



PERGAMON

Information Systems 26 (2001) 205–236



www.elsevier.com/locate/infosys

Data warehouse process management

Panos Vassiliadis^{a,*}, Christoph Quix^b, Yannis Vassiliou^a, Matthias Jarke^{b,c}

^a *Department of Electrical and Computer Engineering, Computer Science Division, National Technical University of Athens, Iroon Polytechniou 9, 157 73, Athens, Greece*

^b *Informatik V (Information Systems), RWTH Aachen, 52056 Aachen, Germany*

^c *GMD-FIT, 53754 Sankt Augustin, Germany*

Abstract

Previous research has provided metadata models that enable the capturing of the static components of a data warehouse architecture, along with information on different quality factors over these components. This paper complements this work with the modeling of the dynamic parts of the data warehouse. The proposed metamodel of data warehouse operational processes is capable of modeling complex activities, their interrelationships, and the relationship of activities with data sources and execution details. Moreover, the metamodel complements the existing architecture and quality models in a coherent fashion, resulting in a full framework for quality-oriented data warehouse management, capable of supporting the design, administration and especially evolution of a data warehouse. Finally, we exploit our framework to revert the widespread belief that data warehouses can be treated as collections of materialized views. We have implemented this metamodel using the language Telos and the metadata repository system ConceptBase. © 2001 Elsevier Science Ltd. All rights reserved.

Keywords: Data warehousing; Process modeling; Evolution; Quality

1. Introduction

Data warehouses (DW) integrate data from multiple heterogeneous information sources and transform them into a multidimensional representation for decision support applications. Apart from a complex architecture, involving data sources, the data staging area, operational data stores, the global data warehouse, the client data marts, etc., a data warehouse is also characterized

by a complex lifecycle. In a permanent design phase, the designer has to produce and maintain a conceptual model and a usually voluminous logical schema, accompanied by a detailed physical design for efficiency reasons. The designer must also deal with data warehouse administrative processes, which are complex in structure, large in number and hard to code; deadlines must be met for the population of the data warehouse and contingency actions taken in the case of errors. Finally, the evolution phase involves a combination of design and administration tasks: as time passes, the business rules of an organization change, new data are requested by the end users,

*Corresponding author. Tel.: +30-1-772-1402; fax: +30-1-772-1442.

E-mail address: pvassil@dbnet.ece.ntua.gr (P. Vassiliadis).

new sources of information become available, and the data warehouse architecture must evolve to efficiently support the decision-making process within the organization that owns the data warehouse.

All the data warehouse components, processes and data should be tracked and administered via a *metadata repository*. In [1], we presented a metadata modeling approach which enables the capturing of the *static* parts of the architecture of a data warehouse. The linkage of the architecture model to quality parameters (in the form of a quality model) and its implementation in the metadata repository ConceptBase have been formally described in [2]. Vassiliadis et al. [3] present a methodology for the exploitation of the information found in the metadata repository and the quality-oriented evolution of a data warehouse based on the architecture and quality model. In this paper, we complement these results with metamodels and support tools for the *dynamic* part of the data warehouse environment: the *operational data warehouse processes*. The combination of all the data warehouse viewpoints is depicted in Fig. 1.

In the three phases of the data warehouse lifecycle, the interested stakeholders need information on various aspects of the examined processes: *what* are they supposed to do, *how* are they implemented, *why* are they necessary and how do they affect other processes in the data warehouse [4,1]. Like the data warehouse architecture and quality metamodels, the process metamodel

assumes the *clustering of their entities in logical, physical and conceptual perspectives*, each assigned with the task of answering one of the aforementioned stakeholder questions. In the rest of this section we briefly present the requirements faced in each phase, our solutions and their expected benefits.

The *design and implementation* of operational data warehouse process is a labor-intensive and lengthy procedure, covering 30–80% of effort and expenses of the overall data warehouse construction [5,6]. For a metamodel to be able to efficiently support the design and implementation tasks, it is imperative to capture at least two essential aspects of data warehouse processes, *complexity of structure and relationship with the involved data*. In our proposal, the logical perspective is capable of modeling the structure of complex activities and capture all the entities of the widely accepted Workflow Management Coalition Standard [7]. The relationship of data warehouse activities with their underlying data stores is taken care of in terms of SQL definitions.

This simple idea reverts the classical belief that *data warehouses are simply collections of materialized views*. In previous data warehouse research, directly assigning a naïve view definition to a data warehouse table has been the most common practice. Although this abstraction is elegant and sufficient for the purpose of examining alternative strategies for view maintenance, it is incapable of capturing real world processes within a data warehouse environment. In our approach, we can deduce the definition of a table in the data warehouse as the outcome of the combination of the processes that populate it. This new kind of definition complements existing approaches, since our approach provides the operational semantics for the content of a data warehouse table, whereas the existing ones give an abstraction of its intentional semantics.

The conceptual process perspective traces the reasons behind the structure of the data warehouse. We extend the demand-oriented concept of *dependencies* as in the Actor-Dependency model [4], with the supply-oriented notion of *suitability* that fits well with the redundancy found often in

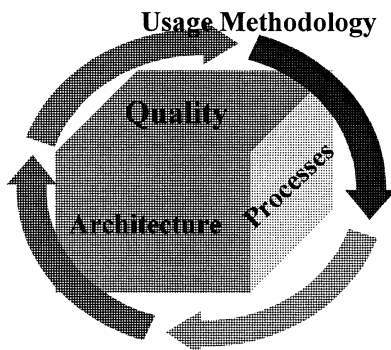


Fig. 1. The different viewpoints for the metadata repository of a data warehouse.

data warehouses. As an another extension to the Actor-Dependency model, we have generalized the notion of role in order to uniformly trace any person, program or data store participating in the system.

By implementing the metamodel in an object logic, we can exploit the query facilities of the repository to provide the support for consistency checking of the design. The deductive capabilities of ConceptBase [8] provide the facilities to avoid assigning manually all the interdependencies of activity roles in the conceptual perspective. It is sufficient to impose rules to deduce these interdependencies from the structure of data stores and activities.

While the design and implementation of the warehouse are performed in a rather controlled environment, the *administration* of the warehouse has to deal with problems that evolve in an ad-hoc fashion. For example, during the loading of the warehouse contingency treatment is necessary for the efficient administration of failures. In such events, the knowledge of the structure of a process is important; moreover, the specific *traces of executed processes* are also required to be tracked down—in an erroneous situation, not only the causes of the failure, but also the progress of the loading process by the time of the failure must be detected, in order to efficiently resume its operation. Still, failures during the warehouse loading are only the tip of the iceberg as far as problems in a data warehouse environment are concerned. This brings up the discussion on *data warehouse quality* and the ability of a metadata repository to trace it in an expressive and usable fashion. To face this problem, the proposed process metamodel is explicitly linked to our earlier quality metamodel [2]. We complement this linkage by mentioning specific quality factors for the quality dimensions of the ISO 9126 standard for software implementation and evaluation.

Identifying erroneous situations or unsatisfactory quality in the data warehouse environment is not sufficient. The data warehouse stakeholders should be supported in their efforts to react against these phenomena. The above-mentioned suitability notion in the conceptual perspective of the process metamodel allows the definition of recovery actions

to potential errors or problems (e.g., alternative paths for the population of the data warehouse) in a straightforward way, during runtime.

Data warehouse evolution is unavoidable as new sources and clients are integrated, business rules change and user requests multiply. The effect of evolving the structure of the warehouse can be predicted by tracing the various interdependencies among the components of the warehouse. We have already mentioned how the conceptual perspective of the metamodel traces interdependencies between all the participants in a data warehouse environment, whether persons, programs or data stores. The prediction of potential impacts (whether of political, structural, or operational nature) is supported by this feature in several ways. To mention the simplest, the sheer existence of dependency links forecasts a potential impact in the architecture of the warehouse in the presence of any changes. More elaborate techniques will also be provided in this paper, by taking into account the particular attributes that participate in these interdependencies and the SQL definitions of the involved processes and data stores. Naturally, the existence of suitability links suggests alternatives for the new structure of the warehouse. We do not claim that our approach is suitable for any kind of process, but focus our attention to the internals of data warehouse systems.

This paper is organized as follows: In Section 2 we present the background work and the motivation for this paper. In Section 3 we describe the process metamodel and in Section 4 we present its linkage to the quality model. In Section 5 we present how the metadata repository can be used for the determination of the operational semantics of the data warehouse tables and for evolution purposes. In Section 6 we present related work and Section 7 presents issues for future research.

2. Background and motivation

In this section we will detail the background work and the motivation behind the proposed metamodel for data warehouse operational processes.

2.1. Background work for the metamodel: the quest for formal models of quality

For decision support, a data warehouse must provide high quality of data and service. Errors in databases have been reported to be up to 10% range and even higher in a variety of applications. Wang et al. [9] report that more than \$2 billion of US federal loan money had been lost because of poor data quality at a single agency; manufacturing companies spend over 25% of their sales on wasteful practices, service companies up to 40%. In certain vertical markets (e.g., the public sector) data quality is not an option but a constraint for the proper operation of the data warehouse. Thus, data quality problems seem to introduce even more complexity and computational burden to the loading of the data warehouse. In the DWQ project (Foundations of Data Warehouse Quality [10]), we have attacked the problem of quality-oriented design and administration in a formal way, without sacrificing optimization and practical exploitation of our research results. In this subsection we summarize our results as far as needed for this paper.

In [1] a basic metamodel for data warehouse architecture and quality has been presented as in

Fig. 2. The framework describes a data warehouse in three perspectives: a conceptual, a logical and a physical perspective. Each perspective is partitioned into the three traditional data warehouse levels: source, data warehouse and client level.

On the metamodel layer, the framework gives a notation for data warehouse architectures by specifying meta-classes for the usual data warehouse objects like data store, relation, view, etc. On the metadata layer, the metamodel is instantiated with the concrete architecture of a data warehouse, involving its schema definition, indexes, table spaces, etc. The lowest layer in Fig. 2 represents the actual processes and data.

The quality metamodel accompanying the architecture metamodel [2] involves an extension of the Goal-Question-Metric approach [11]. The metamodel introduces the basic entities around quality (including Quality Goals, Quality Queries and Quality Factors), the metadata layer is customized per warehouse with quality scenarios, and the instance layer captures concrete measurements of the quality of a given data warehouse.

The static description of the architecture (left part of Fig. 2) is complemented in this paper with a metamodel of the dynamic data warehouse

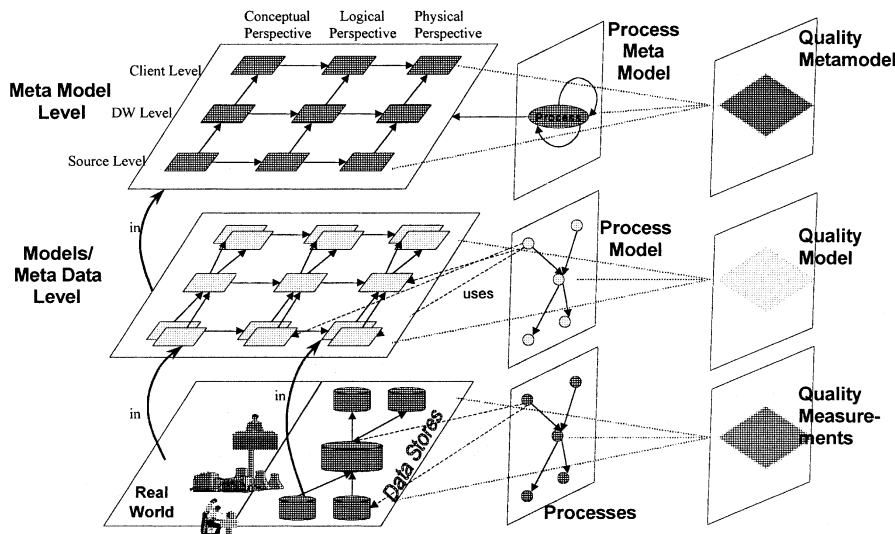


Fig. 2. Framework for Data Warehouse Architecture [1].

operational processes. As one can notice in the middle of Fig. 2. we follow again a three-level instantiation: a *Process Metamodel* deals with generic entities involved in all data warehouse processes (operating on entities found at the data warehouse metamodel level), the *Process Model* covers the processes of a specific data warehouse by employing instances of the metamodel entities, and the *Process Traces* capture the execution of the actual data warehouse processes happening in the real world.

In [3], the Goal–Question–Metric (GQM) methodology has been extended in order (a) to capture the interrelationships between different quality factors with respect to a specific quality goal, and (b) to define an appropriate lifecycle that deals with quality goal evaluation and improvement. The methodology comprises a set of steps, involving the *design* of the quality goal, the *evaluation* of the current status, the *analysis* and *improvement* of this situation, and finally, the *re-evaluation* of the achieved plan. The metadata repository together with this quality goal definition methodology constitutes a decision support system which helps data warehouse designers and administrators to take relevant decisions, in order to achieve a reasonable quality level which fits the best user requirements.

