

© 1995 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

## Support for Collaborative, Integrated Software Development

John C. Grundy<sup>†</sup>, Warwick B. Mugridge<sup>††</sup>, John G. Hosking<sup>††</sup> and Robert W. Amor<sup>††</sup>

<sup>†</sup>Department of Computer Science  
University of Waikato  
Private Bag 3105, Hamilton, New Zealand  
jgrundy@cs.waikato.ac.nz

<sup>††</sup>Department of Computer Science  
University of Auckland  
Private Bag, Auckland, New Zealand

### Abstract

*A new model for supporting collaborative software development with shared, multiple textual and graphical views is presented. Multiple views of software development can be synchronously, semi-synchronously and asynchronously edited by different developers. View versions can be incrementally merged, and view updates broadcast to other developers and incrementally incorporated as required in their alternative versions. The model is illustrated by its use in a software development environment for an object-oriented language.*

### 1. Introduction

Software systems are growing ever larger and more complex. Two related approaches to managing this complexity are integrated software development environments (ISDEs) and programming environments which facilitate Computer Supported Cooperative Work (CSCW).

ISDEs support multiple tools, and multi-view editing in these environments allows developers to work with software components at different levels of abstraction, using different representations [10]. For example, analysis and design views might support graphical construction and representation of the high-level aspects of a program, while textual views might support detailed implementation. Consistency management is required to keep all of these views consistent under change.

Large software systems require the collaboration of multiple developers [9, 12]. Support for collaboration can be provided by two types of tool: version control systems, which allow alternative designs to be created

and merged asynchronously; and synchronous editors, which allow concurrent manipulation of a system by two or more collaborators [9]. Ideally environments should support both synchronous and asynchronous modes for all types of system data.

We describe a new model for constructing collaborative, multi-view ISDEs. Different environment tools and multiple views of the software under construction are integrated by broadcasting descriptions of changes on views via shared data repositories. Asynchronous collaborative software development is facilitated by sharing multiple versions of software components between developers. Both fine-grained and coarse-grained versions of software components and views are supported, with incremental version revision and merging. Synchronous collaboration, where developers interact with “what-you-see-is-what-I-see” views, and semi-synchronous collaboration, where developers are unobtrusively informed of changes other developers are making to views, are supported. These use real-time broadcasting of change descriptions between environments and incremental change description presentation and version merging.

Section 2 discusses related collaborative ISDE research. Section 3 describes the user’s perspective of our collaborative ISDE for object-oriented software development. Section 4 discusses the software architecture this environment is based on and Section 5 describes the implementation of this architecture as a reusable object-oriented framework. Section 6 summarises the contributions of this research and outlines possible future research directions.

## 2. Collaborative environments

SCCS [16] supports version control for text source code files, with different versions regenerated and merged asynchronously. This technique does not work well for more structured information, such as diagrams, and the (manual) version merging process can be tedious and error-prone.

Mercury [7] supports a restricted form of collaborative programming, with module interface changes broadcast between users. Program module versioning and multiple views are not supported. [12] describes a programming environment which supports manual, asynchronous version control for multiple developers working on a shared Smalltalk program. Mjølner/ORM [9] uses a fine-grained version control system to support both asynchronous and semi-synchronous editing. All of these environments provide only one textual, structure-edited view of program code.

Garden [14] supports software development via multiple textual and graphical views. Garden allows developers to share software components and their views via an object-oriented database. FIELD environments

[15] allow Unix tools to be integrated by selective broadcasting of editing changes. Neither support true collaborative, multi-view editing.

Collaborative document editors support synchronous, collaborative work on a shared document [1]. They do not usually support asynchronous editing and version control, as the users are assumed to be working on the same document. For example, ConversationBuilder [8] provides flexible, active support for cooperative work activities, and facilitates text-based tasks which can be constantly changing.

[11] aims to support lazy consistency management for cooperative software development, where changes to software are broadcast before the changes are committed. [2] supports multiple viewpoints for software development. Inconsistency management, where multiple views of software development remain inconsistent for some length of time, is supported via logical predicates describing inconsistencies.

Dora [13] provides multiple textual and graphical views of software development, but without support for propagating “partial” view updates between analysis, design and implementation views.

