



White Paper
Intel® Core™ Microarchitecture

Introducing the 45nm Next-Generation Intel® Core™ Microarchitecture

Contents

New Innovations and Enhancements Deliver Higher Performance and Energy Efficiency.....	3
Intel® Core™ Microarchitecture	3
Intel’s 45nm High-k Metal Gate Process Technology	4
Penryn, the Next-Generation Intel® Core™2 Processor Family	4
Making Software Run Faster	4
New Intel® SSE4 Instructions	4
Larger, Enhanced Intel® Advanced Smart Cache	5
Higher Speed Cores and System Interface	5
Enhanced Intel® Virtualization Technology	5
Super Shuffle Engine.....	5
Fast Radix-16 Divider	5
Store Forwarding	6
Improved Operating System (OS) Synchronization Primitive Performance	6
Improving Energy Efficiency	6
Deep Power Down Technology	6
Enhanced Intel® Dynamic Acceleration Technology	6
Coming in 2008: Intel’s Next-Generation Microarchitecture	7
The Beat Goes on With 32nm Silicon Process Technology	7

New Innovations and Enhancements Deliver Higher Performance and Energy Efficiency.

In the second half of 2007, Intel will begin production of the next-generation Intel® Core™2 processor family codenamed "Penryn." The Penryn processor family is based on our industry-leading 45-nanometer (nm) High-k metal gate silicon technology and our latest microarchitecture enhancements. This next evolution in Intel® Core™ microarchitecture builds on the tremendous success of our revolutionary microarchitecture (currently used in both the Intel® Xeon® and Intel® Core™2 processor families) and marks the next step in Intel's rapid cadence for delivering a new process technology with enhanced microarchitecture or an entirely new microarchitecture every year.

With more than 400 million transistors for dual-core processors and more than 800 million for quad-core, the 45nm Penryn family introduces new microarchitecture features for greater performance at a given frequency, up to 50-percent larger L2 caches, and expanded power management capabilities for new levels of energy efficiency. The Penryn family also includes nearly 50 new Intel® SSE4 instructions for speeding up the performance of media and high-performance computing applications.

The Penryn family will feature new dual-core desktop processors, quad-core desktop processors, quad-core server processors and dual-core mobile processors.

Intel® Core™ Microarchitecture

Intel first introduced Intel Core microarchitecture in 2006 in the Intel Core 2 processor family manufactured with our 65nm silicon process technology. The first generation of this multi-core optimized microarchitecture extended the energy-efficient philosophy first delivered in the mobile microarchitecture of the Intel® Pentium® M processor and enhanced it with many new, leading-edge microarchitecture innovations for industry-leading performance, greater energy efficiency and more responsive multitasking.

Intel Core microarchitecture innovations include:

- Intel® Wide Dynamic Execution
- Intel® Intelligent Power Capability
- Intel® Advanced Smart Cache
- Intel® Smart Memory Access
- Intel® Advanced Digital Media Boost

Processors based on Intel Core microarchitecture have delivered record-setting performance on leading industry benchmarks for desktop, mobile and mainstream server platforms. (See www.intel.com/performance.) For instance, 65nm Quad-Core Intel® Xeon® processors provide 2.5x the performance of previous server solutions.¹ On the desktop, Intel Core 2 Duo processor-based systems provide up to 40 percent more performance with lower energy consumption.² And on the go, Intel Core 2 Duo mobile processor-based laptops provide up to 2x the performance in multitasking, as well as greater energy efficiency to enable longer battery life.³

Intel's 45nm High-k Metal Gate Silicon Technology

In January 2007, Intel introduced one of the biggest advancements in fundamental transistor design in 40 years – the use of dramatically different transistor materials (a new material combination of hafnium-based high-k (Hi-k) gate dielectrics and new metal materials for the gate) to build the hundreds of millions of microscopic 45nm transistors inside the next generation of the company's Intel Core 2 processor family. By using the element Hafnium, a metal that significantly reduces electrical leakage and provides the high capacitance necessary for good transistor performance, Intel can continue its record-breaking PC, laptop and server processor performance while reducing the amount of electrical leakage from transistors that can hamper chip and PC design, size, power consumption and costs.

