Summary of Potential Choices for Synchronization Transport in Residential Ethernet

Geoffrey M. Garner
SAMSUNG Electronics (Consultant)

IEEE 802.3 ResE SG
2005.05.19

gmgarner@comcast.net

Outline

- Introduction
- Basic two-way time stamp scheme
- Variations/choices
  - Use one-way time stamp scheme with less frequent two-way exchange; obtain delay from two-way exchange and assume delay is fixed until next two-way exchange
  - Instantaneous phase adjustments at intermediate nodes
  - Instantaneous phase and frequency adjustments at intermediate nodes (with instantaneous frequency adjustments possibly less frequent)
  - Filtered phase adjustments at intermediate nodes, using digital filter running at local clock rate
  - Full phase-locked loops (PLLs) at intermediate nodes (i.e., filtered phase and frequency adjustments)
  - Use of transparent clock nodes
    - End-to-end versus peer-to-peer
    - Whether or not to adjust rate of local oscillator in transparent clock and, if so, whether to do filtering
  - Time stamp reflects current time versus delay by some number of frames
  - Time stamp reflects local free-running clock time versus latest corrected time based on most recent time stamps and possible filtering)

Introduction

- Numerous discussions and presentations in the Residential Ethernet (ResE) SG have identified the need for providing synchronization to the endpoint nodes where time-sensitive services are recovered
  - Jitter and wander requirements have been identified for digital audio, compressed digital video, and uncompressed digital video applications; see [1] for details
  - Inter-stream synchronization requirements have been identified as well; see [1] and [2]
  - Simulation results in [3] indicate that application jitter and wander requirements will not be met if competing time-sensitive traffic streams are transported over current Ethernet bridges with priority classes and timing is recovered by filtering the stream of packet arrivals
- The ResE SG discussions and presentations to date have mainly focused on transporting synchronization using schemes based on time-stamping
  - These schemes use the same principles that are used in IEEE 1588 [4] and Network Time Protocol (NTP) [5]; however, not all the schemes are identical to these protocols in detail

Introduction (Cont.)

- The schemes all employ the same basic two-way time stamp procedure
  - However, a number of variations have been suggested or implied
  - Eventually, the ResE SG must choose which variation(s) to use
- This presentation summarizes the variations discussed so far (or implied), as an aid in helping the ResE SG make the appropriate choices

