

METHOD AND APPARATUS FOR RECEIVING, STORING, AND PRESENTING MULTIMEDIA PROGRAMMING WITHOUT INDEXING PRIOR TO STORAGE

BACKGROUND

1. Field

5 The present disclosure relates to a method and apparatus for improved digital recording and presentation of broadcast information. Specifically, the present disclosure relates to a method and apparatus for receiving, storing, and presenting broadcast information in real-time and time-shifted modes of operation.

2. Related Art

10 Digital data recorders, such as digital video recorders (DVR's) have been known since at least 1992. Standard DVR's permit users to record broadcast information to a storage device for later playback. Typically, DVR's enable time-shifted (trick-play) modes of operation that are similar to functions found on video cassette recorders, with which most users are familiar. For example, DVR's may have functions such as "pause," "rewind,"
15 "fast-forward," "skip," and "slow motion." One of the first commercially available DVR's was the MediaStream system developed and marketed by Media4, now part of EchoStar Communications Corporation. In April 1996, Media4 introduced the MediaStream receiver, which was a Digital Video Broadcasting-compliant satellite receiver system with integrated DVR functions. The MediaStream system was designed
20 to both record and present programs simultaneously, allowing one program to be both recorded and presented. The MediaStream receiver system demultiplexed a Moving Pictures Group (MPEG) transport stream that contained one or more television programs, for example, and filled separate video packetized elementary stream (PES) and audio PES buffers. The data contained in the buffers was written to disk for later playback in either
25 normal or trick-play mode. The MediaStream system did no intelligent parsing of the input to generate an index to aid in trick-play modes of operation, but merely performed a "brute-force" search of data stored on the hard disk when performing those functions.

Many methods and systems have been developed for creating indices by intelligently parsing broadcast input streams and using index information generated during input to

later find and play back appropriate frames of data. One of the earliest of these systems is described in two patents assigned to Imedia Corporation, U.S. Patent Nos. 5,949,948 and 6,304,714, both to Krause et al. These patents disclose a set-top DVR system for simultaneous presentation and recording of compressed digital data. For example, U.S.

5 Patent No. 5,949,948 discloses a start-code detector for detecting the beginning of video I-frames in an MPEG data stream, an indexing system that correlates I-frames with addresses in memory, and a trick-play system that searches the index information to determine which frames to play back in trick-play operations. Similarly, U.S. 5,614,940, to Cobbley et al., of Intel, discloses a set-top system that can convert broadcast
10 information to a digital format, generate during input various index data relating to the content of the broadcast information, store both the compressed broadcast data and the related index data, and then retrieve the broadcast data for playback (in normal or trick-play mode) based upon the corresponding index information. Similar, front-end, input-side intelligent parsing and index-based searching methods are disclosed in U.S. Patent
15 No. 5,956,716 to Kenner, et al., U.S. Patent No. 5,659,539 to Porter, et al., U.S. Patent No. 6,167,083 to Sporer, et al., and U.S. Patent No. 5,577,190 to Peters.

A later recording system, developed at TiVo Inc. and described in the specification of U.S. Patent No. 6,233,389, to Barton, et al., also employed a specific type of intelligent parsing/indexing during input and prior to storage of the broadcast information on a
20 storage device. The system described in that patent employs a special circuit called a "Media Switch" that generates indices and fills separate appropriate buffers with specific data. The disclosed "Media Switch" mediates between the central processing unit (CPU), storage device, and memory and thus off-loads the intensive index-based processing of the input stream from the CPU to a separate device. Also in the Barton, et al., system, a
25 software "source object" converts the data into data streams and fills a buffer that is assigned by a central software "transform object" that is responsible for overall control of buffer assignment. The software "transform object" then writes the data to a hard disk. The software "transform object" is also responsible for reading data from the hard disk, filling buffers with the data, and assigning the filled buffers to a software "sink object" for
30 later decoding and playback.

These earlier systems may be inefficient and overly complicated in some operational settings. Such systems require intensive processing during input of the entire set of broadcast data. Given the high throughput required for modern DVR functions, the processing power required during input in such systems may tax the CPU or, in the case of the system of Barton, et al., require specialized hardware and software. Moreover, since much of what is recorded will not be played back in anything other than standard mode, the processing power required, and the memory required to store related index information, may be largely wasted. A more robust, cheaper, and less complicated system is needed.

