StakeNet: Using Social Networks to Analyse the Stakeholders of Large-Scale Software Projects

Soo Ling Lim*
Dept. of Computer Science
University College London
United Kingdom
s.lim@cs.ucl.ac.uk

Daniele Quercia
MIT SENSEable City Lab
Cambridge
USA
quercia@mit.edu

Anthony Finkelstein
Dept. of Computer Science
University College London
United Kingdom
a.finkelstein@cs.ucl.ac.uk

ABSTRACT

Many software projects fail because they overlook stakeholders or involve the wrong representatives of significant groups. Unfortunately, existing methods in stakeholder analysis are likely to omit stakeholders, and consider all stakeholders as equally influential. To identify and prioritise stakeholders, we have developed StakeNet, which consists of three main steps: identify stakeholders and ask them to recommend other stakeholders and stakeholder roles, build a social network whose nodes are stakeholders and links are recommendations, and prioritise stakeholders using a variety of social network measures. To evaluate StakeNet, we conducted one of the first empirical studies of requirements stakeholders on a software project for a 30,000-user system. Using the data collected from surveying and interviewing 68 stakeholders, we show that StakeNet identifies stakeholders and their roles with high recall, and accurately prioritises them. StakeNet uncovers a critical stakeholder role overlooked in the project, whose omission significantly impacted project success.

Categories and Subject Descriptors

D.2.8 [**Software Engineering**]: Requirements/Specifications

General Terms

Design, Experimentation

Keywords

stakeholder analysis, social networks, recommender systems

1. INTRODUCTION

Stakeholder analysis is the process of identifying stakeholders – individuals or groups that can influence, or be influenced by a software project – and prioritising them based on their influence in the project. For groups, we also need to identify suitable representatives to act on their behalf.

Many projects fail because they overlook stakeholders or involve individuals who lack knowledge, time or interest for adequate project involvement. Lack of user involvement is the main cause of project failure, and success is rare [20]. Reports suggest that 34% of projects succeed in 2004, 35% in 2006, and 32% in 2009 [21]. Overlooking stakeholders is possibly the most common mistake in development efforts [8]. As requirements are elicited from stakeholders, overlooking stakeholders gives rise to missing requirements, which leads to building the wrong product. For example, one company found that all the change requests for their software during the first year of operation came from stakeholders' needs that they had overlooked [6]. In some cases, the project can be delayed or cancelled: another company cancelled their project after realising that they overlooked a stakeholder, which would have cost them $20 million [8]. It is no wonder that finding stakeholders and the right representatives form 86% of developers' concern on stakeholder analysis [1].

Most, if not all, projects have resource limitations and timing constraints, with different stakeholders having different levels of influence. As such, it is necessary to prioritise stakeholders so as to focus our analysis on high-influence stakeholders while still being aware of low-influence ones, especially for large-scale software projects with hundreds or even thousands of stakeholders.

Although there is substantial literature on stakeholder analysis [1, 6, 8, 9, 15, 19], we find two problems. First, most methods treat all stakeholders as equally influential and all group representatives as equally suitable. Also, some methods are likely to overlook stakeholders, and others return "non-stakeholders". Second, the existing literature is almost entirely qualitative. Without empirical evaluations on real projects, we are uncertain how well a method performs against another. It is surprising given the centrality of stakeholders to software engineering that empirical studies have not been done before.

We address these problems by making two contributions.

- We propose StakeNet, a method to identify and prioritise stakeholders using social networks.
- We conducted an empirical study of the stakeholders of a large-scale software project at University College London (UCL). Our study is one of the first empirical studies on requirements stakeholders and we regard this as a major contribution of our work. We evaluate StakeNet against the existing method used in the

*Also with the University of New South Wales and NICTA.

Citation: S. L. Lim, D. Quercia, and A. Finkelstein. StakeNet: Using social networks to analyse the stakeholders of large-scale software projects. In Proceedings of the IEEE 32nd International Conference on Software Engineering, 2010 (in press).

project based on how well they identify and prioritise stakeholders. We analyse the empirical results with the contextual data gathered from interviewing the stakeholders, observing their engagement with the project, and reviewing project documentation.

The paper is organised as follows. Section 2 describes our case study. Section 3 reviews existing methods. Section 4 introduces StakeNet; Section 5 evaluates it. Section 6 identifies the limitations of our study, and Section 7 concludes.

2. THE RALIC PROJECT

The RALIC project was initiated to replace the existing access control systems at UCL and consolidate the new system with library access and borrowing, hence the name RALIC: *R*eplacement *A*ccess, *L*ibrary and *I*D *C*ard. RALIC was a combination of development and customisation of an off-the-shelf system. The project duration was 2.5 years and the system is now in deployment.

We select RALIC as our case study for two reasons. First, we have access to the stakeholders and project documentation as the system is developed, deployed, and maintained at UCL. Second, RALIC has a large and complex stakeholder base with more than 60 stakeholder groups, which is substantially larger than existing empirical studies in the area. Approximately 30,000 students, staff, and visitors use the system to enter buildings, borrow library resources, use the fitness centre, and gain IT access. Besides all UCL faculties and academic departments, RALIC also involves other supporting departments such as the Estates and Facilities Division that manages UCL's physical estate, Human Resource Division that manages staff information, Information Services Division, Library Services, Security Services, and so on, which must be identified and prioritised.

We use RALIC as a running example in the following sections to discuss existing methods and introduce StakeNet.

3. EXISTING METHODS

We classify existing approaches into four types: semi-structured, checklist-based, interviews, and search. We also broaden our search for related work to other areas.

Semi-structured approaches by Cockburn [6], Gause and Weinberg [8] largely form the basis of existing practice. They identify stakeholders by considering categories such as those who interact directly with the system (e.g., the students and the security guards) and those with interests (e.g., the government who imposes data protection policies). As the categories are broad, stakeholders such as the maintenance team may be overlooked.

Checklist-based approaches by Alexander and Robertson [1, 17] map generic stakeholder roles (e.g., maintenance operator and regulator) to project-specific stakeholder roles (e.g., the maintenance team and data protection officer). As the roles in the checklist are generic, project-specific roles such as the corporate communications division, which acts as a regulator that imposes rules on the branding and design of access cards, may be overlooked.

Interviews by Pouloudi and Whitley [15] consist of three steps: (1) identify generic stakeholder roles and the stakeholder (e.g., the sponsor is Alice, the director of Estates); (2) interview the stakeholder to learn about other stakeholders or stakeholder roles (e.g., Alice suggests the corporate communications division); and (3) add the newly identified stakeholders and their roles to the stakeholder list, and repeat Step 2 to interview them. Interviews are thorough but time-consuming for large projects with many stakeholders.

