

MMA Memo 164: MMA Computing Working Group Report

11 Nov, 96

Steve Scott, Darrel Emerson, Rick Fisher, Mark Holdaway, Jill Knapp, Lee Mundy, Remo Tilanus, Mel Wright

1.0 Overview

The variability of both calibrators and the atmosphere at millimeter and sub-millimeter wavelengths and the potential complexity of the MMA require that the computer control system for the MMA be much more sophisticated than present day instruments. We propose an instrument where the astronomer will interact with the instrument by specifying scientific goals rather than instrumental parameters and the system will intelligently select appropriate parameters. The instrument will typically produce images rather than visibilities.

A system with a state-side control center and wide bandwidth link to the operations support base near the array is optimal for remote access to the instrument by astronomers and engineers. The center will serve as the central hub for interaction with the instrument and the archival database. Almost all observing will be done remotely through the center. The overall computer system architecture will be a distributed hierarchical system, similar to many modern telescope systems (GBT, Gemini, VLT). This design would utilize both embedded microprocessors and more general purpose machines to produce an integrated system. It should be possible to present any aspect of the hardware remotely to allow full fault diagnosis and trouble shooting.

The MMA has the capability to produce data at a rate that would exceed data reduction capabilities. This potential problem has been addressed with data rate limitations, data condensation, and an imaging pipeline. Archive space and data processing resources are treated as finite resources, much like present VLBI tapes. Integration times of around 10 seconds will be the norm, but higher time resolution will be available at the expense of channels and baselines. The 10 second integration time will provide the capability for data editing, self-calibration, and other calibration procedures with reasonable time resolution. The data will flow from the realtime system into a calibration and condensation routine that will apply some calibration and will average the data into frequency bins that match the science requirements. Some of these operations are irreversible. Visibility data will then be written to a database and, in parallel, put into a pipeline to produce images. Dirty images from a subset of the data will be sent to the control center in realtime to aid in data evaluation. The MMA will differ from other previous synthesis imaging instruments in that a concerted effort will be made to automatically produce images that satisfy the scientific requirements for as many of the projects as possible. This will increase the scientific output of the instrument and contribute toward ease of use.

The proposed instrumental specifications that affect computing follow. See Appendix I for details of assumptions and their implications.

1. Minimum interferometric (phase-switched) integration time: 70 msec
2. Minimum single dish (non phase-switched) integration time: 3 msec

3. Maximum data rate: 10 MBytes/sec
4. Average data rate: 1 MByte/sec
5. Setup time: 50 msec

2.0 System Architecture

The basic access model assumed for the instrument involves parallel data streams to operators, observers, engineers and support staff with few restrictions placed on their locations. This model will greatly increase the number of experts available to the array because they can be virtually anywhere, free of the burden of travel. Remote observing, sometimes in a cooperative manner between several astronomers, is viewed as a common way of observing. To support this access model, a wideband data link from somewhere in the U.S. to the operations support base near the array is essential for support of remote observing. Judging from a single quote from AT&T of \$750K/year for T1 bandwidth to Chile, the cost of bandwidth to the array is high enough that data sent over the link will have to be carefully chosen to minimize size and then redistributed from a state-side control center over less expensive networks to satisfy the requirements for parallel access. The link itself must be fast enough to provide images for monitoring observations and for other diagnostics. It appears to be prohibitively expensive to move all the UV data and images over a data link, so CDROMs or tapes will be made on site and sent to the control center for archiving.

The state-side control center serves as the terminus for the data link, the location of the archival database and as the hub for access to the array. The only technical requirement to access the array or the archival database will be an adequate bandwidth link to this center. The links to the center may be the Internet or perhaps an intranet between NRAO and some universities and research centers. Dial-up digital links to the control center with adequate bandwidth will certainly be an option at some point in the operation of the MMA.

