Revisiting Requirements Production[1]

ORLENA GOTEL† and ANTHONY FINKELSTEIN‡

†Department of Computing
Imperial College of Science, Technology and Medicine
London SW7 2BZ
‡Department of Computer Science
City University
London EC1V 0HB

Synopsis

Key information relating to requirements production is either lost or unrecoverable as a consequence of conventional requirements engineering practices. This includes, most significantly, information about the individuals and groups from whom requirements were originally elicited and by whom they were subsequently refined. We outline an approach to render such information traceable, which anchors requirements and development information in the contribution structure from which it arose, and enables the workings of the requirements production process to be revisited in later phases of development. We demonstrate how use of the approach can prevent the premature freezing of requirements elicitation and allows us to establish the consequences of organisational change.

1 Introduction

Systems and software continue to be built which fail to meet user and customer requirements on their delivery. This situation is due to a combination of factors like: inadequate requirements elicitation; the inability to transcribe elicited requirements into a

[1]Note that this is an extended version of a position paper presented at the workshop *Requirements Elicitation for Software-Based Systems* (RESS) held at Keele University in July 1994.

tangible or representative form; the difficulty in reconciling diverse requirements; and the different interpretations given to requirements throughout their use. Current research in requirements engineering therefore attempts to address such issues by: extending the repertory of techniques that can be used to uncover different types of requirements [1, 2]; examining how data gathered using novel techniques for elicitation, like ethnography and scenarios, can best be presented and used to inform the requirements engineering process [3, 4]; providing strategies to deal with multiple, possibly conflicting, viewpoints with regard to requirements [5, 6]; and either developing specialised languages in which to describe requirements less ambiguously or providing more context for documented requirements [7, 8].

Although the above advances in requirements elicitation, negotiation strategies, and requirements description can offer more assurance that user and customer requirements are obtained, agreed upon, and recorded, they offer no guarantee that these requirements will drive subsequent development, eventually get met, or that any additions or later changes can be handled and taken into account. Unless they are coupled with techniques which enable these requirements to be considered and reconsidered throughout the entire development process, as they either emerge or evolve, systems and software will still be delivered which fail to meet them. Such techniques need to establish and maintain some sort of connection between the information that has been elicited from users and customers, the requirements which have been derived from this information, and the subsequent artifacts in which these requirements have been distributed. In other words, the ability to obtain and meet user and customer requirements depends on them being *traceable* from their origin and throughout their project life.

In this paper, we describe why providing for *requirements traceability* is not only a major sub-goal of requirements engineering, but is crucial for promoting on-going requirements elicitation and analysis. We explain how endeavours to increase the potential for this requirements traceability primarily involve uncovering and recording evermore

comprehensive information about the requirements engineering process. We then argue that the most fundamental information to uncover and record for such purposes is that which can identify the human sources of requirements, requirements-related information, and requirements-related work. Although people are the ultimate baseline to trace back to when requirements need to be examined or re-worked, the strive to supplant the need for human contact with extensive rationale, decision records, and the like, does not prepare appropriate foundations to actively assist the identification of those who have contributed along the way in various capacities. We outline an approach, based on modelling the *contribution structure* underlying requirements, which has been developed to focus on such information. We then indicate how the resulting contribution structure can be used, in conjunction with conventional forms of *artifact-based* requirements traceability, to achieve selective forms of *personnel-based* requirements traceability. Finally, we describe how we have demonstrated the feasibility of the approach in practice, point to some outstanding issues, and finish with a critical evaluation.

2 Traceability of requirements production

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) [9]. More specifically, *pre-requirements traceability* is concerned with requirements production and refinement, whilst *post-requirements traceability* is concerned with requirements deployment and use. These two basic types of requirements traceability revolve around a documented specification of requirements, or equivalent baseline, and are illustrated in Figure 1. It is the former type that is of interest for the purposes of this paper. This is because the objective of pre-requirements traceability is to render the workings of the requirements engineering process explicit and, by so doing, provide the ability to describe and follow those aspects of a requirement's life prior to its inclusion in the requirements specification.

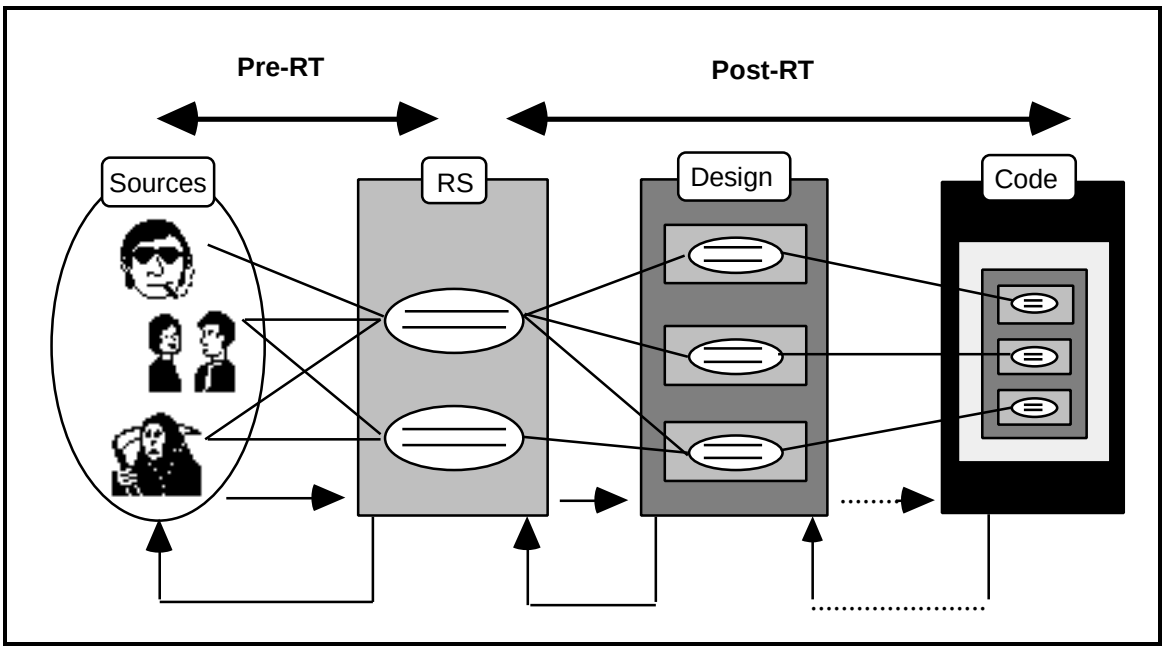

Figure 1: *A simplified diagram to show pre-requirements traceability and post-requirements traceability. For clarity, we only show vertical forms of requirements traceability through a linear progression of subsequent documents. We ignore multiple subsequent documents, the many intermediate documents that are likely to be present, and horizontal forms of requirements traceability between versions. Note the way in which requirements information is distributed and merged in successive representations. Note also the added complication of iteration and change propagation.*

2.1 Pre-requirements traceability

An empirical analysis of the requirements traceability problems commonly experienced in practice indicated that the majority of the problems currently attributed to requirements traceability were in fact due to inadequate pre-requirements traceability [9]. Moreover, many of these more specific problems were found to arise when practitioners were unable to trace particular information that they required about the production and refinement of requirements prior to their occurrence in baselined requirements specifications. For instance, the origins of the dispersed needs, and how these needs have come to be integrated in the requirements specification. These problems are compounded because

there is a prevailing tendency to prematurely commit to requirements and subsequently "black box" them when they are documented. The absence of production details renders these requirements closed to later interrogation and so fails to support an exploratory approach to requirements engineering in which requirements can emerge and evolve over time.

Requirements engineering techniques and tools which are able to couple details about the requirements elicitation and analysis activities with their end products are therefore highly desirable. They can provide the information that is needed to achieve artifact-based forms of pre-requirements traceability. The requirements for improving the availability of requirements information for the purposes of pre-requirements traceability, along with an account of the many ways in which progress towards meeting these requirements are being sought, are summarised in [9] and discussed more fully in [10].

2.2 Traces to the human sources

Although techniques and tools which capture comprehensive information about the requirements engineering process are needed to improve pre-requirements traceability, they typically aim to remove the need for human contact through the information they generate. With this as an objective, they do not actively provide support for human contact as an eventuality. However, a significant finding from the empirical work we referred to earlier was the extreme importance that practitioners attached to personal contact and informal communication, particularly with those who had been involved in requirements production and refinement. This was found to be essential, not only to cope with those situations in which information they required about requirements production was absent, but also to consolidate, supplement, or question any information which was available, to carry out validation and verification of requirements with stakeholders, amongst numerous other activities.

The ability to augment any information that is obtained about the requirements engineering process with details of those agents who have contributed to its production is therefore one way to support this somewhat evident and fundamental working practice. This can provide the potential to extend conventional forms of artifact-based pre-requirements traceability with the forms of personnel-based pre-requirements traceability that practitioners were found to need to identify human sources. Currently, information which serves to identify such sources is either: overlooked in practice; points to the broad notion of a document "author", without taking into account different areas and types of authorship; not maintained over time as information gets refined and used elsewhere by others; or is not oriented to the specific traceability purposes for which such information was found to be required. Not only do a number of contentious "political" issues surround the keeping of such information, but many problems stem from the typically coarse-grain and snapshot nature of labelling the authors. To address these issues, we have developed an approach based on modelling and using the contribution structure underlying requirements. This was first described in [11] and is the subject of Section 3.

