MMA Memo 194
Astronomical Capabilities of the Current Design for the Millimeter Array Correlator

Michael P. Rupen
National Radio Astronomy Observatory
Socorro, NM 87801
&
Ray Escoffier
National Radio Astronomy Observatory
Charlottesville, VA 22903
20 January 1998

Abstract

We summarize the astronomical capabilities of the current correlator design, with comments on size limitations and expansion possibilities. This is a companion document to *MMA Memorandum 166* (Escoffier 1997).

1 Introduction

A fairly detailed design for the Millimeter Array (MMA) correlator was proposed by Escoffier (1997), in *MMA Memorandum 166*. In this memorandum we set forth more explicitly the astronomical capabilities of that correlator, concentrating on the number of channels (spectral points) produced. We also comment on the extent to which that design might be expanded if necessary (e.g., to accommodate more antennas).

For concreteness we assume throughout that linear polarizations (X and Y) are recorded.

1.1 Nomenclature

Following a suggestion by B. Clark (1997, priv. comm.) we use the following terminology. The numbers given here are for reference only, to make a somewhat confusing discussion more concrete, and are based on the current notional design of the MMA system.

"IF" is a terribly confusing term, and we will avoid it whenever possible. Insofar as it must be used, an *IF* is a wire coming out of a receiver package. In the current design for the MMA, each wire carries an 8 GHz bandwidth. There is provision for switching halves of the IFs (that is, a 4 GHz bandwidth) independently, and each half is sometimes also called an IF. The current plan has a total of 16 GHz being sent down from each antenna to the correlator, for a total of 4 independent (4 GHz) IFs.

A *baseband* (BB) is the signal presented to a sampler. In the MMA case, each BB has a bandwidth of 2 GHz or less, and can be flexibly positioned within an IF band, or switched between IFs. The canonical design has 4 pairs of basebands, with the two basebands of each pair having the same frequency but different (linear) polarizations.

A *channel* is the resolution element of frequency, and is also referred to as a *spectral point*. In the canonical design each channel comes out of the correlator as an eight-byte complex number (four bytes real, four bytes imaginary).

2 Correlator Specifications

The correlator design proposed in MMA Memo 166 provides the following capabilities.

2.1 Maximum Total Bandwidth

The maximum total bandwidth processed is 16 GHz per antenna. ¿From the point-of-view of the correlator this can either be 16 GHz bandwidth in a single polarization (e.g., X), or 8 GHz in each of two polarizations (X & Y).

2.2 Basebands per Antenna

There are 8 samplers per antenna, each capable of digitizing a baseband signal up to 2 GHz wide. For the correlator it doesn't matter whether these are 8 independently-tunable IFs which are each 2 GHz wide, or 2 independently-tunable IFs each 8 GHz wide which are split up just before the samplers; the analog IF processing system will of course impose its own constraints. Following the conventions of §1.1, these 8 "chunks" of 2 GHz are referred to hereafter as *basebands* or BBs.

2.3 Number of Channels (Spectral Points)

The correlator design provides 2048 hardware lags for every baseline (1024 lead and 1024 lag multipliers). These lags can be apportioned amongst different basebands and polarization products (up to a maximum of 16 ways) in powers of 2. Each spectrum computed, corresponding to a single baseband/polarization product, would then have between 2048 and 128 lags, resulting in from 1024 to 64 channels (spectral points). This allows one to trade polarization products or basebands for channels (frequency resolution) in a fairly flexible way. One can also obtain an increase in the number of channels by reducing the input bandwidth into a sampler.

When used at the maximum bandwidth, full polarization, the correlator provides 64 spectral points (channels) across each of the 16 products (2 cross-hand and 2 parallel-hand polarization products for each of 4 BB pairs) for every baseline. There would thus be 256 channels across the entire 8 GHz band, corresponding to a resolution of 31.25 MHz.

Polarization products can be traded for channels (on a BB-by-BB basis). Thus one can double the number of channels (halve the frequency resolution) by producing only the parallel-hand products (XX, YY); or even gain a factor of 4, by producing only one of the parallel-hand products (e.g., XX). Similarly, using fewer BBs increases the number of channels proportionately.

