System X - A review
CS466 Project 2 Fall '04

Instructor: Prof. Mitchell Theys

Hareesh Nagarajan
Dept. of Computer Science
University of Illinois at Chicago
hnagaraj@cs.uic.edu

Raj Bharath Swaminathan
Dept. of Computer Science
University of Illinois at Chicago
rswamina@cs.uic.edu

Abstract

System X is a super computer conceived in February 2003 by a team of Virginia Tech faculty and administrators and represents what can happen when the academic and IT organizations collaborate [1]. As of December 1, 2004 the System X is the 7th fastest super computer in the world [2]. It is capable of performing 12.25 trillion floating point operations a second. In this paper we will review the processor, the interconnection network fabric used and the cooling system used in the System X.

1 System X Specifications

- **Nodes** 1100 Apple XServe G5 2.3 GHz dual processor cluster nodes (4 GB RAM, 80 GB S-ATA HD)
- **Primary Communication** 24 Mellanox 96 port InfiniBand switches (4X InfiniBand, 10 Gbps)
- **Secondary Communication** 6 Cisco 4506 Gigabit Ethernet switches
- **Cooling** Liebert X-treme Density System cooling
- **Software** Mac OS X, MVAPICH, XLC & XLF
- **Current Linpack** Rpeak = 20.24 Teraflops, Rmax = 12.25 Teraflops

2 History

In early 2003 the construction of the System X began when 1100 Power Mac G5 were put into racks. In parallel the *Terascale Computing Facility* (TCF) at *VTech* setup a team that began writing software for the super computer. This included writing device drivers, hand optimizing numerical libraries and code porting. The system was finally built in September 2003 at a relatively low cost of $5.2 million. But unfortunately the system was not ready to perfom scientific computation as Error Correcting Codes (ECC) RAM was required and the G5 did not support it. So in early 2004 the team, with the help of Apple, upgraded each of the systems to an XServe G5. The upgrade costed $600,000. So at $5.8 million, the System X officially became the cheapest super computer in

the world.

2.1 The Upgrade

Before the upgrade the System X clocked at 9.8 teraflops. The upgrade saw nearly 25% improvement in the system performance. According to [3] the new XServe servers (custom built by Apple Inc.) are about 15% faster than their counterpart, the G5-desktop machines. This makes the new System X operate about 20 percent faster than it was before the upgrade improving the performance of the system by almost 2 teraflops. The extra 5-percent performance boost can be attributed to optimized software.
Typically the System X runs several projects simultaneously, each tying up 400 to 500 processors for research into weather and molecular modeling.

3 XServe G5

3.1 Introduction

System X consists of a cluster of 1100 XServe G5 nodes each composed of dual processors amounting to a total of 2200 processors. The XServe G5 is a 1U rack mount server built for intensive server tasks and high performance computing. It uses the processing performance of dual PowerPC G5's combined with high bandwidth, server-optimized system I/O and fast internal storage. Refer figure 1

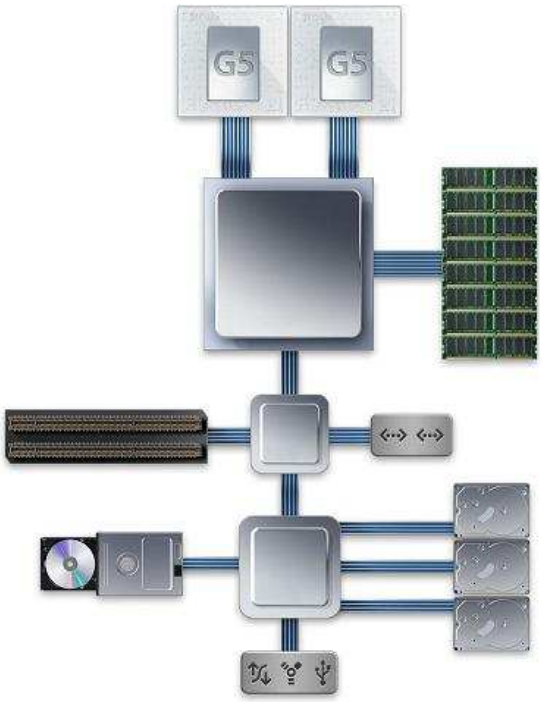

Figure 1: XServer G5

3.2 Key Features

3.2.1 Dual 2GHz PowerPC G5 Processors

According to [10] PowerPC G5 Processors are designed for high performance general purpose processing. Multiple pipelined execution units, branch prediction, and a SIMD, or vector processing (Altivec) unit, combine to allow up to 215 in-flight instructions. With each clock cycle, up to eight instructions can be fetched from the direct-mapped 64K L1 instruction cache, broken down, and dispatched into the execution units, while 32K of write-through, two-way associative L1 data cache can fetch up to eight active data streams, which are loaded into data registers behind the execution units. Different types of instructions are processed concurrently by the execution units, which include two floating-point units, two integer units, two

