1.1 Moore's Law: A Path Going Forward

William M. Holt

Executive Vice President, General Manager,
Technology and Manufacturing Group, Intel, Hillsboro, Oregon

1. Moore's Law Guides the Semiconductor Industry

Moore's Law has served as the guiding principle for the semiconductor industry for 50 years. The societal impact brought about by continually increasing the capability, affordability, and availability of integrated circuits is astonishing. Increasing computing power, increasing energy efficiency, and decreasing size of ICs have revolutionized existing industries and enabled new ones. While concerns over the future of Moore's Law have been present from the beginning, the challenges have changed over time. Currently, the challenges of scaling and increasing cost are the focus of concerns and doubts over the vitality of Moore's Law going forward. However, the industry continues to drive progress on many fronts. Continued advances in CMOS technology such as the introduction of 3D transistors provide increasing capabilities over a broader range of use. Novel 3D-heterogeneous-integration schemes and novel package technologies will further extend product benefits. At the same time, new memory technologies provide opportunities to fundamentally change memory hierarchy and bandwidth to resolve challenges in processor performance and power. In addition to the scaling of CMOS technology beyond 14nm, there are leading technology options on the horizon beyond CMOS with potential design benefits that can advance Moore's Law well into the future.

2. Economics is the Driving Factor

In 2015, Gordon Moore reflected on his 1965 paper, "The message I was trying to get across was that integrated circuits were the road to less-expensive electronics. It really evolved from being a measure of what goes on in the industry to something that more or less drives the industry." [1] Fundamental driving factors of Moore's Law are cost and capability. The drivers to enhance capability have changed over time as focus changed from DRAM to high-performance processing to power efficiency and lower-power SoCs, but the need to be cost effective remains.

Cost-per-transistor is an effective way to measure cost while accounting for differences in features. Cost-per-transistor can be determined by multiplying cost-per-area by area-per-transistor as shown in Figure 1.1.1 [2]. Cost-per-area has increased with successive technology generations; most recently, this increase has accelerated. Cost increases are driven by the need for more-complex processing, more-advanced tools, and increased numbers of process steps required to achieve the feature size as well as to incorporate new materials and architectures to achieve performance and energy-efficiency goals. In order to offset these cost increases, corresponding increases in density are required. Recently, higher-than-historical-density increases have been achieved through architectural innovations (such as 3D transistors) and implementing design rule and process improvements (such as multi-patterning lithography). This density acceleration has enabled the continuation of the historical cost-per-transistor reduction trend as shown in Figure 1.1.1. The resulting decreases in cost-per-transistor can be allocated as cost reduction or as capability increase, depending on the targeted application.

Technology-node timing has also been discussed in the context of Moore's Law. The timing of node introduction has not been uniform over the history of Moore's Law, with some node introductions earlier and some later [3]. In the 1970s, the interval between node introductions was commonly three years, then varying from one to five years prior to a concerted effort in the 1990s to achieve a two-year cycle. In addition, node conversion and technology introductions have not been simultaneous across the industry. While node timing will continue to vary, the driving force for node introduction continues to be the ability to deliver higher capability and better power performance cost effectively over time.

3. Silicon Technology Scaling

Transistor scaling continues to provide higher performance, lower power, and lower cost-per-transistor. Comparing an Intel® Core™ i5 processor on Intel's 14nm process to the 4004, the first commercially available microprocessor, shows a 3,500× increase in performance, 90,000× increase in energy efficiency, and 60,000× decrease in price per transistor [4].

Innovations in process and devices have been essential to continuing Moore's Law as shown in Figure 1.1.2 [5]. Process technology has implemented continual transitions from bipolar to MOSFETs, to CMOS to voltage scaling, to power-efficient scaling, and to System-on-Chip design. Significant processing innovations in CMOS include tungsten plugs, trench isolation, CMP (chemical mechanical polishing), copper interconnects, strained silicon, high-κ/metal gates and FinFETs. The introduction of strained silicon improved drive current. High-κ/metal gates reduced current leakage and heat. FinFETs addressed limitations of electrostatics and short-channel effects.

A snapshot of current silicon-technology scaling can been seen using Intel's 14nm process as an example. The 14nm process incorporates second-generation FinFETs, air-gapped interconnects, self-aligned double patterning, and a 0.05μm^2 SRAM cell size. Figure 1.1.3 shows the improvement in PMOS I_{dsat} and I_{dlin} for Intel's 14nm technology node compared to that for 22nm. Note that I_{dsat} is improved 15% for NMOS and 41% for PMOS over 22nm. Air gaps at critical performance layers provide a 17% improvement in capacitance. A thick top metal is used for improved on-die power distribution. Fin pitch, a key measure of transistor density for FinFETs, is scaled to 42nm, maintaining the historical 0.7× scaling trend from 22nm [6] [7].

While there are many measures of the improved capability that advancing technology provides, possibly the most important is the trend in improving power efficiency. This can be best seen by graphing energy times delay (CV^2 * CV/I = C^2V^3/I) over several generations. Figure 1.1.4 shows improvements in power efficiency with successive generations. Scaled process technologies provide a combination of higher performance and lower power.

An example of how this translates into the product can be seen with Core™ M and fifth generation of Core™ processors on 14nm process technology (code named Broadwell). Over a 2× reduction in total dissipated power (TDP), 60 to 80% higher graphics performance and a 60% reduction in SoC idle power versus the 4th generation Core™ product was achieved using key design optimizations and capabilities provided by the new process technology [8].

