Paper 290

SAS® Software and the Performance Effects of Parallel Architectures

Keith F. Olsen, IBM Corporation, Poughkeepsie, NY
James T. West, IBM Corporation, Austin, TX

ABSTRACT

There are a number of different parallel architectures: parallel hardware implementations, parallel file systems, parallel application enablement software, and parallel databases. Tests were performed to evaluate and demonstrate the positive peformance effects of each of these architectures when used with the SAS System. A brief description of these tests, and their results, are presented in this paper.

INTRODUCTION

There are are a large number of SAS customers running many, if not all, the SAS products on many different computer architectures. As parallel architectures become more commonplace, from the desktop to super-computing, it will be advantageous to understand how those parallel architectures can be used by SAS customers. It will also be of benefit for SAS customers to understand the performance effects of using SAS on those parallel architectures, so as to be able to correctly build and run their SAS environments.

To this end, this paper will discuss actual test environments where different parallel architectures (Symmetric Multi Processing and Massively Parallel Processing) are used and the studied performance characteristics are discussed. Also tested and discussed are parallel software environments, from third party software developers, on those different hardware architectures.

Also investigated and reported will be the use of Relational DataBase Management Systems (RDBMS), flat files and SAS data structures. Performance testing that was done included data from these different sources. The file system and database configurations will be described, and the performance implications discussed..

When discussing performance implications, there are typically two scenarios that must be addressed. The first is the ability to run large single jobs in acceptable timeframes. As jobs get more complex, one tries to use parallel architectures to contain the more complex jobs in the same time slot. The second scenario is the ability to support multiple jobs and/or users, still maintaining acceptable response times for those users. Both the large single job and the large user community scenario tests will be discussed, and the performance effects of different parallel architectures will be described.

HARDWARE

Two parallel hardware architectures will be discussed: Symmetric Multi-Processors (SMP), and Massively Parallel Processors (MPP). SMP and MPP architectures are fundamentally different in that SMP is a shared-everything architecture, while MPP is fundamentally shared-nothing. What is becoming more commonplace, however, is that MPP architectures are now mixed in that some of the nodes of the MPP machines are SMP nodes. The ability to use this mixed architecture model is tested, and the peformance results are discussed.

The tests were conducted on the RS/6000 SP, a scaleable, multi purpose parallel computer at IBM's Teraplex Center in Poughkeepsie, NY. The Teraplex center is used for testing customer's real world environments, with the focus on large scale problems and typically involving large system configurations.

The RS/6000 SP is an MPP machine, but the nodes supported are both uniprocessor and symmetric multi-processor types. The RS/6000 SP has a high speed interconnect called the SP switch, which can support data transfer rates of 150MB/sec. This interconnect is not shared, so as more nodes are used the data transfer rates scales with the number of nodes using it.

One of the main benefits of the SP Switch is its scaleability. In order to achieve a positive performance gain, all bottlenecks must be removed. The SP Switch is a first step in that it removes any network bottlenecks.

PARALLEL FILE SYSTEMS

With the advent of parallel file systems, and their ability to support multiple users accessing data remotely , it was appropriate that the performance was measured and discussed. Parallel file system architecture is described and compared/contrasted to NFS with regard to performance characteristics.

IBM's new General Parallel File System (GPFS) is tested in a SAS environment, and the results of those tests are presented and analyzed. The GPFS configuration is also presented, as well as recomendations on how to set up and use GPFS in a SAS environment, based on those tests.

PARALLEL APPLICATION ENABLEMENT

There are software companies that have developed parallel software environments which can be used with SAS applications, Torrent Systems and Ab Initio. Tests were conducted with Torrent Systems' **Orchestrate** and **Orchestrator for the SAS System** products to show the ability of those software products to increase performance of both single large applications, as well decrease the time it took to extract large amounts of data from a relational database. These performance studies will be described and the data presented.

TESTS

There were a number of different tests that were run to understand and measure the performance characteristics and effects of parallel architectures. These tests could basically be grouped into three separate categories:

Parallel File System - AIX's General Parallel File System (GPFS) performance characteristics were tested and measured for multi-user environments

Parallel SAS applications - scalability was tested by measuring the run times of a SAS customer application that had been parallelized using third party parallel software on different hardware configurations

Data extract - measurements were performed on data extraction from a parallel DB2 database using SAS Access and the Orchestrate product. Multiple configurations were tested.

