

# How Dataflow Diagrams Impact Software Security Analysis: an Empirical Experiment

Simon Schneider\*, Nicolás E. Díaz Ferreyra\*, Pierre-Jean Quéval†, Georg Simhandl†,  
Uwe Zdun†, Riccardo Scandariato\*

\*Hamburg University of Technology, *firstname.lastname@tuhh.de*

†University of Vienna, *firstname.lastname@univie.ac.at*

**Abstract**—Models of software systems are used throughout the software development lifecycle. Dataflow diagrams (DFDs), in particular, are well-established resources for security analysis. Many techniques, such as threat modelling, are based on DFDs of the analysed application. However, their impact on the performance of analysts in a security analysis setting has not been explored before. In this paper, we present the findings of an empirical experiment conducted to investigate this effect. Following a within-groups design, participants were asked to solve security-relevant tasks for a given microservice application. In the control condition, the participants had to examine the source code manually. In the model-supported condition, they were additionally provided a DFD of the analysed application and traceability information linking model items to artefacts in source code. We found that the participants ( $n = 24$ ) performed significantly better in answering the analysis tasks correctly in the model-supported condition (41% increase in analysis correctness). Further, participants who reported using the provided traceability information performed better in giving evidence for their answers (315% increase in correctness of evidence). Finally, we identified three open challenges of using DFDs for security analysis based on the insights gained in the experiment.

**Index Terms**—security, analysis, dataflow diagrams, microservices, model-based, empirical experiment

## I. INTRODUCTION

Dataflow Diagrams (DFDs) are integral to many software security analysis techniques. For instance, they are required by prominent security assessment techniques [1]–[4]. They are also the program representation chosen in many model-based security approaches [5]–[12]. As such, they can help software engineers build more secure software systems. However, their impact on security analysis by their mere provision has not been investigated before to the best of our knowledge. DFDs offer a high-level yet detailed representation of applications’ architecture. Enriched with annotations representing, e.g., employed security mechanisms, they offer easy accessibility of the architectural security. We hypothesize that providing users with a DFD of an application enables them to analyse the application’s security properties with higher correctness.

To investigate this hypothesis, we conducted an empirical experiment. This paper reports on our findings. The experiment was performed with master students who solved tasks related to software security analysis activities. For this, they received the source code and a textual description of an open-source microservice application, and six tasks to answer. We chose microservice applications as the target of analysis because the

distributed nature of this architectural style poses additional challenges in terms of cognitive load to security analysts. Systems following the microservice architecture split their codebase into multiple independent microservices, where each fulfils a part of the business functionality and can be independently developed and deployed [13], [14]. The resulting codebase can be challenging to oversee in analysis scenarios.

The experiment followed a within-groups design. In the model-supported condition, participants received a DFD and traceability information of the analysed application in addition to the source code and textual description provided in the control condition. We chose DFDs created by *Code2DFD*, a tool for the automatic extraction of DFDs from source code [15], as these contain extensive security annotations. We infer insights on the impact of DFDs on security analysis by comparing the participants’ performance in analysis correctness, correctness of evidence, and time in the two conditions. Specifically, we answer the following research questions:

**❓ RQ 1:** Do security-annotated architectural models of applications, specifically DFDs, support developers in the security analysis of the applications?

A crucial part of security analysis is identifying and localizing security (and other) features in source code. To assess whether models with extensive security annotations can support developers and security experts in this activity, the participants in our experiment solved tasks that require the identification of implemented security mechanisms and other relevant system properties. We quantified their answers and compared the scores between the two conditions. Additionally, we asked them about the perceived usefulness and analysed the answers.

**❓ RQ 2:** Does access to and use of traceability information improve the ability of the security analysts to provide correct evidence for their analysis findings?

Traceability information establishes the validity of model items by referencing corresponding artefacts in the source code. It can provide value for security analysis, since in scenarios such as software security assessment or certification, evidence has to be given for the reached findings. The participants provided evidence for their answers in the form of locations in the source code. We examined its correctness and compared the scores of those participants who reported using the traceability information frequently and those who did not.

**RQ3:** What is the experience in using security-annotated DFDs for security analysis, specifically concerning the usefulness and accessibility?

Users' acceptance of offered tools and techniques is crucial for using resources such as DFDs for security analysis. Thus, the participants' perceived usefulness of the DFDs is essential to judge their suitability for real-world application. Further, the information presented by DFDs has to be conveyed to the users efficiently and accessible. To judge this aspect, we asked the participants about their experience using the DFDs in the model-supported condition and analysed their responses.

**RQ4:** What are the open challenges of using security-annotated DFDs in the context of security analysis?

Based on the insights gained during the empirical experiment and the analysis of the results, we identified and formulated open challenges that should be addressed in future work.

The rest of this paper is structured as follows: Section II introduces the used DFDs; Section III describes the experiment's design, i.e., the methodology; the results are presented in Section IV and discussed in Section V; Section VI describes limitations of this work; Section VII presents related work; and Section VIII concludes the paper.

## II. CHARACTERISTICS OF THE USED DFDs

Since no standard specification for DFDs exists, the various styles of DFDs found in the model-based literature differ in their characteristics. Here, *style* refers to the types of model items, their presentation and richness of detail, and the scope of considered system components. All DFD styles share the four base item groups *external entities*, *data flows*, *processes*, and *data stores* [16]. Many approaches include additional model items to increase the models' expressivity [1], [9], [17]. In our experiment, we used DFDs from a dataset published by Schneider et al. [18]. An example is shown in Figure 1. The style is the same as those generated by an automatic extraction approach by Schneider and Scandariato [15]. Two of the DFDs' properties stand out compared to other styles, the included *annotations* and the *traceability information* for model items. Annotations in the DFDs provide information about the corresponding system's security and other properties. The annotations represent implemented security mechanisms (e.g., encryption or authorization mechanisms), deployment information (e.g., ports or addresses), or other system properties (e.g., used technologies and frameworks). Annotations are associated with model items from the four base item groups mentioned above. That means that base model items (such as, e.g., a service) are augmented with the annotations. Figure 1 shows examples of this (see the *--annotation--* and *Key: Value* annotations). Traceability information links model items to source code by pointing to code snippets that prove the existence of the model item. Figure 2 shows as an example the traceability information for the annotation *authorization\_server* as well as a screenshot of the target of the contained URL, i.e., the line of code on GitHub.

