

THE THERMAL COST OF PERFORMANCE

ELECTRIC-ENERGY EFFICIENCY SERVES AS ONE MEASURE OF HOW FAR THE ELECTRONICS INDUSTRY HAS COME. EXPLORE HOW LIGHTING, MEASUREMENT INSTRUMENTATION, AND AUDIO AMPLIFICATION HIGHLIGHT THE THERMAL CHALLENGES THAT ENGINEERS FACE.

Few aspects of the electronics industry offer milestones that mark the entirety of *EDN*'s half-century of publication. For example, with the exception of only a few limited niches, we no longer use thermionic vacuum tubes (glass FETs, for the modernists) nor are Bakelite boxes now much in vogue. On the other end of the historic interval, though the digital abstraction that dominates our modern design practice was known at the time, its practical application was only barely evident in 1956. Indeed, circuit architectures that are now as common as hands simply could not have been imagined five decades ago. As for the means of their physical implementation, what now conveniently fits into those common hands could not then have been realized on the footprint of a typical house ... if at all.

One of the few themes that does connect the dots from the time of this magazine's comparatively ancient beginnings to the current day is our use of electrical energy. In particular, our electric-energy efficiency—largely a measure of how much of the stuff we must convert into heat in the process of completing a useful task—serves as one measure of how far we've come as an industry.

A proper recounting of our industry's progress in this regard over the half-century span of *EDN*'s existence would result in a book-length work; a summary would occupy the whole of the current issue. Instead, let's take a glimpse into that progress. Analog Devices' Barrie Gilbert offers three applications—lighting, measurement instrumentation, and audio amplification—to demonstrate some of the challenges practitioners of electronics design must face.

As with any complex discipline, electronics designers build on what precedes them. That task, however, is not unidisciplinary. On the contrary, it requires an understanding of materials, processes, device behaviors, topological idioms, and system structures as they pertain to parametric performance.

Paradoxically, perhaps, significant improvement over prior art sometimes requires

a dramatic departure from the foundation practice that calls into question our assumptions, habits, and design prejudices. As Gilbert points out, after a century of hot-wires in glass bottles, they are still the dominant interior light source. The contender that holds the greatest promise as a replacement technology looks nothing like the object it will replace, save the presence of a window and a pair of contacts for electrical connection. Just as un-

likely and just as true is the fact that the replacement technology did not derive from lighting-device engineering but evolved from a path that begins with signal-processing and ends through a branch of materials science and III-Vs semiconductor processes.

Unfortunately, as a general approach to our business, waiting for a century of progress is not a winning strategy. Accelerating that progress requires a mindfulness of fundamental physics and a willingness to challenge the conventional conclusions built upon those axioms: The occasional heretical thought is good for the designer's soul. In keeping with our theme, Gilbert asks why there is a thermal cost—a use of electrical power—to perform functions at all and, in the asking, suggests that we consider

those issues that set the minimum thermal cost of performance.

Lastly, for those of you who do not typically toil at IC design in high-speed processes, Gilbert gives insight into one challenge that SOI (silicon-on-insulator) semiconductor processes pose to the circuit designer: There is not only a thermal cost of performance, but also a performance cost of thermals—in this case, one of the fundamental underpinnings of traditional IC-design practice. **EDN**

**50TH ANNIVERSARY ONLINE
THE CELEBRATION CONTINUES**

50 For more on the thermal cost of performance, read Aengus Murray's article on the evolution in the motion-control sector, which has resulted in significant energy savings. Visit www.edn.com/50th.

AUTHOR'S BIOGRAPHY

Joshua Israelsohn is director, technical information at International Rectifier Corp. Formerly, he worked as EDN's analog editor, and he continues to contribute a regular column to EDN. You can reach him at jisrael1@irf.com.

MINIMUM ENERGY AND POWER DEMANDS IN ANALOG ICs AND THE IMPACT OF SELF-HEATING EFFECTS IN TINY TRANSISTORS

By *Barrie Gilbert, Analog Devices*

You need a certain minimum energy to perform any practical operation. Raising a 12-oz can of beer from belly to lips consumes about 1J. The energy you need to execute a one-time function in electron-based systems (more familiarly stated as the power you need to repeat it—or sustain it continuously) is generally not so clear-cut. It comes down to a matter of where you stand along history's slender arrow, and the state of the art in the many relevant and intertwining technologies of an era.