PARALLEL FILE SYSTEM

GPFS

AIX provides for a file system spanning nodes, called General Parallel File System (GPFS). GPFS uses a software layer called Virtual Shared Disk (VSD) which allows disks on remote nodes to look local to that node. GPFS also uses a client server relationship, with the GPFS server being on the same node as the VSD, and the GPFS clients being on the same or different nodes. In our tests, the GPFS clients were on different nodes. This allowed the GPFS server nodes to be dedicated to I/O, where the client nodes were dedicated to running the SAS applications.

HARDWARE

The configuration for the GPFS tests consisted of eight High nodes. Each high node was an 8 processor SMP, configured as follows:

8 - 604e PowerPC processors
2 GB memory
2 - 2.2GB internal disk drives
4 - Serial Storage Architecture (SSA) adapters
16 - 4.5GB SSA drives

SOFTWARE

AIX V4.2.1
PSSP 2.4
SAS System V6.12

The tests used two file systems of three different types, and compared the results of each. The three file system types were local, NFS mounted, and AIX's General Parallel File System.

The local file systems, DATAFS and TEMPFS, each consisted of 6 - 4.5GB disks, with a stripe size of 64K. The TEMPFS filesystem served as the SASWORK directory. The DATAFS filesystem served as the directory to which the created files were written.

The local file systems were NFS mounted across the SP switch for the NFS tests. The NFS mounted file systems were identical to the locally attached file systems.

See Fig. 1 for a diagram of the Local and NFS configurations.

Fig. 1
Local and NFS Configuration

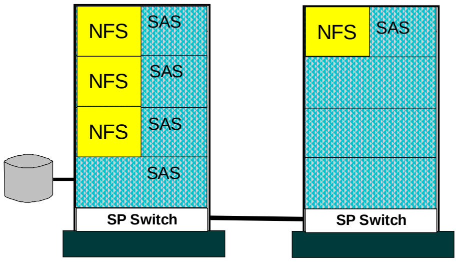

The GPFS file system tests were run twice. The first tests were run with the GPFS file system defined on just one node, identical in every way to the local and NFS file systems. Then, using a GPFS function, 3 nodes were added to the configuration so that a 4 node, 24 disk file system was created for both TEMPFS and DATAFS. The tests which ran 16 concurrent jobs used this 4 node GPFS configuration.

See Fig. 2 for the GPFS configuration.

Fig. 2
4 Node GPFS Configuration

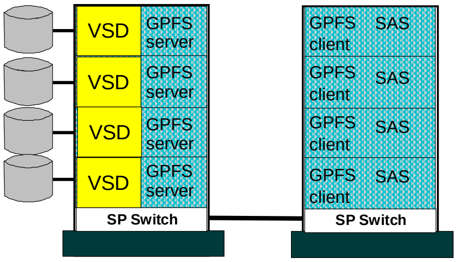

TESTS

The tests executed were supplied by the SAS Institute. The tests were run against 1990 US Census data. There were text files for each of the 50 US states. Many of the tests were run against just the California text file, which was about one tenth the size of the total census data. This was done in order to run many iterations of the tests during the allocated test period. There were a few tests run with the total data, and those test results were consistent with the findings from the California data tests.

There were four tests used:

Test 1 - Reads data from a text file, creates 2 output files (Household and Person data) **I/O intensive**

Test 2 - Sorts the Household and Person files, creates indexes that are needed to join the files together **Memory intensive**

Test 3 - Collects frequency on one and multiple columns, also does summarization **I/O intensive**

Test 4 - Collects statistics on the Household file - **CPU intensive**

RESULTS

The results of the tests demonstrated two important points: that GPFS functions well with the SAS System, and that it can be of great benefit for certain environments.

The following four figures show the response times for the four census tests. Each graph has three sets of bars: the first group indicates the time for 1 job running on Local, NFS, and GPFS file systems. The second group indicates average times for 4 concurrent jobs running, 1 on each of 4 separate nodes, for NFS and GPFS file systems. The third group indicates the average times for 16 concurrent jobs running, 4 on each of 4 separate nodes, for NFS and GPFS.