4. Scaling of Analog Circuits, Mixed-Signal Circuits, and SRAMs

Despite a general trend toward making integrated analog components more "digital", high-quality analog features must be maintained in scaled CMOS process nodes. Traditionally analog components such as PLLs, I/Os, and thermal sensors have gradually converted some analog functions to digital as the power and area cost of logic has reduced with scaling. But most of these digitally-assisted analog architectures still rely on high-quality analog building blocks such as linear amplifiers, DACs, and regulators. As an example, within Intel's 22nm and 14nm technology nodes, transistors and passives were engineered to maintain or improve analog capability while dramatically scaling analog area. The conversion to a FinFET architecture reversed the trend of degrading transistor intrinsic gain (g_m*r_o) and V_t variation starting in 22nm and continuing into 14nm as shown in Figure 1.1.5 [6]. The 14nm semi-digital delay-locked loop described in [9] illustrates how selecting a co-optimized combination of core analog components (delay line, regulator, R-C filter, etc.) along with digital control and calibration, results in better performance, power, and area at scaled process nodes.

In some ways, blocks that are considered traditionally "analog" such as serial I/Os, push the process performance harder than the CPU core or memory. Per-pin I/O data rates double about every four years to keep pace with aggregate system bandwidth requirements [10]. By contrast, microprocessor clock rates have increased relatively slowly over the past decade, instead, emphasis has been on more-power-efficient parallel architectures. As a result, today's serial I/O pin rates such as 8Gb/s PCIe Gen3 or 25Gb/s Ethernet now far outpace the core clock rate. Even the transfer rate of memory interfaces such as DDR will exceed the core frequency within a few years. As shown by [11] and [12] (Figures 1.1.6, 1.1.7), high-speed interfaces continue to benefit from process scaling in terms of density and power. Designing high-speed interfaces in scaled processes clearly requires careful consideration of layout effects within the transistors and interconnect stack. But the continued improvement in area density and power is essential to keep pace with aggregate system bandwidth requirements, which are growing even faster than per-pin bandwidth [12].

SRAM remains the workhorse as the embedded memory for all VLSI applications. Continuous voltage scaling for power efficiency has created a significant challenge in SRAM design to achieve lower operating voltage. The minimum operating voltage of SRAMs is directly determined by the variation

 978-1-4673-9467-3/16/$31.00

control of underlying transistor technologies. Because of the significant improvement in transistor variation in the most advanced 14nm FinFET [6], a much improved SRAM V_{CCmin} has been achieved as shown in Figure 1.1.8 [7].

5. Memory and 3D Integration

Memory has always been an essential component of high-performance energy-efficient computing across entire platforms. To offset fundamental limitations of various memory technologies, memory hierarchy was established to address the overall system-level memory needs (as shown in Figure 1.1.9 [13]. The hierarchy often consists of on-die SRAM as high-speed cache memory, off-chip DRAM as main memory, and high-density NAND as storage, supplemented by a hard-disk drive. But, with ever-increasing demand for memory bandwidth to support new applications, such as high-resolution graphics and cloud computing, the traditional memory solution is no longer sufficient. Recently, in-package memory has proven to be an effective solution to bring high-memory bandwidth (>100MB/s) directly to compute engines [14]. Various memory technologies, including enhanced DRAMs and logic-based high-performance eDRAMs, have been used to serve as the memory core. Currently, a technology breakthrough in high-density nonvolatile memory, called 3D XPoint™ [13], will provide a new way forward in terms of memory-hierarchy optimization, as this technology has blurred the traditional boundary between memory and storage with DRAM-like speed and NAND-like density. Accordingly, a system-level optimization will be needed to take full advantage of new memory technologies moving forward.

While monolithic integration driven by Moore's Law continues to provide the fundamental path to achieve higher performance and lower cost, in-package and 3D integration can also assist in providing a good balance in achieving performance and size improvements for various applications. 3D die-stacking has now been widely adopted by memory manufacturers as an economical solution to achieve high memory density. But logic-to-logic and/or logic-to-memory integration still remains an open field. Silicon interposers have been explored to provide multi-chip integration where data bandwidth is crucial. However, the cost overhead of the conventional silicon interposer scheme has been rather prohibitive for broad adoption of this technology. A new chip-level integration scheme, Embedded Multi-chip Interconnect Bridge (EMIB) shown in Figure 1.1.10 [15], has been developed to provide a much improved cost-performance trade-off. In this approach, a tiny silicon bridge is embedded in the package substrate to provide a dense chip-to-chip interconnect for high data bandwidth while keeping the cost adder to a minimum.

As new memory technologies emerge and new chip-to-chip and 3D integration schemes are introduced, it will become more important to re-optimize the overall system-level architecture and configurations to achieve optimal performance and economic benefits.

6. Future CMOS Transistor Technologies

Transistors have been the core of Moore's-Law semiconductor technology for over half a century. During this time, the transistor has evolved with technologies over the years, from bipolar to silicon-gate PMOS, through NMOS to CMOS. Recently, the continuing pursuit for better electrostatics has led to a major revolution with the introduction of the 3D transistor or FinFET device. Further improvement can potentially be achieved using a gate-all-around or Si nano-wire structure. The ever-increasing demand for faster switching has also driven the need for better conduction materials, such as Ge-channel for PMOS and III-V channel for NMOS or Ge for both NMOS and PMOS channels [16]. Carbon nanotubes (CNT) and 2D semiconductor materials are showing early potential for transistors. (Figure 1.1.11). However, these are still MOSFET transistors and will eventually be a limiter in how low the supply voltage can scale.

