



(11)

EP 3 324 311 A1

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
23.05.2018 Bulletin 2018/21

(51) Int Cl.:
G06F 17/30^(2006.01)

(21) Application number: 17202770.8

(22) Date of filing: 14.01.2011

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB
GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO
PL PT RO RS SE SI SK SM TR

(30) Priority: 15.01.2010 US 688316

(62) Document number(s) of the earlier application(s) in accordance with Art. 76 EPC:
11702533.8 / 2 524 328

(71) Applicant: Ab Initio Technology LLC
Lexington, MA 02421 (US)

(72) Inventors:
• SCHECHTER, Ian
Sharon, MA Massachusetts 02067 (US)

- ALLIN, Glenn John
Arlington, MA Massachusetts 02476 (US)
- WHOLEY, III Joseph Skeffington
Belmont, MA Massachusetts 02478 (US)


(74) Representative: Herrmann, Daniel
Boehmert & Boehmert
Anwaltspartnerschaft mbB
Pettenkoferstrasse 22
80336 München (DE)

Remarks:

This application was filed on 21-11-2017 as a divisional application to the application mentioned under INID code 62.

(54) MANAGING DATA QUERIES

(57) One method includes receiving a database query, receiving information about a data source (102) such as a database table (324a, 324b, 324c) in data storage populated with data elements, producing a structural representation (334) of the database table that includes a formatted data organization reflective of the database table and is absent the data elements of the database table, and providing the structural representation and the database query to a plan generator (318) capable of producing a query plan (314). Another method includes receiving a query plan (314) from a plan generator (318), the plan representing operations for executing a database query on a data source (102) such as a database table (324a, 324b, 324c), and producing a dataflow graph (312) from the query plan.



Description

BACKGROUND

[0001] This description relates to managing data queries.

[0002] Data can be stored in databases and arranged in various forms such as database tables. A database table can contain a set of data having a common theme or purpose. The arrangement of a database table can be defined by a database schema, and multiple database tables can have similar or identical arrangements. Further, the contents of a database and its associated database tables can change over time as data is adjusted, appended or deleted. Various techniques can be used to transfer data into and out of a database and to manipulate the data in the database.

SUMMARY

[0003] In one aspect, in general, a method of preparing a database query for use by a data management system includes receiving a database query, receiving information about a data source, such as a database table, from a data storage, the database table being populated with data elements, producing a structural representation of the database table wherein the structural representation includes a formatted data organization reflective of the database table and is absent the data elements of the database table, and providing the structural representation and the database query to a plan generator capable of producing a query plan representing operations for executing the database query on the database table.

[0004] Aspects can include one or more of the following features.

[0005] The database query includes an SQL query.

[0006] Producing the structural representation of the database table includes determining data types from data structures in the database table, and assigning the data types to corresponding data structures in the structural representation.

[0007] Producing the structural representation of the database table includes determining data storage sizes of the data elements in the database table, and assigning the calculated data storage sizes to the structural representation.

[0008] Receiving the information about the database table includes identifying the database table based on a data source registry.

[0009] The data source registry contains data formats, keys, and indices of data sources.

[0010] In one aspect, in general, a method of generating a dataflow graph representing a database query includes receiving a query plan from a plan generator, the query plan representing operations for executing a database query on a database table, and producing a dataflow graph from the query plan, wherein the dataflow graph includes at least one node that represents at least

one operation represented by the query plan, and includes at least one link that represents at least one dataflow associated with the query plan.

[0011] Aspects can include one or more of the following features.

[0012] The database query includes an SQL query.

[0013] The dataflow graph includes at least one node receiving a flow of data from at least one input dataset, and at least one node providing a flow of data to at least one output dataset.

[0014] The method further includes providing the database table to the dataflow graph, and executing the dataflow graph.

[0015] The method further includes receiving the output of the database query from the dataflow graph.

[0016] The plan generator includes a query optimizer.

[0017] Generating a specification of a dataflow graph includes mapping query operations of the query plan to components of the dataflow graph.

[0018] The query plan includes data types of parameters in the database query.

[0019] The database query contains a reference to an executable function, and the specification of a dataflow graph includes a component representing the executable function.

[0020] In one aspect, in general, a method of generating a dataflow graph representing a database query includes receiving an SQL query, receiving information about a database table from a data storage, the database table being populated with data elements, producing a structural representation of the database table wherein the structural representation includes a formatted data organization reflective of the database table and is absent the data of the database table, providing the structural representation and the SQL query to a plan generator

capable of producing a query plan representing operations for executing the SQL query on the database table, receiving a query plan from the plan generator, and producing a dataflow graph from the query plan, wherein the dataflow graph includes at least one node that represents at least one operation represented by the query plan, and includes at least one link that represents at least one dataflow associated with the query plan.

