

TAINT ANALYSIS FOR GRAPH APIS FOCUSING ON BROKEN ACCESS CONTROL

LEEN LAMBERS ^a, LUCAS SAKIZLOGLOU ^a, TAI SIYA KHAKHAROVA ^a,
AND FERNANDO OREJAS ^b

^a Brandenburg University of Technology, Cottbus, Germany
e-mail address: <name>.<surname>@b-tu.de

^b Universitat Politècnica de Catalunya, Barcelona, Spain
e-mail address: orejas@lsi.upc.edu

ABSTRACT. A Graph API models the data managed by a web application based on the graph data model, i.e., data objects are represented by nodes and relationships by edges. This rather novel type of APIs presents new challenges when it comes to properly securing the APIs against the usual web application security risks, e.g., broken access control. A prominent security testing approach is taint analysis, which traces tainted, i.e., security-relevant, data from sources (where tainted data is inserted) to sinks (where the use of tainted data may lead to a security risk), over the information flow in an application.

We present the first systematic approach to *static* and *dynamic* taint analysis for Graph APIs focusing on broken access control. The approach comprises the following. We taint nodes of the Graph API if they represent data requiring specific privileges in order to be retrieved or manipulated, and identify API calls which are related to sources and sinks. Then, we statically analyze whether a tainted information flow between API source and sink calls occurs. To this end, we model the API calls using graph transformation rules. We subsequently use Critical Pair Analysis to automatically analyze potential dependencies between rules representing source calls and rules representing sink calls. We distinguish direct from indirect tainted information flow and argue under which conditions the Critical Pair Analysis is able to detect not only direct, but also indirect tainted flow. The static taint analysis (i) identifies flows that need to be further reviewed, since tainted nodes may be created by an API call and used or manipulated by another API call later without having the necessary privileges, and (ii) can be used to systematically design dynamic security tests for broken access control. The dynamic taint analysis checks if potential broken access control risks detected during the static taint analysis really occur. We apply the approach to a part of the GitHub GraphQL API. The application illustrates that our analysis supports the detection of two types of broken access control systematically: the case where users of the API may not be able to access or manipulate information, although they should be able to do so; and the case where users (or attackers) of the API may be able to access/manipulate information that they should not.

1. INTRODUCTION

A *Graph API* models the data managed by a web application based on the graph data model, i.e., data objects are represented by nodes and relationships by edges. While initially they emerged to facilitate data management in social networks, Graph APIs are currently used by an increasing number of software applications in the area of service-oriented and

cloud computing [ZZX⁺23, QFGR23]. A prominent specification and query language for Graph APIs is *GraphQL* [GFa]. GraphQL was introduced by Facebook in 2012 as part of the development of mobile applications, and is now used in various web and mobile applications, e.g., GitHub, Pinterest, Paypal, etc.

Security [Gra07] is particularly important in the area of web and mobile applications. Graph APIs are still relatively new, which means that a wide range of methods and tools for security analyses are not yet available for them [QFGR23, PMY23]. *Taint analysis* is a typical representative of such security analyses. It is a particular instance of data flow analysis [PY07] and systematically tracks certain user inputs through the application. *Broken access control* is considered as one of the most prominent security vulnerabilities for web applications [Tea]. Access control enforces policies such that users cannot act outside their intended permissions. Failures typically lead to unauthorized information disclosure, modification, or destruction of all data or performing a business function outside the user’s limits. Typical access control failures occur when an API can be accessed with (i) missing or misunderstood access control, or (ii) wrongly implemented access control. Such failures occur in practice, as demonstrated by issues in the GitHub community forum, e.g., issues 110618, 106598, 85661 [Gitb].

In this paper, we present a *taint analysis for Graph APIs* focusing on *broken access control* (BAC). Specifically, our approach builds on the concept of taint analysis, which propagates taint labels through the implementation of a program, to trace the flow of sensitive data through the Graph API. Graph transformation is a technique dedicated to formalizing graph manipulations and is therefore a natural foundation for formalizing a Graph API. Specifically, we formalize both a Graph API and the taint analysis using concepts from graph transformation [Roz97, EEPT06, HT20] and rely on related, formally founded tool support to perform part of the analysis. A *security analyst* can use the taint analysis to systematically and soundly check for BAC vulnerabilities. To the best of our knowledge, no other approach to checking for BAC vulnerabilities simultaneously supports key feature of our approach: use of taint analysis, native graph formalization, and sound detection—see discussion in Section 5. We illustrate the taint analysis based on GraphQL, but our approach is applicable to any kind of Graph APIs that can be formalized using graph transformation.

The presented taint analysis comprises three segments: the setup, the static analysis, and the dynamic analysis. During *setup*, the security analyst can taint nodes of the Graph API that represent sensitive data. This Graph API is represented formally using graph transformation rules typed over a type graph. The security analyst thus taints nodes in the type graph, representing node types entailing sensitive data that may be processed by the Graph API. From these tainted nodes, we can characterize graph transformation rules as Graph API *sources*, where tainted data are inserted, i.e., created by the rule, and *sinks*, where tainted data are used, i.e., changed by the rule, and may thus lead to a security risk. The *static analysis* supports the security analyst in *validating* the access control policies, i.e. finding missing or misunderstood access control policies. This is achieved by automatically generating pairs of source and sinks that may be risky in terms of leaking sensitive data and thus represent potential BAC vulnerabilities. The security analyst can then systematically check if the policy description indicates that it properly secures these risky flows. Subsequently, the *dynamic analysis* supports the security analyst in *verifying* that the access control policies are implemented correctly. If the security analyst expects a risky flow based on the results of the static analysis, then they can systematically try to uncover a corresponding BAC vulnerability in the Graph API implementation. We demonstrate that

the static analysis is *sound* w.r.t. finding potential *direct* BAC vulnerabilities, i.e, when a sink directly follows a source, as well as a type of *indirect* ones, i.e, when other data manipulations occur between source and sink. However, since the static analysis is incomplete, we integrate it with a subsequent *complete* dynamic analysis.

Graph transformation is a formal paradigm with many applications. A *graph transformation system* is a collection of graph transformation rules that, in union, serve a common purpose. In this paper, a *graph transformation system* is used to formally describe a Graph API. For many applications (see [LST⁺18] for a survey involving 25 papers), it is beneficial to know all conflicts and dependencies that can occur for a given pair of rules. A conflict is a situation in which one rule application renders another rule application inapplicable. A dependency is a situation in which one rule application needs to be performed such that another rule application becomes possible. We will use *dependency analysis* to discover possible tainted information flows that may lead to broken access control vulnerabilities.

This article is an extension of a paper published in ICGT '24 [LSAWK24]. The article introduces the following extensions. First, it extends the formalization to support a certain type of indirect vulnerabilities (see Theorem 3.13), whereas the paper supported only direct vulnerabilities. Moreover, we have added a number of insightful examples illustrating the limitation of the paper formalization as well as the new formal results. As a second extension, the article introduces a new type of coverage, namely *role coverage* (see Definition 3.18), which addresses whether derived tests cover the various roles in an access control policy. Finally, the article introduces a new application of the presented taint analysis—see Section 4.4. This application reproduces the issue 106598 reported in the GitHub community forum, addresses a larger part of the GraphQL schema language, and increases the level of real-world detail covered by the examples. Additionally, we have made the following improvements and changes: we elaborated on preliminaries regarding the dependency analysis—see Section 2.3; and we restructured the section describing the applications of the taint analysis to GitHub, i.e, Section 4, to render it more informative; we have added a summary of the GitHub access control policy, an outline of the application setting, as well as a discussion of the applications.

The rest of this paper is structured as follows: Section 2 revisits preliminaries. Section 3 presents taint analysis, i.e, the setup, the static, and the dynamic analysis, in further detail, based on a basic running example. Section 4 applies the analysis to a part of the GitHub GraphQL API [Gitc]. Section 5 is dedicated to related work while Section 6 to conclusion.

2. PRELIMINARIES

First, we review the basics of Graph APIs and GraphQL, a prominent Graph API discussed further in the paper; in the same section, we introduce the running example used in the remainder. Then we present the basics of access control. Finally, we recall the double-pushout (DPO) approach to graph transformation as presented in [EEPT06]. We reconsider dependency notions on the transformation and rule level [LEO08, LBO⁺18, LST⁺18] for the DPO approach, including dependency reasons.

2.1. Graph APIs and Running Example. A Graph API models the data managed by a web application based on the graph data model, i.e., data objects are represented by nodes and relationships by edges. Graph APIs emerged to facilitate data management in social networks which similarly relied on graph-structured data. Two familiar examples of Graph APIs are the *Facebook API* [Met] and the *GitHub GraphQL API* [Gitc] (henceforth referred

```

type User {
  username: String!
  projects: [Project]!
  repos: [Repository]!
}

type Repository {
  id: ID!
  name: String!
  description: String!
  url: String!
  owner: Repository!
}

type Issue {
  id: ID!
  body: String!
  repo: Repository!
}

type Project {
  name: String!
  id: ID!
  body: String
}

type Query {
  getProject(projectId: ID!): Project
}

type Mutation {
  createRepo(newName: String!, newDescription: String!, repoURL: String!): ID
  updateRepo(repoId: ID!, newName: String!, newDescription: String!): Boolean
  createProject(projectId: ID!, body:String!, newName:String!): ID
  deleteProject(projectId: ID!): Boolean
  createIssue(repoId: ID!, newBody: String!): ID
  updateIssue(repoId: ID!, issueId, ID!, newBody: String!): Boolean
  deleteIssue(repoId: ID!, issueId: ID!): Boolean
}

```

FIGURE 1. The GraphQL schema for the running example.

to as GitHub API). The GitHub API is based on the *GraphQL* specification and query language that describes the capabilities and requirements of graph data models for web applications [GFa]. In the remainder, we focus on APIs which are based on GraphQL.

GraphQL uses a *schema*, i.e., a strongly typed contract between the client and the server, to define all exposed types of an API as well as their structure. Schemata of GraphQL APIs are either publicly available or can be discovered using distinguished *introspection* queries. An example of a schema is shown in Figure 1. The schema captures the GraphQL API of a simplified version of a platform that allows developers to manage their code as well as their development efforts, similar to GitHub. The schema is defined in the GraphQL schema language [GFb] and contains the following object types: the **User**, which represents a user of the platform; the **Repository**, which represents a code repository hosted on the platform; the **Issue**, which represents a task related to the development work being performed in a repository; and the **Project**, which represents a task-board that can be used for planning and organizing development work. Attributes and relationships are captured by fields within types, where squared brackets refer to a list of objects. The exclamation mark renders a field non-nullable.

GraphQL defines two types of operations that can be performed on the server: *queries* are used to fetch data, whereas *mutations* are used to manipulate, i.e, create, update, or delete data. Queries and mutations are also defined in the schema, within the distinguished

types **Query** and **Mutation**, respectively. Figure 1 shows a minimal list of mutations for the platform in our example: the mutations *createRepo*, *updateRepo* create and update an instance of a **Repository**, respectively.

2.2. Access Control Policies. In various types of software systems, e.g., financial and safety-critical, the adequate security of the information and the systems themselves constitutes a fundamental requirement. *Access control* refers to ensuring that users are allowed to perform activities which may access information or resources only when they have the required rights. An access control *policy* is a collection of high-level requirements that manage user access and determine who may access information and resources under which circumstances [HKY17]. Correspondingly, a policy is *broken* when these requirements are not fulfilled.

While there are various methods of realizing policies, we focus on the *role-based* method. In a role-based policy, a security analyst defines roles, i.e., security groups, whose access rights are defined according to roles held by users and groups in organizational functions; the analyst assigns a role to each user who, once authorized, can perform the activities permitted by said role. In the remainder, we assume that, besides the security analyst, users cannot perform administrative actions, i.e, cannot change their own role. A role-based access control policy is broken when a user performs an activity that is not permitted by their role.

A requirement of a role-based policy for our running example could state that *a user may only update the information of a repository, if that user is the owner of the repository*. In practice, the Graph API of the platform should ensure that the query *updateRepo* can only be executed by a user who is authorized to update the repository, i.e, who has the role *owner*. An update of the repository information by a *collaborator*, i.e, a user who does not belong to this group, would break the policy.

2.3. Graph Transformation and Static Dependency Analysis. Throughout this paper we always consider graphs (and graph morphisms) typed over some fixed type graph TG via a typing morphism $type_G : G \rightarrow TG$ as presented in [EEPT06]. A graph morphism $m : G \rightarrow H$ between two graphs consists of two mappings m_V and m_E between their nodes and edges being both type-preserving and structure-preserving w.r.t. source and target functions. Note that in the main text we denote inclusions by \hookrightarrow and all other morphisms by \rightarrow .

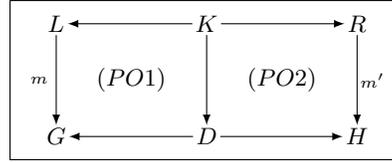
Graph transformation is the rule-based modification of graphs. A *rule* consists of two graphs: L is the left-hand side (LHS) of the rule representing a pattern that has to be found to apply the rule. After the rule application, a pattern equal to R , the right-hand side (RHS), has been created. The intersection K is the graph part that is not changed; it is equal to $L \cap R$ provided that the result is a graph again. The graph part that is to be deleted is defined by $L \setminus (L \cap R)$, while $R \setminus (L \cap R)$ defines the graph part to be created. We consider a graph transformation system just as a set of rules.

A *direct graph transformation* $G \xrightarrow{m,r} H$ between two graphs G and H is defined by first finding a graph morphism m of the LHS L of rule r into G such that m is injective, and second by constructing H in two passes: (1) build $D := G \setminus m(L \setminus K)$, i.e., erase all graph elements that are to be deleted; (2) construct $H := D \cup m'(R \setminus K)$. The morphism m' has to be chosen such that a new copy of all graph elements that are to be created is added. It has been shown that r is applicable at m iff m fulfills the *dangling condition*. It is satisfied if all adjacent graph edges of a graph node to be deleted are deleted as well, such that D becomes

a graph. Injective matches are usually sufficient in applications and w.r.t. our work here, they allow explaining constructions with more ease than for general matches. In categorical terms, a direct transformation step is defined using a so-called double pushout as in the following definition. Thereby step (1) in the previous informal explanation is represented by the first pushout (PO1) and step (2) by the second pushout (PO2) [EEPT06].

Definition 2.1 ((read/write) rule and transformation). A *rule* r is defined by $r = (L \xleftarrow{le} K \xrightarrow{ri} R)$ with L, K , and R being graphs connected by two graph inclusions and $K = L \cap R$. A rule r is a *read rule* if le and ri are identity morphisms. A rule r is a *write rule* if it is not a read rule.

A direct transformation $G \xrightarrow{m, m', r} H$ applying r to G consists of two pushouts as depicted on the right. Rule r is applicable and the injective morphism $m : L \rightarrow G$ ($m' : R \rightarrow H$) is called a *match* (resp. *co-match*) if there exists a graph D such that (PO1) is a pushout.



