




Massive Parallelism for Mission-Critical Applications

Advanced Explicitly Parallel Instruction Computing
(EPIC) Architecture



Steve Undy
Intel Corporation



Intel® Itanium® 9500³ processor series, codenamed Poulson, is the latest Intel Itanium processor in a long line of ground breaking designs. Optimized for Explicitly Parallel Instruction Computing (EPIC) principles, Intel Itanium processor 9500 series' advanced EPIC Architecture can be best summarized as exploiting parallelism on all levels: pipeline, core, thread, memory, pipeline and instructions. End-users are now able to extract more inherent parallelism in their code than ever before to deliver a new level of performance, while benefiting from the mainframe-class RAS features to deliver an always-on experience in their mission-critical enterprise.



Adopting an Advanced EPIC Architecture

Intel Itanium processor 9500 series represents a near clean-sheet redesign of the Intel Itanium cores to support an unprecedented amount of instruction-level parallelism in its main execution pipeline. It can execute up to 12 instructions each cycle in 4 instruction bundles. It has 2 memory execution units, 2 general purpose integer units, 2 ALU units, 2 floating point units, 3 branch units and 1 NOP unit. The Intel Itanium bundle template determines which units are candidates for executing each instruction. The hardware algorithm used to disperse the incoming instructions into each of the 12 execution unit pipelines is simple, deterministic and efficient – allowing compilers to exactly control execution resources. To support 12-wide issue, the register files have 12 read and 12 write ports.

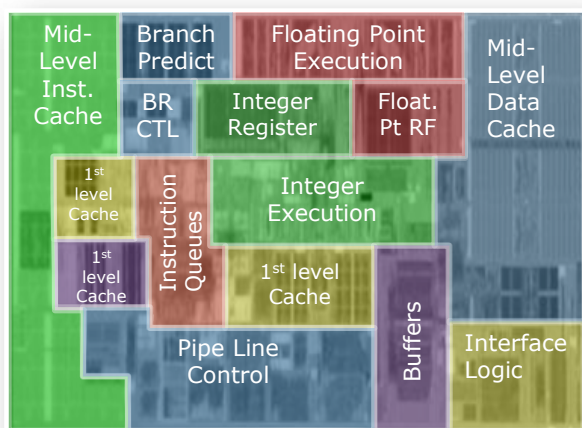


Figure 1 Intel Itanium processor 9500 series core floorplan.
New microarchitecture features an 11-stage pipeline and architectural extensions.

The refreshed microarchitecture also allowed a focus on power efficiency. The power aware design of Intel Itanium processor 9500 series was essential to being able to double core count and operating frequency while simultaneously reducing maximum package power to achieve a factor of three power efficiency advantage over the previous Itanium processor design.

Intel Itanium processor New-Instructions Architectural Extensions

Intel Itanium processor 9500 series adds a set of new instructions that extends the Itanium architecture. It adds integer multiply instructions and a count-leading-zero instruction. It adds an instruction to provide better OS control of thread behavior. It adds and extends instructions that provide more detailed data access hints as well as new user-controlled register file to control those hints. This allows compilers much finer grained control of data cache and TLB policies. It also adds an instruction for multi-line software prefetches. All of these new instructions are motivated by the desire to increase performance, both single-thread and multi-thread.

Memory Parallelism

Intel Itanium processor 9500 series also focuses on increasing memory parallelism by addressing throughput and queuing in the memory subsystem. The core has additional queuing for pending memory operations tweaked for throughput. Queue sizes were increased and the scheduler was changed to focus on performance and power.

Another key improvement on memory parallelism is the ability to avoid pipeline hazards by executing data prefetch operations to move data in advance of use between the various levels of caches. By providing extra hooks to the compilers to control caching policies in addition to the software and hardware prefetchers in the memory pipeline, Intel Itanium processor 9500 series can control explicit data and control speculation mechanisms, and enable its prefetchers to use an adaptive algorithms to conserve bandwidth as much as possible to help relieve potential pipeline bottlenecks.

Core Parallelism

Probably the most obvious form of parallelism Intel Itanium processor 9500 series supports is core-level parallelism. The processor has eight cores per socket connected to eight 4MB last level cache modules via a ring interconnect. The ring interconnect is capable of 700 GB/s of aggregate bandwidth. The ring caches are connected using QPI protocols to the two on-die memory controllers and a ten port router. The router


WHAT IS EPIC?

EPIC or Explicitly Parallel Instruction Computing, represents a paradigm shift in the development of instruction set architectures. Instead of placing the main burden of extracting parallelism and performance on the underlying computing hardware, a synergy is developed between the software ecosystem and the hardware implementation. This allows compilers - which have full access to the program source code - and the processors - which have full access to run-time information as a program executes - to each be optimized for what each does best. In order to do this, the instruction set provides a rich set of features for software to optimally control the low-level hardware resources. This most notably includes the ability for compilers to specify, schedule and exploit the many forms of parallelism inherent in user programs.

ports connect to six QPI interfaces to reach other, external, processor sockets and devices. Each processor is capable of 128 GB/s bandwidth between sockets and 45 GB/S per second of bandwidth to the memory modules.