[0021] In one aspect, in general, a system for preparing a database query for use by a data management system includes a data storage containing information about a database table, the database table being populated with data elements, and a computer system configured to receive a database query, produce a structural representation of the database table wherein the structural representation includes a formatted data organization reflective of the database table and is absent the data elements of the database table, and provide the structural representation and a database query to a plan generator capable of producing a query plan representing operations for executing the database query on the database table.

[0022] In one aspect, in general, a system for generating a dataflow graph representing a database query

includes a computer system configured to receive a query plan from a plan generator, the query plan representing operations for executing a database query on a database table, and produce a dataflow graph from the query plan, wherein the dataflow graph includes at least one node that represents at least one operation represented by the query plan, and includes at least one link that represents at least one dataflow associated with the query plan.

[0023] In one aspect, in general, a system for preparing a database query for use by a data management system includes means for storing information about a database table, the database table being populated with data elements, and means for processing database queries, the processing including receiving a database query, producing a structural representation of the database table wherein the structural representation includes a formatted data organization reflective of the database table and is absent the data elements of the database table, and providing the structural representation and the database query to a plan generator capable of producing a query plan representing operations for executing the database query on the database table.

[0024] In one aspect, in general, a system for generating a dataflow graph representing a database query includes means for processing a query plan, the processing including receiving a query plan from a plan generator, the query plan representing operations for executing a database query on a database table, and producing a dataflow graph from the query plan, wherein the dataflow graph includes at least one node that represents at least one operation represented by the query plan, and includes at least one link that represents at least one dataflow associated with the query plan.

[0025] In one aspect, in general, a computer-readable medium storing a computer program for preparing a database query for use by a data management system includes instructions for causing a computer to receive a database query, receive information about a database table from a data storage, the database table being populated with data elements, produce a structural representation of the database table wherein the structural representation includes a formatted data organization reflective of the database table and is absent the data elements of the database table, and provide the structural representation and the database query to a plan generator capable of producing a query plan representing operations for executing the database query on the database table.

[0026] In one aspect, in general, a computer-readable medium storing a computer program for generating a dataflow graph representing a database query includes instructions for causing a computer to receive a query plan from a plan generator, the query plan representing operations for executing a database query on a database table, and produce a dataflow graph from the query plan, wherein the dataflow graph includes at least one node that represents at least one operation represented by the query plan, and includes at least one link that represents

at least one dataflow associated with the query plan.

DESCRIPTION OF DRAWINGS

5 [0027]

FIG. 1 is a block diagram of a dataflow graph.

FIG. 2 illustrates a database query and a dataflow graph.

10 FIG. 3 is an overview of a database system and associated components.

FIG. 4 represents the execution of a dataflow graph.

15 FIG. 5 is a database table and a structural representation of a database table.

FIG. 6 illustrates a database query and a dataflow graph each containing an executable function.

FIGS. 7-8 are flowcharts for operations of a database system.

20 DESCRIPTION

1 Databases, Queries, and Graphs

[0028] A database management system handles data 25 stored in one or more data sources. Such data sources can take one or more forms such as database tables, which can be collections of data organized into data structures such as rows and columns. In one construct, each row represents a record of data, and each column represents a field within each of the rows.

[0029] The information contained in a database can be 30 accessed and processed (e.g., modified) using database queries. A database query is a set of instructions describing a subset of the database contents and actions to take upon the data in that subset. For example, some database systems perform database queries written in a dedicated database query language such as Structured Query Language (SQL). In these database systems, an SQL query is the primary instrument for manipulating the contents of the database.

[0030] In some implementations, database queries and other computations associated with a database management system in a graphical representation. For example, data may be represented as passing through a 35 collection of operations, referred to as a dataflow. In one arrangement, a dataflow may be provided through a directed graph, with components of the computation being associated with the vertices of the graph and dataflows between the components corresponding to links (arcs, edges) of the graph. A graph is a modular entity and may be connected to or combined with other modular graphs. Each graph can be made up of one or more other graphs, and a particular graph can be a component in a larger graph. A graphical development environment (GDE) provides a user interface for specifying executable graphs and defining parameters for the graph components. A 40 system that implements such graph-based computations is described in U.S. Patent 5,966,072, EXECUTING

COMPUTATIONS EXPRESSED As GRAPHS, which is incorporated herein by reference in its entirety.

