

**SOFTWARE VERIFICATION RESEARCH CENTRE  
SCHOOL OF INFORMATION TECHNOLOGY  
THE UNIVERSITY OF QUEENSLAND**

**Queensland 4072  
Australia**

**TECHNICAL REPORT**

**No. 98-10**

**Supporting Fine-grained Traceability  
in Software Development  
Environments**

**Peter Lindsay and Owen Traynor**

**July 1998**

**Phone: +61 7 3365 1003**

**Fax: +61 7 3365 1533**

**<http://svrc.it.uq.edu.au>**

This report is an expanded version of the paper which appeared in B. Magnusson (ed), *Proceedings 8th International Symposium on System Configuration Management*, Brussels, Springer Verlag LNCS 1439, July 1998, pages 133-139.

**Note:** Most SVRC technical reports are available via anonymous ftp, from `svrc.it.uq.edu.au` in the directory `/pub/techreports`. Abstracts and compressed postscript files are available via `http://svrc.it.uq.edu.au`

# Supporting Fine-grained Traceability in Software Development Environments

Peter Lindsay and Owen Traynor

## Abstract

This paper describes the facilities currently available to support auditing and traceability within a system which provides fine-grained configuration and version management. We contend that the relationship between the configuration management system and the underlying version control system is a critical factor which governs many aspects of the facilities supporting traceability. The model of traceability is formally specified relative to our configuration and versioning models.

## 1 Introduction

Managing and controlling change is a critical part of software engineering. Software components typically pass through many different versions during both the initial development of a system and the ongoing maintenance of the system once deployed. The facilities that are available in such systems for tracing the evolution of requirements through the design, coding, validation and verification stages are especially critical. Such facilities are particularly important in the maintenance phase and, where conformance to standards is required, in demonstrating that such standards have been met. Conformance to such standards is not only a requirement of the initial system development, but also an ongoing requirement throughout the lifetime and evolution of a software system.

The history that documents the evolution of a software system is, in essence, the embodiment of that system. We believe that an accurate account of that history is critical in assessing the worth of the deployed system and in ensuring that subsequent developments of that system are made in a coherent and consistent fashion [2]. The history of a system also provides a wealth of information regarding design decisions and implementation choices [10, 15]. Access to a clear and concise account of such information may go a long way to reducing the effort required to rework or redevelop systems in the context of changes to requirements. As a management tool, this information provides valuable feedback regarding the design choices made, and processes followed, in the construction of a system.

We believe that traceability facilities that allow the documentation of a system at a finer level of granularity than the tradition build, or baseline models, provide access to essential information that is often lost in these traditional approaches. Such facilities also provide the information needed to reduce the effort required to rework or redevelop systems in the context of changes to requirements.

### 1.1 Software Configuration Management (SCM)

As well as providing the framework within which developers work to construct consistent system “builds”, SCM provides the mechanisms needed to demonstrate traceability between the built system, the design, the requirements documents, and other tools and artifacts of a development (such as compilers and test reports). SCM provides the means for recording and controlling the “configuration” of versions of documents associated with software development, including inter- and intra-document dependencies. Regulatory and standards authorities have long recognised the importance of reliable SCM mechanisms, especially in the development of high integrity software systems [3, 8, 9].

An important supporting technology for SCM is *version control*, which concerns storage and retrieval of different versions of development components. Most version control systems attempt to maintain a record of the changes (“deltas”) between different versions of components. This provides the basis for tracing the evolution of a system through its lifetime. The majority of software development companies currently use version control facilities such as RCS [16] or SCCS [11] to manage their documents and software, but such facilities operate at an inadequately coarse level of granularity (typically, whole documents or whole modules) and fall far short of users’ desires.

In defining a coherent framework within which we can provide useful traceability functions, we will see that the core support for SCM (including version control) has substantial impact of the amount of effort needed to implement traceability support. We will focus on a simple example to illustrate these issues (a document conformance system) and illustrate the benefits accrued from our configuration management models from the traceability viewpoint.

### 1.2 Configuration Management (CM) for Formal Methods

*Formal Methods* of software development have particular needs in relation to CM. Formal Methods are based on the use of mathematically precise definitions of development components and their relationships, together with the use of mathematical analysis techniques – including theorem proving – for establishing correctness. The fact that individual development components have mathematical meaning makes it possible to formally verify that desired relationships hold

within and between development components. In contrast to traditional development methods, cross-development configuration consistency can be defined precisely and at fine levels of granularity [13].

Consistent with this observation is our view of traceability. Our traceability model allows us to track, at fine levels of granularity, the changes that a system has undergone that moves the system through its evolving, consistent, versions. Since we are working in a context where relationship and dependencies are formally modelled, we can use this as a basis for defining formal models of traceability. Such an approach offers a great deal in the context of high integrity system development. As well as being able to trace the development of systems, we would like to be able to trace (in isolation as far as possible) the development history of individual fine-grained artifacts. In critical system development it is important to be able to trace the evolution of individual safety requirements right through the design to final implementation [9].

A very substantial side effect of our model is that of the traceability framework we propose, is that, as well as providing facilities to trace the evolution of the development of a system in terms of the individual paragraphs of a requirements document, or the declarations of a formal specification, it may be used to assess the impact of a change to a given requirement. Having a detailed history of the development of an individual requirement, for example, also provides a basis for assessing the impact of changing that requirement.

