Visualizations to Support Interactive Goal Model Analysis

Jennifer Horkoff
Department of Computer Science,
University of Toronto, Canada
jenhork@cs.utoronto.ca

Eric Yu
Faculty of Information,
University of Toronto, Canada
yu@ischool.utoronto.ca

Abstract— **While goals models can visually present alternate ways for achieving goals and how actors depend on each other, it can be challenging to follow the reasoning through complex paths in the model. In earlier studies which tested the utility of procedures for guiding analysts to perform interactive forward and backward reasoning through i* models, we uncovered several difficulties faced by study participants. Users often have trouble choosing suitable starting points for analyzing the model, and in understanding conflicts among alternatives. We have recently developed visualization mechanisms to alleviate these difficulties. Specifically, roots and leaves in the model are automatically detected and highlighted as suggested starting points for analysis. Goals within a conflicting path are highlighted during analysis. The visualization mechanisms were tested with users in five follow-up studies. The results suggest several further visualization mechanisms which could support analysis.**

Keywords-Goal Modeling, Model Analysis, Model Visualizations

I. INTRODUCTION

Goal modeling has been proposed as a tool for "early" stages of system analysis (Early RE) [1], where domain information is often incomplete, difficult to quantify, and where contact with the stakeholders is ideally frequent. Once a goal model reaches a certain threshold of completeness, it is possible to apply systematic analysis, analyzing the effects of alternative requirements. However, applying and understanding the results of analysis over complex visual models is difficult. Analysis over goal models can be better supported by applying various visualization mechanisms.

Most goal model analysis procedures proposed to facilitate "what if?"-type analysis over goal models have focused on fully-automated analysis (for e.g. [2] [3]). Previous work by the authors has argued that the incomplete and "fuzzy" nature of Early RE calls for methods which supplement the model with stakeholder knowledge, encouraging stakeholder participation in the analysis process and iteration over captured knowledge [4] [5] [6]. Such procedures are "interactive" in that they continually prompt for "human judgment" in order to potentially resolve conflicts or promote partial evidence. The interactive nature of the procedure means that visualizations are particularly important in helping users understand the model and analysis. In this work we introduce and study visual interventions aimed to aid users in interactive goal model analysis.

Prior to this study, ten individual and one group case study were conducted in order to test the effects of a guided interactive model analysis procedure, compared to ad-hoc (non-systematic) analysis [7]. Qualitative analysis over the results revealed several challenges experienced by the users when attempting to interactively analyze a goal model. In the current work, we focus on two of these challenges:

a) Starting points for analysis: users had difficulty knowing where or how to start analysis in the models. Depending on the direction of interactive analysis (forward from alternatives to target goals, or backward from target goals to alternatives) model leaves or roots often act as suggested starting points for analysis. However, results from the prior studies revealed that users often have difficulty identifying leaves and roots in agent-oriented goal models. This can be attributed to certain visual characteristics of goal models such as unclear link directions, the presence of actor boundaires, and non-standard layout when compared to typical tree graphs. In this work, we provide a feature which automatically highlights model leaves and/or roots.

b) Understanding conflicts: when the interactive backward procedure revealed conflicts in the models (inability to find a permissible assignment of analysis labels), users had difficulty understanding the reasons behind the conflicts. In the backward analysis direction, the interactive procedure described in [6] finds answers to analysis questions by translating the model into a satisfiability (SAT) problem. When model conflicts are discovered, it finds information about these conflicts by finding the UNSATisfiable core – a list of SAT clauses which are involved in the conflict. Presenting this information to the user in a form which is understandable to users presents a challenge. In this work, we use the information in the UNSAT core to find model elements in the conflict path and elements which are the "sources" of the conflict, highlighting both types of elements in the model.

We studied the effectiveness of these visualizations by performing five follow-up studies using participants of the previous studies in [7]. Users were asked to repeat previous analysis over models, with the analysis tool (OpenOME [8]) supporting the new interventions. Results of the studies are analyzed qualitatively to gauge the utility of the interventions, suggest improvements to the new visualizations and find future directions for visual support of interactive goal model analysis.

This paper is organized as follows: Section II describes necessary background, including goal models, interactive goal model analysis, and the motivating studies; Section III better motivates the challenges currently experienced by analysis users; Section IV describes the visualization techniques; Section V describes follow-up studies applying the new visualizations; Section VI provides and discusses results; Section VII outlines related work; while Section VIII provides conclusions and future work.

II. BACKGROUND

A. Goal Models

Although several goal model frameworks and approaches have been introduced, we use the i* Framework as an example framework after [5] [6]. The i* Framework provides a graphical depiction of system actors, intentions, dependencies, responsibilities, and alternatives. The social aspect of i* is represented by actors, including agents and roles, and the associations between them. Actors depend upon each other for the accomplishment of tasks, the provision of resources, the satisfaction of goals and softgoals. Softgoals are goals without clear-cut criteria for satisfaction. Actors have boundaries containing the intentions of an actor: desired goals and softgoals, tasks to be performed, and resources available.

The relationships between intentions inside an actor are depicted with Decomposition links, showing the elements which are necessary in order to accomplish a task; Means-Ends links, showing the alternative tasks which can accomplish a goal; and Contribution links, showing the effects of softgoals, goals, and tasks on softgoals. Positive/negative contributions representing evidence which is sufficient to satisfy/deny a softgoal are represented by Make/Break links, respectively. Contributions with positive/negative evidence that is not sufficient to satisfy/deny a softgoal are represented by Help/Hurt links. More information on the i* syntax can be found in [9]. An example i* model with a legend is shown in Fig. 1. A more complex example used in one of the case studies which shows some of the visual complexities associated with i* can be found in Fig. 2.

In order to more precisely define the visualizations introduced in this work, we follow the formalization of the i* framework presented in [6]. The definitions use the → notation to represent relationships between elements, so if $(i_1,i_2) \in R$ we write this as $R{:}i_1 \rightarrow i_2$.

