What is claimed is:

1. A computer-implemented method of generating a computer program based on an SQL query, including:

receiving a SQL query, including a reference to a first data set stored at a first data source, and including a reference to a second data set stored at a second data source different from the first data source;

determining that the SQL query includes two or more commands, the commands including a first union-type operation, and a first aggregation operation, and determining that the SQL query describes that the first union-type operation shall be applied to at least a portion of data from the first data set, and applied to at least a portion of data from the second data set;

determining that the SQL query describes that the first aggregation operation shall be applied to data resulting from the first union-type operation;

generating, on a first system, a computer program based on the SQL query, the generating including generating at least a first module of the computer program representing a second system, the first module configured for: transmitting instructions that direct a second system to perform a second aggregation operation on a portion of data from the first data set, and receiving, from the second system, a particular data set that is a summary of the portion of data from the first data set; and

generating at least a second module of the computer program configured for: receiving output originating from the first module, and performing a second union-type operation on the particular data set and a summary of a portion of data from the second data set.

2. The method of claim 1 in which the first system includes a graph-based processing system, and the second system includes a relational database management system.

3. The method of claim 1 in which the computer program is further configured for: transmitting instructions that direct a third system to perform a second

aggregation operation on a portion of data from the second data set; and receiving, from the third system, a particular data set that is the summary of the portion of data from the second data set.

4. The method of claim 1 in which the computer program includes: a first component representing the union-type operation of the SQL query that is applied to the data sets that are the summary of the portion of data from the first data set and the summary of the portion of data from the second data set, where the summaries each contain fewer records than their respective corresponding portions of data, and a second component representing the first data set, where the second component, when executed, causes the instructions to be transmitted.

5. The method of claim 4, the computer program including a third component representing the second data set, and a fourth component representing a third aggregation operation, where output of the third component flows to the fourth component, and output of the fourth component flows to the first component.

6. The method of claim 4 in which the output of the second component includes the particular data set that is the summary of the portion of data from the first data set.

7. The method of claim 4 in which the component representing the first data source transmits the instructions, wherein the instructions are carried out before the first system causes data to be transmitted to a system executing the computer program.

8. The method of claim 1 in which the instructions include a SQL statement that includes a command corresponding to the second aggregation operation.

9. The method of claim 1 in which generating the computer program based on the SQL query includes generating a query plan based on the SQL query; adding, to the query plan, the second aggregation operation that is applied to data from the first data

source; determining that the first system includes functionality for carrying out the second aggregation operation; appending, to a portion of the query plan representing the first data set, instructions to carry out the second aggregation operation; and generating the computer program based on the query plan.

10. The method of claim 1 in which the first aggregation operation is a minimum, maximum, sum, or count operation, and the second aggregation operation is of the same type of aggregation operation as the first aggregation operation.

11. The method of claim 1 in which the first aggregation operation is an average operation, and the second aggregation operation is a sum operation or a count operation.

12. The method of claim 1, including determining the second aggregation operation based on first aggregation operation.

13. The method of claim 12 in which determining the second aggregation operation based on first aggregation operation includes determining a sequence of operations that, when carried out, produces output equivalent to output of the first aggregation operation, and selecting the second aggregation operation from the sequence of operations.

14. The method of claim 1 in which the first union-type operation is of the same type to the second union-type operation.

15. The method of claim 1, wherein generating the first module includes configuring parameters of the first module, the parameters being values that affect the behavior of the modules of the generated computer program.

16. A system for generating a computer program based on an SQL query, including: one or more processors configured for performing the steps of the method of any one of claims 1 to 15.

17. A non-transitory computer-readable medium storing executable code that, when executed, causes a computer system to carry out operations for generating a computer program based on an SQL query, including performing the steps of the method of any one of claims 1 to 15.

18. A system for generating a computer program based on an SQL query, including:

means for receiving a SQL query, including a reference to a first data set stored at a first data source, and including a reference to a second data set stored at a second data source different from the first data source;

means for determining that the SQL query includes two or more commands, the commands including a first union-type operation, and a first aggregation operation, and means for determining that the SQL query describes that the first union-type operation shall be applied to at least a portion of data from the first data set, and applied to at least a portion of data from the second data set;

means for determining that the SQL query describes that the first aggregation operation shall be applied to data resulting from the first union-type operation;

means for generating, on a first system, a computer program based on the SQL query, the generating including generating at least a first module of the computer program representing a second system, the first module configured for: transmitting instructions that direct a second system to perform a second aggregation operation on a portion of data from the first data set, and receiving, from the second system, a particular data set that is a summary of the portion of data from the first data set;

and means for generating at least a second module of the computer program configured for: receiving output originating from the first module, and performing a second union-type operation on the particular data set and a summary of a portion of data from the second data set.