

Troubleshooting Collaborative Ontology Design

Angela Gomes², Adriana S. Vivacqua¹, Ana Cristina B. Garcia^{1,2}

¹ADDLabs, UFF – Fluminense Federal University, Rio de Janeiro, Brazil

²Computing Institute, UFF – Fluminense Federal University, Rio de Janeiro, Brazil
avivacqua@addlabs.uff.br, bicharra@ic.uff.br, angela@addlabs.uff.br

Abstract

This paper presents observations from engineering ontology design sessions. We followed the entire design process of a group of engineers building an ontology for root cause analysis to be applied to petroleum pump failure diagnosis. The group was composed of experts with same background, who work in different locations of the same company. Even though they apparently share similar analysis methods, during meetings discrepancies emerged. Although they officially agreed to each element, they didn't accept the final representation whenever the concept was not fully discussed. Given the problems identified in the final design, we look back at the meetings and attempt to identify what led to these problems. This enables us to draw a set of guidelines for dealing with design meeting problems. Our observations point towards a relationship between an actor's behavior and future artifact acceptance; a set of actions that interrupts or brings back group attention; and a possible metric for evaluating group design meetings. We believe our findings may guide software developers to tools for supporting group design.

Keywords: CSCW, Extreme collaboration, Ontologies, design meetings, Collaborative design, Internet.

1. Introduction

Nowadays, organizations must deal with complex problems, which demand discussion and decision making by a group of people. These individuals usually come together to discuss problems and explore possible solutions, focusing on the best ones.

This also holds true in collaborative design meetings, where a group of experts comes together to create a solution to a design problem. Large projects, such as oil platform design or urban planning, frequently require a team working together, extensive discussions around possible solutions and their evaluation, enumeration of pros and cons and selection of the most appropriate alternatives. Because design is an open-ended activity (i.e., there is no right answer, but there are incorrect alternatives), this process becomes even more important.

Extreme collaboration [2] is a design technique that emphasizes work in an electronic and social environment to maximize communication and information flow. Design teams work in war rooms, where all designers are collocated and may interact as needed. In extreme collaboration, participants have access to networked computers and all the necessary information, and a coordinator leads the group in its exploration of the design space.

We have been successfully applying this technique for ontology design for over a year now. In our design meetings, engineers come together to create an ontology for a given domain and a given purpose. Recently, one of our designs was poorly accepted, despite having gone through the full process and been discussed and agreed upon by the different stakeholders. This led us to an investigation of possible causes for the acceptance problem.

Despite its importance, collaborative design is still poorly supported by computational tools. Manipulation of shared artifacts, generation and evaluation of alternatives and extensive discussion suffer from distance. With our investigations into the design process, we expect to generate requirements and specifications for design support tools, especially where solution generation and evaluation come into play.

Face-to-face interactions are very important for collaborative design: collocated engineers become immediately aware of specification changes, can quickly question or adapt to changing requirements, and can easily interact with one another whenever necessary. However, distance and travel costs make the investigation of alternatives to collocated meetings a necessity.

We are currently engaged in an effort to design a remote meeting environment to support ontology design. To that end, we have been conducting activity-theory-based analyses of meetings, to better understand what the actions and operations are and how they combine to reach the desired outcomes. Our approach is to analyze the process in reverse: for a given (finished) artifact, we analyze its acceptance levels and problems found at the final delivery presentation and move back through meetings verifying what went wrong for each problematic element.

This paper is organized as follows: in the next section we discuss ontology construction, followed by a

description of the methodology used. We present our observations in section 4 and finalize with a discussion in section 5.

2. Ontology Construction as Design

The design activity consists of defining a specification for an artifact that delivers a set of desired functionalities. Both for physical artifacts such as a building and abstract artifacts such as a domain ontology, there is no unique result. Consequently, different designers may reach different designs (sometimes multiple designs). Collaborative design may lead to better results due to the synergy among the participants. However, conflicts often emerge during design sessions, which may produce delays. On the other hand, conflict can be productive, as discussions promote a broader investigation of the design space.

Project deadlines frequently constrain the viable amount of discussion (and exploration of the domain), while participants' individual levels of acceptance of the group decision expand it. The lower the acceptance level, the more discussion and exploration will happen. In this fashion, the design process hangs on a balance of time versus acceptance levels, which is usually managed by the group's coordinator or meeting facilitator.

