POSITION STATEMENT

The Contribution Structures of Networked Organisations

Orlena Gotel & Anthony Finkelstein

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

Abstract

Networked organisations promote distributed and dynamic working arrangements. One of the largest problems they face is the reconfiguration of their personnel, a problem compounded by changing roles and commitments. This problem is especially rife when systems and software are engineered both in and for such organisations. In this paper, we argue that the enabling technologies underpinning current requirements engineering practices need to be re-examined if they are to support development in and for networked organisations. We focus our concern on requirements traceability and describe a particular need to keep track of those people with a stake in the requirements, as well as in the wider requirements engineering process. We outline an approach for doing this based on modelling the contribution structure underlying requirements. We explain why such an approach needs to be adopted when requirements engineering is conducted within a disperse and heterogeneous setting. We also explain how use of such an approach provides an infrastructure through which networked systems and software can be developed and maintained. We then indicate how knowledge of contribution structures can assist with organisational re-engineering and give a brief account of our related research interests.

1. Introduction

The phrase "networked organisation" conjures up an image of a working arrangement in which the organisation's members are distributed in some way, but where they ultimately work together to achieve its goals. For the purposes of this paper, we focus on a physical notion of distribution and consider the implications both in and for the systems and software development process. More specifically, we consider the problem of physically distributed stakeholders in and for requirements engineering. Note that these stakeholders encompass the customers and end-users of a proposed system or piece of software. They also include the requirements engineers and developers. Although interrelated, we do not examine the problems arising from conceptual distribution in this paper. These problems are the remit of work on requirements viewpoint analysis, negotiation, reconciliation and so forth.

In pursuing our focus, we suggest that today's requirements engineering practices are increasingly likely to involve distributed sets of engineers engaged in analysing and specifying requirements gathered from distributed sets of end-users. Staff turn-over, coupled with shifting roles, working interrelations and commitments, is likely to become commonplace in the organisation of tomorrow. As a consequence, we anticipate that the problem of keeping track of evolving requirements will become even more prevalent than it presently is. This will have knock-on effects regarding the quality of the systems and software that can be delivered. This is because requirements traceability provides one of the most powerful mechanisms through which systems and software can be built to account for needs and tested to ascertain conformance. Currently, the techniques traditionally used for conducting and co-ordinating the requirements engineering process tend to be highly dependent upon centralised repositories and forms of management. They rarely exploit any true form of collaborative technology. They will therefore need to adapt if they are to support these changing organisational demands.

In this paper, we summarise what the major problems with requirements traceability currently are and then explain why it is so crucial to keep track of the people who have been involved in the requirements engineering process. We claim that the identification and location of such people is a non-trivial issue, one which is exacerbated by networked approaches to working. We outline an approach to do this which is based on modelling the contribution structure underlying requirements. We then describe how contribution structures can help to address the kinds of problem that arise in a distributed requirements engineering setting. Finally, we point to a number of other benefits that accrue and delineate our research agenda.

2. Contribution structures for 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) [Gotel & Finkelstein 1994]. The two main types of requirements traceability are illustrated in Figure 1. 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. This is because it provides a way to establish and maintain a connection between the information gathered from end-users and customers, the requirements which have been derived from this information by developers, and the subsequent project artifacts in which these requirements have been disseminated and addressed.

