Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831

Prices available from (615) 576-8401, FTS 626-8401

Available to the public from
National Technical Information Service
US Department of Commerce
5285 Port Royal Rd
Springfield, VA 22161

NTIS price codes
Printed copy: A05
Microfiche copy: A01

SAND90-2229
Unlimited Release
Printed November 1990

**Simulation Technology Research Division
Assessment of the
IBM RISC SYSTEM/6000 Model 530 Workstation**

Greg D. Valdez
Applied Methods, Inc.
Albuquerque, NM

John A. Halbleib, Ronald P. Kensek
and Leonard J. Lorence
Simulation Technology Research Division
Sandia National Laboratories
Albuquerque, NM

ABSTRACT

A workstation manufactured by International Business Machines Corporation (IBM) was loaned to the Simulation Technology Research Division for evaluation. We have found that these new UNIX workstations from IBM have superior cost to performance ratios compared to the CRAY supercomputers and Digital's VAX machines. Our appraisal of this workstation included floating-point performance, system and environment functionality, and cost effectiveness. Our assessment was based on a suite of radiation transport codes developed at Sandia that constitute the bulk of our division's computing workload. In this report, we also discuss our experience with features that are unique to this machine such as the AIX operating system and the XLF Fortran Compiler. The interoperability of the RS/6000 workstation with Sandia's network of CRAYs and VAXs was also assessed.

MASTER

ACKNOWLEDGMENTS

There are a number of people whom we would like to acknowledge for their important contributions to this project. Steve Bauer, our main contact at IBM, coordinated the efforts of IBM's RS/6000 support team. His timely provision of much needed documentation and information expedited the evaluation process. We are indebted to Michael Martinez and Bill King for providing tireless on-site installation and valuable assistance. Our thanks also to Glenn Bryans, Lisa Martin and Paul Gooderham for fielding our compiler questions and recommendations, and to Craig Myers and Ernie Milton for their suggestions.

Our many thanks to IBM's local branch management, led by Ted Hobbs and Bob Freese, for allowing us unrestricted access to the organizations under Brian Casto and Fred Peterssen. Their cooperation and enthusiasm are greatly appreciated. We also thank the support teams from IBM's Dallas, Austin, and Toronto facilities for their cordial and quick response.

The authors also wish to thank Phillip Campbell for carefully reviewing this report. His insightful suggestions added to the reports clarity, and direction.

Finally, we are grateful for the advice and encouragement of Jim Lee during this work.

This work was supported by the U. S. Department of Energy under Contract DE-AC04-76DP00789.

DISCLAIMER

This report does not constitute an expressed, or implied recommendation or endorsement by Sandia National Laboratories of the IBM RS/6000 workstations. It is intended simply as feedback. It must be clear that Sandia takes no responsibility whatsoever with regard to the performance, acquisition, or use of these products.

TABLE OF CONTENTS

<u>SECTION</u>	<u>PAGE</u>
EXECUTIVE SUMMARY.....	8
1.0 INTRODUCTION.....	9
2.0 CODE MIGRATION EXPERIENCES.....	13
2.1 ITS.....	13
2.2 CEPXS/ONELD.....	16
3.0 RESULTS.....	19
4.0 OTHER CONSIDERATIONS.....	25
4.1 RISC Implementation.....	25
4.2 Review of Current Workstation Technology.....	27
4.3 AIX Version 3 Operating System.....	30
4.4 Connectivity.....	32
4.5 Security and Data Integrity.....	34
4.6 Ergonomic Design.....	36
4.7 Customer Support.....	37
4.8 Future Directions.....	38
5.0 CONCLUSIONS.....	41
6.0 REFERENCES.....	43
 APPENDICES	
A. MACHINE DESCRIPTION.....	47
B. BENCHMARK CODES DESCRIPTION.....	51
C. OPEN SOFTWARE FOUNDATION.....	53

EXECUTIVE SUMMARY

In this report, we evaluate whether the new IBM RISC SYSTEM/6000 Model 530 POWERstation is a suitable platform for performing the computationally intensive tasks of our division. In particular, we compare this computer with the mainframe computers that are used at Sandia, which are the Cray and the VAX. In addition to timing comparisons with these mainframe computers, we also discuss in this report our experience migrating codes to the machine, and connectivity issue between this workstation and the mainframe. We also discuss various features of the IBM workstation such as the operating system, the Fortran compiler, and system tools.

We conclude that the IBM workstation can effectively compete with the Cray and the VAX. We have found the IBM workstation to be at least 4 times faster than the 8650 VAX, and at most, 50% slower than the Cray X-MP/416. For some cases, this IBM workstation approaches 80 to 90 percent of the Cray performance. While we have not evaluated other RISC-based computers such as Sun, HP/Apollo and DEC workstations, the RS/6000 has a big advantage over these machines on floating-point performance.

Our experience with our related interaction with IBM personnel, at both local and out-of-state, is that IBM is quick to respond and is eager to compete in the open arena.

Simulation Technology Research Division
Assessment of the
IBM RISC SYSTEM/6000 Model 530 Workstation

1.0 INTRODUCTION

The rapid advancement in desktop computing technology has shifted the laboratory's emphasis and future directions toward distributed computing. In the Simulation Technology Research Division at Sandia National Laboratories, we are considering the purchase of a UNIX workstation suitable for an individual scientist or group of scientists, to help with the development and testing of codes and the interpretation of data from production runs. Since workstation technology for scientific computation is evolving so rapidly, the choices between different platforms are not clear. Recognizing the danger of relying on manufacturer claims and counter claims, we decided that a hands-on type of evaluation would be valuable.

We determined a set of quantitative variables to consider regarding hardware specifications for a workstation. Based on our experience and relative costs incurred within our current computing environment, the main criterion of any workstation that we acquire is that the workstation should offer superior performance compared to the VAX 8650 and should approach the performance of the CRAY XMP. To accomplish this, the workstation should:

- have a large virtual memory addressing capability,
- process at least 25 million instructions per second (MIPS),
- execute at least 10 million floating-point operations per second (MFLOPS),
- have an I/O bus that transfers at least 5 million words per second,

- have an internal bus that can move data to and from memory at speeds of at least 60 million words per second,
- be capable of having at least 128 Megabyte of system memory and at least 2 Gigabyte of fixed disk storage,
- have a communications throughput of at least 10 million bits per second, and
- have a color graphics display with a pixel resolution of at least 1280 x 1024.

In addition to the requisite UNIX software, the workstation should possess a robust operating system that includes an efficient set of compilers, a highly-integrated windowing environment, and utilities to simplify and expedite system management and administration. The platform of choice should also exhibit capability for improvements and expandability to be competitive with supercomputer floating-point performance over the next 5 years, and have a payback period no longer than 2 years.

These new UNIX workstations from IBM have been shown [SIMM, BROW2, DONG] to exceed cost/performance of the CRAY supercomputers and Digital's VAX machines, the current computer workhorses at Sandia. Published specifications also seem to indicate that the IBM workstation meets our stringent requirements better than comparable workstations from Sun, HP, and DEC.

Hence we obtained on loan a pre-release version of an International Business Machines Corporation (IBM) workstation for Beta-test evaluation. We benefitted from this arrangement by being able to test a platform in our environment before committing resources. In return, IBM received our unbiased assessment of the RS/6000's performance, feedbacks on implementation and documentation errors, and insights on security issues.

Our appraisal of this workstation included floating-point performance, system and environment functionality, and cost effectiveness. Our assessment was based on a suite of radiation transport codes developed at Sandia that constitute the division's computing workload. The interoperability of this workstation with Sandia's network of CRAYs and VAXs was also assessed.

