

# An Extensible Framework for Architecture-Based Data Flow Analysis for Information Security

Nicolas Boltz\*, Sebastian Hahner\*,  
Christopher Gerking, and Robert Heinrich

Karlsruhe Institute for Technology (KIT)  
{boltz,hahner,gerking,heinrich}@kit.edu

**Abstract.** The growing interconnection between software systems increases the need for security already at design time. Security-related properties like confidentiality are often analyzed based on data flow diagrams (DFDs). However, manually analyzing DFDs of large software systems is bothersome and error-prone, and adjusting an already deployed software is costly. Additionally, closed analysis ecosystems limit the reuse of modeled information and impede comprehensive statements about a system’s security. In this paper, we present an open and extensible framework for data flow analysis. The central element of our framework is our new implementation of a well-validated data-flow-based analysis approach. The framework is compatible with DFDs and can also extract data flows from the Palladio architectural description language. We showcase the extensibility with multiple model and analysis extensions. Our evaluation indicates that we can analyze similar scenarios while achieving higher scalability compared to previous implementations.

**Keywords:** Data Flow Diagram · Software Architecture · Security.

## 1 Introduction

As our modern world becomes increasingly digitized, the integration of various digital services into our daily lives has become more prevalent. To enhance the quality of service, a growing amount of data is stored and processed, e.g., online shops utilizing purchase history data for recommendations. The seamless exchange of such collected data between different services or systems is a common practice. In scenarios like online shopping, sensitive information like payment details and customer addresses are involved. Consequently, security becomes a central concern in designing and building such software-intensive systems.

Information security has several definitions, e.g., as the CIA triad of confidentiality, integrity, and availability, or in ISO 27000 [17]. More recent legal regulations, like the General Data Protection Regulation (GDPR) [11], define information security more broadly. For modern systems, changes and reconfiguration in the context, environment, or internal structure might occur frequently

---

\* Both main authors contributed equally.

[36]. Since the protection goals are highly dependent on the system under consideration, the protection goals that must be addressed may also change. In addition to the CIA goals, other protection goals might be considered, like privacy, authenticity, non-repudiation, accountability, and auditability. A system violating confidentiality or privacy can cause costly fines, as seen in the case of H&M [16] or British Airways [4]. However, identifying such violations can be difficult, because the interconnected software systems represent complex networks of data flows. Hence, a holistic and scalable approach to analyzing them is required.

Data flow analyses based on source code, e.g., JOANA [33], KeY [1], or CodeQL [9], cannot consider context information, such as deployment. However, such information can be essential for information security, e.g., whether the application is deployed to an external cloud provider or not. In addition, source code analyses cannot be used in early design phases because of their need for existing source code. Analyzing the system during design time is beneficial because fixing issues in later phases is usually more costly [31]. Seifermann et al. [28, 30] proposed an architecture-based data flow analysis to analyze software systems for confidentiality violations. Their approach considers additional context information, such as the deployment, enabling software architects to analyze confidentiality during early design phases. However, the original Prolog-based implementation of Seifermann et al. [30] is hard to maintain and has a high resource demand, which severely limits the applicability for large software systems. Although they already used a model of a data flow diagram (DFD) [10] as an intermediate representation during their analysis, they did not continue to follow the idea of using DFDs as the primary model artifact. With appropriate tool support, DFDs represent a powerful and commonly used mechanism for threat analysis [3] that helps in correctly identifying security-related issues [24].

In this paper, we present an extensible analysis framework centered around our previously presented new implementation of the aforementioned approach to data flow analysis [27]. Our framework addresses shortcomings regarding the limited input capabilities, the limited intermediate use of DFDs, and problems with maintainability, scalability, and extensibility:

- C1** We propose a novel DFD metamodel. In contrast to the previous DFD model [30] we do not consider DFDs as intermediate system representations but as primary software architecture modeling artifacts. As part of our framework, we provide means to manually define DFDs as well as automatically derive them from the architecture description language Palladio Component Model (PCM) [21] and other third-party diagram representations [23, 25].
- C2** We present a new Java-based implementation of the analysis approach of Seifermann et al. [30], which is based on a newly developed internal data structure and alleviates the need to mix Java with other technologies like Prolog. In addition, we provide new forms of input, e.g., DFDs defined with our metamodel (**C1**), and a domain-specific language (DSL) that enables software architects to define constraints or queries for the analysis. We demonstrate how our analysis framework can be applied and extended for other security concerns, e.g., regarding the GDPR [6], or uncertainty [12].

This paper is structured as follows: Section 2 introduces our data flow analysis framework. Section 3 describes our new DFD metamodel (**C1**). In Section 4, we describe the analysis (**C2**), and in Section 5, we showcase existing extensions. Section 6 presents our evaluation and Section 7 concludes the paper.

## 2 Overview of the Data Flow Analysis Framework

In this section, we summarize our framework which is explained in more detail hereafter. Figure 1 gives an informal overview of the structure and the dependencies between the different parts of our data flow analysis framework. We highlight analyses and editors with a bold border; all other rectangles represent models.

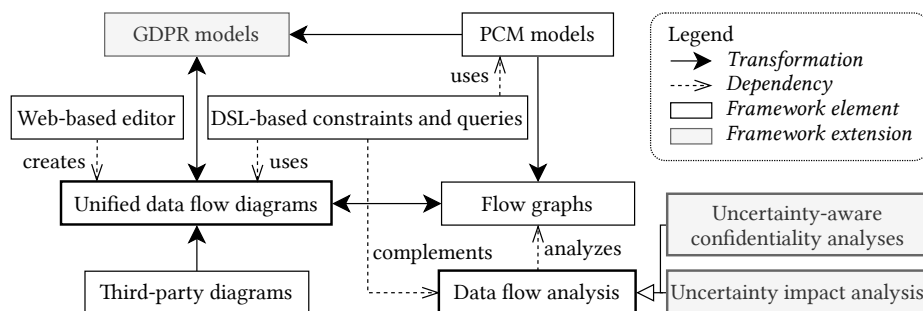


Fig. 1. Informal overview of the structure of the data flow analysis framework.

