

DATA COMPRESSION METHOD AND APPARATUS, AND COMPUTER DEVICE

TECHNICAL FIELD

[0001] The present specification relates to the field of data processing
5 technologies, and in particular, to a data compression method and apparatus, and a
computer device.

BACKGROUND

[0002] Data compression can be understood as representing more information
through less coding. During data compression, a proper compression coding rule
10 needs to be selected to obtain better compression efficiency. Currently, there are
increasingly more compression coding rules, and a relatively large computation cost
needs to be paid to select a proper compression coding rule.

SUMMARY

[0003] To alleviate a problem in related technologies, the present specification
15 provides a data compression method and apparatus, and a computer device.

[0004] According to a first aspect of the implementations of the present
specification, a data compression method is provided, including: obtaining a
to-be-compressed object; searching a recommendation record for whether a
recommended compression coding rule that meets a compression rate condition exists,
20 the recommendation record being used to record a compression coding rule of a
historical compressed object and corresponding compression rate information, and the
historical compressed object being of the same type as the to-be-compressed object; if
the recommended compression coding rule that meets the compression rate condition
exists, compressing the to-be-compressed object by using the recommended
25 compression coding rule; and if the recommended compression coding rule that meets
the compression rate condition does not exist, starting a regular compression coding
process to obtain estimated compression rates of a plurality of compression coding
rules for the to-be-compressed object, selecting a target compression coding rule

based on at least the estimated compression rates, and compressing the to-be-compressed object by using the target compression coding rule.

5 [0005] Optionally, the to-be-compressed object includes to-be-compressed data units obtained by splitting to-be-compressed data, and the historical compressed object includes other historical compressed data units obtained by splitting the to-be-compressed data.

10 [0006] Optionally, the to-be-compressed object includes to-be-compressed data units obtained by splitting to-be-compressed data in a data table, and the historical compressed object includes historical compressed data units corresponding to the data table.

[0007] Optionally, the to-be-compressed object includes a column of to-be-compressed data in a data table, and the historical compressed object includes historical compressed data in the same column as the column of to-be-compressed data in the data table.

15 [0008] Optionally, the to-be-compressed object includes a duplicate of a data table, and the historical compressed object includes a main data table corresponding to the duplicate of the data table.

20 [0009] Optionally, the compression rate condition includes at least: compression rate information of the recommended compression coding rule is higher than a specified threshold.

[0010] Optionally, the method further includes: after the compressing the to-be-compressed object, updating the recommendation record based on a compression coding rule used for the to-be-compressed object.

25 [0011] Optionally, the updating the recommendation record based on the compression coding rule used for the to-be-compressed object includes: obtaining an actual compression rate of the to-be-compressed object, and updating the recommendation record based on at least the actual compression rate and the compression coding rule used for the to-be-compressed object.

30 [0012] Optionally, the method further includes: obtaining access requirement information of the to-be-compressed object, where the access requirement information is related to decompression efficiency of the to-be-compressed object; and the compression rate condition includes: the compression rate information of the recommended compression coding rule matches the access requirement information.

[0013] Optionally, the selecting the target compression coding rule based on at

least the estimated compression rates includes: selecting the target compression coding rule based on the estimated compression rates and the access requirement information.

5 **[0014]** Optionally, the access requirement information of the to-be-compressed object is determined by obtaining historical access data of one or more of the to-be-compressed object or the historical compressed object.

[0015] Optionally, the historical access data includes a historical access frequency.

10 **[0016]** Optionally, the updating the recommendation record based on at least the actual compression rate and the compression coding rule used for the to-be-compressed object includes: updating the recommendation record based on the actual compression rate, the access requirement information, and the compression coding rule used for the to-be-compressed object.

15 **[0017]** Optionally, compression rate information of the compression coding rule includes a confidence level indicating the actual compression rate of the compression coding rule; and the updating the recommendation record based on the actual compression rate, the access requirement information, and the compression coding rule used for the to-be-compressed object includes: if the compression coding rule used for the to-be-compressed object is a recommended compression coding rule recorded in the recommendation record, and the actual compression rate of the to-be-compressed object matches the access requirement information, increasing a confidence level of the recommended compression coding rule; otherwise, reducing a confidence level of the recommended compression coding rule; and if the compression coding rule used for the to-be-compressed object is different from a recommended compression coding rule recorded in the recommendation record, and the actual compression rate of the to-be-compressed object matches the access requirement information, replacing the recommended compression coding rule in the recommendation record with the compression coding rule used for the to-be-compressed object.

30 **[0018]** According to a second aspect of the implementations of the present specification, a data compression apparatus is provided, the apparatus including: an acquisition module, configured to obtain a to-be-compressed object; a searching module, configured to search a recommendation record for whether a recommended compression coding rule that meets a compression rate condition exists, the

recommendation record being used to record a compression coding rule of a historical compressed object and corresponding compression rate information, and the historical compressed object being of the same type as the to-be-compressed object; a first compression module, configured to: if the recommended compression coding rule that
5 meets the compression rate condition exists, compress the to-be-compressed object by using the recommended compression coding rule; and a second compression module, configured to: if the recommended compression coding rule that meets the compression rate condition does not exist, start a regular compression coding process to obtain estimated compression rates of a plurality of compression coding rules for
10 the to-be-compressed object, select a target compression coding rule based on at least the estimated compression rates, and compress the to-be-compressed object by using the target compression coding rule.

[0019] According to a third aspect of the implementations of the present specification, a computer device is provided, including a memory, a processor, and a
15 computer program that is stored in the memory and that can run on the processor, the processor implementing the implementations of the above data compression method when executing the program.

[0020] The technical solutions provided in the implementations of the present specification can have the following beneficial effects:

[0021] In the implementations of the present specification, because the
20 compression coding rule of the historical compressed object and the corresponding compression rate information are recorded in the recommendation record, the recommendation record can be searched, based on the compression rate condition, for whether a proper compression coding rule exists. If no proper compression coding rule exists, the regular compression coding process is started to compress the
25 to-be-compressed object. If the proper compression coding rule exists, the compression coding rule can be directly used for compression. As such, there is no need to consume time and resources to compute a proper compression coding rule for a current to-be-compressed object, thereby significantly improving compression
30 efficiency.

