



DESIGNLINES | POWER MANAGEMENT DESIGNLINE

Power management in mobile devices—A view of energy conservation—Part IV

By Findlay Shearer 05.05.2008 0

Review: [Part I](#), [Part II](#), and [Part III](#).

Low Power Process and Transistor Technology

The genesis of low power electronics can be traced to the invention of the bipolar transistor in 1947. Elimination of the suppressive requirements for several watts of filament power and several hundred volts of anode voltage in vacuum tubes in exchange for transistor operation in the tens of milliwatts range was a breakthrough of unmatched importance in low power electronics. The capability to fully exploit the superb low power assets of the bipolar transistor was provided by a second breakthrough, the invention of the integrated circuit in 1958 (Figure 8).




*Figure 8. First IC Made by Jack Kilby of Texas Instruments in 1958
(Source: www.ti.com)*

Although far less widely acclaimed as such, a third breakthrough of indispensable importance to modern low power digital electronics was the CMOS integrated circuit announced in 1963.

Collectively, these three inventions demonstrated the critical conceptual basis for modern low power electronics.

Process Technology Scaling

The semiconductor industry has been built on the benefits of traditional CMOS scaling. Classical scaling shrinks voltages along with the lithographic dimensions, and the results were that for a given die size, scaling delivered more gates, switching faster, at a constant power.

Scaling advanced CMOS technology to the next generation improves performance, increases transistor density, and reduces power consumption. Technology scaling typically has three main goals:

1. Reduce gate delay by 30%, resulting in an increase in operating frequency of about 43%.
2. Double transistor density.
3. Reduce energy per transition by about 65%, saving 50% of power (at a 43% increase in frequency).

Unfortunately, some parameters have hard limits in scaling. A key component of the performance of CMOS transistors is the drive current or " I_{on} ." This performance is proportional to the gate overdrive, or basically the supply voltage minus threshold voltage (V_t). If we continue scaling the supply voltage then the threshold voltage must also scale to keep up with the gate overdrives or to keep the gate overdrive up. CMOS transistors are not perfect switches. They have an off-state current that rises exponentially as the threshold voltage is reduced and so while lithographic dimensions continue to scale, voltage scaling itself has slowed down. The new realities of scaling are impacting both the static and dynamic power associated with products.

Scaling Theory

Scaling a technology reduces gate delay by 30% and the lateral and vertical dimensions by 30%. Therefore, the area and fringing capacitance, and consequently the total capacitance, decrease by 30% to 0.7.

Power Consumption

A chip's maximum power consumption depends on its technology as well as its implementation. According to scaling theory, a design ported to the next-generation technology should operate at 43% higher frequency.

If the supply voltage remains constant (constant voltage scaling), the power should remain the same. On the other hand, if the supply voltage scales down by 30% (constant electric field scaling), the power should decrease by 50%:

$$V_{DD} = 1 \text{ (constant voltage scaling)}$$
$$\text{Power} = C * V^2 * f = 0.7 * 1 * (1/0.7) = 1$$

$$V_{DD} = 0.7 \text{ (constant electric field scaling)}$$

$$\text{Power} = C \cdot V^2 \cdot f = 0.7 \cdot 0.72 \cdot (1/0.7) = 0.5$$

Unfortunately, logic designs ported to next-generation technologies with constant voltage scaling do not show a decrease in power; the power remains constant. On the other hand, logic designs ported to technologies using constant electric field scaling decrease in power. This is consistent with scaling theory.

However, the power dissipation of a chip depends not only on its technology, but also on its implementation, i.e. on size, circuit style, architecture, operation frequency, and so on.

Constant electric field scaling (supply voltage scaling) gives the lower energy delay product (ignoring leakage energy) and hence is preferable. However, it requires scaling threshold voltage (V_t) as well, which increases the sub-threshold leakage current, thus increasing the chip's leakage power.

Scaling Paradox

A consequence of scaling is rather undesirable, but until recently it has not been a particularly negative feature; the standby current density increases exponentially as the length scale is decreased. Figure 9 illustrates the passive power trend based on sub-threshold currents




Figure 9. Active (Dynamic) versus Static Leakage[2]

calculated from the industry trends of V_t , all for a junction temperature $T_j = 25^\circ\text{C}$. For reference, the active power density is copied onto this scale to illustrate that the sub-threshold component of power dissipation is emerging to compete with the long-battled active power component for even the most power tolerant, high-speed CMOS applications.

