
Overview of recent supercomputers

Aad J. van der Steen
high-performance Computing Group
Utrecht University
P.O. Box 80195
3508 TD Utrecht
The Netherlands
steen@phys.uu.nl
www.phys.uu.nl/~steen

Abstract

In this report we give an overview of high-performance computers which are currently available or will become available within a short time frame from vendors; no attempt is made to list all machines that are still in the research phase. The machines are described according to their architectural class. Shared and distributed-memory SIMD and MIMD machines are discerned. The information about each machine is kept as compact as possible. Moreover, no attempt is made to quote price information as this is often even more elusive than the performance of a system. In addition, some general information about high-performance computer architectures and the various processors and communication networks employed in these systems is given in order to better appreciate the system information given in this report.

This document reflects the technical state of the supercomputer arena as accurately as possible. However, the author nor NCF take any responsibility for errors or mistakes in this document. We encourage anyone who has comments or remarks on the contents to inform us, so we can improve this report.

NCF, the National Computing Facilities Foundation, supports and furthers the advancement of technical and scientific research with and into advanced computing facilities and prepares for the Netherlands national supercomputing policy. Advanced computing facilities are multi-processor vectorcomputers, massively parallel computing systems of various architectures and concepts and advanced networking facilities.

Contents

1	Introduction and account	3
2	Architecture of high-performance computers	6
2.1	The main architectural classes	6
2.2	Shared-memory SIMD machines	8
2.3	Distributed-memory SIMD machines	9
2.4	Shared-memory MIMD machines	11
2.5	Distributed-memory MIMD machines	12
2.6	ccNUMA machines	15
2.7	Clusters	16
2.8	Processors	17
2.8.1	AMD Opteron	17
2.8.2	HP/Compaq Alpha EV7	19
2.8.3	Hewlett Packard PA-RISC 8800	20
2.8.4	IBM POWER4+	21
2.8.5	IBM POWER5	22
2.8.6	IBM BlueGene processor	23
2.8.7	Intel Itanium 2	23
2.8.8	Intel Pentium 4/Xeon Nocona	24
2.8.9	MIPS R16000	26
2.8.10	Sun UltraSPARC IV	27
2.9	Networks	28
2.9.1	Infiniband	29
2.9.2	Myrinet	30
2.9.3	QsNet	31
2.9.4	SCI	31
3	Recount of (almost) available systems	33
3.1	System descriptions	33
3.1.1	The Bull NovaScale.	33
3.1.2	The C-DAC PARAM Padma.	34
3.1.3	The Cray MTA-2.	35
3.1.4	The Cray SX-6.	36
3.1.5	The Cray Inc. X1	36
3.1.6	The Cray Inc. XD1	37
3.1.7	The Cray Inc. XT3	38
3.1.8	The Fujitsu/Siemens PRIMEPOWER.	39
3.1.9	The Hitachi SR11000.	40
3.1.10	The HP 9000 Superdome.	41
3.1.11	The HP Integrity Superdome.	42
3.1.12	The HP/Compaq GS series.	43
3.1.13	The HP/Compaq AlphaServer SC45 series.	44
3.1.14	The IBM eServer p690.	45

3.1.15	The IBM BlueGene/L.	46
3.1.16	The NEC SX-8 series.	47
3.1.17	The NEC TX7 series.	48
3.1.18	The SGI Altix 3000 series.	49
3.1.19	The SGI Origin 3900.	50
3.1.20	The Sun Fire E25K.	51
3.1.21	The Unisys Server ES7000.	52
4	Systems disappeared from the list	53
4.1	Disappeared machines	53
5	Systems under development	59
5.1	Cray Inc.	59
5.2	Hewlett-Packard/Intel	59
5.3	IBM	60
5.4	SGI	60
5.5	Sun	60
6	Glossary of terms	61
6.1	Glossary	61
	Acknowledgments	66
	References	67

1 Introduction and account

This is the 15th edition of a report in which we attempt to give an overview of high-performance computer systems that are commercially available or are expected to become available within a short time frame (typically a few months to half a year). We choose the expression “attempt” deliberately because the market of high-performance machines is highly volatile: the rate with which systems are introduced — and disappear again — is very high and therefore the information will probably be only approximately valid. Nevertheless, we think that such an overview is useful for those who want to obtain a general idea about the various means by which these systems strive at high-performance, especially when it is updated on a regular basis.

We will try to be as up-to-date and compact as possible and on these grounds we think there is a place for this report. At this moment systems are disappearing from the market in a certain sense balance against the ones that are newly appearing. This is because the spectrum of integrated systems has lost a few members but also some (few) new systems have appeared. Besides that, an enormous growth in the use of clusters, some very large, can be observed. A larger amount of new systems may be expected in the next few years because of the renewed interest in computer architectures both on the processor level and the macro-architecture level. This new interest was sparked by the introduction of the Earth Simulator system in Japan which by TOP500 standards, see [44], is at present still the most powerful system in the world. As a result a new discussion emerged in the USA to look at new (and old) architectures as the gap between Theoretical Peak Performance and application performance for systems born from the ASCI (see [3]) initiative has grown steadily and for many users of such systems to an unacceptable level. Programs like the DARPA-funded HPCS program should curb this trend which cannot be done by staying on the same track of SMP-clustered RISC processor technology without further enhancements in memory access and intra- and inter-node bandwidth.

The majority of systems still look like minor variations on the same theme: clusters of RISC-based Symmetric Multi-Processing (SMP) nodes which in turn are connected by a fast network. Culler *et.al.* [8] consider this as a natural architectural evolution. However, it may also be argued that the requirements formulated in the ASCI program has steered these systems in this direction and that this will change in the coming years for the reasons given above.

The supercomputer market is a very dynamic one and this is especially true for the cluster world that have emerged at a tremendous rate in the last few years. The number of vendors that sell pre-configured clusters has boomed accordingly and, as for the last few issues, we have decided *not* to include such configurations in this report: the speed with which cluster companies and systems appear and disappear makes this almost impossible. We will briefly comment on cluster characteristics and their position relative to other supercomputers in section 2.7 though.

For the tightly-coupled or “integrated” parallel systems, however, we can by updating this report at least follow the main trends in popular and emerging architectures. The details of the systems be reported do not allow the report to be shorter than in former years: between 60–70 pages.

As of the 11th issue we decided to introduce a section that describes the dominant processors in some detail. This seems fit as the processors the heart of the systems. We do that in section 2.8. In addition, as in the 13th issue on we include a section that discusses specific network implementations, being also constituents of primary importance apart from the general discussion about communication networks.

The rule for including systems is as follows: they should be either available commercially at the time of appearance of this report, or within 6 months thereafter. This excludes some interesting cluster systems at the Sandia, Los Alamos, and Lawrence Livermore National Laboratories in the USA (all with measured performances in the range of 6–19 Tflop/s) and the Japanese Earth Simulator system (with a performance

around 40 Tflop/s) because they are not marketed or represent standard cluster technology (be it on a grand scale).

The rule that systems should be available within a time-span of 6 months is to avoid confusion by describing systems that are announced much too early, just for marketing reasons and that will not be available to general users within a reasonable time. We also have to refrain from including all generations of a system that are still in use. Therefore, for instance, we do not include the IBM SP1 or the Cray T90 series anymore although some of these systems are still in use. Generally speaking, we include machines that are presently marketed or will be marketed within 6 months. To add to the information given in this report, we quote the Web addresses of the vendors because the information found there may be more recent than what can be provided here. On the other hand, such pages should be read with care because it will not always be clear what the status is of the products described there.

Some vendors offer systems that are identical in all respects except in the clock cycle of the nodes (examples are the SGI Origin3000 series and the Fujitsu PRIMEPOWER). In these cases we always only mention the models with the fastest clock as it will be always possible to get the slower systems and we presume that the reader is primarily interested in the highest possible speeds that can be reached with these systems.

Until the eighth issue of this report we ordered the systems by their architectural classes as explained in section 2.1. However, this distinction became more and more artificial as is explained in the same section. Therefore all systems described are simply listed alphabetically. In the header of each system description the machine type is provided. There is referred to the architectural class for as far this is relevant. We omit price information which in most cases is next to useless. If available, we will give some information about performances of systems based on user experiences instead of only giving theoretical peak performances. Here we have adhered to the following policy: We try to quote *best measured performances*, if available, thus providing a more realistic upper bound than the theoretical peak performance. We hardly have to say that the speed range of supercomputers is enormous, so the best measured performance will not always reflect the performance of the reader's favorite application. In fact, when the HPC Linpack test is used to measure the speed it is almost certain that for the average user the application performance will be significantly lower. When we give performance information, it is not always possible to quote all sources and in any case if this information seems (or is) biased, this is entirely the responsibility of the author of this report. He is quite willing to be corrected or to receive additional information from anyone who is in the position to do so.

Although for the average user the appearance of new systems in the last years tended to become rapidly more and more alike, it is still useful to dwell a little on the architectural classes that underlie this appearance. It gives some insight in the various ways that high-performance is achieved and a feeling why machines perform as they do. This is done in section 2 which will be referred to repeatedly in section 3 in the description of the machines.

Up till the 10th issue we included a section (4) on systems that disappeared from the market. We reduced that section in the printed and PostScript versions because it tends to take an unreasonable part of the total text. Still, because this information is of interest to a fair amount of readers and it gives insight in the field of the historical development of supercomputing over the last 14 years, this information will still be available in full at URL <http://www.phys.uu.nl/~steen/web05a/gone.html>. In section 5 we present some systems that are under development and have a fair chance to appear on the market. Because of the addition of the section on processors that introduces many technical terms, also a glossary is included.

The overview given in this report concentrates on the computational capabilities of the systems discussed. To do full justice to all assets of present days high-performance computers one should list their I/O performance and their connectivity possibilities as well. However, the possible permutations of configurations even for one model of a certain system often are so large that they would multiply the volume of this report, which we tried to limit for greater clarity. So, not all features of the systems discussed will be present. Still we think (and certainly hope) that the impressions obtained from the entries of the individual machines may be useful to many. We also omitted some systems that may be characterised as "high-performance" in the fields of database management, real-time computing, or visualisation. Therefore, as we try to give an overview for the area of general scientific and technical computing, systems that are primarily meant for database retrieval like the AT&T GIS systems or concentrate exclusively on the real-time user community, like Concurrent Computing Systems, are not discussed in this report. Furthermore, we have set a threshold of about 80 Gflop/s for systems to appear in this report as, at least with regard to theoretical peak performance, single

CPUs often exceed 2 Gflop/s although their actual performance may be an entirely other matter.

Although most terms will be familiar to many readers, we still think it is worthwhile to give some of the definitions in section 2 because some authors tend to give them a meaning that may slightly differ from the idea the reader already has acquired.

Lastly, we should point out that also a web version is available. The URLs are:

www.arcade-eu.org/overview (Europe).

www.phys.uu.nl/~steen/web05a/overview.html (Europe).

www.euroben.nl/reports/web05a/overview.html (Europe).

www.netlib.org/utk/papers/advanced-computers/ (USA).

From this issue on we will *attempt* to keep the web version up-to-date by refreshing the contents more frequently than once a year. So, the printed version may lag a little behind the web version over the year.

2 Architecture of high-performance computers

Before going on to the descriptions of the machines themselves, it is important to consider some mechanisms that are or have been used to increase the performance. The hardware structure or *architecture* determines to a large extent what the possibilities and impossibilities are in speeding up a computer system beyond the performance of a single CPU. Another important factor that is considered in combination with the hardware is the capability of compilers to generate efficient code to be executed on the given hardware platform. In many cases it is hard to distinguish between hardware and software influences and one has to be careful in the interpretation of results when ascribing certain effects to hardware or software peculiarities or both. In this chapter we will give most emphasis on the hardware architecture. For a description of machines that can be considered to be classified as “high-performance” one is referred to [8, 39].

2.1 The main architectural classes

Since many years the taxonomy of Flynn [14] has proven to be useful for the classification of high-performance computers. This classification is based on the way of manipulating of instruction and data streams and comprises four main architectural classes. We will first briefly sketch these classes and afterwards fill in some details when each of the classes is described separately.

- **SISD** machines: These are the conventional systems that contain one CPU and hence can accommodate one instruction stream that is executed serially. Nowadays many large mainframes may have more than one CPU but each of these execute instruction streams that are unrelated. Therefore, such systems still should be regarded as (a couple of) SISD machines acting on different data spaces. Examples of SISD machines are for instance most workstations like those of Hewlett-Packard, IBM, and SGI. The definition of SISD machines is given here for completeness’ sake. We will not discuss this type of machines in this report.
- **SIMD** machines: Such systems often have a large number of processing units, ranging from 1,024 to 16,384 that all may execute the same instruction on different data in lock-step. So, a single instruction manipulates many data items in parallel. Examples of SIMD machines in this class are the CPP Gamma II and the Quadrics Apemille which are not marketed anymore since about a year.

Another subclass of the SIMD systems are the vectorprocessors. Vectorprocessors act on arrays of similar data rather than on single data items using specially structured CPUs. When data can be manipulated by these vector units, results can be delivered with a rate of one, two and — in special cases — of three per clock cycle (a clock cycle being defined as the basic internal unit of time for the system). So, vector processors execute on their data in an almost parallel way but only when executing in vector mode. In this case they are several times faster than when executing in conventional scalar mode. For practical purposes vectorprocessors are therefore mostly regarded as SIMD machines. An example of such systems is for instance the NEC SX-8B.

- **MISD** machines: Theoretically in these type of machines multiple instructions should act on a single stream of data. As yet no practical machine in this class has been constructed nor are such systems easy to conceive. We will disregard them in the following discussions.
- **MIMD** machines: These machines execute several instruction streams in parallel on different data. The difference with the multi-processor SISD machines mentioned above lies in the fact that the instructions and data are related because they represent different parts of the same task to be executed. So, MIMD systems may run many sub-tasks in parallel in order to shorten the time-to-solution for the main task to be executed. There is a large variety of MIMD systems and especially in this class the Flynn taxonomy proves to be not fully adequate for the classification of systems. Systems that behave

very differently like a four-processor NEC SX-8 vector system and a thousand-processor IBM p690 fall both in this class. In the following we will make another important distinction between classes of systems and treat them accordingly.

- **Shared-memory systems:** Shared-memory systems have multiple CPUs all of which share the same address space. This means that the knowledge of where data is stored is of no concern to the user as there is only one memory accessed by all CPUs on an equal basis. Shared memory systems can be both SIMD or MIMD. Single-CPU vector processors can be regarded as an example of the former, while the multi-CPU models of these machines are examples of the latter. We will sometimes use the abbreviations SM-SIMD and SM-MIMD for the two subclasses.
- **Distributed-memory systems:** In this case each CPU has its own associated memory. The CPUs are connected by some network and may exchange data between their respective memories when required. In contrast to shared-memory machines the user must be aware of the location of the data in the local memories and will have to move or distribute these data explicitly when needed. Again, distributed-memory systems may be either SIMD or MIMD. The first class of SIMD systems mentioned which operate in lock step, all have distributed memories associated to the processors. As we will see, distributed-memory MIMD systems exhibit a large variety in the topology of their interconnection network. The details of this topology are largely hidden from the user which is quite helpful with respect to portability of applications. For the distributed-memory systems we will sometimes use DM-SIMD and DM-MIMD to indicate the two subclasses.

As already alluded to, although the difference between shared and distributed-memory machines seems clear cut, this is not always entirely the case from user's point of view. For instance, the late Kendall Square Research systems employed the idea of "virtual shared-memory" on a hardware level. Virtual shared-memory can also be simulated at the programming level: A specification of High Performance Fortran (HPF) was published in 1993 [21] which by means of compiler directives distributes the data over the available processors. Therefore, the system on which HPF is implemented in this case will look like a shared-memory machine to the user. Other vendors of Massively Parallel Processing systems (sometimes called MPP systems), like HP and SGI, also support proprietary virtual shared-memory programming models due to the fact that these physically distributed memory systems are able to address the whole collective address space. So, for the user such systems have one *global address space* spanning all of the memory in the system. We will say a little more about the structure of such systems in section 2.6. In addition, packages like TreadMarks ([1]) provide a "distributed shared-memory" environment for networks of workstations. A good overview of such systems is given at [10].

Distributed processing takes the DM-MIMD concept one step further: instead of many integrated processors in one or several boxes, workstations, mainframes, etc., are connected by (Gigabit) Ethernet, or other, faster networks and set to work concurrently on tasks in the same program. Conceptually, this is not different from DM-MIMD computing, but the communication between processors can be much slower. Packages that initially were made to realise distributed computing like PVM (standing for Parallel Virtual Machine) [15], and MPI (Message Passing Interface, [24, 25]) have become *de facto* standards for the "message passing" programming model. MPI and PVM have become so widely accepted that they have been adopted by all vendors of distributed-memory MIMD systems and even on shared-memory MIMD systems for compatibility reasons. In addition, there is a tendency to cluster shared-memory systems by a fast communication network to obtain systems with a very high computational power. E.g., the NEC SX-8, and the Cray X1 have this structure. So, within the clustered nodes a shared-memory programming style can be used while between clusters message-passing should be used. It must be said that PVM is not used very much anymore and that MPI has more or less become the *de facto* standard.

For SM-MIMD systems we mention OpenMP [30, 7], that can be used to parallelise Fortran and C(++) programs by inserting comment directives (Fortran 77/90/95) or pragmas (C/C++) into the code. OpenMP has quickly been adopted by the all major vendors and has become a well established standard for shared memory systems.

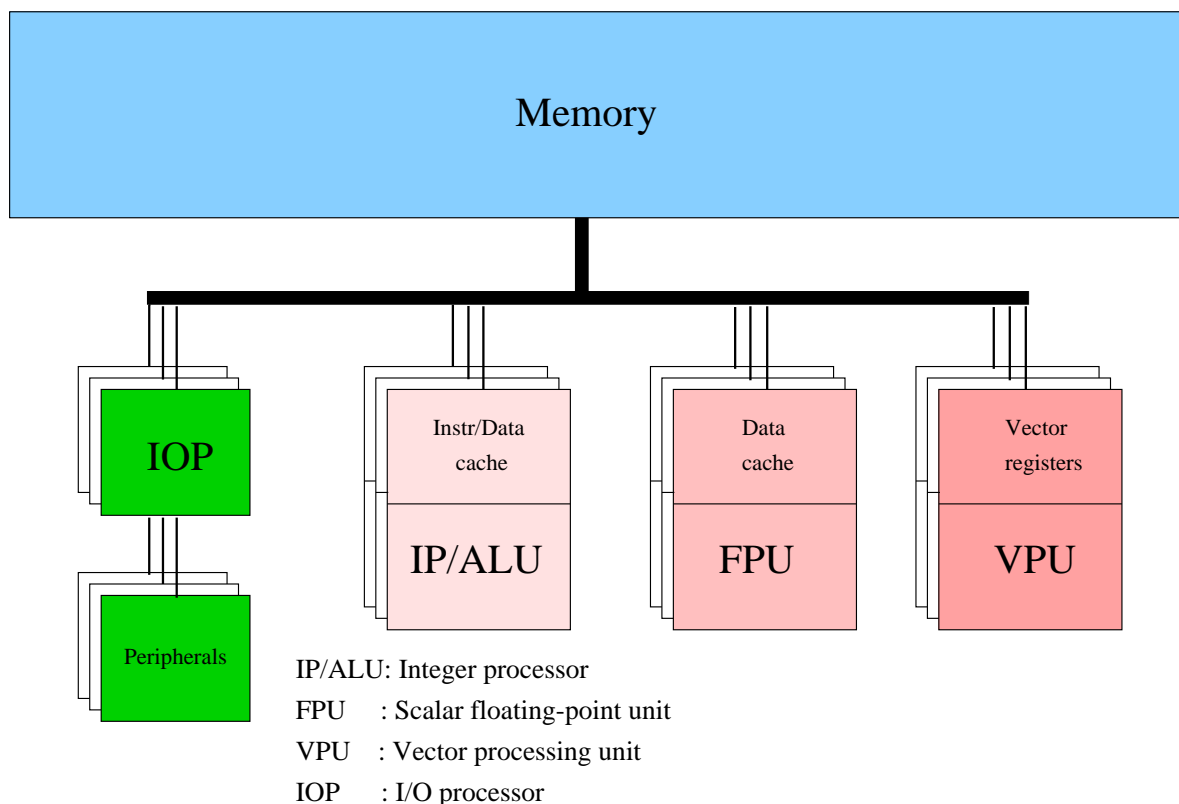


Figure 2.1: Blockdiagram of a vector processor.

2.2 Shared-memory SIMD machines

This subclass of machines is practically equivalent to the single-processor vectorprocessors, although other interesting machines in this subclass have existed (viz. VLIW machines [36]). In the block diagram in Figure 2.1 we depict a generic model of a vector architecture. The single-processor vector machine will have only one of the vectorprocessors depicted and the system may even have its scalar floating-point capability shared with the vector processor (as was the case in some Cray systems). It may be noted that the VPU does not show a cache. The majority of vectorprocessors do not employ a cache anymore. In many cases the vector unit cannot take advantage of it and execution speed may even be unfavourably affected because of frequent cache overflow.

Although vectorprocessors have existed that loaded their operands directly from memory and stored the results again immediately in memory (CDC Cyber 205, ETA-10), present-day vectorprocessors use vector registers. This usually does not impair the speed of operations while providing much more flexibility in gathering operands and manipulation with intermediate results.

Because of the generic nature of Figure 2.1 no details of the interconnection between the VPU and the memory are shown. Still, these details are very important for the effective speed of a vector operation: when the bandwidth between memory and the VPU is too small it is not possible to take full advantage of the VPU because it has to wait for operands and/or has to wait before it can store results. When the ratio of arithmetic to load/store operations is not high enough to compensate for such situations, severe performance losses may be incurred. The influence of the number of load/store paths for the dyadic vector operation $c = a + b$ (a , b , and c vectors) is depicted in Figure 2.2. Because of the high costs of implementing these data paths between memory and the VPU, often compromises are sought and the full required bandwidth (i.e., two load operations and one store operation at the *same* time) Only Cray Inc. in its former Y-MP, C-series, T-series employed this very high bandwidth. Vendors rather rely on additional caches and other

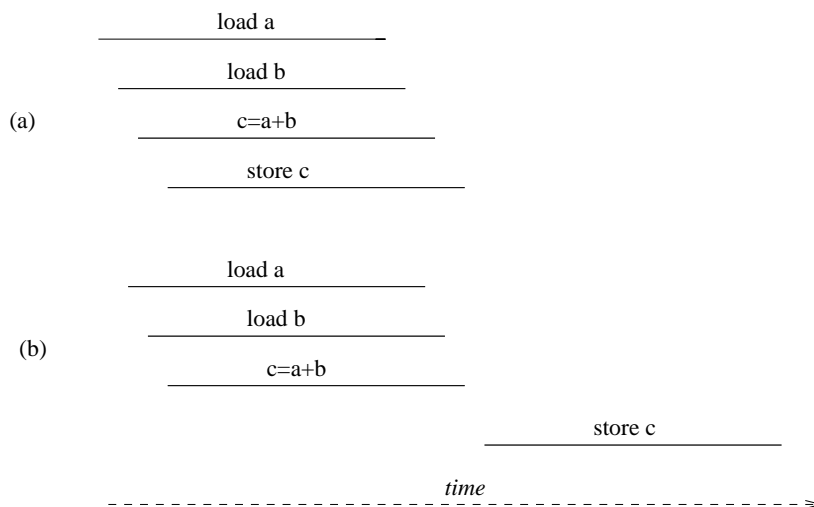


Figure 2.2: Schematic diagram of a vector addition. Case (a) when two load- and one store pipe are available; case (b) when two load/store pipes are available.

tricks to hide the lack of bandwidth.

The VPUs are shown as a single block in Figure 2.1. Yet, there is a considerable diversity in the structure of VPUs. Every VPU consists of a number of vector functional units, or “pipes” that fulfill one or several functions in the VPU. Every VPU will have pipes that are designated to perform memory access functions, thus assuring the timely delivery of operands to the arithmetic pipes and of storing the results in memory again. Usually there will be several arithmetic functional units for integer/logical arithmetic, for floating-point addition, for multiplication and sometimes a combination of both, a so-called compound operation. Division is performed by an iterative procedure, table look-up, or a combination of both using the add and multiply pipe. In addition, there will almost always be a mask pipe to enable operation on a selected subset of elements in a vector of operands. Lastly, such sets of vector pipes can be replicated within one VPU (2 up to 16-fold replication occurs). Ideally, this will increase the performance per VPU by the same factor provided the bandwidth to memory is adequate.

2.3 Distributed-memory SIMD machines

Machines of this type are sometimes also known as *processor-array* machines [18]. Because the processors of these machines operate in lock-step, i.e., all processors execute the same instruction at the same time (but on different data items), no synchronisation between processors is required. This greatly simplifies the design of such systems. A *control processor* issues the instructions that are to be executed by the processors in the processor array. All currently available DM-SIMD machines use a front-end processor to which they are connected by a data path to the control processor. Operations that cannot be executed by the processor array or by the control processor are offloaded to the front-end system. For instance, I/O may be through the front-end system, by the processor array machine itself or by both. Figure 2.3 shows a generic model of a DM-SIMD machine of which actual models will deviate to some degree. Figure 2.3 might suggest that all processors in such systems are connected in a 2-D grid and indeed, the interconnection topology of this type of machines always includes the 2-D grid. As opposing ends of each grid line are also always connected the topology is rather that of a torus. This is not the only interconnection scheme: They might also be connected in 3-D, diagonally, or more complex structures.

