

**(12) STANDARD PATENT**  
**(19) AUSTRALIAN PATENT OFFICE**

**(11) Application No. AU 2010208112 B2**

(54) Title  
**Processing data using vector fields**

(51) International Patent Classification(s)  
**G06F 7/00** (2006.01)

(21) Application No: **2010208112** (22) Date of Filing: **2010.01.29**

(87) WIPO No: **WO10/088523**

(30) Priority Data

(31) Number **61/148,888** (32) Date **2009.01.30** (33) Country **US**

(43) Publication Date: **2010.08.05**

(44) Accepted Journal Date: **2015.05.28**

(71) Applicant(s)  
**Ab Initio Technology LLC**

(72) Inventor(s)  
**Gould, Joel;Studer, Scott**

(74) Agent / Attorney  
**Pizzeys, PO Box 291, WODEN, ACT, 2606**

(56) Related Art  
**US 2008/0256014 A1**  
**US 8069129 B2**  
**US 2007/0050340 A1**  
**US 2006/0294150 A1**

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization  
International Bureau

(43) International Publication Date  
5 August 2010 (05.08.2010)

(10) International Publication Number  
WO 2010/088523 A1



(51) International Patent Classification:  
*G06F 7/00* (2006.01)

(21) International Application Number:  
PCT/US2010/022593

(22) International Filing Date:  
29 January 2010 (29.01.2010)

(25) Filing Language:  
English

(26) Publication Language:  
English

(30) Priority Data:  
61/148,888 30 January 2009 (30.01.2009) US

(71) Applicant (for all designated States except US): **AB INITIO TECHNOLOGY LLC** [US/US]; 201 Spring Street, Lexington, Massachusetts 02421 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **GOULD, Joel** [US/US]; 50 Thomberry Road, Winchester, Massachusetts 01890 (US). **STUDER, Scott** [US/US]; 27 Pillsbury Lanc, Georgetown, Massachusetts 01833 (US).

(74) Agents: **HENNESSEY, Gilbert, H.** et al.; Fish & Richardson P.C., P.O. Box 1022, Minneapolis, MN 55440-1022 (US).

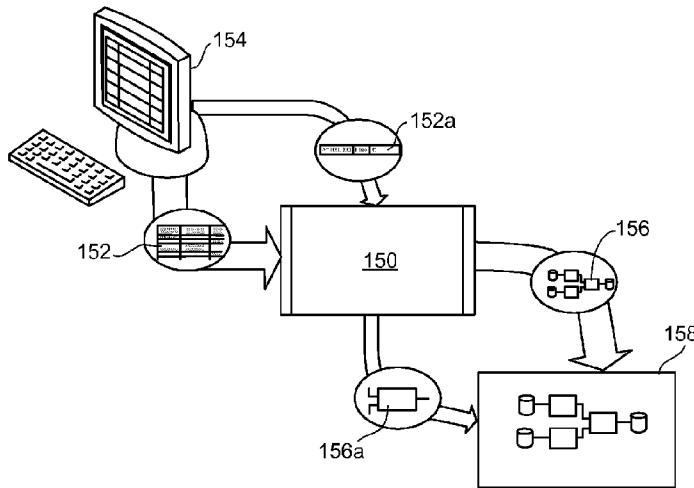
(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

- with international search report (Art. 21(3))
- with amended claims (Art. 19(1))

(54) Title: PROCESSING DATA USING VECTOR FIELDS



(57) Abstract: Disclosed is a method including receiving a rule (152) having at least one rule case (210a-h) for producing an output value based on one or more input values, generating a transform (112, 156) for receiving data from an input dataset (100) and transforming the data based on the rule (152) including producing a first series of values (813) for at least one output variable (508, 801) in an output dataset (120), at least one value in the first series of values (813) including a second series of values (816), and providing an output field corresponding to the at least one output variable (508, 801) in the output dataset (120) for storing the second series of values (816).

FIG. 2

## **PROCESSING DATA USING VECTOR FIELDS**

### **CROSS-REFERENCE TO RELATED APPLICATIONS**

This application claims priority to U.S. Application Serial No. 61/148,888, filed on January 30, 2009, incorporated herein by reference.

### **BACKGROUND**

This description relates to processing data using vector fields.

Some computing systems provide an interface for specifying rules that are used for automated decision making in various data processing applications. Decisions associated with processing data representing credit card transactions or airline frequent flyer programs, for example, may be governed by a given set of rules. In some cases, these rules are described in human-readable form. The computing system may provide an interface for a user to define or edit these rules, and then incorporate the rules into a data processing system.

### **SUMMARY**

In one aspect, in general, a method includes receiving a rule having at least one rule case for producing an output value based on one or more input values, generating a transform for receiving data from an input dataset and transforming the data based on the rule including producing a first series of values for at least one output variable in an output dataset, at least one value in the first series of values including a second series of values, and providing an output field corresponding to the at least one output variable in the output dataset for storing the second series of values.

Aspects can include one or more of the following features.

The transform can be included in a component of a graph-based application represented by a graph, with vertices in the graph representing components, and directed links between vertices in the graph represent flows of data between components.

A first graph component including the transform can provide a flow of data to the transform from the input dataset.

The first graph component can be an executable computation component, and the graph can include a second graph component that is a data storage component representing the input dataset.

Producing a first series of values for at least one variable in an output dataset can include producing rows for an output table, each row defining a record having values for a set of variables including the output variable.

Providing an output field for storing the second series of values can include providing an array for storing a predetermined number of the second series of values, the predetermined number being a default number that is modifiable to a user-specified number. The output field can include a cell in a table.

Receiving the rule can include receiving at least a row of a rule table, the row corresponding to a rule case, and having an output including one or more or a combination of the input values, a predetermined value, or a value computed from one or more of the input values.

The rule case can include one or more of: having an input value equal to a threshold, having an input value above a threshold, having an input value below a threshold, having an input value belonging to a set of values, having an input value matching a pattern of values, having a relationship to another input value, having a relationship to an output value of another set of rules, or having a relationship to a value in a memory.

The input dataset can include records having values for scalar variables and vector variables. At least one of the records can include an array for storing a predetermined number of records, the predetermined number being a default number that is modifiable to a user-specified number. At least one of the records includes an internal reference table to define key relationships to sub-records in the at least one of the records.

The method can also include, in response to a rule, producing the second series of values for the output variable in the output dataset based on the key relationships in the internal reference table.

The method can also include, in response to a rule case in a rule, triggering the rule case to produce a value for the output variable in the output dataset. Triggering a

rule case can include triggering the rule based on a scalar value in the input dataset satisfying the at least one rule case in the rule.

Triggering a rule case can include triggering the rule based on each value in a vector in the input dataset satisfying the at least one rule case in the rule.

Triggering a rule case can include triggering the rule case based on an output of an aggregate function applied to a vector in the input dataset satisfying the at least one rule case in the rule.

Generating the transform can include converting each of a plurality of rule cases in the rule to a logical expression to form a plurality of logical expressions, and compiling the plurality of logical expressions into computer-executable code.

Compiling the plurality of logical expressions can include one or more of combining expressions, optimizing individual expressions, and optimizing groups of expressions.

In another aspect, in general, a computer-readable medium storing a computer program for updating a component in a graph-based computation having data processing components connected by linking elements representing data flows includes instructions for causing a computer to receive a rule having at least one rule case for producing an output value based on one or more input values, generate a transform for receiving data from an input dataset and transforming the data based on the rule including producing a first series of values for at least one output variable in an output dataset, at least one value in the first series of values including a second series of values, and provide an output field corresponding to the at least one output variable in the output dataset for storing the second series of values.