[0022] It should be understood that the general descriptions above and the detailed descriptions below are merely examples for description, and cannot limit the present specification.

BRIEF DESCRIPTION OF DRAWINGS

[0023] The accompanying drawings, which are incorporated in and constitute a part of the present specification, illustrate implementations consistent with the present specification and, together with the specification, serve to explain the principles of the present specification.

[0024] FIG. 1 is a flowchart illustrating a data compression method according to an example implementation of the present specification;

[0025] FIG. 2 is a schematic diagram illustrating a data compression method according to an example implementation of the present specification;

10 [0026] FIG. 3 is a hardware structural diagram illustrating a computer device in which a data compression apparatus is located according to an example implementation of the present specification; and

[0027] FIG. 4 is a block diagram illustrating a data compression apparatus according to an example implementation of the present specification.

DETAILED DESCRIPTION

15 [0028] Embodiments are described in detail herein, examples of which are presented in the accompanying drawings. When the following description refers to the accompanying drawings, unless specified otherwise, same numbers in different accompanying drawings represent same or similar elements. Implementations described in the following example embodiments do not represent all implementations consistent with the present specification. On the contrary, the implementations are merely examples of apparatuses and methods that are described in the appended claims in details and consistent with some aspects of the present specification.

20 [0029] The terms used in the present specification are merely for illustrating specific implementations, and are not intended to limit the present specification. The terms "a" and "the" of singular forms used in the present specification and the appended claims are also intended to include plural forms, unless otherwise specified in the context clearly. It should be further understood that the term "and/or" used in the present specification indicates and includes any or all possible combinations of one or more associated listed items.

30 [0030] It should be understood that although terms "first", "second", "third", etc.

may be used in the present specification to describe various types of information, the information is not limited to the terms. These terms are merely used to differentiate between information of the same type. For example, without departing from the scope of the present specification, first information can also be referred to as second
5 information, and similarly, the second information can be referred to as the first information. Depending on the context, for example, the word "if" used herein can be explained as "while", "when", or "in response to determining".

[0031] For a relational database, data query processing and data storage are always two core elements. Generally, more central processing unit (CPU) and
10 memory resources need to be consumed for data query processing, and data storage mainly consumes resources of storage hardware such as a disk. However, development of the storage hardware currently does not keep pace with development of other system hardware. A space growth speed of the storage hardware always fails to keep up with a storage requirement that a user wants to store a wider range of data
15 for longer time. In addition, an access delay of the storage hardware currently always lags behind a CPU and a memory by at least one order of magnitude. In conclusion, the space and the access delay have become two major problems that hinder the rapid development of database storage.

[0032] To alleviate the above problems, all database manufacturers in the industry
20 actively introduce a general data compression technology, and compress data at a page/block level based on storage of each database. As such, not only occupation of storage space is greatly reduced, but also a storage access delay is improved because an amount of data input/output (IO) is decreased during scanning and query of a large range of data. However, a cost is consumption of more CPU and memory resources in
25 a compression process and a decompression process. From the perspective of the database as a whole, the introduction of the compression technology can make overall resource use more balanced, and more demanding service requirements of the user can be met under the same hardware condition.

[0033] Compression can be understood as representing more information through
30 less coding. A common compression method in row-column hybrid compression of the database is to compress different data by using different coding rules based on data distribution.

[0034] A theoretical limit of data compression is corresponding information entropy, that is, if a larger amount of data information is obtained, a final data

compression rate is higher. In the general data compression technology, to support more use scenarios, a sliding dictionary is generally established only for byte stream for coding compression, and a specific range of context information is used. However, the relational database naturally has more background knowledge for data in the database. For example, a type of a data column can be used to predict value range distribution of data, and data in the same column has stronger clustering. Therefore, increasingly more database manufacturers split data based on a compression unit, a column storage format is used for storage inside each compression unit, and a context relationship in a column and a data relationship between columns are used to find the most matching coding rule for compression, for example, dictionary coding, run-length coding, and incremental coding. Compared with general compression, these coding rules can provide higher compression rates by using embedded-in data information, and using embedded-in storage formats also enables the database to perform query directly based on coded data. In particular, for a database of a log-structured merge-tree (LSM tree) architecture, because stored data needs to be continuously compacted, and a compaction operation involves overall rewriting of the data, row-column hybrid coding compression can be naturally added to a data reshuffling process.

[0035] Although a higher compression rate can be brought when the database internally supports row-column hybrid coding compression, a complex problem of selecting and detecting a compression coding rule is accompanied. The database needs to traverse column data in all compression units to extract summary data, evaluate a final compression size based on different compression rules, and determine a finally selected coding rule. The entire coding rule detection process needs to consume a large amount of precomputation and matching comparison. In addition, currently, the compression coding rule is generally detected only for all data contexts in the same compression unit. Once a range is extended, detection complexity exponentially increases.

[0036] Coding compression of each database is basically implemented based on a row-column hybrid storage mode. A specific number of data record tuples are selected as an overall compression unit, and all data in the compression unit is stored by columns. The database performs traversal scanning on each column of data to obtain basic distribution and a data feature of the column of data, and then selects a proper rule for coding compression based on a specific data feature. Different databases

respectively implement different coding compression rules. Generally, dictionary coding, run-length coding, incremental coding, numeric coding, and partial open-source general compression coding are basically included. In the present implementation, a speed of selecting the compression coding rule is relatively slow.

5 Before determining a compression coding rule that is finally selected for a column of data, the database needs to undergo several stages: data scanning analysis, precomputation of compression rates of different coding rules, coding rule selection, and data compression. A relatively large computation cost needs to be paid in each stage.

