

Domain-Driven Development with Ontologies and Aspects

*Submitted for Domain-Specific Modeling workshop at OOPSLA 2005
Latest version of this paper can be downloaded from <http://phrubby.com>*

Pavel Hruby
Microsoft
Frydenlunds Allé 6
2950 Vedbaek, Denmark
+45 29229183
phrubby@acm.org

ABSTRACT

Common task during developing software applications in specific domains is to identify application objects and their behavior. In this paper, we illustrate that the application objects can be derived from ontological categories for the domain, and the behavior can be modeled as aspects that cross-cut the application objects. We further describe that categories of aspects can be identified for the domain, and the software applications can be developed by merely configuring the object and aspects categories.

1. INTRODUCTION

In developing domain-specific model, the designer has option of using knowledge about the domain, in addition to user requirements and principles of object-oriented design. We will illustrate that using formally specified domain knowledge in software design leads to a specific structure of the model. The model addresses two orthogonal concerns; the general concepts and categories originating in the domain knowledge and the specific behavior originating in user requirements.

In this document, domain-driven development is understood to be the development of software applications in the scope of a specific domain or an application area. Examples of the domains are; the business domain, interactive systems domain and sales domain, see Figure 2. We call the entities of the application model application objects. Examples of application objects in the business domain are the sales order, invoice and shipment. The application object categories are called the metaclasses of the application objects, the entities describing characteristics of the application objects. Examples of application object categories in business domain are economic resource, event, agent, claim, contract and commitment. Instances are the runtime manifestations of application objects.

Figure 1 illustrates the application object Cash, which has a metaclass Economic Resource (an application object category). At runtime, the instance of Cash is an object that represents, for example, the real cash contained in a wallet.

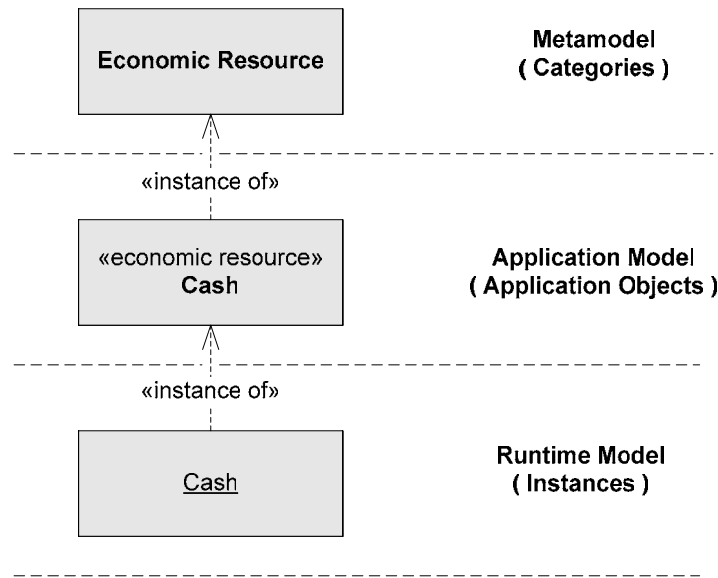


Figure 1. Category, Application Object and Instance

The term aspect is used to denote a concept for a module of functionality that cross-cuts application objects. There are several approaches for modeling and implementing the cross-cutting concerns, such as aspect-oriented programming, composition filters, multidimensional separation of concerns and UML collaborations. In this paper we do not discuss any of these approaches specifically, and we stay at the conceptual level.

The rest of the paper is structured as follows: the section on application objects describes the entities derived from domain ontology. The section on application behavior describes those concerns not covered by the ontology and modeled as aspects. The section on domain-driven development summarizes the central thinking around this paper and illustrates an example of an application configured using ontologies and aspects. The questions and answers section provides feedback to the reviewers of this paper.

2. APPLICATION OBJECTS

2.1 Level of Specificity of Modeling Entities

Figure 2 illustrates application objects at various levels of specificity; the horizontal dimension illustrates the domains the application objects can be applied to. The horizontal dimension in Figure 2 is not ordered; the intervals illustrate sizes and overlaps of different domains. General modeling languages such as UML, describing the model in terms of objects, classes and methods, contain very few semantics about real world concepts. On the other hand, their modeling scope is a domain of all object-oriented systems. Specific modeling languages that use entities such as invoice and shipment contain detailed and exact semantics and information about intentions of the model. This information can be used, for example, to validate the model against domain rules and automatically translate this model to the models in other domains. The trade-off is that the area of applicability of a specific language is more restricted than the area of applicability of a general modeling language.

