

The mysteries of goal decomposition

Scott Munro¹ Sotirios Liaskos² Jorge Aranda³

¹ Department of Philosophy, York University
Toronto, Canada
scmunro@yorku.ca

² School of Information Technology, York University
Toronto, Canada
liaskos@yorku.ca

³ Department of Computer Science, University of Victoria
Victoria, Canada
jaranda@uvic.ca

Abstract. Goal decomposition structures lie at the heart of goal modeling languages such as *i**. High-level goals of stakeholders are recursively decomposed into lower level ones and eventually into leaf level tasks to be performed by agents. The decomposition structure can also develop through a bottom up approach whereby higher-level goals are introduced as justifications for existing low-level ones. The very concept of decomposition, however, both as process and as artefact is rarely questioned in requirements engineering. In this paper, we argue that it may be of value to give a closer look into goal decomposition and clarify what we actually know about it and what is yet to be understood. We report on an on-going effort to identify empirical work on decomposition coming from various research fields, hoping to find such evidence. We then pose some research questions that we believe need to be pursued in order to improve our understanding of goal decomposition.

Key words: requirements engineering, goal modeling, i-star

1 Introduction

Goal decomposition is central in goal modeling. High-level goals of stakeholders are recursively decomposed into lower level ones and eventually into leaf level actions, a subset of which is potentially to be performed by a machine agent. The reverse bottom-up process of discovery of high-level goals as explanations of existing tasks complements the top-down one. The result, the goal decomposition model, is a central piece of a goal model such as an *i** strategic rationale diagram. The decomposition structure serves the purpose of connecting stakeholder desires with system functions and has been found to serve many benefits [14].

But where do decompositions come from and how exactly are they used? While the literature offers a wealth of case studies illustrating the benefits of using goal models, proposing general processes for developing models, or discussing meta-models and ontological treatments, little seems to be known about

goal decomposition as a cognitive process. If we knew more about the nature of decompositions we would be able to allow for more natural and systematic ways of identifying and using such, increasing the quality and usefulness of goal models such as i^* models.

In this paper, inspired from some early experimental results on the matter of goal decomposition, we describe some highlights from our on-going work to understand what the literature in a variety of fields has to say about decomposition. We then describe the particular questions we are hoping an empirical research program on decomposition could attempt to answer.

The presentation is organized as follows. We provide an overview of our research goals and questions (Section 2), present some highlights from the literature we have been studying (Section 3) and then describe what we have learned and how we intend to proceed (Sections 4 and 5).

2 Objectives of Research

Central to the emergence of the goal-oriented paradigm is the fact that it offers a clean approach to connect the problem domain with the solution domain through recursive decompositions of expressions of the former (goals) into expressions of the later (actions). The literature proposes several approaches and techniques for developing such decompositions. For example, KAOS offers AND-decomposition patterns based on temporal semantics of the goals to be decomposed [3], while Rolland et al. [12] propose scenario-goal coupling to guide decomposition, an idea also used in Liaskos et al. later [8]. More general model development and enrichment methods, by e.g. informing development from other sorts of models have been proposed, e.g. [6].

However, although in abstract terms we can think of a decomposition task as a complete and sufficient activity to manage goals in order to make them more easy to tackle, in practical terms, we have people doing the decomposing and people that need to judge whether the decomposed goals fit the higher-level goal exactly - i.e. when a decomposition is “good”. Thus, the result of goal decomposition activities is inextricably linked to cognitive and psychological considerations and is influenced by the mental process that is followed (or lack thereof). Our experience in developing goal models, indeed suggests that not only different approaches and different people produce different decomposition models, but that understanding what the ideal model for a particular situation is can be very puzzling.

Our early results from a pilot experimental study were illuminating as to how different modelers, when given the same exactly input produce different decompositions. In this study we invited a number of senior students of Information Technology to develop goal models for solving a particular problem. The students had similar backgrounds and no working knowledge of goal modeling. They were all shown the same short video presentation on goal decomposition – as a top-down process. Then they were given a short textual description of a problem (a version of the classic Meeting Scheduling problem) and about 40 minutes to




Fig. 1. Different Modelers Produce Different Models

develop goal models identifying alternative solutions for the set problem. The problem statement was illuminating as to what concerns and alternatives were to be addressed (e.g. different communication media). The participants produced different models. The sizes ranged from 18 to 58 nodes and from 7 to 22 decompositions, often very different in structure and approach. Figure 1 depicts parts of two such models. These particular models exhibit a great degree of similarity (relative to other possible pairs). They are not the same however: they arrive to a slightly different set of leaves (relative to the other pairs in our sample) and they have some structural and semantic (wrt. the goal descriptions) dissimilarities. Their differences are instructive. Modeler A lists three mechanisms to obtain information from participants; modeler B includes four, and these are slightly different (both include phone and email, A includes paper, B includes ‘in person’ requests, and B also includes checking the agenda of the participant—a task that A believes is disconnected from the process of talking to participants). Notice that, for B, “talking to participants” is expressed as “asking for availability”. Also, modeler A does not consider a follow-up as part of their goal model. Finally, the top goals are different. B is only concerned with gathering participant availability, while A’s top goal is to find available dates (despite the fact that the two means by which A plans to fulfill that goal would not necessarily lead to its satisfaction).

