Extended Requirements Traceability: Results of an Industrial Case Study

Orlena Gotel & Anthony Finkelstein

Department of Computer Science
City University
Northampton Square
London EC1V 0HB
[olly | acwf]@soi.city.ac.uk

Abstract

Contribution structures offer a way to model the network of people who have participated in the requirements engineering process. They further provide the opportunity to extend conventional forms of artifact-based requirements traceability with the traceability of contributing personnel. In this paper, we describe a case study that investigated the modelling and use of contribution structures in an industrial project. In particular, we demonstrate how they made it possible to answer previously unanswerable questions about the human source(s) of requirements. In so doing, we argue that this information addresses important problems currently attributed to inadequate requirements traceability.

1. Introduction

The inability to answer questions regarding the human source(s) of requirements information has been found to result in claims of *requirements traceability* problems [5]. An approach to address this problem was presented in [6]. This paper describes a case study designed to evaluate, through demonstration, whether use of the approach helps answer practitioner's outstanding questions and, in so doing, alleviates an important class of requirements traceability problems. The case study is based on material gathered from a real industrial project over a period of three years.

In Section 2, we explain what requirements traceability is and describe the underlying reason for long-term requirements traceability problems. We provide examples of the kind of questions that are problematic to answer as a consequence. We then outline an approach to address this fundamental problem and summarise how it is anticipated to provide answers that satisfy practitioners' needs. In Section 3, we describe the case study material used to validate our claim. Since the approach was not in existence at the onset of the project, its requirements did not drive the data gathering. From the extensive records that were maintained, we only summarise that data pertinent to the approach. In Section 4, we demonstrate how the approach was applied in a post-hoc manner to this data, thereby revealing information about the project's evolution. In Section 5, we show how this information makes it possible to answer questions about the project regarding *involvement*, *responsibility*, *ramifications*, *change notification* and *working relationships*. Due to the manner in which the approach was applied, we are only in a position to validate the feasibility of the approach and the usefulness of the information it provides in a historical and subjective manner. Another repercussion of not being on the project's critical path is that we can only suggest how the information could be used, say to assist the maintenance process. Based on our experiences and practitioners' comments, we highlight some outstanding issues in Section 6, then make recommendations for uptake.

2. Contribution structures for requirements traceability

In this section, we describe what requirements traceability is, why it is important and what the problems with it are. We then outline an approach to address a fundamental problem that currently makes it difficult to recover information about the human source(s) of requirements information.

2.1. Requirements traceability

Requirements traceability refers to the ability to describe and follow the life of a requirement in both a forwards and backwards direction (i.e., from its origins, through its development and specification, to its subsequent deployment and use, and through periods of on-going refinement and iteration in any of these phases). It is considered the primary technique to help with many project-related activities, like ensuring that systems and

software conform to their changing requirements, but is commonly cited as a problem area by practitioners.

Although the number of tools that claim to support requirements traceability is growing, some more recent ones being described in [10, 11, 14, 15], the schemes that need to be established prior to their use have received rather less attention. With few exceptions, examples being the requirements traceability models of the U.S. DoD [8, 9] and the requirements traceability meta models arising from the NATURE project [12], endeavours to improve the potential for requirements traceability have mostly involved uncovering and recording as much information as possible about the requirements engineering process, then linking it in interesting ways for trace retrieval. This can lead to an unwieldy mass of unstructured and unusable data without some *a priori* discrimination concerning the type of requirements information that practitioners are likely to need access to.

Following an empirical study reported in [5], we argued that the most fundamental information to record for relieving *long-term* requirements traceability problems was that which identified the *human source(s)* of requirements information. We found that, what are perceived to be requirements traceability problems tend to arise when practitioners are unable to answer questions about the personnel who had been involved in the production and refinement of requirements. This is because people are frequently the ultimate baseline whenever requirements need to be re-examined or re-worked. Examples of such questions are given in Figure 1. However, we also reported that information concerning the human source(s) is often neglected in the strive to replace informal social contact with exhaustive documentary records.

"Who has been involved in the production of this requirement and how?"

"Who was originally responsible for this requirement, who is currently responsible for it and at what points in its life has this responsibility changed hands?"

"At what points in this requirement's life have the working arrangements of all those involved been changed?"

"Who needs to be involved in or informed of changes proposed to this requirement?"

"What are the ramifications, regarding the loss of requirements-related knowledge, if a specific individual or group leaves a project?"

Figure 1: *Personnel-related requirements traceability questions.*

2.2. Contribution structures

We described an approach to address this more focal problem underlying long-term requirements traceability problems in [6]. In that paper, we also described a prototype tool to support the approach and provided scenarios of use. Formalisation of the approach, plus the inferences it supports, can be found in [4].

We summarise the main steps of the approach in Figure 2. The approach is based on modelling the *contribution structure* underlying requirements. This reflects the network of people who have contributed to the artifacts produced in the requirements engineering process. The information in this model makes it possible to extend conventional forms of *artifact-based* requirements traceability with a form of *personnel-based* requirements traceability. We claim that this new dimension can help elicit answers to the above questions and so alleviate an important class of requirements traceability problem.

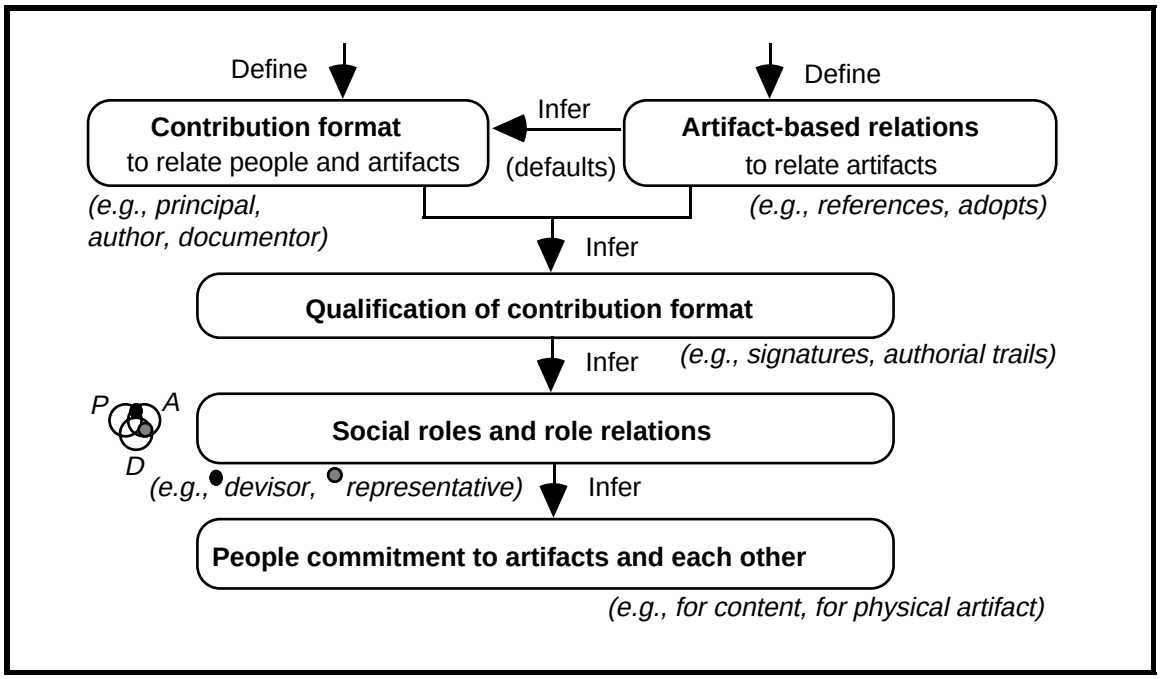

Figure 2: *The main steps of the approach.*

Working through Figure 2, minimal semantics are given to the artifact-based relations ordinarily put in place for requirements traceability purposes. Based on the notion of *communicative function*, an artifact can be said to *reference* or *adopt* the content of a linked artifact, the distinction being whether or not the content in the source and target overlaps. A record of the people who

contributed to an artifact's production is also maintained in its *contribution format*. Based on Goffman's work on the nature of participation in social encounters [3], this structure delineates the *principal*, *author* and *documentor* of an artifact. As described in [6], these categories have been motivated by sociolinguistic and textual linguistic theories, and chosen for their analytic potential. Together, the information they provide can be used to infer details about social roles, role relations and commitments.

3. Case study

In this section, we give details of the company, project and participants of the case study. We describe the data we gathered and our method for so doing.

3.1. Project

The project came from a small communications company employing about twenty-five people. It runs many projects concurrently, providing software and procedural solutions to communications-related problems. The original objective of this project was to supply a dedicated communication service to complement a customer's disaster recovery programme. The project was initiated in February 1992 and went live at the end of March 1992.

