Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

ARC Encodes Digital Video

New Video Subsystems Exploit VRaptor Media Architecture

By Tom R. Halfhill {10/15/07-01}

It seems that no electronic technology fully arrives until it's fully portable. Radio was popular since the 1920s but didn't become ubiquitous until transistorized receivers appeared in the 1950s. Telephony was born in the 1870s, but it didn't become universal until wireless phones transcended landlines in the 1990s. Phonographs begat boom boxes, Walkmans, and then iPods; mainframes eventually led to laptops. Now it is television's turn.

Sure, there have been portable TVs since the 1950s, if anything with a handle bolted on top can be called portable. But the confluence of color LCD screens, highly integrated SoCs, cheaper memory, and better batteries is putting digital video into hands, pockets, and purses. Some mobile devices can merely play video, but a growing number can also record full-motion video. However, digital-video recording is a greater challenge than playback. Video-compression standards tend to be asymmetrical—requiring more processing power for compression than for decompression—unless they compromise video quality.

To capitalize on this growing market, ARC International has introduced five new members of the ARC Video Subsystem, a family of digital-video encoders and decoders. ARC licenses these subsystems to customers as soft intellectual property (IP) for integration in SoCs. The ARC Video Subsystem builds on the ARC VRaptor Media Architecture introduced in 2006. VRaptor, in turn, is based on an ARC 700 32-bit embedded-processor core augmented with instruction-set extensions, SIMD media processors, communication channels, special acceleration logic, and optimized software codecs for popular audio/video standards. Until now, ARC's preconfigured subsystems could handle video decoding but not the more challenging task of encoding. Even so, ARC says its earlier audio/video IP already reaps more than half the company's revenue—a surprising statistic, considering that the company's main business has always been licensing general-purpose processor cores.

Licensees can integrate any member of the ARC Video Subsystem into a custom chip designed for mobile digital-video products. As synthesizable IP, the subsystems are compatible with virtually any IC fabrication process at any foundry. In addition to the decode-only subsystem introduced last year (now called the AV 401), ARC is rolling out five new subsystems for both video encoding and decoding: the AV 402V, AV 404V, AV 406V, AV 407V, and AV 417V. All are low-power solutions for various segments of the mobile market, including cellphones, personal media players, camcorders, toys, and mobile TVs. They are also suitable for some tethered products, such as TV set-top boxes, webcams, digital-video recorders, surveillance cameras, and automotive entertainment systems.

ARC provides several prewritten audio and video codecs for the ARC Video Subsystem, which is crucial, because such codecs are difficult to write. Video encoders include H.264 Baseline Profile, H.263, MPEG-4 Simple Profile, and MPEG-4 Advanced Simple Profile. Video decoders include H.264 Baseline Profile, VC-1, MPEG-4 Simple Profile, MPEG-4 Advanced Simple Profile, MPEG-2, and Motion-JPEG (M-JPEG). There's also a still-image encoder for JPEG and still-image decoders for JPEG, GIF, PNG, and TIFF. Audio codecs include MP3, Windows Media, AAC, aacPlus v2, AC3, and Dolby Digital Plus.

However, ARC's prewritten video codecs currently don't support resolutions beyond D1 or Standard Definition

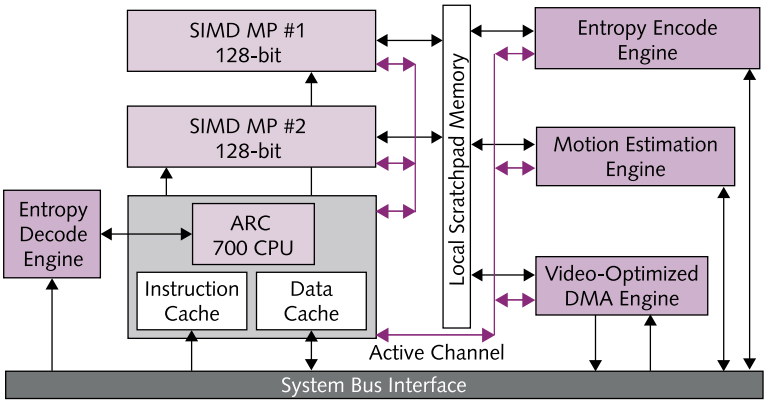

Figure 1. ARC Video 417V block diagram. The AV 417V is the most powerful member of the ARC Video Subystem family. Essentially, other members are subsets of this complex IP block. The ARC 700 processor core gets a great deal of help from SIMD engines and dedicated accelerators to compress and decompress digital video. (Purple lines indicate active-channel interconnects among components.)

(SD) in NTSC or PAL format (maximum 720 x 576 pixels, interlaced). They don't encode or decode High Definition (HD) video. Nor do they decode H.264 video using the higher-quality High Profile or the Main Profile—they are limited to Baseline Profile, which is more error tolerant. Those limitations exclude ARC's preconfigured IP from the rapidly growing HDTV market and from applications requiring H.264 Main Profile in SD resolution. Nevertheless, for the codecs they support, ARC's subsystems perform high-quality encoding. The AV 417V encodes video at higher quality (a higher peak signal-to-noise ratio) than the H.264/AVC standard reference encoder.

ARC says the most powerful implementation of the ARC Video Subsystem (the AV 417V) could handle H.264 Main Profile and HD resolution if someone further optimizes the codecs and cranks up the clock frequency. It's a future possibility. For now, HD is not imperative for the mostly mobile products that ARC is targeting. But the market is moving quickly toward HD, and ARC recently acquired a company that prepares for future developments. (See the sidebar, "Future Directions: Mobile HD Video.")