Unified data flow diagrams (DFDs) [28] play a central role in our framework (**C1**). These diagrams can be manually defined using a web-based editor or transformed from third-party diagram representations. Additionally, we enable the extraction of data flows from software architecture models described using PCM [21]. To unify the analysis process (**C2**), we transform DFD or PCM models into a set of Directed Acyclic Graphs (DAGs), called *transpose flow graphs (TFGs)*. Each vertex represents one individual data processing operation from either a DFD or PCM model, which simplifies the analysis. Using label propagation on TFGs, the analysis finds violations of predefined constraints that reflect information security objectives, e.g., confidentiality requirements. Constraints and queries can be specified using a domain-specific language (DSL), referencing PCM and DFD models. Framework extensions are possible by transforming into our DFDs or TFGs or by inheriting from the data flow analysis. GDPR models [6] are an example of the former while uncertainty impact analysis [12] and uncertainty-aware confidentiality analyses [8, 14, 35] showcase the latter. The framework is tool-supported and available as open source<sup>1</sup>. This includes modeling support, automated model transformations, and a DSL-supported analysis<sup>2</sup>. We also provide a dataset [7] including all tooling, code artifacts, and evaluation data.

<sup>1</sup> See <https://dataflowanalysis.org> or <https://github.com/DataFlowAnalysis>

<sup>2</sup> Video demonstration available [27]: <https://youtube.com/watch?v=q3WJsMyqJcA>

### 3 Modeling and Deriving Data Flow Diagrams

While Seifermann et al. [28] created a unified DFD notation, they only used it as an intermediary representation for their data flow analysis [30]. However, as DFDs are an established software architecture representation [3] and are widely used to analyze various types of data security [2, 30, 32, 34], we present an explicit DFD metamodel that can also be used as input for our analysis. The study of Bernsmed et al. [3] concludes, that, while DFDs are good for evaluating security, there exist challenges in preparing DFDs. Especially tooling that improves the effort of creating and maintaining DFDs is missing. In this section, we present our DFD metamodel (C1). We also provide tooling centered around our metamodel, which aims to aid in the creation of new DFDs, the import of already existing DFD notations [23, 25], and the automated derivation and visualization of DFDs from system architecture models like the PCM [21].

#### 3.1 Unified Data Flow Diagram Metamodel

DFDs, as proposed by DeMarco [10], can be represented as DAGs showing the data flow and processing in software systems. *Nodes* in these graphs represent *External* entities like users, *Processes* that can alter data, or *Stores* like databases, connected by *Flows* of data. Seifermann et al. [30] extend the notation by integrating several strands of work from different research groups into one unified metamodel. Figure 2 shows our metamodel that aligns with the unified DFD notation. It is split into the so-called *Data Dictionary* [10] and the DFD. The *Data Dictionary* does not directly depend on a modeled system and can thus be reused while DFD elements are specific to a certain system.

The central part of the unified notation is the representation of behavior and characteristics as first-class entities. *Labels* represent characteristics in the DFD, e.g., specifying the sensitivity of data, or the role of a user. They can either be defined as a characteristic of a *Node* or as a characteristic of data flowing between *Nodes*. *Labels* are grouped in *LabelTypes*. The *Behavior* of *Nodes* defines which *Labels* flow from one *Node* to the next via the connecting *Flow*. It is made up of *Pins* and *Assignments*. Input *Pins* represent required interfaces and output *Pins* represent provided interfaces of nodes. If a node has a certain *Behavior*, it also has the corresponding input and output *Pins*. A *Flow* connects two *Nodes* by connecting an output *Pin* of the source *Node* to an input *Pin* of the destination *Node*. *Assignments* define which labels flow out of a node. They reference input and output *Pins* of their corresponding *Behavior* and aggregate all *Labels* of the data flowing in through the input *Pin*. By evaluating a logical statement defined in the assignment, it is determined how the incoming *Labels* are changed and passed on via the referenced output *Pins*, e.g., the encryption of data can be represented by an *Assignment* that adds an *encrypted* label to the flowing data.

For assignments, we define two subclasses: *Assignment* which contains a freely definable logical *Term* that is evaluated to decide if a set of *Labels* is applied to the output *Pin*. The *ForwardingAssignment* does not define a logical term but specifies that all *Labels* that flow into the input pins are combined and

