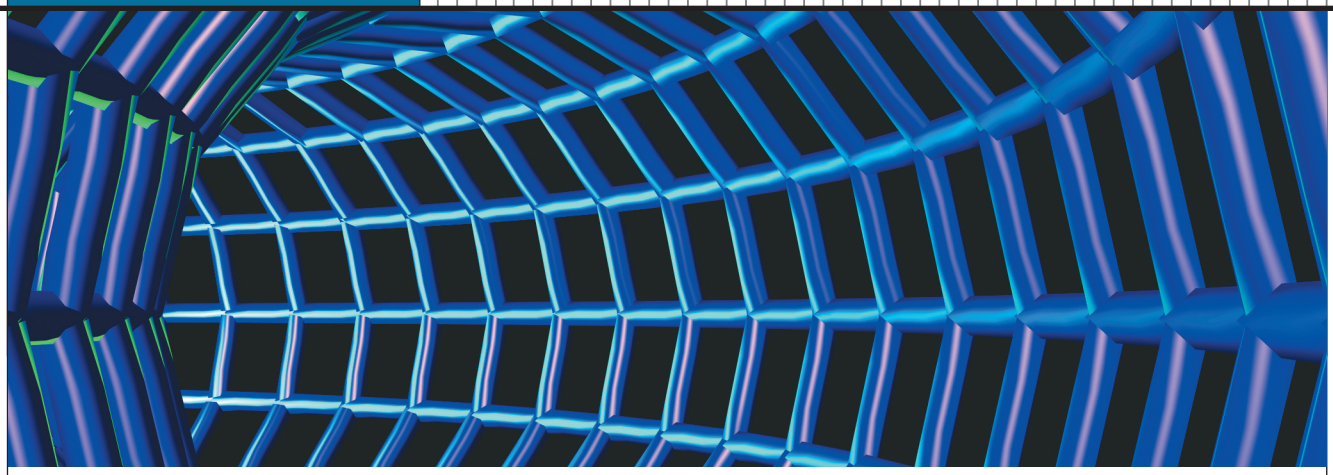

Microprocessor Scaling: What Limits Will Hold?

Jacek Izydorczyk and Michael Izydorczyk, *Silesian University of Technology*

For 40 years, semiconductor technology has followed Moore's law, but continuing that pattern will require a breakthrough in energy-efficient design. Eventually the industry must also seek an entirely new paradigm, such as creating reliable systems from unreliable components.

Nature imposes fundamental limits on computer design in the form of thermodynamic laws and restrictions from quantum mechanics. From the first integrated circuit (IC) design to modern systems on a chip, designers have had to account for these in some form. In addition, technological evolution has been governed by Moore's law—IC cost decreases, and the number of transistors on an IC doubles approximately every two years.

Lately, however, people have begun to question Moore's law and the validity of long-accepted technology development metrics in future efforts. Will the pattern slow or escalate? We believe that technology, in terms of bits processed per second per square centimeter of IC real estate, will develop exponentially in the next four decades. The first 20 years or so will be dedicated to finding new materials, devices, and circuits to enhance energy-efficient data processing, but the decades that follow must focus on new design paradigms, such as creating reliable and predictable systems from unreliable components. Toward the end of the next 40 years, we also expect energy-efficient analog processing, which would enable computing machines with an architecture that is remarkably similar to the human nervous system.

Our research and the results of other experiments have led us to conclude that the capabilities of future technology will depend not on the size of integrated devices but on the IC's ability to drain dissipated heat. Thus, we see two main problems: how to reduce the energy needed to process one bit of data and how to drain heat from the chip. To seek answers, we've examined Moore's law, the hierarchy of limits on technology development, thermodynamic laws, and information theory. Our exploration has led us to make some interesting predictions.

PROGRESS ACCORDING TO MOORE

Many of our deductions are based on the past evolution of complementary metal-oxide semiconductor (CMOS)

 Published by the IEEE Computer Society

technology according to established laws, primarily Moore's law. "The Birth of Moore's Law" sidebar describes the story behind this enduring principle.

Initially, in the 1960s Gordon E. Moore looked at semiconductor improvement in four ways:[1]

- size of the silicon wafers,
- overall miniaturization of the technology,
- imagination of those allocating devices on the wafer, and
- innovation in circuit construction.

