Original 45nm Intel® Core™2 Processor Performance

Asim Nisar, Mobility Group, Intel Corporation
Mongkol Ekpanyapong, Mobility Group, Intel Corporation
Antonio C Valles, Software Solution Group, Intel Corporation
Kuppuswamy Sivakumar, Server Platform Marketing Group, Intel Corporation

Index words: PenrynΔ, 45nm Core 2 processor, Performance, SSE4.1, EDAT

Citations for this paper: Asim Nisar, Mongkol Ekpanyapong, Antonio C Valles, Kuppuswamy Sivakumar "Original 45nm Intel® Core™2 Processor Performance" Intel Technology Journal. http://www.intel.com/technology/itj/2008/v12i3/1-paper/1-abstract.htm (October 2008).

ABSTRACT

The 45nm Intel® Core™2 family of processors, codename PenrynΔ, improves upon the performance of Intel Core 2 processors through new microarchitecture features, a larger cache, new instructions, and enhanced power- and thermal-management schemes. This paper presents measured performance data that show the microarchitectural benefits of the Penryn family of processors on key applications and benchmarks. In addition, this paper showcases performance improvements achieved by new SSE4 instructions on a variety of media, imaging, and 3D workloads. The Penryn family of processors also introduced new power- and thermal-management schemes. This paper discusses performance improvements achieved by these enhanced thermal-management features in thermally limited platforms such as mobile thin and light and small form-factor computers.

INTRODUCTION

Performance is an integral part of product definition and success. Intel sets very aggressive performance targets to deliver products with compelling performance to the end user. While considerable effort is placed on functional validation of Intel® processors, Intel also employs significant time and effort to ensure that the processor performance meets expectations at every stage of the product development cycle from concept to silicon arrival to product launch. All design decisions are weighed against performance impact, and appropriate tradeoffs are made. As a result of this extensive effort, Intel delivered a product with record-breaking performance on a wide range of client and server applications.

In this paper, we present information on performance delivered by products based on the 45nm Intel® Core™2 family of processors, codename PenrynΔ. Please see [1] for a detailed architectural description of some of the new microarchitectural features. We begin with an overview of major performance features and then provide an in-depth discussion of measured performance improvements on a wide range of mobile and desktop products. We conclude by presenting performance and energy-efficiency improvements achieved on server platforms built with Penryn processors.

PENRYN MICROARCHITECTURE ENHANCEMENTS

The PenrynΔ family of processors is the next generation of Intel® processors based on the Intel® Core™2 microarchitecture, implemented on Intel's 45nm, Hi-k metal gate process technology. Frequency improvements, within existing power and thermal envelopes, over previous-generation processors, a larger L2 cache, microarchitectural enhancements, and improvements in power- and thermal-management schemes deliver improved performance per watt and energy efficiency for a broad range of client and server applications. The Penryn family of processors also added 47 new SSE4 instructions that can improve the performance of audio, video, image-editing applications, video encoders, 3-D applications, and games.

Microarchitecture enhancements that improve performance in the Penryn family of processors include the following:

- **Larger Cache**: Penryn processors include up to a 50 percent larger L2 cache with a higher degree of associativity that further improves the hit rate, maximizing its utilization. Dual-core Penryn processors feature up to a 6-MB L2 cache and quad-core processors up to a 12-MB L2 cache.

- **Faster Divider:** Penryn processors provide faster divider performance, roughly doubling the divider speed over previous generations through the inclusion of a new, faster divide technique called Radix 16.

- **Super Shuffle Engine:** Shuffles (the repositioning of bits) is a common operation in image- and video-editing applications. By implementing a full-width, single-pass, 128-bit-wide shuffle unit, a processor from the Penryn family of processors can perform full-width shuffles in a single cycle and is 3 times faster than previous-generation processors. The Super Shuffle Engine improves the performance of Intel Streaming Single Instruction Multiple Data (SIMD) Extensions (SSE), Streaming SIMD Extensions 2 (SSE2), Supplemental Streaming SIMD Extensions 3 (SSSE3), and Streaming SIMD Extensions 4 (SSE4) instructions, and this will benefit a wide range of applications including imaging and video applications, games, 3D modeling, and high-performance computing.

- **Inclusion Filter**: An Inclusion Filter was added in the Penryn family of processors to enhance the existing inclusion logic that was limiting server performance.

- **Renamed RSB:** The Renamed Return Stack Buffer (RRSB) increases return prediction accuracy and improves performance.

- **CLI STI Performance Tuning:** In the Penryn family of processors, the Clear Interrupt Flag (CLI) and Set Interrupt Flag (STI) macroinstructions were optimized to perform an execution pipeline serialization only when a new IF value is consumed and only if the new value is not yet updated, instead of post-serializing on every CLI or STI. This improves throughput of CLI-STI pairs by 2.5 times over previous-generation technology.

- **Enhanced Intel® Dynamic Acceleration Technology (EDAT):** EDAT is a power-management feature added to mobile processors that improves energy efficiency by dynamically increasing the performance of active core(s) when not all cores are utilized.

- **Enhanced Intel® Virtualization Technology:** Virtualization partitions or compartmentalizes a single computer so that it can run separate operating systems and software. This virtual partitioning better leverages multi-core processing power, increases efficiency, and cuts costs by letting a single machine act as many virtual 'mini' computers. The Penryn family of processors speeds up virtual machine transition (entry exit) times by an average of 25 percent to 75 percent. This is all done through microarchitecture improvements and requires no virtual machine software changes.

NEW INSTRUCTIONS (SSE4.1)

While many of the microarchitecture enhancements in the PenrynΔ family of processors can be utilized without recompilation, media-related kernels will achieve the maximum performance and power-efficiency gains by recompiling with the Intel compiler and or manually optimizing code, using the new SSE4.1 instructions introduced in the Penryn family of processors.