The ARC-funded *Fine-Grained Configuration Management (FGCM)* project at the SVRC is establishing a framework for fine-grained configuration management. The framework builds on a programme of work carried out by PhD student Kelvin Ross under the supervision of the first author, investigating the application of SCM techniques to formal development [12, 13, 14]. The aim of the framework is to allow developers to support their correctness claims with evidence that, not only have the individual components of a system been developed correctly, but that the combination and integration of these components has been done in a consistent manner and that the final result is derived from consistent, complete and up-to-date development components. The framework is intended to apply not only to specifications, designs and programs, but also to fine-grained development components such as the specification components, reviews, change requests, refinements, design decisions, test sets, theories and proofs that are generated as part of the development process [6].

### 1.3 High integrity software engineering

We consider the definition of a coherent framework, within which configuration and version management can be carried out, as an important prerequisite in the development of trusted and cost-effective environments for the development of critical software. The processes that define the development activities in such trusted environments must be based on sound underlying technology and models that allow the impact of any development step (in terms of the consistency of

the relationships between the underlying artifacts) to be accurately assessed. Existing Software Engineering Environments (SEEs) use relatively untrusted standard version-management technology in the development of critical systems; this is clearly a weak link since these technologies have no coherent formal basis for consistency checking.

The need for careful control of the development process must be balanced against the need for flexibility. Users will not accept a development process that is overly constraining. Similarly, it is vitally important for encouraging industrial uptake that Formal Methods be adaptable to different situations, project structures and so on, without sacrificing the trustworthiness of the environments.

Our approach has been to define configuration consistency models (or *configuration models*, for short) which define the key configuration items, the relationships between them, and the consistency and completeness conditions desired for the configuration. In our approach, configuration models would form the core part of SEEs, with development processes defined relative to the core models. This means that the consistency and completeness of a development could be established largely independently of the development process applied, giving flexibility and trustworthiness in the one framework [4, 5, 7].

In the spirit of these configuration models, we take a similar approach when defining the notions of traceability. Our trace models are defined relative to the core configuration models. Again, the form of this definition allows us to establish consistency criteria for our framework. Ensuring that such consistency criteria are met, is an important issue, especially in the context of high-integrity development.

## 1.4 This report

This report gives preliminary conclusions from a case study in adding fine-grained versioning, configuration control and traceability to a simple document conformance system. We concentrate here of presenting our results from the traceability viewpoint. Section 2 describes the general framework within which we are working and introduces the example that is used in the rest of the paper. Sections 3–7 formally specify the data model on which the framework is based. Conclusions and future work are presented in section 8.

# 2 The case study

## 2.1 Documents

We start by felling some trees to better see the forest. We shall consider part of a development consisting simply of two documents – called A and B here, for short – which consist of sets of requirements and which are expected to

conform with one another in some way. For example, A might be a Software Requirements Specification (SRS) and B an Architectural Requirements Specification, describing a module structure and how the requirements from the SRS are allocated to modules: see Fig. 1.

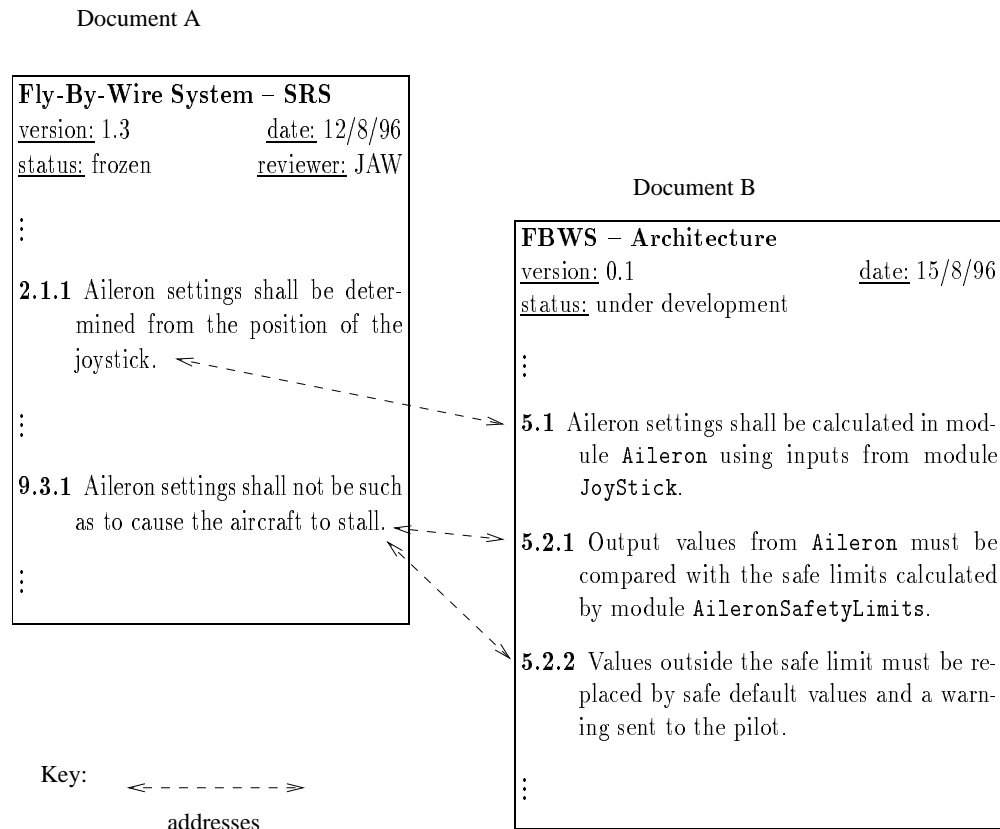


Figure 1: An example of conforming documents, illustrating how requirements in one are addressed by requirements in the other.