We successfully converted and tested, on the model 530 machine, the ITS and CEPXS/ONELD code packages. This midrange workstation, which starts at a non-GSA price of \$42,000, features a 25 MHz CPU that produces 34.5 MIPS and 10.9 MFLOPS performance ratings. Detailed descriptions of this model and other members of the RISC SYSTEM/6000 series (referred to as RS/6000 in this report) are included in Appendix A. The RS/6000 workstations operate under AIX, which is IBM's version of the UNIX operating system. AIX is POSIX standard¹ compatible, and contains support for both the Berkeley and AT&T implementations of UNIX. X Windows Version 11 Release 3 is supported along with the C, Bourne, and Korn command shells. Although the RISC System/6000 has only been recently announced, more than 600 software applications have already been ported to this platform. These include DISSPLA, Frame Maker, Saber-C, visualization software such as PV-Wave and DATA VISUALIZER, Matlab, Mathematica, MASS-11, WordMARC Composer Plus, WordPerfect, SAS Statistical Software, ANVIL-5000, NASTRAN, PATRAN, SLATEC, IMSL and NAG Fortran Libraries, Glockenspiel C++, ORACLE, etc. The number of available third party software is projected by IBM to exceed 1500 by the end of the year [UPDA5].

1. The POSIX standard, IEEE 1003.1-1988, is a set of programmer's interfaces and standards for UNIX and UNIX-like systems. It's not an operating system itself; instead, it specifies how programs and other systems can connect with an operating system.

This report describes the migration of our applications to the IBM model 530 workstation, and assessment of its suitability for our present and anticipated computing needs.

2.0 CODE MIGRATION EXPERIENCES

The two radiation transport code packages that represent the bulk of our division's workload were successfully ported to the IBM RS 530 workstation. These are the Integrated TIGER Series (ITS) of Monte Carlo codes [HALB], and the CEPXS/ONELD discrete ordinates code package [LORE]. Also ported was the UPEML update code [UPEML] that expedites development and maintenance of the above code packages. The major codes that were benchmarked are further described in Appendix B.

2.1 ITS

When porting the ITS codes to the model 530 workstation, we encountered problems with the Beta-release versions of the AIX operating system and the XLF Fortran Compiler. For example, we ran into obscure diagnostics such as "memory fault creation error", and "fatal error: signal 9 received". We were told that these indicated some sort of segmentation violation. Even with the help of IBM personnel, we could never, with this early version of the compiler and operating system, circumvent these problems, nor could we precisely identify what caused them. We suspect that the inability to change the default data size from 1024 bytes to 4096 bytes could have contributed to the problem. More informative compile and run-time diagnostics would immensely aid in resolving these problems.

The early versions of the AIX operating system also unpredictably froze the screen when using the hypertext information utility, INFO Explorer. This required a hard boot of the system for each occurrence. The only hardcopy documentation available for our Beta-test were draft copies of the XLF Fortran and C language user's guide and reference manuals. Hence, we were forced to rely on the software-based documentation in order to effectively use the system. Consequently, our initial impressions

were not favorable due to the lack of dependable documentation. Early releases of the operating system also showed noticeable signs of degradation in processing when InfoExplorer was activated. While these and other initial difficulties were bothersome, this experience did enable us to become familiar with the intricacies of the machine through our frequent contact with IBM. Undoubtedly, our initiative provided valuable feedback to them. It turns out that we were exercising features of the machine that had not, at that time, been adequately tested.

As expected, when we obtained a newer and more stable version of the AIX operating system and Fortran compiler, most of our initial porting problems were solved.² To successfully compile, we also had to resolve collisions between system callable routines and some of our subprograms. Although a tedious task, this only entailed changing the names of, and calls to, the local subprograms, START and ABORT.

We also had to resolve unsatisfied external routines in our programs for successful compilation. Other machine platforms did not seem to balk at this. We had used the trick, on the VAXs and Crays, of deliberately calling a non-existing routine to obtain a meaningful traceback of the error's origin. The IBM Fortran implementation provides a way to obtain these tracebacks only through the use of system signal handler routines. Admittedly, this approach, although more robust, places the ultimate responsibility on the developer to trap and process all abnormal interrupts. For example, the XLF Fortran compiler does not provide any diagnostic when a division by zero occurs. If the programmer fails to use the interrupt utilities properly, the code will not crash when such an overflow condition is encountered, which may lead to erroneous results.

2. The final version of the AIX operating system and compiler we tested was the 9013 release version.

The optimizing features of the Fortran compiler proved to be very efficient as our benchmark results attest. The compiler's "instruction scheduling," a process in which the compiler examines the code flow and determines where optimization can be taken to streamline pipeline execution, is highly maximized to make sure that data are kept in registers, instead of in memory, as much as possible. Since memory fetches are costly, the less of it required the better.

However, some lengthy ITS code subroutines failed to optimize. Since the compiler did not provide helpful diagnostics, we resorted to the use of the NOOPT in-line compiler directive in each of the subprograms to narrow down the search for routines that failed to optimize. Needless to say, this was a lengthy process. We were able to locate the problematic routines, but due to time constraints, we were not able to track down the causes. Even though the ITS codes were not fully optimized, the ITS codes on the IBM workstation are still at least 4 times faster than they were on the VAX 8650 computer. It is conceivable that time differences between the VAX and the IBM can be further increased if full optimization is achieved on the RISC workstation.

Crucial to the machine implementation of a Monte Carlo code is the behavior of the random number generator. Evaluation of the default random number generator on the IBM machine showed that it is as robust as the VAX random number generator, i.e., it has similar cycle length and range, including undesirable generation of zeroes. (The Cray random number generator does not generate zeroes and presumably has a larger cycle length.) One of the more important functionalities we require is access to the "seed" which generated a particular random number set. This is necessary for debugging and for restart runs. Our users have not been able to access the random number seed in versions of the ITS codes that have been ported to the SUN/4 workstation. We have been told that the only way to obtain the seed on the SUN is through a

user-written C program which is then linked with the main program [PER]. The random number generator for the HP workstations presents a similar problem. Therefore, versions of our code ported to these platforms are limited to production calculations since restart and debugging activities associated with code development are severely constrained. However, IBM has made accessing the random number seed very convenient by storing it in the /RANDDAT/ common block. We find this issue to be a very strong point of IBM's random number implementation, and an important consideration for any stochastic code development and application. However, we highly recommend a version of a random number generator that excludes zero since zero random numbers are problematic in some Monte Carlo applications.³

We are aware that mixed-language programming is becoming more acceptable these days especially within the UNIX community. We ourselves have developed and linked C programs with our codes when applicable. However, due to the universal appeal of our codes and the type of users we serve, we find it necessary to adhere to the purity of the Fortran language as much as possible to maintain code portability, enhance user access, and encourage appropriate code maintenance.

2.2 CEPXS/ONELD

The next group of tests involved CEPXS/ONELD, a one-dimensional coupled electron-photon discrete ordinates code package [LORE]. This suite of codes consists of the cross section generating code, CEPXS; the preprocessor code, PRE1D; the discrete ordinates code, ONELD; and the postprocessor code, POST1D.