Heavy-Duty Extensions and Accelerators

All digital video is compressed to save bandwidth during transmission and to save memory during storage. This is true whether the transmission medium is a wireless broadcast network (terrestrial or satellite), a landline cable-TV network (copper or fiber), a conventional telephone network (landline or cellular), or a point-to-point packet network such as the Internet (wired or wireless). Transmitting full-motion video without compression simply requires too much bandwidth, and bandwidth remains a scarcer resource than the processing power required for compression and decompression. Historically, the cost of processing power falls much faster than the cost of network bandwidth.

Even when network bandwidth is not a factor—such as when playing prerecorded video on a home TV—digital video is always compressed to fit the playback media. Video CDs, particularly popular in Asia, can store a full-length feature film on a single CD, but their MPEG-1 compression is inferior to the MPEG-2 compression on higher-capacity DVDs. Even DVDs are inadequate for storing movies recorded in HD, so the industry is rolling out the larger-capacity HD-DVD and Blu-ray disc formats. Another reason for compression is to reduce the amount of temporary storage required during digital-video processing. Playback systems must temporarily buffer several frames of video when applying various steps of the compression or decompression algorithms.

It's important to note that video quality varies among compression algorithms, even when they follow a particular compression standard. The standards leave much room for differentiation, as long as the algorithms encode the video as a defined stream that any standard decoder can interpret. Therefore, saying that a particular encoder supports H.264 Baseline Profile (for example) doesn't necessarily indicate whether its compression quality is good or bad. Some H.264 Baseline encoders are better than others. ARC says its programmers spent many hours optimizing the codecs for higher quality, which is the reason its H.264 Baseline encoder exceeds the peak signal-to-noise ratio of the H.264/AVC reference standard.

When not taking shortcuts that reduce quality, video-compression algorithms tend to be asymmetric—compression requires more processing than decompression. Asymmetric compression yields higher compression ratios. The prevailing assumption has been that playback is more common than recording and therefore should be less expensive. However, that assumption is questionable in view of current trends. Today, even some cellphones are gaining the ability to record full-motion video, and video recording is also a common feature in digital still cameras. True, many of those cameras record video using M-JPEG, a symmetrical whole-frame compression method. Nevertheless, the most popular video-compression standards tend to be asymmetric in their higher-quality modes, and the imbalance imposes an extra burden on systems designed for video encoding.

That burden is readily apparent when studying the ARC Video Subsystem. Without help from extensions and accelerators, even ARC's most powerful processor, the ARC 700, is woefully inadequate for digital video. No shame on ARC—the same is true of virtually any 32-bit general-purpose embedded processor. Fortunately, ARC's processors are designed to be highly configurable for specific applications. Even so, the breadth of ARC's VRaptor Media Architecture is stunning. ARC has added nearly 500 new instructions, about four times the number of instructions in the ARC 700's base instruction-set architecture (ISA). And even those extensions

Feature	ARC AV 417V	ARC AV 407V	ARC AV 406V	ARC AV 404V	ARC AV 402V
Hardware Features					
ARC 700 CPU	1	1	1	1	1
SIMD Media Processors	2	1	1	1	1
Entropy Encode Engine	1	1	1	1	1
Entropy Decode Engine	1	1	—	1	—
Motion Estimation Engine	1	1	1	—	—
Video DMA Engine	1	1	1	1	1
Scratchpad RAM	32K	n/a	n/a	n/a	n/a
L1 Cache (I+D)	32K + 32K	n/a	n/a	n/a	n/a
Subsystem Size (TSMC, Virage libs)	4.95mm² (90nm G) 2.35mm² (65nm GP)	n/a	n/a	n/a	n/a
Video Features					
Max Video Resolution	720 x 576 (D1 PAL)	720 x 576 (D1 PAL)	720 x 576 (D1 PAL)	720 x 576 (D1 PAL)	352 x 288 (CIF PAL)
Encoder Resolutions	D1, CIF	D1, CIF	D1, CIF	CIF	CIF
Decoder Resolutions	D1, CIF	D1, CIF	CIF	D1, CIF	CIF
Applications					
Target Specifications	Higher-quality D1 encode, decode	Smaller-die D1 encode, decode	Low-power D1 encode	Low-power CIF encode	Low-power D1 decode
Sample Applications	High-end handsets, PMP, STB, DVR, camcorders	Midrange handsets, PMP, STB, DVR, mobile TV	Webcams, surveillance, cars, sensors	Low-end handsets, PMP, IPTV, mobile TV	Low-res handsets, PMP, cars, toys, image sensors

Table 1. The ARC Video Subsystem is available in six configurations targeting different segments of the digital-video market; this table shows only the five new configurations capable of video encoding. The AV 417V is the most powerful member of the family and is also the largest and most power hungry. Other configurations are subsets of the AV 417V that require less silicon and power. For competitive reasons, ARC is withholding some of their specifications. (n/a: data not available.)

weren't enough to get the job done, because ARC has added some dedicated accelerators, too.

