

Engineering portable power

Improving system efficiency is one way to extend the run time of portable products. By **Rick Zarr**.

Providing mobility in a product can provide additional revenue and open new markets beyond the scope of fixed applications. A good example of this is the portable ultrasound market.

It wasn't long ago that a person had to go to a clinic for an ultrasound imaging procedure. In most developed countries, this is usually not a problem. However, in remote villages and towns having the ability to bring equipment to the patient provides much improved health care. There are many challenges in trading off weight, size and operation time when designing mobile devices. With the efficiency of power conversion routinely exceeding 90%, many engineers are going back to the drawing board to find areas of their designs that can yield efficiency improvements in various functions to lower the overall energy consumption.

It is often said that, when looking for gains, the most obvious or easiest opportunities should be examined first. When power conversion efficiencies hovered around 60% to 75%, the first large gains were made by converting from linear to switching regulators; a move which greatly improved the overall system efficiency. Now, with the availability of integrated high efficiency switching regulators, engineers must look beyond the power conversion for improvements.

Size, weight, heat dissipation and cost are all drivers for the mobile market and often drive the decision process. Currently, batteries are the weak link in the system and have not kept up with developments in areas such as

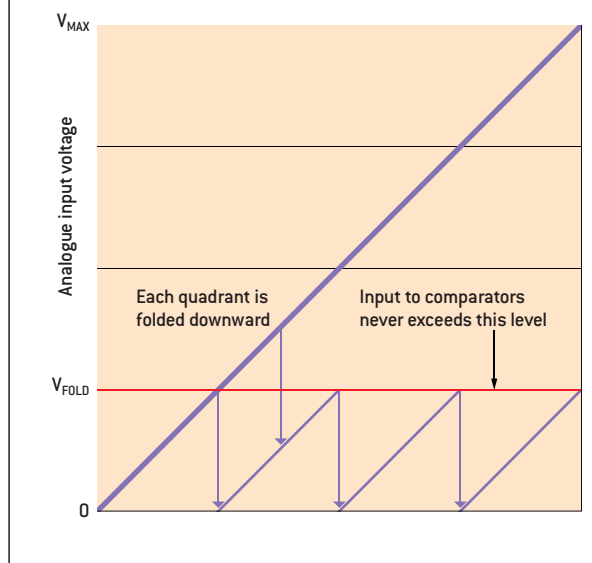
semiconductor processes. With modern power supplies providing excellent efficiencies, the next prospect for reducing losses is in the system architecture itself.

In the last few years, Intel and other cpu manufacturers began to realise that running the cpu faster was not necessarily the best way to get more performance. Their main concern was the heat generated by the processor, as well as the peripheral component dynamic requirements. By migrating to multiple core architectures and providing operating systems that are multicore aware, much larger gains in performance (and lower power dissipation) were realised.

Just as the cpu vendors stopped participating in megahertz wars to provide more performance, mobile product designers should step back and take another look at the way functions are performed.

Analogue to digital conversion is an area that has seen such architectural changes. For example, National Semiconductor pioneered the integrated folding converter that not only runs extremely fast (Gsample/s), but which also uses the least amount of energy in the process. Traditional flash type converters were limited by the number of comparators that could be integrated. The number of comparators required in a flash a/d conversion is a function of the number of output bits ($2^{n\text{-bits}}$). For example, a 10bit flash converter would require 1024 comparators plus thermometer code to Gray code to binary conversion electronics and a very precise and uniform resistor divider ladder.

Fig 1: Foldback converter approach



ZARR: "THERE ARE MANY CHALLENGES ... WHEN DESIGNING MOBILE DEVICES."



Folding converters take an entirely different approach by using a small number of comparators (32 to 64 typically) and 'folding' the input range to always stay within the limits of the comparator network (see figure 1). The trick is to compensate for increased integral and differential non linearity caused by the folding process.

This architecture represented a new way to solve an old problem and greatly reduced the amount of energy required to perform the function. In this way, the energy consumed was reduced from tens of watts to roughly 3W for the PowerWise ADC10D1000, a dual 10bit device capable of converting at gigasamples per second – a major power saving in portable imaging, radar and software defined radio systems.

