

Export/Import of VLBI Data from/to eMERLIN

Bryan Anderson
Version 1.2, April 21, 2006

0 Revision History

| Version | Date | Changes/Notes | Author |
|---------|------------|---|-------------|
| 1.0 | 2006-04-19 | Initial release | B. Anderson |
| 1.1 | 2006-04-20 | Minor corrections & amplification of details in the proposals | B. Anderson |
| 1.2 | 2006-04-21 | More minor corrections | B. Anderson |

1 Introduction

This review examines the ways in which astronomical data can be exported from eMERLIN to VLBI processors and how similar data might be imported to eMERLIN from an external telescope. Since astronomical data are routinely exported from MERLIN as part of our VLBI activities, export is regarded as a higher-priority activity than import. The eMERLIN correlator is being implemented for an eight-telescope array and only one or two of the eight telescopes might normally be external to eMERLIN. An expansion beyond the eight would require a very significant increase in hardware and thus cost and is not being contemplated.

The goals are for data from multiple eMERLIN telescopes to be available for export, for the exported data to be in existing VLBI formats and for there to be potential for expansion in bit rate and formats if the needs arise. The export might be via either disk-based recordings or directly via the net. It would be the costs of the exports that might limit the number of telescopes exported. In either case, the exported data should be as free as possible from eMERLIN-specific artefacts so that there should be no need for software or hardware modifications at the correlator to handle them.

The goals for import are to make it possible to connect an external telescope into eMERLIN via a connection over an existing, non-dedicated network. The data need to be correlatable and the bit rates should be as high as the network can provide. In the first instance, the connection of the Onsala telescope into eMERLIN at up to 4Gbps is considered. The import facility will act as a test bed for future eVLBI systems and protocols whilst adding useful capabilities to eMERLIN.

In both these endeavours, the aims are to minimise development effort by using, as far as possible, existing hardware and intellectual property, and by trying to ensure that any developments have the potential to be re-usable in future activities. Both export and import need to be as non-intrusive as possible and not necessitate any physical reconnections of signal paths in eMERLIN.

The eMERLIN system already has provision for export and import of astronomical data via *VSI-H-like* interfaces to and from the station boards (SB) though the design has not progressed beyond the inclusion of devices and connectors for this purpose. The pcb artwork provides the appropriate connections but this review is the first opportunity to consider the specifications for the firmware. In Section 2, the capabilities and constraints of these interfaces and the eMERLIN correlator are described in more detail.

In the shorter term, it is expected that export of data through disk-based, VLBI recording systems are likely to be wanted so Section 3 is a short summary of the provisions and connection capabilities of existing and projected VLBI systems.

The conclusions of this review are in Section 4. The main conclusions are that viable, off-the-shelf hardware and design tools exist for most of what needs to be done. Additional expertise in signal processing/high-speed digital design is required for the firmware but many of the components of hardware, firmware and software are reusable between the export and import applications and, indeed, other projects so that significant economies are possible.

2 eMERLIN Constraints

The principles on which the eMERLIN correlator, a WIDAR correlator, operate and its broad characteristics are described in references [1] and [2]. The correlator is being developed by Peter Dewdney's group lead by Brent Carlson at the Dominion Radio Observatory at Penticton, Canada. Its functionality is a subset of that of the much larger correlator that they are building for the EVLA. It normally accepts two IF bands 2.048GHz wide sampled at 4.096GHz to 3 bits of precision or two IF bands 512MHz wide sampled at 1.024GHz to 8 bits of precision, splits each of the bands into 16 sub-bands up to 128MHz wide, resamples the results with either 4-bit or 8-bit precision and then correlates them. The main signal processing flows in a *station board* (SB) are shown in Figure 1.