Definition: i*model. *An i* model is a tuple <I, R, A>, where I is a set of intentions, R is a set of relations between intentions, and A is a set of actors. Each intention maps to one type in {Softgoal, Goal, Task, Resource}. Each relation maps to one type in {R^{me}, R^{dec}, R^{dep}, R^{c}}, means-ends, decomposition, dependency, and contribution links, respectively. R^{c} can be broken down into a further set {R^{m}, R^{hlp}, R^{u}, R^{hrt}, R^{b}} corresponding to make, help, unknown, hurt, and break, respectively. R^{dep} and R^{c} are binary relations (one intention relates to one intention, $R: I \rightarrow I$), while the remaining relationships, R^{me} and R^{dec}, are (n+1)-ary (one to many intentions relate to one intention), $R: I \times \ldots \times I \rightarrow I$.*

Analysis labels in are used in i* to represent the degree of satisfaction or denial of an intention. Following [10], the (Partially) Satisfied label represents the presence of evidence which is (insufficient) sufficient to satisfy an intention. Partially denied and denied have the same definition with respect to negative evidence. Conflict indicates the presence of positive and negative evidence of roughly the same strength. Unknown represents the presence of evidence with an unknown effect. We use the formal definition of analysis predicates from [6]:

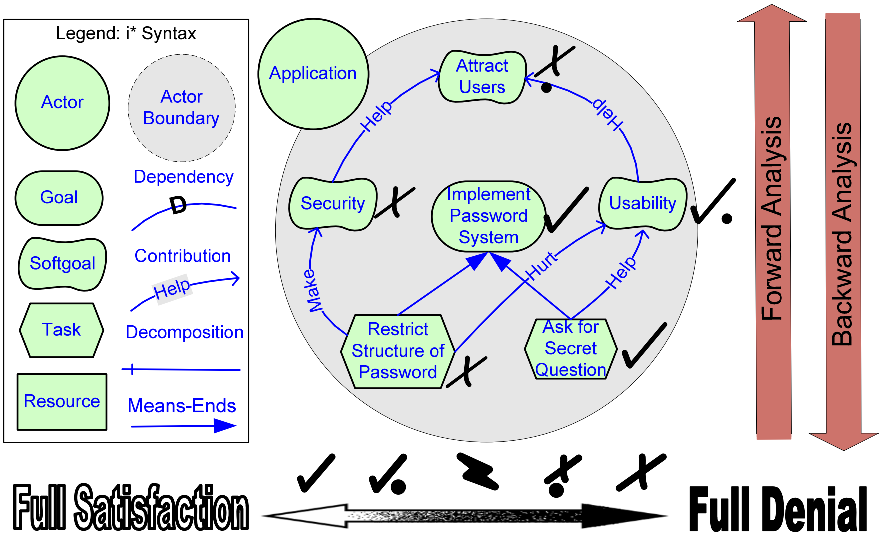

Figure 1. Example i* Model for a Simple Application

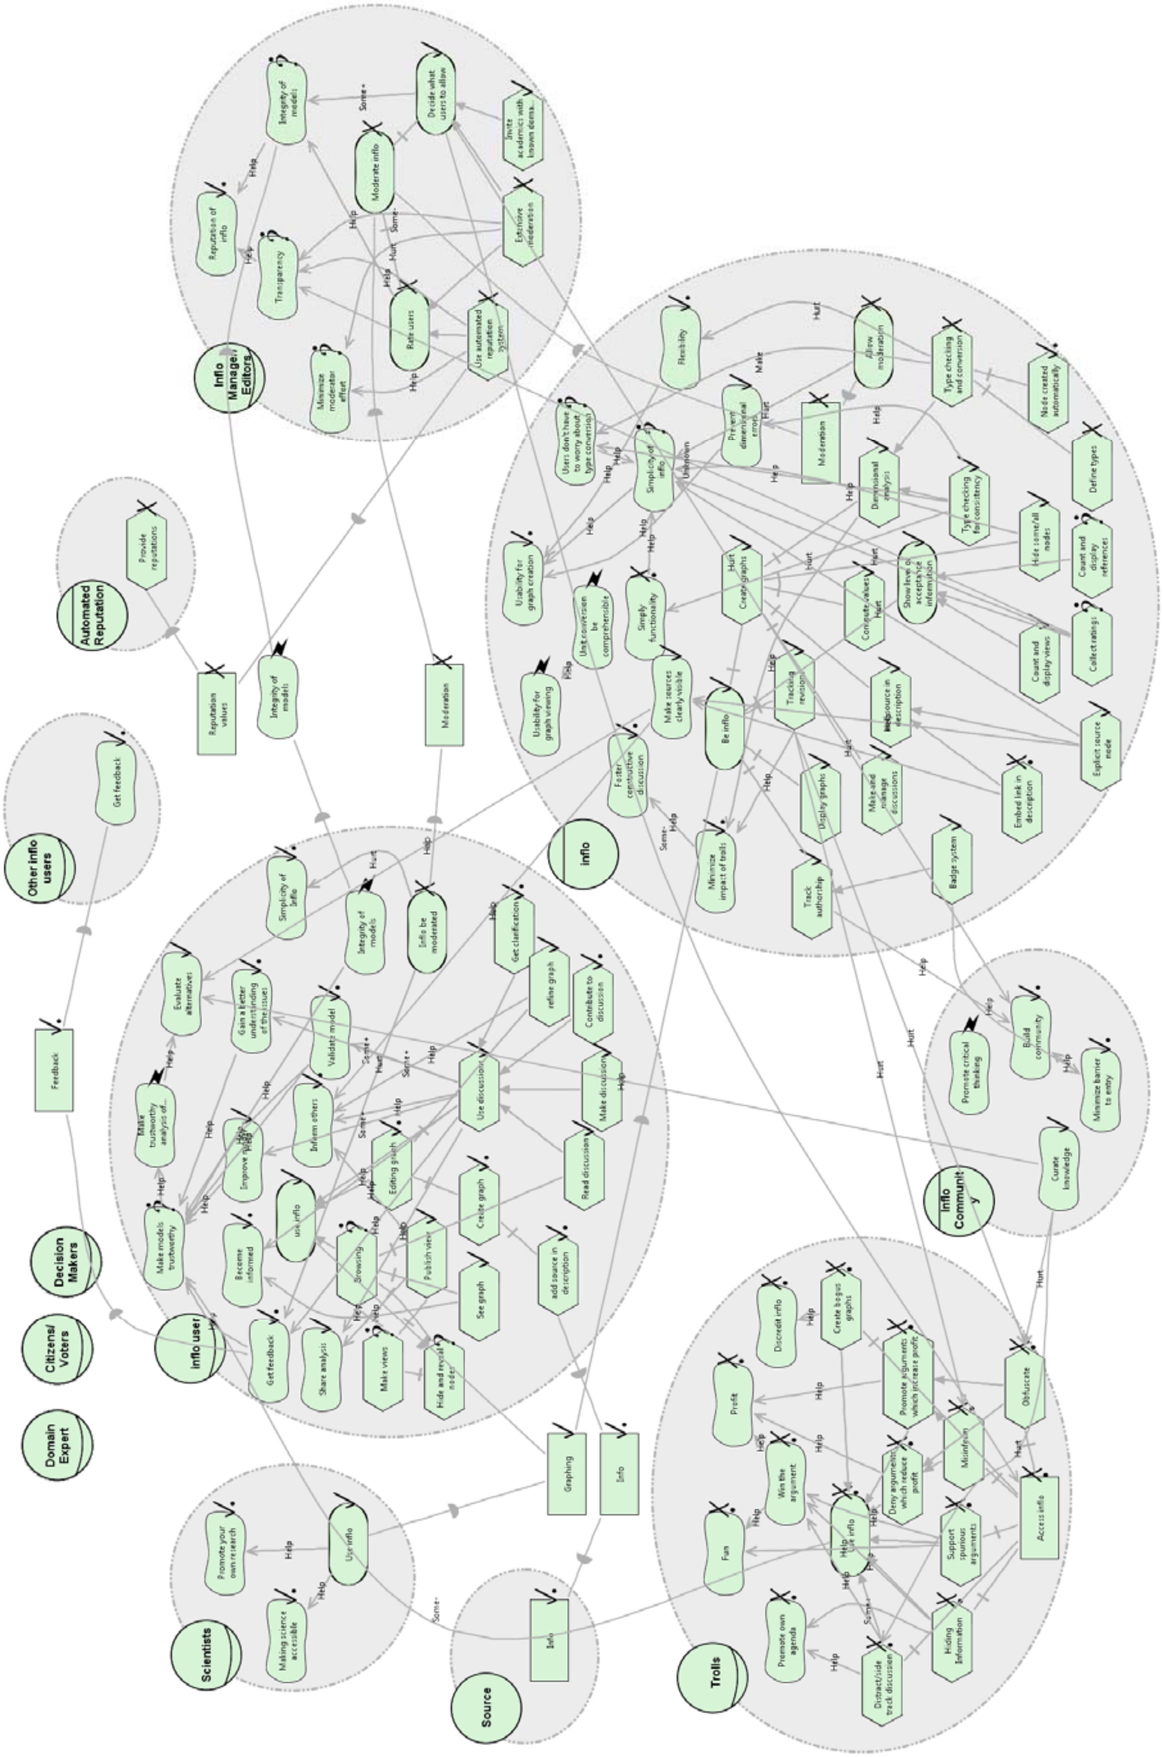

Figure 2. High-level view of the i* Model Resulting from the inflo Case Study

Definition: analysis predicates. *We express agent-goal model analysis labels using a set of predicates, V, over i ∈ I. Each v(i) ∈ V maps to one of {S(i), PS(i), C(i), U(i), PD(i), D(i)} where S(i)/PS(i) represents full/partial satisfaction, C(i) represents conflict, U(i) represents unknown, and D(i)/PD(i) represents full/partial denial.*

B. Interactive Goal Model Analysis

The framework for interactive goal model analysis currently provides two types of analysis procedures: forward (from alternative solutions to goals) and backward (from goals to solutions).

The forward analysis procedure starts with an analysis question of the form "How effective is an alternative with respect to goals in the model?" The analysis starts by assigning qualitative evaluation labels to intentions related to the analysis question. These values are propagated along links in the forward direction (i* links are directed) using defined rules. The nature of a Dependency indicates that if the intention depended upon (dependee) is satisfied then the intention depended for (dependum) and intention depending on (depender) will be satisfied. Decomposition links depict the intentions necessary to accomplish a task, indicating the use of an AND relationship, selecting the "minimum" value amongst intentions in the relation, ordered from satisfied to denied. Similarly, Means-Ends links depict the alternative tasks which are able to satisfy a goal, indicating an OR relationship, taking the maximum value.

The procedure adopts Contribution link propagation rules from the NFR procedure [10]. Positive values (✓, ✓•) propagated through positive links (Make, Help) produce positive values, weakened with the latter link. Positive values propagated through negative links (Break, Hurt) propagate full or partial negative values (✗,✗•). Links in i* are symmetric: negative values propagated through positive links produce negative values, and negative values propagated through negative links produce positive values.

The procedure prompts for interactive input when human judgment is needed to combine multiple incoming conflicting or partial values to determine the resulting label for a softgoal. Judgments rely on participant domain knowledge to compensate for the inherent incompleteness of high-level, social models.

