Through the Looking Glass—Part 2 of 2

Trend Tracking for ISSCC 2013

The intent of the second part of this article is to share a sampling of the views held by the diverse group of experts represented by the International Solid-State Circuits Conference (ISSCC) 2013 Energy Efficient Digital; High Performance Digital; Technology Directions; Imagers, MEMS, Medical, and Displays (IMMD); and Memory program subcommittees. We published the findings of the subcommittees on Analog, Data Converters, RF, Wireline, and Wireless subcommittees in the Winter 2013 issue of this magazine.

Energy-Efficient Digital

The energy efficiency of digital circuits becomes increasingly important as larger and larger numbers of transistors are integrated on a single chip. Demand for ubiquitous mobile functionality for enhanced productivity, a better social-networking experience, and improved multimedia quality continues to drive technological innovation toward energy- and cost-efficient implementations. While the performance of embedded processors has increased to meet the rising demands of general-purpose computation, dedicated multimedia accelerators provide dramatic improvements in performance and energy efficiency for specific applications.

Digital Object Identifier 10.1109/MSSC.2013.2254632
Date of publication: 17 June 2013

Energy harvesting is another area of growing importance, leading to technologies that leverage nonvolatile logic-based systems-on-chips (SoCs) for applications that do not have a constant power source or for handheld devices with very limited battery capacity.

Technology scaling continues to be exploited to deliver designs capable of operating at lower voltages, resulting in reduced energy per operation, as well as reducing the area required to implement specific functions. Correspondingly, at ISSCC 2013, processors unveiled are built on a variety of technologies, with best-in-class results as measured by integration scale, performance/watt, and integration functionality. These include a few industry-first implementations demonstrated in various technologies ranging from 0.13 μm down to 28 nm bulk and SOI CMOS technologies.

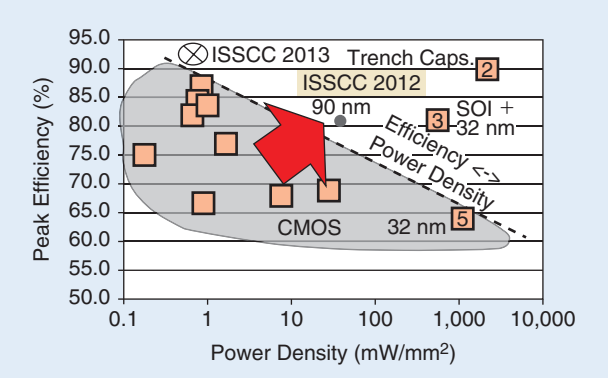

FIGURE 1: The development of integrated power conversion: Each year, more and more are integrated in standard CMOS technologies, optimizing efficiency versus power density. At ISSCC 2013, these trends are evident in a shift to higher performance, as shown by the arrow directed to the upper right. (Adapted, courtesy of ISSCC 2012 Paper 5.4.)

Emerging medical applications require a significant reduction in standby power, compared to state-of-the-art commercial processors. This has driven the exploration of new leakage-reduction techniques in both logic and on-chip memories, targeting orders-of-magnitude reduction in leakage currents. Fast wake-up time requirements drive the need for rapidly saving and restoring the processor state.

Figure 5 illustrates the main trends of energy relevant aspects of feature phones and smartphones. In the late 1990s, a GSM phone contained a simple RISC processor running at 26 MHz, supporting a primitive user interface. After a steady increase in clock frequency to roughly 300 MHz in the early 2000s, there was a sudden spurt toward 1 GHz and beyond. Moreover, following trends in laptops and desktops, processor architectures have become much more advanced, and recent smartphones incorporate dual- and even quad-core processors, running up to 2 GHz. Battery capacity, mostly driven by the required form factor, as well as thermal limits, imply a power budget of roughly 3 W for a smartphone. From this power budget, the RF power amplifier (for cellular communication) and the display are major drains. Overall, digital power consumption ranges from 2 W (peak) to 1 W (sustained). Thus, energy efficiency has become the

main challenge in designing application processors, graphics processors, media processors (video, image, audio), and modems (cellular, WLAN, GPS, Bluetooth). For video and image processing, the trend has been toward dedicated optimized hardware solutions. Some new areas where dedicated processors are particularly needed include gesture-based user interfaces and computational imaging. For all digital circuits, the limited power budget leads to the use of more-fine-grained clock gating, various forms of adaptive voltage-frequency scaling, a variety of body-bias schemes, and elaborate power-management strategies.

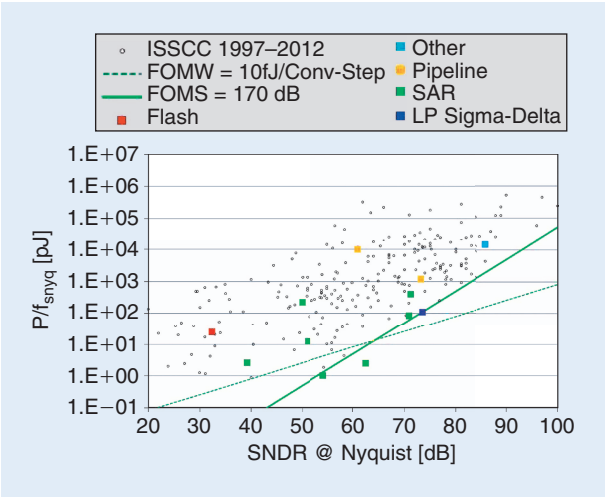

FIGURE 2: Power efficiency versus SNDR (highlighting ISSCC 2013 results).

Figure 6 shows the evolution of bit rates for wired and wireless links over time. Interestingly, cellular links, wireless LAN, as well as short links, consistently show a 10× increase in bit rate every five years, with no sign of abatement. Thus, with essentially constant power and thermal budgets, energy efficiency has become a central theme in designing the digital circuits involved in signal processing. Historically, CMOS feature sizes have halved every five years. For a brief period in the 1990s, CMOS scaling (aka Dennard scaling) provided a 2^3 (α^{-3}) increase in energy efficiency every five years, almost matching the required 10× power reduction. But during the past decade, CMOS scaling offers only a roughly 3× improvement in energy efficiency every five years. The resulting ever-widening gap has led to alternative approaches to improving energy efficiency, namely, new standards, smarter algorithms, more efficient digital signal processors, highly optimized accelerators, smarter hardware-software partitioning, as well as the power-management techniques mentioned previously.

Thus, we see a potential for a host of new applications of lower- and lower-power technology created in response to the recognition of the universal importance of efficient local processing of signals in diverse omnipresent areas.

High-Performance Digital

The relentless march of process technology brings more integration and performance to digital systems each year. At ISSCC 2013, for example, IBM's System z processor leads the charge with a 2.75 B transistor chip, operating at 5.7 GHz.

The chip complexity chart in Figure 7 shows the trend in transistor integration on a single chip over the past two decades. While the 1 billion transistor integration mark was achieved some years ago, we now commonly see processors with beyond 2 B transistors on a single die.

