

# 1 CrySL: An Extensible Approach to Validating the 2 Correct Usage of Cryptographic APIs

3 **Stefan Krüger**

4 Paderborn University, Germany  
5 stefan.krueger@uni-paderborn.de

6 **Johannes Späth**

7 Fraunhofer IEM  
8 johannes.spaeth@iem.fraunhofer.de

9 **Karim Ali**

10 University of Alberta, Canada  
11 karim.ali@ualberta.ca

12 **Eric Bodden**

13 Paderborn University & Fraunhofer IEM, Germany  
14 eric.bodden@uni-paderborn.de

15 **Mira Mezini**

16 Technische Universität Darmstadt, Germany  
17 mezini@cs.tu-darmstadt.de

---

## 18 — Abstract —

19 Various studies have empirically shown that the majority of Java and Android apps misuse  
20 cryptographic libraries, causing devastating breaches of data security. It is crucial to detect such  
21 misuses early in the development process. To detect cryptography misuses, one must first define  
22 secure uses, a process mastered primarily by cryptography experts, and not by developers.

23 In this paper, we present CRYSL, a definition language for bridging the cognitive gap between  
24 cryptography experts and developers. CRYSL enables cryptography experts to specify the secure  
25 usage of the cryptographic libraries that they provide. We have implemented a compiler that  
26 translates such CRYSL specification into a context-sensitive and flow-sensitive demand-driven  
27 static analysis. The analysis then helps developers by automatically checking a given Java or  
28 Android app for compliance with the CRYSL-encoded rules.

29 We have designed an extensive CRYSL rule set for the Java Cryptography Architecture (JCA),  
30 and empirically evaluated it by analyzing 10,000 current Android apps. Our results show that  
31 misuse of cryptographic APIs is still widespread, with 95% of apps containing at least one misuse.  
32 Our easily extensible CRYSL rule set covers more violations than previous special-purpose tools  
33 with hard-coded rules, with our tooling offering a more precise analysis.

34 **2012 ACM Subject Classification** Security and privacy → Software and application security,  
35 Software and its engineering → Software defect analysis, Software and its engineering → Syntax  
36 & Semantics

37 **Keywords and phrases** cryptography, domain-specific language, static analysis

38 **Digital Object Identifier** 10.4230/LIPIcs.ECOOP.2018.10

39



© Stefan Krüger and Johannes Späth and Karim Ali and Eric Bodden and Mira Mezini;  
licensed under Creative Commons License CC-BY

32nd European Conference on Object-Oriented Programming (ECOOP 2018).

Editor: Todd Millstein; Article No. 10; pp. 10:1–10:28

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

40 **1 Introduction**

41 Digital devices are increasingly storing sensitive data, which is often protected using cryp-  
42 tography. However, developers must not only use secure cryptographic algorithms, but also  
43 *securely* integrate such algorithms into their code. Unfortunately, prior studies suggest that  
44 this is rarely the case. Lazar et al. [22] examined 269 published cryptography-related vulner-  
45 abilities. They found that 223 are caused by developers misusing a security library while only  
46 46 result from faulty library implementations. Egele et al. [13] statically analyzed 11,748 An-  
47 droid apps using cryptography-related application interfaces (Crypto APIs) and found 88%  
48 of them violated at least one basic cryptography rule. Chatzikonstantinou et al. [12] reached  
49 a similar conclusion by analyzing apps manually and dynamically. In 2017, VeraCode listed  
50 insecure uses of cryptography as the second-most prevalent application-security issue right  
51 after information leakage [11]. Such pervasive insecure use of Crypto APIs leads to dev-  
52 astating vulnerabilities such as data breaches in a large number of applications. Rasthofer  
53 et al. [31] showed that *virtually all* smartphone apps that rely on cloud services use hard-  
54 coded keys. A simple decompilation gives adversaries access to those keys and to all data  
55 that these apps store in the cloud.

56 Nadi et al. [27] were the first to investigate why developers often struggle to use  
57 Crypto APIs. The authors conducted four studies, two of which survey Java developers  
58 familiar with the Java Crypto APIs. The majority of participants (65%) found their re-  
59 spective Crypto APIs hard to use. When asked why, participants mentioned the API level  
60 of abstraction, insufficient documentation without examples, and an API design that makes  
61 it difficult to understand how to properly use the API. A potential long-term solution is to  
62 redesign the APIs such that they provide an easy-to-use interface for developers that is se-  
63 cure by default. However, it remains crucial to detect and fix the existing insecure API uses.  
64 When asked about what would simplify their API usage, participants wished they had tools  
65 that help them automatically detect misuses and suggest possible fixes [27]. Unfortunately,  
66 approaches based solely on specification inference and anomaly detection [34] are not viable  
67 for Crypto APIs, because—as elaborated above—most uses of Crypto APIs are insecure.

68 Previous work has tried to detect misuses of Crypto APIs through static analysis. While  
69 this is a step in the right direction, existing approaches are insufficient for several reasons.  
70 First, these approaches implement mostly lightweight *syntactic checks*, which yield fast  
71 analysis times at the cost of exposing a high number of false negatives. Therefore, such  
72 analyses fail to warn about many insecure (especially non-trivial) uses of cryptography. For  
73 instance, applications using password-based encryption commonly do not clear passwords  
74 from heap memory and instead rely on garbage collection to free the respective memory  
75 space. Moreover, existing tools cannot easily be extended to cover those rules; instead they  
76 have cryptography-specific usage rules *hard coded*. The Java Cryptography Architecture  
77 (JCA), the primary cryptography API for Java applications [27], offers a plugin design that  
78 enables different providers to offer different crypto implementations through the same API,  
79 often imposing slightly different usage requirements on their clients. Hard-coded rules can  
80 hardly possibly reflect this diversity.

81 In this paper, we present CRYSL, a definition language that enables cryptography experts  
82 to specify the secure usage of their Crypto APIs in a lightweight special-purpose syntax. We  
83 also present a CRYSL compiler that parses and type-checks CRYSL rules and translates  
84 them into an efficient, yet precise flow-sensitive and context-sensitive static data-flow ana-  
85 lysis. The analysis automatically checks a given Java or Android app for compliance with  
86 the encoded CRYSL rules. CRYSL was specifically designed for (and with the help of) cryp-

87 tography experts. Our approach goes beyond methods that are useful for general validation  
88 of API usage (e.g., tpestate analysis [3, 7, 28, 8] and data-flow checks [2, 5]) by enabling  
89 the expression of domain-specific constraints related to cryptographic algorithms and their  
90 parameters.

91 To evaluate CRYSL, we built the most comprehensive rule set available for the JCA  
92 classes and interfaces to date, and encoded it in CRYSL. We then used the generated static  
93 analysis `COGNICRYPTSAST` to scan 10,000 Android apps. We have also modelled the existing  
94 hard-coded rules by Egele et al. [13] in CRYSL and compared the findings of the generated  
95 static analysis (`COGNICRYPTCL`) to those of `COGNICRYPTSAST`. Our more comprehensive rule  
96 set reports 3× more violations, most of which are true warnings. With such comprehensive  
97 rules, `COGNICRYPTSAST` finds at least one misuse in 95% of the apps. `COGNICRYPTSAST` is  
98 also highly efficient: for more than 75% of the apps, the analysis finishes in under 3 minutes  
99 per app, where most of the time is spent in Android-specific call graph construction.

100 In summary, this paper presents the following contributions:

- 101 ■ We introduce CRYSL, a definition language to specify correct usages of Crypto APIs.
- 102 ■ We encode a comprehensive specification of correct usages of the JCA in CRYSL.
- 103 ■ We present a CRYSL compiler that translates CRYSL rules into a static analysis to find  
104 violations in a given Java or Android app.
- 105 ■ We empirically evaluate `COGNICRYPTSAST` on 10,000 Android apps.

106 We have integrated `COGNICRYPTSAST` into crypto assistant `COGNICRYPT` [20] and have  
107 open-sourced our implementation and artifacts on GitHub. `COGNICRYPTSAST` is available  
108 at <https://github.com/CROSSINGTUD/CryptoAnalysis>. The latest version of the CRYSL  
109 rules for the JCA can be accessed at <https://github.com/CROSSINGTUD/Crypto-API-Rules>.

110

## 111 **2 Related Work**

112 Before we discuss the details of our approach, we contrast it with the following related lines  
113 of work: approaches for specifying API (mis)uses, approaches for inferring API specifica-  
114 tions, and previous approaches for detecting misuses of security APIs. Our review of these  
115 approaches shows that existing specification languages are not optimally suited for defining  
116 misuses of Crypto APIs. Additionally, automated inference of correct uses of Crypto APIs is  
117 hard to achieve, and existing tools for detecting misuses of Crypto APIs are limited mainly  
118 because they have hard-coded rule sets, and support for the most part lightweight syntactic  
119 analyses.

### 120 **2.1 Languages for Specifying and Checking API Properties**

121 There is a significant body of research on textual specification languages that ensure API  
122 properties by means of static data-flow analysis. Tracematches [3] were designed to check  
123 tpestate properties defined by regular expressions over runtime objects. Bodden et al. [8, 10]  
124 as well as Naeem and Lhoták [28] present algorithms to (partially) evaluate state matches  
125 prior to the program execution, using static analysis.

126 Martin et al. [24] present Program Query Language (PQL) that enables a developer to  
127 specify patterns of event sequences that constitute potentially defective behaviour. A dy-  
128 namic analysis (i.e., tracematches optimized by a static pre-analysis) matches the patterns  
129 against a given program run. A pattern may include a fix that is applied to each match

130 by dynamic instrumentation. PQL has been applied to detecting security-related vulnerabil-  
 131 ities such as memory leaks [24], SQL injection and cross-site scripting [23]. Compared to  
 132 tracematches, PQL captures a greater variety of pattern specifications, at the disadvantage  
 133 of using only flow-insensitive static optimizations. PQL serves as the main inspiration for  
 134 the CRYSL syntax. Other languages that pursue similar goals include PTQL [16], PDL [26],  
 135 and TS4J [9].

136 We investigated tracematches and PQL in detail, yet found them insufficiently equipped  
 137 for the task at hand. First, both systems follow a black-list approach by defining and  
 138 finding incorrect program behaviour. We initially followed this approach for crypto-usage  
 139 mistakes but quickly discovered that it would lead to long, repetitive, and convoluted  
 140 misuse-definitions. Consequently, CRYSL defines desired behaviour, which in the case of  
 141 Crypto APIs leads to more compact specifications. Second, the above languages are general-  
 142 purpose languages for bug finding, which causes them to miss features essential to define  
 143 secure usages of Crypto APIs in particular. The strong focus of CRYSL on cryptography  
 144 allows us to cover a greater portion of cryptography-related problems in CRYSL compared  
 145 to other languages, while at the same time keeping CRYSL relatively simple. Third, the  
 146 CRYSL compiler generates state-of-the-art static analyses that were shown to have better  
 147 performance and precision than other approaches [37], lowering the threat of false warnings.

## 148 2.2 Inference/Mining of API-usage specifications

149 As an alternative to specifying API-usage properties manually, one can attempt to infer  
 150 them from existing program code. Robillard et al. [33] surveyed over 60 approaches to API  
 151 property inference. As this survey shows, however, all but two of the surveyed approaches  
 152 infer patterns from client code (i.e., from applications that use the API in question). When  
 153 it comes to Crypto APIs, however, past studies have shown that the majority of existing  
 154 usages of those APIs is, in fact, insecure [13, 12, 35]. Another idea that appears sensible at  
 155 first sight is to infer correct usage of Crypto APIs from posts on developer portals such as  
 156 StackOverflow. However, recent studies show that the “solutions” posted there often include  
 157 insecure code [1].

