



Serializing java objects in plain code

Julian Wachter^a, Deepika Tiwari^{b,*}, Martin Monperrus^b, Benoit Baudry^b

^a Karlsruhe Institute of Technology, Germany

^b KTH Royal Institute of Technology, Sweden

ARTICLE INFO

Editor: Lingxiao Jiang

Keywords:

Code
Serialization
Objects on disk
Runtime
Java

ABSTRACT

In managed languages, serialization of objects is typically done in bespoke binary formats such as Protobuf, or markup languages such as XML or JSON. The major limitation of these formats is readability. Human developers cannot read binary code, and in most cases, suffer from the syntax of XML or JSON. This is a major issue when objects are meant to be embedded and read in source code, such as in test cases.

To address this problem, we propose plain-code serialization. Our core idea is to serialize objects observed at runtime in the native syntax of a programming language. We realize this vision in the context of Java, and demonstrate a prototype which serializes Java objects to Java source code. The resulting source faithfully reconstructs the objects seen at runtime. Our prototype is called PRODJand and is publicly available. We experiment with PRODJto successfully plain-code serialize 174,699 objects observed during the execution of 4 open-source Java applications. Our performance measurement shows that the performance impact is not noticeable. Through a user study, we demonstrate that developers prefer plain-code serialized objects within automatically generated tests over their representations as XML or JSON.

1. Introduction

Serialization consists in representing runtime data in a format that can be transmitted (Godefroid et al., 2020) or stored (Liu et al., 2020). Developers serialize objects either through native programming language support, or with third-party libraries that support serialization of objects in different formats (Viotti and Kinderkhedra, 2022). The serialization format may be binary, such as Protocol Buffers (Google, 2023; Eddebuettel et al., 2014), or use a markup language such as JSON (Harand et al., 2021) or XML (Bray et al., 1997).

Serialization has many use cases, such as sharing data among services that are external (Neumann et al., 2018) or internal (Pourhabibi et al., 2020) to a software system. Serialized objects can also be used within tests, to reproduce failures (Artzi et al., 2008), or to initialize test inputs to a target state (Kim et al., 2022; Tiwari et al., 2022; Alshahwan et al., 2024).

The problem space of object serialization has three dimensions: correctness, performance, and readability. A few studies have focused on the correctness of serialization techniques (Viotti and Kinderkhedra, 2022), including assessing the behavior of serialization tools through formal proofs (Ye and Delaware, 2019), or back-to-back testing (Harand et al., 2021). Performance is a major challenge for serialization, and several works have focused on optimizing it in the context of wireless communication (Wong et al., 2004), microservices data exchange

(Pourhabibi et al., 2020), or efficient test execution (Gligoric et al., 2011).

However, readability is an aspect of serialization that has been neglected in previous work. Producing readable serialized objects is the problem we address in this paper. This is of utmost importance when serialized objects are consumed by humans. For example, when serialized objects are included in test cases as part of test inputs (Artzi et al., 2008; Kim et al., 2022; Tiwari et al., 2022; Alshahwan et al., 2024), readability is essential to let developers understand the tests. Clearly, mainstream serialization techniques fall short with respect to readability. For example, Listing 1 shows an example of a test case generated by the automated test generation tool Pankti (Tiwari et al., 2022), a tool that creates tests from objects at runtime. In Listing 1, the test inputs are in the XML format, as is the case for the data used to populate the receiver and key variables. XML strings arguably makes the code hard to read, making it challenging to understand the test's intention.

To solve the problem of unreadable serialized data, such as JSON/XML test fixtures, we propose the novel concept of *plain-code serialization*, which consists in serializing an object directly into code statements. By serializing objects to plain code, the serialized representation becomes readable, using a syntax that is both concise and natural to developers. We illustrate this with the semantically equivalent test of Listing 2, which contains the same serialized objects as Listing 1 represented as plain code, instead of XML. In addition, the compiler gives us

* Corresponding author.

E-mail address: deepikat@kth.se (D. Tiwari).

```

1 @Test
2 public void testGetNameAsString() {
3     // Arrange
4     COSDictionary receiver = deserializeFromXML("""
5     <org.apache.pdfbox.cos.COSDictionary>
6     <direct>false</direct>
7     <needToBeUpdated>false</needToBeUpdated>
8     <items class="\ org.apache.pdfbox.util.SmallMap\ ">
9     <mapArr>
10    <org.apache.pdfbox.cos.COSName>
11    <direct>false</direct>
12    <name>BaseFont</name>
13    <hashCode>-1657267488</hashCode>
14    </org.apache.pdfbox.cos.COSName>
15    <org.apache.pdfbox.cos.COSName>
16    <direct>false</direct>
17    <name>Times-Roman</name>
18    <hashCode>1058166902</hashCode>
19    </org.apache.pdfbox.cos.COSName>
20    </mapArr>
21    </items>
22 </org.apache.pdfbox.cos.COSDictionary>
23 """);
24 COSName key = deserializeFromXML("""
25 <org.apache.pdfbox.cos.COSName>
26 <direct>false</direct>
27 <name>BaseFont</name>
28 <hashCode>-1657267488</hashCode>
29 </org.apache.pdfbox.cos.COSName>
30 """);
31 // Act
32 String actual = receiver.getNameAsString(key);
33 // Assert
34 assertThat(actual).isEqualTo("Times-Roman");
35 }

```

Listing 1. Serialized objects in XML are not meant for human consumption, they are not readable. In this paper, we aim to serialize objects in plain code instead, as demonstrated later in Listing 2 and Listing 9

```

1 @Test
2 @DisplayName("getNameAsString(org.apache.pdfbox.cos.COSName)")
3 public void testGetNameAsString_ProDj() {
4     // Arrange
5     COSDictionary receiver = this.createCOSDictionary();
6     COSName key = COSName.BASE_FONT;
7     // Act
8     String actual = receiver.getNameAsString(key);
9     // Assert
10    assertThat(actual).isEqualTo("Times-Roman");
11 }
12
13 COSDictionary createCOSDictionary1() {
14     // Recreated from trace
15     COSDictionary receiver = new COSDictionary();
16     COSName key = COSName.TYPE;
17     COSBase value = COSName.FONT;
18     receiver.setItem(key, value);
19     COSName key0 = COSName.SUBTYPE;
20     COSBase value0 = COSName.TYPE1;
21     receiver.setItem(key0, value0);
22     COSName key1 = COSName.BASE_FONT;
23     receiver.setName(key1, "Times-Roman");
24     COSName key2 = COSName.ENCODING;
25     COSBase value1 = COSName.WIN_ANSI_ENCODING;
26     receiver.setItem(key2, value1);
27     return receiver;
28 }

```

Listing 2. Automatically generated test containing plain-code serialized runtime objects `receiver` and `key`. Compared to the XML object representations in Listing 1, this test is arguably more readable and maintainable.

strong guarantees about the static correctness of the serialized objects. Deserialization then simply means executing the code. In the context of tests, this completely removes dirty data like XML strings, thus making the generated tests more readable and self-contained (Afshan et al., 2013).

In this paper, we propose two algorithms for plain-code serialization. First, our structure-based serialization strategy relies on static analysis of

the source code to identify a minimal sequence of constructor and setter calls, that are able to initialize an object to a desired state. Second, the trace-based serialization strategy records constructor and method invocations and subsequently generates a corresponding plain-code version that replays these invocations to construct and initialize an object. We prototype the concept in Java, in a tool called PRODJ that serializes JVM objects in plain code.

We demonstrate the feasibility of plain-code serialization by evaluating PRODJ with 4 well-known Java projects. PRODJ successfully serializes 174,699 Java objects observed at runtime, producing Java code that can reconstruct them. We find that the objects are faithfully reconstructed through a full-cycle of plain-code serialization and deserialization. Through a user study with 17 developers, we also demonstrate that the use of plain-code serialized objects contributes to the readability of automatically generated tests, relative to XML- or JSON-serialized objects.

We summarize our contributions as follows:

- The concept of plain-code serialization: producing source code statements that reconstruct an object to a target state, contrasting with binary (e.g. Protobuf) or textual format (e.g. JSON) serialization.
- A publicly available prototype implementation of PRODJ for future research on advanced serialization in the JVM.¹
- An evaluation of PRODJ on 4 notable open-source projects in Java, involving 174,699 serialized objects, assessing the feasibility of plain-code serialization.

PRODJ and our experimental data are available in our replication package.²

2. Plain-code serialization

In this section, we describe our key conceptual contribution: serializing runtime objects in plain code. This is done by monitoring an application as it executes, capturing objects, and translating them to plain Java code.

2.1. Concept & overview

Plain-code serialization means representing a runtime object with the concrete syntax of a programming language. In this context, *serialization* means producing code from an object, such as a sequence of statements, that can be compiled and executed. *Deserialization* means executing the plain-code representation of the runtime object, an operation we call in this paper *object reconstruction*.

The reconstruction operation for any object o must conform to the invariant presented in Eq. (1).

$$\forall o, reconstruct(o) \equiv o \quad (1)$$

$$reconstruct(x) = evaluate(plain_code_serialize(x)) \quad (2)$$

The equivalence relationship (\equiv) means that the object after reconstruction is in the same state as the one observed at runtime. While Equation 1 applies to any serialization library, Eq. (2) is specific to plain-code serialization, with *reconstruct* being the evaluation of the plain-code serialized representation of x obtained from the function *plain_code_serialize*.

