# Draft Month 7 Demo

Mh: Jan Wagner, Guifre Molera, Jouko Ritakari, Ari Mujunen,
Jb: Paul Burgess

14th November 2006

**Abstract**

For the EU EXPReS project a test was planned on 20th October 2006 to demonstrate the use of PC-EVN computers and Internet data transfer to obtain fringes. Thanks to the collaboration of the institutes involved in the experiment (Metsähovi, Onsala, Jodrell Bank and Jive) it was possible to achieve realtime telescope data transfer over the Internet and correlate selected parts of PC-EVN data to get corresponding fringes. This demonstrates the possibility of low-cost high throughput realtime VLBI data with PC-EVN microcomputers and the Internet, and advances the possibilities in further EXPReS development. To the EU the test demonstrates that EVN observatories are now capable of eVLBI observations using the low-cost MRO 'pcevn' hardware design, in addition to existing Mark5 systems, which have been in use for some time.
[1]

## 1 Executive summary

The objective of the Month 7 demo was to prove that eVLBI fringes can be achieved with normal microcomputers and normal Internet at comparably low cost. The demonstration succeeded beyond expectations, the speed achieved was better than in the previous high-profile Internet eVLBI tests, and results surpassed initial Month7 basic goals. Furthermore, the demonstrated technology is easily scalable. This could result in 2-3-fold increase in sensitivity.

Several small problems were encountered. Fortunately fairly simple workarounds for them exist, so this do not effect future production eVLBI.

## 2 Test setup

The original plan was to stream 512 Mbit/s realtime data from Jodrell Bank, Onsala and Metsähovi to three computers in JIVE and correlate the data as soon as possible.

Because of several constraints (the three antennas did not have a common frequency because the Onsala 22 GHz receiver was being upgraded, Metsähovi does not have a 5 GHz receiver and JIVE did not have more than one microcomputer available for the demo) the

[1]This work has received financial support under the EU FP6 Integrated Infrastructure Initiative contract number 026642, EXPReS.

experiment was modified so that Jodrell Bank and Onsala would stream realtime data to Metsähovi. For the demonstration this was a better setup, since the Metsähovi link distance was longer and the connections were not as fast as in Central Europe.

The final guideline schedule on Friday, 20th of October 2006 was:

1. Part A - 0900 -1000 UT. Onsala and Jb record data to disk 512Mbit/s later transfer files to JIVE for correlation

2. Part B1 - 1030 - 1110. Jodrell and/or Onsala try remote recording at 256Mbit/s to disk at Metsähovi

3. Part B2 - 1120 - 1200 Jodrell and/or Onsala try remote recording at 512Mbit/s to disk at Metsähovi

4. Part C - 1230 - 1330 Jodrell and/or Onsala try remote recording at 256Mbit/s to disk at JIVE. Other station records to local disk and transfers later. Data converted at JIVE for correlation.

Selecting sources for the observation and creating a schedule for the stations was done by (JIVE?).

To allow PC-EVN data recording and streaming, Jb and On were equipped by Mh with PCI VSI Boards and a VSI Converter. Jb and On built their own test PC systems, and installed and customized the Mh VSIB reference system based on Debian 3.1 Linux, with precompiled patched kernel and tools necessary for the Month 7 demonstration.

Regarding network connectivity, Jb and On stations had one dedicated test PC each, connected directly to the Internet through a 1G fibre, with appropriate routing changes made by Jb to allow access from Mh, too. JB used lightpath dedicated links provided by UKLight, Surfnet and Geant. Mh used two of their dedicated test PCs behind a 10G fibre, with a connection to FUNET. It was known that at the time of the experiment there still was a 2.5Gbit/s bottleneck in the Metsähovi link, between the university campus and the center of Helsinki. Jv provided access to a shared-use computer behind a 1G fibre. JIVE also provided their correlator facility and Mark5's for the fringe checks.

All five test PCs had fast RAID disks that were adequate for recording observation data at high data rates.

## 3 Pre-experiment network tests and preparations

A couple of days before the experiment we tested the capacity of the backbone network by first testing the connection with the iperf tool, testing the PC-EVN system performance in both radio observatories at On and Jb, and finally by streaming realtime VSI-H test data from the PC-EVN systems to Mh using realtime Tsunami protocol, a protocol extended for realtime eVLBI by Metsähovi. The Tsunami transfers succeeded consistently at 720 Mbit/s (one station at a time), which gave a comfortable margin above the 512 Mbit/s target.