In the ideal case the modeling entities should be exactly at the level of specificity that entirely covers the domain of the software application. We will show in the next section that the domain objects at the right level of specificity correspond to ontological categories for the domain. For example, for the business domain, such objects correspond to categories specified by the REA (Resources, Events Agents) ontology [4, 5].

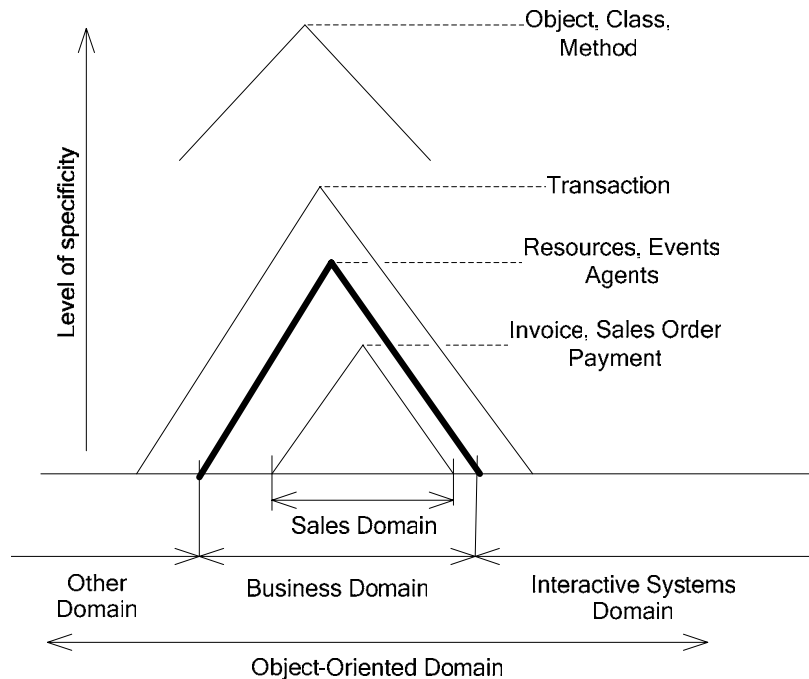


Figure 2. Levels of Specificity and Covered Domain

2.2 Ontology and Domain-Driven Development

“An ontology is an explicit specification of a conceptualization” [7]. Ontological categories define the concepts that exist in the domain, as well as relationships between these concepts. For object-oriented applications, domain ontology defines the metamodel for application models in this domain. Ontological categories correspond to the metaclasses for application objects. For example, the economic agent specified by the REA ontology is a metaclass for objects such as customer and vendor. Economic event is a metaclass for objects such as sale and payment receipt. Economic resource is a metaclass for objects such as money and item. Likewise, relationships between ontological categories become metarelations for the relationships between application object, such as stockflow is a metarelationship for relationships called outflow and inflow; duality between economic events is a metarelationship for the reconciliation relationship between Sale and Payment Receipt.

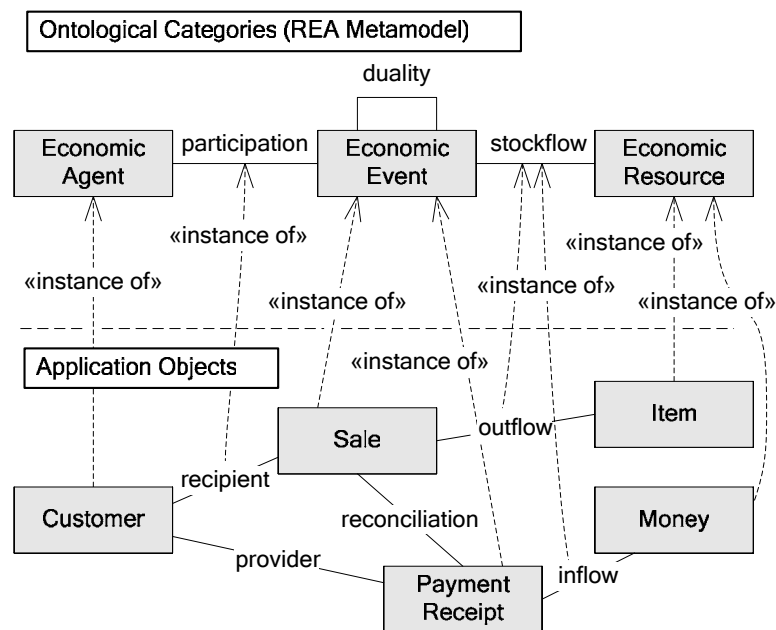


Figure 3. Application Model and its REA Metamodel

3. APPLICATION BEHAVIOR

3.1 Functionality of Domain Objects

The previous section illustrated that structure of a software application can be derived from ontological categories that apply to the application's domain. However, to build a useful software application, the mere structure of domain objects is not sufficient. Domain objects need functionality that is often not specified by the ontological categories, but is required by the application's users. For example, the REA ontology does not specify how to determine the identity of business objects, or how to create financial reports. However, functionality such as serial numbers and accounts are essential for the users of business applications.

