

A Survey on Low-Power Techniques with Emerging Technologies: From Devices to Systems

PIERRE-EMMANUEL GAILLARDON, EPFL

EDITH BEIGNE and SUZANNE LESECQ, CEA-LETI, Minatec Campus

GIOVANNI DE MICHELI, EPFL

Nowadays, power consumption is one of the main limitations of electronic systems. In this context, novel and emerging devices provide new opportunities to extend the trend toward low-power design. In this survey article, we present a transversal survey on energy-efficient techniques ranging from devices to architectures. The actual trends of device research, with fully depleted planar devices, tri-gate geometries, and gate-all-around structures, allows us to reach an increasingly higher level of performance while reducing the associated power. In addition, beyond the simple device property enhancements, emerging devices also lead to innovations at the circuit and architectural levels. In particular, devices whose properties can be tuned through additional terminals enable a fine and dynamic control of device threshold. They also enable designers to realize logic gates and to implement power-related techniques in a compact way unreachable to standard technologies. These innovations reduce power consumption at the gate level and unlock new means of actuation in architectural solutions like adaptive voltage and frequency scaling.

Categories and Subject Descriptors: B.7.1 [Integrated Circuits]: Types and Design Styles—VLSI (*very large scale integration*)

General Terms: Design, Performance

Additional Key Words and Phrases: Low-power techniques, UTBB FDSOI, vertically stacked nanowires, arithmetic logic, power gating, DVFS, AVFS

ACM Reference Format:

Pierre-Emmanuel Gaillardon, Edith Beigne, Suzanne Lesecq, and Giovanni De Micheli. 2015. A survey on low-power techniques with emerging technologies: From devices to systems. *ACM J. Emerg. Technol. Comput. Syst.* 12, 2, Article 12 (August 2015), 26 pages.

DOI: <http://dx.doi.org/10.1145/2714566>

1. INTRODUCTION

Power consumption has become the most important question in a wide range of electronic systems. As demand towards an increasingly higher level of integration and performance continues to grow, the semiconductor industry manufactures devices with dimensions of few tens of nanometers. However, while the reduction of device dimensions increases the computing density, that is, the maximally possible number of computations per unit area and time, it also increases leakage power much faster than active power, assuming identical applications and implementations [Kyung and Yoo 2011].

Recently, power consumption has become limited by many constraints. First, the world's *information-communications-technologies* (ICT) ecosystem is nowadays

This work is supported by the European Research Council under grant ERC senior NANOSYS ERC-2009-AdG-256810.

Authors' addresses: P.-E. Gaillardon (corresponding author), Integrated Systems Laboratory, Swiss Federal Institute of Technology, 1015 Lausanne, Switzerland; email: pierre-emmanuel.gaillardon@epfl.ch; E. Beigne and S. Lesecq, CEA-LETI, Minatec Campus, 38054 Grenoble, France; G. De Micheli, Integrated Systems Laboratory, Swiss Federal Institute of Technology, 1015 Lausanne, Switzerland.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2015 ACM 1550-4832/2015/08-ART12 \$15.00

DOI: <http://dx.doi.org/10.1145/2714566>

consuming 10% of the world electricity generation [Mills 2013]. Second, while considering mobile market, the workload of cell phones increases while keeping a fixed battery capacity. Third, with the transition towards more sustainable green computing, the environmental impact has to be considered in order to reduce the CO₂ emissions generated by IT systems.

To reduce the impact of power while boosting the level of performance, technological developments introduced many innovations during the past decade. At the advanced technology nodes, *fully depleted silicon-on-insulator* (FDSOI) *field effect transistors* (FETs) push the limits of planar transistor geometries and also introduce novel means of leakage control through an additional biasing given by a back gate terminal [Beigné et al. 2013]. In addition, *fin-based FETs* (FinFETs) provide an alternative to planar devices in order to build higher-performance low-power SoCs [Auth et al. 2012; Jan et al. 2012]. To push device performance even further, vertically stacked *silicon nanowire FETs* (SiNWFETs) with gate-all-around control are considered the natural extension of FinFETs [Bangsaruntip et al. 2009] and lower the leakage floor of the device. In addition to their good electrostatic control as compared to FinFETs, *multiple-independent-gate* (MIG) SiNWFETs are also promising, thanks to their enhanced set of functionality [De Marchi et al. 2012; Zhang et al. 2013] that translates to novel opportunities at the circuit level, pushing forward the power reduction. Hence, emerging technologies partially answer the problem of power-related issues and therefore require a holistic approach. Energy-aware design is the design of a system to meet a given performance constraint with the minimum energy consumption. Low-power design can be achieved at every design level from device to architectural level.

In this survey article, we review different opportunities brought by emerging technologies' low-power digital designs. Emerging devices bring a higher level of performance for a reduced power impact, thanks to structural innovations and materials. Furthermore, they also bring an enhanced set of functionality, such as embedded threshold control or higher logic expressiveness. New functionalities lead to a novel offer at the design level. In particular, we review techniques that exploit novel device properties to create simple arithmetic logic gates and embedded power gating techniques, thereby realizing datapath structures in a more power-efficient way. Finally, we also consider low-power techniques at the architectural level, such as dynamic voltage and frequency scaling, and we discuss their applications in light of novel technologies.

The remainder of the article is organized as follows. In Section 2, we give the necessary background about power consumption in digital circuits. In Section 3, we survey recent innovations introduced at the device level. In Section 4, we discuss the impact at circuit level of emerging devices with enhanced functionalities. In Section 5, we review some common architectural-level techniques for power reduction and we draw some architectural perspectives. In Section 6, we conclude.

2. NECESSARY BACKGROUND

In this section we first generally discuss the nature of power consumption in digital circuits. For a complete review on power expressions in electronic circuits, we refer the interested reader to Kyung and Yoo [2011]. The power consumption of a chip can be decomposed according to both active and standby phases. Active phase refers to that period of time when the circuit is normally operating and produces a meaningful output. The standby phase corresponds to the nonactive phase. During the active phase, we identify two power contributions: the dynamic power (P_{dyn}) and the static power (P_{stat}):

$$P = P_{dyn} + P_{stat}.$$

The dynamic power is consumed only when a transistor is switching, that is, when a gate is switching from 0 to 1 or from 1 to 0, and charges or discharges the load

capacitance. Static power corresponds to the consumption of the transistor during the remaining time.

2.1. Dynamic Power

Dynamic power can be divided in two components: switching power (P_{SW}) and short-circuit power (P_{SC}):

$$P_{dyn} = P_{SW} + P_{SC}.$$

The switching power corresponds to the amount of power delivered from the supply voltage to charge and discharge the load capacitance when the transistors switch. It is modeled by the following relation:

$$P_{SW} = \alpha \cdot F \cdot C_L \cdot V_{dd}^2,$$

where C_L is the load capacitance which models the gate capacitances of the fanout logic gates, the wire capacitance and the intrinsic output capacitance of the gate itself; F is the clock frequency; V_{dd} is the power supply; and α is the activity factor, that is, the probability of the output making a pair of rising and falling transition during a single clock cycle.

Short-circuit power is caused by the short-circuit current that flows during a transition, that is, when both the pull-up and pull-down networks of a gate are turned on for a short transient period of time during switching. In practice, P_{SC} is a small proportion of the total dynamic power [Nose and Sakurai 2000]

$$P_{SC} \approx 10\% \cdot P_{dyn}$$

and will therefore be neglected in the following.

2.2. Static Power

The static power is a result of device leakage current which originates from various phenomena [Roy et al. 2003], such as subthreshold conduction, gate direct tunneling current, junction tunneling leakage, *gate induced drain leakage* (GIDL), hot carrier injection current, or punchthrough current. However, since the introduction of high- κ dielectrics that reduced the gate direct tunneling, the main contribution is the subthreshold leakage current. This leakage occurs when gate-to-source voltage of a transistor is below the threshold voltage, that is, when the transistor is supposed to be turned off. It increases exponentially with decreasing threshold voltage and increasing temperature [Kim et al. 2003]. We can therefore express static power using the relation:

$$P_{stat} \propto \beta \cdot V_{dd} \cdot e^{-\frac{V_{th}}{\gamma \cdot V_T}},$$

where β and γ are experimentally derived constants, V_{dd} the power supply, V_{th} the threshold voltage, and V_T the Boltzmann thermal voltage that depends linearly on the temperature.

Note that, in the active phase of the circuit, the static power contribution has also a transient contribution that corresponds to the charge of the internal gate capacitances. Indeed, even if it does not involve any transistor switching, this transient period is different than the one obtained during the standby phase where all parasitic charges are stable.

2.3. Motivation

In the following, we consider principally the switching power as a main contributor of dynamic power and the subthreshold leakage power for the static power in both standby and active modes. To efficiently manage power, the objective is to reduce the

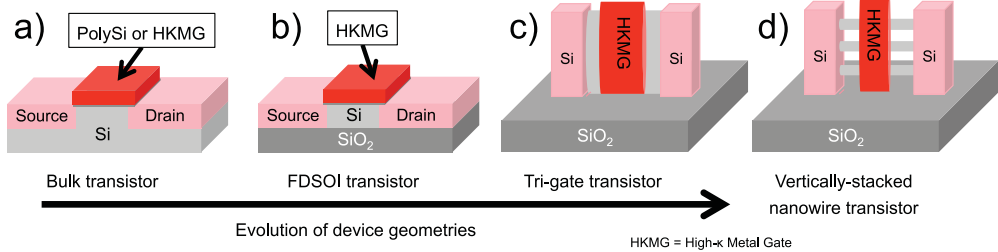


Fig. 1. Electrostatic control improvements by structural innovations: (a) Bulk transistor; (b) FDSOI planar transistor; (c) tri-gate 3D transistor; (d) vertically stacked nanowire transistor.

total switching power (C_L) but also to dynamically control the circuit voltage supply (V_{dd}), clock frequency (F), and threshold voltage (V_{th}). Therefore, we review the different techniques employed to control these two power contributors from three complementary angles: devices, circuits, and architectures.