And finally, halving the bandwidth of each BB increases the number of channels proportionately, up to a maximum of a factor 32.[1]

A few examples should clarify all this.

A. *16 GHz, single polarization:* If the IF system could produce 16 GHz of one polarization, say X, in eight 2 GHz basebands, the 2048 correlator lags would be allocated evenly amongst those 8 BBs, giving a total of 1024 channels (128 per BB). The resolution would then be 16 GHz/1024 channels= 2 GHz/128 channels= 15.625 MHz/channel.

B. *8 GHz:* Here one has 4 BB pairs, each covering 2 GHz in dual polarization.

   - *Full pol'n products:* For full polarization information, the 2048 correlator lags are divided among 4 BB pairs times 4 polarization products= 16 different spectra, giving 128 lags (64 spectral points/channels) per 2 GHz spectrum. Each 8 GHz spectrum is covered by $4 \times 64 = 256$ channels, yielding a resolution of 31.25 MHz/channel.
   - *XX only:* Producing only 1 polarization product gains a factor 4 in the number of channels, because the 2048 lags are now divided only among four 2 GHz basebands, yielding 512 lags (256 channels) per 2 GHz spectrum. The full 8 GHz is covered by $4 \times 256 = 1024$ channels, each with a resolution of 7.8125 MHz.

C. *4 GHz, full pol'n products:* In this case one could split the total bandwidth either between 2 BB pairs, each 2 GHz wide; or 4 BB pairs, each 1 GHz wide.

   - *2×2 GHz BB pairs:* With four polarization products the 2048 hardware lags are divided into 8 spectra, giving 256 lags (128 channels) per 2 GHz bandwidth, for a resolution of 15.625 MHz over the full 4 GHz bandwidth.
   - *4×1 GHz BB pairs:* With four polarization products the 2048 hardware lags are divided into 16 spectra, giving 128 lags (64 channels) per 1 GHz bandwidth; but halving the bandwidth of each baseband allows one to use the extra correlators (see above) to obtain twice as many channels. So one winds up with 256 lags (128 channels) per 1 GHz bandwidth, for a resolution of 7.8125 MHz over the full 4 GHz bandwidth.

   In other words, because of the ability to recirculate signals when the samplers are run below their maximum rate of 4 Gsamples/sec, one always gains by splitting a given total bandwidth among the maximum number of basebands. Whether one would always want to do this, and to what extent this might argue for making the correlator smaller by employing more than eight basebands, is discussed in §3.4.

[1]The clock rate of the correlators is 125 MHz (see the discussion in *Memo 166*). Since they cannot keep up with the 4 GHz rate of the samplers, we need 32 parallel correlators. When the sample rate goes to 2 GHz (1 GHz bandwidth), we need only 16 parallel correlators; the extra 16 can then be used to generate twice the number of lags. Obviously this recirculation is ultimately limited by the number of parallel correlators. The full factor 32 that comes from this can be gained by matching the sampler rate to the clock speed (4 GHz/32 = 125 MHz $\rightarrow$ 62.5 MHz bandwidth per baseband).