Time critical processor intensive chores will be handled by embedded microprocessors closely integrated with the hardware. Higher level tasks of coordination, integration, data reduction and human interface will be handled by workstation style technology. Parallel processing will be used to handle requirements of spatial diversity and raw processing power (see the pipeline requirements). Communication and interfacing in a distributed system can be an important design decision that must be made early.

Maintaining high efficiency from the array will require continuous monitoring of the hardware and evaluation of fault conditions. Diagnosis of problems can be aided by software that understands the interconnections and flow of signals within the instrument, including reconfiguration.

3.0 Observing Modes and Configurability

There will be 6 different modes of observing, or data integration, supported by the software. The phased array mode will only be implemented if the hardware component is included in the instrument. The number of phased arrays may be constrained by the hardware but not by the software.

1. Array synthesis imaging
2. Array On The Fly synthesis mosaicing
3. Single dish On The Fly mapping (continuum and spectral line)
4. Single dish position switched integration

5. Single dish frequency switched (with optional position switching) integration
6. Phased arrays

The ability to configure different portions of the array into sub-arrays that can use different observing modes is useful for several types of experiments and essential for testing. An extreme possibility is using the array as 40 independent single-dish telescopes, so there should be no restrictions on sub-arrays from the standpoint of the software. Any sub-array can be used for any of the 6 different modes of observing. Another configuration concept is a telescope cluster which is defined as a group of telescopes where no correlation products are formed within the cluster. There will be no limits imposed on the number of clusters. Clustering is only valid for the two synthesis observing modes.

4.0 Data Collection

Data collection and processing estimates require that some assumptions be made about progress in future computing capabilities. We have conservatively estimated that technological progress will increase processing speeds and storage densities by the year 2005 by a factor of approximately 20 over resources commonly in use today (assumed to be of 1995 vintage).

The minimum interferometric integration time of 70 msec is driven by solar and flare star observations but constrained by the need to have a full Walsh phase switch cycle and to do the FFT's on every integration. At short integration times, the number of channels available is greatly reduced to stay within the maximum data rate. It is possible to fine tune this number if necessary, as phase switching details are not yet complete.

The minimum single dish integration time of 3 msec is driven by on-the-fly (OTF) continuum mapping. The goal is to Nyquist sample the beam at the targeted mapping rate of a degree per second. At 300 GHz this will be satisfied by samples every 3 msec. Spectral line OTF experience at Kitt Peak shows that there are still gains to be made by sampling faster than 100 msec. Therefore OTF mapping can be done in either spectral line and/or total power allowing a trade off to be made of the sample rate for channels to stay within the maximum data rate of the instrument.

The peak data rate was chosen to satisfy the science requirements specified at Tucson. This rate of 10MB/sec (36GB/hr, 0.9TB/day) will pose no problem at the data recording stage as it can be done with today's technology. The observer must again trade off time resolution for channels and baselines to stay within this limit. This limit is soft and can be increased if warranted. Parallel techniques can be used for the FFTs, calibration and recording if necessary.

The average data rate specification of 1MB/sec was chosen to provide average data sets of a size that have a reasonable chance of being reduced. For example, this rate would be met by experiments with a 10 second integration time and 3200 channels. A 4 hour project collecting data at the average rate would produce a 14GB data set which is equivalent to a 700MB data set using today's technology. Administrative implementation of the average data rate may require that data rate be considered in the proposal evaluation stage. The software will have to enforce the awarded data rate as an integral part of the approved project specification. The limitations imposed by this average rate will encourage wise use of the correlator, conserve archive space, and are consistent with the goal of having the array produce a stack of usable images. Increasing this instrumental specification will have a strong effect on the already difficult problem of the pipeline and data reduction.