Intel works closely with industry partners including independent software vendors (ISVs) to understand their performance needs and to improve their applications' performance. The Penryn family of processors' new instructions, SSE4, are a customer-driven response to improve performance on audio-, video-, and image-editing applications, video encoders, 3-D applications, and games. In this section we discuss performance results achieved by using the SSE4 instructions.

Intel® HD Boost technology

Intel HD Boost, the combination of SSE4 instructions and the Penryn family of processors' Super Shuffle Engine, can provide large speedups on a wide range of applications. The following instructions in particular can provide significant benefits to video, imaging, and audio applications.

- There are twelve new integer format conversions that can perform a conversion such as Byte-> Double-Word in one cycle with one instruction.

- The new MPSADBW instruction performs eight sums of absolute differences (SAD) in one instruction. This is twice what the SSE2 PSAD instruction can do.

- The new PHMINPOSUW instruction can be used to perform a horizontal minimum search to locate a minimum unsigned word in an XMM register or a __m128 data type.

The MPSADBW and PHMINPOSUW SSE4 instructions can be used to significantly improve motion vector search algorithms (also known as block matching) used in motion estimation for video applications. An Intel whitepaper [2] showcases how to use these two instructions for block matching. The whitepaper reports a 1.6× to 3.8× performance improvement (see (Figure 1)).

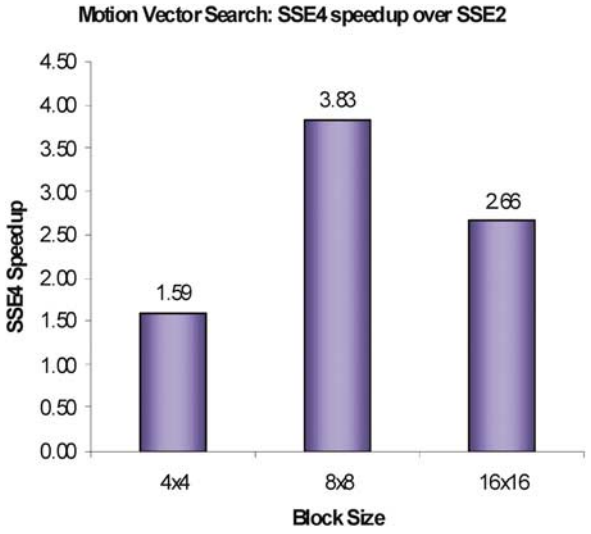

Figure 1: ***SSE4.1 function-level speedups to motion vector search, also known as block matching, used in motion estimation.***

The integer format conversions are commonly used in imaging and video applications. For example, they can be used when converting RGBA from four bytes to four floats prior to computation on a pixel. One SSE4 convert instruction can do the same thing as four SIMD instructions did previously, as shown.

```
SSE2:
pmovd xmm0, m32
pxor xmm7, xmm7
punpcklbw xmm0, xmm7
punpcklwd xmm0, xmm7
cvtdq2ps xmm0, xmm0

SSE4:
pmovzxbd xmm0, m32
cvtdq2ps xmm0, xmm0
```

Conditional moves, blends, early outs

Branches have always been one of the limitations of SIMD code. SSE4 provides new instructions (six Blend instructions plus a PTEST instruction) that can be used to replace either some branches or existing lengthy SIMD code written to get around branches.

The Blend instructions can be used to replace conditional move flows. For example, the PBLENDVB instruction can replace the PAND PANDN POR instructions commonly used in conditional moves where masks are created from a comparison instruction. Another SSE4 instruction, PTEST, can be used as an early out. It is able to compare the entire 128-bit register in one pass. This instruction can be used for conditions that are meant to be infrequent such as divide-by-zero exceptions. One of the benefits of these new instructions is that they provide the compiler more vectorization opportunities; that is, they provide more opportunities to optimize the high-level code by compiling it to use the SIMD instructions.

However, the real benefit of the Blend and PTEST instructions is when multiple branches in a loop can be replaced with multiple Blend and PTEST instructions. The Mandlebrot [3] code shown in Figure 2 is an example that demonstrates how multiple branches can be replaced with multiple PTEST and Blend instructions. In the SSE4 implementation (Figure 3) notice the use of two PTEST instructions:

```
if(_mm_test_all_ones(_mm_castps_si128(vmask)))
if(_mm_test_all_zeros(_mm_castps_si128(vmask),
   _mm_castps_si128(vmask)))
```

and 3 Blend instructions:

```
sx = _mm_blendv_ps(x + sx*sx − sy*sy,sx,vmask);
sy = _mm_blendv_ps(y + _F_TWO_*old_sx*sy,
                   sy,vmask);
iter = I32vec4(_mm_blendv_epi8(iter + _I_ONE_,
                iter,_mm_castps_si128(vmask)));
```

```
void mandelbrot_C()
{
  int i,j;
  float x,y;
  for (i=0,x=-1.8f;i<DIMX;i++,x+=X_STEP) {
      for (j=0,y=-0.2f;j<DIMY/2;j++,y+=Y_STEP) {
          float sx,sy;
          int iter = 0;
          sx = x;
          sy = y;
          while (iter < 256)
          {
              if (sx*sx + sy*sy >= 4.0f)
                break;

              float old_sx = sx;
              sx = x + sx*sx - sy*sy;
              sy = y + 2*old_sx*sy;
              iter++;
          }
          map_C[i][j] = iter;
      }
  }
}
```

