Multi-facet CLEANing in Obit

W. D. Cotton, October 2, 2009

***Abstract*—This memo describes an improved multi-facet technique for imaging and deconvolving large data-sets. The use of "three dimensional" imaging has long been a performance bottleneck due to the need for sequential CLEANs of the facets in the deconvolution phase. Reprojecting the facets during griding onto a common plane with a common image grid allows "two dimensional" imaging and a multi–facet CLEAN. A three tiered CLEAN cycle for the autoWindow technique is described and then applied to a multi–facet CLEAN using two dimensional imaging. This method has been implemented in Obit and the results of several comparisons are given. On a large test in which the data–set was too large to fit in disk cache, multi–facet CLEAN reduced the run time from 29 to 7.6 hours. The multi-facet CLEAN has an obvious extension for further performance enhancements using cluster computers. AutoCentering, self–calibration, and peeling are all supported using this new technique.**

***Index Terms*—interferometry**

I. INTRODUCTION

ONE of the standard solutions to the "W" problem, that images are flat but the sky isn't (see [1] for details), is to use multiple facets to tile the region of the sky to be imaged. The different facets are then deconvolved from the interferometer response using CLEAN. The good dynamic range faceted CLEAN developed by Eric Greisen in AIPS and adopted in Obit does the CLEANing using only a single facet at a time. Using the "3D" faceting, each facet is tangent to the sky at its center resulting in each facet being in a different plane from all others. Since the various facets will overlap, a given source may appear in multiple facets. Furthermore, the sidelobes of a strong source may appear in adjacent facets. This limits high dynamic range imaging to the ability to subtract the response of a source in one facet from another by a subtraction of its response from the visibility data and then re-imaging. Thus, CLEAN deconvolution is limited to one facet per major cycle.

The serial CLEANing of facets has a number of drawbacks. First, it is computationally inefficient, each facet CLEANed requires a pass through the data to make the facet and then a read and write of the full data set to subtract the model. Since only a single facet is imaged at a time, the ability to use parallelization in the form of multiple cores and cluster nodes is reduced.

Another drawback to the facet at a time CLEANing is the potential for a degraded CLEAN. The details of the flux density remaining in other facets is unknown with the potential for mapping side-lobes into false structure.

A solution this these problems is to re-project the individual facets onto a common tangent plane and a common grid. The re-projection technique has been discussed by Leonid Kogan and others [2] and can be done in the griding process. This adds negligible cost to the griding and allows using a common grid for all facets. The instrumental response ("dirty beam") is a slow function of observing angle so the dirty beam determined in one facet is still a reasonable approximation in adjacent facets. This allows a joint deconvolution of multiple facets; this technique is explored further in this memo. The technique described here has been implemented in the imaging software in the Obit package [3] [1].

II. TWO VS. THREE DIMENSIONAL IMAGING

In the following "three dimensional imaging" (3D) refers to the case of using facets tangent to the celestial sphere at their midpoints. Since they are not coplanar, they consist of pixels separated in three dimensions. The term "two dimensional imaging" (2D) refers to the case in which the facets have been projected onto a common plane, tangent to the celestial sphere at the pointing center of the data. The pixels in this case are all coplanar and thus can be described by two dimensional coordinates. The coordinates in each facet use a common reference position. In Obit, the projection of the facets onto the common plane uses the technique of Kogan [2] as implemented in AIPS. The projection is done in the griding of the visibility data and adds a negligible cost.

A. 2D Common Grid

To take full advantage of the coplanarity of the 2D imaging, it is desirable that the various facets also have pixels located on a common grid of positions. This can be achieved by adjusting the shift applied to each facet such that the central pixel in the facet is an integral number of cell spacings from the reference position. This modification to the shift need not exceed 1/2 cells in either dimension.

B. Final CLEAN Model

Using 2D imaging, all facets are on a common grid and with a common reference position; the offsets from the reference position in the CLEAN components tables for each facet are in the same coordinate system. This means that a combined CLEAN components table can be attached to the final "Flattened" image. The coordinate systems used for the 3D imaging facets are all different and the CLEAN components cannot be so readily combined. The loss of the CLEAN model on the final image is a major drawback of the 3D technique.

