



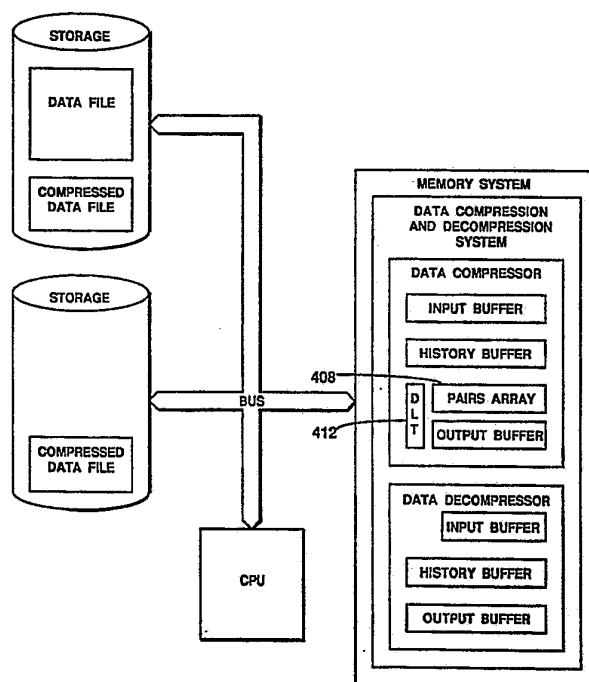
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification⁵ : G06F	A2	(11) International Publication Number: WO 93/17378 (43) International Publication Date: 2 September 1993 (02.09.93)
(21) International Application Number: PCT/US93/01553 (22) International Filing Date: 19 February 1993 (19.02.93) (30) Priority data: 07/839,958 21 February 1992 (21.02.92) US (71) Applicant: SALIENT SOFTWARE, INC. [US/US]; 124 University Avenue, Suite 300, Palo Alto, CA 94301 (US). (72) Inventor: CHAMBERS, Lloyd, Lamont, IV ; 328 Waverly, Apt. 2, Menlo Park, CA 94025 (US). (74) Agent: FLIESLER, Martin, C.; Fliesler, Dubb, Meyer & Lovejoy, Four Embarcadero Center - Suite 400, San Francisco, CA 94111-4156 (US).		(81) Designated States: AU, CA, JP, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>Without international search report and to be republished upon receipt of that report.</i>

(54) Title: METHOD AND APPARATUS FOR LOCATING LONGEST PRIOR TARGET STRING MATCHING CURRENT STRING IN BUFFER

(57) Abstract

An apparatus and a method finds a target string in a buffer, which target string matches a subsequent, current string to a greatest length. An array of array entries each uniquely identifies a value and a location of a respective data byte pair in the buffer. The array entries are sorted primarily upon their data pair values and secondarily upon their data pair locations. A table is indexable by each respective data byte pair in the buffer, to locate a first array entry for that data byte pair. To find a longest matching target string, the first two bytes of the current string are used as an index into the table, and the given table entry is used as an index into the sorted array, and the given array entry is used as an index to a first target string in the buffer. Each subsequent array entry having the same data byte pair value is used to locate a next target string. A longest matching string is determined from among these target strings. The location and length of the longest matching string are returned as a result.



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	FR	France	MR	Mauritania
AU	Australia	GA	Gabon	MW	Malawi
BB	Barbados	GB	United Kingdom	NL	Netherlands
BE	Belgium	GN	Guinea	NO	Norway
BF	Burkina Faso	GR	Greece	NZ	New Zealand
BG	Bulgaria	HU	Hungary	PL	Poland
BJ	Benin	IE	Ireland	PT	Portugal
BR	Brazil	IT	Italy	RO	Romania
CA	Canada	JP	Japan	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	SK	Slovak Republic
CI	Côte d'Ivoire	LJ	Liechtenstein	SN	Senegal
CM	Cameroon	LK	Sri Lanka	SU	Soviet Union
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	MC	Monaco	TG	Togo
DE	Germany	MG	Madagascar	UA	Ukraine
DK	Denmark	ML	Mali	US	United States of America
ES	Spain	MN	Mongolia	VN	Viet Nam
FI	Finland				

-1-

METHOD AND APPARATUS FOR LOCATING
LONGEST PRIOR TARGET STRING
MATCHING CURRENT STRING IN BUFFER

LIMITED COPYRIGHT WAIVER

A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

CROSS-REFERENCE TO RELATED APPLICATION

This application is a continuation-in-part of application 07/759,226 entitled "Fast Data Compressor with Direct Lookup Table Indexing into History Buffer", which was filed September 13, 1991 by Lloyd Lamont Chambers IV, which is commonly assigned with the present application, and which is incorporated herein by reference as if fully set forth herein.

BACKGROUND OF THE INVENTIONField of the Invention

The present invention relates generally to string matching techniques, and more specifically to an improved method and apparatus for finding, in a buffer, a target data string which matches a current data string to a length greater than or equal to all other target data strings which are located in the buffer prior to the current string. This improved method and apparatus may advantageously be used in a data compressor, for example.

Description of the Background

Data compression is useful in a variety of circumstances. Compressed data occupy less storage on a permanent storage medium. Compressed data take less time to transfer over a communication link. Compressed data in network packets have an improved effective data content, increasing the effective packet content and the network throughput. Pages containing compressed data in a virtual memory system can be swapped in and out more quickly.

Data compressors take advantage of repeated strings in a series of data. Second and subsequent instances of a given string are represented by a relatively short reference to a previous instance of the string. Prior techniques, such as Lempel-Ziv compressors, maintain "dictionaries" of previously-encountered strings of bytes. In order to determine whether a given string has previously been encountered, prior techniques generally employ some sort of hashing function to access the dictionary. Hashing consumes computational resources, and slows down the compression of data.

The new method and apparatus of data compression and decompression described in the cross-referenced

application "Fast Data Compressor with Direct Lookup Table Indexing into History Buffer" eliminates the need for a dictionary and also eliminates the need for hashing. A direct lookup table (DLT[]) is indexed or addressed by the first two bytes of a current data string in a history buffer (HB[]). A given entry in the DLT[] contains an offset or address in the HB[], at which a most recent prior string was found which began with the same two bytes with which the current string begins. When a previous matching string exists, the data compressor determines the length of the matching string, and outputs a vector reference back to the previous string. The DLT[] is then updated to reflect the newly-found current string. A subsequent string will later be referenced back to this new string, rather than to the previous string.

This method is extremely fast, offering a substantial improvement over its predecessors. More significantly, perhaps, is the fact that the disclosed decompressor can decompress the resulting compressed data at dramatically improved speeds, as is detailed in the cross-referenced application. However, the compressor of the cross-referenced application does not always obtain a theoretically maximum compression ratio.

For example, consider the string "ABCDABCABCD". According to the compressor of the cross-referenced application, the first "ABCD" bytes will be output as literal data, and the DLT[] entry addressed by "AB" (i.e. DLT["AB"]) will point to the first instance of "AB". Then, when the second instance of "ABC" is encountered, a vector will be outputted referring back to the first instance, with an indicated duplicate length of three (second "ABC" matches first "ABC"). The DLT["AB"] entry will be

updated to point to the second, most recent instance of "AB". Subsequently, the third instance of "ABC" will be found, and will be matched back against the second instance of "ABC", where the entry DLT["AB"] is then pointing. A vector will be outputted referring back to the second "ABC", indicating a length of three (third "ABC" matches second "ABC"). Finally, the last "D" will be output as a literal.

It is desirable that this last, literal "D" be avoided. What is needed, therefore, is an improved string matching apparatus and method which finds optimal previous target strings having maximum matching string lengths with a current string. For example, it is desirable that such an improved data compressor reference the second "ABCD" back to the first "ABCD", rather than merely referencing the third "ABC" back to the second "ABC" and then outputting an excess "D" literal.

20 SUMMARY OF THE INVENTION

Therefore, the present invention is an improved apparatus and method for locating a longest prior target string matching a current string in a data buffer.

25 The string matching apparatus and method maintains a sorted pairs array (SP[]) which references all instances of byte pairs, not only a most recently found pair. The apparatus and method maintain a direct lookup table (DLT[]) which, rather than indexing the HB[], addresses the SP[].

30 The first two bytes of the current string (i.e. the current data pair) are used to address the DLT[]. The DLT[current pair] entry points to the first SP[] entry that identifies an HB[] location at which the same two bytes are found. In other words,

35

the addressed DLT[] entry points to the first prior matching string of length at least two.

The apparatus and method determine matching string lengths for each SP[] entry pointing to target strings which match the current string, and determine which of those target strings matches the current string to the greatest length.

By this method and apparatus, it is guaranteed that a vector output to represent the current string will have near-optimal efficiency, or in other words, will avoid the most literal outputs.

In one embodiment, the apparatus and method are part of a data compression system. The data compression system compresses data from the HB[] to an output buffer, maintaining the DLT[] indexable by data pairs from the HB[]. The DLT[] entries point into the SP[], which indexes all data pairs in the HB[]. The SP[] permits rapid access to all prior strings which match the current string to a length of at least two bytes. Once the optimally matching target string is found, any encoding scheme whatsoever may be employed. Thus, the present invention is a "front end" for a data compressor. In one mode, the encoding scheme disclosed in the cross-referenced application is employed as the "back end" of the data compressor.

In an alternative embodiment, no DLT[] is required. The DLT[] is relatively large, taking 128 kilobytes for example. If this amount of memory is not available, the SP[] may be indexed by any conventional means, such as by binary searching for a first entry which points to a matching target string. This may be slower than directly accessing this first entry via the DLT[], but will nevertheless take less memory.

In various modes which depend upon the selected "back end" encoding portion of the data compressor, certain optional enhancements may be employed upon the SP[]. For example, SP[] entries which point so far back in the HB[] as to be non-referenceable from the current string (i.e. the chosen vector size cannot include a sufficiently large offset), those SP[] entries may be removed or destroyed from the SP[]. Alternatively, the appropriate DLT[] entry may simply be updated to access a next matching string after the non-accessible strings.

BRIEF DESCRIPTION OF THE DRAWINGS

Figs. 1-3 illustrate data compression performed in an archival system, a communication system, and a network system.

Fig. 4 illustrates uncompressed, compressed, and decompressed data files, and the input and output buffers, history buffers, direct lookup table, and sorted pairs array employed by the present invention.

Fig. 5 is a flowchart depicting high-level operation of the present invention.

Fig. 6 is a flowchart depicting operation of one mode of encoding a literal byte, as described in the cross-referenced application.

Fig. 7 illustrates one exemplary content of the history buffer, with the resulting content of the pairs array before and after sorting, and the resulting content of the direct lookup table.