However, after Moore and Robert Noyce founded Intel and began churning out microprocessors and semiconductor memories, Moore changed the way he viewed improvement. At first, progress was largely a matter of rapidly making smaller and smaller components, but in an extremely short time, the entire semiconductor industry went digital and the market opened for advanced ICs, which changed the design focus. Improvement began to be measured in the number of transistors that could be integrated on an IC, which in turn depended on reducing fabrication dimensions. This new design focus slowed the improvement rate, and Moore ventured a prediction: the number of integrated transistors would double about every 26 months.[2] Moore's prediction might have remained within Intel's walls were it not for publications like *Science* and *Scientific American*, which popularized his and Royce's articles.[1]

In retrospect, it's hard to say with certainty if Moore was prescient or if he was merely outlining a development roadmap with an eye toward the already fierce market competition and capacity for technical innovation. Whether or not it was a self-fulfilling prophecy, Moore's law has held since it first became known as such. Only recently have people begun to ask how much longer this law will remain valid.

THE BIRTH OF MOORE'S LAW

In 1954, Gordon Earle Moore was a doctoral candidate in physical chemistry at the California Institute of Technology and wrote his dissertation on nitrogen's infrared spectroscopy in physical chemistry. Shortly after, William Shockley hired Moore to work at Shockley Semiconductor,[1] which was his first hands-on experience with semiconductor devices. Moore and his colleagues, including Robert Noyce, designed the technology to produce silicon transistors. Unfortunately, Shockley's research was aimed at improving an exotic four-layer diode, a vision that clashed with that of his younger coworkers.

Consequently, in 1957, Moore, Noyce, and six others (the infamous Traitorous Eight) broke away and started up Fairchild Semiconductor, with Moore quickly becoming the company's main technologist and a key manager. The goal was to produce transistors from silicon instead of the then popular germanium, with the idea that the devices would be disposable after a certain period.

In 1960, the company built a circuit with four transistors on a single silicon wafer, creating the first silicon IC. By the mid-1960s, ICs were competing with circuits assembled from discrete elements, so IC manufacturers began concentrating on minimizing the fabrication cost of a single device inside the IC. Because almost all IC production cost stemmed from circuit development, the cost of fabricating a single device decreased as the number of devices on the IC increased.

However, with more devices per IC came the problem of surface area on the silicon wafer: the larger the surface, the greater the exposure to defects and the lower the yield. Consequently, the goal became to find an optimal IC size that guaranteed a minimal single device cost at an acceptable performance level. As companies began fabricating more ICs in line with this goal, Moore noted a pattern, which he first documented in 1965.[2] The minimal device cost decreased for circuits that contained more transistors, and the number of transistors on one circuit had doubled every year since the construction of the first IC in 1960. Moore later revised the doubling period to every two years, a pattern that remains valid today and is well known as Moore's law.

References

1. **C. Lécuyer and D.C. Brock, "Biographies: Gordon Earle Moore," *IEEE Annals of the History of Computing*, vol. 28, no. 3, 2006, pp. 89-95.**
2. **G. Moore, "Cramming More Components onto Integrated Circuits," *Electronics*, vol. 38, no. 8, 1965, pp. 114-117.**

HIERARCHY OF LIMITS

To answer that question, we turned to the hierarchy of constraints on computing machines proposed by J.D. Meindl et al. in the late 1980s.[3] At the top are fundamental limits that incorporate the most general principles of efficiency; at the bottom are circuit limits specific to CMOS technology.

Fundamental limits

Any reasonable computing machine must be in thermodynamic balance with its environment, following certain laws that govern its efficient use of heat or other energy forms. Entropy, which is the disorder that thermodynamics or statistical mechanics research must consider, is inherently the same for all kinds of computing systems. The law of information theory, for example, governs information transmission, limiting the efficiency of information processing, such as the amount of energy to process one bit of data.

Other examples of these limits are that signals cannot propagate at a speed greater than the speed of light in a vacuum, and the machine's smallest parts can't violate quantum mechanics. Thus, particles must obey the uncertainty principle; electrons are limited by the Pauli exclusion principle; and so on.

