

Description

This invention relates to methods for correcting multiple byte errors in data read from a magnetic medium, compatible with data recorded in a count-key-data (CKD) or fixed block (FB) environment with inter-record gaps, and capable of being realised in an embodiment which can correct conventional errors using two-level error correction code (ECC) syndromes as well as correcting occasional long bursts of consecutive errors due to a major signal perturbation.

It is known to correct conventional errors in data read from a magnetic medium by use of a pointer, and check bytes, and one or more parity checks. These error correction code (ECC) error detection and correction functions are carried out on both write and read operations. These methods for accomplishing the error correction vary, depending on where the error is, the correction capability available in the hardware, system error recovery procedures, and the like.

The IBM 3380 Models J and K disk files use a two-level ECC to correct errors and are capable of correcting any single burst error involving two successive bytes (16 bits) in each sub-block read from the disk and any error involving four successive bytes (32 bits) in one of the sub-block in a block. However, there are instances where the duration of a perturbation and resultant error burst may be 16 bytes or more. This is longer than that of a "conventional" error burst and typically so long that it will exceed the correction capability of the two-level ECC heretofore available without creating a very wasteful design.

While the IBM 3850 Mass Storage Device is capable of correcting long-burst errors in data read from a magnetic tape, the method and apparatus used are not capable of correcting both conventional and long-burst errors in a disk file using a two-level ECC.

EP 0,240,363 describes a correction method for correcting burst errors that may occur in product-coded data. The product code is generated by considering the data as a two-dimensional matrix having error correcting codes defined over each row and over each column, with the corresponding generated check bytes included in the matrix. Data is stored as a sequential stream, with the check bytes for these error correcting codes being interleaved with the data. The technique described for correcting burst errors uses the row error correcting codes to identify the occurrence of a burst error and the column error correcting codes to provide correction of the affected bytes of data.

There is a need for some arrangement capable of correcting both conventional errors and long-burst errors of 16 or more consecutive bytes due to a major perturbation of the signal during reading of data from a magnetic medium, such as a magnetic disk.

The present invention provides a method of correcting multiple byte errors in data retrieved from a

magnetic medium on which data is recorded in variable length blocks that comprise sub-blocks of data bytes and, for each sub-block, check bytes providing error correction code (ECC) for which ECC syndromes are generated during reading, the method comprising the steps of: (a) writing at the end of each block a sequence of N sequential parity check bytes representing an interleaved parity check on data and ECC; (b) during reading, after the generation of the ECC syndromes, generating parity syndromes from the data, ECC and parity check bytes as read; the method being characterised by the steps of: (c) the parity check bytes being written at step (a) in the gaps between adjacent blocks; (d) in response to a long-burst error condition, generating a pointer to the first of the N consecutive bytes in a block that are potentially in error due to the long-burst error; (e) using any ECC syndromes which are not affected by the N bytes identified by the pointer, correcting correctable errors in all sub-blocks not affected by these identified bytes; (f) adjusting the parity syndromes according to the errors thus corrected; (g) using the adjusted parity syndromes, correcting the errors in the N bytes indicated by the pointer; (h) adjusting any unused ECC syndromes according to the errors corrected by the adjusted parity syndromes; and (i) using the adjusted ECC syndromes, correcting all correctable errors then remaining.

Towards this end, embodiments of the invention are disclosed which are capable of correcting not only conventional errors covered by the two-level ECC but also a long burst of up to N consecutive byte errors in a magnetic disk medium on which data is recorded in form of sub-blocks within a block of variable or fixed length. These blocks comprise data bytes and check bytes and include ECC for which ECC syndromes are generated during reading. At the end of each block a sequence of N sequential parity check bytes is written. Then, during reading, after the generation of the ECC syndromes, parity syndromes are generated by comparing parity check bytes computed from the data bytes and check bytes as read with the parity check bytes as written. In response to a long-burst error, a pointer is generated to the first of N consecutive bytes in a block that could have been influenced by the error. Then those ECC syndromes not dependent on the N bytes identified by the pointer are used to correct any correctable errors. Following this, the parity syndromes are adjusted according to the errors thus corrected. The adjusted parity syndromes are then used to correct correctable errors in the N bytes indicated by the pointer. Any unused ECC syndromes are adjusted according to the errors thus corrected by the adjusted parity syndromes; whereupon the adjusted ECC syndromes are used to correct all correctable errors then remaining.

