Through the Looking Glass—The 2015 Edition

Trends in Solid-State Circuits from ISSCC

Siva G. Narendra, Laura C. Fujino, and Kenneth C. Smith

The 2015 ISSCC Program Committee
(From left) Jae-Youl Lee, Tsung-Hsien Lin, David Stoppa, Joo Sun Choi, Steve Kosonocky, John Trnka, Dan Friedman, Alison Burdett, Aarno Parssinen, Anantha Chandrakasan, Hoi-Jun Yoo, Stefan Rusu, Kevin Zhang, Makoto Ikeda, Andreia Cathelin, Boris Murmann, Laura Fujino, Axel Thomsen, and Eugenio Cantatore.

The International Solid-State Circuits Conference (ISSCC) is the flagship conference of the IEEE Solid-State Circuits Society. This year, for ISSCC, the theme is "Silicon Systems: Small Chips for Big Data." Big data is enveloping us: it is being generated by the Internet-of-Things (IoT), health care, and the Web, changing our society and our individual lives. Small silicon chips enable these changes through data sensing, gathering, processing, storing, and networking through wireless and wireline connectivity. Recent silicon-system technologies, including ultra-low-power systems, high-performance circuits and systems, wireless power and data transmission, and three-dimensional (3-D) IC structures, will open the door to big data applications. Moreover, big data applications

Digital Object Identifier 10.1109/MSSC.2014.2375071
Date of publication: 11 February 2015

1943-0582/15©2015IEEE

such as health care, machine learning, and sensor systems will challenge designers to consider new system architectures requiring advances in circuits and technology. ISSCC 2015 showcases novel circuit and system solutions that open new vistas for society, providing opportunities for new lifestyle transformations, all driven by big data technology.

The conference covers a spectrum of design approaches in various technical areas and advancements broadly categorized into analog systems; analog-digital data conversion systems; communication systems; digital systems; and innovative topics such as micromachines and microelectromechanical systems (MEMS), imagers, sensors, biomedical devices, as well as forward-looking solutions that may be several years away from becoming commercial.

Annually, the ten technical subcommittees of ISSCC update their analysis of industry trends for the benefit of the community at large. This article summarizes some of these views in select technical areas. A more comprehensive trends document will be available at www.isscc.org.

This year, for ISSCC, the theme is "Silicon Systems: Small Chips for Big Data."

Analog

Subcommittee Chair: Axel Thomsen, Silicon Laboratories, Austin, Texas, USA

The efficient control, storage, and distribution of energy are worldwide challenges and increasingly important areas of analog-circuit research. While manipulation and storage of information is efficiently performed digitally, the conversion and storage of energy is fundamentally performed with analog circuits. Thus, the key technologies for power management are predominantly analog. Digital systems such as microprocessors with multiple cores can reduce their overall power consumption by dynamically scaling operating voltage and frequency in response to time-varying computational demands. For this purpose, dc-dc voltage converters are embedded alongside the digital circuitry, driving research into the delivery of locally regulated power supplies with increasing efficiency, decreasing die area, and increasing output power, while involving no external components. As seen in Figure 1, these trends are captured by movement toward the top-right as the output-power level increases.

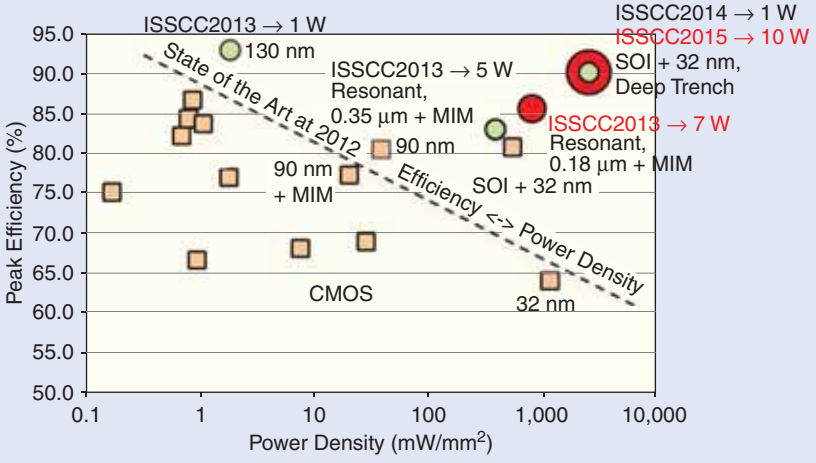

FIGURE 1: A comparison of integrated switched-capacitor power converters from ISSCC papers showing peak efficiency versus power density and output power. As shown in the top-right quadrant, recent advances achieve much higher power density and output power without sacrificing efficiency.

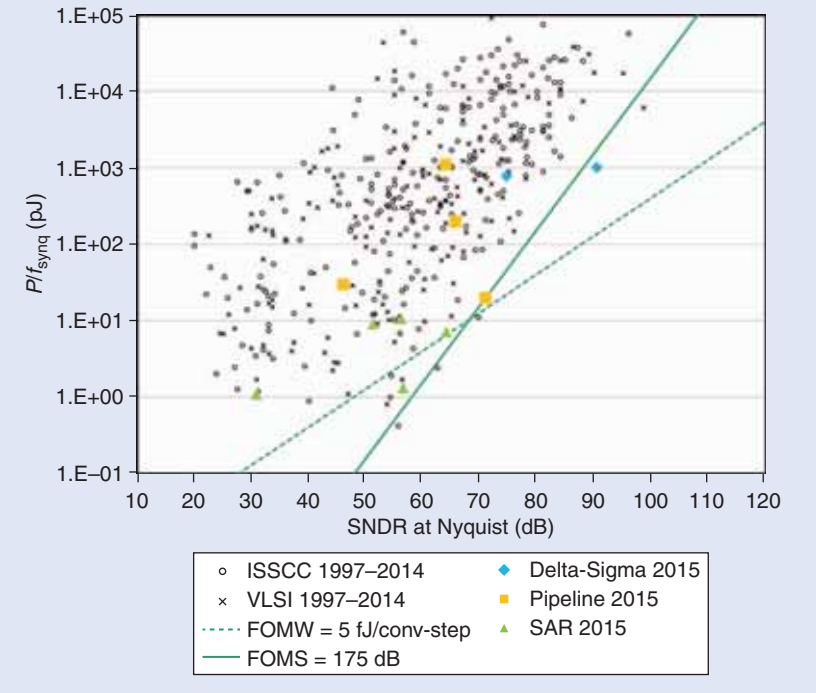

FIGURE 2: ADC power efficiency (P/f_{snyq}) as a function of SNDR.

Annually, the ten technical subcommittees of ISSCC update their analysis of industry trends for the benefit of the community at large.