## 2.2 The dependency relation

The dependency relation to be managed is the relationship which describes how requirements in B address requirements in A. In general, such a relationship is many-to-many: e.g. a single SRS requirement may be addressed in different places in the architecture document, and a single architectural requirement may address multiple SRS requirements. The *addresses* relation is analogous in some ways to the ‘is up to date with respect to’ notion for traditional SCM systems, as

used by the Unix MAKE facility for example. Note however that for informal objects such as requirements, the relation is user-determined and cannot be automated.

A *conformance matrix* (or traceability table) is a common means of indicating the relationship between two development documents. For each item in one of the development documents, the matrix lists the corresponding items in the other document, typically by paragraph or section number. In our case, the conformance matrix will record where each individual requirement in the SRS is addressed in the architecture document: see Fig. 2.

SRS (v1.3)	Architecture (v0.1)
⋮	
2.1.1	..., 5.1, ...
⋮	
9.3.1	..., 5.2.1, 5.2.2, ...
⋮	

Figure 2: Part of the conformance matrix for the example above.

### 2.3 Version and configuration management

For simplicity, we use a simple “version tree” model as the basis for version control, for objects of all granularities: paragraphs, documents, matrices, etc. This part of the model is reasonably generic, but could easily be replaced by something more appropriate if desired.

The configuration management problem for the case study is to manage the “configuration” consisting of the two documents and their conformance matrix. We say the conformance matrix is *complete* if each of A’s requirements are addressed somewhere in B. As part of the case study, we shall assume that at any time the conformance matrix may be incomplete, but it is always *correct*: i.e., that if the matrix says that  $r$  addresses  $t$ , then this relationship was determined by a user and the requirements have not changed in the interim.

It is not reasonable to expect to be able to develop generic version and configuration management functionality – especially at fine levels of granularity – because of the widely varying requirements of different applications, development methodologies, company policies, etc. However, as we shall show, it is possible to develop a reasonably generic framework in which generic definitions and functionality are supplemented and instantiated by application-specific *version and configuration management (V&CM) policies*.

Let us suppose, for the purposes of the case study, that the following V&CM policy is in operation for coarse-grained configuration items:



- all (and only) frozen versions of documents will be stored in the project archive;
- document A can be frozen at any time;
- document B can be frozen only after a review confirms that the appropriate version of A has been frozen and the conformance matrix is complete.

As the case study progresses, we shall introduce V&CM policies appropriate to the finer-grained objects involved.

## 2.4 Change tracking

In our model we require that modifications to (evolution of) individual requirements are carried out within a pre-defined framework. Experience suggests that the following set of change types is a useful factoring of concerns:

**add:** create a new paragraph with no prior history

**split:** create a number of new paragraphs by splitting an existing paragraph

**combine:** create a new paragraph by combining existing paragraphs

**delete:** delete a paragraph

**replace:** replace existing paragraphs by a new paragraph

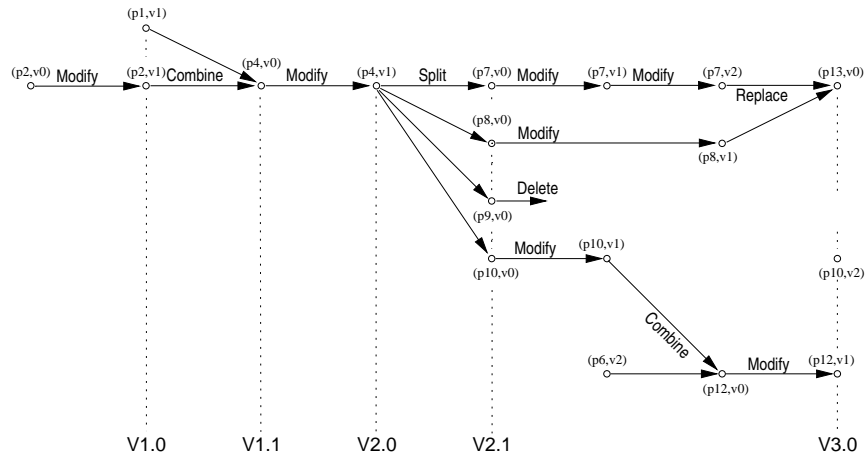
**modify:** modify the paragraph without changing its meaning

We shall extend our base model for documents with information relevant to how the current version of the document has changed since its last major version. The major versions of documents define baselines relative to which these *delta lists* (or change logs) are constructed. Note that in principle the modelling of the documents themselves need not be changed: rather the versioning models that are inherited are extended with the enriched concept of change.

Fig. 3 illustrates one kind of traceability (“forwards tracing”) that will be possible as a result of our approach: given a specific version (p2,v0) of a requirement, report how the requirement changed subsequently. Fig. 4 illustrates backwards tracing: given a specific version (p13,v0) of a requirement, report the evolution that resulted in the requirement.