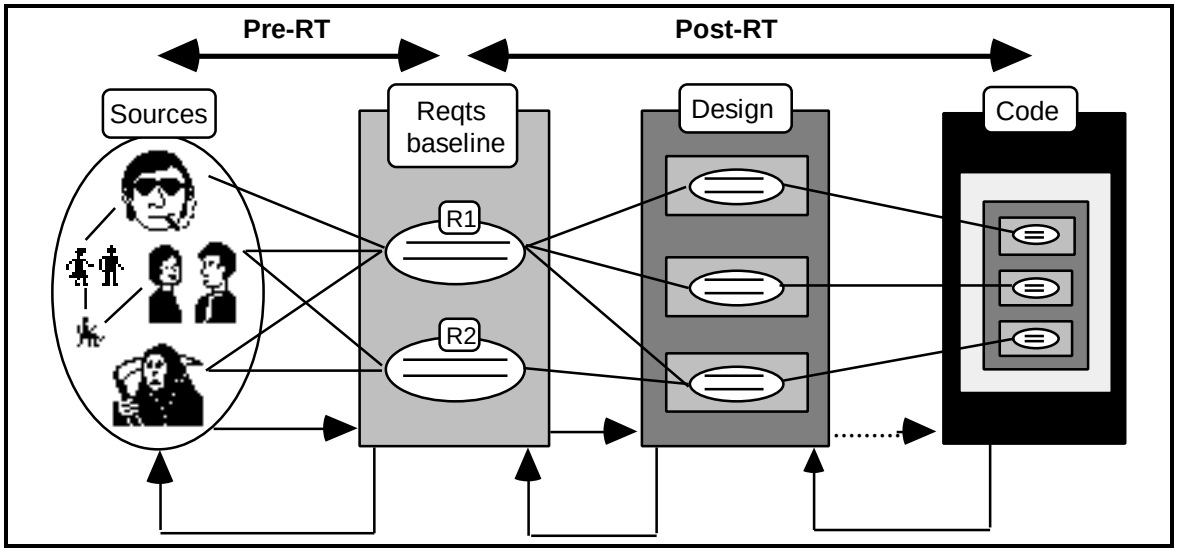

Figure 1: Pre-requirements traceability and post-requirements traceability (simplified).

Endeavours to increase the potential for requirements traceability have mostly involved uncovering and recording ever more comprehensive information about the requirements engineering process, then linking this information together in interesting ways. Through our empirical studies, we found that the most fundamental information to uncover and record for relieving long-term requirements traceability problems is that which can identify the human sources of requirements, requirements-related information and requirements-related work [*op. cit.*]. This is because people are considered the ultimate baseline whenever requirements need to be re-examined or re-worked. However, we also found that such information tends to be discarded in the strive to replace the need for human contact with exhaustive documentation. In [Gotel & Finkelstein 1995], we therefore outlined a dedicated approach which extends conventional forms of *artifact-based* requirements traceability with a form of *personnel-based* requirements traceability to address this focal problem. These distinctions are illustrated in Figure 2. The approach is based on modelling the contribution structure underlying requirements and has been fully documented in [Gotel 1995; Gotel & Finkelstein 1996ab].

In summary, the approach involves minimal extensions to artifact-based requirements traceability to augment their traces with personnel. These extensions take the form of:

(a) the semantic classification of the artifact-based relations ordinarily put in place for requirements traceability purposes; with
(b) the linking of the artifacts produced to a record of the people who have contributed to their production in various ways.

Together, this extra information can be used to reveal further attributes about the contributions and their contributors. It can also be used to infer details about social roles, role relations and commitments. The contribution structure is therefore described by the overall system of people involved in the production of requirements, along with the numerous relations they are involved in. The main steps of the approach are listed in Figure 3 for clarification, but described no further here.

