Business Metadata for the Data Warehouse*

Weaving Enterprise Goals and Multidimensional Models

Veronika Stefanov and Beate List
Women's Postgraduate College for Internet Technologies
Institute of Software Technology and Interactive Systems
Vienna University of Technology
{stefanov, list}@wit.tuwien.ac.at

Abstract

Enterprise organizations use Data Warehouses (DWHs) to analyze their performance. Performance is judged regarding the achievement of goals.

DWH data models are well established. There exist numerous domain-specific modeling approaches. Enterprises also often model their goals in terms of formal or semi-formal goal models.

The problem is that these two aspects - the Data Warehouse and the Enterprise Goals - are described separately and not related to each other. We identify a need for combining these two aspects. If their relationship is made explicit, it can be used to enhance the way users access and interpret data in the DWH.

To address this limitation, in this paper we introduce a weaving model between enterprise goals and DWH data. Thus we present a domain-specific application of model weaving to an aspect of enterprise computing. We describe metamodels for both aspects as well as the weaving links between them, which allows to show the aspects separately but also in combination. We furthermore illustrate how to use the weaving links to create business metadata. Business metadata can be used in the DWH to describe the business context and implications of the data to the users, but is usually not available in today's DWHs. We apply our approach to a sample situation, which is used as a running example in the paper.

1 Introduction

Data Warehouse (DWH) systems represent a single source of information for analyzing the status, the behavior and the development of an organization [17]. Analysts and decision makers analyze this information with regard to the goals and the strategy of the organization and use it to recognize warning signs or trends and to decide on future investments.

Today's DWHs provide their users with very powerful tools for analyzing large amounts of data, but supporting the actual data interpretation by decision makers, i.e. with business metadata, has not yet been a primary concern in research. Business metadata describes the business context of the data, its purpose, relevance, and potential use [26], thus supporting data interpretation.

The context information needed for meaningful data interpretation, e.g. about the goals of the enterprise and the metrics used to evaluate them, is usually available in enterprise models or other repositories within the organization. But the link between the organizational information and the DWH is not easily accessible to DWH users and remains mostly implicit knowledge. If knowledge about the business context is left to chance, data analysis is bound to miss important points, and large parts of the potential of the DWH are wasted.

Models for DWH data structures are well researched and established (see Sect. 2). There exist numerous domain-specific modeling approaches to describe the data structures on the conceptual, logical and physical level. Enterprise organizations using a DWH are aware of these models and rely on them to design and describe their DWH.

To define and analyze their objectives and goals, enterprise organizations use different types of goals models (see Sect. 3). For the goals, metrics are derived to measure the achievement of the goals. There exist many modeling approaches for enterprise goals as well as for metrics.

The problem that we address here in this paper is that these two aspects - the Data Warehouse and the Enterprise Goals - are described separately and not related to each other. We identify a need for combining these two aspects.

*This research has been funded by the Austrian Federal Ministry for Education, Science, and Culture, and the European Social Fund (ESF) under grant 31.963/46-VII/9/2002.

If their relationship is made explicit, it can be used to enhance the way users access and interpret data in the data warehouse.

The goal of this paper is to address this limitation by

- making the relationship between the DWH data and the goals of the organization visible and accessible by
- showing how the enterprise goals and metrics are mirrored in the DWH data model
- using this information to support and improve data interpretation
- enriching the DWH with business metadata that explains the relevance and business context of the data

To achieve these goals, we employ modeling techniques. We propose a weaving model (see Sect. 4) to link an enterprise goal model with the DWH data model. The weaving links then can be used to provide business metadata to the DWH users. Business metadata informs users about the organizational context and implications of what they are analyzing [26].

Thus we apply model weaving to the domain of data warehousing. Our approach offers the following contributions:

- It makes the implicit relationships between the data in the DWH and the business goals visible and accessible.
- By relating the measures in the DWH to the overall organizational strategy and enterprise goals, decision makers can better interpret the enterprise performance, and understand the implications.
- Creating the weaving links is a comparatively small investment for valuable metadata that gives meaning to DWH measures.
- The enterprise goals serve as a "single source of information" to avoid
- DWH requirements analysis and (re-)design are notoriously challenging tasks, because the business context of a DWH is difficult to extract from user interviews and practically impossible to store directly in the multidimensional data structures. Weaving an enterprise goal model with the data model makes context information accessible, and does so without disrupting the involved models.
- The weaving model can be used for model validation, as it identifies missing or superfluous tables and measures as well as omissions in the goal model.
- The link to enterprise goals and business strategy can help to evaluate DWH investments, and to justify them.

