



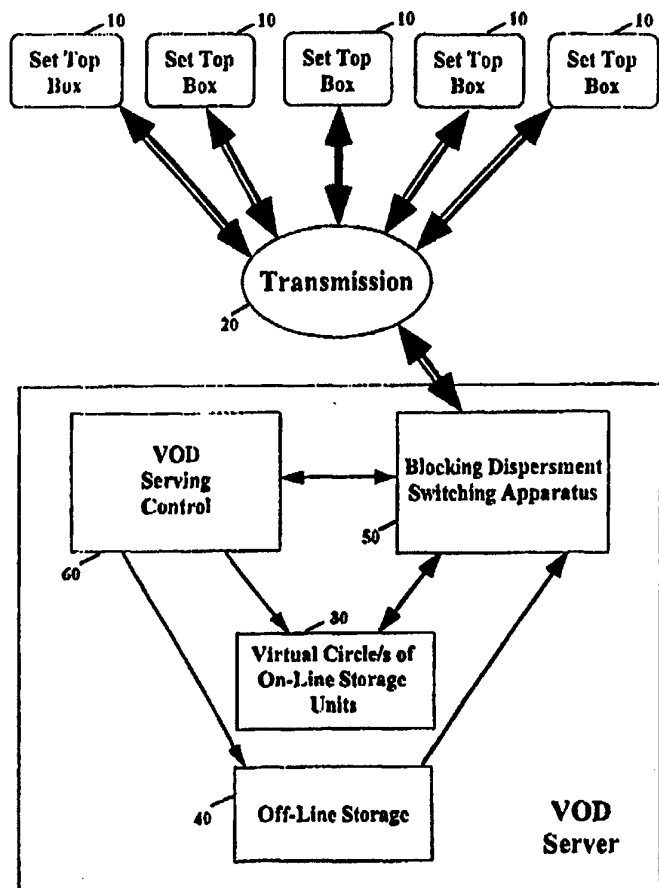
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification <sup>6</sup> : <b>H04N</b></p>	<p><b>A2</b></p>	<p>(11) International Publication Number: <b>WO 95/26095</b> (43) International Publication Date: 28 September 1995 (28.09.95)</p>
<p>(21) International Application Number: PCT/US95/03237 (22) International Filing Date: 13 March 1995 (13.03.95) (30) Priority Data: 108,952 13 March 1994 (13.03.94) IL (71) Applicant (for all designated States except US): SHAKED, Amir [US/US]; 50 Eldorado Boulevard, Plainview, NY 11803 (US). (71)(72) Applicant and Inventor: SANDLERMAN, Nimrod [IL/IL]; 11 Hamatmid Street, Ramat Gan 52493 (IL). (74) Agents: GALLOWAY, Peter, D. et al.; Ladas &amp; Parry, 26 West 61st Street, New York, NY 10023 (US).</p>		<p>(81) Designated States: JP, US, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).  <b>Published</b> <i>Without international search report and to be republished upon receipt of that report.</i></p>

(54) Title: VIDEO SYSTEM

(57) Abstract

This invention discloses a high throughput video serving system operative to serve a multiplicity of set-top boxes with a corresponding multiplicity of selected video items via transmission equipment comprising an on-line storage subsystem comprising a first plurality of on-line storage units arranged in at least one virtual circle and storing a second plurality of video items, each video item comprising a sequence of video blocks, each video block being accessible within a single system tick, wherein, for each individual video item, sequential blocks are stored in sequential on-line storage units in each virtual circle and an access system connecting the on-line storage subsystem to the multiplicity of set-top boxes via the transmission equipment, wherein an individual set-top box is initially connected to the on-line storage unit which contains the first block of its corresponding selected video item and wherein the individual set-top box is subsequently connected to a sequential on-line storage unit at each sequential system tick.



**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

## VIDEO SYSTEM

## FIELD OF THE INVENTION

The present invention relates to video storage and serving systems.

## BACKGROUND OF THE INVENTION

There is a growing demand in the video market for high throughput video systems. Such systems are necessary in implementing VOD (video on demand) systems, video home shopping, video mail, medical image distribution and archiving, corporate video archiving and other sophisticated video applications.

## SUMMARY OF THE INVENTION

The present invention seeks to provide an improved high throughput video serving system.

There is thus provided in accordance with a preferred embodiment of the present invention a high throughput video serving system operative to serve a multiplicity of set-top boxes with a corresponding multiplicity of selected video items via transmission equipment including an on-line storage subsystem including a first plurality of on-line storage units arranged in at least one virtual circle and storing a second plurality of video items, each video item including a sequence of video blocks, each video block being accessible within a single system tick, wherein, for each individual video item, sequential blocks are stored in sequential on-line storage units in each virtual circle, and an access system connecting the on-line storage subsystem to the multiplicity of set-top boxes via the transmission equipment, wherein an individual set-top box is initially connected to the on-line storage unit which contains the first block of its corresponding selected video item and wherein the individual set-top box is subsequently connected to a sequential on-line storage unit at each sequential system tick.

Further in accordance with a preferred embodiment of the present invention each on-line storage unit includes a disk.

Still further in accordance with a preferred embodiment of the present invention the plurality of on-line storage units includes at least one electronic memory.

Additionally in accordance with a preferred embodiment of the present invention the system also

includes an off-line storage subsystem storing additional video items which are transferable to the on-line storage units.

Yet further in accordance with a preferred embodiment of the present invention the access system includes first and second switching arrays, each switching array including a plurality of switching elements, each on-line storage unit communicating with one of the switching elements in the first switching array, each set-top box communicating with one of the switching elements in the second switching array, each of the switching elements in the first switching array communicating with all of the switching elements in the second switching array.

There is further provided in accordance with a preferred embodiment of the present invention an information serving system operative to serve a user-selected one of a plurality of data items, each data item including a first sequence of data elements and a subsequent second sequence of data elements, the system including a first information storage system with a first block size storing the first sequence of data elements for each data item, and a second information storage system with a second block size, larger than said first block size, storing the second sequence of data elements for each data item, and a user request processor operative to serve the first sequence of data elements and, immediately afterward, the second sequence of data elements, for a selected data item.

Further in accordance with a preferred embodiment of the present invention each on-line storage unit includes a RAID system.

Still further in accordance with a preferred embodiment of the present invention the at least one virtual circle includes a plurality of virtual circles.

Yet further in accordance with a preferred

embodiment of the present invention the plurality of virtual circles includes a first virtual circle of on-line storage units of a first type and a second virtual circle of on-line storage units of a second type.

Additionally in accordance with a preferred embodiment of the present invention the plurality of virtual circles includes a first virtual circle of on-line storage units storing video items having a first bit rate and a second virtual circle of on-line storage units storing video items having a second bit rate.

Further in accordance with a preferred embodiment of the present invention the plurality of on-line storage units within an individual virtual circle is operative to store video items of more than one bit rate.

There is additionally provided in accordance with a preferred embodiment of the present invention a high throughput video serving method operative to serve a multiplicity of set-top boxes with a corresponding multiplicity of selected video items via transmission equipment including providing an on-line storage subsystem including a first plurality of on-line storage units arranged in at least one virtual circle, storing a second plurality of video items, each video item including a sequence of video blocks, each video block being accessible within a single system tick, wherein, for each individual video item, sequential blocks are stored in sequential on-line storage units in each virtual circle, and connecting the on-line storage subsystem to the multiplicity of set-top boxes via the transmission equipment, including connecting an individual set-top box initially to the on-line storage unit which contains the first block of its corresponding selected video item and connecting the individual set-top-box subsequently to a sequential on-line storage unit at each sequential system tick.

Further in accordance with a preferred embodi-

ment of the present invention there is provided an information serving method operative to serve a user-selected one of a plurality of data items, each data item including a first sequence of data elements and a subsequent second sequence of data elements, the method including storing the first sequence of data elements for each data item in a first information storage system with a first block size, and storing the second sequence of data elements for each data item in a second information storage system with a second block size, larger than said first block size, and serving the first sequence of data elements and, immediately afterward, the second sequence of data elements, for a selected data item.

Additionally in accordance with a preferred embodiment of the present invention there is provided a high throughput information serving system operative to serve a multiplicity of users with a corresponding multiplicity of selected information items via transmission equipment including an on-line storage subsystem including a first plurality of on-line storage units arranged in at least one virtual circle and storing a second plurality of information items, each information item including a sequence of information blocks, each information block being accessible within a single system tick, wherein, for each individual information item, sequential blocks are stored in sequential on-line storage units in each virtual circle, and an access system connecting the on-line storage subsystem to the multiplicity of users via the transmission equipment, wherein an individual user is initially connected to the on-line storage unit which contains the first block of its corresponding selected information item and wherein the individual user is subsequently connected to a sequential on-line storage unit at each sequential system tick.

Further in accordance with a preferred embodiment of the present invention each on-line storage unit

includes a disk.

Still further in accordance with a preferred embodiment of the present invention the plurality of on-line storage units includes at least one electronic memory.

Yet further in accordance with a preferred embodiment of the present invention the system also includes an off-line storage subsystem storing additional information items which are transferable to the on-line storage units.

