

Extended Requirements Traceability: Results of an Industrial Case Study

Orlena Gotel & Anthony Finkelstein
Department of Computer Science
City University
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 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, based on modelling the *contribution structure* underlying requirements, was presented in [6]. This paper describes a case study designed to evaluate, through demonstration, whether use of the approach helps answer these outstanding questions and, in so doing, alleviates an important class of requirements traceability problems. The case study is based on a real industrial project.

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 for practitioners to answer as a consequence. We then outline an approach to address this fundamental problem and summarise how it is anticipated to provide answers to these questions. In Section 3, we describe the case study material we gathered and used to validate our

claim. Since the approach did not exist at the outset of the case study project, its requirements did not drive the data gathering and its use was not on the project's critical path. In Section 4, we demonstrate how the approach was applied to this data in a post-hoc manner. This application means we are only in a position to validate the feasibility of the approach and the usefulness of the information it provides, say to reveal information about the project's evolution and assist the maintenance process, in a subjective and historical manner. We do this in Section 5, where we show how this information makes it possible to answer questions regarding *involvement, responsibility, ramifications, change notification* and *working relationships*. Based on our experiences and practitioner comments, we highlight some outstanding issues and make recommendations in Section 6.

2: Contribution structures for 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 a 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, 13, 14], 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 and for what purposes.

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 often considered the ultimate baseline whenever requirements need to be re-examined or re-worked. Examples of these problematic questions include:

- (1) Who has been involved in the production of this requirement and how?
- (2) 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?
- (3) At what points in this requirement's life have the working arrangements of all involved been changed? Accordingly, within the remit of which groups do decisions about this requirement lie?
- (4) Who needs to be involved in, or informed of, any changes proposed to this requirement?
- (5) What are the ramifications, regarding loss of project knowledge, if a specific individual or group leaves?

2.2: Contribution structures

We described an approach to address this more focal problem in [6]. The approach is based on modelling the contribution structure underlying requirements. This model reflects the network of people who have contributed to the artifacts produced in the requirements engineering process. In [6], we also described how the

approach can be implemented and gave scenarios of use. Formalisation of the approach and the inferences it supports can be found in [4]. We only summarise the main steps of the approach in Figure 1 and below.

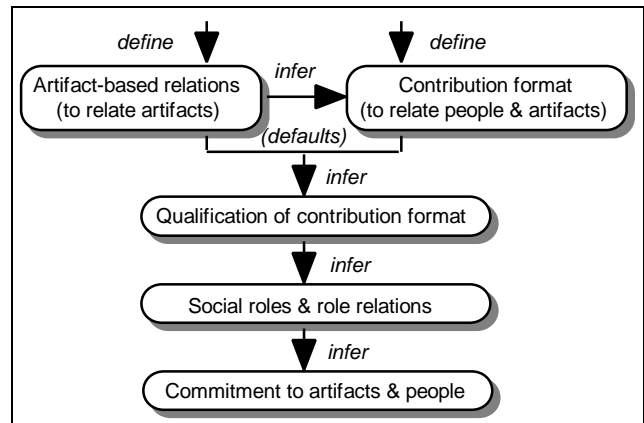


Figure 1: Main steps of the approach.

Working through Figure 1, minimal semantics are given to the artifact-based relations ordinarily put in place for requirements traceability. For example, based on the notion of *communicative function*, an artifact can either *reference* or *adopt* the content of a linked artifact, depending on whether or not their content overlaps. A record of the people who contributed to an artifact's production is maintained in its *contribution format*. For example, based on Goffman's work on the nature of participation in social encounters [3], this can delineate the *principal*, *author* and *documentor* of an artifact. These categories were chosen for their analytic potential to infer details about the social roles, role relations and commitments of those involved. The information provided from these steps makes it possible to extend conventional forms of *artifact-based* requirements traceability with a form of *personnel-based* requirements traceability.