A large reduction in data volume can be achieved if the data are channel averaged from the raw correlator resolution to match the proposed science. Several of the calibration steps (some irreversible) must first be performed on the data. This requires treating the software errors as if they were hardware errors - if they seriously compromise the data then the observations must be repeated. The calibrations to be applied are:

1. Atmospheric phase correction
2. Passband
3. Elimination of "end" channels from the spectrum
4. Channel averaging

Work on the atmospheric phase correction is still in progress, but it is possible that corrections will be needed in realtime on timescales of 1 second. These may be applied to the LO or alternately to the baseline visibilities as they are being integrated. It may be necessary to consider recording both the atmospheric corrected data and the uncorrected data and allow the observer to choose from the two data sets on a per baseline basis in the reduction. This would potentially double the data rate specification.

Phase calibration from the phase reference source will be applied to the data before recording and the correction recorded; this is a reversible step. Fast phase calibration cycles can be done faster than the data recording integration time, if necessary, with the constraint that there be an integral number of cycles per recorded integration.

The minimum setup time required is driven by the fast phase calibration scheme, band changes, single dish frequency switching, and wideband frequency synthesis within a band. Both fast calibration and band changes require slewing the telescope and are therefore limited to a fraction of a second. Frequency switching is desired at rates of several Hz so reasonable efficiency requires a setup time of around 50 msec. The wideband synthesis (stepping the LO by many GHz on successive integrations to cover the entire receiver band) is useful for solar work and could benefit from running as fast as possible. A 50 msec setup time gives a 60% duty cycle at the fastest sample rate.

5.0 Near Real-Time Astronomical Feedback

While observing, the observer or array operator requires some near real-time feedback on the observations to verify that the array is operating correctly, that the observational parameters have been chosen wisely, and to ensure that the scientific objectives specified in the proposal (i.e. noise level) are being met. This feedback is essential for modifying the observing schedule by either human intervention or an automated procedure. The near real-time feedback should be fast, should be able to overcome the low signal to noise inherent in short integrations, and should result in compact data products which can be quickly transmitted to distant observers. Near real-time feedback might include:

- time series or spatial distributions of monitor or calibration data (for example, the atmospheric phase or bandpass solutions).
- preliminary astronomical images that are produced quickly with enhanced SNR by averaging selected spectral channels and tapering in the Fourier plane (no deconvolution).
- preliminary astronomical images that are updated with time to show the improved signal to noise obtained from longer integration.
- in mosaicing, total power images of a few selected spectral channels.
- spectra at a single point in a dirty image.

All of these "quick look" examples of near real-time feedback represent computing and data transmission requirements which are small compared to that of the production imaging system designed to produce final images from the data. A single workstation can probably meet the near real-time demands.

6.0 Imaging Pipeline Requirements

It is clear from the volume of data produced by the MMA that an automated imaging pipeline which would provide final or near final images for the majority of the scientific projects is desirable to optimize both computing and astronomer efficiency. As the majority of spectral line projects will probably be limited by noise rather than systematic errors, many projects are well suited to the automated pipeline described here. There is an additional need for off-line computing to produce images for projects that are too demanding for the pipeline.

The magnitude of the pipeline requirements has been estimated by benchmarking the pipeline tasks on simulated single channel MMA data. The benchmarking machine was a mid-level workstation, a 60 MHz Sparc 20. To scale the benchmarks to the full MMA pipeline, processor power has been scaled using the 20 fold increase previously mentioned and the number of channels and integration times chosen based on data rate limitations and on the brightness sensitivity of the array. The results are given in Appendix II and show that the pipeline will require approximately 40 workstations or their multiprocessor equivalents. The total pipeline cost of about \$800k is obtained by scaling the \$20k cost of the benchmarking machine. The overall computing demand is very high largely because of multiple spectral channels. Imaging algorithms operate on a single channel at a time, and the single channel imaging problems are quite modest and well suited to individual processors. Hence, the pipeline requires a data server that makes the appropriate spectral channels' data available to each processor.

The pipeline specifications are quite preliminary and are likely to change. The performance-to-cost ratio may improve more than the conservative assumption of 20 before the MMA is operational, and mosaicing and deconvolution algorithms may become more efficient. At this time, the computing requirements placed on the pipeline do not seem to be too extreme, and the entire pipeline will likely cost under one million dollars.

