H.264 video compression standard.
New possibilities within video surveillance.

AXIS®
COMMUNICATIONS

Table of contents

1. Introduction

The latest video compression standard, H.264 (also known as MPEG-4 Part 10/AVC for Advanced Video Coding), is expected to become the video standard of choice in the coming years.

H.264 is an open, licensed standard that supports the most efficient video compression techniques available today. Without compromising image quality, an H.264 encoder can reduce the size of a digital video file by more than 80% compared with the Motion JPEG format and as much as 50% more than with the MPEG-4 Part 2 standard. This means that much less network bandwidth and storage space are required for a video file. Or seen another way, much higher video quality can be achieved for a given bit rate.

Jointly defined by standardization organizations in the telecommunications and IT industries, H.264 is expected to be more widely adopted than previous standards.

H.264 has already been introduced in new electronic gadgets such as mobile phones and digital video players, and has gained fast acceptance by end users. Service providers such as online video storage and telecommunications companies are also beginning to adopt H.264.

In the video surveillance industry, H.264 will most likely find the quickest traction in applications where there are demands for high frame rates and high resolution, such as in the surveillance of highways, airports and casinos, where the use of 30/25 (NTSC/PAL) frames per second is the norm. This is where the economies of reduced bandwidth and storage needs will deliver the biggest savings.

H.264 is also expected to accelerate the adoption of megapixel cameras since the highly efficient compression technology can reduce the large file sizes and bit rates generated without compromising image quality. There are tradeoffs, however. While H.264 provides savings in network bandwidth and storage costs, it will require higher performance network cameras and monitoring stations.

2. Development of H.264

H.264 is the result of a joint project between the ITU-T's Video Coding Experts Group and the ISO/IEC Moving Picture Experts Group (MPEG). ITU-T is the sector that coordinates telecommunications standards on behalf of the International Telecommunication Union. ISO stands for International Organization for Standardization and IEC stands for International Electrotechnical Commission, which oversees standards for all electrical, electronic and related technologies. H.264 is the name used by ITU-T, while ISO/IEC has named it MPEG-4 Part 10/AVC since it is presented as a new part in its MPEG-4 suite. The MPEG-4 suite includes, for example, MPEG-4 Part 2, which is a standard that has been used by IP-based video encoders and network cameras.

Designed to address several weaknesses in previous video compression standards, H.264 delivers on its goals of supporting:

> Implementations that deliver an average bit rate reduction of 50%, given a fixed video quality compared with any other video standard
> Error robustness so that transmission errors over various networks are tolerated
> Low latency capabilities and better quality for higher latency
> Straightforward syntax specification that simplifies implementations
> Exact match decoding, which defines exactly how numerical calculations are to be made by an encoder and a decoder to avoid errors from accumulating

H.264 also has the flexibility to support a wide variety of applications with very different bit rate requirements. For example, in entertainment video applications—which include broadcast, satellite, cable and DVD—H.264 will be able to deliver a performance of between 1 to 10 Mbit/s with high latency, while for telecom services, H.264 can deliver bit rates of below 1 Mbit/s with low latency.

3. How video compression works

Video compression is about reducing and removing redundant video data so that a digital video file can be effectively sent and stored. The process involves applying an algorithm to the source video to create a compressed file that is ready for transmission or storage. To play the compressed file, an inverse algorithm is applied to produce a video that shows virtually the same content as the original source video. The time it takes to compress, send, decompress and display a file is called latency. The more advanced the compression algorithm, the higher the latency, given the same processing power.
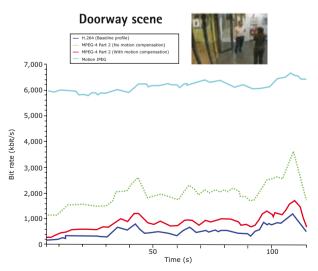
A pair of algorithms that works together is called a video codec (encoder/decoder). Video codecs that implement different standards are normally not compatible with each other; that is, video content that is compressed using one standard cannot be decompressed with a different standard. For instance, an MPEG-4 Part 2 decoder will not work with an H.264 encoder. This is simply because one algorithm cannot correctly decode the output from another algorithm but it is possible to implement many different algorithms in the same software or hardware, which would then enable multiple formats to be compressed.

