

Designing for Performance and Stamina - Processing at Smart Speed

Power consumption must be as important a design consideration for mobile communications devices as speed and ability.

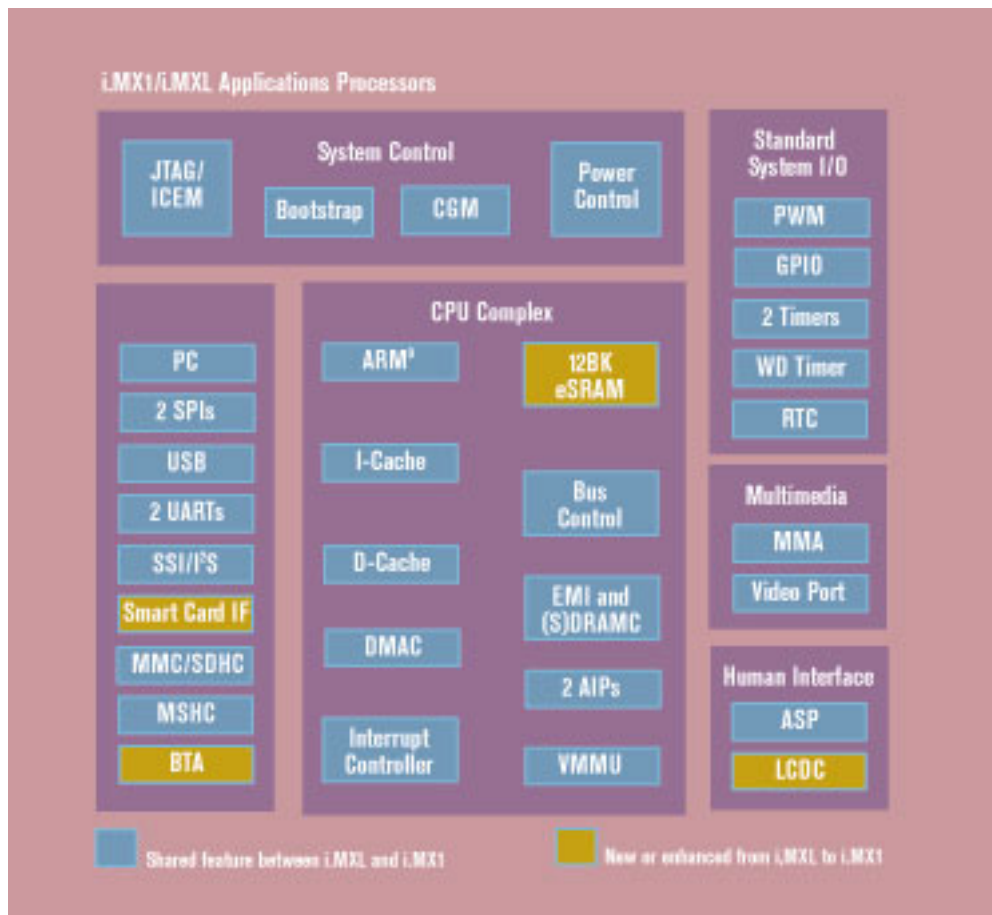
By Kathleen Wiggernhorn, Motorola, Wireless and Broadband Systems Group
As we all take our first steps into a brave new world of third-generation mobile, we are participating in a revolutionary new way of conducting business, where mobility no longer limits our productivity. We are no longer tethered to a power outlet or confined to a few square feet of office space. We see ourselves as having all the power and sophistication of a hard-wired home or office right in the palms of our hands. However, we also know mobile systems have an Achilles heel – a mobile power source. No matter how mobile devices become, the battery will always limit the amount of work performed on them. That is why low power consumption must be as important a design consideration for mobile communications devices as speed and ability, resulting in what we call performance with stamina – the ability to support a broad range of demanding applications while having the battery life to make them useful.



It's Not About the Megahertz

Key to ensuring performance with stamina is the processor at the center of the device. More than anything else, that is key to reducing the time to perform the tasks required, while extending the time you have to get the work done. To listen to some vendors offering processors to the mobile market, one would think this problem can be solved by throwing raw megahertz at the processors. But that theory has flaws — improved performance is not about just cranking up the clock and getting more megahertz through the CPU. It's about using the megahertz you have to get more work from the entire system — without wasting any time or energy in the process. Over-clocking the processor does nothing to mitigate a slow memory system. Processors should be engineered to do the same work quickly at low clock speeds; and low clock speeds translate into critical power savings. Recent studies suggest that the different design approaches can improve performance by up to 70 percent, as well as save 23 percent in power consumption.