In August 1992, the idea of developing a generic service was discussed. Six versions of a requirements and design specification were drawn up throughout September 1992. These were then abandoned until new staff were employed at the end of October 1992 to develop and market the service. Following much staff turn-over, the generic service did not go live until February 1994. Between October 1992 and February 1994, the specification evolved into an operational service, an operations manual and a high-level manager's guide. Since February 1994, the generic service and accompanying documentation has undergone continuous modification to account for the requirements of new customers.

Most of the artifacts produced during the project were informal and paper-based. All that remains within the company today is an early specification, an up-to-date operations manual, an up-to-date manager's guide, contracts with customers and miscellaneous correspondence. Requirements traceability has not been maintained. Those still involved in the project no longer remember from where or from whom the various aspects of the current service have been derived. Some problems have resulted from this loss of information but,

because the project is restricted in scope, and because the team is small and exhibits some staff continuity, these have not been critical to maintainability.

3.2. Data gathered

The work that occurred from the initial discussion about providing a generic service, through to the sixth version of the requirements and design specification, was followed closely. We observed all the meetings that took place, made notes and audio recordings, and collected photocopies of any tangible artifacts produced. We also participated in some aspects of the process. During this time, a detailed picture of what had happened when developing the initial customer-specific service was reconstructed with those who had been involved.

From the end of October 1992, we maintained a record of the main artifacts produced due to the specification, many of which were in production for months. For the purposes of the case study, our definition of “artifact” applied to single physical documents. Not only does this promote identification, but it enables us to examine the viability of the approach at a coarse level of granularity prior to introducing further complexity. Some of the traceability implications that arise from decomposing a single physical artifact into a number of components are mentioned in Section 5. During this period, we also maintained a record of the people involved in the production and distribution of these artifacts. However, we were unable to maintain a full record of the peripheral artifacts, like the notes made during meetings with customers and so forth.

Fifty-eight people contributed directly to the project. These included individuals and groups from within the company and from outside. We use alphabetic identifiers when we refer to these people below. One hundred and sixty-six main artifacts were produced in the project. We use numeric identifiers when we refer to these artifacts below.

3.3. Project phases

(1) The artifacts to do with the initial customer-specific service. Twenty-three artifacts were produced between February and March 1992.

(2) The artifacts to do with developing the baseline specification for the generic service. Sixty-five artifacts were produced from August to September 1992.

(3) The artifacts to do with developing the initial generic service. Thirty-nine artifacts were produced from October 1992 to July 1993.

(4) The artifacts to do with extending the generic service to address new customer requirements. Thirty-nine artifacts were produced from September 1993 to June 1995.

4. Application of approach

In this section, we outline how the approach was applied. Based on the data we had gathered, key project participants were tasked to reconstruct the main artifact-based relations and to give them minimal semantics. Under similar conditions, they were also tasked to reconstruct the contribution format for each artifact. We then examined what could be inferred about social roles, role relations and commitments as a consequence.

4.1. Artifact-based relations

For each project phase, its artifacts were numbered according to production order, then temporal relations were established between them to reflect the order in which they had been produced. This ordering is suggested in Figure 3. The nature of these relations was further clarified, based on [1], as: `before(1,2)`; `during(4,3)`; `equal(11,12)`; `meets(16,17)`; etc. The coarse flow-down of information amongst these artifacts was also established. This is also shown in Figure 3. Note that, no attempt was made to determine all the possible relations, like the many transitive or internal relations for instance.

An example of the minimal semantics given to these artifact-based relations is shown in Table 1. In the first column of this table, we list the original reason given for the relation by participants. In the second column, we categorise the nature of this relation according to standard classifications of cohesion and coherence relations [2]. Based on this, we list the broad communicative function between the two related artifacts in the third column. This function encapsulates the finer purpose of the relation, helps to eliminate individual differences in classification, so provides what we refer to as its *minimal semantics*.

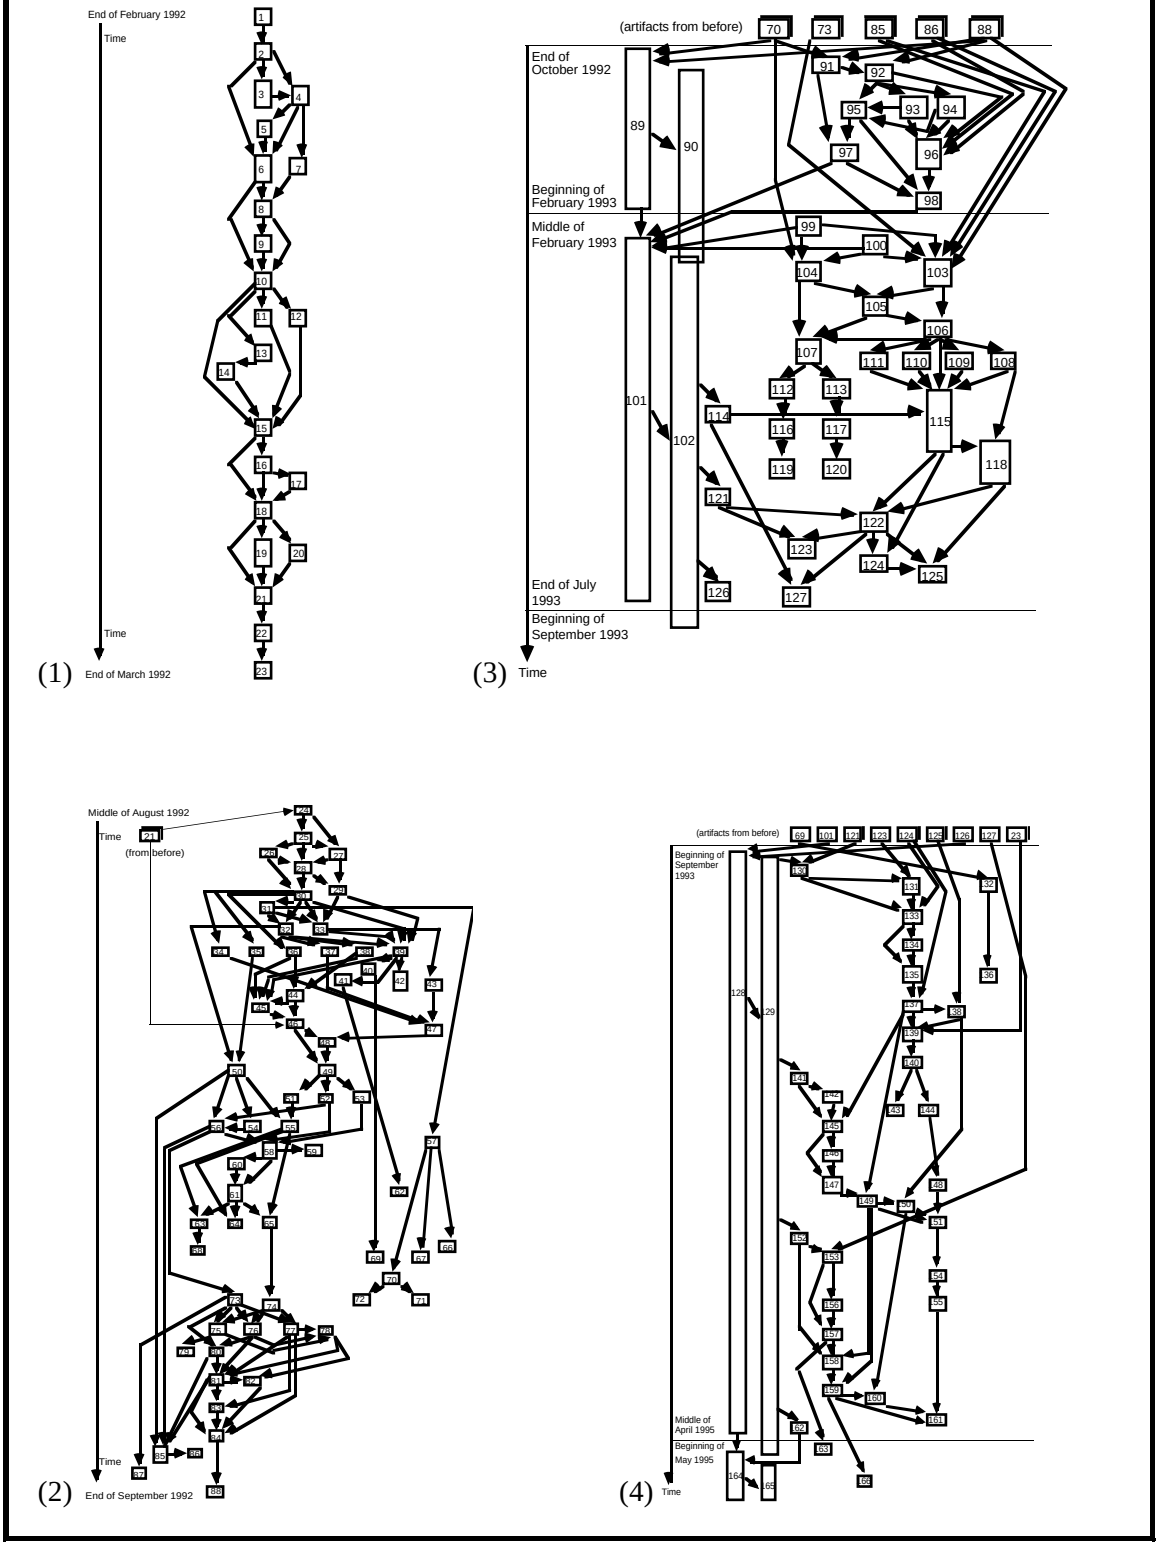

Figure 3: *The order in which artifacts were produced and the predominant artifact-based relations between them. Note that, elongated artifacts, like 102 and 128, represent all the correspondence between the company and potential customers; correspondence was difficult to keep track of. Note that, the slightly enlarged artifacts, like 115 and 147, indicate that other artifacts were produced whilst in production.*