We may notice that a read rule is not really a transformation rule, in the sense that if we apply it to a graph G , the result is again G , i.e. it causes no transformation on it. We use read rules to model operations that just read data on a graph, without transforming it. In particular, when applying the rule $r = (L \xleftarrow{le} L \xrightarrow{ri} L)$ to G via a morphism m , we consider that some given user is reading the data included in $m(L)$.

Given a pair of transformations, a *produce-use dependency* [EEPT06] occurs if the co-match of the first transformation via rule r_1 produces a match for the second transformation that was not available yet before applying the first transformation. More precisely, the *creation graph* C_1 includes all the elements (nodes and edges) that are produced by r_1 . In particular, C_1 is the smallest subgraph of R_1 that includes $R_1 \setminus K_1$. Notice that, in general, $R_1 \setminus K_1$ is not a graph, because it may include edges whose source or target is in K_1 , i.e., C_1 would include, in addition to the edges in $R_1 \setminus K_1$, their source and target. These additional nodes are called *boundary nodes* and summarized in the *boundary graph* B_1 , i.e., $B_1 = C_1 \cap K_1$. Then, the *dependency reason* $C_1 \xrightarrow{q_1} S_1 \xrightarrow{q_2} L_2$ comprises elements produced by the first and used by the second rule giving rise to the dependency. The formal construction of this dependency reason is depicted in Figure 2, where technically square (1) is an initial pushout and square (2) is the pullback of $m'_1 \circ c_1$ and m_2 , i.e. S_1 contains all elements from C_1 and L_2 that are mapped to the same element in H_1 . In [LBK⁺19] it is explained in detail how to obtain, using *static dependency analysis*, such dependency reasons for a pair of rules for which transformation pairs in produce-use dependency exist. Each transformation pair in produce-use dependency comes with a unique dependency reason (completeness).

Definition 2.2 (produce-use dependency, dependency graph). Given a pair of direct transformations $(t_1, t_2) = (G \xrightarrow{m_1, m'_1, r_1} H_1, H_1 \xrightarrow{m_2, m'_2, r_2} H_2)$ applying rules $r_1 : L_1 \xleftarrow{le_1} K_1 \xrightarrow{ri_1} R_1$ and $r_2 : L_2 \xleftarrow{le_2} K_2 \xrightarrow{ri_2} R_2$ as depicted in Fig. 2. Square (1) in Figure 2 can be constructed as the initial pushout over morphism ri_1 . It yields the *boundary graph* B_1 and the *creation graph* C_1 . The *transformation pair* (t_1, t_2) is in *produce-use dependency* if there does not exist a morphism $x_{21} : L_2 \rightarrow D_1$ such that $h_1 \circ x_{21} = m_2$. The *dependency graph* $DG(\mathcal{R})$ for a set of rules \mathcal{R} consists of a node for each rule in \mathcal{R} and an edge (r_1, r_2) from r_1 to r_2 if there exists a transformation pair $(t_1, t_2) = (G \xrightarrow{r_1} H_1, H_1 \xrightarrow{r_2} H_2)$ in produce-use dependency.

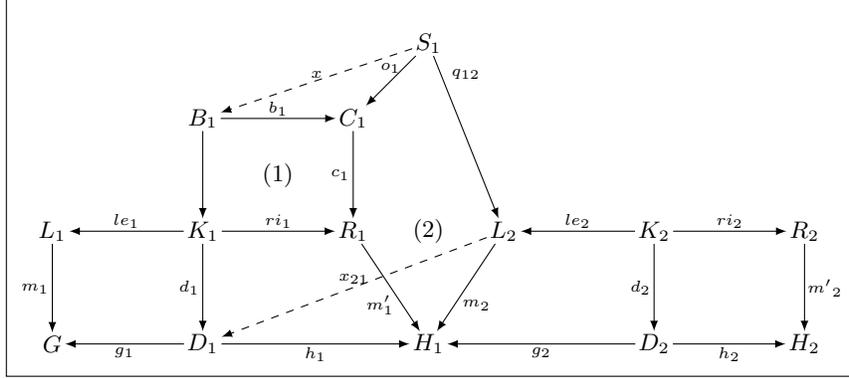


FIGURE 2. Illustration of dependency and dependency reason.

For a given graph transformation system, a *static dependency analysis* technique is a means to compute a list of all pairwise produce-use dependencies. Inspired by the related concept from term rewriting, *critical pair analysis* (CPA, [Plu94]) has been the established static conflict and dependency analysis technique for three decades. CPA reports each dependency as a critical pair, that is, a minimal example graph together with a pair of rule applications from which a dependency arises. The *multi-granular static dependency analysis* [BLST17, LBO⁺18]) supports the computation of dependencies for graph transformation systems on three granularity levels: on binary granularity, it reports if a given rule pair contains a dependency at all; on coarse granularity, it reports *minimal dependency reasons*; fine granularity is, roughly speaking, the level of granularity provided by critical pairs. Coarse-grained results have been shown to be more usable than fine-grained ones in a diverse set of scenarios and can be used to compute the fine-grained results much faster [LST⁺18]. We use the coarse granularity also in the following.

3. TAINT ANALYSIS FOR GRAPH APIS

This section presents our main contribution, the taint analysis for Graph APIs focusing on broken access control. The analysis assumes the following input: a *Graph API schema* (or a Graph API endpoint from which the schema can be introspected—see Section 2.1), comprising object and edge types as well as API calls, i.e., queries and mutations, and a *role-based access control policy* description (henceforth referred to as policy), consisting of high-level access control requirements like the one mentioned in Section 2.2. The analysis comprises three segments: the setup, the static analysis, and the dynamic analysis—see Figure 3. The *setup* segment comprises the derivation of a *tainted Graph API*, i.e., a Graph API schema with *tainted nodes* as well as *source and sink graph transformation rules* derived from the policy and the Graph API schema; we describe the setup in detail in Section 3.1. Once the setup is complete, the *static analysis* can be performed. The output of the static analysis is a set of *rule dependencies* which indicate potential occurrences of a broken policy. We present the formal concepts of the static analysis in Section 3.2. The completion of the static analysis enables the execution of the *dynamic analysis*. The dynamic analysis relies on the output of the static analysis and itself generates a *test case report* which either verifies implementations of policy rules or exposes rules where the policy is broken. We present the formal concepts of the dynamic analysis in Section 3.3.

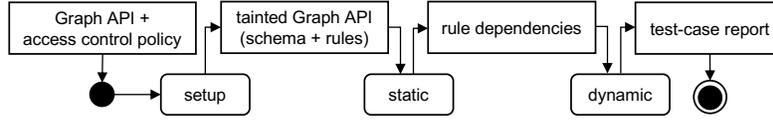


FIGURE 3. The segments of the presented approach, illustrated as an activity diagram.

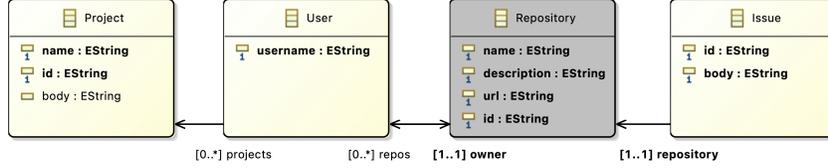


FIGURE 4. The tainted type graph, where `Repository` is a tainted node (in gray).

In the following, illustrations of type graphs and graph transformation rules are mostly based on the well-known *Henshin* project [ABJ⁺10], which is also used for tool support by the static analysis.

3.1. Setup. As mentioned in Section 1, the analysis relies on a formal description of a Graph API based on a graph transformation system, and focuses on broken access control. Therefore, our approach to taint analysis relies on the security analyst (i) obtaining both a representation of the GraphQL schema as a type graph TG and a representation of the API, i.e, all possible API calls, as a set of graph transformation rules \mathcal{R} typed over TG ; (ii) having access to a role-based policy specification for the Graph API; (iii) identifying nodes in the Graph API which represent sensitive data to be checked for broken access control by the taint analysis, as well as the rules in \mathcal{R} which create and manipulate those data. We formally define these concepts below.

3.1.1. Graph APIs. We formally represent a Graph API by a set of graph transformation rules \mathcal{R} typed over TG . We represent each API call as a direct graph transformation. A sequence of API calls, denoted here as API execution, is a graph transformation of arbitrary length. It starts with some initial graph G_0 , representing the data present in the API prior to the execution.

Definition 3.1 (graph API, call, execution). A *Graph API* with a schema TG is a set of graph transformation rules $\mathcal{R} = \{r_i | i \in I\}$ typed over TG . A *Graph API execution* consists of a (possibly infinite) graph transformation sequence $(t : G_0 \xrightarrow{r_1} G_1 \xrightarrow{r_2} G_2 \xrightarrow{r_3} G_3 \dots)$, starting with an initial graph G_0 , such that each r_i is an element of \mathcal{R} . A *Graph API call* is a direct transformation within a Graph API execution. We denote with $Sem(\mathcal{R})$ the set of all Graph API executions, and with $Sem_{call}(\mathcal{R})$ the set of all Graph API calls.

An illustration of a type graph for the platform of the running example is shown in Figure 4. The type graph corresponds to the entity types in the schema defined in the GraphQL schema language in Figure 1. Based on the mutations defined within the distinguished type `Mutation`

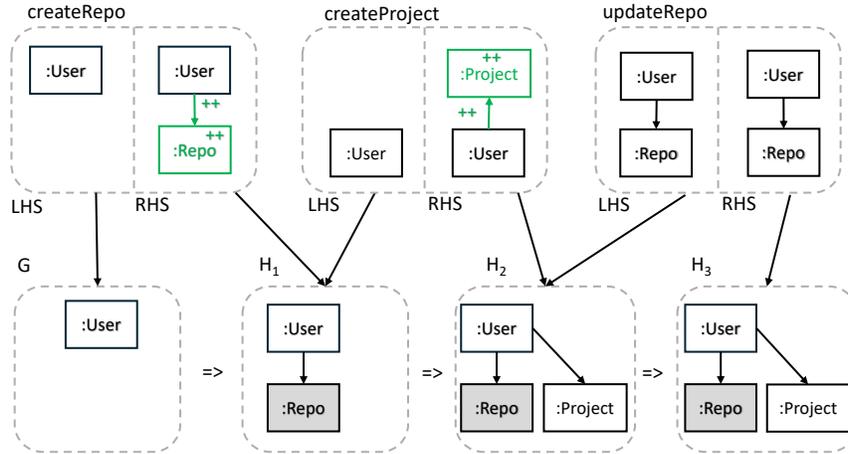


FIGURE 5. Graph API execution consisting of a sequence of three Graph API calls using the plain notation; the calls *createRepo*, *createProject*, *updateRepo* are depicted on top, whereas the graphs G , H_1 , H_2 , H_3 at the bottom depict the results following the execution of each call

in the GraphQL schema, we also illustrate a Graph API according to Definition 3.1, i.e, a set of transformation rules typed over the type graph, in Figure 6a. We show an example of a Graph API execution comprising three calls, i.e, rules, in Figure 5. In the remainder, we illustrate rules using two different notations: the *plain* notation (seen in Figures 5, 8, 9, and 10) illustrates the LHS and RHS as separate graphs and omits certain information, e.g., attributes; the *compact* notation (seen in Figures 6 and 13 and generated by Henshin) embeds the LHS and RHS of a rule in one graph and does not omit any information. Note that figures using the plain notation, show rules on top, while the execution, i.e, graphs before and after the application of a rule, are shown at the bottom. In the Graph API execution in Figure 5, a user creates a repository that is updated after the user creates a project and applies the rules *createRepo*, *createProject* and *updateRepo*.

In this section, we will use a minimal API consisting of the rules *createRepo* and *updateRepo* (see Figure 6a for a detailed illustration) to explain the static and dynamic analysis. We explain the analysis for the overall API in Section 4.

3.1.2. Access control. In a Graph API, access control, i.e, the fact that certain API calls that may create or manipulate sensitive data can only be executed by users with sufficient privileges, can be implemented via a role-based policy. In that case, users of the Graph API can only execute certain calls if they belong to a group with the required privileges. For instance, according to the GitHub access control policy, an API user with the role (or, in GitHub terms, *permission level*) *owner* has full control over the repository; an API user with the role *collaborator* has only read and write access to the code and cannot, for example, delete the repository [Githd].

We formalize role-based access control policy and API calls whose access control is role-based below. The policy formalization is deliberately kept high-level, since we want to support cases where only rather informal policy descriptions are given. We merely assume that a security analyst is able to decide based on their expectations and on the policy

description for each API call executed by some user with a specific role if said call is allowed or not. We call the policy formalization *policy oracle*. For a call executed with a role, the policy oracle evaluates to true (false) if access (no access) should be granted. Below, we also present the formalization of the policy implementation. This formalization adheres to the same signature of the policy oracle and is similarly kept high-level, since we want to support APIs regardless of the specifics of their access control implementation. The policy implementation decides whether access is granted to the actual implementation of the call and is later used to formalize the notion of broken access control vulnerability.

Definition 3.2 (role-based Graph API call, policy oracle, policy implementation). Given a Graph API \mathcal{R} with schema TG . A *role specification* (RO, \leq_{RO}) consists of a finite set of roles RO and a partial order $\leq_{RO} \subseteq RO \times RO$.

A *role-based Graph API execution* $(t : G_0 \xrightarrow{r_1} G_1 \xrightarrow{r_2} G_2 \xrightarrow{r_3} G_3 \dots, role_t)$ is a Graph API execution such that each direct transformation in t is executed with a specific role as given in $role_t : \mathbb{N} \rightarrow \mathcal{P}(RO)$, where $role_t(i)$ refers to the subset of roles of the i -th direct transformation $G_i \xrightarrow{r_{i+1}} G_{i+1}$. A *role-based Graph API call* is a direct transformation within a role-based Graph API execution. We denote with $Sem^r(\mathcal{R})$ the set of all *role-based Graph API executions*, and with $Sem^r_{call}(\mathcal{R})$ the set of all *role-based Graph API calls*.

A *role-based access control policy oracle* $ACP = ((RO, \leq_{RO}), P)$ consists of a role specification (RO, \leq_{RO}) and a *policy oracle* P , i.e, a total mapping $P : Sem^r_{call}(\mathcal{R}) \rightarrow \{true, false\}$.

