Harman International

Specifying SRP Acceptable Latency
(Talker vs. Listener)

Craig Gunther (craig.gunther@harman.com)

March 2014, Beijing

Purpose of presentation

Define an enhancement to the SRP protocol that allows a "network" flexibility in creating routes when latency greater than the SRP *AccumulatedLatency* is acceptable.

For example, when SPB sends Talker Advertisements across the Shortest Path to a Listener the *AccumulatedLatency* will be the "best" latency that can be provided by the network.

When a reservation is eventually requested by the Listener this enhancement would allow a PCE to create a non-Shortest Path route to avoid links with insufficient bandwidth for the stream, thus avoiding a "failed" reservation.

In addition this enhancement would allow a PCE to select and reroute existing streams in order to make bandwidth available on a given link for other streams with more stringent latency requirements.

Example 1: Audio to Four Speakers

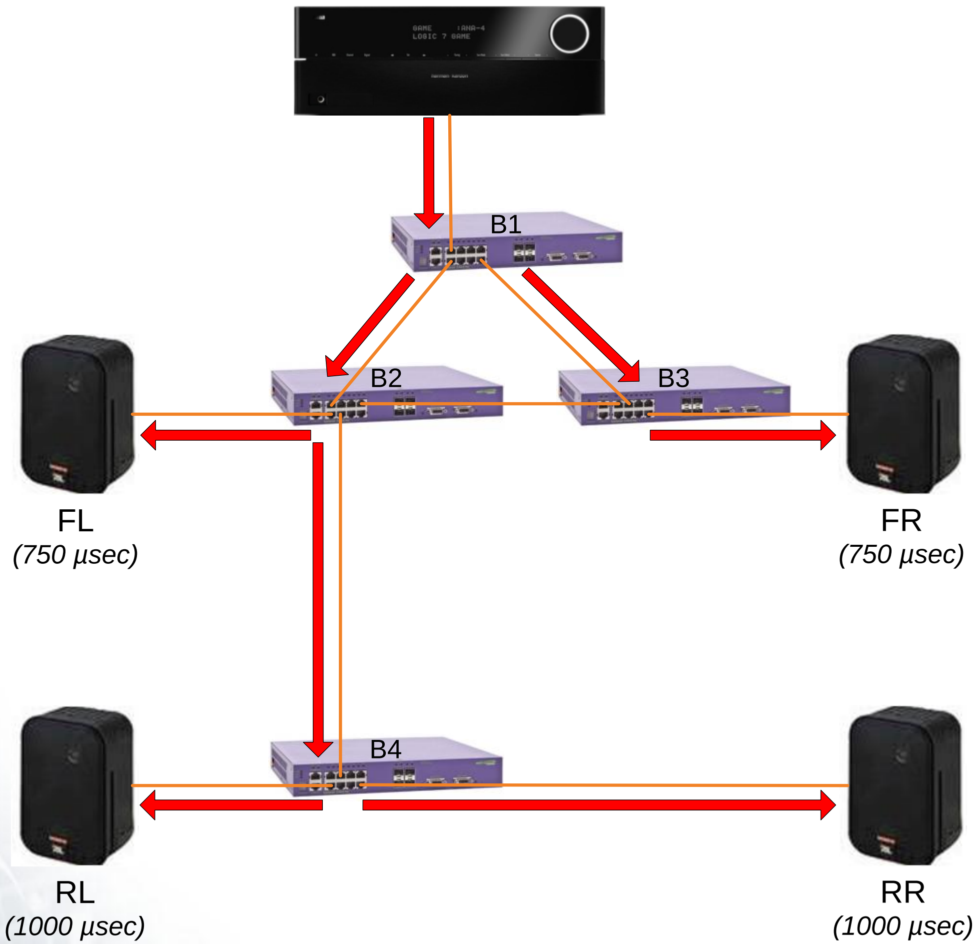

In this example the Talker is going to transmit a 4-channel **audio** stream along the **red** shortest path to each Listener.

Before the Talker begins streaming, a higher layer protocol (1722.1) would gather the `AccumulatedLatency` (Q:35.2.2.8.6) from each of the Listeners. The stream transport protocol (1722) in the Talker would then be configured with a presentation time offset of 1000 µsec, which matches the maximum reported `AccumulatedLatency` from all Listeners*. This will cause the FL and FR speakers to buffer 250 µsec of streaming data so the **audio** played from all four speakers will be synchronized.

* Implicit in this example is the fact that the controlling device knows when all Listeners have reported their latency.