2.3 Artifact-based and personnel-based requirements traceability

Artifact-based requirements traceability makes use of the relations which exist between requirements artifacts themselves, like between a requirement in a requirements specification and its high level design in a design specification, or between alternative versions of the same requirement. It can be considered the conventional notion of requirements traceability since it allows us to answer questions like: "What is the ultimate source and eventual realisation of this requirement?"; "From which previous requirement has this requirement been derived?"; and "Which artifacts are impacted if this requirement is removed and how?"

Personnel-based requirements traceability makes use of the relations which exist between agents and requirements artifacts, as a consequence of the former directly contributing to

the latter in some way, so highlights the various working relations which have been formed between all the contributing agents. It therefore allows us to answer questions like: "Who has been involved in the production of this requirement, how, and on what other requirements have the same agents been involved?"; "At what points in this requirement's life have the working arrangements of all those involved been changed and with what implications?"; and "What are the ramifications, with respect to the loss of possible requirements-related knowledge, if a specific individual or group leaves a project? Who would be the best back-up source of information here?"

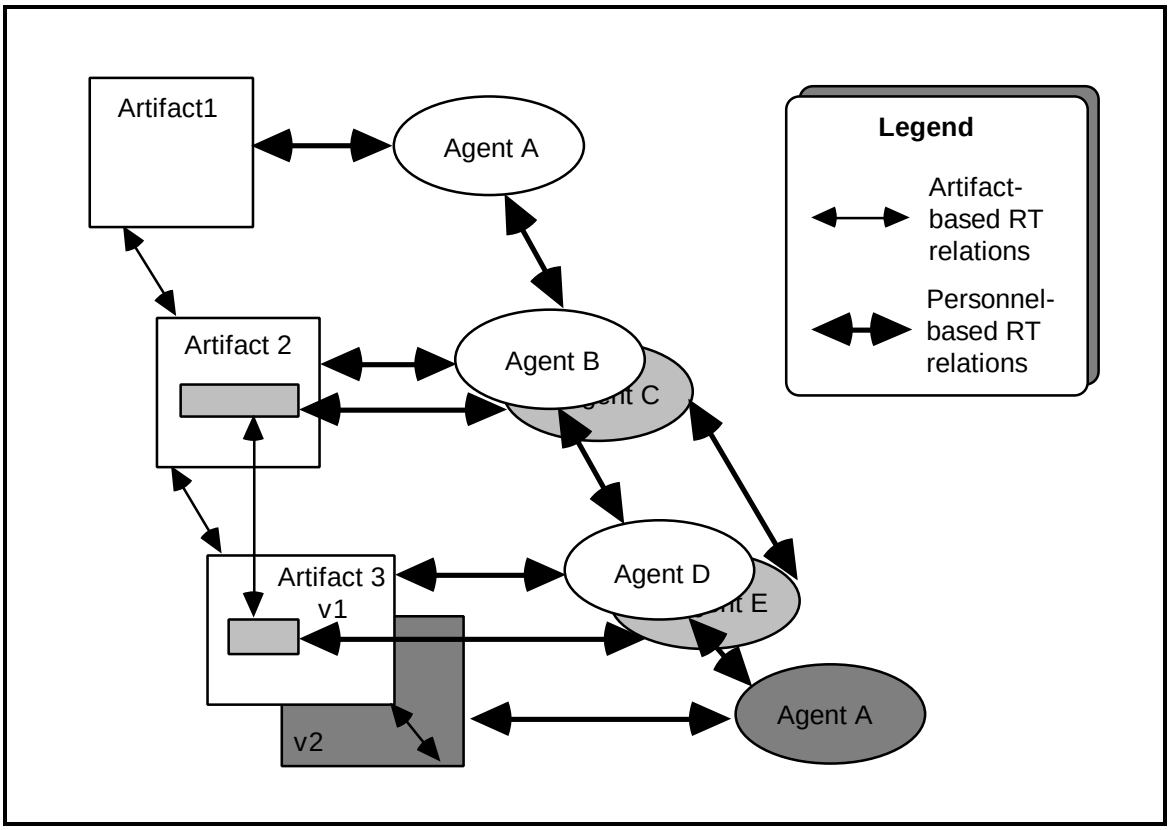

Figure 2: *Artifact-based requirements traceability is provided through the relations that are commonly established between requirements artifacts to depict requirements flow-down in contemporary requirements traceability schemes. Personnel-based requirements traceability is provided through the addition of a relation, established between agents and artifacts, to denote those agents who have contributed to an artifact's production.*

For clarification, the distinction between artifact-based and personnel-based requirements traceability is illustrated in Figure 2. Together, these provide a comprehensive approach to requirements traceability, particularly if they account for both pre-requirements traceability and post-requirements traceability. In turn, this comprehensive approach can better support on-going requirements elicitation and analysis.

3 An approach to model contribution structures

We use the term "contribution structure" to refer to the overall system of agents who have directly contributed to the tangible artifacts produced in the requirements engineering process. This is a subset of the full social structure underlying the process as it only accounts for the first quadrant of the table shown in Figure 3. Extending the coverage to account for the fact that the agents found in each of the quadrants frequently coincide and directly influence each other in critical ways would be essential to achieve more extensive forms of personnel-based requirements traceability. This the subject of future work.

	Production roles	**Reception roles**
Participant roles	Those agents **directly** involved in producing artifacts in requirements engineering	Those agents who make use of the artifacts produced in requirements engineering for whom they **have** been explicitly produced
Non-participant roles	Those agents **indirectly** involved in producing artifacts in requirements engineering	Those agents who make use of the artifacts produced in requirements engineering for whom they **have not** been explicitly produced

Figure 3: *Partitioning the social structure underlying the requirements engineering process, as derived from [12].*

The main steps of our approach to model contribution structures are illustrated in Figure 4 and summarised below. Fuller details and examples, along with a model-based specification

of the approach, its operation, and how it achieves personnel-based requirements traceability, are given in [10].

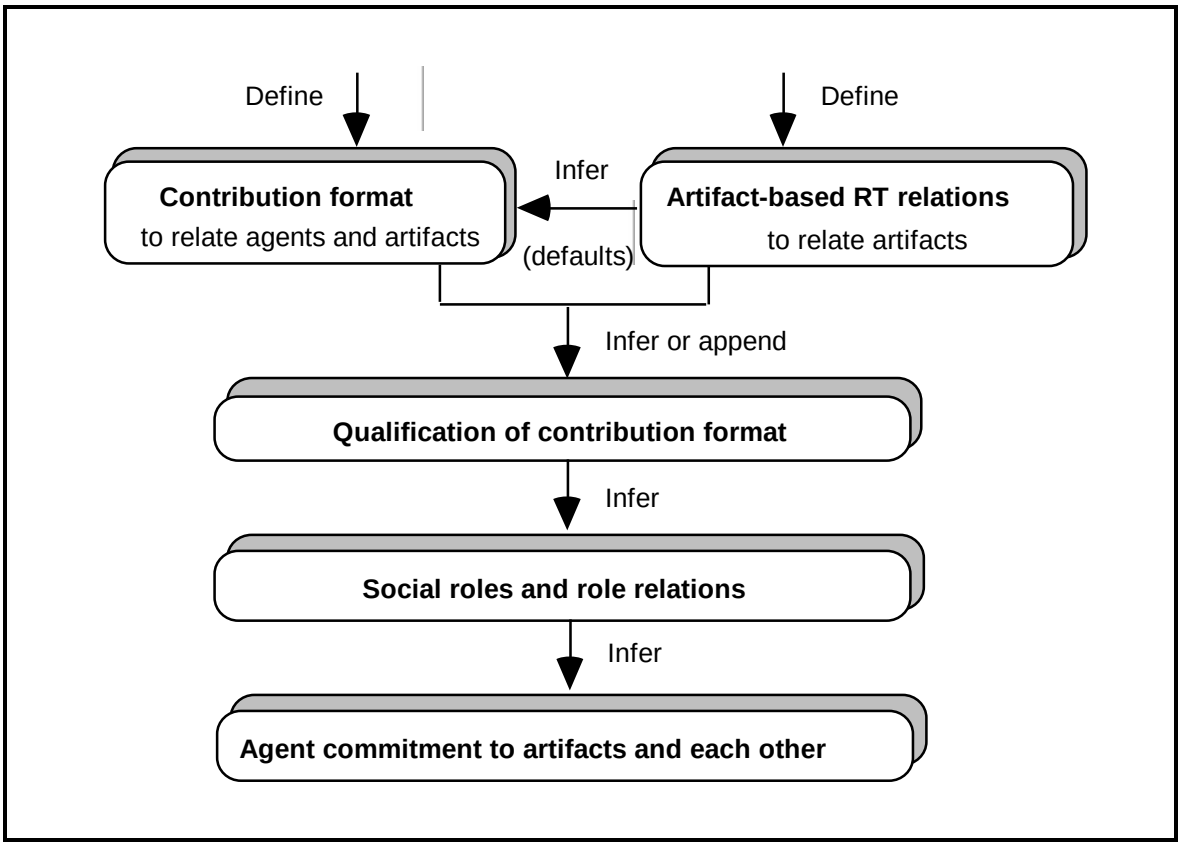

Figure 4: *The main steps of the approach. The approach involves minimal extensions to conventional artifact-based forms of requirements traceability to augment their traces with contribution structures. These extensions take the form of the semantic classification of different types of artifact-based relation and the addition of contribution relations between agents and artifacts. This rudimentary input information can then be used to provide additional attributes about the contributions and contributors, and also to infer details about social roles, role relations, and commitments.*

3.1 Contribution format and qualification

If the relation between agents and artifacts were defined using a singular type of relation depicting the notion of "contribution", this would lead to a flat and coarse model of the contribution structure, as shown in Figure 5. A relation which does not distinguish

different types and degrees of contribution would result in a model which offers little real ability to reason about the contributors and their contributions. This would therefore only provide support for crude forms of personnel-based requirements traceability. Instead, how this relation is described needs to reflect the fact that many agents may be party to the production of an artifact, and further needs to differentiate the various ways in which they have each contributed. In addition, the scheme chosen to delineate the nature of the contribution relations needs to supply the building blocks with which to model a progressively more granular and layered contribution structure as requirements evolve and are changed by others.

