DocNum: T10/01-060R0
Author: Eric Anderson
Title: Minutes of SBP-3 Working Group
ftp://ftp.t10.org/t10/document.01/01-060r0.pdf

Minutes of the SBP-3 Working Group meeting, January 24-25, 2001.
Outrigger Hotel, Waikoloa Beach, Hawaii

Attendees:

Eric Anderson	Apple	ewa@apple.com
Bruce Fairman	Sony	bruce.fairman@am.sony.com
Lance Flake	Maxtor	lance_flake@maxtor.com
John Fuller	Sony	jfuller@computer.org
David Hunter	Sony	david.hunter@am.sony.com
Peter Johansson	Congruent Software	PJohansson@ACM.org
Dan Knudson	Microsoft	dankn@microsoft.com
Ron Roberts	Adaptec	ron_roberts@adaptec.com
Bill Russell	Canon	bill.russell@cis.canon.com
Scott Smyers	Sony	scott.smyers@am.sony.com
Thomas Thaler	BridgeCo	thomas.thaler@bridgeco.net

Johansson called the meeting to order and 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"

Johansson noted that at least three official positions were required for the working group: Secretary, Project Leader, and Editor. Johansson volunteered to act as Project Leader and Editor. Anderson volunteered to act as Secretary. Hunter volunteered to act as backup secretary.

Johansson noted that every meeting would begin with a call for patents.

The following paragraph about ANSI/T10 patent policy is copied from recent 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.

Action: Johansson to check with chair of T10 to see if any patents were disclosed for SBP-2.

Action: All present requested to disclose any patents granted that apply to SBP-2 or proposed SBP-3 changes.

Johansson suggested an informal liaison to IEEE P1394.1, because the group would like to make SBP3 work in a bridged 1394 environment.

Fuller suggested a liaison with IEEE 1394.3 because that standard is built atop SBP-2. Fuller noted that Greg LeClair of Epson is the chair and might fill this role.

Russell suggested a liaison with the Printer Working Group and noted that Lee Farrell might fill this role.

Johansson asked Smyers if the 1394 Trade Association Audio Visual Working Group (of which Smyers is chair) would provide a liaison. Smyers said that he would like to fill this role himself, and added that Fuller or Hunter could stand in for him as needed, particularly during absences anticipated through July 2001.

Johansson reviewed the Project Proposal, including the eight items from the recommended scope.

Action: (scope item b) Johansson to revise the draft to clarify that ORBs can carry 16+ byte CDBs.

Action: (scope item e) Johansson to add the bidirectional ORB directly to the draft rather than write a proposal. Johansson suggested that use of the bidirectional ORB for unidirectional transfer be allowed.

Action: (scope items f and g) Johansson to incorporate 1394a changes and SBP-2 clarifications into the draft.

Most present were interested in isochronous and bridge compatibility work, with a slight preference to doing the isochronous work first.

Anderson presented document 01-057r0, the Fast Start proposal, and briefly explained it.

ftp://ftp.t10.org/t10/document.01/01-057r0.pdf

Fast Start enables an initiator to send to the target an ORB pointer, the associated ORB, and (optionally) part or all of the associated page table, in a single packet. The proposal improves start-up time for idle targets, with full backwards compatibility to SBP-2, and is completely optional in both the initiator and the target.

Anderson also showed slides from the following as background.

ftp://ftp.t10.org/t10/document.00/00-328r0.pdf

After discussion, Anderson volunteered to modify his proposal to include the following changes: 8 bits of quadlet length expressing the largest supported Fast Start packet, with 0x0 indicating max_rec, and 0xff being reserved. Rename the SBP_Features key to Fast_Start, keep it as 0x3E. Add a note about bidirectional ORB use with Fast Start being (yet) undefined. The field name fast_start_offset is desirable if it will fit in the figure.

Johansson led a review of the initial SBP-3 draft document, highlighting the changes from SBP-2.

ftp://ftp.t10.org/t10/drafts/sbp3/sbp3r01.pdf

Action: Johansson to mention 1394.3, MMC, and RBC in a new Bibliography section.

Action: Anderson to ask on reflector for any additional references.

A lengthy discussion ensued regarding the impact of changing the Unit Directory Version field (IEEE P1212r section 7.7.11; formerly known as Unit_SW_Version) in SBP-3 from the value 0x010483 used in SBP-2, as a way of indicating support for SBP-3. Anderson and Fuller observed that existing operating system support for SBP-2 would be broken by such a change. The possibility of using two version keys in a single Unit Directory was discussed, but the impact and precedent of doing so was unclear. The IEEE P1212r instance directory service was seen as capable of solving the problem, but also as perhaps needlessly complicated in this case.

Action: Johansson to describe a 1212-based solution to identifying backwards-compatible revisions of existing spec_ID/version tuples so that changing the Version ID in SBP-3 can be avoided.