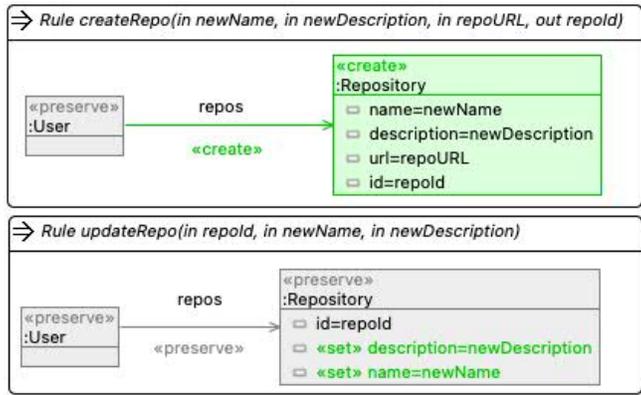
A *role-based access control policy imlementation* $ACP^I = ((RO, \leq_{RO}), P^I)$ consists of a role specification (RO, \leq_{RO}) and a *policy implementation* P^I , i.e, a total mapping $P^I : Sem^r_{call}(\mathcal{R}) \rightarrow \{true, false\}$ describing for a specific implementation I of the Graph API calls if access is actually granted or not, respectively.

Note that each call in an execution t may be executed with a different role, which mirrors the typical interaction with a Graph API in practice—see Section 4.1.

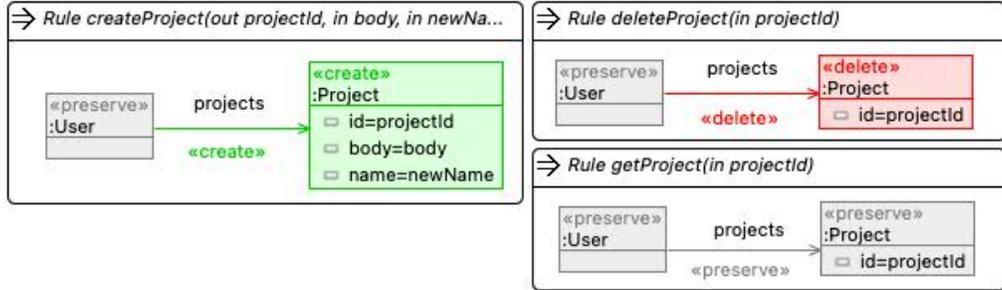
As discussed in Section 2.2, a broken access control vulnerability occurs if policy requirements about user access to information and resources are not fulfilled. In our running example, such a vulnerability occurs if users without the proper privileges get access or are able to manipulate specific elements in the repositories, e.g., issues or projects, or repositories themselves; this would result to a data leak or even the loss of valuable information in the code repository. Formally, a *broken access control vulnerability* is captured by a role-based Graph API execution containing a call that breaks the access control policy. This either happens if the call is executed with the specified role in a concrete policy implementation P^I , although this should not be possible according to the policy oracle P , or also if the call is not allowed with the specified role in a concrete policy implementation P^I , although this should be possible according to the policy oracle P .

Definition 3.3 (Broken Access Control (BAC) vulnerability). Given a graph API \mathcal{R} with schema TG , a role-based access control policy oracle $ACP = ((RO, \leq_{RO}), P)$, and a policy implementation $P^I : Sem^r_{call}(\mathcal{R}) \rightarrow \{true, false\}$ associating the actual policy result to each API call. A *Broken Access Control (BAC) vulnerability* is represented by a role-based Graph API execution $(t : G_0 \xrightarrow{r_1} G_1 \xrightarrow{r_2} G_2 \xrightarrow{r_3} G_3 \dots, roles_t)$ entailing an API call $c : G_i \xrightarrow{r_{i+1}} G_{i+1}$ such that $P(c) \neq P^I(c)$.

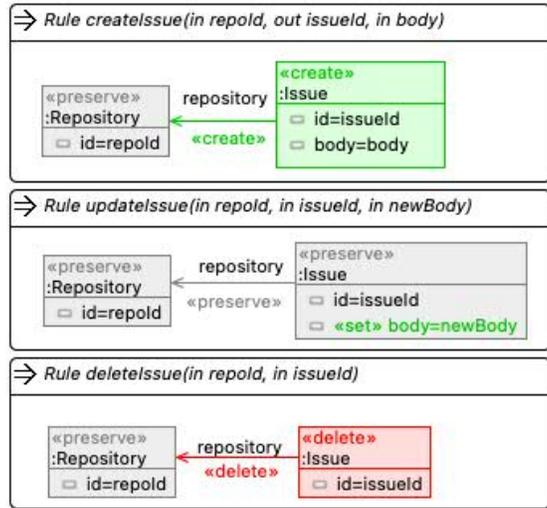
See Figure 5 again for an example of a Graph API execution. Assume that the first and third call are executed by a *user* as an owner of the repository and the second call by a



(A) Rules related to Repository.



(B) Rules related to Project.



(C) Rules related to Issue.

FIGURE 6. The sources (*create* rules) and sinks (*update*, *get*, *delete* rules) used in the running example and Section 4.

collaborator. For such a role-based Graph API execution, the security analyst may expect that access is granted for each of the calls. So the policy oracle would map to true for each of the calls in the role-based Graph API execution. If the actual implementation of the Graph API execution raises a BAC exception and does not allow one of the calls (while being executed with the specified roles), we have a BAC vulnerability. In this case, the policy implementation is too conservative w.r.t. the expectations of the security analyst for this role-based Graph API execution to access control. Imagine now that a repository can only be updated by a user of the repository with the role *owner* or *collaborator*, but not by any other arbitrary user. Then, if the actual implementation of the Graph API execution does not raise a BAC exception when executing this third call with a role other than *owner* or *collaborator*, this again represents a BAC vulnerability. In this case, the access control policy implementation is too liberal w.r.t. the expectations of the security analyst for this role-based Graph API execution, and we have discovered a data leak.

We say that an access control policy oracle is *stable under shift* if for each pair of sequentially independent calls, the policy oracle mapping for each of the calls remains identical when swapping the calls. Recall that a pair of graph transformation steps $G \xrightarrow{r} G' \xrightarrow{r'} G''$ is *sequentially independent* if their order can be reversed, i.e., $G \xrightarrow{r'} H \xrightarrow{r} G''$ without affecting the end result G'' .

Definition 3.4 (policy stable under shift). A role-based access control policy oracle $ACP = ((RO, \leq_{RO}), P)$ is *stable under shift* if for each pair of sequentially independent calls $G \xrightarrow{r} G' \xrightarrow{r'} G''$ executed with roles $ro, ro' \in \mathcal{P}(RO)$ respectively, it holds that $P(G \xrightarrow{r} G') = P(H \xrightarrow{r} G'')$ and $P(G' \xrightarrow{r'} G'') = P(G \xrightarrow{r'} H)$, the latter executed with roles $ro, ro' \in \mathcal{P}(RO)$ respectively. Analogously, a policy implementation P^I is stable under shift, if this property also holds for the policy implementation.

As policy oracles and implementations discussed in this article are assumed to disallow administrative actions (see Section 2.2), they are stable under shift, i.e., swapping the order of two calls does not change the policy oracle mapping for said calls. Definition 3.4 frames this characteristic in terms of graph transformation and is further used in a relevant proof later in the paper—see Theorem 3.13. A policy that would allow administrative actions, i.e., a call could assign the *collaborator* role to a user, could possibly not satisfy this notion of stability, e.g., whether such an administrative call precedes or succeeds another call could change the policy mapping for said calls.

3.1.3. *Tainting*. The idea of the taint analysis is to systematically trace the flow of sensitive data through the Graph API with the aim of detecting potential BAC vulnerabilities. The aim of the static taint analysis is to detect potential Graph API interactions that may not be properly secured. The aim of the dynamic taint analysis is to systematically further test these risky interactions to detect if BAC vulnerabilities arise indeed.

To this extent, the security analysis first needs to taint those nodes in the schema representing sensitive data objects. A BAC vulnerability is related to such a tainted node if a Graph API call introduces an instance node of a tainted node, and a later call reads or manipulates this node without the necessary privileges; a direct BAC vulnerability occurs when the introducing call is directly followed by the reading or manipulating call. Based on the schema knowledge about tainted nodes, we also derive a tainted Graph API. The

tainted API distinguishes source rules (creating sensitive data) from sink rules (reading or manipulating sensitive data). Tainted nodes and tainted graph APIs form the basis for the automatic static analysis. It will detect potential interactions between pairs of source and sink rules representing risky flows w.r.t. passing sensitive data. In this way, the static analysis will allow us to find all potential direct BAC vulnerabilities related to tainted nodes. Moreover, we will argue under which conditions we can also find specific indirect BAC vulnerabilities. We formalize these tainting concepts for our approach below.

Definition 3.5 (tainted type graph, tainted nodes and their relation to a (direct) BAC vulnerability). Let $TG = (V_{TG}, E_{TG}, s_{TG}, t_{TG})$ be a type graph. Then, $TG_t = (TG, T)$ is a *tainted type graph* with $T \subseteq V_{TG}$ the set of *tainted nodes*.

Given a Graph API \mathcal{R} with schema TG , a *role-based access control policy oracle* $ACP = ((RO, \leq_{RO}), P)$ and policy implementation $ACP^I = ((RO, \leq_{RO}), P^I)$. The role-based Graph API execution $(t : G_0 \xrightarrow{r_1} G_1 \xrightarrow{r_2} G_2 \xrightarrow{r_3} G_3 \dots, roles_t)$ represents a BAC vulnerability related to a tainted node s from T if it entails a call $c : G_i \xrightarrow{(m_{i+1}, r_{i+1})} G_{i+1}$ such that a node n with $type_V(n) = s$ in the image of m_{i+1} exists, leading to $P(c) \neq P^I(c)$, and n does not belong already to G_0 . A BAC vulnerability $(t, roles_t)$ is called *direct* if the node n with $type_V(n) = s$ is created in the call $G_{i-1} \xrightarrow{(m_i, r_i)} G_i$ in t directly before $c : G_i \xrightarrow{(m_{i+1}, r_{i+1})} G_{i+1}$.

See Figure 4 for an example of a tainted schema, where we taint the node **Repository**—tainting is illustrated by the gray background color of the node in question. In this case, this node is tainted because we want to further analyze if users without the proper privileges are able to manipulate created repositories. Assume that a BAC vulnerability arises in the implementation of the exemplary role-based Graph API execution in Figure 5, since it is possible to execute the third call successfully as a user with a role that does not have the required privileges. Then we say that this is a BAC vulnerability related to the tainted repository node R . This vulnerability is *not a direct vulnerability*, as the **Repository** node has been created in the first call and not in the previous, i.e., second, call.

Finally, to obtain a tainted graph API, we define *source rules*, i.e., rules that create instances of a tainted node in a schema, and *sink rules*, i.e., rules that read or manipulate instances of tainted nodes.

Definition 3.6 (source rule, sink rule, tainted graph API). Given a tainted type graph $TG_t = (TG, T)$ and a Graph API \mathcal{R} with schema TG . A rule $r = (L \xleftarrow{le} K \xrightarrow{ri} R)$ from \mathcal{R} is a *source rule* if there exists a node n in $R \setminus L$ such that $type_{R,V}(n) \in T$. A rule $r = (L \xleftarrow{le} K \xrightarrow{ri} R)$ from \mathcal{R} is a *sink rule* if there exists a node n in L such that $type_{L,V}(n) \in T$. A *tainted Graph API* $(TG_t, \mathcal{R}, src, sink)$ consists of a tainted type graph TG_t , a Graph API typed over TG , and a corresponding set of source and sink rules, src and $sink$.

As a minimal example of a tainted Graph API comprising a set of source and sink rules, see the rules *createRepo* and *updateRepo* in Figure 6a. The rule *createRepo* creates a tainted node, i.e., an instance of **Repository**, and is therefore a source rule, whereas *updateRepo* refers to instances of **Repository** and is therefore a sink rule. As we will explain in the next section, our static analysis identifies such pairs of source and sink rules such that the security analyst can review the policy description to check if the access to sensitive (i.e. tainted) nodes is properly secured.

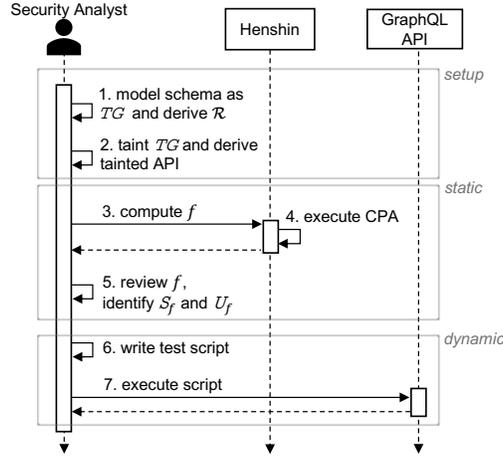


FIGURE 7. An overview of the sequence of steps in the presented approach.

3.2. Static Analysis. Figure 7 shows an overview of the sequence of steps in the presented approach to taint analysis. The setup segment comprised steps 1 and 2. The next segment is the static analysis, which is based on the output of the setup. The static analysis comprises the following steps: the provision of the schema and tainted Graph API obtained during setup by the security analyst (step 3 in Figure 7); the execution of the actual static analysis over the input provided by the security analyst, which is performed based on available tool support (step 4); and the review of output of the analysis by the security analyst (step 5) so that potential BAC vulnerabilities are identified due to an incomplete or incorrect access control policy. We describe this segment in detail below.

Owing to the graph-transformation-based formalization, the static analysis can automatically determine potential dependencies between the rules for a tainted Graph API by using CPA—see Section 2.3. In particular, we use the coarse granularity level of the *multi-granular static dependency analysis* to determine all minimal *dependency reasons*. The *tainted information flow* for a tainted Graph API is then a symbolic yet complete description of all possibilities to use a tainted node directly after it has been created. In the running example, a dependency occurs when a `Repository` is created, which is afterward updated. The unique dependency reason between rules `createRepo` and `updateRepo` in Figure 6a consists of the overlap of the complete right-hand side of the rule `createRepo` with the left-hand side of the rule `updateRepo`.

Definition 3.7 (tainted information flow). Let $(TG_t, \mathcal{R}, src, sink)$ be a tainted Graph API with $TG_t = (TG, T)$. We denote with $src(s)$ ($sink(s)$) the subset of rules in src ($sink$) creating (using) a node of type s from T , respectively. Then, the *tainted information flow* of the API is given by $f = \cup\{DR(r_{src}, r_{sink}) \mid r_{src} \in src(s), r_{sink} \in sink(s), s \in T\}$ with $DR(r_{src}, r_{sink})$ the set of dependency reasons for the rule pair (r_{src}, r_{sink}) .

Figure 11 shows the tainted information flow for the tainted Graph API from the running example—see dependency between `createRepo` and `updateRepo`.

The following theorem states that each BAC vulnerability related to a tainted node (see Definition 3.5) is represented by a sequence of API calls containing a source rule application creating the tainted node followed at some point by a sink rule which uses the tainted node

in a manner that violates the policy. We also say that such a BAC vulnerability entails a tainted flow instance.