National Radio Astronomy Observatory, 520 Edgemont Rd., Charlottesville, VA, 22903 USA email: bcotton@nrao.edu

---

[1] http://www.cv.nrao.edu/∼bcotton/Obit.html

III. MODIFICATION OF THE AUTOWIN ALGORITHM

Multi-faceted imaging is most necessary when imaging wide fields of view. In these cases, the deconvolution may not be strongly constrained in the sense that the number of degrees of freedom in the data may not greatly exceed the number of degrees of freedom (independent beam areas) in the image. The undesirable effects include the creation of fictitious structures, “CLEAN bias” and a spurious reduction of the “noise” in the image. In this case, the “autoWindow” technique described in [3] is applicable.

In the original implementation of autoWindowing, a single new box could be added to the CLEAN window of one facet each major cycle. The minor cycle CLEANing was stopped when the peak residual dropped significantly below the peak CLEANable [2] pixel in the facet outside the current CLEAN window.

This technique was improved in [4] to introduce an intermediate level of image based CLEAN cycles to allow multiple CLEAN boxes to be added to a facet in a given major cycle (see Section V). This technique is preferable to picking multiple boxes at a time as it readily distinguishes between sources and sidelobes.

For multi-facet CLEANing, the limit for minor cycle CLEANing is modified to include the consideration of other facets. The limiting peak residual in the minor cycle CLEANing is set to slightly below the maximum CLEANable pixel not in the current windows of any facet being jointly CLEANed. This value is adjusted downward by five times the RMS but to not less than half of the RMS residual. Furthermore, the minimum residual level is adjusted to be not below 70% of the peak CLEANable pixel minus three times the RMS in all facets not being included in the current major cycle CLEAN. The lower threshold for facets not in the current CLEAN is because the estimate of the peak CLEANable flux density may be based on an out–of–date image of the field. Note: the inclusion of facets not included in the current CLEAN is also applied in the case of 3D imaging. This technique results in a better quality deconvolution as it does not allow CLEANing on one facet to the point CLEAN is modeling side–lobes from a source in another facet.

IV. PARALLEL IMAGING

For data–sets larger than will fit in cache memory, it is desirable that multiple facets can be imaged in a single read of the data. For such large data-sets, the time required for the I/O exceeds that needed to make a single facet. Griding several facets at a time is more efficient on these large data sets as it reduces the total I/O. A parallel technique has been implemented in AIPS and a similar implementation made in Obit [5]. The current implementation in Obit is to image up to a maximum number of images that take no more than a estimated 1 GByte of buffer/work space. The griding of each facet is multi-threaded.

For 3D imaging, it is desirable to produce multiple facets each major cycle even though only one will be used as the pre–imaging estimate of the best facet to CLEAN next may turn out not to be best after the image is actually made. In this case, the “second best” facet needs to be reimaged and so on. If the “second best” has already been reimaged, it does not have to be done again. Tests in [5] suggested that with current technology, griding two facets took about the same time as one pass through non memory resident data. Thus, for 3D imaging, up to two facets are imaged at a time.

V. THREE LEVEL MULTI-FACET CLEAN

The current implementation in Obit of CLEAN for both 2D and 3D imaging now has three CLEANing cycles:

1) Minor cycle
   In the innermost loop, picking CLEAN components, only the brightest residuals and a limited region, the “Beam Patch” of the dirty beam is used. Since the residuals of the various facets being CLEANed are all on the same grid, the most significant pixels in all the facets being jointly deconvolved are included in the search list. The size of the Beam Patch and the limiting brightness residual considered are determined from a histogram of the side-lobe levels of the dirty beam from the field with the maximum pixel value and the histogram of the pixel values of all facets being CLEANed jointly. The residual list is searched for the largest absolute value residual and a fraction (loop gain) of this value is used as the flux density of the CLEAN component. The response to this component at the locations of the other residuals is estimated using the Beam Patch corresponding to the facet in which the component was found and then subtracted from all residuals. This cycle is repeated until either a maximum number of CLEAN components is located, or, the peak residual drops below either the target for the CLEAN, or, below a value estimated from the maximum pixel not included. This process is a modified version of that described by Clark [6]. An SDI implementation of this cycle is also implemented for imaging sources with very extended emission.
2) autoWindow cycle
   This cycle only has the potential for multiple loops when autoWindowing is enabled. If the minor cycle terminated because the residuals reached the level set by pixels not included, a middle, image plane, CLEAN cycle is used. In this cycle, the response to recently picked CLEAN components is subtracted from all pixels in all of the facets being jointly CLEANed using the Beam Patch of the facet in which the components were found. Then, the facet with the maximum CLEANable pixel is examined to see if a new box needs to be added to the CLEAN window. A further minor cycle may then be initiated. This cycle is terminated when either no new boxes are added, or, the minimum flux density due to the non-inclusion of a pixel does not drop by at least 10%, or, a maximum of 10 iterations is reached, or, the completion of the CLEAN. The image subtraction is multi-threaded

---

[2] “CLEANable” here means a pixel inside the outer window but outside any unboxes. The outer window defines the area to be CLEANed and inside of which any geometric (“w”) distortions are deemed acceptable.

reducing the computational cost when many facets are being jointly CLEANed.

3) Major cycle
In each outer cycle, the responses to the recently found CLEAN components are estimated and subtracted from the residual visibility data-set. A selected subset of the facets is then re-imaged and a new set of CLEAN components found. This cycle terminates when either a maximum number of CLEAN components is found, or, all facets are deemed to have been adequately CLEANed.

VI. AUTOCENTERING

One exception to the use of a common grid for all facets is the "autoCenter" technique [7]. This technique is to align the grids of individual facets to the peaks of very strong unresolved sources. The grids of these facets will not, in general, be aligned with the common grid. The autoCenter technique can be adapted to the multi-facet CLEAN technique by having two versions of the autoCentered facets. The first is one with the facet grid aligned with the source; this is used in a single facet-at-a-time CLEAN until the peak residual drops below a threshold level. The second version of the facet, with its grid aligned to the common grid, is used after the bulk of the bright source has been removed. This second version of the facet can be used in a multi-facet CLEAN. The version of the autoCentered facet aligned with the common grid is formed by shifting the autoCentered version using an FFT/phase ramp/FFT method. This shift will be no more than half a pixel in each axis so both versions of the image cover the same area.

VII. OBIT IMPLEMENTATION

The modification to 3D imaging as well as an implementation of a 2D multi-facet imaging were made in Obit, primarily in the ObitDConCleanVis class. Most of the complications in the implementation involve a number of decisions about how many facets to image, which ones to CLEAN jointly, how deep to CLEAN in the various cycles and when is it all done. The 2D multi-facet CLEANing option is enabled in user software by setting the "do3D" parameter to False.

