A View-Oriented Approach to System Modelling based on Graph Transformation*

Gregor Engels[1], Reiko Heckel[2], Gabi Taentzer[2], Hartmut Ehrig[2]

[1] Leiden University, Dept. of Computer Science, P.O. Box 9512, NL-2300 RA Leiden, The Netherlands
engels@wi.leidenuniv.nl
[2] Technical University of Berlin, Dept. of Computer Science, Franklinstrasse 28/29, D-10587 Berlin, Germany
{ehrig, reiko, gabi}@cs.tu-berlin.de

Abstract. The idea of a combined reference model- and view-based specification approach has been proposed recently in the software engineering community. In this paper we present a specification technique based on graph transformations which supports such a development approach. The use of graphs and graph transformations supports an intuitive understanding and an integration of static and dynamic aspects on a well-defined semantical base. On this background, formal notions of view and view relation are developed and the behaviour of views is described by a loose semantics. We define a construction for automatic view integration which assumes that the dependencies between different views are described by a reference model. The views and the reference model are kept consistent manually, which is the task of a model manager. All concepts and results are illustrated at the well-known example of a banking system.

keywords specification language, view, viewpoint, view integration, software process, graph transformation systems

1 Introduction

The most challenging issue of software engineering still is the question how to master the complexity of the development of large software systems. Currently, a variety of approaches try to solve certain aspects of this problem.

One important approach is to reuse well-established pieces of specifications, documentations, and/or software, while developing a new system. While in the beginning the reuse idea was restricted to often needed classes, in the meantime it has become clear that reuse should be tackled on a much greater scale by specialising integrated networks of classes, so-called frameworks [LRP+96]. Thus, a major research as well as development field is currently the definition of frameworks for various application domains.

Another important approach is based on the observation that, due to the size and the diversity of the planned software system, teams of concurrently working application engineers are needed for the realisation of a software system. For instance, during the requirements specification phase, a team of application engineers with different skills and backgrounds is splitted into subgroups. Each subgroup specifies only that aspect of an interactive software system, which is later seen and used by a certain type of user (role). Thus, modularisation concepts are required, which allow to compose a complete and consistent specification out of possibly overlapping pieces.

In the data base world, but also in the field of software and requirements engineering, one way to obtain this modularisation is the concept of views and view integration. In the data base world it is standard to distinguish between a conceptual model and several external models, which are considered to be individual views of the data base. Each view is a restriction of the conceptual model - the total community user view - to just that portion of interest to that particular user (cf. [Dat79]).

In the software engineering field, the view-oriented approach is known by the notion of viewpoints (cf. [FKN+92]). In contrast to the view integration approach, here is no common integrated model intended. The basic idea is to monitor the relationships between different viewpoints, to detect inconsistencies and to resolve them by interactive support of the user. Relationships between different viewpoints are inferred by the use of common names.

This implies that the different application engineers agree on a certain vocabulary for a specific problem domain before they start to develop their own viewpoint. As all notions within a problem domain are somehow related, a more suited starting point than a long list of notions is a so-called reference model for a problem domain, where basic notions and their interrelations are fixed. This idea of a combined reference model- and view-based specification approach was especially proposed by B. Balzer during his keynote speech at SP'96 (Software Process) (cf. [Bal96]).

In this paper we present a specification technique based on graph transformations which supports such a development approach. We explain this approach informally in Section 2. In Section 3, 4, and 5 we present the formal base of our approach together with illustrating examples. The basic notions of graphs and graph transformation for the modelling of static and dynamic aspects of software systems are presented in Section 3. In Section 4 we give a definition for views of graph transformation systems. Since all different views are required to be based on a common reference model, we are able to present in Section 5 a general construction for view integration. This can be considered as an automatic view integration. Finally in Section 6 we summarise the main ideas and discuss some remaining open problems.

* Research partially supported by the German Research Council (DFG), the TMR network GETGRATS, and the ESPRIT Basic Research Working Group APPLIGRAPH

2 Concept of Views and View Integration Using Graph Transformations

Graph grammars and transformations have been introduced as a generalisation of Chomsky grammars on one hand and of term rewriting systems on the other hand about 25 years ago. Meanwhile there is a well-established theory of graph transformations (see e.g., [Roz97]) which has a number of applications to system modelling and software engineering (cf. [Nag96,Zam96,AE96,SWZ95]) based on concrete specification languages and supporting tools (cf. [Sch91,LB93]). The main idea of our specification approach is to model object structures and their interrelationships by graphs and modifying operations by graph transformations. In particular our approach is based on typed graph transformation systems [CEL+96,HCEL96,Rib96] which allow to define a set of graphs by a type (scheme) graph together with type-consistent operations on these graphs. Compared to currently popular object-oriented modelling techniques, typed graph transformation systems really support an integrated modelling of static and dynamic aspects, which goes much further than the use of common names. Nevertheless, it shall be noted that graph transformations in its pure form are not object-oriented. There are, however, class-based extensions (see e.g., [Tae96,Wag97]).

This paper applies typed graph transformation systems for defining the concept of a view that models a certain aspect of the complete system. Thus, a view specifies only partially the structure of the system's state and analogously only partially, what the effect of an operation is. It may be that a view operation, being executed on the system's state, has to be concurrently coupled with operations of other views to ensure a consistent system's state transition. Thus, a view specifies only what at least has to happen on a system's state. In this sense, the semantics of a view can only be a loose one, in contrast to the semantics of the complete model.

The overall specification approach can be sketched as follows (cf. Figure 1). Starting with a common reference model, each application engineer develops his

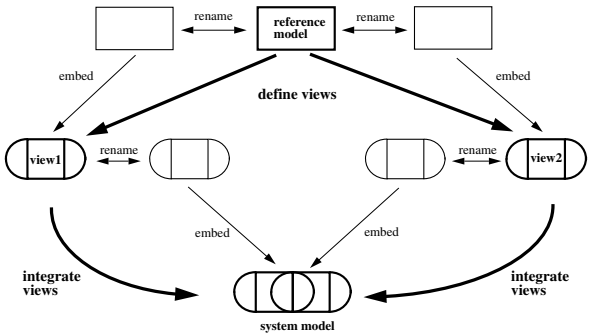

Fig. 1. A view-oriented approach to system modelling.

own viewpoint by extending and refining the reference model appropriately. In the case that different names for the same concept have been used, a renaming step has to be executed by the application engineer. We will explain later that technically spoken, a (partial) specification is called a view on another specification, if a renamed version of the first can be embedded into the second. In a following step, all these views have to be integrated or synchronised in a common system model. We will show in this paper that by following our specification approach, this integration step can be done automatically. This is mainly due to the fact that we assume that the definition of all views is based on the same common reference model.