## 2.5 Requirements tracing

The second major form of traceability we provide is concerned with tracking dependencies – an important part of high integrity development and one of the main mechanisms required in development audit and evaluation. There are many ways in which one might want to trace the evolution of an individual requirement through a development. For our case study, the kinds of checks one might want to apply include the following:



**Key:** V1.1 Version number of Document A  
 (p3,v0) PVersion, i.e. (paragraph Id, version number)

Figure 3: Forwards tracing of the evolution of a given paragraph version

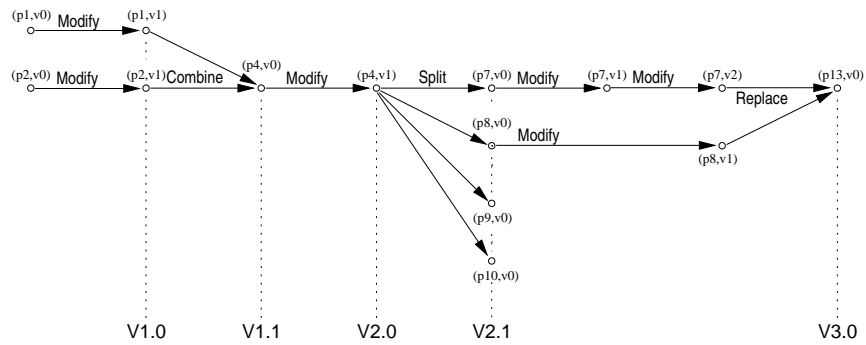


Figure 4: Backwards tracing of the evolution of a given paragraph version

- given a requirement in A, find which requirements in B address it;
- given a B requirement, find which A requirements it addresses;
- find which A requirements have not yet been addressed;
- find which B requirements are extraneous (i.e., do not address an A requirement);
- report the evolution of a given requirement (i.e., the version of the in which it originated, and how it changed in subsequent versions of the document).

### 3 Version control

Our simple model for version control is based on *forests* of version trees and parameterised over the type *Type* of objects being placed under version control. The attributes of version forests are defined as follows:

- Each node in a version tree is labelled with a unique *version label*, from the set *VLabel*.
- The mapping *deref* ‘dereferences’ version labels, yielding the content of a particular version.
- The mapping *parent* returns the (unique) parent of a non-root node. The mapping is acyclic.
- The set *frozen* contains the labels of the frozen nodes.

$VForest[Type]$ $nodes : \mathbb{P} VLabel$ $deref : VLabel \leftrightarrow Type$ $parent : VLabel \leftrightarrow VLabel$ $frozen : \mathbb{P} VLabel$ $roots : \mathbb{P} VLabel$ <hr/> $nodes = \text{dom } deref$ $\text{dom } parent \cup \text{ran } parent \cup frozen \subseteq nodes$ $roots = nodes \setminus \text{dom } parent$ $\forall v : VLabel \bullet (v, v) \notin parent^+$
---

As a V&CM policy decision, parents will be required to be frozen:

$\text{ran } parent \subseteq frozen$
---------------------------------------

## 4 Paragraphs

### 4.1 Paragraph version management

Paragraphs will be identified by a label from the type  $PId$ . A specific *paragraph version* consists of the paragraph's identifier and a specific version label:

$$PVersion ::= PId \times VLabel$$

The collection of *paragraph changes* is modelled as follows:

$$PChange ::= \begin{array}{l} add\langle PVersion \rangle \\ delete\langle PVersion \rangle \\ split\langle PVersion \times PVersion^+ \rangle \\ combine\langle PVersion^+ \times PVersion \rangle \\ derive\langle PVersion^+ \times PVersion \rangle \\ replace\langle PVersion^+ \times PVersion \rangle \\ modify\langle PId \times VLabel \times VLabel \rangle \end{array}$$

The following defines a predicate for checking whether a change affects a given paragraph version:

$$\begin{array}{l} \hline isChangedBy : PVersion \leftrightarrow PChange \\ \hline \neg isChangedBy(opv, add(pv)) \\ isChangedBy(opv, delete(pv)) \Leftrightarrow opv = pv \\ isChangedBy(opv, split(pv, pvs)) \Leftrightarrow opv = pv \\ isChangedBy(opv, combine(pvs, pv)) \Leftrightarrow opv \in \text{ran } pvs \\ \neg isChangedBy(opv, derive(pvs, pv)) \\ isChangedBy(opv, replace(pvs, pv)) \Leftrightarrow opv \in \text{ran } pvs \\ isChangedBy(opv, modify(a, u, v)) \Leftrightarrow opv = (a, u) \end{array}$$

The following predicate checks whether a change creates a given paragraph:

$$\begin{array}{l} \hline isCreatedBy : PVersion \leftrightarrow PChange \\ \hline isCreatedBy(npv, add(pv)) \Leftrightarrow npv = pv \\ \neg isCreatedBy(npv, delete(pv)) \\ isCreatedBy(npv, split(pv, pvs)) \Leftrightarrow npv \in \text{ran } pvs \\ isCreatedBy(npv, combine(pvs, pv)) \Leftrightarrow npv = pv \\ isCreatedBy(npv, derive(pvs, pv)) \Leftrightarrow npv = pv \\ isCreatedBy(npv, replace(pvs, pv)) \Leftrightarrow npv = pv \\ isCreatedBy(npv, modify(a, u, v)) \Leftrightarrow npv = (a, v) \end{array}$$

A check that two changes don't create – or try to change – the same thing:

$$\begin{array}{|l}
\hline
\textit{noninterfering} : P\textit{Change} \leftrightarrow P\textit{Change} \\
\hline
\textit{noninterfering}(c_1, c_2) \Leftrightarrow \\
\quad \forall pv : P\textit{Version} \bullet \\
\quad \quad \neg (\textit{isCreatedBy}(pv, c_1) \wedge \textit{isCreatedBy}(pv, c_2)) \wedge \\
\quad \quad \neg (\textit{isChangedBy}(pv, c_1) \wedge \textit{isChangedBy}(pv, c_2))
\end{array}$$

A *change set* is a set of noninterfering changes:

$$P\textit{ChangeSet} == \{cs : \mathbb{P} P\textit{Change} \mid \forall c_1 \neq c_2 \in cs \bullet \textit{noninterfering}(c_1, c_2)\}$$