Additionally in accordance with a preferred embodiment of the present invention the access system includes first and second switching arrays, each switching array including a plurality of switching elements, each on-line storage unit communicating with one of the switching elements in the first switching array, each user communicating with one of the switching elements in the second switching array, each of the switching elements in the first switching array communicating with all of the switching elements in the second switching array.

There is further provided in accordance with a preferred embodiment of the present invention an information serving system operative to serve a user-selected one of a plurality of data items, each data item including a first sequence of data elements and a subsequent second sequence of data elements, the system including a first information storage system with a first block size storing the first sequence of data elements for each data item, and a second information storage system with a second block size, larger than said first block size, storing the second sequence of data elements for each data item, and a user request processor operative to serve the first sequence of data elements and, immediately afterward, the second sequence of data elements, for a selected data item.



Further in accordance with a preferred embodiment of the present invention each on-line storage unit includes a RAID system.

Still further in accordance with a preferred embodiment of the present invention the at least one virtual circle includes a plurality of virtual circles.

Additionally in accordance with a preferred embodiment of the present invention the plurality of virtual circles includes a first virtual circle of on-line storage units of a first type and a second virtual circle of on-line storage units of a second type.

Further in accordance with a preferred embodiment of the present invention the plurality of virtual circles includes a first virtual circle of on-line storage units storing information items having a first bit rate and a second virtual circle of on-line storage units storing information items having a second bit rate.

Additionally in accordance with a preferred embodiment of the present invention the plurality of on-line storage units within an individual virtual circle is operative to store information items of more than one bit rate.

## BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be understood and appreciated from the following detailed description, taken in conjunction with the drawings in which:

Fig. 1A is a simplified block diagram of a simplified block diagram of a high throughput VOD (video on demand) system constructed and operative in accordance with a preferred embodiment of the present invention;

Fig. 1B is a more detailed block diagram of a preferred implementation of the apparatus of Fig. 1A;

Fig. 2 is a simplified block diagram of a virtual circle of on-line storage units;

Fig. 3 is a simplified block diagram of a variation on the apparatus of Fig. 2;

Fig. 4 is a simplified block diagram of switching apparatus 50 constructed and operative in accordance with a preferred embodiment of the present invention;

Fig. 5 is a detailed schematic illustration of output switch array 110;

Fig. 6 is a simplified block diagram of the VOD serving control unit 60 of Figs. 1A and 1B;

Fig. 7 is a simplified flowchart of a preferred method of loading a video item from the off-line storage unit 40 of Figs. 1A and 1B onto the on-line storage unit 30 of Figs. 1A and 1B;

Fig. 8A is a simplified flowchart of a method for performing the system tick selection step of Fig. 7, in which the first suitable system tick is located;

Fig. 8B is a simplified flowchart of another method for performing the system tick selection step of Fig. 7, in which the system tick having the best load dispersion within a predetermined time period is located;

Fig. 9 is a simplified flowchart of a preferred method for playing a video item stored in the on-line

storage unit 70 of Fig. 3;

Fig. 10 is a simplified flowchart of a preferred method by which the off-line drive transmits a video item from off-line storage 40 to on-line storage unit 30, in accordance with the virtual circular structure of on-line storage unit 30;

Fig. 11 is a modification of Fig. 4 illustrating the connection between switching apparatus 50 and the on-line storage units of on-line storage apparatus 30 for the embodiment in which there are two virtual circles in on-line storage apparatus 30;

Fig. 12 is a simplified diagram of switching apparatus 50 constructed and operative in accordance with another embodiment of the present invention;

Fig. 13 is a detailed illustration of the second switching array 110 of Fig. 12; and

Fig. 14 is a simplified block diagram of two of the controllers 80 of Fig. 3 with their associated on-line storage units 70 and their redundancy connections to their respective on-line storage units 70.

## DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT

Reference is now made to Fig. 1A which is a simplified block diagram of a high throughput VOD (video on demand) system constructed and operative in accordance with a preferred embodiment of the present invention. Reference is also made to Fig. 1B which is a more detailed block diagram of a preferred implementation of the apparatus of Fig. 1A. Reference is further made to Fig. 2 which is a simplified block diagram of a virtual circle 29 of on-line storage units. One or more such virtual circles may be included in the on-line storage apparatus of Fig. 1A - 1B.

The apparatus of Figs. 1A -1B and 2 is operative to serve a multiplicity of set-top boxes 10 installed in association with a multiplicity of video items selected by the users via transmission equipment 20. It is appreciated that the term set-top box, used throughout this application, represents a video terminal, such as a television set, PC's (personal computers), multimedia workstations, medical imaging equipment, etc. It is also appreciated that more than one user may elect to view the same video item at one time. Typically, although not necessarily, different viewers viewing the same video item do not begin viewing at the same time and/or request different pauses while viewing. Therefore, viewers electing the same video item typically although not necessarily are each viewing a different portion within the same video item.

The apparatus of Figs. 1A - 1B includes on-line storage apparatus 30, off-line storage apparatus 40, blocking dispersement switching apparatus 50 and VOD serving control apparatus 60.

The on-line storage apparatus 30 of Figs. 1A - 1B is illustrated in more detail in Fig. 2. As shown in

Fig. 2, the on-line storage apparatus 30 includes a first plurality of on-line storage units 70, illustrated as ellipses and labeled 1 to 24 in Fig. 1B, associated with switching apparatus 50 of Figs. 1A - 1B, typically via a corresponding first plurality of controllers 80, labeled I, II, III..XII. The on-line storage unit 70 may be based on hard disks or on RAID systems or on any combination thereof.

For example, each on-line storage unit 70 may comprise one or more hard disks or one or more Concept 500 RAID systems, commercially available from Storage Concepts, Irvine, CA, USA and each controller 80 may comprise a conventional PC equipped with an ATM board such as the ForeRunner ESA 200, commercially available from Fore Systems. Another way to implement a RAID-based on-line storage unit 70 and associated controller 80 is by means of a plurality of hard disks, a PC equipped with one or more ULTRA 124F boards, commercially available from UltraStor Corporation, Irvine, CA, USA and an ATM board such as the ForeRunner ESA 200.

Alternatively, each on-line storage unit 70 may comprise a memory unit, such as a PC equipped with the ForeRunner ESA 200 ATM board and a RAM or Flash Memory.

Typically, the on-line storage apparatus 30 of Figs. 1A and 1B includes at least one virtual circle of hard disks or RAID systems or at least one virtual circle of memory units, or both.

The on-line storage units 70 are arranged in the virtual circle 29, Fig. 2, and are operative to store a second plurality of video items, such as movies, referenced A, B, C, D, .... Each video item comprises a sequence of video blocks, typically of equal size, containing video information of equal time length, referenced 1, 2, 3, 4, ..., each video block being accessible within a single system tick; a system tick is defined hereinbelow. For example, video item A comprises

the following sequence of video blocks: A1, A2, A3, A4, ....

The entire system is synchronized according to a fixed system clock or a "system tick". The period between two consecutive system ticks is typically of the order of 1 second. All the video blocks, namely A1, A2, A13, B1, B2, C11, etc., have a video time length of exactly "1 tick period". All terminals which are connected to their respective storage units 70 are serviced within one tick period. This means that each storage unit 70 reads and transmits the required video block, as illustrated in Fig. 2, to its respective video terminal or terminals. At the next tick (at the end of the present tick period) all terminals are switched to the next storage unit 70 to receive the next video block, namely A2, A14, B1, B13 etc. of controller II.

For each individual video item, sequential blocks are stored in sequential on-line storage units 70 in the virtual circle 29. The sequentiality of the on-line storage units 70 refers to their sequence in a particular direction around the virtual circle 29. In the illustrated embodiment, the sequentiality is defined as clockwise, as may be appreciated, for example, by an inspection of the on-line storage units 70 in which video blocks A1, A2, A3, A4, ... are stored.

The virtual circle arrangement of the on-line storage units 70 may, for example, be implemented in the software of VOD serving control unit 60 of Figs. 1A and 1B.

The first blocks A1, B1, C1, ... of the video items A, B and C may be stored in randomly selected respective ones of the on-line storage units 70 or alternatively may be stored in sequential respective ones of the on-line storage units 70. In the illustrated embodiment, the first blocks A1, B1, C1, ..., are stored in sequential on-line storage units.

The switching apparatus 50 connects the on-line storage apparatus 30 to the multiplicity of set-top boxes 10 via the transmission equipment 20. An individual set-top box 10, serving a user who has selected a video item (for example, video item A) is initially connected to the on-line storage unit 70 which contains the first block of the selected video item, i.e. block A1. The individual set-top box 10 is subsequently connected to a sequential on-line storage unit 70 at each sequential system tick. In other words, a set-top box 10 serving a user who has selected video item A is connected, at a first system tick following selection of video item A, to the on-line storage unit 70 containing A1. At the second, third, fourth, etc. system ticks immediately following the first system tick, the same user is connected to the on-line storage units 70 containing A2, A3, A4, etc. which, as may be appreciated from Fig. 2, are arranged sequentially in the direction of the virtual circle 29.

