ALMA Correlator Output Data and Computer Processing Rates

Jim Pisano

July 1, 1999

1. Introduction

This document provides a preliminary estimate of the data flow out of the ALMA Correlator Long Term Accumulator (LTA) and computational needs to process the resulting data flow.

2. Overview

The various criteria that are used in these calculations include:

- the data rates of the LTA that Chuck Broadwell provided (see Appendix below).
- the maximum data rates are used which is the case when all antennas are producing both cross-correlation and auto-correlation results.
- there are three data processing functions: a VanVleck correction, a Hanning window correction, and a digital Fourier transform. The VanVleck correction was taken from code used in the MAC 12-meter spectrometer. The Hanning window function was taken from code used by Ray Escoffier for his digital filter simulations. The FFT used is MIT's FFTW (see http://theory.lcs.mit.edu/~fftw/ for details).
- the following procedure was used to empirically determine a FLOP (floating point operation):
    - created a test program that performed the three processing functions: VanVleck correction, Hanning windowing function, and FFT. Two lag arrays were created as test input. One had 512 elements and the second had 16,384 elements which correspond to the smallest and largest number of lags per correlator array (which can relate to a baseband pair). These input data were generated from a digital filter simulation program written by the author.
    - compiled the FLOP rate test program using gnu compiler version 2.8.1 with general optimizations (**-O** flag) enabled.
    - executed the FLOP rate test program on the SGI computer (argus) at AOC using the Workshop Debugger (cvd) which has the ability to count FLOP's for a given function. Care was taken to ensure that no FLOP's were used for test set up, only for the relevant calculations.
- raw power spectrum information in decibels is produced, i.e., overhead to create FITS-compatible data sets is ignored.

The results of the FLOP's rate test program are presented in tables 1 and 2 in § 3.

3. Discussion

First and foremost, all of these rates are based on current, preliminary designs of the LTA (and the correlator cards). As these designs are still in development, these estimates may change in the future.

The two tables show the data and FLOPS (floating point operations per second) rates for two sample spectral sizes. Table 1 demonstrates the data and FLOPS rates for a 256-point spectra (512 total results of leads and lags) that correspond to the ALMA correlator's 2 GHz bandpass (maximum bandpass and minimum spectrum size) while Table 2 shows the rates for a 8,192-point spectra (16,384 total results of leads and lags) with a 62.5 MHz bandpass (minimum bandpass and maximum spectrum size). These values are for one 2 GHz (or 62.5 MHz) baseband pair which can be combined to provide, for example, a 1024-point spectrum for a total bandwidth of 8 GHz (4 x 2-GHz baseband pairs). The total FLOPS rates for the full 8 GHz and 250 MHz bandwidths are shown in the last columns of Table 1 and Table 2.

The top part of each table is based on data rate calculations for 2 through 64 antennas provided by Chuck Broadwell and represents the number of results to be processed to produce raw spectral points. The following descriptions apply to the columns:

- **Cross Results Quantity** is: $N_c = N_b \times (N_{leads} + N_{lags})$ where N_c is the number of cross results, N_{leads}, N_{lags} are the number leads and lags and N_b is the number of baselines: $N_b = N_a \times (N_a-1) / 2$ where N_a is the number of antennas.
- **Self Results Quantity** is: $N_s = N_a \times N_{lags}$ where N_s is the number of self results, N_a and N_{lags} are the number of antennas and lags.
- **Self/Cross Results per second** are the self and cross results rate based on the 16.77712 (.048576 * 16) millisecond minimum correlator long term integration cycle: $R_c = N_c / 0.01677712$ *secs,*
  $R_s = Ns / 0.01677712$ *secs.*
- **Total Cross + Self Results per second** represent the total data output rate in 4-byte LTA results with the auto results duplicated: $R_t = R_c + R_s \times 2$, where R_t is the total data output rate – see below for details.
- **Number of 256 (or 8K) spectra/sec** represents the total output rate as sets of the spectrum size: $R_{sp} = R_t / (N_{leads} + N_{lags})$. This rate can also be represented by: $R_{sp} = (N_b + N_a) / 0.01677712$ *secs.*

An important point to note deals with the number of lead and lag results for cross-correlations versus auto-correlations. The cross results contain both lead and lag quantities while the auto results only contain lags requiring duplication to provide the leads. This point is taken into account in the Total Cross + Self Results per second column thus changing slightly from the true output of the LTA to be the true input to the data processing calculations. This duplication has a greater impact on the small antenna cases, but becomes insignificant for the large antenna cases as the number of cross products dominate.

The bottom part of each table shows the floating point rate calculations for the three data processing steps: FFT, VanVleck correction, and Hanning window. These FLOPS rates are the products of the number of spectra per second and the number of FLOP's for each set of results of a given size (either 512 or 16,384), the latter of which are shown in each column header. The Total MFLOPS column shows the total number of MFLOPS (10^6 floating point operations per second) to perform these calculations for the given data rates. The final column shows the FLOPS rates for all 4 baseband pairs together and is simply 4 times the MFLOPS in the previous column.

