



## CommsDesign

# IP model shift: from blocks to app-specific subsystems

[Ron Wilson](#)

Oct 04, 2004 (9:00 AM)

URL: <http://www.commsdesign.com/showArticle.jhtml?articleID=49400219>

San Mateo, Calif. — The embedding of digital media capability into just about everything consumers touch is changing the way system designers use intellectual property. In the process, it could be turning the concept of IP reusability on its head.

Whereas once every attempt was made to cover as wide a range of applications as possible with a single IP block that was both function-specific and application-independent, today's IP is increasingly tied to specific apps, and in some cases even to specific application software. The reasons for the change lie hidden in the development cost structures of system OEMs and in the changing face of design teams, as designers with scant signal-processing experience confront the need to include media codecs and processors in a vastly expanded range of platforms.

"We are watching an evolution in IP," said Derek Meyer, recently appointed vice president of marketing at ARC International. "IP is increasingly no longer just building blocks; it comes in assembled and verified subsystems."

Design-outsourcing firms have spotted the trend. "We have seen our practice evolving from the role of system optimizers to that of enablers and now to co-developers," said Tapan Joshi, vice president of marketing at outsource vendor eInfochips Ltd. (Ahmedabad, India). "Often, we are bringing not just signal-processing expertise but reference designs, middleware, protocol stacks and RTOS configurations to the table."

Today, work with a client often begins during the product definition phase and continues through design, verification and even into production, said eInfochips CTO Upendra Patel. "We are often working with teams that have strong application domain knowledge but not signal-processing expertise," he said. "For them, this level of support is necessary."

Recent announcements make it clear that the need for such support is influencing the strategies of core IP vendors.

ARC, for one, last week announced ARCSound, a complete IP package for audio processing that's based on an ARC 600 core but includes specific core extensions, memory architecture, development tools, middleware and application code for standard audio codecs, reference implementations, relevant peripheral cores, tools and an optimized real-time OS.

ARC is in the process of obtaining compliance certifications on the complete package and is in a foundry relationship to have a 180-nanometer silicon implementation available for evaluation,

supplementing the more traditional, FPGA evaluation kit. The package will be available only as a turnkey audio processor core.

ARM Ltd. and MIPS Technologies Inc., meanwhile, last week announced their own moves toward application-specific support. MIPS introduced a single-instruction, multiple-data (SIMD) instruction-set extension aimed at audio and voice-over-Internet Protocol (VoIP). ARM rolled Neon — a SIMD architectural extension, architecture-specific compiler and set of open application programming interfaces directed specifically at streaming digital content. Both companies said broader architectural support is in the works.

"We are indeed seeing IP becoming more application-specific," affirmed Will Strauss, president of analysis firm Forward Concepts (Tempe, Ariz.). "And the reason is signal processing.

"As systems OEMs move digital media into their products, they are encountering an entirely unfamiliar technology. Until very recently, only a select group of EEs and applied-math majors learned anything about numerical algorithms in school. You could get a PhD in computer science without ever studying them.

"So the generation of microprocessor-oriented EEs and C programmers who inhabit most design teams don't have the background for media signal processing. In the face of all these media standards, they expect to have all the algorithms — verified and certified — delivered to them along with the RTL. The leading DSP vendors, like Texas Instruments, recognized this and have been supplying that level of applications support, and now the IP vendors who want to serve this market are having to follow. You can't just ship these guys a RISC core and a C compiler."

### **All media, everywhere**

Strategists in the consumer and communications industries think digital media will permeate almost every aspect of electronics in coming years. One of the most radical scenarios for this view is presented by Philips Semiconductors executive vice president for technology and strategy Theo Claasen. Claasen sees the consumers and knowledge workers of the near future living enveloped in media: messages, both personal and commercial, electronic transactions, audio, still images and video.

"Devices that touch the individual will be multifunction," Claasen declared. "Think of three settings: a leisure-time setting in your living room, a mobile setting when you're out and about or in your office, and an automobile setting. We believe people will want access to all of their media in each of those settings, with the user interface and the presentation always adjusted [for] the circumstances.

"For instance, in your living room you might want to watch a movie on a home theater system with full high-fidelity surround sound. But sitting in your car waiting to pick up the kids, you might want to see the same movie on a small screen. Or you may want to preview the movie on your cell phone to see if it's something you want to rent that evening."

