IEEE P802.3 10GBASE-T Minutes

Study Group Interim Meeting
September 17th – 18th, 2003
Hotel La Fonte, Meeting Room E
Portonovo, Italy

Prepared by:
Sanjay Kasturia
Sanjay@teranetics.com
Teranetics

INTRODUCTION

This document is organized as follows:

1. Summary of meeting
2. Meeting agenda
3. General presentations and minutes
4. Motions & associated discussion
5. Other useful information for 10GBASE-T participants

SUMMARY OF MEETING

The meeting was convened at 8:40am, September 17th, 2003. Mr. Bradley Booth (Brad), the 10GBASE-T study group chairman, opened the meeting with a discussion of the agenda and goals for this meeting reiterating the importance of getting to the taskforce stage by the end of this year before the currently approved extensions expire. Brad reviewed the proposed agenda and after it was approved, moved efficiently through the administrative items.

Eleven technical presentations were originally planned for this meeting but one was withdrawn because the presenter, Vivek Telang, was unable to attend the meeting. After the technical presentations, Brad took a number of informal straw polls to gauge study group opinions and supports on a number of items. The meeting adjourned on the first day after the straw polls.

On September 18th, the second day of the meeting, 13 motions were discussed and voted upon. This included approval of the minutes of the last 10GBASE-T meeting at the July 2003 Plenary. The voting of the study group participants on a few key motions is listed in the table below. The study group voting revealed overwhelming support for moving ahead with setting a 100m range target for operation of 10GBASE-T over 100m of CAT-7 (Class F) and over 55m or longer of CAT-6 (Class E). The group was divided on support for requiring the PHY to operate over shorter distances of CAT-5e (Class D).

A detailed accounting of all the motions is available later in this document. The 10GBASE-T study group meeting adjourned at 1pm on Thursday September 18th, 2003. After the meeting adjourned, a subgroup had discussions on preparing tutorials on 10GBASE-T for the November 2003 Plenary to help prepare 802.3 members for the voting that would happen then.

Motion	Yes	No	Abstain	Result
Class E & F motion listing a target reach with one PHY of 100m for Class F and 55m or more for Class E	33	0	2	Pass
Class D Motion requiring the PHY to support >20m over Class D	20	10	8	Fail
PAR	Unanimous			Pass

AGENDA

- **Welcome and Introductions**
- **Goals for this Meeting**
- **Reflector and Web**
- **Ground Rules**
- **IEEE**
  - **Structure**
  - **Bylaws and Rules**
  - **Call for Patents**
  - **IEEE Standards Process**
- **Presentations**

-**Cabling in Internet Data Centers** Luc Adriaenssens, Avaya
-**Does Cat5e always have insertion loss margin?** Luc Adriaenssens, Avaya
-**Extending Category 5e Limits (UL Report and EMC Data)** Terry Cobb, Avaya
-***Loop Reach Feasibility: Class D and Class E UTP Cables*** Vivek Telang, Cicada -- withdrawn
-**Cat6 Shielded Channel Performance** Sterling Vaden, Superior Modular Products
-**Feasibility Study for 10G Over 100m Cat7 and <100m UTP** Scott Powell, Broadcom
-**EMI Considerations for 10GBASE-T,** Larry Cohen SolarFlare
-**Lower Complexity Architectures for Implementing 10GBT XTalk Cancellers and Equalizers FIRs** Sanjay Kasturia, Jose Tellado, Teranetics
-**A Feasibility Study for 10Gbps over Class D, E and F** Takeshi Nagahori, NEC
-**10GBASE-T Market Potential & Technical Feasibility on Installed Cabling by 2005** Shadi AbuGhazaleh, Hubbell Premise Wiring, Sterling Vaden, Superior Modular Products, Valerie Rybinski Hitachi Cable Manchester
-**10GBASE-T Objectives** Chris Di Minico, MC Communications

- **Liaison Letter(s)**
- **Future Meetings**
- **Motion Madness**
  - **Thirteen motions including approval of the Minutes (July '03)**

GENERAL PRESENTATIONS & MINUTES

The meeting was called to order by the 10GBASE-T Study Group Chair, Bradley Booth (Brad) at 8:40am. Jeff Warren did not make it to the meeting due to weather conditions near his home and Sanjay Kasturia volunteered to be the recording secretary for this meeting. Brad asked for approval of the proposed agenda to which call Luc Adriaenssens (Luc) brought up the subject of liaison letters. Brad and Luc agreed to discuss liaison letters towards the end of the presentation session. Subsequent to that, Ron Nordin moved to approve the agenda. Bruce Tolley seconded the motion and the motion was passed by acclamation. The agenda is also available online at: http://www.ieee802.org/3/10GBT/public/sep03/agenda_1_0903.pdf

Brad listed initial goals of study group at the beginning of study group formation and emphasized that this meeting would focus on further refining and developing the Objectives, 5 Criteria and PAR so they could be sent to 802.3 before the November 2003 Plenary meeting. If all went well at the November Plenary meeting, there would be a motion to approve on the Thursday of that meeting. Subsequent to approval at 802.3,

these would go to the 802.3 executive committee on the Friday of the plenary meeting. For this to go smoothly, Brad said that we in the study group should prepare the rest of 802 by presenting tutorials on 10GBASE-T and that planning for the tutorials was scheduled for the afternoon of September 18th after Motion Madness which was scheduled for the morning of September 18th.

Brad then moved on to administrative items. He provided directions on how to subscribe to the 10GBASE-T reflector and also recommended that you subscribe to the 802.3 reflector in addition to the 10GBASE-T reflector. Ground rules for the group were reviewed and Brad stated that two sets of votes would be taken on each motion - the votes of all attendees, which would be recorded as study group votes and votes of the currently eligible 802.3 voters. The voting results would be passed on to 802.3.

Brad clarified that participants can discuss pricing if it is publicly available information. Brad reviewed the IEEE structure. NesCom approves the PAR; RevCom approves the final Standards proposal based on procedural checks rather than content of the specific Standard. Brad presented links to IEEE, the 802 and the 802.3 rules. Brad reviewed the IEEE-SA standards board bylaws on Patents in Standards. After that Brad called for patents at 9:00am. Brad reviewed inappropriate topics for IEEE working group meetings.