An ontology is a description of a domain, constructed for a given purpose by a group of stakeholders that will use it. The ontology construction process is a knowledge creation activity, through which known information is transformed into a final representation. It is a group design process, involving several people who bring distinct knowledge and viewpoints to the process. Given the multiple perspectives, this process frequently involves negotiation of scope, term definition and usage. In this setting, the knowledge acquisition activity becomes a design activity and a knowledge engineer becomes a facilitator of a group activity.

In extreme collaboration environments, the meeting coordinator elicits information from experts and submits it to discussion. He or she asks questions that lead to the introduction of new concepts and challenges the concepts introduced. When the ontology is starting to stabilize, the coordinator should steer the group towards decision making.

Given that experts didn't have much time to spend on collocated meetings, we introduced take-home activities into the equation. Activities that could be isolated were isolated and assigned to participants as activities to be performed at home. This would help the process progress in between sessions. Therefore, the coordinator must plan individual activities to keep the specialists in touch with the ontology construction activity and create engagement with the task. This creates a sense of

ownership which helps increase acceptance of the ontology at a later date.

3. Methodology

In this section we briefly present activity theory and the method used for meeting analysis through video reviews.

3.1. Activity Theory

Activity theory is a framework to help describe how work is done. The framework takes an individual (subject) as the starting point for analysis. The *subject* manipulates an *object* using *tools*, to reach a desired *outcome*. In the framework, work is analyzed in three hierarchical levels: at the topmost level, an activity is motivated by a goal and executed by a community of people. This activity is broken down into actions, executed either by an individual or by a group, that reach specific goals. Each action is accomplished through operations, which are deeply ingrained into actors and are executed almost automatically, given certain conditions. They can be either automated by a machine or executed "without thinking" by a human.

Activity theory has been successfully applied in a number of situations [1]. However, activity theoretic studies usually follow the route of observing the actions and operations contribute to the execution of the activity. We take the reverse approach: given an evaluation of the outcome of the activity, we extract the problematic points and go back to study what went wrong. We expect this line of analysis will shed more light on the process, emphasizing the operations that work against the completion of the activity. This, in turn, should help us design ways to prevent them from happening.

3.2. Video Reviews and Coding

Every design meeting that happens in our facility is completely videotaped and transcribed. These tapes and transcripts can be used as a source for analysis. We use dialogue analysis and coding as our instruments. Every meeting is being coded according to a number of tags, defined from an initial observation of meeting situations.

Initially, tapes were reviewed to generate a set of tags for meeting situations. After verification that these would cover most situations, we started a review of the videotapes and transcripts, coding each contribution to the discussion in turn. This allows us to perform calculations on tags occurrence and to look for correlations between these and events. An ontology of the tags used can be seen in Figure 1.

