WHITE PAPER

Texas Instruments

By
Ajit Rao,
Senior Member, Technical Staff;
Soyeb Nagori,
Software Architect, Multimedia Codecs;
Texas Instruments

Optimizing Video Encoders with TI DSPs

Executive Summary

Digital video compression involves sacrificing some degree of picture quality for a lower bit rate that facilitates transmission and storage. Compression not only requires a high level of performance from the processor but also some versatility in design, since different types of video applications have different sets of requirements for resolution, bandwidth and resiliency. The extended flexibility provided by Texas Instruments' video encoders address these differences and take full advantage of the options offered by advanced video compression standards to help system developers optimize their products. Operating on TMS320C64x+™ digital signal processors (DSPs), DaVinci™ technology video processors and OMAP™ mobile media processors, TI's encoders offer easy-to-use application programming interfaces (APIs) in a standard system interface that help simplify video integration and optimization in a variety of applications.

The Need for Flexible Optimization

Digital video continues to extend visual communication to an ever-larger range of applications, and an increasing number of developers are becoming involved in creating new video systems or enlarging the capabilities of existing designs. Video not only demands very high performance but the wide variety of video applications requires the performance to be optimized depending on the particular system. Programmable, high-performance DSPs and platforms from TI allow developers to adapt their software readily to application requirements that can vary widely in terms of transmission bandwidth, storage, image specifications and quality requirements.

The inherent structure and complexity of video encoding and decoding (codec) algorithms drives the optimization approach. Encoders are particularly important because they must adapt to the application and they represent a major portion of the heavy processing load of video applications. The goal for video compression is to encode digital video using as few bits as possible while maintaining acceptable visual quality. While encoders are based on the mathematical principles of information theory, they may still require implementation trade-offs, which can be quite complex. TI encoders are highly configurable, providing an easy-to-use system interface and performance optimization for a wide range of video applications. Even if the developer is new to DSPs or video algorithms, TI video encoders offer a number of techniques that can make the optimization process easier and more effective in terms of compression and picture quality.

Options in Compression Standards

Raw digital video requires a lot of data to be transmitted or stored. For example, in the absence of any compression, standard-definition (SD) digital TV in North America typically requires a data rate of over 165 Mbps or more than 100 GB required to store a standard movie. Even lower-resolution CIF video, which is often used in

streaming, requires over 36.5 Mbps, which is much more than can be sustained on most of today's broadband networks. An advanced video codec such as H.264/MPEG-4 AVC can achieve compression ratios of between 60:1 and 100:1 with sustained throughput. This makes it possible to squeeze video with a high data rate through a narrow transmission channel and store it in a limited space.

Like JPEG for still images, the widely used ITU and MPEG video encoders employ a combination of discrete transform coding (DCT) or something similar, quantization and variable-length coding to compress macro-blocks in key frames also referred to as *intra*-frames. Once the algorithm has established a baseline intra-coded (I) frame, a number of subsequent predicted (P) frames are created by only coding the difference in visual content or *residual* between each of them. This *inter*-frame compression is achieved using a technique called motion compensation. First the algorithm estimates where the macro-blocks of the current frame were in the previous frame. Then it subtracts and compresses the residual.

Improved motion compensation techniques, along with other techniques, have doubled the compression ratio from MPEG-2 to MPEG-4, but the performance requirements have also increased. The motion estimation (ME) stage, which creates the motion vectors that describe where each of the blocks has moved, is usually the most computation-intensive stage of the algorithm. Figure 1 shows the flow of a generic motion compensation-based video codec. Typically, the macro-blocks contain four 8×8-pixel luminance blocks and two 8×8-pixel chrominance blocks (YCbCr 4:2:0), though other video formats are specified.

Figure 2 on page 3 shows a P frame (right) and its reference (left). Below the P frame, the residual (black) shows how little remains to be encoded once the motion vectors (blue) have been calculated.