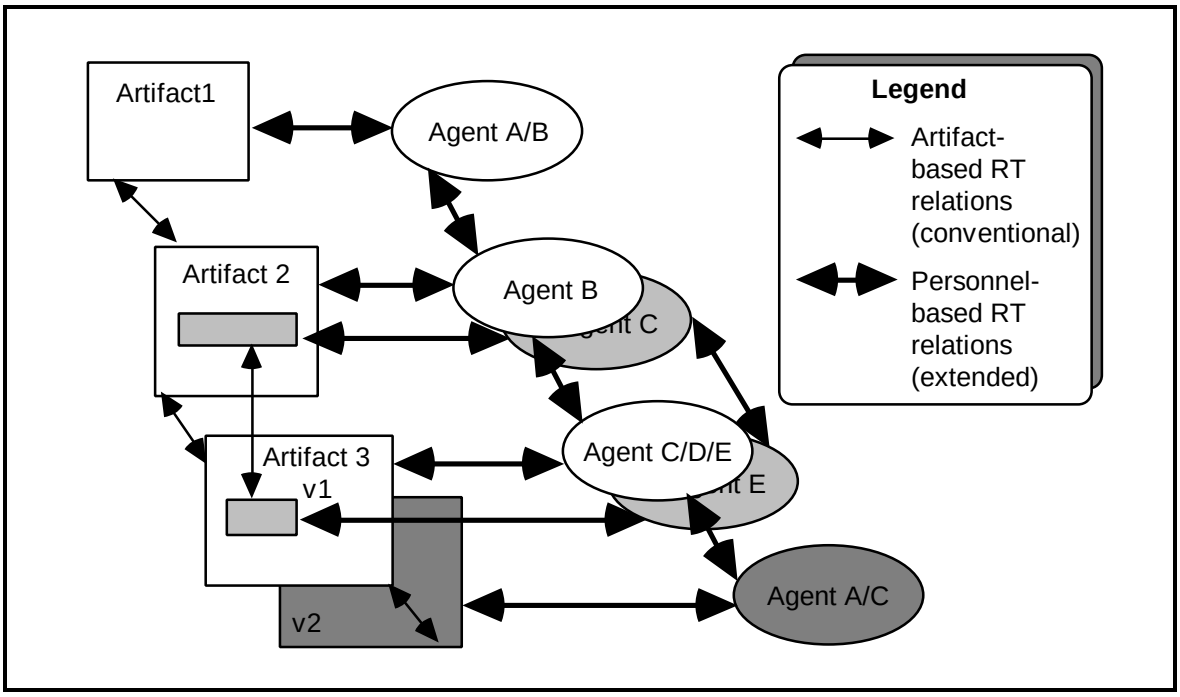

Figure 2: Artifact-based and personnel-based requirements traceability.

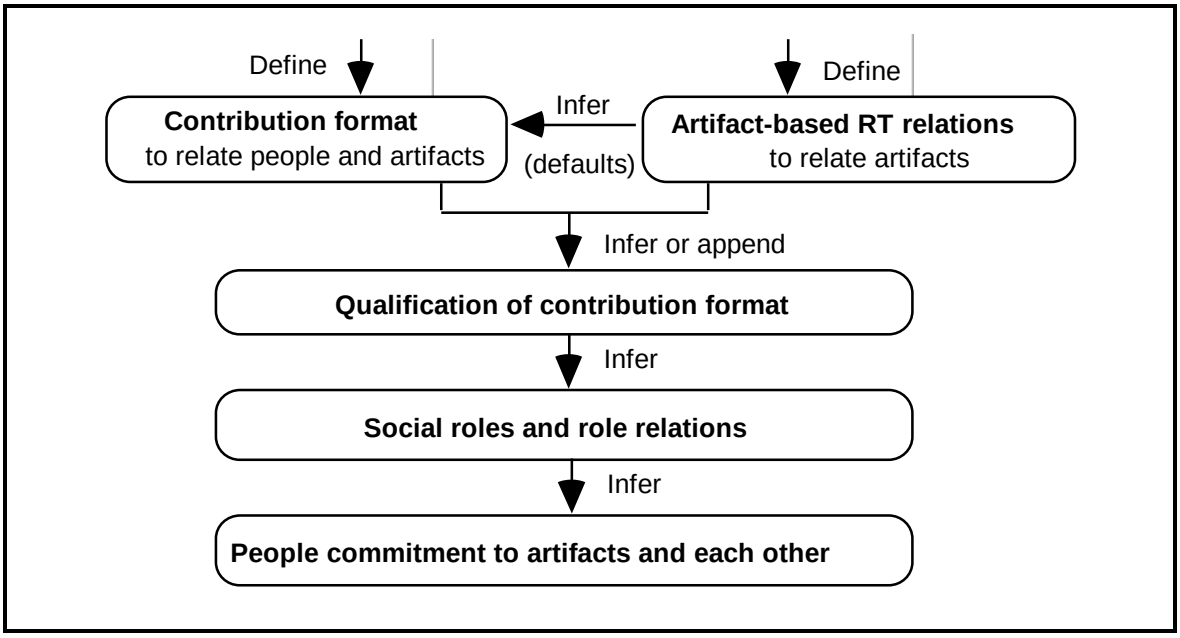

Figure 3: Steps of the approach.

