

EXPR_eS/FABRIC at JBO

Report to EXPR_eS Board ESPOO, Finland 29th May 2007


Ralph Spencer

The University of Manchester



Outline

- FABRIC:
 - Protocols
 - 4 Gbps eVLBI data into e-MERLIN correlator
- EXPReS at JBO:
 - Connection of 4 e-MERLIN telescopes to JIVE
 - 10 Gbps network: 'last mile' connectivity



JRA Fabric: Protocols for high speed data transfer

- ESLEA funded RA – Stephen Kershaw, contributed work until end June 2007, then 1 yr on ESPReS.
 1. Work on constant bit rate data movement over TCP/IP
 2. Simon Casey – PhD student working on VLBI_UDP
 3. VSI-E, DCCP
 4. Connectivity tests and 4 Gbps flows

1. TCP tests – TCP-delay

- TCP is a reliable protocol
 - Data will arrive at the destination...
 - ... in the order sent...
 - ... eventually ...
 - There are no guarantees about the timeliness of data delivery
 - This is a problem for constant bit rate data

Why are we looking at TCP? Because UDP at high data rates on a packet switched production network is a no-no!

Message size: 1448 Bytes

Wait time: 22

Data Rate: 525 Mbit/s

Route:

Manchester – Ams - JB

RTT 27.2 ms

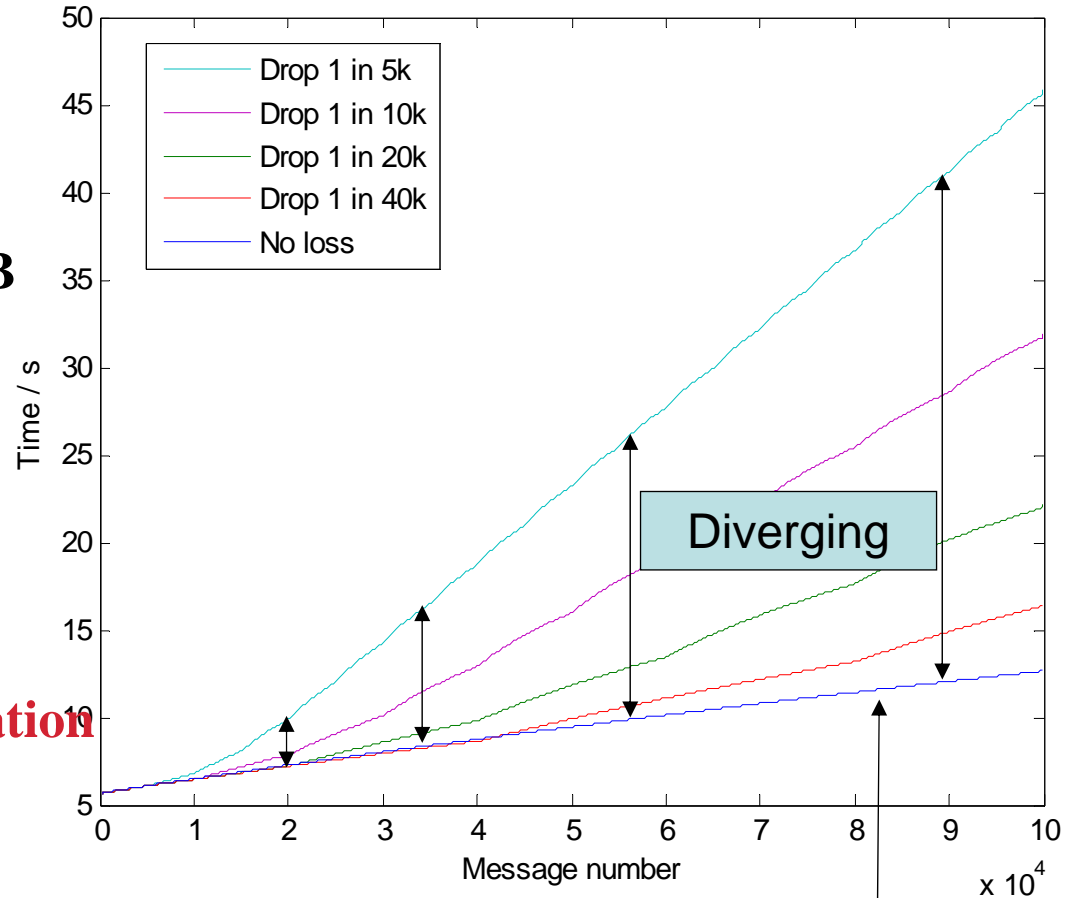
TCP buffer 1.8 MB

Vary packet drop rate

Curves diverge

- Delay depends on duration of data stream**

Effect of loss rate on message arrival time



Diverging

Timely data arrival

Message size: 1448 Bytes

Wait time: 22

Data Rate: 525 Mbit/s

Route:

Manchester - JIVE

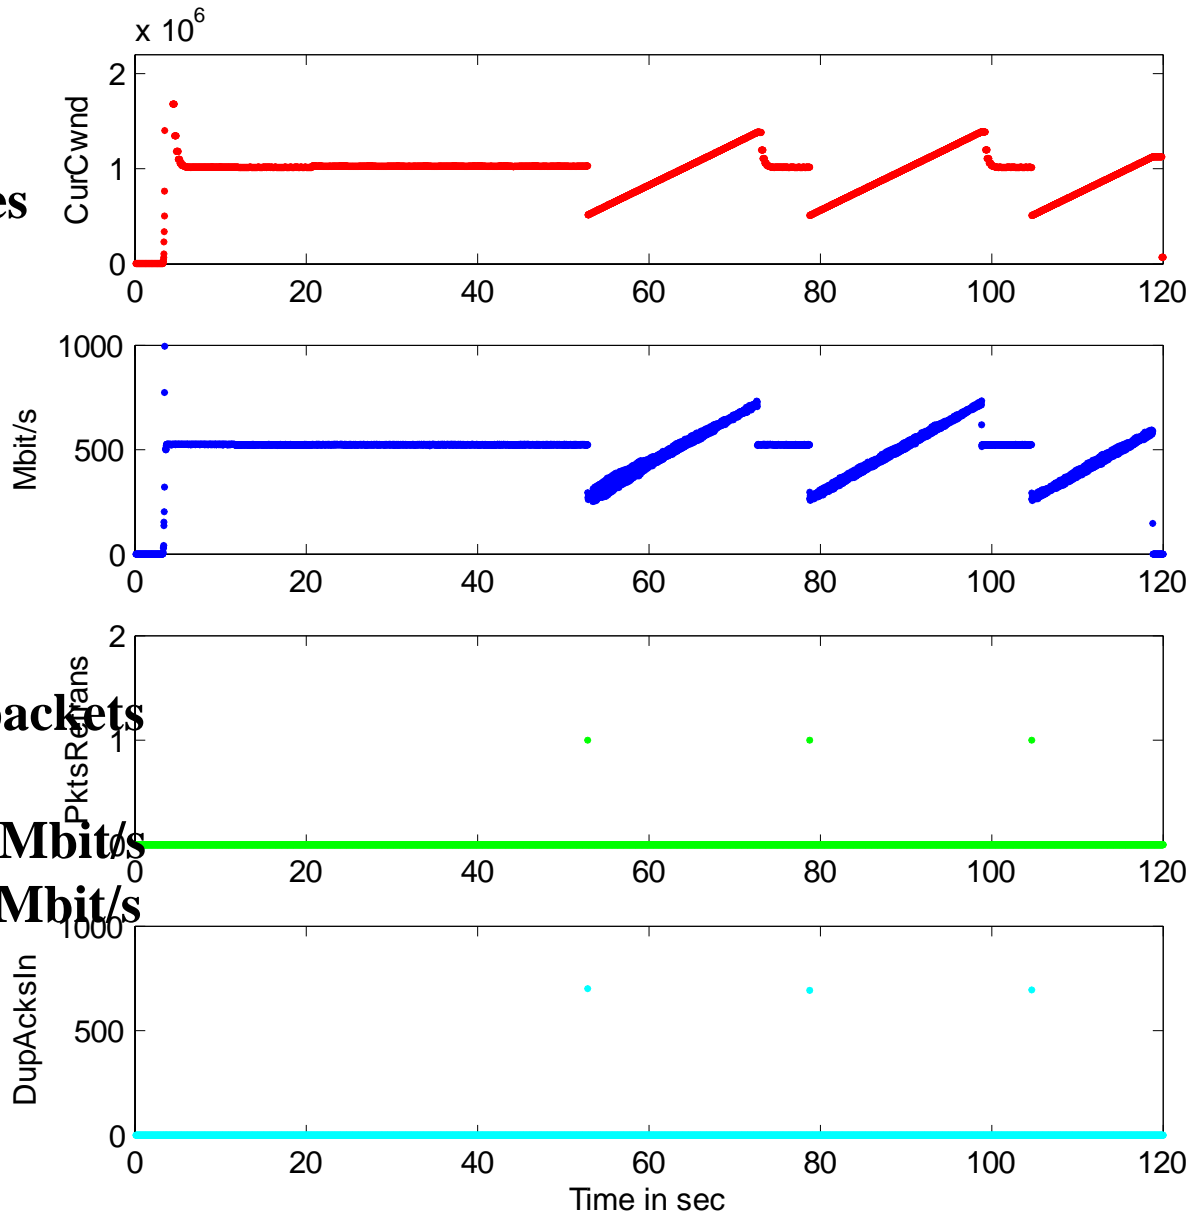
RTT 15.2 ms

TCP buffer 160 MB

Drop 1 in 1.12 million packets

Peak throughput ~ 734 Mbit/s

Min. throughput ~ 252 Mbit/s



TCP: Conclusions

- It is possible to use TCP as a transport protocol for constant bit-rate data transfer e.g. real-time eVLBI data
 - Far from ideal
 - Requires 'spare' bandwidth
 - Unfavourable scaling

Maybe DCCP can help, but more work required

2. VLBI_UDP and correlator tests

Simon Casey*, Ralph Spencer*, Richard Hughes-Jones#,
Matthew Strong*, Paul Burgess*, Arpad Szomoru[§]