The following sections describe how the DWH data (Sect. 2) and the enterprise context (Sect. 3) are modeled. Section 4 introduces the weaving model that connects the two models, and illustrates the resulting business metadata with an example. Related work is discussed in Sect. 5.

2 Multidimensional Data Models

The main data model in Data Warehousing is the multidimensional model, also called star schema [8]. It is meant to provide intuitive and high performance data analysis [17].

DWH applications involve complex queries on large amounts of data, which are difficult to manage for human analysts. Relational data models "are a disaster for querying because they cannot be understood by users and they cannot be navigated usefully by DBMS software" [17]. In Data Warehousing, data is often organized according to the multidimensional paradigm, which allows data access in a way that comes more natural to human analysts. The data is located in n-dimensional space, with the dimensions representing the different ways the data can be viewed and sorted (e.g. according to time, store, customer, product, etc.).

A multidimensional model, also called star schema or fact schema, is basically a relational model in the shape of a star (see Fig. 1 for an example). At the center of the star there is the *fact* table. It contains data on the subject of analysis (e.g. sales, transactions, repairs, admissions, expenses, etc.). The attributes of the fact table (e.g. cost, revenue, amount, duration, etc.) are called *measures*. The spokes/points of the star represent the *dimensions* according to which the data will be analyzed. The dimensions can be organized in hierarchies that are useful for aggregating data (e.g. store, city, region, country). Stars can share dimensions, thus creating a web of interconnected schemas that makes drill-across operations possible.

There are many approaches to modeling the multidimensional data structures of data warehouses [1, 6, 30], some of which are object-oriented models or based on the Unified Modeling Language (UML) [1, 21, 28].

For weaving enterprise goals (as part of the enterprise model, see Sect. 3) with the DWH, we need a data model that supports weaving. We choose the object-oriented approach first presented in [28] and then further developed to a UML profile in [19]. A UML Profile is a domain-specific extension to the UML modeling language. This profile adapts the UML class diagram for multi-dimensional modeling, i.e. the base class of the stereotypes is Class. It allows to create detailed models of the conceptual characteristics of multidimensional data models. Figure 2 shows the main elements of the UML Profile and their relationships as a metamodel, and Fig. 1 shows an Expenses fact from modeled with the profile.

The Expenses fact has four dimensions: *Time*, *Account* (e.g., IT, Marketing, etc.), *Scenario* (e.g., Actual, Forcast) and *Store*. The levels of the dimensions are only shown for the store dimension. Each entry in the fact table contains information about a single expense incurred. In this example there is only one measure that can be analyzed for each expense: the *amount*. Aggregations such as "total value of

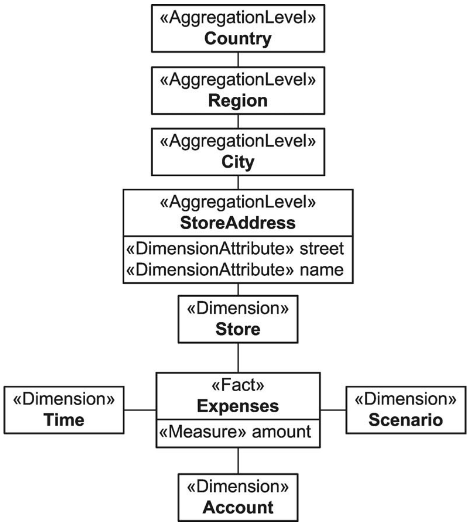

Figure 1. Example using the UML Profile for Multidimensional Data Models [19]. Aggregation levels are only shown for the store dimension.

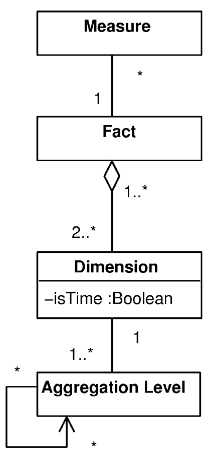

Figure 2. Core elements of the UML profile for modeling multidimensional data models, and their interrelationships (cf. [19])

an account in all stores in one year" become possible by selecting aggregation levels from the dimensions. Several such facts can be connected by sharing the same dimensions, creating a more complex multi-cube model.

The elements shown in Fig. 2 allow to model any number of Fact tables. Each *Fact* table can have any number of optional *Measures* and must have at least two *Dimensions*

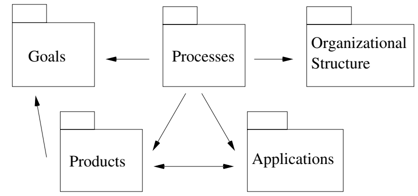

Figure 3. Example of a generic Enterprise Model

connected to it, at least one of which is usually a *Time* dimensions. Dimensions may be shared between facts and have one or more *Aggregation Levels*, which form the aggregation hierarchy.

