P. Jönsson, and C. Wohlin, "Understanding Impact Analysis: An Empirical Study to Capture Knowledge on Different Organisational Levels", IEEE Proceedings of International Conference on Software Engineering and Knowledge Engineering", pp. 707-712, Taipei, Taiwan, 2005.

Understanding Impact Analysis: An Empirical Study to Capture Knowledge on Different Organisational Levels

Per Jönsson and Claes Wohlin
School of Engineering, Blekinge Institute of Technology
PO-Box 520, SE-372 25, Ronneby, Sweden
per.jonsson@bth.se, claes.wohlin@bth.se

Abstract

Change impact analysis is a crucial change management activity that previously has been studied much from a technical perspective. In this paper, we present a systematic interview-based study of a non-technical aspect of impact analysis. In the study, we have investigated how potential issues and uses of impact analysis are viewed by industrial experts at three organisational levels, based on Anthony's decision-making model: operative, tactical and strategic. The results from our analyses show that on the whole, agreement on both issues and uses was large. There were, however, some differences among the levels in terms of issues. Thus, we conclude that it is both relevant and important to study impact analysis on different organisational levels.

1. Introduction

Change impact analysis (IA) is a crucial part of change management and requirements engineering (RE), as all software systems are exposed to changing requirements. Bohner and Arnold define IA as "*...identifying the potential consequences of a change, or estimating what needs to modified to accomplish a change.*" [2]

Research about IA is commonly found in the field of software maintenance, although IA doubtlessly plays an important role during the entire product cycle. For example, Lindvall coined Requirements-Driven Impact Analysis (RDIA) to denote the activity of identifying the impact of new requirements on an existing system [8].

The gross of IA research concerns the development of methods and algorithms for supporting and automating the analysis, or the adaptation of existing methods in new contexts. To our knowledge, there is little research about more non-technical aspects of the subject, such as process and organisational aspects. In our experience, IA is, as a part of the change management process, heavily dependent on organisational support and stakeholder views.

In this paper, we present an empirical study of the views of IA on different organisational levels. In exploring the uses (application areas) and issues of IA at a software development company, we identified three levels with different foci: one with technical focus, one with resource focus and one with product focus. Looking at management science, we found that these levels mapped well to the decision-making model originally defined by Anthony (see, for example, [1] or [9]), where decisions are categorised as *operative*, *tactical* or *strategic*.

In order to understand how potential issues and uses associated with IA are seen on the three organisational levels, we interviewed 18 employees at the company mentioned above, in their roles as industrial experts. Our hypothesis is that people on different organisational levels see IA differently, and consequently have little awareness of issues and uses on other levels.

Gathering knowledge is an important step towards being able to overcome differences. Therefore, our contribution does not only lie in the study of an organisational aspect of IA, but also in the systematic method for collecting, extracting and prioritising knowledge pertaining to issues and uses of IA.

The paper is structured as follows. Section 2 covers related work. Section 3 describes the design of the study, whereas Section 4 details how the study was carried out. Section 5 presents results, which are subsequently analysed and discussed in Section 6. Finally, conclusions are drawn in Section 7.

2. Related Work

Aurum and Wohlin tie RE activities to decision-making models, arguing that RE is a decision-intensive process [1]. They suggest that studying decision-making within RE helps organisations structure their RE decisions better and, ultimately, produce software with higher quality. We mean

that the same argument holds for IA, due to the strong connection between IA and RE.

Several researchers report on differences between managers and engineers in the context of software process improvement (SPI), for example concerning views of software quality [5], use of formal routines to transfer knowledge and experience [4] and how they rate factors affecting an SPI goal [6]. These examples demonstrate the relevance in studying different organisational levels.

Some uses of IA are mentioned in the literature, for example estimating resource needs, assessing system impact, weighing change proposals against each other and finding the overlap of parallel changes [2, 12]. Issues are mentioned as well, for example lack of automation and tools, insufficient traceability, documentation that is not updated, inconsistent models and high time-consumption [2, 3]. In this paper, we explicitly focus on both uses and issues.

3. Method

This section describes our research setting, introduces the three organisational levels and presents our method for carrying out the study. The method consists of three main steps: (1) interviews with employees, (2) results triangulation and filtering, and (3) prioritisation of the results.

3.1. Research Setting

The study was conducted at a large Swedish software development company operating on a world market. The company is one of the leaders within its domain, and has several large international customers. The organisation can be characterised as a matrix organisation, where functional areas and projects are separated [10]. Such an organisation generally fosters exchange of knowledge and experience between workers belonging to the same functional area, but may also induce conflict due to the fact that workers have several different managers.

