

A Verification Approach for Applied System Security

Achim D. Brucker¹, Burkhard Wolff²

¹ Information Security, ETH Zürich, ETH-Zentrum, CH-8092 Zürich, Switzerland
e-mail: brucker@inf.ethz.ch

² Universität Freiburg, George-Köhler-Allee 52, D-79110 Freiburg, Germany
e-mail: wolff@informatik.uni-freiburg.de

Published online: 25 January 2005

Abstract. We present a method for the security analysis of realistic models over off-the-shelf systems and their configuration by formal, machine-checked proofs. The presentation follows a large case study based on a formal security analysis of a CVS-Server architecture.

The analysis is based on an abstract architecture (enforcing a role-based access control), which is refined to an implementation architecture (based on the usual discretionary access control provided by the POSIX environment). Both architectures serve as a skeleton to formulate access control and confidentiality properties.

Both the abstract and the implementation architecture are specified in the language Z. Based on a logical embedding of Z into Isabelle/HOL, we provide formal, machine-checked proofs for consistency properties of the specification, for the correctness of the refinement, and for security properties.

Keywords: verification, security, access control, refinement, POSIX, CVS, Z

security architecture stems from the fact that the administration of authentication and access control is done via CVS itself; i.e. the authentication table is accessed and modified via standard CVS operations.

This work emerged from our own experiences with setting up a CVS-Server for more than 80 users worldwide. Besides overcoming a number of security problems (see, e.g. <http://www.cvshome.org/dev/security9706.html>), we had to develop an improved CVS-Server configuration described in [1] meeting two system design requirements: first, we had to provide a configuration of a CVS-Server that enforces a role-based access control [13]; second, we had to develop an “open CVS-Server architecture”, where the repository is part of the shared filesystem of a local network and the server is a regular process on a machine in this network. While such an architecture has a number of advantages, the correctness and trustworthiness of the security mechanisms become a major concern. Thus, we decided to apply formal modeling and analysis techniques to meet the challenge.

In this paper, we present the method we developed for analyzing the security problems of complex systems such as the CVS-Server and its configuration. As a result, we provide the following contributions:

1. a modeling technique that we call *architectural modeling*, which has an abstraction level in-between the usual behavioral modeling used in protocol analysis and code verification,
2. a technique to use system architecture models for defining security requirements,
3. the presentation of the mapping from security requirements to concrete security technologies as a data refinement problem,

1 Introduction

These days, the *Concurrent Versions System* (CVS) is a widely used tool for version management in many industrial software development projects, and plays a key role in open source projects usually carried out by highly distributed teams [3,4]. (See <http://www.cvshome.org>.) CVS provides a central database (the *repository*) and means to synchronize local modifications of partial copies (the *working copies*) with the repository. The repository can be accessed via a network; this requires a security architecture establishing authentication, access control and non-repudiation. A further complication of the CVS

4. mechanized proof-techniques for refinements and security properties over system transitions, and
5. reusable models for widely used security technologies.

In particular, we provide means to model a certain type of security policies, and show how security analysis can be performed not only on the abstract, but also on the concrete level.

The paper is organized as follows: After introducing some background material, e.g. CVS, our chosen specification formalism Z and the architectural modeling style, we present the model of the abstract system architecture. We proceed with the model of the POSIX filesystem as an infrastructure for the implementation architecture, and present the implementation architecture itself. Then, we describe the refinement relation between the system architecture and the implementation architecture, and the analysis of security properties at the different layers based on formal proofs in an interactive theorem prover.

2 Background

2.1 The CVS Operations

For the purpose of this paper, it is sufficient to mention only the most common CVS commands (initiated by the client). These are: `login` for client authenticating, `add` for registering files or directories for version control, `commit` for transferring local changes to the repository, and `update` for incorporating changes from the repository (e.g. fetching the latest version from the repository) into the working copy. Additionally, CVS provides functionality for accessing the history, for branching, for logging information which is out of the scope of this paper, and a mechanism for conflict resolution (e.g. merging the different versions) which is only modeled as an abstract operation. Further, in order to facilitate both the refinement and the security analysis, we will include in our CVS model a operation which is strictly speaking not part of CVS but of the operating system: the operation `modify`. This operation models changes of the working copy, e.g. by editing a file.

2.2 Z and Isabelle/HOL- Z

As our specification formalism, we chose Z [16] for the following reasons: first, Z fits our modeling problem since the complex states of our components suggest to use a formalism with rich theories for data-structures. Second, syntax and semantics of Z are specified in an ISO-standard;¹ for future standardization efforts of operating system libraries (e.g. similar to the POSIX [17] model in

Sec. 3.3.2), Z is therefore a likely candidate. Third, Z comes with a data-refinement notion [16, p. 136], which provides a correctness notion of the underlying “security technology mapping” between the two architectures and a means to compute the proof obligations. We assume a rough familiarity with Z (the interested reader is referred to excellent textbooks on Z such as [16, 18]).

As our modeling and theorem-proving environment, we chose Isabelle/HOL- Z [2], which is an integrated documentation, type-checking, and theorem-proving environment for Z specifications, which is built on top of Isabelle/HOL. Isabelle [9] is a *generic* theorem prover, i.e. new object logics can be introduced by specifying their syntax and inference rules. Isabelle/HOL is an instance of Isabelle with Church’s *higher-order logic* (HOL) [7], a classical logic with equality. Isabelle/HOL- Z is a conservative embedding of Z into HOL (which is semantically isomorphic to Z). As a result, Isabelle/HOL- Z combines up-to-date theorem proving technology with a widespread, standardized specification formalism, and powerful documentation facilities.

2.3 Architectural Modeling

As a means to identify conceptual entities of the problem domain and to structure the overall specification, we found it useful to describe the *architecture* of the system on several abstraction layers. Following Garlan and Shaw’s approach [6, 15], architectures are composed by *components* (such as *clients*, *servers* or *stores* like the filesystem) and *connectors* (like *channels*, *shared variables*, etc). In this terminology it is straight forward to make the mentioned architectures more precise (as implementation architecture, we present the intended “open server architecture”, see Fig. 1). We assume for each operation (such as `add`) a shared variable as connector that keeps all necessary information that goes to and from the components. This paves the way to formalize this architecture by describing the transition relation of the combined system by the parallel composition of the local transition relations of the components synchronized over the corresponding shared variable. Since such transition relations can be represented in Z by *operation schemas*, we can thus define, for example:

$$\begin{aligned}
 \text{CVS_add} &= \text{Client_add} \\
 &\quad \wedge \text{Server_add} \setminus \text{add}_{\text{shared variable}}
 \end{aligned}$$

where \wedge is the schema-conjunction and \setminus the hiding operator (i.e. an existential quantifier). Throughout this paper, we will only present combined operation schemas and model properties over the transitive closure of their transition relations.

2.4 Architecture Refinement

When analyzing security architectures one can separate an *abstract security architecture* (see Sec. 3.2), which

¹ Z formal specification notation $\hat{\Lambda}$ syntax, type system and semantics, 2002. ISO/IEC 13568:2002

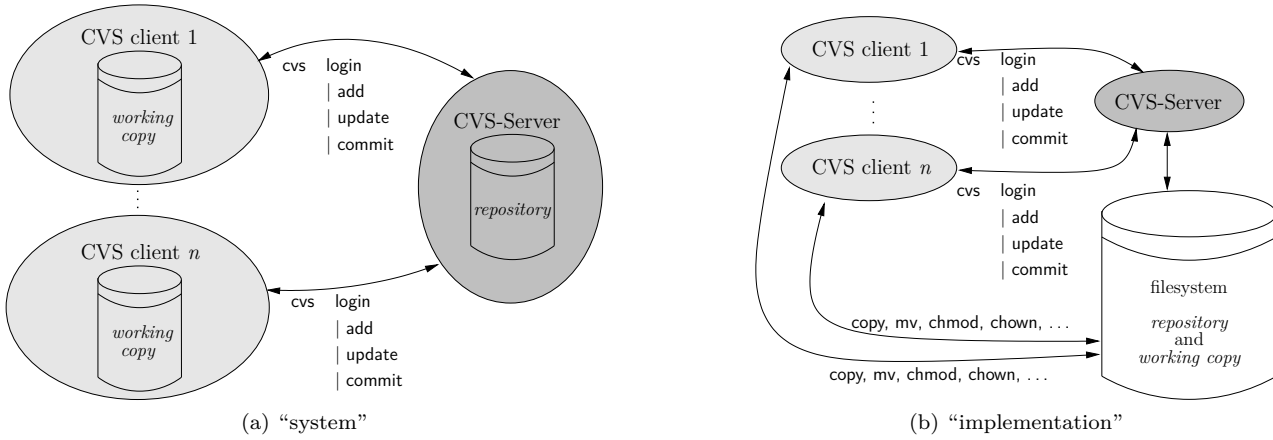


Fig. 1: The different CVS-Server architectures

is merely a framework for describing the security requirements, from an *implementation architecture* (see Sec. 3.3) where a mapping to security mechanisms is described (see Fig. 2). By connecting the abstract and the concrete layer formally, it is possible to reason about safety and security properties on the abstract level. Such a connection between abstract and more concrete views on a system and their semantic underpinning is well-known under the term *refinement*, and security technology mappings can be understood as a special case of this. Various refinement notions have been proposed [18,11]; in our setting, we chose to use only a simple data refinement notion following Spivey [16].