3. *Explicitly checking for zeros adds unnecessary and costly execution time penalties.*

Porting CEPXS and ONELD to the IBM workstation presented some difficulties because of the strict adherence of the IBM XLF Fortran compiler implementation to the ANSI 77 standards [ANSI]. Of course, these problems are not all attributed to IBM's compiler since ONELD contains several non-standard Fortran instructions. However, the XLF compiler needs to comply with generally accepted ANSI extensions. For example, although it is a non-ANSI standard practice, most compilers will allow initialization of variables declared in named COMMON within the same subprogram. IBM forces one to abide with the standard which requires that declaration and initialization of these types of variables and array elements must be within a BLOCK DATA subprogram. IBM is aware of the inconvenience this presents and will relax the restriction in the final released version of the Fortran compiler.

We were also disappointed by the compiler's anemic supply of built-in utility and service functions. For instance, a very useful function that is used to compute the internal address of a storage element, usually called LOC, was sorely missed. We were forced to write a similar function in C to process ONELD's memory allocation routines. In addition, IBM's timing routines, SECOND and MCLOCK, when compared to similar VAX, Cray, and Sun provisions, are quite limited. To their credit, however, IBM was more than eager to provide assistance to meet our specific needs. Along these lines, we have found IBM to be accommodating as they strive to provide a Fortran compiler that is amenable to the majority of technical users. Recognizing that the VAX/VMS Fortran compiler is some sort of a de facto standard, IBM has gone to the extent of gathering input from Beta-test sites and evaluators in order to include some of the more useful VAX Fortran extensions in the final version of the compiler.

3.0 RESULTS

Our numerical processing evaluation procedure differs from an earlier Los Alamos study [SIMM]. In that report, CPU performance was measured through a set of portable benchmarks that are representative of the scientific workload for the whole laboratory. This extensive evaluation included:

- vector operation performance as a function of vector length for several different memory access patterns,
- comparison with other RISC machines such as the Sun SPARCstation 4/330 and DECstation 3100,
- comparison with mid-range vector computers, the Convex C-240 and the Floating Point Systems FPS-500 FPX,
- and comparison with the CRAY YMP representing a high-end vector machine.

It is important to note that although the IBM RS/6000 machines do not have vector capability, their superscalar architecture allows the processor to execute up to four instructions per cycle, including a floating-point instruction capable of completing a double-precision 64-bit multiply and add (two operations) in one cycle. Therefore, this machine is capable of executing up to five operations per clock cycle. This enhanced capability, which is unique among workstations, enables the RS/6000 system to approach the performance of vector processors.

In our study, we concentrated our efforts on applications that are currently relevant to the needs of the Simulation Technology Research Division at Sandia National Laboratories. Even though our tests were more local in nature, they are nonetheless computationally intensive. In addition to timing comparisons, our assessment also included evaluation of the Fortran compiler, the operating system, the development environment functionality, connectivity, and system tools. We

also considered, in our assessment, the future directions for the RS/6000 and software availability.

For the ITS standard test problems, Table 1 shows that the IBM workstation is about 4 to 5 times faster than the VAX, and about 36% to 48% slower than the Cray XMP. The same comparison holds for the CEPXS test runs on the VAX and the workstation. In comparison with the Cray, a significant result shows that the workstation can approach within 90% of the supercomputer's performance. Similarly, the ONELD code performed extremely well on the IBM workstation, about 17 times faster than the VAX. Since ONELD is a highly vectorizable code and the VAX does not have vector processing capability, this result is not that surprising. What is significant, however, is the ability of the workstation to stay abreast of the CRAY XMP performance. This shows the power of superscalar architecture that IBM has implemented. In some of our other ITS and CEPXS test problems, the IBM workstation consistently approached 80% of the Cray performance and exceeded that of the VAX by a factor of six.

One of the more intriguing results from our evaluation involved the CEPXS/ONELD code. This multigroup deterministic code typically requires more memory than a stochastic code with the same functionality. While these requirements are easily met on the Cray, we have had occasional difficulty obtaining sufficient memory on some extremely large CEPXS/ONELD problems on the VAX. To assess the capability of the IBM workstation to handle problems with large memory requirements, we manufactured CEPXS test problems of arbitrarily large size. Not only were we able to run larger problems on the workstation than was possible on the VAX, but, surprisingly, we were able to link larger executables on the IBM workstation than was possible on the Cray XMP. One would certainly expect the VAX, because of its virtual memory architecture, to handle such a configuration, especially when our allocated paging space has been set to the maximum allowable.

Table 1. ITS Ta-A1 Test Case for 10 K Histories and 10 Batches. CPU sec Per Batch and Machine Run Time Ratios.

	ACCEPT	ACCEPTM	ACCEPTP	CYLTRAN	CYLTRANM 2	CYLTRANP	TIGER	TIGERP
IBM 530	115.980	139.053	127.302	85.137	99.278	99.551	66.962	75.444
VAX 8650	480.651	698.208	550.856	386.156	513.468	457.010	262.471	315.848
Cray XMP-416	69.428	83.116	74.820	53.281	56.172	64.022	35.037	43.951
VAX / IBM	4.14	5.02	4.33	4.54	5.17	4.59	3.92	4.19
Cray / IBM	0.59	0.60	0.59	0.64	0.57	0.64	0.52	0.58

Table 2. CEPXS/ONELD Test Case Run times (CPU sec) and Machine Run Time Ratios.

	CEPXS		CEPXS	
	TEST #1	TEST #2	TEST #1	TEST #2
IBM 530	31	180	33	ONELD
VAX 8650	158	762	578	
Cray XMP-416	28	---	15	
VAX / IBM	5.10	4.23	17.52	
Cray / IBM	0.90	----	0.46	

Another reason for our ability to run large problems is an important innovation that IBM has introduced. The AIX version 3.0 operating system has been improved to take advantage of the new hardware. The new operating system has a 52-bit virtual address space that lets the system manage programs and data sets of virtually unlimited size. A 52-bit virtual addressing scheme allows for more than 4000 trillion-bytes (TB) of system-wide memory. By comparison, systems using 44-bit addressing, allows for 17 TB, while 48-bit addressing yields 281 TB of virtual memory. In addition, a feature called "logical volumes" allows a single file system to be allocated across several physical disk drives [LEVITT]. These capabilities, as we have proved, will be useful for handling large executables and the increasingly large volumes of data produced by our transport codes. The fact that a workstation can actually surpass a Cray supercomputer in this regard is surprising. Since the SUN/4 workstation architecture does not have as large a virtual addressing mode because of poorly designed Memory Management Unit (MMU) hardware [IRLAM], and has not demonstrated a logical volume capability, we do not expect it to perform nearly as well as IBM's RS/6000 in this area.

Although we did not compare IBM's model 530 with similar workstations from other vendors, published results for the RS/6000 family in terms of such traditional benchmarks as Dhrystones and Whetstones show that the IBM machine is about 3 to 5 times faster than the competition [SIMM, SOUT, DVOR, BROW, CADR, ENTI, HASK, WILS]. Dr. Jack Dongarra, author of the widely-used LINPACK routines [LINPK] and distinguished computer scientist and mathematician for the University of Tennessee and Oak Ridge National Laboratory, noted that although the RS/6000 workstation he tested did not run as fast as the Cray YMP supercomputer, in its price category it is "an order of magnitude" faster than existing competitive machines [UPDA1]. His benchmarks [DONG] and test results from the Workstation

Laboratories [CRN] substantiate the power of the RS/6000 workstation.

The power of the IBM workstations has also been demonstrated in another study performed at Sandia National Laboratories. Preliminary compute-intensive ray tracing comparisons conducted by the Radiation Effects and Reactor Design Division (6452) showed that IBM's entry model (320) surpassed the VAX 8650 about 3.5 times, the Sun SPARC 1 by a factor of 3, the NeXT workstation by about 13 times and the PS/2 Model 80 personal computer by 30 times [GRIF]. Their results further showed that the model 320 approached 70% of the Cray 1/S run-time and 60% of the Cray X-MP 416. Comparisons with the model 530 are more dramatic, posting times within 90% and 80% of the two Cray supercomputer performances.