In another aspect, a system includes a means for receiving a rule having at least one rule case for producing an output value based on one or more input values, a processor configured to generate a transform for receiving data from an input dataset and transforming the data based on the rule including producing a first series of values for at least one output variable in an output dataset, at least one value in the first series of values including a second series of values, and a means for providing an output field corresponding to the at least one output variable in the output dataset for storing the second series of values.

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

#### **DESCRIPTION OF DRAWINGS**

Fig. 1 is a schematic depicting an example transform.

Fig. 2 is an example transform generator.

Figs. 3 and 4 are example rule sets.

Fig 5 is an example Fire-Many rule set.

Fig. 6, 7, and 8 are example output, rule, and result tabs.

Fig. 9 is a schematic depicting computation of scalars and vectors.

Fig. 10A and 10B show an example input record having record vectors.

#### **DESCRIPTION**

A business rule can be expressed as a set of criteria that can be used to, for example, convert data from one format to another, make determinations about data, or generate new data based on a set of input data. For example, in FIG. 1, a record 102 in a flight reservation system indicates a passenger's name 104, miles 106 the passenger has flown in the current year, class 108 of the passenger's ticket, and the passenger's current row 110 in an airline. A business rule may indicate that such the passenger should be classified within boarding group "1," e.g., group 118. A business rule is generally easy for a human to understand, e.g., "first class passengers are in group 1," but may need to be translated into language that a computer can understand before it can be used to manipulate data. Accordingly, to implement the business rule, a transform 112 is generated to receive an input record, e.g., record 102, from one or more data sources, e.g., input dataset 100, and produce an output record, e.g., record 114, indicating the passenger's name 104 and group 118, into an output dataset 120. Input and output datasets are also referred to as data streams.

To simplify creation of a transform 112 for non-technical users, typically an editor tool (not shown) is provided to input a set of business rules, referred to as a rule set, or a

set of rules, in a format familiar to the users. The set of rules, in turn, instructs a computer system to generate the transform 112 which further instructs the computer system what to do with input dataset 100, and what to produce into output dataset 120. A rule or rule set that corresponds to a single transform can include one or more rule cases that compute different values for a rule set's output variables depending on an input record. When a rule case in a rule is triggered, the rule, and more particularly, the rule case, is regarded to be fired. For example, only one rule case in a rule can be fired. In some examples, more than one rule case in a rule can be fired. In some examples, when a rule case is fired, the entire rule can be regarded as being fired. In some implementations, a rule case or rule is triggered or fired if, for example, an input scalar or vector value in an input dataset satisfies one or more conditions in the rule case or rule. A rule set can also include other rules sets. The other rule sets can produce values for additional or alternative output variables. For example, a rule set can directly contain or indirectly refer to other rule sets, referred to as "included" rule sets.

An example transform generation system is shown in FIG. 2. A generator 150 receives as input a rule set 152 from an editor 154 and generates a transform 156. The generated transform 156 may be provided to a graph-based computation system 158 as a component to be used in a graph or as an entire graph itself, depending on the system's architecture and the purpose of the transform and the business rules. The graph-based computation system 158 can provide a computation environment that allows a programmer to build a graph-based application by using components as building blocks. A graph-based application is often represented by a directed graph, with vertices in the graph representing components (either data storage components or executable computation components), and the directed links or "edges" in the graph representing flows of data between components. A dataflow graph (also called simply a "graph") is a modular entity. Each graph can be made up of one or more other graphs, and a particular graph can be a component in a larger graph.

The generator 150 can be, for example, a compiler, a custom-built program, or a graph-based computation configured using standard tools to receive the rule set 152 and output the transform 156. Any technique for producing, and subsequently updating the transform 156 known to those skilled in the art can be used to generate transform 156.

For example, a technique for producing transforms is described in U.S. Patent Application No. 11/733,434, entitled “Editing and Compiling Business Rules,” filed April, 10, 2007, and incorporated herein by reference in its entirety.

In some examples, the transform 156 generates only one value for an output variable corresponding to an input record 102. In such a scheme, a rule set can fire at most only once. Accordingly, some problems, e.g., data quality problems, may not be easily implemented using the transform 156. In some examples, output variables in an output dataset 120 can include “Write-Once Outputs.” In general, “Write-Once Outputs” are output variables that are typically written to once for a given input record, and store only one value for the given input record. Rule sets that produce such variables are called “Fire-Once” rules.

In some examples, a “Fire-Many” rule can produce “accumulator” output variables, e.g., variables that are capable of receiving a series of values for a given input record, instead of only one value. A “Fire-Many” rule would fire for every rule case within a rule set that is triggered for that input record, and not just, for example, the first rule case that is triggered.

In some examples, a rule set can be entered in a tabular (or “spreadsheet”) format, as shown in FIG. 3, with rows and columns that intersect in cells. Trigger columns 202, 204, 206, 208 in table 200 correspond to criteria for available input data values, and rows 210a-h correspond to rule cases, i.e., sets of criteria that relate to the available input data values. A cell at the intersection of a trigger column and the applicable rule case row 210n contains a criterion for that trigger column and rule case. A rule case 210n applies to a given record, e.g., 102 in FIG. 1, if data values of the record 102, for each trigger column in which the rule case has criteria, meets the triggering criteria. If a rule case 210n applies, output is generated based on one or more output columns 212. As described above, in general, a rule case that has all of its input relationships satisfied may be referred to as “triggered,” and the rule set is referred to as “fired.” Each output column 212 corresponds to a potential output variable, and the value in the corresponding cell at the intersection of the column 212 and the applicable rule case row 210n determines the output, if any, for that variable. In some examples, the cell can contain a

value that is assigned to the variable or it can contain an expression that is evaluated to generate the output value, as discussed below. In some examples, there may be more than one output column, though only one is shown in FIG. 3.

There may be several different types of trigger columns, including columns that correspond to a variable, columns that contain expressions but are calculated once and then treated like variables, and columns that only contain expressions. Columns that only contain expressions are in some respects simpler than those corresponding to or treated as variables. Such trigger columns can contain, for example, one of the following types of cell values for defining trigger column criteria:

- An expression. The condition will be considered to be true if the evaluation of the expression evaluates to a non-zero, or non-NULL value.
- The keyword “any,” or an empty string. The condition is always true. Each empty cell in a trigger column is equivalent to one explicitly containing the keyword “any.”
- The keyword “else.” The condition is true if none of the cells above the cell containing “else” is true, in rows where all cells to the left are identical.
- The keyword “same”. The condition is true if the cell above is true.

Columns that correspond to a variable (column variables) can have two types of cells. One type of cell is an expression cell. Those cells behave exactly like cells in a column that contains only expressions, described above. However, the keyword “this” can be used in the expression to refer to the column variable. The other type of cell is a comparison value. An example grammar for comparison values is as follows:

```

comparison_value ::= compound_value ( “or” compound_value )*
compound_value ::= simple_value ( “and” simple_value )*
simple_value ::= [ “not” ] ( value_expression | simple_function |
membership_expr )
value_expression ::= [ operator ] value_element
operator ::= “>” | “<” | “>=” | “<=” | “!=” | “=” | “equals”
value_element ::= constant | constant | variable | “(“expression “)”
simple_function ::= “is _null” | “is _blank” | “is _valid” | “is _defined” | “is _bzero”

```

```
membership_expr ::= "in" "[" value_element ( ( "," | "to" | "or" ) value_element )* "]"
```

where a "\*" means a term is repeated zero or more times.

