

Feature-Family-Based Reliability Analysis of Software Product Lines

André Lanna^a, Thiago Castro^{a,b}, Vander Alves^{a,d}, Genaina Rodrigues^a,
Pierre-Yves Schobbens^c, Sven Apel^d

^a*Computer Science Department, University of Brasília. Campus Universitário Darcy
Ribeiro - Edifício CIC/EST, 70910-900, Asa Norte, Brasília -DF, Brazil*

^b*Systems Development Center, Brazilian Army. QG do Exército - Bloco G - 2o. andar,
Setor Militar Urbano, Brasília - DF, Brazil*

^c*PReCISE, NaDI, Faculty of Computer Science, University of Namur. Rue
Grandgagnage 21, 5000 Namur, Belgium*

^d*Department of Informatics and Mathematics, University of Passau Innstr. 33, 94032,
Passau, Germany*

Abstract

Context: Verification techniques are being applied to ensure that software systems achieve desired quality levels and fulfill functional and non-functional requirements. However, applying these techniques to software product lines is challenging, given the exponential blowup of the number of products. Current product-line verification techniques leverage symbolic model checking and variability information to optimize the analysis, but still face limitations that make them costly or infeasible. In particular, state-of-the-art verification techniques for product-line reliability analysis are enumerative which hinders their applicability, given the latent exponential blowup of the configuration space.

Objective: The objectives of this paper are the following: (a) we present a method to efficiently compute the reliability of all configurations of a compositional or annotation-based software product line from its UML behavioral models, (b) we provide a tool that implements the proposed method, and (c) we report on an empirical study comparing the performance of different reliability analysis strategies for software product lines.

Method: We present a novel *feature-family-based* analysis strategy to compute the reliability of all products of a (compositional or annotation-based) software product line. The *feature-based* step of our strategy divides the behavioral models into smaller units that can be analyzed more efficiently.

The *family-based* step performs the reliability computation for all configurations at once by evaluating reliability expressions in terms of a suitable variational data structure.

Results: Our empirical results show that our feature-family-based strategy for *reliability* analysis outperforms, in terms of time and space, four state-of-the-art strategies (product-based, family-based, feature-product-based, and family-product-based) for the same property. It is the only one that could be scaled to a 2^{20} -fold increase in the size of the configuration space.

Conclusion: Our feature-family-based strategy leverages both feature- and family-based strategies by taming the size of the models to be analyzed and by avoiding the products enumeration inherent to some state-of-the-art analysis methods.

Keywords: Software Product Lines, Software Reliability Analysis, Parametric Verification

1 Introduction

Achieving a high quality, low costs, and a short time to market are the driving goals of software product line engineering. A software product line [11] is created to take advantage of the commonalities and variabilities of a specific application domain, by reusing artifacts when instantiating individual software products (a.k.a. *variants* or simply *products*). A domain variability is expressed in terms of features, which are distinguishable characteristics relevant to some stakeholder of the domain [13]. Nowadays software product line engineering is widely accepted in both industry [46, 31] and academia [1, 11, 26, 38].

Quality assurance of product lines has drawn growing attention [32, 42]. Particularly, model checking techniques for product lines explore the space of all products in a product line by searching for execution states where functional [7, 8, 9] or non-functional [17, 21, 28, 34, 39] properties are violated [5]. Nevertheless, employing model checking techniques to verify product lines is a complex task, posing a twofold challenge [8]: (1) the number of variants may grow exponentially with the number of features, which gives rise to an *exponential blowup* [10, 9, 4, 1]; and (2) model checking is inherently prone to the *state-explosion problem* [3, 5]. Therefore, model checking all products of a product line is often not feasible in practice [42].

In previous work, model checking techniques have been applied to analyze

22 probabilistic properties of product lines, in particular, reliability [21, 39, 34].
23 These approaches attenuate the complexity of analyzing probabilistic prop-
24 erties by exploiting, to some extent, reuse in modeling and analysis. On the
25 one hand, non-compositional techniques exploit commonalities across prod-
26 ucts resulting into a single model representing the variability and the behavior
27 of the product line as a whole (covering the behaviors of all products), but
28 it may not scale due to the large state space of models generated by this
29 modeling approach [39, 34]. On the other hand, a compositional alternative
30 is to create and analyze isolated models for each feature and then evaluate
31 them jointly for each configuration [21]. This approach is space-efficient, but
32 faces an exponential blowup by enumerating all valid configurations, which
33 leads to time scalability issues. In essence, both approaches have limita-
34 tions in reusing analysis effort in product lines. As a result, state-of-the-art
35 verification techniques for product-line reliability analysis are enumerative
36 (product-based), which hinders their applicability, given the latent exponen-
37 tial blowup of the configuration space. Consequently, unwanted redundant
38 computational effort is wasted on modeling and analyzing product line’s mod-
39 els [21].

40 As our key contribution, we present a strategy to efficiently compute
41 the reliability of all products of both compositional and annotation-based
42 product lines, without enumerating and analyzing each of these products.
43 Our strategy employs a divide-and-conquer approach in which pre-computed
44 reliabilities of individual features are combined to compute the reliability of
45 the whole product line in a single pass. In a nutshell, in the first step, a
46 feature-based analysis is applied to build a probabilistic model per feature
47 and to analyze each such model using a parametric model checker, returning
48 expressions that describe the reliability of features. Parameters in a feature’s
49 reliability expression represent the reliabilities of other features on which it
50 depends at runtime. In the second step, our strategy performs a family-based
51 step to evaluate each expression in terms of Algebraic Decision Diagrams [2]
52 that are used to encode the knowledge about valid feature combinations
53 and the mapping to their corresponding reliabilities. Since our strategy is
54 a combination of feature-based and family-based analyses, it is effectively a
55 feature-family analysis strategy [42], being the first of its kind for reliability
56 analysis.

57 We implemented our approach in the tool REANA (which stands for **Re-**
58 **liability Analysis**), whose source code is publicly available as free and open-

59 source software¹. The tool takes as input a set of UML behavioral mod-
60 els annotated with reliability information and a feature model of a product
61 line, and it outputs the reliability values for the valid configurations (i.e.,
62 products) of this product line. To evaluate the time-space complexity, we
63 performed 120 experiments to empirically compare our feature-family-based
64 analysis strategy with the following state-of-the-art strategies [42]: product-
65 based, family-based, feature-product-based, and family-product-based. We
66 implemented these alternative strategies as variations of REANA and used
67 them to analyze twenty variants of each of six publicly available product-line
68 models: a system for monitoring an individual’s health [39], control systems
69 for mine pumps [29] and lifts [37], an email system [43], inter-cloud configu-
70 ration [19], and a game [43]. These product lines have been used widely as
71 benchmarks; they have configuration spaces of different sizes, ranging from
72 dozens to billions of billions of products.

73 Our experiment consisted of progressively increasing the number of fea-
74 tures and the size of the behavioral models for each of the product lines,
75 analyzing each of the evolved product lines with all analysis strategies. Our
76 results indicate that the feature-family-based strategy has the best perfor-
77 mance in terms of time and space, being the only one that could be scaled
78 to a 2²⁰-fold increase in the size of the configuration space for *reliability*
79 analysis when compared to four state-of-the-art strategies for the same prop-
80 erty: product-based, family-based, feature-product-based, and family-prod-
81 uct-based.

82 In summary, the contributions of this paper are the following:

- 83 • We introduce a novel feature-family-based strategy for reliability anal-
84 ysis that analyzes each feature in isolation and combines the resulting
85 pieces of information to compute the reliability of a given product line
86 (Section 3);
- 87 • We provide a novel tool, called REANA, implementing such feature-
88 family-based strategy, to carry out the analysis of reliability of a prod-
89 uct line from its UML behavioral diagrams and its feature model (Sec-
90 tion 4.1);
- 91 • We report on an empirical study comparing the performance of our

¹<https://github.com/SPLMC/reana-spl/>

92 feature-family-based strategy to other state-of-the-art analysis strate-
93 gies, implemented as an extension of our REANA tool (Section 4.3).

94 Supplementary material, including the REANA tool and its extensions
95 (which include all evaluation strategies considered in this work), as well as
96 models used in our empirical evaluation and respective experimental results
97 are publicly available for replication purposes at [http://splmc.github.io/
98 scalabilityAnalysis/](http://splmc.github.io/scalabilityAnalysis/).

99 2. Background

100 In this section, we provide an overview of fundamental concepts related to
101 our work and a running example to guide the presentation of our approach
102 in later sections. We assume the reader is familiar with software product
103 lines [11, 38] and discrete-time Markov chains (DTMC) [3].

104 2.1. Reliability Analysis and FDTMC

105 Probabilistic verification techniques have been used in the past to substi-
106 tute the concept of absolute correctness by bounds on the probability that
107 certain behavior may occur. Based on probabilistic models, it is possible
108 to specify probabilistic system behavior due to, e.g., intrinsically unreliable
109 hardware components and environmental characteristics. Reliability can be
110 defined as a probabilistic existence property [22], in the sense that it is given
111 by the probability of eventually reaching some set of *success* states in a prob-
112 abilistic behavioral model of a system. (In our setting, we define success to
113 mean that all tasks of interest have been accomplished as intended.)

114 Discrete-time Markov Chain (DTMC) is a well-known formalism to model
115 such probabilistic behavior. In a DTMC, the reachability probability is de-
116 fined as the sum of probabilities for each possible path that starts in an
117 initial state and ends in a state belonging to the set of target states [3].
118 Thus, to compute reliability, we label success states with the atom “*suc-*
119 *cess*” and compute the reachability probability of success states, expressed
120 as $P_{=?}[\diamond \text{“success”}]$ in the query language of the PARAM model checker [24].

121 To analyze the behavior of a product line, it is useful to embed its in-
122 herent variability in such a probabilistic model. A possible approach is to
123 use *parametric DTMCs* (PDTMC) [15], which augment DTMCs with transi-
124 tion probabilities that can be expressed as variables. A PDTMC is a DTMC
125 whose probability matrix takes values from a set X of strictly positive param-
126 eters. A PDTMC gives rise to a family of DTMCs by instantiating the formal

127 parameters to values with an instantiation function $\kappa : \mathbb{Q}_+ \cup X \mapsto [0, 1]$. For
 128 a parametric DTMC D_X and an instantiation function κ , $\kappa(D_X)$ denotes the
 129 DTMC whose probability matrix is given by instantiating D_X 's formal pa-
 130 rameters. For PDTMCs, the reliability analysis problem can be solved by
 131 a *parametric probabilistic reachability* algorithm [23], which outputs a ratio-
 132 nal expression (a fraction of two polynomials) on the same variables as the
 133 ones in the input parametric model. The idea behind this technique is that
 134 evaluating the variables in the rational expression yields the reliability value
 135 of the DTMC that would be obtained by an equivalent evaluation of the
 136 variables in the PDTMC. However, this behavioral representation does not
 137 take a variability model (e.g., a feature model) into account, and thus is not
 138 sufficient for representing *possible* behavior in a product line (i.e, behavior of
 139 actual products).

140 Featured Discrete-time Markov Chains (FDTMC) [39] are probabilistic
 141 models that properly handle product-line variability. They can be thought as
 142 DTMCs that, instead of transition probabilities, have transition *probability*
 143 *profiles*. These profiles are functions $\llbracket FM \rrbracket \rightarrow [0, 1]$ that map a configuration
 144 to a probability value, where $\llbracket FM \rrbracket$ denotes the set of valid configurations of
 145 the feature model FM . Rodrigues et al. [39] proposed a method to encode
 146 an FDTMC as a PDTMC, enabling its analysis by off-the-shelf parametric
 147 model checkers. In the present work, we leverage the view of Rodrigues et al.
 148 [39] of FDTMCs as PDTMCs for the purpose of compositional reliability
 149 analysis.

150 2.2. Software Product Line Analysis

151 Several analysis techniques have been proposed by researchers for soft-
 152 ware product lines, each one taking a particular property into account. To
 153 help researchers and practitioners understand the similarities and differences
 154 among such techniques, Thüm et al. [42] propose a classification of the exist-
 155 ing techniques, which we follow in this work. In our context, a *product-based*
 156 reliability analysis operates only on derived (non-variable) UML behavioral
 157 models, whereas the variability model may be used to generate the models.
 158 As it is a brute-force strategy, it is only feasible for product lines with few
 159 products. In contrast, the *family-based* strategy for reliability analysis oper-
 160 ates over variant-rich UML behavioral models and incorporates the knowl-
 161 edge about valid feature combinations. In a *feature-based* analysis strategy,
 162 the reliability of UML behavioral models related to each individual feature
 163 is analyzed in isolation from the others, i.e., interactions among features and

164 the knowledge about valid feature combinations are not incorporated into
165 the analysis.

166 Other evaluation strategies may be formed by combining two or more
167 strategies aforementioned [42]. For instance, a *feature-product* analysis con-
168 sists of a feature-based analysis step followed by a product-based analysis,
169 such that the result of the feature-based analysis is reused by the product-
170 based analysis. In the context of reliability, the reliability of UML behavioral
171 models related to each feature is first evaluated in isolation and then the anal-
172 ysis result is reused when enumerating and evaluating the reliability of each
173 non-variant UML behavioral model of the product line.

174 Although other combined evaluation strategies are possible, the afore-
175 mentioned strategies suffice as contrast to our proposed strategy. For more
176 information regarding the remaining strategies, please refer to Thüm et al.
177 [42].