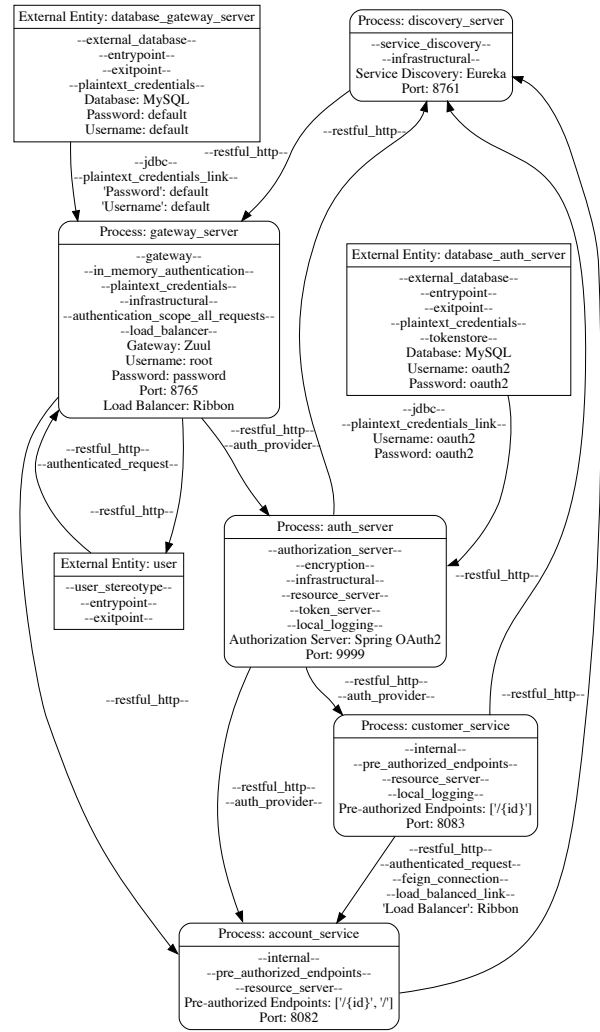


Fig. 1. Example DFD provided to participants in model-supported condition.

```

1  "authorization_server": {
2    "file": "github.com/[...]",
3    "line": "18",
4    "span": "(0,25)"
5  }

```

```

17  @Configuration
... 18  @EnableAuthorizationServer
19  public class OAuth2Config extends Auth

```

Fig. 2. Screenshots of an example traceability information for the annotation *authorization\_server* (top) and the target of the URL (bottom).

## III. STUDY DESIGN

We designed and conducted an empirical experiment to answer the formulated research questions. We consulted established sources of guidelines for empirical research in the design of the experiment ([19]–[21]). All materials as well as the results are available in the replication package [22].

### A. Setup

The experiment followed a within-groups design, where participants participated in both conditions. The study was performed in two 90-minutes lab sessions of a master's course

at a university in subsequent weeks. The participants were randomly assigned to one of the two groups, G1 and G2. A different application was given as the target of analysis in the two weeks to mitigate learning effects, *App 1* in week 1 and *App 2* in week 2. The two sessions were structured as follows:

	Week 1	Week 2
	App 1	App 2
G1	control condition	model-supported condition
G2	model-supported condition	control condition

In the *model-supported condition*, the participants performed tasks with access to a DFD and corresponding traceability information, whereas in the *control condition*, they performed a similar set of tasks without this additional support. They were supervised during the sessions and were discouraged from talking to each other about the experiment. Google Forms surveys were used to provide the tasks and gather the answers.

### B. Tasks

Table I lists the analysis tasks given to the participants. They were chosen to resemble common security analysis activities and relate to architectural security rules for microservices [23]. The first two tasks are not specific to security but are relevant nevertheless since they foster a required comprehension of the analysed system. The tasks cover three different kinds of questions, the participants had to find:

- general information about single services (tasks 1 & 2)
- information about security mechanisms of single services (tasks 3 & 4)
- information about system-wide security mechanisms (tasks 5 & 6)

For all tasks, the participants were asked to provide evidence for their answers via a reference to the code.

After completion of the technical tasks, the participants were also posed an open question about their experience with using the resources they had been provided. In the first week, they were further asked questions concerning their expertise.

As the target of evaluation for the security analysis, we chose two open-source applications from a list published by Schneider et al. [18]. The applications are referred to as App1<sup>1</sup> and App2<sup>2</sup>. These applications were selected based on two properties: (i) high architectural similarity between the two in order to enable an accurate comparison of participants' performance between the two sessions, and (ii) sufficient architectural complexity to allow insightful and relevant tasks. The two applications exhibit a high degree of similarity concerning their architecture, size, and used technologies. They incorporate some of the most prevalent microservice patterns and employ widely adopted technology solutions for Java-based microservice development. For instance, they both realize an API Gateway with Zuul, authentication with OAuth 2.0, a load balancer with Ribbon, and service discovery with Eureka. Consequently, the applications are a suitable representation of open-source microservice applications developed in Java.

<sup>1</sup>[github.com/anilallewar/microservices-basics-spring-boot](https://github.com/anilallewar/microservices-basics-spring-boot)

<sup>2</sup>[github.com/piomin/sample-spring-oauth2-microservices/tree/with\\_database](https://github.com/piomin/sample-spring-oauth2-microservices/tree/with_database)

### C. Provided Application Artefacts

All participants were provided the source code and a basic textual description of the analysed application. The textual description is an explanatory document that was created based on the code and information provided by the developers of the applications. To mitigate a potential influence on the experiment's outcome based on different qualities of the textual descriptions of the two applications, they were created in identical structure and contain the same basic information about the corresponding application. They illustrate the basic architectural design of the applications. We remark that some tasks could be answered based on these documents. The code was provided via GitHub. Specifically, the applications' repositories were forked to remove the original documentation, which could otherwise influence the outcome.

In addition to the code and textual description, the participants in the model-supported condition received the DFD and the traceability information of the application to be analysed.

### D. Participants

The experiment's participants comprised 24 students enrolled in a master's level software security course at *Hamburg University of Technology*. They originate from various disciplines, all incorporating computer science to a large degree. The students were informed about the empirical experiment two months in advance and were invited to participate. Figure 3 shows the participants' programming knowledge (a), experience with reading Java code (b), and work experience (c). This information was reported in a self-assessment at the end of the first week's session. Based on the results ("intermediate level" being the answer most often given to the question of programming skills; little experience with reading Java code; and an average work experience of 1.1 years), we deduced that the participants were on average advanced beginners in software development and had moderate experience in analysing Java code. Accordingly, they represent well the target population of the experiment. The goal of the experiment was to investigate the effect of architectural models on the performance of users with low expertise in software security analysis, for example novice developers. The metrics in Figure 3 fit well to such users. Furthermore, using students as proxies for the target population is a common practice in empirical software engineering and has been shown to be a suitable method [19], [24]–[26].

1) *Incentives*: The students were generally incentivized to participate in the course's lab sessions, independent of the empirical experiment. For this, they were rewarded with a 5% bonus on their final exam, granted upon attending all of the lab sessions in the semester except for one (a common practice at the university). Further, they were encouraged to participate because the sessions were relevant for the final exam, and the participants could hone the required skills there. The lab sessions where the empirical experiment was conducted were akin to all other lab sessions in these regards. Consequently, the experiment's impact on students' grades was consistent

TABLE I  
TASKS FOR APP 2. SERVICE NAMES AND NUMBER OF CONNECTIONS SLIGHTLY DIFFER FOR APP 1.

ID	Task Description
1	What is the port number of the microservice <i>gateway_server</i> ?
2	What library is used to implement the API Gateway (service <i>gateway_server</i> )?
3	There are two outgoing connections from the service <i>customer_service</i> . To which services are they connected?
4	Are these outgoing connections encrypted, i.e., sent with HTTPS?
5	Do all business logic services check whether incoming requests are authorized / authenticated?
6	Which service handles the authorization?