See [19] for the additional, more detailed elements not shown here, such as the attributes of the dimensions, as well as a number of logical constraints on the model elements [19]. They are defined in the Object Constraint Language (OCL) [22] and cover aspects of multidimensional data models such as ensuring that the classification hierarchies have the form of directed acyclic graphs, as well as additivity rules, derived measures, and degenerate dimensions.

3 Enterprise Models

3.1 Modeling the structure, behavior and goals of an enterprise organization

An enterprise model formally represents the basic building blocks of an organization, its structure, behavior and goals. It is usually organized into several aspects [31] that can be modeled individually but also related to each other. The Architecture of Integrated Information Systems (ARIS) [27] is a typical example for such an enterprise model.

Figure 3 shows the outline of a generic enterprise model, organized into five aspects: The enterprise strives to achieve goals, acts through processes, has an organizational structure, produces products and uses applications. In the enterprise model, an organization chart can be used to describe the organizational structure, i.e. the dependencies between the departments, groups and roles that exist within the organization. Similarly, a business process model describes the structure of business processes with control flow, inputs and outputs. The products, applications, and strategic goals can also be modeled separately. An overview model would then connect these model to show for example how processes fulfill goals, are performed by organizational roles, fall into the responsibility of departments, and use applications to produce products for other departments.

3.2 Enterprise Goals

A core part of every enterprise model is the goal model. “Increase market share” or “reduce operating costs” are typical enterprise goals. Goals form the basis for business decisions and the way a company does business. What is relevant and important for business performance measurement can be read directly from the enterprise goal model. They govern the design of business processes and the way the organization behaves. Nevertheless, a goal model is basically very simple, and enterprise goals are long term goals that should remain stable a lot longer than business processes, role definitions, and operating rules. Therefore, they provide excellent metadata for a DWH.

Based on the description of the goals, the enterprise derives metrics that measure the level of achievement of the goals and indicate the performance of the enterprise. These metrics are not identical but closely related to the measures in the DWH. In the early 1990s, business goals and strategy experienced a revival in theory and practice with approaches like the Balanced Scorecard (BSC) [13]. The BSC’s focus is on vision and strategy. Companies define goals from various perspectives, which are then translated into measures. The BSC does not mention the behavior which will lead to the fulfillment of a goal. Rather, people are assumed to be guided by the measures they have to fulfill. Measures and not the desired operations are communicated to the employees. The goals and measures give the long-term focus, while the conditions in which people operate are constantly changing.

In the Goal Question Metric Approach [3], originally aimed at software quality improvement, measurement and evaluation is based on a three-level hierarchy in which the goals (of an organization) form the first, the conceptual level. Goals are the starting point of all measurement activities and provide the organizational context according to which the measurement data can be interpreted.

Different kinds of goals, including enterprise goals, are often used in software engineering for requirements elicitation. For example, the i* Methodology [32] provides an agent- and intention-oriented way to analyse requirements for software (and other) systems. The focus of i* is on interaction between autonomous agents, their actions and strategies.

For enterprise goals in particular, there is often a distinction between three levels of goals: strategical, tactical and operational. In order to be able to transform high level enterprise goals of the strategic level via tactical level goals to every-day operational goals, a goal is decomposed via a *causal transformation* or *operationalization* into one or more subgoals, which in turn can be further decomposed, thus creating a hierarchy (cf. [18]).

3.3 Enterprise Goal Metamodel (EGM)

The Enterprise Goal Metamodel presented here incorporates features from a number of existing goal modeling approaches (cf.[13–15, 18, 32]) It is aimed at providing a sufficiently detailed and comprehensive, yet concise model of the main concepts that are needed to model the context of a DWH.

Figure 4 shows the Enterprise Goal Metamodel (EGM). For sake of clarity and readability, it is shown as two separate graphics: Fig. 4(a) explains goal decomposition hierarchies and relationships between goals and Fig. 4(b) shows the other elements of the metamodel related to goals. The model uses the notation of the UML 2.0 class diagram [23].

Figure 5 shows example goals for Fig. 4(a). In the EGM, a *Goal* may participate in a goal hierarchy via a *Goal Decomposition*. The goal decomposition connects a higher-level goal with a number of lower-level subgoals. A goal may have only one set of subgoals but may participate itself as a subgoal in more than one goal hierarchy. Therefore it can be related to only one goal decomposition in the role of a *satisfied* goal but to many in the role of a *satisfier* goal. The goals “reduce out-of-stock” and “increase freshness” in Fig. 5 are subgoals of “satisfy customers”. From the viewpoint of the “AND” goal decomposition in Fig. 5, the four lower-level goals are satisfier goals, and “satisfy customers” is the satisfied goal. The type of a goal decomposition is either *AND* or *OR*, depending on whether all or only some of the subgoals have to be satisfied to satisfy the upper level goal.