The *search* method by Sharp *et al.* [19] has two steps: (1) identify initial stakeholder roles from users, developers, legislators, and decision-makers; (2) for each stakeholder role R, identify other roles who interact with R, and repeat Step 2 for the newly identified roles. For example, students interact with staff, and staff interact with department administrators. Hence, staff and department administrators are stakeholder roles. *Search* identifies many roles but may return roles of non-stakeholders [19], such as prospective students who interact with department administrators.

In risk management, Woolridge *et al.* [23], and Glinz and Wieringa [9] prioritise stakeholders based on the risk incurred by neglecting them. As prioritisations are done from an individual's perspective, they can be biased for large projects in which no individual can have a global perspective. In software engineering, Mockus and Herbsleb [12] use data from change management systems to quantify development experience and suggest experts. Nevertheless, the suggestions focus on developers, who are a subset of stakeholders. In requirements management, Damian *et al.* [7] use social network analysis to explore collaboration and awareness among project team members. The focus is on their interaction rather than their identification. In natural resource management, Prell et al. [16] use social network analysis to inform their decision on which stakeholders to involve in environmental decision-making. Nevertheless, their analysis is qualitative without explicitly prioritising the stakeholders.

To summarise, existing methods have two shortcomings. First, some overlook stakeholders and others return non-stakeholders. Second, they assume stakeholders are equally influential by returning an unprioritised list of stakeholders and their roles.

An ideal method should identify and prioritise stakeholders and their roles from a global perspective. In doing so, it should extend the qualitative analysis of Prell et al.'s work [16] into *quantitative* analysis, and improve Woolridge *et al.*, Glinz and Wieringa's work [9] to be *independent* of the individual doing the analysis, and *scalable* for large projects. These are the features we set out to achieve with StakeNet described in the next section.

4. STAKENET

The idea behind StakeNet is to be open and inclusive, so that each stakeholder participates in the stakeholder analysis process. As stakeholders are socially related to one another, we can identify and prioritise them using their relations. StakeNet asks stakeholders to recommend other stakeholders, builds a social network with stakeholders as nodes and their recommendations as links, and prioritises stakeholders with social network measures. StakeNet does so in 6 steps (Figure 1) and relies on the concepts in Table 1.

Step 1: Determine the project scope. This is a preparatory step adopted from existing requirements engineering literature (e.g., [17]). Scope describes the boundary of the project so that we know which stakeholders should be involved. For example, the scope of the RALIC project includes installing new card readers throughout the university but excludes changing the existing systems in student residences. Hence, the Estates director is a stakeholder but the Student Residences director is not.

Figure 1: StakeNet's 6 steps.

Table 1: StakeNet Concepts

Concept	Definition
Scope	The work required for completing the project successfully [17].
Stakeholder	An individual or a group who can influence or be influenced by the success or failure of a project [14].
Stakeholder role	A part the stakeholder plays in the project.
Stake	An interest, investment, share, or involvement in the project, as in hope of gain.
Salience	The level of influence a stakeholder has on the project [11]. Stakeholders with high salience are crucial to project success; stakeholders with low salience have marginal impact.

Step 2: Based on the project scope, identify an initial set of stakeholder roles. Using the first step of the *search* method discussed in Existing Methods (Section 3), we identify initial stakeholder roles from the predefined categories of users, developers, legislators, and decision-makers. Users include students, staff, and security guards; developers include software vendor, and interface developer; legislators include data protection officer; and decision-makers include the director of Estates.

Step 3: For each role, find the stakeholders(s). If an individual role is suggested (e.g., director of Estates), the stakeholder is the person taking up the role (e.g., Alice), unless she nominates someone else. If a group is suggested, we use default representatives if they already exist (e.g., the group 'staff' is represented by the staff union representatives). Otherwise, we ask people who hold the role to nominate a representative (e.g., for the group 'security guards', we ask security guards to nominate their colleagues).

Step 4: For each stakeholder, get their stake and recommendations. The stake explains the particular way the stakeholder influences the project, or is influenced by it. For example, students depend on the system to access university buildings and library resources, and the funding director determines the project budget. A recommendation is a triple <stakeholder, stakeholder role, salience>, where salience is a number on an ordinal scale (e.g., 1–5). Salience comes from the literature on stakeholder management [11], where stakeholders are prioritised based on their possession of three attributes that constitute the stakeholder's salience: the *power* to influence the project, the *legitimacy* and *urgency* of their claims. For example, Alice makes a recommendation <Bob, Library, 4>.

To have a broad coverage of stakeholders and their roles, we repeat Step 4 for each new stakeholder, so that recommended stakeholders are, in turn, asked for further recommendations. For example, Bob makes a recommendation <Carl, Students, 1>, Carl then makes a recommendation <Dave, Students, 5>, and so on. The procedure is also known as the *snowballing technique* because the group of identified stakeholders builds up like a snowball rolled down a hill, which results in a well-connected network [18]. Eventually, few additional stakeholder roles are identified in each round of interviews. The snowballing stops when no additional roles are identified in one round of interviews.

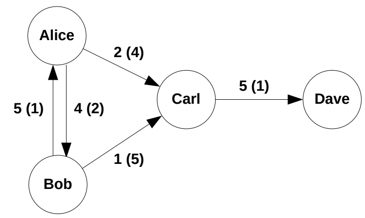

Figure 2: A network of stakeholders.

Stakeholders may make an incomplete recommendation <stakeholder role, salience> if they are unaware of actual stakeholders. If the role has existing stakeholders, then the recommendation is connected to *all* existing stakeholders. For example, Alice makes a recommendation <Students, 2>. Since Carl and Dave are recommended as students, we link them to Alice's recommendation. Otherwise, we go back to Step 3 to find stakeholders.

Step 5: Draw the social network. We draw a social network with the stakeholders as nodes, and their recommendations as directed edges: A links to B if A believes B to be a stakeholder. The edge weights is either the *salience* or the *reversed order of salience*, depending on the social network measure we use in Step 6. For now, we accept that the *reversed order of salience* is the salience subtracted from the upper bound of *maximum salience + 1* [4]. We explain this further when we introduce the measures. Figure 2 illustrates an example stakeholder network where Alice recommends Bob as salience 4 and Carl as salience 2; Bob recommends Alice as salience 5 and Carl as salience 1; Carl recommends Dave as salience 5; and Dave does not recommend anyone. The values in brackets are the *reversed order of salience* where the *maximum salience* is 5.

