

THE PROPOSED GBT CORRELATOR

Carl Heiles

1. 'ORDINARY' SPECTRAL LINE OBSERVATIONS.
 - 1.1. BRIEF SUMMARY OF THE CORRELATOR'S CAPABILITIES.
 - 1.2. SPECIFIC QUESTIONS AND ISSUES.
 - 1.2.1. THE FIRST LOCAL OSCILLATOR.
 - 1.2.2. THE NUMBER OF OPTICAL FIBERS.
 - 1.2.3. THE NUMBER OF CORRELATOR INPUTS AT 'NARROW BANDWIDTHS'.
 - 1.2.4. THE NUMBER OF INPUTS AT 1 GHZ BANDWIDTH
 - 1.2.5. THE NUMBER OF INDEPENDENT LOCAL OSCILLATORS FOR THE 32 INPUTS.
 - 1.2.6. AVAILABLE BANDWIDTHS.
2. 'PULSAR' SPECTRAL OBSERVATIONS.
3. APPENDIX. CORRELATOR HARDWARE ESSENTIALS: AN ASTRONOMER'S DESCRIPTION.

1. 'ORDINARY' SPECTRAL LINE OBSERVATIONS.

By the word, 'ordinary', we mean observations not requiring more than the basic time resolution of the correlator, which is 16.777216 msec.

This section summarizes the properties of the correlator that have been settled on by a number of astronomical considerations and agreed to by the Scientific Working Group (SWG).

1.1. BRIEF SUMMARY OF THE CORRELATOR'S CAPABILITIES.

SAMPLERS: The correlator has two types of sampler: fast (2 GHz clock rate) and slow (125 MHz clock rate). The correlator has four input bandwidths: 1 GHz, 250 MHz, 62.5 MHz, and 15.625 MHz. Normally, the two larger bandwidths are used with the fast samplers and the two smaller bandwidths with the slow samplers. However, one may use the fast samplers with the 62.5 MHz bandwidth, which would allow double-Nyquist sampling (see below). There are 8 fast and 32 slow samplers; each sampler services one input to the correlator, and provides one spectrum.

BASIC CORRELATOR CONFIGURATION: The correlator is constructed in quadrants; each quadrant is fed from the samplers by a switching network. Each quadrant provides a total of NQ channels, where NQ depends on bandwidth as described below. The quadrants have the capability of operating independently. Thus, the quadrants can be run at different bandwidths and with different numbers of simultaneous inputs. In the extreme case when all 8 inputs in a quadrant are used, the number of channels per spectrum is NQ/8. Alternatively, the quadrants can run as a single gigantic correlator with 4NQ channels.

NUMBER OF CORRELATOR CHANNELS: For the 1 GHz and 250 MHz bandwidths, the total numbers of correlator channels $4NQ$ are 16384 and 65536, respectively. For the small bandwidths, $4NQ = 262144$ (256K) channels. These channels can be divided up among the various inputs. Thus, with 32 inputs and 62.5 MHz bandwidth, the number of channels per input is $256K/32 = 8K = 8192$. (There is one exception to this rule: for 250 MHz bandwidth with 8 inputs, intricacies of the correlator design prevent there from being more than 4096 channels, while one would have expected 8192.

The correlator can perform cross-correlation for the purpose of polarimetry. For example, with two inputs from a single feed one produces two auto- and two cross-correlation functions to obtain the four Stokes parameters, and of course each cross product requires its own set of channels; for example, to obtain complete polarization information at 62.5 MHz bandwidth with 16 dual-polarized feeds, one has 32 inputs and obtains 64 spectra, each with 4096 channels.

DOUBLE NYQUIST AND 9-LEVEL SAMPLING: This is a 3-level correlator. Traditionally, correlators are operated at the Nyquist rate, at which the sampling rate is twice the bandwidth. With such operation, a 3-level correlator increases the noise by about 24% relative an analog correlator. At bandwidths smaller than the maximum sample rate, 'double Nyquist' sampling (sampling rate is 4 times the bandwidth) reduces sacrifice to about 11%.

With the fast samplers, double Nyquist sampling is possible at 250 MHz and 62.5 MHz bandwidths. Note that it is possible to operate the 'slow' 62.5 MHz bandwidth with the fast samplers, but then one is of course restricted to only 8 inputs. With double Nyquist sampling at these bandwidths, there is a penalty of a factor of two in total number of channels. Thus $4NQ = 65536/2 = 32768$ for 250 MHz and $4NQ = 262144/2 = 131072$ for 62.5 MHz.