Fig. 8 is a flowchart depicting operation of one mode of encoding a vector, as described in the cross-referenced application.

Fig. 9 is flowchart depicting operation of one mode of a decompressor, as described in the cross-referenced application.

Figs. 10 and 11 respectively show the improved data compressor of the present invention embodied as software in a floppy disk or the like and as firmware in a logic device such as a ROM or the like.

Fig. 12 illustrates one mode of a compressed data file, as described in the cross-referenced application.

Fig. 13 illustrates a compressed data file, specifically pointing out the block table entries, as described in the cross-referenced application.

Fig. 14 is a flowchart illustrating operation of the present invention to fill the pairs array from the history buffer.

Fig. 15 is a flowchart illustrating high-level operation of the present operation to sort the contents of the pairs array to form the sorted pairs array.

Fig. 16 is a flowchart depicting operation of the present invention to sort the pairs array by a hybrid sort methodology, including a preliminary incomplete insertion sort and a recursive quicksort.

Fig. 17 is a flowchart depicting operation of the present invention to perform an insertion sort upon the pairs array.

Fig. 18 is a flowchart depicting operation of the present invention to fill the direct lookup table from the sorted pairs array.

Fig. 19 is a flowchart depicting the operation of the present invention to find a best matching target string, using the direct lookup table and the sorted pairs array.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

A fundamental portion of the core of the present invention -- the string matching methodology -- is described below in section IX with reference to Fig. 19. However, sections I to VIII provide information which the reader may find helpful in understanding the context of the invention, and illustrate one means and method of carrying out the invention as a whole. These other portions of the invention are discussed below prior to a discussion of the string matching, to provide the reader a more chronologically accurate depiction of operation of the invention as a whole.

I. EXEMPLARY ENVIRONMENTS

Fig. 1 shows the present invention as utilized in a data compression system in a microcomputer, for example. Fig. 2 shows the present invention as utilized in data compression in a computer-to-computer modem transfer system, for another example. Fig. 3 shows the present invention as utilized in data compression in a network packetizing apparatus, for yet another example. In each of those embodiments, the data compressor includes a buffer means, an SP[] array means, a DLT[] means, and an output means. The invention may be practiced using any appropriate digital logic system, such as a computer, a "smart" modem, a network controller or node, a hard disk drive, a microprocessor, or the like.

II. DATA STRUCTURES EMPLOYED BY THE INVENTION

Fig. 4 illustrates the present invention as utilized in a data compression/decompression system 400, such as is generally shown in Figs. 1-3. A

data file 402 is read into the system, optionally via one or more input buffers 404, as is well-known. Data from the file are loaded into a history buffer HB[] 408 in blocks of a fixed, given size M. In one mode, the HB[] includes eight kilobytes of data.

5 A sorted pairs array (SP[]) 408 includes M-1 entries 408_0 to 408_{M-2} . Each entry corresponds to an instance of a data pair which exists in the HB[]. Duplicate instances of a given data pair which are found in the HB[] will have corresponding individual entries in the SP[].

10 Each SP[] entry includes two items (not shown in Fig. 4): a copy of an individual instance of a data pair, and a location of that instance in the HB[]. In one mode, the location is stored as an offset, in bytes, from the beginning of the HB[]. In the mode where the HB[] can include eight kilobytes of data (eight kilobytes = 2^{13} bytes), a two-byte offset is used (to hold thirteen offset bits - with three unused bits). Thus, in the one mode, each SP[] entry comprises four bytes.

20 The DLT[] 412 includes 256^N entries 412_0 to 412_{P-1} , where N is a number of consecutive data bytes from the HB[] which are used to index into the DLT[]. In one mode, N = 2, and the DLT[] includes $256^2 =$ sixty-four kiloentries. Thus, for each possible data pair value in the HB[], there is one uniquely corresponding DLT[] entry. Hence, the DLT[] will typically be much larger than either the HB[] or the SP[].

30 Each DLT[] entry includes one item: a location of an SP[] entry. In one mode, the location is stored as an offset, in bytes, from the beginning of the SP[]. Thus, with a maximum SP[] size of $(2^{13} - 1)$ entries * 4 bytes = 32,764 bytes, fifteen bits are required to address the SP[], and each DLT[] entry

35

comprises two bytes. Thus, in the described embodiment, the DLT[] comprises 256^2 entries * 2 bytes = 128 kilobytes.

5 III. EXEMPLARY CONTENTS OF THE DATA STRUCTURES

Fig. 7 illustrates one example of data which may (partially) fill the HB[]: the data "ABUGANDANANTSANGANTHEMS". The twenty-three data bytes reside in the HB[] from location HB[0] to location HB[22], inclusive.

Each pair of bytes at location HB[J] and HB[J+1] comprise a data pair. To fill the SP[], beginning at location J = 0, the pairs are loaded into the pairs array SP[]. Thus, location SP[0] is filled with an entry identifying the pair "AB" from location HB[0], and the location "0" of that pair in the HB[]. Then, SP[1] is filled with "BU" and "1", and so forth. Note that, while there are I = 23 data bytes, there are only I - 1 = 22 data pairs, because the final data byte does not have a subsequent item with which to form a pair.

Once the SP[] is filled, it is sorted. The sort is primarily upon the data byte content of the entries, and secondarily upon the address content of the entries. Thus, the entry for "AB" is found in the SP[] prior to the entry for "UG"; and, within the entries for "AN", the entries are found in numerical order of their SP[] address, "4" before "7" etc.

Once the SP[] is sorted, its contents are used to fill the DLT[]. First, the entire DLT[] is "zeroed out" to some predetermined value. In one mode, this value is -1, represented by its two's-complement hexadecimal equivalent, "FFFF". This number is used to signify empty DLT[] entries, because no entry in

the SP[] will be found before the first entry, SP[0].

After the DLT[] is initialized, selected ones of its entries are set. Each set DLT[] entry identifies a data pair *actually* found in the HB[]. In one mode, each set DLT[] entry more specifically identifies a data pair which is found least twice in the HB[]. Thus, set DLT[] entries identify duplicated strings, which are the potential areas where compression may occur. The choice as to which DLT[] entries should be filled may be made simply by consulting the sorted SP[], and looking for sequential SP[] entries having identical data pair items. In one mode, the DLT[] entry is set to point to the first occurrence of such a repeated SP[] entry. In Fig. 7, for example, DLT["AN"] is set to point to SP[1], which in turn points to HB[4], where the first occurrence of "AN" is found.

20 IV. EXEMPLARY MATCHING METHODOLOGY

Fig. 7 further illustrates a general methodology for using the DLT[] and the SP[] to find an optimally matching target string. In one sense, "optimally" is defined as meaning simply a previous string which matches a current string to a length greater than or equal to all other previous strings. Further refinements of the definition will be understood with reference to enhanced optimizations of the invention, to be described below.

30 Consider, for example, that the current string begins at HB[16], starting with the pair "AN". There are four previous matching target strings, all beginning with the pair "AN". As seen in the sorted SP[], these are found beginning at HB[4], HB[7], 35 HB[9], and HB[13].

According to the method of the cross-referenced application, the DLT[] (which in that case directly addresses the HB[]) would point to HB[13] when the current string is at HB[16]. The third characters of the strings beginning at HB[13] and HB[16] do not match, being "G" and "T", respectively. Thus, the former method would yield a matching string length of two.

By way of contrast, the present invention correctly determines that a matching string of length three is found at HB[9], as there is a string "ANT" at both HB[9] and HB[16].

When the string matching function begins at the new current string at HB[16], the present invention consults DLT["AN"] to find the first instance of a string beginning with "AN". Alternatively, the invention can operate without a DLT[], making reference from the HB[] into the SP[], via any conventional searching technique, rather than via a DLT[]

The invention consults HB[4] (i.e. HB[SP[DLT["AN"]]]) and determines that the matching string length there is two. The invention then checks the next SP[] entry, SP[2], to determine whether there is another "AN" string in the HB[]. SP[2] indicates that there is, so the invention consults the HB[] at the location HB[7] as indicated by SP[2], and finds another string matching to length two. The invention then finds a length three string at HB[9] through SP[3], and another length two string at HB[13] through SP[4]. By recording the location of the largest-yet string along the way, the invention finds the longest matching string, at HB[9].

The result of this portion of the invention is simply an identification of the optimal target

string. In the example, the result would be that "a length three string is found at HB[9]". This result can be used for whatever purposes it is needed. Typically, it will be used in a data compressor to maximize the compression ratio. However, other embodiments are contemplated, such as use in "grep" searching and the like.

V. DATA COMPRESSION EMBODIMENT

Fig. 5 illustrates in general terms the operation of a data compressor constructed using the data accessing and string matching techniques of the present invention. Operation of the data compressor begins at node 502, and proceeds with the outputting of a block header at block 504. Such a block header is described in the cross-referenced application, and in Fig. 12.

In Fig. 5, operation then continues at node 506, at which the invention fills the SP[] with unsorted data pairs as described above. This method is shown in more detail in Fig. 14, to be described below. In Fig. 5, operation continues at node 508 with the sorting of the SP[], performed by the operations shown in Fig. 15, to be described below. Then, the DLT[] is initialized at node 510 as shown in Fig. 18, to be described below. This completes the "setup" portion of the invention.

The data compression begins at block 512, which fetches a next pair from the HB[].

At conditional block 514, the data compressor checks whether the current data pair is one of a duplicated pair in the HB[]. This is done by consulting the DLT[]. As explained above, if there were two or more instances of the pair, the given DLT[] entry will not be "FFFF". In this case, operation passes to the string matching routine

shown generally at node 516 and in more detail in Fig. 19, to be described below. This completes the core string matching portion of the invention.

5 If the string matching routine determines that the matching string length is zero, as determined by conditional block 518, operation passes to node 520. This condition may occur, for example, if a matching string was found, but it was too far back for the back end compressor to reference by a predetermined
10 offset size. Also, if conditional block 514 determined that the pair was not repeated in the HB[], operation will pass immediately to node 520.

Node 520 outputs a literal to the compressed data file, by any preferred encoding. Fig. 6 depicts the
15 encoding method taught by the cross-referenced application.

If, at conditional block 518, the matching string length is not zero, operation passes to node 522, which outputs a vector according to any preferred
20 encoding. Fig. 8 depicts the encoding method taught by the cross-referenced application.

After the literal or vector is outputted, operation continues at conditional block 524, which checks whether compression of the HB[] is complete.
25 If not, operation returns to block 512, which fetches the next pair of data bytes in the HB[]. If the entire HB[] has been processed, the block header is updated in block 526, as appropriate with the preferred file format.

