

DATABASE ACCESS METHODS AND APPARATUSES

TECHNICAL FIELD

[0001] One or more implementations of the present specification relate to the field of computer technologies, and in particular, to database access methods and apparatuses.

BACKGROUND

[0002] To respond to a user's database access request, after receiving the user's database access request, a server usually formulates a corresponding access policy and solution for the access request. The access policy and solution here can be referred to as an execution plan. Then, the server executes the formulated execution plan to access the database, and returns a desired result to the user.

[0003] It is worthwhile to note that, in ideal situations, a corresponding execution plan can be formulated for each access request. However, in the process of formulating an execution plan, cumbersome operations such as statistics calculation and sampling usually need to be performed on data in the database. In other words, the formulation process is time-consuming, usually causing the user's access time to be excessively long. Therefore, to optimize the user's access time, the execution plan is currently formulated in two methods. The first method is to perform the process of formulating an execution plan once for the same type of access request. In other words, access requests of the same type share one execution plan. The second method is to formulate different execution plans based on different parameters that are included in access statements corresponding to the access requests. In short, different execution plans are formulated by using different methods, causing different impact on database access.

SUMMARY

[0004] One or more implementations of the present specification describe database access methods and apparatuses, to improve efficiency and accuracy of a user's access to a database.

[0005] According to a first aspect, a database access method is provided, including the following: receiving a database access statement; parsing the database access statement to determine a target user parameter included in the database access statement; determining a statement type of the database access statement; obtaining a
5 pre-stored execution plan group corresponding to the database access statement of the statement type, where execution plans in the execution plan group are distinguished by using different user parameters; selecting a target execution plan from the execution plan group based on the target user parameter; and executing the target execution plan.

10 [0006] According to a second aspect, a database access apparatus is provided, including the following: a receiving unit, configured to receive a database access statement; a parsing unit, configured to parse the database access statement received by the receiving unit, to determine a target user parameter included in the database access statement; a determining unit, configured to determine a statement type of the
15 database access statement; an acquisition unit, configured to obtain a pre-stored execution plan group corresponding to the database access statement of the statement type determined by the determining unit, where execution plans in the execution plan group are distinguished by using different user parameters; a selection unit, configured to select a target execution plan from the execution plan group obtained by
20 the acquisition unit, based on the target user parameter; and an execution unit, configured to execute the target execution plan selected by the selection unit.

[0007] One or more implementations of the present specification provide database access methods and apparatuses. The database access methods and apparatuses include the following: receiving a database access statement; parsing the database
25 access statement to determine a target user parameter included in the database access statement; determining a statement type of the database access statement; obtaining a pre-stored execution plan group corresponding to the database access statement of the statement type; selecting a target execution plan from the obtained execution plan group based on the target user parameter; and executing the target execution plan. In
30 other words, in the implementations of the present specification, database access statements of the same statement type share one execution plan group, and execution plans in the same execution plan group are distinguished by using different user parameters, thus improving accuracy of a user's access to a database. In addition, the quantity of formulated execution plans can be significantly reduced. When the

quantity of execution plans is reduced, efficiency of the user's access to the database can be improved, and computer resources can be reduced.

BRIEF DESCRIPTION OF DRAWINGS

- [0008] To describe the technical solutions in the implementations of the present specification more clearly, the following briefly describes the accompanying drawings needed for describing the implementations. Clearly, the accompanying drawings in the following description show merely some implementations of the present specification, and a person of ordinary skill in the art can still derive other drawings from these accompanying drawings without creative efforts.
- 5 [0009] FIG. 1 is a schematic diagram illustrating an application scenario of a database access method, according to the present specification;
- [0010] FIG. 2 is a flowchart illustrating a method for determining an object type group, according to the present specification;
- [0011] FIG. 3 is a flowchart illustrating a database access method, according to an implementation of the present specification;
- 15 [0012] FIG. 4 is a schematic diagram illustrating a plan space exploration process, according to the present specification;
- [0013] FIG. 5 is a flowchart illustrating a database access method, according to another implementation of the present specification; and
- 20 [0014] FIG. 6 is a schematic diagram illustrating a database access apparatus, according to an implementation of the present specification.