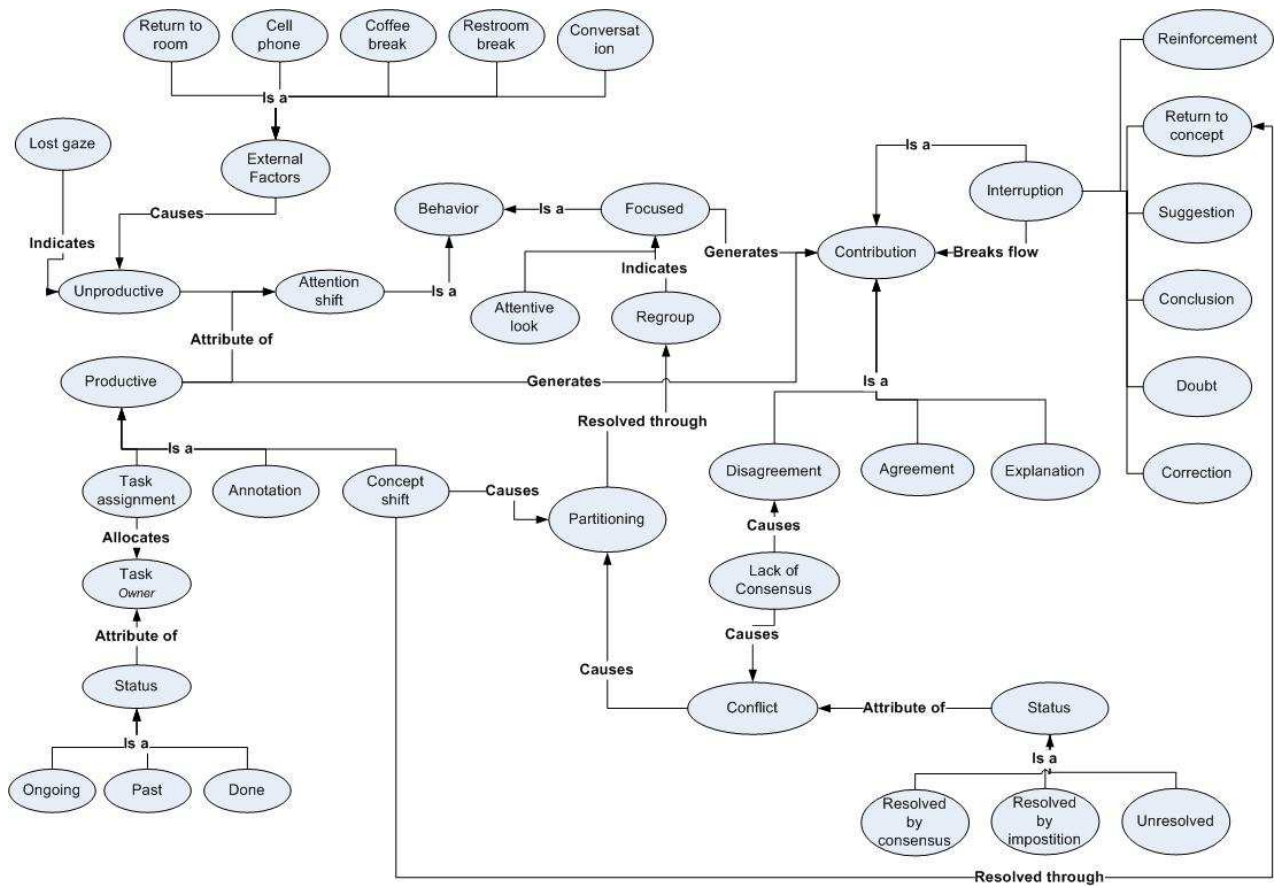


Figure 1: Tag ontology

The ontology describes each participant’s behavior as focused or unfocused (shift in attention). Unfocused behavior can be of two natures: unproductive (does not contribute to the ongoing activity) or productive (even though it doesn’t actively contribute to the discussion, it is related to the activity and may help at a later stage.) The indication of unproductive attention shift most frequently noticed was lost gaze, or when the participant stopped directing his vision to the main discussion. A number of factors causes the shift to an unproductive state. Unfocused productive behaviors include working on a separate task, making annotations or branching the discussion into a different concept.

Focused behaviors generate contributions to the discussion, which may be explanations of concepts or ideas or statements of agreement or disagreement. Interruptions cut the meeting flow and add to the discussion. Disagreements stem from a lack of consensus, which causes conflicts that may also lead to group partitioning. Partitioning is resolved by a regroup, with members going back to the theme at hand.

3.3. Setup

To conduct this research, we took video recordings of design meetings for a finished project. The recordings have been fully transcribed, which makes it

easy to search for key phrases and particular discussions.

We have the final products, namely, the ontology (represented as a set of concepts and relations) and the concept definitions. These are accompanied by a set of annotated slides, which were presented to the client at a final meeting. During this meeting, clients pointed out a number of problems and did not accept the final design, requesting a number of changes. A list of problems pointed out by the clients is also part of this study. The following steps are being conducted in this research:

1. Identify problem concepts: concepts are listed, tagged as problematic or non-problematic and a list of problems noted is written down next to each one.
2. Search for discussion of these concepts: each problematic concept is searched for in the transcripts (we look for where it was first discussed and for every discussion following that event.)
3. Analyze discussion: once found, each discussion involving the concept is tagged according to the ontology shown above. We seek to identify the conditions that might have caused the decline in acceptance of the final ontology.
4. Analyze non-problematic concepts: once problematic concepts have been analyzed,

we turn to non-problematic concepts, noting what was different in their discussions.

5. Draw guidelines: given the observations, we will be designing a framework for the conduction of extreme design meetings and a guide to avoiding the pitfalls identified.

