

P1394a meeting minutes
July 30, 1996
San Jose CA

The meeting was called to order at 1 pm by Dick Scheel. Dick announced that he was acting as substitute chair for this meeting due to Peter Johansson's travel. The attendees introduced themselves. A list of attendees is included at the end of these minutes.

Dick announced that the PAR for this group was initially rejected by the IEEE. They gave certain specific changes that were needed, which Peter will complete and resubmit. Since Peter was not present, Dick did not have information about the changes other than Peter's comment that they only involved the style of the PAR rather than the technical content. Until the PAR is resubmitted and approved this group will continue to operate as an unsanctioned working group, but will still operate itself under IEEE rules.

The minutes of the last meeting (5/24/96) were approved as written.

Action items from the last meeting were reviewed.

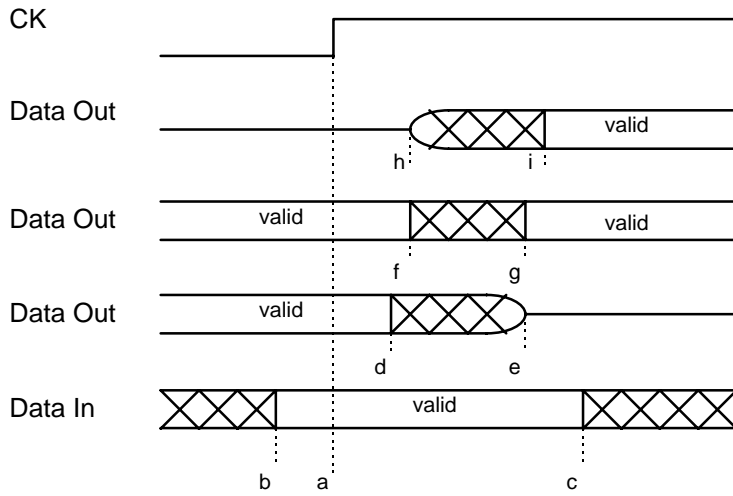
1. Mike Teener to find out the patent status regarding loop detection and parsing.
 - Mike was not present to report
2. Peter Johansson to distribute proposed rules / bylaws.
 - None yet, and Dick said that IEEE only suggests using bylaws, but doesn't require them.
3. Lou Fasano to check with TI on the patent status regarding bus hold / isolation.
 - Lou reported that he has inquired. Larry Blackledge of TI has not been able to give a final answer to Lou, but the general indication from TI is that they will probably go along with IEEE fair licensing policies.
- There were no reports on the following items since Mike and Peter were not present.
4. Mike Teener to find someone to give a presentation on 800 Mbps.
5. Mike Teener and Peter Johansson to seek a "sponsor" for the 800 Mbps work.
6. Peter Johansson to make a first draft P1394a document available for review by the end of June.
7. Peter Johansson to find a source for the IEC PCR document.
8. Peter Johansson to look for information about P1394.2 performance vs 800 Mbps.
9. Mike Teener and Max Bassler to get task forces started.
10. Peter Johansson to communicate to other interested parties that this group strongly recommends that all AV cables be capable of 400 Mbps.

It was reported that the 800 Mbps task group met the previous week. No details were given.

The group then discussed what is needed to change Annex J of the IEEE 1394-1995 standard from informative to normative. Table J-22, "AC Timing" needs to be reworked. It is too vague, and has some inaccuracies. Multiple companies that manufacture Phy and Link chips need to agree on the timing values, to assure that Phys and Links from different companies will work together.

It was noted that although capacitive isolation is acceptable for most uses, transformer isolation would be required for medical applications. Burke Henehan will try to find a specification for the maximum allowable leakage in medical applications.

In an attempt to improve table J-22, Bill Duckwall drew a timing diagram to illustrate the timing parameters that should be defined in the table. This applies equally to LREQ, DATA, and CONTROL signals.



		min	max	units
Data Out				
tristate -> data	a-h	nnn		ns
	a-l		nnn	ns
data -> data	a-f	nnn		ns
	a-g		nnn	ns
data -> tristate	a-d	nnn		ns
	a-e		nnn	ns
Data In				
setup	b-a	nnn		ns
hold	a-c	nnn		ns
Clock				
duty cycle		nnn	nnn	%
pk to pk jitter			nnn	ns

The table in the standard should contain correct numbers in each place that nnn occurs in the above table. As noted above, manufacturers of Link and Phy chips need to discuss what the numbers should be to insure interoperability between different manufacturer's chips. We need to consider what conditions need to be specified.

***** Editor's note: The above diagram and table should be carefully reviewed since they were drawn by the editor rather than Bill (or some other hardware person). *****

An interoperability test was proposed, where all Link and Phy chip manufacturers would bring their chips and different combinations would be operated together. It was decided that this should more properly be done in the 1394 Trade Association rather than the IEEE.

Lou Fasano brought up the question of whether a normative Annex J should specify timings. No answer was noted.

How do we address the interface for future higher speeds, such as S800, S1600, ...? Either use this same model with 16, 32, ... data bits between the Phy and Link, or come up with a new model.