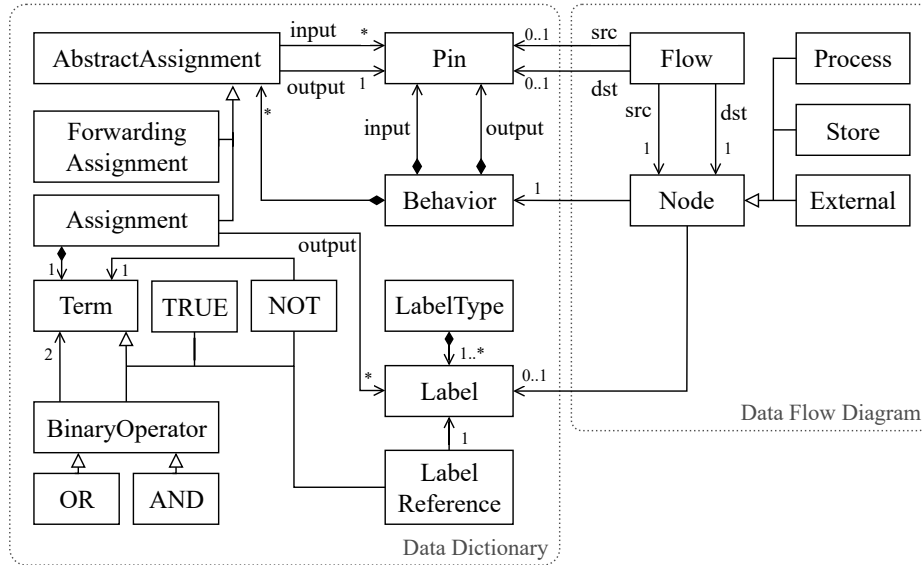


Fig. 2. Metamodel of data flow diagrams and data dictionaries.

directly forwarded to the output pin. The logical terms can be nested with binary operators *AND* and *OR* and negated with *NOT* to express different statements. *LabelReferences* are evaluated by checking if the referenced *Label* flows into the node through one of the input pins of the *Assignment*. In this case, the *LabelReference* evaluates to true, otherwise to false. The *Assignments* of a *Behavior* are ordered. If a *Behavior* contains multiple *Assignments*, first all *ForwardingAssignments* are evaluated and the *Labels* for each output *Pin* are saved. Other *Assignments* add or remove labels for their specific output *Pin*, depending on if their *Term* evaluates to true or false. Once all *Assignments* are evaluated, the *Labels* flow to the next *Node*.

### 3.2 Manually Defining Data Flow Diagrams

Manual ways to define DFDs that go beyond drawing on either paper or in software are limited. With our approach, we therefore offer a ready-to-use web editor to manually define DFDs and means to import DFDs from other notations.

Our web-based editor uses a notation that is compatible with the unified DFD notation from Figure 2. We also incorporated the concept of the data dictionaries. The graphical syntax follows earlier definitions of DFDs [10, 30]. Figure 3 shows the editor with an exemplary DFD. The toolbar on the right allows the creation of the three node types, data flow edges, and input and output pins via drag and drop. The *Label Types* field allows the creation of label types and corresponding labels. Created labels can be annotated to a node by drag and drop. Double-clicking on an output pin opens an editor for specifying assignments for the corresponding output pin. Assignments can be defined in

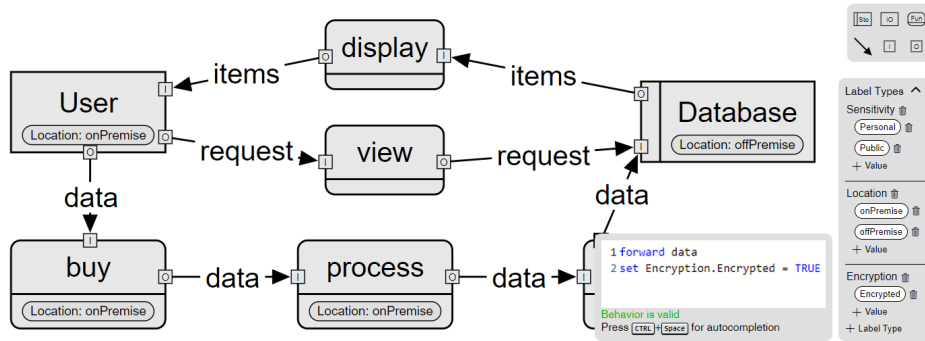


Fig. 3. Screenshot of the web-based editor showing the DFD of a simplified online shop.

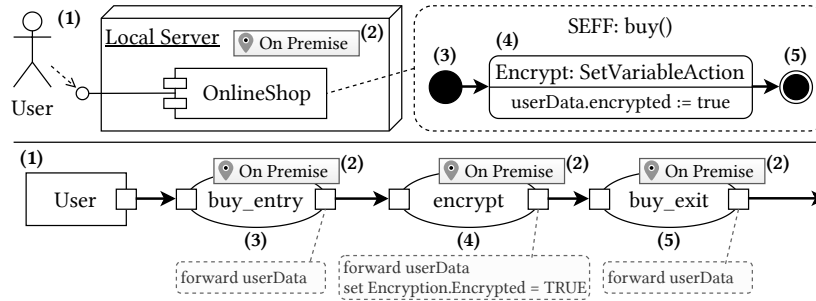
textual form using a DSL. The *forward* keyword is used to forward all labels of the corresponding input. The *set* keyword is used to define an output label and a logical term, similar to the DFD metamodel. Incoming data can be referenced via the name of the incoming edge. Labels are referenced by label type and label name. Assignments are automatically syntax-checked, and issues are reported to the user. Additionally, our web editor supports highlighting in different colors and providing tooltips for nodes. This can be used to, e.g., visualize analysis results and provide additional information regarding identified security violations.

The manually created DFDs can be exported as JSON files. To integrate the editor into our framework, we offer tooling that converts the JSON files of the web editor into an instance of our DFD metamodel. The editor is implemented in TypeScript and uses the open-source diagramming framework Eclipse Sprotty. To ease the adoption of our approach, we additionally created extensible tooling for generating instances of our DFD metamodel from various inputs. At the time of writing, we support DFD notations in PlantUML and two different types of JSON notations. To showcase this functionality, we have processed all security-enriched DFDs of the microSecEnD dataset of Schneider et al. [25]. The resulting instances of our DFD metamodel can be found in our dataset [7].