Benefits. Plain code serialization has two major benefits: readability and correctness checking. First, the serialized objects are readable for developers, using the programming language syntax for which they have been trained for. Second, when serializing to a statically typed language, plain-code serialization provides guarantees about the serialized data. In this case, the compiler checks for the validity of the object with the full

¹ <https://github.com/ASSERT-KTH/prodj>

² <https://doi.org/10.5281/zenodo.7902538>

```

1 public class Monkey {
2     private int age;
3     public Color eyeColor;
4     public Habitat habitat;
5
6     public Monkey(int age) { this.age = age; }
7
8     public Monkey(int age, Color eyeColor, Habitat habitat) {
9         this.age = age;
10        this.eyeColor = eyeColor;
11        this.habitat = habitat;
12    }
13
14    public void setEyeColor(Color eyeColor) { this.eyeColor = eyeColor; }
15
16    public String eatBanana(int dozens) {
17        int quantity = dozens * 12;
18        return "Eating " + quantity + " bananas";
19    }
20 }
21
22 public enum Color { RED, GREEN, BLUE; }
23
24 public class Habitat {
25     private String coordinates;
26     private double area;
27
28     public Habitat(String coordinates) { this.coordinates = coordinates; }
29
30     public void grow(double newArea) {
31         double newValue = this.area + newArea;
32         this.area = newValue;
33     }
34 }

```

Listing 3. The example Monkey class to be serialized, and its associated types

power of the type system, capturing domain-specific correctness that one typically does not have with binary or markup serialization.

In this paper, we devise two complementary strategies to map runtime objects to plain code: structure-based serialization and trace-based serialization. For structure-based object serialization, we statically analyze a class to identify all possible ways for setting each field of the object, such as using constructors or setters. However, an object may not be fully reconstructible using the structure-based strategy. This is due to the presence of internal, mutable state in object-oriented programs that are not exposed through setters. The second strategy, trace-based object serialization, overcomes this limitation. It is based on recording a sequence of events that modify the state of an object at runtime. This recorded sequence of constructors and mutating methods, is then used to reconstruct the object to its target state.

We now illustrate how these two strategies work and complement each other. Consider the Monkey and Habitat classes in Listing 3. Monkey is serializable using the structure-based strategy. The three possible reconstruction plans for the Monkey class are presented in Listing 4. To create the Java code for a given Monkey object, the $\$age\$, \$eyeColor\$, and \$habitat\$$ meta-variables are replaced with actual values from runtime. However, note that the Habitat class in Listing 3 contains the private field `area`. This field is part of the internal state of Habitat and, according to the information hiding principle, not directly exposed. Therefore, we use trace-based object serialization for Habitat. For example, recording the constructor invocation of Habitat, as well as an invocation of its `grow` method would allow us to reconstruct a Habitat object with coordinates "42, 42" and an area of 2.0. In the following subsections, we present the structure-based and trace-based object reconstruction strategies in detail.

2.2. Structure-based serialization (SB)

The core idea of the structure-based object serialization strategy is to statically analyze a class and build a reconstruction plan. This is done by searching the class for statement which set the value of fields, such as constructor or setter method calls. The plan is instantiated at runtime for a given object.

2.2.1. Synthesizing reconstruction plans

A **reconstruction plan** for a class is a template with actions that express how an object of the class is instantiated. The plan contains meta-variables which are populated with the values observed at runtime.

```

1 // Reconstruction plan 1
2 Monkey monkey = new Monkey($age$, $eyeColor$, $habitat$);
3
4 // Reconstruction plan 2
5 Monkey monkey = new Monkey($age$);
6 monkey.setEyeColor($eyeColor$);
7 monkey.habitat = $habitat$;
8
9 // Reconstruction plan 3
10 Monkey monkey = new Monkey($age$);
11 monkey.eyeColor = $eyeColor$;
12 monkey.habitat = $habitat$;

```

Listing 4. Three possible reconstruction plans for Monkey

When instantiated with a captured runtime object, the reconstruction plan is a complete, executable code snippet. Eq. (3) presents the formulation of object reconstruction with structure-based serialization, as a specialization of idiomatic object reconstruction presented in Eq. (1).

$$evaluate(instantiated_reconstruction_plan(o)) \equiv o \quad (3)$$

An **action** within a reconstruction plan is a statement that sets zero or more fields within an object. The 5 action types used for structure-based object reconstruction are denoted as SB in Table 1. Each action can be constructing (α) or field-setting (β_f sets a field f). Constructor calls and calls to factory methods are constructing actions, whereas field assignments and setter calls are field-setting actions. We analyze a target class to extract all possible actions that can be combined in different ways to automatically synthesize reconstruction plans. Each action has one or more meta-variables. A **meta-variable** is a placeholder for a field that is being set through the action. When a plan is instantiated to serialize a runtime object, each meta-variable is replaced with the actual value of the field.

Let us illustrate the generation of a reconstruction plan using the Monkey class in Listing 3. A Monkey has an integer value for `age`, and an `eyeColor`, which takes possible values from an enum. The Monkey also has a `Habitat`. Each field of Monkey has a corresponding meta-variable, namely $\$age\$, \$eyeColor\$, and \$habitat\$. There are three possible reconstruction plans for Monkey, which we present in Listing 4. The most idiomatic way to create a Monkey is using the first plan, which contains a single constructing action. It creates a Monkey by calling the three-argument `age`, `color`, and `habitat` constructor. The second and third plans consist of one constructing and two field-setting actions. The second plan calls the `age` constructor, followed by the `eyeColor` setter, and finally directly assigns the `habitat` field. The last plan has actions to call the `age` constructor, and directly assign both the `eyeColor` and `habitat` fields.$

2.2.2. Optimizing readability

As is the case for Monkey, we enumerate all possible reconstruction plans for a target class. However, we are interested in the shortest and most idiomatic plan, to reconstruct an object using the minimum number of statements (Bach et al., 2020). To quantify how idiomatic a reconstruction plan is, we first assign each action a pre-defined integer cost value. For example, calling a constructor has a very low cost, while directly setting a field has a higher cost. We then select the reconstruction plan with the lowest cumulative cost. We formulate this optimization as a constraint problem, shown in Equation 4. In the constructed problem, every action i is assigned a boolean selection variable a_i . If the model value for a_i is true, the action is selected. If it is false, the action will not be part of the reconstruction plan. If an action sets a field f , it is added to the set β_f , and if it is a constructing action, also to α .

The cost of an action i is denoted by $cost_i$, and the global optimization objective is modeled with Eq. (4a). Eq. (4b) ensures that at least one constructing action is selected; Eq. (4c) ensures that there is a single constructing action; and Eq. (4d) means that all fields are assigned at least once. If we find a solution to the constraint problem, all used

Table 1

The different types of actions for plain-code serialization, and the categories they belong to. α denotes the set of construct-ing actions, β_f the set of actions that set field f . SB refers to structure-based object reconstruction, TB to trace-based object reconstruction.

Action Type	Category	Explanation	Used by
call constructor	α	constructs an instance using a constructor, potentially with parameters	SB, TB
call factory method	α	creates an instance by calling a factory method	SB
use enum constant	α	uses an enum instance identified by its name	SB
call method	β_f	sets field f by calling a method with parameters	SB, TB
assign field	β_f	sets field f by directly assigning it a value	SB
use standard charset	α	creates a Java charset by using one of the constants in <code>StandardCharsets</code>	TB
use static field	α	uses a public static field whose value is reference equal to the object under construction	TB
use object reference	α	saves the object id, the object is then reconstructed using the strategy in Section 2.3	TB

```

1 // Serialized using the "use enum constant" action type
2 Color color = Color.BLUE;
3 // Serialized using the "use object reference" action type
4 Habitat habitat = /* trace based, see subsection 2.3.3 */;
5 // Serialized using the "call constructor" action type
6 Monkey monkey = new Monkey(6, color, habitat);

```

Listing 5. Structure-based plain-code serialization of a `Monkey` from Listing 3

actions are deterministically combined to produce a reconstruction plan. The plan will be also applied repeatedly for recursive structures.

$$a \in \{0, 1\}^n \sum_{i=1}^n a_i \cdot cost_i \quad (4a)$$

$$a_1 \vee a_2 \vee \dots \vee a_i \quad a_1, \dots, a_i \in \alpha \quad (4b)$$

$$a_1 \wedge a_2 \quad a_1, a_2 \in \alpha, a_1 \neq a_2 \quad (4c)$$

$$a_1 \vee a_2 \vee \dots \vee a_i \quad f \in fields, a_1, \dots, a_i \in \beta_f \quad (4d)$$

2.2.3. Instantiating reconstruction plans

Our goal is to plain-code serialize objects that are observed at runtime. We instrument the specific points in the program where we want to serialize an object. This instrumentation adds instructions that capture values at runtime, which are used to set the meta-variables within the reconstruction plan of the object. At runtime, when serialization is requested, the monitoring instructions are executed, resulting in objects serialized in plain code via instantiated reconstruction plans. For example, we can add instructions at the beginning of a specific method in order to serialize its receiver and parameters (Tiwari et al., 2022).

Finally, the instantiated reconstruction plans are converted to a plain-code representation by iterating over the actions (Table 1) and converting them one-by-one. A “call constructor” action first plain-code serializes each argument of the constructor and then emits a constructor call statement. Similarly, the arguments to the “call factory method” and “call method” actions are plain-code serialized before emitting a method call statement. “Assign field” plain-code serializes the value of the field and emits a field assignment statement, while “use enum constant” emits a read to the matching constant. We illustrate this procedure in Listing 5 by reconstructing a `Monkey` using the first reconstruction plan from Listing 4, which is the most idiomatic. For our `Monkey`, this means emitting a code statement that calls the three-argument constructor. The first argument, `age`, is a primitive and is directly inserted. The other two arguments are recursively plain-code serialized, stored as variables and referenced in the generated constructor call. The “use enum constant” action is used when plain-code serializing the `eyeColor` of the `Monkey`. Lastly, the “use object reference” defers serialization to the next strategy, trace-based plain-code serialization, and is explained in Section 2.3. The `habitat` field of the `Monkey` is handled by this action type.

```

1 public void grow(double newArea) {
2   + emitMethodStartEvent("Habitat#grow(double)", this, newArea);
3   double newValue = this.area + newArea;
4   + emitFieldSetEvent(this, this.area, newValue);
5   this.area = newValue;
6   + emitMethodEndEvent("Habitat#grow(double)");
7 }

```

Listing 6. Instrumentation for monitoring events within the trace-based object reconstruction strategy, the injected code is highlighted in green.

