Real Time Data Transfer for Very Long Baseline Interferometry

Simon Casey, Richard Hughes-Jones, Stephen Kershaw, Ralph Spencer, Matt Strong
The University of Manchester, UK

***Abstract*—This paper examines the use of the Internet in real time Very Long Baseline Interferometry (VLBI) high resolution observations in radio astronomy. Owing to the long distances between telescopes, VLBI data have traditionally been recorded to magnetic tape, and now onto hard disks. With the emergence of high data rate network links the VLBI community is beginning to transfer astronomical data via the academic internet. The current system uses TCP/IP and we report on the way TCP transports the constant bit-rate VLBI data. We also give details of a new, UDP based VLBI data transfer system, VLBI_UDP, and results of recent tests, including multiple simultaneous line-rate UDP data flows lasting several hours running over the academic internet and lightpaths. Finally, we give examples of successful real-time astronomical observations using e-VLBI with data transfer rates of up to 256 Mbit/s.**

***Index Terms*—, e-VLBI, CBR, TCP/IP, UDP/IP Gigabit Ethernet.**

I. INTRODUCTION

Very Long Baseline Interferometry (VLBI) [1] is an aperture synthesis technique that utilizes radio telescopes from around the world, combining astronomical data in order to achieve high angular resolution observations. The telescopes observe the same cosmic radio source simultaneously and the use of stable atomic clocks means that signals can be added coherently. The interferometer technique allows signals from each pair of telescopes to be multiplied together in a correlator to recover the cross-correlation signal out of the noise.

The angular structure of radio sources can be obtained by Fourier inversion of the correlated data taken over a number of baselines. The number of Fourier components is effectively increased (and hence improved imaging fidelity) by means of Earth rotation, where the projected baseline length varies with time, tracing out part of an ellipse over ~12 hours [2], and the angular resolution is inversely proportional to the maximum length of the baseline, so trans-world arrays produce images showing the finest detail.

The sensitivity, or signal to noise ratio, is proportional to $\sqrt{B\tau}$ where B is the bandwidth and τ is the integration time, so bandwidth is as important as time, giving an obvious need for high data rates.

VLBI data have traditionally been recorded locally at the telescopes on magnetic tape, but now the MarkV [3] hard disk system is used. With the emergence of high data rate network links the VLBI community is beginning to transfer astronomical data via the academic internet.

e-VLBI has a number of advantages, in particular the ability to check that everything is working immediately rather than the wait of several weeks or even months for tapes to be shipped and correlated. This results in a significant increase in reliability for VLBI operations. In addition, the technique lends itself to rapid reaction observations where subsequent detailed observations on rapidly varying objects can be decided upon within hours rather than weeks after initial measurements are made

II. THE E-VLBI NETWORK IN EUROPE

Figure 1 illustrates how the European radio telescopes used in these tests link to the correlator at JIVE in the Netherlands using the National Research Networks (NRN) of each country and GÉANT2, the European research network backbone. Currently, most of the instruments use the packet switched academic network, but dedicated end-to-end lightpaths are now becoming available across the infrastructure, and one of the aims of the EXPReS project [4] is to install lightpaths to the European telescopes. At present, Jodrell Bank connects to JIVE using two 1 Gigabit lightpaths over UKLight and Surfnet; and JIVE is connected to Amsterdam with eight 1 Gigabit lightpaths from Surfnet.

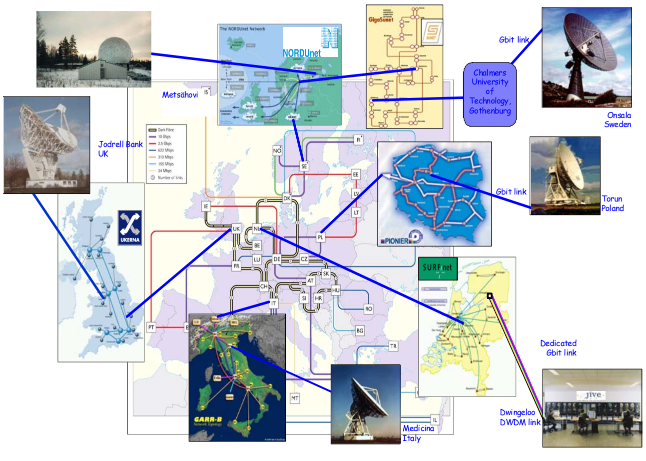

Figure 1. Part of the VLBI network in Europe used in the transport tests.

e-VLBI is now able to perform routine real time VLBI measurements at data rates of 256 Mbps and 512 Mbps tests are underway.