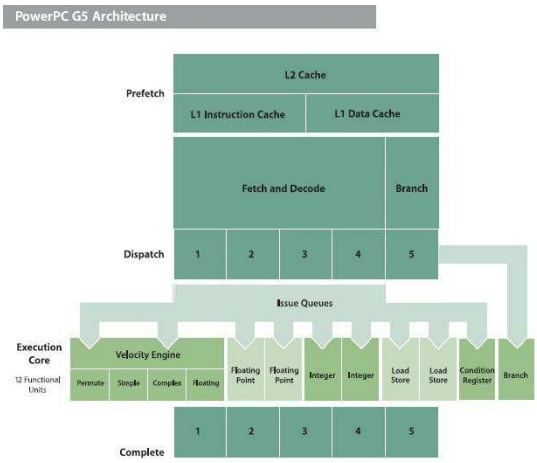

Figure 2: Power PC G5

load/store units, a condition register unit, a branch prediction unit, and a vector processing unit. This dual-pipeline 128-bit vector engine performs SIMD processing, applying a single instruction to multiple data simultaneously, and uses a set of 162 specialized SIMD instructions for optimal performance. Refer figure 2.
More details are provided in section 4.

3.2.2 Gigahertz front side bus

Each PowerPC G5 has a dedicated front side bus, dubbed the industry's fastest, that transfers instructions and data at rates of up to 8GB/s.for a combined throughput of up to 16GB/s. Conventional bidirectional buses carry data to and from the processor over the same link, incurring delays when the bus switches direction and while the processor and the memory controller negotiate use of the bus. Dual-channel unidirectional buses enable data to flow to and from the processor simultaneously, eliminating negotiation overhead and more than doubling the effective data rate. With a clock speed of 2GHz, the front-side bus of the 90nm IBM PowerPC 970FX is theoretically capable of operating at up to 1GHz, for an aggregate bandwidth of up to 10GBps. This type of bus architecture achieves its highest throughput only when the number of reads and writes are fairly well balanced. Bidirectional bus architecture, as seen on Intel IA-64 and AMD Athlon processors, achieves a lower peak throughput, but it can deliver its peak throughput in either direction, making it better suited for applications that perform mostly reads or writes.

3.2.3 Advanced ECC memory technology

A 128-bit memory controller speeds data in and out of main memory at up to 6.4GB/s. Each Xserve G5 on System X supports 8GB of fast, 400MHz RAM with Error Correction Code (ECC) protection. The system controller uses this ECC data To identify single-bit errors and corrects them on the fly, preventing unplanned system shutdowns and protecting data integrity. In the rare event of multiple-bit errors, the system controller detects the error and triggers a system notification to prevent bad data from corrupting further operations.

3.2.4 High-bandwidth I/O

Dual onboard Gigabit Ethernet interfaces and FireWire 800 ports offer high performance connectivity.

3.2.5 Innovative hardware monitoring

Dedicated monitoring hardware integrates with industry-leading software for remote monitoring of one or many Xserve G5 systems.
These features gives each XServe G5, a High-

Density Cluster Node Configuration where one dual processor node can execute over 9 billion double-precision floating-point operations per second, or 9 gigaflops per U.

4 The PowerPC G5 processor

System X is powered by the PowerPC G5 Processor. In this section we will describe the G5 processor.

4.1 Introduction

According to [9] the PowerPC G5 also known as the *PowerPC 970FX* is the result of a long-standing partnership between Apple and IBM. In 1991, they co-created a PowerPC architecture that could support both 32-bit and 64-bit instructions. Leveraging this design, Apple went on to bring 32-bit RISC processing to desktop and portable computers, while IBM focused on developing 64-bit processors for enterprise servers. The PowerPC G5 represents a convergence of these efforts: Its design is based on the PowerPC instruction set, as well as the award-winning IBM POWER Architecture.

4.2 PowerPC Architecture

The PC in PowerPC stands for performance computing. Descended from the POWER architecture, it was introduced in 1993 and was designed from the beginning to run on a broad range of machines, from battery-operated handhelds to supercomputers and mainframes.

Born of an alliance between Apple, IBM, and Motorola (also known as the AIM alliance), the PowerPC was based on POWER[1], but with a number of differences. For instance, PowerPC is open-endian, supporting both big-endian and little-endian memory models, where POWER had been big-endian. The original PowerPC design also focused on floating-point performance and multiprocessing capabilities. Still, it did and still does include most of the POWER instructions. Many applications work on both, perhaps with a recompile to make the transition. While IBM and Motorola develop their chips separately, at the user level, all PowerPC processors run the

