

**MASSACHUSETTS INSTITUTE OF TECHNOLOGY**  
**HAYSTACK OBSERVATORY**  
*WESTFORD, MASSACHUSETTS 01886*

*Telephone:* 978-692-4764  
*Fax:* 781-981-0590

May 22, 2004

TO: Distribution  
FROM : David Lapsley  
SUBJECT: 5 May 2004 e-VLBI telecon summary

Attendees:

Bill Fink, Lee Foster, Paul Lang, Mary Shugrue – GSFC  
Terry Gibbons – Lincoln Laboratory  
Dennis Baron – MIT  
Kevin Dudevior, Hans Hinteregger, David Lapsley, Arthur Neil, Alan Whitney – Haystack

This telecon is one of an ongoing series of telecons to prepare for 10 gigabit/sec e-VLBI demonstrations between NASA GSFC and MIT Haystack Observatory using a combination of network facilities including all or part of Glownet, Bossnet, ISI-E, SuperNet, Max and GSFC/HECN.

**ACTION ITEMS ARE HIGHLIGHTED IN RED.**

Bossnet

David Lapsley: Terry, Russ, Jerry, Tom and I have been discussing off-line for the last week or so. Jerry had got a good deal from Movaz on Optical Add Drop Multiplexers (Ray Express's). Same functionality as JDSU equipment, but a lot of additional functionality. Jerry able to get them at a cheaper price than the JDSU equipment. Currently getting down to the final engineering of the upgrade solution. Basically the same as the JDSU solution, but with Movaz equipment. At Haystack go from Juniper M10 to Ray Express, shifts wavelength from 1310 nm to an ITU grid frequency. Go into Cisco 1500 at this end, then to 230 Congress St through amplifiers, back up to Lincoln Labs to Cisco 1500 there which then feeds into Ray Express. This Ray Express converts wavelength back to 1310 nm and then feeds it back into another Ray Express which would convert it up to another wavelength that could be transmitted over Bossnet. From there it could go into a Firstwave optical switch or into one of Terry's home grown transponders. Then across Bossnet to Eckington where there is another Firstwave switch and then into the DRAGON network. That is currently where we are at the moment. The last email that went around from Jerry was trying to arrange a meeting with Movaz engineers and with engineers from Lincoln Labs to do the final engineering. Terry did you have anything to add to this?

Terry Gibbons: this was pretty much it. I expect to have a Firstwave switch at Lincoln labs. Hesitation is only that is being provided by a different program. So if there is one at Eckington, there should be one at Lincoln Labs as well.

---

Alan Whitney: we are looking forward to getting this equipment installed. Things still have to settle down on the Washington end with the reconfiguration that Jerry is doing for DRAGON.

Alan: Lee or Mary are there any configuration changes down at Goddard that we should be aware of.

Lee Foster: none that they I am aware of.

Mary Shugrue: none that I can think of.

Alan: any more information on the 10 Gbps testing going on down at Goddard?

Mary: Paul and Bill were working on this. Nothing new from them to report.

Lee: I haven't seen any memos on this.

#### e-VLBI Experiments

Kevin Dudevoir: Seems to be progress in Hawaii. Testing again this week. Reconfiguration at Hawaii. They have a link from Kokee to Makaha ridge and from Makaha ridge down to PMRF. They have been working on the second link to PMRF. Now seeing symmetric path from PMRF to Kokee. 80-90 Mbps UDP with very low loss rates. Still seeing better but poor TCP performance. 10 Mbps TCP performance. Could be a tuning issue at Haystack end. Link much improved going to Kokee. Used to be 10 Mbps. Now seeing close to 155 Mbps. Main problem is that link has been down quite a bit. Roger Hall should have it back up and stable for us to use by the end of this week. He has asked me to do some additional testing. Should have it back by the end of the week. He also has some suggestions on how we could improve it, but it could cost money. It is fiber to Makaha ridge and then an additional radio link to PMRF. A lot of competitive traffic. Competition is on radio link. Appears to be congestive losses rather than losses due to radio link. TCP performance still looks bad despite good UDP performance (0.0001%). Often 0 loss over long testing intervals (TCP rate limited at 60 Mbps). We should be able to use UDP based transport protocol (such as Tsunami or UDT). Germany is the same status. Up and ready to go. Still at 30 Mbps.

David: Last experiment to Onsala was on the 26th April. This was the second experiment. Slightly different from previous two. Decided to look at multiple sources rather than just single source with 20 minutes scans. Basically schedule we had setup had 45 scans, 3-4 minutes per scan. Signal to Noise Ratios were a lot less than for the previous experiment, which means that is harder to get fringes. Able to get fringes for about 30 minutes. Took a while to get transport of data up and running. Once this was up and running everything went smoothly. This was real-time with no disc buffering at all. Data from Onsala and Westford into the correlator. It was nice to see the scans coming in constantly and generating new scans plots with minimal intervention from the operator. Of course, that was after Mike had spent 6 hours putting together the correlator setup. There was a lot of work to get it to that point. We did get 64 Mbps from Onsala and were able to get an autocorrelation for that. Weren't quite able to get a cross correlation. Next time we have a correlation, which is in June, we should be able to bump the rate up.

