Processor Selection for Embedded Systems

Most embedded systems need some type of processor, but which one would you select? Would you go for the one that is familiar because it is what 'we have always used,' or would you benchmark them with your application code first? Well, it could be a mix of both and then some others, which this article will help you determine

PANKAJ V. AND DILIN ANAND

When we began working on this article, the first question in our mind was about what made it worthwhile to spend time on this. A few questions asked to some of my knowledgeable friends in system design gave us our first answer for the story.

For starters, it is the best way to reduce power consumption. Simply moving a generation ahead on a manufacturing process node gives you vastly reduced power consumption figures, or more power for the same performance—if that is what you want. This, in turn, reduces the heat generated by the processor, which means smaller heat sinks that allow you to get into an even smaller casing, or even go fanless.

On the other hand, you could select the processor such that it can handle all your number crunching and polygon-rendering madness. This lets you ensure a smooth UI, get new functionality in or go for cost-effectiveness.

Overall, the right selection helps you enhance the user experience by reducing the chance of system overheating like what happened in some of the recent mobile devices, or any other untoward event that might make your customer curse the designer.

Where do we begin

Apart from which central processing unit (CPU) to select, there is also the problem of whether you should have gone for a graphics processing unit (GPU) rather than a CPU. So let us find out which one works best for you.

A CPU is like an executive that keeps on switching between tasks, directing IOs and peripherals to perform their operations and handling virtual memory functions. The GPU, on the other hand, is like a factory employee that can handle repetitive tasks with flair. It cannot switch tasks as fast as the CPU, but it has a lot more arithmetic logic units (ALUs) that make it better at performing mathematical tasks.

"Most CPUs have an execution unit with branch or loops capability, which is used for controlling other logic blocks. GPUs are meant for highly parallel calculation, where data streams are operated upon, not for control application like a normal CPU is," explains Satish Bagalkotkar, president and CEO, Synapse Design Automation.

The spokesman goes on to explain that GPUs are optimal for tight code with high parallelisation since they have hundreds of simple execution units with SIMD (single instruction, multiple data) capabilities. Algorithms that can be parallelised easily are best implemented on GPUs. A GPU normally has smaller memory availability in normal implementation than a CPU, so a GPU is probably not well suited for a data-intensive application.

Typically, a GPU assists video rendering for a CPU, but they also have some additional applications. Earlier servers did not need GPUs as they were managed over a text-based interface. However, firms have already announced GPU-based servers. Some of the key markets where CPU and GPU coexist are mobile devices such as smartphones and tablets, financial trading, video processing, medical imaging, etc. Some of the high-end GPUs with a control CPU are extensively used in industrial simulations, games, entertainment and video processing. There is also an increased trend of high-performance GPU adoption in the finance area.

How does a GPU get more work done, whereas CPUs run at a higher frequency? We guess an example could be a case of BitCoin mining. The four cores of a CPU like the Intel i7-4770k, albeit running at a higher clock of 3.5 GHz, will be able to execute far less instructions per clock than a GPU like AMD Radeon R9 290X, which has 2816 stream processors at a comparatively lower clock of 1000 MHz.

How to select a GPU

Within GPUs, you can choose between solutions that have more complex shaders and those with simple shaders—the difference being that the GPUs with simple shaders will pack a lot more into the same die space. But which one do you need?

"It is really application dependent. The real differentiator is whether the shader implements the functions and calculations needed by the application being discussed. If the shader does not implement the specific instructions needed, the complexity is irrelevant. From the design point of view, a simple shader is easy to implement and replicate to meet the performance," says Satish.

How to select a CPU

The most important consideration for selecting a CPU is the application itself. Each embedded application comes with a specific set of requirements.

Nate Srinath, founder and director, Inxee Technologies, explains, "The first step is to identify the domain of the application. These domains are nicely coupled with various industry verticals such as automotive, industrial automation, medical electronics, defense electronics, consumer electronics, mobiles, etc. Once the domain is identified, the next step is to identify the application requirement in terms of functionality, complexity, performance, cost, area, environmental details, manufacturing details, and power needs."

Processor performance. Usually, we start this section talking about the clock speed, but we have learned that the instructions per second, operations per clock and the efficiency of the computation units are also equally important things to be considered.