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 communications company employing about twenty-five people. The company runs many projects concurrently, providing solutions to various communications-related problems. The objective of this particular project was to supply a dedicated communications 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 for other customers was discussed. Six versions of a requirements and design specification were drawn up throughout September 1992. These were then abandoned until the end of October 1992, when new staff were employed 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 its 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, customer contracts and miscellaneous correspondence. Requirements traceability has not been maintained. Those still involved in the project are no longer aware of 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 has been small and exhibited some staff continuity, these have not been critical to its maintainability and success.

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, took 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 this specification. We also maintained a record of the people involved in the production and distribution of these artifacts.

One hundred and sixty-six main artifacts were produced in the project. These relate to four main phases:

- (1) Development of the customer-specific service (twenty-three artifacts between February and March 1992).
- (2) Development of the baseline for the generic service (sixty-five artifacts from August to September 1992).

- (3) Development of the initial generic service (thirty-nine artifacts from October 1992 to July 1993).

- (4) Extension of the generic service to address new requirements (thirty-nine artifacts from September 1993 to June 1995).

For the purposes of the case study, our definition of "artifact" applied to single physical documents. This was to promote identification and to enable us to examine the viability of the approach at a coarse level of granularity before introducing further complexity.

Fifty-eight people contributed directly to the project. These included individuals and groups from within the company and from outside. To maintain confidentiality, we use alphabetic identifiers when we refer to these individuals and groups in the remainder of this paper.

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 some semantics. They were also tasked to reconstruct the contribution format for each artifact, prompted by contextual material. We then applied the last three steps of the approach to examine what could be inferred about the project and its social roles, role relations and people's commitments.

4.1: Artifact-based relations

For each project phase, its artifacts were numbered according to production order. The temporal relations between them was then clarified, based on [1]. The coarse flow-down of information and influence amongst these artifacts was also established. These orderings for the artifacts produced in phase one are shown in Figure 2.

Table 1 shows how semantics were assigned to these relations. From the original reason provided for the relation by participants, the nature of the relation was categorised according to classifications of *cohesion* and *coherence* [2]. Based on this classification, its broad communicative function was identified as either *referencing* or *adopting*. Although the more detailed semantics have implications for selective traceability, consensus was found difficult to establish at that level, whilst easier to agree at the coarser level.

Figure 2a highlights the adopts relations of Table 1. Since these tend to capture *parent-child* or *predecessor-successor* relations, they provide for what we regard as conventional forms of artifact-based requirements traceability. Figure 2b goes on to illustrate how the

references relations of Table 1 provide additional contextual information that is often not integrated and used for requirements traceability purposes.

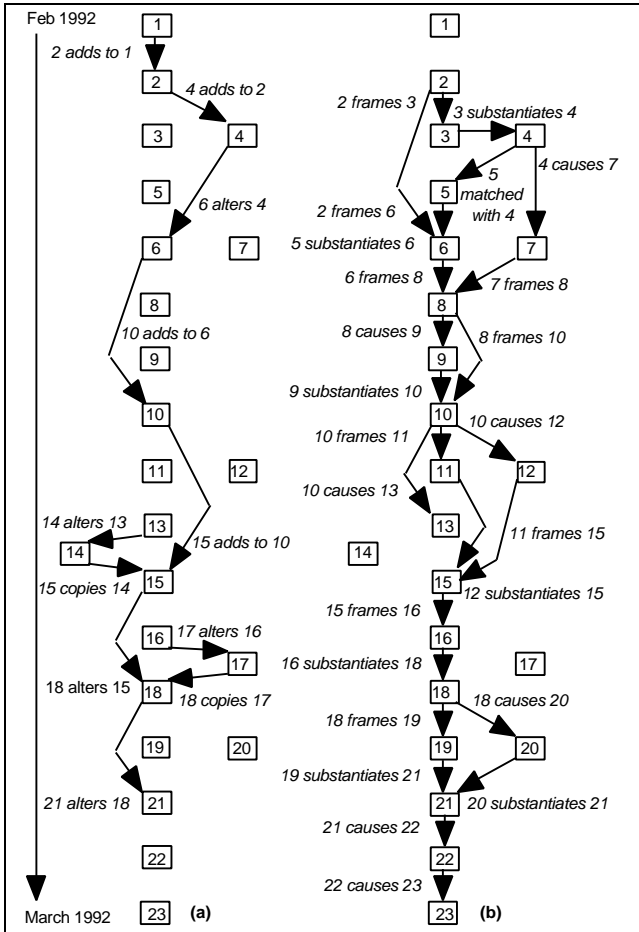


Figure 2: Relations of phase 1: (a) adopts relations - arrows suggest flow-down of content; (b) references relations - arrows suggest direction of influence.