---

David: another non-real-time experiment has been completed. One of the regular monthly ones, T2028 was completed the week before last as well. These experiments are continuing. An additional station was added to that. This was Tsukuba in Japan. They are connected at 622 Mbps. So that is quite a high speed connection. Currently we are still using bbftp we were transferring from 4 separate machines at the same time and getting an aggregate of 60 Mbps. One of the things that we would like to do is move away from using bbftp and use something a bit more aggressive without being too aggressive. That is what we will be doing in the future. The 622 Mbps is the link that Tsukuba has in to the GEMnet network and then across to the US. The pipe to the US is 2 x 2.5 Gbps. The net transmission rate was 60 Mbps using TCP. One reason we are not seeing higher transmission rates is that the boxes we are using over there are FreeBSD boxes. The tuning that is slightly different from tuning Linux boxes. One of the parameters that needed to be tuned had to be done at compile time - it controls the total amount of memory available to TCP socket buffers. Wasn't able to do that. Didn't have time or expertise at the other end to compile the kernel and bring it back up again. So we had to make the best use we could of the available memory, which was quite low. I believe it was less than 1 MB per machine for the send side buffer.

### Performance and Testing

Kevin: I am in the process of adding two new stations to the performance monitoring: Kokee and Germany. He is in the process of doing this. The only problem is that stations are quite hesitant to have testing during their experiments, so we are trying to work around this.

### Miscellaneous

Kevin: We have lost connectivity to Kame at ISI-E. This is one of the main servers that we use for these experiments while we are circumventing Bossnet. I haven't been able to log into it.

Paul Lang: would it be possible to get a 10 Gbps link from Haystack to ISI over Movaz. I have some ideas for doing this. I have been talking with Henry Orejuela about the use of the external wavelength module (about \$5K). You would also need a \$2k filter in addition. The only need for using this is to make sure that nobody plugs in and additional wavelength and messes things up. Found some DWDM Xenpacks and SFPs and are interested in going this route. Also interested in getting 10 Gbps transponder (\$22K each). Our basic concern is getting to McClean. Pat is also involved with this budget related discussion. DWDM Xenpack price was \$7500 which is cheaper than the \$8000 for the ZR Xenpack. Price for DWDM GigE SFP is quite reasonable. Haven't found a DWDM XFP which is 10 Gbps. DWDM Xenpack was from Opnext. We will try to get it off "Soup" (government agencies can buy from here). If you use external wavelength module on the Movaz, the DWDM Xenpack would be one way of providing your own DWDM wavelength 100 GHz spacing. Henry also had some other issues he wanted to clarify about using an additional wavelength. Could potentially save an additional \$10k.

Hans Hinteregger: is it possible to use 1310 nm 40 km reach XFP.

Paul: Movaz uses 1310 nm wavelength for their management. Jerry seemed a bit concerned about

---

adding other wavelengths although 1310 nm was less of a concern.

Bill Fink: Jerry doesn't want to add a MUX/DEMUX.

Paul: the MUX/DEMUX would be needed to strip out a wavelength before it goes into the switch. Because it doesn't understand the 1310 nm. Jerry didn't want to add a MUX/DEMUX between the Movaz equipment. This would add additional complications.

Bill: 1310 nm probably wouldn't make the distance.

Hans: there are 40 km XFPs available for about \$1500.

Paul: Jerry had indicated that when you get to the higher speeds like 10 Gbps even if distance is not a problem, there could be dispersion and might have to add extra equipment within the link to take care of that problem as well. We have the Force10 E300 hooked up at the MAX and Bill was starting to run some tests. Started to hook it up just over a week ago. The E300 with 10 Gbps across 1310 nm mux-ed in with our coarse CWDM equipment. They have some VLAN loops and are maxing out the 10 Gbps link between the 2 of them.

Bill: with TCP streams we have about 95% utilization of the link and with UDP streams have about 99% utilization bi-directional across the 10 Gbps link. A couple of weeks ago, I managed with our testing of the 10 Gbps Intel NIC to bump our throughput up to the 5 Gbps rate over the 10 Gbps card. That was with a dual Xeon 3.06 processor server.

Alan: has any of the 10 Gbps testing been reported in memos?

Bill: I had sent out an email, not sure if it has been forwarded but I will forward it to you.

Paul: I am quite interested in hooking up to ISI and would be glad to work with them on a solution.

Bill: we are trying to improve the connection between GGAO and Haystack up to 10 Gbps.

Paul: might be fairly doable at least from our end to ISI.

David: this would be great. The only thing that might limit this is whether it is possible to get that rate over Bossnet. Getting the cards shouldn't be a problem. When he had spoken to Terry before Terry had reservations about how high it is possible to push the bandwidth on Bossnet. I think that is the only problem.

Paul: is there a possibility of feeding a 10 Gbps DWDM 100 GHz spacing into Bossnet?

David: I would have to check with Terry on that. Technically I can't see any problem.

Paul: if there is any particular wavelength, let us know. We are about to order this soon. After that it

---

will take 4–6 weeks to go through our procurement system. But we need to get things sorted out.

David: would you like to take this offline and shoot an email around talking about this?

Paul: sure.

Next telecon is scheduled for Monday, 24th May 2004 at 2 pm EDT.

---

cc: Steve Bernstein, LL  
Jim Calvin, LL  
Rick Larkin, LL  
Lorraine Prior, LL  
Peter Schulz, LL  
Leslie Weiner, LL  
Herbert Durbeck, GSFC  
Bill Fink, GSFC  
Lee Foster, GSFC  
Pat Gary, GSFC  
Andy Germain, GSFC  
Chuck Kodak, GSFC  
Kevin Kranacs, GSFC  
Paul Lang, GSFC  
Aruna Muppalla, GSFC  
Mary Shugrue, GSFC/ADNET  
Bill Wildes, GSFC  
Dan Magorian, UMCP  
Tom Lehman, ISI-E  
Jerry Sobieski, MAX  
Guy Almes, Internet2  
Charles Yun, Internet2  
Richard Crowley, Haystack  
Kevin Dudevoir, Haystack  
Hans Hinteregger, Haystack  
David Lapsley, Haystack  
Jason Soohoo, Haystack  
Arthur Niell, Haystack  
Joe Salah, Haystack