



# A generic and customizable framework for the design of ETL scenarios

Panos Vassiliadis<sup>a</sup>, Alkis Simitsis<sup>b</sup>, Panos Georgantas<sup>b</sup>, Manolis Terrovitis<sup>b</sup>, Spiros Skiadopoulos<sup>b</sup>

<sup>a</sup>Department of Computer Science, University of Ioannina, Ioannina, Greece

<sup>b</sup>Department of Electrical and Computer Engineering, National Technical University of Athens, Athens, Greece

---

## Abstract

Extraction–transformation–loading (ETL) tools are pieces of software responsible for the extraction of data from several sources, their cleansing, customization and insertion into a data warehouse. In this paper, we delve into the logical design of ETL scenarios and provide a generic and customizable framework in order to support the DW designer in his task. First, we present a metamodel particularly customized for the definition of ETL activities. We follow a workflow-like approach, where the output of a certain activity can either be stored persistently or passed to a subsequent activity. Also, we employ a declarative database programming language, LDL, to define the semantics of each activity. The metamodel is generic enough to capture any possible ETL activity. Nevertheless, in the pursuit of higher reusability and flexibility, we specialize the set of our generic metamodel constructs with a palette of frequently used ETL activities, which we call *templates*. Moreover, in order to achieve a uniform extensibility mechanism for this library of built-ins, we have to deal with specific language issues. Therefore, we also discuss the mechanics of template instantiation to concrete activities. The design concepts that we introduce have been implemented in a tool, ARKTOS II, which is also presented.

© 2004 Elsevier B.V. All rights reserved.

*Keywords:* Data warehousing; ETL

---

## 1. Introduction

Data warehouse operational processes normally compose a labor-intensive workflow, involving

data extraction, transformation, integration, cleaning and transport. To deal with this workflow, specialized tools are already available in the market [1–4], under the general title *Extraction–Transformation–Loading* (ETL) tools. To give a general idea of the functionality of these tools we mention their most prominent tasks, which include (a) the identification of relevant information at the source side, (b) the extraction of this information,

---

*E-mail addresses:* [pvassil@cs.uoi.gr](mailto:pvassil@cs.uoi.gr) (P. Vassiliadis), [asimi@dbnet.ece.ntua.gr](mailto:asimi@dbnet.ece.ntua.gr) (A. Simitsis), [pgeor@dbnet.ece.ntua.gr](mailto:pgeor@dbnet.ece.ntua.gr) (P. Georgantas), [mter@dbnet.ece.ntua.gr](mailto:mter@dbnet.ece.ntua.gr) (M. Terrovitis), [spiros@dbnet.ece.ntua.gr](mailto:spiros@dbnet.ece.ntua.gr) (S. Skiadopoulos).

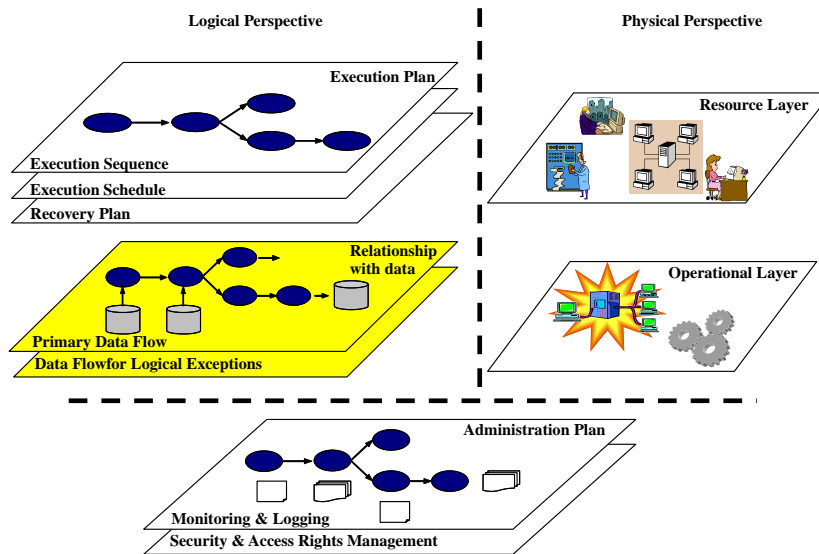


Fig. 1. Different perspectives for an ETL workflow.

(c) the customization and integration of the information coming from multiple sources into a common format, (d) the cleaning of the resulting data set on the basis of database and business rules, and (e) the propagation of the data to the data warehouse and/or data marts.

If we treat an ETL scenario as a composite workflow, in a traditional way, its designer is obliged to define several of its parameters (Fig. 1). Here, we follow a multi-perspective approach that enables to separate these parameters and study them in a principled approach. We are mainly interested in the *design* and *administration* parts of the lifecycle of the overall ETL process, and we depict them at the upper and lower part of Fig. 1, respectively. At the top of Fig. 1, we are mainly concerned with the static design artifacts for a workflow environment. We will follow a traditional approach and group the design artifacts into *logical* and *physical*, with each category comprising its own *perspective*. We depict the logical perspective on the left-hand side of Fig. 1, and the physical perspective on the right-hand side. At the logical perspective, we classify the design artifacts that give an abstract description of the workflow environment. First, the designer is responsible for

defining an *execution plan* for the scenario. The definition of an execution plan can be seen from various perspectives. The *execution sequence* involves the specification of which activity runs first, second, and so on, which activities run in parallel, or when a semaphore is defined so that several activities are synchronized at a rendezvous point. ETL activities normally run in batch, so the designer needs to specify an *execution schedule*, i.e., the time points or events that trigger the execution of the scenario as a whole. Finally, due to system crashes, it is imperative that there exists a *recovery plan*, specifying the sequence of steps to be taken in the case of failure for a certain activity (e.g., retry to execute the activity, or undo any intermediate results produced so far). On the right-hand side of Fig. 1, we can also see the physical perspective, involving the registration of the actual entities that exist in the real world. We will reuse the terminology of [5] for the physical perspective. The *resource layer* comprises the definition of roles (human or software) that are responsible for executing the activities of the workflow. The *operational layer*, at the same time, comprises the software modules that implement the design entities of the logical perspective in the real world.

In other words, the activities defined at the logical layer (in an abstract way) are materialized and executed through the specific software modules of the physical perspective.

At the lower part of Fig. 1, we are dealing with the tasks that concern the administration of the workflow environment and their dynamic behavior at runtime. First, an *administration plan* should be specified, involving the notification of the administrator either on-line (monitoring) or off-line (logging) for the status of an executed activity, as well as the security and authentication management for the ETL environment.

We find that research has not dealt with the definition of data-centric workflows to the entirety of its extent. In the ETL case, for example, due to the data centric nature of the process, the designer must deal with the *relationship of the involved activities with the underlying data*. This involves the definition of a *primary data flow* that describes the route of data from the sources towards their final destination in the data warehouse, as they pass through the activities of the scenario. Also, due to possible quality problems of the processed data, the designer is obliged to define a *data flow for logical exceptions*, i.e., a flow for the problematic data, i.e., the rows that violate integrity or business rules. It is the combination of the execution sequence and the data flow that generates the semantics of the ETL workflow: the data flow defines what each activity does and the execution plan defines in which order and combination.

In this paper, we work in the internals of the data flow of ETL scenarios. First, we present a metamodel particularly customized for the definition of ETL activities. We follow a workflow-like approach, where the output of a certain activity can either be stored persistently or passed to a subsequent activity. Moreover, we employ a declarative database programming language, LDL, to define the semantics of each activity. The metamodel is generic enough to capture any possible ETL activity; nevertheless, reusability and ease-of-use dictate that we can do better in aiding the data warehouse designer in his task. In this pursuit of higher reusability and flexibility, we specialize the set of our generic metamodel constructs with a palette of frequently used ETL

activities, which we call *templates*. Moreover, in order to achieve a uniform extensibility mechanism for this library of built-ins, we have to deal with specific language issues: thus, we also discuss the mechanics of template instantiation to concrete activities. The design concepts that we introduce have been implemented in a tool, ARKTOS II, which is also presented.

Our contributions can be listed as follows:

- First, we define a formal metamodel as an abstraction of ETL processes at the logical level. The data stores, activities and their constituent parts are formally defined. An activity is defined as an entity with possibly more than one input schemata, an output schema and a parameter schema, so that the activity is populated each time with its proper parameter values. The flow of data from producers towards their consumers is achieved through the usage of *provider relationships* that map the attributes of the former to the respective attributes of the latter. A serializable combination of ETL activities, provider relationships and data stores constitutes an ETL scenario.
- Second, we provide a reusability framework that complements the genericity of the metamodel. Practically, this is achieved from a set of “built-in” specializations of the entities of the metamodel layer, specifically tailored for the most frequent elements of ETL scenarios. This palette of template activities will be referred to as *template layer* and it is characterized by its extensibility; in fact, due to language considerations, we provide the details of the mechanism that instantiates templates to specific activities.
- Finally, we discuss *implementation* issues and we present a graphical tool, ARKTOS II that facilitates the design of ETL scenarios, based on our model.

This paper is organized as follows. In Section 2, we present a generic model of ETL activities. Section 3 describes the mechanism for specifying and materializing template definitions of frequently used ETL activities. Section 4 presents ARKTOS II, a prototype graphical tool. In Section 5, we survey related work. In Section 6, we make a

general discussion on the completeness and general applicability of our approach. Section 7 offers conclusions and presents topics for future research. Short versions of parts of this paper have been presented in [6,7].

## 2. Generic model of ETL activities

The purpose of this section is to present a formal logical model for the activities of an ETL environment. This model abstracts from the technicalities of monitoring, scheduling and logging while it concentrates on the flow of data from the sources towards the data warehouse through the composition of activities and data stores. The full layout of an ETL scenario, involving activities, recordsets and functions can be modeled by a graph, which we call the *architecture graph*. We employ a uniform, graph-modeling framework for both the modeling of the internal structure of activities and the ETL scenario at large, which enables the treatment of the ETL environment from different viewpoints. First, the architecture graph comprises all the activities and data stores of a scenario, along with their components. Second, the architecture graph captures the data flow within the ETL environment. Finally, the information on the typing of the involved entities and the regulation of the execution of a scenario, through specific parameters are also covered.

### 2.1. Graphical notation and motivating example

Being a graph, the architecture graph of an ETL scenario comprises nodes and edges. The involved data types, function types, constants, attributes, activities, recordsets, parameters and functions constitute the nodes of the graph. The different kinds of relationships among these entities are modeled as the edges of the graph. In Fig. 2, we give the graphical notation for all the modeling constructs that will be presented in the sequel.

*Motivating example:* To motivate our discussion, we will present an example involving the propagation of data from a certain source  $S_1$ , towards a data warehouse DW through intermediate recordsets. These recordsets belong to a data

staging area (DSA)<sup>1</sup> DS. The scenario involves the propagation of data from the table PARTSUPP of source  $S_1$  to the data warehouse DW. Table DW.PARTSUPP (PKEY, SOURCE, DATE, QTY, COST) stores information for the available quantity (QTY) and cost (COST) of parts (PKEY) per source (SOURCE). The data source  $S_1$ .PARTSUPP (PKEY, DATE, QTY, COST) records the supplies from a specific geographical region, e.g., Europe. All the attributes, except for the dates are instances of the Integer type. The scenario is graphically depicted in Fig. 3 and involves the following transformations.

1. First, we transfer via FTP\_PS1 the snapshot from the source  $S_1$ .PARTSUPP to the file DS.PS1\_NEW of the DSA.<sup>2</sup>
2. In the DSA, we maintain locally a copy of the snapshot of the source as it was at the previous loading (we assume here the case of the incremental maintenance of the DW, instead of the case of the initial loading of the DW). The recordset DS.PS1\_NEW (PKEY, DATE, QTY, COST) stands for the last transferred snapshot of  $S_1$ .PARTSUPP. By detecting the difference of this snapshot with the respective version of the previous loading, DS.PS1\_OLD (PKEY, DATE, QTY, COST), we can derive the newly inserted rows in  $S_1$ .PARTSUPP. Note that the difference activity that we employ, namely Diff\_PS1, checks for differences only on the primary key of the recordsets; thus, we ignore here any possible deletions or updates for the attributes COST, QTY of existing rows. Any not newly inserted row is rejected and so, it is propagated to Diff\_PS1\_REJ that stores all the rejected rows. The schema of Diff\_PS1\_REJ is identical to the input schema of the activity Diff\_PS1.

<sup>1</sup>In data warehousing terminology a DSA is an intermediate area of the data warehouse, specifically destined to enable the transformation, cleaning and integration of source data, before being loaded to the warehouse.

<sup>2</sup>The technical points, like FTP, are mostly employed to show what kind of problems someone has to deal with in a practical situation, rather than to relate this kind of physical operations to a logical model. In terms of logical modelling this is a simple passing of data from one site to another.

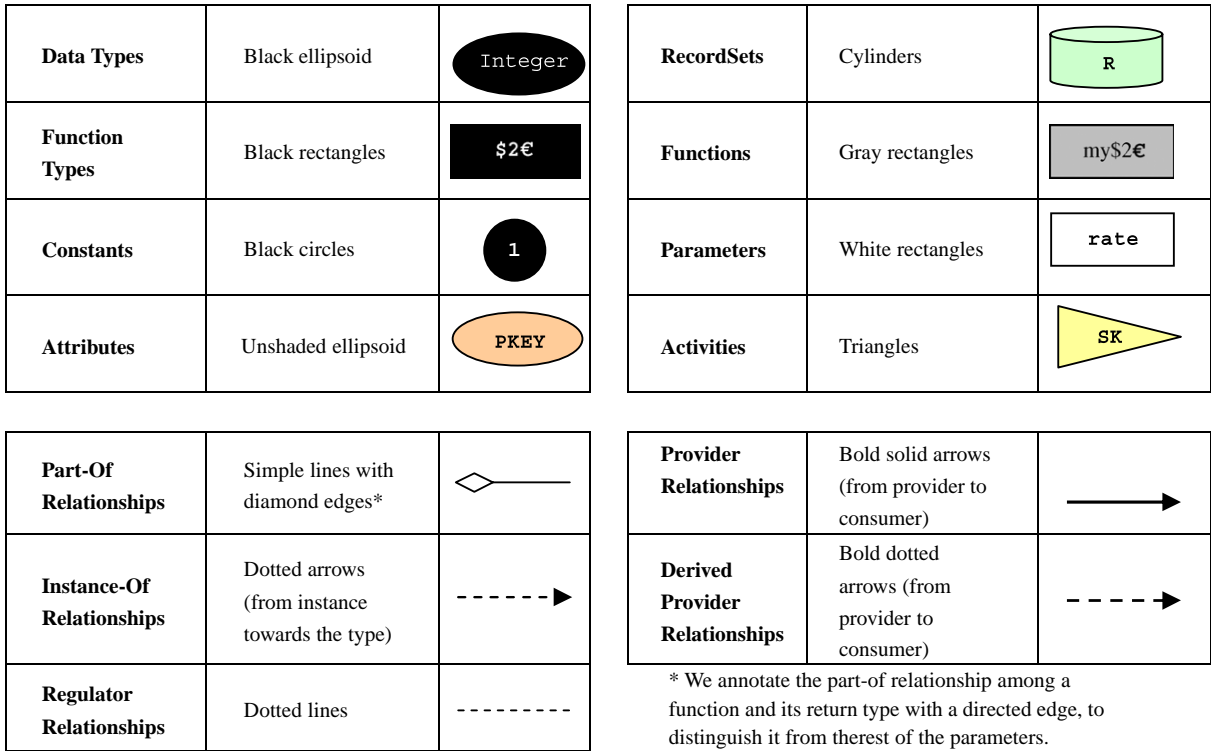


Fig. 2. Graphical notation for the architecture graph.

