

SOURCE CODE TRANSLATION

CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims priority to U.S. Application Serial No. 61/912,594, filed on December 6, 2013.

5

BACKGROUND

This description relates to source code translation, and in particular, to translation of source code specified in one or more original software programming languages to one or more other, different software programming languages.

In the field of software development, software engineers can choose to
10 develop software in one or more of a number of different programming languages. At the time of this writing, some examples of modern programming languages conventionally used by developers are Java, C#, and C++. In general, each programming language has its advantages and disadvantages and it is the job of the software engineer to take these advantages and disadvantages into account when
15 choosing an appropriate programming language for a given application.

Over the years, the state of the art in programming languages has advanced, causing certain early programming languages to become less used, no longer supported, and/or obsolete. Some examples of such early programming languages are Basic and Fortran. Still, source code written in those early programming languages,
20 often referred to as “legacy” code, commonly remains in production for many years due to its adequate performance. However, when such legacy code ceases to adequately function and changes to the code become necessary, it can be difficult to find software engineers who have the necessary skills to update the legacy code.

For this reason, source-to-source compilers have been developed, which
25 receive a first software specification specified in a first programming language as input and generate a second software specification specified in a second, different programming language as output. Such source-to-source compilers are used to translate legacy code into modern programming languages, which are more easily edited by software engineers skilled in the use of modern programming languages.

SUMMARY

A technical problem that is solved involves converting between a software specification containing source code in a procedural language and a software specification containing source code in a language that is not restricted to procedural programming constructs but operates using a different modality. For example, instead of execution being driven solely by control explicitly passing between different procedures, the language may operate in a modality that involves data flowing between different programming entities to drive execution, alone or in combination with explicit control flow. Conversion of source code between languages with such fundamental differences involves more than mere transliteration between different styles of languages. For systems with source code in multiple languages, another technical problem that is solved involves providing source code for a new system that incorporates features of those multiple languages into a different language.

In one aspect, in general, a method for software specification translation includes: receiving a first software specification specified in a first programming language (e.g., COBOL); receiving a second software specification specified in a second programming language (e.g., COBOL); receiving a third software specification specified in a third programming language (in some embodiments, the third programming language being different from the first and second programming languages, e.g., JCL), the third software specification defining one or more data relationships between the first software specification and the second software specification; forming a representation of the first software specification in a fourth programming language (e.g., a dataflow graph) different from the first, second, and third programming languages; forming a representation of the second software specification in the fourth programming language, analyzing the third software specification to identify the one or more data relationships; and forming a combined representation of the first software specification and the second software specification in the fourth programming language including forming connections in the fourth programming language between the representation of the first software specification in the fourth programming language and the representation of the second software specification in the fourth programming language according to the identified one or more data relationships.

Aspects can include one or more of the following features.

The first programming language is a procedural programming language.

The fourth programming language enables parallelism between different portions of a software specification.

The fourth programming language enables a plurality of types of parallelism including: a first type of parallelism enabling multiple instances of a portion of a software specification to operate on different portions of an input data stream; and a second type of parallelism enabling different portions of a software specification to execute concurrently on different portions of the input data stream.

The second programming language is a procedural programming language.

The second programming language is the same as the first programming language.

The one or more data relationships between the first software specification and the second software specification include at least one data relationship that corresponds to the first software specification receiving data from a first dataset and the second software specification providing data to the first dataset.

The fourth programming language is a dataflow graph-based programming language.

The connections in the fourth programming language correspond to directed links representing flows of data.

The first software specification is configured to interact with one or more datasets, each dataset having an associated dataset type of a plurality of dataset types in the first software specification, and the second software specification is configured to interact with one or more datasets, each data set having an associated type of the plurality of dataset types in the second software specification, the method further including: processing the first software specification, the processing including: identifying the one or more datasets of the first software specification, and for each of the identified one or more datasets, determining the associated type of the dataset in the first software specification; and forming a representation of the first software specification in the fourth programming language, including, for each of the identified one or more datasets, forming a specification of the dataset in the fourth programming language, the specification of the dataset in the fourth programming language having a type corresponding to the associated type of the dataset in the first programming language; wherein at least one of the specifications of the one or more datasets in the fourth programming language has: an input dataset type or an output dataset type; processing the second software specification, the processing including: identifying the

one or more datasets of the second software specification and for each of the identified one or more datasets, determining the associated type of the dataset in the second software specification; and forming a representation of the second software specification in the fourth programming language, including, for each of the identified
5 one or more datasets, forming a specification of the dataset in the fourth programming language, the specification of the dataset in the fourth programming language having a type corresponding to the associated type of the dataset in the first programming language; wherein at least one of the specification of the one or more datasets in the fourth programming language enables: an input function or an output function.

10 Forming the combined representation includes at least one of: forming one or more connections to replace connections between the specifications of the one or more datasets of the second software specification in the fourth programming language enabling input functions and the representation of the second software specification in the fourth programming language with connections between the
15 representation of the first software specification in the fourth programming language and the representation of the second software specification in the fourth programming language; or forming one or more connections to replace connections between the specification of the one or more datasets of the first software specification in the fourth programming language enabling input functions and the representation of the
20 first software specification in the fourth programming language with connections between the representation of the second software specification in the fourth programming language and the representation of the first software specification in the fourth programming language.

The method further includes: preserving the one or more datasets of the first
25 software specification in the fourth programming language enabling output functions in the representation of the first software specification in the fourth programming language, or preserving the one or more datasets of the second software specification in the fourth programming language enabling output functions in the representation of the second software specification in the fourth programming language.

30 The first software specification includes one or more data transformation operations and analyzing the first software specification includes identifying at least some of the one or more data transformation operations and classifying the identified data transformation operations into a corresponding data transformation type of the fourth programming language, and forming the representation of the first software

specification in the fourth programming language includes, for each of the identified data transformation operations, forming a specification of the data transformation operation in the fourth programming language, the specification of the data transformation operation in the fourth programming language enabling a data transform operation corresponding to the data transformation type of the identified data transformation operation in the first programming language.

At least one of the specifications of the one or more datasets in the fourth programming language has a read-only random access dataset type.

Determining the associated type of the dataset in the first software specification includes analyzing parameters of dataset definitions and commands that access the dataset.

The parameters include one or more of a file organization associated with the dataset, an access mode associated with the dataset, a mode used to open the dataset, and input-output operations.

The method further includes: storing the combined representation of the first software specification and the second software specification in a storage medium.

The first software specification defines one or more data processing operations that interact with one or more datasets, and the second software specification defines one or more data processing operations that interact with one or more datasets.

The third software specification defines one or more data relationships between the one or more datasets of the first software specification and the one or more datasets of the second software specification.

In another aspect, in general, software is stored in a non-transitory form on a computer-readable medium, for software specification translation. The software includes instructions for causing a computing system to: receive a first software specification specified in a first programming language; receive a second software specification specified in a second programming language; receive a third software specification specified in a third programming language different from the first and second programming languages, the third software specification defining one or more data relationships between the first software specification and the second software specification; form a representation of the first software specification in a fourth programming language different from the first, second, and third programming languages; form a representation of the second software specification in the fourth programming language; analyze the third software specification to identify the one or

more data relationships; and form a combined representation of the first software specification and the second software specification in the fourth programming language including forming connections in the fourth programming language between the representation of the first software specification in the fourth programming language and the representation of the second software specification in the fourth programming language according to the identified one or more data relationships.

