Tests on a High Frequency EVLA TeraByte Simulated Dataset

W. D. Cotton, June 5, 2009

***Abstract*—This memo describes test processing of a high frequency EVLA terabyte scale simulated VLA A configuration dataset using Obit on a fast workstation. Both an 8 GHz bandwidth continuum test at 40 GHz and a 1792 channel line test were performed. The continuum test took less than 13 hours to run and the line test 43.3 hours. A good workstation would be able to process similar data in of the order of a few times the observe time.**

***Index Terms*—interferometry, performance**

I. Introduction

Large datasets will be a prominent feature of the EVLA when it comes on line. Some of the impact at low frequencies was explored in Obit Development Memo 8[1] [1]. and Obit Development Memo 11[2] [2]. This memo is a follow up on the development in the latter using Obit [3] [3] on a terabyte scale dataset simulating VLA A configuration observations testing the possibility of using a workstation for this scale project.

II. Simulated Terabyte Dataset

The simulated dataset contains 53 x 10 min scans on ”Target” plus calibration scans on ”Cal” and ”Amp” spread over 9 hours. There are 1792 channels divided into 32 ”IFs” of 56 channels equally spaced covering the 8 GHz from 40 to 48 GHz. The uv coverage used was that of the VLA A configuration. The Target model was derived from a deep 1.4 GHz B Array image with 131 point sources with the peak 20 cm flux densities and a distribution of locations derived from the 20 cm results but scaled by the ratio of frequencies. There are a total of 11,610,300 x 1 second integrations. The total size of this dataset (with visibilities as a triplet of floats) is 930 GByte. 50 mJy Gaussian noise per correlation was added; the model has no frequency dependence. There is 127.4 mJy total flux density in the model.

III. Processing Tests

Processing followed that used in Obit Development Memo 11. The tests included a broadband continuum imaging case using all the data and a “spectral line” case where each channel was individually imaged and then the channel images accumulated into an image cube.

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

[1]http://ftp.cv.nrao.edu/NRAO-staff/bcotton/Obit/HundredGB.pdf
[2]http://ftp.cv.nrao.edu/NRAO-staff/bcotton/Obit/QBand100GB.pdf
[3]http://www.cv.nrao.edu/∼bcotton/Obit.html

A. Continuum

External calibration was applied and then a single 3600x3600 0.015” pixel facet was formed and deconvolved using a maximum of 5000 iterations of ”Cotton-Schwab” CLEAN or a limiting CLEAN flux density of 50 microJy.

B. Spectral Line

External calibration was applied and then a 3600x3600x1792 spectral cube was formed. Each channel image was deconvolved using a maximum of 2500 iterations of ”Cotton-Schwab” CLEAN or a limiting CLEAN flux density of 50 microJy. The processing consisted of:

1) **scatter (SplitCh)**
   Calibrate the data and split into 32 ”IF” uv data files.
2) **image (Imager**
   Multiple parallel streams imaging and deconvolving.
3) **gather (MCube)**
   Combine images into single cube.
4) **cleanup**
   Delete temporary files.

With minor modifications, this procedure can spread computing of a spectral cube over nodes of a cluster.

C. Multistream Tests

The line processing tests in [2] found that it was more efficient to use 4 parallel streams with 2 CPU cores each rather than using 8 streams with a single core. In that test, the bulk of the portion of the problem being dealt with at any time largely fit in memory for 4 but not for 8 streams; this makes actual I/O performance in the 4 stream case less critical. With the much larger data set being considered here, the problem is larger than the 8 GByte memory even for a single processing stream, actual I/O performance is important.

Simple tests were performed on single stream, two parallel streams and four parallel streams - each with two cores allocated. The tests consisted of runs in which 10 or more channels were processed in each of the parallel streams; this should be sufficiently long (at least half an hour) for an accurate measurement. The single stream used data on the fastest disk available, the other two tests split data between the fast disk and a slower(?), but larger RAID disk. The results are shown in Table I which gives the measured average real time to process a single channel as well as the estimated time to process the full cube in that mode. Four stream processing is the clear choice.