Other investigators have found the RS/6000 workstation line to be faster than Cray's supercomputer on specific scientific tests. The tests performed at Lawrence Livermore Laboratory showed that the IBM POWERstation 530 performed 10% to 50% faster than the Cray X-MP 4/16 when running applications that are scalar processing intensive [DALY]. These findings indicate that the \$42,000 IBM machine offers performance similar to the \$23 million Cray Research machine at approximately one six-hundredth the cost. Of course, this result nowhere implies that the same performance ratios apply to applications that are vector-processing intensive or to applications that utilize multi-processing features in which supercomputers are still superior. However, parity with supercomputer vector capabilities may not be that far off [JEFF].

We recognize that these results do not prove that most applications will run 4 to 6 times faster on the RS/6000 model 530 workstation when compared to the VAX, or that the model 530 is equivalent to the Cray (although, total throughput just might if turn around times are factored in), or that it runs 3 to 5

times faster than other UNIX workstations. However, they strongly indicate that the RS/6000 model 530 workstation is a very fast computer, presenting a viable alternative to mainframes, and that it is more than competitive with other RISC platforms.

4.0 OTHER CONSIDERATIONS

Is being competitive enough? With the stream of emerging technologies, wouldn't it be prudent to think that Sun, DEC, HP or other RISC workstation manufacturers will be able to leap-frog IBM? After all, they have been in the workstation market longer than IBM. So why not stick with one of the "proven" platforms? Another common concern is whether the IBM RS/6000 is a true RISC machine. These are representative of questions we ask ourselves.

4.1 RISC Implementation

RISC, which stands for Reduced Instruction Set Computer, was invented in 1975 by John Cocke, a research fellow at IBM's Thomas J. Watson Research Center. The counterpart of RISC are processors known as Complex Instruction Set Computers (CISC). A good example of a CISC processor is Digital's VAX computers. It is known to have one of the most complex instruction sets. The VAX processor has more than 300 instructions, operates on 5 different integer types, 4 floating-point types, and a host of other data types including strings and queues. It has 16 general-purpose registers and 21 different addressing modes [BERG].

Robert Dewar and Matthew Smosna, professors at the NYU Courant Institute of Mathematical Sciences, have written a newly released book discussing the differences and merits of widely used microprocessor chips [DEW]. In the book they noted that IBM defines RISC, not as a "reduced set of instructions", but a "set of reduced instructions". The first definition, adhered to "strictly" by Sun and DEC, suggests that the number of instructions needs to be reduced. The second definition, however, implies that it is not the number of instructions that need to be reduced, but the complexity of the instructions in the

instruction set. This means that less complex instructions result in increased speed in processing these instructions, providing faster overall throughput.

The authors of the book further state that the goal of RISC is to execute one instruction per clock cycle. Consequently, this requires not only simplified instructions, but also optimal use of the processor's instruction pipeline. The implementation of RISC in the RS/6000 takes full advantage of the latter requirement. Most other RISC chips have been characterized as containing a reduced number of instructions.

To give credence to this, the authors argue that the various RISC architectures they compared differ greatly in the extent to which they hide or expose the instruction pipeline. For example, the MIPS chip, based on Stanford University's RISC research, relies on software conventions rather than hardware interlocks for handling what is called the "load delay slot." In fact, they noted that the acronym, MIPS, originally stood for "Microprocessors without Interlocked Pipeline Stages." The i860 chip by Intel is even more non-transparent since programming it looks as though it consists almost entirely of pipeline manipulation. These authors also say that "It is interesting to note that even SPARC, despite its general commitment to very simple instructions, supports a rather high-level notion of tagged data, with an add instruction that checks the tag. Though quite welcome in a LISP implementation, it is a slight departure from the 'true' path of RISC" [DEW2].

On the other hand, the architecture used in IBM's RISC System/6000 is a lot easier to deal with. According to Dewar and Smosna, the RS/6000 chip has an advantage over other chips since "with just a little knowledge of what is going on--basically little more than the rule that you should not use results you just computed--the programmer can write code that compiles in a highly efficient manner, without needing a sophisticated

optimizing compiler." Our experience with the RS/6000 530 certainly agrees with theirs.

Hence, IBM has not redefined nor violated the principles of RISC, but rather, improved on its implementation.

4.2 Review of Current Workstation Technology

At comparable clock rates, Sun's SPARC processor runs 2 to 8 times slower than the RS/6000. Sun could match IBM's fixed point performance by doubling its current clock rates (the Sparc 2 due in December is suppose to do this [GROS]), but it still can't match floating-point performance without a major overhaul in its architecture. It is well known that in its SPARC line, Sun uses two incompatible floating-point coprocessors. The SPARCstation 1 uses the Weitek 3170 and the SPARCservers use Texas Instruments' TI 8847 FPU and a TI floating-point controller. In order to execute the same code across the two floating-point architectures, Sun specifies a set of floating-point instructions as part of the SPARC architecture. Each SPARC-compatible floating-point unit must therefore include a controller to arbitrate between the SPARC floating-point instructions and the format used by the chip itself. Even though the translation is transparent it still incurs non-negligible overhead.

Since Sun is constrained to maintain its current instruction set so that existing SPARC software can run on new machines, they will most likely rely on third party hardware developers for floating-point processors to boost performance. This is quite an expensive option which would also require extensive code changes in present applications. Furthermore, because of system overhead, an add-on processor will not perform as fast as an integrated floating-point unit. Sun may also develop their own "pipelined" floating point processor but this approach will be difficult and would take considerable resources.

Hewlett-Packard is one of few workstation vendors whose main line is built on the Complex Instruction Set Computer (CISC) architecture. To protect their large investments on CISC-based computers, this will likely remain as HP's position in the highly competitive UNIX workstation market. Thus, Hewlett-Packard's current architecture (and perhaps future ones) are hard-pressed to compete with the IBM RS/6000. For example, doubling their current microprocessor clock rate to 30 MHz achieves about 2/3 of IBM's fixed point performance and only 1/3 of its floating-point performance.

The merger with Apollo provides HP with RISC-based superscalar workstations. However, if HP attempts to incorporate superscalar attributes into the HP 9000 series, the new workstations will not maintain binary compatibility in the next iteration for their "Precision Architecture" and will therefore incur costly code rewrites, recompilation and distribution. However, because of HP's extensive lead and expertise in instrumentation control and data acquisition, it is still the preferred technical platform for such applications.

Digital Equipment's current line of RISC workstations lacks high-performance graphics options, has no expansion options, and depends on microprocessor technology from MIPS computers. The DEC machines, which are not priced significantly lower, fail to match IBM's RS/6000 family in raw power.

Our first-hand impressions corroborates those of industry analysts who almost unanimously agree that IBM has succeeded in its attempt to compete with more established workstation vendors by producing a line of machines that operate nearly 3 times as fast as their closest competitors on numerically intensive computations, while offering competitive pricing [BROW2, CRN, HASK, SOUT, SYSI].