Brad reviewed the Standards process and where we stand on it, starting from the initial meeting in Nov 2002 at Hawaii. Brad said we tried to get 802.3 conditional approval in July but failed and the primary focus of this meeting was to:

- Hear presentations concerning:
  - Objectives, 5 Criteria and PAR
  - Suggested amendments to above
- Refine and enhance:
  - Objectives
  - 5 Criteria (if required)
  - PAR (if required)
- Plan for November 2003 Tutorial

The Objectives stay with the study group and help guide the Task force when we get to that stage. The PAR goes to the Standards Executive Committee (SEC) and then NesCom and the 5 Criteria goes to the SEC.

PAR approval initiates the start of task force meetings. The task force would discuss proposals and continue discussing proposals till it feels ready to move from proposals to a draft. Draft 1 gets worked on by the Task force till it passes task force review. After task force review it is forwarded for the 802.3 working group ballot. Approval of the working group ballot leads to Draft 2.0. After approval by >75% and successful handling of all comments, Draft 3.0 is generated. Draft 3 gets forwarded to sponsored ballot when there are no more technical comments. CX-4 is currently in sponsored ballot. RevCom approval leads to a standard. The slides used for this discussion are available online at: http://www.ieee802.org/3/10GBT/public/sep03/agenda_1_0903.pdf

Objectives stay within the group and help guide the group. The PAR is the 40,000 ft view. Study groups only exist for 6 months unless extended. Extensions can be requested and are voted on by 802.3, and ratified by SEC. We have already asked for two extensions and are coming up on one year for the study Group.

Luc asked if any study group had ever been refused an extension – Brad could not recall one. Tom Dineen commented that the group would move forward if a large enough group of people was interested in proceeding further.

Brad went through the structure of a PAR. Brad reviewed the 5 Criteria:

5 Criteria

• Broad Market Potential
– Broad set(s) of applications
– Multiple vendors, multiple users
– Balanced cost (LAN vs. attached stations)
• Compatibility with IEEE Std. 802.3
– Conformance with CSMA/CD MAC, PLS
– Conformance with 802.2
– Conformance with 802 Functional Requirements
• Distinct Identity
– Substantially different from other 802.3 specifications
– One unique solution for problem
– Easy for document reader to select relevant spec
• Technical Feasibility
– Demonstrated system feasibility
– Proven technology, reasonable testing
– Confidence in reliability
• Economic Feasibility
– Cost factors known, reliable data
– Reasonable cost for performance
– Total installation costs considered
• As per 5criteria_2_0703.pdf

Brad reviewed the Objectives:

• Preserve the 802.3/Ethernet frame format at the MAC Client service interface
• Meet 802 Function Requirements, with the possible exception of Hamming Distance
• Preserve min. and max. frame size of current 802.3 Std.
• Support full duplex operation only
• Support star-wired local area networks using point-to-point links and structured cabling topologies
• Support a speed of 10.000 Gb/s at the MAC/PLS service interface
• Select copper media from ISO/IEC 11801:2002, with any appropriate augmentation to be developed through work of 802.3 in conjunction with SC25/WG3
• Support 100 m over 4-connector structured 4-pair, twisted-pair copper cabling
• To not support 802.3ah (EFM) OAM unidirectional operation
• Support coexistence with 802.3af
• Support Clause 28 auto-negotiation

Brad reviewed the 10GBASE-T timeline.

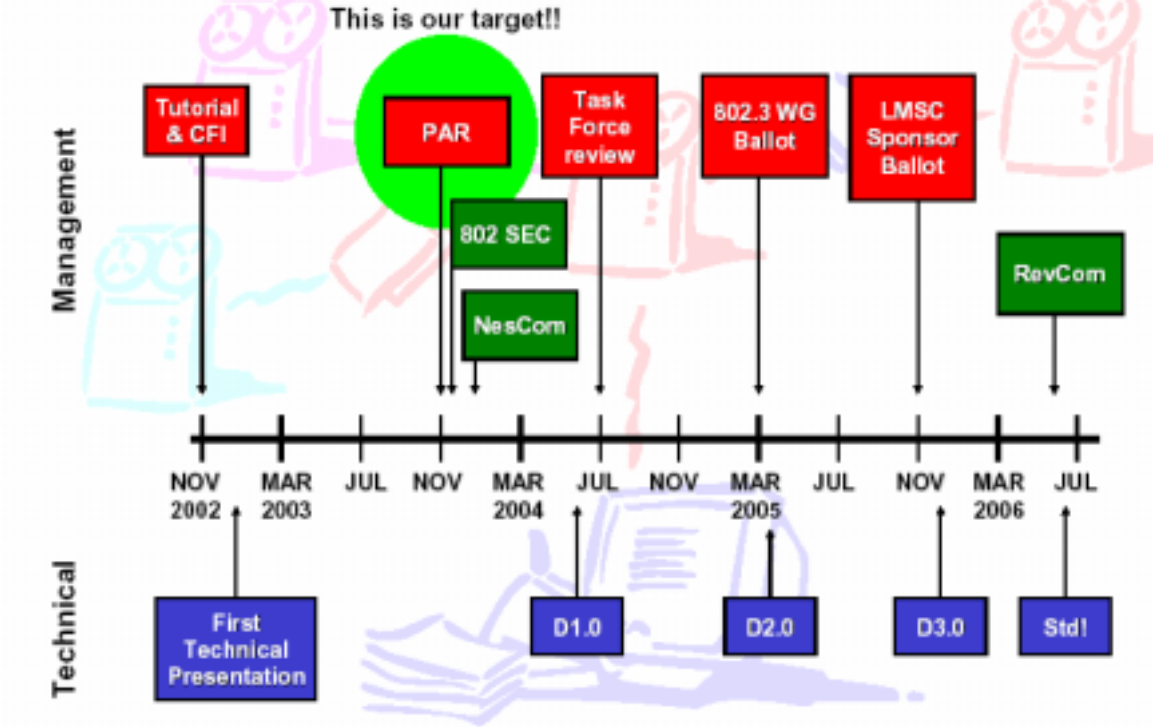

Immediate target is to get the PAR approved through 802 SEC and NesCom.