Peter Teng brought up some requests from the "Open HCI" group. The first is to include two new Phy register bits in P1394a. They would use the two bits that are currently reserved. One bit would be a port disable bit, and the other would be used for loop healing. The port disable bit would make the port look like an open circuit to a port on the other end of the connected cable, as if the cable had been unplugged. The loop healing process is unspecified for now, and might be able to be implemented with only the disable bit. Bill Duckwall warned that the Apple Phy design (which is being used under license by many chip manufacturers) uses these two reserved Phy register bits for diagnostic purposes, storing the speed of the connected port. It is possible that an additional phy register would need to be defined to hold the new bits requested by the Open HCI group.

The Open HCI group also requested that we specify a different behavior of a bus when more than 63 nodes are plugged into the bus. Currently, the bus does not start up. The requested behavior is that when the self ID process reaches phy id = 63, additional nodes would continue to be numbered 63. At the end of self ID, the bus would start up, but all nodes with phy id = 63 would not respond to link-on or phy-config packets. An additional change that would need to go along with this is that the root would need to be the first to do self ID, rather than last. This would prevent the root from being disabled.

The group discussed what the consequences would be of this change. Current equipment may just take the last self ID packet as the ID of the root. Also, isochronous resource manager selection among multiple candidates is done by choosing the last self ID packet (highest id) among the candidates, so this would have a problem when more than 63 nodes are present. The consensus of the group was that there are too many problems with this change. Peter Teng will take this response back to the Open HCI group. More discussion on the reflector is encouraged to see if there is a good solution to this.

Regarding the request for defining the two new bits in the phy registers, the group was basically in favor of the idea as long as any possible problems with existing phy designs are worked out. Burke Henahan will try to get authorization to disclose the way these two bits are used in the Apple phy design.

Calto Wong asked the group to look at the interpretation of the 1394-1995 standard regarding the effect of bus reset on certain CSR's. Several of the CSR's do not have directly specified effects, although there is language at the beginning of the CSR section of the standard (8.3.2.1) that implies that these CSR's should revert to their "initial values". In particular, should the bus_id field of the NODE_IDS register be reset to 3FF on a bus reset? The answer is yes, but the wording in the standard is not clear.

Another item to be more clear about in the standard is in the isochronous resource manager BANDWIDTH_AVAILABLE and CHANNELS_AVAILABLE registers. The standard states that command reset returns these to a state indicating that all bandwidth and all channels are free. The document should be clear about this very dangerous effect of command reset. Perhaps this working group should review all CSR's to see if some should not be reset by a command reset.

It was pointed out that there is a possible problem with access to these registers anyway: If the register is changed by a successful lock transaction but the response packet is corrupted, then the register will have been changed but the initiator will think that it hasn't been. If the initiator sees a split transaction timeout on this lock transaction, it could try reading the register to see if the expected effect occurred, but this would be unreliable because another node may alter the register between the lock and read transactions (either making a success look like a failure or a failure look like a success). There does not appear to be any reliable solution to this.

The general situation is that any acknowledge packet or response pack may be corrupted and the initiator of the transaction will not know what happened.. Bill VanLoo said that the P1394.2 group has a mechanism to deal with this, but he'll have to look at it to see if an equivalent is possible in 1394-1995. Calto pointed out that if the low communication layers are able to detect anything that could help the upper layers & application be aware of this sort of occurrence, then information about it should be passed up to those layers.

Calto will put something on the reflector about this. He will include: a) a warning of what can happen, b) a solicitation of ideas to deal with these cases, and c) a solicitation of additional bad things that can happen that people should be warned about.

Peter Teng brought up another request from the Open HCI group. They propose that transmission of response packets not obey fairness (all other arbitration rules are still obeyed). This way fairness is handled on the basis of complete transactions, not separately for the request and response. This would be particularly helpful in the case of a server that receives many requests from many clients. Currently the server could receive a request from every client during a single fairness interval, but could only respond to one client per fairness interval. Under the new proposal, each client could have one complete transaction per fairness interval. The consensus was that this seems like a good idea, but that it should receive wider review to look for potential loopholes. We will give this more thought and discuss it on the reflector.

Calto asked that the minutes contain a reminder to Peter Johansson to write a clarification for the 1394 standard regarding the fact that bus reset should reset the bus_id field of the NODE_IDS register to 3FF.

Next Meeting:

The next meeting was previously scheduled for mid September, but the group felt that that would be too soon. Peter Johansson and Bill Duckwall should get together to discuss how certain pieces of the document revisions will be written, and then discuss a meeting date on the reflector.

Attendees:

Dan Bezzant	Cirrus Logic	510-623-8300	bezzant@cirrus.com
Bill Duckwall	Firefly	408-429-9845	p1394@aol.com
Lou Fasano	IBM Microelectronics	914-892-8904	lfasano@vnet.ibm.com
Peter Teng	National Semiconductor	408-721-6594	peter@jedi.nsc.com
Kaz Abe	NEC	415-528-5904	kabe@SJ-PCEG.ccgw.nec.com
Calto Wong	Philips	914-945-6382	cxw@philabs.research.philips.com
Dick Scheel	Sony	408-955-4295	dicks@lsi.sel.sony.com
Bill Van Loo	Sun Microsystems	415-786-6870	bvanloo@sun.com
Mike Winchell	Symbios Logic	470-225-4807	Michael.Winchell@Symbios.COM
Richard Mourn	Symbios Logic	709-533-7421	richard.mourne@symbios.com
Burke Henahan	TI	214-480-2603	bhen@msg.ti.com
Henry Angulo	TI	214-480-2117	angu@msg.ti.com