Analysts say that the price-performance ratios of IBM's RS/6000 is better compared with more expensive workstations from Stardent Computers, Inc. and MIPS Computers for compute-intensive processing [BROW2]. They found that the basic MIPS architecture supporting a 60 MHz clock does not match the floating-point power of the RS/6000 Model 540 running at half the clock rate. Stardent relies on a specialized vector processor, and MIPS cranks up clock rates while IBM employs leisurely clock rates with its superscalar architecture to achieve record floating-point performance in a workstation. The authors of this report [BROW2] further projected that at equivalent 50 MHz fabrication levels, the IBM superscalar architecture would overwhelm the SPARC alternative, and exceed the MIPS architecture performance by at least 44 percent.

Clearly, first generation architecture, such as SPARC, and faster clock rates can not effectively compete with IBM's second generation superscalar architecture implementation. If Sun, HP or DEC attempt to match the superior superscalar performance of second generation RISC machines, such as IBM's, they are likely to encounter incompatibility problems with their first generation RISC machines.

Our evaluation of the RS/6000 model 530 has convinced us that this new series of workstations presents substantial improvements over other workstations, and promises to deliver on future expectations that will allow us to expand our present modeling capabilities within the Simulation Technology Research Division. It is the most cost effective platform for meeting our division's current and anticipated needs for computationally intensive calculations.

4.3 AIX Version 3 OPERATING SYSTEM

AIX is a multi-tasking demand-paged, virtual memory operating system [UPDA4, UPDA7]. Its database management capability features logical disk volumes, object-oriented architecture, single-level storage and dynamic configurations, and a hardware assist feature to speed up the processing of large database applications in virtual memory [CER].

According to a leading reviewer [EUNI], IBM's AIX Version 3 constitutes the most advanced UNIX alternative available today when compared to AT&T's System V Release 4 (SVR4), Berkeley's BSD and Carnegie Mellon's Mach flavors of UNIX. They believe that IBM has a strong commitment to fundamentally improve the kernel and layered services of UNIX, thereby fixing the traditional weaknesses of UNIX as a production operating system. They point out that AIX's enhancements extend UNIX's reach beyond the desktop confines of technical and engineering departments. Furthermore, they believe that as the chosen core operating system of the Open Software Foundation (OSF)⁴, AIX aims to deliver everything SVR4 has today, as well as multiprocessor support, B-level security, and fundamental improvements to the kernel technology that should prepare it to evolve gracefully over the next decade.

IBM's active involvement with OSF, a non-profit research and development organization dedicated to develop and deliver an open software environment based on standards for the benefit of the information processing industry, is a strong indication of their commitment to open systems. Ironically, Sun Microsystems Inc. is not a member of OSF. We are further encouraged by IBM adopting an open systems strategy on their RISC platforms, a departure from

4. See Appendix C.

their past practices of pushing "proprietary standards". This move, indicative of their commitment to the future of the UNIX community, relieves our concern that there may be inadequate third-party add-in hardware to support the IBM RS/6000 workstations.

For brevity, we have not included in-depth discussion of the significant improvements that AIX have introduced. The interested reader is directed to the reference [EUNI] for more details.

To enhance ease of use, performance, and reliability, AIX offers many improvements such as physical disk space management, advance file system, program management facilities, extended real time support and enhanced virtual memory. The new UNIX implementation also provides improved system management and network install facilities.

Version 3 provides the System Management Interface Tool (SMIT), an interactive system management facility. This menu-driven system uses system data to determine the state of the system and generates menu options. It facilitates the management of installation, devices, storage, security, users, communication applications, spoolers, problem determination, system performance and policy and system environments.

InfoExplorer, a hypertext database search and retrieval system, provides access to both text and graphic on-line documentation. The user can activate visible hypertext links and access related or prerequisite information. Users can also customize the presentation, set bookmarks, annotate the documentation with local site-specific information, view graphics with zoom/pan capabilities, print information, or copy information to files for private use. The Hypertext Information Base contains documentation on general user tasks, system management tasks, communication tasks and commands. It can be implemented in a single-user workstation or a multi-user server

environment. The Hypertext Information Base can reside on a fixed disk or on CD-ROM. When shipped on CD-ROM, additional graphics files are included. From the CD-ROM Hypertext Information Base, text and graphics can be viewed on graphics displays, and text can be read on ASCII displays. On a graphics display, the database can be accessed via AIXwindows. In addition to CD-ROM support, another attractive feature is the RS/6000's support for digital audio tape for documentation, software distribution and data storage.

Use of shared libraries is enhanced by support of page mapped loading, dynamic load-time binding, and a common linkage convention for inter-language calling. Dynamic binding enables program updating, since individual modules can be added or replaced without requiring the total program to be rebound or relinked. Communications and network computing support includes TCP/IP, X25, Network File System (NFS), and Network Computing System.

4.4 CONNECTIVITY

The IBM workstations are designed to work in heterogeneous, open systems networks with workstations from IBM and other manufacturers. To show the connectivity of their workstations, IBM has set up a laboratory in their Austin facility to continuously exhibit and test interoperability with a diverse number of platforms and operating systems such as their own ES/3090 supercomputers and series 9370 mainframes, VAX and MITEK minicomputers, personal computers, and other workstations from SUN Microsystems, Hewlett-Packard, Digital Equipment Corp. and NeXT Inc. The ability of these systems to interact with one another and other vendors is possible through the ongoing support of IBM's Token-Ring Network protocols and through its support of industry standards such as X Windows, TCP/IP, NFS and Ethernet. IBM is now involved in developing the Fiber Distributed Data

Interface (FDDI), the ANSI X359.3 draft standard for 100 Mbit/sec fiber-optic LANs. The RS/6000 workstations already include a high throughput optical interface to take advantage of this next level of communications bandwidth. This, again, is indicative of IBM's commitment to the UNIX workstation community.

At Sandia, we were the first to test the RS/6000's connectivity to the Cray supercomputers running UNICOS, Cray's variant of UNIX. After sorting out gateway, router and administrative hurdles, we were successful in connecting to the Cray 1/S over 1 MBit Ethernet lines. We were also able to allow seamless communication between the workstation and the Cray by creating and compiling an AIX terminal definition file on the Cray. The process of NFS-mounting directories residing on the Cray onto the workstation was relatively easy on the IBM workstation. This makes user files on the Cray transparently available to a user on the workstation.

Interoperability with the VAX operating under VMS was not as smooth. We had to rely on the ATE terminal emulation program provided with AIX. Although much improved compared to other VT-100 terminal emulators on other workstation platforms, it still felt clumsy. Logging on to the VAX and executing commands did not present any difficulty, but editing on the VAX through the workstation was almost useless because of the inappropriate keyboard mapping. We hope IBM provides a better VT100 terminal emulation program with the final release of its operating system. However, this may not be necessary if VMS 5.1, which supports a Wollongong implementation of TCP/IP, is installed⁵ or VMS is replaced with ULTRIX, DEC's UNIX implementation, on our VAX machines. Nevertheless, since most of the VAX clusters at Sandia are not expected to migrate to ULTRIX any time soon, VMS terminal

5. Prior to this report's release, but after the Model 530 evaluation period, Wollongong TCP/IP was installed on the Area-4 unclassified VAX cluster.

emulation should be addressed. We expect third party developers to get involved in this issue. As far as networking, a key characteristic of IBM's RISC System/6000 family is the capability to link up to 250 workstations to a single server.

Overall, we were satisfied with the connectivity of the IBM workstation to our mainframes at Sandia. We believe that the RS/6000 family of computers will enable us to do more whether operating stand alone, interconnected with a mainframe, or sharing the resources of a distributed computing environment.

