From: OUTBAX::VAX3::RPERLEY 10-MAY-1988 14:39
To: ABRIDLE
Subject: Use of the Toronto Cray.

From Thursday, April 28 to Friday, May 6, I was in Toronto using the Cray-XMP to analyze VLA data. Robert Braun joined me from April 30 to May 5. The Cray is run the the Ontario Centre for Large Scale Computing (CLSC). NRAO has been given an account by the CLSC of 100 hours of computing.

Our goals were:

(1) To deconvolve, using `Maximum Entropy' methods (i.e., VM or VTESS), three 4K x 4K images of Cas A, representing the remnant in 1983, 1985 and 1987.

(2) To attempt self-calibration on the data bases, should the first goal be met sufficiently early.

We took 20 tapes to Toronto, containing FITS images of all the dirty maps and beams (Stokes `I', `Q', `U'), with blocking factor = 10, for each of the three epochs, and the 'XY' sorted databases. These had to be in VAX BACKUP format, since each of the three databases contained about 350 MBytes, about 3 times too large for FITS (which cannot span tapes).

Upon my arrival, a meeting was held with Phil Kronberg, his two programmers (Patricia Monger and Laura Carriere), and the staff of the CLSC The expression on the faces of the CLSC staff upon being informed that we wanted to load a single 350 MByte file onto the Cray was worthy of permanent record. The difficulties they quoted were sufficient to relegate this job to lower priority. It was reported that the status of AIPS was uncertain, but that WERONG, the package produced by Bob Sault should work fine. I then decided that the first thing we would do would be to attempt a deconvolution using this package, then using the AIPS VTESS task. I thought it would be interesting to compare the results.

WERONG requires disk FITS files on the front-end VAX. We thus copied our tape FITS to disk, but soon found that WERONG could not read them. The problem seemed to have something to do with the image header, and we attempted various actions to produce a file acceptable to WERONG. None of these was successful. We then tried reading the tapes directly onto the Cray, with the idea of copying them to the 750. There were numerous problems here as well with the proper use of tape drives. By the time these were solved, another successful route was found. We did not pursue this route further. The third route tried was to read the tapes onto the 750 through AIPS via IMLOD, then re-translating to the disk FITS through FITTP. This route eventually worked, but not without considerable trouble. When IMLOD reads the first file of a tape, it commands the tape drive to backspace to the BOT. On the CLSC 750, this caused the tape drive to hang, and since there is only one tape drive, and the machine is busy serving a lot of other users, we were `frozen out' for the entire rest of the day (until the powers-that-be decided that rebooting the machine was the only solution). It became clear, after numerous trials, and consultations with Eric and Kerry, that the CLSC tape drive can't, or won't, backspace. The offending line of code in `IMLOD' was removed, and no further problems were encountered.

So a route to producing acceptable images was found, but two full days were used up in finding it. I will also comment here that an overloaded 750 is not a suitable machine for running IMLOD and FITTP for images of such a size. The CLSC will soon be installing an 8600 series machine to replace the 750, so the throughput problem here should be reduced. We never did find out

precisely why the FITS images we brought were unacceptable to WERONG. Clearly this is an important point.

While we were waiting for our images to be translated to useable form, Patricia and I tried to start up AIPS, with the intent of attempting a deconvolution of the images we had loaded directly onto the Cray. This was a complete failure. It appears that certain important files (such as the password file and the message file) were missing, having been backed up on tape. The system is supposed to prompt the operators to load tre required tape, but they claimed no such command wassever received. Patricia was somehow able to fool AIPS into starting, but we then found that VM wouldn't start, probably due to more missing files. By this time, it was late at night on the third day, and our first WERONG deconvolution was starting. We elected to put the AIPS route on the `back burner', until we sorted out the WERONG problems.

After numerous false starts, a fully deconvolved image was produced. This was stunningly successful by most measures. It is important to note that we didn't get it until midnight Monday, over 4 days after I arrived. The various problems in getting the data into useable form, followed by learning how to set up the Cray job (which took numerous attempts) took up that length of time. The first deconvolution also took an enormous amount of computer time - more than 15 hours. However, this was soon found to be due to improper use of the SSD, and Patricia, after consultation with the CLSC staff, soon had this time drastically reduced.

From this point on, the sailing was much smoother. Since our primary goal was to come away with useable deconvolved images, we went into a sort of production mode, loading data onto the 750, translating to useable FITS files, followed by deconvolution by WERONG. We deconvolved four images (the 1987 data were used to make two dirty maps and beams, according to data quality). We also elected to continue deconvolution of the 1985 data, since its convergence was slower than either 1983 and 1987. Various unexpected problems (called `features' by programmers) were encountered, but in the end, successful deconvolutions were made of all maps by Wednesday evening. In fact , the procedure became so smooth by that time that external assistance (i.e., by staff, or by Phil's programmers) was no longer needed.

I have indicated above that the throughput improved dramatically after Patricia had properly `tuned' the Cray. The numbers she attained are: For 24 interations of a 4K x 4K image, about 900 seconds of CPU are required, plus about 1 hour of IO-WAIT time (since the data and scratch files exceed the SSD memory). Total, about 1.5 hours real time on an empty machine. The machine is actually idle, or lightly loaded most of the time, so the actual throughput was as short as 3 hours. The current SSD has 16 MWords (=128 Mbytes) capacity , and is soon to be upgraded to 64 Mwords. This will allow a 4K deconvolution to be fully resident in the SSD, and the experts believe the deconvolution time to drop to 15 - 30 minutes. This increased SSD will not be enough to contain 8192 x 8192 images, but it is felt that the deconvolution time should still be roughly 4 times the current, i.e., perhaps 6 hours. This sounds very attractive to me, since Cas A really needs images of this size (using 4K loses 10 to 20% of our data), and I think everyone agrees that attempting such a deconvolution on our Convex is truly futile.

Since the first goal was not attained until the day before we left, there was no time to pursue the AIPS route. Phil Kronberg has set this as his highest priority, and we have left behind a copy of our dirty map and beam for his staff to use and to compare agains the WERONG result.