Leveraging sophisticated strategies to lower leakage and manage voltage, variability, and aging has bolstered the continuing reduction in total power dissipation. These strategies are helping rein in the increase in energy demands from PCs, servers, and similar systems. As power

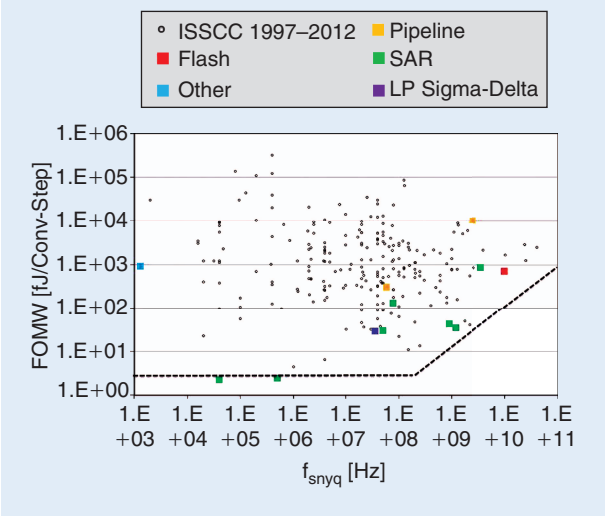

FIGURE 3: FoM (energy per conversion step) versus Nyquist bandwidth for various converters.

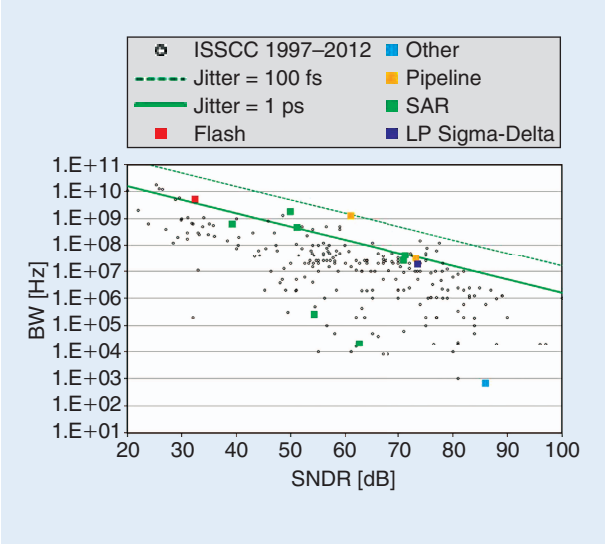

FIGURE 4: Bandwidth versus SNDR.

reduction becomes mandatory in every application, the trend toward maintaining near-constant clock frequencies also continues as shown in the frequency trends plot in Figure 8. This will yield solutions with lower cost and lower cooling demands, resulting in greener products for the future.

As well, processor designers are choosing to trade off performance by lowering supply voltage. The performance loss of reduced voltage and clock frequency is compensated by further increased parallelism. Processors with more than eight cores are now commonplace. This year, at ISSCC 2013, a 24-core processor was presented as noted in the core-count trend chart in Figure 9.

In addition to the trend toward integrating more cores on a single chip, single packages with multiple die are appearing. At ISSCC 2013, IBM presented a multichip module with six CPUs and two embedded DRAM cache chips. As well, dedicated coprocessing units for graphics and communications are now commonly integrated on these complex systems-in-package (SiP). Design of these SoCs and SiPs requires broad collaboration across multiple disciplines including circuits, architecture, graphics, process technology, package, system design, energy efficiency, and software. New performance- and power-efficient computing techniques continue to be introduced for targeted critical applications, such as floating point and SIMD.

As technology continues to scale to finer dimensions, large caches are being integrated within microprocessor die. Figure 10 shows the general trend of large cache integration.

Methods for communication within and between die are becoming increasingly important. These are being driven by two trends: 1) three-dimensional (3-D) integration is increasingly common and 2) interconnect delay becomes more dominant as processes scaled down. Designs emphasizing on-die interchip transport constitute a recent trend (see also wireline).

At ISSCC 2013, another trend in evidence was the continued emergence of all-digital phase-locked loops (PLLS) and delay-locked loops (DLLs) that better exploit nanometer-feature-size scaling while reducing power and area costs. Due to the application of highly innovative architectural and circuit design techniques, the features of these "all-digital" PLLs and DLLs have improved significantly over the recent past. Figure 11 shows the jitter performance versus energy cost for PLLs and MDLLs.

Overall, digital processors continue to grow in complexity, while more circuits are implemented using digital techniques to cope with variability and ease scaling to finer geometries.

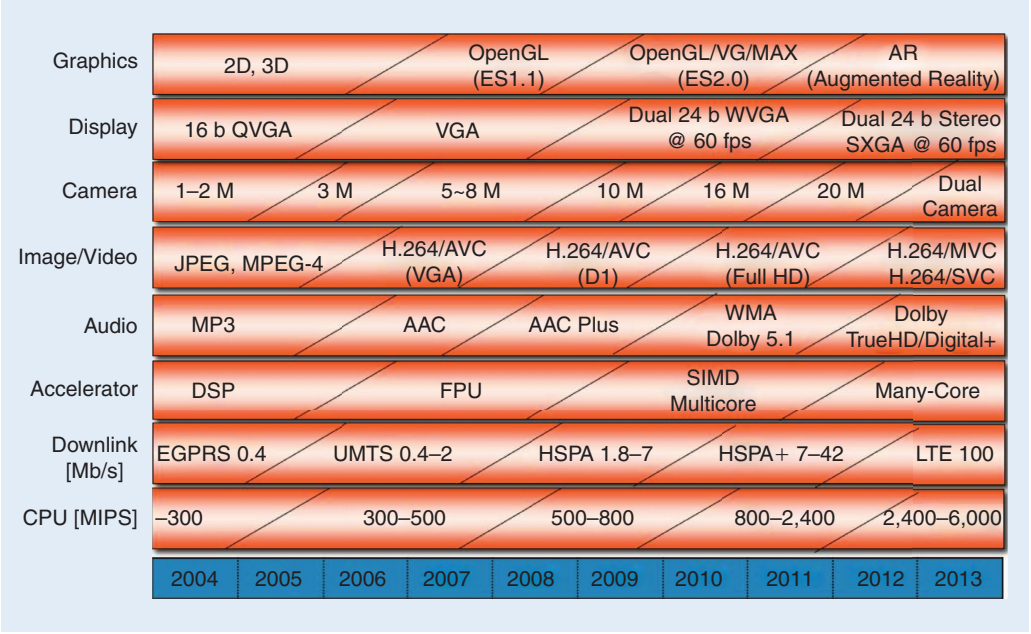

FIGURE 5: Developments in application processors for smartphones.