ESLEA conference
March 2007

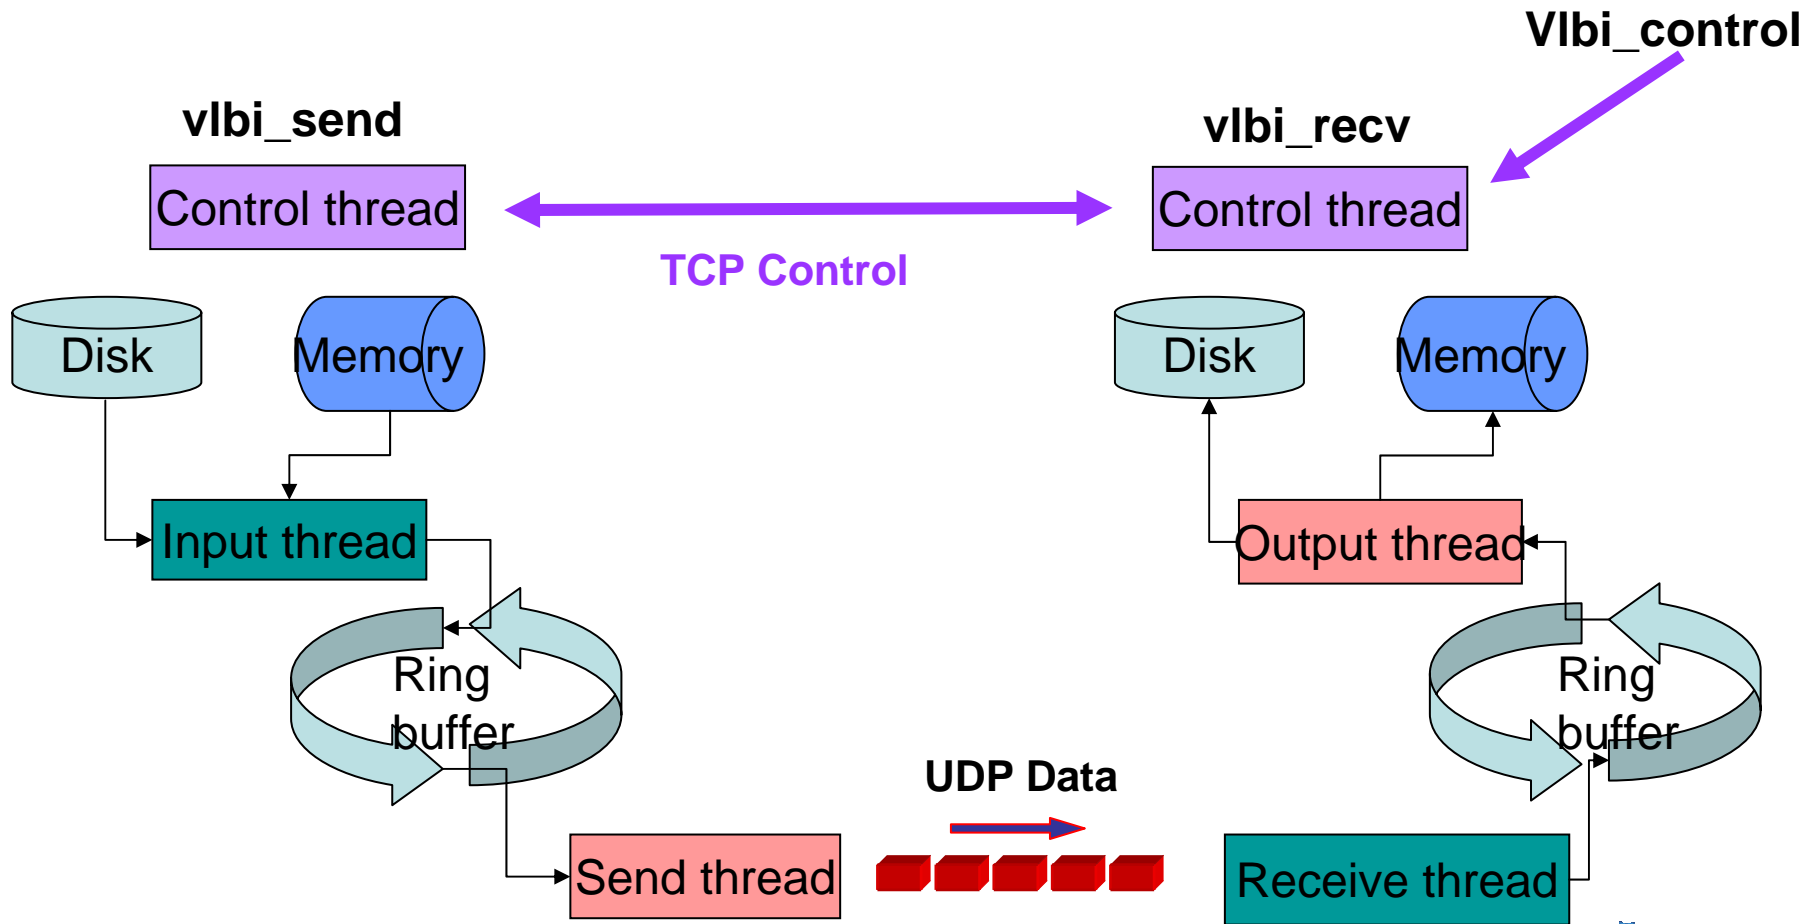
*Jodrell Bank Observatory, The University of Manchester; #HEP Group, The University of
Manchester; [§]Joint Institute for VLBI in Europe (JIVE), NL



JOINT INSTITUTE FOR VLBI IN EUROPE



VLBI_UDP architecture




Ultimate Goals

- What is the maximum tolerable loss rate?
- Can packets be lost selectively (e.g. not the headers) without seriously affecting correlation?
- Implement findings into VLBI_UDP
 - Can be used to selectively drop packets if congestion present yet keep correlation running
- Put VLBI_UDP onto Mk5s !