Fig. 3. Participants' (a) programming skills, (b) experience in reading Java code, and (c) work experience as developers. All self-reported.

with that of the other lab sessions in this course. No other incentives were pledged or given.

2) *Preparation*: To prepare the participants, a 90-minute lecture before the lab sessions was dedicated to introducing them to the topic (available in this paper's replication package [22]). The lecture covered key concepts relevant to the experiment. The primary focus was on the origin of software vulnerabilities and methods for detecting them. The lecture also encompassed topics such as DFDs, microservice architectures, and security considerations in microservice applications. Following this lecture, the students were expected to possess the required knowledge to undertake the experiment. Their attendance at the lecture was recorded and was a prerequisite for participating in the experiment.

3) *Informed consent and ethical assessment*: All participants read and signed an informed consent form before the experiment, informing them that they are the subjects of an empirical experiment, that they participate voluntarily, that they do not have to expect any negative consequences whatsoever if they do not participate, and that they can retract their consent at any time. To ensure the experiment's ethical innocuity, it was assessed by the *German Association for Experimental Economic Research e.V.* before execution. A description of the planned experiment and its design was approved under grant nr. *2pxo1bap*. The certificate can be accessed via <https://gfew.de/ethik/2pxo1bap>.

### E. Measurement

To evaluate the participants' performance, three metrics were introduced. The *analysis correctness* represents the ability to provide correct answers to the tasks. The *correctness of evidence* measures whether the evidence that the participants provided as support for their answers points to a code snippet that justifies their answer. Both are numerical scores derived from the participants' responses. Additionally, we measured the time spent on solving the tasks. The three metrics (analysis

correctness, correctness of evidence, and time) were calculated for each participant in both conditions separately.

1) *Analysis Correctness*: We quantified the given answers concerning the analysis correctness by manually checking the participants' responses. To remove subjectivity, we created a reference solution that was used to check the answers. It was created prior to the execution of the experiment. The DFDs and source code of the applications have been analysed to create the reference answers, which were afterwards confirmed by conducting technical documentation of the code libraries used in the applications, information provided by the developers of the applications, and other typical online resources. This process was performed by the first author and validated afterwards by two additional authors. After the experiment, the participants' responses were mapped via the reference solution to a table indicating correct and incorrect answers. Each response was examined manually, compared against the reference solution, and true answers were marked in the table. We further reviewed answers that did not match the reference solution to check whether they were correct. For this, various typical online resources were conducted to verify whether the specific answer applies to the task. Each correctly given answer gives a score of 1. There is a peculiarity for some tasks. Task 3 asks for a list of connections between services, and tasks 4 and 5 ask whether a property applies to each item on a given list. Consequently, these tasks each required multiple distinct responses. All responses were checked individually. Then, to allow a more detailed and nuanced evaluation, we converted the results to scores. A score of 0 was assigned for no correct responses, a score of 1 was given for partially correct responses (meaning that some but not all responses of a task were correct), and a score of 2 was awarded when all responses were correct. With three tasks giving a maximum of one point and three tasks giving a maximum of two points, the overall highest achievable score in analysis correctness is 9.

2) *Correctness of Evidence*: The traceability information that is contained in the used DFDs constitutes a reference solution for quantifying the correctness of the evidence given by the participants. Each given evidence was checked manually for matches to this reference solution. Here, we employed some tolerance in accepting evidence as correct. For example, when participants referred to a block of code slightly larger than the lines of code needed to prove an answer, we still accepted this as correct (e.g., referring to a method consisting of some lines of code instead of referring to a single line of code in that method). We carried out a further validation check, similar to the quantification of the analysis correctness. Each provided evidence that differed from the reference solution was checked manually whether it supported the given answer or not. The first author carried out the above steps. As for the analysis correctness, each correct evidence gives a score of 1. Again, for tasks 3, 4, and 5, the multiple distinct responses were converted into a score of 0, 1, or 2 for no correct, partially correct, and all correct responses, respectively.

3) *Time*: To measure the time spent on solving the tasks, the participants were asked to record the current time when starting and finishing to work on the tasks. We calculated the time metric based on these answers (i.e., the period of time between the start and finish of solving the tasks).

4) *Reported Usefulness*: The DFDs' usefulness as reported by the participants was assessed via the open question about the participants' experience in using the DFDs that was posed after the technical tasks. We qualitatively analysed all answers' general intents (positive/negative feedback) and identified recurring topics manually. This analysis was performed by the first author and verified by two further authors.

## F. Statistical Tests

Throughout the analysis, the difference of scores between two groups was checked for statistical significance with a Wilcoxon-Mann-Whitney test. Before, a Shapiro-Wilke test was used to verify that the data does not follow a normal distribution. Hence, no parametric tests could be used. The assumptions for the Wilcoxon-Mann-Whitney test (the samples are independent, random, and continuous and the sample size is sufficiently large) are met in our experiment.

## IV. RESULTS

### A. Analysis Correctness

Figure 4 presents the average score in analysis correctness that the participants achieved in the two conditions. The figure shows that the participants performed better in the model-supported condition, both overall and in every individual task. For all tasks, the average score is 6.75 out of a possible 9 in the model-supported condition compared to 4.79 in the control condition, a 41% higher average. The applied statistical test (compare Section III-F) indicates a statistically significant difference between the two conditions' average scores in analysis correctness overall ( $p = 0.0025$ ). These results provide the following answer to RQ1:




Fig. 4. Comparison across the two treatments of the participants' average score in analysis correctness. Per task and overall.

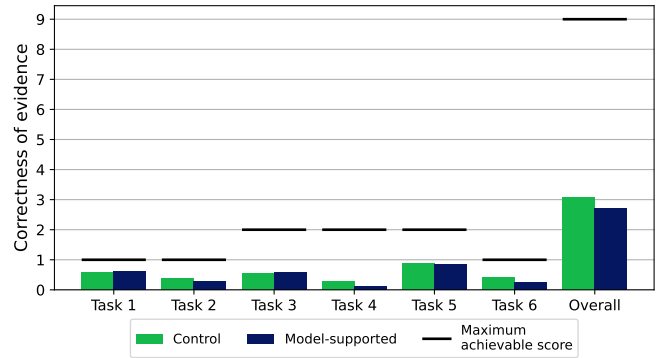


Fig. 5. Comparison across the two treatments of the participants' average score in correctness of evidence. Per task and overall.

**❗ RQ 1:** In the context of our experiment, providing a security-annotated DFD of the system to be analysed improved participants' analysis correctness in solving security analysis tasks. We observed a statistically significant ( $p = 0.0025$ ) improvement of 41% on average.

For some individual tasks, the difference in the average scores is only marginal (task 2: 0.42 vs. 0.46  $\hat{=}$  10% improvement in model-supported condition; task 6: 0.75 vs. 0.79  $\hat{=}$  5.6% improvement). In Section V, we discuss whether the nature of the tasks might be an indication of the extent to which a DFD improves the score in analysis correctness. However, even though the improvement is not statistically significant for all individual tasks (statistical significance only for task 3 with a p-value of 0.0003), Figure 4 clearly shows a trend of improved performance in the model-supported condition.