In another aspect, in general, a computing system for software specification translation includes: an input device or port configured to receive software specifications, the software specifications including: a first software specification specified in a first programming language; a second software specification specified in a second programming language; a third software specification specified in a third programming language different from the first and second programming languages, the third software specification defining one or more data relationships between the first software specification and the second software specification; and at least one processor configured to process the received software specifications, the processing including: forming a representation of the first software specification in a fourth programming language different from the first, second, and third programming languages; forming a representation of the second software specification in the fourth programming language; analyzing the third software specification to identify the one or more data relationships; and forming a combined representation of the first software specification and the second software specification in the fourth programming language including forming connections in the fourth programming language between the representation of the first software specification in the fourth programming language and the representation of the second software specification in the fourth programming language according to the identified one or more data relationships.

Aspects can include one or more of the following advantages.

Converting a program based on identifying certain data relationships between different program specifications enables formation of a combined specification that can be more efficiently executed in various contexts, such as in data processing systems. For example, by converting a program written in one or more procedural programming languages into a dataflow graph representation, component parallelism, data parallelism, and pipeline parallelism are enabled. For component parallelism, a dataflow graph includes multiple components interconnected by directed links

representing flows of data (or “dataflows”) between the components, and components in different parts of the dataflow graph are able to run simultaneously on separate flows of data. For data parallelism, a dataflow graph processes data divided into segments (or “partitions”) and multiple instances of a component are able to operate
5 on each segment simultaneously. For pipeline parallelism, components in a dataflow graph that are connected by a dataflow link are able to run simultaneously as the upstream component adds data onto that dataflow and the downstream component receives data from that dataflow.

Converting a program (or specifications of at least some portions of the
10 program) written in a procedural programming language into a dataflow graph representation of the program may enable the execution of different components of the dataflow graph representation on different servers.

Intermediate datasets which may be required by a program written in a procedural programming language (due to its non-parallel nature) can be eliminated
15 from the dataflow graph by converting to a dataflow graph representation of the program and replacing the intermediate datasets with flows of data. In some examples, the intermediate datasets are taken out of the path of data flowing through the dataflow graph and are preserved to ensure that any other programs using the datasets can still access the data included in the dataset. In some examples,
20 eliminating intermediate datasets can reduce storage and I-O traffic requirements.

Converting a program written in one or more procedural programming languages into a dataflow graph representation enables visualization of data lineage through the program.

The dataflow programming language is agnostic to database type. Thus,
25 converting a program written in a procedural programming language into a dataflow graph representation of the program may enable use of the program with types of databases that were not originally supported by the program as written in the procedural programming language. That is, approaches may abstract inputs and outputs in the code (e.g., JCL/COBOL code) into flows which can be connected to
30 many different types of sources and sinks (e.g., queues, database tables, files, etc.).

Converting a program written in a procedural programming language into a dataflow graph representation of the program may enable the use of reusable user defined data types. This is advantageous when compared to some procedural programming languages such as COBOL which makes no clear distinction between

data types (i.e., metadata) and storage allocation but instead combines the two in the Data Division. The approaches described herein extract metadata from COBOL source code and create reusable data types (e.g., DML data types) and type definition files. The reusable data types and type definition files can be used for storage
5 allocation at the top of a procedural transform as well as for port and lookup file record definitions. In some examples, the extracted data types (e.g., data type metadata from COBOL) in conjunction with the extracted datasets (e.g., dataset metadata from JCL) can also be used to consolidate partial descriptions of data (i.e., partial metadata) from multiple programs that access the same dataset into a
10 comprehensive description of the data.

Converting a program written in one or more procedural programming languages into a dataflow graph representation enables simplified editing of the program thorough a dataflow graph based graphical development environment.

Other features and advantages of the invention will become apparent from the
15 following description, and from the claims.

DESCRIPTION OF DRAWINGS

FIG. 1 is a block diagram of a system including a software translation module.
FIG. 2 is a schematic example of a software specification.
FIG. 3 is a block diagram of a top-level software translation module.
20 FIG. 4 is a block diagram of a software translation module.
FIG. 5 is a table of dataset functions and their associated combinations of dataset organization types, access modes, and open modes.
FIG. 6 is a first procedural transform.
FIG. 7 is a second procedural transform.
25 FIG. 8 is a third procedural transform.
FIG. 9 is a dataflow graph representation of a program.
FIG. 10 illustrates the creation of a composite dataflow graph.
FIG. 11 is an operational example of the top-level software translation module of FIG. 3.
30 FIG. 12 is an operational example of the software translation module of FIG. 4.
FIG. 13 is a composite dataflow graph.

DESCRIPTION

FIG. 1 shows an example of a data processing system 100 in which programs can be translated using the source code translation techniques described herein. The translated programs can be executed to process data from a data source 102 of the data processing system 100. A translation module 120 accepts a first software specification 122 in one or more procedural programming languages as input and processes the software specification 122 to generate a composite dataflow graph representation 332 of the first software specification 122 in a dataflow based programming language. The dataflow graph representation 332 of the first software specification 122 is stored in a data storage system 116, from which it can be presented visually within a development environment 118. A developer 120 can verify and/or modify the dataflow graph representation 332 of the first software specification 122 using the development environment 118.

The system 100 includes a data source 102 that may include one or more sources of data such as storage devices or connections to online data streams, each of which may store or provide data in any of a variety of formats (e.g., database tables, spreadsheet files, flat text files, or a native format used by a mainframe). An execution environment 104 includes a loading module 106 and an execution module 112. The execution environment 104 may be hosted, for example, on one or more general-purpose computers under the control of a suitable operating system, such as a version of the UNIX operating system. For example, the execution environment 104 can include a multiple-node parallel computing environment including a configuration of computer systems using multiple central processing units (CPUs) or processor cores, either local (e.g., multiprocessor systems such as symmetric multi-processing (SMP) computers), or locally distributed (e.g., multiple processors coupled as clusters or massively parallel processing (MPP) systems, or remote, or remotely distributed (e.g., multiple processors coupled via a local area network (LAN) and/or wide-area network (WAN)), or any combination thereof.

The loading module 106 loads the dataflow graph representation 332 into an execution module 112, from which it is executed for processing data from the data source 102. Storage devices providing the data source 102 may be local to the execution environment 104, for example, being stored on a storage medium connected to a computer hosting the execution environment 104 (e.g., hard disk drive 108), or may be remote to the execution environment 104, for example, being hosted on a

remote system (e.g., mainframe 110) in communication with a computer hosting the execution environment 104, over a remote connection (e.g., provided by a cloud computing infrastructure). The dataflow graph representation 332 being executed by the execution module 104 can receive data from a variety of types of systems that
5 may embody the data source 102, including different forms of database systems. The data may be organized as records (also called “rows”) having values for respective fields (also called “attributes” or “columns”), including possibly null values. When first reading data from a data source, the dataflow graph representation 332 typically starts with some initial format information about records in that data source. In some
10 circumstances, the record structure of the data source may not be known initially and may instead be determined after analysis of the data source or the data. The initial information about records can include, for example, the number of bits that represent a distinct value, the order of fields within a record, and the type of value (e.g., string, signed/unsigned integer) represented by the bits.