158 In result, one can only conclude that automated mining of API-usage specifications is  
 159 very challenging for Crypto APIs, if it is possible at all. In the future, we plan to investigate  
 160 a semi-automated approach in which we use automated inference to infer at least partial  
 161 specifications, but directly in CRYSL, that security experts can then further correct and  
 162 complete by hand.

## 163 2.3 Detecting Misuses of Security APIs

164 Only few previous approaches specifically address the detection of misuses of *security* APIs.  
 165 CRYPTOLINT [13] performs a lightweight syntactic analysis to detect violations of exactly  
 166 six hard-coded usage rules for the JCA in Android apps. Those six rules, while important  
 167 to obey for security, resemble only a tiny fraction of the rule set we provide in this work. It  
 168 is also hard to specify and validate new rules using CRYPTOLINT, because they would have  
 169 to be hard-coded. Unlike CRYPTOLINT, CRYSL is designed to allow crypto experts to also  
 170 express comprehensive and complex rules with ease. In Section 8, we extensively compare  
 171 our tool COGNICRYPT<sub>SAST</sub> to CRYPTOLINT.

172 Another tool that finds misuses of Crypto APIs is Crypto Misuse Analyzer (CMA) [35].  
 173 Similar to CRYPTOLINT, CMA’s rules are hard-coded, and its static analysis is rather basic.

```

1  SecretKeyGenerator kG = KeyGenerator.getInstance("AES");
2  kG.init(128);
3  SecretKey cipherKey = kG.generateKey();
4
5  String plaintextMSG = getMessage();
6  Cipher ciph = Cipher.getInstance("AES/GCM");
7  ciph.init(Cipher.ENCRYPT_MODE, cipherKey);
8  byte[] cipherText = ciph.doFinal(plaintextMSG.getBytes("UTF-8"));

```

■ **Figure 1** An example illustrating the use of `javax.crypto.KeyGenerator` to implement data encryption in Java.

174 Many of CMA’s hard-coded rules are also contained in the CRYSL rule set that we provide.  
 175 Unlike COGNICRYPT<sub>SAST</sub>, CMA has been evaluated on a small dataset of only 45 apps.

176 Chatzikonstantinou et al. [12] manually identified misuses of Crypto APIs in 49 apps  
 177 and then verified their findings using a dynamic checker. All three studies concluded that  
 178 at least 88% of the studied apps misuse at least one Crypto API.

179 None of the previous approaches facilitates rule creation by means of a higher-level  
 180 specification language. Instead, the rules are hard-coded into each tool, making it hard for  
 181 non-experts in static analysis to extend or alter the rule set, and impossible to share rules  
 182 among tools. Moreover, such hard-coded rules are quite restricted, causing the tools to have  
 183 a very low recall (i.e., missing many actual API misuses). CRYSL, on the other hand, due  
 184 to its Java-like syntax, enables cryptography experts to easily define new rules. The CRYSL  
 185 compiler then automatically transforms those rules into appropriate, highly-precise static-  
 186 analysis checks. By defining crypto-usage rules in CRYSL instead of hard-coding them, one  
 187 also makes those rules reusable in different contexts.

### 188 3 An Example of a Secure Usage of Crypto APIs

189 Throughout the paper, we will use the code example in Figure 1 to motivate the language  
 190 features in CRYSL. The code in this figure constitutes an API usage that according to the  
 191 current state of cryptography research can be considered secure. Lines 1–3 generate a 128-  
 192 bit secret key to use with the encryption algorithm AES. Lines 5–7 use that key to initialize  
 193 a Java `Cipher` object that encrypts `plaintextMSG`. Since AES encrypts plaintext block by  
 194 block, it must be configured to use one of several *modes of operation*. The mode of operation  
 195 determines how to encrypt a block based on the encryption of the preceding block(s). Line 6  
 196 configures `Cipher` to use the Galois/Counter Mode (GCM) of operation [25].

197 Although the code example may look straightforward, a number of subtle alterations  
 198 to the code would render the encryption non-functional or even insecure. First, both  
 199 `KeyGenerator` and `Cipher` only support a limited choice of encryption algorithms. If the  
 200 developer passes an unsupported algorithm to either `getInstance` methods, the respective  
 201 line will throw a runtime exception. Similarly, the design of the APIs separates the classes  
 202 for key generation and encryption. Therefore, the developer needs to make sure they pass the  
 203 same algorithm (here "AES") to the `getInstance` methods of `KeyGenerator` and `Cipher`.  
 204 If the developer does not configure the algorithms as such, the generated key will not fit  
 205 the encryption algorithm, and the encryption will fail by throwing a runtime exception.  
 206 None of the existing tools discussed in Section 2.3 are capable of detecting such functional  
 207 misuses. Moreover, some supported algorithms are no longer considered secure (e.g., DES  
 208 or AES/ECB [15]). If the developer selects such an algorithm, the program will still run  
 209 to completion, but the resulting encryption could easily be broken by attackers. To make

```

METHOD :=
  methname(PARAMETERS)

PARAMETERS :=
  varname , PARAMETERS
  varname

TYPES :=
  QualifiedClassName , TYPES
  TYPE

CONSTANTLIST :=
  constant , CONSTANTLIST
  constant

AGGREGATE :=
  label | AGGREGATE
  label ;

EVENT :=
  AGGREGATE
  label : METHOD
  label : varname = METHOD

PREDICATE :=
  predname(PARAMETERS)
  predname(PARAMETERS) after EVENT

PREDICATES :=
  PREDICATE ; PREDICATES

```

*A: B = C(D) — a single event with label A consisting of method C, its parameter D, and return object B*

■ **Figure 2** Basic CRYSL syntax elements.

210 things worse, the JCA, the most popular API, offers the insecure ECB mode by default (i.e.,  
 211 when developers request only "AES" without specifying a mode of operation explicitly).

212 To use Crypto APIs properly, developers generally have to take into consideration two  
 213 dimensions of correctness: (1) the functional correctness that allows the program to run  
 214 and terminate successfully and (2) the provided security guarantees. Prior empirical studies  
 215 have shown that developers, for instance by looking for code examples on web portals such  
 216 as StackOverflow [14], frequently succeed in obtaining functionally correct code. However,  
 217 they often fail to obtain a secure use of Crypto APIs, primarily because most code examples  
 218 on those web portals provide “solutions” that themselves are insecure [14].

|                             |   |
|-----------------------------|---|
| <b>SPEC TYPE;</b>           |   |
| <b>OBJECTS</b>              |   |
| OBJECTS :=                  |   |
| OBJECT ; OBJECTS            | <i>A ; B — a list of objects A and B</i>  |
| OBJECT ;                    | <i>A — a list of the single object A</i>  |
| OBJECT :=                   |   |
| TYPE varname                | <i>A B — object B of Java type A</i>  |
| <b>EVENTS</b>               |   |
| EVENTS :=                   |   |
| EVENT ; EVENTS              | <i>A ; B — a list of events A and B</i>   |
| EVENT ;                     | <i>A — a list of the single event A</i>   |
| <b>FORBIDDEN</b>            |   |
| FMETHODS :=                 |   |
| FMETHOD ; FMETHODS          | <i>A ; B — a list of forbidden A and B</i>  |
| FMETHOD ;                   | <i>A — a list of the single forbidden method A</i>  |
| FMETHOD :=                  |   |
| methname(TYPES) => label    | <i>A(B) =&gt; C — a forbidden method named A with parameter of Type B and replacement C</i> |
| <b>ORDER</b>                |   |
| USAGEPATTERN :=             |   |
| USAGEPATTERN , USAGEPATTERN | <i>A , B — A followed by B</i>  |
| USAGEPATTERN   USAGEPATTERN | <i>A   B — A or B</i>   |
| USAGEPATTERN ?              | <i>A? — A is optional</i>   |
| USAGEPATTERN *              | <i>A* — 0 or more As</i>  |
| USAGEPATTERN +              | <i>A+ — 1 or more As</i>  |
| ( USAGEPATTERN )            | <i>(A) — grouping</i>   |
| AGGREGATE                   |   |
| <b>CONSTRAINTS</b>          |   |
| CONSTRAINTS :=              |   |
| CONSTRAINT ; CONSTRAINTS    |   |
| CONSTRAINT => CONSTRAINT    | <i>A =&gt; B — A implies B</i>  |
| CONSTRAINT                  |   |
| CONSTRAINT :=               |   |
| varname in { CONSTANTLIST } | <i>A in {1, 2} — A should be 1 or 2</i>   |
| <b>REQUIRES</b>             |   |
| REQ_PREDICATES :=           |   |
| PREDICATES                  |   |
| <b>ENSURES</b>              |   |
| ENS_PREDICATES :=           |   |
| PREDICATES                  |   |
| <b>NEGATES</b>              |   |
| NEG_PREDICATES :=           |   |
| PREDICATES                  |   |

■ **Figure 3** CRYSL rule syntax in Extended Backus-Naur Form (EBNF) [6].

## 219 **4** CrySL Syntax

220 As we discuss in Section 2.2, mining API properties for Crypto APIs is extremely challenging,  
 221 if possible at all, due to the overwhelming number of misuses one finds in actual applica-  
 222 tions. Hence, instead of relying on the security of existing usages and examples, we here  
 223 follow an approach in which cryptography experts define correct API usages manually in a  
 224 special-purpose language, CRYSL. In this section, we give an overview of the CRYSL syntax  
 225 elements. A formal treatment of the CRYSL semantics is presented in Section 5. Figure 2  
 226 presents the basic syntactic elements of CRYSL, and Figure 3 presents the full syntax for  
 227 CRYSL rules. Figure 4 shows an abbreviated CRYSL rule for `javax.crypto.KeyGenerator`.

### 228 **4.1 Design Decisions Behind CrySL**

229 We designed CRYSL specifically with crypto experts in mind, and in fact with the help of  
 230 crypto experts. This work was carried out in the context of a large collaborative research  
 231 center than involves more than a dozen research groups involved in cryptography research.  
 232 As a result of the domain research conducted within this center, we made the following  
 233 design decisions when designing CRYSL.

234 **White listing.** During our domain analysis, we observed that, for the given Crypto APIs,  
 235 there are many ways they can be misused, but only a few that correspond to correct  
 236 and secure usages. To obtain concise usage specifications, we decided to design CRYSL  
 237 to use white listing in most places (i.e., defining secure uses explicitly, while implicitly  
 238 assuming all deviations from this norm to be insecure).

239 **Typestate and data flow.** When reviewing potential misuses, we observed that many of  
 240 them are related to data flows and typestate properties [38]. Such misuses occur because  
 241 developers call the wrong methods on the API objects at hand, call them in an incorrect  
 242 order or miss to call the methods entirely. Data-flow properties are important when  
 243 reasoning about how certain data is being used (e.g., passwords, keys or seed material).

244 **String and integer constraints.** In the crypto domain, string and integer parameters are  
 245 ubiquitously used to select or parametrize specific cryptography algorithms. Strings are  
 246 widely used, because they are easily recognizable, configurable, and exchangeable. How-  
 247 ever, specifying an incorrect string parameter may result in the selection of an insecure  
 248 algorithm or algorithm combination. Many APIs also use strings for user credentials.  
 249 Those credentials, passwords in particular, should not be hard-coded into the program's  
 250 bytecode. A precise specification of correct crypto uses must therefore comprise con-  
 251 straints over string and integer parameters.