Once the procedure has finished interactive propagation, the final analysis values for the intentions of each actor are examined in light of the original question. By looking at the degree of satisfaction or denial of key intentions, an assessment is made as to whether the alternative would work in the domain.

An example forward evaluation result can be seen in Fig. 1. Here we have asked the question "What if we Ask for Secret Question but do not Restrict Structure of Password?" Analysis labels are propagated and human judgment is needed to resolve incoming evidence for Usability and then Attract Users. The final results show that Attract Users is partially denied and the alternative is not viable. More information concerning the procedure can be found in [4] and [5].

Backward analysis allows users to ask questions of the form "Is it possible for a set of intentions to be satisfied? If so, how?". The procedure uses the same propagation rules as in the forward procedure, but now propagates evidence both forward and backward. The backward propagation is implemented via a formalization of i* expressed in conjunctive normal form (CNF) and passed to a SAT solver. Human judgment is needed for intentions which have conflicting analysis values assigned. In the backward procedure judgment takes the form "I want Intention X to have the value V. Give me a combination of values for the contributing intentions that would result in the target value". Users are shown a list of contributing intentions and their associated links, and then are expected to choose a value for each contributing intention. The user can say "No Combination" if no combination of values would produce the target value. The procedure is iterative in that it repeatedly calls the SAT solver until a satisfying assignment is found and no more human judgment is needed. Each iteration involves more human judgment questions. When judgments produce conflicting results, the procedure "backtracks", re-asking the last round of questions involved in the conflict. Presenting the results of this conflict in an understandable manner is one of the central aims of the current work. A solution may not be found if the model has a conflict and does not require human judgment, or if the user cannot enter human judgments which produce a solution.

We can consider an example backward analysis question over the model in Fig. 1. Here, we ask the question "Is it possible for Attract Users to be at least Partially Satisfied while Implement Password System is Satisfied?" These target lables (✓•, ✓) are placed on the model. The procedure first asks human judgment for Attract Users, asking for a combination of values for Usability and Security which would produce a ✓• value for Attract Users. If the user enters ✓• for both, the procedure then asks for a judgment for Usability. The user provides a combination of labels for Restrict Structure of Password and Ask for Secret Question which would produce ✓• for Usability, say ✗ for the former and ✓ for the latter. The procedure would then find a conflict as Restrict Structure of Password must be ✗ from the judgment for Usability and ✓• through propagation through Security from the judgment for Attract Users (we enforce that leaf intentions must not be both (partially) satisfied and (partially) denied. The procedure backtracks to the judgment for Usability. If the user has no more viable combinations they enter "No Combination" and the procedure backtracks to the initial judgment for Attract Users. If the user now enters ✓ for Security and ↯ (conflict) for Usability, then ✓ for both tasks in the next judgment aiming for ↯ for Usability, then the procedure finds a solution which does not require further human judgment. The final result is that in order for Attract Users to be ✓•, both solutions must be implemented. More information can be found in [6].