Two change sets are disjoint if their elements are pairwise noninterfering:

$$\begin{array}{|l}
\hline
\textit{disjoint} : P\textit{ChangeSet} \leftrightarrow P\textit{ChangeSet} \\
\hline
\textit{disjoint}(cs_1, cs_2) \Leftrightarrow \forall c_1 \in cs_1, c_2 \in cs_2 \bullet \textit{noninterfering}(c_1, c_2)
\end{array}$$

## 4.2 Paragraph collections

A collection of *specific paragraph versions*, in which each paragraph is represented at most once, is modelled as follows:

$$P\textit{Collection} == P\textit{Id} \multimap V\textit{Label}$$

The following function applies a change to a paragraph collection, if it makes sense to do so:

$$\begin{array}{|l}
\hline
\textit{applyChange} : P\textit{Collection} \times P\textit{Change} \multimap P\textit{Collection} \\
\hline
a \notin \text{dom } pc \Rightarrow \textit{applyChange}(pc, \textit{add}(a, v)) = pc \cup \{a \mapsto v\} \\
pv \in pc \Rightarrow \textit{applyChange}(pc, \textit{delete}(pv)) = pc \setminus \{pv\} \\
pv \in pc \wedge \forall i : 1..n \bullet a_i \notin \text{dom } pc \wedge \forall j : 1..i-1 \bullet a_j \neq a_i \\
\quad \Rightarrow \textit{applyChange}(pc, \textit{split}(pv, ((a_1, v_1), \dots, (a_n, v_n)))) \\
\quad = (pc \setminus \{pv\}) \cup \{i : 1..n \bullet a_i \mapsto v_i\} \\
\text{ran } pvs \subseteq pc \wedge a \notin \text{dom } pc \\
\quad \Rightarrow \textit{applyChange}(pc, \textit{combine}(pvs, (a, v))) = (pc \setminus \text{ran } pvs) \cup \{a \mapsto v\} \\
\text{ran } pvs \subseteq pc \wedge a \notin \text{dom } pc \\
\quad \Rightarrow \textit{applyChange}(pc, \textit{derive}(pvs, (a, v))) = pc \cup \{a \mapsto v\} \\
\text{ran } pvs \subseteq pc \wedge a \notin \text{dom } pc \\
\quad \Rightarrow \textit{applyChange}(pc, \textit{replace}(pvs, (a, v))) = (pc \setminus \text{ran } pvs) \cup \{a \mapsto v\} \\
pc(a) = u \Rightarrow \textit{applyChange}(pc, \textit{modify}(a, u, v)) = pc \oplus \{a \mapsto v\}
\end{array}$$

The clauses in the above definition are intended to be exhaustive: i.e.,  $\textit{applyChange}(pc, c)$  is not defined if it is not covered by one of the clauses above.

The above function extends in the obvious way to a function for applying a set of changes, if it is possible to do this in an unambiguous manner:

---

$applyChanges : PCollection \times PChangeSet \mapsto PCollection$

$applyChanges(before, cs) = after \Leftrightarrow$   
 $\exists cseq : seq PChange; pcseq : seq PCollection \bullet$   
 $\#pcseq = \#cseq + 1$   
 $ran cseq = cs$   
 $head pcseq = before \wedge last pcseq = after$   
 $\forall i : 1 .. \#cseq \bullet applyChange(pcseq(i), cseq(i)) = pcseq(i + 1)$

Note that the order in which the changes are applied is not necessarily uniquely determined; however, the result *is* required to be uniquely determined. It is difficult to define the precondition explicitly, but the requirement that the change set be noninterfering is easy to check and takes care of most problems.

### 4.3 The complete paragraph-history data model

The following data type specification models a complete collection of paragraph version trees and their derivation (via a set of changes). For genericity, the paragraph contents are taken to be of a given generic type *Type*. For convenience, we include an auxiliary predicate *okPVersion* for checking whether a given paragraph version is present in the collection.

---

$ParaHistory[Type]$   
 $ptree : Pid \mapsto VForest[Type]$   
 $pchanges : PChangeSet$   
 $okPVersion : \mathbb{P} PVersion$

---

$okPVersion = \cup \{a \in dom ptree \bullet \{v \in ptree(a).nodes \bullet (a, v)\}\}$   
 $\forall pv \in okPVersion \bullet \exists_1 c \in pchanges \bullet isCreatedBy(pv, c)$

The invariant says that the changes are required to be pairwise noninterfering and that each paragraph version has a unique point of creation. Further properties can be added to the invariant to reflect the fact that the paragraph history has been derived in a well-defined manner.

The following V&CM policy will be applied to paragraphs:

- When paragraphs are first created, their initial versions are assigned to root nodes.
- If a change is ever applied to a paragraph, the changed version must be frozen.
- When paragraphs are modified, the new version becomes a child of the old version.

$\forall (a, v) \in \text{okPVersion}; c \in \text{pchanges} \bullet$ $\text{isCreatedBy}((a, v), c) \Rightarrow v \in \text{ptree}(a).\text{roots}$ $\forall (a, v) \in \text{okPVersion}; c \in \text{pchanges} \bullet$ $\text{isChangedBy}((a, v), c) \Rightarrow v \in \text{ptree}(a).\text{frozen}$ $\forall a : PId; u, v : VLabel \bullet \text{modify}(a, u, v) \in \text{pchanges} \Rightarrow$ $\text{okPVersion}(a, u) \wedge \text{okPVersion}(a, v) \wedge \text{ptree.parent}(v) = u$
---

## 5 Documents

### 5.1 Document versions

For the purposes of this paper, each version of an evolving document is modelled as a collection of specific paragraph versions and a *delta set* of the paragraph changes that have been made since the document was last frozen. For convenience, we include an auxiliary variable *pids* representing the set of identifiers of paragraphs in the document.