Application functionality is not specified by domain ontologies for a good reason. Domain ontologies specify the structure of concepts that can be applied to all systems in the domain. Domain ontologies attempt to find the minimal, yet complete set of concepts covering the domain.

The functionality of application objects usually differs from one system to another because of specific user requirements, local conventions, as well as several other reasons. For example in business applications, some objects need human readable serial numbers, such as customers and products; some do not, such as order lines. Financial reporting depends on local legislation, lines of business and reporting usually varies from one application to another, reflecting the fact that every company is somehow different than the other. A complete list of functionality of the domain objects probably cannot be specified in general for the whole domain; users of software applications would always need new features or new versions of existing features, which cannot be foreseen by those who create the ontology.

3.2 Cross-Cutting Domain Objects

Ontological categories determine one dimension of decomposition of the domain. The other dimension of decomposition is the application functionality. In the following paragraph we will show that in many cases the modules of application functionality are not localizable into a single application object.

In the business domain, for example, the serial number of an item is an attribute of the item object. The serial number is usually not a random number. The item **serial number** is determined by a **serial number setup**, which is an attribute of a group of the economic resources, to which the number series is applied. Thus, the object representing the item group contains rules specifying things such as the format of the serial number, whether a serial number should be unique, how does it depend on previous numbers of the series, rules determining whether serial numbers of deleted items can be reused, and other similar rules. The **number series module** cross-cuts two domain objects, the item object and the item group object, and the number is constructed by mutual collaboration between the part that resides on the item and the part that resides on the item group. It is useful to think of the number series as a single module, but this module cross-cuts two application objects.

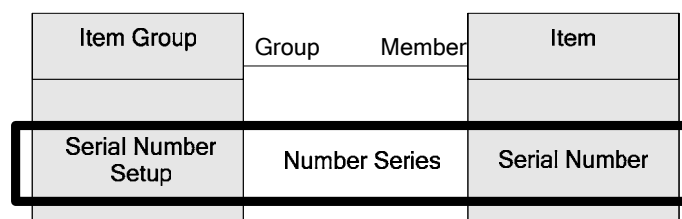


Figure 4. Number Series Cross-Cuts Application Objects

Aspect-oriented programming [9] is one of the approaches, and an additive convention of thought on how to deal with cross-cutting concerns in a modular way. In the scope of domain-driven development, it is useful to think about entities derived from ontological categories as objects and about functionality of the software applications as aspects.

This separation of concerns also determines a mechanism of how to add the new features to the software application without changing its fundamental structure. The objects corresponding to ontological categories determine the fundamental structure of the software application and aspects provide the specific functionality.

3.3 Aspect Categories

In the section about domain objects we have shown that ontological categories correspond to metaclasses of the application objects. A similar approach can be also applied to the entities in aspect dimension.

In this section we describe metaclasses for application aspects. For example, we have shown that the number series is an aspect in the application model. However, the fundamental purpose of the number series is to give the application object unique identity. Therefore, we can think of the number series as a specific instance of a more general aspect category called

identification. Other instances of the identification aspect category are the name, phone number, e-mail address, URL (Uniform Resource Locator), GUID (Globally Unique Identifier) and ISBN (International Standard Book Number).

The identification aspect category specifies a concept of giving identity to something. Identity is not inherently part of the objects and things. People often refer to the real or imaginary things by their names. As the names are not necessarily unique within the application scope, the things are given numbers. Generally, real and imaginary things have one or more given identifiers, so that they can be referred to by using their identifiers.