### 3.3 Automatically Deriving Data Flows from Architectural Models

Besides the manual modeling of DFDs, our framework also supports the automated extraction of data flows from the architecture description language PCM [21]. We choose PCM as it has already been used by previous data flow analysis approaches [29, 30]. However, the described concept of data flow extraction is also applicable to other modeling languages like UML.

Figure 4 shows a simplified PCM model of an online shop. It comprises information about components (e.g., *Online Shop*), resources (e.g., *Local Server*), deployment, usage, and system behavior as Service Effect Specification (SEFF) [21]. The model is annotated with confidentiality-related labels that represent characteristics of data storage like *On Premise* and data processing like the encryption of *userData* in the *SetVariableAction*. In the lower half, we show the



**Fig. 4.** Simplified PCM model of the online shop example and the corresponding data flow with annotated node labels, data labels, and numbered transformation traces.

extracted data flow. We annotate numbers to represent the transformation traces from PCM to DFD. Note, that this is only a simplified example; realistic software systems contain more than one data flow and several hundred nodes [12].

Every action in the usage and system behavior is transformed into one DFD node. This includes calls from the user, external calls between components, start and end nodes, and internal data processing nodes. The nodes' pins correspond to the in and outgoing data types, e.g., *userData*. For every node, we perform a lookup of node labels, which can be annotated, e.g., to resources, or usage scenarios. An exemplary lookup in the PCM model goes from the *encrypt* node to the *Online Shop* component via the deployment to the *Local Server* resource which is annotated with *On Premise*. Additionally, we convert the modeled system behavior to assignments of our DFD metamodel like the encryption of *userData* can be expressed. The default case is the forwarding of labels.

The transformation considers all information that is relevant for security analysis, e.g., data processing and characteristics. Other information is not transformed, e.g., components and servers do not cause additional elements in the DFD. This enables a system view from the perspective of the data which is especially suitable for properties like confidentiality [29]. However, we store all traces to the originating PCM elements during the transformation. This enables the evaluation of advanced queries and constraints in the data flow analysis.

## 4 Data Flow Analysis

The original Prolog-based analysis of Seifermann et al. [30] realized the extraction of data flows and propagation of labels by first transforming the PCM models to an explicit DFD metamodel notation, then transforming the DFD elements to Prolog. Data flow constraints were checked by defining Prolog queries that are unique to the modeled system. As one DFD element with characteristic labels is transformed into multiple Prolog statements, the Prolog code grows exponentially with the model size. The exponential growth results in high demand for memory, as the whole Prolog program needs to be fully loaded by the Prolog interpreter. As the analysis is made up of multiple chained transformations and intermediate model representations, the maintenance of the analysis was made even harder.

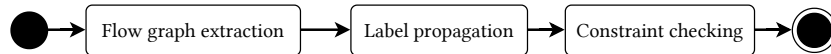
Additionally, the approach of modeling data flows via logical statements in Prolog can lead to increased runtimes: Due to the lazy evaluation of Prolog, the Prolog-based analysis needs to reevaluate the characteristic labels of nodes for each different constraint. For cases where very few nodes need to be evaluated, this might be an advantage. However, in using the analysis, the use case rarely occurs. For most constraints, like Role-based Access Control (RBAC), the node and data characteristic labels need to be evaluated at each node.

Due to the aforementioned reasons, we chose to implement the data flow analysis in Java and made the analysis more extensible as a central part of our framework. In this section, we first provide a general overview of the architecture of the analysis and provide a more detailed technical description of the extraction of data flows into flow graphs, label propagation, and constraint definition.

#### 4.1 Architecture Overview

Our data flow analysis follows the general architecture of the Prolog-based data flow analysis of Seifermann et al. [30]. Figure 5 shows the analysis steps and their sequential order as an activity diagram. Initially, the input models are loaded and references between model elements are resolved. This is done automatically by the Eclipse Modeling Framework (EMF). Using the information from the models and annotations, we extract a set of *transpose flow graphs (TFGs)* that each represents one unambiguous flow of data to a data sink in the modeled software architecture, i.e., the transpose rooted directed graph, where the root is a single data sink. The extraction starts at each identified data sink and follows the modeled flow of data in the opposite direction. Afterward, we transpose the graph to represent data flows between the vertices of the graph, so each TFG connects one or multiple data sources with a single data sink. Each vertex represents one individual data processing step. If the analysis encounters an ambiguity in the data flow of the current element, it is resolved by creating copies of the current TFG, for each of the possible flows. After all TFGs are extracted, we first evaluate the node characteristic labels of the vertices. Afterward, we propagate the data characteristic labels along the edges of the TFGs. Starting with the sink vertices, we calculate the data characteristics flowing into the current vertex by recursively evaluating the behaviors of the previous vertices in the TFG and tracing back the results. How sinks are identified, how characteristic labels and vertex behavior are specified, and how they are evaluated, is specific to the input model type, e.g., DFD or PCM.

Using a set of fully propagated TFGs, data flow constraints can be checked. For example, by comparing propagated data characteristics with specified node characteristics, as described in Section 3.



**Fig. 5.** Analysis architecture as performed key activities.



## 4.2 Flow Graph Extraction

We specify extraction logic for creating TFGs and specific subclasses of *vertex* for each element that represents a data flow node in DFDs and PCM models.

For the DFDs described in Subsection 3.1, sinks are nodes that either have no outgoing flows or nodes whose assignments for an output pin are independent of all its input pins. Starting with these, the analysis performs a depth-first search over the DFD and creates vertices for each node. Ambiguities in the data flow exist if two or more flows point to the same input pin. To resolve the ambiguity, the analysis creates a copy of the current TFG for each path to the pin.

