Hi Bill,

I was just talking w/ Fritz Stauffer and he mentioned talking to you about
correlators & Beowulfs. I will be out in Soc. next week for an ALMA meeting
on Mon. & Tues. and I'd be happy to discuss the little bit that I know
about Beowulfs.

For the processing of lags from the ALMA correlator, I'm planning on using
a Beowulf cluster. The correlator LTAs will use FPDPs, which can deliver
100 MB/s each, do send the lags out w/ 4 per quadrant for a total of 16 for
the final 64 antenna system. My idea is to have one compute node per FPDP
w/ each PC having PCI-to-PMC adapter which can hold a PMC FPDP board, an
example of these 2 boards can be found at

http://www.vmetro.com/adapters/pmc-into-pci-adapter.htm and
http://www.transtech-dsp.co.uk/io/pmc-fpdp.htm

The compute node would be basically the most powerful PC that I can get.
I'm looking at 1+ GHz Pentium III or AMD Athlon K-7. For example, Gateway
gave me a quickie quote of an 1 GHz Xeon P3 for ~$1200. Although there are
some concerns about memory bus bandwidth w/
PC's PCI bus (100 - 133 MHz may not be fast enough).

For more information about Beowulfs see http://www.beowulf.org/.

One thing that I haven't come across is the use of Beowulf clusters in
embedded systems like what we're interested in. As far as I can tell,
this is new territory. There is another engineer here in C'ville,
Richard Bradely, who is also interested in developing an embedded Beowulf
cluster for some work that he's doing w/ the GBT. Finally, I think the
folks at DRAO are looking into using a Beowulf for their ACSIS correlator.
Since you're already working w/ Brent Carlson, you might want to ask him
about the computer system -- I do have the diagrams that you sent Chuck
the other day & it looks like they're using an embedded CPUs to do the FFTs.

I'd like to meet w/ you in Soc. next week. I'll probably be in the
ALMA meetings most of the day, but I should have a little time to talk.

Cheers,

-Jim