"The most important parameter in my mind is cost vs performance ratio for a processor. The next important parameter will be the longevity of the product guaranteed by the manufacturer. These two parameters will ensure the cost-effectiveness and the longevity of the product getting designed using the processor," advises Shanmugasundaram M., associate director-PES at Happiest Minds Technologies.

The instructions per second give you an idea of the computation power for large computers, whereas a higher operation per clock means a computer that might have lower clock speed could still perform competitively. Current processors have multiple cores and GPUs on the same die, which provides enhanced performance without sacrificing on power consumption or thermal design power (TDP). Benchmarking your shortlisted processors with a trial run of the application code is a good way to see if it meets your requirements.

The hardware acceleration feature in a processor is useful to speed up tasks with less software involvement. "It enables the multiprocessing option. The desired performance may be achieved by either a processor with higher processing power or with a slower processor accompanied by

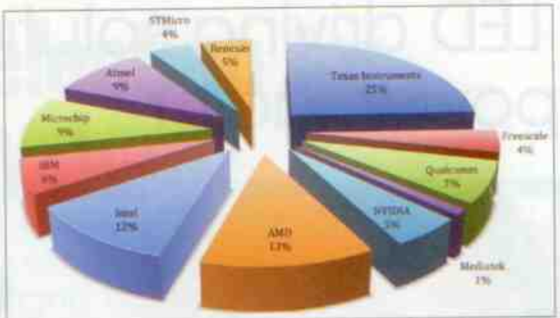

Pie chart of voting done amongst our online readers. The question was—which vendor would you buy a processor from, for your next embedded system?

a set of hardware accelerators. Both approaches have their pros and cons. The merits and demerits of each of the approach would have to be evaluated against assigned weightage for the product requirement," explain Vijay Bharat S. and Sachidananda Karanth-lead architects, Mistral Solutions.

"We have to consider that by doing hardware acceleration of certain modules, expandability might become a trade-off. Consider the example of today's processor chipsets having hardware acceleration for H.264 video paired with a ARM9/11 processor core for other processing. These chipsets are not usable for upgrading to the next generation codec—HEVC, because of the trade-off done when the CPU selected was considering H.264 hardware acceleration use-case only," explains Saravanan T.S., marketing, Semicon & Multimedia, Tata Elxsi.

Industry requirements. Such industries as automotive and medical have a different set of requirements for electronics designed for them. In a car, mission critical systems, such as the engine and safety systems, have stringent standards that are defined by organisations like the International Standards Organisation (ISO) and the Automotive Electronics Council (AEC). The factories have to be ISO certified, while the components and systems made from these factories have to adhere to AEC guidelines. Examples are the ISO/TS16949 safety standard and the AEC-Q reliability guidelines, which include AEC-Q100, Q101 and Q200 documents.

"Additionally, operating temperature range matters since many embedded applications are deployed in harsh environments where temperature shocks are present and thermal performance of processors (and other components) can make or break the solution," explains T. Anand, managing director of Knewron.

The quality management standard for medical electronics is set by the ISO 9001:2008 and ISO 13485:2003 standards. In the US, the FDA 21 CFR Part 820 is a standard pertaining to good manufacturing processes for medical devices. The devices are also classed separately depending on their target application—toothbrushes go into Class I, infusion pumps and stents go into Class II while an implantable heart pump would go in Class III. The ISO 14971 standard specifies the process of risk management for medical device manufacturers.

Certain embedded applications like medical electronics require high levels of safety. "Applications like a media player require hardware acceleration to play media files. Applications in wireless communications require high levels of security. Storage and networking applications require high expandability. In essence, the priority of the functional requirements is driven by the domains and the application itself. What is important to an application is not necessarily important to another application," adds Nate Srinath.

Support. The peripherals integrated on the processor are important as well, as these are key drivers of the overall bill of material (BOM) cost of the end product—higher the overlap between the peripherals required by the application and the peripherals integrated on the processor, the more optimal will be the BOM cost apart from the benefit of simplified board design complexity. "Given the fast-paced nature of today's electronics market, one needs to also consider the software support (on an average, it is estimated that 70 per cent of the effort