### B. Correctness of Evidence

Figure 5 presents the average score in correctness of evidence achieved by the participants. There are only small differences in the average scores for the model-supported and control condition. With an average of 3.08 out of a possible 9 in the control condition compared to an average of 2.71 in the model-supported condition (-12%), the participants performed better in the control condition, albeit without statistical significance ( $p = 0.52$ ). Task 4 has the lowest average correctness of evidence of the individual tasks (average of 0.29 out of 2 in the control condition and 0.13 in the model-supported condition),

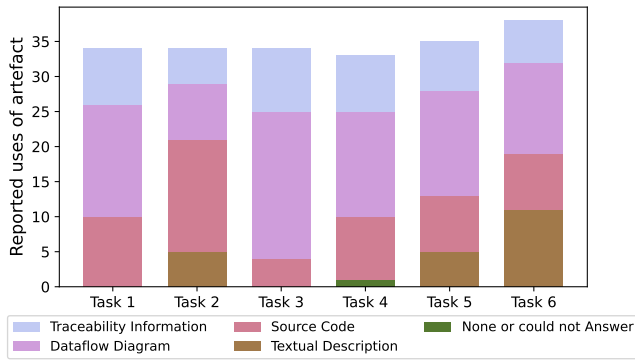


Fig. 6. Reported usage of provided artefacts per task (in the model-supported condition, where all artefacts were available).

while task 1 has the highest average (0.58 out of 1 in the control condition and 0.63 in the model-supported condition).

### C. Use of Provided Artefacts

After each task, the participants were asked to name all provided artefacts that they used to answer the task. Figure 6 shows the reported usage per task in the model-supported condition. We focus on this condition because, there, the participants had access to all artefacts. Note that the participants had the possibility to name multiple artefacts per task. Overall, the DFD was used the most, with a total of 88 answers. The source code was used 55 times, the traceability information 43 times, and the textual description 21 times. The numbers show, that the participants did not solely rely on the provided DFD, but instead also referred to the source code in many cases and the textual descriptions in some cases. This could indicate, that the participants verified information they found in the DFD by checking the corresponding part of the code.

### D. Influence of Use of Artefacts on Scores

We investigated whether the participants' usage of the provided artefacts had an influence on their performance. Although they were provided more artefacts in the model-supported condition, this does not necessarily mean that they used them all. The answers to the tasks could be found with more than one of the artefacts. Thus, the influence of single artefacts on the performance is not necessarily reflected in the comparison of outcomes between the two conditions. For example, participants in the model-supported condition could have not used the provided DFD to answer the tasks. Consequently, we compared the average scores in analysis correctness and in the correctness of evidence between two groups of participants for each artefact. To the group *Using Artefact* we assigned all participants that reported using the artefact in more than 50% of the tasks (4 or more). The group *Not Using Artefact* contains those participants who reported using it less (3 or fewer). We considered only the outcomes from the participants in the model-supported condition, since only here they had access to all artefacts. The grouping and analysis were done separately for each artefact, thus, the cardinality and members of the groups differ between artefacts.

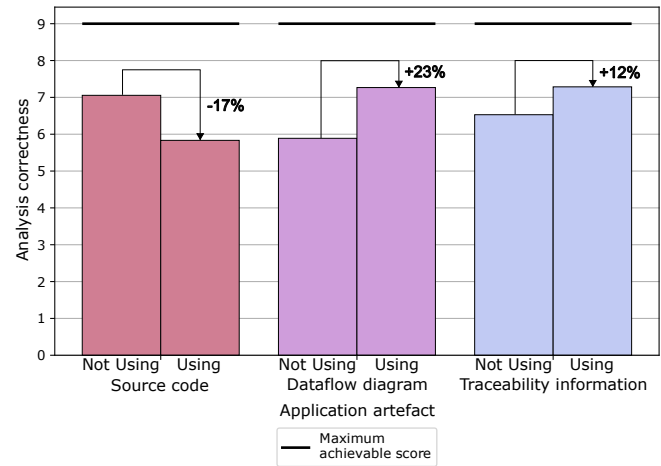


Fig. 7. Average scores in analysis correctness of those participants that reported using an artefact in more than 50% of the tasks (Using Artefact) and those that reported using it less (Not Using Artefact).




Fig. 8. Average scores in the correctness of evidence of those participants that reported using an artefact in more than 50% of the tasks (Using Artefact) and those that reported using it less (Not Using Artefact).

Figure 7 presents the average scores in analysis correctness of the two groups per application artefact. The figure shows that the Using Artefact group performed better compared to the Not Using artefact group for the artefacts DFD (+23% in score) and traceability information (+12%), while they performed worse for the source code (-17%). Figure 8 presents the results of this analysis for the correctness of evidence. For the artefact source code, the Using Artefact group achieved a 15% higher score in correctness of evidence than the Not Using Artefact group. For the use of DFD, they performed worse than the Not Using Artefact group (-20%). The highest difference, however, is seen in the traceability information. The Using Artefact group achieved a 315% higher average score in correctness of evidence than the Not Using Artefact group. We answer RQ2 based on these results since they distinguish between the use of the DFD and traceability information. In the results above, this distinction could not be made because the traceability information is an integral part of the DFDs and its isolated impact on the analysis could not be measured.

**❶ RQ 2:** Using traceability information significantly improved the correctness of evidence given for answers. On average, participants that used this artefact in more than half of the tasks achieved a 315% higher correctness of evidence compared to participants that used it less than that.

### E. Time

All participants were able to complete the tasks in the allotted 90 minutes. Their average time to complete all tasks was 34 minutes in the control condition and 35 minutes in the model-supported condition. No notable difference was observed. To examine a possible correlation between performance and time spent to finish the tasks, we also created a scatter plot visualizing their scores against their time. No correlation between scores and time could be visually identified.

### F. Perceived Usefulness and Usability of DFDs

The answers given to the open question at the end of the analysis sessions provide insights into the participants' perceived usefulness of the DFDs. The question asked about positive or negative observations during the experiment. For participants in the model-supported condition, it explicitly mentioned the usefulness of the DFDs and traceability information.

Out of 23 answers given by participants in the model-supported condition, two were negative, stating that thorough documentation would be preferred and that the DFD was “*a little bit hard to understand at first*”. Three answers listed both positive and negative experiences, where the negative points were two mentions that finding implementation details was hard (both participants reported using the traceability only in one task) and one that the participant lacked domain knowledge. A further 14 answers were predominantly positive.

*“Dataflow Diagrams were incredibly helpful, and all questions were answered almost completely from it.”*

Of the 23 answers, 9 mentioned specific beneficial scenarios for the use of DFDs. The ability to provide an overview of the system was mentioned 8 times, the benefit of referring to the important places in source code and use of the models as interface to the code was mentioned 3 times, and the reduction of the required domain knowledge was mentioned once.