10 **[0037]** In view of this, the implementations of the present specification provide a data compression solution. The solution in the implementations relates to a recommendation record. The recommendation record is used to record a compression coding rule of a historical compressed object and corresponding compression rate information. The historical compressed object is of the same type as a
15 to-be-compressed object. Therefore, in the implementations of the present specification, when the to-be-compressed object is compressed, the recommendation record can be first searched, based on a compression rate condition, for whether a proper compression coding rule exists. If no proper compression coding rule exists, a regular compression coding process is started to compress the to-be-compressed
20 object. If the proper compression coding rule exists, the compression coding rule can be directly used for compression. As such, there is no need to consume time and resources to compute a proper compression coding rule for a current to-be-compressed object, thereby significantly improving compression efficiency.

[0038] The implementations of the present specification are described in detail
25 below. FIG. 1 is a flowchart illustrating a method according to an example implementation of the present specification. The method includes the following steps.

[0039] Step 102: Obtain a to-be-compressed object.

[0040] Step 104: Search a recommendation record for whether a recommended compression coding rule that meets a compression rate condition exists, the
30 recommendation record being used to record a compression coding rule of a historical compressed object and corresponding compression rate information, and the historical compressed object being of the same type as the to-be-compressed object.

[0041] Step 106: If the recommended compression coding rule that meets the compression rate condition exists, compress the to-be-compressed object by using the

recommended compression coding rule.

[0042] Step 108: If the recommended compression coding rule that meets the compression rate condition does not exist, start a regular compression coding process to obtain estimated compression rates of a plurality of compression coding rules for the to-be-compressed object, select a target compression coding rule based on at least the estimated compression rates, and compress the to-be-compressed object by using the target compression coding rule.

[0043] The present implementation is applicable to various data compression scenarios, and the to-be-compressed object can be various types of objects such as a data table, a video file, an audio file, or an image. No limitation is imposed in the present implementation. The historical compressed object is an object that is of the same type as the to-be-compressed object and that has been compressed before the to-be-compressed object is compressed. For different to-be-compressed objects, that the historical compressed object is of the same type as the to-be-compressed object can have a plurality of implementations. For example, the to-be-compressed object and the historical compressed object can be data that is generated at different times and that is of different versions, for example, a program source file. Alternatively, the to-be-compressed object and the historical compressed object can be data in the same format, for example, data in a video format. Alternatively, the to-be-compressed object and the historical compressed object can be child data that belongs to the same parent data, for example, a plurality of pieces of child data obtained by splitting original data, or a plurality of data units that belong to the same data table.

[0044] For example, assume that to-be-compressed object Data-A1 currently needs to be compressed. In the first time of compression, no historical compressed object that is of the same type as the to-be-compressed object can be used for reference, the recommendation record is empty, and no recommended compression coding rule that meets the compression rate condition exists. As such, the regular compression coding process is started to obtain the estimated compression rates of the plurality of compression coding rules for the to-be-compressed object, the target compression coding rule is selected based on at least the estimated compression rates, and the to-be-compressed object is compressed by using the selected target compression coding rule.

[0045] After the first time of compression is completed, the recommendation record is updated based on a compression coding rule used for to-be-compressed

object Data-A1. For example, the current compression can be recorded in the recommendation record. Specifically, the target compression coding rule and compression rate information of the target compression coding rule can be written in the recommendation record. In some examples, the estimated compression rate of the compression coding rule on the to-be-compressed object is not much different from an actual compression rate, and the compression rate information can be information related to the estimated compression rate. In some other examples, the estimated compression rate of the compression coding rule on the to-be-compressed object may be different from an actual compression rate to a degree. The actual compression rate may be lower than the estimated compression rate or higher than the estimated compression rate. Therefore, in the present implementation, the actual compression rate of the to-be-compressed object can be further obtained, and the recommendation record is updated based on at least the actual compression rate and the compression coding rule used for the to-be-compressed object. In this case, the compression rate information can be information related to the actual compression rate. The actual compression rate can be obtained through computation after Data-A1 is compressed, for example, the actual compression rate is determined based on a ratio of an actual size of the to-be-compressed object after compression to an original size of the to-be-compressed object. The compression rate information can be directly a value of a compression rate, or can be related information obtained by converting a compression rate by using a specified method as required.

[0046] Subsequently, because new data is inserted into Data-A1 or for another reason, new Data-A2 is generated based on Data-A1. In the current compression process, Data-A2 is used as a new to-be-compressed object, and Data-A2 is of the same type as Data-A1. For a process of using Data-A1 as the historical compressed object, references can be made to the compression process of Data-A1.

[0047] In a compression process of Data-A2, a compression coding rule can be selected with reference to compression of Data-A1. For example, the compression coding rule for compressing Data-A1 and corresponding compression rate information are recorded in the recommendation record, and it can be determined, based on the compression rate condition, whether to select the compression coding rule of Data-A1 for coding. If the compression coding rule of Data-A1 is selected, Data-A2 can be directly compressed by using the compression coding rule. The compression rate condition in the present implementation can be flexibly configured as required. For

example, some to-be-compressed objects have a specific requirement on the compression rate, and different compression coding rules achieve different compression effects for the compressed object. In the present implementation, the compression rate condition can be set to select a proper recommended compression coding rule. For example, the compression rate condition can include: compression rate information of the recommended compression coding rule is higher than a specified threshold. A specific threshold can be flexibly configured based on actual needs. No limitation is imposed in the present implementation.

[0048] It can be understood that, in the compression process of Data-A2, because the compression coding rule of the historical compressed object can be selected for compression, there is no need to consume time and resources to compute a proper compression coding rule for a current to-be-compressed object, thereby significantly improving compression efficiency.

[0049] Certainly, if the compression coding rule of Data-A1 does not meet the compression rate condition of Data-A2, the regular compression coding process can also be stated for Data-A2 to obtain the estimated compression rates of the plurality of compression coding rules for the to-be-compressed object, the target compression coding rule is selected based on at least the estimated compression rates, and the to-be-compressed object is compressed by using the target compression coding rule.