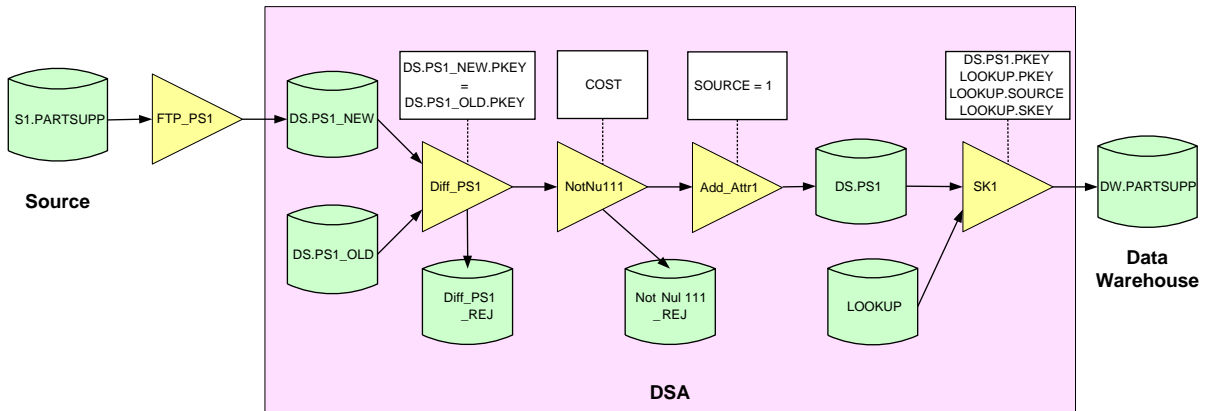


Fig. 3. Bird's-eye view of the motivating example.

3. The rows that pass the activity *Diff\_PS1* are checked for null values of the attribute *COST* through the activity *NotNu111*. Rows having a NULL value for their *COST* are kept in the

*Diff\_PS1\_REJ* recordset for further examination by the data warehouse administrator.

4. Although we consider the data flow for only one source, namely *S1*, the data warehouse can

clearly have more sources for part supplies. In order to keep track of the source of each row entering into the DW, we need to add a ‘flag’ attribute, namely SOURCE, indicating the respective source. This is achieved through the activity `Add_Attr1`. We store the rows that stem from this process in the recordset `DS.PS1` (PKEY, SOURCE, DATE, QTY, COST).

5. Next, we assign a surrogate key on PKEY. In the data warehouse context, it is common tactics to replace the keys of the production systems with a uniform key, which we call a *surrogate key* [8]. The basic reasons for this replacement are performance and semantic homogeneity. Textual attributes are not the best candidates for indexed keys and thus, they need to be replaced by integer keys. At the same time, different production systems might use different keys for the same object, or the same key for different objects, resulting in the need for a global replacement of these values in the data warehouse. This replacement is performed through a lookup table of the form `L` (PRODKEY, SOURCE, SKEY). The SOURCE column is due to the fact that there can be synonyms in the different sources, which are mapped to different objects in the data warehouse. In our case, the activity that performs the surrogate key assignment for the attribute PKEY is SK1. It uses the lookup table `LOOKUP` (PKEY, SOURCE, SKEY). Finally, we populate the data warehouse with the output of the previous activity.

The role of rejected rows depends on the peculiarities of each ETL scenario. If the designer needs to administrate these rows further, then he/she should use intermediate storage recordsets with the burden of an extra I/O cost. If the rejected rows should not have a special treatment, then the best solution is to be ignored; thus, in this case we avoid overloading the scenario with any extra storage recordset. In our case, we annotate only two of the presented activities with a destination for rejected rows. Out of these, while `NotNull1_REJ` absolutely makes sense as a placeholder for problematic rows having non-acceptable NULL values, `Diff_PS1_REJ` is presented for demonstration reasons only.

Finally, before proceeding, we would like to stress that we do not anticipate a manual construction of the graph by the designer; rather, we employ this section to clarify how the graph will look once constructed. To assist a more automatic construction of ETL scenarios, we have implemented the ARKTOS II tool that supports the designing process through a friendly GUI. We present ARKTOS II in Section 4.

## 2.2. Preliminaries

In this subsection, we will introduce the formal modeling of data types, data stores and functions, before proceeding to the modeling of ETL activities.

*Elementary entities:* We assume the existence of a countable set of *data types*. Each data type  $T$  is characterized by a name and a domain, i.e., a countable set of values, called  $\text{dom}(T)$ . The values of the domains are also referred to as *constants*.

We also assume the existence of a countable set of *attributes*, which constitute the most elementary granules of the infrastructure of the information system. Attributes are characterized by their name and data type. The domain of an attribute is a subset of the domain of its data type. Attributes and constants are uniformly referred to as *terms*.

A *schema* is a finite *list* of attributes. Each entity that is characterized by one or more schemata will be called *structured entity*. Moreover, we assume the existence of a special family of schemata, all under the general name of *NULL schema*, determined to act as placeholders for data which are not to be stored permanently in some data store. We refer to a family instead of a single NULL schema, due to a subtle technicality involving the number of attributes of such a schema (this will become clear in the sequel).

*Recordsets:* We define a *record* as the instantiation of a schema to a list of values belonging to the domains of the respective schema attributes. We can treat any data structure as a recordset provided that there are ways to *logically*

restructure it into a flat, typed record schema. Formally, a recordset is characterized by its name, its (logical) schema and its (physical) extension (i.e., a finite set of records under the recordset schema). If we consider a schema  $S = [A_1, \dots, A_k]$ , for a certain recordset, its extension is a mapping  $S = [A_1, \dots, A_k] \rightarrow \text{dom}(A_1) \times \dots \times \text{dom}(A_k)$ . Thus, the extension of the recordset is a finite subset of  $\text{dom}(A_1) \times \dots \times \text{dom}(A_k)$  and a record is the instance of a mapping  $\text{dom}(A_1) \times \dots \times \text{dom}(A_k) \rightarrow [x_1, \dots, x_k]$ ,  $x_i \in \text{dom}(A_i)$ . In the rest of this paper we will mainly deal with the two most popular types of recordsets, namely *relational tables* and *record files*. A *database* is a finite set of relational tables.

*Functions.* We assume the existence of a countable set of built-in system *function types*. A function type comprises a name, a finite list of *parameter data types*, and a single *return data type*. A *function* is an instance of a function type. Consequently, it is characterized by a name, a list of input parameters and a parameter for its return value. The data types of the parameters of the generating function type also define (a) the data types of the parameters of the function and (b) the legal candidates for the function parameters (i.e., attributes or constants of a suitable data type).

### 2.3. Activities

Activities are the backbone of the structure of any information system. We adopt the WfMC terminology [9] for processes/programs and we will call them *activities* in the sequel. An activity is an amount of “work which is processed by a combination of resource and computer applications” [9]. In our framework, activities are logical abstractions representing parts or full modules of code.

The execution of an activity is performed from a particular program. Normally, ETL activities will be either performed in a black-box manner by a dedicated tool, or they will be expressed in some language (e.g., PL/SQL, Perl, C). Still, we want to deal with the general case of ETL activities. We employ an abstraction of the source code of an activity, in the form of an LDL statement. Using

LDL, we avoid dealing with the peculiarities of a particular programming language. Once again, we want to stress that the presented LDL description is intended to capture the semantics of each activity, instead of the way these activities are actually implemented.

An *elementary activity* is formally described by the following elements:

- *Name*: A unique identifier for the activity.
- *Input schemata*: A finite set of one or more input schemata that receives data from the data providers of the activity.
- *Output schema*: A schema that describes the placeholder for the rows that pass the check performed by the elementary activity.
- *Rejections schema*: A schema that describes the placeholder for the rows that do not pass the check performed by the activity, or their values are not appropriate for the performed transformation.
- *Parameter list*: A set of pairs which act as regulators for the functionality of the activity (the target attribute of a foreign key check, for example). The first component of the pair is a name and the second is a schema, an attribute, a function or a constant.
- *Output operational semantics*: An LDL statement describing the content passed to the output of the operation, with respect to its input. This LDL statement defines (a) the operation performed on the rows that pass through the activity and (b) an implicit mapping between the attributes of the input schema(ta) and the respective attributes of the output schema.
- *Rejection operational semantics*: An LDL statement describing the rejected records, in a sense similar to the output operational semantics. This statement is by default considered to be the complement of the output operational semantics, except if explicitly defined differently.

There are two issues that we would like to elaborate on, here:

- *NULL schemata*: Whenever we do not specify a data consumer for the output or rejection schemata, the respective NULL schema

(involving the correct number of attributes) is implied. This practically means that the data targeted for this schema will neither be stored to some persistent data store, nor will they be propagated to another activity, but they will simply be ignored.

- *Language issues:* Initially, we used to specify the semantics of activities with SQL statements. Still, although clear and easy to write and understand, SQL is rather hard to use if one is to perform rewriting and composition of statements. Thus, we have supplemented SQL with LDL [10], a logic programming, declarative language as the basis of our scenario definition. LDL is a Datalog variant based on a Horn-clause logic that supports recursion, complex objects and negation. In the context of its implementation in an actual deductive database management system, LDL++ [11], the language has been extended to support external functions, choice, aggregation (and even, user-defined aggregation), updates and several other features.

#### 2.4. Relationships in the architecture graph

In this subsection, we will elaborate on the different kinds of relationships that the entities of an ETL scenario have. Whereas these entities are modeled as the nodes of the architecture graph, relationships are modeled as its edges. Due to their diversity, before proceeding, we list these types of relationships along with the related terminology that we will use in this paper. The graphical

notation of entities (nodes) and relationships (edges) is presented in Fig. 2.

*Part-of relationships.* These relationships involve attributes and parameters and relate them to the respective activity, recordset or function to which they belong.

*Instance-of relationships.* These relationships are defined among a data/function type and its instances.

*Provider relationships.* These are relationships that involve attributes with a provider–consumer relationship.

*Regulator relationships.* These relationships are defined among the parameters of activities and the terms that populate these activities.

*Derived provider relationships.* A special case of provider relationships that occurs whenever output attributes are computed through the composition of input attributes and parameters. Derived provider relationships can be deduced from a simple rule and do not originally constitute a part of the graph.

In the rest of this subsection, we will detail the notions pertaining to the relationships of the Architecture Graph; the knowledgeable reader is referred to Section 2.5 where we discuss the issue of scenarios. We will base our discussions on a part of the scenario of the motivating example (presented in Section 2.1), including activity SK<sub>1</sub>.

*Data types and instance-of relationships:* To capture typing information on attributes and

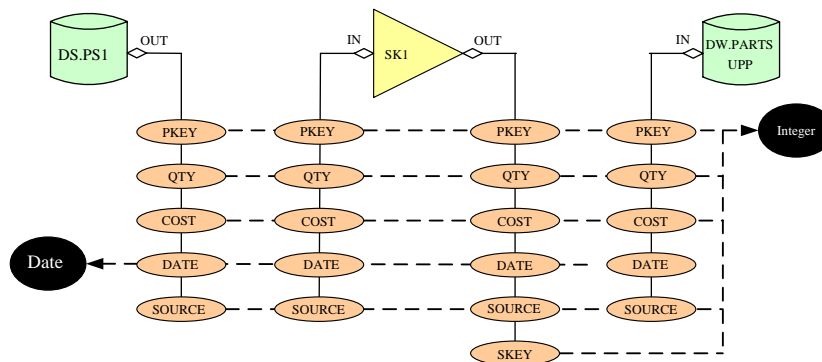


Fig. 4. Instance-of relationships of the architecture graph.



functions, the architecture graph comprises data and function types. Instantiation relationships are depicted as dotted arrows that stem from the instances and head toward the data/function types. In Fig. 4, we observe the attributes of the two activities of our example and their correspondence to two data types, namely integer and date. For reasons of presentation, we merge several instantiation edges so that the figure does not become too crowded.

*Attributes and part-of relationships:* The first thing to incorporate in the architecture graph is the structured entities (activities and recordsets) along with all the attributes of their schemata. We choose to avoid overloading the notation by incorporating the schemata per se; instead we apply a direct part-of relationship between an activity node and the respective attributes. We annotate each such relationship with the name of the schema (by default, we assume a IN, OUT, PAR, REJ tag to denote whether the attribute belongs to the input, output, parameter or rejection

schema of the activity, respectively). Naturally, if the activity involves more than one input schemata, the relationship is tagged with an  $IN_i$  tag for the  $i$ th input schema. We also incorporate the functions along with their respective parameters and the part-of relationships among the former and the latter. We annotate the part-of relationship with the return type with a directed edge, to distinguish it from the rest of the parameters.

Fig. 5 depicts a part of the motivating example. In terms of part-of relationships, we present the decomposition of (a) the recordsets DS.PS<sub>1</sub>, LOOKUP, DW.PARTSUPP and (b) the activity SK<sub>1</sub> and the attributes of its input and output schemata. Note the tagging of the schemata of the involved activity. We do not consider the rejection schemata in order to avoid crowding the picture. Also note, how the parameters of the activity are also incorporated in the architecture graph. Activity SK<sub>1</sub> has five parameters: (a) PKEY, which stands for the production key to be replaced, (b) SOURCE, which stands for an integer

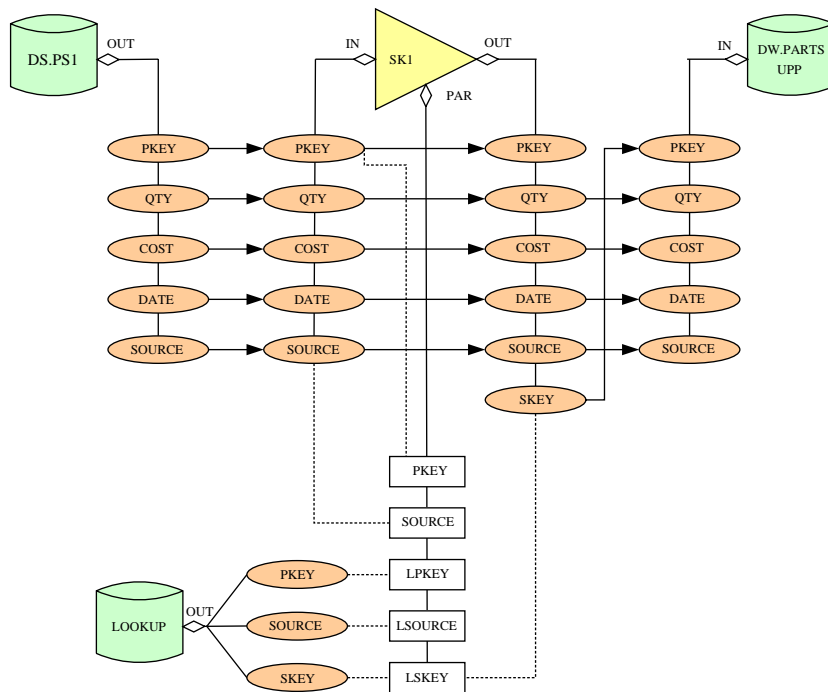


Fig. 5. Part-of regulator and provider relationships of the architecture graph.

value that characterizes which source's data are processed, (c) LPKEY, which stands for the attribute of the lookup table which contains the production keys, (d) LSOURCE, which stands for the attribute of the lookup table which contains the source value (corresponding to the aforementioned SOURCE parameter), (e) LSKEY, which stands for the attribute of the lookup table which contains the surrogate keys.

