



Published on *Machine Design* (<http://machinedesign.com>)

Embedded Systems go on a Power Diet

by

Created 03/20/2008 - 03:00

Electronics makers trim their need for power to help extend battery life and meet new efficiency standards.

Robert Repas
ASSOCIATE EDITOR

Leland Teschler
EDITOR

It used to be the case that embedded systems just had to be small enough to fit in the space available, and powerful enough to handle the processing job at hand. Not anymore. Increasingly, embedded electronics are designed with power consumption in mind. The reason, of course, is a combination of new “green” power-efficiency standards and an explosion in battery- powered uses where operating life carries a premium.

One efficiency metric kicking in this year applies to computers and servers. The 80 Plus standard dictates these devices have an energy efficiency of 80% or greater at 20, 50, and 100% of rated load with a true power factor of 0.9 or greater. The only way to reach this level of operation is with a power-factor controller and by using some clever phase manipulation in the ac/dc power converter.

Another area drawing attention is that of soft switching, more commonly known as zero-voltage or zero-current switching. Soft switching has long been viewed as a way to cut down on the generation of electromagnetic interference. In an energy-efficient world, however, it is a way to reduce the amount of power dissipated in semiconductor switches.

Innovation can also be found in the area of microcontrollers. The easiest way to lower power consumption in computer chips is to put the device into a sleep mode when it isn't busy. But chipmakers are striving to lower power consumption even when chips are sleeping. Another technique is to scale the speed of the clock so chips operate only as fast as they need to for the job at hand.

The battery budget of a baby camarasaurus

The offspring of Caleb Chung, who also created Furby, Pleo is said to be the most-sophisticated and lifelike toy ever released. Indeed, many argue whether the term toy should apply at all. They feel artificial life form (alf) might more suitably describe Pleo's status. But toy or alf, the 22 in. from nose-to-tail-tip Pleo is an engineering marvel.

From the time of its hatching, Pleo is constantly in motion whenever it's awake, lifting its head to look around, tail swinging back and forth, while it trundles along on its four stocky legs. Put another way, one or more of its 14 servomotors is always operating and consuming power.

Obviously, to maintain the illusion of a baby camarasaurus, it can't trail any power-supply wires like some stray umbilical cord. Thus energy management became a key issue with efficiency foremost. The design goal was to get several hours of "life" from a single charge of its batteries.

One aspect that underwent major design changes between concept and final production was the gearing mechanisms that handle the baby dino's motions. The original gears were noisy, an indication that they weren't meshing properly and had poor power-transfer efficiency. Ugobe turned to **Kleiss Gears** for an upgrade. Kleiss engineers used computer modeling to look at how much each gear interacted with its neighbors, then redesigned the gear shape for the best transfer. Some looked more like petaled daisies than gears, but the new gears ran quietly and efficiently.

Pleo carries a nickel-metal-hydride (NiMH) battery for its higher power density compared to nickel-cadmium. While lithium-ion batteries have an even higher power density, there was concern with using them in a toy that might see rough handling. One of the prototype Pleos caught fire because of a loose wire. Ugobe designers say this, coupled with news reports of exploding Li-ion batteries in laptops and cell phones, drove home the importance of safety. Even so, the battery-pack temperature is monitored and Pleo "goes to sleep" if it gets too high.

To conserve space and power, servomotors were pressed into performing more than one duty. For example, the eye-blink motor also opens and closes the mouth. This does limit performance slightly. In the case of the eye and mouth motor, the mouth cannot open if the eyes are shut. So Pleo can't utter any blood-curdling yells with its eyes closed.

Even with all of the energy-saving tricks, it soon became apparent that current battery technologies couldn't deliver the operating-time goal. Acknowledging that fact, a late redesign shifted from a battery totally sealed within the body cavity to one employing a replaceable battery pack. As one pack dies, it's removed and replaced with a fully charged pack that keeps Pleo alive and well for another hour.

Ugobe continues to explore more efficient microprocessors, motors, and sensors to reduce Pleo's energy appetite, if only to let the little dino do more.

HERE COME EFFICIENCY REGS

If you've bought a PC in the last several months, you may already have been affected by the 80 Plus spec. 80 Plus is now part of the Energy Star computer specification. Manufacturers figure 80 Plus supplies are about 33% more efficient than prestandard units. Moreover, they drastically cut down the harmonic distortion induced in the utility lines, thus increasing the life of the distribution transformers in the utility system.

Today, the typical electronic power supply is built with a pulse-width-modulated (PWM) topology. The idea is to rectify ac to dc, then use a PWM circuit to produce pulsed dc at a frequency much higher than that of the ac mains. The high-frequency pulsed-dc is then filtered to produce a constant dc for powering the load.