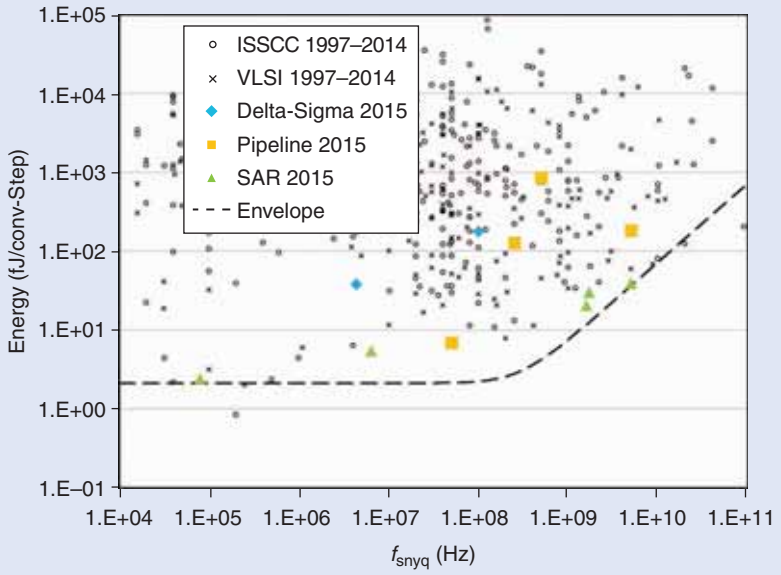

FIGURE 3: The reported ADC energy use per conversion step plotted against Nyquist frequency.

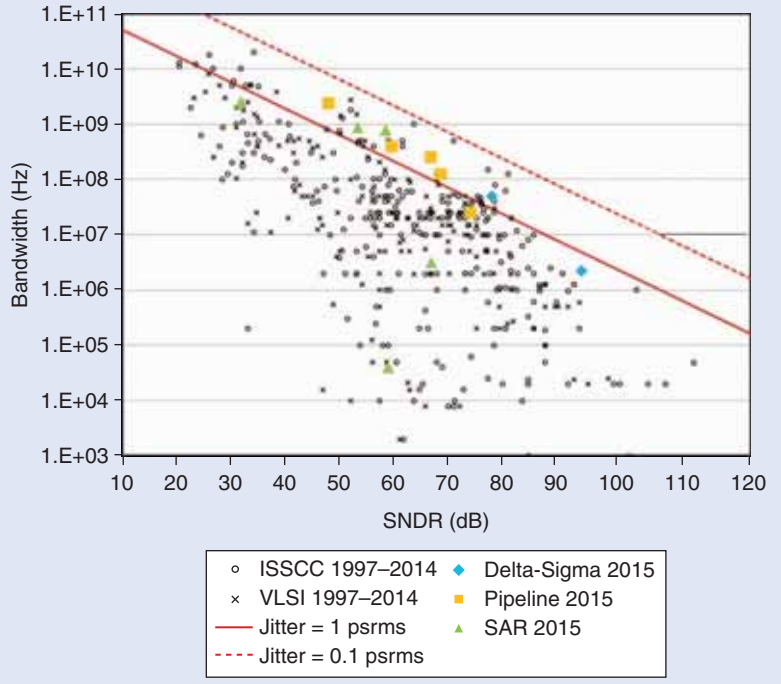

FIGURE 4: The ADC bandwidth (Nyquist frequency) plotted against SNDR.

Data Converters

Subcommittee Chair: Boris Murmann, Stanford University, Stanford, California, USA

Data converters serve as key building blocks in virtually all electronic systems and bridge the analog physical world to the digital circuitry dominant in modern ICs. Key metrics, such as linearity, bandwidth, and power efficiency, continue to be the dominant drivers for innovation, as evidenced by the data converters presented at ISSCC 2015. Also, for the first time, we see a converter in 14-nm FinFET technology. This design leverages the integration density of this technology to realize a process-voltage-temperature (PVT) tolerant time-to-digital converter using 2^{14} delay elements.

Figure 2 shows an overview of analog-to-digital converter (ADC) power efficiency expressed as power dissipated relative to the effective Nyquist rate (P/fsnyq) and as a function of signal-to-noise and distortion ratio (SNDR). For low-to-medium-resolution converters, energy is primarily expended to quantize the signal; the overall efficiency of this operation is typically measured by the energy consumed per conversion and quantization step. The dashed trend-line represents a benchmark of 5-fJ/conversion-step. Higher-resolution converters face the additional burden of overcoming circuit noise, necessitating a different benchmark proportional to the square of signal-to-noise ratio, represented by the solid line. Contributions at ISSCC 2015 are highlighted by the colored squares representing various converter architectures, and contributions from previous years are marked by smaller markers. Delta-Sigma, pipeline, and successive approximation register ADCs at various SNDR design points continue to push the limits of energy efficiency. (Note that on this chart, a lower P/fsnyq metric represents greater efficiency.) Figures 3 and 4 also show overviews of the reported ADC energy use per

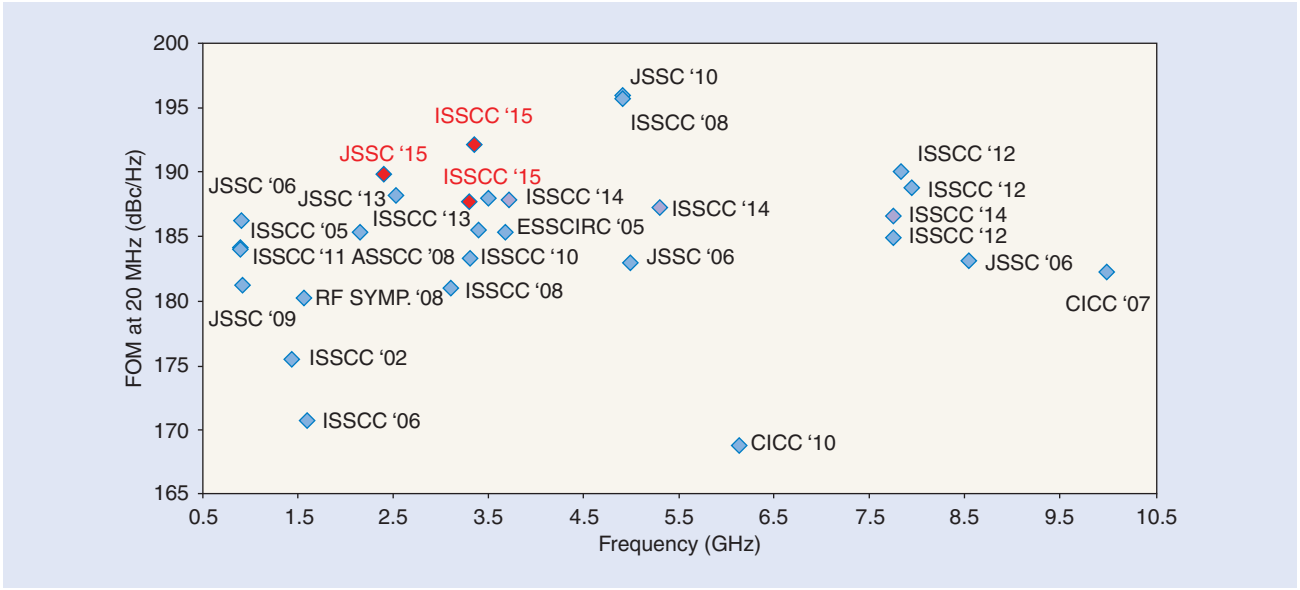

FIGURE 5: An oscillator phase-noise FOM at 20-MHz offset frequency versus oscillation frequency.