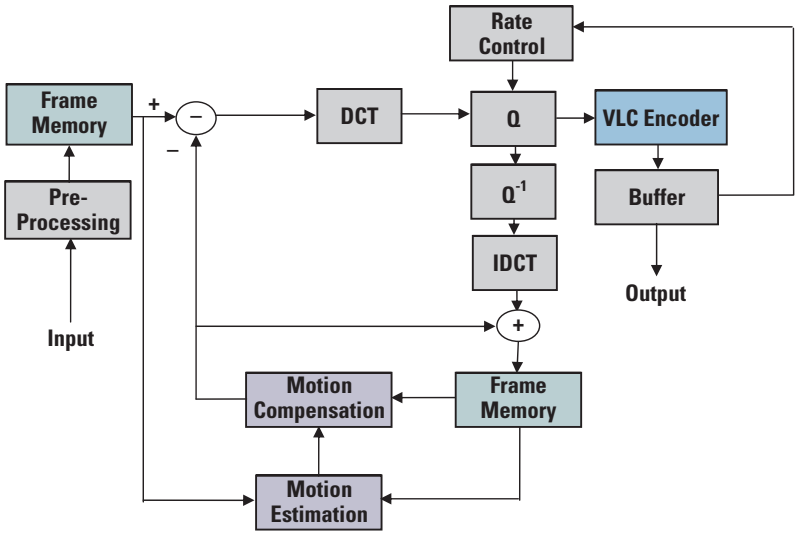

Figure 1. Generic motion compensation-based video encoder.

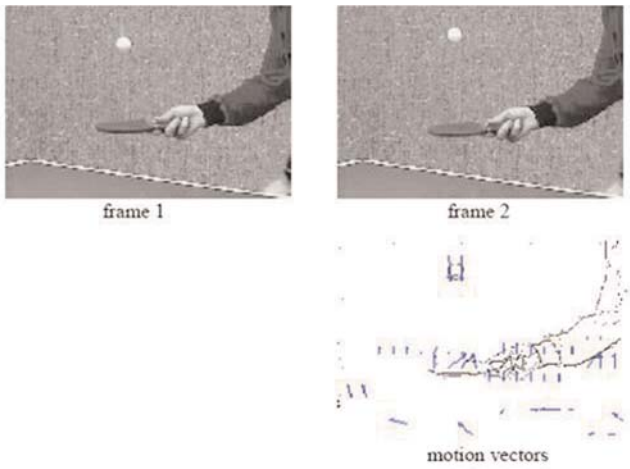

Figure 2. Frame motion vectors and residual.

Video compression standards specify only the bit-stream syntax and the decoding process, leaving a large scope for innovation within the video encoders. For example, in the ME stage, the ways that the motion vectors describe block movement are standardized, but there are no constraints on what techniques the encoder can use to determine the vectors. Rate control is another area with significant latitude for innovation, allowing the encoder to assign the quantization parameters and thus "shape" the noise in the video signal in appropriate ways. In addition, the advanced H.264/MPEG-4 AVC algorithm adds flexibility and functionality to earlier standards with multiple options for macro-block size, quarter-pel (pixel) resolution for motion compensation, multiple-reference frames, bi-directional prediction (B frames) and adaptive loop de-blocking. Important additions in the new standard are shown in Figure 3 on the following page.

Varied Application Requirements

Video application requirements can vary enormously. The range of features in advanced compression standards offer a large potential for trading off the options in order to balance complexity, delay and other real-time constraints. Consider, for instance, the different requirements for video surveillance, video conferencing and video phones, video mobile phones and digital video recorders (DVRs).