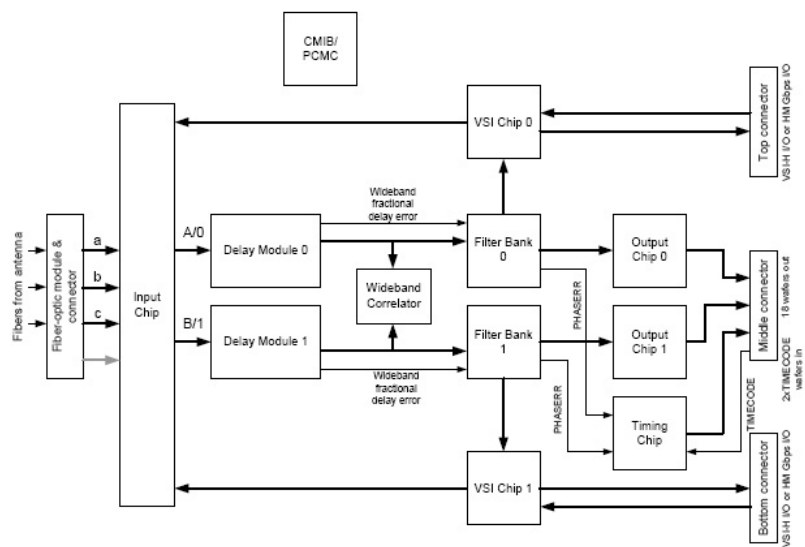


Figure 7-1 Simplified Station Board block diagram.

Figure 1: Simplified diagram of an SB from Carlson's *Programmers Guide*.

The main signal paths are horizontally through the centre of the diagram for each of the two input bands. An input selector is followed by pairs of system delays of up to 0.25 seconds and then 16-channel FIR filterbanks through which the signals are chained from chip-to-chip. Each filter can be used to narrow its output band down to the Hz level and the results, resampled to 4 or 8 bits, are available at the middle connector for correlation. The astronomical data stream continuously through the system and are accompanied by

data-validity and time-tick streams.

As far as this report is concerned, the things to note are the horizontal paths at the top and bottom of the figure. *VSI chips*, Xilinx Virtex-4 SX35 devices, can accept data from each filter bank, process the data in some way and route the results to a pair of connectors at the right-hand side. Likewise a path exists back from the connectors through the VSI chips to the input selector. Whilst the signal connectors are not directly VSI-H compatible, simple wiring harnesses would make them so.

The sub-band connections between the filterbanks and the VSI chips are six differential signals clocked at up to 256MHz. The six signals are a clock, a time-tick and four data samples. In 8-bit mode, the least-significant and most-significant nibbles are transmitted alternately. The total data bandwidth from the filterbanks to the VSI chips is thus $2 \times 16 \times 4 \times 256$ (polarisations/subchannels/bits/clock), i.e. 32Gbps. The *VSI connectors* can support 2×32 bits at 64Mhz for 4Gbps and have the potential, in non-VSI mode, at up to 256MHz clock rate, for 16Gbps depending on the quality of the pcb traces and connecting cable. The combination of these features give lots of scope for export and import of data.