TABLE I
STREAM TIMINGS

No. Streams	Real time/channel/stream min.	Estimated total hrs
1	2.73	80.6
2	3.28	49.0
4	5.65	42.2

TABLE II
TIMINGS

Process	Real Time hr	Comments
Continuum	12.9	CPU/real = 2.6
SplitCh	1.4	
Imager	41.5	Slowest of four streams
MCube	0.264	
cleanup	0.126	
total	43.3	

IV. TIMING TESTS

The timing tests used mortibus, the Obit development machine in Charlottesville. Mortibus has dual quad core Xenon processors for a total of 8 cores, a clock speed of 3 GHz, 8 GByte memory and two fast disk RAID systems. 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

For the continuum test, a single facet was formed and deconvolved; this took 2 major cycles (the image has a really nice dirty beam). For the line processing case, four stream processing with two cores per stream was used based on the results given in Section III-C. CLEANing of each line channel took 4 major cycles. The timing results are given in Table II.

Even though each of the line processing streams had equal amounts of data on each of the two output disks, there was still a significant dispersion in the run times for the streams; the fastest took 40.3 hr and the slowest (given in Table II) took 41.5 hours. Examining the results more closely, these difference become more curious. Table III shows the details of individual runs of Imager on blocks of 56 channels. During these measurements, there were four parallel streams running, each using one of disks /export/data_2 or /export/data_3. The afternoon/early evening run time using /export/data_3 was 30% longer that after midnight run at which point performance using this disk was comparable to using /export/data_2. The current conjecture that this anomaly was caused by hardware problems on /export/data_3 which was logging errors. If so, these problems probably added at least an hour to the total run time.

TABLE III
IMAGER TIMINGS

Avg. time	Stream	disk	Real time hrs
0/15:41:25	1	/export/data_2	4.42
0/20:09:41	1	/export/data_2	4.54
1/00:42:11	1	/export/data_2	4.56
0/16:27:28	2	/export/data_3	5.95
0/22:05:09	2	/export/data_3	5.30
1/02:59:59	2	/export/data_3	4.51
0/15:41:27	3	/export/data_2	4.42
0/20:09:47	3	/export/data_2	4.54
1/00:42:16	3	/export/data_2	4.54
0/16:27:39	4	/export/data_3	5.95
0/22:06:43	4	/export/data_3	5.35
1/03:03:08’	4	/export/data_3	4.54

V. DISCUSSION

The continuum imaging problem was so simple that the bulk of the time was taken in the initial data manipulation, applying calibration and selection, uniform (robust) weighting and conversion to a single Stokes dataset from which to form residuals. The heavy computing portion of the process used the 8 cores making this part of the processing sufficiently fast that the scalar, I/O intensive portion of the process dominated. Apparent disk hardware problems slowed the line imaging test, and possibly the continuum test, by an unknown amount.

Processing of high frequency data should not require more than a hefty workstation for datasets of this size (typical? for A array). Simple continuum imaging is of order the observe time. In reality, wideband continuum imaging, as in this test, will have to account for varying spectral index and perhaps curvature across the field as well as the frequency dependence of the primary beam size. Since the techniques for these are not yet fully worked out, it is not possible to make a timing test but will likely take several times longer than the results presented here. The spectral line reduction took longer that the continuum image but is still only a few times that of the observe time.

ACKNOWLEDGMENT

The author thanks Tracy Halstead, Josh Malone and Dan Klopp for their assistance.

REFERENCES

[1] W. D. Cotton, “Testing Obit with a 100 GByte simulated dataset,” *Obit Development Memo Series*, vol. 8, pp. 1–8, 2009.
[2] ——, “Testing Obit with a 100 GByte Simulated Dataset,” *Obit Development Memo Series*, vol. 11, pp. 1–2, 2009.
[3] W. D. Cotton, “Obit: A Development Environment for Astronomical Algorithms,” *PASP*, vol. 120, pp. 439–448, 2008.