Brad reviewed Liaison letter & ad hoc reports.
A liaison letter was received from TIA TR42 in response to our letter and that we had not responded to it in the July meeting.

Brad listed the presentations scheduled and mentioned that Vivek Telang had pulled his presentation because he was unable to come to the meeting. The next meeting is in Nov 10-13 Albuquerque, NM, the Hyatt is full and they have a waiting list.

Brad talked about a possible interim meeting in January 2003 possibly in Vancouver – Brad needs to check the logistics.

Presentations and associated discussion

Presentations started at 9:25 am on Wednesday, September 17th.

Cabling in Internet Data Centers

Luc Adriaenssens reviewed his first presentation "Cabling in Internet Data Centers".
Luc displayed some pictures of data centers and also some cabling shipment forecasts in the US. The data presented does not cover cables other than copper twisted pairs. In calculating the cumulative profile, Luc presented two sets of data – one assuming no replacement of older cable and one assuming 100% replacement. Luc ended by making the case that the study group should focus on CAT-6 rather than CAT-5.
Shimon Muller questioned Luc as to when the pictures were taken. Shimon's point was that the cabling looks as shown in Luc's pictures when it has just been put in or just

cleaned up. Shimon's contention is that a year after the cabling has been put in place, it no longer looks as organized as in the picture and the picture is not typical through the life of the data center. Equipment is often moved and the cabling tends to get disorganized over time, especially close to the equipment and the patch panels.

Brad sought clarification of Luc's claim that most people are putting in CAT-6 because they want more reliability. Brad questioned whether the same drive to reliability would lead data center managers to move quickly to CAT-7 if it looked like 10G was coming down the road. Luc commented that the change from CAT-5 to CAT-6 is small relative to the change from CAT-6 to CAT-7. Luc mentioned that CAT-7 will have an adverse impact on cabling density and on the installation process, especially the cable termination procedures. Brad asked Luc's opinion on whether data centers would switch to screened or otherwise improved CAT-6 e.g. screened CAT-6. Discussion – Luc said yes, based on a desire by data centers to differentiate themselves from others.

Does CAT-5e always have insertion loss margin

Luc Adriaenssens went through his second presentation – "Does CAT-5e always have insertion loss margin." The presentation showed that cable can have low 0.4dB margin over the specified loss at 100 MHz. At 550 MHz the margin could be 2.25 dB or 3.8 % of the loss. Loss is dominated by copper loss at lower frequencies. At 100 MHz 90% is copper loss, 10% is material loss. At 400 MHz about 20% of the loss is material loss and material loss goes up with temperature. The cable in the plot shown will have less margin at 40degrees C. Insertion loss at higher frequencies increases more than the insertion loss increase at lower frequencies with temperature.

Scott Powell asked whether the material loss was linear with length and Luc said it was. George Zimmerman asked whether Luc had measured the specific cable at higher temperatures. Luc said he had not, but could do so if George was interested. Brad asked if cable actually sees 40 degrees C in data centers. Luc said that data centers typically see 20degrees C plus or minus 2degrees. The higher temperatures are more likely in warehouses.

Bruce Tolley suggested holding the worst case assumptions. Luc brought up the issue that someone has to step up to certify or guarantee cable and the guarantees only apply to the worst case performance. Chris Di Minico said that he has seen cable that is better. Chris claimed that in the 1000BASE-T case we did not make the worst case assumptions on all parameters. There was some debate on this with Sailesh Rao claiming that most parameters were scaled to the worst case limit during the 1000BASE-T standardization exercise.

A question was of Luc as to what is reasonable considering that 10GBASE-T will not be widely deployed initially. Luc suggested that we take the worst case performance at 20degrees C.

Extending Category 5e limits (UL report and EMC data)

The next presentation, "Extending Category 5e limits (UL report and EMC data)," was by Terry Cobb of Avaya. Terry showed a UL report to corroborate data he presented at

the last meeting. He clarified that the specific cable sample under discussion was 40m long. Terry stated that the energy lost in the notch is being radiated. Terry pointed out that the spike in emission will cause the cable to fail emission requirements. Shimon Muller of Sun Microsystems asked if the problem could show up in CAT-6. Terry said he hadn't seen the problem in samples of CAT-6. Luc commented that CAT-6 is built to much tighter manufacturing tolerances and is designed to operate to higher frequencies and is less likely to have similar problems. Larry Cohen of Solarflare asked what source was used for the measurements. Chris Di Minico of MC Communications said that the problem was well known, well understood and he would provide a reference to a paper that discussed the problem and what caused it. Chris said that problem was identified in 1991 and is now well understood and is not a mystery and that people know how to avoid the problem. Luc clarified that the sample Terry presented is relatively new cable and that it meets the CAT-5e spec. Terry mentioned that they have characterized CAT-6 from a number manufacturers – probably from all manufactures and have not seen this problem.

CAT-6 shielded channel performance

The next presentation was by Sterling Vaden and was titled "Cat6 shielded channel performance." He presented data of a setup where the cable and patch cords are shielded on each pair with foil. He showed a picture comparing the cable to CAT-5. The CAT-6 shielded cable was 0.25in in diameter whereas CAT-5 was 0.23in in diameter. Luc commented that foil shielding of the pairs increases the loss of the cable and compensatory changes to correct the loss will lead to a thickening of the cable. Sterling showed a picture of cable arrangement for testing. He used six-cable bundles with tape tying the cables together every meter. All parameters including alien next were measured. Cables not connected to sources were common mode terminated. Baluns used were only rated to 400Mhz. Sterling said he had detailed measurement data on a CD and would make it available to anyone interested. The specific cable had 23awg wire. The power sum ANEXT for all 24 pairs was substantially lower that the SolarFlare proposal except at 500 MHz where it gets close to the SolarFlare proposed specification. There was discussion and speculation on the cause of the increase near 500MHz. The CAT-6 connector did not have each pair individually shielded. Sterling's conclusions were that CAT-6 shielded (SSTP) channel supports 20Gbit (or better) Shannon capacity over a 100m channel. Alien NEXT and Alien FEXT are at or below measurable levels for much of the frequency range. Connectors are connected to frame ground or the rack. Near end and Far end were on different racks and the racks were not grounded. Cable shields were connected to connector ground. Both ends were connected to the same network analyzer and no isolation capacitors were used so there was a ground path between both racks through the test equipment. Kevin Brown of Broadcom asked what the implications were of the data presented. Sterling suggested that SSTP CAT-6 should be added to the list of cables being considered. Kevin asked whether Sterling was recommending that this cable be listed as a desired medium for 10GBASE-T. Luc questioned whether this cable met all existing CAT-6 specifications. This cable is much smaller in diameter than CAT-7. CAT-7 is typically 0.4in in diameter. Antoine Tazbaz of CDT/Nordx asked why Sterling did not go for CAT-6 with an overall foil shield rather than one with foil on each pair. Sterling commented that he could not find that kind of cable – it is harder to manufacture. Next question from Antoine was how much of this kind of cable was installed and what