DESCRIPTION OF IMPLEMENTATIONS

- [0015] The following describes the solutions provided in the present specification with reference to the accompanying drawings.
- 25 [0016] Database access methods provided in the present specification can be applied to a scenario shown in FIG. 1. In FIG. 1, a client can send a database access request to a server. The database access request can be implemented by using a database access statement (such as Structured Query Language (SQL)). The server in FIG. 1 can include a plan cache. The plan cache can be storage space obtained by
- 30 dividing a memory of the server, and is used to store an execution plan generated by

the server. The execution plan here can be generated by the server based on the previously described database access statement. In addition, the server can access the database. The server accesses the database by executing an execution plan.

5 [0017] The database in FIG. 1 can include multiple objects. The object can be an access entity that may exist in a database access statement, such as a table, a view, etc. In the present specification, for example, the object is a table.

[0018] It is worthwhile to note that, before the database access method provided in the present specification is performed, the objects in the database can be grouped to determine multiple object type groups. In some implementations, a method for determining the multiple object type groups can be shown in FIG. 2. In FIG. 2, the determining method can include the following:

[0019] Step S210: Select multiple objects with the same object pattern from a database.

15 [0020] A table is used as an example for description. The multiple objects with the same object pattern can be multiple tables having the same column data type and the same index pattern. The same index pattern can mean that the quantity of indexes created for multiple tables are the same, and the columns for which indexes are created are the same.

20 [0021] For example, for two tables user_01 and user_02 created by using the following SQL statements, it is considered that the two tables have the same object pattern.

[0022] create table user_01 (id integer primary key, gmt_create date, balance_01 number(10,2) index user_01_gmt_blanche on (gmt_create, blance_01));

25 [0023] create table user_02 (id integer primary key, gmt_create date, balance_02 number(10,2) index user_02_gmt_blanche on (gmt_create,blance_02));

[0024] It can be understood from the previous SQL statements that user_01 includes three columns in total, and the data types of the three columns are "integer", "date", and "number (10,2)" respectively. In addition, an index is created for user_01, and a column for which the index is created is column 3. Similarly, it can be understood that user_02 includes three columns in total, and the data types of the three columns are "integer", "date", and "number (10,2)" respectively. In addition, an index is created for user_02, and a column for which the index is created is column 3. Therefore, it can be considered that user_01 and user_02 are two tables having the same object pattern.

[0025] Step S220: Group the multiple objects.

[0026] It is worthwhile to note that, in the present specification, the grouping of objects is a progressive process. For example, for the multiple objects selected in step S210, any two objects can be first grouped, that is, a similarity between the two objects is calculated. When the calculated similarity is less than a predetermined threshold d , the two objects are categorized into one group. When the calculated similarity is not less than a predetermined threshold d , the two objects are categorized into two groups. Then, a third object of the multiple objects is grouped. For example, when the previously described two objects are categorized into one group, a similarity between the third object and one of the two objects can be calculated. When the calculated similarity is less than the predetermined threshold d , the third object is also categorized into the previous one group. If the calculated similarity is not less than the predetermined threshold d , a similarity between the third object and the other object is calculated. If the calculated similarity is less than the predetermined threshold d , the third object is also categorized into the previous one group. Otherwise, the third object is categorized into another group.

[0027] When the previously described two objects are categorized into two groups, a similarity between the third object and one of the two objects can be calculated. If the calculated similarity is less than the predetermined threshold d , the third object is categorized into the group in which the one object is located. If the calculated similarity is not less than the predetermined threshold d , a similarity between the third object and the other object is calculated. If the calculated similarity is less than the predetermined threshold d , the third object is categorized into the group in which the other object is located. Otherwise, the third object is categorized into another group. By analogy, the process is repeated until the previously described multiple objects are all categorized into corresponding groups.

[0028] In summary, the object grouping method provided in the present specification is performed in turn for each object. To be specific, when an object is to be grouped, the object can be successively compared with multiple objects in multiple existing groups in terms of similarity (i.e., the similarity is calculated). When a similarity between the object and a certain object in any group satisfies a need, the comparison is stopped, and the object is categorized into the group. Otherwise, the comparison continues. If no object "is similar to" the object in the end (i.e., the similarity is less than the predetermined threshold d), the object is categorized into a

separate group, and then the similarity comparison process continues for the other objects.