[0050] After Data-A2 is compressed, the recommendation record can be further updated based on compression of Data-A2. For example, in an actual service, the compression coding rule of original Data-A1 may not be applicable to Data-A2 because new data is inserted into Data-A1. Therefore, in the present implementation, an actual compression rate of Data-A2 can be obtained, and the recommendation record is updated based on at least the actual compression rate and the compression coding rule used for Data-A2. For example, the recommendation record is updated to the compression coding rule used for Data-A2, and the compression rate information is updated based on the actual compression rate of Data-A2.

[0051] In an actual compression service scenario, a problem of a large amount of to-be-compressed data may exist. Therefore, the to-be-compressed data can be split into a plurality of to-be-compressed data units as required. In the service scenario, the data compression method in the present implementation can be used to better improve compression efficiency. For example, the to-be-compressed object can include to-be-compressed data units obtained by splitting to-be-compressed data, and the

historical compressed object includes other historical compressed data units obtained by splitting the to-be-compressed data.

[0052] For example, as shown in FIG. 2, Data-B can be split into a total of four to-be-compressed data units from Data-b1 to Data-b4.

5 **[0053]** For example, for Data-b1, in the first time of compression, no historical compressed object can be used for reference, a recommendation record is empty, and no recommended compression coding rule that meets a compression rate condition exists. As such, a regular compression coding process is started to obtain estimated compression rates of a plurality of compression coding rules for the to-be-compressed
10 object, a target compression coding rule is selected based on at least the estimated compression rates, and the to-be-compressed object is compressed by using the selected target compression coding rule.

[0054] After the first time of compression is completed, the current compression can be recorded in the recommendation record. For example, the target compression
15 coding rule and compression rate information of the target compression coding rule are recorded in the recommendation record. Subsequently, in a compression process of Data-b2, a compression coding rule can be selected with reference to compression of Data-b1. For example, the compression coding rule for compressing Data-b1 and corresponding compression rate information are recorded in the recommendation
20 record, and it can be determined, based on the compression rate condition, whether to select the compression coding rule of Data-b1 for coding. If the compression coding rule of Data-b1 is selected, Data-b2 can be directly compressed by using the compression coding rule. After Data-b2 is compressed, it is determined, based on compression, whether to update the recommendation record. Subsequently, for
25 Data-b3 and Data-b4, the same process can also be used to implement fast compression and update the recommendation record.

[0055] In some other examples, a compression scenario of a data table is used as an example. In a distributed relational database of an LSM tree architecture, etc., a data table may undergo compaction a plurality of times, and new data obtained each
30 time of compaction needs to be compressed. In the solution in the present implementation of the present specification, there is no need to repeatedly estimate compression rates of a plurality of compression coding rules, and the recommendation record can be constructed by using a compression knowledge of the historical compressed object. In subsequent other compression processes, a compression coding

rule can be quickly obtained based on the recommendation record.

[0056] For example, to-be-compressed data in tableC currently needs to be compressed. Because there is a large amount of to-be-compressed data in tableC, the to-be-compressed data in tableC is split into a total of five to-be-compressed data units from table-c1 to table-c5 based on a specified data unit size.

[0057] For table-c1, in the first time of compression, no historical compressed object can be used for reference, and a recommendation record is empty. As such, a regular compression coding process is started to obtain estimated compression rates of a plurality of compression coding rules for the to-be-compressed object, a target compression coding rule is selected based on at least the estimated compression rates, and the to-be-compressed object is compressed by using the selected target compression coding rule.

[0058] After the first time of compression is completed, the current compression can be recorded in the recommendation record. For example, the target compression coding rule and compression rate information of the target compression coding rule are recorded in the recommendation record. Subsequently, in a compression process of table-c2, a compression coding rule can be selected with reference to compression of table-c1. For example, the compression coding rule for compressing table-c1 and corresponding compression rate information are recorded in the recommendation record, and it can be determined, based on the compression rate condition, whether to select the compression coding rule of table-c1 for coding. If the compression coding rule of table-c1 is selected, table-c2 can be directly compressed by using the compression coding rule. After table-c2 is compressed, it is determined, based on compression, whether to update the recommendation record. Subsequently, for table-c3 and table-c4, the same process can also be used to implement fast compression and update the recommendation record.

[0059] In some other examples, a data table generally includes a plurality of columns of data having different attributes, the columns of data may vary greatly, and different compression coding rules can be used for different columns of data in the data table. For example, in a user data table, a column of data about a user age is integer data, and a column of data about a user name is string data. A proper compression coding rule can be selected for each of the two columns of data. Based on this, the to-be-compressed object in the present implementation can include a column of to-be-compressed data in a data table, and the historical compressed object

can include historical compressed data in the same column as the column of to-be-compressed data in the data table. Based on this, in the present implementation, column data can be distinguished. A compression coding rule recorded in the recommendation record corresponds to the column data. After a certain column of data in the data table is compressed, when the column of data in the data table subsequently needs to be compressed after the column of data is updated, the column of data can be compressed based on a compression coding rule of historical compressed data in the same column in the data table, thereby improving compression efficiency.

[0060] In some other examples, the to-be-compressed object includes a duplicate of a data table, and the historical compressed object includes a main data table corresponding to the duplicate of the data table.

[0061] Similar to the examples above, the duplicate of the data table is totally the same content as the main data table, and the duplicate of the data table is consistent with the main data table. Generally, after the main data table is updated, the duplicate of the data table is updated based on the update operation of the main data table. If the main data table involves compression storage after update, compression storage also needs to be performed for the duplicate of the data table after update. Based on this, in a compression process of the duplicate of the data table, the duplicate of the data table can be directly compressed by using a compression coding rule of the main data table. Optionally, when the duplicate of the data table is compressed, the recommendation record is searched for the compression coding rule of the main data table and corresponding compression rate information, and the compression coding rule of the main data table is directly used for compression. Certainly, in an actual service, an actual compression rate of the compression coding rule selected for the main data table may not be high. When the recommendation record is searched, it can be determined, as required, that no recommended compression coding rule that meets the compression rate condition exists in the recommendation record, to start a regular compression coding process to compress the duplicate of the data table.

