Shared-Memory Vector Systems Compared

Robert Bell
Deputy Manager, Bureau of Meteorology / CSIRO
High Performance Computing and Communications Centre
Melbourne Australia
Robert.Bell@cmis.csiro.au

Guy Robinson
Research Liaison/MPP Specialist
Arctic Region Supercomputing Center
Fairbanks Alaska USA
robinson@arsc.edu

ABSTRACT: *The NEC SX-5 and the Cray SV1 are the only shared-memory vector computers currently being marketed. This compares with at least five models a few years ago (J90, T90, SX-4, Fujitsu and Hitachi), with IBM, Digital, Convex, CDC and others having fallen by the wayside in the early 1990s. In this presentation, some comparisons will be made between the architecture of the survivors, and some performance comparisons will be given on benchmark and applications codes, and in areas not usually presented in comparisons, e.g. file systems, network performance, gzip speeds, compilation speeds, scalability and tools and libraries.*

1. Introduction

HPCCC

The Bureau of Meteorology and CSIRO in Australia established the High Performance Computing and Communications Centre (HPCCC) in 1997, to provide larger computational facilities than either party could acquire separately. The main initial system was an NEC SX-4, which has now been replaced by two NEC SX-5s. The current system is a dual-node SX-5/24M with 224 Gbyte of memory, and this will become an SX-5/32M with 224 Gbyte in July 2001.

The systems are used by the Bureau for weather forecasting, and by CSIRO for a variety of research: including climate modelling, air quality forecasting, ocean modelling, polymer simulation, antennae design, water percolation in soils, exploration and mining, and industrial fluid flow modelling.

NEC was chosen in a competitive tender, providing the best performance/price ratio for the Bureau operational models, while providing the ease of use of a shared memory system for multi-processing. The systems have to provide immediate running for the Bureau operational jobs, while supporting a heavy load of research work. Recently, processor utilisation has averaged 95% on the larger node, and 92% on the smaller node.

The SXes are supported by data storage systems – SAM-FS for the Bureau, and DMF on a Cray J90se for CSIRO.

ARSC

The mission of the Arctic Region Supercomputing Centre is to support high performance computational research in science and engineering with an emphasis on high latitudes and the Arctic. ARSC provides high performance computational, visualisation and data storage resources for researchers within the Department of Defence, UAF and other academic institutions and government agencies. ARSC is located on the University of Alaska Fairbanks campus. Researchers at the University of Alaska Fairbanks make significant contributions to science on state, national and international levels using ARSC resources and talent.

ARSC operates a 272-processor 450 Mhz CRAY T3E900 system with 68 Gbyte of memory and 522 Gbyte of disk storage and a 32-processor CRAY SV1 parallel vector

system with 4 Gigaword of memory and 2 Tbyte of disk storage. This has been upgraded to 500MHz processors and will shortly be upgraded to full SV1ex capability. All SV1 benchmarks in this paper were run on SV1e processors with SV1 memory. (To be clear . ARSC also maintains a number of on campus access labs which host visualisation hardware including an Immersadesk and several high-end SGI visualisation servers.

Specialists at ARSC provide support in the use of these systems. ARSC's unique relationship with the University facilitates collaborative research opportunities for academic and government scientists. These areas of research include ice, ocean, and atmospheric coupled models, regional climate modelling, global climate change; permafrost, hydrology, and Arctic engineering; magnetospheric, ionospheric and upper atmospheric physics; vulcanology and geology; petroleum and mineral engineering; and Arctic biology. ARSC runs a data storage system using DMF.

Cray and NEC

When the HPCCC was seeking its core systems, the NCAR procurement of 1996 was a hot political issue, and a 454% tariff was imposed on NEC in the US. Outside the US, NEC was making in-roads into Cray's former territory, offering similar peak performance to a T90, at about 1/3 the price. Certainly, Cray's offer to the HPCCC fell short of NEC's offer on peak performance and the sheer scale of the hardware. Software was another matter however.

From 1996 to 1999, the path of Cray has been varied, and the focus on vector systems perhaps became less intense. However, since the acquisition of Cray by Tera, there is again a commitment to vector systems, with the SV1ex being the latest offering. NEC has about 100 vector systems installed worldwide, and has done well, particularly in Europe. The NEC SX-5 processors currently have a peak speed of 10 Gflop/s, and a maximum memory per node of 256 Gbyte, while the SV1ex processors have a peak speed of only 2 Gflop/s and a maximum memory per node of 32 Gbyte. However, the SX-5 peak is achieved by having 16 vector pipes, while the SV1 has only two, but four SV1 processors can be combined in a Multi-Streaming Processor (MSP) with a peak speed of 8 Gflop/s, giving a comparable processor speed to the SX-5.