With the slow samplers, double Nyquist sampling is possible at 31.25 MHz bandwidth and below with no penalty of any kind whatsoever (!). This remarkable fact is owing to a design feature in the basic correlator chip. Thus, double Nyquist sampling is the normal mode of operation at 15.625 MHz bandwidth.

Nine-level sampling/correlation reduces the signal/noise sacrifice to essentially zero. The slow sampler is capable of providing 9 levels. Thus, for both of the small bandwidths, one can gain signal/noise and perform 9-level sampling with single (or double at 15.625 MHz) Nyquist sampling. However, there is a penalty: a four-fold reduction in number of correlator channels. Thus, for 9-level sampling, $4NQ$ drops to $262144/4 = 65536$.

1.2. SPECIFIC QUESTIONS AND ISSUES.

The SWG considered a number of specific questions regarding capabilities of the correlator. We summarize the questions and conclusions below.

1.2.1. THE NUMBER OF FIRST LOCAL OSCILLATORS.

Each first l.o. can feed one or more receivers that cover a single r.f. band. We are providing two first l.o.'s. With two first l.o.'s, no more than two r.f. bands can be used simultaneously. This was not considered to be a serious limitation, particularly because increasing the number of l.o.'s either temporarily, for experimental purposes, or as a standard modification in the future requires no fundamental changes in system design. Having only a single first l.o. at the beginning was considered a satisfactory compromise, if economically necessary.

1.2.2. THE NUMBER OF OPTICAL FIBERS.

After conversion from r.f. to i.f. by the first l.o., the i.f. signals are carried from the front end to the control room by optical fibers. There are 8 optical fibers, each carrying a single i.f. with up to about 8 GHz bandwidth. In the initial stages of GBT operation, we do not plan to multiplex more than one receiver per fiber. Thus, the 8 fibers allow a total of 8 separate receivers. Normally, these 8 might be two frequencies, two feeds at each frequency, two polarizations for each feed.

The restriction of 8 was not considered to be serious until the advent of large-scale multibeam systems. Having only 4 optical fibers at the beginning was considered a satisfactory compromise, if economically necessary.

1.2.3. THE NUMBER OF CORRELATOR INPUTS AT 'NARROW BANDWIDTHS'.

We are providing 32 narrow-band inputs, with the possible option of increasing this number to 64 in the future. Increasing this to 128, as a provision for the future, is possible but technically awkward, and was not considered important by the SWG. The 8 fibers, each of which carries a single receiver's i.f., might seem to impose a natural limit of 8 on the number of correlator inputs. However, there can be more than one correlator input per optical fiber, because within each fiber's i.f. there can be many more than one spectral line. Our decision for 32 inputs is derived from several discussions involving simultaneous coverage of spectral lines and multibeaming, and the possible increase to 64 inputs in the future may become appropriate when serious focal-plane arrays come into operation at the GBT. A smaller number of narrow-band inputs at the beginning would be considered a satisfactory compromise, if economically necessary.

1.2.4. THE NUMBER OF CORRELATOR INPUTS AT 'WIDE BANDWIDTHS'.

The maximum number of wide-band correlator inputs is 8. Scientific considerations dictate that it is important to have all 8 inputs. A smaller number at the beginning would be considered a satisfactory compromise, if economically necessary.

1.2.5. THE NUMBER OF INDEPENDENT LOCAL OSCILLATORS FOR THE 32 INPUTS.

Very important from the economic standpoint is the question of how many l.o.'s we require to service the correlator inputs. An l.o. is required to convert an i.f. signal coming down the optical fiber to the correlator input frequency. The i.f. frequencies on the fibers lie in the range of about 1 to 10 GHz, and the l.o.'s must cover this range, which makes them very expensive.

Nevertheless, the SWG felt quite strongly that it is important to have 16 independent, completely flexible local oscillators. This provides one l.o. for every 2 slow samplers, which is reasonable because most observers want to observe the two polarizations from a single feed simultaneously to enhance signal/noise. The engineering department is continuing to search for more economical alternatives to these l.o.'s. A smaller number of l.o.'s at the beginning would be considered a satisfactory compromise, if economically necessary.

1.2.6. AVAILABLE BANDWIDTHS.

The SWG agreed that the four bandwidths we plan are sufficient. Providing additional bandwidths in the future, either narrower ones or more closely spaced ones, is an easy addition should it prove necessary.