It is possible to exclude processors in the array from executing an instruction on certain logical conditions, but this means that for the time of this instruction these processors are idle (a direct consequence of the SIMD-type operation) which immediately lowers the performance. Another factor that may adversely affect the speed occurs when data required by processor i resides in the memory of processor j — in fact, as

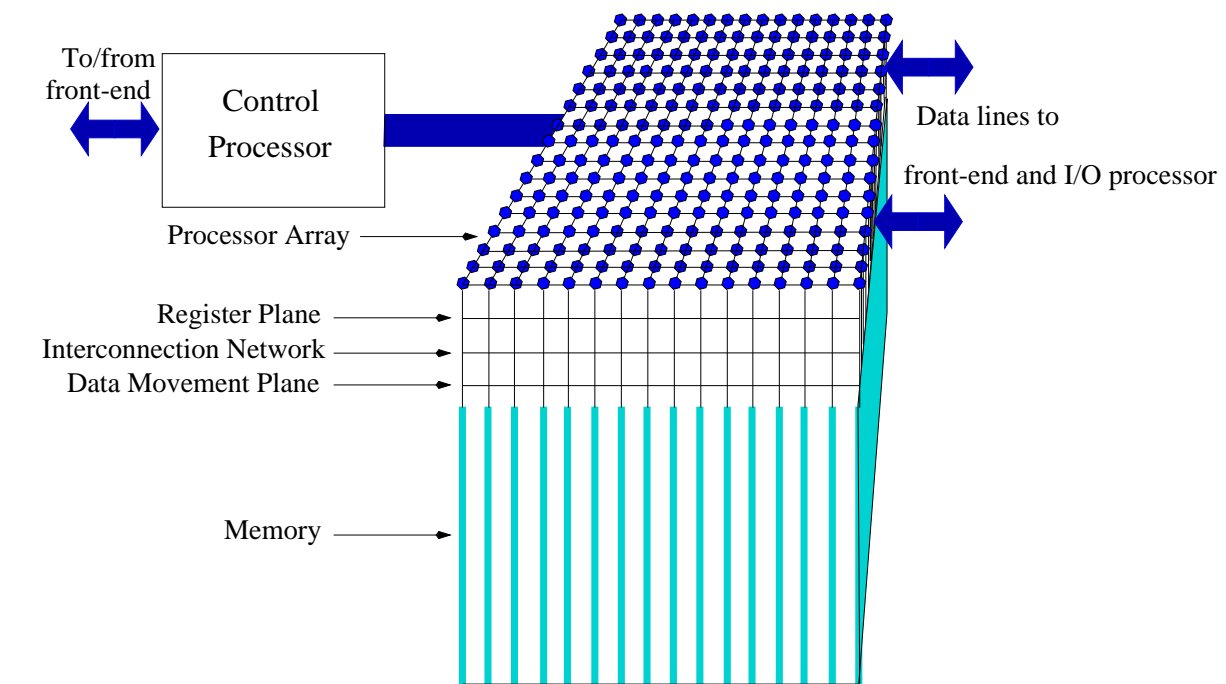


Figure 2.3: A generic block diagram of a distributed-memory SIMD machine.

this occurs for all processors at the same time, this effectively means that data will have to be permuted across the processors. To access the data in processor j , the data will have to be fetched by this processor and then send through the routing network to processor i . This may be fairly time consuming. For both reasons mentioned DM-SIMD machines are rather specialised in their use when one wants to employ their full parallelism. Generally, they perform excellently on digital signal and image processing and on certain types of Monte Carlo simulations where virtually no data exchange between processors is required and exactly the same type of operations is done on massive datasets with a size that can be made to fit comfortably in these machines.

The control processor as depicted in Figure 2.3 may be more or less intelligent. It issues the instruction sequence that will be executed by the processor array. In the worst case (that means a less autonomous control processor) when an instruction is not fit for execution on the processor array (e.g., a simple print instruction) it might be offloaded to the front-end processor which may be much slower than execution on the control processor. In case of a more autonomous control processor this can be avoided thus saving processing interrupts both on the front-end and the control processor. Most DM-SIMD systems have the possibility to handle I/O independently from the front-end processors. This is not only favourable because the communication between the front-end and back-end systems is avoided. The (specialised) I/O devices for the processor-array system is generally much more efficient in providing the necessary data directly to the memory of the processor array. Especially for very data-intensive applications like radar and image processing such I/O systems are very important.

A feature that is peculiar to this type of machines is that the processors sometimes are of a very simple bit-serial type, i.e., the processors operate on the data items bitwise, irrespective of their type. So, e.g., operations on integers are produced by software routines on these simple bit-serial processors which takes at least as many cycles as the operands are long. So, a 32-bit integer result will be produced two times faster than a 64-bit result. For floating-point operations a similar situation holds, be it that the number of cycles required is a multiple of that needed for an integer operation. As the number of processors in this type of systems is mostly large (1024 or larger, the Quadrics Apemille was a notable exception, however), the slower operation on floating-point numbers can be often compensated for by their number, while the cost per processor is quite low as compared to full floating-point processors. In some cases, however, floating-point

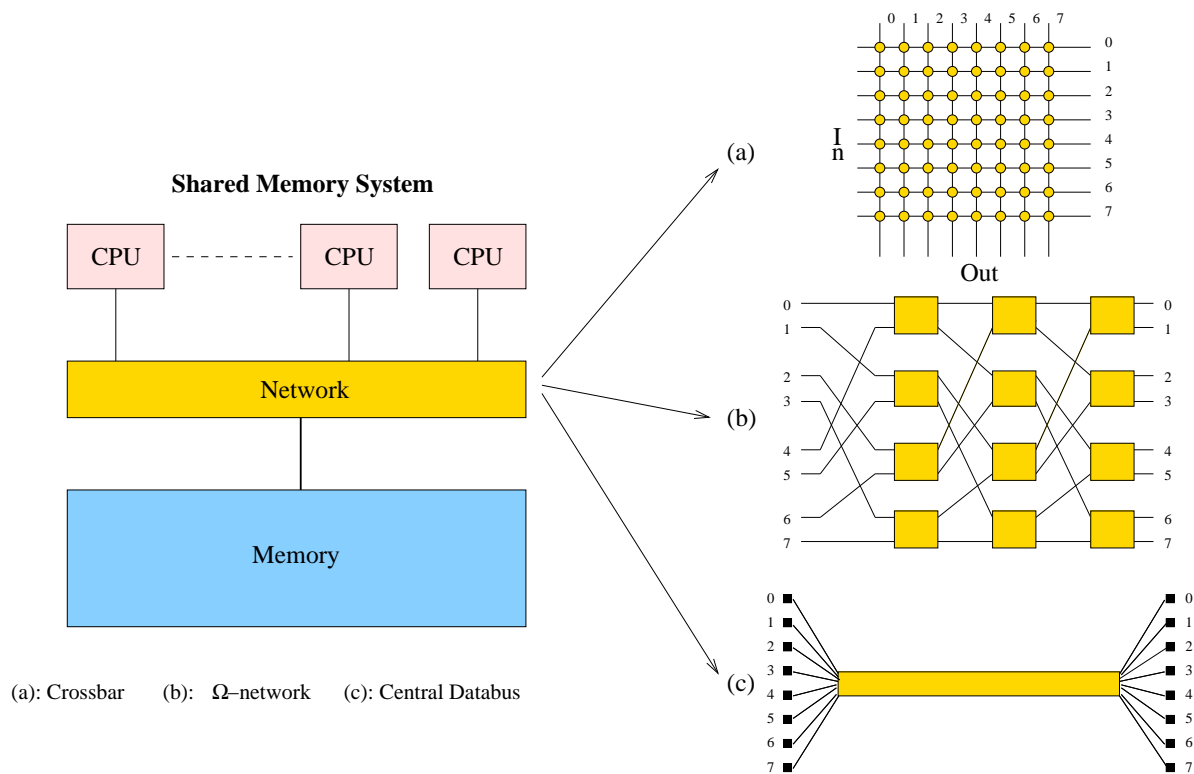


Figure 2.4: Some examples of interconnection structures used in shared-memory MIMD systems.

coprocessors are added to the processor-array. Their number is 8–16 times lower than that of the bit-serial processors because of the cost argument. An advantage of bit-serial processors is that they may operate on operands of any length. This is particularly advantageous for random number generation (which often boils down to logical manipulation of bits) and for signal processing because in both cases operands of only 1–8 bits are abundant. As the execution time for bit-serial machines is proportional to the length of the operands, this may result in significant speedups.

2.4 Shared-memory MIMD machines

In Figure 2.1 already one subclass of this type of machines was shown. In fact, the single-processor vector machine discussed there was a special case of a more general type. The figure shows that more than one FPU and/or VPU may be possible in one system.

The main problem one is confronted with in shared-memory systems is that of the connection of the CPUs to each other and to the memory. As more CPUs are added, the collective bandwidth to the memory ideally should increase linearly with the number of processors, while each processor should preferably communicate directly with all others without the much slower alternative of having to use the memory in an intermediate stage. Unfortunately, full interconnection is quite costly, growing with $\mathcal{O}(n^2)$ while increasing the number of processors with $\mathcal{O}(n)$. So, various alternatives have been tried. Figure 2.4 shows some of the interconnection structures that are (and have been) used.

As can be seen from the Figure, a crossbar uses n^2 connections, an Ω -network uses $n \log_2 n$ connections while with the central bus there is only one connection. This is reflected in the use of each connection path for the different types of interconnections: for a crossbar each data path is direct and does not have to be shared with other elements. In case of the Ω -network there are $\log_2 n$ switching stages and as many data items may have to compete for any path. For the central data bus all data have to share the same bus, so n data items may compete at any time.

The bus connection is the least expensive solution, but it has the obvious drawback that bus contention may occur thus slowing down the computations. Various intricate strategies have been devised using caches associated with the CPUs to minimise the bus traffic. This leads however to a more complicated bus structure which raises the costs. In practice it has proved to be very hard to design buses that are fast enough, especially where the speed of the processors has been increasing very quickly and it imposes an upper bound on the number of processors thus connected that in practice appears not to exceed a number of 10-20. In 1992, a new standard (IEEE P896) for a fast bus to connect either internal system components or to external systems has been defined. This bus, called the Scalable Coherent Interface (SCI) should provide a point-to-point bandwidth of 200-1,000 MB/s. It has been used in the HP Exemplar systems, but also within a cluster of workstations as offered by SCALI. The SCI is much more than a simple bus and it can act as the hardware network framework for distributed computing, see [22].

A multi-stage crossbar is a network with a logarithmic complexity and it has a structure which is situated somewhere in between a bus and a crossbar with respect to potential capacity and costs. The Ω -network as depicted in figure 2.4 is an example. Commercially available machines like the IBM eServer p690, the SGI Origin3000, and the late Cenju-4 use(d) such a network structure, but a number of experimental machines also have used this or a similar kind of interconnection. The BBN TC2000 that acted as a virtual shared-memory MIMD system used an analogous type of network (a Butterfly-network) and it is quite conceivable that new machines may use it, especially as the number of processors grows. For a large number of processors the $n \log_2 n$ connections quickly become more attractive than the n^2 used in crossbars. Of course, the switches at the intermediate levels should be sufficiently fast to cope with the bandwidth required. Obviously, not only the *structure* but also the *width* of the links between the processors is important: a network using 16-bit parallel links will have a bandwidth which is 16 times higher than a network with the same topology implemented with serial links.

In all present-day multi-processor vectorprocessors crossbars are used. This is still feasible because the maximum number of processors in a system is still rather small (32 at most presently). When the number of processors would increase, however, technological problems might arise. Not only it becomes harder to build a crossbar of sufficient speed for the larger numbers of processors, the processors themselves generally also increase in speed individually, doubling the problems of making the speed of the crossbar match that of the bandwidth required by the processors.

Whichever network is used, the type of processors in principle could be arbitrary for any topology. In practice, however, bus structured machines do not have vector processors as the speeds of these would grossly mismatch with any bus that could be constructed with reasonable costs. All available bus-oriented systems use RISC processors. The local caches of the processors can sometimes alleviate the bandwidth problem if the data access can be satisfied by the caches thus avoiding references to the memory.

The systems discussed in this subsection are of the MIMD type and therefore different tasks may run on different processors simultaneously. In many cases synchronisation between tasks is required and again the interconnection structure is very important here. Most vectorprocessors employ special communication registers within the CPUs by which they can communicate directly with the other CPUs they have to synchronise with. The systems may also synchronise via the shared memory. Generally, this is much slower but it can still be acceptable when the synchronisation occurs relatively seldom. Of course, in bus-based systems communication also has to be done via a bus. This bus is mostly separated from the data bus to ensure a maximum speed for the synchronisation.

2.5 Distributed-memory MIMD machines

The class of DM-MIMD machines is undoubtedly the fastest growing part in the family of high-performance computers. Although this type of machines is more difficult to deal with than shared-memory machines and DM-SIMD machines. The latter type of machines are processor-array systems in which the data structures that are candidates for parallelisation are vectors and multi-dimensional arrays that are laid out automatically on the processor array by the system software. For shared-memory systems the data distribution is completely transparent to the user. This is quite different for DM-MIMD systems where the user has to

distribute the data over the processors and also the data exchange between processors has to be performed explicitly. The initial reluctance to use DM-MIMD machines seems to have been decreased. Partly this is due to the now existing standard for communication software ([15, 24, 25]) and partly because, at least theoretically, this class of systems is able to outperform all other types of machines.

The advantages of DM-MIMD systems are clear: the bandwidth problem that haunts shared-memory systems is avoided because the bandwidth scales up automatically with the number of processors. Furthermore, the speed of the memory which is another critical issue with shared-memory systems (to get a peak performance that is comparable to that of DM-MIMD systems, the processors of the shared-memory machines should be very fast and the speed of the memory should match it) is less important for the DM-MIMD machines, because more processors can be configured without the afore mentioned bandwidth problems.

Of course, DM-MIMD systems also have their disadvantages: The communication between processors is much slower than in SM-MIMD systems, and so, the synchronisation overhead in case of communicating tasks is generally orders of magnitude higher than in shared-memory machines. Moreover, the access to data that are not in the local memory belonging to a particular processor have to be obtained from non-local memory (or memories). This is again on most systems very slow as compared to local data access. When the structure of a problem dictates a frequent exchange of data between processors and/or requires many processor synchronisations, it may well be that only a very small fraction of the theoretical peak speed can be obtained. As already mentioned, the data and task decomposition are factors that mostly have to be dealt with explicitly, which may be far from trivial.

It will be clear from the paragraph above that also for DM-MIMD machines both the topology and the speed of the data paths are of crucial importance for the practical usefulness of a system. Again, as in the section on SM-MIMD systems, the richness of the connection structure has to be balanced against the costs. Of the many conceivable interconnection structures only a few are popular in practice. One of these is the so-called hypercube topology as depicted in Figure 2.5 (a).

A nice feature of the hypercube topology is that for a hypercube with 2^d nodes the number of steps to be taken between any two nodes is at most d . So, the dimension of the network grows only logarithmically with the number of nodes. In addition, theoretically, it is possible to simulate any other topology on a hypercube: trees, rings, 2-D and 3-D meshes, etc. In practice, the exact topology for hypercubes does not matter too much anymore because all systems in the market today employ what is called “wormhole routing”. This means that a message is sent from node i to node j a header message is sent from i to j , resulting in a direct connection between these nodes. As soon as this connection is established, the data proper is sent through this connection without disturbing the operation of the intermediate nodes. Except for a small amount of time in setting up the connection between nodes, the communication time has become virtually independent of the distance between the nodes. Of course, when several messages in a busy network have to compete for the same paths, waiting times are incurred as in any network that does not directly connect any processor to all others and often rerouting strategies are employed to circumvent busy links.

Another cost-effective way to connect a large number of processors is by means of a *fat tree*. In principle a simple tree structure for a network is sufficient to connect all nodes in a computer system. However, in practice it turns out that near the root of the tree congestion occurs because of the concentration of messages that first have to traverse the higher levels in the tree structure before they can descend again to their target nodes. The fat tree amends this shortcoming by providing more bandwidth (mostly in the form of multiple connections) in the higher levels of the tree. One speaks of a N -ary fat tree when the levels towards the roots are N times the number of connections in the level below it. An example of a quaternary fat tree with a bandwidth in the highest level that is four times that of the lower levels is shown in Figure 2.5 (b).

A number of massively parallel DM-MIMD systems seem to favour a 2- or 3-D mesh (torus) structure. The rationale for this seems to be that most large-scale physical simulations can be mapped efficiently on this topology and that a richer interconnection structure hardly pays off. However, some systems maintain (an) additional network(s) besides the mesh to handle certain bottlenecks in data distribution and retrieval [19]. Also on IBM’s BlueGene systems this philosophy has been followed.

A large fraction of systems in the DM-MIMD class employ crossbars. For relatively small amounts of processors (in the order of 64) this may be a direct or 1-stage crossbar, while to connect larger numbers of nodes




Figure 2.5: Some often used networks for DM machine types.

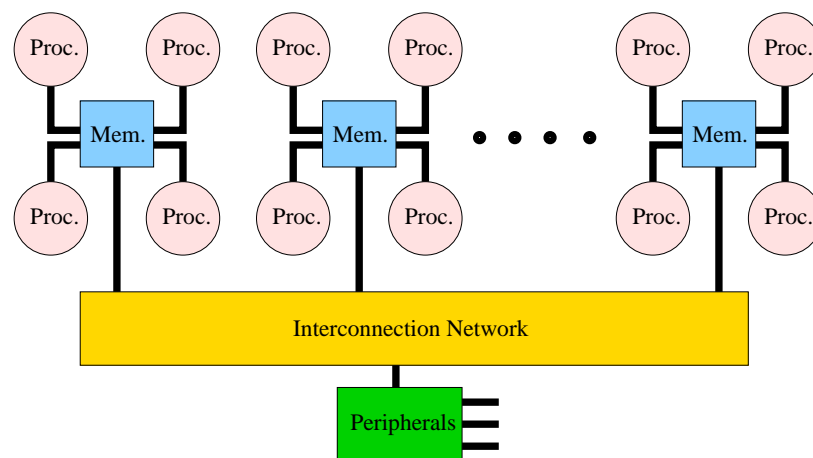


Figure 2.6: Block diagram of a system with a “hybrid” network: clusters of four CPUs are connected by a crossbar. The clusters are connected by a less expensive network, e.g., a Butterfly network.

multi-stage crossbars are used, i.e., the connections of a crossbar at level 1 connect to a crossbar at level 2, etc., instead of directly to nodes at more remote distances in the topology. In this way it is possible to connect in the order of a few thousands of nodes through only a few switching stages. In addition to the hypercube structure, other logarithmic complexity networks like Butterfly, Ω , or shuffle-exchange networks are often employed in such systems.

As with SM-MIMD machines, a node may in principle consist of any type of processor (scalar or vector) for computation or transaction processing together with local memory (with or without cache) and, in almost all cases, a separate communication processor with links to connect the node to its neighbours. Nowadays, the node processors are mostly off-the-shelf RISC processors sometimes enhanced by vector processors. A problem that is peculiar to this DM-MIMD systems is the mismatch of communication vs. computation speed that may occur when the node processors are upgraded without also speeding up the intercommunication. In some cases this may result in turning computational-bound problems into communication-bound problems.

2.6 ccNUMA machines

As already mentioned in the introduction, a trend can be observed to build systems that have a rather small (up to 16) number of RISC processors that are tightly integrated in a cluster, a Symmetric Multi-Processing (SMP) node. The processors in such a node are virtually always connected by a 1-stage crossbar while these clusters are connected by a less costly network. Such a system may look as depicted in Figure 2.6. Note that in Figure 2.6 all CPUs in a cluster are connected to a common part of the memory. This is similar to the policy mentioned for large vectorprocessor ensembles mentioned above but with the important difference that all of the processors can access all of the address space if necessary. The most important ways to let the SMP nodes share their memory are S-COMA (Simple Cache-Only Memory Architecture) and ccNUMA, which stands for Cache Coherent Non-Uniform Memory Access. Therefore, such systems can be considered as SM-MIMD machines. On the other hand, because the memory is physically distributed, it cannot be guaranteed that a data access operation always will be satisfied within the same time. In S-COMA systems the cache hierarchy of the local nodes is extended to the memory of the other nodes. So, when data is required that does not reside in the local node’s memory it is retrieved from the memory of the node where it is stored. In ccNUMA this concept is further extended in that all memory in the system is regarded (and addressed) globally. So, a data item may not be physically local but logically it belongs to one shared address space. Because the data can be physically dispersed over many nodes, the access time for different data items may well be different which explains the term non-uniform data access. The term “Cache Coherent” refers to the fact that for all CPUs any variable that is to be used must have a consistent value. Therefore,

is must be assured that the caches that provide these variables are also consistent in this respect. There are various ways to ensure that the caches of the CPUs are coherent. One is the *snoopy bus protocol* in which the caches listen in on transport of variables to any of the CPUs and update their own copies of these variables if they have them. Another way is the *directory memory*, a special part of memory which enables to keep track of all the copies of variables and of their validness.

Presently, no commercially available machine uses the S-COMA scheme. By contrast, there are several popular ccNUMA systems (Bull NovaScale, HP Superdome, and SGI Altix3000) commercially available. An important characteristic for NUMA machines is the *NUMA factor*. This factor shows the difference in latency for accessing data from a local memory location as opposed to a non-local one. Depending on the connection structure of the system the NUMA factor for various parts of a system can differ for part to part: accessing of data from a neighbouring node will be faster than from a distant node in which possibly a number of stages of a crossbar must be traversed. So, when a NUMA factor is mentioned, this is mostly for the largest network crosssection, i.e., the maximal distance between processors.

For all practical purposes we can classify these systems as being SM-MIMD machines also because special assisting hardware/software (such as a directory memory) has been incorporated to establish a single system image although the memory is physically distributed.

2.7 Clusters

The adoption of clusters, collections of workstations/PCs connected by a local network, has virtually exploded since the introduction of the first Beowulf cluster in 1994. The attraction lies in the (potentially) low cost of both hardware and software and the control that builders/users have over their system. The interest for clusters can be seen for instance from the IEEE Task Force on Cluster Computing (TFCC) which reviews on a regular basis the current status of cluster computing [45]. Also books how to build and maintain clusters have greatly added to their popularity [43, 35]. As the cluster scene becomes relatively mature and an attractive market, large HPC vendors as well as many start-up companies have entered the field and offer more or less ready out-of-the-box cluster solutions for those groups that do not want to build their cluster from scratch.

The number of vendors that sell cluster configurations has become so large that it is not possible to include all these products in this report. In addition, there is generally a large difference in the usage of clusters and their more integrated counterparts that we discuss in the following sections: clusters are mostly used for *capability computing* while the integrated machines primarily are used for *capacity computing*. The first mode of usage meaning that the system is employed for one or a few programs for which no alternative is readily available in terms of computational capabilities. The second way of operating a system is in employing it to the full by using the most of its available cycles by many, often very demanding, applications and users. Traditionally, vendors of large supercomputer systems have learned to provide for this last mode of operation as the precious resources of their systems were required to be used as effectively as possible. By contrast, Beowulf clusters are mostly operated through the Linux operating system (a small minority using Microsoft Windows) where these operating systems either miss the tools or these tools are relatively immature to use a cluster well for capacity computing. However, as clusters become on average both larger and more stable, there is a trend to use them also as computational capacity servers. In [41] is looked at some of the aspects that are necessary conditions for this kind of use like available cluster management tools and batch systems. In the same study also the performance on an application workload was assessed, both on a RISC (Compaq Alpha) based configuration and on Intel Pentium III based systems. An important, but not very surprising conclusion was that the speed of the network is very important in all but the most compute bound applications. Another notable observation was that using compute nodes with more than 1 CPU may be attractive from the point of view of compactness and (possibly) energy and cooling aspects, but that the performance can be severely damaged by the fact that more CPUs have to draw on a common node memory. The bandwidth of the nodes is in this case not up to the demands of memory intensive applications.

Fortunately, there is nowadays a fair choice of communication networks available in clusters. Of course 100 Mb/s Ethernet or Gigabit Ethernet is always possible, which is attractive for economic reasons, but has the drawback of a high latency ($\approx 100\mu\text{s}$). Alternatively, there are for instance networks that operate from user

space, like Myrinet [26], Infiniband, and SCI [22]. The first two have maximum bandwidths in the order of 250 MB/s and 850 MB/s, respectively, and a latency in the range of 7–9 μ s. SCI has a bandwidth of 400–500 MB/s theoretically and a latency under 3 μ s. The latter solution is more costly but is nevertheless employed in some cluster configurations. The network speeds as shown by Myrinet, and, certainly, QsNet and SCI is more or less on par with some integrated parallel systems as discussed later. So, possibly apart from the speed of the processors and of the software that is provided by the vendors of DM-MIMD supercomputers, the distinction between clusters and this class of machines becomes rather small and will undoubtedly decrease in the coming years.

2.8 Processors

In comparison to 10 years ago the processor scene has become drastically different. While in the period 1980–1990, the proprietary processors and in particular the vectorprocessors were the driving forces of the supercomputers of that period, today that role has been taken on by common off-the-shelf RISC processors and Intel x86 compatible processors. In fact there are only two companies left that produce vector systems while all other systems that are offered are based on RISC CPUs or x86-like (except the Cray MTA-2). Therefore it is useful to give a brief description of the main processors that populate the present supercomputers and look a little ahead to the processors that will follow in the coming year.

The RISC processor scene has shrunked drastically in the last year. Although the HP/Compaq EV7 and the HP PA-RISC processors are still produced, systems containing the former chip will only be marketed until 2006 due to HP's decision to replace these processors by the Itanium line of processors. The PA-RISC 8800 is still made for the Superdome 9000 machine but will also disappear shortly, although officially a PA-RISC 8900 still will be produced.

Although at this moment the RISC processors still play a role many of these will disappear from the scene: the withdrawal of systems containing the Alpha, PA-RISC and MIPS processors is announced all in favour of the Itanium processor product line. The disappearance of RISC processor families demonstrates a trend that is both worrying and interesting: worrying because the diversity in the processor field is decreasing severely and, with it, the choice for systems in this sector. On the other hand there is the trend to enhance systems having run-of-the-mill processors with special-purpose add-on processors in the form of preconfigured FPGAs or DSPs because their possibilities in performance, price level, and ease of use has improved to a degree that they offer attractive alternatives for certain application fields.