*Parameters and regulator relationships:* Once the part-of and instantiation relationships have been established, it is time to establish the regulator relationships of the scenario. In this case, we link the parameters of the activities to the terms (attributes or constants) that populate them. We depict regulator relationships with simple dotted edges.

In the example of Fig. 5 we can also observe how the parameters of activity  $SK_1$  are populated through regulator relationships. The parameters *in* and *out* are mapped to the respective terms through regulator relationships. All the parameters of  $SK_1$ , namely PKEY, SOURCE, LPKEY, LSOURCE and LSKEY, are mapped to the respective attributes of either the activity's input schema or the employed lookup table LOOKUP. The parameter LSKEY deserves particular attention. This parameter is (a) populated from the attribute SKEY of the lookup table and (b) used to populate the attribute SKEY of the output schema of the activity. Thus, two regulator relationships are related with parameter LSKEY, one for each of the aforementioned attributes. The existence of a regulator relationship among a parameter and an output attribute of an activity normally denotes that some external data provider is employed in order to derive a new attribute, through the respective parameter.

*Provider relationships:* The flow of data from the data sources towards the data warehouse is performed through the composition of activities in a larger scenario. In this context, the input for an activity can be either a persistent data store, or another activity. Usually, this applies for the output of an activity, too. We capture the passing of data from *providers* to *consumers* by a *provider*

*relationship* among the attributes of the involved schemata.

Formally, a *provider relationship* is defined by the following elements:

- *Name:* A unique identifier for the provider relationship.
- *Mapping:* An ordered pair. The first part of the pair is a term (i.e., an attribute or constant) acting as a provider and the second part is an attribute acting as the consumer.

The mapping need not necessarily be 1:1 from provider to consumer attributes, since an input attribute can be mapped to more than one consumer attributes. Still, the opposite does not hold. Note that a consumer attribute can also be populated by a constant, in certain cases.

In order to achieve the flow of data from the providers of an activity towards its consumers, we need the following three groups of provider relationships:

1. A mapping between the input schemata of the activity and the output schema of their data providers. In other words, for each attribute of an input schema of an activity, there must exist an attribute of the data provider, or a constant, which is mapped to the former attribute.
2. A mapping between the attributes of the activity input schemata and the activity output (or rejection, respectively) schema.
3. A mapping between the output or rejection schema of the activity and the (input) schema of its data consumer.

The mappings of the second type are internal to the activity. Basically, they can be derived from the LDL statement for each of the output/rejection schemata. As far as the first and the third types of provider relationships are concerned, the mappings must be provided during the construction of the ETL scenario. This means that they are either (a) by default assumed by the order of the attributes of the involved schemata or (b) hard-coded by the user. Provider relationships are depicted with bold solid arrows that stem from the provider and end in the consumer attribute.

Observe Fig. 5. The flow starts from table  $DS.PS_1$  of the data staging area. Each of the attributes of this table is mapped to an attribute of the input schema of activity  $SK_1$ . The attributes of the input schema of the latter are subsequently mapped to the attributes of the output schema of the activity. The flow continues to  $DW.PARTSUPP$ . Another interesting thing is that during the data flow, new attributes are generated, resulting on new streams of data, whereas the flow seems to stop for other attributes. Observe the rightmost part of Fig. 5 where the values of attribute  $PKEY$  are not further propagated (remember that the reason for the application of a surrogate key transformation is to replace the production keys of the source data to a homogeneous surrogate for the records of the data warehouse, which is independent of the source they have been collected from). Instead of the values of the production key, the values from the attribute  $SKEY$  will be used to denote the unique identifier for a part in the rest of the flow.

In Fig. 6, we depict the LDL definition of this part of the motivating example. The three rules correspond to the three categories of provider

relationships, previously discussed: the first rule explains how the data from the  $DS.PS_1$  recordset are fed into the input schema of the activity, the second rule explains the semantics of activity (i.e., how the surrogate key is generated) and, finally, the third rule shows how the  $DW.PARTSUPP$  recordset is populated from the output schema of the activity  $SK_1$ .

*Derived provider relationships:* As we have already mentioned, there are certain output attributes that are computed through the composition of input attributes and parameters. A *derived provider relationship* is another form of provider relationship that captures the flow from the input to the respective output attributes.

Formally, assume that (a) *source* is a term in the architecture graph, (b) *target* is an attribute of the output schema of an activity  $A$  and (c)  $x, y$  are parameters in the parameter list of  $A$  (not necessarily different). Then, a derived provider relationship  $pr(\text{source}, \text{target})$  exists iff the following regulator relationships (i.e., edges) exist:  $rr_1(\text{source}, x)$  and  $rr_2(y, \text{target})$ .

```

addskKey_in1(A_IN1_PKEY,A_IN1_DATE,A_IN1_QTY,A_IN1_COST,A_IN1_SOURCE)
  ds_ps1(A_OUT_PKEY,A_OUT_DATE,A_OUT_QTY,A_OUT_COST,A_OUT_SOURCE),
  A_OUT_PKEY=A_IN1_PKEY,
  A_OUT_DATE=A_IN1_DATE,
  A_OUT_QTY=A_IN1_QTY,
  A_OUT_COST=A_IN1_COST,
  A_OUT_SOURCE=A_IN1_SOURCE.

addskKey_out(A_OUT_PKEY,A_OUT_DATE,A_OUT_QTY,A_OUT_COST,A_OUT_SOURCE,A_OUT_SKEY)
  addskKey_in1(A_IN1_PKEY,A_IN1_DATE,A_IN1_QTY,A_IN1_COST,A_IN1_SOURCE),
  lookup(A_IN1_SOURCE,A_IN1_PKEY,A_OUT_SKEY),
  A_OUT_PKEY=A_IN1_PKEY,
  A_OUT_DATE=A_IN1_DATE,
  A_OUT_QTY=A_IN1_QTY,
  A_OUT_COST=A_IN1_COST,
  A_OUT_SOURCE=A_IN1_SOURCE.

dw_partsupp(PKEY,DATE,QTY,COST,SOURCE)
  addskKey_out(A_OUT_PKEY,A_OUT_DATE,A_OUT_QTY,A_OUT_COST,A_OUT_SOURCE,A_OUT_SKEY),
  DATE=A_IN1_DATE,
  QTY=A_IN1_QTY,
  COST=A_IN1_COST
  SOURCE=A_IN1_SOURCE,
  PKEY=A_IN1_SKEY.

```

**NOTE:** For reasons of readability we do not replace the  $A_{in}$  attribute names with the activity name, i.e.,  $A_{OUT\_PKEY}$  should be  $diffPS1\_OUT\_PKEY$ .

Fig. 6. LDL specification of the motivating example.

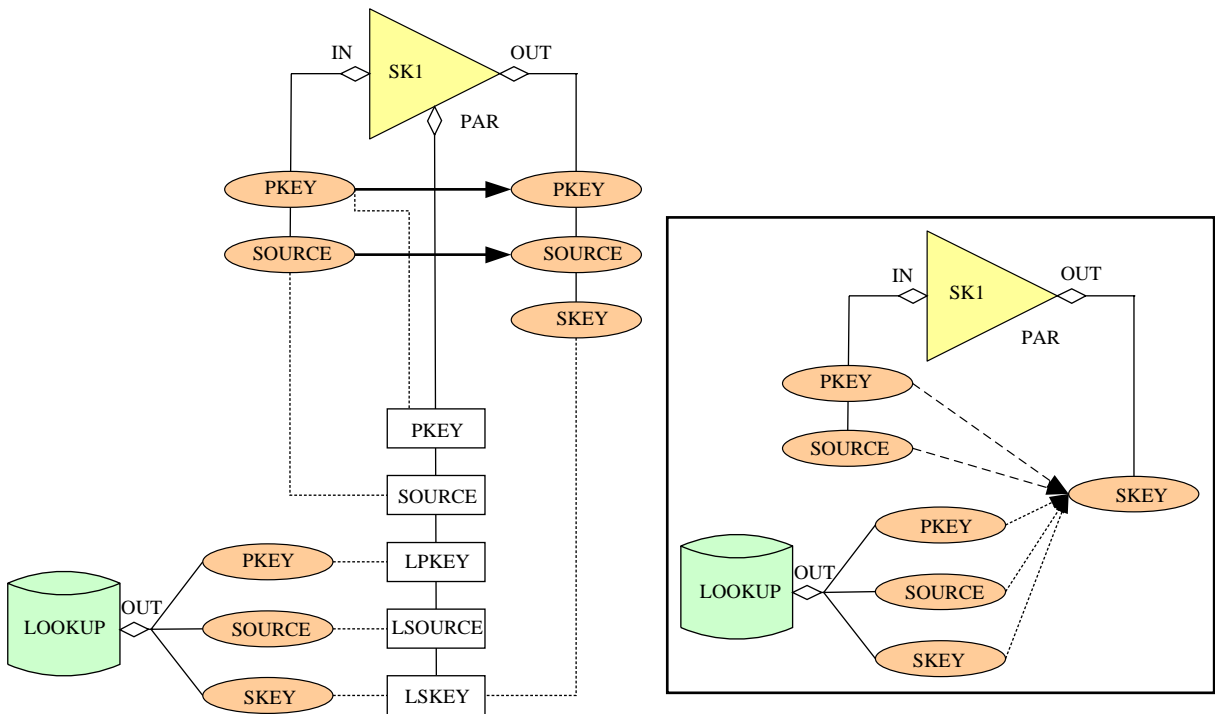


Fig. 7. Derived provider relationships of the architecture graph: the original situation on the left and the derived provider relationships on the right.

Intuitively, the case of derived relationships models the situation where the activity computes a new attribute in its output. In this case, the produced output depends on all the attributes that populate the parameters of the activity, resulting in the definition of the corresponding derived relationship.

Observe Fig. 7, where we depict a small part of our running example. The left side of the figure depicts the situation where only provider relationships exist. The legend in the right side of Fig. 7 depicts how we compute the derived provider relationships between the parameters of the activity and the computed output attribute SKEY. The meaning of these five relationships is that  $SK_1.OUT.SKEY$  is not computed only from attribute  $LOOKUP.SKEY$ , but from the combination of all the attributes that populate the parameters.

One can also assume different variations of derived provider relationships such as (a) relation-

ships that do not involve constants (remember that we have defined *source* as a term); (b) relationships involving only attributes of the same/different activity (as a measure of internal complexity or external dependencies); (c) relationships relating attributes that populate only the same parameter (e.g., only the attributes  $LOOKUP.SKEY$  and  $SK_1.OUT.SKEY$ ).

### 2.5. Scenarios

A scenario is an enumeration of activities along with their source/target recordsets and the respective provider relationships for each activity. An ETL scenario consists of the following elements:

- *Name*: A unique identifier for the scenario.
- *Activities*: A finite list of activities. Note that by employing a list (instead of e.g., a set) of activities, we impose a total ordering on the execution of the scenario.

	Entity	Model-specific	Scenario-specific
Built-in	Data Types	$D^I$	$D$
	Function Types	$F^I$	$F$
	Constants	$C^I$	$C$
User-provided	Attributes	$\Omega^I$	$\Omega$
	Functions	$\Phi^I$	$\Phi$
	Schemata	$S^I$	$S$
	RecordSets	$RS^I$	$RS$
	Activities	$A^I$	$A$
	Provider Relationships	$Pr^I$	$Pr$
	Part-Of Relationships	$Po^I$	$Po$
	Instance-Of Relationships	$Io^I$	$Io$
	Regulator Relationships	$Rr^I$	$Rr$
	Derived Provider Relationships	$Dr^I$	$Dr$

Fig. 8. Formal definition of domains and notation.

- *Recordsets*: A finite set of recordsets.
- *Targets*: A special-purpose subset of the recordsets of the scenario, which includes the final destinations of the overall process (i.e., the data warehouse tables that must be populated by the activities of the scenario).
- *Provider relationships*: A finite list of provider relationships among activities and recordsets of the scenario.

In our modeling, a scenario is a set of activities, deployed along a graph in an execution sequence that can be linearly serialized. For the moment, we do not consider the different alternatives for the ordering of the execution; we simply require that a total order for this execution is present (i.e., each activity has a discrete execution priority).

In terms of formal modeling of the architecture graph, we assume the infinitely countable, mutually disjoint sets of names (i.e., the values of which respect the unique name assumption) of column *model-specific* in Fig. 8. As far as a specific scenario is concerned, we assume their respective finite subsets, depicted in column *scenario-specific* in Fig. 8. Data types, function types and constants are considered *built-in*'s of the system, whereas the rest of the entities are provided by the user (*user provided*).

Formally, the *architecture graph* of an ETL scenario is a graph  $G(V,E)$  defined as follows:

$$V = DUFUCU\Omega\phi USRSUA$$

$$E = PrUPoUIoURrUDr.$$

In the sequel, we treat the terms architecture graph and scenario interchangeably. The reasoning for the term ‘architecture graph’ goes all the way down to the fundamentals of conceptual modeling. As mentioned in [12], conceptual models are the means by which designers conceive, architect, design, and build software systems. These conceptual models are used in the same way that blueprints are used in other engineering disciplines during the early stages of the lifecycle of artificial systems, which involves the creation of their architecture. The term ‘architecture graph’ expresses the fact that the graph that we employ for the modeling of the data flow of the ETL scenario is practically acting as a blueprint of the architecture of this software artifact.

Moreover, we assume the following *integrity constraints* for a scenario:

*Static constraints:*

- All the weak entities of a scenario (i.e., attributes or parameters) should be defined within a *part-of* relationship (i.e., they should have a container object).
- All the mappings in provider relationships should be defined among terms (i.e., attributes or constants) of the same data type.

*Data flow constraints:*

- All the attributes of the input schema(ta) of an activity should have a provider.