Material limits

These limits arise from the properties of the semiconductors, metals, and insulators used in IC production, which constrain a computing machine's dimensions and

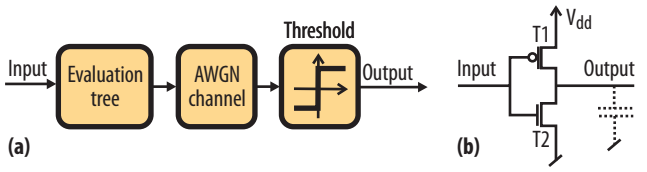

Figure 1. Limits of logic gates. In this model of (a) a three-stage logic gate and (b) a CMOS inverter, the machine could do computations without energy dissipation, but it must spend energy to read out the result under the constant impact of thermal fluctuations.

processing power. Researchers agree that material limits arise from four obstacles: switching energy, transit time, thermal conductance, and dopant fluctuation.

Device limits

Device limits can be traced to the continued use of bulk metal-oxide-semiconductor field-effect transistors (MOSFETs) as switching devices. Problems with current leakage and capacitive coupling between a gate and a channel are motivating researchers to study alternative solutions, such as a fully depleted silicon-on-insulator MOSFET with multiple gates.

Circuit limits

CMOS logic circuits primarily dictate circuit limits, and these circuits still suffer from latency and signal contamination when interconnected. Other problems include performance fluctuation and yield, both of which are of extreme interest in light of emerging logic technologies. So far, most researchers don't believe that these alternative technologies have much promise to improve performance.[4]

One potential limit buster is the introduction of material with a high dielectric constant, or high-κ, for gate insulation. Although some researchers didn't believe the material would be introduced in the next CMOS generation,[5] toward the end of 2007 Intel started commercial production of the 45-nm Penryn processor based on high-κ technology.[6] The Intel Core i3, i5, and i7 microprocessors have successfully implemented this processor.

CHALLENGES AT THE BOTTOM

Generally, the higher the limit in the hierarchy, the more chance that technological progress will bust it. Thus, we believe that lucrative breakthroughs won't be at the higher levels. Rather, we concur with Richard Feynman, who remarked "there's plenty of room at the bottom," meaning that big profits are at the lowest level, where creative design working within the laws of thermodynamics and quantum mechanics can reap rewards—and where energy efficiency should be the primary goal.

Logic gates

Computations are thermodynamically irreversible because computers inevitably dissipate energy; a microprocessor *must* transfer heat to the ambient environment. Overall heat transfer is strictly related to temperature, a core concept in thermodynamics. Entropy, which is basically the amount of data processed, is also a key idea in thermodynamics. In logic gate description in thermodynamics, as well as in information theory, entropy influences the limits that designers must consider.

To illustrate the limits imposed by logic gates, we offer the model in Figure 1a,[7] which has three cascaded stages. The first stage is the evaluation tree, which is a mechanism to compute the logic function's value without energy dissipation.

The next stage is the additive white Gaussian noise (AWGN) channel, which describes the process of reading the output from the evaluation tree. The AWGN adds Gaussian, mutually uncorrelated noise values to the data transferred from input to output. Information theory treats the AWGN channel as a basic model of a real system that transfers data in the presence of noise.

The third stage is a threshold element that restores binary data read out from the evaluation tree. We include this stage only to complete our examination of thermodynamic principles. We don't view it as necessary to a logic-gate model because performing this function immediately after receiving input isn't an optimal strategy.

Figure 1b shows a basic CMOS inverter. Even this basic form reveals elements of the model in Figure 1a. The induction of a channel under transistor T1's gate and the depletion of that channel under transistor T2's gate determine a function's value. Alternatively, the value is computed by the voltage applied to the transistors' gates.

It's possible to read the logic machinery's state only when the current flows through one of the transistors, either T1 or T2, and loads or unloads a parasitic output capacitor. Charging the capacitances leads to an energy accumulation, which discharging dissipates.

Although common sense says that energy dissipation is mandatory, in reality it isn't essential for CMOS logic, a principle illustrated in adiabatic logic gates. These gates are an example of how changes in the circuit alone, without new devices or novel materials, can greatly reduce energy dissipation. Adiabatic logic has certain drawbacks, and it isn't certain how prevalent this kind of circuit will become. However, it is proof that the charging-discharging cycle for energy dissipation isn't a requirement.

