AIPS MEMO NO. 44

# Benchmarking AIPS on a VAX-8600 with FPS-120B Array Processor

Donald C. Wells, Gary A. Fickling

National Radio Astronomy Observatory [1]
Edgemont Road, Charlottesville, VA 22903-2475
(804)296-0211, FTS=928-1271, TWX=910-997-0174

19 April 1986

Abstract

Performance measurements are presented for the PFT benchmarking test executing on a DEC VAX-8600 with FPS AP-120B array processor. The results are compared with the VAX-8600 alone, the VAX-780+120B, the Alliant FX/1, and the Convex C-1. Overall, the 8600+120B combination is about 1.4x faster than the 780+120B for a typical AIPS job mix, and about half as fast as the Convex C-1. The Alliant FX/1 has about 85% of the speed of the 8600+120B for this AIPS job mix. Due to 4x faster scalar speed, an 8600+AP can do more scalar timesharing work while keeping the array processor busy.

CONTENTS



[1] NRAO is operated by Associated Universities, Inc., under contract with the U. S. National Science Foundation.

## 1 INTRODUCTION AND ACKNOWLEDGEMENT

During 1985 NRAO used the PFT benchmarking and certification package to measure the performance of a considerable number of computers. Results of these measurements have been presented in the AIPS newsletter and in two memos [2]. A notable hardware combination not represented in the trials was the DEC VAX-8600 plus FPS AP-120B array processor. The first astronomy-site installation of an 8600 plus 120B was at the University of Toronto; the PFT test was executed on their AIPS on 20 November 1985, and the results are presented in this report.

The authors are grateful to their friends at the University of Toronto for their cooperation and assistance in performing this trial. In particular, permission was granted by Prof. Ernie Seaquist, Laura Carriere assisted in the actual performance of the trial, and the VAX system manager was very cooperative in arranging for exclusive use of the machine.

## 2 SOME DETAILS OF THE TRIAL

The production version of AIPS on the Toronto VAX-8600 was 15APR85, exactly as in NRAO's 8600 trial on 29 April 1985, and exactly as in the Convex and Alliant trials on 13 September and 24 November. Therefore, NRAO's PFT trials of 1985 all executed the same version of AIPS, used the same VLA test data, and compared computed images against the same master test images.

The Toronto 8600 had 12.0 MB of memory. Three disks were available; the POPS adverb values used were "INDISK=OUTDISK=1; BADDISK=0", which kept all of the PFT images on disk 1, but allowed tasks to use the other two disks for scratch files. Note that the 8600 trial on 29 April used only <u>one</u> disk, and so scratch files contended with data files for access to the single disk drive.

NRAO was given one hour exclusive use on the Toronto 8600; miscellaneous daemons were still alive (for example, the machine was on a DECNET with four other nodes), but logins were disabled. Effectively, PFT had no competition. The test was run on the AIPS1 terminal of the system, at 4800 baud. Afterward, the accounting and message files were listed and analyzed in the usual fashion.

---

[2] Results are available for the following machines:
<u>VAX-780</u>, <u>780+120B</u>, <u>VAX-8600</u> (AIPS Memo No. 36, "Certification and Benchmarking of AIPS on the VAX-8600", 24 June 1985)
<u>IBM-4341</u>, <u>Masscomp MC-500</u>, <u>ModComp+120B</u> (<u>AIPSLETTER</u>, Vol. V, No. 3, 15 July 1985, p. 4, "Portability Column, CPU/OS Combinations")
<u>Alliant FX/1&</u> <u>FX/n</u>, <u>Convex C-1</u>, <u>Cray X-MP</u> (AIPS Memo No. 38, "Certification and Benchmarking of AIPS on the Convex C-1 and Alliant FX/8", 24 December 1985)

## 3 NOTES FOR THE PFT TRIALS

When examining Table 1 below the reader will want to keep the following facts in mind:

1. The 780 and 8600 AIPS installations (operating under VMS) make use of "asynchronous I/O" [3], whereas this feature was not available in either the C-1 or the FX/1 at the time of the trials.

2. The CPU/Real ratios in Table 1 are all just the CPU time in that table divided by the corresponding real time (note that when an AP is involved, the CPU time includes no contribution from the AP itself except for the handler). The ratios indicate the extent to which a task exhibits un-overlapped I/O or operating system overhead. Some tasks, a good example is MX with pseudo-AP on the 780, manage to almost completely hide heavy I/O activity behind their CPU operations, whereas others, VM and UVMAP for example, have rather low CPU/Real ratios. Some tasks are I/O dominated and exhibit quite low ratios; UVSRT and SUBIM are good examples (UVSRT does a disk merge sort and SUBIM is effectively a file copy operation). I/O dominated tasks like CNVRT, COMB, SUBIM, and UVSRT will gain little if any advantage from vectorization.

