White Paper on Video Technologies for Mars Airplane

By John F. McGowan, Ph.D.
Desktop Video Expert Center
NASA Ames Research Center
Moffett Field, CA 94035

Introduction

This white paper discusses options for the low-resolution video system for the proposed Mars Airplane mission. This mission would send a small propeller driven airplane to Mars to fly down the Valles Marineris canyon for approximately 30 minutes. A number of scenarios have been proposed for the details of the mission. The airplane would carry two cameras, a high-resolution scientific camera to take a small number, tentatively 20, high-resolution pictures and a low-resolution camera to provide video.

Given the requirement of *television-quality full motion video* for the proposed Mars Airplane mission, the use of MPEG digital video was suggested, either MPEG-1 at 352 x 240 pixels at 30 frames per second, or MPEG-2 at 352 by 240 pixels at 30 frames per second in 4:2:0 video format, using the default quantization matrices, giving a bit rate of 1 Megabit per second.

MPEG-2 has some error control and recovery features not in the MPEG-1 standard at negligible additional bit-rate. Given the noisy deep space environment, the use of MPEG-2 at the MPEG-1 parameters – 352 by 240 by 30 frames per second in 4:2:0 format - may be preferable. This is known as MPEG-2 Main Profile at Low Level.

In the event that the requirement for full motion video is replaced with a slide show at low frame rates, it may be preferable to use the JPEG still image compression standard. At sufficiently low frame rates, there is no advantage and some disadvantages to using MPEG instead of JPEG still image compression. Many JPEG encoder chips exist as well as freely available encoder and decoder software. A discussion of the slide show option is included later in this white paper.

Initial Proposal

The initial proposal consisted of a low resolution NTSC color video camera using CCD technology and an MPEG encoder board with an MPEG encoder chip, some RAM buffers, and miscellaneous support chips.

To avoid confusion, the original proposal for the Mars Airplane low-resolution video system is included as an appendix, Appendix A, to this document. There were many educated guesses on the power, volume, and weight of this system in that document. The specifications for a commercial product, the Hitachi MPEG Camera that incorporates the components discussed has since been located. This provides more accurate estimates of the power, volume, and weight using commercial off the shelf (COTS) components.

Power, Weight, and Volume

The initial, very rough estimate, was:

TOTAL WEIGHT: Using consumer/PC circuit technology: Total weight should NOT exceed 1 KG (kilogram).
TOTAL SIZE: 100 mm by 200 mm by 10 mm
TOTAL POWER DISSIPATION: > 14.6 Watts (everything above added). Guess would be 15-20 Watts.

Please see Appendix A for details of this estimate.

The figures for the Hitachi MPEG Camera are:

TOTAL WEIGHT: 540 g (0.54 KG) including a battery pack and hard disk drive.
TOTAL SIZE: 83.8 mm by 142.2 mm by 55.9 mm (3.3" by 5.6" by 2.2")
TOTAL POWER DISSIPATION: 6.5 Watts

The Hitachi camera includes a number of components such as a microphone and LCD display that would not be needed on the Mars Airplane.

These weights do not incorporate the effects of space hardening. The simplest form of space hardening appears to be enclosing the components in a shielding box of aluminum that will absorb ionizing radiation, shield RF interference, and provide mechanical support and protection. A 100 mil (2.5 mm) thick shielding box for the video package would weigh about 300 grams (400 square centimeters, 100 mil aluminum sheet weights 0.68 grams/cm^2). A 400 mil (10 mm) thick shielding box for the video package would weigh about 1.2 Kilograms.

Reasons for Using MPEG

- Numerous verified (tested and debugged) real-time MPEG encoder and decoder chips and chip designs from multiple manufacturers exist. Philips, C-Cube Microsystems, IBM Microelectronics, and several other firms market real-time MPEG encoder chips. In many cases, these chips have been used in successful consumer and professional video products. SGS Thomson, for example, has shipped over 4 million MPEG-1 decoder chips. Selecting MPEG does not lock NASA into a single manufacturer or a single design.
- Successful products use MPEG at 352 by 240 pixels at 30 frames per second in 4:2:0 format using the default quantization matrices, meaning 1 Megabit per second, to provide television video on 13 inch and larger televisions. These products include the VideoCD family of video players and CD-I players. Thus a high level of confidence exists that MPEG at 352 by 240 by 30 frames per second, 4:2:0 format, at 1 Megabit/second can provide full motion video that human viewers will find acceptable.
- MPEG is good compression. The underlying compression algorithm outperforms most other image and video compression technologies. Few compression technologies surpass the performance of MPEG compression. The other technologies, notably the discrete wavelet transform, are not standardized. All working implementations are proprietary. Very few hardware implementations exist. Very few engineers and developers with the appropriate expertise exist.
- MPEG is preferable to the ITU-T H.261, H.263, and H.263+ videoconferencing standards because MPEG can achieve better compression of full motion video using bi-directionally predicted encoding of frames, known as B pictures. In video-conferencing the delay introduced by the encoding needs to be kept small, ideally less than 100 milliseconds, to achieve real-time face-to-face interaction. Because of this, certain encoding methods such as B pictures, which can achieve higher compression ratios, were not used in H.261. Since the Mars Airplane involves one-way transmission of video back to Earth, these constraints do not apply.
- MPEG is preferable to the ISO JPEG still image compression standard for full motion video because MPEG uses motion estimation and compensation to exploit the small differences between successive frames to achieve greater compression ratios, 50-70:1, than with JPEG, 15-30:1 for 352 by 240 pixel color frames where the 4:2:0 format is exploited to downsample the color components.
- C programming language implementations of MPEG encoders and decoders, both MPEG-1 and MPEG-2, are publicly available.
- Documents detailing the MPEG standard exactly are available from ISO for a fee, which although large for a private individual, would be tiny for a NASA development effort.
- Over a dozen in depth books on the MPEG standard have been published.
- A substantial community of hardware and software developers with experience implementing MPEG encoders and decoders exists.

- ♦ A substantial community of hardware and software developers with experience implementing solutions that use MPEG technology exists.