The clear benefits of using a GPFS file system are demonstrated in the third group. The average response times of 16 concurrent jobs for GPFS are significantly better than

those for the NFS file system. These results are logical in that we have scaled up the number of disks, adapters and processors to handle the I/O portion of the workload for the GPFS tests, while the I/O portion of the workload for the NFS tests all have to be handled by 1 node.
As the workload becomes more I/O intensive, this benefit will be more pronounced.

Fig. 3
Test 1 results

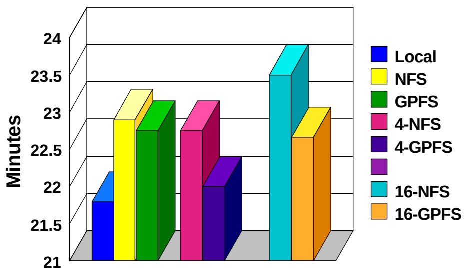

Fig. 4
Test 2 results

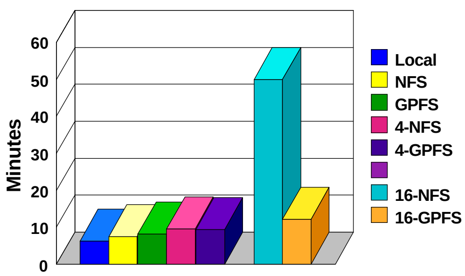

Fig. 5
Test 3 results

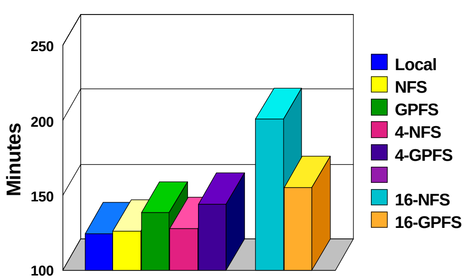

Fig. 6
Test 4 results

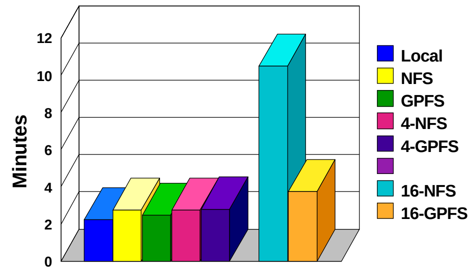

PARALLEL APPLICATION ENABLEMENT

HARDWARE

The configuration for the Parallel Application Enablement and the Parallel Database tests were the same. Four High nodes and thirty-two thin nodes were used in 1, 4, 8, 16 and 32 node configurations. For tests involving just the SMP nodes, 1,2,4, and 8 processors were used.
A description of the nodes:

HIGH NODES
8 - 604 PowerPC processors
2 GB memory
2 - 2.2GB internal disk drives
4 - Serial Storage Architecture (SSA) adapters
32 - 4.5GB SSA drives

THIN NODES
1 - 120 MHz POWER2Superchip processors
1 GB memory (in 28 nodes) 512MB in 4 nodes
2 - 2.2 GB internal disk drives
1 - SSA adapter
8 - 4.5 GB SSA drives

SOFTWARE

AIX V4.2.1
PSSP 2.3
SAS System V6.12
SAS Access V6.12
Orchestrate V 3.0
Orchestrator for the SAS System V3.0
DB2/UDB Version 5.0

TESTS

The parallel SAS application tests consisted of using Torrent Systems' Orchestrate and Orchestrator for the SAS System software to parallelize a real world customer application. Without changing the customers application code, Orchestrate manages the SP nodes, partitions the input data, and runs the application against the multiple partitions of data simultaneously to speed up the execution time of the whole job. Orchestrate has the capability to know how many nodes to utilize, and manages those nodes automatically. (A configuration file needs to be set up initially.)

The tests were run on both the high node configuration (1,2,4 and 8 processors of a high node), and the thin node configuration (1,2,4,8,16 and 32 thin nodes).

