Video compression

MPEG ENCODING

Reference: http://www.cs.cf.ac.uk/Dave/Multimedia/PDF/12_MPEG.pdf by Prof David Marshall and Dr Kirill Sidorov

Video compression

We need to compress video (more so than audio/images) in practice since:

1. Uncompressed video data are huge. In HDTV, the bit rate easily exceeds 1 Gbps — big problems for storage and network communications.
   E.g. HDTV: 1920 x 1080 at 30 frames per second, 8 bits per YCbCr (PAL) channel = 1.5 Gbps.
2. Lossy methods have to be employed since the compression ratio of lossless methods (e.g. Huffman, Arithmetic, LZW) is not high enough for image and video compression.

Video Compression: MPEG

- We first concentrate on some basic principles of video compression:
  - Earlier H.261 and MPEG 1 and 2 standards.
- with a brief introduction of ideas used in new standards such as H.264 (MPEG-4 Advanced Video Coding).
- Image, video, and audio compression standards have been specified and released by two main groups since 1985:
  - ISO International Standards Organisation: JPEG, MPEG.
  - ITU International Telecommunications Union: H.261–264.

Compression Standards

Whilst in many cases one of the groups have specified separate standards there is some **crossover** between the groups. E.g .:

- JPEG issued by ISO in 1989 (but adopted by ITU as ITU T.81)
- MPEG 1 released by ISO in 1991,
- H.261 released by ITU in 1993 (based on CCITT 1990 draft). CCITT stands for Comit´e Consultatif International T´el´ephonique et T´el´egraphique whose parent organisation is ITU.
- H.262 (better known as MPEG 2) released in 1994.
- H.263 released in 1996 extended as H.263+, H.263++.
- MPEG 4 released in 1998.
- H.264 releases in 2002 to lower the bit rates with comparable quality video and support wide range of bit rates, and is now part of MPEG 4 (Part 10, or AVC – Advanced Video Coding).

How to Compress Video?

Basic Idea of Video Compression:

Exploit the fact that adjacent frames are similar.

Spatial redundancy removal — intraframe coding (JPEG)

NOT ENOUGH BY ITSELF?

Temporal — greater compression by noting the temporal coherence/incoherence over frames. Essentially we note the difference between frames.

Spatial and temporal redundancy removal — intraframe and interframe coding (H.261, MPEG).

Things are much more complex in practice of course.

How to Compress Video?

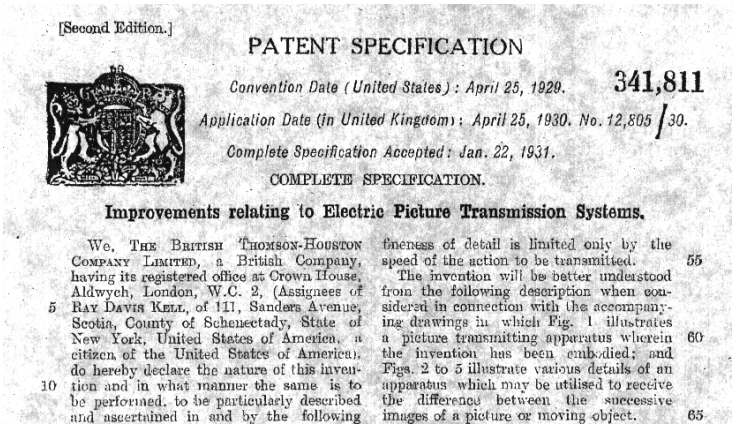

[Second Edition.]

PATENT SPECIFICATION

Convention Date (United States): April 25, 1929. 341,811

Application Date (in United Kingdom): April 25, 1930. No. 12,805/30.

Complete Specification Accepted: Jan. 22, 1931.

COMPLETE SPECIFICATION.

Improvements relating to Electric Picture Transmission Systems.

We, THE BRITISH THOMSON-HOUSTON COMPANY LIMITED, a British Company, having its registered office at Crown House, Aldwych, London, W.C. 2, (Assignees of RAY DAVIS KELL, of 111, Sanders Avenue, Scotia, County of Schenectady, State of New York, United States of America, a citizen of the United States of America), do hereby declare the nature of this invention and in what manner the same is to be performed, to be particularly described and ascertained in and by the following

fineness of detail is limited only by the speed of the action to be transmitted.

The invention will be better understood from the following description when considered in connection with the accompanying drawings in which Fig. 1 illustrates a picture transmitting apparatus wherein the invention has been embodied; and Figs. 2 to 5 illustrate various details of an apparatus which may be utilised to receive the difference between the successive images of a picture or moving object.

"It has been customary in the past to transmit successive complete images of the transmitted picture." . . . "In accordance with this invention, this difficulty is avoided by transmitting only the difference between successive images of the object."

Simple Motion Example

Consider a simple image of a moving circle. Lets just consider the difference between 2 frames. It is simple to encode/decode:

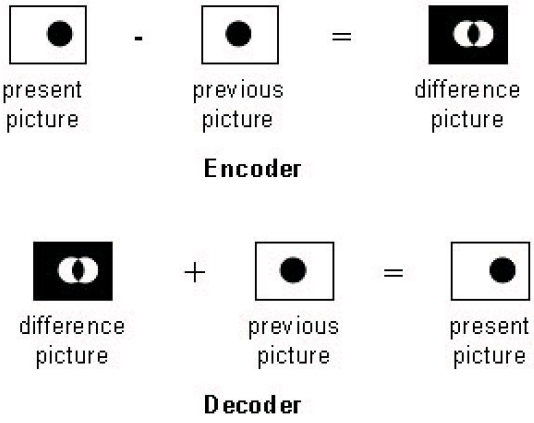

Estimating Motion of Blocks

We will examine methods of estimating motion vectors in due course.