178 2.3. Running Example

179 To illustrate the concepts presented throughout this paper, we introduce
180 an example of a simple product line within the medical domain, for which
181 reliability is considered the major requirement [25]: the Body Sensor Net-
182 work (BSN) product line is a network of connected sensors that capture vital
183 signs from an individual and send them to a central system to analyze the
184 collected data and identify critical health situations [39]. This product line
185 has software components that interpret data provided by the sensors and
186 analyze an individual’s health situation, as well as components for data per-
187 sistence in a database or memory. The set of possible configurations for this
188 product line is defined by its feature model (Figure 1), in which wireless sen-
189 sors are grouped by the feature *Sensor*, software components for interpreting
190 health information are grouped by the feature *SensorInformation*, and the
191 alternatives for data persistence are grouped by the feature *Storage*.

192 To continuously monitor an individual’s health situation, the BSN prod-
193 uct line has a control loop comprised of four activities: capture data coming
194 from sensors, process information about the health condition, identify health
195 goal changes, and reconfigure the system if necessary. This control loop rep-
196 represents the coarse-grained behavior of the BSN product line and it is modeled
197 by the activity diagram shown in Figure 2a, with each activity being repre-
198 sented in detail by a sequence diagram involving the components and their
199 behavior. Therefore, every product instantiated from the BSN product line
200 executes this control loop and, whenever the individual’s health condition





Figure 1: BSN-SPL Feature Model


201 changes and this triggers a quality-of-service goal change, another product
 202 is instantiated from this product line with the desired behavior to reach
 203 the desired quality-of-service goal. The sequence diagrams play the role of
 204 representing the behavioral variability where necessary, by means of guard
 205 conditions involving the presence of features (a.k.a *presence conditions* [14]).

206 For instance, Figure 2b presents an excerpt of the sequence diagram asso-
 207 ciated with the activity *system identifies situation* (Figure 2a). This activity
 208 consists of processing and persisting data regarding the individual’s health
 209 condition, in particular sensor information, represented by feature *SensorIn-*
 210 *formation* and its child features in Figure 1. Figure 2b depicts the behavior
 211 associated with the computation and persistence of the individual’s oxygena-
 212 tion. Such behavior is defined by the messages exchanged between five soft-
 213 ware components, whose roles are data processing (*Oxygenation*) and per-
 214 sistence (*Persistence*, *SQLite* and *Memory*—*Persistence* dispatches calls
 215 to the concrete persistence engines), and components for communication and
 216 coordination (*Bus*). Each message is named according to its task and has an
 217 associated probability value **prob** to represent the reliability of the channel
 218 between the components comprising the interaction. The reliability is given
 219 by the product of (a) the probability that the required message arrives at
 220 the receiver component and (b) the receiver component’s reliability (i.e., the
 221 probability that it performs the required task without failure). For the BSN
 222 product line, we assume that all channels have reliability 0.999.

223 The guard condition at the top level of the sequence diagram presented
 224 in Figure 2b is the atomic proposition *Oxygenation*. This means that the



(a) Activity Diagram representing the control loop of BSN-SPL



(b) Sequence diagram (excerpt) associated with the activity *system identifies situation*, for processing and persisting Oxygenation information.

Figure 2: Behavioral diagrams for BSN-SPL

225 enclosed behavior is associated with the presence of the *Oxygenation* feature
 226 in a given configuration. This behavior, in turn, has two variants, accord-
 227 ing to the chosen mechanism for data persistence. The optional fragment
 228 whose guard condition is *SQLite* models the behavior of persisting data in a
 229 database whenever feature *SQLite* is part of a configuration. Likewise, the
 230 optional fragment associated to the presence of the feature *Memory* (i.e., the
 231 fragment with the *Memory* guard) models persistence on secondary memory.

Intuitively, the reliability of the BSN-SPL in terms of the UML behav-

ioral diagrams shown in Figure 2 is defined by the probability of reaching the final elements of both activity (Figure 2a) and sequence (Figure 2b) diagrams without any error occurrence. This probability is given by the serial execution of the behavioral elements along the possible paths from the first until the final element in both diagrams. In Figure 2a, for instance, there are two possible executions leading to the end state: the first one considers that a reconfiguration is necessary to accomplish a new QoS goal, whereas the other bypasses the reconfiguration activity. The reliability for such diagram is the sum of the probabilities of both executions, considering that the reliability of each individual activity is represented by a variable named after its configuration parameter. Thus, assuming that the decision to reconfigure the BSN is taken 50% of the times, the reliability computed for the model represented in Figure 2a is given by

$$R(BSN) = rCapture \cdot rSituation \cdot rQoSGoal \cdot 0.5 \\ + rCapture \cdot rSituation \cdot rQoSGoal \cdot rReconfiguration \cdot 0.5$$

Similarly, the reliability of the sequence diagram in Figure 2b is given by the probability that all messages are transmitted and processed without errors (the probability for any such message is noted in the corresponding arrow). The reliability of the *Oxygenation* fragment is then given by

$$R(Oxygenation) = 0.999 \cdot 0.999 \cdot 0.999 \cdot 0.999 \\ \cdot rSQLite \cdot rMemory \cdot 0.999 \cdot 0.999 \\ = 0.999^6 \cdot rSQLite \cdot rMemory$$

Similar to activities in the computation of $R(BSN)$, the reliability values of the fragments associated to the features *SQLite* and *Memory* are represented by variables. The reliability of each of these inner fragments is computed in the same fashion, leading to

$$R(SQLite) = R(Memory) = 0.999 \cdot 0.999 = 0.999^2$$

232 Although the reliabilities of the inner fragments are constant, we are not
 233 able to inline these values into the expression for $R(Oxygenation)$. Indeed,
 234 according to the feature model in Figure 1, features *SQLite* and *Memory*
 235 are alternative, meaning that *exactly one* of them is ever present in a given
 236 configuration. Thus, we leverage variables in the reliability expression to also

237 encode product-line variability: whenever *SQLite* is present and *Memory* is
238 absent, for instance, we evaluate *rSQLite* as $R(\textit{SQLite})$ and *rMemory* as 1.

239 Note that the dynamic behavior of the BSN does not affect our approach
240 to reliability analysis, since we only consider the execution of tasks up to
241 reconfiguration (Figure 2a). Moreover, our approach is entirely based on
242 design-time artifacts. For a deeper discussion on how the BSN is engineered
243 for reconfiguration and how the reliability computation affects this dynamic
244 behavior, please refer to the work by Pessoa et al. [36]

245 3. Feature-Family-based Reliability Analysis

246 In this section, we present our approach to evaluate the reliability prop-
247 erty of product lines following a feature-family-based strategy [42]. It consists
248 of three key steps, as shown in Figure 3.

249 First, the transformation step maps UML behavioral diagrams with vari-
250 ability into a graph structure called Runtime Dependency Graph (RDG),
251 whose nodes represent the behavioral fragments and store corresponding
252 FDTMCs (i.e. the probabilistic behavioral model), meanwhile the edges
253 represent the runtime dependencies between such models. Next, the feature-
254 based evaluation step analyzes each FDTMC with respect to a reliability
255 property, with the support of a parametric model checker. Each FDTMC
256 is analyzed in isolation, by abstracting the existing runtime dependencies
257 as parameters. This results in rational expressions [23] (hereafter referred
258 to simply as *expressions*), each giving the reliability of an FDTMC as a
259 function of the reliabilities of the FDTMCs on which it depends. Lastly,
260 the family-based evaluation step follows a topological sorting of the runtime
261 dependency graph, computing the reliability value of each configuration by
262 evaluating the expression in each node and reusing the evaluation results pre-
263 viously computed for the nodes on which it depends. This step also considers
264 the variability model of the product line in question to prune invalid config-
265 urations. The following subsections describe these steps in detail, guided by
266 the example of Section 2.3.

267 3.1. Transformation

268 To perform reliability analysis of a given product line, our approach first
269 composes its inherent variability and probabilistic behavior into a Runtime
270 Dependency Graph (RDG), which is then used for analysis in further steps.
271 The probabilistic behavior can be derived from UML behavioral models,




Figure 3: Feature-family-based approach for efficient reliability analysis of product lines.

272 representing the runtime interactions between software components, enriched
 273 with reliability information for such interactions. Next, we provide details
 274 on the behavioral models, the RDG, and the transformation of the former
 275 into the latter.

276 3.1.1. Behavioral Models

277 In our approach, the coarse-grained behavior of a product line is repre-
 278 sented by a UML activity diagram, with each activity being refined into a se-
 279 quence diagram [39]. The activity diagram is useful for representing whether
 280 the activities are performed in a sequential or parallel manner, whereas se-
 281 quence diagrams represent how the probabilistic behavior of the interactions
 282 between software components varies according to the configuration space of
 283 the product line. To represent probabilistic behavior, each message in a se-
 284 quence diagram is annotated with a probability value that represents the reli-
 285 ability of the channel—i.e., the probability that the interaction succeeds—by

286 using the UML MARTE profile [35] (e.g., *prob* tags in Figure 2b).

287 As an example, Figure 2a shows a UML activity diagram describing, at
288 a high level, the behavior of all products of the BSN product line. The
289 behavior corresponding to the activity *system identifies situation* is modeled
290 by an associated sequence diagram, partially depicted in Figure 2b.

291 Without loss of generality, behavior variability is defined by *behavioral*
292 *fragments*, each of which can be an activity diagram (that has an associated
293 sequence diagram), a sequence diagram, or an optional combined fragment
294 within a sequence diagram such that this fragment has a guard condition
295 denoting presence condition [14]. These conditions are propositional logi-
296 cal statements defined over features, that denote the set of configurations
297 for which the guarded behavior is present. Optional combined behavioral
298 fragments can be nested, which allows representing behavioral variability at
299 several levels.

300 Note that the behavioral variability expressed by optional fragments may
301 be implemented in two distinct ways: 1) in case the fragment’s guard condi-
302 tion is expressed by an atomic proposition (i.e., a single feature), the feature
303 may be implemented in its own module, which characterizes a compositional
304 product line; 2) if the guard condition is a propositional formula compris-
305 ing two or more features, such tangled behavior can be implemented in an
306 annotation-based style by using, for example, the `#ifdef` and `#endif` macros
307 of the C preprocessor. Therefore, our approach can be applied to analyze
308 both compositional and annotation-based software product lines.

309 The sequence diagram shown in Figure 2b presents three behavioral frag-
310 ments whose presence conditions are the atoms `Oxygenation`, `Memory`, and
311 `SQLite`. The outermost behavioral fragment represents the optional behav-
312 ior for processing the oxygenation information in the BSN product line, and
313 it varies according to two nested behavioral fragments. These latter are op-
314 tional combined fragments related to the features `SQLite` and `Memory` of
315 the feature model in Figure 1 and, jointly with this model’s constraints,
316 ultimately represent alternative behavior for data persistence.

317 3.1.2. Runtime Dependency Graphs

318 A Runtime Dependency Graph (RDG) is a behavioral representation for
319 variable systems, which combines the configurability view of a product line
320 (expressed by presence conditions) with its probabilistic behavior (expressed
321 by FDTMCs). Formally, it can be defined as follows.

322 **Definition 1** (RDG). *A Runtime Dependency Graph \mathcal{R} is a directed acyclic*
 323 *graph $\mathcal{R} = (\mathcal{N}, \mathcal{E}, x_0)$, where \mathcal{N} is a set of nodes, $\mathcal{E} \subseteq \mathcal{N} \times \mathcal{N}$ is a set of*
 324 *directed edges that denote a dependency relation, and $x_0 \in \mathcal{N}$ is the root*
 325 *node with in-degree 0. An RDG node $x \in \mathcal{N}$ is a pair $x = (m, p)$, where m*
 326 *is an FDTMC representing a probabilistic behavior and p is a propositional*
 327 *logic formula that represents the presence condition associated with m .*

328 To build an RDG for a software product line, we extract the configura-
 329 bility and probabilistic information only from the UML behavioral diagrams,
 330 such that each RDG node is associated with an FDTMC derived from a
 331 behavioral fragment and its presence condition. Since we consider that the
 332 UML activity diagram represents the product line’s *coarse-grained behavior*
 333 executed by all products and, each activity is further refined (detailed) into
 334 its respective sequence diagram. Thus, the behavioral variability is not con-
 335 sidered at the representation at system level, which implies its related RDG
 336 nodes have *true* as presence condition (i.e., it is satisfied for all products).
 337 Edges represent dependencies between nodes, which are due to refinement or
 338 nesting relations between the respective behavioral fragments. RDG nodes
 339 that do not depend on any other node are called *basic*. The ones with depen-
 340 dencies are called *variant* nodes, which are represented with outgoing edges
 341 directed to the RDG nodes on which they depend.

342 The structure of UML sequence diagrams is tree-like, which suggests a
 343 tree could be a better model of their dependencies. Nonetheless, applications
 344 sometimes have behavioral fragments replicated throughout UML models.
 345 For instance, the data persistence behavior in Figure 2b is present in all
 346 fragments that denote sensor information processing. In our approach, re-
 347 dundant fragments are represented by a single RDG node, with as many
 348 incoming edges as its number of replications. When performing this reuse,
 349 the resulting graph will be acyclic, because the original UML model is a finite
 350 hierarchy.

