

CLAIMS

1. A data processing system, comprising:
 - at least one computer hardware processor; and
 - at least one non-transitory computer readable storage medium storing processor executable instructions that, when executed by the at least one computer hardware processor, cause the at least one computer hardware processor to perform:
 - obtaining information specifying a dataflow graph, the dataflow graph comprising a plurality of nodes and a plurality of edges connecting the plurality nodes, the plurality of edges representing flows of data among nodes in the plurality of nodes, the plurality of nodes comprising:
 - a first set of one or more nodes, each node in the first set of nodes representing a respective input dataset in a set of one or more input datasets;
 - a second set of one or more nodes, each node in the second set of nodes representing a respective output dataset in a set of one or more output datasets; and
 - a third set of one or more nodes, each node in the third set of nodes representing at least one respective data processing operation;
 - obtaining a first set of one or more processing layouts for the set of input datasets;
 - obtaining a second set of one or more processing layouts for the set of output datasets; and
 - determining a processing layout for each node in the third set of nodes based on the first set of processing layouts, the second set of processing layouts, and one or more layout determination rules including at least one rule for selecting among processing layouts having different degrees of parallelism.
2. The data processing system of claim 1, wherein the set of input datasets includes a first input dataset and a second input dataset, wherein the first set of nodes includes a first node associated with the first input dataset and a second node associated with the second input dataset, and wherein obtaining the first set of one or more processing layouts

comprises:

obtaining a first processing layout for reading data from the first input dataset, the first layout at least partially specifying a first degree of parallelism; and

obtaining a second processing layout for reading data from the second input dataset, the second processing layout at least partially specifying a second degree of parallelism different from the at least partially specified first degree of parallelism.

3. The data processing system of claim 1,

wherein the first set of nodes includes a first node associated with a first input dataset in the set of input datasets,

wherein the second set of nodes includes a second node associated with a first output dataset in the set of output datasets,

wherein obtaining the first set of one or more processing layouts comprises obtaining a first processing layout for reading data from the first input dataset, the first processing layout at least partially specifying a first degree of parallelism, and

wherein obtaining the second set of one or more processing layouts comprises obtaining a second processing layout for writing data to the first output dataset, the second processing layout at least partially specifying a second degree of parallelism different from the at least partially specified first degree of parallelism.

4. The data processing system of claim 1, wherein the set of output datasets includes a first output dataset and a second output dataset, wherein the second set of nodes includes a first node associated with the first output dataset and a second node associated with the second output dataset, and wherein obtaining the first set of one or more processing layouts comprises:

obtaining a first processing layout for writing data to the first output dataset, the first layout at least partially specifying a first degree of parallelism; and

obtaining a second processing layout for writing data to the second output dataset, the second processing layout at least partially specifying a second degree of parallelism different from the at least partially specified first degree of parallelism.

5. The data processing system of claim 1, wherein determining a processing layout

for each node in the third set of nodes comprises:

determining an initial processing layout for one or more nodes in the third set nodes in a forward pass performed starting from nodes in the first set of nodes in accordance with structure of the dataflow graph and using the first set of processing layouts and the one or more layout determination rules; and

determining a final processing layout for each node in the third set of nodes in a backward pass performed starting from nodes in the second set of nodes in accordance with structure of the dataflow graph and using the second set of processing layouts and the one or more layout determination rules.

6. The data processing system of claim 1, wherein the third set of nodes comprises a first node, wherein the plurality of edges comprises a first edge between the first node and a second node preceding the first node in the dataflow graph, and wherein determining the initial processing layout for each node in the third set nodes comprises:

determining a first initial processing layout for the first node based on a processing layout determined for the second node.

7. The data processing system of claim 6, wherein the plurality of edges comprises a second edge between the first node and a third node preceding the first node in the dataflow graph, wherein a third initial processing layout is associated with the third node, and wherein determining the first initial processing layout for the first node comprises:

selecting one of the second and third processing layouts as the first processing layout.

8. The processing system of claim 7, wherein the second initial processing layout specifies a first degree of parallelism, wherein the third initial processing layout specifies a second degree of parallelism different from the first degree of parallelism, and wherein the selecting comprises:

selecting the second processing layout when the first degree of parallelism is greater than the second degree of parallelism; and

selecting the third processing layout when the first degree of parallelism is less than the second degree of parallelism.

9. The data processing system of claim 7, wherein the second initial processing layout and the third initial processing layout each specifies a same degree of parallelism, wherein the first edge represents a data flow for a first number of data records, wherein the second edge represents a data flow for a second number of data records, and wherein the selecting comprises:

selecting the second processing layout when the first number of data records is greater than the second number of data records; and

selecting the third processing layout when the first number of data records is less than the second number of data records.

