

Maximum Data Rates from GMRT Correlator

crs/26.6.91

(earlier versions are obsolete)

INTRODUCTION

In view of the difficulties faced in routing signals from 16 FFT engines into a single MAC card to accommodate 32 ASICs per card AND support (RR,LL) mode, we have decided to double the number of MAC cards with 16 ASICs per card, but include buffers within the MAC card in order to facilitate a time-resolution of 4 ms without sacrificing bandwidth and without any need for losing data between FFT cycles. The details of the multiplier card will be provided in a separate document, but we summarise below the main features relevant for the present purpose.

In the revised design of GMRT multiplier cards, there will be 16 ASICs in each multiplier card and, in addition, there will be 32 buffers (18-bit wide) which can accommodate one complex number (equivalent to two 18-bit words) from each of the 16 ASICs in the card. During the 4-tick pause, all the ASICs will download word from identical address pairs into these buffers. The buffers can be individually addressed as 18-bit words and read through a pair of 18-bit busses available in the MAC backplane. Each bus will access upto 5 MAC cards, requiring a total of $32 \times 5 = 160$ operations for reading all the buffers (one complex word from all the MAC cards.) This has to be done in the available 512 clock ticks at 32.25 MHz, corresponding to a 10 MHz operation of the backplane bus. In practice, read-out will be at 16 MHz, requiring 4 clock ticks to read one pair of complex words. These will then be reduced from 36-bits to 32 bits per complex word, and written into the buffers provided in the MAC control card. In the initial implementation, this will simply be a change from (15,6,15) to (13,6,13), thus losing two significant bits from both real and imaginary parts. The results will be stored in the MAC control card in RAMs organised as 32-bit words, to enable an external circuitry access a 32-bit complex word in one operation.

In order to simplify the memory arbitration, the memory in the MAC control card will be organised into 4 memory banks of 16k x 32 bits. During any given FFT cycle, only two of these banks will be in use; the other two will be free to be read by external circuitry.

ACCESSING STA RESULTS

The output of data from the GMRT correlator has four distinct stages. The first stage will occur during the 4-tick pause in each FFT cycle of 516 ticks, and the other stages will occur during the active period of duration 512 ticks in each FFT cycle.

The first stage is essentially an internal operation where every ASIC in the multiplier cards will output one STA result, a complex word of 36 bits into separate buffers (2 x 18 bits per ASIC) provided for each multiplier chip.

In the second stage, one STA result (36-bit complex) from every ASIC is loaded into the appropriate MAC control card, converted to a 32-bit complex word - (13,13,6) format - and stored in appropriate locations in 32-bit memory banks provided in each MAC control card. There will be 4 memory banks in each MAC control card, but only two will be active during any given FFT cycle. The other two banks in each MAC control card can be freely accessed by any external circuitry.

In the third stage, the STA results are read from the memory banks in the MAC control cards into the long-term-accumulator (LTA). The LTA will provide various options of adding the data. In particular, they can be added across the spectral channels, thus reducing the number of spectral channels over the given band, or/and add contiguous samples of each spectral channel to provide any desired integration. The net result of these options is to effectively reduce the data rate to a maximum limit of 1 MB/s currently imposed on the data acquisition from the entire correlator system. In our initial design, the LTA will be multi-cpu array, probably based on the transputers. A possibility being considered is to have the processor within the MAC control card. This will be adequate for reading the STA results at an average rate of once every 128 ins, but the option will be left open for a more powerful system to access the results at much faster rate in some future version.

In the fourth stage, the results from LTA will be accessed by the data acquisition computer (SUN 1e single-board-computer based on SPARC chip) which will provide additional options of selecting! adding data to further modify spectral resolution or integration time. At this stage, it should also be possible to correct the correlations for individual gains and relative instrumental phases at each antenna, and, if necessary, allow on-line flagging of data to the extent feasible. However, such options are only likely to be provided in the software after some initial experience with the array. After these operations, the maximum data rate from the entire correlator system is expected to be below 128 kB/s, which can be sustained on an Exabyte tape archival. The data will be routinely recorded on Exabyte tapes during the data acquisition. In addition, the data are simultaneously transferred to the on-line file server which will may also store them in its disk. Any computer networked to the file server should be able to access the visibilities after they are registered at the file-server.