351 Figure 7a illustrates an excerpt of the BSN product line’s RDG that rep-
 352 represents the behavioral fragment of Figure 2b. As the fragments related to
 353 the features *SQLite* and *Memory* are nested inside the fragment related to
 354 feature *Oxygenation*, the RDG for this fragment represents the dependencies
 355 between their respective nodes. The behavioral fragment related to *Oxygena-*
 356 *tion* is part of the sequence diagram representing the behavior of the activity
 357 *system identifies situation*. Therefore, this relation is also represented by the
 358 edge from the node *rSituation* to the node *rOxygenation*. For brevity, we

```

1 RDGNode transformAD(ActivityDiagram ad) {
2     RDGNode root = new RDGNode(ad.id);
3     root.model = adToFDTMC(ad);
4     root.presenceCondition = true;
5     for (Activity act : ad.activities) {
6         root.addDependency(transformSD(act.sequenceDiagram));
7     }
8     return root;
9 }

```

Listing 1: Activity diagram transformation

359 do not represent the internal structure of the nodes and the remaining RDG
360 nodes (indicated by ellipses in Figure 7a).

361 3.1.3. From Behavioral Models to RDG

362 The transformation from behavioral models to an RDG can be described
363 at two abstraction levels: the RDG topology and the generation of proba-
364 bilistic models. Listings 1 and 2 both depict the transformation process from
365 the topological point of view. Note that this step relies on uniquely gener-
366 ated identifiers for the behavioral models, which are then used as identifiers
367 for the respective RDG nodes.

368 The process starts by calling the `transformAD` method (Listing 1), passing
369 as argument the single activity diagram that embodies the coarse-grained
370 behavior of the product line. This method creates the *root* node (Line 2),
371 setting its presence condition to `true` (i.e., the overall behavior must always
372 be present; Line 4). The root’s probabilistic model is then generated by
373 processing the input diagram with the `adToFDTMC` method (Line 3), to which
374 we will come back later. We then create an RDG node for each sequence
375 diagram that refines an activity (denoted by the property `act.sequenceDi-`
376 `agram`), subsequently creating edges that mark them as dependencies of the
377 *root* node (Line 6). Note that the *root* node is the only RDG node created
378 by the `transformAD` method, so the root’s FDTMC models the behavior
379 represented by the activity diagram.

380 The creation of RDG nodes for sequence diagrams is similar: the method
381 `transformSD` (Listing 2) takes a behavioral fragment as input and then cre-
382 ates a new RDG node whose FDTMC is derived by the `sdToFDTMC` method
383 (Line 4). In this case, since behavioral fragments encode variability, their

```

1 RDGNode transformSD(BehavioralFragment sd) {
2     RDGNode thisNode = new RDGNode(sd.id);
3     thisNode.presenceCondition = sd.guard;
4     thisNode.model = sdToFDTMC(sd);
5     for (BehavioralFragment frag : sd.optFragments) {
6         thisNode.addDependency(transformSD(frag));
7     }
8     return RDGNode.reuse(thisNode);
9 }

```

Listing 2: Sequence Diagram transformation

384 guard is assigned as the presence condition of the newly created node (Line 3).
385 As with refined activities, we create RDG nodes for nested behavioral frag-
386 ments and set them as dependencies of the node at hand (Line 6).

387 The reuse of behavior briefly mentioned in Section 3.1.2 is performed by
388 calling the static method `RDGNode.reuse` (Listing 2, Line 8). This function
389 maintains a registry of all RDG nodes created, and then searches among
390 them for one that we consider *equivalent* to the one just created. This notion
391 of equivalence is comprised of three conditions: (a) equality of presence con-
392 ditions; (b) equality of FDTMCs; and (c) recursively computed equivalence
393 of dependencies.

394 At the abstraction level of generating probabilistic models, the transfor-
395 mation of activity and sequence diagram elements into FDTMCs consists
396 of applying transformation templates for each considered behavioral element
397 represented on such diagrams. These templates are depicted by the UML
398 behavioral element being transformed (left-hand side of the dashed line in
399 Figures 4 and 5) and by its resulting probabilistic structure (right-hand side).

400 Figure 4 shows the templates for transforming an activity diagram into
401 an FDTMC. The initial node of the activity diagram becomes the first state
402 in the FDTMC and thus it is labeled as *init* (Figure 4a). Each activity ab-
403 stracts behavior that is modeled with more detail in an associated sequence
404 diagram. Accordingly, we abstract the reliability of an activity as a param-
405 eter that acts as a placeholder for the reliability of the corresponding sequence
406 diagram. Therefore, each activity is represented in an FDTMC by the struc-
407 ture depicted in Figure 4b, where the upper edge denotes the reliability value
408 of the associated sequence diagram (the parameter *rActivity*) and the lower




Figure 4: Templates for transforming activity diagram elements into FDTMCs.

409 edge denotes the probability of failure ($1 - rActivity$, the complement of the
 410 success probability).

411 A decision node in an activity diagram denotes a choice between alter-
 412 native behaviors, each one represented by an outgoing transition directed
 413 to another activity diagram element (Figure 4c). Each transition has an
 414 associated guard condition that must be satisfied to allow the execution of
 415 its subsequent behavior. This decision is taken at runtime, but a domain
 416 expert is able to define the probability for each alternative. Therefore, the
 417 transformation of a decision node results into an FDTMC structure com-
 418 prised of a state with as many outgoing transitions as the number of the
 419 direct subsequent elements of the decision node. Each outgoing transition
 420 has a probability value assigned by the domain expert, and these probabilities
 421 must sum up to 1².

²States without variability are regular DTMC states, so the stochastic property holds:

422 A merge node denotes a place where different branches of an activity di-
 423 agram join just before the execution of the next element proceeds. For each
 424 merging branch, there is an incoming edge directed to the merge node, and
 425 only one outgoing edge indicating the execution may proceed. The transfor-
 426 mation of a merge node results into an FDTMC structure consisting of two
 427 states and one edge, as shown in Figure 4d. The first created state repre-
 428 sents a synchronization point for a number of previous branches, and the edge
 429 to the second state (with probability 1.0) indicates that the execution can
 430 proceed. Lastly, the final node represents the coarse-grained execution have
 431 sucessfully reached its end. Since the reliability is given by the probability
 432 of a behavioral execution without errors occurrences, the transformation of a
 433 final node becomes a single FDTMC state labeled as *success*, with a reflexive
 434 edge whose probability is 1.0 (indicating it is an absorbing state), as shown
 435 in Figure 4e.

436 The sequence diagram elements considered by our approach are messages
 437 (synchronous or asynchronous) and combined fragments for representing the
 438 optional, alternative, and loop fragments. The optional combined fragment
 439 is used uniformly for representing the variation points of a product line, as
 440 its semantics allows representing behavioral fragments that may comprise
 441 a product (or not), according to its guard condition. Hence, whenever an
 442 optional fragment occurs within a behavioral fragment (sequence diagram
 443 or any other combined fragment), it represents a software product line vari-
 444 ability (i.e. its condition denotes a presence condition statement) and it is
 445 transformed into an FDTMC structure comprised of three states and two
 446 edges, as illustrated in Figure 5. Accordingly, we abstract the reliability
 447 of the optional combined fragment’s content by the parameter $rFragment$
 448 which acts as a placeholder for the reliability of the whole combined fragment.
 449 The first edge is annotated with $rFragment$ for representing the reliability
 450 values the fragment may assume, while the second edge is annotated with
 451 $1 - rFragment$ for representing the probability of failure occurrences.

452 Transformations of the remaining sequence diagram elements (synchronous,
 453 asynchronous and reply messages, and alternative and loop combined frag-
 454 ments) are performed according to Ghezzi and Sharifloo [21], except that our
 455 approach does not use alternative fragments to represent variation points re-

the probability of transitioning to a successor state must be 1, meaning that these transi-
 tions are the only possible events [3].

456 lated to alternative features. In our method, the behavioral variability is
457 addressed uniformly by the optional fragment whose guard condition is ex-
458 pressed by a propositional logical formula denoting its presence condition
459 statement. Such formula indeed expresses any kind of features relations, in-
460 cluding OR and alternative features. In Section 3.3, we explain how the evalu-
461 ation of a optional combined fragment with an arbitrarily associated presence
462 condition statement is guided and constrained by the feature model’s rules.

463 When the loop fragment is transformed into an FDTMC, it results into
464 a structure that express the probabilistic conditions of an iteration. Both
465 first and last states have two outgoing edges that denote the probability of
466 executing (by the `loop` variable) and skipping (by the complement `1-loop`)
467 the iteration behavior. The FDTMC representing the iteration behavior is
468 represented between the first and last states.

469 The transformation of synchronous, asynchronous, and reply messages re-
470 sults into a structure comprised of three states and two edges. The first edge
471 denotes the success probability of sending the message, while the complement
472 edge denotes its failure probability [21]. The difference between the message
473 types expresses the operational semantics of each message. The synchronous
474 message denotes that the sender component holds its execution while it waits
475 the call’s answer that comes back by its associated reply message. In another
476 way, in an asynchronous message the sender component continues its execu-
477 tion just after sending the message to the called component and it does not
478 wait for a reply message.

479 Since the UML sequence diagram does not have a final element (as the
480 end node represents in a UML activity diagram), the execution of a sequence
481 diagram or an optional combined fragment is considered successful whenever
482 the last element is reached and executed accordingly. As our approach con-
483 siders that an FDTMC has a single and absorbing `error` state, when the
484 last FDTMC’s state is reached, it is ensured that no errors occurred during
485 the behavioral execution, including the execution of the last sequence dia-
486 gram element. Thus, when our approach transforms an sequence diagram or
487 behavioral fragment and there is no remaining element, the last state in the
488 FDTMC is labeled as *“success”*.

489 As an example, Figure 7a shows an excerpt of the RDG corresponding
490 to the UML activity and sequence diagrams depicted in Figures 2a and 2b
491 such there is an RDG node for each kind of behavioral fragment found on
492 both figures. Note that whenever a behavioral fragment (activity or sequence
493 diagrams and optional combined fragment) has to be transformed, its RDG





Figure 5: Transformation of optional combined fragment into FDTMC.


494 node and an edge are created to accommodate its FDTMC and represent
 495 the behavioral dependency, respectively. The node labeled $rRoot$ is the root
 496 node of this RDG. The FDTMC assigned to this node (Figure 6a) is built
 497 by applying the transformation rules in Figure 4 to the activity diagram in
 498 Figure 2a. The decision node in this activity diagram gives rise to the bold
 499 and dashed transitions in Figure 6a, representing the *yes* and *no* branches.

500 The RDG node $rSituation$ represents the sequence diagram depicted in
 501 Figure 2b, corresponding to the activity *System identifies situation* of BSN’s
 502 control loop (Figure 2a). Since this activity is performed by all products, its
 503 presence condition is **true**. The node’s FDTMC, depicted in Figure 6b, is
 504 obtained from the sequence diagram according to the transformation tem-
 505 plate in Figure 5 and the templates defined by Ghezzi and Sharifloo [21].
 506 The outgoing edges of the node $rSituation$ in Figure 7a correspond to its
 507 dependency on the availability of sensor information—one RDG node per
 508 optional behavioral fragment. (Most of the RDG nodes corresponding to
 509 such behavioral fragments are omitted for brevity).


510 The node labeled $rOxygenation$ in Figure 7a represents the behavior in
 511 the behavioral fragment whose presence condition is **Oxygenation** (Figure
 512 2b). The corresponding FDTMC, presented in Figure 6c, is built by apply-
 513 ing the transformation rules described in Section 3.1 in a stepwise fashion.
 514 Since the behavioral fragment consists of four messages, followed by two op-




(a) FDTMC of the control loop of BSN-SPL.



(b) FDTMC of Situation sequence diagram.



(c) FDTMC of Oxygenation sequence diagram.




(d) FDTMC of SQLite/Memory sequence diagram.

Figure 6: Resulting FDTMCs. Error transitions are omitted for brevity.


515 tional combined fragments (with presence conditions `SQLite` and `Memory`)
 516 and other two messages (all messages having reliability 0.999), its result-
 517 ing FDTMC comprises a sequence of four transitions with probability 0.999,
 518 two transitions with their probabilities represented by parameters ($rSQLite$
 519 and $rMemory$), and other two transitions with probability 0.999. The node
 520 $rOxygenation$ depends on two basic RDG nodes, $rSQLite$ and $rMemory$, cor-
 521 responding to the nested behavioral fragments whose presence conditions are
 522 `SQLite` and `Memory`, respectively. Since both fragments have similar behav-

523 ior (two sequential messages, each with reliability 0.999) their corresponding
 524 FDTMCs are equal (Figure 6d).

525 Finally, the approach relies on the divide-and-conquer strategy to decom-
 526 pose behavioral models. During the transformation of a behavioral fragment
 527 into a FDTMC, whenever another behavioral fragment is found, an RDG
 528 node is created with a parent-child dependency relation with the parent's
 529 RDG node. The way a software product line is decomposed results into a
 530 tree-like RDG if there is no behavioral fragment being reused. Otherwise,
 531 an RDG node representing a reused behavior fragment will have as many
 532 incoming edges as the times the fragment is reused. In this specific case, the
 533 structure of the resulting RDG will not be tree-like (that is why the RDG is
 534 a directed acyclic graph, in general).



(a) RDG nodes.



(b) Dependencies between expressions.

Figure 7: RDG excerpt for the BSN product line.

535 *3.2. Feature-based Analysis*

536 The role of the feature-based analysis step is to analyze the FDTMC
 537 for each RDG node in isolation, abstracting from the dependencies to other
 538 RDG nodes. That is, instead of evaluating a potentially intractable FDTMC
 539 for the product line as a whole, we perform multiple evaluations of smaller
 540 models, one per feature.

