

Catering for complexity

Application class processors are finding their way into the embedded sector as system complexity increases. By **Graham Pitcher**.

Over the years, ARM has developed a broad portfolio of processor cores, each with the ability to be customised to meet the needs of particular users. On that basis, you could be excused for thinking that there is a core for every application; but that's not the case and is something which was highlighted recently with the introduction of the Cortex-A32 core.

But one thing which stands out is the A32 is being targeted at the embedded world; previously, members of the Cortex-A portfolio were generally used where the application requirements – and A stands for application – were more demanding. Embedded systems have, in general, been the province of the Cortex-M range, with M standing for microcontroller.

Ian Smythe, director of product marketing for ARM's CPU group, explained the thinking behind the launch. "The embedded market is pretty varied, with a wide range of products. At one end, there are the high A class cores, such as the A72. These run applications on complex SoCs. Often algorithmic, with the need for lots of memory and GPUs, these apps need 64bit embedded processing.

"At the other end are 32bit cores which only cost a few cents. These have very small memories and dedicated use cases – sitting in sensors, for example. While billions of devices have shipped featuring M class cores – appropriate to small memory footprints and RTOSs – the embedded market is pushing to places which need support for a more complex OS."


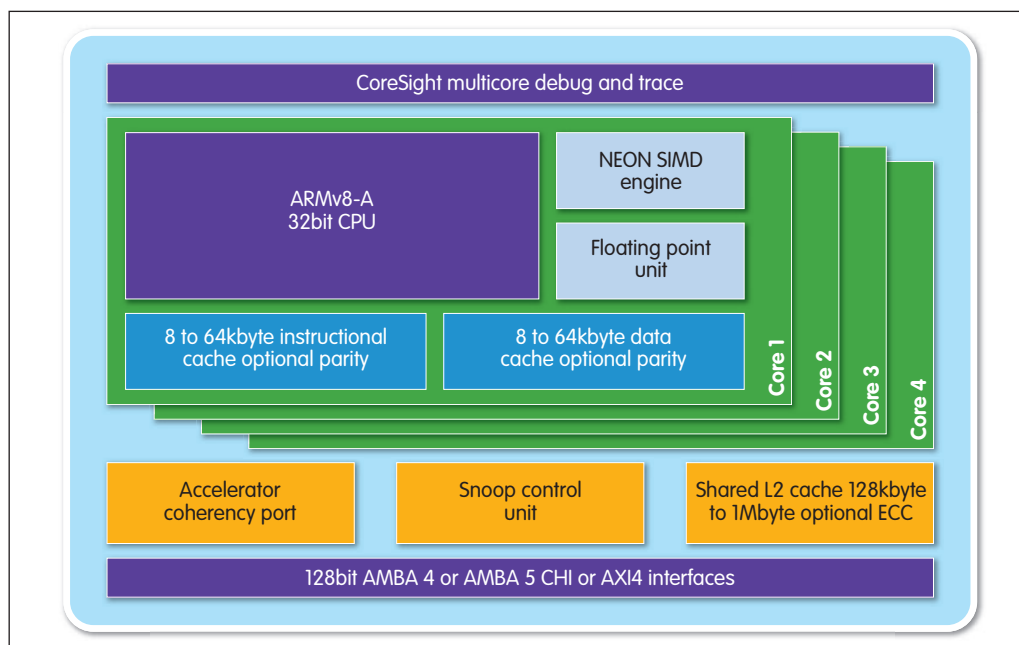
Richard York, ARM's VP of embedded marketing, added: "The Cortex-M and -A worlds are moving closer. They started touching each other a couple of years ago and now overlap. The Cortex-M7 has a lot of flash, a lot of performance and is a useful way to move upwards. But applications are pushing downwards.

"When we launched the M7, people said there was an overlap. That's a good thing; it means designers don't have to jump into a different world if they don't want to."

