ISSCC 2014
TRENDS

CONDITIONS OF PUBLICATION

PREAMBLE

The Trends document serves to capture the trends within each technology track covered by the scope of ISSCC 2014 in February in San Francisco

High resolution figures are available along with the press kit. For access please contact the press coordinator.

OBTAINING COPYRIGHT to ISSCC 2014 press material is EASY!

You are welcome to use this material, copyright- and royalty-free, with the following understanding:

That you will maintain at least one reference to ISSCC 2014 in the body of your text, ideally retaining the date and location. For detail, see the FOOTNOTE below.

That you will provide a courtesy PDF of your excerpted press piece and particulars of its placement to press_relations@isscc.net

FOOTNOTE

- From ISSCC's point of view, the phraseology included in the box below captures what we at ISSCC would like your readership to know about this, the 61st appearance of ISSCC, on February 9th to the 13th, 2014, in San Francisco.

This and other related topics will be discussed at length at ISSCC 2014, the foremost global forum for new developments in the integrated-circuit industry. ISSCC, the International Solid-State Circuits Conference, will be held on February 9-13, 2014, at the San Francisco Marriott Marquis Hotel.

ISSCC Press Kit Disclaimer
The material presented here is preliminary.
As of November 1, 2013, there is not enough information to guarantee its correctness.
Thus, it must be used with some caution.

HISTORICAL TRENDS IN TECHNICAL THEMES

Analog Systems

(Analog Subcommittee, Data Converters Subcommittee)

Analog – 2014 Trends

SC Power Converters

Subcommittee Chair: *Axel Thomsen, Silicon Laboratories, Austin, TX*

The efficient control, storage, and distribution of energy are worldwide challenges, and are increasingly important areas of analog circuit research. While the manipulation and storage of information is efficiently performed digitally, the conversion and storage of energy must fundamentally be performed with analog systems. As a result, the key technologies for power management are predominantly analog. For example, there is much interest in wireless power transmission for battery charging applications, ranging from mobile handsets to medical implants: Increased efficiency in wireless power transmission is enabling faster charging over longer distances. There is also an explosion of technologies that permit energy to be collected from the environment via photovoltaic, piezoelectric, or thermoelectric transducers. A significant focus here is on analog circuits that are able to harvest sub-microwatt power levels from energy sources at tens of millivolts, to provide autonomy for remote sensors or to supplement conventional battery supplies in mobile devices. To achieve this, extremely low power must be consumed by the attendant analog circuits so that some energy is left over to charge a battery or supercapacitor. Similarly, the power consumption of analog instrumentation amplifiers, oscillators, and audio power amplifiers is being scaled down to meet the demands of these low-power systems. Fast power-up and -down is also desired from these circuits to permit high energy-efficiency during intermittent operation. Together, these technologies will permit devices to be powered indefinitely from sustainable sources, opening the door to ubiquitous sensing, environmental monitoring, and medical applications.

Analog circuits also serve as bridges between the digital world and the analog real world. Just like the bridges in our roads, analog circuits are often bottlenecks and their design is critical to overall performance, efficiency, and robustness. Nevertheless, digital circuits such as microprocessors drive the market; thus, semiconductor technology has been optimized relentlessly over the past 40 years to reduce the size, cost, and power consumption of digital circuits. Analog circuitry has proven increasingly difficult to implement using these modern IC technologies. For example, as the size of transistors has decreased, the range of analog voltages they can handle has decreased, and the variation observed in their analog performance has increased.

These aspects of semiconductor technology explain two key divergent trends in analog circuits. One trend is to forgo the latest digital IC manufacturing technologies, instead fabricating analog circuits in older technologies, which may be augmented to accommodate the high voltages demanded by increasing markets in medical, automotive, industrial, and high-efficiency lighting applications. Other applications dictate full integration of analog and digital circuits together in our most modern digital semiconductor technologies. For example, 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 can be embedded alongside the digital circuitry, driving research into the delivery of locally-regulated power supplies with high efficiency and low die area, but without recourse to external components. These trends are captured by movement towards the top-right in the plot below.

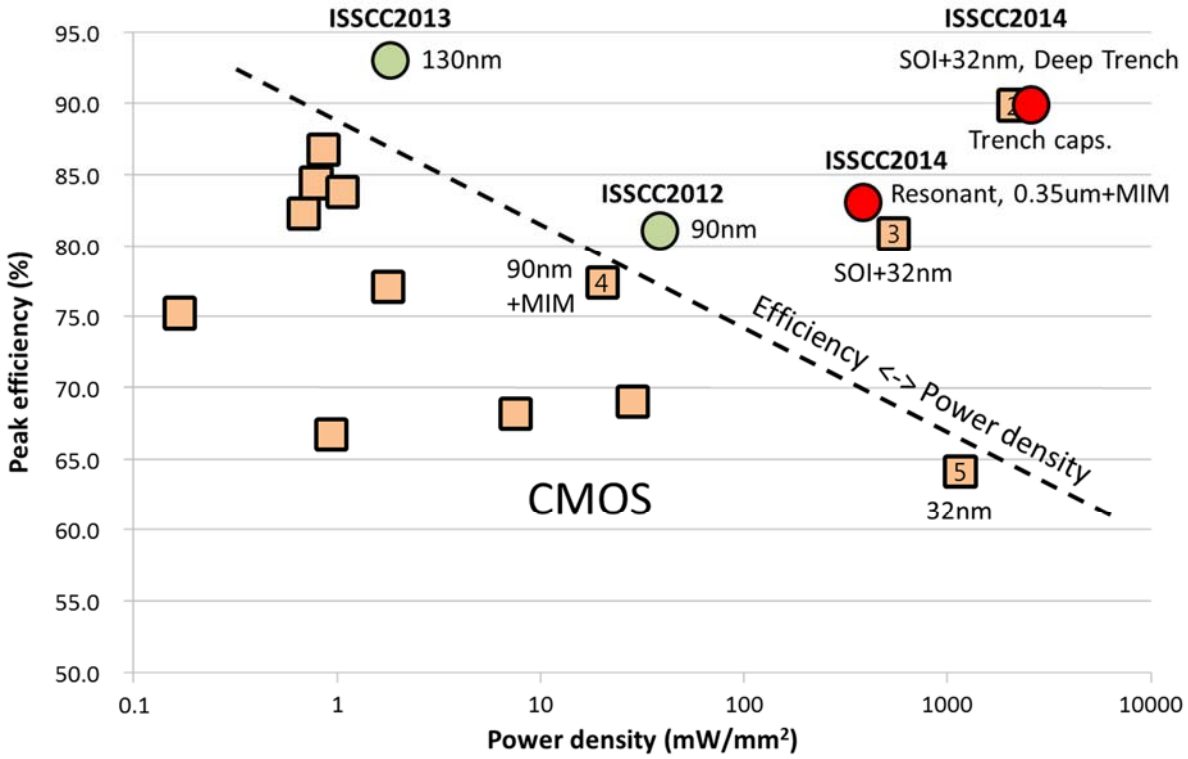

Comparison of Integrated Switched-Capacitor Power Converters showing Peak Efficiency vs Power Density. Recent advances achieve much higher power density without sacrificing efficiency (see the top right quadrant).

Data Converters – 2014 Trends

Subcommittee Chair: *Boris Murmann, Stanford University, Stanford, CA*

Data converters serve as key building blocks in virtually all electronic systems, and serve to bridge the analog physical world to the digital circuitry prevalent in modern integrated circuits. Key metrics such as signal-to-noise ratio, bandwidth, and power efficiency continue to be the dominant drivers for innovation, as evidenced by the data converters presented at ISSCC 2014.

The first figure below is a survey of ADC power efficiency expressed as power dissipated relative to the effective Nyquist rate (P/f_{snyq}), 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 efficiency of this operation can be measured as the energy consumed per conversion step. The dashed trend-line represents a benchmark of 10fJ/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 shown by the solid line. Contributions at ISSCC 2014 are indicated by the colored dots representing various converter architectures with contributions from previous years denoted by smaller dots. (Note that a lower P/f_{snyq} metric represents a more efficient circuit.) Several new SAR-based converters at various SNDR design points continue to push the limits of energy efficiency.

The second figure shows energy per conversion step vs. the Nyquist sampling rate. This figure elucidates the difficulty of maintaining good efficiency at higher speeds of operation. Nevertheless, advances in circuit innovations embodied in leading-edge technologies have resulted in new benchmarks in energy efficiency across the entire spectrum of conversion rates. Most notably, we are seeing for the first time converters with good efficiency at speeds of several tens of gigasamples per second.

The final chart plots achieved bandwidth as a function of SNDR. Sampling jitter or aperture errors make the combination of high resolution and high bandwidth a particularly difficult task. Nonetheless, in 2014, we see many examples setting a new standard in this metric utilizing several different converter architectures.