DEVICE SUMMARY TABLE

Manufacturer	IBM
Part Number	PPC970FX6SB-CCB
Description	PowerPC970FX Processor
Package Type	576-Ball BGA
Die Markings	© 2003 IBM "M"
Die Size	9.3mm x 7mm = 65.1 mm²
Number of Metal Levels	10
Number of Poly Levels	1
Die Thickness	780 µm
Interconnect Level Width/Pitch (µm)	
Metal Level - M1	0.15 / 0.25
Metal Level - M2	0.20 / 0.30
Metal Level – M3	0.20 / 0.30
Poly Layer - P1	0.05 / 0.09
Process Type	Strained Silicon/SOI
Process Generation	90nm
Feature Measured to Determine Process Generation	Gate Length (Logic, TEM)
Transistor Gate Length	48nm
SRAM Cell Size	1.1 µm²

Figure 3: Device Summary Table

[1]POWER stands for Power Optimization with Enhanced RISC and is the main processor in many IBM servers, workstations, and supercomputers. Descended directly from the 801 CPU, it is a 2nd-generation RISC processor. Introduced in 1990 to power the RS, or RISC System/6000 UNIX workstations (now called the eServer and pSeries), POWER exists in iterations from POWER1, POWER2, POWER3, and the current top end, POWER4 is the father of the PowerPC 970 processor

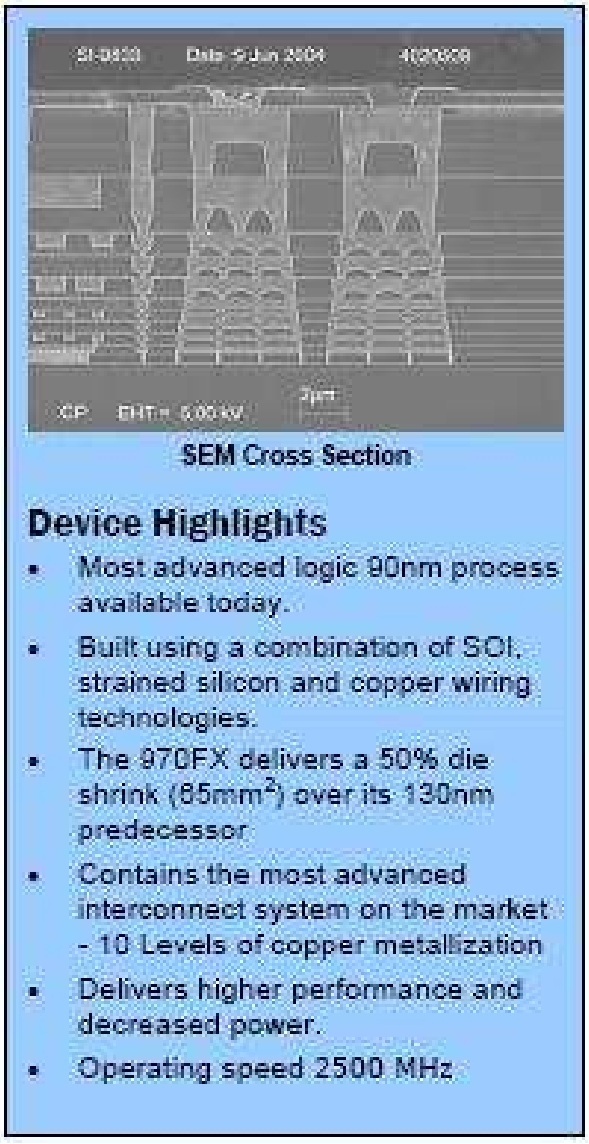

Figure 4: Device Highlights

same core PowerPC instruction set, ensuring full ABI compatibility for the software products that run on them.

Aside from compatibility, one of the best things about the PowerPC architecture is that it is open: it specifies an instruction set architecture (ISA) that allows anyone to design and fabricate PowerPC-compatible processors; and source code for software modules developed in support of PowerPC is freely available. Finally, the small size of the PowerPC core leaves a great deal of room on each die for additional components, from added cache to coprocessors, allowing for an amazing amount of design flexibility.

4.3 PowerPC 900 family

The 64-bit PowerPC 970, a single-core version of the POWER4, can process 200 instructions at once at speeds of up to 2 GHz and beyond – all while consuming just tens of watts of power. Its low power consumption makes it a favorite with notebooks and other portable applications on the one hand, and with large server and storage farms on the other. Its 64-bit capability and single instruction multiple data (SIMD) unit accelerate computationally intensive workloads such as multimedia and graphics. It is used in Apple desktops, Apple Xserve servers, imaging applications, and – increasingly – in networking applications. The Apple Xserve G5 features the first use of the new PowerPC 970FX – the first chip made using both strained silicon and SOI technologies together, enabling the chip to run at even greater speeds with even less power consumption.