Mild critique about the accessibility of the DFDs or traceability information was raised in 4 responses, for example:

*“[...] the transfer from the DFD to the traceability information could be made easier by clickable links in the DFD [...]”*

In summary, the statements made by participants in the model-supported condition include descriptions of the general usefulness of the DFDs, of benefits in finding implementation details via the model items and traceability information, and of their usefulness for architectural details and providing an overview. The positive feedback outweighed the few negative comments. Most participants reported the DFDs to be of help in the analysis and to be accessible to use.

Of the answers given after the control condition, only one was positive, stating that the textual description was helpful.

Four others referred to the DFDs (these answers were given in the second week, the participants had thus already performed the session with the DFD), stating that, in comparison, the lack of a DFD was an obstacle during the analysis. Specifically, they raised concerns about the correctness of their given answers and stated that finding the required features directly in source code was challenging.

*“The traceability file and the DFD were a big help last time, this time I wasn't really sure if I even answered correctly and didn't really know if the evidence I gave was correct. [...]”*

Six further answers of participants in the control condition in the first week (and thus without the comparison to the model-supported condition) reported negatively about their experience in the experiment. Specifically, they mentioned a lack of expertise, uncertainty about the given answers, and general difficulties in answering the tasks. Interestingly, two participants criticized the lack of a “*CFG*” or “*some kind of map of the architecture*”. This could have been sparked by the introductory lecture where DFDs were addressed but is still seen as an interesting comment. The obstacles reported by the participants in the control condition give further weight to the positive feedback of those in the model-supported condition.

Based on a qualitative analysis of the participants' statements, we can cautiously judge the perceived usefulness and accessibility of the DFDs to answer RQ3:

**❶ RQ 3:** In our experiment, the perceived usefulness and accessibility reported by the participants varied from very positive feedback to mild critiques reporting some confusion. Overall, the statements focussed on usefulness and were predominantly positive.

### G. Open Challenges of DFDs

The above observations of the quantitative and qualitative results allowed us to distill a number of open challenges of DFDs, i.e., current obstacles that would increase the DFDs' positive impact further if solved. Although these challenges were not explicitly investigated as independent variables in our experiment, they became evident from the results of the experiment, explicit answers given by participants, and observations made during the analysis of the tasks.

**Open Challenge 1: Understandability of Models** The participants in our experiment performed with statistical significance better in the model-supported condition and they reported a generally good accessibility. Nevertheless, concerns were raised about the understandability of the models. Some participants commented, that they did not understand the model initially or that they did not know what some annotations mean. A more usable model representation of software systems should consider the accessibility for human users, especially those with lower domain knowledge.

**Open Challenge 2: Presenting Missing Features** The DFDs in their current form do not support the explicit presentation of the absence of features or properties. In the context of security analysis, these could be security mechanisms that are

not implemented by a given application. To enable more comprehensive analysis and increase users' trust, it is important to show that such mechanisms were investigated and are not implemented in the analysed application. In this context, the challenges are to prove the absence, to decide what features to consider, and how to convey this information to the user. We see this open challenge as the hardest one to solve, both conceptually and practically.

### **Open Challenge 3: Accessibility of Traceability Information**

The quantitative results of our experiment show, that the traceability information has a positive impact on the correctness of evidence provided for answers to the tasks. While this is an expected observation, multiple participants also mentioned the usefulness of the traceability information for navigating the source code. However, it was also mentioned in some answers that the connection to the source code was difficult to follow. Also, the traceability information was not used by everyone even when it was provided. We conclude that the ease of use can be improved and that navigating the links to source code should be simplified. This challenge is of more practical nature and can likely be solved with some clever engineering.

**❗ RQ 4:** We identified three open challenges of DFDs (understandability of models, presenting missing features, and accessibility of traceability information). If any of these are solved, the positive impact of DFDs on security analysis can be expected to further increase.

## V. DISCUSSION

At the heart of the conducted experiment lay the question of the impact of providing DFDs and traceability information on the participants' performance. The results presented in Section IV indicate an overall positive impact on the analysis correctness. The scores improved with statistical significance in the model-supported condition. Figure 7 emphasizes this finding. Participants in the model-supported condition who reported using the provided DFD in more than half of the tasks had a 23% higher score on average than those who reported using it less. A 12% higher average score for participants using the traceability information is further proof of the usefulness of the DFDs, since the traceability information is one of their core features. The observed 17% lower score in analysis correctness for participants who reported using the source code in more than half of the tasks was an unexpected outcome at first sight. A closer look at the usage of the source code as artefact revealed, that out of 55 responses that mentioned using the source code, 34 (this corresponds to 62%) did not use the DFD or traceability information in conjunction. In other words, the source code was predominantly not used alongside the models, but instead as the only artefact to answer a task. Consequently, in our experiment, many participants who reported using the source code could also be described as not using the provided models. With this re-phrasing, the results are another indication of the models' positive impact.

Looking at the individual tasks, the increase in scores differed between them. This suggests the question of which

type of tasks the DFDs have the most impact on, and how exactly they impact different types of tasks. We investigated whether the nature of the tasks could be an explanation for the observations, i.e. whether the type of task can indicate how the score is influenced. We found that the DFDs impacted the analysis tasks in our experiment in different ways. They are described in the following. Please refer to Table I for the tasks.

*Providing an Overview:* Tasks 1, 2, and 6 have fairly simple answers in comparison to the other tasks. The answer for task 1 (in which the analysis correctness improved by 33% in the model-supported condition) could be found at two places in the code, either a deployment file indicating the container's port, or a configuration file indicating the service's port. Both answers were accepted as correct. In the DFD, the port is shown as annotation to the corresponding node. Interestingly, the wrong answers given are one of two options. One is the port number of a different microservice, which likely showed up when searching for "port" with GitHub's search function. The other is the port of a database that is only visible in the code as part of the database's URL. How this answer was reached by participants is puzzling. For task 6, the improvement of the average score was the lowest of all tasks (0.75 in the control condition and 0.79 in the model-supported condition; 5.6% increase). The task has the overall best average scores, likely, because the authorization service's name ("auth\_server") hints towards the answer of which of the services handles the authorization. Task 2 could be answered based on the textual description, on the Java annotation that implements the API gateway in code, or on an annotation in the DFD. A 10% improvement in average score in analysis correctness was observed from the control condition (0.42) to the model-supported condition (0.46). The answers lead us to believe that the question might have been formulated such that participants did not fully understand it. Many of the wrong answers in both conditions stated the used framework (Spring) instead of the library that was asked for (Zuul). Further, this task had the lowest reported number of usages of the DFD as well as traceability information (compare Figure 6).

The answers and evidences indicate, that DFDs are helpful in providing an overview and presenting the answers to simple questions such as the port number of a microservice. Evidently, finding *any* port in the code is a simple task in many systems' codebases, however, the answers suggest that finding the *correct* one can be challenging. Likely, this is heightened by the complexity that the microservice architecture adds to an application's codebase due to its decoupling. The answers given by participants in the model-supported condition further emphasize this quality of DFDs to provide an overview of the important system components (compare Section IV, where this was the benefit most often mentioned by participants). Simultaneously, for simple tasks with a fairly easy answer, good coding practice such as choosing descriptive identifiers seems to support the analysts well and there is no pressing need to provide a DFD. Whether this holds true in the analysis of larger applications should be investigated in future work. The results of task 2 indicate problems in the DFDs'



accessibility. The presented information seems to not be self-explanatory enough for the participants to answer this task reliably, even when the information is contained in the DFDs.