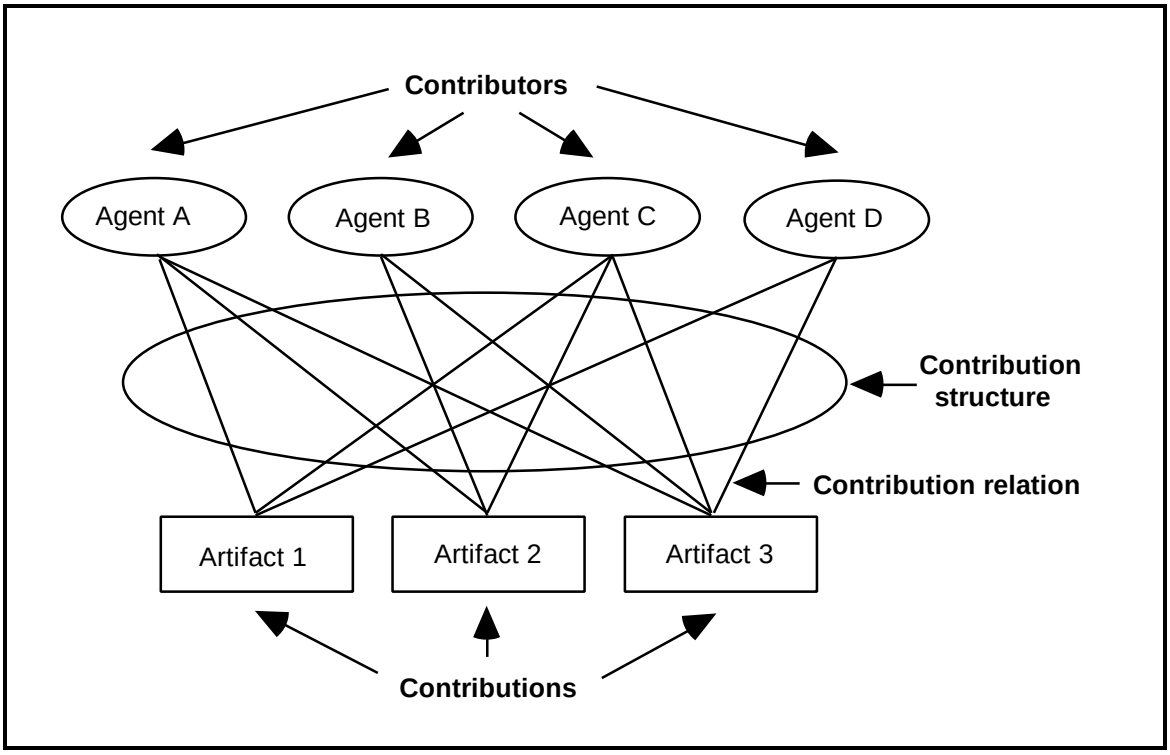

Figure 5: *A flat and coarse model of the contribution structure. Such a network of interrelations could be determined as a consequence of listing all the so-called "authors" associated with the formal documents produced in the requirements engineering process.*

We therefore use the concept of a *contribution format* to delineate three fundamental capacities in which agents can contribute to artifacts, either as their *principal*, their *author*, and/or their *documentor*, as illustrated in Figure 6. This scheme is derived from Goffman's

notion of a "production format" which he uses to deconstruct the broader notion of "speaker" into its underlying constituents [12]. The principal is the agent whose position and/or belief is established by the information in the artifact, so is the agent who motivated its production, is committed to what it expresses, and is responsible for its effect or consequences. The author is the agent who put together and organised the information expressed in the artifact, so is the agent responsible for its content and structure. The documentor is the agent who recorded or transcribed the data present in the artifact, so is the agent responsible for its physical and presentational aspects.

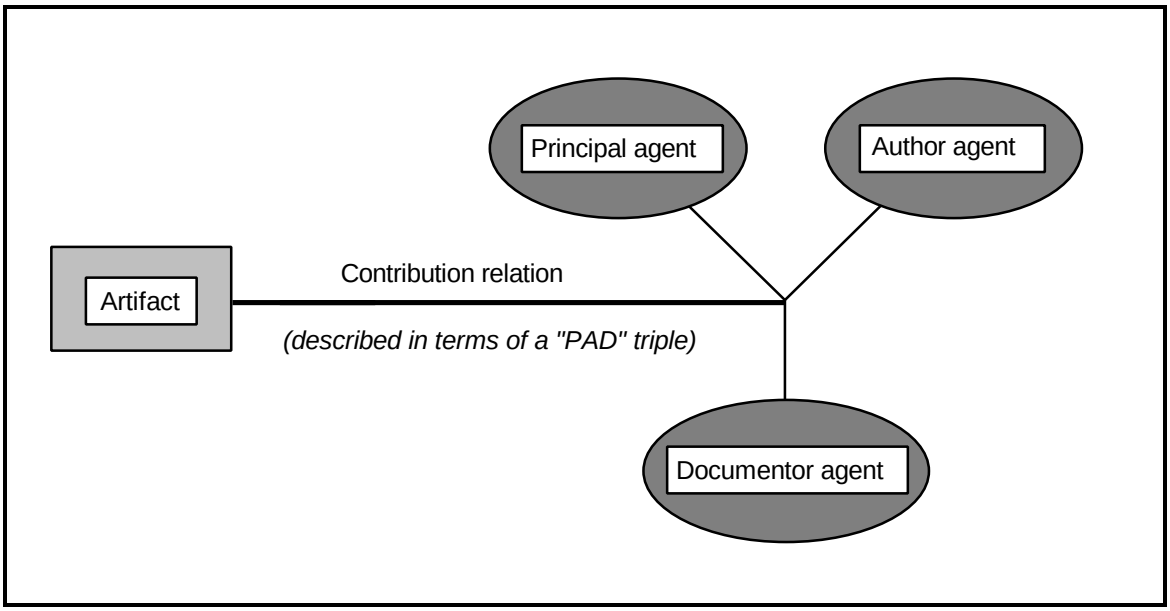

Figure 6: *The three capacities which account for the different ways in which agents can directly contribute to requirements artifacts. Together, these clarify the broader notion of "contribution". Note that the basic capacities in which agents contribute could be determined and labelled to account for different projects and organisational settings. The scheme we have described is only provided as an example.*

By discriminating more than simply a relation of contribution we can analyse the different types of contribution made by agents. The additional structure provided by specifying the contribution format of an artifact, as shown in Figure 7, begins to provide the basis for more selective forms of personnel-based requirements traceability. For example, the

principal contribution relations could be used to retrieve all the principals in an artifact-based trace of a requirement's production, and further be used to uncover information about the structure of authority and power in the process. A comparison of the different types of contribution relation attached to each agent could be used to find out about the various ways in which different agents, say with the same job description, contribute throughout the process and how this changes over time.

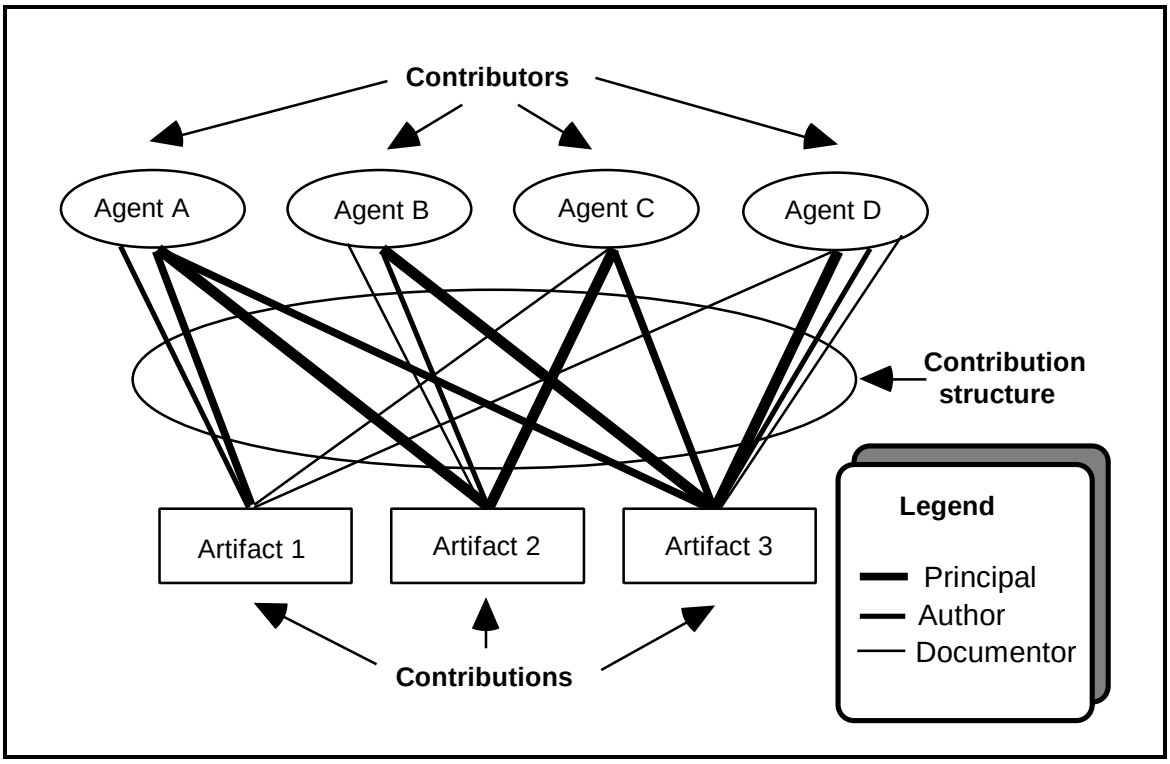

Figure 7: *Using the three capacities of the contribution format to organise the model of the contribution structure. This model therefore discriminates the basic types of contribution and exhibits some structure. For Artifact 1, Agent A is its principal, Agent A is its author, and Agents C and D are its documentors. For Artifact 2, Agents A and C are its principals, Agent B is its author, and Agent B is its documentor. For Artifact 3, Agents A, B, C and D are its principals, Agent D is its author, and Agent D is its documentor.*