The present invention will be described further by way of example with reference to an embodiment

thereof, as illustrated in the accompanying drawings in which:-

Fig. 1 illustrates a two-level ECC format;

Fig. 2 illustrates how each parity byte is computed from the consecutive ordered sets of N (assume as 16) interleaved data bytes and check bytes in a sub-block to form N parity groups;

Fig. 3 illustrates an embodiment of the logic for generating the parity bytes in writing and the parity byte syndromes in reading; and

Figs. 4A, 4B, 4C and 4D illustrate various error conditions under which imaged errors are correctable.

One method implementing the invention is shown applied to a two-level ECC system of the type disclosed in U.S. Patents 4,525,838, 4,703,485 and 4,706,250 constituting improvements thereto. These patents disclose a two-level multibyte error correcting system operable to correct up to t_1 errors in each sub-block of a block of data and t_2 errors in the block, where $t_2 > t_1$.

In the two-level ECC system as described in U.S. Patent 4,706,250 and illustrated in Fig. 1, data is recorded in blocks of variable (or fixed) length each comprising a plurality of sub-blocks, preferably of the same length except possibly the last. Each sub-block comprises a variable (or fixed) number of data bytes and six sub-block check bytes S for first level error correction. At the end of the last sub-block in the block are four CRC bytes C for data integrity checking, followed by two block check bytes B for second level error correction.

According to the disclosed arrangement, a sequence of N parity check bytes P is written at the end of each block in the gap between it and the succeeding block; i.e. in the inter-block gap. The parity byte P_0 immediately follows the last byte of the area; i.e., B_1 , the last block check byte. As illustrated in Fig. 2, these parity check bytes P make an interleaved parity check across N consecutive ordered groups of data bytes, check bytes and CRC bytes; i.e., all bytes in each ordered group are N bytes apart, and all bytes in the block are counted in determining the displacement in the count/key/data (CKD) area to which the parity bytes are appended. The number N is preselected to equal the longest anticipated burst of consecutive errors so that no single long-burst error can affect more than one byte in any one of the N parity groups. As customary in parity checks, all bytes are exclusively OR'd (XOR'd) As used in this specification and in the claims, the term "byte" is defined as comprising one or more bits.

A pointer is generated in response to an error burst of up to N consecutive bytes to indicate the first of N consecutive bytes in a block that could have been affected by the long-burst. The pointer may be an after-the-fact indication. For example, if it could be sensed that a long-burst error had started no more

than ten byte times earlier, then the appropriate earliest byte indication would be ten bytes previous to the byte being read at the time the error was sensed.

In Fig. 2, it is assumed for sake of illustration that N is 16, and that parity bytes P_0 through P_{15} are therefore written in order immediately following the last byte (block check byte B_1) of the block. Thus, the parity group including P_0 is the result of XOR'ing bytes D_2 , D_{18} , S_0 of sub-block X and D_{10} and D_{26} of sub-block Y; and the parity group including P_{11} is the result of XOR'ing bytes D_{13} , D_{29} of sub-block X and D_5 , D_{21} of sub-block Y and C_1 .

Fig. 3 depicts an exclusive OR gate (XOR) 20, a parity shift register 21 and multiplexers 22 and 23 each of which is one byte wide. Register 21 has a length of N stages (assumed, as already stated, as 16).

Assume initially that shift register 21 has been reset to zero in conventional manner. To write on a magnetic medium, multiplexers 22 and 23 are conditioned to connect Write Data line W to XOR 20 and gate all bytes BL of each block via Data In line DI and multiplexer 23 to the Output for writing on a selected track of a direct access storage device (not shown). Thus Data In includes all data bytes followed by appropriate check bytes through B_1 as illustrated in Fig. 1. At the end of Data In, multiplexers 22 and 23 are conditioned to connect line Φ to XOR 20 and gate the N parity bytes P_0 through P_{15} from the parity shift register 21 to the Output via line SR.

