

# Deductive Verification of Active Objects with **Crowbar**

Eduard Kamburjan<sup>a</sup>, Marco Scaletta<sup>b</sup>, Nils Rollshausen<sup>b</sup>

<sup>a</sup>*University of Oslo, Oslo, Norway*

<sup>b</sup>*Technische Universität Darmstadt, Darmstadt, Germany*

---

## Abstract

We present **Crowbar**, a deductive verification tool for the Active Object language **ABS**. **Crowbar** implements novel specification approaches specifically for distributed systems. For user interaction, counterexamples are presented as executable programs. **Crowbar** has a modular structure to explore further approaches, and was applied in the largest Active Objects verification study.

*Keywords:* Deductive Verification, Symbolic Execution, Active Objects

---

## 1. Introduction

2 Deductive verification of functional properties is a static analysis tech-  
3 nique that uses *program logics* to verify the behavior of programs against  
4 user-provided specification. Provers implementing such program logics, for  
5 example the **KeY** [1] prover, have been successfully used to detect bugs in well-  
6 tested libraries [2, 3] for sequential programs. **KeY** implements *heavyweight*  
7 *symbolic execution (SE)*, where the program is transformed into first-order  
8 formulas by keeping track of symbolic values throughout program execution.  
9 Heavyweight symbolic execution is able to deal with unbounded systems by  
10 additional user-provided specification, for example through additional invari-  
11 ants that have to be preserved by loops.

12 For distributed systems, more complex specification languages than for  
13 sequential programs are needed. The specific concurrency model plays a ma-  
14 jor role in the design of specification languages and program logics, and in this  
15 work we present a heavyweight symbolic execution tool for Active Objects.

---

*Email addresses:* `eduard@ifi.uio.no` (Eduard Kamburjan),  
`scaletta@cs.tu-darmstadt.de` (Marco Scaletta),  
`nils.rollshausen@stud.tu-darmstadt.de` (Nils Rollshausen)

16 Active Objects [4] are an object-oriented, actor-based concurrency model for  
17 distributed systems, developed specifically with a focus on analyzability and  
18 implemented in the **ABS** language [5]. **ABS** has been applied as a modeling tool  
19 to a multitude of domains, ranging from cloud systems [6, 7, 8, 9], over rail-  
20 way operations [10] and memory systems [11] to computational biology [12].

21 A prototypical extension of **KeY**, **KeY-ABS** [13], showed that the general  
22 ideas of heavyweight symbolic execution indeed carry over to **ABS** [14]. How-  
23 ever, new specification approaches [15, 16, 17], program logics [18, 19], inte-  
24 gration with static analyses [20] and further case studies [21] have revealed  
25 that the structure of **KeY** with its tight integration of Java-specifics is not  
26 suited to handle (a) the specification languages needed for Active Objects  
27 and (b) parts of the **ABS** language that clash with Java specifics. For exam-  
28 ple, the functional sublanguage of **ABS**, exceptions and method contracts are  
29 not supported. Consequently, **KeY-ABS** supports only object invariants and a  
30 small subset of **ABS** – namely the one which has similar semantics to Java.

31 **Crowbar** is the an alternative system, implemented from scratch and  
32 based on the Behavioral Program Logic (BPL) [18]. It covers full **coreABS**<sup>1</sup>,  
33 delegates part of the proof obligation to external static analyses and uses the  
34 newly developed specification approaches for Active Objects, such as coop-  
35 erative contracts [17]. **Crowbar** is not interactive, i.e., the user cannot see  
36 the proof in the program logic. To realize feedback, it instead integrates a  
37 counterexample generation that presents failed proof branches completely in  
38 terms of the program: the counterexample is output as an **ABS** program that  
39 is executable and contains only the statements needed to reproduce a run  
40 described by the failed proof branch. Thus, the user is not exposed to the  
41 underlying program logic. To accommodate the need for quick prototyping  
42 of new specification approaches, **Crowbar** has a modular structure suited for  
43 prototypical exploration of novel approaches: it is easy to add new BPL  
44 calculi. **Crowbar** was used for the biggest Active Object verification case  
45 study [22], that goes beyond the capabilities of prior systems in (1) language  
46 coverage, (2) complexity of specification, and (3) lines of code.

47 *Targeted Problem..* **Crowbar** can verify functional properties of Active Ob-  
48 jects specified by cooperative method contracts [17] and local session types for  
49 Active Objects [15], both formalized in the Behavioral Program Logic [18]. It  
50 implements delegating parts of the proof obligation to external static analy-

---

<sup>1</sup>**ABS** without its extensions for, e.g., timed models.

51 ses and has a modular structure that enables easy extendability to implement  
52 further behavioral specifications, which we describe in more details below.

53 *Structure.* This article is structured as follows. Sec. 2 introduces Active  
54 Objects, ABS, the used specification languages, symbolic execution and gives  
55 an example. Sec. 3 describes the structure of Crowbar. Sec. 4 describes the  
56 implementation in more detail, the aforementioned case study and compares  
57 Crowbar with KeY-ABS and other related tools. Sec. 5 gives an example of the  
58 usage of Crowbar and the counterexample generator before Sec. 6 concludes.

59 For the technical documentation we refer to the documentation of Crowbar  
60 at <https://github.com/Edkamb/crowbar-tool/wiki>, the theoretical back-  
61 ground is referred to in the corresponding parts of Sec. 2.

## 62 2. Background

63 To verify the safety of a program one must specify its expected behavior  
64 and translate the specified program into a set of *proof obligations*. If all proof  
65 obligations can be discharged, then the program is *safe*, where under safety  
66 we understand (local) partial correctness [23]: if every process terminates,  
67 then the program behaves as specified and no process throws an exception.

68 For distributed systems, with their inherit non-determinism, it is crucial  
69 that the specification is *modular*. This means that changes in one part of  
70 the program should not invalidate all proof obligations. Thus, the design of  
71 specification techniques is a balancing act between strong abstraction with  
72 high modularity and weak abstraction with high expressive power [5]. ABS  
73 and Crowbar are specifically designed to explore this trade-off by integrating  
74 different approaches in one system to compare them in practice. Crowbar is  
75 also modular in its approach to integration and allows to interact with type  
76 systems and static analyses.