In Edison's time, the minimum necessary temperature-rise of frail loops of tungsten wire in glass bottles dictated the "house current" power drain for the drawing-room chandelier. Because most of their output power fell into the infrared region, these glowing wires mainly warmed the ladies' wigs and gentlemen's bald pates. Today, we have a long list of devices for lighting homes and workplaces; they are rugged, durable, and efficient. Yet, 100 years later, the use of "hot-wires" still tops the list.

The direct conversion of an electron source to visible light—exploiting new properties of materials and devices—has made great advances in recent years. Ultimately, the electrical power necessary to generate a continuous photon-flux density limits the efficiency of these devices. No technology permits us to realize its theoretical potential; nonetheless, electron-to-photon efficiency in semiconductor emitters, such as recent white-light LEDs, is rapidly climbing toward its asymptote. These solid-state photon sources are poised to eclipse the wire in a bottle as certainly as transistors ousted vacuum tubes: tentatively at first, inevitably in the end.

But even the most ingenious technologies must also be cost-effective. "Sure" says the skeptic, "those new LED lamps are really bright, and they run stone cold! But are they as cheap as a six-pack of 60-watters from Wal-Mart? If they were to slash lifetime to one-tenth, would they cost one-tenth as much? I've heard that CPU manufacturers play that game. They choose either their 'three-year process,' whose narrow interconnects eventually fail as electromigration creates lateral filaments of metal that short to adjacent traces, or their 'seven-year process,' whose wider metal and spacing rules increase the lifetime of the product but at the cost of a general increase in the *inertia* of the interconnects and a larger die."

POWER TRADE-OFFS

Questions about the minimum energy necessary to perform a unit function (such as a single AND decision) or how much continuous power you must supply to a functional block for it to perform a certain repetitive function are among the most intriguing topics in looking toward the future of electronic signal processing. They are readily tractable in the domain of binary signaling. The energy needs in executing logical functions are usually couched as some voltage V_0 on nodal capacitance C , being CV_0^2 . A gate output must swing to its only other value, V_1 . For a capacitance of 50 fF to swing through 1V, an energy source must provide 50 fJ (femtojoules) at each rising or falling edge. When this operation repeats at a clock rate of 1 GHz (2×10^9 edges per sec), an average current of $50 \text{ fJ} \times (2 \times 10^9) / \text{sec}$ results; that is, each "action node" sips 100 μA of continu-

ous current. A million elements of this sort will happily drink 100A all day long.

In the analog-IC domain, people rarely ask questions about such issues as the minimum energy for a given function; when they do, the approach is often hopelessly academic, purely theoretical, totally out of touch with the practical world, and thus useless for most mere mortals. From an engineer's perspective, many important questions of this genre remain unanswered. However, it is also apparent that analog signals have vastly more variety and complexity of form. They receive support from deeply recursive meshes of plesiolinear elements, often deliberately using the specific nonlinearities of special elements. Each of these elements has desired or incidental inertia (energy-storage aspects) and boasts an imposingly lengthy list of parameter values. It is hardly surprising that, after an hour or two, minimum-energy considerations end up in one's recycling basket.

Analog-IC designers from 1960 to 1980, who predominantly based their designs on junction-isolated, bipolar-junction-transistor processes, judged and juggled many trade-offs, but, with a few obvious exceptions, power efficiency rarely concerned them. The emphasis was largely on maximizing performance until it was comfortably beyond the competitive limit. Power consumption was whatever you needed to meet those objectives. Few in the industry appreciated the growing importance of ICs that frugally used power. Most regarded *low-power design* as a sideshow, useful for providing thesis projects and interesting enough to justify an occasional specialist session at the ISSCC (International Solid-State Circuits Conference). Today, low-power design is at center stage.

Preconceptions about the power necessary to achieve Function X arise from the norms of present-day designs. For example: What is the minimum power a circuit must dissipate to measure a voltage applied to a probe tip? Reviewing prevalent instrumentation techniques, you might mentally list the power each major section consumes, starting with some sort of input range selector and buffer. Then you'd move on to the ADC—perhaps one of the old dual-slope variety, a charge-dispensing voltage-to-frequency converter, or a modern sigma-delta type—and its indispensable voltage-reference cell. Finally, you'd consider the matter of display elements—Nixie tubes; seven-segment, 30-ft-high Times Square illuminators; rolling metal flaps; or LEDs?