Step 6: Prioritise stakeholders by applying social network measures to the network. Given the social network in Step 5, we apply social network measures to prioritise the stakeholders. Four measures account for the whole network: *betweenness centrality*, *closeness centrality*, and *PageRank*; the remaining three account just for single nodes and their own connections: *degree centrality*, *in-degree centrality*, and *out-degree centrality*.

In Step 5, the edge weights for the social network depends on the social network measure. *PageRank*, *degree*, *in-degree*, and *out-degree* interpret edge values directly as strength of recommendation, hence they use salience as edge weights. In contrast, *betweenness* and *closeness centralities* interpret edge values as *lengths*. The interpretation means that *lower edge values* correspond to *shorter distances* between nodes, which imply *stronger ties* and higher salience [4]. Hence, high salience is reflected as low edge value, and low salience as high edge value. This explains why these measures use the *reversed order of salience* as weights.

As each measure ranks the nodes in the network differently, they prioritise different kinds of stakeholders. We explain how the measures prioritise a stakeholder S as follows.

- **Betweenness centrality** [3] ranks S based on S's ability to act as a broker between disparate groups of stakeholders. This measure sums the number of shortest paths between other pairs of stakeholders that pass through S. In Figure 2, Carl is ranked the highest because Alice and Bob both need to go through him to get to Dave. He is followed by Alice, as Bob can get to Carl directly or go through Alice, with equal distance of 5. Bob and Dave share the lowest rank as they do not appear between the shortest paths of other nodes.
- **Closeness centrality** [18] ranks S based on the inverse average shortest-path distance from S to all other reachable stakeholders. This measure prioritises stakeholders who reach others in the network quickly. For example, Carl ranks the highest because he reaches Dave with a high inverse average distance. He is followed by Alice ($\frac{3}{2+4+5}$), Bob ($\frac{3}{1+5+6}$), and Dave (0).
- **PageRank** [5] ranks S in terms of S's relative importance to all other stakeholders. It is used by Google to rank pages on the Internet. This measure is recursive in that stakeholders who are strongly recommended by many salient stakeholders are salient, and the recommendations of a highly salient stakeholder deserve more weight, which, in turn, makes their recommended stakeholders salient. Alice and Dave both have a recommendation of 5, but Alice ranks higher as she is recommended by Bob, who ranks higher than Carl. Carl has the lowest rank due to the low salience recommendations from Alice and Bob.
- **Degree centrality** [18] ranks S based on the number of incoming and outgoing recommendations S has and the weights of the recommendations. Alice ranks the highest as she has the most connections with most weights: $5 + 4 + 2$, followed by Bob: $5 + 4 + 1$, Carl: $1 + 2 + 5$, and Dave: 5.
- **In-degree centrality** [18] ranks S based on the number of stakeholders that recommend S and the recommendation weights. Dave and Alice rank the highest as they have a recommendation of weight 5, followed by Bob with a recommendation of weight 4, and Carl with two recommendations of weights 1 and 2.
- **Out-degree centrality** [18] ranks S based on the number of recommendation S makes and the weights of the recommendations. Alice and Bob rank the highest as they recommend two other stakeholders with a total weight of 6, followed by Carl: 5, and Dave: 0.

For each measure, StakeNet produces two lists (Figure 3):

- **Prioritised stakeholder roles**: The stakeholder roles are prioritised by the highest score of the stakeholder. If a stakeholder role appears in several rows with different stakeholders, we take the row with the highest rank.
- **Prioritised stakeholders**: For stakeholder roles with more than one stakeholder, the stakeholder with a higher rank is returned first.

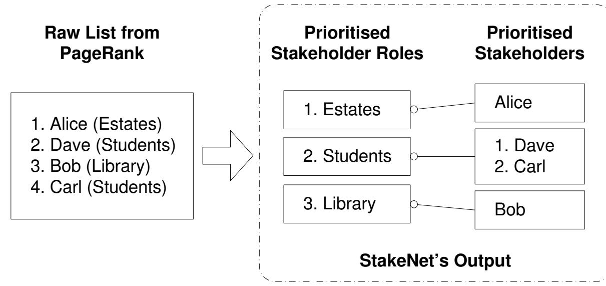

Figure 3: StakeNet's output using *PageRank*.

We first apply each measure independently. Then, we combine them to produce an improved prioritisation.

5. EVALUATION

The goal of StakeNet is to identify and prioritise stakeholders and their roles. Our research questions ask how well StakeNet identifies and prioritises stakeholders and their roles, whether the combination of social network measures produce better prioritisation, whether different methods for collecting recommendations affect the results, how well StakeNet performs against individual stakeholders, and whether stakeholders are motivated to make recommendations.

To evaluate StakeNet, we apply it to the RALIC project (Section 5.1). Then, we build the actual prioritised list of RALIC stakeholders and their roles, which we call the ground truth (Section 5.2). We also collect the list of stakeholders produced by the project team in a stakeholder elicitation workshop, where they used both the *semi-structured* and *checklist-based* approaches described in Section 3 to identify all the possible stakeholders. The outcome is an unprioritised list of 18 stakeholders and 28 roles, which we call the existing method list. Our evaluation compares the existing method list and the StakeNet lists against the ground truth (Section 5.3).

5.1 Applying StakeNet to RALIC

We applied StakeNet's six steps to the RALIC project described in Section 2. We used **Steps 1–3** as discussed in the previous section to arrive with the initial stakeholders and their roles. For the remaining three steps, we conducted a survey, which we report as follows.

Step 4. We emailed the stakeholders for a face-to-face survey to collect their recommendations separately. Each session took an hour on average, starting with the survey and ending with a semi-structured interview where we ask them for any relevant contextual information that can support our analysis, such as an account of their project engagement, the rationale behind their recommendations, and their confidence in their recommendations.

To start off the survey, we gave the respondents a cover sheet describing the survey purpose to identify RALIC stakeholders, the definition of stakeholders and salience as per Table 1, and the questionnaires. To prompt for recommendations, we provided respondents with six types of stakeholders as follows.