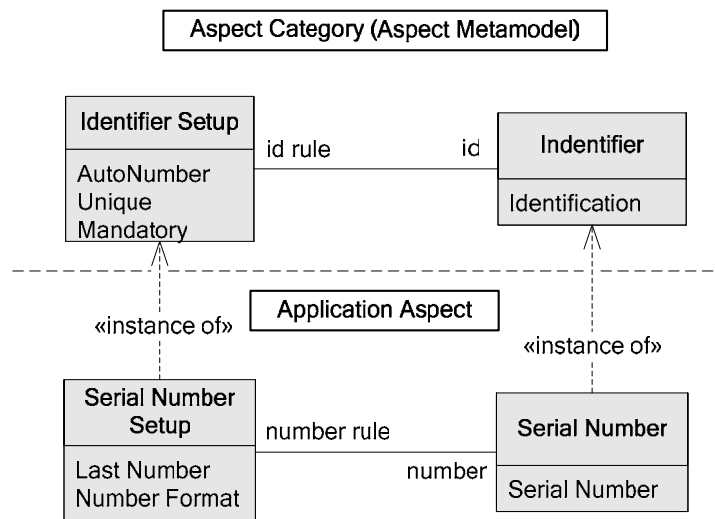


Figure 5. Application Aspect and its Metamodel

Figure 5 illustrates the identification aspect category at the aspect metaclass level and its instantiation in the number series aspect in the application model. Below we describe the identification aspect in more detail to illustrate how the aspect category is specified.

The identification aspect category at the metaclass level consists of two parts: **Identifier Setup** which defines the name of a type of identification. Identifier Setup has the following attributes: **AutoNumber** - a Boolean function that indicates whether the identifier can be automatically generated by the system or not. **Unique** is a Boolean function that indicates whether the identifier is required to be unique or not. **Mandatory** is a Boolean function that indicates whether the identifier must be defined or can be undefined.

The **Identifier** part of the aspect category specifies the data type of the identification, such as a string or a number.

The application level contains aspect parts in which the parameters of the aspect categories have been set. For example, **Serial Number Setup** is an instance of the Identifier Setup that has automatically generated numbers. The numbers are unique and mandatory. The **Serial Number** (an instance of the Identifier) contains attributes for storing the last used number in the series and specifies the identification format; that allows the Serial Number be a combination of numbers and characters.

Other examples of aspect categories in the business domain are the address, account and posting. Their details, along with other aspect categories, have been described as behavioral business patterns [8].

The purpose of the **address aspect category**, see Figure 6, is to specify geographical locations of objects, as well as navigable routes between the locations. The address aspect has four components; the start and destination locations, the actual locations which determines the actual location of some application object and the route, which keeps track of the historical changes of the actual location. The start and destination location are usually configured on economic agents, the route is usually configured on an economic event and the actual location is usually configured on a resource.

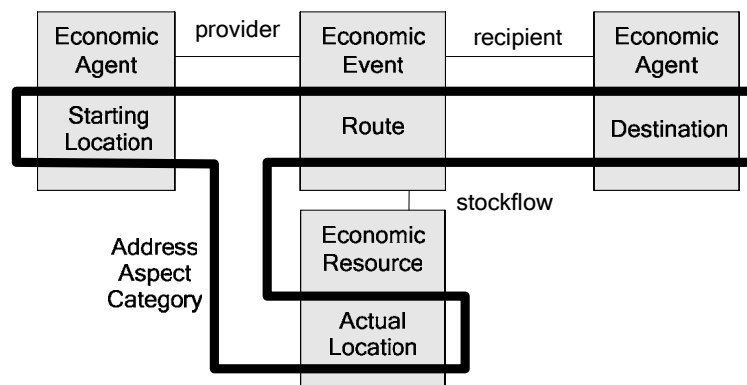


Figure 6. Address Aspect Category

The purpose of the **posting aspect category**, see Figure 7, is to keep track of transactions represented by changes of some application objects. The components of the posting aspect are the entry, which persists the application object and makes it immutable and a number of dimensions that describe the information to be registered with each entry. The entry is typically configured on an economic event or commitment and the dimensions are configured on economic agents, resources, their types and groups.

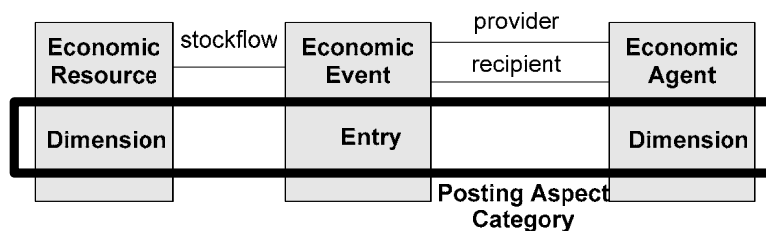


Figure 7. Entry Aspect Category

The purpose of the **account aspect category**, see Figure 8, is to represent aggregated data about entries. The components of the account aspect are the account, which represents the aggregated amount, and one or more entries, which represent the values that increase or decrease the total amount. The account is usually configured on an agent or resource type, and the entries on economic events and commitments.

