Localization approaches based on Ethernet technology

Kees den Hollander, GM Garner, Feifei Feng, Paul Jeong, Eric H.S. Ryu

Contact: Eric H.S. Ryu

eric_ryu@samsung.com

Abstract

This document describes two methods to determine the distance from one end-station to another end-station in an Ethernet based packet network.

The first method relies on special treatment in Ethernet Bridges of an Ethernet ping packet that is used to measure round-trip delay. The ping packet is "cut through" at each intermediate station to ensure minimal delay. Any packet currently in transmission is pre-empted, and retransmitted after completion of transmission of the ping or response. Hardware at the receiver ensures minimal delay at the receiver between the receipt of the ping and the sending of the response.

The second method uses the timing information that is exchanged between 802.1AV bridges according to the proposed 802.1as timing protocol and accumulates this information to determine to total distance between end-stations.

Introduction

Future converged home networks will carry high value applications that need Digital Rights Management (DRM). Such applications include the replay and rendering of time-sensitive digital video and digital audio content. With these applications, a Digital Media Player (e.g. a display) may be connected to a Digital Media Server via a network and use a subscription process, among other things, to set up a connection and reserve the necessary bandwidth between Server and Player. In some networked applications, the owner of the content that is being replayed may require that the application be restricted to a limited geographic area (e.g., a residence)

For example, the owner of the content on a DVD may require that, when the DVD is played, the content be displayed only in the limited local area of the user (e.g., the user's residence) and not be transported beyond this area (e.g., to another residence). To enforce this, the Media Server where the content originates must determine whether the Media Player is within a threshold distance of the sender.

In current schemes, the Media Server sends a ping packet to the Media Player and measures the round-trip time to determine the distance between Player and Server:

- The Server send a ping and records the time it sends the packet
- The Player responds with a ping response packet
- The Server records the time is receives the response, and computes the round-trip delay

If the round-trip delay exceeds a specified threshold, the Server concludes that the Player is beyond the threshold distance and does not set up a connection to transmit the content.

Problem of current delay measurement method

In the current application of this method, the ping and ping response may experience significant queuing delay at intermediate stations. Even if the ping or response are given the highest priority, the prioritization in current networks is non-preemptive; a ping or response may experience significant delay due to waiting for a frame that is already in transmission when the ping or response arrives to complete. For example, the transmission delay for a maximum sized Ethernet frame (1518 bytes, including Ethernet overhead) on 100 Mbit/s Ethernet is 121 us. With a few hops, delays due to transmission and queueing can approach a sizeable fraction of a ms (and more than a ms for many hops).

Additional delay may be experienced at the Player between arrival and sending of ping response if the receiver does not have dedicated hardware or give highest pre-emptive priority to the processing of the ping. The delay due to signal propagation in a typical LAN or WAN that uses full-duplex, point-to-point links is 8 ns/m. Therefore, the delay due only to signal propagation for a distance of 30 – 100 m is 240 – 800 ns, i.e., less than 1 ms. This means that the delays due to queuing and waiting for frames currently being transmitted to complete can vastly exceed the delays due to propagation over the threshold distance.