### 77 2.1. Active Objects and ABS

78 Active Objects are an actor-based, object-oriented concurrency model  
79 which realizes strong encapsulation. At its core, an Active Object program  
80 consists of a set of objects, which communicate with each other using asyn-  
81 chronous method calls, and futures to retrieve the return value of a method  
82 call. In each object, at most one process is active at any time, meaning that  
83 there is no interleaving within an objects. The objects are preemption-free:

84 a process must explicitly deschedule itself before another can become active.  
85 Objects can be created at runtime.

86 In more details, it is based on the following features for concurrency:

87 **Strong Encapsulation.** Every object is strongly encapsulated at runtime,  
88 such that no other object can access its fields, not even objects of the  
89 same class.

90 **Asynchronous Calls with Futures.** Each method call to another object  
91 is asynchronous and generates a future. Futures can be passed around  
92 and are used to synchronize on the process generated by the call. Once  
93 the called process terminates, its future is *resolved* and the return value  
94 can be retrieved. We say that the process *computes* its future.

95 **Cooperative Scheduling.** At every point in time, at most one process is  
96 active per object and a running process cannot be interrupted unless it  
97 *explicitly* releases the object. This is done either by termination with a  
98 **return** statement or with an **await** *g* statement that waits until guard *g*  
99 holds. A guard polls whether a future is resolved or whether a boolean  
100 condition holds.

101 Concurrent systems are challenging for SE, but Active Objects allow to per-  
102 form *local* SE on single methods through their strong encapsulation and de-  
103 coupling of caller and callee processes. Special care, however, has to be taken  
104 to keep track of futures and correct handling of state when using **await**. We  
105 introduce the ABS language using an example to demonstrate these features.

106 *ABS*. The Abstract Behavioral Specification language (**ABS**) is an implemen-  
107 tation of Active Objects with a rich toolkit of static analyses. It supports  
108 extensions of Active Objects such as product lines [24], timed [25], hybrid [26],  
109 resource [27] models or open systems [28]. The fragment which only imple-  
110 ments the core Active Objects features is referred to as **coreABS**. **Crowbar**  
111 only supports **coreABS**<sup>2</sup>. Beyond the object-oriented language, **ABS** has a sim-  
112 ple (no function passing) functional sublanguage for data processing. For a  
113 detailed tutorial we refer to the online material of **ABS**<sup>3</sup>.

---

<sup>2</sup>It also supports product lines in the sense that after a variant is generated, **Crowbar** can verify it. It does not perform a family-based analysis [29].

<sup>3</sup><http://abs-models.org>

114 *2.2. Specification*

115 Specification of Active Objects must take concurrency into account, even  
116 method contracts require special attention due to, among others, the time gap  
117 between calling a method and starting its execution. For **Crowbar** we support  
118 state object invariants, behavioral method contracts and local session types.  
119 **Crowbar** specifications are part of the input **ABS** file using *Spec-annotations*:  
120 Each **ABS** statement and definition **s** can be prefixed with an annotation  
121 using the **[T: e]s** syntax, where **T** is a type and **e** an expression. **Spec** defines  
122 the data type for specifications, which must be provided as the expression  
123 of the annotation. **Spec-annotations** are ignored at runtime. Additionally,  
124 loops are annotated with loop invariants using **WhileInv**.

125 *Behavioral Method Contracts.* A method can be annotated with **Ensures** and  
126 **Requires** as post and preconditions. In interfaces, these specifications can  
127 only contain parameters (and **result**, the special variable for the return  
128 value). For an example, we refer to Fig. 1.

129 Following the principles behind the Java Modeling Language (JML) [30],  
130 **Crowbar** supports **old** to refer to the pre-state of the method and **last** to  
131 refer to the pre-state of the last suspension, i.e., the state before the last  
132 **await** statement was executed, or the pre-state of the method if no **await**  
133 was executed yet. Additionally, **Crowbar** supports the post and preconditions  
134 at **await** suspension statements, as well as **Succeeds** and **Overlaps** context  
135 sets [17]: if the precondition of a method contains assertions about the heap,  
136 then it is not clear which method is responsible to establish it — due to  
137 the concurrency model, the caller has no control over the fields of the callee  
138 object. Context sets specify for a method **m** the following: the methods in  
139 **Succeeds** *must* have run and *must establish* the precondition as their post-  
140 condition. The methods in **Overlaps** *may* have run and *must preserve* the  
141 precondition.

142 *Object Invariants.* Objects are annotated with creation conditions (**Requires**  
143 ) and object invariants (**ObjInv**). A creation condition describes the param-  
144 eters of the constructor (analogously to asynchronous contracts), while the  
145 object invariant has to hold after the constructor terminate and whenever a  
146 process is scheduled or descheduled.

147 *Local Session Types.* Lastly, **Crowbar** supports a variant of local session  
148 types. A local session type for **ABS** is represented as a string and specifies calls,

149 synchronization, sequence, repetition and alternative. Alternative is speci-  
150 fied using `+`, repetition with `*`, suspension with `Susp( $\varphi$ )` (where  $\varphi$  specifies  
151 the state before suspension, and analogously for calls `!`) and synchronization  
152 with `Get(e)`, where `e` is the targeted expression. Session types specify how  
153 *roles* in a protocol communicate. The mapping of roles to fields is specified  
154 with `[Spec: Role("name", this.field)]`. For details we refer to [18, 31, 15].

155 For example, the following expresses that the statement `s` first calls  
156 method `m` on role `f` and then `n` on role `g`. Finally, the last action is a **return**  
157 in a state where `result == 0` holds (i.e., the return value is zero). No other  
158 communication, synchronization or suspension happens. This is annotated  
159 using `[Spec: Local("f!m.g!n.Put(result == 0)")]`. The systems are inde-  
160 pendent: it is not necessary to specify a local type.

### 161 2.3. Examples

162 *Example Contracts and Invariants.* We give a short example on **ABS** and its  
163 specification below. For brevity's sake, we only give asynchronous method  
164 contracts and the object invariant.

165 Consider Lst. 1. The class **Monitor** has two fields: `s` which is a server that  
166 is monitored and `beats`, a counter for successful requests to `s`. The method  
167 `heartbeat` sends a request to itself (l. 7) by an asynchronous method call  
168 (by using `!`). Afterwards, the return value of the call is retrieved (l. 8, using  
169 **get**). This blocks the process until the `HttpRequest` process has terminated.  
170 No other process can run on this object until this happens. If the request  
171 was successful, `beats` is increased by 1. Method `reset` waits *without blocking*  
172 until the number of success reaches a passed threshold and resets `beats`.  
173 Synchronous calls are possible (l. 10) on **this**.