percentage of the installed base consisted of this kind of cable. Terry Cobb raised a concern that going from individually shielded pairs to a common shield at the connector would cause a mode discontinuity and create significant problems.

Meeting called to order after a break at 11:20am

Feasibility Study for 10G Over 100m CAT-7 and <100m UTP

Presentation by Scott Powell of Broadcom – titled: "Feasibility Study for 10G Over 100m Cat 7 and <100m UTP." Scott stated that the presentation would cover 1) ADC requirements for 100m CAT-7 cable, 2) Maximum cable reach on CAT-5e/6 and 3) Power estimation of a single PHY that could operate over both channels.
Results are dependent on the assumptions and Scott outlined the assumptions. The architecture assumed is 10-PAM which is not necessarily the ideal solution but is the one that has been discussed the most. Scott's results assume ideal analog hybrid/echo cancellation and digital cancellation of echo, NEXT and FEXT with no round-off noise and Digital equalization with DFSE (no round-off noise). Ideal transmit and receive amplifiers (no noise or distortion). All cable parameters were scaled to worst case limit lines except for alien NEXT which was kept to actual measurements since they already touch the limit line.
On CAT-7 the main challenge is the ADC. Scott showed a plot of max cable reach relative to the ADC ENOB. 100m on CAT-7 requires ENOB of 10.5 to 11.5 bits assuming a system design margin between 3dB and 6dB.
On CAT-5e and CAT-6 UTP, with 3dB design margin the plots of reach vs. ENOB of ADC show that the curves flatten out because very quickly the alien crosstalk becomes the limiting factor. CAT-5e maxes out at 20m and CAT-6 UTP reach maxes out at 30m. Models used for CAT-6 were four connector measurements. For CAT-5e, the two connector measurements were used.
Scott then moved on to power estimation. He said that this is a difficult thing to do but did estimate it. The power is estimated for a PHY that does 100m on CAT-7 and 30m on CAT-6. CAT-7 drives ADC spec and echo canceller complexity. 30m CAT-6 drives NEXT/FEXT canceller complexity. Sailesh Rao asked what was assumed for the receive filter and questioned Scott's computation of the ADC ENOB requirement and challenged Scott's plots. Scott presented data on high speed high performance ADCs and used that to come up with a number of 4w for the requisite ADC. Tom Dineen questioned Scott's assumption about what will happen in 5 years saying that feasibility should be demonstrated sooner than five years. Sailesh asked what the peak to peak of the signal would be at the ADC input. Scott set the full scale to get a clipping probability of 10 raised to -15. Sailesh suggested that this was too large a scale and that higher clipping probabilities could be tolerated because of the error correction capability of the trellis code in the system. Scott then went through his estimate of the digital power consumption and came up with an estimate of 5.5 watts for the digital implementation of the echo/next cancellers and assumed 1.7 watts for the DFE etc and 1.5W for the FEXT canceller and came up with 8.8w. He expects this to scale down by a factor of 4 over two generations of process technology and be down at 2.2w in five years. This brings the total power per port at about 10 to 18 watts in five years. Eliminating the requirement to go

30m over CAT-6 would save about 1 watt of power per port. Luc questioned the appropriateness of Scott's assumption of 833 MHz for the symbol rate since the distance target on CAT-5E that Scott considered was 30m. Shimon asked why the range in power consumption estimates was so large and whether his power estimates were worst case estimates. Scott said that his power estimates were worst case. Nariman Yousefi of Broadcom added that the range of power variation would be between 10 to 20% over temperature and process. George Zimmerman suggested that over sampling could be used to improve the ENOB of ADC since many of the ADCs that Scott used as a reference for his ADC power estimate were capable of operating at much higher speeds than were needed for this application.
Shimon said that 10W would be ok for many applications. Bruce Tolley also said that the power range seemed reasonable and had confidence that this was not out of line with the introduction of new technology in the past. Bruce mentioned that 1000BASE-T initially was at 7 watts and then down to 5watts soon after that. In comparison, 10GBASE-T was 10x the speed of 1000BASE-T and coming in at possibly 2x the initial 1000BASE-T parts. Bruce also said that he assumed that actual progress on power reduction would be faster than Scott's assumptions. Scott commented that his current estimates were for five years out in the future.

Analysis of EMI measurements on UTP cable

The next presentation was by Larry Cohen and was titled "Analysis of EMI measurements on UTP cable". The objective of Larry's study was to check suitability of media for wideband signaling. He examined different media, different connectors, patch panels, other components and actual installation configurations. Larry summarized that the dominant radiation was from common mode signals. At lower frequencies, the common mode radiation could be greater than 100 times the differential mode radiation. At higher frequencies this ratio drops. The transmit power was 10dB and 12 db was added to the measurements to account for the peak measurements. Larry also did measurements using a current probe that scanned the whole length of the sample cable. Luc suggested that the presence of the current probe could itself affect the measurements. Larry showed that the current probe measurements correlated very well with the radiation/EMI measurements. The current probe measurements are very pessimistic measurements of the EMI. There was detailed discussion on the configuration with Terry Cobb cautioning. Larry said that the EMI is very dependent on the specific termination used to launch the signal on to the wire. Most of the emission happens in the first 10m to 15 m from the launch. There was discussion on the different terminations. The Bob Smith termination terminates the common mode standing waves resistively. Larry was using a different termination than the Bob Smith termination. The cable itself is fine. The common mode radiation is dominant. Line Interface Design is the dominant issue. Sharp bends create resonant effects. Sterling Vaden asked if Larry had looked at variability of radiation based on sharp bends. Larry said he had not measured the impact of sharp bends yet but had looked at other discontinuities. There was a fair amount of discussion and Larry mentioned that there was a lot of variability based in lengths of the launch patch cords. There was debate between Luc and Larry about the impact of the twists in the UTP on the EMI radiation. Scott Powell asked whether the transformers and terminations

needed for 10GBASE-T would be significantly more expensive than those for 1000BASE-T. Larry said that the cost would not be very much higher.