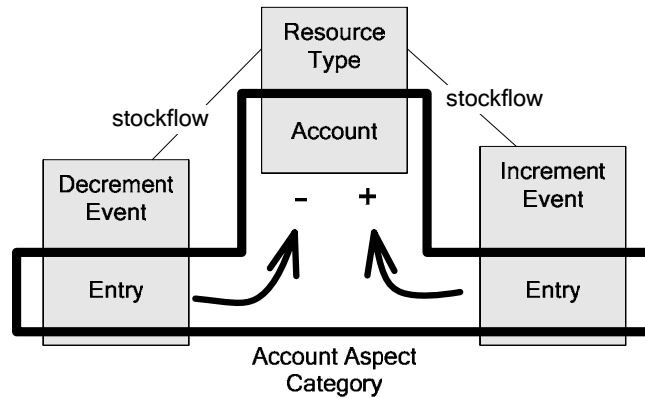


Figure 8. Account Aspect Category

3.4 Number of Aspect Categories

According to our experience the number of aspect categories for a domain is roughly at the same level as the number of ontological categories. For example, the latest version of the REA ontology [5] describes 37 ontological categories; 23 corresponding to metaclasses and 14 to metarelations. A functionality of the CRM (customer relationship management) business application can be fully covered by 18 aspect categories.

4. DOMAIN-DRIVEN DEVELOPMENT

We have shown that a domain-specific model can be decomposed along two dimensions: the object dimension that reflects the ontological categories of the domain and the aspect dimension that reflects the behavior, which the software application must have in order to be useful, see Figure 9. We have also shown that the components both in the object dimension and the aspect dimension can be specified at two levels of specificity; the level of ontological categories or aspect patterns, and the level of application objects and application aspects.

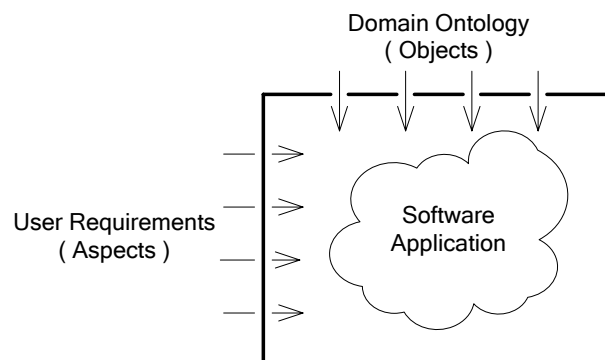


Figure 9. Objects, Aspects and Domain-Driven Development

Figure 9 illustrates the key message of this paper, which is, using domain ontologies to determine the application object model leads to a software architecture with two orthogonal dimensions.

In this section we illustrate how this fundamental idea can be used to develop a software application in a specific domain.

4.1 Application Configuration

If a software provider develops several applications as part of a product line, or develops multitude of very different applications, which all belong to one domain, some form of reuse of the common functionality is desirable. The architecture described in this paper allows for implementing the generalized application functionality in the aspect categories, and develop software applications by configuring the aspect categories with application objects.

The configured software application is an application that conforms to the ontology for the particular domain and also contains specific functionality that meets user's needs. As the application objects are determined by the domain ontology, the process of creating an application model consists of assigning application aspects to application objects. This process is outlined in Figure 10.

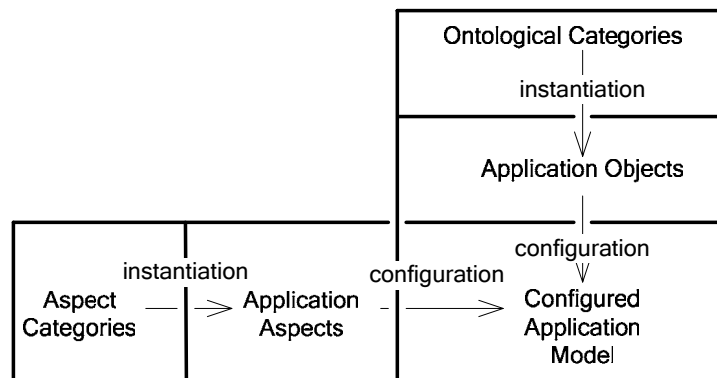


Figure 10. Application Configuration

An example of a business application configured in this way is illustrated in Figure 11. This application is a model of a simple sales module.

The ontological categories in this application are Economic Agent, Economic Event and Economic Resource; their instances are application objects Customer, Sale, Payment, Item and Cash.