15 The dataflow graph representation 332 may be configured to generate output data, which may be stored back in the data source 102 or in the data storage system 116 accessible to the execution environment 104, or otherwise used. The data storage system 116 is also accessible to the development environment 118. The development environment 118 is, in some implementations, a system for developing applications as
20 dataflow graphs that include vertices (representing data processing components or datasets) connected by directed links (also called “flows,” representing flows of work elements, i.e., data) between the vertices. For example, such an environment is described in more detail in U.S. Publication No. 2007/0011668, titled “Managing Parameters for Graph-Based Applications,” incorporated herein by reference. A
25 system for executing such graph-based computations is described in U.S. Patent 5,966,072, titled “EXECUTING COMPUTATIONS EXPRESSED AS GRAPHS,” incorporated herein by reference. Dataflow graphs made in accordance with this system provide methods for getting information into and out of individual processes represented by graph components, for moving information between the processes, and
30 for defining a running order for the processes. This system includes algorithms that choose interprocess communication methods from any available methods (for example, communication paths according to the links of the graph can use TCP/IP or UNIX domain sockets, or use shared memory to pass data between the processes).

1 Software Specification

In some examples, the first software specification 122 is specified using one or more procedural text based programming languages such as C, C++, Java, C#, IBM's Job Control Language (JCL), COBOL, Fortran, Assembly, and so on. For some of the examples below, the software specification 122 includes a batch processing script written using the JCL scripting language and a number of programs written using the COBOL programming language. The JCL script references the COBOL programs and imposes a decision-based controlled flow of execution. It should be appreciated that the first software specification 122 is not limited to a combination of JCL and COBOL programming languages, and that this combination of programming languages is simply used to illustrate one exemplary embodiment of the translation module 120.

Referring to FIG. 2, a schematic view of one example of the software specification 122 of FIG. 1 includes a JCL script 226 including a number of steps 230, some of which execute COBOL programs 228. Other possible steps of the JCL script 226 are omitted to simplify this description. Each step in the JCL script that executes a COBOL program specifies the name of the COBOL program (e.g., COBOL1) and the datasets on which the COBOL program operates. For example, step 3 of the JCL script executes the COBOL program called "COBOL1" on the "DS1.data" and "DS2.data" datasets. In the JCL script 226, each dataset associated with a given COBOL program is assigned a file handle (also referred to as a "DD name"). For example, in FIG. 2, "DS1.data" is assigned the file handle "A" and "DS2.data" is assigned the file handle "B." Each of the COBOL programs 228 includes source code (written in the COBOL programming language) for operating on the datasets specified by the JCL script 226. The file handle (i.e., the DD name) for a given dataset is an identifier that is used by both the JCL script 226 and the code of the COBOL program to identify the dataset.

In operation, a conventional job scheduler running on, for example, an IBM mainframe computer accesses the JCL script 226 and sequentially (i.e., one at a time) executes the steps 230 of the script according to a control flow defined by the JCL script 226. In general, any COBOL programs that access input or output datasets do so by reading from or writing to a storage medium storing the datasets (e.g., a storage medium of the data source 102 or data storage system 116, such as a hard disk drive,

referred to simply as “disk”). In general, each COBOL program executed by the JCL script 226 reads all of its input data from disk and writes all of its output data to disk before passing control back to the JCL script 226. Consequently, any steps that rely on the output of previous steps for input data generally must read the input data from disk.

2 Translation Module

Referring to FIG. 3, one example of the translation module 120 of FIG. 1 receives the software specification 122 including the JCL script 226 and the COBOL programs 228 referenced by the JCL script 226 as input and processes the software specification 122 to generate a composite dataflow graph 332 that implements the same functionality as the first software specification 122 and is usable by the execution environment 104 of FIG. 1. The translation module 120 includes a COBOL to dataflow graph translator 334 and a composite graph synthesizer 336.

Very generally, the COBOL to dataflow graph translator 334 receives the COBOL programs 228 as input and converts each COBOL program into a separate dataflow graph representation 338 of the COBOL program. As is described in greater detail below, each dataflow graph representation 338 of a COBOL program includes a dataflow graph component referred to as a “procedural transform” component and zero or more datasets and/or other dataflow graph components. The procedural transform component includes ports such as input ports and output ports for connecting the procedural transform component to the datasets and other components of the dataflow graph representation 338 of the COBOL program and performs some or all of the functionality of the COBOL program. In some examples, the dataflow graph representation of the COBOL program includes dataflow graph components that are analogous to commands present in the COBOL program. In some examples, dataflow graph representations 338 of COBOL programs can be implemented as “sub-graphs” which have input ports and output ports for forming flows between the dataflow graph representation 338 instances of the COBOL programs and other dataflow graph components (e.g., other dataflow graph components of the composite dataflow graph 332 of FIG. 3).

The JCL script 226 and the dataflow graph representations 338 of the COBOL programs are provided to the composite graph synthesizer 336, which analyzes the JCL script 226 to determine dataflow interconnections between the COBOL programs

and any other components. The composite graph synthesizer 336 then synthesizes the composite dataflow graph 332 by joining the input/output ports of the dataflow graph representations of the COBOL programs 338 using flows according to the determined dataflow interconnections. The composite graph synthesizer 336 determines the
5 dataflow interconnections between the COBOL programs by identifying “intermediate” datasets that are written to by an earlier step of the JCL and read by a later step of the JCL. In some examples, the intermediate datasets can be eliminated and replaced by data flows between the components in the composite dataflow graph 332. Due to pipeline parallelism, significant performance improvements can be
10 achieved by allowing data to flow directly between components without performing the intermediate steps of writing to and reading from disk. It is noted that the term “eliminated” as used above does not necessarily mean that the intermediate dataset is deleted. In some examples, the intermediate dataset is taken out of the path of data flowing through the dataflow graph but is still written to disk to ensure that other
15 programs (e.g., those executed from other JCL scripts) depending on the intermediate dataset can still access its data. Where intermediate files can be entirely eliminated (because the JCL deletes them after their use), the dataflow graph representation will also lower storage capacity requirements.

In some examples, the sequential nature of certain steps in the JCL code can
20 be ignored, yielding component parallelism in the composite dataflow graph 332. In other examples, for steps where the output of one step is provided as an input to another step, the sequential nature of the steps is preserved by connecting the respective components for the steps using a flow, resulting in pipeline parallelism.

2.1 COBOL to Dataflow Graph Translator

Referring to FIG. 4, a detailed block diagram of an implementation of the
25 COBOL to dataflow graph translator 334 receives a number of COBOL programs 228 as input and processes the COBOL programs 228 to generate a number of dataflow graph representations 338 of the COBOL programs. The COBOL to dataflow graph translator 334 includes a COBOL parser 440, an internal component analyzer 444, a
30 dataset function analyzer 442, a metadata analyzer 441, a SQL analyzer 443, a procedure division translator 445, and a sub-graph synthesizer 446.

Each COBOL program 228 is first provided to the COBOL parser 440 which parses the COBOL program 228 to generate a parse tree. The parse tree generated by

the COBOL parser 440 is then passed to the internal component analyzer 444, the dataset function analyzer 442, the metadata analyzer 441, and the SQL analyzer 443.

The internal component analyzer 444 analyzes the parse tree to identify program processes that have an analogous dataflow graph component in the dataflow graph programming language (e.g., internal sort). Some examples of COBOL operations that can be converted to dataflow graph components are “internal sort” and “internal recirculate” (temporary storage) operations. An internal sort operation corresponds to a component with an input port receiving a flow of unsorted data and an output port providing a flow of sorted data, with the input and output ports linked to a main component, as described in more detail below. An internal recirculate operation corresponds to an intermediate file that is first sequentially written in its entirety, and then read in its entirety within a COBOL program. The output of the dataset function analyzer 444 is an internal components result 448 including a listing of the identified operations along with their corresponding locations in the COBOL parse tree.