With more information about the nature of the contributions in each of the three capacities, we can describe more intricate contribution structures and, in turn, enhance the personnel-based requirements traceability provided. Although numerous sets of attributes could be

suggested for qualifying each of the capacities, the most useful are likely to be those that can be automatically obtained as a by-product of the approach and the basic information it collects. For instance, a signatory-related set of attributes can be used to qualify each principal contribution relation. Details about whether an artifact has been approved, not approved, or is pending approval, can be identified automatically if formal procedures are in place to circulate and sign-off requirements artifacts. Such qualification can help to identify those points in a project where requirements have become stabilised or rejected, under who's authority, and further used to highlight those agents who frequently contribute at such key stages.

3.2 Social roles, role relations and commitments

When an agent contributes to an artifact in one of the above capacities, they also act in a *social role* with respect to this artifact and its other contributing agents. So, although it may be clear who the documentor of a requirements artifact is, whether they are documenting for themselves or on behalf of others, and how this changes with the development of the artifact concerned, is not immediately clear. Following Levinson's extensions to Goffman's work, described in [13], we therefore distinguish between the basic capacities in which agents can contribute and the derived social roles that can be determined from these. This distinction is exemplified in Figure 8.

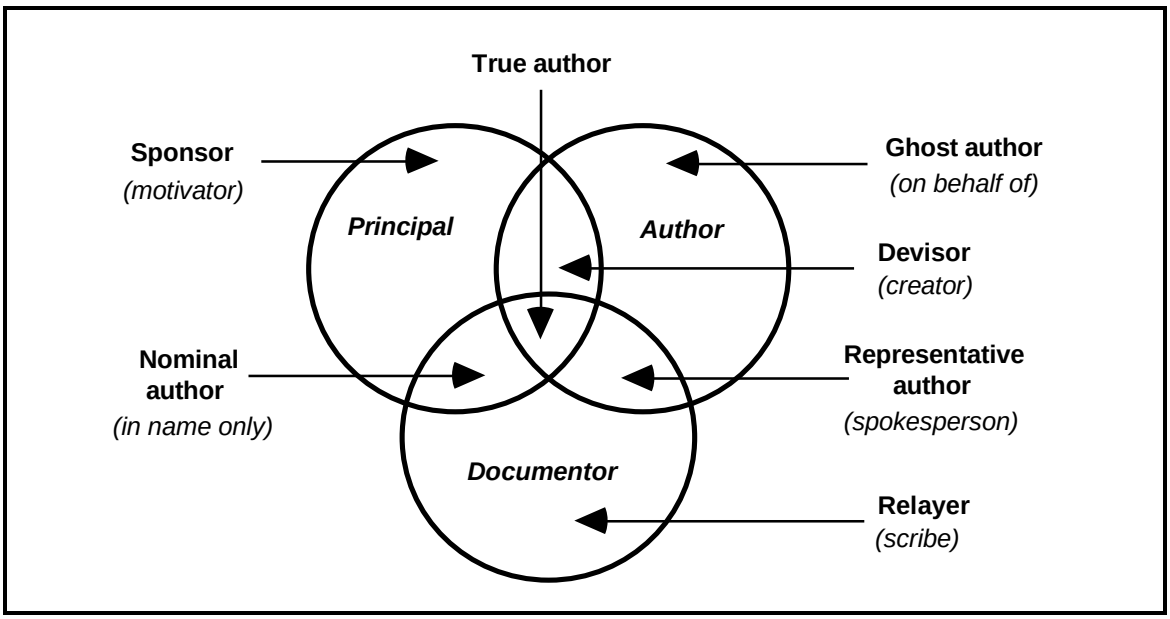

Figure 8: *An example set of social roles, adapted from [13]. Note that the labels that we give to these social roles are meant to be illustrative as opposed to definitive. The arrangement and categorisation could be adapted to correspond with those commonly used in different projects and organisations.*

Although there have been many suggestions as to the kinds of role that need to be assumed by practitioners in systems and software development, such roles tend to be institutional, coarse-grained, prescriptive, and statically assigned. In contrast, social roles are defined relative to specific social situations and relational ties to other agents, so are assigned interactionally and dynamically. They are somewhat akin to Banton's notion of *transient roles* as they can adapt to account for the changing alignment of agents to artifacts, as well as to each other, throughout the development process [14].

An analysis of the social roles assumed in the requirements engineering process provides a way to get a handle on the social relations developed and sustained between the contributing agents and is crucial for developing the model of the contribution structure and the forms of personnel-based requirements traceability it makes possible. Furthermore, the relations that exist between agents when assuming these social roles provides information about the *role relation*s that have been formed in practice. The different types of role

relation that are prominent in a project tell us something about how the agents involved work together. The different agents that assume the social roles in these role relations tells us something about which specific agents work together and in what ways. Such information is crucial for understanding the workings of the underlying contribution structure since, although the same two agents can interact together in many social roles, what is done in one role relation often has repercussions in many others.

Social roles provide an even more effective way to deal with discriminating the types of contribution made by agents, as shown in Figure 9. This is because they characterise the total nature of an agent's contribution with respect to an artifact, as well as to the other agents contributing to the same artifact, so provides the basis for increasingly granular and selective forms of personnel-based requirements traceability. For example, the social roles could be used to retrieve those artifacts in an artifact-based trace of a requirement's production in which selective agents have contributed in a devisor/relayer role relation, and further be used to signal any artifacts in this trace where this role relation has been reversed.

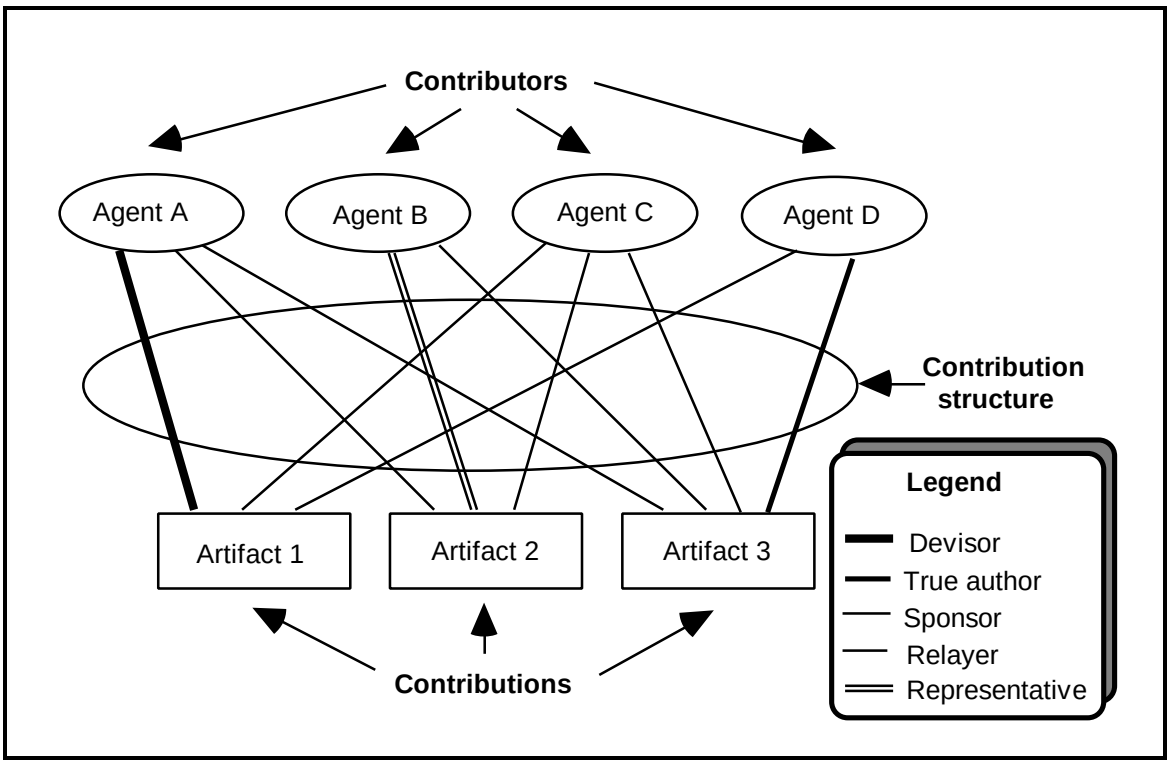

Figure 9: *Using the seven social roles to organise the model of the contribution structure. This model therefore discriminates both the basic and derived types of contribution and exhibits an even finer structure. For Artifact 1, Agent A is its devisor, and Agents C and D are its relayers. For Artifact 2, Agents A and C are its sponsors, and Agent B is its representative. For Artifact 3, Agents A, B, C and D are its sponsors, and Agent C is its true author.*

In addition, knowledge of an agent's social role with respect to an artifact tells us about those aspects of the artifact that they can be called to account for, so about *individual* and *collective commitment*. Such information is useful for filtering the retrieval of agent sources to reflect particular types of query or change proposal. It can be used to indicate which agents to involve or inform about changes to certain aspects of a requirement. It can also help to directly locate the primary source of motive behind an artifact or subsequent sources of its changing content. This makes the contact point for specific aspects of an artifact explicit and traceable as the artifact evolves. The role relations that arise as a consequence of agents jointly contributing to an artifact's production tells us about the ensuing *social commitment* that is formed between agents, something that is rarely captured

by formal organisational structures and pre-assigned project roles. The type and intensity of these role relations, and how they vary with respect to different artifacts or over time throughout a project, can thereby provide useful material for the analysis of informal working arrangements.