Video surveillance can require compression within digital cameras for DVRs, network transmission (IP netcams) or concentration units taking feed from analog cameras. Even with a high level of compression, storing the vast amount of visual information generated is often the key problem. One solution is to keep only the frames containing pertinent activity, such as someone entering or leaving through a secure door. The software used to ensure this may access the frame motion information to determine if any activity is occuring. A high magnitude in motion vectors signifies suspicious activity is occurring and the frame

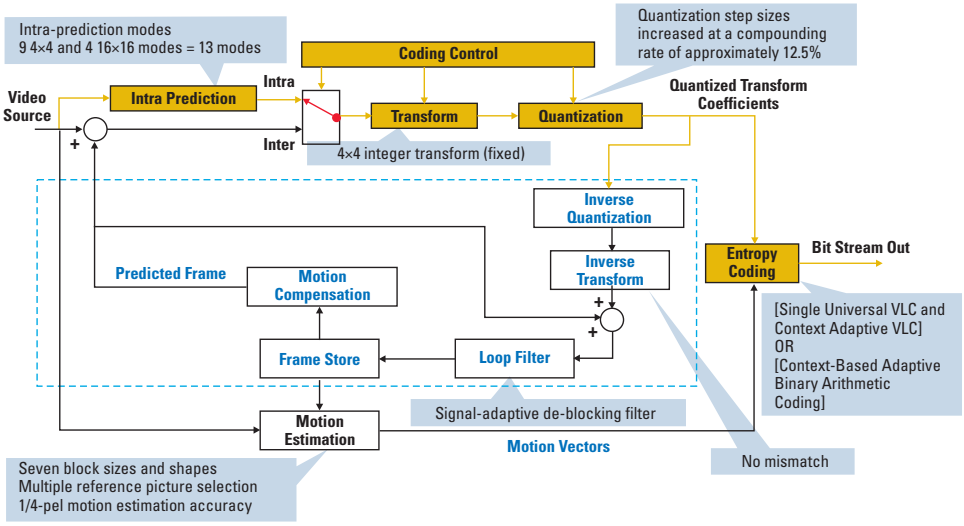

Figure 3. H.264/MPEG-4 AVC key features.

is immediately stored. All TI encoders available on TI devices based on DaVinci™ technology (TMS320DM644x, TMS320DM643x, TMS320DM647 and TMS320DM648 processors) provide access to motion vectors, a feature that not all competitors possess.

Surveillance software may also need to access the encoder parameters to compensate for dawn, dusk or other condition changes. In addition, having the ability to simultaneously encode two different rates can be an advantage with surveillance, since it allows the system to display video at one rate on a monitor while storing it on a disk at another rate. The ability to dynamically trade-off a single high-quality channel with two low-quality channels allows the amount of information stored to be greater when something happens in front of one of the cameras. These changes, like compression itself, must occur in real time in order to capture events as they occur.

Video conferencing and video phones. The most important issue with video-phone and video-conferencing applications is typically the transmission bandwidth, which can range from tens of kilobits per second up to multi-megabits per second depending on the link. In some cases, the bit rate is guaranteed but with the Internet and many intranets, bit rates are highly variable. As a result, encoders used in video conferencing frequently have to address the delivery requirements of different types of links and adapt in real time to changing bandwidth availability. The transmitting system, when it is advised of reception conditions via a reverse channel or RTCP acknowledgement, should adjust its encoded output continually so that the best possible video is delivered with minimal interruption. When delivery is poor, the encoder may respond by reducing its average bit rate, skipping frames or changing the group of pictures (GoP) structure which represents the mix of I

and P frames. Since I frames are not as heavily compressed as P frames, sending fewer I frames requires less bandwidth overall. Since the visual content of a video conference does not change frequently, it is usually acceptable to send fewer I frames than it would be to send with, say, TV or a movie.

H.264 uses an adaptive in-loop de-blocking filter that operates on the block edges to smooth the video for current and future frames. The de-blocking filter significantly improves the subjective quality of video encoded by H.264, especially at low bit rates. On the other hand, turning off the filter can reduce the computational complexity but increase the amount of visual data at a given bit rate, as can changing the ME resolution from quarter-pel to half-pel or more. In some cases, it may be necessary to sacrifice the higher quality of de-blocking and fine ME resolution in order to reduce the computational complexity of encoding.