→ The results indicate, that DFDs serve as a means to “navigate the jungle” that is the application’s codebase. They provide an overview of the application’s architecture and (security and other) features. At the same time, well-chosen identifiers in code can support the solving of simple analysis tasks and the DFDs add less value in this scenario.

*Reducing Required Domain Knowledge:* To answer tasks 3 and 5 in the control condition, some domain knowledge was needed to correctly grasp the functionality of the relevant code. Task 3 required the participants to identify three outgoing connections (for App 1, two for App 2) of a microservice. One is a direct API call implemented with Spring Boot’s `RestTemplate`, another a registration with a service discovery service, and the third a registration with a tracing server (similar for App 2). Some domain knowledge about these technologies or Java was required to identify them. With the DFD at hand, answer the task came down to identifying the correct node in the diagram and noting the three nodes to which there was an information flow. To answer task 5 without the DFD, participants had to check whether three services (for App 1, two for App 2) refer to the authorization service in a configuration file under an `authorization` section. In the DFD, a connection to the authorization server indicated this. Again, knowledge about Spring or Java made it easier to find the correct answers without the support of the DFDs.

Task 3 showed the biggest impact of the models, with a doubled average score in analysis correctness (0.875 in control condition and 1.75 in model-supported condition; 100% increase). While this task was more difficult to answer than the others without a DFD and the required domain knowledge, the magnitude of the difference is still substantial. For task 5, the average score in analysis correctness was 1.29 in the control condition and 1.58 in the model-supported condition, a 23% increase. The differences show how the DFDs reduce the domain knowledge required for analysis activities. However, we hypothesize, that the participants without the DFD could answer the task simply by identifying the keyword “authorization” in the configuration files without checking if the implementation is correct and behaves in the way that is asked for. We believe, that this led them to achieve an average score without the DFDs that is still high. Given the scenario in which they solved the tasks (empirical experiment, where answers are expected), this was likely sufficient evidence for them to answer, independent of whether their domain knowledge was profound enough to fully understand the workings.

→ Our interpretation of the results is that DFDs are especially helpful in scenarios where a lack of domain knowledge about the analysed application’s framework, libraries, etc. hinders the identification of features and system components. The DFDs’ ability to shed light on properties shaded by a curtain of domain knowledge seems to be one of their core virtues.

*Indicating Absence of Features:* Despite the open challenge 2 (presenting missing features in the DFDs, see Section IV), the results also indicate that the DFDs in their current form already support users in answering tasks concerning the absence of features in the code. Task 4 was different from the other ones in that the challenge lay not in finding an artefact in the code but instead the absence of it. The task asked for the presence of encryption in two connections (for App 1, three for App 2) between services. The correct answer to all of them was “No”. The average score in analysis correctness was 0.83 in the control condition and 1.33 in the model-supported condition out of a possible score of 2 (60% increase). The difficulty in this task also became apparent when looking at the results for the evidence. The participants achieved an average score in correctness of evidence of 0.042 in both conditions.

→ Although the DFDs still face the open challenge of presenting missing features, their current form already supports users in answering tasks that require identifying the absence of features in code.

In summary of the discussion of the results, we see that the DFDs had a positive impact on the scores in different types of tasks. Specifically, they provide an overview of the analysed application, they reduce the required domain knowledge, and they can indicate the absence of features in the application. The highest increase in scores is seen for tasks where some domain knowledge was needed to answer them without the DFDs. The only task where the improvement of the analysis correctness in the model-supported condition was neglectable was a simple task where descriptive identifiers in code indicated the answer.

## VI. THREATS TO VALIDITY

*Internal validity:* With a large group of university students as participants, collaborations during or between the sessions and resulting cross-contamination cannot be ruled out completely. As mitigation, we strictly discouraged collaborations and conversations about the study and supervised the analysis sessions. Learning effects or the possibility of preparing for the tasks were mitigated with the employed within-groups design where the scenarios switched over the two sessions and with the use of different applications. With 90-minute long sessions, experimental fatigue is limited. The random assignment to the groups G1 and G2 limits selection bias. Some of the analysed data (timestamps, experience, resource usage) is self-reported, and we have to rely on its correctness. The encouragement of positive as well as negative feedback and the often-repeated reassurance of full anonymity of the answers were used to increase the reliability of the data. By making participation voluntary and using only the standard incentive for attending the lab sessions, it is possible that we have attracted mainly students who show high motivation and are at the top of their class. This could have had distorting effects on the results and could not be reasonably mitigated.

*External validity:* The conclusions drawn in this paper might not entirely map to other scenarios or populations. The tasks used as examples of security analysis activities might differ

from real-world use cases and thus influence the shown effects. Further, the experiment focused on microservice applications written in Java. We chose Java applications because it is the most used programming language for open-source microservice applications. The analysis of systems that follow a different architectural style or are written in another programming language could show other outcomes. The number of participants (24) is relatively small. We chose robust statistical methods that are suitable for the sample size and discussed the impact of the participants' experience and the choice of tasks. The participants' expertise in security analysis is rather low. Thus, the effects described in this paper might not be observed for other users, e.g., with more experience. However, the use of DFDs is not confined to security experts, hence rendering the participants a suited population for the experiment. Finally, experiments with practitioners instead of students could lead to different results, however, it is a common practice and has been shown to produce valid results as well (see Section III-D).

*Construct validity:* We measured the participants' performance in terms of correctness and time, which are common and objective metrics for such experiments. They relate to the practical use-case of the investigated effects directly. The analysis correctness is crucial in security analysis to ensure accurate security evaluations and, consequently, secure systems. The time serves as a measure of productivity and efficiency. Other constructs were disregarded but could be suited as well.

*Content validity:* The tasks concerned the key security mechanisms implemented in the analysed applications. These or similar tasks would be part of a real-world security analysis. However, other tasks might also be important in this context.

*Conclusion validity:* The responses to the tasks were given in free-text fields. Although we did not identify such ambiguities in quantifying the responses, it is possible that some answers were phrased in a way that was interpreted incorrectly. A more restrictive way of collecting the answers could have increased the conclusion's validity.

## VII. RELATED WORK

Although DFDs are used for different aspects of security analysis, no related work could be found that investigates their direct impact on the correctness of the analysis. Publications for other model types exist. For example, a considerable body of empirical research on Unified Modeling Language (UML) diagrams has been published [27]. A number of experiments have been conducted to investigate whether users' comprehension of the modelled systems increases with UML diagrams. Gravino et al. [28], [29] observed a positive impact of the models, while experiments by Scanniello et al. [30] did not show such an improvement (the authors attribute this to the type of UML diagrams, which had little connection to the code since they had been created in the initial requirements elicitation phase of the development process). In an experiment by Arisholm et al. [31], code changes performed by participants with access to UML documentations showed significantly improved functional correctness. Other researchers investigated the impact of specific properties of UML diagrams

on users' comprehension. For example, Cruz-Lemus et al. [32], [33], Ricca et al. [34], and Staron et al. [35], [36] found that stereotypes (which are similar to annotations in DFDs in our experiment) increased users' efficiency and effectiveness in code comprehension. Some publications found alternative model representations to yield better comprehension among participants in empirical experiments: Otero and Dolado [37] reported that OPEN Modelling Language (OML) models led to faster and easier comprehension than UML diagrams, while Reinhartz-Berger and Dori [38] reported Object-Process Methodology (OPM) models to be better suited than UML diagrams for modelling dynamic aspects of applications.