Orthogonally to the goal hierarchy, goals can be seen to *influence* each other in various ways. The fulfillment of one goal might be detrimental to another goal, or the goals may be related to each other in such a way that if one of them is satisfied, this also supports the other goal. Therefore, there are two influencing relationships between goals: *support* and *conflict*. Both may occur between any number of goals, e.g. a goal can support several goals and conflict with others at the same time.

Figure 4(b) (illustrated with values in Table 1) shows that for each *Metric* assigned to a goal it is necessary to define a *Target Value*. Because target values usually change over time, a *Time Frame* for each combination of metric and target value is necessary. Goals can be *relative* goals (“increase the value by x”), or absolute goals (“the value should be x”), indicated by the attribute *isRelative* (shown in Fig. 4(a)). This influences the semantics of the timeframe: For a relative goal and its metric, it means that the change is to be achieved during this time, whereas for an absolute goal it indicates the validity period of the target value. A metric can be constrained with *Parameters*, which define the scope: The goal “reduce inventory cost” has none, “reduce inventory cost of top-20 products” has one and “reduce

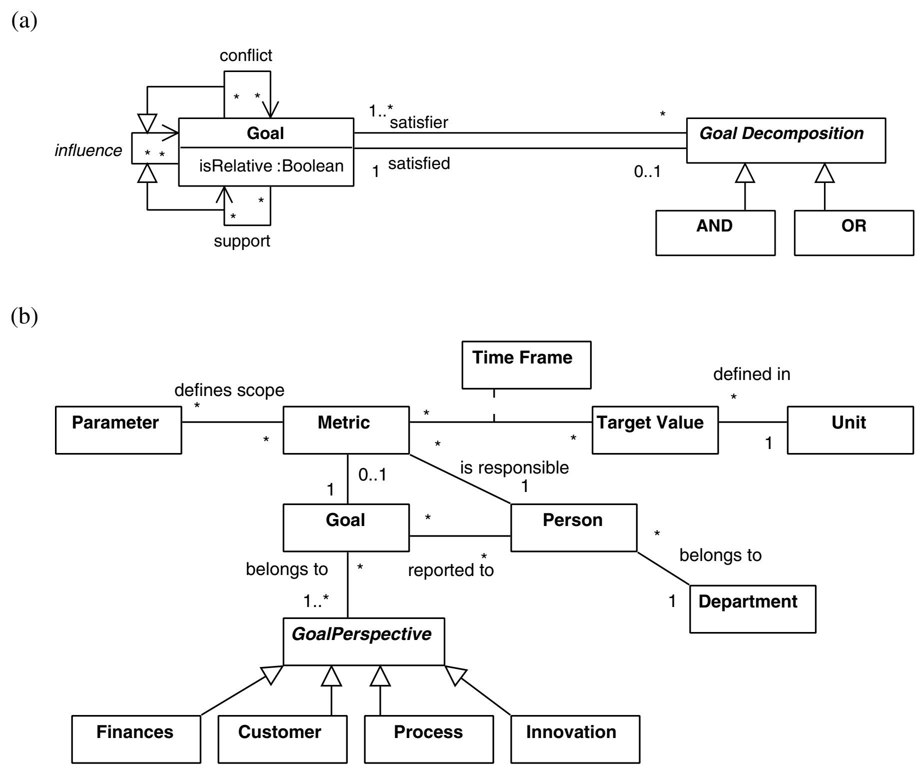

Figure 4. Metamodel for describing (a) goal decomposition hierarchies and relationships between goals and (b) the details related to enterprise goals

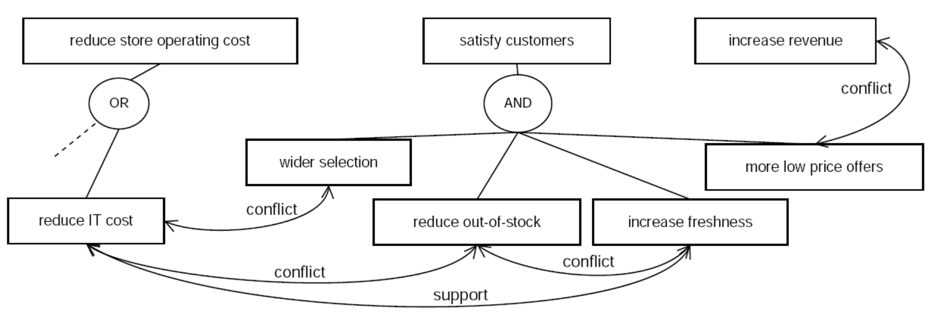

Figure 5. Sample goal hierarchy model corresponding to the metamodel in Fig. 4(a)