Any suitable programming language or syntax may be used. Examples include C, Java, DML, or Prolog. The column variable is compared against the comparison value according to the operator, function, or membership expression. In the example of FIG. 3, the first two columns 202 and 204 contain comparison values with the ">=" operator. Accordingly, the criteria is met if the value for that column is greater than or equal to the corresponding number. If there is no operator, as in the "Class of Seat" column, then "equals" is assumed. A constant can be any legal constant in whatever programming language or syntax is used in the underlying system. An expression is any legal expression in the language being used that returns a compatible datatype that will be compared against the column variable. In some examples, expressions inside comparison values are enclosed in parenthesis to avoid ambiguity.

In the example of FIG. 3, the first row 210a has criteria in only one column, 202, which indicates that if the total number of frequent flier miles for a traveler is greater than 1,000,000, then that rule case applies regardless of what value any other columns may have. In that case, the "Boarding Group" output variable for that user is set to group 1. Likewise, the second rule case 210b indicates that any flier in first class is in group 1. In some examples, the rules are evaluated in order, so a traveler having over 1,000,000 miles and a first class ticket will be in group 1, but only the first rule case 210a will be triggered.

The rule cases 210a-h (FIG. 3) can also be represented as individual simple rules, each in their own table, as shown in FIG. 4. Rules 220a-d corresponds to rows 210a-d of FIG. 3, respectively, while rule 220e has four rule cases corresponding to rows 210e-h together. A user could create these individual rules separately, rather than generating the entire table shown in FIG. 3. Each rule case contains a value (at least implicitly) for every trigger column and a value for every output column (the value can be blank, i.e., effectively set to "any"). When multiple rules generate the same output, the rules are ordered and they are considered in order until a rule case in one rule triggers on the inputs and generates an output. If no rule case in a rule triggers, the next rule that produces the

same output is processed. If no cases in any rule trigger for an output, a default value is used.

In some examples, a user interface of the editor tool can be used to graphically identify cells that contain expressions. Accordingly, a user can understand the difference between an expression that will be evaluated to true or false on its own and an expression that returns a value that is compared against the column variable. When the user is typing, he can indicate that a particular cell is to be an expression cell by, for example, typing an asterisk at the beginning.

For columns that correspond to output variables, the cells can contain one of the following:

- A value. The value that will be assigned to the output variable
- An expression. The value of the expression is assigned to the output variable. If the expression evaluates to NULL then the field gets the NULL value, unless the output field is not-nullable. In which case, an error is generated.
- The keyword “null”. If the output field is nullable, then the field will be assigned NULL. Otherwise, an error is generated.
- An empty string. If the output field has a default value, then the default value is assigned. Otherwise, the cell is treated as if it contains the keyword “null”.
- The keyword “same”. The output field is assigned the same value computed in the cell above.

In addition to expressions, users can be allowed to attach comments to any cell in the rule, which can be displayed in response to user interaction (e.g., clicking or “hovering” a pointer).

In some implementations, a rule set, e.g., the rule set shown below in Table 1, can include multiple rule cases that generate multiple output records for a single input record.

Trigger: Automobile Option	Trigger: Budget	Output: Trim Level
Honda S2000	$\geq 37000$	S2000 CR
Honda S2000	clsc	S2000

Honda Accord Coupe	$\geq 29000$	Accord Coupe EX-L V-6
Honda Accord Coupe	$\geq 26000$	Accord Coupe EX-L
Honda Civic Sedan	$\geq 24000$	Accord Coupe EX
Honda Element	any	Accord Coupe
...	...	...

Table 1

The rule set above considers a family's automobile options in view of the family's budget, and outputs a trim level for the automobile. In some examples of such a rule set (referred to as a "normalize rule set"), at least one of the output values is identified as a key output value, e.g., "S2000 CR." When the rules that compute the key output value "S2000 CR" are evaluated, the rule case (Automobile Option: Honda S2000 and Budget:  $\geq 37000$ ) that triggered on the input data record to generate the output value "S2000 CR" is noted. The rule set is then evaluated again with the previously-triggered rule case (Automobile Option: Honda S2000 and Budget:  $\geq 37000$ ) disabled to see if any other rule cases trigger and produce an output value. The process described above is repeated until no additional rule cases are triggered. Each output value is stored as a separate output record. In some examples, rule cases are grouped, such that if one triggers, others in its group are also disabled on the next iteration for the same input record.

In some examples, the transform corresponding to the normalize rule set can use two stages of processing. First, an input record is read and a count is computed, e.g., by calling a "length" function. The count corresponds to a number of output records that will be generated. Then, another function, i.e., "normalize" function, is called for each output record. The normalize function receives a copy of the input record and a current index from the count produced by the length function and produces output values into different output records. For example, if the input record had a family size of four (4) and a budget of \$20,000, the transform generates three output records, one for each of the three suggested cars (Accord Sedan, Civic, and Element).

In some implementations, the transform calculates all possible values for the Automobile Option, using the "length" function so that the number of output records is

known. Once the transform has calculated all possible Automobile Output values, the transform can then call the “normalize” function as many times as there are output records, to assign values to each of the output records.

In some implementations, instead of the two stage processing described above, the transform can calculate all possible values for the Automobile Option by calling the “normalize” function directly several times until there are no more values to compute.

FIG. 5 is an example rule set 500 for generating multiple values 504 for an output variable 508. A user may be interested in knowing all of the reasons why a specific vehicle is considered invalid, not just the first reason. In some examples, as shown in FIG. 6, a first step is for the user to specify, using an output tab 600 in a user interface of the editor, that the rule set 500 produces multiple output values 504.

As such, the user indicates that the output variable 508 “Name Validation Message” is an accumulator variable for receiving a series of values 504. The Output Type 604 corresponding to the output variable 508 changes to indicate “accumulator” 608.

In some examples, scalar values corresponding to the output variable 508 can be “accumulated” for use with “score-card” style rule sets. A score card style rule set refers to a type of business rule where a user indicates a positive or negative score to be included into a rules value. Accordingly, rather than storing values corresponding to the output variable 508 as an output vector, a sum of the values that are accumulated corresponding to the output variable 508 is stored as a scalar value.

In some examples, the accumulator output variable 508 maps to a variable length vector or array for each record in the output dataset. As such, if the output variable 508 is treated as an array, the user can specify a size for the output variable 508. The user can specify a length of the output variable 508 by changing the Max Count 612 parameter. Accordingly, field 614 indicates that the output variable 508 is treated as an array for receiving a certain number (e.g., 20) of values. In some examples, in the absence of a user-specified size, by default, the output variable 508 can receive unlimited number of values. As such, the Max Count 612 parameter indicates, for example, “unlimited.” In some examples, to help distinguish accumulator type output variables from write-once

type output variables, the editor can prohibit users from editing the Max Count 612 parameter for a write-once variable. In some examples, if the user switches from an accumulator output variable to a write-once output variable, the editor can clear the Max Count 612 parameter.

FIG. 7 is an example rule tab 700 showing a fire many rule set, e.g., “Validate Person.” Accumulator output variables 708 are visually distinguished from write-once outputs 712 by the annotation “Write-Once Outputs” or “Accumulator Outputs.” In addition, various other annotations are possible. For example, a type of rule set, i.e., a “Fire-Many Rule” (rule which produce accumulator outputs), or a “Fire-Once Rule” (rule which produces a scalar output) may be indicated at the top 704 of the rule tab 700, or a vertical annotation 712 on one side indicates “Fire Once” or “Fire Many.” In some examples, different icons may be used for fire once and fire many rules. In some examples, all of the rule cases that fired may be highlighted for inspection by the user.