3. LOW-POWER HIGH-PERFORMANCES DEVICES: TOWARDS THIN DEVICES

In recent years, a large range of device innovations has been introduced to push performance while carefully controlling the power budget. In particular, the semiconductor industry moved from basic planar bulk silicon *field effect transistors* (FETs) to advanced structures such as *fully depleted silicon-on-insulator FETs* (FDSOI) and *fin-based FETs* (FinFETs), as depicted in Figure 1, thereby improving the intrinsic device properties. In this section, we review innovations brought by the device community to improve the energy efficiency of the fundamental switches.

3.1.1. Fully Depleted Thin Planar Devices. In the era of classical scaling, transistor performance primarily improved as a result of dimensional scaling up to the use of fully depleted SOI, where the conducting silicon channel is reduced to a thin layer of intrinsic silicon. In the past decade, performance has progressed through the introduction of transistor architecture innovations, including strained silicon [Ghani et al. 2003] and high- κ metal-gate technologies [Mistry et al. 2007]. Keeping the pace towards more device control, current innovations have now introduced electrostatic means of tuning the device properties.

Ultra-Thin Body and Box (UTBB). FDSOI is an advanced planar technology [Liu et al. 2010]. Figure 2 depicts a cross-section of a UTBB FDSOI device and highlights the main technological features. UTBB FDSOI takes profit of typical advanced planar techniques such as high- κ metal gate and reduced-resistance raised source and drain access. The channel is intrinsic, that is, without global bulk doping or pocket implants. *Shallow trench isolation* (STI) is used to electrically isolate the devices.

Avoiding doping steps makes the process simpler, cost efficient, and less prone to variability than bulk. In standard fully depleted SOI technology, the silicon film (where the active area of the transistor is located) is thinned down to approximately one-third of the minimum gate length value [Mazuré et al. 2010]. Such thin channel enables a good electrostatic control with, for instance, a low *drain induced barrier lowering* (DIBL) value [Khakifirooz et al. 2012] by cutting deep field lines originating from the drain. In a 28nm UTBB FDSOI node, the silicon film thickness T_{SI} is approximately 8–9nm [Liu et al. 2010].

The *buried oxide* (BOX) is thinned down to 25nm, leading to a good trade-off between drain/source-to-substrate parasitic capacitance and body factor, that is, the influence of the back plane on the device. Indeed, a thin BOX enables a back-plane effect. A back plane, either *n* type or *p* type, is implemented underneath the BOX and allows, by




Fig. 2. Cross-section of a UTBB-FDSOI transistor.

electrostatic control over the channel called *back bias* (BB), to improve the *short channel effect* (SCE). Note that a hybrid FDSOI/bulk technology is feasible after removing the BOX. The BOX creates a back interface which can be thought of as a back gate and whose polarization value modifies the *threshold voltage* (V_{th}) of the transistors. The *body factor*, representing the V_{th} sensitivity to the back-gate voltage, is equal to 85mV/V for a 25nm BOX. The circuit-level innovations brought by UTBB FDSOI multi- V_{th} will be detailed in the following section.

From a variability perspective, FDSOI transistors have undoped channels. Therefore, this removes the major source of variability in current deep-submicron technology, that is, the *random dopant fluctuation* [Li et al. 2010], which positively impacts the variability records. Note that the back plane is doped but has no impact on variability since the back plane only plays a support role.

3.1.2. Nonplanar 3D Devices and Improved Channel Control. Subsequent to the reduction of semiconducting film thickness, control of the channel using multigates, that is, gate regions controlling different parts of the channel, was expected to lead to improved performance at lower supply voltage and significantly reduced short channel effects [Choi et al. 2001]. More precisely, controlling a fin-shaped channel on three sides hence appears a reliable way to provide scaled devices with high-performance and low-power capabilities [Auth et al. 2012; Jan et al. 2012]. The introduction of the tri-gate device marked the end of the planar device era. A typical tri-gate FinFET device structure is depicted in Figure 3. The conducting channel is formed by a thin silicon “fin” connected to larger source/drain regions acting as supporting pillars. The gate is wrapped over the channel and controls the three-side of the fin. In addition to the effects of standard device engineering, such as strained silicon and high- κ metal gate, the improved gate control of the tri-gate structure can be seen in the steep subthreshold, typically around $\sim 70\text{mV/dec}$, the very low *drain induced barrier lowering* (DIBL), around 50mV/V obtained for minimum gate length devices and the reduced leakage floor to 10pA μm [Jan et al. 2012]. For *p*-type devices, nonpure silicon source/drain materials such as SiGe are used to induce compressive stress in the channel and to get similar properties as *n*-type devices [Auth et al. 2012].

Keeping the pace towards reduction of transistor dimensionality, *nanowires* (NWs) are expected to further push the scalability of electronic devices. Indeed, *silicon nanowires* (SiNWs) are considered a very promising nanodevice technology for low-power systems because of their ultimate mono-dimensional semiconductor properties combined with the vast experience and investment in Si technologies. Fabrication technologies for SiNW have been the object of recent investigation through bottom-up




Fig. 3. Tri-gate fin-based transistor conceptual sketch.




Fig. 4. Vertically stacked nanowires transistor conceptual sketch.

[Xiang et al. 2006; Yu and Lieber 2010; Heinzig et al. 2011; Yan et al. 2011] and top-down [Dupré et al. 2008; Bera et al. 2006] approaches.

Top-down, that is, lithography-based, SiNW technologies are credible for *very large-scale integration* (VLSI) circuits because of their accurate control of device size and geometry [Bangsaruntip et al. 2009; Fang et al. 2007] as well as of the proven capability of Si technology to scale up to billions of devices per chip. Figure 6 depicts a vertically stacked nanowire structure. In this structure, the channel is formed by a collection of nanowires that provide a high on-current and steep subthreshold slope (around 64mV/dec [De Marchi et al. 2012; Ernst 2013]) with very low leakage properties. The control gate is realized in a *gate-all-around* (GAA) fashion and ensures perfect electrostatic control over the channels.

Vertically stacked GAA SiNWs exploit the vertical dimension in a way reminiscent of tri-gate transistors [Auth et al. 2012] and represent a natural evolution of these structures (Figure 4). They provide better geometry for electrostatic control over the channel and, consequently, superior scalability properties [Dupré et al. 2008; Bangsaruntip et al. 2009].

The pace towards improved electrostatic properties enables the introduction of novel conduction phenomena that can help further push the performance of the doped source drain transistors beyond the usual limits. Indeed, there recently have been many efforts initiated to study *tunnel FETs* (TFETs). The tunnel FET has been identified as a promising candidate for low-power electronics as it can overcome the MOS limitations coming from the non-scalability of the inverse *subthreshold slope* (SS) [Chang et al. 2010].




Fig. 5. Vertically stacked nanowire-based tunnel FET conceptual sketch.




Fig. 6. Three-independent-gate vertically stacked nanowire transistor conceptual sketch.

As demonstrated in Moselund et al. [2011], nanowires are ideal to achieve a good tunnel FET as it requires the highest degree of electrostatic control. In particular, the key parameters to achieve good performance are the tunnel junction abruptness, the effective band gap at the tunneling junction, gate control over the channel, and overall device geometry. A conceptual sketch of a tunnel FET realization over a vertically stacked nanowire structure is depicted in Figure 5.

3.2. Functionality-Enhanced Devices: An Alternative to Moore’s Law

During the last four decades, improvements on the device level have been considerable and successful in following the pace of Moore’s Law. However, with the recent deceleration of pure scaling, new approaches must be explored.

Recently highlighted by many researchers as an alternative approach to the initial Moore’s Law, the question of scaling should be considered in a more generalized way than a simple reduction of dimensions. In particular, it is worth identifying new devices “with scaling understood in the most generic sense of increasing computational performance (function) per unit area” [Bernstein et al. 2010]. Here, we focus on the opportunities given by *multiple-independent-gate* (MIG) devices to enrich the functionalities of elementary switches.

Multi-independent-gate (MIG) devices are transistors whose electrostatic properties are dynamically controlled via additional gate terminals. MIG devices have been successfully fabricated using carbon nanotube [Lin et al. 2005], graphene [Harada et al. 2010], and *silicon nanowire* (SiNW) [Heinzig et al. 2011; De Marchi et al. 2012] technologies. As a form of natural evolution of the FinFET structure, vertically stacked




Fig. 7. Conceptual structure of a vertically stacked TIG SiNWFET and associated symbols.

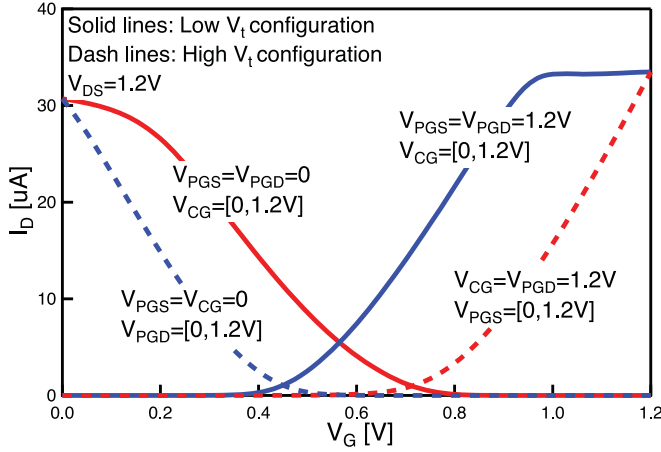


Fig. 8. Simulated triple-independent-gate vertically stacked nanowire transistor characteristic.

SiNWs are a promising platform for MIG-controllable polarity devices thanks to their high I_{on}/I_{off} ratio and CMOS-compatible fabrication process [De Marchi et al. 2012].

Among this family, we can emphasize the *double-independent-gate* (DIG) [De Marchi et al. 2012] and *three-independent-gate* (TIG) [Zhang et al. 2013] SiNWFETs. DIG and TIG FETs consist of three vertically stacked SiNWs with three separated gated regions. The device structure is depicted in Figure 6.

