
SoC drawer: SoCs and the digital content revolution

Opportunities and challenges for secure digital video, audio, and content delivery acceleration

Skill Level: Introductory

[Sam Siewert \(Sam.Siewert@Colorado.edu\)](mailto:Sam.Siewert@Colorado.edu)

Adjunct Professor

University of Colorado

17 Jan 2007

SoC (system-on-a-chip) architectures could significantly accelerate digital video processing and enable the digital video revolution. Sam Siewert offers an overview of digital video processing and emergent applications in the video realm and shows how SoCs can uniquely accelerate processing. If you're an SoC architect, developer, Power Architecture™ platform software developer, or anyone creating digital video applications and services, this article is for you.

The *SoC drawer* series covers emergent applications relevant to SoC implementation, architectural concepts, and development methods. This article provides an overview of digital video applications and services that can be uniquely accelerated with an SoC architecture. The series aims to provide system architects with a starting point and some tips to make SoC design, implementation, and test easier.

Digital video is an exciting SoC application space and growing market. This article looks at the most important digital video services and applications, and then closely examines basic digital video encoders and transcoders as well as methods for protecting streaming content using digital rights management and conditional access. Understanding basic principles of digital video and some of the current challenges facing support for ubiquitous streaming video applications is critical for SoC architects working in this space.

Digital media application opportunities for SoC designs

Digital media, including video, audio, and real-time rendered graphics, collectively offer an enormous market for SoC designs. The growing ubiquity of video and audio devices, ranging from hand-held audio/video players to home entertainment systems and large entertainment and information networks, makes digital video one of the more opportunity-rich applications of SoCs. Some of the most interesting and rapidly growing applications are:

- **Video/audio codecs and transcoders:** This includes a wide variety of digital video and audio acceleration tools to compress, decompress, encode, decode, and transcode content. *Transcoding* provides the ability to convert formats -- from MPEG-2 content to MPEG-4, for example. Most often, the term *codec* refers to digital content compression and decompression, but it can also simply mean the encoding of analog data into a digital format or vice versa.
- **Digital rights management (DRM):** DRM encompasses encryption and decryption of content to protect the content from misuse or pirated redistribution, along with content watermarking and fingerprinting to assist law enforcement in tracking illegally copied content.
- **Video/audio authoring tools and special effects:** Digital media authoring tools have become available to the masses through the recent massive reduction in costs of computing systems, which in turn is being powered by ever more powerful SoCs in embedded devices and personal computers. The ability to author content at home or in a small business setting has created a digital content revolution and a second wave of Internet growth. This phenomenon is well known, but in case you've just returned from a trip to Mars, check out YouTube and *The Long Tail* by Chris Anderson of *Wired*. (See [Resources](#) for links to both.)
- **Geographical information systems (GIS) and photogrammetry:** Streaming media can of course include a combination of video, audio, and graphics, and one of the great examples of such a combination is the real-time global geographical information system. Google Earth and Microsoft® Virtual Earth are examples of this emergent form of streaming content.
- **Computer vision, image processing, and video surveillance:** In the

January 2007 issue of *Scientific American*, Bill Gates prognosticated that a robotics revolution will soon bring changes as sweeping as those wrought by the home PC. If this is to happen, much of this new technology will rely upon the internetworking of large sensor networks in the home, and would likely include significant computer vision and advanced forms of sensing similar to human capabilities. If Bill is right, this will take quite a few real-time sensor and image processing SoCs.