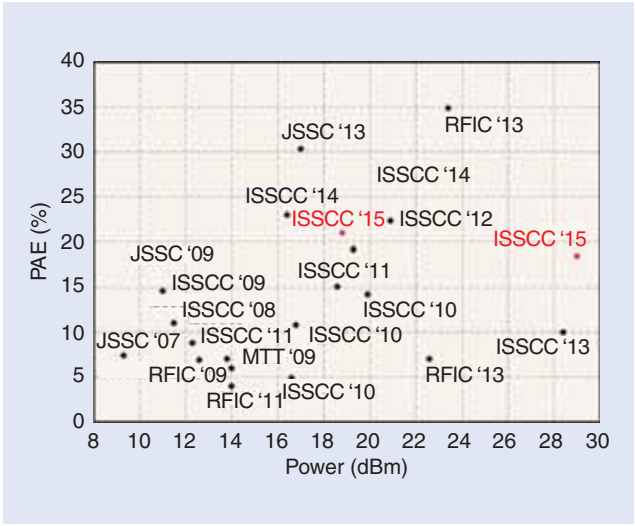

FIGURE 6: The PAE plotted against output amplifier power.

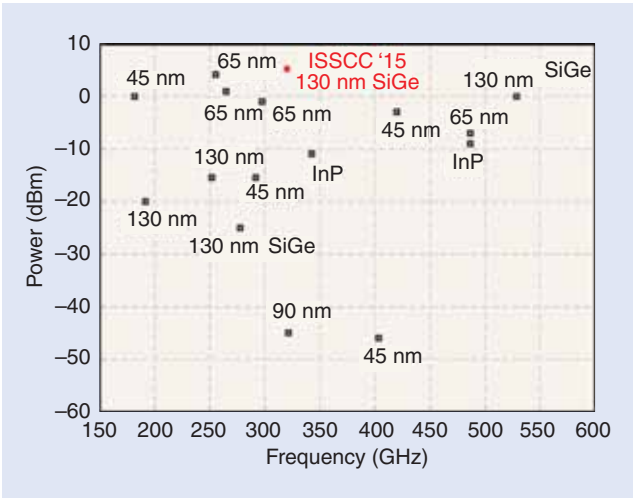

FIGURE 7: Output power versus frequency for mm-wave and sub-mm-wave sources.

conversion step and bandwidth plotted against SNDR, respectively.

Communication Systems—RF

Subcommittee Chair: Andreia Cathelin, STMicroelectronics, Crolles, France

One of the common yet critical building blocks in a radio frequency (RF) communication system is the voltage-controlled oscillator (VCO) frequency generation circuit. The race to reduce phase noise, chip area, and power consumption for synthesizers and VCOs is never ending. Utilizing circuit techniques to suppress the impact of flicker noise on phase noise, as well as employing a capacitance-scaling technique to decrease phase-looked loop (PLL) filter size are becoming popular. Very-low phase noise is achieved, and the VCO figure of merit (FOM) has reached a new record. The trend toward better performing VCOs is shown in Figure 5, where the VCO FOM versus oscillation frequency is continuously improving. In an attempt to remove the bulky quartz crystal, a highly stable thin-film-based reference-frequency generator is reported, attaining a stability of ±3 ppm from 0 to 90 ˚C.

Reported amplifier power-added efficiency (PAE) is plotted against output amplifier power and summarized in Figure 6. At the terahertz frontier, a phase-locked-based transmitter array has been demonstrated with sufficient output power beyond 300 GHz to enable numerous practical terahertz applications that require coherent radiation. The trend in output power versus frequency in

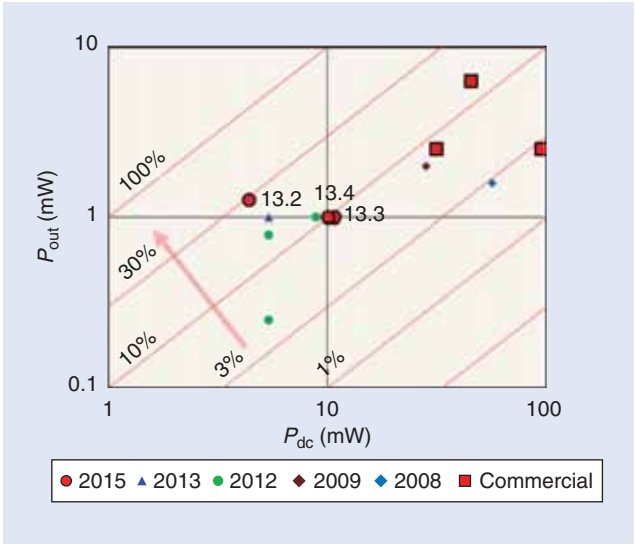

FIGURE 8: Ultra-low-power 2.4-GHz wireless transmit efficiency. The arrow shows the desired trends. All symbols except the red squares represent ISSCC papers. The ISSCC 2015 paper numbers are indicated next to the red circles.

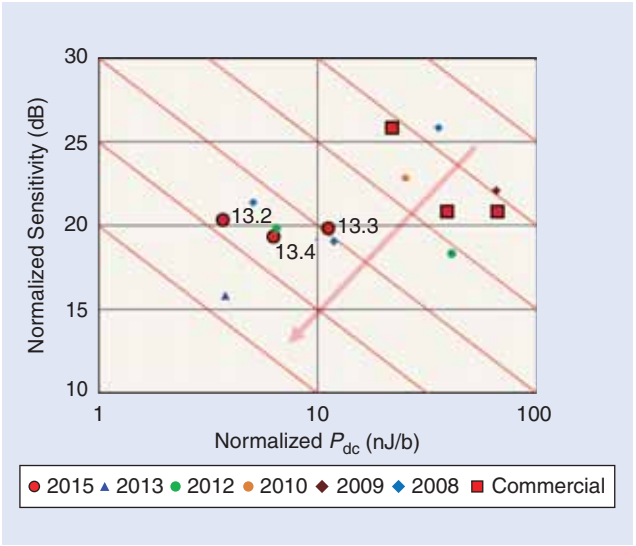

FIGURE 9: Ultra-low-power 2.4-GHz wireless receiver sensitivity trends. Sensitivity is normalized to bandwidth at 2.4-GHz carrier. The lines represent constant performance. The triangle is transceiver only, and the red dots and squares represent full SOCs.

mm-wave and sub-mm-wave sources is seen in Figure 7.

