

# Differentiate or die

Developing an ultra low power microcontroller using commoditised technology still leaves room for innovation. By **Philip Ling**.

The overriding demand – or ‘megatrend’ – for ultra low power technology is well established and, as with any megatrend, the reasons why are numerous. Predominantly, portability drives the need for low power in the high volume consumer market, but for the rest of us, portability isn’t necessarily the issue. Outside consumerism, the demand for ultra low power is dependent upon three Cs: convenience, longer operation between charges or replacing batteries; consumption, lower power devices mean lower energy bills; and continuity, trends don’t become megatrends overnight.

While technology developed for the thinner end of the market traditionally trails the consumer curve, the need for ultra low power has permeated most applications and many technology companies are focusing on the more stable and accessible applications; penetrating the supply chain is simpler in markets that aren’t dominated by a few household names. Volumes may not be as great, but there are many more emerging applications in the ‘light industrial/commercial’ market, such as metering, environmental monitoring and general white goods.

This sector remains the domain of the microcontroller, rather than high performance SoCs targeting ‘convergence’. The momentum behind ultra low power helps ensure demand stays high in these applications, too. But it also ensures competition is high amongst providers, as leading semiconductor vendors strive to win and protect market share by continually bringing ultra low power solutions to market.

For established companies with product portfolios founded on in house technology,

such as Microchip, the move to ultra low power isn’t necessarily a differentiating feature. Whilst claiming to offer one of the industry’s lowest mcu standby currents, this is a development of Microchip’s existing product range and therefore offers greater choice.

But Microchip is almost in a minority, in terms of designing devices and making them itself. While in the competitive market of microcontrollers this model still prevails, the future landscape may be substantially different. The inexorable shift towards 32bit processing has allowed ARM’s range of Cortex-M cores to continue its penetration of the mcu product sector. Leading developers, such as STMicroelectronics, NXP, Renesas and Toshiba, have adopted the ARM core and/or are beginning to standardise on ARM’s technology.

Many are also outsourcing manufacture to foundries, taking advantage of recently developed ultra low power processes, which offer leading edge performance on more affordable ‘behind the curve’ process nodes. The question becomes, therefore, in this homogeneous domain is it still possible to

differentiate a product using pure-play fabrication and commercially available third party IP, particularly when it comes to power?

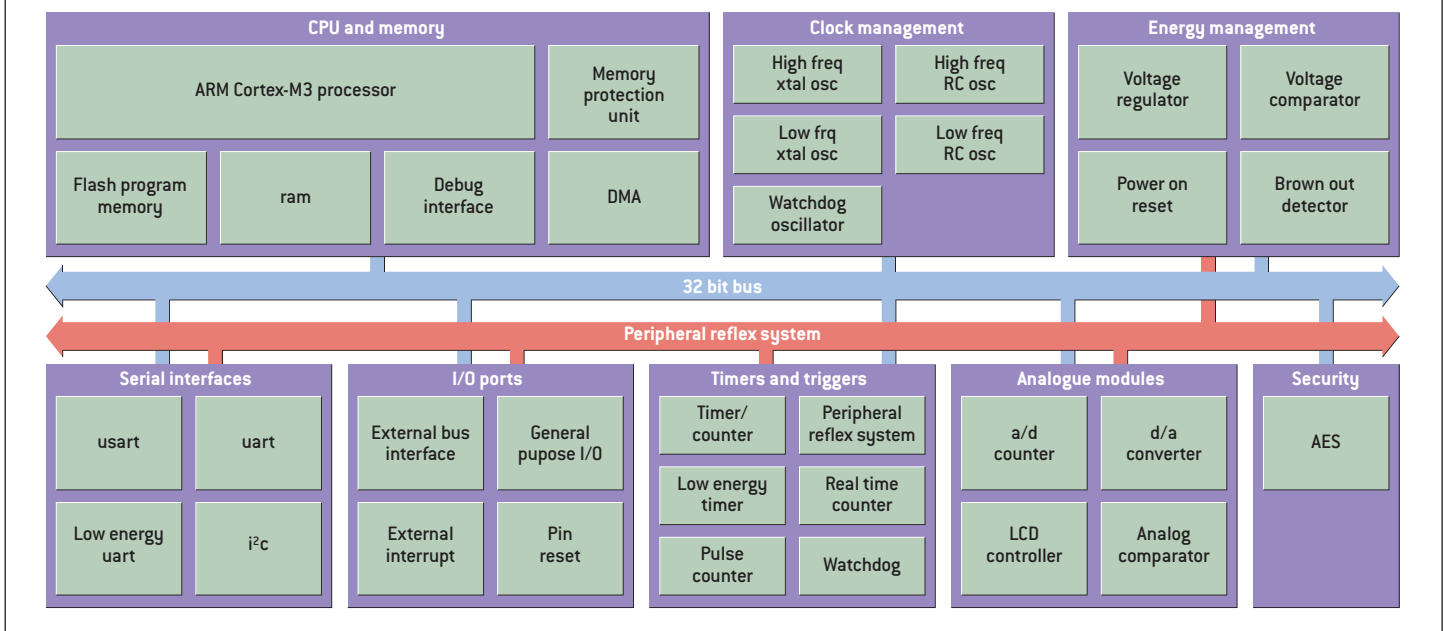
One company believes it is and has recently announced its first product, which fits squarely in the light industrial/commercial sector, targeting applications that spend much of their time doing nothing and must therefore achieve extended operation on a single battery – ultra low power.

