On-The-Fly Mosaicing

M.A. Holdaway and Scott M. Foster
National Radio Astronomy Observatory
Socorro, NM 87801

October 19, 1994

Abstract

In mosaic observations consisting of many (250-10,000) interferometric pointings, non-adjacent pointings may have significantly different (u, v) coverages and hence different synthesized and restoring beams. The advantages to having similar synthesized beams for each pointing include: the image is easier to interpret, deconvolution errors will have similar character across the image, a single clean beam can be used for the entire image, residuals are in the same units across the mosaic, and the inexpensive linear mosaic algorithm can be used. In order to make the beams for all pointings similar to the 1% level, each pointing in the mosaic must be observed every ~12 minutes. For mosaics in which beam shape differences of 5% can be tolerated, each snapshot must be observed every 45-60 minutes. Mosaics smaller than ~250 pointings can achieve very similar (u, v) coverage and beams similar to 1% with the traditional stop-and-go mosaic observing strategy with on-source integrations of about 3 seconds per snapshot with 1 second setup time. Mosaics smaller than ~900 pointings can achieve 5% beam differences with stop-and-go mosaicing. Larger mosaics may require continuous slewing across the source, reading out the data on time scales which are short compared to the beam crossing time. We call this technique on-the-fly (OTF) mosaicing. The Ekers and Rots short spacing scheme places restrictions on the maximum averaging time in terms of the beam-crossing time. The short integration times required by OTF mosaicing place extreme demands on the correlator and on-line system. The mosaicing algorithms currently used will work with OTF mosaicing, but the assumed rotationally-symmetric primary beam must be replaced with the effective primary beam, which is the true primary beam convolved with a 1-D boxcar. The very high data rates may require rethinking the typical data processing pathway.

1 Mosaicing and (u, v) Coverage

1.1 Introduction

Given the MMA's small primary beam (32″ at 230 GHz), mosaics with very many pointings will be quite common. For example, a 10,000 pointing mosaic at 230 GHz will cover a 27′ square patch on the sky, producing an image 1600 1″ pixels across. The processing for such a mosaic is possible even with today's computing environment, and even larger mosaics will be commonplace for the MMA.

We have already laid out much of our view of how interferometric mosaicing observations and imaging should proceed (Cornwell, Holdaway, and Uson, 1993). The interferometric array should be designed to yield essentially complete effective (u, v) coverage. The interferometer will observe many sky positions separated by $\lambda/2D$ on the sky, covering the object of interest plus a guard band of blank sky (or at least a guard band of sky which is not of primary interest). As in single dish imaging, systematic errors such as pointing errors or changing atmospheric opacity can be better dealt with if several passes over the region of interest are made.

We define the hit time t_{hit} to be the time between snapshots of a given pointing on the sky. For a regular, repeating pattern of pointings, the maximum time between two pointings on the sky will be $t_{hit}/2$. Mosaic images made with a finite amount of time spent on each pointing will have different (u, v) coverages for each pointing, with the magnitude of the differences in (u, v) coverage increasing with the hit time. The hit time will depend on the number of pointings in the mosaic, the slew and settle down time, and the fraction of time spent integrating on source.

Even for zero hit time (ie, a magically fast array), differing projection effects at the different sky positions will lead to differing (u, v) coverage. For millimeter wavelength mosaics, this is a small effect which dominates only when $t_{hit} \lesssim \theta_p/15/\sqrt{N_p}$ seconds, where θ_p is the angular distance between pointings in arcseconds and N_p is the total number of pointings in the mosaic. Hence, the difference in the (u, v) coverage of different pointings in a mosaic is expected to be dominated by the change in the azimuth and elevation which occurs over t_{hit} rather than by the change in azimuth and elevation due to different sky coordinates.

1.2 Why Do We Want Similar (u, v) Coverage Across the Mosaic?

The differing (u, v) coverage presents at least three kinds of difficulties:

- The (u, v) sampling envelope will change across pointings, resulting in a restoring beam which changes over the mosaic image. Differing beam shapes will hinder the interpretation of the image and will put the residuals across the image in different units.