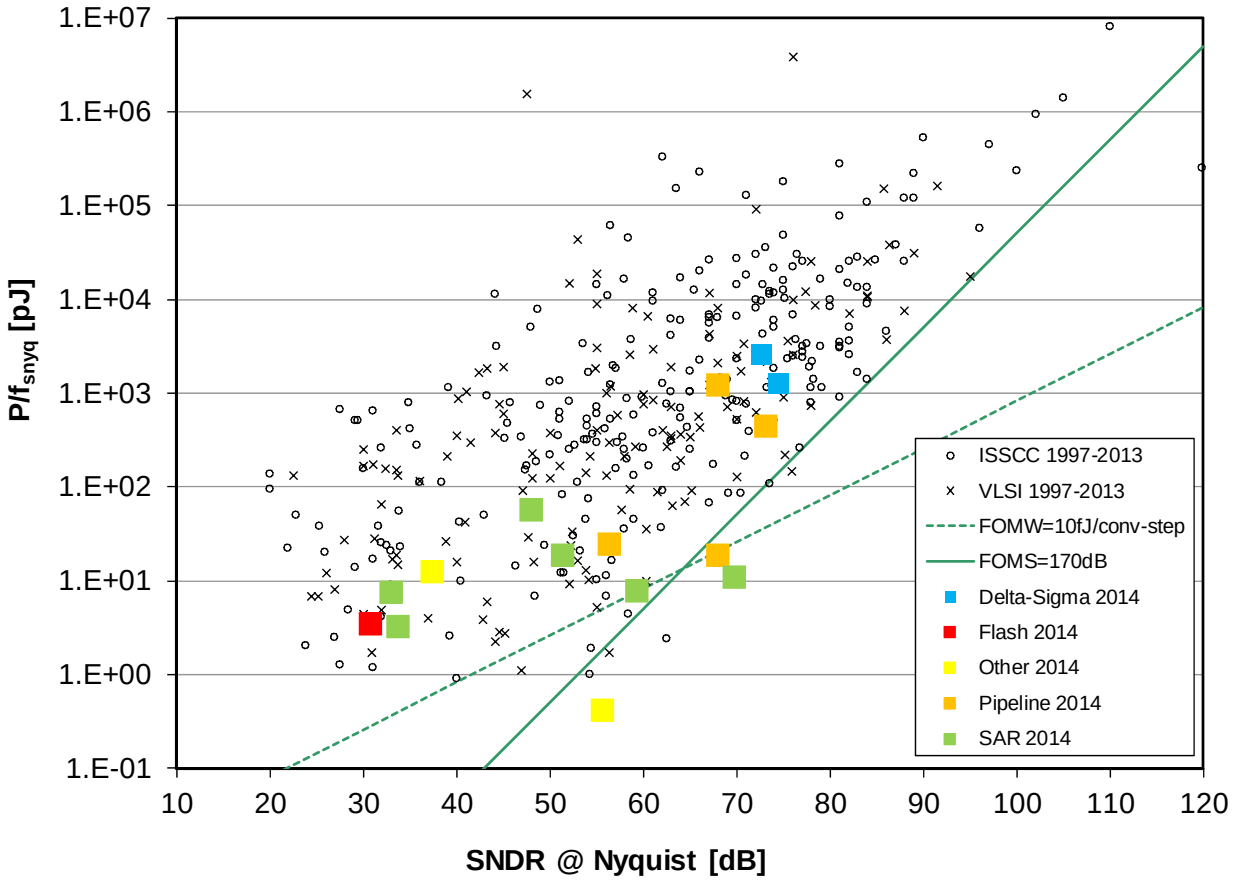

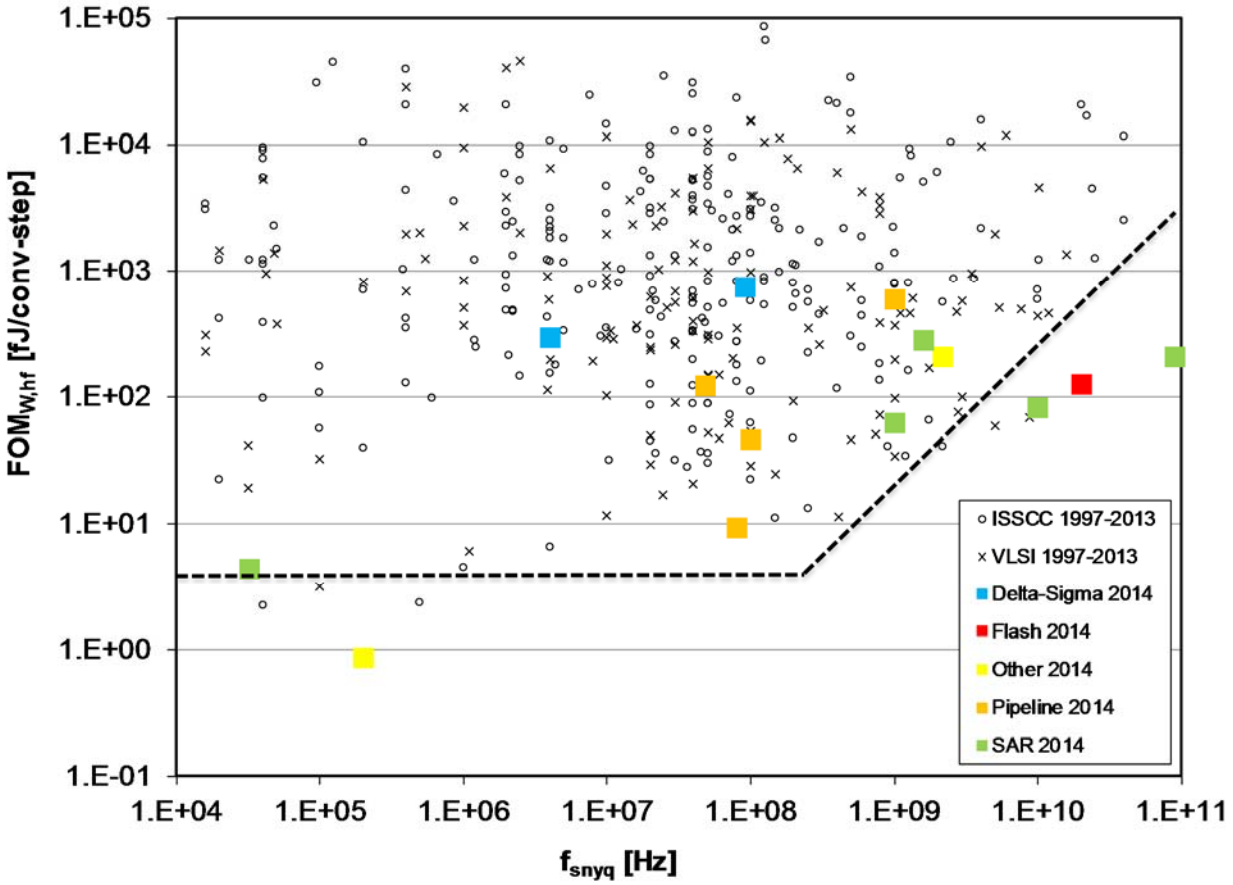

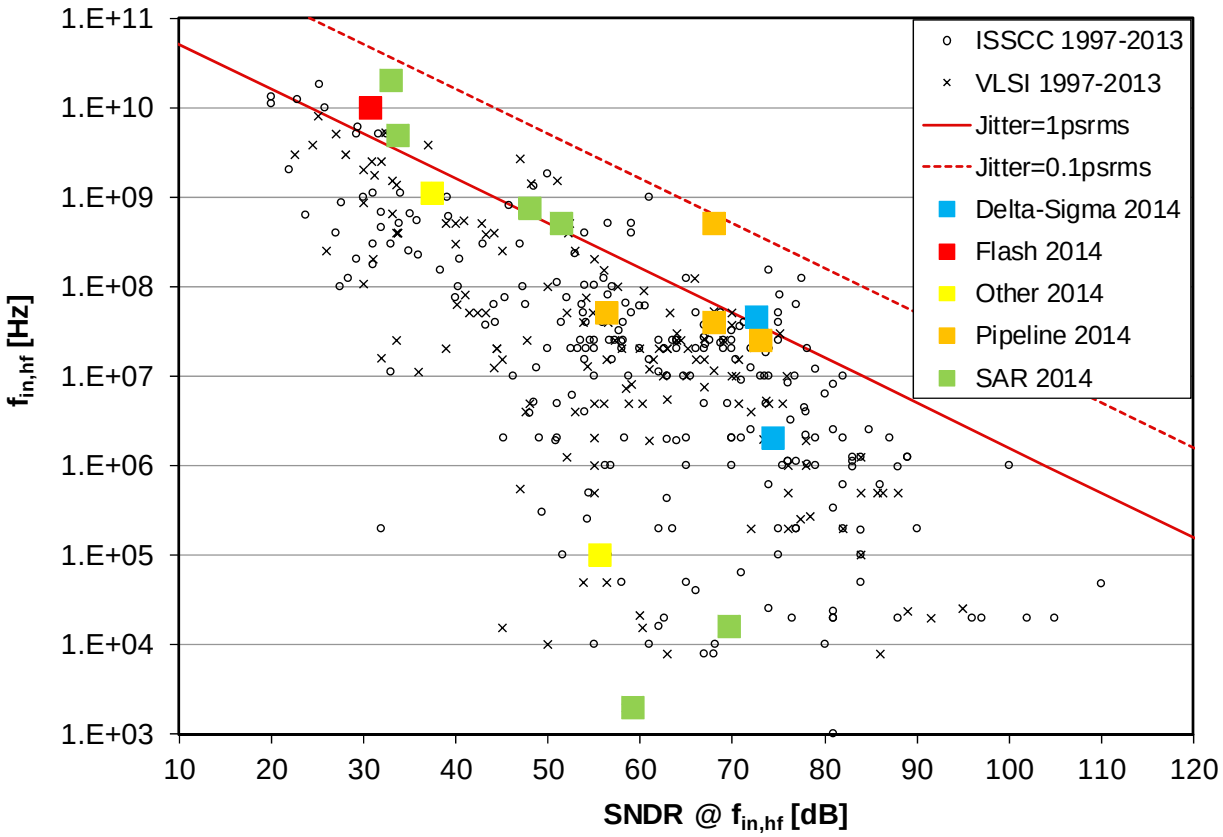

HISTORICAL TRENDS IN TECHNICAL THEMES

COMMUNICATION SYSTEMS

(RF SUBCOMMITTEE, WIRELESS SUBCOMMITTEE, WIRELINE SUBCOMMITTEE)

RF Subcommittee – 2014 Trends

Subcommittee Chair: *Andreia Cathelin, STMicroelectronics, Crolles, France*

Introduction:

This year, ISSCC 2014 will show increased innovation, integration, and technical maturity across the RF frequency bands from a few GHz to above 500GHz. In addition, 28nm CMOS technology is being adopted for the first time to wireless transceivers. This exposition outlines the emerging RF trends that will be displayed at ISSCC 2014: the drive toward increasing levels of integration will be presented. This trend can be seen in all areas of RF design, from cellular and wireless sensors to mm-Wave and imaging.

In cellular applications, efforts continue to increase the efficiency of integrated transmitters and PAs, moving towards multipath solutions, such as Doherty or Envelope Elimination and Restoration topologies. In addition to that, the power transfer to the antenna is maximized through closed-loop architectures that detect the effective RMS output-power transmitted, or allow tunable matching with the antenna. On the receiver side, the **multipath approach,** a well-established strategy to reduce noise of the input termination (using a noise cancelling topology), is appearing in the design of beam-forming and blocker-tolerant receivers, defining a new path towards the next generation of multi-standard transceivers.

The frequency synthesizer is still one of the most power-hungry elements in the entire transceiver; besides the continuing attempt to reach the best power efficiency, **mitigation of frequency pulling** and coexistence of multiple local oscillators are emerging as key aspects for the next generation of carrier aggregation systems. As well, **Multiplying Delay Locked Loops (MDLLs)** are becoming popular in wireless applications in the attempt to replace the large PLLs based on LC-tank oscillators. This year, at ISSCC 2014, for the first time, a fractional MDLL will be presented showing performance in terms of phase-noise comparable to the state-of-the-art traditional PLL-based approaches.

In mm-Wave designs, now that complete systems can be easily integrated, **the frequency of operation is moving rapidly towards THz**, where imaging applications are waiting for fully integrated solutions. 65nm CMOS is establishing itself as the most used technology for RF even above 100GHz, where, however, several designs still use SiGe for its higher f_t, f_{max},and breakdown voltage. Overall, solutions presented at ISSCC 2014 confirm that RF devices will continue to see larger levels of integration at the chip and package level for years to come. Over the past decade, the papers submitted to ISSCC have indicated clear trends in the continuing push to higher frequencies of operation in CMOS and BiCMOS. Thus, this year, the psychological barrier of 0.5THz has been overcome, looking towards THz applications! With the consolidation of mm-Wave designs, multipath solutions are emerging above 100GHz with phase array designs that take advantage of shorter wavelengths for compact single-chip solutions.

These trends are exemplified by presentations at ISSCC 2014 as described below:

Complexity and Maturity in the mm-Wave and sub-mm-Wave Ranges:

The high cutoff frequency of bipolar transistors and highly downscaled MOS transistors enables the realization of circuits and systems operating in the mm-Wave range. In the past few years, high-data-rate communication in the 60GHz band and automotive radar around 77GHz have garnered much attention. The 100GHz barrier for the operating frequency of silicon circuits was broken a few years ago. Whereas initially elementary building blocks like a VCO and an amplifier operating above 100GHz have been realized, we now witness the trend of increasing complexity in circuits operating above 100GHz, aiming at imaging and spectroscopy.

While the integration level in these domains is high, we see an improvement in the performance of individual building blocks: the output power of mm-Wave and sub-mm-Wave sources and PAs increases, (see Figures) and VCOs are operating at ever-increasing frequencies with a higher tuning range. This year, PAs are especially designed to minimize AM-PM distortion, and to be compliant with wideband modulations. In addition, the frequency of operation has increased above 300GHz while maintaining efficiency and output power comparable with the state-of-the-art.

Co-Existence and Efficiency for Cellular Applications:

Recevier Linearization: In the past few years there has been increasing interest in techniques to improve the linearity of receivers. Improved linearity will ease the requirements on the RF filtering of out-of-band blockers which can then be accomplished by adopting frequency-translated BPF and noise cancelling techniques.

Efficiency: At ISSCC 2014, PA efficiency improvements demonstrated will directly impact the battery life in portable applications. These techniques include antenna tuning, envelope elimination and restoration. and the Doherty architecture.

Digitally-Assisted RF: The trend towards digitally-assisted RF continues and is increasingly applied in mm-Wave chips. Such calibration techniques are being demonstrated in order to improve the overall performance of the transceiver by reducing the impact of analog impairments at the system level. This year, a new concept for improving the problematic VCO pulling effect will be presented.

VCOs: Without degredation of the phase-noise figure-of-merit (FOM), several techniques have been introduced to achieve wideband tuning in a compact area, and to overcome the problems of frequency pulling. The next Figure shows trends in VCO performance for some of the most significant VCOs published in the past decade.

PLLs: New architectures have been proposed to improve PLL performance. These techniques include divider-less PLLs, sub-sampling Fractional PLLs, MDLLs, and direct-digital frequency synthesizers.

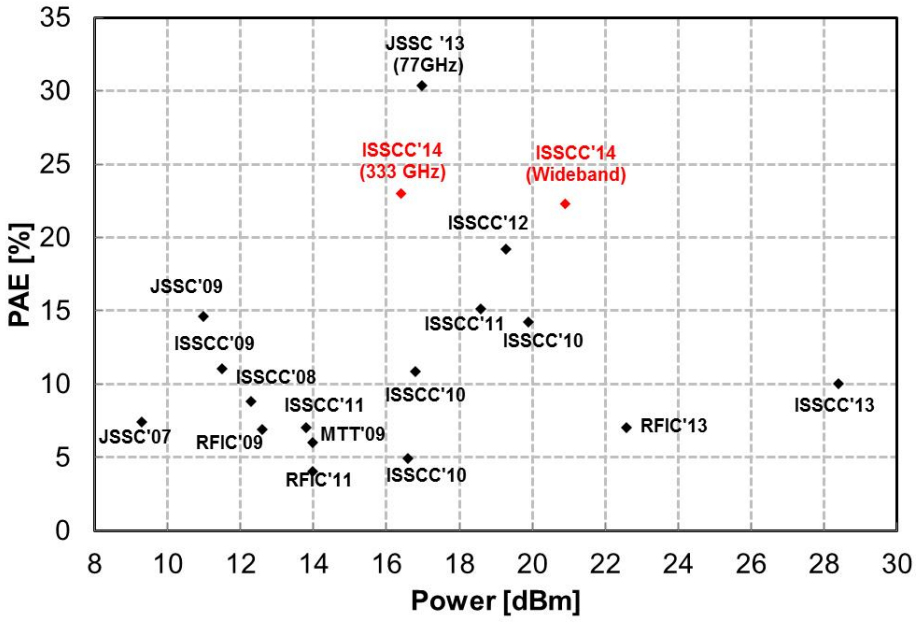

PAE (%) vs. output power for recent submicron mm-Wave CMOS PAs
The circuits to be revealed at ISSCC 2014 operate at very high frequency and/or are wideband compared to previous designs.