[0062] In an actual service, an evaluation standard of the compression coding rule is single. The compression coding rule is mainly selected based on the final actual compression rate, without comprehensively considering additional overheads brought by different compression coding rules for different data access modes in case of actual database service load. For example, assume that a data compression rate is relatively

high. In this case, a corresponding data decompression process consumes more time because a compression rate is negatively correlated with decompression efficiency. If data needs to be accessed frequently, a data compression rate is relatively high, and the data needs to be decompressed in each time of access, relatively large additional decompression overheads are inevitable. Based on this, in the present implementation, access requirement information of the to-be-compressed object is obtained. The access requirement information can be related to decompression efficiency. For example, a higher access requirement indicates higher corresponding decompression efficiency. The access requirement information can be configured by a service party, can be configured by a data access party, or can be determined by collecting historical access data of one or more of the to-be-compressed object or the historical compressed object. For example, assume that the historical compressed object is frequently accessed. In this case, a data access requirement is relatively high, and a corresponding decompression efficiency requirement is relatively high. Based on this, a relatively high access requirement can be set. If the historical compressed object is rarely accessed, a data access requirement is relatively low, and a corresponding decompression efficiency requirement is relatively low. Based on this, a relatively low access requirement can be set. Optionally, historical access information of the to-be-compressed object can be obtained based on the historical compressed data of the same type as the to-be-compressed object. Optionally, the historical access information can include an operation type of the data, for example, adding, deleting, modifying, or query, or can be the number of operations performed on the data within a specific time period, or can be a historical access frequency.

[0063] Correspondingly, the compression rate condition includes: the compression rate information of the recommended compression coding rule matches the access requirement information, so that a proper compression coding rule can be selected based on a compression rate requirement and the access requirement information, instead of only a single factor such as the compression rate, and final compressed data can further meet the access requirement of the service.

[0064] Based on this, when the access requirement information is obtained, the recommendation record can be further updated based on the actual compression rate, the access requirement information, and the compression coding rule used for the to-be-compressed object. For example, in an initial stage of compression, the compression coding rule used for the to-be-compressed object starts to be written into

in the recommendation record. Subsequently, as content of the to-be-compressed data changes or access of the to-be-compressed data changes, the compression coding rule used for the to-be-compressed object may subsequently change. Based on this, the recommendation record needs to be updated dynamically. Optionally, in the present
5 implementation, the compression rate information of the compression coding rule is implemented by using a confidence level, and the confidence level represents the actual compression rate of the compression coding rule, so that reliability of the compression coding rule can be dynamically adjusted in a continuous compression process. Optionally, the updating the compression coding rule and the corresponding
10 compression rate information in the recommendation record based on the actual compression rate and the access requirement information in the present implementation can include: if the compression coding rule used for the to-be-compressed object is a recommended compression coding rule recorded in the recommendation record, and the actual compression rate of the to-be-compressed
15 object matches the access requirement information, increasing a confidence level of the recommended compression coding rule; otherwise, reducing a confidence level of the recommended compression coding rule; or if the compression coding rule used for the to-be-compressed object is different from a recommended compression coding rule recorded in the recommendation record, and the actual compression rate of the
20 to-be-compressed object matches the access requirement information, replacing the recommended compression coding rule in the recommendation record with the compression coding rule used for the to-be-compressed object.

[0065] It can be learned from the above implementation that when the to-be-compressed object is compressed by using the recommended compression
25 coding rule recorded in the recommendation record, after compression ends, whether the recommended compression coding rule is proper is checked by using the actual compression rate and the access requirement information. If the actual compression rate of the to-be-compressed object matches the access requirement information, the confidence level of the recommended compression coding rule is increased, so that
30 the compression coding rule can be continued to use in subsequent compression. Otherwise, the confidence level of the recommended compression coding rule is reduced, so that when the confidence level of the recommended compression coding rule is lower than a threshold, the regular compression coding process can be started to obtain a new compression coding rule. If there is no proper compression coding

rule in the recommendation record, the regular compression coding process is started to obtain a new compression coding rule. The new compression coding rule is different from the recommended compression coding rule recorded in the recommendation record. Similarly, after compression ends, whether the recommended
5 compression coding rule is proper is checked by using the actual compression rate and the access requirement information. If the actual compression rate of the to-be-compressed object matches the access requirement information, the recommended compression coding rule in the recommendation record can be replaced with the compression coding rule used for the to-be-compressed object, so that a
10 newest proper compression coding rule can be recorded in the recommendation record for use in subsequent compression. If the actual compression rate of the to-be-compressed object does not match the access requirement information, the compression coding rule is not applicable, and the recommendation record may not be updated.

15 **[0066]** The following describes the data compression method in the present specification again by using an implementation. In a distributed database of an LSM tree architecture, etc., relatively more data compaction requirements exist. A relational database of the LSM tree architecture is used as an example. The LSM tree has a multi-layer structure similar to a tree structure where a data size at an upper layer is
20 smaller than a data size at a lower layer. First, a C0 layer resident in memory stores all recently written data tuple records. The memory structure is orderly, can implement in-place update at any time, and can be queried at any time. The remaining C1 to Cn layers are all resident on disk.

[0067] A data write process is as follows: When a data write operation comes, the
25 data write operation is first added to a write-ahead log (that is, a log recorded before actual writing), and then added to the C0 layer. When data at the C0 layer reaches a specific size, the C0 layer and the C1 layer are compacted. This is similar to merge sort. The process is compaction. Data at a new-C1 layer obtained through compaction is sequentially written into the disk to replace the original old-C1 layer. When the C1
30 layer reaches a specific size, compaction with a lower layer continues. After compaction, all old files can be deleted, and new files are left.

[0068] It can be learned that when compaction is performed, data needs to be reshuffled and rewritten into a new storage file, and the new storage file needs to be compressed.