C. *Qualitative Studies*

Prior to this investigation of visualization mechanisms, eleven studies were conducted to test the effects of interactive goal model analysis. Ten studies involved individuals analyzing models. Six of the participants analyzed complex models made by others in the domain of ICSE Greening [11]. The other four analyzed simpler models about student life they created themselves. Half of the participants in each group used the procedures described in Section II, implemented in the OpenOME tool, while the other half analyzed the model in an ad-hoc way (either propagating manually or just by examining the model). Participants were graduate or undergraduate students in Computer Science or Informatics who had some previous experience with i*. The individual studies were performed with a facilitator present, observing results and guiding study steps. The eleventh study involved a group of grad students and a professor who were implementing a modeling tool (inflo) to support modeling and discussion of "back-of-the envelope" calculations. In this study the facilitator acted as the modeler, modeling input provided by the participants. The individual studies lasted two hours per participant, while the group study occurred over three two-hour sessions.

The studies aimed to collect evidence to support or deny hypotheses concerning model iteration and elicitation prompted by systematic analysis. Further analysis uncovered additional theories concerning the effects of interactive model analysis. Results showed that although both systematic and ad-hoc analysis can promote model iteration and further elicitation, systematic analysis increased the coverage and consistency of the analysis. Further information on study design and results can be found in [7].

III. Motivation: Making Analysis of i* Models Easier to Perform and Comprehend

In addition to the findings described in the previous section, results of the case studies showed that users have several issues applying analysis and understanding analysis results. Although these issues may be addressed by increased training in i* and interactive analysis, enhanced visualization techniques can also help. We focus on two issues which appeared prominently in those prior studies: starting analysis and understanding model conflicts.

Participants often had difficulty knowing how or where to start analysis. In cases where they were given analysis questions to answer using the model, they had problems finding where to place the labels to reflect the analysis questions. In cases where they were encouraged to come up with their own analysis questions, they had difficulty knowing how to start analysis. In order to help participants, a suggested analysis methodology was developed (see [5] for the full methodology). Based on experience from previous case studies [12] [13], the methodology suggested to start forward analysis by identifying leaf intentions in the model and to start backward analysis by identifying root intentions. These intentions would serve as suggested starting points, although users were encouraged to add other starting points, or to make changes to the model if current model roots or leaves did not seem to be sensible analysis starting points.

Four participants applied this methodology in the individual studies. All participants had at least some problems correctly identifying leaf or root intentions in the models. Although this seems like a simple task for users with training in the use of graphs as computational data structures, certain characteristics of i* models make this task challenging, namely:

- some i* links do not have an obvious direction (for e.g. dependency links),
- the presence of actors and actor boundaries can often cause users to ignore links coming into or out of an actor, and
- as i* models often contain both actors and cycles, they are typically not laid out in similar manner as a typical graph (roots on top, leaves on bottom).

Fig. 2 shows a high-level view of an example i* model used in six of the individual case studies. Note the potential difficulty in identifying model leaves and roots.

Participants also had difficulty understanding conflict results in backward analysis. In this procedure, when a satisfying assignment cannot be found for the CNF formula used to represent the model and current judgments, information about the conflict derived from the UNSAT core is presented to the user. Such information would potentially guide their decisions while backtracking over human judgment thus far. During the studies, the user was presented with a list of intentions involved in the conflict. Results showed that this was insufficient to understand where in the model the conflict occurred, or why.

For example, a conflict occurred during the analysis of Fig. 2. This figure shows the state of the model at the time of the conflict. The user is provided with text ouput listing the intentions involved in the conflict, in this case: Simple functionality, Type checking for consistency, Flexibility, Type checking and conversion, use inflo, Graphing, Be inflo, Create graphs, Node created automatically, Define types, Dimensional analysis, Usability for graph creation, and Simplicity of inflo. In this particular case two leaf intentions, Dimensional analysis and Type checking and conversion are both partially satisfied and not partially satisfied (*PS* and *not PS*), causing a conflict. The reader can see that even with the ability to zoom in and out to explore the model, locating the intentions and understanding the causes of the model conflict is difficult.

An alternative form of output would be a list of the clauses involved the UNSAT core, as in Fig. 3. However, we do not wish to assume that i* analysis users have a background in logic. In the next section we explore visualizations to improve the comprehensibility of conflicts.

IV. Visualizations to Support i* Analysis

A. *Leaf and Root Intention Highlighting*

In order to address difficulties in starting analysis, we have added the ability to automatically highlight leaf and root intentions in the model. We provide a formal definition of leaf and root in the following:

The following intention clauses are conflicting:
not PS(Simply functionality) OR PD(Type checking for consistency)
not PS(Flexibility) OR PD(Type checking and conversion)
not PS(use inflo) OR PS(Graphing)
not PS(Graphing) OR PS(Be inflo)
not PS(Be inflo) OR PS(Create graphs)
not PS(Type checking and conversion) OR PS(Node created automatically)
not PS(Type checking and conversion) OR PS(Define types)
not PD(Type checking and conversion) OR PD(Node created automatically) OR PD(Define types)
not PS(Create graphs) OR PS(Dimensional analysis)
not PS(Dimensional analysis) OR PS(Type checking for consistency) OR PS(Type checking and conversion)
S(use inflo)
PS(Usability for graph creation)
not S(use inflo) OR PS(use inflo)
not PD(Define types) OR not PS(Define types)
not PD(Type checking for consistency) OR not PS(Type checking for consistency)
not PD(Node created automatically) OR not PS(Node created automatically)
not PS(Usability for graph creation) OR PS(Simplicity of inflo)
not PS(Usability for graph creation) OR PS(Flexibility)
not PS(Simplicity of inflo) OR PS(Simply functionality)

Figure. 3 Example Clauses in the UNSAT core for a Conflict During Backward Analysis

Definition: leaf or root intention. *An intention i ∈ I is a leaf if there does not exist any relation, r ∈ R such that r: I → i or r: I x… x I → i , it is a root if there does not exist any relation, r ∈ R such that r: i → I or r: i x … x I → I.*

In our OpenOME implementation the user can select the "Mark Model Leaves" or "Mark Model Roots" options. In each option root and leaf intentions appear blue or bright green, respectively. Fig. 4 shows a model used in one of the individual studies with leaves and roots highlighted.