3.3 Artifact-based requirements traceability relations

If one requirements artifact is a subsequent specialisation of another, it seems reasonable to assume that some responsibility for the resulting artifact will often be retained by the original contributor(s). Our approach therefore needs to account for the relations that exist within and between the artifacts themselves when modelling the contribution structure, particularly if it is to distinguish original artifacts from ones that have been duplicated or derived from others, to account for the linked and embedded nature of contributions.

We suggest that there are three broad categories of artifact-based relation. These describe alternative artifact-based structures, as suggested in Figure 10, and provide the basis for different types of artifact-based requirements traceability. *Temporal relations* reflect the chronological order in which requirements artifacts have been produced and provide the means to trace requirements history. *Developmental relations* reflect the logical order in which requirements artifacts have been produced and provide the means to trace requirements flow-down. *Auxiliary relations* reflect many additional types of order between requirements artifacts and provide supplementary forms of requirements traceability. In this paper, we are mainly concerned with the relations of the third category. We suggest that it is these relations that have the most subtle, though crucial, effects on the contribution structure. Also, as there is no well established set of such relations in use for requirements traceability purposes, we can propose a set of auxiliary relations which subsumes the first two categories that are generally used to provide conventional forms of requirements traceability. The two types of auxiliary relation we are most concerned with are what we call *containment* relations and *connectivity* relations.

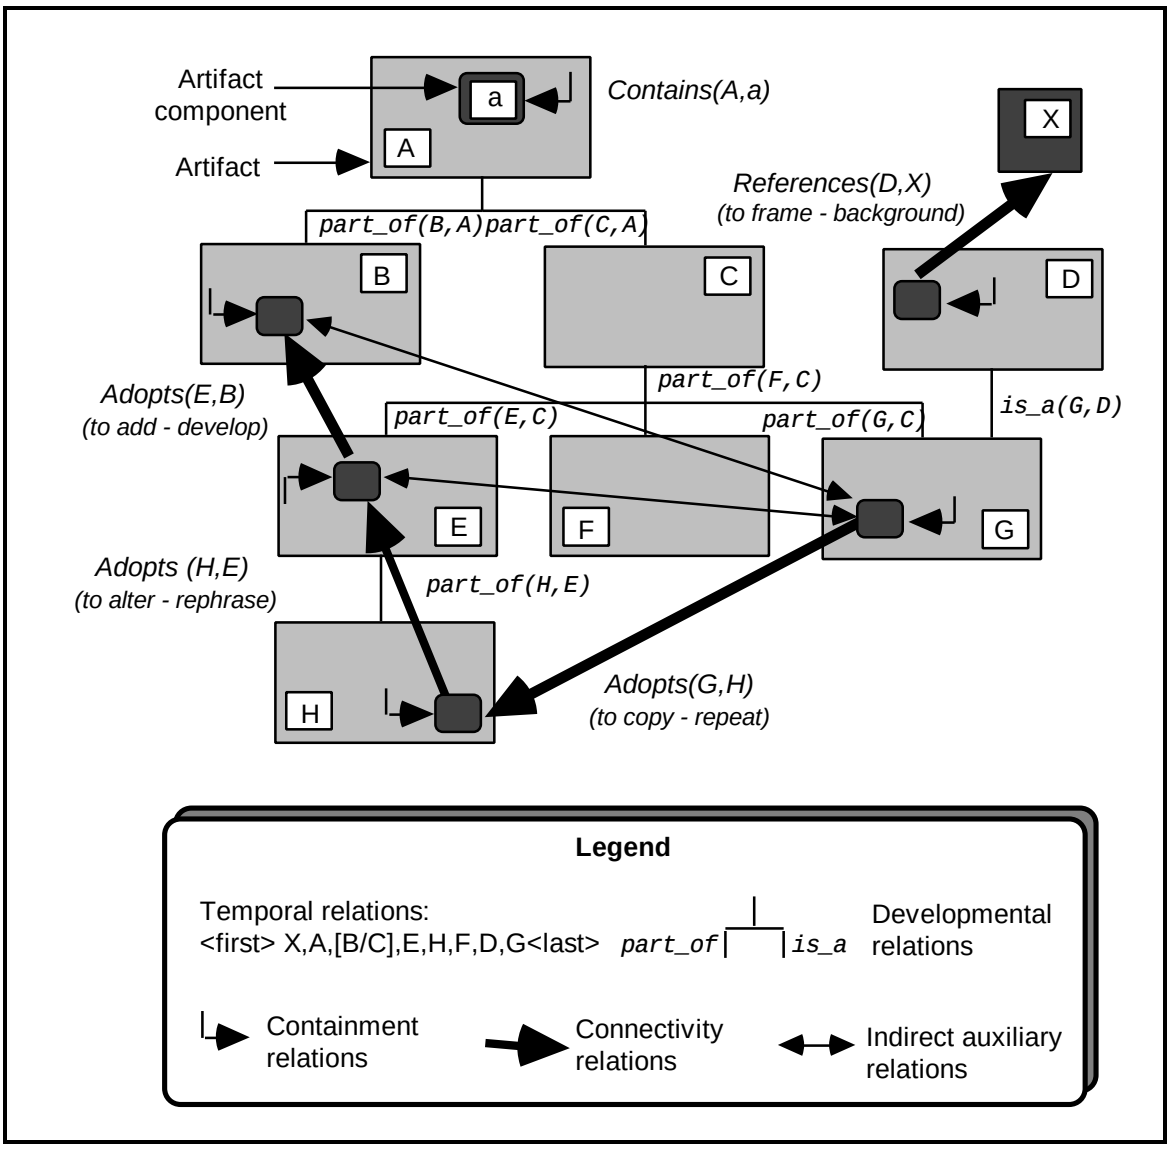

Figure 10: *Different artifact-based structures and forms of artifact-based requirements traceability.*

Containment relations delimit the internal components of composite artifacts. They thereby make the task of assigning the contribution format much easier. Though clearly a composite artifact may have different agents acting in identical capacities with respect to its components, it can be a default assumption that they are the same until declared otherwise. These relations lead to a layering effect and enable multiple contribution formats to be defined, interrelated, and managed. Through their use, areas of contribution can become more finely delineated over time as changes are made by various agents and as artifacts are versioned, partitioned, or used elsewhere.

A useful set of connectivity relations is provided by describing the broad *communicative function* that a connectivity relation serves between two artifacts. A connectivity relation can function either to *reference* or to *adopt* the content of another artifact, and this distinction is expanded upon in Figure 11. Those relations which function to "reference" exist when the physical content of the source and target artifacts does not overlap. Here, information in the target is not integral to information in the source, but is either subordinate, superordinate, or coordinate to it. The referenced artifact may be explicitly signalled in the source, perhaps by cross references, keywords, or synonyms, else is implicitly signalled. Those relations which function to "adopt" exist when the physical content of the source and target artifact overlaps in some way. Here, information in the target is integrated into information in the source, either exactly, inexactly, in full, or in part. The information in the source is also either a static or dynamic version of that in the target. Such distinctions, if taken into account when modelling the contribution structure, signal where responsibility for future maintenance and decisions about requirements artifact lie. Accordingly, they can be used to filter the types of personnel-based traces to be retrieved, like all the artifacts in an artifact-based trace of one requirement's production trail which have been referenced by a specific agent in another's.