BRIEF SUMMARY

The present invention seeks to reduce the identified problems of prior methods and systems.

According to a first aspect of the present invention, there is provided a method of presenting stored multimedia programming, which programming had not been indexed prior to its storage, the method comprising: providing statistical information being of the type that can be generated during normal presentation of the programming; receiving a seek input; loading a portion of the stored multimedia programming, the size of the portion being determined by the statistical information, from a starting position based on the seek input; and, if analysis indicates that the loaded programming is sufficient for special presentation of the programming, making the special presentation.

Embodiments of the invention improve upon earlier methods as parsing, separating, transforming or other processing functions of received broadcast data are not required before storage of the data. With methods of the invention, statistical and probabilistic algorithms are utilized to search for and keep track of the programming data when presenting such data from storage.

Thus, the received programming may be written directly to a storage device using an asynchronous, single buffer read/write process. Upon invocation of trick-play modes of operation, search operations may be performed based upon statistics or dynamically generated during presentation operations or received with the broadcast data. Upon

normal presentation, statistics may be generated to determine the ideal number of frames to skip, the number of bytes to seek, and the size of data files to read from storage during trick-play operation.

Methods of the invention may utilize algorithms and operations for dynamically
5 determining any required skip, seek, and read values which minimize the use of system resources. In this way, data may be stored more efficiently and economically, and the trick-play operations can be performed in a smoother, more nuanced manner with the application of appropriate probabilistic algorithms.

Other features and advantages of the methods and systems described herein will become
10 apparent from the following detailed description, when it is considered in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWING FIGURES

Figure 1 is a block diagram illustrating a DVR system described herein.

Figure 2 is a block diagram of the memory allocation in a DVR system described herein.

15 **Figure 3** is a block diagram of the memory allocation in a single buffer memory described herein.

Figure 4 is a flow diagram depicting the storage of data in a DVR system described herein.

20 **Figure 5** is a logic diagram of the functions of a record thread application described herein.

Figure 6 is a logic diagram for presentation of stored data in a system described herein.

DETAILED DESCRIPTION

As a general matter, it is understood that “intelligent analysis,” “intelligent parsing” or “indexing” of programming data refers to analyzing the data to extract information
25 therefrom. Where the data is presented in a stream, for example, of video and audio data, the analysis may associate video- or audio-specific information, such as frame

presentation time information, with system-specific information, such as position in a stored data file.

The following description sets forth numerous examples of methods and systems described herein for the storage and presentation of multimedia programming, without the need for indexing prior to storage. It should be recognized, however, that such description is not intended as a limitation on the scope of the present invention, but is instead provided as a description of exemplary embodiments.

I. System for Presentation and Storage

With reference to **Figure 1**, a digital video recorder (DVR) embodiment **100** receives an encrypted, scrambled or clear broadcast signal from a direct broadcast satellite (DBS) system through satellite receiver **110**. In alternate embodiments, additional or different broadcasting sources and formats may be used such as off-air or terrestrial transmissions, or cable television (TV). Typically, the broadcast signals are television or other multimedia programming signals. The signals may be high definition (HD) television, standard definition (SD) television, audio-only signals or other signals. Such a signal may contain numerous frequency bands of data, with each band containing numerous TV content programs (e.g., CNN®, HBO®, etc.). In a preferred embodiment, the signal received from receiver **110** contains multimedia programming transmitted as an MPEG transport stream, though alternate formats, such as analog TV formats, may be used. An MPEG transport stream (transport stream) contains packets of video-only, audio-only, or other data. Each packet may have associated header information, including packet identification (PID) information. PID information may identify the data type of the packet (video, audio, other) and the content programming associated with the packet, as well as other information. Video and audio packets encoded in the transport stream may also contain presentation timestamp (PTS) data that allows for synchronization of the packets. The packets may also contain start code data that identifies the beginning of a video or audio frame.