To read data from the magnetic medium and calculate the parity byte syndromes, multiplexers 22 and 23 are conditioned to connect Read Data line R to Data In line DI and gate all bytes BL of each block to the Output via line DI. Thereupon, multiplexers 22 and 23 are conditioned to gate the N parity bytes P_0 through P_{15} from line R to XOR 20 and gate the output of XOR 20 to the Output via line X.

The resultant output is the BL bytes of the block with the 16 parity byte syndromes appended.

In operation, assume that during reading, there are errors detected, but none are long-burst errors. Under this condition, error correction and verification will be effected conventionally, as explained in the cited prior art using the sub-block check, block check and CRC syndromes. At the sub-block level up to t_1 errors can be corrected and $t_1 + c$ errors detected. Thereafter, at the block level, up to t_2 errors are correctable in any one of the sub-blocks of the block.

If there are more errors than can be corrected, correction is abandoned and the error is treated as uncorrectable. Such structure and mode of operation form no part of the present invention.

Note that in the absence of a long-burst error, the parity bytes need not be read. Hence, under normal conditions, with no long-burst error, preparations for processing the next block can be started without reading the parity bytes. Conventionally this process-

ing is done as the read head passes over the inter-block gap - the area that is left unrecorded for this purpose. According to a feature of this invention the parity bytes will be recorded in this inter-block gap and will be read only if long-burst error recovery is required. Thus long-burst error protection is obtained without increasing the length of the medium required for the check bytes and the inter-block gap.

Assume now that, during reading, there are errors detected and that the pointer indicates a long-burst error. Sub-block check, block check and CRC syndromes are calculated as in the absence of a long-burst error. Also, the parity bytes are read and the parity byte syndromes are calculated as described above.

According to the disclosed arrangement, if the earliest byte at which the error could have started (and hence the earliest byte which could have been affected by the error) was at displacement Z (See Fig. 4A), then the following steps are performed:

(1) First level correction is applied, as necessary, to all sub-blocks which could not have been affected by the long-burst error; that is, sub-blocks which have no data or check bytes in the displacement range Z to Z+N. This correction is done as described in the cited prior art.

(2) Since the CRC and block check byte syndromes have been calculated from erroneous data, they need to be adjusted for each first level correction thus made. This adjustment is described in the cited prior art. Since the parity byte syndromes have also been calculated from erroneous data, they need to be adjusted for each first level correction thus made. This adjustment is done by XOR'ing the error pattern for each byte corrected into the parity byte syndrome for the parity group to which the corrected byte belongs.

(3) Next, long-burst errors are corrected as follows: The parity byte syndromes are overlaid on the data bytes and check bytes covering the N bytes following the byte at displacement Z. This range of displacements is referred to as the parity overlay range OR (Fig. 4A), and is N bytes in length. Each byte in this range is corrected by XORing the corresponding overlaid parity byte syndrome. Appropriate sub-block check byte, block check byte and CRC syndromes are adjusted for these changes in the same way that they were adjusted for changes resulting from first level error correction.

(4) Next, the long-burst-affected sub-blocks are processed through the two-level ECC algorithm. Note that any conventional errors in these sub-blocks outside the parity overlay range have contributed to the parity byte syndromes. When the parity byte syndromes were overlaid in step 3, these errors in the parity byte syndromes created errors in the parity overlay range OR as images

of the conventional errors. Figs. 4A-4D illustrate various error conditions under which imaged errors are correctable:

Fig. 4A shows a conventional error e in the affected sub-block which becomes a parity error in the corresponding parity group and becomes an image error e' in the long burst area in the overlay range OR. This error pair, the conventional error e and its image error e' in the same sub-block, is correctable by second level decoding in the two-level ECC.

Fig. 4B shows the case where the conventional error e and its image error e' appear in adjacent sub-blocks. In this case the error pair is corrected by the first-level processing of the two adjacent sub-blocks.

Fig. 4C shows the case where conventional errors e₁ and e₂ in the parity bytes appear as image errors e'₁ and e'₂ in the same sub-block. These errors are corrected by decoding the affected sub-block through second level correction.

Fig. 4D shows the case where conventional errors e₁ and e₂ in the parity bytes appear as image errors e'₁ and e'₂ in adjacent sub-blocks. These errors are corrected by decoding the affected sub-blocks through first level correction.