The RISC processors generally have a clock frequency that is lower than that of the Intel Pentium 3/4 processors or the corresponding AMD Intel look-alikes. However, they have a number of facilities that put them ahead in the speed of floating-point oriented applications. Firstly, all RISC processors are able to deliver 2 or more 64-bit floating-point results in one clock cycle. Secondly, all of them feature out-of-order instruction execution, which enhances the number of instructions per cycle that can be processed (although the newer AMD processors also have 2-way floating-point instruction issuing and out-of-order execution, they are somewhat limited by their adherence to the Intel x86 instruction set). The bandwidth from the processor to the memory, in case of a cache miss, used to be larger than that of the Intel(-like) processors. For the AMD Opteron, however, this is not the case anymore, because of its on-chip memory controller and HyperTransport bus from memory to CPU. Notwithstanding the commonalities between the various RISC processors, there are also differences in instruction latencies, number of instructions processed, etc., which we will address below. We provide block diagrams for each of the processors to give a schematic idea of their structure. However, these figures do not reflect the actual layout of the devices on the respective chips.

2.8.1 AMD Opteron

The Opteron (long known by its code name Hammer) is the newest processor from AMD and the successor of the Athlon processor. The first versions have become available early 2003. As it is, like the Athlon, a clone with respect to Intel's x86 Instruction Set Architecture, it will undoubtedly frequently be used in clusters. Therefore we discuss this processor here although it is not used presently in integrated parallel systems. The Opteron processor has many features that are also present in modern RISC processors: it supports out-of-order execution, has multiple floating-point units, and can issue up to 9 instructions simultaneously.

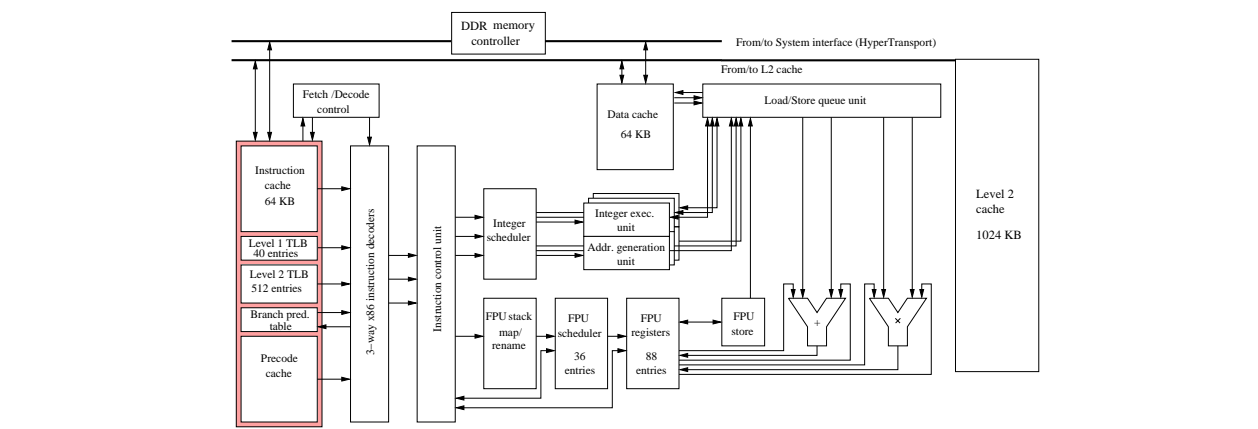


Figure 2.7: *Block diagram of the AMD Opteron processor.*

In fact, the processor core is very similar to that of the Athlon processor. A block diagram of the processor is shown in Figure 2.7.

It shows that the processor has three pairs of Integer Execution Units and Address Generation Units that via a 24-entry Integer Scheduler takes care of the integer computations and address calculations. Both the Integer Scheduler and the Floating-Point Scheduler are fed by the 96-entry Instruction Control Unit that receives the decoded instructions from the instruction decoders. An interesting feature of the Opteron is the pre-decoding of x86 instructions in fixed-length macro-operations, called RISC Operations (ROPs), that can be stored in a Pre-decode Cache. This enables a faster and more constant instruction flow to the instruction decoders. In comparison to the Athlon the instruction decode pipeline has been deepened by two stages to 12 to enable a higher yield of ROPs. Like in RISC processors, there is a Branch Prediction Table assisting in branch prediction.

The floating-point units allow out-of-order execution of instructions via the FPU Stack Map & Rename unit. It receives the floating-point instructions from the Instruction Control Unit and reorders them if necessary before handing them over to the FPU Scheduler. The Floating-Point Register File is 88 elements deep which approaches the number of registers as is available on RISC processors¹.

The floating-point part of the processor contains three units: a Floating Store unit that stores results to the Load/Store Queue Unit which also performs some miscellaneous floating-point operations, and Floating Add and Multiply units that can work in superscalar mode, resulting in two floating-point results per clock cycle. Because of the compatibility with Intel's Pentium 4 processors, the floating-point units also are able to execute Intel SSE2/3 instructions and AMD's own 3DNow! instructions. However, there is the general problem that such instructions are not accessible from higher level languages, like Fortran 90 or C(++). Both instruction sets are meant for massive processing of visualisation data and only allow for 32-bit precision to be used.

Due to the shrinkage of components the chip now can harbour the secondary cache of up to 1024 KB with an 8-cycle latency and the memory controller. This, together with a significantly enhanced memory bus can deliver up to 6.4 GB/s of bandwidth, an enormous improvement over the former memory system. This memory bus, called HyperTransport by AMD, is derived from licensed Compaq technology and similar to that employed in HP/Compaq's EV7 processors. It allows for "glueless" connection of several processors to form multi-processor systems with very low memory latencies.

The clock frequency is in the range of 1.8–2.2 GHz which makes the Opteron an interesting alternative for many of the RISC processors that are available at this moment. Especially the HyperTransport interconnection possibilities makes it highly attractive for building SMP-type clusters. A very large (over 11,000 processor) Opteron cluster, the Red Storm system, will be built at Sandia National Laboratory by Cray Inc. which will use the HyperTransport facility because it is an open standard that makes the connection with the Cray network much simpler.

¹For the x86 instructions 16 registers in a flat register file are present instead of the register stack that is usual for Intel architectures.




Figure 2.8: Block diagram showing the functional units in a Compaq EV7 processor (a). The MVI/PLZ unit performs population counts and trailing or leading (integer) zero counts. In (b) the chip layout for the complete processor is shown.

2.8.2 HP/Compaq Alpha EV7

Since the beginning of 2003 in Compaq machines like the AlphaServerSC, GS1280, and various cluster systems can be fitted with the Alpha EV7 processor. The latest implementation is the EV79, made in in $0.13 \mu\text{m}$ technology. Because of the EV7 structure the macro-architecture of these systems may also significantly change (see below). The core of the EV7 processor is almost identical to that of the EV6x architecture and is depicted in Figure 2.8a.

A notable fact is that there are *two* duplicate integer register files both with 80 entries, that each service a set of integer functional units called cluster 0 and cluster 1, respectively, by Compaq. The four integer Add/Logical units can exchange values in one cycle if required. Although this is not shown in the diagram, the integer multiply is fully pipelined. The two integer clusters and the two floating-point units enable the issuing of up to 6 instructions simultaneously. The two load/store units draw on a 64 KB instruction and a 64 KB data cache that are both 2-way set-associative and have a latency of 2 cycles. Four instructions can be accepted for (speculative) processing. Of the 80 integer and 72 floating-point registers 41 in both register files can hold speculative results. The out-of-order issuing of instructions is supported via an integer queue of length 20 and a floating-point queue with 15 entries. However, as the integer processing clusters do not contain the same functional units, the issuing of integer instructions cannot all be scheduled dynamically. Those instructions that need to execute in a particular unit (e.g., an integer multiply that is only available in cluster 0) are scheduled statically. As soon as an instruction is issued or is terminated due to mis-speculation it is removed from the queue and can be replaced by another instruction. Instruction fetching is governed by the branch predictor. This hardware contains global and local prediction tables and Branch History Tables (BHTs) to train the predictor in order to obtain an optimal instruction fetch to the instruction cache and registers.

The feature density now used is $0.13 \mu\text{m}$ instead of $0.18 \mu\text{m}$ which enables the location on the chip of a 1.75 MB secondary cache which is 7-way set-associative and has a 12-cycle latency. Also there are 2 memory controllers on chip. The largest difference with its predecessor, the EV68, is that there are 4 dual channels (North, East, South, West) from the chip to interconnect it with neighbouring chips at a bandwidth of 6.4 GB/s per single channel for what HP calls “switchless SMP processing” and is, as the name suggests, well-suited to build SMP nodes with low memory latency. The layout of the complete chip is shown in Figure 2.8b.

The path to memory is implemented by 4×5 RAMBUS links as the systems is fitted with RAMBUS memory

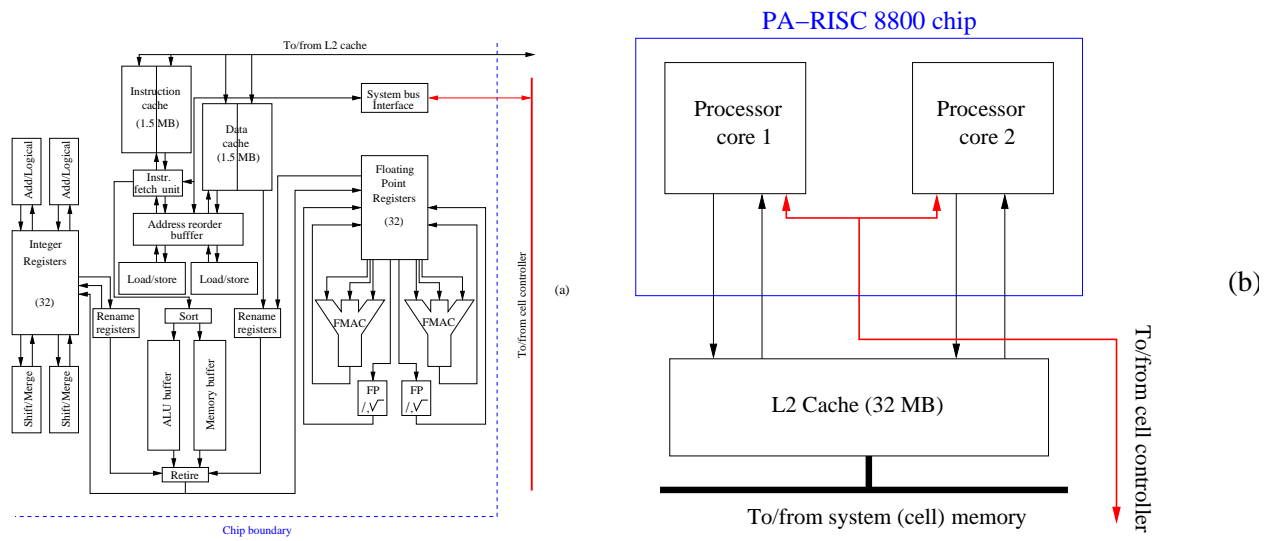


Figure 2.9: Block diagram showing the functional units in a HP PA-RISC 8800 CPU (a) and embedding on chip (b).

yielding a bandwidth of 6 GB/s per memory controller, so 12 GB/s in total. The direct I/O dual link from the chip also has a bandwidth of 1.6 GB/s. The fastest EV79 implementation is shipped at a clock frequency of 1.3 GHz and is expected to still rise slightly in later implementations.

2.8.3 Hewlett Packard PA-RISC 8800

The computational power for the Hewlett Packard systems, like the Superdome is delivered by the PA-RISC-8800 chip. Since February this year it replaces the PA-RISC 8700+ which was as the naming suggests essentially the same as the PA-8700 just has a higher clock rate, 875 MHz instead of 750 MHz. The PA-8800 is different: it follows the trend to put two CPUs on a chip. This could be done because of the denser technology: 130 nm instead of 180 nm. This also enabled to raise the clock cycle to 1 GHz (there is also a slower 900 MHz variant). Furthermore, the bandwidth has considerably increased: from 1.6 GB/s to 6.4 GB/s per chip. Note that this is identical to the frontside bus speed of the Itanium 2 (see section 2.8.7) which enables HP to exchange PA-RISC chips for Itanium 2 chips. The CPU cores on the chip are almost shrunken down versions of the former PA-8700(+). There is one difference though: the L1 instruction cache has doubled from 750 KB to 1.5 MB in two parts like the data cache was already implemented.

As there are now 2 CPUs on a chip the net bandwidth increase per CPU is a factor of 2. The larger HP systems like the Superdome is built from 4-processor cells and the chips are commonly connected to the system memory in a cell and to the cell controller that makes a cell into an SMP node and takes care of the communication with other cells, if present. The block diagram of a CPU and the embedding of the CPUs in the PA-8800 chip are shown in Figure 2.9. A peculiarity of the PA-8x00 chips was the absence of a secondary cache. This was compensated for by a large L1 cache is implemented: 1.5 MB instruction cache and 1.5 MB data cache. Both are 4-way set associative. The absence of the L2 cache has been amended in the PA-8800. There is now a large 32 MB L2 cache off-chip.

From the PA-8600 on the shrinking of the logic has allowed to put the L1 caches on-chip. The latency of the caches is two cycles. To ensure data to be shipped to the registers every cycle, the load/store units work “out-of-phase”. So, one unit loads from odd banks of the data cache while the other loads from the even banks. The Address Reorder Buffer sets the priority for the loads and tries to load from the alternate halves every cycle.

Like all advanced RISC processors the PA-8800 has out-of-order execution, the sequence of instructions being determined by the instruction reorder buffer (IRB) which contains an ALU buffer that drives the computational functional units and a memory buffer that controls the load/store units. When speculative

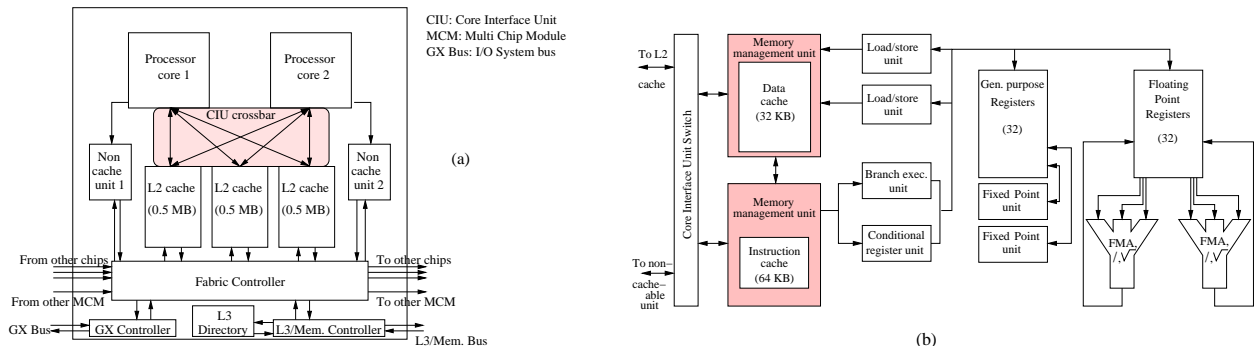


Figure 2.10: Diagram of the IBM POWER4+ chip layout (a) and of the processor core (b).

branches have been mis-predicted the dependent instructions are retired from the IRB and new candidate instructions replace them. Branch prediction is controlled through the branch history table (BHT) but, in addition to this dynamic branch prediction, a static branch prediction can be performed at the compiler level or by execution traces of former executions of a program. The BHT was rather small before the PA-8600 but is now enlarged significantly to 2048 entries to get better prediction results. Also the Translation Lookaside Buffer (a component of the load/store units not shown in Figure 2.9) was enlarged to 160 entries for a more effective address translation. Also there is a pre-fetch capability in the PA-8800 from the data cache.

As can be seen in Figure 2.9a, there are 2 floating-point units which each can deliver 2 flops per cycle but only when the operation is in the `axpy` form $x = x + \alpha y$. This is called a Floating Multiply Accumulate instruction (FMAC) by HP. At a clock frequency of 1 GHz this leads to a theoretical peak performance of 4 Gflop/s/CPU, so, 8 Gflop/s/chip. However, when the operations occur in another order or with another composition, 1 flop per cycle per floating-point unit can be executed with a correspondingly lower flop rate.

According to some of HP's roadmaps one new generation of the PA-8x00 family is still projected: the PA-8900 that will be on the market concurrently with the IA-64 Itanium Montecito. The signals are somewhat confusing in that respect: sometimes is also stated that the PA-RISC family will stop at this latest PA-8800 chip (which we think to be more likely).

2.8.4 IBM POWER4+

In the systems that feature as IBM's supercomputer line, the p690 series the nodes still contain the POWER4+ chip as the computational engine although the successor, the POWER5 chip already is available in other server lines of IBM. Undoubtedly the POWER5 will eventually appear in the p690-type systems but presently this is not yet the case. So, we discuss both chips, the POWER4+ and the POWER5 as the successor that will replace it.

At the time of writing, the clock frequency of the POWER4+ is in the range 1.5–1.9 GHz. It is a boost of the original POWER4 chip and has no new features other than the higher clock frequency. The chip size has become so large (or rather the feature size has become so small) that IBM places now two processor cores on one chip as shown in Figure 2.10a. The chip also harbours 1.5 MB of secondary cache divided over three modules of 0.5 MB each. The L2 cache module are connected to the processors by the Core Interface Unit (CIU) switch, a 2×3 crossbar with a bandwidth of 40 B/cycle per port. This enables to ship 32 B to either the L1 instruction cache or the data cache of each of the processors and to store 8 B values at the same time. Also, for each processor there is a Non-cacheable Unit that interfaces with the Fabric Controller and that takes care of non-cacheable operations. The Fabric Controller is responsible for the communication with three other chips that are embedded in the same Multi Chip Module (MCM), to L3 cache, and to other MCMs. The bandwidths at 1.3 GHz are 13.6, 9.0, and 6.8 GB/s, respectively. The chip further still contains a variety of devices: the L3 cache directory and the L3 and Memory Controller that should bring down the off-chip latency considerably, the GX Controller that responsible for the traffic on the GX bus. This bus transports data to/from the system and in practice is used for I/O. The maximum size of the L3 cache is huge: 32 MB. Some of the integrated devices, like the Performance Monitor, and logic for error detection

and logging are not shown in Figure 2.10a.

A block diagram of the processor core is shown in Figure 2.10b. In many ways the POWER4+ processor core is similar to the former POWER3 processor: there are 2 integer functional units instead of 3 (called Fixed Point Units by IBM) and instead of a fused Branch/Dispatch Unit, the POWER4 core has a separate Branch and Conditional Register Unit, 8 execution units in all. Oddly, the instruction cache is two times larger than the data cache (64 KB direct-mapped vs. 32 KB two-way set associative, respectively) and all execution units have instruction queues associated with them that enables the out-of-order processing of up to 200 instructions in various stages. Having so many instructions simultaneously in flight calls for very sophisticated branch prediction facilities. Instructions are fetched from the Instruction Cache under control of the Instruction Fetch Address Register which in turn is influenced by the branch predict logic. This consists of a local and a global Branch History Table (BHT), each with 16 K entries and a so-called selector table which keeps track of which of the BHTs has functioned best in a particular case in order to select the prediction priority of the BHTs for similar cases coming up.

Unlike in the POWER3, the fixed point units performs integer arithmetic operations that can complete in one cycle as well as multi-cycle operations like integer multiply and divide. There are no separate floating-point units for operations that require many cycles like divisions and square roots. All floating-point operations are taken care of in the FP units and, like in the HP PA-8800, there is an instruction to accommodate the `axpy` operation, called Fused Multiply Add (FMA) at IBM's which could deliver 2 floating-point results every cycle. This brings the theoretical peak performance at 7.6 Gflop/s at the current clock frequency. Like in the HP processor, the composition of the floating-point operations should be such that the units have indeed enough FMAs to perform. Otherwise the performance drops by a factor of 2.

Although here the dual core version of the chip is described that is positioned for general processing, also a single core version is marketed that is recommended for HPC use. The reason is that in this case the bandwidth from the L2 cache does not have to be shared between the CPUs and a contention-free transfer of up to 108.8 GB/s can be achieved while in the dual core version a peak bandwidth of 163.2 GB/s is to be shared between both CPUs.

It is interesting to see that presently three vendors (AMD, HP/Compaq, and IBM, although HP is not using it at present) have facilities that enable glueless coupling of processors although the packaging and implementation is somewhat different. All implementations allow for low-latency SMP nodes with a considerable number of processors stimulating the trend to build parallel systems based on SMP nodes.

2.8.5 IBM POWER5

As remarked in section 2.8.4 the POWER5 chip will replace the POWER4+ in the near future and we therefore present some features of this chip here. Presently it runs at the same speed as the fastest POWER4+: 1.9 GHz. Like the POWER4(+) the POWER5 has two CPUs on a chip. These CPU cores have a similar structure as shown in Figure 2.10b. There is however a significant difference between the chips of the POWER4 and the POWER5. Because of the higher density on the chip (the POWER5 is built in 130 nm technology instead of 180 nm used for the POWER4+) more devices could be placed on the chip and they could also be enlarged. First the L2 caches of two neighbouring chips are connected and the L3 caches are directly connected to the L2 caches. Both are larger than their respective counterparts of the POWER4: 1.875 MB against 1.5 MB for the L2 cache and 36 MB against 32 MB for the L3 cache. In addition the speed of the L3 cache has gone up from about 120 cycles to 80 cycles. Also the associativity of the caches has improved: from 2-way to 4-way for the L1 cache, from 8-way to 10-way for the L2 cache, and from 8 to 12-way for the L3 cache. A big difference is also the improved bandwidth from memory to the chip: it has increased from 4 GB/s for the POWER4+ to ≈ 16 GB/s for the POWER5. We show the embedding of the POWER5 in Figure 2.11. The better cache characteristics leads to less waiting time for regular data access: evaluation of a high order polynomial and matrix-matrix multiplication attain 90% or better of the peak performance while this is 65–75% on the POWER4+ chip. There is another feature of the POWER5 that does not help for regular data access but which can be of benefit for programs where the data access is not so regular: Simultaneous Multithreading (SMT). The POWER5 CPUs are able to keep two process threads at work at the same time. The functional units get instructions for the functional units from any of the two threads whichever is able to fill a slot in an instruction word that will be issued to the functional units. In this way a larger fraction of the functional units can be kept busy, improving the overall efficiency. For

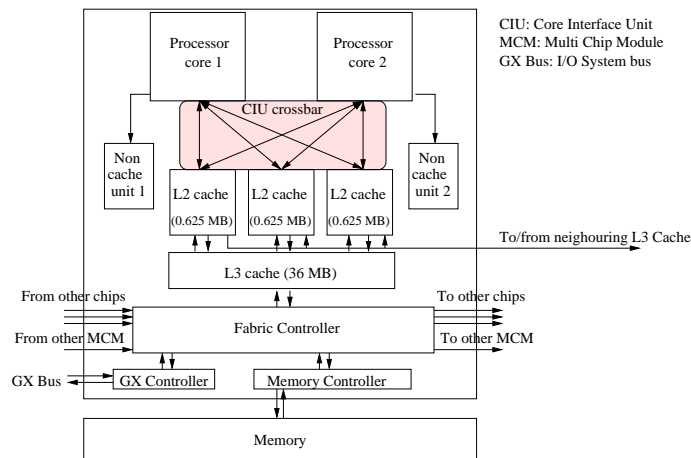


Figure 2.11: *The IBM POWER5 chip layout and embedding.*

very regular computations single thread (ST) mode may be better because in SMT mode the two threads compete for entries in the caches which may lead to trashing in the case of regular data access. Note that SMT is somewhat different from the “normal” way of multi-threading. In this case a thread that stalls for some reason is stopped and replaced by another process thread that is awoken at that time. Of course this takes some time that must be compensated for by the thread that has taken over. This means that the second thread must be active for a fair amount of cycles (preferably a few hundred cycles at least). SMT does not have this drawback but scheduling the instructions of both threads is quite complicated.

2.8.6 IBM BlueGene processor

This processor is in fact a modified PowerPC 400 processor which is made especially for the IBM BlueGene family. It runs presently at a speed of 700 MHz. Because of its rather specialised nature we will discuss when describing the BlueGene system proper in section 3.1.15.

2.8.7 Intel Itanium 2

The Itanium 2 is a representative of Intel’s IA-64 64-bit processor family and as such the second generation. Its predecessor, the Itanium, came out in 2001, but has not spread widely, primarily because the Itanium 2 would follow quickly with projected performance levels up to twice that of the first Itanium. The first Itanium 2 implementation ran at 0.8–1.0 GHz and has been followed quickly by a technology shrink (code name Madison) to the present-day Itanium 2 with clock frequencies in the range 1.3–1.6 GHz.

The Itanium family of processors has characteristics that are different from the RISC chips presented elsewhere in this section. A block diagram of the Itanium 2 is shown in 2.12.

