Using DaVinci Technology for Digital Video Devices

Deepu Talla and Jeremiah Golston
Texas Instruments

It's not practical for one SoC processor to fit all equipment and meet today's aggressive performance, power, and cost targets. The DaVinci integrated portfolio of processors, software, and tools offers support for developing a broad spectrum of optimized digital video equipment.

Developing innovative and cost-efficient digital video products requires systems that encompass open and flexible system-on-chip (SoC) processors, software, and development tools. Manufacturers ship billions of DV products each year, including portable media players, digital media adaptors, IPTV set-top boxes (STBs), digital cameras, solid-state video recorders, multimedia camera phones, digital TVs, automotive infotainment, video security equipment, IP network cameras, and network video for emerging applications.

Many of these products have processor requirements similar to today's PCs, but their power and cost metrics are an order of magnitude apart.

CHANGING NEEDS

Devices have a variety of uses and price points, depending on computing strategies, level of integration, and design methodologies. For instance, the need for a high-level operating system such as Linux or Windows CE typically drives selection of a general-purpose processor. The extent of programmability, field upgradability, feature future-proofing, and need for universal video decoding or a closed system drive selection of digital signal processors (DSPs), programmable accelerators, and fixed-function hardware.

Whether the device is portable or wall-plugged or must be in always-on mode drives the selection of process technology, power domains, and clocking schemes. Target system bill-of-materials and equipment form factor drive the level of system integration, including analog, storage, connectivity, and packaging. Product life-cycle and operating conditions drive design methodologies such as process selection, testing, ability to quickly spin a follow-on device, and optimizations such as speed and power binning. And finally, time to market drives deployment collateral such as software, tools, and reference designs.

For the past decade, SoC technology has been the primary answer for the DV market. SoCs typically have multiple processors, programmable engines and accelerators, fixed-function hardware, peripheral interfaces, and analog components, all integrated in the same piece of silicon to address the need for high performance, low power, and low system cost. A majority of these SoCs are based on closed architectures, giving developers few implementation options.

However, increasingly complex DV applications require design flexibility for greater customization and advanced feature updates. Manufacturers often must use the same system platform across a range of products for specific markets, or combine different applications in the same system, such as a security camera with object recognition, or an IPTV STB with an integrated videophone. As manufacturers produce more feature-rich, multiple-application products, developers increasingly need SoC processors designed with open architectures to meet the market's versatile, rapidly changing requirements.