Earlier this year, Cray and NEC announced an agreement under which Cray would re-sell and support NEC SX-5 and successor systems, and seek to have the tariff barrier lifted in the USA, making the comparison in this paper rather timely for those in the USA who now have an interesting choice of vector systems.

2. Description of systems - hardware

In the following specifications, comparisons will be made between the Cray SV1, and the current model SX-5, rather than the older model as installed at the HPCCC

(which has a lower clock speed and lower peak performance of 8 Gflop/s). Performance results later will be given for the HPCCC machine. An SV1 can be configured with 4-32 processors. Each CPU has 2 add and 2 multiply functional units which can each return one result per clock cycle. Processors can be clocked at 300 MHz or in the latest SV1ex model at 500 MHz giving peak processor performance of 1.2 or 2 Gflop/s. An additional feature of the SV1 is that 4 processors can be combined to create an MSP. The four processors act as one, with interrupts disabled, which permits vectors to flow continuously. Configuration between single processor and MSP operation is dynamic.

Both the NEC SX-5 and the Cray SV1 are shared memory vector supercomputers. The following table shows some of the features of the SX-5 and SV1 processors.

System	Clock speed MHz	vector length	vector pipes	peak speed Gflop/s	Max CPUs
SX-5/A	312.5	256 or 512	16	10	16
SV1	300	64	2	1.2	32
SV1ex	500	64	2	2	32
SV1ex MSP	500	64	8	8	6

Table 1. Processor comparison

(There are other model SX-5s, with half the peak speed.) The numbers of processors shown above are for a single system with shared memory. Details of the memory are shown below.

System	Maximum memory Gbyte	max SSD Gbyte	Bandwidth per CPU Gbyte/s
SX-5	256	nil	80
SV1e	32	96	3.2

Table 2. Memory comparison

The SX-5 memory (64 Mbit 47 ns SDRAM) is divided into a large number of banks – 16384 for the 128 Gbyte system at the HPCCC, with 1 Tbyte/s bandwidth. (Many banks are needed to get the speed from the slower memory in the SX-5 compared with the SX-4.) The SV1 memory architecture is a uniform access, shared central memory. Capacity ranges from a minimum of 4 Gbyte to a maximum of 32 Gbyte, with the option of up to 96 Gbyte of SSD on the SV1ex. To move data between CPU registers and memory via the cache two data paths are provided. In any given clock cycle two read or one read and one write can be active; if there are no reads only one write is possible. The SV1 cache size is 256 kbyte, is 4 way set associative, write allocate, write through and with a least recently users (LRU) replacement strategy. Cache line size is 8 bytes, or 1 word. Cache coherency is achieved in software. The bandwidth is

6.4Gbyte/s between main memory and cache, and 14.4Gbyte/s between cache and processor on the SV1ex.

The SX-5 provides IEEE 64- and 32-bit floating point arithmetic, while the SV1 provides Cray floating point arithmetic.

Both systems provide hardware performance registers. The Cray SV1 has 32 performance counters provided in 4 groups of 8. Only one group can be active at any one time. Group0 provided floating point and memory reference data for both scalar and vector operations. Group2 data covers detailed information on processor memory references and include a memory conflict count. Group3 covers vector and scalar components of floating point operations and other vector functional units. The NEC performance counters cover similar areas, but are in a single group.

Multiple systems can be combined to form clusters. Up to 32 SX-5s can be connected through a proprietary crossbar switch providing 16 Gbyte/s per node. The switch allows MPI jobs to access another node's memory directly across the crossbar, and allows various aspects of a single system image to be implemented. SV1 systems can be connected through the GigaRing I/O network at speeds of 1.2 Gbyte/s per connection. The SX-5 has up to four I/O processors capable of 3.2 Gbyte/s each. The SV1 has a distributed I/O architecture, with specialised nodes on the GigaRing network.

Both the SX-5 and SV1 are air-cooled. The SX-5/16A at the HPCCC weighs about 8.5 tonne (excluding disc), consumes a peak of 113 kW of power, and takes a lot of floor space. The SV1 typically occupies about 5 standard racks, and is on wheels.