This basic assumption is in line with above mentioned current approaches in the object-oriented world, where also reference models in the form of domain-specific frameworks are regarded as the desired starting point for any new software development project. But in addition and in contrast to such a framework-based specification approach, we allow that the framework (or reference model) is specialised concurrently by several views.

Following such a view-based specification approach, different forms of possible inconsistencies can be distinguished.

(i) The same concept, e.g. operation, is specified in two different views by using different names.

(ii) The same names are used in two different views denoting semantically different concepts.

In particular, the first form (i) of inconsistency has extensively been investigated in database research, as it is one of the problems which have to be solved during scheme integration (cf. [BLN86]). Instead of trying to identify dependencies between different names, we start with a common reference model of names and their interrelations. In the case that different views want to share the same name for the same concept and this name is not yet contained in the reference model, the reference model has to be extended. In this situation, a model manager is required who mediates between the different view designers and extends the reference model appropriately.

In the second case (ii), two solutions are possible: The two names are kept distinct within the overall specification (for instance, by qualifying them with the view name) or the two names are even rejected by the model manager.

While the above explained two forms of inconsistencies relate to static inconsistencies between specification documents, a third form of inconsistency may occur during executing (or enacting) the system.

(iii) Execution of a view operation violates the constraints defined by another view.

This means that two different views overlap in their specification of the desired system's behaviour. In this case, the two views have to be synchronised to achieve a consistent system's behaviour specification.

Different solutions for (iii) can be distinguished. The viewpoint approach (cf. [FKN+92]) follows an algorithmic approach by checking the effect of operations and triggering update operations to end in a consistent result state. Other

specification approaches, like e.g. Z (cf. [Jac95]), follow a descriptive approach, where the application engineer has to integrate different view specifications in an overall specification by additional inter-view constraints. In our approach, we follow a constructive approach, where different views are automatically integrated. This means that two operations from two different views are merged into one operation in the resulting overall system specification. The common underlying reference model indicates and identifies the overlapping part.

This discussion shows that the presented automatic view integration relies on a series of prerequisites. In particular, a human being in the role of a model manager is required to guide and support the integration process. Nevertheless, the construction of the integrated specification will be done automatically.

3 Graph Transformation for System Modelling

In this section, we explain how rule-based graph transformations can be used to model the static and dynamic aspects of software systems in a formal and integrated way. The main concepts are illustrated by a small banking example.

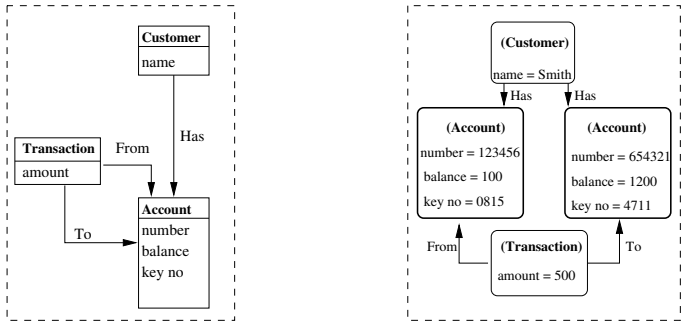

Fig. 2. Example of scheme and instance graphs

Graphs. Graphs and diagrams are often used in software engineering for visualising complex structures. We only mention Entity-Relationship (ER) diagrams and instances in data modelling or class and object diagrams in object-oriented design. Formally, a *graph* consists of a set of vertices V and a set of edges E such that each edge e in E has a source and a target vertex $s(e)$ and $t(e)$ in V, respectively.

Both in ER modelling and OO design graphs occur on two levels, as *scheme graphs* (ER diagram, class diagram) and their *instance graphs* (ER instance, object diagram). Scheme graphs impose structural constraints on its instances by requiring that each instance can be mapped to its scheme in a structure-preserving way. This mapping also provides vertices and edges of the instance graph with their *types*, i.e., the vertices and edges of the scheme graph.

Example. A sample pair of scheme and instance graphs is shown in Figure 2. The scheme graph on the left contains the main object and relationship types.

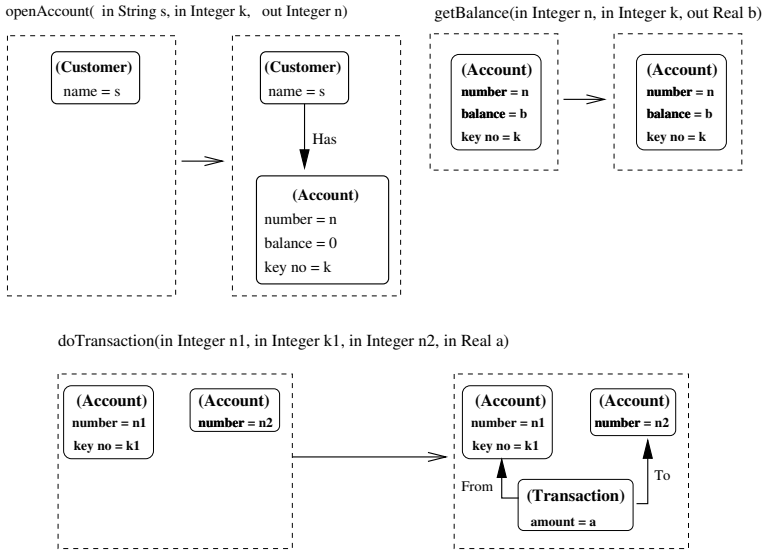

Fig. 3. Graph transformation rules for opening a new account, getting the balance, and starting a transfer transaction.

Object types are *Customer*, *Account*, and *Transaction*. Customers have a name and are linked by a *Has* relationship to their accounts. Accounts have an account number for identification, a key number for authorised access and, of course, a balance. Transactions are requests for transferring money between accounts. On the right side of Figure 2, an instance of this scheme is shown. It represents a toy state of the banking system where a customer holds two accounts with an ongoing transaction. △

Rule-Based Graph Transformation. State changing operations on graphs are modelled by graph transformations which are specified by *graph transformation rules* $r : L \to R$. They consist of a rule name r and two instance graphs L and R, called *left-* and *right-hand side*, which represent a part of the system's state before and after the operation, that is, the pre– and postcondition of the operation. We assume that the intersection $L \cap R$ is a graph, called the *interface* of r. It contains those items that are read but not deleted by the operation.