2.5 Security Models vs. Security Technologies

Many security models distinguish between *objects* (e.g. data) and *subjects* (e.g. users). Using *role-based access control* (RBAC) [13] one assigns each subject at least one role (e.g. the role “administrator”), and access of objects is granted or denied by the role a subject is acting in. Further, roles can be hierarchically ordered, e.g. subjects in the role “administrator” are allowed to do everything other roles are allowed to. Our CVS-Server uses such a *hierarchical RBAC* model.

In an RBAC model, the decision, which roles may have access to which objects is done during system design and cannot be changed by regular users. In con-

trast, in a *discretionary access control* (DAC) model, every object belongs to a specific subject (its *owner*), and the owner is allowed to change the access policies at any time, hence “discretionary”. For example a DAC implementation that also allows grouping users is the Unix/POSIX filesystem layer [17] access control.

Based on a DAC that supports groups, one can ‘implement’ an RBAC model by a special setup [12]. We use a similar technique to implement a hierarchic RBAC model for our CVS-Server on top of the POSIX filesystem layer, which is described in Sec. 3.3.3. However, we will analyze the concrete form in which DAC is implemented in POSIX and not a conceptual model thereof.

3 The CVS-Server Case Study

The Z specification of the CVS-Server [1] consists of more than 120 pages, and the associated proof scripts are about 13000 lines of code. The organization of the Z-sections follows directly the overall scheme presented in Fig. 3. The Z-sections *AbsState* and *AbsOperations* are describing the abstract system architecture of the client and the server components. The Z-section *SysConsistency* contains the consistency conditions (conservatism of axiomatic definitions, definedness of applications, non-blocking operation schemas) of the system architecture. This is mirrored at the implementation architecture level by the structures *FileSystem*, *CVS-Server*, and *ImplConsistency*. The Z-section *Refinement* contains the usual abstraction predicates relating the abstract and the concrete states, and also the proof obligations for this refinement. The security properties, together with the corresponding proof obligations, are defined in the sections *SysArchSec* and *ImplArchSec*.

3.1 Entities of the Security Model

Following the standard RBAC model, we introduce abstract types for CVS clients (users) *Cvs_Uid*, permissions




Fig. 2: Refining Security Architectures




Fig. 3: Organizing the Specification into Z-sections

Cvs_Perm (which are isomorphic to roles in our setting), and CVS passwords Cvs_Passwd used to authenticate a CVS client for a permission:

$$[Cvs_Uid, Cvs_Perm, Cvs_Passwd]$$

Permissions are hierarchically organized by the reflexive and transitive relation cvs_perm_order (over permissions Cvs_Perm) with cvs_adm as greatest element:

$$\begin{array}{l} \hline cvs_adm, cvs_public : Cvs_Perm \\ cvs_perm_order : Cvs_Perm \leftrightarrow Cvs_Perm \\ \hline cvs_perm_order = cvs_perm_order^* \\ \forall x : Cvs_Perm \bullet (x, cvs_adm) \in cvs_perm_order \\ \forall x : Cvs_Perm \bullet (cvs_public, x) \in cvs_perm_order \\ \forall x : Cvs_Perm \bullet (x \neq cvs_adm) \Rightarrow \\ \quad (cvs_adm, x) \notin cvs_perm_order \\ \forall x : Cvs_Perm \bullet (x \neq cvs_public) \Rightarrow \\ \quad (x, cvs_public) \notin cvs_perm_order \\ \forall x : Cvs_Perm \bullet \exists y : Cvs_Perm \bullet \\ \quad (x, y) \in cvs_perm_order \end{array}$$

We turn now to the security entities and mechanisms of the CVS-Server and the clients: first we have to model the working copies and the repositories as maps assigning *abstract names* Abs_Name to *data* Abs_Data (both types are abstract in our model):

$$\begin{array}{l} [Abs_Name, Abs_Data] \\ ABS_DATATAB \equiv Abs_Name \mapsto Abs_Data \end{array}$$

A CVS-Server provides an authorization table, which is used to control access within the repository. The server stores for each file in the repository the required permission. These tables are modeled as follows:

$$\begin{array}{l} AUTH_TAB \quad \equiv Cvs_Uid \times Cvs_Passwd \\ \quad \mapsto Cvs_Perm \\ ABS_PERMTAB \equiv Abs_Name \mapsto Cvs_Perm \end{array}$$

Clients possess in their working also a table that assigns to each abstract name a CVS client and another

map that associates each CVS client to the password previously used during the cvs login procedure. The interplay of these tables will be discussed later, here we just define them:

$$\begin{array}{l} ABS_UIDTAB \equiv Abs_Name \mapsto Cvs_Uid \\ PASSWD_TAB \equiv Cvs_Uid \mapsto Cvs_Passwd \end{array}$$

3.2 The System Architecture

In this section, we give a brief overview on how we model the system architecture, which is divided into: the state of the server (including the repository), the state of the client (including the working copy), and a set of CVS operations working over both of them.

It is a distinguishing feature of a CVS-Server to store the authentication data inside the repository such that it can be accessed and modified with CVS operations. This implies certain formal prerequisites: we require an abstract name $abs_cvsauth$ to be associated with data, which can be converted into an authentication table via a postulated function $authtab$.

$$\begin{array}{l} abs_cvsauth : Abs_Name \\ abs_auth_of : Abs_Data \mapsto AUTH_TAB \\ abs_data_of : AUTH_TAB \mapsto Abs_Data \\ authtab : ABS_DATATAB \mapsto AUTH_TAB \\ \hline ran(abs_data_of) \subseteq dom(abs_auth_of) \\ \forall x : dom abs_auth_of \bullet \\ \quad abs_data_of(abs_auth_of x) = x \\ \forall x : AUTH_TAB \bullet \\ \quad abs_auth_of(abs_data_of x) = x \\ \forall r : ABS_DATATAB \bullet abs_cvsauth \in dom(r) \Rightarrow \\ \quad authtab(r) = abs_auth_of(r abs_cvsauth) \end{array}$$

Modeling the server's state as a Z schema is straight forward. The state contains the repository rep and the map $rep_permtab$ containing the required permissions for each file. Accessing the authentication table inside rep will require to have the role cvs_adm . *RepositoryState* is modeled as follows:

RepositoryState

```

rep : ABS_DATATAB
rep_permtab : ABS_PERMTAB

abs_cvsauth ∈ dom rep
dom rep = dom rep_permtab
rep_permtab(abs_cvsauth) = cvs_adm
rep(abs_cvsauth) ∈ dom abs_auth_of

```

The state of the client component contains the working copy wc , the wc_uidtab assigning a CVS client to each file and a password table abs_passwd with credentials (passwords) used in previous CVS login operations (abs_passwd models the file `.cvspass`). Thus, for any data in the working copy and whenever an access to it may be processed, an individual role may be generated and validated by the server with respect to its current repository state. Further, there is a set of abstract names $wfiles$ which is used as filter in `update` and `commit` operations. This filter corresponds to the concept of the *working directory* in the implementation, i.e. the effects of these operations are restricted to files stored within the working directory:

ClientState

```

wc : ABS_DATATAB
wc_uidtab : ABS_UIDTAB
abs_passwd : PASSWD_TAB
wfiles :  $\mathbb{P}$  Abs_Name

```

In what follows, we define the abstract CVS operations that model combined state transitions of the client and the repository. Due to space reasons, we only present `login` and `commit`.

The `login` operation simply stores the authentication data on the client-side. This is used to authenticate a CVS user for a permissions of the client. The Δ and Ξ notation is used in Z to import the schemas in two variants; one variant as a copy, the other by replacing all variables by corresponding stroked variables (e.g. wc') describing the successor state. The Ξ also introduces equalities enforcing that the components of the previous state are equal to the post state components.

abs_login

```

 $\Delta$  ClientState
 $\Xi$  RepositoryState
passwd? : Cvs_Password
uid? : Cvs_Uid

(uid?, passwd?) ∈ dom(authtab rep)
abs_passwd' = abs_passwd  $\oplus$  {uid?  $\mapsto$  passwd?}
wc' = wc
wc_uidtab' = wc_uidtab
wfiles' = wfiles

```

The `commit` (ci) operation usually takes a set of files as arguments (here denoted by $files?$). The case that no arguments may be passed is modeled by the possibility to set $files?$ to the set of all files ABS_NAME .