- Different details in the (u, v) coverage lead to slightly different deconvolution errors across the image. This may hinder efforts to quantify or even correct for the deconvolution errors.

- The linear mosaic algorithm, which is about 20 times faster than the nonlinear mosaic algorithm, makes the approximation that the (u, v) coverage of each pointing is identical. A larger difference in the (u, v) coverage will result in larger errors in the linear algorithm so that it can be applied to only lower SNR data. If we can make the differences in (u, v) coverages among pointings small enough, we may be able to use the linear mosaic algorithm to image essentially all of our mosaic observations.

Hence, it would be a good thing to find an efficient means of creating similar (u, v) coverage for the different pointings in a mosaic.

1.3 How Much Does The Beam Shape Change Over a Mosaic?

We have performed numerical simulations of mosaics using the straw man MMA compact array with elongation 1.2 for declinations 0° through 50°; and a compact array with elongation 1.8 (to reduce shadowing) for declinations −40° through 0° and 50° through 80°. For an hour angle range of ±1.5 hours and declination intervals of 10°, the (u, v) coverage was calculated for one sky pointing with snapshots at times $t_o + it_{hit}$ and another sky pointing with snapshots at times $t_o + (i + 0.5)t_{hit}$, for integral i, with t_{hit} ranging between 5.625 minutes (32 snapshots per pointing) and 180 minutes (one snapshot per pointing) in factors of two.

Rather than study the detailed differences in the (u, v) coverage, we focused on the differences in the fit beam shapes, which reflect differences in the envelopes of the (u, v) coverage. This is a reasonable approach since many of the reasons for obtaining similar (u, v) coverages pertain primarily to the beam shapes and sizes. We made naturally weighted synthesized beams for both pointings. We then performed a nonlinear Gaussian fit to the inner lobe of the synthesized beams. In order to get a a very accurate Gaussian fit, our pixel size oversampled the beam by a factor of about 10. We evaluated the difference in the beams in two ways:

- the normalized difference in the beam areas, given by

  $$(B_{maj1}B_{min1} - B_{maj2}B_{min2})/B_{maj1}B_{min1},$$

  is appropriate for gauging effects such as the changing flux scale across images, and is plotted as a function of t_{hit} for several declinations in Figure 1.

- the normalized integrated absolute difference between the beams (including beam orientation), given by

  $$\int(|beam1 - beam2|)/\int(beam1)$$

  gauges the overall difference between the beams, and is plotted as a function of t_{hit} for several declinations in Figure 2.

-

We do not present detailed imaging arguments for an acceptable level of beam difference across the mosaic. We expect that the accuracy of quantities derived from mosaic images will degrade gracefully with increasing beam differences. We make the somewhat arbitrary suggestion that the beam-shape differences across the mosaic image be limited to 1% for high quality mosaics and 5% for moderate quality mosaics.

1.4 How Can We Get More Similar (u, v) Coverages?

We see several means by which the (u, v) coverage can be made more similar:

1. The array configuration can be optimized to minimize the differences in (u, v) coverage. Ge (1992) demonstrated that an array stretched in the north-south direction leads to very different (u, v) envelopes for observations on either side of the zenith. Optimizing the array can only take us so far, and this is only one optimization criteria among many.

2. (u, v) points outside a common envelope can be flagged, resulting in similar restoring beams but different detailed (u, v) coverage. This approach will sacrifice both resolution and sensitivity and is probably unacceptable.

3. Large mosaics can be scheduled in blocks of the same LST over several days, thereby reducing the hit time. Very large mosaics probably will need to be scheduled over several days over a small hour angle range ($\pm$2 hours).

4. Observing strategy could be modified to minimize the (u, v) differences: for example, including snapshots symmetrically spaced in hour angle and staying near the zenith will improve the problem. The effectiveness of symmetrical hour angle coverage is limited unless short integration times can be used.