Figure 2: ***C implementation of Mandlebrot.***

```
__declspec(align(16)) float _INIT_Y_4[4] = {0,Y_STEP,2*Y_STEP,3*Y_STEP};
F32vec4 _F_STEP_Y(4*Y_STEP);
I32vec4 _I_ONE_ = _mm_set1_epi32(1);
F32vec4 _F_FOUR_(4.0f);
F32vec4 _F_TWO_(2.0f);

void mandelbrot_F32vec4() {
  int i,j;
  F32vec4 x,y;

  for (j=0,x=F32vec4(-1.8f);j<DIMX;j++,x+=F32vec4(X_STEP)) {
    for (j=0,y=F32vec4(-0.2f)+*(F32vec4*)_INIT_Y_4;j<DIMY/2;j+=4,y+=_F_STEP_Y) {
      F32vec4 sx,sy;
      I32vec4 iter = _mm_setzero_si128();
      int scalar_iter = 0;
      sx = x;
      sy = y;
      while (scalar_iter < 256) {
        int mask = 0;
        __m128 vmask = _mm_cmpnlt_ps(sx*sx + sy*sy,_F_FOUR_);
        if (_mm_test_all_ones(_mm_castps_si128(vmask)))
          break;

        F32vec4 old_sx = sx;
        if (_mm_test_all_zeros(_mm_castps_si128(vmask),_mm_castps_si128(vmask))) {
          sx = x + sx*sx - sy*sy;
          sy = y + _F_TWO_*old_sx*sy;
          iter += _I_ONE_;
        }
        else {
          sx = _mm_blendv_ps(x + sx*sx - sy*sy,sx,vmask);
          sy = _mm_blendv_ps(y + _F_TWO_*old_sx*sy,sy,vmask);
          iter = I32vec4(_mm_blendv_epi8(iter + _I_ONE_, iter,_mm_castps_si128(vmask)));
        }
        scalar_iter++;
      }
      _mm_storeu_si128((__m128i*)&map_SSE4[j][i],iter);
    }
  }
}
```

Figure 3: ***SSE4 (using F32VEC4) implementation of Mandlebrot.***

By using the new SSE4 instructions on the Mandlebrot code, the Mandlebrot performance improves by 2.8 times over the C implementation.

Graphics building blocks

SSE4 instructions can be used to speed up graphical applications such as games. The DPPS DPPD instruction can be used to speed up collision detection and common vector matrix operations such as vector normalization. A detailed example of collision detection and usage guidelines of the DPPS DPPD instructions is discussed in [1]. The example showcases a 1.5x speedup in collision detection by using the DPPS instruction and the EXTRACTPS instruction.

A common problem in graphics applications is 'Data Swizzling' or converting from an Array-of-Structures (AOS) data layout implementation to a more SIMD-friendly Structures-of-Array (SOA) data layout in order to use SIMD. Users have to weigh the cost of these conversions before deciding if it is worth using SIMD. By using the INSERTPS instruction, the data-swizzling operation on the next-generation, Penryn microprocessors now take fifteen cycles per four vertices, down from 23 cycles on the Intel 65-nm Core 2 Duo microprocessors, codename Merom. [5] (see Figure 4).

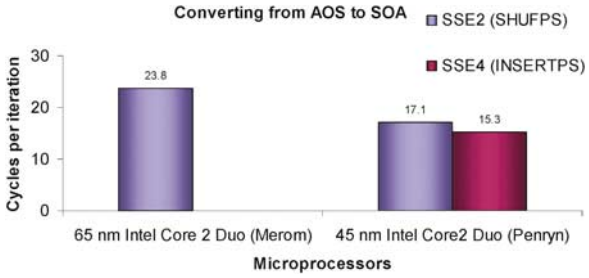

Figure 4: ***Data swizzling improvements per four vertices (one iteration converts four vertices).***

Another potential game improvement is the streaming load instruction MOVNTDQA. This instruction provides a fast method to execute a 16-byte aligned load from Write Combining (WC) memory, such as graphics memory, with a non-temporal hint such that the cache is not polluted.

This instruction can provide a $5\times$ to $7\times$ [6] memory throughput performance increase.

Intel tools

Intel's Integrated Performance Primitives (IPPs), Version 6.0 has over a thousand functions optimized with SSE4. The average speedup across all SSE4-optimized IPP functions vs. SSSE3-optimized IPP functions is $1.12\times$. Figure 5 shows which IPP categories have been optimized with SSE4 and their SSE4 speedup over SSSE3.

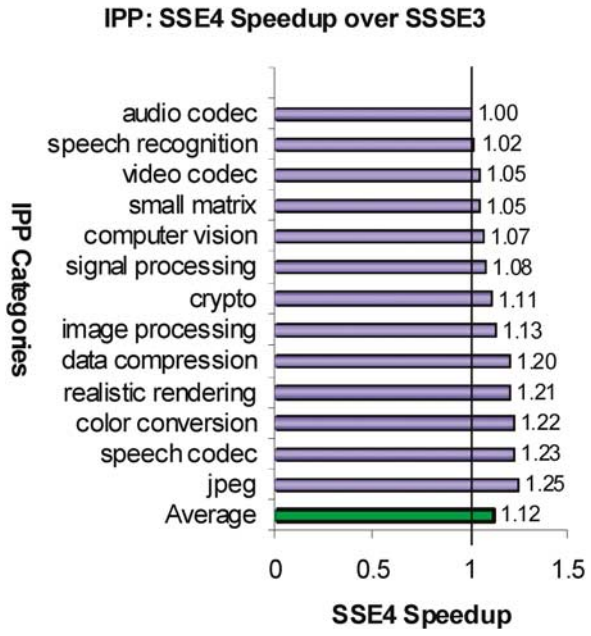

Figure 5: ***Intel Performance Primitive categories and their SSE4 speedups over SSSE3 versions.***

SSE4 instructions also enhance the compiler's ability to vectorize certain loops. Vectorization is when the compiler optimizes a loop to use SIMD instructions including SSE4 instructions. The Intel Compiler, Version 10.0 and later, can be used with the QxS compiler flag to generate SSE4-optimized code specifically for the Penryn family of processors.