- Resulting from the previous requirement, if some attribute is a parameter in an activity A, the container of the attribute (i.e., recordset or activity) should precede A in the scenario.
- All the attributes of the schemata of the target recordsets should have a data provider.

Summarizing, in this section, we have presented a generic model for the modeling of the data flow for ETL workflows. In the next section, we will proceed to detail how this generic model can be accompanied by a customization mechanism, in order to provide higher flexibility to the designer of the workflow.

### 3. Templates for ETL activities

In this section, we present the mechanism for exploiting template definitions of frequently used ETL activities. The general framework for the exploitation of these templates is accompanied with the presentation of the language-related issues for template management and appropriate examples.

### 3.1. General framework

Our philosophy during the construction of our metamodel was based on two pillars: (a) *genericity*, i.e., the derivation of a simple model, powerful to capture ideally all the cases of ETL activities and (b) *extensibility*, i.e., the possibility of extending the built-in functionality of the system with new, user-specific templates.

The genericity doctrine was pursued through the definition of a rather simple activity metamodel, as described in Section 2. Still, providing a single metaclass for all the possible activities of an ETL environment is not really enough for the designer of the overall process. A richer “language” should be available, in order to describe the structure of the process and facilitate its construction. To this end, we provide a palette of *template* activities, which are specializations of the generic metamodel class.

Observe Fig. 9 for a further explanation of our framework. The lower layer of Fig. 9, namely *schema layer*, involves a specific ETL scenario. All the entities of the schema layer are *instances* of the classes Data Type, Function Type,

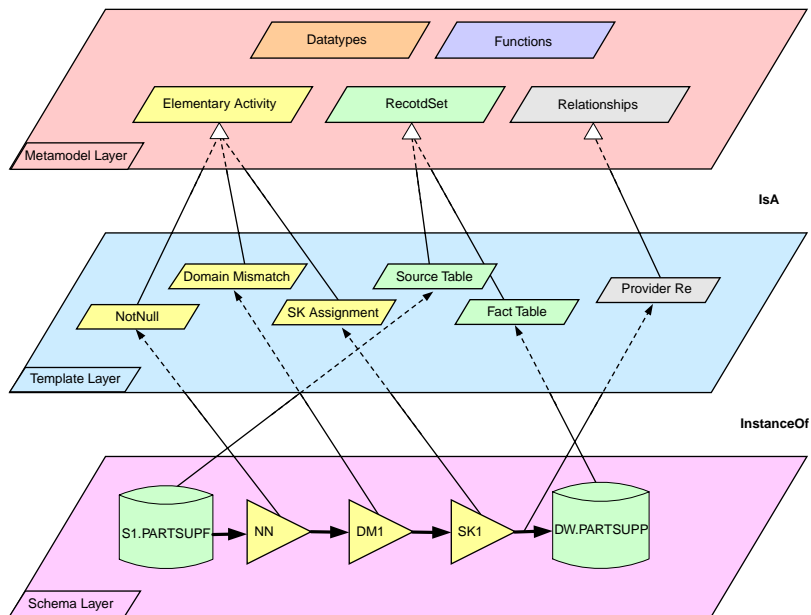


Fig. 9. The metamodel for the logical entities of the ETL environment.

Elementary Activity, RecordSet and Relationship. Thus, as one can see on the upper part of Fig. 9, we introduce a meta-class layer, namely *metamodel layer* involving the aforementioned classes. The linkage between the metamodel and the schema layers is achieved through instantiation (InstanceOf) relationships. The metamodel layer implements the aforementioned genericity desideratum: the classes which are involved in the metamodel layer are generic enough to model any ETL scenario, through the appropriate instantiation.

Still, we can do better than the simple provision of a metalayer and an instance layer. In order to make our metamodel truly useful for practical cases of ETL activities, we enrich it with a set of ETL-specific constructs, which constitute a subset of the larger metamodel layer, namely the *template layer*. The constructs in the template layer are also meta-classes, but they are quite customized for the regular cases of ETL activities. Thus, the classes of the template layer are specializations (i.e., subclasses) of the generic classes of the metamodel layer (depicted as *IsA* relationships in Fig. 9). Through this customization mechanism, the designer can pick the instances of the schema layer from a much richer palette of constructs; in this setting, the entities of the schema layer are instantiations, not only of the respective classes of the metamodel layer, but also of their subclasses in the template layer.

In the example of Fig. 9 the concept DW.PARTSUPP must be populated from a certain source  $S_1$ .PARTSUPP. Several operations must intervene during the propagation. For instance in Fig. 9, we check for null values and domain violations, and we assign a surrogate key. As one can observe, the recordsets that take part in this scenario are instances of class RecordSet (belonging to the metamodel layer) and specifically of its subclasses Source Table and Fact Table. Instances and encompassing classes are related through links of type InstanceOf. The same mechanism applies to all the activities of the scenario, which are (a) instances of class Elementary Activity and (b) instances of one of its subclasses, depicted in Fig. 9. Relationships do not escape this rule either. For instance, observe how the provider links from the concept  $S_1$ .PS toward the concept DW.PARTSUPP are related to class Provider Relationship through the appropriate InstanceOf links.

As far as the class Recordset is concerned, in the template layer we can specialize it to several subclasses, based on orthogonal characteristics, such as whether it is a file or RDBMS table, or whether it is a source or target data store (as in Fig. 9). In the case of the class Relationship, there is a clear specialization in terms of the five classes of relationships which have already been mentioned in Section 2 (i.e., Provider, Part-Of, Instance-Of, Regulator and Derived Provider).

<b>Filters</b>	<b>Unary operations</b>	<b>Binary operations</b>
- Selection ( $\sigma$ )	- Push	- Union (U)
- Not null (NN)	- Aggregation ( $\gamma$ )	- Join ( $\bowtie$ )
- Primary key violation (PK)	- Projection ( $\Pi$ )	- Diff ( $\Delta$ )
- Foreign key violation (FK)	- Function application (f)	- Update Detection ( $\Delta$ UPD)
- Unique value (UN)	- Surrogate key assignment (SK)	
- Domain mismatch (DM)	- Tuple normalization (N)	
	- Tuple denormalization(DN)	
	<b>File operations</b>	<b>Transfer operations</b>
	- EBCDIC to ASCII conversion (EB2AS)	- Ftp (FTP)
	- Sort file (Sort)	- Compress / Decompress (Z/dZ)
		- Encrypt / Decrypt (Cr/dCr)

Fig. 10. Template activities, along with their graphical notation symbols, grouped by category.

Following the same framework, class `Elementary Activity` is further specialized to an extensible set of reoccurring patterns of ETL activities, depicted in Fig. 10. As one can see on the top side of Fig. 9, we group the template activities in five major logical groups. We do not depict the grouping of activities in subclasses in Fig. 9, in order to avoid overloading the figure; instead, we depict the specialization of class `Elementary Activity` to three of its subclasses whose instances appear in the employed scenario of the schema layer. We now proceed to present each of the aforementioned groups in more detail.

The first group, named *filters*, provides checks for the satisfaction (or not) of a certain condition. The semantics of these filters are the obvious (starting from a generic selection condition and proceeding to the check for null values, primary or foreign key violation, etc.). The second group of template activities is called *unary operations* and except for the most generic push activity (which simply propagates data from the provider to the consumer), consists of the classical aggregation and function application operations along with three data warehouse specific transformations (surrogate key assignment, normalization and denormalization). The third group consists of classical *binary operations*, such as union, join and difference of recordsets/activities as well as with a special case of difference involving the detection of updates. Except for the aforementioned template activities, which mainly refer to logical transformations, we can also consider the case of physical operators that refer to the application of physical transformations to whole files/tables. In the ETL context, we are mainly interested in operations like *transfer operations* (ftp, compress/decompress, encrypt/decrypt) and *file operations* (EBCDIC to ASCII, sort file).

Summarizing, the metamodel layer is a set of generic entities, able to represent any ETL scenario. At the same time, the genericity of the metamodel layer is complemented with the extensibility of the template layer, which is a set of “built-in” specializations of the entities of the metamodel layer, specifically tailored for the most

frequent elements of ETL scenarios. Moreover, apart from this “built-in”, ETL-specific extension of the generic metamodel, if the designer decides that several ‘patterns’, not included in the palette of the template layer, occur repeatedly in his data warehousing projects, he can easily fit them into the customizable template layer through a specialization mechanism.

### 3.2. Formal definition and usage of template activities

Once the template layer has been introduced, the obvious issue that is raised is its linkage with the employed declarative language of our framework. In general, the broader issue is the usage of the template mechanism from the user; to this end, we will explain the substitution mechanism for templates in this subsection and refer the interested reader to [13] for a presentation of the specific templates that we have constructed.

A *template activity* is formally defined by the following elements:

- *Name*: A unique identifier for the template activity.
- *Parameter list*: A set of names which act as regulators in the expression of the semantics of the template activity. For example, the parameters are used to assign values to constants, create dynamic mapping at instantiation time, etc.
- *Expression*: A declarative statement describing the operation performed by the instances of the template activity. As with elementary activities, our model supports LDL as the formalism for the expression of this statement.
- *Mapping*: A set of bindings, mapping input to output attributes, possibly through intermediate placeholders. In general, mappings at the template level try to capture a default way of propagating incoming values from the input towards the output schema. These default bindings are easily refined and possibly rearranged at instantiation time.

The template mechanism we use is a substitution mechanism, based on macros, that facilitates the



automatic creation of LDL code. This simple notation and instantiation mechanism permits the easy and fast registration of LDL templates. In the rest of this section, we will elaborate on the notation, instantiation mechanisms and template taxonomy particularities.

### 3.2.1. Notation

Our template notation is a simple language featuring five main mechanisms for dynamic production of LDL expressions: (a) *variables* that are replaced by their values at instantiation time; (b) a function that returns the *arity* of an input, output or parameter schema; (c) *loops*, where the loop body is repeated at instantiation time as many times as the iterator constraint defines; (d) *keywords* to simplify the creation of unique predicate and attribute names; and, finally, (e) *macros* which are used as syntactic sugar to simplify the way we handle complex expressions (especially in the case of variable size schemata).

*Variables:* We have two kinds of variables in the template mechanism: *parameter variables* and *loop iterators*. Parameter variables are marked with a @ symbol at their beginning and they are replaced by user-defined values at instantiation time. A list of an arbitrary length of parameters is denoted by @<parameter name>[ ]. For such lists, the user has to explicitly or implicitly provide their length at instantiation time. Loop iterators, on the other hand, are implicitly defined in the loop constraint. During each loop iteration, all the properly marked appearances of the iterator in the loop body are replaced by its current value (similarly to the way the C preprocessor treats #DEFINE statements). Iterators that appear marked in loop body are instantiated even when they are a part of another string or of a variable name. We mark such appearances by enclosing them with \$. This functionality enables referencing all the values of a parameter list and facilitates the creation of an arbitrary number of pre-formatted strings.

*Functions:* We employ a built-in function, `arityOf(<input/output/parameter schema>)`,

which returns the arity of the respective schema, mainly in order to define upper bounds in loop iterators.

*Loops:* Loops are a powerful mechanism that enhances the genericity of the templates by allowing the designer to handle templates with unknown number of variables and with unknown arity for the input/output schemata. The general form of loops is

```
[(simple constraint)]{(loop body)},
```

where simple constraint has the form:

```
<lower bound> <comparison operator> (iterator)
<comparison operator> <upper bound>.
```

We consider only linear increase with step equal to 1, since this covers most possible cases. Upper bound and lower bound can be arithmetic expressions involving `arityOf()` function calls, variables and constants. Valid arithmetic operators are +, -, /, \* and valid comparison operators are <, >, =, all with their usual semantics. If lower bound is omitted, 1 is assumed. During each iteration the loop body will be reproduced and at the same time all the marked appearances of the loop iterator will be replaced by its current value, as described before. Loop nesting is permitted.

*Keywords:* Keywords are used in order to refer to input and output schemata. They provide two main functionalities: (a) they simplify the reference to the input output/schema by using standard names for the predicates and their attributes, and (b) they allow their renaming at instantiation time. This is done in such a way that no different predicates with the same name will appear in the same program, and no different attributes with the same name will appear in the same rule. Keywords are recognized even if they are parts of another string, without a special notation. This facilitates a homogenous renaming of multiple distinct input schemata at template level, to multiple distinct schemata at instantiation, with all of them having unique names in the LDL program scope. For example, if the template is expressed in terms of two different input schemata `a_in1` and `a_in2`, at instantiation time they will be renamed to

Keyword	Usage	Example
a_out a_in	A unique name for the output/input schema of the activity. The predicate that is produced when this template is instantiated has the form: <unique_pred_name>_out (or, _in respectively)	difference3_out difference3_in
A_OUT A_IN	A_OUT/A_IN is used for constructing the names of the a_out/a_in attributes. The names produced have the form: <predicate unique name in upper case>_OUT (or, _IN respectively)	DIFFERENCE3_OUT DIFFERENCE3_IN

Fig. 11. Keywords for templates.

dm1\_in1 and dm1\_in2 so that the produced names will be unique throughout the scenario program. In Fig. 11, we depict the way the renaming is performed at instantiation time.

*Macros:* To make the definition of templates easier and to improve their readability, we introduce a macro to facilitate attribute and variable name expansion. For example, one of the major problems in defining a language for templates is the difficulty of dealing with schemata of arbitrary arity. Clearly, at the template level, it is not possible to pin-down the number of attributes of the involved schemata to a specific value. For example, in order to create a series of names like the following

```
name_theme_1,name_theme_2,...,
name_theme_k
```

we need to give the following expression:

```
[iterator<maxLimit]
{name_theme$iterator$}
[iterator = maxLimit]
{name_theme$iterator$}
```

Obviously, this results in making the writing of templates hard and reduces their readability. To attack this problem, we resort to a simple reusable macro mechanism that enables the simplification of employed expressions. For example, observe the

definition of a template for a simple relational selection:

```
a_out ([i<arityOf (a_out)] {A_OUT_$i$,}
[i = arityOf (a_out)] {A_OUT_$i$,} <-
a_in1 ([i<arityOf (a_in1)]
{A_IN1_$i$,} [i = arityOf (a_in1)]
{A_IN1_$i$,}),
expr ([i<arityOf (@PARAM)]
{@PARAM[$i$,}],
[i = arityOf (@PARAM)]
{@PARAM[$i$,}],
[i<arityOf (a_out)]
{A_OUT_$i$ = A_IN1_$i$,}
[i = arityOf (a_out)]
{A_OUT_$i$ = A_IN1_$i$,}
```

As already mentioned at the syntax for loops, the expression

```
[i<arityOf (a_out)] {A_OUT_$i$,}
[i = arityOf (a_out)] {A_OUT_$i$,}
```

defining the attributes of the output schema a\_out simply wants to list a variable number of attributes that will be fixed at instantiation time. Exactly the same tactics apply for the attributes of the predicate names a\_in1 and expr. Also, the final two lines state that each attribute of the output will be equal to the respective attribute of the input (so that the query is safe), e.g., A\_OUT\_4 = A\_IN1\_4. We can simplify the definition of the template by allowing the designer

to define certain macros that simplify the management of temporary length attribute lists. We employ the following macros:

```

DEFINE INPUT_SCHEMA AS
  [i<arityOf(a_in1)] {A_IN1_$i$,}
  [i = arityOf(a_in1)] {A_IN1_$i$}
DEFINE OUTPUT_SCHEMA AS
  [i<arityOf(a_out)] {A_OUT_$i$,}
  [i = arityOf(a_out)] {A_OUT_$i$}
DEFINE PARAM_SCHEMA AS
  [i<arityOf(@PARAM)] {@PARAM[$i$],}
  [i = arityOf(@PARAM)] {@PARAM[$i$]}
DEFINE DEFAULT_MAPPING AS
  [i<arityOf(a_out)]
    {A_OUT_$i$ = A_IN1_$i$,}
  [i = arityOf(a_out)]
    {A_OUT_$i$ = A_IN1_$i$}

```

Then, the template definition is as follows:

```

a_out(OUTPUT_SCHEMA) <-
  a_in1(INPUT_SCHEMA),
  expr(PARAM_SCHEMA),
  DEFAULT_MAPPING

```

### 3.2.2. Instantiation

Template instantiation is the process where the user chooses a certain template and creates a concrete activity out of it. This procedure requires that the user specifies the schemata of the activity and gives concrete values to the template parameters. Then, the process of producing the respective LDL description of the activity is easily automated. Instantiation order is important in our template creation mechanism, since, as it can easily be seen from the notation definitions, different orders can lead to different results. The instantiation order is as follows:

1. Replacement of macro definitions with their expansions.
2. `arityOf()` functions and parameter variables appearing in loop boundaries are calculated first.
3. Loop productions are performed by instantiating the appearances of the iterators. This leads to intermediate results without any loops.