0018-9162/07/$25.00 © 2007 IEEE

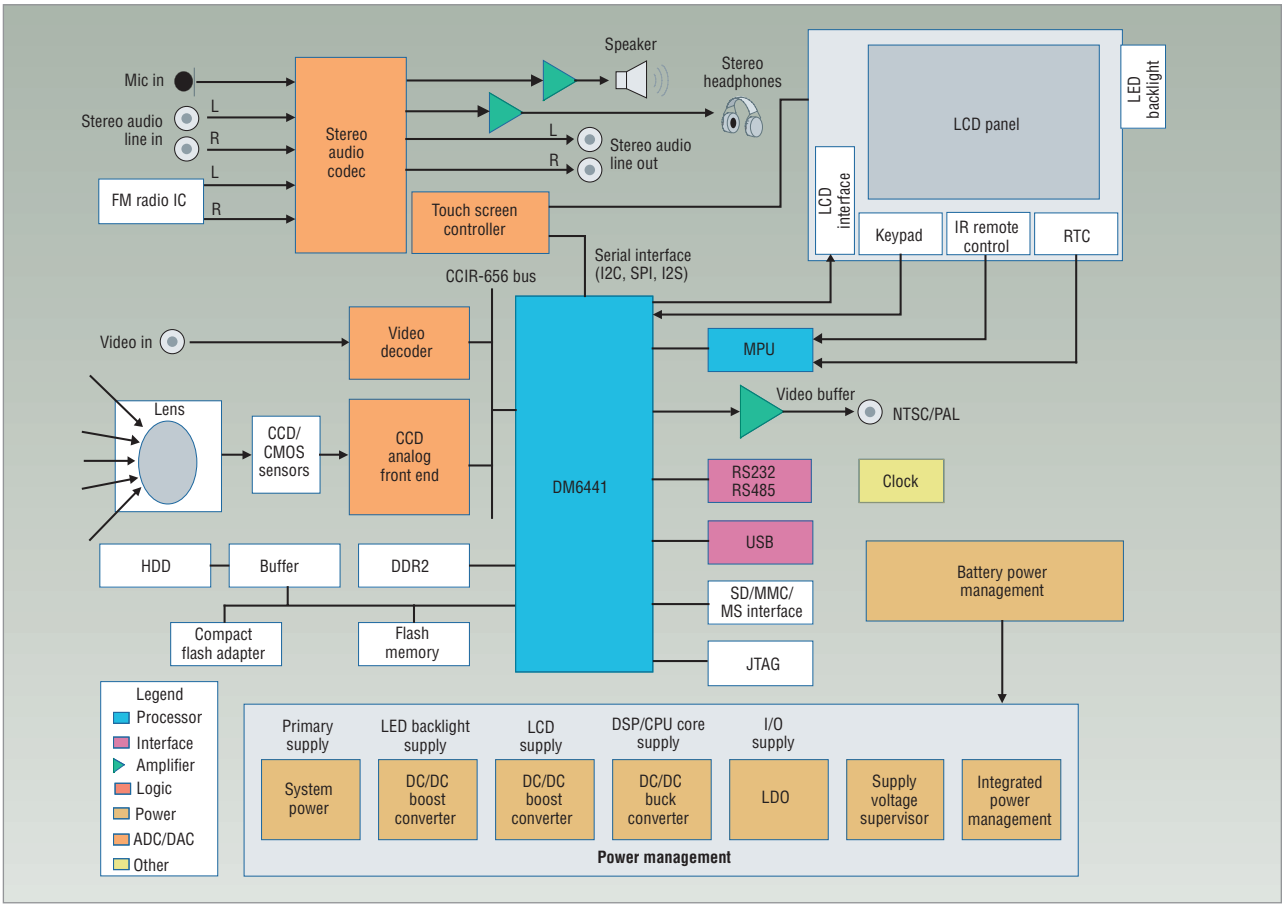

Figure 1. Portable media player system block diagram. The ability to decode multiple standards and resolutions at very low power is a key feature given the wide variety of content sources. Processor flexibility is therefore extremely important.

DAVINCI TECHNOLOGY

It's not practical for one SoC to fit all equipment and meet today's aggressive performance, power, and cost targets. The DaVinci integrated processor portfolio offers software, tools, and support for developing a broad spectrum of optimized DV equipment (http://ti.com/corp/docs/landing/davinci/index.html). DaVinci processors are optimized to match price, performance, and feature requirements for specific DV equipment.

DSP processors

DaVinci uses scalable and programmable DSP-based media processors, including ARM and DSP-centric SoCs with accelerators and peripherals. Designers can use the range of flexible IP blocks to tailor a SoC to meet the target video-equipment requirements. A general-purpose processor is needed in equipment that either has a high-level operating system or a significant amount of system and control code. DaVinci processors use a selection of ARM-based general-purpose processor solutions.

The system needs DSPs to implement flexible video- and image-processing tasks. An integrated TMS320C64x+ DSP core achieves this functionality. Furthermore, programmable and fixed-function video and imaging coprocessors are available to accelerate key algorithms. Video and imaging content-creation equipment, such as digital and video-security cameras and multimedia camera phones, require flexible image signal-processing blocks for processing raw video and image data from charge-coupled device/complementary metal-oxide semiconductor (CCD/CMOS) sensors.

Display-processing IP can range from 2D onscreen display plus a standard-definition (SD) video encoder and LCD controller to a high-definition TV (HDTV) back-end processor that contains advanced video-processing techniques for deinterlacing, color reproduction, and enhancement.

Other attributes

Sourcing and synching data to the various processing elements on the SoC requires a flexible, scalable, and powerful direct-memory access engine coupled with an efficient dynamic RAM controller merged with innovative bus architectures and topologies. Analog IP such as USB 2.0 and video and audio components are keys to low system cost and small form factors. Furthermore, peripherals such as card and memory interfaces, security

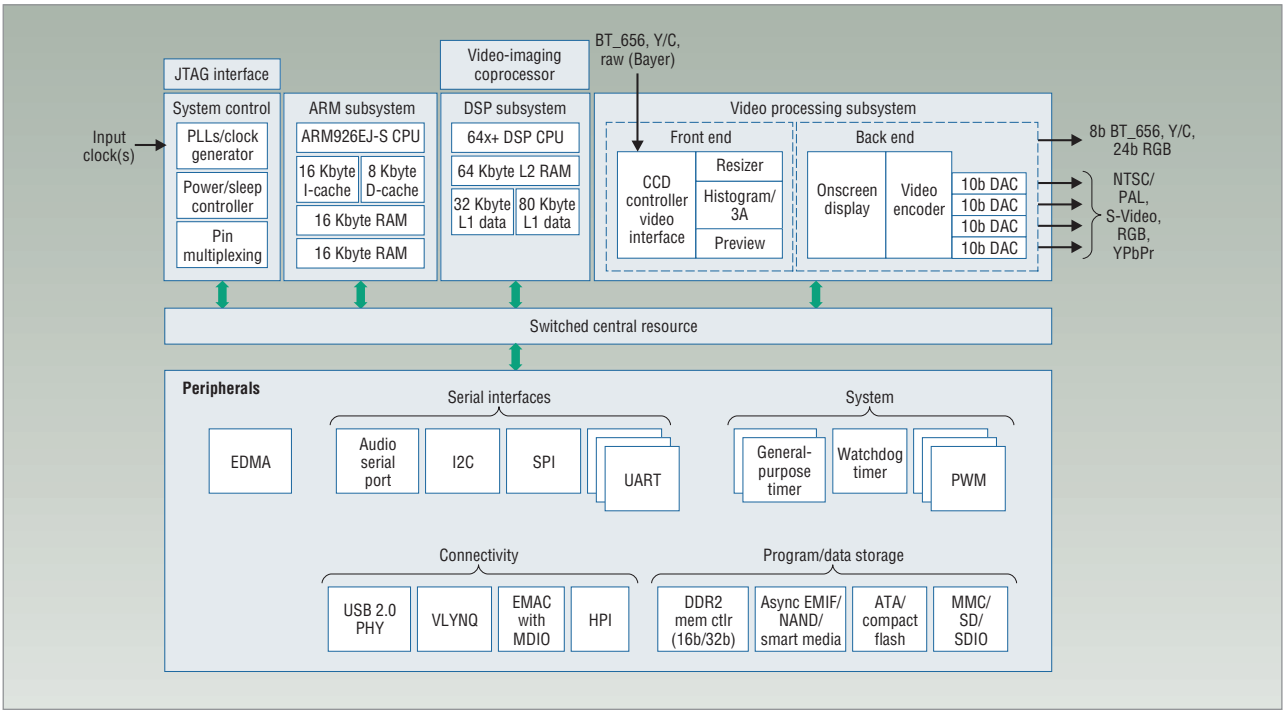

Figure 2. DM6441 block diagram. High integration of key functions, analog, and peripherals requires fewer external components, leading to a low system bill-of-materials cost and a better form factor.

