eVLBI Research in the European VLBI Network

Steve Parsley, JIVE – parsley@jive.nl

eVLBI Research at JIVE and in the EVN is centred on the Proof-of-Concept project set up at the beginning of 2003. The project is supported by DANTÉ (operators of GÉANT, the pan-European R&E network) and a number of European National Research and Education Networks (NREN).

The following organisations are participating:

– DANTÉ/GÉANT:	Pan-European Research Network
– GARR:	Italian NREN
– UKERNA:	UK NREN
– PSNC:	Polish NREN
– DFN:	German NREN
– SURFnet:	Dutch NREN
– KTHNOC/NORDUnet:	Nordic NREN
– University of Manchester:	Network application software
– JIVE:	EVN Data Processor
– Westerbork telescope:	Netherlands
– Onsala Space Observatory:	Sweden
– MPIfR:	Germany
– Jodrell Bank:	UK
– TCfA:	Poland
– CNR IRA:	Italy

For each group the objectives are as follows:

For the EVN:

- Evaluate feasibility of eVLBI
  - Interfaces: network, correlator, field system
  - Network installation cost and timescale
  - Data rates
  - Logistics

- Standards
  - Network protocols
  - Parameter tuning for high data rates
  - Procedures at telescope and correlator

- New Capabilities
  - Higher data rates
  - Improved reliability
  - Quick response to transient phenomena.

For the NRENs:

- To see significant network usage with multiple Gbit streams converging on JIVE from diverse European sources.
- Up to 5 telescope sites supported (6 including Westerbork), for the period up until the end of the current GÉANT contract (Nov 2004). Minimum is three telescopes (not including Westerbork)
- Must be seen to enable new science and not just solve an existing data-transport problem
- Real-time operation is seen as the ultimate aim, buffered operation accepted only as a development stage
- Support will be:
  - best effort IP service transiting NRENs and GÉANT
  - no significant GÉANT upgrades (at least initially)
  - aim to support 512Mbps and 1Gbps modes of operation
  - use existing NREN access ports
  - limited resilience.