4. All the rest parameter variables are instantiated.
5. Keywords are recognized and renamed.

We will try to explain briefly the intuition behind this execution order. Macros are expanded first. Step (2) proceeds step (3) because loop boundaries have to be calculated before loop productions are performed. Loops on the other hand, have to be expanded before parameter variables are instantiated, if we want to be able to reference lists of variables. The only exception to this is the parameter variables that appear in the loop boundaries, which have to be calculated first. Notice though, that variable list elements cannot appear in the loop constraint. Finally, we have to instantiate variables before keywords since variables are used to create a dynamic mapping between the input/output schemata and other attributes.

Fig. 12 shows a simple example of template instantiation for the *function application* activity. To understand the overall process better, first observe the outcome of it, i.e., the specific activity which is produced, as depicted in the final row of Fig. 12, labeled *keyword renaming*. The output schema of the activity, `fa12_out`, is the head of the LDL rule that specifies the activity. The body of the rule says that the output records are specified by the conjunction of the following clauses: (a) the input schema `myFunc_in`, (b) the application of function `subtract` over the attributes `COST_IN`, `PRICE_IN` and the production of a value `PROFIT`, and (c) the mapping of the input to the respective output attributes as specified in the last three conjuncts of the rule.

The first row, *template*, shows the initial template as it has been registered by the designer. `@FUNCTION` holds the name of the function to be used, `subtract` in our case, and the `@PARAM[ ]` holds the inputs of the function, which in our case are the two attributes of the input schema. The problem we have to face is that all input, output and function schemata have a variable number of parameters. To abstract from the complexity of this problem, we define four macro definitions, one for each schema (`INPUT_SCHEMA`, `OUTPUT_SCHEMA`, `FUNCTION_INPUT`) along with a macro for the mapping of input to output attributes

<b>Template</b>	<pre> DEFINE INPUT_SCHEMA AS [i&lt;arityOf(a_in1)](A_IN1_\$i\$,) [i=arityOf(a_in1)] {A_IN1_\$i\$}  DEFINE OUTPUT_SCHEMA AS [i&lt;arityOf(a_in)](A_OUT_\$i\$,) [i=arityOf(a_out)]{A_OUT_\$i\$}  DEFINE FUNCTION_INPUT AS [i&lt;arityOf(@PARAM)+1]{@PARAM[\$i\$],}  DEFINE DEFAULT_MAPPING AS [i&lt;arityOf( a_out )] {A_OUT_\$i\$= A_IN1_\$i\$,} [i=arityOf( a_out )] {A_OUT_\$i\$= A_IN1_\$i\$} [i=arityOf( a_out )] {A_OUT_\$i\$= A_IN1_\$i\$}  a_out(OUTPUT_SCHEMA) &lt;- a_in1(INPUT_SCHEMA), @FUNCTION (FUNCTION_INPUT, @FunOutFIELD ),@OUTFIELD=@FunOutFIELD, DEFAULT_MAPPING.         </pre>
<b>Macro Expansion</b>	<pre> a_out([i&lt;arityOf(a_in)+1]{A_OUT_\$i\$,} OUTFIELD) &lt;- a_in([i&lt;arityOf(a_in)]{A_IN_\$i\$,}[i= arityOf(a_in)]{A_IN_\$i\$}), @FUNCTION([i&lt; arityOf(@PARAM[\$i\$])+1]{@PARAM[\$i\$],} OUTFIELD), [i&lt;arityOf(a_in)] {A_OUT_\$i\$=A_IN_\$i\$,} [i=arityOf(a_in)] {A_OUT_\$i\$=A_IN_\$i\$}.         </pre>
<b>Parameter instantiation</b>	<pre> @FUNCTION=f1 @PARAM[1]=A_IN_2 @PARAM[2]=A_IN_3         </pre>
<b>Loop productions</b>	<pre> a_out(A_OUT_1, A_OUT_2, A_OUT_3, OUTFIELD)&lt;- a_in(A_IN_1, A_IN_2, A_IN_3), @FUNCTION(@PARAM[1],@PARAM[2],OUTFIELD), A_OUT_1=A_IN_1,A_OUT_2=A_IN_2,A_OUT_3=A_IN_3.         </pre>
<b>Variable Instantiation</b>	<pre> a_out(A_OUT_1, A_OUT_2, A_OUT_3, OUTFIELD)&lt;- a_in(A_IN_1, A_IN_2, A_IN_3), f1(A_IN_2, A_IN_3,OUTFIELD), A_OUT_1=A_IN_1, A_OUT_2=A_IN_2, A_OUT_3=A_IN_3.         </pre>
<b>Keyword Renaming</b>	<pre> myFunc_out(PKEY_OUT, COST_OUT, PRICE_OUT, PROFIT)&lt;- myFunc_in(PKEY_IN, COST_IN, PRICE_IN), subtract(COST_IN, PRICE_IN, PROFIT), PKEY_OUT=PKEY_IN, COST_OUT=COST_IN, PRICE_OUT=PRICE_IN.         </pre>

Fig. 12. Instantiation procedure.

(DEFAULT\_MAPPING). The second row, *macro expansion*, shows how the template looks after the macros have been incorporated in the template definition. The mechanics of the expansion are straightforward: observe how the attributes of the output schema are specified by the expression `[i<arityOf(a_in)+1]{A_OUT_$i$,}OUTFIELD` as an expansion of the macro `OUTPUT_SCHEMA`. In a similar fashion, the attributes of the input schema and the parameters of the function are also specified; note that the expression for the last attribute in the list is different (to avoid repeating an erroneous comma). The mappings between the input and the output attributes are

also shown in the last two lines of the template. In the third row, *parameter instantiation*, we can see how the parameter variables were materialized at instantiation. In the fourth row, *loop production*, we can see the intermediate results after the loop expansions are done. As it can easily be seen, these expansions must be done before `@PARAM[]` variables are replaced by their values. In the fifth row, *variable instantiation*, the parameter variables have been instantiated creating a default mapping between the input, the output and the function attributes. Finally, in the last row, *keyword renaming*, the output LDL code is presented after the keywords are renamed. Keyword instantiation

is done on the basis of the schemata and the respective attributes of the activity that the user chooses.

### 3.2.3. Taxonomy: simple and program-based templates

The most commonly used activities can be easily expressed by a single predicate template; it is obvious, though, that it would be very inconvenient to restrict activity templates to single predicates. Thus, we separate template activities in two categories, *simple templates*, which cover single-predicate templates and *program-based templates* where many predicates are used in the template definition.

In the case of *simple templates*, the output predicate is bound to the input through a mapping and an expression. Each of the rules for obtaining the output is expressed in terms of the input schemata and the parameters of the activity. In the case of *program templates*, the output of the activity is expressed in terms of its intermediate predicate schemata, as well as its input schemata and its parameters. Program-based templates are often used to define activities that employ constraints like *does-not-belong*, or *does-not-exist*, which need an intermediate negated predicate to be expressed intuitively. This predicate usually describes the conjunction of properties we want to avoid, and then it appears negated in the output predicate. Thus, in general, we allow the construction of a LDL program, with intermediate predicates, in order to enhance intuition. This classification is orthogonal to the logical one of Section 3.1.

*Simple templates*: Formally, the expression of an activity which is based on a certain simple template is produced by a set of rules of the following form:

```
OUTPUT() < – INPUT(), EXPRESSION, MAPPING
```

where `INPUT( )` and `OUTPUT( )` denote the full expression of the respective schemata; in the case of multiple input schemata, `INPUT( )` expresses the conjunction of the input schemata. `MAPPING` denotes any mapping between the input, output, and expression attributes. A default mapping can be explicitly done at the template level, by

specifying equalities between attributes, where the first attribute of the input schema is mapped to the first attribute of the output schema, the second to the respective second one and so on. At instantiation time, the user can change these mappings easily, especially in the presence of the graphical interface. Note also that despite the fact that LDL allows implicit mappings by giving identical names to attributes that must be equal our design choice was to give explicit equalities in order to support the preservation of the names of the attributes of the input and output schemata at instantiation time.

To make ourselves clear, we will demonstrate the usage of simple template activities through an example. Suppose, thus, the case of the `Domain Mismatch` template activity, checking whether the values for a certain attribute fall within a particular range. The rows that abide by the rule pass the check performed by the activity and they are propagated to the output.

Observe Fig. 13, where we present an example of the definition of a template activity and its instantiation in a concrete activity. The first row in Fig. 13 describes the definition of the template activity. There are three parameters; `@FIELD`, for the field that will be checked against the expression, `@Xlow` and `@Xhigh` for the lower and upper limit of acceptable values for attribute `@FIELD`. The expression of the template activity is a simple expression guaranteeing that `@FIELD` will be within the specified range. The second row of Fig. 13 shows the template after the macros are expanded. Let us suppose that the activity named `DM1` materializes the templates parameters that appear in the third row of Fig. 13, i.e., specifies the attribute over which the check will be performed (`A_IN_3`) and the actual ranges for this check (5, 10). The fourth row of Fig. 13 shows the resulting instantiation after keyword renaming is done. The activity includes an input schema `dm1_in`, with attributes `DM1_IN_1`, `DM1_IN_2`, `DM1_IN_3`, `DM1_IN_4` and an output schema `dm1_out` with attributes `DM1_OUT_1`, `DM1_OUT_2`, `DM1_OUT_3`, `DM1_OUT_4`. In this case the parameter `@FIELD` implements a dynamic internal mapping in the template, whereas the `@Xlow`, `@Xigh` parameters provide values for constants. The mapping from

Template	<pre> DEFINE INPUT_SCHEMA AS   [i&lt;arityOf(a_in1)] {A_IN1_\$i\$,}   [i=arityOf(a_in1)] {A_IN1_\$i\$}  DEFINE OUTPUT_SCHEMA AS   [i&lt;arityOf(a_in)] {A_OUT_\$i\$,}   [i=arityOf(a_out)] {A_OUT_\$i\$}  DEFINE DEFAULT_MAPPING AS   [i&lt;arityOf( a_out )] {A_OUT_\$i\$= A_IN1_\$i\$,}   [i=arityOf( a_out )] {A_OUT_\$i\$= A_IN1_\$i\$}  a_out(OUTPUT_SCHEMA) &lt;-   a_in1(INPUT_SCHEMA),   @FIELD &gt;=@Xlow,   @FIELD &lt;= @Xhigh,   DEFAULT_MAPPING.         </pre>
Full Definition of Template	<pre> a_out([i&lt;arityOf(a_out)]{A_OUT_\$i\$,}[i=arityOf(a_out)]{A_OUT_\$i\$})&lt;-   a_in1([i&lt;arityOf(a_in1)]{A_IN1_\$i\$,}[i=arityOf(a_in1)]{A_IN1_\$i\$}),   @FIELD &gt;=@Xlow,   @FIELD &lt;= @Xhigh,   [i&lt;arityOf( a_out )] {A_OUT_\$i\$= A_IN1_\$i\$,}   [i=arityOf( a_out )] {A_OUT_\$i\$= A_IN1_\$i\$}.         </pre>
Parameter instantiation	<pre> @FIELD=A_IN_3 @Xlow =5 @Xhigh = 10         </pre>
Example	<pre> dm1_out(DM1_OUT_1, DM1_OUT_2, DM1_OUT_3, DM1_OUT_4) &lt;-   dm1_in(DM1_IN_1, DM1_IN_2, DM1_IN_3, DM1_IN_4),   DM1_IN_3 &gt;=5, DM1_IN_3&lt;=10,   DM1_OUT_1=DM1_IN_1,   DM1_OUT_2=DM1_IN_2,   DM1_OUT_3=DM1_IN_3,   DM1_OUT_4=DM1_IN_4.         </pre>

Fig. 13. Simple template example: domain mismatch.

the input to the output is hardcoded in the template.

*Program-based templates:* The case of *program-based templates* is somewhat more complex, since the designer who records the template creates more than one predicate to describe the activity. This is usually the case of operations where we want to verify that some data do not have a conjunction of certain properties. Such constraints employ negation to assert that a tuple does not satisfy a predicate, which is defined in a way that it requires that the data that satisfy it have the properties we want to avoid. Such negations can be expressed by *more than one rules*, for the same predicate, that each negates just one property according to the logical rule  $\neg(q \wedge p) \equiv \neg q \vee \neg p$ . Thus, in general, we allow the construction of a LDL program, with intermediate predicates, in order to enhance intuition. For example the *does-not-belong* rela-

tion, which is needed in the *Difference* activity template, needs a second predicate to be expressed intuitively.