We show here the fastest version that also harbours the largest L3 cache of 9 MB. Figure 2.12 shows a large amount of functional units that must be kept busy. This is done by large instruction words of 128 bits that contain 3 41-bit instructions and a 5-bit template that aids in steering and decoding the instructions. This is an idea that is inherited from the Very Large Instruction Word (VLIW) machines that have been on the market for some time about ten years ago. The two load/store units fetch two instruction words per cycle so six instructions per cycle are dispatched. The Itanium has also in common with these systems that the scheduling of instructions, unlike in RISC processors, is not done dynamically at run time but rather by the compiler. The VLIW-like operation is enhanced with predicated execution which makes it possible to execute instructions in parallel that normally would have to wait for the result of a branch test. Intel calls this refreshed VLIW mode of operation EPIC, Explicit Parallel Instruction Computing. Furthermore, load instructions can be moved and the loaded variable used before a branch or a store by replacing this piece of code by a test on the place it originally came from to see whether the operations have been valid. To keep track of the advanced loads an Advanced Load Address Table records them. When a check is made about

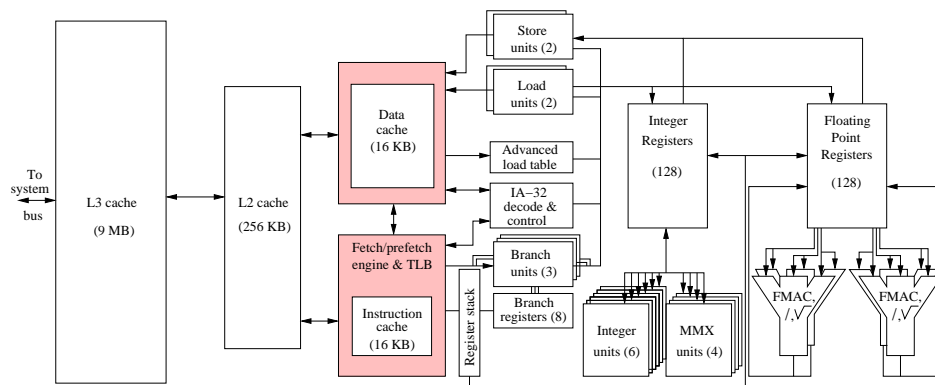


Figure 2.12: Block diagram of the Intel Itanium 2 processor.

the validness of an operation depending on the advanced load, the ALAT is searched and when no entry is present the operation chain leading to the check is invalidated and the appropriate fix-up code is executed. Note that this is code that is generated at compile time so no control speculation hardware is needed for this kind of speculative execution. This would become exceedingly complex for the many functional units that may be simultaneously in operation at any time.

As can be seen from Figure 2.12 there are four floating-point units capable of performing Fused Multiply Accumulate (FMAC) operations. However, two of these work at the full 82-bit precision which is the internal standard on Itanium processors, while the other two can only be used for 32-bit precision operations. When working in the customary 64-bit precision the Itanium has a theoretical peak performance of 6 Gflop/s at a clock frequency of 1.5 GHz. Using 32-bit floating arithmetic, the peak is doubled. In the first generation Itanium there were 4 integer units for integer arithmetic and other integer or character manipulations. Because the integer performance of this processor was modest, 2 integer units have been added to improve this. In addition four MMX units are present to accommodate instructions for multi-media operations, an inheritance from the Intel Pentium processor family. For compatibility with this Pentium family there is a special IA-32 decode and control unit.

The register files for integers and floating-point numbers is large: 128 each. However, only the first 32 entries of these registers are fixed while entries 33–128 are implemented as a register stack. The primary data and instruction caches are 4-way set associative and rather small: 16 KB each. This is the same as in the former Itanium processor. However, speed of the L1 cache is now doubled to full speed: data and instructions can now be delivered every clock cycle to the registers. Further more the secondary cache has been enlarged from 96 KB to 256 KB and it is 8-way set-associative. Moreover, the L3 cache is moved onto the chip and is no less than up to 9 MB. This cache structure greatly improves the bandwidth to the processor core, on average by a factor of 3. This does even more for the performance improvement than the increase in clock speed from 800 MHz to 1.6 GHz. Also the bandwidth from/to memory has increased by more than a factor of 3. The bus is now 128 bits wide and operates at a clock frequency of 400 MHz, totaling to 6.4 GB/s in comparison to 2.1 GB/s for its predecessor. As the frontside bus will not be enlarged, this speed will not scale with upcoming increases of the processor speed and so again may frustrate overall speedup in the near future.

The introduction of the first Itanium has been deferred time and again which quenched the interest for use in high-performance systems. With the availability of the Itanium 2 in the second half of 2002 the adoption has sped up. Apart from HP/Compaq also Bull, Fujitsu, NEC, SGI, and Unisys are offering now multiprocessor systems with this processor and the Alpha, PA-RISC, SPARC, and MIPS processors as employed by HP/Compaq, Fujitsu, and SGI respectively will be phased out in the next few years.

2.8.8 Intel Pentium 4/Xeon Nocona

Although Pentium processors are not applied in integrated parallel systems these days, they play a major role in the cluster community as most compute nodes in Beowulf clusters are of this type. Therefore we

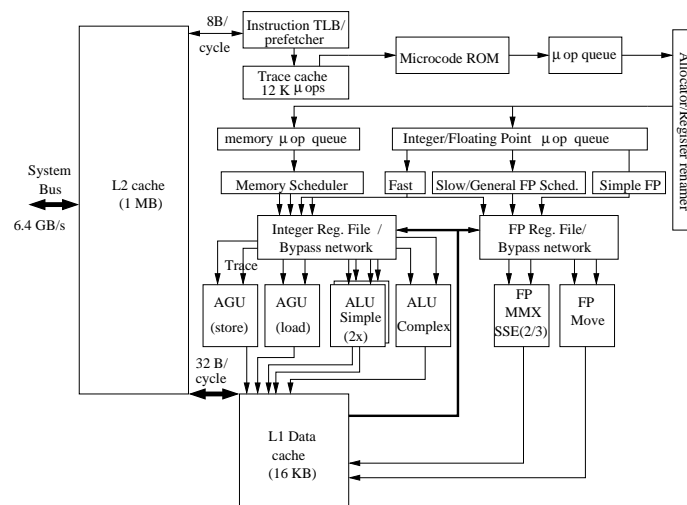


Figure 2.13: *Block diagram of the Intel Xeon Nocona processor.*

briefly discuss also this type of processor. In fact, mostly the Xeon processor is used as it can accommodate a larger cache than the Pentium 4 and has some provisions for use in a multiprocessor environment. As virtually all Intel-based clusters use 2-processor nodes, the Xeon is the most appropriate processor. Apart from the differences mentioned here the architecture of the Pentium 4 and the Xeon are identical.

Intel only provides scant technical information on its processor. Therefore, a rough block diagram of the P4 processor can only be synthesized from various sources. It is shown in Figure 2.13. We show here the Xeon variant with the large secondary cache that with the other additional features of the Nocona fits on chip because of the advanced 90 nm technology used to fabricate the chip. There are a number of distinctive features with respect to the earlier Pentium generations. There are two main ways to increase the performance of a processor: by raising the clock frequency and by increasing the number of instructions per cycle (IPC). These two approaches are generally in conflict: when one wants to increase the IPC the chip will become more complicated. This will have a negative impact on the clock frequency because more work has to be done and organised within the same clock cycle. Very seldomly chip designers succeed in raising both clock frequency and IPC simultaneously. Also in the Pentium 4 this could not be done. Intel has chosen for a high clock speed (initially about 40% more than that of the Pentium III with the same fabrication technology) while the IPC decreased by 10–20%. This still gives a net performance gain even if other changes would have been made to the processor. To sustain the very high clock rate that the present processors have, currently up to ≈ 3.8 GHz, a very deep instruction pipeline is required. The instruction pipeline has no less than 31 stages, where the Pentium III had 10. Although this favours a high clock rate, the penalty for a pipeline miss (e.g., a branch mis-predict) is much heavier and therefore Intel has improved the branch prediction by increasing the size of the Branch Target Buffer from 0.5 to 4 KB. In addition, the Pentium 4 has an execution trace cache which holds partly decoded instructions of former execution traces that can be drawn upon, thus foregoing the instruction decode phase that might produce holes in the instruction pipeline. The allocator dispatches the decoded instructions, “micro operations”, to the appropriate μop queue, one for memory operations, another for integer and floating-point operations.

Two integer Arithmetic/Logical Units are kept simple in order to be able to run them at twice the clock speed. In addition there is an ALU for complex integer operations that cannot be executed within one cycle. There is only one floating-point functional unit that delivers one result per cycle. However, besides the normal floating-point unit, there also are additional units that execute the Streaming SIMD Extensions 2 and 3 (SSE2/3) repertoire of instructions, a 144-member instruction set, that is especially meant for multimedia, and 3-D visualisation applications. The length of the operands for these units is 128 bits. The Intel compilers have the ability to address the SSE2/3 units. This makes it in principle possible to achieve a two times higher floating-point performance.

The Xeon Nocona boast so-called Hypertreading: with the processor two threads can run concurrently under some circumstances. This may for instance be used for speculative execution of `if` branches. Experiments

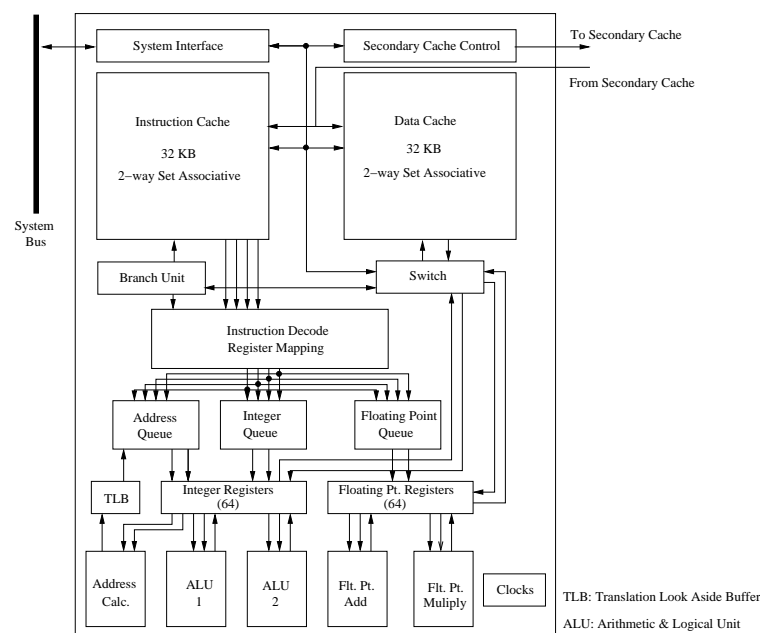


Figure 2.14: Block diagram of the MIPS R16000 processor.

have shown that up to 30% performance improvements can be attained for a variety of codes. In practice the performance gain about 3–5%, however.

The primary cache was quite small by today's standards: 8 KB. This has been doubled to 16 KB since the so-called Prescott implementation of the processor, however, at the cost of a higher latency in shipping data to the functional units. Where it was 2 cycles before, it has now increased to 3 cycles.

The secondary cache has a size of 1 MB for the Nocona processor. It has a wide 256-bit bus run at 800 MHz, which amounts to a bandwidth of 6.4 GB/s. This enables the use of fast DDR-II memory making the memory-CPU speed gap somewhat smaller.

The largest difference with the former processors, however, is the ability to run (and address) 64-bit codes, thereby following AMD, in fact copying the approach used in the AMD Opteron and Athlon processors. The technique is called Extended Memory 64 Technology (EM64T) by Intel. In principle it uses “unused bits” from in the instruction words of the x86 instruction set to signal whether an 64-bit version of an instruction should be executed. Of course some additional devices are needed for operating in 64-bit mode. These include 8 new general purpose registers (GPRs), 8 new registers for SSE2/3 support, and 64-bit wide GPRs and instruction pointers.

It will depend heavily on the quality of the compilers whether they will be able to take advantage of all the facilities present in the Nocona processor (Intel claims that a 30% performance improvement is possible). If they can, the processor could be an interesting basis for HPC clusters.

2.8.9 MIPS R16000

The essentials of the MIPS R1x000 series of processors have not changed since the introduction of the first in this family, the R10000. The current processor that is at the heart of the SGI Origin3000 series is the R16000. The R16000 is similar to the preceding R14000(A) except for the clock cycle: this is presently 800 MHz and as such the lowest of all RISC processors employed in high-performance systems. A block diagram of this processor is given in Figure 2.14. The R16000 is a typical representative of the modern RISC processors that are capable of out-of-order and speculative instruction execution. Like in the Compaq Alpha processor there are two independent floating-point units for addition and multiplication and, additionally, two units that

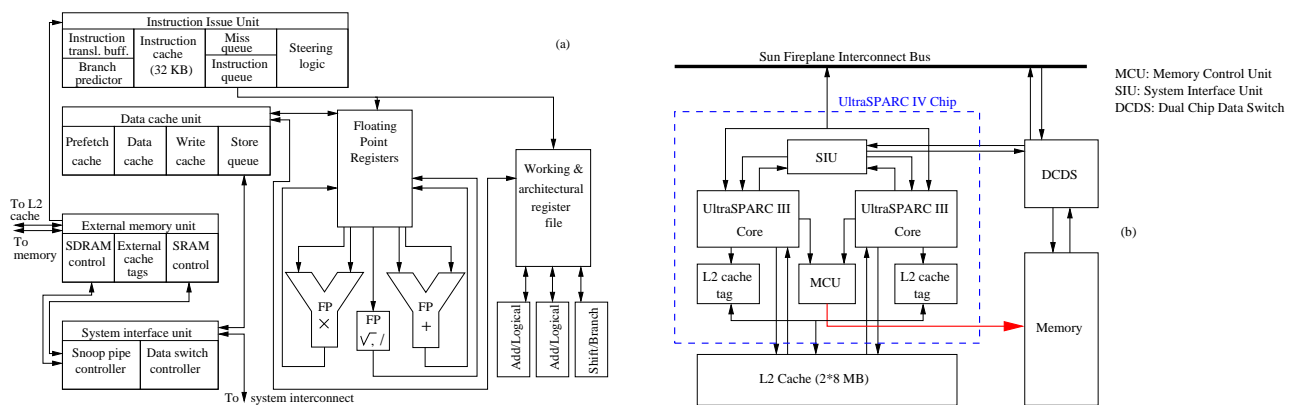


Figure 2.15: Block diagram of the UltraSPARC IV processor core and chip layout.

perform floating division and square root operations (not shown in Figure 2.14). The latter, however, are not pipelined and with latencies of about 20–30 cycles are relatively slow. In all there are 5 pipelined functional units to be fed: an address calculation unit which is responsible for address calculations and loading/storing of data and instructions, two ALU units for general integer computation and the floating-point add and multiply pipes already mentioned.

The level 1 instruction and data caches have a moderate size of 32 KB and are 2-way set-associative. In contrast, the secondary cache can be very large: up to 16 MB. Both the integer and the floating-point registers have a physical size of 64 entries, however, 32 of them are accessible by software while the other half is under direct CPU control for register re-mapping.

The clock frequency of the MIPS R1x000 processors have always been on the low side. The first R10000 appeared at a frequency of the 180 MHz while in the R16000 the clock cycle is started at 700 MHz and has risen slightly during to 800 MHz at present. With the 800 MHz frequency the theoretical peak performance is 1.6 Gflop/s. Because of the independent floating-point units without fused multiply-add capabilities often a fair fraction of that speed can be realised. There also have been made some improvements with respect to the earlier chips: the bus speed has been doubled from 100 MB/s to 200 MB/s and the L1 cache that ran at a 2/3 speed in the predecessor R12000 has been sped up to full speed from the R14000 on.

The R16000 is built in advanced 0.11 μm technology and it has at the present 800 MHz clock frequency an extremely low power consumption: only 20 Watt, several factors lower than that of the other processors discussed here. SGI keeps the clock frequency intentionally as low as possible to enable to build “dense” systems that can accommodate a large amount of processors in a small volume.

Recently it has become evident that SGI will migrate quickly from its MIPS platform, the Origin3000 series entirely to the Itanium-based Altix family. Therefore it is doubtful whether successors of the R16000, the R18000 and the R20000 that were originally projected for mid-2004 and later will ever be realised.

2.8.10 Sun UltraSPARC IV

The UltraSPARC IV is the fourth generation of the UltraSPARC family. Like the HP PA-RISC 8800 and the IBM POWER4/5 also Sun has put two processor cores on a chip. The CPU cores are in fact slightly modified UltraSPARC III processors. It is built in 0.13 μm CMOS technology at a clock frequency that is currently 1.2 GHz. UltraSPARCs are used in all SUN products from workstations to the heavy Fire 3800-15K servers. We show a block diagram of the processor core and its embedding in the UltraSPARC IV chip in Figure 2.15.

The processor is characterised by large large amount of caches of various sorts as can be seen in Figure 2.15a. The Data Cache Unit (DCU) contains apart from a 4-way set associative cache of 64 KB also a write and a pre-fetch cache, both of 2 KB. All these L1 caches operate at half speed: loads and stores from the processor can be done in 2 cycles. The prefetch cache is independent from the data cache and can load data when this is deemed appropriate. The write cache defers writes to the L2 cache and so may evade unnecessary writes of individual bytes until entire cache lines have to be updated. The Instruction Issue

Unit (IIU) contains the 32 KB 4-way set associative instruction cache together with the instruction TLB which is called Instruction translation buffer in SUN's terminology. The IIU also contains a so-called miss queue that holds instructions that are immediately available for the execute units when a branch has been mis-predicted. Branch prediction is fully static in the UltraSPARC-III. It is implemented as a 16 KB table in the IIU that is pipelined because of its size.

The Integer Execute Unit (IEU) has two Add/Logical Units and a branch unit. Integer adds and multiplies are pipelined but the divide operation is not. It is performed by an Arithmetic Special Unit (not shown in the figure) that does not burden the pipelines for the ALUs. The integer register file is effectively divided in two and is called the Working and Architectural Register File by SUN. Operands are accessed and results stored in the working registers. When an exception occurs, the results to be undone in the working registers are overwritten by those from the architectural file. One of the enhancements with respect to the original UltraSPARC III design is the adding of hash indexing for the write cache. This should decrease the number of write misses and thus leave more write store bandwidth for results that need storing.

The floating-point unit (FPU) has two independent pipelined units for addition and multiplication and a non-pipelined unit for floating division and square-root computation that require in the order of 20–25 cycles. The FPU also contains graphics hardware (not shown in Figure 2.15) that shares the pipelined adder and multiplier with general 64-bit calculations. For the chips delivered at 1.2 GHz, the theoretical peak performance is 2.4 Gflop/s per processor core. It is expected that the UltraSPARC-IV technology can be shrunk to reach a clock frequency that is slightly more by the end of its life cycle. In the UltraSPARC IV the FPUs are enhanced by adding hardware support of handling for IEEE 754 floating-point errors (which can be very costly otherwise when properly handled).

The size of the L2 cache is 2×8 MB each part being targeted by one of the processor cores. The L2 cache lines have been shortened from 512 B to 128 B in the UltraSPARC IV to reduce data contention. As is evident from Figure 2.15b the Memory Control Unit (MCU) is on chip as well as the L2 cache controller (in the MCU) and the L2 cache tags. This shortens the latency of accesses from both memory levels. In addition, both controllers communicate with the System Interface Unit (SIU), also on-chip to keep in touch with the snoop pipe controller in the SIU. The processor has been built with multi-processing in mind and the snoop controller keeps track of data requests in the whole system to ensure coherency of the caches when required.

The UltraSPARC IV is around since February 2004. Sun refers to having the two processor cores on a chip and running one execution thread on each of them as Chip Multithreading (CMT). This is not quite what one would normally would understand as multi-threading because one would then expect more execution threads per processor core. So, the CMT terminology is somewhat confusing and one would hope that Sun will drop it in favour of the common use of the term.

At the moment it is not entirely clear what lays ahead with respect to the SPARC family. On one hand the cooperation with Fujitsu-Siemens has intensified and Sun will market their systems with the SPARC64 implementation of Fujitsu-Siemens. The SPARC64 line is the "classic" in the sense that it is a standard RISC design be it that the end of 2004 dual core versions will reach the market. On the other hand Sun has shelved its own plans to produce UltraSPARC V and VI processor by April 2004 in favour of processor designs with many (≥ 8) processor cores, each capable of handling several execution threads. The first of these processors, named Niagara, is in the testing phase. It has 8 cores where each core can execute 4 threads. It is probably ready for the market by the end of 2005. Sun wants to pursue this kind of processing, dubbed "throughput processing", also in next generations: the so-called Rock processor will contain more cores and more threads per core. It is expected that some kind of specialisation in the cores will be introduced because it would be difficult to manage such massive execution thread bundles optimally.

2.9 Networks

Fast interprocessor networks are together with fast processors the decisive factors for both good integrated parallel systems and clusters. In the early days of clusters the interprocessor communication, and hence the scalability of applications, was hampered by the high latency and the lack of bandwidth of the network (mostly Ethernet) that was used. This situation has changed very much and to give a balanced view of the possibilities opened by the improved networks a discussion of some of these networks is in order. The more

so as some of these networks are, or have been employed also in “integrated” parallel systems. Of course Gigabit Ethernet is now amply available and with a maximum theoretical bandwidth of 125 MB/s would be able to fulfill a useful role for some applications that are not latency-bound in any way. But with latencies in the order of somewhat less than 100 μ s and in-switch latencies of 35–45 μ s the applicability is too restricted to be the network of choice (except perhaps for price reasons).

We restrict ourselves here to networks that are independently marketed as the proprietary networks for systems like those of IBM and SGI are discussed together with the systems in which they are incorporated. We do not pretend to be complete because in this new field players enter and leave the scene at a high rate. Rather we present main developments which one is likely to meet when one scans the high-performance computing arena.

A complication with the fast networks offered for clusters is the connection with the nodes. Where in integrated parallel machines the access to the nodes is customised and can be made such that the bandwidth of the network matches the internal bandwidth in a node, in clusters one has to make do with the PCI bus connection that comes with the PC-based node. The type of PCI bus which ranges from 32-bit wide at 33 MHz to 64-bit wide at 66 MHz determines how fast the data from the network can be shipped in and out the node and therefore the maximum bandwidth that can be attained in internode communication. In practice the available bandwidths are in the range 110–480 MB/s. Since 1999 PCI-X is available, initially at 1 GB/s, in PCI-X 2.0 also at 2 and 4 GB/s. Coupling with PCI-X is now common in PC nodes that are meant to be part of a cluster. Recently PCI Express became available. This provides a 200 MB/s bandwidth per data lane where 1 \times , 2 \times , 4 \times , 8 \times , 12 \times , 16 \times , and 32 \times multiple data lanes are supported: this makes it amply sufficient for the host bus adapters of any communication network vendor. So, for the networks discussed below often different bandwidths are quoted, depending on the PCI bus type and the supporting chip set which bandwidth could be attained. Therefore, when speeds are quoted it is always with the proviso that the PCI bus of the host node is sufficiently wide/fast.

2.9.1 Infiniband

Infiniband is rapidly becoming a widely accepted medium for internode networks. The specification was finished in June 2001. From 2002 on a number of vendors has started to offer their products based on the Infiniband standard. A very complete description (1200 pages) can be found in [34] Infiniband is employed to connect various system components within a system. Via Host Channel Adapters (HCAs) the Infiniband fabric can be used for interprocessor networks, attaching I/O subsystems, or to multi-protocol switches like Gbit Ethernet switches, etc. Because of this versatility, the market is not limited just to the interprocessor network segment and so Infiniband is expected to become relatively inexpensive because a higher volume of sellings can be realised. The characteristics of Infiniband are rather nice: there are product definitions both for copper and glass fiber connections, switch and router properties are defined and for high bandwidth multiple connections can be employed. Also the way messages are broken up in packets and reassembled as well as routing, prioritising, and error handling are all described in the standard. This makes Infiniband independent of a particular technology and it is, because of its completeness, a good basis to implement a communication library (like MPI) on top of it.

Conceptually, Infiniband knows of two types of connectors to the system components, the Host Channel Adapters (HCAs), already mentioned, and Target Channel Adapters (TCAs). The latter are typically used to connect to I/O subsystems while HCAs does more concern us as these are the connectors used in interprocessor communication. Infiniband defines a basic link speed of 2.5 Gb/s (312.5 MB/s) but also a 4 \times and 12 \times speed of 1.25 GB/s and 3.75 GB/s, respectively. Also HCAs and TCAs can have multiple ports that are independent and allow for higher reliability and speed.

Messages can be sent on the basis of Remote Memory Direct Access (RDMA) from one HCA/TCA to another: a HCA/TCA is permitted to read/write the memory of another HCA/TCA. This enables very fast transfer once permission and a write/read location are given. A port together with its HCA/TCA provide a message with a 128-bit header which is IPv6 compliant and that is used to direct it to its destination via cut-through wormhole routing: In each switching stage the routing to the next stage is decoded and send on. Short messages of 32 B can be embedded in control messages which cuts down on the negotiation time for control messages.

Infiniband switches for HPC are offered with 8–128 ports and for the moment always at a speed of 1.25 GB/s.




Figure 2.16: An 8×16 Clos network using 8 and 16 port crossbar switches to connect 64 processors.

The switches can be configured in any desired topology but in practice a fat tree topology is almost always preferred (see Figure 2.5b, section 2.5). It obviously depends on the quality of the MPI implementation put on top of the Infiniband specifications how much of the raw speed can be realised. A Ping-Pong experiments on Infiniband-based clusters with different MPI implementations has shown bandwidths of around 850 MB/s and an MPI latency of $< 7 \mu\text{s}$ for small messages. The in-switch latency is typically about 200 ns.

Presently, the price per port is still somewhat higher than for that of Myrinet, the market leader. However, when Infiniband would take on, the price can drop significantly and become a serious competitor for Myrinet (see 2.9.2. Because of the recent availability of PCI Express, 12-wide Infiniband connections would make technical sense. Whether it would also be commercially viable remains to be seen.

2.9.2 Myrinet

Myrinet is the present market leader in fast cluster networks. The Myricom company which sells Myrinet started in 1994 with its first Myrinet implementation, [27], as an alternative for Ethernet to connect the nodes in a cluster. Apart from the higher bandwidth, around 100 MB/s at that time, the main advantage was that it entirely operated in user space, thus avoiding Operating System interference and the delays that come with it. This meant that the latency for small messages was around 10–15 μs . Latency and bandwidth compared nicely with the proprietary networks of integrated parallel systems of Convex, IBM, and SGI at the time. Although such a network came at a non-negligible cost, in many cases it proved a valuable alternative to either an Ethernet connected system or an even costlier integrated parallel system.