In the PCM, sinks are represented by the last element in usage scenarios. As the information regarding data flows is distributed across all PCM models, the analysis has to iterate over them and resolve relationships between elements. The analysis creates a vertex for all elements that can be annotated with a node characteristic label, that specifies data flow behavior, or that joins the control flow after a branch. The latter also creates a new TFG. Calls to Service Effect Specifications (SEFFs) that are defined in interfaces are also handled separately: For each call of these SEFFs, a *calling* and *returning* vertex is created, which enclose the data flows, i.e., the vertices, which make up the SEFF internally.

## 4.3 Label Propagation

We individually propagate the characteristic labels for each TFG. Each vertex references the input model element it represents and contains all logic regarding the calculation of node and data characteristic labels. First, we calculate the node characteristic labels and store them in the corresponding vertex. Starting from the sink of the TFG, we calculate the data characteristic labels that represent the output of the vertex and also store them in the corresponding vertex. This is achieved by recursively calling the calculation logic of all previous vertices and using the hereby calculated output labels as input. Note, that we do not consider cycles in the propagation logic because TFGs represent DAGs.

For our DFD metamodel, the calculation of node characteristic labels is trivial, as DFD nodes already contain these labels, and vertices directly represent nodes. During the calculation of data characteristic labels, each vertex first recursively evaluates its input, as described above, iterating the DFD nodes that are connected by input pins. After the input has been evaluated the labels are aggregated and saved as output of the corresponding output pin.

In PCM, node characteristic labels are directly annotated to PCM elements like resource containers or usage scenarios. For the calculation of node characteristic labels, the vertex iterates over the relationships of the PCM element it represents and stores the annotated characteristic labels relevant to the vertex. For the calculation of data characteristic labels, the PCM-specific vertices use the output of the previous vertex in the TFG. In contrast to our DFDs, the PCM does not support the definition of multiple individual data flows between two nodes that each represents a separate flowing data variable. Rather, one flow between two vertices in the TFG encapsulates all data flowing between two nodes. To

evaluate the input, the vertices filter the variables with their data characteristics to only include variables that are in the scope of the element represented by the vertex. To calculate the output data characteristic labels, the vertex evaluates stochastic expressions that are used in the PCM to define propagation behavior.

#### 4.4 DSL-based Constraint Checking

To help in the specification and checking of constraints and queries, we define a simple domain-specific language (DSL). We follow the general structure of the DSL by Hahner et al. [15], which was defined for the original analysis of Seifermann et al. [30] but simplify the approach by implementing it in Java and fitting it to our new implementation of the analysis.

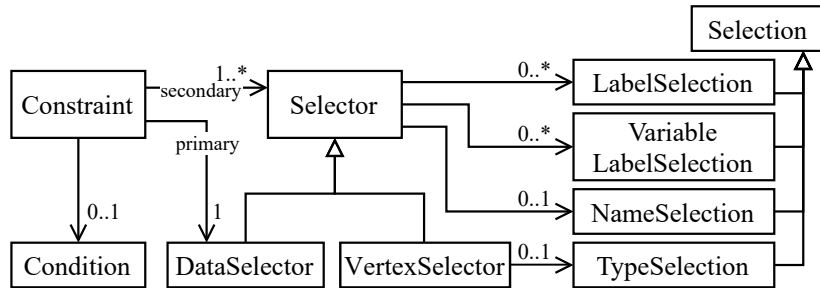


Fig. 6. Metamodel showing the abstract syntax of the DSL for the data flow analysis.

Figure 6 shows the abstract syntax of our DSL. A *Constraint* is made up of primary and secondary selectors, as well as an optional condition. *Selectors* are either specific for data or vertices. *VertexSelectors* match the properties of the vertices themselves, while *DataSelectors* match the propagated data characteristic labels of each vertex. They contain a set of *Selections* that each represent a property. A *Selection* can for example define a characteristic label or the name of a vertex or data. The *VariableLabelSelection* does not reference a specific label but defines a variable that contains all labels of a given label type that are present at either the vertex or data. These variables can be compared in the *Condition* of the constraint using set theory. Executing the constraint searches all TFGs in the modeled software architecture using the flow graph extraction, propagates all labels, and tests each vertex. The selectors return all vertices in a TFG that match the properties defined by its selections. Constraints define a *never flows* relationship between the primary *DataSelector* and secondary selectors. The results of the primary and secondary selectors represent violations. If a condition is defined, it is evaluated in addition. In this case, the violations are the results of primary and secondary selectors, for which the condition evaluates to true.

Listing 1 demonstrates the concrete syntax of our DSL for the online shop example from Subsection 3.2. We provide a *builder* to set up the analysis with required inputs, which is simplified in line 1. We define a constraint using our

---

```

1  var analysis = new DataFlowAnalysisBuilder().build(); // simplified
2
3  var constraint = new Constraint()
4                    .ofData()
5                    .withLabel("Sensitivity", "Personal")
6                    .withoutLabel("Encryption", "Encrypted")
7                    .neverFlows()
8                    .toVertex()
9                    .withLabel("Location", "offPremise")
10                   .create();
11
12 var violations = constraint.execute(analysis);

```

---

**Listing 1.** Code snippet showing a DSL constraint for a simplified online shop.

DSL, starting in line 3. For our example, we specify that *personal data* (line 5) that is *not encrypted* (line 6) should *never flow* to vertices that are *off-premise* of the online shop (line 9). We execute the constraint in line 12. After the execution, the variable *violations* contains a list of all constraint-violating vertices within the modeled software architecture. If no violation has been found, the list is empty.