7. Beyond CMOS Devices

New devices that do not depend on traditional carrier transport have emerged over the past fifteen years or so. They have shown early promise as a way to help further advance Moore's Law in the future. Two, worthy of further consideration are the tunnel FETs (TFETs) and spintronic devices.

In order to lower power consumption in future technologies, there is a need to lower supply voltage (V_{DD}) to reduce switching energy (~CV_{DD}^2) while keeping leakage currents low. However, when CMOS supply voltage is scaled aggressively to sub-0.5V, performance suffers significantly since the physics-limited subthreshold slope does not permit scaling of threshold voltage (V_t) without unacceptable increases in leakage. In contrast, the tunnel FET transistor can use steeper subthreshold slopes to lower V_t, thus enabling supply voltage operation lower than that for CMOS.

Since the first few experimental TFETs have demonstrated subthreshold slope (SS) steeper than the MOSFET limit (SS = 60mV/decade at room temperature) [17 to 19], researchers have tried to improve the steepness of SS and increase the TFET on-current. Although TFETs can be built using conventional Si as channel material, Si realizes only a very low on-current and also the steep slope occurs at very low current levels, due to its large indirect bandgap and carrier mass. Thus, III-V materials have attracted interest as a TFET channel option due to improved material suitability (low direct bandgap and low carrier mass). The prospects for TFETs are further improved by using a heterojunction to lower the effective tunneling barrier [20] (Figure 1.1.12). The first experimental sub-60mV/decade SS III-V TFETs were demonstrated in 2011 [19]. Comparison to theory shows that the devices can be improved when scaled by removing parasitic currents and improving electrostatics by using a thin body structure [21].

One of the main challenges of realizing steep-SS TFETs is low defect (trap) density and thin-body geometry requirements. A comprehensive study of TFET geometry and defect effects has been carried out with experimentally-calibrated models [22]. Intrinsic-material band-to-band tunneling properties extracted from the experimental data and body thickness were found to be very critical parameters (Figure 1.1.13). Whereas existing bulk-material quality is found sufficient, oxide interface state density not exceeding $10^{12}cm^{-2}$ is required to realize steep SS. For a Ge TFET with body thickness >20nm, and oxide thickness >1nm, steep-SS is not expected even with an ideal oxide (Figure 1.1.14). Steep SS can be realized only for an aggressively scaled body ~5nm, due to strong double-gate control of the channel.

Comparing future N-TFETs to Si MOSFETs at L_G=13nm for various TFET material options has shown that each TFET has advantages over a MOSFET at different current levels (Figure 1.1.15) [23]. Circuit simulations using atomistic device models of a nanowire with L_G=13nm (ITRS 2018 technology node) were used to compare power-performance of CMOS and TFET logic. The results were also compared after adding the effects of device variation for MOSFETs and TFETs. Due to their steeper I_D-V_G curve characteristic, TFETs are more susceptible to higher leakage-current variation while for the MOSFET drive-current variation is larger. When variations are included, TFET logic still shows 54% better energy-efficiency than CMOS for the same performance (delay=40ps) (Figure 1.1.16). Clearly, there is potential in TFETs, but there also remain challenges that will require innovation for solution. [24]

Another class of beyond CMOS devices [25], is called spintronic logic. These devices use magnetization of a nanoscale ferromagnet to hold the logic state. Some of the device concepts are based on current-controlled switching by spin torque, which is the effect of spin polarized electrons changing magnetization as in Figure 1.1.17. Spintronic devices differ in geometric structure and the way the spin torque switching works: "All-spin" logic [26] utilizes spin-polarized current produced by one nanomagnet to switch the next. In "charge-spin" logic [27] magnetization controls charge current via tunneling magnetoresistance (TMR). This current is used to switch another nanomagnet using the Spin Hall Effect (SHE). Both "domain-wall" logic [28], and "mLogic" [29], are based on the TMR effect, as well, but the charge current moves domain walls with spin torque and thus switches magnetization. "Spin-torque-oscillator" (STO) logic [30] is based on synchronization of phases of STOs.

Other device concepts are based on voltage-controlled switching of magnetization using magnetoelectric (ME) effects. Such effects are, for example, magnetic exchange bias created by an adjacent anti-ferromagnet, change of magnetic anisotropy by strain of an adjacent piezoelectric material, or voltage controlled surface anisotropy as in Figure 1.1.17. In the spin majority gate [31, 32] magnetizations of three inputs are switched by the ME effect, which results in domain walls propagating in ferromagnetic wires. These domain walls compete to determine magnetization direction at a single output. Spin wave devices [33] form a majority gate as well, but the logic states of inputs propagate as spin waves in ferromagnetic wires and switch the nanomagnet at the output.

These new devices have demonstrated excellent potential in addressing many key challenges of conventional scaling, including switching-energy efficiency and leakage power. In general, ME devices have about two orders-of-magnitude smaller switching energy and similar switching delay to that for spin-transfer-

torque devices. However, while magnetoelectric devices have smaller switching energy, they are about two orders-of-magnitude slower in their current form than conventional CMOS at the same technology node [34]. The energy advantage of spintronics stems from the lower switching voltage and the collective nature of switching magnetization rather than numerous individual electrons. As the demand for lower power continues, a place for these devices may emerge.

Another attractive feature of spintronics is that its elements are non-volatile (that is the computation state remains unchanged when the power is switched off). This may open up a new way to architect future computing for a broad range of applications at very low energy, well below that achievable in CMOS.

