CALIFORNIA INSTITUTE OF TECHNOLOGY
VLBA Correlator Memorandum

TO: VLBA correlator group DATE: 31 July 1985

FROM: Steve Kator SUBJECT: The Data Concentrator

The Data Concentrator is intended to provide an inexpensive, off-the-shelf, easily-understood solution to the problem of connecting the sixteen Elementary Correlator Array crates to the Aptec I/O Controller in the Transform Output Processor. The problem arises because the Aptec was designed primarily to connect to a small number of high-speed devices, and the ECAs constitute a moderately large number of medium-speed devices. Specifically, each ECA will be outputing data at a rate of up to 205 kbytes/sec (dumping at 8 Hz, 4 bytes/sample, 400 elementary correlators, 8 real plus 8 imaginary lags each), with the stipulation that the aggregate data rate from all ECAs will not exceed 819 kbytes/sec, corresponding to the case of all ECAs dumping at 2 Hz.

A brief overview of the Aptec will illuminate the need for the Data Concentrator. The Aptec chassis has 15 slots, 13 of which are available for interface controllers or high-speed memory cards. The bus to which these cards connect is allegedly capable of sustaining a transfer rate of 24 Mbytes/sec. The interface controllers come in two flavors: the Device Interface Adapter (DIA) and the Device Interface Processor (DIP). The DIA is the plain version, with one DR11-W port capable of sustaining a data rate of 3 Mbytes/sec. The DIP is a user-configurable version, which in its stock form has up to three parallel protocol ports (such as DR11-W), and can support an aggregate data rate of 9 Mbytes/sec. Each device or device controller (such as a disk drive controller) typically requires its own interface controller card in the Aptec.

A brute-force solution to the interconnection problem is to plug in 16 of Aptec's DIAs, each connected to a VMEbus DR11-W board in each of the ECAs. For this you would need an extender chassis with the basic Aptec to hold the DIAs, the interfaces to the other devices on the Aptec, and the requisite memory cards. Obviously this would be gross overkill on I/O bandwidth. The relative hardware cost of this approach would be 16 times $7000 per DIA + $15000 (approx.) for the extender chassis, or $127,000 total. For the purposes of comparison I am not considering the costs of the basic Aptec configured for our purposes, the VMEbus interfaces, and the cabling, as these items are required regardless of the means used to interconnect the system.

Another all-in-the-Aptec approach would be to purchase a DIP and add some custom circuitry. Aptec sells special DIP boards that have unassigned (bare) areas for just such a purpose. The custom circuitry would consist of a 16-into-1 multiplexed parallel port (probably 16 bits wide), controlled by the processor on the DIP board. This would likely cost the least in hardware, but it is not an

"off-the-shelf" solution. There are many unknowns involved, namely: How do you write the real-time code to control the multiplexer? How do you sychronize it? Is there sufficient room for chips on the DIP board, and if not, how do you get more? It is not comforting that the Aptec salesman was not aware of any other customer who had implemented such a thing. A rough estimate of the hardware costs is $10,000 for the DIP and perhaps $3000 for the additional chips, sockets, and fabrication.

Hence the Data Concentrator appears as a less-expensive alternative to the brute-force approach, and a more "off-the-shelf" one than that mentioned above. The DC would consist of a VMEbus crate containing a 68000 processor board, multiple parallel interface boards to connect to the ECAs, and one DR11-W board to connect to a DIA in the Aptec. Several options exist for the connections to the ECAs. The most straightforward one (which was costed for the Architecture Report) would require one parallel interface for each of the ECAs. The relative hardware cost is $23,200 for the DC crate plus $7000 for the DIA. Another option would be to use the IEEE-488 bus (HP-IB) to connect together four ECAs at a time, such that each quadrant of the Correlator Electronics had its own HP-IB connecting it to the DC. Most currently-available VMEbus HP-IB boards will accomodate 500 kbytes/sec, so that no part of such a system would be close to its performance limit when the correlator is dumping at 2 Hz or less. For higher dump rates there would be a constraint that, say, only two channels per quadrant may be active when dumping at 4 Hz, and only one active when dumping at 8 Hz. The relative hardware cost for this approach is about $13,000 plus $7000 for the DIA, assuming the HP-IB boards cost $1500 each. In both cases, transfers within the DC would be under DMA control.

The unknown variable in this is development cost. It is my belief (or bias, if you will) that any approach that substitutes software running on standard single-board microcomputers for custom hardware of moderate to high complexity will prove to be easier in the long run to develop, modify, and maintain. Of course, this does not hold in the cases where very simple hardware will suffice, and there are many tasks that must be done in hardware because of speed considerations. I maintain that the effort required to design the DC software will compare favorably to that required to design the microcode for any all-in-the-Aptec scheme, especially considering that the DC software environment should be identical to that in several other places in the correlator where VMEbus processor boards are used, such as the ECAs, the SMPs, and the PWG. The question of whether it would be less expensive to develop custom hardware for this application rather than employ a mostly software-driven system is still an open one and well could be debated for many correlator lifetimes.