The side regions are called the *polarity gates* (PG) while the central region is tied to the *control gate* (CG). The CG enables the flow of carriers as in conventional transistors. The *polarity gate at source* (PG_s) and *polarity gate at drain* (PG_D) tune the Schottky barriers at the source and drain junctions, respectively. When tied together, we obtain a unique PG terminal, and therefore a DIG device that controls the channel carrier's type ($V_{PG} = V_{DD} \rightarrow n$ -type, $V_{PG} = V_{SS} \rightarrow p$ -type). Indeed, at high PG bias, the channel energy is lowered and the depletion regions at Schottky junctions are thinned compared to the conduction band, allowing electrons to flow through the device. On the contrary, when low PG bias is applied, the channel energy is increased and thinning of the depletion regions at Schottky junctions is observed with respect to the valence band, enabling the flow of holes in the SiNW. When separately controlled, that is, within a TIG FET, the set of possible functionalities is larger with an additional control of the device threshold voltage (V_{th}) and the ability to realize two transistors in a unique device [Zhang et al. 2013]. The symbols associated with different configurations are depicted in Figure 7.

Figure 8 shows the I-V characteristic for a 22nm device simulated using a TCAD model [Zhang et al. 2013]. In addition to the dynamic polarity configuration, we observe the dual- V_{th} characteristic of TIG with a threshold difference of about 0.3V. In conventional dual- V_{th} technology, high- V_{th} devices achieve lower leakage current but also reduce I_{ON} compared to low- V_{th} devices. However, with TIG SiNWFETs, the I_{ON} of

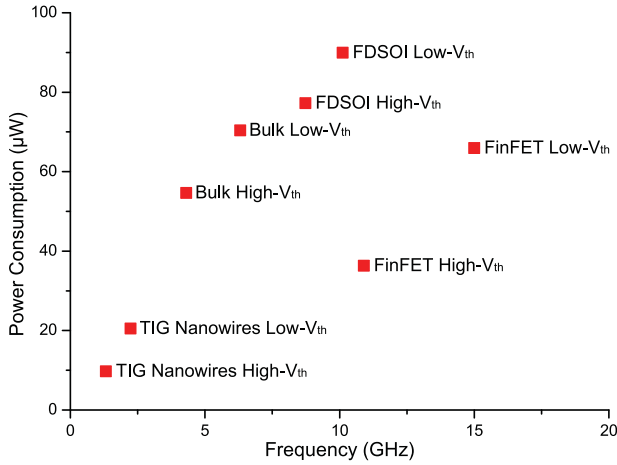


Fig. 9. Simulated power consumption/frequency of an 11-stage ring oscillator implemented with 28nm bulk, 28nm FDSOI, 22nm FinFET, and 22nm TIG NWFETs in high- V_{th} and low- V_{th} options.

both high- and low- V_{th} options are the same, and a leakage floor of less than $5\text{pA}/\mu\text{m}$ is reached. Note that DIG behavior can be derived from conditions where $PG_S = PG_D$. We refer the interested reader to De Marchi et al. [2012] and Zhang et al. [2013] for more details about the physics of DIG and TIG SiNWFETs, respectively.

3.3. Discussions

Device evolution allows the semiconductor industry to keep the pace towards higher performance, less short channel effects, and ultimately a reduced leakage floor. In order to better assess the different performance/power capabilities of the different reviewed technologies, Figure 9 shows the frequency of operation along with the power consumption of an 11-stage ring oscillator realized with 28nm bulk, 28nm FDSOI, 22nm FinFETs, and 22nm polarity-controllable nanowire FETs in both high- and low- V_{th} options. Exploiting improved electrostatic control over the channel structure, these devices are capable of either increasing the performance level, as compared to bulk technologies, while keeping the power consumption under control or drastically reducing the power budget. Device-level optimizations, that is, high- V_{th} or low- V_{th} options, can be done for the different technologies discussed earlier in order to trade off between speed and power consumption.

In addition to technology boosters, we recently observed growing interest in devices whose properties can be fine-tuned through external electrostatic control such as UTBB FDSOI or MIG FETs. In the evaluation of Figure 9, this extra control flexibility is not leveraged. In the next section, we will see how these tuning knobs can be exploited at the circuit level.

From an integration perspective, it is worth pointing out that compatibility of emerging technologies with baseline planar bulk technology is of tremendous importance. Indeed, the semiconductor industry is not able to afford a massive paradigm change in fabrication techniques due to the large investments done on bulk CMOS over past decades. The different technologies described in this article rely on standard integration schemes and are fully compatible with standard nano-fabrication techniques. In particular, while FDSOI and FinFETs are already in industrial production, advanced nanowire-based devices exploit the same geometries and process techniques, making them natively compatible with large-scale integration. Furthermore, their technology compatibility makes possible heterogeneous co-integration of different technologies

within the same chip, such as nonvolatile flash gate stacks with standard digital transistors [Xuan et al. 2003].

4. CIRCUIT-LEVEL OPPORTUNITIES FOR LOW-POWER SYSTEMS

Innovations brought by advanced and emerging devices greatly contribute to the management of power issues in advanced system design. They also introduce new opportunities that can be exploited at the circuit level. In this section, we first discuss multi- V_{th} capabilities given by the additional electrostatic control in UTBB FDSOI and MIG FETs. Then, we explore the ability of MIG FETs to create simpler and higher-performing logic gates and also to implement low-cost embedded power gating techniques.

4.1. Ultra-Wide-Voltage-Range Design through Additional Electrostatic Control

In addition to traditional multi- V_{th} tuning as already used in bulk technology, UTBB FDSOI enables a fine V_{th} adjustment thanks to its back gate. The V_{th} of UTBB FDSOI devices can be described by Beigné et al. [2013]:

$$V_{th} \propto V_{F0} + r \cdot (V_{B0} - V_B)$$

with V_{F0} and V_{B0} the fixed contributions to V_{th} coming from the front face and back face, respectively, r the body factor; and V_B the back biasing voltage. The gate stack and the Si-film counter doping (front face) affect V_{F0} , as in regular devices. However, UTBB FDSOI can be further tuned with a coarser-grain adjustment of the back face, defined by V_{B0} , and fine-grain electrostatic control through the back biasing voltage. V_{B0} mainly depends on the back-plane work function (typically from 4.1eV to 5.1eV). The importance of the back-face effect depends on the body factor, which can be described as the capacitance ratio between back and front faces [Noel et al. 2011]. The body factor for both n - and p -type achieves more than 85mV/V [Planes et al. 2012], which is equivalent to bulk technology. However, the body factor in UTBB FDSOI technology does not degrade with the scaling and can be maintained or increased by reduction of BOX thickness at each technology node. Given that V_B can vary over a 2V amplitude, the V_{th} value can be greatly modified dynamically, enabling circuit designers to define several operating points. The large range of operation brought by this added flexibility enables *ultra-wide-voltage-range* (UWVR) operation. Effectively, UTBB FDSOI is able to span a large set of operating points ranging from high performance to low power [Beigné et al. 2013], thereby bringing a new offer for dynamic tuning of performance. A dynamic tuning of performance is of interest for a large range of design targets, ranging from logic gates [Noel et al. 2011] to embedded memories [Thomas et al. 2012]. Such techniques were recently employed in a low-power *digital signal processing* (DSP) design [Wilson et al. 2014], demonstrating the large interest in the approach.

4.2. Towards Novel, Fast and Power-Efficient Implementation of Logic Functions

Traditional unipolar FETs can implement an inverter with a single loaded device. However, with novel MIG-SiNWFETs, a 2-input XOR function is realized in a single device [Ben Jamaa et al. 2009], enabling compact implementation opportunities for arithmetic- and XOR-intensive logic. In the rest of this section, we discuss the design of complex arithmetic logic with controllable-polarity transistors.

4.2.1. XOR-Based Logic. As depicted by the I/V curves in Figure 8, the intrinsic ability of MIG devices to see their polarity reconfigured embeds a bi-conditional behavior. Table I summarizes the XNOR conduction states of a single device to illustrate this property.

The opportunity to have compact XOR-based logic gates with controllable polarity transistors was first exploited in Ben Jamaa et al. [2009]. In particular, the 2-input XOR

Table I. Abstraction of the Conduction Mode of a DIG Device (i.e., when $P_{G_S} = P_{G_D} = P_G$)

V_{PG}	V_{CG}	Transistor State
0	0	On (p-type)
0	1	Off (p-type)
1	0	Off (n-type)
1	1	On (n-type)




Fig. 10. (a) Static 2-input XOR gate; (b) transmission gate 3-input XOR gate.




Fig. 11. (a) MUX-XNOR generalized arithmetic structure; (b) 3-input majority operator.

gate from Ben Jamaa et al. [2009] is reported in Figure 10(a). Note that the equivalent CMOS realization employs $2 \times$ more devices [Rabaey et al. 2003]. The transmission gate configuration in Figure 10(a) enables a full-voltage swing path between signal output and the power rails while embedding the XNOR logical connective. Extending the logic design style from static to pass-transistor, a 3-input XOR realization as introduced in Zukovski et al. [2011] is obtained and depicted by Figure 10(b). These two implementations, in addition to their compactness, show improved power capabilities. Indeed, less transistors are switching during a transition, thereby reducing the dynamic power consumption.

4.2.2. MAJ/MUX-Based Logic. Devices with controllable polarity enable not only efficient XOR-intensive logic but also compact logic gates based on the majority voting operation. The 4-controllable-polarity FETs configuration used in previous logic gates is generalized in the MUX-like structure, as depicted by Figure 11(a). Its functionality corresponds to a multiplexer driven by a XNOR signal, namely $A \oplus B$, selecting between two external signals G and F .

In particular, one can exploit this property to extend the range of realized functions. In Turkyilmaz et al. [2013], a 4-transistor 3-input majority logic gate is proposed and reported here in Figure 11(b). Note that in static CMOS, the same gate has 10 devices




Fig. 12. Full adder with 8 controllability devices.