- **How many facets to image in parallel?**
  At the beginning and end of each CLEAN all initial dirty and final residual images need be formed and it would be best if all could be made in a single read of the data. However, each facet takes a nontrivial amount of memory for data buffers and work space. Multi-threading increases these requirements. It is necessary to keep all the arrays in physical rather than virtual memory so there are practical limits to the number of parallel images that can be made. The compromise in Obit is to estimate the number of images that require the usage of 1 GByte of memory. This is based on the size of the visibility records, the size of the facets and the number of parallel threads used in imaging. For imaging with ionospheric corrections this is currently limited to a single facet at a time.
- **How many and which facets to image?**
  Early in the CLEAN only a small number of facets with the brightest emission need be imaged but further into the CLEAN, many facets may have comparable brightness emission and more facets should be imaged. Several statistical quantities are kept which help in making this decision. The most critical of these is the "quality" measure which, following the AIPS usage, is $0.95 * max\_res$ + $0.05 * avg\_res$ where max_res is the maximum absolute value CLEANable pixel value in the facet and avg_res is the average CLEANable pixel value in the facet. Facets with estimated quality measures within 30% of the best facet up to the maximum number of parallel images are imaged.
- **Single facet autoCenter vs Multi-facet CLEANing?**
  AutoCentered facets have grids aligned to the peak of the source and not the common grid; thus, these facets cannot be used in a joint CLEAN. However, each autoCentered facet has a version aligned with the common grid. Whenever the peak residual exceeds 10% of the autoCenter threshold, a single facet CLEAN is done using the autoCentered facet. Below this level the aligned version of the autoCenter facet is allowed to be included in a joint CLEAN.
- **How many facets to include in a joint CLEAN?**
  If an autoWindow facet with a peak flux density greater than 10% of the autoWindow cutoff is the best facet to CLEAN, then only this facet is CLEANed. In other cases, all freshly reimaged facets with a quality within 50% of the highest facet and for which CLEANing is not yet finished are included in the CLEAN.
- **When to add a new window box?**
  A box is added to the CLEAN window of the facet with the highest CLEANable pixel if this peak in not already in a box or unbox if the ratio of the peak exceeds the facet RMS by at least a factor 4 for a small source or at least a factor of 3 a more extended source (box radius $>$ 4 pixels).
- **When is CLEAN done?**
  The CLEAN is considered finished when either a maximum number of CLEAN components is found, or, all facets are deemed to have been adequately CLEANed. This latter criterion is that either the maximum CLEANable pixel in the facet is below a given minimum, or, the maximum is outside the current window but the SNR is so low (peak/RMS$<$5) that a new window was not added.

Most of the compromises made in the various decisions are intended to favor the large data case processed on multi–core computers when all data do not reside in memory. These changes to CLEAN are included in all self calibration and/or Peeling operations.

VIII. TEST COMPARISONS

Testing was performed on the Obit Development machine, mortibus, in Charlottesville. This machine has dual quad core Xenon processors for a total of 8 cores, a clock speed of 3 GHz, 8 GByte memory and a fast disk RAID system. Disk

/export/data_2 is a RAID 0 system using two 15K RPM disks on a Dell controller while /export/data_3 is a RAID 5 system using five 7.2K RPM disks on an Areca controller

Two sets of tests were performed, one with a real VLA data set that fits entirely in memory and the second, a simulated EVLA data-set which is larger than fits in memory. Each of the test data sets were imaged and CLEANed using Obit task Imager with both 2D (multi-facet CLEAN) and 3D (single facet CLEAN) imaging. All tests were allowed to use all 8 cores on the computer.

A. *Small Problem*

The "small problem" consists of imaging a set of VLA data obtained in a moderately deep VLA B configuration survey. This data consists of 567,531 visibilities with seven spectral channels in two "IFs" with both RR and LL correlations at 1.4 GHz. representing approximately eight hours of observation.

These data were imaged over a field of view of half–width 0.75° and a maximum of 15,000 clean components. The imaging used 116 facets including outlying sources and resulted in a final image 3600×3600 pixels in size. The results of the Small test timing comparisons are given in Table I. A portion of the image is shown in Figure 1. The two images derived are virtually indistinguishable.

The parallel imaging produced a maximum of 54 facets in each pass through the data. The 2D imaging made a total of 796 facet images while the 3D imaging made only 568. This accounts for most of the difference in CPU time used but since many of these imagings were done in parallel, the increased efficiency of the multi–threading resulted in a much shorter run time for the 2D imaging.

B. *Large Problem*

The simulated data-set used in the tests described in [8] was used to test the performance for a data-set larger than will fit in memory. Data samples were 2 second integrations and contained 1024 spectral channels divided among 32 ”IFs” and spanning from 1.4 to 1.9 GHz. This generated a total of 2,323,269 visibility records of which 2,003,508 were on the ”Target”; derived from a sky model similar to that derived from the small data set presented above. This data-set is approximately 100 GByte in size but the Stokes I portion is only a quarter of this. This is still larger than the 8 GByte of memory available. The simulation used the sky model derived from the data used in the small test.

The imaging for these tests used a field of view of half–width 0.35° requiring 42 facets including outliers. The final images are 1679×1679 pixels. The comparison of the 2D and 3D imaging are given in Table I. The parameters of the two executations of Imager were identical except for the value of do3D.