Although the PowerPC 970 is a part that would cost considerably less to manufacture and sell, its performance actually exceeds the

POWER4 processor in many areas. The reason for this apparent paradox is that the POWER4 processor had been designed for the high cost, continuous availability server market, and in some areas, performance had been traded off to obtain near-absolute reliability guarantees. As an example, in an article published in Microprocessor Reports in 1999, IBM described its use of thicker gate oxides in the POWER4 processor to obtain a failure rate that is two orders of magnitude better than comparable processors from most other manufacturers. The cost of the thicker oxides is the reduced drive current of the transistor and consequently slower switching speeds of the transistors on the POWER4 processor. In the case of the PowerPC 970, the processor does not need to meet similar reliability requirements as the POWER4 processor, and as a consequence, circuit and process technology can be tweaked to obtain higher performance by trading away the near-absolute reliability required by the POWER4 processor.

Like the POWER4, the POWER5 unifies the POWER and PowerPC architectures. Released in 2004, it features communications acceleration, chip multiprocessing, and simultaneous multithreading (SMT). Recent internal testing shows that POWER5-based eServer systems are expected to offer four times the performance of the first POWER4-based servers. The introduction of these systems could bring the next likely upgrade to System X

4.4 The device

Refer figures 3 and 4 which were obtained from [14].

4.5 Design Philosophy

4.5.1 Pipeline Depth

The PowerPC 970 processor uses a "wide and deep" design and has a large number of functional units and a relatively long pipeline structure. There are 9 pipeline stages devoted to instruction fetch and decode, 5 to 13 pipeline stages are used for the out of order execution units. Simple integer instructions can execute in 5 stages, whereas more complex vector floating point instructions may take as many as 13 stages to complete execution. The long 9 stage instruction and fetch decode pipeline is a result of some of the older PowerPC instructions which are complex, and require an instruction cracking process whereby some complex instructions are cracked into simpler, more RISC-like instructions. These simpler instructions are then sent to the dispatch and execution units in the processor core. Although this process superficially resembles the micro-op or risc-op decoding steps seen in various x86 processors, the PowerPC instruction set is not nearly as complex as the venerable x86 ISA, and only few instructions in the PowerPC ISA needs to be cracked into simpler instructions for execution.

The P4, in comparison, takes a "narrow and deep" approach that involves moving a large number of instructions onto the chip and pushing them through fewer functional units in more, faster cycles. The P4’s 20-stage pipeline can hold up to 126 instructions in various stages of execution. When you combine the P4’s high clock speed with its very large instruction window, you get a processor that can grind through a large volume of instructions in a rapid-fire, serial manner. The 970 uses an extremely wide execution core and a 16-stage (integer) pipeline that, while not as deep as the

P4's, is nonetheless built for speed. The 970 can have a whopping 200 instructions on-chip in various stages of execution, a number that dwarfs the P4's 126-instruction one. The reason that the 970's pipeline is a little shorter than the P4's is because the 970's pipeline lacks the "drive" stages, which the P4 inserts in order to allow signals to propagate across the chip. Inserting whole pipeline stages just to account for wire delay is necessary only if you plan to push a design to very high clock speeds. The PowerPC 970's GHz rating, therefore is at a much lower number than the P4's.This difference in clockspeed headroom reflects a fundamental difference in the approaches of the PPC 970 and the Pentium 4. As is evidenced by the enormous power requirements that go along with its high clock speed, the P4 is really aimed at single-processor desktop systems. Intel sells the P4 Xeon for use in 2-way and 4-way server setups, but the price and power consumption of such machines restrict them primarily to the server closet. The PowerPC 970, on the other hand, is designed from the ground up with multiprocessing in mind - IBM intends to see the 970 used in 4-way or higher desktop SMP systems. So instead of increasing the performance of desktop and server systems by using a few narrow, very high clockspeed, high power CPUs IBM would rather see multiple, slower but wider, lower power CPUs ganged together via very high-bandwidth connections.

For more details refer [11], [12] and [13].

4.5.2 Branch prediction

The 970's instruction fetch logic fetches 8 instructions per cycle from the L1 I-cache into an instruction queue, and on each fetch the front end's branch unit scans these eight instructions in order to pick out up to two

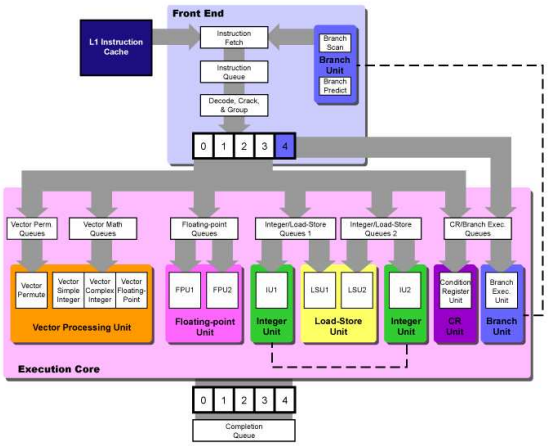

Figure 5: PowerPC G5 Execution Core

branches. If either of the two branches is conditional, then the branch prediction unit will predict the condition's outcome (taken or not taken) and/or its target address using one of two branch prediction schemes.