[0031] Referring to FIG. 1, an example of a dataflow graph 101 includes an input component 102 providing a collection of data to be processed by the executable components 104a- 104j of the dataflow graph 101. In some example example, the input component 102 is a data source that can comprise data records associated with a database system or transactions associated with a transaction processing system. The data records may be stored in various forms such as a database table, for example. Each executable component 104a - 104j is associated with a portion of the computation defined by the overall dataflow graph 101. Work elements (e.g., individual data records from the data collection or database table) enter one or more input ports of a component, and output work elements (which are in some cases the input work elements, or processed versions of the input work elements) typically leave one or more output ports of the component. In dataflow graph 101, output work elements from components 104e, 104g, and 104j are stored in output data components 102a- 102c.

[0032] Some graph-based database systems are used to process database queries. For example, a database query can be applied to one or more database tables to extract an identified subset of the database table contents, for example, for processing in a dataflow. In some implementations, a graph-based database system accepts and executes database queries in the form of dataflow graphs. Other database systems may use other types of database queries.

[0033] In some cases, one or more database tables are moved to a graph-based database system from another kind of database system that uses SQL. The other database system may have many SQL queries already written that are incompatible with the graph-based database system. In some implementations, the SQL queries can be converted to database queries compatible with the graph-based database system.

[0034] FIG. 2 shows an example of a database query 200 written in SQL. This example database query 200 is intended to operate on database tables 202, 204 managed by a graph-based database management system. The graph-based database management system can recognize the data in the database tables 202, 204 because the database tables are in a format native to the graph-based database management system. The database tables 202, 204 might contain data obtained from other database tables, for example, database tables originating from a database system that recognizes SQL queries and does not use dataflow graphs.

[0035] However, the graph-based database management system does not have built-in functionality for processing SQL queries, so a graph-based database query can be produced to emulate the SQL database query 200. The graph-based database query is recognizable by the graph-based database system. For example, the database query 200 can be converted 206 from

an SQL query into a dataflow graph 208. The dataflow graph 208 operates on the database tables 202, 204 by accepting them as input, and provides the execution results of the database query 200 as output.

5

2 Query Plans

[0036] Some database systems carry out database queries such as SQL queries based on a query plan (also sometimes called an explain plan). A query plan is a description of the database operations that may be performed if the database query is executed. The query plan may describe one possible arrangement of database operations, even if other arrangements of the operations or a different set of operations would accomplish the same result.

[0037] To provide such query plans, a database system may include a query plan generator (also sometimes called a query planner). For example, the query plan generator can produce a query plan when a database query is being executed, or the query plan generator can generate a query plan before any decision about executing the query is made.

[0038] In some arrangements, database operations may be executed in various orders while still providing equivalent outputs. As such, the query plan generator may have functionality that determines an optimal query plan. For example, an optimal query plan could be the query plan that describes the arrangement of database operations for executing the database query in the least amount of time, or using the least amount of database resources such as data storage space, or otherwise accomplishing the database query within constraints that have been identified by the database system. The query plan generator's functionality for determining an optimal query plan may include functionality that scores or ranks many possible query plans, and may also include functionality that rearranges possible query plans to an optimal or efficient configuration.

[0039] A single database query can be executed multiple times, and each execution could have a unique optimal query plan. For example, the data within a database table could change between two executions of a database query. In this example, the operations described in the query plan that was generated for the first execution of the database query may need more or less execution time during the second execution of the database query than during the first execution. In this case, a different arrangement of operations may be better suited to the second execution of the database query, for example, a different arrangement of the same operations, or an arrangement of different operations. A query plan optimized for the second execution of the database query can be generated for that execution, taking into account the momentary state of the database table.

3 System Overview

[0040] A query plan generator can be used in producing a graph-based database query that emulates another kind of database query such as an SQL query. FIG. 3 shows a database query management system 300 for preparing a database query 302 for execution on a database management computer system 304. The database management computer system 304 shown includes a graph execution engine 306 that handles database operations implemented as dataflow graphs. The database query management system 300 also includes a graph generation computer system 308 having a graph generation engine 310 that can build a dataflow graph 312 from a description of operations to be performed by the dataflow graph. For example, the description of operations could be a query plan 314.

[0041] The database query management system 300 also includes a query planning computer system 316 that executes a query plan generator 318. The query plan generator 318 can be any query plan generator that produces a query plan from a database query, and need not be designed with any functionality related to dataflow graphs or graph generation. Further, the database query management system 300 also includes a database computer system 320 having a database 322 in a data storage (e.g. a hard drive, optical disc, etc.) and containing one or more database tables 324a, 324b, 324c.