We find these slight differences quite curious: why do they occur and what is the impact of each choice? Which of the produced decomposition models is better? An answer to this would assume a clear criterion of “betterness” of one decomposition over the other. It appears that one needs to quantify into measurable criteria the supposed reasons for using goal models, such as pertinence and sufficient completeness of the resulting requirements, comprehensibility of the decomposition per se, its appropriateness in explaining requirements to users or in allowing better readability of complex requirements documents [14].

But even if we are able to perform “betterness” comparisons does this give us a deep understanding of the *nature* of decomposition? What we find central in understanding goal decomposition models is a study of the very low-level human activities that relate to creating and using decompositions. How do humans produce decompositions? How do they read and understand them? What aspects influence decomposition development and understanding and how? And are those differences across modelers truly problematic, or should they be seen

as an opportunity for them to learn more about the subtleties of their domain and about the perceptions of their peers [5]? In our endeavour to shed some light on these questions and inform further experimentation, we began by looking for relevant empirical evidence in a variety of research areas outside requirements engineering – we present some examples below.

3 Understanding Decomposition

Our preliminary literature search was performed in an attempt to isolate any empirical work done concerning decomposition. Much of the literature we found is on the benefit of pre-constructed decompositions. Armstrong et al., for example [1], are frequently cited as empirical evidence for the usefulness of decomposition as a tool for assisting human judgment. In this study, a number of students were split into two equal groups and assigned the task of solving a problem. The problem was either presented as a single non-decomposed question, e.g. *“how many packs (rolls) of Polaroid color films do you think were used in the United States in 1970?”* or presented as a decomposed set of questions: *“a) how many people do you think were living in the U.S. in 1970?” “b) in 1970 what do you think was the size of the average family living in the U.S.?” “c) in 1970 what percentage of the families do you think owned cameras?”* etc. It was found that more accurate answers were produced by the students who received the decomposed problem. A similar study was performed by Dennis et al. [4] in which multipart questions were presented as a whole, e.g., *“what do you feel are the most important outputs from, inputs to, and data elements in the proposed computer system for Ace Video Rental?”* Or presented one at a time as individual questions, e.g., *“what do you feel are the most important inputs to the proposed [...] ?”*. It was found that more ideas were generated by the groups that received the individual decomposed questions. Neither study, however, speaks about the generation of the decomposition and what the impact to finding the right solution is when the decompositions are improper (incomplete, misleading etc.).

Lyness & Cornelius [9] performed a complex experiment in which students were asked to judge the quality of hypothetical college professors using either decomposed or non-decomposed methods. The non-decomposed method was to give a single rating out of 7 on a Likert scale. The decomposed method had the students rating for many (3, 6, or 9) dimensions given to them, e.g., knowledge of subject, grading philosophy, and testing procedures. These were then placed in weighted combinations. It was found that the subjects using the weighted decomposed method offered more reliable results. In terms of goal decomposition, that could potentially imply that breaking down and aggregating satisfaction criteria could lead to more accurate satisfaction assessment. Note, however, that presence of pre-existing decomposition is assumed.

Decomposition has been shown to be coarsely effective in this same sense in other areas as well. Hertz [7] shows it in capital investments, Polya [10] as a basis for mathematic problem solving, and Raiffa [11] in management science. These works however focus mostly on the application of decomposition to these areas

rather than offering strong empirical evidence. Our understanding so far is that research on the matter has surprisingly been abandoned in the recent years.

4 Conclusions

The empirical results that decomposition appears to help correct problem solving and assessment can be argued to support the utility of goal models in requirements engineering, as well. Thus top-level decomposition may guide correct identification of lower-level ones and eventually of leaf level tasks. A fundamental difference however is that in goal modeling it is assumed that the same agent produces the decomposition and continues with solving the problem (i.e. introduce more decompositions or operationalizations) while in the experimental studies an authoritative high-level decomposition is considered to be a given. In that regard, assuming that assistance in solving a problem also implies assistance in understanding an existing solution, the experimental results may also suggest that decompositions appropriately prepared by analysts are effective in communicating their solutions to e.g. clients and other stakeholders. Of course, the nature of the problem at hand (e.g. to find an answer to a unique question vs. to find an optimal set of requirements) seems to require consideration when utilizing those results.