As the lithography pushes forward, the product designer must devise new strategies to cope with the interference of passive power, which pushes for higher V_t (and thus higher V_{DD}) versus active

power, which demands lower V_{DD} and thus lower V_t . This results in fragmentation of device design points that address these conflicting needs in wafer fabrication.

The approach outlined in **Table 2** allows the product designer flexibility to choose the best device match for active and passive power versus performance. Products that are very sensitive to passive power, such as portable and handheld devices, may sacrifice some performance to enable higher V_t . If these designs require higher performance, they are forced to sacrifice some switching power by use of correspondingly higher V_{DD} as well.

Table 2.2: Wafer Suppliers Offer a Variety of Process Nodes to Meet Different Applications [2]

Application	High Performance	1.2-V logic	1.5-V logic	Low Power	Interface
V_{DD} (V)	1.2	1.2	1.5	1.2	2.5
T_{ox} (nm)	1.8	2.2	2.2	2.2	5
$I_{D,OFF}$ (nA/ μ m)	10	3	6	0.05	0.01

Other applications may be challenged to inexpensively conduct heat generated by active power away from the integrated circuits and thus favor lower- V_{DD} devices with low V_t and higher passive power. Thus, the variety of threshold voltages and power-supply voltages offered in 130nm technology has expanded to address these diverse needs.

Managing power as technology is scaled is a continual interaction between the product design needs and the physical constraints of the process technology. Figure 10 gives a good indication of those relationships.

Given a speed constraint, the goal is to find is the optimal supply voltage and the transistor threshold voltage while minimizing the power. Since these choices impact tradeoffs between static and dynamic power, then the application-dependent power constraints need to enter into that optimization level. Once a performance target has been set, a simple back-of-the-envelope analysis shows that the optimal supply voltage and the optimal threshold voltage choices imply a static power that is about 30% of the total power. The best choice depends on the average active power, which is strongly dependent on the circuit activity.




Figure 10. Threshold Voltage Tradeoff (Source www.freescale.com)

Transistors and Interconnects

Planar Transistor

Planar transistors have been the core of integrated circuits for several decades, during which the size of the individual transistors has steadily decreased. Rapid and predictable miniaturization was predicted by Moore's law [3]; this has allowed the semiconductor industry to make new products with added functions with each new generation of technology (Figure 11).

This predictable scaling is now reaching its limit [4] and has forced the industry to look to novel transistor architectures beyond the 45 nm technology node. At such sizes, planar transistors suffer

from the undesirable short channel effects, in particular “off-state” leakage current, which increases the standby power required by the mobile device.




Figure 11. Moore's Law [3]

In all these years of digital CMOS innovation and scaling, we have only scratched the surface of the semiconductor substrate. The Planar CMOS device, the workhorse of digital applications used in modern electronic systems, has a channel only on the surface of the silicon. These devices have a single gate on the surface of the silicon to modulate the channel on the surface of the semiconductor. Scaling of these planar devices has now begun to hit its limits for power, noise, reliability, parasitic capacitances, and resistance. New device architectures using multiple sides of the semiconductor, not just the planar surface, offer a path to overcome these performance limits. In addition, these non-planar CMOS devices enable new circuits previously not possible with single gate CMOS devices [5"7].

CMOS Switch

The fundamental function of a transistor in a digital system is a switch: to conduct as much current as possible when on, and to shut down when off. The limits of planar CMOS technologies make this fundamental operation impractical as gate lengths and supply voltages are scaled down. The current in the on state is reduced when device sizes are scaled down, due to reduced mobility of the electrons and parasitic resistances (among other effects). The leakage current increases when the device is turned off; this substantially higher leakage can drain batteries quickly, making many mobile applications difficult to engineer (Figure 12).




Figure 12. Ideal Switch Compared to a Good Switch Compared to a Leaky Switch

The fundamental limiting factors to scaling a single gate planar CMOS transistor are the leakage through the gate, and the effect of the drain taking control of the channel making it difficult to control the switch using the gate (Figure 13); this is known as short channel effects.




Figure 13. Planar MOSFET (Source: www.freescale.com)

In a multi-gate device, the channel of the device is controlled (gated) by gates from multiple sides and the body of the device where the channel is formed is made ultra-thin, so that the gate bias controls the channels more efficiently from multiple sides (Figure 14).




Figure 14. Tri-Gate Transistor, Gates Surround the Silicon Channel on Three of Four Sides (Source: www.intel.com)

These multi-gate devices have a single gate electrode, but this single gate electrode wraps around many sides and controls the channel from multiple sides. These have various names, such as, FinFET, Tri-GATE, IT Gate, etc. In all these devices, a single gate electrode controls the channel from multiple sides yielding better control of the device and lower leakage when it is shutdown and conducts more current when turned on.