In operation, first-level correction as described above is applied iteratively to the long-burst-affected sub-blocks until each has been corrected or found non-correctable at first level. If one sub-block remains uncorrected after this processing, the block check bytes can be used and second level correction applied as described in the cited prior art.

(5) After all corrections are applied, the correction is verified from the condition of the adjusted CRC syndromes as described in the cited prior art.

The error-correcting capability of the algorithm in this invention can be generally summarised as follows:

A. When long-burst error correction is not required, the correction capability of the two-level code is unaffected.

B. When long-burst error correction is required, the errors can be corrected provided that all of the following conditions are satisfied:

(i) The long-burst is not longer than the parity byte overlay range.

(ii) The conventional errors in each sub-block are within the first level capability of the two-level ECC.

(iii) If the long-burst error affects more than one sub-block, only one of these long-burst-affected sub-blocks has conventional errors.

(iv) If the parity bytes are affected by conventional errors, the long-burst-affected sub-blocks do not have any conventional errors.

Note that N may exceed the length of one sub-block without impairing correction capability.

It should also be noted that the two-level ECC structure disclosed in the cited patents could be re-

structured to provide N-way (instead of two-way) interleaving of codewords; but, using this approach, the cost to enable correction of up to N consecutive bytes in error in a long-burst error would be prohibitive. With the disclosed arrangement, however, which involves appending the N parity bytes, up to N consecutive bytes in error can be corrected more cheaply and expeditiously; and the selected value of N can be increased without significant increase in complexity and cost.

While the invention has been described in connection with a preferred embodiment thereof, it will be understood that the foregoing and other changes may be made in the arrangement herein disclosed. The embodiment illustrated is therefore to be considered merely illustrative and the invention is not to be considered limited except as specified in the claims.

Claims

1. A method of correcting multiple byte errors in data retrieved from a magnetic medium on which data is recorded in variable length blocks that comprise sub-blocks (X, X+1) of data bytes (D) and, for each sub-block, check bytes (S) providing error correction code (ECC) for which ECC syndromes are generated during reading, the method comprising the steps of:
 - (a) writing at the end of each block a sequence of N sequential parity check bytes ($P_0 - P_{15}$) representing an interleaved parity check on data and ECC;
 - (b) during reading, after the generation of the ECC syndromes, generating parity syndromes from the data, ECC and parity check bytes as read;
 the method being characterised by the steps of:
 - (c) the parity check bytes being written at step (a) in the gaps between adjacent blocks;
 - (d) in response to a long-burst error condition, generating a pointer to the first (Z) of the N consecutive bytes in a block that are potentially in error due to the long-burst error;
 - (e) using any ECC syndromes which are not affected by the N bytes identified by the pointer, correcting correctable errors in all sub-blocks not affected by these identified bytes;
 - (f) adjusting the parity syndromes according to the errors thus corrected;
 - (g) using the adjusted parity syndromes, correcting the errors in the N bytes indicated by the pointer;
 - (h) adjusting any unused ECC syndromes according to the errors corrected by the adjusted parity syndromes; and
 - (i) using the adjusted ECC syndromes, correcting all correctable errors then remaining.
2. A method as claimed in claim 1, wherein if a long burst error is not detected the parity check bytes are not read, and the parity syndromes are not generated at step (b).
3. A method as claimed in either preceding claim, wherein the long-burst error condition does not exceed N bytes.
4. A method as claimed in any preceding claim, wherein variable length blocks comprise a variable number of data bytes (D) and a preselected number of check bytes (C).
5. A method as claimed in any preceding claim, wherein N exceeds the total number of bytes in a sub-block.
6. A method as claimed in any preceding claim, wherein the error correction code (ECC) for which ECC syndromes are generated during reading, can be used to correct up to a preselected number of errors in each sub-block during processing at a first level of error correction and additional errors in each block during processing at a second level of error correction, and the data bytes and ECC check bytes are interleaved, and the N sequential parity check bytes create N parity groups from interleaved data bytes and check bytes; the method comprising the steps of:
 - generating the parity syndromes for burst error correction during first level correction;
 - using the ECC syndromes which are not dependent on the N bytes identified by the pointer during first level correction;
 - after performing all possible first level correction and burst error correction, adjusting second level ECC syndromes for all errors thus far corrected; and
 - using the adjusted second level ECC syndromes to correct all correctable errors then remaining.
7. A method as claimed in claim 6, including, responsive to an error (e) in one of the sub-blocks that results in a parity error in the corresponding parity group that in turn results in an image error (e') in that parity overlay range N bytes in length as measured from the earliest byte (Z) which could have been affected by the error, correcting the error (e) and its image error (e') by decoding during second-level correction.
8. A method as claimed in either claim 6 or claim 7, including, responsive to an error (e) and its image error (e') in adjacent sub-blocks, correcting both errors by first level correction of the adjacent sub-blocks.