252 **Tool-independent semantics.** We equipped CRYSL with a tool-independent semantics (to  
 253 be presented in Section 5). In the future, those semantics will enable us and others to  
 254 build other or more effective tools for working with CRYSL. For instance, in addition to  
 255 the static analysis the CRYSL compiler derives from the semantics within this paper, we  
 256 are currently working on a dynamic checker to identify and mitigate CRYSL violations  
 257 at runtime.

258 Our desire to allow crypto experts to easily express secure crypto uses also precludes us  
 259 from using existing generic definition languages such as Datalog. Such languages, or minor  
 260 extensions thereof, might have sufficient expressive power. However, following discussions  
 261 with crypto developers, we had to acknowledge that they are often unfamiliar with those  
 262 languages' concepts. CRYSL thus deliberately only includes concepts familiar to those de-



263 velopers, hence supporting an easy understanding. We next explain the elements that a  
 264 typical CRYSL rule comprises.

## 265 4.2 Mandatory Sections in a CrySL Rule

266 To provide simple and reusable constructs, a CRYSL rule is defined on the level of individual  
 267 classes. Therefore, the rule starts off by stating the class that it is defined for.

268 In Figure 4, the **OBJECTS** section defines three objects<sup>1</sup> to be used in later sections of  
 269 the rule (e.g., the object `algorithm` of type `String`). These objects are typically used as  
 270 parameters or return values in the **EVENTS** section.

271 The **EVENTS** section defines all methods that may contribute to the successful use of a  
 272 `KeyGenerator` object, including two *method event patterns* (Lines 17–18). The first pat-  
 273 tern matches calls to `getInstance(String algorithm)`, but the second pattern actually  
 274 matches calls to two overloaded `getInstance` methods:

```
275 ■ getInstance(String algorithm, Provider provider)
276 ■ getInstance(String algorithm, String provider)
```

277 The first parameter of all three methods is a `String` object whose value states the algorithm  
 278 that the key should be generated for. This parameter is represented by the previously defined  
 279 `algorithm` object. Two of the `getInstance` methods are overloaded with two parameters.  
 280 Since we do not need to specify the second parameter in either method, we substitute it with  
 281 an underscore that serves as a placeholder in one combined pattern definition (Line 18). This  
 282 concept of method event patterns is similar to pointcuts in aspect-oriented programming  
 283 languages such as AspectJ [19]. For CRYSL, we resort to a more lightweight and restricted  
 284 syntax as we found full-fledged pointcuts to be unnecessarily complex. Subsequently, the  
 285 rule defines patterns for the various `init` methods that set the proper parameter values  
 286 (e.g., `keysize`) and a `generateKey` method that completes the key generation and returns  
 287 the generated key.

288 Line 30 defines a usage pattern for `KeyGenerator` using the keyword **ORDER**. The usage  
 289 pattern is a regular expression of method event patterns that are defined in **EVENTS**. Al-  
 290 though each method pattern defines a label to simplify referencing related events (e.g., `g1`,  
 291 `i2`, and `GenKey`), it is tedious and error-prone to require listing all those labels again in  
 292 the **ORDER** section. Therefore, CRYSL allows defining *aggregates*. An aggregate represents  
 293 a disjunction of multiple patterns by means of their labels. Line 19 defines an aggregate  
 294 `GetInstance` that groups the two `getInstance` patterns. Using aggregates, the usage pat-  
 295 tern for `KeyGenerator` reads: there must be exactly one call to one of the `getInstance`  
 296 methods, optionally followed by a call to one of the `init` methods, and finally a call to  
 297 `generateKey`.

298 Following the keyword **CONSTRAINTS**, Lines 33–35 define the constraints for objects  
 299 defined under **OBJECTS** and used as parameters or return values in the **EVENTS** section. In  
 300 the abbreviated CRYSL rule in Figure 4, the first constraint limits the value of `algorithm` to  
 301 "AES" or "Blowfish". For each algorithm, there is one constraint that restricts the possible  
 302 values of `keysize`.

303 The **ENSURES** section is the final mandatory construct in a CRYSL rule. It allows CRYSL  
 304 to support rely/guarantee reasoning. The section specifies predicates to govern interac-  
 305 tions between different classes. For example, a `Cipher` object uses a key obtained from a

<sup>1</sup> As the example shows, in CRYSL, **OBJECTS** also comprise primitive values.

```

9  SPEC javax.crypto.KeyGenerator
10
11  OBJECTS
12    java.lang.String algorithm;
13    int keySize;
14    javax.crypto.SecretKey key;
15
16  EVENTS
17    g1: getInstance(algorithm);
18    g2: getInstance(algorithm, _);
19    GetInstance := g1 | g2;
20
21    i1: init(keySize);
22    i2: init(keySize, _);
23    i3: init(_);
24    i4: init(_, _);
25    Init := i1 | i2 | i3 | i4;
26
27    GenKey: key = generateKey();
28
29  ORDER
30    GetInstance, Init?, GenKey
31
32  CONSTRAINTS
33    algorithm in {"AES", "Blowfish"};
34    algorithm in {"AES"} => keySize in {128, 192, 256};
35    algorithm in {"Blowfish"} => keySize in {128, 192, 256, 320, 384,
36      448};
37
38  ENSURES
39    generatedKey[key, algorithm];

```

■ Figure 4 CRYSL rule for using `javax.crypto.KeyGenerator`.

```

39  SPEC javax.crypto.Cipher
40
41  OBJECTS
42    int encmode;
43    java.security.Key key;
44    java.lang.String transformation;
45    ...
46
47  EVENTS
48    g1: getInstance(transformation);
49    ...
50    i1: init(encmode, key);
51
52    ...
53
54  REQUIRES
55    generatedKey[key, alg(transformation)];
56
57  ENSURES
58    encrypted[cipherText, plainText];

```

■ Figure 5 CRYSL rule for using `javax.crypto.Cipher`.

■ **Table 1** Helper Functions in CRYSL.

| Function   | Purpose                                    |
|--|--|
| <code>alg(<i>transformation</i>)</code>              | Extract algorithm/mode/padding             |
| <code>mode(<i>transformation</i>)</code>             | from <code>transformation</code> parameter |
| <code>padding(<i>transformation</i>)</code>          | of <code>Cipher.getInstance</code> call.   |
| <code>length(<i>object</i>)</code>                   | Retrieve length of <i>object</i>           |
| <code>nevertypeof(<i>object</i>, <i>type</i>)</code> | Forbid <i>object</i> to be of <i>type</i>  |
| <code>callTo(<i>method</i>)</code>                   | Require call to <i>method</i>              |
| <code>noCallTo(<i>method</i>)</code>                 | Forbid call to <i>method</i>               |

306 **KeyGenerator**. The **ENSURES** section specifies what a class guarantees, presuming that the  
 307 object is used properly. For example, the **KeyGenerator** CRYSL rule in Figure 4 ends with  
 308 the definition of a *predicate* `generatedKey` with the generated key object and its corres-  
 309 ponding algorithm as parameters. This predicate may be *required* (i.e., relied on) by the  
 310 rule for **Cipher** or other classes that make use of such a key through the optional element  
 311 of the **REQUIRES** block as illustrated in Figure 5.

312 To obtain the required expressiveness, we have further enriched CRYSL with some  
 313 simple built-in auxiliary functions. For example, in Figure 5, the function `alg` extracts  
 314 the encryption algorithm from `transformation` (Line 55). This function is necessary, be-  
 315 cause `generatedKey` expects only the encryption algorithm as its second parameter, but  
 316 `transformation` optionally specifies also the mode of operation and padding scheme (e.g.,  
 317 Line 6 in Figure 1). For instance, `alg` would extract "AES" from "AES/GCM" or from  
 318 "AES/CBC/PKCS5Padding". Table Table 1 lists all of these functions. Note the last two  
 319 functions `callTo` and `noCallTo` may seem redundant to the **ORDER** and **FORBIDDEN** (see Sec-  
 320 tion 4.3) sections because they appear to fulfil the same purpose of requiring or forbidding  
 321 certain method calls. However, these two functions go beyond that because they allow for  
 322 the specification of conditional forbidden and required methods.

### 323 4.3 Optional Sections in a CrySL Rule

324 A CRYSL rule may contain optional sections that we showcase through the CRYSL rule for  
 325 **PBEKeySpec**. In Figure 6, the **FORBIDDEN** section specifies methods that must *not* be called,  
 326 because calling them is always insecure. **PBEKeySpec** derives cryptographic keys from a  
 327 user-given password. For security reasons, it is recommended to use a cryptographic salt for  
 328 this operation. However, the constructor `PBEKeySpec(char[] password)` does not allow  
 329 for a salt to be passed, and the implementation in the default provider does not generate  
 330 one. Therefore, this constructor should not be called, and any call to it should be flagged.  
 331 Consequently, the CRYSL rule for **PBEKeySpec** lists it in the **FORBIDDEN** section (Line 72).  
 332 In the case of **PBEKeySpec**, there is an alternative secure constructor (Line 68). CRYSL  
 333 allows one to specify an alternative method event pattern using the arrow notation shown  
 334 in Line 72. With **FORBIDDEN** events, CRYSL's language design deviates a bit from its usual  
 335 white-listing approach. We made this choice deliberately to keep specifications concise.  
 336 Without explicit **FORBIDDEN** events, one would have to simulate their effect by explicitly  
 337 listing all events defined on a given type except the one that ought to be forbidden. This  
 338 would significantly increase the size of CRYSL specifications.

339 In general, predicates are generated for a particular usage whenever it does not use any  
 340 **FORBIDDEN** events, its regular **EVENTS** follow the usage pattern defined in the **ORDER** section,

```

59 SPEC javax.crypto.spec.PBEKeySpec
60
61 OBJECTS
62   char[] pw;
63   byte[] salt;
64   int it;
65   int keylength;
66
67 EVENTS
68   create: PBEKeySpec(pw, salt, it, keylength);
69   clear: clearPassword();
70
71 FORBIDDEN
72   PBEKeySpec(char []) => create;
73   PBEKeySpec(char [],byte [],int) => create;
74
75 ORDER
76   create, clear
77   ...
78
79 ENSURES
80   keyspec[this, keylength] after create;
81
82 NEGATES
83   keyspec[this, _];

```

■ **Figure 6** CRYSL rule for `javax.crypto.spec.PBEKeySpec`.

and if the usage fulfils all constraints in the **CONSTRAINTS** section of its corresponding rule. `PBEKeySpec`, however, deviates from that standard. The class contains a constructor that receives a user-given password, but the method `clearPassword` deletes that password later, making it no longer accessible to other objects that might use the key-spec. Consequently, a `PBEKeySpec` object fulfils its role after calling the constructor but only until `clearPassword` is called.

To model this usage precisely, CRYSL allows one to specify a method-event pattern using the keyword **after** (Line 80). If the respective method is called, a predicate is generated. Furthermore, CRYSL supports invalidating an existing predicate in the **NEGATES** section (Line 83). The last call to be made on a `PBEKeySpec` object is the call to `clearPassword` (Line 76). Additionally, the rule lists the predicate `keyspec[this, _]` within the **NEGATES** block. Semantically, the negation of the predicates means the following. A final event in the **ORDER** pattern, in this case a call to `clearPassword`, invalidates the previously generated `keyspec` predicate(s) for `this`. Section 5.2.2 presents the formal semantics of predicates.