<i>Document</i>
$\text{pcoll} : PCollection$ $\text{delta} : PChange.Set$ $\text{pids} : \mathbb{P} PId$
$\text{pids} = \text{dom pcoll}$

### 5.2 The complete document-history data model

The complete history of a document consists of a complete paragraph history together with a collection of document versions: see Fig. 5.

<i>DocHistory</i>
$\text{ParaHistory}[Text]$ $\text{docs} : VForest[Document]$
$\forall V_1 \neq V_2 \in \text{docs.nodes} \bullet \text{disjoint}(\text{docs.deref}(V_1).\text{delta}, \text{docs.deref}(V_2).\text{delta})$ $\text{pchanges} = \cup \{ V \in \text{docs.nodes} \bullet \text{docs.deref}(V).\text{delta} \}$ $\forall V \in \text{dom docs.roots} \bullet D.\text{pcoll} = \text{applyChanges}(\emptyset, D.\text{delta})$ $\text{where } D = \text{docs.deref}(V)$ $\forall V \in \text{dom docs.parent} \bullet$ $\text{newdoc.pcoll} = \text{applyChanges}(\text{olddoc.pcoll}, \text{newdoc.delta})$ $\text{where newdoc} = \text{docs.deref}(V), \text{olddoc} = \text{docs.deref}(\text{docs.parent}(V))$

The invariant says:

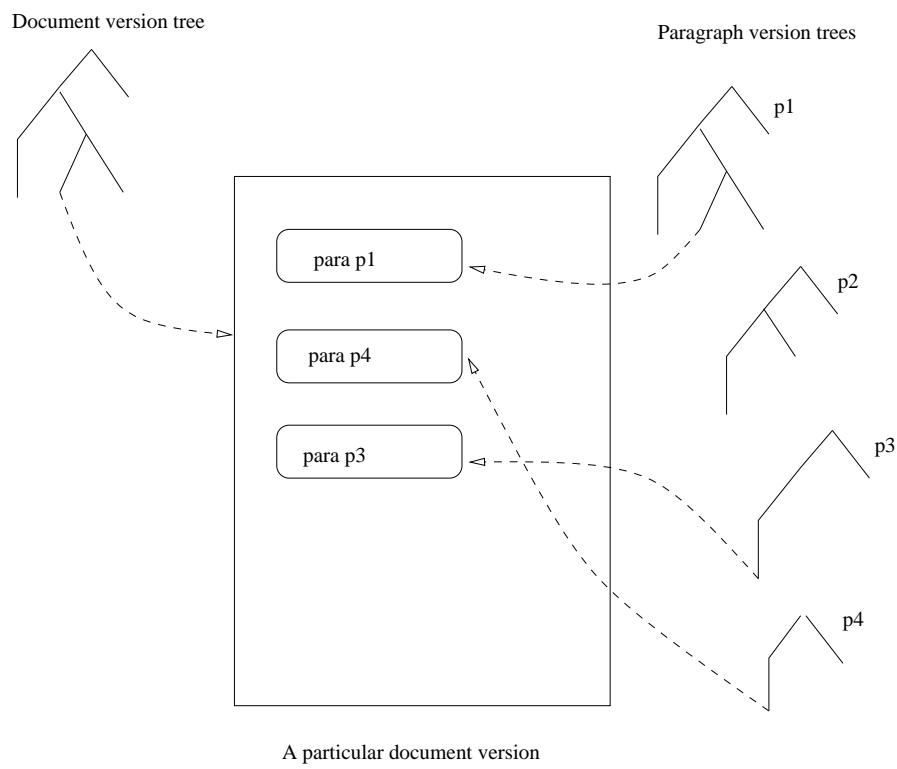


Figure 5: Each version of a document contains a specific set of paragraphs.



- The individual delta sets are all disjoint.
- The paragraph history is in step with the document history, in the sense that it has been derived from the changes recorded in the delta sets of the various document versions.
- For any root version of the document, the delta set records that version's evolution from the null (empty) document.
- For any non-root version of the document, the delta set records that version's evolution from its parent version.

The following V&CM policy will be applied to documents: when a document version is frozen, all its corresponding paragraph versions should also be frozen.

$$\forall V \in docs.frozen \bullet \forall (a, v) \in docs.deref(V).pcoll \bullet v \in ptree(a).frozen$$

## 6 Tracing evolution of paragraphs

This section shows that the above model is rich enough to allow the evolution of individual paragraphs to be traced (backwards and forwards) through a document's history.

### 6.1 Forwards tracing

The following function defines the set of paragraphs which result from a change to a given paragraph:

$$givesRiseTo : PVersion \times PChange \leftrightarrow \mathbb{P} PVersion$$

$$\begin{aligned} \text{dom } givesRiseTo &= isChangedBy \\ givesRiseTo(pv, delete(pv)) &= \emptyset \\ givesRiseTo(pv, split(pv, pvs)) &= \text{ran } pvs \\ i \in 1..n \Rightarrow givesRiseTo(pvs(i), combine(pvs, pv)) &= \{pv\} \\ i \in 1..n \Rightarrow givesRiseTo(pvs(i), replace(pvs, pv)) &= \{pv\} \\ givesRiseTo((a, u), modify(a, u, v)) &= \{(a, v)\} \end{aligned}$$

(Note that the above function does not include paragraphs that are “derived” from the given paragraph, since the latter are not strictly changes to the given paragraph.)