Reasons for Using 352 by 240 by 30 frames per second in 4:2:0 format

The Mars Airplane mission will have a very limited bandwidth from Mars to Earth. For example, 1 Megabit per second would exceed any bandwidth every achieved to Mars. To set the scale, uncompressed NTSC video is about 168 Megabits/second. MPEG-2 Main Profile Main Level 720 by 480 by 30 frames per second in 4:2:0 format is between 4 and 8 Megabits/second. The Mars Airplane will need the lowest bit-rate while retaining acceptable video quality.

NTSC analog video has a horizontal resolution equivalent to 352 pixels. A single field of NTSC analog video is 240 lines. Thus a 352 by 240 frame of video is equivalent to a single field of analog NTSC video. Repeating the frame can provide the other field. This is how the VideoCD players and other MPEG-1 consumer products work. 30 frames per second is the frame rate of NTSC analog video and provides smooth full-motion video for most natural scenes. The 4:2:0 format is matched to NTSC color which down-samples the color signal in an analog fashion. Of the 6 MHz used for a color NTSC signal, 4 MHz is the black and white signal and 2 MHz is the color signal.

Several successful products exist using these parameters. From both empirical and theoretical reasons there is a high level of confidence that these parameters provide acceptable video. They probably are close to or at the low end of acceptability.

Constant Bit-rate Operation

MPEG Encoders typically operate in two possible modes. Variable Bit-rate (VBR) encoding encodes at a constant video quality. With VBR scenes that are more difficult for MPEG to encode will have a higher bit-rate. Constant Bit-rate (CBR) encoding encodes at a constant bit-rate. With CBR scenes that are more difficult for MPEG to encode will have lower video quality. This means various visual artifacts such as blocking, blurring, and ringing may occur in scenes or parts of scenes that are harder for MPEG to encode.

Because of the tight bit-rate constraints of various parts of the communication systems on the Mars Airplane, it seems likely that the MPEG encoder should be run in Constant bit-rate mode. An MPEG encoder capable of CBR operation should be selected.

Mars Airplane Flight Plan

In current discussions, the Mars Airplane will deploy from the aeroshell somewhere in the Marineris Valles canyon. Then it will drop and fly toward the canyon wall. Typical airspeed will be 540 km/hour (150 m/second). The plane will bank as it approaches the canyon wall and fly parallel to the wall at a distance of 500 meters and an altitude of 500 meters. Because of the science requirements, the maximum distance from the canyon wall is 1000 meters. The canyon wall is reportedly 5 to 9 km in height. The canyon width varies. The widest part is 500 kilometers. In general, the far wall is expected to be several kilometers from the Mars Airplane.

Because MPEG exploits the small differences between successive frames to achieve high compression and encodes translation motion better than other kinds of motion – such as stretching, rotation, and skewing – the flight plan affects the compression performance.

Except during the banking maneuver, the motion will be horizontal and vertical translation and horizontal and vertical stretching. The stretching occurs because the distance from the camera to the objects imaged changes significantly during the flight.

The flight plan will not affect the performance of a still image compression method such as JPEG.

Low Resolution Camera Optics

For the purposes of this white paper, a lens that gives a horizontal angle of view of 90 degrees will be assumed. With MPEG-1 352 by 240 pixel spatial resolution, this means that a pixel corresponds to 0.25568 degrees.

MPEG and Camera Orientation

MPEG uses motion estimation and motion compensation to achieve very high compression ratios, much higher than possible with still image compression methods such as JPEG. MPEG can achieve compression ratios of 50-70 whereas JPEG can achieve ratios of 15-30 for 352 by 240 pixel color frames exploiting the 4:2:0 format in which the color components are downsampled. Motion estimation refers to the encoding stage. Motion compensation is the complementary decoding stage. MPEG motion estimation and motion compensation models translation across the image, such as an object flying across the image or the camera panning. MPEG motion estimation does not model and therefore cannot compress as well other types of motion such as rotation, stretching, or skewing of objects. Stretching will occur when the camera zooms in on an object of flies toward an object rapidly, something that will occur during the Mars Airplane mission. The near canyon wall and the canyon floor will approach the camera as the plane flies down the canyon.

Most movies and television, which MPEG was designed to compress, contain primarily translation. Objects rarely fly toward or away from the camera. Zooms are used infrequently in movies and television. MPEG does not compress scenes with zooming as well as scenes with translation.

With the camera pointed forward, there will be more stretching than with the camera pointed to the side, toward the canyon wall. The camera pointed forward seems to be preferable from aesthetic and technology demonstration points of view.

At 30 frames per second, the plane will travel 5 meters between two successive frames. With the camera pointed forward and using a 90 degrees horizontal angle of view, the closest object visible to the camera, the near canyon wall, will be 707 meters from the plane. The closest object moves 0.4 degrees between successive frames, 1-2 pixels. The worst possible stretching would be 1-2 pixels. In fact, most of the motion will be simple translation. Rough calculations suggest that a 16 by 16 pixel macroblock will stretch by less than 1 pixel. MPEG motion estimation will probably work well for 30 frames per second video. With 30 frames per second video, camera orientation is probably not an issue.

However, one option is to use a slide show at one frame per second. In this case, the camera orientation may be an issue.

At 1 frame per second, the plane will travel 150 meters between two successive frames. With the camera pointed forward, the closest visible object, the near canyon wall at 707 meters, will move about 11.9 degrees between successive frames, about 48 pixels. Rough calculations suggest that a 16 by 16 pixel macroblock, the relevant unit for motion estimation, will stretch by 4-5 pixels. Motion estimation may fail.

At 1 frame per second, with the camera pointed to the side, looking at the canyon wall, the stretching is expected to be about 2-3 pixels. Motion estimation may still fail. It is more likely to work with this camera orientation.

An option is to simulate the video for the Mars mission either by flying a plane in a canyon on Earth or by using a scale model and evaluate the performance of MPEG compression on this test video. A simple, inexpensive simulation would be to carry an NTSC video camera down a hall near one wall at a distance and speed scaled to imitate the plane in the canyon. Then encode the video using MPEG. The one frame per second situation could be emulated by sub-sampling the captured video before feeding it to the MPEG encoder.

