

Motivation Matters in the Traceability Trenches

Patrick Mäder¹, Orlena Gotel² and Ilka Philippow¹

¹Department of Software Systems
Ilmenau Technical University, Germany

patrick.maeder|ilka.philippow@tu-ilmenau.de

²Department of Computer Science
Pace University, New York, USA

ogotel@pace.edu

Abstract

Reports from the field are few and far between when it comes to traceability. As a community, we know little more about the traceability practice in companies today than we did a decade ago. This paper reports on findings from a practitioner survey designed to get a high-level update on traceability practice and problems. What emerges is the importance of the prevailing motivation underlying traceability adoption in an organization and we characterize this in four ways. We use these perspectives to discuss our findings and their implications.

Keywords: Requirements traceability, traceability practice, traceability problem, traceability process.

1 Introduction

Two major studies have been conducted on traceability practices in industry ([3], [4]), but their results date back over ten years. Much research has since been conducted on traceability that references these studies as rationale for the problems being tackled. However, we know very little about the traceability practice as it occurs in companies today. Is it widely undertaken and used? Are the problems experienced the same as a decade ago? What, if anything, has changed and where does the research emphasis now need to lie?

This paper reports on an exploratory study of the traceability practice and problems within ten companies based predominantly in Germany. In addition to a status update, our work can be seen as a scoping exercise for preparing broader field studies. The paper provides relevant details about each of the ten cases under study, discusses our findings and lists some implications for the traceability community. In the course of analyzing the data, it became clear that traceability was being done in four different environments, each with their own underlying motivation: (1) regulated - to satisfy industry-imposed or company-imposed regulations; (2) sub-contractor - to align with the processes and demands of customers; (3) consultant - to support in-

ternal or external company-specific efforts to improve processes; and (4) enthusiast - due to an energized individual with knowledge and a passion to learn. We use these categories to characterize the practices and problems experienced in the ten cases.

The paper is structured as follows. In Section 2, we summarize key prior traceability practitioner surveys. Our research questions are listed in Section 3 and we describe the study designed to begin to examine these questions in Section 4. We profile the ten companies involved in our study in Section 5. Our findings are given in Section 6 and the limitations of our study are considered in Section 7. We finish the paper in Section 8 with conclusions and recommendations for future work.

2 Related Work

Gotel and Finkelstein report the findings of a year-long empirical study into traceability conducted in 1992 [3]. The study involved around one hundred software development practitioners, holding a variety of positions within a large organization, with experience ranging between 0.75 and 30 years on a variety of project types. A questionnaire gathered general data on projects and traceability (55 returned), and a second questionnaire was tailored to the specific working area of the subject (33 returned). In addition, five focus group sessions were conducted with thirty-seven practitioners to consolidate the data, and independent observation of practical requirements gathering and development workshops took place. An extensive technique and tool survey occurred in parallel. One major finding was that support for post-requirements traceability existed and the problems lay more with the process and people. They found multiple perspectives on what traceability was expected to enable and on the problems experienced, conflicts particularly evident between those parties responsible for establishing traceability and those parties using it (not always identical people). Further, they found that pre-requirements traceability was in need of more attention, so the need to integrate a wider variety of data in traceability, such as source material and

the people involved in the project. The authors were concerned with understanding and exposing the scope of the problem area; they did not report about the actual traceability practice of their subjects.

Ramesh and Jarke report on a large practitioner study of traceability where the data collection took over three years during the 1990's [4]. The authors conducted a pilot study with fifty-eight masters in information technology students to create an initial traceability meta-model and inform the design of the study. They also performed a survey of tools. The main study consisted of thirty focus group discussions, each with about five people from twenty-six companies. Their primary focus was on the types of traceability link used in current and ideal practice. The study comprised two phases: the first revolved around agreeing a traceability meta-model and the second on defining reference models for other practitioners to use. Again, the everyday traceability practice of their subjects was not exposed.

Arkley and Riddle report on a survey of nine software projects, small to multinational in scope, undertaken using questionnaires and interviews [2]. The authors identified three issues related to traceability: the necessity for extra entry data when using traceability tools; a lack of understanding on how to employ traceability information; and the lack of perceived direct benefits to the main development process. The authors also did not discuss the actual traceability practice of their subjects.

Ahmad and Ghazali report on interviews conducted with fifteen practitioners and on the analysis of project documentation from three different IT companies [1]. Their subjects had six to ten years of practical experience in developing small projects. Their major finding was that the subjects perceived pre-requirements traceability to be more beneficial than post-requirements traceability. Again, the authors did not detail the traceability practice of their subjects.