The above is applicable to any procedural languages where a statement or a sequence of statements and/or operations can be identified, where the statements and/or operations perform a particular transformation on a series of records in a flow that correspond to a component or sub-graph that receives the flow at an input port and provides transformed records from an output port.

The dataset function analyzer 442 analyzes the parse tree to identify all of the data sources and sinks (e.g., datasets) that are accessed (e.g., opened, created, written to, or read from) by the COBOL program 228 and determine a type associated with the dataset for the COBOL program. To do so, the dataset function analyzer 442 identifies and analyzes COBOL statements (e.g., OPEN, READ, WRITE, DELETE, etc.) that access the dataset. In some examples, the possible types that can be associated with a dataset include: INPUT, OUTPUT, LOOKUP, and UPDATABLE LOOKUP. COBOL definitions specify a handle to or the path of the dataset, a file organization of the dataset, and an access mode for the dataset, with additional information such as file open mode(s) determined from Input-Output statements.

Possible dataset file organizations include: SEQUENTIAL, INDEXED, and RELATIVE. A dataset with SEQUENTIAL organization includes records that can only be accessed sequentially (i.e., in the order that they were originally written to the dataset). A dataset with INDEXED organization includes records that are each

associated with one or more indexed key(s). Records of an INDEXED dataset can be randomly accessed using a key, or sequentially from any given position in the file. A dataset with RELATIVE organization has record slots that are numbered with positive integers, with each slot either marked as empty or containing a record. When a file
5 with RELATIVE organization is read sequentially, empty slots are skipped. The records of a RELATIVE file can be accessed directly using the slot number as a key. The notion of 'file position' is common to the three file organizations.

Possible access modes include: SEQUENTIAL, RANDOM, and DYNAMIC. SEQUENTIAL access mode indicates that records in the dataset are accessed
10 sequentially in entry, ascending, or descending key order. RANDOM access mode indicates that records in the dataset are accessed using a record identifying key. DYNAMIC access mode indicates that records in the dataset can be accessed directly using a record identifying key, or sequentially from any selected file position.

Possible open modes include: INPUT, OUTPUT, EXTEND, and I-O. INPUT
15 open mode indicates that the dataset is opened as an input dataset. OUTPUT open mode indicates that an empty dataset is opened as an output dataset. EXTEND open mode indicates that a dataset including preexisting records is opened as an output dataset to which new records are appended. I-O open mode indicates that the dataset open mode supports both input and output dataset operations (regardless of whether
20 such operations are present in the program).

The dataset function analyzer 442 applies the following set of rules to the file organization, access mode, and open mode of the COBOL dataset access command to determine the function associated with the dataset for the COBOL program:

- **OUTPUT** datasets are datasets with SEQUENTIAL, INDEXED, or
25 RELATIVE organization, SEQUENTIAL, RANDOM, or DYNAMIC access mode, and OUTPUT or EXTEND open mode.
- **INPUT** datasets are datasets with INDEXED, RELATIVE, or SEQUENTIAL organization, SEQUENTIAL access mode, and INPUT open mode.
- **LOOKUP** datasets are datasets with INDEXED or RELATIVE
30 organization, RANDOM or DYNAMIC access mode, and INPUT open mode.

- **UPDATEABLE LOOKUP** datasets are datasets with INDEXED or RELATIVE organization, RANDOM or DYNAMIC access mode, and I-O open mode.

In some examples, an “effective open mode” of a file can be determined by counting the actual Input and Output operations for the file. For example, if a file is opened in I-O mode, but has only WRITE operations and no READ or START operations, the “effective open mode” can be reduced to EXTEND.

Referring to FIG. 5, a table 501 lists the different combinations of dataset organization, access mode, and open mode along with the dataset function associated with each combination.

Referring again to FIG. 4, the output of the dataset function analyzer 442 is a dataset functions result 450 that includes a listing of all datasets accessed by the COBOL program along with their associated function in the COBOL program.

The metadata analyzer 441 analyzes the parse tree to extract metadata and to create reusable data types (e.g. DML data types) and type definition files. The reusable data types are distinct from storage allocation in the COBOL program. The output of the metadata analyzer 441 is a data types result 447.

The SQL analyzer 443 analyzes the parse tree to identify embedded structured query language (SQL) code (or simply “embedded SQL”) in the COBOL program. Any identified embedded SQL is processed into Database Interface Information 449. A database application programming interface (API) for accessing a database may provide primitives that can be used within the Database Interface Information 449. In some examples, the inclusion of these primitives avoids the need to access a particular database using a particular schema to compile portions of the embedded SQL (e.g., into a binary form that is operated on using binary operations). Instead, some of the efficiency that would be provided by such compilation can be traded off for flexibility in being able to interpret the embedded SQL at runtime using the appropriate API primitives placed within the Database Interface Information 449, potentially using a different database and/or schema as needed.

The parse tree for the COBOL program is then provided to the procedure division translator 445 along with the internal components result 448, the dataset functions result 450, the data types result 447, and the Database Interface Information result 449. The procedure division translator 445 analyzes the parse tree to translate

the COBOL logic into a “procedural transform” dataflow graph component 452. In general, the procedural transform dataflow graph component 452 is a container type component that contains some or all of the COBOL logic associated with the COBOL program and has input and output ports to accept input data and provide output data
5 from the component, respectively. In cases where the COBOL code includes code from a different programming language (e.g., SQL code is identified by the SQL analyzer 443 and provided in the Database Interface Information result 449) the procedure division translator 445 uses the Database Interface Information result 449 to generate the appropriate representation of that embedded code within the
10 procedural transform dataflow graph component 452. In some examples, the procedure division translator 445 uses a database API to generate the appropriate representation of embedded code. In other examples, embedded SQL Tables and Cursors are replaced with Input Table components, thereby replacing FETCH operations with calls to read_record(port_number) as is done for files.

15 In some examples, the procedure division translator 445 only generates a file including Data Manipulation Language (DML) code which represents the procedural logic of the COBOL program. The sub-graph synthesizer 446 generates the procedural transform dataflow component that uses the file generated by the procedure division translator 445.

20 It is noted that FIG. 4 and the above description relate to one possible order of operation of the internal component analyzer 444, the dataset function analyzer 442, the metadata analyzer 441, and the SQL analyzer 443. However, the order of operation of the analyzers is not limited to the order described above and other orders of operation of the analyzers are possible.

25 Referring to FIG. 6, one simple example of a procedural transform component 554 titled “COBOL2” (i.e., the result of translating the COBOL program executed at step 5 of the JCL script 226 of Fig. 2) has an input port 556 labeled “in0”, an output port 560 labeled “out0”, and a lookup port 562 labeled “lu0”. It is noted that lookup datasets are not necessarily accessed via a port on the component but may instead be
30 accessed using a lookup dataset API. However, for simplification of the description, lookup datasets are described as being accessed via a lookup port.

Each of the ports is configured to be connected to their respective datasets (as identified by the JCL script 226) through a flow. In some examples, the developer can view and edit the DML translation of the COBOL code underlying the procedural

transform component 554 by, for example, shift double clicking on the component or hovering over the component until an information bubble appears and clicking on a ‘transform’ link in the information bubble.