2.3. Trace-based serialization (TB)

Trace-based serialization consists in monitoring an application at runtime, in order to record an ordered sequence of *events* happening on a given object. After the execution finishes, the event sequence is converted to actions per Table 1 and source code is generated from these actions.

2.3.1. Capturing events

An *event* of interest is the occurrence of an operation which mutates an object. The considered events are constructor calls (`ConstructEvent`), field assignments (`FieldSetEvent`), or impure method calls updating a field. Method call events are split into a start (`MethodStartEvent`) and an end event (`MethodEndEvent`), allowing other events to happen in-between. We capture all calls to methods with at least one field assignment.

We denote all potential events to reconstruct an object o as $event_sequence(o)$. Based on this notation, we specialize the general object reconstruction formulation, Eq. (1), for trace-based reconstruction as follows:

$$evaluate(event_sequence(o)) \equiv o \quad (5)$$

Given a type we want to serialize, we first identify all types needed for its reconstruction because they are used either as fields, constructor parameters, or method parameters. Then, we instrument the constructors, accessible methods, and field assignments for all these types. For example, in the `Habitat` class (Listing 3), we instrument the constructor on line 28, as well as the `grow` method on lines 30 – 33. We show the resulting instrumented code for the `grow` method in Listing 6. When the instrumented `grow` is invoked at runtime, it emits multiple events, ordered by a monotonically increasing counter representing logical time.

The first event, `MethodStartEvent`, marks the start of a method call, and contains all information necessary to convert it to a method call action: the receiver object, the fully qualified name of the invoked method, as well as all arguments to the method call. As method invocations can be nested, and methods can be recursive, we also track when a method invocation ends through the emission of a `MethodEndEvent` when the method returns. The start and end events share a common unique ID, which is used to pair them up. We also assign unique identifiers to each object. This ensures that all method invocations, constructor calls, and field assignments for an object can be associated. The object ID is used

```

1 // Constructor call action
2 Habitat habitat = new Habitat("59.34743,18.07378");
3 // Method call action
4 habitat.grow(42);

```

Listing 7. Trace-based plain-code serialization of a Habitat from Listing 3

whenever an object is referenced within an event. In addition to the object ID, every reference also stores the current logical time, as an object may have different states at different points in time. We also see on line 5 of Listing 6 that the `grow` method updates the value of the `area` field. This corresponds to a `FieldSetEvent` event, which is emitted whenever a field of an object is reassigned inside a method. We assume that the principle of encapsulation is followed, and fields are either all public for simple data classes, or not mutated from outside the class itself. Within the `FieldSetEvent`, we store the ID of the receiver object, as well as the new and old values of the field being updated. Using `FieldSetEvents`, we can determine whether a given method invocation mutated the receiver object, and would need to be replayed during object reconstruction. Besides capturing method invocations and field assignments, we also instrument all constructors. Whenever a constructor finishes executing, a `ConstructEvent` is emitted. This event contains all necessary information to create a constructor call action, including the fully qualified name of the invoked constructor, as well as all of its arguments. Additionally, the event also contains the ID of the newly created object, allowing later events to refer to it. Thus, exercising the instrumented program with a workload emits events. The resulting captured event sequence is persisted on disk.

2.3.2. Mapping events to actions

After execution has concluded, the persisted event sequence is analyzed by iterating over it and mapping each event to an action of Table 1 if appropriate. Whenever a `ConstructEvent` is found, the ID of the created object is saved and a “call constructor” action is created. The same action type is used by the structure based serialization, as shown in Table 1. Every field assignment or method call within the constructor is ignored, as they would already be executed when calling the constructor. Next, every pair of `MethodStartEvent` and `MethodEndEvent` is associated with the receiver and discarded if it did not mutate the receiver. A method invocation is considered mutating, if a `FieldSetEvent` occurs on the receiver, or a field within it, before the corresponding `MethodEndEvent`. As this requires each object to know exactly what its fields are at each point in time, the initial field values set by the constructor are stored in the `ConstructEvent` and updated using the encountered `FieldSetEvents`. A “call method” action (Table 1) is created for each mutating method invocation. After the whole event sequence is processed, every object should now be associated with at least one constructor call action and zero or more method call actions.

Whenever the plain-code serialization of an object is requested, we follow the same procedure as for structure-based serialization. We illustrate this by plain-code serializing an instance of the `Habitat` class from Listing 3. The resulting code is shown in Listing 7. First, the constructor call action is converted to code by plain-code serializing all constructor arguments. In our example, this means serializing the `coordinate` parameter of the constructor. Next, a constructor call is emitted and the newly created object is stored in a local variable. This is done on line 2 of Listing 7. Then, all method call actions up until the requested logical time are also converted to code by plain-code serializing the parameters and then calling the method on the created object. For our `Habitat` this is a single method call, `grow(42)`, on line 4.

2.3.3. Object reconstruction mixing

At runtime, one may plain-code serialize only with pure structure-based or with pure trace-based serialization. However, this is subopti-

```

1 // Color uses the structure-based strategy
2 Color color = Color.BLUE;
3 // Habitat uses the trace-based strategy
4 Habitat habitat = new Habitat("59.34743,18.07378");
5 habitat.grow(42);
6 // Monkey uses the structure-based strategy
7 Monkey monkey = new Monkey(6, color, habitat);

```

Listing 8. Plain code serialization of a Monkey, mixing structure- and trace-based serialization

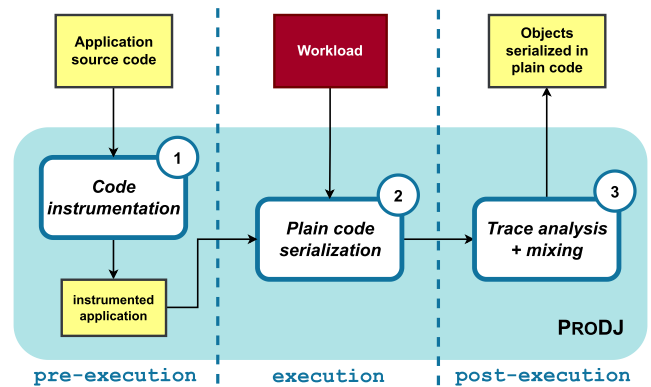


Fig 1. PRODJ uses plain-code serialization for capturing objects at runtime.

mal. In fact, a single object may be plain-code serialized using a combination of both structure- and trace-based strategies. We call this “object reconstruction mixing”.

For example, while the primary object might be serialized with the trace-based strategy, the arguments of the corresponding constructor or method call actions may be serialized using the structure-based strategy. Likewise, an object might be serialized using the structure-based strategy, but its fields may require trace-based serialization.

We show such a case in Listing 8, where we are constructing a `Monkey`. `Monkey` is serialized using structure-based serialization, and its `habitat` field is serialized using trace-based serialization. As the captured events are mapped to code offline, but the plain-code serialized `habitat` field has to be inserted into a specific spot in the generated code, the structure-based serialization strategy uses the “use object reference” (from Table 1) marker action. This action emits an `id@logical – time` reference, which is replaced by the trace-based serialized value after the event sequence conversion is finished. Besides this action, Table 1 also contains two more actions, “use static field” and “use standard charset”. These are actions that can not be done structurally, but also do not need the full trace-based serialization machinery. For the “use static field” action type, we look at all static fields of the class we want to serialize and detect if any of these store a reference-equal object. If such a field can be found, we emit a field read instead of tracing the constructor and subsequent mutating method calls. A similar, but more specific, optimization is realized by the “use standard charset” action type. This action type replaces references to Java’s standard charsets with the canonical named constants, and serves as an example of integrating domain logic into the serialization process. To further facilitate extensibility, such custom adapters and action types can be provided as plugins.

2.4. PRODJ: Plain-code serialization in java

We implement a prototype of plain-code serialization in the context of Java, in a tool called PRODJ. The PRODJ workflow is presented in Fig. 1. We now describe each of its phases, and also provide details on its implementation.

2.4.1. Pre-execution

Per Figure 1, the first phase within PRODJ occurs offline, before the application is launched. During this phase, PRODJ takes as input the source of a Java application, and identifies the types associated with each serialization point of interest. Next, PRODJ instruments each of these points for runtime monitoring, so that the corresponding objects can be plain-code serialized. For example, instructions can be added at the beginning of a method to plain-code serialize its receiver and parameter objects, including the objects within them. As detailed previously, PRODJ instruments associated types such that reconstruction plans are synthesized for types eligible for structure-based serialization, and events are captured for associated types that must be trace-based serialized. This first phase of PRODJ results in an instrumented version of the application, which includes additional instructions required for plain-code serialization.

2.4.2. Execution

The second phase of PRODJ occurs online during the execution of the Java application, as illustrated in Fig. 1. During this phase, the application instrumented by PRODJ is exercised with a workload. The workload may trigger the instrumentation instructions, and consequently, the plain-code serialization of the corresponding objects. If an object is serialized using the structure-based strategy, its instantiated reconstruction plan is embedded directly in the log. For an object serialized using the trace-based strategy, PRODJ logs the sequence of captured events. To summarize, as the application executes under a workload, PRODJ plain-code serializes the objects that are specified at specific instrumentation points.