174 Specifications are annotations of the form `[Spec: KIND(e)]`, where `KIND`  
175 is the used specification pattern. Fig. 1 gives an overview over the keywords  
176 for the specification patterns and where to annotate them. In the exam-  
177 ple, a creation condition and an object invariant are specified for the class  
178 **Monitor**. They state that the passed server must be non-null and stays non-  
179 null throughout execution (l. 1). The method contract states that the `beats`  
180 field is increased. This is not the case, and we return to the fault in Sec. 5.

181 *Example Local Session Types.* Fig. 2 shows an example how the specifications  
182 work together when using local session types. The class **C** uses a protocol with  
183 three roles, that are declared in the class header using `Role`. Additionally,  
184 an object invariant is used to declare that all the fields are non-null. The

```

1 [Spec:Requires(this.s != null)] [Spec:ObjInv(this.s != null)]
2 class Monitor(Server s) {
3   Int beats = 0;
4   [Spec:Ensures(this.beats >= old(this.beats) &&
5     result == this.beats)]
6   Int heartbeat() {
7     Fut<Int> req = s!httpRequest();
8     Int status = req.get;
9     if(status == 200) { this.beats = this.beats + 1;}
10    else { this.handleError(); }
11    return this.beats;
12  }
13  Unit reset(Int i){ await this.beats == i; this.beats = 0; }
14  Unit handleError() { this.beats = 0; /* ... */ }
15 }

```

Listing 1: An example ABS program with specification.

185 sole shown method, `getExpLocalAliasing`, is using a type where as the first  
186 action the `client` is called on method `a`. The connection of the field `this.`  
187 `c` and the role `client` is established through the aforementioned class-level  
188 specification. The invariant is needed to show that no exception is thrown.  
189 The other actions specify a synchronization on the future stored in variable  
190 `f` and a termination without a specific post-condition.

#### 191 2.4. Symbolic Execution

192 Symbolic execution describes the execution of a program (or statement)  
193 with *symbolic values*. A symbolic value is a placeholder and can be described  
194 by condition collecting during the symbolic execution. Heavyweight sym-  
195 bolic execution is used as a *proof strategy* for a sequent calculus of first-order  
196 dynamic logics in, e.g., JavaDL [1], ABSDL [32] or DTL [33] and has success-  
197 fully applied to discover highly involved bugs in non-concurrent libraries of  
198 mainstream languages [2, 3]. One of the shortcomings of symbolic execution  
199 with dynamic logics is that they first fully symbolically execute the program  
200 and then evaluate the post-condition. For distributed systems the specifica-  
201 tion, however, often contains a temporal element and can be partially checked  
202 already *during* symbolic execution.

203 The Behavioral Program Logic (BPL) [18] is a generalization of dynamic

```

1 [Spec: Role("server", this.s)][Spec: Role("client", this.c)]
2 [Spec: Role("database", this.d)]
3 [Spec: ObjInv(this.s != null && this.c != null && this.d != null)]
4 class C(Server s, Client c, Database d) {
5   [Spec:Local("client!a(true).Get(f).Put(true)")]
6   Unit getExpLocalAliasing() {
7     Fut<Int> f = this.c!a();
8     Fut<Int> sth = f;
9     Int a = sth.get;
10  }
11 ...
12 }

```

Figure 2: An example ABS program with a local session type.

204 logic. It uses *behavioral* modalities, which we informally introduce now,  
205 to enable such symbolic execution strategies. BPL is based on behavioral  
206 specifications  $\mathbb{T} = (\alpha_{\mathbb{T}}, \tau_{\mathbb{T}})$ . Set  $\tau_{\mathbb{T}}$  is the set of terms of the specification,  
207 and function  $\alpha_{\mathbb{T}}$  provides the semantics: it maps elements of  $\tau_{\mathbb{T}}$  to trace  
208 formulas. Behavioral specification are referred to in the logic using behavioral  
209 modalities, which have the form  $[s \Vdash^{\alpha_{\mathbb{T}}} \tau]$ , with  $\tau \in \tau_{\mathbb{T}}$  being a specification  
210 term. Its semantics expresses partial correctness: a state  $\sigma$  satisfies the  
211 modality, if every trace  $\theta$  of a normally terminating run of  $s$  from  $\sigma$  is a  
212 model for the trace formula  $\alpha_{\mathbb{T}}(\tau)$ . We omit  $\alpha$  from examples for brevity.

213 A sequent has the form  $\Gamma \Rightarrow \{U\}[s \Vdash^{\alpha} \tau], \Delta$ , where  $\Gamma$  and  $\Delta$  are two sets  
214 of formulas, and represents a symbolic state, where  $s$  is the statement left to  
215 symbolically execute,  $U$  is the state update (i.e., a syntactic representation  
216 of accumulated substitutions [34]), and  $\tau$  is the specification term, e.g., the  
217 post-condition, and  $\bigwedge \Gamma \wedge \neg \bigvee \Delta$  describes the accumulated knowledge and  
218 path condition. An example rule for symbolic execution in sequent calculi  
219 is the following rule, that expresses a split over the branches of a branching  
220 statement for post-conditions  $\varphi$ .

$$221 \frac{\Gamma, \{U\}e \Rightarrow \{U\}[s_1 \Vdash \varphi], \Delta \quad \Gamma, \{U\}\neg e \Rightarrow \{U\}[s_2 \Vdash \varphi], \Delta}{\Gamma \Rightarrow \{U\}[\mathbf{if}(e) \ s_1 \ \mathbf{else} \ s_2 \ \Vdash \varphi], \Delta}$$

222 Behavioral specifications separate *syntax* of the specification ( $\tau$ ) and its  
223 semantics as a trace specification ( $\alpha$ ). This distinction enables behavioral



224 symbolic execution: the *design* of  $\tau$  can now aim to have a simple proof  
 225 calculus that is not restricted by the structure of the logic underlying  $\alpha$ .  
 226 Furthermore,  $\tau$  can serve as an interface to external analyses.

227 Established calculi for dynamic logics perform symbolic execution by re-  
 228 ducing the statement inside a modality without considering the specification.  
 229 In contrast, for each step in a behavioral type system [35], the statement is  
 230 matched with the current specification. BPL combines both: a logical frame-  
 231 work with a behavioral type-style calculus, which we call *guided SE*.