FIG. 8 is an example results tab 800 showing contents of the accumulator output variable 801, “Validation Message.” As shown, the output variable 801 can assume a first series of values 813 for each record, and at least one of the values of the first series of values 813 (e.g., the value corresponding to “TANGELA SCHEPP”) can assume a second series of values 816 that are displayed as a collection of comma separated values. In some examples, a user can “hover” a mouse pointer over an accumulator output value to uncover a tool tip showing a list of accumulated values. In some examples, when performing a test including, for example, benchmark data, an output can be marked as being different if a vector in the benchmark data differs from the vector in the output in any way. For example, differences can include, the benchmark vector having a different number of items than the output vector, the benchmark vector having items in a different order than the output vector, and individual items within each of the vectors being different.

In operation, an accumulator output variable is used for receiving multiple output values produced by a Fire-Many rule set as described below. For example, consider the following rule set shown in Table 2:

Trigger: Budget	Trigger: Family Size	Output: Automobile Option
$\geq 35000$	$\leq 2$	Honda S2000
$\geq 22000$	$\leq 2$	Honda Accord Coupe
$\geq 20000$	$\leq 4$	Honda Accord Sedan
$\geq 15000$	$\leq 4$	Honda Civic Sedan
$\geq 20000$	$\leq 6$	Honda Element
$\geq 28000$	$\leq 7$	Honda Odyssey
$\geq 50000$	$\leq 4$	Acura RL

Table 2

The rule set above considers, for example, a family size of 4 and a budget of \$20,000, to suggest three cars (Accord Sedan, Civic and Element). Accordingly, in this case, an output variable “Automobile Option” in an output dataset is deemed to be able to receive multiple values. Each rule case in the rule set is evaluated and any time a rule case triggers, a value from the rule set above is added to the accumulator output variable.

The triggers in the rule set above can be any scalar variables (non-vectors) including input values, lookups and other output values. In some examples, an output variable can compute another output variable. In some examples, only a non-vector output can be used as a trigger. In some examples, it is possible to indirectly use one accumulator output variable to compute another accumulator output variable by using the aggregation functions. For example, consider the following rule set shown in Table 3:

Trigger	Output: Family Members
is_alive	Self
is_married and not is_separated	Spouse
has_baby	Baby
has_teenage_girl	Daughter
has_teenage_boy	Son

Table 3

The rule set above computes an accumulator output variable called “Family Members.” Now, consider the following rule set shown in Table 4:

Output: Family Size
count_of( Family Members )

Table 4

The rule set in Table 4 computes a scalar (non-vector) called “Family Size,” using an aggregation function. Accordingly, first, an output vector is computed that includes a list of all our family members. Then, a count function counts the number of people in the list. The count is then used as input to compute a list of automobiles.

FIG. 9 illustrates an example implementation using scalars and vectors to compute values for other scalar and vectors using an accumulator output variable. As shown, S1, S2 and S3 represent scalar variables. V1 and V2 represent vector variables. S1 is used to compute S2; then S2 is used to compute four different values of V1. Then all four values of V1 are used to compute S3 (e.g., through the use of an aggregation function). Finally, S3 is used to compute three values of V2.

In some implementations, the editor can produce validation errors when the user attempts to carry out any of the following example actions: Marking an output as an accumulator when the type of the field in any of the datasets is anything other than a variable length vector; mark an output as “write-once” when the type of the field in any of the datasets is a vector; provide a default value for an accumulator (in an implementation in which only write-once outputs can have default-values), use an accumulator output as a comparison trigger column; mix accumulator and write-once outputs within a single rule; and input a value other than unlimited or a positive number in the Max Count parameter of an accumulator output variable.

In some examples, input records can include vectors. FIG. 10A is an example format of an input record 950 that includes at least two vectors records, i.e., driver record vector 952, and vehicle record vector 954. FIG. 10B shows example data 956 for the input record 950.

An aggregation function can be included in a rule set to convert the record vectors 952, 954 into scalars. For example, the rule set can include a specification “Age of youngest driver.” In some implementations, the specification is expressed as “minimum (Driver Age),” or a data manipulation language (DML) function such as “do\_minimum (in0.drivers, ‘age’)” can be used. In response to the rule set, a scalar value is produced, e.g., 21 (from Pebbles’ record in FIG. 10B) In some examples, in operation, a function can loop through all the records in the driver record vector 952 to find the minimum value for driver age.

Considering another example, the specification in a rule set can be “Number of points plus one for the youngest male driver.” The specification can be expressed as “minimum (Driver Age, Driver Sex = Male, Driver Points + 1).” In response to this rule set, a scalar value is produced, e.g., 14 (from BamBam’s record). In some implementations, the scalar values can be assigned to intermediate or output variables, which are scalars.

In some examples, a rule can be written for each element in a record vector. For example, consider the following rule set shown in Table 5:

Vehicle has Air Bag (trigger)	Vehicle has Seat Belts (trigger)	Value Adjustment (output)
no	No	0
no	Yes	100
yes	No	150
yes	Yes	300

Table 5

The specification in the rule set of Table 5 is “For each car, compute the adjustment to the car’s value, which is 100 if the car has seat belts, 150 if the car has air bags, and 300 if the car has both.” As shown, the output variable “Value Adjustment” is a vector variable. In response to the above rule, a vector, e.g., [0, 300, 100] is produced. In some examples, in operation, the rule set is executed multiple times, once for every record in the vehicle record vector 954.

In some examples, the rule set can also reference scalar values, or other vectors as long as the vectors are of the same length. For example, consider the following rule set shown in Table 6:

Vehicle Age (Input)	Adjusted Value (Output)
> 2	Vehicle Value + Value Adjustment + Geographic Adjustment - 50
else	Vehicle Value + Value Adjustment + Geographic Adjustment

Table 6

The specification in the rule set of Table 6 is “For each car, compute the adjusted value, which is the sum of the car’s value, its value adjustment and the geographic risk.

Subtract 50 if the car is older than 2 years.” In this rule, “Adjusted Value” is a vector variable. Accordingly, to avoid a runtime error due to unequal vector lengths, the vector variable “Value Adjustment” is of same length as the vehicle record vector 954. In response to this rule set, a vector, e.g., [1030, 1880, 1330] is produced.

In some examples, when XML records are complex, a single input record can be used to represent many logical records by relating them with key relationships. For example, each vehicle sub-record in the vehicle record vector 954 can include a foreign key, e.g., “driver,” to relate to a matching key in the driver record vector 952, e.g., “name.” In this manner, the record vectors 952, 954 can be implemented as look-up file, or an internal reference table. For example, an internal reference table associated with the vehicle record vector 954 can be as follows:

Primary Driver Name (primary key)

Primary Driver Age

Primary Driver Sex

Primary Driver Points

Accordingly, internal reference tables, can be created for each input record by treating the sub-records in the record vectors as records in the internal reference tables.

In operation, consider for example, a rule set shown in Table 7:

output: Age of Policy Driver
Primary Driver Age ( Policy Driver )

Table 7

The specification in the rule set of Table 7 is “Compute the Age of the Policy Driver, which is the Primary Driver Age found by using the value of Policy Driver as the key for the associated internal reference table.” The specification returns the value in the

Primary Driver Age column, which is then assigned to the output variable, “Age of Policy Driver.” “Age of Policy Driver” is a scalar value. In another example, consider the rule set shown in Table 8 below:

Age at Purchase (output)
Primary Driver Age – Vehicle Age

Table 8

The specification in the rule set of Table 8 is “Compute the Age at Purchase, which is the difference between the vehicle’s age and the age of the vehicle’s primary driver.” For purposes of illustration, assume that the look-up key is assigned “Vehicle Primary Driver” by default. The output variable “Age at Purchase” is a vector variable. Accordingly, in response to the above rule, [31, 19, 27] is produced.