Throughout our models, we encourage the use of *templates* for process, quality and architecture objects. This is especially apparent in the metadata layer where an abstract specification of architecture, process or quality objects, originally coming from the designers of the data warehouse, can be properly specialized by the data warehouse administrator at runtime. For example, a high-level specification of a chain of activities involving extraction, checking for primary and foreign keys and final loading in the warehouse, can be customized to specific programs, later in the construction of the warehouse. Practical experience has shown that this kind of templates can be reoccurring in a data warehouse architecture, i.e., several tables can be populated through very similar programs (especially if the data are coming from the same source). The interested reader is referred to [12,13] for examples of such templates.

2.2. The 3 perspectives for the process model

Our process model (cf. Fig. 3) follows the same three perspectives as the architecture model, since the perspectives of the process model operate on objects of the respective perspective of the architecture model. As mentioned in [4], there are different ways to view a process: *what* steps it consists of (logical perspective), *how* they are to be performed (physical perspective) and *why* these steps exist (conceptual perspective). Thus, we view a data warehouse process from three perspectives: a central *logical* part of the model, which captures the basic structure and data characteristics of a process, its *physical* counterpart which provides specific details over the actual components that execute the process and a *conceptual* perspective which abstractly represents the basic interrelationships between data warehouse stakeholders and processes in a formal way.

Typically, the information about how a process is executed concerns stakeholders who are involved in the everyday use of the process. The information about the structure of the process concerns stakeholders that manage it while the information relevant to the reasons behind this structure concerns process engineers who are involved in the monitoring or evolution of the

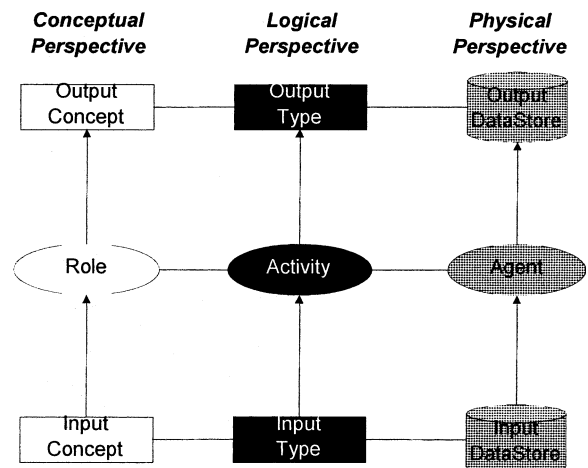


Fig. 3. The reasoning behind the 3 perspectives of the process metamodel.

process environment. Often, all these roles are covered by the data warehouse administration team, although one could also encounter different schemes.

Another important issue shown in Fig. 3 is that we can observe a data flow in each of the three perspectives. In the logical perspective, the modeling is concerned with the functionality of an activity, describing what this particular activity is about in terms of consumption and production of information. In the physical perspective, the details of the execution of the process are the center of the modeling. The most intriguing part, though, is the conceptual perspective covering why a process exists. This can be either due to necessity reasons (in which case, the receiver of information depends on the process to deliver the data) and/or suitability reasons (in which case the information provider is capable of providing the requested information).

2.3. Complexity and traces

Data warehouse operational processes are quite complex, in terms of tasks executed within a single process, execution coherence, contingency treatment, etc. A process metamodel should be able to capture this kind of complexity. In Fig. 4 the data warehouse refreshment process is depicted, as described in [14]. The refreshment process is composed of activities, such as *Data Extraction*, *History Management*, *Data Cleaning*, *Data Integration*, *History Management*, *Update Propagation* and *Customization*. Each of these activities could be executed on a different site. The activities are interlinked through rules, denoted by arrows; in a real-world case study in banking, no less than 17 kinds of knowledge sources determined this process [15]. The gray background in Fig. 4 implies that there is a composition hierarchy in the set of data warehouse operational processes. In fact, the need to isolate only a small subset of the overall processes of the warehouse is common. Any metamodel must be suitable to support zooming in and out of the process structure, in order to achieve this functionality.

The most common reason for this kind of inspection is to avoid or recover erroneous

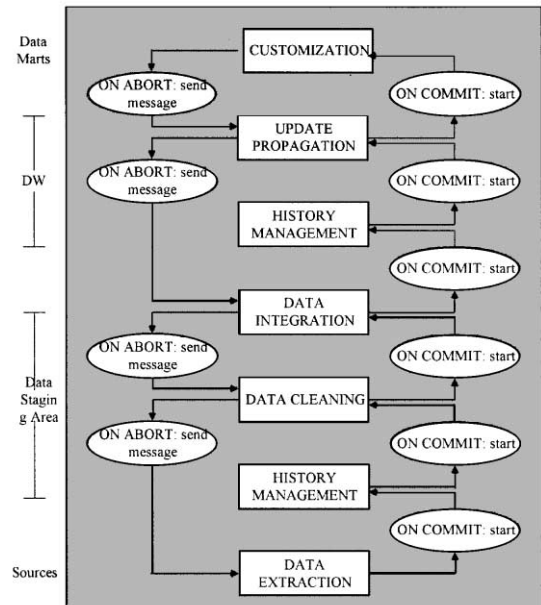


Fig. 4. The Data Warehouse Refreshment process [14].

execution during runtime. The structure of a process is important; moreover the specific traces of executed processes should be tracked down, too. If the repository is able to capture this kind of information, it gains added value since: ex ante the data warehouse stakeholders can use it for design purposes (e.g., to select the data warehouse objects necessary for the performance of a task) and ex post, people can relate the data warehouse objects to decisions, tools and the facts which have happened in the real world [16].

2.4. The data oriented nature of operational data warehouse processes

Data warehouse activities are of data-intensive nature in their attempt to push data from the sources to the tables of the data warehouse or the client data marts. We can justify this claim by listing the most common operational processes:

- *data extraction processes*, which are used for the extraction of information from the legacy systems;

- *data transfer (and loading) processes*, used for the instantiation of higher levels of aggregation in the data warehouse with data coming from the sources or lower levels of aggregation;
- *data transformation processes*, used for the transformation of the propagated data to the desired format;
- *data cleaning processes*, used to ensure the consistency of the data warehouse data (i.e., the fact that these data respect the database constraints and the business rules);
- *computation processes*, which are used for the derivation of new information from the stored data (e.g., further aggregation, querying, business logic, etc.).

To deal with the complexity of the data warehouse loading process, specialized *Extraction-Transformation-Loading (ETL)* tools are available in the market. Their most prominent tasks include:

- the identification of relevant information at the source side,
- the extraction of this information,
- the customization and integration of the information coming from multiple sources into a common format,
- the cleaning of the resulting data set, on the basis of database and business rules, and
- the propagation of the data to the data warehouse and/or data marts.

According to a study for Merrill Lynch [5], ETL and Data Cleaning tools cover a labor-intensive and complex part of the data warehouse processes, estimated to cost at least one third of effort and expenses in the budget of the data warehouse. Demarest [6] mentions that this number can rise up to 80% of the development time in a data warehouse project. Still, due to the complexity and long learning curve of these tools, many organizations turn to in-house development to perform ETL and data cleaning tasks.

2.5. Case study example

To motivate the discussion, we use a part of one of our real-world case studies [17]. The organization collects various data about the annual

activities of all the hospitals of a particular region. The source of data, for our example, is a COBOL file, dealing with the annual information by class of beds and hospital (here we use only three classes, namely A, B and C). It yields a specific attribute for each type of class of beds. Periodically, the COBOL file is transferred from the production system to the data warehouse and stored in a “buffer” table of the data warehouse, acting as mirror of the file.

Then, the tuples of the buffer table are used by computation procedures to further populate a “fact” table inside the data warehouse. Finally, several materialized views are populated with aggregate information and used by client tools for querying.

The entity *Type* denotes the logical schema for all kinds of data stores. Each *Type* is characterized by a name and a set of *Fields*. In our example, we assume the following four *Types*: *CBL*, *Buffer*, *Class.info* and *VI*. The schemata of these types are depicted in Fig. 5. There are four atomic *Activities* in the data warehouse: *Loading*, *Cleaning*, *Computation* and *Aggregation*. The *Loading* activity simply copies the data from the *CBL* Cobol file to the *Buffer* type. *H_ID* is an identifier for the hospital and the three last attributes hold the number of beds per class. The *Cleaning* activity deletes all entries violating the primary key constraint. The *Computation* activity transforms the imported data to a different schema; the date is converted from American to European format and the rest of the attributes are converted to a combination (*Class_id*, *#Beds*). For example, if (03,12/31/1999,30,0,50) is a tuple in the *Buffer* table, the respective tuples in the *Class_Info* table are {(03,31/Dec/1999,A,30), (03,31/Dec/1999,C,50)}. The *Aggregation* activity produces the sum of beds by hospital and year. The combination of all the aforementioned activities is captured by the composite activity *Populate VI*.

3. The metamodel of data warehouse operational processes

We start the presentation of the metamodel for data warehouse operational processes from the

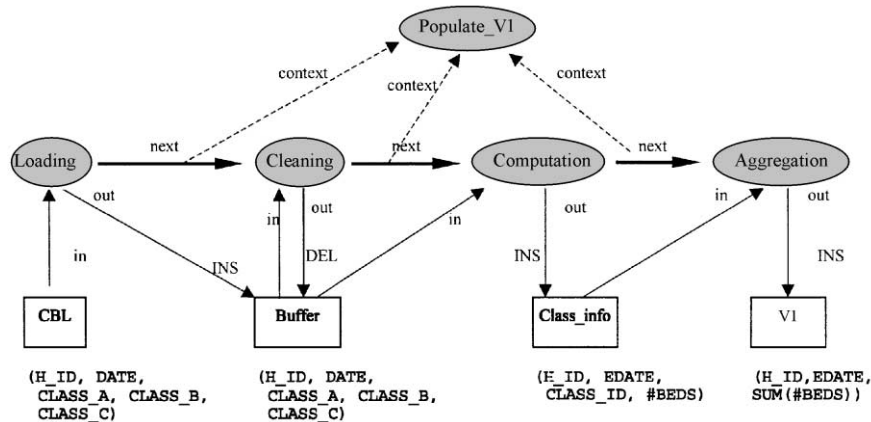


Fig. 5. Motivating example.

logical perspective, to show how the metamodel deals with the requirements of structure complexity and capturing of data semantics in the next two sections. Then, in Sections 3.3 and 3.4 we present the physical and the conceptual perspectives. In the former, the requirement of trace logging will be fulfilled too. The full metamodel is presented in Fig. 6.

For the implementation, we have used a meta-database as a repository for meta-information of the data warehouse components. The architecture, quality and process models are represented in Telos [18], a conceptual modeling language for representing knowledge about information systems. A prototype was implemented in the object-oriented deductive database system ConceptBase [16], that provides query facilities, and a language for constraints and deductive rules. The implementation of the process metamodel in ConceptBase is straightforward. Thus we choose to follow an informal, bird's-eye view of the model, for reasons of presentation.

3.1. Complexity of the process structure

Following the Workflow Coalition [7], the main entity of the logical perspective is *Activity*. An activity represents a unit of “work which is processed by a combination of resource and computer applications”. Activities can be complex, as captured by the specialization of *Activity*, namely *CompositeActivity*. This gives the possi-

bility of Zooming in and out of the repository. The components of composite activities are *Process Elements*. Class *ProcessElement* is a generalization of the entities *Activity* and *TransitionElement*. A transition element is employed for the interconnection of activities participating in a complex activity. The attribute *Next* captures the sequence of events. Formally, a *Process Element* is characterized by the following attributes:

- *Name*: to uniquely identify the *ProcessElement* within the extension of its class.
- *Next*: a *ProcessElement* which is next in the sequence of a composite activity, characterized by:
 - *Context*: Since two activities can be interrelated in more than one complex DW process, the context of this interrelationship is captured by the relevant *CompositeActivity* instance.
 - *Semantics*: denotes whether the next activity in a schedule happens upon successful termination of the previous activity (*COMMIT*) or if a contingency action is required (*ABORT*).