3. Requirements engineering in and for networked organisations

We first consider the situation where the requirements engineering process itself is conducted within a networked organisation. This situation is depicted in Figure 4 and is exemplified by those project teams which include home workers and a rapidly changing mixture of contractors. The main problems with such requirements engineering arrangements are ones of communication and co-ordination. A particular issue is managing the requirements baseline when faced with continued contributions to be integrated. When inconsistencies are recognised, or when further information is needed for clarification, access to those responsible for the contributions is often essential. However, rediscovering who the right people are is generally a problem. Similarly, without an awareness of the social structure that has been woven during the on-going process, it is problematic to identify the relevant people to whom subsequent changes should be requested or communicated, or to suggest ways to handle a changing workforce. Although requirements

traceability can provide a framework through which all the tangible contributions can be interrelated, only by also modelling the contribution structure underlying each of these interrelated contributions can we begin to address these problems.

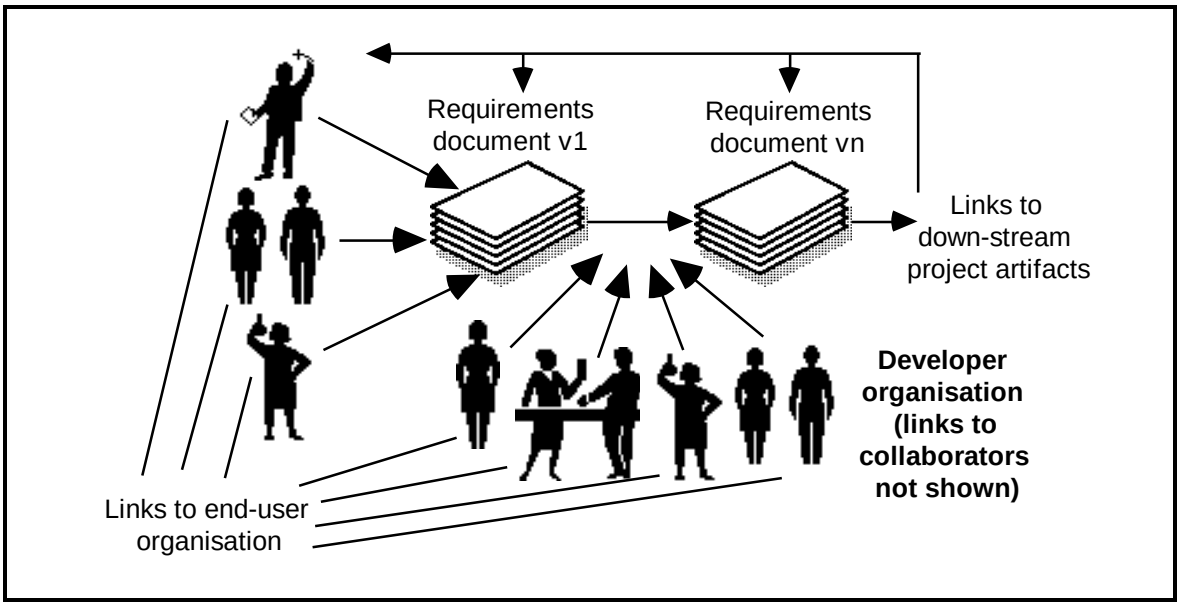

Figure 4: Requirements engineering conducted by distributed developers (simplified).