5. The mosaicing process could be sped up by decreasing the integration time per pointing, scanning through all pointings many times. Using numerical simulations on the MMA antenna design, Cheng (1993) estimates the settle down time (settle down to $1''$ pointing) to be roughly 1 second, so it would not be palatable to spend less than 3 seconds on each pointing. From the numerical simulations, a hit time of 12 minutes results in $\sim$1% beam differences. Hence, a 240 pointing mosaic ($4'$x$4'$ at 230 GHz) is the largest mosaic which would maintain $\sim$1% beam differences over the image. A hit time of 45 minutes results in $\sim$5% beam differences, which limits the mosaic size to 900 pointings ($8'$x $8'$ at 230 GHz). Many extragalactic objects can be imaged with mosaics of this size or smaller, but many targets in the Galaxy will require much larger mosaics. Such an observing strategy would require the correlator to dump all visibilities about once per second.

6. Mosaicing could be accelerated further by not stopping for each pointing, but continuously scanning over the region. This technique, which we call on-the-fly (OTF) mosaicing, could handle very large mosaics with several thousand pointings, but introduces several issues which we will deal with in the following sections.

Our working hypothesis is that stop-and-go mosaicing will work well for many small size mosaics, large mosaics will require either OTF mosaicing or stop-and-go mosaics covering several days, and very large mosaics will require OTF mosaicing covering several days.

2 On-The-Fly Mosaicing

OTF mosaicing is sufficiently different from the proposed observing strategies that it will significantly impact the design of the MMA in a number of areas. The advantages and disadvantages of this observing approach include:

- many complete passes over the region of interest can be made, as in single dish imaging, reducing the effects of systematic errors such as pointing errors and changing opacity.

- there is no loss in sensitivity due to either flagging outlying (u, v) data or due to antenna move time.

- since the (u, v) coverage for every point will be very similar, the "pointing axis" of a multi-dimensional data "cube" will be virtually perpendicular to the u, v, and ν axes, opening the possibility of new kinds of data editing, display, and imaging algorithms for mosaic data.

- the same old mosaicing algorithms will work for OTF mosaic data, but the Ekers and Rots scheme limits the integration time to well under a half beam crossing time.

- the control and monitor system will be pushed very hard to make 40 antennas point together to $1''$ while slewing.

- stop and go pointing every few seconds will probably lead to more antenna maintenance headaches than will continuous scanning.

- shorter integrations will result in a much higher data rate, impacting the correlator design and the data archive mechanism (see Tables 1 and 2).

- imaging computers will need to have more disk space to deal with the extra data, but less cpu time will be required to image the data. Calibration and gridding will need to be done "on-the-fly".

2.1 Imaging Details of On-The-Fly Mosaicing

Over an integration time, the measured visibilities will be

$$V_{ij} = \int_{t_o}^{t_o+\Delta t} dt \int dx A(\mathbf{x} - \mathbf{x}_p(t)) I(\mathbf{x}) e^{i2\pi \mathbf{ux}} / \Delta t \tag{1}$$

Assuming the (u, v) values do not change over the integration, we have generic mosaicing with an effective primary beam

$$A_{eff}(\mathbf{x}) = \int_{t_o}^{t_o+\Delta t} dt A(\mathbf{x} - \mathbf{x}_p(t)) / \Delta t \tag{2}$$

If $\mathbf{x}_p(t)$ changes linearly with time, the effective primary beam will be

$$A_{eff}(\mathbf{x}) = A(\mathbf{x}) * \Pi(x/(2v_{slew}\Delta t)), \tag{3}$$

where Π is a boxcar function: $\Pi(y) = 1$ for $|y| < 1/2$, $\Pi(y) = 0$ for $|y| > 1/2$. In the Fourier plane, each visibility then samples the nominal (u, v) point convolved with the Fourier transform of the effective primary beam, which is

$$a_{eff}(\mathbf{u}) = a(\mathbf{u}) \cdot sinc(\mathbf{u} \cdot (2\mathbf{v}_{slew}\Delta t)), \tag{4}$$

where $a(\mathbf{u})$ is the Fourier transform of the primary beam, or the autocorrelation of the antenna illumination pattern. The MMA is implicitly relying upon the scheme of Ekers and Rots (1979)

N_p	Linear Size, ′	Dump time, s	Data rate, MBytes/s
250	4.3	0.96	3.5
1000	8.6	0.24	14
4000	17.1	0.06	56
16000	34.3	0.015	230

Table 1: How fast do we need to mosaic to get 1% beam differences? Square mosaic sizes at 230 GHz, correlator dump times, and data rates for 2 stokes, 512 spectral channel (compressed) OTF mosaic data as a function of the number of pointings N_p.

