From scurry@sadira.gb.nrao.edu Fri Jul 23 11:05:14 1993
Date: Fri, 23 Jul 93 12:59:32 EDT
From: scurry@sadira.gb.nrao.edu (SHIRLEY CURRY)
To: gbtswg@sadira.gb.nrao.edu
Sender: gbtswg-request@sadira.gb.nrao.edu
Content-Length: 24048

The Science Working Group will have a meeting July 27, 4:00 EDT. The call-in number is 913-749-9248, the ID # is R302.

This memo is background information for the SWG meeting scheduled for 27 July 1993 and provides information on the current version of the proposed correlator. During the SWG meeting we need to discuss, in particular, the specific questions and issues in section 1.2. Note particularly section 1.2.8.

OUTLINE OF THIS MEMO

****

1. 'ORDINARY' SPECTRAL LINE OBSERVATIONS.

By the word, 'ordinary', we mean observations not requiring time resolution faster than about 16.8 msec.

1.1. BRIEF SUMMARY OF THE CORRELATOR'S CAPABILITIES IN ITS FINAL, ULTIMATE VERSION.

In this ordinary mode of observing, the basic time resolution of the correlator is 16.777216 msec. There are two bandwidth modes, the 'small bandwidth' and 'large bandwidth' modes.

For small bandwidths, 62.5 MHz and below, there are 256K channels and up to 32 inputs. This number of 32 inputs specifies the number of analog devices--samplers and input filters--that feed the correlator. The digital section of the correlator can handle up to 64 inputs, and increasing the number of inputs will be retained as an option for the future. The Scientific Working Group (SWG) meeting of 11 May discussed which bandwidths should be provided and settled on 62.5 and 15.625 MHz as being essential. Additional bandwidths require the construction of additional analog devices to feed the correlator, but not the modification of the digital section of the correlator, and can be provided rather easily as options for the future.

For large bandwidths, up to 1 GHz maximum, the number of channels is equal to 256K * (62.5 MHz/bandwidth). The digital section of the correlator can handle up to 8 inputs. We intend to provide samplers and filters for all 8 inputs.

This is a 3-level correlator. Normally, correlators are operated at the Nyquist rate, at which the sampling rate is twice the bandwidth. With such operation, a 3-level correlator sacrifices about 28% signal/noise. At bandwidths smaller than the maximum sample rate, 'double Nyquist' sampling (sampling rate is 4 times the bandwidth) reduces the signal/noise sacrifice to about 8% (8% is a guesstimate); this will be possible below (but not at) 1 GHz bandwidth in the wide band mode and below (but not at) 62.5 MHz in the narrow band mode. Double Nyquist sampling normally requires sacrificing half the channels. However, a unique feature in the chip design allows double Nyquist sampling with NO SACRIFICE OF ANY KIND for bandwidths less than 62.5 MHz. Thus, double Nyquist sampling will be the usual case below (but not at) 62.5 MHz bandwidth. For the small-bandwidth mode--including 62.5 MHz bandwidth--one can go further in signal/noise and perform 9-level sampling with single (or double below 62.5 MHz) Nyquist sampling; 9-level sampling would reduce the signal/noise sacrifice to essentially zero, but at the cost of a four-fold reduction in number of correlator channels.

1.2. SPECIFIC QUESTIONS AND ISSUES.

1.2.1. THE FIRST LOCAL OSCILLATOR. Initially, the receivers will cover r.f. frequencies up to 50 GHz; eventually, we will operate at frequencies up to 115 GHz. Accordingly, the first l.o. must cover an enormous range. Such oscillators are very expensive.

At first, we will have only two first l.o.'s. Each l.o. can feed one or more receivers that cover a single r.f. band. With only two first l.o.'s, no more than two r.f. bands can be used simultaneously. Is this a serious limitation? Note that all feeds will see the sky simultaneously, but each feed will be centered at a different sky position. Some types of program--for example, surveys whose importance is measured in area of sky and frequency range covered and not in terms of what particular positions are actually observed--might wish to use as many different r.f. bands, simultaneously, as possible, and having only two first l.o.'s would be a limitation.

