

**CONGRUENT SOFTWARE, INC.**  
**98 Colorado Avenue**  
**Berkeley, CA 94707**  
**(510) 527-3926**  
**(510) 527-3856 FAX**

FROM: Peter Johansson  
TO: NCITS T10  
DATE: September 19, 2000  
RE: SBP-2 study group

---

I convened a study group during the T10 Huntington Beach meetings to review the current state of the art and to explore the possibility of enhancement or extension to Serial Bus Protocol 2 (SBP-2). The following items were on the informal agenda discussed by the group:

- a) Reduced start-up latency from idle condition
- b) Explicit description of how to transport 16-byte or larger CDBs
- c) SBP modifications necessary for IEEE P1394.1 environment (Serial Bus bridges)
- d) Isochronous facilities in SBP
- e) Bi-directional data transfer (new ORB type)
- f) DMA architecture for InfiniBand™

Two of the possible modifications to SBP were sufficiently minor and sufficiently agreed that little time was spent discussing them: b) explicit description of how to transport 16-byte or larger CDBs and e) bi-directional data transfer (new ORB type). In fact, the first requires no modification to SBP-2, just a better informative description. The second was agreed in order to harmonize SBP-2 with anticipated changes to SAM.

The remainder of the topics deserve more extensive explication. When an accompanying T10 document is available, its number is indicated; these documents may be downloaded from <http://www.t10.org>.

### **Reduced start-up latency from idle condition [Anderson / 00-328r0]**

Eric Anderson discussed experience gained by Apple Computer in their use of SBP-2 for mass storage (disks). SBP-2 exhibits desirable behavior under heavy workloads: the more a target is kept occupied with work and the less often it becomes idle, the greater the opportunities for overlap between ORB fetches and the flow of data to and from the medium. There is also little or no task initiation latency when new work is presented to the target.

In contrast, when a target becomes idle there is significant startup latency to commence a new task. Eric observed that serialized access to mass storage devices common to many desktop environments results in execution of a single task by the SBP-2 target before it becomes idle.

His proposal, 00-328r0, suggested a method by which all of the information necessary to commence a new task could be communicated to an idle target within a single Serial Bus block write request. Significant improvement in efficiency would result.

There was consensus within the study group that this proposed extension to SBP-2 merits the creation of a standards project.

### **SBP modifications necessary for IEEE P1394.1 environment [Johansson / 00-330r0]**

Peter Johansson presented a brief tutorial, 00-330r0, on some of the ways in which the current protocol will not be adequate to operate across Serial Bus bridges specified by IEEE P1394.1, Draft Standard for Serial Bus Bridges.<sup>1</sup> Readers should consult the slides for details; the salient reasons are summarized below:

- a) Remote time-out. Remote requests that transit one or more bridges between the requester and responder cannot (by IEEE P1394.1 definition) complete within a local bus split time-out. Unless a bridge-aware device distinguishes between local and remote transaction requests and applies different time-out limits, correct operation cannot be guaranteed.
- b) Virtual node ID. Remote nodes are addressed by 16-bit virtual IDs whose properties differ from the familiar local node ID. Local node IDs are invalidated by a bus reset; virtual node IDs are not necessarily invalidated by a bus reset. A bridge-aware device must recognize messages transmitted by bridge portals in order to correctly use virtual node IDs.
- c) Synchronization between initiator and target. Existing SBP-2 devices rely upon bus reset to synchronize state between themselves. This is reliable because bus reset is simultaneously observed by all devices on a single, local buses. When initiator and target are on different buses, separated by one or more bridges, not all bus resets will be mutually observed and none will be observed simultaneously. To the extent that SBP-2 relies upon bus reset, modifications must be devised to permit correct operation in the presence of bridges.

There was consensus within the study group that this work should be part of the new standards activity.

### **Isochronous facilities in SBP**

No formal presentation was made to the study group. Peter Johansson recounted some of the history of isochronous facilities in SBP-2. They had been present in Revision 3, the subject of the first T10 letter ballot on SBP-2 and had been deemed substantially complete and correct by the technical committee at the time. Technical changes were made to the draft as a result of ballot comments and Revision 3e contained those changes. However, at the March 19, 1998 T10 plenary a motion was made and approved to remove all isochronous facilities from SBP-2.<sup>2</sup>

The focal point of the discussion was, "*Should isochronous facilities be reintroduced into SBP?*". John Fuller voiced the opinion that if they had remained in the approved standard, they would be in use today. He was less sanguine as to whether or not isochronous facilities

<sup>1</sup> Working group documents may be obtained from <http://grouper.ieee.org/groups/1394/1>. A user name and password are necessary to retrieve the draft standards; they are "P1394" and "High-Perf", respectively (both are case sensitive).

<sup>2</sup> The most recent SBP-2 draft that retained all isochronous facilities specified by the working group may be downloaded from <http://www.t10.org/ftp/t10/drafts/sbp2/sbp2r03c.pdf>.

would be widely used if reintroduced now. In the two years since the approval of SBP-2, experience has been gained with "AV disks", mass storage devices that implement the 1394 Trade Association AV/C command set and permit isochronous streams to be recorded to and played back from disk media. Several mass storage vendors supply these AV/C disks to customers in the consumer electronics market.