The aspect categories in this application are Identification, Account, Address and Posting. Instances of the Identification aspect are Name, Item Number, Customer Number and Transaction ID. Instances of the Account aspect are Inventory Account, Bank Account, Customer Account and Cash Account. Instances of the Address aspect are Billing Address and Shipping Address. Instances of the Posting aspect are G/L (General Ledger) Entry and Inventory Entry.

The choice of the aspect categories is determined by user's needs. Other configurations of the sales process in the software applications for different users would contain a different set of aspect categories.

The Configured Application Model illustrated in Figure 11 contains the Application Objects with Application Aspects. The Customer object contains the aspects Name and Number (identification aspects), Customer Account (account aspect), Billing Address and Shipping Address (address aspects). The Sales object contains the aspects Transaction ID (identification aspect) and G/L Entry (posting aspect). Payment Receipt contains the aspects Transaction ID and G/L Entry (posting aspect). The Item object contains the Item Number (identification aspect), and the Money object contains the Bank Account and Cash Account aspects.

It has been noted that aspects typically cross-cut two or more application objects. For the sake of simplicity, the cross-cutting is not shown in the model in Figure 11, because some parts of the illustrated aspects would reside on the objects that are not shown in the model of the configured software application in Figure 11. For example, Identification Setup (the other part of the Identification aspect) would be present in the objects Customer Type, Sale Type, Payment Receipt Type and Item Type. The other part of Billing Address would be present in the Invoice object; the other part of Shipping Address would be present in the Shipment Object. These objects are not illustrated in Figure 11, as the model would become too large and would obscure the main idea we want to describe in this paper.