After Larry's presentation the meeting adjourned for lunch. Meeting resumed at 2:45pm

Lower Complexity Architectures for Implementing 10GBT XTalk Cancellers and Equalizer FIRs

The next presentation was by Jose Tellado and was titled: "**Lower Complexity Architectures for Implementing 10GBT XTalk Cancellers and Equalizers FIRs**." Jose listed three techniques that can be used to reduce implementation complexity of long FIR filters and provided estimates of the reduction for one of the techniques. Scott Powell questioned the saving suggesting that a closer look needs to be taken at the details in that counting operations may not be an apples to apples comparison since the direct form implementation of the echo and next cancellers may require lower precision multipliers than a transform domain implementation. There was some discussion on the latency introduced by this method.

A Feasibility Study for 10Gbps over Class D, E and F

The next presentation was by Takeshi Nagahori and the title of the presentation was "A Feasibility Study for 10Gbps over Class D, E and F". Mr. Nagahori walked through a system model and then showed results that Class D would support 60m with 5dB of alien next mitigation and a 3dB system margin. On Class E he showed an achievable distance of 70m assuming, once again, 5dB of NEXT mitigation and 3dB of system margin. Based on the system model he also presented a summary of AFE requirements. His requirements chart showed and ADC requirement of 9.5bit ENOB and suggested an implementation based on 8-way interleaved 9.5bit 100Ms/s ADCs. Mr. Nagahori's analysis showed that the total jitter budget is 3ps RMS. Mr. Nagahori also summarized his estimates of chip power consumption. In conclusion, Mr. Nagahori suggested a transceiver that would cover 100m over Class F, 70m over Class E and 60m over Class D. The power consumption in 90nm is estimated at 11.7 watts and the power is estimated to be 6.2W in 65nm technology. There was discussion of the significant differences in the results presented in this presentation relative to the presentation made earlier in the day by Scott Powell. One discrepancy was on the required ENOB. For CAT-7 Scott estimated 10.5 bit ENOB whereas Takeshi's presentation indicated 9.5bits. The other significant difference was in the power consumption. The most significant difference in power consumptions related to the power of the ADC. Takeshi assumed multiple ADCs operating in parallel. His basis for computing power consumption was an ADC core developed by NEC. There was discussion if Takeshi's power numbers counted the power required by the sample and hold and it was pointed out that the sample and hold would need an input bandwidth that exceeded 400Mhz even though the component ADCs only operated at 100Mhz. Takeshi confirmed that his estimates included estimated power for the sample and hold.

10 GBASE-T Market Potential & Technical Feasibility on Installed Cabling by 2005

The next presentation was by Valerie Rybinsky and was titled "**10GBASE-T Market Potential & Technical Feasibility on Installed Cabling by 2005.**" The presentation represented work by Valerie, Shadi AbuGhazaleh and Sterling Vaden. Valerie listed their assumptions in computing Shannon capacity and combined the capacity requirements together with a model of the installed base (covering cable types and cable lengths) to show that the majority of the installed base will support 10G. In her results the assumption was that CAT-5E capacity would exceed 18Gbps for lengths less than 40m. The corresponding ranges were 76m for CAT-6, 71m for CAT-6 scaled and 65m for scaled worst case CAT-6. Luc questioned some of the results and suggested there might be some differences between Valerie's results and results on capacity presented at earlier meetings. Nariman asked Valerie to clarify what the installed base covered. Valerie responded that the installed base data came from measurements collect by the Ad Hoc cabling group. Chris clarified that some of the measurements exceeded in some places the extension of the limit lines. Valerie next went over the length distribution histograms for CAT-5 and CAT-6. From her results it seems that 44% of the CAT-5e installations would meet the calculated length requirement. The corresponding number for CAT-6 was about 80%. Valerie said the while 100% of CAT-7 installations would support 10G they were a very small percentage of the installed base. Valerie's recommendation was that the study group provides cabling performance requirements and leaves it to the cabling industry to come up with a UTP product that would meet these. Bruce Tolley mentioned that instructions from 802.3 at the San Francisco meeting were to come up with cable and specific distance targets. Shimon reiterated that we rephrase the objectives in terms of distance target rather than percentage of installed base. He suggested that percentage of installed base was going to be hard to measure. The % coverage is an appropriate metric in the five criteria. There are significant differences in the range estimates from Valerie and from Scott's presentation. Scott pointed out that some of the difference came from the fact that Valerie came at it from Shannon Capacity and Scott's results were based on MMSE analysis. This difference explains part of the difference in results but the larger difference comes from differences in assumptions of alien NEXT.

The meeting was adjourned for a break and resumed at 4:30pm.

10GBASE-T Objectives

The next talk was by Chris Di Minico on **10GBASE-T Objectives**. He presented the consensus of a group of twelve people (listed in his presentation). Chris suggests modification of the current objectives to specify distances and now suggests 100m on four pair Class F balanced copper cabling, 55m to 100m on four pair class E balanced copper cabling and at least 45m to 100m on four pair class D balanced copper cabling. The presentation also suggests that objectives be modified by specifying that the PHY meet CISPR/FCC Class A emission requirements. Chris handed of the presentation to George Zimmerman to explain the rationale behind the specific distance targets. George walked through the assumptions behind the numbers and provided capacity based justification for the ranges. He also provided SNR margin to hit the 1e-12 BER with a