But look again at that question: As a design objective, it is incomplete. Consider what's missing: Is the source a pure-dc voltage, V_x (all t)? Or, are you chasing $V_x(t)$, a complex waveform? If you are, do you wish to determine its upper peak, lower peak, or both? Do you need to know its mean value, its rms value, and its long-term statistics? Will the touch of that probe tip seriously affect this voltage—possibly annihilating it? Should V_x be a few electrons stored on the subfemtofarad capacitance of some fragile nanogizmo? What

AN IC DESIGNER'S
RELENTLESS DEMAND
FOR CLARITY AND
COMPLETENESS IN
THE OBJECTIVES IS AN
ESSENTIAL PRECUR-
SOR TO THE EVENTUAL
SUCCESS OF ANY
NEW PRODUCT.

accuracy do you need? How long can you wait for a result?

An IC designer's relentless demand for clarity and completeness in the objectives is an essential precursor to the eventual success of any new product. The refining of objectives does not have to come from a formal *proposal*: Experience in your domain is invariably enough to recognize an incomplete specification and fill in the gaps. Here, in tightening the net, you might expand the question as follows: Using the most efficient technologies, what minimum

operating power does a handheld DVM need to visibly display the value of fixed dc voltages, from a source of less than 100V, with an error of less than 0.1%, allowing 3 sec of processing time?

With this much information, and using today's low-inertia IC processes and zero-power (although not zero-energy) LCDs, we are now listing microwatts rather than milliwatts of total power—tiny, but not zero. So you ask: Why not? Does the function of converting a voltage to a visible number fundamentally require the expenditure of any power at all? Why? We can accept that that circuit needs a lump of energy whenever you request a reading to change the state of hundreds of elements. System inertia (due to the charge-based nature of the transistors, the capacitance of the display elements, sometimes stray inductances, and other factors) is unavoidable.

But, suppose the requester really meant: "What power do you need to display the value of a fixed dc voltage for just one reading?" The thinking about this teaser is now more closely bounded, and in turn, the options that spring to mind become more specific, keyed to our familiar technologies. We wonder what to use for a voltage reference. It could still be a bandgap cell, operating at an internal bias of only 1 nA. And we can stomach this much wastage, because the READ button is pressed only once, starting the 3-sec measurement. For a given topology, the reference's native noise is non-negotiable, having its roots in transistor shot noise and the resistors' thermal noise. This characteristic does not mean that every bandgap topology will exhibit the same noise at this bias level.

The curious and industrious may wish to attempt the following exercises: First, determine what reference-noise spectral density is commensurate with a reading error of 0.1% that you attain in a 3-sec interval. (Assume no 1/f component.) Second, determine how you can reduce this noise without increasing the bias current. Third, determine how you can add the 3-sec time-out feature with the same stricture. Then, using the information you collect, calculate the one-shot energy usage. Finally, determine how you can trim the voltage to less than 0.1% absolute error, over a "handheld" temperature range—say, -5 to $+45^\circ\text{C}$ —and for all process corners.

Of course, a bandgap isn't the only choice. You could, import the pristine electrochemical voltage of an NBS

(National Bureau of Standards) West-on cell into your handheld device and store it in an IC analog memory using well-known floating-gate techniques.

Next, consider the ADC. What power will generating the clock require? Do you even need a clock? The ADC can be asynchronous. It accumulates a debt of energy during its 3 sec of activity. But, if you take only a single reading, the power averages to zero (or to some tiny value if you take a reading once every blue moon). What about that display: How much power does it need, if any? Using LCD light modulators requires another lump of energy to charge its capacitive elements but with an average power approaching zero.

Oh! The objectives conveniently fail to mention that those dc voltages fall in the range of 100V to 1 kV. Does this news affect the design of the low-power dc voltmeter that's beginning to take shape in your head? Can you still make this voltmeter work from a 1.5V supply while measuring 1 kV with no increase in consumption? Can you imagine a way of dispensing with electronics altogether in meeting that objective?