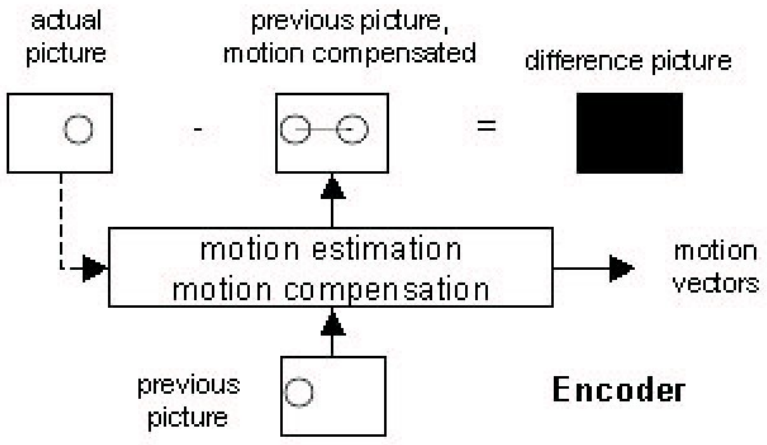

Decoding Motion of Blocks

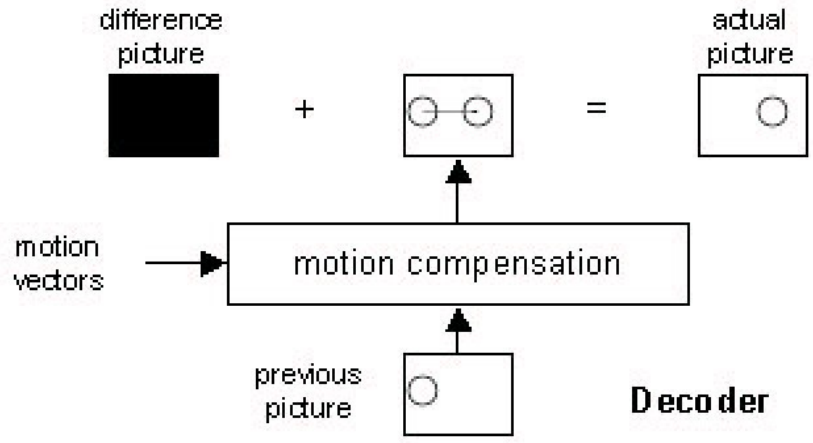

Why is this a better method than just frame differencing?

Motion Estimation Example

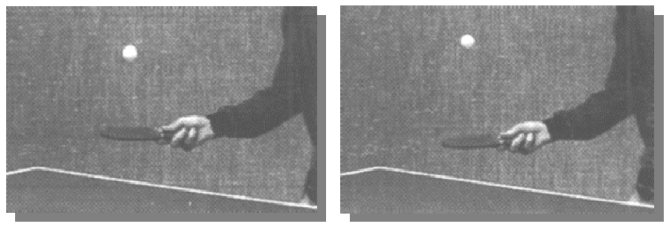

Frame N

Frame N+1

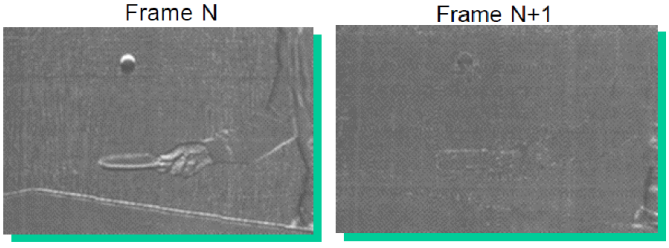

Difference Frame
Without Motion Prediction

Difference Frame
With Motion Prediction

How is Motion Compensation Used?

Block Matching:

MPEG-1/H.261 relies on block matching techniques.

For a certain area (block) of pixels in a picture:

Find a good estimate of this area in a previous (or in a future!) frame, within a specified search area.

Motion compensation:

Uses the motion vectors to compensate the picture.

Parts of a previous (or future) picture can be reused in a subsequent picture.

Individual parts spatially compressed — JPEG type compression.

Any Overheads?

- Motion estimation/compensation techniques reduces the video bitrate significantly
  but
- Introduce extra computational complexity.
  - Decoder needs to buffer reference pictures — backward and forward referencing.
  - Delay.

Lets see how such ideas are used in practice.

Overview of H.261

- Developed by CCITT in 1988-1990 for video telecommunication applications.
- Meant for videoconferencing, videotelephone applications over ISDN telephone lines. Baseline ISDN is 64 kbits/sec, and integral multiples (p×64).
- Frame types are CCIR 601 CIF (Common Intermediate Format) (352x288) and QCIF (176x144) images with 4:2:0 subsampling.
- **Two frame types:**
  - Intraframes (I-frames) and Interframes (P-frames).
- I-frames use basically JPEG — YUV (YCrCb) and larger DCT windows, different quantisation. I-frames provide us with a refresh accessing point — key frames.
- P-frames exploit differences from previous frame (predicted), so frames depend on each other.

H.261 Group of Pictures

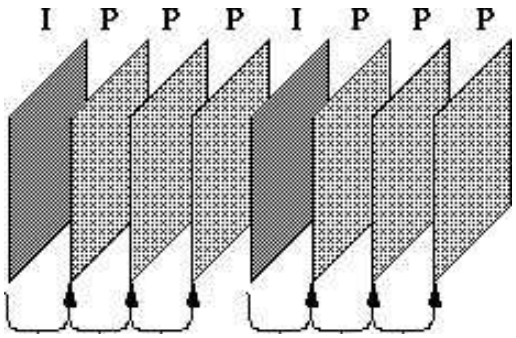

Why this cannot be too large?

We typically have a group of pictures — one I-frame followed by several P-frames — a group of pictures.

Number of P-frames followed by each I-frame determines the size of GOP — can be fixed or dynamic.

Intra-frame Coding

- Various lossless and lossy compression techniques use — like JPEG.
- Compression contained only within the current frame.
- Simpler coding — not enough by itself for high compression.
- Can't rely on intra frame coding alone not enough compression
- However, can't rely on inter frame differences across a large number of frames
  - So when errors get too large — start a new I-frame.

Intra-frame Coding (Cont.)

Intra-frame coding is very similar to JPEG:

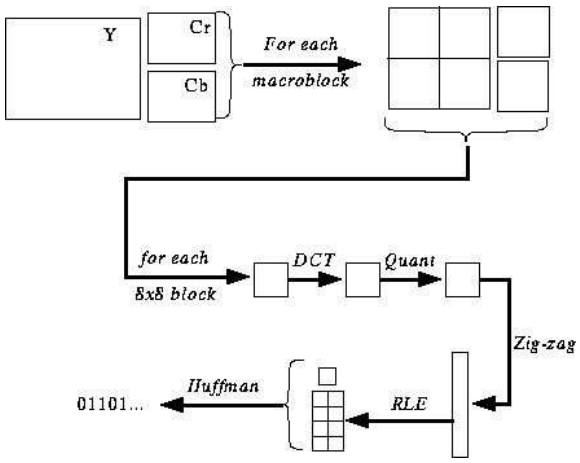

Intra-frame Coding (Cont.)

A basic intra-frame coding scheme is as follows:

- Macroblocks are typically 16x16 pixel areas on Y plane of original image.
- A macroblock usually consists of 4 Y blocks, 1 Cr block, and 1 Cb block. (4:2:0 chroma subsampling)
  - Eye most sensitive to luminance, less sensitive to chrominance.
  - We operate on a more effective color space: YUV (YCbCr) colour which we studied earlier.
  - Typical to use 4:2:0 macroblocks: one quarter of the chrominance information used.
- Quantization is by constant value for all DCT coefficients.
- I.e., no quantization table as in JPEG.

What's 4:2:0?