Communication Systems—Wireless

Subcommittee Chair: Aarno Pärssinen, University of Oulu, Espoo, Finland

The vision for wireless sensor networks (WSNs) and the IoT continues to force a more mature integration strategy. The commercial landscape is focusing these technologies toward Bluetooth low energy, while still maintaining flexibility by having the ability to function as multimode radios. ISSCC 2015 features some of the first published comprehensively integrated wireless-sensor nodes, including a microcontroller unit (MCU), power management, digital baseband modem, and wireless transceivers. The key feature and challenge of these wireless radios continues to be the ability to function at the lowest power possible, while being robust to the presence of other wireless signals. For most WSN nodes, transmission is a prominent operation. Therefore, transmit-power efficiency is a key metric that determines the evolution of these sensor nodes in terms of power consumption. The development trend for transmit-power efficiency of wireless sensor nodes is shown in Figure 8.

As shown in Figure 9, the receiver sensitivity as a function of power consumption in complete systems on a chip (SoCs) has almost reached the levels of a stand-alone wireless transceiver.

While an exponential increase in the data rate of cellular devices is still ongoing, another important trend is toward ever-higher levels of SoC integration. This trend in integration strategy is characterized in Figure 10; progress has been multifold: transitioning from BiCMOS to CMOS, to baseband-modem integration, and finally this year, to 3G PA integration on the same die. PA integration, as well as elimination of the front-end SAW filter, significantly advances the trend toward system simplification and bill of materials reduction in Figure 10.

Communication Systems—Wireline

Subcommittee Chair: Daniel Friedman, IBM T.J. Watson Research, Yorktown Heights, New York, USA

Over the past decade, wireline input–output (I/O) has been instrumental in enabling the incredible scaling of computer systems, ranging from handheld electronics to

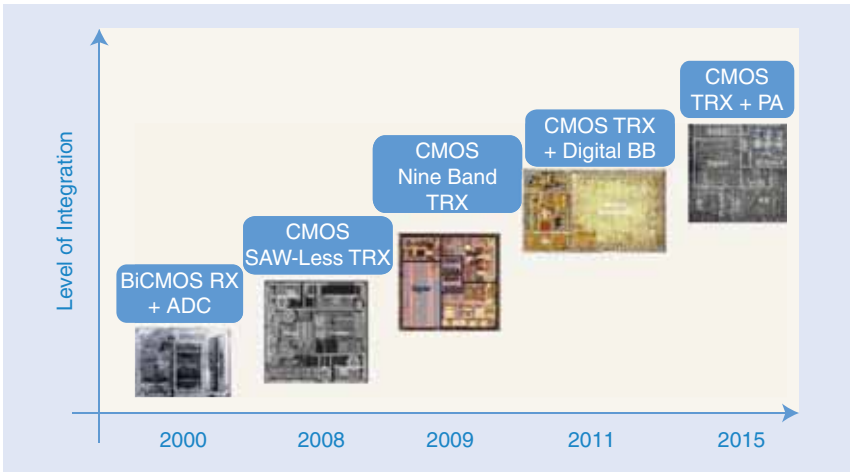

FIGURE 10: Trends in integration strategies.

supercomputers. During this time, aggregate I/O bandwidth requirements have increased at a rate of approximately two to three times every two years. Demand for bandwidth is driven by applications including memory, graphics, chip-to-chip fabric, backplane, rack-to-rack, and local area networks. In part, this increase in bandwidth is enabled by expanding the number of I/O pins per component. As a result, I/O circuitry consumes an increasing amount of area and power on today's chips.

At the same time, increasing bandwidth has also been enabled by rapidly accelerating the per-pin data rate. Figure 11 shows that the per-pin data rate has approximately doubled every four years across a variety of diverse I/O standards ranging from double data rate (DDR) to graphics to high-speed Ethernet.

Figure 12 shows that the data rates for published transceivers have kept pace with these standards, enabled in part by process-technology scaling. However, continuing along this rather amazing trend for I/O scaling will require more than just transistor scaling. Significant advances in both energy efficiency and signal integrity must be made to enable the

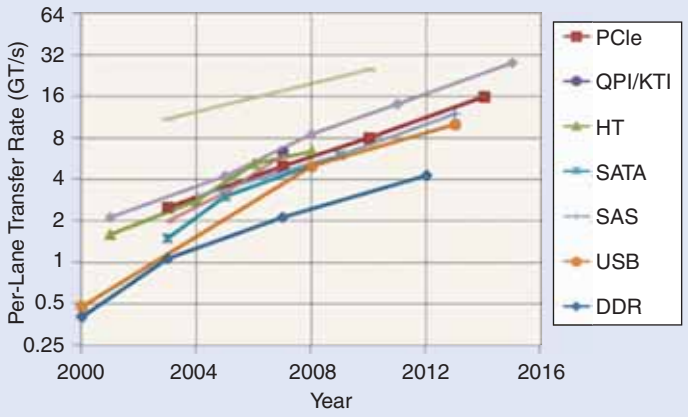

FIGURE 11: The per-pin data rate versus year for a variety of common I/O standards.

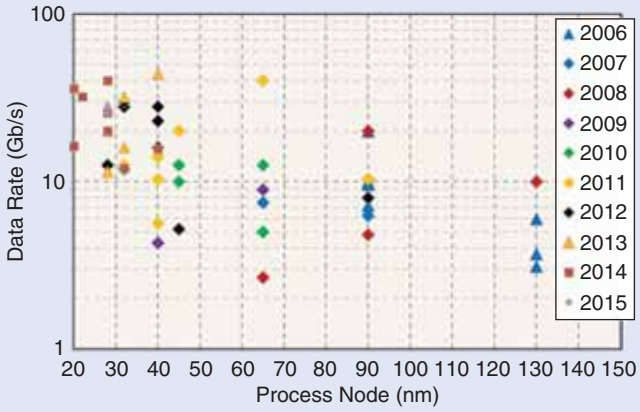

FIGURE 12: The data rate versus process node and year.