8

(1) Informal description of relation given by practitioner	(2) Nature of relation in terms of cohesion and coherence relations	(3) Broad communicative function of relation gives minimal semantics
2 qualifies 1	2 adds to 1	2 adopts 1
2 is the reason for 3	2 frames 3	3 references 2
4 defines 2	4 adds to 2	4 adopts 2
2 is the reason for 6	2 frames 6	6 references 2
3 assists with 4	3 substantiates 4	4 references 3
5 is compared with 4	5 matched with 4	5 references 4
6 refines 4	6 alters 4	6 adopts 4
5 assists with 6	5 substantiates 6	6 references 5
7 responds to 4	4 causes 7	4 references 7
6 is the reason for 8	6 frames 8	8 references 6
7 is background for 8	7 frames 8	8 references 7
9 is a result of 8	8 causes 9	9 references 8
9 assists with 10	9 substantiates 10	10 references 9
10 elaborates 6	10 adds to 6	10 adopts 6
8 is background for 10	8 frames 10	10 references 8
10 is the reason for 11	10 frames 11	11 references 10
12 replies to 10	10 causes 12	12 references 10
13 replies to 10	10 causes 13	13 references 10
15 extends 10	15 adds to 10	15 adopts 10

Table 1: *The nature of some of the main artifact-based relations of phase 1.*

For phase one, Figure 4 (a) illustrates those relations likely to be represented in conventional forms of artifact-based requirements traceability. This is because they tend to capture *parent-child* relations, or *predecessor-successor* relations, so provide requirements history and flow-down [13, 16]. Figure 4 (b) illustrates those additional relations the approach is concerned with, namely the references relations of Table 1. It indicates the wealth of contextual information often not integrated and used actively for requirements traceability purposes.

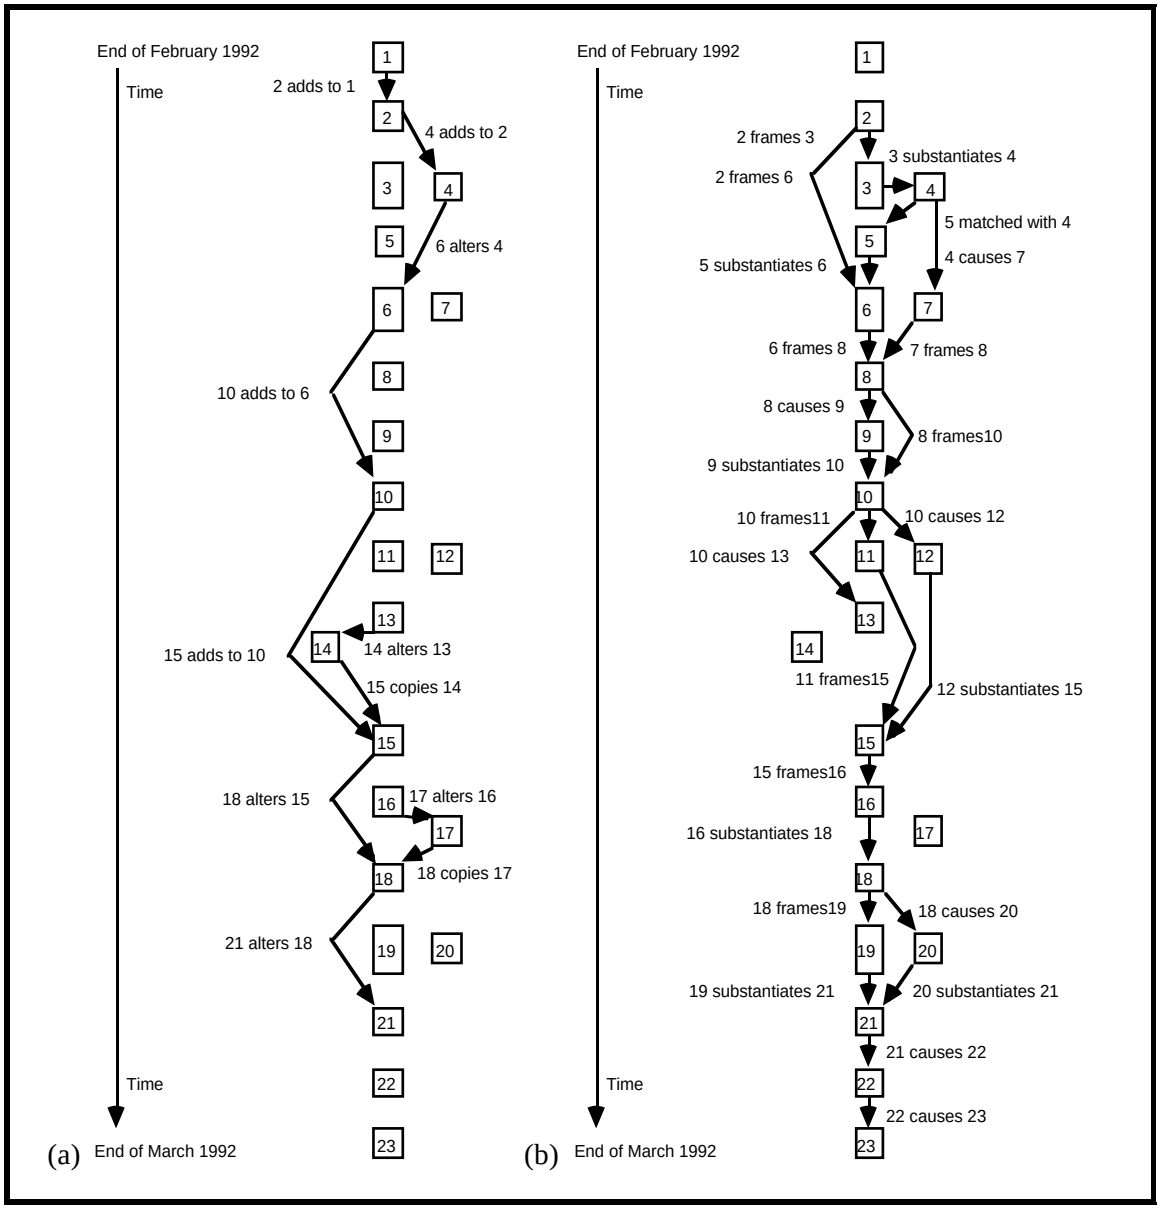

Figure 4: *(a) Adopts - the artifact-based relations typically maintained for requirements traceability. The arrows here suggest the flow-down of artifact content. (b) References - the artifact-based relations that further provide context. The arrows here suggest the direction of influence between artifacts.*

4.2. Contribution format

For each project phase, the contribution format was established for each artifact. Using our example scheme, this indicated the individuals or groups who contributed in the capacities of principal, author and documentor. The contribution formats allocated for some of the artifacts produced in phase one are shown in Table 2. The people to whom an artifact was either passed or copied is also shown in this table for completeness.

Artifact	Principal	Author	Documentor	Distributed to
1	BH	BI	AW	AT
2	AT	BB={AW/AV/AT /AR/AX/AU}	BB={AW/AV/AT/AR/ AX/AU}	AA
3	AA	AA/AE	AA	None
4	AA	AA/AE	AU	BI/BB={AW/AV/AT/ AR/AX/AU}
5	AA	AA/AQ/AP/BB={AW /AV/AT/AR/AX/AU}	AA	None
6	AA	AA/AT	AA	AE
7	BH	BI	BL	AA
8	AA	AA/AE	AA	None
9	AA	AA/BB={AW/AV/AT/ AR/AX/AU}	AA	None
10	AA	AA	AA	AE/BI/BB={AW/AV/ AT/AR/AX/AU}

Table 2: *The contribution formats for artifacts 1 to 10. Note that, AA/AE means joint contributors. Note that, group descriptors are decomposed into their members.*

4.3. Qualification

Each contribution format was qualified to provide further details about contributions and contributors. As an example, we explain how the authorial status was qualified and highlight the use of this information. Artifact 106 was the first version of the operations manual for the generic service. It was authored by AI. From artifact 106, the artifact-based relations can be used to trace all the paths back to the original author(s). They can also be used to trace forwards to locate the author(s) of all the artifacts arising from it. Details about how each progressive author made use of the previous author's contribution can thereby be uncovered. Part of this authorial trail is illustrated in Figure 5.