The experimental realization of many of these spintronic devices has only been demonstrated recently, while others are still in simulation stage. Spintronics is rapidly increasing in its breadth and depth of novel approaches that use magnetism for logic and memory.

These new magnetic devices have demonstrated excellent potential in addressing many key challenges of conventional scaling, including switching-energy efficiency and device-leakage power. While such new devices are encouraging in the search for advancing Moore's Law well beyond the current technology, Si-based CMOS technology will continue to have a significant lead in overall performance and product-level integration. Thus, incremental enhancements in current Si-based technology will likely dominate the mainstream semiconductor industry until new materials and/or devices reach sufficient maturity to meet high-volume manufacturing needs.

8. Conclusion

Semiconductors continue to be the foundation for computing and communications solutions, the basis of the Internet of Everthing, and the primary driver in the future of electronics applications. Moore's Law has led to evermore-powerful smart phones, tablets, personal computers, and data centers. It has enabled computing to become a seamless and powerful force in our homes, offices, cars, factories, and much more. Much has been written about the end of Moore's Law. More recently, speculation has focused on the economic end of Moore's Law. Gordon Moore initially projected 10 years of visibility. [35] Over fifty years later, the Moore's Law horizon remains around 10 years. Moore's Law was never guaranteed. It has thrived and will continue to do so as the result of continuous innovation, rigorous planning, and technology execution. Even though it is getting more expensive to build wafers, improvements in density can provide real cost reduction at the most fundamental level, and this economic benefit drives the ability to continue investing in Moore's Law. Innovations have driven Moore's Law through numerous technological transitions and will continue to power us into the future of CMOS and beyond. As long as there is a cost benefit and rich options for future innovations there is no reason to predict an early end!

Acknowledgments:

The author thankfully acknowledges significant contributions from Ian Young, Kevin Zhang, and Doris Burrill during the preparation of this paper. The author sincerely appreciates the breadth of technical information and material that was received from many colleagues at Intel Corporation.

References:

[1] Gordon Moore, "50 Years of Moore's Law", Intel.com video, April, 2015 http://www.intel.com/content/www/us/en/silicon-innovations/50-years-of-moores-law-video.html.

[2] W. Holt, Intel Investor Meeting, November 20, 2014 http://intelstudios.edgesuite.net/im/2014/index.html

[3] R. Goodall, et al., "Long-Term Productivity Mechanisms of the Semiconductor Industry", Electrochemical Society Proceedings, p.133, 2002.

[4] "Celebrating 50 Years of Moore's Law – Whatever Has Been Done Can Be Outdone", Intel.com, April 2015, http://download.intel.com/newsroom/kits/ml50/pdfs/moores-law-50-years-infographic-entire.pdf.

[5] M. Bohr, "Moore's Law Challenges Below 10 nm", IEEE International Solid-State Circuits Conference panel session, 2015.

[6] S. Natarajan, et al., "A 14nm Logic Technology Featuring 2nd-Generation FinFET, Air-Gapped Interconnects, Self-Aligned Double Patterning and a 0.0588 µm2 SRAM Cell Size," IEDM Tech. Dig., pp. 15-17, 2014.

[7] E. Karl, et al., "A 0.6V, 1.5GHz 84MB SRAM Design in 14nm FinFET CMOS Technology", ISSCC 2015

[8] A. Nalamalpu et al, "Broadwell: A family of IA 14nm processors", 2015 Symposium on VLSI Circuits Digest of Technical Papers, pp. C314-C315, 2015

[9] A. Elshazly, et al., "A 2GHz-to-7.5GHz Quadrature Clock-Generator Using Digital Delay Locked Loops for Multi-Standard I/Os in 14nm CMOS," IEEE Symp. VLSI Circuits, pp. 52-53, June 2014.

[10] ISSCC 2015 Trends [Online], http://isscc.org/doc/2015/isscc2015_trends.pdf

[11] R. Dokania, et al., "A 5.9pJ/b 10Gb/s Serial Link with Unequalized MM-CDR in 14nm Tri-Gate CMOS," ISSCC Dig. Tech. Papers, pp. 184-185, 2015.

[12] J. Kim, et al., "A 16-to-40Gb/s Quarter-Rate NRZ/PAM4 Dual-Mode Transmitter in 14nm CMOS," ISSCC Dig. Tech. Papers, pp. 60-61, 2015.

[13] R. Crooke, "Non Volatile Memory", Intel Data Center Day, August 7, 2015 http://www.intc.com/events.cfm

[14] F. Hamzaoglu, et. al., "A 1Gb 2Ghz Embedded DRAM in 22nm Tri-Gate CMOS Technology", ISSCC Dig. Tech. Papers, pp. 230-231, 2014

[15] S. Rikhi, "Leading at the edge of Moore's Law with Intel Custom Foundry", p.32, IDF, 2014

[16] R. Kim, U. E. Avci, and I. A. Young, CMOS Performance Benchmarking of Si, InAs, GaAs, and Ge Nanowire n- and pMOSFETs with LG=13 nm Based on Atomistic Quantum Transport Simulation including Strain Effects, IEEE IEDM, Dec., 2015

[17] Appenzeller, J.; Yu-Ming Lin; Knoch, J.; Zhihong Chen; Avouris, P., "Comparing carbon nanotube transistors - the ideal choice: a novel tunneling device design," Electron Devices, IEEE TED, vol.52, no.12, pp.2568,2576, Dec. 2005