A computer generated simulation might be useful. MPEG usually compresses computer generated imagery better than natural imagery. A computer generated animation of the plane flying on Mars will probably give too good an estimate of MPEG performance. The use of video from a plane flight in a real canyon or a scale model will probably provide a more accurate simulation than a computer generated animation.

Still image compression methods such as JPEG do not use motion estimation and will work equally well or poorly independent of camera orientation.

Reasons for Using a Hardware Video Encoder

The Mars Airplane will need to compress the low resolution video in real-time. There is no way to carry sufficient memory for 30 minutes of uncompressed video on the airplane. At 352 by 240 by 30 frames per second this is about 73 Megabits/second (color) or 24.3 Megabits/second (grayscale). Thirty minutes of uncompressed video would require 131.4 Gigabits (color) or 43.8 Gigabits (grayscale) of storage, presumably RAM, on the Mars Airplane. With 40-130 Gigabits of storage, the Airplane could then compress the video at a slower rate for transmission, assuming the Airplane survived landing. The uncompressed video rates also far exceed the 1-2 Megabit/second rates possible between the Mars Airplane and a local relay station. Thus, some kind of real-time compression of the video will be required.

Achieving high compression ratios requires processing power. In general, greater compression requires greater processing. Algorithms that can achieve higher compression ratios require more processing than algorithms that achieve lower compression ratios.

Real time encoding of MPEG-1 352 by 240 by 30 frames per second in 4:2:0 format requires 5,000 million instructions per second (MIPS). A general purpose CPU would require a clock speed of at least 5 GHz to implement MPEG-1 encoding in software in real-time. Another way of putting this is that a 1 GHz CPU, probably the current state of the art, could encode 6 frames per second using MPEG-1.

Real time encoding of MPEG-2 Main Profile at Main Level 720 by 480 by 30 frames per second in 4:2:0 format requires about 20,000 MIPS. A general purpose CPU would require a clock speed of at least 20 GHz to implement MPEG-2 Main Profile Main Level encoding in software in real-time. A 1 GHz CPU, probably the current state of the art, could encode 1 ½ frames per second using MPEG-2 Main Profile Main Level.

Video and image compression based on the Discrete Wavelet Transform (DWT) is similar to MPEG in complexity and processing requirements.

Regardless of video compression algorithm selected, the Mars Airplane will almost certainly require a hardware encoder for full motion video.

Alternative Compression Methods

Other International Standards

International Standards for still image and video compression from the International Organization for Standardization (ISO) and the International Telecommunications Union (ITU) share a most of the advantages of the MPEG standard.

The ITU-T H.261, H.263, and H.263+ videoconferencing standards are very similar to MPEG and have many of the same advantages and drawbacks. These standards use the block Discrete Cosine Transform and motion compensation as does MPEG. As mentioned in the **Reasons for MPEG** section, MPEG provides superior compression to these standards.

The ISO JPEG still image data compression standard shares many advantages with MPEG. Like MPEG, JPEG uses the block Discrete Cosine Transform for image compression. JPEG lacks the motion compensation and estimation that allow MPEG to achieve much higher compression of video. However, if the data is collected at low frame rates, instead of 30 frames per second, the advantage of MPEG may disappear. In this case, JPEG would be preferable to MPEG.

Most video and image compression technologies are comparable to or inferior to MPEG, H.26x, and JPEG in performance. These technologies are not standardized, are usually proprietary, and expertise is very hard to locate. There is no reason to use these technologies *unless they already exist as space qualified systems.*

Wavelet Compression

One technology outperforms the block Discrete Cosine Transform based international standards – MPEG, H.26x, and JPEG. This is video compression based on the Discrete Wavelet Transform. Several proprietary commercial products exist in software. These include Intel Indeo 5.10, VDONet’s VDOWave, VxTreme’s video codec, Summus’s video e-mail product, and Infinop’s Lightning Strike video. Studies indicate that Intel Indeo 5.1 and VDONet’s VDOWave outperform MPEG.

A large number of, usually black and white, still image wavelet compression prototypes in software, developed by researchers, exist. The most widely distributed is the Wavelet Image Construction Kit. This outperforms JPEG compression.

Wavelet video or image compression *might* achieve a factor of two improvement, at best, in the bit-rate required for the video over MPEG or JPEG. A wavelet-compressed video of similar quality to the 1 Megabit per second MPEG video *might* require only 500 Kilobits per second.

It is doubtful that wavelet video compression could be done in software. Such video compression probably incorporates something similar to motion estimation, which is computationally intensive. A wavelet video compressor has similar or greater processing requirements to an MPEG encoder. An MPEG encoder requires 5,000 MIPS for 352 by 240 by 30 frames per second in 4:2:0 format video, well beyond the capabilities of general purpose processors.

It is possible that wavelet still image compression, a slide-show option, could be done in software. As mentioned a large number of research prototypes exist that could be adapted to this purpose.

Color and Bit-rates

MPEG is color video. In 4:2:0 digital video, the color components are sub-sampled by 2 in the vertical and horizontal directions before encoding. This means that for a 352 by 240 frame, the luminance (black and white) component is not sub-sampled. 352 by 240 pixels are encoded. The two color components are sub-sampled, meaning only 176 by 120 pixels are encoded for each of the two components. This means that for every 4 black and white pixels, there are 2 color pixels. Naively, two thirds of the MPEG is black and white, not one third. Of the nominal 1 Megabit per second, 0.666 Megabits per second is the black and white signal.

In fact, this understates the compression of the color. MPEG uses a different quantization step size for the color signals, meaning the color is more compressed than the black and white. It is probable that 80 percent of the 1 Megabit per second of MPEG is the black and white signal.

Human vision has very low resolution for color. The black and white, or luminance, resolution is much higher. Most analog and digital video systems exploit this to reduce analog or digital bandwidth requirements. For example, analog color NTSC video uses 4 MHz for the black and white signal and 2 MHz for the color signal.

Thus, dropping the color signal buys little in bit-rate. Current MPEG encoder chips are quite sophisticated and some may support a non-standard black and white mode for some reduction in bit-rate.

Bit-rates and Video Quality

With all video compression technologies, there is a trade-off between bit-rate and perceived quality. In the technical video compression literature this is known as the rate-distortion function or curve. Generally, the lower the bit-rate, the greater the "distortion" or "error", the lower the "quality". At some point the quality becomes so poor the video is unacceptable to human viewers.