In large scale asic or system on ic designs, architecture is similarly

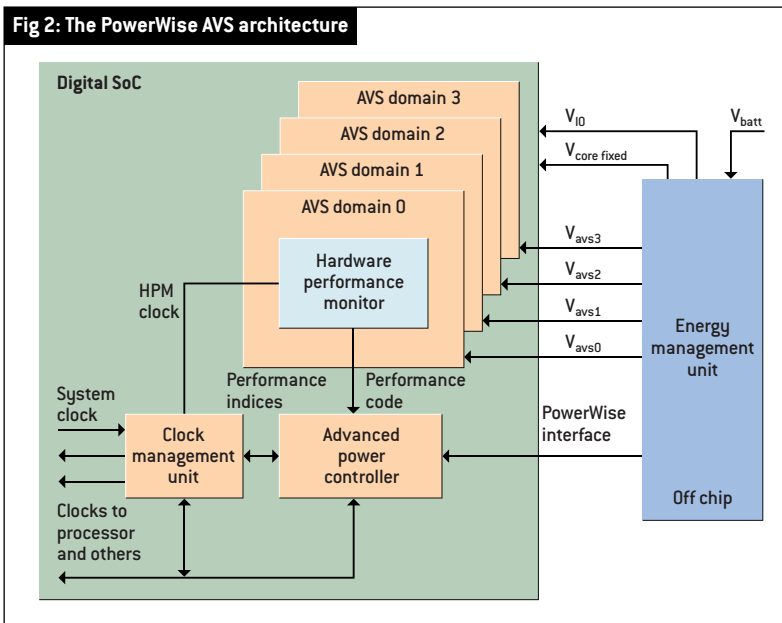
important. One issue that continues to show up, even with shrinking process geometries, is the dynamic and static losses associated with CMOS transistors. Examination of equation 1, the energy equation for CMOS, shows a frequency dependent dynamic term and a static leakage term. As processes continue to shrink in scale, issues emerge with both components.

yields, devices will all consume the same amount of energy. A solution to this problem is to alter the architecture to adapt to the environment of the device. A technology known as Adaptive Voltage Scaling (AVS) does just this.

AVS works by including a digital subsystem that monitors the condition of a device – it is synthesised along with the application digital logic – and alters

provide isolated voltage islands and frequency domains. Where a previous design may have had a single voltage supplying all of the core logic, a new low power design may incorporate multiple voltage islands where the clock domains are constrained limiting the dynamic requirements. Furthermore, these islands can use voltage scaling techniques or simply use a lower core voltage due to slower timing constraints.

Fig 2: The PowerWise AVS architecture



$$E = [\alpha f_{\text{clk}} V^2 + V_{\text{leak}}] \cdot t_{\text{task}} \quad (1)$$

Capacitive loading and shoot through currents might be reduced, but the number of devices on a chip has increased causing higher dynamic power consumption per chip. Static losses caused by sub threshold leakage, drain-source extension leakage and electron tunneling, as well as short channel effects such as Drain Induced Barrier Lowering, are becoming a much greater problem with large digital ASICs.

When designing large digital systems, timing must be maintained over the entire range of operation, which includes supply voltage, process and temperature variation. This design constraint sets the power consumption to the worst possible level: even at more modest temperatures or faster process

the supply voltage to various islands inside the chip dynamically. As performance requirements change, the AVS logic internal to the chip sends updates to an external power management device called the Energy Management Unit (EMU) to increase or decrease the supply voltage to that voltage island. As shown in equation 1, the dynamic term provides the largest gain since it is a function of the square of the supply voltage. Even though the static term is a linear function of supply voltage, significant savings can be realised here as well due to reduced leakage currents.

To maximise energy savings, architectural design is once again critical. For AVS or other voltage scaling techniques to work to their fullest, system designers must reconsider the way they divide up the functionality and

Left: Adaptive Voltage Scaling allows the chip's supply voltage to be varied as performance requirements change.

Right: National Semiconductor's PowerWise technology has allowed the energy consumed by its ADC10D1000 a/d converter to be reduced from tens of watts to roughly 3W



With an ever increasing demand for portability, especially in the field of medicine, communications and defence, engineers need to look beyond the power converters for large gains in system efficiency. Reviewing the architecture for novel and sometimes less obvious methods to perform a function may yield much higher efficiency improvements – especially with power converters now routinely providing efficiencies in excess of 90%. Battery technology will eventually catch up with gains made in processes and IC design, but until engineers have higher energy densities, system efficiency will provide a solution for extending run time and reducing thermal dissipation.

Author profile:

Rick Zarr is PowerWise technologist with National Semiconductor.