[0069] During compaction and writing, data is split into independent compression units based on a fixed size specified when the user creates a table, and row-column hybrid coding compression is performed on data in each compression unit.

5 **[0070]** In addition, in a distributed database scenario, each data table has a plurality of data duplicates. Compaction times of the main data table and the duplicate of the data table need to be staggered to alternately provide a service to ensure service access availability and stability of the user.

10 **[0071]** For example, a statistical information module configured in a database can obtain access data of a home table. An access requirement of the data table can be quickly determined based on the access data. For example, a larger number of data point query operations and a larger number of data table access times indicate a higher access delay requirement of the data table based on an operation type of historical access and a historical access frequency. Storage resources saved by compressing the data cannot compensate for additional computing resource overheads, and a
15 compression coding rule with a relatively low compression rate needs to be used, to ensure an access speed of the user. A lower number of data point query operations and a lower number of data table access times indicate that the user does not perform query much and has a lower access delay requirement. In this case, a compression rule with a relatively high compression rate can be used.

20 **[0072]** In a compaction process of a data table, a large number of data units need to be compressed, and a recommendation record that includes a current best compression coding rule is maintained. After a compression coding rule is selected for the to-be-compressed object for compression each time, compression of the compression coding rule is updated to the recommendation record, and a newest
25 compression coding rule in a current state is maintained in the recommendation record.

[0073] Before each to-be-compressed data unit is compressed, the data table accepts two pieces of input information: access requirement information of the data table and a newest recommendation record.

30 **[0074]** When the recommendation record is empty or a confidence level of the compression coding rule in the recommendation record does not reach a specified threshold, a regular compression coding process is started to detect a to-be-compressed data unit. Data feature analysis is performed on each column of data by using the regular compression coding process. Compression rates are sorted based

on estimated compression rates of different compression coding rules. A compression coding rule that meets the access requirement of the data table is finally selected, as a compression coding rule used for current compression, from the plurality of compression coding rules with different estimated rates based on the access requirement information of data table. The to-be-compressed data unit may include a plurality of columns of different types of data. A corresponding compression coding rule can be selected for each column of data, that is, the to-be-compressed data unit can correspond to two or more compression coding rules.

[0075] When a confidence level of the compression coding rule in the recommendation record exceeds a specified threshold, the compression coding rule in the recommendation record can be directly used as a compression coding rule used for a current to-be-compressed data unit. The to-be-compressed data unit may include a plurality of columns of different types of data, and a corresponding compression coding rule can be selected for each column of data. In some examples, compression coding rules can be selected for some data columns from the recommendation record, and some data columns cannot. In this case, a regular compression coding process can be started for detection and selection.

[0076] After one or more compression coding rules required by the current to-be-compressed data unit are selected, data in the to-be-compressed data unit is scanned by columns, and starts to be compressed by using the selected compression coding rules. After compression of each data column is completed, actual compression rate information of each data column is obtained.

[0077] After overall compression is completed for the data unit, the recommendation record is updated based on the actual compression rate information of each data column:

[0078] if the compression coding rule used for the to-be-compressed object is a recommended compression coding rule recorded in the recommendation record, and the actual compression rate of the to-be-compressed object matches the access requirement information, increasing a confidence level of the recommended compression coding rule; otherwise, reducing a confidence level of the recommended compression coding rule; or

[0079] if the compression coding rule used for the to-be-compressed object is different from a recommended compression coding rule recorded in the recommendation record, and the actual compression rate of the to-be-compressed

object matches the access requirement information, replacing the recommended compression coding rule in the recommendation record with the compression coding rule used for the to-be-compressed object.

5 **[0080]** For a compression process of the duplicate of the data table, data compression can be performed on the duplicate of the data table directly by using a compression coding rule of the main data table having the same data.

10 **[0081]** It can be learned from the above implementation that, in the data compression method in the present implementation of the present specification, a semi-supervised learning method is used, so that more historical context information is fully used to adjust the compression coding rule in the recommendation record timely.

15 **[0082]** For the main data table, in a compaction process, the regular compression coding process needs to be started for a to-be-compressed data unit in an earlier stage of compression to detect each compression coding rule, to obtain a newest recommendation record through training. After the confidence level exceeds the specified threshold, a subsequent to-be-compressed data unit can start to be compressed directly by using the compression coding rule in the recommendation record, and complex rule detection is not required. Because the data in the duplicate of the data table is exactly the same as that in the main data table, the duplicate of the data table can be directly compressed by using a final recommended coding rule of the main data table, so that detection overheads of all compression coding rules can be eliminated.

20 **[0083]** In the present implementation, during selection of a compression coding rule corresponding to each to-be-compressed data unit, an access requirement of a data table to which the to-be-compressed data unit belongs is optimized, and selection of a final compression coding rule is adjusted based on the access delay requirement. In addition, selection of the compression coding rule corresponding to each to-be-compressed data unit is not merely based on data in the current unit. Instead, an optimal compression coding rule is comprehensively considered by using actual compression rate results of more compression units.

30 **[0084]** The above implementation proposes the data compression method for the distributed relational database of the LSM tree architecture. A semi-supervised multi-level feedback mechanism is implemented in conjunction with the row-column hybrid compression coding rule of the database. Specific selection of a compression

coding rule of each data unit is corrected by using the access requirement information of the data table, and a proper compression rule can be continuously learned in the compression process to update the recommendation record, thereby accelerating coding rule selection in the entire compaction process. For the data table, a detection speed of the compression coding rule can be greatly improved based on another data duplicate in a distributed database cluster and a local historical compression data knowledge, and a proper compression coding rule is selected adaptively based on an access requirement feature of data, so that a compression rate is improved while a data access speed is improved as much as possible.

[0085] In particular, for the duplicate of the data table, the recommendation record of the main data table can be directly used, so that a rule recommendation and training process that consumes computing resources can be skipped. Generally, more context information can be fully used to accelerate and optimize detection of the compression coding rule.