The NEC vector processor uses many pipes to achieve high peak speed. It has a huge processor to memory bandwidth. However, compared with the earlier SX-4, it is more of a specialised vector processor. The number of vector pipes was doubled, and the peak vector speed increased by a factor of 5, while the scalar speed increased by a factor of only 2.5, and some vector start-up speeds increased by even less.

Key kernel performance.
To gain high performance from the SX-5, a high degree of vectorisation, and long vectors are needed. As well, the SX-4 could do two vector loads and a vector store simultaneously, but the SX-5 cannot overlap stores with loads.

So, with the SX-4 it was relatively simple to demonstrate peak performance. Here are figures for the percentage of peak speed achieved (for sufficiently long vectors) for two simple loops of length 2^{20}.

Loop	$a_i = \alpha b_i + c_i$		$a_i = \alpha b_i + \gamma$	
	Mflop/s	% of peak	Mflop/s	% of peak
SX-4	1335	66.7	2002	100
SX-5	3997	50.0	5350	66.9
Y-MP	194	58.2	289	86.6
J90se	102	51.2	170	84.9
SV1e	102	5.1	152	7.6
SV1e – 1024	392	19.6	408	20.4

Table 3. Kernel comparison

The SX-4 delivers a higher proportion of peak speed than the SX-5. The first SV1e results show poor performance because cache is not effective with long vectors. Another test with a shorter loop length gives better results, but still only around 20% of peak. Peak speed appears to be getting harder to obtain on all architectures.

3. Software

3.1 Operating System
The SUPER-UX operating system is based on System V release 3, with some release 4 features, and BSD additions. It is not fully Posix compliant, nor does it conform to other UNIX standards. It is not fully 64-bit. For example, the shells do 32-bit arithmetic, and the cpio command failed to handle files larger than 2 Gbyte, and returned an incorrect block count when the archive was larger than 2 Gbyte, but these problems were later corrected.

UNICOS is based on System V release 4, and conforms to Posix, XPG4 and other standards, and includes BSD and AT&T extensions. It is fully 64-bit capable (or at least 48-bit), and has been re-written from the ground up to support high performance. For example, it appears that buffer sizes are automatically selected to give good performance. As an example, consider the times for a cat command which joined two 160 Mbyte and one 80 Mbyte file to form a 400 Mbyte file. (These were not run in dedicated mode, so the elapsed time is not particularly significant).

Machine	User CPU s	System CPU s	Elapsed time s
J90se DD-308	0.025	6.61	136.
SX-5 MRFS	0.39	64.43	89.
SV1e /tmp2	0.031	7.61	30.

Table 4. Timing of cat command

The SX-5 runs were done in a Memory-resident File System, something the large memory of the SX-5 makes feasible! Note that the system time for this command processing 800 Mbyte is over a minute for the SX-5 – nearly ten times as long as on the J90se.

UNICOS is multi-threaded. SUPER-UX is not, which has implications for scalability for larger numbers of processors, and leads to higher system time.

SUPER-UX supports the hardware, and contains many features for high performance, such as support for disc striping and the MRFS. UNICOS also contains these features.

3.2 File systems

SUPER-UX supports the Supercomputer File System (SFS) and the Hybrid extension (SFS/H). These support striping and various options for caching, etc. Speeds of 75 Mbyte/s are achieved on HiPPI discs. UNICOS supports the NC1 file system. This provides high speed on a number of devices. The file systems are not limited to 2 Gbyte or 2 billion files.

However, the NC1 file system has one crucial advantage – it supports primary and secondary partitions with different allocation unit sizes, and UNICOS intelligently works out when to use the right partition size. SUPER-UX file systems have a fixed cluster or allocation unit size for each file system. This is severely limiting. If you want high performance, then cluster sizes of over 1 Mbyte are recommended, and you can't support many files. If you want to support many files, then you need a smaller cluster size (4, 128, 256 kbyte, etc), and the performance is poor. SFS supports a re-allocation facility for large-cluster file systems, where small files can be packed into a large cluster, but this is supported only for file systems smaller than 62 Gbyte.

Neither NC1 nor SUPER-UX supports dynamic re-configuration of file systems – if they are to be expanded or shrunk or moved, the file systems have to be dumped and reloaded. Neither support log-structuring or snapshot capabilities.