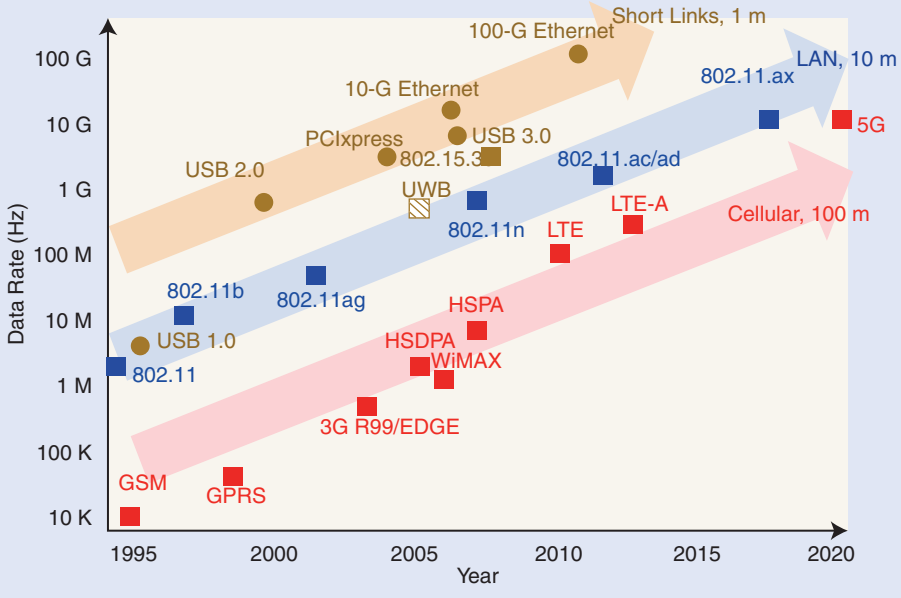

FIGURE 13: Trends in data-rate frequencies for various communication channels.

ISSCC 2015 showcases novel circuit and system solutions that open new vistas for society.

next generation of low-power and high-performance computing systems. Figure 13 summarizes trends in data rates for various communication channels.

Digital Systems—Energy-Efficient Digital

Subcommittee Chair: Stephen Kosonocky, AMD, Fort Collins, Colorado, USA

The demand for ubiquitous mobile functionality to achieve enhanced productivity, a better social-networking experience, and improved multimedia quality continues to drive innovation in SoCs, which are also constrained by a need to improve battery life and reduce cost. While the performance of embedded application processors has increased to meet the rising demands of general-purpose computing, dedicated multimedia accelerators are necessary to provide dramatic improvements in performance and energy efficiency for emerging applications. At the other side of the spectrum, sensor nodes for the IoT require low-energy wireless and sufficient computational capabilities.

Technology scaling continues to be exploited to deliver designs capable of operating at lower voltages, resulting in reduced energy per operation, as well as lowering the area required to implement specific functions. Processors that will be unveiled at ISSCC 2015 are built in a variety of technology nodes, with best-in-class results accomplished with higher integration, and improved performance per watt as seen in Figure 14. These are demonstrated in various process nodes ranging from 130- to 14-nm FinFET, as well as low-power FD-SOI CMOS technologies.

Semiconductor chips used for the IoT and other embedded application areas demand energy-efficient active operation while also requiring ultra-low power state retention to enable event-driven operation with fast wake up. Typically, an energy-efficient MCU is used as the main processing element for local processing and system control of these applications. Over time, the microcontroller energy efficiency, shown in Figure 15, has continuously improved through a combination of technology and design techniques.

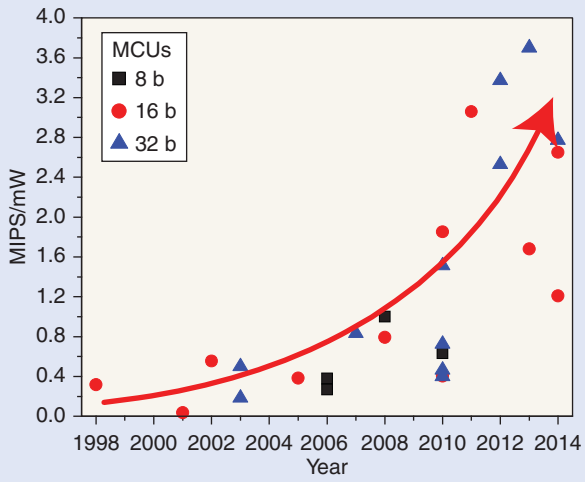

FIGURE 14: The evolution of energy efficiency in commercial microcontrollers.

Digital Systems—High-Performance Digital

Subcommittee Chair: Stefan Rusu, Intel, Santa Clara, California, USA

The relentless march of process technology brings more integration and energy-efficient performance to mainframes, enterprise, and cloud servers.

At ISSCC 2015, superlatives abound: IBM details its high-frequency eight-core, 16-thread System z mainframe

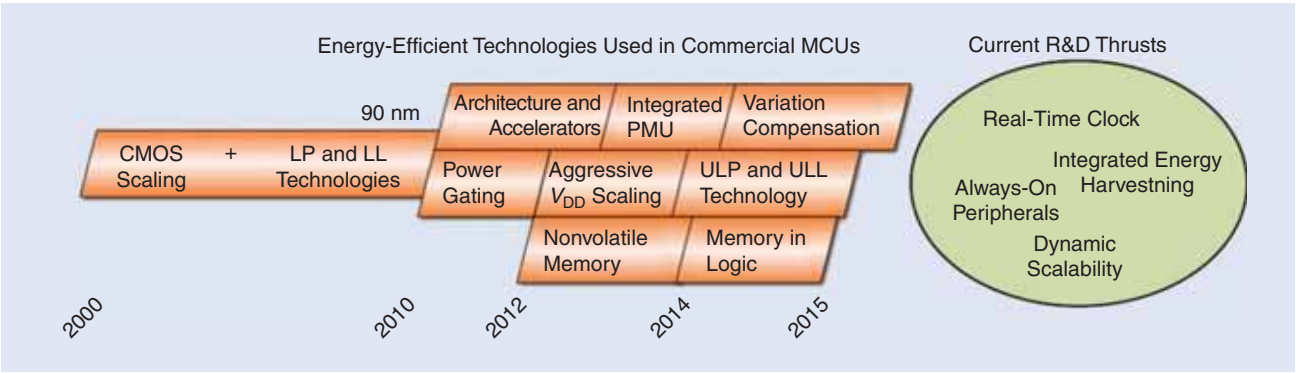

FIGURE 15: Energy-efficiency techniques exploited in microcontroller design.

processor in 22-nm SOI with 64 MB of eDRAM L3 cache and 4-MB/core eDRAM L2 cache. Oracle details its SPARC M7 processor with 32 S4 cores, a 1.6-TB/s bandwidth 64-MB L3 cache, and a 0.5-TB/s data bandwidth on-chip network to deliver more than three times the throughput compared to its predecessor. It also includes 280 SerDes lanes that support up to 18-Gb/s line rate and 1-TB/s total bandwidth. Intel details its next-generation Xeon processor, which supports 18 dual-threaded 64-b Haswell cores, 45-MB L3 cache, 4 DDR4-2133MHz memory channels, 40 8-GT/s PCIe lanes, and 60 9.6-GT/s QPI lanes. It has 5.56-B transistors in Intel's 22-nm trigate HKMG CMOS, achieving a 33% performance boost over previous generations.