It also implies that only thermal noise disturbs this cycle and that a certain amount of this noise is unavoidable. The AWGN channel in Figure 1a accounts for the idea that an infinitesimally low voltage ramp cannot charge capacitances. The AWGN channel not only introduces

noise but also reflects the band limit dictated by the current flow's dynamics.

Thermal noise and energy dissipation

To understand why thermal noise is unavoidable, consider the simple one-dimensional model of an *n*-channel field-effect transistor (FET) in Figure 2. A FET is the basic device for current CMOS technology and likely to remain a building block in future technologies such as carbon nanotubes. The two energy wells in the figure denote the source and drain, which hold the electrons responsible for conductivity. The gate is a high-energy region between the source and drain. Electrons are subject to Fermi-Dirac statistics—a law that governs how a crowd of electrons must behave. Below the Fermi level, electrons are crowded, filling almost all possible energy states; above the Fermi level, many states are empty and the electrons are free.

Dopant concentration in the source and drain region is usually high enough (10^{15} to 10^{20} cm^{-3}) that the Fermi level is either just below the bottom of the conduction band or even a little above it. If the potential energy of electrons in the gate region is sufficiently large, the source and drain are separated.

Decreasing the energy barrier in the gate region allows some electrons to surmount the barrier and thus be anywhere in the transistor. These escaped electrons have an energy level high enough to generate quantum leakage through the barrier. As Figure 2 shows, if a voltage difference (U) exists between the source and drain, high-energy electrons will enter the drain region. Once they interact with other electrons and the crystal lattice, which lowers their energy level, they'll remain in that region. Even if the electrons scatter free over the barrier, some mechanism is still needed to dissipate their energy so that they remain in the drain region. Otherwise, electrons reflect from the rear drain wall, much like a Ping-Pong ball bounces off a hard floor. At first the ball is attracted to the hard floor by gravity, but then it consistently bounces high after colliding with it.

Trapping electrons in the drain region in this way is an effective strategy when an energy increase caused by the move from source to drain significantly changes the average number of electrons in the state more than once. To illustrate, assuming that Maxwell statistics and the transfer of one electron charge ($e \approx 1.602 \times 10^{-19}$ C) from the source to the drain are sufficient to change the gate's state. The energy dissipated is then

$$eU = k\mathrm{T}\ln 2 = \varepsilon_b,$$

where k is the Bolzmann constant of approximately 1.38×10^{-23} J/K and T is the FET temperature in Kelvin. Thus, the equation above represents the ultimate Landauer limit. The lower limit of the energy spent to process one bit of data at room temperature ($T = 300$ K) is extremely small ($\varepsilon_b \approx 17.9$ meV, where 1 eV $= 1.602 \times 10^{-7}$ pJ).

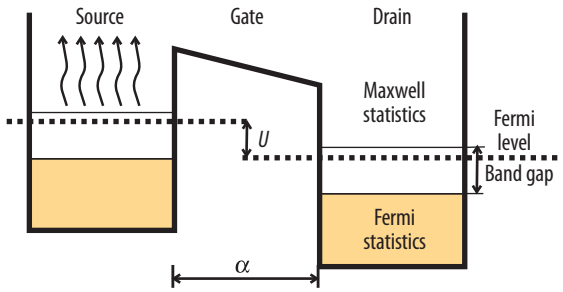

Figure 2. Simple model of an *n*-channel field-effect transistor (FET). The source and drain hold the electrons responsible for conductivity. Below the Fermi level, electrons fill almost all possible energy states. Above the Fermi level, many states are empty and the electrons are free, described by Maxwell statistics. If a voltage difference (U) exists between the source and the drain, high-energy electrons will enter the drain region, where they will remain (localized).

Quantum mechanics restrictions

Transistor dimensions are restricted by the product of particle momentum dispersion and particle position dispersion, which in turn is bounded by the uncertainty principle. Thus, any attempt to localize the particle in the tightly bounded position can enlarge the particle momentum disturbance. When a single electron is confined in the potential well, the uncertainty principle manifests itself as the constant product of an energy gap between the permitted energy levels and squared well width. Energy levels inside a wide well are very close, and a scattered particle causes them to merge into a single band.