IP, networking interfaces, and communication ports are integrated into the SoC.

DaVinci software includes configurable frameworks that popular operating systems present via published APIs for rapid software implementation with DaVinci processors. The software framework abstracts the processor's inner workings and dramatically reduces research and development, system architecture, implementation, and test requirements. The system also delivers standard video, imaging, audio, and speech codecs along with the processor. These include standards such as H.264, MPEG-2 and -4, Windows Media Video 9 (WMV 9), JPEG, Advanced Audio Coding, MP3, G.7xx, and other popular formats.

The DaVinci technology portfolio includes several operating system ports appropriate for different applications, including open source Linux and real-time operating systems for specific equipment. Developers need a full set of cost-appropriate tools and kits for a variety of application spaces and designs, including low-cost starter tools, complete development kits, and even reference designs to speed design and development. Developers can use the ARM/DSP integrated development environment and open source and DSP tools to program in a standard environment.

PORTABLE MEDIA PLAYERS

Portable media players can record and play back audio and video images from a hard drive, flash memory, camera, or Internet download. Figure 1 shows the system block diagram of a typical player. The ability to decode multiple standards and resolutions at very low power is a key feature given the wide variety of content sources. Processor flexibility is therefore extremely important.

A TV, DVD player, or STB can provide the recording source. The encoding video and audio standard is typically fixed to one standard. Many portable media players can capture images and video as well, requiring an integrated ISP for low system cost. The player's display ports can be a combination of SDTV, LCD, and HDTV.

A rich and user-friendly GUI is common. Recently, high-end players have additional functions such as 3D graphics, navigation capabilities, and terrestrial and mobile broadcast reception. Developers must implement all these features at a very low power and system cost to enable mass-market success.

Dual-core architecture

The TMS320DM6441 DaVinci processor was developed to address the portable media player market. As Figure 2 shows, the DM6441's dual-core architecture provides both DSP and reduced-instruction-set computer technologies, incorporating a high-performance TMS320C64x+ DSP core and an ARM926EJ-S core. The ARM926EJ-S is a 32-bit RISC processor core that performs 32-bit or 16-bit instructions and processes 32-bit, 16-bit, or 8-bit data. The ARM core incorporates data and program memory management units with table lookaside buffers and separate 16-Kbyte instruction and 8-Kbyte data caches.

Running a high-level operating system such as Linux requires an ARM9 class of processor with caches and memory-management units. The TMS320C64x+ is based on an enhanced version of the second-generation high-performance, advanced very-long-instruction-word C64x architecture offering solutions to high-performance DSP programming challenges.

The DSP core possesses the operational flexibility of high-speed controllers and the numerical capability of array processors. The C64x+ DSP core processor has 64 general-purpose registers of 32-bit word length and eight highly independent functional units—two multipliers for a 32-bit result and six arithmetic logic units. The eight functional units include instructions to accelerate the performance in video and imaging applications.

The DSP core can produce four 16-bit multiply-accumulates per cycle, or eight 8-bit multiply-accumulates per cycle. The DSP core uses a two-level cache-based architecture. The sizing of the caches and static RAM at level 1 and level 2 has been chosen to maximize a portable media player's system performance given the sharing of the dynamic RAM for several processing elements and IP blocks.

Designers can configure L1 and L2 memory as mapped memory, cache, or combinations of the two.

The L1 program cache is a 32-Kbyte direct-mapped cache, and the L1 data memory is 80 Kbytes, out of which up to 32 Kbytes can be configured as a two-way set-associative cache. The L2 memory/cache consists of a 64-Kbyte memory space that's shared between program and data space. Designers can configure L1 and L2 memory as mapped memory, cache, or combinations of the two.

The DM6441 also includes a video/imaging coprocessor to offload many common video- and image-processing tasks from the DSP core, making more DSP MIPS available for differentiated and application- or customer-specific video and imaging algorithms.

Dual-power modes

With dual-power modes, the DM6441 can run at full speed (513 MHz DSP and 256 MHz ARM) at 1.2 volts or in power-reduced mode (405 MHz DSP and 202 MHz ARM) at 1.05 volts. The dual-power modes allow reduced power consumption during operations, such as audio-only playback, when the device doesn't need the DSP's full capabilities. Additional clock-gating capabilities provide a mechanism for turning off peripherals that aren't in use.

Isolating the power domains of the DSP and ARM cores has made it possible to power down the DSP core individually during nonvideo operations. The combination of dual-power mode, clock gating, and isolated power domains can yield overall power savings up to 35 percent over devices that don't employ these schemes.

