DocNum: T10/01-139r0
Author: Eric Anderson
Title: Minutes of SBP-3 Working Group
ftp://ftp.t10.org/t10/document.01/01-139r0.pdf

Minutes of the SBP-3 Working Group meeting, April 26-27, 2001.
Marriott City Center Hotel, “Rogue” room, Portland, Oregon

Attendees:

Eric Anderson	Apple	ewa@apple.com
Firooz Farhoomand	Panasonic	firoozf@ix.netcom.com
Lee Farrell	Canon	lfarrell@cissc.canon.com
Lance Flake	Maxtor	lance_flake@maxtor.com
John Fuller	Sony	jfuller@computer.org
Andy Green	Oxford Semiconductor	andy.green@oxsemi.com
Paul Grunwald	Philips	paul.grunwald@philips.com
David Hunter	Sony	david.hunter@am.sony.com
Peter Johansson	Congruent Software	PJohansson@ACM.org
Walt Jones	Microsoft	waltj@microsoft.com
Rob Lash	Maxtor	rob.lash@quantum.com
Chuck Rice	HP	chuck_rice@hp.com
Shigeru Ueda	Canon	ueda.shigeru@canon.co.jp

The following agenda was presented by Johansson. In the minutes that follow, the start of discussion of items listed below is denoted by the index number listed within square brackets, such as [4.1]. Note that these references do not always appear in order, and may not signify the conclusion of discussion of a previous agenda item.

1. Introductions and procedures
   1.1 T10 Membership and voting
   1.2 Document naming conventions
   1.3 Two-week rule
   1.4 Meeting fees
   1.5 Approval of prior minutes
2. Call for patents
3. Informal liaison
   3.1 IEEE P1394.1 [Johansson]
   3.2 IEEE P1394.3 [Johansson]

3.3 1394 Printer Working Group [Farrell]
3.4 1394 Trade Association [Fuller] (delete from future agendas)
4. Prior action items
4.1 Modify draft in accordance with Project Proposal item e) [Johansson]
4.2 P1212 Revision entry in unit directory [Johansson]
4.3 Configuration ROM (feature control, instance directories) [Johansson]
4.4 Incorporate 01-070r0 into draft [Johansson]
4.5 Request AV/C expert to review Annex H [Hunter]
4.6 Request AV/C expert to define track metadata [Hunter]
5. Review of changes in working draft
6. Old business
6.1 Divergence between SBP-2 and PPDT [Johansson]
6.2 RBC Access For AV/C Data Interchange in SBP-3 [Flake]
7. New business
7.1 Isochronous charter
7.2 Stream command block ORB [Johansson]
7.3 Isochronous data format [Green]
7.4 Bi-directional ORB [Johansson]
7.5 16-byte CDBs [Johansson]
7.6 Data buffer and page table in different nodes [Anderson]
8. Meeting schedule
June 5 - 6 (Chicago or Toronto)
July 17 - 18 (Colorado Springs, CO)
August 22 - 23 (Cupertino, CA)
October 2 - 5 (east coast)
November 6 - 7 (Monterey, CA)
December (location not yet determined)
9. Review of action items
9.1 Add configuration ROM examples for instance directories [Johansson]
9.2 Modify GET NODE HANDLE to permit release [Johansson]
9.3 Operational description of login (bus reset) [Johansson]
9.4 Remove EUI-64 from login request [Johansson]
10. Adjournment

[1] Johansson called the meeting to order and added several items of New Business to the agenda, as reflected above.

[1.2] Johansson commented that an automated system had been created for the allocation of T10 document numbers, and the subsequent submission of documents for posting. Anderson noted that he had used the system to obtain the document number for these minutes:

http://www.t10.org/members/ad.htm

[1.3] Johansson briefly reviewed the two-week rule, explaining that it did not prevent the discussion of documents posted less than two weeks before a meeting.

[2] Johansson reviewed general T10 policies and procedures. In general, attendance and participation at T10 ad hoc meetings (such as this one) is open to both visitors and T10 members. When formal votes are taken, either in an ad hoc meeting or in the T10 plenary, one vote is permitted each organization, to be cast by its principal representative or designated alternative. A two-week rule is in effect: No matter may be voted on unless notice was given at least two weeks prior. Documents to be voted on must have been posted two weeks prior to the vote. The two-week rule can be waived if nobody objects. Announcements of new documents and meetings must be posted to the T10 email reflector; all other business can be conducted on the working group reflector.

