Options for a VLBA 4-Gbps Upgrade

Jonathan D. Romney
NRAO, Socorro

2006 November 16

0 Introduction

The "VLBA 4-Gbps Upgrade" discussed in this memorandum, part of the larger VLBA Sensitivity Upgrade Project, comprises the replacement of all VLBA systems downstream from the IFs. The primary objective is to enable wideband throughput at a 4 Gbps sustained aggregate data rate. This rate would exploit fully the two 500-MHz IF channels that are generally available, with Nyquist sampling at 2 bits per sample. A 2011completion date is targeted.

Primary constraints are the exclusion of any modifications to the existing receivers and LO/IF system, to limit the overall cost, and a requirement that the upgrade shall not impede progress on the EVLA project.

0.1 Overall Specifications

Input interfaces: two IF channels, each spanning the 500 – 1000 MHz band.

Sampled channels (equivalent to the "baseband channels" in the original VLBA):

- Independently tunable channels: 1 – 16.
- Independent channel bandwidths: 62.5 kHz – 256 MHz.
- Maximum sampled data rate: 4 Gbps.

Output interface formats: FITS-IDI, FITS-UV, EVLA/ALMA, …

0.2 Major Sections

The upgrade will comprise three major sections: the existing samplers, BBCs, and formatter will be replaced by a Sampler/Filter module; the Mark 5 Data Transmission System will be upgraded or replaced (but is expected to remain a disk-based system); and the Correlator will be replaced.

Several options exist for each section. With a few exceptions, the choice among these options can be made independently for each section. The following sections enumerate the options for each section, and discuss their advantages and disadvantages.

1 Options for the Sampler/Filter Module

1.1 EVLA Sampler & WIDAR Station Boards

These units, widely separated and connected by the DTS in the EVLA, in principle could be co-located and connected directly. Their designs already exist, although it might be most effective to replace the optical interfaces to and from the DTS. The EVLA versions support substantially greater bandwidths than can be used for the VLBA 4-Gbps Upgrade, and the unusable capacity might be quite expensive. It may be feasible to mitigate this effect by fractional population of parts.

1.2 EVN/Noto “Digital BBC” (DBBC)

The DBBC is being designed by an EVN group led by Gino Tuccari of the Istituto di Radioastronomia in Italy. It is a very flexible, FPGA-based approach, based on a stack of various numbers of several types of functional boards, with a common data bus. The basic processing unit supports both multiple, independently-tunable, digitally down-converted bands, and equi-spaced (polyphase-filtered) channels. As of the end of 2006, Version 1.1 was undergoing final tests, and at least three such units had been ordered and were under construction. A planned Version 2.0 upgrade would use more modern, faster FPGAs.

This system appears to be appropriate for the VLBA 4-Gbps application, although the minimum channel bandwidth of 250 kHz is four times wider than the specification. A 4-Gbps throughput would require the maximal system supported by the Version 1.1 architecture. Such a system is believed to cost about $85K per station at the current Euro exchange rate.

1.3 Haystack/Berkeley “Digital Backend” (DBE) & Haystack/NRAO Collaboration

Haystack Observatory and the UC Berkeley CASPER group developed a prototype DBE, which includes an existing stand-alone sampler module, and a polyphase filter bank implemented on one of the CASPER general-purpose boards using their high-level FPGA logic definitions. A preliminary design, recently completed, is limited to coarsely-tuned wideband channels. It is well matched to observing modes typically used in geodesy and high-sensitivity continuum astronomy, and to the Haystack Mark 4 correlator, but appears not to support the narrowband spectroscopic observations that are among the VLBA’s most important scientific applications.

A major revision of this preliminary design is beginning at Haystack, and could become a collaborative development with NRAO, with the scope expanded to include specifications tailored to VLBA requirements. NRAO’s contribution would derive from the substantial experience in signal processing with state-of-the-art digital hardware that our engineering staff acquired in designing the EVLA. Indeed, such a joint project was identified as the most promising opportunity for institutional collaboration during a recent Haystack-NRAO technical meeting. It is believed that such a collaborative effort could be mounted without adverse impact on the EVLA project. Preliminary estimates indicate that the VLBA 4-Gbps application could be supported at a cost of about $12 K per station.

2 Options for the Data Transmission System

2.1 Extended Mark 5 System

The Mark 5A recording system currently used by the VLBA is limited to a 512-Mbps throughput by the VLBA formatter, and to 256-Mbps playback per station for arrays of more than 10 stations. Mark 5B is now becoming available, and would raise that throughput to 1 Gbps, but those recordings could not be processed on the existing VLBA correlator without an upgrade that now appears unwarranted in view of the planned replacement of the correlator as discussed in §3. Depending on the timing of the various individual upgrades, this might impact the transition to the full 4-Gbps capability. Mark 5B+, expected to become available in the near future, will record at 2 Gbps, although playback will be limited to 1 Gbps. No known current data-transmission system is capable of 4-Gbps throughput.