Basic Time Stamp Scheme

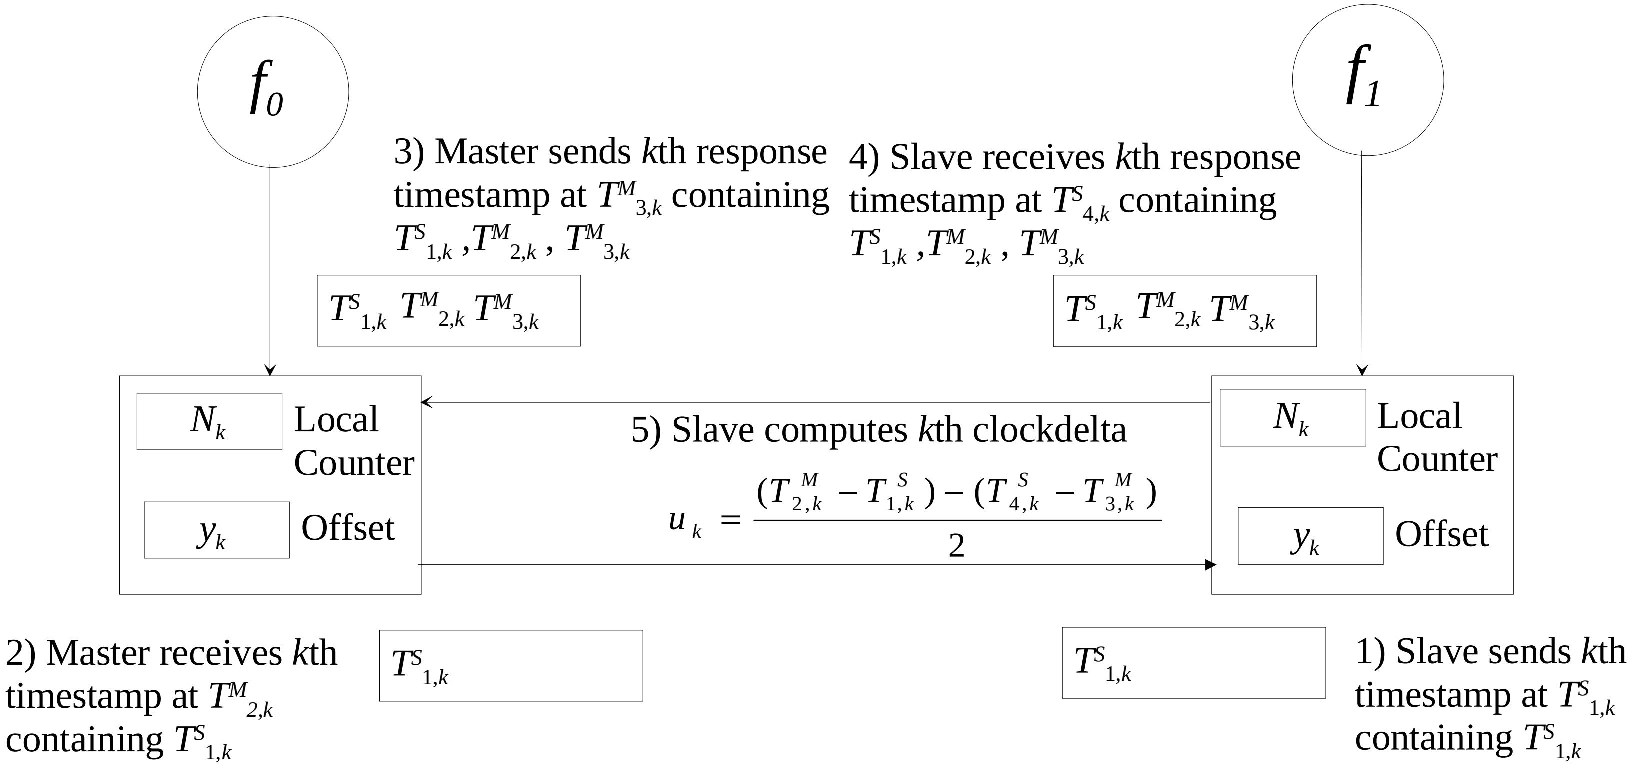

6) Slave computes current offset y_k in terms of current and possibly past clockdelta's u_k

Basic Time Stamp Scheme (Cont.)

- The basic two-way scheme describes information exchanged between a master and slave
  - Allows correction (clockdelta) to be calculated as indicated on previous slide
  - Can also obtain estimate of message delay between the nodes (assumes delay is the same in both directions)

$$\text{delay}_k = \frac{(T_{2,k}^M - T_{1,k}^S) + (T_{4,k}^S - T_{3,k}^M)}{2}$$

- Clockdelta and delay information is used to obtain actual time adjustment
- The cumulative adjustments must be propagated downstream from the grandmaster to all the clocks synchronized to the grandmaster
- There are a number of variations on how the clockdelta and delay information is actually used to calculate the time adjustments at each node
  - These variations are summarized in the following slides
- It is generally assumed that a filtering function (e.g., smoothing PLL) is needed at least at the endpoint where the timing of the Ethernet client (e.g., MPEG-2 packets input to the MPEG-2 decoder, digital audio data stream input to audio codec) is recovered
  - Some of the variations may also use filtering functions at the intermediate nodes

Variations/Choices

1) Use one-way time stamp scheme with less frequent two-way exchange; obtain delay from two-way exchange and assume delay is fixed until next two-way exchange
2) Instantaneous phase adjustments at intermediate nodes
3) Instantaneous phase and frequency adjustments at intermediate nodes (with instantaneous frequency adjustments possibly less frequent)
4) Filtered phase adjustments at intermediate nodes, using digital filter running at local clock rate
5) Full phase-locked loops (PLLs) at intermediate nodes (i.e., filtered phase and frequency adjustments)
6) Use of transparent clock nodes
   a) End-to-end versus peer-to-peer
   b) Whether or not to adjust rate of local oscillator in transparent clock and, if so, whether to do filtering
7) Time stamp reflects current time versus delay by some number of frames
8) Time stamp reflects local free-running clock time versus latest corrected time based on most recent time stamps and possible filtering)

1) One-way and Less Frequent Two-Way Exchanges

- The basic scheme uses only two-way time stamp exchanges
- Since delay is expected to change slowly (if at all), it is possible to do less frequent two-way exchanges in which both clockdelta and delay are obtained
- Between the current and next two-way exchange, more frequent one-way time stamps are sent from the master to the slave
- The slave obtains clockdelta by adding the most recent delay estimate to the received time stamp
- If the frequency of two-way exchanges is much less than the frequency of one-way time stamps, this variation reduces the bandwidth consumed by the messages by approximately a factor of two
    - Assumes the one-way and two-way time stamps are the same size