All switching of connections of individual set-top boxes 10 from one on-line storage unit 70 to the adjacent storage unit 70 around the virtual circle 29 take place synchronously, at the beginning of each system tick. The scans of all on-line storage units 70 begin within the current system tick and terminate before the next system tick. The term "scan" of an individual on-line storage unit 70 refers to the execution of the list of I/O requests, for that system tick.

Each I/O request may comprise a request for retrieval of a video block stored in the individual on-line storage unit 70 or a request to load a video block stored in the off-line storage apparatus 40 onto the individual on-line storage unit 70. The video blocks handled by each I/O request are preferably all the same size and contains the same length of video playing time.

All on-line storage units 70 execute, at most, a maximum number of I/O requests per scan. Typically, the

on-line storage units 70 all have the same maximum bit rate which determines the maximum number of I/O requests per scan.

Typically, all video items are transmitted at the same bit rate. Alternatively, the bit rates at which the various video items are transmitted may all be integer multiples of a basic bit rate. In this case, each I/O request comprises a request for retrieval or loading of a block of the same basic block size but for the different bit rates, different numbers of I/O requests are carried out each system tick. For higher bit rates, more I/O requests are carried out each system tick. Bit rates that are not exact multiplications of the basic bit rate are also possible, however, this will reduce the overall system throughput.

Having excess memory buffers, preferably in all on-line storage controllers 80 or in the specific set-top box 10, enables the system to perform the requested I/Os to the same on-line storage units within several system ticks instead of the same system tick.

Although, for simplicity, only 12 on-line storage units 70 are illustrated in Fig. 2, in fact, any suitable number of on-line storage units 70 may be provided, such as hundreds or thousands of units. For example, if there are 10,000 users and hard disks capable of serving 5 users simultaneously are employed as on-line storage units, then 2000 on-line storage units 70 may be employed. If RAID systems, each including 5 hard disks each capable of serving 5 users simultaneously are employed as on-line storage units 70, then 400 on-line storage units may be employed.

Reference is now made to Fig. 3 which is a variation on the apparatus of Fig. 2. Fig. 3 is similar to Fig. 2 except that, in order to completely utilize the bandwidth of each port 75, 76, 77, 78 of the switching apparatus 50, more than one on-line storage unit 70 is



associated with each port of the switching apparatus 50, via a suitable controller 80 such as a PC equipped with one or more ULTRA 124F boards and with an ATM board such as the ForeRunner ESA 200. In the context of the present invention, bandwidth is defined as an amount of information which can be transferred by a port during a system tick.

Reference is now made to Fig. 4 which is a simplified block diagram of switching apparatus 50 constructed and operative in accordance with a preferred embodiment of the present invention which is useful when the number of controllers 80 exceeds the number of input ports of conventional switches. The switching apparatus of Fig. 4 employs  $2n$  conventional switches to obtain a switching capacity which is  $n$  times the capacity of each of the conventional switches.

For simplicity, the switching apparatus 50 is shown with 64 inputs, labeled 1 to 64 in the drawing, accepting video input from 64 respective controllers 80 whose associated on-line storage units 70 are arranged in a virtual circle from 1 to 64.

The switching apparatus 50 includes an input switch array 100 receiving video input from the on-line storage units, 70 typically via the controllers 80, and an output switch array 110, illustrated in detail in Fig. 5. Output switch array 110 receives input from the input switch array 100 and transmits this input to the set-top boxes via the transmission equipment 20 of Figs. 1A and 1B.

The input switch array 100 includes 8 input switches 120 and the output switch array 110 includes 8 output switches 130. Each of switches 120 and 130 typically comprises an ATM switch having 16 155Mbps trunks which are utilized as 8 input trunks and 8 output trunks. For example, a suitable ATM switch is the ATOM 1E6, commercially available from Telematics Internation-

al, Fort Lauderdale, FL, USA.

Each on-line storage unit 70 communicates with one of the input switches in the input switch array 100 via its associated controller 80. Preferably, the first on-line storage unit 70 and every eighth on-line storage unit 70 thereafter (units 1, 9, 17, ..., 57) communicate with the first input switch SW1. The second on-line storage unit 70 and every eighth on-line storage unit 70 thereafter (units 2, 10, ..., 58) communicate with the second input switch SW2, and so on, so that the eighth on-line storage unit 70 and every eighth on-line storage unit 70 thereafter (units 8, 16, ..., 64) communicate with the eighth input switch SW8. Each set-top box 10 communicates, via transmission 20, with one of the output switches in the output switch array 110. Each of the input switches communicates with all of the output switches, as illustrated in Fig. 1B.

It is appreciated that the terms "input" and "output" as employed herein refer to the dataflow of the video items. However, other information flows in the opposite direction, i.e. from the output side of the switching apparatus 50 to the input side thereof. This information may include user controls such as selections of video items, pause commands, and resume commands; as well as loading of video items from off-line storage apparatus 40 onto the on-line storage apparatus 30.

Although the switching apparatus of Fig. 4 is a blocking switch, this feature does not substantially adversely affect the performance of the system due to the virtual circle structure of the apparatus of Fig. 2. Although the switching system can be blocked, the blocking of the system is "dispersed" throughout the system such that the connection between each individual input switch and each individual output switch, if blocked at an individual system tick, becomes available within at most 7 system ticks, in the illustrated embodiment.

Fig. 6 is a simplified block diagram of the VOD serving control unit 60 of Figs. 1A and 1B and is self-explanatory.

Fig. 7 is a simplified flowchart of a preferred method of loading a video item from the off-line storage unit 40 of Figs. 1A and 1B onto the on-line storage unit 70, via the on-line storage apparatus 30 of Figs. 1A and 1B.

When a request for a video item is made, the system enters the procedure for loading a video item at step 200 in Fig. 7. The system asks in step 210 if there is sufficient on-line storage space available in on-line storage apparatus 30. If there is no storage space available the system will deny the request, as in step 220. If the answer to step 210 is positive, the off-line storage device 40 is instructed to operate a loading device, such as a jukebox-type apparatus, as known in the art, to load the requested video item, such as a video tape cassette or a CD-ROM device, on to its associated drive unit. In step 230, an on-line unit 70 is selected. The selected unit 70 will contain the first video block of the video item to be loaded. In this implementation the selected unit 70 is the one following the first unit of the video item that was previously loaded. Step 240 assigns the correct system tick for commencing the loading of the video item. This step is described in more detail, hereinbelow, with respect to Figs. 8A and 8B. In step 250, the off-line drive of the off-line storage unit 40 is instructed to commence transmission in to the selected on-line storage unit 70, at the selected system tick.

Reference is now made to Figs. 8A and 8B which are simplified flowcharts of two alternative methods for performing the system tick selection process of step 240 of Fig. 7. In the flowcharts of Figs. 8A and 8B the following abbreviations are used:

STE = slot table entry, where the slot table is

defined as the lookup table which lists the available throughput capacity for storage unit 70 for the next system tick;

OSU = on-line storage unit 70;

SLOTS = maximum number of slots per storage unit;

TICKS = number of system ticks to wait for available throughput at First On-Line Storage unit;

BW(n,m) = used bandwidth between input switch element (n) associated with on-line storage element of unit 70 (i) and output switch element (m) associated with a user channel (k);

MBW = maximum possible bandwidth between each two switch elements.

In step 300, a play or load of a video item, A, B., is requested. In step 310 the parameter OSU is set to the first on-line storage unit 70 as selected by the procedure described in Fig. 7 and parameter TICKS is set equal to 1, denoting one tick delay. In step 320, the system requests if there is a slot available from or to the on-line storage unit 70 OSU. If the answer is in the affirmative the system requests to know if there is sufficient bandwidth available between an input switch element n and an output switch element m for transferring information; n is connected to the on-line storage unit OSU and output switch element m is connected to the subscriber at channel k. If the answer, to either step 320 or 330, is negative, the system looks for an available throughput to or from on-line storage unit OSU at the next system tick. This is done by incrementing TICKS by 1 and interrogating the load in the current tick at the previous storage unit 70, by decreasing OSU. If the answers to steps 320 and 330 are both affirmative, the slot to on-line storage unit OSU and the bandwidth between the switch elements n and m are assigned to channel k. The slot at STE(OSU) is allocated an the value

of  $BW(n,m)$  is incremented by 1. The task is completed in step 360.

Fig. 8B illustrates a proposed alternative method of determining the system tick. This method has not been fully developed or tested and may contain errors.

Reference is now made to Fig. 9 which is a simplified flowchart of a preferred method for playing or projecting a video item stored in the on-line storage units 30 of Fig. 3. The table of on-line video items A, B.. contains the names of the video items with their respective first on-line storage unit 70 where the first block of the video title, say A1 of video item A and the total number of blocks, A1, A2.. for the video item are stored. In the flowchart of Fig. 9 the following abbreviations are used:

OSU = On-line storage unit;

First OSU = the OSU where the first block is stored;

TICKS = required delay;

Video Item (i,j) = block number j of video item i at on-line storage unit (OSU).