[0029] In some implementations, the similarity can be determined by calculating the Euclidean distance between two objects. For example, the similarity between two tables is calculated. Three attributes can be predetermined: the quantity of rows r , the quantity of unique values ndv , and an average row length rl . Then, when the similarity between any two tables is to be calculated, attribute values of the previously described three attributes of the two tables can be first determined. Finally, the similarity between the two tables is calculated based on the attribute values of the three attributes of the two tables. For example, a formula for calculating the similarity between any two tables can be as follows:

$$similarity = \sqrt{(r_1 - r_2)^2 + (ndv_1 - ndv_2)^2 + (rl_1 - rl_2)^2} \quad (\text{Formula 1})$$

[0030] In the formula, similarity represents the similarity between the two tables; r_1 , ndv_1 and rl_1 represent the quantity of rows, the quantity of unique values and the average row length of one table, respectively; r_2 , ndv_2 and rl_2 represent the quantity of rows, the quantity of unique values and the average row length of the other table, respectively.

[0031] Certainly, in practice, other attributes such as the quantity of columns can also be defined, which is not limited in the present specification.

[0032] It can be understood that, step S210 and step S220 can be repeated until each object in the database is categorized into a corresponding group. After each object in the database is categorized into a corresponding group, the groups obtained by categorization can be used as multiple object type groups.

[0033] Step S230: Establish a mapping relationship between each object in the database and the multiple object type groups.

[0034] For example, the previously described mapping relationship can be shown in Table 1.

Table 1

Object	Object type group
user_01	group_01
user_02	group_01
dept_01	group_02

...	...
-----	-----

[0035] In Table 1, user_01 and user_02 belong to the same object type group (group_01). In other words, object patterns of user_01 and user_02 are the same and a similarity between user_01 and user_02 is less than a predetermined threshold d.

[0036] In the present specification, objects having the same object pattern and having a similarity less than the predetermined threshold d are categorized into the same object type group. As such, in the case of database division and table division, multiple physical sub-tables of the same logical table can be aggregated, thus significantly reducing the quantity of generated execution plans, and reducing computer resources.

[0037] It is worthwhile to note that, for the previously described multiple object type groups, it can be periodically and automatically detected whether object patterns of objects in each object type group are the same. When the object patterns are the same, for each object, similarities between the object and other objects in the object type group are calculated. When none of the similarities between the object and other objects is less than a predetermined threshold, an object type group that the object belongs to is re-determined. Thus, accuracy of each object group in the database is ensured, and accuracy of database access can be further improved. In addition, the process of automatically detecting accuracy of an object group can reduce dependence on a user and intrusion, thus increasing stability of a database system.

[0038] It is worthwhile to further note that, the policy of determining the object type group based on the object pattern and the similarity in the present specification can alleviate the problem of poor stability and usability of the database caused by designating the object type group by the user.

[0039] It can be understood that, the database access method provided in the present specification can be performed after the objects in the database are categorized into corresponding object type groups.

[0040] FIG. 3 is a flowchart illustrating a database access method, according to an implementation of the present specification. The method can be performed by a device having a processing capability, i.e., a server, a system, or an apparatus, such as a server in FIG. 1. As shown in FIG. 3, the method can include the following:

[0041] Step S310: Receive a database access statement.

[0042] Here, the database access statement sent by a client can be received by a

server. The database access statement can be an SQL statement. For example, the database access statement can be as follows: select * from user_01 a, dept_01 b where a.id = b.id and a.id = 1.

5 [0043] Step S320: Parse the database access statement to determine a target user parameter included in the database access statement.

[0044] When the database access statement is an SQL statement, the SQL statement can be parsed based on a syntax of the SQL statement. The previously described target user parameter can be a fixed value in the SQL statement. The target user parameter can include an integer, a floating-point number, a string, etc.

10 [0045] In the previous example, the target user parameter included in the SQL statement is 1.

[0046] Step S330: Determine a statement type of the database access statement.

[0047] In some implementations, the database access statement can be first parameterized. Then, the statement type is determined based on the parameterized database access statement. The step of parameterizing the database access statement can include the following:

[0048] Step a: Extract a corresponding object set from the database access statement.

20 [0049] An object in the object set can be an access entity that may exist in the database access statement, such as a table, a view, etc. When the database access statement is an SQL statement, the SQL statement can be parsed based on a syntax of the SQL statement to extract the previously described object set.

[0050] For example, the database access statement is as follows: select * from user_01 a, dept_01 b where a.id = b.id and a.id = 1. The extracted object set can be as follows: O={ user_01,dept_01 }.

[0051] Step b: For an object in the object set, determine an object type group that the object belongs to, and obtain a corresponding group name.