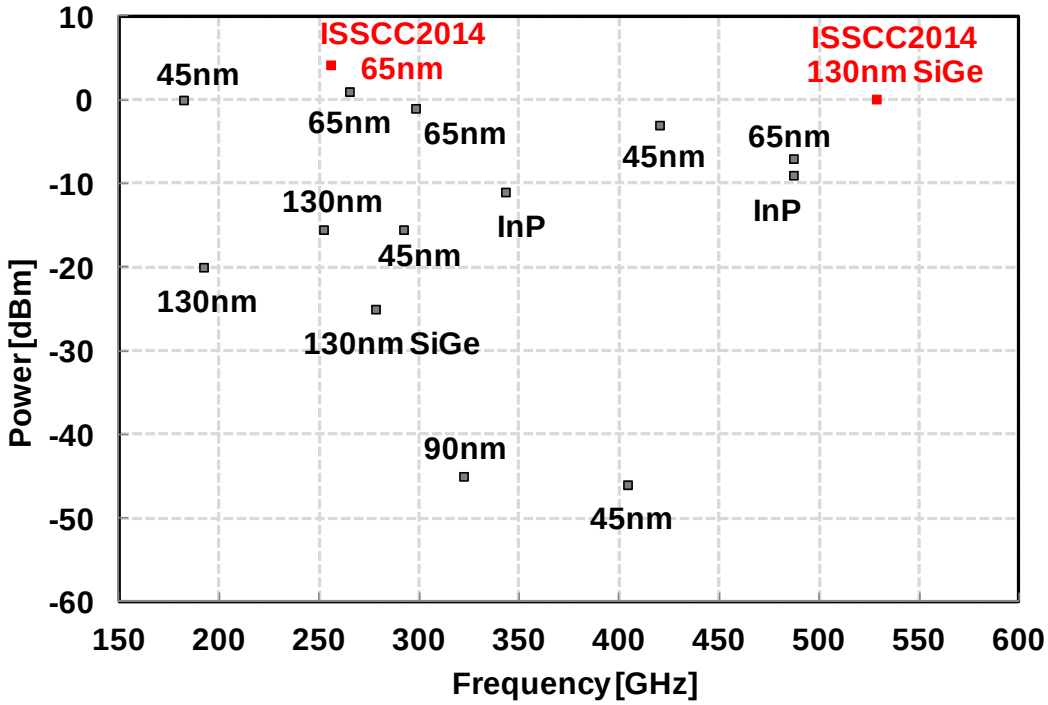

Output power versus frequency for mm-Wave and sub-mm-Wave sources.
Record output power levels are being revealed at ISSCC 2014.

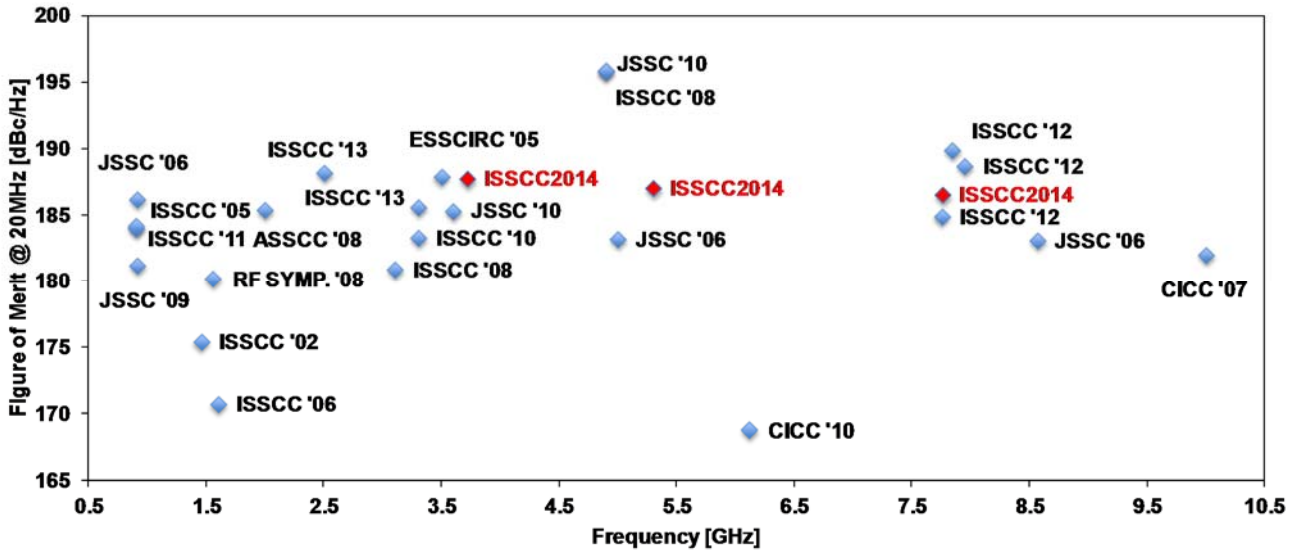

Phase-noise FOM at 20MHz offset frequency versus oscillation frequency. While conventional FOMs remain at competitive levels, the circuits revealed at ISSCC 2014 are digitally assisted to reduce the impact of analog impairment effects on performance.

Wireless – 2014 Trends

Subcommittee Chair: *Aarno Pärssinen, Broadcom, Finland*

The exponential increase of data-rate of cellular devices is a major driver of new wireless standards as reflected in the first Figure. The latest generation of cellular wireless standards, such as Long Term Evolution (LTE), includes multiple frequency bands, a wide range of channel bandwidths, and variousduplexing and diversity schemes, as well as new features, such as carrier aggregation. Fundamental trade-offs in data-rate, power consumption, functionality, and cost, continue to drive research and innovation. Despite the added complexity, the power consumption of cellular transceivers continues to improve thanks to semiconductor technology scaling and the exploitation of techniques, such as envelope tracking in the transmitter. Circuit innovations in various aspects such as highly-linear receiver front-ends and digital calibration enable inclusion of more features in a compact complex multi-function System-on-a-Chip (SoC).

The exponential increase in data-rate is also evident in wireless connectivity devices, including Wireless Local Area Networks (WLAN), as shown by the maximum achievable data rates of wireless-connectivity standards in the next Figure . The latest WLAN SoCs in the traditional 2.4/5GHz frequency bands achieve over 1Gb/s data-rate due to increased channel bandwidth, complex modulation formats, and larger numbers of antenna elements, in a Multi-Input-Multi-Output (MIMO) configuration. Wireless connectivity at a higher carrier frequency of 60GHz enables even higher data rates, thanks to larger available bandwidth. The mm-wavelength of these signals enables compact realization of phased arrays, leading to more robust and high-data-rate wireless links. At ISSCC 2014, with a 60GHz carrier frequency, commercial CMOS chipsets achieve 4.6Gb/s over 10m links, while CMOS front-end solutions reach 28Gb/s for proximity wireless communications.

A major vision for the future of electronics includes Wireless Sensor Networks (WSN) and the Internet-of-Things (IoT), where ubiquitous devices and sensors are embedded in, and harvest energy from, the environment, while staying connected through low-power low-cost short-range wireless. An intriguing application is the Body Area Network (BAN,) where a wirelessly-connected network of sensor nodes spread over the body, monitors human health. While early implementations of such radios used proprietary protocols, many newer implementations apply standards such as Zigbee and IEEE 802.15.6. The key challenge in these low-power radios is maintaining low energy consumption while preserving robustness to environmental change, especially due to the presence of other wireless signals.

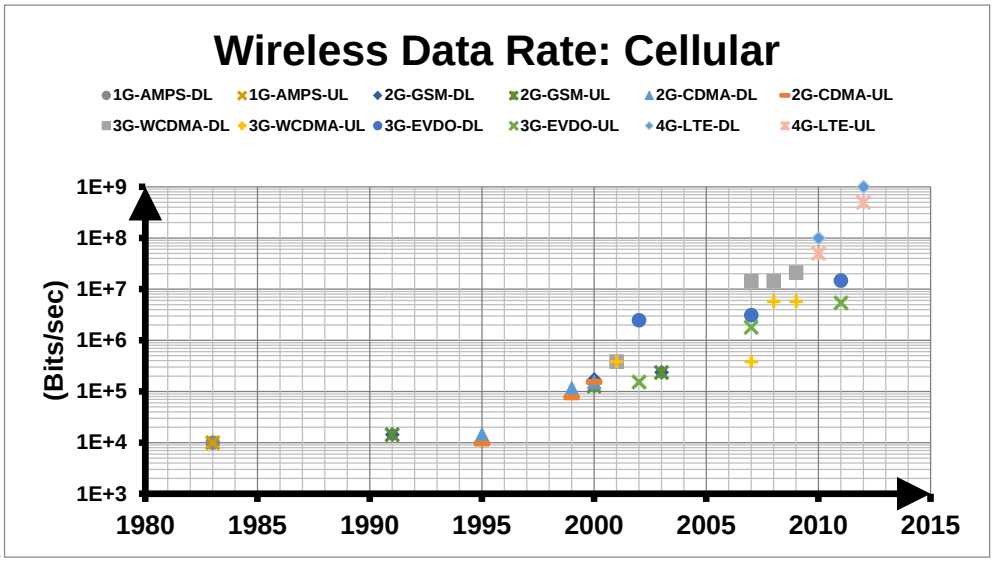

Data Rate of Cellular Standards

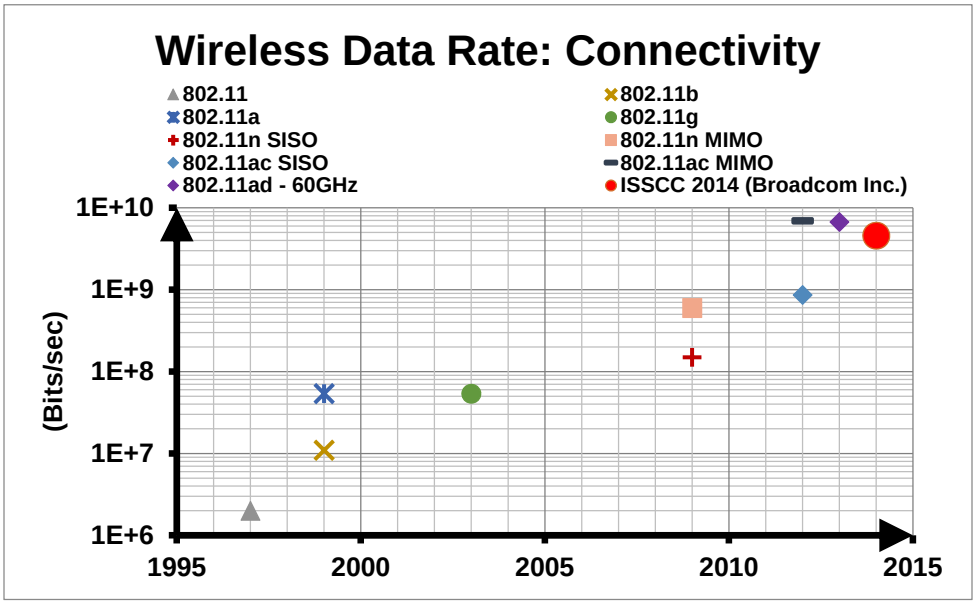

Data Rate of Wireless Connectivity Standards (802.11x). Here, the red dot represents a state-of-the-art circuit, to be described at ISSCC 2014.

Subcommittee Wireline – 2014 Trends

Subcommittee Chair: *Daniel Friedman, IBM T.J. Watson Research, Yorktown Heights, NY*

Over the past decade, wireline I/O has been instrumental in enabling the incredible scaling of computer systems, ranging from handheld electronics to supercomputers. During this time, aggregate I/O bandwidth requirements have increased at a rate of approximately 2× to 3× every 2 years. Demand for bandwidth is driven by applications including memory, graphics, chip-to-chip fabric, backplane, rack-to-rack, and LAN. 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. Increasing bandwidth has also been enabled by rapidly accelerating the per-pin data-rate. The first Figure shows that per-pin data-rate has approximately doubled every four years across a variety of diverse I/O standards ranging from DDR to graphics to high-speed Ethernet. Keeping pace with these standards, the second Figure shows data-rates for published transceivers enabled in part by process-technology scaling. However, continuing with this rather amazing trend for I/O scaling will require more than just transistor size reduction: Significant advances in both energy efficiency and signal integrity must be made to enable the next generation of low-power and high-performance computing systems.

Energy Efficiency and Interconnect Density:
Power consumption for I/O circuits is a first-order design constraint for systems ranging from cell phones to servers. As the pin count and per-pin data-rate for I/Os has increased on a die, so has the percentage of total power that they consume. Technology scaling enables increased clock and data-rates, and offers some energy-efficiency improvement, especially for digital components. Many recent advances have reduced power for high-speed link components through circuit innovation, including low-power RX equalization (DFE, CTLE), CMOS and resonant clocking, low-swing voltage-mode transmitters and links with low-latency power-saving states. This year, ISSCC 2014 includes the lowest-reported long-range 28Gb/s transceiver, achieving 20pJ/b [2.1], and 2 DFE receivers operating at 0.25pJ/b [2.4, 2.5].