B. Conflict Highlighting

Using the information provided by the UNSAT core during a conflict, we highlight intentions in the model involved in the model conflict. The visualization differentiates between intentions on the path involved in the conflict and intentions which were the "logical source" for the highlight, i.e. for which an analysis predicate is both true and not true (for e.g. PS(i) and not PS(i), i ∈I). We formally define these concepts in the following:

Definition: model conflict. *When the SAT solver used in the backward analysis procedure cannot find a solution over a CNF representing a model, there is a conflict. Specifically, in all possible assignment of variables, for one or more intentions, i ∈I, both v(i) and not v(i) hold, v(i) ∈V.*

Definition: UNSAT core from backward analysis. *An UNSAT core is an unsatisfiable subset of clauses, c ∈ C, in an unsatisfiable CNF representing a model. Each clause c ∈ C is a disjunction of one or more analysis predicates over intentions, v(i), for v ∈V and i ∈ I.*

For example, Fig. 3 provides an UNSAT core for an example conflict in Fig. 2. More information concerning the construction of the CNF formalism can be found in [6].

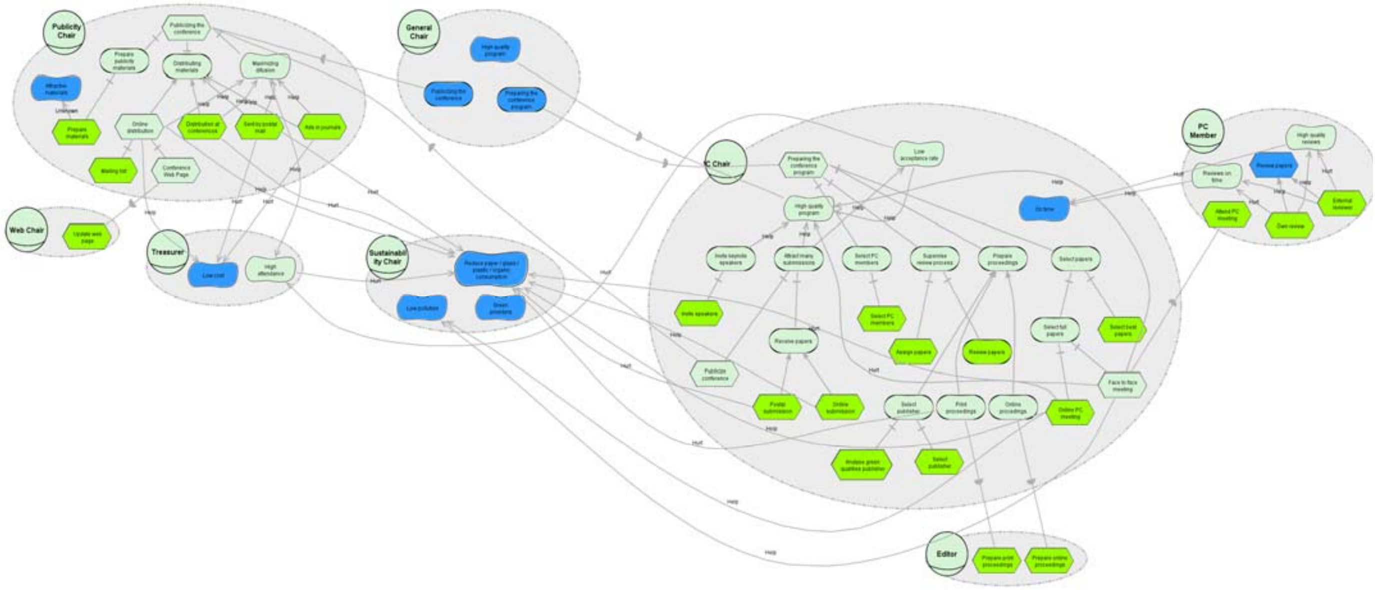

Figure 4. Model from [11] Used in an Individual Study with Roots and Leaves Highlighted

Definition: intention involved in conflict*. An intention, i ∈ I, is involved in the conflict when it appears in one of the clauses, c ∈ C, in the UNSAT core.*

For example, in the UNSAT core in Fig. 3 Simplicity of inflo and Graphing are intentions involved in the conflict.

Definition: logical source of conflict*. An intention, i ∈ I, is a logical source of a conflict when:*

a) It is involved in the conflict, and

b) An analysis predict is both true and not true for an intention, according to the UNSAT core, i.e. for v(i), v ∈ V, both v(i) and not v(i) hold.

In Fig. 3, the sources of the conflict are Dimensional analysis and Type checking and conversion, as mentioned.

Definition: root clauses in UNSAT core. *Clauses c ∈ C with only one predicate , v(i), v ∈ V, i ∈ I. These clauses are a subset of the initial target labels for backward analysis.*

For example, in Fig. 3, Usability for graph creation and Use inflo are root clauses. Our OpenOME implementation parses the UNSAT core using a recursive procedure starting at the root clauses of the core, traversing towards the logical sources of the conflict. Intentions involved in the conflict are collected along the way.

When a conflict occurs during backward analysis, our implementation highlights all intentions involved in the conflict as orange and intentions which are the logical sources of a conflict in red. An example can be seen in Fig. 5, where the conflict previously described in Fig. 2 is now shown through model highlighting. The highlighting is removed as the procedure backtracks to acquire more human judgments.

In addition to conflict highlighting, users are presented with a list of the intentions involved in the highlight along with the analysis value that these intentions would be assigned in order to produce the conflict. These values are extracted from the UNSAT core by storing the analysis value assigned to each predicate, with the logical sources having several conflicting values. Example output of this type corresponding to Fig. 5 is shown in Fig. 6.