N_p	Linear Size, ′	Dump time, s	Data rate, MBytes/s
250	4.3	3.6	0.93
1000	8.6	0.9	3.7
4000	17.1	0.225	15
16000	34.3	0.056	60

Table 2: How fast do we need to mosaic to get 5% beam differences? Square mosaic sizes at 230 GHz, correlator dump times, and data rates for 2 stokes, 512 spectral channel (compressed) OTF mosaic data as a function of the number of pointings N_p.

to measure baselines of up to one dish diameter shortwards of the shortest projected physical spacings. The sinc function degrades the sensitivity to the (u, v) information at the edge of the antenna illumination autocorrelation far away from the nominal (u, v) sampling point. In order to minimize this loss in sensitivity, the first null of the sinc function should be well outside the null in $a(\mathbf{u})$. In the image plane, the slew distance should be much less than the primary beam size (we will say 1/3 of the beam size).

The requirement that the beams are constant to within 1% across the mosaic demands a hit time of about 12 minutes. Together with the 1/3 beam sampling requirement, we can specify the required correlator dump rates and data rates as a function of mosaic size, which are presented in Table 1. Note that for 230 GHz mosaics larger than 16000 pointings, the differences in (u, v) coverage due to different sky positions across the mosaic become comparable to the differences due to 12 minute hit times. Table 2 reports the correlator dump rates and data rates for the less stringent requirement that the beams be similar to 5%.

It is probable that we will not be able to read out the averaged visibilities at the same mean sky positions (to the specified pointing precision of $1''$) on each pass through the mosaic, but we will probably the mean pointing to suitable precision. The mosaicing will be much faster if we can treat the data from the multitude of slightly different mean pointing positions as a single pointing. Since the pointings are fully sampled in sky position space, we should be able to interpolate the data between adjacent pointings to obtain the data that would have been measured on an idealized grid. This is analogous to the problem of single dish on-the-fly mapping, except that the interpolated data cubes require further processing to form an image.

Since the data rates are incredibly high in OTF mosaicing, we should rethink the data reduction path. The original data from the correlator need to be archived, but dealing directly with the original data should probably be avoided as much as possible. The original, unedited data from the correlator are written to the archival medium and simultaneously passed to a reduction computer. The reduction computer will perform on-line editing, calibration, and gridding to the Fourier hypercube (axes u, v, ν, sky pointing). On-line self-calibration is even possible after enough data has been taken to form an image for a few bright channels. Mosaicing algorithms would operate on the calibrated, gridded data. "On-the-fly" reweighting in the Fourier plane will be possible. Such a scheme will make the data processing manageable since the Fourier hypercube can be filled in real time, and mosaicing operations downstream of the hypercube are similar for OTF mosaicing and stop-and-go mosaicing. Since the data are calibrated and gridded online, it is easy to make images while observing in order to further guide the observations. To illustrate the savings in disk space, consider a 1000 pointing mosaic with 1% beam differences taken over 4 hours. The 14 MBytes/s of raw data from the correlator amount to 200 GBytes, independent of configuration. The approximate size of a Fourier hypercube using compressed data would be

$$N_p N_{channels} N_{grid}^2/2 \cdot 4 Bytes,$$

where N_{grid} is the number of pixels across the Fourier grid, which is about 64 for the MMA D array and 256 for the MMA C array. The gridded data from a 1000 pointing, 512 channel mosaic in the D array would be 8 GBytes, independent of observing time. The main drawback to this data reduction path is that it does not accommodate any sort of recalibration since all time information is thrown away. One must go back to the original data on the archive medium to recalibrate.