Example 2: Audio + Video to Four Speakers

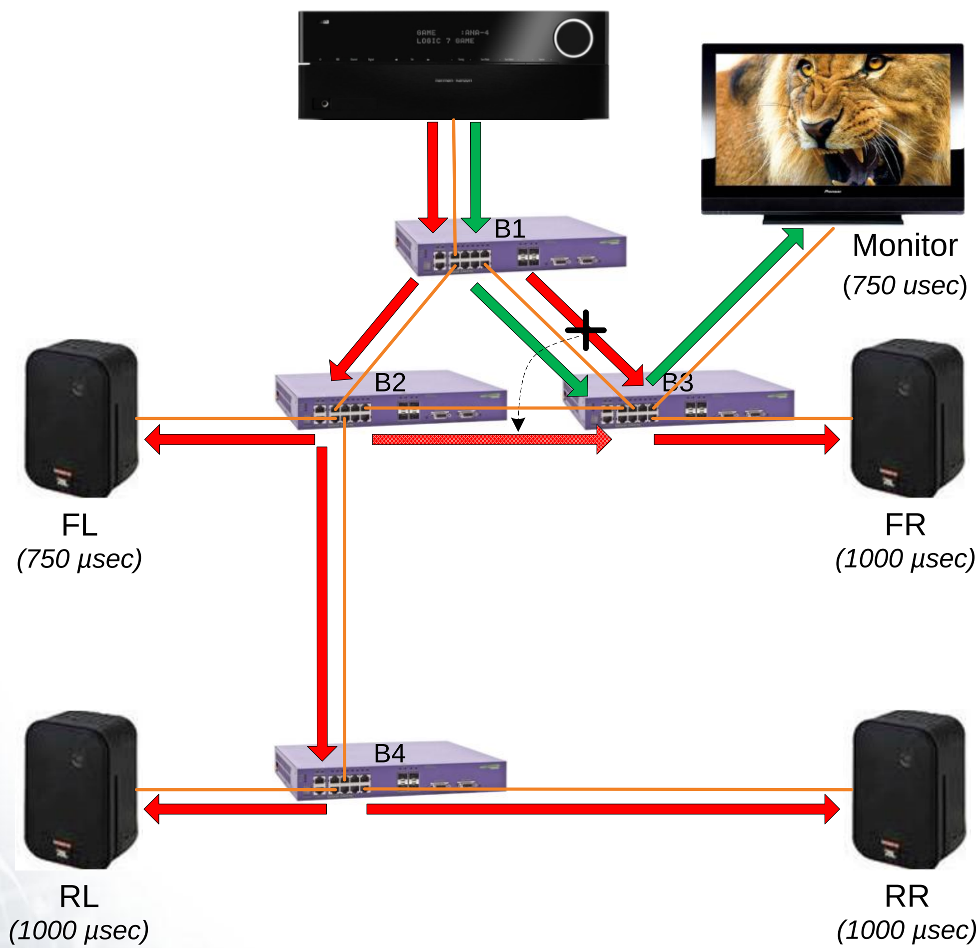

This example adds a **video** stream.

Assume that the B1-B2 and B1-B3 links have insufficient bandwidth to support both the **audio** and **video** streams. Therefore the "network" realizes the **video** stream must replace the B1-B3 **audio** stream.

Since the **audio** stream path to the FR speaker has now been removed a different path is selected via B2-B3. Human observation shows this is acceptable since the latency to FR does not exceed the latency to RL or RR. Therefore this configuration will still work with a presentation time offset of 1000 µsec as specified in Example 1.

How can the "network" know this path selection is acceptable?

Network Path Rerouting

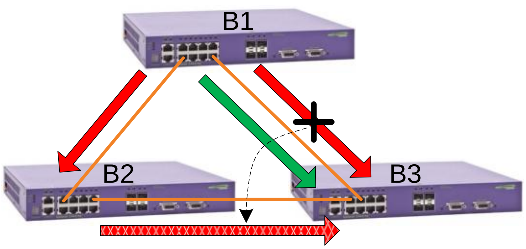

Currently SRP (Q-2011) specifies that a Talker Advertise reports the `AccumulatedLatency` which represents the maximum delay a streaming packet would experience along the associated path.

This enhancement would allow Talkers and/or Listeners to inform the network of the maximum acceptable latency for a stream. Examples 1 and 2 would specify this at 1000 µsec, which is equal to the maximum of the `AccumulatedLatency` reported at FL, FR, RL, and RR (Listeners). The proposal is to do this via a new `AcceptableLatency` field added to the Talker and/or Listener messages, but which one?

Armed with this `AcceptableLatency` the "network" will now have enough information to perform the necessary path rerouting.