Table 1. Details regarding Fig. 4(b) for the example goals from Fig. 5

Name	Value
Goal name	**reduce "out of stock"**
Goal perspective	Customer
Reported to + contact info	Sales Dept., Ms. Baker, Tel. 5800193
Metric name	number of days in month with no items on stock
Metric target value + unit	$< 1\%$
Metric responsible + contact info	Mr. Jones, Tel. 5800234
Conflicting/supporting goals	conflicts with "increase freshness" and "reduce IT cost"
Goal this goal satisfies	"satisfy customers
Goal name	**increase freshness**
Goal perspective	Customer
Reported to + contact info	Mr. Groop, groop@dpt.company.com, Ms. Fitt, Tel. 5800156
Metric	avg. time in warehouse for product group "fresh goods"
Metric target value + unit	< 8 hours
Metric parameter(s)	Warehouse, product group
Metric responsible + contact info	Mr. Stephens, Tel. 5800655
Conflicting/supporting goals	conflicts with "reduce out of stock", supports "reduce IT cost"
Goal this goal satisfies	"satisfy customers"

inventory cost of top-20 products in region x" has two parameters. Goals can be *reported to* a *Person* belonging to a *Department*. For each metric there has to be a *responsible* Person. To indicate the general focus of the goals, they are assigned to *Perspectives*. These perspectives are generic and can be adapted to the analysis needs of the company. Figure 4(b) shows four perspectives according to the Balanced Scorecard. Person and Deptartment are part of the organizational aspect.

4 The Weaving Model

In order to gain business metadata, we need to link the enterprise goals to the DWH data. For this task, we have chosen to employ the technique of model weaving [5]: An additional model is created between the two metamodels that are to be linked. This so-called weaving model consists of typed weaving links that describe the relationships between two or more model elements from the respective metamodels. Advanced modeling tools such as the ATLAS Model Weaver [2, 9] (available as an Eclipse [11] plug-in) support model weaving. The models concerned have to be based on a common high-level modeling approach, e.g., be based on MOF [24] as their common meta-meta-model.

Model weaving superficially resembles techniques used in ETL or EAI, but the intention behind it is different. A weaving link does not necessarily imply that the two elements it connects should in some way be transformed from one into the other. Rather, it simply indicates that the two elements share some semantic link, e.g. "lies in the responsibility of", "is measured by", "affects", etc. Still, weaving can of course be used as a preliminary step to transformation, by indicating transformation sources and targets and then using the weaving model as an input for a transformation language. In this paper it is employed for annotating the DWH data with business metadata, and therefore does not imply transformation.

We choose model weaving because it offers a number of advantages: By adhering the "Everything is a model" principle [4], we can capture practically *all* information in terms of models, also the relationships and correspondences between models. This makes it possible to store, analyze and edit the links with modeling tools. Weaving avoids having one large meta-model "for everything", but instead keeps the individual meta-models separate, easy to handle and focused on their domains, while at the same time they are interconnected into a "lattice of metamodels" [4].

Figure 6 shows the weaving model linking the data (right) and enterprise goal (left) metamodels presented in Sect. 2 and 3. It consists of three links: two binary and one ternary link.

The first link is between the Parameter describing the focus or scope of the metric (e.g., Region is the parameter when a value is given by *by region*) in the EGM and an Aggregation Level (e.g. per Month on the Time Dimension or per Region on the Store dimension). These are similar concepts, which can be easily mapped. The corresponding dimension to a level is provided by the data model.

Weaving links can connect more than just one element on either side. The second, most complicated link in this weaving model connects a Metric with a Measure and optional Aggregation Levels. A metric roughly corresponds to