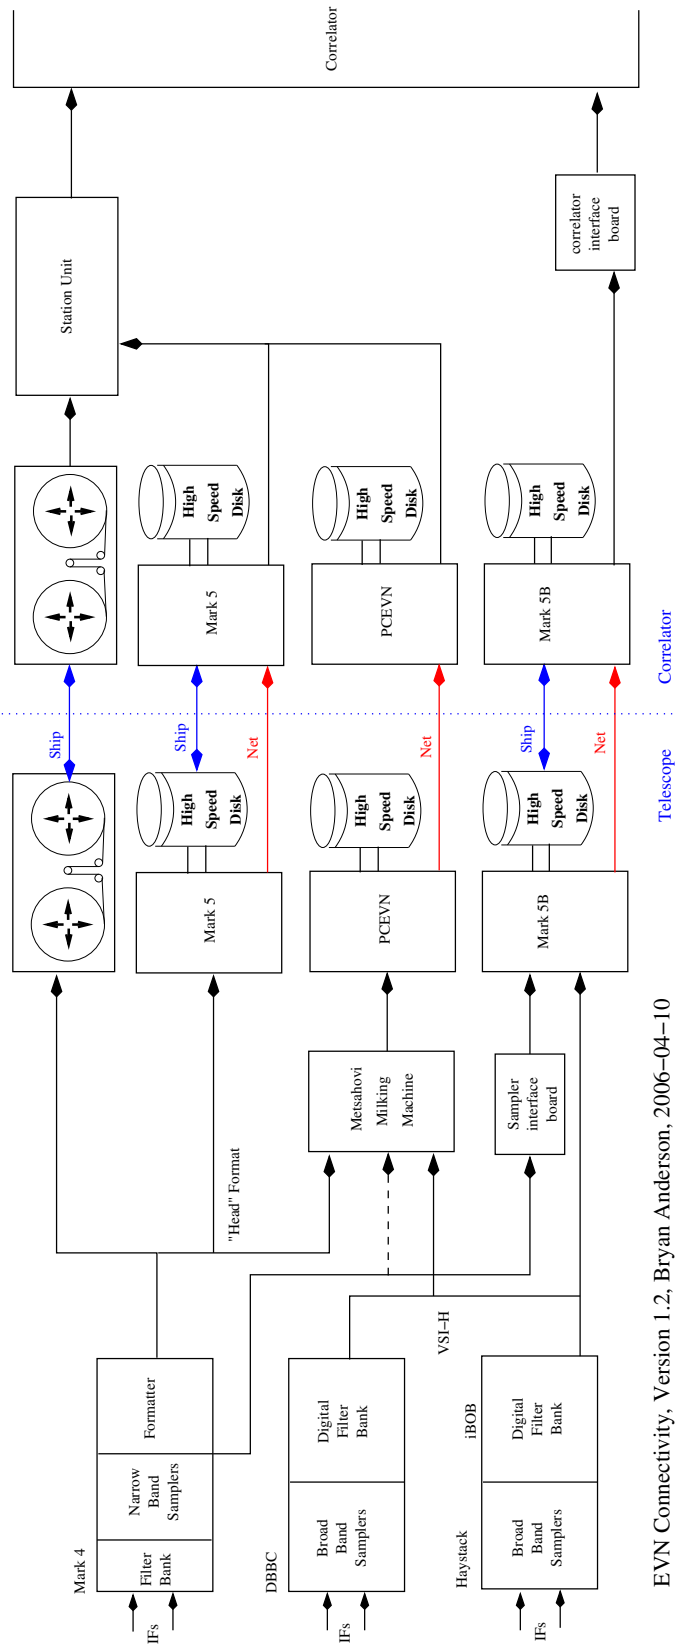
To export both hands of polarisation, it may be necessary to use both VSI connectors because there are no cross-over links between the bottom filterchain and the top VSI chip or vice versa. This feature may make a direct connection from an SB to a VSI-H-compatible external device less attractive. There are also differences in timing between sub-bands of a few clock cycles due to pipeline delays in a chain. These differences can and will be absorbed in the VSI chips.

Note that signals enter the filter chains after having a model delay applied. If eMERLIN is concurrently making observations and trying to export data then the exported data will either have the eMERLIN delay model applied or else a set of expensive, independent SBs will need to be run in parallel with the eMERLIN ones. The VSI chips do not have sufficient internal RAM to correct for the eMERLIN delay so an external solution is preferred.

Another problem is that the VLBI and the preferred eMERLIN frequency channelisations are not identical: the latter would normally be much broader. This latter aspect is not so serious because there is a lot of signal-processing capability in the VSI chips that can be used to provide whatever channelisation is desired. The WIDAR principle also requires the LO frequencies of each telescope to be offset from each other in distinct multiples of, say, 10kHz. For narrow-band spectral observations, it is desirable for these offsets to be removed before the final frequency channelisation is performed. A suitable place to do this would also be in the VSI chips.