30 [0052] When multiple object type groups are predetermined, and a mapping relationship between each object in the database and multiple object type groups is established, the object type group that the object in the object set belongs to can be determined based on the mapping relationship. In the previous example, based on the content shown in Table 1, the object type group that user_01 belongs to can be determined as group_01, and the object type group that dept_01 belongs to can be determined as group_02.

[0053] Step c: Replace the object in the database access statement with the corresponding group name, and replace the target user parameter with a wildcard.

[0054] The wildcard here can refer to any one or more characters that are different from other characters or strings in the database access statement. The character can be a special character such as "@", "?", or "%". For example, the wildcard is "?". In the previous example, user_01 is replaced with group_01, dept_01 is replaced with group_02, and 1 is replaced with "?". Then the SQL statement is converted into the following string: select * from group_01, group_02 where group_01.id = group_02.id and group_01.id = ?.

[0055] It can be understood that, the previously described string is a statement type of the database access statement.

[0056] Step S340: Obtain a pre-stored execution plan group corresponding to the database access statement of the statement type.

[0057] In some implementations, the execution plan group can be obtained from a plan cache of the server. The plan cache can pre-store multiple execution plan groups corresponding to different statement types. Each execution plan group is identified by using a corresponding hash value, and execution plans in the execution plan group are distinguished by using different user parameters.

[0058] For example, content of the plan cache can be shown in Table 2.

Table 2

Hash value	Statement type	Execution plan group		
XXXX	select * from group_01, group_02 where group_01.id = group_02.id and group_01.id = ?	Group 1	plan1	1
			plan2	2
			plan3	3
YYYY	select * from group_01, group_03 where group_01.id = group_03.id and group_01.id = ?	Group 2	plan4	≤10
			plan5	>10
ZZZZ	select * from group_01, group_02 where group_01.id = group_02.id and group_01.age = ?	Group 3	plan6	≤18
			plan7	(18,50)
			plan8	≥50
...	...	...		

[0059] Certainly, in practice, the previously described plan cache does not have to include the statement type, or can further include other content, which is not limited in the present specification.

[0060] After the statement type of the database access statement is determined, a hash value corresponding to the statement type can be calculated based on a predetermined hash function. Then, the corresponding execution plan group is obtained based on the calculated hash value.

[0061] It is worthwhile to note that, the execution plan group in the plan cache can be continuously generated and stored by the server in response to a user's database access request. The process of generating the previously described execution plan group can also be referred to as a plan space exploration process. The plan space exploration process can be shown in FIG. 4.

[0062] In FIG. 4, p1 and p2 are two different points in the parameter space, and are used to represent two different user parameters. When receiving a certain statement type of a database access statement that includes p1, the server can generate a corresponding execution plan based on p1. The process of generating the execution plan can be illustrated as dividing an area in the parameter space shown in FIG. 4, and the area obtained by division can also be referred to as a plan space. Then, when receiving the statement type of the database access statement that includes p2, the server can calculate a distance Δp between p2 and p1 (for example, an Euclidean distance). Assuming that Δp is greater than a distance threshold D, the corresponding execution plan can be generated based on p2, that is, another plan space is divided in the parameter space shown in FIG. 4. It is worthwhile to note that, a new execution plan is no longer generated when the distance Δp between p_i and p1 (or p2) is less than D, where p_i is a user parameter included in the database access statement of the statement type that the server receives at one time. In other words, p_i and p1 (or p2) share an execution plan. It can be understood that, as user parameters increase, the whole parameter space will be divided into multiple different plan spaces gradually. When some plan spaces overlap, two plan spaces are combined into one. Finally, the whole parameter space will be divided, and the plan space exploration process ends. It can be understood that, after the plan space exploration process ends, the multiple execution plans generated form one execution plan group.

[0063] Certainly, in practice, p1 and p2 can also be two different numerical ranges in the parameter space, which is omitted here for simplicity in the present

specification.

[0064] It is worthwhile to note that, FIG. 4 shows only the plan space exploration process in a one-dimensional space. When the database access statement includes multiple different user parameters at the same time, such as n different user parameters, the parameter space forms an n-dimensional space. The plan space exploration process is similar in the n-dimensional space. For example, n different user parameters can be represented as an array (a1, a2, a3, ..., an). Then during division of the plan space, assume that a corresponding plan space has been divided for a certain point (b1, b2, b3, ..., bn) in the parameter space. Then a distance Δp between the array and the point can be calculated based on the following formula.

$$\Delta p = \sqrt{(a_1 - b_1)^2 + (a_2 - b_2)^2 + (a_3 - b_3)^2 + \dots + (a_n - b_n)^2} \quad (\text{Formula 2})$$

[0065] Then, whether a new plan space is divided can be determined based on a magnitude relationship between the distance threshold and the distance Δp calculated based on formula 2. The determining method is the same as the previous description, and details are omitted here for simplicity.