A *TransitionElement* is a specialization of *ProcessElement*, used to support scheduling of the control flow within a composite activity. This extra functionality is supported by two mechanisms. First, we enrich the *Next* link with more meaning, by adding a *Condition* attribute to it. A *Condition* is a logical expression in Telos denoting

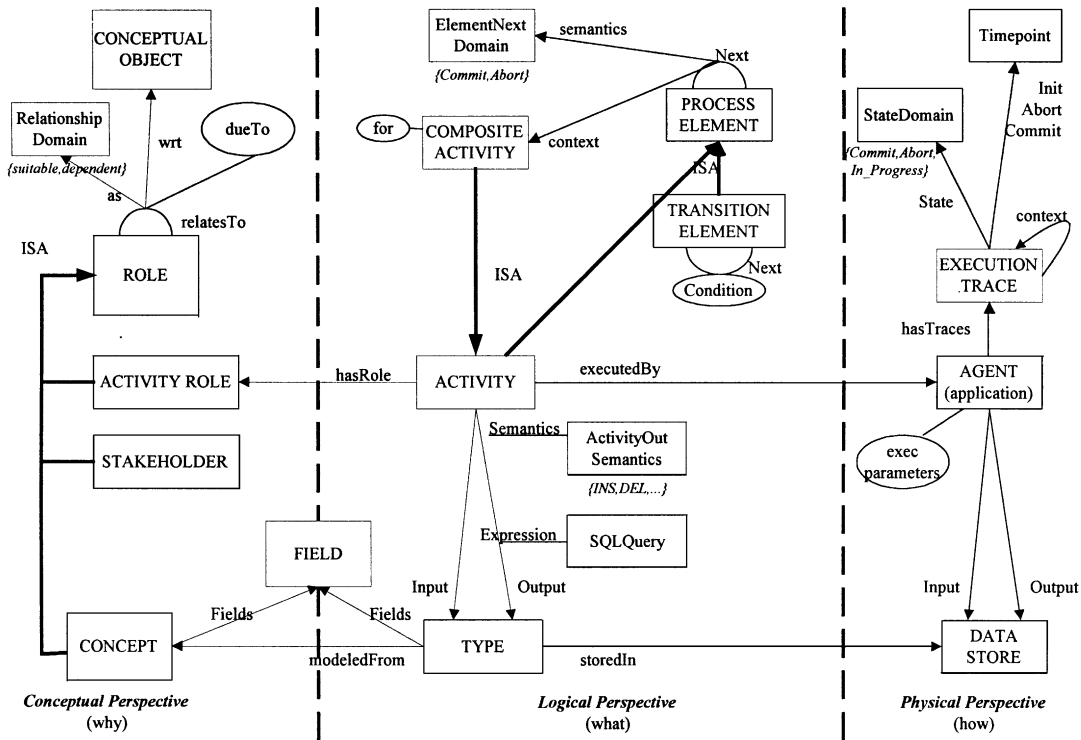


Fig. 6. The metamodel for data warehouse operational processes.

that the firing of the next activity is performed when the *Condition* is met. Second, we specialize the class *TransitionElement* to four subclasses (not shown in Fig. 6), capturing the basic connectives of activities [7]: *Split_AND*, *Split_XOR*, *Join_AND*, *Join_XOR*. The semantics of these entities are the same as the ones of the WfMC proposal. For example, the *Next* activity of a *Join_XOR* instance is fired when (a) the *Join_XOR* has at least two activities “pointing” to it through the *Next* attribute and only one *Next* Activity (well-formedness constraint), (b) the *Condition* of the *Join_XOR* is met and (c) at least one of the “incoming” activities has *COMMIT* semantics in the *Next* attribute. This behavior can be expressed in Telos with appropriate rules. Similarly, *Split_AND* denotes a point in the process chain where more than one concurrent execution threads are initiated, *Split_XOR* denotes a point where exactly one execution thread is to be initiated (out of many different alternatives) and *JOIN_XOR*

acts as a rendezvous point for several concurrent threads, where the execution of the process flow is paused until all incoming activities have completed their execution.

The WfMC proposes two more ways of transition between *Activities*. *Dummy activities* perform routing based on condition checking, they are modeled as simple *Transition Elements*. *LOOP* activities are captured as instances of *CompositeActivity*, with an extra attribute: the *for* condition. Fig. 7 shows the modeling of a composite activity, composed of two sub-activities, where the second is fired when the first activity commits and a Boolean condition is fulfilled. *PI-Committed?* is a transition element.

3.2. Relationship with data

A *Type* denotes the schema for all kinds of data stores. Formally, a *Type* is defined as a specialization of *LogicalObject* with the following attributes:

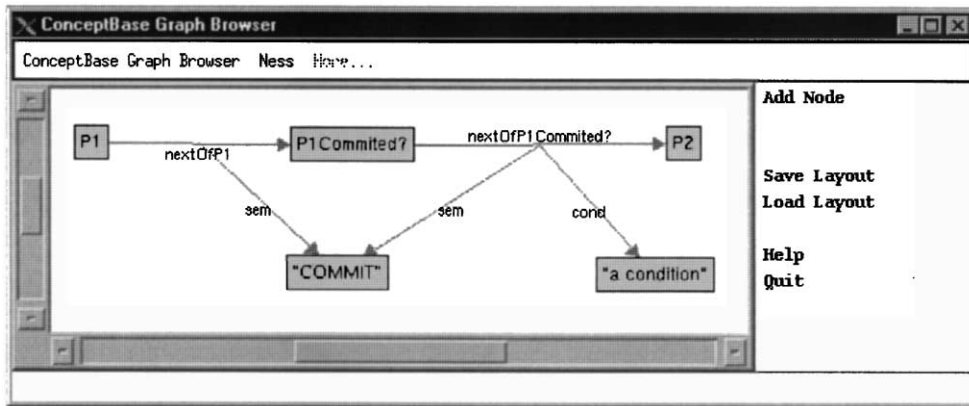


Fig. 7. Example of a composite activity: ON COMMIT(P1), IF <condition> START.

- *Name*: a single name denoting a unique *Type* instance.
- *Fields*: a multi-value attribute. In other words, each *Type* has a name and a set of *Fields*, exactly like a relation in the relational model.
- *Stored*: a *DataStore*, a physical object representing software used to manipulate stored data (e.g., a DBMS). This attribute will be detailed in the description of the physical perspective.
- *ModeledFrom*: a *Concept*, an object in the conceptual perspective, providing a description of the type in a more user-friendly manner. A *Concept* is the generalization of *Entities* and *Relationships* in the Entity-Relationship model (not depicted in Fig. 6).

Any kind of physical data store (multidimensional arrays, COBOL files, even reports) can be represented by a *Type* in the logical perspective. For example, the schema of multidimensional cubes is of the form $[D_1, \dots, D_n, M_1, \dots, M_m]$ where the D_i represent dimensions (forming the primary key of the cube) and the M_j measures [27]. Cobol files, as another example, are records with fields having two peculiarities: nested records and alternative representations. One can easily unfold the nested records and choose one of the alternative representations.

Each *Activity* in a data warehouse environment is linked to a set of incoming and outgoing types. We capture the relationship of activities to data

by expressing the outcome of a data warehouse process as a function over its input data stores. This function is captured through SQL queries, extended with functions. An *Activity* is formally characterized by the following attributes:

- *Name, Next*: inherited from *Process Element*.
- *Input*: multi-valued *Type* attribute modeling all data stores used by the activity to acquire data.
- *Output*: single-valued *Type* attribute. This attribute models the data store or report where the activity outputs data. The *Output* attribute is further explained by two attributes:
 - *Semantics*: a single value belonging to the set $\{Insert, Update, Delete, Select\}$ (captured as the domain of class *ActivityOutSemantics*). A process can either add, delete, or update the data in a data store. Also it can output some messages to the user (captured by using a “Message” *Type* and *Select* semantics).
 - *Expression*: a single SQL query (instance of class *SQLQuery*) to denote the relationship of the output and the input types, possibly with functions.
- *ExecutedBy*: a physical *Agent* (i.e., an application program) executing the *Activity*. More information on agents will be provided in the description of the physical perspective.

- *HasRole*: a conceptual description of the activity. This attribute will be properly explained in the description of the conceptual perspective.

Fig. 8 shows how various types of processes can be modeled in our approach.

To gain more insight in the proposed modeling approach, consider the example of Fig. 5. The expressions and semantics for each activity are listed in Fig. 9. All activities append data to the involved types, so they have *INS* semantics, except for the cleaning process, which deletes data, and thus has *DEL* semantics. We do not imply that everything should actually be implemented using

the employed queries, but rather that the relationship of the input and the output of an activity is expressed as a function, through a declarative language such as SQL.

3.3. The physical perspective

While the logical perspective covers the structure (“what”) of a process, the *physical perspective* covers the details of its execution (“how”). Each process is executed by an *Agent* application program. Each *Type* is assigned to a *DataStore* (providing information for issues like table spaces, indexes, etc.). An *Agent* can be formally described

Activity	Semantics	Expression
Extract & Load	INS/UPD	SELECT * FROM <In>
Transformation	INS/UPD	SELECT * FROM <In> WHERE attr _i = f(attr ₁ , ..., attr _n)
Cleaning	DEL/UPD	- Primary Key: SELECT <P.K.> FROM <IN> GROUP BY <P.K.> HAVING COUNT(*) > 1 - Foreign Key: SELECT <P.K.> FROM <IN> WHERE <F.K.> NOT IN (SELECT <F.K> FROM <TARGET>) - Any other kind of query
Computation	INS/UPD	Any kind of query
Messaging	SEL	Any kind of query

Fig. 8. Examples of *Output* attribute for particular kinds of activities.

Attribute	Expression	Semantics
Loading.Out:	SELECT H_ID, DATE, CLASS_A, CLASS_B, CLASS_C FROM CBL	INS
Cleaning.Out:	SELECT H_ID, DATE, CLASS_A, CLASS_B, CLASS_C FROM BUFFER B1 WHERE EXISTS (SELECT B2.H_ID, B2.DATE FROM BUFFER B2 WHERE B1.H_ID = B2.H_ID AND B1.DATE = B2.DATE GROUP BY H_ID, DATE HAVING COUNT(*) > 1)	DEL
Computation.Out:	SELECT H_ID, EUROPEAN(DATE) AS EDATE, 'A', CLASS_A AS #BEDS FROM BUFFER WHERE CLASS_A <> 0 UNION SELECT H_ID, EUROPEAN(DATE) AS EDATE, 'B', CLASS_B AS #BEDS FROM BUFFER WHERE CLASS_B <> 0 UNION SELECT H_ID, EUROPEAN(DATE) AS EDATE, 'C', CLASS_C AS #BEDS FROM BUFFER WHERE CLASS_C <> 0	INS
Aggregation.Out:	SELECT H_ID, EDATE, SUM(#BEDS) AS SUM_BEDS FROM CLASS_INFO GROUP BY H_ID, EDATE	INS

Fig. 9. Expressions and semantics for the activities of the example.

as follows:

- *Name*: to uniquely identify the *Agent* within the extension of its class.
- *Execution_parameters*: a multi-value string attribute capturing any extra information about the execution of an agent.
- *In, Out*: physical *DataStores* communicating with the *Agent*. The types used by the respective logical activity must be stored within these data stores.
- *HasTraces*: a multi-value attribute capturing all the execution traces of the *Agent*.

An *Execution Trace* is an entity capturing details of activity execution from the proper *Agent*:

- *TraceID*: a unique identification number to uniquely identify each *Execution Trace*.
- *State*: a single value belonging to the domain of class *AgentStateDomain* = {*In_Progress, Commit, Abort*}.
- *Init_time, Abort_time, Commit_time*: timestamps denoting the *Timepoints* when the respective events have occurred.
- *Context*: another *Execution Trace*. This attribute is used in the case where a script (*Agent*) is simply the coordinator script for the execution of several other *Agents*, i.e., in *Composite Activities*. In this case, the *Execution Trace* of the *Agent* of the *Composite Activity* defines the context for the execution of the coordinated agents.

The information of the physical perspective can be used to trace and monitor the execution of data warehouse processes. Fig. 10 sketches the trace information after a successful execution of the process described in Fig. 5. We show the relationship between the logical and the physical perspective by linking each logical activity to a physical application program. Each agent has a set of execution traces. We can see that the trace of composite activity *Populate VI* has as *Init time* the time of the execution of the first sub-activity and *Commit time* the completion time of the last activity. Also, this composite activity defines the context of the execution of its sub-activities: this is captured by properly populating the attribute *Context*.

3.4. The conceptual perspective

A major purpose behind the introduction of the conceptual perspective is to help stakeholders understand the reasoning behind decisions on the architecture and physical characteristics of data warehouse processes. Our modeling approach captures dependency and suitability relationships among the basic conceptual entities to facilitate the design, administration and evolution of the data warehouse.

Each *Type* in the logical perspective is the counterpart of a *Concept* in the conceptual perspective. A concept represents a class of

Activity Name	Populate VI	Loading	Cleaning	Computation	Aggregation
Agent Name	Schedule.exe	sqlldr80.exe	Clean.pql	Comp.pql	Aggr.pql
Execution Parameters	--	Parfile=param.par	--	--	--
HasTraces	{..., 100, ...}	{..., 101, ...}	{..., 102, ...}	{..., 103, ...}	{..., 104, ...}
TraceID	100	101	102	103	104
State	COMMIT	COMMIT	COMMIT	COMMIT	COMMIT
Init time	31/12/99:00:00	31/12/99:00:00	31/12/99:01:00	31/12/99:01:15	31/12/99:03:00
Abort time	--	--	--	--	--
Commit time	31/12/99:03:30	31/12/99:01:00	31/12/99:01:15	31/12/99:03:00	31/12/99:03:30
Context	--	100	100	100	100

Fig. 10. Trace information after a successful execution of the process of the example.

real-world objects, in terms of a conceptual metamodel, e.g., the Entity-Relationship or UML notation. Both *Types* and *Concepts* are constructed from *Fields* (representing their attributes), through the attribute *fields*. We consider *Field* to be a subtype both of *LogicalObject* and *ConceptualObject*. The central conceptual entity is the *Role* which generalizes the conceptual counterparts of activities, stakeholders and data stores. The class *Role* is used to express the interdependencies of these entities, though the attribute *RelatesTo*. *Activity Role*, *Stakeholder* and *Concept* are specializations of *Roles* for processes, persons and concepts in the conceptual perspective. Formally, a *Role* is defined as follows:

- *RelatesTo*: another *Role*.
 - *As*: a single value belonging to the domain of class *RelationshipDomain* = {*suitable*, *dependent*}.
 - *Wrt*: a multi-valued attribute including instances of class *ConceptualObject*.
 - *dueTo*: text string attribute, documenting extra information on the relationship of two roles.