structured to provide N-way (instead of two-way) interleaving of codewords; but, using this approach, the cost to enable correction of up to N consecutive bytes in error in a long-burst error would be prohibitive. With the disclosed arrangement, however, which involves appending the N parity bytes, up to N consecutive bytes in error can be corrected more cheaply and expeditiously; and the selected value of N can be increased without significant increase in complexity and cost.

While the invention has been described in connection with a preferred embodiment thereof, it will be understood that the foregoing and other changes may be made in the arrangement herein disclosed. The embodiment illustrated is therefore to be considered merely illustrative and the invention is not to be considered limited except as specified in the claims.

Claims

1. A method of correcting multiple byte errors in data retrieved from a magnetic medium on which data is recorded in variable length blocks that comprise sub-blocks (X, X+1) of data bytes (D) and, for each sub-block, check bytes (S) providing error correction code (ECC) for which ECC syndromes are generated during reading, the method comprising the steps of:
 - (a) writing at the end of each block a sequence of N sequential parity check bytes ($P_0 - P_{16}$) representing an interleaved parity check on data and ECC;
 - (b) during reading, after the generation of the ECC syndromes, generating parity syndromes from the data, ECC and parity check bytes as read;
 the method being characterised by the steps of:
 - (c) the parity check bytes being written at step (a) in the gaps between adjacent blocks;
 - (d) in response to a long-burst error condition, generating a pointer to the first (Z) of the N consecutive bytes in a block that are potentially in error due to the long-burst error;
 - (e) using any ECC syndromes which are not affected by the N bytes identified by the pointer, correcting correctable errors in all sub-blocks not affected by these identified bytes;
 - (f) adjusting the parity syndromes according to the errors thus corrected;
 - (g) using the adjusted parity syndromes, correcting the errors in the N bytes indicated by the pointer;
 - (h) adjusting any unused ECC syndromes according to the errors corrected by the adjusted parity syndromes; and
 - (i) using the adjusted ECC syndromes, correcting all correctable errors then remaining.
2. A method as claimed in claim 1, wherein if a long burst error is not detected the parity check bytes are not read, and the parity syndromes are not generated at step (b).
3. A method as claimed in either preceding claim, wherein the long-burst error condition does not exceed N bytes.
4. A method as claimed in any preceding claim, wherein variable length blocks comprise a variable number of data bytes (D) and a preselected number of check bytes (C).
5. A method as claimed in any preceding claim, wherein N exceeds the total number of bytes in a sub-block.
6. A method as claimed in any preceding claim, wherein the error correction code (ECC) for which ECC syndromes are generated during reading, can be used to correct up to a preselected number of errors in each sub-block during processing at a first level of error correction and additional errors in each block during processing at a second level of error correction, and the data bytes and ECC check bytes are interleaved, and the N sequential parity check bytes create N parity groups from interleaved data bytes and check bytes; the method comprising the steps of:
 - generating the parity syndromes for burst error correction during first level correction;
 - using the ECC syndromes which are not dependent on the N bytes identified by the pointer during first level correction;
 - after performing all possible first level correction and burst error correction, adjusting second level ECC syndromes for all errors thus far corrected; and
 - using the adjusted second level ECC syndromes to correct all correctable errors then remaining.
7. A method as claimed in claim 6, including, responsive to an error (e) in one of the sub-blocks that results in a parity error in the corresponding parity group that in turn results in an image error (e') in that parity overlay range N bytes in length as measured from the earliest byte (Z) which could have been affected by the error, correcting the error (e) and its image error (e') by decoding during second-level correction.
8. A method as claimed in either claim 6 or claim 7, including, responsive to an error (e) and its image error (e') in adjacent sub-blocks, correcting both errors by first level correction of the adjacent sub-blocks.