V. Visualizations Applied

In order to test the practical utility of the new visualizations, we performed five follow-up studies using participants from the initial eleven studies described in [7]. Three participants repeated analysis over small models they

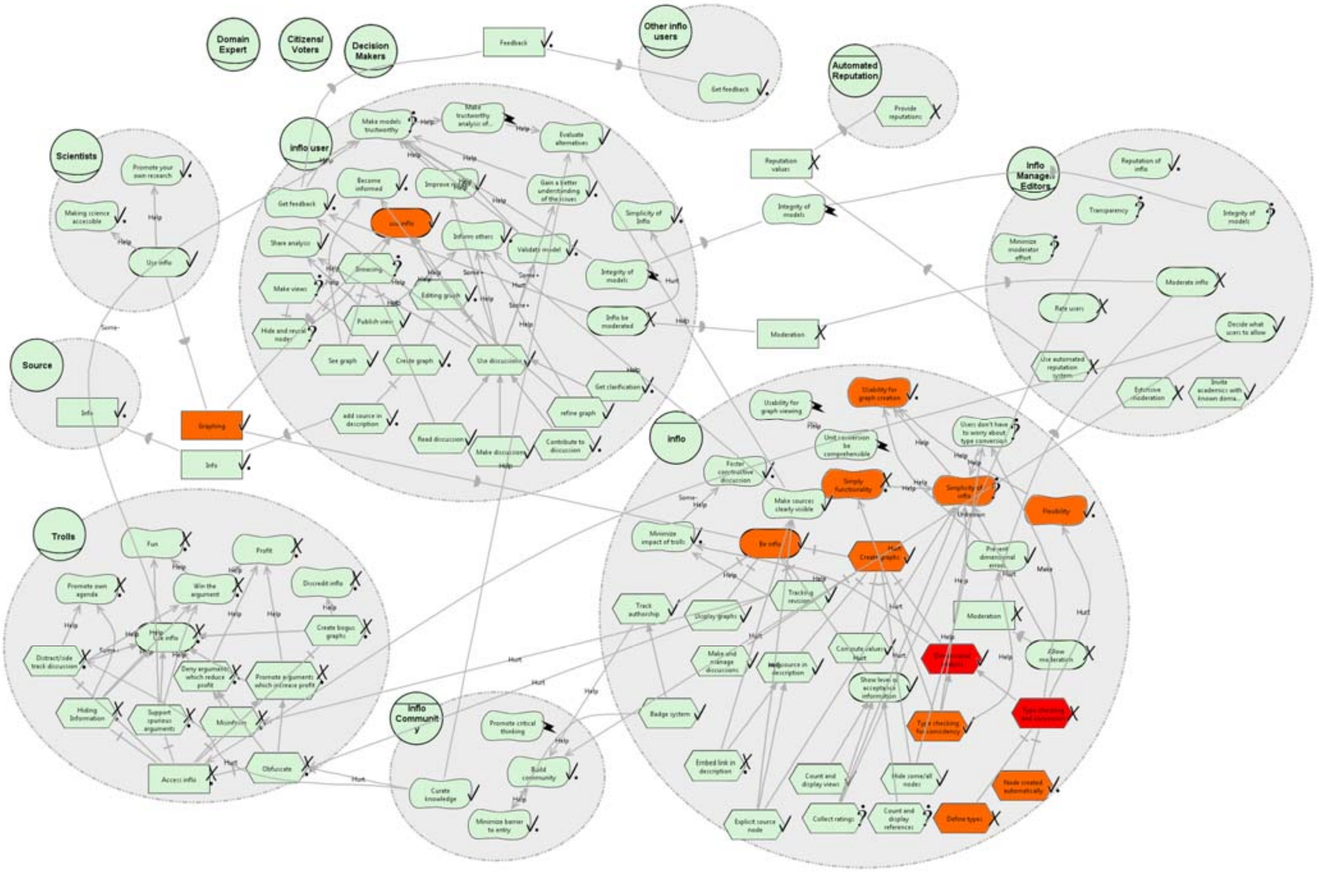

Figure 5. High-level view of Conflict Highlighting in the i* Model Resulting from the inflo Case Study

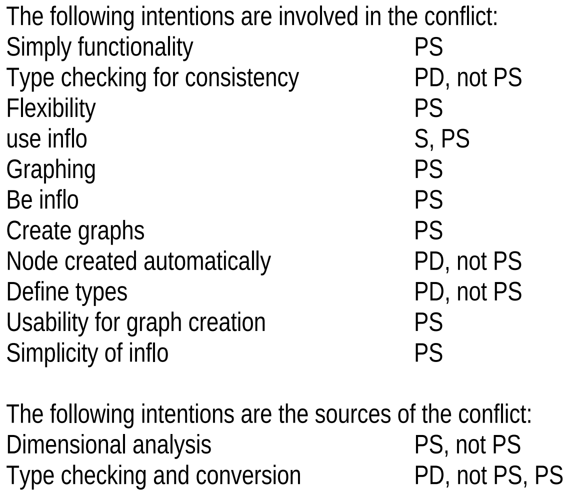

The following intentions are involved in the conflict:

Simply functionality	PS
Type checking for consistency	PD, not PS
Flexibility	PS
use inflo	S, PS
Graphing	PS
Be inflo	PS
Create graphs	PS
Node created automatically	PD, not PS
Define types	PD, not PS
Usability for graph creation	PS
Simplicity of inflo	PS

The following intentions are the sources of the conflict:

Dimensional analysis	PS, not PS
Type checking and conversion	PD, not PS, PS

Figure 6. Clauses in the UNSAT core for the Conflict During Backward Analysis shown in Fig. 5

had created, one participant analyzed a large model created by others (Fig. 4) and the last participant analyzed the model created in the inflo case study, Fig. 2 and 6, (the participant had helped create this model). Each session lasted 30 minutes to an hour. Participants were specifically asked to comment on the new interventions: Do the leaves/roots highlighted in the model make sense? Can you understand why there is a conflict? Participants were paid $20 for their time, and results were not shared with anyone who could affect course or academic standings. Session audio and screen movement were recorded. A facilitator was present in each session directing study steps and making observation notes. Data collected was analyzed qualitatively, as in the previous studies, classifying observations into related categories or theories. Threats to the validity of the study designs are discussed in the next section.

VI. Results and Discussion

In the following section we summarize results from the validation studies, including suggested improvements to the new visualizations. Threats to the validity of the studies are summarized and alternative study designs are considered.

A. *Leaf and Root Intention Highlighting*

Generally, reaction to root and leaf highlighting was positive, with participants understanding the results of the automatic highlighting. A few of the highlighted roots and leaves were surprising to participants, but upon examining the links in the model more closely, they were able to determine why specific intentions were leaves or roots. The participants' surprise at the identification of some leaf/root intentions can be attributed to the difference between global and local roots and leaves. The definition in Section IV.A describes global model leaves and roots, but it is also possible to identify intentions which are only leaves and roots inside of an actor, ignoring incoming or outgoing links. Often participants thought that these intentions should also be considered leaves and roots.

Because we encouraged participants to use the automatic highlighting only as a suggested starting point for analysis, in a few cases participants decided to add initial analysis values to local leaves and roots in addition to global leaves/roots. Future improvements could highlight both global and local leaves and roots automatically, perhaps using different visual cues to distinguish between global and local.