Capture and display

The DM6441 includes a video-processing subsystem with two configurable video-imaging peripherals: one video-processing front-end input used for video capture, one video-processing back-end output used for display. The front-end input performs the ISP functions and consists of a CCD/CMOS controller, a preview engine, a histogram module, a resizer, and an H3A model for autoexposure, white balance, and focus. The controller can interface to common video decoders, MOS sensors, and CCDs.

The preview engine is a real-time image-processing device that converts CMOS sensor or CCD raw imager data from an RGB Bayer pattern to YUV 4:2:2. The histogram and H3A modules provide statistical information on the raw color data for use by the DM6441. The resizer accepts image data for separate horizontal and vertical resizing from 1/4x to 4x.

The video-processing back end consists of an onscreen display engine and a video encoder. The onscreen display engine can handle two separate video windows and two separate bitmap/graphical windows. Other configurations include two video windows, one bitmap window, and one attribute window, allowing up to eight levels of alpha blending. This allows the mixing of multiple video and graphic windows on the same screen, enhancing the GUI.

The video encoder provides four analog digital-analog converters that run at 54 MHz, providing a means for composite National Television Standards Committee/Phase Alternating Line (NTSC/PAL) video, S-Video, or component video output. The video encoder also provides up to 24 bits of digital output to interface to RGB-888 devices. The digital output can provide 8/16-bit BT.656 output or CCIR.601 with separate horizontal and vertical syncs, supporting various LCD monitors and also HDTV connection.

Peripheral set

Integrating the right peripherals on the SoC reduces the overall system cost. The peripheral set includes

- two configurable video ports; 10/100 Mbps Ethernet media-access control with a management data I/O module;
- an interintegrated circuit bus interface;
- one audio serial port;
- two 64-bit general-purpose timers, configurable as two independent 32-bit timers;
- one 64-bit watchdog timer;
- up to 71 pins of general-purpose I/O with programmable interrupt/event-generation modes, multiplexed with other peripherals;