3. Tasks APCLN, APRES, ASCAL, MXMAP, MXCLN, UVMAP, and VM are the "AP-tasks" in the PFT trial (i.e., they utilize the AP on the 780+AP and 8600+AP). They vectorize on the C-1 and FX/1.

4. Task ASCAL is the most highly vectorized AIPS task; it is also especially indicative of vector sine/cosine performance. The FPS 120B array processor uses special lookup tables to evaluate sines and cosines; this makes it unusually effective for ASCAL, compared to its peak floating point pipeline capability. In Table 1, the ASCAL entry is notably faster on the 8600+AP than it is on the FX/1.

5. Task VM (maximum entropy image deconvolution) only uses the "AP" for 2-D FFTs; the timings in Table 1 probably under-represent its

[3] The term "asynchronous I/O" refers to DMA (direct memory access) transmission from a disk controller into a double buffer in the address space of the program, with the transmission occurring concurrently with CPU execution in other parts of the address space; this gives optimum real-time and CPU-time performance. Unix systems traditionally offer only "synchronous I/O", in which execution of a program stops while any I/O transmission into the program's address space is in progress; this degrades real-time performance. In addition, Unix systems traditionally use a form of "move-mode I/O" in which all transmissions are made by copying from a buffer in the operating system's space to the buffer in the program's space; this degrades CPU-time performance, even though read-ahead and write-behind in the OS buffer helps real-time performance.

ultimate performance on the vector-register machines.

6. The FX/1 timings in Table 1 are from the November trials, using Alliant's improved I/O system.

7. The reader should always remember an important fact about attached array processors: while the AP is computing the host CPU is free to process scalar programs. The concurrent scalar and vector processing significantly improves the total throughput of the system, especially with a timesharing load combined with vector processing.

8. The Toronto 8600 tested on 20 November was noticeably slower for the scalar tasks (CNVRT, COMB, SUBIM, UVSRT) than was the 8600 tested at DEC's Marlborough lab on 29 April; the differences are modest but appear statistically significant. The authors have no explanation for this.

9. The authors consider the MXCLN timings to be the best single AIPS performance index, because task MX in its cleaning (deconvolution) mode does a little bit of everything (gridding, transforming, cleaning, and UV-subtracting).

## 4 THE AIPS PFT BENCHMARKING RESULTS

The results of the PFT trial appears in Table 1 below, along with comparison measurements made on several other machines: 8600, 780+AP, FX/1, and C-1. The data for the "naked" 8600 and for the 780+AP are taken from AIPS Memo No. 36 (see footnote [2]). For the C-1, the data come from Table 1 of AIPS Memo No. 38, while the FX/1 data [4] are from the 24 November measurements quoted in Table 3 of that memo (on p. 22).

---

[4] The FX/1 tested by NRAO was an FX/8 configuration with only 1 CE active. Alliant's "FX/1" model is a special configuration which can only have one CE, and which has a different arrangement of cache memory from the FX/8. The authors assume that the CPU performance of the two 1-CE configurations is approximately equal, and that I/O performance is mainly determined by the peripherals, rather than the CPU. Therefore, NRAO's "FX/1" measurements should be useful as an <u>approximate</u> guide to the performance of the <u>real</u> FX/1.

Table 1

VAX-8600+FPS-120B Benchmarks, with comparison
to several other machines, all times in seconds

| | 8600 | | 780 + AP | | 8600 + AP | | FX/1 ‹1› | | C1 ‹2› | |
|---|---|---|---|---|---|---|---|---|---|---|
| date: | 29 Apr 85 | | 12 Apr 85 | | 20 Nov 85 | | 24 Nov 85 | | 13 Sept 85 | |
| TASK | REAL | CPU | REAL | CPU | REAL | CPU | REAL | CPU | REAL | CPU |
| AIPS | 5371 | ? | 2306 | ? | 2133 | ? | ‹3› | 131.2 | 845 | 117.5 |
| CNVRT | 4 | 1.7 | 8 | 5.9 | 6 | 1.6 | 4 | 2.3 | 3 | 1.9 |
| COMB | 13 | 4.9 | 32 | 22.7 | 16 | 5.0 | 8 | 6.9 | 6 | 5.4 |
| SUBIM | 13 | 2.1 | 14 | 6.4 | 20 | 1.7 | 3 | 2.8 | 3 | 2.1 |
| UVSRT | 27 | 5.1 | 39 | 20.8 | 28 | 5.5 | 20 | 10.9 | 9 | 6.6 |
| APCLN | 963 | 931.0 | 340 | 144.4‹4› | 205 | 39.3 | 255 | 242.4 | 133 | 122.3 |
| APRES | 35 | 23.7 | 60 | 32.5‹4› | 32 | 9.0 | 21 | 16.8 | 13 | 10.1 |
| ASCAL | 1088 | 1058.0 | 154 | 35.2‹4› | 150 | 9.1 | 395 | 390.6 | 107 | 103.7 |
| MXMAP | 71 | 47.7 | 106 | 47.4‹4› | 70 | 13.2 | 56 | 49.3 | 31 | 27.2 |
| MXCLN | 2340 | 2275.0 | 712 | 228.3‹4› | 522 | 61.0 | 515 | 500.2 | 236 | 221.0 |
| UVMAP | 59 | 39.7 | 81 | 41.5‹4› | 54 | 11.3 | 46 | 37.0 | 23 | 18.3 |
| VM | 116 | 71.6 | 195 | 126.7‹4› | 110 | 33.3 | 81 | 61.9 | 45 | 30.0 |
| ALL‹5› | 4729 | 4460.5 | 1741 | 676.6‹4› | 1213 | 190.0 | 1404 | 1321.1 | 609 | 548.6 |