2. PULSAR SPECTRAL OBSERVATIONS.

It is possible to start the correlator at a precise time by feeding it a pulse. The design may allow achieving time resolution as good as 1 microsec with the sacrifices of (1) incomplete coverage of a pulsar period, (2) a reduction in the number of channels per spectrum, and (3) a reduction in the number of spectra for bandwidths larger than 62.5 MHz. We are currently examining this and other possibilities for high time resolution spectra, and the final capabilities may depart from those outlined herein either for better or worse. Below, and in the appendix, we provide a brief explanation of how the correlator achieves such time resolutions.

Consider operating with bandwidth 62.5 MHz. The basic integration time of the correlator is 16.8 msec. One can obtain better time resolution by blanking individual chips. This blanking can be done with an externally generated signal, for example a signal in phase with a pulsar. If one wished to cover a pulsar pulse completely with the best possible time resolution, then one would divide the pulsar period into 256 equal parts. (Alternatively, one can use fewer parts and get worse resolution).

Suppose, for example, the pulse period is 2.56 msec. Then the time resolution would be $2.56 \text{ msec}/256 = 10 \text{ microsec}$. One could obtain shorter time resolution by sacrificing the full, complete coverage of the pulsar pulse; for example, one could obtain 3.33 microsec resolution by restricting one's coverage to 1/3 the pulse period. The blanking signals simply turn on each of the 256 chips for the appropriate time interval and turn it off the rest of the time. Thus, during each 16.8 msec we obtain 256 separate spectra, each timed according to the externally-applied blanking signal; the blanking signal time-slices the spectra according to pulsar phase. At the end of the 16.8 msec interval we transfer the results to the long-term-accumulator (LTA), which is an addressable integrating memory for the correlator chips; the location in the LTA to which each spectrum is transferred can be specified by either an internal calculation or an externally-generated signal. Thus the 256 separate spectra are time-sliced within the LTA according to pulsar phase. After accumulating these 16.8 msec integrations for some desired time, we transfer the results to the control computer.

At 62.5 MHz bandwidth, the above technique provides 256 (or fewer, in steps of factor-of-two) spectra, each 1024 (or more, in steps of factor-of-two) channels. For larger bandwidths, each spectrum is 1024 channels but the number of spectra is only 256 times (62.5 MHz/bandwidth) (or fewer, in steps of factor-of-two).

3. APPENDIX

CORRELATOR HARDWARE ESSENTIALS: AN ASTRONOMER'S DESCRIPTION

This appendix is a distillation and translation of Ray Escoffier's July 1993 Long Term Accumulator (LTA) memo and subsequent, ongoing discussions. Thus, the descriptions given here may evolve with time. Furthermore, the correlator has many possible modes of operation, which are initiated by software. Not all will be made available for the user. Generally speaking, features not mentioned in the main body of this memo will not be made available unless astronomers specifically request them in subsequent written documents.

1. THE BASIC CORRELATOR CONFIGURATION.

The correlator consists of four identical quadrants. Each quadrant contains four correlator boards; each correlator board contains 16 chips; each chip contains 1024 channels and can clock at 125 MHz. Thus, for input bandwidths of 62.5 MHz (sample rate of 125 MHz, the chip clock rate), we obtain the full complement of channels: 4 quadrants times 4 correlator boards times 16 chips times 1024 channels = 262144 channels. For higher sample rates, individual chips are fed at their maximum rate of 125 MHz and each correlation function is spread over more than one chip. This makes the obtainable size of the final correlation function $NC = (62.5 \text{ MHz/bandwidth}) \text{ times } 262144$. At the maximum bandwidth of 1 GHz, we obtain a total of $NC = 16384 \text{ channels--}4096$ per quadrant.

The quadrants are identical and can be operated independently or together. Thus, each quadrant can run at its own bandwidth and can be divided into sections, independently of the others. Each quadrant can be divided into 16, 8, 4, 2, or 1 equal pieces and one could have as many as 64 independent correlators by dividing each quadrant into 16 pieces; this sets the maximum possible number of inputs to the correlator. With 64 independent correlators at 62.5 MHz bandwidth, one obtains 64 independent spectra, each with $262144/64 = 4096$ channels. Alternatively, one can do the equivalent of hooking one or more quadrants together in series with the extreme being a single gigantic correlator having 262144 channels for bandwidths less than 62.5 Mhz (in the hardware, this is accomplished by redirecting the output of the RAM described below).

There are two types of sampler, the fast (2 GHz sampler) and the slow (125 MHz) sampler.