## 5 Analysis Framework Extensions

We demonstrate the extensibility of our framework with several related work [5, 6, 8, 9, 12, 14, 26, 35] that is either compatible to or already using our approach.

Boltz et al. [5, 6] showcase the extension of both modeling and analysis for data protection and privacy. As shown in Figure 1, they provide a GDPR metamodel and transformations from PCM and to and from our DFD metamodel. Regarding the consideration of uncertainty within the software architectural design and system environment, multiple black-box and white-box extensions exist. Walter et al. [35] use the data flow analysis as black-box together with PerOpteryx [20] for design space exploration regarding confidentiality under structural uncertainty. Other white-box extensions analyze access control under uncertainty [8] or trace confidentiality violations to related uncertainty sources [14]. Our framework is also used in an uncertainty impact analysis [12] that predicts the impact of uncertainty on confidentiality based on the extracted data flows and a classification of uncertainty regarding confidentiality [13].

## 6 Evaluation

In our evaluation, we compare our new Java-based analysis to the Prolog-based analysis of Seifermann et al. [30]. The primary goals of this evaluation were to assess the accuracy and scalability of both analyses and to show that our Java-based analysis not only maintains the core functionalities of the Prolog-based analysis but also improves execution times and resource efficiency. Due to the lack of support for our new DFD metamodel in the Prolog-based analysis, our evaluation only focuses on PCM model instances. The evaluation of our analysis with a focus on the DFD metamodel or the extensions from Section 5 are considered potential future work.

## 6.1 Evaluation Design

To compare accuracy, we check whether both analyses correctly identify violations across various case study-based PCM models. To ensure a good base for comparison, we utilize the same case study-based models employed by Seifermann et al. [30] for evaluating the accuracy of the Prolog-based approach. The selected case studies use the default call return semantics of the current stable PCM version. We executed both analyses with semantically equivalent constraint queries, using the count of accurately identified violations as the evaluation metric.

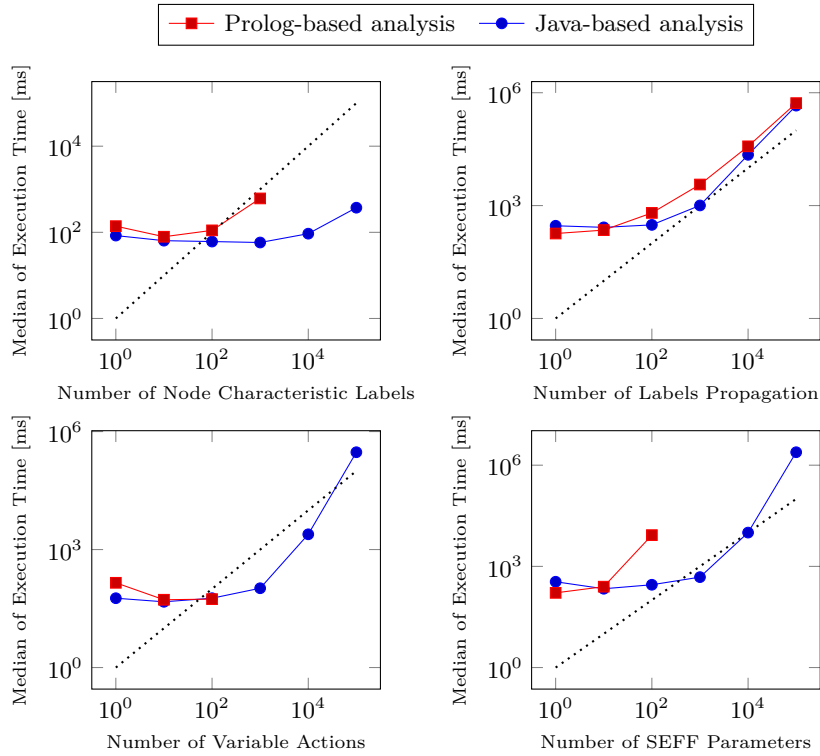
To examine and compare scalability, we measured the full execution time of both analyses while analyzing models of increasing size. To isolate the impact of distinct model features on scalability, we generated individual minimal models incrementally increasing the number of node characteristic labels, characteristic label propagations, variable actions, or SEFF parameters. We chose these elements, as they have the highest impact on either the length of Prolog code or Java loop iterations, depending on the analysis. Each analysis was executed with a constraint designed to detect a violation at each node, thus ensuring a worst-case execution time scenario for both analyses. For each run, we increase the model feature under consideration by the power of ten, starting at  $10^0$  and ending with  $10^5$ . We conducted each test 10 times and calculated the median execution time to mitigate outliers or measurement anomalies. The analyses were performed on a dedicated VM equipped with 4 AMD Opteron 8435 cores, 97 GB RAM, running Debian 11 with OpenJDK 11/17.

## 6.2 Evaluation Results

In terms of accuracy, both analyses successfully identified the 42 violations present in the case study-based models without returning any false positives. Table 1 shows the results of the accuracy evaluation and size of analyzed models. As both analyses performed the same, we assume, that our Java-based analysis is functionally equivalent to the Prolog-based analysis, when analyzing models using the call return semantics of the PCM.