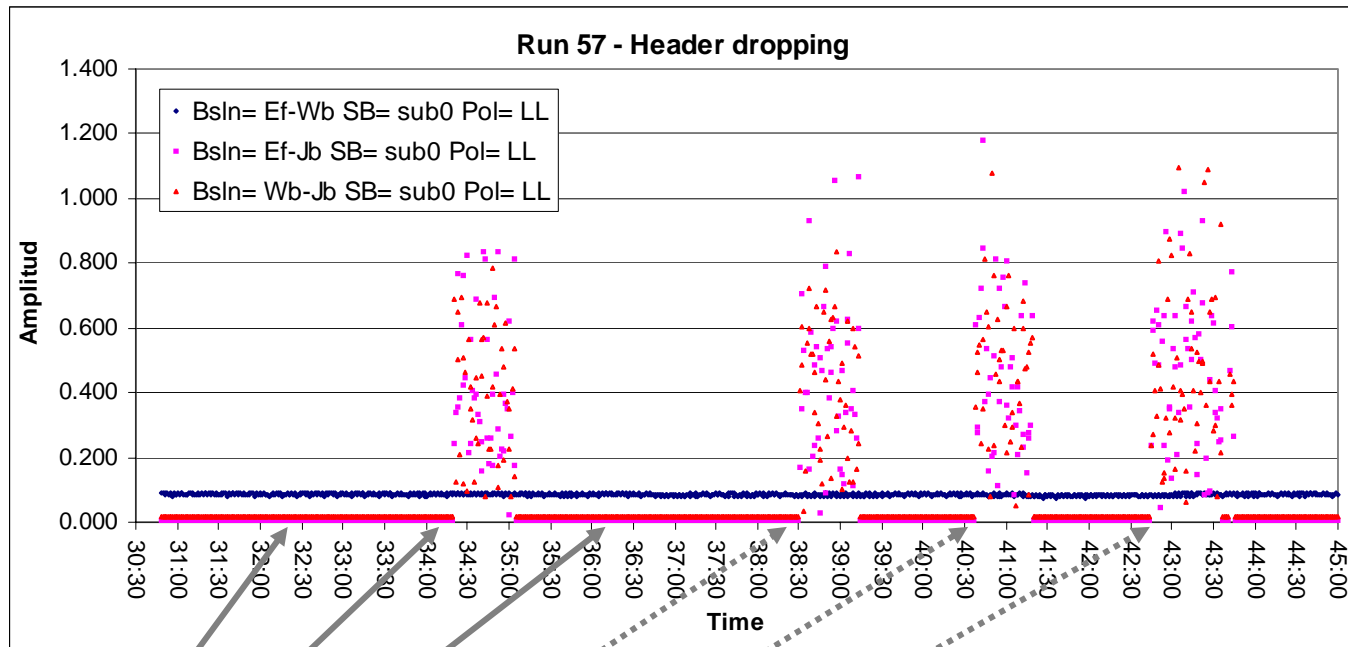
Correlation tests at JIVE

- Take pre-recorded VLBI data
- Use VLBI_UDP in disk-disk mode
 - Write to local file instead of sending to network
- Filter each packet through a dropping function
 - Algorithm can be selected by user and parameters varied
- Place modified data file back on to disk pack
 - Play into correlator

Fringe amplitudes (December 2006)



Results – Header dropping



<i>Minutes in file</i>	<i>Header drop pattern</i>
2-4	Drop---Keep---Keep
4-6	Drop---Keep---Drop---Keep
6-8	Drop---Drop---Keep---Keep
8-10	Drop---Drop---Keep
10-12	Drop---Drop---Drop---Keep
12-14	Drop---Drop---Drop---Drop---Keep



Conclusions

- Use of Mark5 special ‘fill pattern’ to indicate invalid data gives more predictable results
 - Invalid data not correlated
- Station units able to cope with at least 1 in 3 missing headers
- High loss rates (10’s%) should be achievable without loss of synchronisation
- Selective packet dropping is possible

3. VSI-E and DCCP

- VSI-E – now have code thanks to Haystack staff, visit at end of April. Compiles!
- DCCP – enables user to define the congestion control, e.g, can maintain rate for small packet loss, turn off if high etc. Tested at UCL, crashed on 64 bit machines, development continuing, hoping to test with VLBI data soon.