Since then hardware upgrades and software improvements have made Myrinet the network of choice for many cluster builders and until recently there was hardly an alternative when a fast, low-latency network was required.

Like Infiniband, Myrinet uses cut-through routing for an efficient utilisation of the network. Also RDMA is used to write to/read from the remote memory of other host adapter cards, called Lanai cards. These cards interface with the PCI-X bus of the host they are attached to. The latest Myrinet implementation only uses fibers as signal carriers. This gives a high flexibility in the connection and much headroom in the speed of signals but the fiber cables and connectors are rather delicate which can lead to damage when cluster nodes have to be serviced.

Myrinet offers ready-made 8–256 port switches. The 8 and 16 port switches are full crossbars. In principle all larger networks are build form these using a Clos network topology. An example for a 64-port systems is shown in Figure 2.16. A Clos network is another example of a logarithmic network with the maximum bi-sectional bandwidth of the endpoints. Note that 4 ports of the 16×16 crossbar switches are unused but other configurations need either more switches or connections or both.

Myricom provides benchmark measurements of its networks at [26] which for the newest Myrinet2000 switches give quite good results: ≈ 230 MB/s for a Ping-Pong experiment with latencies around 7 μs for small messages with the MPI version that Myricom distributes with its networks.

2.9.3 QsNet

QsNet is the main (and only) product of Quadrics, a company that initially started in the 1990s as the British firm Meiko that made parallel systems called the Computer Surface 1 and 2. In the CS-2 a very good logarithmic network was included that was made in a separate product after Meiko closed down and the network part was taken over by the Italian company Alenia and put in the independent company Quadrics. The success came when Digital/Compaq chose QsNet as the network for its large high-performance computer products, the AlphaServer SC line. Some very large configurations of these machines were sold, e.g., at the Pittsburgh Supercomputing Centre and the French CAE. QsNet proved to be a very fast and reliable network and since about two years QsNet is also offered for cluster systems.

Like Infiniband and Myrinet the network has effectively two parts: the ELAN interface cards, comparable to Infiniband Host Bus Adapters or Myrinet's Lanai interface cards, and the Elite switch, comparable to an Infiniband switch/router or a Myrinet switch. The topology that is used is a (quaternary) fat tree like in most Infiniband switches, see Figure 2.5b, section 2.5 for an example. The ELAN card interfaces with the PCI-X port of a host computer.

Ready-made Elite switches for clusters come in two sizes: with 16 and 128 ports but nothing in between. Of course one can put together networks for other sizes but this goes at the cost of compactness and speed. A difference with the other switches lies in the providing *two* virtual bi-directional channels per link. Since the end of 2003 Quadrics sells its second generation QsNet, QsNet II. The structure, protocols, etc., are very similar to those of the former QsNet but much faster: where in [32] a speed of 300 MB/s for an MPI Ping-Pong experiment was measured while QsNet II has a link speed of 1.3 GB/s. In this case the PCI-X bus speed is a limiting factor: it allows for somewhat more than about 900 MB/s. However, with the advent of PCI Express this limitation may be over. Also the latency for short messages has improved from $\approx 5 \mu\text{s}$ to $3 \mu\text{s}$. Furthermore, the switch supports two priority levels which greatly helps in a fair distribution of message packets. This is more than Myrinet provides but less than the 16 priority levels of Infiniband.

Like Infiniband, QsNet(II) has RDMA capabilities that allow to write/read from/to remote memory regions on the ELAN cards. This can however be extended to memory regions of the host processor itself. So, in principle, one would be able to view a QsNet-connected system as a virtual shared memory system. As yet, this has not been realised nor on integrated systems, nor on clusters. Nevertheless, it could be an attractive alternative for the much slower memory page-based virtual shared memory systems like TreadMarks [1]. A Cray-style `shmem` library is offered that enables one-sided communication via `put` and `get` operations.

2.9.4 SCI

SCI, standing for Scalable Coherent Interface, is the oldest of the networks discussed here. SCI became an IEEE/ANSI standard by October 1992. It was born as a reaction to the limitations that were encountered in bus-based multiprocessor systems in these days and a working group from various vendors and from universities tried to devise a solution that would do away with these limitations once and for all. A nice introduction into the rationale and the design choices for SCI can be found in [16]. Further discussions can be found in [22].

SCI was largely designed to avoid the usual bus limitations which resulted in a ring structure to connect the hosts. In fact pairs of rings are used to enable reverse signaling between a send and receive host. This is because of one of the design principles of SCI: signals flow continuously in one direction enabling a very high signal rate without noise interference. A consequence is that the ring is always active, sending zero-length payload messages when no actual messages have to be transferred. The information length in a packet has a fixed sizes of 0, 16, 64, and 256 bytes and a header of 16 bytes or 32 bytes while the packet is closed by a 2-byte error correcting CRC code. By having the error correcting code at the end of the packet it is immediately known whether the data are corrupted or not and immediate action can be taken. The limited packet format enables fast reception and checking of the packets.

A special feature of SCI is its ability keep the caches of the processors it connects to coherent. This is also a consequence of its design history in the sense that SCI should be able to replace buses in multi-processor systems where such buses via a snoopy-bus protocol (see section 2.6) should keep the processor caches up-to-date or *coherent*. So, like with QsNet, one can use it to implement virtual shared memory. Unlike with QsNet, this is (have been) indeed been done: the late Convex Exemplar systems used SCI to connect its




Figure 2.17: *SCI networks arranged as 1-D, 2-D, and 3-D toruses.*

nodes, as did Data General. Presently the NUMA-Q systems of former Sequent, now IBM, use it as a memory interconnect medium. In clusters SCI is always used as just an internode network. Because of the ring structure of SCI this means that the network is arranged as a 1-D, 2-D, or 3-D torus as shown in Figure 2.17.

The torus network has some risks in comparison to other networks like the fat tree or the Clos network: when a node adapter fails it incapacitates all the nodes that share the ring on which it lies. So, for instance, Dolphin networks, one of the SCI vendors provides software to reconfigure the torus in such a way that the minimal number of nodes become unavailable. Furthermore, it is not possible to add or remove an arbitrary number of nodes in the cluster because of the torus topology.

Bandwidths are reported for the SCI-based clusters: up to about 320 MB/s for a Ping-Pong experiment over MPI with very low latencies of 1–2 μ s for small messages.

3 Recount of (almost) available systems

In this section we give a recount of all types of systems as discussed in the former section and that are marketed presently or will appear within 6 months from now. When vendors market more than one type of machine we will discuss them in distinct subsections. So, for instance, we will discuss Cray systems under entries, X1 and XD1 because they have a very different structure.

At the time of writing Compaq is now fully integrated into HP. Therefore the earlier Compaq systems are now listed under “HP/Compaq” to distinguish between the Alpha-based product line stemming from Compaq and the PA-RISC-based line originating from HP.

3.1 System descriptions

The systems are presented alphabetically. The “Machine type” entry shortly characterises the type of system as discussed in the former chapter: Processor Array, ccNUMA, etc.

3.1.1 The Bull NovaScale.

Machine type: ccNUMA system.

Models: NovaScale 5080, 5160.

Operating system: Linux, WindowsServer 2003, GCOS 8

Connection structure: Full crossbar.

Compilers: Intel’s Fortran 95, C(++).

Vendors information Web page: www.bull.com/novascale/

Year of introduction: 2002.

System parameters:

Model	NovaScale 5080	NovaScale 5160
Clock cycle	1.5 GHz	1.5 GHz
Theor. peak performance	48 Gflop/s	96 Gflop/s
Main Memory	4–64 GB	4–128 MB
No. of processors	8	16
Communication bandwidth		
Point-to-point	6.4 GB	6.4 GB
Aggregate peak	12.8 GB	25.6 GB

Remarks:

The availability of the Itanium 2 has spurred some vendors that are traditionally not active in the HPC business to try their hand in this area. One of these is Bull that markets its NovaScale ccNUMA SMPs with up to 16 nodes. The NovaScale systems are built from standard Intel Quad Building Blocks (QBBs) each housing 4 Itanium 2 processors and a part of the memory. The QBBs in turn are connected by Bull’s proprietary FAME Scalability Switch (FSS) providing an aggregate bandwidth of 25.6 GB. For reliability reasons a NovaScale 5160 is equipped with 2 FSSes. This ensures that when any link between a QBB and a switch or between switches fails the system is still operational, be it on a lower communication performance level. As each FSS has 8 ports and only 6 of these are occupied within a 5160 system, the remaining ports can be used to couple two of these systems thus making a 32-processor ccNUMA system. Larger configurations can be made by coupling systems via QsNet II (see 2.9.3). Bull provides its own MPI implementation which

turns out to be very efficient (see “Measured Performances” below and [42]).

A distinctive feature of the NovaScale systems is that they can be partitioned such that different nodes can run different operating systems and that repartitioning can be done dynamically. Although this is not particularly enticing for HPC users, it might be interesting for other markets, especially as Bull still has clients that use their proprietary GCOS operating system.

A smaller system, the NovaScale 4040, with 4 processor is also available as a departmental server. As Bull employs the Itanium 2, the Fortran 95 and C compilers from Intel are automatically available. Bull’s documentation gives no information about other HPC software that might be available but it should have all third-party software that has been ported to the Itanium 2 platform.

Measured Performances:

In the spring of 2004 rather extensive benchmark experiments with the EuroBen Benchmark were performed on a 16-processor NovaScale 5160 with the 1.3 GHz variant of the processor. Using the EuroBen benchmark, the MPI version of a dense matrix-vector multiply was found to be 13.3 Gflop/s on 16 processors while both for solving a dense linear system of size $N = 1,000$ and a 1-D FFT of size $N = 65,356$ speeds of 3.3–3.4 Gflop/s are observed (see [42]).

3.1.2 The C-DAC PARAM Padma.

Machine type: RISC-based Distributed-memory multi-processor.

Models: C-DAC PARAM Padma.

Operating system: AIX, Linux.

Connection structure: Clos network.

Compilers: Fortran 90, C, C++.

Vendors information Web page: www.cdacindia.com/html/parampma.asp

Year of introduction: 2003.

System parameters:

Model	PARAM Padma
Clock cycle	1 GHz
Theor. peak performance	
Per Proc.	4 Gflop/s
Maximal	1.13 Tflop/s
Memory	500 GB
No. of processors	248
Communication bandwidth	
Aggregate	4 GB/s
Point-to-point	312 MB/s

Remarks:

The PARAM Padma is the newest systems made by the Indian C-DAC. It is built somewhat asymmetrically from 54 4-processor SMPs and 1 32-processor node. All nodes employ 1 GHz IBM POWER4 processors. As an interconnection network C-DACs own PARAMnet-II is used for which a peak bandwidth of 2.5 Gb/s (312 MB/s) is given with a latency for short messages of $\approx 10\mu s$. The network is build from 16-port PARAMnet-II switches and has a Clos64 topology, very similar to the structure used by Myrinet (see 2.9.2). No MPI results over this network are available.

C-DAC has already a long tradition of building parallel machines and it has always provided its own software to go with them. Therefore, the Padma comes with Fortran 90, C(++), MPI, and a Parallel Files System.

Measured Performances:

The Padma performs at 532 Gflop/s with the HPC Linpack Benchmark ([44]) for a linear system of size $N = 224,000$ on a 62-node machine with a theoretical peak of 992 Gflop/s. That amounts to an efficiency of 53.6% for this benchmark.

3.1.3 The Cray MTA-2.

Machine type: Distributed-memory multi-processor.

Models: MTA-2 x , $x = 16, 32, \dots, 256$.

Operating system: Unix BSD4.4 + proprietary micro kernel.

Compilers: Fortran 77/90, ANSI C, C++.

Vendors information Web page: www.cray.com/products/systems/mta/

Year of introduction: 2001.

System parameters:

Model	MTA- x
Clock cycle	—
Theor. peak performance	—
Per Proc. (64-bits)	≥ 750 Mflop/s
Maximal	≥ 192 Gflop/s
Main Memory	≤ 1 TB
No. of processors	16–256

Remarks:

The exact peak speed of the MTA systems cannot be given as the clock cycle is a well kept secret on the data sheets of the Cray MTA. Only some lower bounds of the peak performance are quoted. Systems with a 200 MHz and 227 MHz clocks exist, but the latter cannot be the target clock frequency because it is lower by a factor of about 3 with respect to the speed consistent with that in the Cray documentation.

Let us look at the architectural features: Although the memory in the MTA is physically distributed, the system is emphatically presented as a shared-memory machine (with non-uniform access time). The latency incurred in memory references is hidden by *multi-threading*, i.e., usually many concurrent program threads (instruction streams) may be active at any time. Therefore, when for instance a load instruction cannot be satisfied because of memory latency the thread requesting this operation is stalled and another thread of which an operation can be done is switched into execution. This switching between program threads only takes 1 cycle. As there may be up to 128 instruction streams per processor and 8 memory references can be issued without waiting for preceding ones, a latency of 1024 cycles can be tolerated. References that are stalled are retried from a retry pool. A construction that worked out similarly was to be found in the late Stern Computing Systems SSP machines (see in section 4).

The connection network connects a 3-D cube of p processors with sides of $p^{1/3}$ of which alternately the x - or y axes are connected. Therefore, all nodes connect to four out of six neighbours. In a p processor system the worst case latency is $4.5p^{1/3}$ cycles; the average latency is $2.25p^{1/3}$ cycles. Furthermore, there is an I/O port at every node. Each network port is capable of sending and receiving a 64-bit word per cycle which amounts to a bandwidth of 5.33 GB/s per port. In case of detected failures, ports in the network can be bypassed without interrupting operations of the system.

Although the MTA should be able to run “dusty-deck” Fortran programs because parallelism is automatically exploited as soon as an opportunity is detected for multi-threading, it may be (and often is) worthwhile to explicitly control the parallelism in the program and to take advantage of known data locality occurrences. MTA provides handles for this in the form of compiler directives, library routines, including synchronisation, barrier, and reduction operations on defined groups of threads. Controlled and uncontrolled parallelism approaches may be freely mixed. Furthermore, each variable has a full/empty bit associated with it which can be used to control parallelism and synchronisation with almost zero overhead. An interesting use of this property is made in DNA sequence alignment (see [4]).

An MTA-2 system with 28 processors (instead of the normal 32) was installed at the Naval Research Lab, USA, in 2002.

Measured Performances:

The company has also delivered a 16-processor system to the San Diego Supercomputing Center. This system runs at a clock cycle of 4.4 ns instead of the planned 3 ns. Consequently, the peak performance of a processor is 450 Mflop/s. Using the EuroBen Benchmark ([37]) a performance of 388 Mflop/s out of 450 Mflop/s was found for an order 800 matrix-vector multiplication, an efficiency of 86%. On 4 processors a speed of 1 Gflop/s out of 1.8 Gflop/s was found, an efficiency of 56% on the same problem. For 1-D FFTs

up to 1 million elements a speed of 106 Mflop/s was found on 1 processor and the about the same speed on 4 processors due to an insufficient availability of parallel threads.

Also limited performance data from the NAS parallel benchmark, [29], are given in [5] in which efficiencies of over around 90% were attained on 4 processors with speeds of 700–750 Mflop/s for the CG, FT, and MG kernels.

More recent performance experiments are described in [2] that especially address the scalability for application codes.

3.1.4 The Cray SX-6.

The Cray SX-6 is in fact the NEC SX-6 as marketed by Cray in the USA. It is the predecessor of the NEC SX-8 and very similar to that system but with processors that have peak performance of 9 instead of 16 Gflop/s. Also the maximal multi-frame configuration only supports 1024 instead of 4096 processors. For the description of the NEC SX-8, see section 3.1.16.

3.1.5 The Cray Inc. X1

Machine type: Shared-memory multi-vectorprocessor.

Models: Cray X1 (cluster).

Operating system: UNICOS/mp (Cray Unix variant).

bf Connection structure: Modified 2-D torus (see remarks).

Compilers: Fortran 95, C, C++, Co-Array Fortran, UPC.

Vendors information Web page: www.cray.com/products/x1/

Year of introduction: 2000.

System parameters:

Model	Cray X1
Clock cycle	800 MHz
Theor. peak performance	
Per Proc.	3.2/12.8 Gflop/s
Maximal	819 Gflop/s
Memory	256–512 GB
No. of processors	16–64 (MSP, see below)
Memory bandwidth	34.1 GB/s
Cache–CPU	76.8 GB/s
Communication bandwidth	
Aggregate	400 GB/s/64 MSPs

Remarks:

The hardware structure of the Cray X1 is almost identical to that of the former Cray SV1ex (see 4.1). The clock frequency was raised from 500 to 800 MHz and the amount of memory and the bandwidth to the processors are increased. Each processor board contains 4 CPUs that can deliver a peak rate of 4 floating-point operations per cycle, amounting to a theoretical peak performance of 3.2 Gflop/s per CPU. However, 4 CPUs can be coupled *across* CPU boards in a configuration to form a so-called Multi Streaming Processor (MSP) resulting in a processing unit that has effectively a theoretical peak performance of 12.8 Gflop/s. The reconfiguration into MSPs and/or single CPU combinations can be done dynamically as the workload dictates. The vector start-up time for the single CPUs is smaller than for MSPs, so for small vectors single CPUs might be preferable while for programs containing long vectors the MSPs should be of advantage. MSP mode is regarded as the standard mode of operation. This is also visible in the processor count given in the data sheets of Cray: the maximum number within one cabinet is 64 MSP processors. This is equivalent to 256 SSP processors. In the present Cray optimisation documentation it is said that the Cray Programming Environment is as yet not optimised for SSP processing which is not to say that suitable programs would not run efficiently in SSP mode.

The relative bandwidth both from memory to the CPU boards and from the cache to the CPUs has improved

much in comparison to the predecessor SV1ex: from memory to the CPU board 5.3 8-byte operands can be transferred. From the cache the peak bandwidth to the CPUs is 12 8-byte operands, enough to sustain dyadic operations. The cache structure is rather complex: each of the 4 SSPs on a board have their own 16 KB 2-way set-associative L1 data and instruction cache. The L1 data cache only stores scalar data. The L2 cache is 2MB in size and is shared by the SSP processors on the CPU board.

New features that are less visible to the user are adherence to the IEEE 754 floating-point standard arithmetic and a new vector instruction set that can make better use of the new features like caches and addressability and synchronisation of remote nodes. This is because every cabinet can be regarded as a node in a cluster of which a maximum of 64 can be configured in what is called a modified 2-D torus topology. Cray itself regards a board with 4 MSPs as a “node”. Each node has two connections to the outside world. Odd and even nodes are connected in pairs and the other connection from the board is connected via a switch to the other boards. Thus requiring at most two hops to reach any other MSP in the cabinet. The aggregate bandwidth in such a fully populated cabinet is 400 GB/s. Multi-cabinet configurations are extending the 2-D torus into a 3-D torus structure much like the late Cray T3E 4.1 Latency and bandwidth data for point-to-point communication are not provided but various measurements in an MPI environment have been done, see the Measured Performances below.

On a 4-SSP CPU board `OpenMP` can be employed. When accessing other CPU boards one can use Cray’s `shmem` library for one-sided communication, MPI, Co-Array Fortran, etc.

Measured Performances:

In [44] a speed of 5895 Gflop/s is reported for solving a 494,592-order linear system on a 504-(MSP)processor machine at Oak Ridge National Laboratory. This amounts to an efficiency of 91%.

A more extensive evaluation of the ORNL system is reported in [11]. Here a point-to-point bandwidth of 13.9 GB/s was found within a 4-MSP node and 11.9 GB/s between nodes. MPI latencies for small messages were 8.2 and 8.6 μ s, respectively. For `shmem` and Co-Array Fortran the latencies were only 3.8 and 3.0 μ s.

3.1.6 The Cray Inc. XD1

Machine type: Distributed-memory multi-processor.

Models: XD1.

Operating system: Linux (kernel 2.4.21 with Cray HPC enhancements).

bf Connection structure: Variable (see remarks).

Compilers: Fortran 95, C, C++.

Vendors information Web page: www.cray.com/products/xd1/

Year of introduction: 2004.

System parameters:

Model	Cray XD1
Clock cycle	2.2 GHz
Theor. peak performance	
Per Chassis	52.8+ Gflop/s (see remarks)
Per Rack	663+ Gflop/s (see remarks)
Memory	
Per Chassis	96 GB
Per Rack	1.2 TB
No. of processors	
Per Chassis	12
Per Rack	144
Communication bandwidth	
Point-to-point	≤ 2.9 GB/s
Aggregate, per chassis	96 GB/s

Remarks:

The Cray XD1 is a product that was originally developed by Octigabay until this company was taken over by Cray. A distinctive factor in the Octigabay systems was the possibility to add FPGAs (see Glossary)

to the compute boards of the systems to accelerate algorithms that are of special interest to the user, like massive FFTs or DNA sequence alignments. Hence the plus symbols in the entries for the Theoretical Peak Performance in the System Parameters list above. Cray turned the system into a product by adding its special communication networking capability to connect the compute boards and the nodes, called “chassis” by Cray by means of its proprietary Rapid Array Network.

The general structure of an XD1 is as follows: one chassis houses up to 6 compute cards. Each compute card has 2 AMD Opterons at 2.2 GHz and one or two RapidArray Processors (RAPs) that handle the communication. The two Opterons on a card are connected via AMD’s HyperTransport with a bandwidth of 3.2 GB/s forming a 2-way SMP. Because of the high bandwidth of the HyperTransport bus the memory access does not suffer from using two processors on a board, unlike in most 2 processor/node clusters. Optionally an application acceleration processor (FPGA) can be put onto a compute board. With 2 RAPs/board a bandwidth of 8 GB/s (4 GB/s bi-directional) between boards is available via a RapidArray switch. This switch has 48 links of which half is used to connect to the RAPs on the compute boards within the chassis and the others can be used to connect to other chassis. Twelve chassis fit into a standard rack and because of the number of free links per RapidArray switch the chassis in two racks may be connected directly. Of course larger configurations can be put together by connecting the links in a more sparsely connected network, like a 3-D torus or a fat tree.

The RAPs offload the Opteron processors from communication tasks and have hardware support for MPI, Cray-style `shmem`, and Global Arrays (a virtual shared memory system). The communication characteristics for MPI via the RapidArray network as stated by Cray are impressive: 2.9 GB/s for long messages and a 1.6 μ s latency for small messages.

An extra feature of the Cray-enhanced Linux OS is the synchronisation of tasks in the system. The random scheduling of tasks within the system (by the OS or otherwise) can result in large latencies (see [31]) that may be detrimental to the MPI performance. By task synchronisation this problem can be evaded.

Measured Performances:

The Cray XD1 is still new. Evaluation of 6 chassis test is presently being performed at ORNL. See the website of Tom Dunigan [12].

3.1.7 The Cray Inc. XT3

Machine type: Distributed-memory multi-processor.

Models: XT3.

Operating system: UNICOS/lc, Cray’s microkernel Unix.

bf Connection structure: 3-D Torus.

Compilers: Fortran 95, C, C++.

Vendors information Web page: www.cray.com/products/xt3/

Year of introduction: 2004.

System parameters:

Model	Cray XT3
Clock cycle	2.4 GHz
Theor. peak performance	
Per processor	4.8 Gflop/s
Per Cabinet	460.8 Gflop/s
Max. Configuration	147 Tflop/s
Memory	
Per Cabinet	\leq 768 GB
Max. Configuration	196 TB
No. of processors	
Per Cabinet	96
Max. Configuration	30,508
Communication bandwidth	
Point-to-point	\leq 7.6 GB/s
Bisectional/cabinet	333 GB/s

Remarks:

The Cray XT3 is the commercial spinoff of the 10,000+ processor Red Storm machine, built by Cray for Sandia Laboratories. The structure is similar, be it that there are no provisions are made to have a “classified” and an “unclassified” part in the machine. The basic processor in a node, called PE (Processing Element) in Cray jargon, is the AMD Opteron 100, at 2.4 GHz. Cray has chosen for this uniprocessor version of the chip because of the lower memory latency (about 60 ns) in contrast to the SMP-enabled versions that have a memory latency that can be up to 2 times higher. Per PE up to 8 GB of memory can be configured, connected by a 6.4 HyperTransport to the processor. For connection to the outside world a PE harbours 2 PCI-X busses, a dual-ported FiberChannel Host Bus Adaptor for connecting to disk, and a 10 GB Ethernet card.

The Opteron was also chosen because of the high bandwidth the relatively ease of connecting the processor of to the network processor, Cray’s SeaStar chip. For the physical connection another HyperTransport channel at 6.4 GB/s is used. The SeaStar has 6 ports with a bandwidth of 7.6 GB/s each (3.8 GB/s, incoming and outgoing). Because of its 6 ports the natural interconnection mode is therefore a 3-D torus.

Like for the earlier Cray T3E (see 4), Cray has chosen to use a microkernel approach for the compute PEs. These are dedicated to computation and communication and are not disturbed by other OS tasks that can seriously influence the scalability (see [31]). For tasks like communicating with users, networking, and I/O special PEs are added that have versions of the OS that can handle these tasks.

The XT3 is obviously designed for a distributed memory parallel model, supporting Cray’s MPI 2.0 and its one-way communication `shmem` library that date back to the Cray T3D/T3E systems but is still popular because of its simplicity and efficiency. The system comes in cabinets of 96 PEs, including service PEs. For larger configurations the ratio of service PEs to compute PEs (generally) can be lowered. So, a hypothetical maximal configuration of 30,508 PEs would need only 106 service PEs.

Measured Performances:

The Cray XT3 is quite new and as yet no independent performance results are available.

3.1.8 The Fujitsu/Siemens PRIMEPOWER.

Machine type: RISC-based shared-memory multi-processor.

Models: PRIMEPOWER 900, 1500, 2500.

Operating system: Solaris (Sun’s Unix variant).

Connection structure: Crossbar.

Compilers: Fortran 90, OpenMP, C, C++.