4.5 SECURITY AND DATA INTEGRITY

In our present mode of operations, however, an important consideration is network security. We have a constant need to access the Cray supercomputers through the Supercomputing Secure Network (SSN). The increased functionality that a UNIX operating system provides carries with it responsibility and commitment. To take advantage of a distributed computing environment - i.e. more efficient interconnectivity, localized computing resources, uniform user interface, and transparent file sharing - implies that most users will control their own workstation accounts. Consequently, this raises the question of verification and accountability. Faced with potential security ramifications, DOE has required some type of network and user node authentication system for secure UNIX networking. A network user authentication system developed as part of MIT's Athena Project, Kerberos, was chosen by DOE as the preferred method of packet labeling and user authentication. Network utilities that have been converted to use Kerberos are said to be "kerberized".

IBM has ported Kerberos to the Andrew File System (AFS) developed by Carnegie Mellon University, but not to Sun's Network File System (NFS). NFS is the network file system that will be supported at Sandia. IBM, having funded part of the Athena

Project, has kerberized the communication utilities on their AIX version for the RT and PS/2 computers. Therefore, we do not anticipate any difficulty in having AIX 3.0 modified to understand Kerberos. However, the kerberization of NFS for AIX 3.0 (and a commitment to the kerberization of all future versions of AIX that we may implement) remains a very important criterion that, if not met, may prevent us from acquiring IBM workstations. Although stop-gap measures can be implemented, they are temporary and add another layer of complexity to an already complex network environment. We wait to see if IBM rises to the challenge.

It is important to note that AIX Version 3 for the RISC SYSTEM/6000 meets the National Computer Security Center (NCSC) Trusted Computing Base Level C2 security requirements. In the future, IBM intends to provide functions designed to meet the mandatory access control policies defined by the current criteria for B1 level systems and submit these to the NCSC for B level evaluation

In addition to passwords, physical security includes a hardware lockable console and keyboard. Since we do not anticipate using these machines for classified work, these measures should suffice.

To maintain data integrity, IBM has included an integrated error-logging facility that captures error data from hardware and software in a standardized format. This feature, called "synchronous exception handling", is a mechanism that provides for the simultaneous reporting of exception events within the bus cycle causing the exception. This then enables the operating system to resolve the exception more quickly and easily. Errors are automatically recorded when they occur to enable analysis as part of the problem determination process. Concurrent error notification and reporting enables background notification of multiple services when an error is logged and enables the reporting of each error as it is logged. IBM also says that the

system dump utility now includes the traditional kernel data plus some other data. The dump can be triggered by exception and formatted in either batch or interactive mode through the crash command. A dump can be initiated remotely. The System Event Performance Trace has a menu interface for starting and stopping the trace and generating trace reports. Process names are printed, elapsed time is reported between events, and the output is column aligned. In addition, memory scrubbing is systematically performed to avoid double bit errors. As an added reliability, all processor chips engage in extensive self-test sequences when the machine is powered on.

AIX also has a storage management utility that automatically extends storage for file systems. The storage management utility can extend storage across fixed disk boundaries to support larger file systems, maintain 2 or 3 copies of mirrored files on different disks, and automatically reallocate bad blocks. The file system implements database techniques when updating file system metadata to increase file system reliability. A mechanism is provided such that users do not have to rebuild a file system if the system crashes or the user has not initially allocated enough space. IBM claims that this reduces recovery time by limiting checking to the number of changed file system control structures rather than the entire file system. A mechanism called "bit steering" is implemented by adding redundant data and address parity bits for improved system reliability and data integrity.

4.6 Ergonomic Design

While technical performance is what characterizes a workstation system, it is not the only significant consideration. What makes a difference, we believe, is the usability of a system as a whole - support, system ergonomics, and the flexibility of

its integrated environment to be tailored to an individual user's preference.

For example, the IBM workstation is very quiet and never showed signs of overheating even after several days of continued operation. IBM ascribes this to a new breakthrough in their cooling system which includes a high-volume fan coupled with an array of 2-inch heat sinks to keep the CPU chips cool. The 19" color monitor, however, generated a noticeable amount of heat that is typical of large color screen displays. The windowing environment is well integrated providing relatively easy access for customization, and the clarity of graphics on the 19-inch high performance color monitor is superb. Software to simulate PC DOS, the widely known operating system on IBM compatible personal computers (PC), proved very helpful in moving files between the workstation, VAX, and other PC's.

Our criticism of the workstation's ergonomic design are few. Two that stand out are an oversensitive mouse pointing device and a noisy quarter-inch tape drive. Although these are minor concerns, they can be quite irritating. It would be wise for IBM to pay more attention to such details and continue improving the overall ergonomic design of the RS/6000.

4.7 CUSTOMER SUPPORT

IBM has also implemented a wide range of support which is one of the industry's best. The one-year warranty includes 24-hour, seven-day-a-week availability of hardware service. Full software service is included in the licensing fee. In addition, System Xtra, a fee service providing comprehensive support after installation, is an option. This includes toll-free telephone hot line support for hardware and software problems, on-site installation of software maintenance, and remote services which allows customers to dial in from terminals to receive information

from IBM's system. An IBM field engineer is also available to come to the customer's site to configure the network, install machines, and make sure all systems are integrated, whether or not they are made by IBM.

A program called AIXpert is also available to software developers for information on AIX Version 3, and to provide the necessary resources to assist in the porting process and two-way communication with IBM.

4.8 FUTURE DIRECTIONS

IBM claims that they are committed to doubling the performance of the RS/6000 family of workstations every twelve months, while more or less retaining price brackets [UPDA8].⁶ By 1993, these machines' performance may well exceed 120 MIPS with high end members hovering around 200 MIPS. The capability for multi-processing is also built into the RS/6000 architecture and hardware technology. The first workstations to take advantage of this could arrive in the market as early as 1993 (Parallel processing applications could be of interest to us by then.) By 1994 system performance yield over 1 BIPS (1,000 MIPS) can be reasonably anticipated, with comparable levels of floating-point performance to be competitive with Cray-type supercomputers available at that time [JEFF].

The main driving forces behind this leap are the scalability of the RS/6000 architecture and the breakneck pace in which higher-density submicron semiconductor chips are developed. The current state-of-the-art dynamic random access memory (DRAM) chip is the 4-million-bit chip. IBM has fabricated a 16-million bit

6. The processor chip on the RS/6000 is located on a plug-in card. This provides the potential for an on-site CPU change and upgrade as faster processor chips and boards become available.

DRAM, and as of this January, they have started in the development of a 64 megabit DRAM chip that they expect to be a world standard in the mid-1990's [UPDA3]. IBM is very likely to succeed in this venture since, among workstation vendors, they are the only ones who have the resources to develop and manufacture their own computer chips. It is clear that IBM is positioning the RS/6000 to establish an entirely new line of IBM products ranging from the workstation to the supermainframe level [CER].

5.0 CONCLUSIONS

Our investigations, consistent with other studies, showed that the IBM POWERstation model 530 is at least 4 times faster than the 8650 VAX, and at most 50% slower than the CRAY XMP/416 supercomputer. For some cases, we found that this IBM workstation approaches 80 to 90% of the Cray performance. The architecture of the RS/6000 not only emphasizes performance and upgradability, but security, error detection, and error recoverability as well.

We believe that the RISC System/6000 offers a competitive price-performance ratio compared to mainframe computers from Cray and Digital's VAX. If present performances are any indication of what is to be expected, we are very optimistic on the possible impact these machines can have on our present applications and future plans. It may be possible to justify the investment on these workstations by anticipated reductions in, for example, our yearly Cray-time fees.