232 For example, the local session type given above is expressed with the  
 233 modality  $[s \Vdash \mathbf{f}!m.\mathbf{g}!n. \downarrow \mathbf{result} == 0]$ . One (slightly prettified) rule for  
 234 session types is the following. The first premise checks that the role and field  
 235 coincide, the second premise *reduces the type* during the symbolic execution  
 236 step. Note that this premise contains no modality – we call such branches  
 237 *side-branches*. We stress that the conclusion of the rule has to match (1) the  
 238 call in specification and statement and (2) the method names both *syntacti-*  
 239 *cally*. If matching fails, SE stops.

$$240 \frac{\Gamma \Rightarrow \{U\} \mathbf{this.f} \doteq \mathbf{r}, \Delta \quad \Gamma \Rightarrow \{U\} [s \Vdash L], \Delta}{\Gamma \Rightarrow \{U\} [\mathbf{this.f}!m(); s \Vdash \mathbf{r}!m.L], \Delta}$$

### 241 3. Software Framework

242 At its core, **Crowbar** is a *heavyweight* symbolic execution (SE) engine, i.e.,  
 243 it uses contracts and loop invariants to build a SE tree that abstracts from *all*  
 244 possible runs. As discussed the used program logic allows for *guided SE*: the  
 245 specification is used to guide the construction of the SE tree. Construction  
 246 may abort if there is not possible that any further execution may satisfy  
 247 the specification. For example, if the specification expresses that the first  
 248 interaction is a call to a method  $m$ , but the first call is to a method  $n$ , then  
 249 guided SE will immediately abort the proof. The leaves of the SE tree, after  
 250 SE successfully finished, are logical formulas that are passed to SMT-solvers.

251 Fig. 3 illustrates guided SE using our example for Session Types. The  
 252 node marked with (1) is not generated, as the type system ensures that **this.f**  
 253 is always non-**null** (due to being annotated with **NonNull**). The other dashed  
 254 nodes are omitted because the method is specified to make two calls, but  
 255 executes three. It does not check any steps after the second call as these are  
 256 already following a wrong execution path. We omitted (a) the update, (b) the  
 257 collected path condition and (c) the role check branches in the illustration.

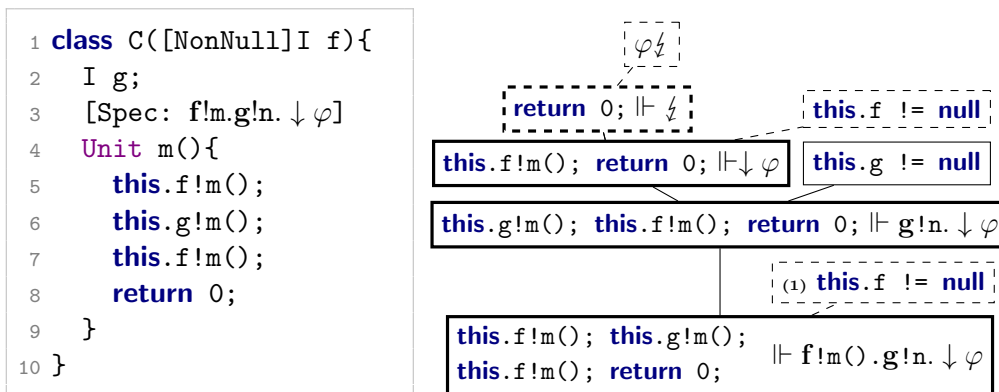


Figure 3: Illustration of guided SE. Verification of the method to the left results in a symbolic execution tree where the dashed nodes are not generated.

### 258 3.1. Software Functionalities

259 **Crowbar** implements the three specification paradigms described in Sec. 2  
260 and the **assert** statement for **ABS**: *object invariants* [36], predicates that have  
261 to hold at every point a process gains or loses control over its object, *co-*  
262 *operative contracts* [17], a generalization of method contracts to distributed  
263 systems, and *local Session Types* [37, 16, 15], a protocol language for allowed  
264 communication actions.

265 It generates proof obligations in BPL and has two additional mechanisms  
266 to interact with the outside: (1) *Counterexample Generation*: If a proof obli-  
267 gation fails, all failed proof branches are translate back into an **ABS** program,  
268 using the values extracted from the SMT solver proof attempt at the leaf.  
269 This allows the user to investigate the failure without being exposed to the  
270 underlying program logic. (2) *Static Nodes*: Cooperative Contracts and Ses-  
271 sion Types rely on additional mechanisms to guarantee safety of composition  
272 (propagation for contracts and projection for Session Types). These mecha-  
273 nisms are external to the program logic and **Crowbar**, thus, outputs a *static*  
274 *node* to communicate that the program is safe, if these external conditions  
275 hold.

276 A verification attempt with **Crowbar** outputs either (1) yes, (2) yes with  
277 external condition or (3) no with counterexample (if generation succeeds).

278 **Crowbar** supports pre-/post-conditions for the functional sublanguage.  
279 **Crowbar** also integrates results directly from the **ABS** compiler: the AST has  
280 nullability annotations, and expressions are not checked for **null**-access if the  
281 type system already ensures this.




Figure 4: Structure of Crowbar.

282 *3.2. Software Architecture*

283 **Crowbar** has a pipeline setup with front-end, middle-end and back-end, as  
 284 shown in Fig. 4.

285 *Front-end.* The front-end uses the ABS parser to generate an annotated AST  
 286 and extracts, per method, one statement  $s$  and one specification term  $\tau$  for  
 287 this method. The statement is translated into an internal representation  
 288 (IR) to normalize the AST. For example, each call has a target variable. The  
 289 specification term language depends on the specification mechanism chosen  
 290 by the user. The front-end furthermore sets up a program repository to  
 291 manage the specification and connects the IR with the original.

292 *Middle-end.* The middle-end implements SE using the chosen set of rules.  
 293 When SE finishes, all leaves of the SE trees are either *static nodes*, discussed  
 294 above, *logical nodes*, which are first-order formulas whose validity ensures  
 295 that this branch is safe. Guided SE is realized by matching the current  
 296 specification term on the current program and reducing *both* in one SE step.

297 *Back-end.* The back-end of **Crowbar** passes logical nodes to external SMT-  
 298 LIB solvers. If all logical nodes can be proven to be valid, then the program is  
 299 considered safe up to external restrictions, which are output as static nodes.  
 300 If one of the logical nodes fails to be proven valid, **Crowbar** attempts to  
 301 construct a *counterexample program* [38] by extracting value from the coun-  
 302 terexample model output by the SMT solver and reconstructing a minimal  
 303 runnable program directly from the path of the SE tree taken to this leaf. If  
 304 a SMT model is not available, counterexample generation fails. Generating  
 305 the SMT-LIB input requires the repository to ensure correct typing.

