


RTMP's
copy



SA9-93-001

1

DATA COMPRESSION APPARATUS

DESCRIPTION

This invention relates to the compression and decompression of data in data storage systems, and more particularly to the parallel execution of a sequential data compression algorithm.

It is known that the efficiency with which data can be compressed and decompressed depends primarily on the buffer size and encoding implementation used. Executing a compression/decompression algorithm using software is slow and therefore not suited to high speed or real time applications. Executing the algorithm using hardware requires an amount of hardware that varies according to the degree of parallelism of the implementation technique employed. If too much hardware is required, it may be difficult to integrate the data compression algorithm into a controller.

A paper by Lempel and Ziv entitled "A Universal Algorithm for Sequential Data Compression", published in IEEE Transactions on Information Theory, May 1977, at pp. 337-343, describes an algorithm for efficiently compressing data.

This Lempel-Ziv 1 (LZ1) algorithm is a sequential algorithm that compresses strings of binary data of variable length into a fixed length compressed binary format. It is implemented using a history buffer that contains the most recent bytes or words of a file in the correct sequence. Methodically, by repeated executions of a basic routine, new bytes are read as long as the sequence of incoming bytes is matched by a sequence in the history buffer, thereby generating a sequential stream of data. Since each incoming byte is sequentially compared with each byte in the history buffer, a significant amount of computation time is required, making this technique unsuitable for real time applications.

Commonly assigned U.S. serial number 07/807,007, filed December 31, 1991, (Docket AT991-030) (EPA 546863) describes a typical implementation of the LZ1 algorithm and then cites a number of patents (not deemed material to the present invention) which cover techniques of improving

the speed with which the LZ1 algorithm is executed or the amount of compression achieved.

This cited commonly assigned application describes a fully parallel architecture that implements the LZ1 algorithm in hardware. With a
5 content addressable memory (CAM) serving as a history buffer, each incoming byte is compared simultaneously with all the bytes in the history buffer. This fully parallel hardware approach desirably provides the fastest execution of the LZ1 algorithm.

However, it requires a separate comparator for each distinct buffer
10 position (i.e., CAM address) and can only achieve the maximal efficiency (speed/hardware performance) when the history buffer is full; namely, after an initial loading period for each sector or input data field of the data storage medium. Therefore, if the sector is approximately the same size as the history buffer, the fully parallel implementation will
15 require many redundant operations.

Since the size of a device controller chip is essentially the same as that of the chip needed to implement the purely parallel compression, a parallel compression chip cannot efficiently be used to perform compression in a device controller. The principal use for this fully
20 parallel approach is for host data compression, where the compression chip is located in the host controller.

There is a need for a data compression/decompression apparatus and method which implements the LZ1 algorithm by use of a modular architecture that:

- 25 1. Divides the history buffer into a plurality of blocks, compares all bytes in a block in parallel, and scans the blocks sequentially;
2. Enables a designer to select any speed ranging from the slow sequential execution of the LZ1 algorithm to the above-described optimal
30 parallel implementation by selecting a desired degree of parallelism in order to limit hardware costs to the needs of a particular application;
3. Is especially suitable for applications in which the data compression is performed in a device controller, where the required execution speed is approximately an order of magnitude less than that
35 required to perform compression in a host controller; and

THIS SHEET HAS BEEN
REPLACED

SA9-93-001

3

4. Is especially advantageous when an input data sector and history buffer contain approximately the same number of bytes.

Accordingly, the invention provides data compression apparatus comprising:

- 5 an array of i identical slice units, each of which stores j symbols to define j separate blocks comprising one symbol from each slice unit, the symbols in each slice unit being separated by exactly i symbols;
 i comparators, each associated with a slice unit, to compare each symbol in a string of i incoming symbols in parallel with the symbols
10 previously stored in the slice units to identify matching sequences of symbols; and

means for controlling the comparators to scan symbols against the slice units in parallel but in each of said blocks sequentially.

- The parameters i and j can be selected to limit the number of said
15 comparators required to achieve a desired degree of efficiency in executing the algorithm based upon a trade-off of algorithm execution speed versus hardware cost.

- It is preferred that each slice unit provides an output signal indicating whether a matching sequence is identified. The output signal
20 identifies the last byte so far of a previously stored sequence that matches the most recent sequence in the string of incoming symbols. The apparatus further includes a priority encoder for calculating from said output signals each j, i address in which a matching sequence is identified, and means for selecting one of the j, i addresses for which a
25 matching sequence is identified. This strategy overcomes the problem that when scanning in parallel, more than one matching sequence may be found.