> A stakeholder can be someone who (1) *finances* the system, (2) *makes decision* about the development of the system, (3) *develops* the system,

Table 2: RALIC Project Scope

Item	Description
1	Replace magnetic swipe card readers with proximity readers
2	Source and install access card printers
3	Decide on card design and categories
4	Define user groups and default access rights
5	Interface the access control system with the library, human resource, student, and visitor systems
6	Issue new cards to the UCL community
7	Replace the library access control system
8	Use new cards at the UCL fitness centre

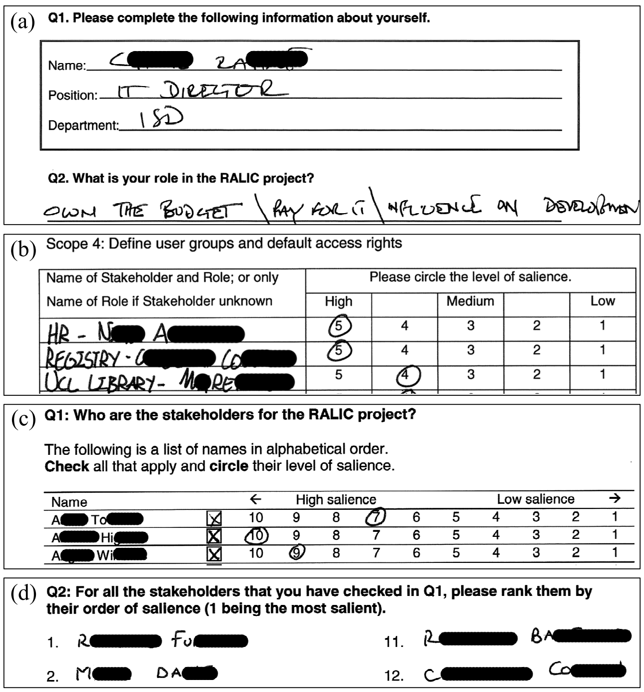

(a) Q1. Please complete the following information about yourself.

Name: C[redacted] ZA[redacted]

Position: IT DIRECTOR

Department: ISD

Q2. What is your role in the RALIC project?

OWN THE BUDGET \ PAY FOR IT \ INFLUENCE ON DEVELOPMEN

(b) Scope 4: Define user groups and default access rights

Name of Stakeholder and Role; or only Name of Role if Stakeholder unknown	High		Medium		Low
HR - N[redacted] A[redacted]	(5)	4	3	2	1
REGISTRY - C[redacted] CO[redacted]	(5)	4	3	2	1
UCL LIBRARY - M[redacted]	5	(4)	3	2	1

(c) Q1: Who are the stakeholders for the RALIC project?

The following is a list of names in alphabetical order.
Check all that apply and **circle** their level of salience.

Name		← High salience							Low salience →		
A[redacted] To[redacted]	☒	10	9	8	(7)	6	5	4	3	2	1
A[redacted] Hi[redacted]	☒	(10)	9	8	7	6	5	4	3	2	1
A[redacted] Wi[redacted]	☒	10	(9)	8	7	6	5	4	3	2	1

(d) Q2: For all the stakeholders that you have checked in Q1, please rank them by their order of salience (1 being the most salient).

1. R[redacted] FU[redacted] 11. R[redacted] BA[redacted]
2. M[redacted] DA[redacted] 12. C[redacted] CO[redacted]

Figure 4: Excerpts of stakeholder details, OpenR, ClosedR, and individual prioritisations.

(4) *imposes rules* on the development or operation of the system, (5) *uses* the system or its output, or (6) *threatens* the success of the system [1].

We collected their recommendations using questionnaires that consist of 4 parts[1] (Figure 4):

(a) **Stakeholder details**: Respondents provide their name, position, department, and role in the project.

(b) **Open-ended recommendations (OpenR)**: For each item in the RALIC project scope summarised in Table 2, respondents make recommendations in the form of <stakeholder, stakeholder role, salience>. OpenR must be completed before moving on to the closed-ended recommendations.

(c) **Closed-ended recommendations (ClosedR)**: Given the project scope, respondents select stakeholders from a checklist of names and circle their salience. The names in the checklist belong to stakeholders with the initial stakeholder roles. We also include the names of other UCL staff to introduce noise. There is an option to check *others* for suggesting stakeholders not in the checklist. We emphasised to the respondents that OpenR and ClosedR are separate questionnaires, and once they start on ClosedR, they cannot return to OpenR. We administered two distinct questionnaires to measure the effect of different survey methods on our results. We expect open-ended ones to be tedious to complete and closed-ended ones to be restrictive.

(d) **Individual prioritisations**: Respondents provide a prioritised list of who they think are the 20 most salient stakeholders. These lists enables us to compare individual prioritisations against the collective prioritisation from StakeNet.

[1]For the complete questionnaires, refer to http://www.cs.ucl.ac.uk/research/StakeNet/.

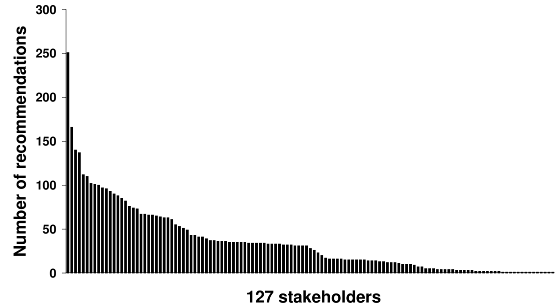

Figure 5: Longtail graph for OpenR.

We started with 15 roles and 22 stakeholders, surveyed 68 respondents and arrived with 1,789 open-ended recommendations, 839 closed-ended recommendations, and 57 individual prioritisations. The 68 respondents consist of 20 developers, 19 representatives of different user groups, 14 managers, 13 directors, and 2 legislators. From OpenR, we identified 127 stakeholders and 70 stakeholder roles, and from ClosedR, 76 stakeholders and 39 stakeholder roles. StakeNet expands role recommendations into their respective stakeholders in Step 4, giving us a total of 4,463 recommendations in OpenR.

Interestingly, the network of stakeholders resulting from OpenR is scale-free [2], i.e., the distribution of the number of edges is a power law as illustrated in Figure 5. A scale-free network suggests that OpenR is able to build a complete picture of the social network of stakeholders. In contrast, the network resulting form ClosedR is not scale-free: the distribution of the number of edges shows a long tail but is not a power law. This suggests that ClosedR builds only a partial view of the network of stakeholders.

Step 5. We build two networks, one from OpenR and one from ClosedR. As OpenR takes recommendations for eight separate scope items, we combine the recommendations to build a network for the whole project as follows. If A recommends B in N number of scope items, a combined recommendation is thus A recommends B as salience N. By combining the recommendations this way, we assume the scope items have equal weights. The maximum salience is 8 because there are 8 scope items.

Step 6. Finally, we apply social network measures to both the OpenR and ClosedR networks.

5.2 Building the Ground Truth

The ground truth is the complete and prioritised list of stakeholders ranked by their involvement in the project. After deploying the system, stakeholders who were overlooked would be uncovered from change requests, and the stakeholders' salience can be observed. We build the ground truth based on RALIC's project documentation, interviews, and observations in four steps.