Our sense is that many questions regarding decomposition are yet to be addressed. For example, why does decomposition of a problem aid in solving it? Is some particular sort of decomposition a natural problem solving method, while other sorts less enabling or even obstructing of the problem-solving process? Are there other representations that are not hierarchical and aid effectiveness even more? What is “effectiveness” after all, especially in less precise problems, which are typically the requirements engineering ones?

5 On-going and Future Work

We intend to continue our effort to better define the empirical research agenda on decomposition. Our understanding so far suggests three fronts on which such study needs to be performed. Firstly, define more rigorously the various uses of decomposition, which may be competing with respect to a “goodness” measure. For instance, the use of decomposition by one person as an aid to solve a problem (e.g. the way pen and paper are an aid to performing complex multiplications) may result in a model that is not optimal for communicating a solution to somebody else. Secondly, understand the role of the process and the agent in the decomposition: what is the difference between a structured approach [12, 3, 8] and a free-form approach, with respect to identified betterness criteria? How do agent characteristics affect the result and/or understanding thereof? Thirdly, how does our acquired understanding of decompositions relate to i^* or other goal modeling meta-models? Is i^* 's means-ends and e.g. KAOS's OR-decomposition ontologically the same and if not what is the impact of their differences in developing them? Is OR-decomposition truly a decomposition (as AND-decomposition is)

or a mere way to express alternative AND-decompositions – which effectively necessitates investigating the two as separate kinds of concepts? Efforts to clarify those aspects are well under way (e.g. [2, 13]) further motivating the question of what empirical research program can inform and be informed by these works.

References

1. J. S. Armstrong, W. B. Denniston, and M. M. Gordon. The use of the decomposition principle in making judgments. *Organizational Behavior and Human Performance*, 14(12):257–263, 1975.
2. Carlos Cares and Xavier Franch. A metamodelling approach for i^* model translations. In Haralambos Mouratidis and Colette Rolland, editors, *Advanced Information Systems Engineering*, volume 6741 of *Lecture Notes in Computer Science*, pages 337–351. Springer Berlin / Heidelberg, 2011.
3. Robert Darimont and Axel van Lamsweerde. Formal refinement patterns for goal-driven requirements elaboration. In *Proceedings of the 4th ACM SIGSOFT symposium on Foundations of Software Engineering (SIGSOFT '96)*, pages 179–190, New York, NY, USA, 1996. ACM Press.
4. A. R. Dennis, J. S. Valacich, T. Connolly, and B. E. Wynne. Process structuring in electronic brainstorming. *Information Systems Research*, 14(12):268–277, 1996.
5. Steve Easterbrook, Eric Yu, Jorge Aranda, Yuntian Fan, Jennifer Horkoff, Marcel Leica, and Rifat Abdul Qadir. Do viewpoints lead to better conceptual models? an exploratory case study. In *Proceedings of the 13th IEEE International Requirements Engineering Conference (RE'05)*, pages 199–208, 2005.
6. Xavier Franch, Gemma Grau, Enric Mayol, Carme Quer, Claudia P. Ayala, Carlos Cares, Fredy Navarrete, Mariela Haya, and Pere Botella. Systematic construction of i^* strategic dependency models for socio-technical systems. *International Journal of Software Engineering and Knowledge Engineering*, 17(1):79–106, 2007.
7. D. B. Hertz. Risk analysis in capital investment. *Harvard Business Review*, pages 86–106, 1964.
8. Sotirios Liaskos, Alexei Lapouchnian, Yijun Yu, Eric Yu, and John Mylopoulos. On goal-based variability acquisition and analysis. In *Proceedings of the 14th IEEE International Requirements Engineering Conference (RE'06)*, Minneapolis, Minnesota, September 2006. IEEE Computer Society.
9. K. S. Lyness and Cornelius E. T. A comparison of holistic and decomposed judgment strategies in a performance rating simulation. *Organizational Behavior and Human Performance*, 29(1):21–38, 1982.
10. G. Polya. *How to Solve It*. Princeton: Princeton University Press, 1945.
11. H. Raiffa. *Decision analysis: Introductory Lectures on Choices Under Uncertainty*. Addison-Wesley, 1968.
12. Colette Rolland, Carine Souveyet, and Camille Ben Achour. Guiding goal modeling using scenarios. *IEEE Transactions on Software Engineering*, 24(12):1055–1071, 1998.
13. Angelo Susi, Anna Perini, John Mylopoulos, and Paolo Giorgini. The Tropos metamodel and its use. *Informatica*, 29:401–408, 2005.
14. Axel Van Lamsweerde. Goal-oriented requirements engineering: A guided tour. In *Proceedings of the Fifth IEEE International Symposium on Requirements Engineering (RE'01)*, page 249, 2001.