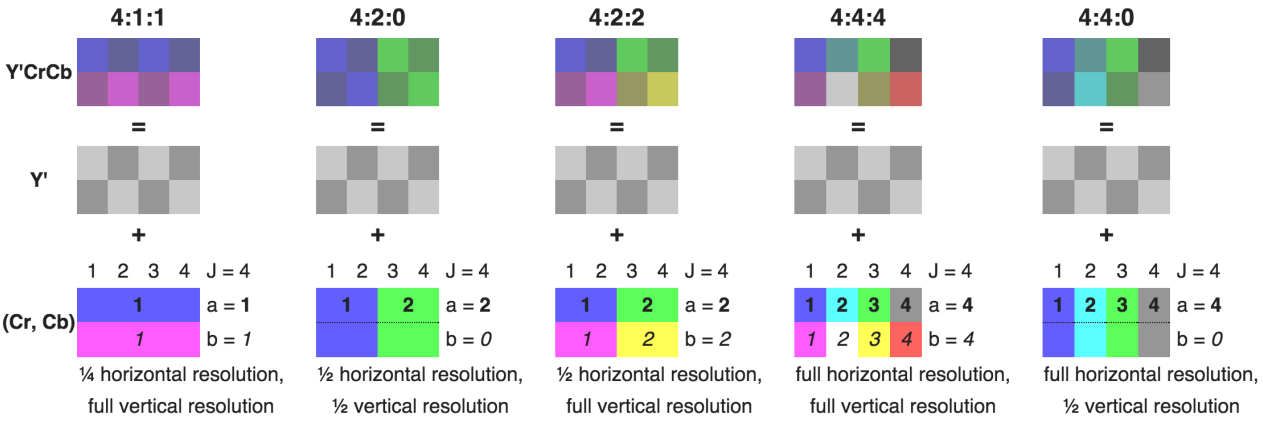

Inter-frame (P-frame) Coding

- Intra frame limited to spatial basis relative to 1 frame. Considerably more compression if the inherent temporal basis is exploited as well.

BASIC IDEA:

- Most consecutive frames within a sequence are very similar to the frames both before (and after) the frame of interest.
- Aim to exploit this redundancy.
- Use a technique known as block-based motion compensated prediction.
- Need to use motion estimation.

Inter-frame (P-frame) Coding (Cont.)

P-coding can be summarised as follows:

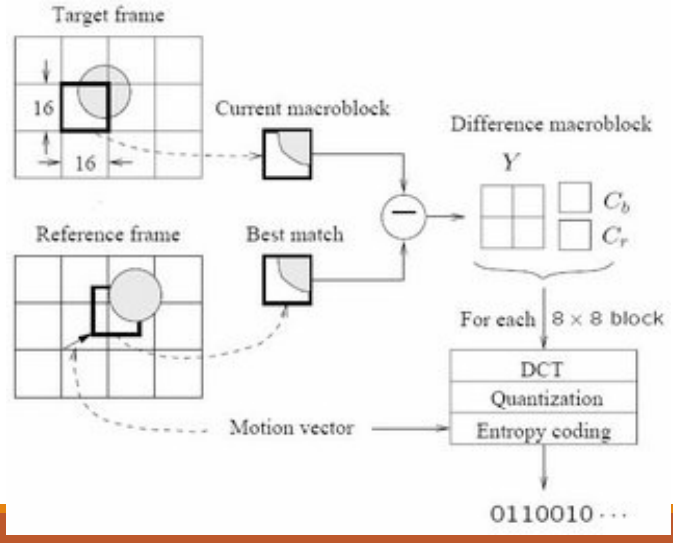

Inter-frame (P-frame) Coding (Cont.)

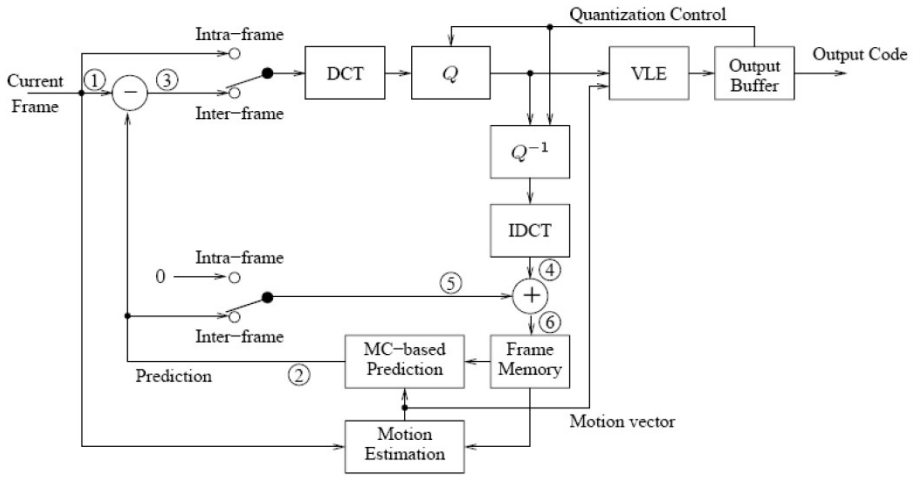

(a) Encoder

Inter-frame (P-frame) Coding (Cont.)

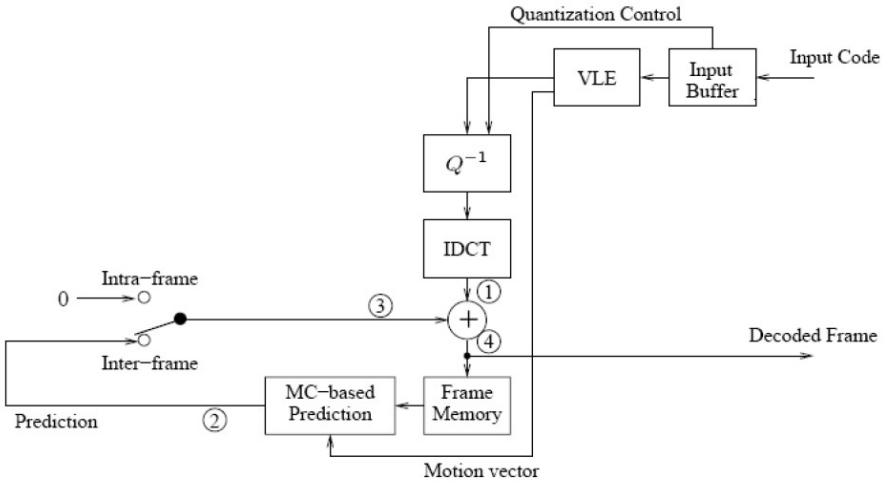

(b) Decoder

Motion Vector Search

So we know how to encode a P-block. How do we find the motion vector?

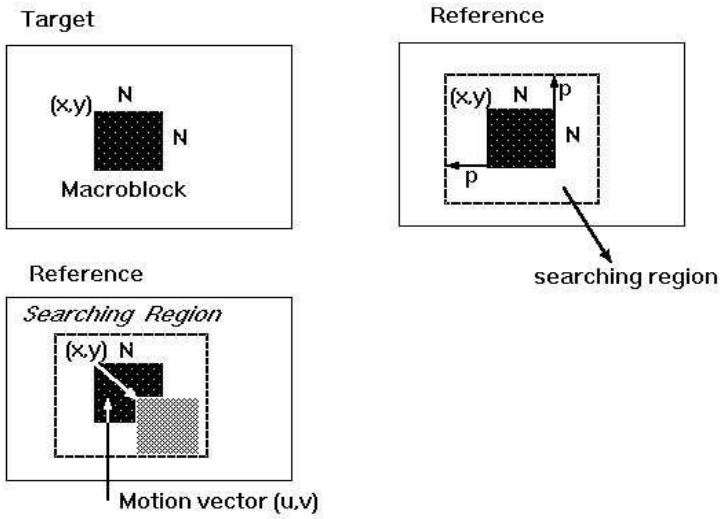

Motion Estimation

- The temporal prediction technique used in MPEG video is based on motion estimation.