4. 4 Gbps tests

- Intel 2.3 GHz twin core machine installed in GEANT PoP (Telecity, Docklands, London) 15 May, Test network: London-Prague-London 4 Gbps, no packet loss!

EXPR_eS/FABRIC and e-MERLIN

- FABRIC:
 - Demonstration of >1 Gb/s real time data transfer and correlation
 - Onsala → e-MERLIN correlator at ~4 Gb/s
- EXPR_eS:
 - SA1 Integration of e-MERLIN and e-EVN
 - 4 e-MERLIN telescopes → JIVE at 1 Gb/s
 - SA2 NetworkConnection
 - JBO to Manchester
 - Manchester to JIVE

JRA FABRIC Exporting & Importing eMERLIN Data Using iBOBs



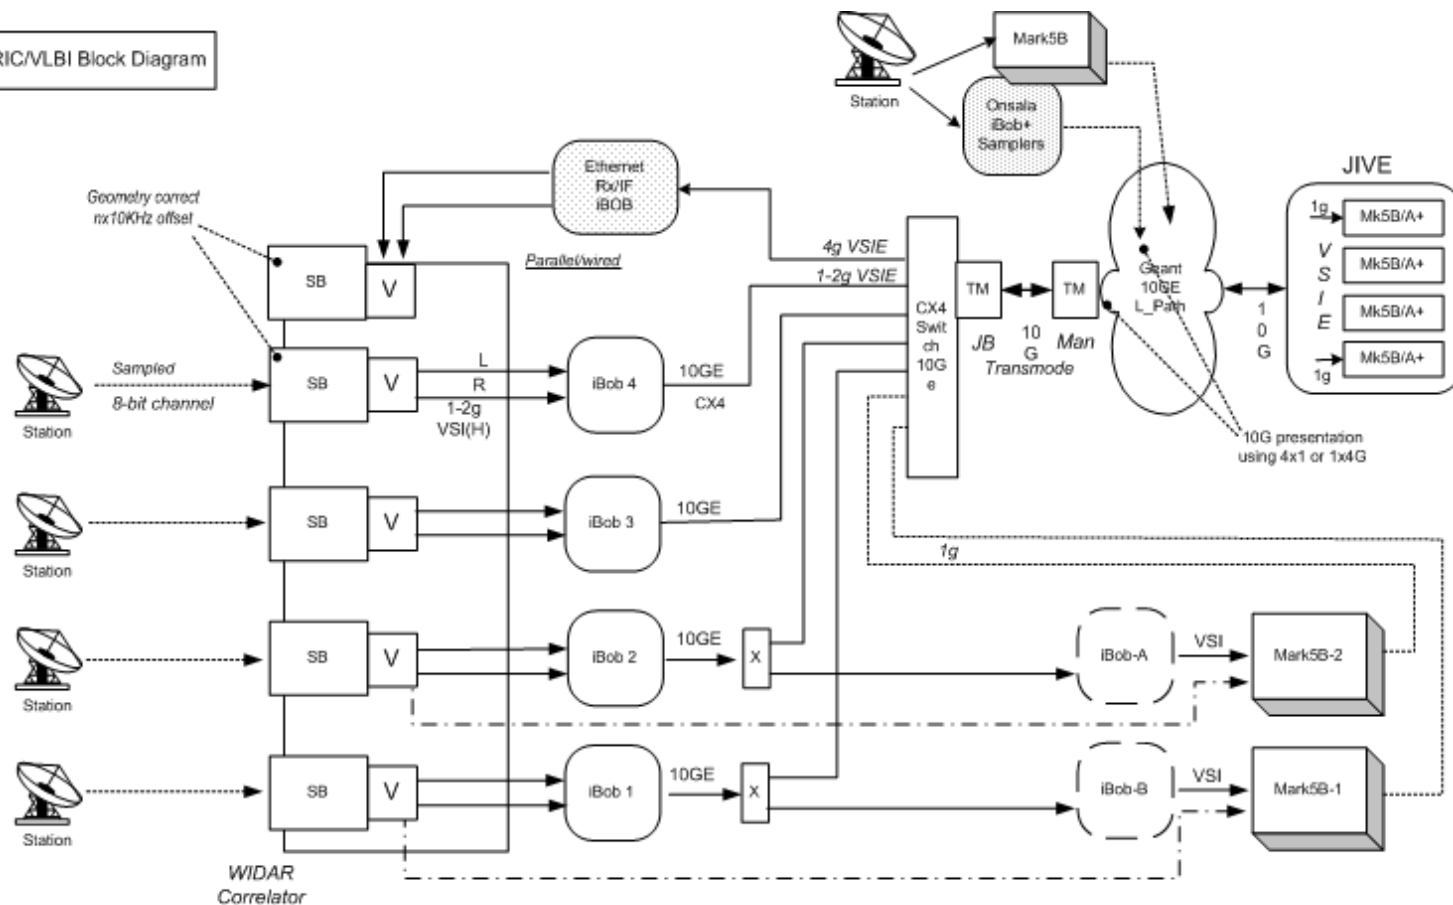
Engineer: Dr Jonathan Hargreaves, appointed Nov 2006