It should be noted that the analyses are ongoing, and that we have thus far analyzed approximately one third of the concepts on the ontology (the ones that presented problems). Meeting data is shown in Figure 2

EVD	Nome	Tempo	Funç	Atrib	Tipo	Operação	Transcrição	Material	Observações
1	Roberto	00:11:26	Visual	Verbal	Fala em defeito menor				
2	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
3	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
4	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
5	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
6	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
7	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
8	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
9	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
10	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
11	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
12	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
13	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
14	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
15	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
16	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
17	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
18	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
19	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
20	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
21	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
22	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
23	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
24	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
25	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
26	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
27	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
28	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
29	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
30	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
31	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
32	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
33	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
34	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
35	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
36	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
37	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
38	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
39	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
40	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
41	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
42	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
43	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
44	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
45	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
46	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
47	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
48	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
49	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				
50	Roberto	00:11:41	Visual	Verbal	Fala em defeito menor				

Figure 2: Meeting data

4. Meeting Observations

The initial step in the project was to frame the ontology design activity in terms of the activity theory framework. The activity structure of ontology design can be seen in Figure 3. The *activity* in this analysis is ontology design, and the ontology is the desired *outcome*. The design process transforms raw knowledge into a structured representation. The *subjects* are the experts involved in the meetings and discussions, including the knowledge engineer. The *object* is the experts' knowledge, including documents and tacit knowledge they have at their disposal. These they work on and transform into the ontology. The *tools* are flipcharts, shared displays, computers and pens and paper.




Figure 3: Structure of the ontology design activity

Activities are driven by goals. The main goals identified involve iteratively expanding and reducing the search space. To that end, the group cycles through actions involving the generation, organization and evaluation of alternatives and making decisions. Table 1 shows the goals and associated activities.

Goals	Actions
Increase coverage	Generate new solutions and alternatives
Adjust model	Organize search space, evaluate alternatives, decide about elements
Consider elements	Evaluate alternatives
Validate representation	Decide about elements

Table 1: Ontology design activities and corresponding goals

To measure influences on acceptance levels, levels of acceptance were qualified as high (complete acceptance, no complaints), medium (partial acceptance, some complaints and changes requested) or low (non-acceptance) for the final design and at each stage of the design. Participation was also qualified in three levels, and measured for each discussion: participative (participant fully engaged in discussions), passive-attentive (participant participated occasionally, but paid attention to the discussion the rest of the time), passive-inattentive (participant did not participate, seemed uninterested.)

5. Discussion

The goal of our research is to map conditions that should be avoided in meetings, and to provide guidelines for the meeting coordinator. Our measurement is the level of acceptance of the final design.

Our initial observations lead us to believe that attentive behavior, group splits with rejoin on the same topic, the level of discussion and distribution of participation have an impact on final design acceptance, individual satisfaction and rework. At this time, we are exploring factors that lead to increased acceptance of the final deliverable.

5.1. Attention and Group Partitioning

We have identified a number of operations that contribute to the outcome, but the negative one that appears most frequently is lack of attention. However, we also noticed that lack of attention can be of two natures: productive or unproductive. While attention shifts are usually bad both for the meeting and the outcomes, productive attention shifts shouldn't necessarily be avoided. Sometimes, a user would shift his/her attention to a different task (such as checking documentation) and later return to the topic with new information or argumentation to contribute. An example is shown in Figure 4.



Figure 4: Example of a productive attention shift.

On the other hand, unproductive attention shifts are usually caused by external distractions (phone ringing, incoming mail or pressing matters external to the meeting) or by boredom (losing interest). Both problems could be overcome by creating more engaging meeting dynamics in order to capture participants' attention. Sometimes attention shifts could lead to partitioning of the group, and parallel discussions would ensue.

Group partitioning may happen when attention shifts happen. Partitioning due to unproductive attention shifts (e.g., side conversations about unrelated topics) should be avoided. The normal tendency would be to attempt to rejoin the group immediately, but productive group divisions are not necessarily bad for the group: in some situations, while part of the group focused on the main topic, a sub-group discussed related issues, exploring part of the design space separately and coming back to the group with new information that could help the group in its task. An example of group partitioning is shown in Figure 5.




Figure 5: An example of group partitioning

The tricky issue in both situations is determining when the subgroup or individual has strayed too far off topic and should be steered back. While a certain level of digression is helpful, going too far from the original subject may not be.

5.2. Mental Models

What we perceived was that the design process was in fact a process of mental model alignment. Individuals

came in with their own mental models and perceptions of what the design should be. During the process, they first presented their beliefs (fashioned according to their mental models), then discussed these beliefs, in order to find a common ground to construct a joint model that represents the perceptions of all participants.

At each challenge an individual mental model could be altered, though reflection and the introduction of new information. Throughout the process, would align their models with the group model. The process inevitably involved initial resistance to change, but after extensive discussion ended with some sort of agreement.

Our observations are that the farther the individual mental model is from the group model at the end, the more difficult it will be to gain acceptance at the deployment stage. While we do not have an infallible way to change an individual's mental model, it is possible to elicit individual mental models by requesting private or anonymous opinions. If the models are too far apart, more discussion is necessary to reach an acceptable compromise position.