9. A method as claimed in any of claims 6 to 8, including, responsive to two errors (e_1 , e_2) in the parity bytes that create corresponding image errors (e_1' , e_2') in adjacent sub-blocks, correcting the errors by first level correction of the adjacent sub-blocks.

10. A method as claimed in any of claims 6 to 9, including, responsive to two errors (e_1 , e_2) in the parity bytes that create corresponding image errors (e_1' , e_2') in the same sub-block, correcting the errors by decoding during second-level correction.

Patentansprüche

1. Eine Methode zur Korrektur von Mehr-Byte-Fehlern in Daten, die von einem magnetischen Medium wiedergewonnen werden, auf dem Daten in Blöcken von variabler Länge aufgezeichnet werden, welche Unterblöcke (X , $X+1$) von Datenbytes (D) enthalten und, für jeden Unterblock, Prüfbytes (S), die einen Fehlerkorrektur-Code (ECC) bereitstellen, für den während der Leseoperation ECC-Synndrome erzeugt werden, wobei die Methode folgende Schritte umfaßt:

(a) am Ende jedes Blocks, Schreiben einer Sequenz von N sequentiellen Paritäts-Prüfbytes ($P_0 - P_{16}$), die eine verschachtelte Paritäts-Prüfung der Daten und des ECC darstellen;

(b) während der Leseoperation, nach der Erzeugung der ECC-Synndrome, Erzeugen von Paritäts-Syndromen aus den gelesenen Daten-, ECC- und Paritäts-Prüfbytes;

wobei die Methode durch folgende Schritte gekennzeichnet ist:

(c) bei Schritt (a), Schreiben der Paritäts-Prüfbytes in die Zwischenräume zwischen benachbarten Blöcken;

(d) in Antwort auf eine Bedingung mit langer Fehlerkette, Erzeugen eines Zeigers zum ersten (Z) von N aufeinanderfolgenden Bytes in einem Block, die aufgrund der langen Fehlerkette potentiell fehlerhaft sind;

(e) unter Verwendung aller ECC-Synndrome, die nicht von den vom Zeiger identifizierten N Bytes betroffen sind, Korrigieren von korrigierbaren Fehlern in allen Unterblöcken, die nicht von diesen identifizierten Bytes betroffen sind;

(f) Anpassen der Paritäts-Synndrome entsprechend den so korrigierten Fehlern;

(g) unter Verwendung der angepaßten Paritäts-Synndrome, Korrigieren der Fehler in den N Bytes, die von dem Zeiger angezeigt wurden;

(h) Anpassen aller nicht verwendeten ECC-Synndrome entsprechend den Fehlern, die durch die angepaßten Paritäts-Synndrome korrigiert wurden; und

(i) unter Verwendung der angepaßten ECC-Synndrome, Korrigieren aller korrigierbaren noch verbleibenden Fehler.

2. Eine Methode nach Anspruch 1, bei der, wenn eine lange Fehlerkette nicht erkannt wird, die Paritäts-Prüfbytes nicht gelesen werden und die Paritäts-Synndrome bei Schritt (b) nicht erzeugt werden.

3. Eine Methode nach jedem der vorangehenden Ansprüche, bei der die Bedingung mit langer Fehlerkette N Bytes nicht überschreitet.

4. Eine Methode nach jedem der vorangehenden Ansprüche, bei der Blöcke von variabler Länge eine variable Anzahl von Datenbytes (D) und eine vorgewählte Anzahl von Prüfbytes (C) umfassen.

5. Eine Methode nach jedem der vorangehenden Ansprüche, bei der N die Gesamtzahl von Bytes in einem Unterblock überschreitet.