Now we address the core of our hierarchic RBAC model of the system architecture, the *has_access* predi-

cate. As a prerequisite, we define the shortcut `is_valid_in` for checking that a CVS client, together with a credential (password), represents a valid role with respect to the current repository:

```

_ is_valid_in _ : (Cvs_Uid  $\times$  Cvs_Password)
 $\leftrightarrow$  ABS_DATATAB

 $\forall$  role : Cvs_Uid; pwd : Cvs_Password;
rep : ABS_DATATAB  $\bullet$ 
(role, pwd) is_valid_in rep
 $\Leftrightarrow$  (role, pwd) ∈ dom(authtab(rep))

```

Further, the *has_access* predicate ensures `is_valid_in` and that the permissions resulting from these credentials are sufficient to access the requested file according to the role hierarchy:

```

has_access _ :  $\mathbb{P}$ (ABS_PERMTAB
 $\times$  ABS_DATATAB  $\times$  PASSWD_TAB
 $\times$  Abs_Name  $\times$  Cvs_Uid)

 $\forall$  rep_pt : ABS_PERMTAB; rep : ABS_DATATAB;
pwtb : PASSWD_TAB; file : Abs_Name;
role : Cvs_Uid  $\bullet$ 
has_access(rep_pt, rep, pwtb, file, role)
 $\Leftrightarrow$  (role, pwtb(role)) is_valid_in rep
 $\wedge$  (rep_pt(file), authtab(rep)(role, pwtb(role)))
 $\in$  cvs_perm_order

```

The `commit` operation consists of the construction of a new repository rep' and a new table with required permissions $rep_permtab'$ which were constructed via the override operator \oplus from previous states of these tables. For rep' , three cases can be distinguished: (i) either a file in the repository does not occur in the working copy, then it is unchanged, or (ii) it occurs in the working copy but not in the repository, then it is copied provided a valid permission is available in the wc_uid_tab of the working copy, or (iii) the file exists both in working copy and repository, then the working copy file overrides the repository file whenever the client has access:

abs_ci

```

 $\Xi$  ClientState
 $\Delta$  RepositoryState
files? :  $\mathbb{P}$  Abs_Name

(wfiles  $\cap$  files?)  $\subseteq$  dom wc
rep' = rep  $\oplus$  ({n : wfiles  $\cap$  files? | n  $\notin$  dom rep
 $\wedge$  n ∈ dom wc_uidtab
(wc_uidtab(n), abs_passwd(wc_uidtab n))
is_valid_in rep}  $\triangleleft$  wc)
 $\oplus$  ({n : wfiles  $\cap$  files? | n ∈ dom rep
 $\wedge$  n ∈ dom wc_uidtab
 $\wedge$  has_access(rep_permtab, rep,
abs_passwd, n, wc_uidtab(n))
}  $\triangleleft$  wc)
rep_permtab' = rep_permtab  $\oplus$  {n : wfiles  $\cap$  files? |
n  $\notin$  dom rep  $\wedge$  n ∈ dom wc_uidtab
 $\wedge$  (wc_uidtab(n), abs_passwd(wc_uidtab n))
∈ dom(authtab rep)  $\bullet$ 
n  $\mapsto$  authtab(rep)(wc_uidtab(n),
abs_passwd(wc_uidtab n))}

```

The table $rep_permtab'$ is extended by permissions for files that are new in the repository (based on the permissions used for committing these files). Further, the table wc_uid_tab is updated by the add operation, which we omit here.

In addition to these abstract models of the CVS operations, we provide a modify operation which explicitly models interactions of users with their files via modifying the files of the working copy of the client state.

3.3 The Implementation Architecture

The implementation architecture of CVS-Server is intended to model realistically the security mechanisms used to achieve the security goals formalized in the previous system architecture. Therefore, it captures the relevant operating system environment methods, i.e. POSIX methods in our case, for accessing files and changing their access attributes. We derived our POSIX model by formalizing the specification documents [17] and detailed system descriptions [5] and by validating it by carefully chosen tests and by inspections of critical parts of the system sources. In this POSIX model, the *CVS Filesystem* will be embedded, i.e. a repository is described as some area in the filesystem, where file attributes are set in a suitable way.

3.3.1 Modeling Basic Data Structures

We declare basic abstract sorts for POSIX user IDs, group IDs, data (file contents left abstract in this model), elementary filenames and file paths.

$$\begin{array}{l} [Uid, Gid, Data, Name] \\ Path \equiv \text{seq Name} \end{array}$$

We assume a static table $groups$ that assigns to each user a set of groups he belongs to. We also describe a special user ID $root$, modeling the system administrator. As we will show later, all security goals can only be achieved for all users except $root$, because $root$ is allowed to do (almost) everything.

$$\left\{ \begin{array}{l} groups : Uid \rightarrow \mathbb{P} Gid \\ root : Uid \end{array} \right.$$

3.3.2 Modeling the POSIX Filesystem Access Control

Within POSIX, every file belongs to a unique pair of owner (user) and group, and file access is divided into access by the *user* (owner), the *group* or *other* (world). The POSIX *discretionary access control* (DAC) distinguishes access for reading (r), writing (w), and executing (x). We also model the “set group id” (sg) on directories, which affects the default group of newly created files within that directory (see [5] for more technical details about the Unix/POSIX DAC):

$$Perm ::= ru \mid wu \mid xu \mid rg \mid wg \mid xg \mid ro \mid wo \mid xo \mid sg$$

The filesystem consists of a map from a file path to file content (which is either *Data* for regular files or *Unit* for directories²) and of file attributes (assigning to each file or directory the permissions³, the user ID of the owner and the group it belongs to). Our concept of file attributes may easily be extended by adding new components to its records.

$$\begin{array}{l} Unit ::= Nil \\ FILESYS_TAB \equiv Path \leftrightarrow (Data + Unit) \\ FILEATTR \equiv [perm : \mathbb{P} Perm; uid : Uid; gid : Gid] \\ FILEATTR_TAB \equiv Path \leftrightarrow FILEATTR \end{array}$$

We use type sums for modeling the *FILESYS.TAB* which are not part of the Z standard. Type sums can simulate enumerations in Z free type definitions on the fly. The two functions $Inl : X \rightarrow (X + Y)$ and $Inr : Y \rightarrow (X + Y)$ are provided for building type sums.

For testing if a directory contains a specific entry (either a file or a directory) we provide the function is_in . Further, we provide functions that test for regular files (is_file_in) and for directories (is_dir_in); their definitions are straight forward:

$$\begin{array}{l} _is_in _ : Path \leftrightarrow (Path \leftrightarrow (Data + Unit)) \\ _is_dir_in _ : Path \leftrightarrow (Path \leftrightarrow (Data + Unit)) \\ _is_file_in _ : Path \leftrightarrow (Path \leftrightarrow (Data + Unit)) \\ \hline \forall fs : (Path \leftrightarrow (Data + Unit)); f : Path \bullet \\ \quad (f \text{ is_in } fs) \Leftrightarrow f \in \text{dom } fs \\ \forall fs : (Path \leftrightarrow (Data + Unit)); d : Path \bullet \\ \quad (d \text{ is_dir_in } fs) \Leftrightarrow (d \text{ is_in } fs) \\ \quad \wedge (\exists u : Unit \bullet fs(d) = Inr(u)) \\ \forall fs : (Path \leftrightarrow (Data + Unit)); f : Path \bullet \\ \quad (f \text{ is_file_in } fs) \Leftrightarrow (f \text{ is_in } fs) \wedge \neg (f \text{ is_dir_in } fs) \end{array}$$

At this point we are ready to model the filesystem state, which mainly describes the map of (name) paths to their attributes. As mentioned, we require that all defined paths must be “prefix-closed”, i.e. all prefix paths must be defined in the filesystem (thus constituting a tree) and point to directories.

$$\begin{array}{l} \text{--- } FileSystem \text{ ---} \\ files : FILESYS_TAB \\ attributes : FILEATTR_TAB \\ \hline \forall p : \text{dom } files \bullet (p = \langle \rangle) \\ \quad \vee (front(p) \text{ is_dir_in } files) \\ \text{dom } files = \text{dom } attributes \end{array}$$

In addition to the filesystem state, we introduce a state schema *ProcessState* for client related information,

² We do not consider *special files*, like devices, named pipes or process files.