5.3. Change of Heart

One interesting case was that of participant R. This participant had medium levels of participation throughout (attentive at least, participative in many occasions), and agreed with decisions during the meetings. However, at the final presentation, he/she had several complaints about the final design, causing not only acceptance problems, but also discomfort with other participants, who couldn't understand why the sudden change of heart.

In our study, it becomes clear that participant R did not change his/her mental model throughout the process. His/her behavior was indicative of "giving up": after a short time arguing, he/she lost interest and agreed with the group in order to dismiss the discussion. At the end of the process, he/she stated his real opinion and disagreed with the end result.

This type of situation could be avoided by encouraging further discussion, especially from the particular participant. Formally requesting an opinion or feedback, especially in between meetings or anonymously might yield the true individual opinions that didn't surface during the meeting.

It should be noted that this participant had recently undergone training and had a theoretic background that others lacked. Additionally, he/she was more the type of person who places value on detailed descriptions, and wanted to include much more detail in the final ontology than the rest of the group. However, the design of an ontology involves the determination of the appropriate level of abstraction and detail for the model. Thus, the final construct might not have been as detailed as participant R would have liked. This type of conflict should be elicited and participants should understand

what the appropriate level of detail is at the beginning of the activity.

5.4. Requirements and Techniques

Given our initial observations, we draw the following requirements for a support system:

- Increase participation;
- Monitor discussion, to detect productive or unproductive attention shifts;
- Check the level of exploration of the design space;
- Check if the discussion was evenly distributed between participants;
- Check participation levels per participant.

Based on our observations and perceived requirements, we are designing a method to evaluate decision quality, using random walk theory [4]. In this technique, we view the discussion as a form of random walk, where every statement is more or less related to the previous one. Using this parallel, we can apply complex systems analysis methods to analyze if the discussion presents persistent or anti-persistent behavior, which tells us whether it is proceeding in the right direction, how far it has moved from the original topic and its tendencies for divergence or convergence. At the initial stage, we will apply these methods to the tags assigned to the discussion entries. However, the method is being designed for application directly to the discussion text.

5.5. Design Acceptance

Cross-referencing participation levels and acceptance models, we can draw a few observations about the expected behavior of participants regarding corporate acceptance.

Participative members who displayed high acceptance of the artifact will function as corporate champions of the technology, disseminating the new ontology into organizational culture. With these participants, there is a high probability that the design will be successfully adopted in the organization. These people have their mental models aligned with the group model (represented by the ontology). We have already identified two such cases in our study.

Attentive participants with high acceptance also lead to good chance of organizational deployment. These individuals have accepted the group model, even though there is no evidence of a change in individual mental model.

Inattentive participants who present high acceptance will not be champions or seek to propagate the

technology, but they should also not work against the adoption of the ontology.

Participative members with low acceptance levels may become problematic agents at the end of the project. Any failure during the test period will lead to the individual abandoning the project and the individual may try to work against organizational adoption, but arguments are weakened by the fact that he/she participated in the discussions and made compromises at that point. This individual did not align his/her mental model during discussions, and this shows at this stage. We have identified one such individual in this project.

Passive participants with low acceptance will also avoid usage, and may create problems if usage is imposed. Inattentive participants who do not accept the technology will avoid adoption, but may use it through imposition. These individuals may create problems during the adoption phase, resisting usage and working against the project (we have identified two such cases in this project).

Even though not yet complete, our studies already yield some interesting observations, which we believe these studies will help form a basis for future work on design meetings. Some new techniques are already being designed based on these findings. Despite the existence of other studies that describe factors that lead to production losses in meetings and how to identify them (eg., [3]), these did not pertain to the design meeting context and there was no cross-reference with acceptance of the final design, which is part of our study.

Acknowledgements

This project is financially supported by Petrobras. We also thank the engineers for their participation in the design meetings.

References

- [1] Kuutti, K. Activity Theory as a potential framework for human computer interaction research. In B. Nardi (ed.): Context and Consciousness: Activity Theory and Human Computer Interaction, Cambridge: MIT Press, 1995, pp. 17-44.
- [2] Mark, G. Extreme Collaboration. In Communications of the ACM, 45, 6, ACM Press, (2002), 89-93.
- [3] Westley, F., Waters, J.A., 1988. Group Facilitation skills for managers. Management Education and Development, vol. 19, pp. 134-143.
- [4] Pemantle, R. A survey of random processes with reinforcement. Probability Surveys, Vol. 4 (2007) 1-79