## 5 CrySL Formal Semantics

### 5.1 Basic Definitions

A CRYSL rule consists of several sections. The **OBJECTS** section comprises a set of typed variable declarations  $V$ . In the syntax in Figure 3, each declaration  $v \in V$  is represented by the syntax element `TYPE varname`.  $M$  is the set of all resolved method signatures, where each signature includes the method name and argument types. The **EVENTS** section contains elements of the form  $(m, v)$ , where  $m \in M$  and  $v \in V^*$ . We denote the set of all methods referenced in **EVENTS** by  $M$ . The **FORBIDDEN** section lists a set of methods from  $M$  denoted by their signatures; forbidden events cannot bind any variables. The **ORDER** section specifies

364 the usage pattern in terms of a regular expression of labels or aggregates that are in  $M$ ,  
 365 i.e., over the defined **EVENTS**. We express this regular expression formally by the equivalent  
 366 non-deterministic finite automaton  $(Q, M, \delta, q_0, F)$  over the alphabet  $M$ , where  $Q$  is a set of  
 367 states,  $q_0$  is its initial state,  $F$  is the set of accepting states, and  $\delta : Q \times M \rightarrow \mathcal{P}(Q)$  is the  
 368 state transition function.

369 The **CONSTRAINTS** section is a subset of  $\mathbb{C} := (\mathbb{V} \rightarrow \mathcal{O} \cup \mathcal{V}) \rightarrow \mathbb{B}$  (i.e., each constraint is  
 370 a boolean function), where the argument is itself a function that maps variable names in  $\mathbb{V}$   
 371 to objects in  $\mathcal{O}$  or values with primitive types in  $\mathcal{V}$ .

372 A CRYSL rule is a tuple  $(T, \mathcal{F}, \mathcal{A}, \mathcal{C})$ , where  $T$  is the reference type specified by the **SPEC**  
 373 keyword,  $\mathcal{F} \subseteq \mathbb{M}$  is the set of forbidden events,  $\mathcal{A} = (Q, M, \delta, q_0, F) \in \mathbb{A}$  is the automaton  
 374 induced by the regular expression of the **ORDER** section, and  $\mathcal{C} \subseteq \mathbb{C}$  is the set of **CONSTRAINTS**  
 375 that the rule lists. We refer to the set of all CRYSL rules as **SPEC**.

376 Our formal definition of a CRYSL rule does not contain the sections **REQUIRES**, **ENSURES**,  
 377 and **NEGATES**. Those sections reason about the interaction of predicates, whose formal treat-  
 378 ment we discuss in Section 5.2.2.

## 379 5.2 Runtime Semantics

380 Each CRYSL rule encodes usage constraints to be validated for all runtime objects of the  
 381 reference type  $T$  stated in its **SPEC** section. We define the semantics of a CRYSL rule in  
 382 terms of an evaluation over a runtime program trace that records all relevant runtime objects  
 383 and values, as well as all events specified within the rule.

384 ▶ **Definition 1 (Event)**. Let  $\mathcal{O}$  be the set of all runtime objects and  $\mathcal{V}$  the set of all primitive-  
 385 typed runtime values. An *event* is a tuple  $(m, e) \in \mathbb{E}$  of a method signature  $m \in \mathbb{M}$  and  
 386 an *environment*  $e$  (i.e., a mapping  $\mathbb{V} \rightarrow \mathcal{O} \cup \mathcal{V}$  of the parameter variable names to concrete  
 387 runtime objects and values). If the environment  $e$  holds a concrete object for the **this** value,  
 388 then it is called the event's *base object*.

389 ▶ **Definition 2 (Runtime Trace)**. A *runtime trace*  $\tau \in \mathbb{E}^*$  is a finite sequence of events  $\tau_0 \dots \tau_n$ .

390 ▶ **Definition 3 (Object Trace)**. For any  $\tau \in \mathbb{E}^*$ , a subsequence  $\tau_{i_1} \dots \tau_{i_n}$  is called an *object*  
 391 *trace* if  $i_1 < \dots < i_n$  and all base objects of  $\tau_{i_j}$  are identical.

392 Lines 1–2 in Figure 1 result in an object trace that has two events:


393  $(m_0, \{algorithm \mapsto \text{"AES"}, \text{this} \mapsto o_{kg}\})$   
 394  $(m_1, \{algorithm \mapsto \text{"AES"}, keySize \mapsto 128, \text{this} \mapsto o_{kg}\})$   
 395

396 where  $m_0$  and  $m_1$  are the signatures of the **getInstance** and **init** methods of the  
 397 **KeyGenerator** class. For static factory methods such as **getInstance**, we assume that **this**  
 398 is bound to the returned object. We use  $o_{kg}$  to denote that the object  $o$  is bound to the  
 399 variable **kG** at runtime.

400 The decision whether a runtime trace  $\tau$  satisfies a set of CRYSL rules involves two  
 401 steps. In the first step, individual object traces are evaluated independently of one another.  
 402 Yet, different runtime objects may still interact with each other. CRYSL rules capture this  
 403 interaction by means of rely/guarantee reasoning, implemented through predicates that a  
 404 rule ensures on a runtime object. These interactions between different objects are checked  
 405 against the specification in a second step by considering the predicates they require and  
 406 ensure. We first discuss individual object traces in more detail.

$$\begin{aligned}
 sat^o &: \mathbb{E}^* \times \text{SPEC} \rightarrow \mathbb{B} \\
 [\tau^o, (T^o, \mathcal{F}^o, \mathcal{A}^o, \mathcal{C}^o)] &\rightarrow sat_F^o(\tau^o, \mathcal{F}^o) \wedge \\
 &\quad sat_A^o(\tau^o, \mathcal{A}^o) \wedge \\
 &\quad sat_C^o(\tau^o, \mathcal{C}^o)
 \end{aligned}$$

■ **Figure 7** The function  $sat^o$  verifies an individual object trace for the object  $o$ .



■ **Figure 8** The state machine for the CRYSL rule in Figure 4 (without an implicit error state).

### 407 5.2.1 Individual Object Traces

408 The sections **FORBIDDEN**, **ORDER** and **CONSTRAINTS** are evaluated on individual object traces.  
 409 Figure 7 defines the function  $sat^o$  that is true if and only if a given trace  $\tau^o$  for a runtime  
 410 object  $o$  satisfies its CRYSL rule. This definition of  $sat^o$  ignores interactions with other  
 411 object traces. We will discuss later how such interactions are resolved. In the following, we  
 412 assume the trace  $\tau^o = \tau_0^o, \dots, \tau_n^o$ , where  $\tau_i^o = (m_i^o, e_i^o)$ . To illustrate the computation, we will  
 413 also refer to our example from Figure 1 and the involved rules of **KeyGenerator** (Figure 4)  
 414 and **Cipher** (Figure 5). The function  $sat^o$  is composed of three sub-functions:

#### 415 5.2.1.1 Forbidden Events ( $sat_F^o$ )

Given a trace  $\tau^o$  and a set of forbidden events  $\mathcal{F}$ ,  $sat^o$  ensures that none of the trace events is forbidden.

$$sat_F^o(\tau^o, \mathcal{F}^o) := \bigwedge_{i=0..n} m_i^o \notin \mathcal{F}^o$$

416 The CRYSL rule for **KeyGenerator** does not list any forbidden methods. Hence,  $sat^o$   
 417 trivially evaluates to true for object **kG** in Figure 1.

#### 418 5.2.1.2 Order Errors ( $sat_A^o$ )

The second function checks that the trace object is used in compliance with the specified usage pattern (i.e., all methods in the rule are invoked in no other than the specified order). Formally, the sequence of method signatures of the object trace  $m^o := m_0^o, \dots, m_n^o$  (i.e., the projection onto the method signatures) must be an element of the language  $\mathcal{L}(\mathcal{A}^o)$  that the automaton  $\mathcal{A}^o = (Q, \mathbb{M}, \delta, q_0, F)$  of the **ORDER** section induces. By definition of language containment, after the last observed signature of the trace  $m_n^o$ , the corresponding state of the automaton must be an accepting state  $s \in F$ . This definition ignores any variable bindings. They are evaluated in the second step.

$$sat_A^o(\tau^o, \mathcal{A}^o) := m^o \in \mathcal{L}(\mathcal{A}^o)$$

420 Figure 8 displays the automaton created for `KeyGenerator` using the aggregate names as  
 421 labels. State  $0$  is the initial state, and state  $3$  is the only accepting state. Following the code  
 422 in Figure 1 for the object `kG` of type `KeyGenerator`, the automaton transitions from state  $0$   
 423 to  $1$  at the call to `getInstance` (Line 1). With the calls to `init` (Line 2) and `generateKey`  
 424 (Line 3), the automaton first moves to state  $2$  and finally to state  $3$ . Therefore, function  
 425  $sat_{\mathbb{A}}^o$  evaluates to true for this example.

### 426 5.2.1.3 Constraints ( $sat_{\mathbb{C}}^o$ )

The validity check of the constraints ensures that all constraints of  $\mathbb{C}$  are satisfied. This check requires the sequence of environments  $(e_0^o, \dots, e_n^o)$  of the trace  $\tau^o$ . All objects that are bound to the variables along the trace must satisfy the constraints of the rule.

$$sat_{\mathbb{C}}^o(\tau^o, \mathbb{C}^o) := \bigwedge_{c \in \mathbb{C}^o, i=0 \dots n} c(e_i^o)$$

427 To compute  $sat_{\mathbb{C}}^o$  for the `KeyGenerator` object `kG` at the call to `getInstance` in Line 1,  
 428 only the first constraint has to be checked. This is because the corresponding environment  
 429  $e_1^o$  holds a value only for `algorithm`, and the other two constraints reference other vari-  
 430 able names. The evaluation function  $c$  returns true if `algorithm` assumes either "AES" or  
 431 "Blowfish" as its value, which is the case in Figure 1. The computation of  $sat_{\mathbb{C}}^o$  for Lines 2–3  
 432 works similarly.

## 433 5.2.2 Interaction of Object Traces

434 To define interactions between individual object traces, the **REQUIRES**, **ENSURES**, and **NEGATES**  
 435 sections allow individual CRYSL rules to reference one another. For a rule for one object to  
 436 hold at any given point in an execution trace, all predicates that its **REQUIRES** section lists  
 437 must have been both previously *ensured* (by other specifications) and not *negated*. Predicates  
 438 are *ensured* (i.e., generated) and *negated* (i.e., killed) by certain events. Formally, a predicate  
 439 is an element of  $\mathbb{P} := \{(name, args) \mid args \in \mathbb{V}^*\}$  (i.e., a pair of a predicate name and a  
 440 sequence of variable names). Predicates are generated in specific states. Each CRYSL rule  
 441 induces a function  $\mathcal{G}: S \rightarrow \mathcal{P}(\mathbb{P})$  that maps each state of its automaton to the predicate(s)  
 442 that the state generates.

443 The predicates listed in the **ENSURES** and **NEGATES** sections may be followed by the term  
 444 **after**  $n$ , where  $n$  is a method event pattern label or aggregate. The states that follow  
 445 the event or aggregate  $n$  in the automaton generate the respective predicate. If the term  
 446 **after** is not used for a predicate, the final states of the automaton generate (or negate) that  
 447 predicate (i.e., we interpret it as **after**  $n$ , where  $n$  is an event that leads to a final state).