Thread Parallelism

Like its predecessors, Intel Itanium processor 9500 series processors are multi-threaded and support the Intel® Hyperthreading Technology¹. There has been significant advancement in Intel Itanium processor 9500 series to add new multi-threading optimizations, including the dual-domain multithreading capability. In the dual domain, the front-end and main pipelines are independently threaded. Each pipeline uses independent and separate algorithms to switch between threads. Many of the core structures have been split into separate per-thread resources. This includes the instruction buffer, the data TLBs, and the hardware page walker. The result is that disparate instruction threads can now run on different parts of the pipeline, to improve performance even on legacy software without requiring costly recompilation and application re-qualification.



Front-end handles instruction fetch, branch prediction, register renaming, and sorting 2 bundles per cycle into instruction buffer

Back-end handles read from instruction buffer at 4 bundles per cycle, execute instructions, and accessing data cache

Figure 2 Dual domain multi-threading. Intel Itanium processor 9500 series enables independent front-end and back-end pipeline operations.

Pipeline Parallelism

Intel Itanium processor 9500 series incorporates several major pipelines that operate independently from one other and are separated by decoupling buffers, representing another major form of parallelism. The front-end pipeline fetches instructions, performs branch prediction, partially decodes instructions and renames registers. After flowing down the front-end pipeline, 6 instructions per cycle are placed into a 192-entry instruction buffer. The instruction buffer is divided into 6 logical queues corresponding to execution unit type. The main pipeline reads instructions out of this buffer and executes them in the 12 execution units described above. This allows the core to run at increased frequencies as well as providing the ability to provide hardware recovery of errors.

Instruction Level Parallelism (ILP)

The EPIC philosophy in general and Intel Itanium are specifically designed to optimize the location of the interface between software and hardware. In addition to the fine-grained control compilers have, the new Intel Itanium processor 9500 series architecture provides hint instructions to communicate additional information from software to hardware. Branch prediction instructions are used to control branch prediction hardware and instruction prefetching. Data prefetch instructions are used to generate explicit prefetches of data prior to use. Data instructions can encode hint completers to indicate speculation and locality information.

Further, the Intel Itanium processor architecture defines a rich set of resources for performance monitoring. This is a powerful framework to allow detailed measurement and profiling of hardware behavior. In essence, this forms a mechanism to “close the loop” by allowing hardware to communicate back to the compilers and application developers.

Conclusion

Intel Itanium 9500 Processor's advanced EPIC architecture increases parallelism at all levels, and provides a robust hardware platform with a rich set of architectural extension to support an unprecedented level in instruction throughput. Given the mission-critical space of the systems that Intel Itanium processors are targeted for, performance gains also cannot be made at the expense of reliability, availability and serviceability. In fact, many of the same microarchitecture features that allow better performance and power efficiency also lead to better error recoverability. With Intel Itanium processor 9500 series, customers can realize a higher level of performance and utilization for their most demanding business workloads, while maintaining the world-class mission critical capabilities that is the signature of the Intel Itanium product line.

For more information on
Intel® Itanium® processor 9500 series,
visit www.intel.com/itanium

INTEL® ITANIUM® PROCESSOR 9500 SERIES

Intel Itanium processor 9500 series contains 8 redesigned cores, 32MB of last-level cache, 4 full-width and 2 half-width Quick Path Interconnect (QPI) links, 2 directory caches and 2 on-die memory controllers connecting to external memory modules using 4 SMI links. Containing 3.1 billion transistors fabricated in a 32 nm process, it runs at a maximum thermal power of 170W and supports the construction of 8 socket “glueless” systems as well as much larger configurations.

Intel Itanium processor 9500 series provides greater than 2-fold increase of performance over the previous Intel Itanium processor, and offer new key performance and reliability features:

- Intel Itanium processor New-Instructions²
- Intel Hyper-Threading Technology¹, enhanced with Dual-domain multithreading
- Intel Instruction Replay Technology²

¹ Hyper-Threading Technology requires a computer system with an Intel® processor supporting HT Technology and a HT Technology enabled chipset, BIOS and operating system. Performance will vary depending on the specific hardware and software you use. See http://developer.intel.com/products/ht/Hyperthreading_more.htm for additional information.

² Some features are supported fully in silicon and are performed automatically and transparently. Others require additional support from the firmware, platform or OS and may not be supported in all systems.

³ Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See http://www.intel.com/products/processor_number for details.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit Intel Performance Benchmark Limitations at <http://www.intel.com/performance/resources/limits.htm>

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked “reserved” or “undefined.” Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.


The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web site at www.intel.com.

Copyright © 2012 Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

*Other names and brands may be claimed as the property of others.

Printed in USA

1012/LYC/DCSG/XX/PDF

 Please Recycle

328102-001US