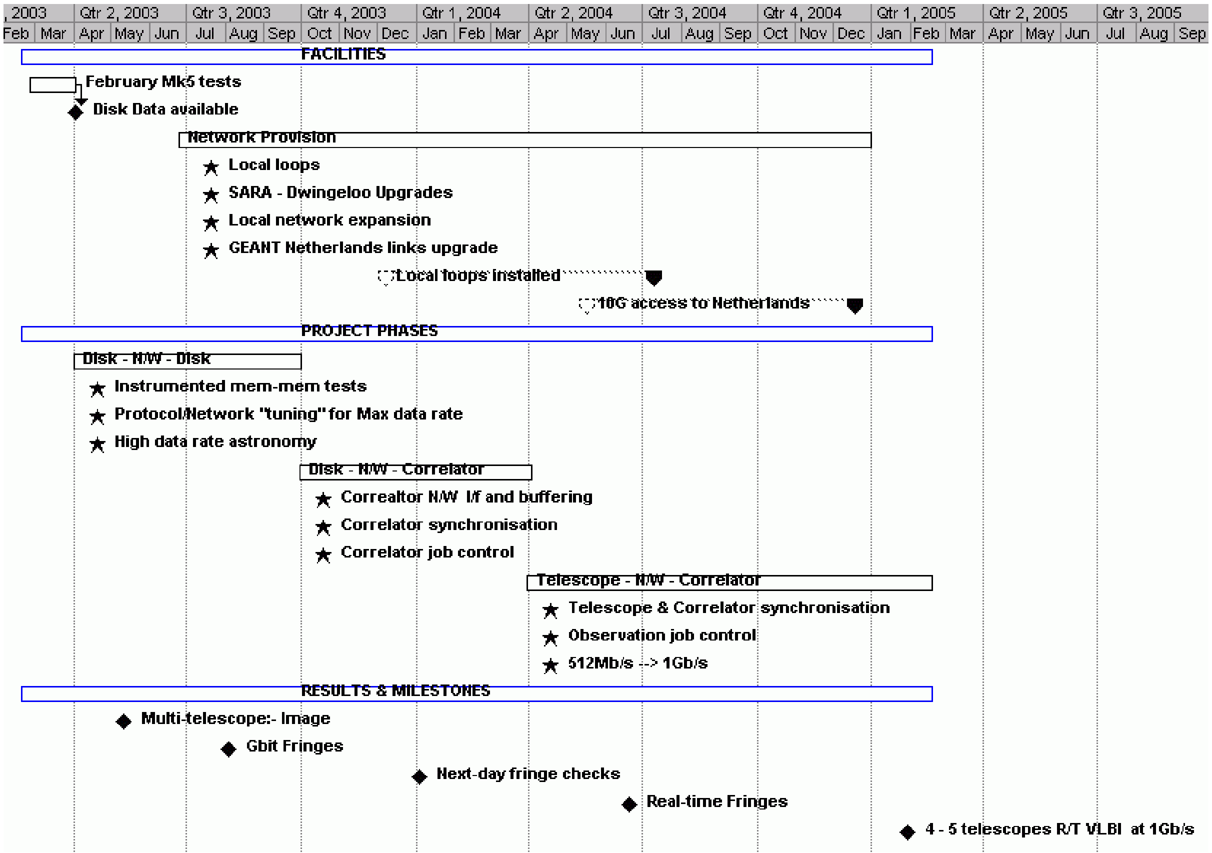

Figure 1: Project Time-line

Project status

In the EVN, deployment of Mark 5 has provided a uniform interface to telescopes, correlator and the networks. The ability to record data on disks enables various eVLBI tests to be performed on and off-line without disruption to normal operations and, in some cases, ahead of the provision of a high bandwidth connection. In general, the growth of network facilities has met or exceeded to targets set. The DWDM connection to JIVE has been upgraded to six, Gigabit Ethernet lines and the network infrastructure within JIVE/Astron has been upgraded to bring these lines to the correlator. Connections into the Netherlands were upgraded to 10Gb/s well ahead of schedule. Progress with local connections to telescopes has been slower but toward the end of 2003 1Gb/s connections to Westerbork and Onsala came on line. Jodrell Bank were already equipped with a 150Mb/s connection and these three stations were able to take part in a series of eVLBI tests. Through these tests and some laboratory work at JIVE, good progress has been made with understanding the network interface and configuration issues, and the evaluation of optimum data transfer parameters.

Milestones

- May 2003: First use of FTP for VLBI session fringe checks.
- September 2003: Mark5 - Mark5 data transfer between Bologna and JIVE.
- October 2003: First light on Westerbork Gigabit connection.
- November 2003: International baseline, eVLBI fringes, only 15 minutes after observations were made.
- November 2003: Onsala Space Observatory connected at 1Gb/s.
- January 2004: First European eVLBI image (see Figures 2 & 3).