3.3 Compilers, tools, libraries, multi-processing support

Both vendors provide compilers for C/C++ and Fortran. Cray supports full Fortran 95, while NEC expects to reach that standard with the June release. Both vendors' Fortran compilers provide support for automatic vectorisation and parallelisation, and for parallelisation driven by compiler directives. Both vendors provide support for MPI, OpenMP and HPF for parallelisation.

In general, the NEC compilers are behind the Cray ones in development - the Cray Fortran compiler superseded the Fortran 77 compiler in 1997. The NEC transition was not until 1999 with the SX-5, and the HPCCC has over the last year provided 20% of the fault reports on the Fortran 90 compiler. NEC provides cross-compilers and a development, debugging and tuning environment (called PSUITE) which runs on workstations (including Linux).

Both vendors provide tools to access the hardware performance registers. Both compilers provide a significant amount of information regarding performance. Gaining information from the compiler about vectorisation and the parallelisation of loops is very important for scientific users.

HPMFLOP permits a Cray centre to monitor user code performance. Quoting from the man pages.
The hpmflop command reports the average megaflops achieved by user processes, based on statistics gathered by special instrumentation from the Hardware Performance Monitor (HPM). These statistics are optionally gathered on a site-wide basis for all normally terminated user programs on the machine. Review of the output of this command allows the site to determine which users and which of their programs might be using large amounts of CPU time with low Megaflop rates. Thus, these users can be contacted and encouraged to optimise their programs. ARSC takes exactly this policy and by contacting users has made improvements to performance and other aspects of users activity on the system.

NEC provides a similar capability for users with the setting of environment variables, and the site capability with the SystemScope utility.

Other utilities provide profiling. UNICOS comes with Totalview for debugging. SUPER-UX includes pdbx, with Totalview available in a beta release from a third-party vendor.

Both vendors provide standard libraries with BLAS, LAPACK, etc. NEC is now making available an enhanced library as a separate chargeable item with SX-5 and multi-processing support.

3.4 3rd party applications

In the past, Cray Research listed about 500 codes in its Applications Software Directory. The current Cray WWW page lists only three – AMBER, Gaussian and MSC.Nastran. SGI's current WWW page lists 3600, with a sub-category "Cray Products" which lists 323 applications. NEC's WWW page lists 84, a few of which are NEC products. Both CSIRO and ARSC are seeing a lessening demand for packaged software, but a greater community involvement in software.

3.5 Network issues

UNICOS supports HiPPI and 100baseT connections. A Gigabit Ethernet is supported through a third-party product connected to a HiPPI node. SUPER-UX supports HiPPI and 100baseT and Gigabit Ethernet connections. Jumbo frame support is now available.

Testing of file transfers across the HiPPI links at the HPCCC were often limited by the underlying file systems. The best transfer rates found between an SX-5 and a J90se

were about 38 Mbyte/s (from a MRFS to /dev/null, using tuned ftp).

When the HPCCC went into production on its earlier SX-4 (32 CPUs), it found that many of the multi-processor operational jobs, despite having the highest priority, were showing very poor scalability, with the elapsed times for runs often being double the time in dedicated mode. After much investigation, it was found that the HiPPI channels on the machine were mapped to fixed CPUs. The channel connecting the SX-4 to a GRF-400 and then to the Bureau's Ethernet network was tied to CPUs 22 and 23, and these showed very high system time. Furthermore the scheduler favoured the higher numbered CPUs for multi-processor jobs. When the scheduler assigned say 16 CPUs to a job including CPUs 22 and 23, then network traffic (with 1500 byte MTUs) would generate a high rate of interrupts on these CPUs. We found that often the other 14 CPUs would have finished a nicely load-balanced parallel region, but CPUs 22 and 23 were chugging along suffering high rates of interrupt until the next barrier was reached. We swapped channels on the SX-4 so that the CPUs servicing the Ethernet were mapped to low numbered CPUs. Later tests with a loopback on an SX-5 showed that a CPU would be saturated at about 8 Mbyte/s. The HPCCC is now working toward connecting front-ends and the Bureau's file servers via HiPPI, to overcome the interrupt problem. Clearly, an integrated front-end to handle network traffic would be desirable.

3.6 peripherals - disc, tape

Cray supports its own disc products, via the GigaRing I/O subsystem. These are usually connected with Fibre Channel. NEC supports SCSI, HiPPI and Fibre Channel discs, all RAIDed.