Multi-Gate Transistors

In a multi-gate device, the channel is surrounded by several gates on multiple surfaces, allowing more effective suppression of “off-state” leakage current. Multiple gates also allow enhanced current in the “on” state, also known as drive current. These advantages translate to lower power consumption and enhanced device performance. Non-planar devices are also more compact than conventional planar transistors, enabling higher transistor density which translates to smaller overall microelectronics.

Dozens of multi-gate transistor variants may be found in the literature. In general, these variants may be differentiated and classified in terms of architecture (planar versus non-planar design) and number of channels/gates (2, 3, or 4).

Planar Double-Gate Transistors

Planar double-gate transistors employ conventional planar (layer-by-layer) manufacturing processes to create double-gate devices, avoiding more stringent lithography requirements associated with non-planar, vertical transistor structures. In planar double-gate transistors the channel is sandwiched between two independently fabricated gate/gate oxide stacks. The primary challenge in fabricating such structures is achieving satisfactory self-alignment between the upper and lower gates.

FinFETs

The term FinFET was coined by University of California, Berkeley researchers to describe a non-planar, double-gate transistor built on an SOI substrate, based on the earlier DELTA (single-gate) transistor design. The distinguishing characteristic of the FinFET is that the conducting channel is wrapped around a thin silicon “fin,” which forms the body of the device. The dimensions of the fin determine the effective channel length of the device.

In the technical literature, FinFET is used somewhat generically to describe any fin-based, multi-gate transistor architecture regardless of number of gates.

A 25-nm transistor operating on just 0.7 V was demonstrated in December 2002 by Taiwan Semiconductor Manufacturing Company. The “Omega FinFET” design, named after the similarity between the Greek letter “Omega” and the shape in which the gate wraps around the source/drain structure, has a gate delay of just 0.39 picosecond (ps) for the N-type transistor and 0.88ps for the P-type transistor (Figure 15).




Figure 15. A double-gate FinFET Device (Source: (www.answers.com))

Tri-Gate Transistors

Tri-gate or 3-D is the term used to describe non-planar transistor architecture planned for use in future microprocessor technologies. These transistors employ a single gate stacked on top of two vertical gates allowing for essentially three times the surface area for electrons to travel. It has been reported that their tri-gate transistors reduce leakage and consume far less power than current transistors.

In the technical literature, the term tri-gate is sometimes used generically to denote any multi-gate FET with three effective gates or channels (Figure 16).




Figure 16. SEM View of SRAM Tri-Gate Transistors (Source: www.intel.com)

Inverted T Gate Transistors

Multi-gate device architectures are rapidly evolving. The FinFET, with all its advantages, still has a significant drawback: the region between the fins is not used as part of the switch. A new family of devices called ITFET that uses both the vertical and horizontal regions of silicon has been proposed and demonstrated. The ITFET has both vertical and horizontal thin body regions shaped like an “inverted T.” The ITFET offers maximum surface area utilization on the wafer for the channel and allows optimization of crucial circuit elements such as the SRAM-based cache that is ubiquitous in all modern digital CMOS products [8, 9] (Figure 17).

Bottom pedestal offers mechanical support for thin and high aspect ratio FIN providing better scaling and yield improvement. Also allows for device width “tuning” instead of FinFET's discrete width increment which is particularly good for SRAM design. In addition, better source/drain contact that results in reduced parasitic resistance.

Process Techniques

The multi-gate device architecture requires two basic technologies that are substantially new compared to existing processes:

1. A process technology to make very thin silicon body of the order of 20 nm.
2. A process to fabricate identical gates on at least two sides of this very thin silicon.

Various process technologies have been proposed to fabricate such a structure (Figure 18). While many processes have been identified to make a very thin silicon channel, a process that easily allows gates on both sides of this channel that are aligned to each other has been provided only on the structure now called FinFET, Tri-Gate etc. Currently, most research




*Figure 17. The ITFET has a Channel that is Shaped Like an Inverted T
(Source: www.freescle.com)*

efforts to make multi-gate devices involve these FinFET device structures. While devices with sub-20 nm silicon body and gates less than 40 nm have been demonstrated, there are still manufacturing challenges to make a product with millions of such transistors [10"14].