- Preferably the apparatus further includes means for inhibiting application of the algorithm to any blocks which do not contain symbols. The advantage of this is that it is not necessary to reset the blocks for
30 each new set of data, since the algorithm will not work on those that do not yet contain the new data. This is further accomplished by the use of a single latch in each of the slice units, said latches being serially connected to constitute in combination a shift register; and means for resetting said latches and hence said shift register before symbols are
35 stored in said array.

This Sheet Has Been Replaced



THIS SHEET HAS BEEN
REPLACED

SA9-93-001

4

Preferably the apparatus further includes means for causing matching sequences and nonmatching sequences of symbols to be stored in said array, and more particularly means for writing the incoming symbols serially into successive symbol positions in the array until all array
5 positions are filled, and then replacing the oldest symbol string in the array with the incoming symbol.

The invention also provides a method of compressing data comprising partitioning an array into i identical slice units, each of which stores j symbols to define j separate blocks comprising one symbol from
10 each slice unit, the symbols in each slice unit being separated by exactly i symbols;

comparing, using i comparators each associated with a slice unit, each symbol in a string of i incoming symbols in parallel with the symbols previously stored in the slice units to identify matching
15 sequences of symbols; and

controlling the comparators to scan symbols against the slice units in parallel but in each of said blocks sequentially.

Thus a sequential data compression algorithm can be executed in a manner that is especially suitable for use where data compression is
20 required in a device (as distinguished from host) controller. A history buffer comprises an array of i identical horizontal slice units. Each slice unit stores j symbols to define j separate blocks in which the symbols in each slice unit are separated by exactly i symbols. Symbols in a string of i incoming symbols are compared by i comparators in
25 parallel with symbols previously stored in the slice units to identify matching sequences of symbols. A control unit controls execution of the sequential algorithm to condition the comparators to scan symbols in parallel but in each of the blocks sequentially and cause matching sequences and nonmatching sequences of symbols to be stored in the array.
30 The parameters i and j are selected to limit the number of comparators required to achieve a desired degree of efficiency in executing the algorithm based upon a trade-off of algorithm execution speed versus hardware cost. A priority encoder calculates from signals output by the slice units each j, i address in which a matching sequence is identified,
35 but it outputs the address of only one (such as the smallest) of these addresses. Incoming symbols are written serially into successive symbol positions in the buffer until all buffer positions are filled and then

This Sheet Has Been Replaced



SA9-93-001

5

the oldest symbol string in the buffer is replaced with the incoming symbol.

The data compression method involves providing a history buffer partitioned into j blocks, each capable of storing i consecutive symbols, such that symbols in adjacent blocks at corresponding symbol positions are i symbols apart; providing i comparators; broadcasting in parallel each symbol of a string of i incoming symbols to said i comparators; writing the incoming symbols serially into successive symbol positions in the buffer until all buffer positions are filled; and then replacing the oldest symbol string in the buffer with the incoming symbol.

An embodiment of the invention will now be described in detail by way of example with reference to the following drawings:

Fig. 1 is a block diagram of a data processing system.

Fig. 2 is a schematic diagram of a data compression/decompression apparatus and including an array of identical horizontal slice units and a priority encoder.

Fig. 3 is a schematic diagram showing in detail the configuration of each of the horizontal slice units shown in Fig. 2.

Fig. 4 depicts a circuit logic diagram of a simplified version of the priority decoder shown in Fig. 2.

As illustrated in Fig. 1, a data processing system comprises a host computer 10 which transmits data to and receives data from a device controller 12 that includes compression/decompression apparatus 14. Apparatus 14 comprises a compression engine 15 and a decompression engine 16 which can be invoked, respectively, to compress and decompress data. Engine 15 provides compressed data output to an input/output (I/O) device 18, such as a disk drive comprising a plurality of disks. Device 18 provides compressed data input to the decompression engine 16 of device controller 12.

As illustrated in Fig. 2, data compression/ decompression apparatus 14 comprises a control unit 20, a history buffer 22, and a priority encoder 24. Input/output (I/O) bus 11 conveys data to and from control unit 20. The data is in the form of "symbols", a term used generically to connote bytes, halfwords, words, or any other preselected number of



bits. However, to facilitate understanding, the data is herein assumed to be in the form of bytes.

As illustrated, history buffer 22 consists of an array of 128 identical horizontal slice units HS0 to HS127, each of which stores four bytes, for thereby partitioning the buffer into four 128-byte stored blocks. This modular architecture creates a history buffer 22 capable of storing 512 (128 x 4) consecutive bytes in which the bytes in each HS unit are separated by exactly 128 bytes (i.e., one block size) and each HS unit stores four bytes, one from each block.