Intel's new silicon process formula will increase transistor switching speeds, enabling higher core and bus clock frequencies, and thus more performance in the same power and thermal envelope. This, in turn, will help extend Moore's Law (a high-tech industry axiom that transistor counts double about every two years to deliver ever more functionality at exponentially decreasing cost) well into the next decade.

Compared to 65nm technology, Intel's hafnium-based 45nm Hi-k silicon process technology will provide the following product benefits:

- Approximately twice the transistor density (for smaller chip sizes or increased transistor counts)
- Approximately 30-percent reduction in transistor-switching power
- Greater than 20-percent improvement in transistor-switching speed or a greater than 5 times reduction in source-drain leakage power
- Greater than 10 times reduction in transistor gate oxide leakage for lower power requirements and increased battery life

Intel's January 2007 demonstration of the world's first 45nm Hi-k processor underscored its process technology lead of more than two years over the rest of the semiconductor industry. According to Intel co-founder Gordon Moore, "The implementation of high-k and metal materials marks the biggest change in transistor technology since the introduction of polysilicon gate MOS transistors in the late 1960s."

Penryn, the Next-Generation Intel® Core™ 2 Processor Family

Penryn, the first family of processors based on Intel's new 45nm Hi-k silicon technology, makes good use of the additional transistors this technology can pack into a chip. This 45nm Hi-k next-generation Intel Core 2 and Intel Xeon processor families will deliver many new architectural features and advancements to make software run faster and improve energy efficiency.

Making Software Run Faster

The Penryn family includes an extensive array of microarchitecture improvements that improve performance across a broad range of software.

New Intel® SSE4 Instructions

The Penryn family includes the Intel® Streaming SIMD Extensions 4 (SSE4) instructions. Intel SSE4 instructions are the most significant media instruction set architecture advancement since 2001. This new instruction set extends the Intel® 64 architecture® instruction set architecture to better take advantage of Intel's next-generation 45 nm silicon manufacturing process and expand the performance and capabilities of Intel® Architecture. Intel SSE4 instructions deliver further performance gains for SIMD (single instruction, multiple data) software and will enable Penryn microprocessors to deliver superior performance and energy efficiency to a broad range of 32- and 64-bit software. Applications that will benefit include those involving graphics, video encoding and processing, 3-D imaging, and gaming. The instructions will also benefit high-performance applications like audio, image and data compression algorithms, as well as many more.

The Penryn family's implementation of Intel SSE4 will improve performance by:

- Adding support for two different vectored 32-bit integer multiply operations
- Introducing 8-bit unsigned min/max operations, plus 16-bit and 32-bit signed and unsigned versions

- Introducing features to improve the compiler's ability to vectorize integer and single-precision code more efficiently
 - Blends, Tests and Rounds, and sign/zero extensions, are straightforward replacements for existing lengthy operations
 - Inserts, Extracts are building blocks to gathers (lookups), scatters, strided loads, and stride stores
- Adding highly specialized operations that can provide significant application level gains in:
 - Video encode acceleration functions
 - Floating-point dot product operation (important in gaming and 3D content creation)
 - Streaming load instruction (important for video processing, imaging, and applications that share data between the graphics processor and processor)

The performance gains are dramatic. For instance, the Intel SSE4 streaming load instruction improves the bandwidth for reading data from a graphics frame buffer. By fetching a full cache line (64 bytes at a time as opposed to 8 bytes and keeping it in a temporary buffer), this instruction can enable an up to 8X theoretical improvement in read bandwidth.

Larger, Enhanced Intel® Advanced Smart Cache

Penryn processors include a 50-percent larger L2 cache with a 24-way associativity to further improve the hit rate and maximize utilization. Dual-core Penryn processors will feature up to a 6 MB L2 cache and quad-core processors up to a 12 MB L2 cache. These large caches improve performance and efficiency by increasing the probability that each execution core can access data from a higher performance, more efficient cache subsystem.

Penryn family caches also contain an *enhanced cache line split loads capability*. A split load occurs when a data value is read and part of the data is located in one cache line and part in another. Reading data from two cache lines is several times slower than reading data from a single cache line even if the data is not otherwise properly aligned. The Penryn family's enhanced cache line split loads greatly improves performance by speculatively dispatching both halves of a split load potentially ahead of other loads or stores. This can speed up the performance of certain applications that perform data scans, such as video motion estimation.