Theorem 3.8 (BAC vulnerability and tainted flow instance). *Let $(t : G_0 \xrightarrow{r_1} G_1 \xrightarrow{r_2} G_2 \xrightarrow{r_3} G_3 \dots, roles_t)$ be a BAC vulnerability related to a tainted node s . Then, the BAC vulnerability entails a source rule application via some rule r_{src} in $src(s)$ creating a node n with $type_V(n) = s$ and a sink rule application via a sink rule r_{sink} in $sink(s)$ using n . We call this subsequence of t a tainted flow instance.*

Proof. There exists an API call $c : G_i \xrightarrow{(m_{i+1}, r_{i+1})} G_{i+1}$ in t such that $P(c) \neq P^I(c)$ for a node n in the image of m_{i+1} with $type_V(n) = s$ (Definition 3.5). We thus have that r_{i+1} equals a sink rule r_{sink} in $sink(s)$, since the preimage of n in the left-hand side of the rule r_{i+1} then also has type s (Definition 3.6). Since the node n is not already contained in G_0 (Definition 3.5), there must be a call $G_j \xrightarrow{(m_{j+1}, r_{j+1})} G_{j+1}$ with $0 \leq j < i$ creating n in t . This means that r_{j+1} equals a source rule r_{src} in $src(s)$, since it creates node n of tainted type s (Definition 3.6). Thus we can rewrite t into $G_0 \dots \xrightarrow{(m_{j+1}, r_{j+1})} G_{j+1} \xrightarrow{r_{j+2}} \dots G_i \xrightarrow{(m_{i+1}, r_{i+1})} G_{i+1} \dots$ with $r_{j+1} = r_{src}$ and $r_{i+1} = r_{sink}$. \square

Based on the definition of a tainted flow instance, we can now technically define a potential BAC vulnerability: such a vulnerability is an API execution that contains a tainted flow instance, although it has not yet been verified whether the vulnerability is covered by the policy oracle implementation, i.e, whether for each call c in the API execution $P(c) = P^I(c)$.

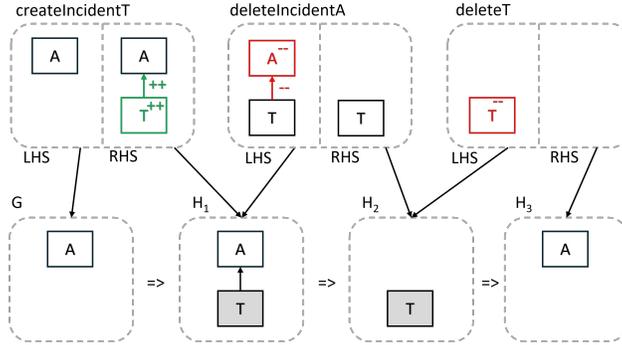
Based on Theorem 3.8, we can distinguish BAC vulnerabilities with atomic and non-atomic, i.e, more complex, tainted flow instances. The former occurs when the sink rule application violating the policy oracle is caused by a source rule application without the occurrence of any intermediate applications of some sink or source rule.

Definition 3.9 (atomic tainted flow instance). Given a BAC vulnerability $(t : G_0 \xrightarrow{r_1} G_1 \xrightarrow{r_2} G_2 \xrightarrow{r_3} G_3 \dots, roles_t)$ related to a tainted node s such that $t : G_0 \dots \xrightarrow{(m_{j+1}, r_{j+1})} G_{j+1} \xrightarrow{r_{j+2}} \dots G_i \xrightarrow{(m_{i+1}, r_{i+1})} G_{i+1} \dots$ with $G_i \xrightarrow{(m_{i+1}, r_{i+1})} G_{i+1}$ violating the policy oracle and $r_{j+1} = r_{src}$ and $r_{i+1} = r_{sink}$. If all intermediate rule applications between $G_j \xrightarrow{(m_{j+1}, r_{j+1})} G_{j+1}$ and $G_i \xrightarrow{(m_{i+1}, r_{i+1})} G_{i+1}$ in t are of rules other than r_{src} and r_{sink} , then we have a BAC vulnerability containing an *atomic tainted flow instance*.

A security vulnerability detection technique is *sound* for a category of vulnerabilities if it can correctly conclude that a program has no vulnerabilities of that category. In the following theorem, we argue that a static analysis technique computing the tainted information flow is sound, i.e, it detects all *direct* potential BAC vulnerabilities related to tainted nodes.

Theorem 3.10 (soundness for direct vulnerabilities). *Given a tainted type graph $TG_t = (TG, T)$, a tainted Graph API $(TG_t, \mathcal{R}, src, sink)$, and a tainted information flow $f = \cup\{DR(r_{src}, r_{sink}) \mid r_{src} \in src(s), r_{sink} \in sink(s), s \in T\}$. Given a direct BAC vulnerability $(t : G_0 \xrightarrow{r_1} G_1 \xrightarrow{r_2} G_2 \xrightarrow{r_3} G_3 \dots, roles_t)$ related to a tainted node s from T , then $DR(r_{src}, r_{sink})$ in f is non-empty for some r_{src} in $src(s)$ and r_{sink} in $sink(s)$.*

Proof. Based on Theorem 3.8 and given that t is a direct vulnerability (Definition 3.5) we can rewrite t into $G_0 \dots \xrightarrow{(m_i, r_i)} G_i \xrightarrow{(m_{i+1}, r_{i+1})} G_{i+1} \dots$ with $r_i = r_{src} \in src(s)$ and



(A) Potential indirect vulnerability with forbidden incident structure deleted first.

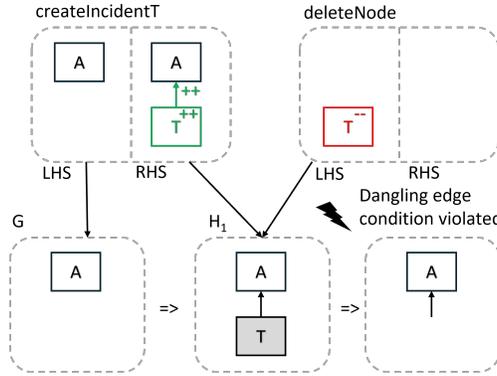
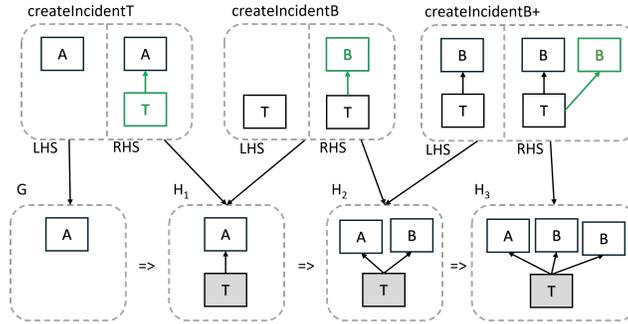
(B) No potential direct vulnerability because of dangling edge to node A .

FIGURE 8. Illustrations for Example 3.11.

$r_{i+1} = r_{sink} \in sink(s)$. As dependency reasons are complete [LBK⁺19], there exists a dependency reason d in $DR(r_{src}, r_{sink})$ such that a node m of type s (instantiated to node n in t) is created by r_{src} and used by r_{sink} as described in d . \square

Our detection technique is not sound for *indirect vulnerabilities*, i.e., situations where a tainted information flow occurs because of applications of intermediate rules between a source rule and a sink rule. Unsoundness may occur for two reasons—of which we give examples in Example 3.11 and Example 3.12. First, the source rule may create a tainted node together with some incident graph structure; this incident graph structure needs to be deleted by an intermediate rule in order for the sink rule that deletes the tainted node to become applicable. The second reason is that an intermediate rule may create an incident graph structure to the tainted node thereby enabling the application of a sink rule requiring this incident graph structure.

Example 3.11 (potential indirect vulnerability due to incident structure deleted). The toy example in Figure 8a illustrates a potential indirect vulnerability, since the rule $deleteT$ deletes the tainted node T created by the rule $createIncidentT$. The rule $deleteT$ is enabled by the application of the intermediate rule $deleteIncidentA$. As the sink rule $deleteT$ is not directly applicable after the source $createIncidentT$ (as this would create a dangling edge incident to node A —see Figure 8b), there is no critical pair, i.e., dependency reason, between the two rules.



(A) Potential indirect vulnerability with required incident structure created first.

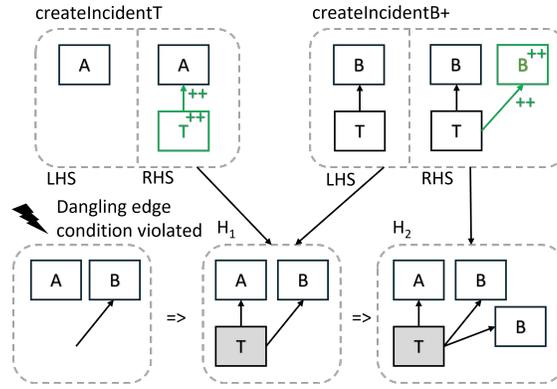
(B) No potential direct vulnerability because of dangling edge to node B .

FIGURE 9. Illustrations for Example 3.12.

Example 3.12 (Potential indirect vulnerability due to incident structure created). The toy example in Figure 9a illustrates a potential indirect vulnerability, since the rule $createIncidentB+$ uses the tainted node T created by rule $createIncidentT$. The application of $createIncidentB+$ is enabled by the intermediate application of $createIncidentB$. As the sink rule $createIncidentB+$ is not directly applicable after the application of the source rule $createIncidentT$ (as this would create a dangling edge incident to node A —see Figure 9b), there is no critical pair, i.e., dependency reason, between the two rules.

The examples above illustrate that critical pairs (or dependency reasons) are effective at illustrating direct dependencies only. As soon as dependencies arise because of the application of more than one rule these compact dependency representations are not complete anymore.

The following theorem proves that under certain conditions our analysis is still sound w.r.t. finding *both direct and indirect vulnerabilities with an atomic flow instance*. This is achieved by *reducing* indirect vulnerabilities to direct ones, provided that these two conditions are met: we can shift all intermediate rule applications between a source and sink rule application to some application before the source or after the sink rule application; and, moreover, the access control policy oracle is stable under shift—see Definition 3.4. A static check of whether shifting of rules can be achieved by checking the sequential independence of pairs of rules—see Definition 3.4; if a rule pair is sequentially independent, then the application of both rules would have the same result regardless of the order; thus, the

transformations can be switched. Sequential independence can be computed by the binary level of the *multi-granular static dependency analysis* [BLST17, LBO⁺18].

Theorem 3.13 (soundness for vulnerabilities). *Given a tainted Graph API $(TG_t, \mathcal{R}, src, sink)$ with $TG_t = (TG, T)$ and tainted information flow $f = \cup\{DR(r_{src}, r_{sink}) \mid r_{src} \in src(s), r_{sink} \in sink(s), s \in T\}$ such that the following holds for each (r_{src}, r_{sink}) : (1) for each rule r from $\mathcal{R} \setminus \{r_{sink}\}$ it holds that the rule pair (r_{src}, r) is sequentially independent or (2) for each rule r from $\mathcal{R} \setminus \{r_{src}\}$ it holds that the rule pair (r, r_{sink}) is sequentially independent.*

Given a policy oracle $ACP = (RO, \leq_{RO}, P)$ and policy implementation P^I which are stable under shift and a BAC vulnerability with atomic flow instance represented by $(t : G_0 \xrightarrow{r_1} G_1 \xrightarrow{r_2} G_2 \xrightarrow{r_3} G_3 \dots, roles_t)$ related to a tainted node s from T , then $DR(r_{src}, r_{sink})$ in f is non-empty for some r_{src} in $src(s)$ and r_{sink} in $sink(s)$.

Proof. Assume the BAC vulnerability $(t : G_0 \xrightarrow{r_1} G_1 \xrightarrow{r_2} G_2 \xrightarrow{r_3} G_3 \dots, roles_t)$ related to a tainted node s from T . If the BAC vulnerability is direct, then, based on Theorem 3.10, it follows that $DR(r_{src}, r_{sink})$ in f is non-empty for some r_{src} in $src(s)$ and r_{sink} in $sink(s)$.

Let the BAC vulnerability be indirect. Then, according to Theorem 3.8, we can rewrite t into $G_0 \dots \xrightarrow{(m_{j+1}, r_{j+1})} G_{j+1} \xrightarrow{r_{j+2}} \dots G_i \xrightarrow{(m_{i+1}, r_{i+1})} G_{i+1} \dots$ with $r_{j+1} = r_{src} \in src(s)$ and $r_{i+1} = r_{sink} \in sink(s)$. Now assume that condition (1) holds, then we know that for each rule r from $\mathcal{R} \setminus \{r_{j+1}\}$ it holds that (r_{j+1}, r) is sequentially independent. Because we have a vulnerability with an atomic flow instance, we know that all intermediate rule applications are applications of some rule r from $\mathcal{R} \setminus \{r_{j+1}, r_{i+1}\}$. We can thus obtain from t a transformation $t' : G_0 \dots G \xrightarrow{r_{j+1}} G' \xrightarrow{r_{i+1}} G'' \dots$ with $roles_{t'}$. This is because we can stepwise switch toward the front each intermediate rule application with the application of the source rule r_{j+1} . Since the policy is stable under shift, from Theorem 3.10 it follows now for $(t', roles_{t'})$ that $DR(r_{src}, r_{sink})$ in f is non-empty for some r_{src} in $src(s)$ and r_{sink} in $sink(s)$.

If condition (2) holds, we argue similarly, but instead switch each intermediate rule application of a rule other than the source rule toward the back. \square

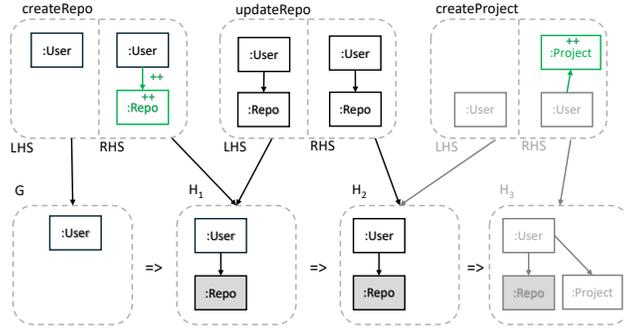
As illustrated by Example 3.11 and Example 3.12, the conditions from Theorem 3.13 are not met in all cases. Moreover, indirect vulnerabilities can also occur because of the intermediate rules being themselves source and sink rules. The development of a new, sound static detection technique for all indirect dependencies constitute part of our future plans; such a technique requires an extension of the theory of critical pairs.

The following example shows that the conditions from Theorem 3.13 applies to a part of our running example, illustrating how indirect vulnerabilities can be reduced to direct ones.