The population we wish to generalise to is industrial software development experts in general rather than just within the company. We believe this to be possible due to the fact that the company deals with large-scale software development and is ISO 9000 certified. It can be assumed that the challenges and characteristics of the development work exist in other companies of the same magnitude as well, and to some extent also in smaller companies.

3.2. Organisational Levels

As mentioned earlier, the three organisational levels we identified map well to the decision-making model defined by Anthony. The model differentiates between decisions at three levels as follows [9]:

- *Strategic* decisions have typically large scope, large impact and long-term perspective. They concern organisational or product-related goals and strategies.
- *Tactical* decisions concern planning of time and resources to reach strategic goals, and are often made by middle management. They have smaller scope and impact, and shorter time horizon, than strategic decisions.
- *Operative* decisions are made when realising the project according to the plan, and are often of technical nature.

3.3. Interview Design

The interview instrument contained seven main topics, of which each was associated with one or more open questions. The instrument in its entirety can be obtained from the authors by request. In this paper, we focus on two of the main topics: potential issues and uses. Each of these topics consisted of only one question, as follows:

- *Which potential issues are associated with performing impact analysis?*
- *Which potential uses does impact analysis have?*

Note that we asked about potential issues and uses, rather than actual ones. The reason for this was to avoid limiting the generalisability of the results by extracting company-specific issues and uses only.

The remaining topics were more qualitative in their nature, and were intended both for providing a context for and for collecting hidden or implicit knowledge about the issues and uses. We did not intend for the participants to prioritise during the interviews, as we expected each of them to see only a subset of the possible issues and uses.

In order to ensure the appropriateness and clarity of the questions, the interview instrument was developed in close cooperation with the company where the study was conducted.

3.4. Results Triangulation and Filtering

A triangulation and filtering scheme was designed in order to get as complete lists as possible of both issues and uses. The scheme involved three information sources: (1) the lists generated in the interviews, (2) qualitative information from the interviews, and (3) information from the literature.

By using information from all interview topics, it would be possible to extract both explicit and implicit knowledge about issues and uses, and by collecting information from

the literature, we would be able to add issues and uses of which the participants were not aware.

The filtering part was intended to remove redundancies and inconsistencies in the lists by merging similar items together, and by discarding items that were not directly related to IA.

3.5. Prioritisation

In order to get prioritised lists of both issues and uses, a post-test was designed as a follow-up to the interviews. In the post-test, the participants should state the distribution of their decisions on the decision levels. Based on this, it would be possible to deduce their organisational levels. For example, a participant making mostly strategic decisions would be regarded as belonging to the strategic organisational level. This scheme was used since, as Aurum and Wohlin also point out [1], Anthony's decision levels are not entirely orthogonal.

For uses, the participants should prioritise such that the use with the highest priority would be the one most relevant to the organisation *if it was realised*. For issues, the participants should prioritise such that the issue with the highest priority would be the one most critical to the organisation *if it existed*.

We chose the organisational perspective based on our initial hypothesis. However, we were also interested in knowing if the participants would prioritise differently from an individual perspective. Trying to maintain a balance between collecting much information and keeping the post-test short, we added the individual perspective to the prioritisation of issues only.

To account for the problem that the first prioritisation of issues could affect the second (due to maturation effects), a two-group design was used for the post-test, such that half of the participants should prioritise from the organisational perspective first, and the other half from the individual perspective first.

When prioritising both issues and uses, the participants should assign weights to the items, such that the weights should sum to 1 000. Thus, each weight could be seen as a certain percentage of the total importance (i.e., criticality for issues and relevance for uses) of all items. This is also known as the Hundred-Dollar method [7]. Advantages of this method are that it is easy to learn and use, and that the resulting weights are on a ratio scale.

4. Operation

In this section, we describe relevant parts of the operation of the study based on the design presented in the previous section.

4.1. Organisational Levels

We did not know the organisational levels of the participants prior to the interviews, since we chose not to map roles to levels directly. During the interview round, we let a manager estimate the levels of the participants, based on the descriptions of Anthony's decision levels as provided by Ngo-The and Ruhe [9]. The estimated levels were used to determine when it was likely that enough interviews had been performed to cover all levels. It was in the interest of the company to keep the number of interviews down to save resources.