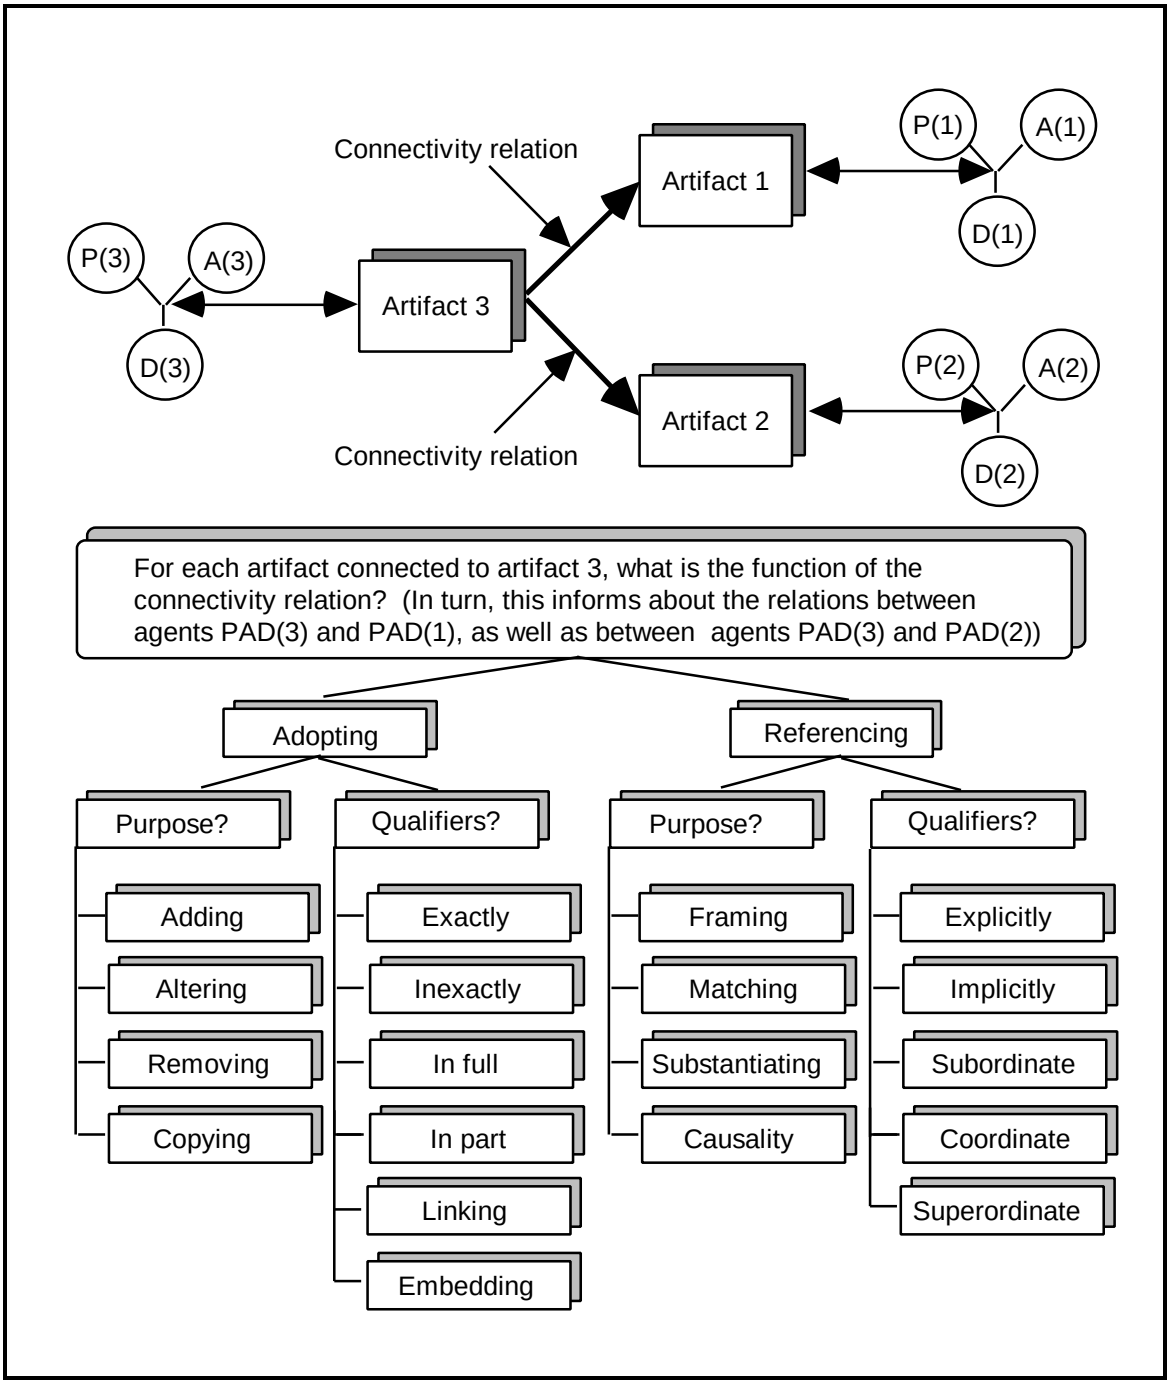

Figure 11: *Decomposing the nature of connectivity relations. For instance, a references relation exists between two requirements artifacts when the purpose of the connectivity relation is either: to frame; to match; to substantiate; or to show causality. Furthermore, a relation that serves to substantiate exists when the information in the target illustrates or supports that in the source. We therefore use the underlying coherence and cohesion relations that exist amongst the requirements artifacts to help determine the nature of connectivity relations.*

4 What the approach should deliver

The primary objective of modelling the contribution structure underlying requirements is to make it possible to trace the human sources of requirements, requirements-related information, and requirements-related work, by using personnel-based requirements traceability in conjunction with more conventional forms of artifact-based requirements traceability. In Section 2.3, we provided some examples of the kinds of question that personnel-based requirements traceability is expected to provide answers to. These were currently found to be problematic to answer in practice and, where left unanswered, were found to eventually lead to claims of requirements traceability problems. We briefly describe how such questions can be readily addressed following application of the approach.

4.1 Involvement

> "*Who has been involved in the production of this requirement, how, and on what other requirements have the same agents been involved?*"

To identify all those agents who have been involved in the production of a requirement, we would need to trace the full chain of contribution behind its production process. This can be done using the artifact-based requirements traceability relations to trace the requirement back to its origin and by extracting the corresponding contribution formats along the way. The ultimate origin of a requirement would be reached when it is not possible to trace back to any other artifacts through further artifact-based requirements traceability relations, though a more relaxed origin would be one which has no further outgoing "adopts" connectivity relations, even though it might reference other artifacts for background information. As a requirement may have multiple origins, it would be essential to trace all the contributing branches and extract the contribution formats from all the relevant artifacts. Dependent on the level of specificity required, we can carry out this personnel-

based requirements traceability at various levels of detail, like selectively tracing the chain of principals, authors, or documentors involved in the production of a requirement. We can identify whether particular agents have always contributed as documentors or principals throughout a requirement's production. Furthermore, with a record of all the contributors and the manner in which they have contributed, we can compare the contributor profiles behind different requirement production trails to uncover whether certain agents consistently supply requirements or push them through the requirements engineering process.

4.2 Working arrangements

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

To identify changes or trends in the working arrangements of agents involved in a requirement's production, by which we mean the alignments of those agents that jointly contribute to the artifacts concerned, again we would need to make use of the full chain of contribution behind its production process. This question can be answered, not only by examining changes with respect to those agents assuming the specific capacities of the contribution format for each artifact in the production trail, but more finely by examining changes in the social roles and role relations of these contributors. We can identify whether a particular agent has consistently been a sole ghost author on behalf of a specified other, which agents only contribute when they collaborate with certain others, whether the role relations between all involved stabilise or fluctuate with a requirement's production, and determine who was originally responsible for different aspects of the requirement and examine how this responsibility has changed hands with its subsequent development. In addition, we can assess the impact of changing working arrangements in the requirements engineering process on various attributes of the resultant requirements specification, like its quality, amount of technical content, and so forth.

4.3 Ramifications

> "*What are the ramifications, with respect to the loss of possible requirements-related knowledge, if a specific individual or group leaves a project? Who would be the best back-up source of information here?*"

To assess the potential effects of staff turnover, we would need to identify all the contributions of the agent intending to leave, including details about the capacities in which they contributed and their social relations with respect to other agents. With such information, it is then possible to examine whether there are agents who have directly collaborated with the agent in question, so who are likely to possess adequate experiential knowledge to that which may be lost. Otherwise, it is possible to assess the most suitable alternative agents to take over the agent in question's outstanding commitments. Through the artifact-based requirements traceability relations, we can also identify those agents who's contributions the agent in question has made use of, revealing possible social contacts for the agent who is to take over. In the same way, we can identify those agents who make use of the agent in question's contributions, so identify those agents who may have outstanding requirements for communication and future contact with the agent who is to take over. With the visibility of all such information, we would be in a more informed position from which to assess the amount of up-heaval that would be caused by staff turnover, to identify those areas where a replacement agent would need de-briefing and with whom, and thereby to determine how best to smooth-over the transition process.

5 Demonstration of the approach

There are a number of levels through which we have demonstrated the approach. Firstly, we developed a model-based specification of the approach to show the feasibility of its underlying concepts, a detailed requirements specification to outline what is needed to

implement the approach, and a prototype tool to illustrate how these requirements and their underpinning specification could be met in a practical way. Using this tool, we then carried out use case scenarios to examine how the approach could be applied, and also to assess the potential of its personnel-based requirements traceability. Finally, we conducted a case study using data from a real industrial requirements exercise to uncover the benefits, drawbacks, and other issues arising from the application of the approach in a practical setting. Each of these is summarised below, thought the reader is referred to [10] for more details.

5.1 Tool support

We have developed a prototype tool in which conventional forms of artifact-based requirements traceability can be extended with associated contribution structures. A schematic of the tool is given in Figure 12. Requirements artifacts are held in a marked up form in an on-line project repository which manages the artifact-based requirements traceability relations we identified in Section 3.3. The project repository also maintains any information provided about the agents participating in a project.

The *traceability extension tool* provides a hypertextual interface to the information held in the project repository and enables details about artifacts, agents, and their various interrelations to be interactively described for a project. Such details are recorded in the descriptive markup of the artifacts concerned. It provides a way to extract that information required to model the contribution structure from this markup and places it in the factbase of the contribution structure manager. It also provides facilities to interrogate the current state of the contribution structure, defined by the information held in this factbase, and to display the results. The traceability extension tool therefore acts mainly as the medium through which the contribution structure can be defined, maintained, and used.