Vendors information Web page: www.fujitsu.com/services/computing/server/unix/index.html

www.fujitsu.com/en/index.html

Year of introduction: 2002.

System parameters:

Model	PRIMEPOWER 1500	PRIMEPOWER 2500
Clock cycle	1.89 GHz	1.82 GHz
Theor. peak performance		
Per Proc. (64-bits)	3.78 Gflop/s	3.64 Gflop/s
Maximal	121.0 Gflop/s	466 Gflop/s
Memory/node	≤ 4 GB	≤ 4 GB
Memory/maximal	≤ 128 GB	≤ 512 GB
No. of processors	4–32	8–128
Communication bandwidth		
Point-to-point	—	—
Aggregate	17.2 GB/s	133 GB/s

Remarks:

We only discuss here the PRIMEPOWER 1500 and 2500 as the smaller models have the same structure but less processors (maximally 16 in the 900 model). In many respects this machine is akin to the SUN Fire E25K, section 3.1.20. The processors are 64-bit Fujitsu implementations of SUN’s SPARC processors, called

SPARC64 V processors and they are completely compatible with the SUN products. Fujitsu's SPARC64 processors have extensive RAS features for all CPU components, caches, registers, execution units, etc., that make them very reliable. The processors are available in a 1.89 GHz and 1.82 GHz for the 1500 and 2500 models respectively. Also the interconnection of the processors in the PRIMEPOWER systems is like the one in the Fire E25K: a crossbar that connects all processors at the same footing, i.e., it is *not* a NUMA machine. In fact, in the PRIMEPOWER 2500 the communication is maintained by two crossbars each at a speed of 66.5 GB/s. If one of the crossbars fails it is possible by reconfiguration to operate at half the speed. Unfortunately, there is no sound technical information available beyond the data sheets that are provided via Fujitsu's web site. These data sheets omit almost all information about the bandwidth of the interconnect be it point-to-point, bi-sectional, or aggregate. Only for the PRIMEPOWER 2500 an aggregate bandwidth is quoted: a very respectable 133 GB/s stemming from the two 66.5 GB/s crossbars. Judging from the available information the 1500 system is more positioned as a commercial server than as a high-performance computer while the 2500 model is also targeted at the HPC market. Fujitsu offers the possibility to link many PRIMEPOWER 2500s together by a very fast (4×4 GB/s) optical link to build very large configurations.

Measured Performances:

In [44] a cluster of 18 fully configured PRIMEPOWER 2500s with the older 1.3 GHz processors is used to solve a linear system of order $N = 658,800$. This yields a performance of 5.4 Tflop/s with an efficiency of 45% on 2,304 processors.

3.1.9 The Hitachi SR11000.

Machine type: RISC-based distributed-memory multi-processor.

Models: SR11000 H1.

Operating system: AIX (IBM's Unix variant).

Connection structure: Multi-dimensional crossbar (see remarks).

Compilers: Fortran 77, Fortran 95, Parallel Fortran, C, C++.

Vendors information Web page: www.hitachi.co.jp/Prod/comp/hpc/SR_e/11ktop_e.html

Year of introduction: 2003.

System parameters:

Model	SR11000 H1
Clock cycle	1.7 GHz
Theor. peak performance	
Per Proc. (64-bits)	109 Gflop/s
Maximal	27.8 Tflop/s
Memory/node	≤64 GB
Memory/maximal	16.4 TB
No. of processors	4–256
Communication bandwidth	
Point-to-point	12 GB/s (bidirectional)

Remarks:

The SR11000 is the fourth generation of distributed-memory parallel systems of Hitachi. It replaces its predecessor, the SR8000 (see 4). Presently only one model is available, the SR11000 H1 in contrast to the SR8000 of which no less than 4 models existed with different processor clock speeds and bandwidths.

The basic node processor is a 1.7 GHz POWER4+ from IBM. Unlike in the former SR2201 and SR8000 systems no modification of the processor its done to make it fit for Hitachi's Pseudo Vector Processing, a technique that enabled the processing of very long vectors without the detrimental effects that normally occur when out-of-cache data access is required. Presumably Hitachi is now relying on advanced prefetching of data to bring about the same effect.

The peak performance per basic processor, or IP, can be attained with 2 simultaneous multiply/add instructions resulting in a speed of 6.8 Gflop/s on the SR11000. However, 16 basic processors are coupled to form one processing node all addressing a common part of the memory. For the user this node is the basic computing entity with a peak speed of 108.8 Gflop/s. Hitachi refers to this node configuration as

COMPAS, Co-operative Micro-Processors in single Address Space. In fact this is a kind of SMP clustering as discussed in sections 2.1 and 2.6. In contrast to the preceding SR8000 does not contain a an SP anymore, a system processor that performed system tasks, managed communication with other nodes and a range of I/O devices. These tasks are now performed by the processors in the SMP nodes themselves.

The SR11000 has a multi-dimensional crossbar with a single-directional link speed of 12 GB/s. Also here IBM technology is used: the IBM Federation Switch fabric is used, be it in a different topology than IBM does for its own p695 servers (see 3.1.14). From 4–8 nodes the cross-section of the network is 1 hop. For configurations 16–64 it is 2 hops and from 128-node systems on it is 3 hops.

Like in some other systems as the Cray XD1 (3.1.6), AlphaServer SC, (3.1.13), and the late NEC Cenju-4, one is able to directly access the memories of remote processors. Together with the fast hardware-based barrier synchronisation this should allow for writing distributed programs with very low parallelisation overhead.

Of course the usual communication libraries like PVM and MPI are provided. In case one uses MPI it is possible to access individual IPs within the nodes. Furthermore, in one node it is possible to use OpenMP on individual IPs. Mostly this is less efficient than using the automatic parallelisation as done by Hitachi's compiler but in case one offers coarser grained task parallelism via OpenMP a performance gain can be attained. Hitachi provides its own numerical libraries to solve dense and sparse linear systems, FFTs, etc. As yet it is not known whether third party numerical libraries like NAG and IMSL are available.

It is expected that sometime in 2005 a POWER5-based SR11000 will come onto the market.

Measured Performances:

The SR11000 was introduced by the end of 2003. A few systems have been sold in Japan in the mean time but as yet no performance results are published for the system.

3.1.10 The HP 9000 Superdome.

Machine type: RISC-based ccNUMA system.

Models: HP 9000 Superdome.

Operating system: HP-UX (HP's usual Unix flavour).

Connection structure: Crossbar.

Compilers: Fortran 77, Fortran 90, Parallel Fortran, HPF, C, C++.

Vendors information Web page: www.hp.com/products1/servers/scalableservers/superdome/

Year of introduction: 2000, 2004 with PA-RISC 8800.

System parameters:

Model	HP 9000 Superdome
Clock cycle	1 GHz
Theor. peak performance	
Per Proc. (64-bits)	8 Gflop/s
Maximal	512 Glop/s
Memory/node	≤64 GB
Memory/maximal	1 TB
No. of processors	≤64
Communication bandwidth	
Aggregate (global)	64 GB/s
Aggregate (cell-backplane)	8 GB/s
Aggregate (within cell, see below)	16 GB/s

Remarks:

The Superdome replaced the Exemplar V2600 system which has been withdrawn by HP (see 4). The connection structure of the Superdome has significantly improved over that of the former V2600. The Superdome has a 2-level crossbar: one level within a 4-processor cell and another level by connecting the cells the crossbar backplane. Every cell connects to the backplane at a speed of 8 GB/s and the global aggregate bandwidth for a fully configured system is therefore 64 GB/s.

As said, the basic building block of the Superdome is the 4-processor cell. All data traffic within a cell

is controlled by the Cell Controller, a 10-port ASIC. It connects to the four local memory subsystems at 16 GB/s, to the backplane crossbar at 8 GB/s, and to two ports that each serve two processors at 6.4 GB/s/port. As each processor houses two CPU cores the available bandwidth per CPU core is 1.6 GB/s. Like the SGI Altix systems (see section 3.1.18), the cache coherency in the Superdome is secured by using directory memory. The NUMA factor for a full 64 processor systems is by HP's account very modest: only 1.8.

The PA-RISC 8800 processors run at a clock frequency of 1 GHz. As each processor contains two processor cores (see section 2.8.3) which in turn contain 2 floating-point units that are able to execute a combined floating multiply-add instruction, in favourable circumstances 8 flops/cycle can be achieved and a Theoretical Peak Performance of 8 Gflop/s per processor can be attained. This amounts to a peak speed of 512 Gflop/s for a full configuration.

Because a shared-memory parallel model is supported over the entire system, OpenMP can be employed on the total of 64 processors (128 CPU cores).

The Superdome can be partitioned in different complexes that run with different processors, e.g., the Itanium 2. In that case the same backplane can be used but the cells are of a different type. In theory one therefore can have a mixed HP 9000 Superdome and an Integrity Superdome (see below, section 3.1.11).

Measured Performances:

From the new model with the dual-core PA-RISC 8800 processors no performance results (in the HPC realm) are known to the author, the system is on the market from April 2004. In [44] a speed of 756 Gflop/s is reported for solving a full linear system of unspecified size. This result is achieved on an older 8-way coupled system with a total of 512 PA-RISC 8700+ processors at 875 MHz. As the theoretical peak performance of such a cluster is 1792 Gflop/s the efficiency is 42%.

3.1.11 The HP Integrity Superdome.

Machine type: RISC-based ccNUMA system.

Models: HP Integrity Superdome.

Connection structure: Crossbar.

Operating system: HP-UX (HP's usual Unix flavour), Linux.

Compilers: Fortran 77, Fortran 90, HPF, C, C++.

Vendors information Web page: www.hp.com/products1/servers/integrity/superdome_high_end/

Year of introduction: 2004.

System parameters:

Model	Integrity Superdome
Clock cycle	1.5 GHz
Theor. peak performance	
Per Proc. (64-bits)	6 Gflop/s
Maximal	768 Glop/s
Memory/node	≤128 GB
Memory/maximal	≤1 TB
No. of processors	≤128
Communication bandwidth	
Aggregate (global)	64 GB/s
Aggregate (cell-backplane)	8 GB/s
Aggregate (within cell, see below)	16 GB/s

Remarks:

The Integrity Superdome is HP's investment in the future for high-end servers. Within a timespan of a few years it should replace the PA-RISC-based HP 9000 Superdome. HP has anticipated on this by giving it exactly the same macro structure: cells are connected to a backplane crossbar that enables the communication between cells. For the backplane it is immaterial whether a cell contains PA-RISC or Itanium processors. For a discussion of the macro architecture see the Remarks part of section 3.1.10. Also within a cell the structure of the system is very similar: an identical Cell Controller ASIC as in the HP 9000 Superdome is

the central hub that controls local memory access, I/O requests, and communication to the other cells via the backplane.

The difference lies in the processor, the 1.5 GHz Itanium 2. HP managed to make a special processor socket that allows 2 Itaniums to be placed in one socket. Therefore 8 Itaniums in total can be housed in the four sockets on the cell board. As each Itanium contains 2 floating-point units that are able to execute a combined floating multiply-add instruction, in favourable circumstances 4 flops/cycle can be achieved and a Theoretical Peak Performance of 6 Gflop/s per processor can be attained. This amounts to a peak speed of 768 Gflop/s for a full configuration.

Because the similarity between the both Superdome systems is so large, the Integrity Superdome has the same ccNUMA characteristics as its relative. It therefore supports OpenMP over its maximum of 128 processors. As the Integrity Superdome is based on the Itanium 2 for which much Linux development is done in the past few years, the system can also be run with the Linux OS. In fact, because the machine can be partitioned, it is possible to run both Linux and HP-UX in the different complexes of the same machine. The remark as for the HP 9000 Superdome applies: in different complexes one can employ cells with different types of processors, making the system a hybrid Integrity and HP 9000 Superdome.

Measured Performances:

In [44] a speed of 1049 Gflop/s is reported for solving a full linear system of unspecified size. This result is achieved on a complex of systems with a total of 220 processors. As the theoretical peak performance of such a cluster is 1320 Gflop/s the efficiency is 79%.

3.1.12 The HP/Compaq GS series.

Machine type: RISC-based SMP system.

Models: AlphaServer GS80, GS160, GS320, GS1280.

Connection structure: 2-D torus.

Operating system: Tru64 Unix (Compaq's flavour of Unix), OpenVMS.

Compilers: Fortran 77, Fortran 90, HPF, C, C++.

Vendors information Web page: h18002.www1.hp.com/alphaserver/sc_gs.html

Year of introduction: 1999.

System parameters:

Model	GS1280
Clock cycle	1.3 GHz
Theor. peak performance	
Per Proc.	2.6 Gflop/s
Maximal	166 Gflop/s
Memory	≤ 512 GB
No. of processors	≤ 64
Memory bandwidth	
Processor/memory	12 GB/s
Communication bandwidth	
Point-to-point	6.4 GB/s
Aggregate/processor	25.6 GB/s
Aggregate (global)	1638 GB/s

Remarks:

We only discuss the largest (fastest) member of the GS family here as the others do not qualify anymore with respect to the speed boundary we set (80 Gflop/s). The GS series is a family of ccNUMA servers with currently the fastest Alpha EV7 processor available at 1.3 GHz, the EV7z. The systems are build with drawers that house 4 dual-processor cards. A fully configured GS1280 harbours 8 drawers.

Since the system is based on the EV7 which has four direct communication links on chip with a bandwidth of 6.4 GB/s each (see section 2.8.2), the natural way of linking the processors is in a 2-D torus structure. Practical measurements have shown that the NUMA factor for a 16-processor system is quite low: only 1.6, see [23].

The EV7 at 1.3 GHz has a theoretical peak performance of 2.6 Gflop/s. Because of its exceptional high memory bandwidth it may be expected to be able to perform well on a large variety of applications showing in a higher overall efficiency even though its peak performance is lower than that of many other processors. The maximum amount of memory is 512 GB, 8 GB/processor, enough to qualify it for high-end technical/scientific applications for the number of processors available.

Because of its ccNUMA character, users can employ `OpenMP` for shared-memory parallelisation on the GS systems to up to 64 processors in the GS1280. Of course also MPI can be used along with the full range of Compaq compilers.

Measured Performances:

For the EV7-based GS1280 as yet no meaningful performance results are known.

3.1.13 The HP/Compaq AlphaServer SC45 series.

Machine type: RISC-based SMP-clustered DM-MIMD system.

Models: AlphaServer SC45.

Operating system: Tru64 Unix (Compaq's flavour of Unix).

Compilers: Fortran 77, Fortran 90, HPF, C, C++.

Connection structure: Fat Tree.

Vendors information Web page: h18002.www1.hp.com/alphaserver/sc/

Year of introduction: 2002.

System parameters:

Model	AlphaServer SC
Clock cycle	1.25 GHz
Theor. peak performance	
Per Proc.	2.5 Gflop/s
Maximal	10 Tflop/s
Memory	≤ 8 TB
No. of processors	≤ 4096
Memory bandwidth	
Processor/memory	1.33 GB/s
Between cluster nodes	280 MB/s

Remarks:

The AlphaServer SC is the very high end of HP's AlphaServer line (SC stands for SuperComputer). The system is typical for the present development of SMP-based clustered systems. In the SC system the basic SMP node is the Compaq ES45, a 4-CPU SMP system with the Alpha 21264a (EV68) as its processor. The clock frequency is 1.25 GHz. The SMP node has a crossbar as its internal network with an aggregate bandwidth of 5.2 GB/s (1.33 GB/s/processor). This is sufficient to deliver 1.064 byte/clock cycle to each processor in the node simultaneously.

Within a node the system is a shared memory machine that allows for shared-memory parallel processing, for instance by using `OpenMP`. When more than four processors are required, one has to use a message passing programming model like MPI, PVM, or HPF (Compaq is one of the few companies that still provides its own HPF compiler).

For communication between the SMP nodes the SC uses QsNet, a network manufactured by QSW Limited. In fact QsNet is the follow-on of the network employed in the former Meiko CS-2 systems (see section 2.9.3). The network has the structure of a fat tree, is based on PCI technology, and has a point-to-point bandwidth of 280 MB/s. Because of its fat tree structure the bandwidth in the upper level of the network is 340 MB/s sustained. The peak bandwidth is, according to the documentation, "500 MB/s per server" without further specification which looks impressive but is not very informative. QSW claims a very low latency of 5 μ s for MPI messages.

Measured Performances:

In [44] a performance of 13.88 Tflop/s on a 2-way cluster of fully configured AlphaServer SC45s (8192 processors) was reported solving a dense linear system of order 633,000 with an efficiency of 68.0

3.1.14 The IBM eServer p690.

Machine type: RISC-based distributed-memory multi-processor.

Models: IBM eServer p695.

Operating system: AIX (IBM's Unix variant), Linux.

Connection structure: Ω -switch.

Compilers: XL Fortran (Fortran 90), (HPF), XL C, C++.

Vendors information Web page:

www-1.ibm.com/servers/eserver/pseries/hardware/highend/p4_.html.

Year of introduction: 2004 (16/32-CPU POWER4+ SMP).

System parameters:

Model	eServer p690
Clock cycle	1.9 GHz
Theor. peak performance	
Per Proc. (64-bits)	7.6 Gflop/s
Per 16-proc. HPC node	121.6 Gflop/s
Per 32-proc. Turbo node	243.2 Gflop/s
Maximal	124.5 Tflop/s
Memory/node	≤ 1 TB
Memory/maximal	512 TB
No. of processors	8–16,384
Communication bandwidth	
Node-to-node (bi-directional)	2 GB/s

Remarks:

The eServer p690 is the successor of the RS/6000 SP. It retains much of the macro structure of this system: multi-CPU nodes are connected within a frame either by a dedicated switch or by other means, like switched Ethernet. The structure of the nodes, however, has changed considerably, see 2.8.4. Up to four Multichip Modules (MCMs) are housed in a node totaling 16 or 32 CPUs in a node depending on whether the dual or single core version of the chip is used. For high-performance Computing IBM recommends to employ the 16 CPU, single core, nodes because a higher effective bandwidth from the L2 cache can be expected in this case. For less data intensive work that primarily uses the L1 cache the difference would be small while there is a large cost advantage using the 32-CPU so-called Turbo nodes. Also POWER4+ based systems with processors at clock frequencies 1.9 GHz of 1.3, 1.5, and 1.7 GHz are offered. Apart from the processor speed they are identical to the fastest variant discussed here. The p690 is accessed through a front-end control workstation that also monitors system failures. Failing nodes can be taken off line and exchanged without interrupting service.

The so-called Federation switch is the fourth generation of the high-performance interconnects made for the p690 series. The Federation switch is, like its predecessors, an Ω -switch as described in section 2.4. It has a bi-directional link speed of 2 GB/s and an MPI latency of 5–7 μ s. Although we mention only the highest speed option for the for inter-node communication here, there is a wide range of other options that could be chosen instead, e.g., Gbit Ethernet is also possible.

Applications can be run using PVM or MPI. IBM used to support High Performance Fortran, both a proprietary version and a compiler from the Portland Group. It is not clear whether this is still the case. IBM uses its own PVM version from which the data format converter XDR has been stripped. This results in a lower overhead at the cost of generality. Also the MPI implementation, MPI-F, is optimised for the p690-based systems. As the nodes are in effect shared-memory SMP systems, within the nodes OpenMP can be employed for shared-memory parallelism and it can be freely mixed with MPI if needed. In addition to its own AIX OS IBM also supports some Linux distributions: both the professional versions of RedHat and SuSe Linux are available for the p690 series.

The standard commercial models that are marketed contain up to 128 nodes. However, on special request

systems with up to 512 nodes can be built. This largest configuration is used in the table above (although never a system of a size exceeding 128 nodes has been sold yet). A POWER5-based system p695 system might come onto soon the market but no definite plans in this direction are known.

Measured Performances:

In [44] a performance of 6188 Gflop/s for a 1600 processor system with the slightly slower 1.7 GHz variant of the processor is reported for solving a dense linear system of order $N = 355,000$ yielding an efficiency of 57%. A system with 8 Turbo nodes was reported to obtain a speed of 737 Gflop/s out of 1331 Gflop/s on a linear system of size 285,000, an efficiency of 55%. As this type of application primarily operates from the L1 cache, the more or less similar efficiencies are as expected.

3.1.15 The IBM BlueGene/L.

Machine type: RISC-based distributed-memory multi-processor.

Models: IBM BlueGene/L.

Operating system: Linux.

Connection structure: 3-D Torus, Tree network.

Compilers: XL Fortran 90, XL C, C++.

Vendors information Web page:

www-1.ibm.com/servers/deepcomputing/.

Year of introduction: 2004.

System parameters:

Model	BlueGene/L
Clock cycle	700 MHz
Theor. peak performance	
Per Proc. (64-bits)	2.8 Gflop/s
Maximal	367/183.5 Tflop/s
Memory/card	512 MB
Memory/maximal	≤ 16 TB
No. of processors	$\leq 2 \times 65,536$
Communication bandwidth	
Point-to-point (3-D torus)	175 MB/s
Point-to-point (Tree network)	350 MB/s

Remarks:

The BlueGene/L is the first in a new generation of systems made by IBM for very massively parallel computing. The individual speed of the processor has therefore been traded in favour of very dense packaging and a low power consumption per processor. The basic processor in the system is a modified PowerPC 400 at 700 MHz. Two of these processors reside on a chip together with 4 MB of shared L3 cache and a 2 KB L2 cache for each of the processors. The processors have two load ports and one store port from/to the L2 caches at 8 bytes/cycle. This is half of the bandwidth required by the two floating-point units (FPUs) and as such quite high. The CPUs have 32 KB of instruction cache and of data cache on board. In favourable circumstances a CPU can deliver a peak speed of 2.8 Gflop/s because the two FPUs can perform fused multiply-add operations. Note that the L2 cache is smaller than the L1 cache which is quite unusual but which allows it to be fast.

The packaging in the system is as follows: two chips fit on a compute card with 512 MB of memory. Sixteen of these compute cards are placed on a node board of which in turn 32 go into one cabinet. So, one cabinet contains 1024 chips, i.e., 2048 CPUs. For a maximal configuration 64 cabinets are coupled to form one system with 65,356 chips/130,712 CPUs. In normal operation mode one of the CPUs on a chip is used for computation while the other takes care of communication tasks. In this mode the theoretical peak performance of the system is 183.5 Tflop/s. It is however possible when the communication requirements are very low to use both CPUs for computation, doubling the peak speed; hence the double entries in the System Parameters table above. The number of 360 Tflop/s is also the speed that IBM is using in its marketing material.

The BlueGene/L possesses no less than 5 networks, 2 of which are of interest for inter-processor communication: a 3-D torus network and a tree network. The torus network is used for most general communication patterns. The tree network is used for often occurring collective communication patterns like broadcasting, reduction operations, etc. The hardware bandwidth of the tree network is twice that of the torus: 350 MB/s against 175 MB/s per link.

At the time of writing this report no fully configured system exists yet. One such system should be delivered to Lawrence Livermore Lab by the end of this year. A smaller system of around 34 Tflop/s peak will be delivered at ASTRON, an astronomical research organisation in the Netherlands for the synthesis of radio-astronomical images.

Measured Performances:

Recently IBM has reported to have attained a speed of 70.7 Tflop/s on the HPC Linpack benchmark in solving a linear system of size $N = 933,887$, see [44].

3.1.16 The NEC SX-8 series.

Machine type: Shared-memory multi-vector processor.

Models: SX-8B SX-8A, SX-8xMy, $y = 2, \dots, 512$.

Operating system: Super-UX (Unix variant based on BSD V.4.3 Unix).

Compilers: Fortran 90, HPF, ANSI C, C++.

Vendors information Web page: www.hpce.nec.com/572.0.html

Year of introduction: 2004.

System parameters:

Model	SX-8B	SX-8A	SX-8xMy
Clock cycle	2 GHz	2 GHz	2 GHz
Theor. peak performance Per Proc. (64-bits)	16 Gflop/s	16 Gflop/s	16 Gflop/s
Maximal Single frame:	64 Gflop/s	128 Gflop/s	—
Multi frame:	—	—	90.1 Tflop/s
Main Memory, DDR2-SDRAM	32–64 GB	32–128 GB	≤16 TB
Main Memory, FCRAM	16–32 GB	32–64 GB	≤8 TB
No. of processors	1–4	4–8	8–4096

Remarks:

The SX-8 series is offered in numerous models but most of these are just frames that house a smaller amount of the same processors. We only discuss the essentially different models here. All models are based on the same processor, an 8-way replicated vector processor where each set of vector pipes contains a logical, mask, add/shift, multiply, and division pipe (see section 2.2 for an explanation of these components). As multiplication and addition can be chained (but not division) the peak performance of a pipe set at 2 GHz is 4 Gflop/s. Because of the 4-way replication a single CPU can deliver a peak performance of 16 Gflop/s. The official NEC documentation quotes higher peak performances because the peak performance of the scalar processor (rated at 4 Gflop/s, see below) is added to the peak performance of the vector processor to which it belongs. We do not follow this practice as a full utilisation of the scalar processor along with the vector processor in reality will be next to non-existent. The scalar processor that is 2-way super scalar and at 2 GHz has a theoretical peak of 4 Gflop/s. The peak bandwidth per CPU is 64 B/cycle. This is sufficient to ship 8 8-byte operands back or forth and just enough to feed one operand to each of the replicated pipe sets. Unlike from what one would expect from the naming the SX-8B is the simpler configuration of the two single-frame systems: it can be had with 1–4 processors but is in virtually all other respects equal to the larger SX-8A that can house 4–8 processors. There is one difference connected to the maximal amount of memory per frame: NEC now offers the interesting choice between the usual DDR2-SDRAM or FCRAM (Fast Cycle Memory. The latter type of memory can a factor of 2–3 faster than the former type of memory. However, because of the more complex structure of the memory, the density is about two times lower. Hence that in the systemparameters table, the entries for FCRAM are about two times lower than for

SDRAM. The lower bound for SDRAM in the SX-8A and SX-8B systems are the same: 32 GB. For the very memory-hungry applications that are usually run on vector-type systems, the availability of FCRAM can be beneficial for quite some of these applications.