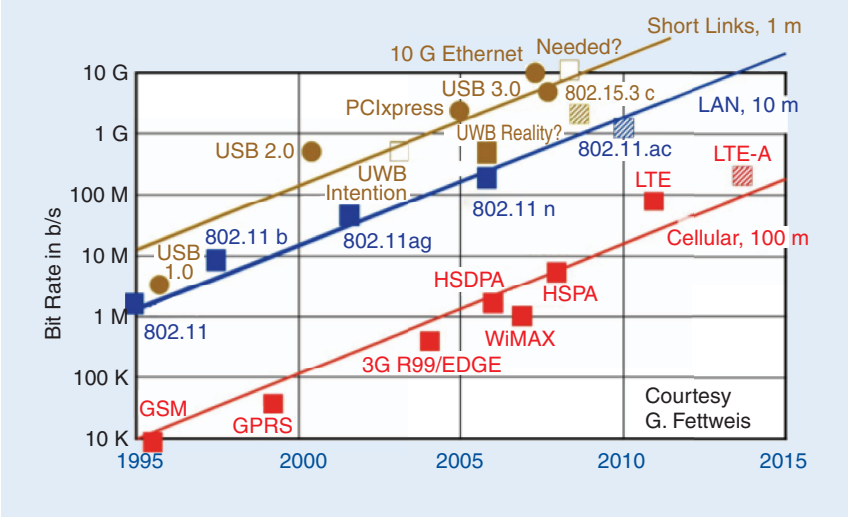

FIGURE 6: Wireless and wired data rates over time. (Image courtesy of G. Fettweis.)

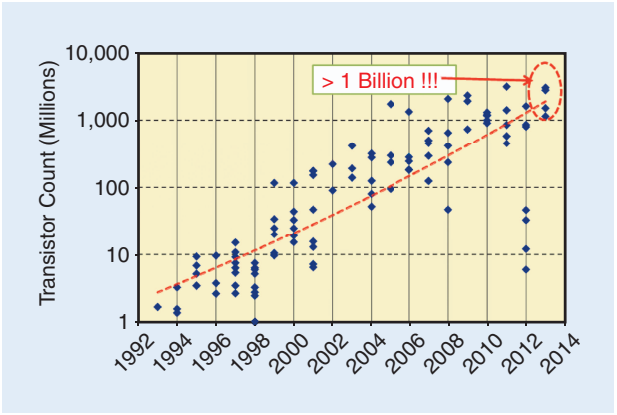

FIGURE 7: Chip complexity.

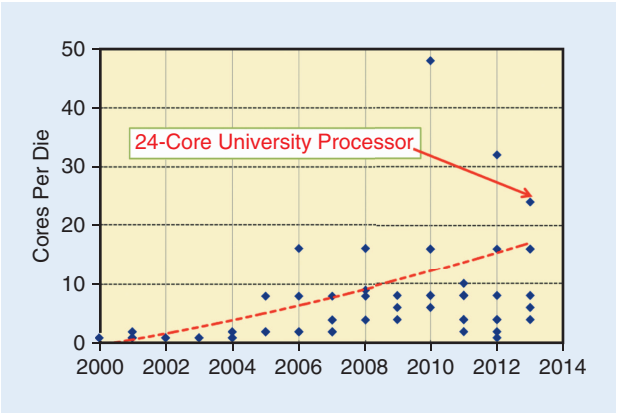

FIGURE 9: Core count.

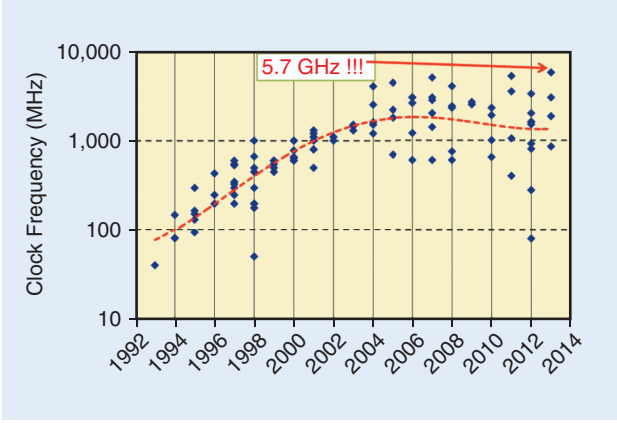

FIGURE 8: Clock frequency.

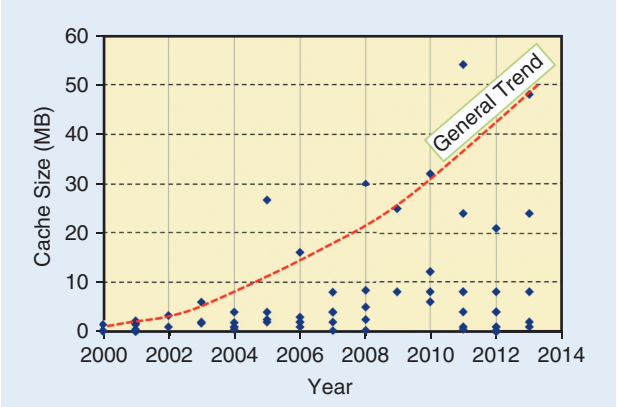

FIGURE 10: Total on-die cache.

IMMD

The common thread that unites the constituents of IMMD is a concern for interfaces with the physical world—both the human body and its environment. Conventionally, we consider this broad area in terms of both tool and task: imagers, MEMS, medical, and displays, as discussed below.

Imagers

Since 2010, there has been growth beyond expectations in the adoption of mobile devices, such as smartphones and tablets, that has induced larger volumes of CMOS image sensor chips. Both the resolution and miniaturization races are ongoing, and performance metrics are becoming more stringent. In addition to conventional pixel shrinkage, a "more than Moore" trend is increasingly evident: Resolution of over 20 Mpixels is commercially available for mobile devices employing enhanced small-size pixels. As a consequence of innovative readout and ADC architectures embedded at the column and chip levels, data rates approaching 50 Gb/s, and a noise floor below single electron have been demonstrated. In addition to conventional applications, ultra-low-power vision sensors, 3-D, high-speed, and multispectral imaging, are front-running emerging technologies.

Back-side illumination (BSI) is now the mainstream technology for high-volume, high-performance mobile applications. While 1.12 μm BSI pixels are currently available, the industry is potentially moving toward 0.9 μm pixel pitch and below. Additional innovative technologies outside of traditional scaling include advanced 3-D stacking of a specialized image-sensor layer on top of deep-submicron digital CMOS (65 nm 1P7M) using through-silicon vias (TSVs) and microbumps.

The importance of digital signal processing technology in cameras continues to grow to mitigate sensor imperfections and noise and to compensate for optical limitations. The level of sensor computation is increasing to thousands of operations per pixel, requiring high-performance and low-power digital signal processing solutions. In parallel with these efforts is a trend through out the image-sensor industry toward higher levels of integration to reduce system costs.

Ultra-low-power vision sensors are being reported in which more programmability and computation is performed at the pixel level to extract scene information, such as object features and motion.