30 Then, operation continues at conditional block 528, which checks whether the block was compressed. As is known, incompressible data actually expands upon attempted compression, owing to flag bit penalties associated with identifying literals, and
35 the like. If the block was not effectively compressed, operation passes to block 530, which

simply copies the input to the output, or in other words, copies the HB[] to the output buffer. In either case, operation continues at block 532 which once again updates the block header as appropriate.

5 Operations end at node 534.

VI. FILLING THE SP[]

Fig. 14 shows one mode of filling the SP[] from the HB[], starting at node 1402 and continuing at block 1404. At block 1404, the filling starts at HB[0] and SP[0].

10

Then, at block 1406, an SP[] entry is set. The offset item of the SP[] entry is simply set to offset of the current HB[] location. The data pair item is a copy of the current data pair. A t decision block 1408, the method determines whether all HB[] pairs have been loaded into the SP[]. If not, operation returns to block 1406 to load a next pair and offset. Otherwise, operation is complete at node 1410.

15

20

VII. SORTING THE SP[]

Fig. 15 illustrates, at a high level, the general method of sorting the SP[] after the SP[] entries are made. Operation passes from node 1502 to block 1504. The preferred method of sorting the SP[] is with a recursive quicksort, for purposes of speed. However, it has been observed that a quicksort may, on occasion, use up an excessive amount of a host processor's available stack space, owing to excess recursion -- if the input data are too large, for example. Therefore, at block 1504, the method initializes a recursion level counter to one, signifying the first level of recursion.

25

30

Then, operation passes to node 1506, where the hybrid quicksort method of Fig. 16 performs a rough

35

sort on the SP[]. Then, at node 1508, the insertion method of Fig. 17 performs a final, more refined insertion sort on the SP[]. Sorting ends at node 1510.

5 Fig. 16 shows the hybrid quicksort utilized to rough sort the SP[]. As is known in the art, quicksort operates recursively upon increasingly smaller "chunks" of the entire sort data. Operation begins at node 1602 and passes to conditional block
10 1604. At conditional block 1604, the quicksort checks whether two parameter conditions are met. First, the quicksort checks whether the chunk size is within an acceptable range. This check relates to the hybrid nature of the quicksort. It has been
15 observed that for sorting a "chunk" of a relatively small number of data, quicksorting is slower than certain other methods of sorting. At block 1604, if the chunk size is below a predetermined number, the present instantiation (recursion) of the Fig. 16
20 method terminates at ending node 1632. In one mode, the minimum chunk size upon which quicksorting is performed is ten SP[] entries.

The second check which is performed at conditional block 1604 is a check of the current level of
25 recursion. This prevents system crashes caused by stack overflow, in the event of "runaway" recursion such as may result from an execution error. If recursion has descended below a predetermined number of instantiations, execution of the current
30 instantiation terminates at ending node 1632. In one mode, the maximum recursion count is twenty.

If both tests at conditional block 1604 are met, operation continues at block 1606, which swaps the values at the beginning and end of the current
35 chunk.

Then, operation passes to node 1608 which represents the insertion sort method of Fig. 17. However, node 1608 does not perform a full insertion sort over the given chunk. Rather, it performs a limited insertion sort, which has been observed to speed up the overall sorting operation. This improves the quicksort's chances of selecting a "good" pivot.

After the limited insertion sort is performed upon the chunk, operation continues at block 1610, which selects the pivot. In blocks 1610 and 1612, the quicksort swaps the first and last elements in the chunk.

Then, operation passes to conditional block 1614, which checks whether the pivot element is greater than the left element of the chunk. If so, operation passes to block 1616, which advances the left pointer one element to the right, i.e. toward the pivot. This only advances the pivot, and does not move data. Data are moved at block 1612 only. Then, operation returns to conditional block 1614 for further testing.

Eventually, the left pointer will identify an element which is not less than the pivot, and operation will pass to conditional block 1618. Conditional block 1618 and block 1620 operate similarly to conditional block 1614 and block 1616, moving the right pointer leftward toward the pivot. When the right pointer identifies a data element which is not greater than the pivot element, operation passes to conditional block 1622.

Conditional block 1622 checks whether the right element is equal to the pivot element. If so, the right pointer is advanced one position past the pivot to the left, at block 1624. Regardless, operation then passes to conditional block 1626.

Conditional block 1626 checks whether the left and right pointers have passed each other. If not, operation passes back to block 1612, which swaps the new left and right elements, and checking begins again. If, at conditional block 1626, the left and right pointers have passed each other, the quicksort recursively calls itself at node 1628. Upon recursive entry, the new chunk size and recursion count will, of course, be checked as before.

When operation eventually returns from the new recursive instantiation of the quicksort, operation passes to block 1630, which moves the starting point to the right, to the point where the left pointer currently points. From there, operation passes back (non-recursively) to conditional block 1604, where operations continue as before. Thus, when viewed at a high level of abstraction, rough sorting of the HB[] very generally proceeds from HB[0] toward HB[M-1].

It will be understood that, depending upon the data and upon the predetermined recursion count limit, each given instantiation of the quicksort method may perhaps not fully sort the data in its particular chunk. Thus, after node 1506 in Fig. 15, the quicksort will be fully complete, but the SP[] may not be entirely sorted. That is why the insertion sort is invoked at node 1508. This is a full insertion sort over the entire SP[]. With the data in a *nearly sorted* configuration, the insertion sort will finish the sorting very rapidly, and the SP[] will be fully sorted, in a "pair-major, offset-minor" manner.

Fig. 17 illustrates the methodology of the insertion sort. Operation begins at node 1702, and passes to block 1704. At block 1704, the insertion sort finds the smallest element in the SP[]. With

the SP[] nearly sorted by the quicksort, the smallest SP[] item will be found very near the beginning of the SP[]. Thus, the insertion sort will not have to spend much time looking for the smallest element. Once it is found, the smallest element is swapped with the first SP[] entry, at SP[0].

Then, operation passes to conditional block 1706, which checks whether the insertion sort is complete. If so, the insertion sort terminates at node 1716. Otherwise, operation continues at block 1708. At block 1708, the insertion sort advances to a next starting point. This takes advantage of the fact that the smallest element has already been placed at the beginning of the SP[]. After that is done, there is no need to look at that location for the second, third, etc. smallest items. Thus, the insertion sort will only start looking for a next smallest item after all previously-found smallest items. This, together with the nearly sorted contents of the SP[], ensure that the insertion sort will never have to search far for a next smallest element.

Also at block 1708, the insertion sort starts at this new starting point (which is where the next smallest item will be placed). Then, operation passes to conditional block 1710, where the insertion sort checks whether the SP[] entry at the new starting point is smaller than that at the next subsequent location. If so, operation passes to block 1712, which moves the SP[] entry at the new starting point one position toward the end of the SP[]. Then, operation returns to the conditional block 1710. If, at the conditional block 1710, the subsequent SP[] entry is not less than that at the new starting point, operation passes to block 1714,

where the insertion sort moves the newly-found smaller item into SP[new start]. Then, operation returns to conditional block 1706, and the insertion sort continues.

5 It will be understood that any useful sorting method may be employed. The hybrid quicksort/insertion sort depicted has been found to produce the most satisfactory sorting speeds, owing to the particular sizes of the data structures
10 employed.

VIII. INITIALIZATION AND LOADING OF THE DLT[]

Fig. 18 shows the method whereby the DLT[] is initialized and loaded, once the SP[] is fully
15 sorted. Operation begins at node 1802, and passes to block 1804. At block 1804, all DLT[] entries are initialized to the reserved value, such as FFFF_{hex}. Then, operation continues at block 1806, where the method begins at SP[0].

20 Then, operation passes to conditional block 1808, which checks whether the entire SP[] has been processed. If so, operation terminates at node 1818. Otherwise, operation passes to block 1810, which causes the method to check a next pair of
25 SP[] entries the next time around the loop.

Then, operation continues at conditional block 1812, which checks whether the data pair in the next SP[] entry is the same as that in the current SP[] entry. In other words, conditional block 1812
30 checks whether two consecutive SP[] entries have the same data pair value. If not, this identifies a non-repeated string in the HB[], and operation returns to conditional block 1808 to check a next set of pairs.

35 Otherwise, operation passes to conditional block 1814, which checks whether the DLT[] entry indexed

by the repeated data pair value is still zeroed to the reserved value. In other words, conditional block 1814 checks whether the next SP[] entry is exactly the second SP[] entry to contain the given data pair value. If not, the DLT[] entry has already been set to point to the first SP[] entry containing the given data pair value. In other words, the DLT[] entry already accesses the first instance of the repeated string. In this case, operation passes back to conditional block 1808 to check a next pair of SP[] entries on the next loop.

Otherwise, operation passes to block 1816, which sets the DLT[] entry indexed by the current data pair to point to the first SP[] entry having the given data pair value. Thus, points the DLT[] entry accesses the first instance of the repeated string. In block 1816, the method advances to check a next pair of DLT entries, and operation returns to conditional block 1808.

IX. FINDING THE OPTIMALLY MATCHING STRING

Fig. 19 illustrates the methodology by which the present invention locates, in a buffer, a prior target string which matches a current string to a greatest length. Operation passes from a starting node 1902 to block 1904, where the method sets a matched length counter to zero, and finds the first pair in the HB[] which pair matches the current pair in the HB[]. This first pair is found as HB[SP[DLT[HB[current location]]]. In other words, the current pair is used as an index into the DLT[], and that DLT[] entry points to the first SP[] entry for a matching string in the HB[].

Next, operation continues at conditional block 1906, which checks whether the potential target string begins before the start of the HB[]. This

would indicate an error condition. If so, operation terminates at ending node 1938. Otherwise, operation passes to conditional block 1908. At conditional block 1908, the method checks whether the current string is within a predetermined maximum history length back from the current string.

The predetermined history length need not be related to the size of the HB[]. Rather, the predetermined history length defines a maximum remembered "history" which the method will maintain. For example, the method may use an eight kilobyte HB[], but retain only a most recent four kilobytes of history. This allows, in effect, a "sliding" history, such as is well-known in the art. Early in the HB[], a vector may refer back as far as the beginning of the HB[] to reference the current string to a prior target string; later on in the HB[], a vector may only be allowed to reach a certain distance back from the current pair. The particular magnitude chosen for the predetermined history length will depend upon a variety of considerations, such as the chosen encoding method for vectors (so that the invention does not find target strings which are unreferenceable by a data compressor utilizing the invention, for example).