Referring to FIG. 7, another example of a procedural transform component 664 illustrates a situation where a COBOL program titled “COBOL1” (i.e., the COBOL program executed at step 3 of the JCL script 226 of Fig. 2) includes a sort command in its code. In this situation, the internal component analyzer 448 identifies the sort command and passes information related to the sort command to the procedure division translator 445. The procedure division translator 445 uses the information from the internal component analyzer 448 to replace the sort command in the code associated with the procedural transform 664 with an interface to a specialized internal sort sub-graph. The sub-graph synthesizer 446 uses the sort information created by 448, and creates an output port, out1, from the procedural transform 664 for providing the data to be sorted to the internal sort dataflow sub-graph component 666 and an input, in1, for receiving the sorted data from the internal sort dataflow sub-graph component 666.

Referring to FIG. 8, another similar example of a procedural transform including a sort command is illustrated. In this example, rather than creating a single procedural transform having an output for providing data to be sorted and an input for receiving the sorted data, two procedural transforms are created. A first procedural transform 768 of the two procedural transforms has an output for providing the data to be sorted and a second procedural transform 770 of the two procedural transforms has an input for receiving the sorted data. As is illustrated, in some examples a sort dataflow component 766 can be automatically connected between the two procedural transforms 768, 770 by the sub-graph synthesizer 446. In other examples, a sort dataflow component 766 can be manually connected between the two procedural transforms 768, 770.

2.2 Sub-Graph Synthesizer

Referring again to FIG. 4, the procedural transform 452 for the COBOL program is passed to the sub-graph synthesizer 446 along with the internal components result 448 the dataset functions result 450, the data types result 447, and the Database Interface Information result 449. The sub-graph synthesizer 446 uses the inputs to generate a dataflow graph representation 338 for the COBOL program

228. Very generally, for each COBOL program 228, the sub-graph synthesizer 446 creates a dataflow graph including the procedural transform for the COBOL program 228, the datasets associated with the COBOL program 228, and any internal components identified by the internal components analyzer 444. The sub-graph synthesizer 446 then uses the internal components result 448 and the dataset functions result 450 to appropriately connect flows between the datasets, the internal components, and the procedural transform 452. The sub-graph synthesizer 446 uses the data types result 447 to describe the data flowing through component ports. Referring to FIG. 9, one example of a dataflow graph representation 838 for the exemplary COBOL program titled COBOL1 includes a procedural transform 864 having an input port labeled in0 connected by a flow to an input file with the file handle "A" associated with dataset DS1.data, an output port labeled out0 connected by a flow to an output file with the file handle "B" associated with dataset DS2.data, and output and input ports, out1 and in1 connected by flows to an internal sort component 866.

2.3 Composite Graph Synthesizer

Referring back to FIG. 3, the dataflow graph representations 338 of the COBOL programs are then passed to the composite graph synthesizer 336 along with the JCL script 226. By analyzing the order of execution of the COBOL programs in the JCL script 226 along with the functions of the datasets associated with the COBOL programs, the composite graph synthesizer 336 connects the dataflow graph representations of the COBOL code into a single composite dataflow graph 332.

For example, referring to FIG. 10, a dataflow graph representation of the COBOL program titled COBOL2 reads from an input file "C" associated with dataset DS2.data at an input port labeled in0, enriches the data by accessing a lookup file "D" associated with DS3.data at a lookup port lu0, and writes to an output file "E" associated with dataset DS4.data at an output port labeled out0. A dataflow graph representation of the COBOL program titled COBOL3 reads from two input datasets: "F" associated with DS4.data at an input port labeled in0 and "G" associated with DS5.data at an input port labeled in1 and writes to an output dataset "H" associated with DS6.data at an output port labeled out0. The composite graph synthesizer 336 merges the JCL script 226 information with the information derived by the translation of the COBOL programs to determine that COBOL2 is executed before COBOL3,

and that DS4.data is output by COBOL2 and input by COBOL3, so that the output port labeled out0 of COBOL2 can be connected by a flow to the input port labeled in0 of COBOL3, thereby eliminating the need for COBOL3 to read dataset DS4.data from disk. FIG. 10 illustrates an exemplary composite dataflow graph 932 with a flow
5 connecting the output port of COBOL2 labeled out0 and the input port of COBOL3 labeled in0 through a replicate component 933. The replicate component 933 writes data into DS4.data on disk but also passes the data directly to the input port of COBOL3 labeled in0 via a flow. In this way, COBOL3 can read data flowing from COBOL2 without having to wait for the dataset DS4.data to be written to disk, and
10 the data stored in DS4.data, which is not deleted by the JCL script 226 is available to other processes.

In some examples, if a JCL procedure does not delete an intermediate dataset (e.g., file) after it is created, it is possible that the dataset is used by some other process running in the execution environment. In examples where this is the case, the
15 intermediate dataset is preserved in the dataflow graph representation of the JCL procedure (e.g., by using a replicate component as is described above). In some examples, if the JCL procedure does delete the intermediate dataset after it is created, the intermediate dataset is completely eliminated in the dataflow graph representation of the JCL procedure, and no Replicate component is needed for it.

20 In some examples, the metadata of ports that are connected by a flow, as described above for the COBOL2 and COBOL3 dataflow graphs, might not be identical, because the first software specifications used alternate definitions for the same dataset. The Composite Graph Synthesizer 336 can then insert a Redefine Format component on the connecting flow. The presence of such Redefine Format
25 components can later be used to consolidate dataset metadata. Metadata information is derived by the Metadata Analyzer 441 for each Dataflow Graph 338.

3 Exemplary Operation

Referring to FIG. 11, a simple operational example of the translation module
120 receives the JCL script 226 and the four COBOL programs 228 of FIG. 2 as input
30 and processes the inputs to generate a composite dataflow graph 332.

In a first stage of the translation process, the COBOL programs 228 are provided to the COBOL to Dataflow Graph Translator 334, which processes each of the COBOL programs to generate a dataflow graph representation 338a-d of the

COBOL program. In the second stage, the JCL script 226 and the dataflow graph representations 338a-d of the COBOL programs are provided to the composite graph synthesizer 336, which processes the JCL script 226 and the dataflow graph representations 338a-d of the COBOL programs to generate the composite dataflow graph 332.

Referring to FIG. 12, the COBOL to dataflow graph translator 334 processes each of the COBOL programs 228 using the COBOL parser 440, the internal component analyzer 444, the dataset function analyzer 442, the metadata analyzer 441, and the SQL analyzer 443. The outputs generated by the COBOL parser 440, the internal component analyzer 444, the dataset function analyzer 442, the metadata analyzer 441, and the SQL analyzer 443 are provided to the procedure division translator 445, and together with its output to the sub-graph synthesizer 446, which generates the dataflow graph representations 338a-d for each of the COBOL programs.

For the COBOL1 program executed at step 3 of the JCL script 226, the internal component analyzer 444 identified that the program includes an internal sort component. The dataset function analyzer 442 identified that the COBOL1 program accesses one input dataset, "A" and one output dataset, "B". The identified internal sort component, the datasets, and their relationships to the procedural transform for the COBOL1 program are reflected in the dataflow graph representation 338a of the COBOL1 program.

For the COBOL2 program executed at step 5 of the JCL script 226, the internal component analyzer 444 did not identify any internal components and the SQL analyzer 443 did not identify any embedded SQL code. The dataset function analyzer 442 identified that the COBOL2 program accesses one dataset, "C" as an input dataset, another dataset, "E" as an output dataset, and another dataset, "D" as a lookup dataset. The identified datasets and their relationships to the procedural transform for the COBOL2 program are reflected in the dataflow graph representation 338b of the COBOL2 program.