Broadcast multimedia programming is received at receiver **110** and forwarded to input section **120** of DVR **100**. The signal may be a modulated broadcast signal spanning a broadcast frequency band. Receiver **110** may translate the signal it receives to an

intermediate frequency before forwarding it to DVR 100. Tuner 121 of section 120 tunes the signal received from receiver 110 to a frequency range (channel) that contains content programming of interest. Input section 120 may also contain a demodulator 122 that demodulates the broadcast signal to produce a demodulated transport stream. Section 120
5 may also contain a demultiplexor 123 that filters the transport stream according to programming-specific PID's to produce a transport stream that contains only packets associated with the content programming of interest. In one embodiment, demultiplexor 123 may produce a separate video-only packetized elementary stream (PES) and a separate audio-only PES stream. In another embodiment, a single transport stream is
10 produced with interleaved video and audio data. Demultiplexor 123 may also filter out the other (e.g., non-video and non-audio) data packets for use in DVR 100. Input section 120 may also perform additional functions such as error correction, descrambling, decryption, analog-to-digital conversion or a number of other basic signal processing functions.

The MPEG transport stream outputted from section 120 may be routed to a display section
15 130 for immediate presentation in real time. Display section 130 contains at least an MPEG video decoder 131 and an MPEG audio decoder 132. Display section 130 may further contain digital-to-analog converters, encoders, additional decoders, video or audio filters, and/or memory buffers, as needed for delivery to a television 140 or other display device.

20 The MPEG transport stream outputted from section 120 may also be routed to a storage device, such as hard disk 150, for later presentation or for presentation in other than real time. In a preferred embodiment, program logic uses a single buffer for transfer to hard disk 150, without the use of additional buffers. Preferably, the transport stream received from section 120 is written onto hard disk 150, without first analyzing or indexing MPEG
25 video and/or audio frame information, as an MPEG transport stream file (TSP file). In another embodiment, the MPEG information is stored as a PES file or other suitable file format. By doing so, the MPEG transport stream is efficiently stored for later use without employing significant system resources. Time sequence, PTS, start code or other embedded MPEG frame information need not be analyzed, indexed or otherwise
30 correlated with system-specific information, such as TSP file position, prior to storage. Separate TSP files may be maintained for each separately recorded content program or for

each separate recording session. Hard disk 150 is connected to display section 130 to provide both contiguous and non-contiguous presentation of any content program stored as a TSP file on hard disk 150. Section 120 is capable of simultaneously outputting to both display section 130 and hard disk 150 for simultaneous storage and presentation of content programming.

DVR 100 also includes at least one processor 160 and at least one system RAM module 170. Program logic, such as record logic, normal playback logic or trick playback logic necessary for the operation of DVR 100 may be executed on processor 160 in conjunction with RAM module 170. In alternate embodiments, separate processors and separate RAM modules may be employed for the functions of input, storage, display, and/or other functions of DVR 100. In one embodiment, DVR 100 is a system operating on a Linux operating system. In alternate embodiments the DVR may be a system operating on a UNIX, Windows, Mac OS, or other operating system. DVR 100 may comprise multiple input sections, display sections, storage devices, processors and RAM modules. In this way, DVR 100 may accommodate a number of signal sources and display and record a number of content programs, simultaneously or separately.

II. Storage Operations

In one embodiment, the recording program logic operates using a single memory buffer, having a fixed memory address that can be accessed asynchronously by both a record driver and a record thread application. This single buffer, also referred to as a record buffer, is filled by a record driver. The data in the single buffer is then moved in a single operation from the single buffer to hard disk 150 by a record thread application.

Preferably, recording program logic is not flow controlled and the record driver and the record thread application write or read to or from the single buffer independently, without either application having control over the other. More preferably, the single record buffer is a circular buffer. Use of a single buffer eliminates the need for transfer between two or more separate buffers, which may conserve processor and other system resources. A single buffer method may further increase system efficiency by eliminating the need for communication between a record driver and a record thread application. Preferably, DVR

100 employs one record driver and one record thread application for each tuner in the DVR.

Figure 2 is the system memory allocation in one embodiment of DVR **100**. System drivers, including the record driver(s), may occupy driver space **210**. Single buffer memory **220** is reserved by the Linux kernel at system startup. In a preferred embodiment, single buffer memory **220** is assigned a fixed memory address during each system startup. System memory also contains a Linux kernel and user space **230** that may contain a playback buffer or other user-specific elements. In one embodiment, the Linux kernel reserves a fixed address for memory **220**. Single buffer memory **220** may be logically divided into one or more transport buffers (not shown). In a preferred embodiment having fixed memory addresses for single buffer memory **220**, the record driver(s) in the driver space **210** calculates a transport buffer memory address using a hard-coded offset from the top of the kernel memory. The hard-code offset is determined by the `high_memory` symbol exported by the Linux kernel in space **230**. The record thread application(s) residing in space **230** uses actual hard-coded transport buffer address(es), as well.