Reflector information (SBP3@isg.apple.com):

Subscribing:
email requests@isg.apple.com w/subject "subscribe sbp3"

Help?:
email requests@isg.apple.com w/subject "help"

The following paragraph about ANSI/T10 patent policy is copied from past T10 Plenary minutes:

A document is available from ANSI, "Procedures for the Development and Coordination of American National Standards", at no charge. This document is also on the web at http://www.ncits.org/help/ansi_sdo.html. Section 1.2.11 contains the ANSI patent policy. Amy Marasco manages patent issues for ANSI and can be contacted at amarasco@ansi.org or 212-642-4954. Gene Milligan prepared a useful “Handy dandy Technical Committee's Patents Guide“, which is available at ftp://ftp.t10.org/t10/document.99/99-291r0.pdf.

[3.1] Johansson noted that the 1394.1 ballot pool is still forming, and is open until May 5.

[3.2] Johansson commented that the Ballot Review Committee for 1394.3 (also known as Peer to Peer Data Transport (PPDT), a transport layer based on SBP-2) had not yet convened.

[3.3] Farrell reported that no 1394-related activity had recent happened in the Printer Working Group (PWG), and added that the group is awaiting the completion of 1394.3.

[3.4] Fuller stated that the 1394 Trade Association had recently met in Singapore, and had done nothing in particular pertaining to SBP-3.

Johansson led a review of old business, as follows.

[4.1] Johansson noted that item e), Bidirectional ORBs, had not been incorporated into the SBP-3 draft, but that a presentation on this topic would be made later in the meeting.

[4.2] Johansson said that his action item to discuss a Revision entry in the Unit Directory with the P1212r Ballot Review Committee had not been completed, but that topic was now being discussed on the 1212 reflector. Johansson noted that as a fallback plan, SBP-3 could define its own revision key if 1212 did not standardize a suitable one.

[4.3] Johansson reported that he had documented the use of Instance Directories in the latest draft of SBP-3, and noted that the draft recommends their use. The related action item to perform editorial cleanup regarding the applicability of keys in Logical Unit Directories and Unit Directories might still be incomplete.

Action item: [9.1] Johansson to add examples of Instance Directories to the SBP-3 draft.

[1.5] Hunter moved to approve the minutes of SBP-3 from March 6-7, 2001:

ftp://ftp.t10.org/t10/document.01/01-101r0.pdf

Fuller seconded.
Motion passed with none opposed.

[5] Johansson lead a review of the latest draft of SBP-3, revision 1b:

ftp://ftp.t10.org/t10/drafts/sbp3/sbp3r01b.pdf.

Johansson noted that section 4.7 now incorporates text from document 01-070 regarding Node Handles as discussed at the previous meeting:

ftp://ftp.t10.org/t10/document.01/01-070r0.pdf

Johansson remarked that additional text would probably be needed before bridges were fully accommodated.

Johansson showed that in section 5.1.4 (Table 2), a new function value had been defined for GET NODE HANDLE, which is further defined in section 5.1.4.5. Johansson observed that how a Node Handle is deallocated had not been clearly defined, though a Logout would probably deallocate Node Handles. Fuller endorsed having a RELEASE NODE HANDLE function.

Action item: [9.2] Johansson to add a RELEASE NODE HANDLE directly to the draft, possibly replaced by some kind of unified MANAGE NODE HANDLE operation to conserve function codes.

Johansson reviewed Figure 22 (Login ORB), which had been modified to support node handles. An EUI 64 value can be placed in the password field when using a bridge-aware login, to ensure that the target establishes a login with the correct initiator node. A new "aware" field in the fifth quadlet indicates the initiator's bridge awareness.

The group discussed the impact of placing the EUI 64 value in the two quadlets presently used for a direct or indirect password, and how this would affect the ability to use passwords. Fuller noted that if the EUI 64 was moved into a new structure that also contained the password, and was referenced by a pointer in the password field, that the association of the EUI 64 with the login request might no longer be infallible. Fuller asked if there was really a problem that needed to be solved in associating the login with a particular EUI 64.