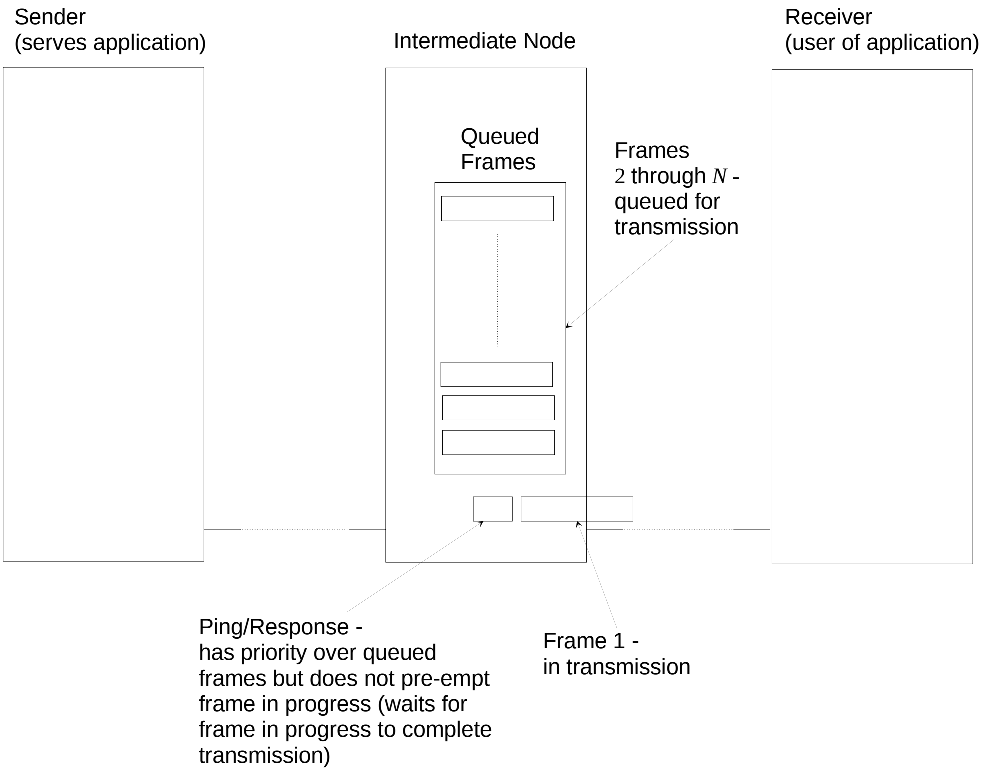

Current forwarding mechanism

To allow for delays due to queueing and completion of transmission of frames in progress (when a ping or response arrives at a station), the delay threshold must be set larger than the amount due to propagation, probably by a factor of 100 or more. The result is to greatly decrease the reliability of the method; setting the delay threshold this large to avoid denying service to legitimate users in heavily loaded networks means that Players outside the distance threshold may receive contents if the network is sufficiently lightly loaded.

Approach #1

The proposed method eliminates the variable queuing and transmission delay for the ping and ping response by cutting-through the ping and ping response at intermediate stations without any buffering: Each bit of the ping or response is transmitted as soon as it is received; the only delay is that due to the receiver, clock recovery, ping recognition and transmitter circuits.
Giving the ping and ping response pre-emptive priority over all other frames, requires that a copy of each regular frame must be retained in the output queue until transmission of the frame is completed in order to prevent loss of frames. If a frame is pre-empted, transmission of the full frame is repeated after completion of transmitting the ping or ping response.

In AV bridges dedicated hardware at the MAC layer is needed to perform the following functions:

- Examine incoming packet
- If it is neither a ping nor a response, handle it as a normal frame (traffic, management, synchronization, etc.)
- If it is a ping or response not destined for this station, cut it through to the appropriate output port, pre-empting any non-ping or non-response frame currently being transmitted on that port
- Repeat the transmission of that frame after transmitting the ping or response
- Do not release the current frame from its queue until its transmission is completed
    - There is no limit on the number of attempts in transmitting the current frame (because the frame can be pre-empted only by a ping or response frame

At the Player end-station we also need dedicated hardware to cut-through the ping and change it to a ping response on-the-fly. The following operations are required:

- Change destination address to source address and vice-versa,
    - Can do this in cut-through manner easily in Ethernet by placing response packet source in source address field as ping source field is arriving; then place ping source in response destination field as destination field is transmitted
- Change packet type to response
    - Can do this in cut-through manner easily in Ethernet because the packet type (response) is known
- Respond to the challenge in the ping using a security key that is shared with the Server

The Server end-station can calculate the distance between itself and the Player from the round trip delay of the ping and decided whether this distance is within the desired threshold.

As discussed below, the cut-through operation minimizes the delay of the ping through an intermediate bridge, but the bridge transit delay is still significant compared to the propagation delay on the links between bridges. It is however a value with little variation for which a lower and an upper bound could be specified for AV bridges. By using, for instance, the proposed Linktrace protocol of IEEE802.1ag the number of bridges on the path between Server and Player can be determined by the Server end-station. The measured RTT can then be corrected for the number of bridges times the minimum delay per bridge to provide a more accurate measurement of the true distance between Server and Player.