The procedure of playing a video item begins in step 400. In step 410 the system interrogates the table of On-Line Video Items and finds the on-line storage unit 70 where the video item starts, namely unit OSU and the length of the video item, labeled LENGTH. The parameter LENGTH is the number of blocks that comprise the video item. In 420 the system calls the procedure described in either Figs. 8A or 8B to allocate the system ticks for the playing or projection procedure of the video item i. The system tick allocation procedure returns the number of system ticks to wait until the required throughput is available for on-line storage unit OSU. In steps 430 and 440 the system waits the required number of system ticks. After the required time has

elapsed the system proceeds to step 450 and instructs the on-line storage unit OSU to transmit video item (i;j) to channel k. As each video item (i,j) is transmitted the parameter LENGTH is decremented by 1, in step 450. This action is repeated in steps 460, until all the video items (i,j) are transmitted by checking if LENGTH=0, step 460. If LENGTH equals 0 the video item has been transmitted and the playing/projection is completed, step 470. If the answer to step 460 is negative, the system proceeds to step 480. After the transmission of each video block (i,j), the system waits one tick unit, the parameters OSU and j are increased by 1, step 480, and the unit 70 proceeds around the virtual circle 29, as illustrated in Fig. 2.

Fig. 10 is a simplified flowchart of a preferred method by which an off-line drive transmits a video item from off-line storage 40 to on-line storage apparatus 30, in accordance with the virtual circular structure of on-line storage unit 30.

The system after allocating in Fig. 7 the on-line storage unit 70 where loading should commence and the system tick is defined when throughput is available to the desired storage unit 70, the transfer of a video item begins, step 500. In step 510, the first block i is written to the first of on-line storage unit 70, via on-line storage apparatus 30. The procedure is repeated, as in steps 520 and 530, until all the video item has been transferred. At the transfer of each video block, the block number i and block storage unit number j, are incremented by 1. At each stage of the transfer video block i is written to on-line storage unit j. On completion of the transfer of the video item, the number of blocks for the given item is returned and stored in the VOD Control unit, as illustrated in Fig. 6, step 540. The video item transfer ends at step 550.

Reference is now made briefly to Fig. 11. As

described above, the on-line storage apparatus 30 of Figs. 1A and 1B may include two or more virtual circles 31 and 32, preferably one virtual circle of hard disks or RAID systems and preferably one virtual circle of memory units. Fig. 11 is a modification of Fig. 4 illustrating the connection between switching apparatus 50 and the on-line storage units of on-line storage apparatus 30 for the embodiment in which there are two virtual circles 31 and 32 in on-line storage apparatus 30.

Fig. 11 also illustrates use of redundancy to prevent malfunctioning of the entire system due to malfunctioning of a single one of the input switches or a single one of the controllers 80. The redundancy feature is expressed that each on-line storage unit 70 and its corresponding controller 80 are interconnected amongst the other elements of the virtual circle by a chain-like formation, as illustrated in Fig. 11. It is appreciated that the two-circle feature and the redundancy feature are independent. In other words, use of the redundancy feature illustrated in Fig. 11 is not limited to the two circle embodiment but instead is equally applicable, for example, to the embodiment of Fig. 4.

According to a preferred embodiment of the present invention, the memory virtual circle is not separate but rather is parallel to the disk storage virtual circle. Each controller 80 includes memory over and above the memory required for the buffers such that the number of buffers in the memory of each controller 80 is greater than the maximum allowed number of disk I/O's per scan. Popular video items which are viewed simultaneously more than a certain number of times are called to memory and stored there, even if they are not required for the next scan. I/O requests for a video item which is stored in memory in its entirety are not taken into account when computing the limitation on requests

per scan.

Preferably, the controllers 80 are redundantly associated with the on-line storage units. Specifically:

a. At least one of the controllers 80 is redundant (124) in that it is not intended to serve any on-line storage unit or units unless one of the non-redundant controllers fails.

b. Each controller 80 is connected not only to the on-line storage unit or units which it is intended to serve but also to the on-line storage unit or units served by a second controller, so that the first controller can substitute for the second controller, if it fails. The extra connections between "backing up" controllers and "backed up" controllers typically forms a chain back to the redundant controller, as illustrated in Fig. 11, such that each time a controller A fails, its "backing up" controller B substitutes for it, whereas the "backing up" controller C of controller B substitutes for B, and so on, back to the redundant controller. Fig. 11 shows 8 chains going back to 8 redundant controllers which are depicted as blank (numberless) ovals 124.

The "backing up" controller and the controller it backs up are preferably connected to different input switching elements, as illustrated in Fig. 11, such that the redundancy of controllers provides a solution for inoperative input switching elements as well as for inoperative controllers.

Reference is now made to Fig. 12 which is a simplified diagram of switching apparatus 50 constructed and operative in accordance with another embodiment of the present invention in which the number of input switch elements, labeled 1, 2..32, does not equal the number of input ports per input switch element. Fig. 13 is a detailed illustration of the second switching array 110 of Fig. 12.

It is appreciated that the user may not receive



immediate serving of the video item he has selected. If there is no available throughput to the on-line storage unit 70 storing the first block of the video item he has selected, the system typically waits until the relevant on-line storage unit 70 has available throughput and only then connects the user to that on-line storage unit. Throughput availability is then guaranteed for the on-line storage units storing the second, third, fourth, ... last blocks of the selected video item, assuming that the user does not request any pauses. If the user requests a pause and then requests that the video item be resumed, there may not be available throughput to the on-line storage unit storing the next block of the selected video item. Therefore, the system typically waits until the relevant on-line storage unit 70 has available throughput and only then connects the user to that on-line storage unit.

EXAMPLE: Referring back to Fig. 2 assume that a user wishes to view video item A and that the following conditions hold:

a. Each input switching element can simultaneously transmit up to 3 transmissions to each output switching element. In the illustration, 3 connections between individual input switching elements and output switching elements are indicated by circles, sharp-edged squares and rounded-edged squares, respectively, where the connection which is required to allow the user to view video item A is that indicated by the sharp-edged square. In other words, for each individual on-line storage unit, there can be, at most, only 3 circles, 3 sharp-edged squares and 3 rounded-edged squares.

b. Each on-line storage unit can support up to 8 I/O requests per tick. In other words, for each individual on-line storage unit, the total number of circles, sharp-edged squares and rounded-edged squares must be no more than 8.

The user wishing to view video item A cannot do so during the current system tick because the on-line storage unit I which stores A1, the first block in video item A, is already handling 8 I/O requests.

The same user also cannot be serviced in the next system tick because the on-line storage unit XII which is adjacent, counterclockwise along the virtual circle, to the on-line storage unit which stores A1, already carries 3 I/O requests via the sharp-edged square connection. Even if storage unit XII was not carrying 3 I/O requests via the sharp-edged square connection, but nonetheless 3 such I/O requests were served via the sharp-edged square connection from any storage unit XII at that system tick, the user could not have been serviced in that system tick.

The user can be serviced in the following system tick because:

- a. the on-line storage unit XI is handling less than 8 I/O requests; and
- b. the on-line storage unit XI is carrying less than 3 I/O requests via the sharp-edged square connection.

Reference is now made to Fig. 14 which is a simplified block diagram of two of the controllers 80 of Fig. 3 which are redundantly associated with the on-line storage units 70 as described above with reference to Fig. 11.

It is appreciated that the apparatus shown and described herein is suitable not only for video applications but also for slower bit rate applications such as voice applications and data acquisition applications such as telemetry. For slower bit rate applications, the system tick is typically relatively long, and therefore the response time to a user request is also relatively long. To reduce the response time, more than one virtual

circle of on-line storage units may be provided. For example, one virtual circle of small block-size rapid tick rate on-line storage units may be provided for storing the first block A1, B1, C1, ..., of each data item A, B, C, .... A second virtual circle of medium block-size medium tick-rate on-line storage units may be provided for storing the second block A2, B2, C2, ..., of each data item A, B, C, .... Finally, a third virtual circle of large block-size slow tick-rate on-line storage units may be provided for storing the remaining blocks A3 onward, B3 onward, C3 onward, ... of each data item A, B, C, ....

When a user requests a certain data item, such as item B, the request is typically processed simultaneously for all three virtual circles. Due to the rapid tick rate of the first virtual circle, the first block B1 is provided to the user almost immediately. The second block B2 is ready for the user immediately after serving of block B1 and the block sequence from the third block B3 onward is ready for the user immediately after serving of block B2. It is appreciated that various features of the invention which are, for clarity, described in the contexts of separate embodiments may also be provided in combination in a single embodiment. Conversely, various features of the invention which are, for brevity, described in the context of a single embodiment may also be provided separately or in any suitable subcombination.

It will be appreciated by persons skilled in the art that the present invention is not limited to what has been particularly shown and described hereinabove. Rather, the scope of the present invention is defined only by the claims that follow:

## CLAIMS

1. A high throughput video serving system operative to serve a multiplicity of set-top boxes with a corresponding multiplicity of selected video items via transmission equipment comprising:

an on-line storage subsystem comprising a first plurality of on-line storage units arranged in at least one virtual circle and storing a second plurality of video items, each video item comprising a sequence of video blocks, each video block being accessible within a single system tick, wherein, for each individual video item, sequential blocks are stored in sequential on-line storage units in each virtual circle; and

an access system connecting the on-line storage subsystem to the multiplicity of set-top boxes via the transmission equipment, wherein an individual set-top box is initially connected to the on-line storage unit which contains the first block of its corresponding selected video item and wherein the individual set-top-box is subsequently connected to a sequential on-line storage unit at each sequential system tick.

