

VC² - Providing Awareness in Off-The-Shelf Version Control Systems

Daniel Machado¹

Nuno Preguiça²

Carlos Baquero¹

J. Legatheaux Martins²

¹ DI/CCTC, Universidade do Minho

² CITI/DI, FCT, Universidade Nova de Lisboa

ABSTRACT

Version control systems have been used to help groups of people working at the same or distributed sites to cooperatively create documents. In particular, these systems are very popular in distributed collaborative software development. However, even using these systems, users often perform concurrent changes that require manual conflict resolution. Important causes for this situation are the lack of mutual awareness and coordination, among developers, and reluctance to commit unstable modifications. The paper addresses this problem by providing a tool that integrates with off-the-shelf version control systems and monitors filesystem accesses to relevant files in order to enhance the awareness among developers. With VC² users can be aware of uncommitted changes made by remote users; receive request to commit their own changes; be advised to update their local versions. While the final decision is always under user control, the team is made aware of the level of risk when delaying commits and updates.

INTRODUCTION

The development of team projects often requires the use of tools to aid the coordination and synchronization of work between members. There is a subgroup of these tools known as Version Control (or Revision Control) Systems, many of them are widely established and used, specially in software development projects. These systems divide in two main subgroups: those using the client-server model (CVS, SVN, ...) and those using the distributed model (Bazaar, SVK, BitKeeper, ...).

In the client-server model, the files of a project are located on a server repository. Each user can checkout her local copy from the repository, work locally in the copy and submit changes back to the server. At any time she can also check the server for new versions of files submitted by other users.

In the distributed model each user has her own local repository with a copy of the project. She may submit changes to her local repository, and at any time, two or more users may choose to merge their local copies.

Motivation

While working in a cooperative project under version control tools, users end up concurrently updating the same files. In a client-server system, one of the users

will commit first to the central server and the second one will have to incorporate these changes into her files before committing successfully. Version control systems have mechanisms to automatically merge these concurrent changes when different areas of the files are modified. When the same areas of the files are modified, the user has to manually solve the detected conflicts before being able to commit her changes to the central server.

The sole usage of version control tools is not enough to avoid this problem. In typical usage scenarios, these situations occur with some frequency. However, solving conflicts may be problematic, specially if the users have made a big number of changes that have not been committed for a long time. These problems may lead users to be reluctant to engage in cooperative projects and even to avoid parallel development [8].

To minimize these problems, users could follow some “best-practices” guidelines, such as regularly keeping their local copies up-to-date and committing as soon as possible, as this approach decreases the probability of conflicts [5]. The problem is that these guidelines are often contrary to good working practices for cooperative projects – e.g. users may only want to commit changes that leave the project in a consistent state, therefore delaying their commits for long periods of time such as days or weeks [13].

The goal of this work is to provide awareness of other users’ activities between synchronization points, helping users to coordinate their work and avoid the need for time consuming conflict resolution tasks. Unlike previous works, where awareness is provided in an application that is different from the application used to edit the shared documents [3, 12] or that require specific editors [14, 11], our solution allows users to continue using their preferred unmodified editors. Awareness information is provided as users access files with any application. Moreover, our solution also does not require any additional infrastructure, building on the existing version control system to propagate the needed information.

The remainder of this paper is organized as follows: Related work and the semantics of version control systems are reviewed in the next sections. Section 4 presents the proposed architecture and Section 5 illustrates a typical use case. Implementation details and evaluation are

reviewed in Section 6. Discussion and future work in Section 7 is followed by Conclusions.

RELATED WORK

Awareness information has long been identified as important for the success of cooperative activities by providing users with an understanding of other users' activities [6]. In groupware systems, a large number of awareness tools have been proposed for synchronous collaboration, such as multi-user scrollbars [1], telepointers and radar views [9], remote screen view [17], etc. These tools allow a user to have some information about the current activities of other users, but they are not appropriate for asynchronous collaboration.

Version control systems (e.g. CVS [4], SVN [16] and Bazaar [2]) provide the basic support for team projects, by maintaining and controlling the evolution of file versions for the project files. Basic support for awareness is available, with users being able to include comments when they commit their changes. CVS also includes an additional mechanism that can be used to provide awareness information: CVS watches. When using CVS watches, users must announce their intents of modifying a file beforehand (by executing a special command). Users can register their interest on specific files and be notified by email when someone announces the intent of modifying it. Some tools, such as Eclipse's Team CVS, automatically issues the required CVS commands when the user starts editing a file and it allows a user to check which users are modifying some file.