The basic premise:

- Consecutive video frames will be similar except for changes induced by objects moving within the frames.
- Trivial case of zero motion between frames — no other differences except noise etc.
- Easy for the encoder to predict the current frame as a duplicate of the prediction frame.
- When there is motion in the images, the situation is not as simple.

Example

The problem for motion estimation to solve is:

How to adequately represent the changes, or differences, between these two video frames.

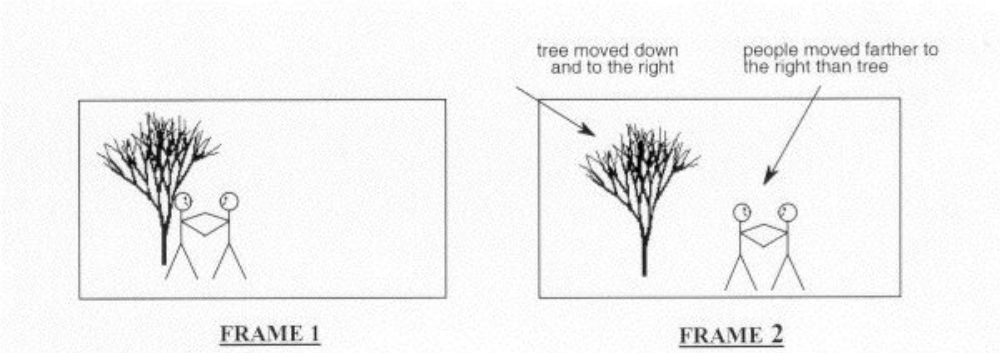

Solution

A comprehensive 2-dimensional spatial search is performed for each luminance macroblock.

- Motion estimation is not applied directly to chrominance in MPEG
- MPEG does not define how this search should be performed.
- A detail that the system designer can choose to implement in one of many possible ways.
- Well known that a full, exhaustive search over a wide 2-D area yields the best matching results in most cases, but at extreme computational cost to the encoder.
- Motion estimation usually is the most computationally expensive portion of the video encoding.

Motion Estimation Example

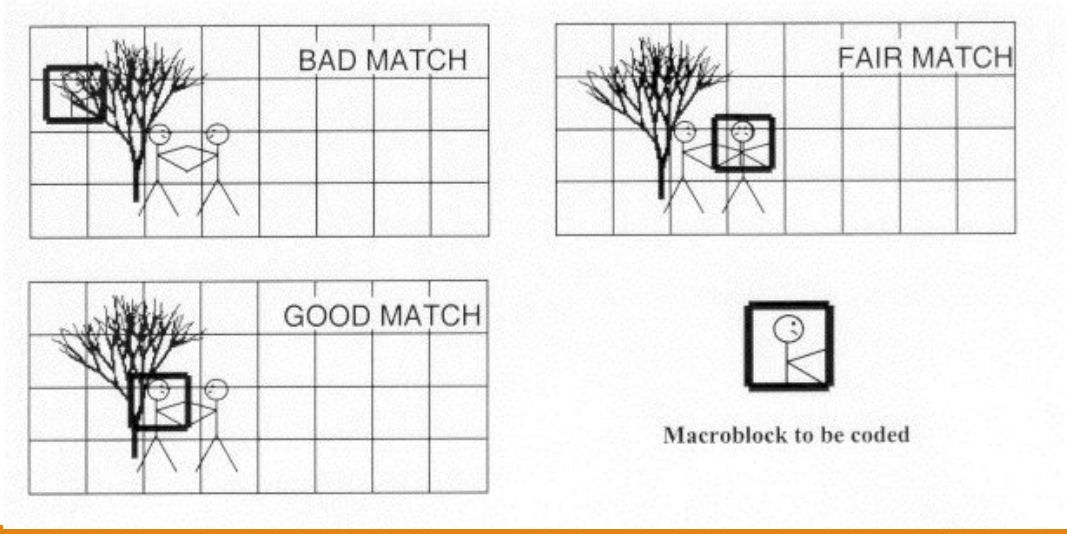

Motion Vectors, Matching Blocks

Previous figure shows an example of a particular macroblock from Frame 2 of earlier example, relative to various macroblocks of Frame 1:

- The top frame has a bad match with the macroblock to be coded.
- The middle frame has a fair match, as there is some commonality between the 2 macroblocks.
- The bottom frame has the best match, with only a slight error between the 2 macroblocks.
- Because a relatively good match has been found, the encoder assigns motion vectors to that macroblock,

Final Motion Estimation Prediction (Cont.)

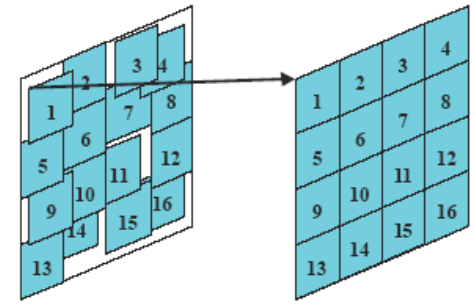

The predicted frame is subtracted from the desired frame,

Leaving a (hopefully) less complicated residual error frame which can then be encoded much more efficiently than before motion estimation.