In a single frame of the SX-8A models fit up to 8 CPUs. Internally the CPUs in the frame are connected by a 1-stage crossbar with the same bandwidth as that of a single CPU system: 64 GB/s/port. The fully configured frame can therefore attain a peak speed of 128 Gflop/s.

In addition, there are multi-frame models (SX-8xMy) where $x = 8, \dots, 4096$ is the total number of CPUs and $y = 2, \dots, 512$ is the number of frames coupling the single-frame systems into a larger system. There are two ways to couple the SX-8 frames in a multi-frame configuration: NEC provides a full crossbar, the so-called IXS crossbar to connect the various frames together at a speed of 8 GB/s for point-to-point unidirectional out-of-frame communication (1024 GB/s bisectional bandwidth for a maximum configuration). Also a HiPPI interface is available for inter-frame communication at lower cost and speed. When choosing for the IXS crossbar solution, the total multi-frame system is globally addressable, turning the system into a NUMA system. However, for performance reasons it is advised to use the system in distributed memory mode with MPI.

For distributed computing there is an HPF compiler and for message passing optimised MPI (MPI/SX) is available. In addition for shared memory parallelism, OpenMP is available.

Measured Performances:

The NEC SX-8 was announced in November 2004. Some very early benchmark results with the HPC Challenge benchmarks can be found at [20] for a 6-processor SX-8A located at the Institute of Laser Engineering in Osaka, Japan.

3.1.17 The NEC TX7 series.

Machine type: Shared-memory SMP system.

Models: TX7 i9010, i9510.

Operating system: Linux, HP-UX (HP's Unix variant).

Connection structure: Crossbar.

Compilers: Fortran 95, HPF, ANSI C, C++.

Vendors information Web page: www.hpce.nec.com/465.0.html

Year of introduction: 2002.

System parameters:

Model	i9010	i9510
Clock cycle	1.5 GHz	1.5 GHz
Theor. peak performance		
Per Proc. (64-bits)	6 Gflop/s	6 Gflop/s
Maximal	96 Gflop/s	192 Gflop/s
Main Memory (per frame)	≤64 GB	≤128 GB
No. of processors	16	32

Remarks:

The TX7 series is offered in 4 models of which we only discuss the 2 largest ones. The TX7 is another of the Itanium 2-based servers (see also 3.1.1, 3.1.18, and 3.1.21) which recently appeared on the market. The largest configuration presently offered is the TX7/i9510 with 32 1.5 GHz Itanium 2 processors. NEC had already some experience with Itanium servers offering 16-processor Itanium 1 servers under the name Asuza. So, the TX7 systems can be seen as a second generation.

The processors are connected by a flat crossbar. NEC still sells its TX7s with the choice of processors that Intel has: 1.3, 1.4, and 1.5 GHz processors with L3 caches of 3–6 MB depending on the clock frequency (see 2.8.7 for a full description).

Unlike the other vendors that employ the Itanium 2 processors, NEC offers its own compilers including an HPF compiler which is probably available for compatibility with the software for the NEC SX-8 because it is hardly useful on a shared-memory system like the TX7. The software also includes MPI and OpenMP. Apart from Linux also HP-UX is offered as an Operating System which may be useful for migration of

HP-developed applications to a TX7.

Measured Performances:

In the spring of 2004 rather extensive benchmark experiments with the EuroBen Benchmark were performed on a 16-processor TX7 i9010 with the 1.5 GHz variant of the processor. Using the EuroBen benchmark, the MPI version of a dense matrix-vector multiply was found to be 14.5 Gflop/s on 16 processors while both for solving a dense linear system of size $N = 2,000$ and a 1-D FFT of size $N = 16,384$ speeds of 3.8–4.1 Gflop/s are observed (see [42]).

3.1.18 The SGI Altix 3000 series.

Machine type: RISC-based ccNUMA system.

Models: Altix 3300, Altix 3700.

Operating system: Linux.

Connection structure: Crossbar, hypercube (see remarks).

Compilers: Fortran 95, C, C++.

Vendors information Web page: www.sgi.com/servers/altix/.

Year of introduction: 2002.

System parameters:

Model	Altix 3700
Clock cycle	1.5 GHz
Theor. peak performance	
Per Proc. (64-bits)	6 Gflop/s
Maximal	1.5 Tflop/s
Memory	≤512 GB
No. of processors	4–256
Communication bandwidth	
Point-to-point	1.6 GB/s
Aggregate peak/64 proc. frame	44.8 GB/s

Remarks:

The structure of the Altix 3700 is very similar to that of the SGI Origin systems (see 3.1.19). The smaller variant of the system, the Altix 3300 is not discussed here. Like the Origin systems the Altix has so-called C-bricks that contains board four Itanium 2 processors, 2 memory modules, two I/O ports, and two ASICs called SHUBs. Each SHUB connects two a memory module, an I/O port, and a shared path to two processors. In addition the 2 SHUBs are connected to each other by 6.4 GB/s link. The bandwidth of the memory modules and the I/O ports to the SHUBs are 10.2 and 2.4 GB/s, respectively. For the connection to the other bricks the same routers and network as in the Origin 3000 systems are used: the so-called Numalink3 network with a bi-section bandwidth of 25.6 GB/s. Like the Origin, the Altix is a ccNUMA system which means that the address space is shared between all processors (although it is physically distributed and therefore not uniformly accessible). Note that the bandwidth within the nodes is higher than for the off-board connections. On the boards the new Numalink4 technology is employed.

SGI does not provide its own suite of compilers. Rather it distributes the Intel compilers for the Itanium processors. Also the operating system is Linux and not IRIX, SGI's own Unix flavour. SGI is developing its cluster file system CXFS to run on Linux and will be available shortly.

The 64-processor frames can again be coupled with Numalink3 connections, making them effectively a cluster of Altix systems. Up to 4 frames can be presented in a single-system image making it into a 256-processor system with a peak performance of 1.5 Tflop/s. So OpenMP programs with up to 256 processes can be run. On larger configurations, because Numalink allows remote addressing, one can apart from MPI also employ the Cray-style `shmem` library for one-sided communication. It is expected that SGI will extend the number of processors within a single system image in the very near future.

Measured performances:

In the TOP 500 list, [44], a complex of 10160 Altix processors connected by Infiniband attained a speed of

51.78 Tflop/s solving a 1,290,240-order linear system. The efficiency for this complex is 85%.

3.1.19 The SGI Origin 3900.

Machine type: RISC-based ccNUMA system.

Models: Origin 3200, Origin 3400, Origin 3800.

Operating system: IRIX (SGI's Unix variant).

Connection structure: Crossbar, hypercube (see remarks).

Compilers: Fortran 77, Fortran 90, HPF, C, C++, Pascal.

Vendors information Web page: www.sgi.com/products/servers/origin/3000/overview.html.

Year of introduction: 2003 (With new R16000-based nodes).

System parameters:

Model	Origin3900
Clock cycle	800 MHz
Theor. peak performance	
Per Proc. (64-bits)	1.6 Gflop/s
Maximal	819 Gflop/s
Memory	≤1 TB
No. of processors	16–512
Communication bandwidth	
Point-to-point	1.6 GB/s
Aggregate peak	717 GB/s

Remarks:

By July 2000 SGI has passed from its Origin2000 series to its new Origin3000 series comprised of the Origin3200, Origin3400, and Origin3800 models. The smaller intermediate models have been dropped for the Origin 350 midrange server. In the high-end range the Origin3900 is now the only representative. Many of the characteristics of the Origin2000 have been retained of which the most important is its ccNUMA character. The processor used is presently the MIPS R16000, a direct successor of the R14000s in the Origin2000 systems. The R16000 is identical to the R14000 processor, save for the clock frequency that has gone from 600 MHz to 800 MHz in the latest version.

SGI has further modularised the Origin3900 in comparison with its predecessor. A system contains so-called C-bricks, CPU boards with 2–4 processors and a router chip connecting the on-board memory with the processors to router boards called R-bricks for communication with the rest of the system and to I-bricks that contain disks, PCI expansion slots, etc. and that together make up the I/O sub-system of the machine. Recently an ultra-dense Cx-brick has been introduced that packs 16 R16000s in a 4U-height brick. This is possible because of the very low power consumption of the MIPS processors. Using Cx bricks therefore can greatly diminish the amount of floorspace needed.

The basic hardware bandwidth within a C-brick is 1.6 GB/s from the router chip to one pair of CPUs, 3.2 GB/s from memory to the router chip (2×1.6 GB/s full duplex). The same bandwidth is available for inter-node communication. The off-board I/O bandwidth is 2.4 GB/s (2×1.2 GB/s full duplex). The R-brick can be connected to 16 C-bricks and it has 8 ports to connect it to other R-bricks. So, 128 C-bricks or 512 processors can maximally be interconnected in this way.

The machine is a typical representative of the ccNUMA class of systems. The memory is physically distributed over the node boards but there is one system image. Because of the structure of the system, the bi-sectional bandwidth of the system remains constant from 8 processors on: 210 GB/s. This is a large improvement over the earlier Origin2000 systems where the bi-sectional bandwidth was 82 GB/s.

Parallelisation is done either automatically by the (Fortran or C) compiler or explicitly by the user, mainly through the use of directives. All synchronisation, etc., has to be done via memory. This may cause potentially a fairly large parallelisation overhead. Also a message passing model is allowed on the Origin using the optimised SGI versions of PVM, MPI, and the Cray-specific `shmem` library. Programs implemented in this way will possibly run very efficiently on the system.

A nice feature of the Origins is that it may migrate processes to nodes that should satisfy the data requests of these processes. So, the overhead involved in transferring data across the machine are minimised in this

way. The technique is reminiscent of the late Kendall Square Systems although in these systems the data were moved to the active process. SGI claims that the time for non-local memory references is on average about 2 times longer than for local memory references, an improvement of 50% over the Origin2000 series.

Measured Performances:

No performance figures for the 800 MHz-based systems are available but in [44] the performance of the solution of a linear system of unspecified order is quoted for a complex of 512-processor machines with the earlier 600 MHz R14000 processors. A speed of 405.6 Gflop/s was found, an efficiency of 66%.

3.1.20 The Sun Fire E25K.

Machine type: RISC-based shared-memory multiprocessor.

Models: Sun Fire E25K.

Operating system: Solaris (Sun's Unix flavour).

Connection structure: Crossbar (see remarks).

Compilers: Fortran 77, Fortran 90, C, C++.

Vendors information Web page: www.sun.com/servers/highend/sunfire_e25k/index.xml

Year of introduction: 2004.

System parameters:

Model	Sun Fire E25K
Clock cycle	1.2 GHz
Theor. peak performance	
Per Proc. (64-bits)	4.8 Gflop/s
Maximal	345.6 Gflop/s
Main Memory	≤ 576 GB
No. of processors	≤ 72
Communication bandwidth	
Aggregate	172.8 GB/s

Remarks:

In the E25K is the successor of the 3800-15K system. It employs the newest UltraSPARC IV processors at a clock frequency of 1.2 GHz. The structure of the systems is exactly the same as for its predecessor. In fact, the backplane is identical and one can turn one system into the other by just exchanging the processor boards. The processor/memory boards are plugged into a backplane that is an 18×18 flat crossbar. Each board contains four 1.2 GHz UltraSPARC IV processors and a maximum of 32 GB of memory. The maximum number of processors is therefore 72. Where in the 3800-15K there is the possibility to use slots of an additional crossbar to plug in extra 2-CPU boards to extend the computational power at the cost of I/O capacity, this possibility is *not* offered for the E25K system. Because of the flat crossbar memory access is uniform and the aggregate bandwidth of the crossbar is 172.8 GB/s. This is equivalent to 2.4 GB/s/processor or 2 B/cycle. So, an 8-byte operand needs 4 cycles to be shipped to the processor.

The main (almost only) difference with its predecessor is that the E25K system employs the dual-core UltraSPARC IV processors (see section 2.8.10). This formally doubles the peak performance per processor, although because of bandwidth constraints the effective performance increase will be lower. Sun maintains a somewhat confusing terminology with respect to multi-threading: it uses the term to express the fact that the two processor cores on an UltraSPARC IV chip each do their processing independently. It calls this chip multi-threading (CMT). Yet the cores are not capable of switching between process threads as what is normally understood by multi-threading.

The E25K is a typical SMP machine with provisions for shared-memory parallelism with OpenMP over the full range of processors in the system.

Measured Performances:

In [44] a speed of 891.4 Gflop/s is reported for a 7-way cluster with a total of 672 processors in solving a dense linear system of unspecified size. The efficiency for this problem is 74%.

3.1.21 The Unisys Server ES7000.

Machine type: RISC-based shared-memory multiprocessor.

Models: ES7000 Orion 440.

Operating system: Microsoft Windows Server 2003, UnitedLinux.

Compilers: Fortran 90, C, C++.

Vendors information Web page: www.unisys.com/products/es7000__linux/hardware/orion__440.htm

Year of introduction: 2004.

System parameters:

Model	Orion 440
Clock cycle	1.5 GHz
Theor. peak performance	
Per Proc. (64-bits)	6.0 Gflop/s
Maximal	192 Gflop/s
Main Memory	≤512 GB
No. of processors	≤32

Remarks:

We just mention the Orion 440 here for completeness' sake. Normally Unisys is not concerned with the HPC market but recently Unisys expressed interest in stepping in. As can be seen from the list of characteristics above, the systems are presently offered with a Windows OS, and Linux. The system has a flat crossbar between the standard Intel 4-CPU Quad Building Blocks which makes it very similar to those of Bull NovaScale and the NEC TX-7. No further information is available, including performance results.

4 Systems disappeared from the list

As already stated in the introduction the list of systems is not complete. On one hand this is caused by the sheer number of systems that are presented to the market and are often very similar to systems described above (for instance, the Volvox system not listed was very similar but not equivalent to the former C-DAC system and there are numerous other examples). On the other hand, there are many systems that are still in operation around the world, often in considerable quantities that for other reasons are excluded. The most important reasons are:

- The system is not marketed anymore. This is generally for one of two reasons:
 1. The manufacturer is out of business.
 2. The manufacturer has replaced the system by a newer model of the same type or even of a different type.
- The system has become technologically obsolete in comparison to others of the same type. Therefore, listing them is not sensible anymore.

Below we present a table of systems that fall into one of the categories mentioned above. We think this may have some sense to those who come across machines that are still around but are not the latest in their fields. It may be interesting at least to have an indication how such systems compare to the newest ones and to place them in context.

It is good to realise that although systems have disappeared from the section above they still may exist and are actually sold. However, their removal stems in such cases mainly from the fact that they are not serious candidates for high-performance computing anymore.

The table is, again, not complete and admittedly somewhat arbitrary. Furthermore, we have reduced the contents of this section (in comparison to the former version) to systems that very probably are still actually in use in this printed (PostScript) version.

A more comprehensive, “full” version is available in the web version of this document at:

www.phys.uu.nl/~steen/web04/gone.html

This full information may be of interest for a subset of readers that want to know about the historical development of supercomputing over the last 11 years.

The data are in a highly condensed form: the system name, system type, theoretical maximum performance of a fully configured system, and the reason for their disappearance is given. The arbitrariness lies partly in the decision which systems are still sufficiently of interest to include and which are not.

We include also both the year of introduction and the year of exit of the systems when they were readily accessible. These time-spans could give a hint of the dynamics that governs this very dynamical branch of the computer industry.

4.1 Disappeared machines

Machine: Avalon A12.

Year of introduction: 1996.

Year of exit: 2000.

Type: RISC-based distributed memory multi-processor, max. 1680 processors.

Theoretical Peak performance: 1.3 Tflop/s.

Reason for disappearance: Avalon is not in business anymore.

Machine: Cambridge Parallel Processing DAP Gamma.

Year of introduction: 1986.

Year of exit: 1995.

Type: Distributed-memory processor array system, max. 4096 processors.

Theoretical Peak performance: 1.6 Gflop/s (32-bit).

Reason for disappearance: replaced by newer Gamma II Plus series (4.1).

Machine: Cambridge Parallel Processing DAP Gamma II Plus.

Year of introduction: 1995.

Year of exit: 2003.

Type: Distributed-memory processor array system, max. 4096 processors.

Theoretical Peak performance: 2.4 Gflop/s (32-bit).

Reason for disappearance: system became too slow, even for its specialised tasks. (4.1).

Machine: C-DAC PARAM 9000/SS.

Year of introduction: 1995.

Year of exit: 1997.

Type: Distributed-memory RISC based system, max. 200 processors.

Theoretical Peak performance: 12.0 Gflop/s.

Reason for disappearance: replaced by newer OpenFrame series (see below).

Machine: C-DAC PARAM Openframe series.

Year of introduction: 1996.

Year of exit: 1999.

Type: Distributed-memory RISC based system, max. 1024 processors.

Theoretical Peak performance: Unspecified.

Reason for disappearance: The system is not actively marketed anymore by C-DAC.

Machine: C-DAC PARAM 10000 Openframe series.

Year of introduction: 2000.

Year of exit: 2002.

Type: Distributed-memory RISC based system, max. processors: Unspecified.

Theoretical Peak performance: Unspecified.

Reason for disappearance: The system is replaced by the newer C-DAC PARAM Padma, see section 3.1.2.

Machine: Cray T3E Classic.

Year of introduction: 1996.

Year of exit: 1997.

Type: Distributed-memory RISC based system, max. 2048 processors.

Theoretical Peak performance: 1228 Gflop/s.

Reason for disappearance: replaced Cray T3Es with faster clock. (4.1).

Machine: Cray T3E 1350.

Year of introduction: 2000.

Year of exit: 2003.

Type: Distributed-memory RISC based system, max. 2048 processors.

Theoretical Peak performance: 2938 Gflop/s.

Reason for disappearance: Cray does not market the system anymore.

Machine: Cray J90.

Year of introduction: 1994.

Year of exit: 1998.

Type: Shared-memory vector-parallel, max. 32 processors.

Theoretical Peak performance: 6.4 Gflop/s.

Reason for disappearance: replaced by newer Cray SV1 (4.1).

Machine: Cray Y-MP T90.
Year of introduction: 1995.
Year of exit: 1998.
Type: Shared-memory vector-parallel, max. 32 processors.
Theoretical Peak performance: 58 Gflop/s.
Reason for disappearance: replaced by newer Cray SV1 (see below).

Machine: Cray SV-1(ex).
Year of introduction: 2000.
Year of exit: 2004.
Type: Shared-memory vector-parallel, max. 32 processors(per frame).
Theoretical Peak performance: 64 Gflop/s.
Reason for disappearance: replaced by newer Cray X1 (3.1.5).

Machine: Digital Equipment Corp. AlphaServer 8200 & 8400.
Year of introduction: —.
Year of exit: 1998.
Type: Distributed-memory RISC based systems, max. 6 processors (AlphaServer 8200) or 14 (AlphaServer 8400).
Theoretical Peak performance: 7.3 Gflop/s, resp. 17.2 Gflop/s.
Reason for disappearance: replaced by newer AlphaServer SC. (3.1.13).

Machine: Fujitsu AP3000.
Year of introduction: 1996.
Year of exit: 2003.
Type: Distributed memory RISC based system, max. 1024 processors.
Theoretical Peak performance: 614 Gflop/s.
Reason for disappearance: Fujitsu does not market the system anymore in favour of its PrimePower series (3.1.8).

Machine: Fujitsu AP1000.
Year of introduction: 1991.
Year of exit: 1996.
Type: Distributed memory RISC based system, max. 1024 processors.
Theoretical Peak performance: 5 Gflop/s.
Reason for disappearance: replaced by the AP3000 systems (4.1).

Machine: Fujitsu VPP300/700 series.
Year of introduction: 1995/1996.
Year of exit: 1999.
Type: Distributed-memory multi-processor vectorprocessors, max. 256 processors.
Theoretical Peak performance: 614 Gflop/s.
Reason for disappearance: replaced by the VPP5000 series (4.1).

Machine: Fujitsu VPP5000 series.
Year of introduction: 1999.
Year of exit: 2002.
Type: Distributed-memory multi-processor vectorprocessors, max. 128 processors.
Theoretical Peak performance: 1.22 Tflop/s.
Reason for disappearance: Fujitsu does not market vector systems anymore, this machine line is replaced by the PRIMEPOWER series. (3.1.8).

Machine: Hitachi S-3800 series.
Year of introduction: 1993.

Year of exit: 1998.

Type: Shared-memory multi-processor vectorprocessors, max. 4 processors.

Theoretical Peak performance: 32 Gflop/s.

Reason for disappearance: Replaced by the SR8000 (4.1).

Machine: Hitachi SR2201 series.

Year of introduction: 1996.

Year of exit: 1998.

Type: Distributed-memory RISC based system, max. 1024 processors.

Theoretical Peak performance: 307 Gflop/s.

Reason for disappearance: Replaced by the newer SR8000 (4.1).

Machine: Hitachi SR8000 series.

Year of introduction: 1998.

Year of exit: 2000–2003 (different models).

Type: Distributed-memory RISC based system, max. 512 processors.

Theoretical Peak performance: 7.3 Tflop/s.

Reason for disappearance: Replaced by the newer SR11000 (3.1.9).

Machine: The HP Exemplar V2600.

Year of introduction: 1999.

Year of exit: 2000.

Type: Distributed-memory RISC based system, max. 128 processors.

Theoretical Peak performance: 291 Gflop/s.

Reason for disappearance: Replaced by the HP 9000 Superdome. (3.1.10).

Machine: IBM SP1 series.

Year of introduction: 1992.

Year of exit: 1994.

Type: Distributed-memory RISC based system, max. 64 processors.

Theoretical Peak performance: 8 Gflop/s.

Reason for disappearance: Replaced by the newer RS/6000 SP (4.1).

Machine: IBM RS/6000 SP series.

Year of introduction: 1999.

Year of exit: 2001.

Type: Distributed-memory RISC based system, max. 2048 processors.

Theoretical Peak performance: 24 Gflop/s.

Reason for disappearance: Replaced by the newer eServer p690. (3.1.14).

Machine: Meiko CS-2.

Year of introduction: 1994.

Year of exit: 1999.

Type: Distributed-memory RISC based system.

Theoretical Peak performance: 200 Mflop/s per processor.

Reason for disappearance: Quadrics Supercomputers World Ltd. does not market the system anymore. The updated network technology is now offered for other systems like Compaq's AlphaServer SC (see 3.1.13).

Machine: NEC Cenju-3.

Year of introduction: 1994.

Year of exit: 1996.

Type: Distributed-memory system, max. 256 processors.

Theoretical Peak performance: 12.8 Gflop/s.

Reason for disappearance: replaced by newer Cenju-4 series (4.1).

Machine: NEC Cenju-4.

Year of introduction: 1998.

Year of exit: 2002.

Type: Distributed-memory system, max. 1024 processors.

Theoretical Peak performance: 410 Gflop/s.

Reason for disappearance: NEC has withdrawn this machine in favour of a possible successor. Specifics are not known, however.

Machine: NEC SX-4.

Year of introduction: 1995.

Year of exit: 1996.

Type: Distributed-memory cluster of SM-MIMD vector processors, max. 256 processors.

Theoretical Peak performance: 1 Tflop/s.

Reason for disappearance: replaced by newer SX-5 series (see below).

Machine: NEC SX-5.

Year of introduction: 1998.

Year of exit: 2002.

Type: Distributed-memory cluster of SM-MIMD vector processors, max. 512 processors.

Theoretical Peak performance: 5.12 Tflop/s.

Reason for disappearance: replaced by newer SX-6 series (see below).

Machine: NEC SX-6.

Year of introduction: 2002.

Year of exit: 2005.

Type: Distributed-memory cluster of SM-MIMD vector processors, max. 1024 processors.

Theoretical Peak performance: 9.2 Tflop/s.

Reason for disappearance: replaced by newer SX-8 series (3.1.16).

Machine: Quadrics Appemille.

Year of introduction: 1999.

Year of exit: 2004.

Type: Processor array, max. 2048 processors.

Theoretical Peak performance: 1 Tflop/s.

Reason for disappearance: Not marketed anymore.

Machine: Silicon Graphics Origin2000.

Year of introduction: 1996.

Year of exit: 2000.

Type: Shared-memory multi-processor, max. 128 processors.

Theoretical Peak performance: 102.4 Gflop/s.

Reason for disappearance: replaced by the SGI Origin 3000 (3.1.19).

Machine: SUN E10000 Starfire.

Year of introduction: 1997.

Year of exit: 2001.

Type: Shared-memory multi-processor, max. 64 processors.

Theoretical Peak performance: 51.2 Gflop/s.

Reason for disappearance: replaced by the Fire 3800-15K (see below).

Machine: SUN Fire 3800-15K.

Year of introduction: 2001.

Year of exit: 2004.

Type: Shared-memory multi-processor, max. 106 processors.

Theoretical Peak performance: 254 Gflop/s.

Reason for disappearance: replaced by the Fire E25K (3.1.20).

Machine: Thinking Machine Corporation CM-5.

Year of introduction: 1991.

Year of exit: 1996.

Type: Distributed-memory RISC based system, max. 16K processors.

Theoretical Peak performance: 2 Tflop/s.

Reason for disappearance: Thinking Machine Corporation has stopped manufacturing hardware and hopes to keep alive as a software vendor.

5 Systems under development

Although we mainly want to discuss real, marketable systems and no experimental, special purpose, or even speculative machines, we want to include a section on systems that are in a far stage of development and have a fair chance of reaching the market. For inclusion in section 3 we set the rule that the system described there should be on the market within a period of 6 months from announcement. The systems described in this section will in all probability appear within one year from the publication of this report. However, there are vendors who do not want to disclose any specific data on their new machines until they are actually beginning to ship them. We recognise the wishes of such vendors (it is generally wise not to stretch the expectation of potential customers too long) and they will not disclose such information.