Case Study	Prolog-based	Java-based	Components Labels	
ContactSMS [18]	10 violations	10 violations	3	4
FlightControl [30]	0 violations	0 violations	6	6
FriendMap [34]	0 violations	0 violations	5	12
Hospital [34]	0 violations	0 violations	4	12
ImageSharing [30]	0 violations	0 violations	1	9
PrivateTaxi [18]	0 violations	0 violations	13	20
TravelPlanner [18]	32 violations	32 violations	7	8
WebRTC [34]	0 violations	0 violations	20	12

**Table 1.** Accuracy results of both analyses compared and size of the models.



**Fig. 7.** Scalability results of the Prolog-based analysis and the Java-based analysis.

Regarding scalability, we plotted the results of both analyses as line graphs for each examined model feature, shown in Figure 7. Each graph contains data points from both analyses—the Prolog-based analysis (in red) and the Java-based analysis (in blue). Both axes are scaled logarithmically, with the x-axis showing the increasing number of model elements and the y-axis the median execution times in milliseconds. Our evaluation showed that the Prolog-based analysis fails to complete a run for more than 1000 node characteristic labels or 100 for variable actions and SEFF parameters, due to high memory demand (see Section 4). In our tests, the analysis ran in *out of memory errors* or crashed, despite the substantial *97 GB* of available memory. Regarding execution time behavior, while the Prolog-based analysis displayed an exponential increase in execution times or incomplete analysis runs for larger models, our Java-based analysis maintained nearly constant execution times up to  $10^3$  elements for most evaluated cases. When increasing the number of label propagations, the execution time behavior of both analyses is similar. The exponential increase in execution time of the Java-based analysis for larger models can be explained due to inefficiencies in TFG finding, and overhead during label propagation.

Overall, despite the noted increase in execution times for larger models in the Java-based analysis, we consider the time required in all scenarios feasible

for design-time analyses. Our Java-based analysis, compared to the Prolog-based analysis, offers more manageable execution times and the capability to analyze large models, rendering it more suitable for real-world systems.

As parts of our evaluation are based around artificial scenarios and case studies, we discuss the external, internal, and construct validity, as well as reliability of our evaluation, as characterized by Runeson et al. [22]. Our main threat to external validity is the limited generalizability due to the case study-based evaluation. We try to mitigate this threat by using well-known case studies from literature to evaluate and compare accuracy. For the evaluation of scalability, the models were programmatically generated to only scale and focus on individual aspects of the models and analysis. A threat to the internal validity of our evaluation of scalability is that, due to the use of different technologies in both analyses, it was not possible to use the exact same constraints. We mitigate this threat by defining semantically equivalent constraints that find a violation at each node. Our main threat to construct validity of our scalability evaluation is that it does not comprehensively cover all aspects that influence the execution time. We cannot fully mitigate this threat but have chosen the examined aspects based on the execution logic of both analyses and a previous scalability evaluation of Seifermann et al. [30]. To mitigate threats regarding the reliability of our evaluation and to address the lack of replication packages in software architecture research [19], we have published a data set [7]. The dataset contains all raw and compiled code artifacts, as well as an Eclipse-based product that already includes the plugins that make up the framework. The product can be used to model DFD or PCM instances and analyze them using our data flow analysis. We also include the raw results of our scalability evaluation and the used case study models.

## 7 Conclusion and Future Work

In this paper, we have presented our open and extensible framework for data flow analysis. We have introduced a unified DFD metamodel as a primary software architecture modeling artifact and input for our data flow analysis framework. We have described means that we provide to manually define DFDs as well as automatically derive them from the architecture description language PCM and other third-party representations.

Based on the approach of Seifermann et al. [30], we have implemented a Java-based data flow analysis. We described the general architecture of the analysis and provided detailed technical descriptions of the core features. For the analysis, we have defined an extensible intermediate representation of data flows, called transpose flow graphs. We have described how data flows are extracted from input models and how characteristic labels are propagated using our new intermediate representation. To enable the definition of data flow constraints for the analysis, we have defined a new domain-specific language.

We highlight the problems of the Prolog-based analysis of Seifermann et al. [30] and show in our evaluation, that our Java-based analysis is functionally equivalent to the Prolog-based analysis and can analyze larger system models.

In future work, we aim to further enhance the tooling that makes up our framework. We also aim to further work on the various framework extensions, like the data protection [6] and uncertainty analyses [14] and include more cooperation points of our framework, e.g., with continuous security analysis [26]. Lastly, we aim to comprehensively evaluate the overall approach of our framework.

### Acknowledgements

This publication is partially based on the research project SofDCar (19S21002), which is funded by the German Federal Ministry for Economic Affairs and Climate Action. This work was also supported by funding from the topic Engineering Secure Systems of the Helmholtz Association (HGF) and by KASTEL Security Research Labs, the BMBF (German Federal Ministry of Education and Research) grant number 16KISA086 (ANYMOS), and the NextGenerationEU project by the European Union (EU). We like to thank Felix Schwickerath, Tom Hüller, Daniel Huber, Tizian Bitschi, Anne-Kathrin Hermann, and Nils Niehues for their support in the development of the presented work.

