

64-bit CPUs: UltraSPARC-III vs. Intel IA-64

Celeste Pinto

*Departamento de Informática, Universidade do Minho
4710 - 057 Braga, Portugal
celpinto@netcabo.pt*

Abstract. This communication sets an evaluation framework and performs a comparative analysis of the UltraSPARC-III and the IA-64 processor architectures. It starts with an overview of a pure RISC architecture and the new EPIC technology. It then presents the unique combination of innovative features that both manufacturers claim - explicit parallelism, predication and speculative loading – and the goals for their architecture design: to be highly scalable to fill the ever-increasing performance requirements of various server and workstation market segments.

1 Introduction

The UltraSPARC-III (US-III) is the third generation from the UltraSPARC family of Sun and, as one of the last RISC processor families, with full 64-bit precision and addressing range. The processor uses the 64-bit SPARC-V9 architecture and provides the basic hardware foundation on which to build highly effective and truly open solutions to real computing problems [1].

The Itanium, also known as IA-64, is Intel's extension into a 64-bit architecture. The IA-64 processor is based on a set of architectural concepts known as EPIC (Explicitly Parallel Instruction Computing) which is a collection of techniques and an overall design philosophy [2].

The purpose of this paper assignment is to describe with detail the most important of US-III, comparing them with the Intel's Itanium processor at a relatively broad level. It first introduces a summary of each architecture, then it presents a framework comparison and finally it evaluates each architecture on the basis of that framework. Is intended, this way, to confront the features of each processor considering their architectures, RISC and EPIC.

2 Architecture Overview

The US-III processor implements the 64-bit SPARC-V9 RISC architecture. It can sustain the execution of up to four instructions per cycle, even in the presence of conditional branches and cache misses, mainly because the units asynchronously feed instructions and data to the rest of the pipeline. Instructions that are predicted to be executed are issued in program order to multiple functional units, executed in parallel, and for added parallelism can be completed out-of-order. To further increase the number of instructions executed per cycle, instructions from two basic blocks can be issued in the same group.

US-III design includes 14 stage non-stalling pipeline. The pipeline combined with the high system bandwidth allows each core unit to operate at maximum efficiency while retaining a low latency throughput. The processor offer high clock speeds as well as wide superscalar issue to exploit instruction-level parallelism.

The chip supports a 64-bit virtual address space and a 43-bit physical address space. The core instruction set has been extended to include graphics instructions that provide the most common operations related to two-dimensional image processing, three-dimensional graphics and image compression algorithms, and parallel operations on pixel data with 8- and 16-bit components.

The processor offers large Level-1 instruction and data caches, large flexible memory management units, and support for large external cache.

IA-64 is a new advancement enabling processors to break through the limitations of RISC. The EPIC architecture places the emphasis on the ordering of instructions and inclusion from the compiler of parallelism hints. It gives compilers the ability to take advantage of parallel execution. It includes predication and speculation techniques, which address performance losses due to control flow and memory latency, large register files, a register stack and advanced branch architecture. The processor as a 10-stage pipeline, nine execution units, grouped as two integer units, two combination integer-and-load store units, two floating-point units and three branch units.

3 An Evaluation Framework

An evaluation framework introduces the basis features that will be analysed for the processor comparative evaluation:

- **Instruction Set (IS):** the IS reflects the range of commands a CPU can efficiently execute, matching the application specific requirements to the functional units available in the CPU.
- **Reliability:** refers to the capacity to detect and to treat the errors, in a way to increase the performance of the processors. This capacity corresponds to advanced error detection, correction and error removal system.
- **Pipeline:** structured in several stages to increase the CPU throughput. Each pipeline stage performs part of the work necessary to execute an instruction.
- **Branch Handling:** branches are instructions which can alter the flow of instruction execution in a program. Branch prediction aims to reduce the penalty of branches in deep pipelines. Two techniques for handling branches will be presented, such as branch delay slots (relies on programmer/compiler to fill instructions, which depends on being able to find suitable instructions and link resolution delay to a particular pipeline) and predication (conditional execution of instructions). There are two branch prediction strategies such as static (decided before runtime) and dynamic (prediction decisions may change during the execution of the program). Dynamic branch predictors reduce execution time predictability by introducing a history state in the processor.
- **Instruction Level Parallelism (ILP):** ILP allows multiple non-dependent operations to be simultaneously executed and the key performance metric in all ILP processor classes is IPC (instructions per cycle), this is the degree of parallelism achieved.
- **Memory Hierarchy:** it includes main memory and cache memory. A cache is a subset of main memory, used to make memory access faster. There are several cache levels - which can be located, or not, in the same chip – and ways to organize and access data (direct mode, fully associative or n-way set associative). Other relevant aspect in memory performance is memory latency that defines the time between a CPU request for data and the data actually arriving.