3 References

1. Cornwell, Holdaway, and Uson, 1993, A&A **271**, 697.

2. Ge, Jing Ping, 1992, MMA Memo 80, "Further Simulations of Possible MMA Configurations".

3. Cheng, Jingquan, 1993, MMA Memo 101, "Slanted-Axis Antenna Design II".

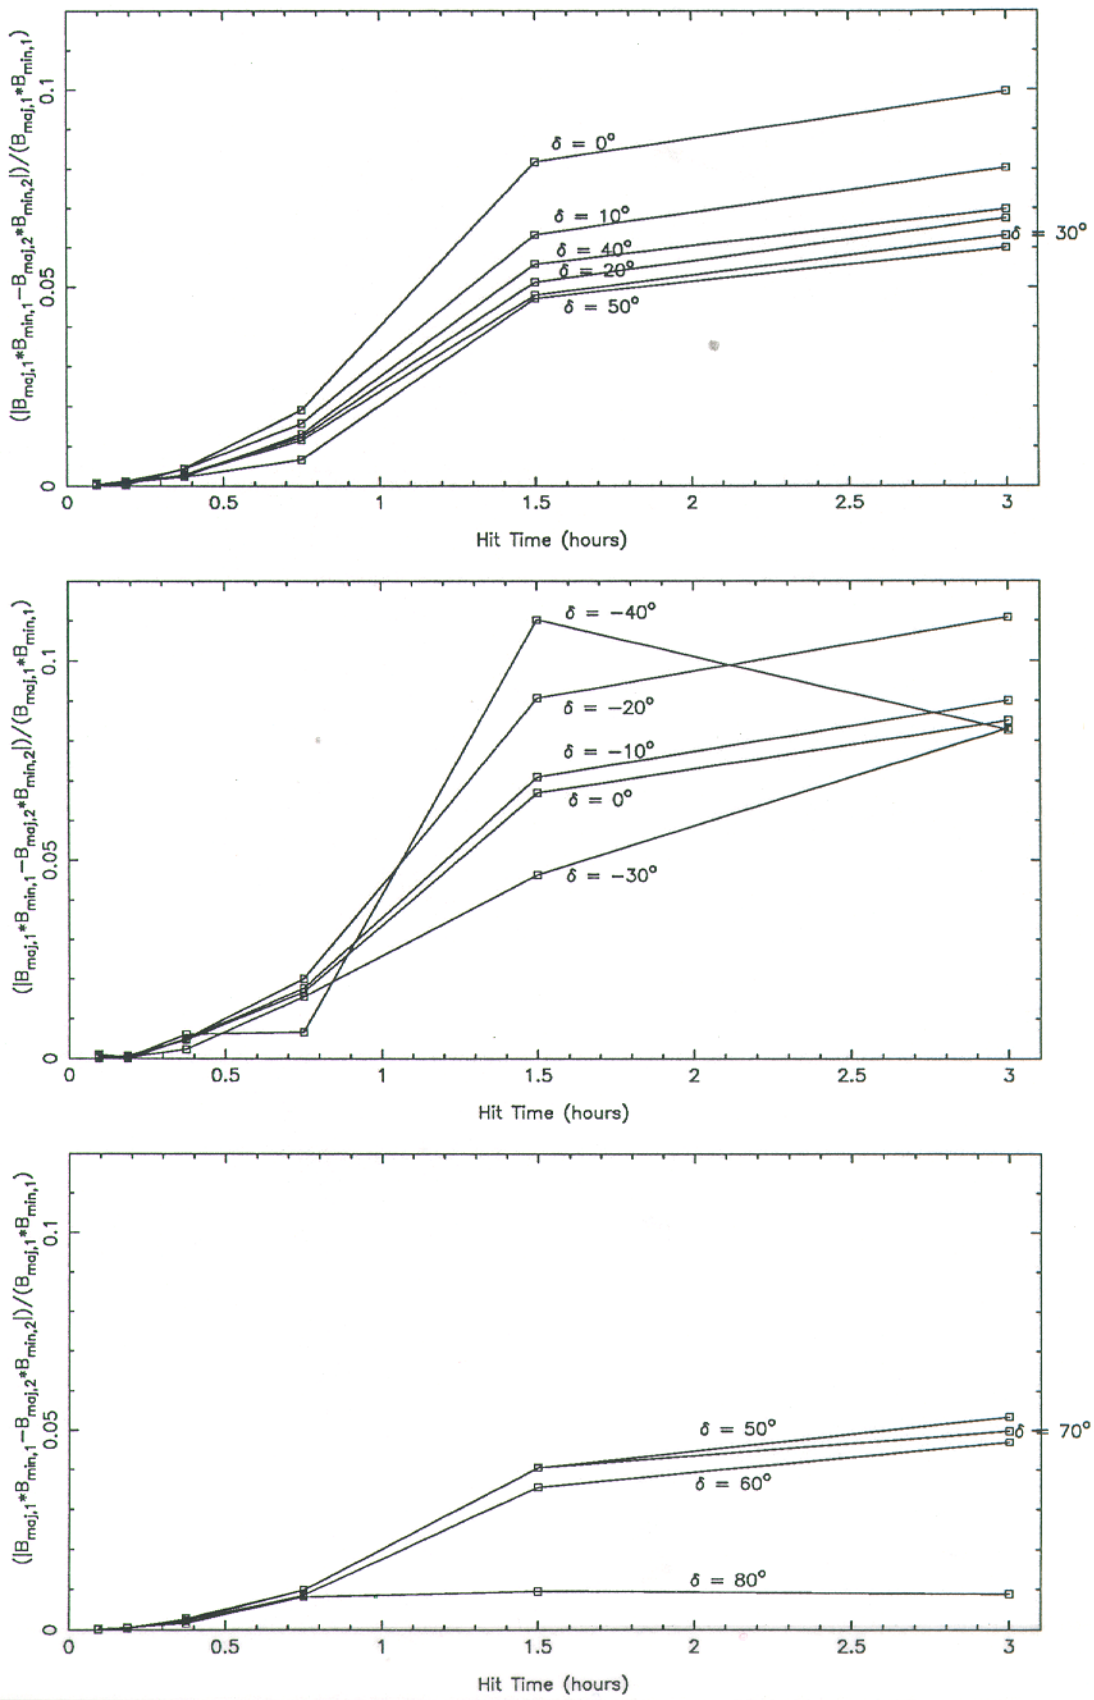

Figure 1: Beam volume differences as a function of declination and snapshot hit time.

