A Simulation model for e-VLBI traffic on network links in the Netherlands - Extended Abstract.

J. Sansa
Kapteyn Astronomical Institute
Postbus 800
9700 AV Groningen
The Netherlands
sansa@astro.rug.nl

A. Szomoru
Joint Institute for VLBI in Europe
Postbus 2
7990 AA Dwingeloo
The Netherlands
szomoru@jive.nl

J.M. van der Hulst
Kapteyn Astronomical Institute
Postbus 800
9700 AV Groningen
The Netherlands
vdhulst@astro.rug.nl

November 28, 2006

Keywords
electronic-Very Long Baseline Interferometry (e-VLBI), radio astronomy, transport protocols, simulation

1 Introduction

The European VLBI Network (EVN) is an array of radio telescopes located throughout Europe and as far away as China and South Africa. These radio telescopes produce data at rates of up to 1 Gbps each. Until recently, these data streams were recorded on tapes, nowadays on hard disk drives, and shipped to the correlator located at JIVE, the Joint Institute for VLBI in Europe, in Dwingeloo, the Netherlands. During the last few years JIVE, in collaboration with the European National Research Networks and the pan-European Research Network GEANT, has worked on a proof-of-concept (PoC) project to connect several telescopes across Europe in real-time to the correlator via the Internet (electronic VLBI or e-VLBI). This project has led to an EC sponsored project called EXPReS, which over the next few years will transform the EVN to a fully functional real-time e-VLBI network. During the PoC project it became clear that in spite of the vast capacity of the connecting networks, the actual transport of large data streams poses quite a challenge.

The Mark5 [1] application that handles e-VLBI data uses the Transport Control Protocol (TCP). By its nature, e-VLBI involves transporting huge amounts of data via the Internet over long distances from geographically dispersed telescopes to one central correlator. TCP is somewhat problematic in combination with long distance high speed links [2, 3, 4]. In order to investigate the e-VLBI data transport characteristics and the data flow limitations we used ns-2 [5, 6] to simulate e-VLBI data flows based on experiments reported in [7]. In these simulations we have gathered statistics on the packet sizes, inter-packet spacings, congestion window (CWND), receive widow (RWND), Round Trip Time (RTT), packet loss and the resulting throughput. In this paper we identify a number of bottlenecks and show the extent to which each limits the e-VLBI flow, as well as the combined effect. We propose an e-VLBI application model, which highlights the traffic generation patterns, background traffic and network limitations of such a network. This model can be used to test improvements of the underlying transport protocols.

2 An e-VLBI Traffic Model

2.1 Data Generation Characteristics

Data generation follows a nearly "constant rate" traffic pattern, which can be varied by adjusting the interval between packets. Packets of a constant size (MSS) are generated with a constant interval. Having observed in previous experiments that the e-VLBI data flow is composed of packets separated by large intervals, we vary the interval between the packets between 0 & 8 ms to find the most representative value. *Figure 1* indicates the achieved throughput for each of the intervals simulated. It shows that for larger intervals less throughput is attained.