Lightfield/plenoptic commercial cameras, which have been available since 2010, are gaining popularity and are being marketed for 3-D imaging and/or all-in-focus

two-dimensional imaging. On-chip stereoscopic vision has been demonstrated through digital micro lenses (DMLs), paving the way to next-generation passive 3-D imaging for mobile and entertainment applications, such as through gesture-control user interfaces.

Significant R&D effort is being spent on active-3-D-imaging time-of-flight (TOF) applications to support requirements from autonomous driving, gaming, and industrial applications, addressing open challenges, such as background light immunity, higher spatial resolution, and longer distance range. Deep-submicron CMOS single-photon avalanche diodes (SPADs) have been developed by several groups in various submicron technologies. They are now capable of meeting the requirements for high-resolution and high-timing accuracy by employing highly parallel time-to-digital converters (TDCs) and small pixel pitch with better fill factor.

Ultra-high-speed image sensors for scientific imaging applications having up to 20 Mf/s acquisition speed have been demonstrated.

Multispectral imaging is gaining a lot of interest from the image-sensor community: several research groups have demonstrated fully-CMOS room-temperature terahertz image sensors, and a hybrid sensor capable of simultaneous visible, IR, and terahertz detection has been reported.

The share of the market for CCDs in machine vision, compact DSCs for security applications, continues to shrink. Only for high-end digital cameras for astronomy and medical imaging do CCDs still maintain a significant market share.

MEMS and Sensors

MEMS inertial sensors are finding widespread use in consumer applications to provide enhanced user interfaces, localization, and image stabilization. Accelerometers and gyroscopes are being combined with 3-D magnetic-field sensors to form nine-degree-of-freedom devices, and pressure sensors will eventually add a tenth degree. The power consumption of such devices is becoming sufficiently low for the sensor to operate continuously, enhancing indoor navigation. There have been further advances in heterogeneous integration of MEMS with interface circuits to support increased performance, larger sensor arrays, reduced noise sensitivity, reduced size, and lower costs.

To address the stringent requirements of automotive, industrial, mobile, and scientific applications, MEMS inertial sensors, pressure sensors, and microphones are becoming more robust against electromagnetic interference (EMI), packaging parasitics, process-voltage-temperature (PVT) variations, humidity, and vibration.

Sensor interfaces achieve increasingly high resolution and dynamic range while maintaining or improving power or energy efficiency. This is achieved through techniques such as zooming, nonuniform quantization, and compensation for baseline values.

New calibration approaches, such as voltage calibration, are being adopted for BJT-based temperature sensors to reduce cost. In addition to thermal-management applications (prevention of overheating in microprocessors and SoCs), temperature sensors are also increasingly co-integrated with other sensors (such as humidity, pressure, and current sensors) and MEMS resonators for cross-sensitivity compensation. Alternative temperature-sensing concepts find their way into applications with specific requirements not easily addressed by BJTs: thermal diffusivity-based sensing for high-temperature applications; thermistor-based and Q-based concepts for in-situ temperature sensing of MEMS devices; and for ultra-low-voltage operation.

MEMS oscillators continue to improve. Phase noise is now low enough for demanding RF applications (12-kHz to 20-MHz integrated jitter is now below 0.5 ps); and frequency accuracy is now better than 0.5 p/m. Consumer applications are adopting such new low-power and low-cost oscillators.

Biomedical

There have been continuous achievements in the area of ICs for neural and biopotential interfacing technologies. Spatial resolution of neural monitoring devices is being reduced, utilizing the benefits of CMOS technology. IC providers are increasing their component offerings toward miniaturization of portable medical devices.

Telemedicine and remote-monitoring applications are expanding with support from IC manufacturing companies. Moreover, the applications of such systems are not limited to

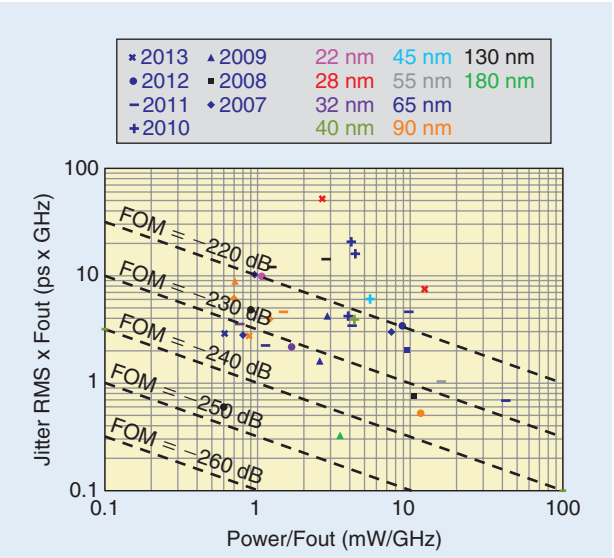

FIGURE 11: PLL and MDLL trends.

Developments reported at ISSCC 2013 continue to present breakthroughs in the broad domain of solid-state circuits and systems.

services targeted for elderly or chronically ill patients; for example, there are several technologies developed to enhance the way clinical trials are conducted by monitoring patient adherence and improving data collection. Low-power WiFi and Bluetooth low energy is emerging as a standard wireless connection between portable communication services and wearable technology.

Smart biomolecular sensing is another major trend that marries the solid-state and biochemical worlds together with the ultimate goal of enabling a more predictive and preventative medical care. With the help of the accuracy and parallelism enabled by CMOS technology, time, cost, and error rate of DNA sequencing may be significantly improved. Direct electronic readout may relax the need for complex biochemical assays. Similar trends are becoming increasingly evident in proteomics and sample preparation.

Even for medical imaging, there is a shift from hospital imaging toward point-of-care and portable devices. A key example is in portable high-resolution ultrasounds in which larger scientific imaging setups are being integrated onto the sensor through process technologies (such as integrated spectral filters, CMUT). Another example is in molecular imaging. The advent of silicon photo multipliers (SiPM), which provide a solid-state alternative to PMTs, enables the realization of PET scanners compatible with MRI, opening the way to new frontiers in the field of cancer diagnostics. More recently, SiPMs realized within deep-submicron CMOS technologies have allowed the integration at pixel- and chip-level extra features, such as multiple timestamp extraction, allowing a dramatic reduction in system cost.

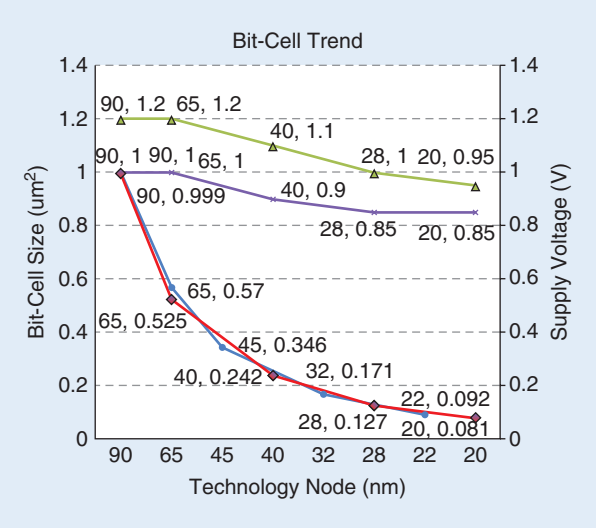

FIGURE 12: SRAM bit-cell size and supply scaling range from major semiconductor manufacturers.