## 306 4. Implementation and Empirical Results

307 *Implementation.* **Crowbar** is implemented in Kotlin in 5500 lines of SLoC (ac-  
308 cording to `cloc` [39]), and a ANTLRv4 parser for local Session Types. The  
309 build system is using gradle, the used testing framework is kotest. User man-  
310 ual and developer documentation on adding a new specification/verification  
311 module is available in the github repository. We performed two evaluations:  
312 For a performance evaluation, we use the `absexamples`<sup>4</sup> repository, where  
313 case studies and examples from the development of ABS are collected. We  
314 have loaded all 647 methods which are fall into the CoreABS fragment and  
315 used Crowbar to verify the default specification. As the default specification  
316 is not meaningful, we additionally adopted two bigger examples to CoreABS  
317 (`WaterTank.abs`, 60 LoC, `chat.abs`, 307 LoC), specified and verified that no  
318 exceptions are thrown (and a simple functional invariant of the water tank:  
319 water level never drops below 0). Benchmarking was run on a 8-core i7-8565U  
320 CPU with 1.80GHz and 32GB RAM on a Ubuntu 20.04.5 laptop. On aver-  
321 age, 10 symbolic execution steps are performed (with 171 symbolic execution  
322 steps for the biggest), which on average needs  $\sim 110ms$  (with  $1650ms$ ).

323 As **Crowbar** requires a fully specified program to return meaningful re-  
324 sults, we performed the following case study for functional correctness for a  
325 qualitative evaluation.

326 *Case Study.* **Crowbar** is used in the biggest verification case study for Active  
327 Objects [22]: A model extracted from C code [40] with 260 lines of ABS code,  
328 with 5 classes (with invariants and creation conditions), 5 interfaces (with 19  
329 method contracts) and 1 function with a contract. The verification succeeds  
330 fully automatically. In contrast, the previously biggest case study [14] has  
331 140 LoC for 1 class (with invariant) and requires manual interaction.

332 *Coverage and Comparison with KeY-ABS.* **Crowbar** covers full `coreABS`, i.e.,  
333 ABS without its extensions for time or variability. Specification in **Crowbar** is  
334 purely in terms of the program, i.e., using expressions, using the specification  
335 patterns described in Sec. 2.

336 Table 1 gives the syntax for specifications and compares **Crowbar** and  
337 KeY-ABS with respect to language and specification coverage. Note that KeY-  
338 ABS does not support several statements of ABS, such as `case`, `throw`, `try`,

---

<sup>4</sup><https://github.com/abstools/absexamples>

Feature	Support	Specification	KeY-ABS	SN <sup>1</sup>
Asynchronous Contr.	Yes	<b>Requires, Ensures</b> on methods in interfaces	Yes <sup>§</sup>	No
Synchronous Contr.	Yes	<b>Requires, Ensures</b> on methods in classes	No	Yes
Cooperative Contr.	Yes	<b>Succeeds, Overlaps</b> on methods and <b>await</b> <b>Resolves</b> on <b>get</b>	No	Yes
Object Invariants	Yes	<b>Requires, ObjInv</b> on classes	Yes <sup>f</sup>	No
Function Contracts	Yes	<b>Requires, Ensures</b> on functions	No	No
Loop Invariants	Yes	<b>WhileInv</b> on loops	No	No
Session Types	Partial <sup>a</sup>	<b>Role</b> on classes, <b>Local</b> on methods	No	Yes
<b>old</b> and <b>last</b>	Yes		No	-
Counterexamples	Yes		No	-
History Specification	Partial <sup>b</sup>		Yes	-
First-order logic Specification	No		Yes	-
Exceptions and <b>assert</b>	Yes		Partial <sup>‡</sup>	-
Functions and ADTs	Yes		No	-

<sup>a</sup> No passive choice and exceptions.    <sup>b</sup> Can be partially encoded by hand.    <sup>1</sup> SN = static nodes.

<sup>§</sup> Encodable.    <sup>†</sup> No creation condition.    <sup>‡</sup> Only null access

Table 1: Overview over feature support in **Crowbar** for coreABS.

339 and **assert**.<sup>5</sup> For a detailed discussions on the limitations on the underly-  
340 ing ABSDL logic we refer to Kamburjan [31, Ch.2]. A further point worth  
341 mentioning is that KeY-ABS does not support specification within the pro-  
342 gram and takes it as additional input. In contrast, **Crowbar** supports all  
343 this and connects specification and program tightly by using annotations for  
344 specifications.

345 *Limitations.* **Crowbar** is limited to partial correctness of coreABS. Further-  
346 more, **Crowbar** does not support passive choice and exceptions in local types,  
347 which are an open research question. Additionally, it does not support first-  
348 order specifications and specification of the history of events, which however  
349 can be added manually by the user: Histories can be added by introducing  
350 a special ADT for events and a field of list type in every class that managed  
351 explicitly. Due to its design, **Crowbar** rely on the counterexample generation  
352 to interact between user and proof, as it relies on an automatic SMT solver  
353 for its backend. Additionally, **Crowbar** can output the symbolic execution  
354 tree, as well as the concrete SMT-LIB output.

---

<sup>5</sup>A reason for the limitations is that it reuses KeY-Java internally, which do not support, for example, functional structures, and encode assumptions that are valid for Java, but not ABS directly in the code. An example for the last point is that in ABS, classes cannot be extended, while Java supports code inheritance.

```

1 class CeFrame {
2   Int beats = 21239; String s = "o_3";
3   Unit ce() { // Snippet from: heartbeat
4     String req = "fut_1";
5     Int status = 5; //Int status = (req).get
6     if((status == 200)){
7       else { // this.handleError();
8         //Assume following effects while blocked:
9         this.beats = 21238;
10    }
11    // Evaluates to: 21238
12    println(toString(this.beats));
13    // Failed precondition:
14    //   (heap.beats >= old.beats)
15    //   ^ (result = heap.beats)
16    // Failed to show the sub-obligations:
17    // (select(anon(heap), this.beats)
18    // >= heap.beats)
19  }
20 }

```

Figure 5: An example for counterexample generation with Crowbar.

355 *Related Tools.* We already discussed KeY-ABS in detail above. Chisel [41] is  
356 a tool for *hybrid* Active Objects that generates proof obligations in  $d\mathcal{L}$  [42],  
357 which in turn has three implementations as KeY-style symbolic execution  
358 engines [43] for a minimal hybrid programming language.