[18] W. Y. Choi; B.-G. Park; J.-D. Lee; T.-J. K. Liu, "Tunneling Field-Effect Transistors (TFETs) With Subthreshold Swing (SS) Less Than 60 mV/dec," IEEE EDL , vol.28, no.8, pp.743,745, Aug. 2007

[19] Dewey, G.; Chu-Kung, B.; Boardman, J.; Fastenau, J. M.; Kavalieros, J.; Kotlyar, R.; Liu, W. K.; Lubyshev, D.; Metz, M.; Mukherjee, N.; Oakey, P.; Pillarisetty, R.; Radosavljevic, M.; Then, H. W.; Chau, R., "Fabrication, characterization, and physics of III–V heterojunction tunneling Field Effect Transistors (H-TFET) for steep sub-threshold swing," IEEE IEDM, pp.33.6.1,33.6.4, 5-7 Dec. 2011

[20] Mohata, D.K.; Bijesh, R.; Mujumdar, S.; Eaton, C.; Engel-Herbert, R.; Mayer, T.; Narayanan, V.; Fastenau, J.M.; Loubychev, D.; Liu, A K.; Datta, S., "Demonstration of MOSFET-like on-current performance in arsenide/antimonide tunnel FETs with staggered hetero-junctions for 300mV logic applications," IEEE IEDM, pp.33.5.1,33.5.4, Dec. 2011

[21] Avci, U.E.; Hasan, S.; Nikonov, D.E.; Rios, R.; Kuhn, K.; Young, I.A., "Understanding the feasibility of scaled III–V TFET for logic by bridging atomistic simulations and experimental results," VLSI Technology (VLSIT) Symp., pp.183,184, June 2012.

[22] Avci, U.E.; Chu-Kung, B.; Agrawal, A.; Dewey; G., Le, V.; Rios, R.; Morris, D.H; Hasan, S.; Kotlyar, R.; Kavalieros, J; Young, I.A., "Study of TFET Non-ideality Effects for Determination of Geometry and Defect Density Requirements for Sub-60mV/dec Ge TFET" IEEE IEDM, pp.34.5.1,34.5.4, Dec. 2015.

[23] Avci, U.E.; Morris, D.H.; Hasan, S.; Kotlyar, R.; Kim, R.; Rios, R.; Nikonov, D.E.; Young, I.A., "Energy efficiency comparison of nanowire heterojunction TFET and Si MOSFET at Lg=13nm, including P-TFET and variation considerations," IEEE IEDM, pp.33.4.1,33.4.4, Dec. 2013.

[24] Avci, U.E.; Rios, R.; Kuhn, K.J.; Young, I.A., "Comparison of power and performance for the TFET and MOSFET and considerations for P-TFET," IEEE-NANO, pp.869-872, Aug. 2011.

[25] D. E. Nikonov and I. A. Young, ``Overview of beyond-CMOS devices and a uniform methodology for their benchmarking,'' Proc. IEEE, vol. 101, no. 12, pp. 2498-2533, Dec. 2013.

[26] B. Behin-Aein, D. Datta, S. Salahuddin, and S. Datta, "Proposal for an all-spin logic device with built-in memory", Nature Nano v. 5, p. 266 (2010).

[27] S. Datta, S. Salahuddin, and B. Behin-Aein, "Non-volatile spin switch for Boolean and non-Boolean logic", Appl. Phys. Lett. 101, 252411 (2012).

[28] J. A. Currivan, Y. Jang, M. D. Mascaro, M. A. Baldo, and C. A. Ross, "Low Energy Magnetic Domain Wall Logic in Short, Narrow Ferromagnetic Wires", IEEE Mag. Lett. (2012).

[29] D. Morris, D. Bromberg, J.-G. Zhu, and L. Pileggi, "mLogic: Ultra-Low Voltage Non-Volatile Logic Circuits Using STT-MTJ Devices", Proc. of DAC 2012, San Francisco, California, USA, June 3-7, 2012.

[30] I. Krivorotov and D. Markovic, presentation, NRI Benchmarking Workshop, Notre Dame University, Aug. 2011.

[31] D. E. Nikonov and G. I. Bourianoff, "Recent progress, opportunities and challenges for beyond-CMOS information processing technologies", ECS Trans., v. 35, no. 2, pp. 43–53, May 2011.

[32] D. E. Nikonov, G.I. Bourianoff, and T. Ghani, "Proposal of a Spin Torque Majority Gate Logic", IEEE Electron. Device Lett., v. 32, n. 8, pp. 1128-30, Aug. 2011.

[33] A. Khitun and K. L. Wang, "Nano scale computational architectures with Spin Wave Bus", Superlatt. and Microstruct. 38, 184 (2005).

[34] D. E. Nikonov and I. A. Young, "Benchmarking of Beyond-CMOS Exploratory Devices for Logic Integrated Circuits", IEEE J. Explor. Comput. Devices and Circuits, v. 1, p. 3, Apr. 2015.

[35] Gordon Moore, "Cramming more components onto integrated circuits", Electronics, Volume 38, Number 8, April 19, 1965.

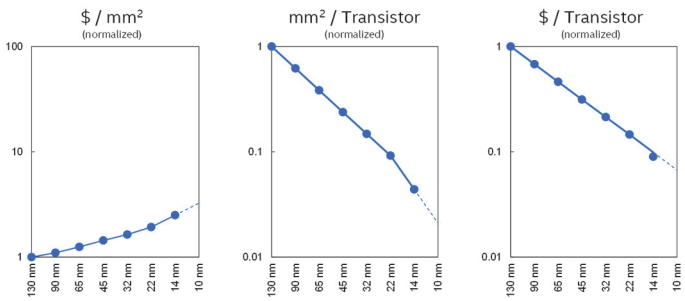

Figure 1.1.1: Cost/Area × Area/Transistor = Cost/Transistor [2].

