Changing Demands Are Changing the Role and Applicability of SoCs and ASICs

Multicore architectures and raw silicon speed are closing the performance gap between general purpose processors, SoCs and ASICs. This makes it possible to do in software all of the functions on a GPP that would have previously required an ASIC.

by Brian Calder, IntervalZero

Companies building embedded systems that depend on SoCs and ASICs have reached a tipping point. They cannot continue to meet their customers' increasingly aggressive product-feature demands and at the same time accommodate equally aggressive product-refresh rate expectations. In fact, SoC and ASIC platform limitations are forcing longer product cycles—not to mention more expensive and higher-risk investments—in order to deliver the required product features.

This emerging dilemma is creating opportunities for companies that can harness the power of x86 and ARM general purpose processors (GPPs) for the applications that traditionally relied on SoCs and ASICs. Indeed more and more SoC manufacturers are integrating GPPs such as ARM and Atom into their solutions.

Over the past two years, *RTC* magazine has referred to innovative approaches that address the critical issues of cost, time and risk as application services platforms or ASPs, a name coined by *RTC* Editor-in-Chief Tom Williams. To date, ASPs have been defined as a new class of integrated circuits that combine a CPU, a standard set of configurable peripherals, and a programmable fabric—all on a single device. An April 2011 article in *RTC* magazine noted that "because they combine logic programming with software programming, these new devices offer design teams the maximum flexibility, enabling users to rapidly develop unique functionality for whatever application they are targeting."

Still, why limit the ASPs to SoCs and ASICs? The definition can be expanded to include a promising—and revolutionary—all-software approach that runs on GPPs like x86 or ARM, and an off-the-shelf OS with real-time capabilities. Any successful, long-term ASP must be able to deliver successive generations of product in shorter product cycles—not longer—and do so while reducing costs. While the question remains as to what the next dominant platform will ultimately look like, there are some key features required from the ASP to deliver the needed cost, time and risk-mitigation breakthroughs.

Context is important to understand why an ASP is important and how a breakthrough is possible. What are the market drivers and technology advances that make a breakthrough in time-to-market possible and what are the key features of the ASP itself?

Multicore: The Platform of the Future

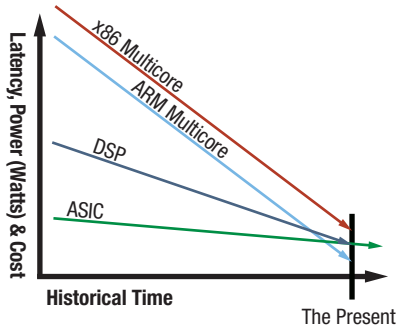

FIGURE 1

Over recent years, the latency and performance gap between general purpose multicore processors such as x8 and ARM has been closing to the point that many of the functions that required the specialized hardware properties of DSPs and FPGAs can now be done in software in the GPP.

Market Drivers

There will always be a market for ASICs or SoCs, but today's economic reali-

ties reduce the number of suitable projects. For example, consumer electronics have a shorter time-to-market, high product volume, and lower cost per unit than corporate electronics. More costly designs that also have less time-in-market for a given product limit the ability for companies to earn a return on new products. Of course, as opportunity for the time-in-market shrinks, so does the feasibility of using an ASIC.

Aberdeen's "Best in Class" study and Design Trends reports, which track ASIC pre-silicon architectural trends, confirm the growing complexity of ASIC chips at all levels of design metrics. Contrast this complexity with the continued decrease in ASIC project starts, and the conclusion is that ASICs may be getting larger in size, but less numerous in unique projects.

Besides the increase in design time and cost for SoC and ASIC, there are other profound technology advances in other parts of the chip industry that threaten the domain of the SoC: multicore GPPs-namely x86 and ARM (Figure 1).

With Intel's relentless drive to double performance every 18 months and with the inclusion of technology like AVX to handle Signal Processing needs, x86-based solutions can now feasibly—actually very favorably—compete with SoC solutions.

Clearly, x86 is not a solution for a camcorder, but for any powered system like X-Ray or shop floor control it is already being used today. As multicore processor power consumption and prices come down by 50 percent every 18 months, it will not be long until SoCs feel even more pressure.

ARM has had profound impact on the SoC and ASIC vendors. By providing a low-power GPP that was openly licensable, ARM started a technology adoption rate that has rarely been seen. With the help of standardized platforms supporting the ARM platform, such as Android, iOS, Linux, etc., ARM has become the de facto standard for all handheld devices. The ASP is actually recognition by the ASIC and SoC vendors that the GPP can handle more and more of the processing. Still, in order to gain the most benefit, an all-software solution running on a GPP best meets the requirements to shorten the time–to-market. The concept is called Hard Real-Time Platform.