2. A system according to claim 1 wherein each on-line storage unit comprises a disk.

3. A system according to claim 1 or claim 2 wherein the plurality of on-line storage units comprises at least one electronic memory.

4. A system according to claim 1 and also comprising an off-line storage subsystem storing additional video items which are transferable to the on-line storage units.

5. A system according to claim 1 wherein the access system comprises first and second switching arrays, each switching array comprising a plurality of switching elements, each on-line storage unit communicating with one of the switching elements in the first switching array, each set-top box communicating with one of the switching elements in the second switching array, each of the switching elements in the first switching array communicating with all of the switching elements in the second switching array.

6. An information serving system operative to serve a user-selected one of a plurality of data items, each data item comprising a first sequence of data elements and a subsequent second sequence of data elements, the system comprising:

a first information storage system with a first block size storing the first sequence of data elements for each data item; and

a second information storage system with a second block size, larger than said first block size, storing the second sequence of data elements for each data item; and

a user request processor operative to serve the first sequence of data elements and, immediately afterward, the second sequence of data elements, for a selected data item.

7. A system according to claim 1 wherein each on-line storage unit comprises a RAID system.

8. A system according to claim 1 wherein the at least one virtual circle comprises a plurality of virtual circles.

9. A system according to claim 8 wherein the plurality of virtual circles includes a first virtual circle of on-line storage units of a first type and a second virtual circle of on-line storage units of a second type.

10. A system according to claim 8 wherein the plurality of virtual circles includes a first virtual circle of on-line storage units storing video items having a first bit rate and a second virtual circle of on-line storage units storing video items having a second bit rate.

11. A system according to claim 1 wherein the plurality of on-line storage units within an individual virtual circle is operative to store video items of more than one bit rate.

12. A high throughput video serving method operative to serve a multiplicity of set-top boxes with a corresponding multiplicity of selected video items via transmission equipment comprising:

providing an on-line storage subsystem comprising a first plurality of on-line storage units arranged in at least one virtual circle;

storing a second plurality of video items, each video item comprising a sequence of video blocks, each video block being accessible within a single system tick, wherein, for each individual video item, sequential blocks are stored in sequential on-line storage units in each virtual circle; and

connecting the on-line storage subsystem to the multiplicity of set-top boxes via the transmission equipment, including connecting an individual set-top box initially to the on-line storage unit which contains the first block of its corresponding selected video item and

connecting the individual set-top-box subsequently to a sequential on-line storage unit at each sequential system tick.

13. An information serving method operative to serve a user-selected one of a plurality of data items, each data item comprising a first sequence of data elements and a subsequent second sequence of data elements, the method comprising:

storing the first sequence of data elements for each data item in a first information storage system with a first block size; and

storing the second sequence of data elements for each data item in a second information storage system with a second block size, larger than said first block size; and

serving the first sequence of data elements and, immediately afterward, the second sequence of data elements, for a selected data item.