448 In addition to states selected as predicate-generating, the predicate is also ensured if the  
 449 object resides in any state that transitively follows the selected state, unless the states are  
 450 explicitly (de-)selected for the same predicate within the **NEGATES** section. At any state that  
 451 generates a predicate, the event driving the automaton into this state binds the variable  
 452 names to the values that the specification previously collected along its object trace.

453 Formally, an event  $n^o = (m^o, e^o) \in \mathbb{E}$  of a rule  $r$  and for an object  $o$  ensures a predicate  
 454  $p = (predName, args) \in \mathbb{P}$  on the objects  $e^o \in \mathcal{O}$  if:

- 455 1. The method  $m^o$  of the event leads to a state  $s$  of the automaton that generates the  
 456 predicate  $p$  (i.e.,  $p \in \mathcal{G}(s)$ ).
- 457 2. The runtime trace of the event's base object  $o$  satisfies the function  $sat^o$ .

```

84  boolean option1 = isPrime(66); //some non-trivial predicate returning
      false
85  byte[] input = "Message".getBytes("UTF-8");
86
87  String alg = "SHA-256";
88  if (option1) alg = "MD5";
89  MessageDigest md = MessageDigest.getInstance(alg);
90
91  if (input.size() > 0) md.update(input);
92  byte[] digest = md.digest();

```

■ **Figure 9** An example illustrating the usage of `java.security.MessageDigest` in Java.

458 3. All relevant **REQUIRES** predicates of the rule are satisfied at execution of event  $n^o$ .

459 For the `KeyGenerator` object `kG` in Figure 1, a predicate is generated at Line 7 because (1)  
 460 its automaton transitions to its only predicate-generating state (state  $\mathcal{3}$  of the automaton  
 461 in Figure 8), (2)  $sat^o$  evaluates to true as previously shown for each subfunction and (3) the  
 462 corresponding CRYSL rule does not require any predicates.

## 463 6 Detecting Misuses of Crypto APIs

464 To detect all possible rule violations, our tool `COGNICRYPTSAST` approximates the evaluation  
 465 function  $sat^o$  using a static data-flow analysis. In a security context, it is a requirement to  
 466 detect as many misuses as possible. One drawback is the potential for false warnings that  
 467 originate from over-approximations any static analysis requires. In the following, we use the  
 468 example in Figure 9 to illustrate why and where approximations are required. We will show  
 469 later in our evaluation that, in practice, our analysis is highly precise and that the chosen  
 470 approximations rarely actually lead to false warnings.

471 The code example in Figure 9 implements a hashing operation. By default, the code  
 472 uses `SHA-256`. However, if the condition `option1` evaluates to true, `MD5` is chosen instead  
 473 (Line 88). The CRYSL rule for `MessageDigest`, displayed in Figure 10, does not allow the  
 474 usage of `MD5` though, because it is no longer secure [15].

475 The `update` operation is performed only on non-empty input (Line 91). Otherwise,  
 476 the call to `update` is skipped and only the call to `digest` is executed, without any input.  
 477 Although not strictly insecure, this usage does not comply with the CRYSL rule for  
 478 `MessageDigest`, because it leads to no content being hashed.

479 To approximate  $sat^o_P$ , the analysis must search for possible forbidden events by first  
 480 constructing a call graph for the whole program under analysis. It then iterates through the  
 481 graph to find calls to forbidden methods. Depending on the precision of the call graph, the  
 482 analysis may find calls to forbidden methods that cannot be reached at runtime.

483 The analysis represents each runtime object  $o$  by its allocation site. In our example,  
 484 allocation sites are `new` expressions and calls to `getInstance` that return an object of a type  
 485 for which a CRYSL rule exists. For each such allocation site, the analysis approximates  $sat^o_A$   
 486 by first creating a finite-state machine. `COGNICRYPTSAST` then evaluates the state machine  
 487 using a typestate analysis that abstracts runtime traces by program paths. The typestate  
 488 analysis is path-insensitive, thus, at branch points, it assumes that both sides of the branch  
 489 may execute. In our contrived example, this feature leads to a false positive: although  
 490 the condition in Line 91 always evaluates to true, and the call to `update` is never actually  
 491 skipped, the analysis considers that this may happen, and thus reports a rule violation.



```
93 SPEC java.security.MessageDigest
94
95 OBJECTS
96   java.lang.String algorithm;
97   byte[] input;
98   int offset;
99   int length;
100  byte[] hash;
101  ...
102
103 EVENTS
104   g1: getInstance(algorithm);
105   g2: getInstance(algorithm, _);
106   Gets := g1 | g2;
107   ...
108   Updates := ...;
109
110   d1: output = digest();
111   d2: output = digest(input);
112   d3: digest(hash, offset, length);
113   Digests := d1 | d2 | d3;
114
115   r: reset();
116
117 ORDER
118   Gets, (d2 | (Updates+, Digests)), (r, (d2 | (Updates+, Digests)))*
119
120 CONSTRAINTS
121   algorithm in {"SHA-256", "SHA-384", "SHA-512"};
122
123 ENSURES
124   digested[hash, ...];
125   digested[hash, input];
```

■ **Figure 10** CRYSL rule for `java.security.MessageDigest`.

492 To approximate  $sat_C^o$ , we have extended the tpestate analysis to also collect potential  
 493 runtime values of variables along all program paths where an allocated object is used. The  
 494 constraint solver first filters out all *irrelevant* constraints. A constraint is irrelevant if it refers  
 495 to one or more variables that the tpestate analysis has not encountered. In Figure 10, the  
 496 rule only includes one internal constraint—on variable `algorithm`. If we add a new internal  
 497 constraint to the rule about the variable `offset`, the constraint solver will filter it out as  
 498 irrelevant when analyzing the code in Figure 9 because the only method this variable is  
 499 associated with (`digest` labelled `d3`) is never called. The analysis distinguishes between  
 500 never encountering a variable in the source code and not being able to extract the values of  
 501 a variable. With the same rule and code snippet, if the analysis fails to extract the value  
 502 for `algorithm`, the constraint evaluates to false. Collecting potential values of a variable  
 503 over all possible program paths of an allocation site may lead to further imprecision. In  
 504 our example, the analysis cannot statically rule out that `algorithm` may be MD5. The rule  
 505 forbids the usage of MD5. Therefore, the analysis reports a misuse.

506 Handling predicates in our analysis follows the formal description very closely. If  $sat^o$   
 507 evaluates to true for a given allocation site, the analysis checks whether all required pre-  
 508 dicates for the allocation site have been ensured earlier in the program. In the trivial case,  
 509 when no predicate is required, the analysis immediately ensures the predicate defined in the  
 510 **ENSURES** section. The analysis constantly maintains a list of all ensured predicates, including  
 511 the statements in the program that a given predicate can be ensured for. If the allocation  
 512 site under analysis requires predicates from other allocation sites, the analysis consults the  
 513 list of ensured predicates and checks whether the required predicate, with matching names  
 514 and arguments, exists at the given statement. If the analysis finds all required predicates,  
 515 it ensures the predicate(s) specified in the **ENSURES** section of the rule.

## 516 **7** Implementation

517 We have implemented the CRYSL compiler using Xtext [17], an open-source framework for  
 518 developing domain-specific languages as well as the CRYSL- parameterizable static analysis  
 519 COGNICRYPT<sub>SAST</sub>. We have further integrated COGNICRYPT<sub>SAST</sub> with COGNICRYPT [20], in  
 520 which it replaces the original code-analysis component.

### 521 7.1 CrySL

522 Given the CRYSL grammar, Xtext provides a parser, type checker, and syntax highlighter for  
 523 the language. When supplied with a type-safe CRYSL rule, Xtext outputs the corresponding  
 524 AST, which is then used to generate the required static analysis.

525 We developed CRYSL rules for all relevant JCA classes in an iterative process. That is, we  
 526 first worked through the JCA documentation to produce a set of rules and then refined these  
 527 rules through selective discussions with cryptographers and searching security blogs and for-  
 528 ums. In total, we have devised 23 rules covering classes ranging from key handling to digital  
 529 signing. All rules define a usage pattern. Some classes (e.g. `IvParameterSpec`) contain  
 530 one call to a constructor only, while others (e.g. `Cipher`) involve almost ten elements with  
 531 several layers of nesting. Fifteen rules come with parameter constraints, eight of which con-  
 532 tain limitations on cryptographic algorithms. The eight rules without parameter constraints  
 533 are mostly related to classes whose purpose is to set up parameters for specific encryptions  
 534 (e.g. `GCMPParameterSpec`). All rules define at least one **ENSURES** predicate, while only eleven  
 535 require predicates from other rules. Across all rules, we have only declared two methods  
 536 forbidden. We do not find this low number surprising as such methods are always insecure

537 and should not at all be part of a security API. If at all, two forbidden methods is too high a  
 538 number. All rules are available at <https://github.com/CROSSINGTUD/Crypto-API-Rules>.

## 539 7.2 CogniCrypt<sub>sast</sub>

540 COGNICRYPT<sub>SAST</sub> consists of several extensions to the program analysis framework Soot [39,  
 541 21]. Soot transforms a given Java program into an intermediate representation that facilit-  
 542 ates executing intra- and inter-procedural static analyses. The framework provides standard  
 543 static analyses such as call-graph construction. Additionally, Soot can analyze a given An-  
 544 droid app intra-procedurally. Further extensions by FlowDroid [5] enable the construction  
 545 of Android-specific call graphs that are necessary to perform inter-procedural analysis.

546 Validating the **ORDER** section in a CRYSL rule requires solving the typestate check  $sat_A^?$ .  
 547 To this end, we use IDE<sup>al</sup>, a framework for efficient inter-procedural data-flow analysis [37],  
 548 to instantiate a typestate analysis. The analysis defines the finite-state machine  $\mathcal{A}^o$  to check  
 549 against and the allocation sites to start the analysis from. From those allocation sites, IDE<sup>al</sup>  
 550 performs a flow-, field-, and context-sensitive typestate analysis.

551 The constraints and the predicates require knowledge about objects and values associated  
 552 with rule variables at given execution points in the program. The typestate analysis in  
 553 COGNICRYPT<sub>SAST</sub> extracts the primitive values and objects on-the-fly, where the latter are  
 554 abstracted by allocation sites. When the typestate analysis encounters a call site that  
 555 is referred to in an event definition, and the respective rule requires the object or value  
 556 of an argument to the call, COGNICRYPT<sub>SAST</sub> triggers an on-the-fly backward analysis to  
 557 extract the objects or values that may participate in the call. This on-the-fly analysis  
 558 yields comparatively high performance and scalability, because many of the arguments of  
 559 interest are values of type **String** and **Integer**. Thus, using an on-demand computation  
 560 avoids constant propagation of *all* strings and integers through the program. For the on-  
 561 the-fly backward analysis, we extended the on-demand pointer analysis Boomerang [36]  
 562 to propagate both allocation sites and primitive values. Once the typestate analysis is  
 563 completed, and all required queries to Boomerang are computed, COGNICRYPT<sub>SAST</sub> solves  
 564 the internal constraints and predicates using our own custom-made solvers.

565 COGNICRYPT<sub>SAST</sub> may be operated as a standalone command line tool. This way, it takes  
 566 a program as input and produces an error report detailing misuses and their locations. How-  
 567 ever, we have further integrated COGNICRYPT<sub>SAST</sub> into COGNICRYPT [20]. COGNICRYPT  
 568 is a Eclipse plugin, which supports developers in using Crypto APIs by means of scenario-  
 569 based code generation as well code analysis for Crypto APIs. In this context, COGNICRYPT  
 570 translates misuses found by COGNICRYPT<sub>SAST</sub> into standard Eclipse error markers.