Different video compression standards utilize different methods of reducing data, and hence, results differ in bit rate, quality and latency.

Results from encoders that use the same compression standard may also vary because the designer of an encoder can choose to implement different sets of tools defined by a standard. As long as the output of an encoder conforms to a standard's format and decoder, it is possible to make different implementations. This is advantageous because different implementations have different goals and budget. Professional non-real-time software encoders for mastering optical media should have the option of being able to deliver better encoded video than a real-time hardware encoder for video conferencing that is integrated in a hand-held device. A given standard, therefore, cannot guarantee a given bit rate or quality. Furthermore, the performance of a standard cannot be properly compared with other standards, or even other implementations of the same standard, without first defining how it is implemented.

A decoder, unlike an encoder, must implement all the required parts of a standard in order to decode a compliant bit stream. This is because a standard specifies exactly how a decompression algorithm should restore every bit of a compressed video.

The graph below provides a bit rate comparison, given the same level of image quality, among the following video standards: Motion JPEG, MPEG-4 Part 2 (no motion compensation), MPEG-4 Part 2 (with motion compensation) and H.264 (baseline profile).

Figure 1. *An H.264 encoder generated up to 50% fewer bits per second for a sample video sequence than an MPEG-4 encoder with motion compensation. The H.264 encoder was at least three times more efficient than an MPEG-4 encoder with no motion compensation and at least six times more efficient than Motion JPEG.*

4. H.264 profiles and levels

The joint group involved in defining H.264 focused on creating a simple and clean solution, limiting options and features to a minimum. An important aspect of the standard, as with other video standards, is providing the capabilities in profiles (sets of algorithmic features) and levels (performance classes) that optimally support popular productions and common formats.

H.264 has seven profiles, each targeting a specific class of applications. Each profile defines what feature set the encoder may use and limits the decoder implementation complexity.

Network cameras and video encoders will most likely use a profile called the baseline profile, which is intended primarily for applications with limited computing resources. The baseline profile is the most suitable given the available performance in a real-time encoder that is embedded in a network video product. The profile also enables low latency, which is an important requirement of surveillance video and also particularly important in enabling real-time, pan/tilt/zoom (PTZ) control in PTZ network cameras.

H.264 has 11 levels or degree of capability to limit performance, bandwidth and memory requirements. Each level defines the bit rate and the encoding rate in macroblock per second for resolutions ranging from QCIF to HDTV and beyond. The higher the resolution, the higher the level required.

5. Understanding frames

Depending on the H.264 profile, different types of frames such as I-frames, P-frames and B-frames, may be used by an encoder.

An I-frame, or intra frame, is a self-contained frame that can be independently decoded without any reference to other images. The first image in a video sequence is always an I-frame. I-frames are needed as starting points for new viewers or resynchronization points if the transmitted bit stream is damaged. I-frames can be used to implement fast-forward, rewind and other random access functions. An encoder will automatically insert I-frames at regular intervals or on demand if new clients are expected to join in viewing a stream. The drawback of I-frames is that they consume much more bits, but on the other hand, they do not generate many artifacts.

A P-frame, which stands for predictive inter frame, makes references to parts of earlier I and/or P frame(s) to code the frame. P-frames usually require fewer bits than I-frames, but a drawback is that they are very sensitive to transmission errors because of the complex dependency on earlier P and I reference frames.