6.0 REFERENCES

- [ANSI] American National Standard Institute, Inc., "American National Standard Programming Language FORTRAN, ANSI X3.9-1978", 1978.
- [BERG] B. A. Bergman, "RISCy Business," Embedded Systems Programming, November 1990.
- [BROW] D. Brown, "IBM's Drive Into Technical/Engineering Markets," Technology Trends, February 26, 1990.
- [BROW2] "Hot Technical Workstations," Technology Trends, April 26, 1990.
- [CADR] "IBM Gets Serious," Computer Aided Design Report, Vol. 10, No. 3, March 1990.
- [CAE] "Workstation Update", Computer-Aided Engineering, July 1990.
- [CER] "INSIDER'S VIEWPOINT: IBM's RS/6000 - A Strategic Report", COMPUTER ECONOMIC\$ Report, Vol. 12, No. 4, April 1990.
- [CRN] "Fastest Machine Ever Tested By Workstation Labs," Computer Reseller News, August 20, 1990.
- [DALY] J. Daly, "Can RS/6000 Outshine Supers?," Computerworld, March 26, 1990, p. 126
- [DEW] R. K. Dewar, M. Smosna, "Microprocessors: A Programmer's View," McGraw-Hill, NY, 1990.
- [DEW2] R. K. Dewar, M. Smosna, "Rating The New RISC Chips: IBM's RS-6000 Microprocessor Is A Step Ahead of the Rest--But For How Long?," UNIX Today, May 14, 1990.
- [DONG] J. J. Dongarra , " Performance of Various Computers Using Standard Linear Equations Software," CS-89-85, Computer Science Department, University of Tennessee, Knoxville, TN; Mathematical Sciences Section, Oak Ridge National Laboratory, Oak Ridge, TN, September 1990.
- [DVOR] J. C. Dvorak, "Inside Track", PC Magazine, May 15, 1990.
- [ENTI] N. Enticknap, "Measuring Up A Star Performer: Benchmark Results for IBM's RS6000 Workstation," Computer Weekly, March 8, 1990
- [EUNI] J. Eunice, D. Brown, "UNIX Forges Ahead," Technology Trends, February 15, 1990

- [GRIF] Private Communication from Pat Griffin, Sandia National Laboratories.
- [GROS] O. Grossman, "Sun Microsystems Redoubles SPARC Effort", PC Week, April 23, 1990.
- [HALB] J. A. Halbleib, R. P. Kensek, T. A. Mehlhorn, G. D. Valdez, "ITS VERSION 3.0: The Integrated TIGER Series of Coupled Electron/Photon Monte Carlo Transport Codes," Sandia National Laboratories, Albuquerque, NM, (To be published, 1990)
- [HASK] D. W. Haskin, "IBM's RISC Lineup: Packed With Power," Digital Review, October 1, 1990.
- [IRLAM] G. Irlam, "A Guide to Sun-4 Virtual Memory Performance," Sun-Spots Digest, Vol. 9, Issue 257, July 1990.
- [JEFF] B. Jeffrey, "What's Ahead for IBM's RISC System/6000," FOCUS, May 1990, p. 11
- [LEVITT] J. Levitt, "A Look At IBM's New Filesystem," UNIX Today, October 1990.
- [LORE] L. J. Lorence, J. E. Morel, G. D. Valdez, "User's Guide To CEPXS/ONEDANT: A One-Dimensional Coupled Electron-Photon Discrete Ordinates Code Package," SAND 89-1661, Sandia National Laboratories, Albuquerque, NM, September 1989.
- [PER] Private communication from N. R. Pereira, Berkeley Research Associates, Springfield, VA, June 12, 1989.
- [SIMM] M. L. Simmons, H. J. Wasserman, "Los Alamos Experiences with the IBM RISC SYSTEM/6000 Workstations," LA-11831-MS, Los Alamos National Laboratory, Los Alamos, NM, March 1990.
- [SOUT] A. Southerton, R. Farrow, "IBM Blazes Into UNIX," UNIX World, April 1990.
- [SYSI] "Big Blue's Percentage RISC," Systems International, April 1990.
- [UPDA1] "Benchmark Substantiates 'Most Powerful Workstation' Claim," UPDATE, March/April 1990.
- [UPDA2] "In LINPACK Analysis, Performance Talks; Hype Walks," UPDATE, March/April 1990.
- [UPDA3] "Coming Soon: the 64-Megabit DRAM," UPDATE, March/April 1990.

- [UPDA4] "These Are 'Breakaway Products'," UPDATE, March/April 1990.
- [UPDA5] "Applications Ported with Panache," UPDATE, March/April 1990.
- [UPDA6] "RISC System/6000 Technology Breaks New Ground," UPDATE, March/April 1990.
- [UPDA7] "AIX V3 Fits RISC System/6000, Other IBM and Non-IBM Systems," UPDATE, March/April 1990.
- [UPDA8] "RISC System/6000: The Workstation for the 90's," UPDATE, May 1990.
- [UPEML] T. A. Mehlhorn, M. F. Young, "UPEML Version 2.0: A Machine-Portable CDC Update Emulator," SAND87-0679, Sandia National Laboratories, Albuquerque, NM, August 1987.
- [WILS] D. Wilson, "Tested Mettle: The IBM RISC SYSTEM/6000 POWERstation 320," UNIX Review, October 1990.

APPENDIX A

MACHINE DESCRIPTION

The RISC SYSTEM/6000 (RS/6000) computers employ a highly concurrent, superscalar (defined below), second generation Reduced Instruction Set Computer (RISC). This series includes the POWERstations and POWERservers, all of which use IBM's POWER architecture (Performance Optimization with Enhanced RISC). What sets this architecture apart from any other is the implementation of RISC as maximized-parallelism of three units: fixed-point, floating-point, and branch, as well as the high bandwidth of the paths between them. The processor contains the 3 processor chips, each handling a specific portion of the instruction set, thus eliminating the time wasted when signals are shuttled back and forth between chips. IBM claims that they can double the performance of the processor every 12 months. The floating point unit performs up to two 64-bit operations simultaneously in a cycle. The branch unit grabs instructions, executing some and passing the rest to either the fixed- or floating-point units. The pipeline implementation (known as "superscalar") allows execution of up to five operations in a single cycle, a new industry standard.

Workstations in the series use either a 64 or 128-bit memory bus that can move data to and from memory at speeds up to 480 Mbytes/second. The new systems also sport IBM's enhanced and proprietary Micro Channel bus, with throughput of up to 40 Mbytes/second. During our evaluation period, the machines comprising the series were capable of throughputs ranging from 27.5 to 41 MIPS.⁷

7. System improvements prior to official market release have increased the MIPS ratings to 29.5 and 44.3.

IBM claims the desktop POWERstation 320 rates at 27.5 MIPS and 7.4 MFLOPS for a system starting at \$12,995. At 20 MHz, the 320 includes a system unit with two memory slots and four I/O slots, 8 Mbytes of RAM (expandable to 32 Mbytes), and 120-Mbyte DBA(direct bus attached)/SCSI storage (expandable to 640 Mbytes).

The POWERstations 520 and 530 are both deskside systems that expand the number of memory slots and 32-bit Micro Channel slots to eight per system. The 520 has the same processor clock speed as the 320, but offers further expandability as the 530. These include maximum capacities of 128 Mbytes of RAM and 2.5 Gbyte of fixed disk storage. The 20-MHz 520 comes standard with 8 Mbytes of RAM and 355-Mbyte hard drive, and the 25-MHz 530 includes a base of 16 Mbytes of memory and a 355-Mbyte hard drive for storage. The availability of 4 MBit CMOS DRAM chips in the fourth quarter of this year will quadruple the maximum memory limit for every member of the RS/6000 series except for the Model 540 which will ship with the 4 MBit chips.