Simply increasing per-pin data-rates with existing circuit architectures and channels is not always a viable path given fixed system power limits. The third Figure shows the energy efficiency (expressed in mW/Gb/s, which is equivalent to pJ/b) as a function of channel loss, for recently reported transceivers. The data indicate that the scaling factor between link power and signaling loss is approximately unity. In other words, the required link power approximately doubles with every additional 6dB of channel loss. At ISSCC 2014, a new coded signaling scheme is presented that transmits 8 bits over 8 wires enabling 96Gb/s at 4.3pJ/b/wire over 15dB loss channels [26.3]. As a result, the pin efficiency is doubled, making the approach suitable for pin-constrained systems such as memory interfaces. In an alternative approach [26.2], the combined optimization of channel and connector, a scalable bi-directional serial transceiver at 32Gb/s/lane achieves an efficiency of 6.3pJ/b while equalizing 17dB channel loss.

Electrical Interconnect:
Some types of channels, especially those related to medium-distance electrical I/O such as in server backplanes, must support high data-rates over high-loss channels. For these links, the key to scaling has been improvements in equalization and clock/data recovery. Recent high-speed transceivers use a combination of fully adaptive equalization methods, including TX FIR, RX CTLE, DFE, as well as RX FIR and/or IIR FFEs. As a result, recent transceivers achieve data-rates above 20Gb/s with channels that have 30dB or more loss. This year, ISSCC 2014 includes a 28Gb/s transceiver operating over a 30dB-loss channel, as well as a paper that demonstrates a 60Gb/s transmitter — the highest transmit data-rate reported to date. As well, ISSCC 2014 includes one of the fastest aggregate links presented to date, specifically, 320Gb/s (20Gb/s × 16) [26.1].

Optical Interconnect:
As the bandwidth demand for traditionally electrical wireline interconnects has accelerated, optics has become an increasingly attractive alternative for interconnects within computing systems. Optical communication offers clear benefits for high-speed and long-distance interconnects. Relative to electrical interconnects, optics provides lower channel loss. Circuit design and packaging techniques that have traditionally been used for electrical wireline are being adapted to enable integrated optical with extremely low power. This trend has resulted in rapid progress in optical ICs for Ethernet, backplane and chip-to-chip optical communication. ISSCC 2014 includes a 2-dimensional (12×5) optical array achieving an aggregate data-rate of 600Gb/s [8.2]. Pre-emphasis using group-delay filtering extends the useful date rate of a 25Gb/s VCSEL to 40Gb/s [8.9]. Additional examples of low-power-linear and non-linear equalizers tackle electronic dispersion compensation in multi-mode and long-haul cables [8.1, 8.3].

Concluding Remarks:
Continuing to aggressively scale I/O bandwidth is both essential for the industry and extremely challenging. Innovations that provide higher performance and lower power will continue to be made in order to sustain this trend. Advances in circuit architecture, interconnect topologies, and transistor scaling are together changing how I/O will be done over the next decade. The most exciting and most promising of these emerging technologies for wireline I/O will be highlighted at ISSCC 2014.

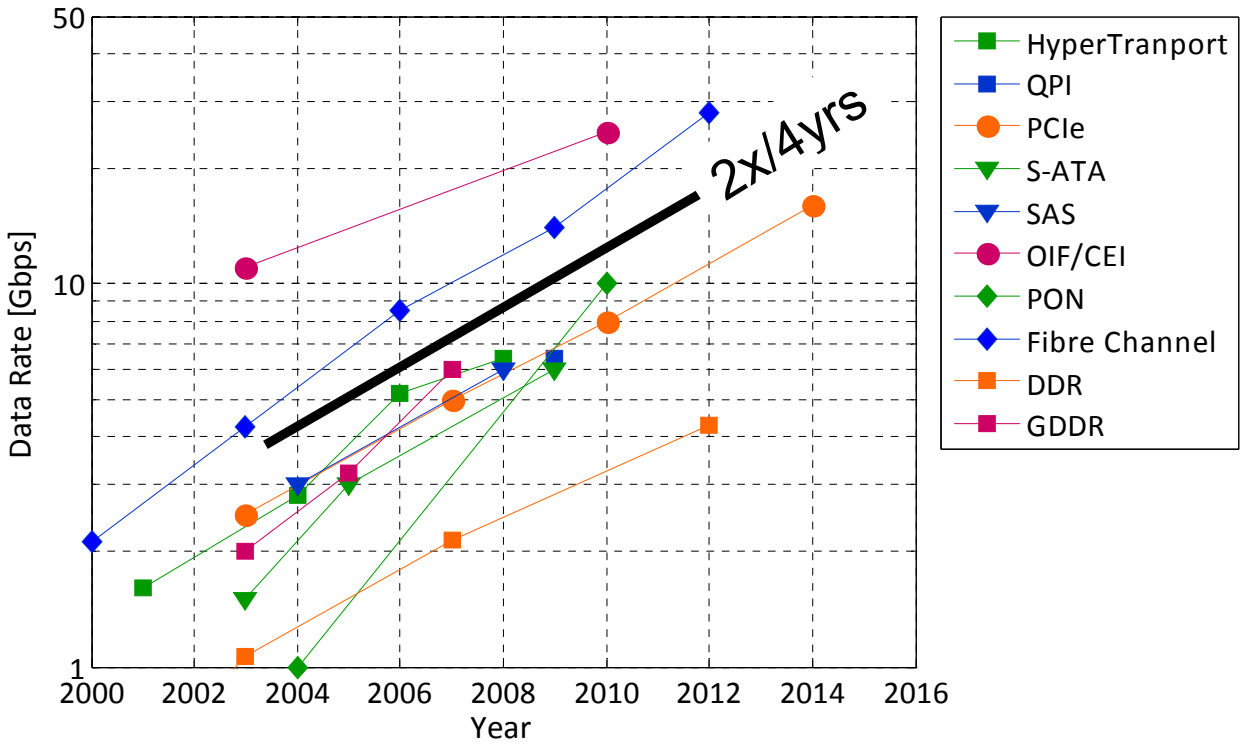

Per-pin data-rate vs. year for a variety of common I/O standards.

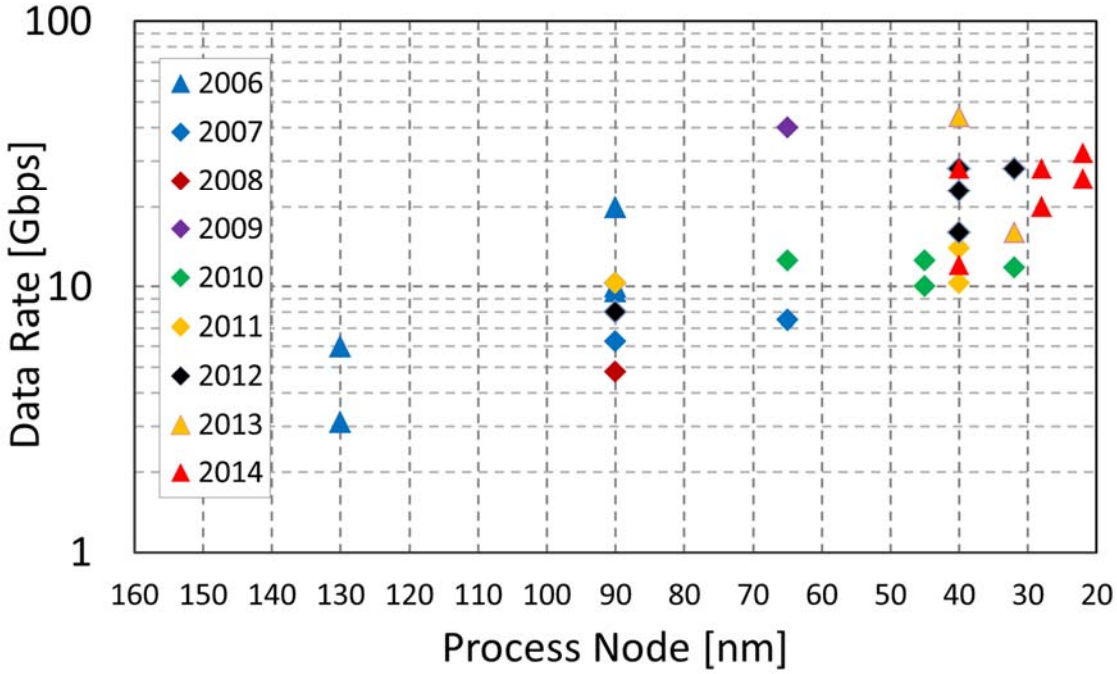

Data-rate vs. process node and year.

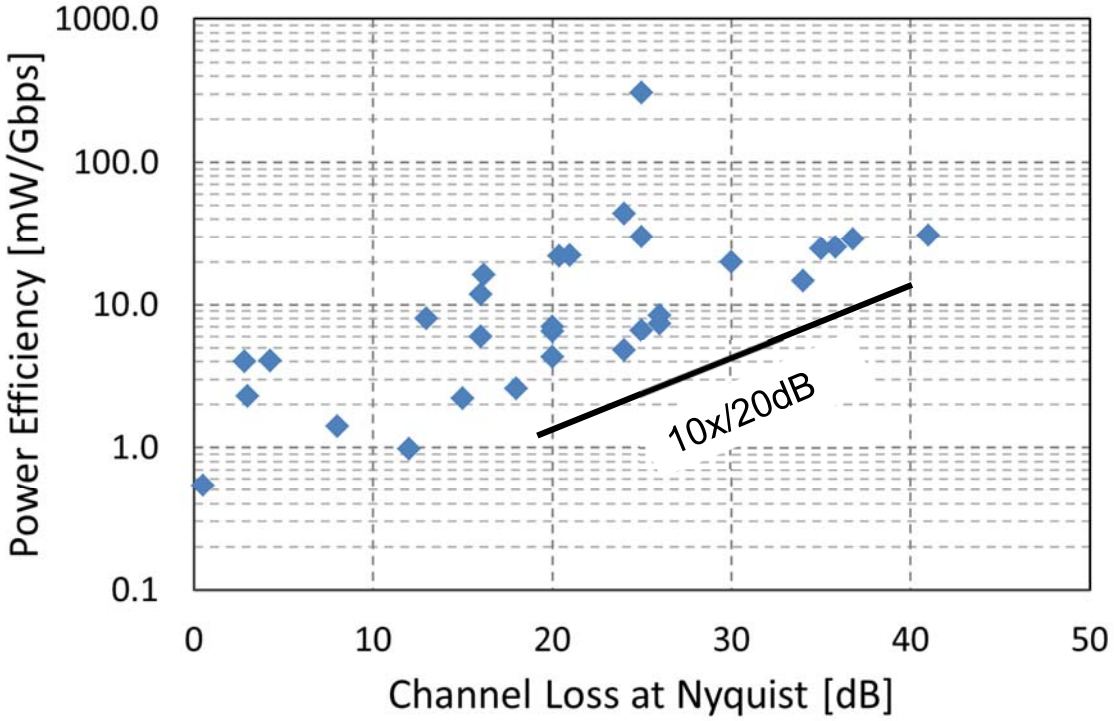

Transceiver Power Efficiency vs. Channel Loss.

HISTORICAL TRENDS IN TECHNICAL THEMES

Digital Systems

(Energy-Efficient Digital Subcommittee
High-Performance Digital Subcommittee
Memory Subcommittee)

Energy Efficient Digital – 2014 Trends

Subommittee Chair: *Stephen Kosonocky, AMD, Fort Collins, CO*

The demand for ubiquitous mobile functionality to achieve enhanced productivity, a better social-networking experience, and improved multimedia quality, continues to drive innovation in systems-on-chip, while also improving battery life and lowering cost. While the performance of embedded 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 of emerging applications. At the other end of the spectrum, sensor nodes for the Internet-of-Things require low-energy wireless and 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 unveiled at ISSCC 2014 are built in a variety of technologies, with best-in-class results accomplished with higher integration, and improved performance-per-watt. These are demonstrated in various processes ranging from 90nm down to 28nm bulk, and SOI CMOS technologies.

Adaptive, near-threshold operation paves the road to further supply voltage reduction. A clock frequency of 460MHz for DSP processing is now possible at a supply voltage of 400mV, with a processor implemented in 28nm UTBB FDSOI. This supply voltage will further decrease in the near future thanks to advanced technology development and novel circuit techniques that compensate for process variations. Explicit trends in voltage and frequency are difficult to quantify because of the variety of applications that are addressed. Consequently, the operations performed by the different systems are not directly comparable in terms of energy or clock frequency.

The first Figure illustrates the major trends in smartphones and tablets relevant to energy-efficient digital circuits. In the late 1990s, a GSM phone contained a simple RISC processor running at 26MHz, supporting a primitive user interface. After a steady increase in clock frequency to roughly 300MHz in the early 2000s, there has been a sudden spurt towards 1GHz and beyond. Moreover, following trends in laptops and desktops, processor architectures have become much more advanced, and recent smartphones incorporate four- and even eight-core processors, running up to 2GHz. Battery capacity, mostly driven by the required form factor, as well as thermal limits, implies a power budget of roughly 3W for a smartphone (including the power amplifier for cellular communications and the display). The available power budget for everything digital has been holding steady in the range of 2W (peak) to 1W (sustained). As a result, 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 towards dedicated, optimized hardware solutions. Some new areas where dedicated processors are particularly needed include gesture-based user interfaces, augmented reality. and computational imaging, to name a few. For all digital circuits, the limited power budget leads to more fine-grained clock gating, various forms of adaptive voltage-frequency scaling, variable device threshold-schemes, and elaborate power and thermal management.