541 For each RDG node $x \in \mathcal{N}$, its FDTMC is subject to parametric proba-
 542 bilistic reachability analysis [24, 20]. This feature-based analysis yields x 's re-
 543 liability as an expression over the reliabilities of the n RDG nodes x_1, \dots, x_n ,
 544 on which it depends. This expression is denoted by a function $[0, 1]^n \rightarrow [0, 1]$,
 545 that is, the computation of a reliability value takes n reliability values as in-
 546 put. Therefore, there is a function $\varepsilon : \mathcal{N} \rightarrow ([0, 1]^n \rightarrow [0, 1])$ that yields
 547 the semantics of the reliability expression for a given RDG node. To remove
 548 possible ambiguities, the order of the formal parameters is determined by
 549 a total order relation over the corresponding RDG nodes x_i (e.g., a lexi-
 550 cographic order over node labels). When analyzing RDG nodes, the same
 551 reliability property of eventually reaching the *success* final state (expressed
 552 by the model checker query expression $P_{=?}[\diamond \text{“success”}]$ —see Section 2.1) is
 553 used for all FDTMCs.

554 Performing feature-based analysis over the RDG, as depicted in Figure 7a,
 555 yields the expressions shown in Figure 7b. These expressions illustrate that
 556 basic nodes have their reliabilities defined in terms of constants, whereas the
 557 reliabilities of variant nodes ultimately depend on the ones of basic RDG
 558 nodes. For the sake of simplicity, we overload the names of RDG nodes in
 559 Figure 7a as variables in the expressions in Figure 7b. This way, we map
 560 each variable to the RDG node whose reliability it represents.

For instance, in Figure 7b, the reliability expression of the node labeled *rOxygenation* is $0.999^6 \cdot rSQLite \cdot rMemory$, since the only path that reaches the **success** state in the corresponding FDTMC (Figure 6c) is a succession of four transitions with probability 0.999, two parametric transitions (*rSQLite* and *rMemory*), and two other 0.999-valued transitions. The reliability expressions of the nodes *rSQLite* and *rMemory* are constant, since these nodes are basic and, thus, their FDTMCs (Figure 6d) have only constant transitions. In this case, the single path to the **success** state in both FDTMCs has a reachability probability of 0.999^2 . Hence, the reliability expressions for the feature-based analysis of the BSN product line are given by a function ε

such that

$$\begin{aligned}\varepsilon(rOxygenation) &= 0.999^6 \cdot rSQLite \cdot rMemory \\ \varepsilon(rSQLite) &= 0.999^2 \\ \varepsilon(rMemory) &= 0.999^2\end{aligned}$$

561 *3.3. Family-based Analysis*

562 A possible next step would be to evaluate the obtained expressions once
 563 for each valid configuration, so that the reliability of every product would
 564 be computed. This enumerative approach would be, in fact, a *product-based*
 565 analysis, yielding an overall *feature-product-based* analysis, similar to the one
 566 described by Ghezzi and Sharifloo [21]. However, evaluating all products
 567 using this approach would be still prone to an exponential blowup, which
 568 would harm scalability.

569 To avoid this problem, we leverage a family-based analysis strategy to
 570 *lift* each expression to perform arithmetic operations over variational data,
 571 with the help of an appropriate *variational data structure* [45]. This way,
 572 we are able to represent all possible values under variation and efficiently
 573 evaluate results, sharing computations whenever possible. The data structure
 574 of choice is the Algebraic Decision Diagram (ADD)³ [27], because it efficiently
 575 encodes a Boolean function $\mathbb{B}^n \rightarrow \mathbb{R}$. This is the same type as a mapping
 576 from configurations to reliability values would have, provided the Boolean
 577 values $b_1, \dots, b_n \in \mathbb{B} = \{0, 1\}$ are taken to denote the presence (or absence)
 578 of the corresponding features $f_1, \dots, f_n \in F$ (where F is the set of features
 579 in the feature model).

580 Given an expression $\varepsilon(x)$, obtained for an RDG node x in the feature-
 581 based step of the analysis (Section 3.2), the reliability ADD $\alpha(x)$ is obtained
 582 by first valuating the parameters x_1, \dots, x_k of the lifted expression with the
 583 ADDs for the reliabilities $\alpha(x_1), \dots, \alpha(x_k)$ of the corresponding nodes upon
 584 which x depends. Then, arithmetic operations are performed using ADD
 585 semantics: for ADDs A_1 and A_2 over k Boolean variables and a binary oper-
 586 ation $\odot \in \{+, -, \cdot, \div\}$, $(A_1 \odot A_2)(b_1, \dots, b_k) = A_1(b_1, \dots, b_k) \odot A_2(b_1, \dots, b_k)$.

587 However, the computation of $\alpha(x)$ must take presence conditions into
 588 account. To accomplish this, we constrain the valuation of a variable x_i with

³ADDs, also called Multi-Terminal Binary Decision Diagrams (MTBDD), generalize Binary Decision Diagrams (BDD) to Real-valued Boolean functions.


589 an ADD $p_x : \llbracket FM \rrbracket \rightarrow \mathbb{B}$ encoding its presence condition, with x ranging over
590 x_1 to x_n , such n is the number of features. This ADD has the property that
591 all configurations $c \in \llbracket FM \rrbracket$ that satisfy x_i 's presence condition evaluate to
592 1, while all others evaluate to 0. The resulting constrained decision diagram
593 φ_{x_i} is given by:

$$\varphi_{x_i}(c) = \begin{cases} \alpha(x_i)(c) & \text{if } p_{x_i}(c) = 1 \\ 1 & \text{otherwise} \end{cases}$$


594 Notice the attribution of 1 to the reliability of a behavior that is absent
595 in a given configuration. The intuition is that, for those configurations that
596 do not satisfy the fragment's guard conditions (i.e., $p_{x_i}(c) = 0$), the behav-
597 ior represented by the optional fragment will not be part of the resulting
598 product's behavior. Since an absent behavioral fragment has no influence
599 on the reliability of the overall system, in practice we can assume 1.0 as its
600 reliability value (i.e., it cannot fail). The ADD φ_{x_i} is obtained by means of
601 the *if-then-else* operator for decision diagrams, and the operational details
602 of this construction are presented in Section 4.1.

603 This method of evaluating the expressions is inherently recursive, since
604 the resulting value of computing the expression for a given RDG node de-
605 pends on the results of computing the expressions for the nodes on which it
606 depends. For example, Figure 7b shows that the expression $\varepsilon(rOxygenation)$
607 is defined in terms of the variables $rSQLite$ and $rMemory$. Thus, before com-
608 puting the lifted counterpart of expression $\varepsilon(rOxygenation)$, it is necessary to
609 compute the lifted counterparts of expressions $\varepsilon(rSQLite)$ and $\varepsilon(rMemory)$.
610 In a brief, the family-based step computes the reliabilities values each RDG
611 node may assume by solving its ε expression using reliabilities values encoded
612 by α for the nodes it depends on. Thus, it follows that the reliability of the
613 product line as a whole is given by the ADD resulting from the computation
614 of $\alpha(rRoot)$, where $rRoot$ is the root RDG node.

615 Naturally, basic nodes are the base case of this recursion, since, by def-
616 inition, they depend on no other node. Figure 8a depicts the ADDs repre-
617 senting the reliability encoding of the RDG nodes $rSQLite$ and $rMemory$,
618 respectively. Each ADD node represents a feature whose continuous out-
619 going edge denotes the feature's presence at the configuration, meanwhile
620 the dashed outgoing edge means the feature is absent. Thus, $\alpha(rSQLite)$
621 encodes that the RDG node $rSQLite$ assumes the reliability value of 0.999^2



(a) ADDs for $rSQLite$ and $rMemory$ nodes, respectively.



(b) ADD for $rOxygenation$ node.

Figure 8: ADDs for the running example.

622 when the feature $SQLite$ is part of the configuration, and assumes the value
 623 1.0 otherwise.

624 Figure 8b shows the reliability encoding computed for the $rOxygenation$
 625 RDG node. Since $\varepsilon(rOxygenation)$ is defined in terms of the variables rep-
 626 resenting the reliabilities of the nodes on which it depends, $\alpha(rOxygenation)$
 627 is computed by assigning the ADDs previously computed to $rSQLite$ and
 628 $rMemory$ to the corresponding variables in $\varepsilon(rOxygenation)$, which is solved
 629 by employing ADD arithmetics. The resulting ADD is constrained to rep-
 630 resent only the reliabilities of valid configurations when it is multiplied by
 631 the ADD representing the feature model’s rules. In fact, all paths leading to

632 non-zero terminal represent valid configurations. In the case that the feature
 633 *Oxygenation* is absent, its influence on the configuration’s reliability is none,
 634 thus $\alpha(rOxygenation)$ assumes the value 1.0. Otherwise, for configurations
 635 containing *Oxygenation* and only one persistence feature (*SQLite* or *Mem-*
 636 *ory*), the corresponding path in the ADD leads to the reliability value 0.999⁸.
 637 Finally, the paths leading to the reliability value 0 represent ill-formed con-
 638 figurations. For example, since *SQLite* and *Memory* are alternative features,
 639 the paths representing that both features are present or absent will lead to
 640 0. All these cases are also represented by the Table 1.

Table 1: Reliability of *Oxygenation* feature.

Configuration (c)	$\alpha(rOxygenation)(c)$
{Oxygenation, SQLite, \neg Memory}	$995*(998/1000)*1/1000 = 0,99301$
{Oxygenation, \neg SQLite, Memory}	$995*1*(998/1000)/1000 = 0,99301$
{Oxygenation, SQLite, Memory}	–

641 4. Evaluation

642 To assess the merits of a feature-family-based strategy, we first highlight
 643 key aspects of its implementation (Section 4.1) and analyze its complexity
 644 (Section 4.2). Then we report on an empirical evaluation (Section 4.3).

645 4.1. Implementation

646 We implemented our approach as a new tool named REANA (**R**eliability
 647 **A**nalysis), whose source code is open and publicly available⁴. REANA takes
 648 as input a UML behavioral model, for example, built using the MagicDraw
 649 tool⁵, and a feature model described in conjunctive normal form (CNF),
 650 for example, as exported by FeatureIDE [41]. It then outputs the ADD
 651 representing the reliability of all products of the product line to a file in
 652 DOT format, and it prints a list of configurations and respective reliabilities.
 653 The latter can be suppressed or filtered to a subset of possible configurations
 654 of interest.

⁴<https://github.com/SPLMC/reana-spl>

⁵<http://www.nomagic.com/products/magicdraw.html>

```

1 ADD evalReliability(RDGNode root) {
2     List<RDGNode> deps = root.topoSortTransitiveDeps();
3     LinkedHashMap<RDGNode, String> expressionsByNode =
4         getReliabilityExpressions(deps);
5     Map<RDGNode, ADD> reliabilities =
6         evalReliabilities(expressionsByNode);
7     return reliabilities.get(root);
8 }

```

Listing 3: REANA’s main evaluation routine

655 REANA uses PARAM 2.3 [24] to compute parametric reachability prob-
656 abilities and the CUDD 2.5.1 library⁶ for ADD manipulation. However, any
657 other tool or library providing the same functionality (e.g., the parametric
658 model checker from Filieri and Ghezzi [20]) could be used too.

659 REANA’s main evaluation routine is depicted in Listing 3. After parsing
660 and transforming the input models into an RDG structure (see Section 3.1),
661 the method `evalReliability` is invoked on the RDG’s root node. Its first
662 task is to perform a topological sort of the RDG nodes, so that it obtains a
663 list in which every node comes after all the nodes on which it (transitively)
664 depends (Line 2). This implements the recursion described in Section 3.3 in
665 an iterative fashion.

666 Then, it proceeds to the analysis of the reliability property in the FDTMC
667 corresponding to each of the nodes (Line 3), with the support of a paramet-
668 ric model checker. Although this step does not depend on the ordering of
669 nodes (because it handles dependencies as variables), it is useful that its out-
670 put respects this order. This way, the resulting reliability expressions (ε in
671 Section 3.2) can be evaluated in an order that allows every variable to be
672 immediately resolved to a previously computed value, thus eliminating the
673 need for recursion and null checking.

674 The third step is to evaluate each reliability expression, which yields an
675 ADD representing the reliability function (α in Section 3.3) for each of the
676 nodes. The evaluation of such reliability ADDs (method `evalReliabilities`
677 in Line 4, Listing 3) invokes, for each node, method `evalNodeReliability`,
678 which we present in Listing 4. It computes the φ functions of a node’s depen-

⁶<ftp://vlsi.colorado.edu/pub/cudd-2.5.1.tar.gz>

```

1 ADD evalNodeReliability(RDGNode node,
2     String reliabilityExpression,
3     Map<RDGNode, ADD> relCache) {
4     Map<String, ADD> depsReliabilities = new HashMap();
5     for (RDGNode dep: node.getDependencies()) {
6         ADD depReliability = relCache.get(dep);
7         ADD presCond = dep.getPresenceCondition();
8         ADD phi = presCond.ifThenElse(depReliability,
9             constantAdd(1));
10        depsReliabilities.put(dep.getId(), phi);
11    }
12    ADD reliability = solve(reliabilityExpression,
13        depsReliabilities);
14    return FM.times(reliability);
15 }

```

Listing 4: Evaluation of the reliability function for a single node

679 dependencies (as in Section 3.3), encoding satisfaction of their presence conditions
680 by means of conditionals in ADD ITE (*if-then-else*) operations (Line 8, List-
681 ing 4). The reliability function of each dependency is looked up in a reliability
682 cache (`relCache`, in Line 6, Listing 4) and is then used as the *consequent*
683 argument of the ITE operator, with the *alternative* argument being the con-
684 stant ADD corresponding to 1.