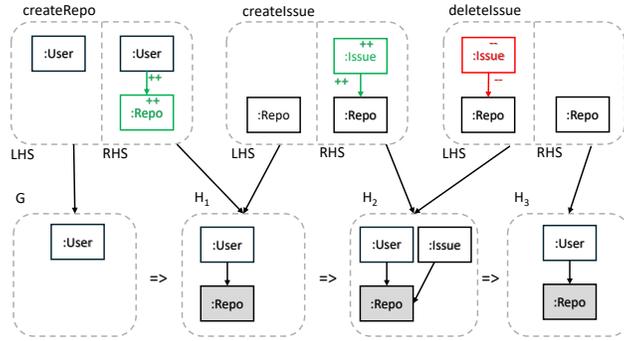
Example 3.14 (potential indirect vulnerability reducible to potential direct vulnerability). Figure 5 illustrates a potential indirect vulnerability: a user creates a repository that is updated after the user creates a project. This potential indirect vulnerability can be reduced to a direct one, since it is possible to switch the intermediate application of *createProject* with the application of *updateRepo*—see Figure 10a.

Figure 10b illustrates another potential indirect vulnerability: a user creates a repository for which an issue is created and afterward this issue is deleted again. We cannot switch the intermediate rule application *createIssue* forward or backward in this sequence. This example illustrates the need for a new, sound static detection technique for all indirect dependencies.

Effectively, each dependency reason in the tainted information flow constitutes a potential vulnerability that calls for further investigation. But the tainted information flow may still



(A) Potential indirect vulnerability reduced to potential direct vulnerability.



(B) Potential indirect vulnerability.

FIGURE 10. Illustrations for Example 3.14.

contain *false positives*, i.e. issues that do not turn out to be actual vulnerabilities. Therefore, a manual review of the tainted information flow by the security analyst (step 5 in Figure 7) as described in the following is necessary.

The analyst reviews each dependency reason between source and sink rules contained in the tainted information flow against the access control policy of the API. The review has one of the following two outcomes. First, if the policy does not properly cover the tainted flow expressed by the dependency reason, it represents an *unsecured tainted flow*; the policy thus needs to be adapted such that proper access control for this flow can be implemented; in this case, the dynamic analysis (see Section 3.3) can be used to prove that the dependency reason leads to a true positive, i.e., a BAC vulnerability, by demonstrating the occurrence of unsecured access to sensitive data. Alternatively, the dependency reason represents a *secured tainted flow*; in this case, the dynamic analysis can be used to verify the correct implementation of the policy.

We formalize these two alternative outcomes below.

Definition 3.15 (secured and unsecured tainted flow). Given a tainted type graph $TG_t = (TG, T)$, a tainted Graph API $(TG_t, \mathcal{R}, src, sink)$, a tainted flow $f = \cup\{DR(r_{src}, r_{sink}) | r_{src} \in src(s), r_{sink} \in sink(s), s \in T\}$ and an access control policy oracle $ACP = (RO, \leq_{RO}, P)$, then f can be divided in two subsets:

- S_f is the set of *secured tainted flows* and consists of all dependency reasons in f for which the security analyst has confirmed the correctness of the policy description based on their manual review.

- $U_f = f \setminus S_f$ the set of *unsecured tainted flows* consisting of all dependency reasons for which the security analyst did not confirm the correctness of the policy description based on their manual review.

Recall the policy requirement stated in Section 2.2 based on the running example: *a user may only update the information of a repository, if that user is the owner of the repository*; additionally, recall the tainted flow of the running example containing the pair $(createRepo, updateRepo)$. Following a manual review, the security analyst can identify that the pair is covered by the policy. Therefore, in this case, $f = S_f = \{(createRepo, updateRepo)\}$.

3.3. Dynamic analysis. The static analysis detects all potential BAC vulnerabilities, which may contain false positives, i.e., dependencies which do not constitute vulnerabilities. Therefore, the static analysis is not *complete*; to achieve completeness, we complement the static analysis with the dynamic analysis. The dynamic analysis conducts actual API executions which validate potential BAC vulnerabilities detected by the static analysis; the dynamic analysis is complete, in that, all vulnerabilities it detects are indeed BAC instances. We describe the dynamic analysis segment of our approach below.

The dynamic analysis relies on the secured flows S_f and unsecured flows U_f identified during the static analysis—step 5 in Figure 7 (left) and Definition 3.15. Drawing on S_f and U_f , the security analyst defines test cases (step 6 in Figure 7) which:

- If a dependency reason belongs to S_f , test whether the policy has been implemented correctly for this dependency reason.
- If a dependency reason belongs to U_f , expose the BAC vulnerability stemming from this dependency reason or discover an implicit undocumented implementation of the policy.

We refer to these test cases as *taint tests*, and formally define them below. We distinguish positive taint tests from negative taint tests. For negative tests we expect a BAC exception to be raised, and for positive tests we expect no exception to be raised. If this is not the case, we have detected a BAC vulnerability.

Definition 3.16 (taint test). Given a tainted Graph API $(TG_t, \mathcal{R}, src, sink)$, with *tainted flow* $f = \cup\{DR(r_{src}, r_{sink}) \mid r_{src} \in src(s), r_{sink} \in sink(s), s \in T\}$ and access control policy oracle $ACP = (RO, \leq_{RO}, P)$, a *taint test* $((t, roles_t), access)$ consists of a role-based Graph API execution $(t, roles_t)$ and an expected outcome represented by the boolean variable $access \in \{true, false\}$.

A taint test $((t, roles_t), access)$ is *positive* (*negative*) in case $access$ equals true (false).

For a systematic method to the definition of test cases, we define the *flow coverage*. This coverage criterion helps with ascertaining whether taint tests covering the tainted flow are available. Flow coverage draws from similar approaches for covering rule dependencies [RKH13, HLG13] of testing.¹

Definition 3.17 (flow coverage). Given a tainted Graph API $(TG_t, \mathcal{R}, src, sink)$, with *tainted flow* $f = \cup\{DR(r_{src}, r_{sink}) \mid r_{src} \in src(s), r_{sink} \in sink(s), s \in T\}$ and access control policy oracle $ACP = (RO, \leq_{RO}, P)$. A taint test $((t, roles_t), access)$ *covers* a dependency reason d in $DR(r_{src}, r_{sink})$ if for $(t : G_0 \xrightarrow{r_1} G_1 \xrightarrow{r_2} G_2 \xrightarrow{r_3} G_3 \dots, roles_t)$ there exists some

¹We focus in this first approach on *covering dependency reasons* for pairs of sink and source rules, instead of covering dependencies between all rule pairs, or even dependency paths in a dependency graph.

$i, j > 0$ with $i < j$ such that $r_i = r_{src}$ and $r_j = r_{sink}$, and t creates and uses elements as described in the dependency reason d .

A set of taint tests T satisfies the (*unsecured, or secured*) *flow coverage* if for each dependency reason in f (U_f , or S_f , resp.) there exists a positive and negative test in T .

We call a taint test covering a dependency reason *minimal*, when it consists of the source rule application directly followed by the sink rule application.

In order to make sure that we test with diverse roles from the access control policy, we formulate an additional coverage criterion. This criterion ensures that, for each role in the access control policy, there exists a positive and negative test such that: (i) both tests include two different API calls (ii) for positive tests, the second call is executed with a more privileged role than the first call or a role with the same privileges; for negative tests; the second call is executed with a less privileged role than the first call.

Definition 3.18 (role coverage). Given a tainted Graph API $(TG_t, \mathcal{R}, src, sink)$, with *tainted information flow* $f = \cup\{DR(r_{src}, r_{sink}) \mid r_{src} \in src(s), r_{sink} \in sink(s), s \in T\}$ and access control policy oracle $ACP = (RO, \leq_{RO}, P)$. A taint test $((t, roles_t), access)$ covers a pair (ro, ro') from $RO \times RO$ if for $(t : G_0 \xrightarrow{r_1} G_1 \xrightarrow{r_2} G_2 \xrightarrow{r_3} G_3 \dots, roles_t)$ there exists some $i, j \geq 0$ with $i < j$ such that $ro \in roles_t(i)$ and $ro' \in roles_t(j)$.

Let ro, ro' be elements from RO . A set of taint tests T satisfies *role coverage* if for each role ro in RO (1) there exists a positive test in T covering some (ro, ro') from \leq_{RO} and (2) there exists a negative test in T covering (ro, ro') with (ro, ro') from $>_{RO}$ except if ro is the least privileged role in ACP .

Based on the flow and role coverage, a security analyst is in position to generate a set of concrete test cases for a Graph API implementation. Their execution (step 7 in Figure 7) comprises the execution of API calls, as prescribed by the graph transformation sequence captured in a taint test. We have realized steps 6 and 7 by capturing these concrete test cases as unit-tests in a script which evaluates tests automatically—see Section 4. During evaluation, we observe if an expected exception is included in the API response, e.g., in GitHub such an exception indicates that the call is attempting to perform an action for which the user executing the call has insufficient privileges.

Regarding the test case evaluation results and their correspondence to taint tests, we distinguish the following cases:

- A *positive* taint test $((t, roles_t), true)$ is *successful* if its execution does *not* lead to a BAC exception.
- A *positive* taint test $((t, roles_t), true)$ *fails* if its execution leads to a BAC exception.
- A *negative* taint test $((t, roles_t), false)$ is *successful* if its execution leads to a BAC exception.
- A *negative* taint test $((t, roles_t), false)$ *fails* if its execution does *not* lead to a BAC exception.

The dynamic analysis is *complete*, in that, if a positive (or negative) taint test fails, then a BAC vulnerability has been detected. Depending on the context in which the security analysis of the Graph API is performed, if the dynamic analysis focuses on flows from S_f , the security analyst should expect to obtain successful positive and negative tests—a test evaluation that would not meet these expectations would indicate an incorrect policy implementation; if the dynamic analysis focuses on flows in U_f , the security analyst should expect to obtain failing positive or negative tests.

As an example, recall the (secured) tainted flow obtained for the running example: $f = \{(createRepo, updateRepo)\}$. A security analyst could satisfy the secured flow coverage by performing a positive and a negative test. The positive test should entail the execution of the query *createRepo* (see Figure 1) as a user with an owner role and subsequently the execution of *updateRepo* as a user with the same role. The negative test should entail the execution of *createRepo* as a user with an owner role and subsequently the execution of *updateRepo* as a user with the collaborator role, i.e, a user who is not an owner. As in this case the flow is secured, the expected outcome is that the positive test should succeed, i.e, the test should lead to no BAC exception and the repository information should be updated; the negative test should also succeed, i.e, the test should lead to a BAC exception and the repository information should not be updated. Assume a hypothetical scenario in which the security analyst concludes that the tainted flow $f = \{(createRepo, updateRepo)\}$ is not secured properly and is thus part of the unsecured tainted flow. This may be the case because no information on securing the call *updateRepo* is found in the policy description. Then, the expected outcome for the negative test would be that it fails, i.e, does not lead to a BAC exception, because indeed no proper access control for this case is implemented. On the other hand, this may be the case if the policy appears to be too severe and thus the expected outcome for the positive test would be that it fails, i.e, does lead to a BAC exception, although it should not.

The running example considers two roles: the owner and the collaborator. Regarding the role coverage, the tests described for the flow coverage above, cover the owner role: they include (i) a positive test with two calls executed by a user with an owner role (ii) a negative test with two calls, the first one executed by a user with an owner role and the second one executed by a less privileged user, i.e, the user with a collaborator role. To satisfy the role coverage, the security analyst could add one more test: a positive test with two calls executed by a user with a collaborator role. The reason the positive test suffices is that the collaborator is the least privileged role, so no negative test using this role can be defined.

4. APPLICATION OF THE TAINT ANALYSIS TO GITHUB

We applied the presented taint analysis to the GitHub API. In particular, we focused on the free version of GitHub and repositories that are owned by personal accounts.

We performed the taint analysis for two examples: the first application of the analysis builds on an *extended version of the running example* (see Section 2.1), and aims to test a basic part of the access permissions of personal accounts [Gita]; the second application reproduces a *real issue with permissions* which has been reported in the GitHub community forum, and therefore includes more details from the GitHub API. In each application, we perform the three segments of our approach—see Figure 3. We precede the presentation of the applications by relevant information on the GitHub access policy and a summary of the application setting. Following the presentation of the applications, we discuss decisions we made during the application and their implications. All artifacts mentioned in this section are contained in the *digital package* accompanying the article [KSL25]. The package also contains instructions on reproducing the applications.

4.1. GitHub Access Control Policy. Each user of GitHub has a personal account. A repository owned by a personal account has two available roles or, as they are referred to in the GitHub documentation, *permission levels*: the repository *owner* and *collaborators*.

Owners have full control over a repository, e.g., they can delete the repository, whereas collaborators are typically only allowed to read and write to the code in a repository. Roles can only be assigned via the GitHub website and not via the API; the GitHub Access Policy does not allow administrative actions—see Section 2.2. When using the GitHub API, user authentication proceeds via *tokens* associated with the personal accounts of users. While there are two types of personal access tokens, the *classic* and the *fine-grained*, in the remainder, unless otherwise indicated, we focus on the classic token—henceforth referred to simply as token. The permissions of a token correspond to the role of the token owner, but can be further limited by so-called *scopes* granted to a token. For example: the scope *repo* grants a collaborator full access to a repository, including read and write privileges as well as access to deployment statuses; the scope *repo_deployment* grants access to the deployment statuses; a token of a collaborator associated only with the *repo_deployment* scope grants the collaborator access to the deployment statuses but no access to the code.

The GitHub API access control policy requires that a user interacting with the API is authenticated using their token. The API allows the execution of calls only if the executing user has a role with sufficient permissions and the used token is associated with the necessary scopes; otherwise the API returns an error. The GitHub policy consists of relevant sections in GitHub’s online documentation which map roles and scopes to possible actions and, thus, implicitly to API calls. As an example, recall the collaborator from above whose token is only associated with the *repo_deployment*; the documentation mentions that this scope only allows access to the deployment statuses without allowing access to the code; in terms of API calls, this means that the token only enables the collaborator to fetch information about the field `deployments` from the object `Repository` in the GitHub GraphQL schema, but not to access other unrelated fields.

4.2. Application Setting. The presented applications focused on the examples and omitted other data on GitHub. In technical terms, the applications assume an empty initial graph (see Definition 3.1), although this is not required by the presented taint analysis; the analysis only requires that the initial graph does not contain any tainted nodes.

The analysis focused on three roles: the *Owner*, who has the owner permission level and the *repo* scope—giving users with this role full control over the repository plus some additional privileges, e.g., deleting the repository; the *Collaborator*, who has the collaborator permission level and the *repo* scope—giving users with this role full control over the repository minus the additional privileges of the owner; and the *NoPe-Collaborator*, who has the collaborator permission level but no scopes—granting users with this role read-only access to public information about the repository, e.g., user profile info. In technical terms, we relied on a role specification where the set of roles RO from our formal definition of a Graph API in Definition 3.2 is equal to $\{Owner, Collaborator, NoPe-Collaborator\}$, while the partial order \leq_{RO} is equal to $\{(NoPe-Collaborator, NoPe-Collaborator), (NoPe-Collaborator, Collaborator), (Collaborator, Collaborator), (Collaborator, Owner), (Owner, Owner)\}$.