The second Figure 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 data-rate every five years, with no sign of abatement. With essentially-constant power and thermal budgets, energy efficiency has become a central theme in designing digital circuits for mobile processing. Historically, CMOS feature sizes have halved every five years. For a brief period in the 1990s, CMOS scaling (a.k.a. Dennard scaling) provided a 2^3 (α^{-3}) increase in energy efficiency every five years, almost matching the required 10×. During the past decade, however, CMOS scaling offers a roughly 3× improvement in energy efficiency every five years. The resulting ever-widening gap has spawned 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 above.

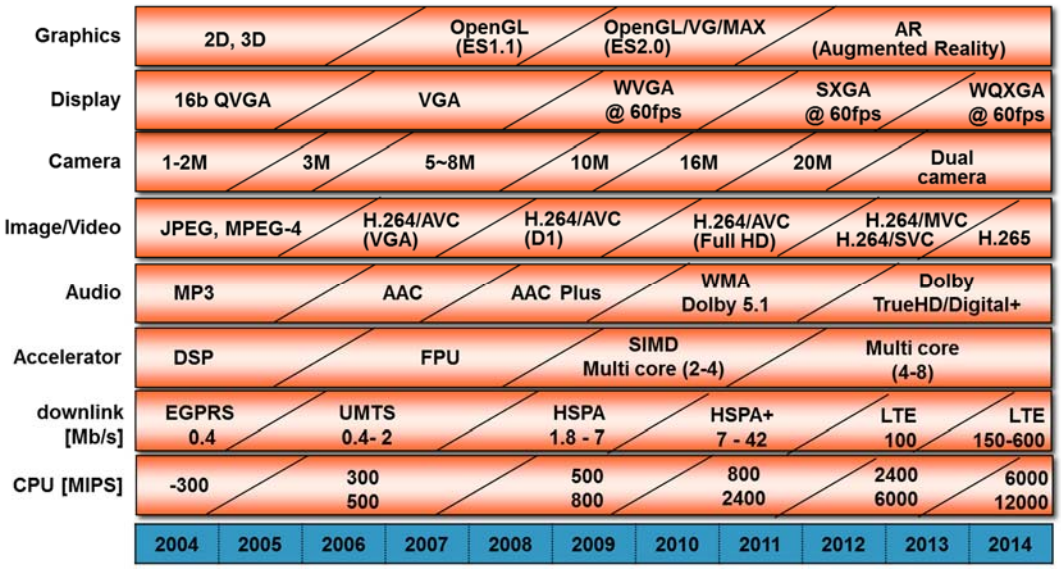

Application processor trends in smart phones.

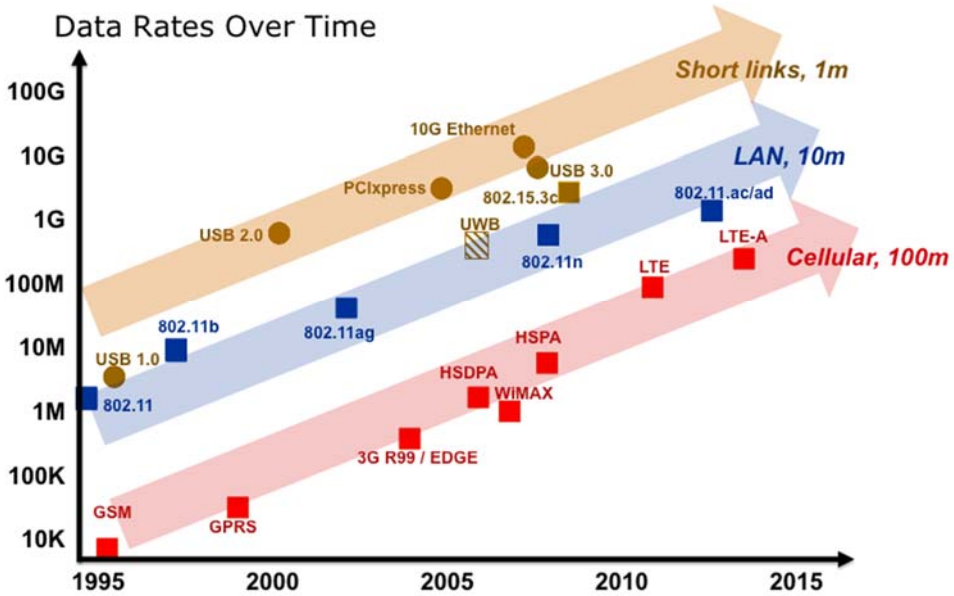

Wireless and wired data-rates over time.

High-Performance Digital – 2014 Trends

Subcommittee Chair: *Stefan Rusu, Intel, Santa Clara, CA*

The relentless march of process technology brings increasing integration and energy-efficient performance to enterprise and cloud servers. ISSCC 2014 features IBM's 12-core, 96-thread POWER8™ processor in 22nm SOI, with 96MB of eDRAM shared L3 cache, all employing 4.2B transistors, and offering up to 2.5× higher socket performance over its 32nm POWER7+™ predecessor . Intel's 15-core, 30-thread next-generation Xeon® server processor in 22nm tri-gate technology with 37.5MB shared SRAM L3 cache integrates 4.31B transistors. Intel's Haswell processor in a 22nm tri-gate process introduces a 128MB multichip package eDRAM L4 cache to boost integrated graphics performance.

The chip complexity chart below 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 4B transistors on a die.

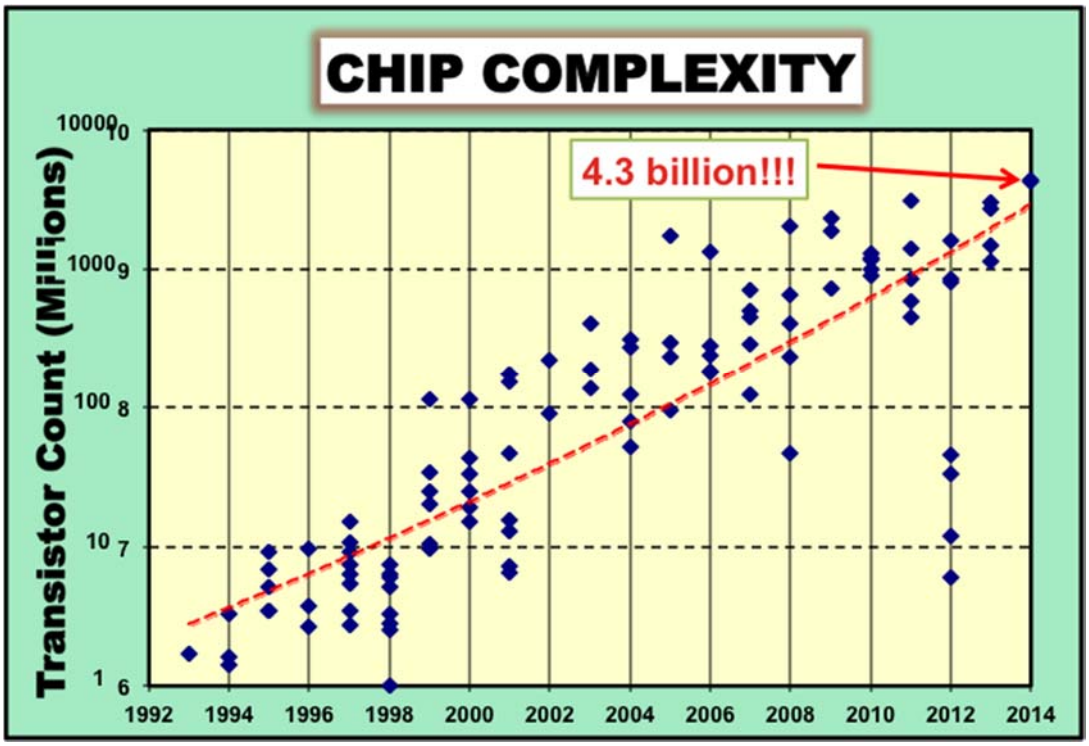

The maximum core clock frequency seems to have saturated in the range of 5 to 6GHz, primarily limited by thermal considerations. Thus, the IBM POWER8™ processor core reports a nominal 5GHz+ operation. The energy-efficient 64b ARM v8 processor for micro-servers operates nominally at 3GHz in 40nm bulk CMOS.