In contrast, energy levels in a very narrow well are wide apart and thus far less likely to overlap, even when the particle interacts with the crystal lattice. A small FET naturally dissipates energy while the electron is localized. An extremely small device—about 5 nm wide—can dissipate the excess energy, but a larger device—about 28 nm wide—translates to an energy gap on the order of the Landauer limit.

To illustrate these restrictions, consider a hypothetical IC with the following characteristics:

- energy dissipated per processed bit is as low as the Landauer limit at room temperature ($\varepsilon_b \approx 17.9$ meV);
- it uses planar technology;
- its modus operandi is one electron device;
- quantization of electron energy allows minimal linear dimensions; and
- it has a switching frequency of 3.5 GHz, consistent with the top commercial microprocessor line.

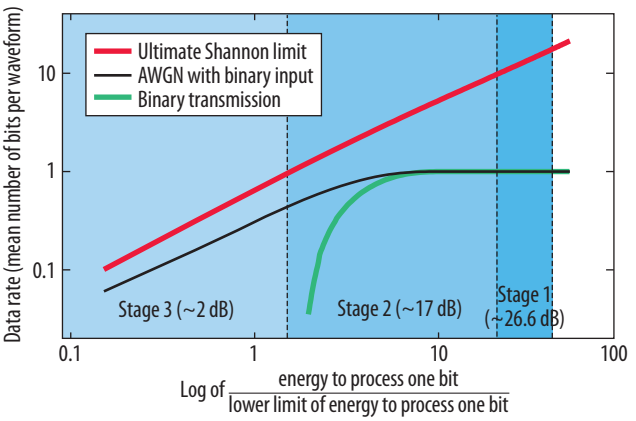

Figure 3. Three evolutionary stages of energy-efficient technology. The horizontal axis represents the energy spent to process one bit of data on a double logarithmic scale moving from the Landauer limit to infinity. The vertical axis represents data rate as the mean number of bits carried by one waveform. The ultimate Shannon limit divides the plane into two domains. In the upper domain, reliable transmission and data processing are impossible. The lower domain contains the additive white Gaussian noise (AWGN) for binary transmission and with binary input.

More than 400 W/cm^2 of heat density must be removed from the surface of this hypothetical chip—which is hard to achieve outside a research lab. Assume that the IC temperature is higher than the ambient temperature and less than 400 K. Theoretically, the best planar-circuit cooling method would be forced liquid flow and boiling, which can drain up to 1 kW/cm^2 of heat from the chip.[8] For 3D structures, the theoretical amount that can be drained is even higher, about 10 kW/cm^2.

However, in practice neither this method nor a heat stream formed from thermoelectric phenomena will make a significant difference in the chip's 300 W/cm^2 border.[8] Radical progress in cooling technology seems unlikely. As the sidebar "The Impracticality of Cryogenics" describes, cryogenics has long been a source of hope, but practical concerns have kept it from commercial development. Energy dissipation must come from another source.

One possibility is to reduce the energy spent processing data. To put the problem in perspective, the Intel Core2 (four cores) has a total power dissipation of 136 W with 820 million switching transistors and a 3.2-GHz switching frequency (www.intel.com/products/desktop/processors/index.htm?iid=processors_body+dt_core). Assuming a conservative average of two transistors to process one bit of data, the value of energy per bit is approximately 647 eV. Obviously this value must greatly decrease, but integrating as many transistors as possible on a chip works against that. With switching frequency unlikely to improve, the most densely integrated circuit will consist of 1,012 switching elements per square centimeter and operate at 8 MHz only.

STAGES OF ENERGY EFFICIENCY

We predict that the work to create energy-efficient technologies will evolve in three stages: a new metric for integration density and processing speed, a new design paradigm that allows the construction of reliable machines from unreliable components, and energy-efficient analog data processing. Figure 3 shows the three stages.

A new density metric

The current estimated lower bound on energy used to process one bit is approximately 647 eV, about 36,000 times higher than the absolute Landauer limit. Unfortunately, to achieve the Landauer limit the design must use advanced coding and decoding algorithms.

At present, digital design typically consists of a binary symmetrical channel organized along the lines of a Gaussian channel and with bipolar binary modulation. The optimal demodulator would be a matched filter, sampler, and threshold element that maximizes the power ratio of useful signal to thermal noise and thus minimizes the probability of transmission error.