The chip complexity chart in Figure 16 shows the trend in transistor integration on a single chip over the past two decades. While the 1 billion transistor-integration threshold was achieved some years ago, we now commonly see processors incorporating more than 5-B transistors on a die.

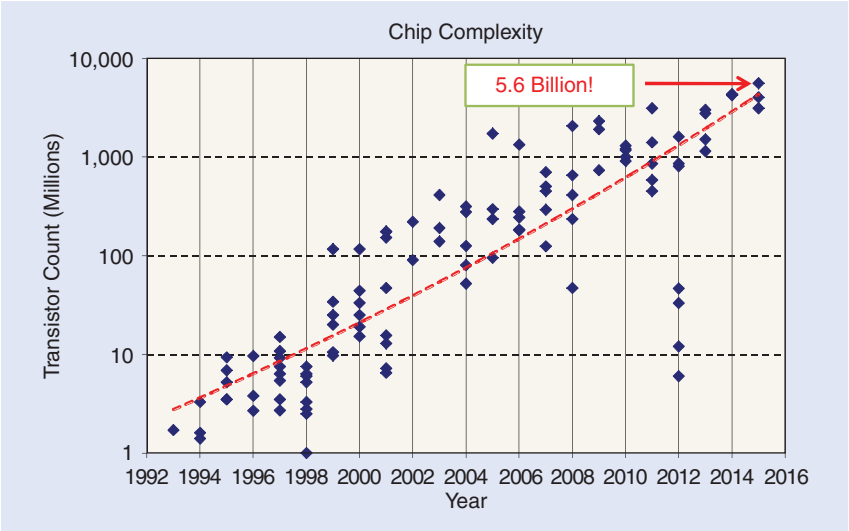

FIGURE 16: Transistor count versus year.

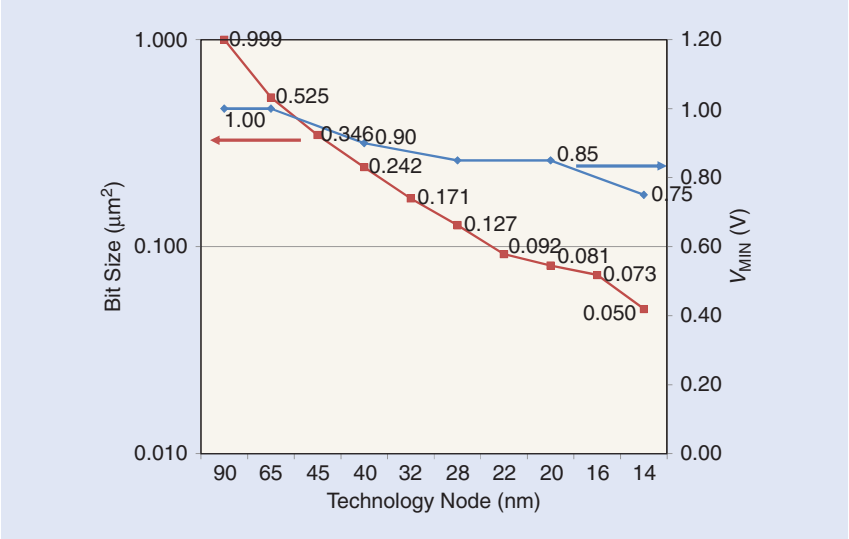

FIGURE 17: Bit cell and V_{DD} scaling trend for SRAM.

Digital Systems—Memory

Subcommittee Chair: Joo Sun Choi, Samsung Electronics, Hwasung, Korea

Mobile products are everyone's companion and need to store and process ever-increasing amounts of data. Progress is possible only by constant improvements in area, power, and performance of volatile and nonvolatile memory (NVM). FinFET technology is now mainstream for embedded SRAM and DRAM, facilitating continued scaling. Improvements in DRAM data rates support the increasing demands of greater data volumes. NAND flash memories have moved from 2 b/cell to 3 b/cell, and 3-D multilayer designs are now typical. Embedded flash, which is essential to IoT and wearable applications, has moved to 28 nm. Among emerging memories, STT-MRAM is the most mature, although ReRAM is quickly catching up.

Some outstanding state-of-the-art paper topics from ISSCC 2015 include

- two 14-nm SRAM bit cells; 0.050 μm^2 (HDC) and 0.058 μm^2 (LVC) capable of achieving 1.5 GHz operation at 0.6V
- a 14-nm FinFET SOI eDRAM with a cell size of 0.01747 μm^2 and 1-ns access time
- 128-Gb, 3 b/cell 32 stacked WL layer 3-D NAND flash running at 1Gb/s I/O rate
- low-power 64-Gb 2b/cell NAND flash manufactured in 15-nm technology
- a 1.1-V, 10-Gb/s/pin transceiver for DRAM interface suitable for use beyond LPDDR4
- a high-speed 1-Mb STT-MRAM using 2T2MTJ cells achieves 3.3-ns access time and a sub-20-nm technology node STT-MRAM uses a high-density 1T1MTJ memory cell
- a 28-nm embedded SG-MONOS FLASH developed for automotive applications.

SRAM

Consumer and computing products in 2015, from smart watches to the cloud, depend on low-power and high-performance embedded SRAM. Challenges for SRAM include V_{MIN}, leakage, and dynamic power reduction. Last

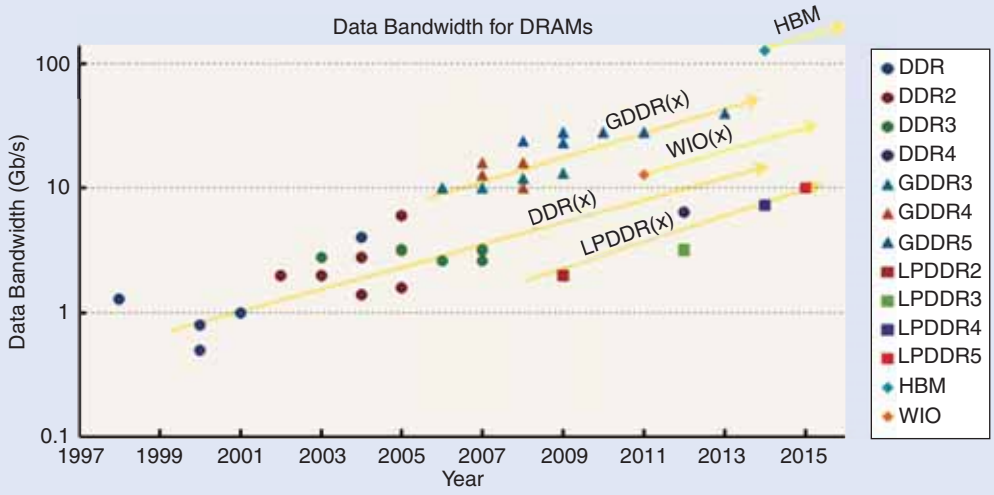

FIGURE 18: The data bandwidth for DRAMs.