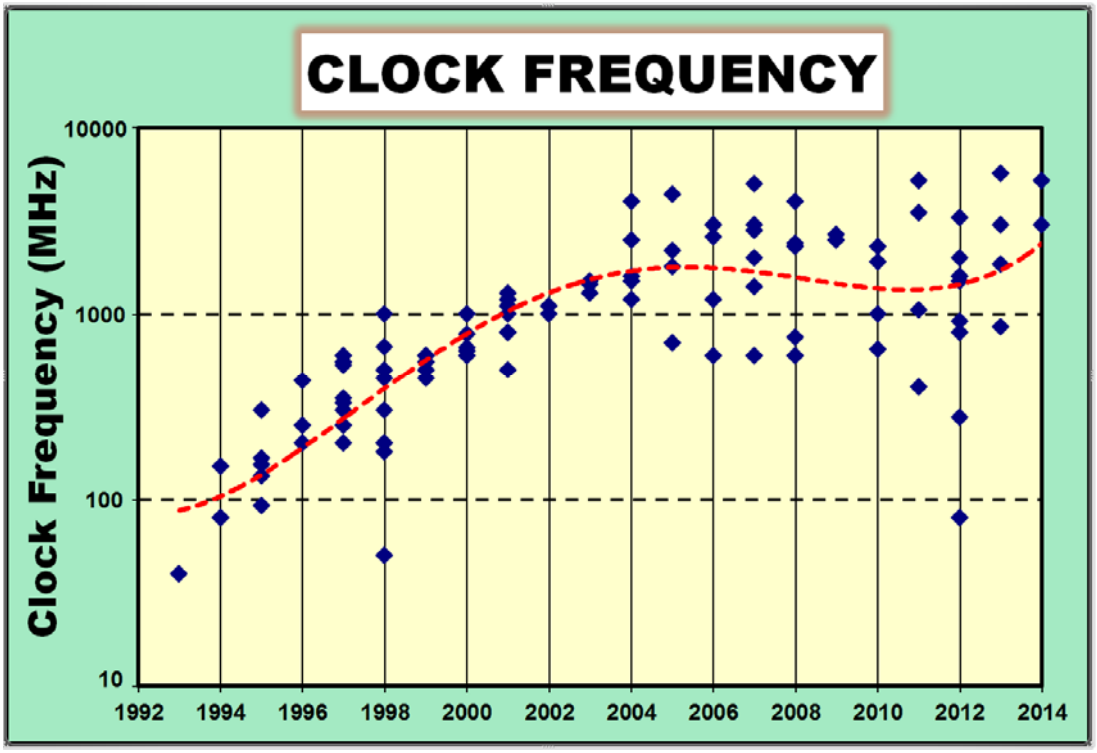

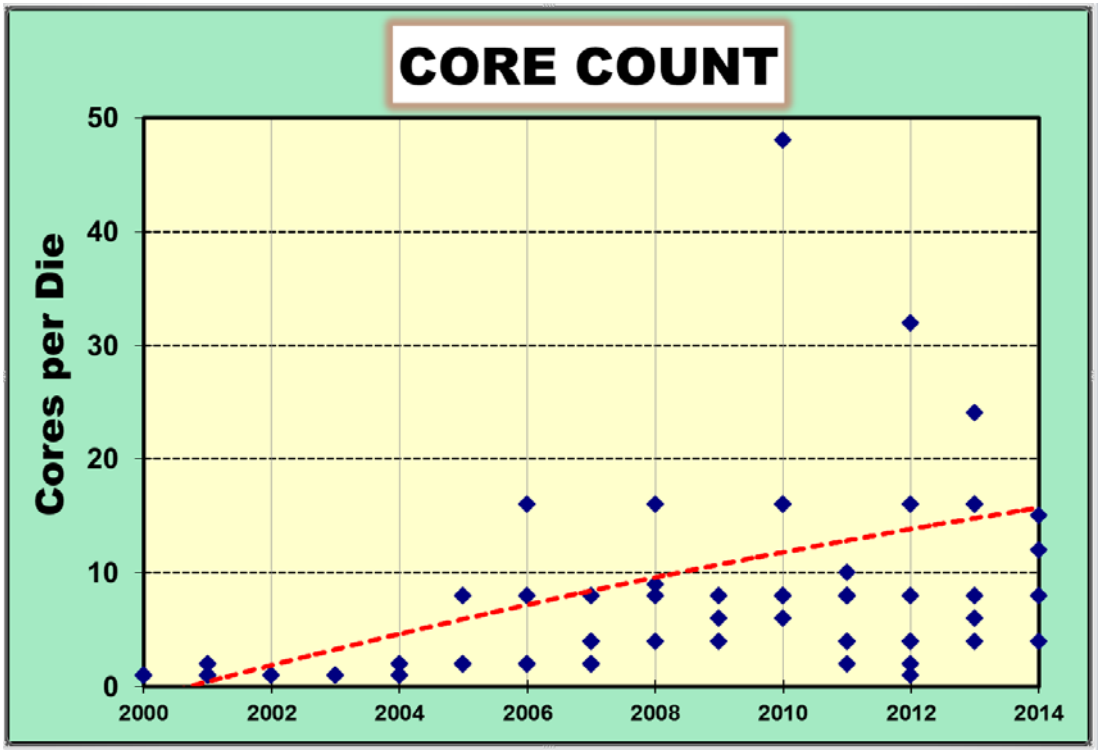

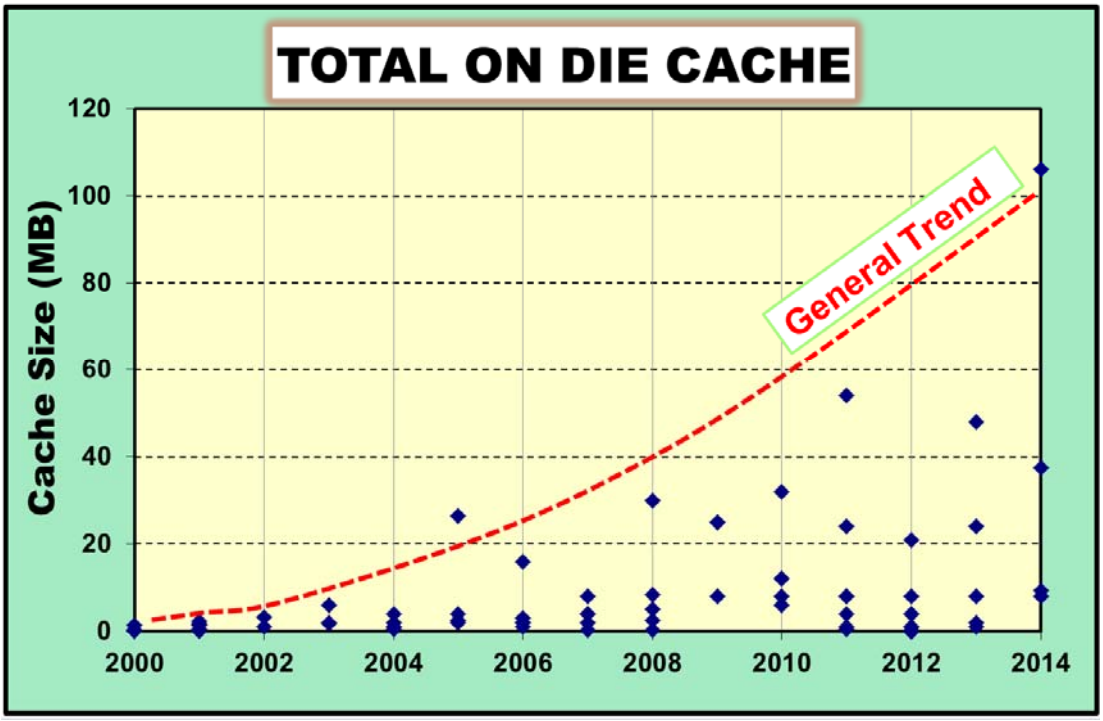

This year, there is renewed focus on improving overall energy efficiency of integrated systems-on-chip via use of more independently controlled voltage and clock domains using smart on-die control. IBM uses many integrated micro-linear regulators on die in the POWER8™, and Intel demonstrates fully-integrated buck voltage regulators on the Haswell. Adaptive clocking is used in IBM's POWER8™, and in AMD's Steamroller chips to reduce the impact of process-voltage-temperature (PVT) variation on energy efficiency and performance. Resonant clocking that also supports dynamic voltage-frequency scaling (DVFS) is used in both of the chips to reduce clock-grid power.

Another trend evident this year is the continued emergence of digital phase-locked loops (PLL) and delay-locked loops (DLL) to better exploit nanometer feature-size scaling, and reduce power and area. Through use of highly innovative architectural and circuit design techniques, the features of these digital PLLs and DLLs have improved significantly over the recent past: The diagram below shows the jitter performance vs. energy cost for PLLs and multiplying DLLs (MDLL).

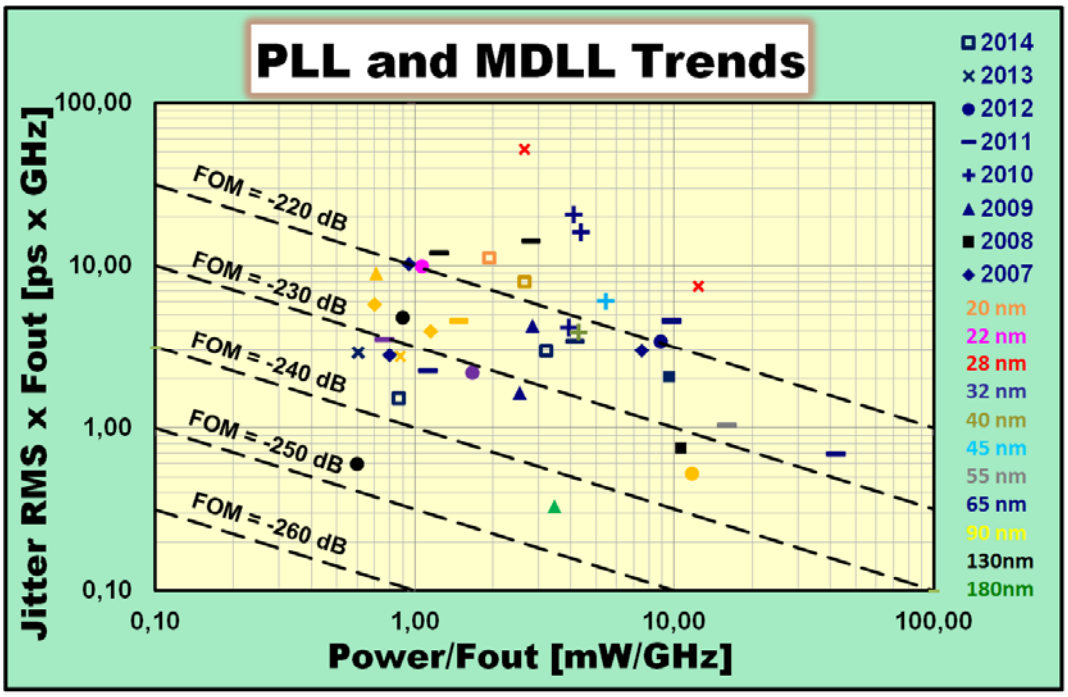

Overall, digital processors continue to grow in complexity, with additional focus on integrated fine-grain power management for better energy efficiency. We observe that traditional analog building blocks are being implemented using digital techniques to cope with variability and ease scaling to finer geometries.

MEMORY – 2014 Trends

Subcommittee Chair: *Kevin Zhang, Intel, Hillsboro, OR*

In memory systems, we continue to see progressive scaling in embedded SRAM, DRAM, and floating-gate-based Flash for very broad applications. However, with ever-increasing scaling challenges in all mainstream memory technologies, smart design algorithms and error-correction techniques are widely adopted to compensate growing device variability. FinFET technology is quickly becoming the mainstream logic technology to provide SRAM scaling, along with various read- and write-assist circuits in 22nm and beyond. Emerging memory technologies continue to make steady progress towards product intercepts, including PCRAM and ReRAM, while STT-MRAM is beginning to become a strong candidate for both stand alone and embedded applications.

Some state-of-the-art developments from ISSCC 2014 include:

- A 1Gb eDRAM using a 22nm tri-gate logic process and capable of being clocked at 2GHz
- A 128Gb 2 bit-per-cell NAND-Flash design using 3-dimensional cell technology with 24-WL stacked layers
- A 128Gb 2 bit-per-cell NAND-Flash design using 16nm planar cell technology
- A 128Mb SRAM designed in 14nm FinFet CMOS using a 6T bitcell with V_{min}-enhancement techniques
- An embedded ReRAM in 28nm capable of working down to 270mV
- A 1Gb 8 Channel 128GB/s High-Bandwidth Memory(HBM) DRAM
- A 3.2Gb/s/pin 8Gb 1.0V LPDDR4 SDRAM with integrated ECC engine

SRAM:

Embedded SRAM continues to be a critical technology enabler for a wide range of applications from high-performance computing and graphics to mobile applications. Challenges for SRAM include V_{min}, leakage and dynamic power reduction, while relentlessly following Moore's law to shrink the area by 0.5× for every generation of technology. As the transistor feature size diminishes below 20nm, device variation has made it very difficult to shrink the bit-cell size at the 50% rate while maintaining or lowering V_{min} between generations. Introduction of high-κ metal-gate (at 45nm) and FinFET or fully-depleted SOI transistors (at 22nm) reduces the V_{th} mismatch and have enabled further device scaling. This year, eDRAM is introduced at 22nm as a means for memory scaling in high-performance CPU designs. Design solutions such as read-/write-assist circuitry and variation-tolerant sensing schemes have been used to improve SRAM V_{min} performance starting at 32nm, and are now ubiquitous in SRAM designs below 20nm. Dual-rail SRAM design emerges as an effective solution to enable dynamic voltage-frequency scaling (DVFS) by decoupling logic supply rails from SRAM arrays, and thus allowing a much wider operating window. The use of assist-circuit techniques, FinFET transistors, and dual-rail architectures has extended the viability of using the high-density 6T SRAM bitcell down to 16nm/14nm designs.

It is important for SRAM to reduce both leakage and dynamic power, keeping products within the same power envelope at the next technology node. Sleep transistors, fine-grain clock gating, and clock-less SRAM designs have been proposed to reduce leakage and dynamic power. Aggressive uses of these techniques are resulting in designs with extremely low leakage which allows SRAM to replace non-volatile memories in selected applications. The Figure shows the bit cell and V_{dd} scaling trend of SRAM from major semiconductor manufacturers.

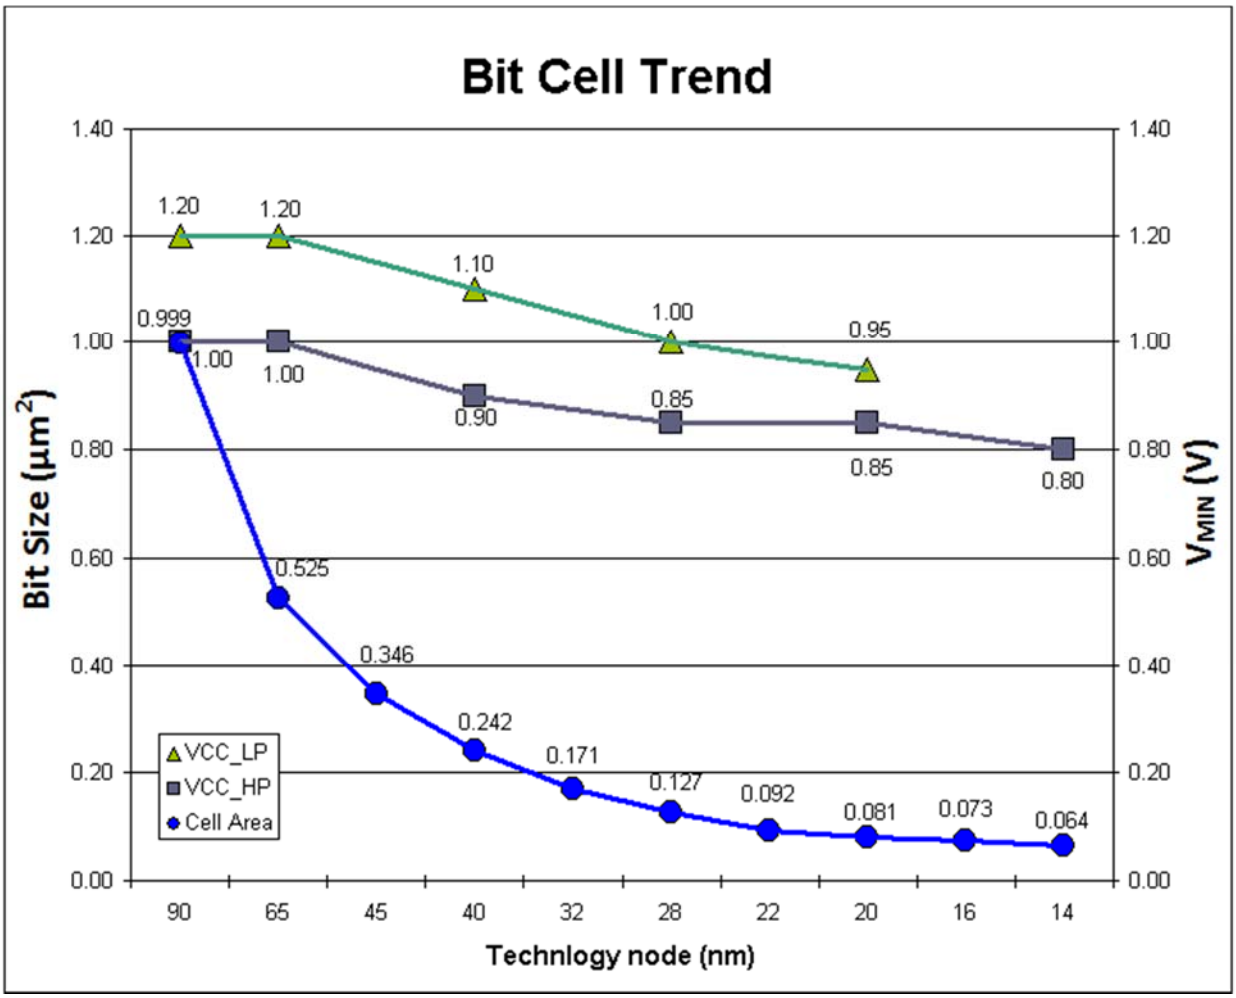

Bit cell and V_{dd} scaling trend of SRAM from major semiconductor manufacturers.