685 After all these functions are computed, they are used to evaluate the lifted
686 reliability expression (Line 12, Listing 4). Whenever a variable appears in
687 this expression, function φ of the corresponding RDG node (on which the
688 current one depends) is looked up in a variable–value mapping, indexed by
689 the node id (`depsReliabilities`).

690 When this evaluation of α is done, it is necessary to consider only the valid
691 configurations for the node at hand by discarding the reliability values of ill-
692 formed products. We represent the feature model’s rules by an ADD where all
693 paths leading to terminal 1 represent a valid configuration, otherwise the path
694 leads to terminal 0. Thus, for the node under evaluation we prune invalid
695 configurations by multiplying its reliability ADD by the one representing the
696 feature-model’s rules (Line 14, Listing 4), so the resulting ADD yields the
697 value 0, for ill-formed products and the actual reliability for the valid ones.

698 All reliabilities computed in this way are progressively added to the reli-

699 ability cache *relCache*. At the end of this loop inside `evalReliabilities`,
700 the cache contains the reliability function for every node and is then returned
701 (Line 4, Listing 3). The reliability of interest is then the one of the root RDG
702 node (the one argument to `evalReliability`, Listing 3), so it is queried in
703 constant time because of the underlying data structure.

704 4.2. Analytical Complexity

705 The overall analysis time is the sum of the time taken by each of the
706 sequential steps in Listing 3. First, the computation of an ordering that
707 respects the transitive closure of the dependency relation in an RDG (Line 2)
708 is an instance of the classical topological sorting problem for directed acyclic
709 graphs, which is linear in the sum of nodes and edges [12].

710 Second, the computation of the reliability expression for an RDG node
711 consists of a call to the PARAM parametric model checker, which requires
712 n calls to cover all nodes (Line 3). The problem of parametric probabilistic
713 reachability in a model of s states consists of $O(s^3)$ operations over polynomi-
714 als, each of which depends on the number of monomials in each operand [23].
715 This number of monomials is, in the worst case, exponential in the number
716 of existing variables. The number of variables for a given node is, in turn,
717 dependent on its number of child nodes and on the modeled behavior (e.g.,
718 if there are loops or alternative paths). Thus, the time complexity of com-
719 puting all the reliability expressions is linear in the number of RDG nodes,
720 but depends on the topologies of the RDG and of the models represented by
721 each of its nodes (we address such dependencies with more details later on).

722 Last, method `evalReliabilities` calls method `evalNodeReliability`,
723 which corresponds to the reliability function α in Section 3.3, once for each
724 node. `evalNodeReliability`'s complexity is dominated by that of ADD
725 operations, which are polynomial in the size of the operands [27]. Indeed,
726 for ADDs f , g , and h , the *if-then-else* operation $\text{ITE}(f, g, h)$ is $O(|f| \cdot |g| \cdot$
727 $|h|)$. Likewise, $\text{APPLY}(f, g, \odot)$, where \odot is a binary ADD operator (e.g.,
728 multiplication), is $O(|f| \cdot |g|)$. Here, $|f|$ denotes the size of the ADD f , that
729 is, its number of nodes. Because of configuration pruning (Section 3.3), all
730 ADD sizes in our approach are bound by $|FM_{ADD}|$ (i.e., the size of the ADD
731 that encodes the rules in the feature model).

732 Since the evaluation of α for a given node comprises a number of oper-
733 ations on the reliability ADDs of the nodes on which it depends (Listing 4,
734 Line 12), we must estimate an upper bound for polynomial arithmetics. If a
735 node identified by x has c children (nodes on which it depends), $f'(x)$ is a

736 polynomial in c variables and it has, at most, e_{max}^c monomials of c variables
737 each, where e_{max} is the maximum exponent for any variable. Each monomial
738 has in turn, at most, $2c$ operations: c exponentiations and c multiplications
739 among variables and the coefficient. Also, no variable can have an exponent
740 greater than the maximum number of transitions between the initial and the
741 success states of the original FDTMC, and this number is itself bound by
742 the number m of messages in the corresponding behavioral model fragment.
743 Thus, the number of ADD operations needed to compute this reliability ADD
744 is $O(c \cdot m^c)$. This leads to an evaluation time of $O(c \cdot m^c \cdot |FM_{ADD}|^2)$.

745 Since the reliability of each RDG node needs to be evaluated exactly
746 once (due to caching), we have n computations of $f(x_i)$, one for each of the
747 n RDG nodes x_i . Hence, the cumulative time spent on reliability functions
748 computation is $O(n \cdot c_{max} \cdot m_{max}^{c_{max}} \cdot |FM_{ADD}|^2)$, where c_{max} is the maximum
749 number of children per node, and m_{max} is the maximum number of messages
750 per model fragment.

751 Although this complexity bound is quadratic in the number of features,
752 the number of nodes in an ADD is, in the worst case, exponential in the num-
753 ber of variables. As the variables in FM_{ADD} represent features, this means
754 $|FM_{ADD}|$ can be exponential in the number F of features. Hence, the worst-
755 case complexity is $O(n \cdot c_{max} \cdot m_{max}^{c_{max}} \cdot 2^{2 \cdot F})$. This worst-case exponential blowup
756 cannot be avoided theoretically, but, in practice, efficient heuristics can be
757 applied for defining an ordering of variables that can cause the ADD's size to
758 grow linearly or polynomially, depending on the functions being represented
759 [3]. Thus, as the growth in the sizes of ADDs varies with the product line
760 being analyzed [30] and is, at least, linear in the number of features, we can
761 also say the best-case time complexity is $O(n \cdot c_{max} \cdot m_{max}^{c_{max}} \cdot F^2)$.

In summary, the time complexity of our feature-family-based analysis strategy lies between $O(n \cdot c_{max} \cdot m_{max}^{c_{max}} \cdot F^2)$ and $O(n \cdot c_{max} \cdot m_{max}^{c_{max}} \cdot 2^{2 \cdot F})$, where n is the number of RDG nodes, c_{max} is the maximum number of child nodes in an RDG node, m_{max} is the maximum number of messages in a behavioral fragment, and F is the number of features of the product line.

762 4.3. Empirical Evaluation

763 Our empirical evaluation aims at comparing our feature-family-based
764 analysis strategy (cf. Section 3) with other state-of-the-art strategies for

765 product-line reliability analysis, as identified by Thüm et al. [42]: product-
766 based, family-based, feature-product-based, and family-product-based. It is
767 expected that our feature-family-based approach performs better than the
768 others, since it (a) decomposes behavioral models into smaller ones and (b)
769 prevents an exponential blowup by computing the reliabilities of all products
770 at once using ADDs. The comparison focuses on the practical complexity of
771 the selected strategies and is guided by the following research question:

- 772 • **RQ1:** How do product-line reliability analysis strategies compare to
773 one another in terms of time and space?

774 To address RQ1, we measured the time and space demanded by each
775 strategy for the analysis of six available software product lines and augmented
776 versions thereof. For the time measure, we considered the wall-clock time
777 spent during analysis after model transformation, including the recording of
778 reliability values for all configurations of a given product line. Transformation
779 time was excluded from this measurement, because all of our implementations
780 of the analysis strategies employ the same transformation routines (using
781 the rules presented in Section 3.1.3). From the transformation step on, the
782 analysis strategies start to differ as each one traverses the resulting FDTMC
783 in its specific fashion. For the space measure, we considered the peak memory
784 usage for each strategy during the evaluation of each product line. This
785 empirical assessment is described in detail in the following subsections.

786 *4.3.1. Subject Systems and Experiment Design*

787 To empirically compare the complexity of the different analysis strategies,
788 we started with the models of six available product lines. Table 2 shows the
789 number of features, the size, and the characteristics of the solution space of
790 each one of these product lines. The solution space is described in terms of
791 the number of activities in the activity diagram and of the total number of
792 behavioral fragments present in the sequence diagrams. The general criterion
793 for choosing these systems was the availability of their variability model.
794 We chose EMail, MinePump, BSN, and Lift due to the fact that they had
795 been commonly used in previous work studying model checking of product
796 lines [7, 8, 39]. We selected InterCloud and TankWar product lines due to
797 the significant size of their configuration spaces.

798 Each of the six original systems was evolved 20 times, with each evolu-
799 tion step adding one optional feature and a corresponding behavioral frag-
800 ment with random messages defining its probabilistic behavior. According

Table 2: Initial version of product lines used for empirical evaluation.

	# Features	# Products	Solution Space’s Characteristics	
			# Activities	# Behavioral fragments
EMail [43]	10	40	4	11
MinePump [29]	11	128	7	23
BSN [39]	16	298	4	15
Lift [37]	10	512	1	10
InterCloud [19]	54	110592	5	51
TankWar [43]	144	4.21×10^{18}	7	81

801 to Section 3.1.1, the name of the newly introduced feature was assigned as
 802 the guard condition of each new behavioral fragment, and each message in a
 803 fragment received a probability value. Thus, each evolution step doubles the
 804 size of the configuration space of the subject product line, with an optional
 805 behavior for the added feature.

806 The independent variable of the experiment is the evaluation strategy
 807 employed to perform the reliability analysis. The dependent variables are the
 808 metrics for time and space complexity. Each subject system was evaluated
 809 by all treatments.

810 We analyzed the outcomes using statistical tests, to properly address out-
 811 lying behavior and spurious results. This way, we are more likely to overrule
 812 factors that affect performance but are difficult to control (e.g., JVM warm-
 813 up time and OS process scheduling). Ideally (i.e., disregarding uncontrollable
 814 factors), we would expect all runs of a given analysis strategy over the same
 815 subject product line to yield the same result. Thus, instead of comparing
 816 isolated runs of different strategies, we compare the inferred distribution of
 817 results of all runs of a strategy to the corresponding distribution for another
 818 strategy. Since there were multiple analysis strategies to compare with, we
 819 did so pairwise with the feature-family strategy, for example, feature-based
 820 with feature-family-based or family-based with feature-family-based.

821 We applied standard statistical tests for equality of the pairs of samples.
 822 The null hypothesis was that both samples come from the same distribution,
 823 while the alternative hypothesis was that one comes from a distribution with
 824 larger mean value than the other. The specific statistical test was the Mann-
 825 Whitney U test whenever one of the samples, at least, was not normally
 826 distributed. Otherwise, we applied the t test for independent samples if the
 827 variances were equal, or Welch’s t test in case of different variances. The

828 significance level for all tests was 0.01.

829 4.3.2. Experiment Setup

830 **Modeling** We implemented each strategy as a variant of REANA, thus
831 relying on the same tools and libraries for parametric reachability checking,
832 ADD manipulation, and expression parsing (see Section 4.1). These REANA
833 extensions are also publicly available at the supplementary Web site⁷. Grad-
834 uate students created the input UML behavioral models using MAGICDRAW
835 18.3 with MARTE UML profile. All models were validated by the authors.

836 **Instrumentation** For this experiment we implemented a tool called
837 SPL-Generator to create valid feature and behavioral models of a product
838 line, according to a set of parameters (more details in Appendix B). This
839 tool was used to create evolution scenarios, in order to assess how each eval-
840 uation strategy behaves with the growth of the configuration space. To ob-
841 tain data regarding analysis time, we used Java’s standard library method
842 `System.nanoTime()` to get the time (with nanoseconds precision) reported
843 by the Java Virtual Machine immediately before and right after REANA’s
844 main analysis routine (Listing 3). The difference between these two time mea-
845 sures is taken to be the elapsed analysis time. Space usage was measured
846 using the maximum resident set size reported by the Linux `/usr/bin/time`
847 tool. This value represents the peak RAM usage throughout REANA’s exe-
848 cution.

849 **Evolution Scenarios** We used our SPL-Generator tool to evolve each
850 software product line we chose as a subject system of our empirical evalu-
851 ation, according to the representation provided in Figure 9. This evolution
852 was accomplished stepwise, and it started with the original feature model
853 (created by FeatureIDE) and behavioral models (created by MagicDraw)—
854 this set of models is hereafter referred to as *original seed* or *seed₀*. At each
855 evolution step *ev_i*, the generator tool doubled the configuration space of the
856 subject system by adding an optional feature in order to generate a new fea-
857 ture model *FM_i* (no cross-tree constraint was added, to avoid constraining
858 configuration space growth). For the newly created feature, the generator
859 tool also creates an optional behavioral fragment comprising 10 messages

⁷<http://splmc.github.io/scalabilityAnalysis/>




Figure 9: Evolution of subject systems accomplished by the SPL-Generator tool

860 randomly generated between 2 lifelines randomly chosen from a set of 10 life-
 861 lines. To establish a relation between the new feature and the corresponding
 862 new behavioral fragment, the fragment’s guard condition is defined as being
 863 the atomic proposition containing the new feature’s name, which character-
 864 izes the evolutions as being compositional. However, it is worth mentioning
 865 that our evaluation method also applies to the analysis of annotation-based
 866 software product lines since it was able to evaluate the original version of
 867 the EMail subject system ($seed_0$ that contains optional fragments expressed
 868 by a conjunction of two features thus, following an annotation-based imple-
 869 mentation) and its evolutions. Each lifeline received a random reliability
 870 value from the range $[0.999, 0.99999]$. The guard condition of the behavioral
 871 fragment received an atomic proposition named after the feature, to relate
 872 the newly created items. The topological allocation method was used by the
 873 generator tool to create the new behavioral model BM_i , so the nesting of se-
 874 quence diagrams follows the feature relations in the feature model. The end
 875 of an evolution step results into a new version of the product line ($seed_i$),
 876 which will be considered as a new $seed$ for the next evolution step. Each
 877 subject system was evolved 20 times, as shown in Figure 9, and all artifacts
 878 are available at the paper’s supplementary site.