Example. The upper left rule in Figure 3 specifies the customer's operation of opening a new account. It requires the customer's name and a key number as input. Then, a *Customer* object with this name is selected from the current state, and a new *Account* object is created together with a *Has* relationship. Balance and key number are set, and an account number is chosen and passed to the customer as output. The *getBalance* rule in the right specifies the operation of reading the balance of an account, and *doTransaction* below creates a transfer transaction. △

More generally, a derivation step $G \overset{r}{\Longrightarrow} H$ from G to H using a rule $r : L \to R$ requires that (a renaming of) L occurs as a subgraph in G. Then, $L \backslash R$ is removed

from G and $R \setminus L$ is added to the result leading to the derived graph H that contains a renaming of R as a subgraph.[1] Hence we denote by $delete(r)$ the part $L \setminus R$ and by $add(r)$ the part $R \setminus L$. The application deletes and creates exactly what is specified by the rule, i.e., there is an implicit *frame condition* stating that everything that is not rewritten explicitly by the rule is left unchanged.

The same rule can also be interpreted in a more loose way. In this case, it specifies only some part or local view of the changes that affect the current state. Since we are interested in the behaviour of views, we introduce the notion of *graph transition* by dropping the above mentioned frame condition. Like a derivation step, a graph transition $G \overset{r}{\rightsquigarrow} H$ from G to H via r requires that L occurs in G. Then, *at least* $delete(r)$ is removed from G and *at least* $add(r)$ is added, but there may be unspecified deletion and addition as well.

For modelling state transitions that are entirely caused by the environment we introduce ϵ-transitions, i.e., transitions using the empty production $\epsilon : \emptyset \to \emptyset$. Such a production specifies no effect, that is, $delete(\epsilon) = add(\epsilon) = \emptyset$. A transition via ϵ allows any change to the current state.

Graph Transformation Systems. A graph transformation system $\mathbf{G} = \langle SG, N, \rho \rangle$ consists of a scheme graph SG, a set of rule names N and a mapping ρ providing for each rule name r a rule $L \to R$ where L and R are instance graphs of SG.

The rules of a graph transformation system can be applied sequentially and/or in parallel leading to sequences of (parallel) derivation steps. These *derivation sequences* form its classical semantics that corresponds to the *closed behaviour* of a non-reactive and fully specified system. The *loose semantics* of

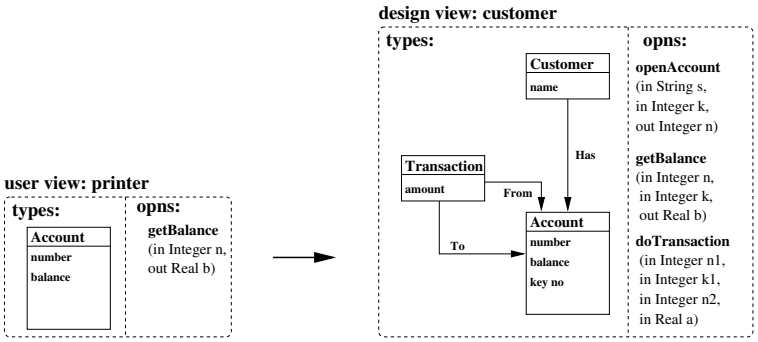

Fig. 4. A graph transformation system modelling the customer's view of a bank with a view relation from an account printer's view.

a graph transformation system $\mathbf{G}$, given by all *transition sequences* in $\mathbf{G}$, represents the *open behaviour* of a system that is embedded in a not completely specified environment.

[1] Formally, the occurrences of L in G and R in H are specified by a graph homomorphism $m : L \to G$ and $m^* : R \to H$, called *match* and *comatch* respectively, that allow the renaming of L and R.