For the COBOL3 program executed at step 6 of the JCL script 226, the internal component analyzer 444 did not identify any internal components and the SQL analyzer 443 did not identify any embedded SQL code. The dataset function analyzer 442 identified that the COBOL3 program accesses two datasets, "F" and "G" as input datasets and one dataset, and "H" as an output dataset. The identified

datasets and their relationships to the procedural transform for the COBOL3 program are reflected in the dataflow graph representation 338c of the COBOL3 program.

For the COBOL4 program executed at step 10 of the JCL script 226, the internal component analyzer 444 did not identify any internal components and the SQL analyzer 443 did not identify any embedded SQL code. The dataset function analyzer 442 identified that the COBOL4 program accesses one dataset, “I” as an input dataset and another dataset, “J” as an output dataset. The identified datasets and their relationships to the procedural transform for the COBOL4 program are reflected in the dataflow graph representation 338d of the COBOL4 program.

Referring again to FIG. 11, the JCL script 226 and the dataflow graph representations 338a-d of the four COBOL programs are provided to the composite graph synthesizer 336 which analyzes the JCL script 226 and the dataflow graph representations 338a-d to connect the dataflow graph representations 338a-d into the single composite graph 332. Referring to FIG. 13, the composite graph for the JCL script 226 and the four COBOL programs 228 of FIG. 2 includes four procedural transforms COBOL1 452a, COBOL2 452b, COBOL3 452c, and COBOL4 452d interconnected by flows. Replicate components 933 are used to set aside (i.e., write as output datasets) a number of intermediate datasets (i.e., DS2.data, DS4.data, and DS5.data) in the composite dataflow graph 332, directly connecting the components using flows.

4 Alternatives

While the above description describes only a limited number of operations and elements of a program written in a procedural programming language being translated into dataflow graph components, in some examples, all of the source code of the original programs (e.g., the COBOL programs) is translated into a dataflow graph representation.

The above-described system can be used to translate a software specification including any combination of one or more procedural programming languages into a dataflow graph representation of the software specification.

In some examples, the above-described translation module may encounter translation tasks that it is unprepared to process. In such examples, the translation module outputs a list of incomplete translation tasks that the developer is able to read and use to manually repair the translation.

While the above description describes certain modules of the COBOL to dataflow graph translator 334 as running in parallel, this is not necessarily the case. In some examples, the metadata analyzer 441 first receives the parse tree from the COBOL parser 440. The metadata analyzer 441 enriches and/or simplifies the parse
5 tree and provides it to the dataset function analyzer 442. The dataset function analyzer 442 enriches and/or simplifies the parse tree and provides it to the SQL analyzer 443. The SQL analyzer 443 enriches and/or simplifies the parse tree and provides it to the internal component analyzer 444. The internal component analyzer 444 enriches and/or simplifies the parse tree and provides it to the procedure division
10 translator 445. That is the components operate on the parse tree serially.

5 Implementations

The source code translation approach described above can be implemented, for example, using a programmable computing system executing suitable software instructions or it can be implemented in suitable hardware such as a field-
15 programmable gate array (FPGA) or in some hybrid form. For example, in a programmed approach the software may include procedures in one or more computer programs that execute on one or more programmed or programmable computing system (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
20 volatile and/or non-volatile memory and/or storage elements), at least one user interface (for receiving input using at least one input device or port, and for providing output using at least one output device or port). The software may include one or more modules of a larger program, for example, that provides services related to the design, configuration, and execution of dataflow graphs. The modules of the program
25 (e.g., elements of a dataflow graph) can be implemented as data structures or other organized data conforming to a data model stored in a data repository.

The software may be provided on a tangible, non-transitory medium, such as a CD-ROM or other computer-readable medium (e.g., readable by a general or special purpose computing system or device), or delivered (e.g., encoded in a propagated
30 signal) over a communication medium of a network to a tangible, non-transitory medium of a computing system where it is executed. Some or all of the processing may be performed on a special purpose computer, or using special-purpose hardware, such as coprocessors or field-programmable gate arrays (FPGAs) or dedicated,

application-specific integrated circuits (ASICs). The processing may be implemented in a distributed manner in which different parts of the computation specified by the software are performed by different computing elements. Each such computer program is preferably stored on or downloaded to a computer-readable storage medium (e.g., solid state memory or media, or magnetic or optical media) of a storage device accessible by a general or special purpose programmable computer, for configuring and operating the computer when the storage device medium is read by the computer to perform the processing described herein. The inventive system may also be considered to be implemented as a tangible, non-transitory medium, configured with a computer program, where the medium so configured causes a computer to operate in a specific and predefined manner to perform one or more of the processing steps described herein.

A number of embodiments of the invention have been described. Nevertheless, 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 following claims. Accordingly, other embodiments are also within the scope of the following claims. For example, various modifications may be made without departing from the scope of the invention. Additionally, some of the steps described above may be order independent, and thus can be performed in an order different from that described.

What is claimed is:

1. A method for software specification translation, the method including: receiving a first software specification specified in a first programming language, wherein the first programming language is a procedural programming language; receiving a second software specification specified in a second programming language, wherein the second programming language is a procedural programming language; receiving a third software specification specified in a third programming language different from the first and second programming languages, the third software specification referencing the first and second software specifications and constraining a flow of execution, and the third software specification defining one or more data relationships between the first software specification and the second software specification, where the data relationships are defined based on references to specified datasets; forming a representation of the first software specification in a fourth programming language different from the first, second, and third programming languages, wherein the fourth programming language is a dataflow graph-based programming language; forming a representation of the second software specification in the fourth programming language; analyzing the third software specification to identify the one or more data relationships; and forming a combined representation of the first software specification and the second software specification in the fourth programming language including forming connections in the fourth programming language between the representation of the first software specification in the fourth programming language and the representation of the second software specification in the fourth programming language according to the identified one or more data relationships.
2. The method of claim 1 wherein the fourth programming language enables parallelism between different portions of a software specification.
3. The method of claim 2 wherein the fourth programming language enables a plurality of types of parallelism including: a first type of parallelism enabling multiple instances of a portion of a software specification to operate on different portions of an input data stream; and a second type of parallelism enabling different portions of a software specification to execute concurrently on different portions of the input data stream.

4. The method of claim 1 wherein the second programming language is the same as the first programming language.

5. The method of claim 1 wherein the one or more data relationships between the first software specification and the second software specification include at least one data relationship that corresponds to the first software specification receiving data from a first dataset and the second software specification providing data to the first dataset.

6. The method of claim 1 wherein the connections in the fourth programming language correspond to directed links representing flows of data.

7. The method of claim 1 wherein the first software specification is configured to interact with one or more datasets, each dataset having an associated dataset type of a plurality of dataset types in the first software specification, and the second software specification is configured to interact with one or more datasets, each data set having an associated type of the plurality of dataset types in the second software specification, the method further including: processing the first software specification, the processing including: identifying the one or more datasets of the first software specification, and for each of the identified one or more datasets, determining the associated type of the dataset in the first software specification; and forming a representation of the first software specification in the fourth programming language, including, for each of the identified one or more datasets, forming a specification of the dataset in the fourth programming language, the specification of the dataset in the fourth programming language having a type corresponding to the associated type of the dataset in the first programming language; wherein at least one of the specifications of the one or more datasets in the fourth programming language has: an input dataset type or an output dataset type; processing the second software specification, the processing including: identifying the one or more datasets of the second software specification and for each of the identified one or more datasets, determining the associated type of the dataset in the second software specification; and forming a representation of the second software specification in the fourth programming language, including, for each of the identified one or more datasets, forming a specification of the dataset in the fourth programming language, the specification of the dataset in the fourth programming language having a type corresponding to the associated type of the dataset in the first programming language; wherein at least one of the specification of the one or more datasets in the fourth programming language enables: an input function or an output function.