[0042] Although separate computer systems are shown for the database management computer system 304, the graph generation computer system 308, the query planning computer system 316, and the database computer system 320, two or more of these computer systems could be the same computer system, or components of the same computer system. All of the computer systems have at least one processor for executing their respective executable components and at least one data storage system. The computer systems can be connected to each other using a computer network such as a local area network (LAN), a wide-area network (WAN), a network such as the Internet, or another kind of computer network.

[0043] To demonstrate the production of a dataflow graph 312 from one or more database queries 302, a database query and a database table 326 are received and processed by the graph generation computer system 308 prior to an execution of the dataflow graph. The graph generation computer system 308 receives the database table 326 from the database computer system 320.

[0044] The database table 326 can take any of several forms. For example, the database table 326 could be a relational database table, a partial database table, a flat file, or another kind of data file or collection of data files. In some examples, the database table 326 could be received in the form of information about the database table, e.g. metadata about the database table, or a description of the database table.

[0045] In some implementations, the database table

326 could be identified by a data registry associated with the database computer system 320 or otherwise accessible to the graph generation computer system 308. The data registry could be in the form of lookup file catalog, for example, which may contain a data file location associated with the database table 326, and primary key and index information associated with the database table. The data registry could also provide information about the data formats for different types of database tables. Further, the data registry could also provide information about how a dataflow graph 312 can access the database table.

[0046] The graph generation computer system 308 also receives a database query 302 to be applied to the database table 326. For example, the database query could be an SQL query. The database query 302 could be received from any number of possible sources. For example, the database query 302 could be received from a user interface 328 where a user 330 has entered the database query. In some examples, the database query 302 is received from a data storage, or the database query is received from a computer network such as the Internet, or the database query is generated based on another previously-received database query.

[0047] In some implementations, the graph generation computer system 308 provides (as represented by an arrow 332) a version of the database table 326 to the query planning computer system 316, produced from information about the database table 326. For example, the version of the database table 326 provided to the query planning computer system 316 could be a structural representation 334 of the database table that is smaller in size than the database table and thus requires fewer computational resources to process. The structural representation 334 of the database table 326 may contain information about the database table, but is absent some or all of the data of the database table. For example, the structural representation 334 of the database table 326 could contain a format that reflects the format of the database table, such as the columns, rows, or fields of the database table. The structural representation 334 of the database table 326 could also contain information about the data, such as data storage sizes of elements in the database table, or the data types of elements in the database table.

[0048] The graph generation computer system provides (as represented by the arrow 332) the database query 302 and the structural representation 334 of the database table 326 to the query planning computer system 316. The query planning computer system 316 executes the query plan generator 318, which produces a query plan optimized for executing the database query 302 over the database table 326. The structural representation 334 of the database table 326 supplies the same information used by the query plan generator 318 as would be supplied by the database table 326 itself, for example, data sizes, data types, and other information about the data contained in the database table. In some

implementations, the structural representation 334 of the database table 326 includes an index of data elements that is used to optimize data lookup and retrieval. The query plan generator 318 can use the index to calculate the speed of identifying and retrieving data elements from the indexed database table 326.

[0049] The graph generation computer system 308 receives a query plan 314 from the query planning computer system 316. The query plan 314 describes an arrangement of database operations that can be used to execute the database query 302 over the database table 326. For example, the operations in the query plan 314 can correspond to nodes of a dataflow graph 312. The query plan 314 can also include information about the data types used by the operations in the query plan. For example, the operations in the query plan could have parameters, and the data types of the parameters could be described in the query plan.

[0050] Once produced, the query plan 314 is provided to the graph generation computer system 308 for dataflow graph 312 production by the graph generation engine 310. The graph generation engine 310 outputs a dataflow graph 312 corresponding to the query plan 314. In some implementations, the dataflow graph 312 has nodes representing operations described in the query plan, and node links representing flows of data between the operations. Because a dataflow graph 312 may be generated for each instance of preparing a database query for execution, the graph generation engine 310 can generate a dataflow graph quickly enough to respond to real-time requests to execute a database query. In some implementations, the graph generation engine 310 can generate a dataflow graph from a query plan in less than one second.

[0051] The graph generation computer system 308 provides (represented by an arrow 336) the dataflow graph 312 generated by the graph generation engine 310 to the database management computer system 304. In some implementations, the graph generation computer system 308 also prepares the database table 326 for use by the database management computer system and provides the prepared database table 338. For example, graph generation computer system 308 might convert the database table 326 from a format used by the graph generation computer system 308 to a format used by the database management computer system 304.