A B-frame, or bi-predictive inter frame, is a frame that makes references to both an earlier reference frame and a future frame.
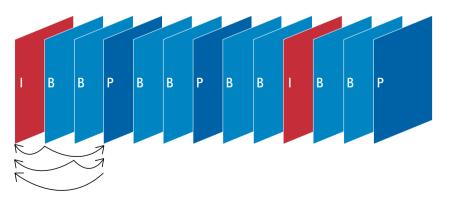

Figure 2.   *A typical sequence with I-, B- and P-frames. A P-frame may only reference preceding I- or P-frames, while a B-frame may reference both preceding and succeeding I- or P-frames.*

When a video decoder restores a video by decoding the bit stream frame by frame, decoding must always start with an I-frame. P-frames and B-frames, if used, must be decoded together with the reference frame(s).

In the H.264 baseline profile, only I- and P-frames are used. This profile is ideal for network cameras and video encoders since low latency is achieved because B-frames are not used.

6.   Basic methods of reducing data

A variety of methods can be used to reduce video data, both within an image frame and between a series of frames.

Within an image frame, data can be reduced simply by removing unnecessary information, which will have an impact on the image resolution.

In a series of frames, video data can be reduced by such methods as difference coding, which is used by most video compression standards including H.264. In difference coding, a frame is compared with a reference frame (i.e. earlier I- or P-frame) and only pixels that have changed with respect to the reference frame are coded. In this way, the number of pixel values that are coded and sent is reduced.

Figure 3.   *With Motion JPEG format, the three images in the above sequence are coded and sent as separate unique images (I-frames) with no dependencies on each other.*
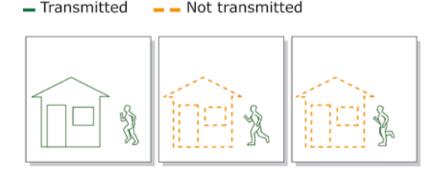

Figure 4.   *With difference coding (used in most video compression standards including H.264), only the first image (I-frame) is coded in its entirety. In the two following images (P-frames), references are made to the first picture for the static elements, i.e. the house, and only the moving parts, i.e. the running man, is coded using motion vectors, thus reducing the amount of information that is sent and stored.*

The amount of encoding can be further reduced if detection and encoding of differences is based on blocks of pixels (macroblocks) rather than individual pixels; therefore, bigger areas are compared and only blocks that are significantly different are coded. The overhead associated with indicating the location of areas to be changed is also reduced.

Difference coding, however, would not significantly reduce data if there was a lot of motion in a video. Here, techniques such as block-based motion compensation can be used. Block-based motion compensation takes into account that much of what makes up a new frame in a video sequence can be found in an earlier frame, but perhaps in a different location. This technique divides a frame into a series of macroblocks. Block by block, a new frame—for instance, a P-frame—can be composed or 'predicted' by looking for a matching block in a reference frame. If a match is found, the encoder simply codes the position where the matching block is to be found in the reference frame. Coding the motion vector, as it is called, takes up fewer bits than if the actual content of a block were to be coded.
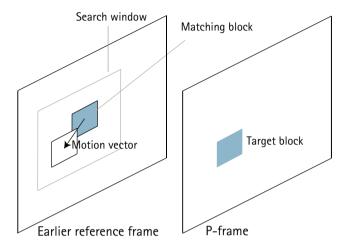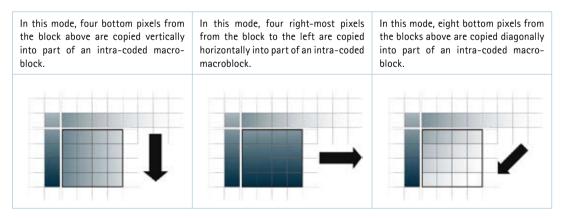