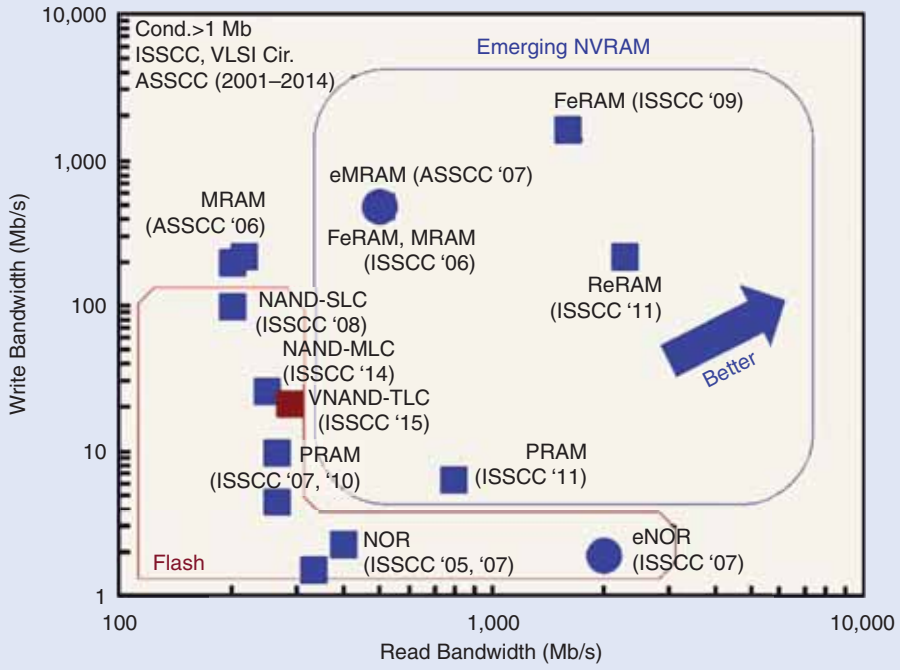

FIGURE 19: Trends in NVM technologies comparing write bandwidths to read bandwidths.

year saw the first introduction of 14-nm/16-nm technology. This year, a 14-nm SRAM using second-generation FinFET technology implements the smallest SRAM bit cell ever reported at just 0.050 μm^2. As the transistor feature size reduced below 20 nm, device variation has made it very difficult to shrink the bit-cell size at the desired 50% rate while maintaining or lowering V_{MIN} between generations. The introduction of high-κ metal-gate (45 nm) and FinFET or fully depleted SOI transistors (22 nm) has reduced V_{TH} mismatch and has enabled further device scaling. Design solutions such as read/write-assist circuitry and variation-tolerant sensing schemes have been used to improve SRAM V_{MIN} performance starting at 32 nm and are now ubiquitous in high-density SRAM designs at 14/16 nm. Dual-rail SRAM design emerges as an effective solution to enable dynamic voltage-frequency scaling by decoupling logic-supply rails from SRAM arrays, thus allowing a much wider operating window. The use of assist-circuit techniques, FinFET transistors, and dual-rail architectures is expected to

extend the viability of the high-density 6-T SRAM bit cell beyond 14 nm. It is important for SRAM to reduce both leakage and dynamic power, keeping products within the same power envelope at the next technology node. Last year, the first FinFET-based eDRAM at 22 nm was introduced. This year, a 14-nm FinFET-based eDRAM is discussed. eDRAM continues to show itself as a desirable way to provide memory scaling in high-performance CPU designs. Figure 17 shows the trends in bit-cell and V_{DD} scaling for SRAMs from major semiconductor manufacturers.

High-Bandwidth DRAM

To reduce the bandwidth gap between main memory and processor performance, DRAM data rates continue to increase at the memory interface with schemes such as DDR(x), LPDDR(x), and GDDR(x), as shown in Figure 18. Notable in Figure 18 is the introduction of the hybrid memory cube.

Nonvolatile Memory

Trends in the density and speed of nonvolatile random access memory (NVRAM) continue to show density and speed advancements. Figure 19 compares NAND/NOR flash write/read bandwidths to emerging NVRAM including magnetic random access memory (MRAM), ferroelectric RAM (FeRAM), resistive RAM (ReRAM or RRAM), embedded MRAM (eMRAM), and phase-change RAM (PRAM). The performance of embedded NOR Flash (eNOR) is also seen in this figure. Figure 20 details the changes in density for these various technologies. It's interesting to note that NAND flash density has leveled off the last few years despite the higher density stored per cell, as shown in Figure 21, including triple-level cell (TLC, eight levels or 3 b) and higher density per mm^2 of chip area. Of further interest is the trajectory of the resistive memory technologies including RRAM and PRAM seen in Figure 21. If this trend continues, it appears that resistive NVM technologies will replace flash memory in the near future.

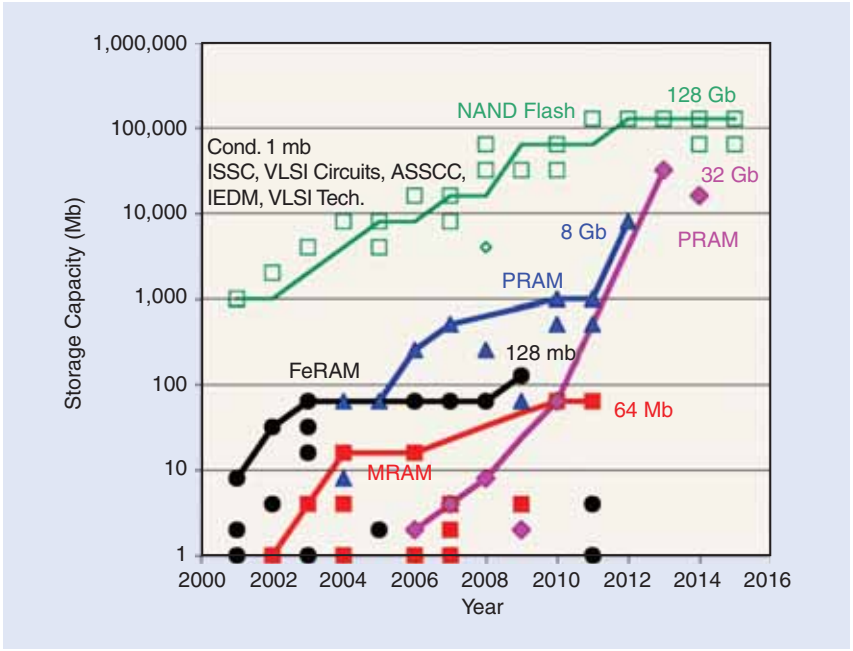

FIGURE 20: Trends in NVM technology capacities.