4.2: Contribution format

The contribution format of each artifact was established to indicate the individuals and/or groups who contributed in the capacities of principal, author and documentor. The contribution formats of some of the artifacts produced in phase one are shown in Table 2.

4.3: Qualification

The capacities of each contribution format were qualified to provide more details about contributions and contributors. As an example, the authorial capacity was qualified to indicate the levels and types of dependency upon other authors, providing a citation-like network.

Following the resulting authorial trails, we are able to see: how each progressive author made use of previous people's contributions; which authors produce the most original artifacts; which authors 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 types of query.

(1) Informal description of relation given by practitioner	(2) Relation in terms of cohesion & coherence	(3) Broad communicative function
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 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: Semantics for the relations of phase 1.

Artifact	Principal	Author	Documentor
1	BH	BI	AW
2	AT	BB={AW/AV/AT/AR/AX/AU}	BB={AW/AV/AT/AR/AX/AU}
3	AA	AA/AE	AA
4	AA	AA/AE	AU
5	AA	AA/AQ/AP/BB={AW/AV/AT/AR/AX/AU}	AA
6	AA	AA/AT	AA
7	BH	BI	BL
8	AA	AA/AE	AA
9	AA	AA/BB={AW/AV/AT/AR/AX/AU}	AA
10	AA	AA	AA

Table 2: Contribution formats for artifacts 1 to 10. AA/AE means person AA & person AE were joint contributors in the given capacity. BB refers to a group, so its members are given in curly brackets.

4.4: Social roles & role relations

The social roles that people assume when contributing to artifacts can be inferred from the information 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 contribute jointly 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, we can also see how they collaborated and whether these role relations have varied or been sustained throughout a project.

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.

Social role of AI	On how many artifacts?	How many on own?	Social roles of people who collaborate with AI & number of times
True author (i.e., contributes as P, A & D)	16	13	BM=true author (x1) AA/AE/AD=ghost author (x1) AA/AE/AD/AT=ghost author (x1)
Nominal author (i.e., P & D)	2	0	AA/AE/AD/AG=ghost author (x1) BB/AP/AQ=ghost author (x1)
Representative (i.e., A & D)	2	0	AD=sponsor (x2)
Ghost author (i.e., A)	2	0	AD=sponsor & BO=relayer (x2)

Table 3: Social roles & role relations for AI.

From these tables, we can see that AI worked on his own on over half 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. Although details delineating the type and content of artifacts have not been included in these tables, it is noteworthy that AD tended to collaborate with AI as a *sponsor* when dealing with customer-related artifacts. In contrast to AI, we can see that AJ worked rarely on his own and collaborated mainly with one or two others. He had a strong dependency on AE as his *ghost author* when they worked together and on many other people as *relayers*, the latter hinting at the need for secretarial support. It is noteworthy here 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 people who collaborate with AJ & number of times
True author (i.e., P, A & D)	7	2	BM=true author (x1) AE=ghost author (x4)
Ghost author (i.e., A)	9	0	AD=sponsor & AP/AS=relayer (x1) AD=sponsor & AP=relayer (x5) AD=sponsor & AR=relayer (x2) AD=sponsor, AF=ghost author & AO=relayer (x1)
Devisor (i.e., P & A)	9	0	AE=ghost author & AW=relayer (x2) AE=ghost author & BQ=relayer (x1) AL=relayer (x4) AM=relayer (x2)
Sponsor (i.e., P)	1	0	AE=true author (x1)

Table 4: Social roles & 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. Since AE had also collaborated with AI as ghost author during the earlier project phase, this collaboration obviously served to maintain some continuity.