3 Research Motivation and Questions

The first author of this paper has drawn upon the findings of the studies summarized in Section 2 to justify his traceability research. Given the age of these findings, he decided to sample a number of companies to get an update as to what is being done in industry today and to see whether the key problems still hold. This was not intended to be a thorough survey of all industries and project types. Rather, to contextualize the above mentioned research effort and to act as a scoping study to inform future surveys of traceability practice and problems. The high-level questions included: Whether – Are practitioners doing traceability? Where – What kinds of companies are doing traceability (i.e., size, domain, project types)? Why – Why is traceability being undertaken? Who – Who is

establishing the traceability, maintaining it and using it? When – When is the traceability established, maintained and used? What – What level and degree of traceability is actually being done? What artifacts are included? What relations are formed? How – Do practitioners define and follow a traceability process? What is the role of tools?

4 Study Design

There are a number of ways to gain answers to our high-level traceability questions. Since our objective was to collect credible data in a short period of time and to inform more substantive empirical studies, we chose to conduct a field survey of ten companies undertaking software systems development. This would give us access to practitioners and companies to clarify points and explore our findings.

4.1 Selection of Cases

Ten business cards were selected at random from amongst thirty cards that had been collected from practitioners within commercial companies over a one year period, from exhibitions, presentations, former contacts, etc. Note that we specifically excluded companies within the military domain to gain insight into everyday commercial practice. Only two of the ten practitioners selected were known to have a prior interest in traceability topics.

The ten practitioners were contacted by phone and all agreed to an interview. This is an observation that we want to highlight: in contrast to what is often reported about the unavailability of concrete data from industry, we found that these practitioners instantly agreed to provide data about their traceability practice and problems. During the interviews, the practitioners were asked to report on a representative project in order to gather specific data on roles, artifacts, processes and tools used. Most reported on their current project. More general questions were answered based on their cumulative project experience.

The participating companies were mostly German, except for one located in the Czech Republic, and each was involved in developments that targeted European or international markets. The size of the participating companies included six medium (50–500 employees) and three large (>500 employees) companies. The only small company (<50 employees) actually reported about a consulting project they were undertaking for a large company and the traceability practices therein. Our cases included a mix of software development offering, including companies who predominantly create end products to sell to a user, who do project development work for other companies supplying a market, or who provide expert advice on processes, techniques and methods. Most worked in the transportation domain (avionics and automotive). Table 1 summarizes the background of the ten companies.

Category	Regulated		Sub-contractor			Consultant		Enthusiast		
Case	5	6	1	2	9	3	7	4	8	10
Employees	>500	>500	50-500	50-500	50-500	>500	<50	50-500	50-500	50-500
Company age (yrs)	>25	20	7	17	10	>25	8	18	11	8
Offering	Product	Product	Project	Project	Product	Consulting	Consulting	Project, Product		
Business area	Automotive	Automotive	Avionic, medical	Avionic	Automotive	Transport., diverse	Automotive	Security	Insurance	Transport.

Table 1. Background of cases (companies)

The subjects we interviewed held the following positions: three system analysts, two consultants, one requirements engineer and four team or project leaders. Six of the subjects reported about automotive projects, two about avionics projects, one about an IT security project and one about an insurance project. In one company, case 9, two additional practitioners wanted to be involved, so we aggregated their data for this case. Table 2 provides information about the experience of the subjects and the duration of the projects they reported on.

	Avg.	Min.	Max.
Systems development experience (yrs)	8.6	2	15
Time working for company (yrs)	4.8	2	10
Duration of reported project (yrs)	2.1	0.75	5

Table 2. Background of subjects

4.2 Process

Four of our interviews took place in the company's offices; the other six took place by phone. We held from one to three interviews with each practitioner, typically lasting between one and two hours. In all cases, a prepared interview protocol was used to ensure that we examined the research questions listed in Section 3 consistently. The interviews were audio recorded and notes were taken. In certain cases, more practitioners from within the company joined a conference call to discuss certain practices and issues, particularly ones associated with tools. We also collected background information about the company, the practitioner and the project being discussed via questionnaire.

5 Categorizing the Cases

To report on the cases, we distinguish four types of environment in which traceability was observed to happen. These distinctions characterize the underlying motivation for doing traceability, and appear to help explain the nature of the practice and problems experienced (see Figure 1). We suggest that the prevailing motivation in the environment establishes a culture that influences how traceability is both established and perceived, and how it can fit in with other

environments, so forms a foundation upon which to analyze specific details of traceability stakeholders and needs.