Displays

The desire to put much higher-resolution and higher-definition displays into mobile applications is one of the display-technology trends that is now opening a full HD smartphone era. High-definition displays of 440 p/i are expected, even for 5-in display sizes. Low-temperature poly silicon (LTPS) technology seems to have more merit than amorphous-silicon thin-film-transistors (a-Si TFTs) technology. But a-Si TFT and oxide TFT technologies supported by compensating driver systems are beginning to compete with it. Very-large-size LCD TVs over 84 in, with UD (3,840 × 2,160) resolution are now the leading entertainment systems. Also opening new opportunities in consumer applications are 55-in AMOLED TVs with full HD resolution.

As touch-screen displays for mobile devices become increasingly thin, capacitive touch sensors move closer to the display. But the resulting in-cell touch displays come with reduced signal levels due to increased parasitics and increased interference from the display and switched-mode chargers. Noise immunity is improved by adopting noise filtering and new signal-modulation approaches.

Thus, overall, we see an expansion of diverse techniques evolving in aid the most important role of modern technology—to serve mankind.

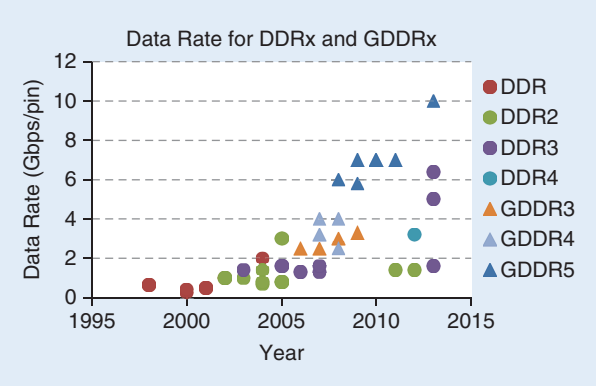

FIGURE 13: Trends in DRAM data rate/pin.

Memory

Development in mainstream memory technologies continues unabated. We see progressive sustained scaling in embedded SRAM, DRAM, and floating-gate-based Flash for very broad applications. However, due to the major scaling challenges in all mainstream memory technologies, we see a continued increase in the use of smart algorithms and error-correction techniques to compensate for increased device variability. In further response to these challenges, we see logic processes adopting FinFET devices along with read- and write-assist circuits in SRAMs. Meanwhile, emerging memory technologies are making steady progress toward product introduction, including PCRAM and ReRAM, while STT-MRAM is beginning to become a strong candidate for both standalone and embedded applications.

SRAM

Embedded SRAM continues to be a critical technology enabler for a wide range

of applications from high-performance computing to mobile utilization. The key challenges for SRAM include V_{CCmin}, leakage and dynamic power reduction while relentlessly following Moore's law to shrink the area by 2× for every technology generation. As the transistor feature size has marched to under 30 nm, device variation has made it very difficult to shrink the bit-cell size at the 2× rate while maintaining or lowering V_{CCmin} between generations. Starting at 45 nm, the introduction of high-k metal-gate technology reduces the V_t mismatch and further enables device scaling by significantly reducing the equivalent oxide thickness. Starting at 22 nm and beyond, new transistors such as FinFETs and fully depleted SOIs are key to enabling the continuous scaling of bit-cell area and low-voltage performance. Design solutions such as read/write-assist circuitry have been used to improve SRAM V_{CCmin} performance starting at 32 nm. New SRAM bit cells with more than six transistors have also been proposed to minimize operating voltage. For example, 8T register file cells have been reported in recent products requiring low V_{CCmin}. Dual-rail SRAM design emerges as an effective solution to enable dynamic voltage-frequency scaling (DVFS) by decoupling the logic supply rail from the SRAM array, thus allowing a much wider operating window. It is important for SRAM to reduce both leakage and dynamic power, keeping products within the same power envelope for succeeding feature-size reductions. Sleep transistors, fine-grain clock gating, and clockless SRAM designs have been proposed to reduce leakage and dynamic power. Redundancy and ECC protection are also keys to ensure yield and reliability when embedded SRAM products go into production. Figure 12 shows the trend in SRAM-bit-cell scaling on the left axis and the trend in SRAM supply-voltage scaling on the right axis, using data from major semiconductor manufacturers.

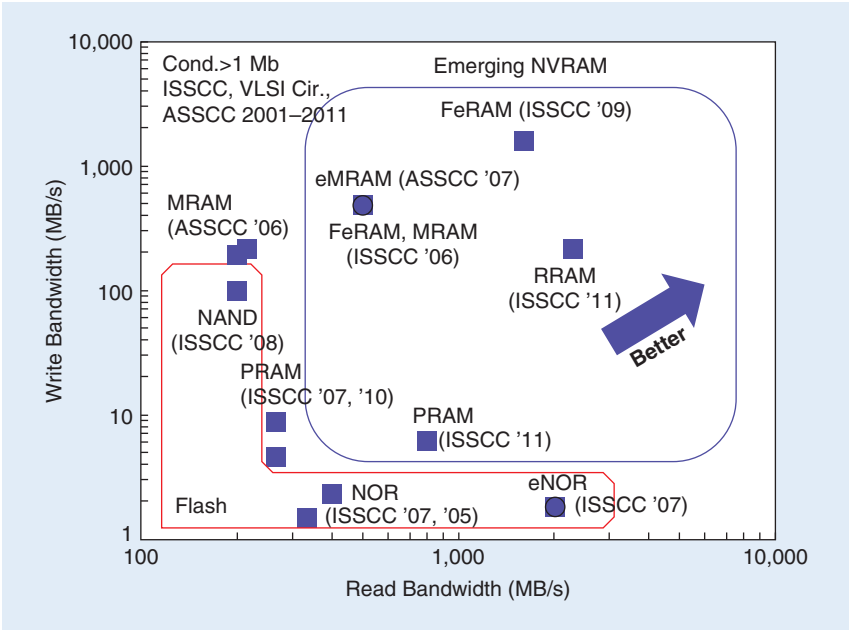

FIGURE 14: Read- and write-bandwidth comparison of nonvolatile memories.

