

# AIPS Benchmarks for the Silicon Graphics Origin200

Athol Kembal and Chris Flatters

January 29, 1997

## Abstract

We have run *AIPS* benchmarks on the Silicon Graphics Origin200, an entry level, shared-memory multiprocessor system. We achieved a single-user AIPSMark<sup>(93)</sup> of 13.7 and were able to run four simultaneous DDTs (on a four-processor system) with only a slight loss of speed.

## 1 Introduction

The Silicon Graphics Origin line of servers are shared-memory, symmetric multiprocessing computers based on the MIPS R10000 processor. Options range from the Origin200, which can be configured with up to 4 CPUs, to the CRAY Origin2000, which can be configured with up to 128 CPUs.

An Origin200 server, belonging to SGI's benchmarking group, was made available to us for testing. This machine was configured with four R10000 processors, running at 180 MHz and 1 GByte of main memory. Each processor had 1 MByte of secondary cache, which is standard for this system. The disk system comprised seven 4.5 GB disks with two SCSI II controllers concatenated in software to form a striped volume of ~ 30 GB.

The tests were run using 15OCT96 with some components from 15APR97 used for the VLB DDT tests. The AIPS executables were compiled on a Power Challenge machine using version 7.0 of the SGI compilers under IRIX 6.2 (the test machine was running IRIX 6.4). The executables were compiled with level-2 optimization (-O2 for version 4 of the MIPS instruction set (-mips4) using 32-bit addressing (-n32).

## 2 Basic DDT Results

The large DDT was run twice, we had exclusive use of the machine for the second run while the machine had one other user during the first run. The timing results for both runs are as follows.

| Test            | Run 1    |              | Run 2    |              |
|-----------------|----------|--------------|----------|--------------|
|                 | CPU time | Elapsed time | CPU time | Elapsed time |
| $T_{large}/sec$ | —        | 302.0        | —        | 292.0        |
| $A_m^{(93)}$    | —        | 13.2         | —        | 13.7         |
| $T_{APCLN}/sec$ | 52.22    | 53.0         | 52.01    | 53.0         |
| $T_{CALIB}/sec$ | 17.63    | 18.0         | 17.25    | 18.0         |
| $T_{MXCLN}/sec$ | 73.27    | 75.0         | 72.77    | 74.0         |
| $T_{VTESS}/sec$ | 20.91    | 30.0         | 20.52    | 22.0         |

The  $u, v$  datasets generated during the DDT were not found to differ significantly from the master datasets. The number of correct bits in the image tests are shown here.

| Image  | Max. Error | RMS Error |
|--------|------------|-----------|
| UVMAP  | 14.8164    | 21.0378   |
| UVBEAM | 20.6107    | 26.0573   |
| APCLN  | 11.7986    | 17.1534   |
| APRES  | 13.7520    | 20.7065   |
| MXMAP  | 16.4649    | 22.5102   |
| MXBEAM | 20.9163    | 27.0501   |
| MXCLN  | 12.2863    | 16.7303   |
| VTESS  | 22.7466    | 30.9920   |

### 3 DDT Throughput Tests

In order to test the amount of degradation we might expect from multiple simultaneous users we ran 4 DDTs in parallel. Each was run from a separate terminal window using a different AIPS user number so that there would be no contention for message file locks. We did not have exclusive use of the machine for this test: one other CPU intensive process was started while the test was running.

| Test            | Results for Each DDT |       |       |       |
|-----------------|----------------------|-------|-------|-------|
|                 | 1                    | 2     | 3     | 4     |
| $T_{large}/sec$ | 357.0                | 354.0 | 329.0 | 314.0 |
| $A_m^{(93)}$    | 11.2                 | 11.3  | 12.2  | 12.7  |
| $T_{APCLN}/sec$ | 53.0                 | 53.0  | 59.0  | 56.0  |
| $T_{CALIB}/sec$ | 19.0                 | 18.0  | 24.0  | 19.0  |
| $T_{MXCLN}/sec$ | 73.0                 | 75.0  | 86.0  | 80.0  |
| $T_{VTESS}/sec$ | 21.0                 | 21.0  | 25.0  | 29.0  |

The worst-case degradation from the single-user time was 22.3%; the mean degradation was 15.9%.

### 4 VLB DDT Results

The performance of **BLING** was tested using simulated data generated by **DTSIM**, as implemented in the run file **VLBDDT** in the 15APR97 version of **AIPS**. For one processor the elapsed time per single baseline in the **VLBDDT** test was 95 seconds. This compares with a previously determined value of 185 seconds on a system with an  $A_m^{(93)}$  of  $\sim 6$ , and is thus consistent with the overall DDT results presented above.

### 5 Acknowledgements

We would like to thank Steve Simonds and Ed Hernandez of the Albuquerque SGI office for making the benchmark tests possible.