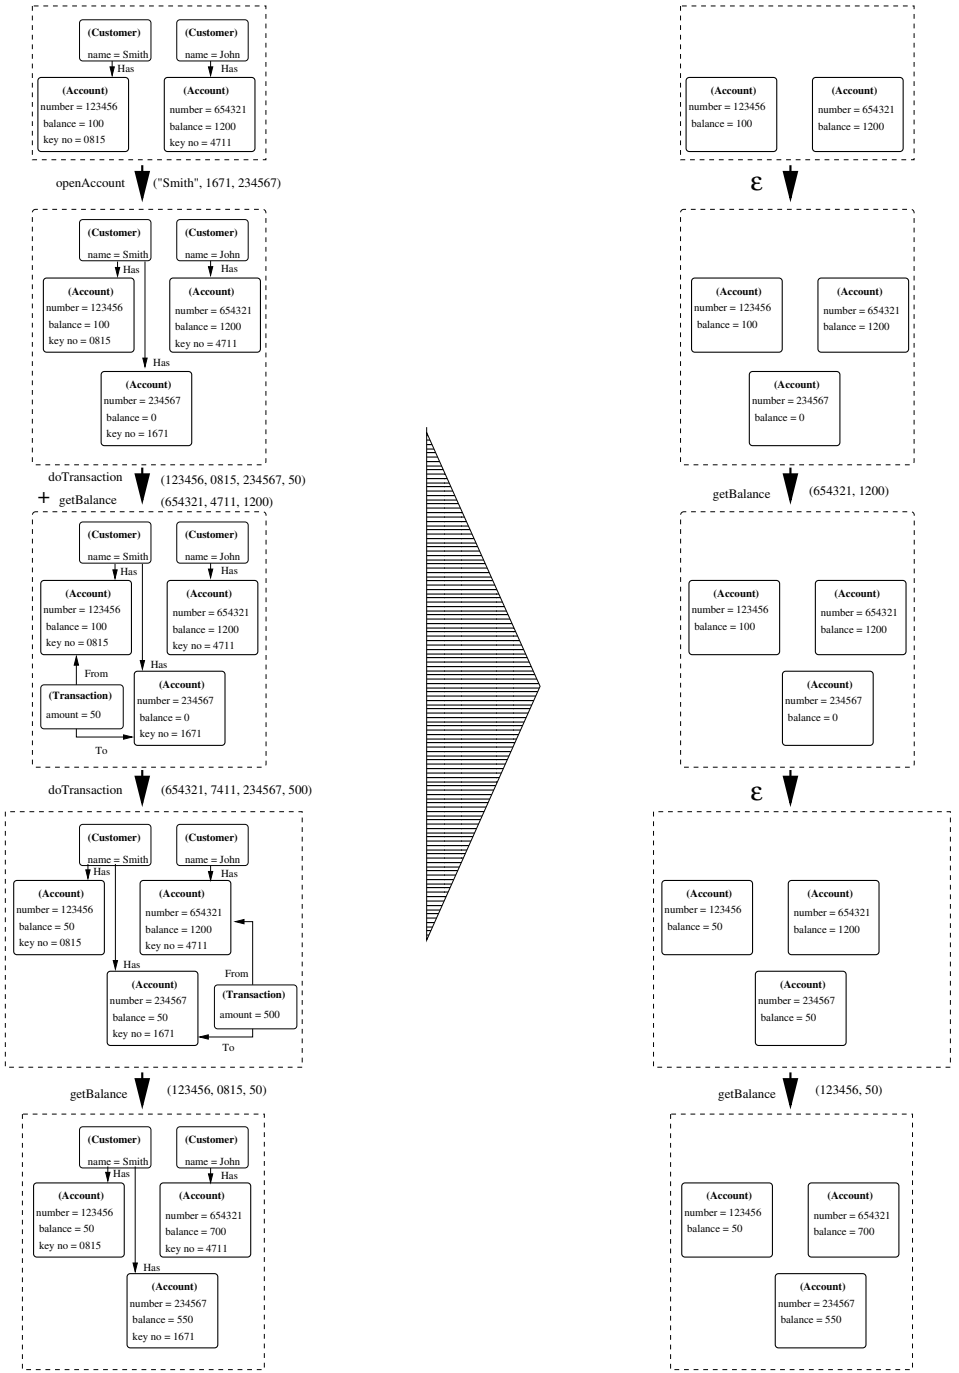

Fig. 5. Transition sequences in the customer's and the printer's view.

Example. A sample graph transformation system modelling a banking system from the customer's point of view is depicted in Figure 4 on the right. The rules are shown in Figure 3. The left-hand side of Figure 5 shows a sample transition sequence modelling the customer's view of some banking operations. After opening the new account "234567" (where the number is provided by the bank), a transfer transaction is ordered by customer Smith. At the same time, customer John asks for the balance of his account. This is modelled by two derivation steps (without unspecified changes) where the second one consists of the parallel application of *doTransaction* and *getBalance*. Thereafter, also customer John starts an order, while the first order is executed by the bank and the result becomes visible for the customers. Thus, the first *Transaction* object disappears without actually being deleted by the *doTransaction* rule, and the balances of the accounts "123456" and "234567" change. These are (from the customer's point of view) unspecified effects, i.e., the third step in the sequence is a true transition. Finally, customer Smith asks for the balance of account "123456" while the second transaction order is executed (which again is a true transition). △

4 Views of Graph Transformation Systems

For the integration of views it is essential to specify the intended correspondences between different views by relating a common reference model to each view by a view relation. A view relation allows the renaming and extension of scheme graphs and rules. After the integration, similar view relations are established between a view and the overall system model.

Renaming. In order to allow, for example, the use of different names for the same operation in different views and the reference model, renaming relations are introduced. A *renaming relation* $\mathbf{G} \overset{ren}{\longleftrightarrow} \mathbf{G}'$ can be seen as a kind of dictionary establishing a one-to-one correspondence between the types, the rule names, and the (vertices and edges of the) rules of two graph transformation systems $\mathbf{G}$ and $\mathbf{G}'$. If x is an item (a type, a rule name, etc.) of $\mathbf{G}$ and x' the corresponding item in $\mathbf{G}'$ we write $x \overset{ren}{\longleftrightarrow} x'$.

Extension. As anticipated above, a view relation shall be composed of a renaming and an extension. The extension of a rule r_0 by another rule r_1 is modelled by the *subrule relation*. The rule $r_0 : L_0 \to R_0$ is a *subrule* of $r_1 : L_1 \to R_1$, written $r_0 \subseteq r_1$ if the effects of applying the latter extend the effects of applying the first. Formally, this means that $L_0 \subseteq L_1, R_0 \subseteq R_1$ (pre- and post-conditions are extended), $delete(r_0) \subseteq delete(r_1)$ (more is deleted by r_1), and $add(r_0) \subseteq add(r_1)$ (more is added by r_1).

A graph transformation system $\mathbf{G}_1$ *extends* another one $\mathbf{G}_0$, written $\mathbf{G}_0 \subseteq \mathbf{G}_1$, if scheme graph and rule names of $\mathbf{G}_0$ are extended, i.e., $SG_0 \subseteq SG_1$ and $N_0 \subseteq N_1$, and for each rule name $r \in N_0$ the associated rule in $\mathbf{G}_0$ is a subrule of the one in $\mathbf{G}_1$, i.e., $\rho_0(r) \subseteq \rho_1(r)$.

View relation. In order to specify the relation between a view and a system model, a *view relation* has to be specified. A view relation $v = (\mathbf{G}_0 \overset{ren}{\longleftrightarrow} \mathbf{G}_0^1 \subseteq \mathbf{G}_1)$ from $\mathbf{G}_0$ to $\mathbf{G}_1$ is a renaming of $\mathbf{G}_0$ such that $\mathbf{G}_1$ is an extension of the renamed system $\mathbf{G}_0^1$. More abstractly, we write $v : \mathbf{G}_0 \to \mathbf{G}_1$ and say that $\mathbf{G}_0$ *is a view of* $\mathbf{G}_1$. View relations may be composed by composing the underlying renamings and extensions in a suitable way. This makes it possible to regard a view $v_0 : \mathbf{G}_0 \to \mathbf{G}_1$ on a view $v_1 : \mathbf{G}_1 \to \mathbf{G}_2$ as a view $v_0; v_1 : \mathbf{G}_0 \to \mathbf{G}_2$.

Let's discuss in more detail the relationship between a graph transformation system $\mathbf{G}_1$ and its view $\mathbf{G}_0$.