The first branch prediction scheme employed by the 970 is the standard BHT-based scheme described above. The 970s BHT has 16K entries–four times the number of entries in the P4's BHT and eight times the number of entries in the G4's BHT. For each of these 16K entries, a 1-bit flag tells the branch predictor whether the branch should be taken or not taken.

The second scheme involves another 16K-entry table called the global predictor table (as opposed to the local predictor table of the first scheme). Each entry in this global predictor table is associated with an 11-bit vector that records the actual execution path taken by the previous 11 fetch groups. The processor uses this vector, which it constantly keeps up-to-date with the latest information, to set another 1-bit flag for the global predictor table that specifies whether the branch should be taken or not taken.

Finally, there's a third 16K-entry table that's used by the 970's front end to keep track of which of the two schemes works best for each branch. When each branch is finally evaluated, the processor compares the success of both schemes and records in this selector table which scheme has done the best job so far of predicting the outcome of that particular branch.

4.5.3 Decode, Dispatch and Issue

The fetch and decode pipeline breaks down instructions to a smaller, simpler format for use inside the CPU. This is similar to the P4, which breaks down each x86 CISC instruction into smaller micro-ops (or "uops"), which more or less resemble the instructions on a RISC machine. The 970 breaks its instructions down into what it calls "IOPs". Like uops on the P4, it is these IOPs that are actually executed out-of-order by the 970's execution core. And also like uops, cracking instructions down into multiple, more atomic and more strictly defined IOPs can help the back end squeeze out some extra instruction-level parallelism (ILP) by giving it more freedom to schedule code. The IOPs are packaged into groups. A "group" consists of five IOPs arranged in program order according to certain rules and restrictions. It is these organized and packaged groups of five IOPs, and not single IOPs in isolation, that the 970 dispatches in-order to the six issue queues in its execution core. Once the IOPs in a group reach their proper issue queues, they can then be issued out of order to the execution units at a rate of 8 IOPs/cycle for all the queues combined. Before they reach the completion stage, however, they need to be placed back into their group so that an entire group of 5 IOPs can be completed each cycle.

The group's five slots must be populated with IOPs in program order, starting with the oldest IOP in slot 0 and moving up to newest IOP in slot 4. Another rule is that all branch instructions must go in slot 4, and slot 4 is reserved for branch instructions only. This means that if the front end cannot find a branch instruction to put in slot 4, then it can issue one less instruction that cycle.

4.5.4 Integer Units

There are two general-purpose IUs that execute almost all integer instructions. The IUs are slightly specialized as one of them is capable of fixed-point divides, and the other can handle special-purpose register (SPR) operations. There is a separate dedicated unit for handling logical operations related to the PPC's condition register: the condition register logical unit (CRU).

4.5.5 Load-Store Units

The 970 has two identical load-store units that execute all of the LOADs and STOREs for the entire chip.

4.5.6 Floating Point Units

There are two FPUs that take a minimum of 6 cycles to finish executing the fastest floating-point instructions. (Some instructions take many more cycles). Single- and double-precision floating-point operations take the same amount of time. The FPUs are fully pipelined for all operations except floating-point divides, which are very costly in terms of cycles and stall both FPUs

4.5.7 Vector Units

The 970 can dispatch four vector IOPs per cycle total to the two vector issue queues–two IOPs per cycle maximum to the 16-entry VPU queue and two IOPs per cycle maximum to the 20-entry VALU queue. Each of the two queues can then issue one vector operation per cycle out-of-order to any of the units attached to it. So this means that the 970 can issue one IOP per cycle to the VPU and one IOP per cycle to any of the VALU's three sub-units.

4.5.8 PowerPC 970FX Microprocessor Overview

The 970FX is a 64-bit PowerPC reduced instruction set computer (RISC) microprocessor with AltiVec extensions. The single-instruction, multiple-data (SIMD) operations that accelerate data intensive processing tasks. This processor is designed to support multiple system configurations ranging from desktop and low end server applications for uniprocessor up through 4-way symmetric multiprocessor (SMP) configurations.

For more details [15].

The PowerPC 970FX RISC Microprocessor is comprised of three main components:

- Core which includes the vector processing execution units (VPU)
- Storage subsystem (STS) which includes the core interface logic, non-cacheable unit, L2 cache, controls, and the bus interface unit
- Pervasive functions[2]

[2]We were unable to find adequate information on this topic

4.6 PowerPC 970FX Functional Units

This section provides an overview of the 970FX microprocessor core, VPU, storage, and bus interface units. It includes a summary and details of key design fundamentals. The information below has been obtained from [16].

4.6.1 Key Design Fundamentals of the Microprocessor Core

Please check the CS466 - Fall 2004, class SWIKI for more details [8].

4.6.2 Detailed Features of the Microprocessor Core

Please check the CS466 - Fall 2004, class SWIKI for more details [8].

5 Internetworking the System X

Before we get into the intricacies of the internetworking mechanism in the System X we will first take a brief excursion into concepts from interconnection networks.