### References

1. Ahrendt, W., *et al.*: Deductive software verification—the key book. Springer (2016)
2. Alshareef, H., *et al.*: Precise Analysis of Purpose Limitation in Data Flow Diagrams. In: ARES, pp. 1–11 (2022)
3. Bernsmed, K., *et al.*: Adopting threat modelling in agile software development projects. *Journal of Systems and Software* 183, 111090 (2022)
4. Beverley-Smith, H., Perowne, C.H., and Kelleher, F.: British Airways Faces Significantly Reduced £20M Fine for GDPR Breach. *The National Law Review*, [www.natlawreview.com/article/british-airways-faces-significantly-reduced-20m-fine-gdpr-breach](http://www.natlawreview.com/article/british-airways-faces-significantly-reduced-20m-fine-gdpr-breach) (visited on 12/11/2023)
5. Boltz, N., Walter, M., and Gerking, C.: Designing Automotive Case Studies for Architectural Security Analyses. (2023)
6. Boltz, N., *et al.*: A Model-Based Framework for Simplified Collaboration of Legal and Software Experts in Data Protection Assessments. *INFORMATIK 2022* (2022)
7. Boltz, N., *et al.*: *Dataset: An Extensible Framework for Architecture-Based Data Flow Analysis for Information Security*. Zenodo, 2024. DOI: 10.5281/zenodo.10794265.
8. Boltz, N., *et al.*: Handling environmental uncertainty in design time access control analysis. In: SEAA, pp. 382–389 (2022)
9. De Moor, O., *et al.*: “.QL: Object-Oriented Queries Made Easy”. In: International Summer School on Generative and Transformational Techniques in Software Engineering. 2008, pp. 78–133.
10. DeMarco, T.: Structure analysis and system specification. In: *Pioneers and Their Contributions to Software Engineering*, pp. 255–288 (1979)
11. General Data Protection Regulation (GDPR), (2016). <https://gdpr.eu/tag/gdpr/>
12. Hahner, S., Heinrich, R., and Reussner, R.: Architecture-Based Uncertainty Impact Analysis to Ensure Confidentiality. In: SEAMS, pp. 126–132 (2023)
13. Hahner, S., *et al.*: A Classification of Software-Architectural Uncertainty regarding Confidentiality. In: ICETE, pp. 139–160 (2023)
14. Hahner, S., *et al.*: Model-based Confidentiality Analysis under Uncertainty. In: ICSA-C, pp. 256–263 (2023)

15. Hahner, S., *et al.*: Modeling Data Flow Constraints for Design-Time Confidentiality Analyses. In: ICSA-C, pp. 15–21 (2021)
16. HmbBfDI: 35.3 Million Euro Fine for Data Protection Violations in H&M’s Service Center, [www.datenschutz-hamburg.de/fileadmin/user\\_upload/HmbBfDI/Presse\\_mitteilungen/2020/2020-10-01-H\\_M.pdf](http://www.datenschutz-hamburg.de/fileadmin/user_upload/HmbBfDI/Presse_mitteilungen/2020/2020-10-01-H_M.pdf) (visited on 12/11/2023)
17. International Organization for Standardization: Information technology — Security techniques — Information security management systems — Overview and vocabulary. Standard ISO/IEC 27000:2018,
18. Katkalov, K.: Ein modellgetriebener Ansatz zur Entwicklung informationsflusssicherer Systeme. doctoralthesis, Universität Augsburg (2017).
19. Konersmann, M., *et al.*: Evaluation Methods and Replicability of Software Architecture Research Objects. In: ICSA, pp. 157–168 (2022)
20. Koziolok, A., Koziolok, H., and Reussner, R.: PerOpteryx: automated application of tactics in multi-objective software architecture optimization. In: QoSA-ISARCS, pp. 33–42 (2011)
21. Reussner, R., *et al.*: Modeling and Simulating Software Architectures – The Palladio Approach. MIT Press, Cambridge, MA (2016). ISBN: 9780262034760
22. Runeson, P., *et al.*: Case Study Research in Software Engineering: Guidelines and Examples. John Wiley & Sons (2012)
23. Schneider, S., and Scandariato, R.: Automatic extraction of security-rich dataflow diagrams for microservice applications written in Java. *Journal of Systems and Software* 202, 111722 (2023)
24. Schneider, S., *et al.*: How Dataflow Diagrams Impact Software Security Analysis: an Empirical Experiment. In: SANER (2024)
25. Schneider, S., *et al.*: microSecEnD: A Dataset of Security-Enriched Dataflow Diagrams for Microservice Applications. In: MSR, pp. 125–129 (2023)
26. Schulz, S., *et al.*: Continuous Secure Software Development and Analysis. In: SSP (2021)
27. Schwickerath, F., *et al.*: Tool-Supported Architecture-Based Data Flow Analysis for Confidentiality. In: arXiv preprint (2023). DOI: 10.48550/arXiv.2308.01645
28. Seifermann, S., *et al.*: A unified model to detect information flow and access control violations in software architectures. In: SECRYPT, pp. 26–37 (2021)
29. Seifermann, S., Heinrich, R., and Reussner, R.: Data-Driven Software Architecture for Analyzing Confidentiality. In: ICSA, pp. 1–10 (2019)
30. Seifermann, S., *et al.*: Detecting violations of access control and information flow policies in data flow diagrams. *Journal of Systems and Software* 184, 111138 (2022)
31. Shull, F., *et al.*: What we have learned about fighting defects. In: METRICS, pp. 249–258 (2002)
32. Sion, L., *et al.*: Solution-aware data flow diagrams for security threat modeling. In: SAC, pp. 1425–1432 (2018)
33. Snelting, G., *et al.*: Checking probabilistic noninterference using JOANA. *it - Information Technology* 56(6), 280–287 (2014)
34. Tuma, K., Scandariato, R., and Balliu, M.: Flaws in Flows: Unveiling Design Flaws via Information Flow Analysis. In: ICSA, pp. 191–200 (2019)
35. Walter, M., *et al.*: Architectural optimization for confidentiality under structural uncertainty. In: ECSA, pp. 309–332 (2021)
36. Weyns, D., *et al.*: Towards a Research Agenda for Understanding and Managing Uncertainty in Self-Adaptive Systems. *SIGSOFT Software Engineering Notes* 48(4), 20–36 (2023)