2.4.3. Post-execution

The third phase of PRODJ occurs offline. During this phase, PRODJ parses the logs captured in the previous phase, i.e., at runtime. Recall that objects of interest can be serialized using both structure- and trace-based strategies. In the former case, PRODJ has already produced plain code instantiated reconstruction plans. In the latter case, PRODJ maps the captured event sequence to actions. In both cases, plain-code representations of associated objects are produced from the actions. Furthermore, PRODJ performs extra steps to enhance the readability and idiomacy of the plain-code objects. This includes de-duplicating redundant objects and outlining statements into helper methods to set up objects. The output of this phase of PRODJ is plain Java representations of objects seen in production.

2.4.4. Implementation

PRODJ relies on Spoon (Pawlak et al., 2015) for source code analysis, in order to identify the associated types of methods involving serialization within the target Java application. PRODJ uses Z3 (De Moura and Björner, 2008) to solve the constraint problem of Eq. (4) and determine the most idiomatic reconstruction plan for types that are structure-based serializable. The event sequences for trace-based serializable objects are logged in JSON. The instrumentation of methods and their associated types is achieved using ByteBuddy.³

3. Experimental methodology

This section describes how we empirically evaluate PRODJ, our prototype of plain code serialization for Java. First, in Section 3.1, we describe an application of PRODJ for the generation of tests from runtime. Next, Section 3.2 introduces the real-world projects we use as the subjects for this evaluation. This is followed in Section 3.3 by a description of the workloads we use with each project for our experiments. Finally, in Section 3.4, we present the research questions we answer with our evaluation, as well as the protocol we use to do so.

Table 2

Projects used in the evaluation of plain code serialization: the exact version (SHA) we use for our experiments, the number of lines of Java code (LOC), the number of serialization points we consider (#SPS), the number of types related to the serialization points (#TYPES), the number of serialization points triggered in production by our workloads (#SPS PROD), and the number of types that we target for plain code serialization at runtime (#TYPES PROD).

PROJECT	SHA	LOC	#SPS	#TYPES	#SPS PROD	#TYPES PROD
PDFBox	8f23f8791	170,823	1390	738	183	188
GraphHopper	f64c57c2e	96,716	945	590	199	114
BroadLeaf	951216e942	198,602	2933	1883	508	444
Gephi	82684f927	130,705	1057	1721	140	225
TOTAL			6325	4932	1030	971

3.1. Use-Case: Test generation

In order to evaluate PRODJ, we envision a use case for it. We monitor the invocation of a target method within an executing application, use PRODJ to capture objects associated with the method, and then automatically generate tests for it with plain-code objects. The use of plain-code serialized objects in tests eliminates the dependence on test fixtures and external resources, which contain objects serialized in JSON (Alshahwan et al., 2024) or XML formats (Tiwari et al., 2022). Instead, generated tests have test inputs that are expressed in the native programming language, while also reflecting actual production states (Wang et al., 2017).

First, PRODJ finds methods within the application that can be suitable candidates for test generation, based on a configurable set of criteria. We refer to a candidate method as a method under test, or an MUT. Next, for each MUT, the three phases of PRODJ described in Section 2.4 are applied. The final output of PRODJ is automatically generated JUnit tests,⁴ which follow the *Arrange-Act-Assert* pattern. In contrast to previous work (Tiwari et al., 2022) where production objects serialized as XML are used in the *Arrange* part, such as in Listing 1, or as JSON within resource files (Alshahwan et al., 2024), PRODJ makes it possible to use plain-code serialized objects. Within the *Act* phase, the MUT is invoked on the receiver object, with plain-code serialized arguments. Finally, the *Assert* phase verifies that the output from the invocation of the MUT within the test is the same as the one observed in production. If the assertion involves an object, the expected value of the assertion is also a plain-code serialized object.

3.2. Subjects

We evaluate PRODJ with 4 open-source Java projects that are popular and active, and span diverse application domains. We present them in Table 2. PDFBox is a library and command-line tool for the manipulation of PDF documents.⁵ Its features include conversion between text and PDF documents, extraction of images, and merging of PDF documents, among others. GraphHopper is an open-source routing application.⁶ used to compute a path between input locations on a map. BroadLeaf is an e-commerce web application.⁷ Gephi is a graphical tool for working with graph data,⁸ including capabilities for visualization and analysis.

For each project in Table 2, column 2 links to the exact version (SHA) we use for our evaluation, and column 3 presents the number of lines of Java source code within each project, calculated with *tokei*. Columns 4 and 5 give the number of methods instrumented by PRODJ within the project (#SPS for ‘serialization point’), and consequently, the number

⁴ <https://junit.org/junit5/>

⁵ <https://pdfbox.apache.org/>

⁶ <https://www.graphhopper.com/>

⁷ <https://www.broadleafcommerce.com/>

⁸ <https://gephi.org/>

³ <https://bytebuddy.net/>

of associated types (#TYPES) that need to be instrumented at this point. For example, PRODJinstruments 2,933 methods within BroadLeaf, involving 1,883 types. For our evaluation, we select and instrument methods within each project that are public, non-abstract, non-static, non-deprecated, contain more than one statement, and are declared in a public, non-anonymous, non-local, and non-deprecated class. Columns 6 and 7 of Table 2 provide the same metrics, but computed after runtime observations, i.e., the number of serialization points and types that are covered at runtime, which we discuss in the next subsection.

3.3. Workloads

The key reason we select the projects presented in Table 2 for evaluating plain-code test generation with PRODJ is that we can design representative and repeatable production workloads for each of them, which is essential for our systematic experimental protocols. We invoke PDF-Box through its command line interface, and use its “render” command to convert a PDF document containing multiple pages into JPEG image files. We run the GraphHopperserver, and query GraphHopper for a route between two coordinates in Berlin.⁹ In BroadLeaf, we add some products to our shopping cart, create a user account, log in, and favorite some products. Next, we browse through our shopping cart and change the quantity of some products, before finally completing our order. We use the graphical interface of Gephi to load the “Les Misérables” dataset,¹⁰ and compute all the statistics on it that Gephi offers. This includes computing the average degree of nodes, the network diameter, as well as operations like connected components or Pagerank.

Our interactions with each project result in the invocation of serialization points and correspondingly, the instantiation of types that are instrumented. The last two columns of Table 2 give those numbers. For example, of the 2,933 methods that are instrumented by PRODJ for BroadLeaf, 508 are triggered by our workload (COLUMN #SPS PROD). These methods become the methods under test. Consequently, 444 unique types involved in serialization are instantiated (column #TYPES PROD). As we shall see in Section 4.1, these serialization points will trigger the serialization of thousands of objects under our workload.

3.4. Research questions

We exercise the 4 software applications introduced in Section 3.2 with the production workloads defined in Section 3.3. PRODJ observes these executions and generates plain-code serialized objects within tests. We answer the following research questions based on this evaluation.

RQ1 (Objects)

To what extent can proDJ serialize runtime objects in plain code?

In order to answer this RQ, we report the number of plain-code serialized objects for each project under the considered workload (#OBJ PLAIN CODE). We also present the number of objects serialized solely through the structure-based strategy (SB), as well as those serialized using the trace-based strategy (TB). The findings from this RQ illustrate plain-code serialization in action, and highlight the ability of PRODJ in achieving plain-code serialization for real-world Java projects.

RQ2 (Accuracy)

To what extent does PRODJ successfully reconstruct production objects from plain code?

For this RQ, we use the tests generated (#TESTS) by PRODJ. In contrast to RQ1, these tests contain assertions that help us assess the correctness of deserialized objects. We execute all generated tests, and report the number of tests that PASS as well as the number of tests that FAIL. This RQ allows us to identify objects for which plain-code serialization comes with strong consistency verified with assertions.

RQ3 (Performance)

What is the serialization performance of PRODJ?

To answer this RQ, we compute the overhead for the three phases of PRODJ, pre-execution, execution, and post-execution. The pre-execution overhead consists of source code analysis and preparation for runtime capturing. We use the operating system clock to measure the duration of this phase and normalize the duration by dividing it by the number of analyzed types. Subsequently, at runtime, PRODJ serializes objects and instruments the bytecode of traced classes. We measure the wall-clock time for serialization and normalize it by the amount of serialized objects. For the ByteBuddy instrumentation overhead we use `async-profiler`.¹¹ The post-execution overhead corresponds to the wall-clock time taken for trace-analysis for object reconstruction, and actual code generation. We also measure the time taken to optimize for readability, such as outlining object reconstruction into helper methods, as well as the eventual writing of the generated tests to disk. The total time is then normalized by the number of created tests.

RQ4 (User Study)

What is the opinion of developers on tests containing plain-code serialized objects?

With this RQ, our goal is to assess the opinion of developers on the representation of serialized objects as plain code, relative to other serialized textual representations, specifically XML and JSON. In order to do this, we first select a set of two generated tests for each of the four projects under study, which contain objects serialized at runtime. Next, we invite 17 developers in our professional network over email, to participate in an online survey. As part of the survey, we present each developer a set of five random pairs of generated tests. Each pair contains two variants of the same test, contrasting XML or JSON with PRODJ-produced plain-code serialized objects. This means that one test in the pair contains either XML or JSON, while the other contains PRODJ-produced plain-code serialized objects. We invite each participant to mark the test in each pair that is more readable, according to their preference. For each pair of tests that receives two responses, we compute the Cohen’s kappa coefficient for agreement. For pairs that receive more than two responses, we report the Fleiss’ kappa (Fleiss and Cohen, 1973). We analyze the responses from this survey and present our findings in Section 4.4. Our replication package contains all details of this survey.