D. *The highest possible frequency resolution:*

- *Narrowing the bandwidth per BB:* The number of channels goes up by a factor two for every halving of the bandwidth per baseband, up to a factor of 32. So, when using all 4 BB pairs and requesting full (four) correlator products, one can have 64 channels covering 2 GHz, or 128 covering 1 GHz, etc., up to 2048 channels covering 62.5 MHz. In this last case one would have full polarization products for each of $4 \times 2048 = 8192$ channels (spectral points) covering a total of $4 \times 62.5 = 250$ MHz (326 km/s at 230 GHz), giving a resolution of 30.5 kHz (0.04 km/s at 230 GHz).
- *Giving up BB pairs:* The number of channels available for each baseband pair goes up by a factor of two for each halving of the number of baseband pairs. So, one could observe with only two baseband pairs, with a bandwidth of 62.5 MHz per BB pair, and obtain 4096 channels per baseband pair (with full polarization products). The total number of channels would still be 8192, but those would cover only $2 \times 62.5 = 125$ MHz, giving twice the frequency resolution of the last case discussed (i.e., 15.3 kHz).
- *Maximizing the number of channels:* What is the largest number of channels the correlator can produce, and what is the corresponding frequency resolution? Consider observing with a single 62.5 MHz baseband, producing only one correlator product (e.g., XX). This gains another factor 8 over the above case, yielding 32,768 channels covering 62.5 MHz, for a frequency resolution of 1.9 kHz. Of course one obtains only a single polarization product and a single baseband, so the total number of spectra (so to speak) has gone down by a corresponding factor. The velocity resolution for this case (at 230 GHz) is 2.5 m/s, over a total of 82 km/s.
- *Even higher resolution:* The highest resolution one can achieve is set by the minimum bandwidth presented to the samplers, which in turn in is given by the narrowest filters available in the IF system; in the current design, this is 31.25 MHz (Webber 1998, priv. comm.). Since the limit of recirculation has already been reached, cutting the bandwidth further does not produce more channels; all one gains is the corresponding increase in resolution. For this minimum bandwidth one still obtains 32,768 channels (assuming only a single polarization product is desired), and the resolution is 0.95 kHz per channel.

  Ignoring the IF system for a moment, the correlator itself could give virtually any desired spectral resolution, albeit over a limited bandwidth. So if one wants, e.g., 1 Hz resolution, one can only get a maximum of 32,768 channels (for one BB, one polarization product), so the total bandwidth covered would be 32,768 Hz. Similarly, if full polarization products are needed (which also requires using two BBs), one could cover only 8,192 Hz at 1 Hz resolution.

E. *1 GHz total bandwidth, full polarization products:* To maximize the number of channels, one would use 4 BB pairs, each covering 250 MHz. The 2048 hardware lags would be split among 16 spectra, while recirculation would increase the number of channels by a factor 2 GHz/250 MHz= 8, yielding 512 channels (spectral points) per 250 MHz, or a total of 2048 channels over 1 GHz. If on the other hand one required the full 1 GHz within a single baseband pair, for instance to avoid calibration difficulties in splitting a single broad line up into several basebands, recirculation would only give a factor 2 (rather than 8), and one would obtain only 512 channels over 1 GHz.

The astute reader will have noticed that all the tradeoffs discussed so far have involved factors of two (halving the bandwidth or the number of basebands; asking for two rather than four correlator products). This is probably not absolutely necessary, but allowing for other than binary trade-offs would force one to support an even larger number of modes, making the correlator even more complex. So far there has been no compelling scientific argument that the additional flexibility would be worth it.

The correlator modes used to process each baseband or baseband pair can be selected independently. Such (sub-)modes should also be powers of 2 (instead of 7 basebands at one resolution and 1 at another, each sub-mode should use 1/8, 1/4, or 1/2 of the correlator), and to avoid complexity the current design envisions a maximum of four different sub-modes in use at the same time. Some examples of this are:

F. Cover 4 GHz with full polarization (2 BB pairs), and another 4 GHz with parallel-hand polarization products only (2 BB pairs). Each BB pair has 2048/4 = 512 hardware lags available.

   - The full polarization BB pairs split those into four polarization products, yielding 128 lags $\rightarrow$ 64 channels (spectral points) over each 2 GHz bandwidth.
   - The parallel-only BB pairs gain a factor 2, yielding 128 channels over each 2 GHz bandwidth.

   So in this mode the correlator would produce full polarization products for 128×31.25 MHz channels, and parallel-hand products for another 256×15.625 MHz channels.

G. Use 3 BB pairs to cover 6 GHz, producing full polarization products; use one of the remaining BBs to cover 500 MHz, producing a single polarization product (e.g., XX).

   - Each of the three 2 GHz BB pairs has 512 lags available, split amongst four polarization products gives the usual 64 channels over each 2 GHz bandwidth.
   - The 500 MHz BB has 512 lags times a recirculation factor of 2 GHz/500 MHz= 4, for a total of 2048 lags available. Producing only a single polarization product, this gives 1024 channels (spectral points) across 500 MHz.