Basic principles of codecs

From analog (NTSC) to digital (ATSC) television standards

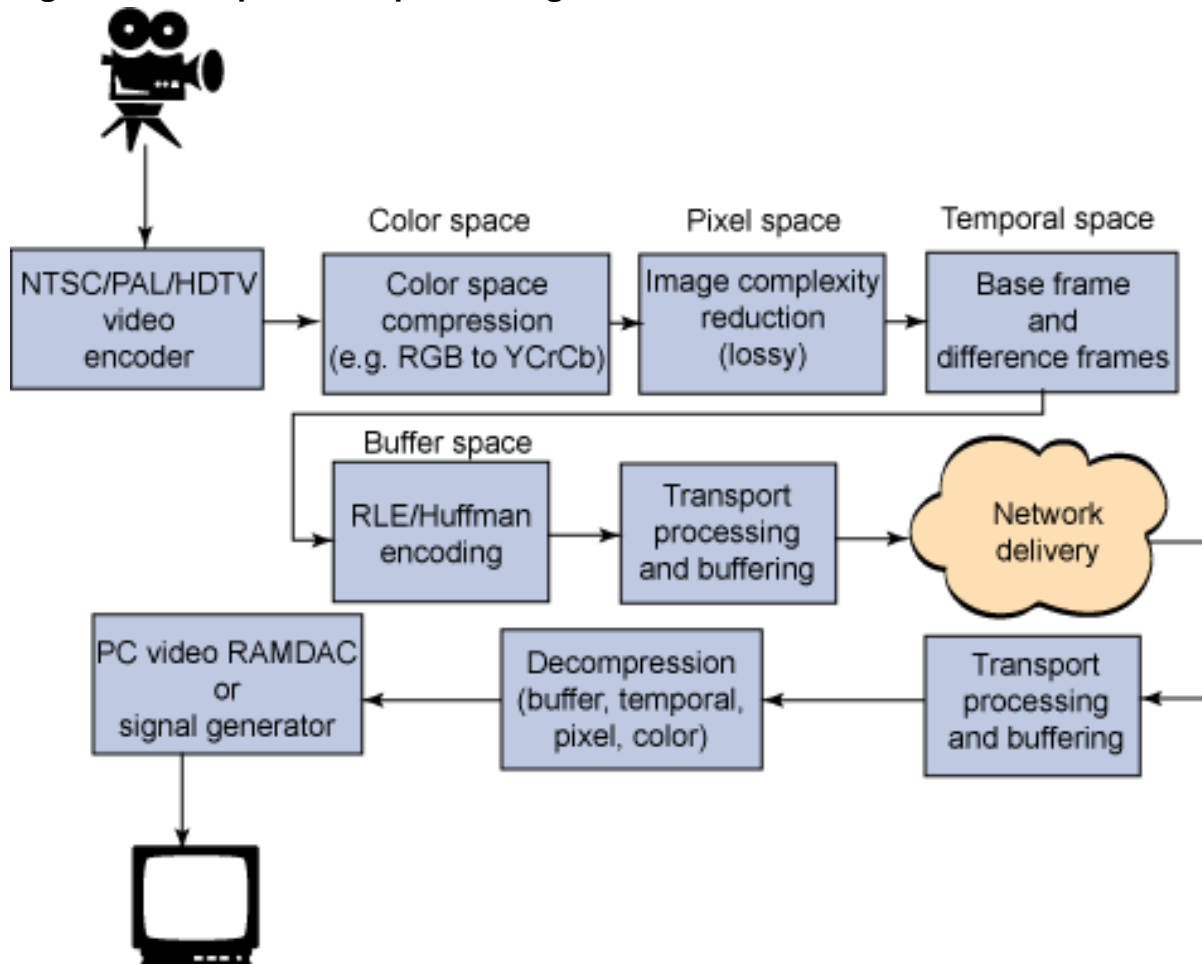
The NTSC (National Television Standards Council) has defined analog television in North America since 1941. Much of Europe, South America, Africa, and Asia uses a similar analog standard, called PAL; a much smaller portion of the world, including France and parts of Africa and Asia, uses the SECAM standard. HDTV (High Definition TV) and EDTV (Enhanced Definition TV) are slowly replacing NTSC. The FCC has mandated a change from analog NTSC modulated broadcast to digital broadcast conforming to a new ATSC (Advanced Television Systems Committee) standard. This mandate has no direct relationship to the emergence of HDTV, but many broadcasters will coordinate the upgrade to digital broadcast with the use of HDTV.