4.5: Commitment

Table 5 indicates the kind of information that can be inferred about the commitments of project contributors, both to artifacts and to other people. We can see that, as AP has predominantly been a relayer (i.e., purely a documentor), she is mainly responsible for physical artifacts. She is only responsible for the content of artifacts when collaborating with others. She has never been responsible for their ultimate effect (i.e., a principal). We can also see the people AP is committed to through their collaboration on artifacts. For instance, we can identify AD and AJ as those with whom AP has collaborated the most often, as well as identify 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 in the wider network of artifacts.

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 type of artifacts and for the same aspects; which people have contributed to those artifacts that are the initial sources of requirements; and so on.

Artifacts AP contribute s to	Aspect of artifact committed to (is the commitment shared with others?)	Other contributors & number of artifacts on which collaborate
5	Content (as one of many contributors)	AD (x6)
22	Physical (on own)	AJ (x6)
31	Physical (on own)	AQ (x4)
41	Content (as one of two) Physical (on own)	AT (x4)
85	Content (as one of two)	AU (x4)
96	Content (as one of many) Physical (as one of two)	AW (x3)
100	Content (as one of many)	AV (x3)
111	Content (as one of many)	AR (x3)
139	Physical (as one of two)	AX (x3)
140	Physical (on own)	AC (x2)
148	Physical (on own)	AD (x2)
151	Physical (on own)	AJ (x2)
154	Physical (on own)	AQ (x1)
155	Physical (on own)	AT (x1)

Table 5: AP's artifact & collaborator commitment stores.

5: Results & discussion

In this section, we demonstrate how the questions of Section 2 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, led to much investigation and many artifacts that later became redundant. It was a requirement pursued throughout phase two of the project and cited in all six versions of the specification. Once dropped in phase three, 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.

Following application of the approach, this requirement was traced back to artifact 27. Note that, we define an “original contribution” to be one that does not depend upon other artifacts for its existence; we do not attempt to measure degrees of originality here. The contribution format at the source shows that AA was writing requirements in the name of a collective. Having delineated the contribution format of internal components, AX can be identified as the member who originated this particular requirement. Knowledge of the source makes it possible to recover AX's original intention, one that was actually misconstrued by AA in the project. Furthermore, we can see how this misconstrued requirement pervaded subsequent artifacts, due to AA's backing and no later 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.

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 and its earlier versions.

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
Document	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
Principal (of referenced artifact)	AJ	AJ	AJ	AI	(a)AI (b)AA
Author (of referenced artifact)	AJ	AJ	AJ	AI	(a)AI (b)AA
Doc (of referenced artifact)	AM	AL	AL	AI	(a)AI (b)AA

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

From Table 6, we can see the transition between AI's original work on the guide and AJ's later work on it. We can also see that AJ only made additions to what AI

originally produced. The working arrangements also changed from AI working on his own to AJ working in conjunction with one other person doing the physical 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 now responsible for all physical aspects of the document. 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 of the guide is related is artifact 108. Inspection of this artifact reveals that AA, as its true author, was originally responsible for the idea to develop this guide.

5.3: Working arrangement & remit

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

In Figure 3, we depict the contributors to the formal versions of the requirements and design specification produced in phase two. 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. 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 3: Changing decision making authority amongst members of the group contributing to the requirements & design specification.

In Table 7, we delineate the social roles of the contributors to the different versions of this specification. From this, we can see the subtle transformation in the role relation between AE and AA as other people became involved. We can also see how the role relations between all those involved became stable with version four. With such information about how group members come together, including how their interrelations change, we can begin to examine the impact of changing working arrangements on different attributes of an evolving artifact, like its attention to technical detail and so forth.