## 571 8 Evaluation

572 We evaluate our implementation COGNICRYPT<sub>SAST</sub> using the following research questions:

573 **RQ1:** What are the precision and recall of COGNICRYPT<sub>SAST</sub>?

574 **RQ2:** What types of misuses does COGNICRYPT<sub>SAST</sub> find?

575 **RQ3:** How fast does COGNICRYPT<sub>SAST</sub> run?

576 **RQ4:** How does COGNICRYPT<sub>SAST</sub> compare to the state of the art?

577 To answer these questions, we applied the generated static analysis COGNICRYPT<sub>SAST</sub>  
 578 to 10,000 Android apps from the AndroZoo dataset [4] using our full CRYSL rule set for  
 579 the JCA. We ran our experiments on a Debian virtual machine with sixteen cores and  
 580 64 GB RAM. We chose apps that are available in the official Google Play Store and

581 received an update in 2017. This ensures that we report on the most up-to-date us-  
 582 ages of Crypto APIs. We make available all artefacts at this Github repository: <https://github.com/CROSSINGTUD/paper-crysl-reproducibility-artefacts>.  
 583

## 584 8.1 Precision and Recall (RQ1)

### 585 Setup

586 To compute precision and recall, the first two authors manually checked 50 randomly selected  
 587 apps from our dataset for typestate errors and violations of internal constraints. To collect  
 588 this random sample, we implemented a Java program that generates random numbers using  
 589 `SecureRandom` and retrieved the apps from the corresponding lines in the spreadsheet con-  
 590 taining the results of analysing the 10,000 apps. We did not check for unsatisfied predicates  
 591 or forbidden events, because these are hard to detect manually—while it may seem simple  
 592 to check for calls to forbidden events, it is non-trivial to determine whether or not such  
 593 calls reside in dead code. We compare the results of our manual analysis to those reported  
 594 by `COGNICRYPTSAST`. The goal of this evaluation is to compute precision and recall of the  
 595 analysis implementation in `COGNICRYPTSAST`, not the quality of our CRYSL rules. We dis-  
 596 cuss the latter in Section 8.4. Consequently, we define a false positive to be a warning that  
 597 should not be reported according to the specified rule, irrespective of that rule’s semantic  
 598 correctness. Similarly, a false negative would arise if `COGNICRYPTSAST` missed to report a  
 599 misuse that, according to the CRYSL rule, does exist in the analyzed program.

### 600 Results

601 In the 50 apps we inspected, `COGNICRYPTSAST` detects 228 usages of JCA classes. Table 2  
 602 lists the misuses that `COGNICRYPTSAST` finds (156 misuses in total). In particular, `COG-`  
 603 `NICRYPTSAST` issues 27 typestate-related warnings, with only 2 false positives. Both arise  
 604 because the analysis is path-insensitive (Section 6). We further found 4 false negatives that  
 605 are caused by initializing a `MessageDigest` or a `MAC` object without completing the opera-  
 606 tion. `COGNICRYPTSAST` fails to find these typestate errors because the supporting off-the-  
 607 shelf alias analysis `Boomerang` times out, causing `COGNICRYPTSAST` to abort the typestate  
 608 analysis without reporting a warning for the object at hand. A larger timeout or future  
 609 improvements to the alias analysis `Boomerang` would avoid this problem.

610 The automated analysis finds 129 constraint violations. We were able to confirm 110  
 611 of them. In the other 19 cases, highly obfuscated code causes the analysis to fail to ex-  
 612 tract possible runtime values statically. For such values, the constraint solver reports the  
 613 corresponding constraint as violated. A better handling of such highly obfuscated code can  
 614 be enabled by techniques complementary to ours. For instance, one could augment `COG-`  
 615 `NICRYPTSAST` with the hybrid static/dynamic analysis `Harvester` [32]. We have also checked  
 616 the apps for missed constraint violations (false negatives), but were unable to find any.

617 **RQ1:** In our manual assessment, the typestate analysis achieves high precision (92.6%) and  
 618 recall (86.2%). The constraint resolution has a precision of 85.3% and a recall of 100%.

■ **Table 2** Correctness of COGNICRYPT<sub>SAST</sub> warnings.

|             | Total Warnings | False Positives | False Negatives |
|-------------|----------------|-----------------|-----------------|
| Typestate   | 27             | 2               | 4               |
| Constraints | 129            | 19              | 0               |
| Total       | 156            | 21              | 4               |

## 8.2 Types of Misuses (RQ2)

### Setup

We report findings obtained by analyzing all our 10,000 Android apps from AndroZoo [4]. We then use the results of our manual analysis (Section 8.1) as a baseline to evaluate our findings on a large scale.

COGNICRYPT<sub>SAST</sub> detects the usage of at least one JCA class in 8,422 apps. Further investigation unveiled that many of these usages originate from the same common libraries included in the applications. To avoid counting the same crypto usages twice, and to prevent over-counting, we exclude usages within packages `com.android`, `com.facebook.ads`, `com.google` or `com.unity3d` from the analysis.

### Results

Excluding the findings in common libraries, COGNICRYPT<sub>SAST</sub> detects the usage of at least one JCA class in 4,349 apps (43% of the analyzed apps). Most of these apps (95%) contain at least one misuse. Across all apps, COGNICRYPT<sub>SAST</sub> started its analysis for a total of 40,295 allocation sites (i.e., abstract objects). Of these, a total of 20,426 individual object traces violate at least one part of the specified rule patterns. COGNICRYPT<sub>SAST</sub> reports typestate errors (**ORDER** section in the rule) for 4,708 objects, and reports a total of 4,443 objects to have unsatisfied predicates (i.e., the object expected a predicate from another object as listed in the **REQUIRES** block of a rule). The analysis also discovered 97 reachable call sites that call forbidden events. The majority of object traces that violate at least one part of a CRYSL rule (54.7%) contradict a constraint listed in the **CONSTRAINTS** section of a rule.

Approximately 86% of these constraint-violations are related to **MessageDigest**. In our manual analysis (see RQ1), 89 of the 110 found constraint violations originated from usages of **MD5** and **SHA-1**. We expect a similar fraction to also hold for the 11,178 constraint contradictions reported over all 10,000 apps. Many developers still use **MD5** and **SHA-1**, although both are no longer recommended by security experts [15]. COGNICRYPT<sub>SAST</sub> identifies 1,228 (10.9%) constraint violations related to **Cipher** usages. In our manual analysis, all misuses of the **Cipher** class are due to using the insecure algorithm **DES** or the **ECB** mode of operation. This result is in line with the findings of prior studies [13, 35, 12].

More than 75% of the typestate errors that COGNICRYPT<sub>SAST</sub> issues are caused by misuses of **MessageDigest**. Our manual analysis attributes this high number to incorrect usages of the method `reset()`. In addition to misusing **MessageDigest**, misuses of **Cipher** contribute 766 typestate errors. Finally, COGNICRYPT<sub>SAST</sub> detects 157 typestate errors related to **PBEKeySpec**. The **ORDER** section of the CRYSL rule for **PBEKeySpec** requires calling `clearPassword()` at the end of the lifetime of a **PBEKeySpec** object. We manually inspected 3 of the misuses and observed that the call to `clearPassword()` is missing in all of them.

Predicates are unsatisfied when COGNICRYPT<sub>SAST</sub> expects the interaction of multiple

656 object traces but is not able to prove their correct interaction. With 4,443 unsatisfied  
 657 predicates reported, the number may seem relatively large, yet one must keep in mind that  
 658 unsatisfied predicates accumulate transitively. For example, if `COGNICRYPTSAST` cannot  
 659 ensure a predicate for a usage of `IVParameterSpec`, it will not generate a predicate for the  
 660 key object that `KeyGenerator` generates using the `IVParameterSpec` object. Transitively,  
 661 `COGNICRYPTSAST` reports an unsatisfied predicate also for any `Cipher` object that relies on  
 662 the generated key object.

663 `COGNICRYPTSAST` also found 97 calls to forbidden methods. Since only two JCA classes  
 664 require the definition of forbidden methods in our CrySL rule set (`PBEKeySpec` and `Cipher`),  
 665 we do not find this low number surprising. A manual analysis of a handful of reports suggests  
 666 that most of the reported forbidden methods originate from calling the insecure `PBEKeySpec`  
 667 constructors, as we explained in Section 4.

668 From the 4,349 apps that use at least one JCA Crypto API, 2,896 apps (66.6%) contain at  
 669 least one tpestate error, 1,367 apps (31.4%) lack required predicates, 62 apps (1.4%) call at  
 670 least one forbidden method, and 3,955 apps (90.9%) violate at least one internal constraint.  
 671 Ignoring the class `MessageDigest`, and hereby excluding MD5 and SHA-1 constraints, 874  
 672 apps still violate at least one constraint in other classes.

673 **RQ2:** Approximately 95% of apps misuse at least one Crypto API. Violating the constraints  
 674 of `MessageDigest` is the most common type of misuse.

### 675 8.3 Performance (RQ3)

#### 676 Setup


677 `COGNICRYPTSAST` comprises four main phases. It constructs (1) a *call graph* using Flow-  
 678 Droid [5] and then runs the actual analysis (Section 6), which (2) calls the *tpestate analysis*  
 679 and (3) *constraint analysis* as required, attempting to (4) *resolve all declared predicates*.  
 680 During the analysis of our dataset, we measured the execution time that `COGNICRYPTSAST`  
 681 spent in each phase. We ran `COGNICRYPTSAST` once per application and capped the time of  
 682 each run to 30 minutes.

683 In Section 8.2, we report that `COGNICRYPTSAST` found usages of the JCA in 4,349 of  
 684 all 10,000 apps in our dataset. If we include in the reporting those usages that arise from  
 685 misuses within the common libraries previously excluded (see Section 8.2), this number rises  
 686 to 8,422. We include the analysis of the libraries in this part of the evaluation because it helps  
 687 evaluate the general performance of the analysis in the worst case when whole applications  
 688 are analyzed.

#### 689 Results

690 Figure 11 summarizes the distribution of analysis times for the four phases and the total  
 691 analysis time across these 8,422 apps. For each phase, the box plot highlights the median,  
 692 the 25% and 75% quartiles, and the minimal and maximal values of the distribution.

693 Across the apps in our dataset, there is a large variation in the reported execution time  
 694 (10 seconds to 28.6 minutes). We attribute this variation to the following reasons. The  
 695 analyzed apps have varying sizes—the number of reachable methods in the call graph varies  
 696 between 116 and 16,219 (median: 3,125 methods). The majority of the total analysis time  
 697 (83%) is spent on call-graph construction. For the remaining three phases of the analysis,  
 698 the distribution is as follows. Across all apps, the resolution of all declared predicates takes



■ **Figure 11** Analysis time (in log scale) of the individual phases of `COGNICRYPTSAST` when running on the apps that use the JCA.

699 approximately a median of 50 milliseconds, and the typestate analysis phase takes a median  
 700 of 500 milliseconds. The median for the constraint phase is 350 milliseconds. Therefore, the  
 701 major bottleneck for the analysis is call-graph construction, a problem orthogonal to the  
 702 one we address in this work. Our analysis itself is efficient and the overall analysis time is  
 703 clearly dominated by the runtime of the call-graph construction.

704

**RQ3:** On average, `COGNICRYPTSAST` analyzes an app in 101 seconds, with call-graph construction taking most of the time (83%).