Although the total data quantity from the LTA does depend on the baseband bandwidth, which can range in six steps from 2 GHz to 62.5 MHz (Table 5 shows this relationship), the total rate of results does not change. This is due to the accumulations performed by the LTA and the trade-off between the number of spectral points versus the long term integration rate – again see Table 5 for this relationship. The LTA will have a basic integration interval of 16 milliseconds in cross-correlation mode, i.e., both cross- and auto-results will dump on intervals that are multiples of 16 milliseconds.

2 GHz Bandwidth w/ 512 results (leads & lags -- 256 spectral points) Integration Time: 16.78 ms						
Antennas	**Cross Results Quantity**	**Cross Results per second**	**Self Results Quantity**	**Self Results per second**	**Total Cross + Self Results per second**	**Number of 256 spectra/sec**
2	512	30,518	512	30,518	91,553	178.81
4	3,072	183,105	1,024	61,035	305,176	596.05
8	14,336	854,492	2,048	122,070	1,098,633	2,145.77
16	61,440	3,662,109	4,096	244,141	4,150,391	8,106.23
32	253,952	15,136,719	8,192	488,281	16,113,281	31,471.25
64	1,032,192	61,523,438	16,384	976,563	63,476,563	123,977.66
	FFT (MIT FFTW) FLOPS 13,288[1]	**VanVleck Correction 11,318[1]**	**Hanning Windowing 14,689[1]**	**Total MFLOPS**	**Total MFLOPS for 8 GHz Bandwidth**	
2	2,376,080	2,023,816	2,626,598	7	28	
4	7,920,265	6,746,054	8,755,326	23	94	
8	28,512,955	24,285,793	31,519,175	84	337	
16	107,715,607	91,746,330	119,072,437	319	1,274	
32	418,190,002	356,191,635	462,281,227	1,237	4,947	
64	1,647,415,161	1,403,179,169	1,821,107,864	4,872	19,487	

Table 1

62.5 MHz Bandwidth w/ 16,384 results (leads & lags -- 8K spectral points) Integration Time 536.9 ms						
Antennas	**Cross Results Quantity**	**Cross Results per second**	**Self Results Quantity**	**Self Results per second**	**Total Cross + Self results per second**	**Number of 8K spectra/sec**
2	16,384	30,518	16,384	30,518	91,553	5.59
4	98,304	183,105	32,768	61,035	305,176	18.63
8	458,752	854,492	65,536	122,070	1,098,633	67.06
16	1,966,080	3,662,109	131,072	244,141	4,150,391	253.32
32	8,126,464	15,136,719	262,144	488,281	16,113,281	983.48
64	33,030,144	61,523,438	524,288	976,563	63,476,563	3,874.30
	FFT (MIT FFTW) FLOPS 420,103[2]	**VanVleck Correction 360,192[2]**	**Hanning Windowing 466,932[2]**	**Total MFLOPS**	**Total MFLOPS for 250 MHz bandwidth**	
2	2,347,508	2,012,730	2,609,186	7	28	
4	7,825,028	6,709,099	8,697,286	23	93	
8	28,170,101	24,152,756	31,310,230	84	335	
16	106,420,383	91,243,744	118,283,093	316	1,264	
32	413,161,486	354,240,417	459,216,714	1,227	4,906	
64	1,627,605,855	1,395,492,554	1,809,035,540	4,832	19,329	

[1] Number of floating point operations for 512 results

[2] Number of floating point operations for 16,384 results

Table 2

4. Comparison with the Test Correlator

The test correlator, a.k.a., the GBT clone correlator, will be used with the prototype array and ALMA control software during the D&D phase of the prototype ALMA correlator. Table 3 provides comparison data and FLOPS rates for the test correlator. These values were derived from "GBT Spectrometer", Volume 1, R. Escoffier, December 1998. The data rate values for the test correlator are based on the following criteria:

- A cross-correlation long term integration from the correlator contains 512 lead and 512 lag results for a total of 1024 results (a 512 point spectrum). An auto-correlation long term integration from the correlator contains 1024 lags which are duplicated to leads for a total of 2048 results (for a 1K spectrum). Note that the test correlator will not be able to provide auto and cross results at each long term integration interval, so, for example, the cross results will be obtained for a given period and then auto results are obtained during the next interval.
- The integration time was set for 100 ms. which is a rate used by the MAC spectrometer at the Tucson 12-meter (private communication with Matt Waddel). The test correlator has a minimum long term integration rate of approximately 80 ms.

Test Correlator						
800 MHz Bandwidth with		1024 Total results (leads & lags)			Integration Time:	100 ms
Antennas	Cross Results Quantity	Cross Results per second	Self Results Quantity	Self Results per second	Total results per second	# of 512/ 1K spectra/ sec
2	1,024	10,230			10,230	9.99
2			2,048	20,460	40,919	19.98
	FFT (MIT FFTW) FLOPS 29,263 / 52,854[3]	VanVleck Correction 22,572 / 45,080[3]	Hanning Windowing 29,263 / 58,452[3]	Total MFLOPS		
2	264,156	225,495	292,338	1	cross correlations	
2	1,056,024	900,699	1,167,872	3	auto-correlations	

Table 3

Since the test correlator must perform cross correlations and auto correlations independently, Table 3 shows the 2 sets of data rates and FLOPS rates on separate rows. Also the spectra rates refer to 512 spectral points for cross products and 1024 spectral points for self products.

