White Paper

Introduction to Intel's 32nm Process Technology

Intel has been in high volume manufacturing on its 32nm process technology with second generation high-k + metal gate transistors since Q4, 2009. This process technology builds upon the tremendously successful 45nm process technology – the first to use high-k + metal gate transistors – introduced a couple of years before.

Building upon the 45nm process technology that led to the launch of the initial Intel® Core™ i7, i5, and i3 processor family based on Intel microarchitecture codenamed Nehalem, Intel began high volume production with its 32nm process technology with second generation high-k + metal gate transistors in late 2009. This 32nm process technology is being used to manufacture the latest versions of processors based on the Intel microarchitecture codenamed Nehalem. Intel offers these 32nm based products across all computer market segments – embedded, mobile, desktop, and server. Intel is a leader in semiconductor manufacturing process technology and was the first company shipping 32nm microprocessors for PCs – and the only company for well over a year. Intel is on track with its cadence of new product innovation – known as the "Tick-Tock" model – that delivers alternating new generations of advanced manufacturing process technology and processor microarchitecture every year: first the "tick" of a new manufacturing process technology applied to the existing processor microarchitecture, then the "tock" of a new processor microarchitecture built on the latest manufacturing process.

To understand the significance of 32nm, we must first take a look back at 45nm process technology and high-k + metal gate transistors:

To better understand the significance of the 32nm process technology it is helpful to look back and revisit the 45nm process technology introduction that occurred in 2007. The Intel 45nm process that enabled Intel to launch the incredibly successful and high performing Intel microarchitecture codenamed Nehalem was the first to introduction high-k + metal gate transistors – a technological breakthrough that enabled higher performing transistors while simultaneously reducing leakage currents. At the introduction of the 45nm process technology Intel promised a fast and meaningful ramp of the 45nm technology.

In fact, the 45 nm production ramp has been the fastest in Intel history. The 45nm processor unit production has ramped twice as fast as the 65nm process technology in its first year. Today, 45nm products are being manufactured across computing segments as represented by single core Intel Atom™ processors, dual core Intel Pentium® processors, Intel Core i7 processors with six cores, and even eight core Intel Xeon® processors are all today manufactured on the 45nm process.

Intel takes another major leap ahead on the 32nm process technology with second generation high-k + metal gate transistors.

The foundation of the 32nm process technology is the second generation high-k + metal gate transistor. The improvements over the first generation high-k + metal gate transistors are many. The equivalent oxide thickness of the high-k dielectric has been reduced from 1.0nm on 45nm process to 0.9nm on the 32nm process, while gate length as been reduced to 30nm. Transistor gate pitch continues to scale 0.7x every two years - with Intel's 32nm process providing the tightest gate pitch in the industry.

The 32nm process also uses the same basic replacement metal gate process flow as Intel's

45nm process technology, enabling Intel to take advantage of an existing highly successful process. These improvements are critical for scaling the size of integrated circuits and increasing transistor performance. The 32nm process technology with second generation high-k + metal gate transistors enables designers to optimize for size, performance and power simultaneously.

The decreased oxide thickness and reduced gate length enables a >22% transistor performance gain in terms of drive current. These transistors provide the highest drive currents and tightest gate pitch reported in the industry. Leakage current can also be optimized for a >5X reduction in leakage over 45nm for NMOS transistors, and >10X reduction in leakage for PMOS transistors. These improvements combine to enable circuits to be designed that are both smaller and have improved energy efficiency. The 32nm process also uses the 4th generation of strained silicon technology for improved transistor performance and energy efficiency – and Intel's engineers have been improving the transistors with each generation.

Intel has already shipped over half a billion processors with high-k + metal gate transistors (using both 45nm and 32nm manufacturing process technology).

The 32nm yield curve mirrors the rapid improvement Intel saw with its previous fastest ramping technology ever – 45nm. Great 32nm process and product health enables Intel to accelerate the planned ramp of 32nm products.

Intel is very proud of the ramp and yields achieved on the 45nm process. Intel was able to achieve a rapid defect reduction with 45nm technology. This was done despite revolutionary features and new industry leading technology. Yield improvements for the 32nm process have matched or exceeded the rate of the 45nm process.

Intel has brought four fabs (semiconductor manufacturing facilities) online over two years to transition the volume processor lineup to 32nm as a result of a $7B investment announced in early 2009. The four fabs are D1D and D1C in Oregon, Fab 32 in Arizona and Fab 11X in New Mexico.

Intel processors based on 32nm are shipping in growing volumes into embedded, mobile, desktop, and server segments – and expected to exceed half of the mainstream processor shipments by the middle of 2011.

A dramatically repartitioned mainstream client: greater performance and lower power via higher integration.

The 32nm client processors got more than just improved performance headroom and smaller die sizes. The initial client Intel Core i7, i5, and i3 processor designs also underwent a dramatic repartitioning of platform capabilities. Previously, mainstream Intel based PCs used a 3-chip combination consisting of (1) a chip containing the processor, (2) a chip containing a memory controller, a graphics controller and/or PCI express interface, a display interface, and – for business clients – the circuits supporting Intel® vPro™ Technology, and (3) a chip supporting primarily I/O functions such as local area network controllers and USB device interfaces.

For the initial 32nm based client processors, the memory controller is integrated with the processing core die in the processor multi-chip package. The graphics and the memory controller circuits are on a 45nm based die in the package with the 32nm based processing core die. The remaining functions from the second chip plus the functions from the third chip are now combined in a second chip such that what used to take three discrete components is now done in two. This simplifies platform design and reduces power requirements for improved efficiency.

The 2nd generation Intel Core processors due in 2011 (based on the Intel microarchitecture code named Sandy Bridge) further integrate functions into the processor and add features to both the processor and the secondary component. The processor is a single 32nm based die with integrated graphics, integrated memory controller – and even some I/O functions integrated – yielding improved performance, power management capabilities, reduced system costs, and platform design flexibility.

And 32nm is only the latest in Intel's continuing march of advancing technology

Even as the 32nm process is ramping to full production volume, the engineers in Intel's development pipeline are busy preparing future process generations. In September 2009, Intel disclosed that it had the world's first working 22nm silicon technology. Specifically, Intel disclosed that it had functional SRAM test chips with 364 Mbits, a whopping 2.9 billion transistors per chip. These SRAMs had the smallest cell in working circuits to date, at 0.092 square microns. The 22nm process is on track for production in the 2nd half of 2011, two years after start of 32nm high volume production.

December, 2010

www.intel.com

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.

Printed in the United States.