SSE4 instructions combined with the Super Shuffle Engine can significantly improve the performance of imaging, video, audio, multimedia, and high-performance computing applications. Users can add these instructions to their applications via assembly code or use Intel tools such as the Intel Compiler 10.0 and IPPs. For detailed information on the SSE4 instructions, including throughput, latency, and optimization guidelines, please see the Intel 64 and IA-32 Architectures Optimization Reference Manual [5] and the instruction manuals [7,8].

INTEL® CORE™ PROCESSOR

Desktop and Mobile

The PenrynΔ family of processors, including dual- and quad-core desktop processors and a dual- and quad-core mobile processor are branded as 'Intel Core processors.'

Desktop and mobile systems built with 45nm Intel® processors, based on Penryn Core architecture, give gamers, researchers, and serious multitaskers a significant performance boost over previous-generation processors. In this section of the paper we present measured performance data on next-generation Intel Core™2 Extreme processors, the new addition to Intel's high-end desktop product line-up. We compare this processor with previous-generation Intel Core 2 Extreme processors on key client benchmarks and real-world applications.

Microarchitectural performance

Improvements

Figure 6 shows a comparison between the Intel Core 2 Extreme QX6850 processor (3.00 GHz, 1333 MHz FSB, 8 MB L2) and the next-generation 45nm Intel Core 2 Extreme QX9650 (3.00 GHz, 1333 MHz FSB, 12 MB L2) on an Asus* P5EX38 board at the same frequency and platform configuration. 45nm quad-core performance is up to 6 percent faster than previous-generation technology on SPEC 2006.

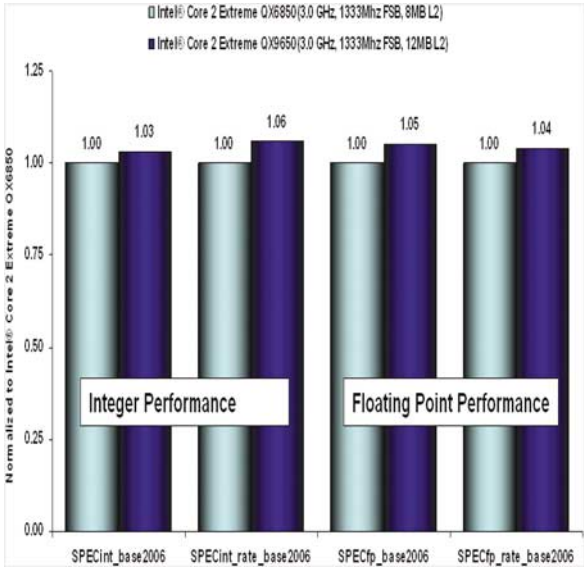

Figure 6: ***Quad-core performance comparison for SPEC CPU2006 at same frequency (estimated SPEC CPU2006 as measured on pre-production systems).***

Video and audio encoding are becoming increasingly important in the world of personal computing. Home-editing of videos and sound recordings are among the popular applications as is standard archiving of DVD material. As shown in Figure 7, the 45nm Intel Core 2 Extreme QX9650 provides a significant boost over previous-generation processors at the same frequency and platform configuration for some of the media-encoding applications. For example, Premiere* Pro CS3 software from Adobe is used to create high-quality visual and editorial effects; it allows users to add color correction, lighting, and other effects such as

audio filters and more, with fast, flexible, built-in tools. As shown in Figure 7, the new Qx9650 is 20 percent faster than the Qx6850 in rendering 210 frames to the disk using this Adobe software. Fathom* is an advanced encoding platform product from Inlet Technologies that is used by media companies to encode content for streaming over the Internet or broadcasting over the air. As shown in Figure 7, Intel measures a 23-percent improvement with new 45nm processors for Fathom to transcode 1080i YV12 high-definition video (HDV) to a 1080i VC1 format. Intel measures a 40-percent improvement for Qx9650 over previous-generation technology for a Pegasys* TMPGenc XPress 4.4 encoder to convert original Variable Bit Rate encoded, 76 second, 29.97fps, 1440×1080 video clips into HDV format MPEG video with 1440×1080 resolution, 29.97 fps, and 25 000 Kbs Constant Bit Rate encoding. Another example is VirtualDub* software, which is a video capture processing utility that uses the DivX* 6.7 software for encoding movies. Virtualdub* 1.7.1 and later with DivX 6.7 are optimized for SSE4 instructions and provide a very noticeable 60-percent performance gain over previous-generation processors that use encoding in SSE2 to convert to the higher-compression DivX format.