The static analysis segment (see Figure 7) is performed based on the *Henshin project*. Specifically, the GraphQL schemata were manually translated into class diagrams, i.e, Ecore files, whereas the rules were captured in Henshin diagrams. The CPA (coarse) was then executed using the relevant Henshin feature, presented in [BLST17].

The dynamic analysis segment entails the creation of Python scripts with positive and negative test cases for the tainted flows obtained by the static analysis. Each test case entails a role-based API execution (see Definition 3.2), i.e, a sequence of API calls realized

	createIssue	updateRepo	createProject	updateIssue	getProject	deleteIssue	createRepo	deleteProject
createIssue				CDR		CDDR		
updateRepo								
createProject					CDR			CDDR
updateIssue								
getProject								
deleteIssue								
createRepo	CDR	CDR						
deleteProject								

FIGURE 11. The output of the dependency analysis (coarse) in Henshin for all the rules in Figure 6; CDR stands for *Create Dependency Reason* and CDDR for *Create-delete Dependency Reason*.

by GraphQL queries and mutations and executed using a given token. The scripts use three tokens, corresponding to the roles in *RO* from above. Positive tests entail the application of source and sink rules by roles with sufficient permissions, whereas negative tests entail the application of a source rule by a role with sufficient permissions and the application of the sink rule by a role with insufficient permissions. The test cases are *minimal*, i.e., an application of a source rule is directly followed by an application of a sink rule; if the application of a sink rule requires a certain setup, e.g., the creation of a project requires a repository to be present, this setup is created prior to the sink rule.

The scripts put test cases inside unit tests. The unit-tests are executed and evaluated based on the Pytest library [Kd]. The execution of the API calls in the unit-tests is performed using the Python GraphQL client [Kri]. The scripts create the necessary resources ad-hoc; after the execution of the unit tests, the scripts delete all created resources, thus returning again to the empty initial graph state from above.

4.3. Application to (Extended) Running Example. We commenced with the setup segment of the analysis and completed it as follows. First, we extended the running example: as presented in Section 2, in the tainted type graph TG_t only a single node was tainted, i.e., the *Repository*, a single source rule was considered, i.e., *createRepo*, and a single sink rule, i.e., *updateRepo*. For the extended example, since the unauthorized manipulation of projects and issues could also compromise the data security of the platform, we also tainted the nodes *Project* and *Issue* in the schema in Figure 4; moreover, according to queries and mutations in the GitHub API schema, we defined source and sink rules for the creation, deletion, and updating of instances of these newly tainted nodes. These new rules, including the previously presented *createRepo* and *updateRepo* rules, are illustrated in Figure 6.

We continued with the static analysis segment. As discussed earlier, we use Henshin for tool support of the static analysis. An illustration of the output by Henshin, i.e., the dependency reasons, of the CPA is shown in Figure 11. Overall, the CPA returned the following dependencies, i.e., tainted information flow: $(createIssue, updateIssue)$, $(createIssue, deleteIssue)$, $(createProject, getProject)$, $(createRepo, updateRepo)$, $(createProject, deleteProject)$, $(createRepo, createIssue)$. Besides the pair $(createProject, deleteProject)$, for which we found no explicit documentation, all the other detected dependencies are addressed by the documentation for repositories, issues, and projects—see [Gita]. Therefore, the pair in

question belongs to the unsecured tainted flow, whereas all other pairs belong to the secured flow—see Definition 3.15.

For the dynamic analysis segment, we created a script containing positive and negative taint tests based on the tainted flows obtained by the static analysis. In terms of Definition 3.16, a taint test consists of an API execution $(t, role_t)$ and its expected outcome *access*. For example, to address the dependency $(createIssue, updateIssue)$, the script contains a (positive) taint test consisting of the API execution $t : G_0 \xrightarrow{r_1} G_1 \xrightarrow{r_2} G_2$ with r_1 the *createIssue* rule, r_2 the *updateIssue* rule, $role_t(0) = Owner$, $role_t(1) = Collaborator$, and $access = true$; the script furthermore contains a (negative) taint test consisting of the API execution $t : G_0 \xrightarrow{r_1} G_1 \xrightarrow{r_2} G_2$ with r_1 the *createIssue* rule, r_2 the *updateIssue* rule, $role_t(0) = Owner$, $role_t(1) = NoPe-Collaborator$, and $access = false$.

The set of test cases in the script satisfies (secured and unsecured) *flow coverage* (see Definition 3.17), as it contains twelve test cases, i.e., a positive and a negative test case per dependency. Moreover, in order to satisfy the *role coverage* (see Definition 3.18), the set contains six more test cases which ensure the criteria of the role coverage are met, e.g., in the $(createIssue, updateIssue)$ test pair from above, the set adds another taint test, consisting of the API execution $t : G_0 \xrightarrow{r_1} G_1 \xrightarrow{r_2} G_2$ with r_1 the *createIssue* rule, r_2 the *updateIssue* rule, $role_t(1) = Owner$, $role_t(2) = Owner$, and $access = true$; thereby, this last test ensures that role coverage is satisfied for the owner role in *RO* in the taint test in question.

For dependencies in the secured flow, all positive and negative tests were successful, i.e., positive tests did not lead to a BAC exception whereas negative tests led to a BAC exception—see Section 3.3. The same was true for the dependency in the unsecured flow which indicates that, although undocumented, access control is implicitly implemented for this dependency by the GitHub API.

4.4. Application to GitHub Permissions Issue. To demonstrate that the proposed taint analysis can be used to test a real vulnerability, we also applied the analysis to an issue with permissions reported on the GitHub community forum—see issue 106598 in [Gitb]. The issue is related to the field `compare` in the object type `Ref`; the `Ref` object represents a branch in Git, and the `compare` field compares the latest commit in the branch represented by a given `Ref` object with the latest commit in another user-provided branch—see Figure 12. The issue is that executing the GraphQL API call that resolves this field, i.e., performs the comparison, only requires that a token has the *repo* scope when the user is authenticated using the OAuth method; if a user is authenticated using a fine-grained token (see Section 4.1), then the scope is not required. Our goal was to apply the analysis to obtain taint tests which would perform the relevant actions, i.e., creation of a `Ref` object and execution of a comparison, and subsequently reproduce the vulnerability by differentiating the authentication method used while executing the API calls in the tests.

The tainted type graph TG_t obtained in the setup segment addressed more aspects of the GitHub schema and adhered to the conventions of the GraphQL schema language. The TG_t contains fields with input parameters, e.g., the `compare` field from above—see Figure 12, that may alter the behavior of the field resolution; in UML terms, these fields resemble operations and have been modeled as such in the TG_t . The GraphQL schema language requires that object types are not used as input parameters; instead, for providing complex, non-scalar input to fields, the schema language defines *input types*, i.e., data structures, which may include object types. Our TG_t for this example fulfills these requirements. As in the GraphQL

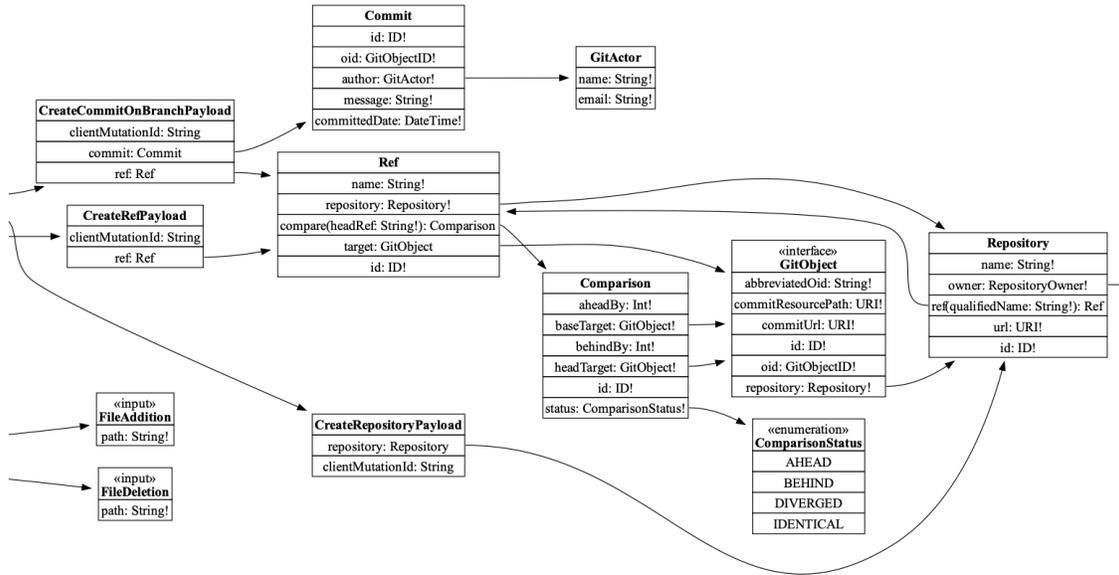
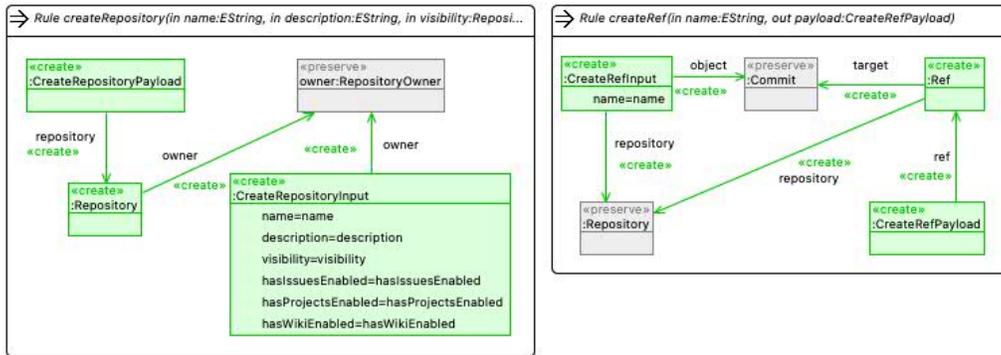


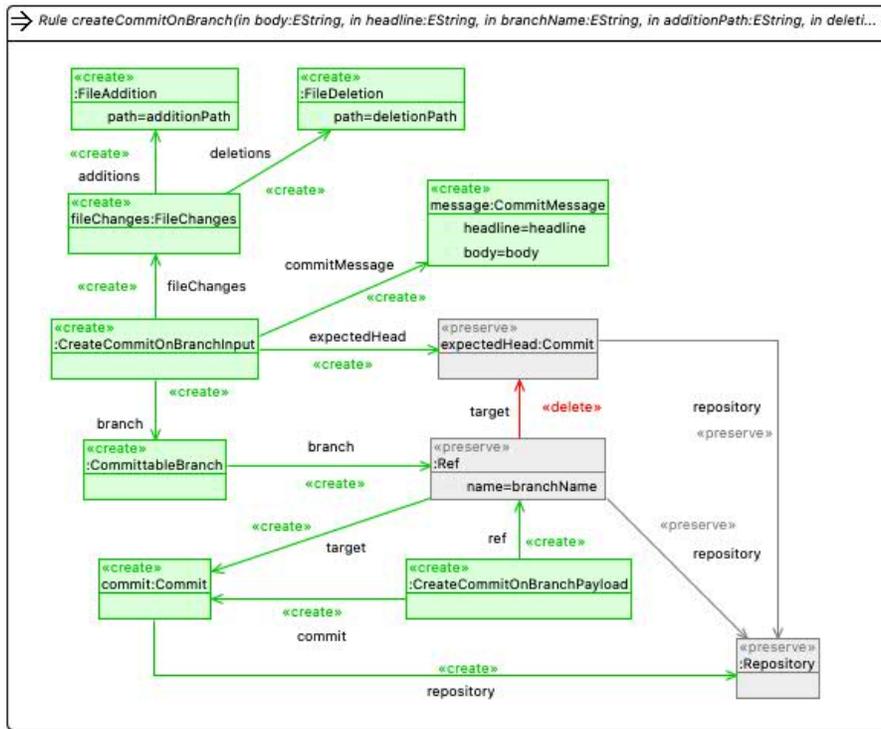
FIGURE 12. A partial graphical representation of the GraphQL schema used for the taint analysis of the GitHub permission issue in Section 4.4. The complete representation can be found in the digital package accompanying the article [KSL25].

schema language, the TG_t contains the two special types **Query** and **Mutation**; the content of these types are operations corresponding to queries and mutations in the GraphQL schema. Regarding the output of queries and mutations, the GitHub schema supports complex data structures by automatically generating *payload* objects which comprise all output values from a query or a mutation. Our TG_t correspondingly introduces payload classes for the output of queries and mutations. The GitHub schema defines enums and custom scalar, i.e. primitive, data types, which are modeled in the TG_t via enum and data type classes, respectively. Finally, the GitHub schema equips many objects with a unique global ID, which can be used to directly retrieve an object and thus to improve performance, e.g., when caching. In the TG_t , the ID fields have been replaced with a many-to-one reference.

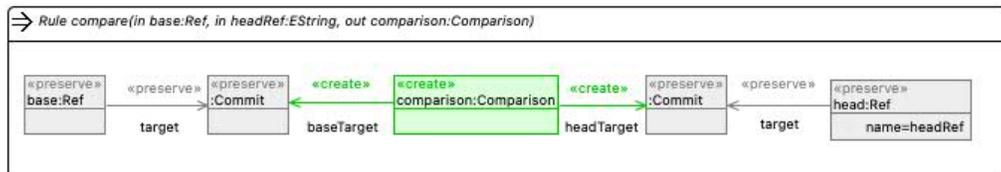
Regarding the set of rules, we focused on the mutations in the GitHub schema that create objects that are relevant to the example, i.e. mutations for creating a repository, a branch, and a commit. Finally, as mentioned above, the behavior of fields with input parameters resembles operations, therefore we also created a mutation for **compare** which, as in the GitHub schema, takes as parameter the branch to be compared to the **Ref** object from which the **compare** is resolved, and returns a **Comparison** object. The rules are illustrated in Figure 13. With respect to modeling of the rules in Henshin, we assume that the input for creating objects is provided as input parameters to the rule; however to represent the execution of a call to the GitHub API accurately, the rules entail the creation of input type objects which use the rule input parameters as fields—as in the GitHub schema. See for example the class **CreateRepositoryInput** in Figure 13a. We tainted the nodes in the TG_t that are relevant to the actions required for a comparison as well as the comparison itself, i.e. the nodes **Repository**, **Ref**, **Commit**, and **Comparison**.



(A) The rules for creating a Repository (left) and a Ref.



(B) The rule for creating a Commit.



(C) The rule for performing a Comparison.

FIGURE 13. The source and sink rules used in the permissions issue example in Section 4.4.

	createRef	compare	createCommitOnBranch	createRepository
createRef		CDR		
compare				
createCommitOnBranch	CDR	CDR(2)	CDDR	
createRepository	CDR			

FIGURE 14. The output of the dependency analysis for the rules in Figure 13.