High-Speed I/O for DRAM

To reduce the bandwidth gap between main memory and processor frequencies, external data rates continue to increase as conventional high-speed wired interface schemes (such as DDRx and

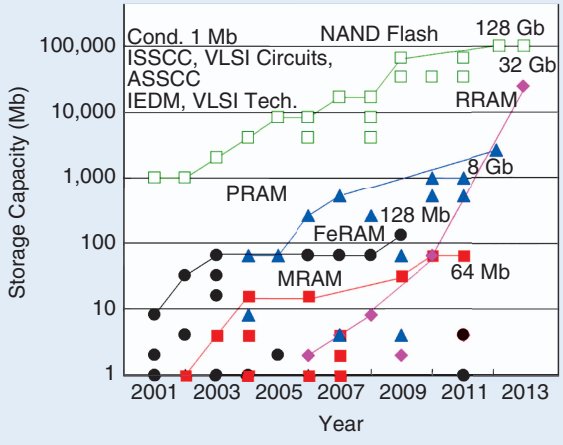

FIGURE 15: Memory capacity of emerging nonvolatile memories.

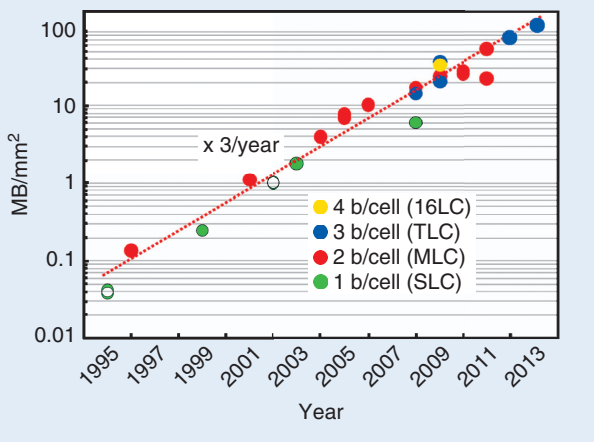

FIGURE 16: NAND Flash memory density.

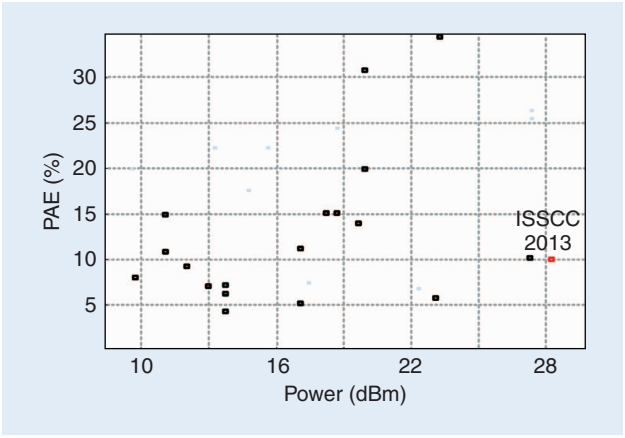

FIGURE 17: PAE versus output power for recent submicron mm-Wave CMOS PAs.

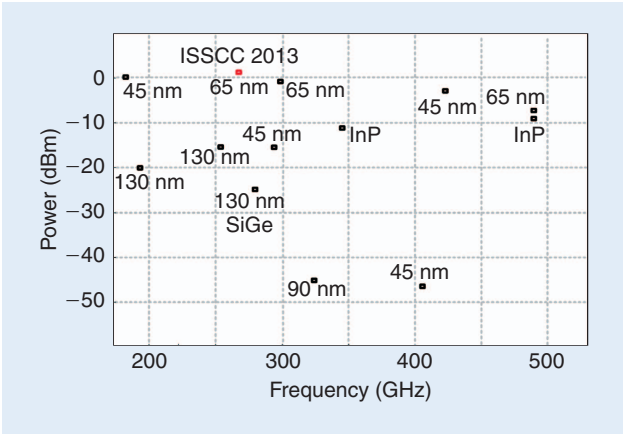

FIGURE 18: Output power versus frequency for mm-wave and sub-mm-wave sources.

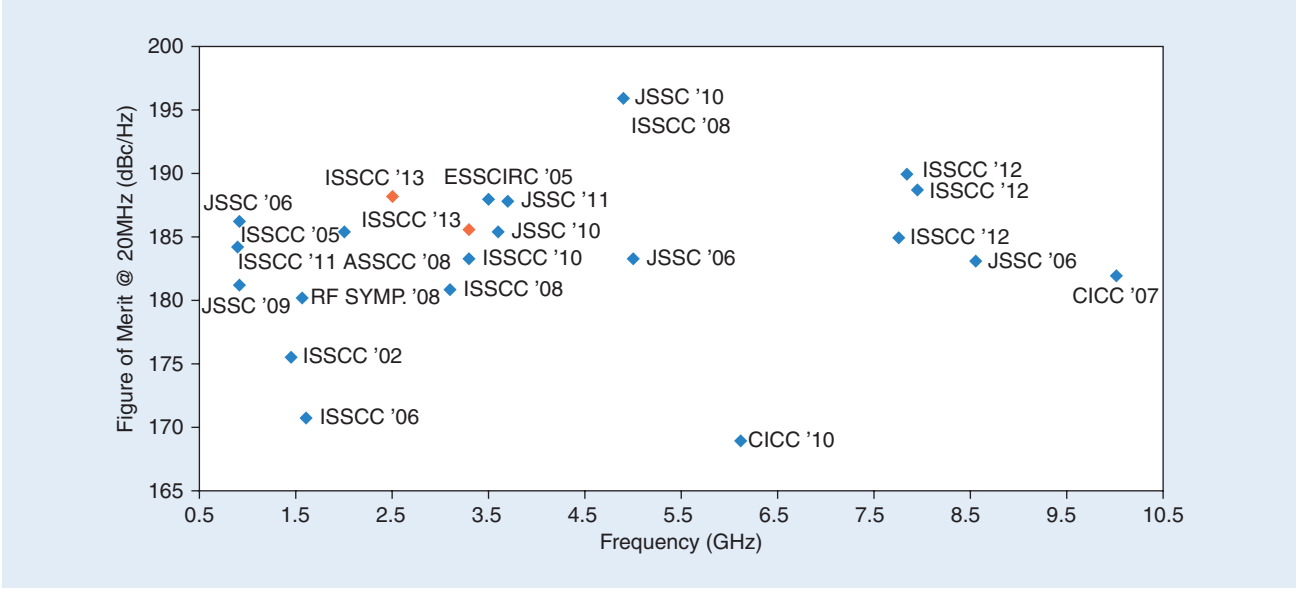

FIGURE 19: Phase-noise FOM at 20 MHz offset frequency versus oscillation frequency.