879 **Measurement Setup** We executed the experiment using twelve Intel
 880 i5-4570TE, 2.70GHz, 4 hyper-threaded cores, 8 GB RAM and 1 GB swap
 881 space, running 64-bit CentOS Linux 7. The experiment environment (i.e.,
 882 the set of tools, product line models, and automation running scripts) was
 883 defined as a Docker⁸ container running 64-bit Ubuntu Linux 16.10, with
 884 access to 4 cores and 6 GB of main memory of the host machine. Each subject
 885 system was evaluated 8 times by each analysis strategy in each machine,
 886 thus summing up 96 evaluations for each pair of subject system and strategy.
 887 Because of the number of evaluations, we set a limit of 60 minutes for analysis

⁸<https://www.docker.com/>

888 execution time, after which the analysis at hand would be canceled. The
889 results were then grouped to perform the time and memory consumption
890 analysis. The evaluations that exceeded the time limit were discarded from
891 the statistical analysis.


892 4.3.3. Results and Analysis

893 Figures 10, 11, 12, 13, 14 and 15 show plots with the mean time and
894 memory demanded to analyze the Email, MinePump, BSN, Lift, InterCloud,
895 and TankWar product lines (and corresponding evolutions), respectively. The
896 horizontal axes represent the number of added features (with respect to the
897 original product line) in the analyzed models. Thus, they range from 0 (the
898 original model) to 20 (last evolution step). The vertical axes represent either
899 the time in milliseconds (in logarithmic scale) or the space in megabytes.

900 The values of the plots are available in Tables A.4 and A.5 of Appendix
901 A. Statistical tests over both time and space data rejected the null hypothesis
902 for all pairs of strategies. Thus, within a significance level of 0.01, we can
903 assume no two samples come from distributions with equal means.

904 Overall, our experiments show with statistical significance that the feature-
905 family-based strategy is faster than all other analysis strategies (as shown
906 in Figures 10a, 11a, 12a, 13a, 14a, and 15a). Regarding execution time,
907 in the worst case, our feature-family-based strategy performed 60% faster
908 than the family-product-based strategy, when analyzing the original models
909 of the Email product line (Figure 10a); in the best case, it outperformed
910 the family-product-based analysis of the BSN product line with 4 optional
911 features added (i.e., its 5th evolution step—Figure 12a) by 4 orders of magni-
912 tude. Such cases are highlighted in yellow in Table A.4. Regarding memory
913 consumption (Figures 10b, 11b, 12b, 13b, 14b, and 15b), the experiment also
914 shows with statistical significance that, in the worst case, the feature-family-
915 based strategy demanded 2% less memory than the family-based strategy
916 when analyzing the original model of the Lift product line; in the best case,
917 it saved around 4,757 megabytes when analyzing the 3rd evolution step of the
918 InterCloud product line. Such cases are highlighted in yellow in Table A.5.


919 Our feature-family-based strategy also scaled better in response to con-
920 figuration space growth in comparison with other strategies. In the worst
921 case, this strategy scaled up to a configuration space one order of magnitude
922 larger than the limit of the nearest scalable strategy (the feature-product-
923 based analysis of the Email, MinePump, BSN, and Lift systems). In the
924 best case, the feature-family-based strategy supported a configuration space



(a) Analysis time.

(b) Demanded memory.


Figure 10: Time and memory required by different analysis strategies when evaluating evolutions of Email System.



(a) Analysis time.

(b) Demanded memory.


Figure 11: Time and memory required by different analysis strategies when evaluating evolutions of MinePump System.



(a) Analysis time.

(b) Demanded memory.


Figure 12: Time and memory required by different analysis strategies when evaluating evolutions of BSN-SPL.




(a) Analysis time.

(b) Demanded memory.

Figure 13: Time and memory required by different analysis strategies when evaluating evolutions of Lift System.



(a) Analysis time. (b) Demanded memory.
 Figure 14: Time and memory required by different analysis strategies when evaluating evolutions of InterCloud System.



(a) Analysis time. (b) Demanded memory.
 Figure 15: Time and memory required by different analysis strategies when evaluating evolutions of TankWar battle game.

925 5 orders of magnitude larger than supported by the feature-product-based
 926 strategy (when analyzing the InterCloud product line). Finally, we highlight
 927 that only our feature-family-based strategy was able to analyze the TankWar
 928 product line, from its original model up to its 9th evolution step. That is,
 929 the feature-family-based strategy was able to analyze the reliability of up to
 930 10^{21} products within 60 minutes.

Table 3: Probabilistic models statistics.

SPL	Feature-*			Family-*		Product	
	# states	# variables	# models	# states	# variables	# states	# models
EMail	12	0.93	14	182	9	115.8	40
MinePump	7.26	0.95	23	289	10	155.5	128
BSN	11.37	1.44	16	238	12	136.56	298
Lift	12.91	0.91	11	153	10	114	512
InterCloud	7.4	0.98	52	437	47	352.25	110592
TankWar	8.30	0.99	79	735	69	≈ 500	4.21×10^{18}

931 4.3.4. Discussion

932 One reason for the feature-family-based strategy being faster than the
 933 alternatives is that it computes the reliability values of a product line by an-
 934 alyzing a small number of comparatively simple models. In contrast, family-
 935 based and family-product-based strategies yield more complex probabilistic
 936 models than the others, trading space for time. The complementary ex-
 937 planation for the performance boost is that the family-based analysis step
 938 leverages ADDs to compute reliability values, which leads to fewer opera-
 939 tions than necessary if these values were to be calculated by enumeration of
 940 all valid product line configurations (cf. Section 4.2).

941 Table 3 shows the average number of states and variables present in the
 942 models created by each analysis strategy⁹, with feature-family-based and
 943 feature-product-based strategies grouped under Feature-*, and family-based
 944 and family-product-based ones grouped under Family-*. Some values are
 945 omitted, because the number of models is always 1 for family-based ap-
 946 proaches, and the number of variables is always 0 for product-based ones.
 947 In this table, all probabilistic models created by Feature-* analyses have,
 948 indeed, fewer states than the ones generated during Family-* and product-

⁹For TankWar, the average number of states in the product-based case is an estimate, because it is impractical to generate all models.

949 based approaches. Feature-based models also have fewer variables than the
950 corresponding family-based ones.

951 The plots of the experiment results reveal some characteristics that de-
952 part from the expected behavior, which we discuss next. First, there is a
953 single data point for the family-based analysis of the BSN product line (Fig-
954 ure 12), despite its analysis time being in the order of seconds (far from reach-
955 ing the time limit). In fact, the family-based strategy was able to analyze
956 BSN’s models up to the 6th evolution step. However, the resulting expression
957 representing the family’s reliability contained numbers that exceeded Java’s
958 floating-point representation capabilities. Thus, converting these numbers to
959 the `double` data type yielded *not a number* (NaN). To the best of our knowl-
960 edge, the overflow of floating-point representation was not reported yet by
961 previous studies addressing reliability analysis of software product lines.

962 The second remarkable characteristic are the plateaus for feature-product-
963 based analysis at the memory plots in Figures 10b, 11b, 12b, and 13b. Our
964 hypothesis is that this behavior is related to the memory management of the
965 Java Virtual Machine (JVM), but a detailed investigation was out of scope.

966 We also note that the plots for feature-family-based analysis are monoton-
967 ically increasing, with two exceptions: a single decrease at the 14th evolution
968 step of the Intercloud product line (Figure 14) and a “valley” from TankWar’s
969 original model to its 4th evolution step (Figure 15). These outliers result from
970 different ordering of variables in ADDs. The inclusion of new variables for
971 the mentioned cases led to a variable ordering that caused a decrease in the
972 number of internal nodes of the resulting ADDs. Thus, the space needed by
973 such data structures was reduced, and so was the time needed to perform
974 ADD operations (which are linear in the number of internal nodes).

975 Moreover, our approach does not constrain the relation between the fea-
976 ture model’s structure and the UML behavioral models implementing the
977 SPL. For instance, the sequence diagram depicted in Figure 2b represents
978 optional behavioral fragments that do not follow the structure of the feature
979 model presented by Figure 1. The *Oxygenation* feature and the *Persistence*
980 features (*SQLite* and *Memory*) are defined in different branches of the feature
981 model, but the behavioral fragments related to them are nested. In general,
982 the guard condition of an optional behavioral fragment is a propositional
983 formula defined over features and can be defined arbitrarily, with no regard
984 to the structure of the feature model.

985 Finally, the effect of having (many) cross-tree constraints in a feature
986 model may affect our evaluation method in a twofold manner. First, by

987 adding cross-tree constraints, the structure of the ADD representing the fea-
988 ture model's rules and the reliabilities values of each node is changed. How-
989 ever, it is not possible to foresee if the number of internal nodes will increase,
990 decrease or stay the same, since this number also depends on the variable
991 ordering. In our implementation, such ordering is defined by an internal
992 heuristic defined by the CUDD library, on which our tool relies (namely,
993 symmetric sifting). The second effect regards to the growth of the configu-
994 ration space. In our experiments, the growth in the configuration space at
995 each evolution step will be less than it is now, which will probably have a
996 positive effect in the scalability of the strategies relying on a product-based
997 step. However, since cross-tree constraints would have a random effect on
998 the assessment, we decided not to add them, so as to have more control over
999 the dependent variables.

1000 *4.3.5. Threats to Validity*

1001 A threat to internal validity is the creation of UML behavioral models of
1002 the product lines by graduate students. To mitigate this threat, the students
1003 received an initial training on modeling variable behavior of product lines.
1004 To validate the accuracy of the produced models, these were inspected by
1005 the authors.

1006 A possible threat to construct validity would be an inadequate definition
1007 of metrics for the experiment. To address this, we tried to rule out imple-
1008 mentation issues such as the influence of parallelism and reporting of results.
1009 Thus, we measured the total elapsed time between the parsing of behavioral
1010 models and the instant the reliabilities were ready to be reported, with all
1011 analysis steps taking place sequentially. In terms of memory usage, we tried
1012 to reduce the influence of garbage collection by measuring the peak memory
1013 usage during execution.

1014 Finally, a threat to external validity arises from the selection of subject
1015 systems. To mitigate this threat, we selected systems commonly used by the
1016 community as benchmarks to evaluate work on model checking of product
1017 lines. To mitigate the risk of our approach not being generalizable, we applied
1018 it to further product lines (InterCloud and TankWar) whose configuration
1019 spaces resemble ones of real-world applications.

1020 5. Related Work

1021 In this section, we discuss related work to our approach, and we highlight
1022 the significant differences. For this purpose we use the classification of Thüm
1023 et al. [42]. Our approach differs from prior work [8, 21, 39] in that (a) it
1024 captures the runtime feature dependencies from the UML behavioral models,
1025 (b) which are enriched with variability information extracted from the feature
1026 model, and (c) we leverage ADDs to compute the reliability of all products
1027 of a product line with fewer operations than an enumeration would require.

1028 5.1. Comparison to a Feature-Product-based Strategy

1029 The evaluation method proposed by Ghezzi and Sharifloo [21] is the clos-
1030 est to our work and, to the best of our knowledge, it represents the state-of-
1031 the-art for reliability evaluation of software product lines. The whole behav-
1032 ior of a product line is modeled by a set of small sequence diagrams arranged
1033 in a tree, where each node has an associated expression resulting from the
1034 analysis performed by a parametric model checker. To compute the reliability
1035 of a product, the tree is traversed in a bottom-up fashion, when each node’s
1036 expression is solved considering the configuration under analysis. The result-
1037 ing value for the root node denotes the product’s reliability. This method
1038 reduces time and effort required for evaluation by employing parametric in-
1039 stead of non-parametric reachability checking of probabilistic models, but it
1040 faces scalability issues as it is inherently enumerative (i.e., the decomposi-
1041 tion tree is traversed for each product). The analysis strategy followed by the
1042 method is *Feature-Product-based*, as it decomposes the behavioral models into
1043 smaller units (feature-based step) and later composes the evaluation results
1044 of each unit to obtain the reliability of a product (product-based step).

1045 Despite the resemblances with this method, our approach presents some
1046 distinguishing characteristics. While Ghezzi and Sharifloo [21] must explore
1047 their decomposition tree each time a configuration is evaluated (thus employ-
1048 ing a product-based analysis as an evaluation step), our approach employs a
1049 family-based evaluation for each RDG node, such that all reliability values it
1050 may assume are computed in a single step. Another difference refers to the
1051 usage of UML sequence diagram elements for representing behavioral vari-
1052 ability. Ghezzi and Sharifloo [21] establish a direct relation from the feature
1053 model’s semantics of optional and alternative features and the semantics of
1054 optional (OPT) and alternative (ALT) combined fragments, respectively. Al-
1055 though such relation is straightforward, it constrains the approach’s expres-

1056 siveness, as only single features can be associated to a combined fragment
1057 (i.e., the combined fragment’s guard condition assumes only atomic proposi-
1058 tions). In contrast, our approach represents behavioral variability uniformly
1059 by the optional combined fragment, with an arbitrary presence condition
1060 as a guard statement. This construct is simpler, because it does not lever-
1061 age alternative fragments, but more expressive, as guards can be defined by
1062 propositional statements.