Step 1. Based on the 6 types of stakeholders introduced in Section 5.1, we *derive subtypes of stakeholders* relevant to the RALIC project. For example, for stakeholders who *finances the development of the system*, the 3 subtypes are stakeholders who *finances the human resources*, *finances the hardware and software*, and *allocates server space*. From the 6 types of stakeholders, we derive 38 subtypes of stakeholders which we assume to be discrete and complete for RALIC.

Step 2. For each subtype, we *find actual stakeholders and their roles* from the project plan, functional specification, meeting minutes, maintenance action items, and post implementation report. We observe stakeholders on their present involvement and learn about their past involvement from interviews. For example, for the subtype *finances human resource*, we find from the meeting minutes that "Miles[2] confirmed Management Systems will provide long term support using existing permanent resources"; and "approval for additional resource granted by Fuller, Estates". Hence, Miles and Fuller are stakeholders, and their roles are Management Systems and Estates, respectively.

Step 3. We *rate each stakeholder* along their involvement in each subtype in terms of High, Medium, and Low. By doing so, we assume that stakeholders' salience is reflected in their project involvement. We rate across four types of involvement: (1) **finance**: a stakeholder who pays more has a higher rating; (2) **management**: a stakeholder who is more accountable has a higher rating; (3) **development**: a stakeholder with more responsibilities has a higher rating; and (4) **usage**: a stakeholder who uses the system more frequently, and is more affected by the system, has a higher rating. For example, in terms of **finance**, Fuller allocated the highest budget for human resources and hardware, hence he gets a High for both the subtypes *finances human resource* and *finances hardware and software*. Miles allocated around half of Fuller's amount for human resources, hence he gets a Medium for the subtype *finances human resource*.

Step 4. We *rank the stakeholders by the sum of their ratings*, from the highest to the lowest. To do so, we convert the ratings into numerical values (High = 3, Medium = 2, Low = 1). In our example, Fuller gets two High ratings so he has the score of $3 + 3 = 6$. He is ranked higher than Miles who has a Medium rating that converts to a score of 2. Overall, the output is a ranked list 85 stakeholders and 62 roles shown partially in Table 3. We then prioritise the stakeholders according to their roles, as in Figure 3.

Table 3: Partial Ground Truth List

Rank	Stakeholder	Role
1	Dawson	Security Services
2	Fuller	Estates and Facilities
3	Payne	Estates and Facilities
4	Crowe	Library Services
5	Cook	Student Registry

5.3 Method and Results

We now ascertain whether StakeNet can identify and prioritise stakeholders and their roles by answering the following research questions.

- **Identifying stakeholder roles:** (1) How many stakeholder roles identified by StakeNet are actual stakeholder roles? (2) How many of all the actual stakeholder roles does StakeNet identify? (Section 5.3.1)
- **Prioritising stakeholder roles:** How accurately do StakeNet's social network measures prioritise stakeholder roles? (Section 5.3.2)
- **Combining measures:** Does the combination of StakeNet's social network measures improve its accuracy in prioritising stakeholder roles? (Section 5.3.3)
- **Identifying and prioritising stakeholders:** (1) For each stakeholder role, how many of all the actual stakeholders are identified by StakeNet? (2) For each stakeholder role with more than one stakeholder, how accurately does StakeNet prioritise the stakeholders? (Section 5.3.4)
- **OpenR vs. ClosedR:** StakeNet uses open-ended recommendations (OpenR). What if it uses a predefined checklist (ClosedR)? (Section 5.3.5)
- **StakeNet vs. individual stakeholders:** In prioritising stakeholder roles, how does StakeNet compare to individual stakeholders? (Section 5.3.6)
- **Survey response:** Are stakeholders motivated to provide recommendations for StakeNet? (Section 5.3.7)

5.3.1 Identifying Stakeholder Roles

Method. We compare the stakeholder roles returned by StakeNet and the existing method against the ground truth, in terms of precision and recall [10]. The precision of identified stakeholder roles is the number of actual stakeholder roles in the set of identified stakeholder roles divided by the total number of identified stakeholder roles.

$$precision = \frac{\{X\} \cap \{GroundTruth\}}{\{X\}}, \tag{1}$$

where X is the set of stakeholder roles identified by StakeNet or the existing method, and *GroundTruth* is the set of stakeholder roles in the ground truth. The recall of identified stakeholder roles is the number of actual stakeholder roles in the set of identified stakeholder roles divided by the total number of actual stakeholder roles.

$$recall = \frac{\{X\} \cap \{GroundTruth\}}{\{GroundTruth\}}, \tag{2}$$

with X and *GroundTruth* same as for *precision*. Both precision and recall range from 0 to 1. Precision of 1 means all the identified roles are actual stakeholder roles. Recall of 1 means all the actual stakeholder roles are identified.

The stakeholder roles returned by the methods can be at a finer, equal, or coarser grain compared to the ground truth. If the returned stakeholder roles are at a finer grain, then we consider the results a match. For example, if ground truth

[2] Names have been changed for reasons of privacy.

Table 4: Identifying Stakeholder Roles

Method	Precision	Recall
Existing method	0.893	0.403
StakeNet	0.903	0.903

returns students, and StakeNet returns graduates, undergraduates, and medical students, we consider that StakeNet returns a stakeholder role that matches the ground truth. Otherwise, if the returned stakeholder roles are at a coarser grain than the ground truth, then we consider the results not a match. If they are of equal grain, we consider each stakeholder role individually. For example, for Security Services, ground truth returns the subgroup card issuer. StakeNet returns the subgroups card issuer and photographers. We consider photographers an error.

Results. StakeNet identifies most of the stakeholder roles in the ground truth with very few errors. From Table 4, StakeNet shows a recall of 90%, 50% higher than the existing method used in the project, and maintains a high precision of 90%. By using the existing method, the project team was only able to determine a limited number of stakeholder roles. As a result, they overlooked both high and low salience stakeholder roles, such as, the UCL Development and Corporate Communications Office that influences the access card design, the network team who provides network connectivity to the access control system, the access card vendor, short course students, external library users (library members not in other UCL systems), and the maintenance team. In contrast, by combining individual views of all stakeholders, StakeNet comes up with a more complete picture. It also successfully identifies the previously mentioned stakeholder roles which were overlooked by the project team.

Interestingly, overlooking external library users impacted RALIC's success. According to the post implementation report, "it was established that there were some 17,000 'external' library members who would consequently require new cards if the existing Library access control system were to be de-commissioned. The Board decided that this objective should be removed from the Project scope as the associated costs and complications of issuing such a high number of cards to non UCL members far outweighed the benefits." Had StakeNet been used, the project team would not need to wait till the end to find that they have overlooked external library users and caused the library to stay with the old access control system.