Figure 1 is a block diagram of the ARC AV 417V, the most powerful member of the ARC Video Subsystem family. At its heart is the ARC 700, the flagship of ARC's configurable processors. (See *MPR 3/8/04-01*, "ARC 700 Aims Higher," and *MPR 6/21/04-01*, "ARC 700 Secrets Revealed.") Assisting the CPU are two 128-bit SIMD media processors, which implement the ISA extensions to the ARC 700. Each media processor uses vector instructions to manipulate 8-, 16-, or 32-bit chunks of data stored in 128-bit-wide vector registers. In addition, there are four special accelerators plus a configurable amount of scratchpad memory (32KB in the high-end AV 417V).

A big advantage of building the ARC Video Subsystem on the ARC 700 is that, in some users' designs, the CPU can serve as the host processor as well as the video processor. Many other digital-video solutions need a separate host processor. The ARC 700 is a relatively powerful CPU with a memory-management unit (MMU), so it can run sophisticated operating systems, such as Linux (the 2.4 and 2.6 kernels). It can also run a smaller real-time operating system (RTOS), such as uClinux (a Linux 2.0 derivative for microcontrollers), Express Logic's ThreadX, open-source uITron (also known as µITron or Mitron), or ARC's own MQX.

Distributed DMA Keeps Data Moving

By subtracting or shrinking various components of the AV 417V, ARC has derived four less powerful configurations of the encoder: the AV 407V, AV 406V, AV 404V, and AV 402V. These derivatives are intended for applications requiring less performance, lower power consumption, and lower cost. Table 1 summarizes the features and target applications of all five configurations. (This table omits the decode-only AV 401V, introduced last year.)

The primary factors distinguishing these configurations from each other are the numbers of SIMD media processors and accelerators assisting the ARC 700. The SIMD media processors implement the VRaptor Media Architecture extensions. The accelerators are dedicated engines created specifically for the ARC Video Subsystem, not extensions to the ARC 700 architecture. They are command driven, not fully programmable, and are connected to the ARC 700 over a 32-bit internal bus. Programmers can issue commands to the engines by using assembly language or intrinsic functions in C/C++.

The video DMA engine is particularly interesting, because it operates independently of the DMA controllers in the other video engines, ARC 700 core, and SIMD media processors. Optimized for handling digital-video datatypes (two-dimensional arrays of 8-, 16-, and 32-bit integers), the video DMA engine fetches video frames from external memory and stashes them in the subsystem's scratchpad memory for processing. The memory interface is 32 bits wide and provides up to 1.2GB/s of memory bandwidth at 300MHz, although the subsystem needs only about 76MB/s for H.264 Baseline encoding. (In some designs, customers can take advantage of the bandwidth headroom by reducing the subsystem's clock frequency, thereby saving power.)

The other accelerators perform common but complex video-compression/-decompression tasks that would overwhelm the ARC 700 (or virtually any other general-purpose

Instruction Type	Instructions	Example	Function
Channel	4	VSENR Ra, Rb, u8	Send message in register Rb and command u8 to channel ID in Ra
Compare	48	VEQW Va, Vb, Vc	Compare each element in vector registers Vb and Vc for equality
Control	8	VRUN Rc	Run instructions from SIMD memory starting from location in Rc
DMA	14	VDIRUN Rb, Rc	Write DMA in reg D13 with contents of regs Rb, Rc and start DMA
Jump	9	VJP Ra	Jump to location in register Ra
Vector LD / ST	71	VLD128r Va, [Rb]	Load vector reg Va from SIMD data memory at byte addr in reg Rb
Logical	23	VAND Va, Vb, Vc	Bitwise logical AND of regs Vb, Vc and store result in reg Va
Mask	4	VGMW<.p> Ia, Vb	Compare elements in vector reg Vb with 0x0 and write 0x1 (if not equal) or 0x0 (if equal) to corresponding bit in scalar reg Ia
Arithmetic	176	VADDW Va, Vb, Vc	Add vectors in regs Vb and Vc and store result in reg Va
Data Move	22	VMVW Va, Vb, simm	Move vector reg Va to vector reg Vb with enables on 16-bit fields, else unchanged
Data Packing	27	VUPSBW Va, Vb	Convert signed bytes to words from lower part of reg Vb and write 128-bit result to vector reg Va
Data Permute	68	VEXCH4 Va, Vb	Exchange higher 64-bit element from vector reg Va with lower 64-bit element from vector reg Vb
App Specific	6	VH264F Va, Vb, Vc	H.264 filter test: vector reg Vb holds pixel values, reg Vc holds parameters, reg Va holds control information
Vector Aux	2	VRGET<.p> Va, [Ib, s12]	Read local system reg at Ib+s12 into vector reg Va
TOTAL	**482**		

Table 2. A sampling of ARC's VRaptor Media Architecture. In all, VRaptor extends the ARC 700 base architecture with nearly 500 new instructions in 14 categories, as well as adding new registers, communication channels, and accelerators. Most instructions perform the specialized tasks of manipulating audio and video data, but VRaptor also has numerous control instructions that allow media processors based on these extensions to function semiautonomously.

processor core). For instance, the motion-estimation engine analyzes the differences between blocks of pixels in successive video frames and calculates whether particular parts of the image are moving. If so, the engine calculates the motion vector of the movement. The compression algorithm conserves storage space and bandwidth by saving the motion vectors instead of the actual pixels. Later, the decompression algorithm reconstructs the pixels by applying the motion vectors to corresponding blocks of pixels in the video frames.