We completed the static segment by performing CPA between source and sink rules. The result of the analysis is shown in Figure 14. The completion of the dynamic segment entailed the creation of a script which creates the necessary setup, i.e., the repository and two branches, and focuses on the dependency relevant to the reported issue, i.e., the creation of a commit followed by an execution of a comparison. The script executes two unit-tests: one is creating two commits and then performs the comparison using the OAuth authentication without the *repo* scope, whereas the other creates two commits and uses the fine-grained token, again without the *repo* scope. The execution of the script confirmed the observation reported in the issue 106598 in the GitHub community forum, i.e., the test using OAuth fails to execute the comparison, whereas the test using the fine-grained token succeeds.

4.5. Discussion. The goal of the extended running example in Section 4.3 was to demonstrate the applicability of the approach as well as to exemplify the formal results, e.g., flow and role coverage. The goal of the reproduction of the permissions issue in Section 4.4 was to obtain a Graph API which covered a larger part of the GraphQL schema language, thereby addressing the translation of more GraphQL concepts into a graph-transformation-based formalism.

Our modeling decisions when translating the GraphQL schema into the artifacts required by the taint analysis, i.e., the type graph and the set of transformation rules, were based on two considerations. First, the translation concerns transferring concepts from an implementation-oriented language into a language of higher abstraction. Therefore, our translation omitted some object types that are closely related to the implementation of applications that would use the API, e.g., types intended for facilitating the display of query results by the front-end, as well as concepts which are closely related to the runtime behavior, e.g., *directives*, that is, annotations that enable the description of alternate runtime execution and type validation behavior in GraphQL. On the other hand, to comply to the GraphQL instructions and GitHub conventions, we modeled input types and payloads. The second consideration was that the translation occurs between a textual language for the description of structures into a graphical language with the same purpose. This consideration motivated our choice to translate fields referring to unique IDs into many-to-one references.

Obtaining the type graph and rules for the permissions issue also paved the way for automating this translation, which constitutes part of our future plans for the taint analysis (see Figure 15 and Section 6). Although the translation for the permissions issue example addresses a large part of the GraphQL schema language, there are other parts which may require formalization extensions, e.g., the translation of fields with input parameters which lead to the creation of objects with other input parameters; in practice, such nested fields resemble the invocation of operations within operations and may thus require the nesting

of rules or the introduction of control flow elements. We plan to realize such extensions by employing relevant Henshin features, e.g., units.

5. RELATED WORK

To the best of our knowledge, there currently does not exist any approach to security testing for Graph APIs with simultaneous support for key features of our approach, i.e, using taint analysis, focusing on broken access control, relying on graphs, and allowing for the sound detection of vulnerabilities. In the following, we first discuss three works that are closely related to our approach, in that, they support at least three of the key features. We then discuss works that focus on either implementing or testing broken access control and rely on graph-based modeling. Subsequently, we discuss other approaches to security testing for access control. We then proceed with a discussion of static and dynamic taint analysis for security testing of mobile and web applications, and conclude the section with a discussion of testing approaches for GraphQL APIs.

The work in [MFBLT08] presents a framework for the specification, deployment, and testing of security policies, based on Model-driven Engineering. The framework relies on a class diagram for representing the system structure and on a novel security metamodel, which may be seen as an informal representation of a typed graph, for enabling the specification of rule-based access control policies. The metamodel enables the specification of policies based on several methods, e.g., role-based or organization-based. The framework supports the automated transformation of such policies into security components. In terms of testing, the authors present a fault model defined at the meta-level that is used for systematically generating test cases using mutation. Moreover, the authors argue that these policies enable intrinsic consistency (via the metamodel-model relationship) as well as manually created conformance checks. In comparison, our work focuses solely on testing for role-based policies. Moreover, our work deliberately assumes the existence of an *informal* policy document and instead requires that the security analyst models API calls, i.e, the manners in which to create and manipulate data in the system, as graph transformation rules. These rules enable a sound and systematic detection of dependencies between rules. Dependencies are then verified by the creation of tests which can be adjusted to satisfy coverage criteria. On the other hand, the generation of tests in our approach is currently manual. The automation of this process, as well as the extension of the presented formalization to support other access policy methods, are planned in our future work.

Another model-based approach is presented in [DLM20]. This work focuses on testing access control policies given in the *eXtensible Access Control Markup Language* (XACML) [OAS]. The approach represents such a policy as a tree that is then used to generate a typed graph, i.e, a model, representing the policy evaluation. Based on this typed graph, the approach systematically generates an oracle as well as tests and can, moreover, assess the coverage of a given test suite. While our approach also relies on an oracle being available, as mentioned in the comparison with the previous work, our work deliberately assumes that the policy is given in natural language. This decision differentiates the two approaches significantly: our approach enables a static analysis, similar to the work in [DLM20], via the modeling of the API calls as graph transformation rules but, moreover, entails the dynamic analysis to verify whether the policy has been implemented correctly.

The work in [MNL16] uses UML notations to specify access control, i.e, the structure of the data is modeled as a class diagram while security rules are modeled as SecureUML

and activity diagrams. The authors present rules that translate the UML notations into a specification in a formal language. The specification can then be systematically checked for consistency and validated. Subsequently, the validated specification can be used for the systematic derivation of a formal, proved access control filter. While this approach is not intended for testing but rather for implementing access control, this approach relates to our work in its modeling of the data as a graph, i.e., a class diagram, as well as in its capability for formal precision.

There exist a number of approaches formalizing, analyzing, and implementing access control policies based on *graph-based representations*. Implementation of access control based on graphs typically entails the modeling of a policy as a graph that is stored in a graph database; the access control mechanism then checks compliance to the policy based on graph queries, e.g., [YASV24, PRLGM12, JK19]. Other works implement mechanisms to control access to a graph database by rewriting queries such that are compliant to the policy before passing them on for evaluation to the database, e.g., [HMA⁺23, BRMI24].

Regarding analysis and formalization, Koch et al. [KMP05] use a sophisticated formalization to compare and analyze different policy models. We decided to equip our taint analysis with a first relatively simple access control policy model. It lowers the barrier for applying our analysis and allow cases, where only natural language descriptions for access control policies are given. We plan to integrate more complex formalizations into our taint analysis in future work. Burger et al. [BJW15] present an approach to detect (and correct) *vulnerabilities in evolving design models* of software systems. Graph transformation is used to formalize patterns used to detect security flaws. This approach is not dedicated specifically to Graph APIs, nor broken access control. The works in [RLFK04] and [AF17] use graph-based approaches to visualize and analyze access control policies, while the work in [MAHK21] extends XACML for graph-structured data. The suitability of graph-based representations for representing policies and enabling access control is also being investigated by recent approaches that realize access control via Machine Learning [YYW⁺23, YCH⁺24].

A number of approaches to testing for broken access control rely on *model-based testing*. These approaches typically generate tests for the access control based on a formal or semi-formal model of the policy or the system. For instance, the works in [MBGM09, FLWL24, PMT08] model the policy as a finite state model, as activity diagrams, and as a metamodel, respectively; the work in [XTK⁺12] models the system as a high-level Petri Net. In comparison, our approach opts for assuming the policy is only informally documented and focuses on dependencies between interactions that may cause vulnerabilities. Moreover, our approach entails the dynamic analysis that aims to verify whether the access control has been implemented correctly.

Several *static and dynamic taint analysis* approaches for security testing of *mobile or web applications* exist, see e.g. [EGC⁺14, KGT⁺15, TPF⁺09]. These approaches also analyze the data flow of security-relevant or private data, but there is currently no approach dedicated to Graph APIs, nor broken access control in particular. Pagey et al. [PMY23] present a security evaluation framework focusing also on access control vulnerabilities, but it is specifically dedicated to software-as-a-service e-commerce platforms.

The following approaches are dedicated to *testing GraphQL APIs*, but do not focus on finding *security vulnerabilities* such as broken access control. Belhadi et al. [BZA24] present a fuzzing approach based on evolutionary search. Karlsson et al. [KCS21] introduce a property-based testing approach. Vargas et al. [VBV⁺18] present a regression testing

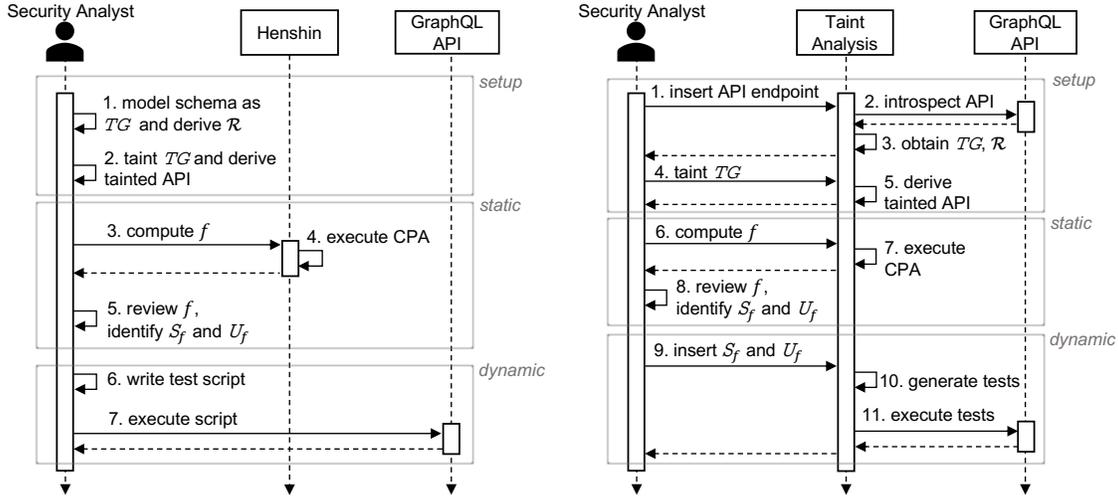


FIGURE 15. The overview of the sequence of steps in the presented approach from Figure 7 (left), and the vision for the automation of the approach (right).

approach. Zetterlund et al. [ZTMB22] present an approach to mine GraphQL queries from production with the aim of testing the GraphQL schema.

6. CONCLUSION AND FUTURE WORK

We have presented a *systematic approach supporting security analysts to test for broken access control vulnerabilities in Graph APIs*. The security testing approach is based on taint analysis and consists of a static and dynamic segment. The static taint analysis primarily aims at validating the access control policies. The dynamic taint analysis aims at finding errors in the implementation of the (validated) access control policies.

Future work comprises *increasing the level of automation* of the taint analysis as illustrated in Figure 15 (right). In particular, it is part of current work to provide automatic translations from Graph APIs such as GraphQL into tainted typed graph transformation systems. Naturally, the validation then still will include a manual review, but the formal specification could at least in parts be mined from the GraphQL implementation. There is also room for improving automation of the dynamic segment of our analysis. Currently, we provided automated test execution scripts, but we are also planning to explore the automation of the test generation step or coverage analysis for existing test suites. Moreover, we want to *evaluate the approach* on further and larger case studies and test its generality, also in the context of other relevant domains such as access control to graph databases. As mentioned in Section 3, we plan to devise a *more fine-grained* sound detection technique for potential *indirect vulnerabilities* related to broken access control relying on a new type of critical pair analysis for indirect dependencies. Moreover, as mentioned in Section 5, a further line of future work is the integration of more *detailed graph-based formalizations* of (other types of) access control. Finally, although for simplicity we omitted attribute manipulation and concepts like inheritance in the current formalization of Graph APIs, it seems feasible to incorporate such more sophisticated concepts based on the state-of-the-art for graph transformation.

ACKNOWLEDGEMENT

The authors would like to thank the anonymous reviewers for their insightful comments and suggestions.

REFERENCES

- [ABJ⁺10] Thorsten Arendt, Enrico Biermann, Stefan Jurack, Christian Krause, and Gabriele Taentzer. Henshin: Advanced Concepts and Tools for In-Place EMF Model Transformations. In Dorina C. Petriu, Nicolas Rouquette, and Øystein Haugen, editors, *Model Driven Engineering Languages and Systems*, Lecture Notes in Computer Science, pages 121–135, Berlin, Heidelberg, 2010. Springer. doi:10.1007/978-3-642-16145-2{_}9.
- [AF17] Sandra Alves and Maribel Fernández. A graph-based framework for the analysis of access control policies. *Theoretical Computer Science*, 685:3–22, July 2017. URL: <https://www.sciencedirect.com/science/article/pii/S0304397516305965>, doi:10.1016/j.tcs.2016.10.018.
- [BJW15] Jens Bürger, Jan Jürjens, and Sven Wenzel. Restoring security of evolving software models using graph transformation. *Int. J. Softw. Tools Technol. Transf.*, 17(3):267–289, 2015. URL: <https://doi.org/10.1007/s10009-014-0364-8>, doi:10.1007/s10009-014-0364-8.
- [BLST17] Kristopher Born, Leen Lambers, Daniel Strüber, and Gabriele Taentzer. Granularity of Conflicts and Dependencies in Graph Transformation Systems. In Juan de Lara and Detlef Plump, editors, *Graph Transformation*, Lecture Notes in Computer Science, pages 125–141, Cham, 2017. Springer International Publishing. doi:10.1007/978-3-319-61470-0_8.
- [BRMI24] Adil Achraf Bereksi Reguig, Houari Mahfoud, and Abdessamad Imine. Towards an Effective Attribute-Based Access Control Model for Neo4j. In Mohamed Mosbah, Tahar Kechadi, Ladjel Bellatreche, and Faiez Gargouri, editors, *Model and Data Engineering*, pages 352–366, Cham, 2024. Springer Nature Switzerland. doi:10.1007/978-3-031-49333-1_25.
- [BZA24] Asma Belhadi, Man Zhang, and Andrea Arcuri. Random testing and evolutionary testing for fuzzing graphql apis. *ACM Trans. Web*, 18(1):14:1–14:41, 2024. doi:10.1145/3609427.
- [DLM20] Said Daoudagh, Francesca Lonetti, and Eda Marchetti. XACMET: XACML Testing & Modeling. *Software Quality Journal*, 28(1):249–282, March 2020. doi:10.1007/s11219-019-09470-5.
- [EEPT06] Hartmut Ehrig, Karsten Ehrig, Ulrike Prange, and Gabriele Taentzer. *Fundamentals of Algebraic Graph Transformation*. Monographs in Theoretical Computer Science. An EATCS Series. Springer-Verlag, Berlin/Heidelberg, 2006. doi:10.1007/3-540-31188-2.
- [EGC⁺14] William Enck, Peter Gilbert, Byung-Gon Chun, Landon P. Cox, Jaeyeon Jung, Patrick D. McDaniel, and Anmol Sheth. Taintdroid: an information flow tracking system for real-time privacy monitoring on smartphones. *Commun. ACM*, 57(3):99–106, 2014. doi:10.1145/2494522.
- [FLWL24] Ao Fan, Li Liao, Lulu Wang, and Bixin Li. Test Case Generation for Access Control Based on UML Activity Diagram. In *2024 IEEE 24th International Conference on Software Quality, Reliability and Security (QRS)*, pages 95–104, July 2024. ISSN: 2693-9177. URL: <https://ieeexplore.ieee.org/abstract/document/10684652>, doi:10.1109/QRS62785.2024.00019.
- [GFa] The GraphQL Foundation. GraphQL | A query language for your API. <https://graphql.org/>. accessed: 2023-12-05.
- [GFb] The GraphQL Foundation. Schemas and Types | GraphQL. <https://graphql.org/learn/schema/>. accessed: 2024-03-08.
- [Gita] Inc GitHub. Access permissions on GitHub. <https://docs.github.com/en/get-started/learning-about-github/access-permissions-on-github>. accessed: 2024-03-08.
- [Gitb] Inc GitHub. GitHub GraphQL API community discussions. <https://github.com/orgs/community/discussions/>. accessed: 2024-03-01.
- [Gite] Inc GitHub. GitHub GraphQL API documentation. <https://docs.github.com/en/graphql>. accessed: 2024-03-01.
- [Gitd] Inc GitHub. Permission levels for a personal account repository. <https://docs.github.com/en/account-and-profile/setting-up-and-managing-your-personal-account-on-github/managing-user-account-settings/permission-levels-for-a-personal-account-repository>. accessed: 2024-03-06.
- [Gra07] Gary Mc Graw. *Software Security: Building Security in*. Addison-Wesley Professional, 2007.