The parallel imaging could make all 42 facets in a single pass. The 2D imaging made a total of 207 facet images while the 3D imaging made 268. Most notably, the serial 3D imaging took 3.7 times longer that the parallel 2D imaging. The resulting images are visually indistinguishable.

IX. POTENTIAL FOR FURTHER PARALLELIZATION

The current implementation makes heavy use of threading for the expensive operations when multiple cores or shared memory processors are available. For very large continuum data sets, a cluster could be efficiently used by splitting the residual visibility data by frequency over the nodes. Since image formation is a linear process, the frequency channels on each node could be imaged and the final dirty/residual image or beam formed by summing and normalizing the images from the various blocks of channels. This requires moving the partial images between nodes but the volume of these should be much less that that of the visibility data.

The CLEAN model visibility subtraction could then be accomplished by copying the CLEAN components to each node and then subtracting them from the visibility data resident on that node. The volume of data needed for the CLEAN components is rather minor.

The bulk of the processing time in these tests was used in the griding of the visibility data onto the grids used to Fourier transform the data; this operation is easily parallelized. The much larger number of images produced in each major cycle of the 2D imaging allows greater opportunity for parallel computing, and in particular, the use of clusters than does the 3D imaging.

X. CONCLUSIONS

The use of multi–facet CLEANing is shown above to have substantially better performance on a multi–core computer than the serial facet CLEANing. One of the major motivations for this technique was to minimize the number of major cycles, hence the I/O needed for large data–sets. However, even in the small test presented for which the data–set easily fit in disk cache (largely eliminating the I/O) the multi–facet CLEAN was substantially faster (1.8×) in real time. Due to the larger number of images formed, the multi–facet CLEAN used more CPU time but made much more efficient use of the 8 cores available through multi–threading to more than make up for this.

The large test presented had sufficient data that it could not reside in memory and thus actual I/O had to be performed. In this case, the 3D serial imaging took 3.7 times longer than the parallel imaging, 29 vs. 7.6 hours. This is a dramatic demonstration of the power of the parallel, multi–facet CLEAN technique.

The multi–facet CLEAN shows much promise for dealing with the very large data–sets to come from the next generation of instruments now under construction (EVLA, ALMA, eMERLIN, LOFAR). Further enhancements in performance can be had by adapting this technique to cluster computers.

REFERENCES

[1] R. A. Perley, "Imaging with Non-Coplanar Arrays," in *Synthesis Imaging in Radio Astronomy II*, ser. Astronomical Society of the Pacific Conference Series, G. B. Taylor, C. L. Carilli, and R. A. Perley, Eds., vol. 180, 1999, pp. 383–+.

[2] L. Kogan and E. Greisen, "Faceted imaging in AIPS," *AIPS Memo*, vol. 113, pp. 1–4, 2009.

[3] W. D. Cotton, "Obit: A Development Environment for Astronomical Algorithms," *PASP*, vol. 120, pp. 439–448, 2008.