Bernsmed et al. [39] presented insights into the use of DFDs in agile teams by triangulating four studies on the adoption of DFDs. In the studies, software engineers were confused about the structure, granularity, and what to include in the models, because no formal specification of DFDs exists. The participants in our experiment also showed some difficulties that could be resolved by a clear definition and well-established documentation of DFDs. Regarding DFDs' structure, Faily et al. [40] argued that they should not be enriched with additional groups of model items, since their simplicity and accessibility for human users might suffer. Instead, they proposed to use them together with other system representations. In contrast, Sion et al. argued in a position paper [41] that using DFDs in their basic form is insufficient for threat modelling. Based on our findings, we argue that adding annotations to DFDs does not impede their accessibility and that security-enriched DFDs are well suited to support security analysis activities.

In conclusion, no publications were found that empirically investigate the impact of DFDs (or other model representations) on the security analysis (or related activities) of microservice applications.

## VIII. CONCLUSION

This paper presents the results of an empirical experiment conducted to investigate the impact of DFDs on software security analysis tasks. DFDs are widely used for security analysis and their varied adoption indicates a high confidence in their usefulness. To the best of our knowledge, the presented results are the first to investigate these assumptions and can confirm a positive impact of DFDs in the given context. We found, that participants performed significantly better concerning the analysis correctness of security analysis tasks when they were provided a DFD of the analysed application. Additionally, traceability information that links model items to artefacts in source code significantly improved their ability to provide correct evidence for their answers. Consequently, this paper serves as a basis for future research on specific applicabilities and properties of DFDs. Further, it can provide guidance in decisions on the adoption of model-based practices.

## ACKNOWLEDGEMENT

This work was partly funded by the European Union's Horizon 2020 programme under grant agreement No. 952647 (Assure-MOSS).