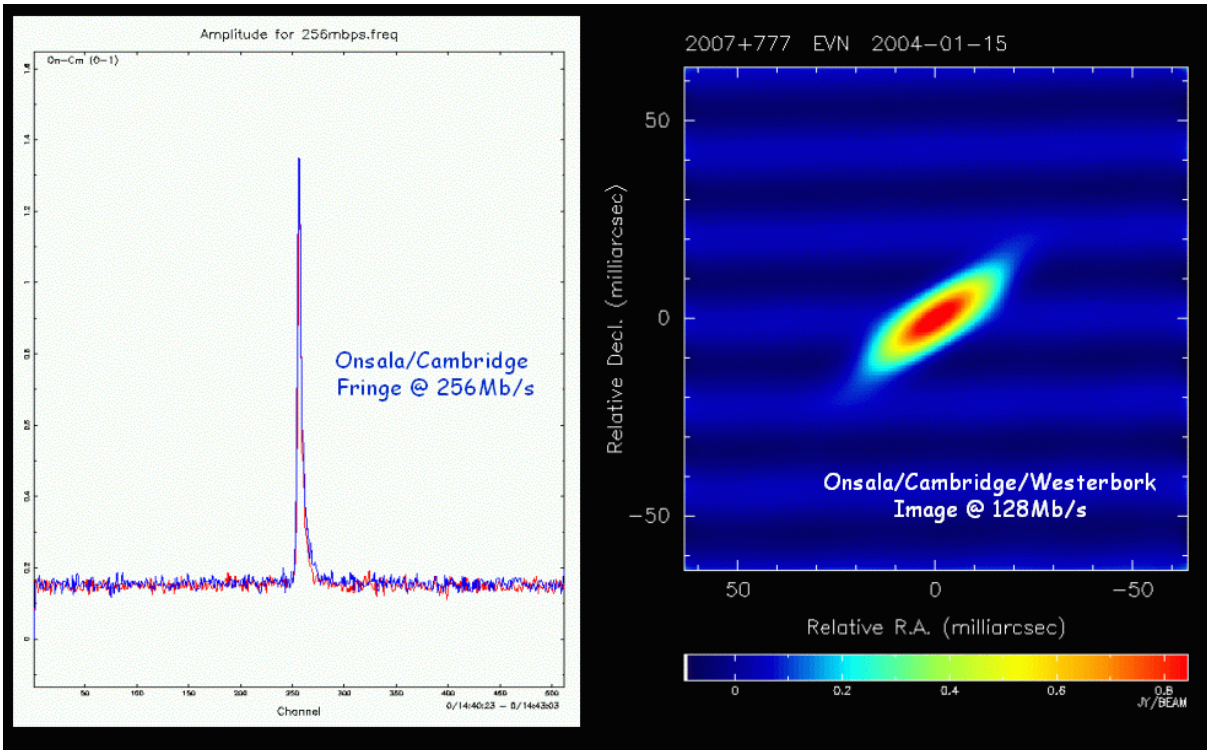

Figure 2. First European eVLBI Image

eVLBI Facilities at JIVE

Two DPUs have been removed from service. In their place two new cabinets house six Mark5s, two with dedicated Station Units and four sharing the SUs in adjacent DPUs. These units are used routinely for disk-based VLBI operations. Several units are also connected via optical-fibre cables to the SURFnet DWDM unit in the nearby, computer facilities room.

Currently six, Gigabit Ethernet lines are available using LX Optics. Further modules can be added to the Cisco 15252 unit to expand to up to sixteen such lines. In addition, a dedicated line from Westerbork has been installed. The Westerbork line is completely isolated from other network elements. A number of experimental environments are provided by these facilities:

- Mark5 - Mark5 via a patch cable for tests excluding any other network data or equipment.
- Mark5 - Mark5 via Amsterdam for longer distance tests away from competing traffic.
- Westerbork - JIVE for fully equipped eVLBI tests, also in isolation from other network traffic.
- Connection via NRENs and GÉANT to international PoPs and connected stations on shared networks.

By the use of these network variants some isolation of variables is possible and bottlenecks can be located.