   In this way the correlator would produce XX, YY, XY, and YX for each of 192×31.25 MHz channels covering 6 GHz, plus e.g. XX alone for another 1024×0.5 MHz channels covering 500 MHz.

H. Suppose one wants to do a survey over 500 MHz producing YY only, while observing one transition in dual polarization over 250 MHz and another over 62.5 MHz; meanwhile one wishes also to zoom in on the central MHz of one of these transitions for an experiment requiring all polarization products. This might be organized as follows:

   - 500 MHz, YY only: use 2 basebands of 250 MHz, each with $256 \times 8 = 2048$ lags (the factor 8 comes from recirculation), producing a YY spectrum with $2 \times 1024 \times 0.24$ MHz channels.
   - 250 MHz, XX & YY: use 1 BB pair with $512 \times 8 = 4096$ lags split between two polarization products, to give two spectra (XX and YY) each having 1024×240 kHz channels (spectral points).
   - 62.5 MHz, XX & YY: use another BB pair with $512 \times 32 = 16,384$ lags, again split between two polarization products. From this pair one obtains XX and YY spectra, each with 4096×15 kHz channels.
   - 1 MHz[2], all four pol'n products: The final BB pair also has $512 \times 32 = 16,384$ lags, since a factor 32 is the limit of the gain for recirculation. Those lags are split between four polarization products, giving 2048 channels (spectral points) over each of the four (XX, YY, XY, YX) 1 MHz spectra, for a spectral resolution of 0.5 kHz.

Table 1 summarizes these examples.

3 Size Limitations and Expansion Possibilities

This section is more tentative, as it relates to possible changes in, rather than simply describing the properties of, the design given in MMA Memo 166. The following should be taken as current thoughts rather than definitive lore.

3.1 Number of Antennas

The number of antennas is hard-wired in from the beginning, and would be very difficult to change after the correlator is built. With the chip design proposed in *Memo 166* the correlator works most naturally in multiples of 8, e.g. 40 and 80 antenna designs are simple scalings of one another, while using the same design for 75 antennas would be inefficient. The design is probably optimal for 64 antennas, with 72 and 80 following in that order. Handling a bandwidth of 16 GHz per antenna for more than 80 antennas would present significant challenges for this design:

- The number of inherently unreliable high-power power supplies becomes even more of a worry with a larger array: the power requirement goes up faster than linear, perhaps as the number of antennas to the 1.75 power.

[2]Note that the current design for the IF system envisions a minimum bandwidth per baseband of 31.25 MHz; see above.

Table 1. Examples of Correlator Modes

	Bandwidth	Corr. Products	$N_{chan} \times \Delta\nu$	Total dv at 230 GHz	Δv at 230 GHz
A	16 GHz ($8 \times$ 2 GHz)	XX	1024×15.6 MHz	20,900 km/s	20 km/s
B	8 GHz ($4 \times$ 2 GHz pairs)	XX,YY,XY,YX	256×31.3 MHz	10,400 km/s	41 km/s
or	8 GHz ($4 \times$ 2 GHz)	XX	1024×7.8 MHz	10,400 km/s	10 km/s
C	4 GHz ($2 \times$ 2 GHz pairs)	XX,YY,XY,YX	256×15.6 MHz	5,200 km/s	20 km/s
or	4 GHz ($4 \times$ 1 GHz pairs)	XX,YY,XY,YX	512×7.8 MHz	5,200 km/s	10 km/s
D1	250 MHz ($4 \times$ 62.5 MHz pairs)	XX,YY,XY,YX	8192×30.5 kHz	330 km/s	40 m/s
D2	125 MHz ($2 \times$ 62.5 MHz pairs)	XX,YY,XY,YX	8192×15.3 kHz	160 km/s	20 m/s
D3	62.5 MHz ($1 \times$ 62.5 MHz)	XX	32768×1.9 kHz	82 km/s	2.5 m/s
D4	32.8 kHz ($1 \times$ 32.8 kHz)	XX	32768×1.0 Hz	43 m/s	0.1 cm/s
E	1 GHz (4×250 MHz pairs)	XX,YY,XY,YX	2048×0.5 MHz	1,300 km/s	0.6 km/s
or	1 GHz ($1 \times$ 1 GHz pair)	XX,YY,XY,YX	512×1.9 MHz	1,300 km/s	2.5 km/s
F	4 GHz ($2 \times$ 2 GHz pairs)	XX,YY,XY,YX	128×31.3 MHz	5,200 km/s	41 km/s
and	4 GHz ($2 \times$ 2 GHz pairs)	XX,YY	256×15.6 MHz	5,200 km/s	20 km/s
G	6 GHz ($3 \times$ 2 GHz pairs)	XX,YY,XY,YX	192×31.3 MHz	7,800 km/s	41 km/s
and	500 MHz (1×500 MHz)	XX	1024×0.5 MHz	650 km/s	0.6 km/s
H	500 MHz (2×250 MHz)	YY	2048×0.2 MHz	650 km/s	0.3 km/s
and	250 MHz (1×250 MHz pair)	XX,YY	1024×0.2 MHz	330 km/s	0.3 km/s
and	62.5 MHz ($1 \times$ 62.5 MHz pair)	XX,YY	4096×15 kHz	82 km/s	20 m/s
and	1 MHz ($1 \times$ 1 MHz pair)	XX,YY,XY,YX	2048×0.5 kHz	1,300 m/s	0.6 m/s