Each *Role* represents a person, program or data store participating in the environment of a process, changed with a specific task and/or responsibility. An instance of the *RelatesTo* relationship is a statement about the interrelationship between two roles in the real world, such as ‘View *VI* relates to table *Class_Info* with respect to the attributes *H_ID*, *EDate* and *#Beds* as dependent *due to* loading reasons’. Note that, since both data and processes can be characterized by SQL statements, their interrelationship can be traced in terms of attributes.

The conceptual perspective is influenced by the Actor Dependency model [4]. In this model, actors *depend* on each other for the accomplishment of goals and the delivery of products. The *dependency* notion is powerful enough to capture the relationships in the context of a data flow, where a data consumer (person, data store or program) depends on the proper function of its data providers, to achieve its mission. Our extension can capture *suitability* as well (e.g., in the case where more than one concepts can apply for the population of the

aggregation, one concept is suitable to replace the other).

In the process of understanding the occurring errors or the design decisions on the architecture of a data warehouse, the conceptual perspective can be exploited in various ways.

1. The *design* of the data warehouse is supported, since the conceptual model serves as a documentation repository for the reasons behind the structure of the data warehouse. The model allows the tracing of the relationships between any pair of persons, programs or data stores. With minimum query facilities of the metadata repository, these interdependencies do not have to be directly stored in all the cases, but can also be computed incrementally, due to their transitivity.
2. The *administration* of the data warehouse is facilitated in several ways. The conceptual model is a good roadmap for the quality management of the warehouse and can act as an entry-point to the logical perspective, since it can enable the user to pass from the abstract relationships of roles to the structure of the system. During runtime, the suitability notion can be used to obtain solutions to potential errors or problems (e.g., alternative paths for the population of the data warehouse) in a straightforward way.
3. Data warehouse *evolution* is supported at two levels. At the *entity* level, the impact of any changes in the architecture of the warehouse can be detected through the sheer existence of dependency links. At the same time, the existence of suitability links suggests alternatives for the new structure of the warehouse. At the *attribute* level, on the other hand, internal changes in the schema of data stores or the interface of software agents can be detected, by using the details of the relationships of the data warehouse roles. For example, the previous statement for the relationship of view *VI* and table *Class_Info* could be interpreted as ‘View *VI* is affected by any changes to table *Class_Info* and especially the attributes *H_ID*, *EDate* and *#Beds*’.

We detail all the aforementioned features and benefits through our example. As we can see in Fig. 11 our conceptual model consists of several entities, including:

- *Aggregation Process* is an *Activity Role*, corresponding to the logical activity *Aggregation*;
- *Information By Class and Date* (for brevity, “*Info By Class*” in Fig. 11) is a *Concept* corresponding to the *Type* “*Class_Info*”;
- *Hospital Information By Date* (for brevity, “*Hospital Info*” in Fig. 11) is also a *Concept* corresponding to the *Type* “*VI*”;
- *Administrator 1*, *Administrator 2* (for brevity, “*Admin 1* and *2*” in Fig. 11) as well as *End User* are the involved *Stakeholders* with obvious roles in the system.

Clearly, *Hospital Information By Date* is an aggregation over *Information By Class and Date* and actually, what the conceptual schema tells us is that it is the *Aggregation Process* that performs the aggregation in practice. Due to the data flow, there exists a strong dependency between the aforementioned roles. *Hospital Information By Date* depends on *Aggregation* since the latter acts as a data pump for the former. *Hospital Information By Date* transitively depends on *Information By Class*. *Administrator 1* depends also on *Information By Class*. Although not depicted in Fig. 11 to avoid overloading the figure, this relationship relies on the idea of political dependency; up to now, it was the *Administrator 1* that

provided the data for this kind of information and in fact, he still does within the data warehouse environment.

Another interesting point shown in Fig. 11 is the idea of *suitability*: according to the stated needs of the users, the concept *Hospital Information By Date* represents information which is *suitable* for the *End User*. It is worth stressing the fact that although some of the entities correspond to processes, other than stakeholders and other than data stores, this does not affect the uniformity and the simplicity of the representation. Also, notice that we do not delve into the fields of conceptual multidimensional aggregation (for example, see [19]) or requirements engineering: the specific approach one can adopt is orthogonal to our modeling.

The notion of suitability helps to support data warehouse evolution. As already mentioned, it is the redundancy in the data warehouse that makes suitability so attractive. Consider the following real-world case, where the information in the final view *VI*, was not 100% consistent. Simple measurements of quality (cf. Section 4) indicated that the COBOL file *CBL* was responsible for this quality problem (upper part of Fig. 12). Thus the original data provider was flawed. Out of the many different choices one could possibly have to resolve the problem, the most suitable one proved to be the exploitation of redundancy. An alternative population scheme for view *VI* used the source file *CBL'* as initial input. *CBL'* is the corresponding data store of the *Concept Information By Department and Date* (“*Info By Dept*” in Fig. 12) and captures the number of beds by department of hospital (instead of class of beds).

Sometimes, suitability in the conceptual model of a data warehouse can be automatically derived from aggregate reasoners and algorithms proposed by previous research on view containment [20–23]. Again, we would like to stress that suitability in our proposal is not restricted to either persons or resources but can uniquely cover all the entities in a conceptual model of a data warehouse.

Apart from the support for evolution of the data warehouse at the entity level, the proposed model is also capable of supporting evolution at the attribute level. As mentioned, the relationship

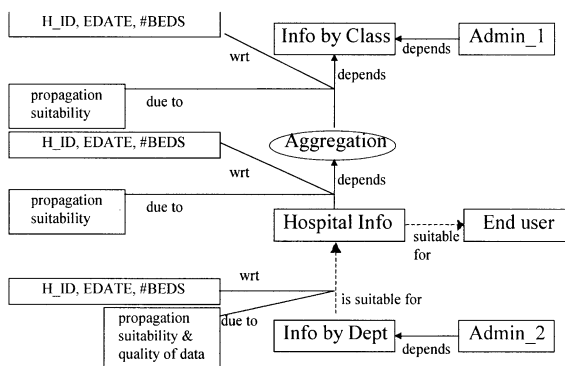


Fig. 11. The conceptual perspective for the example.

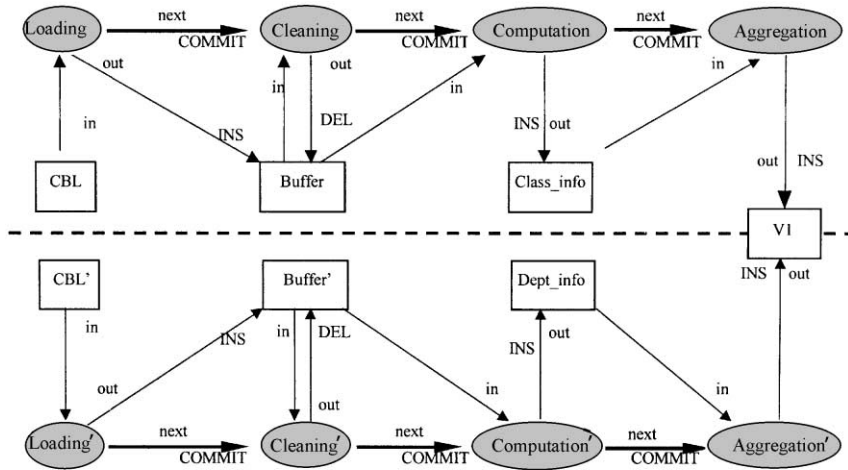


Fig. 12. Data warehouse evolution. Upper part: original problematic configuration. Lower part: the new configuration.

between two roles is normally expressed in terms of fields. In our example, *Aggregation* depends on a subset of the attributes of *Info By Class*, namely *H_ID*, *DATE* and *#Beds*. If the attribute *Class_ID* changes, e.g., due to change of type, removal, or renaming, the final information of the *Concept* “*Hospital Info*” is not affected at all. On the other hand, changes in any of the attributes depicted in Fig. 11 clearly affect the information delivered to end users.

It is interesting to note the political conflict that takes place due to the proposed change. As we can see, removing *Info By Class* from the data flow, automatically affects the Stakeholder *Administrator 1*, who depends on *Info By Class*. A simple query in the metadata repository for the dependents of the entity *Info By Class* could give a warning for the political vibrations coming from such a decision¹. Finally, we should also mention that hidden within the short story that we have just summarized is the idea of quality measurement, which will be detailed in Section 4.

¹We take the opportunity to stress the huge importance of politics in the development lifecycle of a data warehouse. See [6,17] for more details.

3.5. Facilitating data warehouse design through consistency checking in the metadata repository

To ensure the validity of the representation in the metadata repository, consistency checks can be performed during the design of the data warehouse, by defining constraints, rules and views in the language Telos. The following view finds out whether the types used as inputs of an activity are stored in the respective data stores used as inputs of the agent, executed by the activity.

```

QueryClass InconsistentInTypes isA Type
with constraint
  c: $ exists d/DataStore ac/Activity ag/
    Agent
    (ac input this) and (ac executedBy ag)
    and
    (ag input d) and not (this storedIn d) $
end
    
```

Other simple constraints involve the local structure of the process elements. For example, split transition elements must have at least one incoming edge and more than one outgoing edge. The timestamps of the agent should also be consistent with its state. The repository can also be used by external programs to support the

Depender Role	Dependee Role	As	Wrt
Activity	Input types	dependent	schema of input types
Type	Populating activity	dependent	schema of type
Conceptual Object	Conceptual Object	dependent	the set of fields wrt. which the dependee depends (transitivity of dependency)

Fig. 13. Simple rules for the production of role interdependencies.

execution of consistency checking algorithms as proposed in [24,25].

3.6. Repository support for the automatic derivation of role interdependencies

The *Roles* of the conceptual perspective can be directly assigned by the data warehouse administrator, or other interested stakeholders. Still, one could argue that this is too much of an administrative burden, based on the number of relationships that should be traced in the repository. By deductive reasoning in the metadata repository, we can impose simple rules to deduce these interdependencies from the structure of data stores and activities. For example, we can derive dependency relationships by exploiting the structure of the logical perspective of the metadata repository. Also, interdependencies do not have to be directly stored in all the cases, but can also be computed incrementally, due to the transitivity of their nature. In Fig. 13 we show three simple rules which can be used to derive the production of role interdependencies. They can also be implemented in the metadata repository.

4. Process quality

In this section we present how the process metamodel is linked to the metamodel for data warehouse quality proposed in [2]. Moreover, we complement this quality metamodel with specific dimensions for data warehouse operational processes.

4.1. Terminology for quality management

The quality metamodel in [2] customizes the GQM approach of [11] for data warehouse

environments. In this section, we adopt the same metamodel for the operational processes of the data warehouse.

Each object in the data warehouse is linked to a set of quality goals and a set of quality factors (Fig. 15). A *quality goal* is an abstract requirement, defined on data warehouse objects, and documented by a purpose and the stakeholder interested in it, e.g., ‘improve the availability of source S1 until the end of the month in the viewpoint of the data warehouse administrator’. *Quality dimensions* (e.g., ‘availability’) are used to group quality goals and factors into different categories. A *Quality Factor* represents a quantitative assessment of a particular aspect of a data warehouse object, i.e., it relates quality aspects both to actual measurements and expected ranges for these quality values. Finally, the method of measurement is attached to a quality factor through a *measuring agent*.

The bridge between the abstract, subjective quality goals and the specific, objective quality factors is determined through a set of *quality queries*, to which quality factor values are provided as possible answers. Such queries are the outcome of the methodological approach described in [3] which offers template quality factors and dimensions, defined at the metadata level and instantiates them, for the specific data warehouse architecture under examination. As a result of the goal evaluation process, a set of improvements (i.e., design decisions) can be proposed, in order to achieve the expected quality.

4.2. Quality dimensions and factors for data warehouse operational processes

ISO 9126 standard [ISO97] on software implementation and evaluation provides a general understanding of how to measure the quality of software systems. Data warehouses do not stray

from these general guidelines; thus we adopt the standard as a starting point. ISO 9126 is based on six high level quality dimensions (*Functionality, Reliability, Usability, Efficiency, Maintainability, Portability*). Time and budget constraints in the development of a data warehouse cause the addition of *Implementation Effectiveness*. The dimensions are analyzed to several sub-dimensions (Fig. 14).

ISO 9126 does not provide specific quality factors. To deal with this shortcoming, Appendix A gives a set of quality factors customized for the case of data warehouse operational processes. It does not detail the whole set of possible factors for all operational data warehouse processes, but rather, we intend to come up with a minimal representative set. This set of quality factors can be

refined and enriched by the data warehouse stakeholders with customized factors. Once again, we encourage the use of “templates” in a way that fits naturally with the overall metadata framework that we propose.