2) Instantaneous Phase Adjustments at Intermediate Nodes

- In this variation, the time at the slave is adjusted instantaneously by an amount equal to clockdelta, when clockdelta is computed
  - I.e., the phase offset y_k is set equal to clockdelta
  - This variation could be used for both pure two-way and one-way plus two-way schemes
- This scheme has the advantage of being simple
- This scheme has the disadvantage that the instantaneous phase adjustment may be large if the time stamps are not sent sufficiently frequently for the specified clock accuracy
  - E.g., for time stamps sent/exchanged every 10 ms and $\pm$ 100 ppm frequency accuracy, the phase step would be 2 μs
- In addition, the phase error will accumulate over multiple hops
- Eventually must filter the phase error at the endpoint; may need narrow bandwidth filter/PLL

3) Instantaneous Phase and Frequency Adjustments at Intermediate Nodes

- A variation has been proposed [6] where the frequency difference between the master and slave is measured at the slave
- The measurement could be made by counting the number of slave clock cycles between two time stamps received from the master
  - Each time-stamp would indicate the time relative to the master's clock, and therefore their difference would indicate the number of master clock cycles occurring during that interval (the interval would actually be offset by a fixed amount equal to the delay, but this would not affect the calculation for a pure frequency offset)
  - The frequency offset would be computed from the number of slave and master clock cycles
- The slave frequency would be adjusted instantaneously by an amount equal to the measured frequency offset
- The frequency measurement and adjustment would occur less frequently than the instantaneous phase adjustment y_k
  - In [6], the phase adjustments are made every 10 ms, and the frequency adjustments every 100 ms

Instantaneous Phase and Frequency Adjustments at Intermediate Nodes (Cont.)

- This scheme has the advantage of reducing the phase steps when the phase adjustments are made, compared to (2)
- However, this scheme may result in excessive jitter and wander, due to clock phase noise
  - The jitter and wander arise because the frequency offset is computed as a first difference

4) Filtered Phase Adjustments at Intermediate Nodes

- In this scheme, the phase adjustment y_k is computed by filtering the clockdelta sequence
  - This scheme could be used for both pure two-way and one-way plus two-way schemes
- The filter is a digital low-pass filter that is run at a rate tied to the local clock at the slave
- This scheme smoothes the instantaneous phase adjustments of scheme (2)
- Jitter and wander will be reduced compared to scheme (2) by the smoothing
  - Since this is a pure digital filter and not a PLL, there will not be PLL noise generation; any inherent phase noise will be that due to the local clock
  - Jitter and wander will accumulate, but the accumulation will be lower for narrower bandwidth filters and lower gain peaking

5) Full PLLs at Intermediate Nodes

- In this scheme, a full phase-locked loop (sometimes referred to as a servo in the IEEE 1588 literature) is implemented at each intermediate node to filter both phase and frequency adjustments
  - This scheme could be used for both pure two-way and one-way plus two-way schemes
- If PLL gain peaking and noise generation are within limits, jitter and wander accumulation over multiple intermediate nodes can be controlled
  - However, large gain peaking can result in rapid jitter accumulation after traversing some number of nodes
  - In addition, the noise generation requirements may result in expensive oscillators
  - Analysis of the above two points is needed

6) Use of Transparent Clock Nodes

- The master and slave clocks described in the basic time stamp scheme above correspond to Boundary Clocks (BCs) and/or Ordinary Clocks (OCs) in IEEE 1588
- An new clock type, called a Transparent Clock (TC), has been proposed in [7] and also for inclusion in Version 2 of IEEE 1588
  - [7] describes the basic principles of the TC; Version 2 of IEEE 1588 is currently under development
- In the basic time stamp scheme, timing is transported from the grandmaster to the various endpoints where timing is needed
  - In general, the timing transport will traverse one or more BCs
  - At each BC, a clockdelta is computed and used to obtain phase offset (possibly using one of the methods described in previous slides)
  - Timing is then available at each transparent clock, whether or not it actually is needed there
- A TC could replace one or more BCs, i.e., a time stamp between 2 BCs and/or a BC and OC could traverse one or more TCs

Use of Transparent Clock Nodes (Cont.)

- Each TC measures the time it takes for the time stamp to traverse the TC node
  - This is termed the *residential time*
  - The time at which the time stamp was sent is adjusted (incremented) by the residential time at each TC
  - With this adjustment, the basic time stamp exchange process will produce a clockdelta that accounts for variable residential times in the TCs and a delay that reflects only propagation delay and transmission delay on the line
  - Actually, there is a choice of whether to directly adjust the time carried in the time stamp or accumulate the residential time correction in a separate field
    - The current decision in the IEEE 1588 committee is to carry the correction in a separate field

Use of Transparent Clock Nodes (Cont.)