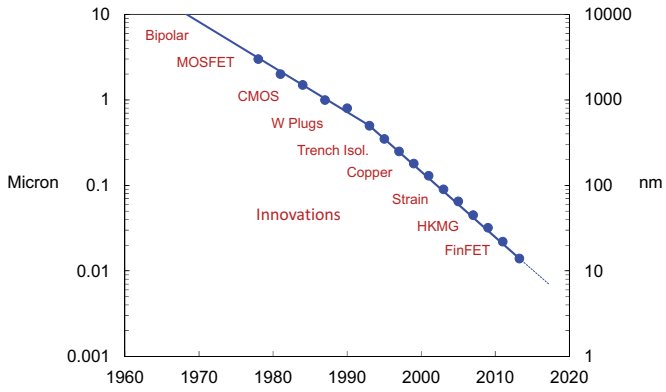

Figure 1.1.2: Process and Device Innovations Essential to Moore's Law [5].

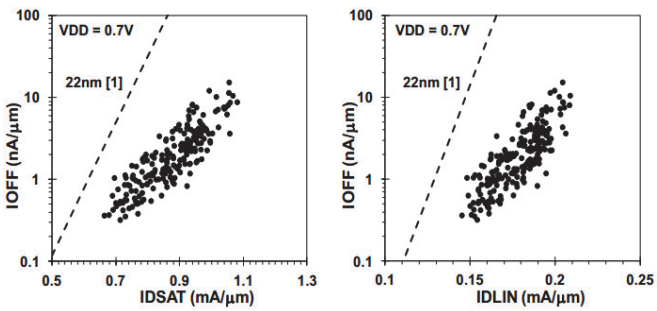

Figure 1.1.3: PMOS I_{dsat} and I_{dlin} curves for 14nm [6].

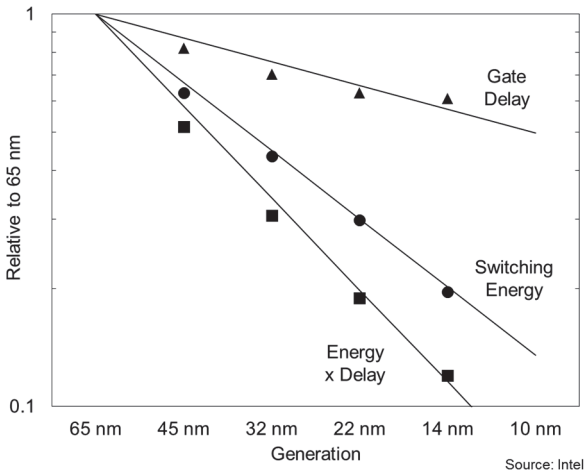

Figure 1.1.4: Generational Technology Benefits.

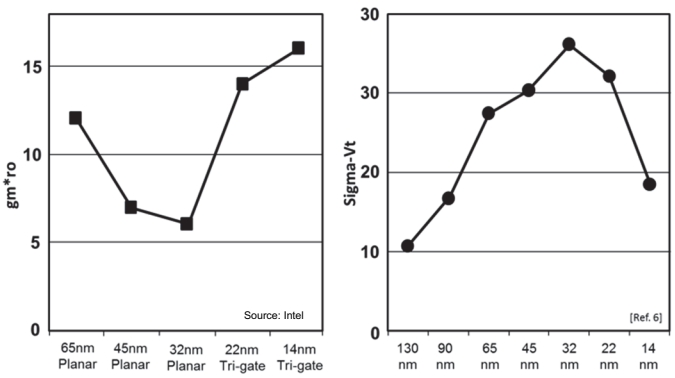

Figure 1.1.5: Intrinsic gain and Sigma-V_t scaling trends through 14nm.

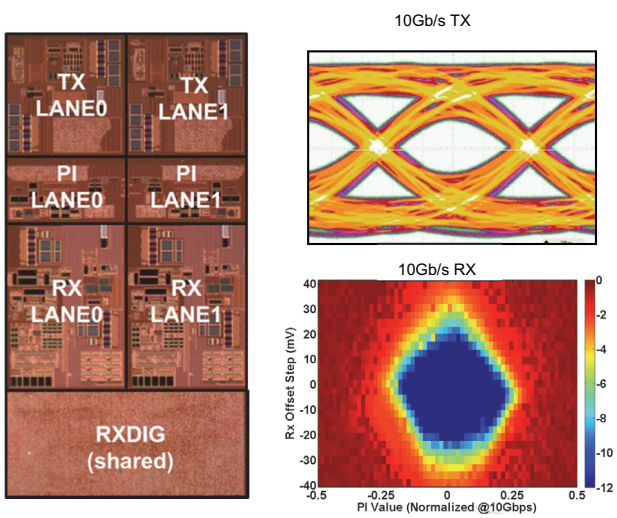

Figure 1.1.6: 10Gb/s serial I/O in 14nm CMOS [11].