[0066] It is worthwhile to further note that, FIG. 4 shows only the process of generating one execution plan group, that is, the database access statements received by the server have the same statement type. When receiving multiple different statement types of database access statements, the server can generate a corresponding execution plan group for each statement type. As such, multiple execution plan groups can be obtained. For details, references can be made to the content shown in Table 2.

[0067] In the present specification, the method for generating execution plan in the plan space exploration process can significantly improve the quality of the execution plan, thus further improving the efficiency of the user's access to the database. In addition, in the present specification, the method for generating the execution plan group for a type of database access statement can prevent a large amount of repeated calculations, thus increasing the speed and efficiency of generating the execution plan. In addition, the method can also significantly reduce the quantity of generated execution plans, thus reducing the occupation of space and reducing the computer resources.

[0068] Step S350: Select a target execution plan from the execution plan group based on the target user parameter.

[0069] Optionally, to improve the quality of the selected execution plan, before step

S350 is performed, it can be determined again whether the object type group that each object in the object set belongs to has changed. Step S350 is performed only if there is no change.

[0070] The previously described determining process can be as follows: for each
5 object in the object set, it is determined whether an object pattern of the object is the same as an object pattern of each object in the object type group that the object belongs to. If the object patterns are the same, a similarity between the object and each object is successively calculated. For example, assume that the object set is {user_01, dept_01}, and a predetermined mapping relationship between an object in
10 the database and an object type group is shown in Table 1. Then, it can be determined whether object patterns of user_01 and user_02 are the same. If the object patterns are the same, a similarity between the two objects can be calculated. In addition, it can be further determined whether object patterns of dept_01 and other objects in group_02 are the same. If the object patterns are the same, similarities between dept_01 and
15 other objects in group_02 can be successively calculated. For the calculation of the similarity, references can be made to formula 1, and details are omitted here for simplicity. When a similarity between each object and any object in the object type group that each object belongs to is less than a predetermined threshold, it is considered that the object type group that each object belongs to has not changed. It is
20 worthwhile to note that, when the object patterns are different, or when the object patterns are the same but none of the similarities between the object and other objects in the object type group is less than the predetermined threshold, the object type group can be re-determined. In the previous example, assume that the object patterns of user_01 and user_02 are different, or that the object patterns are the same, but the
25 similarity between them is not less than the predetermined threshold, and that the object patterns of dept_01 and other objects in group_02 are the same and the similarity between dept_01 and a certain object in the group is less than the predetermined threshold. In such case, the object type group that user_01 belongs to can be re-determined.

30 **[0071]** It can be understood that, when the object type group that a certain object or some objects belong to are re-determined, step S330 and step S340 can be performed again.

[0072] It is worthwhile to note that, for the execution plan group in step S350, when the previously described plan space exploration is finished, the execution plan

group includes all the execution plans. As such, the corresponding target execution plan can be selected based on the target user parameter. In other words, an Euclidean distance between the target user parameter and a user parameter corresponding to each execution plan in the execution plan group is calculated. When an Euclidean distance
5 between the target user parameter and a user parameter corresponding to any execution plan is less than a distance threshold, the any execution plan is used as the target execution plan.

[0073] When the previously described plan space exploration has not been finished, the process of selecting the target execution plan is the same as the previous
10 description. The difference is that because the plan space exploration has not been finished, the target execution plan may not be selected. In other words, there may be a case in which none of Euclidean distances between the target user parameter and user parameters corresponding to various generated execution plans in the execution plan group is less than a distance threshold. In such case, the previously described plan
15 space exploration process can continue. In other words, a corresponding plan space is divided in the parameter space shown in FIG. 4, that is, a new execution plan is generated based on the target user parameter. As such, the new execution plan is added to the execution plan group.

[0074] Step 360: Execute the target execution plan.

20 **[0075]** Here, if a target execution plan is selected, the target execution plan is executed. If no target execution plan is selected, the new execution plan is executed.