DFE based receiver. For 45m of CAT-5e, the margin was 3.8dB. For 55 m of CAT-6, the capacity computed was 19.8 Gb/s with the SNR margin being 3.4dB. The corresponding numbers for 100m of CAT-7 are 28.6 GHz and 16dB. George then reviewed some of Larry's results on EMI and stated that that indicates that EMI emissions will be met but that this is safer with some augmentation of the cable specifications. Luc asked about the assumptions on the transmit power level. George replied 10dBm applied with a flat spectrum. PJ Sallaway of Vativ asked about assumptions behind the margin numbers, specifically the ADC assumptions. George mentioned 8.8 bit is the assumed ENOB though the margin numbers are based on an optimal DFE. George commented that detailed simulations produced results within 1dB of the optimal DFE results. George also mentioned that the cable models used were from the cabling Ad Hoc. George also reviewed Scott Powel's estimate of the power consumption of the AFE for 10GBASE-T. George applied Scott's methodology of starting with the power consumption of stand-alone ADC chips to 1000BASE-T and came up with larger numbers for 1000BASE-T than are actually achievable. Scott challenged George's analysis of 1000BASE-T by pointing out that George had not picked the right parts. 1000BASE-T should require no more than a 125 MHz ADC and an ENOB to 8bits. George suggested an alternative approach to computing the ADC power. He suggested a parallel approach and computed the power for this approach. George referred to a Broadcom presentation that showed a parallel approach to addressing the ADC requirements. Using some data available for 10 bit 120 MHz ADCs, George assumed achievable power today would be 2.88W. George also claimed that process advances should improve this number and suggested <1.5W by 2005.

Based on an existing detailed design George estimated a digital implementation would take 6M gates and is capable of 1.5Teraops. George estimates that 90nm technology will permit implementation with 7Watts.

Motions and associated discussion

Chris Di Minico suggested some straw polls in anticipation of the Motion Madness on Thursday morning. Brad followed up and asked for straw polls on a number of topics.

Straw Poll 1: 802 Functional Requirements

Agree that the objective "Meet 802 Functional Requirements, with the possible exception of Hamming Distance" should be modified to "Meet the requirements of IEEE Std 802-2001".

Results Y: 25 N: 0 Don't care 0

Discussion: The first poll was on the subject of 802 Functional Requirements. There were a lot of questions what the significance of the polls was. Brad mentioned that these polls would help him better prepare for the motions on Thursday.

Straw Poll 2: CISPR/FCC

Agree that requirement to "Meet CISPR/FCC Class A" should be added to the objectives.

Results Y: unanimous by acclamation N: 0 Don't care: 0

Discussion: Luc pointed out that there are differences in CISPR and FCC. Above 230 MHz one is more stringent than the other by a few dB.

Straw Poll 3: BER

Agree that requirement to "Support a BER of 10E-12" should be added to the objectives.
Results Y: passed by acclamation N: 0 Don't care: 0

Straw Poll 4: Cabling Structure

Agree that the objective "Support 100m over a 4-connector structured 4-pair, twisted-pair copper cabling" should be modified to "Support operation over 4-connector structured 4-pair, twisted-pair copper cabling"
Results Y: 28 N: 0 Don't care: 3

Straw Poll 5: Class F

Agree that the requirements to " Define a single 10Gb/s PHY that would support link of at least 100m on four-pair Class F balanced copper cabling"
Results Y: 24 N: 0 Don't care: 8

Straw Poll 6: Class E

Agree that the requirement to "Define a single 10 Gb/s PHY that would support a link of at least 100 m on four-pair Class F balanced copper cabling".
Results Y: 22 N: 5 Don't care: 0

Straw Poll 7: Class D

Agree that the requirements to "Define a single 10Gb/s PHY that would support a link of at least 45 to 100m on four-pair class D balanced copper cabling"
Results Y: 13 N: 12 Don't care: 4

Straw Poll 8: Class D II

Agree that the requirements to "Define a single 10Gb/s PHY that would support a link of at least 40 to 100m on four-pair class D balanced copper cabling"
Results Y: 11 N: 15 Don't care: 1

Straw Poll 9: Class D III

Agree that a requirement for Class D should be added to the Objectives"
Results Y: 14 N: 15 Don't care: 3

802.3 voters: Y: 9 N: 3 Don't care: 2
Discussion: Brad asked for a vote of the currently eligible 802.3 voters. There was discussion whether it should be current 802.3 voters or those who will be eligible in Nov could vote. Brad decided to take count only of currently eligible voters.

After the straw polls, the meeting was adjourned for the day.

Meeting re-convened at 8:50am on Thursday, September 18th by Brad Booth.

Motion 1: Class E & F

Move that the following objective be added:
Define a single 10 Gb/s PHY that would support links of:

- **At least 100 m on four-pair Class F balanced copper cabling**
- **At least 55 m to 100 m on four-pair Class E balanced copper cabling**

Proposed by: Bruce Tolley
Seconded by: Shimon Muller
Technical motion (>75%)
SG vote: S: 33 N: 0 A: 2
Motion Passed
802.3 voters S: 13 N: 0 A: 2

Discussion: Sterling asked to change the length from 55m to either 50m or 60m. Brad requested Sterling propose a friendly amendment to change to 50m. Shimon recommended it be considered unfriendly. Geoff Thompson suggested that the four connector model allow for 2.5 m of jumper length on each end. Luc suggested that 55m would imply more precision that the prior presentations implied and suggested 50m. Geoff Thompson said that if the target is set to 50m, it would leave less than 50m for the structured part. Suggestions were made that the 55m be explained as 50m for the structured part and 5 for cords on each end. Shimon suggested that this was too much detail for the PAR.