Example

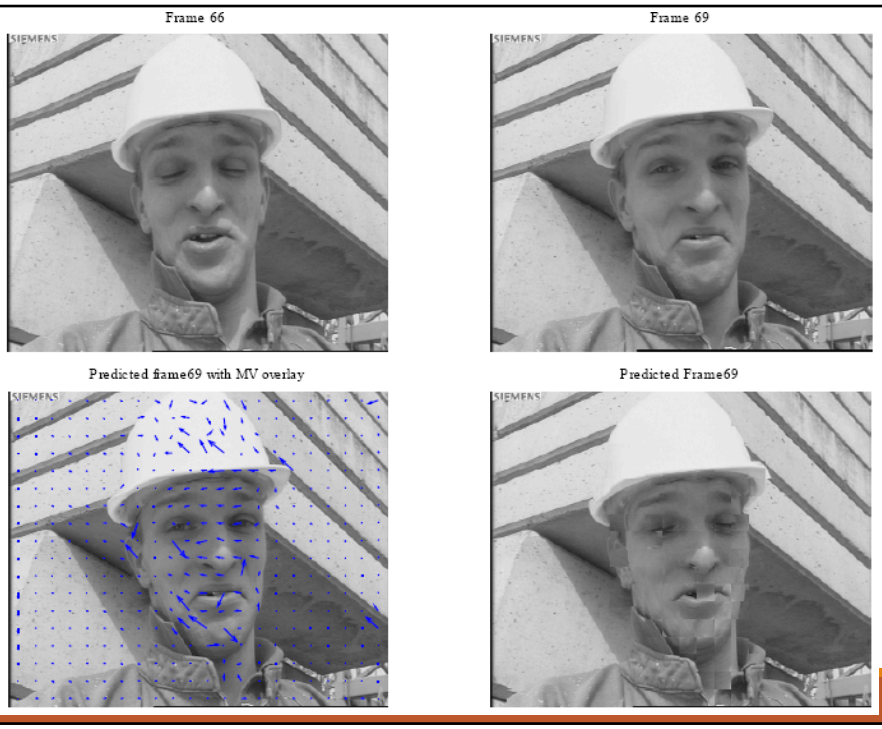

Example

Frame 66

Frame 69

Absolute Difference w/o Motion Compensation

Absolute Difference with Motion Compensation

Further Coding Efficiency

Differential Coding of Motion Vectors

- Motion vectors tend to be highly correlated between macroblocks:
- The horizontal component is compared to the previously valid horizontal motion vector and
- Only the difference is coded.
- Same difference is calculated for the vertical component
- Difference codes are then described with a variable length code (e.g. Huffman) for maximum compression efficiency.

Recap: P-Frame Coding Summary

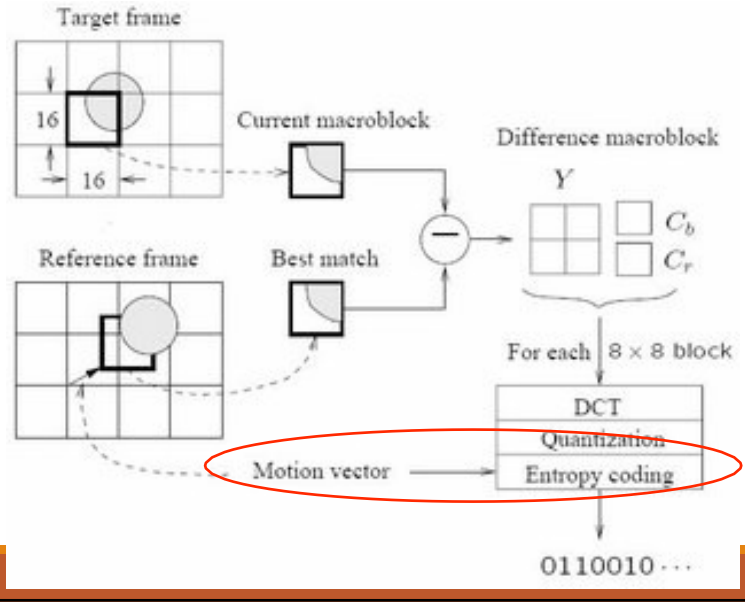

Estimating the Motion Vectors

So how do we find the motion?

Basic Ideas is to search for Macroblock (MB)

- Within a ±n x m pixel search window
- Work out for each window
  - Sum of Absolute Difference (SAD) (or Mean Absolute Error (MAE))
- Choose window where SAD/MAE is a minimum.

If the encoder decides that no acceptable match exists then it has the option of

- Coding that particular macroblock as an intra macroblock,
- Even though it may be in a P frame!
- In this manner, high quality video is maintained at a slight cost to coding efficiency.

Sum of Absolute Difference (SAD) and MAD

SAD is computed by:

$$\mathrm{SAD}(i,j) = \sum_{k=0}^{N-1}\sum_{l=0}^{N-1} |C(x+k, y+l) - R(x+k+i, y+l+j)|$$

- N = size of macroblock window typically (16 or 32 pixels),
- (x, y) the position of the original macroblock, C, and R is the reference region to compute the SAD.
- C(x + k, y + l) — pixels in the macro block with upper left corner (x, y) in the target.
- R(x + k + i, y + l + j) — pixels in the macro block with upper left corner (x + i, y + j) in the reference.

Mean absolute difference (MAD)=SAD/N^2

Sum of Squared Differences (SSD)

Alternatively: sum of squared differences

$$\mathrm{SSD}(i,j) = \sum_{k=0}^{N-1}\sum_{l=0}^{N-1} (C(x+k, y+l) - R(x+k+i, y+l+j))^2$$

Goal is to find a vector (i, j) such that SAD/SSD (i, j) is minimum.

Full Search

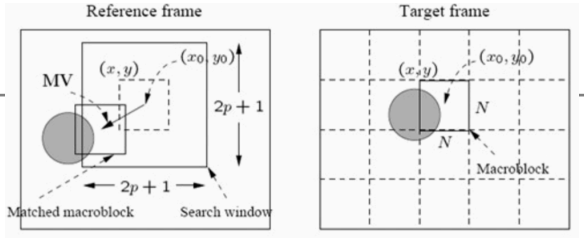

- Search exhaustively the whole (2p + 1) × (2p + 1) window in the reference frame.
- A macroblock centred at each of the positions within the window is compared to the macroblock in the target frame pixel by pixel and their respective SAD (or MAE) is computed.
- The vector (i, j) that offers the least SAD (or MAE) is designated as the motion vector for the macroblock in the target frame.
- Full search is very costly.

Full Search

Advantages:

Guaranteed to find optimal motion vector within search range.

Disadvantages:

Can only search among finitely many candidates. What if the motion is in fractional number of pixels?

High computation complexity

HOW TO IMPROVE?

Accuracy: consider fractional translations.

This requires interpolation (e.g . bilinear in H.263).

Speed: try to avoid checking unlikely candidates.

Bilinear Interpolation

- Bilinear interpolation:

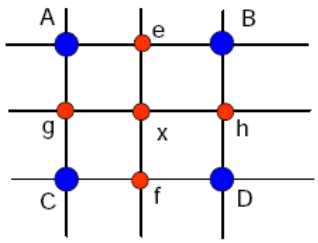

- Original samples:
  - A, B, C, D
- Half-pixel locations:
  - e, f, g, h, x

$$e = \left\lfloor \frac{A+B}{2} + 0.5 \right\rfloor$$

$$f = \left\lfloor \frac{C+D}{2} + 0.5 \right\rfloor$$

$$g = \left\lfloor \frac{A+C}{2} + 0.5 \right\rfloor$$

$$h = \left\lfloor \frac{B+D}{2} + 0.5 \right\rfloor$$

$$x = \left\lfloor \frac{A+B+C+D}{4} + 0.5 \right\rfloor$$

2D Logarithmic Search

An approach takes several iterations akin to a binary search.

Computationally cheaper, suboptimal but usually effective.

Initially only nine locations in the search window are used as seeds for a SAD-based search (marked as '1').

After locating the one with the minimal SAD, the centre of the new search region is moved to it and the step-size ("offset") is reduced to half.

In the next iteration, the nine new locations are marked as '2' and this process repeats.

If L iterations are applied, for altogether 9^L positions, only 9L positions are checked.

2D Logarithmic Search (Cont.)

