# From Moore's Law to Intel Innovation—Prediction to Reality

Radhakrishna Hiremane
Technical Marketing Engineer
Intel Corporation

# Table of Contents

(Click on page number to jump to sections)

# From Moore's Law to Intel Innovation—Prediction to Reality

Radhakrishna Hiremane
Technical Marketing Engineer
Intel Corporation

## *Overview: Bettering Billions of Lives*

Observations only become prophetic when they're profound enough to remember and when history proves them true. Moore's Law, published as a simple observation by Gordon Moore in 1965, has become exactly that: a prophetic observation about integration that has been made a reality—led by Intel's innovative silicon and manufacturing technologies, investments, and research. We've made Moore's Law a reality for nearly 40 years, and will continue to do so well into the future.

For nearly four decades, Intel technologies have touched billions of lives and changed the way the world lives, works and plays. We've been a world leader in silicon process technology, manufacturing technology, and thermal and power management technologies. Intel's researchers and engineers continue to achieve the next evolution of computing architectures and platforms because they also continue to make the vital trajectory predicted by Moore's Law a reality.

"We've used Moore's Law to drive the convergence of computing and communications," said Intel CEO Craig Barrett. "Intel's commitment to Moore's Law now allows us to create integrated platforms that deliver a broad range of capabilities for individuals and organizations that use technology. To realize the full potential of these capabilities, continued innovation and industry cooperation will be more important than ever."

## *Intel Innovations: Past, Present and Future*

If steel was the raw material for the 20th century, silicon is for the 21st century. And the silicon semiconductor industry—led in large part by Intel's technology advances—has delivered a dramatic spiral of rapid cost reduction and exponential value creation that is unequalled in history. Because of the cumulative impact of these spiraling increases in capability, silicon—the raw material of the microprocessor—powers today's economy and the Internet, running everything from digital phones and PCs to stock markets and spacecraft—and enables today's information-rich, converged, digital world.

But if silicon technology's past has been dramatic, its future promises to be even more spectacular and far-reaching. The next several billion users will come from parts of the world where technology has traditionally been less pervasive: Asia, Eastern Europe, Latin America, the Middle East and others areas. These users represent the next growth wave for the technology industry, and—more importantly—an opportunity to make profoundly positive changes to the human condition globally. These changes won't happen without enormous innovation and investment by the technology industry to design next-generation products tailored to the needs of billions of users, both current and future. Silicon technology advances will pave the way.

History shows some remarkable and innovative milestones achieved by Intel since the company was co-founded by Dr. Moore. The invention of the DRAM. The development of the first CMOS silicon process. Breaking the 1-micron transistor gate-length barrier. Then breaking the 100-nm gate-length barrier. Now the world's most advanced, 65-nanometer (nm) silicon process technology. And just recently, Intel's breakthrough silicon photonics announcement of the first continuous wave all-silicon laser using a physical property called the Raman Effect—built using standard volume CMOS silicon.

Fulfilling the predictions of Moore's Law means accomplishing the near impossible—again and again. Intel's silicon chip designers continue to make transistors smaller and smaller, which means reducing process geometries—scaling (or shrinking) the nominal feature size of the devices populating and powering the silicon chip. Scaling the process geometries makes more space available to bring more transistors, as well as to converge different types of devices and functions, onto the chip.

With a technology vision that extends to 2015 and beyond, Intel's innovations in silicon and high-volume production will soon enable not only the next generation of computing, but the cost-effective convergence of computing, communications, consumer electronics, and other exciting new devices with advanced capabilities and uses.

## *Origins of Moore's Law*

Moore, one of the founders of Intel, observed in an article in the April 19, 1965 issue of Electronics magazine that innovations in technology would allow a doubling of the number of transistors in a given space every year (in an update article in 1975, Moore adjusted the rate to every two years to account for the growing complexity of chips), and that the speed of those transistors would increase. What is less well-known is that Moore also stated that manufacturing costs would dramatically drop as the technology advanced.