[0076] In summary, the solution provided in the implementations of the present specification can improve accuracy of the user's access to the database, and improve access efficiency.

25 **[0077]** FIG. 5 is a flowchart illustrating a database access method, according to another implementation of the present specification. As shown in FIG. 5, the method can include the following:

[0078] Step S510: Receive a database access statement.

[0079] Step S520: Parse the database access statement to determine a target user
30 parameter included in the database access statement.

[0080] Step S530: Extract a corresponding object set from the database access statement.

[0081] Step S540: For an object in the object set, determine an object type group that the object belongs to.

- [0082] Step S550: Parameterize the database access statement based on the determined object type group.
- [0083] Step S560: Determine the statement type of the database access statement based on the parameterized database access statement.
- 5 [0084] Step S570: Obtain a pre-stored execution plan group corresponding to the database access statement of the statement type.
- [0085] Step S580: For the object in the object set, determine whether an object pattern of the object is the same as an object pattern of each object in the object type group that the object belongs to. If no, step S590 is performed. If yes, step S5100 is performed.
- 10 [0086] Step S590: Re-determine the object type group that the object belongs to, and return to step S550.
- [0087] Step S5100: Successively calculate a similarity between the object and each object.
- 15 [0088] Step S5110: Determine whether a similarity between the object and any object is less than a predetermined threshold. If no, step S590 is performed. If yes, step S5120 is performed.
- [0089] Step S5120: Select a target execution plan from the execution plan group based on the target user parameter.
- 20 [0090] Step S5130: Determine whether the target execution plan is selected. If yes, step S5140 is performed. If no, step S5150 is performed.
- [0091] Step S5140: Execute the target execution plan.
- [0092] Step S5150: Generate a new execution plan based on the target user parameter.
- 25 [0093] Step S5160: Execute the new execution plan and add the new execution plan to the execution plan group.
- [0094] In the previously described implementations of the present specification, the object type group that the object in the database access statement belongs to is first determined, and then the execution plan is selected or generated, so as to aggregate multiple objects of the same object type (for example, aggregate multiple physical sub-tables of the same logical table), thus sharing the execution plan.
- 30 [0095] Corresponding to the previously described database access method, an implementation of the present specification further provides a database access apparatus. As shown in FIG. 6, the apparatus includes the following: a receiving unit

601, configured to receive a database access statement; a parsing unit 602, configured to parse the database access statement received by the receiving unit 601, to determine a target user parameter included in the database access statement; a determining unit 603, configured to determine a statement type of the database access statement; the
5 determining unit 603 can be configured to: parameterize the database access statement; and determine the statement type based on the parameterized database access statement, where the parameterizing the database access statement can include the following: extracting a corresponding object set from the database access statement; for an object in the object set, determining an object type group that the object
10 belongs to, and obtaining a corresponding group name; and replacing the object in the database access statement with the corresponding group name, and replacing the target user parameter with a wildcard; an acquisition unit 604, configured to obtain a pre-stored execution plan group corresponding to the database access statement of the statement type determined by the determining unit 603, where execution plans in the
15 execution plan group are distinguished by using different user parameters; a selection unit 605, configured to select a target execution plan from the execution plan group obtained by the acquisition unit, based on the target user parameter; and an execution unit 606, configured to execute the target execution plan selected by the selection unit 605.

20 **[0096]** Optionally, the apparatus can further include a grouping unit 607 and an establishment unit 608.

[0097] The selection unit 605 is further configured to select multiple objects with the same object pattern from a database.

[0098] The grouping unit 607 is configured to group the multiple objects selected
25 by the selection unit 605.

[0099] The execution unit 606 is further configured to repeat the previously described selection and grouping steps until all objects in the database are categorized into corresponding groups.

[0100] The establishment unit 608 is configured to use the multiple groups
30 obtained by the grouping unit 607 as the multiple object type groups, and establish a mapping relationship between each object in the database and the multiple object type groups.

[0101] The determining unit 603 can be further configured to: determine the object type group that the object belongs to, based on the mapping relationship.

[0102] Optionally, the selection unit 605 can be configured to: for the object in the object set, determine whether an object pattern of the object is the same as an object pattern of each object in the object type group that the object belongs to; if the object patterns are the same, successively calculate a similarity between the object and each
5 object; and when a similarity between the object and any object is less than a predetermined threshold, select the target execution plan from the execution plan group based on the target user parameter; or when no similarity between the object and each object is less than a predetermined threshold, re-determine an object type group that the object belongs to, and obtain a corresponding target execution plan
10 based on the re-determined object type group.