Below we discuss systems that may lead to commercial systems to be introduced on the market between somewhat more than half a year to a year from now. The commercial systems that result from it will sometimes deviate significantly from the original research models depending on the way the development is done (the approaches in Japan and the USA differ considerably in this respect) and the user group which is targeted.

A development that has shown to be of significance is the introduction of Intel's IA-64 Itanium processor family. Already four vendors are offering Itanium 2-based systems at the moment and it is known that HP will end the marketing of its Alpha and PA-RISC based systems in favour of the Itanium processor family. At the same time SGI will stop the further development of MIPS processor based machines. This means that in a few years only AMD, IBM, Intel, and SUN will produce RISC-like processors for HPC systems. This will make the HPC system field much less diverse and interesting. On the other hand, the shock that was caused in the USA by the advent of the Japanese Earth Simulator system may help in refueling the funding of alternative processor and computer architecture research. Indeed, some initiatives in that direction are already under way but these will not bear real new results in one or two years (except maybe with the IBM Blue Gene, see below).

5.1 Cray Inc.

In the end of 2002 the next generation vector processor, the X1, from Cray Inc. was ready to ship. It built on the technology found in the Cray SV-1s. Cray widely publicises a roadmap of future systems as far as around 2010. It remains to be seen how much can be realised, however, at least 2 of these systems are certain to reach the market in 2005: first the Cray X1e, which is nothing more than the present X1 system in which the clock cycle of the processor is raised from 800 MHz to 1.2 GHz. The other one is the Cray Strider, a commercialised version of the AMD Opteron-based (11,648 processors) Red Storm machine that is presently built by Cray for Sandia Laboratories. There is much interest for this type of system because of the cheap basic processor and the fast network based on AMD's HyperTransport and Cray's SeaStar router ASIC. Further away lies the Black Widow a follow on to the X1e, scheduled for 2006. Recently plans for a new type of system have been disclosed, code name "Rainier". In this system the inter-processor network is the central part and in this network nodes of different type can be mixed, vector type (Black Widow), scalar type (AMD or other), and FPGAs/DSPs for special functions. An upgraded form of the Red Storm network seems a very good basis for such a system but the realisation is still some time away and undoubtedly some technical challenges will be met.

5.2 Hewlett-Packard/Intel

HP and Intel will have a great influence in the next few years with their Itanium processor family. A dual core processor based on the Itanium is already on (or beyond) the drawing board and will hit the market in

a year or two. As dual core processors usually have a relatively poorer performance than their single core equivalents, the performance improvement will not be spectacular. The system architecture will be much more important. Also a diversification of the processors themselves may help to boost the performance. Because of HP/Intel's experience with VLIW processors (as the Itanium essentially is), one might expect that the research will go in the direction of processors with even longer instruction words and possibly including specialised devices for high level operations like FFTs or sparse Matrix-Vector multiplies as well. When and how such improvements would turn up in future systems is however speculative. It will certainly not happen within the next two years. As yet no radically different system architectures are known to be on HP's drawing boards. Instead it may try to penetrate more in the cluster field were it already has installed some large Itanium-based systems.

5.3 IBM

IBM has been working for some years on its BlueGene systems. Of which the first models, the BlueGene/L, will be installed within a few months (see 3.1.15). Other BlueGene follow-ups are planned called the BlueGene/P with a peak speed of 1 Pflop/s, and the BlueGene/Q with a peak speed of 3 Pflop/s, respectively. All these systems are hardly meant for the average HPC user but they may help in finding suitable architectural features for systems for the general market.

Of course the development of the POWER x processors also will make its mark: the POWER5 processor has the usual technology-related advantages over its predecessor, and now it is a subject of research how to couple 8 of them such that a virtual vector processor with a peak speed of 60–80 Gflop/s can be made. This approach is called the ViVA (Virtual Vector Architecture). It is reminiscent of Hitachi's SR811000 processors (which are also POWER5 processors) or the MSP processors in the Cray X1. This road will take some years to go also after the POWER5 processor has become available and will extend to the next generation(s) of the POWER x .

5.4 SGI

Last year it has become known that SGI will stop producing its MIPS-based systems. Therefore, the difference they would like to make with respect to other vendors that also offer Itanium-based systems would have to lie in the macro-architecture of their systems. Improvements can be realised in the speed of the network (Numalink3 to Numalink4 and beyond) and systems with a large amount of processors in a single system image (SSI). In that respect SGI has a track record with its MIPS-based Origin 3000 systems which may be extended for its future Altix x 000 systems where at present SSIs of 256 are realised.

Further in the future SGI seems to have plans that are more or less similar to Cray's Rainier project: coupling of heterogeneous processor sets by its proprietary network, in this case a successor of the NUMalink4 network architecture. Development of such systems is very cost-intensive, so it remains to be seen whether such plans will pass the stage of intentions.

5.5 Sun

Like Cray and IBM Sun has been awarded a grant from DARPA to develop so-called high-productivity systems in DARPA's HPCS program. Up till now Sun has concentrated on developing heavily multi-threaded processors, the first product being the Niagara chip and the next, the still more multi-threaded Rock processor. The first implementation of the Niagara processor is about ready, although not product-ready. The chip harbours 8 CPU cores of which each core is 4-way multi-threaded. Systems based on the Rock processor would not be available before 2008. But also for the Niagara processor no system architecture is detailed yet which makes the future for these rather experimental systems somewhat speculative. For its near future mainstream high-end systems Sun will therefore rely on Fujitsu's SPARC64 processors.

6 Glossary of terms

This section contains the explanation of some often-used terms that either are not explained in the text or, by contrast, are described extensively and for which a short description may be convenient.

6.1 Glossary

Architecture: The internal structure of a computer system or a chip that determines its operational functionality and performance.

Architectural class: Classification of computer systems according to its architecture: e.g., distributed memory MIMD computer, symmetric multi processor (SMP), etc. See this glossary and section 2.1 for the description of the various classes.

ASCI: Accelerated Strategic Computer Initiative. A massive funding project in the USA concerning research and production of high-performance systems. The main motivation is said to be the management of the USA nuclear stockpile by computational modeling instead of actual testing. ASCI has greatly influenced the development of high-performance systems in a single direction: clusters of SMP systems.

ASIC: Application Specific Integrated Circuit. A chip that is designed to fulfill a specific task in a computer system, e.g. for routing messages in a network.

Bank cycle time: The time needed by a (cache-)memory bank to recover from a data access request to that bank. Within the bank cycle time no other requests can be accepted.

Beowulf cluster: Cluster of PCs or workstations with a private network to connect them. Initially the name was used for do-it-yourself collections of PCs mostly connected by Ethernet and running Linux to have a cheap alternative for “integrated” parallel machines. Presently, the definition is wider including high-speed switched networks, fast RISC-based processors and complete vendor-preconfigured rack-mounted systems with either Linux or Windows as an operating system.

Bit-serial: The operation on data on a bit-by-bit basis rather than on byte or 4/8-byte data entities in parallel. Bit-serial operation is done in processor array machines where for signal and image processing this mode is advantageous.

Cache — data, instruction: Small, fast memory close to the CPU that can hold a part of the data or instructions to be processed. The primary or level 1 caches are virtually always located on the same chip as the CPU and are divided in a cache for instructions and one for data. A secondary or level 2 cache is mostly located off-chip and holds both data and instructions. Caches are put into the system to hide the large latency that occurs when data have to be fetched from memory. By loading data and or instructions into the caches that are likely to be needed, this latency can be significantly reduced.

Capability computing: A type of large-scale computing in which one wants to accommodate very large and time consuming computing tasks. This requires that parallel machines or clusters are managed with the highest priority for this type of computing possibly with the consequence that the computing resources in the system are not always used with the greatest efficiency.

Capacity computing: A type of large-scale computing in which one wants to use the system (cluster) with the highest possible throughput capacity using the machine resources as efficient as possible. This may have adverse effects on the performance of individual computing tasks while optimising the overall usage of the system.

- ccNUMA:** Cache Coherent Non-Uniform Memory Access. Machines that support this type of memory access have a physically distributed memory but logically it is shared. Because of the physical difference of the location of the data items, a data request may take a varying amount of time depending on the location of the data. As both the memory parts and the caches in such systems are distributed a mechanism is necessary to keep the data consistent system-wide. There are various techniques to enforce this (directory memory, snoopy bus protocol). When one of these techniques is implemented the system is said to be cache coherent.
- Clock cycle:** Fundamental time unit of a computer. Every operation executed by the computer takes at least one and possibly multiple cycles. Typically, the clock cycle is now in the order of one to a few nanoseconds.
- Clock frequency:** Reciprocal of the clock cycle: the number of cycles per second expressed in Hertz (Hz). Typical clock frequencies nowadays are 400 MHz–1 GHz.
- Clos network:** A logarithmic network in which the nodes are attached to switches that form a *spine* that ultimately connects all nodes.
- Communication latency:** Time overhead occurring when a message is sent over a communication network from one processor to another. Typically the latencies are in the order of a few μs for specially designed networks, like Infiniband or Myrinet, to about 100 μs for (Gbit) Ethernet.
- Control processor:** The processor in a processor array machine that issues the instructions to be executed by all the processors in the processor array. Alternatively, the control processor may perform tasks in which the processors in the array are not involved, e.g., I/O operations or serial operations.
- CRC:** Type of error detection/correction method based treating a data item as a large binary number. This number is divided by another fixed binary number and the remainder is regarded as a checksum from which the correctness and sometimes the (type of) error can be recovered. CRC error detection is for instances used in SCI networks.
- Crossbar (multistage):** A network in which all input ports are directly connected to all output ports without interference from messages from other ports. In a one-stage crossbar this has the effect that for instance all memory modules in a computer system are directly coupled to all CPUs. This is often the case in multi-CPU vector systems. In multistage crossbar networks the output ports of one crossbar module are coupled with the input ports of other crossbar modules. In this way one is able to build networks that grow with logarithmic complexity, thus reducing the cost of a large network.
- Distributed Memory (DM):** Architectural class of machines in which the memory of the system is distributed over the nodes in the system. Access to the data in the system has to be done via an interconnection network that connects the nodes and may be either explicit via message passing or implicit (either using HPF or automatically in a ccNUMA system).
- Dual core chip:** A chip that contains two CPUs and (possibly common) caches. Due to the progression of the integration level more devices can be fitted on a chip. HP, IBM, and Sun make dual core chips and other vendors may follow in the near future.
- EPIC:** Explicitly Parallel Instruction Computing. This term is coined by Intel for its IA-64 chips and the Instruction Set that is defined for them. EPIC can be seen as Very Large Instruction Word computing with a few enhancements. The gist of it is that no dynamic instruction scheduling is performed as is done in RISC processors but rather that instruction scheduling and speculative execution of code is determined beforehand in the compilation stage of a program. This simplifies the chip design while potentially many instructions can be executed in parallel.
- Fat tree:** A network that has the structure of a binary (quad) tree but that is modified such that near the root the available bandwidth is higher than near the leaves. This stems from the fact that often a root processor has to gather or broadcast data to stands for Host Bus Adapter. It is the part in an external network that constitutes the interface between the network itself and the PCI bus of the

compute node. HBAs usually carry a good amount of processing intelligence themselves for initiating communication, buffering, checking for correctness, etc. HBAs tend to have different names in different networks: HCA or TCA for Infiniband, LANai for Myrinet, ELAN for QsNet, etc. all other processors and without this modification contention would occur near the root.

FPGA: FPGA stands for Field Programmable Gate Array. This is an array of logic gates that can be hardware-programmed to fulfill user-specified tasks. In this way one can devise special purpose functional units that may be very efficient for this limited task. As FPGAs can be reconfigured dynamically, be it only 100–1,000 times per second, it is theoretically possible to optimise them for more complex special tasks at speeds that are higher than what can be achieved with general purpose processors.

Functional unit: Unit in a CPU that is responsible for the execution of a predefined function, e.g., the loading of data in the primary cache or executing a floating-point addition.

Grid — 2-D, 3-D: A network structure where the nodes are connected in a 2-D or 3-D grid layout. In virtually all cases the end points of the grid are again connected to the starting points thus forming a 2-D or 3-D torus.

HBA: HBA stands for Host Bus Adapter. It is the part in an external network that constitutes the interface between the network itself and the PCI bus of the compute node. HBAs usually carry a good amount of processing intelligence themselves for initiating communication, buffering, checking for correctness, etc. HBAs tend to have different names in different networks: HCA or TCA for Infiniband, LANai for Myrinet, ELAN for QsNet, etc.

HPF: high-performance Fortran. A compiler and run time system that enables to run Fortran programs on a distributed memory system as on a shared memory system. Data partition, processors layout, etc. are specified as comment directives that makes it possible to run the processor also serially. Present HPF available commercially allow only for simple partitioning schemes and all processors executing exactly the same code at the same time (on different data, so-called Single Program Multiple Data (SPMD) mode).

Hypercube: A network with logarithmic complexity which has the structure of a generalised cube: to obtain a hypercube of the next dimension one doubles the perimeter of the structure and connect their vertices with the original structure.

Instruction Set Architecture: The set of instructions that a CPU is designed to execute. The Instruction Set Architecture (ISA) represents the repertoire of instructions that the designers determined to be adequate for a certain CPU. Note that CPUs of different making may have the same ISA. For instance the AMD processors (purposely) implement the Intel IA-32 ISA on a processor with a different structure.

Memory bank: Part of (cache) memory that is addressed consecutively in the total set of memory banks, i.e., when data item $a(n)$ is stored in bank b , data item $a(n + 1)$ is stored in bank $b + 1$. (Cache) memory is divided in banks to evade the effects of the bank cycle time (see above). When data is stored or retrieved consecutively each bank has enough time to recover before the next request for that bank arrives.

Message passing: Style of parallel programming for distributed memory systems in which non-local data that is required explicitly must be transported to the processor(s) that need(s) it by appropriate send and receive messages.

MPI: A message passing library, Message Passing Interface, that implements the message passing style of programming. Presently MPI is the *de facto* standard for this kind of programming.

Multithreading: A capability of a processor core to switch to another processing thread, i.e., a set of logically connected instructions that make up a (part of) a process. This capability is used when a process thread stalls, for instance because necessary data are not yet available. Switching to another thread that has instructions that can be executed will yield a better processing utilisation.

- NUMA factor:** The difference in speed of accessing local and non-local data. For instance when it takes 3 times longer to access non-local data than local data, the NUMA factor is 3.
- OpenMP:** A shared memory parallel programming model in which shared memory systems and SMPs can be operated in parallel. The parallelisation is controlled by comment directives (in Fortran) or pragmas (in C and C++), so that the same programs also can be run unmodified on serial machines.
- PCI bus:** Bus on PC node, typically used for I/O, but also to connect nodes with a communication network. The bandwidth varies with the type from 110-480 MB/s. Newer upgraded versions PCI-X and PCI Express are (becoming) available presently.
- Pipelining:** Segmenting a functional unit such that it can accept new operands every cycle while the total execution of the instruction may take many cycles. The pipeline construction works like a conveyor belt accepting units until the pipeline is filled and then producing results every cycle.
- Processor array:** System in which an array (mostly a 2-D grid) of simple processors execute its program instructions in lock-step under the control of a Control Processor.
- PVM:** Another message passing library that has been widely used. It was originally developed to run on collections of workstations and it can dynamically spawn or delete processes running a task. PVM now largely has been replaced by MPI.
- Register file:** The set of registers in a CPU that are independent targets for the code to be executed possibly complemented with registers that hold constants like 0/1, registers for renaming intermediary results, and in some cases a separate register stack to hold function arguments and routine return addresses.
- RISC:** Reduced Instruction Set Computer. A CPU with its instruction set that is simpler in comparison with the earlier Complex Instruction Set Computers (CISCs). The instruction set was reduced to simple instructions that ideally should execute in one cycle.
- Shared Memory (SM):** Memory configuration of a computer in which all processors have direct access to all the memory in the system. Because of technological limitations on shared bandwidth generally not more than about 16 processors share a common memory.
- shmem:** One-sided fast communication library first provided by Cray for its systems. However, **shmem** implementations are also available for SGI and HP AlphaServer systems.
- SMP:** Symmetric Multi-Processing. This term is often used for compute nodes with shared memory that are part of a larger system and where this collection of nodes forms the total system. The nodes may be organised as a ccNUMA system or as a distributed memory system of which the nodes can be programmed using OpenMP while inter-node communication should be done by message passing.
- TLB:** Translation Look-aside Buffer. A specialised cache that holds a table of physical addresses as generated from the virtual addresses used in the program code.
- Torus:** Structure that results when the end points of a grid are wrapped around to connect to the starting points of that grid. This configuration is often used in the interconnection networks of parallel machines either with a 2-D grid or with 3-D grid.
- Vector register:** A multiple entry register (typically 128–256) that hold the operands to be processed by a vector unit. Using vector registers controlled by vector instructions guarantees the production of results every cycle for the amount of operands in the register.
- Vector unit (pipe):** A pipelined functional unit that is fed with operands from a vector register and will produce a result every cycle (after filling the pipeline) for the complete contents of the vector register.
- Virtual Shared Memory:** The emulation of a shared memory system on a distributed memory machine by a software layer.

VLIW processing: Very Large Instruction Word processing. The use of large instruction words to keep many functional units busy in parallel. The scheduling of instructions is done statically by the compiler and, as such, requires high quality code generation by that compiler. VLIW processing has been revived in the IA-64 chip architecture, there called EPIC (see above).

Acknowledgments

It is not possible to thank all people that have been contributing to this overview. Many vendors and people interested in this project have been so kind to provide me with the vital information or to correct us when necessary. Therefore, we will have to thank them here collectively but not less heartily for their support.

References

- [1] C. Amza, A.L. Cox, S. Dwarkadas, P. Keleher, H. Lu, R. Rajamony, W. Yu, W. Zwaenepoel, *TreadMarks: Shared Memory Computing on Networks of Workstations*, to appear in IEEE Computer (also: www.cs.rice.edu/~willy/TreadMarks/papers.htm)
- [2] W. Anderson, D.W. Hess, P. Briggs, A. Khoklov, R. Rosenberg, C.S. Hellberg, M. Lanzagorta, *Early Experience with Scientific Programs on the Cray MTA-2*, Proc. SC2003, November 2003, Phoenix, AZ, USA.
- [3] The ASCI program: <http://www.llnl.gov/asci/>.
- [4] S.H. Bokhari, J.R. Sauer, *Sequence Alignment on the Cray MTA-2*, Proc. HICOMB2002, April 2002, Fort Lauderdale, USA.
- [5] L. Carter, J. Feo, A. Snively, *Performance and Programming Experience on the Tera MTA*, Proc. SIAM Conf. on Parallel Processing, March 1999.
- [6] K. Cassirer, B. Steckel, *Block-Structured Multigrid on the Cenju*, 2nd Cenju Workshop, October 1998, Sankt Augustin, Germany.
- [7] R. Chandra, L. Dagum, D. Kohr, D. Maydan, J. McDonald, R. Menon, *Parallel Programming in OpenMP*, Morgan Kaufmann Publishers Inc., January 2001.
- [8] D.E. Culler, J.P. Singh, A. Gupta, *Parallel Computer Architecture: A Hardware/Software Approach*, Morgan Kaufmann Publishers Inc., August 1998.
- [9] J.J. Dongarra, *Performance of various computers using standard linear equations software*, Computer Science Technical Report CS-89-85, Univ. of Tennessee, July 7, 2002 (early version).
- [10] Overview page of Distributed Shared Memory systems: <http://www.ics.uci.edu/~javid/dsm.html>.
- [11] T.H. Dunigan, M.R. Fahey, J.B. White, P.H. Worley, *Early Evaluation of the Cray X1*, Proc. SC2003, November 2003, Phoenix, AZ, USA.
- [12] T.H. Dunigan, ONRL Cray XD1 Evaluation, www.csm.ornl.gov/~dunigan/xd1/.
- [13] Directory with EuroBen results: www.euroben.nl/results.
- [14] M.J. Flynn, *Some computer organisations and their effectiveness*, IEEE Trans. on Computers, Vol. C-21, 9, (1972) 948–960.
- [15] A. Geist, A. Beguelin, J. Dongarra, R. Manchek, W. Jaing, and V. Sunderam, *PVM: A Users' Guide and Tutorial for Networked Parallel Computing*, MIT Press, Boston, 1994.
- [16] D.B. Gustavson, Q. Li, *Local-Area MultiProcessor: the Scalable Coherent Interface*, SCIzzL Report, Santa Clara University, Dept. of Computer Engineering, 1995. Available through: www.scizzl.com.
- [17] J.M.D. Hill, W. McColl, D.C. Stefanescu, M.W. Goudreau, K. Lang, S.B. Rao, T. Suel, T. Tsantilas, R. Bisseling, *BSPlib: The BSP Programming Library*, Technical report PRG-TR-29-9, Oxford University Computing Laboratory, May 1997. (Compressed Postscript with ANSI C examples, 142K; Compressed Postscript with Fortran 77 examples, 141K)

- [18] R. W. Hockney, C. R. Jesshope, *Parallel Computers II*, Bristol: Adam Hilger, 1987.
- [19] T. Horie, H. Ishihata, T. Shimizu, S. Kato, S. Inano, M. Ikesaka, *AP1000 architecture and performance of LU decomposition*, Proc. Internat. Symp. on Supercomputing, Fukuoka, Nov. 1991, 46–55.
- [20] HPC Challenge Benchmark, icl.cs.utk.edu/hpcc/.
- [21] High Performance Fortran Forum, *High Performance Fortran Language Specification*, Scientific Programming, **2**, 13, (1993) 1–170.
- [22] D.V. James, A.T. Laundrie, S. Gjessing, G.S. Sohi, *Scalable Coherent Interface*, IEEE Computer, **23**, 6, (1990), 74–77. See also:
Scalable Coherent Interface: <http://sunrise.scu.edu/>
- [23] D.J. Kerbyson, A. Hoisie, S. Pakin, F. Petrini, H.J. Wassermann, *A performance evaluation on an Alpha EV7 processing node*, Int. J. of High Perf. Comp. Appl., **18**(2), 199–209, Summer 2004.
- [24] M. Snir, S. Otto, S. Huss-Lederman, D. Walker, J. Dongarra, *MPI: The Complete Reference Vol. 1, The MPI Core*, MIT Press, Boston, 1998.
- [25] W. Gropp, S. Huss-Ledermann, A. Lumsdaine, E. Lusk, B. Nitzberg, W. Saphir, M. Snir *MPI: The Complete Reference, Vol. 2, The MPI Extensions*, MIT Press, Boston, 1998.
- [26] <http://www.myrinet.com>
- [27] N.J. Boden, D. Cohen, R.E. Felderman, A.E. Kulawik, C.L. Seitz, J.N. Seizovic, Wen-King Su, *Myrinet – A Gigabit-per-second Local Area Network*, IEEE Micro, **15**, No. 1, Jan. 1995, 29–36.
- [28] W.E. Nagel, *Applications on the Cenju: First Experience with Effective Performance*, 2nd Cenju Workshop, October 1998, Sankt Augustin, Germany.
- [29] Web page for the NAS Parallel benchmarks NPB2: <http://science.nas.nasa.gov/Software/NPB/>
- [30] OpenMP Forum, *Fortran Language Specification, version 1.0*, Web page: www.openmp.org/, October 1997.
- [31] F. Petrini, D.J. Kerbyson, S. Fakin, *The Case of the Missing Supercomputer Performance: Achieving Optimal Performance on the 8,192 Processors of ASCI Q*, Proc. SC2003, November 2003, Phoenix, AZ, USA.
- [32] F. Petrini, W. Feng, A. Hoisie, S. Coll, E. Frachtenberg, *The Quadrics Network: High-performance clustering technology*, IEEE Micro Jan/Febr 2002, 46–57.
- [33] T. Sato, *The Earth Simulator*, Proc. SC2002, Nov. 2002, IEEE Press. (Proceedings only on CD-ROM.)
- [34] T. Shanley, *Infiniband Network Architecture*, Addison-Wesley, Nov. 2002.
- [35] D.H.M. Spector, *Building Unix Clusters*, O’Reilly, Sebastopol, CA, USA, July 2000.
- [36] A.J. van der Steen, *Exploring VLIW: Benchmark tests on a Multiflow TRACE 14/300*, Academic Computing Centre Utrecht, Technical Report TR-31, April 1990.
- [37] A.J. van der Steen, *The benchmark of the EuroBen Group*, Parallel Computing **17** (1991) 1211–1221.
- [38] A.J. van der Steen, *Benchmark results for the Hitachi S-3800*, Supercomputer, **10**, 4/5, (1993) 32–45.
- [39] A.J. van der Steen, ed. *Aspects of computational science*, NCF, The Hague, 1995.
- [40] A.J. van der Steen, *Benchmarking the Silicon Graphics Origin2000 System*, Technical Report WFI-98-2, Dept. of Computational Physics, Utrecht University, The Netherlands, May 1998. The report can be downloaded from: www.euroben.nl/reports/

-
- [41] A.J. van der Steen, *An evaluation of some Beowulf clusters*, Technical Report WFI-00-07, Utrecht University, Dept. of Computational Physics, December 2000. (Also available through [www.euroben.nl, directory reports/.](http://www.euroben.nl/directory/reports/))
 - [42] A.J. van der Steen, *An evaluation of Itanium 2-based high-end servers*, Technical Report HPCG-2004-04, Utrecht University, High Performance Computing Group, November 2004. (Also available through [www.euroben.nl, directory reports/.](http://www.euroben.nl/directory/reports/))
 - [43] T.L. Sterling, J. Salmon, D.J. Becker, D.F. Savaresse, *How to Build a Beowulf*, The MIT Press, Boston, 1999.
 - [44] H.W. Meuer, E. Strohmaier, J.J. Dongarra, H.D. Simon, *Top500 Supercomputer Sites*, 23rd Edition, June 2004, The report can be downloaded from: [www.top500.org/.](http://www.top500.org/)
 - [45] Task Force on Cluster Computing home page: www.clustercomputing.org.