Examples are described further in the text. Note that most examples are split into several lines in this table.

The *Bandwidth* column shows the total bandwidth covered, and how that bandwidth is divided among basebands or baseband pairs.

The *Corr. Products* column shows the polarization products produced by the correlator.

N_{chan} is the number of channels; $\Delta\nu$ is the frequency resolution of each channel.

Total dv at 230 GHz is the total velocity coverage corresponding to the bandwidth (q.v.), assuming a line observed at 230 GHz.

Δv at 230 GHz is the velocity resolution corresponding to $\Delta\nu$ (q.v.), assuming a line observed at 230 GHz.

- The number of signal wires increases with the increase in the number of antennas: a 100-antenna array would require over one hundred thousand 125 MHz interfaces, i.e. 12,800 cables (at 8 interfaces per cable).
- When the number of antennas gets so big that an $N \times N$ array of correlators must be split between two rack bins or even two racks, the number of cables will suddenly jump by a factor of two (because every signal that drives the correlator array must go two places). For this reason the current conceptual design will probably not extend gracefully beyond a 100-antenna array. A 128-antenna array would require a 64×64 matrix of chips to fit on one card, which is not very attractive simply from a physical standpoint, and would probably have a power dissipation that would be difficult to live with. A 100-antenna array would require a 50×50 matrix and is also not pleasant to contemplate.
- As more racks are required, the length of each cable increases, making all cables more difficult to control to ensure proper data capture – the signals from all the antennas must be maintained to within a few nanoseconds, everywhere within the correlator.

While a 100 or 125 antenna correlator may not be impossible, the likelihood of an unreliable system or even an outright failure will increase as the array gets bigger (at a higher than linear rate).

Unfortunately one cannot put off the decision on the array size very long. One would not like to work on the design of a correlator system for more than a few months without knowing the final array size. The number of antennas in an array is fundamental to the design of a correlator. A lot of the very early work on the systems aspect of the correlator design has to do with geometric considerations as to how many what per who (how many antennas per chip, how many chips per card, how many cards per bin, how many bins per rack, how many racks per system). All of these considerations are to some extent interdependent and a good design tries to optimize all of them at the same time to the extent possible. For example, the array size might indicate an advantage of a 3×3 matrix of correlators in the custom correlator chip over a 4×4 matrix. Thus without the final array size, the best custom chip configuration will have to be guessed at. Putting some restrictions on the array size, like 64, 72, 80 or 88 antennas, would help but the exact size would be much better.

3.2 Total Bandwidth

Adding more 2 GHz BBs (samplers) would be relatively simple, increasing the size/complexity of the correlator by the same factor as the increase in the bandwidth. Increasing the bandwidth of the existing BBs would be more difficult to accommodate, and would increase the correlator complexity by more than the factor increase in bandwidth (see also §3.4 below).