- Each byte b in buffer 22 has a unique address which is identified by its block index B and HS index. Thus, for bytes 0-127, the block index B is 0; for bytes 128-255, the block index B is 1; for bytes 256-383, the block index B is 2; and for bytes 384-511, the block index B is 3. The HS index is the byte address mod block size.
- Thus, the address of byte 356 would be [3,100]. Four bytes b , one from each block at the same HS address 0...127, comprise a word W whose address in this case would be (1,L;2,L;3,L;4,L) where L is a specific HS address. With the 4x128-byte buffer illustrated, if the address of a byte is expressed in binary format, the block index can be denoted by the two most significant bits and the HS unit index can be denoted by seven less significant bits.

Referring now to Fig. 3, each HS unit comprises a memory register or cell ME, an address selector cell SE, a comparator cell CO, a multiplexer MU and latches S1, S2, S3. Controller 20 transmits to each HS unit the same 8-bit byte of an incoming symbol string and four signals c1-c4 that control the multiplexer MU of that HS unit. Controller 20 also provides in parallel to each HS unit: signals r1 and r2 for use in resetting latches S1 and S2, respectively; block address selecting signals ad1 and ad2, respectively; a clock signal ck that synchronizes operation of all HS units; and a write enable signal w.

Each HS unit has outputs s1o, s2o, s3o which represent the current states of the respective latches S1, S2, S3. These outputs constitute the inputs s1i, s2i, s3i (see Fig. 2) for the next successive HS unit in a cyclical manner, with the outputs s1o, s3o from HS127 being the inputs

s1i, s3i to HS0. Multiplexer MU uses the control signals c1-c4, and generates an output m.

In operation, assume the content of buffer 22 (i.e., 1L, 2L, 3L, 4L, S1, S2, S3) is undefined. This state occurs whenever a new sector is compressed. Initially, a signal r1 is sent in parallel by control unit 20 to AND gate 42 in all HS units to reset the latches S1 in all HS units. A byte can only be written into a location in the memory cell of a given HS unit under the following conditions:

1. The write enable signal w from control unit 20 must be "1".
2. Latch S1 of this HS unit must be "0" (i.e., in reset state).
3. Input s1i to this HS unit must be "1".

These conditions are satisfied with the AND gate 50 as shown in Fig. 3. When a block is full, the latches S1 of all HS units except the last one (here, HS127) are set (i.e., are "1"). The HS unit with S1 = "0" is the location where the new incoming byte replaces an old byte in buffer 22. However, since the S1 latches of all HS units initially have the value "0", control unit 20 provides a "1" signal via line 44 and OR gate 46 (Fig. 2). This "1" signal is maintained until the S1 latches of all HS units except the last one (here, HS127) are "1". This procedure is repeated (i.e., resetting of the S1 latches in all HS units and maintaining the initializing "1" signal on line 44) until all blocks are full; whereupon the signal on line 44 is switched by control unit 20 to "0".

The latch S1 of each HS unit also provides the output slo of that unit. Consequently, the S1 latches of all HS units act as a 128-bit shift register. This shift register is updated by write enable signal w which is transmitted in parallel to all HS units from control unit 20.

Write enable signal w also causes an incoming byte from line 48 to be written into the memory cell ME of the HS unit. The location is determined by the latches S1 and inputs s1i to be one of the locations 1L, 2L, 3L or 4L as determined by the selector unit SE. One of the advantages of the latches S1 is the fact that the memory cells need not be reset at the beginning of a new sector, thus reducing power consumption. This feature is of practical importance in a device controller implementation where power dissipation is critical due to battery power source limitations.

To initiate a string comparison operation with a new block, control unit 20 sends a signal r2 in parallel to all HS units to reset latch S2 via AND gate 40. The output of s2 controls the selector cell SE in each HS unit. Also, control unit 20 sends two signals ad1 and ad2 to the SE cells of all HS units. These two-bit signals ad1 and ad2 denote adjacent block addresses (e.g., B1 and B2). Depending on the state of latch S2 in each HS unit, either ad1 or ad2 becomes the determining block address ado for that HS unit. Initially because all S2 latches are reset, ado equals ad1 for all HS units. The output of the S2 latch of an HS unit becomes output s2o and then input s2i for the next HS unit.