If the distance (offset) from the current pair back to the beginning of the HB[] is within the allowed history length, at block 1910, the method indicates that any target string starting no earlier than HB[0] is referenceable. Otherwise, at block 1912, the method indicates the earliest HB[] location which is referenceable.

Then, operation passes to block 1914. There, the method prevents referencing of any target string which is subsequent to the current pair in the HB[].

Also at block 1914, the method sets a best matching pair to point to the current pair, sets a best matching value to zero, and makes a copy of the current pair (for non-dereferenced use).

5 Then, operation continues at conditional block 1916, where the method checks the potential matching target string, to see whether it begins with the same pair as the current string. It also ensures that the offset from the current pair back to the
10 potential target string is small enough for the chosen vector referencing method. If not, operation ends at node 1938.

 Otherwise, operation passes to conditional block 1918, which checks whether a next (i.e. third or
15 subsequent) byte in the potential target string matches a byte at the corresponding position in the current string. If so, operation passes to block 1920, which increments pointers which access the current and target strings. Then, operation passes
20 back to conditional block 1918, to check a next byte in the target string against a next byte in the current string.

 If, at block 1918, corresponding bytes in the target and current strings do not match (i.e.
25 processing has advanced beyond the end of matching strings), operation passes to block 1922. At block 1922, the method saves the value of the matched length. Then, operation continues at conditional block 1924, where the method determines whether the
30 just-matched length is greater than or equal to any previously-matched length.

 If not, operation passes to block 1926, where the method selects the next potential target string, and operation passes back to conditional block 1916 to
35 begin the checking of this new target against the current string.

Otherwise, operation passes to conditional block 1928, where the method checks whether the strings overlap. This feature is optional within the bounds of this invention, to allow the use of the same output format as is disclosed in the cross-referenced application. Overlap occurs when the target string extends into the current string. For example, the string "ABCDABCDABX" contains a target string "ABCDAB" which begins at the first position, and which matches to length six the string "ABCDAB" which begins at the fifth position. However, there is an overlap in that the second instance of "AB" is included in both the target and the current string. It has been determined that some compression encoding methodologies operate more efficiently if overlapping is disallowed, such as is shown in Fig. 19.

If overlap is detected, operation passes to block 1930, which shortens the determined length of the target string such that no overlap exists (i.e. the target string will end at the byte immediately before the current pair), and operation returns to conditional block 1924 to check whether the shortened match length is still the longest so far.

If there is no overlap, operation passes from conditional block 1928 to block 1932, where the method indicates that the just-matched target is the best so far, saving indications of where the target begins, and the length to which it matches the current string.

Then, operation passes to conditional block 1934, where the method checks whether the matching target string length is equal to a predetermined maximum matching string length. This value is selected according to the chosen vector encoding method, to prevent identification of a target string longer

than can be encoded. According to the cross-referenced application, the chosen maximum length is twenty bytes. However, in other modes, it could be nearly as large as the history buffer itself. It will be understood that conditional block 1934 could test for "greater than or equal to". However, the particular coding of the blocks 1918 and 1920 allows the use of a simple "equal to" test. In the preferred mode, the code representing blocks 1918 and 1920 has been "loop unrolled", to remove conditional branching instructions which have a known effect of reducing pipeline efficiency. The loop has been unrolled a number of times such that it is guaranteed that the longest string which can be matched is of exactly the maximum allowable string length.

At conditional block 1934, if the matching target string length is not greater than the maximum size, operation returns to block 1926, which selects a next target to match. Otherwise, operation passes to block 1936, where the method saves an indication of the starting position and matching length of the target string. These are the "results" of the method of Fig. 19; finding them is the goal of the core portion of the present invention.

Finally, operation terminates at node 1938.

X. PREFERRED ENVIRONMENT

The preferred environment, in which the present invention is practiced, is an Apple Macintosh computer having a microprocessor selected from the Motorola 680x0 microprocessor family. The invention may, however, be practiced using a wide variety of other systems. The source code, disclosed in the Appendix hereto, is written in a combination of C and assembly language, and has been written

specifically for use with the Symantec Think C 5.0 compiler/assembler. The source code discloses the best mode contemplated for carrying out the invention, and further includes certain optimizations which have been made according to the preferred host processor and compiler/assembler. Those skilled in the art will readily appreciate that these optimizations need not necessarily be used when practicing this invention.

It will be understood that the heart of this invention lies in the methodology of using one or more bytes at a current position in the input data to directly or indirectly access a sorted array of pointers, to locate one or more previously-found instances of a potentially matching target string. While it is contemplated that one desirable method of doing this is to use two bytes at the current location to index a DLT[], and to then use the particular DLT[] entry to index an SP[], other similar methods are within the scope of this invention. For example, it may be desirable, depending upon a particular host computer for example, to use a first and a third byte at the current location, rather than a first and a second. It may further be desirable, depending upon the particular host computer, to use more or less than two bytes as an index. For example, on a host computer having a relatively large memory, it may be more desirable to use three or four bytes as an index.

Other adaptations will be understood by those skilled in the art. For example, the method may be fine-tuned according to the chosen "back-end", be it a data compressor or whatever. Specifically, the method may be adapted to find not just a longest matching string, but a latest one of two or more

such equally longest matching strings. This may enable the vector encoding method to utilize a more efficient vector, such as one having a reduced number of offset bits.

5 It will be further understood that, while the discussed embodiment uses indexing, corresponding embodiments may readily be constructed which use pointer addressing or the like.

10 It will further be understood that the present invention, which is a string location method, may be advantageously employed in areas other than data compression.

15 While the invention has been described with reference to the embodiment of a data compression system, its use in any suitable embodiment is contemplated. Various changes in the form and details discussed herein are possible within the scope of this disclosure, including but not limited to those discussed above. However, the invention is
20 to be afforded protection according to the scope of the following Claims.

CLAIMS

I Claim:

1. A method performed by a digital logic system upon data in a buffer, wherein the data include a plurality of strings, each string beginning with a value at a location in the buffer, each value including one or more bytes, the digital logic system including an array (SP[]) of array entries, each array entry uniquely identifying the value and the location at which a respective one of the strings begins, the method for finding a longest matching target string in the buffer with respect to a current string beginning with a current value at a current location in the buffer, the method comprising the steps of:

accessing one or more array entries in the array responsive to the current value, each of the one or more array entries identifying a value equal to the current value and identifying a location in the buffer prior to the current location, the one or more array entries identifying target strings;

determining, for each of the target strings, a length to which the target string matches the current string;

identifying a given one of the target strings which matches the current string to a greatest length determined by the step of determining a length; and

providing an identification of the given one of the target strings.

2. The method of Claim 1, further comprising the step of:

constructing the array by the steps of,
scanning the buffer to identify each value in
the buffer,
copying, into a next array entry, each identified value and its location, and
sorting the array entries primarily by value and secondarily by location.

3. The method of Claim 2, wherein the buffer includes M bytes, each value includes N consecutive bytes of the M bytes, and wherein:

the steps of scanning and copying identifies M-(N-1) value/location combinations, such that the array includes M-(N-1) array entries.

4. The method of Claim 3, wherein:

N = 2, such that the current value is a current pair of bytes, and the step of accessing one or more array entries uses the current pair of bytes from the buffer to access the array.

5. The method of Claim 2, wherein the digital logic system further includes a table (DLT[]) of table entries, each table entry identifying a given array entry, and wherein the step of accessing one or more array entries comprises the steps of:

using the current value as an index into the table to identify a given table entry; and

using the given table entry as an index into the array to access a given array entry.

6. The method of Claim 5, further comprising the step of:

constructing the table by the steps of,
scanning the array to identify values
5 appearing in the buffer,

for each given value not appearing in the
buffer, invalidating a given table entry which is
accessible by the given value, and

10 for each respective value appearing in the
buffer, validating a given table entry which is
accessible by the respective value.

7. The method of Claim 6, wherein:

the step of invalidating includes setting the
given table entry to index outside the array; and

5 the step of validating includes setting the
respective table entry to index a respective array
entry which identifies a value equal to the
respective value appearing in the buffer.

8. The method of Claim 7, wherein:

5 the step of setting sets the respective table
entry to point to a respective array entry which
further identifies an earliest location of all array
entries identifying the respective value.

9. The method of Claim 8, further comprising the
steps of:

5 updating a given table entry to access a
subsequent array entry, to prevent access to target
strings which are farther back from the current
string than a predetermined maximum offset.

10. The method of Claim 6, wherein:
the step of invalidating invalidates a given table entry for each given value not appearing in the buffer at least twice; and
5 the step of validating validates a respective table entry for each respective value appearing in the buffer two or more times.

11. The method of Claim 10, wherein the steps of validating and invalidating include the step of:
comparing values identified by two or more consecutive array entries, to determine whether a
5 value is repeated in the buffer.

12. The method of Claim 6, wherein the buffer includes M bytes, each value includes N consecutive bytes of the M bytes, and wherein:
the step of constructing the table constructs the
5 table to include 256^N table entries, uniquely accessible by respective N-byte values from the buffer.

13. The method of Claim 12, wherein:
N = 2.

14. The method of Claim 1, wherein:
the steps of accessing one or more array entries and determining a length terminate in response to the step of determining finding a target string
5 having a length equal to a predetermined maximum length.

15. The method of Claim 1, further comprising the step of:
compressing the data from the buffer.

16. The method of Claim 15, wherein:

in the step of providing an identification, the identification may further identify no string, to indicate that no target string matches the current string, indicating that the current string is a first occurrence of a value; and

the method further includes the steps of,

responsive to an identification of no string, outputting the value of the current string as a literal, and

responsive to an identification of a given target string, encoding the given target string as a vector.

17. A method of finding a longest target string which is prior to a current string at a current location in a buffer of data, the method comprising the steps of:

entering, into each of one or more respective array entries of an array (SP[]), an identification of a unique location of respective data in the buffer and a value of the respective data;

entering, into each of one or more respective table entries of a table (DLT[]), a location of an array entry;

responsive to a value of data at a current location in the buffer, accessing a given one table entry;

responsive to a location identified by the given one table entry, accessing one or more given array entries;

responsive to respective locations identified by the one or more given array entries, comparing target strings at the respective locations in the buffer to the current string in the buffer, to identify a given one target string which has a

longest matching string length of the target strings; and

25 providing an identification of the given one target string as the longest target string.

18. The method of Claim 17, further comprising the steps of:

sorting the array according to the identification of the respective data in the array entries.

19. The method of Claim 18, further comprising the steps of:

5 sorting the array according to the identification of the unique locations of the respective data in the array entries.

20. The method of Claim 17, further comprising the steps of:

5 invalidating a table entry corresponding to a data value not appearing in the buffer; and validating a table entry corresponding to a data value appearing in the buffer.