Because everything looked just fine with a 0.0% error rate reported by Tsunami, we at Mh decided to try a very short test to transfer at an aggregate rate of >1.4 Gigabit/s from the two stations simultaneously.

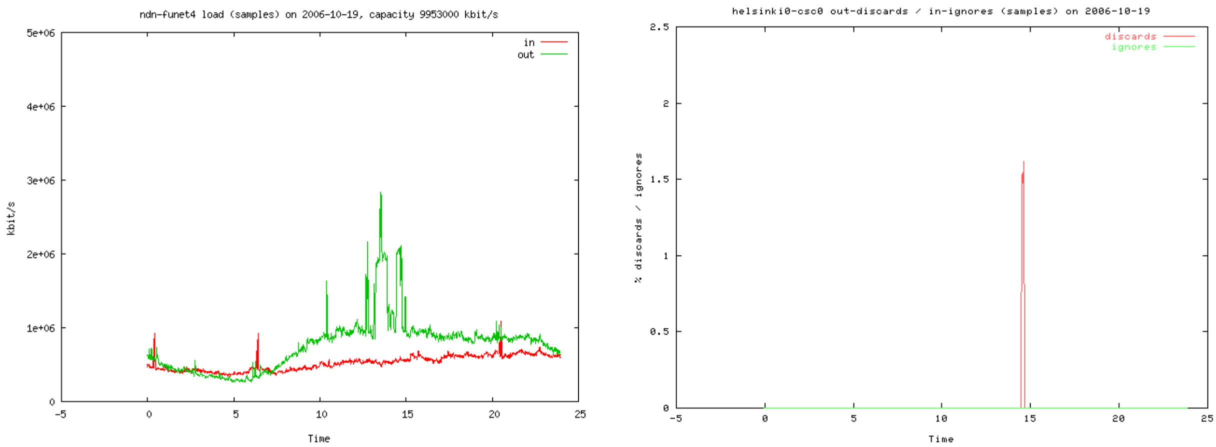


Figure 1: Nordunet traffic rate - Traffic peak error 19/10/2006 - Image courtesy of Nordunet

You may note that we tripled the academic network traffic between Finland and the rest of the world.

Unfortunately the 2.5Gbit/s bottleneck could not tolerate this and caused a small packet loss when we tried to push through 2Gbit/s in addition to the normal traffic. The good thing was that 1Gbit/s traffic was not even noticeable.

The results of the Mh pre-experiment were several, and included not just network traffic results.

First, a number of cosmetic Tsunami issues were found during the tests. This lead to Mh improving the Tsunami programs for final production and then delivering these to all Month7 participants. The Tsunami client was improved such that it could be easily called from a observation schedule shell-script. NTP configuration issues and other small problems found at Jb and On were detected and quickly resolved by fast cooperation with Jb and On staff.

Second, the pre-experiment allowed Mh to correct and fine-tune the documentation and instructions, and to complete the PC-EVN data acquisition related scripts and schedule scripts before passing them on to On, Jb and Jv where they were tested and further customized.

Third, the tests through the current Mh-external bottleneck confirmed that there had been no degradation in available bandwidth compared to earlier tests, and that there was still ample of headroom to achieve the 1 Gbit/s maximum total throughput of the Month7 demonstration.

Fourth, as the pre-experiment progressed, the people at Jb, Jv and On station grew more familiar with the VSIB data capture tools, the VSIC formatting hardware and with using the Tsunami protocol. All additional PC-EVN and VSIC configuration required by the observation schedule and data transfer was worked out, tested, and found operational. This paved the way for a successfull Month7 demo.

## 4 Experiment on 2006-10-20

The experiment was divided into three parts: In part A On and Jb would record locally at 512Mbit/s, in Part B they would stream data first at 256Mbit/s and then 512 Mbit/s to Mh, while simultaneously making backup copies of the data to prove via later file comparison that the transfers were flawless. Finally in Part C either the On or Jb station would stream data to JIVE at 256 Mbit/s, with the other station recording locally to disk.