Thus all latches S2 form a 128-bit shift register whose input (s2i from HSO) is always "1". As matching of the input byte string with the content of the buffer 22 continues, this shift register is filled with "1"s from the bottom, thus causing more and more HS units to select ad2 instead of ad1. This constitutes a virtual shifting of the block boundaries during the matching process. For example, assume there has been a match up to byte 255 in buffer 22 (i.e., up to byte address [2,127]). Then, the next incoming byte has to be compared with byte 256 in the buffer, which has a different block address, namely, [3,0].

The incoming byte is also transmitted via branches of line 48 to a corresponding comparator cell CO in each HS unit. The comparator cells CO in all 128 HS units compare the incoming byte simultaneously with 128 bytes in that particular one of the four blocks chosen by the selector cells SE. If the output byte from one of the memory cells ME matches the incoming byte in line 48, comparator CO will provide a "1" output signal m to multiplexor unit MU.

Multiplexor unit MU is conditioned by control unit 20 into one of four modes or states, as determined by the signals c1-c4. The functions implemented by these signals are as follows:

Signal c1 denotes that the incoming byte will be the first byte of a new matching string and that the block which is currently active (i.e., whose address is given by ad0) is full. In this event, the state of latch S1 can be ignored.

Signal c2 denotes that the incoming byte will also be the first byte of a new matching string but that the currently active block is not yet full (i.e., the contents of some memory locations

in this block are undefined, and therefore a match can only occur when latch S1 is "1".

Signal c3 denotes that the incoming byte is not the first byte of a matching string, and that the state of latch S1 cannot be ignored. This condition can only occur if the whole buffer is full and the shifting area of one block overlaps with the area of another block into which bytes of the matching string are going to be written.

Signal c4 denotes that the incoming byte is not the first byte in a matching string but that the state of latches S1 can be ignored.

If an incoming byte is not the first byte of a matching string, then a continued match requires that the preceding byte must also have been matched, a condition which is denoted by the input signal s3i from the preceding HS unit. Thus, when the multiplexor MU is conditioned by signal c1 or c2, the signal s3i can be ignored because s3i denotes the match of an initial byte of a string. The m signal of comparator cell C0 is selectively ANDed with signal c1 at 52, with signals c2 and s1o at 54, with signals c3 and s1o and s3i at 56, and with signals c4 and s3i at 58. The outputs of the respective AND gates 52, 54, 56 and 58 are ORed at 60, generating the output signal m. If the output signal m is "1", it denotes that a matching string in buffer 22 has been extended to include the most recent incoming byte. The m outputs of all HS units are ORed at gate 62 (Fig. 2) to signal control unit 20 that the present input sequence has found an active match in the currently active block. This output signal m of an HS unit is also input to the latch S3. Therefore, if the output s3o of latch S3 of an HS unit is "1", it means that the buffer location determined by this unit and the address ado is the thus far last byte of a previously stored sequence that matches the most recent input sequence. The s3o outputs of each HS unit are connected to priority encoder 24.

It should be noted that updating of buffer 22 is independent of and can be done concurrently with a string comparison operation.

Priority encoder 24 uses the s3o outputs from the 128 HS units to calculate and encode the ending address of a sequence in history buffer 22 that matches the most recent sequence in the incoming data to be

compressed. In a parallel implementation, such as herein described, more than one sequence in the history buffer 22 may match the incoming sequence. In such event, encoder 24 will provide only a single ending address for a matching sequence.

5 The manner in which this singular ending address is selected will be demonstrated by reference to the simplified priority encoder 24' illustrated in Fig. 4. This encoder encodes 512 bytes, but to simplify the drawing and description assumes there are only 8 HS units, each storing 64 bytes, one from each of 64 blocks, to encode into only three
10 bits (rather than the seven bits that would be necessary to convert the address from 128 HS units of the type shown in Fig. 3.)

As illustrated in Fig. 4, the simplified priority encoder 24' receives as inputs the eight s3o signals from all (now 8) HS units. It then determines the lowest of the eight indices among the HS units for
15 which s3o is "1". This is achieved by the logic illustrated in Fig. 4 which ignores all the s3o signals except the one with the smallest HS address. Thus, it determines a unique ending address of a matched sequence. This address is sent to control unit 20.

Since the decompression algorithm requires the starting address of
20 a matching sequence, it is necessary to calculate this starting address from the ending address provided by the priority encoder 24. Since it is known that the matching sequence started in the currently active block (all S2 latches were reset initially), it is only necessary to calculate the last seven bits of the address. This can be done by simply
25 subtracting the length of the matching sequence; i.e., subtracting the number of bits in this sequence from the ending address and taking the result mod the number of HS units (128 in the embodiment illustrated in Fig. 3). The length of the matching sequence is part of the compressed data and therefore available in the control unit 20 where it is easily
30 obtained with an incrementing counter.