We now consider the situation where the organisation in which a proposed system or software solution is to fit is networked in structure. This situation is depicted in Figure 5 and is exemplified by large international organisations which have regionalised presence for sales, production, marketing and so forth. The main problems here are also ones of communication and co-ordination. For instance, how do we ensure that the needs of all the potential end-users have been elicited and taken into account? How do we detect and deal with changing needs as end-users come and go? How do we get end-users to recognise (let alone appreciate) other perspectives when finalising their needs? Over time, it becomes particularly difficult to keep track of the people from whom requirements were initially elicited and to keep track of who else has since shaped their development and refinement. Although requirements traceability can provide a framework through which to keep track of the original and subsequent contributions, only by also modelling the contribution structure underlying these contributions can we ground requirements sources in the network of people from which they arose.

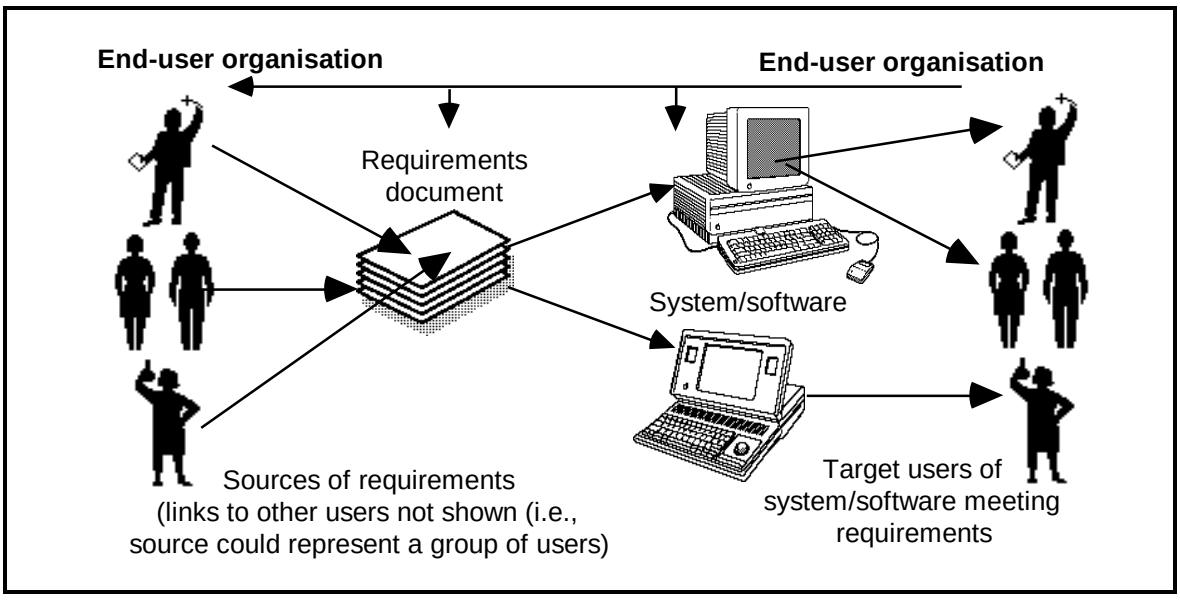

Figure 4: Requirements engineered from distributed users (simplified).

The distinctions we have just made are not so separate in practice. In any requirements engineering process, the developers and end-users would be intertwined in many intricate ways. Contribution structures provide a disciplined way to record and examine the dynamics between all the parties with a stake in the final system or software. It is this combined information, revealing the changing relations and commitments of the distributed stakeholders, that can tell us about how the process is actually carried out. Ways to decentralise this process and introduce collaborative technologies can only be understood once the nature and dynamics of these working practices are clear. In addition, once a system or piece of software has been deployed in the networked organisation for which it has been built, knowledge of the contribution structure underlying the project can provide a mechanism through which problems can feed back into the maintenance process. Not only can these problems be channelled back to the appropriate developers and end-users responsible for the requirement in question, but other affected stakeholders can be identified and informed as necessary. Reaction to trouble reports or changing requirements should be faster and more controlled. Furthermore, this process would provide in-built rationale.