Figure 1. Relation between categories

The regulated – conforms with rules and regulations. Cases 5 and 6 were large companies developing in the automotive domain. In the regulated environment, traceability is strategic to business and the investment is high. Without shared and agreed processes for traceability, both followed and enforced, products cannot enter the market. In order to achieve the required quality, in-house and/or external consultants are employed to define a development process to be followed by all sub-divisions and sub-contractors.

The sub-contractor – agrees to perform services for others. Cases 1 and 2 were sub-contractors working on large avionics projects. Case 9 was a sub-contractor for a number of different automotive manufacturing projects. The sub-contractor not only has to fulfill the regulations of the domain being worked in (if applicable), but usually has to fulfill much more specific process, documentation and tool requirements imposed by the customer (if the customer is regulated or focused on process improvement).

The consultant – gives professional expert advice. Cases 3 and 7 were both developing a traceability methodology for large companies in the automotive domain, case 7 as an external consultant and case 3 as an internal technology department. Given the consultant's job depends upon studying techniques in depth and leveraging past work, the consultant is often better able to develop an abstract traceability process that applies to many development contexts.

The enthusiast - interested and willing to learn. Cases 4, 8 and 10 were found to practice traceability purely through personal interest and choice. These cases were domains in which process improvement was not an organizational concern and where traceability was not mandated.

6 Findings and Observations

Table 3 provides an overview of our main findings and these are elaborated upon in detail below.

Question	Regulated	Sub-Contractor	Consultant	Enthusiast
Whether?	yes			
Where?	mainly transportation			diverse
Why?	maturity, regulations	driven by regulated	part of consulting	benefits
Who?	developer create, developer and project leader use			
When?	in parallel with development			
What?	requirements, tests, design as major artifacts			
How?	very limited definition, less constraints			
	multiple tools (DOORS, modeling, test)			single tool

Table 3. Overview of findings

6.1 Whether, Where and Why?

All the companies had traceability established within their software development processes. From our study, we can report that these were mostly medium to large companies working in the automotive industry. In the interviews, we explored the more specific reasons they were each doing traceability, in terms of the benefits they anticipated. We asked about better impact analysis, documentation of dependencies, navigation between artifacts, coverage analysis, change support and validation support, and also sought additional reasons. While all subjects could state what was expected of traceability, none could articulate what was actually attained.

The regulated were developing large systems and were doing so with a number of specialized divisions and outsourced sub-contractors (e.g., engine control systems and interior security systems, then combined into one car). The diversity, complexity and dependability requirements cause quality problems for the final system that the regulated tries to meet with a mature process that is enforced for all the development stakeholders. The goal to attain maturity certification, like CMMI level 3, was mentioned as a way to institutionalize such processes. If the developed system was to be safety-critical, then additional regulations applied that required traceability by law for that development (e.g., imposed by the European Aviation Safety Agency). The regulated were interested in using traceability to assess the current implementation status of any project and to gain coverage information. They were interested in ways to exchange test results because, while the sub-contractor has to perform component tests, the regulated has to perform integration tests in the final environment. There was also interest in using traceability for impact analyses, but with less priority. Change support and navigation played a less important role.

The sub-contractors specialized in the development of certain components of systems being built by the regulated

(e.g., providing navigation systems for car manufacturers), has to adapt to the processes, tools and artifacts of the regulated, that is their main driver for adopting traceability. The sub-contractors were especially interested in using traceability for coverage analyses and validation support (as per the regulated), but the navigation between related artifacts became equally important. There was interest in using traceability for change support, but with less priority, and impact analysis played an even less important role.

The consultants specialized in creating development processes for the regulated to address either maturity or regulation requirements, or both. The consultants reported validation support as the important benefit of using traceability. Using traceability for impact analyses played a less important role, while change support and coverage analysis were perceived to be even less important.

The enthusiasts were not subject to organizational improvement goals or regulations in their developments. Nevertheless, all three cases reported process quality, maturity and control as their reason for adopting traceability. The subject of case 4 had detailed knowledge of a tool's functionality, coupled with the freedom to explore, so he adapted his process to fit the abilities of the tool, enabled traceability and now maintains much of his project work (traced) within it. The subject of case 8 mentioned the need for his company to stay competitive as the reason for his interest in traceability. He had been developing a concept for issue tracking and selected a tool to support it. As a side-effect, the tool provided traceability between issues. All the enthusiasts reported a traceability solution that included only one development tool (a modeling or issue tracking tool) and that the tool was the trigger for adopting traceability in their projects. The traceability process can easily become what the technologies enable. The enthusiasts used traceability for navigation support between related artifacts and documentation. They perceived coverage analysis as important, but with less priority, and impact analysis, change support and validation support were regarded as less important.