Fig. 13. (a) Double-device functionality merge in TIG SiNWFETs; (b) AOI gate implementation.

in place of 4 [Rabaey et al. 2003]. With different assignments of G and F , it is possible to reproduce 3-input MAJ ($G = C, F = A$), 3-input XOR ($G = C, F = C$), and 2-input XOR ($G = 1, F = 0$) logic gates. Therefore, the MUX-XNOR gate can be seen as a generalized arithmetic gate.

4.2.3. Full Adder. The full adder is a widely used arithmetic circuit that supports the addition of two binary numbers. It is represented by the following 3-input 2-output logic function: $\text{Sum} = A \oplus B \oplus C$ and $C_{\text{OUT}} = \text{MAJ}(A, B, C)$. Controllable-polarity transistors offer an advantageous implementation for both the Sum and C_{OUT} functions, therefore the full adder is competitively realized with 8 devices, input inverters apart, as depicted by Figure 12. The corresponding static (transmission gate) CMOS version has 28 (14) transistors [Rabaey et al. 2003].

4.2.4. Multi-Transistor Merging. Extending the approach to TIG FETs, we note that it is possible to realize two standard unipolar transistors in a unique device. Indeed, by combining the low- V_{th} and high- V_{th} n -FET configurations, applying logic signals on CG and P_G is equivalent to 2 series n -FETs. Similarly, the configuration of 2 series p -FETs is obtained by combining the low- V_{th} and high- V_{th} p -FET configurations. These configurations are illustrated in Figure 13(a).

Even though a specified gate is used for polarization in TIG SiNWFETs, these 2-input configurations efficiently utilize the extra gates to merge two devices in a unique one, thereby mitigating the area overhead compared to conventional CMOS devices.

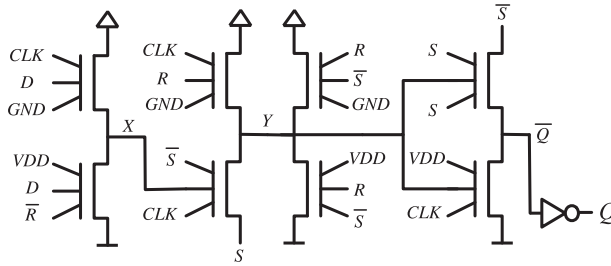


Fig. 14. TSPC flip-flop design using TIG-FETs compactness.

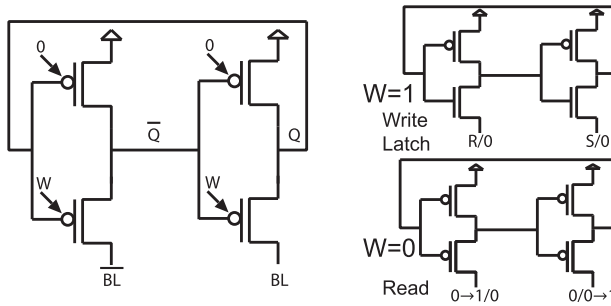


Fig. 15. SRAM circuit structure exploiting DIG FETs: (left) General structure; (right) read and write modes.

In addition, the internal node capacitance between two inputs does not exist in TIG SiNWFETs. This helps to reduce the delay of circuits. Figure 13(b) presents an example of an AOI gate. Its functionality is obtained by applying the configurations of 2 series p -FETs/ n -FETs and a low- V_{th} n -FET. Thus, only 4 transistors are needed for this AOI gate, instead of 6 conventional MOSFETs [Zhang et al. 2013].

4.2.5. Memory Opportunities. As already stated, DIG/TIG FETs enable two major circuit-level improvements: (i) a compact realization of XOR functions and (ii) the merge of two serial transistors in a single device. These two properties can be efficiently leveraged in memory design. First, we report on a *true single-phase clock* (TSPC) design implementation using TIGs [Yuan and Svensson 1997].

Figure 14 shows an FF design build with only 8 transistors as compared to 15 in its traditional CMOS counterpart. By reducing the number of transistors stacked in pull-up and pull-down networks and by using the larger functionality set offered by the controllable polarity, it has been shown in Tang et al. [2014] that the proposed design leads to datapath storage elements with, on average, area and delay savings of 20% and 43%, respectively, compared again to FinFET LSTP transistors at 22nm technological node.

Second, we also highlight the opportunities given for memory array design. In particular, we introduce a *static random access memory* (SRAM), depicted in Figure 15 (left), which consists of four transistors realizing two cross-coupled inverters with special properties [Gaillardon et al. 2014]. First, the bottom transistors are not standard FETs but DIG FETs, where one gate is still connected as in usual inverters while the other provides enhanced controllability. Second, the bottom terminals of the cross-coupled inverters are not grounded, but connected to *bitlines* (BLs). By exploiting the controllability of the bottom FETs, it is possible to let the BLs write/read the cell by directly forcing/sensing the logic value at the output nodes of the cross-coupled inverter.

The proposed cells have 2 operation modes as highlighted in Figure 15 (right). The signal W (write) controls the polarity of the bottom FETs and thus imposes the operation mode. When $W = 1$, the memory cell is in writing or static latch mode. Indeed, if both BLs are grounded, then the cell behaves as a static latch as depicted by Figure 15 (right). When instead the BLs assume nonidentical value in this specific operation mode, the internal nodes of the cross-coupled inverter are forced to assume such values, thereby operating in a similar way as an SR latch. Note that, after a short period (typically few tens of ps), the memory cells naturally stabilize to the written values. When $W = 0$, the cell is in reading mode (Figure 15 (right)). The BLs are initially discharged to ground. Subsequently, the bottom FETs charge the BLs to values stored in the internal nodes. Similarly as in standard SRAMs, the reading process can be speed up by using sense amplifiers and related circuitry. As compared to a standard SRAM cell in 22nm FinFET technology, the proposed cell is 14% smaller and 16% faster.

4.2.6. Discussions. The functionality enhancement of TIG/DIG devices opens new ways of implementing standard logic gates that translates to circuit-level improvements. From a first-order analysis, TIG FETs are, in general, larger (three gate regions) and present higher parasitic capacitances than a unipolar device at the same technology node. Hence, circuits implemented with such transistors would have worsened if no circuit-level opportunities had been applied. In fact, the logic gate compactness compensates for the lack of device performance, thereby leading to gates that are smaller (fewer transistors), faster (fewer transistors per stack), more energy efficient, and less prone to variability. This fact illustrates the large promises behind the functionality enhancement of devices.

4.3. Gate-Level Low-Cost Multi- V_{th} Control

As briefly introduced previously, TIG-FETs can be configured not only in terms of polarity but also in terms of threshold voltage.

The dual- V_{th} characteristic of a TIG SiNWFET is depicted in Figure 8. For a low- V_{th} configuration (solid lines), PG_S and PG_D are biased with the same voltage. In this configuration, the device is switching between on and standard off states [Zhang et al. 2013]. For a high- V_{th} configuration (dashed lines), the device is wired unconventionally as compared to a DIG device. Indeed, fixed bias voltages are now applied to CG and PG_S for p type (CG and PG_D for n type), while a voltage sweep is applied on PG_D (PG_S). Here, the device is switching between on and low-leakage off states. The difference between low and high V_{th} is about 0.3V. Note that I_{on} of both low- and high- V_{th} configurations keeps the same value since they share the same on states.

Such properties are used to create multi- V_{th} circuits in a simplified way. Indeed, traditional multi- V_{th} circuits require extra technological steps to build devices with different threshold voltages, which affects the layout regularity and increases the process costs compared to the single- V_{th} design [Matsukawa et al. 2008].

Here, the same transistors are used for the 2 configurations, leading to a drastic cost reduction.

Figure 16 illustrates two different NAND gate realizations for *high-performance* (HP) and *low-leakage* (LL) applications. In Figure 16(a), the HP gate is obtained by connecting inputs to the CGs of p -FETs. Thus, the performance for pulling the logic gate up is improved by applying the low- V_{th} configuration of the devices (solid line in Figure 8). In contrast, the LL gate (Figure 16(b)) is obtained by controlling the p -FETs from the PG_D s. Leakage power is thereby reduced by forcing the devices into high- V_{th} operation (dashed line in Figure 8). Note that in both HP and LL gates, the PG_S s and CGs of n -FETs are connected to input signals, and therefore delay and leakage in pull-down paths cannot be further tuned.




Fig. 16. NAND gates realization using TIG-FETs: (a) HP configuration; (b) LL configurations.

Extensively studied in Zhang et al. [2013], it is expected that such an approach gets, using electrical simulation, a leakage power reduction of 54% compared to FinFET LSTP transistors at 22nm technology node for a slight area overhead of 28%.

4.4. Embedded Power Gating Techniques

An efficient power gating implementation is also unlocked by the enhanced functionality offered by MIG-FETs. From a system perspective, power gating is a common and effective technique to reduce leakage power in conjunction with multi- V_{th} design. Power gating uses sleep transistors to disconnect the power supply from the rest of the circuit during standby periods. The main drawbacks of power gating are due to the series sleep transistor that: (i) reduces the speed during normal operation and (ii) increases the circuit area.

By exploiting online control of the device polarity, it is possible to create logic gates with power gating capabilities without any series sleep transistors [Amarù et al. 2013]. Based on *differential cascade voltage switch logic* (DCVSL), pull-up devices are not fixed to behave as p type but rather their polarity is online modulated by a sleep signal connected to the polarity gates.

The global concept is depicted in Figure 17. In contrast to traditional sleep-transistor-based approaches, the proposed gates can be power gated with no additional series device, thereby avoiding major performance degradation. Indeed, pull-up DIG devices are not fixed to behave as p type but rather their polarity is online modulated by the sleep signal which is connected to the polarity gates. Together with floating output countermeasures, that is, two n -type devices in parallel to the pull-down networks, the pull-up polarity control automatically provides the demanded disconnection from the power supply during standby mode. In standby mode, that is, when $sleep = 1$, the pull-up devices are switched to n type through the PGs. The CGs are tied to ground by the two additional n -type devices. Therefore both pull-up devices are in the off-state. This provides the desired isolation from the power supply. In active operation mode, that is, when $sleep = 0$, the pull-up devices act as p type. The CGs (connected to the gate outputs) are no longer tied to ground since the two additional n -type devices are in off state. The pull-down networks are now enabled to drive the outputs and close the standard feedback in DCVSL gates. Note that, during active operation mode, the circuit is the same as its nonpower gated versions' exception made for the two parallel n -type devices. The equivalent slowdown due to power gating is here only dependent on the additional drain capacitance carried by the extra parallel n -type device. Instead, in traditional power gating, the slowdown is more marked due to the extra series sleep transistor.