During compression, the entropy-encoder engine compresses data derived from the original pixels into shorter symbol codes. During decompression, the entropy-decoder engine converts the codes in the symbol table back into data that eventually gets reconstructed as pixels. Video compression is lossy, so the reconstituted pixels don't exactly match the original pixels—some quality is inevitably lost. The amount of degradation depends on many factors and varies from one system to another, even when following the same video-compression standard. To accelerate the entropy process, three versions of the ARC Video Subsystem have separate engines for entropy encoding and decoding, and each engine has its own DMA to the system interface.

ARC says it optimized the video subsystems for high-quality encoding and playback, even at the expense of some additional complexity. Simpler solutions often suffer from jerky motion or blocky compression artifacts. In particular, the ARC Video Subsystem is designed to optimize the critical tasks of motion estimation, bit-rate control, inter/intraframe comparisons, macro-block partitioning, and pixel deblocking.

VRaptor Architecture for Media Processing

ARC estimates that without help from custom extensions and accelerators, a generic RISC microprocessor would need to run at 5.0GHz to encode a video stream in MPEG-4 format or at 18.0GHz for H.264 format (in D1 resolution at 30 frames per second). Clearly, those are impractical requirements, especially for an embedded processor aimed at battery-powered mobile systems. It simply takes too many generic RISC instructions to perform the special tasks needed to compress and decompress digital video. Fortunately, ARC's processor cores are designed to be customized for special tasks. Although ARC usually leaves the customizing to licensees, the company sometimes creates its own extensions. As part of the VRaptor Media Architecture, ARC has added hundreds of new instructions, 32 new vector registers (each 128 bits wide), 16 additional scalar registers (each 16 bits wide), and other features.

The total number of new instructions depends on the way the usual variations are counted: we find 482. Most data-manipulation instructions follow the single-instruction multiple-data (SIMD) format—one instruction performs the same operation on two or more operands simultaneously. Video datatypes are usually 8-, 16-, or 32-bit integers, so a 128-bit vector register can hold 4 to 16 operands. In addition to SIMD instructions, which are commonplace in media extensions, VRaptor adds several new shift instructions and flow-control instructions. Furthermore, many VRaptor instructions support predicated execution. Table 2 shows a sample instruction from each category in the VRaptor Media Architecture.

VRaptor is almost self-sufficient. That is, the media-processing units that implement these extensions are loosely coupled to the ARC 700 processor core and can operate in a semiautonomous fashion. This division of labor allows low-level media processing to continue with less attention from the CPU core, freeing the CPU for other tasks. Vectorized predicated execution is a key part of this capability, because it

offloads some flow control from the CPU, eliminates some branches, and allows the low-level media-processing code to make multiple IF-THEN decisions simultaneously.

Predicated execution is a familiar concept implemented in many CPU architectures, including those from ARC, ARM, Intel (Itanium), and Tensilica. A predicated instruction may or may not execute, depending on the state of a status register or some other flag. Among other things, predication reduces the need for compare-and-branch instruction sequences.

Vectorized predicated execution expands on that concept. It allows vector instructions to manipulate multiple operands in a vector register in different ways, depending on predicate flags in another register. One VRaptor instruction can simultaneously perform different operations on as many as eight operands. Without vectorized predicated execution, a program would have to execute a long series of compare-and-branch instructions to test the condition of individual flags, then fork to additional code that manipulates each operand separately. ARC didn't invent vectorized predication, but VRaptor makes effective use of it.

Active Channels Connect the Blocks

With its ARC 700 CPU core, SIMD media processors, video engines, and local memories, any member of the ARC Video Subsystem is almost an SoC in itself, although each is intended to be merely a component of a larger system on chip. As with all such systems, communications among the various logic blocks is critically important. To connect the CPU and other subsystem blocks together, ARC has created point-to-point "active channels" that carry commands as well as data. The commands remotely invoke functions in the blocks. (Note that active channels are internal to the ARC Video Subsystem; externally, the subsystem may connect to other components over any chip-level interconnect the designer prefers.)

In effect, the active commands are remote procedure calls. Each logic block (CPU, SIMD media processor, or accelerator) can send messages to other blocks. First, the sender opens a channel to a receiver by issuing a command. (Each block has a multiplexed input interface so it can receive new messages without closing a previously opened channel to another block.) When the channel is open, the sender can remotely invoke operations in the receiver. VRaptor instructions support all these operations. Input commands to the channels resemble register-move instructions, and the output resembles a register writeback.

A command requires six CPU clock cycles to make a round trip between blocks—in other words, to invoke an operation and return the result. However, channel communications are pipelined, so the subsystem can process messages at the ideal rate of one message per clock cycle. Inevitably, of course, some channel operations will stall. This may happen when the read buffers are empty or the write buffers are full. The ARC 700 discovers these dependencies halfway through its seven-stage pipeline and can put the channels into a lower-power mode until the dependency resolves. Wakeup is instantaneous.

Because the ARC Video Subsystem has its own local memory and DMA, the media processors and accelerators can execute many instructions independently of the CPU. ARC calls it "fire and forget"—the CPU stuffs VRaptor instructions into a special queue, from which the video subsystem fetches and executes them. In fact, the CPU can fetch instructions and store them in local memory, then step aside while the accelerators execute them without bothering the CPU at all. Another feature is similar to recording and playing a macro. A whole series of cached instructions, including branches and loops, can execute on the SIMD processors when the CPU issues a single VRUN instruction. To a large degree, these relationships decouple instruction issue and execution between the CPU and the rest of the subsystem, although the relationship is not fully parallel. The CPU is still responsible for fetching, decoding, and dispatching VRaptor instructions, and those instructions initially share the CPU's pipeline with ordinary instructions.