It is conceivable that we could trade oscillator cost for quality and capability. Two examples: One, somewhat degraded spectral purity would be acceptable for spectral line work but not for VLBI. Two, less precision (fewer digits in the synthesizer) could be used by performing Doppler corrections after observing in the computer; there is, in principle, no sacrifice whatsoever with this technique and it even has advantages when dealing with interference. We are currently considering various options along such lines.

There may be enough money for only one first l.o. at the beginning. How serious is this limitation?

1.2.2. NUMBER OF OPTICAL FIBERS. The i.f. signals are carried from the front end to the control room by optical fibers. There will be 8 optical fibers, each carrying a single i.f. with up to about 8 GHz bandwidth; we do not plan to multiplex more than one receiver per fiber. Thus, the 8 fibers allow a total of 8 separate receivers, which would normally cover two r.f. frequencies (because we will have only two first l.o.'s; see 1.1.1 above), two feeds at each frequency, two polarizations for each feed.

There may be enough money for only 4 fibers at the beginning. How serious a limitation is this?

1.2.3. NUMBER OF CORRELATOR INPUTS AT 'NARROW BANDWIDTHS'. Each correlator input incurs a significant cost: each requires its own input filter (one for each bandwidth) and sampler. 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. We plan on 32 inputs, half the maximum possible number of 64. This number of 32 is derived from several discussions involving multibeaming and simultaneous coverage of spectral lines. One of these discussions occurred during at the last SWG meeting. When serious focal-plane arrays come into operation at the GBT in the future, it may be appropriate to increase the number of inputs to 64. Is there any need to do so now?

If 32 inputs costs too much, how serious is reducing the number of inputs to 16 at the beginning?

It might be possible to increase the maximum possible number of correlator inputs (as an option for the future) from 64 to 128. When using all 128 inputs, 9-level sampling would not be possible. Providing this, even as an option for the future, involves significant design effort and we prefer to exclude this possibility. How important is this?

1.2.4. NUMBER OF INPUTS AT 1 GHZ BANDWIDTH. At 1 GHz bandwidth, the maximum possible number of correlator inputs is 8. 8 inputs is desirable for the most demanding type of observation at 1 GHz bandwidth, searching for high-z emission from galaxies near 50 GHz. The 1 GHz bandwidth corresponds to a velocity coverage of 6000 km/s and we require a velocity resolution of 5 km/s, so we require about 2000 channels (with Hanning smoothing). At 1 GHz bandwidth, the full correlator has a total 16K channels; this leads naturally to 8 inputs.

The fast (2 GHz clock rate) samplers may be expensive. How serious would be limiting the number of 1 GHz inputs to four?

1.2.5. 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 inputs. An l.o. is required to convert an i.f. signal to the correlator input frequency. The i.f. frequencies lie in the range of about 1 to 10 GHz, and the l.o.'s must cover this range, which makes them very expensive; it is important to minimize the number. It is not always necessary to have a separate oscillator for each input. For example, a multibeam study at a single frequency, with many separate correlator inputs at a single i.f. frequency, requires only one oscillator. For multitransition studies, one oscillator per transition is required.

For 'large bandwidths,' we should have at least 4 l.o.'s, which is half the number of fast samplers; this allows a standard mode of observing with two fast samplers per frequency, one for each polarization. Better would be one l.o. per fast sampler, because a user is likely to want to trade sensitivity (looking at only one polarization) for frequency coverage (8 different center frequencies).