Fig. 17. DIG-FETs-based DCVSL style with advanced power gating scheme.

Table II. Circuit Strategies Summary

<i>Gain (vs. FinFET 22nm)</i>	Area	Delay	Power
Arithmetic Logic Compactness	0.98×	1.45×	1.48×
Multi-Transistor Merging	1.33×	1.2×	1.68×
Flip-flop Design (TSPC)	1.25×	1.8×	1.07×
SRAM Design	1.16×	1.19×	1.15×
Multi- V_{th} Control	0.84×	1.02×	2.04×
Embedded Power Gating	1.38×	3.62×	2.23×

Applied to arithmetic- and computation-intensive circuits, it has been shown, using electrical simulations, in Amarù et al. [2013] that such a technique leads to area, delay, and leakage power savings of $1.4\times$, $3.6\times$, and $2.2\times$ on average, respectively, compared to power-gated circuits using FinFET LSTP devices at 22nm technology node.

4.5. Discussions

The additional degrees of freedom brought at the device level by polarity gates of DIG/TIG FETs can be leveraged for many different design targets. A summary of the different circuit strategies and their gain with regard to current state-of-the-art FinFET technologies is given in Table II. Overall, the different techniques take advantage of the natural compactness of multi-independent gate devices to create circuits with a reduced area impact. This area reduction compensates for the impact of larger device footprints for arithmetic implementations and multi- V_{th} control, but also gives significant improvements for all other design targets reviewed in the article, that is, control logic, embedded flip-flops, large memory arrays, and power reduction techniques. In addition, we note that the different circuit techniques target help in either reducing the power numbers or increasing performance for the same power budget. The different techniques can therefore be applied in a holistic design approach to keep the pace towards higher performance with lower power dissipation.

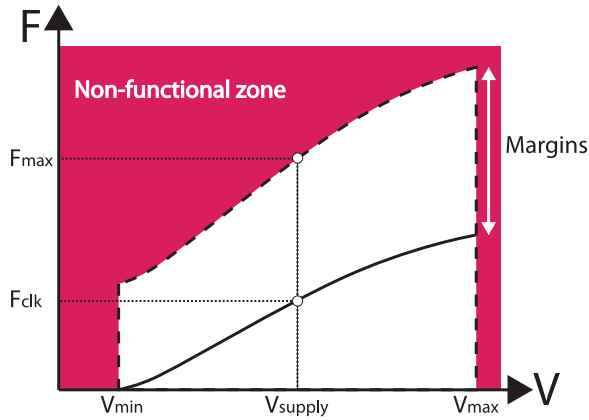


Fig. 18. VF functional domain and design margins.

5. ARCHITECTURAL-LEVEL TECHNIQUES FOR LOW-POWER SYSTEMS

In the previous sections, we showed device- and circuit-level innovations to create low-power systems with emerging technologies. However, considering the problem of power consumption from a holistic approach and integrating architectural-level techniques to further improve power control is today fundamental. In this section, we review current mainstream power management techniques and discuss them in light of advanced technologies.

5.1. Generalities

Architectural solutions to reduce the power consumption of complex systems mainly consist of controlling specific hardware actuators to reach a targeted speed while reducing the power consumption. In other words, architectural solutions try to maximize energy efficiency by always targeting an optimum power/speed functional point with respect to applicative constraints. Even if power consumption reduction relies on technological choices and on software-level application management, developing and integrating specific hardware will provide even more power reduction at the design step. As introduced in Section 2, the objective of efficiently managing power from an architectural perspective consists in dynamically controlling the circuit voltage supply, clock frequency, and threshold voltage of transistors.

A digital block is able to operate in a *voltage/frequency* (VF) range predefined at the design step, which depends on technological and environmental conditions like, for instance, *temperature* (T) variation. As shown in Figure 18, the VF domain is limited to a range $\{V_{\min}, V_{\max}\}$ and to a maximum clock frequency F_{\max} . However, due to worst-case design, the real running clock is considered to be F_{clk} and takes into account all the timing margins.

In this section, we report on architectural-level techniques for power management in processor and multiprocessor systems.

5.2. Dynamic Voltage and Frequency Scaling Techniques

The first widely used technique to reduce the power consumption consists in modifying the circuit's speed by scaling clock frequency on-the-fly. The *dynamic frequency scaling* (DFS) principle is illustrated in Figure 19. Clock frequency is scaled through a clock generator (F actuator) up to the effective applied frequency (F_{eff}) according to speed requirements [Anghel et al. 2011; Nowka et al. 2002].




Fig. 19. DFS principle on digital block.

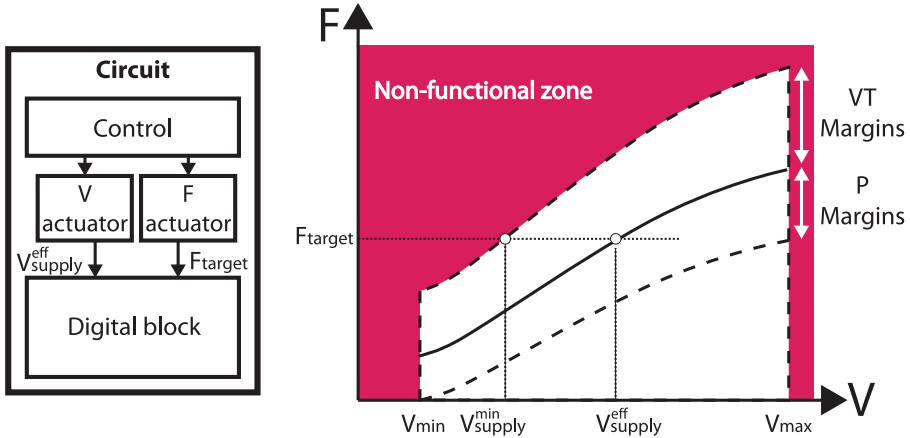


Fig. 20. DVFS principle on digital block.

This technique slightly reduces dynamic power consumption but does not improve the energetic efficiency, as the voltage supply is fixed. To improve the global energetic efficiency of the circuit, it has been proposed to scale both the voltage and frequency using *dynamic voltage and frequency scaling* (DVFS) techniques as shown in Figure 20 [Venkatachalam and Franz 2005; Kolpe et al. 2011]. This technique is also sometimes called *dynamic voltage scaling* (DVS) and aims at applying both V and F online tuning to efficiently minimize power consumption while running at a frequency, F_{target} , in line with the applicative constraints.

Actuators play a major role in architectural-level power reduction techniques. Actuators have to be implemented locally and therefore constrain the integration choices. To provide the clock frequency to a digital block, analog *phase-locked loops* (PLLs) [Maneatis 1996], fully digital PLLs [Staszewski et al. 2005; Javidan et al. 2011] and *frequency-locked loops* (FLLs) [Lesecq et al. 2011] have been used. In the context of dynamic frequency scaling, output frequencies can be scaled very fast (in only a few cycles of the reference clock), making FLL preferred.

To adjust a processor voltage supply and to avoid large embedded DC-DC converters [Reynolds 1997], V_{dd} -hopping has been considered an efficient solution [Onizuka and




Fig. 21. DVFS principle applying globally on an MPSoC and VF domain for core 9.

Sakurai 2002]. V_{dd} -hopping consists in hopping the voltage supply between two or three levels, typically V_{high} and V_{low} . The hopping between voltage levels is done dynamically during operation (typically in less than ten nanoseconds, and correlated to frequency changes) and presents very high power efficiency [Beigne et al. 2008].

In the scope of emerging technologies, designers are facing the increasingly difficult task of designing systems that handle very tight sets of specifications while dealing with numerous uncertain variables. Considering the variations in the fabrication process, temperature, and voltage, it is nowadays extremely challenging to implement systems that can meet these tight specifications while keeping high levels of integration and preserving cost effectiveness. Process variations are caused by limited control over the conditions and characteristics of the fabrication process. This limited control causes variation in the length, width, and threshold voltage of the transistors as well as the absolute values of on-chip resistors and capacitors. In addition to process variations, the system has to meet the required set of specifications for a range of supply voltages and temperatures. Process compensation techniques are widely used today by measuring, after fabrication, device parameters and by modifying the functional point of the circuit accordingly. Indeed, designs have been used to design at worst case, introducing process margins (see P margins in Figure 20) taking into account these variations [Bowman et al. 2002; Kuhn 2011]. In addition to process variation, dynamic variation is occurring in modern SoCs due to temperature gradients [Altet et al. 2001] and voltage drops [Pant and Blaauw 2008]. This second category of uncertainty leads to even larger margins (see VT margins in Figure 20). Note that, using DVFS, it could be possible to keep reducing the operating voltage down to V_{supply}^{min} if the margins were removed, or at least reduced, and therefore obtain a higher energy gain.

5.3. Power Management Techniques for MPSoCs

During the last decade, the limited performance of single processors has led to the development of distributed *multiprocessor SoC* [Wolf et al. 2008] for whom power management techniques have been adapted.

5.3.1. Adaptation of DVFS Techniques to MPSoCs. Considering these multicore architectures and DVFS efficiency, it is natural to adapt DVFS techniques to MPSoC. Due to implementation difficulties such as voltage and frequency islands and verification issues across domains, the most common technique consists of applying a single voltage and frequency point on the entire circuit as shown in Figure 21(a). This solution is




Fig. 22. Fine-grain DFVS in MPSoC.

the simplest as it requires single VF actuators and the full circuit is optimized and verified at a single *process/variation/temperature* (PVT) point. However, there is a loss of energy gain, as each processor has to run at the speed of the most constrained processor inside the MPSoC. In fact, due to previously mentioned variations, the effective performance of each processor can strongly differ. As illustrated in Figure 21, VF global adaptation is performed so that the worst core, marked w , is functional whereas core number 9 could be more energy efficient by reducing its voltage point down to $V_{\text{supply}}^{\text{core9}}$.