The algorithm for decompressing the data, which has been compressed by applicants' modified LZ1 compression algorithm, is inherently sequential and therefore independent of the degree of parallelism used in compression engine 15. However, for efficient hardware usage, the
35 history buffer 22 of compression engine 15 is used. The compressed

incoming data consists of a starting buffer address, a string length number, and a character that is the last symbol in the compressed string. Control unit 20 extracts the buffer content at the initial address and updates buffer 22 with the extracted symbol in the same fashion as
5 compression engine 15. This step is repeated for the length of the compressed string. During this process, the buffer address is constantly incremented in a cyclical fashion. Finally, the last symbol is also sent to host 10, and the control unit 20 is once again ready to receive the data of the next compressed string.

10 Attached hereto for illustrative purposes are APPENDIX A, a program in C language for implementing the LZ1 compression algorithm in accordance with the invention, and APPENDIX B, a program in C language for decompressing data compressed by use of applicants' modified LZ1
15 compression algorithm. The programs in both Appendices are provided with comments to provide additional illustrations of the sequences of steps in implementing the compression/decompression described above.

APPENDIX A - COMPRESSION

```

/***** This program simulates a partitioned data *****/
/***** compression algorithm LZ1 where each *****/
5 *****/ string is followed by a single byte *****/

#include <stdio.h>

#define B1 128 /* number of memory blocks */
#define B2 4 /* number of bytes inside a block */
#define L_K 6
10 #define L 64 /* L=2**L_K, maximum length of a string */
#define C_K 9
#define C 512 /* C=2**C_K, maximum block C = B1*B2 */
#define M_NR 10000
#define MAX(u,v) ((u>v) ? u : v)
15 #define MIN(u,v) ((u<v) ? u : v)

int s1[B1], s2[B1], s3[B1]; /* latches */
int match[B1]; /* intermediate result */
int a1, a2; /* new write position in buffer */
int p1, p2; /* actual buffer position in sweep */
20 int o1, o2; /* first buffer position for sweep */
int z1, z2; /* start position for string */
int buffer[B1][B2]; /* buffer */
int byte; /* input byte */
int oldbyte[2]; /* backup of the first two input bytes */
25 int recbyte; /* backup of the most recent byte */
int l, l_a; /* length values of strings */
int r; /* sweep counter */
int any_mt; /* there is at least one match */
int first, new, end; /* logical variables */
30 FILE *output; /* output file */

long c_input, c_output, c_iter, c_bits, c_cycles;
/* counter for input, output, iterations, output bits and cycles */
long c[L+1]; /* counter array for matches */

/* This function calculates the beginning of a string. */

35 void new_p()
{
/* The position of the pointer is calculated if there is a string match
of at least 2 */

if (new && l>=2) {
40 /* number of the memory block, assuming l < B1. */

/* One location for which a match is found (the first location). */
for (z1=0; !s1[z1]; z1++);

/* The location is calculated backwards to the original pointer position.
*/
45 if (z1<1-F) {

```

```

        z1 = z1+B1-1+1;
    }
    else {
        z1 = z1-1+1;
5      }
    z2 = r;
}

/* This function advances a pointer to a buffer location. */
10 void inc_p(int *p1, int *p2)
{
    *p1 = (*p1+1) % B1;
    if (!(*p1))
        (*p2)++;
15 }

/* This function writes a buffer location and advances the corresponding
   pointer. */

void write_b()
{
20     int i;

    /* The current byte has just been read in, it must be stored in the
       buffer. */
    recbyte      = byte;
    buffer[a1][a2] = byte;

25     /* The shift register for the buffer locations must be advanced. */
    if (s3[B1-2])
        for (i=0; i<B1; i++)
            s3[i] = 0;
    else {
30         for (i=B1-1; i>0; i--)
            s3[i] = s3[i-1];
        s3[0] = 1;
    }

    /* The position of the current buffer pointer is advanced. */
35     inc_p(&a1, &a2);
}

int block_c()
{
40     int i, j;
    int t, k;

    /* Resetting of all the counters. */
    c_input = c_output = c_iter = c_bits = c_cycles = 0;

    /* The first byte is read in. */
    if ((byte = getchar()) == EOF)
45     return 1;
    c_input++;