All of these POWERstations include a 19-inch 1280 x 1024 monochrome display, 16 gray-scale video adapter, Ethernet, keyboard, three-button mouse, AIX, graphical user interface, software service, and a one-year hardware and software warranty. A choice of 16, 19 or 23-inch color monitor can be used to upgrade the monochrome display along with color video adapters capable of 16.7 million colors and 24-bit "true color" graphics.

The RS/6000 model 530 evaluated came equipped with 48 Mbytes of RAM, 8-bit color graphics adapter, 1.4 Gbytes of fixed disk storage, 3-1/2 inch 1.44 Mbyte and 5-1/4 inch 1.2 Mbyte internal diskette drives, internal CD-ROM drive, 150 Mbyte external quarter-inch tape drive, three-button mouse, and a 19" 1280 x 1024 high performance color monitor. The latest AIX 3.0 operating system and Fortran compiler evaluated was the 9013 pre-release version.

We are evaluating the ability of the model 530 POWERstation to be the division's code development platform. In addition, we are also considering acquisition of model 320 POWERstations for less compute intensive applications. A model 530 POWERserver would also be useful in our network to interface with the Internal Secure Network (ISN), maintain software resources, provide additional computational capability, and to serve as a boot node for several diskless workstations which may operate within our network. Machine specific run-time modules and recompilation is unnecessary since the RISC System/6000 family workstations are both source- and object-code compatible.

Both the POWERstations and POWERservers use the new AIX version 3. For a graphical user interface (GUI) running over AIX, IBM offers either AIXwindows Environment/6000 using Motif, or AIX NextStep Environment/6000 using the NextStep GUI from NeXT, Inc. Recent advertisements placed by Sun Microsystems in several computer journals announced the availability of the Open Look Interface (the standard GUI for AT&T's UNIX) on the IBM platform. The family of workstations and servers are also capable of optimized 3-D graphics capabilities.

APPENDIX B
BENCHMARK CODES DESCRIPTIONS

ITS

The Integrated TIGER Series consist of time-independent coupled electron-photon Monte Carlo transport codes that provide a state-of-the-art description of the production and transport of the electron-photon cascade in multimaterial and multidimensional geometries. These continuous-energy Monte Carlo codes combine microscopic photon transport with a macroscopic random walk for electron transport. TIGER, CYLTRAN, and ACCEPT are the base codes of the series and differ primarily in their dimensionality and geometric modeling. TIGER is a one-dimensional multilayer code. CYLTRAN employs a fully three-dimensional description of particle trajectories in a cylindrical geometry, and quite naturally finds application in problems involving electron or photon beam sources. ACCEPT is a general three-dimensional transport code that employs a sophisticated combinatorial geometry scheme.

These base codes were primarily designed for transport at energies that are large compared to the binding energies of the atomic electrons. The "P Code" versions of these ITS codes were developed for low-energy transport applications. In TIGERP, CYLTRANP, and ACCEPTP, more elaborate ionization and relaxation models are used. In the "M Code" versions of CYLTRAN and ACCEPT, collisional transport is combined with transport in macroscopic electric and magnetic fields of arbitrary spatial dependence using a Runge-Kutta-Fehlberg algorithm to integrate the Lorentz force equations.

TIGER is the simplest of the ITS Monte Carlo codes. ACCEPTM is the most complex. An additional member of the ITS code package is the XGEN code which provides the electron-photon cross section information that is needed by the Monte Carlo codes.

CEPXS/ONELD

CEPXS/ONELD is a discrete ordinates code package for one-dimensional coupled electron-photon transport over the energy range of 100 MeV to 1.0 KeV [LORE]. This suite of codes consists of the cross section generating code, CEPXS; the preprocessor code, PRE1D; the discrete ordinates code, ONELD; and the postprocessor code, POST1D.

CEPXS is a cross-section generating code. The cross section data it produces are the multigroup-Legendre expansion coefficients of coupled electron-photon cross sections. These coefficients are suitable for use with standard discrete ordinates codes. Discrete ordinates codes solve the Boltzmann transport equation deterministically (i.e. by finite-difference techniques).

ONELD is a general Boltzmann transport code. While it was designed to expedite electron transport calculations, the code can also be used for neutron and photon transport. The ONELD discrete ordinates code employs linear-discontinuous spatial differencing and S2 synthetic acceleration.

APPENDIX C
OPEN SOFTWARE FOUNDATION

The Open Software Foundation (OSF) is a computer industry first: a non-profit corporation jointly established by the industry's leading vendors as a neutral supplier of open systems software technologies. It's mission is to be the leading supplier of a complete, innovative and equitable open systems environment to the worldwide computer industry.

OSF provides software that addresses 3 key user needs:

- portability -- which allows application software to run on multiple vendors' systems;
- interoperability -- which allows systems from different vendors to work together transparently; and
- scalability -- which allows a software environment to be hosted on the entire range of available hardware platforms, from desktop devices to supercomputers.

To meet these goals, OSF uses an innovative, open process to solicit, acquire, and develop technology. OSF adheres to 7 guiding principles:

- Offerings based on relevant industry standards
- Open process to actively solicit inputs and technology
- Timely, vendor-neutral decision process
- Early and equal access to specifications and software
- Hardware-independent implementations
- Reasonable, stable terms for licensing
- Technical innovation through universities and research

History and Membership

OSF was formed in May 1988 with the support of 7 original sponsors: Apollo Computer Inc., Digital Equipment Corporation, Groupe Bull, Hewlett-Packard, International Business Machines Corporation, Nixdorf Computer AG, N.V., and Siemens AG. Subsequently, Philips Telecommunications and Data Systems and Hitachi Ltd. of Japan also become sponsors.

Since that time, the number of members has surpassed 200. They include systems manufacturers, hardware systems vendors, independent software vendors, semiconductor manufacturers, service companies, industry consortia, educational institutions, public and quasipublic entities, and consulting and research firms. Recent additions to OSF's membership roster include Apple Computer, Microsoft Corporation and NCR. Sun Microsystems and AT&T are not members; they have formed their own consortium called Unix International.⁸

Through its members, OSF has a funding commitment in excess of \$130 million through 1991. Membership is open, and all members may participate in decision-making and take advantage of early access to technologies under development.

OSF Research Institute

More than one-third of OSF members are universities, industry research laboratories, or government research institutions. The

8. The main driving force behind OSF's formation is the intent of AT&T and Sun to control the development of UNIX. AT&T's System V has already built up momentum in the marketplace. If OSF was not formed, Unix International would have succeeded in pushing subsequent releases of the UNIX operating system as proprietary, and would have allowed them to get to market months ahead of competitors.

OSF Research Institute establishes strong links with these research organizations to benefit from their experience and technology. Its mission is three-fold:

- To involve the research community worldwide in the evolution of open standards
- To analyze emerging technologies and identify opportunities for future requests for technology or OSF development
- To build a long-term shared vision and architecture for the future of open systems

The OSF Research Institute serves as the link between the foundation and both university and commercial laboratories around the world. To date, the OSF Research Institute has awarded several grants to research organizations, including MIT Project Athena; Trusted Information Systems; Carnegie Mellon University, the University of Guelph, Ontario; the University of Lowell, Massachusetts; Vrije University, and the University of Calgary, Alberta.

END

DATE FILMED

01 / 30 / 91