- A name x_0 of $\mathbf{G}_0$ may change to x_1 in $\mathbf{G}_1$. In order to represent this relationship, a dictionary $\mathbf{G}_0 \overset{ren}{\longleftrightarrow} \mathbf{G}_1$ is used containing the entry $x_0 \overset{ren}{\longleftrightarrow} x_1$ (and $x \overset{ren}{\longleftrightarrow} x$ for all unchanged names of $\mathbf{G}_0$).
- $\mathbf{G}_0$ may be extended by $\mathbf{G}_1$ by introducing new types or rule names. An item is new in $\mathbf{G}_1$ if it is not listed in the dictionary *ren*. The new types and rules of $\mathbf{G}_1$ are not visible to $\mathbf{G}_0$, i.e., they do not belong to this view.
- A rule r defined in $\mathbf{G}_0$ may be extended in $\mathbf{G}_1$. On the left- and/or right-hand side of r new vertices and edges may be added, while the name of the rule might be the same. This means that the effects of applying r are extended (additional items are deleted and/or added) while the pre- and post-conditions are strengthened.

Example. Figure 4 shows a view *Printer* of the customer's view that shall become a view of the banking system by composition of view relations later on. It models the restricted view of a printer where customers can ask for the balances of their accounts. Such a printer does not know about key numbers and is not able to open new accounts or to order transactions. Hence, the corresponding types and rules are not visible in the printers view. A renaming is not needed.

The printer's view of the sample transition sequence in Figure 5 on the left is shown in the same figure on the right. Recall that, e.g., the step in the user sequence on the left – opening an account – is a derivation step. In the printer's sequence on the right, however, it is seen as an ϵ-transition. The first *getBalance*-transition in the printer's view results from the parallel derivation step using *doTransaction* and *getBalance* in the customer's view. The *doTransaction* rule is hidden in the printer's view but its effects are still visible. △

Hence, a view relation $v : \mathbf{G}_0 \to \mathbf{G}_1$ describes not only a projection of the state graphs of $\mathbf{G}_1$ to $\mathbf{G}_0$ but also a more abstract view of the behaviour of $\mathbf{G}_1$. Notice that derivation sequences (without unspecified effects) are not viewed as derivation sequences in general but as transition sequences, too: The view of a derivation sequence in $\mathbf{G}_1$ may be a transition sequence in $\mathbf{G}_0$.

5 Automatic Integration of Views

In the previous section, view relations were introduced for describing the relationships between views, reference model, and system model. Now, we explain the main technical concept of our approach, the integration of views.

All development starts with the reference model as domain-specific framework. The reference model of the banking system is shown in Figure 8. In the

beginning it contains the basic object and relationship types of the banking system (i.e., without the operation *newAccount* that shall be added later on as extension of the reference model). The views *Customer* and *Clerk* are derived independently of each other from the reference model. This results in a situation where the reference model itself forms a view on the two specifications *Customer* and *Clerk*, which are so-called design views on the complete system model.

When integrating the design views to the system model, we have to know which items in the two views represent the same types and operations. Rather then relying on the names of these items, this correspondence is specified by the reference model, i.e., two items are assumed to represent the same concept if and only if they have a common origin. In fact, since view relations allow the renaming of items, this means that developers are free to choose the names in their view according to the preferences of the particular user group.

A situation of two design views $\mathbf{G}_1$ and $\mathbf{G}_2$ based on a reference model $\mathbf{G}_0$ by view relations v_1 and v_2 is shown in Figure 6. The integration of $\mathbf{G}_1$ and $\mathbf{G}_2$ over $\mathbf{G}_0$ will be done in two steps. First, we have to rename the design views so that the renamed views $\mathbf{G}'_1$ and $\mathbf{G}'_2$ share a name if and only if there is a common origin in the reference model $\mathbf{G}_0$. Then the construction of the integrated system view can be done by componentwise set-theoretical union of $\mathbf{G}'_1$ and $\mathbf{G}'_2$. We assume that the names in the design views $\mathbf{G}_1$ and $\mathbf{G}_2$ are disjoint which can be ensured by qualification with the name of the view (like $\mathbf{G}_i.name$). The view relations $v_i : \mathbf{G}_0 \to \mathbf{G}_i$ are given by $\mathbf{G}_0 \overset{ren_i}{\longleftrightarrow} \mathbf{G}_0^i \subseteq \mathbf{G}_i$ in Figure 6.

Renaming of Views. The system $\mathbf{G}'_1$ and the renaming $\mathbf{G}_1 \overset{ren'_1}{\longleftrightarrow} \mathbf{G}'_1$ are obtained by extending the renaming $\mathbf{G}_0^1 \overset{ren_1}{\longleftrightarrow} \mathbf{G}_0$ to $\mathbf{G}_1$. More precisely, ren'_1 agrees with ren_1 on $\mathbf{G}_0^1$, i.e., $ren'_1|_{\mathbf{G}_0^1} = ren_1$, and is minimal in the sense that nothing else is renamed, i.e., the renaming is the identity on $\mathbf{G}_0 \setminus \mathbf{G}_0^1$. The renamed system $\mathbf{G}'_1$ becomes an extension of $\mathbf{G}_0$. In a similar way we obtain $\mathbf{G}'_2$ with $\mathbf{G}_0 \subseteq \mathbf{G}'_2$ and $\mathbf{G}_2 \overset{ren'_2}{\longleftrightarrow} \mathbf{G}'_2$ by extending the renaming ren_2 to $\mathbf{G}_2$. This renaming is always possible and can be done automatically. Now, the integrated view $\mathbf{G}'_3$ can be constructed as union of the renamed views $\mathbf{G}'_1$ and $\mathbf{G}'_2$ in Figure 6.

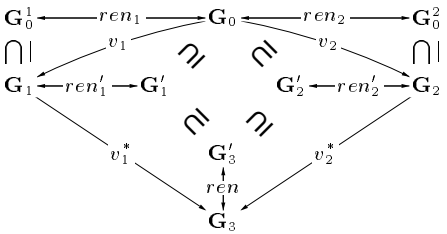

Fig. 6. Integration of the design views $\mathbf{G}_1$ and $\mathbf{G}_2$ to the system model $\mathbf{G}_3$.