```

```

/* The pointer of the buffer is reset and the latches s1 are reset. */
    a1 = a2 = 0;
    end = 0;
    for (i=0; i < B1; i++)
5       s1[i] = s3[i] = 0;

/* Start of a new string: first and new are set, the length l is 0,
   o1 and o2 reflect the current buffer location, start with the current
   buffer location (j = a2) and the number of sweeps r is reset. */

    while (1) {
10       first = new = 1;
        l = 0;
        o1 = a1;
        o2 = a2;
        j = a2;
15       r = a2;

/* Start with a new sweep: all s2 latches are reset (start with the same
   position in all memory banks). */
        while (1) {
            for (i=0; i < B1; i++)
20               s2[i] = 0;

/* Start with a new matching */
            while (1) {

/* The selected content in all memory banks is compared to the current
   byte. This requires a cycle. */
25             c_cycles++;
             c_iter++;

/* The byte in question is compared with the selected content in the
   memory banks. */
            any_mt = 0;
30             for (i=0; i < B1; i++) {

/* Considering the cyclic structure of the buffer. */
                t = (byte == buffer[i][j+s2[i]]);

/* The match is different for the first byte of a prospective string and
   for the remaining bytes. */

35                 if (first && new)
                     match[i] = (t && s3[i]);
                 else if (first)
                     match[i] = t;
                 else
40                     match[i] = (i) ? (t && s1[i-1]) : (t && s1[B1-1]);

/* Has there been any match? */
                any_mt = (any_mt || match[i]);
            }

/* Savings of bytes in temporary buffers */
45             if (new) {

```

```

/* The first or second byte of a prospective string is saved. */
    if (l<=1)
        oldbyte[1] = byte;

/* The byte is written in the buffer and saved in the recent byte buffer.
5 */
    write_b();
}
/* Update the next position of the buffer */
    else
10     inc_p(&p1, &p2);

/* There has been a match */
    if (any_mt) {

/* The match results are stored in the shiftregister s1 */
        for (i=0; i < B1; i++)
15         s1[i] = match[i];

/* We are not at the beginning of the string any more. */
        first = 0;

/* Updating of the s2 cyclic shift array. */
        if (s2[B1-1]) {
20             j++;
            for (i=0; i < B1; i++)
                s2[i] = 0;
        }
        else {
25             for (i=B1-1; i > 0; i--)
                s2[i] = s2[i-1];
            s2[0] = 1;
        }

/* The current string length matches the previous overall string length
30 and becomes the new overall string length. */
        if (!new && (l==l_a))
            new = 1;

/* The next byte is obtained from the input. */
        if (new) {
35             l++;
            if (c_input == C || l == L) {
                end = 1;
                break;
            }
40             c_input++;
            if ((byte = getchar()) == EOF)
                break;
        }

/* The next byte has already been read in. */
45     else {
        l_a++;
/* The next byte is obtained from the old byte buffer. */
        if (l_a == 1)
            byte = oldbyte[1];
50 /* The next byte is obtained from the recent byte buffer. */

```

```

        else if (l_a == 1)
            byte = recbyte;
        /* The next byte is obtained from the buffer, this costs a cycle. */
        else {
5             byte = buffer[p1][p2];
              c_cycles++;
        }
    }
10 /* No match, therefore a new iteration must be started. */
    else
        break;
    /* end inner while */

    /* Either the end of the file is reached, the string has its maximal
15 length, the block size is maximal or the sweeps are completed. */
    if (byte == EOF || l == L || r == 0 || c_input == C) {
        /* The starting pointer of the string is determined if not already done.
        */
        new_p();
20         break;
    }

    /* We can continue with the next sweep. */
    else {
        /* The starting pointer of the string is determined. */
25         new_p();

        /* The binary variable new is reset for the next run */
        new = 0;

        /* The next comparison is the first byte of a string. This byte is
        taken from the oldbyte buffer. */
30         first = 1;
        byte = oldbyte[0];

        /* The length of the temporary string is reset. */
        l_a = 0;
        /* The number of the current sweep is incremented. */
35         r--;
        /* The position of the first byte of the previous string in the buffer
        is taken. */
        p1 = o1;
        p2 = o2;

40 /* The sweep starts with the next position in the memory blocks (cyclic).
        */
        j = r;
    }
    /* end of middle while loop */

45 /* No match, 9 bit output. */
    if (l==0) {
        if (byte != EOF)
            fprintf(output, "0 %d\n", recbyte);
        c_output++;
50         c_bits += 9;
    }