Motion 2 Class D Motion

Move that the following objective be added to the "Define a single 10 Gb/s PHY that would support links of:"

- **At least 20 m over four-pair Class D balanced copper cabling**

Proposed by: Bruce Tolley
Seconded by: N. Yousefi:
Technical motion (>75%)
SG vote: S: 20 N: 10 A: 8
802.3 S: 7 N: 3 A: 1
Motion failed.
Discussion: Terry Cobb questioned how you could write a channel spec for 20m. Antoine Tazbaz (Nordex/CDT) asked if this would be a four connector model. His comment was that if it was 20m over four connector channel it would not be viable. Chris Di Minico

said that Scott Powell's presentation showed that the 20m was achievable. Valerie said that 20m would cover a small percentage of the installed base and should be dropped. Luc said that the majority of cabling in data centers is already Class E. Terry Cobb said he had a hard time believing that people would continue to use CAT-5 in data centers and would not pull it out and replace it. George Zimmerman said that based on data presented, a small percentage of CAT-5 would still be a significant addition to the market potential. Sterling expressed concern about who would support calls when 10GBASE-T does not work on a specific installation. Sterling said that it was not a technical issue but a support issue. Antoine said that 10G would not be needed at the desktop and the effective installed base that would be supported would be down to 1%. Geoff Thompson pointed out that the cabling industry said it would not support augmentation of the Class D channel. Since we don't have an alterative support organization this will have to operate over an unspecified channel and Geoff reiterated that he would not support a motion for operation over an unspecified channel. Shimon Muller brought up the point that data presented yesterday showed that the PHY would operate over some distance over CAT-5 and this distance was likely to exceed 20m and Shimon felt that there are many customers who would have data centers where this distance could be adequate and he should not have to ask these customers to replace the cable. George Eisler of Solarflare said that a similar situation existed in 1000BASE-T in that an installed base had to be recertified. George Eisler said that it hadn't been a problem in the past. Luc said that there was a difference in this case and the cabling industry had stated that they would not support 10G over CAT-5. Antoine said that the practicalities of identifying which cable lengths would be a pain. Shimon said that there are rooms where you can be sure that the length will be less than the specified range over CAT-5. Sterling said that the room size did not guarantee cable length and length would have to be actually measured. Chris pointed out that a similar situation existed in the fiber space for 10G and that the industry figured out how to work the issue. Chris reiterated the importance of protecting the customers' investments in their cabling. Bruce Tolley said that if the cabling industry was not willing to support Class D, he would approach others, possibly the PHY vendors to support this re-qualification. Geoff pointed out that the fiber issue delayed the optical standard by one meeting cycle. Chris said the situation here was better than the situation in fiber. Terry Cobb said that he had presented data that clearly showed that 10G would not work over some Class D cable. In the 1000BASE-T case he said that they had verified that it would work so they had supported it. Bruce Tolley said that if he could not offer some support of the installed base to his customers, he would rather recommend CAT-7 than CAT-6 since CAT-7 is specified and the modified CAT-6 discussed in the presentations has not yet been specified. Chris pointed out that at the shorter distances being discussed the testing should be easy and definitely doable. He said that at the shorter distances there would be no need to try and measure alien crosstalk. Antoine said that he supports the motion technically but feels that the practicality is not viable. Bruce mentioned that 1000BASE-T PHYs support TDR functionality and could self measure the distance. George Zimmerman brought up the point that cable testers that can measure length on CAT-6 should be able to measure the length on CAT-5. Luc pointed out that the cabling industry will not support CAT-5 and said that the liaison letter was very clear on this. Shimon pointed out that there is no elaborate qualification required. George Eisler commented that the IEEE should not have direction dictated by the cabling

industry. Geoff Thompson reiterated George Zimmerman's comment that a tester will not care whether the cable being tested is CAT-5 or CAT-6 but said that IEEE only makes standards based on the specified behavior of cable. Geoff said that when such a situation arises, there are three choices – a) cabling industry will augment the specification, b) IEEE does the specification augmentation itself, or c) we drop the proposal. George Eisler said that if those were the choices and the cabling industry was not stepping up to doing this, he could take this on. Luc said that cooperation with the cabling industry will be required and the situation has to be a win-win situation for all groups. Luc said that in this case they have decided that 5E is obsolete and they will not support it because they will not get a return on their investment. Luc also pointed out that there was a lot more to cabling than people here appreciated. Chris reiterated George Eisler's comment and said that we should not be held hostage to component and cabling requirements. Chris said that the participants in the IEEE had put a lot of work on this and were being held hostage to the cabling guys. Antoine said he wanted to look at this from another angle. He said that there have been a variety of options in the past – CAT-3, CAT-5, CAT-7 and the industry has not always supported the older cable.
Sanjay Kasturia asked for clarification on the comment that the cabling industry will not recover its investment by asking what investment was involved in recertifying CAT-5e for 10G. Luc, Valerie and Sterling said that there were three costs – the characterization cost, the support cost and the customer cost and the last two would dominate in this case. Hugh Barass said that plausibility is all that is required at this stage rather than certainty and we should proceed. Luc asked to call the question and there was no opposition. Brad initiated a count.

Luc asked clarification on why we are taking 802.3 votes since opinions here should not be taken to 802.3 if it doesn't get approved by the study group. Bob Grow clarified that 802.3 members here have every right to take things to 802.3 – approval from the study group is not required. Brad clarified that when he took things to 802.3 in his role as chair of the study group, he would represent the results of the study group vote rather than the 802.3 vote.

Motion 3: Cabling Structure Motion

Move that the following objective "Support 100m over a 4-connector structured 4-pair, twisted-pair copper cabling" be modified to read "Support operation over 4-connector structured 4-pair, twisted-pair copper cabling for the supported distances and Classes".
Proposed by: Chris Di Minico
Seconded by: Bruce Tolley
SG vote: S: 32 N: 0 A: 2
Motion passed
802.3: S: 13 N: 0 A: 2

Discussion: Geoff Thompson pointed out that the cabling structure motion calls for a 4-connector structured model and this seems to imply four connector model for each of the

cabling types. Brad said his interpretation was that one of the cables should support a four connector model. Geoff requested modification of the text to clearly spell this out. Chris, as the mover said he meant the four connector model to all the cable types. Terry Cobb suggested a friendly amendment to the text to read – "Support operation over 4-connector structured 4-pair, twisted-pair copper cabling for the supported distances and Classes." Chris and Bruce agreed that it was a friendly amendment.

Motion 4: CISPR/FCC Motion

Move that the following objective be added:
Meet CISPR/FCC Class A
Proposed by: George Zimmerman
Seconded by: C. Di Minico
Technical (>75%)
SG vote: S: 32 N: 0 A: 3
Motion passed
802.3: S: 14 N: 0 A: 1

Discussion: There was discussion on how the IEEE could ensure compliance. The conclusion was that the lack of contrary evidence is all that would be required.