Annex H was discussed, and no strong reason for its existence was evident. The group elected to delete it. [See below for a reversal of this decision]

Smyers talked about isochronous services and listed two areas of interest:

1. Isochronous Recording Format

Smyers felt that the draft format was probably OK, based on experiments at Sony, but he will study it further.

2. Asynchronous access to Isochronous data

Example: DV editing by a computer. Asynchronous access permits super-realtime copies, with full error detection, unlike isochronous transfer. Sony and Western Digital showed a demo at CES in which a rq_fmt value of 1 was used to indicate a special ORB where the first two bytes of the CDB indicate an AV/C track number (as defined in AV/C command set). The disk appears to an AV/C controller as a fully native AV hard disk, but appears to an SBP-2 initiator as a fully native SBP-2 device. A special SBP-2 initiator can access the AV track data using the special rq_fmt. Smyers noted that using LUNs to select tracks instead of embedded track numbers provides some access control, for example allowing one to deny access to a track that is already in use, but has other disadvantages. Though 16 bits of track selector were implemented for the demonstration, Smyers felt that 32 bits of track selector would be preferable.

A lengthy discussion of CDBs and where to put the track select information followed. The discussion included the merits of AV vs SBP access to track data and computer filesystem data, and how these two kinds of data can coexist on

one drive, as well as the level of interoperability that will or won't result from each of several proposed solutions.

Anderson proposed the following: A disk drive supporting both AV-based access and RBC-based access would have a single partition between RBC space and AV space. This partition can be moved by some presently undefined command; moving it will result in trashing at least the RBC side. Such a disk, if sold pre-installed in a computer such as a Macintosh, would come with a partitioning of 100% RBC and 0% AVC. The same disk, sold built into a TV might be partitioned for 90% AVC and 10% RBC. The ideal partitioning for such a disk when used as portable media by a consumer (not bolted into other equipment) is unknown and may need to be changed by the user depending on their desired usage model. One possibility in this case is to ship the disk as "undecided" and require using some software or 1394-enabled device to "format" it?.

Anderson proposed that computer users would not be allowed to store non-AV data (e.g. spreadsheet files) in the AV partition (if any) of such a disk, but could store AV data in the RBC partition using computer file formats such as QuickTime; consequently users would likely format their disks (built into their computer) for 100% RBC. Critically, a disk formatted in this way could only send and receive AV data if the computer was powered on in order to direct the transfer, because native AV devices would be unable to directly access the data. Computer users who want to use such a disk when it is not connected the computer, or when the computer is turned off, must format at least part of the disk as AV and accept that they cannot directly write to that part of the disk, except when writing AV data, perhaps with help from specific AV-related application software. Such a disk shows up as two units on 1394: an SBP/RBC unit, and an AVC unit.

Anderson concluded that such an arrangement met the needs of existing computer filesystems, while enabling the benefits of combining SBP and AV in a single drive, as long as the computer is turned on and present on the bus. Significantly, the user is not required to make a decision about how much space to allocate to RBC and AV. Anderson and some others present felt strongly that users would be dissatisfied if exposed to a partition between RBC and AV because one partition or the other would fill up first, and the free space in the other partition would not be usable as a supplement to the full partition.

Anderson noted that users with multiple disks can understand that each one represents a fixed amount of storage space that exists independently of other disks, but that a user with a single disk likely would not understand the concept of a partition that restricts their ability to store data on that disk.

Flake explained Maxtor's desire to ship a single product ("SKU"). Anderson responded that there are two usage models for such a disk: Computer storage and AV-accessible storage. The user must make a choice at some point: Either the user must choose among two SKUs, or the user must configure the drive to meet their needs. Configuration could be made semi-automatic, based on how the drive is used, but that may be confusing. Opinions varied as to how auto-configuration might be done.

Johansson asked if it was awkward to send AV/C commands over FCP to the disk intermixed with RBC commands transported by SBP. Anderson agreed that it was awkward but noted that having multiple ways to do the same thing (i.e. transport AV/C commands) was also awkward. Anderson added that with two ways to send the same command there may be many interlock problems to resolve. Johansson asked if transporting AV/C over SBP would be a good compromise? Anderson agreed, noting that though this created two ways to send AV/C commands to the same device, it allowed computer software to use a single transport, SBP, rather than interleaving SBP and AV/C, and that a single transport was highly desirable (as was avoiding FCP whenever possible). The group agreed to restore the AV/C-over-SBP transport annex to the draft.

It was proposed that a new RBC command will be defined to be able to set the RBP/AVC partition point; but that no AVC command will be able do this. It was also suggested for safety that no AVC command be able to delete the RBC partition (e.g. reallocate it as AV space). No decision was reached regarding auto-configuration; Vendors can implement auto-config based on market needs or SBP-2 compatibility desires.

The meeting continued on Thursday, January 25, with all of Wednesday's attendees still present except for Thaler.