3 Existing VLBI Interfaces

Any data exported from eMERLIN should not, in the short term, preclude disk-based recording of data. Figure 2 shows the current systems and, it is hoped, what can connect to what. The information about the Mark 5B and the sampler and correlator interface boards come from [3] and [4] respectively. Another system, PCEVN2, is being/going to be developed but specifications are not yet available. It is expected that it will be VSI-H compatible. It is somewhat disappointing that even the Mark 5B, in a late stage of development, can only support a bit rate of up to 2Gbps and that only for recording.



EVN Connectivity, Version 1.2, Bryan Anderson, 2006-04-10

Figure 2: A connectivity diagram for current and proposed VLBI systems.

What is clear is that it is desirable to provide VSI-H capabilities for eMERLIN export of astronomical data to disk-based systems and that a SB-to-disk-based converter is desirable in order to cope with some of the problems noted in the previous section. The use by the Haystack group of the already-developed iBOB *break-out-board* [5] in their new digital back-end (DBE) [6] is a welcome development. An iBOB provides lots of very useful facilities and reduces development times and costs.

4 Proposed Solutions

4.1 Requirements

A number of additional requirements have been identified, viz.

1. Completion of the firmware for the VSI chips in the SBs. This would encompass the points noted above and the interfaces to the outside world.
2. An external system to interface between the SB connectors and the disk-based system. This interface is desirable because both connectors on the SB are probably required and it would be very expensive and complicated to use two disk-based systems per telescope. It would also provide a means of solving other problems.
3. A means in the external system to remove the eMERLIN delay model from the data. This requires megabytes of RAM and access to the delay model. Exported data would have the delay model subtracted and a positive delay offset added. The result might be data with, say, 2 milliseconds of fixed delay with respect to signal arrival times at a telescope.
4. A means to get data up to 4Gbps from the net into an SB. It would need RAM to buffer arriving packets and deal with transmission impairments such as missing and out-of-order packets.
5. If possible, a means to get up to 4Gbps of data out to the net. Ideally, this would be capable of providing data packets compatible with current and future eVLBI practice.

Given the need to minimise development effort and shrink timescales, it would be profitable to look for an off-the-shelf solution to these needs. The iBOB selected by Haystack for their DBE seems to be a very good match to the problems and is reputed to be available with high-level tools for firmware development. It has multiple high-speed interfaces, on-board RAM and a capable FPGA. Accordingly, the proposed solutions incorporate iBOBs.

4.2 The Proposals

Figure 3 shows what is being proposed. The arrows show the direction of signal flow for data export. The data flows would reverse for imported data and the connection to the disk-based system would not be used. The data flows between the SB and the iBOB would not necessarily be VSI-H compatible since that would limit the flow rates somewhat artificially.

What needs to be done in each VSI chip, apart from data routing and the data buses to and from the iBOB, is to produce 16, potentially-narrower bands from one or more of

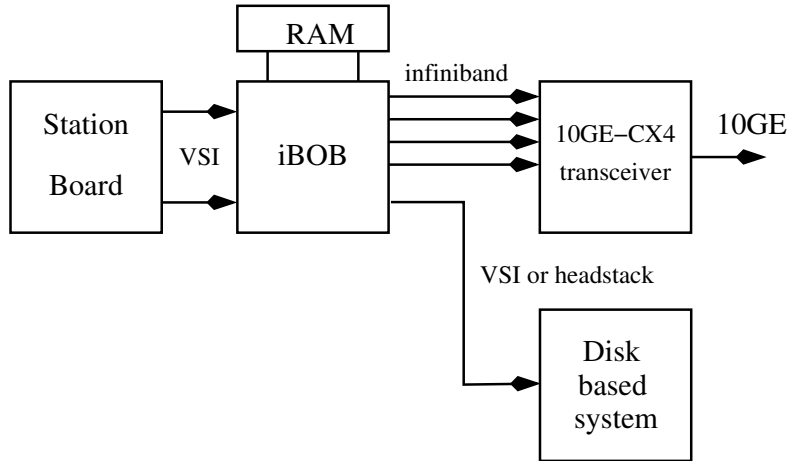


Figure 3: An external interface for eMERLIN using an iBOB.

the incoming 16 bands. Ideally, one would want to perform a frequency conversion first so that the output bands could be taken from anywhere in an incoming band. In fact, the IP to do all this is already in the filter chip FPGAs of the SB in the form of the second-through-fourth filter stages. Since the multiply-accumulate burden is smaller for these stages, it *may* be possible to fit 16 completely-independent such channels into the VSI-chip FPGAs. The details of this need to be investigated but it is likely that much of the needed IP exists and the major burden will be in getting it all into shape.