Cray supports StorageTek, IBM, and Quantum tape drives. NEC supports some StorageTek, Sony and Quantum drives. Both support auto-mounters, though some of the SX capabilities were developed specifically for the HPCCC.

4. Operations

4.1 Resource allocation

In a production environment with a high degree of demand, resource allocation is important. This is particularly the case when there are operational demands, as at the HPCCC for weather forecasting.

Both NEC and Cray provide versions of NQS. An annoying part of the conversion from UNICOS to SUPER-UX was the need to reformat all NQS directives – from the form #QSUB to the cryptic #@$. Both systems provide job checkpoint and restart. However, SUPER-UX does not provide periodic checkpointing to protect jobs from unscheduled interrupts, and checkpoint/restart is less comprehensive and resilient.

UNICOS incorporates the Fair Share Scheduler for apportioning processor time in accordance with a hierarchy of set shares. UNICOS has also provided a number of systems for the political decisions about which job to run next – Fair-share NQS and the Unified Resource Manager.

SUPER-UX has facilities for subdividing resources called Resource Blocks and Resource Sharing Groups. These allow resources such as memory, swap space and numbers of CPUs to be apportioned between groups, with a minimum, intended and maximum value assigned for each resource for each group. The system is rather inflexible. The HPCCC requested a political scheduler in its contract with NEC, and the Enhanced Resource Scheduler (ERS) was developed and is in production there. It provides rapid response to the arrival of operational jobs, and implements the FSS for deciding which jobs to start or hold.

SUPER-UX provides a gang-scheduling feature, which has been found to be effective in preventing large blow-outs in processor time when running multi-processor jobs in busy shared environments. However, it is hard to manage the use of gang scheduling with resource blocks. HPCCC has experienced several occasions when jobs were requesting perhaps four CPUs, and using only one, leading to large wastage.

Both systems provide file quotas. SUPER-UX does not have a User Data Base system.

4.2 Data management, backups

SUPER-UX supports an HSM called SX-BackStore. Enhancement of this product was part of the HPCCC's contract with NEC, and NEC delivered many enhancements (with features beyond DMF in some cases, e.g. stub support). However, there were delays in delivering an acceptable standard of performance for production running, especially for a restore from dump and for file system checking, and the HPCCC has declined to accept SX-BackStore. NEC has recently announced a marketing agreement with Unitree, using IA-64 servers. UNICOS supports DMF, which has been in production use for well over a decade. It is in use at both the HPCCC and the ARSC.

Both systems provide standard dump and restore facilities. However, because of the lack of suitable tape-driving facilities in the early stages, the HPCCC implemented backups to the CrayJ90se using cpio for some of the user file systems on the SXes. Later it was found that the user file systems were too large and busy for the dump command to ever complete.

5. Applications performance

5.1 Compilation

Compilation speed is slow on the SXes. The table below shows a typical sample of compilation speeds for a Fortran program.

System	User CPU s	System CPU s	Elapsed time s
SX-5	29.76	36.57	81
J90se	39.29	16.15	63
SV1e	14.19	11.02	36
Dell 1GHz PIII (cross)	3.42	0.75	16

Table 5. Timing of compilations

Note that compilation on an SX-5 is slower than an old J90se, but compilation on the Intel platform with the NEC cross-compiler is an order of magnitude faster.

5.2 Execution - single CPU - scalar

The results above for compilation speeds give some idea of scalar performance. It is clear that intensive scalar work is not a good use of a supercomputer, but it is still necessary to have a good scalar speed to support the vector processor. This was one of the lessons of the early CDC vector machines, e.g. the Star 100.

Users will persist in compressing files on supercomputers! Here are some results from a gzip command on a 136 Mbyte file.

System	gzip CPU s	gzip elapsed s	gunzip CPU s	gunzip elapsed s
SX-4	1380	1980	127	160
SX-5	660	845	45	88
J90se	2320	2380	230	240
SV1e	81	103	35	44
Sun Ultra-2	225	320	26	65
Dell 1GHz PIII (NFS)	38	260	4.8	252
Dell 1GHz PIII	34	34	4.4	4.6

Table 6. Timing of gzip and gunzip

The supercomputers are not designed for this type of scalar byte-oriented work. The J90se is particularly slow, and the SV1e beats the SX-5. But fastest of all is the Intel machine, and it is faster to gzip on such a machine with files NFS mounted from the J90se than to gzip on the J90se itself.