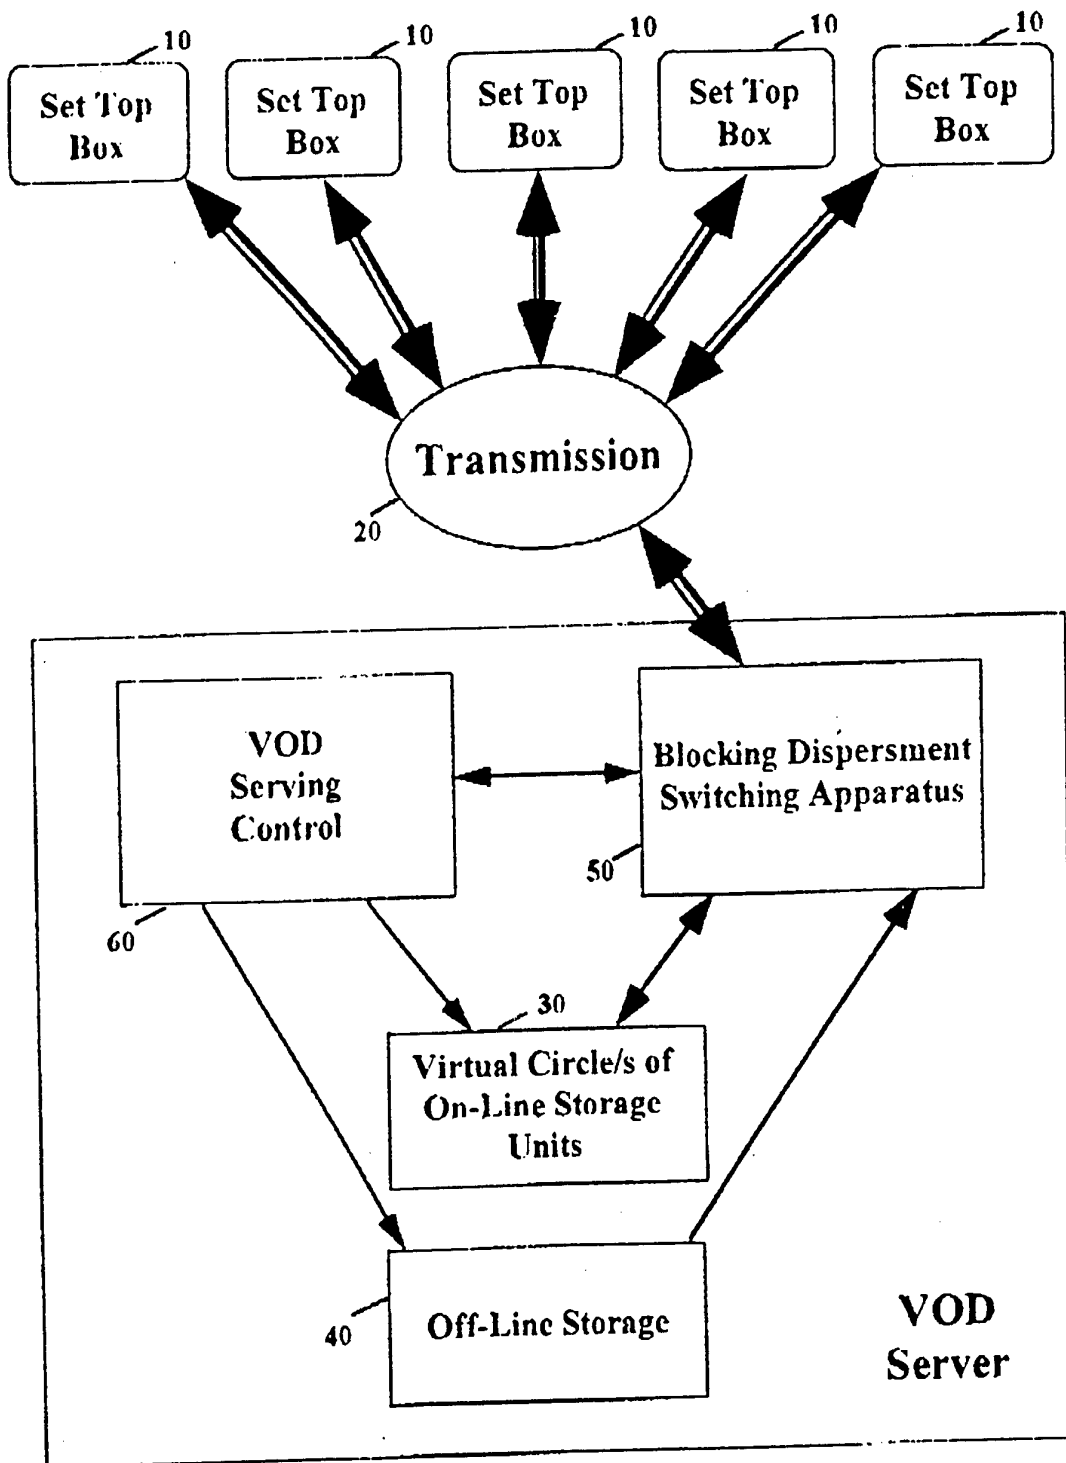


Fig. 1A/14






Fig. 1B/14

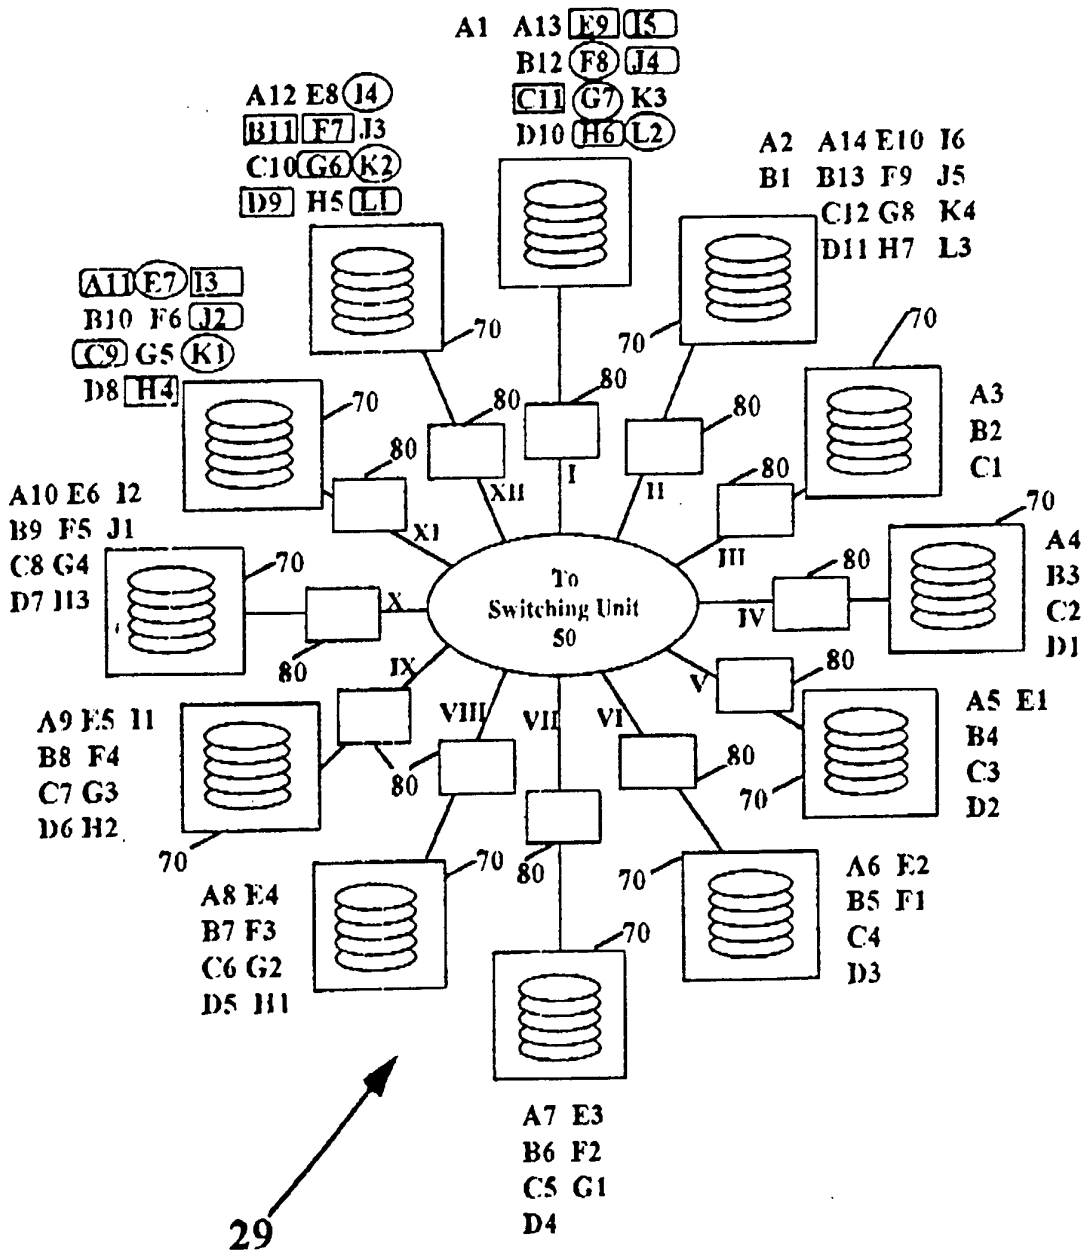


Fig. 2/14




Fig. 3/14




Fig. 4/14




Fig. 5/14

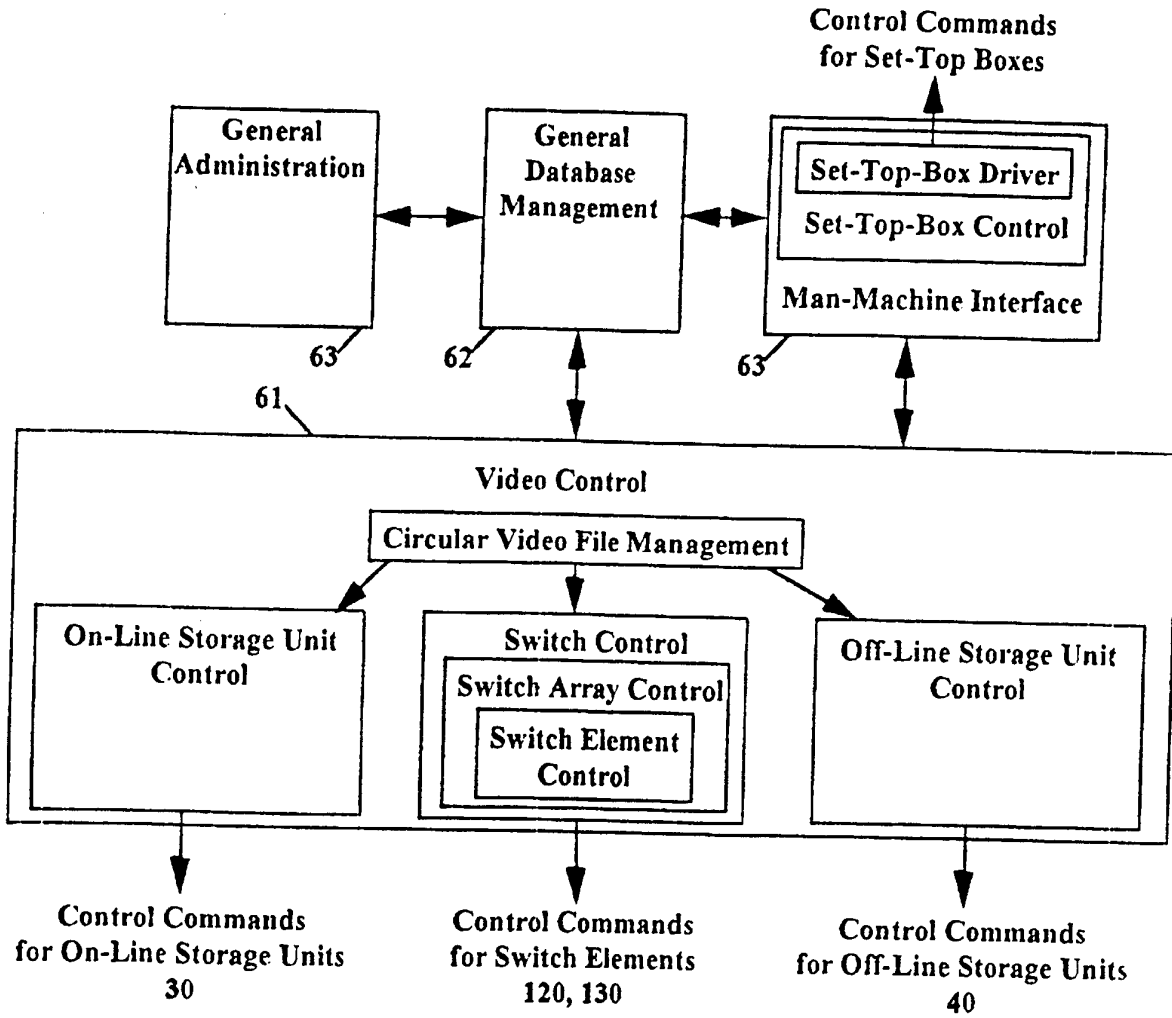