³ The terms *attributes* and *permissions* are used interchangeably.

namely the current user and group ID, the client's umask (which is used to set the initial file attributes on new files) and current working directory (*wdir*). The working directory is often used as an implicit parameter to filesystem and CVS operations:

ProcessState
$uid : Uid$ $gid : Gid$ $umask : \mathbb{P}(Perm \setminus \{sg\})$ $wdir : Path$

As a prerequisite for describing functions that do modifications on the file system, we need to model the POSIX DAC in detail. Therefore we first introduce a function *has_attrib*, which decides whether the attributes (read, write and execute) of a file are set with respect to a specific user (and the groups he is a member of). Within this function, a crucial detail of the POSIX access model is formalized, namely that file access is checked by *sequentially* testing the following conditions (leading to an overall failure if the first condition fails):

1. If the user owns the file, he can only access the file if the access attributes for users grant access.
2. If the user is a member of the group owning the file, he can only access the file if the access attributes for the group grant access.
3. Last, the access attributes for others are checked.

These requirements may lead to some unexpected consequences, e.g. assume a user *u* being member of the group *g* and owner of a file with the permissions $\langle perm == \{rg, ro\}, uid == u, gid == g \rangle$. Curiously, file access will be denied for him, while granted for all others in his group, because the rights specified for the user precede the rights given for the group.

$has_attrib_ : \mathbb{P}(Uid \times Path \times FILEATTR_TAB$ $\quad \times Perm \times Perm \times Perm)$
$\forall uid : Uid; fa : FILEATTR_TAB;$ $pu, pg, po : Perm \bullet \forall p : \text{dom}(fa) \bullet$ $has_attrib(uid, p, fa, pu, pg, po) \Leftrightarrow$ $((uid = root) \vee$ $(\forall m : \mathbb{P} Perm; diruid : Uid; dirgid : Gid $ $\quad \langle perm \equiv m, uid \equiv diruid, gid \equiv dirgid \rangle$ $\quad = fa(p) \bullet$ $\quad (diruid = uid \wedge pu \in m) \vee$ $\quad (diruid \neq uid \wedge dirgid \in groups(uid)$ $\quad \wedge pg \in m) \vee$ $\quad (diruid \neq uid \wedge dirgid \notin groups(uid)$ $\quad \wedge po \in m)))$

Based on *has_attrib* we introduce shortcuts for checking read, write and execute attributes (e.g. *has_w_attrib*) of files and directories as well as definitions for checking the read, write and execute access (e.g. *has_w_access*).

$has_w_attrib_ : \mathbb{P}(Uid \times Path \times FILEATTR_TAB)$ $has_r_attrib_ : \mathbb{P}(Uid \times Path \times FILEATTR_TAB)$ $has_x_attrib_ : \mathbb{P}(Uid \times Path \times FILEATTR_TAB)$

$\forall uid : Uid; p : Path; fa : FILEATTR_TAB \bullet$ $has_w_attrib(uid, p, fa)$ $\Leftrightarrow has_attrib(uid, p, fa, wu, wg, wo)$ \dots

$has_w_access_ : \mathbb{P}(Uid \times Path \times FILEATTR_TAB)$ $has_r_access_ : \mathbb{P}(Uid \times Path \times FILEATTR_TAB)$ $has_x_access_ : \mathbb{P}(Uid \times Path \times FILEATTR_TAB)$

$\forall uid : Uid; p : Path; fa : FILEATTR_TAB \bullet$ $has_w_access(uid, p, fa) \Leftrightarrow$ $(\forall pref : Path pref \text{ prefix } (front\ p) \bullet$ $has_x_attrib(uid, pref, fa))$ $\wedge has_w_attrib(uid, front\ p, fa)$ \dots
--

As an example for our approach to specify POSIX operations, we present the (shortened) file remove specification [17], which corresponds to *unlink()*:

The unlink() function shall fail and shall not unlink the file if:

- A component of path does not name an existing file or path is an empty string.
- Search permission is denied for a component of the path prefix, or write permission is denied on the directory containing the directory entry to be removed.

This text is formalized by a Z operation schema *rm* as follows: The first condition in the body is common for most filesystem operations and requires the path of the file must be a valid one in the filesystem table. The second condition requires that the client has write permissions on the file and the working directory (“*the directory containing the directory entry to be removed*”), which is checked via the *has_w_access* predicate:

rm
$\Delta FileSystem$ $\Xi ProcessState$ $u? : Name$
$(wdir \hat{\cap} \langle u? \rangle) \text{ is_file_in_files}$ $has_w_access(uid, wdir, attributes)$ $\quad \wedge has_w_access(uid, wdir \hat{\cap} \langle u? \rangle, attributes)$ $files' = \{wdir \hat{\cap} \langle u? \rangle\} \triangleleft files$ $\quad \wedge attributes' = attributes$

The definitions for the remaining filesystem operations are similar, see [1] for details.

3.3.3 Mapping CVS Access Control onto POSIX DAC

We turn now to a crucial aspect of the implementation of the security goals by security mechanisms provided from

standard POSIX DAC: any CVS role will be mapped to a particular pair of system *owner* and a set of system *groups*. This mapping has the consequence of an inheritance mechanism for generating default roles when creating new objects in the repository. Additionally, there is a mechanisms to “down-scale” and “up-scale” the permissions in the repository for the CVS administrator (not described here).

For every CVS operation, the server determines the CVS role according to the client’s CVS ID and password. These roles are then mapped to POSIX user and group IDs, and these are compared to the file attributes of the files and directories the operations operates on. This translation is done by the two functions *cvsperm2uid* and *cvsperm2gid*.

$$\begin{array}{l} \text{cvsperm2uid} : \text{Cvs_Perm} \mapsto \text{Uid} \\ \text{cvsperm2gid} : \text{Cvs_Perm} \mapsto \text{Gid} \\ \text{users} : \mathbb{P} \text{Uid} \\ \hline \text{root} \notin \text{ran } \text{cvsperm2uid} \\ \text{ran } \text{cvsperm2uid} \cap \text{users} = \emptyset \\ \text{ran } \text{cvsperm2gid} \cap \bigcup \{x : \text{users} \bullet \text{groups}(x)\} = \emptyset \\ \text{ran } \text{cvsperm2uid} \triangleleft \text{groups} = \{x : \text{Cvs_Perm} \bullet \\ \text{cvsperm2uid } x \mapsto \\ \{c : \text{Cvs_Perm} \mid (c, x) \in \text{cvs_perm_order} \bullet \\ \text{cvsperm2gid } c\}\} \end{array}$$

It is important to notice that CVS IDs (*Cvs_Uid*) are independent of POSIX IDs (*Uid*) and that the POSIX IDs which are used by CVS are disjoint from “normal” POSIX user IDs, i.e. it is impossible to login with such a special POSIX ID.

From these distinctness constraints follows that the POSIX system administrator and the CVS administrator may be different. Moreover, we require that the group table (administrated by the system administrator and nobody else) is compatible with *cvs_perm_order*. These requirements have to be assured during installation of a CVS server.

The CVS repository is a subtree of the normal filesystem; its root is denoted by the absolute path *cvs_rep* and all paths inside the repository are relative to the root *cvs_rep*. Further, the administrative files of CVS are stored in the *CVSROOT* directory, which is a subdirectory of *cvs_rep*, and the file that contains all authentication information is called *cvsauth* and is located inside *CVSROOT*.

$$\begin{array}{l} \text{cvs_rep} : \text{Path} \\ \text{CVSROOT} : \text{Name} \\ \text{cvsauth} : \text{Name} \\ \text{auth_of} : \text{Data} \mapsto \text{AUTH_TAB} \\ \text{data_of} : \text{AUTH_TAB} \mapsto \text{Data} \\ \hline \text{ran } \text{data_of} \subseteq \text{dom } \text{auth_of} \\ \forall x : \text{dom } \text{auth_of} \bullet \text{data_of}(\text{auth_of } x) = x \\ \forall x : \text{AUTH_TAB} \bullet \text{auth_of}(\text{data_of } x) = x \end{array}$$

3.3.4 Modeling the CVS Filesystem

A major design decision for our specification is to enrich the *FileSystem* state by new state components relevant to CVS, or more precisely, the combined client/server component of CVS. In CVS, working copies contain specific attributes assigned to the files; we restrict ourselves to security relevant attributes, i.e. the CVS client ID and password, and the path *rep* where the file is located in the repository. This information is kept in an own table implicitly associated to the working copies.

$$\begin{array}{l} \text{CVS_ATTR} \quad \equiv [\text{rep} : \text{Path}; \text{f_uid} : \text{Cvs_Uid}] \\ \text{CVS_ATTR_TAB} \equiv \text{Path} \mapsto \text{CVS_ATTR} \end{array}$$

Due to the space reasons, we only show some requirements of the combined POSIX and CVS filesystem:

- working copies and the repository are distinct areas of the filesystem.
- the repository contains a special directory that contains the administrative data of CVS. Certain restrictive access permissions must be ensured to this directory and its contents to preserve the system integrity.
- requirements on file attributes within the repository:
 - since the owners of files must be POSIX user IDs that are disjoint from “regular” POSIX user IDs, and the group IDs must be legal with respect to the CVS role hierarchy. This guarantees that regular users only have the rights described by the file attributes for others. Thus, our initial invariant for the base directory of the repository implies that such a user cannot do anything, using only POSIX operations, within the repository.
 - read, write and execute permissions are the same for user and group. Together with our group setup this ensures that the initial CVS role and all roles with higher precedence have the same rights to access that file.