Energy Micro has based its first range of products on ARM’s 32bit Cortex M3 core, yet it claims to achieve significantly lower power than devices from other manufacturers using a similar approach.

The device in question is the first member of the EFM32 family. Energy Micro claims the family – which runs at up to 32MHz and provides 1.25DMIPS/MHz – consumes a quarter of the energy required by existing microcontrollers. Drawing less than 180µA/MHz



Fig 1: Block diagram of the EFM32 microcontroller



while executing code from flash, the EFM32G has a typical standby current consumption of 900nA and, in its deepest sleep mode, a consumption of 20nA.

Just as there are various forces that must come together to create a megatrend, the key to the EFM32's low power is a combination of many factors. Firstly, the fabrication process; leakage affects every single transistor in an integrated device, such that passive leakage is a major contributor to overall power consumption. But, despite using a commodity 180nm ultra low leakage process, the company claims it has improved the standard performance by modifying the way it is used.

Secondly, the core; although it is using a standard Cortex-M3 core, the company has extended the low power modes offered (table 1) to achieve more control over power consumption. Thirdly, the peripherals; by using peripherals that are largely autonomous, the core is able to remain in a deeper sleep mode for more of the time, thereby reducing total power.

**Less processing, lower power**

Rasmus Christian Larsen, support and training officer with Energy Micro, explained that the key to the EFM32's low active power isn't to get the most out of the ARM core, but to actually use the ARM core as little as possible. This may seem strange, given that the Cortex-M family has seen

widespread adoption, thanks largely to its low power/high processing combination. But the point he is making is that if you can execute tasks in hardware, then you are taking no processing power.

To achieve this, Energy Micro has developed all of the EFM32's peripherals in house, with the exception of the direct memory access (DMA) controller. Crucially, while it is the DMA controller that enables the autonomous functionality of the peripherals, it is the flexibility of the peripherals themselves that creates that autonomy.

Larsen described the low level driver library that has been developed for the peripherals, which allows the M3 to configure the peripherals and then leave them to complete their tasks without further intervention. This allows the core to return to other tasks or, more importantly, to go in to a deep, power saving sleep mode. Of course, the peripherals have also been developed to be as energy efficient as possible.

Looking more deeply into the EFM32's architecture, it is clear that the majority of the hardware blocks (Figure 1) have been developed with ultra low power as a priority. The voltage regulators are described as being 'extremely efficient', which in part enables the core to execute directly from the on chip flash memory. In conjunction with that, the energy management unit handles the extended sleep modes, which works cooperatively with the clock

management unit. Here, the use of multiple oscillators running independently allows large parts of the device to run slower or stop altogether. Importantly, the restart time for the oscillators is claimed to be much lower than competing devices, meaning fewer wasted clock cycles and therefore lower active power.

Another crucial element on the clock architecture is the use of asynchronous logic and clock gating throughout the device. Larsen described the use of asynchronous logic and clock gating in the EFM32 as a major differentiator, not only in their use but the way they have been systemically integrated to the overall architecture.

While the methods used in the EFM32 achieve significant reductions in power consumption, there is a potential penalty: applications that depend more heavily on the processor than the peripherals could suffer. It may be for this reason that the Cortex M3 was chosen over another less power efficient core.

However, for other licensees of the same core family, the challenge is to match or better the low power capabilities of the EFM32 – in other words, to differentiate or die.

**ne** For further information on any subject visit: [www.newelectronics.co.uk](http://www.newelectronics.co.uk)