Enhancing the Spectral Performance of the 64-antenna ALMA Correlator

NAASC Memo #114

R. Escoffier (NRAO), R. Lacasse (NRAO), Alejandro Saez (JAO), John Webber (NRAO), Rodrigo Amestica (NRAO), Alain Baudry (U. of Bordeaux)

Date: March 18, 2015

ABSTRACT

A way of enhancing the spectral resolution of the ALMA baseline correlator, and hence the ALMA science capabilities, without expensive investment is presented. The idea is to take advantage of the increase in the performance of digital electronics made since the initial design of the correlator, while retaining the basic infrastructure and most of the current circuit boards.

Introduction

More than 16 years have passed since the initial design of the baseline ALMA correlator and according to Moore's law, the capability of digital logic (at equal cost) should have increased by a factor of more than 250 over this time period. The custom ASIC correlator chips used in the system were manufactured using a 250-nanometer technology, the same used in the earliest generation Pentium 3 microprocessors.

One area of the ALMA correlator that may be able to take advantage of this performance increase but not require prohibitively expensive upgrades in the rest of the ALMA system is spectral resolution.

The initial parameter to be selected is how much of an enhancement in spectral performance to shoot for. Factors of 4, 8, or 16 in resolution are probably possible, but for this memo, as 8 seems to be a good cost/performance trade, an increase of 8 in spectral resolution will be investigated (resulting in 65,536 spectral points per baseline per quadrant). Thus, all modes now supported by the correlator would show an 8-fold improvement in spectral resolution (although technical reasons may limit the improvement Time Division Modes will be able to realize).

A brief strategy to take advantage of modern digital electronics follows.

Concept

The plan would be to discard the current Correlator and Final Adder logic cards. These cards would be replaced by a newly designed Correlator card and an Output Interface card. (The present LTA cards will remain, but only to serve as control cards for the upgraded correlator design.)

Sets of 4 correlator cards would be replaced by a single newly designed plug-in compatible correlator/LTA card (128 required for the entire system). This card would duplicate the current correlator card function for an entire 64X64 antenna matrix (currently requiring 4 correlator cards) but with 8 times as many total hardware correlator circuits, and would incorporate some of the Long Term Accumulator function.

Ethernet outputs from the 32 new correlator cards per quadrant would drive a newly designed Output Interface card. This card would also provide some summation functions for Time Division Modes.

Optimum performance would probably require a complete redesign for all aspects of this new hardware. However, a lot of time and design effort could be saved by copying as much of the present design as possible, just with 8 times as many hardware lags and integration results. This way it could also more easily be made backward compatible to minimize interruption in ALMA operation.

Custom Chip Size Estimate

The new correlator card would require an increase in the level of integration by a factor of 32 (4 cards to one, 8 times as many lags) and should be easily accommodated by the greater than 250 increase in logic circuits per dollar since the original design.

A rough count of flip/flops in the present ALMA custom correlator IC yields about 200,000 f/fs. Currently, one correlator plane (4 correlator cards) contains 256 of these custom chips or about 50 Mf/f, total. Thus a redesigned plane would require 50 X 8 or 400 Mf/fs. Unfortunately, this level of integration is likely to be well beyond the realm of FPGA technology (both in cost and power consumption) so a new custom chip would probably be required.

The size of a new custom correlator chip depends on how many chips per card:

- 16 chips per correlator card requires a custom chip with 25 Mf/f
- 32 chips per card requires 12.5 Mf/f
- 64 chips per card requires 6.25 Mf/f

Which level of integration it would be best to use requires a more detailed design study and cost analysis but a feature size of 65/55 or 28 nm seems to be the most likely option.

Other Performance Considerations

In addition to enhancing the spectral resolution, many observations may also be able to take advantage of better sensitivity. When the highest spectral resolution is not needed, 4-bit X 4-bit and double Nyquist modes could be used in a higher percentage of observations (and still get better spectral resolution than the current system). This increase in observing sensitivity of about 12% and 7%, respectively, would be equivalent to the collecting area of up to 5 additional antennas.

Higher dump rates could also be provided, since a newly installed output computer system would probably be able to accommodate a higher output data rate.

Even higher dump rates could be made available by trading frequency resolution for greater time resolution; for example, keeping only ¼ the spectral points for a 4 times improvement in time resolution.

Backward Compatibility

If desirable, the correlator modification could be made to have the same data output computer interface as present and thus be somewhat backward compatible. In such a case, the spectral resolution improvement could still be available for longer integration times, providing the current maximum output rate is not exceeded.

Such an implementation strategy would allow more time for the design of and software development for a new output computer system to fully support the upgrade.

No modification to the correlator cooling system would be required since the total power used will probably go down somewhat.

Other System Improvements

This upgrade project might provide the opportunity for a few other low cost system improvements. For example, since the data output of the correlator ultimately would be converted to optical signals, moving the CDP cluster to the OSF would allow easier maintenance and better reliability.

Implications for the Computer System

At present, the 64-antenna, 4-quadrant, correlator can produce data output volumes of up to about 16.8 GB/s (134,217,728 sixteen-bit lags every 16 msec) from the correlator cards. After the proposed upgrade, that figure becomes 134 GB/s. For a 50-antenna observation (no ACA antennas), the output rate would be 82 GB/s.