First we discuss the fast samplers. There can be as many as 8 samplers. A sampler does not feed the correlator chips directly, because at 2 GHz the sampler operates 16 times faster than the chips. Rather, a sampler feeds one or more RAM's which dole out the samples to 16 chips, appropriately 'packetized', each at the 16-times slower rate of 125 MHz, the chip clock rate. We use the word 'packetized' because contiguous samples are not multiplexed across chips. Rather, contiguous channels in the correlation function occupy contiguous locations on a chip (this particular point is a crucial element in the conceptual design). There are two RAM memories per quadrant. Each RAM memory is large enough to feed 131072 contiguous samples to each chip. At the 125 MHz clock rate for the chips, the 131072 contiguous samples are transferred to the chip in 1.048576 msec. **THIS TIME INTERVAL, 1.049 MSEC, IS CALLED THE 'BURST PERIOD' AND IS A FUNDAMENTAL DESIGN PARAMETER OF THE CORRELATOR.**

The slow (125-MHz) samplers also feed the RAM's, which feed the correlator appropriately for the selected configuration. In principle, we could provide for future expansion to one sampler for each group of two chips--128 samplers--but this is difficult. Instead, we intend to provide for a maximum future expansion to 64 samplers, one for every four chips. Initially, we plan to provide 32 samplers. At bandwidths narrower than 62.5 MHz, the clock rate remains 125 MHz and samples are duplicated.

2. TIME RESOLUTION AND INTEGRATION.

The correlator chips have enough short-term integration capacity for thousands of burst periods, but are read out every 16 burst periods, or every 16 times $1.049 = 16.777216$ msec. **THIS TIME, WHICH WE DESIGNATE $TR = 16.8 \text{ MSEC}$, IS THE BASIC TIME RESOLUTION OF THE CORRELATOR;** however, better resolution is achievable as described below. Each 16.8 msec, a chip dumps into an integrator called the Long Term Accumulator, or LTA. The LTA can accumulate an almost arbitrarily large integral number of 16.8 msec periods before dumping the integrated correlation functions to the control computer.

Each correlator quadrant has its own LTA. Each LTA contains 1M (1048576) 32-bit words. Each quadrant has 64 chips, or 65536 channels. Thus the LTA contains 16 words per correlator channel. These can be allocated in different ways as required. Normally, the LTA is split into two 8-word portions for the purpose of double-buffering, which allows one to accumulate new data while transferring data to the control computer, but for some purposes a user may wish to use all 16 words for data storage and sacrifice the data transfer time. Beyond this, for example, one can break each portion up into eight parts and accumulate up to eight different spectra as various versions of 'signal' and 'reference', or one can save fewer than the full number of correlator channels and accumulate more than eight different spectra.

Or one can map time into the LTA. The fundamental time resolution of the LTA is $TR = 16.8$ msec. One could integrate for, say, 8 times $TR = 134$ msec and store the result in the first 1/16 of the LTA; integrate for a second 134 msec and store the results in the second 1/16 of the LTA, repeating this process up to 16 separate times to generate 16 independent spectra, each with 65536 channels; then repeat the process to accumulate additional integration time if desired, and finally dump to the host computer. Or one could do the equivalent but generate a larger number of independent spectra by saving only some of the 65536 channels. These examples have time boundaries as multiples of the correlator time resolution TR ; alternatively, one could generate the address bits of the LTA memory from a computation of the phase of a pulsar, thus generating 16 or more spectra in synchronism with the pulse phase.

One can also generate such spectra with time boundaries as submultiples of TR --SHORTER than TR , and then store the results in the LTA according to pulsar phase. One can obtain submultiples of TR because the correlator chips can be blanked, turning them off.

Even though the chips are grouped in units of two, they can be blanked individually (!), so one can break the correlator into $T2$ sections ($= 256, 128, \dots, 2$) and obtain a number of separate spectra. The number of separate spectra is not equal to $T2$. Rather, it is equal to $T2$ times (62.5 MHz/bandwidth); each separate spectrum has $262144/T2$ channels. For example, with $T2 = 256$, one obtains 1024 channels per spectrum; one obtains 256 spectra for 62.5 MHz bandwidth, or 16 spectra for 1 GHz bandwidth. The spectra have time resolution determined by the blanking, which can be provided by an external signal (synchronized with a pulsar, for example), and can be as short as about 1 microsec; with such high time resolution, one does not cover the full period of a pulsar. The integration is accumulated in the LTA, with the address to which each spectrum is written being specified either by an internal computer calculation or an external signal (generated in synchronism with a pulsar, for example). The LTA is dumped to the control computer after a user-selected integral number of TR periods.

3. DUMPING RAW DATA.

An interface will be provided to get the raw results from the correlator chips as they are being dumped into the LTA. This will supply the results from an entire correlator card each burst period of 1.049 msec. In this time, at the 125 MHz clock rate only about 17 bits per channel will be nonzero.