In further discussion, the group agreed that additional text describing how to react to a bus reset during the login process would be valuable for traditional bridge-unaware implementations, as well as in a new bridge-aware environment, and that in either case additional guidance would help implementors to create compliant products. Anderson noted that legacy initiators would probably assume that a login interrupted by a bus reset was invalid, because no response was received. Johansson and Anderson agreed that it was desirable for bridge-aware devices to be able to complete a login despite a bus reset. Fuller commented that a Net Generation change might be fatal to the login process, because the target and initiator could have different views of what had happened.

Anderson favored adding text to define situations in which the initiator might have to take action to recover from a loss of synchronization with the target during the login process, rather than trying to add text to automate such recovery work in the target, because with both sides working autonomously it would be very hard to prove that recovery would always work. Anderson added that if such complications during login are rare, it would be reasonable to burden the initiator

with additional work and/or delays. Fuller noted as an example that the initiator could be required to use the Query Logins function to clear up certain ambiguous situations.

Action item: [9.3] Johansson to prepare a new document with text clarifying how to react to a bus reset during login in a legacy target and initiator, as well as how bridge-aware devices can tolerate bus resets and recover from Net Generation changes during login.

[9.4] Johansson noted that he would remove the EUI 64 value from the Login ORB for now, and reduce the "aware" field from two bits to one, pending further resolution of how the login process would work in a bridge-aware device.

Johansson showed that Figure 23 (Login response) had been updated to return a node handle in the fourth quadlet, if an "aware" login ORB had been issued.

Johansson remarked that the new ORB defined in section 5.1.4.5 for obtaining a Node Handle was only needed if a third-party node (neither the target nor the initiator) is participating in the SBP-3 protocol, such as by holding data buffers accessed by the target. Johansson added that the ORB includes a hint for efficient operation. Anderson asked if text existed to require a target to successfully complete a get node handle ORB even if the hint was missing or invalid. Johansson noted that such text would belong in an operations section such as chapter 8, which had not yet been updated.

Grunwald asked if the reconnect process could be used by an initiator that was uncertain if it was logged in, as an efficient way to resolve such uncertainty. Anderson suggested that the draft does not clearly say if a target should accept and process a reconnect from an initiator that the target thinks is already logged in (and not subject to a pending logout due to a bus reset).

Fuller suggested that a reconnect operation should be allowed to change the Status FIFO address, and added that this would be useful when a BIOS hands over control to an OS, because the initiator wants to keep its login, but the BIOS and the OS may be unable to agree on the memory address of the Status FIFO. Anderson noted that Status FIFO address already present in the reconnect ORB could be used to return status for the reconnect, and could then become the normal Status FIFO address for the login, although making this change without a corresponding enable bit would break legacy initiators. Someone noted that another option was to allow a "replacement" login to overlay an existing login, but this solution could be problematic if SBP allows multiple logins from one initiator to one target, because it was unclear how to distinguish between a new and a "replacement" login. Anderson preferred an option to specify a new Status FIFO for command ORBs using reconnect ORB fields that Johansson noted were

presently reserved. Anderson added that for this feature to be useful, the FIFO replacement must be allowed at times other than immediately after a bus reset, but that it could be required that the target be in a state such as Suspended or Idle in order to change the FIFO address.

Fuller said that he would prepare a document of drop-in changes to define how a modified reconnect ORB could be used to assign a new status FIFO and/or to reconfirm an existing login, and when this would be allowed.

Anderson noted a past proposal to allow multiple logins to a target from one initiator node, which is specifically prohibited by SBP-2, and observed that such multiple logins could be useful in a printer target that accepts multiple jobs simultaneously, which modern printers typically do. This topic was added to the agenda, but was not discussed again.

Johansson noted that (unlike a login ORB) a reconnect ORB had plenty of reserved space where an EUI 64 could be specified in addition to a replacement Status FIFO address, and added that including the EUI 64 might be helpful in avoiding confusion about who is who after a bus reset.