Since packet delivery via the Internet is not guaranteed, video conferencing often benefits from encoding mechanisms that increase error resilience. For instance, progressive strips of P frames can be intra-coded (I strips), as shown in Figure 4. This technique eliminates the need for complete I frames (after the initial frame), reducing the risk that an entire I frame will be dropped and the picture broken up. Also, without the bursts created by I frames, the data flow is steadier. There is a trade-off in compression, though, since the presence of narrow I strips reduces the encoder's ability to exploit spatial redundancy. An estimated two to five percent is lost in bit rate, so it is useful if the encoder can switch this capability on or off as needed for coping with network delivery conditions.

Video mobile phones. In wireless phones with video capabilities, bandwidth is at a premium, even with 3G channels. Processing may also be more limited in these systems since handsets are designed to trade-off low power consumption for performance. The encoder

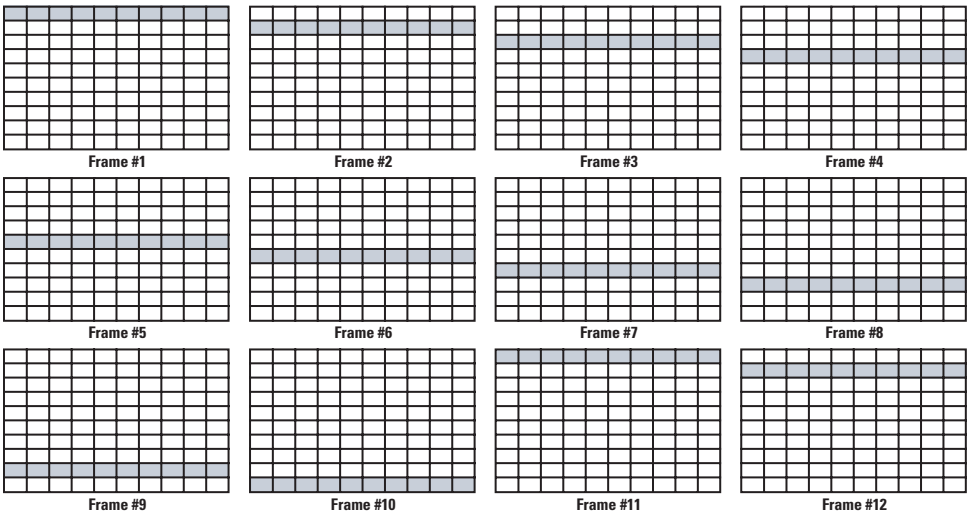

Figure 4. Intra-coded strips in P frames.

at the transmitting end has to take the receiving limitations into account by adjusting the resolution to the small display. For video streaming, a lower frame rate with fewer I frames or I strips is likely, since the picture degradation is not as apparent as it would be on a larger display, while the bandwidth savings are considerable. Video conferencing is an extreme case since the handset needs to encode as well as decode. Due to background that is usually static and little motion in the foreground, traffic may be limited to a single I frame at the beginning of the call, followed only by P frames.

Video recording. Digital video recorders (DVRs) for home entertainment are perhaps the most widely used application for real-time video encoders. For these systems, achieving the best trade-off of storage with picture quality is a significant problem. Unlike video conferencing, which is not delay tolerant, compression for video recording can take place with some real-time delay if sufficient memory is available in the system for buffering. Realistic design considerations mean the output buffer is designed for handling several frames, which is sufficient to keep a steady flow of data to the disk. Under certain conditions, however, the buffer may become congested because the visual information is changing quickly, thus creating a large amount of P frame data. In this case, the encoder should be able to trade-off quality for a lower bit rate. When the congestion has been taken care of, then the quality can be increased again.