359 Outside the KeY-family, the Why3 [44] and Boogie [45] frameworks pro-  
360 vide deductive verification middle-ends based on SE. Their separation of  
361 verification technique and programming language makes them not suited for  
362 our situation: we target a specific concurrency model with a *tight* coupling  
363 of specification and verification.

364 For Active Objects, Rebeca [46] supports model checking [47], which is  
365 limited to bounded systems and supports no modular specifications. Re-  
366 garding deductive verification, KeY-ABS and Crowbar are the only systems  
367 for Active Objects. For actors, there exists a proposed program logic by  
368 Gordon [48] and a number of static analyses for Erlang, for which we refer to  
369 an overview in the work of Bagherzadeh et al. [49]. Note, however, that pure  
370 actors are not object-oriented, and that they do not implement cooperative  
371 scheduling or futures.

## 372 5. Illustrative Example

373 We give an example how to specify, verify and investigate a program and  
374 continue with the code specified in Lst. 1. The following attempts verifica-

375 tion:

```
376 > crowbar example.abs --method Module.Monitor.heartbeat -inv
```

377 The switch `--method` verifies a single method and requires its fully qual-  
378 ified name, in this case `Module.Monitor.heartbeat`. The proof attempt  
379 fails: There is no contract given for `handleError`, so it is not specified how  
380 the method changes the field `heartbeat`. The `-inv` flag activates the coun-  
381 terexample generator, which outputs the counterexample shown in Fig. 5. All  
382 context interactions of the method (calls, suspension, synchronization, etc.)  
383 are removed to ensure that it can be executed on its own. Furthermore, it  
384 adds comments about extracted values and what interactions have been re-  
385 moved. Here, it shows that if the value of `heartbeat` was 21239 before and  
386 `handleError` changes it to 21238, then a part of the post-condition does not  
387 hold. The code is executable: as all context is removed, the programmer may  
388 now examine and manipulate the counterexample using standard debugging  
389 techniques.

## 390 6. Conclusion

391 *Impact.* **Crowbar** is an important step in the empirical research on analy-  
392 sis of Active Objects: it is the first verification tool to cover full **coreABS**  
393 and implements novel specification and verification techniques in a flexible  
394 framework that allows one to investigate further new approaches with little  
395 overhead. In particular, we are now able to verify feature-rich programs, such  
396 as the C extraction case study. For potential impact, **Crowbar** enables us to  
397 tackle the numerous ABS case studies focusing on their specification without  
398 modifying them to fit the very restricted fragment supported by prior tools.

399 *Future Work.* We plan to use **Crowbar** to implement novel heavyweight sym-  
400 bolic execution systems, in particular the probabilistic dynamic logic of Pardo  
401 et al. [50], for which development has started, and a concurrent setting for the  
402 dynamic logic for memory access patterns [51]. As for planned extensions,  
403 we plan to integrate delta-oriented verification [52] next, as well as integrate  
404 first-order specifications and automate the handling of histories.

405 **Bibliography**

- 406 [1] W. Ahrendt, B. Beckert, R. Bubel, R. Hähnle, P. H. Schmitt, M. Ul-  
407 brich (Eds.), *Deductive Software Verification - The KeY Book - From*  
408 *Theory to Practice*, Vol. 10001 of LNCS, Springer, 2016. doi:10.1007/  
409 978-3-319-49812-6.
- 410 [2] S. de Gouw, J. Rot, F. S. de Boer, R. Bubel, R. Hähnle, *Openjdk's*  
411 *java.utils.collection.sort() is broken: The good, the bad and the worst*  
412 *case*, in: *CAV (1)*, Vol. 9206 of LNCS, Springer, 2015, pp. 273–289.
- 413 [3] S. de Gouw, F. S. de Boer, R. Bubel, R. Hähnle, J. Rot, D. Steinhöfel,  
414 *Verifying OpenJDK's sort method for generic collections*, *J. Autom.*  
415 *Reason.* 62 (1) (2019) 93–126.
- 416 [4] F. S. de Boer, V. Serbanescu, R. Hähnle, L. Henrio, J. Rochas, C. C.  
417 Din, E. B. Johnsen, M. Sirjani, E. Khamespanah, K. Fernandez-Reyes,  
418 A. M. Yang, *A survey of active object languages*, *ACM Comput. Surv.*  
419 50 (5) (2017) 76:1–76:39. doi:10.1145/3122848.
- 420 [5] E. B. Johnsen, R. Hähnle, J. Schäfer, R. Schlatte, M. Steffen, *ABS:*  
421 *A core language for abstract behavioral specification*, in: *FMCO*,  
422 Vol. 6957 of LNCS, Springer, 2010, pp. 142–164. doi:10.1007/  
423 978-3-642-25271-6\_8.
- 424 [6] G. Turin, A. Borgarelli, S. Donetti, E. B. Johnsen, S. L. T. Tar-  
425 ifa, F. Damiani, *A formal model of the kubernetes container frame-*  
426 *work*, in: T. Margaria, B. Steffen (Eds.), *ISoLA*, Vol. 12476 of *Lec-*  
427 *ture Notes in Computer Science*, Springer, 2020, pp. 558–577. doi:  
428 10.1007/978-3-030-61362-4\_32.
- 429 [7] J. Lin, I. C. Yu, E. B. Johnsen, M. Lee, *ABS-YARN: A formal framework*  
430 *for modeling hadoop YARN clusters*, in: P. Stevens, A. Wasowski (Eds.),  
431 *FASE*, Vol. 9633 of *Lecture Notes in Computer Science*, Springer, 2016,  
432 pp. 49–65. doi:10.1007/978-3-662-49665-7\_4.
- 433 [8] J. Lin, M. Lee, I. C. Yu, E. B. Johnsen, *Modeling and simulation of*  
434 *spark streaming*, in: L. Barolli, M. Takizawa, T. Enokido, M. R. Ogiela,  
435 L. Ogiela, N. Javaid (Eds.), *AINA*, IEEE Computer Society, 2018, pp.  
436 407–413. doi:10.1109/AINA.2018.00068.