- [HKY17] Vincent C Hu, Rick Kuhn, and Dylan Yaga. Verification and Test Methods for Access Control Policies/Models. Technical Report NIST SP 800-192, National Institute of Standards and Technology, Gaithersburg, MD, June 2017. doi:10.6028/NIST.SP.800-192.
- [HLG13] Stephan Hildebrandt, Leen Lambers, and Holger Giese. Complete specification coverage in automatically generated conformance test cases for TGG implementations. In Keith Duddy and Gerti Kappel, editors, *Theory and Practice of Model Transformations - 6th International Conference, ICMT@STAF 2013, Budapest, Hungary, June 18-19, 2013. Proceedings*, volume 7909 of *Lecture Notes in Computer Science*, pages 174–188. Springer, 2013. doi:10.1007/978-3-642-38883-5_16.
- [HMA⁺23] Daniel Hofer, Aya Mohamed, Dagmar Auer, Stefan Nadschläger, and Josef Küng. Rewriting Graph-DB Queries to Enforce Attribute-Based Access Control. In Christine Strauss, Toshiyuki Amagasa, Gabriele Kotsis, A. Min Tjoa, and Ismail Khalil, editors, *Database and Expert Systems Applications*, pages 431–436, Cham, 2023. Springer Nature Switzerland. doi:10.1007/978-3-031-39847-6_34.
- [HT20] Reiko Heckel and Gabriele Taentzer. *Graph Transformation for Software Engineers - With Applications to Model-Based Development and Domain-Specific Language Engineering*. Springer, 2020. doi:10.1007/978-3-030-43916-3.
- [JK19] Ying Jin and Krishna Chaitanya Kaja. Xacml implementation based on graph database. *Proceedings of 34th International Confer*, 58:65–74, 2019.
- [KCS21] Stefan Karlsson, Adnan Causevic, and Daniel Sundmark. Automatic property-based testing of graphql apis. In *2nd IEEE/ACM International Conference on Automation of Software Test, AST@ICSE 2021, Madrid, Spain, May 20-21, 2021*, pages 1–10. IEEE, 2021. doi:10.1109/AST52587.2021.00009.
- [Kd] Holger Krekel and dev team pytest. Pytest: Helps you write better programs — pytest documentation. <https://docs.pytest.org/en/8.0.x/>. accessed: 2024-03-06.
- [KGT⁺15] Konstantin Kuznetsov, Alessandra Gorla, Iaria Tavecchia, Florian Gross, and Andreas Zeller. Mining android apps for anomalies. In Christian Bird, Tim Menzies, and Thomas Zimmermann, editors, *The Art and Science of Analyzing Software Data*, pages 257–283. Morgan Kaufmann / Elsevier, 2015. URL: <https://doi.org/10.1016/b978-0-12-411519-4.00010-0>, doi:10.1016/B978-0-12-411519-4.00010-0.
- [KMP05] Manuel Koch, Luigi V. Mancini, and Francesco Parisi-Presicce. Graph-based specification of access control policies. *J. Comput. Syst. Sci.*, 71(1):1–33, 2005. URL: <https://doi.org/10.1016/j.jcss.2004.11.002>, doi:10.1016/J.JCSS.2004.11.002.
- [Kri] Justin Krinke. Python-graphql-client: Python GraphQL Client.
- [KSL25] Taisiya Khakharova, Lucas Sakizloglou, and Leen Lambers. Supplement to "taint analysis for graph apis focusing on broken access control" (LMCS journal), January 2025. accessed: 2025-01-15. doi:10.5281/zenodo.14524445.
- [LBK⁺19] Leen Lambers, Kristopher Born, Jens Kosiol, Daniel Strüber, and Gabriele Taentzer. Granularity of conflicts and dependencies in graph transformation systems: A two-dimensional approach. *Journal of Logical and Algebraic Methods in Programming*, 103:105–129, February 2019. doi:10.1016/j.jlamp.2018.11.004.
- [LBO⁺18] Leen Lambers, Kristopher Born, Fernando Orejas, Daniel Strüber, and Gabriele Taentzer. Initial Conflicts and Dependencies: Critical Pairs Revisited. In Reiko Heckel and Gabriele Taentzer, editors, *Graph Transformation, Specifications, and Nets: In Memory of Hartmut Ehrig*, Lecture Notes in Computer Science, pages 105–123. Springer International Publishing, Cham, 2018. doi:10.1007/978-3-319-75396-6_6.
- [LEO08] Leen Lambers, Hartmut Ehrig, and Fernando Orejas. Efficient Conflict Detection in Graph Transformation Systems by Essential Critical Pairs. *Electronic Notes in Theoretical Computer Science*, 211:17–26, April 2008. doi:10.1016/j.entcs.2008.04.026.
- [LSAWK24] Leen Lambers, Lucas Sakizloglou, Osama Al-Wardi, and Taisiya Khakharova. Taint Analysis for Graph APIs Focusing on Broken Access Control. In Russ Harmer and Jens Kosiol, editors, *Graph Transformation*, pages 180–200, Cham, 2024. Springer Nature Switzerland. tex.ids=Lambers_2024_TaintAnalysisGraphAPIsFocusingBrokenAccessControla. doi:10.1007/978-3-031-64285-2_10.

- [LST⁺18] Leen Lambers, Daniel Strüber, Gabriele Taentzer, Kristopher Born, and Jevgenij Huebert. Multi-granular conflict and dependency analysis in software engineering based on graph transformation. In *Proceedings of the 40th International Conference on Software Engineering, ICSE '18*, pages 716–727, New York, NY, USA, May 2018. Association for Computing Machinery. doi:10.1145/3180155.3180258.
- [MAHK21] Aya Mohamed, Dagmar Auer, Daniel Hofer, and Josef Küng. Extended Authorization Policy for Graph-Structured Data. *SN Computer Science*, 2(5):351, June 2021. doi:10.1007/s42979-021-00684-8.
- [MBGM09] Ammar Masood, Rafae Bhatti, Arif Ghafoor, and Aditya P. Mathur. Scalable and Effective Test Generation for Role-Based Access Control Systems. *IEEE Transactions on Software Engineering*, 35(5):654–668, September 2009. URL: <https://ieeexplore.ieee.org/abstract/document/4967616>, doi:10.1109/TSE.2009.35.
- [Met] Inc Meta. Overview - Graph API - Documentation. <https://developers.facebook.com/docs/graph-api/overview/>. accessed: 2024-03-01.
- [MFBLT08] Tejeddine Mouelhi, Franck Fleurey, Benoit Baudry, and Yves Le Traon. A Model-Based Framework for Security Policy Specification, Deployment and Testing. In *Lecture Notes in Computer Science*, pages 537–552. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008. ISSN: 0302-9743, 1611-3349. URL: http://link.springer.com/10.1007/978-3-540-87875-9_38, doi:10.1007/978-3-540-87875-9_38.
- [MNL16] Amel Mammam, Thi Mai Nguyen, and Régine Laleau. Formal Development of a Secure Access Control Filter. In *2016 IEEE 17th International Symposium on High Assurance Systems Engineering (HASE)*, pages 173–180, January 2016. ISSN: 1530-2059. URL: <https://ieeexplore.ieee.org/document/7423149/>, doi:10.1109/HASE.2016.10.
- [OAS] OASIS. eXtensible Access Control Markup Language (XACML) Version 3.0. URL: <https://docs.oasis-open.org/xacml/3.0/xacml-3.0-core-spec-os-en.html>.
- [Plu94] Detlef Plump. Critical pairs in term graph rewriting. In Igor Prívvara, Branislav Rován, and Peter Ruzička, editors, *Mathematical Foundations of Computer Science 1994*, Lecture Notes in Computer Science, pages 556–566, Berlin, Heidelberg, 1994. Springer. doi:10.1007/3-540-58338-6_102.
- [PMT08] Alexander Pretschner, Tejeddine Mouelhi, and Yves Le Traon. Model-Based Tests for Access Control Policies. In *and Validation 2008 1st International Conference on Software Testing, Verification*, pages 338–347, April 2008. ISSN: 2159-4848. URL: <https://ieeexplore.ieee.org/abstract/document/4539561>, doi:10.1109/ICST.2008.44.
- [PMY23] Rohan Pagey, Mohammad Mannan, and Amr Youssef. All your shops are belong to us: Security weaknesses in e-commerce platforms. In *Proceedings of the ACM Web Conference 2023*, WWW '23, page 2144–2154, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3543507.3583319.
- [PRLGM12] Santiago Pina Ros, Mario Lischka, and Félix Gómez Mármol. Graph-based XACML evaluation. In *Proceedings of the 17th ACM symposium on Access Control Models and Technologies, SACMAT '12*, pages 83–92, New York, NY, USA, June 2012. Association for Computing Machinery. URL: <https://dl.acm.org/doi/10.1145/2295136.2295153>, doi:10.1145/2295136.2295153.
- [PY07] Mauro Pezzè and Michal Young. *Software testing and analysis - process, principles and techniques*. Wiley, 2007.
- [QFGR23] Antonio Quiña-Mera, Pablo Fernandez, José María García, and Antonio Ruiz-Cortés. GraphQL: A systematic mapping study. *ACM Comput. Surv.*, 55(10):202:1–202:35, 2023. doi:10.1145/3561818.
- [RKH13] Olga Runge, Tamim Ahmed Khan, and Reiko Heckel. Test case generation using visual contracts. *Electron. Commun. Eur. Assoc. Softw. Sci. Technol.*, 58, 2013. URL: <https://doi.org/10.14279/tuj.eceasst.58.847>, doi:10.14279/TUJ.ECEASST.58.847.
- [RLFK04] Indrakshi Ray, Na Li, Robert B. France, and Dae-Kyoo Kim. Using uml to visualize role-based access control constraints. In Trent Jaeger and Elena Ferrari, editors, *9th ACM Symposium on Access Control Models and Technologies, SACMAT 2004, Yorktown Heights, New York, USA, June 2-4, 2004, Proceedings*, pages 115–124. ACM, 2004. doi:10.1145/990036.990054.
- [Roz97] Grzegorz Rozenberg, editor. *Handbook of Graph Grammars and Computing by Graph Transformations, Volume 1: Foundations*. World Scientific, 1997.

- [Tea] Team OWASP Top 10. A01 Broken Access Control - OWASP Top 10:2021. https://owasp.org/Top10/A01_2021-Broken_Access_Control/. accessed: 2024-03-06.
- [TPF⁺09] Omer Tripp, Marco Pistoia, Stephen J. Fink, Manu Sridharan, and Omri Weisman. TAJ: effective taint analysis of web applications. In Michael Hind and Amer Diwan, editors, *Proceedings of the 2009 ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2009, Dublin, Ireland, June 15-21, 2009*, pages 87–97. ACM, 2009. doi:10.1145/1542476.1542486.
- [VBV⁺18] Daniela Meneses Vargas, Alison Fernandez Blanco, Andreina Cota Vidaurre, Juan Pablo Sandoval Alcocer, Milton Mamani Torres, Alexandre Bergel, and Stephane Ducasse. Deviation Testing: A Test Case Generation Technique for GraphQL APIs. In *International Workshop on Smalltalk Technologies (IWST)*, 2018.
- [XTK⁺12] Dianxiang Xu, Lijo Thomas, Michael Kent, Tejeddine Mouelhi, and Yves Le Traon. A model-based approach to automated testing of access control policies. In *Proceedings of the 17th ACM symposium on Access Control Models and Technologies, SACMAT '12*, pages 209–218, New York, NY, USA, June 2012. Association for Computing Machinery. URL: <https://dl.acm.org/doi/10.1145/2295136.2295173>, doi:10.1145/2295136.2295173.
- [YASV24] Mian Yang, Vijayalakshmi Atluri, Shamik Sural, and Jaideep Vaidya. A Graph-Based Framework for ABAC Policy Enforcement and Analysis. In Anna Lisa Ferrara and Ram Krishnan, editors, *Data and Applications Security and Privacy XXXVIII*, pages 3–23, Cham, 2024. Springer Nature Switzerland. doi:10.1007/978-3-031-65172-4_1.
- [YCH⁺24] Jiao Yin, Guihong Chen, Wei Hong, Jinli Cao, Hua Wang, and Yuan Miao. A heterogeneous graph-based semi-supervised learning framework for access control decision-making. *World Wide Web*, 27(4):35, May 2024. doi:10.1007/s11280-024-01275-2.
- [YYW⁺23] Mingshan You, Jiao Yin, Hua Wang, Jinli Cao, Kate Wang, Yuan Miao, and Elisa Bertino. A knowledge graph empowered online learning framework for access control decision-making. *World Wide Web*, 26(2):827–848, March 2023. doi:10.1007/s11280-022-01076-5.
- [ZTMB22] Louise Zetterlund, Deepika Tiwari, Martin Monperrus, and Benoit Baudry. Harvesting production graphql queries to detect schema faults. In *15th IEEE Conference on Software Testing, Verification and Validation, ICST 2022, Valencia, Spain, April 4-14, 2022*, pages 365–376. IEEE, 2022. doi:10.1109/ICST53961.2022.00014.
- [ZZX⁺23] Neng Zhang, Ying Zou, Xin Xia, Qiao Huang, David Lo, and Shanping Li. Web apis: Features, issues, and expectations – a large-scale empirical study of web apis from two publicly accessible registries using stack overflow and a user survey. *IEEE Transactions on Software Engineering*, 49(2):498–528, 2023. doi:10.1109/TSE.2022.3154769.