One participant stated that leaf/root highlighting was not useful, because of the small size of the model and because they had already successfully identified roots in the last round of the study. We can hypothesize that this visualization is less useful for small, familiar models.

Once leaves and roots were identified by the application, participants had an easier time selecting initial values for analysis when compared to the previous study rounds. Typically, participants went from highlighted intention to intention, deciding what value the intention should have in the current analysis question.

In the inflo case, when leaves or roots were identified, this prompted changes, adding more incoming contributions to some sparsely connected roots. In several of the studies, analysis over the model began with more initial analysis values when compared to the previous study rounds. It can be argued that adding extra constraints or initial values to the analysis produces richer, more useful results over the model.

Finally, an additional unexpected benefit emerged when in a few cases root and leaf highlighting prompted changes to model layout. Roots were moved to the top of actors while leaves were moved to the bottom, if they were not already in these positions. To be fair, most of the layout changes were prompted by the facilitator, but in all cases the participants agreed with these changes, with one participant saying it made the model more organized, making it easier to see the structure of the model.

B. *Conflict Highlighting*

Results concerning conflict highlighting show that this intervention is helpful in understanding model conflicts; however, a considerable knowledge of i* modeling and analysis is needed to completely understand the causes of the conflict. Often the participants were not able to understand the reasons for the conflict on their own, even with highlighting. In these cases the facilitator had to use the highlighting and assigned intention values (Fig. 8) to explain the underlying reason for the conflict to participants. This result echoes the results from [7]: although interactive analysis is helpful, in order to receive the full benefits, including model iteration and further elicitation, relies heavily on i* and analysis experience, or the presence of an experienced facilitator.

Despite the need for a high level of i* knowledge, highlighting of conflict intentions made it much easier for the facilitator to understand and explain conflicts in the model. All participants indicated that conflict highlighting was helpful in some way.

In several cases conflict results revealed interesting tradeoffs in the model, prompting the participant to make tradeoffs in their analysis decisions. For example, in one individual model conflict highlighting revealed a tradeoff between Networking with friends and Get a good job, with choices over Study Hard and Do an internship. In another study, conflict highlighting revealed a tradeoff between Distributing Materials for the Publicity Chair and Low cost for the Treasurer, with the means to distribute materials increasing cost. In this case the participant lowered the analysis target value for low cost from *satisfied* to *partially satisfied*, and selected only one means of Distributing Materials.

In some cases understanding of conflicts prompted changes in the model, although in these cases the changes were suggested by the facilitator. For example, in one individual study, in the first run of the backward analysis, a conflict was found immediately without human judgment. The facilitator suggested that this result was due to the "all or nothing" nature of some decomposition links in the model, suggesting that these links may be better represented by contribution links. In this case the participant agreed and changes were made. The next run of the analysis found a solution in the model.

Results also revealed that the logical source of the conflict (defined in IV.B) is not necessarily equivalent to the "conceptual source". Here the logical source is the most immediate or direct cause of the conflict (e.g. PS and not PS) as reported in the UNSAT core, while the conceptual source is the construct in the model which is the originating source of the conflict, i.e. if this were removed/changed, the conflict would disappear.

We see an example of this in the inflo conflict example in Fig. 5. Here, the conflict sources include several tasks, while the conceptual source is the *hurt* link from Type checking and conversion to Flexibility. In another example in one of the ICSE Greening models (Fig. 4), the logical source of a conflict was Update web page while the conceptual source was the *unknown* link from Prepare materials to Attractive materials which constrained Prepare materials from having a denied value. In these cases it may not make sense to remove or change the conflict source construct in the model – perhaps Type checking and conversion really does *hurt* Flexibility. In such cases a conflict in the domain is revealed. Future work should investigate methods to suggest conceptual sources for model conflicts, perhaps highlighting negative or unknown links along the path of the conflict.

Results also revealed that conflict highlights should be left in the model for a longer duration of time, perhaps until after the user has completed the next backtracking stage. One participant mentioned that they were having difficulty remembering where the previous conflicts occurred when they were making judgments. They suggested leaving "traces" or some type of list of past conflicts.

C. *Threats to Validity*

Although the results of our studies were a useful first step towards testing the utility of the new visualization to support interactive goal model analysis, several threats to validity exist. As the subjects of our study were all students, all with some exposure to modeling in systems development, it is possible that results would not generalize to participants with different backgrounds or experiences. It is possible that the subject matter of the models or the size of the models may have affected results; however, we have tested the visualizations over models in three domains with a variety of sizes. Finally, the study designers and facilitators are the inventors of the analysis procedure and visualizations, possibly introducing bias via their presence in the room, their guidance as facilitators, and through their interpretation of the results. We have tried to minimize bias whenever possible, encouraging participants to be honest, and reporting both positive and negative results.

D. *Study Design*

Instead of performing case studies using qualitative analysis of results, we could have designed and conducted experiments using some sort of quantitative measure. This form of testing would be difficult for two primary reasons. First, it is difficult to find participants who have sufficient knowledge of i* to be able to apply and understand analysis, making it difficult to collect enough data to produce statistically significant results. Second, it is difficult to find trustworthy quantitative measure for the comprehension and utility of our visualizations. Such measures would not tell us why a visualization mechanism was good or bad, or lead us towards more useful designs.

Future work could take an action research approach to testing visualizations. Interactive analysis could be applied to a more realistic problem in an organization, noting any benefits of the visualizations and recording ideas for improvements.

VII. RELATED WORK

Any goal modeling framework which provides a visual syntax could be considered as a type of requirements visualization. However, work which focuses specifically on the visual aspects of goal modeling is especially relevant. Several techniques focus on using goal model visualizations in novel ways without making significant changes to the typical visual representation of goal graphs. For example, Sen and Jain use visualizations of activity cards and compilation tables to allow stakeholders to create and review goal decomposition hierarchies [14]. Work in [15] uses i* models to capture various levels of concerns in dynamically adaptive systems without modification of i* visualization techniques. The approach of Rohleder provides a visualization technique to view the effects of Non-functional Requirements (NFRs) represented as goals on software variants [16].

Other work suggests modifications to the ways in which goal models are typically presented. For example, Ernst et al. uses various visualization interventions (e.g., line thickness, color, size) to visualize levels for software qualities such as trust and feasibility over goal models [17]. Results of quantitative goal analysis from [2] are displayed beside intentions, but are not otherwise used to affect visualization. Although the color intervention is similar to the visualizations used in our work, the meaning of the colors differs.