While these novel devices make progress, new gate materials are also researched. Incorporating these new materials is crucial to gain the maximum benefits out of these novel structures. The use of metal gates instead of conventional polysilicon gates will allow less parasitic resistance and poly depletion effects. Patterning these metal gates on FinFETs with traditional oxides is a challenge, but they have been demonstrated with new process techniques (Figure 19).

Transistor Design

Figure 20 highlights the top five areas of focus on transistor design to improve the power/performance tradeoff. These areas are key for the transistor development delivering on performance improvements while minimizing the leakage currents and parasitics that reduce the performance or increase the power.




Figure 18. Multiple-Gate MOSFETs [15]. (a) Different Gate Configurations for SOI Devices: 1, Single Gate; 2, Double Gate; 3, Triple Gate; 4, Quadruple Gate; 5, New Proposed Pi-Gate MOSFET, (b) Basic Pi-Gate Fabrication Steps (Cross-Section): A, Silicon Island Patterning; B, Shallow Buried Oxide RIE; C, Gate Oxide Growth and Gate Material Deposition and Patterning




Figure 19. Schematic and SEM View of a Tri-Gate Transistor (Source: www.intel.com)

Strained Silicon

When silicon is allowed to deposit on top of a substrate with atoms spaced farther apart than in silicon lattice, the atoms in silicon are naturally stretched to align with the substrate atoms underneath, leading to “strained silicon.” Moving these silicon atoms farther apart reduces the atomic forces that interfere with the movement of electrons through the transistors, subsequently resulting in better chip performance and lower energy consumption, e.g. the electrons can flow 70% faster in strained silicon, leading to chips that can be 35% faster in performance. Strained silicon strategy involves either inserting germanium atoms into silicon lattice or adding intermediate layer of silicon-germanium compound semiconductor.




Figure 20. Top 5 Areas for Transistor Development (Source: www.freescle.com)




Figure 21. Novel Materials and Transistors (Source: T. Ghani, IEDM 2003, p. 978)




Figure 22. Integrated Tri-Gate NMOS and PMOS Transistors Demonstrate Record Drive Current Performance (Source: www.intel.com)

Strained silicon engineering is a technique used to increase the transistor channel mobility of either the nFET or pFET. Figure 21 shows a couple of examples of strain engineering. Figure 22 illustrates that improving the device performance starts with managing the tradeoff between drive current versus leakage current.

Strained silicon allows for substantial improvements in the transistors performance at a constant leakage. The performance improvement can be achieved without increasing the static power or alternatively, static and dynamic power can be decreased by lowering the operating voltage, given a constant performance goal.

High-k Gate Dielectric

The term high-k dielectric refers to a material with a high dielectric constant (k), as compared to silicon dioxide, used in semiconductor manufacturing processes which replaces the silicon dioxide gate dielectric.

Silicon dioxide has been used as a gate oxide material for decades. As transistors have decreased in size, the thickness of the silicon dioxide gate dielectric has steadily decreased to increase the gate capacitance and thereby drive current and device performance.

The problem is that as the oxide gets thinner, the rate of gate leakage tunneling goes up contributing to power dissipation and heat.

As the thickness scales below 2 nm, leakage currents due to tunneling increase drastically, leading to unwieldy power consumption and reduced device reliability. Replacing the silicon dioxide gate dielectric with a high-k material allows increased gate capacitance without the contaminant leakage effects. In addition to its dielectric properties silicon dioxide has an almost defect free dielectric interface which insures good compatibility with the silicon substrate and also the polysilicon gate electrode.

Conventional silicon dioxide gate dielectric structure compared to a potential high-k dielectric structure (Figure 23).




Figure 23. Conventional Silicon Dioxide Gate Dielectric Structure Compared to a Potential High-k Dielectric Structure

The gate oxide in a MOSFET can be modeled as a parallel plate capacitor. Ignoring quantum mechanical and depletion effects from the silicon substrate and gate, the capacitance C of this parallel plate capacitor is given by:

$$C = \frac{\kappa \epsilon_0 A}{t}$$

where

- A is the capacitor area
- κ is the relative dielectric constant of the material (3.9 for silicon dioxide)
- ϵ_0 is the permittivity of free space
- t is the thickness of the capacitor oxide insulator

Since leakage limitation constrains further reduction of t , an alternative method to increase gate capacitance is alter k by replacing silicon dioxide with a high- k material. In such a scenario, a thicker gate layer might be used which can reduce the leakage current flowing through the structure as well as improving the gate dielectric reliability (Figure 24).