6.2 Who, When and What?

Several questions were asked regarding the traceability process of the reported project. None of the cases reported having defined roles or tasks for the creation, maintenance or use of traceability. Since the process the regulated uses is often created and improved by the consultant, and applied by the sub-contractor, we found no major differences between these categories in terms of the fundamental process.

The regulated, sub-contractor and consultant. All these subjects reported that they establish and maintain traceability in parallel to the development activities rather than post development. Only one of the subjects reported that, under time pressure, traceability may have to wait. All the subjects reported that they established at least two

levels of requirements traceability (like high-level and low-level requirements traced to each other). They also traced test cases to requirements or design, and validated them. Except for one of the regulated subjects, all reported tracing between requirements and design artifacts with differing levels of granularity (i.e., to components, classes, methods, functions, etc.) Pre-requirements traceability between requirements and directives was reported by both regulated subjects. One each of the sub-contractor and consultant subjects reported tracing into code. All subjects expressed difficulties in finding the right granularity at which to relate artifacts and sought advice from the interviewer.

In all cases, the person maintaining the artifact was also creating and maintaining the traceability, be this the developer, tester, system analyst or requirements engineer. In addition to the project leader, these are the same roles that were reported to be using the traceability. Three of the subjects mentioned the effort necessary to create and maintain traceability, but none said it was not worth it. Both consultants reported that they had defined a traceability information model for their project. One said that only predefined relations were allowed to be created within projects and that these types were altered according to new needs or insights. The other mentioned that intended traces were defined, but the developer could decide to create additional ones. One regulated subject explained that their in-house tool only allowed for predefined traceability relations and, as a consequence, computed analyses during nightly builds.

The enthusiast. As with all the other subjects, these subjects reported traceability as occurring in parallel with regular development activities. All the subjects reported tracing levels of requirements and two of them traced test cases to requirements or design. Only two subjects reported tracing requirements to design artifacts. Two subjects reported tracing into code. All the subjects emphasized the importance of user interface models in communication with the customer, so they commonly traced to these models. One subject reported tracing to bug tracking tickets and work items. The roles responsible for doing and using the traceability were identical with the categories above.

One subject had defined an information model to visualize intended traceability relations. We saw that model and learned that it was not consistent with the current traceability, that some of the relations were only intended for certain circumstances and others were not shown at all. Furthermore, we learned that they had defined company-specific link types, but the tool being used was not able to constrain these link types to certain artifact types, so they got instantiated inappropriately (e.g., verification relations between bug tickets). One subject mentioned the need for advice on granularity decisions, especially to reduce the overall number of traces. The high effort for traceability creation and maintenance was also emphasized.

6.3 How?

Another topic within our interviews involved exploring how the traceability process was implemented and supported, and the main issues experienced in so doing.

The regulated. Given the strategic importance of traceability, the cost of tools is usually only a minor factor in such environments. Licenses are purchased in volume and budgets are relatively high. As a consequence, expensive tools tend to be used and then their use is also required of sub-contractors. All regulated, sub-contractor and consultant subjects reported DOORS as their requirements management tool, along with use of additional modeling tools. The main issue confronting the regulated was finding a workflow to do traceability over organizational boundaries: How to provide only the important part of the specification to the contractor, while including the traceability? How to create and provide traces between the initial specification and the refined specification, design and implementation that the sub-contractor develops? How to propagate changes to artifacts and maintain consistency between them? Both subjects mentioned the versioning of traceability relations as a major issue they were facing.

The sub-contractor. The sub-contractor has the same problems regarding achieving traceability over organizational boundaries and also reported problems in bridging the gap between specification and design. The sub-contractors were especially frustrated by the weak integration between requirements management and modeling tools in terms of traceability, and the inability to do sophisticated forms of analyses over tool boundaries. They want an efficient way to propagate changes between the specification, design and implementation, and they want ways to perform coverage analyses. Model and traceability exchange between different tools was a real issue. The sub-contractor is often already using a tool other than that required by the regulated and want to save the cost of additional licences and training. They therefore seek ways to exchange data with the tool the regulated requires. Even where the same tool is used, data exchange can still be problematic and provide intellectual property issues.

The consultant. Both subjects reported that they had decided on a whole tool chain for the regulated companies they were working with, consisting of requirements management, modeling and testing tools, and both subjects found the tools to be the main limiting issue regarding the traceability provided. They reported difficulties in customizing tools and the poor integration solutions in terms of traceability. Both expressed a vision of a common repository for all tools in the process and a common, customizable information model integrated with it. Similar to the regulated, the consultant also expressed the need for an efficient and effective versioning of traceability relations.