Figure 3 depicts a specific embodiment of memory allocation in single buffer memory **220**, the single buffer shared asynchronously by both the record driver(s) and the record thread application(s). Memory **220** includes one record information region **310** and a transport buffer for at least each tuner in the system. In a preferred embodiment, a single transport buffer exists for each tuner in DVR **100**. Depicted are Transport Buffer 0 **320**, Transport Buffer 1 **330** and Transport Buffer 2 **340**, though more or fewer transport buffers may be used. In one embodiment, record information region **310** is one page of memory 4096 bytes in size and contains an array of structures where information regarding the position and size of the data written to each transport buffer by the record driver(s) is stored and updated by the record driver(s). The record thread application(s) may access these structures.

Figure 4 depicts handling of a transport stream intended for storage, in DVR **100**. Transport stream data **410** received from input section **120** is written to transport buffer **430** of single memory buffer **220** by record driver **420**. In an alternate embodiment, the data is moved into the transport buffer by the hardware of input section **120**. Record

driver 420 also updates information page 310 to indicate the position of the driver's record pointer. The record driver operates in real time as transport stream data is received from the input section without flow control.

Record thread application 440, which stores a last read position, accesses information page 310 to determine the size of un-written data in buffer 220. Record thread application 440 transfers the un-written data directly to hard disk 150 for storage as a TSP file.

Figure 5 depicts a simplified flow diagram for one embodiment of a record thread application loop 500 operating on DVR 100. The record thread application sleeps on a "Start" semaphore until a master application signals the semaphore allowing the record thread to begin execution, process 510. Once the record thread executes, several local variables may be initialized and the execution loop is entered, process 520. In process 530, the execution loop reads the recording information structure associated with the transport buffer associated with the record thread application to determine how much unread data exists in the transport buffer and the location of that data. Also in process 530, this data is then written to hard disk 150, even if it is determined that zero bytes of data exist. The record thread application transfers the data to hard disk 150 via a direct IO (input/output) transfer, without processing the data through the record thread application. After being written to the hard disk 150, several local variables, such as last read address, may be updated, process 530. In process 540, the record thread may then go into an interruptible sleep, via a timed "Sleep" semaphore. The timeout period of the "Sleep" semaphore may be based on the service type, e.g., SD or HD television, or audio-only. This timed "Sleep" semaphore allows other processes in the system to run. If the "Sleep" semaphore times out, then the record thread executes another loop. If the master application has signaled the record thread to stop execution, then the semaphore does not time out, and the application terminates at process 550. The record thread application operates asynchronously with respect to the record driver. In alternate embodiments, the record thread application may be a record task or a record process application.

III. Presentation from Storage

DVR 100 accommodates several presentation modes for the stored video and audio data.

In one embodiment, presentation modes include forward play, pause, reverse play, slow

motion forward or rewind, fast forward or rewind, and skip forward or back. Using the methods and systems described herein, DVR 100 is able to accommodate these modes without using previously indexed MPEG frame information or the need for specific frame positioning or time sequence information. By avoiding the need to determine time
5 sequence information for all stored video and/or audio data before recording and presentation from storage, system resources are conserved. In one embodiment, presentation from a storage device such as hard disk 150 is performed by reading portions of the stored MPEG transport stream to a read buffer prior to outputting to display section 130. In one embodiment, the read buffer is a circular read buffer. The presentation
10 methods described herein may be employed with video-only data, audio-only data or combined video and audio data.