5.3.2 Prioritising Stakeholder Roles

Method. We compare only the StakeNet lists against the ground truth, but not the existing method list, as it is unprioritised. The accuracy of stakeholder role prioritisation is the similarity between the prioritisation of the identified stakeholder roles and the prioritisation in the ground truth. We use Pearson's correlation coefficient, ρ, to determine the similarity [10]: ρ ranges from -1 to $+1$, -1 means the two lists are negatively correlated, 0 means no correlation, and $+1$ means positively correlated. A positive ρ means that high priorities in the ground truth list are associated with high priorities in the list of identified stakeholders. The closer the values are to -1 or $+1$, the stronger the correlation. To compute ρ, we need lists of the same size. Therefore, we take the intersection between the lists.

Table 5: Prioritising Stakeholder Roles

Measure	Accuracy
Betweenness	0.785
Closeness	0.628
PageRank	0.716
Degree	0.686
In-degree	0.653
Out-degree	0.625

Results. StakeNet prioritises stakeholder roles with high accuracy. According to Table 5, measures that consider the whole network have a higher accuracy in prioritising stakeholder roles than measures that consider direct connections. *Out-degree* has the lowest overall accuracy, as stakeholders who make a lot of recommendations are not necessarily salient stakeholders. *Closeness* also has low accuracy as stakeholders who can reach other stakeholders quickly are not necessarily salient. We also find that *degree* correlates strongly with *betweenness* (measuring the correlation between the lists produced by *degree* and *betweenness* gives us $\rho = 0.831$), a common feature in social networks [13].

Two StakeNet features influence the accuracy of prioritisation. First, StakeNet prioritises stakeholder roles based on the highest rank of the stakeholders with that role, which produces inaccurate prioritisation when a stakeholder has more than one role. For example, a stakeholder with two roles receives a recommendation every time either role is recommended. Hence, both roles receive higher prioritisation than if different stakeholders play the two roles. Second, the expansion of role recommendations into existing stakeholders with the same role also affects the accuracy because respondents who recommend roles appear to make many recommendations. For example, the new data protection officer was ranked first by *out-degree* as he was unfamiliar with actual stakeholders and suggested many roles.

5.3.3 Combining Measures

Method. We perform two preliminary studies to combine different measures. First, we predict the suitable measure for each stakeholder role. Each role has two Boolean attributes gathered from interviews: `is_confident` and `is_community`. `is_confident` indicates whether the stakeholder expresses a lack of confidence in his or her recommendations. `is_community` indicates whether the stakeholder belongs to any cliques within the organisation, as stakeholders in cliques may recommend one another more frequently and skew the results. We train a C4.5 (J48) decision tree classifier [22] to find the measure that produces the closest result to the ground truth for each stakeholder role from the three most accurate and diverse measures: *betweenness*, *in-degree*, and *PageRank*.

Second, we predict the stakeholder role's actual rank in the ground truth. Each role now has five attributes: `is_confident` and `is_community` as before, `has_turnover`, `betweenness`, and `pagerank`. `has_turnover` is a Boolean indicating whether staffing changes occurred among the stakeholders of that role. `betweenness` and `pagerank` are the ranks of the stakeholder role produced by *betweenness* and *PageRank*. We use these measures as they are the most accurate measures. We train a Gaussian Process classifier [22] to predict each stakeholder role's actual rank.

Results. Results from 10-fold cross-validation [22] show that both classifiers produce more accurate prioritisation

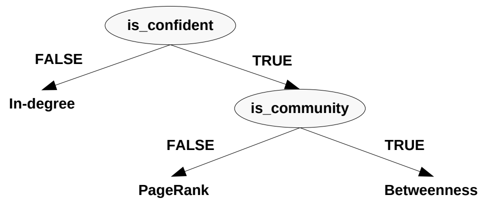

Figure 6: Decision tree classifier.

Table 6: Identifying and Prioritising Stakeholders

Method/Measure	Recall	Accuracy
Existing method	0.465	N/A
Betweenness	0.920	0.517
Closeness	0.907	0.564
PageRank	0.867	0.778
Degree	0.907	0.649
In-degree	0.907	0.659
Out-degree	0.907	0.565

Table 7: OpenR vs. ClosedR

7.1 Identifying Stakeholder Roles

	Precision	Recall
OpenR	0.903	0.903
ClosedR	0.949	0.597

7.2 Prioritising Stakeholder Roles (Accuracy)

	btw	close	pgrank	deg	in-deg	out-deg
OpenR	0.785	0.628	0.716	0.686	0.653	0.625
ClosedR	0.712	0.467	0.767	0.802	0.786	0.713

7.3 Identifying Stakeholders (Recall)

	btw	close	pgrank	deg	in-deg	out-deg
OpenR	0.920	0.907	0.867	0.907	0.907	0.907
ClosedR	0.823	0.839	0.806	0.855	0.806	0.855

7.4 Prioritising Stakeholders (Accuracy)

	btw	close	pgrank	deg	in-deg	out-deg
OpenR	0.517	0.564	0.778	0.649	0.659	0.565
ClosedR	0.427	0.278	0.910	0.956	0.918	0.366

than the most accurate individual measure (*betweenness* with $\rho = 0.785$). The decision tree classifier has an accuracy of $\rho = 0.812$, and the Gaussian Process classifier has a higher accuracy of $\rho = 0.846$. Figure 6 shows the learned decision tree. We learn that *betweenness* avoids stakeholder communities who highly suggest one another from gaining disproportional influence, *PageRank* avoids salient stakeholders recommended only by a few salient stakeholders from losing influence, and *in-degree* avoids salient stakeholders from being penalised for their answers. Nevertheless, other attributes may be more appropriate for other projects and organisations.

5.3.4 Identifying and Prioritising Stakeholders

Method. We compare the lists of stakeholders returned by StakeNet and the existing method against the ground truth, in terms of recall. The recall of identified stakeholders is the number of actual stakeholders identified divided by the total number of stakeholders in the ground truth, for each stakeholder role that appear in both lists, with the following rule. If the ground truth has N stakeholders with the same role, we consider only the first N stakeholders of the method. This is because if the project consults one stakeholder and the method returns the correct stakeholder after two incorrect ones, the project would have consulted the less suitable stakeholder. We use the recall formula (2) in Section 5.3.1, with X as the set of identified stakeholders by StakeNet or the existing method, and $GroundTruth$ as the set of stakeholders in the ground truth.