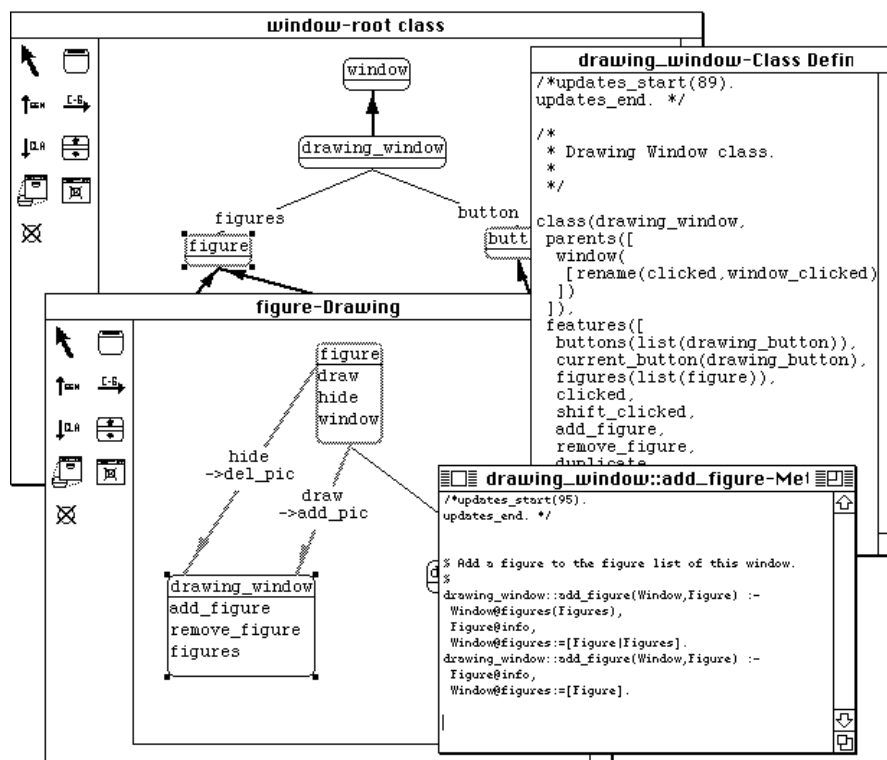


Figure 1. A screen dump from SPE.

## 3. C-SPE

### 3.1. SPE

SPE (Snart Programming Environment) is an ISDE which provides multiple textual and graphical views for

constructing programs in Snart, an object-oriented Prolog [4]. SPE supports integrated analysis, design, implementation, debugging and documentation tools. Figure 1 shows a screen dump from SPE with two graphical views (one for analysis and one for design), and two textual views (a class interface and a method implementation). There is full consistency management between all view types, so changes to one view are

always reflected in other views that share the updated information, no matter how loose the connection between the view representations.

Additional analysis and design views, such as class contract and documentation views, as well as graphical and textual debugging views, are also provided by SPE. SPE supports a rich set of view navigation facilities, utilising hypertext techniques.

Graphical views are interactively edited and are kept consistent with other views by the environment directly updating changed icons. Descriptions of changes affecting graphical view components can also be viewed in dialog boxes.

Textual views are free-edited and parsed. Textual view consistency involves expanding descriptions of changes, called *change descriptions*, into the view's text in a special header annotation. Some changes can be automatically applied by SPE to update the view's text, such as renaming classes and features and adding or deleting features. Other changes represent "partial" changes affecting the view (eg a design level change propagated to an implementation view) which must be implemented manually by the software developer.

Figure 2 shows an example of change descriptions being used to maintain textual view consistency. Users can select some change descriptions, such as those describing feature additions, deletions and renames, and have SPE automatically update the textual code. Other changes, such as the addition of a new client-supplier relationship or a semantic error, must be manually resolved [4].

C-SPE extends SPE to provide a collaborative object-oriented programming environment. SPE automatically generates descriptions of all view and software component updates as change descriptions. C-SPE uses these change descriptions to inform other developers of changes applied at the analysis, design and implementation levels. This is done by broadcasting these change descriptions as they occur to other developers' environments and either storing them in different, shared versions or presenting them to developers.

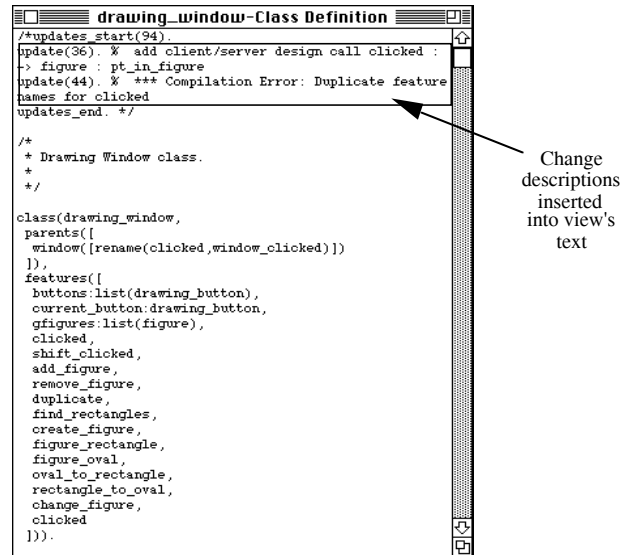


Figure 2. Textual view consistency.

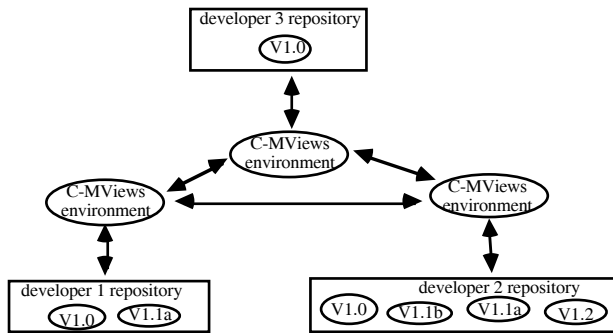
### 3.2. Asynchronous Collaboration

The loosest form of collaborative software development supported by C-SPE involves sharing multiple versions of views, classes and class methods among software developers. After independent update by different developers, alternate versions of the same view or software component may be merged to produce a new version.