The BSCW system [3] is a web-based system that includes a version control system to manage shared files. In this system, when some action (check-out, check-in) is executed an event is recorded. The system can present a list of recent event to users when they connect to the system. In [7], the authors introduce a tool for integrating notification and chat with the CVS system. In this system, users are informed when some user commits changes to a file, with events being propagated using an event-dissemination system. Unlike our system, these systems presents no awareness about modifications before they are committed.

In State Treemap [12], the authors propose an awareness widget that allows users to visualize which files are being concurrently modified (leading to a potential conflict) and which locally modified files have already been committed (leading to a conflict). This widget has been integrated in a platform for supporting virtual teams of architects. The Palantír [14] system provides similar information for files stored in version control systems, relying on an event notification system for propagating information among users. The authors have created wrappers for SVN, RCS and CVS, with events being propagated when edit/update/commit commands are executed. The authors have also developed a plug-in for Eclipse, allowing awareness information to be presented in Eclipse. The Jazz [10] system also provides

similar information, as an extension to Eclipse.

In Gasper [11], the authors propose a generic mechanism for propagating limited information about changes being performed. Awareness information is provided in the editors in the form of annotations — e.g. if some user is modifying a method in a code file, other users could see an annotation about this fact in their user interface.

Several other systems have been designed for providing awareness information in the context of collaborative software development (see [15] for a survey). These systems can be divided in two groups. The first (including basic CVS watches, State Treemap and Palantír) requires the use of an additional tool for checking the awareness information. Besides the problem of convincing users to use an additional tool, this approach has the drawback of requiring users to explicitly check for awareness information when they start editing the shared files (as there is no connection between the editing activity and the tool that provides the awareness information). The second group (including Eclipse's Team CVS watches, Palantír, Jazz and Gasper) provides the awareness information in the context of a specific editor. This approach is interesting but it forces users to use a specific editor. Additionally, it requires the plug-ins to be updated when a new version of the editor is released. In our work, we provide awareness information for any editor, thus allowing users to continue using their preferred editors. Additionally, unlike previous works that require specific support from the version control systems or that rely on an additional infrastructure, we propagate the required information using files stored in the version control system. Thus, our approach can be easily deployed with existing version control system.

THE SEMANTICS OF VERSION CONTROL

In this section we present an overall view of how version control works. At this point we will only consider the centralized model (CVS, SVN, ...) and leave the distributed model as future work. In this presentation, to avoid too much complexity we will not consider the adding and removing of files and folders from the repository and the conflicts that may arise from these operations.

In version control systems, a user must start by checking out files that are stored in the repository, thus creating a local private copy. At any moment, the user may update her local copy against the latest version stored in the repository. The user may also modify her local copy, one or more times. After modifying her files, the user may commit her changes to the repository.

There are two situations that may lead to a conflict if the user modifies her local copy:

- The user modifies a file that was remotely modified

by some user that has not yet committed her changes to the repository.

- The user modifies a file that is outdated against the current version on the repository.

Conversely, conflicts may also arise if a remote user modifies an old version of a file that has been locally modified, either it has already been committed or not.

In order to help the user avoiding the situations that may lead to conflicts, our tool will provide advance warning about conflict-leading actions. Thus, when a users starts accessing a file that has been modified elsewhere, a notification will be presented to the user suggesting an alternative action. For example, if the user accesses a file that is being concurrently modified by some other user, the user is notified of the fact and may ask the other user to update the file.

SYSTEM ARCHITECTURE

Approach

The Version Control Control tool (VC^2) provides awareness for cooperative editing activity by combining filesystem probing and activity dissemination among the group using metafiles stored in the version control system.

Filesystem probing can be achieved by adding code to kernel filesystem routines, with kernel tapping mechanisms such as FUSE [18] and FIST [20], or by subscribing to filesystem events on kernels that support them, e.g. INOTIFY[19]. Since VC^2 notifications do not require blocking or delaying file operations, kernel access is not mandatory and event subscription can suffice. This also means that users are warned of actions that can lead to conflicts but are free to ignore those warnings and have full accesses to the files.