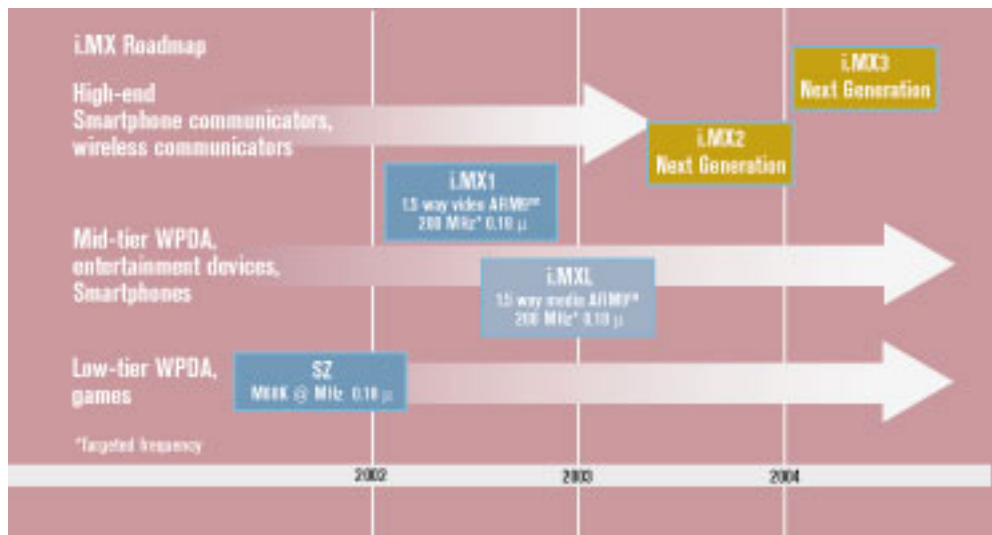
True, high megahertz processors with a very wide, very fast system "bus" can afford to run flat-out, pushing data unencumbered down the pipe. However, in portable wireless systems, busses tend to be much slower and narrower and can only carry so much data at a time. When you increase the CPU clock, you need to feed it with data and instructions. As this data is pushed through the bus and hits a bottleneck — trying to read and write to the external main memory — an over-clocked processor will just burn up excess energy with the extra megahertz and subsequently drain the battery.



Beyond the Processor

Establishing a new processor architecture is very difficult. That's why many vendors, including Motorola, turn to technology based on ARM[®] cores which have gained industry-wide acceptance and offer many advantages. Once the core is chosen, there are a number of design and manufacturing strategies for achieving economical performance. Some vendors, such as Motorola, are very adept at these practices to achieve performance with stamina, including the following:

Dual V_T - Each transistor in a semiconductor design has an associated threshold voltage (V_T) that determines at what voltage level the gate is triggered to open or close. Typically, a lower V_T transistor offers higher performance because the voltage doesn't have to swing as far to trigger the gate. However, the manufacturing techniques necessary to produce low V_T transistors tend to make them higher leakage devices, so the standby current drain is higher than what you want to include in a portable, battery-powered product. In the past, manufacturers had to decide whether they wanted to tune a device for high speed, in which case they would have poor standby current, or design it using higher V_T transistors, which would not run fast but would have less standby current drain. Now, we have the ability to include both high and low V_T transistors in the same chip, using only low V_T transistors for those critical path circuits required to get the speed that we want. The bulk of the device can then use the higher V_T transistors, which have the advantage of lower standby current.



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. By biasing up the substrate voltage to more closely match the voltage on the other side of the transistor, the voltage across the junction is lowered, thus reducing the leakage. This “well-biasing” helps reduce current drain on a transistor level whereas Dual V_T is chip-level technique. Both methods can be used in a single device to minimize standby current drain.

Dynamic Voltage Frequency Scaling (DVFS) - This is simply adjusting the clock speed and power supply on the fly to lower current drain when full speed operation is not required. For instance, there is no reason to run a clock at 200 MHz if the application only requires 50 MHz. Slowing the clock means the operating voltage can be lowered, which, in turn, lessens the demand for battery power.

Direct Memory Access (DMA) - This facilitates data movement with minimal processor intervention. DMA is engineered to allow pathways between peripherals and memory to bypass the processor. It increases efficiency and performance because of short and direct data paths, and it is designed to save the processor from allocating power and performance to this particular task. This means the CPU can either be used to perform other functions, which increases system performance, or, since it has less work to do, it can be slowed down to save power.

Clock Gating - This is an effective strategy for reducing power consumption while maintaining the same levels of performance and functionality. Basically, a circuit uses power when it is being clocked, but leaks a smaller amount of current (measured in micro-amps) when its clock has been gated, or tuned off. By shutting off the clocks of unused portions of the processor, some manufacturers have realized significant power savings during operation.

Partitioning - Many tasks in an embedded system can be implemented in either dedicated hardware or in software on a programmable core. Generally, software is more flexible and cheaper, but hardware is faster and consumes less current. The challenge is partitioning tasks between hardware and software to take best advantage of their attributes to get the fastest yet most efficient solution. A better model for speed and efficiency can be realized by committing intensive machine cycling tasks to hardware accelerators rather than software. However, this must be accomplished in a manner that preserves the flexibility of software where that

flexibility is most important.

For instance, stable, computer-intensive functions required in key applications can be built in hardware to get the performance edge when flexibility is not the major issue. Functions that are less well defined and can change frequently, such as Digital Rights Management where there is no universal standard, the value of software flexibility may outweigh the benefits of dedicated hardware. Intelligent hardware/software partitioning decisions of this sort accomplish a great deal toward high performance, low power results.

Keeping Up

Battery life, or stamina, in the real world is every bit as important to the wireless user as performance. Ask anyone who has lost a connection because his or her mobile phone has run out of battery power. Unfortunately, advances in processor performance, with their associated appetites for power, are far outstripping advances in battery life. Moore's Law states that processor performance doubles every eighteen months. Meanwhile, battery energy density only doubles every ten years. What's more, system performance adheres to another law -Shannon's Law. It shows that communication system performance requirements are doubling every eight-and-a-half months -about twice the rate of processor performance and more than fourteen times the rate of improved battery life. This is why manufacturers, such as Motorola, work so hard to save energy in their designs, from transistor-level to system-level.

We must continue to push the performance limits, however, we cannot afford to emphasize speed over stamina. We know that battery technology will never keep up with us, so it is up to us to close the gap between system power requirements and the availability of battery power.

Source URL (retrieved on 01/31/2015 - 4:09am):

<http://www.wirelessdesignmag.com/product-releases/2003/07/designing-performance-and-stamina-processing-smart-speed>