Put another way, it's all about power, silicon area and the ability to support more software complexity. As York pointed out: "It's about software. We are in a world driven by software, not hardware. While MCUs are useful if you want high levels of integration, if you're working in a world where software efficiency is 'king', you need to be able to embrace CPUs." Step forward cores like the A32.

And one reason why ARM sees the need for cores such as the A32 is the Internet of Things. Charlene Marini, vp of segment marketing, noted: "The IoT is opening opportunities for platforms and A series cores are becoming appropriate. It's difficult to say what will be needed five years

The Cortex-A32 can be synthesised with up to four cores, but its smallest implementation, right, occupies just 0.25mm² of silicon



from now, but nothing we're seeing at the moment points to a dominant core."

This is something which Geoff Lees, general manager of NXP's MCU business, noted a couple of years ago,

when he told New Electronics: “We are beginning to see more IoT applications suited to low end Cortex-A processors, rather than high end Cortex-M cores.”

Smythe said: “Cortex-A cores bring two things to the embedded world. One is the ability to use a rich OS, like Windows or Linux. These use virtual memory, so they need to run on an A class core and we need to support that.

“But A cores also bring performance. Some apps need high levels of performance, while meeting cost, power, efficiency and security targets.”

manager, commented: “While there has been a lot of interest in the v8A architecture, there is a legacy 32bit software ecosystem and, for this sector, the Cortex-A7 is the efficiency standard. We have attempted to give this sector something which is more efficient than the A7, whilst also bringing the latest enhancements to the 32bit world.”

In Dave’s opinion, the A32 is ‘about 25%’ more efficient than the A7 in terms of integer performance per mW consumed. “We’ve achieved this using power reductions, along with

instructions for authentication and protection, it can be coupled with TrustZone CryptoCell-700 series products to enable enhanced cryptographic hardware acceleration and advanced root of trust.


Phil Burr, product marketing manager with ARM’s CPU group, added: “Each ARM core is trying to solve different problems and the A32 builds on that by bringing the v8A architecture into the embedded sector. But the difference between the A32 and other A cores is that the A32 is a 32bit only core. However, it features 100 new 32bit instructions, which bring improvements in cryptography, multimedia and NEON performance.

“It is an option for those designers who would have thought about using an A5 or an A7 in the past. But one of the key things about the A32 is that it can be configured with from one to four cores, so it gives designers the ability to tailor their design to suit the market.”

In its smallest configuration, the Cortex-A32 occupies less than 0.25mm² of silicon, while consuming less than 4mW at 100MHz in a device manufactured on a 28nm process.

Dave concluded: “There will be so many applications for the A32 core; it will be a good candidate for wearable products, but also for things like entry level set top boxes, TVs on a stick and so on.

“Someone building an SoC on a modern process, say 28nm, will say ‘if I combine process advantage with new processor, can create more compelling products’. And that’s what we’re trying to achieve,” he concluded.



So what is the justification for the A32 core when ARM only recently unveiled the A35?

“Most of the embedded software ecosystem is 32bit,” said Smythe, “but there are embedded apps that will need 64bit processing. A32 expands the offering and the v8A based core will meet those demands.

“There are niches where the A35 can handle 32 and 64bit processing; apps where every mW and square millimetre matter. But there’s also the A32, where developers can get additional efficiency and area savings. It has presented ARM with an opportunity to refine the architecture and to create a derivative core.”

Kinjal Dave, a senior ARM product

improvements to the microarchitecture. If you’re looking at apps where battery life really matters, then a 25% improvement is important.”

The Cortex-A32 can be configured in multiple ways, from single to quad core. This, says ARM, makes it scalable enough to serve the smallest and most efficient compute devices through to IoT gateways and industrial compute applications.

“But we have also added other things,” Dave continued. “For example, the A5 and A7 cores don’t support ECC parity and there’s now a crypto block.”

The A32 includes TrustZone technology to provide a security foundation for SoC hardware. While the A32 includes cryptographic

“The embedded market is pushing to places which need support for a more complex OS.”

Ian Smythe