```

```

/* Only 1 byte match, 18 bit output. */
else if (l==1) {
    fprintf(output, "0 %d ", oldbyte[0]);
    if (byte != EOF && !end)
5      fprintf(output, "0 %d\n", recbyte);
    c_output += 2;
    c_bits += 18;
}
/* At least a match of 2 bytes. */
10 else {
    fprintf(output, "1 %d", z1);
    fprintf(output, " %d", z2);
    fprintf(output, " %d\n", l);
    if (byte != EOF && !end)
15      fprintf(output, "0 %d\n", recbyte);
    end = 0;
    c_output += 3;

/* Output depends on the amount of data in the buffer. */
    for (i=0, k=1; k<=B2*B1; i++, k*=2)
20      if (o1+B1*o2 < k) {
          c_bits += 10+i+L_K;
          break;
      }
    }

25    (c[l])++;

    if (byte == EOF || c_input == C)
        break;

    else {
        c_input++;
        if ((byte = getchar()) == EOF) {
30          c_output++;
          (c[0])++;
          c_bits += 9;
          break;
35        }
    }
}

/* end of outer while loop */

return (byte == EOF);
}

40 int main()
{
    float d_input, d_output, d_iter, d_bits, d_cycles;
    long mi_output, mi_iter, mi_bits, mi_cycles;
    long ma_output, ma_iter, ma_bits, ma_cycles;
45    float mi_rat1, mi_rat2, ma_rat1, ma_rat2, h_rat;
    int i;
    int j;

/* Resetting of all the counters. */
    d_input = d_output = d_iter = d_bits = d_cycles = 0;

```

```

mi_output = mi_iter = mi_bits = mi_cycles = M_NR;
ma_output = ma_iter = ma_bits = ma_cycles = 0;
mi_rat1 = mi_rat2 = 2;
ma_rat1 = ma_rat2 = 0;

5   output = fopen("comp.dat", "w");

   for (j=0; j<L; j++)
       c[j] = 0;
   i = 0;

   while (1) {
10      if (block_c())
           break;
       i++;
       d_input += (float) c_input;
       d_output += (float) c_output;
15      d_iter += (float) c_iter;
       d_bits += (float) c_bits;
       d_cycles += (float) c_cycles;
       mi_output = MIN(mi_output, c_output);
       ma_output = MAX(ma_output, c_output);
20      mi_iter = MIN(mi_iter, c_iter);
       ma_iter = MAX(ma_iter, c_iter);
       mi_bits = MIN(mi_bits, c_bits);
       ma_bits = MAX(ma_bits, c_bits);
       mi_cycles = MIN(mi_cycles, c_cycles);
25      ma_cycles = MAX(ma_cycles, c_cycles);
       h_rat = (float) (c_output*9) / (float) (c_input*8);
       mi_rat1 = MIN(mi_rat1, h_rat);
       ma_rat1 = MAX(ma_rat1, h_rat);
       h_rat = (float) (c_bits) / (float) (c_input*8);
30      mi_rat2 = MIN(mi_rat2, h_rat);
       ma_rat2 = MAX(ma_rat2, h_rat);
   }

   if (i<=0)
       printf("The file has no content");
35  else {
       printf("Number of blocks: %d\n", i);
       printf("Inputbytes: %.0f ", d_input);
       printf("Inputbits: %.0f\n", d_input*8);
       printf("Output-9-bits: Summe %.0f ", d_output);
40      printf("Average %.2f ", d_output/i);
       printf("Maximal %d ", ma_output);
       printf("Minimal %d\n", mi_output);
       printf("Outputbits: Summe %.0f ", d_output*9);
       printf("Average %.2f ", d_output*9/i);
45      printf("Maximal %d ", ma_output*9);
       printf("Minimal %d\n", mi_output*9);
       h_rat = (d_output*9) / (d_input*8);
       printf("Comp_ratio1: Average %.4f ", h_rat);
       printf("Maximal %.4f ", ma_rat1);
50      printf("Minimal %.4f\n", mi_rat1);
       h_rat = (d_bits) / (d_input*8);
       printf("Comp_ratio2: Average %.4f ", h_rat);
       printf("Maximal %.4f ", ma_rat2);
       printf("Minimal %.4f\n", mi_rat2);

```

```
printf("Iterations: Summe %.0f ", d_iter);
printf("Average %.2f ", d_iter/i);
printf("Maximal %d ", ma_iter);
printf("Minimal %d\n", mi_iter);
5 printf("Cycles: Summe %.0f ", d_cycles);
printf("Average %.2f ", d_cycles/i);
printf("Maximal %d ", ma_cycles);
printf("Minimal %d\n", mi_cycles);
for (j=0; j<=L; j++)
10     if (c[j])
        printf("Matches for length %d: %d\n", j, c[j]);
    }
    fclose(output);
    return;
15 }
```