8. The method of claim 7 wherein forming the combined representation includes at least one of: forming one or more connections to replace connections between the specifications of the one or more datasets of the second software specification in the fourth programming language enabling input functions and the representation of the second software specification in the fourth programming language with connections between the representation of the first software specification in the fourth programming language and the representation of the second software specification in the fourth programming language; or forming one or more connections to replace connections between the specification of the one or more datasets of the first software specification in the fourth programming language enabling input functions and the representation of the first software specification in the fourth programming language with connections between the representation of the second software specification in the fourth programming language and the representation of the first software specification in the fourth programming language.

9. The method of claim 8 further including: preserving the one or more datasets of the first software specification in the fourth programming language enabling output functions in the representation of the first software specification in the fourth programming language, or preserving the one or more datasets of the second software specification in the fourth programming language enabling output functions in the representation of the second software specification in the fourth programming language.

10. The method of claim 7 wherein: the first software specification includes one or more data transformation operations and analyzing the first software specification includes identifying at least some of the one or more data transformation operations and classifying the identified data transformation operations into a corresponding data transformation type of the fourth programming language, and forming the representation of the first software specification in the fourth programming language includes, for each of the identified data transformation operations, forming a specification of the data transformation operation in the fourth programming language, the specification of the data transformation operation in the fourth programming language enabling a data transform operation corresponding to the data transformation type of the identified data transformation operation in the first programming language.

11. The method of claim 7 wherein at least one of the specifications of the one or more datasets in the fourth programming language has a read-only random access dataset type.

12. The method of claim 7 wherein determining the associated type of the dataset in the first software specification includes analyzing parameters of dataset definitions and commands that access the dataset.

13. The method of claim 12 wherein the parameters include one or more of a file organization associated with the dataset, an access mode associated with the dataset, a mode used to open the dataset, and input-output operations.

14. The method of claim 1 further including storing the combined representation of the first software specification and the second software specification in a storage medium.

15. The method of claim 1 wherein the first software specification defines one or more data processing operations that interact with one or more datasets, and the second software specification defines one or more data processing operations that interact with one or more datasets.

16. The method of claim 15 wherein the third software specification defines one or more data relationships between the one or more datasets of the first software specification and the one or more datasets of the second software specification.

17. The method of claim 1, wherein forming the combined representation includes forming at least one of the connections in the fourth programming language to replace a connection between: (1) a specification of at least one dataset of the first software specification in the fourth programming language enabling input functions and (2) a representation of the second software specification in the fourth programming language.

18. Software stored in a non-transitory form on a computer-readable medium, for software specification translation, the software including instructions for causing a computing system to:

receive a first software specification specified in a first programming language, wherein the first programming language is a procedural programming language; receive a second software specification specified in a second programming language, wherein the second programming language is a procedural programming language; receive a third software specification specified in a third programming language different from the first and second programming languages, the third software specification referencing the first and second software specifications and constraining a flow of execution, and the third software specification defining one or more data relationships between the first software specification and the second software specification, where the data relationships are defined based on references to specified datasets; form a representation of the first software specification in a fourth programming language different from the first, second, and third programming languages, wherein the fourth programming language is a dataflow graph-based programming language; form a representation of the second software specification in the fourth programming language; analyze the third software specification to identify the one or more data relationships; and form a combined representation of the first software specification and the second software specification in the fourth programming language including forming connections in the fourth programming language between the representation of the first software specification in the fourth programming language and the representation of the second software specification in the fourth programming language according to the identified one or more data relationships.

19. The software of claim 18 wherein the one or more data relationships between the first software specification and the second software specification include at least one data relationship that corresponds to the first software specification receiving data from a first dataset and the second software specification providing data to the first dataset.

20. The software of claim 18 wherein the connections in the fourth programming language correspond to directed links representing flows of data.

21. The software of claim 18 wherein the first software specification is configured to interact with one or more datasets, each dataset having an associated dataset type of a plurality of dataset types in the first software specification, and the second software specification is configured to interact with one or more datasets, each data set having an associated type of the plurality of dataset types in the second software specification, the software further including instructions for causing a computing system to: process the first software specification, the processing including:

identifying the one or more datasets of the first software specification, and for each of the identified one or more datasets, determining the associated type of the dataset in the first software specification; and forming a representation of the first software specification in the fourth programming language, including, for each of the identified one or more datasets, forming a specification of the dataset in the fourth programming language, the specification of the dataset in the fourth programming language having a type corresponding to the associated type of the dataset in the first programming language; wherein at least one of the specifications of the one or more datasets in the fourth programming language has: an input dataset type or an output dataset type; process the second software specification, the processing including: identifying the one or more datasets of the second software specification and for each of the identified one or more datasets, determining the associated type of the dataset in the second software specification; and forming a representation of the second software specification in the fourth programming language, including, for each of the identified one or more datasets, forming a specification of the dataset in the fourth programming language, the specification of the dataset in the fourth programming language having a type corresponding to the associated type of the dataset in the first programming language; wherein at least one of the specification of the one or more datasets in the fourth programming language enables: an input function or an output function.

22. The software of claim 21 wherein forming the combined representation includes at least one of: forming one or more connections to replace connections between the specifications of the one or more datasets of the second software specification in the fourth programming language enabling input functions and the representation of the second software specification in the fourth programming language with connections between the representation of the first software specification in the fourth programming language and the representation of the second software specification in the fourth programming language; or forming one or more connections to replace connections between the specification of the one or more datasets of the first software specification in the fourth programming language enabling input functions and the representation of the first software specification in the fourth programming language with connections between the representation of the second software specification in the fourth programming language and the representation of the first software specification in the fourth programming language.

23. The software of claim 22 further including instructions for causing a computing system to: preserve the one or more datasets of the first software specification in the fourth programming language enabling output functions in the representation of the first software specification in the fourth programming language, or preserve the one or more datasets of the

second software specification in the fourth programming language enabling output functions in the representation of the second software specification in the fourth programming language.

24. The software of claim 21 wherein: the first software specification includes one or more data transformation operations and analyzing the first software specification includes identifying at least some of the one or more data transformation operations and classifying the identified data transformation operations into a corresponding data transformation type of the fourth programming language, and forming the representation of the first software specification in the fourth programming language includes, for each of the identified data transformation operations, forming a specification of the data transformation operation in the fourth programming language, the specification of the data transformation operation in the fourth programming language enabling a data transform operation corresponding to the data transformation type of the identified data transformation operation in the first programming language.

25. The software of claim 21 wherein at least one of the specifications of the one or more datasets in the fourth programming language has a read-only random access dataset type.

26. The software of claim 21 wherein determining the associated type of the dataset in the first software specification includes analyzing parameters of dataset definitions and commands that access the dataset.

27. The software of claim 26 wherein the parameters include one or more of a file organization associated with the dataset, an access mode associated with the dataset, a mode used to open the dataset, and input-output operations.

28. The software of claim 18 further including instructions for causing a computing system to store the combined representation of the first software specification and the second software specification in a storage medium.

29. The software of claim 18 wherein the first software specification defines one or more

data processing operations that interact with one or more datasets, and the second software specification defines one or more data processing operations that interact with one or more datasets.