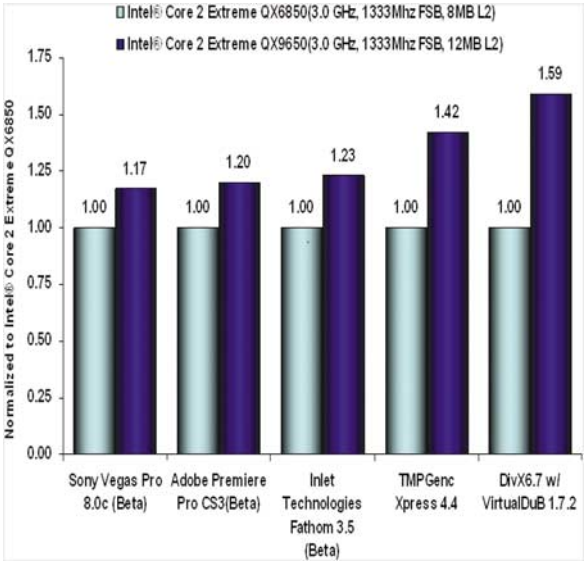

Figure 7: ***Quad-core performance comparison for video encoders at same frequency.***

Figure 8 shows a similar comparison for some of the popular games. The new 45nm Intel® Core™2 Quad processor is roughly 10 percent faster than previous-generation processors at 1024×768. Even when looking at just the two quad-core processors that run at the same FSB and clock speeds, the Intel 45nm Core 2 processors have a clear lead over previous-generation processors. The larger cache, new microarchitectural features, the Penryn high-definition boost, the Super Shuffle unit, and the SSE4 instructions that were discussed in the previous sections all contribute to the increased performance.

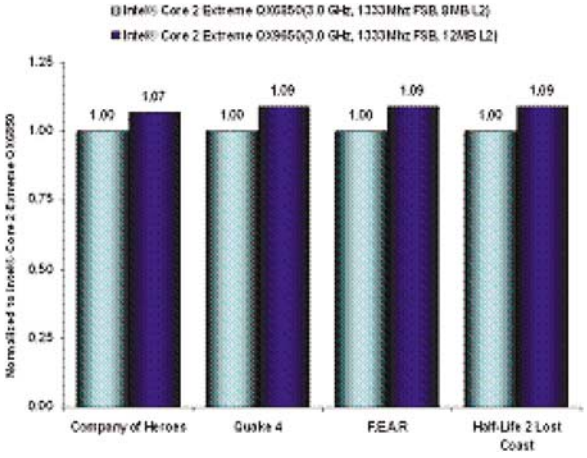

Figure 8: ***Quad-core performance comparison for games at same frequency.***

Frequency and platform improvements

Figures 9 and 10 compare the Intel Core 2 Extreme QX6850 processor (2.93 GHz, 1066 MHz FSB, 8 MB L2) on an Intel 975BX2 board with DDR2 800 RAM with an Intel Core 2 Extreme QX9770 processor (3.20 GHz, 1600 MHz FSB, 12 MB L2 Penryn). Please see Table 1 for a detailed system configuration. This is a more realistic comparison as it takes into account core enhancements, frequency improvements achieved with new 45nm technology, and other platform improvements that were added to support core enhancements. A new 45nm Intel Core 2 Quad-based platform provides double-digit gains on compute-intensive workloads such as SPEC 2006 and the Sysmark 07 Preview that reflect usage patterns of business users in the areas of video creation, E-learning, 3D modeling, and office productivity. Intel measures a 17 percent improvement for Cinebench's multi-threaded rendering test and roughly a 20 percent improvement for Quake 4* and Half Life 2*. A video-encoding application such as DivX and TMPGenc* see a 50 percent to 80 percent gain. The increased frequency and 1600-MHz FSB improves the system bus and memory bandwidth and are the significant contributors to the performance difference. The enhancements in the Penryn family of processors are setting milestones in desktop computing performance.

Table 1: Detailed system configuration for the results shown in Figures 9 and 10.

Processors	Intel® Core™2 Extreme QX6850 8 MB L2, 3.0 GHz, 1333 MHz FSB	Intel® Core™2 Extreme QX9770 12 MB L2, 3.2 GHz, 1600 MHz FSB
Memory	Deluxe Dual channel DS Corsair 2 GB (2 × 1 GB) DDR3-1333 9-9-9-24	Dual channel DS Corsair CM3X1024-1600 C7DHXIN XMP 2 GB (2 × 1 GB) DDR3-1600 7-7-7-20
Graphics card	1 × G8800 GTZPCle graphics	
Motherboard	Asus* P5E3 X38 Deluxe board	
Hard disk	Seagate 320 GB NCQ SATA	
BIOS	Beta 0504, INF:8.4.0.1016, Graphic: NV163.69	

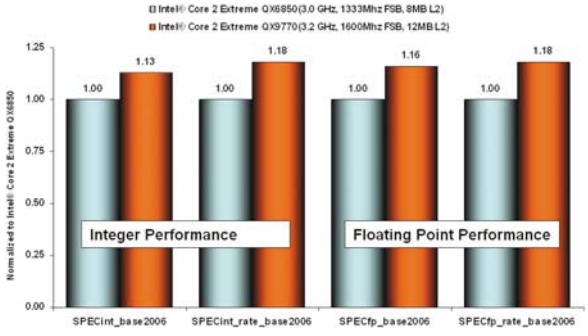

Figure 9: ***QX9770 (45 nm) comparison with QX6850 (65 nm) or SPEC 2006.***

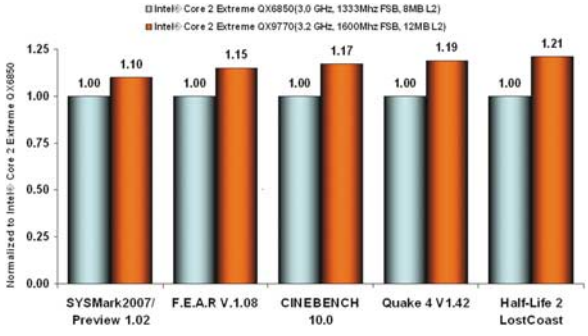

Figure 10: ***QX9770 comparison with QX6850 for games and general purpose applications.***

Enhanced Intel® Dynamic Acceleration Technology

Performance data presented in Figures 6–10 are based on desktop system measurements, but mobile platforms built with 45nm cores will see similar performance improvements over previous-generation platforms. Mobile processors, however, operate at lower frequencies because of tighter power and thermal limitations. In this section we illustrate how 45nm enhancements in Intel® Dynamic Acceleration Technology (DAT) improve mobile platform performance.