With MPEG, there are three main knobs that can be adjusted to change the bit-rate and quality. These are the spatial resolution – the width and height – of the frames, the frame rate, and the quantization step sizes, an overall quality factor.

*Bit-rate of compressed video does **not** scale linearly with spatial resolution and frame rate.* With uncompressed digital video, there is a simple linear relationship between bit-rate and spatial resolution, between bit-rate and frame rate. If the spatial resolution of *uncompressed video* is reduced by a factor of four, halving the width and halving the height of the frames, the bit-rate is quartered. If the frame rate of *uncompressed video* is halved, the bit-rate is halved. MPEG, like many compression schemes, exploits the correlation between adjacent pixels, that is nearby pixels tend to have similar or even identical values. With higher spatial resolution images of the same scene, the pixels tend to be more correlated. Greater compression is possible. As the spatial resolution of the same scene is reduced, the pixels correspond to features that are further apart in the scene, less correlated, and less compression is possible.

MPEG, like many video compression schemes, exploits the small differences between successive frames in a video sequence to achieve substantial compression. MPEG encodes only the change between successive frames. This allows MPEG to achieve compression ratios of 50 to 70 whereas still image encoding alone, such as JPEG, can achieve compression ratios of 15 to 30 for 352 by 240 pixel color frames using the 4:2:0 format in which the color components are downsampled. When frames are close together in time, as with 30 frame per second video, the difference between successive frames is usually small. The encoded change is tiny. However, when the frames are further apart in time, for example 1 frame per second video, the encoded change is larger. This means that the compression ratio that can be achieved becomes less. At some point the difference between successive frames is so large that the encoding becomes pointless. A still image compression scheme such as JPEG might as well be used.

With compressed video caution should be exercised in scaling bit-rates with frame rate and spatial resolution. The relationship is probably roughly linear for small changes such as scaling from 30 frames per second to 25 frames per second, or from 400 by 300 pixels to 360 by 270 pixels. For larger changes, such as extrapolating from 30 frames per second to one (1) frame per second, the prediction will probably be wrong.

A third parameter that can be adjusted is the quantization step size, actually step sizes. The larger the step size, generally the lower the bit-rate and the poorer the video quality. The relationship between the quantization step size and the bit-rate and the video quality is highly non-linear. With MPEG, as the quantization step size is increased, various artifacts such as blocking, blurring, and ringing begin to appear. Eventually the image becomes unrecognizable. Most other lossy compression schemes contain a quality factor or quality factors similar to the quantization step sizes in MPEG.

Frame Rate	**Subjective Quality**
1	Slide Show
5	Slide Show
10	Slide Show or Very Jerky
15	Jerky Video
16	According to Video Folklore, transition to smooth video. The early motion picture industry standardized on 16 frames per second, which seemed to provide smooth motion.
24	Modern Film Rate, smooth video, clearly acceptable
25	PAL (European TV) Rate, smooth video, clearly acceptable
30	NTSC (American TV) Rate, smooth video, clearly acceptable
60	Some High Definition Television Proposals, allegedly looks better than NTSC, clearly acceptable

Spatial Resolution	**Subjective Quality**
88 by 60 (also known as Sub QSIF or SQSIF)	very poor
176 by 120 (also known as QSIF)	mediocre
352 by 240 (also known as SIF)	“similar to VHS videotapes”
720 by 480 (also known as CCIR-601)	“Studio or broadcast quality”
1280 by 720 (from Grand Alliance High Definition Television standard)	“very high quality”
1920 by 1080 (from Grand Alliance High Definition Television standard)	“very very high quality”

Video Compression	**Spatial Resolution**	**Frame Rate**	**Bit-rate**	**Subjective Quality**
H.261	176 by 144	29.97 fps	128 Kbits/second	unacceptable, heavy blocking artifacts
H.261	176 by 144	29.97 fps	384 Kbits/second	allegedly acceptable for videoconferencing
MPEG-1	352 by 240	30	1 Megabit/second	acceptable
MPEG-2 Main Profile/Main Level	720 by 480	30	6 Megabit/second	good quality, acceptable

What Does MPEG Look Like Below 352 by 240 by 30 frames per second at 1 Megabit Per Second

The following frames from a video of the first launch of the Space Shuttle Columbia illustrate what MPEG will look like at lower bitrates.

384 Kbits/second

This is a frame at 160 by 120 (about a quarter of the proposed spatial resolution). This is an image from video of the first launch of the space shuttle Columbia. This is the best it would look at this spatial resolution. This is approximately what encoding 160 by 120 by 30 frames per second at 384 Kilobits /second using MPEG would look like. Keep in mind the target display is a 13” or larger NTSC television set.

256 Kilobits per second
Below is the same frame encoded for a bit-rate of 256 Kilobits per second. This is a 160 by 120 at 30 frames per second video.

128 Kilobits per second
Below is the same frame at 128 Kilobits per second.

64 Kilobits per second
Below is the same frame at 64 Kilobits per second.

Compression Artifacts and Scientific Analysis

In current discussions, the scientific analysis of the mission will use the twenty high resolution images and will not use the low resolution video. The high resolution images would be transmitted either uncompressed or using lossless compression. In lossless compression, the uncompressed image will be bit for bit identical with the original uncompressed image.

MPEG is a *lossy* compression technology. This means that the MPEG frames are not bit for bit identical with the original uncompressed frames. MPEG can introduce a variety of compression artifacts. These include blocking, also known as tiling. The blocking is easily visible in the example pictures in the preceding section. The other common artifact is known as ringing, the Gibbs effect, "mosquito noise", or contouring and typically occurs near sharp edges. Many other compression technologies share these artifacts or exhibit very similar types of artifacts. Sometimes the artifacts are obviously unnatural. *The MPEG artifacts could imitate sedimentation in the images.*

Wireless Communication Channel Issues

The radio communication link between Mars and Earth is expected to have a Bit Error Rate (BER) of 10^-6, meaning one bit error for every million bits transmitted. Compressed video, such as MPEG, is highly sensitive to single bit errors. In extreme cases a single bit error can severely corrupt a half second of video. There are several options to deal with this problem.