ARC provides software codecs for popular audio/video standards and software-development tools for programmers who want to write their own code. Although the tools fall short of a true vectorizing compiler that can automatically extract parallelism from high-level C/C++ code, they do have features that promote parallel programming. For instance, color-coded activity logs show the utilization of various logic blocks, memories, and queues, so programmers can optimize the utilization of those resources. Likewise, code profilers identify hot spots in a program where code optimization or parallelism would improve throughput. Unfortunately, ARC's C/C++ compiler doesn't support VRaptor extensions yet, so programmers must write low-level code for the SIMD processors in assembly language.

Of course, developers who find ARC's prewritten codecs sufficient for their purposes won't have to write low-level code. Another advantage of using ARC's codecs is that some of them implement "dynamic adaptive encoding," a clever technique for which ARC is seeking a patent. This technique looks ideal for mobile video systems that must compromise their performance because of constraints on processing power, network bandwidth, bus bandwidth, memory latency, or power consumption. It's also a technique that's easier to implement with a programmable video encoder like ARC's.

In concept, dynamic adaptive encoding is similar to variable bit-rate encoding, except that it varies compression quality according to the system's performance constraints as well as the texture of the media content. Also, it can vary compression quality by turning particular features (such as pixel deblocking) on or off, not just by varying the bit rate. For example, if a video stream must be compressed without exceeding a certain power limit, the codec can automatically and dynamically adjust the compression parameters and processor clock frequency to achieve that goal. Table 3 shows performance without using this technique—in other words,

Codec	Video Resolution	Video Frame Rate	Video Bit Rate	ARC 700 Clock Rate
H.264 Baseline Profile (Encode)	D1	30 fps	Up to 10Mb/s	200MHz
H.264 Baseline Profile (Decode)	D1	30 fps	1.5Mb/s	160MHz
MPEG-4 SP/ASP (Encode)	D1	30 fps	Up to 8 Mb/s	166MHz
MPEG-4 ASP (Decode)	D1	30 fps	2.0Mb/s	135MHz

Table 3. ARC Video Subsystem performance. Clock frequencies in this table assume the AV 417V is synthesized for TSMC's 65nm GPlus fabrication process. In such a process, the subsystem (as illustrated in Figure 1) would occupy 2.35mm^2 of silicon, including system and memory interfaces. Lower clock speeds are possible by using ARC's dynamic adaptive encoding, which automatically varies the compression quality to meet certain performance constraints.

worst-case (highest-quality) performance. Dynamic adaptive encoding allows real-time compression at much slower clock frequencies if poorer quality is acceptable.

Surging Competition for Digital Video

Digital video is hot, so ARC is entering a crowded marketplace. For those who don't want to develop their own chips, numerous standard parts are available off the shelf. (See *MPR 9/17/07-02*, "Video ICs: A Lively Market.") For those who want to make a chip but would rather avoid some logic synthesis and verification, there are drop-in hard macros. For those who prefer synthesizable IP but don't need full programmability, there are hard-wired accelerators provided in RTL. And for those who want to license a programmable video engine as soft IP, there are alternatives to the ARC Video Subsystem.

ARC's main competitors for licensable processor cores are ARM, MIPS Technologies, and Tensilica. None offers a video subsystem quite like ARC's. ARM does offer a configurable coprocessor engine called OptimoDE that developers can shape into a video engine, but ARM's only preconfigured OptimoDE product is an audio engine. (See *MPR 6/7/04-01*, "ARM's Configurable OptimoDE.") Currently, ARM and MIPS prefer to leave the implementation of video subsystems to their licensing partners.

Tensilica, like ARC, sells user-configurable embedded-processor cores and preconfigured video subsystems based on those cores. The video subsystems are the Diamond Standard VDO Video Engines. Tensilica offers four VDO Engines, but only two are capable of encoding digital video, and they are limited to MPEG-4 Part 2 encoding—they can't encode H.264. However, the VDO Engines can decode H.264 Main Profile streams, as well as other popular standards.

Keeping that limitation in mind, Tensilica's VDO Engines deliver performance similar to that of the ARC Video Subsystem. Tensilica's best VDO Engine is the Diamond 388VDO, which can encode video in MPEG-4 Advanced Simple Profile (ASP) in D1 resolution at 30 frames per second (fps). Its estimated maximum clock speed while performing that task is 185MHz. As Table 3 shows, ARC's AV 417V subsystem can encode an 8Mb/s MPEG-4 ASP stream at 166MHz—slightly better performance. Both companies say their video encoders have enough I/O bandwidth and internal memory to work with affordable external memory, such as SDRAM with bus latencies of 60 cycles or more. However, ARC's subsystem requires less memory bandwidth to encode H.264 Baseline video (76MB/s) than Tensilica's engine needs to encode MPEG-4 ASP video (148MB/s).

Tensilica says the 388VDO Engine occupies 6.6mm^2 of silicon, including memories, when synthesized for speed in TSMC's 90nm-G process. It's 33% larger than the 4.95mm^2 required for ARC's subsystem in the same fabrication process, and ARC didn't need to synthesize for speed. (ARC's size estimate assumes a synthesis target of 200MHz. When synthesized for higher frequencies, it can perform additional tasks, such as dual-channel video processing or heavier host processing.) One reason for the 388VDO Engine's larger area is that it has more scratchpad memory.