The actual level used for a participant was determined by looking at the estimated and reported (i.e., from the post-test) levels, and, if necessary, by matching the participant's work tasks to Anthony's decision levels. As an example, consider a person with mismatching reported and estimated levels. If there is another person with the same role but matching levels, use his or her level for the first person. Otherwise, deduce the level by matching the first person's work tasks to the decision levels.

It was necessary to match work tasks to the decision levels for four of the participants. To increase the certainty in the level assignment, we successfully validated the levels of the remaining participants by matching their work tasks to the decision levels as well.

4.2. Interviews

A pilot interview was conducted at the company in order to measure the interview time and find discrepancies, if any, in the interview instrument. Since it resulted in only minor modifications to the interview instrument, the pilot was included in the analysis.

18 interviews were conducted, including the pilot, in the course of one month. The participants were sampled using *convenience sampling* [11], which in practise means that they were selected based on accessibility and recommendations from people at the company. We sampled based on convenience for two main reasons. First, we did not know prior to the interview which organisational level a person belonged to, and could consequently not sample based on that. Second, we argued that a person would be more committed to participate if he or she had been recommended by someone else.

The interviews were *semi-structured*, meaning that it was not necessary to follow the predefined question order strictly, and that the wording of questions was not seen as crucial for the outcome of the interviews [11]. The participant could speak rather freely, but the interviewer made sure that all questions were answered in one way or another.

The participants were asked if they would accept receiving and answering a post-test where they should prioritise both issues and uses. This was done in order to prepare the participants and increase their commitment towards the post-test.

A great variety of roles were covered in the interviews, including developer, tester, technical coordinator, manager (functional, product and project) and system architect. It should also be noted that the participants in general had been working at the company for a long time, and were thus familiar with processes and routines.

4.3. Prioritisation

Based on the complete lists of issues and uses, we constructed a post-test, as described in Section 3.5. The purpose of the post-test was to let the participants prioritise issues (from two different perspectives) and uses. To avoid maturation effects in the prioritisation of issues, we divided the participants into two groups based on their estimated organisational levels. Persons on each level were split at random between the two groups.

We required that the participants should specify a unique (non-tied) share for the level of their principal decisions. This way, we could deduce a non-ambiguous organisational level.

Since we were interested in potential issues and uses only, we asked the participants to prioritise without regard to actual issues and uses. In other words, we wanted to avoid priorities biased towards actual issues and uses.

5. Results

The mapping of participants to the organisational levels (described in Section 4.1) resulted in eight participants on the operative level, five on the tactical level and five on the strategic level.

The interviews resulted in 18 uses after irrelevant uses were removed and similar uses were merged together. These were subsequently combined with 11 uses found in the literature, which, due to overlap, resulted in 20 uses in total. For issues, there were 25 coming from the interview data. We found six issues mentioned in the literature, but these were already among the 25. Thus, the resulting list contained 25 issues in total.

Because of space constraints, we cannot show all issues and uses. The following list contains the issues (prefix *i*) and uses (prefix *u*) that are relevant for the analysis (see Section 6):
- **i1**: Hard to get resources for performing IA
- **i2**: Lack of time for performing IA
- **i3**: System impact is underestimated or overlooked
- **i4**: Unclear change requests

- **i6**: Analyses are incomplete or delayed
- **i8**: Analyses are too coarse or uncertain
- **i14**: Affected parties are overlooked
- **i15**: Analyses are performed by the wrong persons
- **i16**: Interest-based change request decisions
- **i19**: Not possible to see change request outcome
- **i22**: Cheap, short-term solutions win over good, long-term solutions
- **i23**: High levels specify solutions with too much detail
- **i24**: Hardware and protocol dependencies are difficult to handle for late change requests
- **i25**: Missing relevant structure and documentation to support the analysis
- **u1**: Planning the project with respect to time and cost
- **u2**: Determining cost versus benefit
- **u3**: Deciding whether to accept or reject the change
- **u6**: Understanding technical and market consequences of including or not including the change
- **u8**: Understanding the proposed change
- **u13**: Assessing system impact
- **u14**: Obtaining a new or changed requirements baseline
- **u20**: Revealing synergies and conflicts between change proposals

5.1. Threats to Validity

External validity is concerned with the generalisability of the results [13]. The small sample size and the fact that we used convenience sampling are threats to this type of validity. Nevertheless, as the participants were selected based on recommendations, we believe they were good representatives of their respective organisational levels. Furthermore, the fact that we focused on potential issues and uses rather than actual ones should increase the external validity. Also, the participants covered all uses but two and all issues from the literature, which indicates that their views of IA were not company-specific.