A mechanism for performing this trade-off effectively is by changing the quantization parameter, Qp, on the fly. Quantization, as shown in Figure 1, is one of the last steps in the algorithm for compressing data. As quantization introduces noise, it creates picture distortion in direct proportion to the square of Qp. Increasing Qp reduces the bit rate output of the algorithm but sacrifices picture quality. However, since the change occurs in real-time, it reduces the likelihood of frame skips or picture break-up. When the visual content is changing rapidly, as it is when the buffer is congested, lower image quality is likely to be less noticeable than it is when the content changes slowly. Once the visual content returns to a lower bit rate and the buffer clears, Qp can be reset to a normal value.

Flexibility with TI Encoders

As we have seen, video compression standards give encoding programmers wide latitude in implementing the algorithms. H.264/MPEG-4 AVC provides a number of options that can be used by encoders in adapting to specific application requirements. TI processors are used in a wide range of systems and TI encoders are designed with flexibility to support these options through easy-to-use application APIs. Encoder versions are available that operate on any TMS320C64x+™ DSP or DaVinci™ technology video processor, with or without acceleration from the video and imaging coprocessor (VICP). Non-VICP versions are also available for OMAP™ media processors.

A basic set of APIs with default parameters is used in all the encoders, so that the system interface remains the same, regardless of the type of system. Extended API parameters adapt an encoder to the requirements of specific applications. By default, parameters are preset to high-quality settings; a high-speed preset is also available. All preset parameters can be overridden by the program using extended parameters.

Features of the encoders include support for both YUV 4:2:2 and YUV 4:2:0 input formats, motion resolution down to a quarter-pel, GoP sizes with I frame intervals ranging from every frame to none after the first frame, bit-rate control via Qp, access to motion vectors, de-blocking filter, simultaneous encoding of two channels, I strips in P frames for resiliency, and many other options for a wide range of applications. In motion estimation, TI's encoders allow the search range to be restricted for compliance with some standards, such as H.263. By default, the encoder dynamically and unrestrictedly determines the search range.

TI's encoders allow a wide range of encoding bit rates. However, there are generally "sweet spots" of operation, where output bit rates operate optimally for a given input resolution and frames per second (fps). Developers should be aware of the sweet spots with TI encoders in order to design their systems for the best trade-offs of transmission and picture quality. For DVD-type content using H.264, the main sweet spot correspondences are as follows:

- Below 128 kbps: QCIF at 15 fps
- Between 128 and 256 kbps: QCIF/QQVGA at 30 fps
- Between 256 kbps and 1 Mbps: CIF/QVGA at 30 fps
- Above 1 Mbps: VGA/D1 30 fps

For the same content with MPEG-4 compression, bit rates should increase by 30 percent to allow for the reduced compression efficiency of the MPEG-4 standard. Since video conferencing and surveillance have relatively static visual content, they have a 30 percent lower bit rate need compared to DVD-type content.

Development of the encoders is ongoing, and TI works closely with its customers to make available the features that are needed for their products. Planned enhancements include auto-detection of scene changes in order to place I frames for better picture quality and support for additional intra prediction modes that improve compression. Regardless of what changes are introduced, the encoder APIs will remain simple and straightforward in order to keep video system integration uncomplicated.

Optimization for a World of Video

Digital video continues to spread to new types of systems, and developers need to be aware of the many differences that exist among entertainment, video conferencing, surveillance, mobile and other video applications. In general, compression requirements trade-off bit rate and picture quality, though the many different ways of achieving these trade-offs can be complicated. Since H.264/MPEG-4 AVC provides a number of options that can affect these trade-offs, and since the standard does not specify *how* the encoder determines data such as motion vectors, there is a great deal of latitude in encoder design for optimization. TI encoders use this specification latitude to provide a straightforward interface that provides the flexibility to adapt video compression to the requirements of a wide range of applications. As the world of digital video continues to expand, TI processors and codecs provide high performance, ease of use and versatility that enable developers to optimize their products effectively for a wide variety of application requirements.

Important Notice: The products and services of Texas Instruments Incorporated and its subsidiaries described herein are sold subject to TI's standard terms and conditions of sale. Customers are advised to obtain the most current and complete information about TI products and services before placing orders. TI assumes no liability for applications assistance, customer's applications or product designs, software performance, or infringement of patents. The publication of information regarding any other company's products or services does not constitute TI's approval, warranty or endorsement thereof.

