From bsahr Tue Apr 17 15:30:36 2001
To: bclark@aoc.nrao.edu
Subject: Re: Correlator Interfaces
Cc: ghunt, gvanmoor, ksowinski, bwaters, jbenson, wkoski, gpeck
X-Sun-Charset: US-ASCII
Content-Length: 12406
X-Lines: 239
Status: RO

Barry,

Thank you for the comments on snapshot #1 of the EVLA Architecture and Design document. I will try to respond to some of your comments in this email. Others will require time for additional thought.

I plan to cc a copy of this reply to Wayne Koski and George Peck as it contains a number of points concerning monitor and control which may interest them.

Bill

> From bclark@zia.aoc.NRAO.EDU Thu Apr 12 15:38 MDT 2001
> Date: Thu, 12 Apr 2001 15:38:37 -0600 (MDT)
> From: Barry Clark <bclark@aoc.nrao.edu>
> Subject: Re: Correlator Interfaces
>
> [It might be profitable to set up majordomo lists for EVLA software and
> EVLA correlator].

I plan to work on this item sometime this week or next.
>
> Comments on Document of 01mar26. <- EVLA Arch & Design Snapshot #1
>
>
> p. 5. para 3. With Bill's network design sketch, it was unclear to me
> that there was a roll for an antenna control computer. It seems to me
> we need either a switch to let us talk to the devices directly or a
> control computer (also serving as a router), but not both. Incidentally,
> using ethernet as the antenna bus rather harks back to the old CAMAC
> idea, that you have fairly intelligent "crate controllers" which then
> talk to stupid devices on simple wires. (Side remark: I suspect that
> 10T ether will generate much less interference that 100T.)

Inclusion of both a switch and an antenna control computer was based on conversations with James Robnett and others. There may still be merit to the idea. In snapshot #2 I want to explore the logic of doing so and possible alternatives in more detail. This point has nagged at me also and requires more thought.

I suspect 10T will be the choice for the fieldbus, because it will be sufficient to do the job & because of the interference issue.

(BTW - about 13 years ago I spent a year working with CAMAC modules. Nice hardware. Worked as advertised. The parallel to the present scheme had not occurred to me.)

>
> p. 7, para 2. Doing UIs in a window system may be the right decision, but
> is not to be taken lightly. In my opinion, the VLBA screens interface is
> much more robust than, say, java.swing, and infinitely more easy to
> implement stuff in than directly in X. It gets its robustness by being
> very limited indeed, but it is still adequate for strictly control functions.
> It is totally inadequate for any data display job, other than instantaneous
> values of monitor points, but it would be possible to make the decision
> that all data display (including, eg, plots of monitor data) is to be done
> through AIPS++, and to feel no strong urge that the operator and technician
> screens should be the same interface - they are used by different people.

Using AIPS++ for all data display is an interesting idea. It is noted and
will be investigated.

I feel no strong urge that the operator and technician screens should be
the same interface, but I would like a situation in which the entire
suite of UIs can be run from multiple locations - the control building,
from someone's home, at the antenna.

I suspect that current thinking is to use UIs for display & control, but
little thought has been given to separating the two.
>
> p. 8, para 2. A very interesting question is whether antsol, which now
> runs in the Modcomps, should run under AIPS++. Although I can think of
> good arguments both ways, I am inclined to the affirmative. This has
> repercussions on the properties of the "archive".

This one comes under the heading of "I need to think about it at greater
length".
>
> p. 8, para 3. We need a much more powerful scrip language than the current
> JOBSERVE output cards. We should not feel at all inhibited about thinking
> up a good one - writing a translater for the JOBSERVE cards will be easy.
> (It took me less than a week to write the VLA->VLBA translator obs2crd.)

Agreed. JOBSERVE output cards are the 1st target. Much more will be needed.
>
> p. 10. There is something to be said for providing different objects for
> the different levels of things within the Antenna object, rather than just
> inheriting properties. I now include a lengthy digression about one way to
> do things.
>
> We could have an Antenna object, which is essentially the antenna structure
> we currently have in Modcomp global common, with modest changes, eg to
> include an array of model server outputs. There is only one type of
> Antenna object - it can be sufficiently general to support all four antenna
> types.
>
> [Digression in the middle of a digression. The model server concept is a
> horrible construct, which we are driven to by the fact that CALC is a nice
> piece of software, which of Goddard (understandably) won't support a reentrant
> version.]
>
> At suitable intervals, the Antenna object sends a message, probably consisting
> only of its own handle, to an AntennaInterface object. There is one
> AntennaInterface object per antenna, and these objects are different