From such authorial trails, we can see: which people produce the most original artifacts; which use their own or another's contributions the most often; whose contributions get referenced with the greatest frequency; and so on. We can also begin to assess the influence of a person's authored contributions on the surrounding body of artifacts and on the project as a whole. Such details can help identify those to notify following different types of change or those to contact regarding different queries.

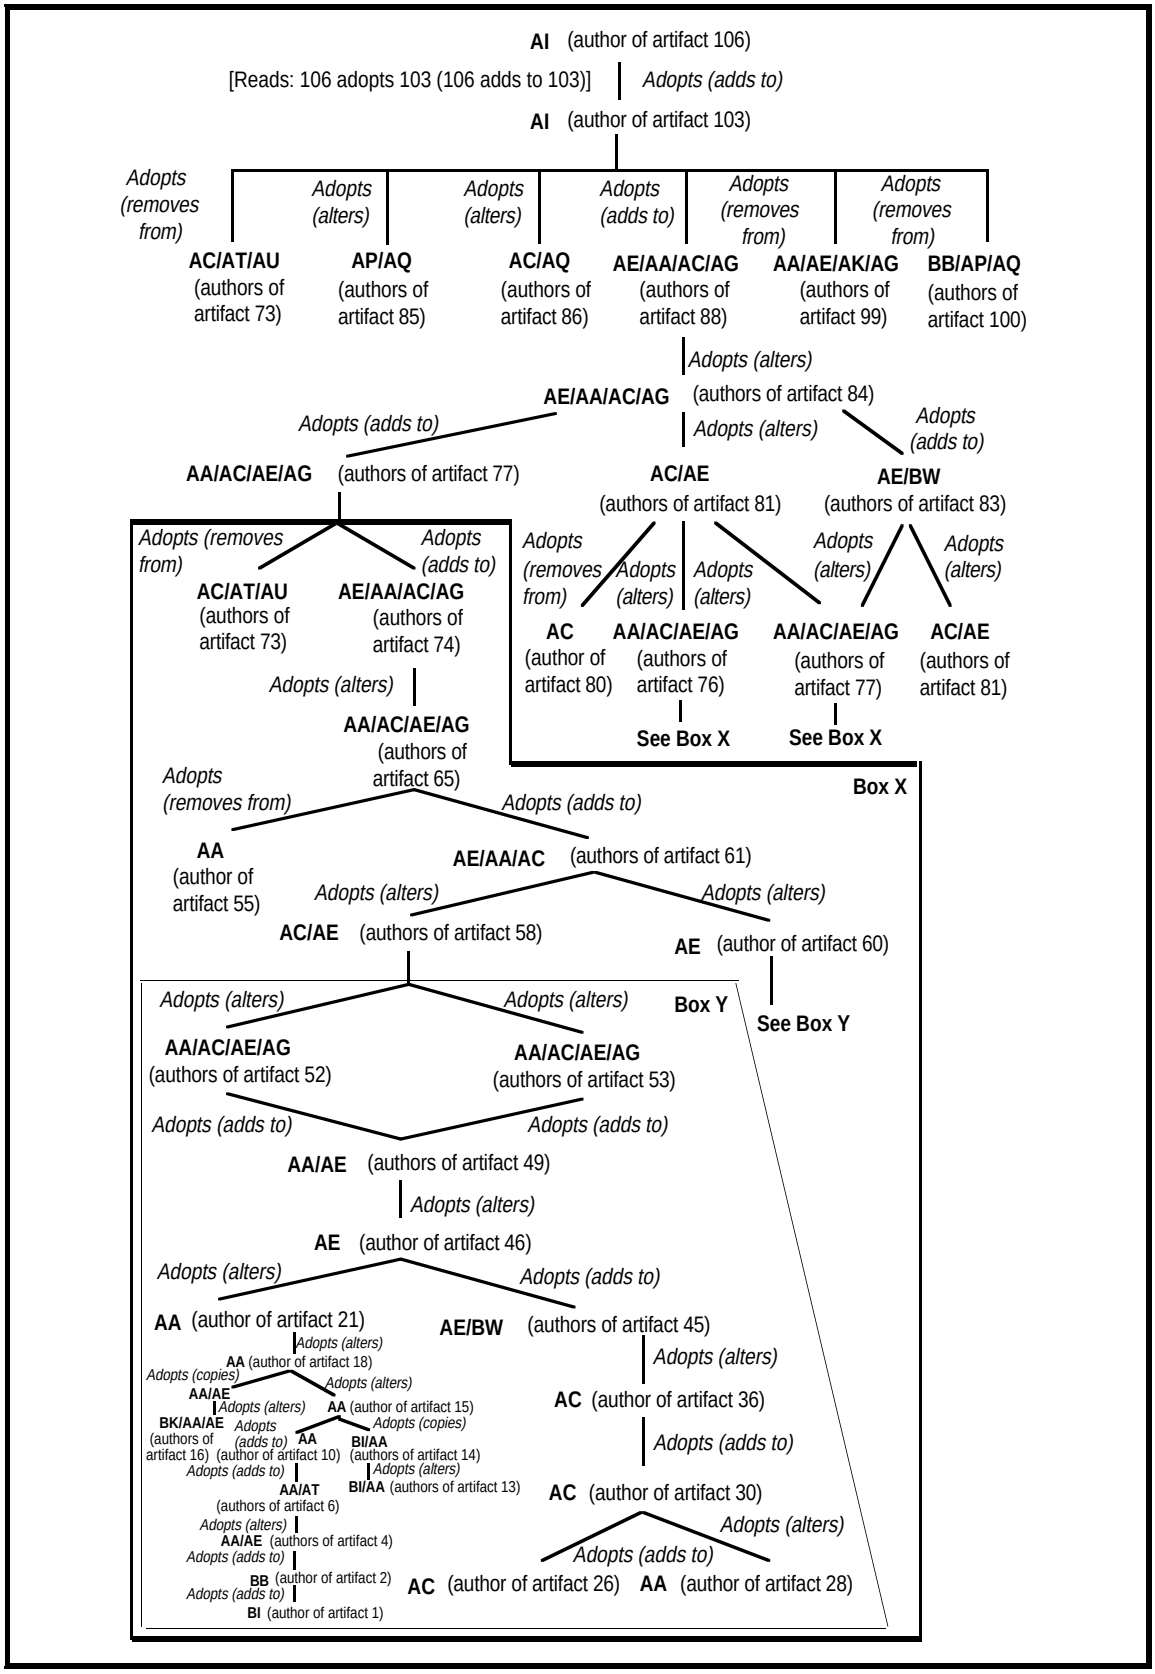

Figure 5: *Authorial trail following backwards requirements traceability from artifact 106. For clarity, it is filtered to highlight those whose contributions were adopted in getting to106.*

4.4. Social roles and role relations

The social roles that people assume when contributing to artifacts can be inferred from the information we have gathered so far. For instance, if a person is both the principal and author of an artifact, they can be said to be its *devisor*. If they are solely the documentor, they can be said to be its *relayer*. The ensuing role relations between people when they jointly contribute to artifacts, say as a *devisor/relayer* pair, reveals more about the underlying contribution structure. Not only can we see whom has collaborated with whom, but we can see how they have collaborated and whether these role relations have varied or been sustained.

To explain the use of such information, we compare the social roles of two of the project leaders. AI was the project leader when artifacts 99 to 127 were produced and a contributor to twenty-two of these. AJ was the project leader when artifacts 128 to 162 were produced and a contributor to twenty-six of these. Their social roles when contributing to these artifacts, as well as their role relations to collaborators, are shown in Tables 3 and 4 respectively.

Social role of AI	On how many artifacts?	How many on own?	Social roles of other contributors who collaborate with AI
True author (i.e., contributes as P, A and D)	16	13	True author = BM (x1) Ghost author = AA/AE/AD (x1) Ghost author = AA/AE/AD/AT (x1)
Nominal author (i.e., PD)	2	0	Ghost author = AA/AE/AD/AG (x1) Ghost author = BB/AP/AQ (x1)
Representative (i.e., AD)	2	0	Sponsor = AD (x2)
Ghost author (i.e., A)	2	0	Sponsor = AD and Relayer = BO (x2)

Table 3: *The social roles and role relations for AI.*

From these tables, we can see that AI worked on his own on over half of the artifacts he contributed to, else he worked with small groups of people. As he worked largely as a *true author*, he was evidently a self-sufficient documentor. It is noteworthy that AD tended to collaborate with AI as a *sponsor* when dealing with customer-related artifacts. In contrast, we can see that AJ worked rarely on his own and mainly collaborated with one or two others. He had a strong dependency on AE as his *ghost author* when working together and on many others as *relayers*, the latter hinting at the need for secretarial support. It is noteworthy that AD was ultimately responsible for about a third of the artifacts that AJ had contributed to.