MPEG video compression standards reduce the amount of data required to transmit or store a video signal by representing certain frames of video as a delta from a previous or subsequent frame. MPEG video generally consists of three major frame types. I-frames,
15 or intra-coded frames, are pictures encoded without referencing any other frame. P-frames, or predictive frames, are pictures encoded by referencing the delta from previous frames. B-frames, or bi-predictive frames, are pictures encoded by referencing the delta from previous and subsequent frames. MPEG-4 specifies an additional intra-coded frame type, the IRD-frame, which may also be used. It is understood that an IRD-frame may be
20 substituted for an I-frame in the methods and systems described herein. To display a complete image, at least one intra-coded frame (I or IRD) must be decoded and presented. MPEG encoded video streams are broadcast in real time at a predetermined frame per second (fps) rate. The fps may vary depending on the content program. For example, the frame rate may be approximately 30 fps (standard television), 24 fps (movies), 25 fps
25 (some foreign content), or other frame rate. MPEG standards may also be used for the compression of audio data into a frame format.

Presentation modes may be conceptually divided into three categories, as provided:

Table 1

Linear(play)	forward slow motion forward pause
contiguous(trick)	rewind fast forward (certain speeds) fast rewind (certain speeds)
non-contiguous(trick)	fast forward (certain speeds) fast rewind (certain speeds) skip forward skip back

Linear(play) is any presentation mode that displays every frame (I, P, and B) in sequential order. Forward mode, also referred to as “normal” play, which is presentation of all the video data at its broadcast fps rate, is a form of linear(play). The terminology “trick” is used to denote any presentation mode that requires either non-contiguous reading from the TSP file (“seeking”) or display of fewer than the total number of picture frames (“skipping”). Contiguous(trick) is any trick mode that loads stored multimedia data contiguously. Non-contiguous(trick) is any trick mode that loads stored multimedia data non-contiguously. Other conceptual divisions of the presentation modes may be employed. In one embodiment, the presentation mode is selected by a user of DVR 100 through the use of a remote control device capable of facilitating user control of DVR 100.

Figure 6 depicts a simplified logic diagram of the steps involved in presentation once the user selects a change in the desired presentation mode, in a preferred embodiment using data stored as a TSP file. Depending on the presentation mode selected, statistical information regarding the video data encoded in an MPEG transport stream may be used to enable presentation without the need to exhaustively analyze the TSP file from the beginning to find the desired presentation location. In process 600, the statistical data is either “spoofed” (i.e., provided) using non-specific, pre-generated video data statistics or the statistical data is generated specifically during normal presentation of the content program. The following table presents video data statistical information that is spoofed or collected, in one embodiment:

Table 2

Statistic:	Description:
total_num_frames	The total number of I, P, and B frames used in the statistics.
total_num_I_frames	The total number of I frames used in the statistics.
avg_frm_size	The average frame size taken from all frames encountered during playback.
I_spacing	The average integer number of frames from one I-frame to the next, also referred to as a group of pictures.
GOP_size	Average group of pictures size. Calculated as an I-frame spacing multiplied by average frame size.
Fps	Output frames per second.

The information collected during normal presentation may or may not be stored in non-volatile memory for later use. In one embodiment, the information is maintained only for the duration of the current presentation session. In another embodiment, the statistical information may be contained in the transport stream, as broadcast. In an embodiment using statistical information broadcasted in the transport stream, the statistical information contained in the transport stream is private data contained in the adaptation field of a transport stream packet.

- At process 610, the system selects the desired presentation mode and sets the number of frames to be skipped. In an embodiment employing a remote control device, user selection of a presentation mode is handled as a user input, from which DVR 100 determines the number of frames to skip. Skipping frames during presentation results in time-shifted display which a user perceives as accelerated display, expressed as multiples of the predetermined play rate (e.g., a presentation speed value). By way of example, if every 8th I-frame (*Nth_I_frame*) is displayed 2 times (*M_repeats*) and an I-frame occurs, on average, every 15th frame (*I_spacing*) in the content program, the user would perceive the presentation as "60x" (*Speed*) the normal rate. The perceived speed of presentation can be determined by the following formula:

$$Speed = (I_spacing) * (Nth_I_frame) / (M_repeats)$$

The product of (*I_spacing*) and (*Nth_I_frame*) determines the number of frames to skip, from the last frame presented for display. Generally, the statistical data from process 600 is used to provide (*I_spacing*). Alternatively, when the desired presentation mode dictates a single frame skipping event, the user perceives a “jump” or single skip forwards or
5 backwards in an otherwise normal speed presentation.