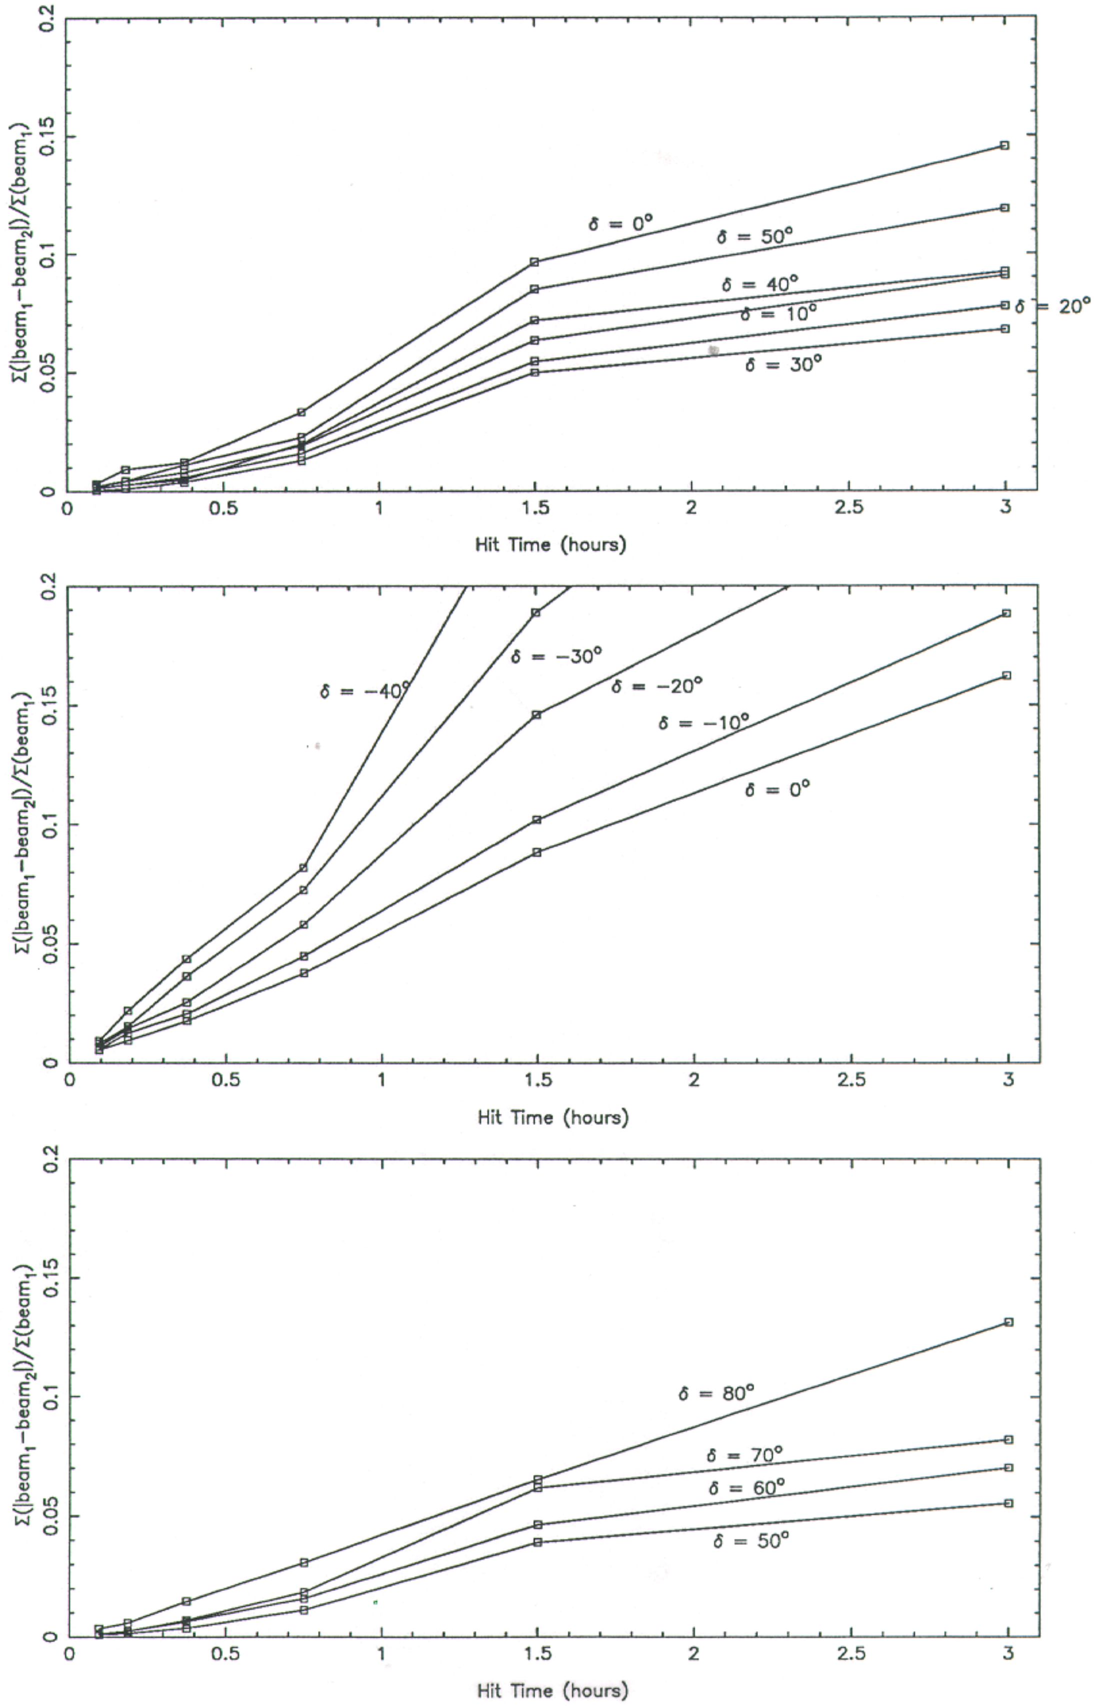

Figure 2: Normalized integrated absolute beam differences as a function of declination and snapshot hit time.