Power-factor correction (PFC) takes place after rectification and before PWM. The basic behavior that PFC corrects is the creation of current spikes on the ac line that result when current begins to conduct through the ac bridge diodes in the power supply. Current conducts to charge up the capacitor that is across the bridge and to power the load. Conduction takes place at a relatively high point on the ac voltage waveform, so the resulting spikes can have a substantial amount of energy.

PFC eliminates the power-line current spikes by pulling current through the bridge diodes at an earlier point in the ac waveform as a means of evening out the power-supply current demand. They accomplish this by incorporating a switched inductor across the bridge diodes. During the first part of the ac waveform, bridge diodes send current to the inductor. The switch then opens, connecting the inductor to the bridge capacitor and load, thus making the energy stored in the inductor available to the load.

The PFC circuit manages the time interval during which the inductor is switched. The timing changes dynamically depending on the instantaneous load. But the point of the process is to present a load to the ac power line wherein the power-supply current demand rises and falls with instantaneous ac-supply voltage, thus simulating a purely resistive load.

In practice, PFC uses two inductors and two switches to ensure the circuit looks resistive under both light and heavy loads. And to keep the supply “green,” interleaved switching features synchronize the PFC and PWM stages and reduce switching noise. At light loads, switching frequency drops to reduce power consumption. And if the supply load is sufficiently light, the PFC will shut off to further cut power drain.

Soft switching is another trick in the bag of switch-mode power-supply designers. The basic idea is to configure the switching transistors so they only turn on or off when there is no voltage applied across their terminals. Also known as zero-voltage switching, this technique has long been applied in solid-state relays and other switching elements to reduce radiated RF noise.

Switching power supplies still use soft switching to reduce noise, but an additional rationale is to reduce the resulting energy lost during switching. The usual way of accomplishing this is through use of circuit resonance that keeps power transistors off until their terminals are at zero volts.

It is easy to understand where the energy goes in the absence of soft switching by examining what happens when the switching transistor changes state. The switching interval is about a half microsecond in typical switchers. In the absence of soft switching, the voltage across the transistor begins to fall at the same time as the current begins to flow. The presence of voltage and a current flow means power gets dissipated within the switch during any period of turn-on or turnoff. The problem is particularly bad during turn-off when the switch will be carrying its full current load.

Energy lost in such switching has become worse in recent years as manufacturers have slowed the rise and fall times of switch-mode supplies to reduce RF interference. Slowing rise and fall times, of course, proportionally boosts the power lost in the transistor switch during transitions.

New power convertor topologies avoid this energy loss typically through use of constant-frequency resonant switching, aka soft switching. The basic idea employs parasitic output capacitance in the switching transistor (usually a MOSFET) and parasitic leakage inductance of the power transformer as a resonant circuit. Electrically, the inductor and capacitor are in series and parallel with each active switch. Additional circuitry is added to losslessly recover the LC energy and send it to either the load or the input.

There are many different circuits to accomplish soft switching. They all employ a special switching sequence optimized to limit energy loss. The overall efficiency improvement is on the order of 2%, accounting for, as an example, a savings of more than 20 W in a 1-kW power supply.

One of the problems inherent in a softswitch scheme is that the capacitive and inductive components involved are temperature sensitive. So digital control is the means used to dynamically monitor operating conditions and optimize circuit operation. “We use transistor dead time to do the soft switching,” says Freescale Senior System and Application Engineer Charlie Wu. **Freescale Semiconductor** is one of the chipmakers that field ICs for managing soft switching in PWM-style supplies. “Based on the load, you need to adjust the dead time dynamically, widening it for a large load, shortening it for a small load. If the dead time stays constant you distort the waveform,” he says.

TIME TO WAKE UP THE DINOSAUR

For a good example of how to conserve power in embedded controls, look no further than the 3.3-lb Pleo. Jammed with 38 sensors to detect light, motion, touch, and sound, the robotic pet carries six processors and 14 servomotors. Pleo's manufacturer, **Ugobe** in Emeryville, Calif., used 32-bit and 8-bit ARM processors from **Atmel Corp.** to control the motors, sense Pleo's surroundings, and make dinosaur noise.

Battery life in Pleo comes at a premium — those servomotors burn power at a healthy clip. So the Atmel processors onboard employ a variety techniques to minimize their own current drain. Perhaps the most obvious way is to just go into sleep mode when they aren't doing anything important. But problems can arise when the circuit wakes up. Specifically, it makes a difference how the circuit gets back into a normal operating state. In Pleo as well as in many other embedded applications, it's important to transition into a full-awake mode as quickly as possible. The end effect otherwise can be, say, a hand tool that doesn't instantly respond when you pull the trigger, or a radio-controlled toy that is slow on the up-take.