Fig. 6/14

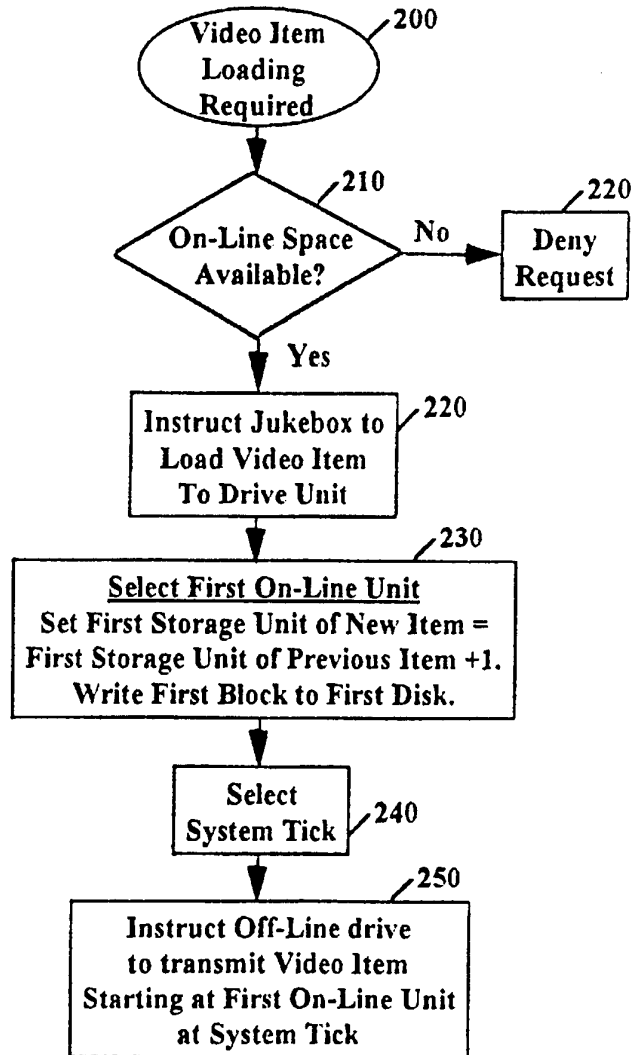


Fig. 7/14

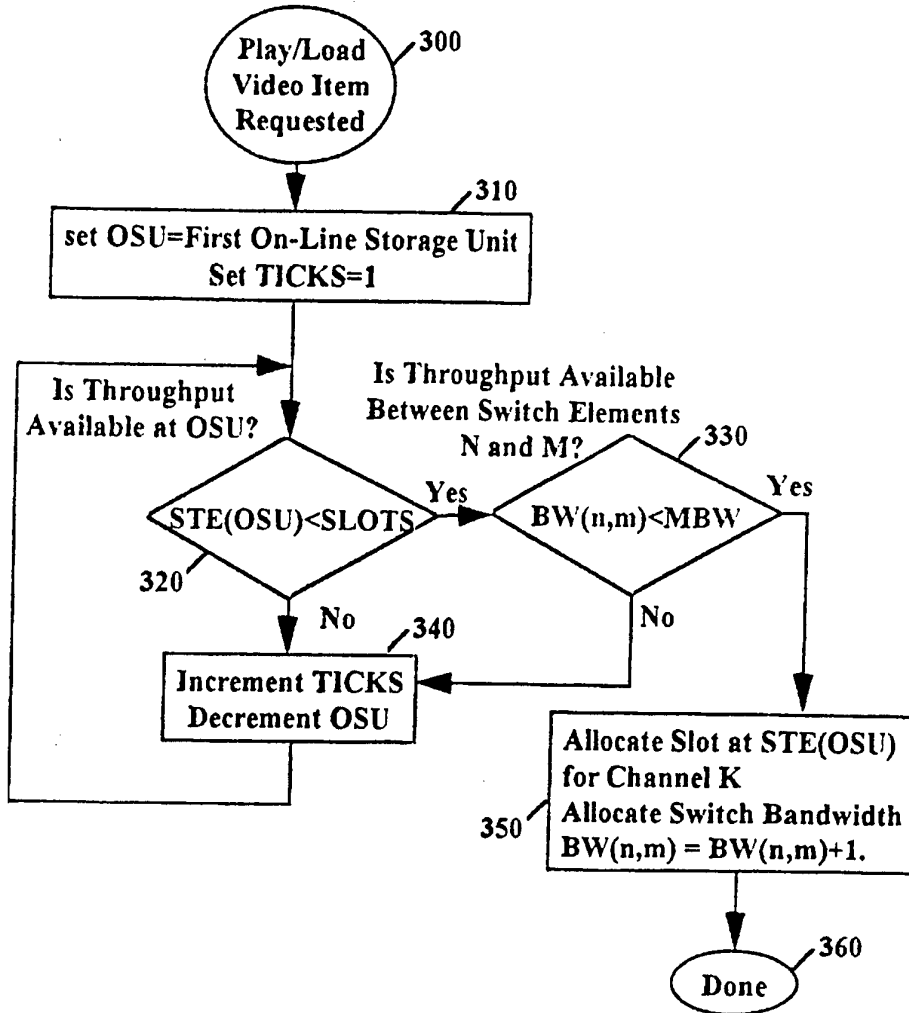


Fig. 8A/14



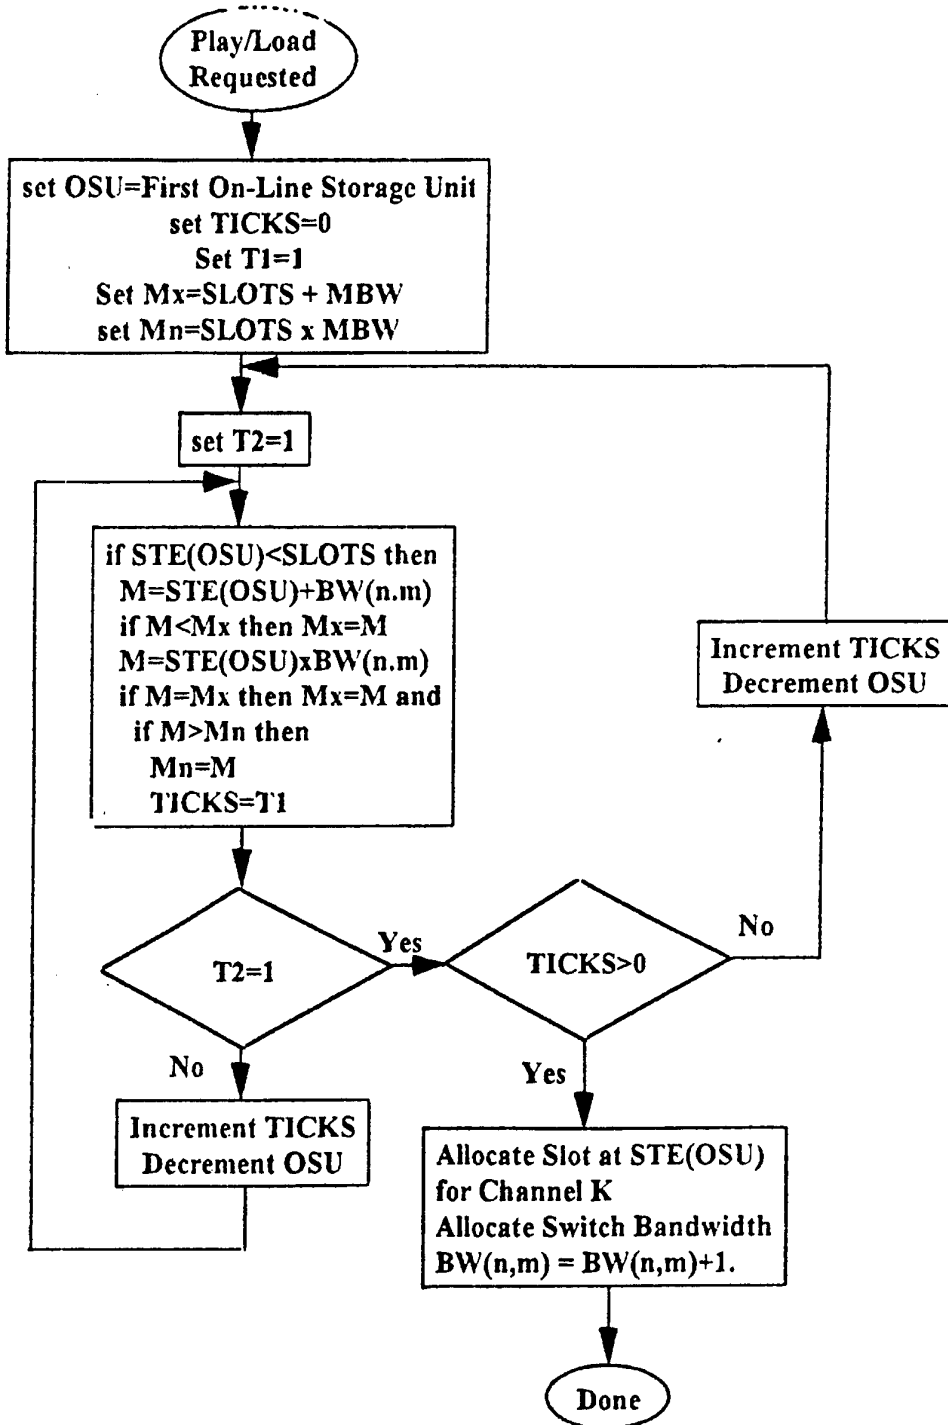


Fig. 8B/14