Motion 5: BER Motion

Move that the following objective be added:
Support a BER of 10^-12 on all supported distances and Classes
Proposed by: C Di Minico
Seconded by: S Muller
Technical (>75%)
SG vote: S: Unanimous by voice N: 0 A: 0
Motion passed
802.3: S: 14 N: 0 A: 1

Motion 6: 802 FR Motion

Move that the objective "Meet 802 Functional Requirements, with the possible exception of Hamming Distance" be deleted.
Proposed by: G. Zimmerman
Seconded by: C Di Minico
Technical (>75%)
SG vote: S: passed by voice vote N: 0 A: 0
Motion passed
802.3: S: 6 N: 0 A: 1

Motion 7: 5 Criteria 802.FR motion

Move that "The proposed standard will conform to the 802 Functional Requirements Document, with the possible exception of the Hamming distance." in *2. Compatibility with IEEE Standard 802.3* section of the 5 Criteria be modified to

read "The proposed standard will conform to the requirements of IEEE Std 802-2001."
Proposed by: S. Muller
Supported by: G. Eisler
Technical (>75%)
SG vote: S: Unanimous by voice N: 0 A: 0
Motion Passed
802.3 S: 13 N: 0 A: 0

Motion 8: PAR motion

Move that the 10GBASE-T Study Group give the 10GBASE-T SG Chair authorization to make the necessary modifications the "Sponsor Date of Request" and the "Contact Information of Official Reporter" fields in the PAR
Proposed by: Hugh Barrass
Seconded by: B. Armijo
Technical (>75%)
SG vote: S: Unanimous by voice N: 0 A: 0
Motion Passed
802.3: S: 9 N: 0 A: 1

Motion 9: Criteria 1 Motion

Move that the word "infrastructure" be changed to be "cabling" in the 5 Criteria
Proposed by: T. Cobb
Supported by: R. Mei
Technical (>75%)
SG vote: S: 13 N: 7 A: 2
Motion failed
802.3: S 3 N: 4 A: 2
Discussion: Terry asked for the change because he said that "infrastructure" implies 100m. George Zimmerman said that this change was unnecessary – 100m had to be called out explicitly and was not implied.

Motion 10: Criteria 2 Motion

Move that the word "structured" be removed from the second paragraph of Criteria #2.
Proposed by: T. Cobb
Supported by: R. Mei
Technical (>75%)
SG vote: S: 7 N: 10 A: 15
Motion failed
802.3: S: 1 N: 7 A: 4

Discussion: George Eisler questioned the need for this change at this point. Terry Cobb pointed out that Structured in TIA implies 100m and that IEEE has typically used cabling. George Zimmerman said that the current criteria have a requirement for 100m and this sentence is correct.

Motion 11: Criteria 3 Motion

Move that the second paragraph of Criteria #3 be modified from "It is the only standard that will use horizontal structured twisted pair cabling as defined in ISO/IEC 11801, offering upgrade paths to 10Gb/s for present Ethernet users connected with copper" to read "It is the only 10 Gb/s standard that will use twisted pair cabling."

Proposed by: Terry Cobb
Supported by: R. Mei
Technical (>70%)
SG Vote: S: 9 N: 14 A: 5
Motion Failed
802.3 S: 1 N: 10 A: 1

Discussion: Terry moved to call to question – 7 for calling 12 opposed. Question not called. Floor was kept open. Bruce Tolley was not sure why Terry is asking for this because Bruce believed we are using terminology that is consistent with instructions that were originally provided to us. Terry clarified that the discussion requiring measurements up to 625 MHz put measurements beyond 600Mhz as required by ISO/IEC. Chris suggested that there was mention of augmentation in the objectives so it should not be listed here. Others mentioned that objectives stay with the study group and criteria go to a different group so it should be called out here. George Zimmerman said that the CAT-7 spec was adequate.

Motion 12: Criteria 4 motion

Move that the second paragraph shown below of Criteria #4 be deleted.

- **The study group acknowledges that 10Gb/s operation is achievable on Class D and Class E cabling and augmentation of their specifications is required to higher frequencies for performance parameters such as insertion loss and the addition of alien crosstalk characterization. The study group also agrees that the 10Gb/s operation is achievable on Class F cabling. The channel models are supported by the measurement of the properties of cables and network hardware in both laboratory and field installations.**

Proposed by: T. Cobb
Supported by: R. Mei
Technical (>75%)
SG vote: S: 10 N: 15 A: 6
Motion failed
802.3 S: 1 N: 9 A: 2

Discussion: Terry said that the reference to Class D should be dropped. Bruce and Shimon said that there is nothing in this statement that commits the group to anything and so it should not be changed. Brad expressed concern that eliminating this paragraph would give the SEC the impression that we are not considering ISO/IEC specifications.

Motion 13: Minutes Motion

Move that the Study Group approve the minutes of the July 2003 meeting as documented in minutes_0703.pdf
Proposed by: B. Armijo
Supported by: R. Mei
Technical (>75%)
SG vote: S: unanimous by voice vote N: 0 A: 0
Motion passed

Closing

Brad asked if there was any other business. Brad reminded Luc that he had raised the topic of responding to the TIA TR42 liaison letter. Luc asked whether we should respond to the liaison letter saying that we wanted to work with them to lock in the cabling specs for 10GBASE-T, for example the enhancements to the CAT-6 spec. Brad asked whether an official response was appropriate or whether informal communication was adequate, as an official response would have to wait until the November meeting and approval from 802.3. By voice vote the group expressed that they would like the TIA to work on an augmented spec. Chris and Valerie will informally communicate this to the TIA TR42 group.

Randy Below moved to adjourn the meeting. Valerie seconded. The group approved by voice vote unanimously. The meeting was adjourned.

Other useful information for 10GBASE-T participants

Other useful Links:
September 2003 Interim week 10GBASE-T Study Group material
5 Criteria (after meeting)
Draft PAR (after meeting)
Draft Objectives
10GBASE-T Study Group Material
all_files.zip

The IEEE 802.3 10GBASE-T Study Group will hold a meeting during the 802 Plenary
November 9-14, 2003 in Albuquerque, NM
Space is limited at the Hyatt, but they do have a waiting list.