## REFERENCES

- [1] L. Sion, K. Yskout, D. Van Landuyt, W. Joosen, Solution-aware data flow diagrams for security threat modeling, in: Proceedings of the 33rd Annual ACM Symposium on Applied Computing, SAC '18, Association for Computing Machinery, New York, NY, USA, 2018, p. 1425–1432. doi:10.1145/3167132.3167285.
- [2] S. Hernan, S. Lambert, T. Ostwald, A. Shostack, Threat modeling-uncover security design flaws using the stride approach, MSDN Magazine (2006) 68–75.
- [3] Microsoft Corporation, Microsoft threat modeling tool 2016 (2016). URL <https://www.microsoft.com/en-us/download/details.aspx?id=49168>
- [4] P. Torr, Demystifying the threat modeling process, IEEE Security & Privacy 3 (5) (2005) 66–70. doi:10.1109/MSP.2005.119.
- [5] M. Abi-Antoun, D. Wang, P. Torr, Checking threat modeling data flow diagrams for implementation conformance and security, in: Proceedings of the 22nd IEEE/ACM International Conference on Automated Software Engineering, ASE '07, Association for Computing Machinery, New York, NY, USA, 2007, p. 393–396. doi:10.1145/1321631.1321692.
- [6] M. Abi-Antoun, J. M. Barnes, Analyzing security architectures, in: Proceedings of the IEEE/ACM International Conference on Automated Software Engineering, ASE '10, Association for Computing Machinery, New York, NY, USA, 2010, p. 3–12. doi:10.1145/1858996.1859001.
- [7] B. Berger, K. Sohr, R. Koschke, Automatically Extracting Threats from Extended Data Flow Diagrams, in: Engineering Secure Software and Systems, Vol. 9639, 2016, pp. 56–71. doi:10.1007/978-3-319-30806-7\_4.
- [8] C. Cao, S. Schneider, N. Diaz Ferreyra, S. Verweer, A. Panichella, R. Scandariato, CATMA: Conformance Analysis Tool For Microservice Applications, in: 2024 IEEE/ACM 46th International Conference on Software Engineering: Companion Proceedings (ICSE-Companion), 2024. doi:10.1145/3639478.3640022.
- [9] K. Tuma, R. Scandariato, M. Balliu, Flaws in Flows: Unveiling Design Flaws via Information Flow Analysis, in: 2019 IEEE International Conference on Software Architecture (ICSA), 2019, pp. 191–200. doi:10.1109/ICSA.2019.00028.
- [10] R. Chen, S. Li, Z. E. Li, From monolith to microservices: A dataflow-driven approach, in: 2017 24th Asia-Pacific Software Engineering Conference (APSEC), 2017, pp. 466–475. doi:10.1109/APSEC.2017.53.
- [11] T. D. Stojanovic, S. D. Lazarevic, M. Milic, I. Antovic, Identifying microservices using structured system analysis, in: 2020 24th International Conference on Information Technology (IT), 2020, pp. 1–4. doi:10.1109/IT48810.2020.9070652.
- [12] S. Li, H. Zhang, Z. Jia, Z. Li, C. Zhang, J. Li, Q. Gao, J. Ge, Z. Shan, A dataflow-driven approach to identifying microservices from monolithic applications, Journal of Systems and Software 157 (2019) 110380. doi:10.1016/j.jss.2019.07.008.
- [13] N. Dragoni, S. Giallorenzo, A. Lluch-Lafuente, M. Mazzara, F. Montesi, R. Mustafin, L. Safina, Microservices: yesterday, today, and tomorrow, Springer International Publishing, 2016, Ch. 12, pp. 195–216. doi:10.1007/978-3-319-67425-4\_12.
- [14] J. Lewis, M. Fowler, Microservices: a definition of this new architectural term, MartinFowler.com 25 (14-26) (2014) 12.
- [15] S. Schneider, R. Scandariato, Automatic extraction of security-rich dataflow diagrams for microservice applications written in java, Journal of Systems and Software 202 (2023) 111722. doi:10.1016/j.jss.2023.111722.
- [16] T. DeMarco, Structure Analysis and System Specification, Springer Berlin Heidelberg, 1979. doi:10.1007/978-3-642-48354-7\_9.
- [17] K. Tuma, R. Scandariato, M. Widman, C. Sandberg, Towards security threats that matter, in: S. K. Katsikas, F. Cuppens, N. Cuppens, C. Lambrinouidakis, C. Kalloniatis, J. Mylopoulos, A. Antón, S. Gritzalis (Eds.), Computer Security, Springer International Publishing, Cham, 2018, pp. 47–62. doi:10.1007/978-3-319-72817-9\_4.
- [18] S. Schneider, T. Özen, M. Chen, R. Scandariato, microsecend: A dataset of security-enriched dataflow diagrams for microservice applications, in: 2023 IEEE/ACM 20th International Conference on Mining Software Repositories (MSR), 2023, pp. 125–129. doi:10.1109/MSR59073.2023.00030.
- [19] B. Kitchenham, S. Pfleeger, L. Pickard, P. Jones, D. Hoaglin, K. El Emam, J. Rosenberg, Preliminary guidelines for empirical research in software engineering, IEEE Transactions on Software Engineering 28 (8) (2002) 721–734. doi:10.1109/TSE.2002.1027796.
- [20] N. Juristo, A. Moreno, Basics of Software Engineering Experimentation, 2001. doi:10.1007/978-1-4757-3304-4.
- [21] C. Wohlin, P. Runeson, M. Höst, M. Ohlsson, B. Regnell, A. Wesslén, Experimentation in Software Engineering, Springer, Germany, 2012. doi:10.1007/978-3-642-29044-2.
- [22] S. Schneider, N. E. Diaz Ferreyra, P.-J. Queval, G. Simhandl, U. Zdun, R. Scandariato, Replication package, 2024. URL [https://github.com/tuhh-softsec/SANER2024\\_empirical\\_experiment\\_DFDs](https://github.com/tuhh-softsec/SANER2024_empirical_experiment_DFDs)
- [23] A. Bambhore Tukaram, S. Schneider, N. E. Diaz Ferreyra, G. Simhandl, U. Zdun, R. Scandariato, Towards a security benchmark for the architectural design of microservice applications, in: Proceedings of the 17th International Conference on Availability, Reliability and Security, ARES '22, Association for Computing Machinery, New York, NY, USA, 2022, pp. 1–7. doi:10.1145/3538969.3543807.
- [24] I. Salman, A. T. Misirli, N. Juristo, Are students representatives of professionals in software engineering experiments?, in: Proceedings of the 37th International Conference on Software Engineering - Volume 1, ICSE '15, IEEE Press, 2015, p. 666–676.
- [25] M. Svahnberg, A. Aurum, C. Wohlin, Using students as subjects - an empirical evaluation, in: Proceedings of the Second ACM-IEEE International Symposium on Empirical Software Engineering and Measurement, ESEM '08, Association for Computing Machinery, New York, NY, USA, 2008, p. 288–290. doi:10.1145/1414004.1414055.
- [26] D. Falesi, N. Juristo, C. Wohlin, B. Turhan, J. Münch, A. Jedlitschka, M. Oivo, Empirical software engineering experts on the use of students and professionals in experiments, Empirical Softw. Engg. 23 (1) (2018) 452–489. doi:10.1007/s10664-017-9523-3.
- [27] D. Budgen, A. J. Burn, O. P. Brereton, B. A. Kitchenham, R. Pretorius, Empirical evidence about the uml: a systematic literature review, Software: Practice and Experience 41 (4) (2011) 363–392. doi:https://doi.org/10.1002/spe.1009.
- [28] C. Gravino, G. Scanniello, G. Tortora, Source-code comprehension tasks supported by uml design models: Results from a controlled experiment and a differentiated replication, Journal of Visual Languages & Computing 28 (2015) 23–38. doi:https://doi.org/10.1016/j.jvlc.2014.12.004.
- [29] C. Gravino, G. Tortora, G. Scanniello, An empirical investigation on the relation between analysis models and source code comprehension, in: Proceedings of the 2010 ACM Symposium on Applied Computing, SAC '10, Association for Computing Machinery, New York, NY, USA, 2010, p. 2365–2366. doi:10.1145/1774088.1774576.
- [30] G. Scanniello, C. Gravino, M. Risi, G. Tortora, G. Doderò, Documenting design-pattern instances: A family of experiments on source-code comprehensibility, ACM Trans. Softw. Eng. Methodol. 24 (3) (may 2015). doi:10.1145/2699696.
- [31] E. Arisholm, L. Briand, S. Hove, Y. Labiche, The impact of uml documentation on software maintenance: an experimental evaluation, IEEE Transactions on Software Engineering 32 (6) (2006) 365–381. doi:10.1109/TSE.2006.59.
- [32] J. A. Cruz-Lemus, M. Genero, D. Caivano, S. Abrahão, E. Insfarán, J. A. Carsí, Assessing the influence of stereotypes on the comprehension of uml sequence diagrams: A family of experiments, Information and Software Technology 53 (12) (2011) 1391–1403. doi:10.1016/j.infsof.2011.07.002.
- [33] M. Genero, J. A. Cruz-Lemus, D. Caivano, S. Abrahão, E. Insfarán, J. A. Carsí, Assessing the influence of stereotypes on the comprehension of uml sequence diagrams: A controlled experiment, in: K. Czarnecki, I. Ober, J.-M. Bruel, A. Uhl, M. Völter (Eds.), Model Driven Engineering Languages and Systems, Springer Berlin Heidelberg, Berlin, Heidelberg, 2008, pp. 280–294.
- [34] F. Ricca, M. Di Penta, M. Torchiano, P. Tonella, M. Ceccato, How developers' experience and ability influence web application comprehension tasks supported by uml stereotypes: A series of four experiments, IEEE Transactions on Software Engineering 36 (1) (2010) 96–118. doi:10.1109/TSE.2009.69.
- [35] M. Staron, L. Kuzniarz, C. Wohlin, Empirical assessment of using stereotypes to improve comprehension of uml models: A set of experiments, Journal of Systems and Software 79 (5) (2006) 727–742, quality Software. doi:10.1016/j.jss.2005.09.014.

- [36] M. Staron, L. Kuzniarz, C. Thurn, An empirical assessment of using stereotypes to improve reading techniques in software inspections, *SIGSOFT Softw. Eng. Notes* 30 (4) (2005) 1–7. doi:10.1145/1082983.1083308.
- [37] M. C. Otero, J. J. Dolado, An empirical comparison of the dynamic modeling in oml and uml, *Journal of Systems and Software* 77 (2) (2005) 91–102. doi:10.1016/j.jss.2004.11.022.
- [38] I. Reinhartz-berger, D. Dori, Opm vs. uml—experimenting with comprehension and construction of web application models, *Empirical Software Engineering* 10 (2005) 57–80. doi:10.1023/B:EMSE.0000048323.40484.e0.
- [39] K. Bernsmed, D. Cruzes, M. Jaatun, M. Iovan, Adopting threat modelling in agile software development projects, *Journal of Systems and Software* 183 (2021) 111090. doi:10.1016/j.jss.2021.111090.
- [40] S. Faily, R. Scandariato, A. Shostack, L. Sion, D. Ki-Aries, Contextualisation of data flow diagrams for security analysis, in: H. Eades III, O. Gadyatskaya (Eds.), *Graphical Models for Security*, Springer International Publishing, Cham, 2020, pp. 186–197.
- [41] L. Sion, K. Yskout, D. Van Landuyt, A. van den Berghe, W. Joosen, Security threat modeling: Are data flow diagrams enough?, in: *Proceedings of the IEEE/ACM 42nd International Conference on Software Engineering Workshops, ICSEW'20*, Association for Computing Machinery, New York, NY, USA, 2020, p. 254–257. doi:10.1145/3387940.3392221.