705

## 706 8.4 Comparison to Existing Tools (RQ4)

### 707 Setup

708 We compare `COGNICRYPTSAST` to `CRYPTOLINT` [13], as we explained in Section 2.3 the most  
 709 closely related tool. Unfortunately, despite contacting the authors we were unable to obtain  
 710 access to `CRYPTOLINT`'s implementation. We thus resorted to reimplementing the original  
 711 rules that are hard-coded in `CRYPTOLINT` as `CRYSL` rules. The fact that all `CRYPTOLINT`  
 712 rules can be modelled in `CRYSL` shows its superior expressiveness.

713 In this section, `RULESETFULL` denotes `COGNICRYPT`'s comprehensive `CRYSL` rules that  
 714 we have created for all the JCA classes, while `RULESETCL` denotes the set of `CRYSL` rules  
 715 that we developed to model the original `CRYPTOLINT` rules. Additionally, `COGNICRYPTSAST`  
 716 denotes our analysis when it runs using `RULESETFULL`, and `COGNICRYPTCL` denotes the  
 717 analysis when it runs using `RULESETCL`.

718 `RULESETFULL` consists of 23 rules, one for each class of the JCA. `RULESETCL` comprises only  
 719 six individual rules, and they only use the sections **ENSURES**, **REQUIRES** and **CONSTRAINTS**. In  
 720 other words, the original hard-coded `CRYPTOLINT` rules do not comprise typestate properties  
 721 nor forbidden methods. For three out of six rules, we managed to exactly capture the  
 722 semantics of the hard-coded `CRYPTOLINT` rule in a respective `CRYSL` rule. The remaining  
 723 three rules (3, 4, and 6 of the original `CRYPTOLINT` rules) cannot be perfectly expressed as  
 724 a `CRYSL` rule, and our `CRYSL`-based rules over-approximate them instead.

725 `CRYPTOLINT` rule 4, for instance, requires salts in `PBEKeySpec` to be non-constant. In  
 726 `CRYSL`, such a relationship is expressed through predicates. Predicates in `CRYSL`, however,  
 727 follow a white-listing approach and therefore only model correct behaviour. Therefore, in  
 728 `CRYSL` we model the `CRYPTOLINT` rule for `PBEKeySpec` in a stricter manner, requiring the  
 729 salt to be not just non-constant but truly random, i.e., returned from a proper random  
 730 generator. We followed a similar approach with the other two `CRYPTOLINT` rules that

731 we modelled in CRYSL. In result, RULESET<sub>CL</sub> is stricter than the original implementation  
 732 of CRYPTO<sub>LINT</sub>. In the comparison of COGNICRYPT<sub>SAST</sub> and COGNICRYPT<sub>CL</sub> in terms of  
 733 their findings, the stricter rules produce more warnings than the original implementation of  
 734 CRYPTO<sub>LINT</sub>. In our comparison against COGNICRYPT<sub>SAST</sub>, this setup favours CRYPTO<sub>LINT</sub>  
 735 because we assume that these additional findings to be true positives. Both rule sets are  
 736 available at <https://github.com/CROSSINGTUD/Crypto-API-Rules>.

## 737 Results

738 COGNICRYPT<sub>CL</sub> detects usages of JCA classes in 1,866 Android apps. For these apps, COG-  
 739 NICRYPT<sub>CL</sub> reports 5,507 misuses, only a third of the 20,426 misuses that COGNICRYPT<sub>SAST</sub>  
 740 identifies using RULESET<sub>FULL</sub>, our more comprehensive rule set.

741 Using COGNICRYPT<sub>CL</sub>, all reported warnings are related to 6 classes, compared to 23  
 742 classes that are specified in RULESET<sub>FULL</sub>. As we have pointed out, CRYPTO<sub>LINT</sub> does not  
 743 specify any tpestate properties or forbidden methods. Hence, COGNICRYPT<sub>CL</sub> does not find  
 744 the 4,805 warnings that COGNICRYPT<sub>SAST</sub> identifies in these categories using RULESET<sub>FULL</sub>.  
 745 Furthermore, while COGNICRYPT<sub>SAST</sub> reports 11,178 constraint violations with the standard  
 746 rules, COGNICRYPT<sub>CL</sub> reports only 1,177 constraint violations. Of the 11,178 constraint  
 747 violations, 9,958 are due to the rule specification for the class `MessageDigest`. CRYPTO<sub>LINT</sub>  
 748 does not model this class. If we remove these violations, 1,609 violations are still reported  
 749 by COGNICRYPT<sub>SAST</sub>, a total of 432 more than by COGNICRYPT<sub>CL</sub>.

750 We compare our findings to the study by Egele et al. [13] that identifies the use of ECB  
 751 mode as a common misuse of cryptography. In that study, 7,656 apps use ECB (65.2% of  
 752 apps that use Crypto APIs). On the other hand, in our study, COGNICRYPT<sub>CL</sub> identified  
 753 663 uses of ECB mode in 35.5% of apps that use Crypto APIs. Although a high number of  
 754 apps still exhibit this basic misuse, there is a considerable decrease (from 65.2% to 35.5%)  
 755 compared to the previous study by Egele et al. [13]. Given that all apps in our study must  
 756 have received an update in 2017, we believe that the decrease of misuses reflects taking  
 757 software security more seriously in today’s app development.

758 Based on the high precision (92.6%) and recall (96.2%) values discussed in **RQ1**, we argue  
 759 that COGNICRYPT<sub>SAST</sub> provides an analysis with a much higher recall than CRYPTO<sub>LINT</sub>.  
 760 Although the larger and more comprehensive rule set, RULESET<sub>FULL</sub>, detects more complex  
 761 misuses, the precise analysis keeps the false-positive rate at a low percentage.

762 **RQ4:** The more comprehensive RULESET<sub>FULL</sub> detects 3× as many misuses as CRYPTO<sub>LINT</sub>  
 763 in almost 4× more JCA classes.

## 764 8.5 Threats to Validity

765 Our ruleset RULESET<sub>FULL</sub> is mainly based on the documentation of the JCA [18]. Although  
 766 we have significant domain expertise, our CRYSL-rule specifications for the JCA are only  
 767 as correct as the JCA documentation. Our static-analysis toolchain depends on multiple  
 768 external components and despite an extensive set of test cases, of course, we cannot fully  
 769 rule out bugs in the implementation.

770 Java allows a developer to programmatically select a non-default cryptographic service  
 771 provider. COGNICRYPT<sub>SAST</sub> currently does not detect such customizations but instead as-  
 772 sumes that the default provider is used. This behaviour may lead to imprecise results because  
 773 our rules forbid certain default values that are insecure for the default provider, but may be  
 774 secure if a different one is chosen.



## 9 Conclusion

In this paper, we present CRYSL, a description language for correct usages of cryptographic APIs. Each CRYSL rule is specific to one class, and it may include usage pattern definitions and constraints on parameters. Predicates model the interactions between classes. For example, a rule may generate a predicate on an object if it is used successfully, and another rule may require that predicate from an object it uses. We also present a compiler for CRYSL that transforms a provided ruleset into an efficient and precise data-flow analysis  $\text{COGNICRYPT}_{\text{SAST}}$  checking for compliance according to the rules. For ease of use, we have integrated  $\text{COGNICRYPT}_{\text{SAST}}$  and with Eclipse crypto assistant COGNICRYPT. Applying  $\text{COGNICRYPT}_{\text{SAST}}$ , the analysis for our extensive ruleset  $\text{RULESET}_{\text{FULL}}$ , to 10,000 Android apps, we found 20,426 misuses spread over 95% of the 4,349 apps using the JCA.  $\text{COGNICRYPT}_{\text{SAST}}$  is also highly efficient: for more than 75% of the apps the analysis finishes in under 3 minutes, where most of the time is spent in Android-specific call graph construction.

In future work, we plan to address the following challenges. We have developed all the rules used in  $\text{COGNICRYPT}_{\text{SAST}}$  ourselves. While we have acquired some deeper familiarity with cryptographic concepts in general and the JCA in particular, we are not cryptographers. Therefore, we are open to and want cryptography experts to correct potential mistakes in our existing rules. We would further encourage domain experts to model their own cryptographic libraries in CRYSL to improve the support in  $\text{COGNICRYPT}_{\text{SAST}}$  and, by extension, COGNICRYPT. CRYSL currently only supports a binary understanding of security—a usage is either secure or not. We would like to enhance CRYSL to have a more fine-grained notion of security to allow for more nuanced warnings in  $\text{COGNICRYPT}_{\text{SAST}}$ . This is challenging because the CRYSL language still ought to be concise. Additionally, CRYSL currently requires one rule per class per JCA provider, because there is no way to express the commonality and variability between different providers implementing the same algorithms, leading to specification overhead. To address this issue, we plan to modularize the language using import and override mechanisms. Moreover, we plan to extend CRYSL to support more complex properties such as using the same cryptographic key for multiple purposes. We will also perform consistency checks for the CRYSL rules. For now, only Xtext-based type checks are performed.

Lastly, we also intend on applying CRYSL in other contexts. One of the authors of this paper has already started to have students implement a dynamic checker to identify and mitigate violations at runtime. While the JCA is indeed the most commonly used Crypto library, other Crypto libraries such as BouncyCastle [29] are being used as well and we will to extend  $\text{COGNICRYPT}_{\text{SAST}}$  to support them. Additionally, we will investigate to which extent CRYSL is applicable to Crypto APIs in other programming languages. At the time of writing, we are exploring CRYSL's compatibility with OpenSSL [30]. We finally aim to examine whether CRYSL is expressive enough to meaningfully specify usage constraints for non-crypto APIs.

## Acknowledgments

This work was supported by the DFG through its Collaborative Research Center CROSSING, the project RUNSECURE, by the Natural Sciences and Engineering Research Council of Canada, by the Heinz Nixdorf Foundation, a Fraunhofer ATTRACT grant, and by an Oracle Collaborative Research Award. We would also like to thank the maintainers of AndroZoo for allowing us to use their data set in our evaluation.