7.0 Observational Scheduling

There are three different methods proposed for scheduling observing time that span the range from those conventionally used today to some new methods that are particularly well suited to the MMA. Atmospheric conditions will almost certainly play an important role in efficiently scheduling the array and is the area explored by the proposed new methods. Phase correction techniques may make a major impact on the ability of the instrument to adapt to the atmosphere but these techniques are still under development.

The most traditional method of single dish mm observing, **interactive**, will give control of the array to an observer for a fixed time interval. The observer uses feedback from the observations to determine the next observation. The second scheduling method is **fixed queue** observing which simply steps through a set of pre-planned observations. These may be fixed in sequence or fixed in absolute or sidereal time. In the context of the MMA, this mode would be used when observing must be done at specified times (solar system events, VLBI, etc.) or when a specific set of observations needs to be done independent of

other factors. It is also the simplest form of automated observing and will therefore play a role as the array is brought online. The scheduling method that can take maximal advantage of changing weather conditions is **flexible** scheduling. It uses the current atmospheric monitors (seeing and opacity) and instrument status information to select from the queue the project which makes the best use of telescope time under these conditions.

Remote observing (either from a state-side control center or local 2500m operations center) is a requirement for any of the scheduling modes as it will reduce travel and is essential for flexible scheduling. The transition from the flex mode into interactive will require that interactive projects be elevated into a "standby" state while awaiting appropriate atmospheric conditions and that the observer be notified to be prepared to observe.

There are innovative yet simple ways to extend the flexible queue that will provide all of the features of the traditional mode as well as some interesting advantages. On most occasions the interactive user simply requires one or more 'checkpoints' during the observations to enable a decision on the further direction of the project. Such requirements for feedback can be incorporated into the flexible queue so that it will stop a project at a predetermined point, continue with the rest of the queue, and make the project ineligible for rescheduling until there is feedback from the observer. This allows indefinite time (minutes, hours, or even days) for head scratching and reflection without pressure. A simple yet powerful example of the use of checkpoints occurs while trying to integrate down to an adequate S/N. After spending only half of the projected integration time, a checkpoint can be scheduled, and spectral or spatial smoothing will allow a decent evaluation of the data. Cases in which either the science target has already been met or which will never achieve their science goal within the allotted time should be obvious. This thus argues for a system of 'mid-way checkpoints' for a whole class of observations, with the potential of greatly enhancing the efficiency of the telescope. The exact mixture of classical interactive and the enhanced flex queue will depend on the success of the implementations and will probably vary over time.

A number of the points raised above may seem fairly unconventional, but most of these techniques (flexible scheduling, remote observing, checkpoints) are in use, being planned or in the process of being implemented at telescopes such as the JCMT, BIMA, OVRO, and the SMA. By the time the MMA begins operation, remote observing using flexible queues should be common for mm telescopes.

8.0 Smart Observing Algorithms and Tools

The variability of the atmosphere and of the calibrators requires the use of sophisticated "smart" algorithms to make observing easier and more efficient. Some of these algorithms will be capable of figuring out optimal observing parameters on the fly and adjusting them accordingly. The goal is to let the experiment be specified in terms of the science instead of instrumental parameters. In all cases it will be possible to override smart parameter settings and run in manual mode. Some suggested concepts, algorithms and tools are:

1. Proposal preparation tool

The user should be able to specify source(s), catalogs, frequencies, spectral resolution, pointing centers; i.e. the astronomically desired parameters. The tool will generate an expert observing file that can then be edited by hand if desired. This file can subsequently be checked and optionally used to generate model beams and images with expected noise levels for several classes of weather conditions. The checking operation will include validation of syntax, evaluation of integration time

on sources and calibrations, and scanning for potential problems such as long slew times or erroneous spectrometer configurations. This output can be used to evaluate weather dependent factors filled in on the proposal form.