To overcome these issues and to manage MPSoC power and speed performance, solutions have been proposed at the system level. Task scheduling [Zhuo and Chakrabarti 2005] and task migration [Coskun et al. 2008] have been proposed to improve the resource usage and limit static power. DVFS software algorithms have been proposed to improve global speed performance [Herbert and Marculescu 2009] while reducing power consumption [Puschini et al. 2008]. Power gating or VF reduction techniques are also applied on inactive cores [Singh et al. 2007; Mahmoodi et al. 2009]. Power gating methods are widely used thanks to their fast and easy implementation. However, they still do not allow the circuits to reach best energy efficiency, as they do not consider variations that modify single-core performance. This point is fundamental while considering that emerging technologies suffer large variations that impact the speed and global power. It appears that an optimum power management technique has to take into account local variations at core level.

5.3.2. Fine-Grain Power Management. The optimum energy efficiency of a complex SoC is reached if each resource or processor is optimized at its own best energy efficiency considering the applicative constraints. Fine-grain power management is thus the best-suited technique to reach optimal global efficiency for distributed architectures, as each resource is able to run at its best VF point, whatever the global power management strategy at system level. Local optimization will depend on the task to be scheduled by a resource. A *globally asynchronous and locally synchronous* (GALS) processor array [Truong et al. 2008] with dynamic supply voltage and dynamic clock frequency circuits managed at processor granularity [Elgebal and Sachdev 2007] is the key for such an integration. As shown in Figure 22, each processor contains an independent stoppable local clock oscillator and communicates through dual-clock FIFOs, enabling an easy frequency scaling at processor level. The processors change their supply voltage by connecting their local power grid to one of two global voltages distributed over the circuit. It is also possible to disconnect the voltage supply of unused processors for leakage power reduction.

Local fine-grain DFVS with more control at unit level [Beigné et al. 2009] can be employed to further reduce the dynamic and static power consumption. To simplify the interface constraints, a *globally asynchronous locally synchronous* (GALS) circuit can be based on a fully asynchronous *network-on-chip* (NoC) [Beigné et al. 2009]. The main objective is to be able to locally generate frequency and power supply and to switch

from one VF point to another during activity without stopping the processor clock. By using a V_{dd} -hopping scheme locally controlled, it is possible to reduce static power consumption by 2 decades in standby mode, while the dynamic power consumption can be reduced up to a factor of 8.

5.3.3. Fine-Grain Adaptive Architectures. As detailed in the previous section, the DFVS technique at fine grain shows good power reduction results but has the drawback of additional hardware that increases total area and reduces power efficiency. This drawback is even more important if we consider that PVT variations should also be monitored in emerging technologies. In fact, DVFS cannot be controlled efficiently if circuit variation is not taken into account. The objective is to improve the DVFS solution by proposing low-cost adaptive techniques called *adaptive voltage and frequency scaling* (AVFS). This power control that alleviates PVT variation has to be implemented in each resource to reach the best power reduction factors.

The introduction of design margins to face PVT variation has strongly limited circuit performance and reduced the energetic efficiency. To improve the power/speed trade-off, these pessimistic margins have to be reduced. For this purpose, each MPSoC resource should be able to work at its own optimal energetic point while remaining functional and, if possible, in an automatic way. Essentially, an adaptive circuit is a circuit aware of any changes in its own characteristics and that has the ability to tune itself back to the desired state accordingly while considering applicative constraints. It requires sensor integration for locally measuring both spatial and temporal variation at the lowest cost. There is also a need for control to re-adjust the local functional VF point of the resource and thus compensate the variation. In the context of advanced technologies, the main limitation is the speed detector implementation that should consider itself the local variation. Therefore, sensors have been extensively studied to accurately monitor the evolution of the circuit [Burd et al. 2000; Nakai et al. 2005; Elgebaly and Sachdev 2007; Das et al. 2009].

Adaptive voltage and frequency scaling (AVFS) solutions are mandatory to reach best energy efficiency in the context of emerging technologies accounting for increased PVT variation and exploit the maximum of technology opportunities. A complete AVFS solution, fitting with MPSoC requirements, is discussed in the next section.

5.4. Perspectives: An Ideal AVFS Architecture Fully Unlocking Emerging Technologies

The proposed AVFS architecture sketch is based on the idea of decentralizing the power control as close as possible to the resource (actuator control, variability sensing, local adaptation). It should enable easy power consumption management while considering dynamic local PVT variation. The proposed AVFS architecture can be defined as a DVFS with local adaptation. The first challenge is to be able to implement it at fine grain with a minimum area overhead. The second challenge is to efficiently monitor PVT variation during resource operation.

Figure 23 illustrates this architecture in which each *island* is adaptive to take into account local in-die variation. Each closed-loop adaptive island consists in:

- actuators* (typically voltage and frequency, but also innovative actuators deriving from emerging devices) applied to the resource;
- variability *sensors* integrated inside the core (black squares); and
- a control block including a *local controller* and an *adjustment block* receiving, respectively, global applicative constraints and local variation measurements.

As each island has its own independent functional point, it is mandatory to have a GALS system for communication, either using a fully asynchronous communication scheme [Beigne et al. 2008] or a classical synchronous one with resynchronization

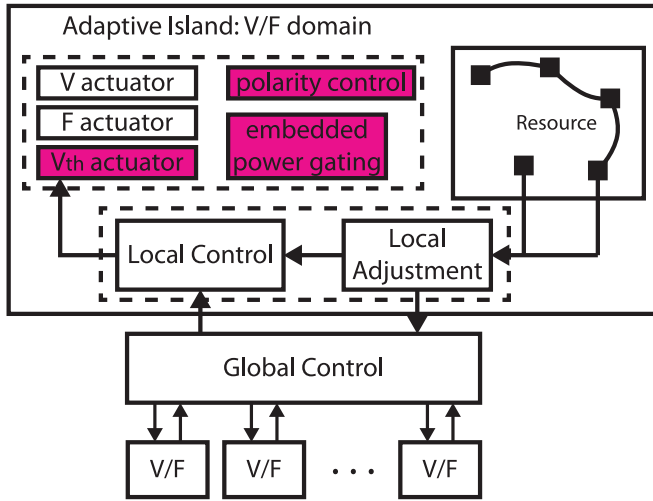


Fig. 23. Main components of the local AVFS architecture.

interfaces. The granularity of this approach depends on the application, but typically in an MPSoC architecture it can be done at processor or cluster levels. If the resource is large enough, the area overhead is reduced but intra-resource variability is not taken into account. The main constraint is thus to propose small hardware at low power cost to instrument the resource. Under these constraints, novel technologies bring a large set of innovations with intrinsic, device-level means of actuation.

Sensors are uniformly integrated inside the block to provide spatial and temporal information of resource variability; they should be thus very small and easy to integrate in a digital flow. They need to measure *process*, *voltage*, and *temperature* (PVT) variation in real time with the best possible accuracy. On one hand, process sensors give fabrication and aging information, but on the other hand, voltage and temperature sensors give a cartography of spatial and temporal gradients over the circuit.

Actuators are applying the best functional *voltage* and *frequency* (VF) point to the circuit. Their main characteristics are to generate quickly, at low-power and low-area cost, the optimal set point chosen by the control part. Many actuators, shown in red in Figure 23, can be envisaged in the context of new and emerging technologies like UTBB-FDSOI Beigne et al. [2013] or MIG FETs [De Marchi et al. 2012; Zhang et al. 2013], where energy efficiency can be further pushed by playing directly on device performance.

The control part gives actuators set points to be applied to the resource. This control is dependent on the local resource variability (from the local adjustment block) and on global applicative constraints linked to the task to be scheduled. This closed loop has to be fast enough to dynamically modify the resource functional point at its target frequency. It receives task deadline and workload information. Workload can be provided during compilation, whereas the real-time deadline constraint is generally known through the *operating system*. The local control will also choose the power supply level so that the resource can run at the targeted frequency while reducing total power consumption. Without variability information, actuators' control should account for design margins. The role of the local adjustment is thus to treat variability information to dynamically adjust VF control and reduce these margins.

Although the basis of this architecture has been demonstrated in literature [Rebaud et al. 2011; Vincent et al. 2012; Agkul et al. 2012], practical realizations fully exploiting

the additional functionalities of emerging devices still need to be explored, thereby opening broad research horizons towards even lower power consumption in complex systems.

6. CONCLUSIONS

Power-related issues are of tremendous importance in current electronic design. In this context, emerging technologies can play a significant role. In this article, we reviewed the actual trends of device research, such as fully depleted planar devices, tri-gate geometries, and gate-all-around structures, that enable an increasingly higher level of performance while reducing the associated power. Beyond simple device property enhancements, emerging devices lead to innovations at circuit and architectural levels. In particular, devices whose properties can be tuned through additional terminals, such as UTBB FDSOI or MIG vertically stacked nanowires, open the way towards fine-grain and dynamic control of device threshold and towards the design of logic gates, as well as power-related techniques, in a compact way unreachable to standard technologies. These innovations reduce power consumption at the gate level and unlock new means of actuation in architectural solutions like adaptive voltage and frequency scaling. Hence, it is of tremendous importance for energy-aware design to be considered in a holistic approach, rather than focusing on specific design levels.

ACKNOWLEDGMENTS

The authors would like to thank Mr. Luca Amarù, Dr. Esteve Amat, Mr. Michele De Marchi, Mr. Lionel Vincent and Mr. Jian Zhang for their help in the figure realization.