Rather than the check-out style of SCCS, C-SPE adopts an optimistic approach to version control. This is especially appropriate when the software under development cannot be easily partitioned between developers [9, 12]. For example, the addition of a single function to an OO software system can lead to changes in several classes.

"Evolution graph" views, similar to those of [9], show the relationships between component and view versions. They are used to graphically specify new versions, alternate versions and alternate merges, and allow developers to view all change descriptions in different versions.

Developers edit their own versions of views and/or the software components rendered in the views asynchronously. Developers may create and modify new versions of a component based on their current version. This "freezes" the current version (ie new changes are locked out), and allows it to be exported for other developers to use. Another developer may subsequently import an exported (frozen) version, and merge it with his/her own version, exporting or further modifying the resulting version.



**Figure 3. Asynchronous view editing.**

Figure 3 illustrates this asynchronous editing and merging approach to collaborative development. Developer 1 copies version V1.0 of a component from developer 3 and creates a new version (denoted by V1.1a). Developer 3 also imports V1.0 and creates a new version (V1.1b), so they can modify it at the same time. After updating their alternatives, developers 1 and 2 freeze their component versions. Developer 2 then takes responsibility for integrating the changes, obtains a copy of version 1.1.a and merges it with version 1.1b. This merged component version is then exported as V1.2, for other developers to use.


Alternatives are merged by having C-SPE apply change descriptions selected from one alternative version to the other alternative's view or component, or reverting to a common ancestor version and applying all change descriptions from both alternatives to this earlier version. Developers can choose a subset of a version's change descriptions to have C-SPE apply. Unlike most other systems, developers can also request groups of change descriptions be applied out of the sequence they occurred, if desired. This is useful when merging two alternatives with large numbers of differences.

C-SPE incrementally applies the change descriptions that are to be merged incrementally into a view, updating the view as each change is applied to it. Developers can also step through each view update as it is made, or ask for a group of updates to be made together. This dynamic merge animation allows developers to more clearly identify the effects of a merge operation in terms of actual changes to views.

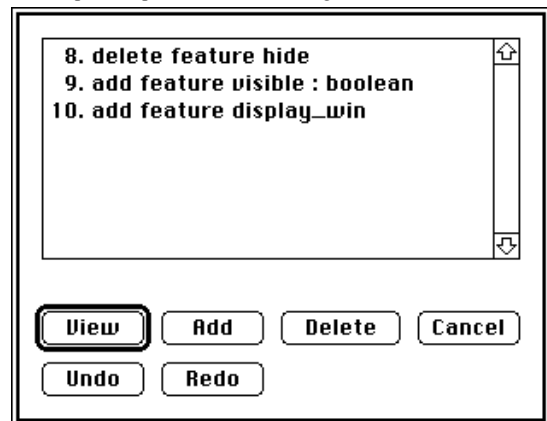
Two alternative versions may contain conflicts that must be resolved when merging them. For example, one version may delete something that the other updates. C-SPE identifies change descriptions it can't merge automatically and informs the developer of the conflict.

For example, in Figure 4 C-SPE must carry out both sets of change descriptions, but developer 2 has deleted feature "hide" while developer 1 has renamed it. This problem is identified by C-SPE and the merging developer informed of the conflict. The developer can then decide which update to allow (if either) or make other changes to reconcile the two alternatives. A similar problem occurs with the renaming of display to display\_win by developer 1 and the addition of

display\_win by developer 2 (a semantic error). This can be resolved if one of the updates is disallowed or the old display method is deleted.



Developer 1's updates on their drawing\_window class alternative



Developer 2's updates on their drawing\_window class alternative

**Figure 4. Merging example.**

C-SPE also identifies semantic errors caused by the merging process. These are presented to the merging developer as semantic error change descriptions, along with any structurally inconsistent merge conflicts. Figure 5 shows the Merge Conflicts dialogue. This shows changes that C-SPE could not carry out ("8" and "10"), and semantic error change descriptions ("6"), generated by merging the versions in Figure 4.

SPE supports free-edited textual views, stored as blocks of text, which are parsed to update information shared by other views. C-SPE supports multiple versions of these text components, but only stores change descriptions generated by the parsing process. Other aspects of the views which are updated between two versions of the same component, such as comments, expressions or code statements, need to be reconciled manually or by using a traditional SCCS-style textual differencing approach.