Eric Anderson indicated that isochronous facilities in SBP-2 might be complementary to those already in AV/C disks. The video editing supported by applications under Mac OS relies on asynchronous access to the material to perform the edits, but obtains the material from an isochronous stream and ultimately places it back on the medium via an isochronous stream. If it were possible to directly create the image of an isochronous stream on the medium with asynchronous access, the time required would be greatly reduced.

Another area of potential interest for isochronous SBP-2 is with other peripherals, such as scanners, printers, FAXes and multifunction devices. Some vendors already combine two functional blocks within their products, such as a scanner and a printer, and use a dedicated Serial Bus for isochronous data transfer. If open and standard isochronous facilities were specified in SBP-2, additional product innovation would be possible. The principal motivation for this class of product is the reduction of buffer space made possible by isochronous data delivery at the device's native speed. An isochronous printer has been demonstrated to yield more than 10 MB in memory savings.

Whatever constituency exists for isochronous SBP-2, it was generally agreed that new isochronous facilities in SBP-2 should be harmonized with the experience gained in the development of AV/C disks. This poses a challenge, in that most of the expertise pertinent to AV/C disks resides not in T10 but in groups such as the 1394 Trade Association. It is unlikely that their members would add T10 meetings to their already crowded calendars.

In order to explore better participation between these groups of experts, Peter Johansson will make a liaison report to the 1394 Trade Association at the upcoming quarterly meeting in Redmond, WA. The results of this SBP-2 study group will be presented and discussion will be made as to how to best encourage full participation. Those present at the SBP-2 study group meeting were receptive to the idea of locating working group meetings at times and places amenable to the consumer electronics participants.

### **DMA architecture for InfiniBand™ [Haydt / 00-329r0]**

This topic is not directly related to SBP-2; instead it explores the suitability of SBP-2 as a prototype for the development of a similar transport protocol for InfiniBand. Rob Haydt, who made the presentation, explained that Microsoft is interested in leveraging the experience gained with SBP-2 and applying it to mass storage devices connected via InfiniBand.

The presentation slides, 00-329r0, investigate the similarities and differences between IEEE 1394 and InfiniBand and suggest that there are enough similarities to make it advantageous to adapt SBP-2 to InfiniBand. Such a protocol might be called InfiniBand Protocol (IBP), although this name was not suggested during the meeting. For the sake of brevity, IBP is used in this summary report to represent a SAM-compliant transport protocol for SCSI commands over InfiniBand.

Some of the study group participants observed that a similar protocol, SCSI VI Protocol (SVP), is under development by T10. When SVP is combined with a mapping of InfiniBand facilities to VI constructs, the resultant protocol is functionally equivalent to the proposed IBP. Rob said that Microsoft's analysis is that the VI architecture seems too cumbersome for an efficient foundation for mass storage on InfiniBand.

There was further discussion about T10 practices with respect to overlapping standards. Peter Johansson said that he believed it was not T10 practice to restrict the development of overlapping standards: so long as adequate resources are committed to a project it may pursue its own destiny in the marketplace. In any event, if work is to proceed within T10 on a proposed IBP standard, the next step would be for someone to bring forward a project proposal at a T10 plenary.

## Conclusions

The overall consensus of the study group was that new Serial Bus Protocol standards activity should be initiated and should address all of the topics mentioned in this report. There was also agreement that it was premature to discuss the venue for the work until after the October meeting of the 1394 Trade Association. The comments and suggestions from the consumer electronics participants might have significant influence on any T10 decision.

Because a new work item proposal contingent upon discussions October 27 in Redmond, WA, would not meet the two-week rule for consideration at the T10 plenary November 2, it may be unrealistic to expect action prior to the January 18, 2001 T10 plenary. However, this conflicts with the 1394 Trade Association meeting in January; I may ask the plenary's indulgence to consider, in November, a new work item proposal with less than two week advance publication.

## Attendance

Eric Anderson	Apple	EWA@Apple.com
Robert Elliott	Compaq	Rob_Elliott@CompuServe.com
Mark Evans	Quantum	Mark.Evans@Quantum.com
Firooz Farhoomand	Panasonic	FiroozF@IX.NetCom.com
John Fuller	Sony	John.Fuller@AM.Sony.com
Edward Gardner	Ophidian Designs	EAG@Ophidian.com
Rob Haydt	Microsoft	RobHay@Microsoft.com
David Hunter	Sony	David.Hunter@AM.Sony.com
Peter Johansson	Congruent Software	PJohansson@ACM.org
Randy Lawson	Texas Instruments	Rlawson@TI.com
Gene Milligan	Seagate	Gene_Milligan@Seagate.com
Chuck Rice	Hewlett Packard	Chuck_Rice@HP.com
Ron Roberts	Adaptec	Ron_Roberts@Adaptec.com