21. The method of Claim 20, further comprising the steps of:

5 invalidating a table entry corresponding to a data value not appearing in the buffer at least twice; and validating a table entry corresponding to a data value appearing in the buffer at least twice.

22. The method of Claim 21, further comprising the steps of:

5 sorting the array primarily according to the data values and secondarily according to the unique locations.

23. The method of Claim 22, wherein:
the current data comprise a pair of bytes
(CURPAIR) in the buffer; and
the accessing steps comprise access by
5 SP[DLT[CURPAIR]].

24. The method of Claim 23, further comprising
the steps of:
updating the table entry DLT[CURPAIR] to index an
array entry at a location after that entered by the
5 entering step, to prevent access to a target string
which is farther back from the current location than
a predetermined maximum offset.

25. The method of Claim 17, further comprising
the steps of:
responsive to the step of providing an indication
of the longest target string, outputting a vector to
5 compress the data in the buffer.

26. A method performed by a digital computing
device upon a buffer HB[], the buffer containing M
bytes of data at buffer locations from HB[0] to
HB[M-1], the method comprising the steps of:
5 A) creating data structures for accessing the
HB[], including the steps of,
1) responsive to the M bytes of data,
creating an array SP[] having M-1 SP[] entries at
locations from SP[0] to SP[M-2], including the steps
10 of, for each entry SP[J],
a) copying into the entry SP[J]
contents of a data pair from buffer locations HB[J]
and HB[J+1], and
b) copying into the entry SP[J] a value
15 of J to identify a buffer location of the data pair,

35

2) sorting the SP[] entries primarily upon their respective data pairs and secondarily upon their respective buffer locations,

3) creating a direct lookup table DLT[] having 256² entries at locations from DLT[0000_{hex}] to DLT[FFFF_{hex}], including the steps of,

20 a) copying into each DLT[] entry a value outside the range 0 to N-2, and

b) responsive to the SP[], for each

25 data pair value XY which was copied into two or more SP[] entries, copying into the entry DLT[XY] a value of a location Z of an entry SP[Z] at which a first instance of the data pair value XY was copied;

B) finding a target string at a location in the

30 HB[] prior to a current location CUR at which a current string begins, including the steps of,

1) obtaining a current data pair CP from locations HB[CUR] and HB[CUR+1],

2) obtaining an SP[] location S from entry

35 DLT[CP],

3) for successive each SP[] location SP[S] containing a data pair equal to CP,

a) obtaining an HB[] location T from entry SP[S],

40 b) determining a given length to which a target string at location HB[T] matches the current string,

c) if the given length is greater than any previously determined length, identifying a

45 longest-yet target string by the given length and the location HB[T], and

d) proceeding to a next location SP[S];

and

C) providing an indication of the identified

50 longest-yet target string as an output of the method.

27. The method of Claim 26, further comprising the steps of:

5 determining that an offset from the current location to a first matching target string is greater than a predetermined maximum offset; and
updating the entry DLT[CP] such that DLT[CP] does not access an SP[] entry corresponding to the first matching target string, including,
10 if there is a second matching target string between the first matching target string and the current string, updating DLT[CP] to access the second matching target string's corresponding SP[] entry, otherwise updating DLT[CP] to a value outside the range 0 to N-2.

28. The method of Claim 26, further comprising the steps of:

5 responsive to the output, encoding the longest matching target string as a vector, to compress the data from the HB[] to an output buffer.

29. An apparatus comprising:

a buffer (HB[]) of data;
a table (DLT[]) of table entries, including means for addressing respective table entries in response to respective data from the buffer;
5 an array (SP[]) of array entries, including means for addressing respective array entries in response to respective table entries from the table;
means, coupled to the table and the array, for
10 accessing target strings in the buffer;
means, coupled to the means for accessing, for comparing each target string against a current string in the buffer, to determine lengths to which the respective target strings match the current
15 string; and

means, coupled to the means for comparing, for determining which target string is a longest matching target string, and for providing an identification of the longest matching target string.

20

30. A data compressor for compressing data from an input buffer (HB[]) to an output buffer, the data compressor comprising:

means for constructing an array (SP[]) of array entries each uniquely identifying a value and a location in the input buffer of a plurality of bytes, including means for sorting the array entries;

5

means for using a current plurality of bytes at a current location in the input buffer to access one or more given array entries which identify a value equal to a current value at a current location in the input buffer;

10

means for comparing each of one or more target strings at locations identified by the one or more given array entries against a current string at the current location in the input buffer, to determine and indicate which if any of the target strings matches the current string to a greatest length;

15

literal means, responsive to an indication that no target string matches the current string, for outputting literal data from the current location in the input buffer; and

20

vector means, responsive to an indication that a given target string matches the current string to a greatest length, for outputting a vector identifying the given target string.

25

31. The data compressor of Claim 30, wherein the means for sorting comprises:

means for sorting primarily on data value and secondarily on data location.

32. The data compressor of Claim 30, wherein the means for sorting further comprises:

means for performing a quicksort.

33. The data compressor of Claim 32, wherein: the means for performing a quicksort performs a hybrid quicksort/insertion sort.

34. The data compressor of Claim 30, wherein the means for comparing comprises:

5 means for halting comparison upon encountering a target string which matches the current string to a predetermined maximum length encodable by the vector means.

35. The data compressor of Claim 30, wherein: the data compressor further comprises means for reading a data file into the input buffer in blocks of a predetermined size; and

5 the means for constructing an array, means for using a current plurality of bytes, means for comparing, literal means, and vector means are responsive to the means for reading to reset operations at each block, to output
10 independently-compressed data blocks.

36. Computer software comprising:

instructions for maintaining a buffer (HB[]) of data;

5 instructions for maintaining a table (DLT[]) of table entries, including instructions for addressing

respective table entries in response to respective data from the buffer;

instructions for maintaining an array (SP[]) of array entries, including instructions for addressing
10 respective array entries in response to respective table entries from the table;

instructions for accessing target strings in the buffer in response to respective array entries from the array;

15 instructions for comparing each target string against a current string in the buffer, to determine lengths to which the respective target strings match the current string; and

instructions for determining which target string
20 is a longest matching target string, and for providing an identification of the longest matching target string.

37. Computer software for finding a longest target string which is prior to a current string at a current location in a buffer of data, the computer software comprising:

5 a routine for entering, into each of one or more respective array entries of an array (SP[]), an identification of a unique location of respective data in the buffer and a value of the respective data;

10 a routine for entering, into each of one or more respective table entries of a table (DLT[]), a location of an array entry;

a routine, responsive to a value of data at a current location in the buffer, for accessing a
15 given one table entry;

a routine, responsive to a location identified by the given one table entry, for accessing one or more given array entries;

20 a routine, responsive to respective locations
identified by the one or more given array entries,
for comparing target strings at the respective
locations in the buffer to the current string in the
buffer, to identify a given one target string which
has a longest matching string length of the target
25 strings; and

a routine for providing an identification of the
given one target string as the longest target
string.

38. Data compression software for compressing
data from an input buffer (HB[]) to an output
buffer, the data compression software comprising:

5 a floppy disk having recorded thereon instructions
for performing data compression, the instructions
including,

instructions for constructing an array (SP[])
of array entries each uniquely identifying a value
and a location in the input buffer of a plurality of
10 bytes, including instructions for sorting the array
entries,

instructions for using a current plurality of
bytes at a current location in the input buffer to
access one or more given array entries which
15 identify a value equal to a current value at a
current location in the input buffer,

instructions for comparing each of one or
more target strings at locations identified by the
one or more given array entries against a current
20 string at the current location in the input buffer,
to determine and indicate which if any of the target
strings matches the current string to a greatest
length,

instructions, responsive to an indication
25 that no target string matches the current string,

for outputting literal data from the current location in the input buffer, and

instructions, responsive to an indication that a given target string matches the current string to a greatest length, for outputting a vector identifying the given target string.