Moore's prediction, now popularly known as Moore's Law, had some startling implications, predicting that computing technology would increase in value at the same time it would actually decrease in cost. This was an unusual idea at the time since, in a typical industry, building a faster, better widget with twice the functionality also usually means doubling the widget's cost. However, in the case of solid-state electronics, the opposite is true: Each time transistor size shrinks, integrated circuits (ICs) become cheaper and perform better.

In 1965, a single transistor cost more than a dollar. By 1975, the cost of a transistor had dropped to less than a penny, while transistor size allowed for almost 100,000 transistors on a single die. From 1979 to 1989, to 1999, processor performance went from about 1.5 million instructions per second (MIPS), to almost 50 MIPS on the i486™, to over 1,000 MIPS on the Intel® Pentium® III. Today's Intel® processors, some topping out at well above 1 billion transistors, run at 3.2 GHz and higher, deliver over 10,000 MIPS, and can be manufactured in high volumes with transistors that cost less than 1/10,000th of a cent.

Intel's innovations in shrinking transistor size while dramatically improving performance mean that advanced technology can be delivered for almost every industry, from computing and communications, to agriculture, mechanics, and medicine. Today's microprocessors power the economy, fuel the growth of the Internet, and run everything from toys to traffic lights. That in turn has touched billions of lives, and transformed the world and the human experience.

## *Changing the World*

At a fundamental level, Moore's Law has become synonymous with technological evolution and the continuous steps made to enable new capabilities and usage models. Intel has continued to imagine what that evolution could be like, and then has focused on delivering those new usage models, new platforms, and new opportunities for better living and business and economic growth.

"Technology leadership begins with a vision—often a bold vision," said Justin Rattner, Intel senior fellow and director of Intel's Corporate Technology Group. "You create that vision, you pursue that vision, and you turn that vision into reality. Sometimes it changes hardly anything and sometimes it changes everything."

In the case of Intel’s innovations, the vision has helped change the world. Without the remarkable pace of Intel technology, the world we live in today would be very different. For example, without Intel’s innovations, it is unlikely that data streaming and video downloads via the Internet would be possible. Video conferencing would be an exotic technology tool, and video editing would require a room full of computers. Cell phones, which require sophisticated buffer processes, would probably still be in the realm of the imagination, and users would still be wrestling with large, heavy, clunky "mobile" phones. Our technologies have also helped enable numerous advances in medicine, such as telemedicine, complex imaging systems, high-performance compute power for DNA and pharmaceutical research.

Intel is even now researching future technologies that could enable a phone that translates languages in real time so you can talk to people in other countries more easily. We’re also studying equipping houses with technologies that can actually help you monitor your health and detect early signs of disease. We’re using the principles of Moore’s Law to make technologies that were once considered science fiction become affordable realities.

Intel's superior research, development, and process technologies have allowed the power of silicon to spread across the globe from cities to farms, offices to homes, and impacted billions of people.

## *Can the Trend Continue?*

The big question is, can this doubling continue and with it ever-increasing functionality and ever-decreasing costs?

Continuing to deliver innovation to make the predictions of Moore’s Law a reality means shrinking the nominal size of the devices that populate the silicon. Skeptics in the industry have believed that going down that path of decreasing transistor sizes would be more and more difficult since, as transistors shrink in size, they consume less power (they scale in voltage), but their leakage current (the continued flow of current even when transistors are “off”) increases. The more transistors there are on a chip, the more power is wasted. Also, as transistor density and speed increase, the chip as a whole consumes more power and generates more heat. Thus, the efficiency of cooling techniques must also increase to dissipate the heat from the increases in device density and current leakage.

The difficulty is that researchers even now are coming up against the physical limits of atomic structure for scaling transistors while still managing power and thermals. This issue—as well as other factors—creates a significant and continuous challenge for the entire silicon industry.