An important ingredient in our final success was the cooperation of

the people involved. Patricia Monger, in particular, worked extraordinary hours (to 4 AM on some days) to sort out the problems we encountered. Phil Kronberg ensured the help of the CLSC staff. I must stess most strongly that the CLSC staff were most helpful and interested. A good example is their willingness to give our first deconvolution the highest priority (10.0, when 4.0 is normal for a job of such length), very much to the detriment of all the other users (at the time, the machine was very busy). It was the results of this first job which allowed Patricia to tune up the system. Wherever we went , we were greeted with interest and helpfulness. This attitude certainly was a major contributor to our final success. We are sorry that we could not test AIPS, but given the difficulties, we had to concentrate our energies on just one route. It would seem reasonable to expect that Phil and Patricia will have AIPS running soon. The date will certainly be sooner if someone is interested in going to Toronto, and sets a firm date.

From: CVAX::RBROWN 3-APR-1989 09:10
To: BBURNS,ABRIDLE,MGOSS,GHUNT,EGREISEN,PVANDENB,RBROWN
Subj: Illinois Supercomputer

Dick Crutcher called. He now has an adjunct appointment at NCSA, a personal budget of $100k and a charge from Smarr to get astronomical imaging projects on the Illinois Cray 2. Wants to talk to us about (1) finding a mechanism for us to send them large users and (2) a mechanism for them to handle the work of the large users. Argues that things have changed recently in that the Cray 2 is a UNIX machine with multiple processors (4 I think). He/they would assign us one of the processors for our users full-time, no competition, which should ease one of our earlier concerns. Issue of AIPS remains, is possible now with UNIX, but is this (all of AIPS) the right approach or should users only use the Cray for the real number crunching tasks, either as extracted from AIPS or as extant in WERONG?

Crutcher wants to talk. Phone meeting with a larger group fine. Is this something we want to discuss? Think about it and let's get together in Socorro and decide how to respond.

From: CVAX::ABRIDLE 7-APR-1989 16:38
To: BBURNS,EGREISEN,ABRIDLE
Subj: Tim on Cray-2

I had an interesting discussion with Tim Cornwell about his experience in the Cray-2 at NCSA, where he's been testing some large-memory things with his SDE code. He told me

(a) Tape handling is bad. He's still not been able to send them either source code or data by tape in any form that they can read.

(b) There are two FORTRAN compilers that allegedly compile Fortran-77 but which both fail on various standard constructs. He says you recompile with the other compiler before debugging -- always. And it's not that the compilers blow up - they quietly produce output code that runs incorrectly!

(c) The local help at NCSA is not helpful. Tim says the calibre of people assigned to assist him has been very low.

In other words, "deja vu all over again". Still. And he's relating the status quo there of a week ago, not a year ago. I'm going to try to get Tim and Larry Rudnick together before Larry leaves, so that Larry hears some of this first-hand from Tim before we get into discussions about AIPS in the Cray at Minnesota. I presume that the closeness of Larry's group to Cray itself might help a bit with compiler bugs and that the local help issue might be different for Larry at Minnesota than for Tim at NCSA. But it might not. I think the important thing will be for everyone involved to be heads-up about the possible problems when we first get Larry and Kerry together, maybe some time next week. Perhaps you guys should also talk to Tim about details before deciding whether or not even to proceed with the Minnesota initiative.

P.S. Tim has gotten around the tape problems exclusively by using ftp. A bit daunting for AIPS ?

From: CVAX::ST%"tcornwel@vax3@VAX1.NRAO" 20-APR-1989 15:01
To: ABRIDLE
Subj: You may be interested in this.

Received: from VAX1.NRAO by NRAO.nrao with DECNET ;
Thu, 20 Apr 89 15:00:51 MDT
Received: from vax3 by VAX1.NRAO with DECNET ;
Thu, 20 Apr 89 11:22:20 MDT
Date: Thu, 20 Apr 89 11:21:46 MDT
From: tcornwel @ vax3
Message-Id: <890420112149.000@vax3>
Subject: You may be interested in this.
To: abridle @ vax3
X-ST-Vmsmail-To: ST%"ABRIDLE"

From: VAX3::TCORNWEL 20-APR-1989 11:21
To: RPERLEY, TCORNWEL
Subj: SDE on the CRAY2

I thought that I should bring you up to date on my efforts to port SDE to the CRAY2. As you know, I have decided to abandon the effort. However, I had a spare day at the site so I decided to have one moresattempt at fixing SDE. Dean was going to carry on anyway and it seemed easier to do it myself rather than answering Dean's questions every day or so (By the way, Dean is being paid to do this, presumably by NCSA). I isolated all the routines which were causing problems in compilation. Basically any routine which does Fortran i/o may be mis-compiled with CFT77, so that the old compiler CFT must be used. In the end, only two packages of routines needed to be compiled with CFT: the user interface and the text i/o routines. Everything else seems to compile ok now with CFT77.

All tasks are now working with the exception of the plotting routines, and a few tasks which have an illegal FORTRAN construct (an internal free-format read/write). In particular UVMAP, VM and CCLEAN now work and can be compared with the Convex C-1 versions. The results are:

	Convex	Cray2
CCLEAN	51.87	11.27
VM	16.81	3.18
UVMAP	26.43	3.82

(CCLEAN is my version of APCLN, VM is like the old AIPS VM, and UVMAP is functionally equivalent to the AIPS UVMAP)

The CRAY2 FFT routine used is their hand-coded routine which takes 0.49 msec/1024 complex FFT versus about 4msec for the Convex handcoded routine.

These numbers are for small jobs, and may be unduly weighted by i/o, but nevertheless the times are disappointing. The real-time response is not noticeably better than that on the CONVEX when the load factor is less than 8, say. We should let Dean have a go at optimizing parts of the code. We may gain something there.

The main work in porting to the CRAY-2 can be summarised:

- Getting the code on the system

- Eliminating names of more than 8 characters
- Getting acceptable screen-editor on system (MicroEmacs!)
- Reworking parts of c/fortran interface
- Discovering CFT/CFT77 bugs (undoubtedly the most painful part)
- Re-writing FITS reader/writer (word-length assumptions snuck in)
- Re-constructing code which was not saved in a system backup.

I have ported SDE to the CONVEX (from the VAX originally), to the SUN 3/60, and to my COMPAQ 386 running Microway Fortran/C. The port to the CRAY2 was by far the most time-consuming of these, and the least fun.