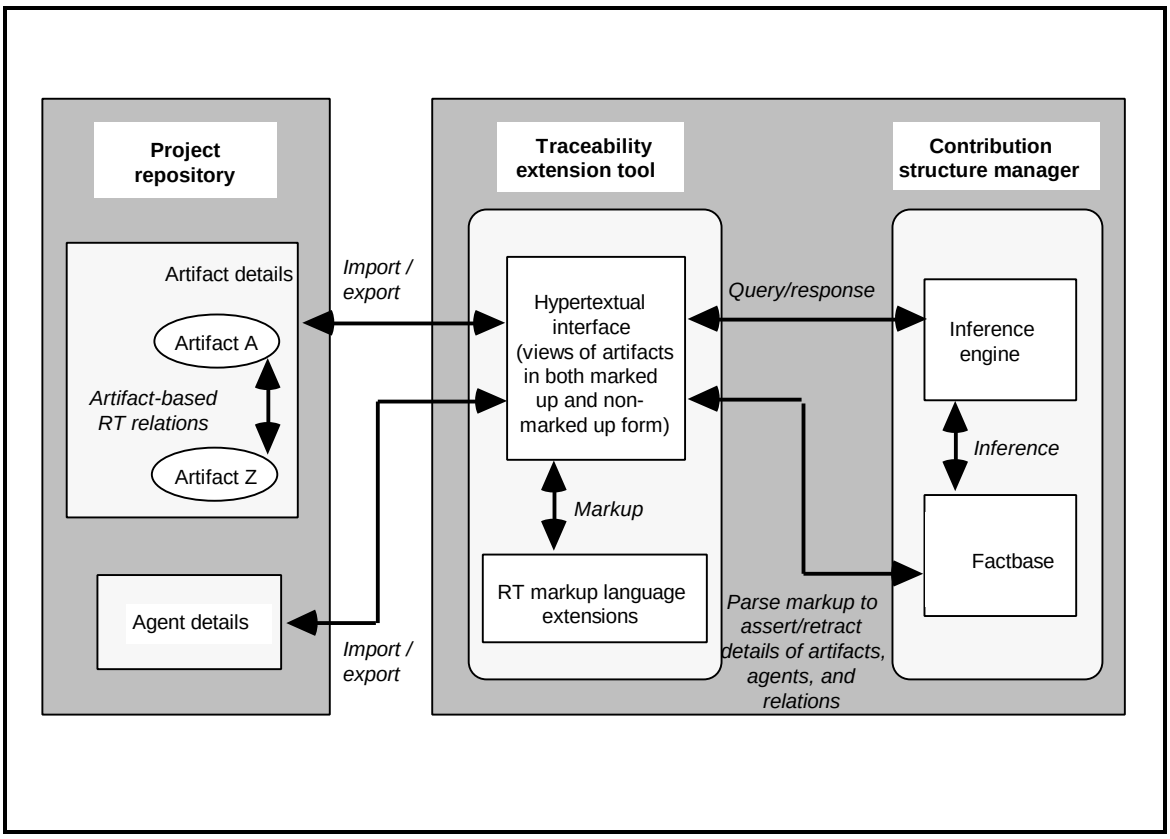

Figure 12: *The architecture of the tool.*

The *contribution structure manager* stores the information described by the markup of the artifacts registered to a project within its factbase. This includes information about the project's agents, artifacts, artifact-based relations, and contribution relations. It further contains rules that use this information to infer details regarding default agent capacities, the social roles of agents, and their commitments. It implements the model-based specification we mentioned earlier and is therefore responsible for defining and maintaining an up-to-date model of the contribution structure. It also implements those operations which provide the ability to query and make use of this model of the contribution structure to achieve diverse forms of personnel-based requirements traceability.

Our experience with building this prototype indicates that the approach could readily be made operational by configuring general-purpose tools to support it or by extensions to commonly used document preparation and database systems. Moreover, where tools to

support requirements traceability are currently used within an organisation, the approach could be supported by minimal extensions to the artifact-based requirements traceability schemes that such tools are typically configured to support.

5.2 Scenarios

We have carried out use case scenarios to show how the approach can be applied through use of the tool, to illustrate the personnel-based requirements traceability, and to see how such extensions can complement conventional forms of artifact-based requirements traceability and inform practice. One such scenario is given in [11]. This highlights the value-added information that can be uncovered about the social dimension underlying requirements production and thereby demonstrates the insight that can be brought to the scenario as a result of modelling its underlying contribution structure.

Our experience with scenarios indicates that, from the provision of rudimentary input data, the approach can be used to construct a rich picture of a potentially complex contribution structure underlying any artifact production process. The added ability to visualise and interrogate this contribution structure can further make explicit information that would otherwise remain hidden or be incorrectly inferred.

5.3 Case Study

We have conducted a case study based on a real, albeit small-scale, industrial requirements engineering project. The project comes from a small company which is in the business of providing both software and procedural solutions to many communications-related problems. The project was initiated in 1992 and, although the intended service is now operational, its requirements continue to evolve to date. It has involved roughly sixty agents and produced over one hundred and sixty six artifacts. In conjunction with practitioners, we applied the approach in a post-hoc manner to all the data that had been collected. We then examined whether our claims as to what the approach should deliver

could actually be met. We found that we were not only able to achieve extensive forms of personnel-based requirements traceability as predicted, but that we were also able to uncover all sorts of additional information about the requirements engineering process and its participants.

Members of the company agreed that the data we revealed about the contribution structure underlying the project rang true and that it would have pointed to the right agents for further contact where problems of misunderstanding and the like surfaced. In particular, they were surprised to see how extensive some of the trails of contributors behind eventual requirements had actually been, and recognised how critical it could be to retrieve such information when considering different types of change. The approach immediately highlighted those problems which arose as a direct consequence of early contributions in this project, especially those related to the contributions of non-central stakeholders and changing team members, which were problems that were only recognised much later on in the project. It further provided information about agent social roles and role relations that could not have been determined from the company's organisational chart or from the project's work allocation timetables. It was considered important to compare this information with such managerial devices to inform how work could be better allocated amongst agents in future projects. Even for this company, a company that was not overly concerned with maintaining requirements traceability, the approach was considered a feasible addition to practice. The personnel-based requirements traceability and the value-added information it yielded was also considered likely to impact practice in many positive ways.

The case study also served to highlight outstanding issues with the approach and the information it provides. The most central of these included: the need to account for the different degrees in which agents contribute; the need to indicate the topics of requirements artifacts to enable more filtered forms of personnel-based requirements traceability; the need to capture information about the undocumented events leading up to an artifact and

account for these in the approach; the sensitivity of the information the approach deals with and the need to exhibit care in its analysis and generalisation; and the time that could be required to both analyse the data provided and act upon it. Such issues point to areas for future work.

6 Evaluation of the approach

As a basis for our evaluation, we have adopted some of the challenges originally set out by Grudin as criteria to examine [15]. One point that is important to keep in mind when we evaluate the approach is that the manner in which the artifact-based requirements traceability relations and contribution relations are to be obtained remains something that can be determined on a project-specific basis. This means that application of the approach could be the distributed responsibility of all the agents involved in a project, the responsibility of a particular agent contributing to each of the artifacts produced, or the responsibility of a dedicated individual or group. Therefore, where organisational policies are already in place as to how artifact-based requirements traceability is to be established and maintained, we anticipate that our extensions would be handled in the same manner.

6.1 Balance between the work involved and benefits reaped

Whatever policy is chosen for collecting the information the approach requires, there is likely to be a disparity between the work involved by the providers and the benefits reaped by the end-users, since not all providers would also be end-users. We have suggested ways to reduce the manual effort necessary to collect the required information, like the inference of default contributors from containment relations, but this is unlikely to eliminate the need for agent involvement. The potential disparity is probably best handled by introducing the dedicated job of a requirements traceability facilitator or by extending an existing job description to cover the responsibilities of provision. For such an individual or group, the work involved would no doubt outweigh the personal benefits reaped from personnel-based requirements traceability, though bring benefits to the larger collective. Those agents that

would now be in a position to perform personnel-based requirements traceability to assist their working practices, those involved in the various phases of a project, would be the primary and direct beneficiaries of the approach. However, as the need for personnel-based requirements traceability is likely to be low at the onset of a project and steadily increase as it progresses, those agents involved early on are likely to get fewer immediate benefits and would be making potentially more work for themselves later on if they were to cooperate. To address such issues, policies would need to be put in place to stress how the information the approach provides is to be used, for how long, and so forth. Project leaders and managers would be indirect beneficiaries of the information the approach provides, since they would be in a position to learn something about the social dimensions underlying their projects. It is exactly these agents that would do little, if any of the work, and would potentially reap lots of value-added information.

6.2 Dependence on a critical mass

Were provision of the information the approach requires the joint responsibility of all involved, or the responsibility of one of the agents in the contribution format of each artifact, the approach would still enable a minimal form of personnel-based requirements traceability if only some of these agents complied. Although this would lead to missing contributor details at periods throughout a requirement's life, the minimal information could still guide the search process and so not cause the approach to fail totally in its main objective. However, some of the value-added information provided would not be so reliable. As mentioned above, cooperation would be essential in the early phases of a project to prepare the groundwork for later personnel-based requirements traceability but, unless more immediate benefits were provided to these agents, the approach could be abandoned by the majority before the real benefits were reaped. Even if provision of this information were the responsibility of a dedicated individual or group, these agents would still depend on some cooperation from all those involved to gather the right information in a timely manner. This problem could be addressed by introducing a project "champion"

charged to ensure support for the approach in the early phases. In contrast, few agents need to make use of the personnel-based requirements traceability that the approach makes possible for it to have been worthwhile. This is because the time saved in identifying those agents that need to be involved in even one change proposal could quite easily be offset against the time taken to establish the potential for this traceability in the first place.