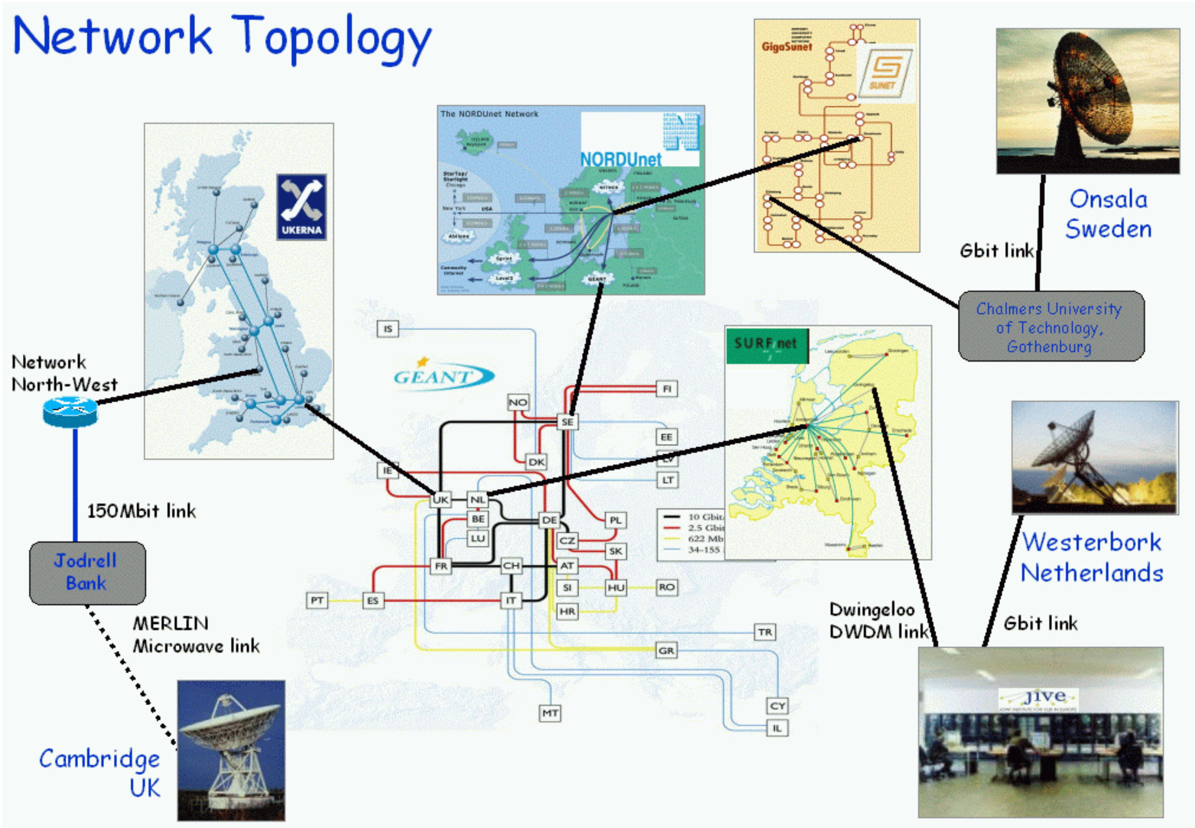

Figure 2. Network Topology for First Image

Figure 4: Mark5 Installation at JIVE

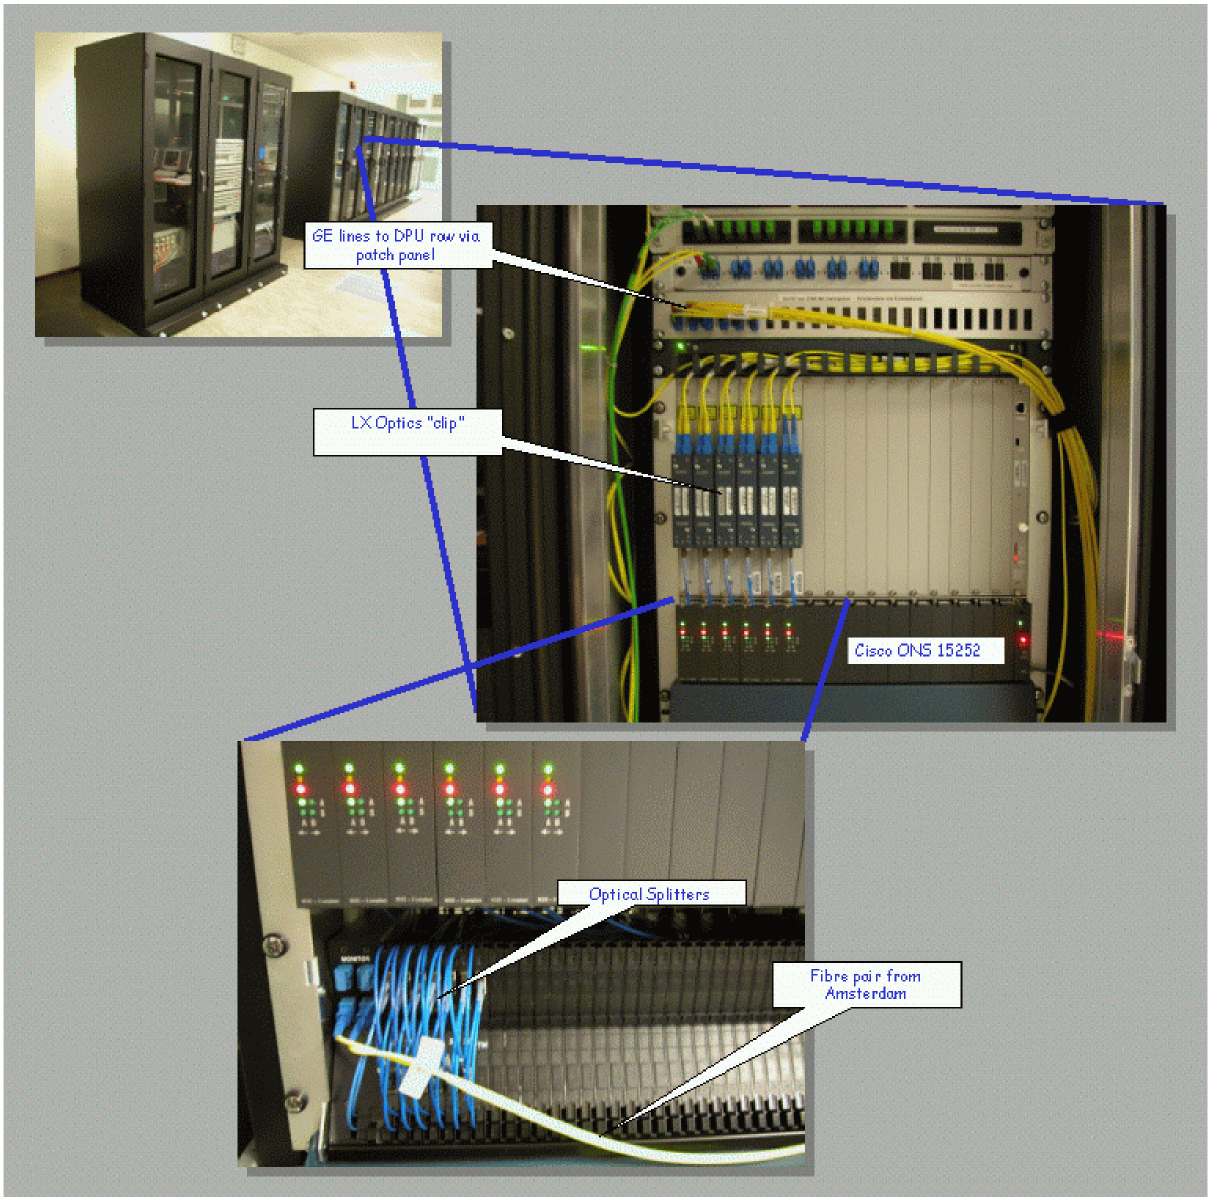

Figure 5: DWDM Equipment

Tests and Results

The built-in transfer processes of the Mark5 (in2net, net2disk etc.) have been used in different combinations for tests involving various levels of disk buffering. Some tests were also performed using TCP and UDP based network performance monitoring programmes. The results of these tests are collected in Table 1. A number of questions arise from these results that are the subject of ongoing investigation:

- Why are data rates lower when disk access in involved? What can be done to improve this?
- JIVE-Amsterdam-JIVE and Westerbork-JIVE connections are private lines, with no competing traffic. Why are they slower than via a patch cable on the bench?
- When the underlying protocol for Mark5 disk-disk transfers is changed from TCP to UDP the data rate goes down. How can this happen?
- Data rates across the European networks vary by a factor of two. Is this typical?

The data rates detailed in Table 1 where achieved after a lengthy investigation to find the optimum values for various parameters. Results of this investigation are given in Table 2. So far the underlying logic of these numbers is not well understood.

	Memory-Memory		Disk2net2Disk		In2net2Disk	In2net2Out
	UDP	TCP	UDP	TCP	TCP	TCP
Bench - patch cable		~960Mb/s J	250Mb/s	544Mb/s J		
JIVE - Ams - JIVE	500Mb/s	360Mb/s	341Mb/s J	456Mb/s J		
Wb - JIVE	867Mb/s	680Mb/s	249Mb/s J	378Mb/s J	256Mb/s	
Bologna - JIVE	670Mb/s	128Mb/s		307Mb/s		
Onsala - JIVE				177Mb/s(!)	256Mb/s	
JIVE – Haystack Internet2 Demo	612Mb/s			71Mb/s		
Jodrell (150Mb link) - JIVE					64Mb/s	64Mb/s

Table 1: Test Results