Example 3: Garage Band (multiple streams)

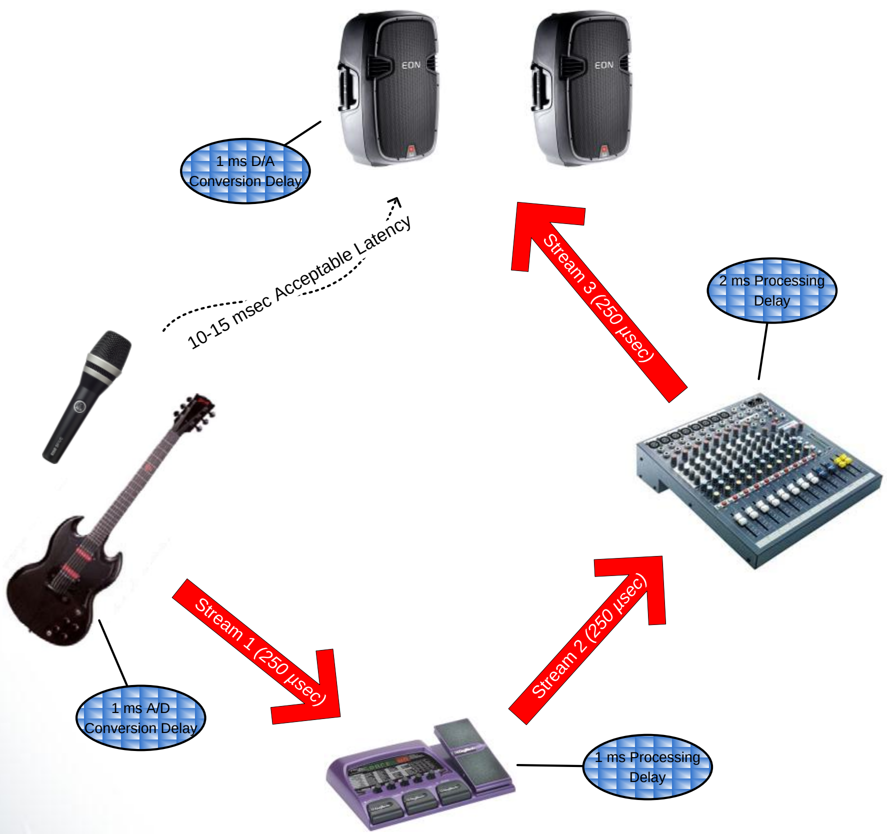

A musician can accept a 10-15 ms delay between playing/singing a note and hearing the amplified sound*.

This illustration shows three AVB streams are required to mix pitch corrected vocals and play them on self-powered speakers.

The total latency in this diagram is 5.75 ms (1 + 0.25 + 1 + 0.25 + 2 + 0.25 + 1). That leaves 4.25 ms (out of 10 ms) for network rerouting. The 250 µsec stream latency numbers are provided by the `AccumulatedLatency` field in the SRP Talker Advertise. Without this 250 µsec "baseline" it would be difficult to derive the `AcceptableLatency`.

The "network" can be allowed to reroute any of the three streams as long as the additional overall `AcceptableLatency` of 4.25 ms is not exceeded.

* See: http://www.ieee802.org/3/re_study/public/200411/beliaev_1_1104.pdf

Example 4: Recording Truck

Live concerts have very stringent audio latency requirements in order to create a consistent sound field. This implies that the `AcceptableLatency` is equal to the maximum `AccumulatedLatency` as shown in Examples 1 and 2.

However, that same audio can be sent to a Recording Truck which has very flexible latency requirements. The audio samples can arrive at the truck late (i.e. after the specified Presentation Time) and still be acceptable since they are being recorded for later playback. Therefore the `AcceptableLatency` can be much greater than the `AccumulatedLatency` if desired; thus allowing the network to route those streams wherever necessary. Similar "tricks" can be used for audio played in overflow areas of concert halls and mega-churches as well as airports, etc.

This is only possible with Listener specified `AcceptableLatency`.

AcceptableLatency Summary

These are the concepts discussed in this presentation:

- Examples 1 & 2 illustrate the need for a new `AcceptableLatency` field in SRP in order to allow the network the flexibility to support more streams.
- Example 3 demonstrates how the `AcceptableLatency` could be related to multiple serial streams. It shows that the current `AccumulatedLatency` in the Talker Advertise is still required.
- Example 4 suggests why Listener-based declarations of `AcceptableLatency` is preferable to Talker-based declarations.
- In order to perform these actions a “network” (PCE?) is required to learn about the SRP stream requirements when it does initial routing and later if it needs to do rerouting to support more streams. This implies a certain knowledge of the network topology.

Harman International

Premium Solutions for Audio and Infotainment