6. Eine Methode nach jedem der vorangehenden Ansprüche, bei der der Fehlerkorrektur-Code (ECC), für den während der Leseoperation ECC-Synndrome erzeugt werden, zur Korrektur einer vorgewählten Anzahl von Fehlern in jedem Unterblock verwendet werden kann, während der Verarbeitung auf einer ersten Fehlerkorrektur-Stufe, und zusätzliche Fehler in jedem Block während der Verarbeitung auf einer zweiten Fehlerkorrektur-Stufe, und die Datenbytes und ECC-Prüfbytes verschachtelt sind, und die N sequentiellen Paritäts-Prüfbytes N Paritäts-Gruppen aus verschachtelten Datenbytes und Prüfbytes erzeugen; die Methode umfaßt hierbei die folgenden Schritte:

Erzeugen der Paritäts-Synndrome zur Korrektur der Fehlerkette während der Korrektur auf der ersten Stufe;

Verwenden der ECC-Synndrome, die nicht von den N Bytes abhängen, die von dem Zeiger während der Korrektur auf der ersten Stufe identifiziert werden;

nach Durchführen aller möglichen Korrekturen auf der ersten Stufe und nach Korrektur der Fehlerkette, Anpassen von ECC-Syndromen der zweiten Stufe für alle bis dahin korrigierten Fehler; und

Verwenden der angepaßten ECC-Synndrome der zweiten Stufe zur Korrektur aller dann noch verbleibenden korrigierbaren Fehler.



7. Une Methode nach Anspruch 6, die, in Antwort auf einen Fehler (e) in einem der Unterblöcke, der einen Paritäts-Fehler in der entsprechenden Paritäts-Gruppe zur Folge hat, der wieder zu einem Abbildungsfehler (e') in diesem Paritäts-Überlagerungsbereich führt, mit einer Länge von N Bytes, gemessen ab dem frühesten Byte (Z), das von dem konventionellen Fehler hätte betroffen sein können, eine Korrektur des Fehlers (e) und seines Abbildungsfehlers (e') durch Decodierung während der Korrektur auf der zweiten Stufe beinhaltet. 5
8. Eine Methode nach jedem der Ansprüche 6 oder 7, die, in Antwort auf einen Fehler (e) und seinen Abbildungsfehler (e') in benachbarten Unterblöcken, die Korrektur beider Fehler durch Korrektur der benachbarten Unterblöcke auf der ersten Stufe beinhaltet. 10 15
9. Eine Methode nach jedem der Ansprüche 6 bis 8, die, in Antwort auf zwei Fehler (e₁, e₂) in den Paritäts-Bytes, die entsprechende Abbildungsfehler (e'₁, e'₂) in benachbarten Unterblöcken erzeugen, die Korrektur der Fehler durch Korrektur der benachbarten Unterblöcke auf der ersten Stufe beinhaltet. 20 25
10. Eine Methode nach jedem der Ansprüche 6 bis 9, die, in Antwort auf zwei Fehler (e₁, e₂) in den Paritäts-Bytes, die entsprechende Abbildungsfehler (e'₁, e'₂) in demselben Unterblock erzeugen, die Korrektur der Fehler durch Decodierung während der Korrektur auf der zweiten Stufe beinhaltet. 30 35

Revendications

1. Procédé de correction d'erreurs sur des multiplets multiples dans des données recherchées à partir d'un support magnétique sur lequel les données sont enregistrées en blocs de longueur variable qui comportent des sous-blocs (X, X+1) de multiplets (D) de données et, pour chaque sous-bloc, des multiplets de contrôle (S) fournissant un code de correction d'erreur (ECC) pour lequel des syndromes ECC sont générés au cours de la lecture, le procédé comportant les étapes suivantes: 40
 - (a) écriture à la fin de chaque bloc d'une séquence de N multiplets séquentiels de contrôle de parité (P0-P15) représentant un contrôle de parité d'imbrication sur les données et l'ECC; 45
 - (b) au cours de la lecture, après la génération des syndromes ECC, génération de syndromes de parité à partir des données, de l'ECC 50

et des multiplets de contrôle de parité tels que lus;

le procédé étant caractérisé par les étapes suivantes:

(c) les multiplets de contrôle de parité étant écrits au cours de l'étape (a) dans les intervalles entre des blocs adjacents;

(d) en réponse à une condition d'erreur à save longue, génération d'un pointeur au premier (Z) des N multiplets consécutifs dans un bloc qui sont potentiellement erronés à cause de l'erreur de save longue;