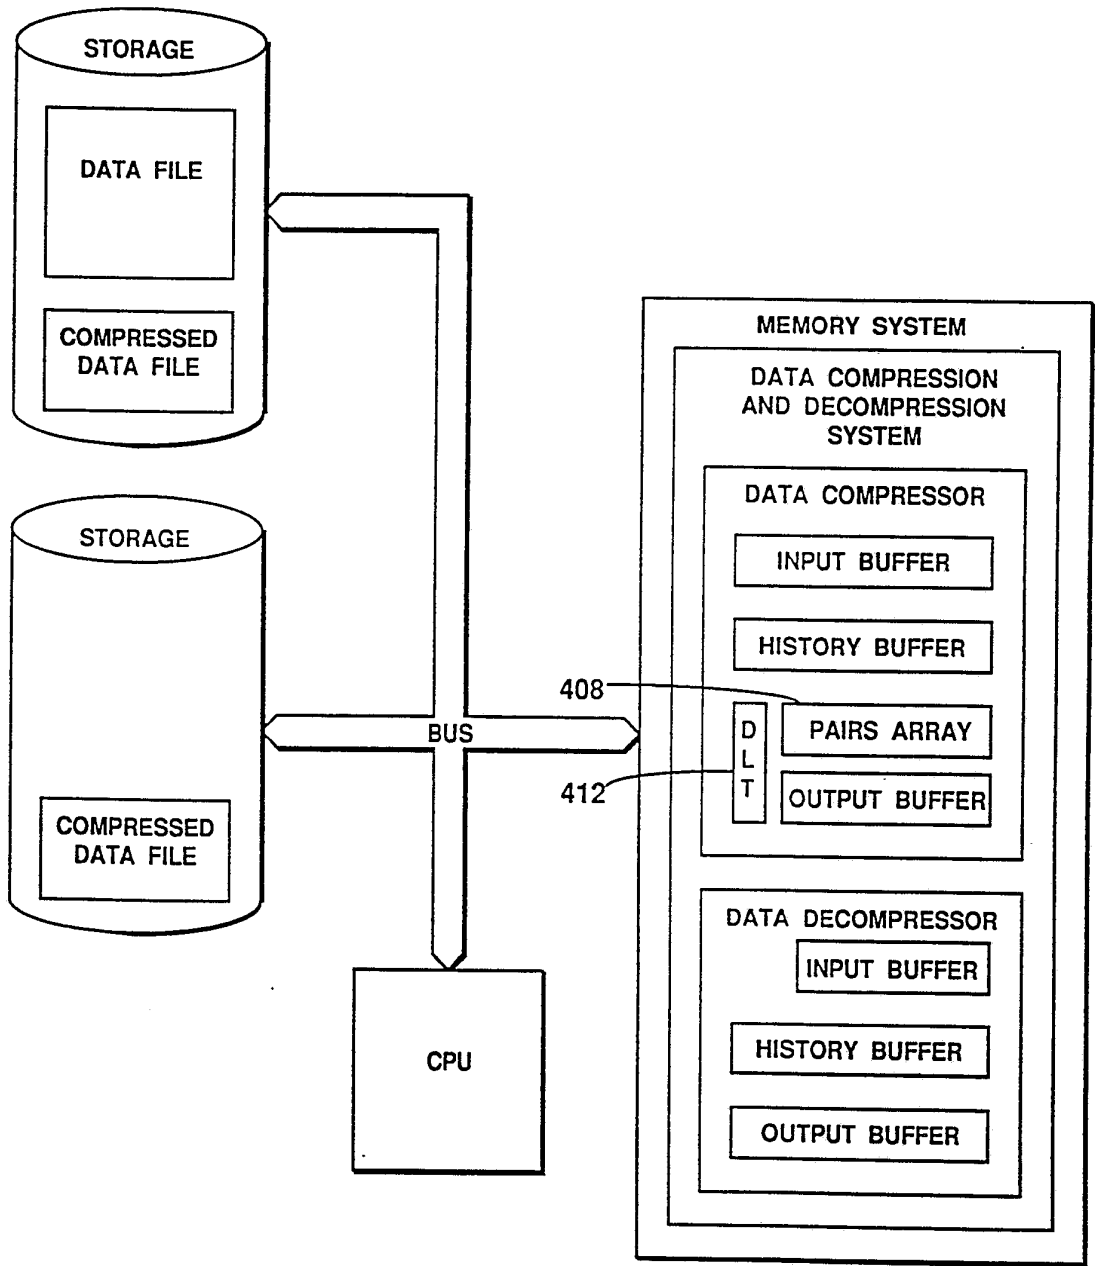


Fig. 1




Fig. 2




Fig. 3


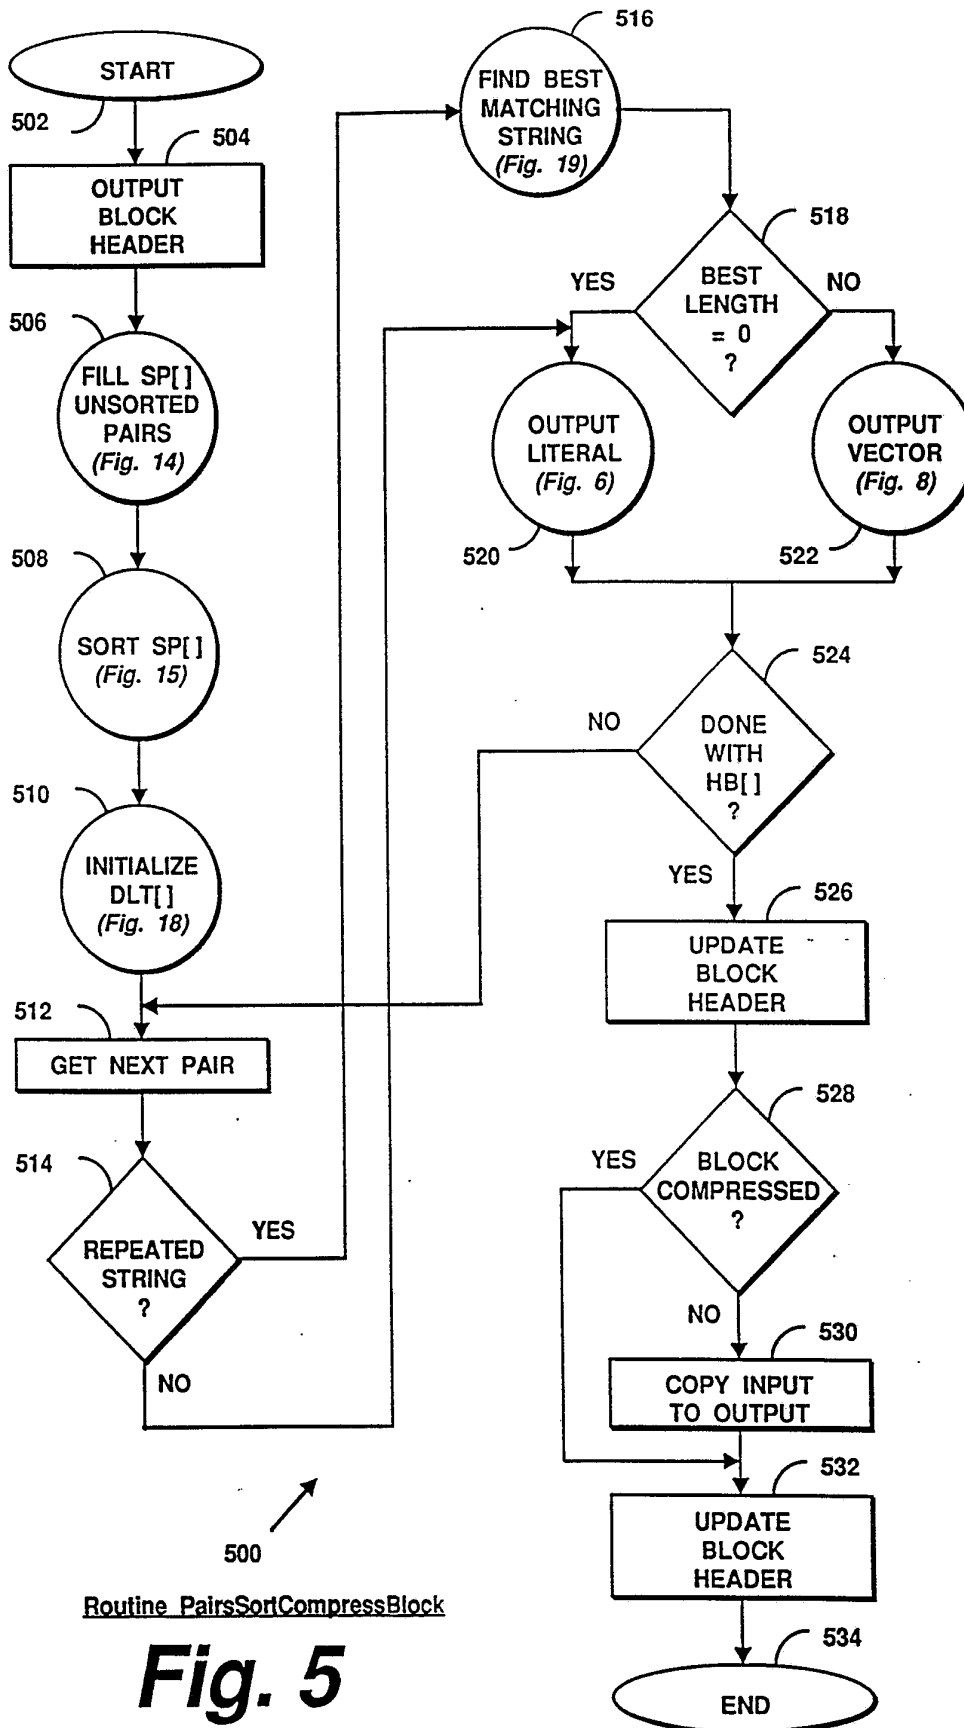


Fig. 4



Routine PairsSortCompressBlock

Fig. 5

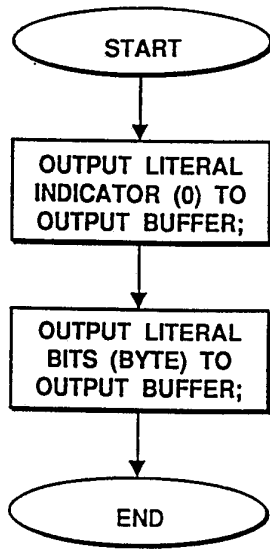