Construction of Integrated View. The integrated view $\mathbf{G}'_3 = \langle SG'_3, N'_3, \rho'_3 \rangle$ is obtained by forming the union of the scheme graphs $SG'_3 = SG'_1 \cup SG'_2$ and of

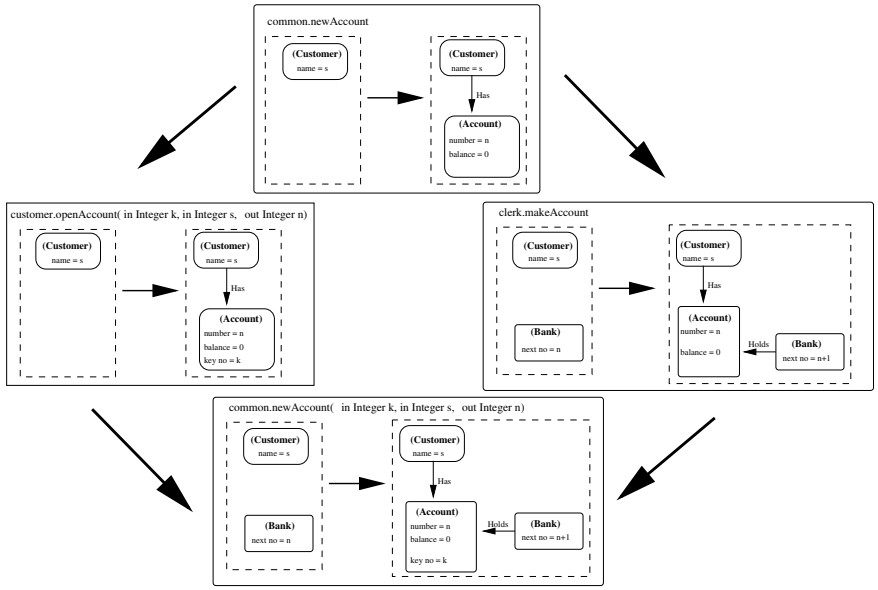

Fig. 7. Synchronisation of the *openAccount* rules.

the sets of rule names $N'_3 = N'_1 \cup N'_2$ of $\mathbf{G}'_1$ and $\mathbf{G}'_2$, respectively. For the rule $\rho'_3(r')$ associated with a rule name $r' \in N'_3$ we distinguish three cases:

- If $r' \in N'_1 \setminus N'_2$ then $\rho'_3(r') = \rho'_1(r')$, i.e., the rule of $\mathbf{G}'_1$ is inherited
- If $r' \in N'_2 \setminus N'_1$ then $\rho'_3(r') = \rho'_2(r')$, i.e., the rule of $\mathbf{G}'_2$ is inherited
- If $r' \in N'_2 \cap N'_1$ then $\rho'_3(r') = L'_1 \cup L'_2 \to R'_1 \cup R'_2$ where $L'_i \to R'_i$ is the rule associated with r'_i in $\mathbf{G}'_i$ for $i = 1, 2$. The new so-called synchronised rule obtained by componentwise union of left- and right-hand sides models the combined effect of applying these two rules simultaneously.

Then, the integrated view $\mathbf{G}'_3$ may be renamed to $\mathbf{G}_3$ via ren. The view relation v_1^* is obtained by composing the view relation $\mathbf{G}_1 \overset{ren'_1}{\longleftrightarrow} \mathbf{G}'_1 \subseteq \mathbf{G}'_3$ with the renaming ren (which is a special view relation as well). In a similar way we obtain view relation v_2^*.

Example. The synchronised rule *common.newAccount* is constructed in Figure 7 as the union of the rules *openAccount* and *makeAccount* of the customer's and the clerk's view. It synchronises the activities that are necessary for creating a new account. The integrated system model *Bank* in Figure 8 also contains the other rules of the customer's and the clerk's view, which are not synchronised. Note that not only the customer's view contains a rule *doTransaction* but also the clerk's view. These two rules are not identified, however, since they have no common source in the reference model. The name conflict is resolved automatically by qualification of the local names with the names of the views. (The qualifications are skipped in the clerk's and customer's view in Figure 8.) On the other hand, the rules *openAccount* and *makeAccount* represent the same op-

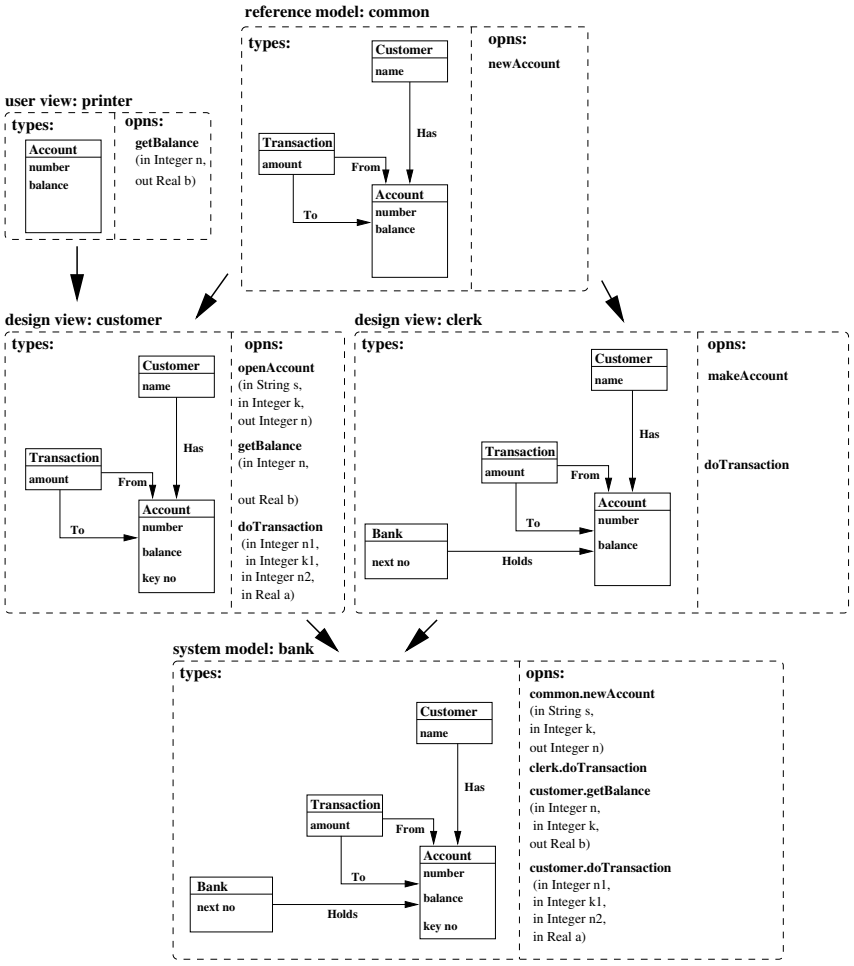

Fig. 8. Integration of the customer's and the clerk's view.

eration (despite their different names) since they both stem from the same rule *common.newAccount*. They are both renamed to *common.newAccount* in the renaming step of the construction. The rule *clerk.doTransaction* is not shown. It describes the execution of a transfer transaction. △