- 437 [9] E. Albert, F. S. de Boer, R. Hähnle, E. B. Johnsen, R. Schlatte,  
438 S. L. T. Tarifa, P. Y. H. Wong, Formal modeling and analysis of re-  
439 source management for cloud architectures: an industrial case study  
440 using real-time ABS, *Serv. Oriented Comput. Appl.* 8 (4) (2014) 323–  
441 339. doi:10.1007/s11761-013-0148-0.
- 442 [10] E. Kamburjan, R. Hähnle, S. Schön, Formal modeling and analysis  
443 of railway operations with active objects, *Sci. Comput. Program.* 166  
444 (2018) 167–193. doi:10.1016/j.scico.2018.07.001.
- 445 [11] N. Bezirgiannis, F. S. de Boer, E. B. Johnsen, K. I. Pun, S. L. T. Tarifa,  
446 Implementing SOS with active objects: A case study of a multicore  
447 memory system, in: R. Hähnle, W. M. P. van der Aalst (Eds.), *FASE*,  
448 Vol. 11424 of *Lecture Notes in Computer Science*, Springer, 2019, pp.  
449 332–350. doi:10.1007/978-3-030-16722-6\\_20.
- 450 [12] C. Consortium, *Compugene*, www.compugene.tu-darmstadt.de (2019).
- 451 [13] C. C. Din, R. Bubel, R. Hähnle, KeY-ABS: A deductive verification tool  
452 for the concurrent modelling language ABS, in: *CADE’25*, Vol. 9195 of  
453 *LNCS*, 2015, pp. 517–526. doi:10.1007/978-3-319-21401-6\\_35.
- 454 [14] C. C. Din, S. L. T. Tarifa, R. Hähnle, E. B. Johnsen, History-based spec-  
455 ification and verification of scalable concurrent and distributed systems,  
456 in: *ICFEM*, Vol. 9407 of *LNCS*, Springer, 2015, pp. 217–233.
- 457 [15] E. Kamburjan, T. Chen, Stateful behavioral types for active objects, in:  
458 *IFM*, Vol. 11023 of *LNCS*, Springer, 2018, pp. 214–235.
- 459 [16] E. Kamburjan, C. C. Din, T. Chen, Session-based compositional anal-  
460 ysis for actor-based languages using futures, in: *ICFEM*, Vol. 10009 of  
461 *Lecture Notes in Computer Science*, 2016, pp. 296–312.
- 462 [17] E. Kamburjan, C. C. Din, R. Hähnle, E. B. Johnsen, Behavioral con-  
463 tracts for cooperative scheduling (2020).
- 464 [18] E. Kamburjan, Behavioral program logic, in: *TABLEAUX*, Vol. 11714  
465 of *LNCS*, Springer, 2019, pp. 391–408.
- 466 [19] R. Bubel, C. C. Din, R. Hähnle, K. Nakata, A dynamic logic with traces  
467 and coinduction, in: *TABLEAUX*, Vol. 9323 of *LNCS*, Springer, 2015,  
468 pp. 307–322.

- 469 [20] E. Kamburjan, Detecting deadlocks in formal system models with con-  
470 dition synchronization, in: AVOCS, Vol. 76 of ECEASST, 2018.
- 471 [21] E. Kamburjan, R. Hähnle, Deductive verification of railway operations,  
472 in: RSSRail, Vol. 10598 of LNCS, Springer, 2017, pp. 131–147.
- 473 [22] E. Kamburjan, N. Wasser, The right kind of non-determinism: Using  
474 concurrency to verify C programs with underspecified semantics 365  
475 (2022) 1–16.
- 476 [23] C. A. R. Hoare, An axiomatic basis for computer programming, *Com-  
477 mun. ACM* 12 (10) (1969) 576–580.
- 478 [24] D. Clarke, R. Muscheci, J. Proença, I. Schaefer, R. Schlatte, Vari-  
479 ability modelling in the ABS language, in: B. K. Aichernig, F. S.  
480 de Boer, M. M. Bonsangue (Eds.), FMCO, Vol. 6957 of Lecture Notes  
481 in Computer Science, Springer, 2010, pp. 204–224. doi:10.1007/  
482 978-3-642-25271-6\_11.
- 483 [25] J. Bjørk, F. S. de Boer, E. B. Johnsen, R. Schlatte, S. L. T. Tarifa, User-  
484 defined schedulers for real-time concurrent objects, *Innov. Syst. Softw.  
485 Eng.* 9 (1) (2013) 29–43. doi:10.1007/s11334-012-0184-5.
- 486 [26] E. Kamburjan, S. Mitsch, M. Kettenbach, R. Hähnle, Modeling and  
487 verifying cyber-physical systems with hybrid active objects, *CoRR*  
488 abs/1906.05704. arXiv:1906.05704.
- 489 [27] R. Schlatte, E. B. Johnsen, E. Kamburjan, S. L. T. Tarifa, Model-  
490 ing and analyzing resource-sensitive actors: A tutorial introduction,  
491 in: F. Damiani, O. Dardha (Eds.), COORDINATION, Vol. 12717 of  
492 Lecture Notes in Computer Science, Springer, 2021, pp. 3–19. doi:  
493 10.1007/978-3-030-78142-2\_1.
- 494 [28] R. Schlatte, E. B. Johnsen, J. Mauro, S. L. T. Tarifa, I. C. Yu, Re-  
495 lease the beasts: When formal methods meet real world data, in: F. S.  
496 de Boer, M. M. Bonsangue, J. Rutten (Eds.), It’s All About Coordina-  
497 tion - Essays to Celebrate the Lifelong Scientific Achievements of Farhad  
498 Arbab, Vol. 10865 of Lecture Notes in Computer Science, Springer, 2018,  
499 pp. 107–121. doi:10.1007/978-3-319-90089-6\_8.