Selecting the presentation mode may set a number of variable flags. In one embodiment, the flags set are TRUE/FALSE binary flags such as: “trick,” “contiguous,” and “forward.” In an embodiment employing a remote control device, user seek input based on selection of a presentation mode is used to set the variable flags. The state of the flags may affect
10 subsequent processing steps. In one embodiment, the “contiguous” flag is set as “TRUE” either if every frame is displayed (e.g., linear(play)) or if the number of frames to be skipped is fewer than or equal to (*I_spacing*). For example, display is “contiguous” when four frames are skipped and an I-frame occurs every 15th frame. In these instances, system efficiency may be optimized by loading data contiguously. Accordingly, certain
15 presentation methods that skip frames will be considered contiguous, while other methods that also skip frames are designated as non-contiguous. Frames to be skipped can be expressed as a positive value if forward=TRUE and as a negative value if forward=FALSE.

At process 620, a recycle operation may be performed on the read buffer, depending on
20 whether the contiguous flag is set. If a contiguous presentation mode is selected, stored data may be loaded contiguously. Accordingly, any portion of the read buffer that is resident but has not been forwarded to the display section (unused) can be recycled for potential use. Recycling conserves system resources by reducing the amount of file data to be read. If a non-contiguous presentation mode is selected, recycle process 620 does
25 not occur and the unused data is cleared from memory or overwritten (flushed).

The size of stored MPEG transport stream file data to be read (read size) is determined at process 630. Read size is determined by the state of the contiguous flag. If the mode is contiguous, read size is equal to the maximum read buffer size minus the recycled data size. For non-contiguous modes, in one embodiment, the read size is twice the average
30 group of picture size, as determined by process 600. Logically, there is a tradeoff in

system efficiency between increasing the read size and the cost that would be incurred by an additional read event, if a complete I-frame cannot be located. By setting read size at twice the GOP size but less than the maximum buffer size, system resources are conserved while maintaining a high probability that a complete I-frame is loaded to the read buffer, while in a non-contiguous mode. Alternatively, the non-contiguous mode read size can be determined using the following formula, wherein *Service_time* is the time required to locate a complete I-frame, $P_{miss}(s)$ is the probability of not locating a complete I-frame and $t_{read}s$ is the time required to perform a read of size s :

$$(Service_time) = (P_{miss}(s) + 1) * (t_{read}s)$$

- 10 Once the curve for $P_{miss}(s)$ is either provided or determined empirically within DVR 100 (e.g., through sequential non-contiguous presentation events) the value of s can be dynamically adjusted to minimize *Service_time*.

Seek position is calculated at process 640. Seek position is determined relative to the current read position in the stored MPEG transport stream file. In one embodiment, the current read position is indicated by a file pointer. For presentation modes in which both the "forward" and "contiguous" flags were set to "TRUE" during process 610, no seeking should occur, as data loading will be performed contiguously. For contiguous rewind (i.e., contiguous=TRUE, forward=FALSE) in systems having a file pointer that only reads forward, seek position is the sum of the recycled data size and the read size determined at 630 (i.e., the maximum buffer size), so that data preceding the current file pointer position will have been placed into the read buffer, after the read event. For non-contiguous modes, a seek vector equal to the product of the frames to be skipped (set at 610), and the average frame size (determined at 600) is calculated. An adjustment equal to half of the GOP size is also determined to increase accuracy. Seek position is calculated based on the following formula, wherein (*Origin*) is the current file pointer position:

$$(Seek_Position) = (Origin) + (Seek_vector) - \frac{(GOP_size)}{2}$$

At process 650, the file pointer seeks to the position determined at 640. A portion of the stored MPEG transport stream file equal to the read size determined at 630 is read into the read buffer, process 660.

At process 670, program logic analyzes the data in the read buffer to determine if a complete I-frame of data is present. In an MPEG transport stream, each packet may be optionally structured with an Adaptation Field. The Adaptation Field may contain transport stream state signaling, stream timing details, transport private data, and/or video splicing information. Transport private data contained within the Adaptation Field may contain access unit (AU) information. Access units are coded representations (e.g., I, B, and P frames) of a unit suitable for display (presentation unit), such as a video frame. Typically, access unit information signals whether an I-frame start is contained within the payload of the transport stream packet. Once an I-frame start is identified, locating another frame start further in the read buffer signals a complete I-frame. If access unit information is not available, program logic can analyze the transport stream payload for start code information, which may signal the start of a video frame. Data immediately following the start code indicates the video frame type (I, P, or B). As with access unit information, locating a subsequent video frame start code after identifying an I-frame start indicates the presence of a complete I-frame in the buffer. Generally, identification of I-frames through use of the Adaptation Field data is less system resource intensive than start code identification. However, start code information is always available while Adaptation Field information is optionally encoded. As described earlier, the Adaptation Field may also contain frame statistical information as private data. In one embodiment, start code identification is used only when Adaptation Field information is unavailable. In another embodiment, start code data is always used either independently or in conjunction with Adaptation Field information.