4.3. Relationships between processes and quality

Quality goals describe *intentions* of data warehouse users with respect to the status of the data warehouse. In contrast, our process model describes *facts* about the current and previous status of the data warehouse and what activities are performed in the data warehouse. However, the reason behind the existence of a process is a quality goal. For example, a data cleaning process is executed in the data staging area in order to

Quality Dimension	Quality Factors
<u>Functionality:</u>	
<i>Suitability</i>	Benchmark scores, number of software requirements met.
<i>Accuracy</i>	Completeness, accuracy and consistency of data (which are the final product).
<i>Interoperability</i>	Number of modules unable to interact with specified systems.
<i>Compliance</i>	Number of modules not compliant with application related standards or conventions or regulations in laws and similar prescriptions.
<i>Security</i>	Number of modules unable to prevent unauthorized access, whether accidental or deliberate, to programs and data.
<u>Reliability:</u>	
<i>Maturity</i>	Frequency of failure by faults in the software.
<i>Fault tolerance</i>	Number of occasions where the software was unable to maintain a specified level of performance (in cases of software faults or of infringement of its specified interface).
<i>Recoverability</i>	Number of occasions where the software was unable to re-establish its level of performance and recover the data affected in the case of a failure. The time and effort needed for this reestablishment.
<u>Usability:</u>	
<i>Understandability</i>	Percentage of acceptance by the users.
<i>Learnability</i>	Percentage of acceptance by the users.
<i>Operability</i>	Existence of time prediction for operation, number of error alerts, available built in functions for ETL tasks
<u>Software Efficiency:</u>	
<i>Time behavior</i>	Response times, processing times and throughput rates.
<i>Resource Behavior</i>	Volume of extracted and loaded data, maximum data size the software can handle.
<u>Maintainability:</u>	
<i>Analyzability</i>	Comment percentage, readability index.
<i>Changeability</i>	Number of logical paths in a module, control flow intersection, GOTO statements, size, etc.
<i>Stability</i>	Same as Changeability
<i>Testability</i>	Same as Changeability
<i>Portability:</i>	Same as Maintainability
<u>Implementation Effectiveness</u>	Man-hours spent, tasks completed.

Fig. 14. Software quality dimensions [26] and proposed quality factors in data warehouse environments.

improve the accuracy of the data warehouse. We have represented this interdependency between processes and quality goals by establishing a relationship between roles and data warehouse objects in the conceptual perspective of the process model (relationship *Imposed On*). This is shown in the upper part of Fig. 15.

Our model is capable of capturing all dependency types mentioned in [4]. *Task dependencies*, where the dependum is an *activity*, are captured by assigning the appropriate role to the attribute *Wrt* of the relationship. *Resource dependencies*, where the dependum is the availability of a *resource*, are modeled when fields or concepts populate this attribute. The relationship *ExpressedFor* relates a role to a high-level quality goal; thus *Goal dependencies*, dealing with the possibility of making a condition true in the real world, are captured from the model, too. *Soft-goal dependencies* are a specialization of goal dependencies, where evaluation cannot be done in terms of concrete quality factors.

The lower part of Fig. 15 represents the relationship between processes and quality on a more operational level. The actions of an agent in the data warehouse *affect* the expected or measured quality factors of some data warehouse objects. For example, a data cleaning process affects the availability of a source: it decreases the amount of time during which it can be used for regular operations. Consequently, this process will affect the quality factors *Average Load*, *CPU state*, *Available Memory* defined on a *Source*. All these quality factors are concrete representations of the abstract notion *Availability*—the relevant quality

dimension. The effect of a data warehouse process must always be confirmed by new measurements of the quality factors. Unexpected effects of data warehouse processes can be detected by comparing the measurements with the expected behavior of the process. The measurement of the quality of the particular agents through their own quality factors is analyzed in [1].

A *Quality Query* provides the methodological bridge to link the high-level, user-oriented, subjective quality goals to the low-level, objective, component-oriented quality factors. The vocabulary (or domain) of quality queries with respect to the process model is the set of data warehouse activities, which can be mapped to reasons (roles) and conditions (of agents) of a specific situation.

Let us return to our hospital example to clarify how the process and the quality metamodels interplay, and how the different perspectives gracefully map to each other. More than one of our experiences in the public sector indicated a need of total quality of data, were no errors were allowed and no information was missing. Thus, the quality goal is ‘100% quality of data delivered to the end users’.

For the purpose of our example, we narrow this high level goal to the subgoal, 100% consistency of the produced information. There are two objects involved in this quality goal, namely the quality dimension *consistency* and the role *end user*. Both entities, as well as the quality goal itself, belong to the conceptual perspective and can be used to explain why the chain of processes exists: to bring clean, consistent information to the involved stakeholders.

According to the GQM paradigm, a good start to examine a situation would be to find out its current status. With respect to the elements of the process model, the basic question over process status is naturally over the *correctness* dimension: are all activities performing as they should? The question, itself belonging to the logical perspective, involves an object of the logical part of the process metamodel: *activities*. If one applies the methodology of [3], this question would directly be analyzed to five consequent questions, each involving one of the activities *Loading*, *Cleaning*, *Computation*, *Aggregation*, and *Populate VI*.

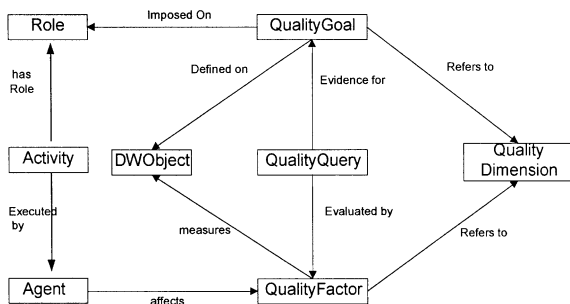


Fig. 15. Relationships between processes and quality.

The actual quality evaluation of the produced software is done in terms of concrete measurements. On the physical perspective (where the quality factors that provide the measurements belong) the measurements involve specific software agents. In our case, the quality factor *Correctness though White Box Software Testing (WBCorrectness* in the sequel) performs this kind of measurements.

The discrimination of logical, conceptual and physical perspectives is proven useful once more, in the quality management of the data warehouse: the quality goals can express “why” things have happened (or should happen) in the data warehouse, the quality questions try to discover “what” actually happens and finally, the quality factors express “how” this reality is measured (Fig. 16). On top of this, we have organized a seamless integration of the process and quality models, by mapping the objects of the same perspectives to each other.

In Fig. 17 we can also see the assignment of quality factors to the various objects at the physical perspective. We assign the quality factor *WBCorrectness* to the software agents and the quality factors *Consistency* and *Completeness* to the data stores. A simple formula derives the quality of the data stores in the latest steps of the data flow from the quality of the previous data stores and software agents. Let us take *Consistency*

for example:

$$\begin{aligned} \text{consistency}(ds) &= \prod \text{correctness}(a) \prod \text{consistency}(ds'), \end{aligned} \tag{1}$$

where *ds* is the data store under consideration, *a* denotes the agents having *ds* as their output and *COMMIT* semantics and *ds'* is any data store different from *ds*, serving as input to the agents *a*.

Clearly, the consistency of the final view *V1* depends on the consistency of all the previous data stores and the correctness of all the involved software agents. Although arguably naive, the formula fitted perfectly in our real world scenario.

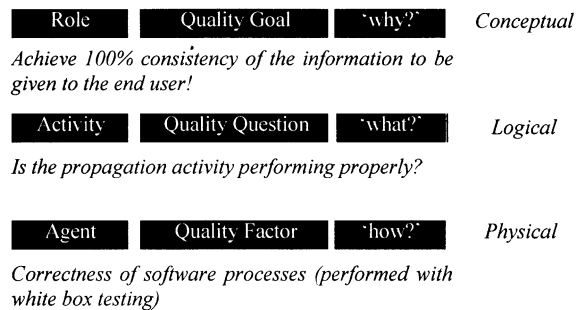


Fig. 16. Perspectives and interrelationships for the process and quality metamodels.

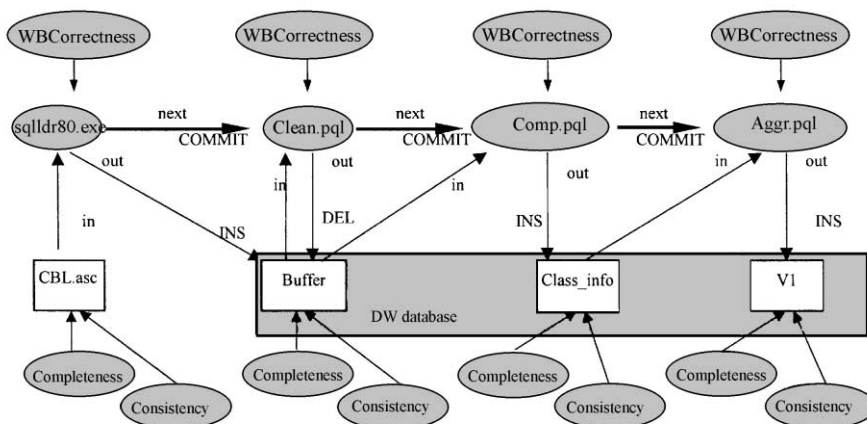


Fig. 17. Application of quality factors to the entities of the example.

4.4. Exploitation of quality modeling in the repository for data warehouse administration

The different instantiations of the quality model can be exploited to assist both the design and the administration of the data warehouse. In [27], it has been shown how the classification of quality dimensions, factors and goals can be combined with existing algorithms to address the data warehouse design problem (i.e., the selection of the set of materialized views with the minimum overall “cost” that fulfill all the quality goals set by the involved stakeholders). As far as the administration of the warehouse is concerned, the information stored in the repository may be used to find deficiencies in a data warehouse.

To show how the quality model is exploited, we take the following query. It returns all data cleaning activities which have decreased the availability of a data store according to the stored measurements. The significance of the query is that it can show that the implementation of the data cleaning process has become inefficient.

```
GenericQueryClass DecreasedAccuracy isA
DWCleaningAgent with parameter
  ds: DataStore
constraint
  c: $ exists qf1, qf2/DataStoreAccuracy
    t1, t2, t3/Commit_Time v1, v2/Integer
    (qf1 onObject ds) and (qf2 onObject ds)
    and
    (this affects qf1) and (this affects
    qf2) and
    (this executedOn t3) and (qf1 when t1)
    and (qf2 when t2) and
    (t1<t2) and (t1<t3) and (t3<t2) and
    (qf1 achieved v1) and (qf2 achieved v2)
    and (v1>v2) $
end
```

The query has a data store as parameter, i.e., it will return only cleaning processes that are related to the specified data store. The query returns the agents which have worked on the specified data store and which were executed between the measurements of quality factors *qf1* and *qf2*, and the measured value of the newer quality factor

is lower than the value of the older quality factor. The attribute *executedOn* of an agent represents the time when this agent was executed.

5. Repository support for data warehouse description and evolution

Summarizing the discussion so far, during the *design* phase, the user can check the consistency of his/her design, to determine any violations of the business logic of the data warehouse, or the respect of simple rules over the structure of the data warehouse schema. During the *administration* phase, we can use the repository to discover quality problems.

In this section, we continue to show how the metadata repository can be exploited in different ways. First, we complement the perception of data warehouses as collections of materialized views with a precise operational description of the content of data warehouse tables. Second, a particular task in the data warehouse lifecycle, *data warehouse evolution*, is examined separately, in order to determine possible impacts, when the schema of a particular table in the data warehouse changes.

5.1. Why data warehouses are not (just) collections of materialized views

Many database researchers have considered data warehouses to be collections of materialized views, organized in strata where the views of a particular stratum are populated from the views of a lower stratum. For example, in [17] where the papers of three major database conferences related to data warehousing, between the years 1995 and 1999, are classified into different categories, almost half of the papers (around 46%) deal with view maintenance and integration issues. The papers on view maintenance have focused on algorithms for updating the contents of a view in the presence of changes in the sources. Papers related to integration have targeted the production of a single interface for the processing of distributed heterogeneous data, along with query processing techniques for that cause and resolution of

conflicts at the schema level. One can observe, thus, that most of the performed research has been dedicated to *what* should be extracted and loaded, instead of *how* this process is actually performed. Practical aspects of extraction, loading and conversion processes, such as scheduling, declarative process definition, or data peculiarities (source format, errors, conversions) are clearly neglected (see [17] for a broader discussion).

In summary, although the abstraction of treating data warehouses as strata of materialized views, has efficiently served the purpose of investigating the issues of (incremental) view maintenance, it lacks the ability to accurately describe the real content of data warehouse tables, due to the fact that all the intermediate processing, transformation and reorganization of the information is systematically absent from most research efforts.

Our modeling approach follows a different path, by treating data warehouse processes as first class citizens. The semantic definitions for data warehouse views are not assigned directly by the designer but result from combining the respective definitions of the data warehouse processes. Thus, we can argue that there are two ways to define the semantics of a table in the data warehouse:

- A *Type* can be defined as a materialized view over its previous *Types* in the data flow. This is the definition of what the *Type* should contain ideally, i.e., in a situation where no physical transformations or cleaning exists.
- A *Type* can be defined as a view again, resulting from the adoption of our modeling approach. In this case, the resulting definition explains how the contents of the *Type* are *actually* produced, due to schema heterogeneity and bad quality of data.

Of course, both kinds of definition are useful but only the former has been taken into consideration in previous research. To complement this shortcoming, the rest of this subsection is dedicated to showing how we can derive view definitions from the definitions of their populating processes.

