


New Microprocessors architectural Approaches - The 'era of the tera'

Dr. Danny Rittman Feb, 2004



Intel's P4 – 1 Layer Poly

In the early 1980s, PCs were basic office productivity tools that offered little more than word processing and basic spreadsheet capabilities. The 1990s witnessed the advent of the graphical user interface, the mouse, and color monitors. That was followed in the mid-90s by the advent of the Internet, multimedia capabilities, joysticks, and speakers for enhanced sound capabilities. That brings us to the modern PC, which features wireless connectivity and enhanced multimedia capabilities.

Historical Processors – The birth of an Intel value


At the time, the 8086 and related 8088 ushered in a new age of computing, many at Intel viewed these two relatively modest chips as a means to buy time until the launch of the powerful 432, a higher-end processor designed to deliver mainframe-like capabilities in silicon. In fact, the two processors bought much more than just time—the enormous success of the 8086 and 8088 chips helped Intel initiate a

succession of processors that incorporated numerous features and steadily increased the processor speed.

Today, Intel® architecture processors—the label given to descendants of the 8086—form the foundation of the computing industry. The 80286, later known as the i286™, debuted in 1982 and broke the 1 MB limit on addressable memory. Three years later in 1985 the advent of the 80386 processor inaugurated a second wave of personal computing. The Intel 486™ processors that followed bumped up the clock speed, utilized five-stage pipeline architecture, and integrated an instruction data cache.

Enter the Pentium® processor—five times more powerful than the original Intel 486™ processor and 300 times faster than the 8088. The Pentium® processor shipped in 1993 and represented a significant departure from the 386 generation. Best of all, this superscalar chip could execute more than one instruction at a time.

The complexity of developing ever faster processors had now become one of Intel's greatest challenges. Intel's first microprocessor, the 4004, ran at 108 kilohertz (108,000 hertz). The Intel® Pentium® 4 processor's initial speed was a blazing 1.5 gigahertz (1.5 billion hertz). If automobile speed had increased similarly over the same period, you could now drive from San Francisco to New York in about 13 seconds.



Intel's 8086



Intel 4004 microprocessor.

Where do we go from here?

What have we learned? For one thing, Moore's law (which states that circuit density increases predictably over time), has been continuous, and as a result, the capabilities of modern PCs have transformed every aspect of our daily lives. So quickly has the amount of data available on the Internet swelled, that new scientific notations might soon be needed to quantify the "sea of digital information" at our fingertips. Probably to the extent, that soon it would touch every human on the planet, 24-seven, in every modality of their lives.

The process of handling that enormous repository of information could be divided into the three phases of "RMS": *recognition, mining, and synthesis*, each of which can be improved with advances in microprocessor architecture.

Recognition is the ability of a computing device to detect an anomaly and alert the user. Such an early warning capability would have been useful by the U.S. State Department, which recognized a terrorist threat only after the tragic events of 9/11. Another advanced application of recognition might be in medicine, where sophisticated diagnostic tools might be used for early detection of a malignancy. Recognition is followed by "mining," wherein a recognized event is processed in the context of an existing database. Using the example of the 9/11 events the system would compare intelligence information to a repository of data on known terrorists.

The final aspect of information processing is "synthesis," wherein information is acted upon to elicit a desired response. Once again drawing on the terrorist example a system might alert airlines and other transportation authorities that an imminent threat exists. In a medical application, synthesis might create a potential prognosis for the condition diagnosed (recognized) based on the medical records (mining).

The 'era of the tera'

The RMS workload of tomorrow's network shapes the architectures developed by the world's microprocessor technology; Intel. Intel is constantly developing new architectures to best meet user needs, and that the process would usher in the "era of the tera," a period during which high-horsepower computers would deliver to users *terabytes* of data at *tera*FLOPS speeds.

The new era would require all the computing power that Intel and others could muster. As evidence of the need for speed, Intel demonstrated in one of the hi tech conferences new "ray tracing" technology that uses high-powered graphics-rendering capabilities to create remarkably realistic graphic images with dynamically accurate lighting, shadows, and metallic reflections of an object in its immediate environment. Developed by Stanford Professor Philipp Slusallek, ray tracing technology is already being used by the automotive industry to design virtual vehicles that are nearly indistinguishable from photos of the final products.

"How many FLOPS could you use for an application like this?" asked Pat Gelsinger Intel Chief Technology Officer at the spring 2004 Intel Developer Forum, of Slusallek, "I can guarantee we can use all the processing power you can create in the next few years," answered Slusallek, adding that the current entry-level requirements for the application are a bank of five dual-processor PCs.

Reaching beyond Moore's Law

Power-hungry applications such as ray tracing will require creative architectural problem-solving beyond the extension of Moore's law. Among the hurdles the company would face are heat management and memory access speeds.

Some of the innovations that would be required are further refinements of Intel® Hyper-Threading technology, which was only a "small first step" in the performance enhancements that are available using multi-thread technology. Future architectures will draw on greater thread levels and can be characterized as "helper threads" that will effectively pre-warm execution sequences, significantly speeding overall performance.

Tomorrow's processor performance will also be enhanced by adaptive features in the architecture and in the platform that would make maximal use of computing power while helping eliminate transistor power "leakage" during computational down periods.

"As Intel quickly approaches billion-transistor processors, we are aggressively developing innovative ways to use the added transistors by applying new architectural approaches that will ultimately result in much more useful computers." This will be done through architectures that can scale and adapt to their computing environments.

Taken in total, new architectural innovations will spur the development of new generations of silicon products, helping "enable the era of the tera."

Reference:

Intel Developer Forum - <http://www.intel.com/idf/us/spr2004/index.htm>

Intel Microprocessors - <http://www.intel.com/research/mrl/research/arch.htm>