5.1 Introduction

Using shared media (broadcast) for the nodes in the system to communicate with each other has the obvious disadvantage of collisions. **Switches** allow communication directly from source to destination, without having intermediate nodes interfere with these signals. The following are some of the different types of switches [4]:

- A **crossbar switch** allows any node to communicate with any other node in one pass through interconnection so that there is no network conflict. These switches usually do not scale well. Refer figure 6.
- An **Omega interconnection** uses less hardware. But this usually means that blocking [3] is more likely. Here the destination address specifies the path. Refer figure 7.
- A **fat tree switch** has more bandwidth added higher in the tree to match the requirements of common communication patterns. This makes the network have constant bandwidth at all levels. It must be noted that the switching nodes do not have full crossbar-like connectivity. Refer figures 8 and 9.

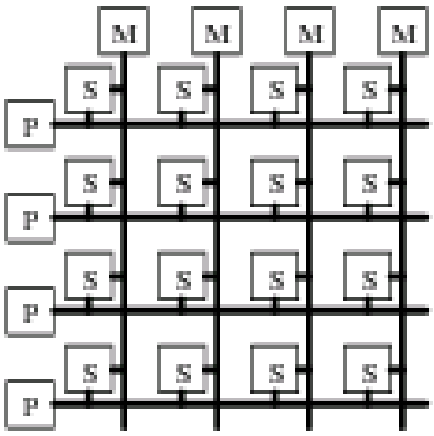

Figure 6: Crossbar switch

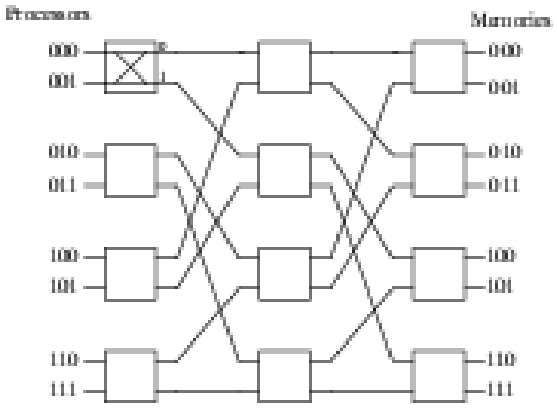

Figure 7: Omega switch

The System X utilizes fat tree or **Constant Bi-Directional Bandwidth** networks to construct a large node count[4] non-blocking switch configurations. Refer figure 10. For switching the System X makes use of **integrated crossbars** with relatively low number of ports to build a non-blocking switch topology that supports a much larger number of endpoints.

As we mentioned in Section 1, System X makes use of 24 Mellanox switches. Each one of these switches has 96 ports. So the maximum number of nodes possible is $96/2 * 24 = 1152$ which is approximately 1100 nodes.

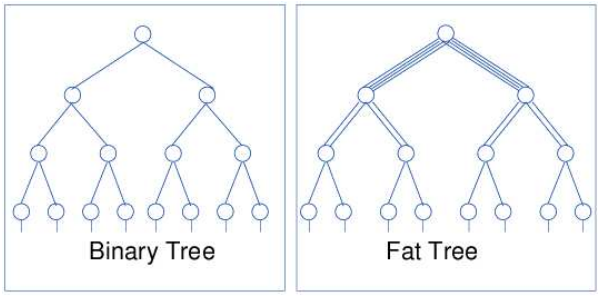

Figure 8: A binary tree and a Fat Tree

5.2 Why use Infiniband?

Infiniband is a switch based serial I/O interconnect architecture operating at a base

[3] also called contention
[4] 1100 nodes

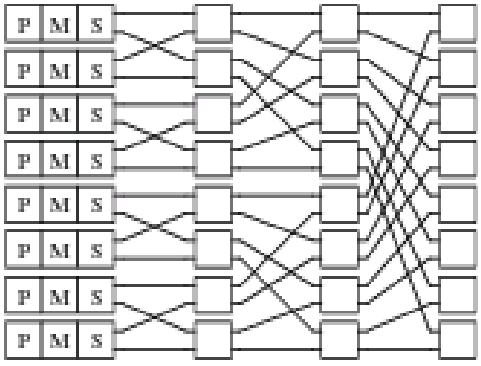

Figure 9: Fat tree switch

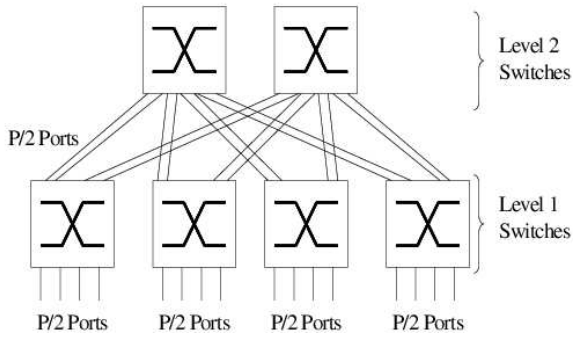

Figure 10: Constant Bi-Directional Bandwidth Network

speed of 10 GB/s in each direction per port. In this section we will see the advantages of Infiniband over Gigabit ethernet. There are two important aspects of constructing CBB networks [5]:

Topology: Topology defines how the network is physically connected together. From a purely physical perspective, point to point networks are virtually identical and any topology can be constructed using similar crossbar switch elements.