> depending on the antenna type. The AntennaInterface object maintains an
> image of the state of its antenna, mostly derived from monitor data, but
> for some items, the image must be based on commands previously sent. When
> it receives the message from the Antenna object, the AntennaInterface
> object converts the Antenna object into a desired antenna state, and then
> compares the desired and actual states, and formats commands that convert
> the actual to desired. Commands are sent as messages to an AntennaMC
> object.
>
> The AntennaMC object handles the actual routing of commands to the hardware
> (and thus is different for each type of antenna). The AntennaMC object
> would also provide monitor data, most reasonably by way of a cache which
> it keeps up-to-date on its own schedule. Technician screens converse
> directly with the AntennaMC object; the other two types of objects are
> creatures of the observing system. The AntennaMC object could keep track
> of who is authorized to send commands, so that observing system commands
> are ignored if a technician is at work.
>
> With this system, source changes are very much optimized. When the master
> sequencer notes that a source change will happen shortly, it
> 1.) Clones the Subarray object and loads the new source information into it.
> 2.) Clones the Antenna objects attached to the subarray. (the antenna
>   (objects know not to send messages to their AntennaInterface because
>   the subarray is marked inactive).
> 3.) The antenna objects are brought up-to-date (CALC etc)
> 4.) At the time of source change, the old Subarray object is marked inactive
>   and the new one marked active.
> 5.) At leisure, the sequencer checks whether the next source change is back
>   to the previous setup; if not, it recycles the surplus Subarray and
>   Antenna objects and starts over again.
>
> [I don't regard the antenna as "having" devices; it has a "state", which is
> comprised of the state of its devices, and a set of rules for changing the
> "state".]
>
> The AntennaInterface object probably also has to send information to the
> correlator.
>
> An interesting question is whether the AntennaMC object needs to know anything
> about the data it is transmitting (eg units, data length) or not.
>
> End of digression.

Liked your digression very much. It is in advance of my own thinking in this area. I view my present responsibility as being the development of an overall view of the software architecture and design - broad brushstrokes covering a multiplicity of areas - the EVLA network, basics of monitor and control, approaches
to UIs, correlator interfaces, etc. More detailed development of an object structure is an item I have deliberately postponed. I have given some thought to monitor data, and agree that one component of the approach is likely to be an object "close" to the actual antenna with a cache kept up-to-date on its own schedule. Your digression on the object structure will be kept, studied, and considered in detail. The efficiencies introduced by the "read ahead" strategy used by the present VLA software demonstrates the benefits of optimizing
source changes.

>
> p. 16, para 0. Special ports into the equipment (using non-system hardware
> or software) are to be heavily discouraged. Such things caused no end of
> problems in the VLBA, both in the correlator and in tape drive control.

Agreed. Steve B was very aware of the difficulties caused by special ports
in the VLBA and gave strong emphasis to this point. I agreed with him
then and still agreee with you now.
>
> p. 17, para 4. If ethernet is the antenna bus, I rather think that there
> would be fairly few devices; eg, a single "crate controller" might control
> all frontends, sending stuff (mostly switches) on a wire per frontend,
> bringing analogs back to the crate for conversion. An ethernet stack
> requires a lot of real estate that a frontend designer might hate to build
> in.

Yes. If ethernet & TCP/IP are used as the fieldbus, it will "move" the
M&C interfaces to the level of significant subsystems that can afford the
overhead of the needed protocol stack.
>
> p. 18, para 1. Technician screens and test stuff should be written to run
> on a system element - either the antenna controller or the control system
> computer - and the technician station should emulate it's behavior. Easier
> than trying to write both at once, and should work fine. Probably the
> antenna controller is the best place - it is a big advantage to be able to
> talk to antennas when the control system is down, but a rather small one to
> be able to talk to it if the antenna controller is down.

I need to make this point with greater clarity in the document. So far,
the logic of the situation seems to me to converge on placing test stuff
in the antenna controller.

>
> p. 18, para 3. I believe the figure is more nearly 250 times the current
> output rate. (I just saw an IBM ad boasting that their delivered base of
> mass storage devices aggregated to 7 petabytes.)

I'll re-examine my figure of X25 as opposed to X250.
>
> p. 19, para 3. The widefield mapping remark is not quite right. The
> number of channels required to avoid delay smearing is independent of
> frequency, and is about the length of the longest baseline divided by
> dish diameter times the fractional bandwidth of the front end, times
> a few (depending on how negligible you want to make the smearing).
> The systems with the largest fractional bandwidth, are, I think, L band
> and X band (or is K greater than X?). L band is likely to be the worst
> problem - you want lots of channels for post facto interference excision.

I need to speak with Ken & others in order to better understand this issue.
>
> p. 20. Projecting average data rates is a very difficult business. I'd
> say current continuum data rates will go up by a factor of 30, to handle
> the extra bandwidth. Data rates for the current sort of line experiments
> will go up by a factor of maybe 8, which they really need, or 16, which
> might be nice. Bottom line of that is about 10 TB per year, way less than
> your numbers. The great unknown is to what extent the increased capabilites

> attract harder observations. Jacqueline van Gorkom has discussed an
> observation that currently looks very attractive, of observing at L band
> a deep field, with maybe 60,000 channels, full polarization, 5 second
> integrations for a few tens of days. Whether it is still scientifically
> attractive ten years from now when we might actually do it is far from clear.
> (ALMA is costing things on the basis of average data rate is .1 maximum,
> just for reference.)

Yes. I think the utilization model is way off. For that matter, the entire approach taken in the final paragraphs of page 20 may be a poor way to arrive at an estimate. If I apply the ALMA figure of .1 of maximum to the max initial data rate of 25Mbytes/sec & 612 observing hours per month, the result is ~ 63 TB/yr, still way over your figure of 10 TB/yr.

I too wonder what effect increased capabilities will have on the type of observations done.
>
Again, Barry, thank you for the response. Your comments go into the growing pile on my desk of material to be included in future versions of the snapshots.

Bill