CHEATIN' THE POWER DEMON

High current consumption is often truly unavoidable, but at other times, engineers widely accept it as the sad truth until a new paradigm appears. Consider an IC-audio-power amplifier. Work backward from the speaker with a load impedance (casually assume it to be a pure resistance) of R_L and the desired maximum rms power, P_{MAX} . You can now calculate the peak output current $\sqrt{(2P_{MAX}/R_L)}$ and the peak output voltage $\sqrt{(2P_{MAX}R_L)}$. The former dictates the minimum size of the output transistors, with margins for process variations; the latter determines the minimum permissible supply voltage after deciding on the output-stage topology—single-sided or bridged—with adequate allowances for headroom. The required breakdown voltage of the transistors and such other considerations as frequency response and dynamic range narrow down the choice of IC process, then the output-stage bias mode (Class A, AB, and others), and finally, other detailed aspects of this fine architecture.

But design basics and trade-offs change over time. When faced with the need to provide high-quality audio from CMOS amplifiers, engineers dusted off and tried a very old idea—Class D. This approach didn't change the essentials of peak load current and voltage, but it drastically impacted other issues of output-stage design and eventually the entire amplifier. Most obviously, the transistors were now operating in the mode CMOS likes best: on/off switching. One consequence of this major difference is that the overall power efficiency becomes much higher. Just as the hot-wire light bulb turns most of the power it consumes into useless heat, so do classic analog-output stages. (It's what those monster-scale heat sinks are for.) Not surprisingly, the process worked, and, after a bit of learning and refining, it worked rather well. The

WHEN FACED WITH THE NEED TO PROVIDE HIGH-QUALITY AUDIO FROM CMOS AMPLIFIERS, ENGINEERS DUSTED OFF AND TRIED A VERY OLD IDEA—CLASS D.

new *binary amplifier* had emerged from one long-ago-discarded and crumbling cocoon.

First, the amplifier received a facelift by combining its core Class-D nature with other lessons from IC-switching-regulator design and sigma-delta data converters. Pulse-density methods replaced its simplistic duty-cycle modulation; the use of pseudostochastic “carrier” frequencies to broadband the EMI spectrum and other proprietary advances further augmented its sophistication. The “analog” audio amplifier has become a very-large-scale-integration digital engine.

THE RELEVANCE OF INERTIA

The ongoing development of IC-fabrication technologies led first to the significant benefits of well-balanced complementary-bipolar processes using standard junction isolation and, later, to significantly faster SOI complementary-bipolar processes using bonded wafers. In the early days, manufacturers made these SOI wafers by bringing a pair of standard 3-in. wafers into intimate contact, whereupon they would voluntarily “weld”—native oxide to native oxide. A laborious process of grinding and polishing removed all but a few microns of silicon from one of these wafers. This layer became the pure-crystal starting material on which to form transistors, starting with epitaxial deposition, followed by masking, ion implantation and drive-in, and, finally, a solitary metal layer for connections. The other wafer became simply a mechanical handle; the thin oxide layer between them became an important insulator.

Today, engineers widely use SOI processes, so these sandwiches are available commercially. SOI provides several crucial advantages. The transistors are true three-terminal devices: The absence of the usual parasitic transistor that the base, collector, and substrate layers form ensures that layout-level latch-up cannot occur. There is zero leakage current from a collector to its substrate layer. The collector-substrate capacitance, C_{JS} , is much smaller, and it is a pure, voltage-independent capacitance, unlike the varactor C_{JS} of a junction-isolated process.

However, these benefits come with one substantial setback: The thermal resistance of these devices, from the intrinsic transistor to the handle wafer, is very high, due mainly to the low conductivity of silicon dioxide (only $1/100$ that of silicon). Thus, self-heating effects are pronounced. The bottom line: It is no longer possible to make the assumption of isothermal operation. This assumption has for decades been critical to monolithic analog design, as has the assumption of reliable matching in the key parameters of a transistor pair (notably, of the V_{BE} via I_S).

AUTHOR'S BIOGRAPHY

Barrie Gilbert is an IEEE fellow and director of the Northwest Labs at Analog Devices. EDN named him Innovator of the Year for 1999.