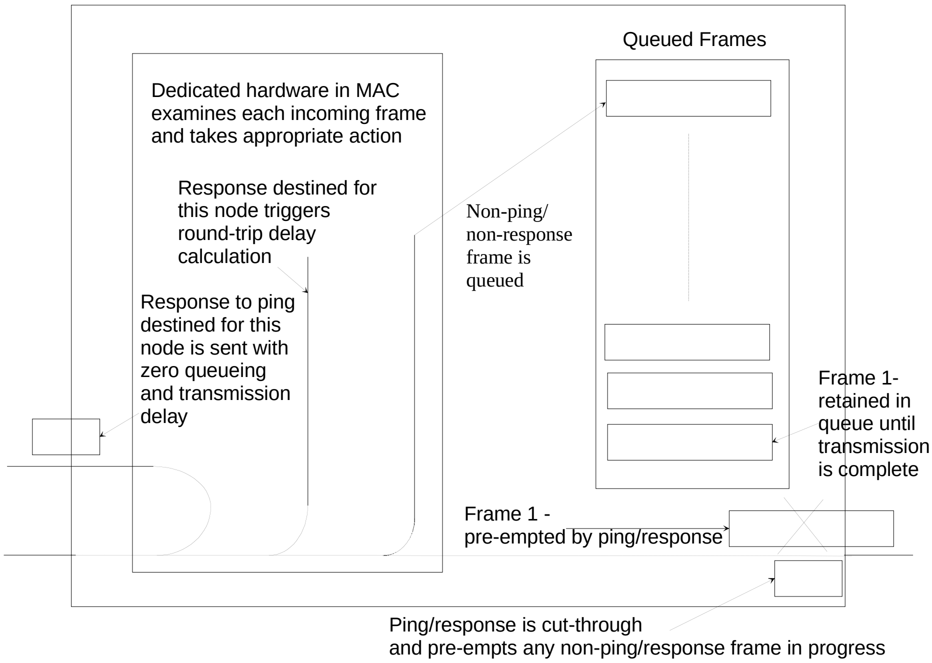

Achievable accuracy using existing Ethernet hardware

100 Mbit/s Ethernet case:

Clock for interface between MAC and MII and between MII and PCS is 25 MHz. If the departure of the ping is measured relative to the MAC/MII interface, the best achievable accuracy is 40 ns. If the arrival of the ping response is measured relative to the MAC/MII interface, the best achievable accuracy is 40 ns. Assuming the measurements are truncated to the nearest 40 ns, the accuracy of the delay measurement is ±40 ns. This neglects any additional delay in the MII, PCS, or PMA layers.

- IEEE 802.3, clause 24.2.2.3 indicates there will be at least 2 nibbles (8 bits) of delay in the PCS in aligning from bits to nibbles; this is equivalent to 80 ns
- IEEE 802.3, clause 24.6 indicates that the delay between MDI and MII must not exceed 14 bits (140 ns) on transmit and 32 bits (320 ns) on receipt

Therefore, while ±40 ns is the best achievable accuracy, the actual accuracy will be lower. Assuming propagation delay of 8 ns/m, an accuracy of ±40 ns corresponds to ±5 m

- A worst-case additional error of 140+ 320 ns = 460 corresponds to 57.5 m

1 Gbit/s Ethernet case:

Clock for interface between MAC and MII and between MII and PCS is 125 MHz. If the departure of the ping is measured relative to the MAC/MII interface, the best achievable accuracy is 8 ns. If the arrival of the ping response is measured relative to the MAC/MII interface, the best achievable accuracy is 8 ns. Assuming the measurements are truncated to the nearest 8 ns, the accuracy of the delay measurement is ±8 ns. This neglects any additional delay in the MII, PCS, or PMA layers

- IEEE 802.3, clause 36.5 indicates that the delay between MDI and MII must not exceed 136 bits (136 ns) on transmit and 192 bits (192 ns) on receipt
- Therefore, while ±8 ns is the best achievable accuracy, the actual accuracy will be lower

Assuming propagation delay of 8 ns/m, an accuracy of ±8 ns corresponds to ±1 m.