The following function extracts each change (and corresponding document version) which occurs “downstream” in a given paragraph's evolution:

$\text{forwardsTrace} : PVersion \times DocHistory \leftrightarrow \mathbb{P}(PChange \times VLabel)$
$(pv, D) \in \text{dom } \text{forwardsTrace} \Leftrightarrow pv \in D.okPVersion$
$\text{forwardsTrace}(pv, D) =$
$\text{if } \exists V \in D.docs.nodes; c \in D.docs.deref(V).delta \bullet \text{isChangedBy}(pv, c)$
$\text{then } \{(c, V)\} \cup \cup \{npv : \text{givesRiseTo}(pv, c) \bullet \text{forwardsTrace}(npv, D)\}$
$\text{else } \emptyset$

(Note that the fact that this definition is well formed depends on an unrecorded assumption about the “well foundedness” of the set of paragraph changes (*pchanges*). We should really add an appropriate property to the configuration invariant on *ParaHistory* to cover this.)

## 6.2 Backwards tracing

The next function defines the set of paragraphs which a given change “uses”:

$\text{usesPVs} : PChange \leftrightarrow \mathbb{P} PVersion$
$\text{dom } \text{usesPVs} = \text{ran } \text{isCreatedBy}$
$\text{usesPVs}(\text{add}(pv)) = \emptyset$
$\text{usesPVs}(\text{split}(pv, pvs)) = \{pv\}$
$\text{usesPVs}(\text{combine}(pvs, pv)) = \text{ran } pvs$
$\text{usesPVs}(\text{derive}(pvs, pv)) = \text{ran } pvs$
$\text{usesPVs}(\text{replace}(pvs, pv)) = \text{ran } pvs$
$\text{usesPVs}(\text{modify}(a, u, v)) = \{(a, u)\}$

The next function finds the change (and corresponding document version) which created a given paragraph version:

$\text{creationPoint} : PVersion \times DocHistory \leftrightarrow PChange \times VLabel$
$(pv, D) \in \text{dom } \text{creationPoint} \Leftrightarrow pv \in D.okPVersion$
$\text{creationPoint}(pv, D) = (c, V) \Leftrightarrow$
$\text{isCreatedBy}(pv, c) \wedge V \in D.docs.nodes \wedge c \in D.docs.deref(V).delta$

The following function extracts the important “creation” steps along the way to arriving at a given paragraph:

$\text{backwardsTrace} : PVersion \times DocHistory \leftrightarrow \mathbb{P}(PChange \times VLabel)$
$(pv, D) \in \text{dom } \text{backwardsTrace} \Leftrightarrow pv \in D.okPVersion$
$\text{backwardsTrace}(pv, D) = \{(c, V)\} \cup \cup \{opv : \text{usesPVs}(c) \bullet \text{backwardsTrace}(opv, D)\}$
$\text{where}(c, V) = \text{creationPoint}(pv, D)$

(A similar remark about well formedness of the *forwardsTrace* definition applies here.)

## 7 Document conformance

The case study is completed by demonstrating how conformance between pairs of documents can be modelled.

A conformance matrix is modelled as a relation between individual paragraphs in the two documents:

$$CMatrix == PId \leftrightarrow PId$$

The whole configuration is modelled as a pair of document histories, a forest of conformance matrix versions, and a relation which records which versions of the three objects make up “legitimate” configurations:

<p><i>DocPair</i></p> <p><math>A, B : DocHistory</math>  <math>cmatrix : VForest[CMatrix]</math>  <math>corres : \mathbb{P}(VLabel \times VLabel \times VLabel)</math></p> <hr/> <p><math>\forall (V_A, V_B, V_C) \in corres \bullet</math>  <math>V_A \in A.docs.nodes \wedge V_B \in B.docs.nodes \wedge V_C \in cmatrix.nodes</math>  <math>\forall (a, b) \in cmatrix.deref(V_C) \bullet</math>  <math>a \in A.docs.deref(V_A).pids \wedge b \in B.docs.deref(V_B).pids</math></p>
--

The *corres* relation records which versions of the three objects make up recognised configurations. The invariant says that

- there are no dangling references in the *corres* relation
- there are no dangling references to paragraph identifiers in the conformance matrices.

The V&CM policy from Section 2 will be strengthened as follows:

- The conformance matrix and B-document should be managed as a single configuration: i.e., there is a one-one correspondence between versions of the two objects.
- Moreover, each version of said configuration should refer to a single version of the A-document.
- If B is frozen then the corresponding A should also be frozen and the conformance matrix should be frozen and complete (i.e., every A-paragraph should be addressed by at least one B-paragraph).

(Note that in multidocument situations there may be a possibility of deadlock here.)

$$\begin{aligned}
& \forall (V_A, V_B, V_C), (V'_A, V'_B, V'_C) \in \text{corres} \bullet V_B = V'_B \Leftrightarrow V_C = V'_C \\
& \text{corres}^\sim \in \text{VLabel} \times \text{VLabel} \leftrightarrow \text{VLabel} \\
& \forall (V_A, V_B, V_C) \in \text{corres} \bullet V_B \in B.\text{docs.frozen} \Rightarrow \\
& \quad V_A \in A.\text{docs.frozen} \wedge V_C \in \text{cmatrix.frozen} \wedge \\
& \quad \text{dom cmatrix.deref}(V_C) = A.\text{docs.deref}(V_A).\text{pids}
\end{aligned}$$

This policy could usefully be strengthened, say in COMPUSEC applications, to say that every B-paragraph should address at least one A-paragraph (e.g. to ensure that no unauthorised functionality has been added).