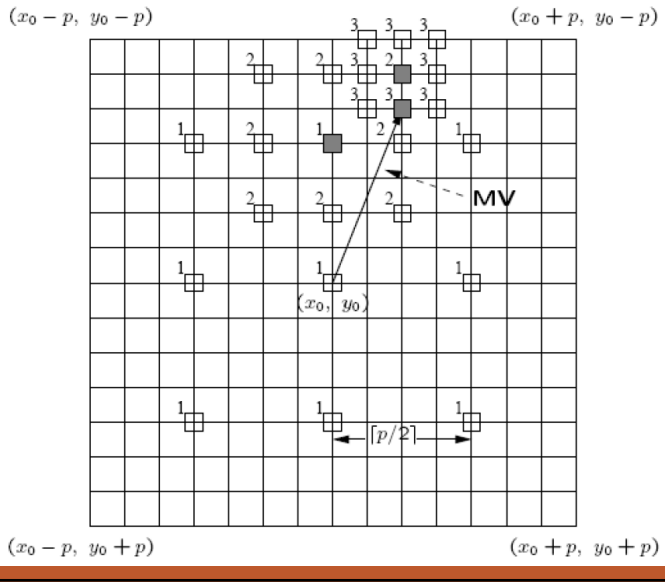

Hierarchical Motion Estimation

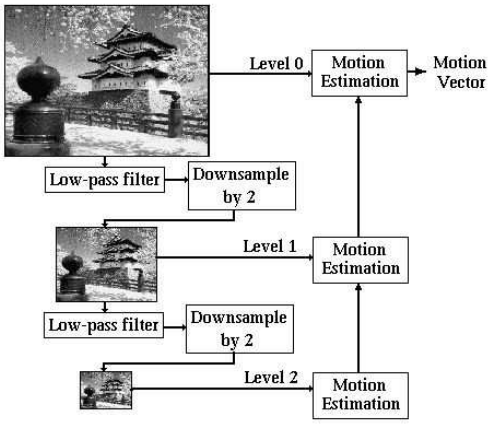

1. Form several low resolution version of the target and reference pictures.
2. Find the best match motion vector in the lowest resolution version.
3. Modify the motion vector level by level when going up.

Hierarchical Motion Estimation

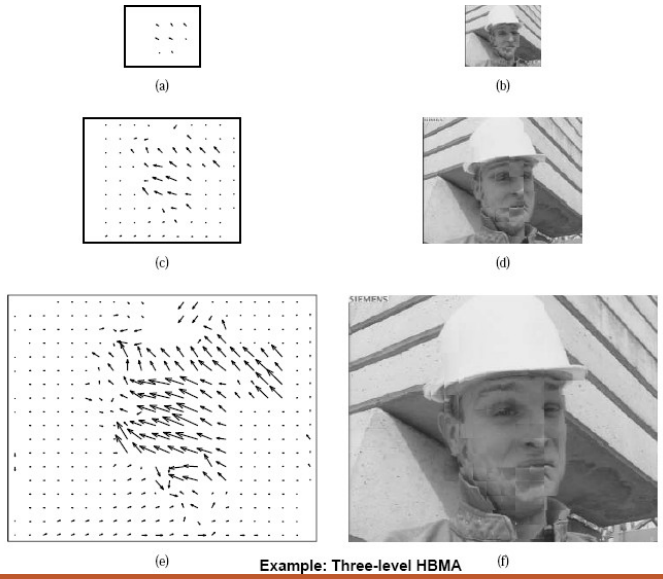

Example: Three-level HBMA

Performance Comparison

Operation for 720x480 at 30 fps (GOPS):

Search Method	p = 15	p=7
Full Search	29.890	6.990
Logarithmic	1.020	0.778
Hierarchical	0.507	0.399

MPEG Compression

MPEG stands for:

- Motion Picture Expert Group — established circa 1990 to create standard for delivery of audio and video
- MPEG-1 (1991). Target: VHS quality on a CD-ROM (320 x 240 + CD audio @ 1.5 Mbits/sec).
- MPEG-2 (1994): Target Television Broadcast. MPEG-3: HDTV but subsumed into an extension of MPEG-2.
- MPEG 4 (1998): Very Low Bitrate Audio-Visual Coding, later MPEG-4 Part 10 (H.264) for wide range of bitrates and better compression quality.
- MPEG-7 (2001) "Multimedia Content Description Interface".
- MPEG-21 (2002) "Multimedia Framework".

Three Parts to MPEG

The MPEG standard had three parts:

Video: based on H.261 and JPEG
Audio: based on MUSICAM (Masking pattern adapted Universal Subband Integrated Coding And Multiplexing) technology
System: control interleaving of streams

MPEG Video

MPEG compression is essentially an attempt to overcome some shortcomings of H.261 and JPEG:

Recall H.261 dependencies:

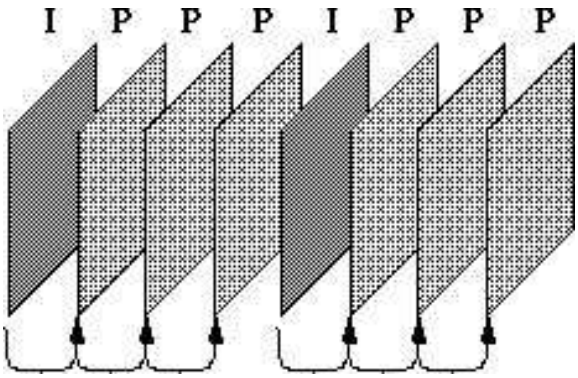

The Need for a Bidirectional Search

The problem here is that many macroblocks need information that is not in the reference frame.

For example:

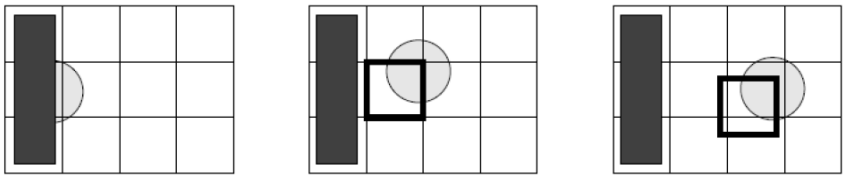

- Occlusion by objects affects differencing
- Difficult to track occluded objects etc.
- MPEG uses forward/backward interpolated prediction.

MPEG B-Frames

- The MPEG solution is to add a third frame type which is a bidirectional frame, or B-frame
- B-frames search for macroblock in past and future frames.
- Typical pattern is IBBPBBPBB IBBPBBPBB IBBPBBPBB. Actual pattern is up to encoder, and need not be regular.

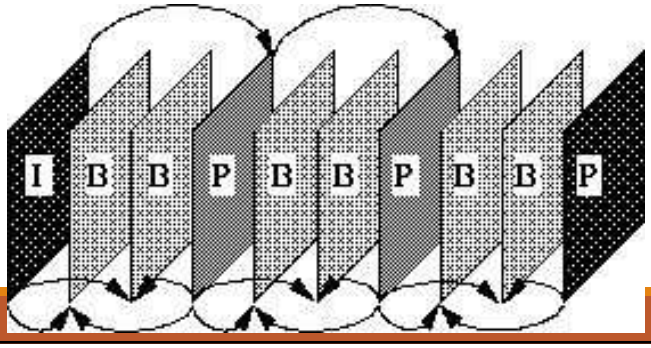

Example: I, P, and B frames

Consider a group of pictures of 8 frames:

Given: I, B, P, B, P, B, P, B, I, B, P, B, P, B, . .