These invariants are formally described in the axiomatic definition:

$$\begin{array}{l} \text{attr_in_rep} _ : \mathbb{P} \text{FileSystem} \\ \text{attr_in_root} _ : \mathbb{P} \text{FileSystem} \\ \text{attr_outside_root} _ : \mathbb{P} \text{FileSystem} \\ \hline \forall \text{fs} : \text{FileSystem} \bullet \text{attr_in_rep}(\text{fs}) \Leftrightarrow \\ (\forall p : \text{dom } \text{fs.files} \mid (\text{cvs_rep } \text{prefix } p) \bullet \\ (((\text{fs.attributes } p).\text{uid}) \in \text{ran } \text{cvsperm2uid} \\ \wedge ((\text{fs.attributes } p).\text{gid}) \\ \in \text{groups}((\text{fs.attributes } p).\text{uid}) \wedge \\ (\text{ru} \in ((\text{fs.attributes } p).\text{perm}) \Leftrightarrow \text{rg} \\ \in (\text{fs.attributes } p).\text{perm}) \wedge \\ (\text{wu} \in ((\text{fs.attributes } p).\text{perm}) \Leftrightarrow \text{wg} \\ \in (\text{fs.attributes } p).\text{perm}) \wedge \\ (\text{xu} \in ((\text{fs.attributes } p).\text{perm}) \Leftrightarrow \text{xg} \\ \in (\text{fs.attributes } p).\text{perm}))) \\ \dots \end{array}$$

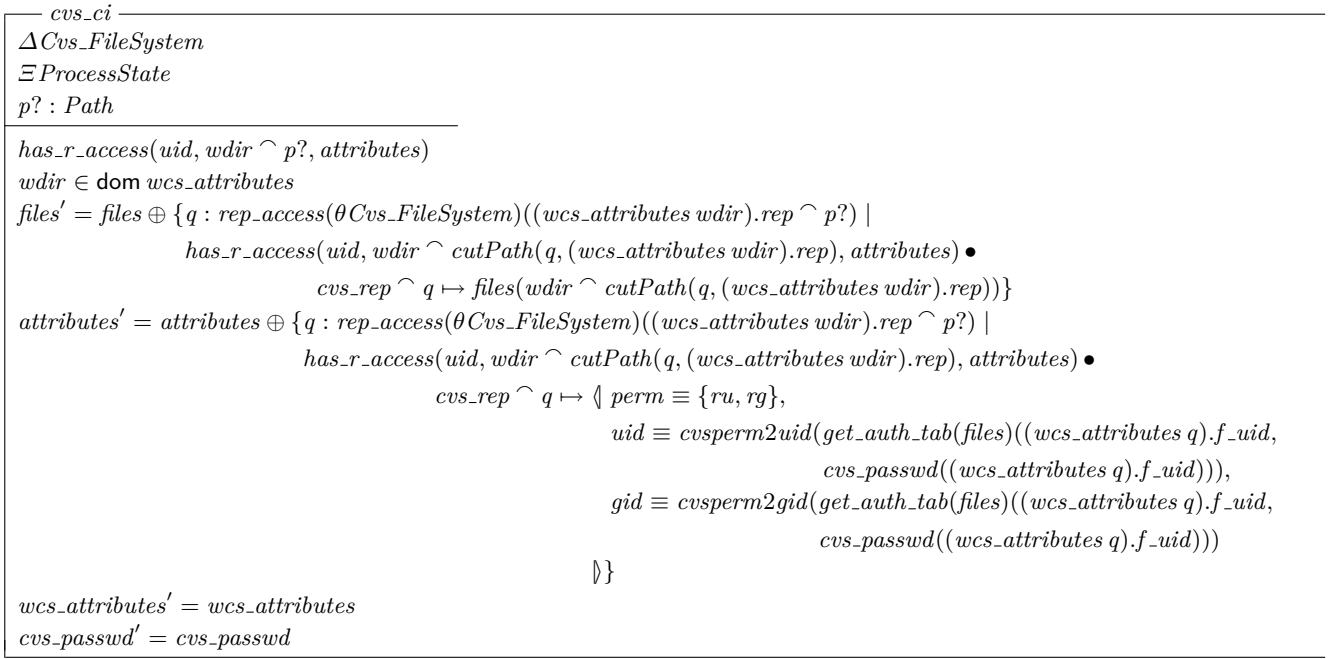
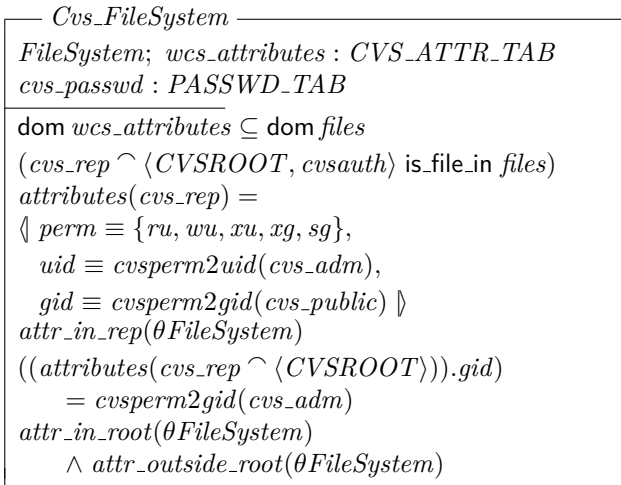


Fig. 4: The specification of the commit command (implementation architecture)

We turn now to a formal description of the repository *within* the filesystem. This invariant of the system is captured in the state schema *Cvs_FileSystem*:




Additionally to *rep_attributes*, we impose similar requirements for the administrative area of the repository by the predicate *attr_in_root*. Further, we describe in the predicate *attr_outside_root* the requirements for the data in the repository, i.e. files that are subject to version control. Both axiomatic definitions are omitted here.

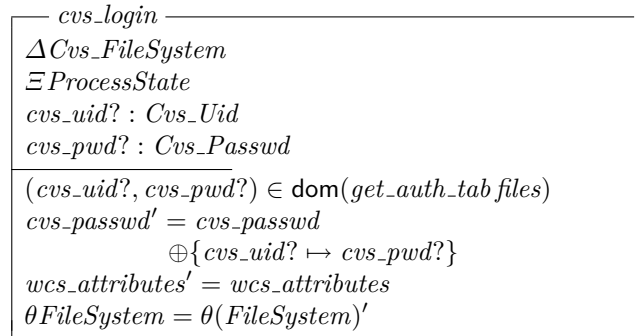
Now we have established a basis for the operations on the combined POSIX and CVS environment. As in Sec. 3.2, we present the *login* and *commit* operations in order to compare the two different architecture levels.

Before we describe the operations of the CVS-Server we need to model the access to the CVS authentication table (*get_auth_tab*) that is part of the *cvs_rep* \wedge *CVSROOT* directory and underlies the standard access discipline of CVS-Server. In particular, the authentica-


tion table is only modifiable by the CVS administrator, but not by any other client of the system.



The *login* operation updates the variable *cvs_passwd*, provided that for the combination of user ID and password the authentication will succeed.



In the *commit* operation, the current working directory *wdir* can be restricted by the parameter *p?* to just one file or directory. All files below *p?* for which the client has access will be committed. We use the function *cutPath* to remove a given prefix from a path.



In contrast to the system architecture specification we also must determine the POSIX file attributes of the files. The particularity of the update and the *commit* operation is the use of *rep_access* which computes the

paths into the repository to which the client has read access according to his CVS role.

$$\begin{array}{l}
 \overline{rep_access : Cvs_FileSystem \rightarrow Path \rightarrow \mathbb{P} Path} \\
 \forall cfs : Cvs_FileSystem; p : Path \bullet \\
 \quad rep_access(cfs)(p) = \{q : Path \mid p \text{ prefix } q \\
 \quad \wedge cvs_rep \wedge q \in \text{dom } cfs.files \\
 \quad \wedge (\exists idpwd : cfs.cvs_passwd \bullet \\
 \quad \quad idpwd \in \text{dom}(get_auth_tab(cfs.files)) \\
 \quad \wedge (has_r_access(cvsperm2uid(\\
 \quad \quad get_auth_tab(cfs.files)(idpwd)), \\
 \quad \quad cvs_rep \wedge q, cfs.attributes) \\
 \quad \vee (has_x_access(cvsperm2uid(\\
 \quad \quad get_auth_tab(cfs.files)(idpwd)), \\
 \quad \quad cvs_rep \wedge q, cfs.attributes) \\
 \quad \wedge cvs_rep \wedge q \text{ is_dir_in } cfs.files))\}
 \end{array}$$

The schema *cvs.ci* (see Fig. 4) models the **commit** command. We require that the client has read access for the file or directory in the current working directory and sufficiently high-ranked role to modify the repository.