More generally, we may have the following situations:

- It may be the case that "semantically the same" concept is described in $\mathbf{G}_1$ and $\mathbf{G}_2$ using different names (like *openAccount* and *makeAccount* above). This relationship between $\mathbf{G}_1$ and $\mathbf{G}_2$ is only understood (and may be taken into account by the integration) if both names have a common source in the common view $\mathbf{G}_0$ (like *newAccount*). This has to be defined in the dictionaries ren_1 and ren_2. Then, the concept occurs only once in the integrated model, under the name of the common view. If this relationship is not specified, the two concepts are considered as unrelated and are kept separately in the integrated model. This illustrates also the difference between a synchronised rule and the parallel application of two rules. A parallel application of *Customer.openAccount* and *Clerk.makeAccount* would create two distinct new *Account* nodes. The synchronised rule *common.newAccount* realizes that, as desired, only one new account node is created.
- On the other hand, the same name may be used in $\mathbf{G}_1$ and $\mathbf{G}_2$ in order to describe "semantically different" concepts (e.g., *doTransaction* in the customer's and the clerk's view). This does not causes any problem in our approach since we assume that the names are qualified, i.e., *Customer.doTransaction* and *Clerk.doTransaction*.
- In order to represent shared knowledge of $\mathbf{G}_1$ and $\mathbf{G}_2$, which is not yet expressed by the reference model, it has to be extended. This should be done by the model manager. If the reference model is used by more than two views, however, this means that the extra information is also propagated to all other views as well. If this is not desired the reference model has to be kept unchanged and an *abstract view* has to be introduced instead which is also based on the reference model and specifies the sharing between $\mathbf{G}_1$ and $\mathbf{G}_2$. The result is a hierarchy of views. Scenarios of more than two views are discussed in a subsequent paper.
- A similar observation holds if a rule of $\mathbf{G}_0$ is extended in $\mathbf{G}_1$ and $\mathbf{G}_2$ with the same intended meaning. Also in this case, the model manager may suggest to lift this extension to the reference model, or in case of other views, to specify the extension by an abstract view instead.

Example. Figure 9 shows a derivation sequence that models the same operations of Figure 5 from the bank's point of view. The *common.newAccount* operation is a synchronised action of a customer and the clerk. The parallel *Customer.doTransaction* and *Customer.getBalance* operation is performed, while the clerk has an idle step. The second *Customer.doTransaction* and *Customer.getBalance* are complemented by two *Clerk.doTransaction* operations, that take over the formerly (in the customer's view) unspecified effects. Hence, the transitions of the system model are obtained in a compositional way by integrating the transitions of the two design views. △

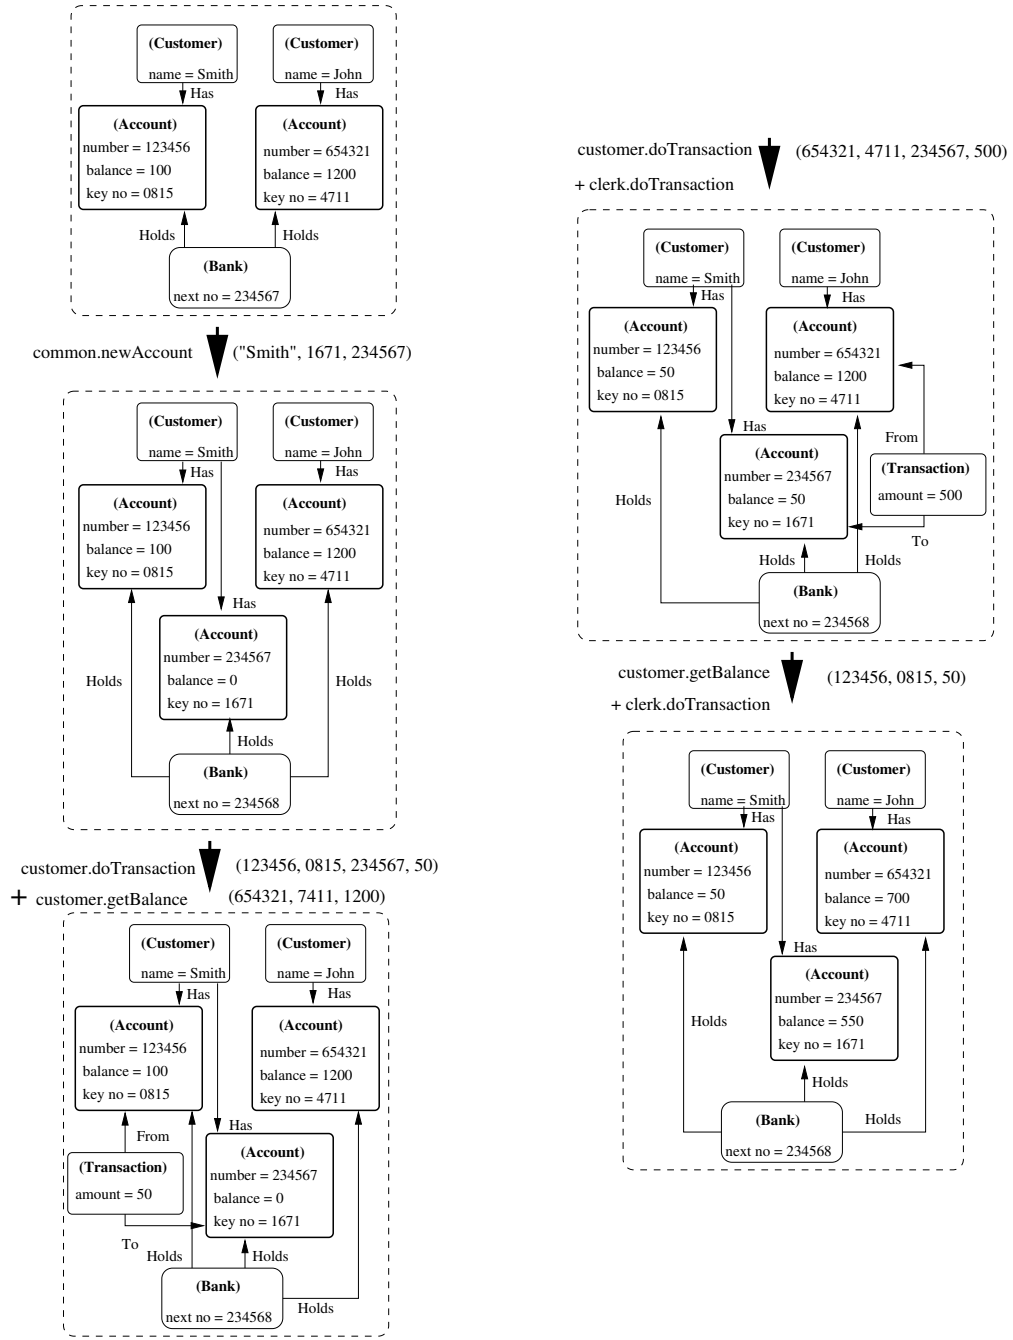

Fig. 9. Derivation sequence in the system view of the bank.