2. Science goal comparison operators

Many traditional interactive observations are implicitly doing tests such as "if SNR in channel 20 is greater than 3 then" or "if integrated flux over channels 6 to 120 is less than 50mj then" to determine the flow of the observations. To support this decision making in an automated way, it must be possible to extract the following from the pipeline data so that it can be used by the scheduling system:

- Flux - integrated, average, max or min
- Signal to noise - average, max or min
- Noise in map

The above quantities are specified over sets of boxes in the image cube as selected by the observer.

3. Calibration interval adjuster

Uses the atmospheric seeing monitor to select how often phase calibration needs to be done.

4. Calibrator selector

At the start of a new source, measures flux of adjacent potential calibrators using the array as 40 separate single dishes at lower frequencies or interferometrically at high frequencies where SNR is a limiting factor.

5. Pointing peak-up routine

Integrates coherently until adequate SNR is achieved and then moves on to the next measurement point. Does multiple points to statistically measure the accuracy of the offsets and automatically quits when the offsets are measured to sufficient accuracy.

9.0 Archival Database

The data will be shipped on tape or CDROM from the array to the control center where the database will be located and then become accessible from there. Items to be stored in an archival database include instrumental parameters, weather, and pointing as well as the headers for the data. The actual visibilities and images may be stored in another file system but pointers to it would be part of the headers stored in the database. The database will be used to provide general access to the data when the proprietary period has expired. To promote this data mining, it must support general queries via a structured interface plus some higher level constructs for ease of use.

10.0 Design and Implementation Issues

The complexity of the MMA and heavy reliance on flexible remote operation create a very challenging software design problem for control and monitoring. A long and iterative design analysis period must be anticipated and begun as early as possible. Ideally, the MMA project will assemble a design team with full-time participation of people from the several most recent telescope control software projects. Experience is the most transferable software product, and the best experience is that which has used modern software design and management techniques on closely related projects. The result will draw heavily on industry standards and incorporate the best features of existing systems and packages such as CORBA, EPICS, and glish. However, each instrument and its requirements are different and new enough to require a fresh design. This must be recognized in the project cost.

One lesson learned from other projects is the importance of a thorough analysis of the data flow from the outset of the project. This will ensure a smooth interface between major components of the system. The

VLT is a good example of the use of this technique. The flow should begin at proposal submission and include proposal review, observing schedules, actual observing logs, data products, and archival data. Implementation should be done with a single database to maintain a consistent view of the data as it moves through different stages in its life cycle.

Appendix I

Data Rate Assumptions and Implications

To evaluate the data rate assumptions a clear definition of a "channel" and the correlator capabilities is needed. A "channel" here is one spectral channel (real & imag from a single sideband - 2 numbers total) on one baseline. There is one channel per sideband from every lag. The per baseline correlator is 1K lags at full clock and 32K lags at the slowest clock. These lags may be split over multiple windows but this will only cause a modest increase in header, not data.

At the Tucson meeting in Oct, '96 the science groups presented their requirements which are now summarized.

- Solar flares and flare stars require temporal resolution of 50-100 msec but with a reduced number of channels.
- Single dish spectral line on-the-fly mapping requires sample rates of 100 Hz.
- On-the-fly interferometric mosaicing at sample rates of 1 sec but restricted to a few thousand channels.
- Fast switching (~10 sec cycle) between source and phase calibrator will be required if a realtime phase correction for atmospheric water vapor is not possible. The system should support integration times of 10 seconds for the full correlator output (the phase calibrator will be wideband only). If shorter cycles are required then multiple cycles can be integrated together to form one recorded integration.