4 Formal Analysis

A formal model, even if successfully type-checked, is in itself not a value of its own: it must be validated, e.g. by testing techniques or by formal proof activities as in our approach. In this section, we present a formal consistency check of the specifications, and we show that the implementation architecture is, in a formal sense, a refinement of the abstract system architecture. We specify and prove security properties of the type “no combination of user-commands will enable a user to write into the repository, except he has the required access rights”.

4.1 Checking the Consistency

Two types of “sanity checks” are useful and have been carried out with HOL-Z [2] routinely:

- definedness checks for all applications of partial functions in their context, as undefined applications usually indicate that some part of the precondition of a schema context is missing, and
- checking the state invariant of all operation schemas; in particular, we require that in a schema, all syntactic preconditions (i.e. the conjuncts in the predicate part that contain occurrences of variables without stroke “/” and “!” suffix) suffice to show that a successor state exists.

Violating these conditions does not result in logical inconsistencies but in unprovable statements or operation definitions with undesired semantical effects.

4.2 Establishing the Refinement

To prove that the concrete implementation architecture correctly implements the abstract system architecture,

we have to define an abstraction schema *R* which relates the components of the abstract state to the components of the concrete state. In particular, we must map abstract names and data to paths and files in the sense of the POSIX filesystem, and the working copies and repositories of the abstract model must be related to certain areas of the filesystem, the authentication tables must be related, the user must not be *root* (the refinement simply does not work otherwise) and the file attributes in the concrete filesystem must be convertible along the mapping discussed in Sec.3.3.3.

Due to limited space, we will only show two constraints of *R* formally. As a prerequisite, let us define a function *Rname2path*, which maps abstract names, to file paths in the implementation model. One constraint is that *abs_cvsauth* is mapped to the right path and that the authentication tables in both models are equal:

$$\begin{aligned}
 Rname2path(abs_cvsauth) &= cvs_rep \\
 &\quad \wedge \langle CVSROOT, cvsauth \rangle \\
 authtab(rep) &= get_auth_tab(files)
 \end{aligned}$$

The last constraint we present here enforces the abstract working copy to have a counterpart in the implementation working copy:

$$Rname2path(\text{dom } wc) = \text{dom } wcs_attributes$$

To verify the refinement relation *R*, following Spivey in [16], we must prove two refinement conditions for each operation on the abstract state and its corresponding operation on the concrete state: Condition (a) ensures that a concrete operation terminates whenever their corresponding abstract operation is guaranteed to terminate, condition (b) ensures that the state after the concrete operation represents one of those abstract states in which the abstract operation could terminate.

As an example of the refinement, we show the instantiation of conditions (a) and (b) for the CVS **login** operation. The refinement conditions, though, as defined in [16], assume that both operations have the same input parameters, but since we define them differently in our two models, we introduce an additional schema *Asm*, which is used to insert further assumptions into the refinement proofs (the effect could also have been achieved by a suitable renaming):

$$\begin{array}{l}
 \overline{Asm} \\
 passwd?, cvs_pwd? : Cvs_Passwd \\
 uid?, cvs_uid? : Cvs_Uid \\
 \hline
 passwd? = cvs_pwd? \\
 uid? = cvs_uid?
 \end{array}$$

In the case of the login operation, these assumptions are simple since the parameters are of the same type but differ in name. Instantiating condition (a) and (b) for the login operation and adding the assumption schema *Asm*

leads to the following two proof obligations:

$$\begin{aligned}
login_a &\equiv \forall ClientState; RepositoryState; \\
&\quad ProcessState; Cvs_FileSystem; \\
&\quad passwd?, cvs_pwd? : Cvs_Passwd; uid?, \\
&\quad cvs_uid? : Cvs_Uid \bullet \\
&\quad Asm \wedge \text{pre } abs_login \wedge R \Rightarrow \text{pre } cvs_login \\
login_b &\equiv \forall ClientState; RepositoryState; \\
&\quad ProcessState; Cvs_FileSystem; \\
&\quad ProcessState'; Cvs_FileSystem'; passwd?, \\
&\quad cvs_pwd? : Cvs_Passwd; uid? \\
&\quad cvs_uid? : Cvs_Uid \bullet \\
&\quad Asm \wedge \text{pre } abs_login \wedge R \wedge cvs_login \\
&\quad \Rightarrow (\exists ClientState'; RepositoryState' \bullet \\
&\quad \quad R' \wedge abs_login)
\end{aligned}$$

The obligations for the other operations are defined analogously. So far, we proved these obligations formally for the refinement of `login`, `add` and `update`. These proofs considerably helped us to identify subtle side-conditions in our model and thus to get our real CVS configuration “right”.

4.3 Security Properties in Architecture Layers

Specifying the security properties motivates a Z-section for the system architecture and one for the implementation architecture, both containing a classical *behavioral* specification. In *SysArchSec* we investigate security properties of the system architecture. In *ImplArchSec* we investigate the same properties and additional ones that are specific to the implementation architecture.

4.3.1 The General Scheme of Security Properties

As an interface between the operation schemas of the two architecture layers and the behavioral part allowing to specify safety properties, we convert suitably restricted operation schemas of both system layers into explicit relations over the underlying state. The purpose of these restrictions is to provide a slot for side-conditions that are related to the security model and not the functional model described in the previous sections:

$$\begin{aligned}
rop_1 &= op_1 \wedge R_1 \\
&\dots \\
rop_n &= op_n \wedge R_n
\end{aligned}$$

where each rop_i represents the operation schema op_i constrained by the restriction schema R_i . Further the schema disjunction *step* represents the overall step relation of the system, which is converted into a transitively closed relation *trans*:

$$\begin{aligned}
step &= rop_1 \vee \dots \vee rop_n \\
trans &= \{step \mid (\theta state, \theta state')\}^*
\end{aligned}$$

In the literature, three types of properties can be distinguished: One may formalize properties over the *set of reachable states*, the *set of possible transitions* or the *set of possible sequences of states (traces)* of a system. While the first two types are only sufficient for classical safety invariants (“something bad will never happen”), the latter two allow for the specification of liveness properties (“eventually something good will happen”). The general scheme for properties over reachable states and possible transitions for safety properties and the schema for liveness properties looks as follows:

$$\begin{aligned}
SP_{RS} &= \forall \sigma : \text{trans}(\text{init}) \bullet P\sigma \\
SP_{RT} &= \forall (\sigma, \sigma') : \text{init} \triangleleft \text{trans} \bullet P(\sigma, \sigma') \\
LP_{RT} &= \forall (\sigma, \sigma') : \text{init} \triangleleft \text{trans} \bullet \\
&\quad \exists (\sigma'', \sigma''') : \text{trans} \bullet P(\sigma, \sigma', \sigma'', \sigma''')
\end{aligned}$$

Note that the reachable states are restricted via the existential image operator or the domain restriction to the states (respectively transitions) reachable from the set of initial states `init`.

4.3.2 An Instance of the General Scheme: *RBAC_write*

We will exemplify the scheme SP_{RT} for a crucial security property, namely “the user may write in the repository only if he has RBAC-permissions”, which we will call *RBAC_write* in the following. Moreover, we will outline the inductive proof.

As a prerequisite, we postulate two arbitrary sets *knows* and *invents*; a client “knows” a set of pairs of roles and passwords, and “invents” only files from a given set of pairs from names to data. We assume *invents* to be closed under the *merge*-operation left abstract in our model.⁴ On this basis, we define a *security policy*, by providing suitable restrictions $op_i R$ for the system operations.⁵ For example, we restrict the add operation to elements in the domain of the *invents*-set, we assume `login` being restricted to roles and passwords the client *knows* set, the modify operation and add being restricted to data the client “invents”. While these restrictions have a more technical nature, a more conceptual restriction of *abs_ci* is as follows: in the role *cvs_adm*, the authentication table may only be altered such that rights are withdrawn, not granted. A typical restriction looks as follows:

$$\begin{aligned}
abs_loginR &\equiv abs_login \\
&\quad \wedge [cvs_uid? : Cvs_Uid; passwd? : \\
&\quad \quad Cvs_Passwd \mid \\
&\quad \quad (cvs_uid?, passwd?) \in Aknows]
\end{aligned}$$

⁴ This is very similar to the concept of abstract crypt-functions and the closures *analz*, *synth* and *parts* in [10]; see discussion

⁵ In practice, such security policies may be based on voluntary self-restrictions of users or enforced by administrative means.