In response to that challenge, Intel is aggressively pursuing research into both conventional and unconventional technologies that will remove any obstacles to enabling the next evolutionary step in computing. This includes introducing many new and exciting technologies and innovations in materials, design, and packaging. For example, at the process level, Intel is addressing current leakage through strained silicon technology, while also investigating novel structures of transistors (such as tri-gate transistors) and unique dielectric materials (such as the dielectric referred to as Hi-K material). Other Intel innovations, such as the multi-core processor, Hyper-Threading Technology† (HT Technology), execution trace cache, and Enhanced Intel® SpeedStep® Technology, address the challenge at an architectural level, while at the same time, add dramatic performance enhancements to Intel-based platforms.

Intel believes the answer to the question of continuing on the trajectory predicted by Moore’s Law lies in addressing power challenges at every level, from silicon to system. This means using a combination of packaging technology, new transistor designs, exciting new lithography improvements, and other breakthrough technologies to deliver a complete solution that will enable evolution for years to come.

## *Technologies for Today and Tomorrow*

Intel's roadmap for researching and developing new process technologies is a long-term vision that seeks to make real the predictions of Moore's Law well into the future. This plan addresses every variable and manufacturing aspect that affects the power equation. It encompasses technologies from conventional CMOS processes to research in unconventional materials, such as carbon nanotubes or carbon nanowires.

For example, two years ago, Intel developed and deployed the first high-volume production 90-nm process technology using strained silicon, reducing the current leakage from transistors by a factor of five (or more, actually) without reducing on-current performance. And Intel is already readying its shift toward 65-nm process technology, with improved, second-generation strained silicon, and transistors whose gate length is just 35 nm—so tiny that about 100 of these gates could fit inside the diameter of a human red blood cell. And beyond that? Intel is already researching 45-nm, 32-nm and 22-nm process technologies.

Some of Intel's latest innovations and breakthrough research areas include:

- Packaging technology, including eliminating the bumps of solder that make the connections between the package and the chip, and so reducing the thickness of the layers and allowing the further shrinking of devices.
- Transistor design, including novel, tri-gate transistors that reduce leakage current in general, and so could reduce power consumption in mobile devices.
- New dielectric materials, such as High-K, which reduces leakage current by a factor of 100 over silicon dioxide.
- Extreme ultra-violet (EUV) lithography, which uses a wavelength of 13.5nm, is expected to enable the printing of features that are 10nm and below.
- Silicon photonics, including the world's first continuous wave silicon laser, which solves the previously insurmountable, two-photon absorption problem.

## *Silicon Integration—the Power to Innovate*

As transistor counts climb, so does the ability to increase device complexity and integrate many new capabilities on a chip. This combination of count, complexity, and convergence creates a richer set of resources to make silicon devices even more capable and improve the flexibility and cost effectiveness with which they can be applied.

"Moore's Law isn't just about more transistors," said Paul Otellini, Intel president and chief operating officer. "It's also about how creatively you use those transistors."

One example of Intel’s creativity is in low-power circuit design techniques, which include adaptive body biasing and sleep-transistor technology. Essentially, by dynamically adjusting the voltage applied to the body of a transistor (bias), researchers can manipulate the threshold voltage at which a transistor turns on. This means that the bias voltage can come under localized control. With localized control, the voltage leakage of specific transistors can be reduced during periods of inactivity, while performance can still be increased during peak use of the system.

Another example of Intel’s creativity is the tri-gate transistor. This is a novel, three-dimensional transistor structure with a depleted substrate. The transistor’s raised, plateau-like structure improves the drive current, while the depleted substrate reduces leakage current when the transistor is in the off state. The result is an overall reduction in power consumption, which in turn results in improved battery life and more compact form factors in mobile devices. It’s another way that Intel is again delivering innovation that will improve life for millions of mobile-device users.

