

A Survey on the Relevance of the Performance of Model Transformations

Raffaella Groner,¹ Katharina Juhnke,¹ Stefan Höppner,¹ Matthias Tichy,¹ Steffen Becker,² Vijayshree Vijayshree,² Sebastian Frank²

Abstract: When we are confronted with performance issues in a general-purpose language, like Java, it is a given to us that we have various tools and techniques at our disposal to help us. But is such support also needed when using model transformation languages? To address this question, we conducted a quantitative online survey as part of a mixed methods study with 84 respondents to our questionnaire. Our results show that a certain performance is desired but not always achieved. The developers would like to improve the performance, but they lack insights on how a transformation is performed. As a first step to mitigate this issue, we compiled a list of information regarding the models used, the transformations applied and their execution deemed to be helpful by the participants. Additionally, we used hypotheses tests to investigate possible influencing factors that cause participants to try to improve the performance of transformations. The main relevant factors found in our study are the satisfaction with the execution time, the size of the models used, the relevance of whether a certain execution time is not exceeded in the average case, and the knowledge of how a transformation engine executes a transformation.

Keywords: Model transformation; Performance; Survey; ATL; Henshin; QVTo; Viatra

Techniques for analyzing performance issues of programs written in a general-purpose language are widely used. Due to the increasing popularity of model-driven software development and the use of models at runtime, the question arises whether such a support is also necessary for transformation languages? Currently, research in this area mainly focuses on improving the execution of transformations, such as the work of Boronat [Bo18]. To address this question, we conducted a quantitative online survey as part of a mixed methods study with 84 respondents to our questionnaire.

Our results show that developers using transformation languages also face performance issues, but currently lack support to examine them properly. For example, only 13 of the 84 participants are always satisfied with the execution time of their transformations. 37 participants are often satisfied and 44 participants are only sometimes or rarely satisfied. Using the Mann-Whitney-U-Test, we investigated the differences between participants who have already tried to analyze or improve the performance and those who have not. The test results show that there is a significant difference between these two groups in four aspects: Participants who have tried to improve the performance tend to 1) be less satisfied with the

¹ Ulm University, Institute for Software Engineering and Programming Languages, James-Franck-Ring, D-89069 Ulm, Germany {raffaella.groner, katharina.juhnke, stefan.hoeppner, matthias.tichy}@uni-ulm.de

² University of Stuttgart, Institute for Software Engineering, Universitätsstraße, D-70569 Stuttgart, Germany {stefen.becker, vijayshree.vijayshree, sebastian.frank}@iste.uni-stuttgart.de

execution time, 2) use larger models, 3) deem it more important that a certain execution time is not exceeded in the average case, and 4) have more knowledge of how a transformation is executed by the engine. In particular, the knowledge about the transformation execution seems to be an essential influencing factor that causes participants to try to improve the performance of transformations, as we also discovered a very high demand for information. This includes overview information, like the number of applications per transformation rule, or detailed information, like which elements of the of the input model were transformed using which transformation rule. In addition, the developers would like to have more information about the structure of the models used, such as the depth of the inheritance hierarchy or the cycles they contain. Developers also want more support, in the form of profilers or analyses. For example, an analysis that helps to determine an optimal order for applying transformation rules was mentioned.

The results of our study show that techniques for analyzing the performance of transformations are needed and give an overview of the desired information that can help to support developers. More details about our study can be found in our corresponding publication, which was accepted by the Journal of Object Technology (JOT) in 2021 [Gr21a].

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Ti 803/4-1 and BE 4796/3-1

Data Availability

In [Gr21b], we published the following data: 1) Documentation of our search for suitable study participants. 2) The questionnaire used. 3) Additional material used for the design of the questionnaire. 4) The anonymized answers to our questionnaire. 5) Documentation of the results of the hypothesis tests and the corresponding SPSS project.

Literaturverzeichnis

- [Bo18] Boronat, Artur: Expressive and Efficient Model Transformation with an Internal DSL of Xtend. In: Proceedings of the 21th ACM/IEEE International Conference on Model Driven Engineering Languages and Systems (MODELS' 18). ACM, New York, NY, USA, S. 78–88, 2018.
- [Gr21a] Groner, Raffaella; Juhnke, Katharina; Götz, Stefan; Tichy, Matthias; Becker, Steffen; Vijayshree, Vijayshree; Frank, Sebastian: A Survey on the Relevance of the Performance of Model Transformations. Journal of Object Technology, 20(2):2:1–27, 2021.
- [Gr21b] Groner, Raffaella; Juhnke, Katharina; Götz, Stefan; Tichy, Matthias; Becker, Steffen; Vijayshree, Vijayshree; Frank, Sebastian: , A Survey on the Relevance of the Performance of Model Transformations: Data of the Participant Search and the Questionnaire. Open Access Repository of the Ulm University, 2021.