4 Architecture Comparison

4.1 Instruction Set

The US-III instruction set has 3 possible instruction formats. It also has 74 possible instructions including the floating-point operations. There is a tagged word format in which the 2 least significant bits serve as flags to indicate the type of object.

The floating-point numbers can be 32 (single), 64 (double), or 128 (quad) bits long. The quad format uses a 112-bit mantissa for applications requiring incredible floating-point precision. The floating-point unit has 32 32-bit non-windowed registers, which must be saved on a per-context basis [6].

US-III has visual and media processing features with Visual Instruction Set (VIS). VIS combines several standard multimedia functions into a set of single instruction sequences. The 64-bit super-scalar nature of the US-III allows the core to process multiple graphic-intensive instructions per global clock cycle. VIS addresses performance concerns by dividing standard floating-point graphic operations across the entire scope of the integer and floating-point execution units. VIS operations rely on integer registers for loading and storing data, while floating-point registers provide power for data manipulation [8].

The VIS uses a series of specialized commands to complete normally CPU intensive graphic operations. Also encoded into the VIS is support for real-time video compression and decompression. The US-III can support various compression formats in hardware, including MPEG-I/II video and JPEG still imaging [5].

Looking at the IA-64 instruction set, its concept of "instruction" is a little different from the US-III. IA-64 instructions are of a fixed length of 128 bit brought together in so called bundles. Each of those two cycle bundles contains three instructions and has a template field which contains details of the bundles instructions dependencies and more information that is needed to know whether, or not, the three instructions in that bundle and subsequent instructions can be executed in parallel. The advantage of bundling instructions this way is that the CPU is explicitly told about possible parallelism. It is the compiler problem to detect and make use of parallelism.

All IA-64 instructions fall into one of four categories: integer, load/store, floating-point, and branch operations [9].

4.2 Reliability

US-III performs advanced error detection and correction system. The L2 cache offers full ECC (error checking and correction) capabilities. The US-III architecture also offers data parity bit protection with the ability to detect and correct single bit errors. An automatic recovery system minimizes the impact of errors, as it can isolate and resolve several common software problems without hard-crashing the system. An error removal system also aids in preventing the propagation of copy-back errors within an SMP configuration [3]. The US-III provides an "up-time bus" consisting of an 8-bit data path for diagnosing system bus errors.

Relatively to the IA-64, Machine Check Architecture (MCA) and Error Correcting Code (ECC) help provide advanced error detection, containment, and correction as well as error logging [4].

4.3 Pipeline

To achieve high clock rates, the new US-III processor execution pipeline is segmented into 14 separate stages, as figure 1 shows.

The instruction issue unit occupies the A through J stages of the pipeline, and the integer execution unit accounts for the R through D stages. The data cache unit occupies the E, C, M, and W stages of the pipe in parallel with integer execution unit stages. The floating-point unit is shown as a side pipeline that parallels the E through D stages of the integer pipeline. The other units of the machine (system interface unit and external memory unit) have internal pipelines but are not considered part of the core processor pipe [1].