Unfortunately, neither company has published power-consumption estimates for video encoding, which would be illuminating. Both companies offer good code-profiling tools, but Tensilica says its C compiler doesn't require using assembly language to perform low-level video operations, although it does rely on intrinsic functions.

H.264 (also known as MPEG-4 Part 10) is regarded as the next-generation digital-video standard, because it produces higher quality at lower bit rates than MPEG-4 does. Tensilica's Diamond VDO Engines don't currently support H.264 video encoding. Tensilica says this is not a technical limitation—the software simply hasn't been written yet. All five new members of the ARC Video Subsystem do support H.264 video encoding, which gives them an edge over Tensilica in recording applications. On the other hand, Tensilica has an edge on the decoding side. Tensilica provides codecs for the Diamond 385VDO and 388VDO Engines that decode H.264 Main Profile, a mode intended for professionally produced entertainment content. Currently, ARC provides codecs for H.264 Baseline Profile, not Main Profile. (Baseline can match the quality of Main, but only at a higher bit rate. Baseline is more error tolerant, which can make it more suitable for ARC's intended mobile applications.)

Of course, it's possible for anyone to design a video subsystem having virtually any capability by independently customizing the configurable processors licensed separately by ARC and Tensilica (and MIPS). The drawback is that starting from scratch requires much more design effort than licensing a preconfigured subsystem. One glance at the complexity of these subsystems will dissuade most customers from rolling their own solution. Fortunately, preconfigured video subsystems are available from other sources as well. Among these alternatives are DSP-based subsystems and hard-wired accelerators. (See the sidebar, "Hard-Wired Video Accelerators for ASICs.")

Future Directions: Mobile HD Video

It seems odd to expect High-Definition (HD) video to appear in mobile consumer devices when most people don't even have HDTV in their living rooms yet. And why would anyone need HD resolution on a small screen? Nevertheless, the trend toward mobile HD video is clear.

For a harbinger of things to come, look no further than the Texas Instruments OMAP3-series cellphone processors. Based on the powerful ARM Cortex-A8 superscalar processor core, the OMAP3430 supports HD-DVD playback, XGA resolution (1,024 x 768 pixels), 24-bit color, and 12Mpixel still imaging. Those capabilities may seem like overkill for the tiny screen of a cellphone, but they might be perfect for portable HD-DVD players, vehicle entertainment systems, camcorders, and digicams. (See *MPR 4/24/06-01*, "OMAP3 Sets Specs for Cellphones.")

ARC is obviously anticipating this future. Although the ARC Video Subsystem is currently limited to SD resolution, last month ARC acquired a Russian software-development company that specializes in multimedia middleware, including codecs for HD video. The Russian company is Alarity, a 40-person shop in St. Petersburg. ARC is paying about $4 million up front and another $1.8 million if Alarity meets certain goals over the next two and a half years.

Alarity's engineers have developed software for multiple HD video standards, high-fidelity audio, image processing, signal processing, encryption, digital watermarking, digital rights management, and audio/video pre- and post-processing. As ARC points out, all this software is required to support HDTV, HD-DVD, Blu-ray DVD, and next-generation Internet Protocol TV (IPTV) systems.

The Russian engineers will remain in St. Petersburg, where ARC has established a multimedia development center. Alarity's cofounder and CTO, Dr. Eugene Metlitski, will head the center. Metlitski has written more than 50 books and papers, and he holds nine computer-science patents. The St. Petersburg center joins a network of ARC development centers in Silicon Valley (San Jose) and England (Cambridge and St. Albans).

CEVA MM2000 Has Proven Performance

One ARC competitor, better known for its DSP cores, is CEVA. CEVA's Mobile-Media2000 (MM2000) is a formidable video subsystem that has been available for about two years. Like ARC's best subsystems, the MM2000 is capable of encoding and decoding H.264 video at 30 fps in D1 resolution, and it's a low-power soft-IP core suitable for the same sorts of mobile applications. Unlike the new ARC subsystems, the MM2000 is based on a DSP architecture, not a RISC architecture, and it's been around long enough to be proven in silicon. CEVA says more than 20 companies worldwide have licensed the MM2000.

CEVA offers two versions: the MM2000-Pro and MM2000-Lite. Only the Pro version supports D1; the Lite version is limited to CIF resolution (352 x 288 pixels, 30 fps). Both can apply JPEG compression to still images at 12Mpixels per second, and both support a range of audio codecs, such as MP3, AAC, and WMA.

At the heart of each MM2000 is the CEVA-X1622 DSP and the CEVA-XS1200 peripheral subsystem. The X1622 is a 16-bit DSP with dual multiply-accumulate (MAC) units, multimedia SIMD instructions, and very long instruction word (VLIW) instructions. It can execute up to eight operations in parallel. Another notable feature is a 32-bit integer unit that offloads control functions and allows the DSP to run a mainstream real-time operating system (RTOS), such as Nucleus (Mentor Graphics), µC/OS-II (Micrium), or Xinu (an open-source RTOS ported to the X1622 by CEVA). This feature can eliminate the need for a separate host processor, which is unusual for a DSP-based solution. As mentioned above, the ARC 700 processor in the ARC Video Subsystem can also serve double-duty as a host processor.

The CEVA-XS1200 peripheral subsystem supports the DSP by providing an interrupt controller, a power-management unit, four time-division multiplex (TDM) ports, and a programmable DMA controller that handles all memory transactions for media data. In addition, the subsystem has bridges to AMBA AHB and APB buses. Offloading such functions as interrupt control and DMA allows the DSP to concentrate on what it does best—signal processing. As a result, the MM2000 delivers admirable performance in very little silicon.