3.3 Resolution

Very high resolution is not very difficult to get, as long as the total bandwidth is not excessive. The maximum number of channels (spectral points) in the current design, for bandwidths below 62.5 MHz, using only one baseband, and requesting only a single polarization product (e.g., XX), is 32,768 (see §2.3, esp. example C). Thus one could achieve 1 Hz resolution over 33 kHz, so long as the appropriate low-pass filter is available at the sampler. Although the current design for the IF system has 31.25 MHz as the narrowest available bandwidth for each baseband, it would be fairly easy to modify that design to allow for narrower bandwidths for one of the baseband pairs (B. Clark 1998, priv. comm.).

3.4 Number of Baseband Pairs

As stated above, the number of basebands in the array (of a given bandwidth) has a linear effect of the correlator. Eight baseband pairs would double the size of the correlator described in *Memo 166.*

On the other hand, increasing the number of baseband pairs, while keeping the total bandwidth constant, might actually make for a smaller correlator, as pointed out by L. D'Addario (1998, priv. comm.). For instance, switching from the current four 2 GHz baseband pairs, to eight 1 GHz baseband pairs, would to zeroth order halve the size of the correlator, because one could use half the number of lags to give the same spectral resolution (see §2.3). The most obvious argument against this is the difficulty of keeping consistent calibration between different basebands, particularly for single-dish data; this implies that the maximum bandwidth of a BB should roughly match the width of the broadest lines that would regularly be observed. Providing more narrow BBs would increase the number of BB converters and samplers, thus increasing the cost, but it would decrease the sampler rate, making them easier to design and build. There may be other problems with this scheme – the analogue correlator on the NRAO 12m, based on a similar idea, had many difficulties – but it should probably be thought about more seriously than it has been.

3.5 Number of Channels

Increasing the number of channels beyond the current design would require a much larger custom chip, which is probably not practical at the start of the MMA project. Possibly one could replace the correlator chips in an interim MMA correlator after a few years with chips that have many more lags, potentially yielding 4 to 16 times more channels. The difficulty with allowing such an expansion path is to design the system downstream of the correlator chips to handle the future expansion in data rate and quantity. One way to do this would be to do little or no processing on the lag data in the correlator chassis itself, doing all post-lag computation in completely separate racks that could grow as advances in computer technology occurred.

Such an expansion would also require that the power consumption per lag be reduced in proportion to the increase in the number of lags. If such a new chip required a different operating voltage, there could be further difficulties.

3.6 Dump Times

The design of MMA memo 166 inherently performs fraction of a millisecond integrations. Hence extremely high dump rates have an effect only on relatively small parts of the system.

It has been suggested (Rupen 1997) that even faster dumps may be needed for the autocorrelations, to allow cancelling out the atmosphere without the need for a chopping secondary. The system can easily be designed to provide very fast dumps on the autocorrelation lags. Processing these fast dumps, and providing very fast dumps of more than just the autocorrelation lags (and processing them), requires a lot of study before any definite statement can be made but the problem is mainly the data rate and computation in a relatively small part of the total system.

4 Acknowledgements

Bryan Butler and John Webber provided useful feedback on the draft document. Barry Clark at least partially cleared up MPR's confusion as to the difference between IFs and basebands. Finally, Larry D'Addario pointed out the confusion between various people's readings of *Memo 166* which led to the present document, and also brought up the question of whether the broadest possible basebands are actually desirable.

5 References

D'Addario, L.R. 1989, *MMA Memo No. 55: Millimeter Array Correlator Cost Equation.*

D'Addario, L.R. 1989, *MMA Memo No. 56: Millimeter Array Correlator: Further Design Details.*

Dowd, A. 1991, *MMA Memo No. 66: MMA Correlator: Some Design Considerations.*

Escoffier, R. 1995, *MMA Memo No. 146: An MMA Lag Correlator Design.*

Escoffier, R. 1997, *MMA Memo No. 166: The MMA Correlator.*

Rupen, M.P. 1997, *MMA Memo No. 192: The Astronomical Case for Short Integration Times on the Millimeter Array.*

Thompson, A.R. 1997, *MMA Memo No. 190: A System Design for the MMA.*