A second aspect concerns the diffusion of information within the developer's group. While other systems choose to use external mechanisms of dissemination or tailor their own dissemination service, in VC^2 we choose not to add any additional mechanism. This is possible because an existing communication mechanism is already available in the version control mechanism.

VC^2 uses the existing version control system to disseminate additional control information across the hosts. It suffices to enclose this information as special metafiles in the project tree and exchange them with the existing commit and update functionality. This also allows reuse of authenticated channels, say SSH, that might already be prepared for communication with the project server. Finally, this approach provides awareness information even if users are not connected at the same time, because the control information is maintained on the server. For example, if a user has started to change her local copy of the file but then she has stopped working and disconnected her computer, the information about this activity is still stored in the server. If another user

(that could have been disconnected) starts accessing the file, he will be notified of the potential problem.

Architecture

The typical setup, when using the tool with client-server version control, comprises a CVS/SVN server in a hosting server machine and two or more users in a given number of other machines – we will refer to these as clients. Client machines must have connectivity and access to the server, either by SSH or specific CVS/SVN ports. There is no need for connectivity among client machines. There is no impact even if they are all hidden by firewalls.

Any client machine will just have to support standard CVS/SVN client tools and, in addition, have installed local support for either FUSE or INOTIFY. The VC^2 approach tolerates the presence of standard CVS/SVN clients that cannot introspect the filesystem and run the tool. As expected, activity on those clients will not be made aware to others and vice-versa.

Each client machine runs a VC^2 daemon that acts on filesystems events, creating separate threads for any event that requires GUI interaction or communication to the server.

The daemon is implemented as a Java user level process. It makes use of a filesystem notification layer that can be provided by either FUSE probes or INOTIFY events. Once aware of all system calls to the filesystem, our tool can check if the accessed files are under control of any version control system, and try to detect behaviors that may potentially lead to conflict scenarios. If such a case is detected it will alert the user suggesting a recommended action to take.

To start using our tool it suffices to have a working setup of the supported version control systems and start the Java VC^2 daemon. At this point the user can check-out her projects to directories inside a controlled local filesystem. After that, every time she opens or closes a file, the system will check if that file is under control of the version control system. All other files are ignored and are transparent to VC^2 .

If a file is under control, our tool will check its status and alert the user when necessary. The status of the file is kept in a metafile which is saved in the repository. The first time the tool checks the status of a file, it creates the associated metafile if none exists. Each time the system consults or changes a metafile, it will update it from the repository and if there are changes it will commit immediately.

At the moment, the only information on the metafile is the number of users with uncommitted changes, and the number of users requesting the commit of those changes. Later in this paper, we will elaborate on additional information that could be used.

When opening a file, our tool will consult the metafile to check for uncommitted changes made by other users. If there are any, it will alert the user of the situation and ask if she wants to request the updating of that file. Also, if the local copy of the file is outdated compared to the version on the repository, it will alert the user and ask if she wants to update to that version.

When closing a file, it will check if there were any changes made, by checking the status against the repository version. If the file was locally changed it will increment the number of changers on the metafile. The file will also be added to a list of uncommitted files, that is regularly checked for requests.

User alerts are implemented by dialog boxes that pop up on the screen, usually with a yes/no question.

As it was described, there is no need for server side intervention on CVS/SVN platforms, allowing the use of any off-the-shelf public server. Additionally, as it relies only on common version control management, it should be immediate to include support for other version control system.

USAGE EXAMPLE

In this section we show a small example of how the tool is used and how it helps the coordination of a team project.

Suppose there are three developers working on an AddressBook application in C++. Alice is developing the user interface, Bob is working mostly on the core of the application, and Charlie is developing a storage module to save application data. The structure of the project is the following:

```
AddressBook/inc/core.h
AddressBook/inc/storage.h
AddressBook/inc/UI.h
AddressBook/src/core.cpp
AddressBook/src/storage.cpp
AddressBook/src/UI.cpp
AddressBook/help.html
```

The developers start by checking out the project from the CVS repository into a local directory, and then begin working on it. Suppose the following action are executed:

1. Bob changes *core.h* and *core.cpp*
2. Alice which has been working on *UI.cpp* now needs to know how to retrieve phone numbers and consults *core.h*. As she opens the file, a popup alerts her that there is one user with uncommitted changes on that file, asking if she wants to send an update request (Figure 1). She agrees.
3. A popup appears on Bob's screen saying that one user is requesting an update on *core.h*, and asking if

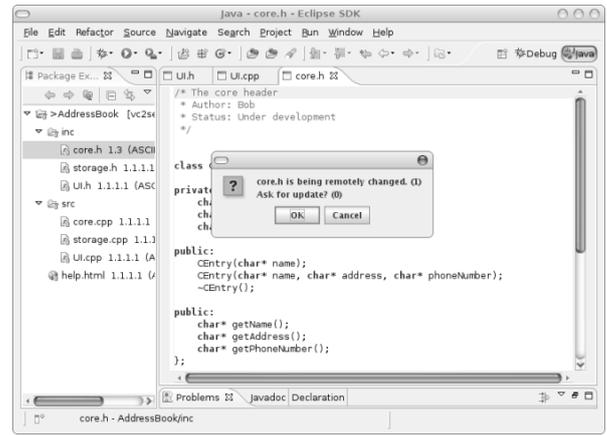


Figure 1. Alice being alerted that the file she wants to access has been modified by some other user. In this case, Alice is using the Eclipse IDE.

he wants to commit (Figure 2). Bob agrees and the file is committed.

4. Alice, who decided to work on something else, now returns to her previous task. As she opens *core.h*, a popup alerts her that her version of the file is outdated, asking if she wants to update. She agrees and the file is now up-to-date.
5. Charlie, who's been working on the storage has committed a new version of *storage.cpp* and is already working on a new one.
6. After running some tests, Bob believes there may be a bug on the storage, and decides to consult *storage.cpp*. He is alerted to the fact that there is a user with uncommitted changes on that file, and agrees to request an update. He is also alerted to the fact that his version of the file is outdated, but decides not to update immediately and waits until Charlie commits his recent changes.
7. Charlie is alerted about one user requesting him to commit his changes on *storage.cpp*, but as he is currently fixing a bug, he decides not to commit until he finishes.

This is a little example of how awareness can help the coordination of software development. Its simplicity may not fully expose the importance that a small increase in awareness can have in team development, but this becomes clear when we consider the number of conflicts that occur in real situations.

IMPLEMENTATION

We have created two implementations of the filesystem layer of our tool using different approaches. The first is based on FUSE [18] and the second one on INOTIFY [19]. Each has its own advantages and disadvantages, that are discussed throughout this section.

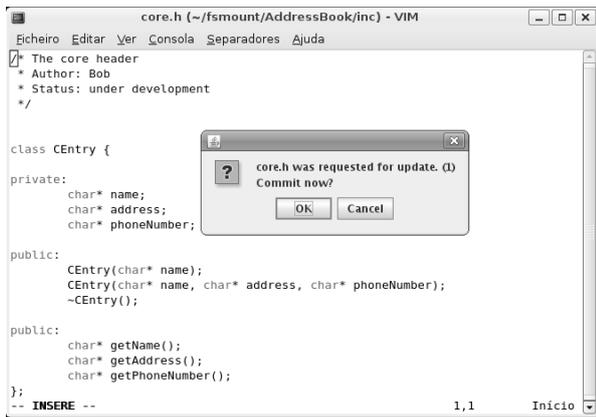


Figure 2. Bob being asked to update his uncommitted version of the file (as a result of Alice’s request). In this case, Bob is using the VIM editor.

Fuse

File System in Userspace (FUSE) is a Unix kernel module that allows non-privileged users to create virtual filesystems that run in user space [18]. The FUSE module intercepts system calls to the filesystem and redirects them to code that runs at user level. There are many projects using FUSE to create a virtual filesystem with different purposes.

This solution has the advantage of intercepting the filesystem calls, allowing to execute code before returning results to the applications. This could be used, for example, to update the local copy of a file before allowing the application to read its contents. These kind of functionality is not used in our tool, as we have decided to just provide awareness information.

Inotify

INOTIFY [19] is a Linux kernel subsystem that provides filesystem event notification (a similar mechanism exists in the Mac OS X system). With INOTIFY it is possible to monitor directories and files for events such as *open*, *close*, *create* or *delete*. An application may register itself to be notified for events occurring inside a set of directories, which is much more efficient than actively searching for changes or interposing code in the execution if the file system calls (as in FUSE).