From: VAX1::MGOSS "Miller Goss" 3-MAY-1989 14:16
To: ABRIDLE,MGOSS
Subj: mike norman

Mike was here last week with Jack Burns. As expected the old issue of collabaration with NCSA came up. I re-iterated that any AIPS issues between Norman - Crutcher should go through the AIPS group in CV. Cornwell and Perley are starting a 3-d clean etc project at NCSA and they have made a small begining. We will see - I am rather sceptical that any large scale project will work out given the past history. But we should keep an open mind. Miller

From: CVAX::ABRIDLE 15-JAN-1990 16:05
To: VAX1::VAX3::RPERLEY,ABRIDLE
Subj: RE: Not for AMAG

I agree with you that I don't see why NRAO supercomputer access should be coupled to AIPS alone. I think it makes sense to run lean, mean standalone codes for the most demanding applications and ship the results back to AIPS for lower-level processing. I think that AIPS should be NRAO's only general purpose package for reduction, but not all data reduction should be done by the general purpose package. I have been at odds with first Eric and then Bill about this. Bill regards the entire supercomputer access project as a waste of time, which is another reason for detaching it from AIPS. I think the supercomputer access probably can and should grow out of Tim Cornwell's algorithm research and codes such as MIRIAD (WERONG is dead so I hear). Then when supercomputers all run UNIX we just slip in a UNIX AIPS where it is needed with minimal extra effort from the AIPS group.

I have talked with Ken K. re the Smarr issue. He has independent evidence of Smarr's exaggeration and political steering activities from the main meeting, and agrees with my suspicions about what is really going on. We felt the next step will be for one of us to draft a page of text that might be the first draft of what the Radio Panel will say about VLA computing as part of the VLA upgrade section of its report. I'll make a draft of that and ship it to you for comment. If the Computer Panel indeed sticks entirely to philosphical issues as Mike Norman said it would, then this approach through the Radio Panel may be all that is needed or desirable. However, I believe that Smarr will not stick to philosophy but will use your panel to promote his own political ends, i.e. attempt to corral VLA/VLBA top-end processing for the Illinois empire. If I am right, further management may be needed within your panel. There seems to be some agreement here that co-operation with supercomputer center*s* PLURAL!! is necessary and important but that monogamy with Herr Smarr is not especially desirable.

From: CVAX::ABRIDLE 16-JAN-1990 12:15
To: RPERLEY,MGOSS,BBURNS,BBROWN,ABRIDLE
Subj: Smarr White Paper & NRAO

Comments on the Smarr Panel "White Paper" as it relates to NRAO computing:

1. p. 4 para.1

It states (ungrammatically) that many of the most exciting and innovative of the possible VLA projects cannot be carried out because of the lack of sufficient *supercomputer* power. There is no clear relationship between need for supercomputer power and innovation or excitement. Some of the supercomputer-demanding projects are merely big. Jacqueline's analysis of extragalactic HI work in our new plan is a good example, showing how the science needs more computing power of all kinds, including interactive small-scale computing. Nobody has argued that the 10% of all projects that need supercomputers are the *best* 10%, just the most troublesome!

p.8 first big para. It is absolute nonsense to say that all VLA snapshot programs require supercomputers. I also think it's dangerous to read about "special problems inherent in the VLA's design" making huge computer resources necessary. It's more a case of advances in the algorithms increasing astronomers' appetite for difficult problems than it is of an intrinsically faulty design!

p.8 second big para. Here's the potential disaster area. We are characterized as having a two-pronged approach: portable software and supercomputer access (not surprisingly, there are both things that Smarr wants to take over!) The main thrust of NRAO's plan is distributed computing at the small supercomputer level and below, down to workstations, shared between the NRAO and its user base at their home institutions. The Smarr document makes no reference whatsoever to the need for a significant (i.e. similar to NCSA's present capacity) computational center at the NRAO. This *must* be corrected before this document gets far outside the panel (it's already been distributed to the chairs of the other panels, maybe more!).

p.8 last para. Notice that the scenario described here for "supermaps" whatever they are, leaves no role at all for the NRAO. It's an interaction entirely between NCSA and the users. This approach will leave the AOC freezing in the dark. It is an example of the basic conflict between the NRAO's direction and Smarr's. Smarr sees the telescopes as a bit supply for his supersupercenter and we see his computers as a cycle shop to be used when needed. He wants to bind the users to the computing center. We, over 30 years, have seen the benefits of close interaction between the NRAO and its users. The form of operation that is suggested here, unless strictly confined to a few truly CPU-intensive projects, could eventually undo the partnership between the NRAO and its users that has made the radio astronomy community in the U.S.A. so successful. I asked these super-hypers what they meant by a "supermap". They mean an ordinary 4k continuum map! Read this and shiver -- it's the thin end of a wedge that leaves the NRAO sitting on a small pile of workstations and the lonely end of a fiber optic cable, while our users are interacting with "contact people" at thi NCSA!! This mode, if not confined to the 10% of all projects that we wish to limit it to, will make the NCSA the "national observatory", not the NRAO. From my conversation with Crutcher, I have little doubt that this is what is intended!

p.14 The reasons for difficulty in porting AIPS to supercomputers are interesting to explore, which is why they don't elaborate on them and instead

leave the impression that there's some basic deficiency in NRAO's coding standard. The problems arise because AIPS asks for efficient use of the resources of the whole computer, which the center bureaucracies are reluctant to provide. AIPS has *not* been in the least difficult to port into supercomputers. It *has* been difficult to port into supercomputer *centers*, with their inflexible rules for the use of resources such as SSD's, tapes, disk residence and archiving systems. Also, our effort in this direction has been limited because the *real-time* performance in supercomputer centers has been poor. Mike Norman admitted to me that this was also true at the NCSA and that they needed to do something about it. The production user does not care about how little CPU time went into their computation, they care about how many ticks on the wall clock elapsed. Supercomputer centers have yet to provide good wall-clock performance in our applications except during "publicity stunts" for which they have harnessed their resources in ways that promote their own aggrandizement! When it comes to routinely supporting production by users, their time-shared performance is unremarkable, and that is why the NRAO has been reluctant to move a big package like AIPS into their environments. This point needs to be made to the panel very forcefully!

p.32 To label this case a "typical" VLA data processing request is absurd. As the last para on this page states, the imaging needs are an *exceptional* case! Note that the NRAO machine proposed in the Array telescope computing plan addendum will comfortably handle all but the A array case at 327 MHz; it is a several-Gflop machine with 4 GBytes of memory!