Let us see in more detail the case of *Difference*. During the ETL process, one of the very first tasks that we perform is the detection of newly inserted and possibly updated records. Usually, this is physically performed by the comparison of two snapshots (one corresponding to the previous extraction and the other to the current one). To capture this process, we introduce a variation of the classical relational difference operator, which checks for equality only on a certain subset of attributes of the input records. Assume that during the extraction process we want to detect the newly inserted rows. Then, if **PK** is the set of attributes that uniquely identify rows (in the role of a primary key), the newly inserted rows can be found from the expression  $\Delta_{\langle \mathbf{PK} \rangle}(R_{\text{new}}, R)$ . The formal semantics of the difference operator are

<b>Template</b>	<pre> DEFINE INPUT_1_SCHEMA AS   [i&lt;arityOf(a_in1)] {A_IN1_\$i\$,}   [i=arityOf(a_in1)] {A_IN1_\$i\$}  DEFINE INPUTS_2_SCHEMA as   [i&lt;arityOf(a_in2)] {A_IN2_\$i\$,}   [i=arityOf(a_in2)] {A_IN2_\$i\$}  DEFINE OUTPUT_SCHEMA AS   [i&lt;arityOf(a_out)] {A_OUT_\$i\$,}   [i=arityOf(a_out)] {A_OUT_\$i\$}  DEFINE DEFAULT_MAPPING AS   [i&lt;arityOf( a_out )] {A_OUT_\$i\$= A_IN1_\$i\$,}   [i=arityOf( a_out )] {A_OUT_\$i\$= A_IN1_\$i\$}  DEFINE COMMON_MAPPING as   [i&lt;arityOf(@COMMON_IN1)] {@COMMON_IN1[\$i\$]= @COMMON_IN2[\$i\$],}   [i=arityOf(@COMMON_IN1)] {@COMMON_IN1[\$i\$]= @COMMON_IN2[\$i\$]}.  a_out(OUTPUT_SCHEMA) &lt;-   a_in1(INPUTS_1_SCHEMA), a_in2(INPUTS_2_SCHEMA),   semijoin(INPUTS_1_SCHEMA),   DEFAULT_MAPPING.  semijoin(INPUTS_1_SCHEMA) &lt;-   a_in1(INPUTS_1_SCHEMA),   a_in2(INPUTS_2_SCHEMA),   COMMON_MAPPING. </pre>
<b>Full Definition of Template</b>	<pre> a_out([i&lt;arityOf(a_out)]{A_OUT_\$i\$,}[i=arityOf(a_out)]{A_OUT_\$i\$})&lt;-   a_in1([i&lt;arityOf(a_in1)]{A_IN1_\$i\$,}[i=arityOf(a_in1)]{A_IN1_\$i\$}),   a_in2([i&lt;arityOf(a_in2)]{A_IN2_\$i\$,}[i=arityOf(a_in2)]{A_IN2_\$i\$}),   ~semijoin([i&lt;arityOf(a_in1)]{A_IN1_\$i\$,}   [i=arityOf(a_in1)]{A_IN1_\$i\$}),   [i&lt;arityOf( a_out )] {A_OUT_\$i\$= A_IN1_\$i\$,}   [i=arityOf( a_out )] {A_OUT_\$i\$= A_IN1_\$i\$} .  semijoin([i&lt;arityOf(a_in1)]{A_IN1_\$i\$,}[i=arityOf(a_in1)]{A_IN1_\$i\$})&lt;-   a_in1([i&lt;arityOf(a_in1)]{A_IN1_\$i\$,}[i=arityOf(a_in1)]{A_IN1_\$i\$}),   a_in2([i&lt;arityOf(a_in2)]{A_IN2_\$i\$,}[i=arityOf(a_in2)]{A_IN2_\$i\$}),   [i&lt;arityOf( @COMMON_IN1 )] {@COMMON_IN1[\$i\$]=@COMMON_IN2[\$i\$],}   [i=arityOf( @COMMON_IN1 )] {@COMMON_IN1[\$i\$]=@COMMON_IN2[\$i\$]} . </pre>
<b>Parameter instantiation</b>	<pre> @COMMON_IN1_1=A_IN1_1 @COMMON_IN1_2=A_IN1_3  @COMMON_IN2_1=A_IN2_3 @COMMON_IN2_2=A_IN2_2  @COMMON_NUM=2 </pre>
<b>Example</b>	<pre> dF1_out(DF1_OUT_1,DFI1_OUT_2, DFI1_OUT_3) &lt;-   dF1_in1(DF1_IN1_1,DFI1_IN1_2, DFI1_IN1_3),   dF1_in2(DF1_IN2_1,DFI1_IN2_2, DFI1_IN2_3),   ~semijoin(DF1_OUT_1,DFI1_OUT_2, DFI1_OUT_3),   DF1_OUT_1=DF1_IN1_1,   DF1_OUT_2=DF1_IN1_2,   DF1_OUT_3=DF1_IN1_3.  semijoin(DF1_IN1_1,DFI1_IN1_2, DFI1_IN1_3) &lt;-   dF1_in1(DF1_IN1_1,DFI1_IN1_2, DFI1_IN1_3),   dF1_in2(DF1_IN2_1,DFI1_IN2_2, DFI1_IN2_3),   DF1_IN1_1=DF1_IN2_1,   DF1_IN1_2=DF1_IN2_2,   DF1_IN1_3=DF1_IN2_3. </pre>

Fig. 14. Program-based template example: Difference activity.

given by the following calculus-like definition:  
 $\Delta_{\langle A_1, \dots, A_k \rangle}(R, S) = \{x \in R \mid \neg \exists y \in S: x[A_1] = y[A_1] \wedge \dots \wedge x[A_k] = y[A_k]\}$ .

In Fig. 14, we can see the template of the Difference activity and a resulting instantiation for an activity named dF1. As we can see we need the *semijoin* predicate so we can exclude all tuples that satisfy it. Note also that we have two different inputs, which are denoted as distinct by adding a number at the end of the keyword *a\_in*.

#### 4. Implementation

In the context of the aforementioned framework, we have implemented a graphical design tool, ARKTOS II, with the goal of facilitating the design of ETL scenarios, based on our model. In order to design a scenario, the user defines the source and target data stores, the participating activities and the flow of the data in the scenario. These tasks are greatly assisted (a) by a friendly GUI and (b) by a set of reusability templates.

All the details defining an activity can be captured through forms and/or simple point and click operations. More specifically, the user may explore the data sources and the activities already

defined in the scenario, along with their schemata (input, output and parameter). Attributes belonging to an output schema of an activity or a recordset can be “drag’n’dropped” in the input schema of a subsequent activity or recordset, in order to create the equivalent data flow in the scenario. In a similar design manner, one can also set the parameters of an activity. By default the output schema of the activity is instantiated as a copy of the input schema. Then, the user has the ability to modify this setting according to his demands, e.g., by deleting or renaming the proper attributes. The rejection schema of an activity is considered to be a copy of the input schema of the respective activity and the user may determine its physical location, e.g., the physical location of a log file that maintains the rejected rows of the specified activity. Apart from these features, the user can (a) draw the desirable attributes or parameters, (b) define their name and data type, (c) connect them to their schemata, (d) create provider and regulator relationships between them, and (e) draw the proper edges from one node of the architecture graph to another. The system assures the consistency of a scenario, by allowing the user to draw only relationships respecting the restrictions imposed from the

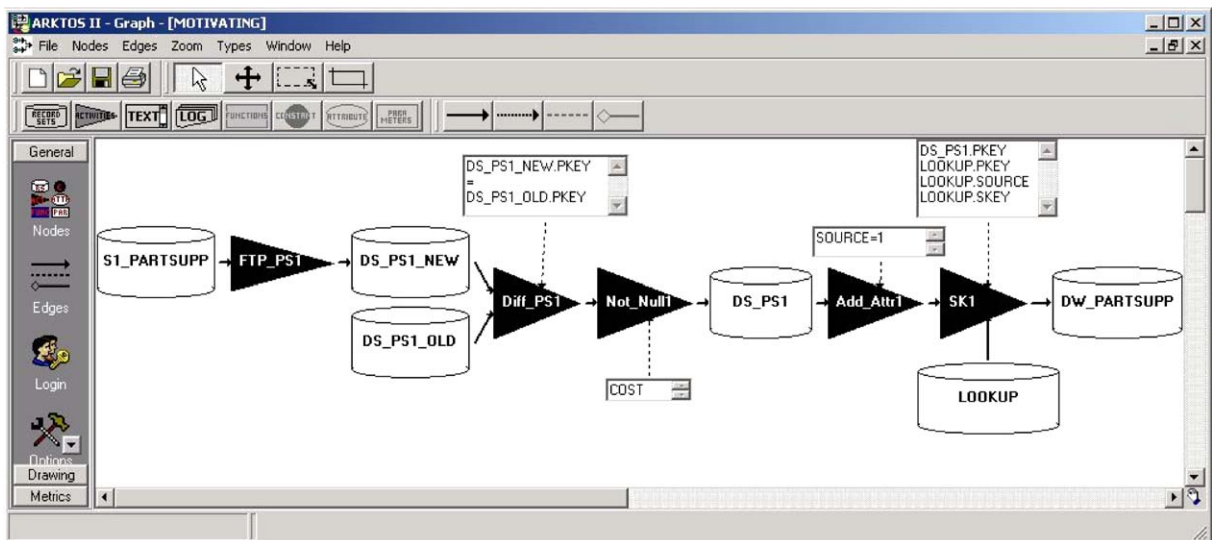


Fig. 15. The motivating example in ARKTOS II.



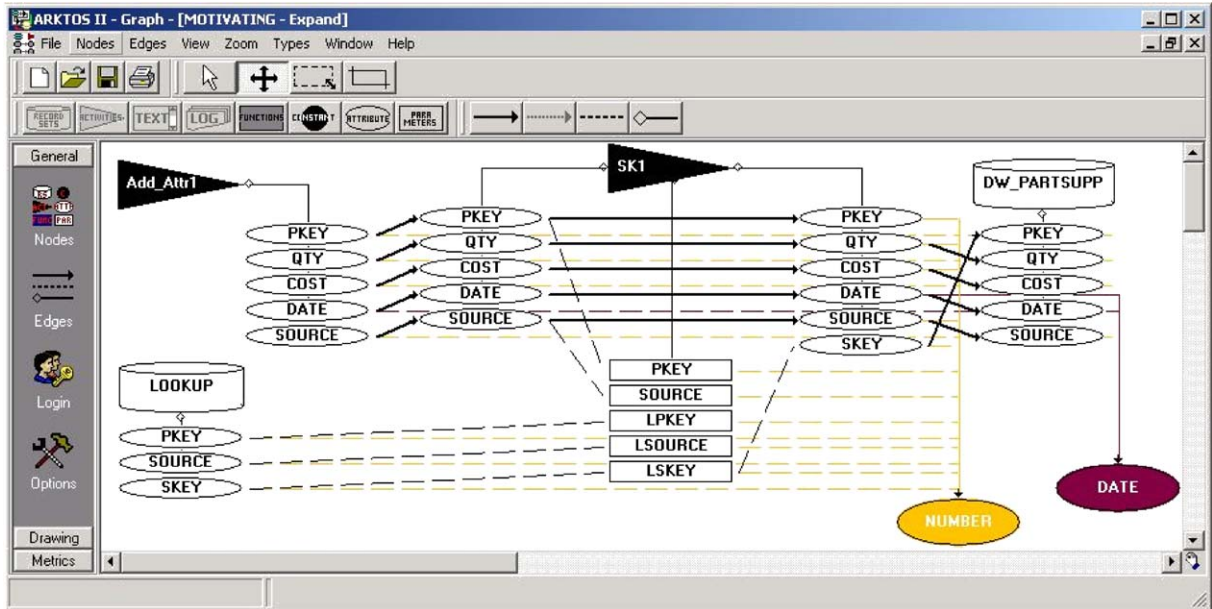


Fig. 16. A detailed zoom-in view of the motivating example.

model. As far as the provider and instance-of relationships are concerned, they are calculated automatically and their display can be turned on or off from an application's menu. Moreover, the system allows the designer to define activities through a form-based interface, instead of defining them through the point-and-click interface. Naturally, the form automatically provides lists with the available recordsets, their attributes, etc. Fig. 15 shows the design canvas of our GUI, where our motivating example is depicted.

ARKTOS II offers zoom-in/zoom-out capabilities, a particularly useful feature in the construction of the data flow of the scenario through inter-attribute "provider" mappings. The designer can deal with a scenario in two levels of granularity: (a) at the *entity* or *zoom-out level*, where only the participating recordsets and activities are visible and their provider relationships are abstracted as edges between the respective entities, or (b) at the *attribute* or *zoom-in level*, where the user can see and manipulate the constituent parts of an activity, along with their respective providers at the attribute level. In Fig. 16, we show a part of the scenario of Fig. 15. Observe (a) how part-of

relationships are expanded to link attributes to their corresponding entities, (b) how provider relationships link attributes to each other, (c) how regulator relationships populate activity parameters and (d) how instance-of relationships relate attributes with their respective data types that are depicted at the lower right part of the figure.

In ARKTOS II, the *customization* principle is supported by the reusability templates. The notion of template is in the heart of ARKTOS II. There are templates for practically every aspect of the model: data types, functions and activities. Templates are extensible; thus, providing the user with the possibility of customizing the environment according to his/her own needs. Especially for activities, which form the core of our model, a specific menu with a set of frequently used ETL Activities is provided. The system has a built-in mechanism responsible for the instantiation of the LDL templates, supported by a graphical form that helps the user define the variables of the template by selecting its values among the appropriate scenario's objects. Another distinctive feature of ARKTOS II is the computation of the scenario's

design quality by employing a set of metrics that are presented in [6], either for the whole scenario or for each activity of it.

The scenarios are stored in ARKTOS II repository (implemented in a relational DBMS); the system allows the user to store, retrieve and reuse existing scenarios. All the metadata of the system involving the scenario configuration, the employed templates and their constituents are stored in the repository. The choice of a relational DBMS for our metadata repository allows its efficient querying as well as the smooth integration with external systems and/or future extensions of ARKTOS II. The connectivity to source and target data stores is achieved through ODBC connections and the tool offers an automatic reverse engineering of their schemata. We have implemented ARKTOS II with Oracle 8.1.7 as basis for our repository and Ms Visual Basic (Release 6) for developing our GUI.

An on-going activity is the coupling of ARKTOS II with state-of-the-art algorithms for individual ETL tasks (e.g., duplicate removal, or surrogate key assignment) and with scheduling and monitoring facilities. Future plans for ARKTOS II involve the extension of data sources to more sophisticated data formats, outside the relational domain, like object-oriented or XML data.

## 5. Related work

In this section, we will report (a) on related commercial studies and tools in the field of ETL, (b) on related efforts in the academia in the issue, and (c) applications of workflow technology in the field of data warehousing.

### 5.1. Commercial studies and tools

In a recent study [14], the authors report that due to the diversity and heterogeneity of data sources, ETL is unlikely to become an open commodity market. The ETL market has reached a size of \$667 millions for year 2001; still the growth rate has reached a rather low 11% (as compared with a rate of 60% growth for year 2000). This is explained by the overall economic downturn environment. In terms of technological

aspects, the main characteristic of the area is the involvement of traditional database vendors with ETL solutions built in the DBMSs. The three major database vendors that practically ship ETL solutions “at no extra charge” are pinpointed: Oracle with Oracle Warehouse Builder [4], Microsoft with Data Transformation Services [3] and IBM with the Data Warehouse Center [1]. Still, the major vendors in the area are Informatica’s Powercenter [2] and Ascential’s DataStage suites [15,16] (the latter being part of the IBM recommendations for ETL solutions). The study goes on to propose future technological challenges/forecasts that involve the integration of ETL with (a) XML adapters, (b) enterprise application integration (EAI) tools (e.g., MQ-Series), (c) customized data quality tools, and (d) the move towards parallel processing of the ETL workflows.