At decision 680, program logic determines whether additional data must be read. Preferably, DVR 100 is a system having a read buffer at least equal to the sum of the maximum group of pictures size and the maximum I-frame size. In such a preferred system, a complete I-frame will with a high probability be located in the read buffer during any contiguous play mode, as the maximum read buffer size is employed. In embodiments having a read buffer less than the sum of the maximum group of pictures

size and the maximum I-frame size, it may be necessary to perform additional recycle processes, additional read processes, and/or to flush at least a portion of the read buffer to locate a complete I-frame. In non-contiguous mode, which does not recycle and reads less than the maximum read buffer size, an additional data read ("append") may be necessary if
5 a complete I-frame of data is not located. An append occurs when the system loops to 660 and reads an additional portion of MPEG transport stream data from storage equal to the calculated read size and appends the new data to the data already loaded in the read buffer. Depending on the size of the read buffer, a flush of at least a portion of the read buffer may be necessary to allow for additional append operations in non-contiguous mode.

10 Program logic will again analyze what the read buffer contains to determine if a complete I-frame is loaded, and if a complete I-frame is not loaded, will perform looping read and analyze processes until a complete I-frame is located.

Once a complete I-frame is located, it is forwarded to display section 130 at process 690 for decoding and display. Display section 130 is capable of outputting video and/or audio
15 signals in a number of formats for presentation on a variety of display devices, such as a television set. In an embodiment wherein the broadcasted signal is an audio-only broadcasted signal, the display device may be a device capable of presenting only audio signals, such as a stereo system. DVR 100 may repeat the presentation process described as needed to create the desired presentation mode. Using the methods disclosed herein, a
20 DVR system may display MPEG transport stream encoded video and/or audio data in multiple presentation modes including at least normal speed, variable speed forward and reverse, and skip forward and reverse without having to linearly analyze the transport stream from the beginning to find the desired picture frame and without having to analyze and index video and/or audio frame information prior to storage of the transport stream.

25 While the present invention has been described with reference to certain preferred embodiments, those skilled in the art will recognize that various modifications and other embodiments may be provided. These and other embodiments are intended to fall within the scope of the present invention. These and other variations upon and modifications to the embodiment described herein are provided for by the present invention, which is
30 limited only by the following claims.

CLAIMS

What is claimed is:

1. A method of presenting a stored multimedia program, the method comprising:
receiving a seek input associated with non-contiguous playback of a multimedia program;
based on statistical info associated with the program, determining a seek position and a data read size associated with the seek input, wherein the statistical information includes at least an average frame size of the multimedia program and an average size of a group of pictures in the multimedia program;
loading a portion of the stored multimedia program, wherein the portion loaded begins at the seek position and is of the data read size;
determining whether the loaded portion of the stored multimedia program includes a complete I-frame of data; and
responsive to determining that the loaded portion includes the complete I-frame, outputting the loaded portion commencing at, and including at least, the I-frame.
2. The method of claim 1, further comprising:
outputting video data in the multimedia program sequentially at a normal presentation rate; and
generating the statistical information based on the normal presentation rate of the multimedia program.
3. A method of claim 1, wherein the statistical information has been pre-generated and is received from a service provider.
4. A method as claimed in any of claims 1 to 3, wherein the data read size is determined based on the average size of the group of pictures in the multimedia program.