DaVinci, OMAP and TMS320C64x+ are trademarks of Texas Instruments. All trademarks are the property of their respective owners.

© 2007 Texas Instruments Incorporated

IMPORTANT NOTICE

Texas Instruments Incorporated and its subsidiaries (TI) reserve the right to make corrections, modifications, enhancements, improvements, and other changes to its products and services at any time and to discontinue any product or service without notice. Customers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All products are sold subject to TI's terms and conditions of sale supplied at the time of order acknowledgment.

TI warrants performance of its hardware products to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are used to the extent TI deems necessary to support this warranty. Except where mandated by government requirements, testing of all parameters of each product is not necessarily performed.

TI assumes no liability for applications assistance or customer product design. Customers are responsible for their products and applications using TI components. To minimize the risks associated with customer products and applications, customers should provide adequate design and operating safeguards.

TI does not warrant or represent that any license, either express or implied, is granted under any TI patent right, copyright, mask work right, or other TI intellectual property right relating to any combination, machine, or process in which TI products or services are used. Information published by TI regarding third-party products or services does not constitute a license from TI to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property of the third party, or a license from TI under the patents or other intellectual property of TI.

Reproduction of TI information in TI data books or data sheets is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations, and notices. Reproduction of this information with alteration is an unfair and deceptive business practice. TI is not responsible or liable for such altered documentation. Information of third parties may be subject to additional restrictions.

Resale of TI products or services with statements different from or beyond the parameters stated by TI for that product or service voids all express and any implied warranties for the associated TI product or service and is an unfair and deceptive business practice. TI is not responsible or liable for any such statements.

TI products are not authorized for use in safety-critical applications (such as life support) where a failure of the TI product would reasonably be expected to cause severe personal injury or death, unless officers of the parties have executed an agreement specifically governing such use. Buyers represent that they have all necessary expertise in the safety and regulatory ramifications of their applications, and acknowledge and agree that they are solely responsible for all legal, regulatory and safety-related requirements concerning their products and any use of TI products in such safety-critical applications, notwithstanding any applications-related information or support that may be provided by TI. Further, Buyers must fully indemnify TI and its representatives against any damages arising out of the use of TI products in such safety-critical applications.

TI products are neither designed nor intended for use in military/aerospace applications or environments unless the TI products are specifically designated by TI as military-grade or "enhanced plastic." Only products designated by TI as military-grade meet military specifications. Buyers acknowledge and agree that any such use of TI products which TI has not designated as military-grade is solely at the Buyer's risk, and that they are solely responsible for compliance with all legal and regulatory requirements in connection with such use.

TI products are neither designed nor intended for use in automotive applications or environments unless the specific TI products are designated by TI as compliant with ISO/TS 16949 requirements. Buyers acknowledge and agree that, if they use any non-designated products in automotive applications, TI will not be responsible for any failure to meet such requirements.

Following are URLs where you can obtain information on other Texas Instruments products and application solutions:

Products		**Applications**	
Amplifiers	amplifier.ti.com	Audio	www.ti.com/audio
Data Converters	dataconverter.ti.com	Automotive	www.ti.com/automotive
DSP	dsp.ti.com	Broadband	www.ti.com/broadband
Interface	interface.ti.com	Digital Control	www.ti.com/digitalcontrol
Logic	logic.ti.com	Military	www.ti.com/military
Power Mgmt	power.ti.com	Optical Networking	www.ti.com/opticalnetwork
Microcontrollers	microcontroller.ti.com	Security	www.ti.com/security
RFID	www.ti-rfid.com	Telephony	www.ti.com/telephony
Low Power Wireless	www.ti.com/lpw	Video & Imaging	www.ti.com/video
		Wireless	www.ti.com/wireless

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265
Copyright © 2007, Texas Instruments Incorporated