III. TRANSPORTING VLBI CONSTANT BIT RATE DATA WITH TCP/IP

A. *The Behaviour of TCP when packets are lost*

Signals from the radio telescopes are digitised at each telescope using precisely synchronised clocks, and then formatted into a stream of constant bit-rate (CBR) data which needs to be transported across the network.

TCP is a byte-stream transport layer protocol that guarantees reliable, in-order, and non-duplicated delivery data from sender to receiver. It uses acknowledgments (ACK) from the receiver to slide a window (Cwnd) over the data to regulate the transmission rate. If a packet has been lost, TCP interprets this as congestion on the network, and the standard congestion avoidance algorithm, known as New Reno, decreases the window by half and then slowly increases it by one packet per round trip time.

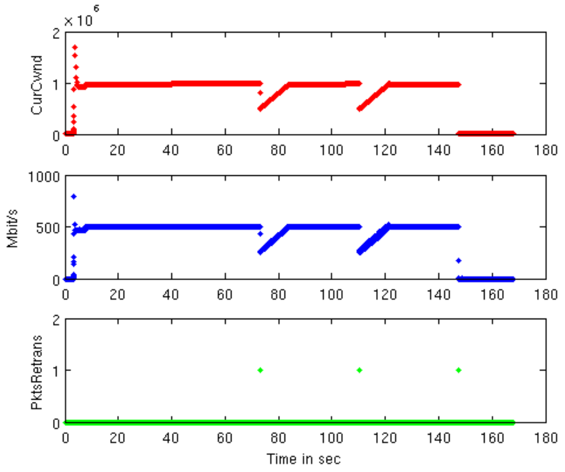

Figure 2. Plots of TCP parameters using Web100 when packet loss is present.
Top: TCP Congestion window
Middle: Achievable throughput Mbit/s
Bottom: Number of packets re-transmitted.

Plots of the TCP parameters taken from the web100 interface to the TCP stack [5] are shown in Figure 2 and demonstrate the congestion avoidance behaviour in response to lost packets. This TCP flow was set up between Manchester and JIVE with a TCP buffer size of 0.9 Mbytes which is equal to the bandwidth-delay product, BDP, for a round trip time, RTT, of 15.2 ms. It is normally recommended to set the TCP buffer size to the BDP [6] to obtain continuous data flows for file transfer applications. The packets were deliberately dropped using a kernel patch in the receiving host. The upper plot in Figure 2 shows Cwnd decreasing by half with the corresponding reduction in the achievable TCP throughput.

Notice that it takes about 20s for the flow to recover to the rate before the packet loss. The bottom plot shows that one packet is re-transmitted for each packet dropped.

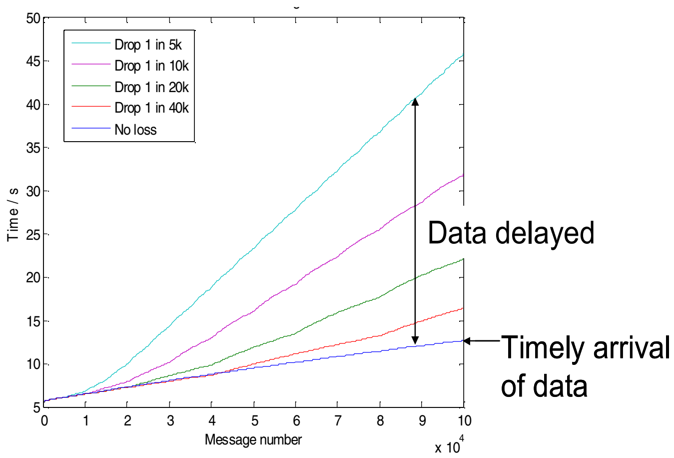

Figure 3. The effect of packet loss on message arrival time..
TCP Buffer size 1.8 Mbytes, RTT 27 ms

The packet re-transmission and the decrease in throughput delay the delivery of data to the receiving application as shown in Figure 3, which shows the arrival time of the messages sent over a path with a RTT of 27 ms using a TCP buffer size set to the BDP. For e-VLBI this delay in the arrival times is undesirable and can lead to loss of correlation.