TIME RESOLUTION

The time resolution is decided by the minimum integration time that can be provided at the output of the LTA. The overall design limitation corresponds to 4 ms or perhaps 1 ins. However, this only corresponds to the capability promised in the first two stages mentioned above. In practice, a severe restriction will arise from the LTA whose initial design may not be able to handle such high data rates and limit the time resolution to 128 ins. Further

limitation will arise from the archival rate if time solution is to be retained without sacrificing spectral resolution.

We have followed a very modular approach in the design such that time resolution can be improved in future designs when fast processors become economical and easily accessible. In the following paragraphs, we will outline the limitations of our design for the first years of operation of GMRT.

OPTION A (The Standard Option) An option which will certainly be available corresponds to reading all the relevant STA results once every 128 ms (131.07 ns to be precise). While we have not finalised our initial design of LTA, it is possible to realise it by a fast processor (e.g. T800) provided within the MAC control card. The data acquisition system planned for initial operation of GMRT cannot sustain this data rate. Hence it will be necessary to reduce the effective data rate (by about a factor of 16) either by adding adjacent spectral channels, thus reducing the spectral resolution, or sacrificing time-resolution by increasing the integration time, or by selecting specific channels for data acquisition, or by allowing dead-time between consecutive records of data. All these are software options which can be provided as necessary in the LTA and the data acquisition system.

OPTION B (low-resolution) We are working on changing the design of address generation for multipliers, wherein we can cycle through fewer (e.g. 16) addresses for each chip, instead of the standard 256. This will be applicable to the entire array and not specific to any subset of baselines. The proposed options are 16, 128 or 256 channels such that the FFT size is precisely twice the number of channels, i.e., 32, 256 or 512 points respectively. The last two correspond to the highest spectral resolution in the polarization and non-polarization mode respectively. Intermediate resolutions can be generated in the LTA by suitable software. At the worst resolution of 16 channels, the effective delay step introduced by the fractional-sample-time~ correction (FSTC) corresponds to 1 ns.

With the option of low-resolution (16 spectral channels), a time resolution of 8 ns will have the same net data rate at the LTA as the 128 ns in the standard mode providing for full spectral resolution.

Further details of options available in the cross-multiplier system will be enumerated in a future document on the design of multiplier system.

DATA ACQUISITION

The data acquisition will be performed on the SPARC-based single board ~xnPuters (Sun le) running the real time operating system VxWorks. Although our initial data acquisition system will have only one of these computers in a VME chassis, future upgrades may have more of similar VME boards to provide faster data acquisition

The current software specs allow a maximum Of 1 MB/s data rate between LTA and the data acquisition system. But the data rate will further be reduced in the data acquisition software depending on the archival media handling capabilities. Presently, we have prescribed a maximum data recording rate to 16k complex channels per second, which corresponds to 128 kB/s without data compression. With suitable data compression, the speed can further be reduced for recording on the archival media (Exabyte tapes). The tape requirement corresponds to about two 8 mm video tapes (Exabyte) per day. Each certified tape costs about \$13

TYPES OF ARCHIVED RECORDS

The data recorded on Exabyte tape will be in compressed format. The first record will be a header record which gives adequate information the structure of the data to follow, a time-reference and a unit of time used in the records. it will be possible to allow multiple data structures by inserting a new header record anywhere in the tape. The end of a session will have a trailer record whose contents are mostly derivable from the contents of most recent header

record, and the time offsets present in every data record.

There is no limit on the number of header/trailer records between two tape filemarks. The tape file-marks are only used to denote end-of- information (equivalent to double eof marks). A record identifier is built into the first byte of every record. The subsequent 3 bytes will be used in the data records to indicate a time offset with reference to reference time in the header/trailer; the unit of time offset is also present in the header/trailer.

DATA COMPRESSION

The data records mentioned above will have the correlation coefficients recorded in Compressed floating point format. The exact format has not been finalised, but I give below 4 possible alternatives as an indication. In two of these formats, a complex number is represented as two floating point numbers (real and imaginary parts); in the other formats,