[0052] Once provided to the database management computer system 304, the dataflow graph 312 is prepared for execution. As shown in FIG. 4, the database management computer system 304 can execute operations of the dataflow graph 312 and use the database table 326 in order to produce results 402 of the database query. The database management computer system 304 provides the database table 326 to one or more nodes 404a, 404b, 404c of the dataflow graph 312 and executes the dataflow graph using the graph execution engine 306. The graph execution engine 306 performs the operations represented by the nodes 404a, 404b, 404c of the data-

flow graph 312, which correspond to database operations for executing the underlying database query. Further, links 408a, 408b, 408c between the nodes represent flows of data between the database operations as the database table is processed. The dataflow graph 312 outputs the results 402 of the database query.

4 Structural Representation

[0053] FIG. 5 shows an example of a database table 500 and a structural representation 502 of the database table (containing none of the data from the database table). The database table 500 has columns 504a-504i that may contain similar or different types of data. The database table 500 also has rows 508a-508e each containing a field corresponding to each column. Each field of a row contains a data element 510a, 510b, 510c of the data type of the corresponding column (e.g. character string, integer, floating point number, etc.). Further, each row 508a-508e has an inherent data storage size 512a-512e. For example, the data storage size 512a-512e might be the amount of storage space used by the data elements of the row when the row resides in data storage such as memory.

[0054] The structural representation 502 of the database table 500 (produced by the graph generation computer system 308 as shown in FIG. 3) has columns 514a-514i identical to the database table, including the same data types as the original columns 504a-504i. The structural representation 502 of the database table 500 also has rows 518a-518e corresponding to the rows 508a-508e of the database table. However, the rows 518a-518e do not contain the data elements 510a, 510b, 510c from the database table 500. Each row 518a-518e is associated with a data storage size value 520a-520e. While the data storage size 512a-512e of a row in the database table 500 could be calculated from the data storage sizes of the individual data elements 510a, 510b, 510c, the data storage size value 520a-520e of each row 518a-518e can itself be a data element stored alongside each row in the structural representation 502 of the database table. The total data storage size of the structural representation 502 of the database table 500 may be a small percentage of the total data storage size of the database table, because the fields 522a, 522b, 522c of the structural representation of the database table are absent the data elements 510a, 510b, 510c from the fields of the database table. In some implementations, the structural representation 502 of the database table 500 may retain some data elements from the database table, for example, the structural representation of the database table may retain data elements corresponding to key values 524, 526 or other data elements used in the structure, organization, or format of the database table. In some implementations, the structural representation 502 of the database table 500 may contain an index or another data structure that provides information used in data retrieval. In some implementations, the structural representation

502 of the database table 500 may contain other statistics about the data contained in the database table.

5 Function Calls

[0055] As shown in FIG. 6, in some implementations, a custom graph function 602 can be embedded in a database query 604. For example, the custom graph function 602 might represent an operation that a dataflow graph 606 can carry out during the execution of the dataflow graph. The custom graph function 602 might not have any functionality within the context of the database query 604, and is placed in the database query to be passed to the generated dataflow graph 606. For example, the custom graph function 602 might be added to a database query 604 to prepare it for conversion to a dataflow graph but replacing or augmenting part of the existing database query. Further, the query plan generator might not have information about the functionality of the custom graph function 602. The custom graph function 602 may be a function that has no equivalent function in the language of the database query 604. In some implementations, the query plan generator may be aware of the input data type or output data type for the custom graph function 602. When the dataflow graph 608 is generated 608, the custom graph function 602 could act as a node 610 of the dataflow graph.

[0056] In some implementations, the custom graph function 602 may be a function used for accessing data in a special or customized format, or the custom graph function may be a function used to execute a regular expression or pattern matching expression, or the custom graph function may be a function implemented by a user, for example, a user of the database management computer system.

6 Flow of Operation

[0057] FIG. 7 is a flowchart 700 showing exemplary operations of the graph generation computer system 308 (shown in FIG. 3). In step 702, the graph generation computer system receives a database query, for example, from a user interface or from data storage. In step 704, the graph generation computer system receives information a database table from a database on a database system, e.g., the graph generation computer system may receive the database table itself from a data storage, or the graph generation computer system may receive metadata about the database table, or the graph generation computer system may receive information about the database table as entered by a user at a user interface. In step 706, the graph generation computer system generates a structural representation of the database table that includes a formatted data organization reflective of the database table and is absent the data of the database table. In step 708, the graph generation computer system provides the structural representation of the database table and the database query to a query planning

computer system for use by a query plan generator. In step 710, the graph generation computer system receives a query plan, produced by the query plan generator, from the query planning computer system. In step 5 712, the graph generation computer system uses a graph generation engine to generate a dataflow graph from the query plan based on the operations described in the query plan. In step 714, the graph generation computer system provides the dataflow graph for execution (e.g. by a graph execution engine of a database management system).