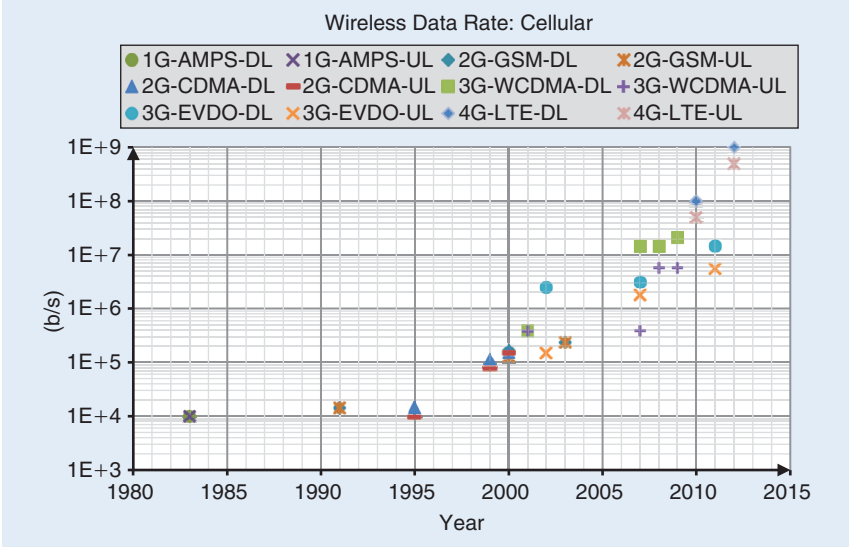

FIGURE 20: Data rates for cellular standards.

GDDRx for DRAM) evolve, as shown in Figure 13. Currently, GDDR5 and DDR4 memory I/Os operate around 7 Gb/s/pin and 3 Gb/s/pin, respectively. To achieve higher data-transfer rates, signal-integrity techniques (such as crosstalk, noise, and skew cancellation) and speed enhancement techniques (such as equalizers and pre-emphasis) have been developed. These advanced techniques have pushed I/O speeds toward 10 Gb/s/pin. Lower power consumption for data center and mobile applications has also been pursued. A near-ground signaling method, termination impedance optimization, decision feedback equalization, and clock-feathering

slew-rate control technologies have been demonstrated to reduce the power dissipation of memory interfaces significantly, while achieving high bandwidth.

Nonvolatile Memories

In the past decade, significant focus has been put on the search for emerging memories to provide a possible alternative to floating-gate nonvolatile memory (NVM). The emerging NVMs, such as phase-change memory (PRAM), ferroelectric RAM (FeRAM), spin-torque-transfer magnetic RAM (STT-MRAM), and resistive memory (ReRAM), are showing the potential to achieve high-cycling capability (operational lifetime) and lower power per bit for both read and write operations. Some commercial applications, such as cellular phones, have recently begun to use PRAM, demonstrating that reliability and cost competitiveness of emerging memories is becoming a reality. Fast write speed and low read-access time are being achieved in many of these emerging memory schemes. At ISSCC 2013, a 32 Gb ReRAM cross-point array was demonstrated in 24-nm technology. Figures 14 and 15 provide a summary of the scaling trends for both bandwidth and density of emerging memories.

NAND Flash Memory

NAND Flash memory continues to advance toward higher density and lower power, resulting in low-cost storage solutions that are enabling the replacement of traditional hard-disk storage with solid-state disks (SSDs). The use of multiple bits per cell has proven to be effective in increasing the density. Figure 16 shows the observed trend in NAND Flash capacities presented at ISSCC over the past 18 years. With scaling, device variability and error rates increase, requiring system designers to develop sophisticated control algorithms to compensate. Some of these algorithms are implemented outside the NAND silicon, in the system memory controller (in particular, the ECC and data management methods), for improved overall reliability. Possible future scenarios include 3-D-stacked NAND vertical gates as a solution to further increase overall NAND-memory density.

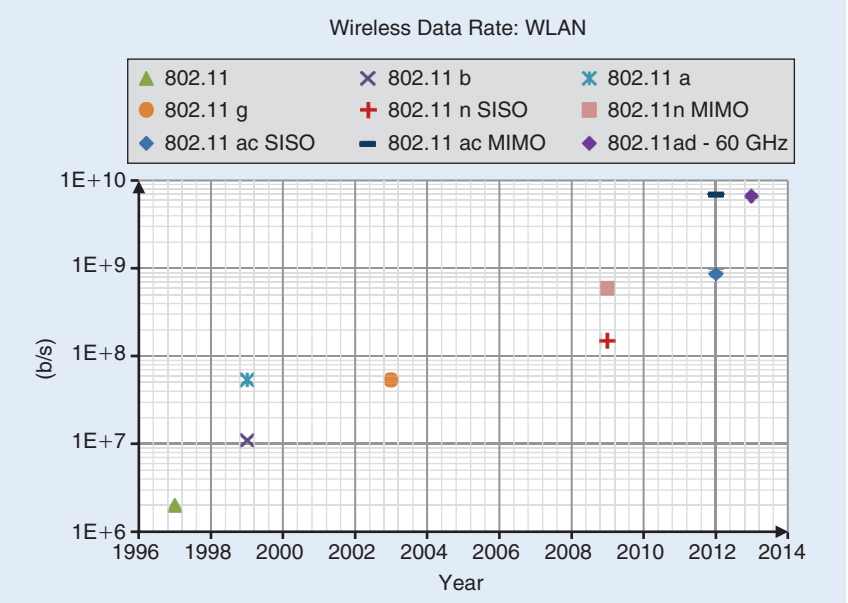

FIGURE 21: Data rate for wireless connectivity standards (802.11x).

Current state-of-the-art results from ISSCC 2013 include:

- 32 Gb ReRAM test chip developed in 24-nm CMOS
- the first ever 128 Gb 3 b/cell NAND Flash design in 20-nm planar-cell CMOS
- a 45 nm, 6 b/cell charge-trapping Flash memory using LDPC-based ECC demonstrates a ten-year error-free operation
- a highly efficient 6.4-Gb/s near-ground single-ended low-common mode transceiver for memory interface
- a highly efficient SRAM operating at 0.6 V used statistically gated sense amplifiers.

Technology Directions (TD)

The role of the Technology-Directions Subcommittee is to identify and encourage developments of potential importance in the ongoing evolution of ISSCC. Characteristically, a wide variety of topics are covered, some new and some continuing, with success achieved by the transference of the emerging technique to one of the evolving mainstream subcommittees. This year at ISSCC 2013, two strong emerging trend directions were visible: in flexible electronics and in NVM.

Large-Area Flexible Electronics

In the field of flexible large-area electronics fabricated at low temperatures, the current focus is now on lowering the cost-per-unit area, rather than on increasing the number of functions-per-unit area that is the focus of crystalline silicon technology, following Moore's law.

A clear breakthrough in research for large-area electronics in the past decade has been the development of thin-film transistor (TFT) processes with an extremely low temperature budget (<150 °C), enabling manufacturing on flexible and inexpensive substrates such as plastic films and paper.

For some time, the materials used for these developments have been carbon-based organic molecules such as pentacene, with properties of p-type semiconductors. More recently, air-stable organic n-type semiconductors and amorphous metal oxides, which are also

n-type semiconductors, have emerged. The most popular metal-oxide semiconductor is amorphous indium gallium zinc oxide (IGZO), but variants exist (zinc oxide, zinc tin oxide, and so on). The mobility of n- and p-type organic semiconductors has reached values exceeding 10 cm^2/Vs, which is already at par or exceeding the performance of that using amorphous silicon. Amorphous metal-oxide transistors have typical charge carrier mobility of 10–20 cm^2/Vs. Moreover, the operational stability of all organic semiconductor materials has greatly improved to a level sufficient to enable commercial applications, especially in combination with large-area compatible barrier layers to seal the transistor stack.