[0103] Optionally, the apparatus can further include the following: a calculation unit 609, configured to: for each object in each object type group, periodically and successively calculate similarities between the object and other objects in the object type group; and the determining unit 603 is further configured to: when none of the
15 similarities between the object and other objects is less than a predetermined threshold, re-determine an object type group that the object belongs to.

[0104] Optionally, the selection unit 605 can be further configured to: calculate an Euclidean distance between the target user parameter and a user parameter corresponding to each execution plan in the execution plan group; and when an
20 Euclidean distance between the target user parameter and a user parameter corresponding to any execution plan is less than a distance threshold, use the any execution plan as the target execution plan.

[0105] Optionally, the apparatus can further include the following: a generation unit 6010, configured to generate a new execution plan based on the target user parameter; and an adding unit 6011, configured to add the new execution plan
25 generated by the generation unit 6010 to the execution plan group, and execute the new execution plan.

[0106] The functions of the functional modules of the apparatus in the previous implementation of the present specification can be implemented by using the steps of
30 the previous method implementation. Therefore, a specific working process of the apparatus provided in an implementation of the present specification is omitted here for simplicity.

[0107] In the database access apparatus provided in some implementations of the present specification, the receiving unit 601 is configured to receive a database access

statement; the parsing unit 602 is configured to parse the database access statement to determine a target user parameter included in the database access statement; the determining unit 603 is configured to determine a statement type of the database access statement; the acquisition unit 604 is configured to obtain a pre-stored
5 execution plan group corresponding to the database access statement of the determined statement type; the selection unit 605 is configured to select a target execution plan from the execution plan group based on the target user parameter; and the execution unit 606 is configured to execute the target execution plan. Thus, accuracy of the user's access to the database and access efficiency can be improved.

10 **[0108]** It is worthwhile to note that, the database access apparatus provided in the implementations of the present specification can be a module or a unit of the server in FIG. 1.

[0109] A person skilled in the art should be aware that in the previously described one or more examples, functions described in the present specification can be
15 implemented by hardware, software, firmware, or any combination thereof. When being implemented by using software, these functions can be stored in a computer-readable medium or transmitted as one or more instructions or code in the computer-readable medium.

[0110] The objectives, technical solutions, and beneficial effects of the present
20 specification are further described in detail in the previously described specific implementations. It should be understood that the previous descriptions are merely specific implementations of the present specification, but are not intended to limit the protection scope of the present specification. Any modification, equivalent replacement, or improvement made based on the technical solutions of the present
25 specification shall fall within the protection scope of the present specification.

CLAIMS

What is claimed is:

1. A database access method, comprising:

receiving (S310) a database access statement;

parsing (S320) the database access statement to determine a target user parameter in the database access statement;

determining (S330) a statement type of the database access statement, wherein determining the statement type comprises

parameterizing (S550) the database access statement to obtain a parameterized database access statement by

(1) extracting (S530) a corresponding object set from the database access statement,

(2) for an object in the object set, determining (S540) an object type group that the object belongs to, and obtaining a corresponding group name, and

(3) replacing the object in the database access statement with the corresponding group name, and replacing the target user parameter with a wildcard, and

determining (S560) the statement type based on the parameterized database access statement;

obtaining (S340) a pre-stored execution plan group corresponding to statement type of the database access statement, wherein execution plans in the execution plan group are distinguished by using different user parameters;

selecting (S350) a target execution plan from the execution plan group based on the target user parameter; and

executing (S360) the target execution plan.

2. The method according to claim 1, further comprising a step of determining multiple object type groups, comprising:

selecting multiple objects with the same object pattern from a database;

grouping the multiple objects;

repeating the previously described selection and grouping steps until all objects in

the database are categorized into corresponding multiple groups; and

using the multiple groups obtained by categorization as the multiple object type groups, and establishing a mapping relationship between each object in the database and the multiple object type groups, wherein

the determining an object type group that the object belongs to comprises:

determining the object type group that the object belongs to, based on the mapping relationship.

3. The method according to claim 1 or 2, wherein the selecting a target execution plan from the execution plan group based on the target user parameter comprises:

for the object in the object set, determining whether an object pattern of the object is the same as an object pattern of each object in the object type group that the object belongs to;

if the object patterns are the same, successively calculating a similarity between the object and each object; and

when a similarity between the object and any object is less than a predetermined threshold, selecting the target execution plan from the execution plan group based on the target user parameter; or

when no similarity between the object and each object is less than a predetermined threshold, re-determining an object type group that the object belongs to, and obtaining a corresponding target execution plan based on the re-determined object type group.