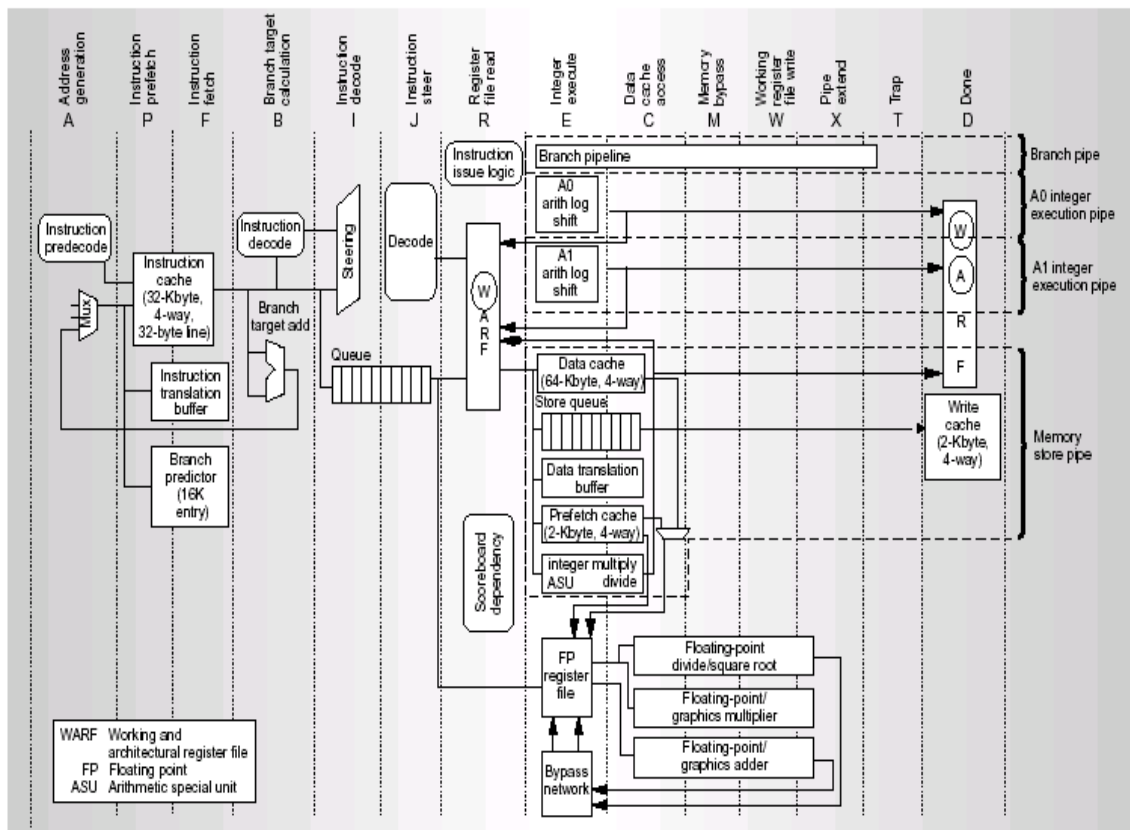


Fig. 1 - UltraSPARC-III instruction pipeline

The IA-64 processor has a 10-stage in-order pipeline with cycle time designed for single cycle ALU (4 ALUs globally bypassed) and low latency from data cache. This processor has a parallel, deep, and dynamic pipeline designed for maximum throughput [4].

4.4 Branch Handling

Branch prediction is fully static in the US-III and use branch delay slots technique for handling branches [3]. Implements a slightly modified Gshare algorithm with 16K saturating 2-bit counters (the three low order index bits into predictor use PC info only) and has 8 cycle misprediction delay to drain stages. While the branch target instructions are being fetched, a branch prediction will be taken, the “fall-through” instructions are

readied for issue in parallel, using a special 4-entry branch miss queue. If branch taken prediction turns out to be wrong, this precaution reduces the penalty for error to just 3 cycles. If the prediction that a branch will be not-taken fail, it must pay a penalty of 7 cycles. When the two misprediction cases are averaged together (weighted by relative frequency), the overall branch misprediction penalty is lowered to about 4.5 cycles [6].

IA-64 uses dynamic branch prediction and use predication to reduce the number of branches in the code. This improve instruction fetching because there are fewer control flow changes, decreases the number of branch mispredict since there are fewer branches, out increase the branch prediction hit since there are less competition for prediction resources; providing software hints for branches to improve hardware use of prediction and prefetching resources; supplying explicit support for software pipelining of loops and exist prediction of counted loops [7].

The major difference from the US-III is that branch triggers are controlled by predicates rather than conditions encode in branch instructions.

4.5 Instruction Level Parallelism

The US-III processor maintains a sustainable issue rate of 4 instructions every clock cycle (the maximum number of instructions that can be fetched from cache in a clock cycle). A total of 16 fetched instructions can be queued up, waiting for an appropriate execution unit to become available. There are six parallel execution units: 2 integer ALUs (both identical), 1 branch unit, 1 load/store unit, and 2 floating-point units (1 for add/subtract, 1 for multiply/divide operations). Since the US-III contains independent execution units for different operations, instructions can be executed concurrently. With this parallel execution technique, a programmer can achieve true ILP [5]. This solution works, but it requires the CPU to do a lot of work.

IA-64 uses a 128-bit encoding, called a bundle, which contains three 41-bit instructions and a 5-bit template field. The template bits help decode and route instructions and indicate the location of stops that mark the end of groups of instructions that can execute in parallel. The process of sending instructions to functional units is called dispersal. Thus the IA-64 can issue a maximum of six instructions per clock cycle. The specific functional unit to which an instruction is sent is determined by its type and its position within the current set of instructions being issued [2].