Construct validity is concerned with the design of the main study instrument and that it measures what it is intended to measure [11]. A threat to this type of validity is that the participants may not have had the desired mindset when prioritising items. As stated in Section 4.3, we asked the participants to prioritise as if neither issues nor uses currently were present, but we could not verify if they adhered to our request.

Internal validity is concerned with the relationship between the treatment and the outcome [13]. The assignment of participants to organisational levels is a threat to this type of validity. We tried to minimise the threat by basing the assignment on several information sources (see Section 4.1). The use of an external source (matching work tasks to decision levels) also strengthens the external validity.

6. Analysis and Discussion

This section describes two separate data analyses, one qualitative and one quantitative (statistical). We also comment briefly on the most interesting results.

6.1. Qualitative Analysis

In the qualitative analysis, we studied two aspects of the prioritised lists of potential issues and uses. First, we looked at the top five placements (which could include more than five items due to tied priorities) for each organisational level, in order to see if there was agreement on the most important items among the levels. Table 1 shows the top five placements for issues from an individual perspective (left), issues from an organisational perspective (middle) and uses (right). Top issues and uses common for two or more levels are displayed in bold text. It is clear that there was much agreement on the most important issues (regardless of perspective) and uses, although there were some differences mostly with respect to issues.

Note that no single issue was considered important by all three levels, while two uses were (u1 and u3). Moreover, the two issue perspectives differed somewhat. For example, issue i4 (*unclear change requests*) was seen as the most critical issue by the operative level from an individual perspective, but was not among the top five from an organisational perspective.

Table 1. Top Five Placements

Issues, ind.			Issues, org.			Uses		
O	**T**	**S**	**O**	**T**	**S**	**O**	**T**	**S**
i4	i24	**i15**	**i3**	**i14**	i6	**u3**	**u3**	**u1**
i22	**i14**	**i1**	i22	**i3**	i1, i15, i23	**u8**	**u8**	**u3**
i14	i3, i1, i19	**i3**	**i14**	**i1**		u13	**u1**	u14
i15		i2	i16	i4	i24, i25	**u2**	**u6**, u20	**u6**
i8		i6	**i15**	**i6**		**u1**		**u2**

Second, we compared the top five placements in each level with the bottom five in the other levels. The intention was to find items that were considered important by one level but unimportant by another level. None of the uses matched this criterion, and only a few issues from both perspectives did. This concurs with the observation that there was much agreement among the levels. Table 2 shows the ranks of the issues for each level, with negative ranks counting from the end. For example, issue i16 (*interest-based change request decisions*) had the fourth highest priority in the operative level, but the fourth lowest in the tac-

tical level. In the strategic level, it had neither a top five nor bottom five placement (empty cell).

6.2. Quantitative Analysis

In the quantitative analysis, we tested each issue and use for departure from normality by using the Shapiro-Wilk test for normality. The test showed that the data in general did not have a normal distribution. Therefore, we used the Kruskal-Wallis test for further analysis. The Kruskal-Wallis test is a non-parametric alternative to ANOVA, and should be used when there are more than two independent groups. At a significance level of 0.05, the test showed the following:

- There were neither significant differences among the three organisational levels for uses nor for issues from an organisational perspective.
- There were significant differences among the levels for two of the issues from an individual perspective. These are displayed in Table 3.

Table 2. Top Five vs. Bottom Five Issues

Perspective	Issue	O	T	S	Issue	O	T	S
Organisational	i16	4	-4		i25	-5		5
Individual	i19	-2	3	-1				

Table 3. Kruskal-Wallis for Issues i4 and i19

Issue	H	Sig.	Issue	H	Sig.
i4	6.934	0.031	i19	9.530	0.009

Outliers were not removed due to the relatively small sample size and the fact that there were many items to prioritise. We did, however, verify that the outliers for issues i4 and i19 were not responsible for the significant differences. There was a larger spread of priorities for issue i4 on the operative level than on the other levels. Similarly, the spread for issue i19 was larger on the tactical level than on the other levels. This also has to do with the fact these two issues were prioritised as important mainly by participants at the operative and tactical levels, respectively (see also Table 1).

6.3. Discussion

The two analyses both indicate that the participants, regardless of organisational level, had a coherent view of what IA could (or should) be used for. Top uses were *planning the project with respect to time and cost* (u1), *deciding whether to accept or reject the change* (u3) and *understanding the proposed change* (u8). We had, however,

expected *assessing system impact* (u13) to have a top five placement as well.