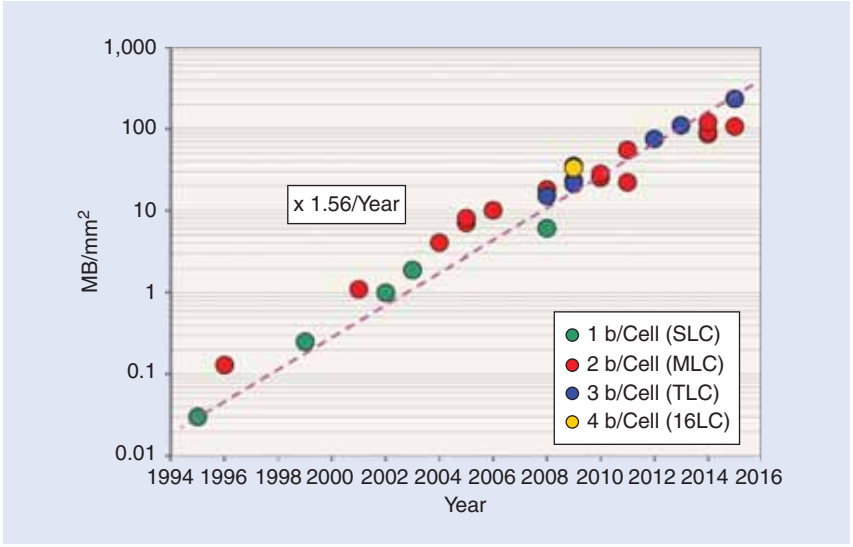

FIGURE 21: Trends in the density of NAND flash memory.

Innovative Topics—Sensors, MEMS, and Medical Devices

Subcommittee Chair: Makoto Ikeda, University of Tokyo, Japan

CMOS temperature sensors continue to improve, with new circuit techniques increasing accuracy and supporting wider applications. Strides in low-power architectures for such sensors continue to eclipse previous results with record-setting efficiency that support battery-operated and mobile applications. Pure digital implementations are often used in thermal monitoring for large SoCs, which is an important application where small sensor area is a key feature.

Current sensors are becoming more integrated and precise. These devices detect the magnetic field around a wire or trace carrying a current, and they are used, for instance, in electrical motor drives, solar power, and battery-charger applications.

Capacitive-to-digital converters are essential sensor interface components, where new circuit technologies result in increased accuracy, smaller die size, and reduced power consumption. Digital-centric implementations enable these improvements to be maintained in more advanced processes.

MEMS inertial sensors (accelerometers and gyroscopes) are key components used in a large variety of consumer products, where ultra-low power consumption is a key requirement. For automotive applications, reduced vibration sensitivity and high precision are additional requirements.

In life-science applications, there is an ever-increasing demand for higher-throughput tools, such as cellular and molecular assays, coming from the field of precision and personalized medicine. As well, for improved capability and quality, medical ultrasound is moving toward 3-D imaging with large arrays. As these arrays increase in size, the number of connections to the front-end processing circuitry and the amount of required signal processing are becoming bottlenecks. To resolve this congestion, MEMS transducers and CMOS technology increasingly allow transducer arrays and their signal processors to be mounted together. On the other hand, smart wearable sensors for medical applications will support continuous remote monitoring with data transmission to centralized analysis systems.

Innovative Topics—Technology Directions

Subcommittee Chair: Eugenio Cantatore, Eindhoven University of Technology, The Netherlands

A breakthrough concept in biomedical electronics has been the development of "anytime and anywhere" human monitoring systems using wearable/implantable devices. This will enable improvements in the quality of life, such as self-health checks, remote medical examination, and constant monitoring for acute diseases.

> ***ISSCC continues to be the premier technical forum for presenting advances and predicting trends in solid-state circuits and systems.***

Key technologies for such systems include

1) small footprint devices and flexible electronics to enhance the comfort of wearable devices
2) high accuracy and low-power monitoring devices
3) low-power communication systems for wireless operation.

To develop devices with a small footprint, highly integrated biomedical SoCs capable of sensor detection, diagnosis, and wireless communication have been developed. Progress in the development of 3-D-integrated electrodes and sensors with MEMS has also contributed to reduce the footprint. To monitor weak vital signs, such as electrocardiogram (ECG) and electroencephalogram (EEG), with a high degree of accuracy, there is clear progress in variation-correction methodologies and techniques to compensate for artifacts due to the subject's movements. Furthermore, state-of-the-art biomedical SoCs are capable of multimodal telemonitoring of various biosignals. To improve the device longevity, SoCs have evolved to include a power-efficient dedicated processor and algorithms for diagnosis, ultra-low-power circuit blocks, low-power wireless circuits, and energy-harvesting technologies. Furthermore, various communication standards are being developed to satisfy the low-power requirements of portable medical systems.

Another trend in technology directions is a move toward lab-on-a-chip or semiconductor systems for diagnosis and disease screening. Silicon solutions are being designed for portable systems to quickly and inexpensively diagnose illness from simple blood tests, to imaging, to spectroscopy. The trend of using emerging technologies such as MEMS or CMOS-CNT (carbon nanotube) arrays to solve medical problems will make it possible to bring traditionally expensive diagnostic tools to rural and remote areas where medical expertise is typically scarce.

Summary

According to the SIA, the semiconductor industry generated US$306 billion in sales in 2013! In this environment, ISSCC continues to be the premier technical forum for presenting advances and predicting trends in solid-state circuits and systems. Beyond this article, a complete trends document will be available at www.isscc.org. These trends are highlighted in papers to be presented at the 62nd ISSCC, 22–26 February, in San Francisco, California. Attendance at ISSCC 2015 is expected to be around 3,000. Corporate attendees from the semiconductor and system industries typically represent around 60% of the attendees. We look forward to seeing you there!

Acknowledgments

The authors would like to acknowledge the ISSCC 2015 Technical Program Committee for providing original content for this article and each of the ten subcommittee chairs as referenced in each of the trends sections for their leadership role. Without their collective efforts, this article would not have been possible.

About the Authors

Siva G. Narendra is with Tyfone in Portland, Oregon.

Laura C. Fujino is with the University of Toronto in Canada.

Kenneth C. Smith is with the University of Toronto in Canada.