Forward Error Correction (FEC)

Both the Grand Alliance High Definition Television (GA-HDTV) and Digital Broadcast Satellite (DBS) use a combination of Trellis Coded Modulation (TCM) and Reed-Solomon codes to protect MPEG-2 streams from errors over wireless channels. DBS seems closer to the Mars Mission. In DBS, MPEG-2 is transmitted from a geostationary satellite in a 37,000 km orbit to an 18 inch receiver dish on Earth. Depending on mode, the error correction adds about 30 percent or about 100 percent to the bitrate.

If forward error correction similar to DBS was used to protect the video stream from channel errors this would probably increase the required bit-rate from 1 Megabit per second to 1.3 or 2 Megabits per second.

Cyclic Redundancy Check

If it is possible to store the video either in the landed Mars Airplane or in a relay station on Mars or orbiting Mars, then the video could be broken into packets with a cyclic redundancy check to detect errors. If errors were detected, then mission control could request retransmission of the corrupted packets.

Live with the Errors

MPEG-2 has some error control, recovery, and concealment features. Using MPEG-2 at Low Level (352 by 240 by 30 frames per second), it may be possible to transmit the video without any additional error protection. This seems very risky. This is not what is done with commercial broadcasts over wireless channels, such as DBS and GA-HDTV.

Some studies suggest that MPEG-2 can handle bit error rates of 10^-6.

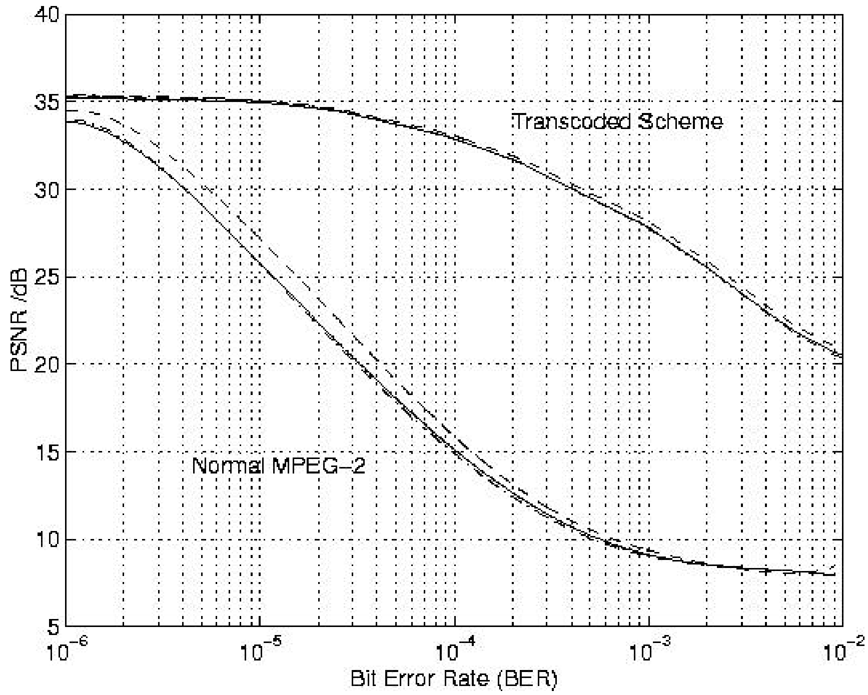

Note: A PSNR (Peak Signal to Noise Ratio) of 25 dB usually indicates poor image quality. A PSNR of 25 dB is usually as low as one can go for acceptable images. 30 dB is usually good.

The figure is from *Bandwidth Efficient Transmission of MPEG-II Video over Noisy Mobile Links* by Robert Swann and Nick Kingsbury, Signal Processing: Image Communications Special Issue on Mobile Image/Video Transmission, Vol. 12, No. 2 April 1998 pp.105-115. The Mars Airplane would seem to qualify as a Noisy Mobile Link.

RF Pickup Issues

The Mars Airplane will carry a radio transmitter to send the collected data back to Earth or to a relay station. This will produce around 15 Watts of radiated RF energy. Circuit boards such as the proposed video package can act as antennas picking up significant noise from the RF transmitter. This problem has occurred in cellular telephone designs, for example.

Enclosing the video package in a shielding box of aluminum may eliminate this problem.

This potential problem needs to be addressed by circuit board design engineers familiar with RF pickup and grounding issues.

Space Hardening Issues

Commercial off-the shelf MPEG chips are fabricated using commercial bulk CMOS semiconductor process technologies and packaged in lightweight plastic. COTS circuit boards are usually fabricated from an epoxy material known as FR-4. It is questionable whether these packages and circuit boards could survive the mechanical and other stresses of a deep space mission.

It is not clear that a commercial CMOS chip such as a standard MPEG encoder could survive the radiation from a trip from Earth to Mars. For example, solar flares and storms can produce bursts of 1 MeV protons and other ionizing radiation. A Mars Mission will reportedly deliver 10-20,000 rads Total Ionizing Dose(TID). This figure apparently assumes 100 mil of Aluminum shielding. The threshold for effects in commercial CMOS is somewhat greater than 500 rads. The radiation hardening literature suggests that

COTS CMOS parts vary significantly in radiation tolerance. Some can withstand 10-20,000 rads and others cannot.

It may be necessary to fabricate the chips in radiation hardened CMOS or SOS (Silicon on Sapphire). Radiation hardened processes reportedly lag behind the commercial bulk CMOS processes in speed and miniaturization. Chips fabricated in radiation hardened processes will probably be larger and slower and require more power. Fabricating an MPEG or other video encoder in a radiation hardened process is probably a big project fraught with risks.

The simplest option for space hardening of COTS electronics appears to be to enclose the circuit board in an aluminum box. The aluminum can shield the ionizing radiation, shield against RF interference, and provide mechanical support and protection for the circuit board. A 100 mil (2.5 mm) thick aluminum shielding box for the video package would weigh about 300 grams. A 400 mil (10 mm) thick aluminum shielding box would weigh 1.2 Kilograms.

A small heater, 5 W, may be required to keep the COTS components warm during the trip from Earth to Mars.