DAT is a power-management feature that can improve system performance by increasing the frequency of active core(s) when at least half of the cores in a multi-core processor are inactive and thermal headroom is available. DAT was introduced in the 65-nm Intel Core 2 mobile processors, but in the Penryn family of processors we further enhanced DAT performance and energy efficiency by reducing the number of transitions in and out of DAT, reducing transition overhead in high-interrupt-rate scenarios. In the Penryn family of processors, we also extended DAT support to quad-core mobile processors. Architectural implementations and more details about these enhancements in Intel's Enhanced Dynamic Acceleration Technology (EDAT) are discussed in [4]. Single-threaded applications running on a dual-core or quad-core processor based on the Penryn family of processors, or two single applications (or a two-threaded application) on a quad-core processor, can take advantage of EDAT.

Figures 11-12 illustrate EDAT performance on an Intel Core 2 Quad processor for SPEC 2006, games, and multimedia applications on pre-production mobile platforms with 2 GB of DDR3 memory, a 1066 MHz FSB, a 120 GB hard disk, and a baseline frequency of 2.4 GHz.

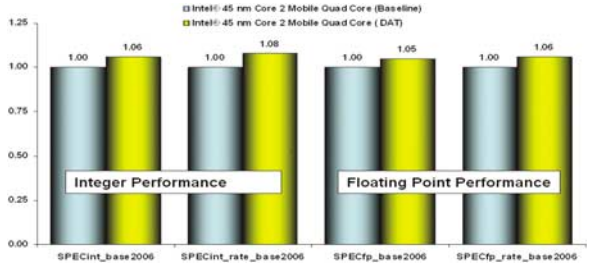

Figure 11: ***EDAT performance improvements (estimated SPEC 2006 as measured on pre-production hardware).***

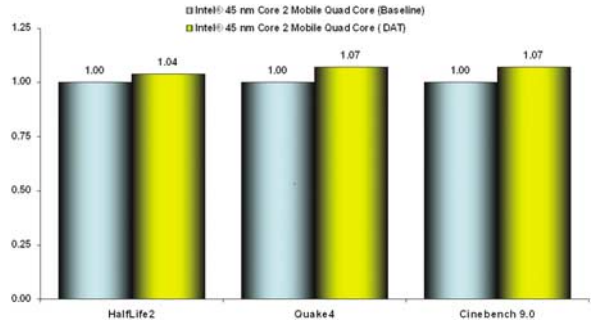

Figure 12: EDAT performance improvements (multi-media and games as measured on pre-production hardware)

To fit QC processors into the mobile Thermal Design Power (TDP) envelope, Intel had to reduce the processor's operating frequency, which in turns gives a wider dynamic range in terms of thermal headroom to operate at EDAT frequencies for one- and two-core operations. The performance gains for EDAT are solely from frequency scaling. The amount of frequency improvement varies from product to product. Single-threaded applications, or two single-threaded applications, or an application with two worker threads with good frequency scaling running on a quad-core processor, can see up to a 10 percent performance boost from EDAT.

INTEL XEON® PROCESSORS

New servers, workstations, and high performance computing (HPC) systems are built with new quad-core Intel Xeon processors 5400 series that are based on the 45nm PenrynΔ core technology. Intel's 45nm technology packs 820 million transistors into the Intel Xeon processor 5400 series. The chip is smaller than the previous-generation Intel Xeon processor 5300 series (214 mm^2 vs. 286 mm^2), which had 582 million transistors. More transistors on new 45nm technology means more capability, performance, and energy efficiency.

Increased performance

The Intel Xeon 5400 series, based on technology from the Penryn family of processors, featuring a larger 12-MB L2 cache, delivered a strong performance gain to the already stable and shipping server platform based on the Intel 5000 series chipset. A drop-in into the existing platform, the 5400 series, added up to a 21 percent performance increase over the previous-generation, quad-core Intel Xeon processor 5300 series for mainstream server benchmarks at the highest frequency level (comparing Xeon X5460 at 3.16 GHz to Xeon X5365 at 3 GHz). Figure 13 shows the comparison on a range of server benchmarks. Figure 14 shows performance comparisons to the previous generation at the same clock frequency (3 GHz) achieving up to a 19 percent performance increase and highlighting the benefits of the improvements in the Penryn family of processors.

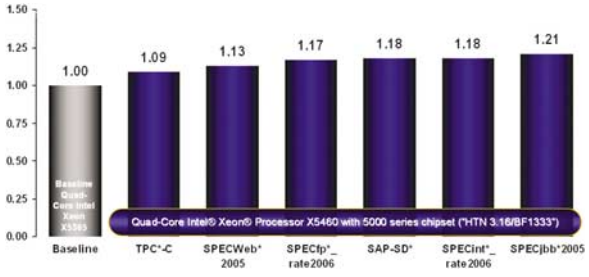

Figure 13: ***Comparison of server benchmarks with those of the previous generation.***

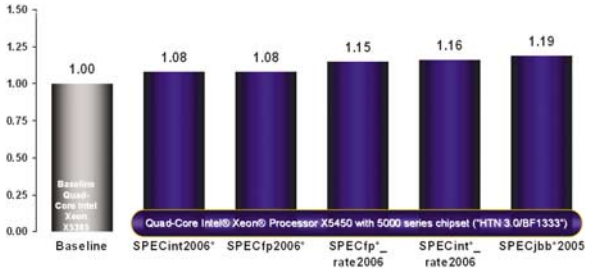

Figure 14: ***Comparison of server benchmarks to those of the previous generation at the same frequency on the same platform.***

The Penryn family of processors' architecture presents several compiler optimization opportunities. These opportunities include tuning for the new ISA, a larger cache, and hardware pre-fetching. The overall gain for the SPEC CPU2006 benchmark suite is shown in Figure 14 as 15–16 percent on the 'rate' benchmark. But gains across the individual components could be as high as 57 percent. Figure 15 shows some of the highlights across the Integer and Floating-point component workloads. All the results shown are on the peak result metric (SPECint_rate2006 and SPECfp_rate2006).