Throughout: these guys should use "hyper" more than "super" because that's what some of this is -- HYPE. What they are disguising is that the evolution of the chip is steadily putting more elements in the same volume at reduced cost so the performance of affordable decentralized computing has a doubling time of about 18 months *at all levels* -- from workstations to what we used to call supercomputers. This means that in the hands of the companies who build user-friendly, industry-standard machines, *the computer power is redistributing itself* closer to the users about as fast as we build networks into the national centers. We are therefore not obliged to place ourselves in thrall to the directors of a few centers, just as we were not obliged to do this on the campus in the 1970's when the Vaxen appeared. What the "hypers" also leave out is the human element, of who the users interact with and why, the value of a radio astronomy "community" interacting with what it sees as *its* national observatory. The NCSA in particular is worried about the consequences of becoming a "cycle shop", servicing only the high-CPU needs of many communities. This is of course precisely the role we want for them if we are to retain our closeness to our own user community! All this talk about "supertelescopes" and "supermaps" and "supernodes" is an attempt to hype a trend that we wish to keep the brakes on -- a shift of emphasis away from the telescopes, which are controlled by the community that built them, toward remote national computing centers, which serve multiple communities as impersonally and inefficiently as the old campus centers used to. The centers want us to be bit suppliers. We want them to be cycle shops.

There is a fundamental conflict here. For the health of the radio community and of its relationship to the national observatory, we *must not* fall victims to these "hypers"!

From: CVAX::ABRIDLE 17-JAN-1990 11:46
To: RPERLEY,MGOSS,ABRIDLE
Subj: CrAIPS

Just to reinforce my point re the misrepresentation of AIPS portability to supercomputers in the current Smarr Panel White Paper - Kerry has returned from Minnesota with excellent results from installing AIPS and the DDT suite in the machines at Cray Research and U.Minn. He went for a week, was up and running at Cray in about half that time, and so had time to get the package going at U.Minn. also. This is under UNICOS. Some details remain to be sorted out re IEEE floating point conversions and the TV tasks, but the point is that this transport took only a few days of work. This emphasizes that the "problems" of AIPS in supercomputers are problems with management of centers, rather than of intrinsic portability. We are close to the point where installation of AIPS under UNICOS willmbe no more of a problem that installation on an Alliant or some other UNIX machine that NRAO does not have in-house.

This of course supports what Kerry has been saying all along, and squares with the fact that he was able to get AIPS going at NRL under COS in only about a week, whereas it took months at the Pittsburgh Supercomputer Center.

From: CVAX::BBURNS 18-JAN-1990 10:18
To: ABRIDLE,RPERLEY,BBROWN,PVANDENB
Subj: aips on cray

From: EXOS 17-JAN-1990 09:01
To: BBURNS
Subj: [EXOS Mail From: "Larry Rudnick" <larry@ast1.spa.umn.edu>] Cray - I don't know whether Bill forwarded this to you...

Return-path: <larry@ast1.spa.umn.edu>
Received: from cv3 (NRAO.EDU) by cvax.CV.NRAO.EDU
id 00003526002 ; Wed, 17 Jan 90 08:59:42 EDT
Return-Path: <larry@ast1.spa.umn.edu>
Received: from uc.msc.umn.edu by NRAO.EDU (3.2/DCW-2n)
id AA22818; Wed, 17 Jan 90 09:03:17 EST
Received: from ast1.spa.umn.edu by uc.msc.umn.edu (5.59/1.14)
id AA08835; Wed, 17 Jan 90 08:01:25 CST
Date: Wed, 17 Jan 90 07:26:05 CST
Message-Id: <9001171326.AA12523@ast1.spa.umn.edu>
Received: by ast1.spa.umn.edu; Wed, 17 Jan 90 07:26:05 CST
To: bburns@nrao.edu

Thanks also for your part in supporting this project.

From: "Larry Rudnick" <larry>
Message-Id: <9001151312.AA07990@ast1.spa.umn.edu>
Received: by ast1.spa.umn.edu; Mon, 15 Jan 90 07:12:55 CST
To: bcotton@nrao.edu
Subject: CrayAips
Cc: khilldru@nrao.edu, larry
Status: R

Dear Bill :

Please wait at least a short time before sending Kerry back to Minnesota so that I can catch up on my sleep!! My first (semi) conscious act this morning was to melt our glass coffee pot by putting the wrong stove element on.

Seriously, it was a most productive visit, due to Kerry's expertise and extreme hard work. 15APR90 is now installed on our XMP in a useable form, with speeds that Kerry will show you (looks like about 200x faster than our Sun 3's, for some programs!). These speeds and the major accomplishment of the basic virtual tv functions, (commemorated with some pictures which we'll get soon) make this a great system to start working on.

When I'm 3/4 coherent, I will send Kerry a note summarizing my understanding of where the project is, what major problems need to be solved, and which things I can take care of or should try. I'll send you a copy.

Thanks for supporting this project. It will make a big difference in the amount of work that we can do, and hopefully will provide NRAO a springboard for other super platforms.

lr

Kerry -
Here is my summary of where I think things stand, and what has to be done. Please add/delete/modify this list, and let me know what I/we should be doing.

KNOWN PROBLEMS/UNFINISHED BUSINESS

FITTP - works only with 32bit integers, ASCII tables (not AN)
real TAPE - not yet implemented
virtual TV - (almost) no interactive functions at present
plotter output (no implementation for Apple LW?)
TPH dead

NOT YET TRIED

TKPL
lots of tasks - (Dickey will be trying some spectral line cal stuff
as soon as he can gear up)

TO BE DONE

clean up disk areas on /m/h2/mg24501 15APR88,kerry,MSCI? - KH
tape backup of 15APR90 on XMP - (tars made, no tape yet) - LR
user comments on installation guide - LR
info. to LR on modifying gridding for MX tests - KH
manual update lesson - KH/LR
disk usage (alternate storage, recompiles thru ZSTRTA?,

elimination of non-XMP needed code,...) - KH/LR

[What a job! - Thanks again for your support, Bob!]