4. Experimental results

This section presents the results from our evaluation of PRODJ with the 4 projects.

4.1. RQ1 (Objects)

Columns 2 to 4 of Table 3 present the results for RQ1. First, we note that PRODJ successfully plain-code serializes production objects for all 4 projects. For instance, PRODJ plain-code serializes 100,830 objects for PDFBox. Overall, PRODJ plain-code serializes 174,699 objects that occur at the considered serialization points. Column 3 highlights that nearly 48% of the objects (83,061) are plain-code serialized using the structure-based strategy. This means that PRODJ synthesized a minimal reconstruction plan for each of these objects, and instantiated the meta-variables within it, at runtime. Next, from column 4, we see that 91,638 objects (52%) were serialized via the trace-based strategy. As described in Section 2.3, this means that PRODJ either found a static field with the same value, or recorded constructor and mutating method calls that can reconstruct the object to its target state.

To illustrate plain-code serialization, Listing 9 shows an excerpt of the serialized version of a `GHRequest` object in GraphHopper. This serialized object is concise, and we note that the variable names refer to the

⁹ <https://bit.ly/3LA2PVu>

¹⁰ <https://github.com/gephi/gephi/wiki/Datasets#social-networks>

¹¹ <https://github.com/async-profiler/async-profiler>

Table 3
Results of plain-code serialization on 4 open-source, real-world Java applications.

PROJECT	RQ1: OBJECTS			RQ2: ACCURACY		
	#OBJ PLAIN CODE	SB	TB	#TESTS	PASS	MEDIAN OBJ PER TEST
PDFBox	100,830	29,385	71,445	489	376 (82%)	4
GraphHopper	40,859	38,325	2534	258	191 (74%)	3
BroadLeaf	30,273	13,929	16,344	904	403 (45%)	2
Gephi	2,737	1,422	1315	82	68 (82%)	2
TOTAL	174,699	83,061	91,638	1,733	1,038	2

Answer to RQ1

PRODJ successfully plain-code serializes thousands of objects (174,699) observed at runtime in 4 real-world Java applications. Both the structure-based and trace-based serialization strategies are useful to translate runtime objects in plain Java.

```

1 // Plain-code serialization based on the trace-based strategy
2 GHRequest object = new GHRequest();
3 // GHPoint uses the structure-based strategy
4 GHPoint points = new GHPoint(52.565499, 13.469925);
5 GHPoint pointso = new GHPoint(52.508281, 13.326416);
6 // The point list uses the structure-based strategy
7 List<GHPoint> points1 = new ArrayList<>(List.of(points, pointso));
8 object.setPoints(points1);
9 ...
10 object.setCurbsides(curbsides);

```

Listing 9. Excerpt of the plain-code serialized `GHRequest` object from GraphHopper, requesting a route between two points. Plain-code serialization uses a combination of the structure- and trace-based strategies.

semantics of the program. The names of the variables, such as `points1`, correspond to the names of the parameters they are passed as, and are used thanks to source code analysis happening together with plain-code serialization.

`com.graphhopper.routing.SPTEntry`, which makes up the shortest-paths tree in graphhopper, is another good example. Its entries form a tree structure, each entry has a pointer to its parent, and information about edges, nodes and weights. PRODJ faithfully recreates the recursive structure, unrolling and wiring up the children. To do this, PRODJ naturally starts with the root, and then creates the other entries in dependency order. We see that the proportion of objects that are eligible for the structure- and trace-based strategies is impacted by the application domain. PDFBox uses constants derived from the PDF specification, such as `Author` and `Colors`. These correspond to static fields, resulting in a majority of objects being eligible for trace-based reconstruction with PRODJ. On the other hand, GraphHopper defines collections that PRODJ can effectively reconstruct with the structure-based strategy. From Table 3, we see that each strategy is applied to roughly half of the objects in BroadLeaf and Gephi, resulting from the interactions with the UI as part of our workload.

4.2. RQ2 (Accuracy)

In this research question we reason about the behavior of deserialized objects by running the tests generated by PRODJ. Columns 5 through 8 of Table 3 present the results for RQ2. PRODJ generates 1,733 test cases in total for our 4 study subjects. Those test cases contain assertions that enable us to assess the validity of deserialized objects.

Per RQ1, the total number of objects plain-code serialized by PRODJ is 174,699 but we use only a fraction of them in the 1,733 generated tests. We de-duplicate identical objects, before including them within the generated tests. We consider an object to be identical to another if the generated statements are the same, modulo identifier names, or refer to the same object (`==`) reconstructed using trace-based serializa-

```

1 @Test
2 @DisplayName("alignOrientation(double, double)")
3 void test774AlignOrientation$double_double() {
4     // Arrange: Recreated with SB
5     AngleCalc receiver = new AngleCalc();
6     double baseOrientation = -1.971212484045175;
7     double orientation = 2.744264992492817;
8     // Act
9     double actual = receiver.alignOrientation(baseOrientation,
10        orientation);
11     // Assert
12     assertThat(actual).usingComparator(doubleNanAwareComparator())
13        .isEqualTo(-3.538920314686769);
14 }

```

Listing 10. A test generated by PRODJ for GraphHopper

```

1
2 @Test
3 @DisplayName("getComponents()")
4 void test557GetComponents() throws IOException {
5     // Arrange
6     PDColor receiver = this.createPDColor();
7     // Act
8     float[] actual = receiver.getComponents();
9     // Assert
10    assertThat(actual).isEqualTo(new float[]{0.917647F, 0.92549F,
11        0.937255F});
12 }
13 PDColor createPDColor() throws IOException {
14     // Recreated with TB
15     COSArray receiver = new COSArray();
16     COSBase objectsList = COSFloat.get("0.917647");
17     COSBase objectsListo = COSFloat.get("0.92549");
18     COSBase objectsList1 = COSFloat.get("0.937255");
19     // Recreated with SB
20     Collection<COSBase> objectsList2 = new ArrayList<>(List.of(
21        objectsList, objectsListo, objectsList1));
22     // Recreated with TB
23     receiver.addAll(objectsList2);
24     PDColorSpace receivero = PDDeviceRGB.INSTANCE;
25     return new PDColor(receiver, receivero);
26 }

```

Listing 11. A test generated by PRODJ for PDFBox. A part of the test has been outlined to the helper method `createPDColor`, that recreates production objects

tion. Furthermore, as trace-based serialization indiscriminately records all instances of the types we monitor, we serialize many objects only for trace-based serialization events. For example, we might record multiple `GHRequest` instances, but only one is actually used in a method under test. We then discard all other instances that are not associated with it. This means that we do not use all of the 174,699 plain-code serialized objects within the 1,733 generated tests. We present the median number of objects that are deserialized in each each generated test, in the last column of Table 3. For example, 55 tests generated for PDFBox use 1 plain-code serialized object, while one generated test uses 122 deserialized objects, resulting from the reconstruction of a `List` object and its constituent elements. Overall, the median number of objects in a test generated for PDFBox is 4.

In Listing 10, we present one test generated by PRODJ for GraphHopper, clearly following the *Arrange-Act-Assert* pattern. This test verifies the behavior of the `alignOrientation` method of the `AngleCalc` class in GraphHopper. PRODJ generates the Java statements from lines 5 to 7 at runtime using the structure-based strategy. Note that PRODJ reuses

Answer to RQ2

Our experiment with using plain-code serialization for testing shows that PRODJfaithfully recreates runtime states for all objects in nearly 60 % of the generated tests. The assertions in this testing context provide soundness evidence.

the parameter names of `alignOrientation` for the local variables in the test on lines 6 and 7. The assertion on line 11 verifies that the `double` value output from the invocation of `alignOrientation` on line 9 matches the value recorded at runtime. We also present one PRODJ-generated test for PDFBox in Listing 11. The method under test, called `getComponents` defined in the class `PDCoLor`, returns a non-null, `float` array output, based on the value of a field called `colorSpace`. Note that PRODJ outlines the recreation of the receiver object of type `PDCoLor` into a separate `createPDCoLor` method. Within this method, PRODJ uses object reconstruction mixing for plain-code serializing `PDCoLor`. Specifically, during the pre-execution phase, PRODJ discovers that the associated type `COsArray` is eligible for trace-based serialization. At runtime, PRODJ records a constructor call event for `COsArray` (recreated on line 15), followed by a call to the method `addAll` (line 22). The only parameter of `addAll` is a list, which PRODJ serializes using the structure-based strategy (line 20). However, the contents of this list are serialized using the trace-based strategy (lines 16 - 18). PRODJ uses the parameter name `objectList` of `addAll` to derive the names of the local variables in the helper method. Note that PRODJ observes the calls to the factory method `COsFloat.get` at runtime and therefore recreates these calls, instead of invoking the `COsBase` constructor directly. The `String` arguments for these `COsFloat.get` method calls are directly inlined, in contrast with the `double` arguments on lines 6 and 7 in the first test of Listing 10. This is because PRODJ tries to inline primitive values, but only if this would not break up the *Arrange-Act-Assert* pattern. Next, on line 23 of Listing 11, a `PDCoLorSpace` object is reconstructed using the runtime trace as being an instance of the RGB color space. Finally, the two-argument constructor call on line 24 reuses the reconstructed objects to recreate the runtime `PDCoLor` object. This object is used within the generated test as the receiver for the invocation of `getComponents` (line 8). Note also the inline creation of the `float` array on line 10 within the assertion, which verifies the equality of the expected and actual `float` array output. We argue that these plain-code serialization and readability optimization strategies employed by PRODJ result in tests that not only reflect production observations, but are also focused, well-structured, an easy-to-read. Each generated test focuses on the invocation of a single target method on the reconstructed receiver object. Consequently, the tests generated by PRODJ do not contain test smells: they are not ‘eager tests’ and do not contain multiple assertions (Garousi and Küçük, 2018). We note that these tests may miss inconsistencies within other fields of the receiver that are not visible from the assertion.