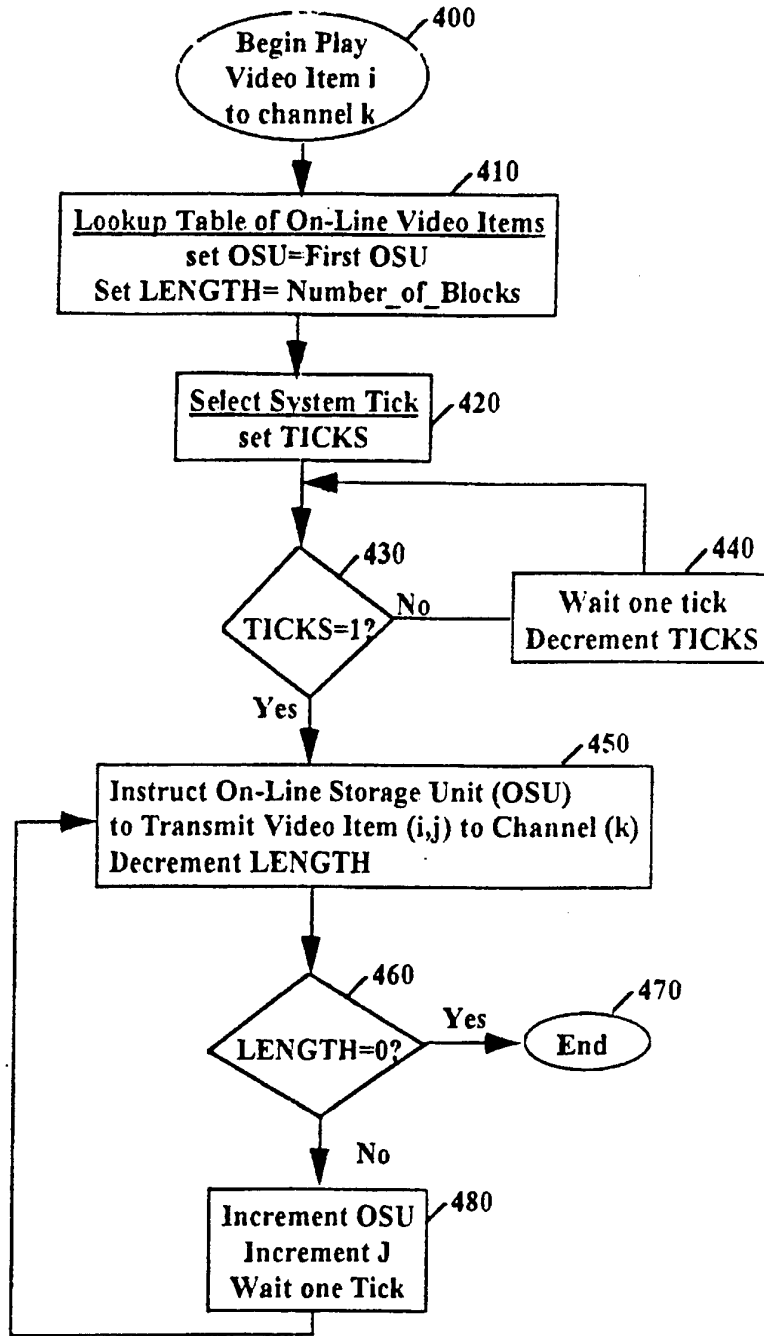


Fig. 9/14

12 / 16

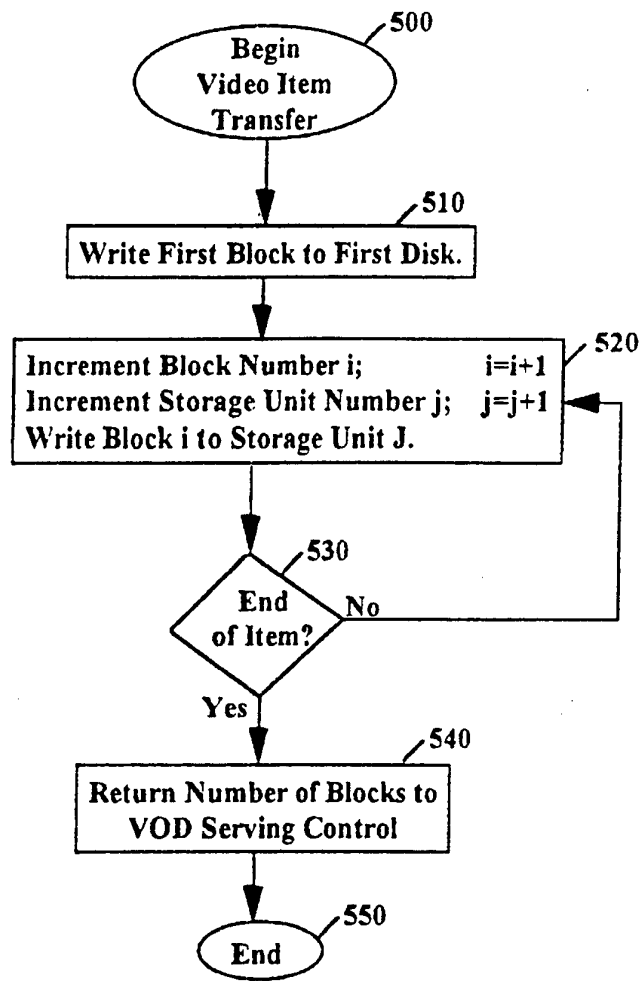


Fig. 10/14




Fig. 11/14




Fig. 12/14




Fig. 13/14

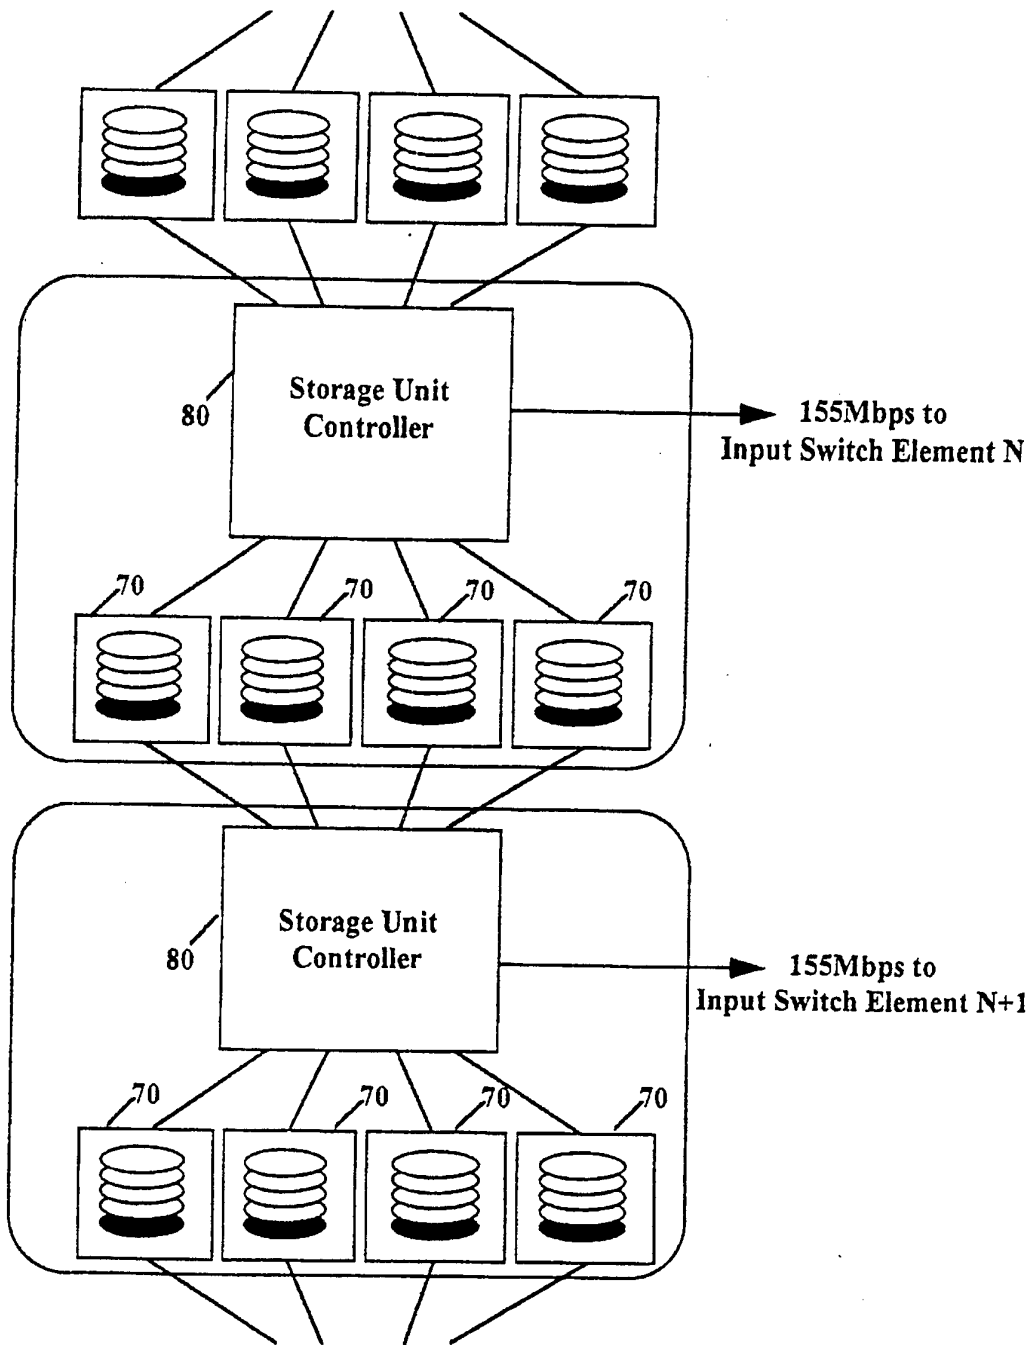


Fig. 14/14