Requirements & design specification	Social roles & role relations of direct contributors
Version 1 (artifact 46)	AE=true author
Version 2 (artifact 49)	AA=devisor, AE=representative
Version 3 (artifact 61)	AC=devisor, AA=ghost author, AE=representative
Version 4 (artifact 74)	AC=devisor, AA/AG=ghost author, AE=representative
Version 5 (artifact 84)	AC=devisor, AA/AG=ghost author, AE=representative
Version 6 (artifact 88)	AC=devisor, AA/AG=ghost author, AE=representative

Table 7: Working arrangements of those contributing to the requirements & design specification.

By extending the analysis of this specification into phase three (not shown above), we are able to see that the ultimate responsibility for the specification passed from AC to AH once AC left the project. Interestingly, it did not pass back to one of those who had held this responsibility earlier on. Whilst AH held this position when contributing to the evolution of the specification, no further contributions were made to it by members of the original team. They only reassembled when AI took over AH's position in phase four. This information helps to explain why development of the specification proceeded successfully in phase four, but was compounded by problems and misunderstandings in phase three.

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 AI left the project in phase three. As of version three, artifact 124, each new version saw the introduction of a new section to add novel features to the generic service. Had a change been proposed to the section introduced in version six, artifact 159, a section that described a new electronic mailbox service to be implemented, we would be able to identify all those who contributed, so able to check who would need to be involved in the change process. Similarly, we would be able to identify all those who made subsequent use of this service in later work, so able to check who would need to be informed of changes. These trails are shown in Figure 4. Only crude details are provided about artifact content in the figure. Trace visualisation is an on-going research issue that is not explored here.

In examining those involved in the production path of the mailbox service, we can see that it arose following a request from a specific customer, CF, in correspondence captured in artifact 152. We can also use

these trails to see that the requirement for a mailbox service was raised earlier in artifact 114, a list of requirements drawn from all the customer correspondence received in phase three. In particular, this requirement had been noted by customer BX, subsequently documented formally in artifact 122, then reported more fully in artifact 127. This report was used as background material when the requirement for the service surfaced again later.

In examining those involved in the usage path of artifact 159, we can see that it is adopted by artifact 166,

and referenced by artifacts 160 and 161. If internal links were present from the section on the mailbox service in artifact 159 to artifact 166, we could see we would need to inform AT and AW of any change. Where project policy is to inform the authors of any artifacts referencing ones that are to be changed, we would also be able to see the need to inform AJ and AF. With knowledge of such trails, different types of change or change proposal can be dealt with in the most desirable way, on a project-specific basis, with automatic notification.