Social role of AJ	On how many artifacts?	How many on own?	Social roles of other contributors who collaborate with AJ
True author (i.e., PAD)	7	2	True author = BM (x1) Ghost author = AE (x 4)
Ghost author (i.e., A)	9	0	Sponsor = AD and Relayer = AP/AS (x1) Sponsor = AD and Relayer = AP (x5) Sponsor = AD and Relayer = AR (x2) Sponsor = AD, Ghost author = AF and Relayer = AO (x1)
Devisor (i.e., PA)	9	0	Ghost author = AE and Relayer = AW (x2) Ghost author = AE and Relayer = BQ (x1) Relayer = AL (x4) Relayer = AM (x2)
Sponsor (i.e., P)	1	0	True author = AE (x1)

Table 4: *The social roles and role relations for AJ.*

There could be many reasons for the subtle differences in how these two people with the same job description worked in the project. AI did not close any sales and focused on developing a marketable service. In contrast, AJ focused on selling what AI had developed and only made subsequent additions to it to account for new customer requirements. Notably, it was with such additions that AE collaborated with AJ as ghost author. This served to maintain some continuity, since AE had also collaborated with AI as ghost author.

4.5. Commitment

Based on the data given in Table 5, we give an example of the kind of information that can be inferred about the commitments of project contributors. We can see that AP is mainly responsible for the physical appearance of artifacts, only responsible for their content when collaborating with others, though never responsible for their ultimate effect. We can also see which other people AP is committed to through their joint contributions to shared artifacts. Here, we can identify AD and AJ as those with whom AP has collaborated the most often, as well as the number and type of artifacts on which they collaborated. By extension, we can examine those people that AP is committed to due to the artifact-based relations that situate her contributions, though not shown in the table.

The intersection and difference between commitments can uncover much interesting information. For example, we can identify: which people have collaborated with specified others the most or least often; which people are committed to the same set of other people; which people have collaborated

with customers; which people are committed to the same artifacts and for the same aspects; and which people have contributed to those artifacts that are the initial sources of requirements.

Contrib to (artifact)	Aspect of artifact committed to	Contrib with (person)	No. artifacts collab on
5	Content (as one of many contributors)	AD	6
22	Physical appearance (on own)	AJ	6
31	Physical appearance (on own)	AQ	4
41	Content (as one of two contributors) Physical appearance (on own)	AT	4
85	Content (as one of two contributors)	AU	4
96	Content (as one of many contributors) Physical appearance (as one of two contributors)	AW	3
100	Content (as one of many contributors)	AV	3
111	Content (as one of many contributors)	AR	3
139	Physical appearance (as one of two contributors)	AX	3
140	Physical appearance (on own)	AC	2
148	Physical appearance (on own)	AD	2
151	Physical appearance (on own)	AJ	2
154	Physical appearance (on own)	AQ	1
155	Physical appearance (on own)	AT	1

Table 5: *AP's artifact and collaborator commitment store.*

5. Results and discussion

In this section, we select some of the questions given in Figure 1 to demonstrate how they can be addressed. We also mention other forms of analysis the approach makes possible. The reader is referred to [4] for a more detailed description and a thorough evaluation.

5.1. Involvement

"Who has been involved in the production of this requirement and how?"

One of the requirements in version two of the requirements and design specification, artifact 49, was to do with security. It was pursued throughout phase two of the project, cited in all six versions of the specification, then dropped in phase three. It led to much investigation and many artifacts that became redundant. Once removed, its impact only surfaced over time. The resulting problems could have been alleviated with knowledge of its original source and of those who had pushed for its concern.

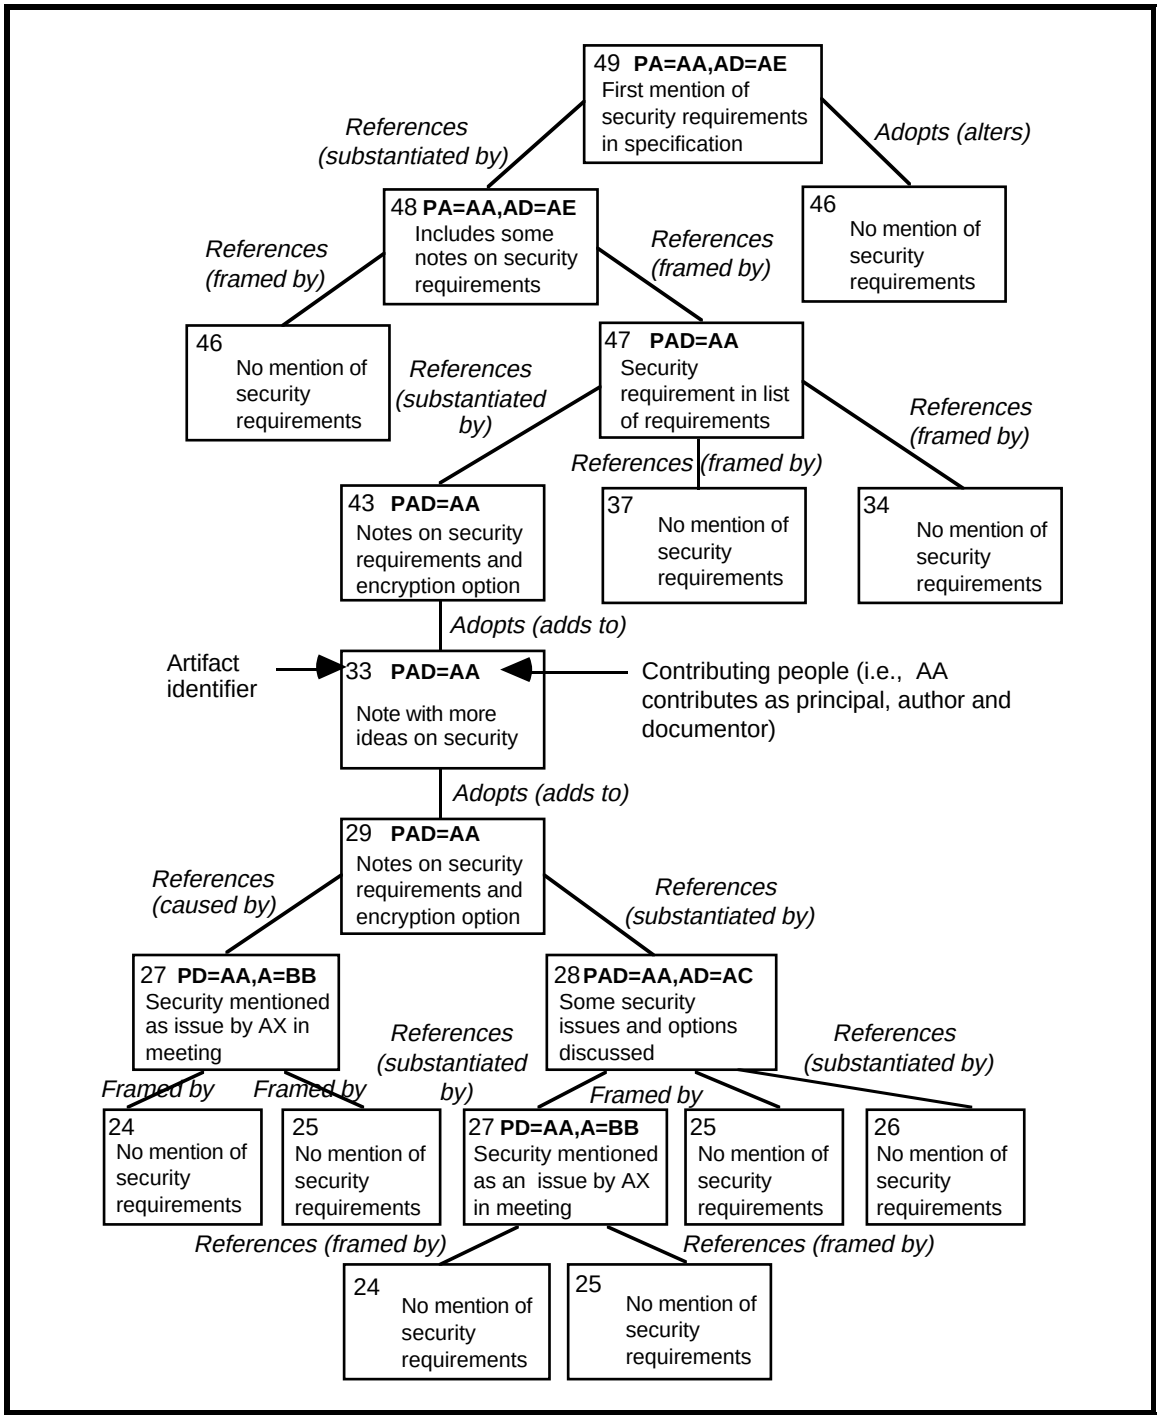

Figure 6: *Who was involved in the production of the security requirement baselined in artifact 49. Note that, only coarse-grain relations have been put in place to link the applicable internal components of artifacts in this figure. Note that, where groups of people have contributed, like BB in artifact 27, the group members contributing to internal components have only been delineated where needed for clarity.*

From Figure 6, we can see that this requirement can be traced back to artifact 27. The contribution format at the source shows that AA was writing requirements in the name of the BB collective. Where contribution formats are provided for the internal components of artifacts, AX can be identified as the specific originator. This makes it possible to recover AX's original

intention, one that was actually misconstrued by AA. Furthermore, we can see how this misconstrued requirement pervades subsequent artifacts due to AA's backing and no subsequent recourse to AX. We can also see which people ended up doing the most redundant work as a consequence. Notably, it was a requirement that dominated many of AA's early and individual contributions. Only AE and AC also contributed to this requirement's evolution before it became baselined in the specification, notably when collaborating with AA.