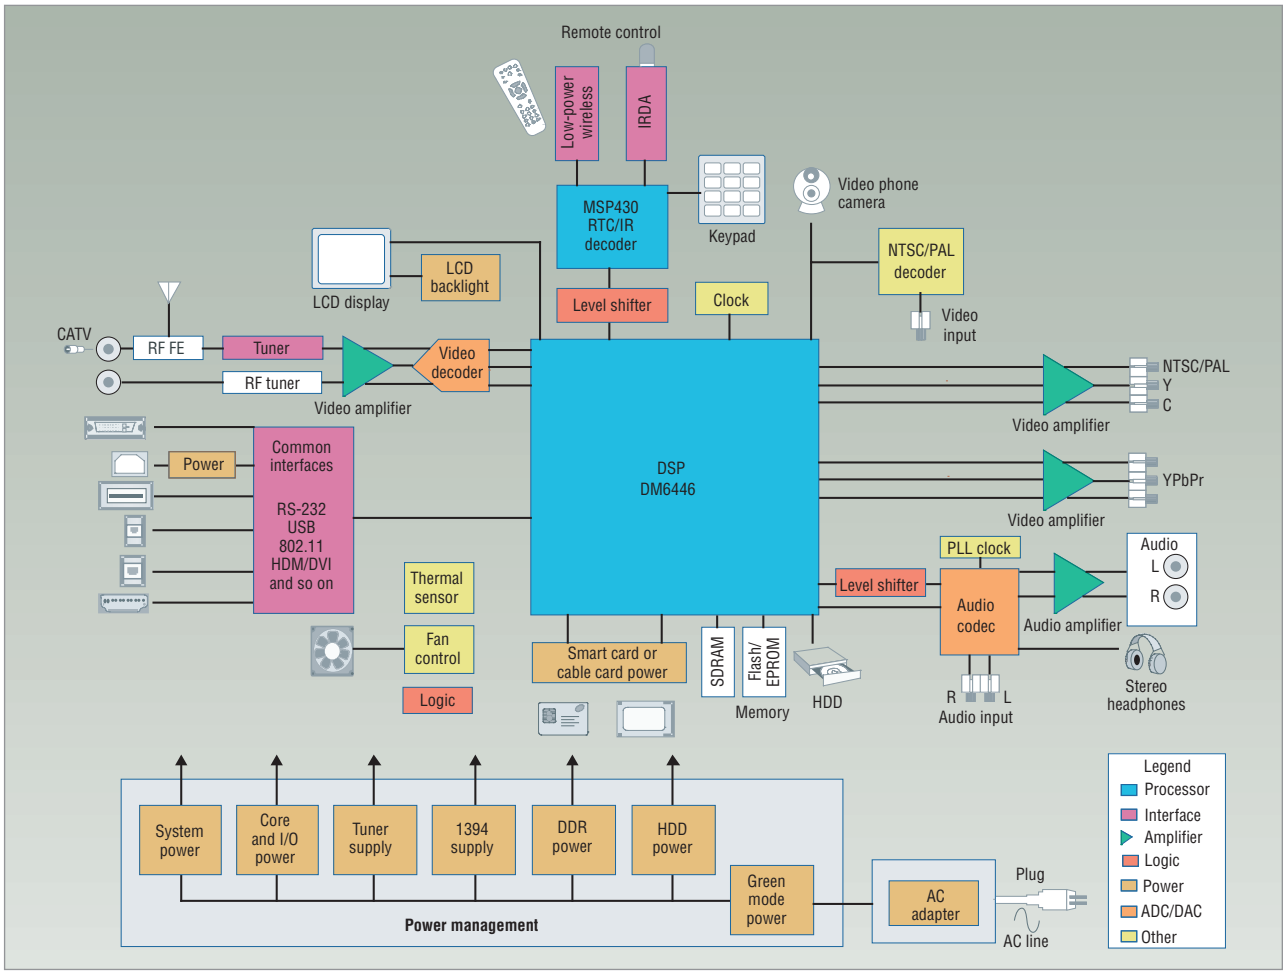

Figure 3. IP STB system block diagram. The DM6446 has an architecture similar to the DM6441 but runs at higher clock rates and includes additional peripherals and features tailored for the IP STB market.

- three universal asynchronous receiver/transmitters with hardware handshaking support on one receiver/transmitter;
- three pulse-width modulator peripherals;
- an asynchronous external memory interface for slower memories and peripherals; and
- a higher-speed synchronous memory interface.

The host port interface, interintegrated circuit, serial peripheral interface, USB2.0, and VLYNQ ports allow DM6441 to easily control peripheral devices and communicate with host processors. The DM6441 also supports Memory Stick and Memory Stick Pro cards, MultiMediaCard/Secure Digital with Secure Digital I/O, and a USB.

IP SET-TOP BOXES

The IP STB is an emerging application for delivering video content including traditional live broadcast (using multicast) as well as video-on-demand (VOD) movies and specialized content over a broadband connection. Telecommunications companies in particular are looking to offer new video services to compete with cable providers' triple-play offerings.

Figure 3 shows a block diagram of an IP STB system using the DaVinci DM6446 media processor. With an architecture similar to the DM6441, the DM6446 runs at higher clock rates and includes additional peripherals and features tailored for the IP STB market.

Enabling advanced services

While video playback is the primary application, service providers also like to offer services such as video telephony. It's also useful to support digital-media-adapter functionality, including digital-photo viewing or Web streaming via a TV display.

The DM6446 media processor supports the basic requirements for IP STB and offers opportunities for advanced services. One of the primary applications for IP STB systems is for decoding advanced media codecs such as H.264 and WMV 9/Video Codec 1.

Early IP STB offerings used the MPEG-2 video format common in digital TV broadcast. However, with maximum guaranteed bandwidth still often less than

1.5 Mbps for many service providers around the world, it was difficult to deliver even full SD video using MPEG-2. This drove IP STB services to become one of the first applications to migrate to H.264 and WMV 9 because the improved coding efficiency was critical for delivering the minimum acceptable video quality within the available bandwidth constraints.

Codec support

Box vendors can leverage the programmability of the DM6446 64x+ DSP and accelerators to support multiple service providers that require different codecs. Initial deployments were based largely on H.264 and Video Codec 1, but trial deployments in China are also starting to use the audio video standard.

Multiformat decoding. Multiformat decoding is also a key ingredient for supporting digital media adapters. Internet streaming across a wide range of content sources requires support for proprietary codecs in addition to more common industry standards. For example, On2 VP6.2 is the codec of choice for various sites using Adobe Flash. Programmable media processors like the DM6446 offer a cheaper embedded alternative to the PC with support for a wide range of current and future proprietary codecs.

The DM6446 supports all of the processing functions required for IP STB, including videophones. The ARM926 host processor can run the necessary networking stacks and system control layers such as Session Initiation Protocol or H.323. The combination of the 64x+ DSP core and the video/imaging coprocessor can perform H.261, H.263, or H.264 video encoding and decoding. The VICP is used as a motion estimation accelerator supporting up to 16 sum-absolute differences per cycle in parallel with the 64x+ DSP, which performs most of the other video-encoding functions.

Optimized rate control. The programmability of the 64x+ lets customers implement optimized rate control tailored for each application. Since latency is critical for video communication, designers can tune codec implementations to minimize the end-to-end latency in encode and decode functions. The DSP also supports flexible combinations of speech codecs (G.711, G.728, and G.722.1) and acoustic echo cancellation. The DSP basic I/O system real-time kernel enables multitasking between video and audio functions.

Integrated peripherals. The DM6446 integrates the key peripherals needed to perform IP STB, videophone, and digital-media-adapter functions with low overall system cost. Integrated peripherals support capturing video from a standard camera, USB camera, or CCD/CMOS sensor. The integrated video back end can drive an LCD or TV with the decoded video. Designers can use the onscreen display hardware to include a self-view of the local video being captured or the reference frame of the encoded video in a picture-in-picture window.

The integrated Ethernet media-access control allows interfacing to an external PHY for the required network connection. An advanced-technology-attachment interface supports including an optional hard drive for trickle-down VOD content. The DM6446 includes options for supporting secure boot and key management, which are critical for conditional access and digital-rights management.

The DSP basic I/O system real-time kernel enables multitasking between video and audio functions.

VIDEO SECURITY

Video security is moving to digital processing by replacing the older analog closed-circuit TV deployments. Low-cost vision sensors combined with DSPs that send continuous information to a security company are becoming more common in homes and offices. Instead of humans continuously monitoring the scenes, digital video security systems have built-in intelligence based on analytics running on the processor, resulting in lower system operating costs and chance of human errors. These analytics tend to be advanced motion and event detection, object detection and tracking, and object-recognition algorithms that require powerful and programmable DSP capability.