Haystack Observatory, the developer of the several Mark 5 systems, remains dedicated to expanding their capability toward greater bandwidths. It can reasonably be anticipated that such further development could lead to a 4-Gbps system by the Upgrade’s 2011 target date. This would allow the VLBA to continue using a familiar system, and would probably preserve the investment (312 K$ at current prices) in disk modules. It might also, however, continue our dependence on Conduant Corporation’s semi-proprietary hardware, requiring at least the new “Amazon” board (already required for Mark 5B+) in place of the StreamStor used in Mark 5A and 5B, and on their marginally supported software. On the other hand, current thinking at Haystack appears to be tending toward the COTS approach described in the next section.

2.2 Totally COTS System

A recording system based entirely on mass-produced commercial equipment and industry standard interfaces should have several advantages over the traditional, highly specialized VLBI systems. Indeed, Mark 5 can be viewed as a transitional stage between these two poles. A COTS approach allows all parts of the system to participate in continuing performance enhancements, following some version of Moore's Law (with possibly differing exponents). And it should do much to minimize the compatibility issues that have plagued VLBI throughout most of its existence.

Credit for the first attempt to develop a recording system along these lines belongs to the PC-EVN system, developed at Metsähovi Observatory in Finland for possible use by the European VLBI Network. PC-EVN was not able to achieve the full 1-Gbps recording specification using a single PC, but the COTS concept remains compelling, and should support substantially wider-band recording using modern components. No specific plans are known to develop a system capable of recording 4 Gbps using this approach, but there is interest across the VLBI community. NRAO probably does not have the necessary expertise, nor the personnel resources, to pursue such a development on our own at this time, but it may be possible to join or organize an international group to do so.

The COTS approach would deliver data in a form more compatible with a software correlator. This connection is one of the few linkages between sections considered for the 4-Gbps Upgrade. A potential disadvantage is that it might be difficult to retrofit the existing disk modules for use with the new system.

2.3 e-VLBI

Pure e-VLBI – direct transmission of data, via wide-area optical networks, from observing stations to the correlator, without intermediate recording or shipment – is an attractive option for future wideband data transmission. It appears unlikely to be a workable approach for 4-Gbps throughput by the Upgrade's 2011 target date. A major obstacle, in the United States, is the high expected cost under current policies in the US telecommunications industry, and the absence of the government subsidies that are prevalent in other countries.

3 Options for the Correlator

3.1 EVLA Correlator

The WIDAR correlator, being built for the EVLA at the Dominion Radio Astronomy Observatory (DRAO) in Penticton, British Columbia, will have a 32-station capacity. The unused capacity beyond the EVLA's 27 antennas could be configured to handle 20 stations at 4 Gbps. VSI-compliant input ports are included on the station boards, although additional chips would have to be added for data format translation. Thus, this system could correlate VLBA observations in parallel with EVLA operations, at a low, and predictable, hardware cost. For that reason, this option is the current default choice for replacing the VLBA correlator.

However, the software costs, and timescale, are both unknown, and there is substantial concern that they may be fairly unfavorable. Allocation of the hardware resources between the EVLA and VLBA applications, without constraining the EVLA, may impose substantial additional complexity on an already complex process. It is also likely that the software development timeline could not accommodate the VLBA application until fairly late.

3.2 Mini-WIDAR

The engineering investment in the EVLA correlator could also be exploited through a small dedicated system using a minimal set of WIDAR boards. Information provided by Brent Carlson of DRAO indicates that 8 station and 4 baseline boards would be required to support 16-station 4-Gbps VLBA observations, assuming standard delay modules were used. (Using special delay modules would reduce the station-board count to 4.) Brent's cost estimate for such a system, based on current costs for the EVLA system, and including rack and power, was about $350 K.

3.3 Software Correlation

A number of software correlators already exist, both for VLBI and for other interferometers. The parallelism of the processing makes such systems natural candidates for operation on computing clusters, and the fundamental throughput restriction by the data transmission system makes VLBI a good candidate for this approach. Although a detailed examination of all candidates has not been done, the DiFX software correlator, written by Adam Deller at Swinburne University, is probably the best option for the VLBA.

Software correlation offers several substantial benefits for the 4-Gbps Upgrade. The extreme flexibility is well matched to the versatility of the VLBA and the range of our users' interests. No development time is required; the software already exists and is available at no cost. The required computing cluster can be procured at anytime, and expanded as necessary, appropriately to the circumstances under which we would start the Upgrade and the final goal. However, the actual cost of the computing cluster required to support the full range of VLBA observations is poorly known.

3.4 Software Correlation, Supercharged

A variant of the preceding option would use the same software system, but target a higher-performance computing cluster, based on a smaller number of cell processors or video cards. It is hoped that the overall cost could be reduced in this way, but the balance between the required number of processors and their unit cost is unknown at present.

3.5 FPGA-Based Design

Intermediate between the ASIC-based hardware exemplified by the WIDAR correlator, and the fully general-purpose hardware that would support a software correlator, this option would use the design tools and standard boards developed by the UC Berkeley CASPER group to achieve a faster development cycle and greater flexibility, but lower performance, relative to the default WIDAR option. No quantitative information is available at present, but we may know more after Walter Brisken's visit to Berkeley in a few months.