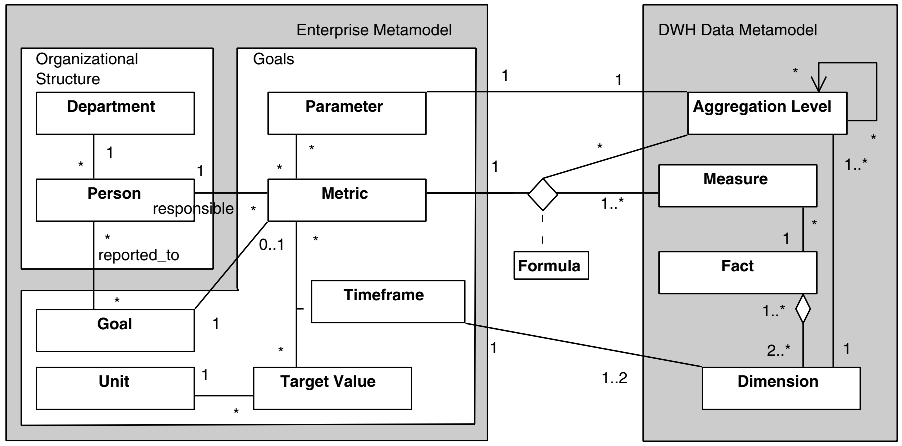

Figure 6. Weaving Model connecting the Goal Model (left) with the Data Model (right). Only a part of the Enterprise Model is shown here for readability. It has more than the two aspects, and the organizational aspects has more than the two elements shown here.

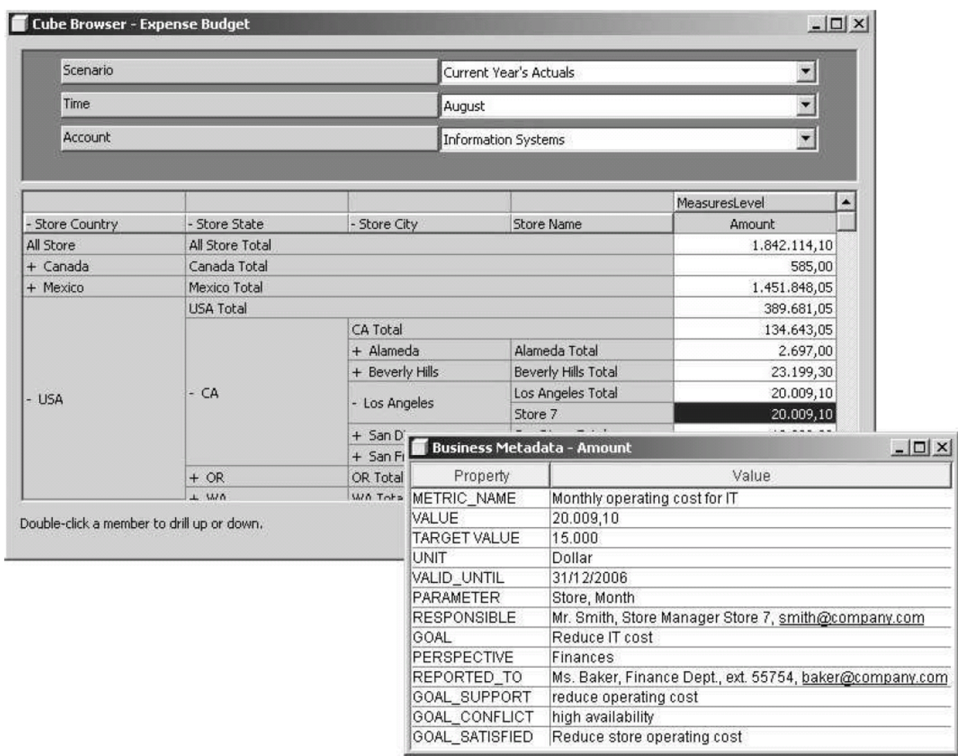

Figure 7. Displaying business metadata connected to fact data

a fact attribute. The fact attribute itself is not aggregated, but the metric can be restricted by parameters: This has to be indicated on the data model side by adding the corresponding aggregation base(s) to the link. Additionally, because the fact attribute itself contains only the absolute value of a variable, while the metric related to it might contain an average, a percentage, or a rate of change, this weaving link can contain a formula (e.g., *(IT cost/total cost)*100* for the *percentage of IT costs*).

Finally, the third link allows us to handle the relationships between the Timeframe of a metric's target value and a Dimension containing time values in the data model[1]. A timeframe is a time period, indicated by start and end point, whereas a time dimension contains single points in time. Therefore the weaving link connects one timeframe with two points on the time dimension. Or if the timeframe has the format of "until x", with one point.

Analysis tools and applications can use the business metadata derived through the weaving links, similarly to technical metadata, to enhance the way users access and interpret data. Where before there were basically only numbers, there now is context and explanation. Through knowing which goal it measures, it becomes clearer what a certain value means and why it is important. The actual values of the measures can be compared to the target values of the metrics. This business metadata can be also incorporated directly into the user interface of analysis tools.

Figure 7 shows how the business metadata can be displayed for an example cube, based on the prototype we are developing. The organizational knowledge captured in the enterprise goal model becomes available to the user. Providing this information to the user directly within the analysis tool helps to improve data interpretation. The business metadata thus increases the usefulness of the data.