In some examples, the look-up key “Vehicle Primary Driver” can be specified explicitly in parentheses as follows “Primary Driver Age (Vehicle Primary Driver) – Vehicle Age.

In some examples, the internal reference tables can be used in aggregation functions. For example, a specification can be “Compute the average over all the vehicles of the age of their primary drivers.” This specification can be implemented by the function, for example, “average (Primary Driver Age (Vehicle Primary Driver)).” In response to this function, a scalar value is produced, e.g., 29.67.

In some implementations, a user can visualize the computations steps in the above rule sets. For example, in testing mode, it may be useful for a user to be able to examine values of interest, e.g., intermediate values of input and output variables (both scalar and vector variables). Various techniques for visualizing the steps known in the art can be used. For example, a pop-up table having a row for each element in the input record vector 952, 954 can be implemented to summarize the intermediate values indicating what items have been filtered out, or computed.

The techniques described above can be implemented using software for execution on a computer. For instance, the software forms procedures in one or more computer

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

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

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

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

claims. For example, a number of the function steps described above may be performed in a different order without substantially affecting overall processing.

Throughout this specification and the claims which follow, unless the context requires otherwise, the word "comprise", and variations such as "comprises" and "comprising", will be understood to imply the inclusion of a stated integer or step or group of integers or steps but not the exclusion of any other integer or step or group of integers or steps.

The reference to any prior art in this specification is not, and should not be taken as, an acknowledgement or any form of suggestion that the prior art forms part of the common general knowledge in Australia.

Other embodiments are within the scope of the following claims.

What is claimed is:

1. A computer-implemented method, including:
  - via a processor, receiving a rule having at least one rule case for producing an output value based on one or more input values,
  - generating a transform for receiving data from an input dataset and transforming the data based on the rule, where transforming the data based on the rule includes producing a first series of values for at least one output variable in an output dataset, and producing at least one value in the first series of values including a second series of values, and
  - providing in the output dataset an output field corresponding to the at least one output variable for storing the second series of values.
2. The method of claim 1 in which the transform is included in a data structure representing a component of an application represented by a graph, with vertices in the graph representing components, and directed links between vertices in the graph representing flows of data between components.
3. The method of claim 2 in which a first graph component including the transform provides a flow of data to the transform from the input dataset.
4. The method of claim 3 in which the first graph component is an executable computation component, and the graph includes a second graph component that is a data storage component representing the input dataset.
5. The method of claim 1 in which producing a first series of values for at least one variable in an output dataset includes producing rows for an output table, each row defining a record having values for a set of variables including the output variable.
6. The method of claim 1 in which providing an output field for storing the second series of values includes providing an array for storing a predetermined number of the

second series of values, the predetermined number being a default number that is modifiable to a user-specified number.

7. The method of claim 1 in which the output field includes a cell in a table.

8. The method of claim 1 in which receiving the rule includes receiving at least a row of a rule table, the row corresponding to a rule case, and having an output including one or more or a combination of the input values, a predetermined value, or a value computed from one or more of the input values.

9. The method of claim 8 in which the rule case includes one or more of: having an input value equal to a threshold, having an input value above a threshold, having an input value below a threshold, having an input value belonging to a set of values, having an input value matching a pattern of values, having a relationship to another input value, having a relationship to an output value of another set of rules, or having a relationship to a value in a memory.

10. The method of claim 1 in which the input dataset includes records having values for scalar variables and vector variables.

11. The method of claim 10 in which at least one of the records includes an array for storing a predetermined number of records, the predetermined number being a default number that is modifiable to a user-specified number.

12. The method of claim 10 in which at least one of the records includes an internal reference table to define key relationships to sub-records in the at least one of the records.

13. The method of claim 12 also including, in response to a rule, producing the second series of values for the output variable in the output dataset based on the key relationships in the internal reference table.

14. The method of claim 1 also including, in response to a rule case in a rule, triggering the rule case to produce a value for the output variable in the output dataset.

15. The method of claim 14 in which triggering the rule case includes triggering the rule case based on a scalar value in the input dataset satisfying the at least one rule case in the rule.
16. The method of claim 14 in which triggering the rule case includes triggering the rule case based on each value in a vector in the input dataset satisfying the at least one rule case in the rule.
17. The method of claim 14 in which triggering the rule case includes triggering the rule case based on an output of an aggregate function applied to a vector in the input dataset satisfying the at least one rule case in the rule.
18. The method of claim 1 in which generating the transform includes converting each of a plurality of rule cases in the rule to a logical expression to form a plurality of logical expressions, and compiling the plurality of logical expressions into computer-executable code.
19. The method of claim 18 in which compiling the plurality of logical expressions includes one or more of combining expressions, optimizing individual expressions, and optimizing groups of expressions.
20. A computer-readable medium, storing a computer program for updating a component in a graph-based computation having data processing components connected by linking elements representing data flows, the computer program including instructions for causing a computer to:
  - receive a rule having at least one rule case for causing a computation system to produce an output value based on one or more input values,
  - generate a data structure including a transform for receiving data into the computation system from an input dataset and transforming the data based on the rule, where transforming the data based on the rule includes producing a first series of values for at least one output variable in an output dataset, and producing at least one value in the first series of values including a second series of values, and

provide, in the output dataset, an output field corresponding to the at least one output variable for storing the second series of values.

21. A system, including:

means for receiving a rule having at least one rule case for producing an output value based on one or more input values,

a processor configured to generate a transform for receiving data from an input dataset and transforming the data based on the rule, where transforming the data based on the rule includes producing a first series of values for at least one output variable in an output dataset, and producing at least one value in the first series of values including a second series of values, and

means for providing, in the output dataset, an output field corresponding to the at least one output variable for storing the second series of values.

22. The method of claim 14 in which triggering the rule case includes satisfying predefined one or more criteria by the input values.

23. The method of claim 1 also including:

triggering two or more rule cases belonging to the rule in response to the input values; and

producing, by the transform, two or more output values each corresponding to the triggered two or more rule cases as the second series of values.

24. The method of claim 23 also including providing to a user an ability to inspect the two or more rule cases that triggered in response to the input values.

25. The method of claim 1 in which the at least one rule case includes one or more criteria which when satisfied by the input values causes production by the transform of one or more output values.

26. The method of claim 1 also including storing a scalar value in the output dataset corresponding to an operation performed on the second series of values.

27. The method of claim 1 also including displaying the second series of values in a tool tip upon a cursor being positioned over the output field.
28. The computer-readable medium of claim 20 in which the transform is included in a component of a graph-based application represented by a graph, with vertices in the graph representing components, and directed links between vertices in the graph representing flows of data between components.
29. The computer-readable medium of claim 28 in which a first graph component including the transform provides a flow of data to the transform from the input dataset.
30. The computer-readable medium of claim 29 in which the first graph component is an executable computation component, and the graph includes a second graph component that is a data storage component representing the input dataset.
31. The computer-readable medium of claim 20 in which producing a first series of values for at least one variable in an output dataset includes producing rows for an output table, each row defining a record having values for a set of variables including the output variable.
32. The computer-readable medium of claim 20 in which the computer program includes instructions for causing the computer to provide an array for storing a predetermined number of the second series of values, the predetermined number being a default number that is modifiable to a user-specified number.
33. The computer-readable medium of claim 20 in which the output field includes a cell in a table.
34. The computer-readable medium of claim 20 in which the computer program includes instructions for causing the computer to receive at least a row of a rule table, the row corresponding to a rule case, and having an output including one or more or a combination of the input values, a predetermined value, or a value computed from one or more of the input values.