A more drastic option may be using circuit boards and packages made from rugged materials such as ceramic.

The most drastic option is fabricating the chips in a radiation hardened semiconductor process.

This potential problem needs to be addressed by engineers or scientists familiar with space hardening issues.

Relevant NASA Technology and Expertise

NASA has substantial resources and expertise in data, image, and video compression scattered across different centers. This may be especially important because space-qualified systems and radiation hardened chips may already exist.

The Jet Propulsion Laboratory appears to have substantial expertise in this area. There are reports that JPL was developing or developed a chip implementing the FBI’s wavelet image compression standard. Eric Majani of the Image Analysis Systems Group has developed E.R.I.C. (Efficient Reversible Image Compression), a wavelet based image compression system based on wavelets. E.R.I.C. is available in software and reportedly supports both lossless and lossy modes. Robert Rice developed the Rice algorithm, a lossless compression algorithm, at JPL in the 1970’s.

Malin Space Science Systems (MSSS), the contractor responsible for the Mars Global Surveyor camera, may have developed or be developing an appropriate camera and video processing system.

Relating to the high resolution camera, a group at Goddard Space Flight Center under the leadership of Dr. Pen-Shu Yeh has developed space-qualified image encoder and decoder chips that have been incorporated in several space missions. These chips implement lossless image compression based on the Rice algorithm. The technology reportedly can achieve about a factor of two lossless compression of image data.

Undoubtedly other expertise exists.

Scope of Fabricating an MPEG Encoder in a New Semiconductor Process

Although several detailed MPEG encoder chip designs in Verilog or VHDL, RTL (Register Transfer Level), and other representations used by VLSI chip designers now exist and NASA could probably license one from someone, a new semiconductor process technology such as radiation hardened CMOS will have different timing behavior. If signals arrive too early or too late, the encoder chip will fail.

New layout masks will have to be "synthesized" from the chip design, the new timings estimated based on the physical length of the pathways in the chip, and the timings fed back into simulations of the chip. Many simulations need to be run to find any cases where the timing fails and the chip fails. This is a substantial project. An educated guess would be twenty (20) design and verification engineers working for six months, followed by fabrication of a chip, which still might not work. A conservative estimate would be to plan for at least one more iteration: twenty engineers for six more months and another fabrication of a chip, hopefully successful.

20 FTE at 125 K/year for one year is $2.5 million.
Unknown fees from the semiconductor fabrication facility for each chip. Guess $500 K

This is about a $3 million project, starting with a working verified MPEG encoder chip design used for a working commercial CMOS product.

Total Cost: $3 million
Total Schedule: 1 year
Deliverable: Radiation Hardened MPEG Encoder Chip

Scope of Developing an MPEG Encoder Chip

This is a very large and risky project, to be avoided.

Detailed documentation of the MPEG standard and C language "behavioral models" are available for verifying a design. There is minimal information available on timing issues in an MPEG encoder or decoder.

A Verilog or VHDL design which matches the C language behavioral models from ISO (educated guess again) probably take 20 VLSI design engineers a year to develop followed by a timing verification and fabrication phase similar to that discussed in **Scope of Fabricating an MPEG Encoder Chip in a New Semiconductor Process**.

20 FTE at 125 K /year for a year (design phase).
20 FTE at 125 K/year for a year (verification phase)
Two design turns at $250 K each, gives $500 K (wild guess)

Total cost: $5.5 million
Total Schedule: two years.
Deliverable: Radiation Hardened MPEG Encoder Chip

Scope of Developing a Hardware Video Encoder with Custom Compression Algorithm

By definition, a custom compression algorithm would lack the documentation and C language prototypes available for MPEG (and also the other International Standards H.261, H.263, H.263+, and JPEG). These would need to be developed.

1 Algorithm Developer at 150K/year for a year (this is a very senior person)
1 Algorithm Implementor in C at 125 K/year for a year

This assumes that the custom compression algorithm uses "off the shelf" algorithm components such as vector quantization or the block Discrete Cosine Transform. There is no basic research into better compression involved. For example, this could be a wavelet based video encoder using the wavelet results in the published literature. If basic research is involved, the time, cost, and risk level will be much higher.

Total Cost: 275K
Total Schedule: 1 year
Deliverables: C language behavioral models of the video encoder and decoder.
Detailed written description of the algorithm for the chip designers, similar to a ISO standard document.

Then, the costs are similar to creating an MPEG encoder chip.

A Verilog or VHDL design which matches the C language behavioral models (educated guess again) probably takes 20 VLSI design engineers a year to develop followed by a timing verification and fabrication phase similar to that discussed in **Scope of Fabricating an MPEG Encoder Chip in a New Semiconductor Process**.

20 FTE at 125 K /year for a year (design phase).
20 FTE at 125 K/year for a year (verification phase)
Two design turns at $250 K each, gives $500 K (wild guess)

Total Cost: $5.775 million
Total Schedule: three years.
Deliverables: Video Encoder Chip (either commercial CMOS or radiation hardened process)

Scope of Developing Custom Compression Software

As discussed earlier, real-time encoding of full-motion video in software is almost certainly impossible, even on 1 GHz processors. However, if a slide show option was selected, such as one frame per second, various software options may be possible.

1 Algorithm Developer at 150K/year for a year (this is a very senior person)
1 Algorithm Implementor in C at 125 K/year for a year

Total Cost: 275K
Total Schedule: 1 year
Deliverable: Portable ANSI C encoder and decoder software.

Slide Show Option

Given the very low bit-rates that appear possible between Mars and Earth and the limited resources available for local storage and relay on Mars at the time of the mission, sending full motion video back to Earth may be impractical. Another option is a slide-show, a very low frame rate such as 1 frame per second. Very low bitrate full motion video (below 1 Megabit per second) is very poor in quality. Thus, a slide show of higher quality images would probably be preferable.

A system for a slide show would be very similar or identical to the proposal. Power, weight, and volume would be the same. The primary difference would be reduced bit-rate. It may be possible to do still image compression for the slide show in software, providing greater flexibility.