Benchmarking

In this section we do some performance analysis on our tool, measuring the overhead of using our tool. For this, we used a project with 100 text files with a total of 1.7 MB (average file size: 17 KB).

As described earlier, each time the user opens a file, our tool will check if it is under version control, and in this the case it will check its status. This check is asynchronous, which means the user can keep opening files while another thread performs the check. We measured the time spent to open and read all the files (using the command: `cat * >/dev/null`) in the following scenar-

Table 1. Benchmarking results (in milliseconds per file) for the 1st run, the mean of 10 runs and standard deviation.

	x_1	\bar{x}	σ
Native	0.206	0.185	0.014
Inotify	0.208	0.220	0.054
Inotify+CVS	0.209	0.299	0.285
Inotify+SVN	0.399	0.383	0.557
Fuse	11.913	11.569	0.203
Fuse+CVS	24.906	13.898	3.688
Fuse+SVN	33.179	14.174	6.502

ios: on the native filesystem; FUSE *vs* INOTIFY; normal files *vs* files under CVS control *vs* files under SVN control. In all cases the repository is on a remote machine. On Table 1, we present the results for the first run, followed by the mean value of 10 runs and the standard deviation.

From the results in Table 1 we can observe that the overhead of using INOTIFY is almost null. Since notifications do not delay the filesystem call flow, the overhead is only due to the increased load on the machine and the running of additional tasks. The system asynchronously verifies if some awareness information must be provided. The delay to provide such information depends on the latency to the repository server, as the client’s daemon must check if there is any recent control information on the server.

It was not unexpected to confirm that there is a some time overhead associated to filesystem interception in FUSE. In particular, on the first access to each file, when the metafiles do not exist and have to be created. It could be argued that the loss of performance caused by the virtual filesystem does not affect the normal development of a project, since the overhead in opening a single file is not perceived by the user in an interactive session (the incurred delay is below 15ms). However, there are still situations where it can have a greater impact. The initial checkout of a large project will be considerably slower on the virtual filesystem.

In the current VC^2 implementation there is no situation in which we opted to delay or deny user actions on files. There are cases in which one could want to do so, for instance if we wanted users to have mandatory interactions with the VC^2 GUI in order to proceed with the opening of files. In such cases, filesystem interception would be the only solution.

Under the present VC^2 interaction model the best solution is clearly to resort to filesystem notifications and avoid the overhead incurred by file system interception.

DISCUSSION AND FUTURE WORK

The simple and generic architecture of our tool makes it easy to extend it with new features. As we mentioned before, there are still many features worth of develop-

ment. In this section we discuss some of these features.

Knowing the number of users that are changing/requesting a file is useful, but this can be further improved by showing a list with these users' names. Along with this list we could also associate messages with these actions by having a small text input field in the popup dialogs. These would be stored on the metafiles and displayed in the alert messages. This can provide a lightweight communication channel with a history that may be integrated in the version control system's logs (as in [7]).

As discussed before, it may not be convenient for developers to immediately commit their source code. This is not the case of documentation and other kind of files. For these we could define user configured properties such as *auto-commit* or *auto-update* that would automatically commit a changed file, or automatically update outdated files, and consequently eliminate most of the need for user interaction.

One of the main goals of our tool is automatic integration with minimum user interaction required. The users do not need to explicitly consult the state of the project's files to receive awareness information. Instead, they work normally and eventually receive alerts for potentially dangerous situations. Still, our alert-based approach is a bit intrusive, creating a new window for each alert. We are studying the possibility of providing an alternative interface, where all alerts could be cumulatively displayed in a single, less intrusive, popup window.

CONCLUSIONS

Conflicts are bound to occur when users engage in collaborative development of code (or text) in standard revision control frameworks. Since conflicts are often troublesome to solve, users try to avoid them by frequently issuing commits or by negotiating in parallel channels, like instant messaging applications and email. These actions are preventive and time consuming, and they are necessary due to the lack of mutual awareness among developers.