Figure 24. Cross-Section of a Nch MOSFET Showing the Gate Oxide Dielectric

Gate Capacitance Impact on Drive Current

The drive current I_D for a MOSFET can be written, using the gradual channel approximation, as:

$$I_D = \frac{W}{L} \mu C_{inv} \left(V_G - V_t - \frac{V_D}{2} \right) V_D$$

where

- W is the width of the transistor channel
- L is the channel length
- μ is the channel carrier mobility (assumed constant here)
- C_{inv} is the capacitance density associated with the gate dielectric when the underlying channel is in the inverted state
- V_G is the voltage applied to the transistor gate
- V_D is the voltage applied to the transistor drain
- V_t is the threshold voltage

It can be seen that in this approximation the drain current is proportional to the average charge across the channel with a potential $V_D / 2$ and the average electric field V_D / L along the channel direction. Initially, I_D increases linearly with V_D and then eventually saturates to a maximum when:

$$V_{D,sat} = V_G - V_t$$

to yield:

$$I_{D,sat} = \frac{W}{L} \mu C_{inv} \frac{(V_G - V_t)^2}{2}$$

The term $(V_G - V_t)$ is limited in range due to reliability and room temperature operation constraints, since too large a V_G would create an undesirable, high electric field across the oxide. Furthermore, V_t cannot easily be reduced below about 200 mV, because T is approximately 25 mV at room temperature. Typical specification temperatures 100 °C could therefore cause statistical fluctuations in thermal energy, which would adversely affect the desired V_t value. Thus, even in this simplified approximation, a reduction in the channel length or an increase in the gate dielectric capacitance will result in an increased $I_{D,sat}$.

Replacing the silicon dioxide gate dielectric with another material adds complexity to the manufacturing process. Silicon dioxide can be formed by oxidizing the underlying silicon, ensuring a uniform, conformal oxide, and high interface quality. As a consequence, development efforts have focused on finding a material with a requisitely high dielectric constant that can be easily integrated into a manufacturing process. Other key considerations include band alignment to silicon (which may alter leakage current), film morphology, thermal stability, maintenance of a high mobility of charge carriers in the channel, and minimization of electrical defects in the film/interface. Materials which have received considerable attention are hafnium and zirconium silicates and oxides, typically deposited using atomic layer deposition.

Rapid and predictable scaling of planar CMOS devices is becoming difficult. New device structures to replace planar CMOS devices are being researched. Multi-gate devices using multiple surfaces are promising continued scaling, and could even make new circuits feasible.

These devices can provide new and better characteristics across all logic, memory, and analog device function. The challenges in making these devices to enter mainstream products are many. However, rapid strides in process, design, and modeling in the last few years has delivered substantial progress.

New Logic Circuits with a Better Switch

Traditional digital CMOS has operated between ON and OFF states, as explained in Figure 12. The excellent I_{on} and I_{off} characteristics of the multi-gate devices allow future scaling of traditional circuits for a few generations. Even this is not sufficient for certain low-power applications such as pacemakers, hearing aids, and other self-powered logic devices. While sub-threshold logic has been proposed as a low-power circuit alternative, it has not been widely used, due in part to the limitations of single gate devices. Multi-gate devices, with their steep turn on characteristics and extremely low leakage characteristics, promise to be ideal to make these systems practical [16].

New Analog Circuits Can Use Multiple Independent Gates

While digital CMOS logic leads the process technology roadmap for computing applications, the communications or RF applications have a substantial mix of analog components that are integrated into the CMOS logic or as standalone products. The double-gate device architecture allows better scaling of these analog applications and new functions that were not possible with single gate transistors. Just as in CMOS logic, the fundamental switch is improved by the double-gate architecture for analog applications.

The double-gate architecture offers better gain and can be used as a better mixer, amplifier or VCO.

Beyond the Transistor

Rotated substrates or junction engineering are examples of developments to deliver the best drive currents while minimizing the off-state leakages. SOI technology brings higher speed at lower power through reduced capacitances and improved performance at a given voltage.

High-k Metal-Insulator-Metal (MIM) capacitor is a decoupling capacitance technology that helps voltage supply integrity. And also, the low-k and ultra-low-k interconnect insulators reduce capacitance and thereby drive down the chip-level power.

Copper Interconnects and Low-k Dielectrics

With continuing device scaling, wiring interconnect becomes increasingly important in limiting chip density and performance. Fundamental changes in interconnect materials are needed with Copper (Cu) replacing aluminum and low permittivity dielectrics replacing silicon dioxide. The integration of these two advanced materials results in significant reduction in signal delay, cross-talk, and power dissipation, enabling the semiconductor industry to continue device scaling.