“In a wake-up cycle, you don't want to reboot the system,” says Atmel Corp. Product Marketing Engineer Jerome Gaysse. “You might want to save the context of the chip before shutdown to come out of standby more quickly.”

Another trick: Some chips have built-in RC oscillators that temporarily serve as fast clocks just during start-up. Once things have settled down, the system clock once again takes over.

And there is more than one kind of low-power mode. The two most common types are power save mode and idle mode. Everything is off during power save except for a clock that keeps track of time. Idle mode is characterized by selectively shutting off parts of the circuitry but with the main parts of the microcontroller still functioning. The differences in power consumption between these modes can be significant. For example, one Atmel controller chip operating at 1.8 V consumes 340 μA when active but only 150 μA in idle mode and 0.65 μA when in power save. Interestingly, the device consumes 0.1 μA when completely powered down. The nonzero current drain is because of leakage currents inherent in the semiconductor process used to make the chip and the geometries involved. In general, the smaller the geometries, the higher the leakage current. Hence another trade-off: Bigger chips with larger device geometries leak less. More-compact chips fit in smaller spaces and may consume less current when active, but power lost to leakage current can be more of a problem.

Clock management is another key tool for saving power. A chip that runs with a slower clock consumes less energy than one running faster. So speed-of-operation is a trade-off against power needs. And often, not all parts of a circuit need to run at the same speed. So it may be possible to slow down some sections of the design while others crank away. Or, designers may only apply the clock signal selectively, temporarily shutting off circuits that aren't in use. “You might save 50% of the power you'd normally use this way,” says Gaysse.

But not all designers know how to design a system so it can switch from low to high frequencies without locking up. “The issue is generally synchronization when you have multiple clock domains,” says Gayse. “You don’t see a lot of architectures with multiple clocks right now, but it is a growing trend.”

“It’s tricky to do scaled clocking,” says Freescale’s Wu. “If modules aren’t structured carefully, they’ll lock up during the transitions.”

Finally, fundamental decisions about the type of processor can impact power consumption. Risc-style architectures have a reputation for being power misers simply because they use no multiply instructions. They handle multiplication and division with sequences of adds and subtracts. In contrast, Cisc architectures with multiplication instructions are relatively power intensive because multiplications involve several steps and use more circuitry than simple adding and subtracting.

Nevertheless, the choice between Risc and Cisc may not be straightforward from the standpoint of power consumption. “Though a Risc processor consumes less, it uses more instructions to do the same operation,” points out Freescale’s Charlie Wu. “That means a Risc machine may need to operate longer to get things done. So sometimes the comparison between Risc and Cisc power dissipation can be misleading.”

MAKE CONTACT

For more info on what’s inside Pleo and a look at a Pleo teardown:


tinyurl.com/yvput3 ^[1]

tinyurl.com/ynoc7w ^[2]

tinyurl.com/yprthk ^[3]


Atmel Corp., atmel.com ^[4]

Freescale Corp., freescale.com ^[5]





Ugobe’s Pleo resembles a baby Camarasaurus. But with his skin off, the toy looks like six Atmel CPUs, 28 sensors, and 14 servomotors mounted on various circuit boards. Power consumption

is dear. Early reviews say the battery only lasts for about an hour of active play before needing a recharge. The Atmel processors inside employ various means of keeping their power needs to an absolute minimum.



Interleaving in pulsewidth-modulated power supplies is a key technique employed to hit the Energy Star 80 Plus spec. The idea is to synchronize the PFC and PWM switching to reduce the harmonic content of the resulting current. The MC56F8013 in this diagram is a Freescale processor often used to implement power factor control.





SAVE POWER - GATE A CLOCK -- Power consumption in CMOS logic circuits is proportional to clock-toggle frequency. So one widely used technique to save power is to disable the clock going to circuits temporarily not in use. As illustrated by this diagram from Atmel, a typical approach is to gate clock signals going to different parts of the system with a sleep-mode disable signal. Analog components and I/O are often gated separately. Some processors also have power-reduction registers (PRRs) built in that can selectively disable sections of the system even when the sleep mode is not engaged. Another similar technique is to gate-out the clock to storage elements whose state has not changed since the last clock cycle, thus avoiding unnecessary updating.

Source URL: <http://machinedesign.com/article/embedded-systems-go-on-a-power-diet-0320>

Links:

- [1] <http://tinyurl.com/yvput3>
- [2] <http://tinyurl.com/ynoc7w>
- [3] <http://tinyurl.com/yprthk>
- [4] <http://atmel.com>
- [5] <http://freescale.com>