5.2. Responsibility

"Who was originally responsible for this requirement, who is currently responsible for it and at what points in its life has this responsibility changed hands?"

Phase three of the project saw the introduction of a manager's guide, its latest version being artifact 160 in phase four. Table 6 shows a subset of the information gathered relating to this artifact.

Manager's guide	**Artifact 160**	**Artifact 150**	**Artifact 138**	**Artifact 125**	**Artifact 118**
Version	5	4	3	2	1
Principal	AJ	AJ	AJ	AI	AI
Author	AJ	AJ	AJ	AI	AI
Documentor	AM	AL	AL	AI	AI
Adopts relations	Adds to 150	Adds to 138	Adds to 125	Alters 118	None
References relations	Matched with 159	Matched with 149	Matched with 137	Matched with 124	(a)Matched with 115 (b)Framed by 108
P (of referenced artifact)	AJ	AJ	AJ	AI	(a)AI (b)AA
A (of referenced artifact)	AJ	AJ	AJ	AI	(a)AI (b)AA
D (of referenced artifact)	AM	AL	AL	AI	(a)AI (b)AA

Table 6: *Changes in responsibility for the manager's guide.*

From Table 6, we can see the transition between AI's original work and AJ's later work. We can also see that AJ has only made additions to what AI originally produced. The working arrangements have also changed from AI working on his own to AJ working in conjunction with one other person doing the documentation. Therefore, although AI was originally responsible for all aspects of the guide, AJ is now responsible for its content and effect, whilst AM is responsible for the physical document. Note that we define an "original contribution" to be one that does not depend upon other artifacts for its existence and that we do not attempt to measure degrees of originality.

Table 6 further shows that the guide has been aligned with versions of the operations manual throughout its evolution, these being artifacts 159, 149, 137, 124 and 115. The only other artifact with which the first version is related is artifact 108. Inspection of this artifact can reveal that AA, as its author, originally suggested its need.

5.3. Working arrangement and remit

"At what points in this requirement's life have the working arrangements of all those involved been changed? Accordingly, within the remit of which groups do decisions about this requirement lie?"

In Figure 7, we depict the contributors to the formal versions of the requirements and design specification produced in phase two of the project. From this, we can see that any decisions about the later versions of the specification lie with AC, AA, AE and AG. However, decisions relating to its earlier versions lie with different subsets of this group at various times. Notably, we can see that AE provides continuity through the evolution of the specification, since he remains its sole documentor and one of its authors.

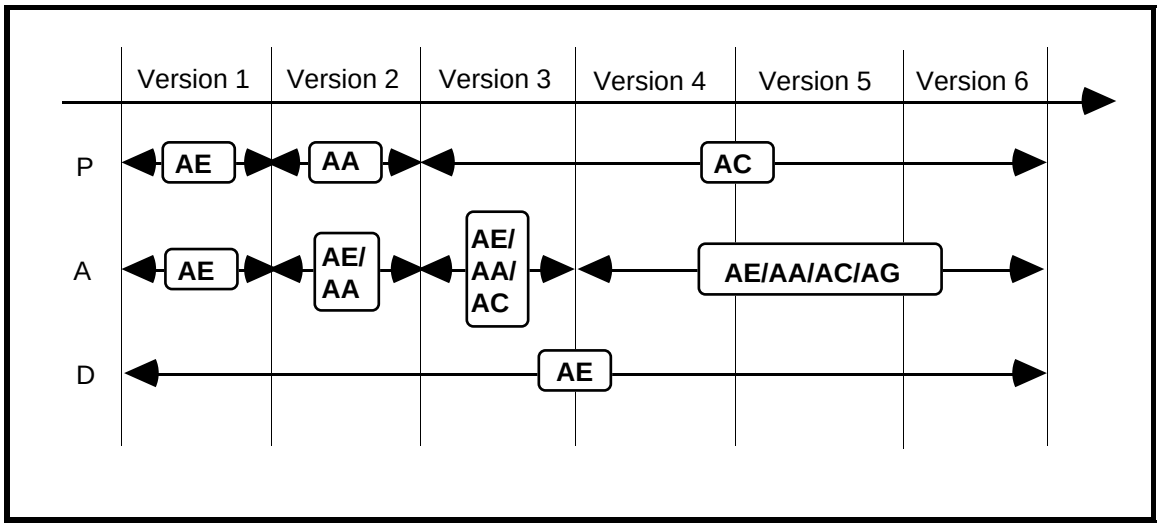

Figure 7: *The changing decision making authority amongst members of the group contributing to the requirements and design specification.*

In Table 7, we further delineate the social roles of the contributors to the different versions of the specification. From this, we can see the subtle transformation in the role relation between AE and AA as other personnel became involved. We can also see how the role relations between all those involved became stable from version four. With such information about how group members have come together, including how their interrelations have

changed, we can begin to examine the impact of changing work arrangements on different attributes of an artifact, like its attention to technical detail.

Requirements and design specification	Social roles and role relations of direct contributors
Version 1 (artifact 46)	AE = True author
Version 2 (artifact 49)	AA = Devisor, AE = Representative author
Version 3 (artifact 61)	AC = Devisor, AA = Ghost author, AE = Representative author
Version 4 (artifact 74)	AC = Devisor, AA/AG = Ghost author, AE = Representative author
Version 5 (artifact 84)	AC = Devisor, AA/AG = Ghost author, AE = Representative author
Version 6 (artifact 88)	AC = Devisor, AA/AG = Ghost author, AE = Representative author

Table 7: *The working arrangements of those contributing to the requirements and design specification.*

By extending the analysis of this artifact into phase three, we note that ultimate responsibility for the specification passed from AC to AH once AC left the project. It did not pass back to one of those who had been its principal earlier on. Whilst AH held this position, no further joint contributions were made by the rest of the original team regarding the specification. They only reassembled once more when AI took over AH's position. It is interesting to point out that development of the specification proceeded successfully in the latter scenario, but was compounded by problems and misunderstandings in the former. Consequently, most of the artifacts produced between October 1992 and February 1993 had little impact on the development of the specification.

5.4. Change notification

"Who needs to be involved in, or informed of, any changes proposed to this requirement?"

Changes were not made to the content of the operations manual after agent AI left the project in phase three. As of version three, artifact 124, only new sections were added to this manual to introduce new features to the generic service. Had a change been proposed to the section introduced in version six, one that specified a new mailbox facility, we could identify all those who had been involved in its original production path to check who would need to be involved in the change process. Similarly, we could identify all those who made subsequent use of it to check who would need to be informed of any changes. These trails are illustrated in Figure 8.