For stakeholder roles with more than one stakeholder, we also prioritise the stakeholders according to their suitability to represent the role. The accuracy of stakeholder prioritisation is the similarity between the ordering of the identified stakeholder and their actual ordering in the ground truth. We use Pearson's correlation coefficient on the stakeholder list for each stakeholder role. The total accuracy is the average accuracy for all such roles. Again, we measure only the StakeNet lists as the existing method list is unprioritised.

Results. StakeNet identifies stakeholders with high recall. *Betweenness* has the highest recall in identifying stakeholders, 45% more than the existing method (Table 6). In line with the literature, these measures prioritise stakeholders whose positions in the network allow them connect to different groups and have more knowledge [18], which are the characteristics of good stakeholders. In the existing method, many roles were identified rather than actual stakeholders. Although some stakeholder roles, such as the software vendor, need not have predefined stakeholders, those involved in requirements elicitation do. If we involve unsuitable stakeholders, we risk eliciting the wrong requirements.

StakeNet prioritises stakeholders with the same role accurately. *PageRank* has the highest accuracy of 78% (Table 6). StakeNet produces high accuracy in prioritising stakeholders as it prioritises them based on their relations with other stakeholders rather than their formal positions. Just from looking at the organisational chart for Security Services, one may assume that the Head of Security Services is the most suitable stakeholder to talk to. Nevertheless, the manager is more involved in RALIC and connected to the other stakeholders, hence is more suitable. StakeNet correctly identifies this. In contrast to high accuracy in prioritising *stakeholder roles*, measures considering network connectivity have low accuracy in prioritising *stakeholders*. Their accuracy is low for roles with recent turnover, as the new stakeholders do not share the same ties in the network as the old.

Besides knowing who the suitable stakeholders are, it is also useful to know about potential problems with a stakeholder's involvement. Inspired from the social network literature [13], we find that comparisons between different measures that are strongly correlated reveal two such problems. First, stakeholders who have a high rank in *degree* but low rank in *betweenness* tend to have high influence but low involvement. An example is the Head of Security Services – many security related issues discussed in meetings required his input but he was absent. Second, stakeholders who have a high rank in *betweenness* but low rank in *closeness* are salient stakeholders who are often out of the loop.

5.3.5 OpenR vs. ClosedR

Method. We compare the lists produced from OpenR against the lists produced from ClosedR, using precision, recall, and accuracy as before.

Results. In identifying stakeholders and their roles, ClosedR has less recall than OpenR, 30% less recall in identifying roles, and an average of 8% less recall in identifying stakeholders (Table 7). Checklists limit the discovery of project-specific stakeholder roles as they constrain survey outcomes around the given options. Although respondents were encouraged to recommend stakeholders not in the checklist, only 34% of them did so, with an average of less than three recommendations each. As fewer new suggestions reduce the likelihood of error, ClosedR has 5% more precision than OpenR in identifying stakeholder roles.

In prioritising stakeholders and their roles, measures that consider direct connections have a higher accuracy for ClosedR. Recommendations revolve around the names in the checklist rather than from the respondents' past interactions. In contrast to OpenR where most measures return the security service owner and the project manager as globally high ranked stakeholder, measures in ClosedR return division heads and provosts. This is because in OpenR, respondents who lack confidence recommended groups, but in ClosedR, they checked the names of division heads and provosts.

5.3.6 StakeNet vs. Individual Stakeholders

Method. We compare the individual prioritisations provided by each stakeholder against the ground truth, in terms of accuracy as in Section 5.3.2. As before, we take the intersection between the lists to produce lists of the same size.

Results. Prioritisations by individual stakeholders are less accurate compared to the collective prioritisation by StakeNet. The 57 individual prioritisations have an average accuracy of $\rho = 0.360$, less than half compared to StakeNet's most accurate prioritisation from *betweenness centrality* ($\rho = 0.785$). Managers have the highest accuracy ($\rho = 0.514$), followed by directors ($\rho = 0.404$), developers ($\rho = 0.317$), users ($\rho = 0.240$), and then legislators ($\rho = 0.163$). The prioritisations are biased by local perspective. Developers tend to prioritise their immediate managers and other developers; managers prioritise other managers.

5.3.7 Survey Response

Method. We calculate the response rate of our survey as the number of stakeholders who responded, over the total number of stakeholders we contacted, expressed as a percentage. We also look at the precision and recall of stakeholder roles as the snowball progresses, to know the quality of the identified roles had we surveyed fewer stakeholders.

Results. Stakeholders were motivated to recommend other stakeholders. Our survey response rate was 81%, which is 30% higher than the weighted average response rate without regard to technique[3] [24]. Those who did not respond to our email request were mostly away, not stakeholders, or have low salience.

During our survey, the recall increases quickly at the start and then stabilises, and the precision lowers consistently throughout (Figure 7). As the snowball progresses towards the boundary of the network, fewer additional stakeholder roles are identified in each round of interviews. StakeNet terminates after Round 4 of surveys, as no new stakeholder roles are identified. After Round 4, new stakeholder roles

[3]In the study by Yu and Cooper, the sample sizes for 497 response rates from various survey methods (e.g., mail surveys, telephone surveys, personal interviews) varied from 12 to 14,785. As such, the response rate averages are weighted by the number of contacts underlying the response rate.

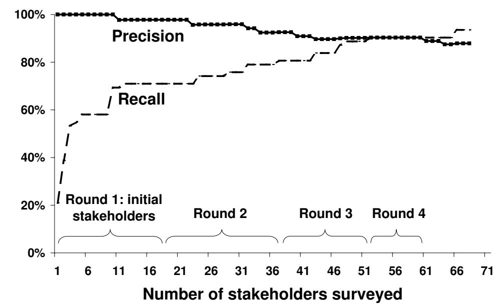

Figure 7: Precision and recall as the snowball rolls.

can still be identified, but at a slower rate. The increase in recall at the end was from an initial stakeholder who was unavailable at the start of our survey. The graph shows that there is always a tradeoff between precision and recall.

Respondents involved in software projects are keen to use StakeNet in their work. During the survey, a system support asked eagerly, "Will UCL use StakeNet?". He explained that many of the change requests he received can be avoided if stakeholders were identified initially. A director recommended StakeNet to be used in his division. He added: "Managers often think in terms of systems and organisation. When we have a function, we only look at the layer in the organisation affected by the function. This limited view often gives us problems in the long run!"