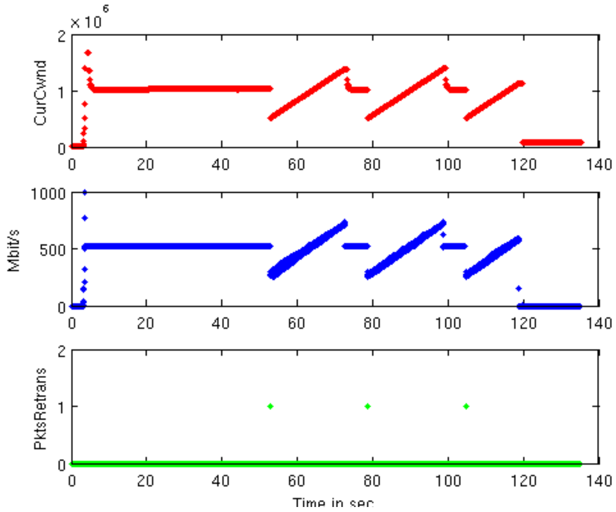

Figure 4. Plots of the TCP parameters on the 15.2 ms RTT link when a 160 Mbyte TDP buffer is used and packet loss is present.
Top: TCP Congestion window
Middle: Achievable throughput Mbit/s
Bottom: Number of packets re-transmitted.

B. *Timely Deliver of CBR data using TCP.*

If the incoming CBR data can be stored in the TCP socket buffer during the loss event and data can be sent on the link faster than the constant bit-rate, then catch-up of the data delivery can occur.

Figure 4 shows that TCP can indeed increase the sending rate under these conditions, however, the TCP buffer was 160Mbyte about 200 times greater than the normally advised value of the BDP. The use of these large buffers introduces temporary delays in the arrival times of the data of many seconds.

IV. MOVING VLBI DATA WITH UDP/IP

The architecture of the program used to move VLBI data using UDP/IP is shown in Figure 5. The operation is as follows: at the sender, the input thread either reads data from a file or generates random data in memory and places these into the ring buffer. The output thread takes packet-sized sections of data from the ring buffer, encapsulates the data in a UDP/IP packet together with an application header containing a sequence number which increments by 1 for each packet sent.

At the receiver, the receive thread places incoming packets directly into the next position in the ring buffer. The sequence number is read from the header and this reveals whether the packet is at the correct position in the buffer. If the sequence number increment is not equal to 1, then the packet is moved forwards or backwards in the buffer as appropriate. The output thread takes data from the ring buffer and writes it to disk or memory.

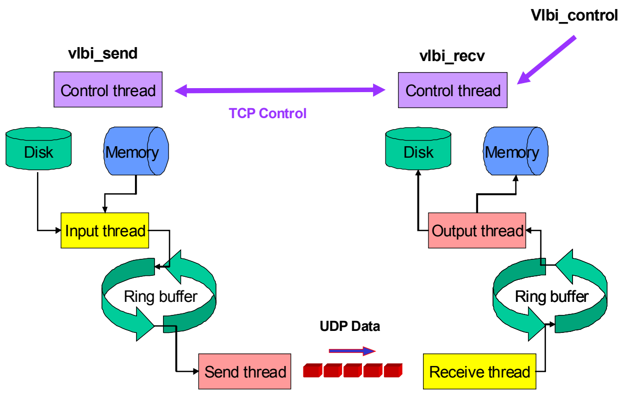

Figure 5. Architecture of the VLBI_UDP program.