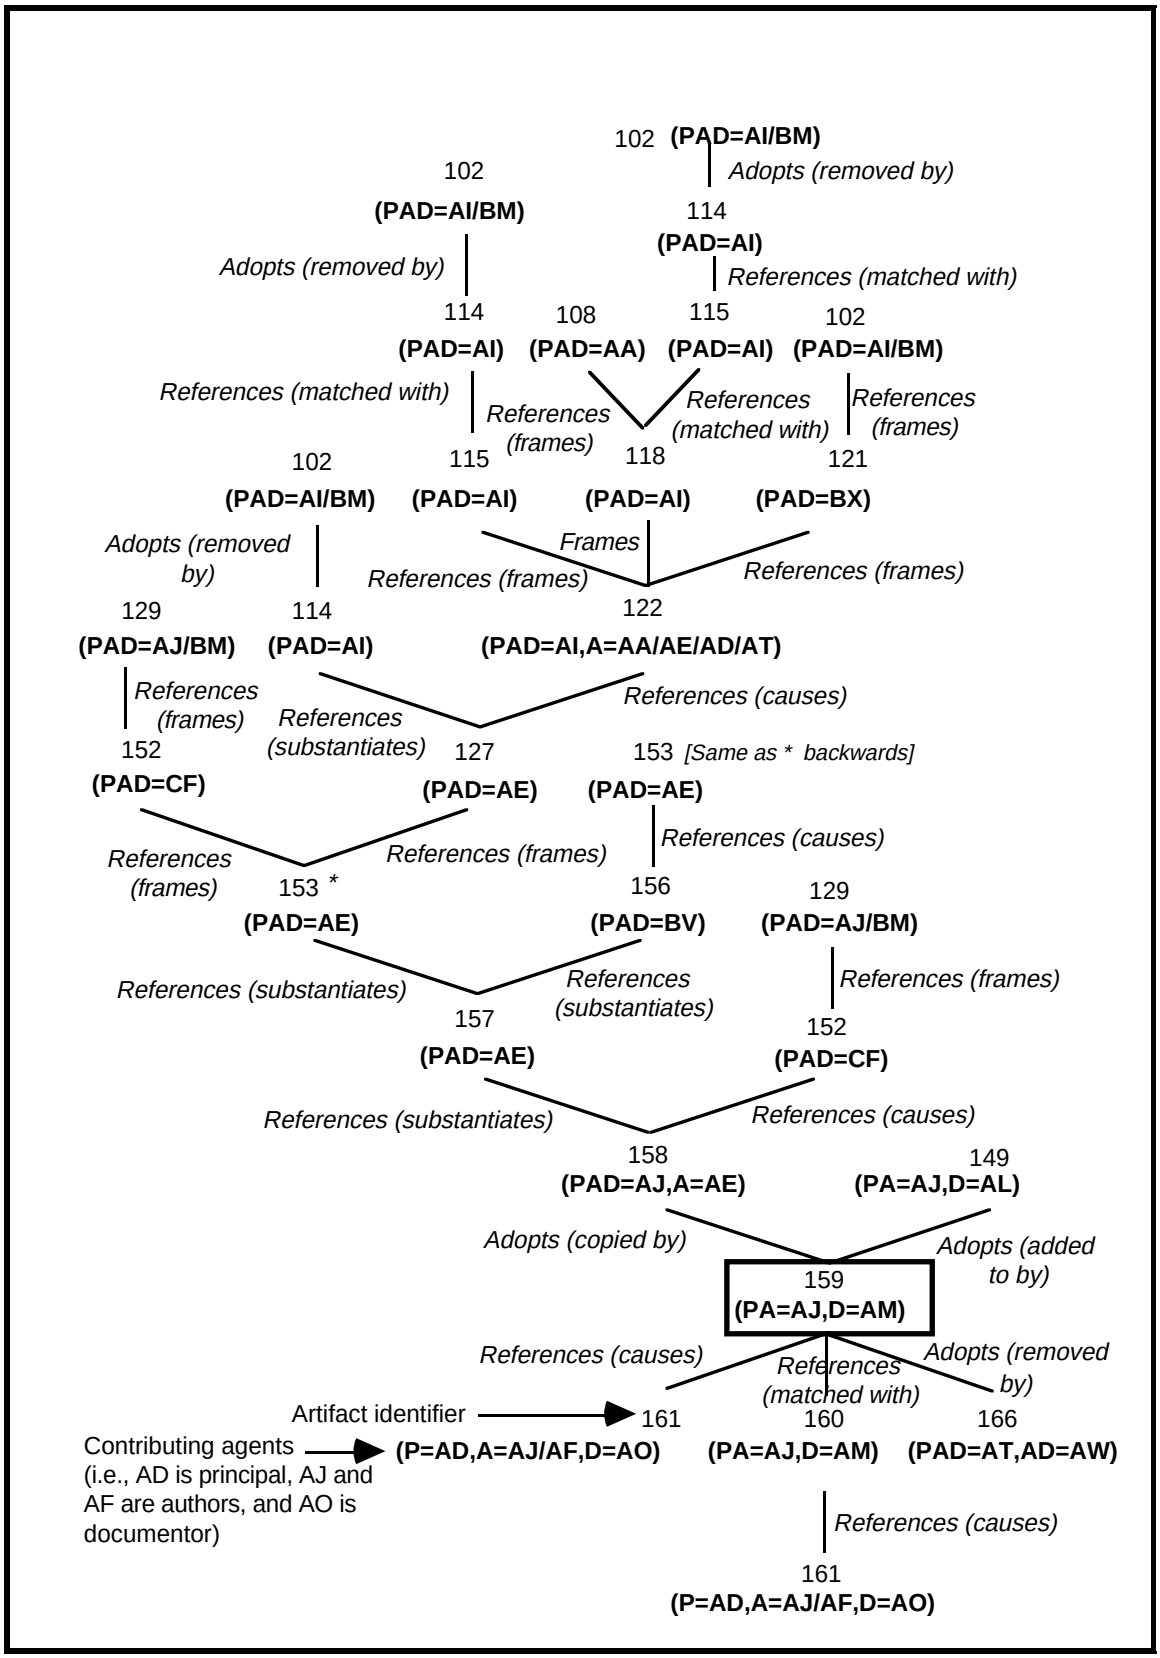

Figure 8: *Who needs to be involved in, and informed of, any changes proposed to the new section in version six of the operations manual, artifact 159. Note that, artifacts 108 and 149 were the earlier versions of the operations manual. We do not pursue their trails in this figure because they did not contain the section under concern.*

20

In examining those involved in the production path of the mailbox section of the artifact, we can see that the additional service arose following correspondence between staff member AJ and a specific customer CF. As noted earlier in Figure 3, artifact 129 is one of the composite artifacts that we used to collate all the correspondence between the company and its potential customers over an extended period in phase four. AJ was the company contact point for this correspondence and BM was the group descriptor for all the customers. Artifact 152 is CF's particular correspondence requesting the mailbox service. We can also see that the requirement for a mailbox service was raised earlier in artifact 114. This artifact listed requirements drawn from all the customer correspondence received in phase three, collated in artifact 102, when AI was the company's contact point. In particular, this requirement had been noted by customer BX, subsequently documented in artifact 122, then reported more fully in artifact 127 after further research.

In examining those involved in the usage trace of the artifact under concern, we can see that artifact 159 is adopted by artifact 166, and referenced by artifacts 160 and 161. If internal links were present from the mailbox section in artifact 159 to artifact 166, we could see that we would need to inform AT and AW of any change. Where project policy is to inform the authors of any artifacts referencing ones which are to be changed, we would be able to see the need to inform AJ and AF.

The project contributors to involve in or inform about change can obviously be determined on a project-specific basis. For example: retrieval can be filtered to inform only the documentors of those artifacts adopting the artifact to which a superficial change is made; or retrieval can be filtered to contact only the principals of any two requirements that need to be reconciled through changes to either. Different types of change or change proposal can thereby be dealt with in the most economical way. Moreover, the people to involve or inform can be determined and notified automatically as changes get proposed or made.

5.5. Ramification

"What are the ramifications, regarding the loss of requirements-related knowledge, if a specific individual or group leaves a project?"

AC left the project at the end of phase two. Before his departure, we can examine which of AC's contributions other project participants are not aware

of, so ensure they are passed on and not lost. We can also identify alternative points of contact for AC's contributions. We list AC's contributions and collaborators in Table 8. We also list those artifacts that adopt or reference AC's contributions to examine their contributors in turn. This can reveal his unused contributions.

Artifact	AC's social role	Other contributors	Adopted by artifacts	Referenced by artifacts
26	True author	None	**30**	**28**
28	Representative	AA = True author	**30**	29
30	True author	None	**34**/**35**/**36**	**31**/32/33/39
31	Devisor	AP = Relayer	None	32/33
34	Devisor	AA = Rep, AE/AG = Ghost author	None	47
35	True author	AA/AE/AG = Ghost author	None	**50**
36	Devisor	AA/AG = Ghost author, AE = Rep	45	44
37	Devisor	AA = Rep, AE/AG = Ghost author	None	47
38	Devisor	AA/AG = Ghost author, AE = Rep	None	44
50	True author	None	**56**	54/55/**85**
51	Devisor	AA = Rep, AE/AG = Ghost author	None	55
52	True author	AA/AE/AG = Ghost author	**58**	**56**
53	Devisor	AA/AG = Ghost author, AE = Rep	**58**	None
56	True author	AE = Ghost author	None	**58**/**73**/**85**
58	True author	AE = Ghost author	60/**61**	**59**
59	True author	AE = Ghost author	None	None
61	Devisor	AA = Ghost author, AE = Rep	**63**/**64**/**65**	None
63	Devisor	AA = Rep, AE/AG = Ghost author	None	68
64	True author	AA/AE/AG = Ghost author	None	None
65	Devisor	AA/AG = Ghost author, AE = Rep	**74**	None
73	True author	AT/AU = Ghost author	**75**/**76**/**77**	**80**/**87**
74	Devisor	AA/AG = Ghost author, AE = Rep	**75**/**76**/**77**	None
75	Devisor	AA = Rep, AE/AG = Ghost author	None	**78**/79
76	True author	AA/AE/AG = Ghost author	**81**	**78**/**80**
77	Devisor	AA/AG = Ghost author, AE = Rep	**81**/83/**84**	**78**
78	Devisor	AA/AG = Ghost author, AE = Rep	**82**	**81**
80	True author	None	**81**	**85**
81	True author	AE = Representative (Rep)	83/**84**	**82**/**85**
82	Devisor	AE = Representative (Rep)	None	**84**
84	Devisor	AA/AG = Ghost author, AE = Rep	**88**	None
85	Nominal author	AQ/AP = Ghost author	None	**86**
86	True author	AQ = Ghost author	None	None
87	True author	AT = Ghost author	None	None
88	Devisor	AA/AG = Ghost author, AE = Rep	None	None