Claasen called such a scenario feasible if three conditions were met: The experience was compelling, the user interface was easy and the price was affordable. "If we can achieve this, the semiconductor content of consumer and communications devices will increase so rapidly that the consumer semiconductor segment will grow faster than the electronics industry for perhaps five to seven years," he said.

### **All media, inside**

But such a future means that media codecs and processors will be going into systems where they have never been before — and integrated by design teams with little or no prior digital media-processing experience. That is leading IP vendors into a radical change in business models.

Reusability has been one of the greatest goals of IP design. In practice, that has meant designing an IP core so that it cleanly performs a range of highly generic functions. The more purely functional the design, the more likely that the core could be used by a number of design teams without alteration to the RTL or, more important, the verification suite.

But for some vendors, the new goal is to provide a complete subsystem, all the way from RTL to application software, that is specific to a particular application and preverified with particular data sets, and that requires little or no application expertise on the part of the user.

"We are seeing the customer base expanding to include customers with no audio integration experience," Meyer said. "Here is a cell phone or digital camera design team suddenly having to produce a standards-compliant audio subsystem, when they have never done anything like that before. They are behind the eight ball on time-to-market."

"There's no question that as these media proliferate, the number of developers working with streaming content is likely to increase rapidly," said Kerry McGuire, product-marketing manager at ARM. "We need to make provision for development teams that don't have resources for hand coding a set of audio or image-processing algorithms for a piece of custom hardware."

A second factor is total development cost. IP has always, at the end of the day, been a tool for reducing development time and cost. When IP meant generic functional blocks, the job of reducing time and cost was limited to making the RTL easy to integrate and verify. But in media markets, development can mean not just silicon design and verification but also such system-level issues as middleware and application development, certification by independent agencies like Dolby, or even presenting the finished system to a set of golden eyes or ears to receive a judgment. IP vendors are being pushed to help with these issues.

### **Rollout particulars**

Both the MIPS and ARM instruction-set extensions provide a SIMD signal-processing environment by partitioning the 32-bit or 64-bit execution pipeline into 8-, 16- or 32-bit slices and executing identical operations in parallel in each slice. Thus, by the addition of a relatively small amount of control logic, the machines become parallel-processing units. The same technique has been used for years in the Intel Architecture MMX extensions and a number of other architectures.

While the ARM and MIPS extensions are just architectural documents — not implementations in silicon or even IP at this point — both vendors say they have plans for extensive application-specific support.

The ARM extensions will be implemented on "a next-generation ARM architecture," McGuire said. It will include a compiler that optimizes code for the SIMD instructions and specific algorithms accessible through a set of OpenMAX APIs. "We intend to support handcrafting of SIMD code for design teams that prefer to work that way, but also precoded libraries of known hot spots and a compiler that is capable of utilizing the SIMD instructions very efficiently," McGuire said.

MIPS intends to provide not just tools and libraries but also applications code and, eventually, certification for audio and VoIP apps. The company is also working to add the necessary new contexts to its embedded Linux to deliver OS support.


Both companies' extensions fit in between the capabilities of their basic CPU cores and media-processing capabilities already available through add-on accelerators. ARM offers the Optimo capability, which derives an accelerator directly from the application code and attaches it to the AXI bus. MIPS has for some time offered a heavier-weight SIMD capability in its MDMX coprocessor.

Both companies said that the lighter-weight extensions embedded in the CPU execution pipelines were tailored to the specific needs of the audio, VoIP and, in ARM's case, image-processing applications for which they were intended.

The narrow focus on specific applications in all these designs permitted significant acceleration to be achieved with minimal additional silicon area. ARC's move involved the addition of multiply-accumulate hardware and a somewhat more aggressive local memory configuration. The ARM and MIPS extensions complicate the pipeline but provide SIMD capability rather than single-execution throughput.

Since both the ARM and MIPS designs are at the instruction-set level at this point, the implications for register set expansion and changes to local memory architectures are not entirely clear. ARM did say that the Neon architecture would have its own private registers.

But ARM Austin design center CTO Gerard Williams said that in the applications ARM had analyzed so far, classical data cache memory organization had proved sufficient to achieve the necessary throughput. Williams added that there are probably application situations in which more-specialized local memory would be necessary.



[Copyright © 2003 CMP Media, LLC](#) | [Privacy Statement](#)