Comparing CEVA With ARC

Because some customers have already brought MM2000-based designs to silicon, CEVA has real-world data. In one instance, a customer is using the MM2000-Pro for H.264 Baseline encoding and decoding at D1 resolution, 30 fps, with simultaneous High-Efficiency AAC audio at 48Kb/s. When decoding, the DSP runs at 214MHz and consumes 49mW; when encoding, it runs at 271MHz and consumes 64mW. Those performance measurements include audio processing, file parsing, file composing, and overhead for miscellaneous additional functions. Note the relatively small differences in clock frequency (+26%) and power (+30%) for encoding versus decoding. Usually, a rule of thumb is that H.264 encoding is twice as taxing as decoding.

ARC's stated performance, based on FPGA testing, compares favorably with CEVA's use-case data. As Table 3 shows, the ARC AV 417V can decode H.264 Baseline video (D1 resolution, 30 fps) at 160MHz and encode at 200MHz. In

Hard-Wired Video Accelerators for ASICs

Video subsystems from ARC and Tensilica are based on those companies' configurable processor cores, which have 32-bit RISC architectures. Numerous extensions—including SIMD instructions and registers—adapt these general-purpose embedded processors for digital-video processing while preserving full programmability. However, one alternative is to use hard-wired subsystems specially designed for digital video. Hard-wired subsystems frequently outperform programmable solutions on specific tasks, but they sacrifice some flexibility.

In this discussion, "hard-wired" simply means the subsystems are command-driven accelerators, not fully programmable coprocessors. They are still soft IP. Usually, third-party vendors provide these accelerators as synthesizable RTL blocks for easy ASIC integration and fabrication in any digital CMOS process. Because they are optimized for audio/video tasks, they are small and efficient, although they usually require a separate host processor for control functions. The accelerators come with prewritten software drivers for audio/video encoding and/or decoding in popular formats.

Two well-known vendors of hard-wired video accelerators are Hantro (a Finnish company recently acquired by On2 Technologies) and Imagination Technologies, based in the U.K. Lesser-known vendors are Global Digital Technologies (GDT), based in Athens, Greece, and Ittiam Systems, headquartered in Bangalore, India.

Hantro offers video encoders and decoders that customers can use separately or together, in combination with a host processor (usually ARM). When combined, the accelerators can share some scratchpad memory to reduce redundancy. The Hantro 6280 Video Encoder and Hantro 8170 Video Decoder provide capabilities similar to the ARC Video Subsystem, including MPEG-4 encoding (Simple Profile), H.264 encoding (Baseline Profile), 16Mpixel JPEG compression/decompression, MPEG-4 decoding (Simple Profile and Advanced Simple Profile), and H.264 decoding (Baseline Profile), among other standards. Hantro supports higher resolutions than ARC does—up to 720-pixel progressive-scan HD (720p) for encoding, and up to 1,080-pixel progressive-scan HD (1080p) for decoding. Hantro's RTL is configurable, so customers can scale down those capabilities to save area and power if desired.

According to Hantro, a full configuration of the 6280 Video Encoder would occupy about 2.0mm^2 of silicon when fabricated in TSMC's 90nm LP process. That area includes 40KB of memory. The 8170 Video Decoder occupies about 2.2mm^2 of silicon when fabricated in 90nm, also including 40KB of memory. When merged together, they share some memory, so total SRAM is only about 55KB, not 80KB.

Hantro estimates power consumption at 29mW to encode an H.264 Baseline Profile video stream in SD resolution at 30 fps. Power consumption increases to 76mW at 720p HD resolution. For the 8170 Video Decoder, Hantro estimates power consumption at only 9.0mW when decompressing H.264 Baseline video in SD resolution at 30 fps. At 720p resolution, power rises to 21mW. Hantro's area and power estimates are very competitive with other solutions described in this article. One reason is that Hantro's accelerators rely on the host processor for some tasks. Another reason is that they cruise at modest clock speeds. The 8170 Decoder can decompress H.264 Baseline video (SD resolution, 30 fps) while running at a mere 40MHz. The 6280 Encoder can compress the same type of video stream while running at only about 100MHz.

Imagination Technologies' latest competing products are the PowerVR VXE250 encoder, which supports video resolutions up to SD, and the PowerVR VXD370 decoder, which supports resolutions up to HD. Fully configured, the VXE250 can encode H.264 Baseline Profile or MPEG-4 video while running at only 100MHz. In TSMC's 90nm LP process, the VXE250 core measures 2.83mm^2 and consumes 42mW when encoding an H.264 stream.

The PowerVR VXD370 decoder can decompress H.264 High Profile, VC-1 Advanced Profile, MPEG-2, and MPEG-4 video at full HD resolution (1080i, 60Hz). It also decompresses M-JPEG video and JPEG still images. When fabricated in TSMC's 90nm LP process, the VXD370 core occupies 4.84mm^2. It can decode H.264 video at up to 50Mb/s in HD resolution (1,920 x 1,080 pixels, 60Hz) while running at 150MHz. Power consumption is 45mW for H.264 HD decoding and 32mW for MPEG-2 HD decoding. When clocked at 267MHz, it can decode two H.264 HD streams simultaneously. At 220MHz, it can decode one H.264 HD stream and one H.264 SD stream simultaneously.