The DaVinci DM643x processors are tailored for cost-sensitive digital-media applications and include features that make them suitable for video security as well as automotive vision applications such as lane departure and collision avoidance, machine-vision systems, robotics, and video telephony.

In contrast to the DM644x, these systems don't contain a general-purpose processor or ARM. The lack of a high-level operating system requirement in the DM643x target digital-video markets, including security devices, results in a lower cost.

The combination of the C64x+ DSP and the coprocessor results in a lower L2 cache size in the DM644x. Doubling the L2 cache size allows for higher performance of the video data analytics, yet reduces the processor cost for the video-security device market. The video-processing subsystem block in the DM643x is reused from the DM644x processors. Relevant peripherals such as a component interconnect interface and a high-end controller area network are added to the DM643x processor while removing the inapplicable elements.

DIGITAL TV

The past few years have seen a rapid migration of TVs to support digital broadcast reception, including the Advanced Television Systems Committee in the US. ATSC digital TVs must be able to receive MPEG-2 trans-

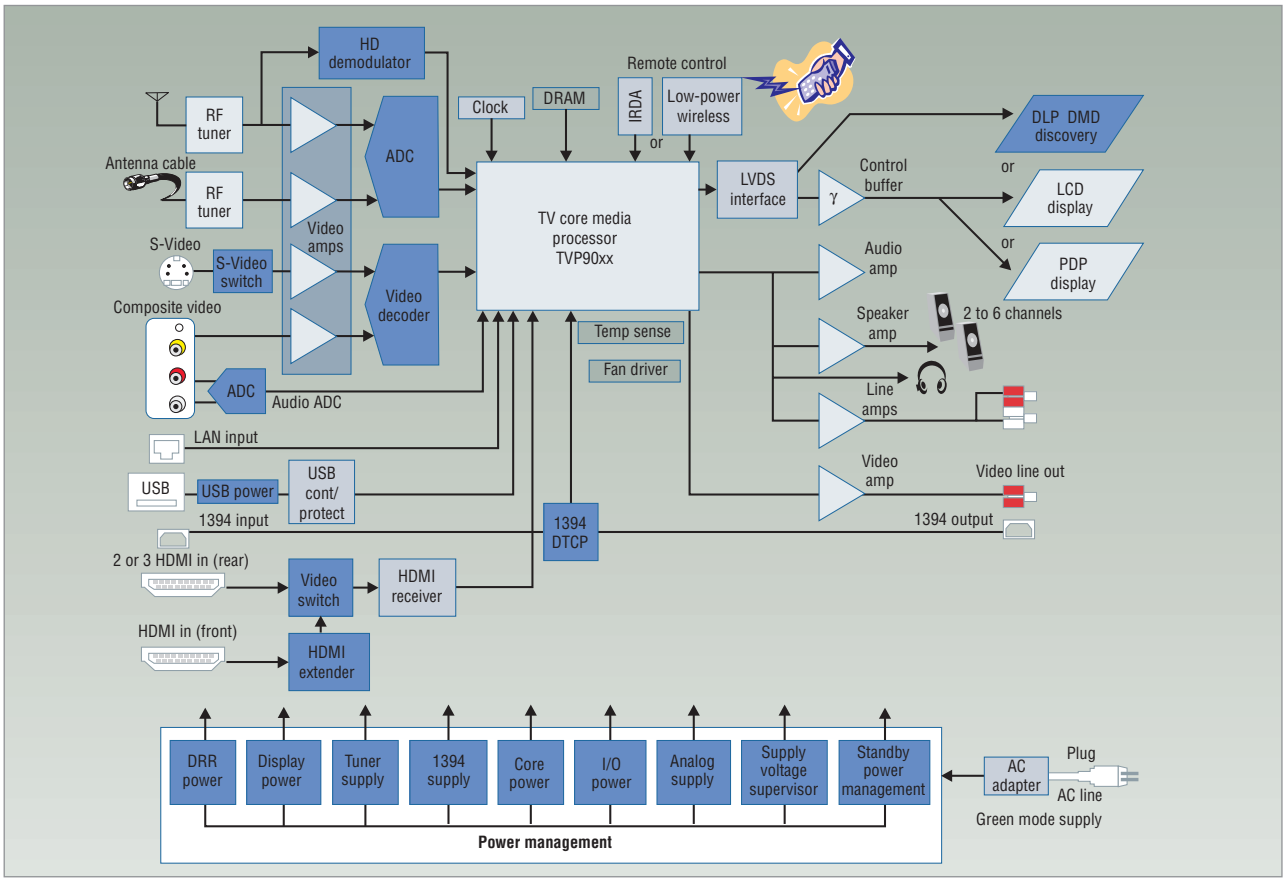

Figure 4. System block diagram of digital TV. The system minimizes overall system bill-of-materials while maximizing video quality.

port stream data from a front-end vertical sideband demodulator and decode MPEG-2 video and 5.1 AC-3 audio. Unlike the previous examples, basic broadcast digital TV solutions are architected with fixed-function video decoders since, for example, ATSC only requires MPEG-2 video.