All test cases generated by PRODJ compile correctly and when we execute them, 1,038 test cases pass, which is a success rate of 59.9%. Specifically, for three of the four case studies, a large majority of the generated tests pass. For PDFBox, 82% of all tests generated pass, the highest among the four projects. That is a strong property. Given the structure of generated tests, it means that the state of an object after calling a method, i.e., after the ‘Act’ phase of the test, is consistent with the state observed in production (the ‘Assert’ phase). While not being a correctness proof, this is strong consistency evidence.

For BroadLeaf, we note that field dependency injection through the Spring framework results in partially initialized objects. These do not properly reflect production state, causing the test cases to fail. The failing tests are excellent indicators of the fundamental challenges of serialization, as discussed in Section 5.

4.3. RQ3 (Performance)

The results for RQ3 are presented in Table 4, with the three column groups representing the three measurement categories, corresponding to each of the PRODJ phases discussed in Section 2.4. We run each experiment three times and present the median measurement from the three runs. Before the application is run, PRODJ analyses the source code. The time to perform this step is presented in the column *Pre-execution*. This includes parsing the source code, building an AST, and extracting all necessary data for structure- and trace-based serialization. We compute the values by taking the total runtime of this step and dividing it by the amount of types in the project. For example, during the pre-execution phase, PRODJ analyses 1,883 types within BroadLeaf, spending 11 ms per type. On the other hand, PRODJ spends 27 ms on each of the 590 types in GraphHopper. This difference is mostly because the analysis overhead per type is less than the parsing overhead. As BroadLeaf is a bigger project, the cost is amortized over a larger number of diverse types.

After the analysis is completed, we start the application with PRODJ attached. The measurements, made when the application is running, are included within the group of columns labeled *Execution*. Column four reports the average serialization duration per serialized object. We see that the serialization duration is relatively similar for all projects, with GraphHopper being slightly more expensive. This is due to the fact that the serializer looks up many nested types for GraphHopper, which requires a full AST traversal.

The last two columns of Table 4 present the time taken by PRODJ for post-production. This duration includes the time spent offline, analysing the event sequence for objects using trace-based serialization in order to reconstruct them, as well as the time taken for readability optimizations, and writing of the post-processed generated tests to disk. The post-execution phase takes under a minute for all projects. Of the four projects, PRODJ spends most time generating tests for PDFBox because the tests contain many mutator calls that are made on objects in PDFBox. Furthermore, the arguments to these calls often consist of structure-based serialized arrays which have multiple entries. As PRODJ needs to build an AST representation of a variable’s type in order to de-duplicate objects, this necessitates a lot of string parsing. The time per test is the highest for Gephi, as many tests PRODJ tries to generate contain objects it is unable to fully reconstruct. Consequently, the test has to be discarded, increasing the time per successfully generated test.

In addition to the performance data presented in Table 4, we use a profiler for a finer-grained analysis of the time spent by PRODJ on its various runtime tasks. According to our profiles, the class instrumentation makes up around 8 % of the PDFBox runtime, but around 72 % of the BroadLeaf runtime. Most of the time for BroadLeaf is spent deciding whether to instrument a class, as resolving the supertype hierarchy thousands of times per class is not efficient. Recall that PRODJ uses plain JSON for event sequences and serializing millions of small objects to JSON and writing them to disk is expensive, so PRODJ offloads this to a separate thread. Whenever an event needs to be written to disk, it is appended to a threadsafe queue instead, which is then drained in batches by a dedicated thread. Whenever the queue is full, for example, when a multi-threaded application calls instrumented methods in tight loops, the application will be temporarily suspended until the writing thread has caught up. This minimizes the impact on the running application. Analyzing the profiler recordings for PDFBox showed that, while 3.5 % of the time was spent writing the event sequence on different threads, only 0.23 % of the runtime was spent in application code persisting events. Similarly, around 6 % of the runtime for GraphHopper was spent writing the event sequence in a different thread, and also only 0.12 % of the runtime was spent in application code persisting events. For BroadLeaf, PRODJ instruments 1197 types, amounting to a total of over 22,000 rewritten field accesses or method calls. PRODJ uses ByteBuddy and its `MemberSubstitution` for this, leading to about 72 % of the runtime being spent in bytecode rewriting.

Table 4
Performance measurements of the different overhead categories.

Project	Pre-execution		Execution		Post-execution	
	Analysis (ms / type)	Total time (sec)	Serialize (ms / object)	Total time (sec)	Generate (ms / test)	Total time (sec)
PDFBox	18	14	0.14	35	105	59
GraphHopper	27	16	0.40	41	65	40
BroadLeaf	11	23	0.78	157	38	40
Gephi	11	21	0.51	159	145	29

Answer to RQ3

The performance evaluation of PRODJ shows that plain-code serialization takes a few milliseconds per serialized object, which is invisible to the user in interactive use-cases. Contrary to traditional serialization techniques, plain-code serialization involves a computation cost before and after execution.

Overall, serializing an object using PRODJ only takes a few milliseconds. Most of the time spent instrumenting the bytecode outside of this can be attributed to inefficiencies in the usage of ByteBuddy and is not a conceptual limitation. Similarly, most of the post-processing time is spent building an AST and converting strings to types.

4.4. RQ4 (User study)

We conduct a user study per the protocol defined in Section 3.4, based on an online survey with developers. The goal of the survey is to assess the opinion of developers regarding the readability of tests containing plain-code serialized objects. This survey received 17 responses from developers who have between 8 and 43 years of experience with programming (median 20), and between 1 and 30 years programming with Java (median 13). We prepare 16 pairs of tests in total, presenting five random pairs to each respondent. We select two tests for each project which have less than 30 lines of Java code, including outlined utility methods. From Table 5, we see that one test each in PDFBox(PDF1) and GraphHopper(GH1) have 23 lines of code (LOC), the maximum among the 8 plain-code tests generated by PRODJ. For the JSON and XML variants of these tests, we collapsed identical lines in the JSON and XML code. The maximum number of lines in the tests with JSON and XML tests is 1794 and 1,784, respectively (both for PDF2 in Table 5). Note from the table that in each of the 16 test pairs, as well as across all the pairs, the number of lines of code for the plain-code test is significantly lesser than its corresponding JSON and XML representation. Each respondent was shown five randomly selected pairs. One test within the pair contains plain-code serialized objects, while the other was the same test, but with either JSON- or XML-serialized objects. We randomized the order in which the plain-code and JSON or XML tests were presented. The respondents were requested to mark which test in each pair they preferred, with respect to its readability. Moreover, 11 of the 17 respondents shared additional qualitative comments about serialization.

Table 5 summarizes the responses from the survey. We find that through 79 of the total 85 votes (92.9%), the developers showed a preference for the tests containing the plain-code serialized objects, regardless of the differences in the number of lines of code between the plain-code and JSON/XML variants. In only 6 cases, developers preferred the versions with XML or JSON representations. For 11 of the 16 test pairs, the preference for plain-code serialized objects was unanimous. In those cases, the agreement measure (Cohen's and Fleiss' kappas) is a perfect 1.0. In the 5 remaining cases, the agreement varies from poor to moderate.

Qualitative statements by the developers confirmed the preference, as the follows:

Answer to RQ4

Our user study demonstrates, through a sound protocol, that developers rate tests containing plain-code serialized objects higher with respect to their readability, in comparison with equivalent tests with XML- or JSON-serialized objects.

I [preferred] the shorter versions because the longer ones contained serialization "noise" that did not seem relevant.

Intuitively, the [plain-code version] is often more readable. However, legacy code may be hard to test as classes could have many dependencies that are hard to construct. In those cases, the readability gap between object construction and deserialization could be narrowed.

A single developer voted for the XML or JSON variants of all presented pairs, hence they are responsible for 5 of the 6 votes for JSON/XML, owing to the higher likelihood of capturing unintentional breakages:

I prefer the unchanging nature of JSON and XML, which do not get refactored automatically in sync with the application code, when the classes under test change.

The test pair PDF2–JSON has the lowest agreement score. This is due to the aforementioned developer who prefers JSON/XML, and another developer who misunderstood the meaning of the test case, per their comment. Note that the latter respondent ranked plain-code tests higher in the other four pairs they were shown.

One developer correctly noted that the tests would be organized differently in practice:

I believe I would store the long text-snippets as variables in another file, instead of having them in the test.

Finally, a respondent balances the requirements of serialization performance versus readability.:

When sending data back and forth over the wire, more efficient formats may be better, like Protobuf [...]. Overall, it really depends on the task at hand.

This is valid point. Our technique is entirely optimized for readability, and we do recommend binary serialization for performance-sensitive use cases.

5. Limitations and threats

During our prototyping, we have encountered fundamental and interesting challenges of serialization, which we now discuss.

Internal classes

Serialization may fail because of JDK-internal classes, such as those loaded by the bootstrap-classloader. When a bootstrap class is instrumented, additional code is inserted and the bootstrap-classloader is used to look up any required PRODJ classes. As the bootstrap-classloader only loads JDK classes, this does not work and the application crashes. It would be possible to include our agent in the bootstrap-classpath, but this sophistication is out of the scope of a research prototype.