Optimizing the decoder will become important. Current logic circuits aren't built around the concept of an optimal decoder. However, the minimal energy needed to process one bit of data *without* coding bits, a value of 1.4 eV, is only weakly related to the heat stream density removed from the chip and the acceptable error ratio.[7] And heat stream density is key to energy efficiency as it cannot exceed 300 W/cm^2.

The heat stream drained from the IC—the energy spent to process one bit, multiplied by the number of bits processed per second—isn't likely to change. For the past two to three decades, a typical general-purpose microprocessor has consistently used slightly more than 1 cm^2 of its silicon wafer. The energy spent to process one bit is now about 647 eV, and number of bits processed per second per 1 cm^2 of IC is now about 10^{18}.

All this argues for a new processing-speed metric. The latest trends show a systematic decrease in area with an asymptotic value of 1 cm^2. Assuming that new value, we believe it's possible to estimate a microprocessor's total processing power by calculating the number of bits processed in one second on 1 cm^2 of wafer space. This value could then replace the number of integrated devices as the metric of integration density and processing speed. If in the near future heat stream density increases to 300 W/cm^2, and assuming that the practical limit of energy dissipated per bit remains 1.4 eV, an integration density of approximately 1.3×10^{21} bits per cm^2 per second should be possible—a 1,300-fold increase in processing power.

In Figure 3, this first evolutionary stage spans about

26.6 dB on the energy scale, and the design preserves the equivalence between waveform and bit, or wave-to-bit correspondence, with a reasonable error rate.

A new design paradigm

If the integration density growth rate that results from Moore's law remains constant, we anticipate about 22 years of uninterrupted development. At the end of this time, we expect to see the next evolutionary stage. Ad hoc methods similar to those described by Robert P. Colwell, chief architect of the P6 microprocessor family, can't alleviate soft errors caused by thermal noise. In *The Pentium Chronicles*, Colwell describes transient errors in the Pentium 4 machines: ". . . forward progress of the engine is self-monitored, and if too much time has elapsed since forward progress was last detected, a watchdog timer will flush the machine and restart it from a known-good spot. . . ."[9]

This brief description underlines the need for a revolution in circuit design. If indeed no device is reliable because of transient errors, then the design paradigm should shift to reliable, predictable systems from unreliable components with unpredictable behavior.

Designers use two fault-tolerant techniques to increase a system's reliability, availability, or dependability. The first is to select highly reliable components. The second is to use protective redundancy, just as communication uses a repeat code. However, neither of these techniques is the path to the envisioned paradigm shift.

A more promising approach is to implement techniques akin to advanced forward-correction codes at all levels of the system hierarchy, starting with circuits that implement base arithmetic operations such as add, shift, complement, and multiply.

Invented codes could have several properties similar to those of communication codes, as well as having many distinct and unique differentiating properties. With these codes, the minimal energy spent processing one bit would be only slightly above the Landauer limit. Eventually, the processing speed could reach 6.6×10^{22} bits per cm^2 per second without breaking the heat dissipation limit of 300 W/cm^2. If the technological growth rate resulting from Moore's law continues, we foresee about 12 years of development.

In Figure 3, this second evolutionary stage spans about 17 dB on the energy scale. Wave-to-bit correspondence gradually relaxes, which probably reflects the greater number of devices processing one bit.

Analog data processing

As Figure 3 shows, below the energy scale of 28.1 meV, further binary data processing is impossible, but analog computation remains feasible. The small distance between the limiting energy of the second stage and the Landauer limit restricts the final exploration of analog computations. Even so, the huge discrepancy between the data rate of binary transmission and the ultimate Shannon limit in the first and second stages can support energy-efficient analog data processing.

Experimental results show that an analog coprocessor used to solve ordinary differential equations dissipates no more than 1 percent of the energy of a general-purpose digital microprocessor and no more than 20 percent of the energy of a digital signal processor. At the same time, computations have been up to 10 times faster.[10] Perhaps this is a harbinger of the analog computing era that so many researchers have predicted.