REFERENCES

- Yeter Akgul, Diego Puschini, S. Lesecq, Ivan Miro-Panades, Pascal Benoit, et al. 2012. Power mode selection in embedded systems with performance constraints. In *Proceedings of the IEEE Faible Tension Faible Consommation (FTFC'12)*.
- Josep Altet, Antonio Rubio, Emmanuel Schaub, Stefan Dilhaire, and Wilfrid Claeys. 2001. Thermal coupling in integrated circuits: Application to thermal testing. *IEEE J. Solid-State Circ.* 36, 1, 81–91.
- Luca Amarù, Pierre-Emmanuel Gaillardon, Jian Zhang, and Giovanni De Micheli. 2013. Power-gated differential logic style based on double-gate controllable polarity transistors. *IEEE Trans. Circ. Syst.* 60, 10, 672–676.
- Ionut Anghel, Tudor Cioara, Ioan Salomie, Georgiana Copil, Daniel Moldovan, et al. 2011. Dynamic frequency scaling algorithms for improving the CPU's energy efficiency. In *Proceedings of the International Conference on Intelligent Computer Communication and Processing (ICCP'11)*.
- Chris Auth, C. Allen, A. Blattner, D. Bergstrom, M. Brazier, et al. 2012. A 22nm high performance and low-power CMOS technology featuring fully-depleted tri-gate transistors, self-aligned contacts and high density MIM capacitors. In *Proceedings of the Symposium on VLSI Technology (VLSIT'12)*.
- Sarunya Bangsaruntip, Guy M. Cohen, Amlan Majumdar, Ying Zhang, S. U. Engelmann, et al. 2009. High performance and highly uniform gate-all-around silicon nanowire MOSFETs with wire size dependent scaling. In *Proceedings of the IEEE International Electron Devices Meeting (IEDM'09)*.
- Edith Beigné, Fabien Clermidy, Sylvain Miermont, Alexandre Valentian, Pascal Vivet, et al. 2008. A fully integrated power supply unit for fine grain power management application to embedded low voltage SRAMs. In *Proceedings of the European Solid-State Circuits Conference (ESSCIRC'08)*.
- Edith Beigné, Fabien Clermidy, Helene Lhermet, Sylvain Miermont, Yvain Thonnart, et al. 2009. An asynchronous power aware and adaptive NoC based circuit. *IEEE J. Solid-State Circ.* 44, 4, 1167–1177.
- Edith Beigné, Alexandre Valentian, Bastien Giraud, Olivier Thomas, Thomas Benoist, et al. 2013. Ultra-wide voltage range designs in fully-depleted silicon-on-insulator FETs. In *Proceedings of the Design, Automation and Test in Europe Conference (DATE'13)*.
- M. Haykel Ben Jamaa, Kartik Mohanram, and Giovanni De Micheli. 2009. Novel library of logic gates with ambipolar CNTFETs: Opportunities for multi-level logic synthesis. In *Proceedings of the Design, Automation and Test in Europe Conference (DATE'09)*.

- Lakshmi Kanta Bera, Hoai Son Nguyen, Navab Singh, Tsung Y. Liow, De-Xiang Huang, et al. 2006. Three dimensionally stacked SiGe nanowire array and gate-all-around p-MOSFETs. In *Proceedings of the IEEE International Electron Devices Meeting (IEDM'06)*.
- Kerry Bernstein, Ralph K. Cavin III, Wolfgang Porod, Alan Seabaugh, and Jeff Welser. 2010. Device and architecture outlook for beyond CMOS switches. *Proc. IEEE* 98, 12, 2169–2184.
- Keith A. Bowman, Steven G. Duvall, and James D. Meindl. 2002. Impact of die-to-die and within-die parameter fluctuations on the maximum clock frequency distribution for gigascale integration. *IEEE J. Solid-State Circ.* 37, 2, 183–190.
- Thomas Burd, Trevor Pering, Anthony Stratakos, and Robert Brodersen. 2000. A dynamic voltage scaled microprocessor system. In *Proceedings of the International Solid-State Circuits Conference (ISSCC'00)*.
- Leland Chang, David Frank, Robert Montoye, S. J. Koester, Brain Ji, et al. 2010. Practical strategies for power-efficient computing technologies. *Proc. IEEE* 98, 2, 215–236.
- Yang-Kyu Choi, Nick Lindert, Peiqi Xuan, Stephen Tang, Daewon Ha, Erik Anderson, Tsu-Jae King, Jeffrey Bokor, and Chenming Hu. 2001. Sub-20nm CMOS FinFET technologies. In *Proceedings of the IEEE International Electron Devices Meeting (IEDM'01)*.
- Ayse K. Coskun, Tajana S. Rosing, Keith A. Whisnant, and Kenny C. Gross. 2008. Static and dynamic temperature-aware scheduling for multiprocessor SoCs. *IEEE Trans. VLSI Syst.* 16, 9, 1127–1140.
- Shidhartha Das, Carlos Tokunaga, Sanjay Pant, Wei-Hsaing Ma, Sudharsen Kalaiselvan, et al. 2009. RazorII: In situ error detection and correction for PVT and SER tolerance. *IEEE J. Solid-State Circ.* 44, 1, 32–48.
- Michele De Marchi, Davide Sacchetto, Stefano Frache, Jian Zhang, Pierre-Emmanuel Gaillardon, Yusuf Leblebici, and Giovanni De Micheli. 2012. Polarity control in double-gate, gate-all-around vertically stacked silicon nanowire FETs. In *Proceedings of the IEEE International Electron Devices Meeting (IEDM'12)*.
- Cécilia Dupré, Alexandre Hubert, S. Becu, Michael Jublot, V. Maffini-Alvaro, et al. 2008. 15nm-diameter 3D stacked nanowires with independent gates operation: Φ FET. In *Proceedings of the IEEE International Electron Devices Meeting (IEDM'08)*.
- Mohamed Elgebaly and Manoj Sachdev. 2007. Variation-aware adaptive voltage scaling system. *IEEE Trans. VLSI Syst.* 15, 5, 560–571.
- Thomas Ernst. 2013. Controlling the polarity of silicon nanowire transistors. *Sci.* 340, 6139, 1414–1415.
- Wei-Wei Fang, Navab Singh, Lakshmi K. Bera, Hoai Son Nguyen, Subhash C. Rustagi, et al. 2007. Vertically stacked SiGe nanowire array channel CMOS transistors. *IEEE Electron. Dev. Lett.* 28, 3, 211–213.
- Pierre-Emmanuel Gaillardon, Luca Amarù, Jian Zhang, and Giovanni De Micheli. 2014. Advanced system on a chip design based on controllable-polarity FETs. In *Proceedings of the Design, Automation and Test in Europe Conference (DATE'14)*.
- Tahir Ghani, Michael Armstrong, Chris Auth, M. Bost, P. Charvat, et al. 2003. A 90nm high volume manufacturing logic technology featuring novel 45nm gate length strained silicon CMOS transistors. In *Proceedings of the IEEE International Electron Devices Meeting (IEDM'03)*.
- Naoki Harada, Katsunori Yagi, Shintaro Sato, and Naoki Yokoyama. 2010. A polarity-controllable graphene inverter. *Appl. Phys. Lett.* 96, 1.
- André Heinzig, Stefan Slesazek, Franz Kreupl, Thomas Mikolajick, and Walter M. Weber. 2011. Reconfigurable silicon nanowire transistors. *Nano Lett.* 12, 1, 119–124.
- Sebastian Herbert and Diana Marculescu. 2009. Variation-aware dynamic voltage/frequency scaling. In *Proceedings of the International Symposium on High Performance Computer Architecture (HPCA'09)*.
- Chia-Hong Jan, Uddalak Bhattacharya, Ruth Brian, Sang-Jun Choi, G. Curello, et al. 2012. A 22nm SoC platform technology featuring 3-D tri-gate and high-k/metal gate, optimized for ultra low power, high performance and high density SoC applications. In *Proceedings of the IEEE International Electron Devices Meeting (IEDM'12)*.
- Mohammad Javidan, Eldar Zianbetov, Francois Anceau, Dimitri Galakok, Anton Kornilenko, et al. 2011. All-digital PLL array provides reliable distributed clock for SoCs. In *Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS'11)*.
- Ali Khakifrooz, Kangguo Cheng, Qing Liu, Toshiharu Nagumo, Nicolas Loubet, et al. 2012. Extremely thin SOI for system-on-chip applications. In *Proceedings of the Custom Integrated Circuits Conference (CICC'12)*.
- Nam Sung Kim, Todd Austin, David Blaauw, Trevor Mudge, Krisztian Flautner, et al. 2000. Leakage current: Moore's law meets static power. *IEEE Comput.* 36, 12, 68–75.
- Tejaswini Kolpe, Antonia Zhai, and Sachin S. Sapatnekar. 2011. Enabling improved power management in multicore processors through clustered DVFS. In *Proceedings of the Design, Automation and Test in Europe Conference (DATE'11)*.