Table 5

Summary of our user study. We select 2 tests from each of the four projects for the study. The second column lists the 16 test pairs we prepare using these tests. Each pair contains one test with plain-code serialized objects and an equivalent test with either JSON or XML representations of the same serialized objects. The number of lines of code in each test in the pair is indicated in the column LOC. We present 5 random test pairs to each of the 17 participants of the study, who vote for tests with plain-code (column 3) or JSON/XML (column 4) serialized objects, based on their readability. The last column presents the Cohen's (C) or Fleiss' (F) kappas of agreement between the participants.

Project	Test pair	LOC			Votes		C or F kappa (agreement)
		Plain-code	JSON	XML	Plain-code	JSON/XML	
PDFBox	PDF1-JSON	23	76	-	8	0	F 1.0 (perfect)
	PDF1-XML		-	65	3	0	F 1.0 (perfect)
	PDF2-JSON	8	1,794	-	4	2	F -0.07 (poor)
	PDF2-XML		-	1,784	6	1	F 0.43 (moderate)
GraphHopper	GH1-JSON	23	349	-	7	1	F 0.5 (moderate)
	GH1-XML		-	259	6	1	F 0.43 (moderate)
	GH2-JSON	18	192	-	2	0	C 1.0 (perfect)
	GH2-XML		-	147	6	0	F 1.0 (perfect)
BroadLeaf	BL1-JSON	8	24	-	4	1	F 0.2 (slight)
	BL1-XML		-	22	9	0	F 1.0 (perfect)
	BL2-JSON	10	20	-	5	0	F 1.0 (perfect)
	BL2-XML		-	17	2	0	C 1.0 (perfect)
Gephi	GEP1-JSON	11	1,128	-	4	0	F 1.0 (perfect)
	GEP1-XML		-	704	6	0	F 1.0 (perfect)
	GEP2-JSON	11	1,128	-	3	0	F 1.0 (perfect)
	GEP2-XML		-	704	4	0	F 1.0 (perfect)
TOTAL	16	112	4,711	3,702	79	6	

Streams

Input streams are problematic as they often read from an ephemeral source, such as a network stream, which might no longer exist when the object is reconstructed. A simple case illustrating this problem is file access or byte array input streams, as can be seen with PDFBox. In this case, PRODJ fails to serialize `java.io.ByteArrayInputStream` and `org.apache.fontbox.afm.CharMetric`.

Runtime code generation

Runtime annotation processing is not handled by PRODJ. In Java, annotations are used for dynamic analysis, and for generating implementations or proxy classes on-the-fly. These classes do not exist on disk, and can therefore not be plain-code serialized. Likewise, for projects such as BroadLeaf that heavily rely on dependency injection through runtime code generation, PRODJ can not faithfully recreate all objects.

Bounded depth

PRODJ uses a bounded serialization approach to limit the size of objects and consequently the size of generated tests. In our experiments, we limit the length of serialized arrays to 25, in order to have readable test cases. Consequently, when a test case asserts behavior that requires more than 25 elements, the assertion fails. For example, 37 test cases generated for PDFBox fail for this reason.

Visibility

The next challenge relates to visibility and encapsulation, as some serialized fields and methods may not be visible in source code. For example, the trace-based reconstruction strategy can capture calls to package-private methods, which cannot be replayed in the test. We observed such a limitation for 22 tests of PDFBox. To address this limitation, PRODJ would need to determine and only use accessible method creating objects of a given type. A similar limitation occurs when fields are set through reflection or native code. Such reconstructions are not supported by PRODJ.

Security Implications

Untrusted object deserialization poses a significant security risk that has been increasingly prevalent in recent years (Sayar et al., 2023b). This holds for plain-code representations of runtime objects observed in the field. An attacker could potentially craft a special object, on which they execute a specific malicious sequence of method calls on the testing or CI machine, which is then captured and replayed by PRODJ. Developers should therefore be mindful of the risks associated with deserialization of production objects. On the positive side, we note that plain code is easier to audit for security than alternative serialization formats.

State

The last major challenge for serialization relates to identify all parts of the state that are related to the execution. For example, 11 tests in PDFBox fail because they try to read a certain file, which does not exist anymore during test execution. PRODJ only recreates states that are local to serialized objects, not the complete *system state*. This is also a potential challenge at reconstruction time, reconstructing a serialized object may incur side effects that we do not handle, such as writing on disk.

Threats to Validity

Threats to the *internal* validity of our findings may arise from the tools that we use to implement PRODJ. We mitigate this threat by relying on open-source state-of-the-art tools to handle instrumentation and code transformation.

External threats to the validity of our findings are related to the artifacts we use for our experiments. We demonstrate the feasibility of plain-code serialization of runtime objects through four large, real-world Java projects exercised with representative field workloads. We envision that our approach will also be applicable to other Java projects.

A threat to the validity of our conclusions from the *user study* of RQ4 arises from the expectation that developers may automatically prefer shorter versions of tests with plain-code serialized objects over their XML or JSON counterparts. We design our user study to minimize this threat and also report the deviations we observe from this outcome. Our key finding is that, even for test pairs with comparable number of lines of code, developers prefer plain-code serialized objects over JSON or XML.

We note that the benefits of plain-code serialization may not *generalize* to all use cases of serialization. For example, inter-process communication involves the transmission of serialized objects that are not necessarily required to be human-readable. Serializing these objects in plain code is also likely to be impractical due to performance costs. On the other hand, where readability is in fact the use case, we believe that our strategy can be adapted to handle plain-code serialization in other programming languages. The key requirement is that their runtime should support introspection, which is the case for most modern languages, including JavaScript and Python.

6. Related work

This section summarizes studies that propose strategies for more efficient serialization, and those that use serialization in the context of testing.

The performance optimization of the (de)serialization mechanism has received much attention in the literature. Runtime code generation can be used to produce code that is tailored to the type being serialized. This has been proposed for programs written in Java (Benani et al., 2004; Aktemur et al., 2005; Horváth et al., 2017), C++ (Tansey and Tilevich, 2008), and Scala (Miller et al., 2013), and for high-performance computing (Tseng et al., 2016). In contrast to these studies targeting performance, PRODJ seeks to optimize readability by expressing objects in idiomatic plain code that is tailored to the project source.

Several studies focus on the construction of an object for use within generated tests. Some approaches solve this problem by suggesting method sequences for setting up test objects. MSeqGen (Thummalapenta et al., 2009) mines source code to propose method call sequences on objects associated with a method under test. These sequences are used within tests generated by Randoop (Pacheco et al., 2007) or PEX (Tillmann and De Halleux, 2008) to increase branch coverage. OCAT (Jaygarl et al., 2010) serializes objects and objects captured by OCAT are recreated through a randomly generated method call sequence within tests generated automatically by Randoop (Pacheco et al., 2007). Seeker (Thummalapenta et al., 2011) uses a combination of static analysis and dynamic symbolic execution to produce sequences of method invocations that initialize test inputs to desired states. Bach et al. (2020) use static analysis to solve the object creation problem. Their approach suggests optimal method call sequences for the creation of test inputs in C++ programs, but does not address the initialization of objects to a target state.

Serialized representations of objects are also directly used within tests. Artzi and colleagues propose Recrash (Artzi et al., 2008), a technique that maintains a shadow stack to store parts of receiving objects and arguments. These partial states are serialized and used to generate unit test cases that can reproduce program crashes. Elbaum et al. (2008) serialize objects associated with invoked MUTs during the execution of system tests, in order to use them within generated unit tests. Pankti (Tiwari et al., 2022) serializes objects associated with MUTs, while monitoring applications in production. The objects are deserialized within generated unit tests to recreate production states. These studies rely on the popular XStream library for serialization to XML, as opposed to plain-code. TestGen by Meta (Alshahwan et al., 2024) serializes runtime objects for the Instagram application, representing them as JSON for use within unit tests. Our key contribution is that we serialize runtime objects in plain-code, combining the complementary structure- and

trace-based strategies. In the technological space, related to our work is a Python tool called fickling (Sultanik, 2023). Instead of directly evaluating a pickled data stream, which could contain malicious executable code, fickling decompiles it into human-readable Python code. This is essential, since security is an important aspect at deserialization time (Gauthier and Bae, 2022; Sayar et al., 2023b).

Closely related to our work is CoDeSe, proposed by Gligoric et al. (2011). CoDeSe, short for CODE-based DEserialization and SERIALization, is an approach to serialize objects as executable Java code focusing on performance. While sharing the same idea of plain-code serialization, our work is very different in the outcome. The objective for serialization with CoDeSe is fast deserialization, and to achieve this, it heavily invokes the Java reflection API and the `Unsafe` class (Evans, 2020) to reconstruct an object and set its fields. We invite the reader to peruse Figure 3b of (Gligoric et al., 2011), which presents the serialized representation of a red-black tree object produced by CoDeSe. The generated code includes calls to `unsafe.allocateInstance(...)` which directs the JVM to create the object in memory, as well as calls to `unsafe.putInt(...)` and `unsafe.putObject(...)` to manipulate its fields. CoDeSe can also generate C code to reconstruct objects, which is then processed through the Java Native Interface (JNI). We argue that this use of reflection and the JNI can be considered as low-level reconstruction. On the contrary, our approach never uses reflection and unsafe calls, and only uses application-specific constructors and methods that are visible in the deserialization scope. The focus of plain-code serialization with PRODJ is idiomacy and readability, as illustrated by the serialized objects embedded within the test presented in Listing 11.

7. Conclusion

In this paper, we have presented the novel concept of plain-code serialization. The key benefits are two-fold. First, the serialized objects are readable by human developers and can be used, for example, in generated tests. Second, deserialization becomes as trivial as executing the statements of the plain-code serialization. We have fully implemented the vision in Java, and used our prototype implementation called PRODJ to serialize more than a hundred thousand objects observed at runtime in four real-world Java applications. Future work can explore the use of advanced impurity analysis to further optimize both structure-based and trace-based plain-code serialization.