From: CVAX::BBURNS 29-JAN-1990 12:08
To: ABRIDLE
Subj: cray rudnick - fyi

From: EXOS 29-JAN-1990 12:06
To: PVANDENB, BCOTTON, BBURNS
Subj: [EXOS Mail From: "Larry Rudnick" <larry@ast1.spa.umn.edu>]

Return-path: <larry@ast1.spa.umn.edu>
Received: from cv3 (NRAO.EDU) by cvax.CV.NRAO.EDU
id 00003EA2002 ; Mon, 29 Jan 90 12:04:54 EDT
Return-Path: <larry@ast1.spa.umn.edu>
Received: from uc.msc.umn.edu by NRAO.EDU (3.2/DCW-2n)
id AA28045; Mon, 29 Jan 90 12:05:12 EST
Received: from ast1.spa.umn.edu by uc.msc.umn.edu (5.59/1.14)
id AA03958; Mon, 29 Jan 90 11:04:45 CST
Date: Mon, 29 Jan 90 11:03:53 CST
Message-Id: <9001291703.AA06330@ast1.spa.umn.edu>
Received: by ast1.spa.umn.edu; Mon, 29 Jan 90 11:03:53 CST
To: pvandenb@nrao.edu
Cc: bburns@nrao.edu, bcotton@nrao.edu, khilldru@nrao.edu

If appropriate, we would like to include in the the next NRAO newsletter some note about the fact that we are now processing data on the Cray X-MP at Minnesota. Could you please pass this message on to whomever is in charge so I can get them a little note before your next deadline?
thanks.
lr

From: VAX1::VAX3::RPERLEY "Rick Perley" 8-FEB-1990 12:30
To: ABRIDLE
CC: RPERLEY
Subj: White Paper letter

COMMENTS ON THE FIRST DRAFT OF THE COMPUTER PANEL REPORT

This draft is, of course, too long and disorganized. But this is quite natural for a first draft which combines contributions from such a diverse group, and I expect the next draft to be much more concise.

A more serious concern is that the draft appears, upon a cursory read (which is what many people will give it), to be mainly a justification of the existence and necessity of supercomputer centers. It seems to say that all that is needed for future computing will be either on your desk, or at a supercomputer center. Any computing task too big for the former will be handled in some behemoth center, through high speed links. There is no mention in the paper of intermediate-scale solutions to computing needs.

I disagree with this emphasis. In my view (and this is a view held within the NRAO), computing will become hierarchical, starting with the powerful workstation on the scientist's desk. This workstation, and many others of whatever organization the scientist is a part of, will be served from a larger `super-minicomputer', which has considerably greater computing power. The great majority of data calibration, imaging, and modeling tasks will be handled by this local hierarchy.
There will remain a few tasks which will exceed the capacity of this grouping, and these we agree must be handled at a national or regional center. The division between local and national computing is not simple to define, for it will depend on many factors, including the available local computing environment and capacity, size of the computing request, degree of support, and perhaps most importantly, nd the patience and personality of the resarcher. Further, the dividing line between local and regional/national will surely be a function of time, as the capacity of affordable machines continues to increase. In summary, I would say that the draft must emphasize the role of local/regional computing, beyond what sits on the researcher's desk.

Allow me to state the preceding point in another way. The `White Paper' uses `super' to modify practically every noun in the report. But there is no reason to believe that the extremely large data reduction and imaging needs which can only be addressed (currently) on supercomputers represents the best science. Super-jobs do not necessarily equal super-science.
As in other fields, some will be better than others. What distinguishes this class is their size -- these problems are merely big. Now the good science will continue to come from data reduction/imaging jobs of all sizes. The great majority of computing jobs in observational astronomy are of the small and medium class. These also represent the majority of good science. And the needs of this majority must be specifically stated in the report. They should not be ignored simply because they are not `super'. The needs of this majority can be covered by local and regional computing. This fits into the hierarchicial scheme I have mentioned above.

Speaking as the representative of a national observatory, I will bring up another clear shortcoming of the `White Paper'. There is virtually no mention of the role of the observatories in computing. Conversely, there is altogether too much emphasis on national supercomputing centers.
The majority of data reduction in radio astronomy

has historically, and quite successfully, been done on the observatory's computers, or at the home institute of the researcher.
I feel that this balance must continue. The central reasons are quality and innovation.

First, quality: We must have fast, local return on our product. By this, I mean that the observatory must be able to see the end product of its instruments quickly and efficiently. It is not sufficient to merely image a calibrator source here and there, for the kinds of problems we must react to are often subtle. These are complicated instruments, and they produce complicated problems. If we lose the tight coupling between data flow and image making, the productivity of the instrument and of the science is guaranteed to decline. The best way, --the only way, -- to keep the coupling tight is to have top-notch staff, and top-notch users (visitors) coming to the observatory and interacting with the data on short timescales. This coupling will not occur if all users stay at home, observing remotely, and it cannot happen if the observatory loses the ability to image quickly (i.e., having less than state-of-the-art computing).

Second- innovation: Imaging quality is very much a factor of algorithm development. And I claim that virtually all progress in this subject has been made at the observatories. The fantastic progress with deconvolution (CLEAN, MEM, Mosaicing) and self-calibration, which has been so instrumental in redefining what we expect from our telescopes, has been accomplished by researchers at this (NRAO) and other observatories and universities. The role of these organizations in the continued development of algorithm research must be emphasized. Continued healthy development of our techniques requires resources to support individuals and working groups in the observatory and university environment, and top-of-the-line computing facilities. This does not mean supercomputers, but rather modern powerful machines of the `super-minicomputer' class. These types of machines fit into the hierarchical approach I mentioned above. Put another way, I will claim that this kind of development cannot be efficiently done at a centralized computer center which is set up to serve a wide range of users and subjects. Ours is a specialized subject, serving a particularly class of instrument. Continued development requires continued tight coupling between the instrument and its engineers and scientists, and the imaging experts whose use the data produced by the instrument. These classes of people must interact personally on a daily basis. Any separation in space and time is guaranteed to hurt the product.

I strongly feel that endorsement of this approach be specifically stated in the report.

I have annotated the White Paper with my detailed comments.

From: CVAX::ABRIDLE 8-FEB-1990 14:58
To: RPERLEY,ABRIDLE
Subj: Comments on your draft letter to Smarr

Here are my comments on your draft to Smarr. The overall form and emphasis is good, and the letter's strongest point is the most important one. My comments are about details, but some of the details are important and should not be glossed over. You may be addressing some of them in the "red ink". But I would recommend putting them also in the letter, as Smarr has a habit of quoting only from the sources that suit him and what you put in the "red ink" only may be conveniently overlooked. I would also encourage sending copies of the letter to the other committee members "FYI", whether or not Larry has asked for this.