- A worst-case additional error of 136+ 192 ns = 328 ns corresponds to 41 m

Achievable accuracy using specialized hardware

If specialized hardware is available to measure the departure time of the ping and the arrival time of the ping response within the PMA layer after conversion to a serial bit stream using a clock at the full line rate, then the arrival and departure times may be resolved to ±1 bit period. For 100 Mbit/s Ethernet, the line rate is 125 MHz and the resulting accuracy is ±8 ns, or ±1 m. For 1 Gbit/s Ethernet, the line rate is 1.25 Gbit/s and the resulting accuracy is ±0.8 ns, or ±0.1 m.

Approach #2

For IEEE802.1as network timing/synchronization schemes that already measure propagation delay between each pair of stations, the propagation delays for the successive links on the path can be accumulated to calculate the distance between two end-stations. In this approach a ping is sent from Server to Player with a field in which each intermediate bridge writes the addition of the received value and the known (from the synchronization scheme) link delay of the link on which the ping has been received.

The procedure is as follows:

- The Server station sends a ping with a propagation delay field that is initialized to zero,
- Each intermediate bridge adds the propagation delay for the incoming link to the value in the propagation delay field,

- The Player station notes whether the accumulated propagation delay is within a set threshold and communicates this back to the Server.

Note that there is no response ping in this approach. The message to communicate to the server that the propagation delay is within the allowed threshold is not time critical. Alternatively also the measured propagation delay may be communicated back to the Server.Also note that the IEEE802.1as synchronization schemes generally assume that the link propagation delay is the same in both directions. The propagation delay of a link is calculated as the average delay of the up- and the down-link, and is therefore only an accurate measurement of the down-link if the link is symmetrical. This seems however to be a fair assumption for full-duplex networks, and AV Bridging is designed for full duplex networks only

Achievable accuracy of Approach #2

The achievable accuracy depends on how accurately the propagation delays are measured by the network synchronization scheme. If the peak-to-peak error in propagation delay measurement is ε and the number of links is N, the maximum peak-to-peak error is Nε. For example, if there are 7 links and propagation delay can be measured to an accuracy of ±40 ns, then the possible measurement error is ±280ns. Assuming a propagation delay of 8 ns/m, this corresponds to ± 35 m.

If the industry can agree on approach number 2, the marginal implementation effort to measure distance accurately will be relatively simple. It is also a scheme that is extremely simple to specify and standardize. The achievable accuracy is by far superior to that of current schemes that rely on a ping at the IP layer. There are possibilities to reach an even higher accuracy than the ± 35 m mentioned above, but this would incur additional complexity, and more stringent requirements should not be imposed lightly for cost-critical home networks.

Security

Approach #1 has the advantage from a security point of view that it is implemented in hardware in intermediate bridges and therefore not easy to tamper with by the general public. The only security association that has to be created is between the Server and the Player, in order for the Player to respond to a challenge that is provided by the Server. Many Key Agreement protocols could be used for this purpose, for instance the IEEE802.1af AKA protocol.

The enhancement to compensate the RTT measurement for the delay of intermediate bridges does however rely on information that is provided by Linktrace messages from intermediate bridges, and this is inherently less secure. To secure this information IEEE 802.1ae provides a standardized way to secure the information that is passed on a Local Area Network. This standard relies on the use of IEEE 802.1af to distribute keys between stations and this level of security does obviously carry complexity. The risks of contents theft will have to be carefully weighted against the costs of a highly secure solution.

Approach#2 does totally rely on the information by intermediate bridges. IEEE802.1ae and .1af are therefore also applicable to ensure that this information is reliable. Costs and operational complexity of a highly secure solution for the user will again have to be weighed against the risk of fraud.

Conclusion

The introduction of new capabilities in Ethernet to support high quality audio-video bridging offers the opportunity to improve the localization accuracy of a Media Player relative to a Server with at least four orders of magnitude. Further discussion is needed in the industry to determine how far we should drive this accuracy and how secure the measurement should be. The latter needs to take the security needs and mechanisms that are considered at the application layers into account. The advantage of applying security at the Ethernet layer is that an integrated solution can be provided for residential users that is also applicable for security and privacy control of wireless connectivity solutions in the residential network.