- Export – Jodrell to JIVE
- Import – Onsala to Jodrell



E-MERLIN---e-VLBI system

eMerlin/FABRIC/VLBI Block Diagram



FULL FILENAME	
F:\PROJECTS\EVVLBI\XPRES07\EM+M5_2_VSD	
TITLE	REVISED
BLOCK DESIGN	13/02/2007
DATE	DRAWN BY
02/10/2006	P. BURGESS

----- Altern. path (VSI) limited length

----- Existing path, <1G

1. iBob A/B possibly replaceable with network switch: conv. 10G->1G

2. X: iBob 1 & 2 patch, use switch vlan, or duplicate stream.

3. 'V' = Xilinx+VSI IP+cableconverter

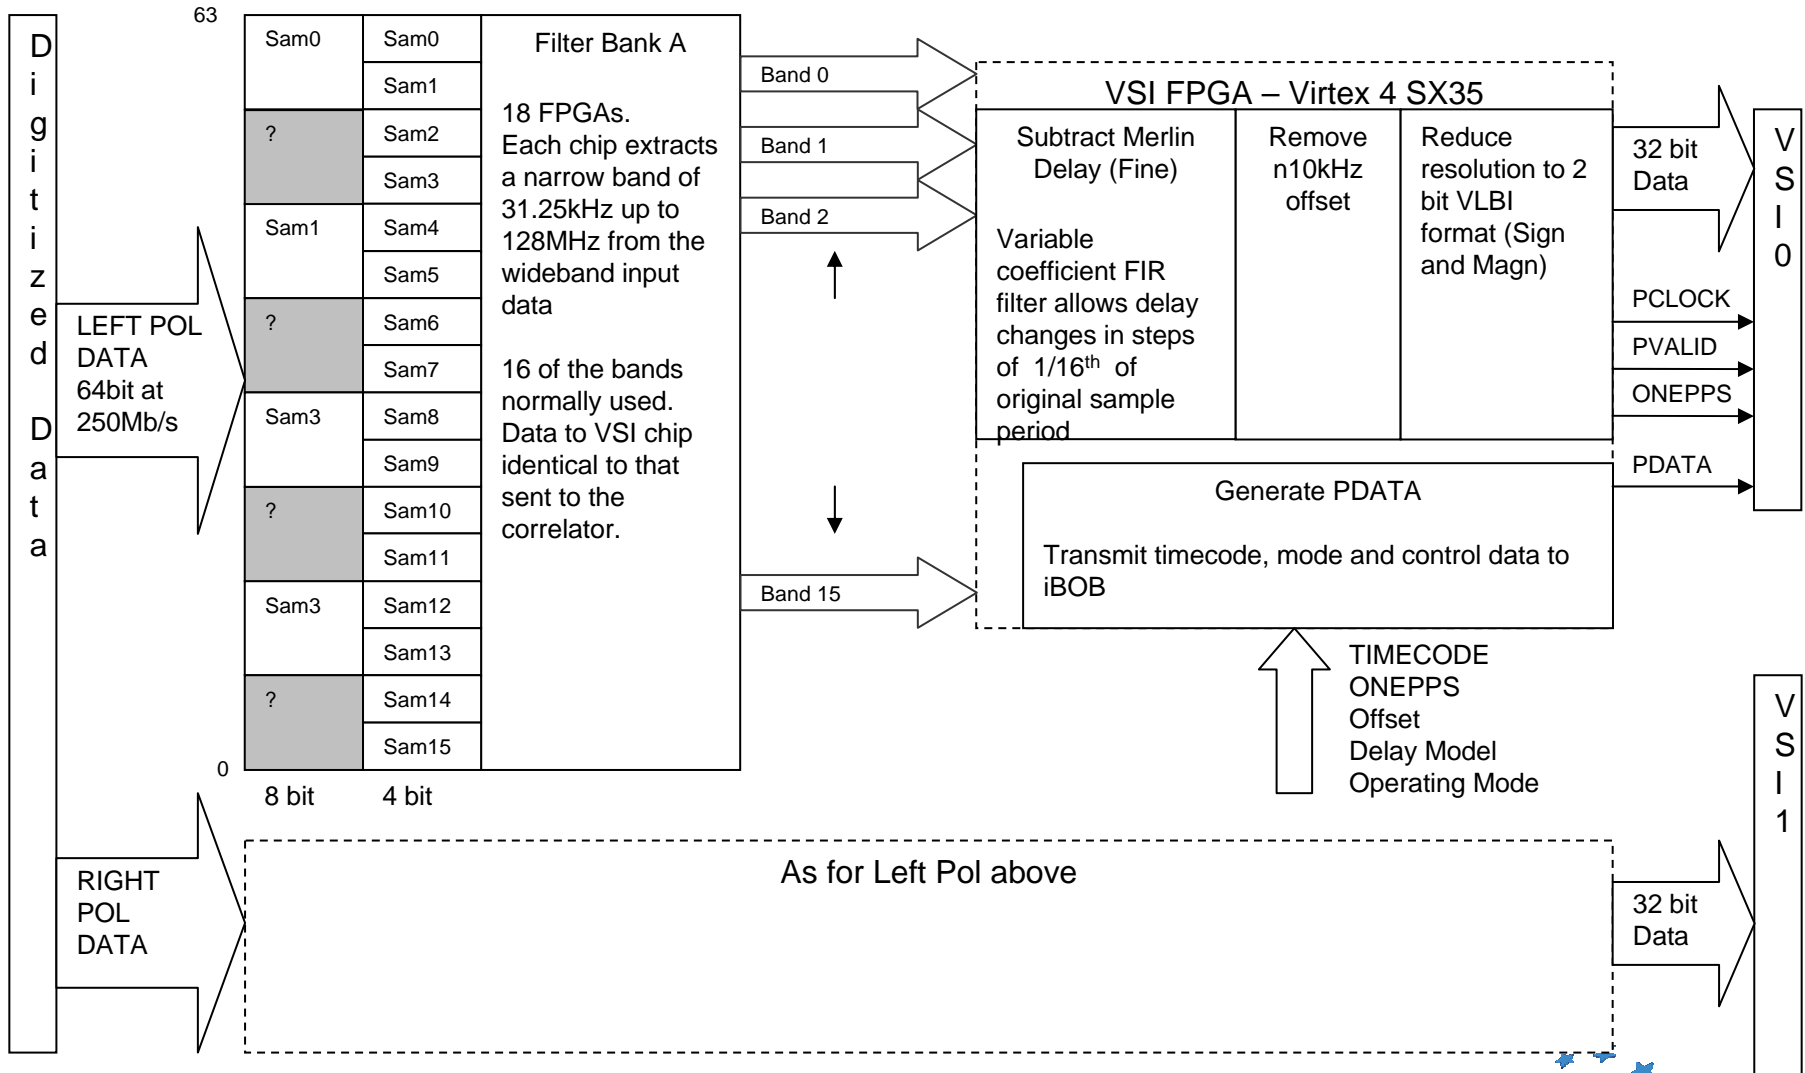
4. iBobs and Mk5's require 1024MHz and 1pps