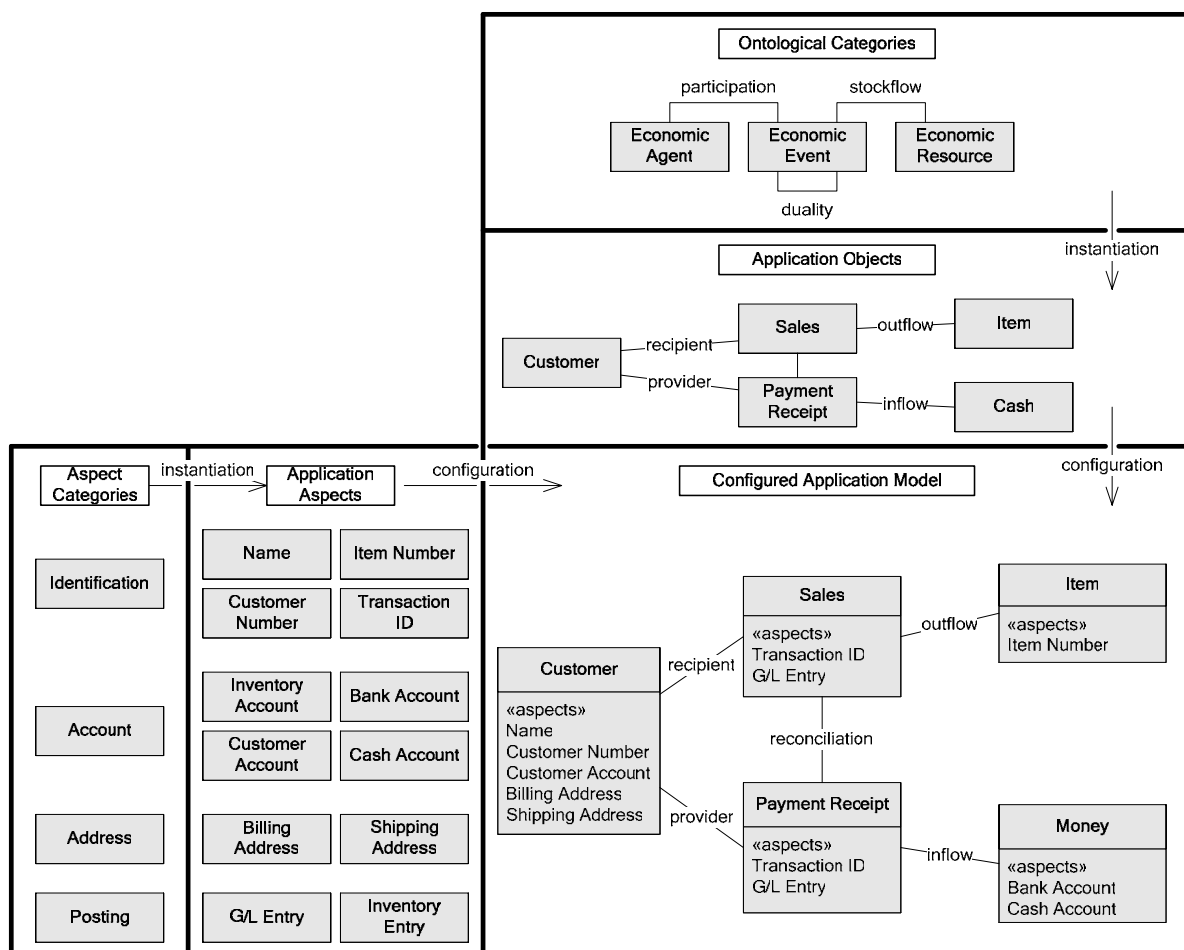


Figure 11. Example of Application Configuration

5. SYSTEM LEVEL ARCHITECTURE

This approach leads to the architecture of a domain-specific system consisting of interconnected components. There is a lower limit determining the minimal size of the component. This limit is derived from domain rules of the ontology. There is no upper limit on the number of interconnected components other than implementation technology. The rest of this section describes it in more detail.

As the ontology designers' aim is to define a minimal set of concepts that describe the domain, there usually is a small number of domain object categories. If an ontology defines domain rules, it is possible to determine a minimal set of domain object categories that must be present in the system in order to satisfy the domain rules. For example, the minimal component in the REA ontology has to include at least one increment and one decrement economic event, each of them having a relationship to an economic resource and each event having two relationships to economic agents. Figure 12 illustrates an example of a minimal REA component; if some of the object categories or relationships are missing, the design would violate the domain rules specified by the ontology.

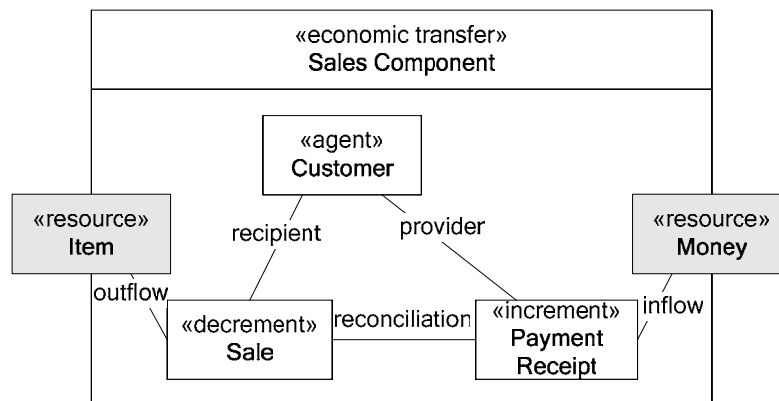


Figure 12. An example of a minimal REA component

The provided and required interfaces of the components, in Figure 12 we use the UML 2.0 port notation [11], are well defined by the shared domain ontology. For example, the ports of the REA components are the economic resources. The economic resources relate the components together into the value chain of the enterprise, see Figure 13.

Each of the components schematically illustrated in Figure 13 has its own, potentially different set of aspects. These components can be implemented as modules of a single ERP system or as several collaborating ERP systems. The components can also be replaced by legacy applications, as long as their ports can be expressed at semantic level in terms of categories of the domain ontology

.

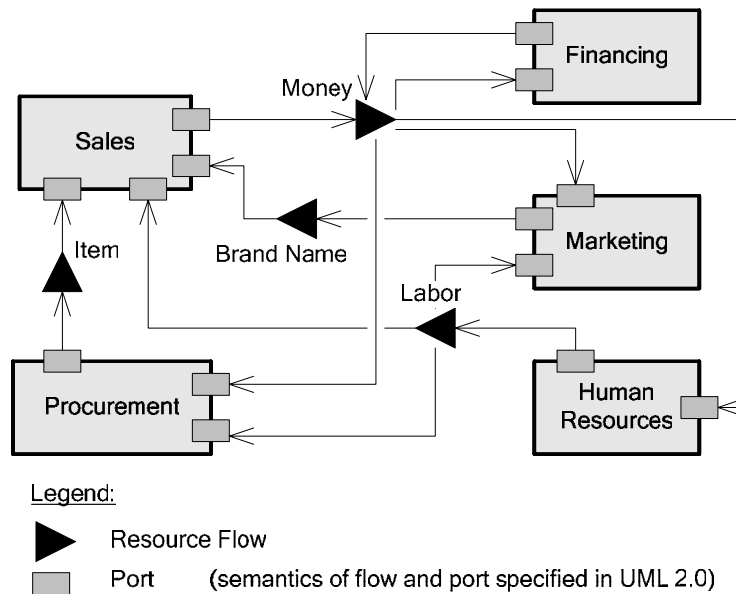


Figure 13. A system of components

Each of the components schematically illustrated in Figure 13 has its own, potentially different set of aspects. These components can be implemented as modules of a single ERP system or as several collaborating ERP systems. The components can also be replaced by legacy applications, as long as their ports can be expressed at semantic level in terms of categories of the domain ontology.