The jUCMNav tool for drawing GRL (Goal-oriented Requirements Language) also makes use of color visualizations. In this case varying shades of red and green are used to represent the level of satisfaction and denial resulting from automated qualitative and quantitative analysis [3]. These visualizations are potentially compatible with those introduced in the current work, assuming that new mechanisms for highlighting leaf /root and conflict intentions can be found (e.g., increasing size, thicker borders).

Further work focuses on improving i* syntax. Moody et al. analyzes i* syntax using principles for visual notation design such as semiotic clarity (one graphical symbol for each semantic concept) [18]. Although suggested changes to i* syntax may improve i* visualizations, our focus is on finding visualizations to support analysis, and not on improving the underlying visual representation of i*.

VIII. Conclusions

Although analysis over goal models can be helpful to better understand and improve models, increasing understanding of the domain, difficulties in model and analysis comprehension exist. We claim that these difficulties can be partially mitigated through visualization techniques. To our knowledge, no other work specifically addresses visualization aspects of goal model analysis.

We have described and tested visualization techniques to address two specific analysis usability issues: finding starting points for analysis and understanding model conflicts. Features allowing highlighting of model leaves and roots are provided. Mechanisms to highlight intentions involved in or acting as the logical source of conflicts are described, including relevant formal definitions of useful concepts.

Results of follow-up studies showed that the presence of a facilitator may be necessary to understand highlighted conflicts, or to identify areas of useful model changes identified through the visualizations. However, the visualizations helped users to more easily identify starting points for analysis, asking more complete analysis questions. Explicitly identifying leaves and roots can help to improve the layout of the model. Conflicts were more easily explained to participants with the presence of highlighted intentions, and interesting tradeoffs in the domain were identified by participants.

Future work should continue to find ways to enhance goal model analysis through visualization techniques. Further visual enhancements can be made specific to the interactive analysis procedures used in this work. For example, when a human judgment question is asked, all intentions involved in the question could be highlighted or otherwise marked in the model. When changes are made to the model or to human judgments, intentions whose analysis values may be affected by the change could be highlighted for the user. We are currently in the process of adding these visualizations to our analysis implementation.

Acknowledgment

Financial support has been provided by the Natural Sciences and Engineering Research Council of Canada and the Ontario Graduate Scholarship Program.

References

[1] E. Yu, "Towards Modelling and Reasoning Support for Early-Phase Requirements Engineering" Proc. 3rd IEEE Int. Symp. on Requirements Engineering (RE'97), IEEE Press, pp. 226-235.

[2] P. Giorgini, J. Mylopoulous, and R. Sebastiani, R, "Goal-Oriented Requirements Analysis and Reasoning in the Tropos Methodology," Engineering Applications of AI, vol. 18(2), 2005, pp. 159-171.

[3] D. Amyot, S. Ghanavati, J.Horkoff, G. Mussbacher, L. Peyton, and E.Yu, "Evaluating Goal Models within the Goal-oriented Requirement Language," Int. J. of Intelligent Sys. (IJIS), vol. 25, issue 8, Aug. 2010, 841-877.

[4] J. Horkoff, E. Yu, "A Qualitative, Interactive Evaluation Procedure for Goal- and Agent-Oriented Models," Proc. International Conference on Advanced Information Systems Engineering Forum (CAiSE Forum'09). CEUR Workshop Proceedings, 2009, pp. 19-24.

[5] J. Horkoff, and E.Yu, "Evaluating Goal Achievement in Enterprise Modeling – An Interactive Procedure and Experiences," Proc. 2nd IFIP WG 8.1 Working Conf. on The Practice of Enterprise Modeling (PoEM'09), Springer, Nov. 2009, LNBIP, vol. 39, pp. 145-160.

[6] J. Horkoff, E. Yu, "Finding Solutions in Goal Models: An Interactive Backward Reasoning Approach" Proc. 29th Int. Conf. on Conceptual Modeling (ER'10), 2010, in press.

[7] J. Horkoff, E. Yu, A. Ghose, "Interactive Goal Model Analysis Applied - Systematic Procedures versus Ad hoc Analysis" submitted to 3rd IFIP Working Conference on The Practice of Enterprise Modeling (PoEM'10) unpublished.

[8] OpenOME, https://se.cs.toronto.edu/trac/ome/wiki.

[9] i* Wiki, http://istar.rwth-aachen.de/.

[10] L. Chung, B.A. Nixon, E. Yu, and J. Mylopoulos, Non-Functional Requirements in Software Engineering. Norwell, MA. Kluwer Academic Publishers, 2000.

[11] J. Cabot, S. Easterbrook, J. Horkoff, J. Mazon, L. Lessard, S. Liaskos, "Integrating Sustainability in Decision-Making Processes: A Modelling Strategy," Proc. 31st International Conference on Software Engineering Compantion Volume, New Ideas and Emerging Results (NIER'09), IEEE Press, 2009, pp. 207-210.

[12] M. Strohmaier, J. Horkoff, E. Yu, J. Aranda, S. Easterbrook, "Can Patterns improve i* Modeling? Two Exploratory Studies," Proc. International Working Conference on Requirements Engineering: Foundation for Software Quality (REFSQ'08), Springer, 2008, LNCS, vol. 5025, pp. 153-167.

[13] J. Horkoff, E. Yu, L Liu, "Analyzing Trust in Technology Strategies," Proc. International Conference on Privacy, Security and Trust (PST'06), ACM, 2006, pp. 21-32.

[14] A. Sen, S. K. Jain, "A Visualization Technique for Agent Based Goal Refinement to Elicit Soft Goals in Goal Oriented Requirements Engineering," Proc. Second International Workshop on Requirements Engineering Visualization (REV'07), IEEE Press, Oct. 2007, pp. 2.

[15] P. Sawyer, N. Bencomo, D. Hughes, P. Grace, H.J. Goldsby, B.H.C. Cheng, "Visualizing the Analysis of Dynamically Adaptive Systems Using i* and DSLs," Proc. Second International Workshop on Requirements Engineering Visualization (REV'07), Oct. 2007, pp. 3.

[16] C. Rohleder, "Visualizing the Impact of 4on-Functional Requirements on Variants : A Case Study," Proc. Third International Workshop on Requirements Engineering Visualization (REV'08), Sept. 2008, pp.11-20.

[17] N. Ernst, Y. Yu, J. Mylopoulos, "Visualizing non-functional requirements," First International Workshop on Requirements Engineering Visualization (REV'06), 2006, pp. 2.

[18] D. L. Moody, P. Heymans, R. Matulevicius, "Improving the Effectiveness of Visual Representations in Requirements Engineering: An Evaluation of i* Visual Syntax," Proc.17th IEEE International Requirements Engineering Conference (RE'09), IEEE Press, Sept. 2009, pp. 171-180.