## 8 Conclusions

### 8.1 Summary

This paper described a case study in fine-grained version and configuration management from the perspective of providing integrated functionality to support auditing and traceability. We demonstrated that fine-grained versioning of versioned objects provides substantially more flexibility than traditional approaches that manage only high-level coarse-grained objects.

There are two dimensions to the benefits accrued by the use of fine-grained versioning models in our case study. The first is due to the fact that we consider the individual components of our systems as first class citizens in the context of configurations. This allows substantial flexibility in the way in which consistency of an overall system is determined, as well as focusing attention on the specific objects undergoing change.

The second benefit comes from the actual versioning of the system components themselves. It allows us to define consistency criteria in terms of the conditions that must be satisfied by the individual components. We can then show that the chosen versioning model actually meets the criteria.

### 8.2 Further Work

During our work a number of other issues arose, particularly with respect to possible extensions to our framework which would allow developers to reduce the impact of change (as opposed to accurately assessing the actual work required to react to a change).

It's clear that, by supporting more sophisticated structuring mechanisms the impact of change could be localised better. We intend to test our hypothesis that the versioning model outlined here will scale to more complex structures, to provide a basis for the flexible construction and maintenance facilities suitable for large-scale development.

The model we have presented is a core model. In addition to this, one would typically define high (user)-level processes based on these models. Whereas consistency constraints for objects within a system are defined by the model, the process which evolves and uses these underlying concepts need not be fixed. Our framework can be used as a basis upon which more sophisticated process models can be developed. Such models can offer context-sensitive guidance to the users of systems, as well as the opportunity to further constrain the way in which a system is used. It is possible to define and reason about intermediate states of configuration consistency, and to offer guidance on how to bring the system back into a consistent state, for example. We have illustrated these ideas on theory management [4, 7].

Finally, we have an ongoing effort in prototyping tools to support our fine-grained configuration management framework using object-oriented database technology [1].

### 8.3 Acknowledgements

The authors gratefully acknowledge the useful contributions of Yaowei Liu and Sabine Sachweh to the work presented here.

## References

- [1] F. Bancilhon, C. Delobel, and G. Harrus. *Building an Object-Oriented Database System: the Story of O2*. Morgan Kaufman, 1992.
- [2] S.C. Choi and W. Scacchi. Assuring the correctness of configured software descriptions. *ACM SIGSOFT Software Engineering Notes*, 14:66–75, 1989. Proc 2nd Intl Workshop on Software Configuration Management.
- [3] IEC. Functional Safety: Safety-Related Systems. Draft International Standard IEC 1508, June 1995.
- [4] P.A. Lindsay. A formal approach to specification and verification of task management in interactive systems. *IEE Proceedings of Software Eng*, 144(4):206–214, August 1997. (Formerly Sw Eng Journal.) Also appears as SVRC TR 97-23.
- [5] P.A. Lindsay, Y. Liu, and O. Traynor. A generic model for fine-grained configuration management including version control and traceability. In *Proc. Australian Software Engineering Conference (ASWEC'97)*, pages 27–36. IEEE Computer Society Press, 1997. SVRC TR 97-45.
- [6] P.A. Lindsay, Y. Liu, and O. Traynor. Managing document conformance: a case study in fine-grained configuration management. *Aust Comp Sci Communications*, 19(1):373–382, 1997. Also appears as SVRC TR 96-20.

- [7] P.A. Lindsay and O. Traynor. Version and configuration management of formal theories. In *Proc. Formal Methods Pacific (FMP'97)*, pages 165–185. Springer Verlag, 1997. Also appears as SVRC TR 97-13.
- [8] U.K. Ministry of Defence. The Procurement of Safety Critical Software in Defence Equipment. Defence Standard 00-55, August 1995. <http://www.dstan.mod.uk/> or <http://www.seasys.demon.co.uk/>.
- [9] U.K. Ministry of Defence. Safety Management Requirements for Defence Systems Containing Programmable Electronics. Second Draft Defence Standard 00-56, August 1996. <http://www.dstan.mod.uk/> or <http://www.seasys.demon.co.uk/>.
- [10] F.A.C. Pinheiro and J.G. Goguen. An object-oriented tool for tracing requirements. *IEEE Software*, pages 52–64, 1996.
- [11] M. J. Rochkind. The source code control system. *IEEE Transactions on Software Engineering*, 1:24–36, 1975.
- [12] K.J. Ross. Models for configuration management of refinement calculus developments. In *Proceedings 7th Refinement Workshop*, pages 1–26. BCS-FACS, July 1996. also appears as SVRC TR 95-8.
- [13] K.J. Ross and P.A. Lindsay. Maintaining consistency under changes to formal specifications. In *Proc. 1st Int. Symp. of Formal Methods Europe (FME'93)*, LNCS 670, pages 558–577. Springer Verlag, 1993. Also appears as SVRC TR 93-3.
- [14] K.J. Ross and P.A. Lindsay. Applying software configuration management techniques to formal development. Technical Report 95-12, Software Verification Research Centre, The University of Queensland, 1995.
- [15] S. Sachweh and W. Schäfer. Version management for tightly integrated software engineering environments. In *Proc. 7th Int. Conf. on Software Eng Environments*, pages 21–31, The Netherlands, 1995. IEEE Computer Society Press.
- [16] W. Tichy. RCS - a system for version control. *Software - Practice and Experience*, 15, 1985.