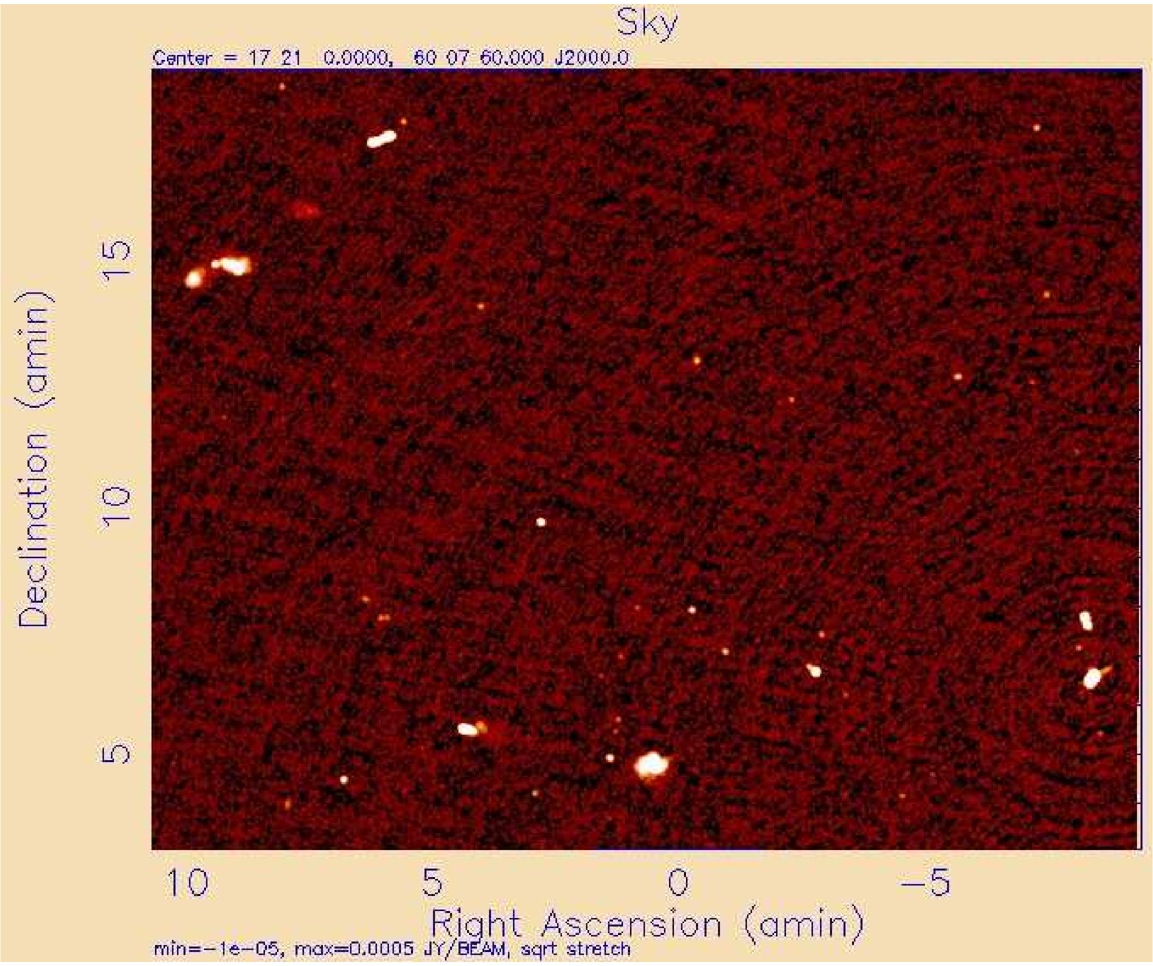

Fig. 1. A section of the image derived in the small data test.

TABLE I
RESULTS FROM IMAGER TESTS

Type	Real min	CPU min	CPU/real	no. components	CLEAN flux mJy	RMS μJy	no. major cycles
2D Small	9.87	49.86	5.05	8861	361.8	20.3	26
3D Small	17.42	40.09	2.30	9934	374.3	20.9	170
2D Large	462.2	2270.2	4.91	1515	41.2	6.89	12
3D Large	1732.3	3781.3	2.18	1611	41.8	6.80	92

[4] ——, “Performanence Enhancement of the autoWindow Technique,” *Obit Developement Memo Series*, vol. 9, pp. 1–3, 2009.
[5] ——, “Parallel Facet Imaging in Obit,” *Obit Developement Memo Series*, vol. 6, pp. 1–3, 2009.
[6] B. G. Clark, “An efficient implementation of the algorithm ’CLEAN’,” *A&A*, vol. 89, pp. 377–+, Sep. 1980.
[7] W. D. Cotton and J. M. Uson, “Pixelization and Dynamic Range in Radio Interferometry,” *A&A*, vol. 490, pp. 455–460, 2008.
[8] W. D. Cotton, “Testing Obit with a 100 GByte simulated dataset,” *Obit Developement Memo Series*, vol. 8, pp. 1–8, 2009.