The aforementioned discussion is supported from a second recent study [17], where the authors note the decline in license revenue for pure ETL tools, mainly due to the crisis of IT spending and the appearance of ETL solutions from traditional database and business intelligence vendors. The Gartner study discusses the role of the three major database vendors (IBM, Microsoft, Oracle) and points that they slowly start to take a portion of the ETL market through their DBMS-built-in solutions.

In the sequel, we elaborate more on the major vendors in the area of the commercial ETL tools, and we discuss three tools that the major database vendors provide, as such two ETL tools that are considered as best sellers. But, we stress the fact that the former three have the benefit of the minimum cost, because they are shipped with the database, while the latter two have the benefit to aim at complex and deep solutions not envisioned by the generic products.

*IBM.* DB2 Universal Database offers the Data Warehouse Center [1], a component that automates data warehouse processing, and the DB2 Warehouse Manager that extends the capabilities of the Data Warehouse Center with additional agents, transforms and metadata capabilities. Data Warehouse Center is used to define the processes that move and transform data for the warehouse. Warehouse Manager is used to

schedule, maintain, and monitor these processes. Within the Data Warehouse Center, the *warehouse schema modeler* is a specialized tool for generating and storing schema associated with a data warehouse. Any schema resulting from this process can be passed as metadata to an OLAP tool. The *process modeler* allows user to graphically link the steps needed to build and maintain data warehouses and dependent data marts. DB2 Warehouse Manager includes enhanced ETL function over and above the base capabilities of DB2 Data Warehouse Center. Additionally, it provides metadata management, repository function, as such an integration point for third-party independent software vendors through the *information catalog*.

*Microsoft.* The tool that is offered by Microsoft to implement its proposal for the Open Information Model is presented under the name of *Data Transformation Services (DTS)* [3,18]. DTS are the data-manipulation utility services in SQL Server (from version 7.0) that provide import, export, and data-manipulating services between OLE DB [19], ODBC, and ASCII data stores. DTS are characterized by a basic object, called a package, that stores information on the aforementioned tasks and the order in which they need to be launched. A package can include one or more connections to different data sources, and different tasks and transformations that are executed as steps that define a workflow process [20]. The software modules that support DTS are shipped with MS SQL Server. These modules include:

- *DTS designer:* A GUI used to interactively design and execute DTS packages,
- *DTS export and import wizards:* Wizards that ease the process of defining DTS packages for the import, export and transformation of data,
- *DTS programming interfaces:* A set of OLE Automation and a set of COM interfaces to create customized transformation applications for any system supporting OLE automation or COM.

*Oracle.* Oracle Warehouse Builder [4,21] is a repository-based tool for ETL and data warehousing. The basic architecture comprises two components, the *design environment* and the

*runtime environment*. Each of these components handles a different aspect of the system; the design environment handles metadata, the runtime environment handles physical data. The metadata component revolves around the metadata repository and the design tool. The repository is based on the Common Warehouse Model (CWM) standard and consists of a set of tables in an Oracle database that are accessed via a Java-based access layer. The front-end of the tool (entirely written in Java) features wizards and graphical editors for logging onto the repository. The data component revolves around the runtime environment and the warehouse database. The Warehouse Builder runtime is a set of tables, sequences, packages, and triggers that are installed in the target schema. The code generator that bases on the definitions stores in the repository, it creates the code necessary to implement the warehouse. Warehouse Builder generates extraction specific languages (SQL\*Loader control files for flat files, ABAP for SAP/R3 extraction and PL/SQL for all other systems) for the ETL processes and SQL DDL statements for the database objects. The generated code is deployed, either to the file system or into the database.

*Ascential software.* DataStage XE suite from Ascential Software [15,16] (formerly Informix Business Solutions) is an integrated data warehouse development toolset that includes an ETL tool (DataStage), a data quality tool (Quality Manager), and a metadata management tool (MetaStage). The DataStage ETL component consists of four design and administration modules: *Manager*, *Designer*, *Director* and *Administrator*, as such a *metadata repository* and a *server*. The DataStage Manager is the basic metadata management tool. In the Designer module of DataStage, ETL tasks execute within individual “stage” objects (source, target and transformation stages), in order to create ETL tasks. The Director is DataStage’s job validation and scheduling module. The DataStage Administrator is primarily for controlling security functions. The DataStage Server is the engine that moves data from source to target.

*Informatica.* Informatica PowerCenter [2] is the industry-leading (according to recent studies

[14,17]) data integration platform for building, deploying, and managing enterprise data warehouses, and other data integration projects. The workhorse of Informatica PowerCenter is a data integration engine that executes all data extraction, transformation, migration and loading functions in-memory, without generating code or requiring developers to hand-code these procedures. The PowerCenter data integration engine is metadata driven, creating a repository-and-engine partnership that ensures data integration processes are optimally executed.

## 5.2. Research efforts

*Research focused specifically on ETL.* The *AJAX* system [22] is a data cleaning tool developed at INRIA France. It deals with typical data quality problems, such as *the object identity problem* [23], *errors due to mistyping* and *data inconsistencies* between matching records. This tool can be used either for a single source or for integrating multiple data sources. *AJAX* provides a framework wherein the logic of a data cleaning program is modeled as a directed graph of data transformations that start from some input source data. Four types of data transformations are supported:

*Mapping transformations* standardize data formats (e.g. date format) or simply merge or split columns in order to produce more suitable formats.

*Matching transformations* find pairs of records that most probably refer to same object. These pairs are called *matching pairs* and each such pair is assigned a similarity value.

*Clustering transformations* group together matching pairs with a high similarity value by applying a given grouping criteria (e.g. by transitive closure).

*Merging transformations* are applied to each individual cluster in order to eliminate duplicates or produce new records for the resulting integrated data source.

*AJAX* also provides a declarative language for specifying data cleaning programs, which consists of SQL statements enriched with a set of specific

primitives to express mapping, matching, clustering and merging transformations. Finally, a interactive environment is supplied to the user in order to resolve errors and inconsistencies that cannot be automatically handled and support a stepwise refinement design of data cleaning programs. The theoretic foundations of this tool can be found in [24], where apart from the presentation of a general framework for the data cleaning process, specific optimization techniques tailored for data cleaning applications are discussed.

Raman et al. [25,26] present the *Potter's Wheel* system, which is targeted to provide interactive data cleaning to its users. The system offers the possibility of performing several algebraic operations over an underlying data set, including *format* (application of a function), *drop*, *copy*, *add* a column, *merge* delimited columns, *split* a column on the basis of a regular expression or a position in a string, *divide* a column on the basis of a predicate (resulting in two columns, the first involving the rows satisfying the condition of the predicate and the second involving the rest), *selection* of rows on the basis of a condition, *folding* columns (where a set of attributes of a record is split into several rows) and *unfolding*. Optimization algorithms are also provided for the CPU usage for certain classes of operators. The general idea behind *Potter's Wheel* is that users build data transformations in iterative and interactive way. In the background, *Potter's Wheel* automatically infers structures for data values in terms of user-defined domains, and accordingly checks for constraint violations. Users gradually build transformations to clean the data by adding or undoing transforms on a spreadsheet-like interface; the effect of a transform is shown at once on records visible on screen. These transforms are specified either through simple graphical operations, or by showing the desired effects on example data values. In the background, *Potter's Wheel* automatically infers structures for data values in terms of user-defined domains, and accordingly checks for constraint violations. Thus users can gradually build a transformation as discrepancies are found, and clean the data without writing complex programs or enduring long delays.

We believe that the AJAX tool is mostly oriented towards the integration of web data (which is also supported by the ontology of its algebraic transformations); at the same time, Potter's wheel is mostly oriented towards an interactive data cleaning tool, where the users interactively choose data. With respect to these approaches, we believe that our technique contributes (a) by offering an extensible framework through a uniform extensibility mechanism, and (b) by providing formal foundations to allow the reasoning over the constructed ETL scenarios. Clearly, ARKTOS II is a design tool for traditional data warehouse flows; therefore, we find the aforementioned approaches complementary (especially Potter's Wheel). At the same time, when contrasted with the industrial tools, it is evident that although ARKTOS II is only a design environment for the moment, the industrial tools lack the logical abstraction that our model, implemented in ARKTOS II, offers; on the contrary, industrial tools are concerned directly with the physical perspective (at least to the best of our knowledge).

*Data quality and cleaning.* An extensive review of data quality problems and related literature, along with quality management methodologies can be found in [27]. A collection of articles on data transformations [28] offers a discussion on various aspects of this research area. A collection of articles on data cleaning [29] (including a survey [30]), provides an extensive overview of the field, along with research issues and a review of some commercial tools and solutions on specific problems, e.g. [31,32]. In a related, still different, context, we would like to mention the IBIS tool [33]. IBIS is an integration tool following the global-as-view approach to answer queries in a mediated system. Departing from the traditional data integration literature though, IBIS brings the issue of data quality in the integration process. The system takes advantage of the definition of constraints at the intentional level (e.g., foreign key constraints) and tries to provide answers that resolve semantic conflicts (e.g., the violation of a foreign key constraint). The interesting aspect here is that consistency is traded for completeness. For example, whenever an offending row is detected over a foreign key constraint, instead of assuming

the violation of consistency, the system assumes the absence of the appropriate lookup value and adjusts its answers to queries accordingly [34].

*Workflows.* To the best of our knowledge, research on workflows is focused around the following reoccurring themes: (a) modeling [5,9,35,36,37], where the authors are primarily concerned in providing a metamodel for workflows; (b) correctness issues [35–37], where criteria are established to determine whether a workflow is well formed, and (c) workflow transformations [35–37] where the authors are concerned on correctness issues in the evolution of the workflow from a certain plan to another.

In the literature there is a standard proposed by the workflow management coalition (WfMC) [9]. The standard includes a metamodel for the description of a workflow process specification and a textual grammar for the interchange of process definitions. A *workflow process* comprises of a network of *activities*, their interrelationships, criteria for starting/ending a process and other information about *participants*, invoked *applications* and relevant *data*. Also, several other kinds of entities which are external to the workflow, such as system and environmental data or the organizational model, are roughly described. In [38] several interesting research results on workflow management are presented in the field of electronic commerce, distributed execution and adaptive workflows. Still, there is no reference to data flow modeling efforts. In [5] the authors provide an overview of the most frequent control flow patterns in workflows. The patterns refer explicitly to control flow structures like activity sequence, AND/XOR/OR split/join and so on. Several commercial tools are evaluated against the 26 patterns presented. In [35–37] the authors, based on minimal metamodels, try to provide correctness criteria in order to derive equivalent plans for the same workflow scenario.

In more than one work [5,36] the authors mention the necessity for the perspectives already discussed in the introduction of the paper. Data flow or data dependencies are listed within the components of the definition of a workflow; still in all these works the authors quickly move on to assume that control flow is the primary aspect of

workflow modeling and do not deal with data-centric issues any further. It is particularly interesting that the [9] standard is not concerned with the role of business data at all. The primary focus of the WfMC standard is the interfaces that connect the different parts of a workflow engine and the transitions between the states of a workflow. No reference is made to business data (although the standard refers to data which are relevant for the transition from one state to another, under the name *workflow related data*).

### 5.3. Applications of ETL workflows in data warehouses.

Finally, we would like to mention that the literature reports several efforts (both research and industrial) for the management of processes and workflows that operate on data warehouse systems. In [39], the authors describe an industrial effort where the cleaning mechanisms of the data warehouse are employed in order to avoid the population of the sources with problematic data in the first place. The described solution is based on a workflow that employs techniques from the field of view maintenance. The industrial effort at Deutsche Bank, involving the import/export, transformation and cleaning and storage of data in a Terabyte-size data warehouse is described in Ref. [40]. The paper explains also the usage of metadata management techniques, which involves a broad spectrum of applications, from the import of data to the management of dimensional data and more importantly for the querying of the data warehouse. A research effort (and its application in an industrial application) for the integration and central management of the processes that lie around an information system is presented in the work of Jarke et al. [41]. A metadata management repository is employed to store the different activities of a large workflow, along with important data that these processes employ.

Finally, we should refer the interested reader to [6] for a detailed presentation of ARKTOS II model. The model is accompanied by a set of *importance metrics* where we exploit the graph structure to measure the degree to which activities/recordsets/attributes are bound to their data providers or

consumers. In [42] we propose a complementary conceptual model for ETL scenarios and in [43] a methodology for constructing it. Ref. [44] abstractly describes our approach of modeling and managing ETL processes.

## 6. Discussion

In this section we would like to briefly discuss some comments on the overall evaluation of our approach. Our proposal involves the data modeling part of ETL activities, which are modeled as workflows in our setting; nevertheless, it is not clear whether we covered all possible problems around the topic. Therefore, in this section, we will explore three issues as an overall assessment of our proposal. First, we will discuss its completeness, i.e., whether there are parts of the data modeling that we have missed. Second, we will discuss the possibility of further generalizing our approach to the general case of workflows. Finally, we will exit the domain of the logical design and deal with performance and stability concerns around ETL workflows.

*Completeness.* A first concern that arises, involves the completeness of our approach. We believe that the different layers of Fig. 1 fully cover the different aspects of workflow modeling. We would like to make clear that we focus on the data-oriented part of the modeling, since ETL activities are mostly concerned with a well-established *automated* flow of cleanings and transformations, rather than an interactive session where *user decisions and actions* direct the flow (like for example in [45]).

Still, is this enough to capture all the aspects of the data-centric part of ETL activities? Clearly, we do not provide any “formal” proof for the completeness of our approach. Nevertheless, we can justify our basic assumptions based on the related literature in the field of software metrics, and in particular, on the method of *function points* [46,47]. Function points is a methodology trying to quantify the functionality (and thus the required development effort) of an application. Although based on assumptions that pertain to the technological environment of the late 1970s,

the methodology is still one of the most cited in the field of software measurement. In any case, function points compute the measurement values based on the five following characteristics: (i) user inputs, (ii) user outputs, (iii) user inquiries, (iv) employed files and (v) external interfaces.

We believe that an activity in our setting covers all the above quite successfully, since (a) it employs input and output schemata to obtain and forward data (characteristics i, ii and iii), (b) communicates with files (characteristic iv) and other activities (practically characteristic v). Moreover, it is tuned by some user-provided parameters, which are not explicitly captured by the overall methodology but are quite related to characteristics (iii) and (v). As a more general view on the topic we could claim that it is sufficient to characterize activities with input and output schemata, in order to denote their linkage to data (and other activities, too), while treating parameters as part of the input and/or output of the activity, depending on their nature. We follow a more elaborate approach, treating parameters separately, mainly because they are instrumental in defining our template activities.