JBO to JIVE

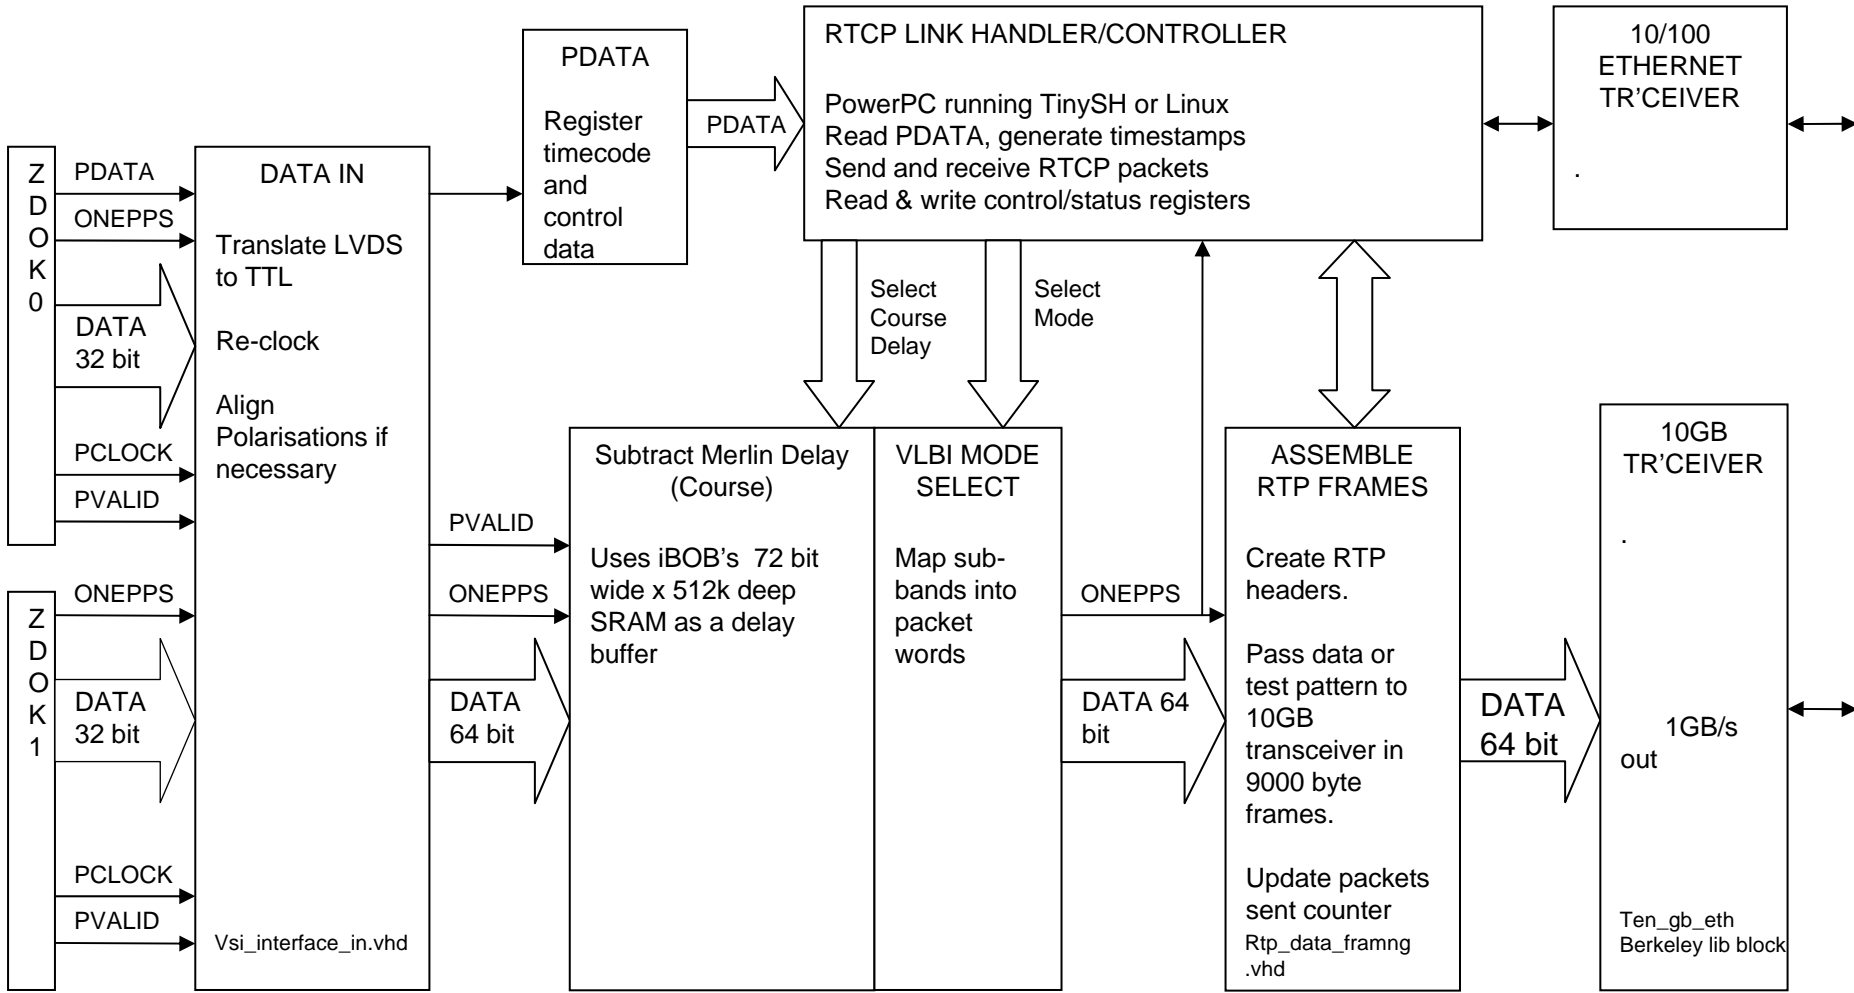
- Export 4 x 1Gb/s from Merlin correlator at Jodrell
- Four iBOBs, each connected to a Station Board → 1Gb/s each
- Input: narrowband data streams from filter bank – identical to data sent to the correlator