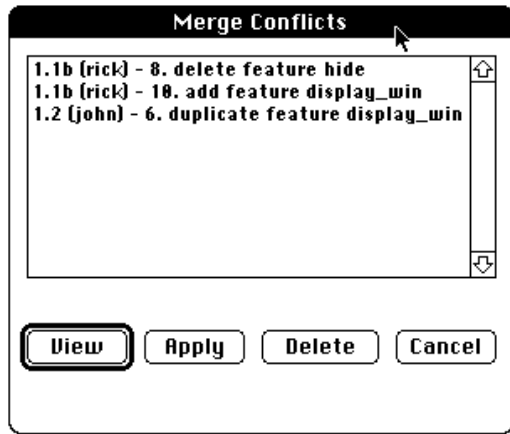


Figure 5. Merge conflicts example.

### 3.3. Synchronous Collaboration

Synchronous collaboration allows two or more developers to simultaneously examine and alter a view, communicating the changes they make between them as they occur. There are two main approaches to handling the integration of the changes made. The first approach is to consider that the developers are communicating and negotiating in order to derive a single result. In this case, it is appropriate that all the developers concerned share a common view, so that a change can be rejected by any participant and thus undone in all of the shared views.

The second approach does not aim for a single result, so that developers may end up with their own distinct versions that reflect their current thinking. In this case, each developer can choose whether or not to accept the changes of others in the collaboration without affecting the others. This results in “semi-synchronous” collaborative editing.

C-SPE supports both types of collaboration. With semi-synchronous collaboration, developers have their own alternative and are incrementally informed of updates other developers are making to their alternative versions. This is done by the receipt of change descriptions, which are then displayed in dialogs or textual views. With synchronous collaboration, developers share the same version and view updates are shown simultaneously in other developers’ views.

Figure 6 illustrates semi-synchronous view editing in C-SPE. The dialog shown allows a developer to view change descriptions of the changes before deciding to apply them, ignore the changes, or view the effect of (some or all of) the changes on their version. Developers can request C-SPE to automatically merge received change descriptions with their current alternative as they arrive. Change descriptions can also be viewed in textual view headers, and some can be selected and automatically applied to the view text by C-SPE.

During synchronous collaboration, developers share the same version of a view. Changes made by one developer are reflected immediately in the view the other developers interact with. Whereas view alternatives edited semi-synchronously may have different layouts and viewed components, synchronously edited views are always identical among all collaborating developers.

Developers can move between semi-synchronous and asynchronous development at will. To move from synchronous to asynchronous collaboration, however, developers must obtain the shared version, and create an alternative of it, before asynchronously editing it.

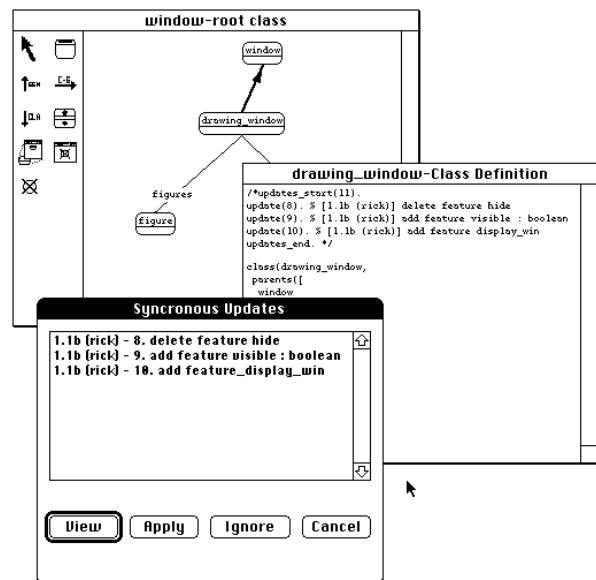


Figure 6. Semi-synchronous editing.