A *codec* most often describes a digital algorithm for compressing, transporting, and decompressing digital video or audio so that it can be distributed over IP (Internet Protocol) networks. A codec may also include analog signal encoding to digital through an ADC (analog to digital converter) -- for example, an NTSC codec can be used to digitize NTSC luminance and chrominance signals into various digital formats, such as RGB (24-bit red, green, and blue) or Y:Cr:Cb (16-bit luminance and red/blue chrominance). Likewise, for audio, a codec might include the analog front-end to record audio using pulse code modulation to sample audio, which may be further converted into compressed formats like MP3 (MPEG-1 Audio Layer 3) that are suitable for network transport.

The key is that a codec transforms analog audio and video information into digital formats that can be processed, managed by software applications, and easily integrated into Internet applications and services. Likewise the codec is fundamental to all audio/video consumer products and to content distribution systems and networks. This is what most of us imagine to be the keystone of the future Internet. When codecs digitize two of the most human information mediums (visual and auditory) and make them available to be time shifted, location shifted, and even augmented with graphics to become virtualized, the future world seems almost

boundless. The Star Trek holodeck is coming -- well, okay, maybe in a few years. Before we get that far ahead of ourselves, though, look at a basic example of a video codec, illustrated in Figure 1.

Figure 1. Example codec processing flow



Most networks for entertainment are still either entirely analog or are hybrid networks, with digital media storage at content head-ends, digital media transport over IP backbones, and conversion to analog near or at the end-user devices. Broadcast television, radio, and traditional cable are still mostly analog. However, in the future we will see digital fiber to the home, digital radio, and more on-demand video and audio over the Internet in real-time streams, as file sharing and as downloads. The codec shown in Figure 1 might even be used in the home: it could simply be used to distribute video from analog devices to a media store and from a media store for viewing and listening on analog equipment.

SoCs offer huge advantages over multichip solutions for home-based media systems, consumer electronics, and the emerging digital video networks

infrastructure in terms of cost, power, time to market using rapid development and reconfigurable SoCs, and overall integration of features. For example, AES-128 encryption on a typical MPEG-2 stream encoding an NTSC signal would bring a microprocessor to its knees if data were to be encrypted with a software-based algorithm. However, an SoC, even a reconfigurable one, can incorporate a crypto IP core to offload software and perform in-data-path encryption of a video stream for conditional access. For example, the Virtex-4 reconfigurable SoC provides a PowerPC® core and plenty of reconfigurable logic to integrate a crypto IP core like the Helion AllianceCORE (see [Resources](#)). This type of integrated software processing and encryption engine in a reconfigurable SoC can help accelerate time to market for content distribution engines used from the head-end all the way down to a set-top box.

Digital rights management

For home videophiles and audiophiles, the quantity and quality of authoring, editing, and distribution tools has skyrocketed. Most users are not too concerned about others obtaining their content. In fact, most of us proudly display content on home Web pages and wish that there were easier ways to share large audio and video files over the Internet with family and friends.

Of course, premium content providers, such as movie or television studios, care deeply about controlling who can view their content and when, since it is their livelihood. In the coming years, individuals may also care much more as content proliferates and the need to share becomes more balanced with the desire to maintain privacy as well. The ability to share easily yet retain privacy so that only select audiences can view home-authored or premium content is the domain of *digital rights management*, or DRM. The DRM space is full of applications that could benefit from SoCs.

Digital rights management includes three main features to protect digital content from misuse:

- **Data encryption for conditional access:** Digital content such as MPEG-2/4 streams and movies, MP3 audio, and multimedia can be encrypted in advance by the content provider prior to delivery for streaming over on-demand or broadcast digital networks. Users of the content must therefore have decryption hardware (most often a set-top box for video on demand, satellite, or cable) and an authenticating smart card providing the decryption key. This ensures that only users who have paid for the content and have a right to use it can decrypt it for consumption.

- **Content watermarks:** Historically, applications that control access to digital content have not used the strongest encryption available. For example, today AES-128 might be used for such an application despite the availability of AES-1024. The encryption is in general strong enough to prevent most illegal use of content; however, it cannot prevent the most determined pirates.