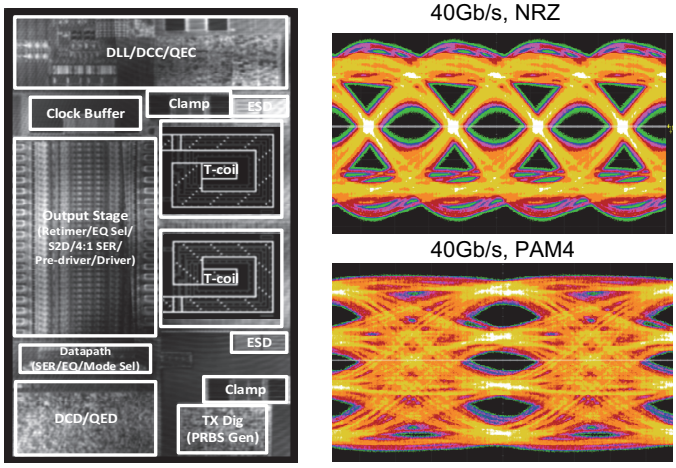

Figure 1.1.7: 16 to 40Gb/s NRZ/PAM4 transmitter in 14nm CMOS [12].

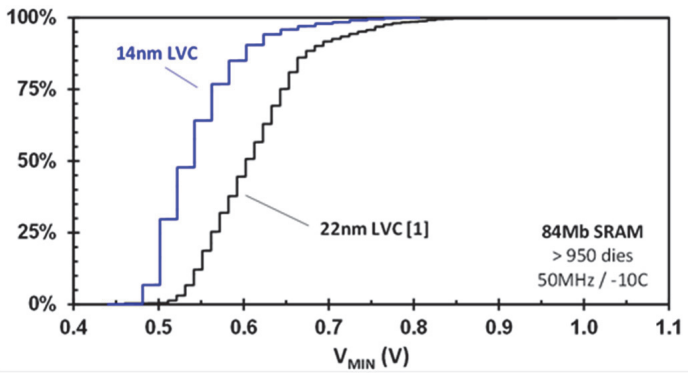

Figure 1.1.8: Improved SRAM V_{CCmin} [7].

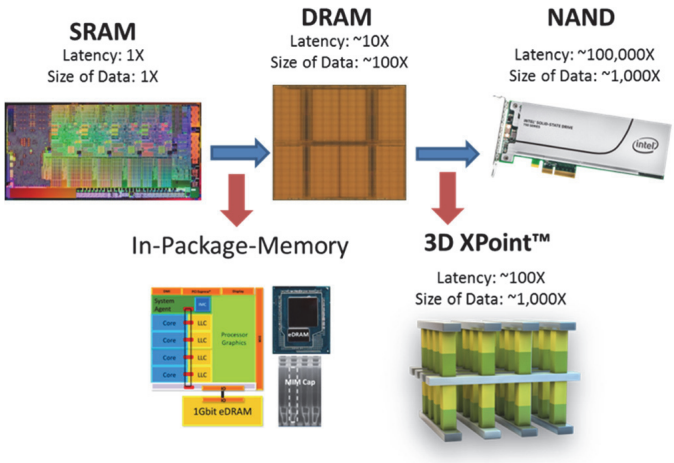

Figure 1.1.9: Memory Hierarchy [13].

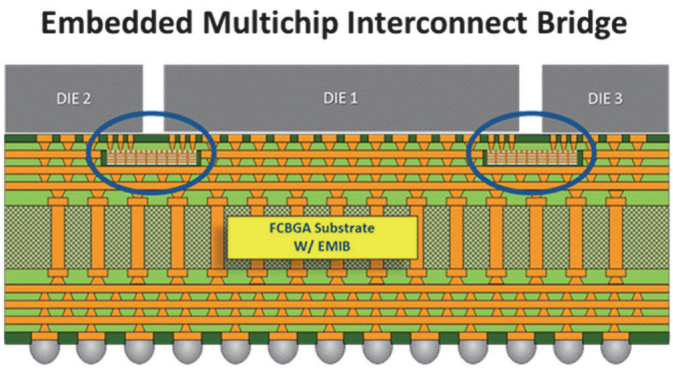

Figure 1.1.10: Embedded Multi-chip Interconnect Bridge (EMIB) [15].

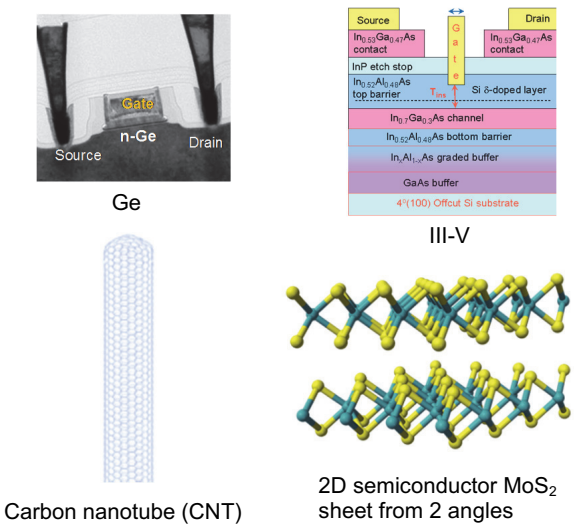

Figure 1.1.11: Transistors with New Materials.

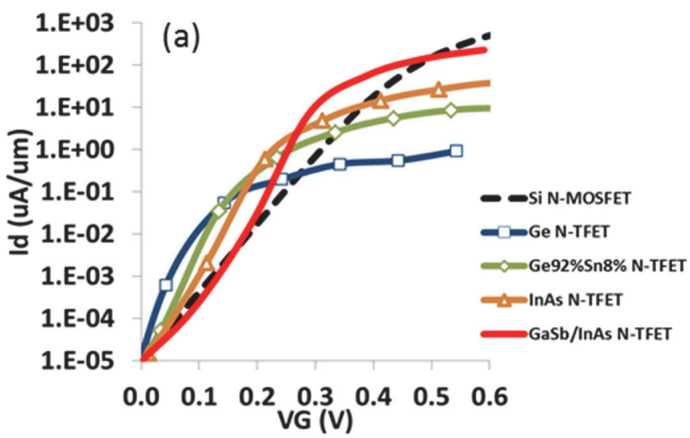

Figure 1.1.12: TFET I_D vs V_{GS} for a range of tunneling junction materials compared to Si MSOFET at L_G=13nm [23].