- Need to remove geometrical delay, added to the Merlin data at the antenna, (fractional delay filter simulated)
- Remove N x 10kHz offset – can switch it off for initial testing
- VLBI modes:
 - 2 pol x 2 bits x 32MS/s x 8 sub-bands
 - 2 pol x 2 bits x 16MS/s x 16 sub-bands
 - Narrowband
- Output in VSI-E 2.7 format

Station Board showing 'VSI' chip



iBOB Configured as 1GB/s Transmitter



Onsala to JBO

- 8 bit iADC samples 2 channels at 1GS/s or 1 channel at 2GS/s
- iBOB reduces data rate to 4Gb/s and transmits to internet
- Another iBOB at Jodrell receives and buffers data against out of order packets and variations in network delay
- Expand RTP header to include 32 bit ‘real time’ in seconds, plus 32 bit ‘sample count’
- No need for RTCP channel, but may need to reduce transmit rate if congestion occurs

Summary

- Protocol work proceeding well though late on original timescale
 - Looking at multiple destination data transfer after July
 - Further multi-Gbps tests as 10 Gbps network gets up
- Work on iBOBs continuing
 - Manufactured, being tested at Berkeley
 - Code being produced and simulated
- 10 Gbps links
 - JBO-Manchester upgrade out to tender
 - Manchester 10 Gbps links expected in Summer (delays in Cienna kit)
 - Pathway through SJ5 being negotiated
 - GEANT connection need clarification
- ESLEA work published by PoS
(http://pos.sissa.it/cgi-bin/reader/index_html.cgi)

Issues!!

- Data transfer protocols – liaise with Posnan re distributed processing –what is required?
- Nx10 kHz can it be done in the correlator? (takes a large part of SB FPGA)
- Onsala connections SUNet, NorduNet and GEANT?
- ADCs and iBOBS DAQ – progress at Helsinki?
- Data format in and out – VSI-E using UDP?

Questions?

- Contact information
Ralph.Spencer@Manchester.ac.uk

- Additional Information
<http://expres-eu.org/>
<http://www.jive.nl/>

- EXPReS is made possible through the support of the European Commission (DG-INFSO), Sixth Framework Programme, Contract #026642