In December 2006, a 3 station e-VLBI experiment was emulated by simultaneously transmitting data from 3 locations over the GÉANT2 network, shown in Figure 1, into PCs at JIVE using VLBI_UDP. For the tests shown in this paper, the following network paths and transmission rates were used to JIVE: a dedicated lightpath from Jodrell Bank at 800 Mbit/s, the packet switched network from Manchester at 600 Mbit/s and the packet switched network from the GARR Point of Presence in Bologna at 400 Mbit/s.

A comparison was made of the performance obtained for each path, the achieved throughputs and packet loss are shown in Figure 6. This shows that all of the paths were able to sustain the required throughput. There was the occasional packet loss on the path from Bologna, but up to 0.5% packet loss on the packet switched path from Manchester. The measurements were made during the transition to Super JANET 5 and the installation of a 10 Gigabit access link to Manchester. The loss was due to the temporary bonding of 1Gigabit links, all being utilised at about 50%. The absence of packet loss on the top plot clearly shows the superior performance of the UKLight lightpath when compared with the packet switched production network.

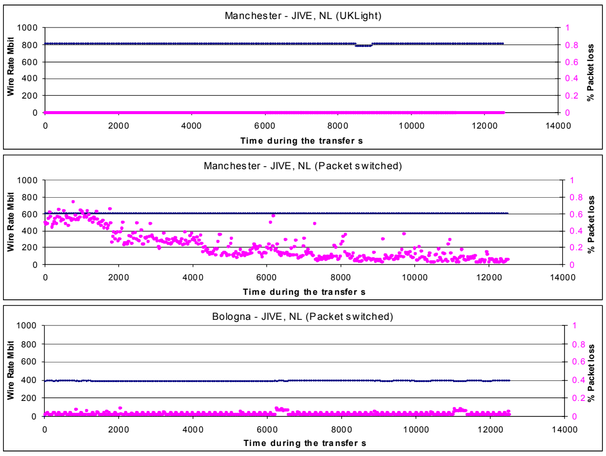

Figure 6. Plots of the throughput and packet loss for the three UDP/IP flows between the telescopes and JIVE.
Top: Jodrell using a dedicated Lightpath
Middle: Jodrell using the packet switched network
Bottom: Bologna using the packet switched network.

V. E-VLBI SCIENCE

The first scientific results from real time e-VLBI were published in 2007 [7][8] and real time e-VLBI science observing sessions have become a regular part of European VLBI operations.

Figure 7 shows an image of Microquasar GRS1915+105 (11 kpc) taken on 21 April 2006 at 5 Ghz using 6 EVN telescopes during a weak flare [7]. This object has jets of material which move away from an accretion disk surrounding a central black hole at velocities close to that of light. The jets were quiescent in these observations.

Figure 8 shows Microquasar Cygnus X-3 (10 kpc) observed on two days, plot (a) was taken on 20 April when the source was in a semi-quiescent state and plot (b) on 18 May 2006. when it was in a flaring state. The core of the source is probably ~20 mas to the North of knot A. [8].

An illustration of the advantages of real time VLBI occurred recently when the first rapid response experiment was undertaken [9]. Here a 6 telescope real time observation was run on 29th Jan 2007, the results were analysed in double quick time, selecting sources for follow up observations on 1stFeb. This kind of operation would be impossible for conventional VLBI. The experiment worked well and we successfully observed 16 sources (weak microquasars), but all were <0.5

mJy – too weak to observe in the follow up run – indicating a perverse universe, however the feasibility of the technique was clearly demonstrated.