10. The data processing system of claim 1,

wherein, during the determining, a first processing layout is determined for a first node in third set of nodes, the first processing layout specifying a first degree of parallelism,

wherein a second processing layout for a second node immediately preceding the first node in the graph specifies a second degree of parallelism different from the first degree of parallelism, and

wherein the processor executable instructions further cause the at least one computer hardware processor to configure at least one node of the dataflow graph to perform at least one repartitioning operation.

11. The data processing system of claim 1,

wherein, during the determining, a first processing layout is determined for a first node in third set of nodes, the first processing layout specifying a first degree of parallelism,

wherein a second processing layout for a second node immediately preceding the first node in the graph specifies a second degree of parallelism different from the first degree of parallelism, and

wherein the processor executable instructions further cause the at least one computer hardware processor to add, to the dataflow graph, a new node between the first node and the second node, the new node representing at least one repartitioning

operation.

12. The data processing system of claim 10, wherein adding the new node associated with the at least one repartitioning operation comprises adding a node associated with a gather operation, a merge operation, or a partition operation.

13. The data processing system of claim 1, wherein the determining comprises determining a first processing layout for a first node in the third set of nodes, the first node representing a first data processing operation, and wherein determining the first processing layout comprises:

- determining a degree of parallelism for performing the first data processing operation; and

- identifying a set of one or more computing devices to perform the first data processing operation in accordance with the determined degree of parallelism.

14. The data processing system of claim 13, wherein determining the first processing layout comprises:

- determining that a single processor is to be used for performing the first data processing operation; and

- identifying a computing device to perform the first data processing operation.

15. The data processing system of claim 1, wherein determining the processing layout is performed using the at least one rule for selecting among processing layouts having different degrees of parallelism.

16. The data processing system of claim 1, wherein the processor-executable instructions cause the at least one computer hardware processor to perform:

- after determining the processing layout for each node in the dataflow graph, executing the dataflow graph in accordance with processing layouts determined for each node in the dataflow graph.

17. The data processing system of claim 1, wherein the processor-executable

instructions cause the at least one computer hardware processor to perform:

- receiving a structured query language (SQL) query;
- generating a query plan from the SQL query; and
- generating the dataflow graph from the generated query plan.

18. The data processing system of claim 1, wherein determining the processing layout for each node in the third set of nodes comprises:

- determining whether adjacent nodes in the third set of nodes were assigned layouts having different degrees of parallelism; and
- when it is determined that the adjacent nodes were assigned layouts having different degrees of parallelism, configuring at least one of the adjacent nodes to perform at least one repartitioning operation and/or adding a new node to the dataflow graph configured to perform the at least one repartitioning operation.

19. At least one non-transitory computer readable storage medium storing processor executable instructions that, when executed by at least one computer hardware processor, cause the at least one computer hardware processor to perform:

- obtaining information specifying a dataflow graph, the dataflow graph comprising a plurality of nodes and a plurality of edges connecting the plurality nodes, the plurality of edges representing flows of data among nodes in the plurality of nodes, the plurality of nodes comprising:

- a first set of one or more nodes, each node in the first set of nodes representing a respective input dataset in a set of one or more input datasets;
 - a second set of one or more nodes, each node in the second set of nodes representing a respective output dataset in a set of one or more output datasets;
- and

- a third set of one or more nodes, each node in the third set of nodes representing at least one respective data processing operation;
- obtaining a first set of one or more processing layouts for the set of input datasets;
- obtaining a second set of one or more processing layouts for the set of output datasets; and

- determining a processing layout for each node in the third set of nodes based on

the first set of processing layouts, the second set of processing layouts, and one or more layout determination rules including at least one rule for selecting among processing layouts having different degrees of parallelism.

20. At least one non-transitory computer readable storage medium storing processor executable instructions that, when executed by the at least one computer hardware processor, the processor executable instructions comprising:

means for obtaining information specifying a dataflow graph, the dataflow graph comprising a plurality of nodes and a plurality of edges connecting the plurality nodes, the plurality of edges representing flows of data among nodes in the plurality of nodes, the plurality of nodes comprising:

a first set of one or more nodes, each node in the first set of nodes representing a respective input dataset in a set of one or more input datasets;

a second set of one or more nodes, each node in the second set of nodes representing a respective output dataset in a set of one or more output datasets; and

a third set of one or more nodes, each node in the third set of nodes representing at least one respective data processing operation;

means for obtaining a first set of one or more processing layouts for the set of input datasets;

means for obtaining a second set of one or more processing layouts for the set of output datasets; and

means for determining a processing layout for each node in the third set of nodes based on the first set of processing layouts, the second set of processing layouts, and one or more layout determination rules including at least one rule for selecting among processing layouts having different degrees of parallelism.