5. A method as claimed in claim 4, wherein the data read size is determined based on twice the average size of a group of pictures in the multimedia program.
6. A method of as claimed in any preceding claim, wherein the seek position is determined based on a current file pointer position, a number of frames to skip for the seek input, the average frame size and a group of pictures adjustment, the adjustment based on the average size of the group of pictures.
7. The method as claimed in any preceding claim, wherein the seek position is determined by calculating a seek vector equal to the product of the frames to be skipped and the average frame size, adding the calculated seek vector to the current file pointer position, and adjusting the result by half of the average size of the group of pictures.
8. A method as claimed in claim 7, wherein the adjustment is determined based on half the average size of the group of pictures.
9. A method as claimed in any of claims 1 to 7, wherein, when the loaded portion does not include a complete frame, the method further comprises loading a second portion of the stored multimedia program.
10. A method as claimed in claim 1, wherein the multimedia program is provided in an MPEG stream comprising video packets of data, the method further comprising:
 - determining a data read size and a starting read position in the stored MPEG stream based on the seek input;
 - loading a portion of the stored MPEG stream, the size of the loaded portion being based on the data read size, and the position of the loaded portion being based on the starting read position;

analyzing the loaded portion to determine if the loaded portion includes a complete MPEG intra-coded frame; and

if the loaded portion includes a complete MPEG intra-coded frame, decoding the MPEG intra-coded frame to provide a video frame for presentation.

11. A method as claimed in any preceding claim, wherein the program has been stored by receiving the program into an input buffer, and storing the program from the input buffer in a storage device without, prior to storage, analyzing data in the program to generate indexing information.

12. A digital video recorder comprising:

a storage device operable to store a multimedia program without indexing;

a user interface operable to receive a seek input associated with non-contiguous playback of a multimedia program;

a processor operable, based on statistical info associated with the program, to determine a seek position and a data read size associated with the seek input, wherein the statistical information includes at least an average frame size of the multimedia program and an average size of a group of pictures in the multimedia program;

a read buffer operable to receive a portion of the stored multimedia program, wherein the portion loaded begins at the seek position and is of the data read size;

the processor operable to determine whether the loaded portion of the stored multimedia program includes a complete I-frame of data, and responsive to determining that the loaded portion includes the complete I-frame, outputting the loaded portion commencing at, and including at least, the I-frame.

13. The digital video recorder of claim 12, wherein the processor is operable to:

output video data in the multimedia program sequentially at a normal presentation rate; and

generate the statistical information based on the normal presentation rate of the multimedia program.

14. A digital video recorder of claim 12, wherein the statistical information has been pre-generated and is received from a service provider.

15. A digital video recorder as claimed in any of claims 12 to 14, wherein the data read size is determined based on the average size of a group of pictures in the multimedia program.

16. A digital video recorder as claimed in claim 15, wherein the data read size is determined based on twice the average size of a group of pictures in the multimedia program.

17. The digital video recorder as claimed in any of claims 12 to 16, wherein the seek position is determined by calculating a seek vector equal to the product of the frames to be skipped and the average frame size, adding the calculated seek vector to the current file pointer position, and adjusting the result by half of the average group of picture size.

18. A digital video recorder as claimed in claim 17, wherein the adjustment is determined based on half the average size of the group of pictures.

19. A digital video recorder as claimed in any of claims 12 to 18, wherein, when the loaded portion does not include a complete frame, the processor further operable to load a second portion of the stored multimedia program.

20. A computer-readable medium carrying instructions for enabling the presentation of a stored multimedia program, the instructions comprising:

receiving a seek input associated with non-contiguous playback of a multimedia program;

based on statistical info associated with the program, determining a seek position and a data read size associated with the seek input, wherein the statistical information includes at least an average frame size of the multimedia program and an average size of a group of pictures in the multimedia program;

loading a portion of the stored multimedia program, wherein the portion loaded begins at the seek position and is of the data read size;

determining whether the loaded portion of the stored multimedia program includes a complete I-frame of data; and

responsive to determining that the loaded portion includes the complete I-frame, outputting the loaded portion commencing at, and including at least, the I-frame.

21. The computer-readable medium of claim 20, wherein the instructions further comprise:

outputting video data in the multimedia program sequentially at a normal presentation rate; and

generating the statistical information based on the normal presentation rate of the multimedia program.

22. The computer-readable medium of claim 20, wherein the data read size is determined based on twice the average size of a group of pictures in the multimedia program.

23. The computer-readable medium as claimed in any preceding claim, wherein the seek position is determined by calculating a seek vector equal to the product of the frames to be skipped and the average frame size, adding the calculated seek vector to the current file pointer position, and adjusting the result by half of the average group of picture size.