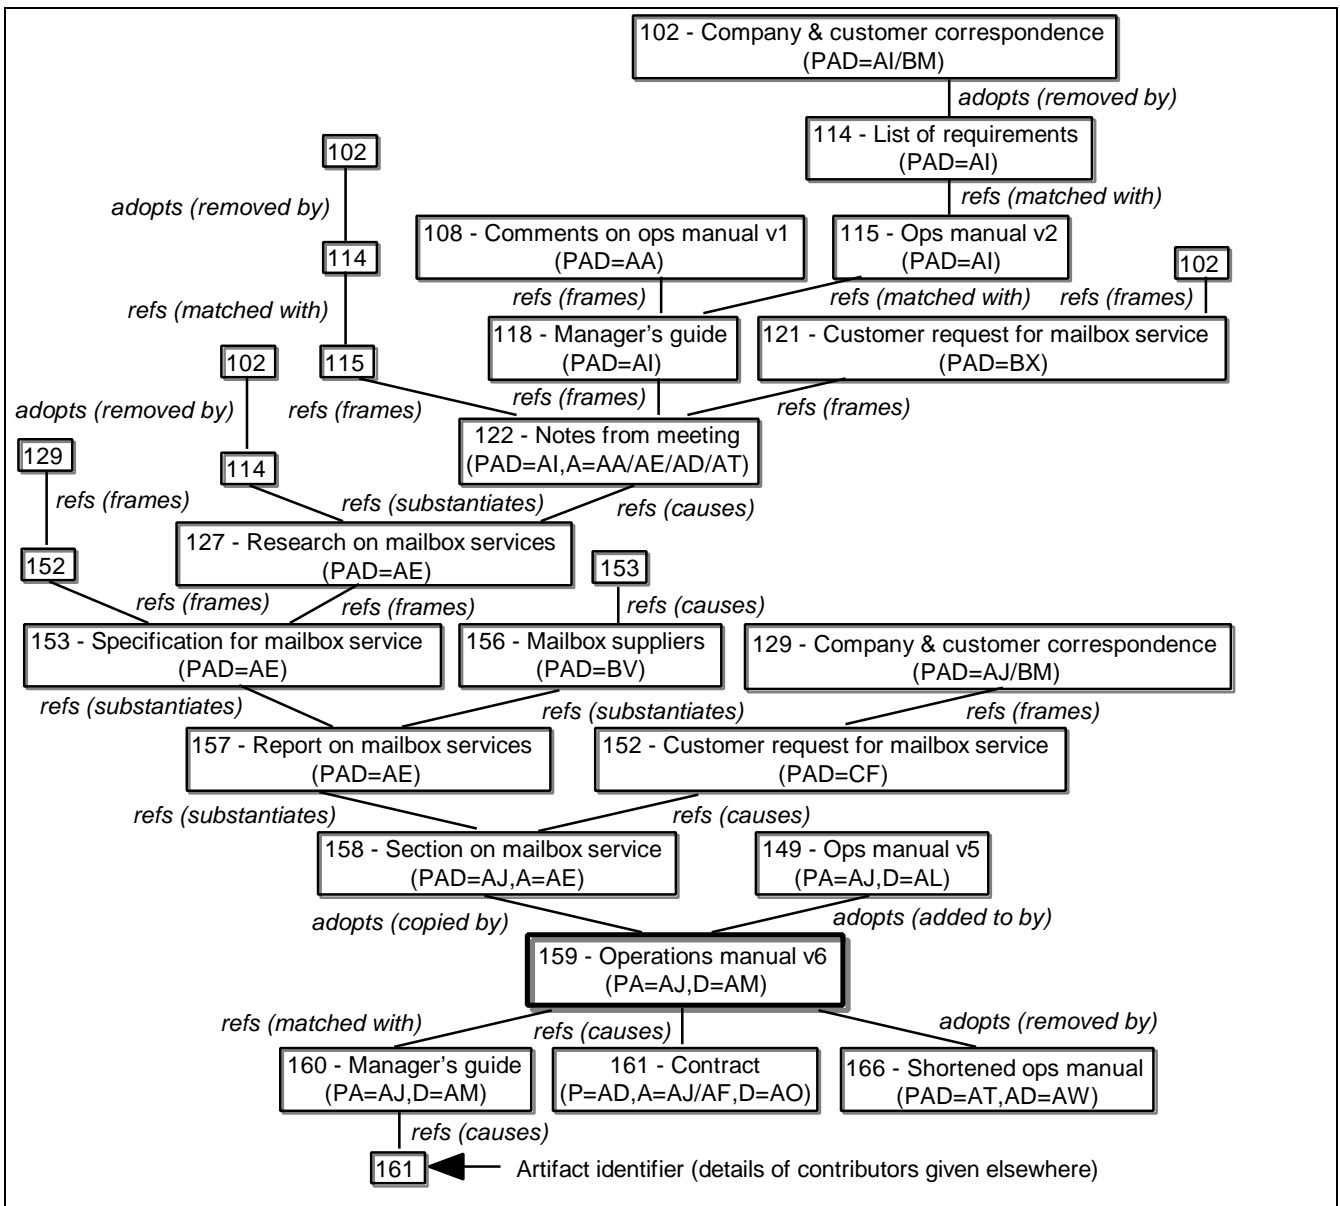


Figure 4: Who needs to be involved in, & informed of, any changes proposed to the new section of the operations manual introduced in version six, artifact 159. (PA=AJ, D=AM means person AJ is principal & author, person AM is documentor.)