The main advantage here is that the hardware does not have to spend effort figuring out which instructions go with each other; the compiler does it. One of the problems with this approach is that it can lead to significant code bloat.

Both US-III and IA-64 processors operate basically in parallel, making use of a number of concurrently working execution units.

4.6 Memory Hierarchy

The memory hierarchy is organized in main memory and two cache units. The first level (L1) consists of four separate caches, two large and two small. The second level (L2) consists of a single large cache. The two large L1 caches hold instructions and data, respectively. The two small L1 caches are a prefetch cache, mainly used by speculatively executed load instructions to hold floating-point data, and a write cache, that serves to greatly reduce the amount of storage bandwidth required. All four L1 caches are on-chip. The large L2 cache is a unified instruction and data cache. The contents of the L2 cache are stored off-chip (in SRAM), but the L2 address tags are kept on-chip for faster access.

Figure 2 shows the memory organization of an US-III processor system and the associated performance strategy [5].

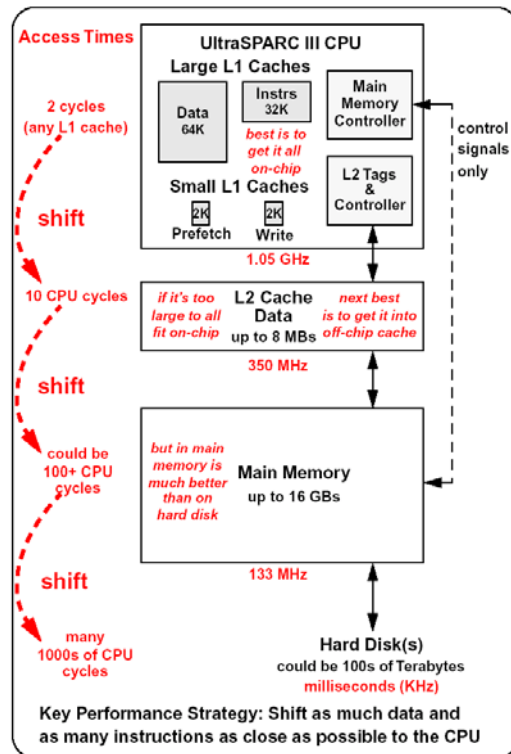


Fig. 2 - UltraSPARC-III Memory Hierarchy

The US-III features 100KB of on-die L1 cache. The 64KB of data caching combined with 32KB of instruction caching would provide enough bandwidth to steadily supply the execution stage of the CPU with a consistent flow of information.

The pre-fetch and write caches work in conjunction with memory and register addressing to buffer incoming and outgoing core data within a high-speed, low latency cache memory area. By utilizing these advanced caching areas, the US-III offers 9.6GB/s of cache bandwidth between each CPU. To further aid in bandwidth, the full-speed L1 cache uses a 4-way set associative architecture. In simple terms, this means the cache will buffer data at four different points within the memory array.

Since data can be found at four different memory locations, the cache controller does not need to expend CPU overhead for address searching the entire cache region before the needed information can be found. While a 1-way associative architecture provides the best performance, a 4-way architecture provides more effective bandwidth with lower clock latency.

The US-III L2 cache is incorporated in an off-die module, which limits the operational clock speed to a fraction of the CPU's core rate. To offset the clock differential, the US-III design incorporates a wide 256-bit data path to the L2. This pipe allows increased bandwidth per global CPU clock cycle [6].

The US-III integrates the system memory controller directly into the CPU's core. Since the controller operates at the core clock rate, this design allows for a peak of approximately 3.2GB/s of CPU-to-memory bandwidth with standard SDRAM. However, the US-III does not incorporate a standard CPU bus. In place of the standard chipset controller, the US-III uses a series of intelligent, high-speed data switches to correctly

correlate and direct the system's data flow. The data switch architecture offers 2.4GB/s (peak) of system level bandwidth.

The IA-64 has three cache levels and a large registry set. The L1 cache is 16 KB, 4-way set-associative and has 2-cycle latency. L2 cache is 96 KB, 6-way set-associative and has 6 or 9-cycle latency for INT or FP. FP loads read from L2 cache instead of L1 because typically have large workloads. L3 cache is 4 MB, off-chip, on-package, 4-ways, and has 21 or 24-cycle (INT or FP) latencies [7]. The IA-64 registry files have a wide range of register, among others, 128 for integers, 128 for floating point, 8 branch registers. All of them are programmer visible and random accessible. This bigger number of registers means fewer loads, stores and recalculations. In addition there are 64 1-bit predicate registers for use with predication [2].