[0058] FIG. 8 is a flowchart 800 showing exemplary operations of the graph generation computer system 308 (shown in FIG. 3). In step 802, the graph generation computer system receives information about a database table, for example, the graph generation computer system may receive a database table previously received from a user interface or from data storage, or may receive metadata describing the database table. In step 10 804, the graph generation computer system generates a new database table having the same column and row format of the received database table. In step 806, the graph generation computer system identifies the data types associated with the received database table by extracting the data types from the columns of the database table. In step 15 808, the graph generation computer system assigns the identified data types to the columns of the new database table. In step 810, the graph generation computer system calculates data statistics about the received database table, for example, the data storage sizes of the data elements of the rows of the received database table. In step 20 812, the graph generation computer system assigns the calculated data statistics to the new database table, which serves as the structural representation of the received database table.

[0059] The database query managing approach described above can be implemented using software for execution on a computer. For instance, the software forms procedures in one or more computer programs that 40 execute on one or more programmed or programmable computer systems (which may be of various architectures such as distributed, client/server, or grid) each including at least one processor, at least one data storage system (including volatile and non-volatile memory and/or storage elements), at least one input device or port, and at least one output device or port. The software may form one or more modules of a larger program, for example, that provides other services related to the design and configuration of computation graphs. The nodes and elements of the graph can be implemented as data structures stored in a computer readable medium or other organized data conforming to a data model stored in a data repository.

[0060] The software may be provided on a storage medium, such as a CD-ROM, readable by a general or special purpose programmable computer or delivered (encoded in a propagated signal) over a communication medium of a network to the computer where it is executed.

All of the functions may be performed on a special purpose computer, or using special-purpose hardware, such as coprocessors. The software may be implemented in a distributed manner in which different parts of the computation specified by the software are performed by different computers. Each such computer program is preferably stored on or downloaded to a storage media or device (e.g., solid state memory or media, or magnetic or optical media) readable by a general or special purpose programmable computer, for configuring and operating the computer when the storage media or device is read by the computer system to perform the procedures described herein. The inventive system may also be considered to be implemented as a computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer system to operate in a specific and predefined manner to perform the functions described herein.

[0061] A number of embodiments of the invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. For example, some of the steps described above may be order independent, and thus can be performed in an order different from that described.

[0062] Various embodiments have been described. The present invention may also (alternatively) be described by the below numbered aspects 1 to 22:

1. A method of preparing a database query for use by a data management system, the method including:

receiving a database query;
 receiving information about a data source from a data storage, the data source being populated with data elements;
 producing a structural representation of the data source wherein the structural representation includes a formatted data organization reflective of the data source and is absent the data elements of the data source; and
 providing the structural representation and the database query to a plan generator capable of producing a query plan representing operations for executing the database query on the data source.

2. The method of aspect 1, wherein the database query includes an SQL query.

3. The method of aspect 1, wherein producing the structural representation of the data source includes:

determining data types from data structures in the data source; and
 assigning the data types to corresponding data structures in the structural representation.

4. The method of aspect 1, wherein producing the structural representation of the data source includes:

determining data storage sizes of the data elements in the data source; and
 assigning the calculated data storage sizes to the structural representation.

5. The method of aspect 1, wherein receiving the information about the data source includes identifying the data source based on a data source registry.

6. The method of aspect 5, wherein the data source registry contains data formats, keys, and indices of data sources.

7. A method of generating a dataflow graph representing a database query, the method including:

receiving a query plan from a plan generator, the query plan representing operations for executing a database query on a data source; and producing a dataflow graph from the query plan, wherein the dataflow graph includes at least one node that represents at least one operation represented by the query plan, and includes at least one link that represents at least one dataflow associated with the query plan.

8. The method of aspect 7, wherein the database query includes an SQL query.

9. The method of aspect 7, wherein the dataflow graph includes:

at least one node receiving a flow of data from at least one input dataset, and
 at least one node providing a flow of data to at least one output dataset.

10. The method of aspect 7, further including:

providing the data source to the dataflow graph; and
 executing the dataflow graph.

11. The method of aspect 10, further including receiving the output of the database query from the dataflow graph.

12. The method of aspect 7, wherein the plan generator includes a query optimizer.

13. The method of aspect 7, wherein generating a specification of a dataflow graph includes mapping query operations of the query plan to components of the dataflow graph.