Anderson mentioned the cycle mark feature and asked if it could be optionally left out? A lengthy discussion of cycle mark and cycle index mark followed, in which Smyers was strongly opposed to supporting the cycle index mark. Anderson was asked what to do about cycle miss if no cycle marks were recorded - how will a disk know how to recreate a data stream that a camcorder will accept if cycle misses are not recorded? Anderson answered that the computer would have to graft in the missing packet.

Discussion turned to the problem of how to find packet boundaries if the cycle index mark feature is not used. No reliable way exists for all protocols, especially those with variable-length packets. A cycle master disruption during recording, such as that caused by joining two busses, will create a time discontinuity in the recorded cycle marks.

Discussion moved to the topic of how the computer would represent a play list for the drive when sending or receiving isochronous data under SBP (ORB-based) direction. It was proposed to use a page-table-like structure rather than the list of full-fledged ORBs once proposed in SBP-2. A problem was mentioned: how to transmit a stream composed of parts recorded at different times, which may have variable channel numbers and packet source IDs. A worse problem was mentioned: The MPEG Transport Stream has embedded timestamps, up to 3 per packet, in addition to CIP headers, which must be made consistent throughout the stream during transmission - How will drive know how to mutate those to create a valid stream? Smyers concluded that SBP/AV drives will need to be built with knowledge of popular transport formats such as DV and MPEG2-TS, and likely there won't be any more formats because the consumer world thinks that MPEG2-TS is the ultimate goal. Smyers added that other protocols such as the IICP protocol can work with such drives as long as they are better designed and don't include embedded absolute timestamps or source IDs.

Anderson withdraw his suggestion to make the cycle mark optional.

Fairman suggested starting the recording of cycle marks at a zero cycle index, and recording increasing cycle marks that are frequency-locked to the actual bus cycle starts, regardless of any time jumps (such as from joining two busses). Anderson noted that if two bussed were joined they will likely not be phase-synchronized, and so the number of cycles lost during the join period may be hard to define, and in addition, opinions may vary among different nodes on the joined bus as to exactly how many cycles occurred.

Smyers indicated he would make a formal presentation regarding asynchronous access to isochronous data at the next meeting.

Flake indicated he would present a competing proposal.

[This became ftp://ftp.t10.org/t10/document.01/01-067r0.pdf]

Johansson gave a presentation on the impact of 1394.1 bridges on SBP-based disk drives. Of particular note, Johansson suggested that it may be appropriate for disks to analyze self-ID packets after each bus reset, to learn if a quarantine exists as a result of the bus reset. Also, a bus reset can no longer cause a task set abort, because an initiator on a remote bus won't know about the reset, and won't take action to requeue the tasks.

Johansson described his "node handle" proposal, in which the target and initiator use synthetic 16-bit nodeID/busID values between themselves to refer to specific nodes. Each side will translate those to actual node ID/busID values even as bus numbers change over time. Node Handles are never used directly in 1394

packet headers, only as encapsulated in SBP data structures. After a bus reset, both sides can update their handles and continue operation, without aborting the task set. Local bus topology analysis and/or rediscovery/reverification, possibly including DEP, will be needed to refresh a handle after a bus reset.

Johansson plans to write up a proposal for node handles. Johansson noted that this behavior is desirable even without bridges, because it allows targets to continue operation after a bus reset without aborting their task set. In such a case, the login would not be lost.

Knudson gave a presentation (DocNum: T10/01-069R0) regarding surprise removal of 1394 storage devices from the 1394 bus. Such removal has really bad consequences if write caching is enabled on the drive. Microsoft would like to enable write caching only on disks that are known to be physically secure (such as inside a computer's enclosure) for maximum performance.

Fuller explained that there is a way to know which 1394 ports on a computer are internal, and that software can track the bus topology to find out which devices are inside the box. Software can then set write caching accordingly. It was also suggested that sophisticated users be allowed to enable write caching in any drive if they agree that they understand the danger that they must not hot-unplug that drive. Russell suggested advocating the use of the topology information Fuller mentioned to system vendors so that computers will be built with this information correctly recorded, allowing operating systems to know which ports are internal.

Knudson discussed the two-host one-drive problem. Microsoft's upcoming Whistler operating system uses the SBP-2 password set feature to make a drive belong to a specific PC. Whistler sets the drive's password to be equal to the drive's GUID. A computer that thinks it owns the drive will use the password; other systems won't use the password and will fail to log in. Consequently the drive will mount only on the correct computer.

A long discussion of other possible solutions ensued. Proposals included tagging the drive with the GUID of its owner, tagging the computer as an owner, and other solutions. Many present were uncomfortable with using the setable password as a way to tag drives that need ownership management.

Anderson called for a straw poll as to the desirability of Microsoft's proposed solution.
9 voted against this solution.
0 voted in favor.

Smyers moved to adjourn.