Now we define the *step*-relation and its transitive closure of the system architecture layer:

$$\begin{aligned} \text{step} &\equiv \text{abs_loginR} \vee \text{abs_addR} \vee \text{abs_ciR} \\ &\quad \vee \text{abs_modifyR} \vee \text{abs_up} \vee \text{abs_cd} \\ \text{AbsState} &\equiv \text{ClientState} \wedge \text{RepositoryState} \\ \text{trans} &\equiv \{\text{step} \bullet (\theta \text{AbsState}, \theta \text{AbsState}')\}^* \end{aligned}$$

Finally, for constructing the proof goal *RBAC_write*, we instantiate the *P* in our schema *SP_{RT}* by:

$$\left[\begin{array}{l} \text{rbac_write_} : \dots \\ \hline \forall \text{rep}, \text{rep}' : \text{ABS_DATATAB}; \\ \forall \text{rptab}' : \text{ABS_PERMTAB} \bullet \\ \quad \text{rbac_write}(\text{rep}, \text{rep}', \text{rptab}') \Leftrightarrow \\ \quad (\forall f : \text{dom } \text{rep}' \bullet \\ \quad \quad (\text{rep}(f) \neq \text{rep}'(f) \\ \quad \quad \wedge (f, \text{rep}'(f)) \in \text{invents}) \\ \quad \Rightarrow (\exists m : \text{knows} \bullet \\ \quad \quad (\text{rptab}'(f), \text{authtab}(\text{rep}')(m)) \\ \quad \quad \in \text{cvs_perm_order})) \end{array} \right.$$

This property reads as follows: whenever there is a change in the repository, and the changed file stems from the users *invents*-set, the user must have valid permissions according to the *RBAC*-model. We observe that *rbac_write* is *true* whenever the repository does not change, i.e. *rbac_write*(*r*, *r*, *rt*) holds.

4.3.3 A Proof-Outline

We will now present an exemplary proof (performed with HOL-Z) for *RBAC_write*. The initial proof goal stating that *RBAC_write* holds is refined by unfolding elementary definitions and simplification of Z notation to the following proof state:

$$\begin{aligned} &\llbracket \sigma_0 = (\text{abs_passwd}, \text{rep}, \text{rep_permtab}, \text{wc}, \\ &\quad \text{wc_uidtab}, \text{wfiles}); \\ &\sigma_1 = (\text{abs_passwd}', \text{rep}', \text{rep_permtab}', \text{wc}', \\ &\quad \text{wc_uidtab}', \text{wfiles}'); \\ &\text{AbsState}\sigma_0; \\ &\text{AbsState}\sigma_1; \\ &(\sigma_0, \sigma_1) : \{\text{step} \bullet (\sigma_0, \sigma_1)\}^* \\ &\rrbracket \Rightarrow \\ &\quad \sigma_0 : \text{init} \\ &\quad \Rightarrow \text{rbac_write}(\text{rep}, \text{rep}', \text{rep_permtab}') \end{aligned}$$

Over this implication, we can now apply an induction rule over the transitive closure:

$$\frac{\begin{array}{ccc} [x \in \text{dom } r] & [y \in \text{ran } r] & \left[\begin{array}{l} Pxy, \\ (x, y) \in r^*, \\ (y, z) \in r \end{array} \right] \\ \vdots & \vdots & \vdots \\ (a, b) \in r^* & Pxx & Pyy & Pxz \end{array}}{Pab}$$

This leads to two base cases and the induction step; both base cases are trivially true due to observation *rbac_write*(*r*, *r*, *rt*). Now the induction steps, which looks after some massage as follows, remains to show:

$$\begin{aligned} &\llbracket \dots \\ &\sigma_{00} = (\text{abs_passwdx}, \text{repx}, \text{rep_permtabx}, \text{wcx}, \\ &\quad \text{wc_uidtabx}, \text{wfilesx}); \\ &\sigma_{01} = (\text{abs_passwdy}, \text{repy}, \text{rep_permtaby}, \text{wcy}, \\ &\quad \text{wc_uidtaby}, \text{wfilesy}); \\ &\sigma_{10} = (\text{abs_passwdz}, \text{repz}, \text{rep_permtabz}, \text{wcz}, \\ &\quad \text{wc_uidtabz}, \text{wfilesz}); \\ &\sigma_{00} : \text{init} \Rightarrow \text{rbac_write}(\text{repx}, \text{repy}, \text{rep_permtaby}); \\ &(\sigma_{00}, \sigma_{01}) : \{\text{step} \bullet (\sigma_0, \sigma_1)\}^*; \\ &(\sigma_{00}, \sigma_{10}) : \{\text{step} \bullet (\sigma_0, \sigma_1)\} \\ &\rrbracket \Rightarrow \sigma_{00} : \text{init} \\ &\quad \Rightarrow \text{rbac_write}(\text{repx}, \text{repz}, \text{rep_permtabz}) \end{aligned}$$

Here, the point of proof refinement is the assumption $(\sigma_{00}, \sigma_{01}) : \{\text{step} \bullet (\sigma_0, \sigma_1)\}^*$, which can be decomposed via the definition of *step* into a disjunction of schemas, where the input variables are existentially quantified. A generic tactic strips away the disjunctions and the existential quantifiers in the assumption. The result is a case split over all operations of the system architecture and universally quantified input parameters of all operations under consideration. Now, the observation is crucial that all operations except *abs_ci* do not change the repository, and, as a consequence of observation *rbac_write*(*r*, *r*, *rt*), imply the truth of the step. We can therefore focus on the case *abs_ci*:

$$\begin{aligned} &\llbracket \dots \\ &(\sigma_{00}, \sigma_{01}) : \{\text{step} \bullet (\sigma_0, \sigma_1)\}^*; \\ &\text{rbac_write}(\text{repx}, \text{repy}, \text{rep_permtaby}); \\ &\text{abs_ci}(\text{abs_passwdy}, \text{abs_passwdz}, \text{filesq}, \text{repy}, \text{repz}, \\ &\quad \text{rep_permtaby}, \text{rep_permtabz}, \text{wcy}, \text{wcz}, \\ &\quad \text{wc_uidtaby}, \text{wc_uidtabz}, \text{wfilesy}, \text{wfilesz}) \\ &(\sigma_{00}) : \text{init}; \\ &\rrbracket \Rightarrow \text{rbac_write}(\text{repx}, \text{repz}, \text{rep_permtabz}) \end{aligned}$$

This is the core part of an invariance proof: the system made a transition from an initial system state (with *repx*) to another (with *repy*) performing an arbitrary combination of operations *and* the system behaved well (i.e. *rbac_write*(*repx*, *repy*, *rep_permtaby*)). Now a commit operation (*abs_ci*) occurs, and the question is if the resulting state (with *repz*) will also fulfill our safety property.

The core of this subproof is, of course, a case distinction following the definition of *abs_ci* shown in Sec. 3.2: a file may be

1. in the repository *and not* in the working copy: then *abs_ci* will change nothing,

2. in the working copy *and not* in the repository: then *abs_ci* will only change the latter if the current credentials are *is_valid_in* which implies *write_correct* as the *rep_permtab* was changed accordingly,
3. both in the working copy *and* the repository: then *abs_ci* will only change the file in the repository if the current credentials allow for *has_access* which implies *write_correct*.

The interested reader may note that the overall scheme of the proof follows the structure of the general scheme of the property descriptions, which allows for automated tactic support that copes with Z-related technicalities, the choice of the inductions, the decomposition of the specification and the systematic derivation of state components remaining invariant. Obviously, there is a high potential of automation for this type of proofs, such that the proof developer may be guided rather automatically to the critical questions in the induction step.

4.3.4 Other Examples

The verification of the analogous property *RBAC_read* is straight forward; files in the working copy of a client are either invented by him (via the operation *modify*) or stem from the repository, where the client knows a password to obtain sufficient permissions.

An important, but quite obvious liveness property in the *LP_{RT}*-scheme is *RBAC_do_write*: Provided the client has access, it can change a file arbitrarily and perform operations leaving the repository changed accordingly; the proof immediately boils down to *abs_ci* which is designed to fulfill this property. At first sight, *RBAC_do_write* looks very similar to *RBAC_write*, however, note that both properties are independent: one could model an absolutely secure CVS-Server that never changes the repository. Such a model trivially fulfills *RBAC_write*, but is ruled out by *RBAC_do_write*.

So far, *RBAC_write* is formalized for a single-user client/server setting. Extending the analysis to a multi-user client/server model only requires simple modifications in the definition of the *step*-relation; via renaming of the working copies and the *invents* and *knows*-sets, instances of *abs_ci*, *abs_up* and *modify* for each client with individual working copy can be generated. Adding suitable restrictions (e.g. *invents* and *knows*-sets must be pairwise disjoint), *RBAC_write* and similar properties remain valid.