35. The computer-readable medium of claim 34 in which the rule case includes one or more of: having an input value equal to a threshold, having an input value above a threshold, having an input value below a threshold, having an input value belonging to a set of values, having an input value matching a pattern of values, having a relationship to another input value, having a relationship to an output value of another set of rules, or having a relationship to a value in a memory.
36. The computer-readable medium of claim 20 in which the input dataset includes records having values for scalar variables and vector variables.
37. The computer-readable medium of claim 36 in which at least one of the records includes an array for storing a predetermined number of records, the predetermined number being a default number that is modifiable to a user-specified number.
38. The computer-readable medium of claim 36 in which at least one of the records includes an internal reference table to define key relationships to sub-records in the at least one of the records.
39. The computer-readable medium of claim 38 in which the computer program also includes instructions for causing the computer to, in response to a rule, produce the second series of values for the output variable in the output dataset based on the key relationships in the internal reference table.
40. The computer-readable medium of claim 20 in which the computer program also includes instructions for causing a computer to, in response to a rule case in a rule, trigger the rule case to produce a value for the output variable in the output dataset.
41. The computer-readable medium of claim 40 in which triggering the rule case includes triggering the rule case based on a scalar value in the input dataset satisfying the at least one rule case in the rule.

42. The computer-readable medium of claim 40 in which triggering the rule case includes triggering the rule case based on each value in a vector in the input dataset satisfying the at least one rule case in the rule.
43. The computer-readable medium of claim 40 in which triggering the rule case includes triggering the rule case based on an output of an aggregate function applied to a vector in the input dataset satisfying the at least one rule case in the rule.
44. The computer-readable medium of claim 20 in which the computer program includes instructions for causing the computer to convert each of a plurality of rule cases in the rule to a logical expression to form a plurality of logical expressions, and compile the plurality of logical expressions into computer-executable code.
45. The computer-readable medium of claim 44 in which compiling the plurality of logical expressions includes one or more of combining expressions, optimizing individual expressions, and optimizing groups of expressions.
46. The computer-readable medium of claim 40 in which triggering the rule case includes satisfying predefined one or more criteria by the input values.
47. The computer-readable medium of claim 20, the computer program further including instructions for causing a computer to:
  - trigger two or more rule cases belonging to the rule in response to the input values;
  - and
  - produce, by the transform, two or more output values each corresponding to the triggered two or more rule cases as the second series of values.
48. The computer-readable medium of claim 47 also including instructions for causing a computer to provide to a user an ability to inspect the two or more rule cases that triggered in response to the input values.

49. The computer-readable medium of claim 20 in which the at least one rule case includes one or more criteria which when satisfied by the input values causes production by the transform of one or more output values.
50. The computer-readable medium of claim 20 also including instructions for causing a computer to store a scalar value in the output dataset corresponding to an operation performed on the second series of values.
51. The computer-readable medium of claim 20 also including instructions for causing a computer to display the second series of values in a tool tip upon a cursor being positioned over the output field.
52. The system of claim 21 in which the transform is included in a component of a graph-based application represented by a graph, with vertices in the graph representing components, and directed links between vertices in the graph represent flows of data between components.
53. The system of claim 52 in which a first graph component including the transform provides a flow of data to the transform from the input dataset.
54. The system of claim 53 in which the first graph component is an executable computation component, and the graph includes a second graph component that is a data storage component representing the input dataset.
55. The system of claim 21 in which producing a first series of values for at least one variable in an output dataset includes producing rows for an output table, each row defining a record having values for a set of variables including the output variable.
56. The system of claim 21 in which providing an output field for storing the second series of values includes providing an array for storing a predetermined number of the second series of values, the predetermined number being a default number that is modifiable to a user-specified number.

57. The system of claim 21 in which the output field includes a cell in a table.
58. The system of claim 21 in which receiving the rule includes receiving at least a row of a rule table, the row corresponding to a rule case, and having an output including one or more or a combination of the input values, a predetermined value, or a value computed from one or more of the input values.
59. The system of claim 58 in which the rule case includes one or more of: having an input value equal to a threshold, having an input value above a threshold, having an input value below a threshold, having an input value belonging to a set of values, having an input value matching a pattern of values, having a relationship to another input value, having a relationship to an output value of another set of rules, or having a relationship to a value in a memory.
60. The system of claim 21 in which the input dataset includes records having values for scalar variables and vector variables.
61. The system of claim 60 in which at least one of the records includes an array for storing a predetermined number of records, the predetermined number being a default number that is modifiable to a user-specified number.
62. The system of claim 60 in which at least one of the records includes an internal reference table to define key relationships to sub-records in the at least one of the records.
63. The system of claim 62 also including, in response to a rule, a means for producing the second series of values for the output variable in the output dataset based on the key relationships in the internal reference table.
64. The system of claim 21 also including, in response to a rule case in a rule, a means for triggering the rule case to produce a value for the output variable in the output dataset.

65. The system of claim 64 in which triggering the rule case includes triggering the rule case based on a scalar value in the input dataset satisfying the at least one rule case in the rule.
66. The system of claim 64 in which triggering the rule case includes triggering the rule case based on each value in a vector in the input dataset satisfying the at least one rule case in the rule.
67. The system of claim 64 in which triggering the rule case includes triggering the rule case based on an output of an aggregate function applied to a vector in the input dataset satisfying the at least one rule case in the rule.
68. The system of claim 21 in which generating the transform includes converting each of a plurality of rule cases in the rule to a logical expression to form a plurality of logical expressions, and compiling the plurality of logical expressions into computer-executable code.
69. The system of claim 68 in which compiling the plurality of logical expressions includes one or more of combining expressions, optimizing individual expressions, and optimizing groups of expressions.
70. The system of claim 64 in which triggering the rule case includes satisfying predefined one or more criteria by the input values.
71. The system of claim 21 also including:
  - means for triggering two or more rule cases belonging to the rule in response to the input values; and
  - means for producing, by the transform, two or more output values each corresponding to the triggered two or more rule cases as the second series of values.
72. The system of claim 71 also including a means operable by a user to inspect the two or more rule cases that triggered in response to the input values.

73. The system of claim 21 in which the at least one rule case includes one or more criteria which when satisfied by the input values causes production by the transform of one or more output values.
74. The system of claim 21 also including a means for storing a scalar value in the output dataset corresponding to an operation performed on the second series of values.
75. The system of claim 21 also including a means for displaying the second series of values in a tool tip upon a cursor being positioned over the output field.