6. HOW THIS APPROACH DIFFER FROM “TRADITIONAL” ASPECT-ORIENTED PROGRAMMING?

In most aspect-oriented approaches, the aspects are used to capture non-functional requirements such as security and persistence. In the ontology-based approach all functional requirements are modeled and implemented as aspects. For example, in Figure 11, the name and address are modeled as aspects, because address and name originate in functional requirements and not in the ontology.

In the ontology-based approach **all** attributes and methods of an application object should be part of an aspect (or several aspects if the implementation allows for that). This is because aspect categories define the semantics of the application object’s features. If there would be an attribute that is not part of the aspect, the semantics of this attribute will be unknown to the business application. As a consequence, this approach forces the application developer to discover the semantic of each attribute that is required by the users, and implement an aspect category for that attribute that captures this semantic.

The non-functional requirements are in this approach modeled as model transformations to other domains, in which the ontology describes the semantics of the non-functional requirement. For example, security will be modeled as a transformation from the business-logic domain characterized by the REA ontology, to the domain in which the basic ontological categories are securable item, user role and security rights. As the transformation rules can be defined at the metamodel level that corresponds to ontological categories, a tool can automatically generate security settings for a specific application business-logic model.

6. SUMMARY

We have illustrated that in the domain-driven development of software applications, application designers can use domain ontologies as one of the sources for creating application models, in addition to traditional analysis based on user requirements. The conceptualized domain knowledge in the form of domain ontology can be used as a metamodel for application models.

However, domain ontologies cannot describe specific functionality and the differences between different applications in a given domain, which originate in user requirements, because the concepts of the domain ontology must be applicable to all systems in the domain. The modules of functionality resulting from the user requirements are often not localizable into application objects that originate in the domain ontology.

To solve this conflict, we have illustrated that the domain-specific software applications lead to the architecture of a component with two orthogonal dimensions. The object dimension represents the categories originating in the domain ontology and the aspect dimension represents the functional modules that originate from user requirements.

This approach determines the software architecture for domain specific software applications, allowing customizations within a given domain and adding new features into the existing components without changing their fundamental structure. As this approach specifies the semantics of ports of the components and subsequently the component connectors, it also determines the software architecture at the system level.

7. REFERENCES

- [1] Czarnecki, K. Eisenecker, U.W: *Generative Programming - Methods, Tools, and Applications*, Addison-Wesley, 2000
- [2] Buschmann, F. et al: *Pattern-Oriented Software Architecture*, Wiley, 1996
- [3] Clarke, S, Baniassad, E.: *Theme: An Approach for Aspect-Oriented Analysis and Design*, ICSE 2004
- [4] Geerts, G. McCarthy, W. *The Ontological Foundation of REA Enterprise Information Systems*, Michigan State University, August 2000
- [5] Geerts, G. McCarthy, W. *An Ontological Analysis of the Economic Primitives of the Extended REA Enterprise Information Architecture*, The International Journal of Accounting Information Systems, 2002, Vol. 3, pp. 1-16
- [6] Generative Techniques in the context of Model Driven Architecture, Workshop at OOPSLA 2002
- [7] Gruber, T. R. *A Translation Approach to Portable Ontology Specifications*, Knowledge Acquisition, 1993.
- [8] Hruby, P. et al: *Model-Driven Development with Business Patterns*, Springer-Verlag, to be published.
- [9] Kiczales, G. et al: *Aspect-Oriented Programming*, in: M. Aksit and S. Matsuoka (Eds.): ECOOP '97 Proceedings, Jyväskylä, Finland, June 1997, Springer-Verlag Berlin Heidelberg 1997.

- [10] Tarr, P. Ossher, H. Harrison, W, Sutton, S. M. Jr. *N Degrees of Separation: Multi-Dimensional Separation of Concerns*, Proceedings 21st International Conference on Software Engineering (ICSE'99), May 1999.
- [11] *UML 2.0 Superstructure Specification*, OMG document ptc/04-10-02, 2004.