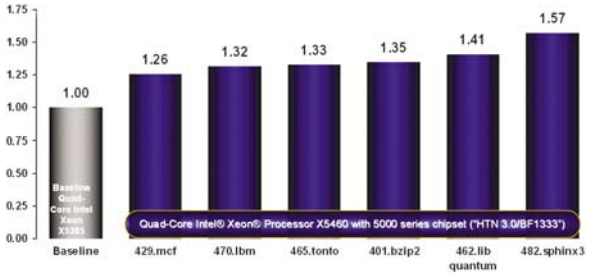

Figure 15: ***Gains on specific SPEC CPU2006 'rate' components at same frequency.***

To leverage the full potential of the 45nm micro-architecture, a new platform targeting the HPC market segment was launched with the new Intel 5400 chipset that could run at the faster FSB speed of 1600 MHz. The additional bandwidth delivered by the platform is critical for the HPC segment. Figures 16a and b show results on key HPC workloads on segments such as manufacturing, financial services, energy, weather and climate modeling, electronic design automation (EDA), and life sciences. It is important to note that the gains achieved by these applications were due to a combination of microarchitectural enhancements and the increased platform bandwidth.

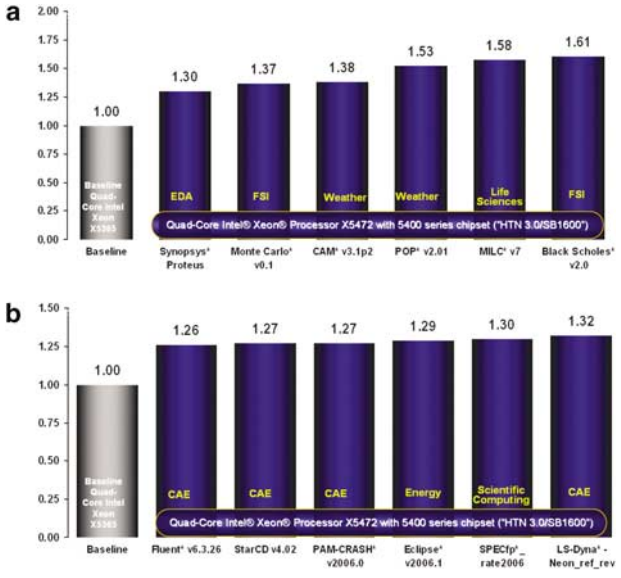

Figure 16: ***(a) HPC benchmarks. (b) HPC benchmarks.***

Improved energy efficiency

Servers based on the quad-core Intel Xeon processor 5400 series also maximize data centers performance and density through improved energy efficiency. As shown in Figure 17, these processors can deliver up to 38 percent more performance per watt in the same platforms and at the same system power level. The platform power in this chart is based on measured average power value at the steady-state window of the benchmark run. For this comparison we used the most energy-efficient mainstream processor SKU from each of the processor families. In this case, that would mean the Intel Xeon E5450 running at 3 GHz compared to the Intel Xeon E5345 running at 2.33 GHz, both at an 80 W TDP rating. The 45nm, Hi-k-based processor also lowers the idle power significantly. Results on the new industry standard SPECpower*s_ssj2008 benchmark, which is the first comprehensive benchmark to measure energy-efficient performance across a load-line including idle power, highlights the energy-efficient performance of the Penryn family of processors. Table 2 shows the top-10 list for this benchmark, all occupied by platforms based on the Penryn family of processors' architecture.

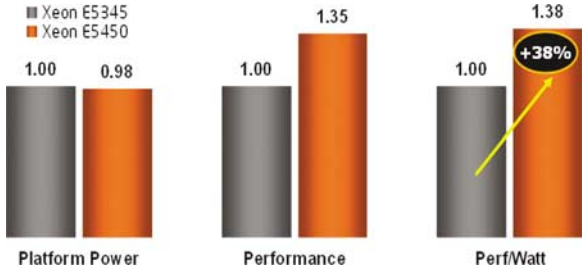

Figure 17: ***Energy efficiency-SPECjbb2005 Benchmark.***

Table 2: Top ten results on the first Industry Standard benchmark for energy efficiency–SPECPower*_ssj2008 (as of June 10, 2008).

Rank	Sponsor	Overall ssj_ops/ Watt	Platform	# of Sockets	Processor(s)	Processor Microarchitecture
1	FSC	1124	TX150 S6	1	Intel® Xeon® X3360	Intel® "Penryn"
2	IBM	1054	X3200 M2	1	Intel® Xeon® X3360	Intel® "Penryn"
3	FSC	1018	TX150 S6	1	Intel® Xeon® X3360	Intel® "Penryn"
4	HP	930	DL180 G5	2	Intel® Xeon® L5420	Intel® "Penryn"
5	IBM	926	X3350	1	Intel® Xeon® X3360	Intel® "Penryn"
6	IBM	913	X3250 M2	1	Intel® Xeon® X3350	Intel® "Penryn"
7	Inspur	910	NF290D2	2	Intel® Xeon® L5420	Intel® "Penryn"
8	IBM	854	X3450	2	Intel® Xeon® E5462	Intel® "Penryn"
9	Dell	800	PE R300	1	Intel® Xeon® L5410	Intel® "Penryn"
10	HP	778	DL180 G5	2	Intel® Xeon® E5450	Intel® "Penryn"

SPECpower results from http://www.spec.org/power_ssj2008/results/power_ssj2008.htmlas of June 10, 2008.