- 500 [29] E. Kuitert, A. Knüppel, T. Bordis, T. Runge, I. Schaefer, Verification  
501 strategies for feature-oriented software product lines, in: P. Arcaini,  
502 X. Devroey, A. Fantechi (Eds.), VaMoS, ACM, 2022, pp. 12:1–12:9.  
503 doi:10.1145/3510466.3511272.
- 504 [30] G. T. Leavens, E. Poll, C. Clifton, Y. Cheon, C. Ruby, D. Cok, P. Müller,  
505 J. Kiniiry, P. Chalin, D. M. Zimmerman, W. Dietl, JML Reference Man-  
506 ual, draft revision 2344 (May 2013).
- 507 [31] E. Kamburjan, Modular verification of a modular specification: Be-  
508 havioral types as program logics, Ph.D. thesis, Technische Universität  
509 Darmstadt (2020).
- 510 [32] C. C. Din, O. Owe, Compositional reasoning about active objects with  
511 shared futures, *Formal Aspects Comput.* 27 (3) (2015) 551–572.
- 512 [33] B. Beckert, D. Bruns, Dynamic logic with trace semantics, in: CADE,  
513 Vol. 7898 of LNCS, Springer, 2013, pp. 315–329.
- 514 [34] B. Beckert, A dynamic logic for the formal verification of java card  
515 programs, in: Java Card Workshop, Vol. 2041 of LNCS, Springer, 2000,  
516 pp. 6–24.
- 517 [35] D. Ancona, V. Bono, M. Bravetti, J. Campos, G. Castagna, P. Deniélou,  
518 S. J. Gay, N. Gesbert, E. Giachino, R. Hu, E. B. Johnsen, F. Martins,  
519 V. Mascardi, F. Montesi, R. Neykova, N. Ng, L. Padovani, V. T. Vascon-  
520 celos, N. Yoshida, Behavioral types in programming languages, *Found.*  
521 *Trends Program. Lang.* 3 (2-3) (2016) 95–230.
- 522 [36] C. C. Din, O. Owe, A sound and complete reasoning system for asyn-  
523 chronous communication with shared futures, *J. Log. Algebraic Methods*  
524 *Program.* 83 (5-6) (2014) 360–383.
- 525 [37] K. Honda, N. Yoshida, M. Carbone, Multiparty asynchronous session  
526 types, in: POPL, ACM, 2008, pp. 273–284.
- 527 [38] N. Rollshausen, Counterexample generation for formal verification of  
528 abs, Bachelor Thesis, TU Darmstadt, available at [https://tuprints.  
529 ulb.tu-darmstadt.de/17856/](https://tuprints.ulb.tu-darmstadt.de/17856/) (2021).
- 530 [39] CLOC tool, version 1.82.

- 531 [40] N. Wasser, A. H. Tabar, R. Hähnle, Automated model extraction: From  
532 non-deterministic C code to active objects, *Sci. Comput. Program.* 204  
533 (2021) 102597.
- 534 [41] E. Kamburjan, From post-conditions to post-region invariants: Deduc-  
535 tive verification of hybrid objects, in: *HSCC'21*, ACM, 2021.
- 536 [42] A. Platzer, The complete proof theory of hybrid systems, in: *LICS*,  
537 IEEE Computer Society, 2012, pp. 541–550.
- 538 [43] S. Mitsch, A. Platzer, A retrospective on developing hybrid system  
539 provers in the KeYmaera family - A tale of three provers, in: *20 Years of*  
540 *KeY*, Vol. 12345 of *Lecture Notes in Computer Science*, Springer, 2020,  
541 pp. 21–64.
- 542 [44] J.-C. Filliâtre, A. Paskevich, Why3 — where programs meet provers, in:  
543 M. Felleisen, P. Gardner (Eds.), *ESOP*, Vol. 7792 of *Lecture Notes in*  
544 *Computer Science*, Springer, 2013, pp. 125–128.
- 545 [45] M. Barnett, B. E. Chang, R. DeLine, B. Jacobs, K. R. M. Leino, Boogie:  
546 A modular reusable verifier for object-oriented programs, in: *FMCO*,  
547 Vol. 4111 of *LNCS*, 2005.
- 548 [46] M. Sirjani, A. Movaghar, A. Shali, F. S. de Boer, Modeling and veri-  
549 fication of reactive systems using rebecca, *Fundam. Informaticae* 63 (4)  
550 (2004) 385–410.
- 551 [47] M. Sirjani, A. Movaghar, A. Shali, F. S. de Boer, Model checking,  
552 automated abstraction, and compositional verification of rebecca mod-  
553 els, *J. Univers. Comput. Sci.* 11 (6) (2005) 1054–1082. doi:10.3217/  
554 jucs-011-06-1054.
- 555 [48] C. S. Gordon, Modal assertions for actor correctness, in: F. Bergenti,  
556 E. Castegren, J. D. Koster, J. Franco (Eds.), *AGERE!@SPLASH*, ACM,  
557 2019, pp. 11–20. doi:10.1145/3358499.3361221.
- 558 [49] M. Bagherzadeh, N. Fireman, A. Shawesh, R. Khatchadourian, Actor  
559 concurrency bugs: a comprehensive study on symptoms, root causes,  
560 API usages, and differences, *Proc. ACM Program. Lang.* 4 (OOPSLA)  
561 (2020) 214:1–214:32. doi:10.1145/3428282.

- 562 [50] R. Pardo, E. B. Johnsen, I. Schaefer, A. Wasowski, A specification logic  
563 for programs in the probabilistic guarded command language, in: IC-  
564 TAC, Vol. 13572 of Lecture Notes in Computer Science, Springer, 2022,  
565 pp. 369–387.
- 566 [51] R. Bubel, R. Hähnle, A. H. Tabar, A program logic for dependence  
567 analysis, in: IFM, Vol. 11918 of Lecture Notes in Computer Science,  
568 Springer, 2019, pp. 83–100.
- 569 [52] M. Scaletta, R. Hähnle, D. Steinhöfel, R. Bubel, Delta-based verifica-  
570 tion of software product families, in: E. Tilevich, C. D. Roover (Eds.),  
571 GPCE, ACM, 2021, pp. 69–82.

572 **Required Metadata**

573 **Current code version**

574 Ancillary data table required for subversion of the codebase. Kindly re-  
 575 place examples in right column with the correct information about your cur-  
 576 rent code, and leave the left column as it is.

<b>Nr.</b>	<b>Code metadata description</b>	<b>Please fill in this column</b>
C1	Current code version	v1.1.2
C2	Permanent link to code/repository used for this code version	<a href="https://github.com/Edkamb/crowbar-tool/releases/tag/v1.1.2">https://github.com/Edkamb/crowbar-tool/releases/tag/v1.1.2</a>
C3	Permanent link to Reproducible Capsule	<a href="https://doi.org/10.24433/CO.6726262.v1">doi.org/10.24433/CO.6726262.v1</a>
C4	Legal Code License	BSD-3-Clause
C5	Code versioning system used	git
C6	Software code languages, tools, and services used	Kotlin, gradle, antlr, github
C7	Compilation requirements, operating environments & dependencies	Z3, Java $\geq$ 1.11
C8	If available Link to developer documentation/manual	<a href="https://github.com/Edkamb/crowbar-tool/wiki">https://github.com/Edkamb/crowbar-tool/wiki</a>
C9	Support email for questions	eduard@ifi.uio.no

Table 2: Code metadata (mandatory)