CRediT authorship contribution statement

Julian Wachter: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization; **Deepika Tiwari:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization; **Martin Monperrus:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization; **Benoit Baudry:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Data availability

No data was used for the research described in the article.

Declaration of competing interests

We declare that we have no conflict of interest.

Acknowledgements

This work has been partially supported by the Wallenberg Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

References

- Afshan, S., McMinn, P., Stevenson, M., 2013. Evolving readable string test inputs using a natural language model to reduce human oracle cost. In: 2013 IEEE Sixth International Conference on Software Testing, Verification and Validation. IEEE, pp. 352–361.
- Aktemur, B., Jones, J., Kamin, S., Clausen, L., 2005. Optimizing marshalling by run-time program generation. In: Proceedings of the International Conference on Generative Programming and Component Engineering (GPCE). Springer, pp. 221–236.
- Alshahwan, N., Harman, M., Marginean, A., Tal, R., Wang, E., 2024. Observation-based unit test generation at meta. In: Proceedings of the International Conference on the Foundations of Software Engineering, pp. 173–184.
- Artzi, S., Kim, S., Ernst, M.D., 2008. Recrash: making software failures reproducible by preserving object states. In: ECOOP 2008—Object-Oriented Programming: 22nd European Conference Paphos, Cyprus, July 7–11, 2008 Proceedings 22. Springer, pp. 542–565.
- Bach, T., Pannemans, R., Andrzejak, A., 2020. Determining method-call sequences for object creation in c++. In: 2020 IEEE 13th International Conference on Software Testing, Validation and Verification (ICST). IEEE, pp. 108–119.
- Bennani, M.T., Blain, L., Courtès, L., Fabre, J., Killijian, M., Marsden, E., Taiani, F., 2004. Implementing simple replication protocols using CORBA portable interceptors and java serialization. In: 2004 International Conference on Dependable Systems and Networks (DSN).
- Bray, T., Paoli, J., Sperberg-McQueen, C.M., Maler, E., Yergeau, F., 1997. Extensible markup language (XML). W. W. W. J. 2 (4), 27–66.
- De Moura, L., Bjørner, N., 2008. Z3: An efficient SMT solver. In: Tools and Algorithms for the Construction and Analysis of Systems: 14th International Conference, TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29–April 6, 2008. Proceedings 14. Springer, pp. 337–340.
- Eddelbuettel, D., Stokely, M., Ooms, J., 2014. RProtoBuf: efficient cross-language data serialization in R. arXiv preprint arXiv:1401.7372 .
- Elbaum, S., Chin, H.N., Dwyer, M.B., Jorde, M., 2008. Carving and replaying differential unit test cases from system test cases. IEEE Trans. Software Eng. 35 (1), 29–45.
- Evans, B., 2020. The Unsafe Class: Unsafe at Any Speed. Available at <https://blogs.oracle.com/javamagazine/post/the-unsafe-class-unsafe-at-any-speed>.
- Fleiss, J.L., Cohen, J., 1973. The equivalence of weighted kappa and the intraclass correlation coefficient as measures of reliability. Educ. Psychol. Meas. 33 (3), 613–619.
- Garousi, V., Küçük, B., 2018. Smells in software test code: a survey of knowledge in industry and academia. J. Syst. Softw. 138, 52–81.
- Gauthier, F., Bae, S., 2022. Runtime prevention of deserialization attacks. In: Pasquale, L., Treude, C. (Eds.), 44th IEEE/ACM International Conference on Software Engineering: New Ideas and Emerging Results ICSE (NIER) 2022.
- Gligoric, M., Marinov, D., Kamin, S., 2011. CoDeSe: fast deserialization via code generation. In: Proceedings of the 2011 International Symposium on Software Testing and Analysis, pp. 298–308. <http://users.ece.utexas.edu/~gligoric/papers/Gligoric-ETAL11CoDeSe.pdf>.
- Godefroid, P., Huang, B.-Y., Polishchuk, M., 2020. Intelligent REST API data fuzzing. In: Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, pp. 725–736.
- Google, . Protocol buffers. 2023 Accessed May 1, 2023. <https://developers.google.com/protocol-buffers/>.
- Harrand, N., Durieux, T., Broman, D., Baudry, B., 2021. The behavioral diversity of java JSON libraries. In: 2021 IEEE 32nd International Symposium on Software Reliability Engineering (ISSRE). IEEE, pp. 412–422.
- Horváth, G., Pataki, N., Balassi, M., 2017. Code generation in serializers and comparators of apache flink. In: Proceedings of the 12th Workshop on Implementation, Compilation, Optimization of Object-Oriented Languages, Programs and Systems, pp. 1–6.
- Jaygarl, H., Kim, S., Xie, T., Chang, C.K., 2010. OCAT: object capture-based automated testing. In: Proceedings of the 19th International Symposium on Software Testing and Analysis, pp. 159–170. http://people.csail.mit.edu/hunkim/images/9/99/Ocat_issta2010.pdf.
- Kim, M., Xin, Q., Sinha, S., Orso, A., 2022. Automated test generation for REST APIs: no time to rest yet. In: Proceedings of ISSTA, pp. 289–301. <https://doi.org/10.1145/3533767.3534401>
- Liu, Z.H., Hammerschmidt, B., McMahon, D., Chang, H., Lu, Y., Spiegel, J., Sosa, A.C., Suresh, S., Arora, G., Arora, V., 2020. Native JSON datatype support: maturing SQL and NoSQL convergence in oracle database. Proceedings of the VLDB Endowment 13 (12), 3059–3071.
- Miller, H., Haller, P., Burmako, E., Odersky, M., 2013. Instant pickles: generating object-oriented pickler combinators for fast and extensible serialization. In: Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages & Applications, pp. 183–202. <http://lamp-epfl.ch/~hmiller/files/oopsla-pickling-moved.pdf>.
- Neumann, A., Laranjeiro, N., Bernardino, J., 2018. An analysis of public REST web service APIs. IEEE Trans. Serv. Comput. 14 (4), 957–970.
- Pacheco, C., Lahiri, S.K., Ernst, M.D., Ball, T., 2007. Feedback-directed random test generation. In: ICSE 2007, Proceedings of the 29th International Conference on Software Engineering. Minneapolis, MN, USA, pp. 75–84.
- Pawlak, R., Monperrus, M., Petitprez, N., Noguera, C., Seinturier, L., 2015. Spoon: a library for implementing analyses and transformations of java source code. Softw. Pract. Exper. 46, 1155–1179. <https://hal.archives-ouvertes.fr/hal-01078532/document>. <https://doi.org/10.1002/spe.2346>
- Pourhabibi, A., Gupta, S., Kassir, H., Sutherland, M., Tian, Z., Drumond, M.P., Falsafi, B., Koch, C., 2020. Optimus prime: accelerating data transformation in servers. In: Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems, pp. 1203–1216.
- Sayar, I., Bartel, A., Bodden, E., Traon, Y.L., 2023b. An in-depth study of java deserialization remote-code execution exploits and vulnerabilities. ACM Trans. Softw. Eng. Methodol. 32 (1), 1–45. <https://doi.org/10.1145/3554732>.
- Sultanik, E., Never a dill moment: Exploiting machine learning pickle files. <https://blog.trailofbits.com/2021/03/15/never-a-dill-moment-exploiting-machine-learning-pickle-files/>. Accessed Sep 27, 2023.
- Tansey, W., Tilevich, E., 2008. Efficient automated marshaling of C++ data structures for MPI applications. In: 2008 IEEE International Symposium on Parallel and Distributed Processing. IEEE, pp. 1–12.
- Thummalapenta, S., Xie, T., Tillmann, N., De Halleux, J., Schulte, W., 2009. MSe-qGen: object-oriented unit-test generation via mining source code. In: Proceedings of the 7th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering, pp. 193–202.
- Thummalapenta, S., Xie, T., Tillmann, N., De Halleux, J., Su, Z., 2011. Synthesizing method sequences for high-coverage testing. ACM SIGPLAN Not. 46 (10), 189–206. <https://web.cs.ucdavis.edu/~su/publications/oopsla11.pdf>.
- Tillmann, N., De Halleux, J., 2008. Pex—white box test generation for .net. In: Tests and Proofs: Second International Conference, TAP 2008, Prato, Italy, April 9–11, 2008. Proceedings 2. Springer, pp. 134–153.
- Tiwari, D., Zhang, L., Monperrus, M., Baudry, B., 2022. Production monitoring to improve test suites. IEEE Trans. Reliab. 71 (3), 1381–1397. <https://doi.org/10.1109/TR.2021.3101318>
- Tseng, H.-W., Zhao, Q., Zhou, Y., Gahagan, M., Swanson, S., 2016. Morpheus: creating application objects efficiently for heterogeneous computing. ACM SIGARCH Comput. Archit. News 44 (3), 53–65.
- Viotti, J.C., Kinderkhedha, M., 2022. A survey of JSON-compatible binary serialization specifications. arXiv preprint arXiv:2201.02089.
- Wang, Q., Brun, Y., Orso, A., 2017. Behavioral execution comparison: are tests representative of field behavior? In: 2017 IEEE International Conference on Software Testing, Verification and Validation (ICST). IEEE, pp. 321–332.
- Wong, E. Y.C., Chan, A. T.S., Leong, H.V., 2004. Xstream: a middleware for streaming xml contents over wireless environments. IEEE Trans. Software Eng. 30 (12), 918–935.
- Ye, Q., Delaware, B., 2019. A verified protocol buffer compiler. In: Proceedings of the 8th ACM SIGPLAN International Conference on Certified Programs and Proofs, pp. 222–233.