Johansson observed that Figure 45 (Configuration ROM hierarchy) had been updated to show Instance Directories. Anderson commented that SBP-2 had not specified the use of Instance Directories, but added that SBP-2 had not prevented devices from having Instance Directories either, and noted that some contemporary SBP-2 products did include Instance Directories. Anderson felt that Figure 45 could be incorrectly interpreted to imply that SBP-2 devices could not have Instance Directories. Johansson agreed to revise the figure and text accordingly. The group made some changes to the final paragraph in section 7 to clarify this point.

Minor changes were made to section 9.1.4 (Fast Start Register) regarding the difference between a packet that was actually sent and what could have been sent in a packet.

Johansson showed changes made to figure 66 and the accompanying text regarding the fetch agent state machine.

After reviewing section E.2, Fuller commented that a target could be required to ignore fast-start packets (and ORB_POINTER writes) when it is not in the SUSPENDED or RESET state, to make it easier to write safe fast-start code on multi-processor system, where code might occasionally try to send a fast-start packet even when a target is already active. Fuller felt it would be cheaper and easier to solve this problem in target hardware than to make software do it correctly. Johansson noted such a guarantee would be hard to use, because a

failed (ignored) fast start packet isn't automatically linked into the active list of ORBs, so software would have to detect the failure and recover. After further discussion, Jones and Anderson felt an SMP OS could probably use fast start as is without holding spin locks for long periods, though both felt that more study was needed.

[4.4] Johansson noted that document 01-070r0 had been incorporate into the draft, and reviewed as part of agenda item [5].

[4.5 and 4.6] Hunter reported no progress on the AV/C action items, which were carried over for discussion next time.

[6.2] Flake described updates he made to the 3/23 version of document 01-067:

ftp://ftp.t10.org/t10/document.01/01-067r1.pdf

Flake noted that the group had previously discussed adding metadata, but that on further review he had decided not to incorporate metadata in his proposal. Flake explained that he had simplified the proposed object management commands so that they could access only RBC objects only, leaving existing AV/C mechanisms as the only way to access information about AV tracks. Fuller and Johansson expressed concern that not having a single service to list all objects represented a step backwards. Anderson noted that using two independent techniques could be difficult because no reliable way existed to synchronize their access or to correlate data between them. Flake commented that the object ID concept is used in both spaces and would prevent confusion about objects visible in both spaces. Fuller felt that both AV/C and RBC types of access should be able to fully enumerate the tracks or objects on a device. Lash said that the existing AV/C Disk subunit model provides a descriptor-based list of tracks. Nobody was certain how AV/C Disk would deal with a foreign or non-AV formatted track, or how a descriptor could express the presence of such a track.

Fuller and Johansson both said that using AV/C commands (perhaps transported by SBP) to list the full contents of a drive might be sufficient, if AV/C descriptors are able to adequately describe non-AV tracks.

Johansson asked the group to agree on a term to describe both AV tracks and RBC partitions (e.g. computer filesystem data). Most present agreed on "extent", defined as "an addressably contiguous area allocated on the drive". Johansson noted that Extents include AV tracks and RBC-acessible partitions. Johansson said he would like to be able to create and delete extents using RBC commands, possibly including the ability to create extents that are AV tracks. Johansson added that to create a first-class AV extent, it might be necessary to express attributes regarding performance when creating the extent. Fuller suggested that

there should be a common naming system for all extents, regardless of how they are accessed.

Lash explained that the AV/C create track command would set aside a fixed amount of space, but added that an AV/C Disk subunit device could also be asked to record without a predetermined bound (other than available space) and it would allocate media as needed. Johansson asked if data requirements could be characterized when a track was created using AV/C. Lash said yes, but noted that the requirements might be implicitly derived from plug control registers or other indirect sources, rather than being explicitly specified. Lash suggested that an app note could be written advising one how to use the existing AV/C Disk design to specify extent properties. Johansson noted that AV data on removable media such as MiniDisc or DVD-RAM should be stored in an open standard format so that vendors can achieve media and device interoperability.

Anderson stated that the existence of protected extents should be visible through all access methods, provided that this is done in a way that does not compromise the security of the protected data. Anderson explained that denying the existence of an extent just because it contains protected data leads to user-interface problems; for example, the user may mistakenly conclude that their data has been lost.