We're optimistic about the prospects for microprocessor scaling and see the possibility of exponential growth and the retention of Moore's law. Although no solution to energy reduction per bit or to cooling is on the horizon, the continuous small upgrades of current CMOS technology are evidence that no one can be certain of the

THE IMPRACTICALITY OF CRYOGENICS

Although cryogenic devices were initially promising, offering the potential for ultrahigh switching speeds, operating large equipment at temperatures close to absolute zero was an unsolvable problem. If the IC is cooled to cryogenic temperatures, fluctuations caused by thermal noises drop drastically, and hence less energy is needed to process one bit of data. Unfortunately, the requirement to keep the circuit at a temperature below the ambient temperature requires some cooling mechanism. A cryogenic computer dissipates little heat, but energy must be dissipated to drain off heat to keep the circuit below ambient temperature. And the efficiency of any cooling machine is worse than the efficiency of a Carnot cycle. Consequently, the total dissipated energy for computation and cooling is no lower than the heat dissipated during computation at ambient room temperature.

Moreover, "cold" electrons are associated with very long waves, and the device size limit is close to what devices are already. The switching time of such a circuit remains at 1.4 ps, but the heat density is only 43 W/cm^2, which allows the construction of an IC with a switching frequency in the hundreds of gigahertz.

These realities make cryogenic technologies very resistant in practical applications, which prompted Niklaus Wirth, Swiss computer scientist and creator of Pascal, to include them in a group of elegant but not too sensible ideas about computing.[1] In his retrospective of cryogenics, Wirth noted that, with the appearance of personal computers, cryogenic dreams were expected to "either freeze or evaporate,"[1] and nothing on the horizon promises to change that outlook.

Reference

1. **N. Wirth, "Good Ideas, through the Looking Glass" *Computer*, vol. 39, no. 1, 2006, pp. 28-39.**

technology's limitations. We remain confident that such upgrades will continue to provide insights into new directions for materials, devices, circuits, downscaling, and interchip communication. As researchers solve new problems, they'll push the limits of microprocessor scaling until at some point a design paradigm shift will occur.

References

1. E. Mollick, "Establishing Moore's Law," *IEEE Annals of the History of Computing*, vol. 28, no. 3, 2006, pp. 62-75.
2. G. Moore, "Progress in Digital Integrated Electronics," *Technical Digest—IEEE Int'l Electron Devices Meeting*, vol. 21, 1975, pp. 11-13.
3. J.D. Meindl, Q. Chen, and J.A. Davis, "Limits on Silicon Nanoelectronics for Terascale Integration," *Science*, vol. 293, no. 5537, 2001, pp. 2044-2049.
4. J.A. Hutchby et al., "Emerging Research Memory and Logic Technology," *IEEE Circuits & Devices Magazine*, vol. 21, no. 3, 2005, pp. 47-51.
5. S. Luryi, J. Xu, and A. Zaslavsky, eds., *Future Trends in Microelectronics*, John Wiley & Sons/IEEE Press, 2007.
6. M.T. Bohr, R.S. Chau, and K. Mistry, "The High-κ Solution," *IEEE Spectrum*, vol. 44, no. 10, 2007, pp. 29-35.
7. J. Izydorczyk, "Three Steps to the Thermal Noise Death of Moore's Law," *IEEE Trans. Very Large Scale Integration (VLSI) Systems*, vol. 18, no. 1, 2010, pp. 161-165.
8. V.V. Zhirnov et al., "Limits to Binary Logic Switch Scaling—A Gedanken Model," *Proc. IEEE*, vol. 91, no. 11, 2003, pp. 1934-1939.
9. R.P. Colwell, *The Pentium Chronicles: The People, Passion, and Politics behind Intel's Landmark Chips,* Wiley/IEEE CS Press, 2005.
10. G.E.R. Cowan et al., "A VLSI Analog Computer/Digital Computer Accelerator," *IEEE J. Solid-State Circuits*, vol. 41, no.1, 2006, pp. 42-53.

Jacek Izydorczyk *is an assistant professor of electronics at Silesian University of Technology, Gliwice, Poland. His research interests include network invariants, circuit simulations, magnetic media modeling, signal processing, speech synthesis, speech modeling, and analog filters. Izydorczyk received a PhD in electronics from Silesian University of Technology. He is a senior member of IEEE. Contact him at jacek.izydorczyk@ieee.org.*

Michael Izydorczyk *is an undergraduate student studying genetics at ETE in Gliwice, Poland. His research interests include computer simulation, computational genomics, and bioinformatics. Contact him at captainoats@o2.pl.*