High-Bandwidth DRAM:

In order to reduce the bandwidth gap between main memory and processor performance, DRAM data rates continue to increase at the memory interface, using schemes such as DDRx, LPDDRx, and GDDRx, as shown in next **Error! Reference source not found.**. Currently, DDR4 and GDDR5 memory I/Os operate around 3Gb/s/pin and 7Gb/s/pin, respectively, which represent aggregate rates of 6GB/s and 28GB/s, respectively. To achieve higher bandwidth of 128GB/s, 4 TSV layers are stacked with 8 channels using a total of 1024 I/Os. This high-bandwidth memory (HBM) device is targeted for high-performance next-generation high-end devices including graphics and network applications. The LPDDR4 SDRAM achieves a data rate of 3.2Gb/s/pin with a 1.0V supply, and has an integrated ECC engine for sub-1V operation. For mobile graphics memory, a GDDR5M achieves low standby power of only 5.4mW using WCK auto-sync mode. Additional circuit techniques including clock timing skewing and error-adaptive duty-cycle correction are applied to improve signal integrity.

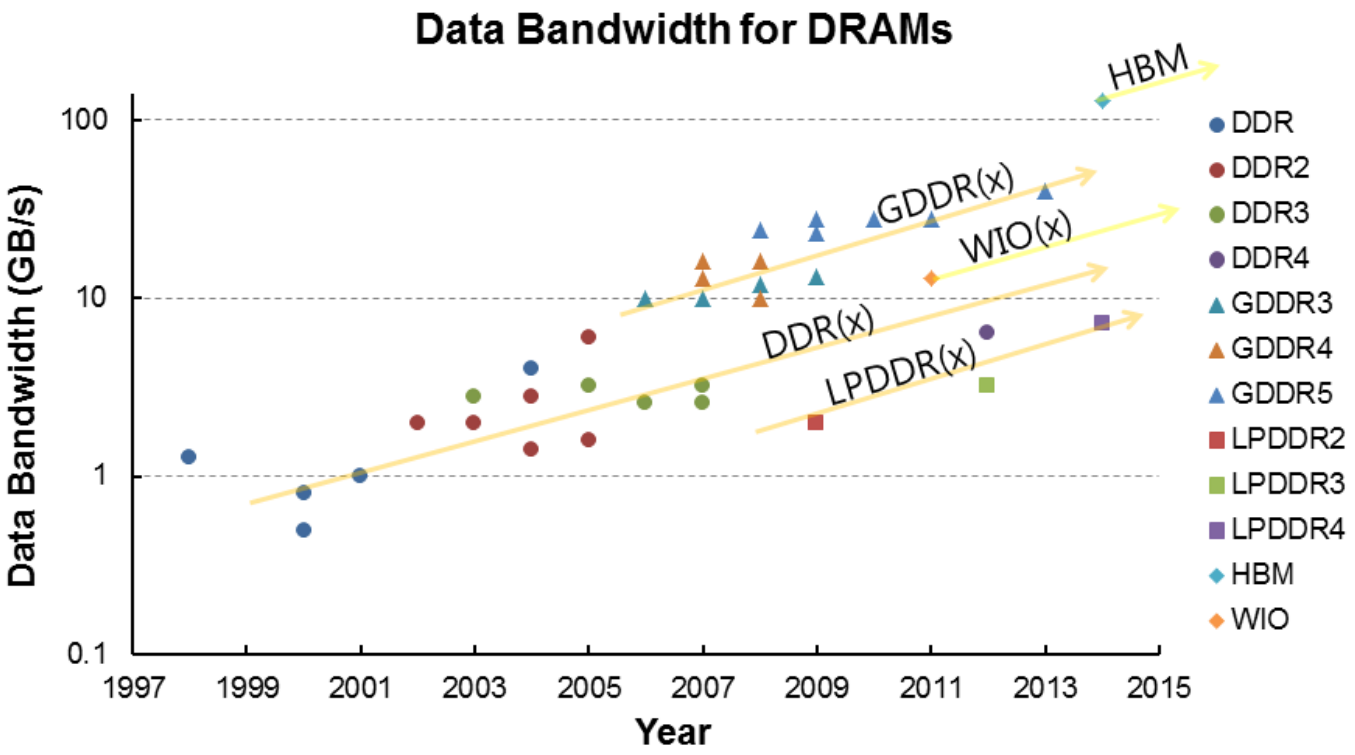

DRAM Data bandwidth trends.

Non-Volatile Memories (NVMs):

In the past decade, significant investment has been put into emerging memories to find an alternative to floating-gate based non-volatile memory. The emerging NVMs, such as phase-change memory (PRAM), ferroelectric RAM (FeRAM), magnetic spin-torque-transfer (STT-MRAM), and Resistive memory (ReRAM), are showing potential to achieve high cycling capability and lower power per bit in read/write operations. Some commercial applications, such as cellular phones, have recently started to use PRAM, demonstrating that reliability and cost competitiveness in emerging memories is becoming a reality. Fast write speed and low read-access time are the potential benefits of these emerging memories. At ISSCC 2014, a high-density ReRAM with a buried WL access device is introduced to improve the write performance and area. The next Figure highlights how MLC NAND Flash write throughput continues to improve. However, while the Figure following shows no increase in NAND Flash density over the past year, recent devices are built with finer dimensions or more sophisticated 3-dimensional vertical bit cells.

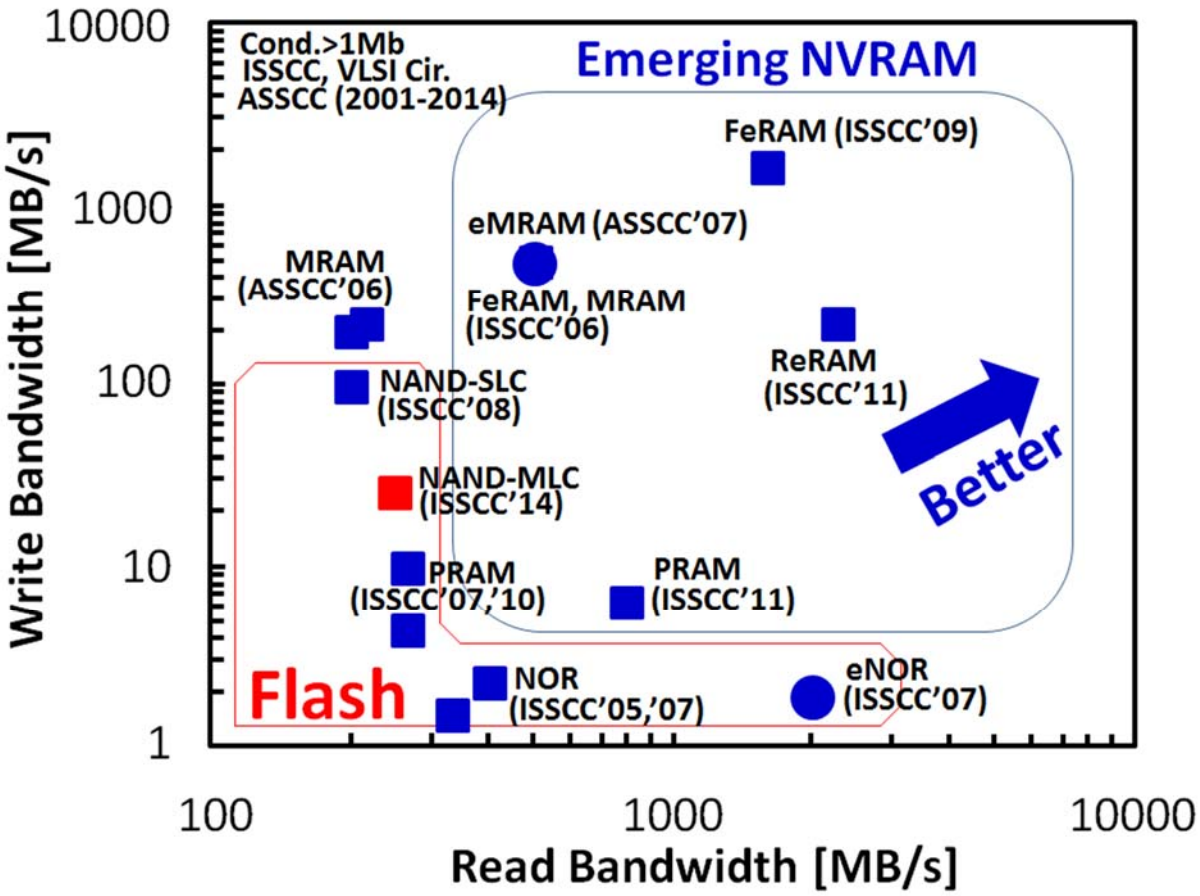

Read/write bandwidth comparison of nonvolatile memories.

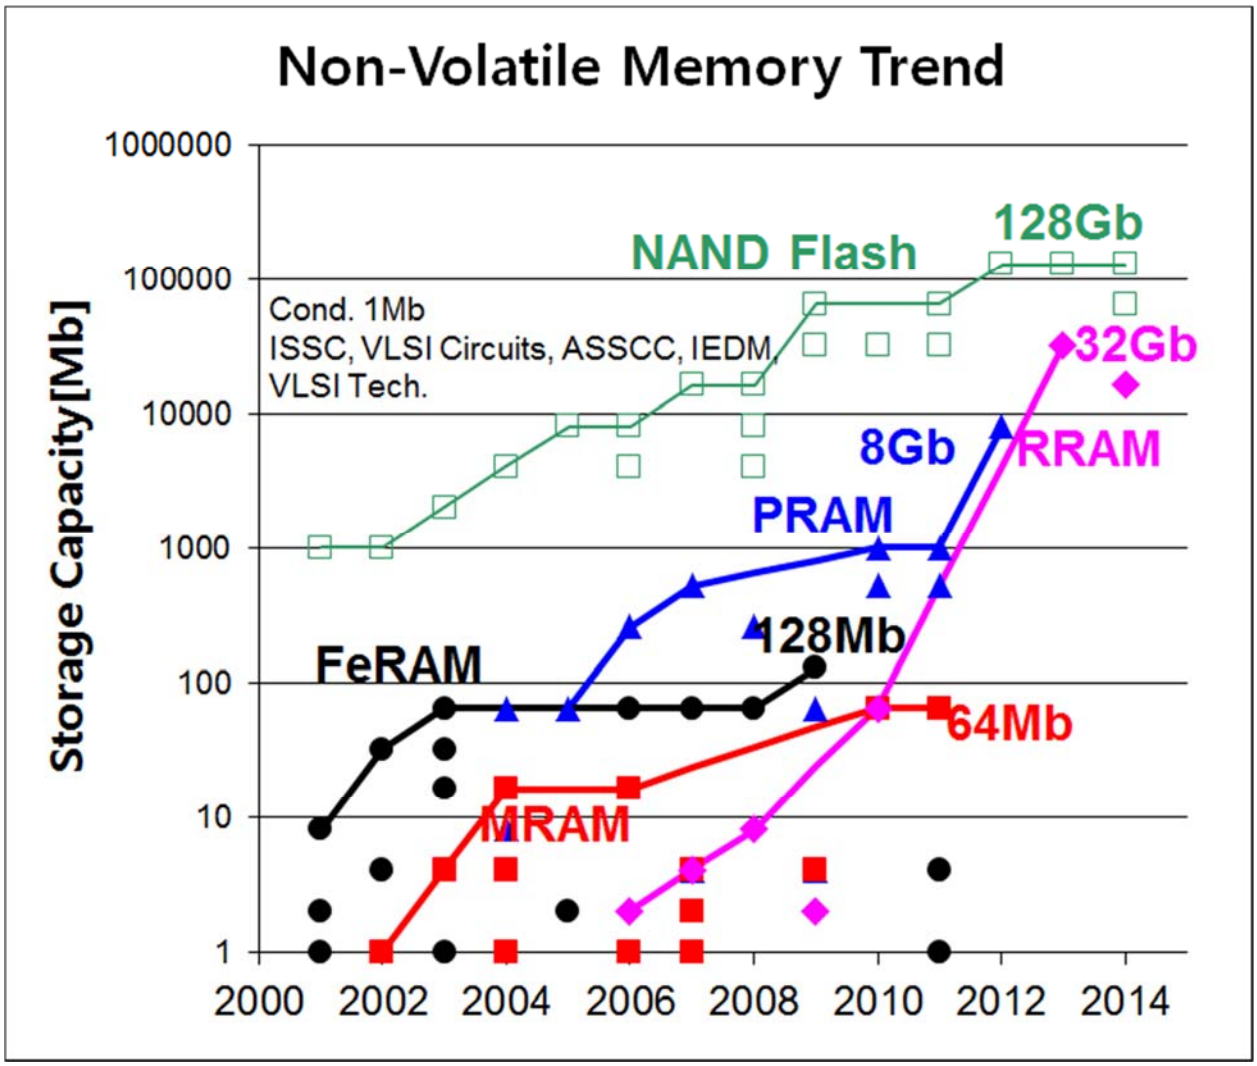

Memory capacity trend of emerging nonvolatile memories.