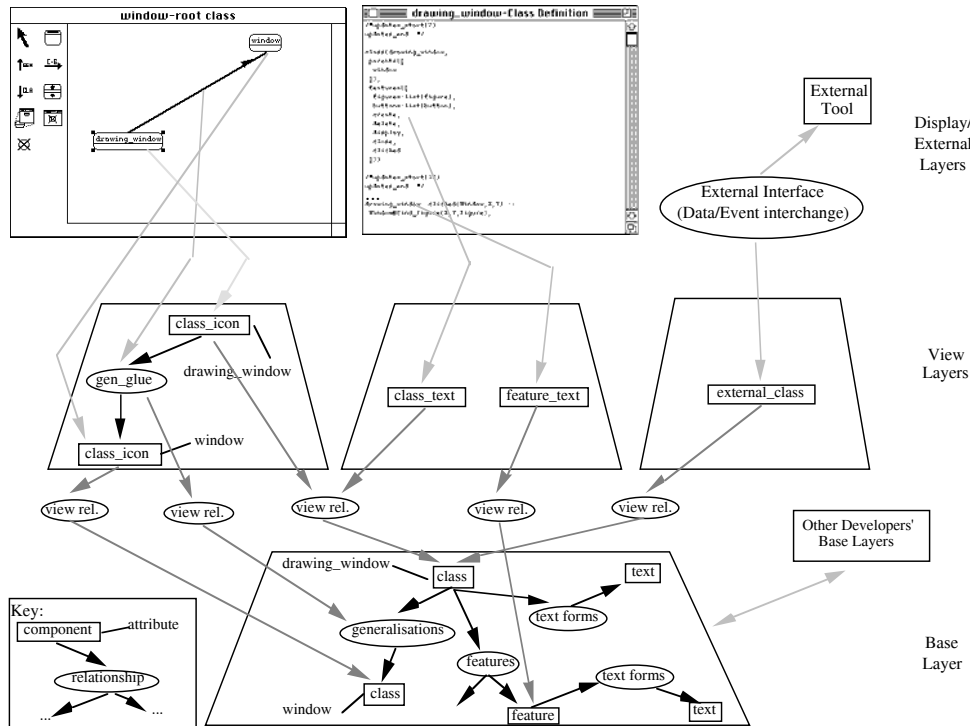


Figure 7. The MViews Architecture.

## 4. C-MViews

We now describe C-MViews, the framework used to construct C-SPE, commencing with a description of the single user MViews framework, and following with extensions to support collaborative environments. Only the MViews framework was modified to produce C-MViews, a framework for building collaborative ISDE. No aspect of SPE was modified to produce C-SPE.

### 4.1. MViews

SPE is implemented as a collection of Smart classes, specialised from the MViews framework [3]. MViews supports the construction of new ISDEs by providing a general model for defining software system data structures and tool views, with a flexible mechanism for propagating changes between software components, views and distinct software development tools.

As shown in Figure 7, ISDE data is described as *components* with *attributes*, linked by a variety of *relationships*. Multiple views are supported by representing each view as a graph linked to the base software system graph structure. Each view is rendered and edited in either a graphical or textual form. Distinct environment tools can be interfaced at the view level (as editors), via external view translators, or multiple base layers may be connected via inter-view relationships, as described in [5].

When a software or view component is updated, a change description is generated. This is of the form UpdateKind(UpdatedComponent, ...UpdateKind-specific Values...). For

example, an attribute update on Comp1 of attribute Name is represented as: update(Comp1, Name, OldValue, NewValue)

All basic graph editing operations generate change descriptions and pass them to the propagation system. Change descriptions are propagated to all related components that are dependent upon the updated component's state. Dependents interpret these change descriptions and possibly modify their own state, producing further change descriptions.

Figure 8 shows an example of change propagation between MViews components used in the construction of SPE: 1) a view component is edited, applying operations to the view component; 2) change descriptions are generated and propagated to all of the view component's relationships; 3) the view relationship translates view component changes into operations on its base component, if the base component is affected by the view change; 4) the base component operations generate change descriptions; 5) these change descriptions are propagated to the base component's relationships; 6) the view relationships translate the base component change descriptions into operations on their view components, if the view components are affected by the change; 7) the updated view components re-render their displays. An external view component sends change messages to an external tool or translates the changed data into changed external tool data.

This change description mechanism supports a diverse range of software development environment facilities, including semantic attribute recalculation, multiple views of a component, flexible, bi-directional textual and graphical view consistency management, a

generic undo/redo mechanism, and component “modification history” information [6].

New software components and editing tools are constructed by reusing abstractions provided by an object-oriented framework. ISDE developers specialise MViews classes to define software components, views and editing tools to produce the new environment. A persistent object store is used to store component and view data.

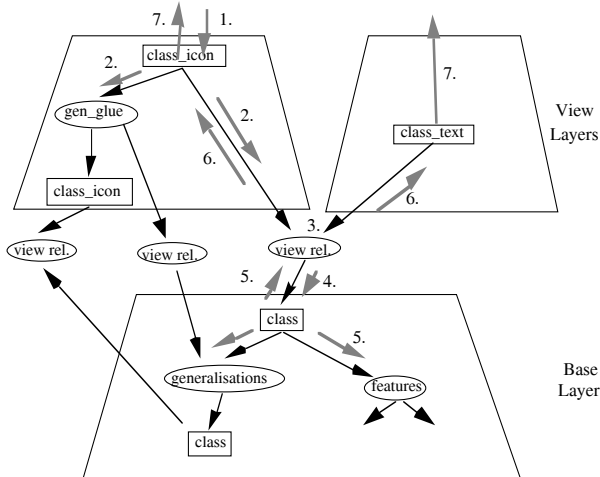


Figure 8. MViews change propagation.

#### 4.2. Multiple Versions

Any component operation generates a change description in MViews. As these change descriptions are first-class objects, they can not only be propagated between different components, but can be stored and broadcast between different developers’ environments. This is the basis for supporting collaborative software development facilities in C-MViews.

Software system components usually have a natural hierarchy, with some components being “composed of” other (sub-)components. For example, an object-oriented program in SPE is made up of several class frameworks, a framework is composed of several classes and class relationships, and a class is composed of various features and inter-class relationships.

Several approaches to managing hierarchical versioning exist: a new version of the whole system can be created whenever any change is made; individual version numbers at a particular level of the component hierarchy can be maintained; or individual versions for any component in the hierarchy can be stored, with a configuration management tool used to reconstruct a system version from lower-level sub-component versions.