It is well-known that security properties are usually not preserved under refinement (see discussion later). The reason is that *implementing* one security architecture by another opens the door to *new* types of attacks on the implementation architecture that can be completely overlooked on the abstract level. For example, on the implementation architecture, it is possible to realize an attack on the repository by combinations POSIX commands such as *rm* and *setumask* etc (see Sec.3.3.2). In principle, our method can be applied for this type of

analysis of the implementation architecture as well. In this setting, the *step*-relation and the *init* is defined as:

$$\begin{aligned} \text{step}_{\text{impl}} &\equiv \text{rm} \vee \text{setumask} \vee \dots \vee \text{chmod} \\ &\quad \vee \text{cvs_login} \vee \dots \vee \text{cvs_update} \\ \text{init}_{\text{impl}} &\equiv \text{ConcState} \\ &\quad \wedge [\text{wcs_attributes} : \text{CVS_ATTR_TAB} \mid \\ &\quad \quad \text{wcs_attributes} = \emptyset] \end{aligned}$$

Although the proofs on the implementation architecture have the same structure as on the system architecture, they are far more complex since concepts such as paths, the distinction between files and directories, and their permissions are involved. Moreover, they require new side-conditions (for example, the refinement can only be established for the case that the user is not root) which were systematically introduced by the abstraction predicate *R*.

On the other hand, the higher degree of detail on the implementation architecture makes a formalization of new types of security properties possible: For example, since the crucial concept *directory* is present on the implementation level and since the existence of files can only be established by having access to all parent directories of a file, one can express confidentiality properties such as “the user can not find out that a file with name *x* exists in some directory of the repository” on this level.

5 Conclusion

5.1 Discussion

We demonstrate a method for analyzing the security in off-the-shelf system components thus made amenable to formal, machine-based analysis. The method proceeds as follows: First, specify the system architecture (as a framework for formal security properties), second, specify the implementation architecture (validated by inspecting informal specifications or testing code), third, set up the security technology mapping as a refinement, and fourth, prove refinements and security properties by mechanized proofs. The demonstration of the method follows a case study of a security problem for a real system, the CVS client/server architecture. We believe that the method is applicable for a wider range of problems such as mission-critical e-commerce applications or e-government applications.

The core of our approach is based on the presentation of the security technology mapping as data refinement problem. In general, it has been widely recognized that security properties can not be easily refined — actually, finding refinement notions that preserve security properties is a hot research topic [8,14]. However, standard refinement proof technology has still its value here since it checks that abstract security requirements are indeed achieved by a mapping to concrete security technology,

and that implicit assumptions on this implementation have been made explicit. Against implementation specific attacks, we believe that specialized security property refinement techniques will be limited to restricted aspects. For this problem, in most cases the answer will be an analysis on the implementation level, possibly by reusing results from the abstract level.

In our approach, the analysis is based on interactive theorem proving while security analysis is often based on model-checking techniques for logics like LTL, the μ -calculus or process algebras like CSP. While these techniques offer a high degree of automation, they possess well-known and obvious limitations: the state-space must usually be finite and in practice be very small, and the analysis tends to be infeasible for many models, in particular those imposed by system specifications. As a consequence, proof engineers tend to develop oversimplified and unsystematically abstracted system models. In contrast, in our approach technical concerns like the size of the system state-space, aesthetic concerns like naturalness of the modeling (in our example, we use architectural modeling) or methodological needs like realistic treatments of system specifications do not represent fundamental obstacles to the analysis. In particular the latter paves the way for the reuse of standard system models like POSIX. Moreover, we have the full flexibility of Z and HOL to express security properties at need.

Our case study shows that the presented technology and method makes the treatment of complex security problems possible. Naturally, the question arises how long the formalization and the proof work took. This question is hard to answer, partly because the method and technical components had been developed during the project, partly because library theorems had to be proven, partly because some contributors needed time to learn Isabelle. The overall case study took about 18 man months, including the development of tool support. A substantial part (about six man months) was the formalization and testing (i.e. reverse engineering) of the system which was done by Achim Brucker and Burkhart Wolff. The proof work was done by Frank Rittinger, Harald Hiss and Burkhart Wolff. Using our tool support, we estimate that someone with experience in theorem proving would be able to solve a similar task (specifying a similarly complex system and proving the core security properties) in less than 10 months. By improving the general technology (e.g. better front-ends and tactic support) a further speed-up by a factor two seems feasible.

5.2 Related Work

Sandhu and Ahn described in [12] a method for embedding role-based access control with the discretionary access control provided by standard Unix systems. Our model used this construction for providing the static roles, but extended it to a dynamic model.

Wenzel developed a specification of the basic Unix functionality, which was done in Isabelle/HOL and is part of the actual Isabelle [9] distribution. On the file system part, only a simple access model, not supporting groups and the concepts of set-id bits, is formalized.

Our behavioral analysis is based on the same foundations as Paulson's inductive method for protocol verification [10]. Beyond the obvious difference, that Paulson research focus is on analysis (the language of protocols is deliberately small and restrictive) and not on modeling, technical differences consist merely in some details: Paulson uses specialized induction schemes which are automatically derived from the protocol-rules; these are considered as inductive rules defining the set of system traces. In contrast, we use standard induction over transitive relations, which leads to a different organization of the specification and the security properties and leads to different tactic support.

5.3 Future Work

In our opinion, amazingly little work has been addressed to the specification of the POSIX interface; due to its often not intuitive features, its importance for security implementations and its high degree of reuse, this is a particularly rewarding target. We believe that our formalization is a starting point for a comprehensive, more complete model of the filesystem related commands.

Clearly, the formal proofs established so far do not represent a *complete* analysis of the (real) CVS-Server. Many more security properties can be formulated, and, by setting up different operation restrictions R_i , "best-practice" security policies can be formally investigated. Moreover, in order to make implementation level security analysis more feasible, it could be highly rewarding to develop techniques and methods to reuse (abstract) system level proofs on the more concrete levels.

References

1. A. D. Brucker, F. Rittinger, and B. Wolff. A CVS-Server security architecture — concepts and formal analysis. Technical Report 182, Albert-Ludwigs-Universität Freiburg, 2002.
2. A. D. Brucker, F. Rittinger, and B. Wolff. HOL-Z 2.0: A proof environment for Z-specifications. *Journal of Universal Computer Science*, 9(2):152–172, 2003.
3. P. Cederqvist et al. *Version Management with CVS*, 2000. <http://www.cvshome.org/docs/manual/>.
4. K. Fogel. *Open source development with CVS*. The Coriolis Group, 1999.
5. Æ. Frisch. *Essential System Administration*. O'Reilly, 1995.
6. D. Garlan and M. Shaw. An introduction to software architecture. In *Advances in Software Engineering and Knowledge Engineering*, pages 1–39. World Scientific Publishing Company, 1993.

7. M. J. C. Gordon and T. F. Melham. *Introduction to HOL*. Cambridge University Press, 1993.
8. J. Jürjens. Secrecy-preserving refinement. In *Formal Methods Europe (FME)*, volume 2021 of *LNCS*. Springer Verlag, 2001.
9. T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer Verlag, 2002.
10. L. C. Paulson. The inductive approach to verifying cryptographic protocols. *Journal of Computer Security*, 6:85–128, 1998.
11. A. Roscoe. *Theory and Practice of Concurrency*. Prentice Hall, 1998.
12. R. Sandhu and G.-J. Ahn. Decentralized group hierarchies in UNIX: An experiment and lessons learned. In *National Information Systems Security Conference*, pages 486–502, 1998.
13. R. S. Sandhu, E. J. Coyne, H. L. Feinstein, and C. E. Youman. Role-based access control models. *Computer*, 29(2):38–47, 1996.
14. T. Santen, M. Heisel, and A. Pfitzmann. Confidentiality-preserving refinement is compositional — sometimes. In *ESORICS*, volume 2502 of *LNCS*, pages 194–211. Springer Verlag, 2002.
15. M. Shaw and D. Garlan. *Software Architecture: Perspectives on an Emerging Discipline*. Prentice Hall, 1996.
16. J. M. Spivey. *The Z Notation: A Reference Manual*. Prentice Hall, 1992. <http://spivey.oriel.ox.ac.uk/~mike/zrm/>.
17. *The Single UNIX Specification Version 3*. The Open Group and IEEE, 2002. This standard supersedes the “Single UNIX Specification Version 2” (Unix 98) and the “IEEE Standard 1003.1-2001” (POSIX.1).
18. J. Woodcock and J. Davies. *Using Z: Specification, Refinement, and Proof*. Prentice Hall International Series in Computer Science. Prentice Hall, 1996. <http://www.usingz.com/>.