1/11

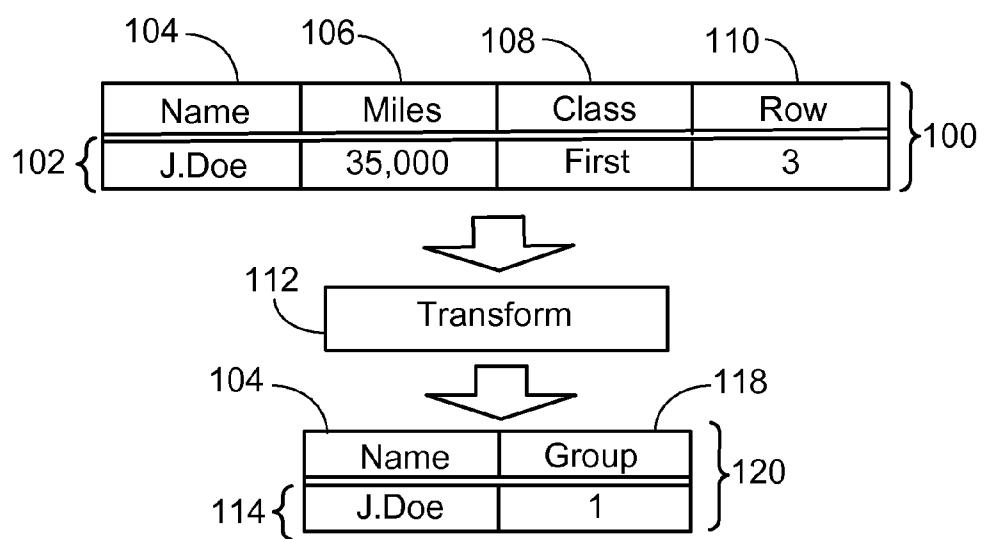


FIG. 1

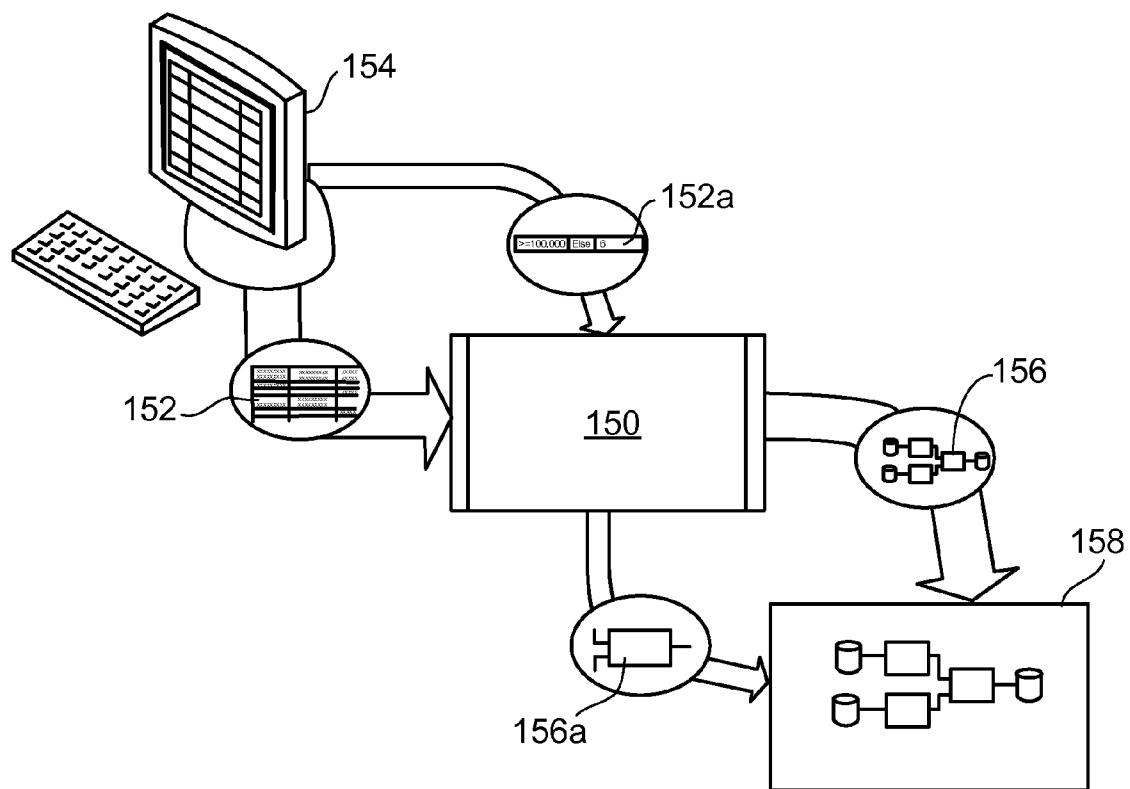


FIG. 2

3/11

The table defines a mapping (200) for boarding groups (152) based on several criteria (202, 204, 206, 208). The columns are:

- Total frequent flyer miles (202)
- Current year frequent flyer miles (204)
- Class of seat (206)
- Row of seat (208)
- Boarding Group (152)

The rows are labeled 210a through 210h. The data is as follows:

Total frequent flyer miles	Current year frequent flyer miles	Class of seat	Row of seat	Boarding Group
210a $\geq 1,000,000$				1
210b		first		1
210c $\geq 100,000$	$\geq$ lat year frequent flyer miles			2
210d		business		2
210e		else	$\leq 10$	2
210f		same	$\leq 40$	3
210g		same	$\leq 50$	
210h		same	else	5