Higher Speed Cores and System Interface

Penryn family processors will run at higher core speeds (greater than 3 GHz for some versions) than previous Intel Core 2 processor families. Front-side bus speeds will be increased up to 1.600 GHz, in addition to the 1.066 GHz and 1.333 GHz now available. This will improve overall system performance.

Enhanced Intel® Virtualization Technology^A

Penryn speeds up virtual machine transition (entry/exit) times by an average of 25 to 75 percent. This is all done through microarchitecture improvements and requires no virtual machine software changes. (Virtualization partitions a computer so that it can run separate operating systems and software in each partition, better leveraging multi-core processing power, increasing efficiency, and cutting costs by enabling a single machine to act as many virtual computers.)

Super Shuffle Engine

Implementing a full-width, single-pass shuffle unit that is 128-bits wide, Penryn processors can perform full-width shuffles in a single cycle. This doubles the speed for most byte, word, or dword SSE data shuffle operations and significantly reduces latency and throughput for SSE2, SSE3 and Intel SSE4 instructions that have shuffle-like operations like pack, unpack and wider packed shifts. This capability will provide a general performance improvement in a broad range of SSE algorithms.

Fast Radix-16 Divider

Penryn processors provide faster divide performance, roughly doubling the divider speed over previous generations for scientific computations, 3D transformations, and other mathematical-intensive functions. The inclusion of a new, fast divide technique called radix 16 speeds division in both floating-point and integer operations. (A radix 4 algorithm computes 2 bits of quotient in every iteration. Increasing to a radix 16 algorithm allows for computing 4 bits in every iteration for a 2X reduction in latency.)

Store Forwarding

To speed up the reading of the result of a “misaligned” store that crosses an 8-byte address boundary and is still in a pipeline, Penryn processors can forward the result of the store to the load immediately rather than waiting for the store to finish and write to memory.

Improved Operating System (OS) Synchronization

Primitive Performance

Certain OSs temporarily block out or “mask” interrupts when starting a critical section of code and needing exclusive access over a resource such as an I/O device. Through faster clear interrupts/set interrupts (CLI/STI) capability, Penryn processors can move into and out of this mode faster, significantly improving performance. What’s more, they can execute locked instructions (such as XCHG, ADD/XADD/NEG/BTS/AND, and CMPXCHG) faster. Penryn processors also feature a faster access of the time stamp counter (read time stamp counter or RDTSC). This can be a frequently invoked function in database or transaction processing-based server workloads.

Improving Energy Efficiency

In addition to Intel 45nm Hi-k silicon technology benefits, the Penryn family builds on the energy-efficiency capabilities of the Intel Core microarchitecture with two important additions: Deep Power Down Technology and Intel® Dynamic Acceleration Technology.

Deep Power Down Technology

This is a radically new and advanced power management state (C-state) that significantly reduces the power of the processor during idle periods so internal transistor power leakage is no longer a factor. This latest processor “sleep” state is the lowest power state a processor can reach and significantly helps extend laptop battery life. It enables Penryn to achieve up to a substantial improvement over the lowest power state of Merom, the previous-generation Intel Core microarchitecture for mobile platforms.

Upon entering Deep Power Down, the processor flushes cache, saves the processor microarchitecture state internally, and shuts off power to cores and L2 cache. While in Deep Power Down, the chipset continues to service memory traffic for input/output (I/O), but doesn’t wake up the processor. When the core is needed, the voltage is ramped up, the clocks turned on, the processor reset, the microarchitecture state is restored, and instruction execution resumed.

The deeper a C-state, the higher the energy cost of the transition to and from this state. Too frequent of transitions to deep C-states can result in a net energy loss. To prevent this, Penryn includes an auto-demote capability that uses intelligent heuristics to determine when idle period savings justify the energy cost of shutting down a processor and restarting it. If it doesn’t, the Deep Power Down request is demoted to C4, a less deep power management state. The result is a power savings appropriate to the probable idle period.

Enhanced Intel® Dynamic Acceleration Technology

To further increase the performance of single-threaded applications, Intel has enhanced the *Intel® Dynamic Acceleration Technology* available in current Intel Core 2 processor families. This feature uses the power headroom freed up when a core is made inactive to boost the performance of another still active core. (Imagine a shower with two powerful shower heads. When one shower head is turned off, the other has increased water pressure, or performance.) If one core is in C3 or deeper C-state, some of the power normally available to that idle core can be applied to the active core while still staying within the thermal design power specification for the processor. This increases the speed at which single-threaded applications can be processed, thus improving the performance of many applications.