To give an intuition of the difference between the two approaches, consider the example of Fig. 5. Ideally, we would like to express the view

VI in terms of the input file *CBL* (or its relational counterpart, table *Buffer*). A simple formula suffices to give this *intentional* semantics:

```
SELECT H_ID, EUROPEAN(DATE) AS EDATE,
CLASS_A+CLASS_B+Class_c AS SUM_BEDS FROM CBL
```

On the other hand, reality is clearly different. It involves the identification of multiple rows for the same hospital at the same time period, and the restructuring of the information to a normalized format before the loading of data in view *VI*. The full expression capturing this operational semantics is definitely more complex than the simple SQL query denoting the intentional semantics.

To construct expressions for data warehouse tables with respect to their operational semantics, we constrain ourselves to the case where we are able to construct an acyclic, partially ordered graph of activities (produced by proper queries in ConceptBase). Thus, we can treat the whole set of data warehouse activities as an ordered list. Mutually exclusive, concurrent paths in the partially ordered graph are treated as different lists (the execution trace determines which list is considered each time).

Furthermore, a set of types belonging to the set *SourceSchema*, denoting all the types found in the data sources, are treated as source nodes of a graph. For the rest of the types, we can derive an SQL expression by using existing view reduction algorithms, such as [28] (corrected with the results of [29–31], to which we shall refer to as [28+]), [32–36]. Our algorithm is applicable to graphs of activities that do not involve updates. In most cases, an update operation can be considered as the combination of insertions and deletions or as the application of the appropriate function to the relevant attributes.

The results of the application of this algorithm to our example are shown in Fig. 18. For convenience, we break composite definitions of table expressions into the different lines of Fig. 19. For example, when the third iteration ($i = 3$) refers to the definition of table *Buffer*, it does so with respect to the definition of line 2 ($i = 2$). The expression of a single type can also be computed locally.

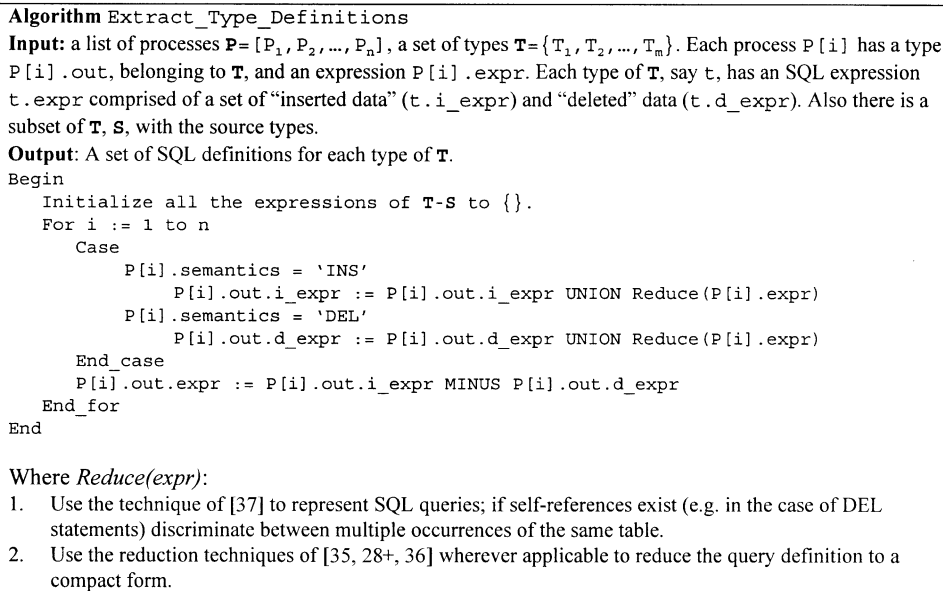


Fig. 18. Algorithm for extracting the definition of a type in the repository.

i	Expression
1	Buffer.expr = Buffer.i_expr:= SELECT * FROM CBL
2	Buffer.d_expr:= (SELECT * FROM CBL C1 WHERE EXISTS (SELECT C2.H_ID, C2.DATE FROM CBL C2 WHERE C1.H_ID=C2.H_ID AND C1.DATE=C2.DATE GROUP BY H_ID,DATE HAVING COUNT(*)>1) Buffer.expr:= (SELECT * FROM CBL) MINUS (SELECT * FROM CBL C1 WHERE EXISTS (SELECT C2.H_ID, C2.DATE FROM CBL C2 WHERE C1.H_ID=C2.H_ID AND C1.DATE=C2.DATE GROUP BY H_ID,DATE HAVING COUNT(*)>1)
3	Class_info.expr = Class_info.i_expr:= SELECT H_ID, EUROPEAN(EDATE) AS EDATE, 'A', CLASS_A AS #BEDS FROM BUFFER WHERE CLASS_A <> 0 UNION SELECT H_ID, EUROPEAN(EDATE) AS EDATE, 'B', CLASS_B AS #BEDS FROM BUFFER WHERE CLASS_B <> 0 UNION SELECT H_ID, EUROPEAN(EDATE) AS EDATE, 'C', CLASS_C AS #BEDS FROM BUFFER WHERE CLASS_C <> 0;
4	V1.expr = V1.i_expr:= SELECT H_ID, EDATE, SUM(#BEDS) AS SUM_BEDS FROM CLASS_INFO GROUP BY H_ID, EDATE

Fig. 19. SQL expressions of the types of the example.

5.2. Repository support for data warehouse evolution

The data warehouse is constantly evolving. New sources are integrated in the overall architecture

from time to time. New enterprise and client data stores are built in order to cover novel user requests for information. As time passes by, users seem more demanding for extra detailed information. Due to these reasons, not only the structure

but also the processes of the data warehouse evolve.

The problem arises to keep the data warehouse objects and processes consistent in the presence of changes. For example, changing the definition of a materialized view in the data warehouse triggers a chain reaction: the update process must evolve (both the refreshment and the cleaning steps), and the old, historical data must be migrated to the new schema (possibly with respect to new selection conditions). All data stores of the data warehouse and client level which are populated from this particular view must be examined with respect to their schema, content and population processes.

In our approach, we distinguish two kinds of impact of a hypothetical change:

- *Direct impact*: the change in the data warehouse object imposes that some action must be taken against an affected object. For example, if an attribute is deleted from a materialized view, then the activity which populates it must also be changed accordingly.
- *Implicit impact*: the change in the data warehouse object might change the semantics of another object, without obligatorily changing the structure of the latter.

Our model enables us to construct a partially ordered graph: for each *Type* instance, say t , there is a set of types and activities, used for the population of t (“before” t), denoted as $\mathbf{B}(t)$. Also, there is another set of objects using t for their population (“after” t), denoted as $\mathbf{A}(t)$. We can recursively compute the two sets from queries on the metadata repository of process definitions. Queries for the successor and after relationships can be defined in a similar way.

Suppose that the final SQL expression of a type t , say e , changes into e' . In the spirit of [38], we can use the following rules for schema evolution in a data warehouse environment (we consider that the changes abide by the SQL syntax and the new expression is valid):

- If the *select clause* of e' has an extra attribute from e , then propagate the extra attribute to the base relations: there must be at least one path from one type belonging to a *SourceSchema* to

an activity whose out expression involves the extra attribute. If we delete an attribute from the select clause of a type, it must not appear in the select clause of the processes that directly populate the respective type, as well as in the following types and the processes that use this type. In the case of attribute addition, the impact is direct for the previous objects $\mathbf{B}(t)$ and implicit for the successor objects $\mathbf{A}(t)$. For deletion the impact is direct for both categories.

- If the *where clause* of e' is more strict than the one of e , the *where clause* of at least one process belonging to $\mathbf{B}(t)$ must change identically. If this is not possible, a new process can be added before t simply deleting the respective tuples through the expression $e' - e$. If the *where clause* of e' is less strict than the one of e , subsumption techniques [20–22,39] determine which types can be used to calculate the new expression e' of t . The *having clause* is treated in the same fashion. The impact is direct for the previous and implicit for the successor objects.
- If an attribute is deleted from the *group by clause* of e , at least the last activity performing a *group-by* query should be adjusted accordingly. All consequent activities in the population chain of t must change too (as if an attribute has been deleted). If this is not feasible we can add an aggregating process performing this task exactly before t . If an extra attribute is added to the *group by clause* of e , then at least the last activity performing a *group-by* query should be adjusted accordingly. The check is performed recursively for the types populating this particular type, too. If this fails, the subsumption techniques mentioned for the *where-clause* can be used for the same purpose again. The impact is direct for both previous and successor objects. Only in the case of attribute addition it is implicit for the successor objects.

Returning to our example, suppose that we decide to remove attribute *CLASS_C* from the table *BUFFER*. This change has the following impacts:

Previous	Loading	SELECT H_ID, DATE, CLASS_A, CLASS_B, CLASS_C FROM CBL
	Cleaning	SELECT H_ID, DATE, CLASS_A, CLASS_B, CLASS_C FROM BUFFER B1 WHERE EXISTS (SELECT B2.H_ID, B2.DATE FROM BUFFER B2 WHERE B1.H_ID = B2.H_ID AND B1.DATE = B2.DATE GROUP BY H_ID,DATE HAVING COUNT(*) > 1)
Next	Computation	SELECT H_ID, EUROPEAN DATE) AS EDATE, 'A', CLASS_A AS #BEDS FROM BUFFER WHERE CLASS_A <> 0 UNION SELECT H_ID, EUROPEAN DATE) AS EDATE, 'B', CLASS_B AS #BEDS FROM BUFFER WHERE CLASS_B <> 0 UNION SELECT H_ID, EUROPEAN DATE) AS EDATE, 'C', CLASS_C AS #BEDS FROM BUFFER WHERE CLASS_C <> 0
	Cleaning	As is, there is no existence of CLASS_C in its definition any more.

Fig. 20. Impact of the removal of attribute CLASS_C from table BUFFER to the definitions of the affected activities.

- The activities belonging to *Previous*(BUFFER), namely *Loading* and *Cleaning*, must change accordingly, to remove CLASS_C from their *select clause*.
- The activities belonging to *Next*(BUFFER), namely *Cleaning* and *Computation* must also change to remove any appearance of CLASS_C in their query expression.

Note that *Cleaning* participates both in the *Previous* and in the *Next* list. Fig. 20 shows the impact of these changes. Attribute CLASS_C is removed from the select list of the *Previous* activities and from the body of the queries of the *Next* list.

Due to the existence of implicit impacts, we do not provide a fully automated algorithmic solution to the problem, but rather, we sketch a methodological set of steps, in the form of suggested actions to perform this kind of evolution. Similar algorithms for the evolution of views in data warehouses can be found in [38,40]. A tool could easily visualize this evolution plan and allow the user to react to it.

6. Related work

In this section we discuss the state of art and practice for research efforts, commercial tools and standards in the fields of process and workflow

modeling, with particular focus on data warehousing.

6.1. Standards

The standard [7] proposed by the *Workflow Management Coalition (WfMC)* 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 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 entities external to the workflow, such as system and environmental data or the organizational model are roughly described.

The *MetaData Coalition (MDC)*, is an industrial, non-profit consortium which aims to provide a standard definition for enterprise metadata shared between databases, CASE tools and similar applications. The *Open Information Model (OIM)* [41] is a proposal (led by Microsoft) for the core metadata types found in the operational and data warehousing environment of enterprises. The OIM uses UML both as a modeling language and as the basis for its core model. The OIM is divided in *packages* extend UML in order to address different areas of information management. The *Database and Warehousing Model* is composed from the *Database Schema Elements* package, the *Data Transformations*

Elements package, the *OLAP Schema Elements* package and the *Record Oriented Legacy Databases* package. The *Database Schema Elements* package contains three other packages: a *Schema Elements* package (covering the classes modeling tables, views, queries, indexes, etc.), a *Catalog and Connections* package (covering physical properties of a database and the administration of database connections) and a *Data Types* package, standardizing a core set of database data types. The *Data Transformations Elements* package covers basic transformations for relational-to-relational translations. It does not deal with data warehouse process modeling (i.e., it does not cover data propagation, cleaning rules, or querying), but covers in detail the sequence of steps, the functions and mappings employed and the execution traces of data transformations in a data warehouse environment.

6.2. Commercial tools

Basically, commercial ETL tools are responsible for the implementation of the data flow in a data warehouse environment which is only one (albeit important) of the data warehouse processes. Most ETL tools are of two flavors: *engine-based*, or *code-generation-based*. The former assumes that all data have to go through an engine for transformation and processing. In code-generating tools all processing takes place only at the target or source systems. There is a variety of such tools in the market; we mention three engine-based tools, from Ardent [42], DataMirror [43] and Microsoft [37,44], and one code-generation-based from ETI [45].

6.3. Research efforts

Workflow modeling: There is a growing research interest in the field of workflow management. Sadiq and Orlowska [24] use a simplified workflow model, based on [7], using *tasks* and *control flows* as its building elements. The authors present an algorithm for identifying structural conflicts in a control flow specification. The algorithm uses a set of graph reduction rules to test the correctness

criteria of *deadlock freedom* and *lack-of-synchronization freedom*. In [25] the model is enriched with modeling constructs and algorithms for checking the consistency of workflow temporal constraints. In [46], the authors propose a conceptual model and language for workflows. The model gives the basic entities of a workflow engine and semantics about the execution of a workflow. The proposed model captures the mapping from workflow specification to workflow execution (in particular concerning exception handling). Importance is paid to task interaction, the relationship of workflows to external agents and the access to databases. Other aspects of workflow management are explored in [47,48]. In [49] a general model for transactional workflows is presented. A transactional workflow is defined to consist of several tasks, composed by constructs like ordering, contingency, alternative, conditional and iteration. Nested workflows are also introduced. Furthermore, correctness and acceptable termination schedules are defined over the proposed model. In [50] several interesting research results on workflow management are presented in the field of electronic commerce, distributed execution and adaptive workflows. A widely used web server for workflow literature is maintained by [51].