*Generality of the results.* A second issue that we would like to bring up is the general applicability of our approach. Is it possible that we apply this modeling for the general case of workflows, instead of applying it simply to the ETL ones? As already mentioned, to the best of our knowledge, typical research efforts in the context of workflow management are concerned with the management of the control flow in a workflow environment. This is clearly due to the complexity of the problem and its practical application to semi-automated, decision-based, interactive workflows where user choices play a crucial role. Therefore, our proposal for a structured management of the data flow, concerning both the interfaces and the internals of activities appears to be complementary to existing approaches for the case of workflows that need to access structured data in some kind of data store, or to exchange structured data between activities.

It is possible, however, that due to the complexity of the workflow, a more general approach should be followed, where activities have multiple

inputs and outputs, covering all the cases of different interactions due to the control flow. We anticipate that a general model for business workflows will employ activities with inputs and outputs, internal processing, and communication with files and other activities (along with all the necessary information on control flow, resource management, etc); nevertheless, we find this to be outside the context of this paper.

*Execution characteristics.* A third concern involves performance. Is it possible to model ETL activities with workflow technology? Typically, the back-stage of the data warehouse operates under strict performance requirements, where a loading time-window dictates how much time is assigned to the overall ETL process to refresh the contents of the data warehouse. Therefore, performance is really a major concern in such an environment. Clearly, in our setting we do not have in mind EAI or other message-oriented technologies to bring the loading task to a successful end. On the contrary, we strongly believe that the volume of data is the major factor of the overall process (and not, for example, any user-oriented decisions). Nevertheless, to our point of view, the paradigm of activities that feed one another with data during the overall process is more than suitable.

Let us mention a recent experience report on the topic: in [48], the authors report on their data warehouse population system. The architecture of the system is discussed in the paper, with particular interest (a) in a “shared data area”, which is an in-memory area for data transformations, with a specialized area for rapid access to lookup tables and (b) the pipelining of the ETL processes. A case study for mobile network traffic data is also discussed, involving around 30 data flows, 10 sources, and around 2TB of data, with 3 billion rows for the major fact table. In order to achieve a throughput of 80 M row/h and 100 M row/day, the designers of the system were practically obliged to exploit low-level OCI calls, in order to avoid storing loading data to files and then loading them through loading tools. With 4 h of loading window for all this workload, the main issues identified involve (a) performance, (b) recovery, (c) day-by-day maintenance of ETL activities, and (d) adaptable and flexible activities.

Based on the above, we believe that the quest for a workflow, rather than a push-and-store paradigm, is quite often the only way to follow.

Of course, this kind of workflow approach possibly suffers in the issue of software stability, and mostly recovery. Having a big amount of transient data, processed through a large set of activities in main memory is clearly vulnerable to both software and hardware failures. Moreover, once a failure has occurred, rapid recovery, if possible within the loading time-window is also a strong desideratum. Techniques to handle the issue of recovery already exist. To our knowledge the most prominent one is the one by Labio et al. [49], where the ordering of data is taken into consideration. Checkpoint techniques guarantee that once the activity output is ordered, recovery can start right at the point where the activity did the last checkpoint, thus speeding up the whole process significantly.

## 7. Conclusions

In this paper, we have focused on the data-centric part of logical design of the ETL scenario of a data warehouse. First, we have defined a formal logical metamodel as a logical abstraction of ETL processes. The data stores, activities and their constituent parts, as well as the provider relationships that map data producers to data consumers have formally been defined. We have also employed a declarative database programming language, LDL, to define the semantics of each activity. Then, we have provided a reusability framework that complements the genericity of the aforementioned metamodel. Practically, this is achieved from an extensible set of specializations of the entities of the metamodel layer, specifically tailored for the most frequent elements of ETL scenarios, which we call template activities. In the context of template materialization, we have dealt with specific language issues, in terms of the mechanics of template instantiation to concrete activities. Finally, we have presented a graphical design tool, ARKTOS II, with the goal of facilitating the design of ETL scenarios, based on our model.

Still, several research issues are still left open, on the grounds of this work. A broad area of research involves the efficient and reliable execution of an ETL scenario. In this context, an obvious issue is the optimization of ETL scenarios under time and throughput constraints. The topic appears interesting, since the frequent usage of functions in ETL scenarios, drives the problem outside the expressive power of relational algebra (and therefore the traditional optimization techniques, used in the context of relational query optimizers). The problem becomes even more complex if one considers issues of reliability and recovery in the presence of failures or even issues of software quality (e.g., resilience to changes in the underlying data stores). Similar results already exist in the context of materialized views maintenance [50,51]. Of course, the issue of providing optimal algorithms for individual ETL tasks (e.g., duplicate detection, surrogate key assignment, or identification of differentials) is also very interesting. In a different line of research one could also work towards providing a general model for the data flow of data-centric business workflows, involving issues of transactions, alternative interfaces in the context of control flow decisions and contingency scenarios. Finally, the extension of ETL techniques for streaming or XML-formatted data is also another interesting topic of future research.

## Acknowledgments

We would like to thank the anonymous reviewers of this paper for valuable comments that improved the overall quality of the paper.

## References

- [1] IBM, IBM Data warehouse manager, available at <http://www-3.ibm.com/software/data/db2/datawarehouse/>.
- [2] Informatica, Power Center, available at <http://www.informatica.com/products/data+integration/powercenter/default.htm>.
- [3] Microsoft, Data transformation services, available at <http://www.microsoft.com>.
- [4] Oracle, Oracle warehouse builder product page, available at <http://otn.oracle.com/products/warehouse/content.html>.



- [5] W.M.P. van der Aalst, A.H.M. ter Hofstede, B. Kiepuszewski, A.P. Barros. Workflow Patterns, BETA Working Paper Series, WP 47, Eindhoven University of Technology, Eindhoven, 2000, available at the Workflow Patterns web site, at tmit <http://www.tm.tue.nl/research/patterns/documentation.htm>.
- [6] P. Vassiliadis, A. Simitsis, S. Skiadopoulos, Modeling ETL activities as graphs, in: Proceedings of the Fourth International Workshop on Design and Management of Data Warehouses (DMDW), pp. 52–61, Toronto, Canada, 2002.
- [7] P. Vassiliadis, A. Simitsis, P. Georgantas, M. Terrovitis, A framework for the design of ETL scenarios, in: Proceedings of the 15th Conference on Advanced Information Systems Engineering (CAiSE '03), pp. 520–535, Klagenfurt/Velden, Austria, 16–20 June, 2003.
- [8] R. Kimbal, L. Reeves, M. Ross, W. Thornthwaite, The Data Warehouse Lifecycle Toolkit: Expert Methods for Designing, Developing, and Deploying Data Warehouses, Wiley, New York, 1998.
- [9] Workflow Management Coalition, Interface 1: Process Definition Interchange Process Model, Document no. WfMC TC-1016-P, 1998, available at <http://www.wfmc.org>.
- [10] S. Naqvi, S. Tsur, A Logical Language for Data and Knowledge Bases, Computer Science Press, Rockville, MD, 1989.
- [11] C. Zaniolo, LDL++ Tutorial, UCLA, <http://pike.cs.ucla.edu/ldl/>, December 1998.
- [12] D. Dori, Conceptual modeling and system architecting, Commun. ACM 46 (10) (2003) 62–65.
- [13] P. Vassiliadis, A. Simitsis, P. Georgantas, M. Terrovitis, S. Skiadopoulos, A generic and customizable framework for the design of ETL scenarios (long version), Technical Report TR-2004-1, Knowledge and Database Systems Laboratory, National Technical University of Athens, available at <http://www.dbnet.ece.ntua.gr/pubs/>.
- [14] Giga Information Group. Market Overview Update: ETL. Technical Report RPA-032002-00021, March 2002.
- [15] Ascential Software Inc., available at: <http://www.ascentialsoftware.com>.
- [16] Ascential Software Products—Data Warehousing Technology, available at: <http://www.ascentialsoftware.com/products/datastage.html>.
- [17] Gartner Inc., ETL magic quadrant update: market pressure increases, Gartner's Strategic Data Management Research Note, M-19-1108, January 2003.
- [18] P.A. Bernstein, T. Bergstraesser, Meta-data support for data transformations using Microsoft repository. Special issue on data transformations, Bull. Tech. Committee Data Eng. 22 (1) (1999) 9–14.
- [19] Microsoft Corp, OLEDB specification, available at <http://www.microsoft.com/data/oledb>.
- [20] C. Graves, M. Scott, M. Benkovich, P. Turley, R. Skoglund, R. Dewson, S. Youness, D. Lee, S. Ferguson, T. Bain, T. Joubert, Professional SQL Server 2000 data warehousing with analysis services, 1st ed, Wrox Press Ltd., 2001.
- [21] Oracle, Oracle 9i Warehouse Builder, Architectural White paper, April 2002.
- [22] H. Galhardas, D. Florescu, D. Shasha, E. Simon, Ajax: An extensible data cleaning tool, in: Proceedings of the ACM SIGMOD International Conference on the Management of Data, pp. 590, Dallas, TX, 2000.
- [23] W. Cohen. Some practical observations on integration of Web information, in: WebDB'99 Workshop in conj. with ACM SIGMOD, 1999.
- [24] H. Galhardas, D. Florescu, D. Shasha, E. Simon, An extensible framework for data cleaning, Technical Report INRIA 1999 (RR-3742).
- [25] V. Raman, J. Hellerstein, Potters Wheel: an interactive framework for data cleaning and transformation, Technical Report University of California at Berkeley, Computer Science Division, 2000, available at <http://www.cs.berkeley.edu/rshankar/papers/pwheel.pdf>.
- [26] V. Raman, J. Hellerstein, Potter's Wheel: an interactive data cleaning system, in: Proceedings of 27th International Conference on Very Large Data Bases (VLDB), pp. 381–390, Roma, Italy, 2001.
- [27] M. Jarke, M. Lenzerini, Y. Vassiliou, P. Vassiliadis, Springer, New York, 2000.
- [28] E. Rundensteiner, Special issue on data transformations, Bull. Tech. Committee Data Eng. 22 (1) (1999).
- [29] S. Sarawagi, Special issue on data cleaning, Bull. Tech. Committee Data Eng. 23 (4) (2000).
- [30] E. Rahm, H. Hai Do, Data cleaning: problems and current approaches, Bull. Tech. Committee Data Eng. 23 (4) (2000).
- [31] V. Borkar, K. Deshmuk, S. Sarawagi, Automatically extracting structure form free text Addresses, Bull. Tech. Committee Data Eng. 23 (4) (2000).
- [32] A. Monge, Matching algorithms within a duplicate detection system, Bull. Tech. Committee Data Eng. 23 (4) (2000).
- [33] A. Calí, D. Calvanese, G. De Giacomo, M. Lenzerini, P. Naggari, F. Vernacotola, IBIS: Semantic data integration at work, in: Proceedings of the 15th International Conference on Advanced Information Systems Engineering (CAiSE 2003), vol. 2681 of Lecture Notes in Computer Science, pp. 79–94. Springer, 2003.
- [34] A. Calí, D. Calvanese, G. De Giacomo, M. Lenzerini, Data integration under integrity constraints, in: Proceedings of the 14th International Conference on Advanced Information Systems Engineering (CAiSE 2002), vol. 2348 of Lecture Notes in Computer Science, pp. 262–279. Springer, 2002.
- [35] J. Eder, W. Gruber, A meta model for structured workflows supporting workflow transformations, in Proceedings of the Sixth East European Conference on Advances in Databases and Information Systems (ADBIS 2002), pp. 326–339, Bratislava, Slovakia, September 8–11, 2002.

- [36] W. Sadiq, M.E. Orłowska, On business process model transformations, 19th International Conference on Conceptual Modeling (ER 2000), Salt Lake City, UT, USA, October 9–12, 2000, pp. 267–280.
- [37] B. Kiepuszewski, A.H.M. ter Hofstede, C. Bussler, On structured workflow modeling, in: Proceedings of the 12th International Conference on Advanced Information Systems Engineering (CAiSE 2000), pp: 431–445, Stockholm, Sweden, June 5–9, 2000.
- [38] P. Dadam, M. Reichert, (eds.), Enterprise-wide and cross-enterprise workflow management: concepts, systems, applications, GI Workshop Informatik'99, 1999, available at <http://www.informatik.uni-ulm.de/dbis/veranstaltungen/Workshop-Informatik99-Proceedings.pdf>.
- [39] M. Jarke, C. Quix, G. Blees, D. Lehmann, G. Michalk, S. Stierl, Improving OLTP Data Quality Using Data Warehouse Mechanisms, Proceedings of 1999 ACM SIGMOD International Conference on Management of Data, Philadelphia, USA, June 1999, pp. 537–538.
- [40] E. Schafer, J.-D. Becker, M. Jarke, DB-Prism: Integrated data warehouses and knowledge networks for bank controlling, Proceedings of the 26th International Conference on Very Large Databases, Cairo, Egypt, 2000.
- [41] M. Jarke, T. List, J. Koller, The challenge of process warehousing. Proceedings of the 26th International Conference on Very Large Databases, Cairo, Egypt, 2000.
- [42] P. Vassiliadis, A. Simitsis, S. Skiadopoulos, Conceptual modeling for ETL processes, in: Proceedings of the Fifth ACM International Workshop on Data Warehousing and OLAP (DOLAP), pp. 14–21, McLean, VA, USA, 2002.
- [43] A. Simitsis, P. Vassiliadis, A methodology for the conceptual modeling of ETL processes, in: Proceedings of the Decision Systems Engineering (DSE '03), Velden, Austria, June 17, 2003.
- [44] A. Simitsis, Modeling and managing ETL processes, in: Proceedings of the VLDB 2003 Ph.D Workshop, Berlin, Germany, September 12–13, 2003.
- [45] F. Casati, S. Ceri, B. Pernici, G. Pozzi. Conceptual Modeling of Workflows, in: Proceedings of the OO-ER Conference, Australia, 1995.
- [46] A.J. Albrecht, Measuring Application Development Productivity, in IBM Applications Development Symposium, Monterey, CA, 1979, pp. 83–92.
- [47] R.S. Pressman, Software Engineering: A Practitioner's Approach, 5th ed, McGraw-Hill, New York, 2000.
- [48] J. Adzic, V. Fiore, Data Warehouse Population Platform, in: Proceedings of the Fifth International Workshop on the Design and Management of Data Warehouses (DMDW'03), Berlin, Germany, September 2003.
- [49] W. Labio, J.L. Wiener, H. Garcia-Molina, V. Gorelik, Efficient resumption of interrupted warehouse loads, in: Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data (SIGMOD 2000), pp. 46–57, Dallas, TX, USA, 2000.
- [50] J. Chen, S. Chen, E.A. Rundensteiner, A Transactional Model for Data Warehouse Maintenance, in: Proceedings of the of ER 2002, LNCS 2503, pp. 247–262, 2002.
- [51] B. Liu, S. Chen, E.A. Rundensteiner, A transactional approach to parallel data warehouse maintenance, in: Proceedings of DaWaK 2002, LNCS 2454, 2002. pp. 307–316.