(e) utilisation de certains syndromes ECC qui ne sont pas affectés par les N multiplets identifiés par le pointeur, correction des erreurs qui peuvent être corrigées dans tous les sous-blocs qui ne sont pas affectés par ces multiplets identifiés;

(f) réglage des syndromes de parité conformément aux erreurs ainsi corrigées;

(g) utilisation des syndromes de parité réglés, correction des erreurs dans les N multiplets indiqués par le pointeur;

(h) réglage de certains syndromes ECC non utilisés conformément aux erreurs corrigées par les syndromes de parité réglés; et

(i) utilisation des syndromes ECC réglés, correction de toutes les erreurs qui peuvent être corrigées restant alors.

2. Procédé selon la revendication 1, dans lequel si une erreur de save longue n'est pas détectée les multiplets de contrôle de parité ne sont pas lus, et les syndromes de parité ne sont pas générés au cours de l'étape (b).
3. Procédé selon l'une quelconque des revendications précédentes, dans lequel la condition d'erreur à save longue n'excède pas N multiplets.
4. Procédé selon l'une quelconque des revendications précédentes, dans lequel les blocs à longueur variable comportent un nombre variable de multiplets de données (D) et un nombre présélectionné de multiplets de contrôle (C).
5. Procédé selon l'une quelconque des revendications précédentes, dans lequel N excède le nombre total de multiplets dans un sous-bloc.
6. Procédé selon l'une quelconque des revendications précédentes, dans lequel le code de correction d'erreur (ECC) pour lequel les syndromes ECC sont générés au cours de la lecture, peuvent être utilisés pour corriger jusqu'à un nombre d'erreurs présélectionné dans chaque sous-bloc au cours du traitement à un premier niveau de 55



- correction d'erreur et des erreurs supplémentaires dans chaque bloc au cours du traitement à un second niveau de correction d'erreur, et les multiplets de données et les multiplets de contrôle ECC sont imbriqués, et les N multiplets séquentiels de contrôle de parité créent N groupes de parité à partir des multiplets de données imbriqués et des multiplets de contrôle; le procédé comportant les étapes suivantes:
- 5 génération des syndromes de parité pour la correction d'erreurs en salve au cours du premier niveau de correction; 10
 - 15 utilisant des syndromes ECC qui ne dépendent pas des N multiplets identifiés par le pointeur au cours du premier niveau de correction;
 - 20 après avoir effectué toutes les corrections possibles de premier niveau et les corrections d'erreurs en salve, réglage des syndromes ECC du second niveau pour toutes les erreurs corrigées jusqu'à; et
 - 25 utilisation des syndromes ECC du second niveau réglé pour corriger toutes les erreurs qui peuvent être corrigées restant alors.
7. Procédé selon la revendication 6, comportant, sensible à une erreur (e) dans l'un des sous-blocs qui conduit à une erreur de parité dans le groupe de parité correspondant qui à son tour conduit à une erreur d'image (e') dans cette plage de parité N multiplets en longueur mesurés à partir du dernier multiplet (Z) qui pourrait avoir été affecté par l'erreur classique, correction de l'erreur (e) et de son erreur d'image (e') par décodage au cours de la correction de second niveau. 30 35
 8. Procédé selon l'une de la revendication 6 ou de la revendication 7, comportant, sensible à une erreur (e) et à son erreur d'image (e') dans des sous-blocs adjacents, la correction des deux erreurs par la correction de premier niveau des sous-blocs adjacents. 40
 9. Procédé selon l'une quelconque des revendications 6 à 8, comportant, sensible à deux erreurs (e1, e2) dans les multiplets de parité qui créent des erreurs images correspondantes (e1', e2') dans des sous-blocs adjacents, la correction des erreurs par la correction de premier niveau des sous-blocs adjacents. 45 50
 10. Procédé selon l'une quelconque des revendications 6 à 9, comportant, sensible aux deux erreurs (e1, e2) dans les multiplets de parité qui créent des erreurs images correspondantes (e1', e2') dans le même sous-bloc, la correction des erreurs par décodage au cours de la correction de second niveau. 55