Display processing

While audio and video decoding has rapidly become a commodity, the differentiating capability in digital TV SoCs is predominantly in the display processing. Content is now available from many sources. Some content is created from lower-quality video recorders, yet consumers expect high quality on their HDTV purchases. The bar for system integration is rapidly rising, including integration of an analog video decoder with optimized 3D Y/C comb filtering to optimize the quality from an analog source.

Figure 4 shows a system block diagram of a family of ATSC digital TV SoCs optimized for integrated digital TV receivers. The TVP90xx minimizes overall system bill-of-materials while maximizing video quality, combining a high level of integration and high-quality display processing.

The display-processing pipeline is the key differentiator for digital TV and includes multiple steps. Designers can use firmware to adjust display-processing parameters to optimize quality across a range of display characteristics. The processing requirements can be in the trillions of operations per second, requiring massive hardware parallelism and highly optimized algorithms.

The TVP90xx includes the same ARM9 core used in the DM6441 and DM6446. The ARM supports a complete ATSC digital TV software stack. MPEG-2 HD video decoding is supported using optimized hardware as opposed to a programmable DSP. Future networked digital TVs that allow Internet content access may migrate to multiformat video engines that require programmability, but for pure broadcast applications, this is the best tradeoff in terms of system cost.

Transport packet processing

A key component in DTV devices is a complete transport packet-processing subsystem. Key TPP functions include separating audio, video, and adaptation field data included in the transport stream and decryption including the Data Encryption Standard/Triple DES and Digital Video Broadcasting common scrambling algorithm associated with the particular service. The TPP also includes hardware support for synchronizing the local clock to reference information tied to the encode source's clock. This allows fine-tuning to handle drift

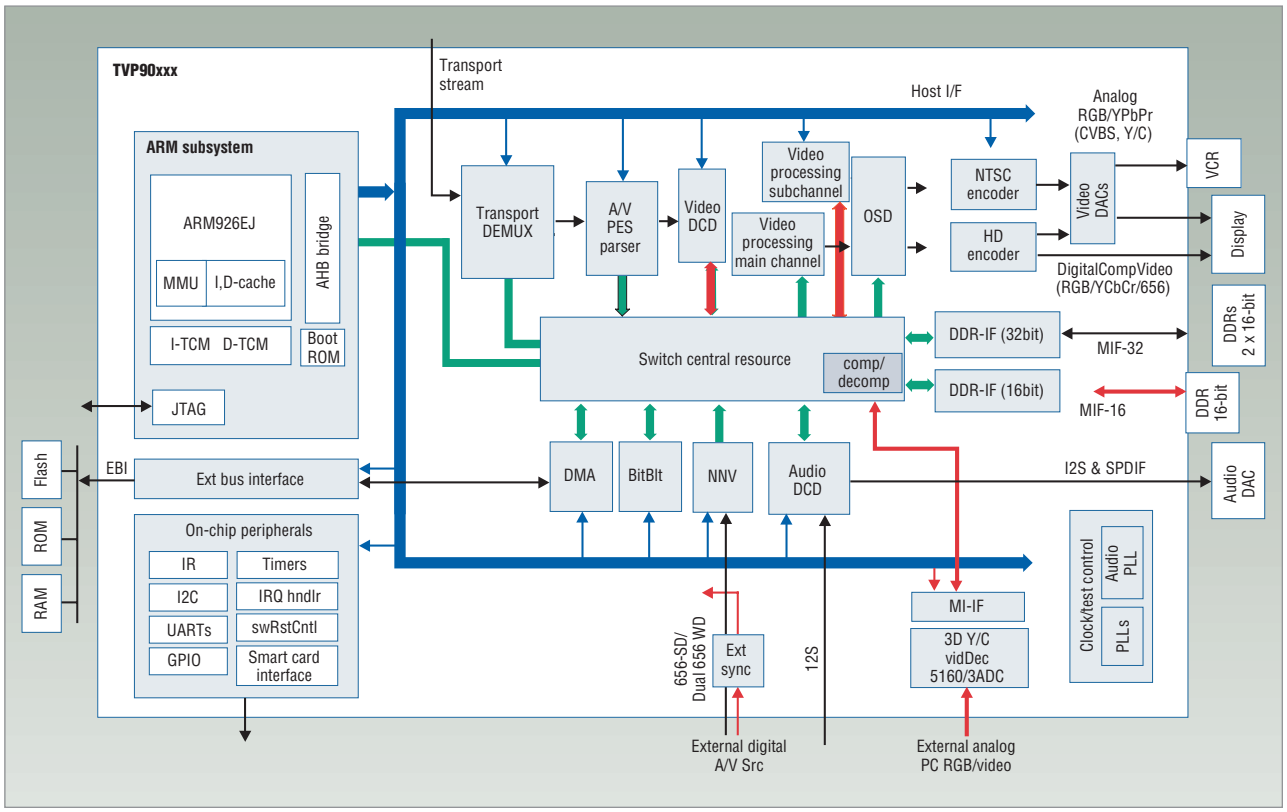

Figure 5. Block diagram of the TVP90xx processor. A 24-bit audio processor supports decoding of audio formats.

between the local and broadcast source clocks without having to drop frames of video or audio samples.

Deinterlacing

Digital TVs are typically progressive, while a significant amount of the broadcast content is encoded in 1080i HD or 480i standard definition (SD). Deinterlacing uses data from multiple fields, with the number of fields depending on the amount of available I/O bandwidth and the target video quality, to create the best representation for the missing line in each field. Many types of deinterlacing algorithms are available, but the bar is increasingly higher in terms of the expected quality. The TVP90xx includes high-quality motion-adaptive deinterlacing hardware supporting output at up to 1080p60.