Table 8: *AC's contributions, collaborators and the subsequent project artifacts that have made use of his contributions. Note that, where AC is a contributor to the artifacts cited in columns four and five, the artifact identifier is given in bold.*

By inspection of Table 8, we can see which of AC's contributions are not used in any way by distinct others. We can see that AE must be aware of AC's individual contribution in artifact 80. This is because he adopted its content in artifact 81 when working in conjunction with AC in both an authorial and documenting capacity. However, since AQ, AP and AT are relatively minor players in the project, we note that the key players may be unaware of artifacts 85, 86 and 87. Pin-pointing such artifacts can signal which of AC's artifacts are

still pending approval for integration into the project's critical path. Furthermore, it can help ensure his outstanding commitments are handed over smoothly before leaving.

As we can see who has contributed with AC, and in what role relations, we can pass on this information if there are later queries about any of his contributions. If a new person is to take over AC's commitments, we can identify AC's long-term, transitory and current collaborators for contact purposes. By indicating those who have made use of AC's contributions, especially in conjunction with AC himself, we can identify those who are likely to have had additional communication with AC about any of his artifacts. Potentially, these people can act as replacement points of contact.

5.6. Further analyses

As a by-product of the approach, it becomes possible to carry out other forms of analysis. These can provide value-added information. For example, the number of contributors to each of a project's artifacts, or its *contribution profile*, can highlight phases of group activity and those artifacts perhaps more prone to later query. Similarly, the number and type of contribution made by specific individuals or groups in a project, or their *contributor profile*, can highlight its driving forces or stable backbone. Although premature to generalise, interesting future work would be to consider the health of a project in terms of such profiles.

5. Conclusions

Although fortunate to have access to high-quality case study material, it has some limitations for demonstrating and evaluating our approach. Requirements traceability was not practiced in the organisation studied, and the development philosophy was informal and unstructured. A different perspective would be obtained by those organisations with some form of requirements traceability or document control already in place, or by those currently experiencing problems arising from inadequate requirements traceability. Similarly, by those organisations running larger projects involving many people and artifacts, or by those with explicit process improvement agendas. A summary of the main issues that arose during the case study, concerning the use of the approach and the information it provides, are given in Table 9. These highlight areas for further research.

Main issues concerning use of the approach	Main issues concerning use of the information the approach provides
Whose job is it to record contributors and to insert artifact-based relations? How much is it feasible to do automatically?	The time to analyse and act upon the data has implications for using the approach during a project. How to make its use transparent in activities like change management?
How to balance the granularity and semantics of artifacts and relations against the complexity of the contribution structure model and the potential of the traceability provided?	Overwhelming analytical opportunities for organisational, project and workflow analyses. What information can best inform practice in particular organisations and projects?
How to account for how an "author" actually contributes when there are many authors?	Sensitivity of information. A need to re-examine organisational cultures and introduce policies?
When should details of the undocumented events that influence an artifact, like informal interactions, be captured and how?	Care in analysis and generalisation. Does a large number of contributions indicate productivity, quality, centrality, etc?
How to balance the work involved versus benefits reaped, dependence on stakeholder buy-in, etc? [7] discusses many such issues in detail.	Not related to other forms of organisational modelling. For example, how could contribution structures be used in the context of the Actor Dependency model of [17].
How to expand the social dimension to account for, say, artifact distribution details (i.e., able to examine who contributed as a consequence)?	No metrics provided. A real advance over current practice? Effective in providing answers to personnel questions during a project? etc.

Table 9: *Outstanding issues and research directions*

Despite the above, members of the company agreed that the data we revealed about the contribution structure underlying the project rang true. It identified the right people to help rectify matters where problems of misunderstanding surfaced, to consider requirements change and to handle staff turn-over. It further provided information about social roles and role relations that could not have been determined from the company's organisational chart or work allocation timetables. This information was considered invaluable to inform how work could be allocated in future projects and to entertain the notion of requirements reuse.

Drawing from this case study, we suggest the approach need not be overly labour-intensive if introduced in a suitable setting and in an appropriate manner. For instance, it is best introduced into those organisations that already practice some form of requirements traceability and are concerned with such issues. Moreover, it would need to be introduced incrementally and as an extension to current requirements traceability schemes, however crude. With small extensions, simply distinguishing basic types of artifact-based relation and contribution, it would be possible to trace those involved in different aspects of a project. This would provide a more comprehensive form of requirements traceability. It would also be possible to reveal the working relations of those involved and so inform practice. Eventually, were such information gathered across projects and organisations, it should be possible to investigate how the social organisation of the requirements engineering process itself could be improved.

Acknowledgements

The authors acknowledge the comments and assistance of colleagues, particularly David Michael, Stephen Morris, Wolfgang Emmerich and George Spanoudakis. They would also like to thank Steve Fickas, Eric Yu and the anonymous referees for their recommendations.

References

[1] Allen, J. F. Maintaining Knowledge about Temporal Intervals, *Communications of the ACM*, Volume 26 (November 1983), pp. 832-843.

[2] De Beaugrande, R. A. and Dressler, W. U. *Introduction to Text Linguistics*, Longman (1981).

[3] Goffman, E. Footing, *Semiotica*, Volume 25 (1979), pp. 1-29.

[4] Gotel, O. C. Z. *Contribution Structures for Requirements Traceability*, Ph.D. Thesis, Imperial College of Science, Technology and Medicine, University of London (August 1995).

[5] Gotel, O. C. Z. and Finkelstein, A. C. W. An Analysis of the Requirements Traceability Problem, *Proceedings of the IEEE International Conference on Requirements Engineering*, IEEE Computer Society Press, Colorado Springs, Colorado (April 1994), pp. 94-101.

[6] Gotel, O. C. Z. and Finkelstein, A. C. W. Contribution Structures, *Proceedings of the Second IEEE International Symposium on Requirements Engineering*, IEEE Computer Society Press, York, U.K. (March 1995), pp. 100-107.

[7] Gotel, O. C. Z. and Finkelstein, A. C. W. Revisiting Requirements Production, *Software Engineering Journal*, Volume 11 (May 1996), pp. 166-182.

[8] Harrington, G. A. and Rondeau, K. M. *An Investigation of Requirements Traceability to Support Systems Development*, Naval Postgraduate School, Monterey, California (September 1993).

[9] Laubengayer, R. C. and Spearman, J. S. *A Model of Pre-Requirements Specification (pre-RS) Traceability in the Department of Defense*, Naval Postgraduate School, Monterey, California (June 1994).

[10] Macfarlane, I. A. and Reilly, I. Requirements Traceability in an Integrated Development Environment, *Proceedings of the Second IEEE*

International Symposium on Requirements Engineering, IEEE Computer Society Press, York, U.K. (March 1995), pp. 116-123.

[11] Pinheiro, F. A. C. and Goguen, J. A. An Object-Oriented Tool for Tracing Requirements, *IEEE Software*, Volume 13 (March 1996), pp. 52-64.

[12] Pohl, K. PRO-ART: Enabling Requirements Pre-Traceability, *Proceedings of the Second IEEE International Conference on Requirements Engineering*, IEEE Computer Society Press, Colorado Springs, Colorado (April 1996), pp. 76-84.

[13] Sodhi, J. *Software Requirements Analysis and Specifications*, McGraw-Hill (1992).

[14] Structured Software Systems Limited. *Cradle: Systems Engineering Guide*, Document RM/CRY/006/01, Issue 1, Product Version 1.8X, 3SL, Barrow-in-Furness, Cumbria, U.K. (February 1995).

[15] TD Technologies, Inc. *SLATE: System Level Automation Tool for Engineers*, Marketing Literature, http://www.slate.tdtech.com (1995).

[16] Watkins, R. and Neal, M. Why and How of Requirements Tracing, *IEEE Software*, Volume 11 (July 1994), pp. 104-106.

[17] Yu, E. S. K. and Mylopoulos, J. Understanding "Why" in Software Process Modelling, Analysis, and Design, *Proceedings of the Sixteenth International Conference on Software Engineering*, IEEE Computer Society Press, Sorrento, Italy (May 1994), pp. 159-168.