4. The method according to claim 2, further comprising:

for each object in each object type group, periodically and successively calculating similarities between the object and other objects in the object type group; and

when none of the similarities between the object and other objects is less than a predetermined threshold, re-determining an object type group that the object belongs to.

5. The method according to claim 1, wherein the selecting a target execution plan from the execution plan group based on the target user parameter comprises:

calculating an Euclidean distance between the target user parameter and a user parameter corresponding to each execution plan in the execution plan group; and

when an Euclidean distance between the target user parameter and a user

parameter corresponding to any execution plan is less than a distance threshold, using the any execution plan as the target execution plan.

6. The method according to claim 1, wherein when the target execution plan is not obtained by selection, the method further comprises:

generating a new execution plan based on the target user parameter; and

adding the new execution plan to the execution plan group and executing the new execution plan.

7. A database access apparatus, comprising:

a receiving unit (601), configured to receive a database access statement;

a parsing unit (602), configured to parse the database access statement received by the receiving unit (601), to determine a target user parameter comprised in the database access statement;

a determining unit (603), configured to determine a statement type of the database access statement, the determining a statement type of the database access statement comprising:

parameterizing the database access statement to obtain a parameterized database access statement, the parametrizing comprising:

extracting a corresponding object set from the database access statement;

for an object in the object set, determining an object type group that the object belongs to, and obtaining a corresponding group name; and

replacing the object in the database access statement with the corresponding group name, and replacing the target user parameter with a wildcard; and

determining the statement type based on the parameterized database access statement;

an acquisition unit (604), configured to obtain a pre-stored execution plan group corresponding to the database access statement of the statement type determined by the determining unit (603), wherein execution plans in the execution plan group are distinguished by using different user parameters;

a selection unit (605), configured to select a target execution plan from the execution plan group obtained by the acquisition unit (604), based on the target user parameter; and

an execution unit (606), configured to execute the target execution plan selected by the selection unit (605).

8. The apparatus according to claim 7, wherein the apparatus further comprises a grouping unit and an establishment unit;

the selection unit is further configured to select multiple objects with the same object pattern from a database;

the grouping unit is configured to group the multiple objects selected by the selection unit;

the execution unit is further configured to repeat the previously described selection and grouping steps until all objects in the database are categorized into corresponding multiple groups;

the establishment unit is configured to use the multiple groups obtained by the grouping unit as the multiple object type groups, and establish a mapping relationship between each object in the database and the multiple object type groups; and

the determining unit is further configured to:

determine the object type group that the object belongs to, based on the mapping relationship.

9. The apparatus according to claim 7, wherein the selection unit is configured to:

for the object in the object set, determine whether an object pattern of the object is the same as an object pattern of each object in the object type group that the object belongs to;

if the object patterns are the same, successively calculate a similarity between the object and each object; and

when a similarity between the object and any object is less than a predetermined threshold, select the target execution plan from the execution plan group based on the target user parameter; or

when no similarity between the object and each object is less than a predetermined threshold, re-determine an object type group that the object belongs to, and obtain a corresponding target execution plan based on the re-determined object type group.

10. The apparatus according to claim 8, wherein the apparatus further comprises:

a calculation unit, configured to: for each object in each object type group,

periodically and successively calculate similarities between the object and other objects in the object type group; and

the determining unit is further configured to: when none of the similarities between the object and other objects is less than a predetermined threshold, re-determine an object type group that the object belongs to.

11. The apparatus according to claim 7, wherein the selection unit is further configured to:

calculate an Euclidean distance between the target user parameter and a user parameter corresponding to each execution plan in the execution plan group; and

when an Euclidean distance between the target user parameter and a user parameter corresponding to any execution plan is less than a distance threshold, use the any execution plan as the target execution plan.

12. The apparatus according to claim 7, wherein the apparatus further comprises:
a generation unit, configured to generate a new execution plan based on the target user parameter; and

an adding unit, configured to add the new execution plan generated by the generation unit to the execution plan group, and execute the new execution plan.

13. A computer-readable medium comprising instructions stored thereon which, when executed by a computing device, cause the computing device to perform the method of any of claims 1 to 6.