The test results for the thin node configuration(shown in Fig. 7), and the high node configuration (shown in Fig. 8) show that the parallel part of the application (middle portion) shrinks as the number of processors is increased, while the serial portions (top and bottom) remain constant.

Fig. 7
Parallel Appl. Speedup (thin)

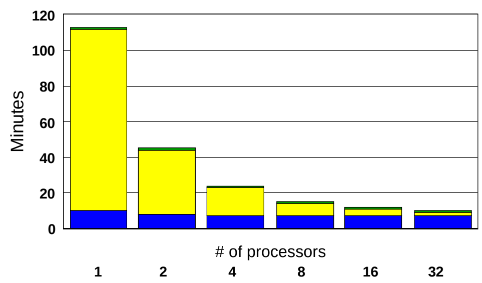

Fig. 8
Parallel Appl. Speedup (High)

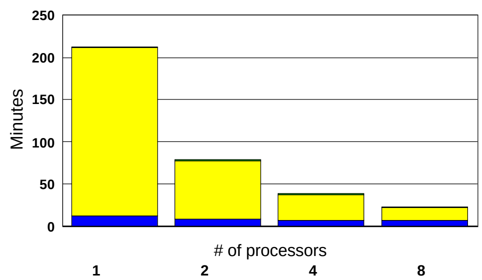

PARALLEL DATABASE

HARDWARE

The parallel database tests were run using the same configuration as the parallel application tests, using DB2/UDB Version 5 as the parallel database software. The database was configured on 1,4, 8, 16 and 32 nodes of the SP.

TESTS

The tests consisted of extracting data out of the DB2 database for the 4 different configurations. The amount of data was consistent in each test: 1.52GB of data was extracted each time. The tests extracted an entire table of 5 million rows, each record being 304 bytes long and composed of 126 fields.

Typically, when data is extracted from a parallel DB2 database, the data is sent to a coordinator node and then fed to the SAS application. This introduces a serial bottleneck, which in fact makes extraction times worse when you add nodes to the database.

In order to eliminate this serial bottleneck, SAS Access was enabled on each database node. Only data local to that node was extracted. Orchestrator for SAS was then used to merge the parallel streams of local data into a single SAS data set.

The results from the parallel data extraction are shown in Fig. 9.

Fig. 9
Parallel Data Extract

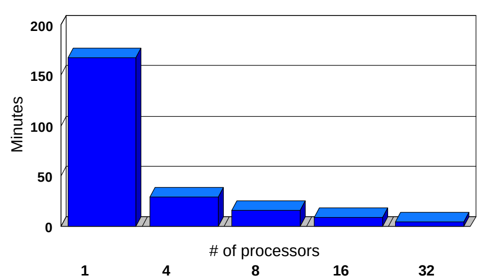

CONCLUSION

The results obtained demonstrated the positive effects on performance that can be attained through the use of different parallel architectures. While all of these parallel architectures may not be applicable to your specific application, the choice exists and greatly increases your ability to eliminate current performance bottlenecks. Performance bottlenecks can be found in I/O subsystems, memory subsystems, and the CPUs themselves.

Utilization of parallel architectures offers the ability to scale up any or all of these subsystems in order to alleviate performance bottlenecks.

The question has now changed from “Can I use a parallel architecture” to “Which parallel architecture will best meet my needs?”

REFERENCES

For more information, there are white papers available on the internet at the following URLs:

SAS and GPFS - A Scaleable Solution

IBM Corporation

http://direct.boulder.ibm.com/bi/press/whitepaper/index.htm

SAS Institute

http://www.sas.com/partners/enterprise/ibm/whitepap.html

Achieving Scaleable Performance for Large SAS Applications & Database Extracts

IBM Corporation

http://direct.boulder.ibm.com/bi/press/whitepaper/index.htm

Torrent Systems

http:// www.torrent.com/whitepapers

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the authors at:

Keith F. Olsen	James T. West
IBM Corporation	IBM Corporation
522 South Road	11400 Burnet Road
Poughkeepsie, NY 12601	Austin, TX 78758
914-433-8443	281-334-2906
914-433-8618	281-334-1973
ko @ us.ibm.com	westj @ us.ibm.com

TRADEMARKS