A. Comments on what's alreody in your letter.
~~~~~~

1. You should specify what you mean by the terms "local" and "regional" as used throughout paragraph 2. Any vagueness here open to misinterpretation.

2. Computer nomenclature - "super-minicomputer" is a late-1970's term that means "Vax-class". "Mini-supercomputer" is a 1980's term for "Convex-class". Both are rotten terms that combine opposites (having "mini" and "super" in the same word is a bit like having matter and antimatter in the same jar). Like matter/antimatter pairs, these terms are also both vanishing. The 1990 NRAO plan points out that the "mini-supercomputers" now overlap the "supercomputers" in performance. The industry now calls C-2 and C-3 class machines "small supercomputers". This is a better term. It allows the modification "large supercomputers" to be used for what you call "behemoths". (I love *that* term, it implies lumbering giants and the trampling of innocent victims, the right image for the bureaucracies that surround these centers!). But in any case, please replace "super-minicomputer", which you *definitely* don't mean, with "small supercomputer", or, if you insist, "mini-supercomputer".

3. Why be so vague in paragraph 2 about fractions of workload that should be done where? Why not give the NRAO plan as an example? 10% of the project load of radio imaging is considered appropriate for large supercomputers. The rest (90%) can and should be done in an equal partnership at the observatories and at university centers in smaller machines. Larry is fond of quoting our statistics when they suit him. Encourage him to quote our side of them, too!

4. In para.4, I would spell out more on the reasons for wanting close interactions between the observatory staff and our users. Point out the human contact factor -- the partnership between the NRAO and NAIC and their users has been an important factor in the success of the U.S. national system for radio astronomy. (It isn't well matched by the optical community, and they know it!) Personal contacts build relationships and collaborations and *community* in ways that can't be emulated within a universe of E-mail and file transfer protocols. This sense of community is just as important in the long run as rapid feedback from the images to the telescope (it addresses a different, but equally important, problem: that of whether the users feel the observatory is "their" observatory too!). Larry is correctly promoting networking and high-speed data interchange as a "brave new world". He
~~~~~~

could actually solve your turn-around problem in the long term, but he can't ever solve the human interaction factor. (In fact, he wishes to pervert it into an interaction with him rather than an interaction with us). The computer panel report will carry more weight if it acknowledges that not all the problems are asociated with bit flow and computing turn-around time -- that there are sound "community" reasons for distributing the distributable part of the workload. This is worth a sentence or two. Ask him for something on this, use mine or yours.

5. "Innovation" paragraph. By "all progress ... has been made at the observatories" you mean "national and *university* observatories", not just NRAO-like organizations. Say so explicitly. Statements like the one you've actually got could make people from, e.g., CalTech very angry!

B. Comments on what's not in your letter yet.
~~~~~~

1. Suggest that he should refer to the entire NRAO plan as an example of a balanced, distributed approach. He's lifting all kinds of details out of it where they suit him (often perverting the context in a subtle way to his own advantage). Let's have a crack at making him pay for that! Ask him to say "The NRAO array telescope computer plan is a good model of a balanced approach to distributed, hierarchical computing. It would exploit the unique merits of workstations, small supercomputers ans large supercomputers in a variety of environments linked by high-speed networks to solve a wide range of computing problems in radio astronomical imaging. At each level, the computing power available would be well matched to the scientific needs of different types of projects, in proportion to the number of such projects." He doesn't have to priorize our proposal relative to any other one to say that. Remember that he *wants* to say something like that about the *supercomputer access only*. Let's turn a little of that pressure back on him!

2. You must mention somewhere that he should delete the ludicrous statement that all VLA snapshot programs need to be reduced in supercomputers. If only because it will make the computer panel look silly if it stays in.

3. If he's going to keep the statements about AIPS being hard to port into supercomputers, he should say why. (AIPS went into the Crays at NRL, Cray Research and Minnesota like a pig into slop, just a few days' work, so the statement is flat wrong as it stands). The only "problem" is that AIPS asks for use of the computer's resources in ways that center bureacracies are reluctant to grant. It's got nothing to do with portability or coding standards, a lot to do with the rigidity of the management of what you call "behemoths", with the philosophy that you can never waste the computer's time but can waste the users' time, etc..... Given the choice of explaining it, he'll probably decide to drop it, i.e. not take a gratuitous and inaccurate sideswipe at AIPS. I'd settle for him dropping this issue, but I'd prefer him to explain it properly!
~~~~~~

From: VAX1::VAX3::RPERLEY "Rick Perley" 8-FEB-1990 17:28
To: ABRIDLE
CC: RPERLEY
Subj: Changes made

Thanks for your comments. I have:

1) Defined local (as the computer network within walking distance), and dropped regional. I kept national.

2) Changed the nomenclature to the way you suggested. No more super-minis, or mini-supers. Just small supercomputers and behemonths.

3) I specifically referred to the 10%/90% as given inthe plan, and also stated that the 90% doesn't HAVE to be done on a national supercomputer. Some can be done on a lightly loaded little machine (especially snapshots).

4) Expanded the close interactions section -- added a new paragraph.

5) Corrected the unintended slight of the university observatories.

6) Directly suggested that the NRAO plan be cited as an example of distributed computing.

I did not:

1) Specifically refer to Larry's misunderstanding of what snapshots have to be reduced on. This correction has been made in red ink on the white paper. I will ensure that this stupid statement does not get into the final report. Remember that the panel meets again next month.

2) Make any comment about porting AIPS. AGain, I have addressed this quite fully in red ink on the white paper.

I wanted to address the more global issues and lackings of the white paper in this letter. These last two issues I believe can be handled within the red-ink method, and at the next panel meeting.

I'll send you a copy of the final letter.\

I'll also send copies of the letter to the computer panel members. The problem is that I don't know who they all are! Larry has not circulated the membership. I have requested the list from Radha (Larry's research student who is doing all the coordination).

Thanks again. This little job is now done. I shall do science tomorrow!

From: CVAX: BBURNS 9-FEB-1990 11:34
To: ABRIDLE RPERLEY,BBURNS
Subj: smarr

I agree with both your comments on Smarr's report

Maybe we should put Larry on the defensive instead of us. Although supercomputing centers are wonderful things for the world there are some things that they don't do very well, although this may be a surprise to Larry. Because they support hundreds of users they do not have the flexibility necessary for the kinds of things necessary to make them attractive for our needs. For example you can not shut them off at will because you want to re IPL because the display is giving problems. You can not let a small group decide how they want to handle tape use policy because whatever rules are made must be suitable for a much larger group. There are lots of examples.
The problem is not that programs like AIPS do not sit well at supercomputer centers but that such centers are not well suited for the kind of things that systems like AIPS do. We have used a number of supercomputing centers and are painfully aware of their limitations.
The problem is that Larry is not because he does not understand the needs of observational astronomy. Larry should be reminded in our response that although he has an impressive list of accomplishments, doing observational astronomy is not one of them and he does not speak in this area with the authority that he does in theoretical work.
We, not Larry, are the experts in this particular area and we insist that as chairman he respect our experience and advise, no matter how damaging to the visions of his ego. - Perhaps he should be cautioned as a theoretician that his model of the world conflicts with the data we have observed.