In the example in Fig. 7, business metadata is derived from the link between metrics and measures. Combined with the two other links (concerning dimensional aggregation and temporal values), the links can also provide insights for DWH (re)design and maintenance, or requirements analysis. The knowledge captured by the weaving model can be exploited by analysis tools (e.g., to offer better navigation, or hints).

5 Related Work

The term "weaving" is also used in aspect-oriented programming, where it denotes the integration of aspects into the base program [16]. See the AOSD Ontology [29] for more general definitions that apply not only to the programming level, but also to modeling.

[1] This can be ensured by an OCL [22] constraint, e.g.
```
Context TimeFrame
inv:  self.dimensions-> forAll(d|d.isTime = true)
```

In [7], Breton and Bézivin apply model weaving to the area of workflow and process modeling. The build-time and the run-time workflow definitions are weaved together to create a binding between definition and execution of the process.

The alignment of models in the Data Warehousing research field is quite young, although there are some very good examples available.

Mazon et al. defined in [20] the application of the MDA framework to DWH repository development, and aligned multidimensional conceptual data models with code. In MDA terms, they aligned a Platform Independent Model (PIM) with a Platform Specific Model (PSM) and defined a transformation between them. Starting with a conceptual model, Mazon et. al developed a logical model using a relational approach to build a star schema. Then, they derive the necessary code to create data structures for the DWH. Our approach can be seen on top of this work targeting the Computation Independent Level (CIM) level, as we align enterprise goals, representing the business requirements as well as context, with the DWH conceptual data model.

Giorgini et al. focus on DWH requirement analysis based on goals in [12]. They derive the data model from the goals, which represent a rather narrow software engineering type of goals. In contrast, we integrate enterprise goals and align the DWH directly with business strategy.

Sarda linked DWH business metadata with technical metadata in [26], in order to provide a better context for decision support. Several business metadata categories like goals, organizational elements, processes, events, measures, etc., and a number of desirable characteristics such as evolution of navigation between metadata and data, are defined. The business metadata is described with UML classes and associations and then linked directly to the technical metadata within the same model. The approach only covers metadata and does not use separate conceptual models of the business context. Additionally, our weaving model is focused on the details of enterprise goals and their measures, rather than on all aspects of an enterprise.

6 Conclusion and Future Work

In this paper we have presented an approach to business metadata that is based on the relationship between the DWH data and the goals of an organization. The enterprise goals and information related to them such as metrics and target values as well as the people and departments involved, are taken from an enterprise model. The business metadata is created by linking this knowledge about the organization to the DWH by means of a weaving model.

The business metadata is then created directly from the weaving model. It improves data interpretation by explaining the relevance and context of the data, whereas the weaving model itself supports DWH requirements analysis, (re)design and evolution by making context visible and accessible. The approach is applied to an example.

Regarding future work, we are currently working on a prototype of a toolkit that supports users in choosing and integrating the enterprise models available to them, through creating the weaving model that links them to the DWH, to finally creating and displaying the business metadata. We are also investigating the use of other aspects in the enterprise model (apart from the enterprise goals) as suitable business metadata for the DWH. The prototype is based on the Eclipse [11] platform. It will make use of the open-source DWH platform Pentaho [25], which is built on Eclipse and combines well-known DWH components such as Mondrian, BIRT, and JFreeReport, on the one hand, and on the otherhand use the the Eclipse Modeling Framework (EMF) [10] and the ATLAS Model Weaver [2]. It forms the basis for future case studies.

References

[1] A. Abelló, J. Samos, and F. Saltor. YAM^2 (Yet Another Multidimensional Model): An Extension of UML. In *IDEAS '02*, pages 172–181. IEEE Computer Society, 2002.

[2] AWM. ATLAS Model Weaver. `http://www.eclipse.org/gmt/amw/` (16.06.2006), 2006.

[3] V. R. Basili, G. Caldiera, and H. D. Rombach. *The Goal Question Metric Appraoch*, volume Encyclopedia of Software Engineering. Wiley, 1994.

[4] J. Bézivin. On the unification power of models. *Software and System Modeling*, 4(2):171–188, 2005.

[5] J. Bézivin, F. Jouault, P. Rosenthal, and P. Valduriez. Modeling in the Large and Modeling in the Small. In *Model Driven Architecture, LNCS 3599*, pages 33–46, 2005.

[6] M. Blaschka, C. Sapia, G. Höfling, and B. Dinter. Finding Your Way through Multidimensional Data Models. In *DEXA '98*, pages 198–203, 1998.

[7] E. Breton and J. Bézivin. Weaving Definition and Execution Aspects of Process Meta-models. In *35th Hawaii Int. Conf. on System Sciences*, page 290, 2002.