6.3 Handling of exceptions

The schemes we have suggested for the artifact-based requirements traceability relations and the capacities of the contribution format are by no means definitive. These could be tailored on a project-specific basis to account for existing requirements traceability practices and any current distinctions used in documentary practices. Use of the approach could further fit in with the work schedules of agents, since no constraints have been imposed on the order in which the information needs to be collected. Even if certain information is omitted, noting a solitary contributor would lead to some form of personnel-based requirements traceability, and later additions and changes could always be made without jeopardising the approach. However, the approach does not guarantee that only complete and accurate information will ever be recorded, though the information required could be subject to some checking as it is input to assist with this issue. With more experience in the application of the approach in practice, we could examine the types of exception most likely to occur when entering the information it requires and when requesting various forms of personnel-based requirements traceability, and so assess how these contingencies could best be accounted for.

6.4 Disruption to the status quo

The approach does not stop agents carrying out personnel-based requirements traceability in the current informal manner. If a subtle reason exists for this practice, perhaps because agents can learn all sorts of extra information by attempting to manually trace those who have been involved, this can still happen where needed. The approach just serves to guide

this practice and, in making the contribution structure visible, also helps to legitimise the artifacts produced. Agendas which need to be hidden for political reasons can therefore remain so, whilst traceability back to contact points for such agendas would at least ensure their involvement in subsequent changes and the like. Whether or not the application of the approach would disrupt existing working practices in an organisation would firstly depend upon whether requirements traceability was already implemented. It would also depend upon whether the basic tenet behind the approach complied with the prevailing organisational culture. Since the approach depends upon agents being open about their contributions, it could be undermined if accountability for one's contributions were used to attribute blame, rather than used to encourage process improvement. The range of information that could be inferred as a by-product of the approach could also be very sensitive in some organisational settings because, although the approach itself does not undermine organisational structures, subversions could be revealed through the information it provides. Furthermore, introducing the potential for personnel-based requirements traceability may upset existing practices, mainly because it would mean that agents remain active in a project for longer than they might have originally anticipated. To address such issues again suggests the need to put policies in place to govern the use of the approach and its results.

6.5 Accessibility from mainstream applications

Organisations do not require a specialised tool to apply the approach and use its results. Therefore, it need not obscure what are considered to be the mainstream activities in systems and software development. As most requirements artifacts are either textual, graphical, or compound in nature, the approach could be implemented as extensions to existing document preparation tools in conjunction with some form of underlying database system. With the introduction of open technologies which allow compound artifacts to be created using different applications, individual preferences would still be supported with regard to the applications used for producing them. As the approach could also be

implemented through simple extensions to document management systems, or to contemporary requirements traceability tools, again it need not obscure the use of the other tools ordinarily used to carry out development activities. Furthermore, integration with those tools which support the communication likely to be required as a consequence of the personnel-based requirements traceability provided would, for example, enable tailored messages to be sent to those agents involved in the production of a requirement when proposing to change it.

6.6 User involvement

The approach is grounded in an understanding of requirements traceability problems and was derived from a specification of requirements to alleviate these. These early phases of our own development process involved many potential users. Although potential users were not directly involved when subsequently developing the approach, we periodically received comments, suggestions, and anecdotes from those that could be considered potential users. However, we were never in a position to examine, from first-hand experience, what actually happens when requirements traceability breaks down in projects. This means that we were unable to determine information like: whether the reason for needing to find the agents involved in a requirement's life influenced how they were identified and located; what information about those involved ordinarily needed to be uncovered; and what generally took place once these agents had been identified and located. Instead, tailoring the basic approach to a specific set of users and their working practices is now the subject of our longer-term action research with practitioners from within a specific organisation.

6.7 Managing uptake and acceptance

We have identified some guidelines for managing the uptake of the approach and for encouraging its continued acceptance. For instance, as we have noted that existing development practices and organisational cultures are important factors that are likely to

influence the uptake of the approach, this implies that the approach is more likely to be taken up in those organisations where an emphasis is increasingly being placed on requirements traceability to improve development processes or where project management concerns are turning towards the personnel involved in projects. It is also more likely to be taken up if it does not require specialised tools and training, so can be implemented as an addition to existing applications and working practices. Policies governing how the information is to be gathered, and how the information it provides is to be used, would be important to establish at the onset of a project to deal with user concerns and so encourage acceptance. Finally, acceptance would be more likely if the approach were first used to address the most pressing problems within particular projects and organisations, so some significant benefits could be immediately demonstrated, before extending its scope to cover more of what we have mentioned. Phasing in the approach would further enable those factors which are likely to promote and hinder its acceptance in each particular organisation or project to be uncovered and taken into account.

6.8 Evaluation and generalisation

As we have not quantitatively evaluated aspects like the usability of the proposed implementation of the approach and the usefulness of the personnel-based requirements traceability it provides, having obtained subjective opinions on such matters from practitioners involved in the case study, we cannot generalise from our findings. Each application of the approach is likely to have some dependence on the background and priorities of the individual organisations, projects, and agents involved. However, we have demonstrated the feasibility of the approach, in that it addresses a real problem in a practical way. We have also identified those factors which are crucial to its success and those factors which are most likely to cause it problems.

7 Summary

In this paper, we have explained why it is important to retain information about those individuals and groups from whom requirements have been elicited and amongst whom they have subsequently been refined. This is because people are the ultimate sources underlying and driving the production of requirements and are frequently essential to revisit to assist with their subsequent elaboration and eventual evaluation. We have introduced the notion of "contribution structure" to anchor requirements artifacts in the social structure that gave rise to them and provide the most fundamental baseline to support on-going requirements elicitation and analysis. We have described an approach to define, maintain, and make use of this model, and demonstrated how it enables conventional forms of artifact-based requirements traceability to be extended with personnel-based requirements traceability. This makes it possible to selectively and rapidly identify the most appropriate agents to involve and inform in situations of change, decision making, requirements query, quality assessment, conflict resolution, and so forth. We have further pointed to the sorts of question about working structures and organisational dynamics that can be answered as a by-product of the approach. Finally, we have provided a critical evaluation of the approach and what it delivers, which also highlights implications for practice and areas for future research.

8 Acknowledgements

The authors acknowledge the comments and assistance of colleagues. The first author would also like to thank David Michael for many helpful discussions and the States of Jersey for financial support whilst the ideas in this paper were being formed.

9 References

[1] GOGUEN, J. A. and LINDE, C.: 'Techniques for Requirements Elicitation'. Proc. IEEE Intl. Symp. on Requirements Engineering, San Diego, California, January 1993

[2] SHAW, M. L. G. and GAINES, B. R.: 'Requirements Acquisition'. Papers of a Workshop on Requirements Elicitation for Software-Based Systems, University of Keele, England, July 1994

[3] GOUGH, P. A., FODEMSKI, F. T., HIGGINS, S. A. and RAY, S. J.: 'Scenarios - an Industrial Case Study and Hypermedia Enhancements'. Proc 2nd IEEE Intl. Symp. on Requirements Engineering, York, England, March 1995

[4] HUGHES, J., O'BRIEN, J., RODDEN, T., ROUNCEFIELD, M. and SOMMERVILLE, I.: 'Presenting Ethnography in the Requirements Process'. Proc 2nd IEEE Intl. Symp. on Requirements Engineering, York, England, March 1995

[5] NUSEIBEH, B., KRAMER, J. and FINKELSTEIN, A.: 'A Framework for Expressing the Relationships Between Multiple Views in Requirements Specifications', *IEEE Trans. on Softw. Eng.*, 1994, **20**, (10), pp.760-773

[6] ROBINSON, W. N. and FICKAS, S.: 'Supporting Multi-Perspective Requirements Engineering'. Proc. 1st Intl. Conf. on Requirements Engineering, Colorado Springs, Colorado, April 1994

[7] GREENSPAN, S., MYLOPOULOS, J. and BORGIDA, A.: 'On Formal Requirements Modeling Languages: RML Revisited'. Proc. 16th Intl. Conf. on Software Engineering, Sorrento, Italy, May 1994

[8] JARKE, M. and POHL, K.: 'Requirements engineering in 2001: (virtually) managing a changing reality', *Softw. Eng. J.*, 1994, **9**, (6), pp.257-266

[9] GOTEL, O. C. Z. and FINKELSTEIN, A. C. W.: 'An Analysis of the Requirements Traceability Problem'. Proc. 1st Intl. Conf. on Requirements Engineering, Colorado Springs, Colorado, April 1994

[10] GOTEL, O. C. Z.: 'Contribution Structures for Requirements Traceability'. Draft Ph.D. Thesis, Imperial College of Science, Technology and Medicine, University of London, August 1995

[11] GOTEL, O. C. Z. and FINKELSTEIN, A. C. W.: 'Contribution Structures'. Proc 2nd IEEE Intl. Symp. on Requirements Engineering, York, England, March 1995

[12] GOFFMAN, E.: 'Footing', *Semiotica*, 1979, **25**, (1), pp.1-29

[13] LEVINSON, S. C.: 'Putting Linguistics on a Proper Footing: Explorations in Goffman's Concepts of Participation', *in* DREW, P. and WOOTTON, A. (Eds.): 'Erving Goffman: Exploring the Interaction Order' (Polity Press and Basil Blackwell, 1988)

[14] BANTON, M.: 'Roles: An Introduction to the Study of Social Relations' (Tavistock Publications Ltd., 1965)

[15] GRUDIN, J.: 'Groupware and Social Dynamics: Eight Challenges for Developers', *CACM*, 1994, **37**, (1), pp.93-105