Integrated Platform Addresses Hard Real-Time System Requirements

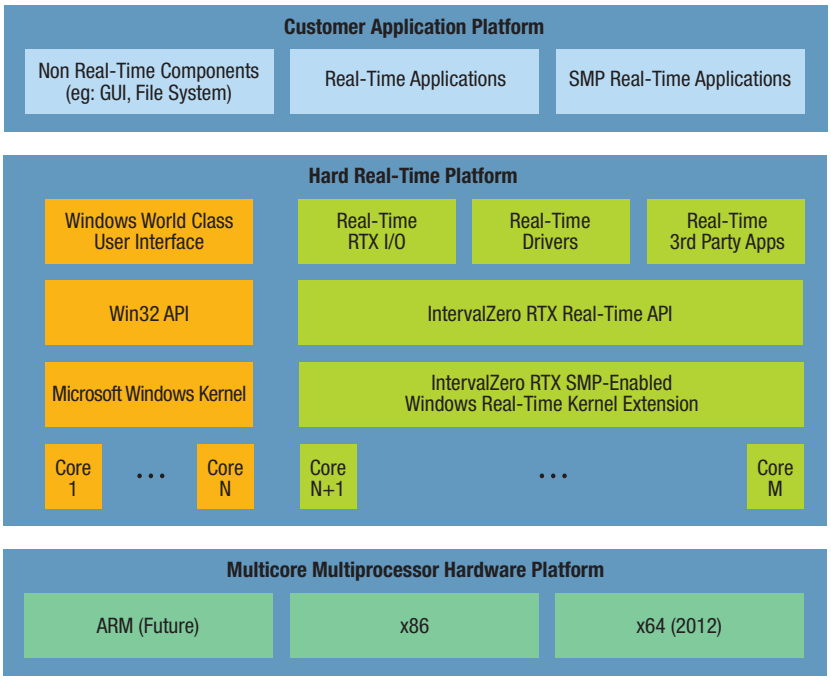

FIGURE 2

A hard real-time platform such as RTX integrates an operating system like Windows with symmetrical multiple processing (SMP) over a number of cores on a GPP so that functions that required specialized hardware components can now be done in software.

Hard Real-Time Platform Enabled by Multicore GPP

A Hard Real-Time Platform is simply any development platform that allows the customer to build a standardized real-time system with standardized hardware. The vision is that the product can be designed entirely in software without respect to the target hardware and then deployed to the appropriately sized target PC hardware, depending on the target application.

For example, a mixing board for the music industry might have 4, 16 or 96 channels. Traditionally, a mixing board would require a specific hardware design and software application to target each product. Today, world-class organizations in the music industry are designing a single, scalable mix engine that can be deployed to a two-core PC for the 4-channel product, a four-core PC for a 16-channel board, and a 24-core PC for the 96-channel mix surface. This approach is available today and has supplanted an SoC. Similar solutions are in other markets, too.

In today's hyper-competitive and wired global economy, standard pre-integrated application development, software-only platforms are increasingly in demand because they allow companies to focus solely on delivering value-added applications. The Android and iPhone platforms are excellent examples. With the supporting infrastructure already in place, smartphone and PDA application developers are using those platforms to build and deploy applications with confidence.

No development resources or time are spent building custom DSP or FPGA hardware or operating systems. The platforms provide all the necessary computing resources, integrate the components and ensure both interoperability and stability. This frees application developers to concentrate on capturing their intellectual property in the application.

Using SoCs to lower the bill of material costs has been the traditional way for manufacturers to increase profits. However, by moving to multicore GPPs, manufacturers are able to fully leverage standardized software to deliver much richer, more desirable products than either SoC

or ASICs can, with only a slightly higher BOM. The GPP-powered smartphones can command a higher average selling price and generate significantly more profit. A win-win: market demands met, profitability increased.

The primary GPP in deployment today is the x86, but it is exciting to think about the possibilities of Windows 8 on ARM if Microsoft creates a standard multicore ARM environment for the Hard Real-Time Platform. An ASP based on an SoC model will have a difficult time matching the relentless reduction in cost for an ARM-based platform that runs only software.

Best Practices

We are witnessing a trend in which devices that were once conceived as ASICs or SoCs are able to run standardized software on GPPs using a Hard Real-Time Platform, a sample architecture of which is shown in Figure 2. In researching alternative approaches, it is advisable to consider five key features that will decrease the time-to-market and reduce overall design cost.

Leveraging a *single integrated development environment* will maximize productivity. Additionally, the use of system-wide debugging directly improves quality and an immersive user interface drives product differentiation. Proprietary hardware limits continuous performance gains. Taking advantage of *COTS multicore PC boards* will drive down computing costs by 25 to 50 percent, and by staying on multicore x86, it is possible to double application performance every 18 months.

It is clear that scalability is vital for expanding product offerings and opening new markets. With a *parameter-driven real-time "engine,"* system builders can move from a 2-core system to a 24-core system without redesign, quickly scaling with x86 to create their next-generation products. *Hardware independence* will also increase the ability to create new products and revenue by re-hosting from Atom to x86, to x64 to ARM. And finally the use of *pre-tested components, drivers and applications* will cut development time and allow engineering teams to focus on adding Intellectual Property.

These five differentiating components comprise a Hard Real-Time Platform, which can be used for developing complex hard real-time systems that run on COTS hardware. An example of this would be IntervalZero's Hard Real-Time Platform, which is based on its SMP-enabled RTX software. RTX extends the Windows operating system to deliver hard real-time and runs on multicore x86.

The Hard Real-Time Platform approach will likely be very interesting to companies that have traditionally relied on SoCs and DSPs to satisfy their hard real-time requirements for systems with a sophisticated user experience. Today these companies are facing significant challenges: the relentless pressure to get to market faster at a lower cost; increased competition from smaller-form-factor systems with higher quality; and the need to open new markets to continue to grow.

It is hardly surprising that CEOs, CFOs and product management are increasingly recognizing the value of replacing their proprietary real-time hardware with standardized software components that can run in real-time as a component of an off-the-shelf PC.

Unlike traditional architectures that rely on two development environments and two subsystems to deliver both the hard real-time and the sophisticated user interface, a Real-Time Platform such as IntervalZero's supports the creation of a single, integrated system that executes a Windows-based HMI/GUI and the real-time system in parallel with SMP architecture on a single PC.

At the outset of this article, we talked about the need for the SoC and ASIC systems builders to be mindful of the gap that is growing between design cycle time (and costs) and the shortening product life cycles. The more the industry can evolve its ASP approach the better off it will be. The winning companies will be those that think big, start small and scale fast. Standardized software on a Hard Real-Time Platform that runs on multicore x86 cores offers tremendous advantages.

IntervalZero
Waltham, MA.
(781) 996-4481.
[www.intervalzero.com].