The iBOB board is used to provide the rest of the functionality. It provides the interfaces to the SB, the disk-based system and the internet, and the internal RAM is used to correct for the model-delays in exported data and to buffer packets of imported data. It may be necessary to make a connection from the Time Code Generator in the correlator to get the delay model to the iBOB. The *infiniband* interfaces are mechanically compatible with the CX4 standard and 10GE-CX4 transceivers are available for a current cost of under \$1k. The required IP for this part may already be available since this does not seem to be an unusual thing to want to do with an iBOB.

4.3 Data Capture at Onsala

There still remains the problem of how compatible data can be generated at the Onsala telescope and put on to the net. The obvious solution follows from the solution described above. It is to use a variant on the Haystack DBE in which the output is taken through the infiniband port to an external 10GE-CX4 transceiver which is used to get the data out to the net. A dual, 8-bit ADC would also need to be acquired from UC Berkeley or their suppliers. Haystack claim a cost of less than \$5k per DBE.

Very little modification of the programming of the FPGA in the iBOB is required apart from the additions for the ethernet port. To make the frequency channelisation more compatible with eMERLIN, the polyphase filterbank/FFT engine would probably need to produce fewer and wider bands but these are just parameter changes. Some decisions would need to be taken as to what data are sent. For example, one could choose to transmit a relatively small number of bands digitised to 8 bits at one extreme to a larger number of bands digitised to 1 bit at the other extreme. This sort of choice is probably already available in the DBE. Similarly, whether data are sent at 256, 512, 1024, 2048 or 4096 Mbps are choices probably already there to some extent but over a smaller range.

5 Conclusions

The natures of the problems have been described and solutions have been outlined. The merits of the solutions are that no hardware development is required, the costs of the hardware components are modest and many of the IP problems are already being tackled elsewhere. It may be possible to share IP development costs, e.g. trade, with others interested in similar applications. Finally, the hardware elements, largely in the form of iBOBs, are reusable in, for example, digital back-ends.

These proposals need to be discussed, agreed and modified, if necessary, in conjunction with our collaborators. The steps then would be: to produce a firm set of requirements; to resolve some open issues about availability of components, tools and IP; to estimate manpower requirements; to cost the developments and develop a plan with timelines.

6 References

1. Efficient wideband digital correlation by Brent R. Carlson and Peter E. Dewdney, *Electronics Letters*, vol. 36, no. 11, 25 May 2000, pp. 987-988.
2. A proposed WIDAR correlator for the EVLA Project by Brent Carlson, http://www.aoc.nrao.edu/evla/geninfo/memoseries/nrc_evla_memo.pdf
3. The Mark 5B VLBI Data System by Alan Whitney, ftp://ivsc.gsfc.nasa.gov/pub/general-meeting/2006/presentations/gm2006_4-06_whitney1.ppt
4. VSI Interfaces for Legacy Systems by Dan Smythe, ftp://ivsc.gsfc.nasa.gov/pub/general-meeting/2006/presentations/gm2006_4-07_smythe.ppt
5. BEE2: A High-End Reconfigurable Computing System by Chen Chang, John Wawrzynek, and Robert W. Brodersen, <http://bwrc.eecs.berkeley.edu/Publications/2005/JOURNALS/c.chang/cheng.bee2.pdf>
6. A Wide-Band VLBI Digital Backend System by Alan Whitney, Shep Doeleman, Brian Fanous, Hans Hinteregger, ftp://ivsc.gsfc.nasa.gov/pub/general-meeting/2006/presentations/gm2006_2-02_whitney1.ppt