5. Conclusion

The approximate computing rate for the maximum data rate for that can come from the prototype correlator with 2 antennas is about 30 MFLOPS while the test correlator can range from about 1 - 3 MFLOPS. The MFLOPS rate increases dramatically for the 64 antenna case, to about 30 GFLOPS (billion floating point operations per second.

The selected microcomputer for the real time control computer is a Motorola MVME 2700-1361 which uses a 266 MHz PowerPC 750 CPU. The estimated SPECfp95 for this computer is 9.4[4]. In its technical description of the PowerPC 750, Motorola states that multiply-add operations are

[3] Number of floating point operations for 1024 or 2048 results

[4] See "MVME2700 Processor Module with MVME761 I/O" at http://www.mcg.mot.com/litfiles/27001.pdf

pipelined to 1 clock cycle[5]. This leads to a value of ≈ 250 MFLOPS. Unfortunately, SPECfp95 rates do not correlate to MFLOPS, so these calculations are only ballpark estimates. Once the VME computer is up and running, timing tests on the same set of calculations used for this paper can be performed resulting in more accurate estimates of the calculation rates. The results of these tests should be available in about a month.

An alternative to a general purpose computer is a DSP card. Chuck Broadwell has been investigating DSP cards to provide an interface between the VME bus and the ALMA prototype LTA. Once such card has 4 TMS320C6201 floating point processors that run at 167 MHz. In product literature the manufacturer states that this board delivers 4 GFLOPS. Thus approximately 8 - 10 of these cards could produce the estimated 30 GFLOPS required in the 64 antenna case.

6. Appendix

29 April, 1999 CMB PRELIMINARY

Jim P. is looking for information on the data rate that will be available from the 2 antenna prototype correlator for the ALMA. This is not the GBT clone, but the first prototype using new ALMA correlator hardware.

At the LTA output, the real-time computer will see at most 2048 multiplier results (1024 leads and 1024 lags) per cross baseline and with four bytes per result from the full system of four baseband pairs. The values listed below are based on the 2048 and 1024 numbers. If fewer results per baseline are required, the rates go down in a linear fashion.

Given N antennas, there are N*(N-1)/2 cross baselines and N self baselines. The following table shows the results rates at the LTA output for cross plus self results, where all results are dumped every 16 milliseconds. This is the fastest rate planned. If the user requests longer LTA accumulation periods, the rates decrease in a linear fashion.

Number Antennas	Number Cross Results	Number Cross Results/second	Number Self Results	Number Self Results/second	Total Self + Cross Results	MTotal Results/ second
2	2,048	122,070	2,048	122,070	4,096	0.244
4	12,288	732,422	4,096	244,141	16,384	0.977
8	57,344	3,417,969	8,192	488,281	65,536	3.906
16	245,760	14,648,438	16,384	976,563	262,144	15.625
32	1,015,808	60,546,875	32,768	1,953,125	1,048,576	62.500
64	4,128,768	246,093,750	65,536	3,906,250	4,194,304	250.000

Table 4

The data in Table 5 were provided by Chuck Broadwell and relate the accumulation mode with the minimum integration time, the number of spectral channels per baseband pair and the bandwidth for each baseband.

[5] See "PowerPC Advanced Information: MPC750 RISC Microprocessor Technical Summary" at http://www.motorola.com/SPS/PowerPC/teksupport/teklibrary/techsum/750ts.pdf

ACCUMULATION MODE NUMBER	MIN INT TIME	SPEC CHANS EACH ARRAY	BANDWIDTH EACH BB
0	16 msec	256	2 GHz
1	16 msec	256	1 GHz
2	16 msec	256	500 MHz
3	16 msec	256	250 MHz
4	16 msec	256	125 MHz
5	16 msec	256	62.5 MHz
6	32 msec	512	1 GHz
7	32 msec	512	500 MHz
8	32 msec	512	250 MHz
9	32 msec	512	125 MHz
10	32 msec	512	62.5 MHz
11	64 msec	1024	500 MHz
12	64 msec	1024	250 MHz
13	64 msec	1024	125 MHz
14	64 msec	1024	62.5 MHz
15	128 msec	2048	250 MHz
16	128 msec	2048	125 MHz
17	128 msec	2048	62.5 MHz
18	256 msec	4096	125 MHz
19	256 msec	4096	62.5 MHz
20	512 msec	8192	62.5 MHz

Table 5

7. Acknowledgments

I am indebted to Chuck Broadwell for his extremely patient assistance in the verification and clarification of most aspects of this paper. Ray Escoffier helped clarify the differences between the test correlator and the ALMA prototype correlator.