21. A method, comprising:

using at least one computer hardware processor to perform:

obtaining information specifying a dataflow graph, the dataflow graph comprising a plurality of nodes and a plurality of edges connecting the plurality nodes, the plurality of edges representing flows of data among nodes in the

plurality of nodes, the plurality of nodes comprising:

- a first set of one or more nodes, each node in the first set of nodes representing a respective input dataset in a set of one or more input datasets;

- a second set of one or more nodes, each node in the second set of nodes representing a respective output dataset in a set of one or more output datasets; and

- a third set of one or more nodes, each node in the third set of nodes representing at least one respective data processing operation;

- obtaining a first set of one or more processing layouts for the set of input datasets;

- obtaining a second set of one or more processing layouts for the set of output datasets; and

- determining a processing layout for each node in the third set of nodes based on the first set of processing layouts, the second set of processing layouts, and one or more layout determination rules including at least one rule for selecting among processing layouts having different degrees of parallelism.

22. A data processing system, comprising:

- at least one computer hardware processor; and

- at least one non-transitory computer readable storage medium storing processor executable instructions that, when executed by the at least one computer hardware processor, cause the at least one computer hardware processor to perform:

- obtaining information specifying a dataflow graph, the dataflow graph comprising a plurality of nodes and a plurality of edges connecting the plurality nodes, the plurality of edges representing flows of data among nodes in the plurality of nodes, the plurality of nodes comprising:

- a first set of one or more nodes; and

- a second set of one or more nodes disjoint from the first set of nodes;

- obtaining a first set of one or more processing layouts for the first set of

nodes; and

determining a processing layout for each node in the second set of nodes based on the first set of processing layouts and one or more layout determination rules including at least one rule for selecting among processing layouts having different degrees of parallelism.

23. The data processing system of claim 22,
wherein the first set of one or more nodes includes:

a set of one or more input nodes, each node in the set of input nodes representing a respective input dataset in a set of one or more input datasets; and

a set of one or more output nodes, each node in the set of output nodes representing a respective output dataset in a set of one or more output datasets,
and

wherein obtaining the first set processing layouts comprises:

obtaining a set of one or more processing layouts for the set of input datasets; and

obtaining a set of one or more processing layouts for the set of output datasets.

24. A data processing system, comprising:

at least one computer hardware processor; and

at least one non-transitory computer readable storage medium storing processor executable instructions that, when executed by the at least one computer hardware processor, cause the at least one computer hardware processor to perform:

obtaining computer code that, when executed by the at least one computer hardware processor, causes the at least one computer hardware processor to execute a database query, wherein the computer code comprises:

a first set of one or more computer code portions each representing a data processing operation for reading in a respective input dataset;

a second set of one or more computer code portions each representing a data processing operation for writing out a respective output dataset;

a third set of one or more computer code portions each representing a respective data processing operation;
obtaining a first set of one or more processing layouts for one or more code portions in the first set of code portions;
obtaining a second set of one or more processing layouts for one or more code portions in the second set of code portions; and
determining a processing layout for each code portion in the third set of code portions based on the first set of processing layouts, the second set of processing layouts, and one or more layout determination rules including at least one rule for selecting among processing layouts having different degrees of parallelism.

25. The data processing system of claim 24, wherein the determining is performed further based on information at least partially indicating an order of execution of at least some computer code portions in the computer code.

26. The data processing system of claim 24,
wherein the computer code is associated with a dataflow graph,
wherein the dataflow graph comprises a plurality of nodes and a plurality of edges connecting the plurality nodes, the plurality of edges representing flows of data among nodes in the plurality of nodes, and
wherein computer code portions in the computer code are associated with respective nodes in the dataflow graph.

27. The data processing system of claim 26, wherein the plurality of nodes comprises:
a first set of one or more nodes, each node in the first set of nodes representing a respective input dataset, wherein each computer code portion in the first set of computer code portions is associated with a respective node in the first set of nodes;
a second set of one or more nodes, each node in the second set of nodes representing a respective output dataset, wherein each computer code portion in the second set of computer code portions is associated with a respective node in the second set of nodes; and

a third set of one or more nodes, each node in the third set of nodes representing a respective data processing operation.

28. The data processing system of claim 24, wherein determining the processing layout is performed using the at least one rule for selecting among processing layouts having different degrees of parallelism.

29. The data processing system of claim 26, wherein the processor-executable instructions cause the at least one computer hardware processor to perform:
after determining the processing layout for each node in the dataflow graph,
executing the dataflow graph in accordance with processing layouts determined for each node in the dataflow graph.

30. The data processing system of claim 24, wherein the processor-executable instructions cause the at least one computer hardware processor to perform:
receiving a structured query language (SQL) query;
generating a query plan from the SQL query; and
generating the computer code from the query plan.