14. The method of aspect 7, wherein the query plan includes data types of parameters in the database query. 5

15. The method of aspect 7, wherein the database query contains a reference to an executable function, and the specification of a dataflow graph includes a component representing the executable function. 10

16. A method of generating a dataflow graph representing a database query, the method including: 15

receiving an SQL query;
 receiving information about a data source from a data storage, the data source being populated with data elements; 15
 producing a structural representation of the data source wherein the structural representation includes a formatted data organization reflective of the data source and is absent the data of the data source; 20
 providing the structural representation and the SQL query to a plan generator capable of producing a query plan representing operations for executing the SQL query on the data source; 25
 receiving a query plan from the plan generator; and
 producing a dataflow graph from the query plan, wherein the dataflow graph includes at least one node that represents at least one operation represented by the query plan, and includes at least one link that represents at least one dataflow associated with the query plan. 30

17. A system for preparing a database query for use by a data management system, the system including: 35

a data storage containing information about a data source, the data source being populated with data elements; and 40
 a computer system configured to receive a database query, produce a structural representation of the data source wherein the structural representation includes a formatted data organization reflective of the data source and is absent the data elements of the data source, and 45
 provide the structural representation and a database query to a plan generator capable of producing a query plan representing operations for executing the database query on the data source. 50

18. A system for generating a dataflow graph representing a database query, the system including: 55

a computer system configured to

receive a query plan from a plan generator, the query plan representing operations for executing a database query on a data source, and produce a dataflow graph from the query plan, wherein the dataflow graph includes at least one node that represents at least one operation represented by the query plan, and includes at least one link that represents at least one dataflow associated with the query plan. 60

19. A system for preparing a database query for use by a data management system, the system including: 65

means for storing information about a data source, the data source being populated with data elements; and
 means for processing database queries, the processing including
 receiving a database query,
 producing a structural representation of the data source wherein the structural representation includes a formatted data organization reflective of the data source and is absent the data elements of the data source, and
 providing the structural representation and the database query to a plan generator capable of producing a query plan representing operations for executing the database query on the data source. 70

20. A system for generating a dataflow graph representing a database query, the system including: 75

means for processing a query plan, the processing including
 receiving a query plan from a plan generator, the query plan representing operations for executing a database query on a data source, and producing a dataflow graph from the query plan, wherein the dataflow graph includes at least one node that represents at least one operation represented by the query plan, and includes at least one link that represents at least one dataflow associated with the query plan. 80

21. A computer-readable medium storing a computer program for preparing a database query for use by a data management system, the computer program including instructions for causing a computer to: 85

receive a database query;
 receive information about a data source from a data storage, the data source being populated with data elements;
 produce a structural representation of the data source wherein the structural representation in-

cludes a formatted data organization reflective of the data source and is absent the data elements of the data source; and provide the structural representation and the database query to a plan generator capable of producing a query plan representing operations for executing the database query on the data source.

22. A computer-readable medium storing a computer program for generating a dataflow graph representing a database query, the computer program including instructions for causing a computer to:

receive a query plan from a plan generator, the query plan representing operations for executing a database query on a data source; and produce a dataflow graph from the query plan, wherein the dataflow graph includes at least one node that represents at least one operation represented by the query plan, and includes at least one link that represents at least one dataflow associated with the query plan.

[0063] It is to be understood that the foregoing description is intended to illustrate and not to limit the scope of the invention, which is defined by the scope of the appended claims.

Claims

1. A computer-implemented method of operating on data from a data source, the operating based on a query that is written in accordance with a database query language, the operating including:

identifying a data source; receiving a query to be applied to the data source; identifying a requesting system (308) other than a system (320) having a relational database that stores data of the data source; generating, by the requesting system, a request to a query planner (318) based on the query; providing, by the requesting system, the request to the query planner; receiving a query plan generated by the query planner based on the request, the query plan including a description of database operations for executing the query on the data source; generating, from the query plan based on the database operations described in the query plan, a dataflow graph that includes:

a first node that is configured to receive input data and that represents at least one operation to be executed on the input data,

wherein the at least one operation is described in the query plan, and at least one link that represents a flow of output data of an operation of the first node being provided as input data to an operation represented by a second node of the dataflow graph;

receiving data from the data source; identifying an executing system (304) other than the system (320) having the relational database that stores data of the data source; and executing, on the identified executing system (304), one or more operations of the dataflow graph using the data of the data source to produce results of the query.

2. The method of claim 1, the method further comprising:

embedding a custom graph function (602) in the query (604), the custom graph function (602) replacing or augmenting a part of the query (302, 604), wherein the custom graph function (602) represents an operation the dataflow graph (100, 312, 606) that carries out during execution of the dataflow graph.

3. The method of claim 2, wherein the custom graph function (602) is a function having no equivalent function in the database query language of the query (302, 604).