C-MViews aims to support all of these approaches by providing a tailorable low-level versioning mechanism, based on stored sequences of change descriptions, called *version records*. Version records can be associated with any C-MViews component or view, and contain a record of changes made to that component

(as change descriptions) since the previous version. These can include change descriptions describing changes to the component itself, its sub-components, or the configuration (version used) of sub-components.

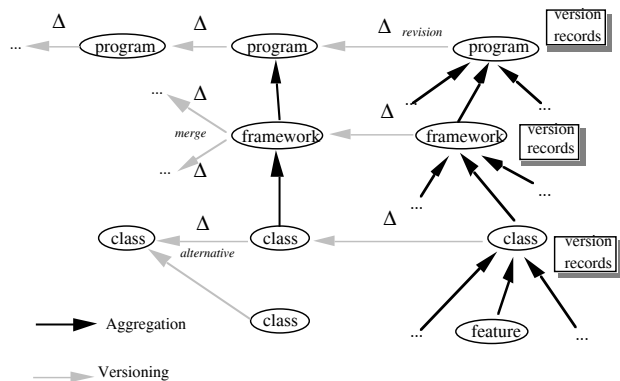


Figure 9. Example of C-SPE versions.

For example, Figure 9 shows how C-MViews version records are used for C-SPE’s version control. C-SPE adopts a component-level versioning approach where each component in the component hierarchy down to the level of individual classes has multiple versions with associated version records.

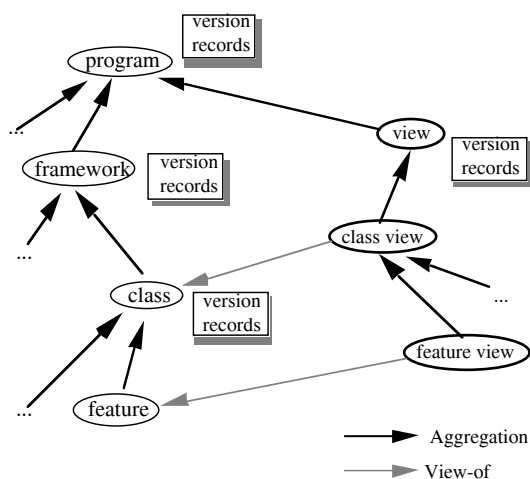
The most recent version of a C-SPE component is either frozen, and hence closed to change, or open, and hence able to be further modified. When a new C-SPE component version is created a new version record for that component is also created. The component’s parent, eg the feature’s class, is also notified of this new sub-component version. If the parent’s version is frozen, a new version of the parent is also created.

For small sub-components it is useful to record changes only in their parents’ version records, for efficiency. For example, in C-SPE, version records are associated with classes but not individual class features. This reduces the number of version records needed, but means versioning only proceeds down to the class level, with no individual feature versions. C-MViews aggregation (part-of) relationships between components automatically propagate a sub-component (eg class feature) change description to its parent component (eg class) which can then store the change description in its version record.

#### 4.3. View Versioning

An important distinction between C-MViews and other ISDE models is its support for view versioning. View versions are kept separate from base component versions, as a view may render several different base components. Changing any base component will thus partially change the view (and vice versa). Views may also change independently of their base components. For example, layout information is view-specific. Figure 10 shows view versioning as used by C-SPE. In this case a single version record is held for the view as a whole

rather than having individual records for view components, such as class icons and connectors.



**Figure 10. C-SPE view versions.**

Views complicate the versioning process, as each view can have multiple versions and each base software component rendered in the view can also have multiple versions. Changing the current version of a view will modify affected base components, which should also cause modifications to other views that render these base components. This can result in large scale changes from simple view version merging or when switching between a previous or subsequent view version.

Merging of two alternatives needs to resolve layout and composition conflicts (which often occur) with underlying base information conflicts (which occur less often). C-MViews currently presents the developer with a list of all conflicting change descriptions for manual resolution.

#### 4.4. Version Merging

Change descriptions stored in a version record, including change descriptions specifying changes to sub-component configurations, are used as deltas to (re)generate previous or subsequent versions. In order to convert one version to another, C-MViews will undo the change descriptions (to go back a version) or apply them (to go forward a version). Previous or subsequent versions may also be cached for more efficient configuration management.

C-MViews based environments, such as C-SPE, provide user interface facilities for capturing information about why a change is made, and to present this information appropriately in views. The capture and presentation of this information needs to be of limited “interference”, as developers typically do not want to supply or see all of it every time they make a minor view modification. C-MViews attempts to overcome some of this interference by allowing an application to capture and present this information on demand via the evolution graph dialogs used to browse version records.

During version merging, any structural update conflicts encountered during merging (such as deleting

components updated in another version) are identified and presented to the merging developer. Semantic errors can be detected by incrementally reevaluating semantic attribute values and constraints after each change description merge. Any “error” change descriptions generated can be presented to the developer, as they indicate a semantic merge conflict has occurred.