6. THREATS TO VALIDITY

Our study is based on a single project, hence there must be some caution in generalising the results to other projects and organisations. As we conducted the survey after project completion, post-project knowledge may influence the results. The respondents may learn about the missing stakeholders during the project and recommend based on the knowledge. Nevertheless, only 15% of the respondents were involved throughout the project due to staff turnover and department restructuring, hence their influence on the overall prioritisation is low. Also, as RALIC is completed four years ago, it is unlikely for recommendations to come from memory or daily conversations. Still, to confirm our results, future work should evaluate StakeNet on different projects in different organisations.

Misunderstandings during the survey can affect recommendations and the results. Some respondents have a different understanding of stakeholders due to multiple existing definitions [19], and some find salience and role difficult to grasp. For roles, respondents who are not actively involved in the project put their job titles instead of why they are stakeholders. Feedback from respondents reveals that the types of stakeholders provided as prompts should have included the categories: *manages the development of the system*, *provides input to the system*, and *maintains the system*.

We clean the data provided by respondents in two ways: (1) merge synonymous stakeholder roles (e.g., research students and PhD students), (2) merge different names referring to the same person (e.g., Nicholas is sometimes recommended as Nic or Nick). Manual merging of roles may be inconsistent and subjective (e.g., one may consider research students and PhD students as different roles), which then influences the results of role prioritisation. Future work should aim for more objective and ideally automated cleaning.

In Step 2 of StakeNet where we identify initial stakeholders, one can say that knowledge about the stakeholders (e.g., from the ground truth or the project documentation) affects the determination of the initial stakeholders, and the quality of the final set of stakeholders. Nevertheless, our initial set is constructed *before* the ground truth. Also, as we use established approaches to identify stakeholder roles from key stakeholder categories (users, developers, legislators, and decision-makers), our approach converges and is relatively insensitive to the initial set.

One can say that the ground truth is biased in our perspective, thus affecting the results of our study. Nevertheless, we argue that it is representative of the actual stakeholder list because we have the global perspective from reviewing project documentation, observing the stakeholders' engagement with the project, and interviewing them. We also validated the ground truth with management level stakeholders. Still, future work should consider alternative ways of constructing the ground truth to increase its objectiveness.

7. CONCLUSION

StakeNet uses social networks to identify and prioritise software project stakeholders and their roles. By applying StakeNet to the UCL access control project, we demonstrate that StakeNet performs better than the existing method used in the project, with a higher precision and recall in identifying stakeholders and their roles. Different social network measures prioritise different kinds of stakeholders, and StakeNet outperforms individual stakeholders.

StakeNet prioritises stakeholders based on the sum of their salience over the project lifecycle. Future work should consider prioritising over time to accommodate changing scope, stakeholders, roles and salience; prioritising across different issues, such as funding, development, and usage; and using multi-dimensional salience, such as power, legitimacy, and urgency. We have developed a software tool that implements StakeNet, which is now used in real software projects. Hence, our next step is to measure StakeNet's performance in practice. Finally, we are strongly persuaded that further empirical studies of stakeholders are called for.

In a broader context, we propose a new methodology that shifts the emphasis from stakeholder analysis run by a small number of experts to a collaborative approach in which all stakeholders have a say.

8. ACKNOWLEDGMENTS

We thank RALIC stakeholders for their participation, Neal Lathia, Licia Capra, Emmanuel Letier, Eric Platon, and Andy Maule for their discussions and feedback, and Jennifer van Heerde for her survey advice.

9. REFERENCES

[1] I. Alexander and S. Robertson. Understanding project sociology by modeling stakeholders. *IEEE Software*, 21(1):23–27, 2004.

[2] A. L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509, 1999.

[3] U. Brandes. A faster algorithm for betweenness centrality. *Journal of Mathematical Sociology*, 25(2):163–177, 2001.

[4] U. Brandes. On variants of shortest-path betweenness centrality and their generic computation. *Social Networks*, 30(2):136–145, 2008.

[5] S. Brin and L. Page. The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems*, 30(1-7):107–117, 1998.

[6] A. Cockburn. *Writing Effective Use Cases*. Addison-Wesley, 2002.

[7] D. Damian, S. Marczak, and I. Kwan. Collaboration patterns and the impact of distance on awareness in requirements-centred social networks. In *Procs. of the 15th Int. Req. Eng. Conf.*, pages 59–68, 2007.

[8] D. C. Gause and G. M. Weinberg. *Exploring Requirements: Quality Before Design*. Dorset House Publishing, 1989.

[9] M. Glinz and R. J. Wieringa. Stakeholders in requirements engineering. *IEEE Software*, 24(2):18–20, 2007.

[10] J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl. Evaluating collaborative filtering recommender systems. *ACM Trans. on Information Systems*, 22(1):5–53, 2004.

[11] R. K. Mitchell, B. R. Agle, and D. J. Wood. Toward a theory of stakeholder identification and salience. *Academy of Management Review*, 22(4):853–886, 1997.

[12] A. Mockus and J. D. Herbsleb. Expertise browser: a quantitative approach to identifying expertise. In *Procs. of the 24th Int. Conf. on Soft. Eng.*, pages 503–512, 2002.

[13] M. E. J. Newman. A measure of betweenness centrality based on random walks. *Social Networks*, 27(1):39–54, 2005.

[14] B. Nuseibeh and S. Easterbrook. Requirements engineering: a roadmap. In *Procs. of the Conf. on the Future of Soft. Eng.*, pages 35–46, 2000.

[15] A. Pouloudi and E. A. Whitley. Stakeholder identification in inter-organizational systems. *European Journal of Information Systems*, 6(1):1–14, 1997.

[16] C. Prell, K. Hubacek, and M. Reed. Stakeholder analysis and social network analysis in natural resource management. *Society & Natural Resources*, 22:501–518, 2009.

[17] S. Robertson and J. Robertson. *Mastering the Requirements Process*. Addison-Wesley, 2006.

[18] J. Scott. *Social Network Analysis: A Handbook*. Sage, 2000.

[19] H. Sharp, A. Finkelstein, and G. Galal. Stakeholder identification in the requirements engineering process. In *Procs. 10th Int. Workshop on Database and Expert Systems Applications*, pages 387–391, 1999.

[20] Standish Group. *The CHAOS Report*, 1994.

[21] Standish Group. *CHAOS Summary 2009*, 2009.

[22] I. H. Witten and E. Frank. *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann, San Francisco, 2005.

[23] R. W. Woolridge, D. J. McManus, and J. E. Hale. Stakeholder risk assessment: An outcome-based approach. *IEEE Software*, 24(2):36–45, 2007.

[24] J. Yu and H. Cooper. A quantitative review of research design effects on response rates to questionnaires. *Journal of Marketing Research*, 20(1):36–44, 1983.