[8] S. Chaudhuri and U. Dayal. An Overview of Data Warehousing and OLAP Technology. *SIGMOD Rec.*, 26(1):65–74, 1997.

[9] M. D. del Fabro, J. Bézivin, F. Jouault, E. Breton, and G. Gueltas. AMW: A Generic Model Weaver. In *IDM'05 Ingénierie Dirigée par les Modèles*, 2005.

[10] Eclipse Foundation. Eclipse Modeling Framework. `http://www.eclipse.org/emf/` (16.06.2006), 2006.

[11] Eclipse Foundation. Eclipse Platform. `http://www.eclipse.org/` (16.06.2006), 2006.

[12] P. Giorgini, S. Rizzi, and M. Garzetti. Goal-oriented requirement analysis for data warehouse design. In *Proceedings DOLAP '05*, pages 47–56. ACM Press, 2005.

[13] R. S. Kaplan and D. P. Norton. The Balanced Scorecard - Measures that Drive Performance. *Harvard Business Review*, 1992.

[14] E. Kavakli. Modeling organizational goals: Analysis of current methods. In *ACM SAC 2004*, pages 1339–1343, 2004.

[15] E. Kavakli and P. Loucopoulos. *Information Modeling Methods and Methodologies*, pages 102 – 124. IDEA Group, 2004.

[16] G. Kiczales, J. Lamping, A. Mendhekar, C. Maeda, C. V. Lopes, J.-M. Loingtier, and J. Irwin. Aspect-Oriented Programming. In *Proceedings of the 11th European Conference on Object-Oriented Programming (ECOOP 1997), LNCS 1241*, pages 220–242. Springer-Verlag, 1997.

[17] R. Kimball, L. Reeves, M. Ross, and W. Thornthwaite. *The Data Warehouse Lifecycle Toolkit*. John Wiley & Sons, Inc., 1998.

[18] P. Loucopoulos and V. Kavakli. Enterprise Knowledge Management and Conceptual Modelling. In *Conceptual Modeling (ER'97)*, pages 123–143. Springer, 1997.

[19] S. Luján-Mora, J. Trujillo, and I.-Y. Song. Extending the UML for Multidimensional Modeling. In *UML '02*, pages 290–304. Springer-Verlag, 2002.

[20] J.-N. Mazon, J. Trujillo, M. Serrano, and M. Piattini. Applying MDA to the development of data warehouses. In *Proceedings DOLAP '05*, pages 57–66, New York, NY, USA, 2005. ACM Press.

[21] T. B. Nguyen, A. M. Tjoa, and R. Wagner. An Object Oriented Multidimensional Data Model for OLAP. In *Web-Age Information Management (WAIM) 2000)*, pages 69–82. Springer-Verlag, 2000.

[22] Object Management Group, Inc. UML 2.0 Object Constraint Language (OCL) Specification. `http://www.omg.org/cgi-bin/apps/doc?ptc/05-06-06.pdf`, 2005.

[23] Object Management Group, Inc. UML 2.0 Superstructure. `http://www.omg.org/cgi-bin/apps/doc?formal/05-07-04.pdf`, 2005.

[24] Object Management Group (OMG). Meta Object Facility (MOF) 2.0 Core Specification. `http://www.omg.org/cgi-bin/apps/doc?ptc/03-10-04.pdf`, 10 2003.

[25] Pentaho Project. Pentaho Open-Source Business Intelligence. `http://www.pentaho.org` (16.06.2006), 2006.

[26] N. L. Sarda. Structuring Business Metadata in Data Warehouse Systems for Effective Business Support. *CoRR, http://arxiv.org/abs/cs.DB/0110020*, 2001.

[27] A.-W. Scheer. *ARIS - Business Process Modeling*. Springer Verlag, 1999.

[28] J. Trujillo, M. Palomar, J. Gómez, and I.-Y. Song. Designing Data Warehouses with OO Conceptual Models. *IEEE Computer*, 34(12):66–75, 2001.

[29] K. van den Berg, J. M. Conejero, and R. Chitchyan. AOSD Ontology 1.0 - Public Ontology of Aspect-Orientation. AOSD-Europe-UT-01 D9, AOSD-Europe, May 2005.

[30] P. Vassiliadis and T. K. Sellis. A Survey of Logical Models for OLAP Databases. *SIGMOD Record*, 28(4):64–69, 1999.

[31] L. Whitman, K. Ramachandran, and V. Ketkar. A taxonomy of a living model of the enterprise. In *WSC '01*, pages 848–855. IEEE Computer Society, 2001.

[32] E. S. Yu. Towards modelling and reasoning support for early-phase requirements engineering. In *Symposium On Requirements Engineering (RE '97)*, pages 226–235. IEEE Computer Society Press, 1997.