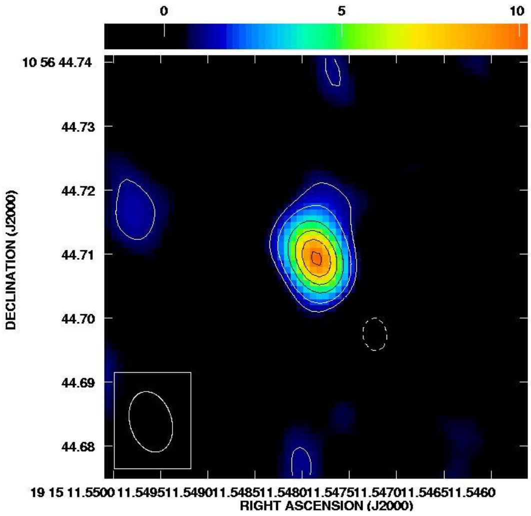

Figure 7 e-VLBI observation of the Microquasar GRS1915+105

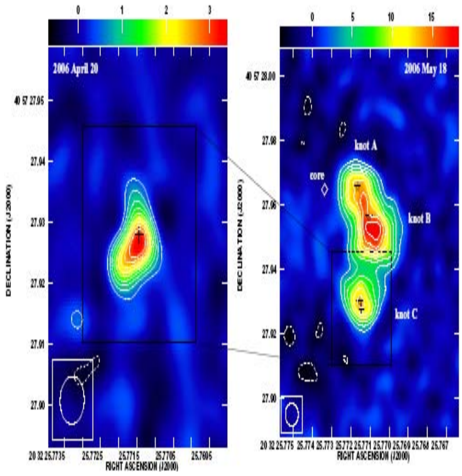

Figure 8 e-VLBI observation of the Microquasar Cygnus X-3
(a) Left plot: source was in a semi-quiescent state.
(b) Right plot: source in a flaring state.

VI. CONCLUSIONS

We have briefly described the network connectivity of the radio telescopes and the correlator in Europe that form e-VLBI and its requirement to move constant bit-rate data.

Investigations of moving CBR traffic using TCP/IP in an environment with packet loss have shown that using the normal buffer settings for TCP, each packet loss event will introduce a delay in the arrival of all subsequent information at the receiving host. For e-VLBI this delay in the arrival times is undesirable and can lead to loss of correlation. If the incoming CBR data can be stored in the TCP socket buffer during the loss event, and data can be sent on the link faster than the CBR rate, then catch-up of the data delivery is possible. However, extremely large buffers are required and there may be temporary delays of many seconds.

Our tests with VLBI_UDP indicate that the required throughput can be achieved along with timely delivery of the data. With the current networks, there is little packet loss in the Academic Internet, making the UDP transport protocol suitable for e-VLBI.

The first scientific real time e-VLBI results were published in 2007 and real time eVLBI science observing sessions at data rates of 256 Mbps have become a regular part of European VLBI operations. 512 Mbps tests are underway.

This work was conducted as part of the ESLEA project [10] as part of protocol research and eVLBI work-packages, in collaboration with the European Commission FP6 EXPReS project.

REFERENCES

[1] J. Moran *Very Long Baseline Interferometer Systems*, Methods of Experimental Physics, 12, 174, 1976.

[2] A. Thompson, J. Moran, and G. Swenson, *Interferometry and Synthesis in Radio Astronomy*, 2nd ed. Wiley, 1991

[3] Alan Whitney, *Mark 5 disk-based gbps vlbi data system*, viewed 19th July 2006, http://web.haystack.mit.edu/mark5/paper.pdf

[4] EXPReS = Express Production Real-time e-VLBI Service Three year project, started March 2006, funded by the European Commission (DG-INFSO), Sixth Framework Programme, Contract #026642 www.express-eu.org .

[5] web100 interface to the TCP stack Web100 Project home page, http://www.web100.org/

[6] TCP tuning information may be found at: http://www.ncne.nlanr.net/documentation/faq/performance.html and http://www.psc.edu/networking/perf_tune.html

[7] Rushton, A.; Spencer, R. E.; Strong, M.; Campbell, R. M.; Casey, S.; Fender, R. P.; Garrett, M. A.; Miller-Jones, J. C. A.; Pooley, G. G.; Reynolds, C.; Szomoru, A.; Tudose, V.; Paragi, Z., *First e-VLBI observations of GRS1915+105,* MNRAS 374, L47, 2007

[8] Tudose, V.; Fender, R. P.; Garrett, M. A.; Miller-Jones, J. C. A.; Paragi, Z.; Spencer, R. E.; Pooley, G. G.; van der Klis, M.; Szomoru, A., *First e-VLBI observations of Cygnus X-3,* MNRAS, 375, L11, 2007

[9] A. Rushton, R. E. Spencer, EVN e-VLBI project RR001, April 2007.

[10] Exploitation of Switched Lightpaths for eScience Applications project (www.eslea.uklight.ac.uk)