To assist with identification of legitimate content and to assist with identification of the source of content, *digital watermarks* can be applied prior to encryption. A digital watermark is a signature that is hidden in the background noise of the audio or video so that it does not inhibit enjoyment of the media, but does clearly identify the source and legitimacy of the content. A watermark could for example include information about when, from where, by whom, and for whom content was originally provided. This would make it quite clear if the content were being broadcast, streamed, or otherwise distributed from an illegitimate source.

- **Content fingerprinting:** A watermark is used to identify content's source and legitimacy. While this is very helpful for distinguishing illegally from legally distributed content, it does not help law enforcement to trace down the content pirate. *Digital fingerprinting* does. Fingerprints are placed on content by means of techniques similar to those for watermarks; however, the information includes much more specific data on the original consumer of the content. For example, in an on-demand streaming system, fingerprints could be applied based upon the exact end user on each stream as it goes out or as it is consumed. In the end, if such a stream is pirated and redistributed, the pirated content can be traced back to a specific end-user's set-top box or personal computer.

Encryption export regulations

Note that U.S. regulations prohibit the export of some encryption engines. A careful understanding of the regulations regarding algorithms and strengths (for example, AES-128 versus AES-1024) that are being incorporated into designs is advised. Of course, this legislation can be and has been challenged through litigation.

To be safe, product developers should carefully consider whether their product will be sold in the public domain and whether the product provides an encryption engine capability that requires export control or simply data privacy and protection that can be exported. For example, the Seagate DriveTrust technology (see [Resources](#)) provides SoC-based encryption in laptop drive controllers, but limits the use of the RSA and AES encryption capability to content on the drive media only; the SoC-based

encryption can't be used to stream encrypted data off of the drive.

After a user of the trusted drive authenticates him or herself with a pass phrase, data is delivered to and from the drive in clear text only. The point of DriveTrust is to protect the data on the drive if a laptop is stolen, preventing unauthorized use of the drive to access or modify data on it. As such, this product is not an encryption engine with export restrictions, and it only provides personal privacy and data protection. Developers of SoCs with encryption should carefully review encryption export regulations.

Protecting content, especially premium content from studios in Hollywood and around the globe, is fundamental to the digital content revolution. Individuals may also have concerns about personal information privacy as well. Clearly, large premium content producers and providers need to carefully protect their content, especially in digital forms that can be reproduced and modified with high fidelity. A scratchy, bootlegged audio tape recording from a concert has little to no resale value and is no threat to premium content providers; but high-definition audio or video is of course a huge risk. Likewise, for an individual, the idea of someone you don't know -- perhaps someone with a criminal background -- stealing all of your family digital photos and video is really frightening. Given both realizations, the race is on to protect individual content and downloaded content (for example, the Seagate DriveTrust technology) and content on networks (for example, on-demand streaming video and audio).

SoCs are uniquely poised to implement conditional access, watermarks, and fingerprinting because of the need for high data rates and hardware-accelerated encryption and hashing algorithms that just can't be cost-effectively implemented in software. Furthermore, while audio fingerprinting is fairly well defined, video fingerprinting and watermarking is an emergent and rapidly changing technology that can benefit from the rapid deployment that reconfigurable SoCs can support.

Streaming content from all to all

Assuming that all goes well with the digital media revolution, SoCs can play a key role in premium content distribution, DRM, and presentation to end users. Likewise, SoCs already play a huge role in consumer electronics, and this role will only continue to grow. The ability to safely, privately, and selectively stream content from anywhere to anywhere will open up a new virtual world, but one with limitations that can be imposed with policy that makes it safer, more profitable, and better targeted to personal interest and taste. The long tail described by Chris Anderson is only the start. Geographically remote families will be closer in the Jetsons sense, finally. The ability to virtually travel to far-off places before spending large sums on real travel will also be a reality, and the ability to navigate and be guided by advanced GIS will

help bridge the real world and the virtual. Continuing education will be possible effectively anywhere and anytime, publishing will be transformed and open to more authors, and virtual presence will become more real; all this will be powered by countless SoCs.