1063 Another major difference concerns the underlying data structure for rep-
1064 resenting the dependencies between behavioral fragments. Ghezzi and Sharif-
1065 floo [21] use a decomposition tree while our approach uses a directed acyclic
1066 graph that allows to represent a group of replicated behavioral fragments by
1067 a single node. This avoids the effort of performing redundant modeling and
1068 evaluation of the replicated model, which is not possible to accomplish in a
1069 tree structure.

1070 A precise comparison of the tool implementing the method proposed by
1071 Ghezzi and Sharifloo and REANA was not possible, since the former is not
1072 publicly available. Nonetheless, the feature-product-based variant of RE-
1073 ANA we created for our experiment closely resembles Ghezzi and Shari-
1074 floo’s approach, the only exception being the parametric model checker of
1075 choice. Empirical results (Section 4.3) show with statistical significance that
1076 the feature-family-based approach performs faster and demands less mem-
1077 ory than REANA’s feature-product-based variant. For the evaluation time,
1078 the feature-family strategy outperformed our feature-product-based strategy
1079 from 2 times (for the original seed of EMail system) up to 4 orders of mag-
1080 nitude (for the 3rd evolution of Intercloud product line). Regarding space,
1081 the feature-family-based strategy required from 2.6% (original seed of Email
1082 system) up to 97% (3rd evolution of InterCloud) less memory. Moreover, the
1083 feature-product-based strategy was not able to analyze the subject system
1084 with the largest configuration space (Tankwar), whereas our feature-family-
1085 based strategy succeeded up to Tankwar’s 9th evolution.

1086 Ghezzi and Sharifloo’s work [21] presents a theoretical analysis of time
1087 complexity, in which the authors devise a formula for computing the time
1088 needed to verify a number of properties for a product line with their approach.
1089 Their model transformation time is not comparable to ours, mainly because
1090 Ghezzi and Sharifloo do not handle activity diagrams in their work, and we
1091 do not handle reward models in ours. Also, both approaches use external
1092 tools with similar capabilities to perform parametric reachability analysis.
1093 In fact, Ghezzi and Sharifloo argue their tool [20] is actually faster than

1094 PARAM, which is used by REANA. Nonetheless, both model checkers could
1095 be used interchangeably, so we omit parametric reachability analysis time.

1096 Because of that, we assume the output expressions from the parametric
1097 reachability phase to be correspondingly equal in both approaches. This way,
1098 the difference between the strategies is isolated in the way they solve each
1099 expression. While Ghezzi and Sharifloo perform a number k of floating-point
1100 operations for each configuration, our approach performs the same number
1101 k of ADD operations, but only once. Since the number of configurations
1102 is $O(2^F)$, the feature-product-based approach performs $O(k \cdot 2^F)$ computing
1103 steps. As no lowest number of steps is possible if one is to compute the
1104 reliability of all possible configurations, the number of computations in the
1105 best case is also $O(k \cdot 2^F)$. In contrast, an operation over ADDs in our
1106 approach comprises $O(2^{2 \cdot F})$ steps in the worst case, but is $O(F^2)$ in the best
1107 case (see Section 4.2). Thus, the feature-family-based approach performs
1108 between $O(k \cdot F^2)$ and $O(k \cdot 2^{2 \cdot F})$ computing steps.

1109 Hence, we conclude that, in the worst case, the upper bound for our
1110 method’s asymptotic complexity is worse than that of Ghezzi and Sharifloo’s,
1111 but its best-case complexity is better, which is consistent with the empirical
1112 findings from the previous section.

1113 5.2. Other Related Work

1114 Rodrigues et al. [39] present and compare three family-based strategies
1115 to analyze probabilistic properties of product lines. Two of them leverage
1116 PARAM as model checker; the third one relies on FDTMCs representing
1117 the behavior of a whole product line by encoding its variability, resulting in
1118 an ADD expressing the reliability values of all configurations. Our feature-
1119 family-based strategy benefits even more from further breaking down prob-
1120 abilistic models. Indeed, the methods by Rodrigues et al. show a time-space
1121 tradeoff, but all of them presented scalability issues even for small product
1122 lines (around 12 features), whereas our approach is able to analyze a product
1123 line with 144 features and about 10^{18} products within reasonable time.

1124 Further research has addressed efficient verification of other non-func-
1125 tional properties of product lines by exploiting family-based analysis strate-
1126 gies [40, 28, 17, 18, 44, 6, 8, 16]. Siegmund et al. [40] propose an approach
1127 for performance evaluation by simulating the behavior of all variants at run-
1128 time from the variability encoded in compile-time. Such simulator is created
1129 from the log of method calls traced by features. Kowal et al. [28] create
1130 a model representing the whole performance variability of a product line

1131 from UML activity diagrams annotated with performance-related annota-
1132 tions. Dubslaff et al. [17, 18] present an approach for modeling dynamic
1133 product lines and performing quantitative analysis of systems endowed of
1134 non-deterministic choices. Given the non-deterministic characteristic of the
1135 systems evaluated by this approach, the authors consider Markov Decision
1136 Processes as the suitable model for representing the model behavior. Simi-
1137 larly, Varshosaz and Khosravi [44] introduce a mathematical model named
1138 Markov Decision Process Family for representing the behavior of a product
1139 line as a whole, as well as a model checking algorithm to verify properties
1140 expressed in probabilistic computation tree logic. Classen et al. [8] estab-
1141 lish the foundations of *Featured Transition Systems* (FTS) to create a model
1142 endowed with features expressions to represent the states variation of the
1143 whole software product line. The authors also present a family-based model
1144 checker [6] that is able to analyze *Linear Temporal Logic* (LTL) properties
1145 of the whole software product line by employing semi-symbolic algorithms
1146 to verify FTSs. All these pieces of work exploit symbolic computation on
1147 a model representing the whole variability of a product line as a better al-
1148 ternative to product-based strategies. Our study supports this conclusion,
1149 especially if a suitable variational data structure (e.g., ADD) is used for such
1150 analysis. However, our results indicate that feature-family-based analysis
1151 further improves performance.

1152 Dimovski et al. [16] also present an efficient family-based technique to
1153 verify LTL properties of a software family. The authors leverage abstract
1154 interpretation to reduce the configuration space of an FTS, so that it can
1155 be verified by off-the-shelf model checkers (i.e., aimed and optimized to an-
1156 alyze single systems). Our method employs a divide-and-conquer strategy
1157 to reduce model size, without changing the configuration space. Moreover,
1158 our analysis method also employs off-the-shelf model checkers, but to ana-
1159 lyze probabilistic properties of software product lines. Therefore, it is worth
1160 investigating the extent to which the technique proposed by Dimovski et al.
1161 [16] can be applied to the verification of PCTL properties. If that is the
1162 case, we conjecture that both strategies could be combined to further reduce
1163 verification effort.

1164 6. Conclusion

1165 We presented a feature-family-based strategy and corresponding tool for
1166 efficient reliability analysis of product lines. Our approach limits the effort

1167 needed to compute the reliability of a product line by initially employing
1168 a *feature-based* analysis to divide its behavioral models into smaller units,
1169 which can be verified more efficiently. For this purpose, we arrange proba-
1170 bilistic models in an RDG, which is a directed acyclic graph with variability
1171 information. This strategy facilitates reuse of reliability computations for re-
1172 dundant behaviors. The *family-based* step comes next when we perform the
1173 reliability computation for all configurations at once by evaluating reliabil-
1174 ity expressions in terms of ADDs. These decision diagrams encode presence
1175 conditions and the rules from the feature model, so that computation is in-
1176 herently restricted to valid configurations.

1177 The empirical evaluation was accomplished by conducting an experiment
1178 to compare our feature-family-based approach with the following evalua-
1179 tion strategies: feature-product-based, family-based, family-product-based,
1180 and product-based. Overall, the results show the product-based had the
1181 worst time and space performance among all strategies, as we expected. The
1182 family- and family-product-based strategies yield more complex probabilistic
1183 models than the other strategies, due to variability encoding in their mod-
1184 els. The product, family-product and feature-product-based approaches were
1185 sensitive to the size of the configuration space of the software product line,
1186 given their inherent enumerative characteristic. Overall, our experiments
1187 show that the feature-family-based strategy is faster than all other analysis
1188 strategies and demanded less memory in most cases, being the only one that
1189 could be scaled to a 2^{20} -fold increase in the configuration space. Such results
1190 suggest that our feature-family-based strategy outperformed the alternative
1191 strategies due to the following: (a) the feature-based step explores a lower
1192 number of simpler models having fewer variables in comparison to family-
1193 based models; and (b) as the family-based step leverages ADD to compute
1194 reliability values, fewer operations are necessary to compute reliability values
1195 in comparison to the enumerative strategies.

1196 As future work, we plan to extend the empirical evaluation to a larger
1197 number of subject systems. Furthermore, the present study investigated the
1198 sensitivity of analysis performance with respect to changes in the size of
1199 the configuration space of the subject product lines. Thus, we also plan to
1200 extend the study so as to evaluate the performance impact of changes in
1201 other characteristics, such as the number of decision nodes and the number
1202 of messages per behavioral fragment.

1203 **Acknowledgements**

1204 We would like to thank the following people for fruitful discussions and
1205 suggestions on how to improve this work: Alexandre Mota, Cecília Rubira,
1206 Azzedine Boukerche, Rodrigo Bonifácio, Marcelo Ladeira, Abílio Oliveira,
1207 Paula Gueiros, and Eneias Silva. Vander Alves would like to thank for the
1208 research grant CAPES ref. BEX 0557-16-1 / Alexander von Humboldt ref.
1209 3.2-1190844-BRA-HFSTCAPES-E. Sven Apel’s work has been supported by
1210 the German Research Foundation (AP 206/4 and AP 206/6).

1211 **Appendix A. Experiment Data**

1212 The following tables present the mean values for analysis time and mem-
1213 ory consumption obtained in our experiment. Values typeset in boldface are
1214 the best values (i.e., the lowest) gathered from the experiments. Cells con-
1215 taining dashes represent unavailable data, meaning that the corresponding
1216 analysis violated the time limit of 60 minutes.

Table A.4: Time in milliseconds (fastest strategy in boldface).

		SPL's evolutions steps											
		0	1	2	3	4	5	6	7	8	9	20	
		10	11	12	13	14	15	16	17	18	19		
Email	<i>Configuration space's order</i>	10 ¹	10 ¹	10 ²	10 ²	10 ²	10 ³	10 ³	10 ³	10 ⁴	10 ⁴		
	Feature-family	183.04	223.78	233.69	259.65	267.32	285.79	341.65	348.46	366.73	433.30		
	Feature-product	370.63	517.67	742.91	1108.95	1659.31	2358.51	3829.95	6919.98	12803.15	25110.63		
	Family	319.72	1167.27	13944.18	154067.34	-	-	-	-	-	-		
	Family-product	293.26	558.77	1095.75	2850.86	9451.57	34704.42	137866.42	562117.02	1607837.42	-		
	Product	2424.77	7387.35	14349.32	29137.45	57575.0	114084.61	275598.17	-	-	-		
	<i>Configuration space's order</i>	10 ⁴	10 ⁴	10 ⁵	10 ⁵	10 ⁵	10 ⁶	10 ⁶	10 ⁶	10 ⁶	10 ⁷	10 ⁷	10 ⁷
	Feature-family	970.69	1613.76	2833.40	5425.14	10838.39	21719.17	44171.89	90015.26	187645.77	667138.0	-	
	Feature-product	50748.90	103510.61	215932.90	456329.22	945445.46	1966865.48	-	-	-	-		
	Family	-	-	-	-	-	-	-	-	-	-		
	Family-product	-	-	-	-	-	-	-	-	-	-		
	Product	-	-	-	-	-	-	-	-	-	-		
	Minepump	<i>Configuration space's order</i>	10 ²	10 ²	10 ²	10 ³	10 ³	10 ³	10 ³	10 ⁴	10 ⁴	10 ⁴	
		Feature-family	261.18	287.24	298.10	330.88	358.81	408.45	485.06	621.01	877.52	1375.22	
Feature-product		895.80	1226.97	1844.15	2624.96	4204.27	8952.61	14037.50	25989.10	51495.22	104090.89		
Family		-	-	-	-	-	-	-	-	-	-		
Family-product		-	-	-	-	-	-	-	-	-	-		
Product		-	-	-	-	-	-	-	-	-	-		
<i>Configuration space's order</i>		10 ⁵	10 ⁵	10 ⁵	10 ⁶	10 ⁶	10 ⁶	10 ⁶	10 ⁷	10 ⁷	10 ⁷	10 ⁸	
Feature-family		2390.78	4445.44	8790.54	17995.17	36593.45	76513.51	168694.38	354887.72	-	-	-	
Feature-product		211806.42	439411.29	905878.46	1876640.52	-	-	-	-	-	-		
Family		-	-	-	-	-	-	-	-	-	-		
Family-product		-	-	-	-	-	-	-	-	-	-		
Product		-	-	-	-	-	-	-	-	-	-		
BSN		<i>Configuration space's order</i>	10 ²	10 ²	10 ³	10 ³	10 ³	10 ³	10 ⁴	10 ⁴	10 ⁴	10 ⁵	
		Feature-family	237.14	253.65	273.01	305.48	321.69	377.40	389.41	462.66	651.84	1032.05	
	Feature-product	991.30	1487.19	2404.18	4312.01	7875.91	14788.91	28881.71	57887.92	117630.81	241553.61		
	Family	1604.07	-	-	-	-	-	-	-	-	-		
	Family-product	3288.70	11543.38	46273.48	187134.89	672512.22	2109118.92	-	-	-	-		
	Product	6696.06	20259.05	43489.98	97280.21	241249.72	519495.92	1217404.53	-	-	-		
	<i>Configuration space's order</i>	10 ⁵	10 ⁵	10 ⁶	10 ⁶	10 ⁶	10 ⁶	10 ⁷	10 ⁷	10 ⁷	10 ⁸	10 ⁸	
	Feature-family	1713.19	3443.81	7332.75	15090.83	31208.01	83984.25	197660.21	528948.31	-	-	-	
	Feature-product	495594.45	1022294.56	2145986.30	-	-	-	-	-	-	-		
	Family	-	-	-	-	-	-	-	-	-	-		
	Family-product	-	-	-	-	-	-	-	-	-	-		
	Product	-	-	-	-	-	-	-	-	-	-		