Coming in 2008: Intel's Next-Generation Microarchitecture

Intel's architecture and silicon technology advancements are based on a rapid cadence that delivers an accelerated pace of innovation in driving processor performance and energy efficiency for the next decade and beyond. Intel calls this cadence the "tick-tock" model of silicon and microarchitecture. Each "tick" represents a new silicon process technology with an enhanced microarchitecture. The corresponding "tock" represents the design of a brand new microarchitecture. The cycle repeats approximately every two years.

The Penryn family, with its Intel 45nm Hi-k silicon technology, is the latest "tick" and includes many microarchitecture innovations to Intel Core microarchitecture. Coming in 2008 is the following "tock," Intel's next brand-new microarchitecture codenamed Nehalem.

Nehalem is a truly dynamic and design scalable microarchitecture enabling it to deliver both performance on demand and optimal price/performance/energy efficiency for each type of platform.

Nehalem's dynamic scalability delivers performance on demand through:

- Dynamically managed cores, threads, cache, interfaces, and power
- Leveraging leading 4-instruction issue Intel Core microarchitecture technology (Intel Core microarchitecture's ability to process up to 4 instructions per clock cycle on a sustained basis as compared to 3 instructions per clock cycle or less for other processors)

- Simultaneous multi-threading (Intel Hyper-Threading Technology) to enhance performance and energy efficiency
- Innovative new Intel® SSE4 and ATA instruction set additions
- Superior multi-level shared cache
- Leadership system and memory bandwidth
- Performance-enhanced dynamic power management

Nehalem's design scalability will enable optimal price/performance/energy efficiency for each market segment through:

- New system architecture for next-generation Intel processors and platforms
- Scalable performance for from one-to-sixteen (or more) threads and from one-to-eight (or more) cores
- Scalable and configurable system interconnects and integrated memory controllers
- High-performance integrated graphics engine for client platforms

The Beat Goes on With 32nm Silicon Process Technology

Next up after Nehalem will be processors based on Intel's upcoming 32nm silicon process technology. This next tick in Intel's rapid cadence of both silicon technology and microarchitecture innovation will further sustain Intel product leadership. For our customers, it will mean remarkable performance and efficiency gains, features and capabilities for years to come.

¹Performance measured using SPECint*_rate_base2000 comparing a Quad-Core Intel® Xeon® processor X5355-based platform to a Dual-Core AMD Opteron® processor Model 2220SE-based platform.

²Performance measured Intel® Core™2 Duo desktop processors compared to Intel® Pentium® D processor 805 on SPECint_base2000 and SPECint_rate_base2000 (2 copies). Actual performance may vary. See <http://www.intel.com/performance> for more information.

³Performance compared to prior-generation Intel® Pentium® M processors. Actual performance may vary. See <http://www.intel.com/performance> for processor and benchmark details.

⁶64-bit computing on Intel architecture requires a computer system with a processor, chipset, BIOS, operating system, device drivers and applications enabled for Intel® 64 architecture. Performance will vary depending on your hardware and software configurations. Consult with your system vendor for more information.

⁴Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM) and, for some uses, certain platform software enabled for it. Functionality, performance or other benefits will vary depending on hardware and software configurations and may require a BIOS update. Software applications may not be compatible with all operating systems. Please check with your application vendor.

[†]Hyper-Threading Technology requires a computer system with an Intel® processor supporting Hyper-Threading Technology and an HT Technology enabled chipset, BIOS and operating system. Performance will vary depending on the specific hardware and software you use. See <http://www.intel.com/info/hyperthreading/> for more information including details on which processors support HT Technology.

All products, platforms, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.

All data is based on comparisons of engineering data sheets or measurements using actual hardware or simulators.

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.

Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications. 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 www.intel.com.

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

Copyright © 2007 Intel Corporation. All rights reserved.

Intel, the Intel logo, Intel. Leap ahead., Intel. Leap ahead. logo, Intel Core, Xeon, and Pentium are trademarks of Intel Corporation in the U.S. and other countries.