- Kelin J. Kuhn. 2011. CMOS scaling for the 22nm node and beyond: Device physics and technology. In *Proceedings of the International Symposium on VLSI Technology, Systems and Applications (VLSITSA'11)*.
- Tadahiro Kuroda, K. Suzuki, Shinji Mita, Tetsuya Fujita, F. Yamane, et al. 1998. Variable supply-voltage scheme for low-power high-speed CMOS digital design. *IEEE J. Solid-State Circ.* 33, 3, 454–462.
- Chong-Min Kyung and Sungjoo Yoo. 2011. *Energy-Aware System Design: Algorithms and Architectures*. Springer.
- Suzanne Lesecq, Diego Puschini, Edith Beigné, Pascal Vivet, and Yeter Akgul. 2011. Low-cost and robust control of a DFLL for multi-processor system-on-chip. In *Proceedings of the IFAC World Congress (IFAC'11)*.
- Yiming Li, Chih-Hong Hwang, Tien-Yeh Li, and Ming-Hung Han. 2010. Process-variation effect, metal-gate work-function fluctuation, and random-dopant fluctuation in emerging CMOS technologies. *IEEE Trans. Electron. Dev.* 57, 2, 437–447.
- Yu-Ming Lin, Joerg Appenzeller, Joachim Knoch, and Phaeton Avouris. 2005. High-performance carbon nanotube field-effect transistor with tunable polarities. *IEEE Trans. Nanotechnol.* 4, 5, 481–489.
- Qing Liu, Atsushi Yagishita, Nicolas Loubet, Ali Khakifrooz, Pranita Kulkarni, et al. 2010. Ultra-thin-body and BOX (UTBB) fully depleted (FD) device integration for 22nm and beyond. In *Proceedings of the Symposium on VLSI Technology (VLSIT'10)*.
- Hamid Mahmoodi, Vishy Tirumalashetty, Matthew Cooke, and Kaushik Roy. 2009. Ultra low-power clocking scheme using energy recovery and clock gating. *IEEE Trans. VLSI Syst.* 17, 1, 33–44.
- John G. Maneatis. 1996. Low-jitter process-independent DLL and PLL based on self-biased techniques. *IEEE J. Solid-State Circ.* 31, 11, 1723–1732.
- Takashi Matsukawa, Kazuhiko Endo, Yongxun Liu, Shinichi Ouchi, Meishoku Masahara, et al. 2008. Dual metal gate FinFET integration by Ta/Mo diffusion technology for Vt reduction and multi-Vt CMOS application. In *Proceedings of the European Solid-State Device Research Conference (ESSDERC'08)*.
- Carlos Mazuré, Richard Ferrant, Bich-Yen Nguyen, Walter Schwarzenbach, and Cécile Moulin. 2010. FDSOI: From substrate to devices and circuit applications. In *Proceedings of the European Solid-State Circuits Conference (ESSIRC'10)*.
- Mark P. Mills. 2013. The cloud begins with coal – Big data, big networks, big infrastructure, and big power: An overview of the electricity used by the global digital ecosystem. http://www.tech-pundit.com/wp-content/uploads/2013/07/Cloud_Begins_With_Coal.pdf?c761ac.
- Kaizad Mistry, C. Allen, Chris Auth, Bruce Beattie, D. Bergstrom, et al. 2007. A 45nm logic technology with high-k+metal gate transistors, strained silicon, 9 Cu interconnect layers, 193nm dry patterning, and 100% Pb-free packaging. In *Proceedings of the IEEE International Electron Devices Meeting (IEDM'07)*.
- Kirsten E. Moselund, Mikael T. Bjork, Heinz Schmid, Hershram Ghoneim, Siegfried Karg, et al. 2011. Silicon nanowire tunnel FETs: Low-temperature operation and influence of high-k gate dielectric. *IEEE Trans. Electron Dev.* 58, 9, 2911–2916.
- Masakatsu Nakai, Satoshi Akui, Katsunori Seno, Tetsumasa Meguro, Takahiro Seki, et al. 2008. Dynamic voltage and frequency management for a low-power embedded microprocessor. *IEEE J. Solid-State Circ.* 40, 1, 28–35.
- Jean-Philippe Noel, Oliver Thomas, Marie-Anne Jaud, Oliver Weber, Thierry Poiroux, et al. 2011. Multi-VT UTBB FDSOI device architecture for low-power CMOS circuit. *IEEE Trans. Electron Dev.* 58, 8, 2473–2482.
- Koichi Nose and Takuyasu Sakurai. 2000. Analysis and future trend of short-circuit power. *IEEE Trans. Comput.-Aided Des.* 19, 9, 1023–1030.
- Kevin J. Nowka, Gary D. Carpenter, Eric W. Macdonald, Hung C. Ngo, Bishop C. Brock, et al. 2002. A 32-bit PowerPC system-on-a-chip with support for dynamic voltage scaling and dynamic frequency scaling. *IEEE J. Solid-State Circ.* 37, 11, 1441–1447.
- Kohei Onizuka and Takayasu Sakurai. 2005. VDD-hopping accelerator for on-chip power supplies achieving nano-second order transient time. In *Proceedings of the Asian Solid-State Circuits Conference (ASSCC'05)*.
- Sanjay Pant and David Blaauw. 2008. Circuit techniques for suppression and measurement of on-chip inductive supply noise. In *Proceedings of the European Solid-State Circuits Conference (ESSCC'08)*.
- Nicolas Planes, Oliver Weber, V. Barral, S. Haendler, D. Noblet, et al. 2012. 28nm FDSOI technology platform for high-speed low-voltage digital applications. In *Proceedings of the Symposium on VLSI Technology (VLSIT'12)*.
- Diego Puschini, Fabien Clermidy, Pascal Benoit, Gilles Sassatelli, and Lionel Torres. 2008. Temperature-aware distributed run-time optimization on MP-SoC using game theory. In *Proceedings of the International Symposium on Very Large Scale Integration (ISVLSI'08)*.
- Jan M. Rabaey, Anantha Chandrakasan, and Borivoje Nikolic. 2003. *Digital Integrated Circuits*. Prentice Hall.

- Bettina Rebaud, Phillippe Maurine, Marc Belleville, Edith Beigne, Christian Bernard, et al. 2011. Timing slack monitoring under process and environmental variations: Application to a DSP performance optimization. *Microelectron. J.* 42, 5, 718–732.
- Scott K. Reynolds. 1997. A DC-DC converter for short-channel CMOS technologies. *IEEE J. Solid-State Circ.* 32, 1, 111–113.
- Kaushik Roy, Saibal Mukhopadhyay, and Hamid Mahmoodi-Meinand. 2003. Leakage current mechanisms and leakage reduction techniques in deep-submicrometer CMOS circuits. *Proc. IEEE* 91, 2, 305–327.
- Harmander Singh, Kanak Agarwal, Dennis Sylvester, and Kevin J. Nowka. 2007. Enhanced leakage reduction techniques using intermediate strength power gating. *IEEE Trans. VLSI Syst.* 15, 11, 1215–1224.
- Robert B. Staszewski, John Walberg, Sameh Rezeg, Chih-Ming Hung, Oren Eliezer, et al. 2005. All-digital PLL and transmitter for mobile phones. *IEEE J. Solid-State Circ.* 40, 12, 2469–2482.
- Xifan Tang, Jian Zhang, Pierre-Emmanuel Gaillardon, and Giovanni De Micheli. 2014. TSPC flip-flop circuit design with three-independent-gate silicon nanowire FETs. In *Proceedings of the International Symposium on Circuits and Systems (ISCAS'14)*.
- Olivier Thomas, Brian Zimmer, Bertrand Pelloux-Prayer, Nicolas Planes, Kaya Can Akyel, et al. 2012. 6T SRAM design for wide voltage range in 28nm FDSOI. In *Proceedings of the IEEE International SOI Conference (SOI'12)*.
- Dean Truong, Wayne Cheng, Tinoosh Mohsenin, Zhiyi, Yu, Toney Jacobson, et al. 2008. A 167-processor 65 nm computational platform with per-processor dynamic supply voltage and dynamic clock frequency scaling. In *Proceedings of the IEEE Symposium on VLSI Circuits (VLSIC'08)*.
- Ogun Turkyilmaz, Fabien Clermidy, Luca Amarù, Pierre-Emmanuel Gaillardon, and Giovanni De Micheli. 2013. Self-checking ripple-carry adder with ambipolar silicon nanowire FET. In *Proceedings of the International Symposium on Circuits and Systems (ISCAS'13)*.
- Vasanth Venkatachalam and Michael Franz. 2005. Power reduction techniques for microprocessor systems. *ACM Comput. Surv.* 37, 3, 195–237.
- Lionel Vincent, Philippe Maurine, Suzanne Lesecq, and Edith Beigné. 2012. Embedding statistical tests for on-chip dynamic voltage and temperature monitoring. In *Proceedings of the Design Automation Conference (DAC'12)*.
- Robin Wilson, Edith Beigne, Philippe Flatresse, Alexandre Valentian, Fady Abouzeid, et al. 2014. A 460MHz at 397mV, 2.6GHz at 1.3V, 32b VLIW DSP, embedding FMAX tracking. In *Proceedings of the IEEE International Solid-State Circuits Conference (ISSCC'14)*.
- Wayne Wolf, Ahmed A. Jerraya, and Grant Martin. 2008. Multiprocessor system-on-chip (MPSoC) technology. *IEEE Trans. Comput.-Aided Des. Integr. Circ. Syst.* 27, 10, 1701–1713.
- Jie Xiang, Wei Lu, Yongjie Hu, Yue Wu, Hao Yan, and Charles M. Lieber. 2006. Ge/Si nanowire heterostructures as high-performance field-effect transistors. *Nature* 441, 489–493.
- Peiqi Xuan, Min She, Bruce Harteneck, Alex Liddle, Jeffery Bokor, et al. 2003. FinFET SONOS flash memory for embedded applications. In *Proceedings of the IEEE International Electron Devices Meeting (IEDM'03)*.
- Hao Yan, Hwan Sung Choe, Sungwoo Nam, Yongjie Hu, Shamik Das, James F. Klemic, James C. Ellenbogen, and Charles M. Lieber. 2011. Programmable nanowire circuits for nanoprocessors. *Nature* 470, 240–244.
- Guihua Yu and Charles M. Lieber. 2010. Assembly and integration of semiconductor nanowires for functional nanosystems. *Pure Appl. Chem.* 82, 12, 2295–2314.
- Jiren Yuan and Christer Svensson. 1997. New single-clock CMOS latches and flip-flops with improved speed and power savings. *J. Solid-State Circ.* 32, 1, 62–69.
- Jian Zhang, Pierre-Emmanuel Gaillardon, and Giovanni De Micheli. 2013. Dual-threshold-voltage configurable circuits with three-independent-gate silicon nanowire FETs. In *Proceedings of the International Symposium on Circuits and Systems (ISCAS'13)*.
- Jianli Zhuo and Chaitali Chakrabarti. 2005. System-level energy-efficient dynamic task scheduling. In *Proceedings of the Design Automation Conference (DAC'05)*.
- Andrew Zukovski, Xuebei Yang, and Kartik Mohanram. 2011. Universal logic modules based on double-gate carbon nanotube transistors. In *Proceedings of the Design Automation Conference (DAC'11)*.

Received December 2013; revised May 2014; accepted November 2014