View update animation during merging is achieved by updating and re-rendering view components for each change description merged. Developers can choose a sequence of change descriptions to merge with another version, and can step through the application of each change as it is applied, seeing the view dynamically updated. Unlike most other systems, developers can even merge version updates out-of-sequence, with C-MViews detecting any structural or semantic conflicts that this produces.

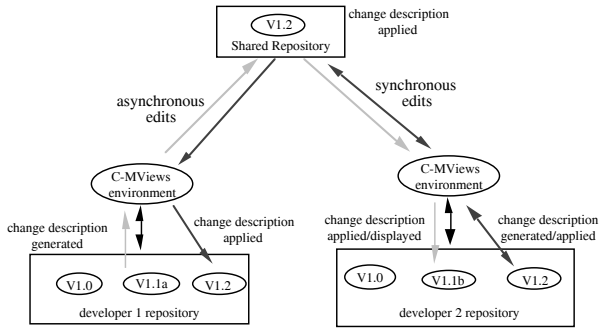
#### 4.5. Synchronous Collaboration

Synchronous and semi-synchronous collaboration are supported by broadcasting change descriptions to other developers’ environments as they are generated, as shown in Figure 11. As all MViews component operations generate change descriptions, C-MViews utilises these to support collaborative software development by broadcasting them to other developers’ environments, as shown in figure 11.

To support semi-synchronous editing, broadcast change descriptions are received by other developers and cached in special version records. They are then presented in dialogs or textual view headers to inform other developers of changes made to other alternatives. These change descriptions can then be incrementally merged with other developers’ view alternatives using the same merging techniques as employed for asynchronous merging.

To support synchronous editing, no developer “owns” the shared version of a view. All updates attempted by developers are sent to a central server which then updates the shared version itself. Generated change description(s) are then broadcast to all collaborating developers’ environments, whose views are then re-rendered to reflect the change. Fine-grained locking is maintained by the server so only one edit of the same view component is accepted at a time. On receiving an update request from a developer, the view component is modified. All other developers’ environments are then informed of the change and their views re-rendered, before the next update on this view component is accepted from any developer.





**Figure 11. Synchronous development.**

#### 4.6. Change Description Broadcasting

Change description broadcasting is handled as part of the general C-MViews mechanism of propagating change descriptions. The current version record for a component records other developers' interest in updates to the component. When a change description is stored in this version record, it is also broadcast to the environments of these other, interested developers. Shared version records used for synchronous editing also co-ordinate updates made to the shared version.

Broadcasted change descriptions are either stored by the current version record of the other developer's component or applied immediately to shared versions. Stored updates are presented to these developers when they work on the version record's component or view. An indication of new change descriptions is usually given by shading icons or changing a menu bar item which results in a context-dependent communication mechanism. This is important in making C-MViews environments useable, so developers aren't inundated with messages at inappropriate times.

Broadcasted change descriptions include who, when and optional why information, which assists developers in understanding why the changes have been made. User-defined change descriptions can also be broadcast to facilitate flexible, context-dependent communication.

C-MViews also timestamps each broadcast change description by attaching a version record ID and unique sequence number. They are then stored in special "broadcasted" and "received" version records in each developers' environment. Thus, no matter whether semi-synchronous view editing is switched on or off for an alternative, changes the developer is not notified about can still be incrementally merged at a later date, using the timestamp information.

This also helps to support fault-tolerance, eg when one developer's network connection or machine fails. Change descriptions cached by the developer's environment and the central server can be rebroadcast when the connection is re-established. Failure of the central server prevents any form of synchronous collaboration from taking place, but as developers have their own alternatives, asynchronous development can continue. We plan to extend C-MViews to support

synchronous collaboration between developers without using the central server, which will improve the robustness of resulting environments.

## 5. Implementation

A prototype of C-MViews has been implemented in Snart and has been used to construct the prototype C-SPE environment. C-MViews extends persistent Snart object stores to support multiple versions of an object, for multiple component versions. Change descriptions are stored in these version records, and include additional information, such as time-stamp, user id, and change reason. Each version record object contains a sequence of change descriptions together with links to its predecessor(s) and successor(s), giving an evolution graph for the component.

C-MViews currently uses a common, shared component repository as a shared object store, and high-performance, single-user repositories. A database server is provided to moderate access to the shared object store. This allows a group of collaborating environments to provide high-speed data storage for each developer's alternatives, supports sharing of these alternatives, and handles change description broadcasting between developers.

A component alternative in one developer's object store may be merged with a version in another's object store. Thus a component's object, its sub-component objects, and its version objects must be copied from one object store to another. As C-MViews knows about the aggregation structures present between software components, it can import and export the sub-components of a component automatically.

The shared repository acts as a form of distributed database by ensuring objects created in any developer's object space are always unique. When editing different alternatives two developers may create different objects which represent the same conceptual view or base component. Subsequent merging of these alternatives will result in redundancy that can be resolved by discarding one object in favour of the other.