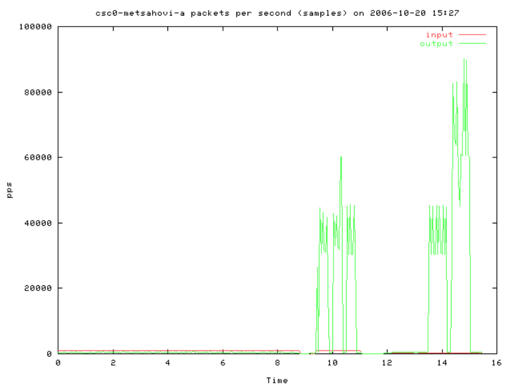


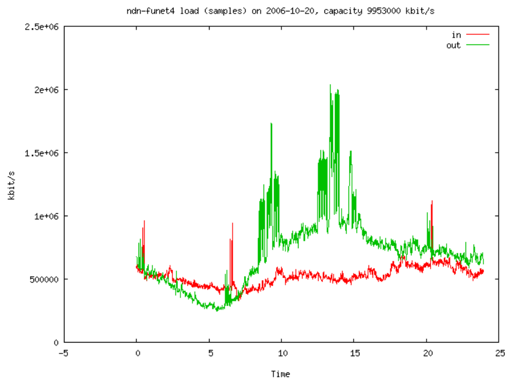


Figure 2: MRO output traffic and the Funet connection rate to Sweden - 20/10/2006 - Image courtesy of CSC, the Finnish IT center for science, operator of the Funet network

Part A succeeded so well that it was decided to abort it already at an early stage and transfer one of the completed scans to JIVE for immediate correlation to check if everything worked. Tranfering the data with Tsunami was successful, at faster than recording speed. Correlating succeeded also, but was delayed because of difficulties transferring the data into the old Mark5 units.

Part B as we can see in the plots above 2 was a complete success at 256Mbit/s speeds. At 512Mbit/s the network link between Jb and Mh shortly slowed down for some reason in scan 3. The transfer time was prolonged since data integrity, in contrast to maintained rate with accepted data loss, was an emphasis in the Month7 test. Thus the scan 3 transfer extended well into the 60 second scan margin, such that scan 4 could not be started in time and was skipped. Fortunately a relatively simple protocol change can prevent this from happening in future transfers. All other 512Mbit/s transfers completed flawlessly.
From the graph2 you can see that the transfer rate fell inexplicably during scan #3 in the 512 Mbit/s part but recovered later. It's also visible that one of the pre-experiment checks suffered from packet-loss (see the spike at 1010) which did not affect the transfer. In the latter parts of the 512 Mbit/s experiment it's also possible to see the effect of a small packet-loss.

In the histogram4 we see a time series of the transfer rate (in Mbps) from the failed scan#3 from Jb to Mh. It is noticeable that the rate is not so smooth around the average

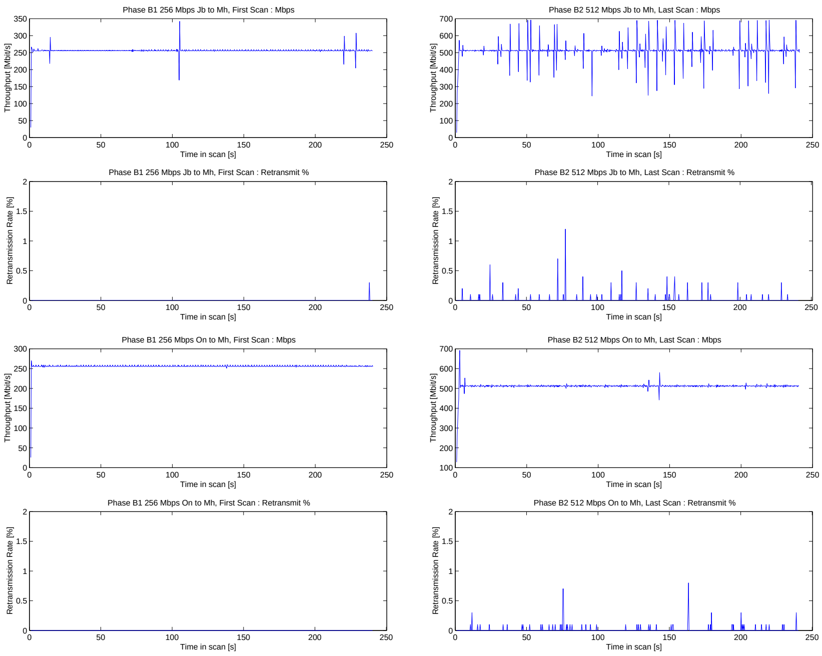


Figure 3: First scan at 256 Mbps and last scan at 512 Mbps

rate of 512 Mbps as in the Figure 3. The link from Mh to On did not exhibit much rate variance at all. We suspect that on the longer Jb Mh link other Internet traffic caused the short drops in link throughput.