## 820 References

- 821 1 Y. Acar, C. Stransky, D. Wermke, C. Weir, M. L. Mazurek, and S. Fahl. Developers  
822 need support, too: A survey of security advice for software developers. In *2017 IEEE*  
823 *Cybersecurity Development (SecDev)*, pages 22–26, Sept 2017. 10.1109/SecDev.2017.17.
- 824 2 Dima Alhadidi, Amine Boukhtouta, Nadia Belblidia, Mourad Debbabi, and Prabir Bhat-  
825 tacharya. The dataflow pointcut: a formal and practical framework. In *Proceedings of*  
826 *the 8th International Conference on Aspect-Oriented Software Development, AOSD 2009,*  
827 *Charlottesville, Virginia, USA, March 2-6, 2009*, pages 15–26, 2009.
- 828 3 Chris Allan, Pavel Avgustinov, Aske Simon Christensen, Laurie J. Hendren, Sascha  
829 Kuzins, Ondrej Lhoták, Oege de Moor, Damien Sereni, Ganesh Sittampalam, and Julian  
830 Tibble. Adding trace matching with free variables to aspectj. In *Proceedings of the*  
831 *20th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems,*  
832 *Languages, and Applications, OOPSLA 2005, October 16-20, 2005, San Diego, CA,*  
833 *USA*, pages 345–364, 2005. 10.1145/1094811.1094839.
- 834 4 Kevin Allix, Tegawendé F. Bissyandé, Jacques Klein, and Yves Le Traon. Androzoos:  
835 collecting millions of android apps for the research community. In *Proceedings of the*  
836 *13th International Conference on Mining Software Repositories, MSR 2016, Austin, TX,*  
837 *USA, May 14-22, 2016*, pages 468–471, 2016.
- 838 5 Steven Arzt, Siegfried Rasthofer, Christian Fritz, Eric Bodden, Alexandre Bartel,  
839 Jacques Klein, Yves Le Traon, Damien Octeau, and Patrick D. McDaniel. Flowdroid:  
840 precise context, flow, field, object-sensitive and lifecycle-aware taint analysis for android  
841 apps. In *ACM SIGPLAN Conference on Programming Language Design and Imple-*  
842 *mentation, PLDI '14, Edinburgh, United Kingdom - June 09 - 11, 2014*, pages 259–269,  
843 2014.
- 844 6 John W. Backus, Friedrich L. Bauer, Julien Green, C. Katz, John McCarthy, Alan J.  
845 Perlis, Heinz Rutishauser, Klaus Samelson, Bernard Vauquois, Joseph Henry Wegstein,  
846 Adriaan van Wijngaarden, Michael Woodger, and Peter Naur. Revised report on the  
847 algorithm language ALGOL 60. *Communications of the ACM*, 6(1):1–17, 1963.
- 848 7 Kevin Bierhoff and Jonathan Aldrich. Modular typestate checking of aliased objects.  
849 In *Proceedings of the 22nd Annual ACM SIGPLAN Conference on Object-Oriented Pro-*  
850 *gramming, Systems, Languages, and Applications, OOPSLA 2007, October 21-25, 2007,*  
851 *Montreal, Quebec, Canada*, pages 301–320, 2007. 10.1145/1297027.1297050.
- 852 8 Eric Bodden. Efficient hybrid typestate analysis by determining continuation-equivalent  
853 states. In *ICSE '10: International Conference on Software Engineering*, pages 5–14,  
854 New York, NY, USA, May 2010. ACM. ISBN 978-1-60558-719-6.
- 855 9 Eric Bodden. TS4J: a fluent interface for defining and computing typestate analyses.  
856 In *Proceedings of the 3rd ACM SIGPLAN International Workshop on the State Of the*  
857 *Art in Java Program analysis, SOAP 2014, Edinburgh, UK, Co-located with PLDI 2014,*  
858 *June 12, 2014*, pages 1:1–1:6, 2014.
- 859 10 Eric Bodden, Patrick Lam, and Laurie Hendren. Partially evaluating finite-state runtime  
860 monitors ahead of time. *ACM Transactions on Programming Languages and Systems*  
861 *(TOPLAS)*, 34(2):7:1–7:52, June 2012.
- 862 11 VeraCode (CA). State of software security 2017. [https://info.veracode.com/  
863 report-state-of-software-security.html](https://info.veracode.com/report-state-of-software-security.html), 2017.
- 864 12 Alexia Chatzikonstantinou, Christoforos Ntantogian, Georgios Karopoulos, and Christos  
865 Xenakis. Evaluation of cryptography usage in android applications. In *International*  
866 *Conference on Bio-inspired Information and Communications Technologies*, pages 83–  
867 90, 2016.

- 868 **13** Manuel Egele, David Brumley, Yanick Fratantonio, and Christopher Kruegel. An em-  
869 pirical study of cryptographic misuse in android applications. In *ACM Conference on*  
870 *Computer and Communications Security*, pages 73–84, 2013.
- 871 **14** Felix Fischer, Konstantin Böttinger, Huang Xiao, Christian Stransky, Yasemin Acar,  
872 Michael Backes, and Sascha Fahl. Stack overflow considered harmful? the impact of  
873 copy&paste on android application security. In *2017 IEEE Symposium on Security and*  
874 *Privacy, SP 2017, San Jose, CA, USA, May 22-26, 2017*, pages 121–136, 2017.
- 875 **15** German Federal Office for Information Security (BSI). Cryptographic mechanisms: Re-  
876 commendations and key lengths. Technical Report BSI TR-02102-1, BSI, January 2017.
- 877 **16** Simon Goldsmith, Robert O’Callahan, and Alexander Aiken. Relational queries over  
878 program traces. In *Proceedings of the 20th Annual ACM SIGPLAN Conference on*  
879 *Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2005,*  
880 *October 16-20, 2005, San Diego, CA, USA*, pages 385–402, 2005.
- 881 **17** Xtext home page. <http://www.eclipse.org/Xtext/>, 2017.
- 882 **18** Oracle Inc. Java Cryptography Architecture (JCA) Reference Guide.  
883 [https://docs.oracle.com/javase/8/docs/technotes/guides/security/crypto/](https://docs.oracle.com/javase/8/docs/technotes/guides/security/crypto/CryptoSpec.html)  
884 [CryptoSpec.html](https://docs.oracle.com/javase/8/docs/technotes/guides/security/crypto/CryptoSpec.html), 2017.
- 885 **19** Gregor Kiczales, Erik Hilsdale, Jim Hugunin, Mik Kersten, Jeffrey Palm, and William  
886 Griswold. An overview of aspectj. *ECOOP 2001—Object-Oriented Programming*, pages  
887 327–354, 2001.
- 888 **20** Stefan Krüger, Sarah Nadi, Michael Reif, Karim Ali, Mira Mezini, Eric Bodden, Florian  
889 Göpfert, Felix Günther, Christian Weinert, Daniel Demmler, and Ram Kamath. Cog-  
890 niCrypt: Supporting Developers in Using Cryptography. In *Proceedings of the 32nd*  
891 *IEEE/ACM International Conference on Automated Software Engineering, ASE 2017,*  
892 *Urbana, IL, USA, October 30 - November 03, 2017*, pages 931–936, 2017.
- 893 **21** Patrick Lam, Eric Bodden, Ondřej Lhoták, and Laurie Hendren. The Soot framework  
894 for Java program analysis: a retrospective. In *Cetus Users and Compiler Infrastructure*  
895 *Workshop (CETUS 2011)*, October 2011.
- 896 **22** David Lazar, Haogang Chen, Xi Wang, and Nickolai Zeldovich. Why does cryptographic  
897 software fail?: a case study and open problems. In *ACM Asia-Pacific Workshop on*  
898 *Systems (APSys)*, pages 7:1–7:7, 2014.
- 899 **23** V. Benjamin Livshits and Monica S. Lam. Finding security vulnerabilities in java ap-  
900 plications with static analysis. In *Proceedings of the 14th USENIX Security Symposium,*  
901 *Baltimore, MD, USA, July 31 - August 5, 2005*, 2005.
- 902 **24** Michael C. Martin, V. Benjamin Livshits, and Monica S. Lam. Finding application  
903 errors and security flaws using PQL: a program query language. In *Proceedings of the*  
904 *20th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems,*  
905 *Languages, and Applications, OOPSLA 2005, October 16-20, 2005, San Diego, CA,*  
906 *USA*, pages 365–383, 2005.
- 907 **25** David A. McGrew and John Viega. The security and performance of the galois/counter  
908 mode (GCM) of operation. In *Progress in Cryptology - INDOCRYPT 2004, 5th In-*  
909 *ternational Conference on Cryptology in India, Chennai, India, December 20-22, 2004,*  
910 *Proceedings*, pages 343–355, 2004.
- 911 **26** Clint Morgan, Kris De Volder, and Eric Wohlstadter. A static aspect language for  
912 checking design rules. In *Proceedings of the 6th International Conference on Aspect-*  
913 *Oriented Software Development, AOSD 2007, Vancouver, British Columbia, Canada,*  
914 *March 12-16, 2007*, pages 63–72, 2007.

- 915 **27** Sarah Nadi, Stefan Krüger, Mira Mezini, and Eric Bodden. Jumping through hoops:  
 916 why do Java developers struggle with cryptography APIs? In *International Conference*  
 917 *on Software Engineering (ICSE)*, pages 935–946, 2016.
- 918 **28** Nomair A. Naeem and Ondrej Lhoták. Typestate-like analysis of multiple interacting ob-  
 919 jects. In *Proceedings of the 23rd Annual ACM SIGPLAN Conference on Object-Oriented*  
 920 *Programming, Systems, Languages, and Applications, OOPSLA 2008, October 19-23,*  
 921 *2008, Nashville, TN, USA*, pages 347–366, 2008.
- 922 **29** Legion of the Bouncy Castle Inc. BouncyCastle, 2018. [https://www.bouncycastle.](https://www.bouncycastle.org/java.html)  
 923 [org/java.html](https://www.bouncycastle.org/java.html).
- 924 **30** OpenSSL. OpenSSL - Cryptography and SSL/TLS Toolkit, 2018. [https://www.](https://www.openssl.org/)  
 925 [openssl.org/](https://www.openssl.org/).
- 926 **31** Siegfried Rasthofer, Steven Arzt, Robert Hahn, Max Kohlhagen, and Eric Bodden.  
 927 (in)security of backend-as-a-service. In *BlackHat Europe 2015*, November 2015.
- 928 **32** Siegfried Rasthofer, Steven Arzt, Marc Miltenberger, and Eric Bodden. Harvesting  
 929 runtime values in android applications that feature anti-analysis techniques. In *Network*  
 930 *and Distributed System Security Symposium (NDSS)*, February 2016.
- 931 **33** Martin P. Robillard, Eric Bodden, David Kawrykow, Mira Mezini, and Tristan Ratch-  
 932 ford. Automated API property inference techniques. *IEEE Transactions on Software*  
 933 *Engineering (TSE)*, 39:613–637, 2013.
- 934 **34** Martin P. Robillard, Eric Bodden, David Kawrykow, Mira Mezini, and Tristan Ratch-  
 935 ford. Automated api property inference techniques. *IEEE TOSEM*, 39(5):613–637, May  
 936 2013. ISSN 0098-5589. 10.1109/TSE.2012.63.
- 937 **35** Shuai Shao, Guowei Dong, Tao Guo, Tianchang Yang, and Chenjie Shi. Modelling ana-  
 938 lysis and auto-detection of cryptographic misuse in Android applications. In *International*  
 939 *Conference on Dependable, Autonomic and Secure Computing*, pages 75–80, 2014.
- 940 **36** Johannes Späth, Lisa Nguyen Quang Do, Karim Ali, and Eric Bodden. Boomerang:  
 941 Demand-driven flow- and context-sensitive pointer analysis for java. In *30th European*  
 942 *Conference on Object-Oriented Programming, ECOOP 2016, July 18-22, 2016, Rome,*  
 943 *Italy*, pages 22:1–22:26, 2016.
- 944 **37** Johannes Späth, Karim Ali, and Eric Bodden. *Ide<sup>al</sup>*: Efficient and precise alias-aware  
 945 dataflow analysis. In *2017 International Conference on Object-Oriented Programming,*  
 946 *Languages and Applications (OOPSLA/SPLASH)*. ACM Press, October 2017. To ap-  
 947 pear.
- 948 **38** Robert E. Strom and Shaula Yemini. Typestate: A programming language concept  
 949 for enhancing software reliability. *IEEE Trans. Software Eng.*, 12(1):157–171, 1986.  
 950 10.1109/TSE.1986.6312929. URL <https://doi.org/10.1109/TSE.1986.6312929>.
- 951 **39** Raja Vallée-Rai, Etienne Gagnon, Laurie J. Hendren, Patrick Lam, Patrice Pominville,  
 952 and Vijay Sundaresan. Optimizing java bytecode using the soot framework: Is it feasible?  
 953 In *Compiler Construction*, pages 18–34, 2000.