Process modeling. Process and workflow modeling have been applied in numerous disciplines. In [16] the authors propose a software process data model to support software information systems with emphasis on the control, documentation and support of decision making for software design and tool integration. Among other features, the model captures the representation of design objects (“what”), design decisions (“why”) and design tools (“how”). A recent overview on process modeling is given in [52], where a categorization of the different issues involved in the process engineering field is provided. The proposed framework consists of four different but complementary viewpoints (expressed as “words”): the *subject* world, concerning the definition of the process with respect to the real world objects, the *usage* world, concerning the rationale for the process with respect to the way the system is used, the

system world, concerning the representation of the processes and the capturing of the specifications of the system functionality and finally, the *development* world, capturing the engineering meta-process of constructing process models. Each world is characterized by a set of *facets*, i.e., attributes describing the properties of a process belonging to it.

Data quality and quality management. There has been a lot of research on the definition and measurement of data quality dimensions [53–56]. A very good review of research literature is found in [9]. Jarke et al., [1] give an extensive list of quality dimensions for data warehouses, and in particular data warehouse relations and data. Several goal hierarchies of quality factors have been proposed for software quality. For example, the GE Model [57] suggests 11 criteria of software quality, while Boehm [58] suggests 19 quality factors. ISO 9126 [26] suggests six basic factors which are further refined to an overall 21 quality factors. In [59] a comparative presentation of these three models is offered and the SATC software quality model is proposed, along with metrics for all their software quality dimensions. In [60] a set of four basic quality dimensions for workflows is suggested also. Variants of the Goal-Question-Metric (GQM) approach are widely adopted in software quality management [11]. A structured overview of the issues and strategies, embedded in a repository framework, can be found in [61]. Jarke et al. [1,10] provide extensive reviews of methodologies employed for quality management, too.

Research focused specifically on ETL. The *AJAX* data cleaning tool developed at INRIA [62] deals with typical data quality problems, such as the object identity problem, errors due to mistyping and data inconsistencies between matching records. *AJAX* provides a framework wherein the logic of a data cleaning program is modeled as a directed graph of data transformations (mapping, matching, clustering and merging transformations) that start from some input source data. *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.

6.4. Relationship of our proposal to state-of-the-art research and practice

Our approach has been influenced by ideas on dependency and workflow modeling stemming from [4,7,16,24,46,63,64]. As far as the standards are concerned, we found both the Workflow Reference Model and the Open Information Model too abstract for the purpose of a repository serving well focused processes like the ones in a data warehouse environment. First, the relationship of an activity with the data it involves is not really covered, although this would provide extensive information of the data flow in the data warehouse. Second, the separation of perspectives is not clear, since the standards focus only on the structure of the workflows. To compensate this shortcoming, we employ the basic idea of the Actor-Dependency model [4] to add a conceptual perspective to the definition of a process, capturing the reasons behind its structure. Moreover, we extend [4] with the notion of suitability. As far as data quality and quality engineering are concerned, we have taken into account most of the previous research for our proposed quality dimensions and factors.

7. Conclusions

In this paper we have described a metamodel for data warehouse operational processes and techniques to design, administrate and facilitate the evolution of the data warehouse through the exploitation of the entities of this metamodel. This metamodel takes advantage of the clustering of its entities in logical, physical and conceptual perspectives, involving a high-level conceptual description, which can be linked to the actual structural and physical aspects of the data warehouse architecture. This approach is integrated with the results of previous research, where data warehouse architecture and quality metamodels have been proposed assuming the same categorization.

The physical perspective of the proposed metamodel covers the execution details of data warehouse processes. At the same time, the logical perspective is capable of modeling the structure of complex activities and capture all the entities of the Workflow Management Coalition Standard. Due to the data oriented nature of the data warehouse activities, their relationship with data stores is particularly taken care of, through clear and expressive semantics in terms of SQL definitions. This simple idea reverts the classical belief that data warehouses are simply collections of materialized views. Instead of directly assigning a naïve view definition to a data warehouse table, we can deduce its definition as the outcome of the combination of the processes that populate it. This new kind of definition does not necessarily refute the existing approaches, but rather complements them, since the former provide the operational semantics for the content of a data warehouse table, whereas the latter give an abstraction of its intentional semantics. The conceptual perspective is a key part of our approach as in the Actor Dependency model [4]. Also, we generalize the notion of role to uniformly capture any person, program or data store participating in the system. Furthermore, the process metamodel is linked to a quality metamodel, thereby facilitating the monitoring of the quality of data warehouse processes and a quality-oriented evolution of the data warehouse.

In the process of taking design decisions or understanding the occurring errors over the architecture of a data warehouse, the proposed metamodel can be exploited in various ways. As far as the design of the data warehouse is concerned, simple query facilities of the metadata repository are sufficient to provide the support for consistency checking of the design. Moreover, the entities of the conceptual perspective serve as a documentation repository for the reasons behind the structure of the data warehouse. Second, the administration of the data warehouse is also facilitated in several ways. The measurement of data warehouse quality, through the linkage to a quality model, is crucial in terms of enabling the desired functionality during the everyday use of the warehouse.

Evolution is supported by the role interdependencies with two ways. As the entity level, the impact of any changes in the architecture of the warehouse can be detected through the existence of dependency links. At the same time, the existence of suitability links suggests alternatives for the new structure of the warehouse. At the attribute level, on the other hand, internal changes in the schema of data stores or the interface of software agents can also be detected by using the details of the relationships of the data warehouse roles.

We have used our experiences from real world cases as a guide for the proposed metamodel. As far as the practical application of our ideas in the real world is concerned, we find that field of ETL and data warehouse design tools is the most relevant to our research. As a partial evaluation of our ideas and to demonstrate the efficiency of our approach, we have developed a prototype ETL tool.

Research can follow our results in various ways. First, it would be interesting to explore automated ways to assist the involved stakeholders (data warehouse designers and administrators) to populate the metadata repository with the relevant information. An example for how this can be achieved is explained in Section 3.4 for the suitability interrelationships. Specialized tools and algorithms could assist in extending this kind of support for more aspects of the proposed metamodel. Also, in this paper we have dealt only with the operational processes of a data warehouse environment. Design processes in such an environment may not fit this model so smoothly. It would be worth trying to investigate the modeling of the design processes and to capture the trace of their evolution in a data warehouse. Finally, we have used the *global-as-view* approach for the definitions of the data warehouse processes, i.e., we reduce these definitions in terms of the sources in the warehouse architecture. We plan to investigate the possibility of using the *local-as-view* approach, which means reducing both the process definitions and the data sources to a global enterprise model. The local-as-view approach appears to be more suitable to environments where a global, corporate view of data is present and thus, provides several benefits that the global-as-view approach lacks [65].

Acknowledgements

This research is sponsored in part (a) by the European Esprit Projects “DWQ: Foundations of Data Warehouse Quality”, No. 22469, and “MEMO: Mediating and Monitoring Electronic Commerce”, No. 26895, (b) by the Deutsche Forschungsgemeinschaft (DFG) under the Collaborative Research Center IMPROVE (SFB 476) and (c) by the Greek General Secretariat of Research and Technology in the context of the project “Data Quality in Decision Support Systems” of the French–Greek Cooperation Programme “Plato”, 2000. We would like to thank our DWQ partners who contributed to the progress of this work, especially Mokrane Bouzeghoub, Manfred Jeusfeld, Maurizio Lenzerini, and Timos Sellis. Many thanks are also due to the anonymous reviewers for their useful comments and the interesting issues they have raised.

We organize the presentation of these quality factors around the lifecycle stage during which they are introduced. The examined lifecycle stages are *requirements analysis*, *system design*, *implementation and testing* (including the quality of the software module itself), *management* and *administration*. For each of these stages, we list the involved stakeholder and its products (depicted in the header of each table). Each of the following tables comprises three columns: the leftmost describes a high level quality dimension, the middle column shows the related quality factors of the dimension and the rightmost column suggests measurement methods for each quality factor. We make a full presentation, including both generic quality factors, applicable to any kind of software module and data warehouse specific quality factors. The former are presented in gray background. To provide the list of these quality factors we have relied on existing research and practitioner publications, market tools and personal experiences [59,67,66,3,1].

Appendix A

In this Appendix, we list quality factors for the software involved a data warehouse environment.

Who: Designer	When: Requirements Analysis	Output: Requirements Document Role Conceptual model
Generic Requirements Quality	Ambiguity	Number of weak phrases
		Number of optional phrases
	Completeness	Number of to-be-announced, to-be-added
	Understandability	Document structure (depth, breadth)
		Readability index
	Volatility	Count of changes / Count of requirements
		Lifecycle stage when change is made
	Traceability	Number of software requirements not traced to system requirements
		Number of software requirements not traced to code and tests
DW Process Consistency	Consistency	Number of contradictions to the conceptual schema of data stores
		Number of contradictions to enterprise model (concepts + constraints)

Who: Designer	When: System Design	Output: Activity model
DW Design Quality	Completeness	Number of modules not specified in terms of data
		Number of modules without care for trace management
		Number of modules without care for error reporting
	Consistency	Number of modules outside a flow sequence
		Number of contradictions to the logical schema of data stores
		Number of contradictions to the physical schema of data stores

Who: Programmer	When: Testing	Output: Software cases
Testing Effectiveness	Correctness	Reason for error
		Time of finding errors
		Time of error fixes & cumulative error curve
		Code location of fault
		Criticality of error
Traceability	Traceability	Number of empty cells in traceability matrix

Who: Programmer	When: Implementation	Output: Software module (agent)
Functionality		
Suitability	Reputation (for purchased SW)	Score in benchmark
		Market share
		Stock price
		Comments from other users
	Traceability	Number of empty cells in traceability matrix of requirements, roles, activities, tests, SW modules
		Number of enterprise rules & constraints tracked in cleaning software
Accuracy		Data Quality (see below)
	Data Completeness	Numbers of loaded records / number of expected records
		Number of null values in a column / number of rows
		Selectivity of extracted rows
		Percentage of involved columns
	Data Accuracy	Performance of statistic check: number of rows contradictory to real world values
	Data Consistency	Number of rows with reference violations / number of rows
		Number of duplicate rows / number of rows
		Number of rows violating business rules / number of rows
		Number of groups of tables with inconsistent granularity for the same information
		Number of rows with different codes for the same object / number of rows
		Number of rows with same codes for different objects / number of rows
		Number of rows having unknown codes / number of rows
	Interoperability	Number of modules unable to interact with specified systems.
	Compliance	Number of modules not compliant with application related standards or conventions or regulations in laws and similar prescriptions.
	Security	Number of modules unable to prevent unauthorized access, whether accidental or deliberate, to programs and data.
Reliability	Maturity	Number of error alerts over time period
	Fault tolerance	Number of occasions where the software was unable to maintain a specified level of performance (in cases of software faults or of infringement of its specified interface).
	Recoverability	Number of occasions where the software was unable to re-establish its level of performance and recover the data affected in the case of a failure. he time and effort needed for this reestablishment.
Usability	Understandability	Percentage of acceptance by the users.
	Learnability	Percentage of acceptance by the users.
	Operability	Existence of time prediction for operation
		Number of error alerts over time period
		Number of built-in transformations in extraction and transformation SW
		Number of constructed transformations in extraction and transformation SW
Software Efficiency	Time behavior	Throughput: maximum (average) propagation time / available time window
		Availability window of each source
		Estimated response time of extraction for each source
		Extraction frequency of each source
	Resource Behavior	History duration for each DW store
		Volume of data extracted and integrated
		Max data size the software can handle / required data size

Maintainability	Analyzability	Comment Percentage
		Readability Index
	Changeability, Stability, Testability	Number of logical paths in a module
		Number of groups of code having more than one entry or exit point
		Number of paths + number of control variables
		Total number of operators and operands / discrete number of operators and operands
		Number of control flow intersections
		Number of lines of dead code
		Number of data items that are not referenced
		Number of GOTO commands
		Number of entry points for a module (having more than one entry points)
		Number of exit points for a module (having more than one exit points)
		Number of recursions
		Size
		Correlation of complexity, size
Portability		Same as Maintainability

Who: Manager	When: Management	Output: Observation of the whole process
Implementation Effectiveness	Resource Usage	Staff hours spent on life cycle activities
	Completion Rates	Tasks completions
		Planned task completions

Who: Administrator	When: DW and Source Administration	Output: Administration activities
Administration Operability	DW Availability	Same as SW efficiency
	Source Availability	Same as SW efficiency
	Error Reporting	Same as SW Reliability