[0086] Corresponding to the implementations of the data compression method, the present specification further provides implementations of a data compression apparatus and a computer device to which the apparatus is applied.

[0087] The implementation of the data compression apparatus in the present specification can be applied to the computer device, for example, a server or a terminal device. The apparatus implementation can be implemented by software, hardware, or a combination of hardware and software. Software implementation is used as an example. As a logical apparatus, the apparatus is formed by reading a corresponding computer program instruction in a non-volatile memory to a memory by a processor of file processing where the apparatus is located. In terms of hardware, FIG. 3 is a hardware structural diagram illustrating a computer device where the data compression apparatus is located in the present specification. In addition to a processor 310, a memory 330, a network interface 320, and a non-volatile memory 340 shown in FIG. 3, a server or electronic device where the apparatus 331 is located in the implementation can usually include other hardware based on an actual function of the computer device. Details are omitted herein for simplicity.

[0088] FIG. 4 is a block diagram illustrating a data compression apparatus according to an example implementation of the present specification. The apparatus includes:

[0089] an acquisition module 41, configured to obtain a to-be-compressed object;

[0090] a searching module 42, configured to search a recommendation record for whether a recommended compression coding rule that meets a compression rate condition exists, the recommendation record being used to record a compression coding rule of a historical compressed object and corresponding compression rate information, and the historical compressed object being of the same type as the to-be-compressed object;

[0091] a first compression module 43, configured to: if the recommended compression coding rule that meets the compression rate condition exists, compress the to-be-compressed object by using the recommended compression coding rule; and

[0092] a second compression module 44, configured to: if the recommended compression coding rule that meets the compression rate condition does not exist, start a regular compression coding process to obtain estimated compression rates of a plurality of compression coding rules for the to-be-compressed object, select a target compression coding rule based on at least the estimated compression rates, and compress the to-be-compressed object by using the target compression coding rule.

[0093] Optionally, the to-be-compressed object includes to-be-compressed data units obtained by splitting to-be-compressed data, and the historical compressed object includes other historical compressed data units obtained by splitting the to-be-compressed data.

[0094] Optionally, the to-be-compressed object includes to-be-compressed data units obtained by splitting to-be-compressed data in a data table, and the historical compressed object includes historical compressed data units corresponding to the data table.

[0095] Optionally, the to-be-compressed object includes a column of to-be-compressed data in a data table, and the historical compressed object includes historical compressed data in the same column as the column of to-be-compressed data in the data table.

[0096] Optionally, the to-be-compressed object includes a duplicate of a data table, and the historical compressed object includes a main data table corresponding to the duplicate of the data table.

[0097] Optionally, the compression rate condition includes at least: compression rate information of the recommended compression coding rule is higher than a specified threshold.

[0098] Optionally, the apparatus further includes an update module, configured to:

after the to-be-compressed object is compressed, update the recommendation record based on a compression coding rule used for the to-be-compressed object.

[0099] Optionally, the update module is further configured to:

5 **[0100]** obtain an actual compression rate of the to-be-compressed object, and update the recommendation record based on at least the actual compression rate and the compression coding rule used for the to-be-compressed object.

[0101] Optionally, the acquisition module is further configured to obtain access requirement information of the to-be-compressed object, where the access requirement information is related to decompression efficiency of the
10 to-be-compressed object; and

[0102] the compression rate condition includes: the compression rate information of the recommended compression coding rule matches the access requirement information.

[0103] Optionally, the second compression module is further configured to:

15 **[0104]** select the target compression coding rule based on the estimated compression rates and the access requirement information.

[0105] Optionally, the access requirement information of the to-be-compressed object is determined by obtaining historical access data of one or more of the to-be-compressed object or the historical compressed object.

20 **[0106]** Optionally, the historical access data includes a historical access frequency.

[0107] Optionally, the update module is further configured to:

[0108] update the recommendation record based on the actual compression rate, the access requirement information, and the compression coding rule used for the to-be-compressed object.

25 **[0109]** Optionally, compression rate information of the compression coding rule includes a confidence level indicating the actual compression rate of the compression coding rule; and

[0110] the update module is further configured to:

[0111] if the compression coding rule used for the to-be-compressed object is a
30 recommended compression coding rule recorded in the recommendation record, and the actual compression rate of the to-be-compressed object matches the access requirement information, increase a confidence level of the recommended compression coding rule; otherwise, reduce a confidence level of the recommended compression coding rule; or

[0112] if the compression coding rule used for the to-be-compressed object is different from a recommended compression coding rule recorded in the recommendation record, and the actual compression rate of the to-be-compressed object matches the access requirement information, replace the recommended compression coding rule in the recommendation record with the compression coding rule used for the to-be-compressed object.

[0113] Correspondingly, the present specification further provides a computer device, including a memory, a processor, and a computer program that is stored in the memory and that can run on the processor, the processor implementing the implementation of the data compression method when executing the program.

[0114] For an implementation process of functions and roles of each module in the data compression apparatus, references can be made to an implementation process of a corresponding step in the data compression method. Details are omitted herein for simplicity.

[0115] Because the implementation of the data compression apparatus corresponds to the implementation of the data compression method, for related parts, references can be made to related descriptions in the method implementation. The apparatus implementation describe above is merely an example. The modules described as separate parts may or may not be physically separate, and parts displayed as modules may or may not be physical modules, may be located in one position, or may be distributed on a plurality of network modules. Some or all of the modules can be selected based on actual needs to achieve the objectives of the solutions of the present specification. A person of ordinary skill in the art can understand and implement the implementations of the present specification without creative efforts.

[0116] Specific implementations of the present specification are described above. Other implementations fall within the scope of the appended claims. In some situations, the actions or steps described in the claims can be performed in an order different from the order in the implementation and the desired results can still be achieved. In addition, the process depicted in the accompanying drawings does not necessarily require a particular execution order to achieve the desired results. In some implementations, multi-tasking and parallel processing can be advantageous.