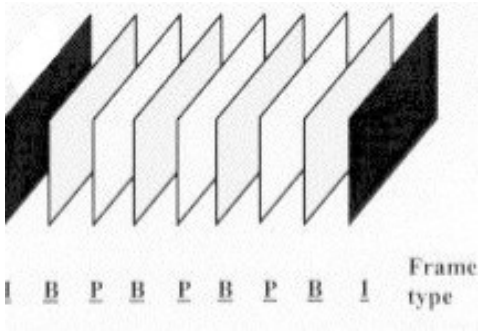

I frames are coded spatially only (as before in H.261).

P frames are forward predicted based on previous I and P frames (as before in H.261).

B frames are coded based on a forward prediction from a previous I or P frame, as well as a backward prediction from a succeeding I or P frame.

Bidirectional Prediction

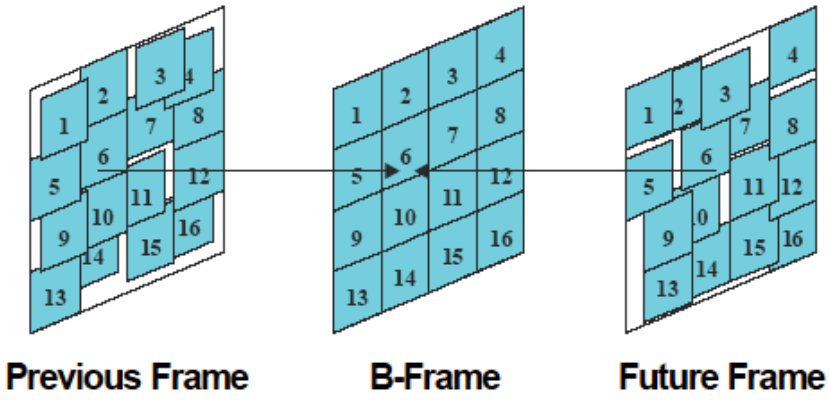

Example: I, P, and B frames (Cont.)

- 1st B frame is predicted from the 1st I frame and 1st P frame.
- 2nd B frame is predicted from the 1st and 2nd P frames.
- 3rd B frame is predicted from the 2nd and 3rd P frames.
- 4th B frame is predicted from the 3rd P frame and the 1st I frame of the next group of pictures.

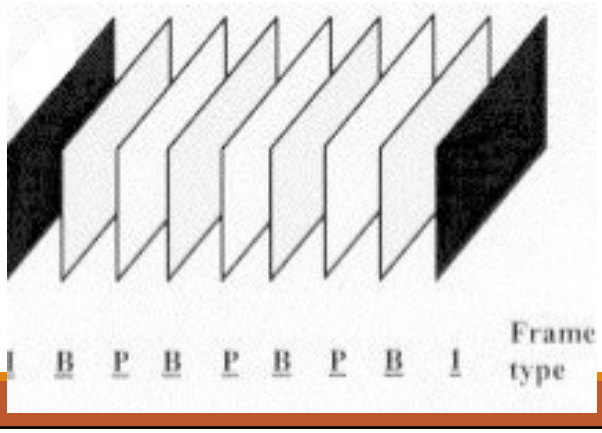

Bidirectional Prediction

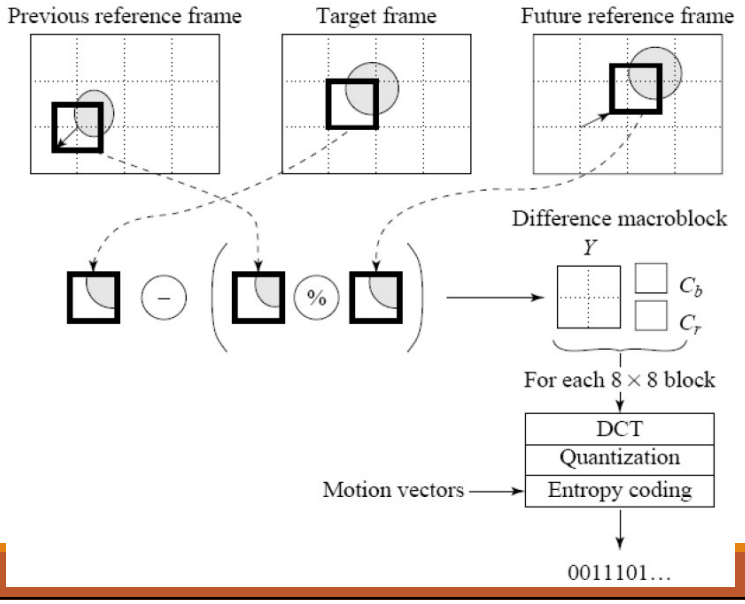

Backward Prediction Implications

Note: Backward prediction requires that the future frames that are to be used for backward prediction be

- Encoded and transmitted first, i.e. out of order.

This process is summarised:

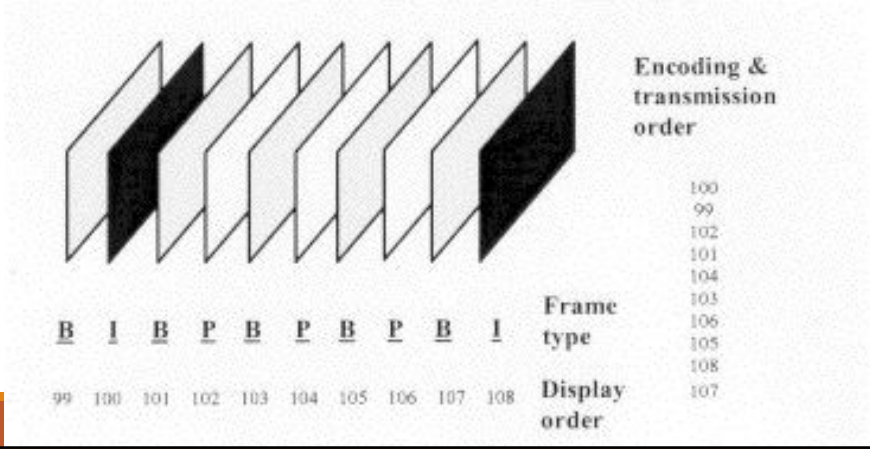

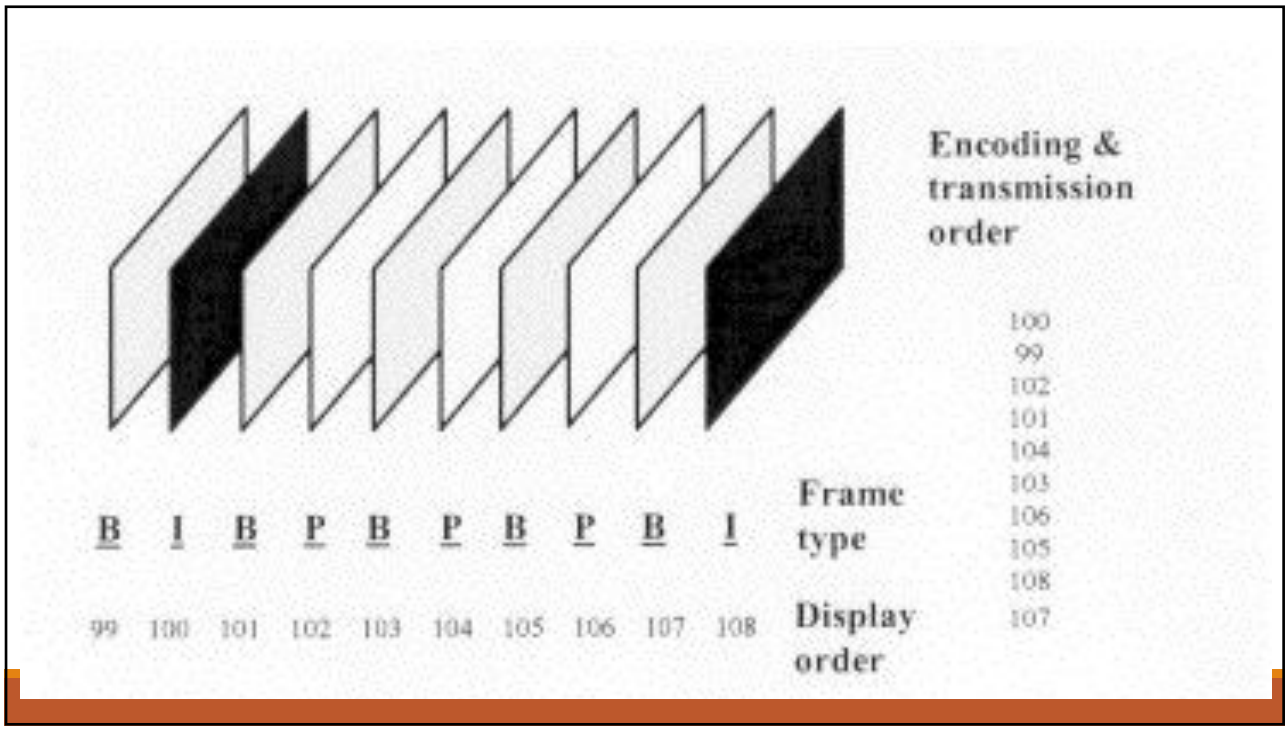

Backward Prediction Implications (Cont.)

Also NOTE:

- No defined limit to the number of consecutive B frames that may be used in a group of pictures.
- Optimal number is application dependent.
- Most broadcast quality applications, however, have tended to use 2 consecutive B frames (I,B,B,P,B,B,P,. . .) as the ideal trade-off between compression efficiency and video quality.
- MPEG suggests some standard groupings.

Advantage of Using B frames

Coding efficiency.

- Most B frames use less bits.
- Quality can also be improved in the case of moving objects that reveal hidden areas within a video sequence.

Disadvantage:

- Frame reconstruction memory buffers within the encoder and decoder must be doubled in size to accommodate the 2 anchor frames.
- More delays in real-time applications.

Frame Sizes

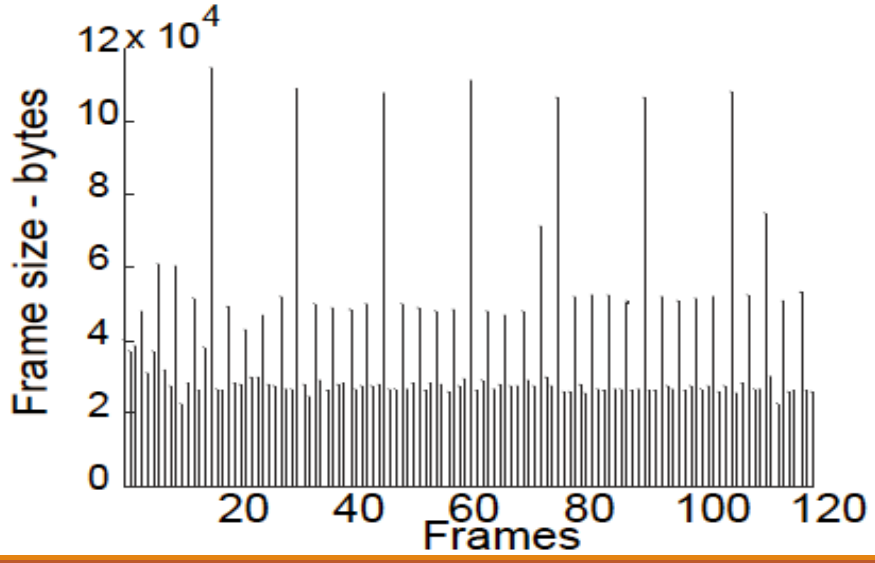

Random Access Points

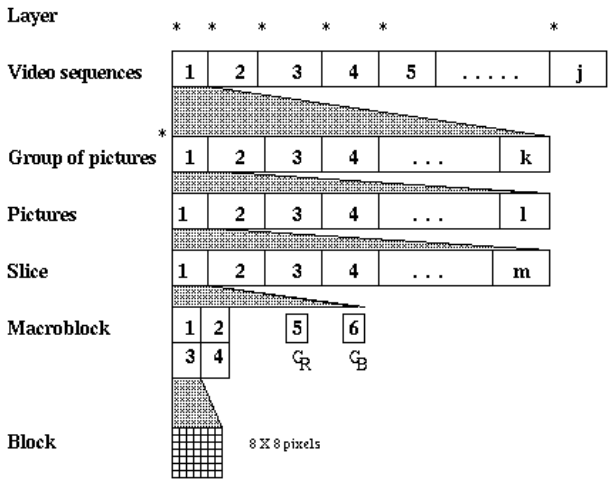

MPEG-2, MPEG-3, and MPEG-4

MPEG-2 target applications:

Level	size	Pixels/sec	bit-rate (Mbits)	Application
Low	352 x 240	3 M	4	VHS tape equiv.
Main	720 x 480	10 M	15	studio TV
High 1440	1440 x 1152	47 M	60	consumer HDTV
High	1920 x 1080	63 M	80	film production

MPEG-2, MPEG-3, and MPEG-4 (Cont.)

MPEG-2 differences from MPEG-1

- Search on fields, not just frames.
- 4:2:2 and 4:4:4 macroblocks
- Frame sizes as large as 16383 x 16383
- Non-linear macroblock quantization factor
- A bunch of minor fixes

MPEG-3: Originally for HDTV (1920 x 1080), got folded into MPEG-2

MPEG-4: very low bit-rate communication (4.8 to 64 kb/sec).
Around objects, not frames.