In the present state of the art, p-type only, n-type only, and complementary technologies are available. For the latter, all-organic implementations are available, but also we see hybrid solutions, featuring the integration of p-type organic with n-type oxide TFTs. At present, most TFTs are still manufactured with technologies from display lines, using subtractive methods based on lithography. However, there is a clear emphasis on the development of additive technologies that could provide higher production throughput, based on different approaches borrowed from the graphic printing world, such as screen and inkjet printing. The feature sizes and spread of characteristics of printed TFT technologies are still larger than those made by lithography, but there is clear progress toward size reduction.

The primary applications for such TFT systems are as backplanes in active-matrix displays, particularly flexible ones. Organic TFTs are well suited for electronic-paper-type displays, whereas oxide TFTs are targeting OLED displays. Furthermore, TFTs on foil are well suited for integration with temperature or chemical sensors to create pressure-sensitive foils, photodiode arrays, antennas, sheets capable of distributing RF power to appliances, energy scavenging devices, and so on, in hybrid integrated systems on foil. Early demonstrations include smart labels, smart shop shelves, smart medical patches, and so on. These are enabled by continuous progress in the complexity of analog TFT circuits targeting the interface with sensors and actuators, to modulate, to amplify, and to convert analog signals, as well as progress in digital TFT circuits and NVM for signal processing and storage.

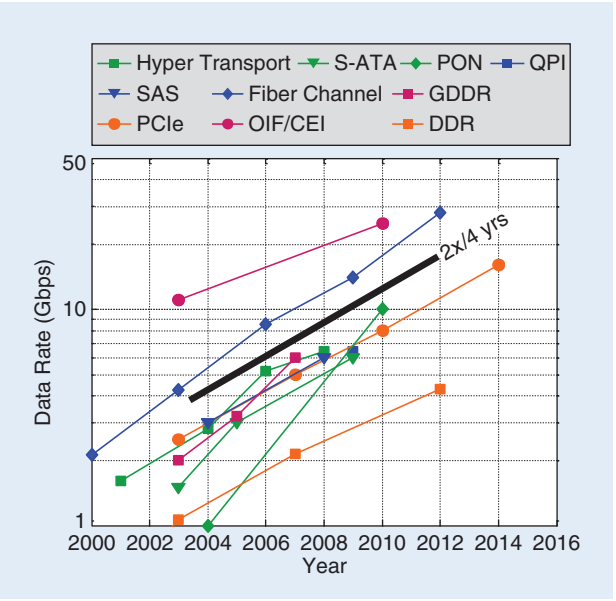

FIGURE 22: Per-pin data rate versus year for a variety of common wireline I/O standards.

Nonvolatile Memory in Logic

With continuing technology scaling, advances in energy-efficient computation will become even more important. Correspondingly, at ISSCC 2013, a new trend can be seen: It involves the integration of NVM with logic for ultra-fast power-down/power-up while maintaining the state of the computation. Such technologies will impact future mobile platforms and industrial applications, making mobile computing truly ubiquitous.

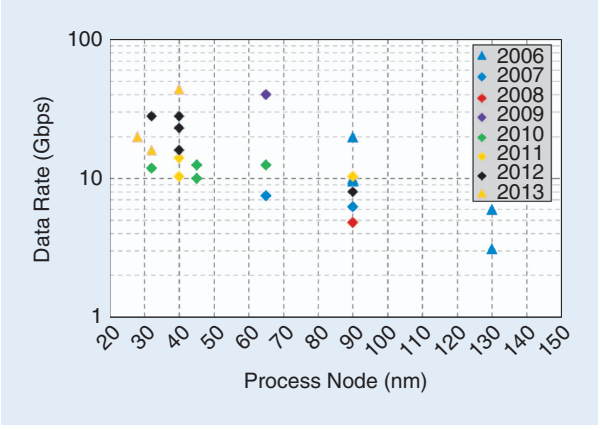

FIGURE 23: Wireline data rate versus process feature size and year.

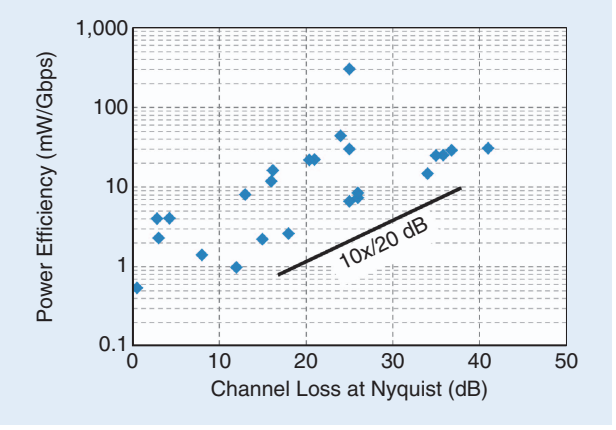

FIGURE 24: Wireline transceiver power efficiency versus channel loss.

Thus, we see two examples of the many new directions in which the electronics industry are likely to flow, as continuously uncovered by technology directions at ISSCC.

Summary

Developments reported at ISSCC 2013 continue to present breakthroughs in the broad domain of solid-state circuits and systems. In this rich environment, presentations at ISSCC characteristically predict ways in which electronics techniques will fulfill the present and future needs of the IEEE Solid-State Circuits Society. In this role, ISSCC continues to present a road map of things to come, both in the immediate future and in the longer term.

Acknowledgments

The authors wish to acknowledge the creators of the original material from which this article has been structured. While this has been largely a team effort by several members of each subcommittee, each has operated under the direction of his or her subcommittee chair, as follows:

- Energy-Efficient Digital: Stephen Kosonocky, Advanced Micro Devices
- High-Performance Digital: Stefan Rusu, Intel
- Technology Directions: Hoi-Jun Yoo, KAIST
- Imagers, MEMS, Medical, and Displays: Roland Thewes, TU Berlin
- Memory: Kevin Zhang, Intel
- Analog: Bill Redman-White, Southampton University
- Data Converters: Boris Murmann, Stanford University
- Radio Frequency: Andreia Cathelin, STMicroelectronics
- Wireless: David Su, Atheros Communications
- Wireline: Daniel Friedman, IBM Thomas J. Watson Research Center.

Our thanks go as well to the ISSCC technical editors for their initial editing and interaction with subcommittee members:

- Jason Anderson, University of Toronto
- Vincent Gaudet, University of Waterloo
- Glenn Gulak, University of Toronto
- James Haslett, University of Calgary
- Kostas Pagiamtzis, Semtech.

—Kenneth C. Smith
University of Toronto, Canada
—Alice Wang
MediaTek, Texas
—Laura C. Fujino
University of Toronto, Canada