References

- [1] M. Jarke, M.A. Jeusfeld, C. Quix, P. Vassiliadis, Architecture and quality in data warehouses: an extended repository approach, *Inform. Systems* 24 (3) (1999) 229–253 (A previous version appeared in the Proceedings of the 10th Conference of Advanced Information Systems Engineering (CAiSE'98), Pisa, Italy (1998)).
- [2] M.A. Jeusfeld, C. Quix, M. Jarke, Design and analysis of quality information for data warehouses, Proceedings of the 17th International Conference on Conceptual Modeling (ER'98), Singapore, 1998, pp. 349–362.
- [3] P. Vassiliadis, M. Bouzeghoub, C. Quix, Towards quality-oriented data warehouse usage and evolution. *Inform. System* 25 (2) (2000) 89–115 (A previous version appeared in the Proceedings of the 11th Conference of Advanced Information Systems Engineering (CAiSE'99), Heidelberg, Germany, 1999, pp. 164–179).
- [4] E. Yu, J. Mylopoulos, Understanding 'why' in software process modelling, analysis and design, Proceedings of the 16th International Conference on Software Engineering (ICSE), Sorrento, Italy, 1994, pp. 159–168.
- [5] C. Shilakes, J. Tylman, Enterprise Information Portals, Enterprise Software Team, 1998, Available at <http://www.sagemaker.com/company/downloads/eip/in-depth.pdf>.
- [6] M. Demarest, The politics of data warehousing, 1997, <http://www.hevanet.com/demarest/marc/dwpol.html>.
- [7] Workflow Management Coalition, Interface I: process definition interchange process model, Document Number WfMC TC-1016-P, 1998, Available at www.wfmc.org.
- [8] M. Jarke, R. Gellersdörfer, M.A. Jeusfeld, M. Staudt, S. Eherer, ConceptBase—a deductive objectbase for meta data management, *J. Intell. Information Systems (Special Issue on Advances in Deductive Object-Oriented Databases)* 4 (2) (1995) 167–192.
- [9] R.Y. Wang, V.C. Storey, C.P. Firth, A framework for analysis of data quality research, *IEEE Trans., Knowledge Data Eng.* 7 (4) (1995) 623–640.
- [10] M. Jarke, M. Lenzerini, Y. Vassiliou, P. Vassiliadis (Eds.), *Fundamentals of Data Warehouses*, Springer, Berlin, 2000.
- [11] M. Oivo, V. Basili, Representing software engineering models: the TAME goal-oriented approach, *IEEE Trans. Software Eng.* 18 (10) (1992) 886–898.

- [12] C. Quix, M. Jarke, M. Jeusfeld, P. Vassiliadis, M. Lenzerini, D. Calvanese, M. Bouzeghoub, Data warehouse architecture and quality model, DWQ Technical Report DWQ-RWTH-002, 1997, Available at <http://www.dbnet.ece.ntua.gr/~dwq>.
- [13] P. Vassiliadis, Data warehouse modeling and quality issues, Ph.D. Thesis, 2000.
- [14] M. Bouzeghoub, F. Fabret, M. Matulovic, Modeling data warehouse refreshment process as a workflow application, Proceedings of the Workshop on Design and Management of Data Warehouses (DMDW'99), Heidelberg, Germany, 1999.
- [15] E. Schäfer, 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 (VLDB), Cairo, Egypt, 2000, pp. 715–718.
- [16] M. Jarke, M.A. Jeusfeld, T. Rose, A software process data model for knowledge engineering in information systems. *Inform. Systems* 15 (1) (1990) 85–116.
- [17] P. Vassiliadis, Gulliver in the land of data warehousing: practical experiences and observations of a researcher, Proceedings of the Second International Workshop on Design and Management of Data Warehouses (DMDW), Stockholm, Sweden, 2000, pp. 12.1–12.16.
- [18] J. Mylopoulos, A. Borgida, M. Jarke, M. Koubarakis, Telos—a language for representing knowledge about information systems, *ACM Trans. Inform. Systems* 8 (4) (1990) 325–362.
- [19] F. Baader, U. Sattler, Description logics with concrete domains and aggregation, Proceedings of the 13th European Conf. on Artificial Intelligence (ECAI-98), Brighton, UK, 1998, pp. 336–340.
- [20] D. Srivastava, S. Dar, H.V. Jagadish, A.Y. Levy, Answering queries with aggregation using views, Proceedings of the 22nd International Conference on Very Large Data Bases (VLDB), Bombay, India, 1996, pp. 318–329.
- [21] A. Gupta, V. Harinarayan, Dallen Quass: aggregate-query processing in data warehousing environments, Proceedings of the 21st Conference on Very Large Data Bases (VLDB), Zurich, Switzerland, 1995, pp. 358–369.
- [22] W. Nutt, Y. Sagiv, S. Shurin, Deciding equivalences among aggregate queries, Proceedings of the 17th ACM SIGACT-SIGMOD-SIGART Symposium, Principles of Database Systems (PODS'98), Seattle, 1998, pp. 214–223.
- [23] S. Chaudhuri, S. Krishnamurthy, S. Potamianos, K. Shim, Optimizing queries with materialized views, Proceedings of the 11th International Conference on Data Engineering (ICDE), Taipei, 1995, pp. 190–200.
- [24] W. Sadiq, M. Orłowska, Applying graph reduction techniques for identifying structural conflicts in process models, Proceedings of the 11th Conference on Advanced Information Systems Engineering, (CAiSE'99), Heidelberg, Germany, 1999, pp. 195–209.
- [25] O. Marjanovic, M. Orłowska, On modeling and verification of temporal constraints in production workflows, *Knowledge Inform. Systems* 1 (2) (1999) 157–192.
- [26] ISO, International Organization for Standardization, ISO/IEC 9126:1991 Information Technology—Software product evaluation—Quality characteristics and guidelines for their use, 1991.
- [27] P. Vassiliadis, Modeling multidimensional databases, cubes and cube operations, Proceedings of the 10th International Conference on Statistical and Scientific Database Management (SSDBM), Capri, Italy, 1998, pp. 53–62.
- [28] W. Kim, On optimizing an SQL-like nested query, *ACM Trans. Database Systems* 7 (3) (1982) 443–469.
- [29] R.A. Ganski, H.K.T. Wong, Optimization of nested SQL queries revisited, Proceedings of ACM SIGMOD International Conference on the Management of Data, San Francisco, CA, 1987, pp. 23–33.
- [30] M. Muralikrishna, Optimization and dataflow algorithms for nested tree queries, Proceedings of the 15th International Conference on Very Large Data Bases (VLDB), Amsterdam, The Netherlands, 1989, pp. 77–85.
- [31] M. Muralikrishna, Improved unnesting algorithms for join aggregate SQL queries, Proceedings of the 18th International Conference on Very Large Data Bases (VLDB), Vancouver, Canada, 1992, pp. 91–102.
- [32] U. Dayal, Of nests and trees: a unified approach to processing queries that contain nested subqueries, aggregates and quantifiers, Proceedings of the 13th International Conference on Very Large Data Bases (VLDB), Brighton, UK, 1987, pp. 197–208.
- [33] S. Chaudhuri, K. Shim, Optimizing queries with aggregate views, Proceedings of the Fifth International Conference on Extending Database Technology (EDBT), Avignon, France 1996, pp. 167–182.
- [34] I. Mumick, S. Finkelstein, H. Pirahesh, R. Ramakrishnan, Magic is relevant, Proceedings of the ACM SIGMOD International Conference on the Management of Data, Atlantic City, NJ, 1990, pp. 247–258.
- [35] H. Pirahesh, J. Hellerstein, W. Hasan, Extensible/rule based query rewrite optimization in starburst, Proceedings of the ACM SIGMOD International Conference on the Management of Data, San Diego, CA, 1992, pp. 39–48.
- [36] A. Levy, I. Mumick, Y. Sagiv, Query optimization by predicate move-around, Proceedings of the 20th International Conference on Very Large Data Bases (VLDB), Chile, 1994, pp. 96–107.
- [37] Microsoft Corp, MS Data Transformation Services, Available at [ww.microsoft.com/sq](http://www.microsoft.com/sq).
- [38] A. Gupta, I. Memick, K. Ross, Adapting materialized views after redefinitions, Proceedings of the ACM SIGMOD International Conference on the Management of Data, pp. 211–222, San Jose, CA, 1995.
- [39] A.Y. Levy, A.O. Mendelzon, Y. Sagiv, D. Srivastava, Answering queries using views, Proceedings of the 14th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS), San Jose, CA, 1995, pp. 95–104.

- [40] Z. Bellahsène, Structural view maintenance in data warehousing systems, *Journées Bases de Données Avancées (BDA'98)*, Tunis, 1998.
- [41] MetaData Coalition, Open Information Model, version 1.0, 1999, Available at www.MDCinfo.com.
- [42] Ardent Software, DataStage Suite, Available at <http://www.ardentsoftware.com/>.
- [43] DataMirror Corporation, Transformation Server, Available at <http://www.datamirror.com>.
- [44] P.A. Bernstein, T. Bergstraesser, Meta-data support for data transformations using microsoft repository, *Bull. Tech. Committee Data Eng.* 22 (1) (1999) 9–14.
- [45] Evolutionary Technologies Intl., ETI*EXTRACT, Available at <http://www.eti.com/>.
- [46] F. Casati, S. Ceri, B. Pemicci, G. Pozzi, Conceptual modeling of workflows, *Proceedings of the 14th International Conference on Object-Oriented and Entity-Relationship Modelling (ER'95)*, Gold Coast, Australia, 1995, pp. 341–354.
- [47] F. Casati, M. Fugini, I. Mirbel, An environment for designing exceptions in workflows, *Inform. Systems* 24 (3) (1999) 255–273.
- [48] F. Casati, S. Ceri, B. Pernici, B. Pozzi, Workflow Evolution, *DKE* 24 (3) (1998) 211–238.
- [49] D. Kuo, M. Lawley, C. Liu, M. Orłowska, A general model for nested transactional workflows, *Proceedings of the International Workshop on Advanced Transaction Models and Architecture*, India, 1996.
- [50] 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.uniulm.de/dbis/veranstaltungen/Workshop-Infomatik99-Proceedings.pdf>.
- [51] R. Klamma, Readings in Workflow Management: Annotated and Linked Internet Bibliography, RWTH, Aachen, <http://sunsite.informatik.rwth-aachen.de/WFBib>.
- [52] C. Rolland, A comprehensive view of process engineering, *Proceedings of the 10th International Conference on Advanced Information Systems Engineering, (CAiSE'98)*, Pisa, Italy, 1998, pp. 1–25.
- [53] Y. Wand, R.Y. Wang, Anchoring data quality dimensions in ontological foundations, *Comm. ACM* 39 (11) (1996) 86–95.
- [54] R.N. Wang, H.B. Kon, S.E. Madnick, Data quality requirements analysis and modeling, *Proceedings of the 9th International Conference on Data Engineering*, IEEE Computer Society, Vienna, Austria, 1993, pp. 670–677.
- [55] R.Y. Wang, A product perspective on total data quality management, *Comm. ACM* 41 (2) (1998) 58–65.
- [56] G.K. Tayi, D.P. Ballou, Examining data quality, *Commun. ACM* 41 (2) (1998) 54–57.
- [57] J.A. McCall, P.K. Richards, G.F. Walters, Factors in software quality, Technical Report, Rome Air Development Center, 1978.
- [58] B. Boehm, Software Risk Management, IEEE Computer Society Press, Los Alamitos, CA, 1989.
- [59] L. Hyatt, L. Rosenberg, A software quality model and metrics for identifying project risks and assessing software quality, *Proceedings of the 8th Annual Software Technology Conference*, Utah, 1996.
- [60] D. Georgakopoulos, M. Rusinkiewicz, Workflow management: from business process automation to inter-organizational collaboration, *Tutorials 23rd International Conference on Very Large Data Bases (VLDB)*, Athens, Greece, 1997.
- [61] M. Jarke, K. Pohl, Information systems quality and quality information systems, in: Kendall, Lyytinen, DeGross (Eds.), *Proceedings of the IFIP 8.2 Working Conference on the Impact of Computer-Supported Technologies on Information Systems Development*, North Holland, Minneapolis, 1992, pp. 345–375.
- [62] H. Galhardas, D. Florescu, D. Shasha, E. Simon, Ajax: an extensible data cleaning tool, *Proceedings of the ACM SIGMOD International Conference on the Management of Data*, Dallas, TX, 2000, 590 pp.
- [63] B. Ramesh, V. Dhar, Supporting systems development by capturing deliberations during requirements engineering, *IEEE Trans. Software Eng.* 18 (6) (1992) 498–510.
- [64] D. Georgakopoulos, M. Hornick, A. Sheth, An overview of workflow management: from process modeling to workflow automation infrastructure, *Distributed Parallel Databases* 3 (1995) 119–153.
- [65] D. Calvanese, G. De Giacomo, M. Lenzerini, D. Nardi, R. Rosati, A principled approach to data integration and reconciliation in data warehousing, *Proceedings of the International Workshop on Design and Management of Data Warehouses (DMDW'99)*, Heidelberg, Germany, 1999.
- [66] The Data Warehouse Information Center, An (Informal) Taxonomy of Data Warehouse Data Errors, Available at <http://www.dwinfocenter.org/errors.html>.
- [67] ViaSoft, Visual recap quality measurements. Viasoft White Papers, 1996, Available at <http://www.viaoft.com/download/white.papers/>.