The fabrication of Cu low-k interconnects requires novel materials and processes, including electroplating Cu, dual damascene structures, chemical/mechanical polishing, ultra-thin barriers, and passivation layers. These novel materials and processes give rise to distinct structure and defect characteristics raising yield and reliability concerns for Cu/low-k interconnects.

As the technology continues to advance, the implementation of porous low-k dielectrics beyond the 65 nm node will bring in new processing and reliability issues, such as pore sealing, etch damage, and ultra-thin barriers. These problems will be discussed together with recent advances in material and process development and reliability improvement for Cu low-k interconnects.

Dual Voltage Threshold Transistor

Each transistor in a semiconductor design has an associated threshold voltage (V_t) that determines at what voltage level the transistor is triggered to open or close. Typically, a lower V_t transistor offers higher performance because the voltage does not have to swing as far to trigger the transistor (Figure 25).




Figure 25. Threshold Voltage Tradeoff (Source: www.freescale.com)

Design libraries include low-threshold-voltage, high-threshold-voltage, and threshold-voltage-with-MTCMOS (multi-threshold CMOS) cells. Multiple-cell libraries help designers to deal with both leakage and dynamic power. One library that is going to work for an entire design, because you have designs that are speed-critical, and, for the areas that are not speed-critical, you want to reduce the leakage.

A multi-cell library typically comprises at least two sets of identical cells that have different threshold voltages. Those with higher threshold voltage are slower but have less leakage; conversely, the cells with lower threshold voltage are faster but leak. Conceding a little **bit** of speed, you get a reduction in leakage. A high-threshold-voltage cell typically has 50% less leakage than a low-threshold-voltage cell with no bad side effects, such as area gain.

For most applications, designers typically use a low-threshold-voltage library for a first pass through synthesis to get maximum performance and meet timing goals. They then determine the critical paths in their design, that is, the path or paths in the design that require the highest performance. They then try to locate areas that do not require low-threshold-voltage cells and swap out low-voltage cells for high-voltage cells to reduce overall power and leakage of the design. This approach represents the most common use of the multi-threshold design technique because most applications have timing as a first requirement, low-threshold-voltage libraries run faster through synthesis, and synthesis tools ultimately produce smaller design areas from these libraries.

Synthesis tools tend to run longer and produce larger design areas when running heavy doses of high-threshold-voltage cells (Figure 26).




Figure 26. Mixing Threshold Voltages to Ensure Performance and Power Targets are Met

However, in some wireless-system applications, power is the main goal, and area increases are less of an issue. In those cases, some designers first run synthesis with high-threshold-voltage cells, find the critical path, and then swap out the high-voltage cells with low-voltage cells until they reach their performance goal.

Active Well Biasing

Essentially, standby current drain is electrons “leaking” through a gate junction of thin oxide within the transistor. The amount of current flow is a function of the voltage across the junction. The greater the difference between voltages on both sides of the junction, the greater the leakage, more current is being drained, even in standby mode.

Well-biasing techniques help control channel leakage currents. Active well bias enhances the energy/performance relationship by manipulating transistor performance real-time. With an active back bias technique low V_t device is used for maximum performance, and then raise the threshold in standby mode, thus providing two performance levels in one transistor.

Active well bias affects the threshold voltage, which in turn affects the sub-threshold leakage currents. If the body of a PMOS device is increased above V_{DD} or the body of an NMOS device is reduced below ground, the device is said to be in back bias and the sub-threshold leakage current can be reduced.

Next: Low Power Packaging Techniques

About the Author

Findlay Shearer holds a B.S.E.E. from Glasgow Caledonian University and a M.S.S.E. and M.B.A. from University of Texas at Austin. He is currently a Senior Product Manager in Freescale Semiconductor, Inc.

Printed with permission from Newnes, a division of Elsevier. Copyright

2008. "Power Management in Mobile Devices" by Findlay Shearer. For more information about this title and other similar books, please visit www.elsevierdirect.com.

RELATED TOPICS: [ANALOG ICs](#), [CONSUMER ELECTRONICS & APPLIANCES](#), [POWER](#)

THE NEXT GENERATION OF ELECTRIFICATION



The Quest for the Ideal Switch

It's time for a new era of switches that improve today's technologies and enable entirely new possibilities in the future.

GET THE WHITEPAPER

PARTNER:  menlo micro