Conclusion

In the end, SoCs seem well poised to accelerate this human need to share high-fidelity content that appeals to our highest bandwidth senses, and to share with managed risk, security, and privacy. The role of the SoC in this revolution is bound only by our ability to implement the services needed to fuel human imagination and the next generation of the World Wide Web.

Resources

Learn

- Wikipedia has an excellent overview of several of the technologies discussed in this article:
 - [NTSC](#)
 - [EDTV](#)
 - [HDTV](#)
 - [MJPEG](#)
 - [Acoustic fingerprinting](#)
 - [MD5 \(Message Digest algorithm 5\)](#)
- [MPEG standards reference software](#): The Moving Picture Expert Group (MPEG) provides reference software for MPEG-1, -2, -4, -7, and -21.
- [MPEG LA](#): Note that usage of MPEG requires licensing.
- [Theora](#): An alternative, fully open standard for digital media codecs.
- [Litigation in progress at the Center for Democracy and Technology](#): This litigation is related to the interpretation of strong encryption export laws. For those not familiar with the CDT, the organization has been in the news regarding numerous privacy issues related to the Internet and homeland security.
- ["Video-server designs for supporting very large numbers of concurrent users,"](#) M. Kumar (*IBM Journal of Research and Development*, 2000): On the server side, IBM has worked to define scalable stream servers such as the system described here.
- ["Introduction to cryptography, from Egypt through Enigma"](#) and ["The right coprocessor can help with encryption,"](#) Sam Siewert (developerWorks, July-August 2006): Get an introduction and in-depth look at cryptography.
- ["Applications for data hiding,"](#) W. Bender, W. Butera, D. Gruhl, R. Hwang, F. J. Paiz, and S. Pogreb (*IBM Systems Journal*, 2000): IBM research on DRM technology.
- ["Media, networks, and content,"](#) A. Lippman (*IBM Systems Journal*, 2000): More musings on the subject.
- [IBM Systems Journal, 1996](#): This edition of the journal contains multiple articles on digital video research from the MIT Media Lab and IBM.

- [The Long Tail](#): For more background on this concept, check out Chris Anderson's blog.

Get products and technologies

- [IBM set-top boxes](#): The Power Architecture includes PowerPC processors and custom MPEG-2 DV chipsets for these products.
- [IBM professional encoders](#): The NV 4xx MPEG-2 encoders from IBM are based on ASIC solutions.
- [IBM high-performance decoders](#): The CS24D/E decoders are also ASICs.
- [MERLOT](#): IBM has created DRM for educational multimedia content with MERLOT (multimedia educational resource for learning and online teaching). In fact, online content is poised to revolutionize publishing and education, but success hinges upon DRM. (Document is a PDF.)
- [IBM DataHiding](#): Another DRM tech from IBM.
- [YouTube](#): In case you're out of touch with viral videos and the high-tech business world at large, this Web-based digital video service provider was recently purchased by Google for 1.65 billion dollars.
- [Helion AllianceCORE](#): Xilinx reconfigurable SoCs with PowerPC cores can integrate numerous encryption engine cores, including the AllianceCORE, which provides AES-128 encryption for conditional access.
- [The OpenCable initiative](#): CableLabs has helped coordinate cable system standards for interoperability through this initiative. Open systems standards for content protection are critical to SoC designs that will integrate well.
- [free-codecs.com](#): Find numerous open source video and audio codecs and transcoders here.
- [MediaCoder](#): This GPL'd transcoder provides transcoding to and from most codec formats with a batch process.
- [DriveTrust technology: A technical overview](#): Find out more about this technology from Seagate.

About the author

Sam Siewert

Dr. Sam Siewert is an embedded system design and firmware engineer who has worked in the aerospace, telecommunications, and storage industries. He also teaches

at the University of Colorado at Boulder part-time in the Embedded Systems Certification Program, which he co-founded. His research interests include autonomic computing, firmware/hardware co-design, microprocessor/SoC architecture, and embedded real-time systems.