The backward compatibility feature of the upgrade design means that the computer system need not be immediately upgraded to obtain at least some advantage in time or frequency resolution; thus the computer system could take a natural evolutionary path.

Performance

For a back end computer system capable of handling 1/N of the 134 GB/s output capacity of the proposed correlator system, an equation of performance can be specified as:

$$R = 4096T/N$$

Where: R is spectral resolution in spectral points per baseline per quadrant and T is time resolution in milliseconds. Extreme examples for, say, N = 16 (or a computer system capable of inputting about 8 GB/s) would be:

- T = 1 msec, R = 256 points
- T = 16 msec, R = 4,096 points
- T = 256 msec, R = 65,536 points

How to Proceed

A study requiring two engineers and a computer professional (perhaps a 0.5 FTE effort) would be needed to ascertain the full feasibility of this correlator enhancement and to come up with reasonable cost and schedule estimates.

An initial survey of possible custom chips to implement an increase of 8 in spectral resolution yielded an estimated total cost of about $2,700,000 for development and all chips needed. All other correlator hardware to complete the upgrade would be small compared to this.

Reference to "ALMA Scientific Specifications and Requirements"/Discussion on spectral resolution: see ALMA EDM CORL-60.00.00.00-0133-A-GEN, A. Wootten, 2014-08-14 (or see below)

```
Summary of discussion:
ALMA Scientific Specifications and Requirements
CORL-60.00.00.00-0133-A-GEN

States:

The most interesting narrow lines are those which have self-absorption which
indicates infall ("infall asymmetry").  In that case the minimum number of
channels across the line needed to model the line properly increases, from
perhaps 2 resolution elements = 1 sigma for a simple gaussian line to something
like 4 resolution elements = 1 sigma for a line with a dip or a red shoulder.
If one applies these criteria to an extreme case (low frequency, heavy
molecule, low temperature), then one should consider e.g.  an HC3N line at 100
GHz, in a cold, slowly contracting starless core with a central temperature of
8 K.  The thermal velocity dispersion would be 0.036 km/s and so the spectral
resolution should be about 0.018 km/s to resolve a gaussian line or 0.01 km/s
```

to resolve a self-absorbed gaussian line [SCI-90.00.00.00-00030-00]. The baseline correlator provides a resolution of 0.011 km/s, meeting this need.

According to ALMA Memo 556, Mode No 12, 2x2 bits, double polarization using TFB provides a bandwidth of 58.6 MHz spanned by 3840 channels, for a spacing of 0.015259 MHz, or .046 km/s at 100 GHz. At Band 6, this would be 0.023 km/s and the resolution would be just capable of resolving the above thermal velocity dispersion (though at the higher frequency).

But, currently, according to the OT at 100 GHz, the highest default resolution available is 0.092 km/s, from the 58.6 MHz bandwidth, with a Hanning smoothed resolution of 30.5 kHz. However, one can select other windows at the Phase 2 stage to improve the resolution by up to 80% as described in 4.4.3 of the Technical Handbook. These other choices are:

* Uniform: 1.2xchannel width (this is what SMA uses, but narrow features will ring)
* Welch: 1.59xchannel width (provides 19dB of ringing suppression)
* Cosine: 1.64xchannel wdith (provides 20dB " ")

With these choices one could obtain 0.021 km/s resolution for instance in the middle of B7.

In fact, proposals seeking to do just the science case in the requirements have been turned down owing to ALMA's inadequate spectral resolution. Insofar as I have been able to ascertain, there has been no use of these other modes implemented at Phase 2.

Correlator Mode 31, providing a bandwidth of 31.25 MHz and 3840 effective channels, should provide twice the resolution of Mode No 12, 2x2 bits, dual polarization, for a resolution of 0.046 km/s at 3mm. Twice Nyquist mode is needed for this correlator mode. The resolution provided at 1.3mm is 0.023 km/s, capable of resolving the thermal line discussed above.

The ALMA twice Nyquist mode has yet to be commissioned, though as noted above users have requested the resolution it provides. According to Rafael Hiriart, twice Nyquist mode was given a low priority in the ICPM3 meeting. It was not requested for Cycle 3. He estimates it would take one month of time. Another high priority upgrade, which would affect all projects, is 3-bit quantization correction, which is on the schedule for Cycle 3. This twice Nyquist high resolution mode is not currently an ALMA science priority.

Escoffier, Lacasse and Saez pointed out that advances in digital electronic performance since the correlator was designed should enable higher spectral resolution to be achievable, perhaps up to a factor of 8 or better. Similarly, owing to the faster electronics, faster dump times might be achieved. When the baseline correlator Frequency Division modes were implemented, the increased data volume had to be offset by poorer time resolution.

The first prototype Band 1 receiver is to be built within the next year. At Band 1, .01 km/s need to resolve the line would not be available from the baseline correlator. In fact, the science case above has not been achieved with ALMA at B3.

Faster dump times might be useful for some applications--a case in mind would be the tests of a pulsar mode for ALMA run last summer toward the SgrA* magnetar.