In this case, time lapse photography might be exploited. The frames could be played back on Earth at 30 frames per second to provide an accelerated video of the plane flying down the canyon. At one frame per second, a thirty minute mission (30) would provide one minute of video. This would only look appealing if the differences between successive frames are relatively small. This depends on the speed of the airplane and how far it is from the objects being imaged.

Bit-rates for Slide Show Using JPEG Still Image Compression

Slide Show Parameters	Compression Ratio	Bit-rate
352 by 240 at 1 frame per second	5:1	400 Kilobits/second
352 by 240 at 1 frame per second	10:1	200 Kilobits/second
352 by 240 at 1 frame per second	20:1	100 Kilobits/second
352 by 240 at 1 frame per second	30:1	66 Kilobits/second
352 by 240 at 1 frame every 2 seconds	5:1	200 Kilobits/second
352 by 240 at 1 frame every 2 seconds	10:1	100 Kilobits/second
352 by 240 at 1 frame every 2 seconds	20:1	50 Kilobits/second
352 by 240 at 1 frame every 2 seconds	30:1	33 Kilobits/second

If the differences between successive frames in the slide show are sufficiently small, then the motion estimation in MPEG may be exploited to achieve greater compression. This depends on how fast the Mars Airplane is flying and how far it is from the objects being imaged. For example, if the Mars Airplane flies at 500 kilometers/hour, then it travels about 140 meters in one second. If the plane is only 100 meters from the objects being imaged, such as the canyon wall, the change between successive frames in the slide show will be quite large. The successive frames may be completely different. The motion estimation will fail to achieve significant improvement in the compression over still image compression. If the plane is 10 km from the object being imaged, the two frames will look very similar. The motion estimation will improve compression significantly over still image compression. The following numbers are rough, representing a best case, since MPEG bit-rates do not scale exactly linearly with frame rate.

The **MPEG and Camera Orientation** section contains a more detailed discussion of the effects of one frame per second video on MPEG encoding.

Bit-rate for Slide Show Using MPEG Compression (Assuming Small Differences in Successive Frames)

Slide Show Parameters	Compression	Bit-rate
352 by 240 at 1 frame per second	Standard MPEG (1 Megabit/second for 30 fps)	33.3 Kilobits/second
352 by 240 at 1 frame every 2 seconds	Standard MPEG (1 Megabit/second for 30 fps)	16.65 Kilobits/second

Open Issues

Although this white paper tries to address a number of important issues, several issues remain unresolved and require further study by experts in relevant areas.

1. Space hardening issues are unresolved. Ideally an engineer or scientist familiar with space hardening issues should contribute to the design.
2. The scope of such possible projects as fabricating an MPEG encoder in a radiation hardened process is an unresolved issue. The author is not a VLSI chip designer. These estimates are educated guesses only.
3. RF pickup issues for the video package are unresolved. Ideally a circuit board designer familiar with RF and grounding issues should contribute to the design.

4. The effects of the noisy channel, the 10^-6 bit error rate, on the mission are unresolved. A communications expert familiar with noisy channels and forward error correction should contribute to the design.
5. The compression advantage of MPEG at the slide show rate of one frame per second (1 fps) is not clear. This should be simulated with video either of a plane in a canyon on Earth or a scale model of the Mars canyon.
6. At low frame rates, such as 1 frame per second, the camera orientation may significantly affect the performance of MPEG or other video compression methods. A camera pointed to the side may perform better than a forward pointing camera. A simulation of the video as in point 5 can resolve this issue.

Appendix A: Initial Proposal

Mars Airplane Video Package

by John F. McGowan, Ph.D.
Desktop Video Expert Center

If space hardening, power, weight, and other space-specific restrictions were not an issue, then I would propose a system based on a video camera and MPEG encoding. MPEG can achieve about 1 Megabits per second bandwidth. As an exercise I've tried to estimate the size, weight, and power requirements of such a system. Given the short time constraint for preparing the proposal there is a lot of educated guesswork in these numbers.

Note that a different system, not based on the widely used MPEG standard, would have many of the components listed below. Most encoders would need some sort of encoder buffer. All systems need a video camera. All systems will need a video system clock from somewhere.

Option A

- NTSC Video Camera Chip (CCD based)
- Lens and camera housing (on outside of Mars Airplane)
- Port/connector to carry analog NTSC from video camera to a video decoder.
- Video decoder (converts the analog NTSC video to 8-bit CCIR-601 digital video)
- Single Chip Real-Time MPEG Encoder
- Clock crystal to provide timing for video package.
- External DRAM or SDRAM for encoder buffers for the MPEG encoder chip.
- Boot ROM with microcode for the Single Chip Real-Time MPEG Encoder
- Printed Circuit Board for components.
- Mechanical supports and housing for the printed circuit board.
- Connector/data and control port to rest of the system.

CCD Camera Chip

Unknown
Maximum Power Dissipation: 2 Watts (????)
Dimensions: 4.8 mm by 3.6 mm by (1mm ????)

Video Decoder

Philips Semiconductors SAA7111 Video Input Processor

Maximum Power Dissipation: 1.26 Watts
Dimensions:14 by 14 by 2.7 mm (Plastic Quad Flat Package Version)

Single Chip MPEG Encoder

Philips Semiconductor SAA6750H Encoder for MPEG2 image recording (EMPIRE)
Maximum Power Dissipation: 2.0 Watts
Dimensions: 28 by 28 by 3.4 mm (Plastic Quad Flat Package Version)

External Encoder Buffer Memory

Philips SAA6750H MPEG Encoder uses four (4) 4 Megabit DRAM chips
Maximum Power Dissipation: 8.0 Watts (4 x 2 Watts – total guess)
Dimensions: ????

Boot ROM for Single Chip MPEG Encoder

Maximum Power Dissipation: 1.0 Watt?
Dimensions: ????

Printed Circuit Board

Dimensions: guess about 4" by 8" or 100 mm by 200 mm., by 10 mm with chips on the board using surface mount technology.

TOTAL WEIGHT: Using consumer/PC circuit technology: Total weight should NOT exceed 1 KG (kilogram).
TOTAL SIZE: 100 mm by 200 mm by 10 mm
TOTAL POWER DISSIPATION: > 14.6 Watts (everything above added). Guess would be 15-20 Watts.