It's often necessary to scale content from the source resolution to the display resolution. The quality of the scaling is particularly important when displaying SD source content on an HD display. The system needs edge enhancement to avoid excessive blurring of detail in the scaling process. It's important to perform the edge enhancement with carefully designed algorithms to avoid undesired flickering.

After format conversion is completed from the source resolution to the display resolution, another critical step is color processing. Key features include noise reduction, black-level and white-level expansion, sharpness enhancement, chroma transient improvement, dynamic-contrast improvement, flesh-tone correction, color-saturation compensation, green and blue boosting, picture controls (brightness, hue, contrast, and saturation), and programmable gamma.

Other attributes

As Figure 5 shows, a 24-bit audio processor integrated into the TVP90xx supports decoding of various audio formats, including 5.1 channel Dolby AC-3 decoding and MPEG-1 and MPEG-2 Layers I and II audio decoding with mono, stereo, joint stereo, and dual channels. The audio processor block contains the arithmetic processing hardware and internal memory necessary for performing audio decoding for MPEG-2, Dolby AC-3, and linear pulse-code modulation modes. The core's programmable read-only memory block stores the algorithm in microcode. This microcode is based on a rich instruction set with additional dedicated assembly instructions to facilitate the audio decode process.

The onscreen display controller and 2D graphics accelerator can support applications with sophisticated GUIs. The TVP90xx includes a digital-video input interface that provides a glueless interface for a variety of digital-video signals. It offers a 20-bit data interface, horizontal sync, vertical sync, clock, and field ID signals. It supports 4:2:2 SD (480i/576i), enhanced definition (480p/576p), and HD (720p/1080i) video inputs. The TVP90xx

includes an NTSC/PAL/Sequential Color Memory video decoder to decode all popular baseband analog video formats into digital component video. The video-decoder digital core utilizes technology for locking to weak, noisy, or unstable signals and can autodetect between broadcast quality and VCR-style video sources.

In addition to digital-video outputs, the video-decoder digital core generates synchronization, blanking, field, active video window, horizontal and vertical syncs, clock, genlock (for downstream video encoder synchronization), host CPU interrupt, and programmable logic I/O signals.

The TVP90xx requires external memory for both the video and audio decoder and the actual application software. A unified double data rate 32-bit memory interface is provided with adequate bandwidth to support high-performance HDTV applications. A second independent 16-bit DDR memory interface is also provided for higher-performance features that exceed the primary 32-bit DDR bandwidth capability.

An NTSC/PAL video encoder is integrated into the TVP90xx. Video output capability is provided for YC/RGB analog/digital output up to 1,080-pixel resolution, including nonstandard resolution support for panel-based DTV receivers.

To further reduce system cost, several peripherals have been integrated into the TVP90xx: two universal asynchronous receiver/transmitter serial data interfaces, programmable infrared input and output ports, a smart card bus to connect peripherals such as display and control panels, five inter-integrated circuits, and erasable programmable ROM.

The TVP90xx offers a scalable software-development platform enabling developers to build DTV receivers for consumer markets using readily available software modules. Software on the TVP90xx is divided into three sections: firmware, API, and user application software. The firmware and APIs are delivered along with the processor, while the manufacturer develops the user application software.

Firmware is used for low-level control of the hardware and bitstream demultiplexing in the TVP90xx. APIs are provided in the form of a runtime support library. The TVP90xx firmware comprises a collection of interrupt service routines and supervisor mode runtime support programs. This software isolates the application software from any direct interaction with the hardware. The user application software resides in external memory and contains the software the application programmer writes. The user application software communicates with the firmware through APIs.

The performance, power, features, and cost points for DV equipment vary, preventing one SoC solution for all devices. For power efficiency, increasing computation requirements are solved with massive parallelism, including optimized subsystems for image and video capture, video compression and decompression, and display processing instead of relying on ever-increasing CPU clock rates.

The mix of programmability and processing complexity in key subsystems varies. Yet, there's enough commonality between these devices to enable sharing of several computing engines, IP blocks, and peripherals. Reuse of hardware and software components can lead to significantly reduced development costs and times.

The bar has been raised for SoC manufacturers to provide a more complete system solution beyond silicon. Software and collateral around the SoC are increasingly important. The market for DV devices and the diversity of equipment are expected to grow significantly over the next decade, and consumer expectations will continually increase from one product generation to the next. System solutions that can hit the right level of optimization across different metrics will separate themselves in the market.

Reference

1. D. Talla et al., "Anatomy of a Portable Digital Media Processor," *IEEE Micro*, vol. 24, no. 2, Mar./Apr. 2004, pp. 32-39.

Deepu Talla *is a senior member of technical staff responsible for product definition for digital imaging and video devices at Texas Instruments. He received a PhD in computer engineering from the University of Texas, Austin. Talla is a member of the IEEE, the IEEE Computer Society, and the ACM. Contact him at deepu@ti.com.*

Jeremiah Golston *is a distinguished member of technical staff and chief technical officer of Texas Instruments' digital imaging and video business. Golston received an MS in electrical engineering from the University of Missouri-Rolla. Contact him at jgolston@ti.com.*