APPENDIX B - DECOMPRESSION

```

/***** This program simulates a partitioned data *****/
/***** decompression algorithm LZ1 where each *****/
5 *****/ string is followed by a single byte *****/

#include <stdio.h>

#define B1 128 /* number of memory blocks */
#define B2 4 /* number of bytes inside a block */
#define L_K 6
10 #define L 64 /* L=2**L_K, maximum length of a string */
#define C_K 9
#define C 512 /* C=2**C_K, maximum block C = B1*B2 */
#define M_NR 10000
#define MAX(u,v) ((u>v) ? u : v)
15 #define MIN(u,v) ((u<v) ? u : v)

int a1, a2; /* new write position in buffer */
int buffer[B1][B2]; /* buffer */
FILE *input, *output; /* input and output for extension */
*/

20 /* This function advances a pointer to a buffer location. */

void inc_p(int *p1, int *p2)
{
    *p1 = (*p1+1) % B1;
    if (!(*p1))
25     (*p2)++;
}

/* This function writes a buffer location and advances the corresponding
pointer. */

void write_b(int byte)
30 {
    /* The current byte has just been read in, it must be stored in the
buffer. */
    buffer[a1][a2] = byte;

    /* The position of the current buffer pointer is advanced. */
35     if ((a1 == B1-1) && (a2 == B2-1))
        a1 = a2 = 0;
    else
        inc_p(&a1, &a2);
}

40 int main()
{
    int i;

```

```
int byte;
int datum1, datum2, datum3;
int flag;

input = fopen("comp.dat", "r");
5 output = fopen("outp.dat", "w");
  al = a2 = 0;
  while (fscanf(input, "%d", &flag) != EOF) {
    fscanf(input, "%d", &datum1);
    if (!flag) {
10      fputc((unsigned char) datum1, output);
      write_b(datum1);
    }
    else {
      fscanf(input, "%d", &datum2);
15      fscanf(input, "%d", &datum3);
      for (i=0; i < datum3; i++) {
        byte = buffer[datum1][datum2];
        fputc((unsigned char) byte, output);
        write_b(byte);
20        inc_p(&datum1, &datum2);
      }
    }
  }
25 fclose(input);
  fclose(output);
  return;
}
```



THIS SHEET HAS BEEN
REPLACED

SA9-93-001

22

CLAIMS

1. Data compression apparatus comprising:
an array (22) of i identical slice units, each of which stores j symbols to define j separate blocks comprising one symbol from each slice
5 unit, the symbols in each slice unit being separated by exactly i symbols;
 i comparators, each associated with a slice unit, to compare each symbol in a string of i incoming symbols in parallel with the symbols previously stored in the slice units to identify matching sequences of
10 symbols; and
means for controlling the comparators to scan symbols against the slice units in parallel but in each of said blocks sequentially.
2. The apparatus of claim 1, wherein each slice unit provides an output signal indicating whether a matching sequence is identified, and
15 said apparatus further includes a priority encoder (24) for calculating from said output signals each j, i address in which a matching sequence is identified, and means for selecting one of the j, i addresses for which a matching sequence is identified.
3. The apparatus of claim 2, wherein said output signal identifies the
20 last byte so far of a previously stored sequence that matches the most recent sequence in the string of incoming symbols.
4. The apparatus any preceding claim, further including means for inhibiting application of the algorithm to any blocks which do not contain symbols.
- 25 5. The apparatus of any preceding claim, including:
a single latch in each of the slice units, said latches being serially connected to constitute in combination a shift register; and
means for resetting said latches and hence said shift register before symbols are stored in said array.

This Sheet Has Been Replaced





THIS SHEET HAS BEEN
REPLACED

SA9-93-001

24

6. The apparatus of any preceding claim, further including means for causing matching sequences and nonmatching sequences of symbols to be stored in said array.

7. The apparatus of claim 6, further including means for writing the
5 incoming symbols serially into successive symbol positions in the array until all array positions are filled; and
then replacing the oldest symbol string in the array with the incoming symbol.

8. A method of compressing data comprising
10 partitioning an array (22) into i identical slice units, each of which stores j symbols to define j separate blocks comprising one symbol from each slice unit, the symbols in each slice unit being separated by exactly i symbols;
comparing, using i comparators each associated with a slice unit,
15 each symbol in a string of i incoming symbols in parallel with the symbols previously stored in the slice units to identify matching sequences of symbols; and
controlling the comparators to scan symbols against the slice units in parallel but in each of said blocks sequentially.

20 9. The method of claim 8, including the step of:
selecting parameters i and j to limit the number of said comparators required to achieve a desired degree of efficiency in executing the algorithm based upon a trade-off of algorithm execution speed versus hardware cost.

This Sheet Has Been Replaced