Our current C-SPE prototype detects structural and semantic conflicts as they occur during the merging process, and presents invalid or error change descriptions to the merging developer. C-MViews currently does not, however, give developers any assistance in rearranging updates to resolve conflicts). C-SPE currently lacks a mechanism for relating different component versions. For example, if changes are made to several classes to implement one new system feature, these version relationships are not documented. It is thus difficult for developers to trace between related updates to different classes and frameworks.

## 6. Conclusions

Our experience in developing integrated software development environments indicates that multiple views

of software development, integrated development tools, and collaborative software development are important when building large software systems. C-MViews supports multiple versions of software components and their textual and graphical views. Tools are integrated at the view or data repository levels. Collaborative development is via asynchronous, semi-synchronous and synchronous editing of multiple views. Semi-automatic version merging is supported, including incremental version merging. Two forms of synchronous collaboration are supported by broadcasting change descriptions between developers' environments.

Our experience with C-SPE suggests that asynchronous development is most useful for low-level design and implementation views, or when major system changes are being carried out. Semi-synchronous development is useful for higher-level collaboration where different alternatives are maintained by each developer. Synchronous development seems most appropriate when high-level designs are being worked on and alternatives are not desired during collaboration.

We are currently working on several improvements to C-MViews and C-SPE. This includes determining how changes broadcast via synchronous and semi-synchronous editing can be most usefully presented to developers. The capture of extra information, particularly a description of why a change was made, usually occurs above the change description level but below the version level. We are experimenting with "development tasks" for each developer and for groups of developers. These associate groups of related change descriptions, and relate change descriptions on different components. This is particularly useful for object-oriented systems, where changes are often made to several classes to provide one new system function. Partial consistency management between higher-level software components, as done for multiple view consistency, would also assist in supporting programming-in-the-large.

## Acknowledgements

The helpful comments of the anonymous reviewers on the draft of this paper are gratefully acknowledged.

## References

- [1] Ellis, C.A., Gibbs, S.J., and Rein, G.L., "Groupware: Some Issues and Experiences," *Communications of the ACM*, **34** (1), 1991.
- [2] Finkelstein, A., Gabbay, D., Hunter, A., Kramer, J., and Nuseibeh, B., "Inconsistency Handling in Multiperspective Specifications," *IEEE Transactions on Software Engineering*, **2** (8), 1994, 569-578.
- [3] Grundy, J.C. and Hosking, J.G., "A framework for building visual programming environments," in *Proceedings of the 1993 IEEE Symposium on Visual Languages*, IEEE Computer Society Press, 1993, 220-224.
- [4] Grundy, J.C., Hosking, J.G., Fenwick, S., and Mugridge, W.B., Chapter 11 in *Visual Object-Oriented Programming*, Burnett, M., Goldberg, A., Lewis, T. Eds, Manning/Prentice-Hall, 1994.
- [5] Grundy, J.C. and Venable, J., "Integrating CASE tools via hierarchical data model integration," Working Paper, Department of Computer Science, University of Waikato, 1994.
- [6] Grundy, J.C. and Hosking, J.G., "Constructing Integrated Software Development Environments with Dependency Graphs," Working Paper, Department of Computer Science, University of Waikato, 1994.
- [7] Kaiser, G.E., Kaplan, S.M., and Micallef, J., "Multiuser, Distributed Language-Based Environments," *IEEE Software*, **4** (11), 1987, 58-67.
- [8] Kaplan, S.M., Tolone, W.J., Bogia, D.P., and Bignoli, C., "Flexible, Active Support for Collaborative Work with ConversationBuilder," in *1992 ACM Conference on Computer-Supported Cooperative Work*, ACM Press, 1992, 378-385.
- [9] Magnusson, B., Asklund, U., and Minör, S., "Fine-grained Revision Control for Collaborative Software Development," in *Proceedings of the 1993 ACM SIGSOFT Conference on the Foundations of Software Engineering*, Los Angeles CA, December 1993, 7-10.
- [10] Meyers, S., "Difficulties in Integrating Multiview Editing Environments," *IEEE Software*, **8** (1), 1991, 49-57.

- [11] Narayanaswamy, K. and Goldman, N., "Lazy consistency: a basis for cooperative software development," in *1992 ACM Conference on Computer-Supported Cooperative Work*, ACM Press, 1992, 257-264.
- [12] Nascimento, C. and Dollimore, J., "A model for co-operative object-oriented programming," *IEE Software Engineering Journal*, **8** (1), January 1993, 41-48.
- [13] Ratcliffe, M., Wang, C., Gautier, R.J., and B.R., W., "Dora - a structure oriented environment generator," *IEE Software Engineering Journal*, **7** (3), 1992, 184-190.
- [14] Reiss, S.P., "GARDEN Tools: Support for Graphical Programming," in *Proceedings of Advanced Programming Environments*, Lecture Notes in Computer Science #244, Springer-Verlag, Trondheim, Norway, June 1986, 59-72.
- [15] Reiss, S.P., "Connecting Tools Using Message Passing in the Field Environment," *IEEE Software*, **7** (7), 1990, 57-66.
- [16] Roekind, M.J., "The Source Code Control System," *IEEE Transactions on Software Engineering*, **1** (4), 1975, 364-370.