[0117] A person skilled in the art can easily figure out another implementation solution of the present specification after thinking over the specification and practicing the present specification herein. The present specification is intended to

cover any variation, use, or adaptation of the present specification, and these variations, uses, or adaptations follow the general principles of the present specification and include common knowledge or conventional techniques that are not disclosed in the technical field of the present specification. The specification and the
5 implementations are merely considered as examples, and the actual scope and the spirit of the present specification are pointed out by the following claims.

[0118] It should be understood that the present specification is not limited to the precise structures that have been described above and shown in the drawings, and various modifications and changes can be made without departing from the scope of
10 the present specification. The scope of the present specification is limited by the appended claims only.

[0119] The descriptions above are merely preferred implementations of the present specification, but are not intended to limit the present specification. Any modification, equivalent replacement, or improvement made without departing from
15 the spirit and principle of the present specification shall fall within the protection scope of the present specification.

CLAIMS

What is claimed is:

1. A data compression method, comprising:
 - obtaining a to-be-compressed object;
 - 5 searching a recommendation record for whether a recommended compression coding rule that meets a compression rate condition exists, the recommendation record being used to record a compression coding rule of a historical compressed object and corresponding compression rate information, and the historical compressed object being of the same type as the to-be-compressed object;
 - 10 if the recommended compression coding rule that meets the compression rate condition exists, compressing the to-be-compressed object by using the recommended compression coding rule; and
 - if the recommended compression coding rule that meets the compression rate condition does not exist, starting a regular compression coding process to obtain
 - 15 estimated compression rates of a plurality of compression coding rules for the to-be-compressed object, selecting a target compression coding rule based on at least the estimated compression rates, and compressing the to-be-compressed object by using the target compression coding rule.
2. The method according to claim 1, wherein the to-be-compressed object
- 20 includes to-be-compressed data units obtained by splitting to-be-compressed data, and the historical compressed object includes other historical compressed data units obtained by splitting the to-be-compressed data.
3. The method according to claim 1, wherein the to-be-compressed object
- includes to-be-compressed data units obtained by splitting to-be-compressed data in a
- 25 data table, and the historical compressed object includes historical compressed data units corresponding to the data table.
4. The method according to claim 1, wherein the to-be-compressed object
- includes a column of to-be-compressed data in a data table, and the historical compressed object includes historical compressed data in the same column as the
- 30 column of to-be-compressed data in the data table.
5. The method according to claim 1, wherein the to-be-compressed object
- includes a duplicate of a data table, and the historical compressed object includes a

main data table corresponding to the duplicate of the data table.

6. The method according to claim 1, wherein the compression rate condition includes at least: compression rate information of the recommended compression coding rule is higher than a specified threshold.

5 7. The method according to claim 1, further comprising:

after the compressing the to-be-compressed object, updating the recommendation record based on a compression coding rule used for the to-be-compressed object.

8. The method according to claim 7, wherein the updating the recommendation record based on the compression coding rule used for the to-be-compressed object
10 includes:

obtaining an actual compression rate of the to-be-compressed object, and updating the recommendation record based on at least the actual compression rate and the compression coding rule used for the to-be-compressed object.

9. The method according to claim 1, 6, or 8, wherein the method further
15 comprises: obtaining access requirement information of the to-be-compressed object, wherein the access requirement information is related to decompression efficiency of the to-be-compressed object; and

the compression rate condition includes: the compression rate information of the recommended compression coding rule matches the access requirement information.

20 10. The method according to claim 9, wherein the selecting the target compression coding rule based on at least the estimated compression rates includes:

selecting the target compression coding rule based on the estimated compression rates and the access requirement information.

11. The method according to claim 9, wherein the access requirement
25 information of the to-be-compressed object is determined by obtaining historical access data of one or more of the to-be-compressed object or the historical compressed object.

12. The method according to claim 11, wherein the historical access data includes a historical access frequency.

30 13. The method according to claim 9, wherein the updating the recommendation record based on at least the actual compression rate and the compression coding rule used for the to-be-compressed object includes:

updating the recommendation record based on the actual compression rate, the access requirement information, and the compression coding rule used for the

to-be-compressed object.

14. The method according to claim 13, wherein compression rate information of the compression coding rule includes a confidence level indicating the actual compression rate of the compression coding rule; and

5 the updating the recommendation record based on the actual compression rate, the access requirement information, and the compression coding rule used for the to-be-compressed object includes:

10 if the compression coding rule used for the to-be-compressed object is a recommended compression coding rule recorded in the recommendation record, and the actual compression rate of the to-be-compressed object matches the access requirement information, increasing a confidence level of the recommended compression coding rule; otherwise, reducing a confidence level of the recommended compression coding rule; and

15 if the compression coding rule used for the to-be-compressed object is different from a recommended compression coding rule recorded in the recommendation record, and the actual compression rate of the to-be-compressed object matches the access requirement information, replacing the recommended compression coding rule in the recommendation record with the compression coding rule used for the to-be-compressed object.

20 15. A data compression apparatus, the apparatus comprising:

an acquisition module, configured to obtain a to-be-compressed object;

25 a searching module, configured to search a recommendation record for whether a recommended compression coding rule that meets a compression rate condition exists, the recommendation record being used to record a compression coding rule of a historical compressed object and corresponding compression rate information, and the historical compressed object being of the same type as the to-be-compressed object;

a first compression module, configured to: if the recommended compression coding rule that meets the compression rate condition exists, compress the to-be-compressed object by using the recommended compression coding rule; and

30 a second compression module, configured to: if the recommended compression coding rule that meets the compression rate condition does not exist, start a regular compression coding process to obtain estimated compression rates of a plurality of compression coding rules for the to-be-compressed object, select a target compression coding rule based on at least the estimated compression rates, and compress the

to-be-compressed object by using the target compression coding rule.

16. A computer device, comprising a memory, a processor, and a computer program that is stored in the memory and that can run on the processor, the processor implementing the method according to any one of claims 1 to 14 when executing the

5 program.