Assuming the data are stored using scaled integers so that each number takes 2 bytes, all of the above requirements can be met with a maximum data rate of 10MBytes/second. If all 780 baselines (or 40 telescopes in the single dish case) are used then the following table holds:

	Min int time	Channels	Data Rate
Solar flares & Flare stars	70msec	224	10MB/sec
Single dish OTF mapping	3msec	420	10MB/sec
OTF mosaicing	1sec	3200	10MB/sec
Fast phase calibration	10sec	32000	10MB/sec

Appendix II

Imaging Pipeline Requirements

The basis for the assessment of the imaging pipeline computing requirements is a series of benchmarks of the reduction time for simulated single channel MMA data. The benchmark platform was a 60 MHz Sparc 20 purchased in 1995. Three problems were considered:

- one hour single pointing observations including imaging and deconvolution
- compact array linear mosaicing for moderate dynamic range observations
- compact array nonlinear mosaicing for high dynamic range observations

Phase and bandpass calibration require much less computing than imaging, and will not be considered further. The extrapolations from the benchmarks to typical MMA projects involve an estimation of the number of channels and integration times for different configurations. The 3200 channels used for the 2 compact MMA arrays are consistent with the average data rate with 10 second integrations. For the larger arrays, the brightness sensitivity of the MMA decreases and wider bandwidths are required in order to detect typical astronomical lines. Hence, we assume the B array will typically image 1600 or fewer channels, and the A array will image 400 or fewer channels. This will partially offset the increased computing requirements of the larger arrays caused by the larger number of pixels across the primary beam.

More detailed justification for these estimates and for the pipeline will be given in an upcoming MMA memo "Computing Requirements of an MMA Imaging Pipeline" (Holdaway, 1996).

A. Single Pointing Observations

The single pointing observations consisted of one hour spectral line observations with the number of channels given above. The timing results for imaging without deconvolution are summarized in Table 1. At 10 s integration times, gridding dominates the compact array imaging. Averaging the data can result in a significant reduction in processing time as is shown in the lower half of the table with 240 s integration times. In the extended arrays, the FFT gains relative importance and is clearly dominant in the A array. The UV cell crossing time will restrict integration times to about 120 s in B array and 30 s in A array but these should not prove to be practical limitations because the gridding is less important than in the compact arrays. Deconvolution requirements are more difficult to estimate due to the strong dependencies on noise level, array configuration, source structure, the fraction of the field which is filled with emission, and the deconvolution algorithm. However, the interplay of these factors will be addressed in detail in the future MMA memo. For the time being, it is assumed that deconvolution will take about 10 times longer than generating the dirty images. With the factor of 10 for deconvolution, the future MMA pipeline of 40 processors can keep up with the data as long as the dirty imaging cpu time c. 2005 for all channels is less than about 4 hours. For D and C array, this will require integrations of about 30 s, B array will require about 20 s, while A array can use 10 s.

Table 1. Imaging 1 hour of data without deconvolution

Array	FFT Size	CPU Time c.1995 for 1 Channel	Number of Channels (N)	CPU Time c.2005 for N channels
10 s int time: dominated by reweighting and gridding				
D (80m)	64x64	226 s	3200	10.0 h
C (200m)	256x256	237 s	3200	10.5 h
B (800m)	512x512	257 s	1600	5.7 h
A (3000m)	2048x2048	558 s	400	3.0 h
240 s int time: dominated by FFT				
D (80m)	64x64	10.3 s	3200	0.46 h
C (200m)	256x256	16.6 s	3200	0.74 h
B (800m)	512x512	33.4 s	1600	0.75 h
A (3000m)	2048x2048	389 s	400	2.2 h

B. Linear Mosaicing Observations

The brightness sensitivity of the MMA in its most compact array permits integrations as short as a few seconds per mosaic pointing. The linear mosaic algorithm, which assumes that each pointing has the same point spread function, can efficiently mosaic MMA data with dynamic range up to about 500:1. With 2 s integrations per pointing and 3200 channels, the pipeline can keep pace with the MMA using 32 of its future processors.

C. Nonlinear Mosaicing Observations

Nonlinear mosaicing will usually only be required for very high dynamic range (greater than about 500:1) observations. At 60 s integration times and 400 channels per pointing (or 1 s and 6 channels per pointing), the pipeline could image as fast as the MMA produces data. This is probably sufficient for the very limited use of the nonlinear mosaic algorithm.