Link Level Forwarding: Forwarding deals with the way individual packets are actually routed through the physical topology and thus effectively define a logical topology that exists on top of the physical one.

While all point to point technologies can implement the same physical topologies, this is not the case with the logical topology created by different forwarding algorithms. In fact, unlike InfiniBand, bridged Ethernet imposes severe constraints on what types of logical topologies are implementable and typically can not make efficient use of the underlying physical topology. In order to overcome this limitation of layer 2 Ethernet switches, it is necessary to use layer 3 (and higher) IP switches. Using layer 3 and above switches severely impact both, the price and performance of the base Ethernet offering.

It is for this reason and because Infiniband has a much lower price to performance ratio that we feel the System X designers chose the Infiniband over Gigabit Ethernet to act as the networking fabric.

The System X makes use of the 4x Infiniband links. Refer figures 11 and 12. Since we were unable to obtain the actual schema of the System X network we would assume that the cluster looks like any other Infiniband cluster such as the one depicted in figure 13.

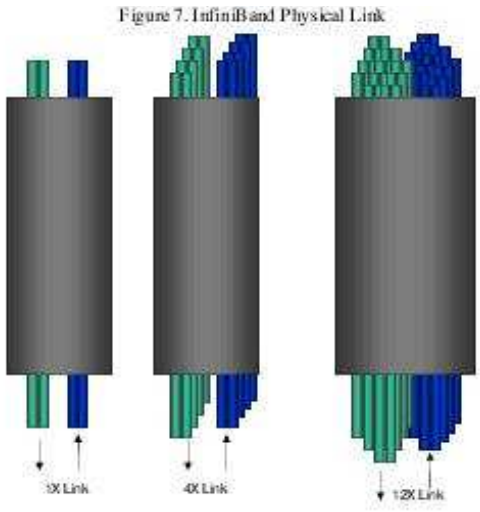

Figure 11: Links

Table 3. InfiniBand Link Rates

InfiniBand Link	Signal Count	Signalling Rate	Data Rate	Fully Duplexed Data Rate
1X	4	2.5 Gb/s	2.0 Gb/s	4.0 Gb/s
4X	16	10 Gb/s	8 Gb/s	16.0 Gb/s
12X	48	30 Gb/s	24 Gb/s	48.0 Gb/s

Figure 12: Data link rates

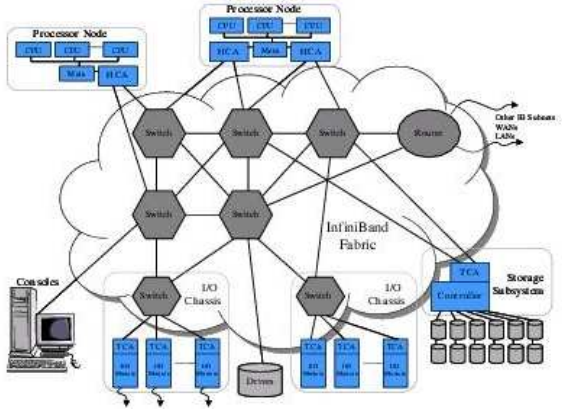

Figure 13: Infiniband System Fabric

6 Cooling system for the processor

The new Apple G5 and XServe have a new liquid cooling system [6]. The figure 14 shows the side view of the liquid cooling system.

1. G5 processor at point of contact to the heatsink.
2. G5 processor card from IBM
3. Heatsink
4. Cooling fluid output from the radiator to the pump
5. Liquid cooling system pump
6. Pump power cable
7. Cooling fluid radiator input from the G5 processor
8. Radiant grille
9. Airflow direction

The figure 15 shows the rear view of the liquid cooling system.