Johansson identified the following key issues to resolve:

How to create extents

How to specify metadata, even if its contents or format aren't defined by SBP

How to return an extent ID and native name when creating an extent

Johansson explained that a "native name" would be, for example, the identifier that is specified by AV/C Disk for a AV extent, as might be referenced by an AV/C "play" command, while an "extent ID" would be a more generic unique identifier not bound to a specific protocol or model (such as AV/C).

Fuller stated that a desired goal is the ability to lay down an AV track asynchronously, perhaps faster than real-time, and then play the track isochronously using only AV/C commands. Rice noted that some devices need media access that is even slower than real time, such as to print a single video frame, and noted that asynchronous access would be helpful for such devices too.

[7.1] Johansson led the group in a discussion of what the goal(s) for Isochronous capabilities in SBP-3 devices should be.

Lash noted that the AV/C Disk model required that the disk manage its own filesystem, whereas on a disk controlled by RBC, the filesystem would be imposed and managed by a computer or other external controller. Lash observed that this difference leads to the partition problem, of how to divide a disk between AV/C and RBC uses.

Anderson said that traditional computer filesystem can't handle dynamic partition sizes, and added that Apple had no immediate plans to change that. Jones suggested that Microsoft had no such plans either. Anderson continued that as a consequence, Apple would be satisfied with a solution in which a partition exists between AVC and RBC spaces, where moving the partition would require erasing at least the RBC portion of the disk. It was noted that this could be accomplished by using the AVC Disk model and designating one or more "AV" tracks as RBC partitions; though the disk might be capable of growing or shrinking those partitions, the computer would never request such a change.

Anderson expressed a desire to read and write first-class AV/C Disk tracks using fully asynchronous access, free of the AV/C command set and FCP transport, provided that doing so could never compromise the security of content-protected data. Anderson added that Apple also wanted the ability to transmit and receive isochronously from or to disk media fully managed by a computer file system, under the control of a computer. Fuller said that Sony also wants at least asynchronous read and write of AV/C Disk tracks.

Anderson explained that computer and consumer products have different cost/performance requirements. Anderson said that to be used inside a computer, disk drives need to be accessible in a computer-native way for high performance. Anderson added that to be useful in a consumer device, drives need to behave like the AV/C Disk model so that they can interoperate with other AV devices, and services such as HAVi. Anderson stated that these two scenarios have substantially different requirements and cannot be unified, but a disk that can do one has most of the smarts and hardware resources needed to do the other, so putting both capabilities in one disk leads to a product with much broader usability for nearly the same cost as making a product for just one scenario or the other.

Johansson commented that some devices such as "DCAM" cameras use a simple isochronous format where plain capture and playback is feasible by a disk with knowledge only of 1394, and not of standards such as AV/C and ISO/IEC 61883. Johansson explained that unlike DV, where CIP headers must be adjusted for retransmission, DCAM can be retransmitted without any knowledge beyond that specified by 1394. Johansson felt that the group should define SBP-3 such that transports such as DCAM that stick with the basic services offered by

1394 could be recorded and retransmitted, and added that by providing such a service, other groups might be motivated to make better use of the basic services offered by 1394, instead of layering on protocol-specific complications.

A lengthy discussion of a common on-media data format ensued. Johansson and Anderson advocated that cycle marks must be visible "in band" so that a computer could asynchronously write a packet stream on media that could be played back correctly (on the right cycles), and so a computer can understand a recorded packet stream in which a packet was not sent on every cycle. Lash and Flake responded that cycle marks must only be stored "out of band", where the disk can see them but asynchronous access cannot, however after extensive discussion both agreed that asynchronous access to "in-band" cycle marks might be desirable. Anderson gave an example of recording a MIDI performance in which most cycle are empty, and no timestamps are embedded, yet a simple playback at 8000 cycle marks per second would faithfully reproduce the stream. Anderson added that even though some formats have been invented that embed timestamps, their existence doesn't prove that all isochronous formats must have embedded timestamps. Anderson repeated Johansson's earlier suggestion that defining a visible cycle mark system that works with plain 1394 and requires no additional embedded structure may enable simpler, lower-cost isochronous protocols than the CIP-based ones that are common today.