Scaling process geometries makes more space available to add more transistors, more functionality, and more integration. Increased transistor budgets also allow engineers to provide new capabilities, such as the *T’s (Hyper-Threading Technology†, Intel® Virtualization Technology, LaGrande Technology, Intel® Active Management Technology, Intel® I/O Acceleration Technology), integrated communications capabilities, and dual- and multi-core platforms. Intel engineers have used the design flexibility provided as a result of increased transistor counts in the incorporation of air gaps, moving parts, antennae, or other new structures to add functionality through silicon, and to develop dual and multi-core processors.

## *Moving into Ultra-Violet*

Extreme ultra-violet lithography, or EUV, is another dramatic step forward in doubling the density of circuitry. Today's most advanced lithography is limited by the wavelength of visible light, which is 400 to 650nm. In contrast, EUV lithography uses a wavelength of 13.5nm, which will enable the printing of features that are 10nm and below (By contrast, Intel’s current volume manufacturing produces feature sizes of 50nm). This leap in feature size reduction will enable Intel to continue to achieve the predictions of Moore's Law.

Of course, there are interesting challenges in making EUV—and other advances—a viable manufacturing technology, For example, EUV light is absorbed by glass, so instead of lenses, mirrors must be used. Also, since EUV light will not pass through a glass mask, a reflective mask must be used to reflect the light in certain regions and absorb it in others, so that a circuit pattern can be effectively transferred to the wafer. Intel's researchers and engineers are already working to address and solve those challenges, and deliver yet another breakthrough technology for advancing the silicon industry.

### *The Lights Are On*

Along with transistor design, packaging, and process technologies, Intel is also looking into unconventional technologies that may offer new solutions and opportunities. For example, the industry is already coming up against the physical (atomic) limits of chip-to-chip electrical signaling. The solution to this challenge may be in silicon photonics, an area that Intel is now aggressively pursuing.

Intel has already created a wide variety of the essential structures needed to move light around a chip just as conveniently as electrons are moved. The key element that had been lacking before was a consistent light source. This is where Intel's dramatic breakthrough of the first continuous wave all-silicon laser using a physical property called the Raman Effect—built using standard volume CMOS silicon—comes in.

Using the power of silicon, Intel researchers have been able to take the functionality of a conventional, bulky Raman laser—which uses glass and is usually about the size of a boom box—and reduce it to a single line on a silicon wafer. You can see a demo of the technological breakthrough on Intel's Silicon Photonics Web site.

This silicon photonics breakthrough could lead to practical and affordable applications for computing and telecom, as well as new medical equipment and sensors, where a tunable silicon laser could replace models that cost tens of thousands of dollars. This could also lead to speedy new optical interconnects for chips and backplanes, because thin optical fibers take up less space than electrical cables, leaving more space to cool down computers and servers.

As with advances in lithography and transistor design, Intel's optical science builds on existing processes, adapting and improving them to handle new advances and scientific breakthroughs. For example, Intel's demonstration wafer, full of silicon lasers, was made using standard CMOS manufacturing in an existing fabrication path. What this means is that the path between the laboratory and factory for this exciting new technology may not be long or painful, as with some unconventional technologies, but amazingly direct and quickly viable.

### *What Will the Future Look Like?*

Four decades of innovation have made the predictions of Moore's Law a reality, but Intel believes the best is yet to come. By 2015, Intel envisions processors with tens to potentially hundreds of cores per processor die. Those cores will be supporting tens, hundreds, or maybe even thousands of simultaneous execution threads.

Intel is even now researching 3-dimensional (3D) die and wafer stacking technologies, which could move device density from hundreds or thousands of pins, to a million or 10 million connections. This is the type of dramatic increase in memory-to-processor connectivity that will be required to deliver the bandwidth needed to support Intel's upcoming many-core architectures.

Intel also expects to see more natural, more human, and more error-tolerant interfaces; personalized, interactive 3D entertainment; and intelligent data management for both home and business applications. There will be predictive modeling for home and business applications; seamless, real-time collaboration between geographic locations; and self-configuring and self-maintaining platforms that make owning and operating devices—desktop, mobile, or personal—more painless.