In an effort to move the enquiry forward JIVE has upgraded two Mark5 units with Intel(R), SE7501BR2, server grade motherboards. This will confirm or otherwise that some of the performance limitations are a function of the underlying PC hardware. Installation of the new motherboard was moderately easy and operation as a normal Mark5 unit has been confirmed. First results of testing with these units will be reported at the TOG.

Parameter settings (disk2net2disk, TCP)

recv/send(B)	socbuf(B)	datarate(Mb/s)
268320	134160	331
	268320	421
	536640	425
	1073280	412
2097152	1048576	522
	2097152	509
	3145728	496
	3997696	453
4194304	1048576	521
	2097152	**544**
	3145728	506
	3997696	524

Correlator Developments for Real-Time
In a real time system data will arrive at the correlator with some delay comprising a network delay and an instrumental delay caused by various bus/memory/peripheral manoeuvres within the interface units. Within Europe the network delay is in the range 10-30ms. If other delays are known and stable, the correlator time reference can be synchronised, with a suitable delay, to UT. The existing input buffer will be able to absorb differential delays across the networks. To achieve this both hardware and software modifications are required at the correlator and this is the subject of one package of work at JIVE. If delays are unpredictable it may be necessary to provide a bigger buffer, and if this is done the synchronisation requirement can be relaxed. The Mark5 has a 512Mbyte buffer in the output path which, at the highest data rate, holds up to four seconds of data; more than enough for our needs. Tests are currently under way at JIVE to characterise the performance of Mark5 in these direct-to-network modes. Of interest are the data rates achievable and the behaviour of the servo mechanism.