Enhanced virtualization

Virtualization partitions or compartmentalizes a single computer so that it can run separate operating systems and software. These partitions can better leverage multi-core processing power, increase efficiency, and cut costs, by letting a single machine act as many virtual 'mini' computers. Consolidating applications onto fewer systems not only results in better multi-core utilization but improves performance density. In the Penryn family of processors, virtual machine transition (entry exit) times show an improvement of between 25 percent and 75 percent. Based on virtualization benchmark results on different VMMs, the Penryn family of processors provide up to a 20 percent performance gain when compared to previous-generation platforms. Figure 18 shows comparisons of different benchmarks such as VMmark and vConsolidate running various VMMs such as VMware ESX Server, Parallels Virtuozzo, and Virtual Iron 4.0.

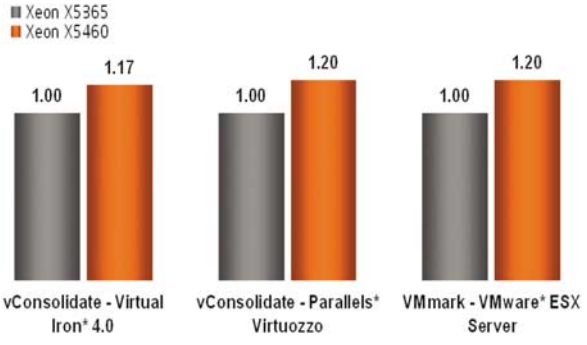

Figure 18: ***Virtualization performance.***

CONCLUSION

The PenrynΔ family of processors brings record levels of performance to the end user through its larger cache, new microarchitectural features, new instructions, and enhanced power- and thermal-management schemes. These processors, manufactured on Intel's 45nm, Hi-k metal gate process technology, not only provide significant performance improvements over the previous-generation processors, but they also provide building blocks for software to be created to take advantage of this power and deliver new levels of functionality to end users.

ACKNOWLEDGEMENTS

The authors thank all the researchers, architects, designers, validators, and software engineers who took the PenrynΔ family of processors from a vision to a real product. Special thanks go to Mike Fard, Jeff Reilly, Ronen Zohar, Daniel Shea, and Eric Palmer who provided performance data and or examples used in this paper.

REFERENCES

[1] Coke, Jim. "Improvements in the Intel® Core™2 PenrynΔ Processor Family Architecture and Microarchitecture." Intel Technology Journal, Volume. 12, Issue. 3, 2008.

[2] Kuah, K. "Motion Estimation with Intel® Streaming SIMD Extensions 4 (Intel® SSE4)," at http://softwarecommunity.intel.com/articles/eng/1246.htm.

[3] Mandelbrot set. http://en.wikipedia.org/wiki/Mandelbrot_set.

[4] Jahagirdar, S. "Power-Management Enhancements in 45 nm Intel®Core™ Microarchitecture." *Intel Technology Journal*, Volume 12, Issue 3, 2008.

[5] *Intel® 64 and IA-32 Architectures Optimization Reference Manual*. http://www.intel.com/products/processor/manuals/.

[6] Jha, A. and Yee, D. "Increasing Memory Throughput With Intel® Streaming SIMD Extensions 4 (Intel(4) SSE4) Streaming Load." At http://softwarecommunity.intel.com/articles/eng/1248.htm.

[7] *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A: Instruction Set Reference, A-M*. At http://www.intel.com/products/processor/manuals/.

[8] *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B: Instruction Set Reference, N-Z*. At http://www.intel.com/products/processor/manuals/.

AUTHORS' BIOGRAPHIES

Asim Nisar is a Senior Architect with Intel's Mobile Microprocessor Group in California, USA. He led the 45nm Intel® Core™2 performance modeling, analysis, performance projections, pre-silicon and post-silicon performance validation activities. Prior to this, he was involved in defining and enhancing several microarchitectural features. He has an M.S. degree from the Georgia Institute of Technology, USA. His e-mail is asim.nisar at intel.com.

Mongkol Ekpanyapong is a Senior Architect with Intel's Mobility Microprocessor Group in California, USA. He received his B.E. degree from Chulalongkorn University, Thailand, his M.E. degree from the Asian Institute of Technology, Thailand, his M.S. and Ph.D. degrees from the Georgia Institute of Technology, USA. Mongkol was involved in performance modeling and validation for the Intel® Core™2 Duo processor architectures. His current focus is on performance modeling for the next-generation Intel microprocessor. His e-mail is mongkol.ekpanyapong at intel.com.

Antonio C Valles is a Senior Software Engineer in Intel, Chandler, AZ focusing on broad and in-depth pre-silicon and early-silicon tests of Intel microprocessors and chipsets. Antonio has created multiple internal pre-silicon and post-silicon tools and kernels for performance analysis and coordinates development of tuning guidelines for the processors. He received his B.S. degree in Electrical Engineering from Arizona State University in 1997. His email is antonio.c.valles at intel.com.

Kuppuswamy Sivakumar (Siva) is a Marketing Manager for Intel's Server Platforms Group in California, USA. He received his B.E. degree from the University of Madras, India, his M.S. degree from the University of Kentucky, Lexington, and his MBA from the University of California, Berkeley. Siva manages performance marketing activities for Intel® Xeon® server products. His e-mail is kuppuswamy.sivakumar at intel.com.