4. The method of any one of claims 1 to 3, wherein generating a request to a query planner based on the query includes producing a structural representation of the data source wherein the structural representation includes a formatted data organization reflective of the data source and is absent data of the data source; and including the structural representation in the request to the query planner.

45. The method of claim 2, 3 or 4, wherein the request comprises the query (302, 604) with embedded custom graph function (602).

6. The method of any one of claims 1 to 5, wherein the first node is associated with information usable by the executing system to invoke executable program code to perform the at least one operation, wherein the first node is associated with information usable by the executing system to make data available to the program code, and wherein the executing of the one or more operations of the dataflow graph includes executing the program code to perform the at least one operation.

7. The method of any one of claims 1 to 6, wherein the data source includes a relational database table.

8. The method of any one of claims 1 to 7, including outputting the results of the query. 5

9. The method of any one of claims 1 to 8, wherein the query plan generator includes a query optimizer.

10. The method of any one of claims 1 to 9, wherein the query plan includes data types of parameters in the query. 10

11. The method of any one of claims 1 to 10, wherein the query contains a reference to an executable function, and the dataflow graph includes a node representing the executable function. 15

12. The method of any one of claims 1 to 11, wherein the dataflow graph is generated by the identified requesting system (308). 20

13. A computer-readable medium storing a computer program including instructions for causing a computer to perform the operations of the method of any one of claims 1 to 12. 25

14. A system for operating on data from a data source, the data stored in a tangible, non-transitory computer-readable medium of the data source, the operating based on a query that is written in accordance with a database query language, the system including:
 a computer system configured for: 35
 identifying a data source;
 receiving a query to be applied to the data source;
 identifying a requesting system (308) other than a system (320) having a relational database that stores data of the data source; 40
 generating, by the requesting system, a request to a query planner (318) based on the query;
 providing, by the requesting system, the request to the query planner; 45
 receiving a query plan generated by the query planner based on the request, the query plan including a description of database operations for executing the query on the data source; 50
 generating, from the query plan based on the database operations described in the query plan, a dataflow graph that includes: 55
 a first node that is configured to receive input data and that represents at least

one operation to be executed on the input data, wherein the at least one operation is described in the query plan, and at least one link that represents a flow of output data of an operation of the first node being provided as input data to an operation represented by a second node of the dataflow graph;

receiving data from the data source;
 identifying an executing system (304) other than the system (320) having the relational database that stores data of the data source; and
 executing, on the identified executing system (304), one or more operations of the dataflow graph using the data of the data source to produce results of the query.

15. The system of claim 14, wherein the executing system (304) and the requesting system (308) are different systems. 20




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8



EUROPEAN SEARCH REPORT

Application Number

EP 17 20 2770

5

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
X	US 2003/093410 A1 (COUCH TANYA [US] ET AL) 15 May 2003 (2003-05-15) * paragraphs [0003], [0006] - [0009], [0013], [0014], [0021], [0024] - [0029], [0034] - [0036], [0042] - [0044] * * figures 1, 3, 4, 7A-7L, 8 * * table 1 * ----- A Maciej Pilecki: "Optimizing SQL Server Query Performance", , 30 November 2007 (2007-11-30), XP055458550, Retrieved from the Internet: URL: https://technet.microsoft.com/en-us/library/2007.11.sqlquery.aspx [retrieved on 2018-03-12] * page 2 * -----	1-15	INV. G06F17/30
			TECHNICAL FIELDS SEARCHED (IPC)
			G06F
The present search report has been drawn up for all claims			
1	Place of search	Date of completion of the search	Examiner
50	The Hague	15 March 2018	Martínez Espuche, F
CATEGORY OF CITED DOCUMENTS			
55	X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document	T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document	

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 17 20 2770

5 This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

15-03-2018

10	Patent document cited in search report	Publication date	Patent family member(s)	Publication date
15	US 2003093410	A1 15-05-2003	NONE	
20				
25				
30				
35				
40				
45				
50				
55				

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 5966072 A [0030]

摘要

一种方法包括：接收数据库查询；接收关于数据存储器中的诸如数据库表（324a、324b、324c）的数据源（102）的信息，所述数据源填充有数据元素；产生所述数据库表的结构表示（334），所述结构表示包括反映所述数据库表的经格式化的数据组织，并且不包括所述数据库表的数据元素；以及向计划生成器（318）提供所述结构表示以及所述数据库查询，所述计划生成器能够产生查询计划（314）。另一方法包括：从计划生成器（318）接收查询计划（314）；所述计划表示用于对诸如数据库表（324a、324b、324c）的数据源（102）执行数据库查询的操作；以及根据所述查询计划产生数据流图（312）。