Forward Planning
In January 2004 NREN and EVN representatives met again to review progress one year on. The overall conclusion was that the project is largely on schedule. Developments for real-time operation are seen as the next priority. Plans are in place to achieve multiple telescope real-time operation by mid-2004. There are good prospects for a Gigabit connection to Torun on this timescale. Other telescopes, including Jodrell Bank and Medicina, do not expect to get a connection until the end of 2004 but DANTÉ and the NRENS are now willing to continue support for another year.
Planning for other activities is as follows:

- Next three months:
  - Understand/remove Mk5 data rate limitations
  - Develop/understand use of UDP for higher rates
  - Higher rate Bologna JIVE using UDP

- Next six months:
  - Further tests between Wb, On & Jb as link speed improves
  - More NREN PoP JIVE tests to push larger data volumes across the European networks

The original goal of six telescopes in real-time at 1Gb/s was recognised as unrealistic for the following reasons:

- Mark 5 is unable to deliver 1 Gb/s to the network
- Gigabit Ethernet is also limited to something less
- VLBI is restricted to fixed octaves- 64 Mb/s, 128 Mb/s, 256 Mb/s, 512 Mb/s, 1 Gb/s
- Many stations are unlikely to get more than one GE connection
- The lines into JIVE are GE

The revised goals are real-time eVLBI at 512Mb/s with at least three telescopes participating. Non real-time tests will be performed to maximise network loading up to the limits of Mark5 and Gigabit Ethernet.

Conclusions and Acknowledgements
The engineers and scientists involved in the creation of high-speed networks recognise VLBI as a genuine high-demand application that can already source very large data flows from all corners of the European continent. The big-science image of VLBI is also very welcome to the administrators and policy makers who need eye-catching publicity to maintain the profile of their product.

The VLBI community meanwhile is enthusiastic but, understandably, reluctant to commit heavily from very limited funds. This is especially problematic in the case of local loop connections. Scheduling of eVLBI tests is also a not easy. During the normal VLBI session such tests are an unwelcome distraction to already busy operations staff. Outside the sessions coordinated reservation of expensive telescope time can be difficult to achieve. At the same time the true nature of the support afforded by the networks compounds any problem. The network providers offer a best-efforts service with limited resilience. This means that no special tuning of the network for VLBI traffic is provided and technical problems have to take their place in the normal service queue, at a low level of priority.

For the above reasons the success of the project depends on the drive, effort and cooperation of individuals in both communities who are willing to work a little outside the normal parameters of their jobs. The experience of the first year of the project tells us that real progress only begins when real data starts to flow across real networks; and the key to this is to have telescopes connected directly to the network. When this begins to happen technical problems can be identified and engaged, and identifiable progress can be made.

The following roll-call identifies some of the individuals and organisations who have made the progress achieved so far possible:

- European research networks:
  - GARR
  - NORDUnet
  - SUNET
  - SURFnet
  - UKERNA
  - DANTÉ
  - GARR

- At the telescopes:
  - Michael Lindqvist, Michael Olberg: Onsala Space Observatory
  - Paul Burgess, Ralph Spencer: Jodrell Bank Observatory
  - Tony Foley, Hanno Holties: Westerbork
  - Giuseppe Maccaferri: Bologna

- At JIVE/ASTRON:
  - Cormac Reynolds, schedules
  - Arpad Szomoru, technical co-ordination
  - Sergei Pogrebenko, science advisor and software correlator
  - Klaas Stuurwold, network facilities and NREN liaison

- Richard Hughes-Jones of Manchester University for network performance measurement software
- Thanks should also go to the Haystack Observatory Mark5 team for the Mark5 itself and helpful support.
- Finally thanks and congratulations to the management and staff of Onsala and Westerbork Observatories for providing the high-bandwidth links to their telescopes that have got this project off the ground.