FIG. 3

4/11

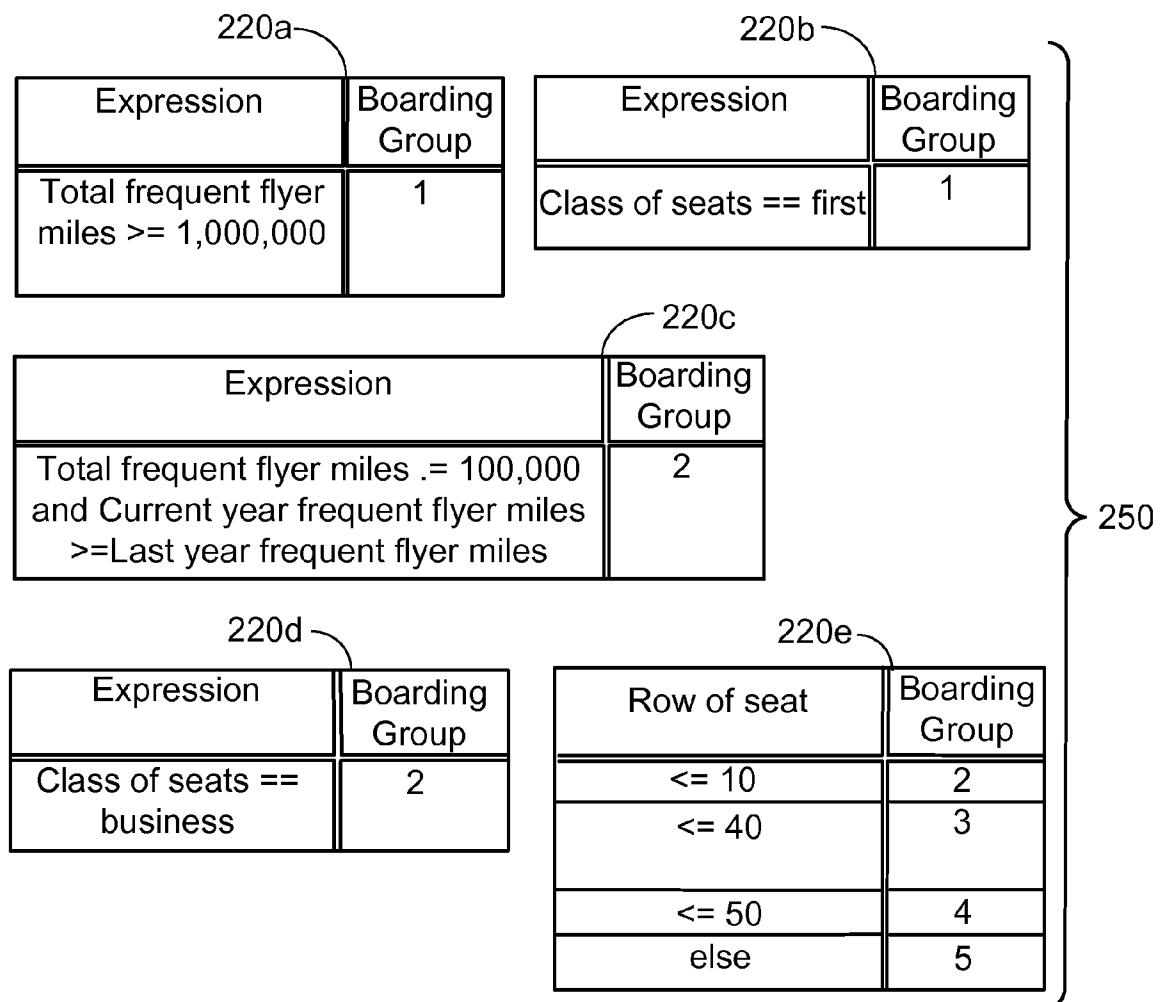


FIG. 4

5/11

508

500

504

Outputs

Triggers

► Year	☒ Vehicle Value	☒ Premium Vehicle	If True	► Valid Vehicle	► Vehicle Validation Message
1 >2004	<15000	No	*Seat Material =	No	"Vehicle value too low to have..."
2	any	any	*Four Wheel...		"Vehicle value too low to have..."
3	any	Yes	*Power Wind...		"Recent premium vehicles al..."
4	any	any	*On Star = Yes		"Vehicle value too low to have..."
5 <2000	any	any	*Ipod Jack = Ye		"Vehicle too old to have had a..."
6 <1990	any	any	*CD Changer = Y		"Vehicle too old to have had a..."
7 <2004	any	any	*Blue Tooth E...		"Vehicle to old to have had a..."
8 else	any	Yes	↓		"Recent premium vehicles ha..."
9 <1995	any	any	*Car Color in M		"Vehicle too old to have had ...."
10 <1900	any	any	any		"Vehicle too old to insure."
11 any	any	any	any	Yes	default value
*					

FIG. 5

6/11

	Name	Type	Name of Constants	Default Value	Output Type	Max Count	Technical Name in Dataset: "Dataset 1"	Descr
1	<input checked="" type="checkbox"/> Name Valid	String	0	"Yes"	Write Once	20	name_valid	
2	<input checked="" type="checkbox"/> Name Validation Message	String	0	null	Accumulator	20	name_valid	
	New variable							

FIG. 6

7/11

700

8 Validate Request and Compute Quote Ruleset

File Edit View Ruleset Tools Window Help

Test Mode: File Test Deployment ▶

Default Deployment

Results R: Validate Vehicle\* DEP: Default Deployment\* R: Validate Person\* Home

Fire Many Rule – “Validate Person” (704) Ruleset – “Validate R

ABOUT

CASES GRID (All true cases will fire)

708

712

Fire Many

Triggers		Accumulator Outputs	
	IF TRUE	Person Validation Message	Times Fired
1	*string_length(First Name)<3	“Invalid first name”	53
2	*string_length(Last Name)<3	“Invalid last name”	11
3	*Year Of Birth<1937	“Invalid too old to insure”	1083
4	*Year Of Birth<1991	“Individual too young to insure”	322
5	*(2007 - Year Of Birth) < Number 0	“Individual is too young to have that”	91
6	any		8440
*			

FIG. 7

8/11

File Test Results				Ruleset – “Validate Request and Compute Quote” (808)			
Validate Request and Compute Quote Ruleset				File Test			
File Edit View Ruleset Tools Window Help				Deployment: Default Deployment ▶ M1 of 100 [Z] [X] ▶			
Home RS: Validate Request and Compute Quote [Z] Results				Test Mode: File Test ▶ Deployment: Default Deployment ▶ M1 of 100 [Z] [X] ▶			
File Test Results				Validation Message			
First Name	Last Name	Val	Validation Message	First Name	Last Name	Val	Validation Message
1 LENARD	SITSLER	Y	“Invalid first name”	1 LENARD	SITSLER	Y	“Invalid first name”
2 TANGELA	SCHEPPI	N	Vehicle too old to have had metallic paint. Vehicle to old to insure. Vehicle too old to have had blue t...	2 TANGELA	SCHEPPI	N	Vehicle too old to have had metallic paint. Vehicle to old to insure. Vehicle too old to have had blue t...
3 KENDALL	ALCANT	N	Recent premium vehicles always have power windows.	3 KENDALL	ALCANT	N	Recent premium vehicles always have power windows.
800	801	816	813	800	801	816	813

8  
E  
G

9/11

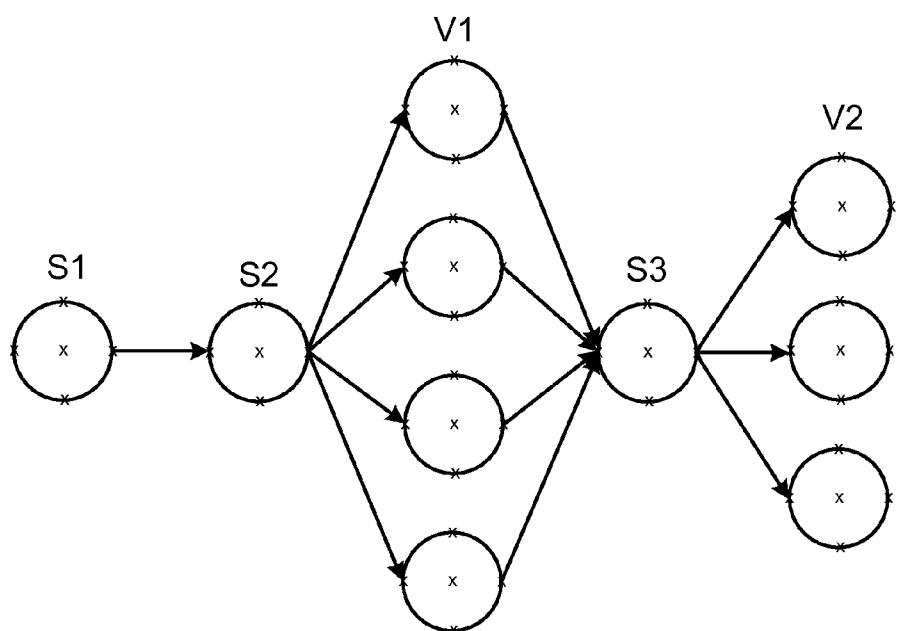


FIG. 9

```
record
  decimal 1 (2) num_drivers; // Number of Drivers
  record
    string (20) name;           // Driver Name
    string (1) sex;             // Driver Sex
    decimal (3) age;            // Driver Age
    decimal (2) points;         // Driver Points
  end driver [num_drivers];
  decimal (2) num_vehicles;    // Number of Vehicles
  record
    decimal (10) value;         // Vehicle Value
    decimal (2) age;             // Vehicle Age
    decimal (1) airbag;          // Vehicle has Air bag
    decimal (1) has_belt;        // Vehicle has Seat Belts
    decimal (20) driver;         // Vehicle Primary Driver
  end driver [num_vehicles];
  decimal (2) geo_risk;         // Geographic Adjustment
end;

952
954
950
```

FIG. 10A

956

```
<record>
  <drivers>
    <driver name="Fred" sex="M" age="35" points="12" />
    <driver name="Wilma" sex="F" age="33" points="9" />
    <driver name="Pebbles" sex="F" age="21" points="14" />
    <driver name="BamBam" sex="M" age="22" points="13" />
  <drivers>
  <vehicles>
    <vehicle value="1000" age="4" air_bag="0" has_belt="0" drivers="Fred" />
    <vehicle value="1500" age="2" air_bag="1" has_belt="1" drivers="Pebbles" />
    <vehicle value="1200" age="6" air_bag="0" has_belt="1" drivers="Wilma" />
  <vehicles>
    <geo_risk value="80" />
  <record>
```

**FIG. 10B**