Said Paolo A. Gargini, Intel fellow and director of Technology Strategy in the Technology and Manufacturing Group, "Moore's Law is important, not so much as a unit of measure in terms of who is leading the pack in silicon technology, but as a milestone, both as a mark of where we've come from in the last 40 years, and as a possibility for where we're headed."

We don't know what the world will look like in 10 years, but at Intel we can imagine the possibilities. And, our focus will remain on researching and delivering the breakthrough technologies that enable the usage models of today and tomorrow. This includes a convergence of computing and communication, increasingly powerful applications for home and business use, and an extension of the benefits of technology across the globe.

## Summary

It's been 40 years since Intel co-founder Gordon Moore made an observation that functions as the "time pacing of technology" [Shona Brown in the Harvard Business Review]. And for most of that time, Intel has researched, developed, and delivered the innovative technologies that have kept evolutions in computing in line with Moore's prediction—and will continue to do so in the future. Our technologies have touched billions of lives and changed the way the world lives, works and plays.

"We envision a world where everyone is touched by our technologies," said Pat Gelsinger, Intel senior vice president, "whether through traditional personal computing or consumer electronics, interactive or proactive devices."

In this article we reviewed some of the important breakthroughs that Intel researchers have developed to remain on the trajectory predicted by Moore's Law. With some of the industry's top minds and research talent, unmatched investment in world-leading fabrication facilities, and the world's most advanced process technology, Intel has the unique ability to translate advanced developments and scientific breakthroughs into next-level performance and high-volume manufacturing. This has allowed Intel to not only maintain the path predicted by Moore's Law, but to extend that trajectory of evolution from computing into communications and consumer electronics.

Silicon process technology continues to be the foundation for Intel's architecture and platform innovation. Intel is taking advantage of the extra transistor budget available through the evolving process technology to provide new capabilities and features such as the *T's, and increase performance through dual- and multi-core computing platforms. Intel's innovation is expanding the reach of silicon process technologies to influence new areas such as standard CMOS radios, silicon photonics and health initiatives—extending its leadership and competitive edge in delivering new and innovative platforms and solutions.

There is little doubt that Intel's performance in technology innovation has transformed the world, but the future is likely to be even more dramatic. Increases in device counts and complexity are enabling remarkable new usage models, bringing technology to industries ranging from toys to kitchen appliances, and from cell phones to security. Improvements in performance and price have made these advanced technologies available to users who would never have dreamed when Gordon Moore wrote his original 1965 article that they'd be using them every day—from singers to senators, cab drivers to architects. This is Moore's Law in action, proven by history and brought to you by Intel, the company that continues to imagine and establish the vision, and then deliver the future.

## More Info

You can learn much more by visiting the following areas of the Intel Web site:

- Moore's Law
- Intel Silicon Technologies
- Platform 2015
- Architecting the Era of Tera

## Author Bio

Radhakrishna (RK) Hiremane is a technical marketing engineer at Intel Corporation. He joined Intel's LAN Access Division in 1999 as a software engineer working on Ethernet device drivers for 10/100 and Gigabit products and has represented Intel in the development of the Uniform Driver Interface specification. As part of the Software Solutions Group, he also worked in the area of software performance analysis and optimization on large SMP Intel® Itanium® processor-based systems running both Linux* and Windows* operating systems. Hiremane holds a bachelor's degree in engineering, electronics and communications from Bangalore University, India, and an M.S. in computer engineering from the University of New Mexico. He is currently pursuing his M.B.A. from Arizona State University.

† Hyper-Threading Technology requires a computer system with an Intel Pentium 4 processor at 3.06 GHz or higher, a chipset and BIOS that utilize this technology, and an operating system that includes optimizations for this technology. Performance will vary depending on the specific hardware and software you use. See http://www.intel.com/info/hyperthreading/ for information.

*—End of Technology@Intel Magazine Article—*