Fig. 6




Fig. 11

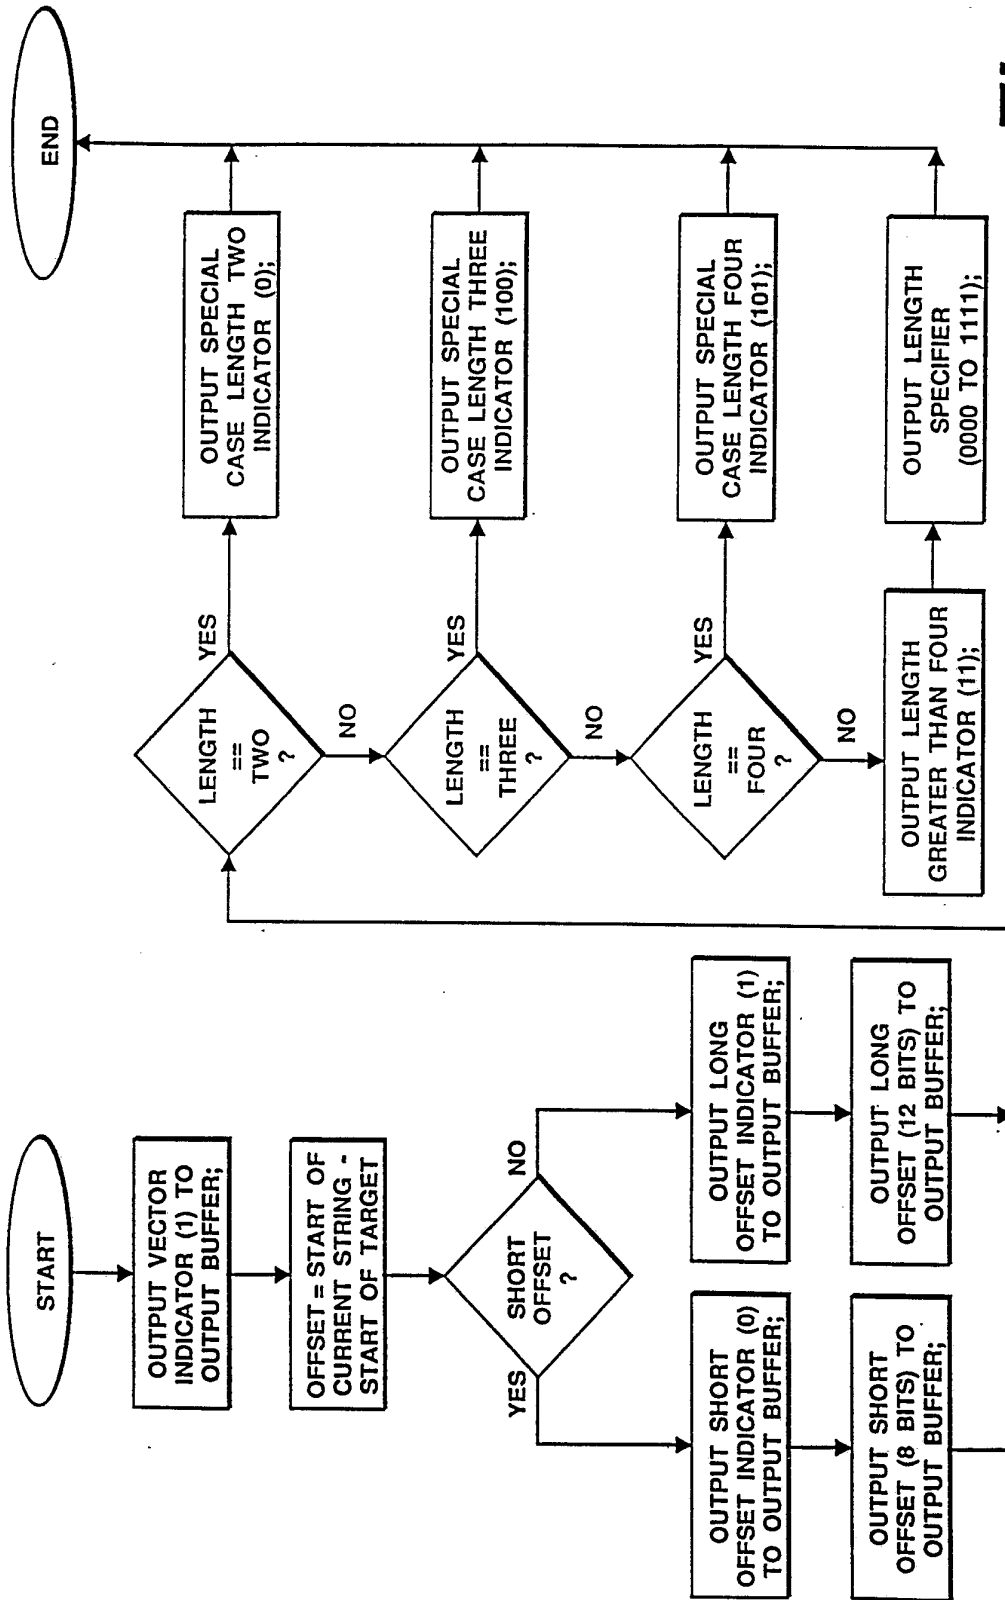


Fig. 8

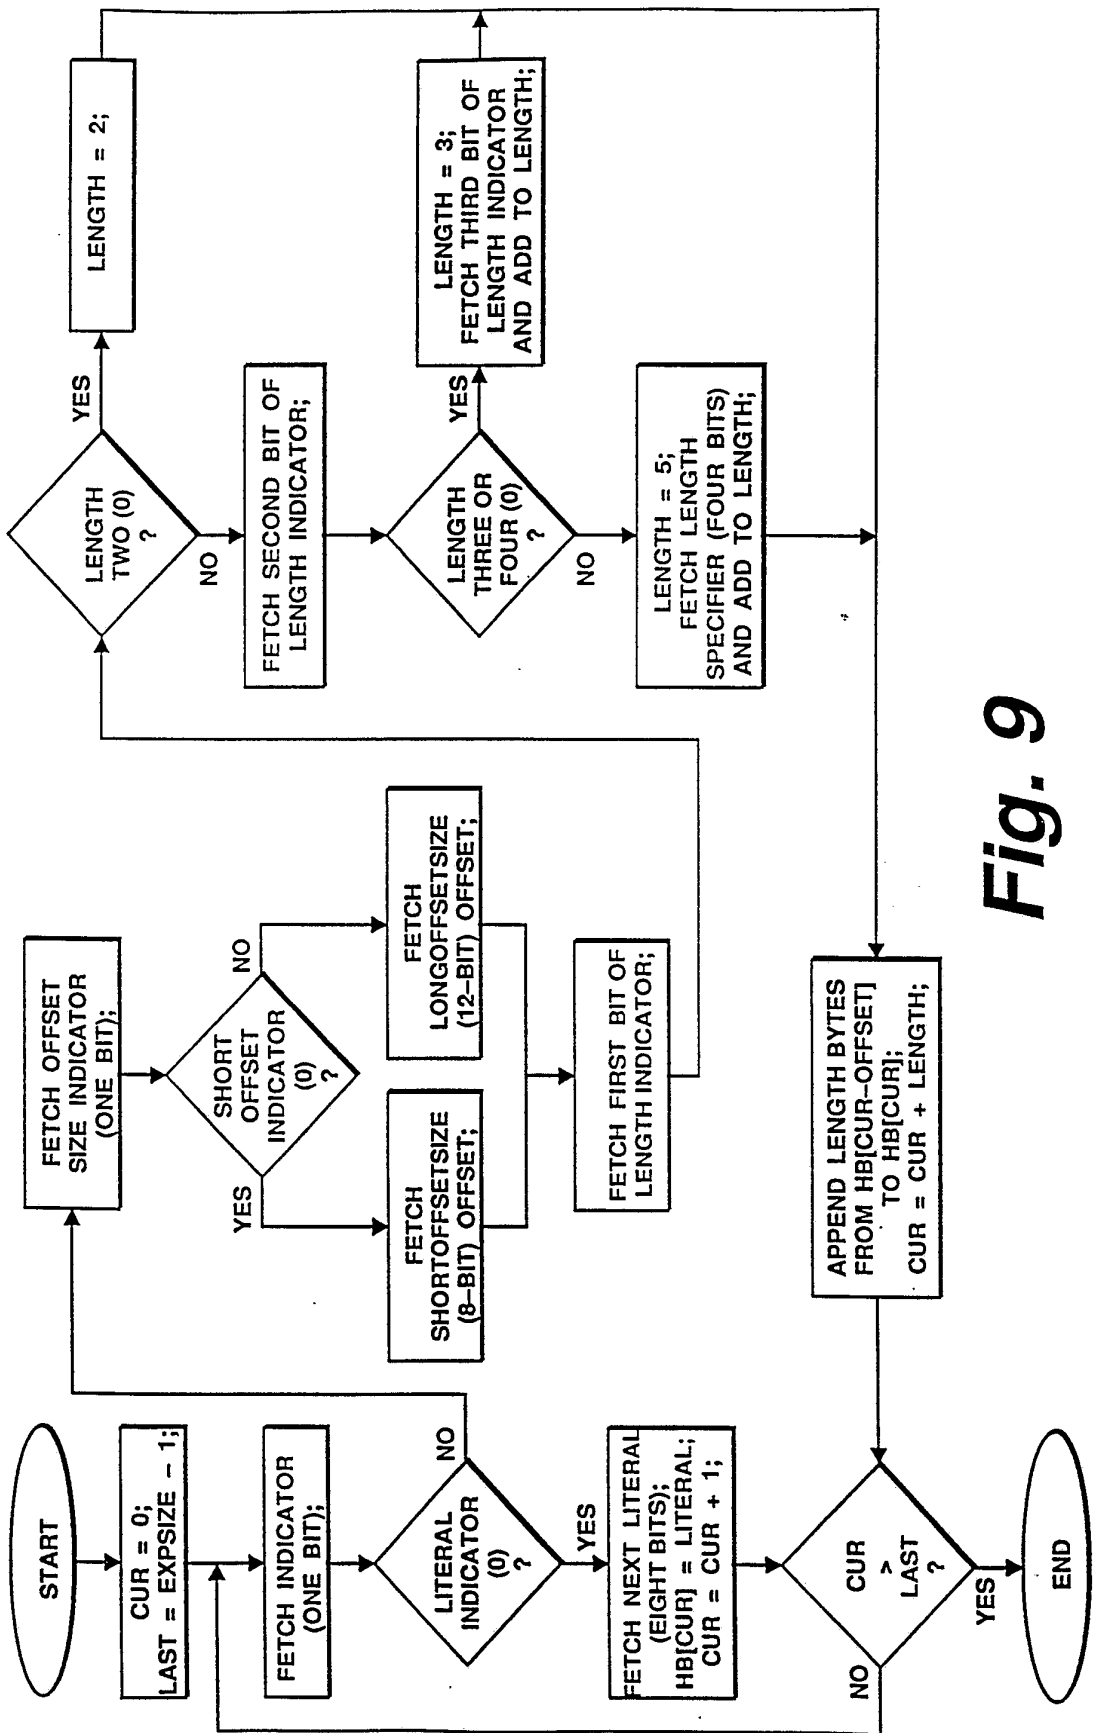


Fig. 9




Fig. 10

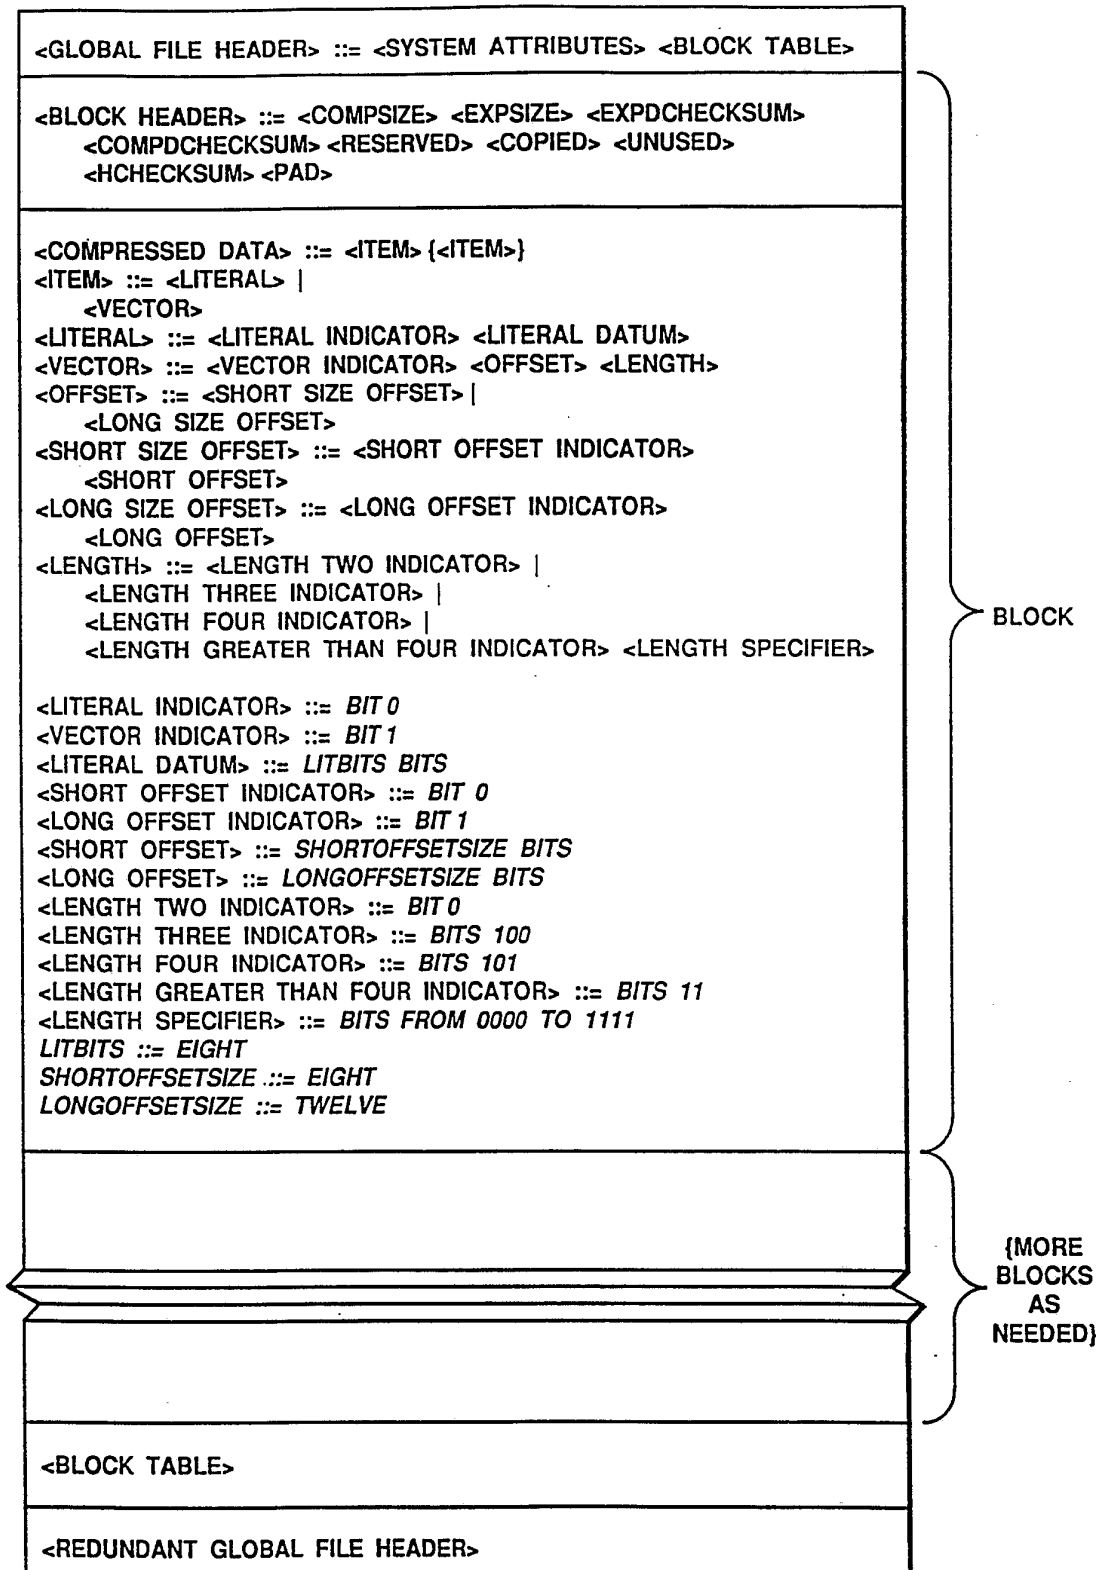


Fig. 12

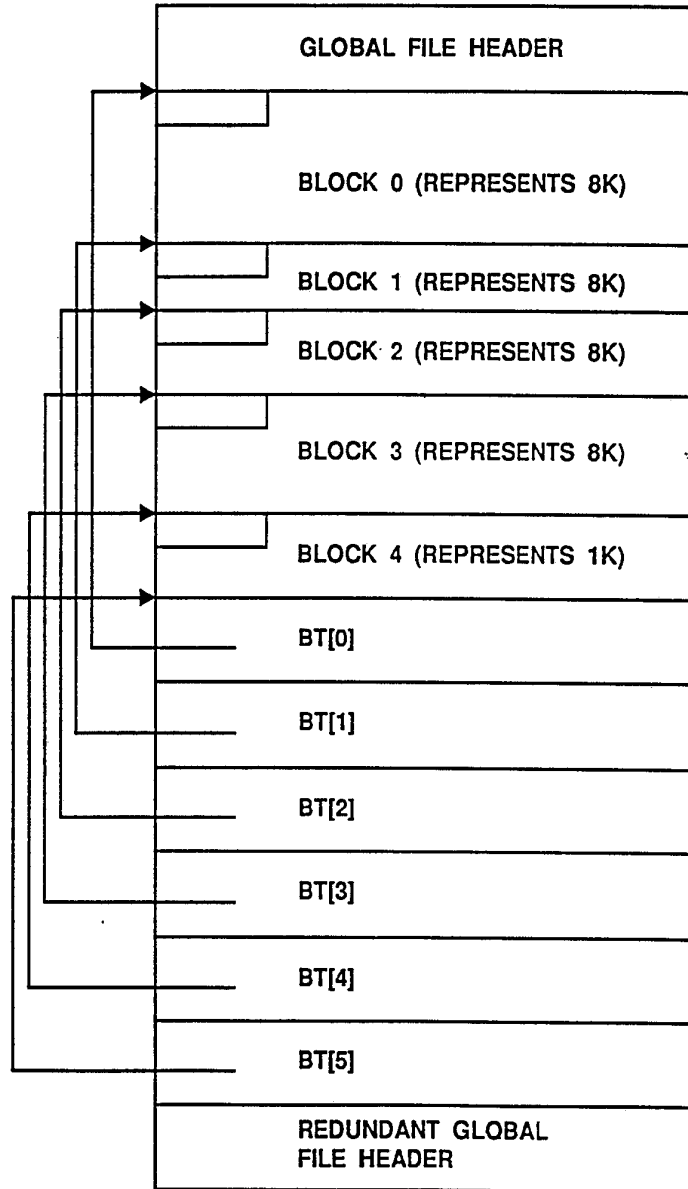