4. Organisational re-engineering

A model of the contribution structure underlying the work of the development organisation in the requirements engineering process enables us to see how this process is actually structured in practice. It uncovers the informal roles, working relations and commitments that have been formed between team members. This information can reveal how members of a development organisation work together, pointing to effective and less effective partnerships, and so inform ways in which the development team could be restructured. If subsequently restructured, contribution structures could also be used to examine whether the proposed re-organisation has actually been taken up in practice. With knowledge of the working structure of the development organisation, it becomes possible to recommend developmental approaches which would best reflect the structure and characteristics of the particular organisation.

A model of the contribution structure underlying the end-user contributions in the requirements engineering process enables us to assess aspects such as requirements coverage. It also allows us to examine how the various needs are interrelated. Where seemingly unrelated stakeholders are found to be related by their shared or overlapping needs, this can reveal subtle information about the end-user organisation and its needs that is not accessible to contemporary requirements elicitation techniques. Such knowledge would again help inform any proposed re-structuring of the organisation. Moreover, it would be possible to examine the extent to which the eventual structure of the system or software reflects the formal or informal structure of the organisation for which it has been developed.

In this way, a model of the full contribution structure which accounts for all the stakeholders involved in the requirements engineering process would provide a firm handle for process improvement. In turn, this knowledge would provide the foundations against which it would be possible to explore how best to exploit collaborative technologies.

5. Towards distributed requirements traceability

One of the challenges we are currently working on is that of distributed requirements traceability. We are particularly interested in maintaining requirements traceability across interrelated projects and programmes, themselves dispersed throughout global enterprises. We believe it will become all the more essential to model contribution structures as a way to provide the firmest of requirements anchors in such environments. The questions become: How can we maintain requirements traceability without centralised repositories or forms of control? What is an appropriate decentralised architecture for requirements traceability?

In [Gotel 1995], we suggested how world wide web technology could provide some support for the above. Not only does the internet provide an ideal medium through which to support distributed working, but agreed document standards mean that requirements artifacts can be shared and contextualised without the kinds of duplication and lengthy circulation times that frequently beset paper and lead to versioning and configuration problems. As it is becoming accepted practice for many requirements documents and related artifacts to be marked up using forms of descriptive markup, we demonstrated how a language similar to the hypertext markup language could be used to embed requirements traceability links with requirements artifacts themselves. We described how these links could be typed to support different forms of artifact-based requirements traceability within a document corpus. We also described how this language could be configured to tag details about contributors to artifact fragments and further used for selective forms of personnel-based requirements traceability. Where such extended forms of requirements traceability are achieved through markup, the artifacts become modular and amenable for reuse. However, the many procedural and management issues that would arise as requirements traceability links and contribution information becomes dispersed between distributed artifacts have yet to be fully investigated.

6. Summary

In this paper, we have explored the use of contribution structures both in and for networked organisations. We have focused on networked development organisations and networked end-user organisations from the perspective of requirements engineering. Notwithstanding the "political" problems, we have argued a need to provide stakeholder visibility if we are to make distributed forms of requirements engineering a reality. In particular, we have described how the modelling of contribution structures underlying requirements can help to address some of the more important problems inherent in a distributed requirements setting.

References

A full bibliography of related work can be found in the references below. In particular, the ICRE paper provides a survey of the state-of-the-art in requirements traceability, whilst the RE 95 paper describes those areas which have influenced our work on contribution structures.

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

[Gotel & Finkelstein 1994] 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*, Colorado Springs, Colorado (April 1994), pp. 94-101.

[Gotel & Finkelstein 1995] Gotel, O. C. Z. and Finkelstein, A. C. W. Contribution Structures, *Proceedings of the 2nd IEEE International Symposium on Requirements Engineering*, York (March 1995), pp. 100-107.

[Gotel & Finkelstein 1996a] Gotel, O.C.Z. and Finkelstein, A.C.W. Revisiting Requirements Production, *Software Engineering Journal* (to appear May 1996).

[Gotel & Finkelstein 1996b] Gotel, O.C.Z. and Finkelstein, A.C.W. *Extending Requirements Traceability Through Contribution Structures* (submitted for publication in March 1996).