Anderson noted that isochronous printing is desirable even though (buffer) memory prices have plunged, because guaranteed delivery allows one to reduce latency in printing, allowing motors to start sooner, because timely data arrival is guaranteed. Anderson added that isochronous delivery is also desirable because asynchronous delivery performance is not guaranteed; devices like streaming tape drives or scanners that are capable of backing up their mechanism to handle transport slowness can give much better overall performance if isochronous (guaranteed) delivery is used so that time-wasting mechanical backups are avoided.

[7.2] Johansson reviewed proposed changes to the stream command block ORB to accommodate a compact "play list" for isochronous data:

ftp://ftp.t10.org/t10/document.01/01-137r0.pdf

Johansson explained that the new proposal was very efficient when fine-grained data is being specified for isochronous transport, because the scatter/gather list is now contiguous in memory and can be efficiently fetched by the target with many fewer read requests. Johansson added that the resource demands on the initiator are also reduced because fewer ORBs are needed.

Anderson asked if a disk could do isochronous playback without cycle marks, by looking at the 1394 isochronous packet headers on the media, and playing one packet per cycle. Johansson answered that this would probably work, and confirmed that the recording format does include packet headers.

The group decided to consider the proposal further before incorporating it into the draft. Johansson said he would call for a vote by email, to motivate feedback and discussion.

Johansson noted that he plans to add text to the draft to allow fast start to be used with stream ORBs. Johansson observed that there's no direct way for a target to indicate that it supports fast start for asynchronous transfer but not for isochronous transfer, although a target could treat unwanted (or unsupported) fast start writes as if they were writes to ORB_POINTER. Discussion did not lead to a conclusion as to how useful fast start would really be with streams.

[7.3] Green led a lengthy discussion of interchange formats, particularly as might apply to a disk that had detailed protocol and format knowledge of one or two selected media types and transports, such as DV.

Johansson showed figure 10 (Stream engine block diagram) and reviewed the basics of what it illustrated. Lash noted that if the stream control agent is bottled up with pending "at time X" commands, and then the initiator wants to insert an immediate command such as "stop" or "pause", it's unclear how to get that new command to the head of the queue. Johansson agreed, and noted that "abort task set" might work, but would probably be a poor way to do this. Further discussion, however, revealed that abort task set might be a reasonable way to deal with this situation.

Johansson proposed that the figure (10) and model be simplified to a single fetch/execution agent by defining simple modes that the stream engine could be in.

[7.4] Johansson explained that other T10 groups have defined bidirectional I/O, and gave a presentation on how SBP-3 might offer such a service:

ftp://ftp.t10.org/t10/document.01/01-138r0.pdf

Jones, Fuller, and Anderson favored the expanded ORB with two descriptors, instead of double indirection. Jones and Anderson agreed with Johansson that the new ORB should be flexible, using direction bits in each data descriptor, with targets required to find the descriptor with the needed direction (in or out) by examining the direction bits. Green asked if the new two-buffer ORB would be standard for SBP-3 targets, and would SBP-3 targets accept only the new kind of

ORB? Johansson and Anderson agreed that the SBP-2 ORB format would remain available in the standard, and it would be up to targets to decide which one (or two) ORB formats to accept.

Flake asked what kind of command would use bidirectional data. Johansson explained that some RAID operations were anticipated to use this. Fuller noted that the 1394.3 effort had previously requested bidirectional ORBs, though they later designed a solution based on unidirectional ORBs.

Johansson noted that the dual-buffer ORB could work with fast start, though placing both page tables into a fast start packet would be awkward. Johansson recommended placing page table entries for at most one of the two descriptors into the fast start packet.

Johansson said he plans to write up more details regarding dual-buffer ORBs.

[7.5] Johansson noted that Annex B could be interpreted as to imply that ORBs transporting SCSI commands were truly limited to 32 bytes in total size, limiting the CDB to 12 bytes. Johansson proposed to update Annex B to clearly show ORBs with (at least) 16-byte CDBs. Green noted that an 18-byte CDB had recently been defined.