Later while examining B1 256 Mbps and B2 512 Mbps Jb scans on their Mark5, Jv found short time gaps in the recorded data. The total amount of data was correct but due to the gaps the scan files covered a longer time, with several gigabytes of unusable data recorded outside the schedule. Jv's information helped to locate a buffer configuration problem at Jb that had caused these gaps in the data and had gone unnoticed in the pre-experiment tests. Jv could still successfully get fringes for the first B1 scan.

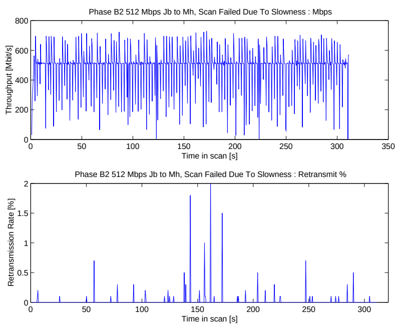


Figure 4: Transfer rate samples from the scan #3

Part C was close to getting cancelled due to the unexpected problems in data transfer from the Jv PC to the Mark5. Luckily Jv managed to complete the transfer just in time before Part C, and Part C could be started.

It was then decided to perform Part C differently than planned. Jv would still do realtime recording with Jb, but additional realtime streaming would be done from On to Mh instead of doing On local disk recording only.

Unfortunately Part C of the experiment failed. The realtime transfer from Jb to Jv met several difficulties with the Tsunami client in JIVE, despite earlier successful transfers of some Part A scans to JIVE by Jv staff. What caused the problems is still unclear [?], but they might be attributed to changes made to the Jv computer's configuration while JIVE was attempting to transfer data from that computer to a Mark5 in the free time slot before Part C.

At Mh, while the Part C On to Mh streamed scans were largely successful, two scans slowed in mid way and were subsequently cancelled, because the corresponding scans at Jv had failed to start at all. Continued local and online attempts to get the Jv scans operational failed. Thus no fringes could be obtained for Part C.

# 5 Conclusions

[– others can write their own conclusions, and someone else will decide about what finally is politically correct to include in the final report –]

The best results of previous high-profile eVLBI experiments can be duplicated using normal commercial off-the-shelf microcomputers. Aside from COTS hardware ensuring low cost and easy upgradeability, the best about this is that the system is scalable, to the limits of network capacity. In a couple of years the network will transfer more than we can capture from the sky, meanwhile we have to install some simple safeguards in our protocols.

Currently the main conclusion however is that the objective of the Month 7 demonstration has been met. The next goals of EVN development in EXPReS appear easy to reach.

Causes of the temporary Internet link slowdowns in part B2 and C may need further investigation. Alternatively, the simple Tsunami provision for coping with such situations should be tested.

In addition, the transfer problems in the slow final few metres from the Jv computer to the Mk5 still need addressing, if current Mk5's should be used in production eVLBI. Barring any signs of cooperation from Haystack to help resolve such issues, it might be very worthwhile to reconsider what actually had already been considered for the Month7 demonstration, that is, perform correlation or fringe checks using the JIVE software correlator instead of Mk5's and the hardware correlator.

So what is next?

It is clear we want to extend the tests to a 1024 Mbps transfer rates, hence upcoming tests are oriented to stream telescope data from PC-EVN computers in JB to Metsähovi.

For that, we expect to parallel up two commodity microcomputers at both ends and use Tsunami in real-time at 512 Mbps between each microcomputer pair.

The tests with PC-EVN have has demonstrated datarates that can easily be achieved with the Tsunami UDP protocol over the Internet with low cost hardare. The same lightweight protocol should be considered for embedded PC-EVN applications that are planned for the future, such as in the iBOB project, instead of designing entirely new protocols. A common fast protocol that has low requirements on hardware would make it easier to develop disributed processing of realtime data in embedded applications, such as an iBOB data packetizer or a computing grid for correlation using inexpensive COTS hardware, etc.