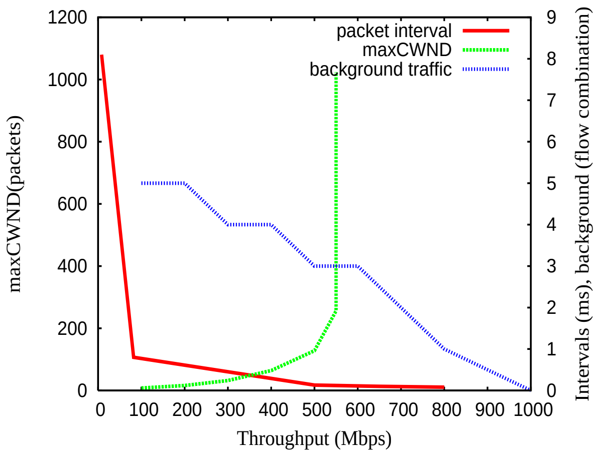

Figure 1: *Impact of each limitation on the achieved throughput*

2.2 Combined Bottleneck Effect

The traffic generation model used to estimate the e-VLBI flow is characterized by a relatively constant data generation, with larger than optimal intervals between the packets. These extra long intervals could be due to the application or the routing process. Receiver limitation is simulated by setting the maximum CWND and RWND to values estimated by the real flow, however we also simulate the flow with varying values of maximum CWND. This is shown in *Figure 1*, which shows that initially the throughput increases with increasing maximum CWND. However beyond a maximum CWND of 256 packets (0.2 Mbytes), the throughput becomes constant due to the large interval between the packets referred to above.
To reproduce the real flow, we set the maximum CWND to 64 packets (0.06 Mbytes) and RWND to the 50 packets (0.05 Mbytes).
Background traffic can vary greatly on a public network, both in amplitude and duration. In general normal internet background traffic consists of many small TCP flows using small window sizes, web traffic which is characterised by using a wide range of window sizes lasting short periods and a few long lived TCP flows with window size comparable to what is achieved by applications such as e-VLBI. *Figure 1* shows the achieved throughput for a number of simulated background traffic combinations. The plot shows, as was to be expected, that the throughput decreases with increasing background traffic.
We simulate background traffic by superposing up to four regular TCP flows, thirty small TCP flows and sixty web sessions in either direction of the e-VLBI flow.

3 Conclusions

In this paper we present an ns-2 based model that simulates e-VLBI traffic and the bottlenecks currently limiting the transfer speeds of astronomical data from radio telescopes across the world over high speed Internet links to the central processing centre in the Netherlands. This model shows that a combination of large intervals between data packets (which may be caused by application limitations or router processes), inefficient receiver hardware and excessive background traffic negatively affect the performance of e-VLBI data transfers.
Future work will include designing data generation models for other e-VLBI transfer modes (involving disk buffering).
The model presented here was based on data traces collected during e-VLBI transfers between two hosts at JIVE. We would like to validate it by studying data traces collected during e-VLBI experiments involving a

number of stations in order to avoid biases in the model due to local network conditions such as hardware and local usage patterns.
Finally, since the main goal of this work is to improve e-VLBI transport, models that shorten the intervals between packets during the lifetime of an e-VLBI data flow are to be explored.

References

[1] Mark5 vlbi data system. Haystack observatory Website. http://web.haystack.mit.edu/mark5/.

[2] S. Floyd. Highspeed tcp for large congestion windows. RFC 3649, Internet Engineering Task Force, 2003.

[3] Gary R. Wright and W. Richard Stevens. *Tcp/Ip Illustrated: The Protocols*, volume 1. Addison-Wesley, 1994.

[4] M. Allman, Paxson V., and W. Stevens. Tcp congestion control. RFC 2581, Internet Engineering Task Force, 1999.

[5] L. Breslau, D. Estrin, S. Fall, J. Heidemann, P. Helmy, H. S. McCanne, K. Varadhan, Y. Xu, and H. Yu. Advances in network simulation. *IEEE Computer*, 55(5):59–67, 2000.

[6] The ns2 simulation. www.isi.edu/nsnam/ns.

[7] J. Sansa, A. Szomoru, and J.M. van der Hulst. On network measurement and monitoring of end-to-end paths used for e-vlbi. *Chapter in Advances in Systems Modeling and ICT Applications ISBN: 13: 978-9970-02-604-3*, 2006.

Vitae

Julianne Sansa holds a BSc. (Maths, Computer Science) and MSc. (Computer Science) from Makerere University. She is also also a Cisco Certified Network Professional having acquired her training from the University of Central England, Birmingham, UK. Since 2001 she has worked with the Faculty of Computing and IT of Makerere University in various capacities. She is currently registered as a Ph.D. Student at the University of Groningen, in the Netherlands and her research interests are Network Quality of Service and Network protocols.

Arpad Szomoru obtained a PhD in astronomy at the University of Groningen, the Netherlands. He worked for three years in the USA at the UCO/Lick Observatory of the University of California in Santa Cruz. He then worked at JIVE, the Joint Institute for VLBI in Europe, for several years as a senior software scientist. During this period he was heavily involved in the deployment and integration of disk-based recording systems and the early development of e-VLBI. Since 2006 he is the head of data processor research and development at JIVE.

Thijs van der Hulst got his PhD in Groningen in 1977 on a study of neutral hydrogen emission from interacting galaxies. He spent 5 years in the USA at the National Radio Astronomy Observatory and the University of Minnesota before returning to the Netherlands, where he joined the staff of ASTRON, the Netherlands Foundation for Radio Astronomy. In 1982 he moved to the Kapteyn Astronomical Institute of the University of Groningen, of which he now is the director. His main field of interest is structure and evolution of galaxies.