Fig. 13




Fig. 14



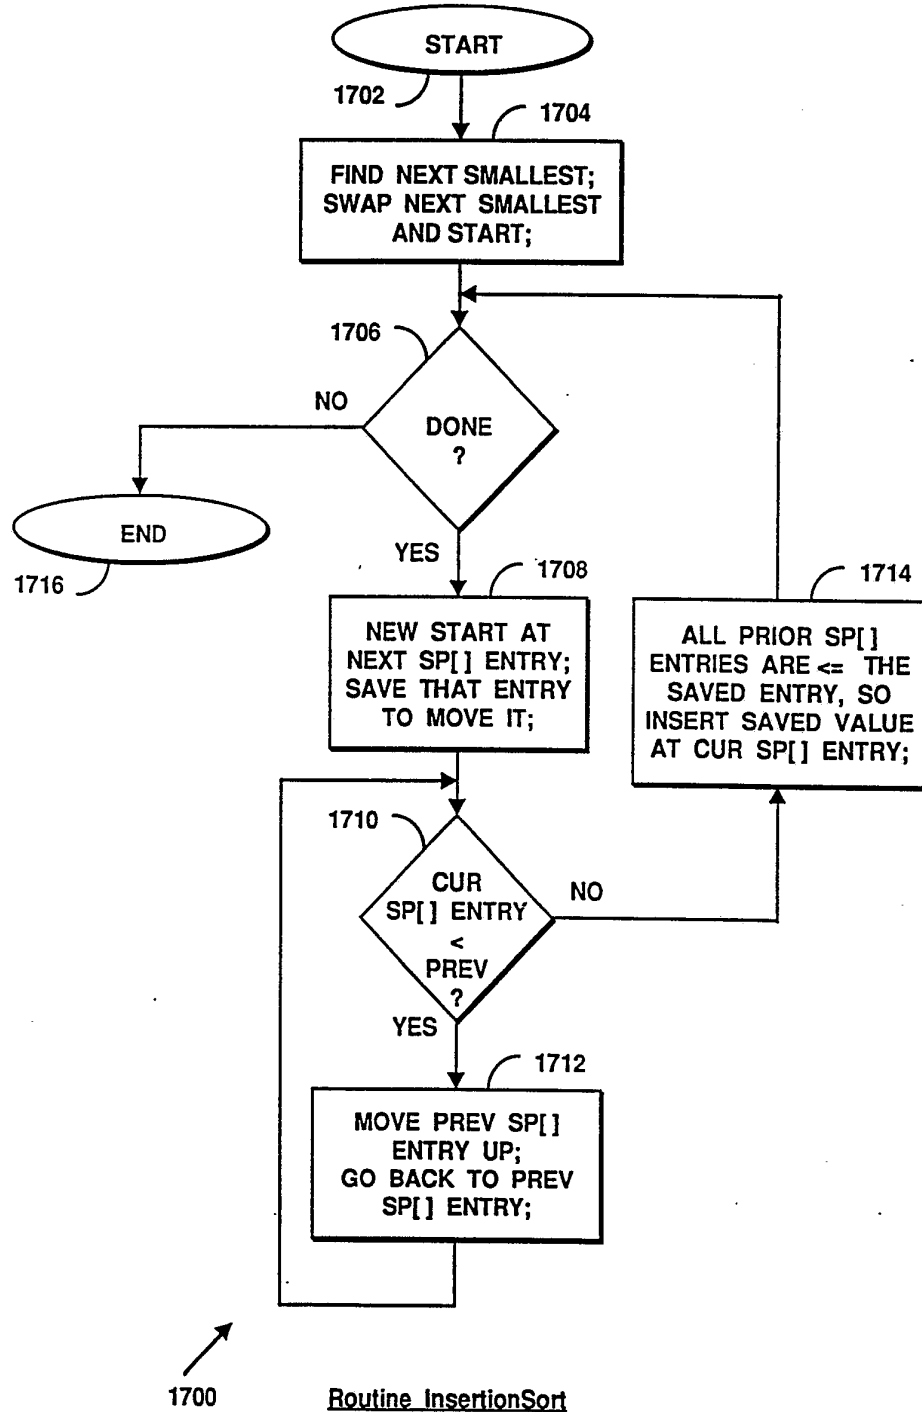


Fig. 15



Routine RealSortPairs

Fig. 16



Routine InsertionSort

Fig. 17





Fig. 18



Routine_MatchBestPair
Fig. 19

1900