5.3 Execution – vector and parallel

Below are some results from executing a code which was designed to try to gain the maximum speed from an SV1 type of processor, i.e. it is vectorisable, but seeks to be cache-friendly. Both vendors' compilers were able to vectorise and automatically parallelise this code.

Mflop/s per CPU				
System	1 CPU	2 CPUs	4 CPUs	1 MSP
SV1 PE 3.4	1031		919	838
SV1 PE 3.5	1060		932	994
SV1e PE 3.4	1751		1449	1345
SV1e PE 3.5	1786		1495	1622
SX-5	3281	3111	2849	

Table 7. Timing of cache_test

There is about a 70% speedup from the SV1 to the SV1ex for single CPU executions, and about 60% for multiple CPUs. The PE 3.5 version provides a modest improvement on most tests, but a 20% improvement for the MSP. The MSP provides about a 10% improvement on 4 non-MSPs on the SV1ex. On this test, the SX-5 single CPU provides less than double the speed of the SV1ex (but the later model would be over double the speed). However, it is interesting to compare the latest MSP performance (6.49 Gflop/s) with the 2-CPU SX-5 performance (6.22 Gflop/s). This puts a 32-SV1ex using MSPs on a par with the SX-5 on this test, but about 20% behind the latest model SX-5.

Speed-ups over the single CPU case were: 3.35 for the SV1e with 4 CPUS, 3.48 for the SX-5, and 3.63 for the MSP. These all show good speed-ups, with the MSP showing the best result.

6. Usability

6.1 Documentation

Documentation of UNICOS is first class, with well-written, comprehensive and consistent man pages, the explain facility, and enhanced WWW-format manuals. SUPER-UX documentation is far poorer, with the language barrier being a difficulty. WWW documents are available, and are of a higher standard than the man pages.

6.2 Error handling

The SUPER-UX operating system is poor at handling exceptions in some cases. For example, the system failed to generate error messages or set status flags when a tar command hit a time limit, and files were lost because the tar archive was assumed to be complete. UNICOS implements a consistent error flagging mechanism, with coded error messages, and the availability on-line of further explanation through the explain utility.

6.3 Vendor suppport

The HPCCC has found that NEC gives excellent support through on-site support staff and the developers in Tokyo. Urgent problems are investigated within hours, and fixes delivered in days. Cray's controlled software release mechanisms provide higher quality control, but slower response. As new technologies are rolled out ARSC has received excellent support from Cray in bringing these into the production environment.

7. Conclusion

7.1 Price

Price comparison is a difficult issue, partly because vendors are loath to disclose prices, and comparable configurations are hard to prepare. A 32-CPU 32 Gbyte plus 32 Gbyte SSD SV1ex with a peak speed of 64 Gflop/s might be about US$3.5M. An 8-CPU 64 Gbyte SX-5 with a peak speed of 80 Gflop/s might be about US$4M. The SV1ex is likely to out-perform the SX-5, and is more cost-effective, but cannot attain the same peak speed or provide the enormous memory.

7.2 Sustained performance measures

Sustained performance is hard to define – it depends on the application. For the vector-cache-friendly code, particularly using the MSPs, the SV1ex gains twice the performance of the SX-5 relative to peak. However, the SX-5 provides high memory bandwidth and will do better for vector-cache-unfriendly code, e.g. with long vectors and little re-use. One of the SX-5 sites has demonstrated performance of over 110 Gflop/s on 15 processors of a 128 Gflop/s SX-5 on a real application.

Conclusion

Both the SX-5 and SV1ex provide impressive vector performance. The SX-5 has a huge memory and huge memory bandwidth. The SV1's caching can make up for its lack of bandwidth for many codes.

The SX-5 can reach greater peak performance in the one system. The SX-5 has a poorer scalar/vector balance than the SV1, and for some codes, such as cache_test, the SV1 MSP can out-perform two SX-5 processors, twice as well as might be expected based on peak speed..

The SX-5 software is far weaker, conforming only to older standards, and being behind in compiler technology, and in documentation.

There is no doubt that the SX-5 can provide spectacular performance and capability – 160 Gflop/s peak around a shared memory of 256 Gbyte is extraordinary. However, for many sites or applications where the capability is not needed, the SV1 can provide more cost-effective throughput, with a better software environment. If a performance level above the SV1 is needed, then the SX-5

can certainly provide that, without having to use distributed-memory programming.