In contrast with the programmable video subsystems described in the main part of this article, the PowerVR VXD370 decoder tolerates very slow external memory with no loss of performance. Imagination Technologies says the VXD370 works with memory latencies up to 128 clock cycles, and a configuration option increases the tolerance to 512 cycles. Although the PowerVR encoder and decoder need a separate host processor, the burden on the CPU is modest—less than 10mips—thanks to a small embedded microcontroller. Both these RTL blocks include internal scratchpad memory: 24KB for the encoder and 32KB for the decoder. The die-area specifications include the memory and the microcontroller. For more information, visit

- *www.hantro.com* and
- *www.imgtec.com*

both cases, the ARC 700 processor gets the job done at a slower clock speed than that of the CEVA DSP, which saves power. Furthermore, ARC says it can handle H.264 bit rates up to 10Mb/s, whereas the MM2000 is limited to 4Mb/s. However, CEVA's performance measurements include simultaneous AAC audio processing, file parsing, and miscellaneous overhead functions, whereas ARC's data is for video processing alone. CEVA's video-only numbers for a 2Mb/s stream are 166MHz for decoding and 226MHz for encoding, which are very close to ARC's measurements.

When fabricated in a 65nm low-power process, the MM2000-Pro occupies 2.9mm^2 of silicon, including 64KB instruction and data caches. ARC says the AV 417V, fabricated in TSMC's 65nm GP process, occupies 2.35mm^2 after place and route, including 32KB instruction and data caches. Doubling the size of ARC's caches to match CEVA's would add about 0.5mm^2 of silicon at 65nm, bringing the AV 417V near the area of the MM2000-Pro, but ARC says 32KB caches are large enough to deliver the stated performance. ARC's caches are configurable, so some customers could shrink them and use the dynamic adaptive encoding technique to achieve real-time encoding in less silicon.

Common wisdom says it's easier to program a general-purpose RISC processor like the ARC 700 than a VLIW DSP like the CEVA-X1622. However, that wisdom doesn't necessarily apply in this case. ARC's numerous VRaptor instructions and accelerators add considerable complexity to the simple RISC model of the ARC 700, and ARC's high-level compilers don't yet support VRaptor without using assembly language or intrinsic functions. Overall, the MM2000's DSP-only architecture looks simpler, and CEVA provides an optimizing C compiler for it. Both ARC and CEVA provide codecs for popular audio/video standards, so most customers should be spared the onerous tasks of pedal-to-the-metal programming. In either case, developers who can't get the codecs they need from ARC or CEVA should take a long, hard look at these architectures before plunging into a project that requires writing new codecs from scratch.

Mobile Video: A Fast-Moving Target

With most of ARC's revenue already coming from digital audio and video, it makes sense for the company to strengthen its ARC Video Subsystem product portfolio. Media processing might even become ARC's specialty, in the same way that ARM has become famous as the cellphone-processor company. The biggest hurdle for ARC, besides the surging competition, is the fast-moving nature of the digital-video market. Many potential customers will prefer to buy an off-the-shelf chip instead of spending 18–24 months and millions of dollars developing and testing their own silicon. SoC projects require big sales volumes and stable products to earn back their enormous investments.

Price & Availability

All five new members of the ARC Video Subsystem family are available for licensing now from ARC International. The new members are preconfigured for different market segments and are called the AV 417V, AV 407V, AV 406V, AV 404V, and AV 402V. They support video resolutions from CIF to D1/SD in NTSC or PAL formats. Each subsystem includes synthesizable Verilog models of the ARC 700 processor core, SIMD media processors, and dedicated accelerators. ARC doesn't publicly disclose licensing fees.

ARC provides optimized software codecs for video and audio encoding and decoding. Encoders for video and still imaging: H.264 Baseline Profile, H.263, MPEG-4 Simple Profile, MPEG-4 Advanced Simple Profile, and JPEG. Decoders for video and still imaging: H.264 Baseline Profile, VC-1, MPEG-4 Simple Profile, MPEG-4 Advanced Simple Profile, MPEG-2, Motion-JPEG, JPEG, GIF, PNG, and TIFF. Audio codecs: MP3, Windows Media, AAC, aacPlus v2, AC3, and Dolby Digital Plus. For more information, visit:

- *www.arc.com/subsystems/video.html*

Even a dream customer like Apple is a risky bet. Apple churns its iPod product line at least once a year, sometimes switching SoC suppliers along the way. Apple's new iPhone, and other smart phones, will likely turn over at the same rate. The only consolation is that Apple pushes millions of units out the door before discontinuing a product. Less successful vendors may sell only a few hundred thousand units—if that. And if those products rely on an SoC that cost $10 million to develop, where's the return on investment? These worries will deter some potential customers from licensing any soft IP, including ARC's.

In ARC's favor, a programmable solution like the ARC Video Subsystem adapts more readily to evolving digital-video standards than fixed-function solutions do. Programmable subsystems can also provide an easier upgrade path to future products by using better codecs. The importance of these advantages depends on how quickly standards evolve and products change. ARC says customers frequently ask for new codecs to support emerging standards and changing products.

ARC's recent acquisition of Alarity shows that ARC is thinking for the future. HD video is finally penetrating U.S. homes, and it seems to be going mobile pretty soon, as well. Meanwhile, ARC can focus its energies on things that never go out of style: efficient processing, lower power consumption, low cost, and easier programmability. ◇