For 'small bandwidths', 8 l.o.'s seems to be an appropriate minimum number. 8 l.o.'s would provide one l.o. for every FOUR slow samplers, which is reasonable because at many frequencies there are two feeds, each of which has two polarizations, and one popular mode of observing will probably be to look at the same frequency with both feeds--i.e. four correlator inputs for each frequency. Better, of course, would be 16 l.o.'s for the small-bandwidth mode. This would provide full spectral capability for the 32 correlator inputs if each frequency had two inputs, one for each polarization.

How important is it to have the full flexibility at narrow bandwidths, which requires 16 l.o.'s? If resources become tight, how bad is it to restrict the number of l.o.'s to four at the beginning?

1.2.6. WHAT BANDWIDTHS SHOULD WE PROVIDE? At present, we envision the correlator bandwidths covering the full range from 1 GHz to 15.625 MHz in incremental steps of a factor of four. However, it may not be very expensive to provide octave bandwidth coverage over this range. Furthermore, we could go even narrower than 15.625 MHz bandwidth if necessary. (But recall that with 32 inputs and 256K channels, the frequency resolution at 15.625 MHz bandwidth is 3.8 kHz with Hanning smoothing--which is pretty good--and that the resolution gets better with fewer inputs).

Should we provide a narrower minimum bandwidth than 15.625 MHz?

Should we provide octave bandwidth coverage for some range of bandwidths?

Note that these capabilities are easy to add in the future because they do not affect the digital section of the correlator.

1.2.7. HOW IMPORTANT IS ENHANCED CAPABILITY AT 250 MHZ BANDWIDTH? At present, the design has two types of sampler: fast (2GHz clock rate) and slow (125 MHz clock rate). We are considering--very reluctantly--the possibility of increasing the slow sampler clock rate sampler to 500 MHz. This would be significantly more expensive and complicated, and might not even be economically feasible with current technology. However, it would allow the number of inputs at 250 MHz bandwidth to be 32 (equal to the number of narrow-band inputs) instead of 8 (the number of 1 GHz inputs). (At the same time, the increased expense might decrease the number of narrow-band inputs!) It would also allow 9-level sampling at 250 MHz, which would increase the signal/noise by about 28% at this bandwidth.

Astronomically speaking, we are guessing that this possibility is not very attractive in view of the small interest in the 250 MHz bandwidth option expressed at the last SWG meeting. Is this correct?

1.2.8. THE GRIM REALITY: A PROBABLE MINIMALIST CONFIGURATION AT THE BEGINNING. The correlator budget remains highly uncertain, both because the design is still in progress and because the cost of the correlator chips is not known. We forsee the probable need to provide a minimilist system at the beginning. Its current design has one first l.o., 4 fibers, 2 second l.o.'s, and 4 correlator inputs. This is very limited compared to what we intend to provide as a final system, and quite limited compared to the compromises suggested in sections 1.2.1 to 1.2.7.

How serious are the limitations imposed by this probable minimalist configuration? If we have enough money to expand it, what are the priorities--what should be expanded first?

2. PULSAR SPECTRAL OBSERVATIONS.

It will be 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 msec/256 = 10 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 desired 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 spectra are time-sliced according to pulsar phase. At the end of the 16.8 msec interval we transfer the results to the 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 * (62.5 MHz/bandwidth) (or fewer, in steps of factor-of-two).

****

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--and 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; the chips are grouped into units of two; 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 * 4 correlator boards * 16 chips * 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 MHz/bandwidth) * 262144. At the maximum bandwidth of 1 GHz, we obtain a total of NC = 16384 channels--4096 per quadrant, 1024 per correlator board.

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 group of two chips works together and is the minimum unit. 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 the chips, appropriately 'packetized', 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; and we plan to actually provide only 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 will be read out every 16 burst periods, or every 16 * 1.049 = 16.777216 msec. THIS TIME, WHICH WE DESIGNATE TR = 16.8 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 will be 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 * 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 submultiplies 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, ..., 2) and obtain a number of separate spectra. The number of separate spectra is not equal to T2. Rather, it is equal to T2 * (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 will be nonzero.