From: CVAX::ABRIDLE 9-FEB-1990 11:54
To: VAX1::VAX3::RPERLEY,ABRIDLE
Subj: RE: Changes made

I'm amused at your remark that Larry didn't give the panel members a list of the other members! Classic "divide and conquer" tactics by a chairman, i.e. how to keep control firmly in the hands of the chair. You really must take the political side of that guy seriously. Keep one hand on your wallet, the other on your pants, and watch the silverware, whenever you deal with him!

From: CVAX::ABRIDLE 28-FEB-1990 16:57
To: BCOTTON,BBURNS,ABRIDLE
Subj: Toronto UNICOS installation

Phil Kronberg called me today. They now have UNICOS up at Toronto and Brian Glendenning is ready to bring up AIPS under it as soon as we could let him have it. I told him that the U.Minn. experiment had been largely successful with the usual residual TV and tape problems, and they are keen to have that version a.s.a.p. Phil emphasized that they don't actually want Kerry at this point, that Brian would be happy and able to do the work there with occasional consults to Kerry if we could let them have the code that Kerry brought back from U.Minn. I also told them there would be a freeze coming on March 15 that was supposed to contain this UNICOS version. Phil said Brian could start sooner, immediately, if the code was available.

Given that brian is going to be working for us (Phil thinks he'll finish in the "late spring" still), I agreed to ask you two "what about letting Brian have the UNICOS version right away?" I presume that Kerry has some housecleaning to do before the March 15 freeze. What's the possibility of him doing that quickly and having the UNICOS code go up to Toronto for Brian to work on it some time next week, for example? It might be useful to have Brian work on this before the release to the other UNICOS sites, especially as Brian will be here later in the summer when we really want to pitch the UNICOS version out to the other centers (and Bob wants to have the CRAIPS workshop, etc. etc).

Please let me know if you think it would be possible, and worth, accelerating the schedule for Toronto given that it's Brian who would be at the other end. I would be in favor of this if it doesn't take Kerry off something too vital in the meanwhile. It might mean thst he has the UNICOS package tidiued up earlier than he otherwise would, but that would free him for other things perhaps.

From: CVAX::BBURNS 29-MAR-1990 17:51
To: RPERLEY,TCORNWELL,ABRIDLE,GHUNT,BBURNS
Subj: draft of nrao newsletter article

Draft for NRAO newsletter 3/29/90

ARRAY COMPUTING - SUPERCOMPUTER OPPORTUNITIES

The NRAO long term plan for VLA and VLBA computing has three distinct components.

1) The NRAO will focus on hardware and software approaches to array telescope data processing that can readly be exported to our users at their home institutions or at regional supercomputer centers;

2) The NRAO will purchase enough "small supercomputers" and imaging workstations to handle the data projected data processing demands of about 40% to 50% of all VLA/VLBA proposals, expecting that a total capacity sufficient to process another 40- 50% of such proposals will also become available to NRAO users at their home institutions;

3) The most cpu-intensive 10% of the projected VLA/VLBA projects will be handled either by processing at supercomputer centers or by being deferred until its processing costs become manageable.

The NRAO's planned use of supercomputers involves both the use of AIPS on these machines and the use of other software. The other software will include comprehensive packages like AIPS, but developed and supported outside the NRAO, or specialized stand alone programs supplied by the NRAO or by others. Examples of the former are the WERONG package (designed by the NCSA for the Cray X-MP and continuum data) and the MIRIAD package (designed by the NCSA for the Cray 2 and spectral-line data). Examples of the latter are the SDE-fly routine for wide field imaging, especially at low frequencies, developed at the NRAO by Tim Cornwell.

In support of AIPS on supercomputers, the NRAO has been working to make AIPS run on Cray hardware under UNICOS. This work has been done in cooperation with Cray Research using their facilities and those provided by University of Minnesota. AIPS currently runs on the CRAY XY-MP under UNICOS with only a few limitations. The system uses the virtual TV with the display facility provided by a SUN workstation. The most serious limitations are lack of interactivity on the display (which can be overcome by running AIPS separately on the workstation) and limited access to tapes. These limitations are considered temporary. The NRAO is distributing a partial UNICOS implementation on the 15APR90 AIPS installation tapes. For information. contact Bob Burns at the NRAO in Charlottesville, bburns@nrao.edu, (804)-296-0229.

The University of Minnesota is currently using the AIPS UNICOS system to do routine processing on a CRAY X-MP. For information, contact Larry Rudnick. He may be contacted at larry@ast1.spa.umn.edu, (612)-624-3396, University of Minnesota, Dept. of Astronomy, 116 Church St. SE, Minneaopolis, MN., 55455. The Minnisota group offers no computer time but can provide useful information on their experience.

The University of Toronto has begun installing the AIPS UNICOS system on their CRAY X-MP and also supports a number of stand alone programs. For

information contact Phil Kronberg. Kronberg may be reached at kronberg@utorphys.bitnet (416)-978-4971, Univ. of Toronto, Dept. of Astronomy, Toronto, ON M5S 1A7, Canada. The Toronto group will not in general provide supercomputer time but would consider some small projects where there would be mutual benefit.

The WERONG, MIRIAD and SDE packages are available now at the NCSA on a "friendly user" basis without the need to apply explicitly for supercomputer time. Those with datasets that cannot be processed without supercomputer power, e.g. 8k by 8k continuum images or spectral line observations that exceed the limits prescribed by the NRAO for reduction on NRAO computers, are invited to contact Richard Crutcher to discuss their program. Crutcher may be contacted at crutcher@atlas.astro.uiuc.edu, (217)-333-9581, National Center for Supercomputing Applications, Beckman Institute for Advanced Science and Technology, Drawer 25, 405 North Mathews Avenue, Urbana, IL 61801.

Several other NSF-funded supercomputer centers have expressed interest in supporting array data processing, and we expect that by summer AIPS will be running at about five Cray sites. We believe that at least one of these will be an NSF funded center that accepts proposals for computer time from the academic community.

The NRAO plans to hold an AIPS supercomputer school/workshop in the fall of this year. This will be held at Cray Research in Minneapolis and will be co-sponsored by Cray Research and the NRAO. Its purpose will be to provide information and training for new AIPS supercomputer sites and to compare experience at existing sites. For information, contact Bob Burns at NRAO.