Note that the chips used are consumer multimedia chips implemented using 0.27 CMOS process technology with a temperature rating of 0 to 70 degrees Centigrade. They need either 3.3V or 5 V to operate. They are packaged in plastic. I would be astounded if they worked on a Mars mission.

Option B

- Single Chip CCD camera with digital output (CCIR-601, RS 422??)
- Lens and camera housing (on outside of Mars Airplane)
- Port/connector to carry uncompressed digital video from CCD camera to MPEG encoder.
- Single Chip Real-Time MPEG Encoder
- Clock crystal to provide timing for video package.
- External DRAM or SDRAM for encoder buffers for the MPEG encoder chip.
- Boot ROM with microcode for the Single Chip Real-Time MPEG Encoder
- Printed Circuit Board for components.
- Mechanical supports and housing for the printed circuit board.
- Connector/data and control port to rest of the system.

Single Chip CCD Camera with Digital Video Output

Unknown
Maximum Power Dissipation: 2 Watts (????)
Dimensions: 4.8 mm by 3.6 mm by (1mm ????)

Single Chip MPEG Encoder

Philips Semiconductor SAA6750H Encoder for MPEG2 image recording (EMPIRE)
Maximum Power Dissipation: 2.0 Watts
Dimensions: 28 by 28 by 3.4 mm (Plastic Quad Flat Package Version)

External Encoder Buffer Memory

Philips SAA6750H MPEG Encoder uses four (4) 4 Megabit DRAM chips
Maximum Power Dissipation: 8.0 Watts (4 x 2 Watts – total guess)
Dimensions: ????

Boot ROM for Single Chip MPEG Encoder

Maximum Power Dissipation: 1.0 Watt?
Dimensions: ????

Printed Circuit Board

Dimensions: guess about 4" by 8" or 100 mm by 200 mm., by 10 mm with chips on the board using surface mount technology.

TOTAL WEIGHT: Using consumer/PC circuit technology: Total weight should NOT exceed 1 KG (kilogram). Somewhat less than Option A because Video Decoder chip omitted.
TOTAL SIZE: 100 mm by 200 mm by 10 mm
TOTAL POWER DISSIPATION: > 13.4 Watts (everything above added). Guess would be 15-20 Watts.

Note that the chips used are consumer multimedia chips implemented using 0.27 CMOS process technology with a temperature rating of 0 to 70 degrees Centigrade. They need either 3.3V or 5 V to operate. They are packaged in plastic. I would be astounded if they worked on a Mars mission.

Other Technical Issues:

1. RF pickup. The RF subsystem may generate pickup electrical signals in the video package. The video package circuit board may act as an antenna receiving RF signals from the radio transmitter in the Mars Airplane.
2. Mechanical supports and housing for the camera and video package circuit board would probably have to be designed and tested.
3. Can 0.25 – 0.5 micron CMOS survive in space environment? I don't know. Virtually all consumer multimedia technology is CMOS.
4. What type of packages, leads, and circuit boards can survive in space?
5. A possible option is to license a verified MPEG encoder design and fabricate the encoder in a semiconductor process technology that can survive space. This is fraught with risks since the timing behavior of the design will probably change with a new process technology. MPEG is very fragile. However, this is an option and would be preferable to designing an MPEG encoder from scratch.

Organizational Issues

1. The video package would require at least one printed circuit board designer with experience building systems to send into space and with RF and grounding issues encountered in radio communication systems, ideally in space. Even if an off-the-shelf video package could be identified, a designer would be needed to test, evaluate, and validate the selection.
2. Most likely a circuit board would need to be designed and built for the video package.
3. The video package would require a mechanical engineer to design and build, or evaluate the mechanical supports and housings. Again this would probably require specialized experience with space systems since consumer devices use plastics that probably would not survive a deep space mission.

NOTE: 1 inch = 25 mm

Appendix B

Summary of Parameters

Parameter	**Value**
Airspeed	150 meters/second
Mars Airplane Altitude	500 meters
Distance to Near Canyon Wall	500 meters
Height of Canyon Walls	5-9 kilometers
Width of Canyon	< 500 kilometers
Camera Angle of View (Horizontal)	90 degrees
Bit Error Rate (BER)	10^-6
Overhead for Forward Error Correction	30 or 100 %
Mars to Earth Bit Rate	8.25 Kilobits/second
MPEG-1 352 by 240 by 30 frames per second	1 Megabit/second
MPEG-2 720 by 480 by 30 frames per second	4-8 Megabits/second
Uncompressed 352 by 240 by 30 frames per second video (RGB)	60.8 Megabits/second
Uncompressed 720 by 480 by 30 frames per second video (RGB)	248.8 Megabits/second
Uncompressed Size of One 352 by 240 Color Frame	2,027,520 bits
Storage Requirement for 30 minutes of MPEG-1 video	1.8 Gigabits
Storage Requirements for 30 minutes of Uncompressed 352 by 240 by 30 frames per second video	109.4 Gigabits
Size of JPEG Compressed 352 by 240 Color Frame (Using 5:1 Compression, Excellent Quality)	405,504 bits
Size of JPEG Compressed 352 by 240 Color Frame (Using 10:1 Compression, Excellent Quality)	202,752 bits
Size of JPEG Compressed 352 by 240 Color Frame (Using 20:1 Compression, Good Quality)	101,376 bits
Size of JPEG Compressed 352 by 240 Color Frame (Using 30:1 Compression, Fair-Poor Quality)	66,000 bits
Processing for MPEG-1 Encoder	5,000 MIPS
Processing for MPEG-2 Encoder	20,000 MIPS
Typical Minimum Acceptable Peak Signal to Noise Ratio (PSNR) for Video or Images	25 dB

Typical Peak Signal to Noise Ratio (PSNR) for Good Quality Video or Images	30 dB or better
Total Ionizing Dose (TID) Earth to Mars (assuming 100 mil Aluminum shielding?)	10-20,000 rads
Threshold for Total Ionizing Dose (TID) effects in commercial CMOS	500 rads
Weight of 100 mil Aluminum Shielding	0.68 grams/cm^2
Weight of 400 mil Aluminum Shielding	2.72 grams/cm^2