- In addition, if one-way plus two-way time stamps are used, the delays between BCs are needed
  - If there is a network reconfiguration (e.g., a new grandmaster), some amount of time will be needed for the new BC to BC (and/or OC) delays to be obtained, because the two-way exchanges are less frequent
  - To decrease the reconfiguration time, it is possible for the TCs to exchange time stamp messages at various intervals to compute the TC to TC link delays
    - These TC to TC link delays can be used in obtaining BC to BC (and/or OC) delays
  - A new, but similar, set of messages can be defined for TC to TC exchanges
    - These are currently referred to in the IEEE 1588 discussions as Adelay and Aresp.
  - TCs that can exchange messages to determine the delay between them are termed *peer-to-peer* TCs
  - TCs that cannot exchange messages but only measure residential times for the BC and OC messages are termed *end-to-end* TCs
  - The current decision in the IEEE 1588 committee is that both types of TCs must be capable of interworking in the same sub-domain

Use of Transparent Clock Nodes (Cont.)

- In the basic TC (whether peer-to-peer or end-to-end), residential time (and time-stamps for peer-to-peer TCs) is measured using a local free-running oscillator
  - This results in an error in the computed residential time equal to the actual frequency offset between the TC and the BC/OC whose time stamp is traversing the TC, multiplied by the residential time
  - Depending on the free-run accuracy and the magnitude of the residential time, this error may be unacceptably large
  - It is possible for an end-to-end TC to adjust its rate to that of the Master Clock whose time stamp is traversing the TC, using the time stamp information
    - The master clock could be a peer-to-peer TC or a BC
  - It is possible for a peer-to-peer TC to adjust its rate to the next upstream peer-to-peer TC (or BC if this is the first peer-to-peer TC following the BC)
  - In principle, any of the schemes (2) – (5) could be used

7) Time Stamp Reflects Current Time Versus Delay by Some Number of Frames

- Both the pure two-way and one-way plus two-way schemes require that the time stamp values reflect as accurately as possible the actual time the time stamp was sent or received
- Any errors in the time stamp values will, in general, result in errors in the clockdelta and computed phase offset (for all the previous schemes)
- Depending on the implementation, it may be possible to make a more accurate measurement of the time the time stamp packet leaves a node if it is not necessary to write the measurement in the timestamp
- To take advantage of this, it has been proposed (see [6]) that time stamp values reflect the times of the respective packet in a previous packet exchange some number of frames in the past, rather than the current packet exchange
  - This is similar to the follow-up message in IEEE 1588
  - This issue also arises for TCs; the current decision in the IEEE 1588 committee is that TCs that are follow-up-capable and those that are not follow-up-capable must be able to operate together in the same network
- Since phase and frequency adjustments are applied at the current time, the method assumes that changes in delay and local clock frequencies are small over one time stamp interval

8) Time Stamp Reflects Local Clock Time Versus Latest Corrected Time

- In the basic time stamp scheme, clockdelta is computed based on the local counters tied to the local, free-running clocks
- Therefore, the computed y_k at the slave represents the time offset between the slave and the master
- We are eventually interested in the time offset between the slave and grandmaster
    - Therefore, the cumulative offset between each successive slave in a clock chain and the grandmaster must be accumulated
- An alternative approach would be to compute the clockdelta at a time step using the local clock plus current offset value
    - In this approach, the computed y_k at the slave represents the time offset between the slave and the grandmaster
- This scheme would not require separately accumulating and transporting the cumulative phase error
- The jitter/wander performance of this scheme must be analyzed

References

1. Geoffrey M. Garner, *End-to-End Jitter and Wander Requirements for ResE Applications*, Samsung presentation at May, 2005 IEEE 802.3 ResE SG meeting, Austin, TX, May 16, 2005.
2. Ralf Steinmetz, *Human Perception of Jitter and Media Synchronization*, IEEE JSAC, Vol. 14, No. 1, January, 1996, pp. 61 – 72.
3. Geoffrey M. Garner and Felix Feng, *Delay Variation Simulation Results for Transport of Time-Sensitive Traffic over Conventional Ethernet*, Samsung presentation at July, 2005 IEEE 802.3 ResE SG meeting, San Francisco, CA, July 18, 2005.
4. IEEE Std™ 1588-2002, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, IEEE, November 8, 2002.
5. Mills, L. David, *A Kernel Model for Precision Timekeeping*, RFC 1589, March, 1994.

References

6. *Residential Ethernet (RE) (a working paper)*, Draft 0.133, maintained by David V. James and based on work by him and other contributors, July 12, 2005, available at http://dvjames.com/esync/dvjReNext2005Jul12.pdf.
7. Sven Nylund and Oyvind Holmeide, *IEEE 1588 Ethernet Switch Transparency – No Need for Boundary Clocks!*, 2004 Conference on IEEE 1588 Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, Gaithersburg, MD, September 27 – 29, 2004, pp. 48 – 55.