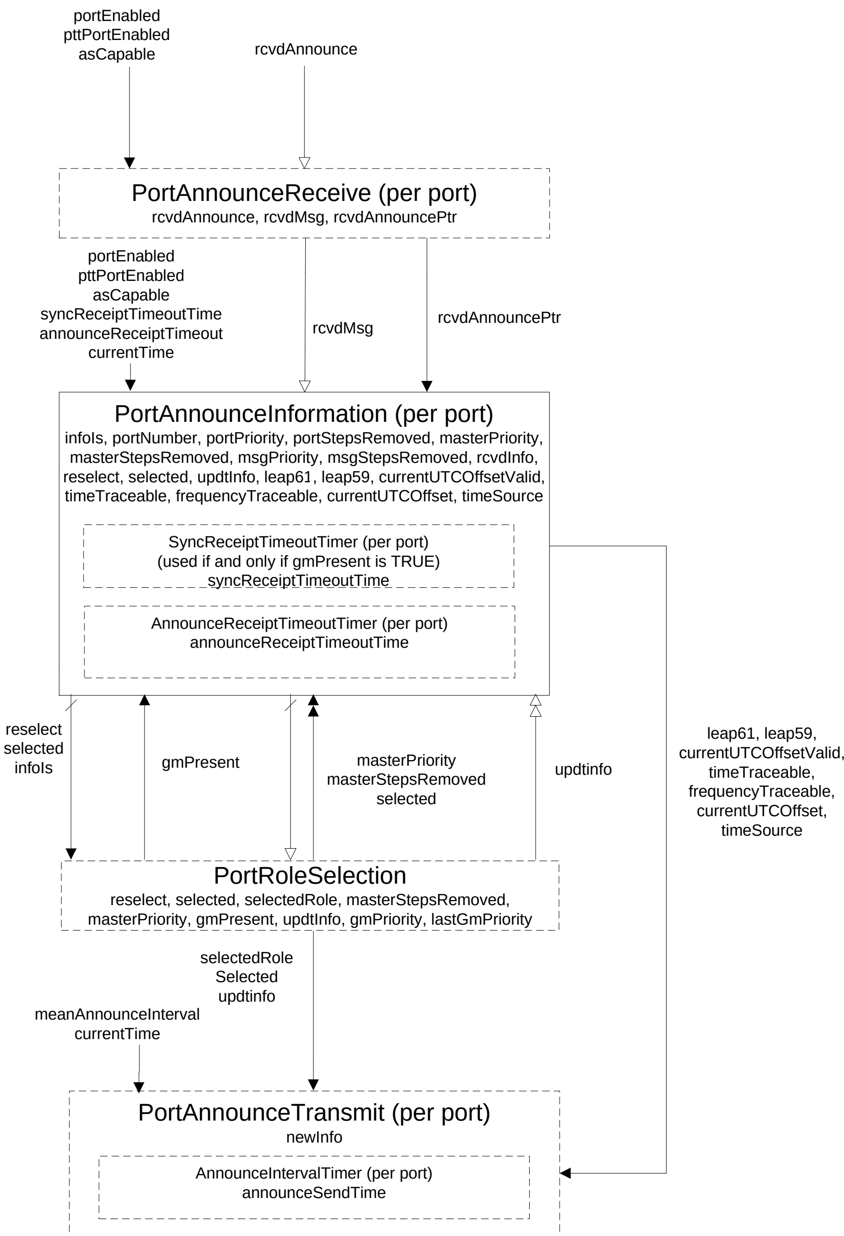

Summary

- Overview of IEEE 802.1AS was presented
- Compatible with IEEE 1588™ – 2008
  - Includes PTP profile
- Specifics chosen for low cost while still meeting performance requirements
- Few options, plug-and-play
- Support added for 802.11, plus annex for CSN
- Alternate BMCA, though very similar to default
  - Simplified to provide for faster convergence
- Unfortunately, test results are not available as of the presentation of this paper. An amendment to this paper will be available in late September, 2008