In this paper we introduce a system that addresses this problem, by greatly improving the level of awareness without forcing the commitment of unstable versions. The solution depicted in VC^2 is general and independent from the actual version control system in use. It currently runs on CVS and SVN and it is straightforward to adapt it to other version control tools. More important, VC^2 does not require any intervention on the server side, thus the solution can be deployed with public servers, such as SourceForge. Our solution also works independently of the editor used by users, thus allowing users to continue using their preferred applications.

REFERENCES

1. Ronald M. Baecker, Dimitrios Nastos, Ilona R. Posner, and Kelly L. Mawby. The user-centred iterative design of

- collaborative writing software. pages 775–782, 1995.
2. Bazaar. Bazaar version control, 2007. <http://bazaar-vcs.org/>.
3. R. Bentley, W. Appelt, U. Busbach, E. Hinrichs, D. Kerr, K. Sikkel, J. Trevor, and G. Woetzel. Basic support for cooperative work on the world wide web. *International Journal of Human Computer Studies: Special issue on Novel Applications of the WWW*, 46(6):827–856, Spring 1997.
4. Per Cederqvist et al. Version management with CVS, 2007. <http://www.cvshome.org/docs/manual>.
5. S. Dekeyser and R. Watson. Extending Google Docs to Collaborate on Research Papers. Technical report, The University of Southern Queensland, Australia, 2006.
6. Paul Dourish and Victoria Bellotti. Awareness and coordination in shared workspaces. In *Proceedings of the 1992 ACM conference on Computer-supported cooperative work*, pages 107–114. ACM Press, 1992.
7. G. Fitzpatrick, P. Marshall, and A. Phillips. CVS integration with notification and chat: lightweight software team collaboration. *Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*, pages 49–58, 2006.
8. R.E. Grinter. Using a configuration management tool to coordinate software development. *Proceedings of conference on Organizational computing systems*, pages 168–177, 1995.
9. Carl Gutwin, Mark Roseman, and Saul Greenberg. A usability study of awareness widgets in a shared workspace groupware system. In *CSCW '96: Proceedings of the 1996 ACM conference on Computer supported cooperative work*, pages 258–267, New York, NY, USA, 1996. ACM Press.
10. Susanne Hupfer, Li-Te Cheng, Steven Ross, and John Patterson. Introducing collaboration into an application development environment. In *CSCW '04: Proceedings of the 2004 ACM conference on Computer supported cooperative work*, pages 21–24, New York, NY, USA, 2004. ACM Press.
11. Claudia-Lavinia Ignat, Gérald Oster, Pascal Molli, and Hala Skaf-Molli. Gasper: A collaborative writing mode for avoiding blind modifications. Research Report RR-6204, LORIA – INRIA Lorraine, may 2007.
12. Pascal Molli, Hala Skaf-Molli, and Christophe Bouthier. State treemap: an awareness widget for multi-synchronous groupware. In *7th International Workshop on Groupware - CRIWG'2001*, Darmstadt, Germany, September 2001.
13. R. Robbes and M. Lanza. Versioning systems for evolution research. *Proceedings of IWPSE*, pages 155–164, 2005.
14. Anita Sarma, Zahra Noroozi, and Andre van der Hoek. Palantír: raising awareness among configuration management workspaces. In *ICSE '03: Proceedings of the 25th International Conference on Software Engineering*, pages 444–454, Washington, DC, USA, 2003. IEEE Computer Society.
15. M.A.D. Storey, D. Čubranić, and D.M. German. On the use of visualization to support awareness of human activities in software development: a survey and a framework. *Proceedings of the 2005 ACM symposium on Software visualization*, pages 193–202, 2005.
16. Subversion. Next-generation open source version control, 2007. <http://subversion.tigris.org/>.
17. Kimberly Tee, Saul Greenberg, and Carl Gutwin. Providing artifact awareness to a distributed group through screen sharing. In *CSCW '06: Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*, pages 99–108, New York, NY, USA, 2006. ACM Press.
18. Wikipedia. Filesystem in userspace — wikipedia, the free encyclopedia, 2007. [Online; accessed 30-May-2007].
19. Wikipedia. Inotify — wikipedia, the free encyclopedia, 2007. [Online; accessed 6-June-2007].
20. Erez Zadok and Jason Nieh. Fist: a language for stackable file systems. In *Proceedings of the Annual Technical Conference on 2000 USENIX Annual Technical Conference*, Berkeley, CA, USA, 2000. USENIX Association.