NAND-Flash Memory:

NAND-Flash memory continues to advance towards higher density and lower power, resulting in low-cost storage solutions that are replacing traditional hard-disk storage with solid-state disks (SSDs). Despite growing difficulties with further down scaling of planar cell technology, the latest 2D NAND Flash technology has scaled from 20nm down to 16nm.The next **Error! Reference source not found.** shows the observed trend in NAND Flash capacities presented at ISSCC over the past 20 years. Along with scaling, to overcome the increased device variability and error rates, sophisticated control algorithms to offset this trend have been implemented outside the NAND silicon in the system memory controller, especially ECC and data management methods. This year, to address the ever-growing scaling challenges, such as manufacturing cost increase and reliability degradation, in conventional 2D NAND Flash memories, 3D stacked NAND vertical gates is being introduced as an alternative technology solution to further increase NAND density, and lower manufacturing cost.

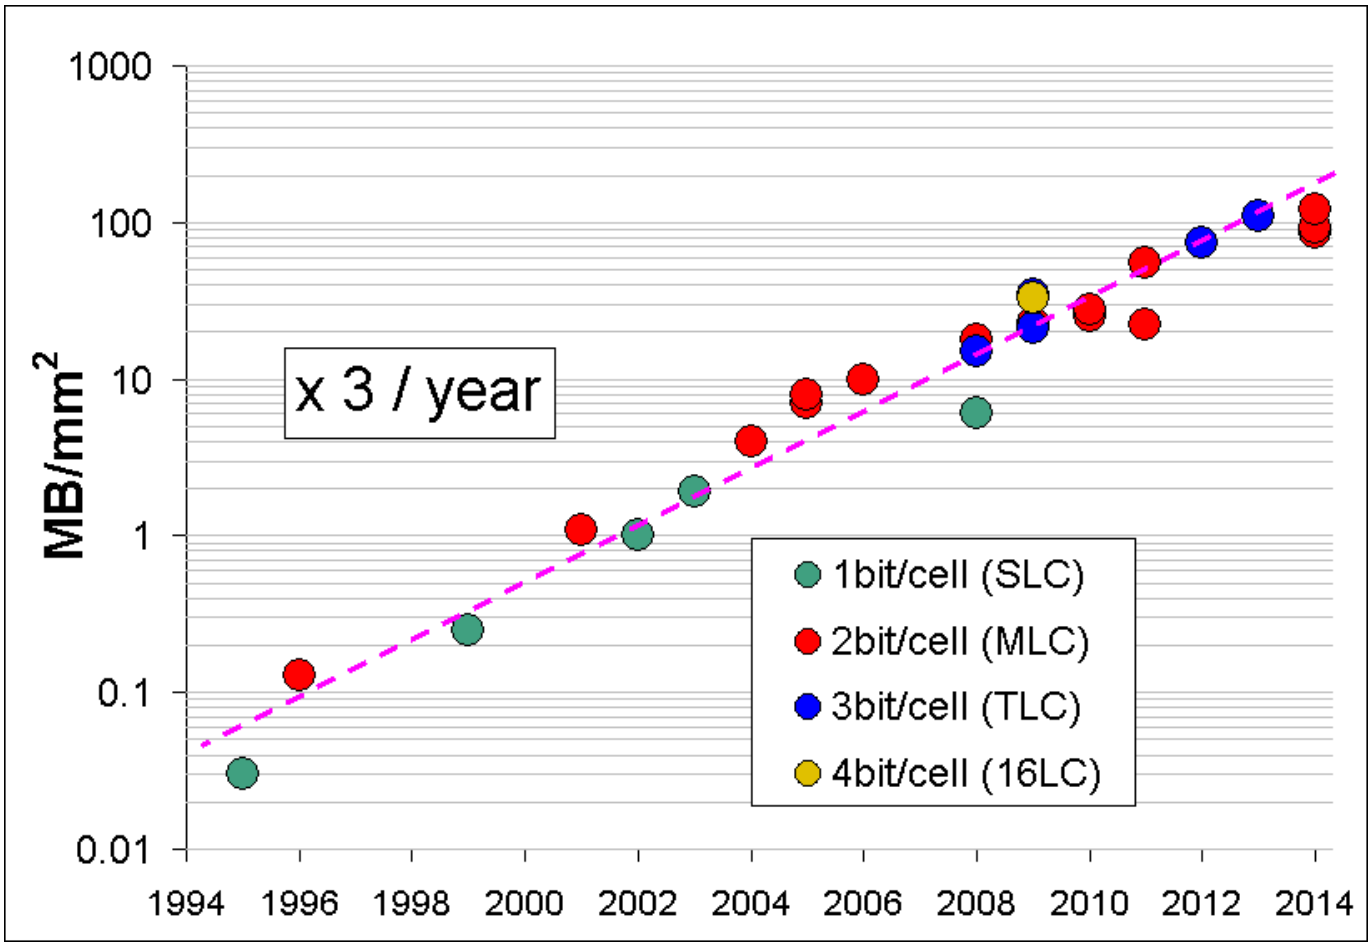

NAND-Flash memory trends.

HISTORICAL TRENDS IN TECHNICAL THEMES

Innovative Topics

(Imagers/MEMS/Medical/Displays Subcommittee
Technology Directions Subcommittee)

Imagers/MEMS/Medical/Displays – 2014 Trends

Subcommittee Chair: *Roland Thewes, TU Berlin, Berlin, Germany*

Imagers:

The CMOS-image-sensor business is one of the fastest-growing segments of the semiconductor industry. Key applications include cell-phone cameras, digital still cameras, camcorders, security cameras, automotive cameras, digital-cinema cameras, and gaming.

The resolution and miniaturization races are ongoing, and while the performance requirements stay constant, pixel size continues to scale down. Images of over 40M Pixels are commercially available. A column-parallel approach based on pipelined and multiple-sampling implementations has become standard for low-power, high-speed, low-noise camera and video applications. Backside illumination is now a mainstream technology for mobile imaging. Wafer stacking of the image array on a CMOS image signal processor will become common.

The importance of digital signal processing technology in cameras continues to grow in order 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.

High Dynamic Range (HDR) is now well established in low-cost consumer imaging. While HDR combines multiple distinct images, new work is progressing with specialized architectures to extend the dynamic range in single exposures, and thus avoid movement artifacts. As well, global shutters are being introduced to avoid movement artifacts.

For precision scientific and medical applications, we now employ single-photon avalanche diode (SPAD) imager arrays. Deep-submicron CMOS SPADs will meet the requirements for high resolution, high accuracy time-to-digital converters (TDCs), and small-pitch imagers with better fill factor.

The market share of CCD imagers continues to shrink and they are now relegated to minor applications. Top-end broadcast and top-end digital cameras have moved from CCD to CMOS sensors.

Sensors & MEMS:

MEMS has now enabled the world's smallest 32kHz ultra-low-power timing sources. Low-power timing has normally been supplied by quartz tuning forks, but miniaturization of that technology is proceeding slowly. MEMS oscillators are available in 1.5×0.8mm^2 chip-scale packages, and consume under 1µA supply current. Temperature compensation provides 3 parts per million accuracy over temperature.

On another front, gesture recognition for computers, tablets, and phones will become very important. This can be implemented using capacitive, image, or ultrasonic sensors, where the latter has a potential power advantage. MEMS transducers will be a key component of such ultrasonic systems, employing specialized driver chips. This technology has now been demonstrated, and can track multiple objects with high resolution.

CMOS temperature sensors continue to improve, with new circuit techniques increasing accuracy and supporting wider applications. Strides in low-power architectures continue to eclipse previous results with record-setting efficiency which support battery-operated and mobile applications. For applications in fine-geometry systems, we are seeing DTMOS-based sensors becoming more accurate, particularly at low supply voltages.

i

Medical:

Medical applications continue to emphasize smaller implantable and wearable devices. While often these must operate at low power levels using miniature battery or energy-scavenging , the required data transmission rates can be large, for instance for neural-recording arrays. To minimize power consumption, and thus minimize size, such devices will require local signal pre-processing and feature extraction.

For improved capability and quality, medical ultrasound is developing toward 3D 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.

Smart wearable sensors for medical applications will support continuous remote monitoring with data transmission to centralized analysis systems. These sensors will adapt their algorithms to match specific user's vital signs. Potentially, closed-loop control will allow therapy to be applied directly, for instance to suppress seizures or arrhythmias.

Displays:

Touch- and gesture-sensing systems for large displays are in development. These will require extensions of capacitive-sensing technology beyond the present state-of-the --art. Applications will require sensing hand gestures at a distance (30cm), and high-resolution touch for writing recognition.

LCD panels with integrated touch sensing are being driven toward thinner and lower-cost single-chip solutions. Robust circuit technology that is immune to display-driver noise is a key design factor.

Higher-resolution and higher-definition displays are being developed for mobile applications. Displays with 440 pixels-per-inch (ppi) are now in production, with new work focused beyond 500 ppi for 6-inch and smaller displays. While low-temperature polysilicon (LTPS) technology seems to be superior to amorphous-silicon (a-Si), traditional a-Si TFT and oxide TFT technologies supported by compensating driver systems are being prepared to compete.

Plastic organic light-emitting diodes (OLEDs) are entering into commercial production for curved smartphones. These will give improved user experiences, being very thin, light, and un-breakable. They are also expected to find applications in wearable displays.

Technology Directions – 2014 Trends

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

The Need for Innovative Systems:

This year, at ISSCC 2014 the theme is, "*Silicon Systems Bridging the Cloud*". System considerations are extremely important for circuit engineers. The system designer defines the parameters such as performance, power, bandwidth, SNR, QoS, etc. to which engineers design their hardware. In addition, the system designer can help the hardware engineer optimize across the entire system. Thus, it is apparent that the next level of innovative technology will come out of a broad understanding of how all parts of the system work together so that engineers can optimize across multiple layers. Correspondingly, we are highlighting two diverse systems which will drive future innovation: Biomedical Systems for Improved Quality of Life, and High-speed Systems for Data Networks.

Innovative Biomedical Systems:

A breakthrough concept in biomedical electronics is the development of "anytime and anywhere" human monitoring systems using wearable/implantable devices. Such devices will enable improved of quality of life, through self-health checks, remote examination by doctors, and constant monitoring for acute diseases.

Key technologies such system include:

(1) Small-footprint devices and flexible electronics to enhance the comfort of wearable devices.
(2) High-accuracy monitoring devices.
(3) Low-power monitoring and communication systems for long-duration autonomous operation.

Correspondingly, highly integrated biomedical SoCs capable of sensor detection, diagnosis and wireless communication have been developed to ensure a small footprint. To this end, development of 3D-integrated electrodes and sensors with MEMS, and new-materials such as organic devices, has also contributed. In order to monitor weak vital-signs (such as electrocardiogram (ECG) and electroencephalogram (EEG)) with a high degree of accuracy, there has beenclear progress in the variation-correction methodologies needed to address environmental change and the misalignment of wearable devices. Furthermore, state-of-the-art biomedical SoCs increasingly embody multimodal bio-monitoring systems for diverse signal detection. To operate reliably over long periods, such SoCs have evolved to contain a power-efficient dedicated processor running algorithms for diagnosis, along with ultra-low-power circuit designs for body-area networking, and energy-harvesting technologies. Illustrating these trends, Session 18 at ISSCC 2014 features the latest in biomedical systems.

Innovative High-speed Data Networks:

In the past few decades, many new applications such as mobile devices and the Internet have driven the growth of high-performance computing systems. The next big application drivers will include cloud computing, Big-Data, and the Internet-of-Things, which will require increasing performance demands on the backbone of the infrastructure. Important trends and challenges of the future that necessitate system innovation include:

(1) Satisfying performance demands that increase at an exponential rate.
(2) Containing the power dissipated in data centers whose electricity and cooling costs are skyrocketing.

To meet these challenges, innovative architectures for the infrastructure (such as servers, network routers, andmobiles) are under development to support increased performance demands. For example, minimizing latency of memory access is critical to network processing performance. For this purpose, novel network-on-chip (NoC) architectures that optimize memory bandwidth and implement advanced memory access protocols/algorithms are needed. Furthermore, low-power processor architectures are needed to contain the power dissipated by high performance systems. As protocols and software can make significant impact, these need to be co-optimized with the architecture.

New circuit technology is also needed to reach the power and performance requirements of data networks. For example, advanced clocking techniques are needed within the required GHz ICs. As well, power-management techniques traditionally used for energy limited applications are being employed within high-performance systems. Furthermore, new packaging techniques, such as 3D interconnects with homogeneous and heterogeneous usage are being investigated for high-performance systems. Session 6 at ISSCC 2014 features innovative high-performance data network systems where challenging system requirements spur the development of novel solutions.