 978-1-4673-9467-3/16/$31.00 ©2016 IEEE

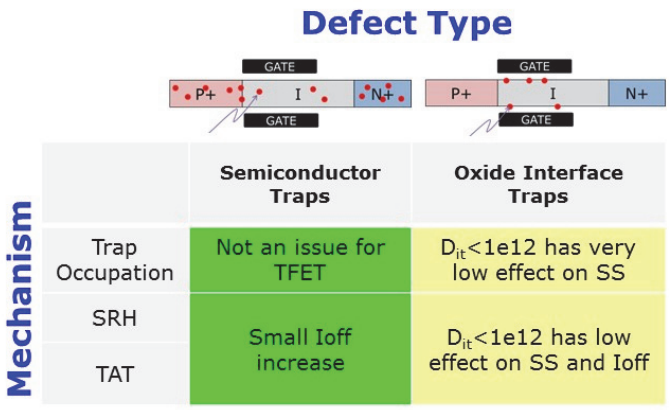

Figure 1.1.13: Summary of possible TFET defects/traps including SRH (Shockley-Read-Hall) and TAT (trap-assist-tunneling) and their impact on the SS I-V performance [22].

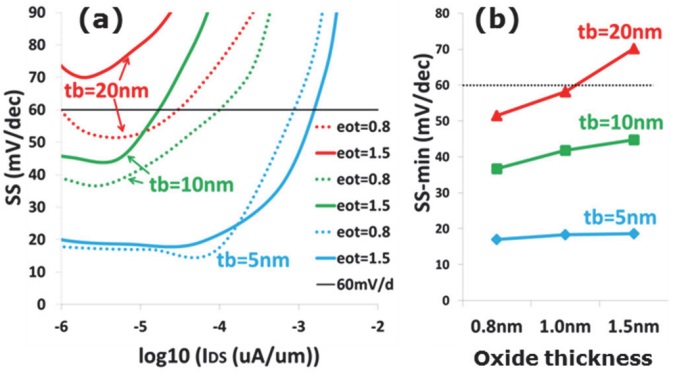

Figure 1.1.14: The effect of Ge TFET geometric dimensions on its SS I-V performance for a long L_G [22].

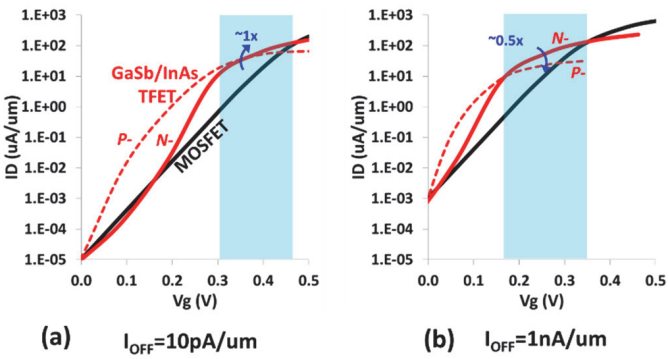

Figure 1.1.15: N-TFET and P-TFET compared to MOSFET devices: I_D vs V_{GS} for two I_{off} targets [23].

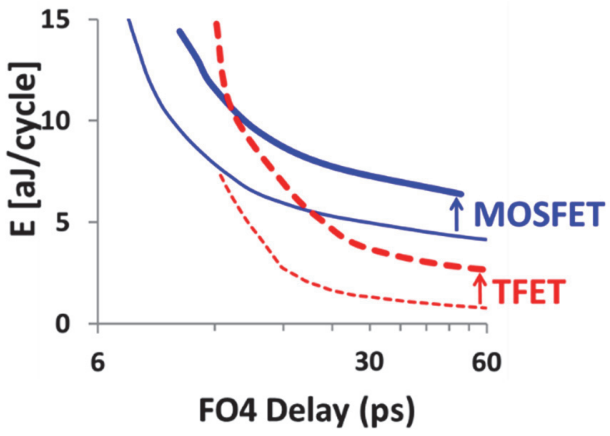

Figure 1.1.16: Energy vs Delay without (thin lines) and with (thick lines) device parameter variations [23].

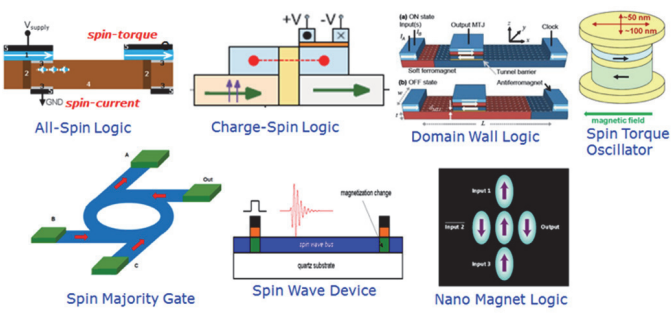

Figure 1.1.17: Schemes for various proposed spintronic logic devices: Spin torque switched devices (top row) and magnetoelectric switched devices (bottom row) [25].