[7.6] Anderson led an ad hoc discussion about ORBs with page tables being able to describe data buffers in a node other than the node holding the page table. Anderson noted that SBP presently requires that if a page table is used, all data buffers described by it are required to exist in the same node as the page table, because the page table does not individually specify the top 16 bits of each data pointer. Anderson reported that a developer has found a possible application of SBP in which an interloper target processes ORBs and reissues them to one or more other targets on a bus. As part of this reprocessing, page tables are written (or rewritten), and whole new ORBs may be generated. The interloper may be willing to provide memory for these structures, but it is unable to use page tables to describe the original data buffers because it cannot allocate memory in the node where the original data buffer is located. It may be impossible to break up the original ORBs such that page tables are no longer needed.

Anderson said that the developer had requested the ability to address data in one node using a page table stored in another node. Anderson observed that in addition to the developer's intended application, the general ability to address multiple nodes could be useful for performing arbitrary scatter/gather I/O using SBP, such as in a distributed multiprocessing system that used 1394 as its interconnect.

Anderson sketched out two ways to support the developer's request. One, called 3PT, stretched out page table entries from two quadlets to three quadlets, containing a 32-bit length entry and a 64-bit data pointer. The other, called Header, added a single 32-bit header to the entire page table to contain 16 bits specifying the node ID (and bus ID) for all data pointers in the page table.

Fuller commented that the interloper capability was compelling, even though it was not the same as the problem of supporting 1394.1 bus bridges. Fuller added that the multiprocessing example was not compelling to him. Flake said that he favored supporting fully-qualified 64-bit addresses and had never really liked specifying addresses by just their low 48 bits, as SBP presently does. Johansson suggested that the existing page table structure could be preserved, and additional information could be stored in a new intermediate structure, pointed to by an ORB, which in turn pointed to a conventional page table. Johansson said he would prefer such an indirection technique. Fuller said he favored using 3PT or Header instead of Johansson's indirection technique, and noted that indirection would cause complications for Fast Start. Anderson commented that when performing large I/Os farmed out to multiple targets, Fast Start might not be important because the I/O would take a long time to complete anyway. Johansson noted that details regarding transfer speed, payload size, and other transfer attributes described in an ORB would become more complicated with additional nodes participating in the execution of an ORB.

Anderson said he would take the group's discussion back to the developer who had proposed the idea for further consideration.

Grunwald asked if anything about SBP or the proposed changes to SBP prevented the use of the "5C" or "DTLA" content protection mechanisms, and Johansson replied that nothing was known to prevent such use. Fuller noted that the details of 5C/DTLA operation are not easily accessible, so few people could know for sure if any problem would exist.

In general discussion, the group realized that an earlier assertion that all asynchronous access to content-protected media should be disallowed was overly simplistic, and that carefully controlled access to content-protected data through asynchronous transfer might be required in order to deliver innovative services in future applications of SBP. An example given was to securely receive a content-protected movie from an asynchronous communication media such as the Internet, and record it onto an SBP-3 disk in a secure way that could later be played back isochronously while preserving the content protection. It was noted that the 5C/DTLA mechanism and rules as written today might have to bend somewhat to permit such an application, even though all copyrights would be respected in such use.

[8] Johansson noted that the next meeting would probably be June 5-6 in Chicago or Toronto.

In continued discussion, Anderson reviewed the dual goals of SBP-driven first-class access to AV data in a disk-managed track that is fully interoperable with AVC devices, as well as SBP-driven isochronous transmission and reception that is not AVC-disk compatible, but which can be 6-1883 compatible if the computer manages plug control registers and other bus resources so that a third node can correspondingly send or receive the isochronous data. Anderson also noted that in his depiction of a partitioned disk, the RBC partition is really just a track managed by the disk like any of the AV tracks, but the fact that the RBC partition is an AVC track is invisible through legacy RBC access. In particular, the track size is fixed, and LBA-based addressing works as it would in an SBP-2 RBC disk drive, so legacy software is unaware of (and unaffected by) the fact that the RBC partition is "managed" to some extent by the disk.

Further discussion involved the creation of an "activate" or "prepare" command that could be part of RBC-2, and transported by SBP. This command would cause an AV track that was previously prepared and populated by (possibly new) RBC commands to become "live", enabling it to act as a first-class AVC track, eligible for real-time playback or recording under entirely AVC-driven control. A related "offline" or "deactivate" command could remove a track from AVC's "world" so that it could be modified by RBC access, provided that no content protection mechanisms could be defeated in the process.

[10] Flake moved to adjourn.