1. Liquid cooling system pump
2. G5 processors
3. Radiator output
4. Radiator
5. Pump power cable
6. Radiator input

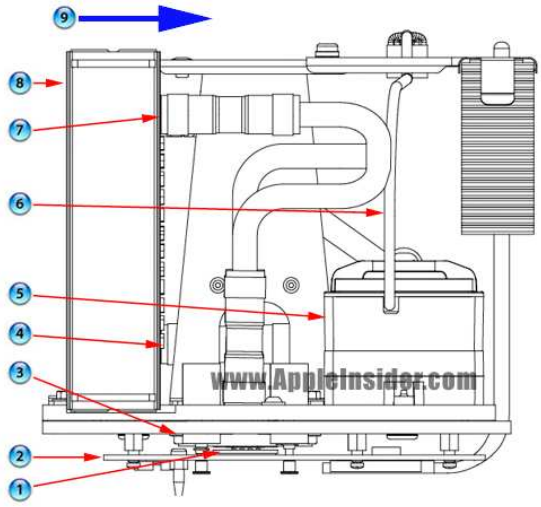

Figure 14: Side view of the liquid cooling system

Figure 16: Liebert X-treme Density System cooling

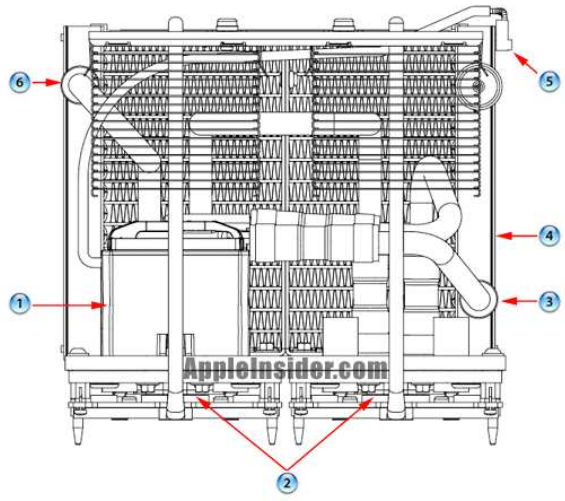

Figure 15: Rear view of the liquid cooling system

7 Cooling system for System X

Lieberts XDR system utilizes a cooling module that is attached to the back door of the computer rack enclosure [7]. Fans in the module move room temperature air from the front of the enclosure, past the equipment in the rack, past a cooling coil and expel it from the back of the unit, chilled to the point where the impact on the room is close to neutral. The XDR system can be configured to take care of uneven heat loads within the room. Refer figure 16.

8 Software Used

- **Operating system** Mac OS X
- **MVAPICH** is a high performance imple-

mentation of MPI-1 over InfiniBand based on MPICH1.

- **Compilers** XL C/C++ Advanced Edition V6.0 for Mac OS X and XL Fortran Advanced Edition for Mac OS X

9 Results and Conclusion

- The PowerPC G5 is designed from ground up for symmetric multiprocessing so this makes it an ideal processor for a super computer such as the System X.
- Dual independent frontside buses allow each processor to handle its own tasks at maximum speed with minimal interruption.
- With sophisticated multiprocessing capabilities built in, Mac OS X and Mac OS X Server dynamically manage multiple processing tasks across the two processors. This allows dual PowerPC G5 systems to accomplish up to twice as much as a single-processor system in the same amount of time, without requiring any special optimization of the application.
- By using Infiniband and highly optimized software for message passing, the System X keeps overheads low and maximizes performance.

10 Trademarks

References

[1] Terascale Computing Facility *http://www.tcf.vt.edu/systemX.html*

[2] top500.org *http://www.top500.org/lists/2004/11/*

[3] Wired.com *http://www.wired.com/news/mac/0,2125,65476,00.html?tw=wn_tophead_2*

[4] Hennessy and Patterson *Computer Architecture: A quantitative approach*

[5] Mellanox 2003 - Infiniband versus the Ethernet *http://www.mellanox.com/technology/shared/IB_vs_Ethernet_Clustering_WP_100.pdf*

[6] Apple Insider 2004 - Apple's new liquid cooling system *http://www.appleinsider.com/article.php?id=508*

[7] Liebert.com *http://www.liebert.com/dynamic/displayproduct.asp?id=1077&cycles=60hz*

[8] CS-466 Swiki - Fall 2004 *http://expedition.cs.uic.edu:8080/cs466/38*

[9] Apple Inc. - PowerPC G5 *http://images.apple.com/powermac/pdf/PowerPCG5_WP_06092004.pdf*

[10] Apple Inc - XServe G5 *http://images.apple.com/server/pdfs/L301323A_XserveG5_TO.pdf*

[11] Arstechnica.com - PowerPC *http://arstechnica.com/wankerdesk/3q02/powerpc.html*

[12] Arstechnica.com - PowerPC
http://arstechnica.com/cpu/02q2/ppc970/ppc970-1.html

[13] Arstechnica.com - PowerPC
http://arstechnica.com/cpu/03q1/ppc970/ppc970-0.html

[14] IBM.com - PowerPC Device Quick Look
http://www-03.ibm.com/technology/power/newsletter/august2004/files/powerpc_device_quick_look.pdf

[15] IBM.com - Understanding 64-bit PowerPC architecture*http://www-128.ibm.com/developerworks/library/pa-microdesign/index.html*

[16] IBM.com - 970FX Processor User Manual
http://www-306.ibm.com/chips/techlib/techlib.nsf/techdocs/AE818B5D1DBB02EC87256DDE00007821/$file/970FX_user_manual_v1.4.pdf