R. Burns
A. H. Bridle

From: CVAX::BBURNS 30-MAR-1990 09:29
To: ABRIDLE
Subj: tim's reply

From: EXOS%"tcornwel@sparc2.AOC.NRAO.EDU" 29-MAR-1990 19:27
To: BBURNS
Subj: Re: input from crutcher which was used for newsletter

Return-path: <tcornwel@sparc2.aoc.nrao.edu>
Received: from sparc2.aoc.nrao.edu by cvax.CV.NRAO.EDU
id 00001E24002 ; Thu, 29 Mar 90 19:27:35 EDT
Received: by sparc2.aoc.nrao.edu (4.1/SMI-DDN)
id AA09411; Thu, 29 Mar 90 17:35:34 MST
Date: Thu, 29 Mar 90 17:35:34 MST
Message-Id: <9003300035.AA09411@sparc2.aoc.nrao.edu>
To: BBURNS@cvax.CV.NRAO.EDU

Hey Bob,

No one has asked me about friendly-user status for SDE at NCSA. What's going on? I think that Crutcher must have jumped the gun on this one.

From: CVAX::ABRIDLE 30-MAR-1990 09:34
To: CVAX::BBURNS,ABRIDLE
Subj: RE: draft of nrao newsletter article

In para after the three-point plan, last sentence (SDE-fly) should start
"An example of the latter is" not "Examples are .."

From: CVAX::BBURNS 30-MAR-1990 12:57
To: ABRIDLE,BBURNS
Subj: most recent - Tim took SDE out of Illinois announcement

Draft for NRAO newsletter 3/30/90

ARRAY COMPUTING - SUPERCOMPUTER OPPORTUNITIES

The NRAO long term plan for VLA and VLBA computing has three distinct components.

1) The NRAO will focus on hardware and software approaches to array telescope data processing that can readly be exported to our users at their home institutions or at regional supercomputer centers;

2) The NRAO will purchase enough "small supercomputers" and imaging workstations to handle the data projected data processing demands of about 40% to 50% of all VLA/VLBA proposals, expecting that a total capacity sufficient to process another 40- 50% of such proposals will also become available to NRAO users at their home institutions;

3) The most cpu-intensive 10% of the projected VLA/VLBA projects will be handled either by processing at supercomputer centers or by being deferred until its processing costs become manageable.

The NRAO's planned use of supercomputers involves both the use of AIPS on these machines and the use of other software. The other software will include comprehensive packages like AIPS, but developed and supported outside the NRAO, or specialized stand alone programs supplied by the NRAO or by others. Examples of the former are the WERONG package (designed by the NCSA for the Cray X-MP and continuum data) and the MIRIAD package adesigned by the NCSA for the Cray 2 and spectral-line data). An example of the latter is the SDE-fly routine for wide field imaging, especially at low frequencies, developed at the NRAO by Tim Cornwell.

In support of AIPS on supercomputers, the NRAO has been working to make AIPS run on Cray hardware under UNICOS. This work has been done in cooperation with Cray Research using their facilities and those provided by University of Minnesota. AIPS currently runs on the CRAY XY-MP under UNICOS with only a few limitations. The system uses the virtual TV with the display facility provided by a SUN workstation. The most serious limitations are lack of interactivity on the display (which can be overcome by running AIPS separately on the workstation) and limited access to tapes. These limitations are considered temporary. The NRAO is distributing a partial UNICOS implementation on the 15APR90 AIPS installation tapes. For information. contact Bob Burns at the NRAO in Charlottesville, bburns@nrao.edu, (804)-296-0229.

The University of Minnesota is currently using the AIPS UNICOS system to do routine processing on a CRAY X-MP. For information, contact Larry Rudnick. He may be contacted at larry@ast1.spa.umn.edu, (612)-624-3396, University of Minnesota, Dept. of Astronomy, 116 Church St. SE, Minneaopolis, MN., 55455. The Minnisota group offers no computer time but can provide useful information on their experience.

The University of Toronto has begun installing the AIPS UNICOS system on their CRAY X-MP and also supports a number of stand alone programs. For

information contact Phil Kronberg. Kronberg may be reached at kronberg@utorphys.bitnet (416)-978-4971, Univ. of Toronto, Dept. of Astronomy, Toronto, ON M5S 1A7, Canada. The Toronto center can provide help and supercomputer time to users on request.

The WERONG and MIRIAD packages are available now at the NCSA on a "friendly user" basis without the need to apply explicitly for supercomputer time. Those with datasets that cannot be processed without supercomputer power, e.g. 8k by 8k continuum images or spectral line observations that exceed the limits prescribed by the NRAO for reduction on NRAO computers, are invited to contact Richard Crutcher to discuss their program. Crutcher may be contacted at crutcher@atlas.astro.uiuc.edu, (217)-333-9581, National Center for Supercomputing Applications, Beckman Institute for Advanced Science and Technology, Drawer 25, 405 North Mathews Avenue, Urbana, IL 61801.

Several other NSF-funded supercomputer centers have expressed interest in supporting array data processing, and we expect that by summer AIPS will be running at about five Cray sites. We believe that at least one of these will be an NSF funded center that accepts proposals for computer time from the academic community.

The NRAO plans to hold an AIPS supercomputer school/workshop in the fall of this year. This will be held at Cray Research in Minneapolis and will be co-sponsored by Cray Research and the NRAO. Its purpose will be to provide information and training for new AIPS supercomputer sites and to compare experience at existing sites. For information, contact Bob Burns at NRAO.

R. Burns
A. H. Bridle

From: CVAX::ABRIDLE 30-MAR-1990 16:16
To: CVAX::BBURNS,ABRIDLE
Subj: RE: most recent - Tim took SDE out of Illinois announcement

If Tim has actually discussed this with Crutcher, that's fine by me of course!

In the U.Minn. para, I suggest that you say, "For information, contact Larry Rudnick at larry@ ...." rather than the present 2-sentence form, and spell Minnesota Minnesota rather than Minnisota in the last sentence!

From: CVAX::BBURNS 7-JUN-1990 18:14
To: CVAX::ABRIDLE,BBURNS
Subj: RE: Cray sites

I assume MIT and Goddard will not make time available to outsiders. San Diego will only bring AIPS up if the see a pool of potential users - so they are interest. In fact their interest in AIPS is just to attract user demand. I am now looking for an observer who in principle will use reasonable amounts of time at San Diego and who will work with them to get AIPS going. I just found this out today.