As much the US-III as the IA-64, both use speculative execution of loads and prefetching techniques to reduce memory latency.

5 Future Trends

Sun has been working on an architecture called MAJC or "Microprocessor Architecture for Java Computing" (it is not a Java chip though, simply designed with the expectation of mostly running Java programs on it) which is designed specifically with multiple CPU cores in one die (up to 1024) as well as hardware support for threading. The successor to the US-III will be released in early 2003. The UltraSPARC IV (US-4) will be offered at speeds beginning at 1.2GHz, the maximum speed of the US-III, and will eventually graduate to 2GHz. The chip will include the capacity for instruction-level parallelism and thread-level parallelism, which will enable the US-4 to combine the functions of rival Intel chip into one single processor that can modify itself between the demand for heavy computing or more simple everyday business functions. Sun's product line road map also revealed that the UltraSPARC-V will be released in the early part of 2005 and will first be offered at 1.8GHz speeds with the capacity to expand to 3.0GHz or higher. Sun will also develop mid-range UltraSPARC derivatives that will include the UltraSPARC-IIIi at 1.1GHz and growing to 1.6GHz, and the UltraSPARC-Vi, with speeds of 1.1GHz to 2GHz [3].

The second IA-64 processor after Itanium is code-named McKinley and it is likely to be faster, smaller, and all-around better than its predecessor. McKinley's L1 caches will be the same size as IA-64, but the L2 cache will grow from 96K to 256K. The L3 cache will get smaller (3M instead of 4M) but move onto the actual chip, not just on the same cartridge. All three cache interfaces will get faster. McKinley shaves one cycle off the L1 cache access time (from two cycles to one), shortens L2 access time by seven cycles (to five), and takes eight cycles off the L3 latency (to 12 cycles). McKinley's system bus will widen to 128 bits and its clock frequency will improve from 133 MHz to 200 MHz. The bus will still be double-pumped (i.e., transferring data on both rising and falling edges of every clock) yielding 6.4GB/sec front-side bus bandwidth. Next up comes Madison, expected to be a 0.13-micron shrink of McKinley, all other things being equal. Deerfield, the fourth member of IA-64's growing family, will also be a 0.13-micron shrink of McKinley, but this time with a smaller 1M L3 cache and yet another new bus interface intended for cheaper systems [9].

6 Conclusions

Both processors aim mainly for server and workstation applications, focusing on improved capabilities and speed. The US-III processor holds advantages in handling branches. Integer operations also go in favour of the US-III because of its faster clock speed, low instruction latency, and adequate number of execution units. In terms of L1, US-III uses 64KB of data caching combined with 32KB of instruction caching while IA-64 has on the Level 1 a cache size of only 16KB. In terms of L2, US-III incorporates 8Mb of L2 cache while the IA-64 is only 96KB in size. Through explicit parallelism, the IA-64 implementation of EPIC technology enables the compiler and hardware to work together efficiently to expose new levels of parallelism and performance. The innovative use of predication and speculation, uniquely combined with explicit parallelism, has allowed EPIC technology to progress well beyond the limitations of traditional architectures. However, significant similarities still exist, both processors use speculative loading to help to overcome memory latency problems but the way the two chips deal with exception errors differs slightly.

References

- [1] Horel, Tim, Lauterbach, Gary, ULTRASPARC-III: Designing Third-Generation 64-Bit Performance, IEEE Micro, (May - June 1999)
- [2] Intel IA-64 Architecture – Software Developer's Manual, Vol. I, Intel, Santa Clara, California, (July 2000)
- [3] Rijk, Chris, Inside Tech: The UltraSPARC III , (February 2000)
- [4] Sharangpani, Harsh, Intel Itanium Processor Microarchitecture Overview, Intel Corporation, (October 1999).
- [5] Richmond, Robert, Sun UltraSPARC-III Platform, Sun Microsystems (October 2000)
- [6] An Overview of the UltraSPARC®III Cu Processor, The UltraSPARC III Processor Moves to Copper Technology, Sun Microsystems, (June 2002)
- [7] Inside the Intel Itanium 2 Processor: an Itanium Processor Family member for balanced performance over a wide range of applications, Hewlett Packard Technical White Paper, (July 2002)
- [8] The VIS™ Instruction Set, Sun Microsystems, Santa Clara, California (June 2002)
- [9] Turley, Jim, Evaluating High-End Microprocessors: 64-Bit Microprocessors, (September 2002)