The enthusiast. The subjects reported problems with visualizing and using traceability within the modeling tools they used. Those that were using issue tracking tools were frustrated by poor integration with the modeling tool they were using. All three subjects reported that they decided to keep all their artifacts in one tool to prevent consistency problems. They did not use an additional requirements management tool. The enthusiast who did not report integration problems kept all relevant artifacts, including most requirements, within the UML tool Enterprise Architect. His complaints were about insufficient functionality to get traceability information from the tool, including poor visualization and no way to find relations between two artifacts that exist over intermediate artifacts.

7 Limitations

This is a small pilot study. We investigated traceability in the field via ten separate cases. These were the reflections of individuals and from the trenches. In most cases, in-depth questions and discussions helped to clarify vague data and confirm assumptions, but we are aware that there can be a gap between telling and doing, and we had no means to determine neither how representative the practitioner was of the company concerned nor the project reported. Reporting on current projects, it is perhaps natural that the practitioners using the traceability at this point are the same as those establishing and maintaining it; we did not gather historical data on perceived value returned from traceability over time on projects. All our cases were located in central Europe, and this might have generated geographical biases, and individuals were selected after some prior contact that made it more likely they would be working with traceability. We expect that the practice of the transportation industry is quite advanced (most of our cases), with process discipline and tool uptake, so we need to complement our findings with other domains. Also, whether the findings extend to small companies and their projects has not been examined. Our findings and recommendations are offered as preliminary.

8 Conclusions and Future Work

This survey paper provides some insight into today's traceability practice and problems. The problems are not unlike those reported in the past, but the role of tools is now more prominent in practice, possibly a result of the increasing complexity of software systems being built. Pre-requirements traceability appears slightly more supported, but there are still multiple expectations as to what traceability should assist with at a practical project level and unproven actual benefits. Similar to Ramesh and Jarke, we found two distinct groups of traceability user, but less in the intended use and practice of traceability, but more in the motivation and organizational setting they practice traceability within. For that reason we think that traceability re-

search in the future should also include specific needs that arise from such settings.

Tools are central enablers, like it or not. Independent from how much traceability is intended by the practitioner, the traceability attained is influenced by the abilities of the tool(s) used. This is even more pronounced than was the case in the 1990s. While practitioners live with what tools provide, they would probably do more sophisticated traceability if it were available - they are asking for better analyses and visualizations. Practitioners, researchers and tool vendors need to work together.

More practitioner guidance and conceptual support. Almost no guidance is available for practitioners to help them establish traceability in their projects. For instance, decisions about granularity are problematic and can lead to either over-complex or inadequate traceability graphs. An issue mentioned by the consultants was the missing support for defining traceability information models and process life-cycles within tools, helping to constrain and check trace possibilities and to enable automated traceability analyses. We need to give practitioners easier means to define and establish their traceability, and make the traceability information model not purely descriptive, but enforced and monitored by the tool. With that support, practitioners will gain value from the definition.

Traceability across boundaries. The regulated and sub-contractors asked for a workflow to enable analyses and change propagation across organizational and regional boundaries. Distributed traceability needs more attention.

Versioning of trace relations. An integrated artifact repository that holds at least the major models of the development process would make the storage and versioning of traceability relations less problematic. This would require an agreed format among tool vendors, something unlikely to be popular. After several mergers, IBM now owns a large number of major CASE tools, so this could be a chance to provide at least a unified repository among these tools.

More empirical studies from the field. We need more data on traceability reality. Our pilot shows that projects for study are available and practitioners are willing to talk.

Acknowledgments The authors would like to thank all the practitioners involved in this pilot study.

References

- [1] A. Ahmad and M. Ghazali. Documenting requirements traceability information for small projects. *IEEE Int'l Multitopic Conference INMIC 2007.*, pages 1–5, Dec. 2007.
- [2] P. Arkley and S. Riddle. Overcoming the traceability benefit problem. In *13th Int'l Req. Eng. Conf.*, pages 385–389, 2005.
- [3] O. Gotel and A. Finkelstein. An analysis of the requirements traceability problem. In *Proc. 1st Int'l Conf. Req. Eng. ICRE94.*, pages 94–101, Colorado Springs, CO, April 1994.
- [4] B. Ramesh and M. Jarke. Toward reference models of requirements traceability. *IEEE TSE*, 27(1):58–93, 2001.