| | CPU/REAL | CPU/REAL | CPU/REAL | CPU/REAL | CPU/REAL |
|---|---|---|---|---|---|
| AIPS | N/A | N/A | N/A | N/A | N/A |
| CNVRT | .43 | .74 | .27 | .58 | .63 |
| COMB | .38 | .71 | .31 | .86 | .90 |
| SUBIM | .16 | .46 | .09 | .93 | .70 |
| UVSRT | .19 | .53 | .20 | .55 | .73 |
| APCLN | .97 | N/A | N/A | .95 | .92 |
| APRES | .68 | N/A | N/A | .80 | .78 |
| ASCAL | .97 | N/A | N/A | .99 | .97 |
| MXMAP | .67 | N/A | N/A | .88 | .88 |
| MXCLN | .97 | N/A | N/A | .97 | .94 |
| UVMAP | .67 | N/A | N/A | .80 | .80 |
| VM | .62 | N/A | N/A | .76 | .67 |
| ALL‹5› | .94 | .60‹6› | .20‹6› | .94 | .90 |

‹1› Alliant FX/1 with revision "B" CE and 8K file system.
‹2› Convex C-1 executing at 9.5 MHz, with 4-way disk stripe.
‹3› tasks initiated standalone to insure a 1 CE environment; AIPS real time is not meaningful. Compiled "vector-only".
‹4› represents VAX/780 CPU time only; similarly for 8600 for these tasks.
‹5› tasks only (i.e., does not include AIPS itself)
‹6› for non-AP tasks only

# 5 CONCLUSIONS DRAWN FROM THE PFT MEASUREMENTS

The following conclusions can be drawn from these data:

1. I/O-limited tasks such as CNVRT, COMB, SUBIM, and UVSRT are generally only moderately faster on the 8600 than on the 780, probably because the two machines use the same disks. Note that the C-1 is slightly slower than the 8600 in CPU time, but 2-3x faster in real-time, which is partly due to the four-way striped disks, but that cannot be the whole explanation, because the FX/1 is also significantly faster than the 8600 in real time, even though it is slower than the C-1 in CPU speed for scalar work.

2. MXCLN is a good vector-processing performance index. The 8600+AP is 1.4x faster than the 780+AP for this task, and the C-1 is 2.2x faster than the 8600+AP. Of course, of the 522 second duration of this calculation, only 60 seconds are consumed by the 8600 itself; much of the remaining 460 seconds are available for overlapped timesharing performance. This extra scalar capability is a real advantage for the 8600+AP in comparisons with the vector-register computers.

3. ASCAL executes in almost exactly the same real time on both the 780+AP and 8600+AP. This is because ASCAL executes almost completely in the APs, which are identical. Of course, the scalar overhead is about 4x less in the 8600. For ASCAL the C-1 is only 1.4x faster than the 8600+AP, while the FX/1 has only 40% of the speed of the AP in this sine/cosine-dominated calculation. [5]

4. The "ALL" entry of Table 1 is a reasonable approximation of a typical AIPS job mix. For this case, the 8600+AP is 1.4x faster than the 780+AP, and the C-1 is 2.0x faster than the 8600+AP.

5. For the "ALL" case, the Alliant FX/1 has 85% of the speed of the 8600+AP; consideration of the relative pricing of these systems suggests that the Alliant FX/1 offers one of the most attractive combinations of price and performance in the current market.

---

[5] in February 1986 a new implementation of the vector sine/cosine routines considerably improved the speed of ASCAL on the C-1.