5.5: Ramification

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

AC left the project at the end of phase two. Before his departure, we can examine which of AC's contributions are unused by other people, so those that the other project participants are probably not aware of. In this way, we can ensure that his outstanding commitments are passed on and not lost, identify alternative points of contact for AC's contributions, so smooth staff turn-over. We list AC's contributions and collaborators in Table 8. We also list those artifacts that adopt or reference AC's contributions in this table to examine their contributors in turn.

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

Table 8: AC's legacy. Where AC is a contributor to the artifacts cited in columns three & four, its identifier is given in bold. (rep = representative, ghost = ghost author.)

By inspection of Table 8, we can see which of AC's contributions are not used in any way by distinct others. Firstly, 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. We can also see that, since AQ, AP and AT are relatively minor players in the project, we might need to alert the key players to artifacts 85, 86 and 87. We can thereby signal which of AC's artifacts are still pending approval for integration into the critical path.

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 concerning his artifacts. Potentially, these people can act as replacement contact points.

5.6: Further analyses

It becomes possible to carry out other forms of analysis as a by-product of the approach. These can provide much value-added information. For example, the number of contributors to each artifact in a project 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 can highlight its driving forces and its stable backbone. Although premature to generalise, interesting future work would be to consider the health of a project in terms of contribution and contributor profiles.

6: Conclusions

Members of the company we studied 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; those to involve in requirements change; how to handle staff turn-over; amongst other things. In particular, it 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.

However, although we were fortunate to have access to high-quality material, the case study has some limitations for demonstrating and evaluating our approach: requirements traceability was not practiced in the organisation studied; the development philosophy was informal and unstructured. A different perspective would no doubt arise in those organisations with some form of requirements traceability or document control already in place, or by those currently experiencing problems caused by 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 suggest areas for further research.

Drawing from this case study, we suggest that the approach is practical and feasible. Furthermore, it need not be overly labour-intensive if introduced in a suitable setting and in an appropriate manner. For instance, if introduced into organisations that already practice some form of requirements traceability, incrementally and as an extension to current requirements traceability schemes. Even with crude extensions distinguishing basic types of artifact-based relation and contribution, it becomes possible to trace those involved in different aspects of a project and to reveal their working relations. This provides for a more comprehensive form of requirements traceability that is able to answer many problematic and outstanding questions. Eventually, were such information gathered across projects and organisations, it would become possible to investigate how the organisation of the requirements engineering process itself impacts practice. This information could be used as a basis for process and quality improvement programmes.

Main issues concerning use of the approach	Main issues concerning use of the information the approach provides
Whose job is it to record contributors & to insert artifact-based relations? How much is it feasible to do automatically?	The time to analyse & act upon the data has implications for use during a project. How to make its use transparent in activities like change management?
Balancing the granularity & semantics of artifacts & relations against the complexity of the contribution structure modelled & the traceability provided.	Overwhelming analytical opportunities for organisational, project & workflow analyses. What information can best inform practice in particular organisations & projects?
A need to account for how an "author" actually contributes	Sensitivity of information indicates a need to re-examine

when there are many authors.	organisational cultures & to introduce use policies.
When should details of the undocumented events that influence an artifact, like informal interactions, be captured & how?	A need to take care in analysis & generalisation. Does a large number of contributions really indicate productivity, quality, centrality, e.t.c.?
How to balance the work involved versus the benefits reaped? How to ensure commitment? e.t.c.	Integration with other forms of organisational modelling (e.g., how could it be used in the context of the Actor Dependency model [15]?)
A need to expand the social dimension. How to account for artifact distribution details, so we can examine who contributes as a consequence?	No metrics provided. Is it a real advance over current practice? Is it cost-effective in providing answers to personnel questions during a project?

Table 9: Outstanding issues & research directions.

Acknowledgements

The authors acknowledge the comments and assistance of colleagues, particularly David Michael, Wolfgang Emmerich, Stephen Morris and George Spanoudakis. They would also like to thank the company of the case study. Finally, thanks to 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] 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).
- [14] TD Technologies, Inc. *SLATE: System Level Automation Tool for Engineers*, Marketing Literature, <http://www.slate.tdtech.com> (1995).
- [15] 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.