

RBR Study group meeting, teleconference
2004Aug17
13:00-15:55 PST
Dial-in Number: 1-702-835-5000 (Las Vegas, Nevada)
Minutes by David V. James

David V. James	JGG	dvj@alum.mit.edu
David Gustavson	SCIzzL	dbg@SCIzzL.com
Tom Mathey	Self	tmathey@concentric.net

The Acting Working Group Chair (DVJ) chaired the meeting.
Pre-meeting activities included an outside on-the-lawn BBQ lunch.
Follow-on discussions were held indoors, to withdraw from the heat.

No 30-day notice was provided, so all decisions are subject to review and change at the first official meeting.

**** Next meeting ****

The target date for the next meeting is mid to late October.
That fits between the 802.3 interim meeting in Ottawa and the following Plenary in San Antonio.

As such, we can avoid conflicts with both meetings.

The Chair noted that coordination with Leonard Tsai has been difficult. Although we appreciate his kind offer to host the next meeting, we may have to find a Sponsor that travels in a more predictable or less frequent fashion. Otherwise, it's very difficult to plan the meeting time with sufficient time for others to arrange travel.

**** Summary of 802.3 activities ****

Tom Mathey provided an overview of 802.3 activities, with the disclaimer that he wasn't able to track all activities in detail.

Link security and frame extension activities were discussed, but seemed to have little effect on P1796. Tom noted that the frame extensions might be difficult, and some members had noted difficulties with various protocols as a result of frame-length extensions.

10G Ethernet is progressing. One vendor had requested a complete solution be proposed before details are voted, since individual decisions might not be consistent with a well-integrated whole.

**** Prototypes discussion ****

Tom Mathey noted that credibility would be improved if the P796 solution could fit into an ASIC, preferably an ASIC size below the knee of the curve (i.e. a cost effective size).

**** RBR properties ****

Tom Mathey and David Gustavson noted that RBR needs to have well defined properties that distinguish it from other efforts. Slides should clearly illustrate the following:

- 1) What was removed from RPR, when creating RBR.
 - * Simplifications and option elimination listing needed.
- 2) What has been added to RBR, based on experience with IEEE 1394 and IEEE 1596?
 - * Lossless transient congestion management, in 1394 & 1596
 - * Efficient request/response transactions, "
 - * Real time delivery, based on 1394 concepts.

**** Lossless congestion management (LCM) ****

We discussed lossless congestion management (LCM) is more detail.

Objective: Do not lose frames in transient congestion conditions, such as many processors accessing one memory semaphore.

Reason: The overhead of invoking end-to-end frame-loss recovery protocols (in terms of extra transmission latencies and bandwidth) is large, more than offsetting modest overheads that might be associated with LTCM.

Strategy: Limited busy-retry (LBR). The local source receives an ack, which is either busy or done. If the local ack is busy, the request is retried at the proper time and (in the absence of transmission errors) will succeed.

Overhead: A bandwidth loss of 50% for LCM (during congestion) is acceptable, since end-to-end retry overhead is much larger.

Coding: This should probably be an option, with a distinct code. No such code is needed for classA bandwidth, since this capacity is pre-reserved and such real-time traffic cannot afford an additional retry latency.

The rationale for LBR is the lack of other viable alternatives. Any sort of pre-reservation flow control is very hard to manage, when one local destination may be hit with frames sourced by many (up to 254) stations.

We noted that LBR is likely to be more effective on RBR, than on standard Ethernet, due to two factors:

- 1) Each RBR "ring" can consist of multiple stations, while Ethernet has only open two-station rings. As a result, some of the congestion management (resolving conflicts with other ring-local sources) occurs naturally as part of the baseline RBR ringlet bandwidth sharing protocols.
- 2) The LBR approach allows each "busy retry" decision to be made in a frame-contents dependent fashion. For example, the local destination could limit busy's to frames with congested-port destinations or lower-class markings.

**** Simplicity ****

Tom Mathey noted that all frames should have the same format. The design of high-speed sequencers is greatly complicated if each frame type is processed differently. Don't put the CRC in frame-type dependent places, for example.

David James noted that the RPR draft (from which large portions of RBR have been formed) has some "interesting" frame-type distinctions, including parity, partial coverage, and different header lengths. There would be a small overhead if these were to be eliminated (minimum frames would grow from 16 bytes to 20 bytes in size).

Tom noted this modest increase in size would be well worth the associated hardware savings. The extra bandwidth penalty was compared to almonds and cashews eaten by rodents.

There was general concurrence on this observations, and David James offered to update the current RBR draft accordingly.

**** Draft status ****

The current RBR draft has progressed rapidly, based on extensive simplification editing and inputs from this meeting. HOWEVER, this should NOT be considered as unchangeable or even vote-required to change. Its simply a working document useful for communicating between non-attending members.

The Acting Chair (David James) noted that many MSC standards differ

RbrMinutes2004Aug17.txt

from 802 standards, with respect to working documents. The presence of a working draft (or several drafts, in the presence of conflicts) in the MSC represents a useful communication tool, with more detail than can be provided in slides. Documenting two options in one draft is perfectly acceptable, since it allows both to be understood in detail and in context, before binding votes are taken.

This is different than 802, where slide presentations receive binding votes and the text is then written.

Either technique works, but there is danger if the RBR working draft is miss-interpreted as binding content. For now, the entire RBR draft is truly SUBJECT TO CHANGE.

The intent of the current draft is to help provide a common base of understanding. Since many things are different from 802.17 RPR (the source of much material), it was deemed unwise (and possibly copyright sensitive) to distribute unchanged 802.17 RPR drafts as background material.

Secretary for the day,
DVJ