Figure 5. *Illustration of block-based motion compensation*

7. Efficiency of H.264

H.264 takes video compression technology to a new level. With H.264, a new and advanced intra prediction scheme is introduced for encoding I-frames. This scheme can greatly reduce the bit size of an I-frame and maintain a high quality by enabling the successive prediction of smaller blocks of pixels within each macroblock in a frame. This is done by trying to find matching pixels among the earlier-encoded pixels that border a new 4x4 pixel block to be intra-coded. By reusing pixel values that have already been encoded, the bit size can be drastically reduced. The new intraprediction is a key part of the H.264 technology that has proven to be very efficient. For comparison, if only I-frames were used in an H.264 stream, it would have a much smaller file size than a Motion JPEG stream, which uses only I-frames.
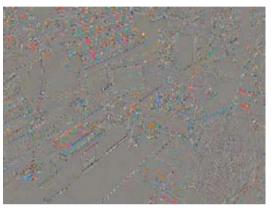
In this mode, four bottom pixels from the block above are copied vertically into part of an intra-coded macro-block.	In this mode, four right-most pixels from the block to the left are copied horizontally into part of an intra-coded macroblock.	In this mode, eight bottom pixels from the blocks above are copied diagonally into part of an intra-coded macro-block.

Figure 6. *Illustrations of some of the modes that intra prediction can take in coding 4x4 pixels within one of the 16 blocks that make up a macroblock. Each of the 16 blocks in a macroblock may be coded using different modes.*

Original source image

Intra predicted image

Residual image

Output image

Figure 7.   *The above images illustrate the efficiency of H.264's intra prediction scheme, whereby the intra predicted image is sent for "free". Only the residual content and the intra prediction modes need to be coded to produce the output image.*

Block-based motion compensation—used in encoding P- and B-frames—has also been improved in H.264. An H.264 encoder can choose to search for matching blocks—down to sub-pixel accuracy—in a few or many areas of one or several reference frames. The block size and shape can also be adjusted to improve a match. In areas where no matching blocks can be found in a reference frame, intra-coded macroblocks are used. The high degree of flexibility in H.264's block-based motion compensation pays off in crowded surveillance scenes where the quality can be maintained for demanding applications. Motion compensation is the most demanding aspect of a video encoder and the different ways and degrees with which it can be implemented by an H.264 encoder can have an impact on how efficiently video is compressed.

With H.264, typical blocky artifacts—seen in highly compressed video using Motion JPEG and MPEG standards other than H.264—can be reduced using an in-loop deblocking filter. This filter smoothes block edges using an adaptive strength to deliver an almost perfect decompressed video.

Figure 8. *Blocky artifacts in the highly compressed image at left are reduced when a deblocking filter is applied, as seen in the image at right.*

8. Conclusion

H.264 presents a huge step forward in video compression technology. It offers techniques that enable better compression efficiencies due to more accurate prediction capabilities, as well as improved resilience to errors. It provides new possibilities for creating better video encoders that enable higher quality video streams, higher frame rates and higher resolutions at maintained bit rates (compared with previous standards), or, conversely, the same quality video at lower bit rates.

H.264 represents the first time that the ITU, ISO and IEC have come together on a common, international standard for video compression. Due to its flexibility, H.264 has been applied in diverse areas such as high-definition DVD (e.g. Blu-ray), digital video broadcasting including high-definition TV, online video storage (e.g. YouTube), third-generation mobile telephony, in software such as QuickTime, Flash and Apple Computer's MacOS X operating system, and in home video game consoles such PlayStation 3. With support from many industries and applications for consumer and professional needs, H.264 is expected to replace other compression standards and methods in use today.

As the H.264 format becomes more broadly available in network cameras, video encoders and video management software, system designers and integrators will need to make sure that the products and vendors they choose support this new open standard. And for the time being, network video products that support both H.264 and Motion JPEG are ideal for maximum flexibility and integration possibilities.

31669/EN/R1/0803

About Axis Communications

Axis is an IT company offering network video solutions for professional installations. The company is the global market leader in network video, driving the ongoing shift from analog to digital video surveillance. Axis products and solutions focus on security surveillance and remote monitoring, and are based on innovative, open technology platforms.

Axis is a Swedish-based company, operating worldwide with offices in 18 countries and cooperating with partners in more than 70 countries. Founded in 1984, Axis is listed on the OMX Nordic Exchange, Large Cap and Information Technology. For more information about Axis, please visit our website at www.axis.com.