Similarly, both analyses show that participants on different levels saw different issues as critical (individual perspective), whereas on the whole, the awareness of each other's potential issues was large (organisational perspective). We were surprised that issue i19 (*not possible to see change request outcome*) was only considered critical by the tactical level from an individual perspective.

7. Conclusions

In this paper, we have presented an empirical study of views of impact analysis (IA) at three different organisational levels: operative, tactical and strategic. In the study, we interviewed 18 employees, representing industrial experts, at a large software development company in order to understand how potential issues and uses of IA are seen at the three levels.

The qualitative analysis shows that there were some differences among the levels with respect to top issues, regardless of whether the participants prioritised from an individual or an organisational perspective. On the whole, however, there was much agreement among the levels for both issues and uses. In other words, people at the different organisational levels seemed to mostly view IA in the light of what was important for all levels, not just their own. This means that our initial hypothesis, that we expected the views to diverge, cannot be fully supported.

The statistical analysis reveals some differences for issues as well, but only for the individual perspective. *Unclear change requests* (issue i4) was seen as more critical by the operative level, which is reasonable given that people at this level are the ones responsible for implementing changes. *Not possible to see change request outcome* (issue i19) was seen as more critical by the tactical level, which may be related to a need for the ability to follow up estimates.

We see the following practical implications of the results:

- The fact that there were some differences supports the relevance of looking at organisational levels when studying IA. We see Anthony's decision-making model as a good basis for the levels.

- Knowledge of issues associated with IA allows for focused improvement of the change management processes.

- The issues and uses presented in this paper may serve as a foundation for future research about the practical implications of the role of IA in change management.

With respect to the last bullet, we want to emphasise the importance of studying the non-technical aspects of IA in order to better understand how it is used by software practitioners in requirements engineering and change management contexts.

As can be seen in Section 5, neither the uses nor the issues are strictly orthogonal. The hierarchical dependencies that exist among them can potentially be used in future analyses to learn more about how the three organisational levels relate to each other.

8. Acknowledgements

9. References

[1] Aurum, A. and Wohlin, C., "The Fundamental Nature of Requirements Engineering Activities as a Decision-Making Process", in Information and Software Technology, 45:945–954, 2003.

[2] Bohner, S. A. and Arnold, R. S., Software Change Impact Analysis, IEEE Computer Society Press, 1996.

[3] Clements, P., Bachmann, F., Bass, L., Garlan, D., Ivers, J., Little, R., Nord, R. and Stafford, J., Documenting Software Architectures: Views and Beyond, Addison Wesley, 2003.

[4] Conradi, R. and Dybå, T., "An Empirical Study on the Utility of Formal Routines to Transfer Knowledge and Experience", in Proc. of ESEC/SIGSOFT FSE, 2001, pp. 268–276.

[5] Hall, T. and Wilson, D., "Views of Software Quality: a Field Report", in IEE Proc. on Software Engineering, 144:111–118, 1997.

[6] Karlström, D., Runeson, P. and Wohlin, C., "Aggregating Viewpoints for Strategic Software Process Improvement — a Method and a Case Study", in IEE Proc. on Software Engineering, 149:143–152, 2002.

[7] Leffingwell, D. and Widrig, D., Managing Software Requirements — A Unified Approach, Addison Wesley, 1999.

[8] Lindvall, M., An Empirical Study of Requirements-Driven Impact Analysis in Object-Oriented Systems Evolution, Ph.D. thesis, Linköping University, Sweden, 1997.

[9] Ngo-The, A. and Ruhe, G., "Decision Support in Requirements Engineering", in C. Wohlin and A. Aurum (eds.), Engineering and Managing Software Requirements, Springer-Verlag, to be published.

[10] Nicholas, J. M., Project Management for Business and Technology — Principles and Practice, 2nd ed., Prentice-Hall, 2001.

[11] Robson, C., Real World Research, 2nd ed., Blackwell Publishing, 2002.

[12] Sneed, H. M., "Impact Analysis of Maintenance Tasks for a Distributed Object-oriented System", in Proc. of the Intl. Conf. on Software Maintenance, Florence, Italy, Nov. 2001, pp. 190–195.

[13] Wohlin, C., Runeson, P., Höst, M., Ohlsson, M. C., Regnell, B. and Wesslén, A., Experimentation in Software Engineering: An Introduction, Kluwer Academic Publishers, 2000.