6 Conclusion

In this paper we have presented a view-oriented approach to concurrent system modelling with the following basic features:

- Separate viewpoints of a system are represented by different views sharing a common reference model.
- Each view is represented in an intuitive graphical way, which integrates static and dynamic aspects of the system, and has a sound theoretical base.
- The views are kept consistent by a model manager by extending the reference model whenever new dependencies occur in the development process.
- Using the reference model, consistent views can be integrated automatically.
- In case of more than two views, additional abstract views may have to be introduced leading to a hierarchy of views.
- The specification approach is based on typed graph transformation systems with a rich mathematical theory.
- In addition to the classical semantics of graph transformation systems based on derivation sequences, a new loose semantics is considered which is able to model an open behaviour of a view.

Despite the benefits mentioned above, our view-oriented approach has still to be fully exploited and systematically applied to examples of realistic size. Reference models for different application areas have to be developed. Moreover, views are usually more complex than flat graph transformation systems, i.e., we have to extend our approach by horizontal structuring techniques as considered in [HCEL96] and more powerful operations like rule-based transactions. Also, we have to consider different kinds of constraints, especially temporal logic formulas, in order to specify and verify important (e.g. safety critical) system properties (see [HEWC97]). Last but not least our view-oriented technique must be supported by a specification language based on typed graph transformation systems and corresponding tools. A good candidate for this is PROGRES [Sch91] where a related construction is already used for merging different versions of a document [Wes91]. The relationship of this construction and our approach to view integration has still to be investigated.

References

[AE96] M. Andries and G. Engels. A hybrid query language for the extended entity relationship model. *Journal of Visual Languages and Computing*, 7(3):321–352, 1996. Special Issue on Visual Query Systems.

[Bal96] R. Balzer. Current state and future perspectives of software process technology. Keynote Speech, Software Process (SP 96), Brighton, 2-6 December 1996.

[BLN86] C. Batini, M. Lenzerini, and S. Navathe. A comparative analysis of methodologies for database schema integration. *ACM Computing Surveys*, 18(4):323–364, 1986.

[CEL+96] A. Corradini, H. Ehrig, M. Löwe, U. Montanari, and J. Padberg. The category of typed graph grammars and their adjunction with categories of derivations. In *5th Int. Workshop on Graph Grammars and their Application to Computer Science, Williamsburg '94, LNCS 1073*, pages 56–74, 1996.

[Dat79] C.J. Date. *An Introduction to Database Systems.* Addison-Wesley, 1979.

[FKN+92] A. Finkelstein, J. Kramer, B. Nuseibeh, M. Goedicke, and L. Finkelstein. Viewpoints: A framework for integrating multiple perspectives in system development. *International Journal of Software Engineering and Knowledge Engineering*, 2(1):31–58, March 1992.

[HCEL96] R. Heckel, A. Corradini, H. Ehrig, and M. Löwe. Horizontal and vertical structuring of typed graph transformation systems. *Math. Struc. in Comp. Science*, 6(6):613–648, 1996. Also techn. report no 96-22, TU Berlin.

[HEWC97] R. Heckel, H. Ehrig, U. Wolter, and A. Corradini. Integrating the specification techniques of graph transformation and temporal logic. In *Proc. of MFCS'97, Bratislava*, 1997. To appear.

[Jac95] D. Jackson. Structuring Z specifications with views. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 4(4):365–389, October 1995.

[LB93] M. Löwe and M. Beyer. AGG — an implementation of algebraic graph rewriting. In *Proc. Fifth Int. Conf. Rewriting Techniques and Applications, LNCS 690*, pages 451–456. Springer Verlag, 1993.

[LRP+96] T. Lewis, L. Rosenstein, W. Pree, A. Weinand, E. Gamma, P. Calder, G. Andert, J. Vlissides, and K. Schmucker. *Object-oriented Application Frameworks.* Prentice-Hall, 1996.

[Nag96] M. Nagl, editor. *Building Tightly Integrated Software Development Environments: The IPSEN Approach.* LNCS 1170, Springer, 1996.

[Rib96] L. Ribeiro. *Parallel Composition and Unfolding Semantics of Graph Grammars.* PhD thesis, TU Berlin, 1996.

[Roz97] G. Rozenberg, editor. *Handbook of Graph Grammars and Computing by Graph Transformation, Volume 1: Foundations.* World Scientific, 1997.

[Sch91] A. Schürr. Progress: A vhl-language based on graph grammars. In *4th Int. Workshop on Graph Grammars and their Application to Computer Science, LNCS 532.* Springer, 1991.

[SWZ95] A. Schürr, A.J. Winter, and A. Zündorf. Graph grammar engineering with PROGRES. In W. Schäfer and P. Botella, editors, *5th European Software Engineering Conference (ESEC'95), Sitges*, pages 219–234. Springer LNCS 989, September 1995.

[Tae96] G. Taentzer. Modeling dynamic distributed object structures by graph transformation. *Object Currents*, 1(12), Dec. 1996. http://www.sigs.com/publications/docs/oc/9612/oc9701.f.taentzer.html.

[Wag97] A. Wagner. *A Formal Object Specification Technique Using Rule-Based Transformation of Partial Algebras.* PhD thesis, TU Berlin, 1997.

[Wes91] B. Westfechtel. Structure-oriented merging of revisions of software documents. In P. Feiler, editor, *Proc. 3rd Int. Workshop on Software Configuration Management (SCM3) New York*, pages 68–79. ACM Press, 1991.

[Zam96] A. Zamperoni. GRIDS - graph-based integrated development of software: Integrating different perspectives of software engineering. In *Proc. 18th International Conference on Software Engineering (ICSE)*, pages 48–59. IEEE CS Press, March 25-29 1996.