30. The software of claim 29 wherein the third software specification defines one or more data relationships between the one or more datasets of the first software specification and the one or more datasets of the second software specification.

31. The software of claim 18, wherein forming the combined representation includes forming at least one of the connections in the fourth programming language to replace a connection between: (1) a specification of at least one dataset of the first software specification in the fourth programming language enabling input functions and (2) a representation of the second software specification in the fourth programming language.

32. A computing system for software specification translation, the computing system including: an input device or port configured to receive software specifications, the software specifications including: a first software specification specified in a first programming language, wherein the first programming language is a procedural programming language; a second software specification specified in a second programming language, wherein the second programming language is a procedural programming language; a third software specification specified in a third programming language different from the first and second programming languages, the third software specification referencing the first and second software specifications and constraining a flow of execution, and the third software specification defining one or more data relationships between the first software specification and the second software specification, where the data relationships are defined based on references to specified datasets; and at least one processor configured to process the received software specifications, the processing including: forming a representation of the first software specification in a fourth programming language different from the first, second, and third programming languages, wherein the fourth programming language is a dataflow graph-based programming language; forming a representation of the second software specification in the fourth programming language; analyzing the third software specification to identify the one or more data relationships; and forming a combined representation of the first software specification and the second software specification in the fourth programming language including forming connections in the fourth programming language between the representation of the first software

specification in the fourth programming language and the representation of the second software specification in the fourth programming language according to the identified one or more data relationships.

33. The computing system of claim 32 wherein the fourth programming language enables parallelism between different portions of a software specification.

34. The computing system of claim 33 wherein the fourth programming language enables a plurality of types of parallelism including: a first type of parallelism enabling multiple instances of a portion of a software specification to operate on different portions of an input data stream; and a second type of parallelism enabling different portions of a software specification to execute concurrently on different portions of the input data stream.

35. The computing system of claim 32 wherein the second programming language is the same as the first programming language.

36. The computing system of claim 32 wherein the one or more data relationships between the first software specification and the second software specification include at least one data relationship that corresponds to the first software specification receiving data from a first dataset and the second software specification providing data to the first dataset.

37. The computing system of claim 32 wherein the connections in the fourth programming language correspond to directed links representing flows of data.

38. The computing system of claim 32 wherein the first software specification is configured to interact with one or more datasets, each dataset having an associated dataset type of a plurality of dataset types in the first software specification, and the second software specification is configured to interact with one or more datasets, each data set having an associated type of the plurality of dataset types in the second software specification, the computing system further configured to: process the first software specification, the processing including: identifying the

one or more datasets of the first software specification, and for each of the identified one or more datasets, determining the associated type of the dataset in the first software specification; and forming a representation of the first software specification in the fourth programming language, including, for each of the identified one or more datasets, forming a specification of the dataset in the fourth programming language, the specification of the dataset in the fourth programming language having a type corresponding to the associated type of the dataset in the first programming language; wherein at least one of the specifications of the one or more datasets in the fourth programming language has: an input dataset type or an output dataset type; process the second software specification, the processing including: identifying the one or more datasets of the second software specification and for each of the identified one or more datasets, determining the associated type of the dataset in the second software specification; and forming a representation of the second software specification in the fourth programming language, including, for each of the identified one or more datasets, forming a specification of the dataset in the fourth programming language, the specification of the dataset in the fourth programming language having a type corresponding to the associated type of the dataset in the first programming language; wherein at least one of the specification of the one or more datasets in the fourth programming language enables: an input function or an output function.

39. The computing system of claim 38 wherein forming the combined representation includes at least one of: forming one or more connections to replace connections between the specifications of the one or more datasets of the second software specification in the fourth programming language enabling input functions and the representation of the second software specification in the fourth programming language with connections between the representation of the first software specification in the fourth programming language and the representation of the second software specification in the fourth programming language; or forming one or more connections to replace connections between the specification of the one or more datasets of the first software specification in the fourth programming language enabling input functions and the representation of the first software specification in the fourth programming language with connections between the representation of the second software specification in the fourth programming language and the representation of the first software specification in the fourth programming language.

40. The computing system of claim 39 further configured to: preserve the one or more datasets of the first software specification in the fourth programming language enabling output functions in the representation of the first software specification in the fourth programming language, or preserve the one or more datasets of the second software specification in the fourth

programming language enabling output functions in the representation of the second software specification in the fourth programming language.

41. The computing system of claim 38 wherein: the first software specification includes one or more data transformation operations and analyzing the first software specification includes identifying at least some of the one or more data transformation operations and classifying the identified data transformation operations into a corresponding data transformation type of the fourth programming language, and forming the representation of the first software specification in the fourth programming language includes, for each of the identified data transformation operations, forming a specification of the data transformation operation in the fourth programming language, the specification of the data transformation operation in the fourth programming language enabling a data transform operation corresponding to the data transformation type of the identified data transformation operation in the first programming language.

42. The computing system of claim 38 wherein at least one of the specifications of the one or more datasets in the fourth programming language has a read-only random access dataset type.

43. The computing system of claim 38 wherein determining the associated type of the dataset in the first software specification includes analyzing parameters of dataset definitions and commands that access the dataset.

44. The computing system of claim 43 wherein the parameters include one or more of a file organization associated with the dataset, an access mode associated with the dataset, a mode used to open the dataset, and input-output operations.

45. The computing system of claim 32 further configured to store the combined representation of the first software specification and the second software specification in a storage medium.

46. The computing system of claim 32 wherein the first software specification defines one or more data processing operations that interact with one or more datasets, and the second software specification defines one or more data processing operations that interact with one or more datasets.

47. The computing system of claim 46 wherein the third software specification defines one or more data relationships between the one or more datasets of the first software specification and the one or more datasets of the second software specification.

48. The computing system of claim 32, wherein forming the combined representation includes forming at least one of the connections in the fourth programming language to replace a connection between: (1) a specification of at least one dataset of the first software specification in the fourth programming language enabling input functions and (2) a representation of the second software specification in the fourth programming language.

49. A computing system for software specification translation, the computing system including: means for receiving software specifications, the software specifications including: a first software specification specified in a first programming language, wherein the first programming language is a procedural programming language; a second software specification specified in a second programming language, wherein the second programming language is a procedural programming language; a third software specification specified in a third programming language different from the first and second programming languages, the third software specification referencing the first and second software specifications and constraining a flow of execution, and the third software specification defining one or more data relationships between the first software specification and the second software specification, where the data relationships are defined based on references to specified datasets; and means for processing the received software specifications, the processing including: forming a representation of the first software specification in a fourth programming language different from the first, second, and third programming languages, wherein the fourth programming language is a dataflow graph-based programming language; forming a representation of the second software specification in the fourth programming language; analyzing the third software specification to identify the one or more data relationships; and forming a combined representation of the first software specification and the second software specification in the fourth programming language including forming connections in the fourth programming language between the representation

of the first software specification in the fourth programming language and the representation of the second software specification in the fourth programming language according to the identified one or more data relationships.

50. The software of claim 49 wherein the fourth programming language enables parallelism between different portions of a software specification.

51. The software of claim 50 wherein the fourth programming language enables a plurality of types of parallelism including: a first type of parallelism enabling multiple instances of a portion of a software specification to operate on different portions of an input data stream; and a second type of parallelism enabling different portions of a software specification to execute concurrently on different portions of the input data stream.

52. The software of claim 49 wherein the second programming language is the same as the first programming language.