continued in the next page

continued from last page

		SPL's evolutions steps											
		0	1	2	3	4	5	6	7	8	9	20	
		10	11	12	13	14	15	16	17	18	19		
Lift	<i>Configuration space's order</i>	10 ²	10 ³	10 ³	10 ³	10 ³	10 ⁴	10 ⁴	10 ⁴	10 ⁵	10 ⁵		
	Feature-family	140.32	169.51	188.56	199.95	223.20	266.20	339.85	472.01	601.46	1021.76		
	Feature-product	1160.78	1786.75	1289.76	4281.11	7739.10	14769.15	29418.50	60785.39	127344.46	266609.58		
	Family	358.06	625.16	2606.86	18223.06	-	-	-	-	-	-	-	
	Family-product	1413.96	3462.52	10777.60	42156.47	167837.42	453830.76	1870142.24	-	-	-	-	
	Product	-	-	-	-	-	-	-	-	-	-	-	
	<i>Configuration space's order</i>	10 ⁵	10 ⁶	10 ⁶	10 ⁶	10 ⁶	10 ⁷	10 ⁷	10 ⁷	10 ⁸	10 ⁸	10 ⁹	
	Feature-family	3114.06	5728.56	10895.96	21081.40	48194.05	100755.69	215756.37	-	-	-	-	
	Feature-product	555525.51	1136506.33	2457317.34	-	-	-	-	-	-	-	-	
	Family	-	-	-	-	-	-	-	-	-	-	-	
	Family-product	-	-	-	-	-	-	-	-	-	-	-	
	Product	-	-	-	-	-	-	-	-	-	-	-	
	InterCloud	<i>Configuration space's order</i>	10 ⁵	10 ⁵	10 ⁵	10 ⁵	10 ⁶	10 ⁶	10 ⁶	10 ⁷	10 ⁷	10 ⁷	
		Feature-family	671.54	717.7	794.95	880.11	922.98	994.65	1126.91	1315.89	1742	2544.49	
Feature-product		407702.34	861181.31	1752682.98	3394277.78	-	-	-	-	-	-	-	
Family		-	-	-	-	-	-	-	-	-	-	-	
Family-product		-	-	-	-	-	-	-	-	-	-	-	
Product		-	-	-	-	-	-	-	-	-	-	-	
<i>Configuration space's order</i>		10 ⁸	10 ⁸	10 ⁸	10 ⁸	10 ⁹	10 ⁹	10 ⁹	10 ¹⁰	10 ¹⁰	10 ¹⁰	10 ¹¹	
Feature-family		4074.43	4280.4	4568.7	5344.4	3936.76	6719.68	12829.35	25588.69	67156.86	-	-	
Feature-product		-	-	-	-	-	-	-	-	-	-	-	
Family		-	-	-	-	-	-	-	-	-	-	-	
Family-product		-	-	-	-	-	-	-	-	-	-	-	
Product		-	-	-	-	-	-	-	-	-	-	-	
TankWar		<i>Configuration space's order</i>	10 ¹⁸	10 ¹⁸	10 ¹⁹	10 ¹⁹	10 ¹⁹	10 ²⁰	10 ²⁰	10 ²⁰	10 ²¹	10 ²¹	
		Feature-family	6643.88	3588.49	2734.86	2966.2	2902.18	3079.4	4221.14	8012.	17096.88	160259.19	
	Feature-product	-	-	-	-	-	-	-	-	-	-	-	
	Family	-	-	-	-	-	-	-	-	-	-	-	
	Family-product	-	-	-	-	-	-	-	-	-	-	-	
	Product	-	-	-	-	-	-	-	-	-	-	-	
	<i>Configuration space's order</i>	10 ²¹	10 ²¹	10 ²²	10 ²²	10 ²²	10 ²³	10 ²³	10 ²³	10 ²⁴	10 ²⁴	10 ²⁴	
	Feature-family	-	-	-	-	-	-	-	-	-	-	-	
	Feature-product	-	-	-	-	-	-	-	-	-	-	-	
	Family	-	-	-	-	-	-	-	-	-	-	-	
	Family-product	-	-	-	-	-	-	-	-	-	-	-	
	Product	-	-	-	-	-	-	-	-	-	-	-	

Table A.5: Space in megabytes (smallest footprint in boldface).

		SPL's evolutions steps											
		0	1	2	3	4	5	6	7	8	9	20	
		10	11	12	13	14	15	16	17	18	19		
Email	<i>Configuration space's order</i>	10 ¹	10 ¹	10 ²	10 ²	10 ²	10 ³	10 ³	10 ³	10 ⁴	10 ⁴		
	Feature-family	113.70	113.84	113.93	114.30	114.45	114.33	114.52	114.91	115.86	117.64		
	Feature-product	117.22	144.30	186.59	269.67	475.61	738.99	1136.73	2359.24	2839.02	2842.46		
	Family	116.97	125.48	136.57	196.99	-	-	-	-	-	-		
	Family-product	120.25	157.90	235.41	510.41	827.79	722.88	1037.62	1501.80	3231.31	-		
	Product	122.65	231.84	272.04	277.98	310.59	309.06	327.65	-	-	-		
	<i>Configuration space's order</i>	10 ⁴	10 ⁴	10 ⁵	10 ⁵	10 ⁵	10 ⁶	10 ⁶	10 ⁶	10 ⁷	10 ⁷	10 ⁷	
	Feature-family	130.65	146.25	174.93	287.00	489.00	839.80	1523.88	3041.86	5807.80	7223.00	-	
	Feature-product	2849.01	2878.10	2927.46	3158.43	3367.68	4181.64	-	-	-	-	-	
	Family	-	-	-	-	-	-	-	-	-	-	-	
	Family-product	-	-	-	-	-	-	-	-	-	-	-	
	Product	-	-	-	-	-	-	-	-	-	-	-	
	MinePump	<i>Configuration space's order</i>	10 ²	10 ²	10 ²	10 ³	10 ³	10 ³	10 ³	10 ⁴	10 ⁴	10 ⁴	
		Feature-family	113.51	114.05	114.41	114.34	114.8	115.61	116.47	118.48	129.12	133.96	
Feature-product		210.97	333.42	504.98	743.93	1319.2	2400.89	2841.77	2844.10	2851.49	2879.39		
Family		-	-	-	-	-	-	-	-	-	-		
Family-product		-	-	-	-	-	-	-	-	-	-		
Product		-	-	-	-	-	-	-	-	-	-		
<i>Configuration space's order</i>		10 ⁵	10 ⁵	10 ⁵	10 ⁶	10 ⁶	10 ⁶	10 ⁶	10 ⁷	10 ⁷	10 ⁷	10 ⁸	
Feature-family		162.48	265.31	390.03	705.39	1165.72	2224.17	4011.27	6921.67	-	-	-	
Feature-product		2914.44	2971.55	3378.84	3789.31	-	-	-	-	-	-	-	
Family		-	-	-	-	-	-	-	-	-	-	-	
Family-product		-	-	-	-	-	-	-	-	-	-	-	
Product		-	-	-	-	-	-	-	-	-	-	-	
BSN		<i>Configuration space's order</i>	10 ²	10 ²	10 ³	10 ³	10 ³	10 ³	10 ⁴	10 ⁴	10 ⁴	10 ⁵	
		Feature-family	114.05	114.30	114.52	114.56	114.83	115.37	115.32	116.97	120.09	134.23	
	Feature-product	339.91	490.76	737.50	1716.41	2379.06	2837.60	2843.36	2850.78	2874.49	2923.93		
	Family	156.54	-	-	-	-	-	-	-	-	-		
	Family-product	493.99	841.31	1171.71	1153.13	2189.89	3263.80	-	-	-	-		
	Product	320.43	335.18	339.40	352.72	327.95	440.60	446.75	-	-	-		
	<i>Configuration space's order</i>	10 ⁵	10 ⁵	10 ⁶	10 ⁶	10 ⁶	10 ⁶	10 ⁷	10 ⁷	10 ⁷	10 ⁸	10 ⁸	
	Feature-family	148.34	186.03	348.58	588.99	1043.94	2225.13	4640.35	7130.79	-	-	-	
	Feature-product	3005.12	3234.19	3821.39	-	-	-	-	-	-	-	-	
	Family	156.54	-	-	-	-	-	-	-	-	-	-	
	Family-product	-	-	-	-	-	-	-	-	-	-	-	
	Product	-	-	-	-	-	-	-	-	-	-	-	

continued in the next page

continued from last page

		SPL's evolutions steps											
		0	1	2	3	4	5	6	7	8	9	20	
		10	11	12	13	14	15	16	17	18	19		
Lift	<i>Configuration space's order</i>	10 ²	10 ³	10 ³	10 ³	10 ³	10 ⁴	10 ⁴	10 ⁴	10 ⁵	10 ⁵		
	Feature-family	113.77	113.85	114.37	114.27	114.56	115.23	116.87	119.88	120.72	134.41		
	Feature-product	292.23	507.43	757.52	1539.24	2399.04	2838.97	2840.51	2859.23	2907.41	2993.05		
	Family	116.54	122.63	136.83	177.02	-	-	-	-	-	-		
	Family-product	272.85	506.39	1277.44	1296.95	1551.49	2440.83	2669.75	-	-	-		
	Product	-	-	-	-	-	-	-	-	-	-		
	<i>Configuration space's order</i>	10 ⁵	10 ⁶	10 ⁶	10 ⁶	10 ⁶	10 ⁷	10 ⁷	10 ⁷	10 ⁸	10 ⁸	10 ⁷	
	Feature-family	203.42	319.45	539.66	826.57	1791.86	3230.47	6324.48	-	-	-	-	
	Feature-product	3199.10	3489.45	4644.73	-	-	-	-	-	-	-	-	
	Family	-	-	-	-	-	-	-	-	-	-	-	
	Family-product	-	-	-	-	-	-	-	-	-	-	-	
	Product	-	-	-	-	-	-	-	-	-	-	-	
	InterCloud	<i>Configuration space's order</i>	10 ⁵	10 ⁵	10 ⁵	10 ⁵	10 ⁶	10 ⁶	10 ⁶	10 ⁷	10 ⁷	10 ⁷	
		Feature-family	119.44	119.87	127.68	127.79	136.03	132.63	136.3	143.96	152.61	175.78	
Feature-product		3071.58	3158.59	3602.16	4884.81	-	-	-	-	-	-		
Family		-	-	-	-	-	-	-	-	-	-		
Family-product		-	-	-	-	-	-	-	-	-	-		
Product		-	-	-	-	-	-	-	-	-	-		
<i>Configuration space's order</i>		10 ⁸	10 ⁸	10 ⁸	10 ⁸	10 ⁹	10 ⁹	10 ⁹	10 ¹⁰	10 ¹⁰	10 ¹⁰	10 ¹¹	
Feature-family		224.06	223.7	251.48	275.81	237.67	378.87	635.76	1102.48	2628.21	-	-	
Feature-product		-	-	-	-	-	-	-	-	-	-	-	
Family		-	-	-	-	-	-	-	-	-	-	-	
Family-product		-	-	-	-	-	-	-	-	-	-	-	
Product		-	-	-	-	-	-	-	-	-	-	-	
TankWar		<i>Configuration space's order</i>	10 ¹⁸	10 ¹⁸	10 ¹⁹	10 ¹⁹	10 ¹⁹	10 ²⁰	10 ²⁰	10 ²⁰	10 ²¹	10 ²¹	
		Feature-family	286.99	258.64	246.91	256.09	253.42	271.44	295.76	407.85	622.82	4104.61	
	Feature-product	-	-	-	-	-	-	-	-	-	-		
	Family	-	-	-	-	-	-	-	-	-	-		
	Family-product	-	-	-	-	-	-	-	-	-	-		
	Product	-	-	-	-	-	-	-	-	-	-		
	<i>Configuration space's order</i>	10 ²¹	10 ²¹	10 ²²	10 ²²	10 ²²	10 ²³	10 ²³	10 ²³	10 ²⁴	10 ²⁴	10 ²⁴	
	Feature-family	-	-	-	-	-	-	-	-	-	-	-	
	Feature-product	-	-	-	-	-	-	-	-	-	-	-	
	Family	-	-	-	-	-	-	-	-	-	-	-	
	Family-product	-	-	-	-	-	-	-	-	-	-	-	
	Product	-	-	-	-	-	-	-	-	-	-	-	

1217 Appendix B. SPLGenerator tool

1218 To increase the number of subject systems and inspect how each evalua-
1219 tion strategy behaves with the growth of the configuration space, we imple-
1220 mented a product-line generator tool called SPL-Generator¹⁰, which is able
1221 to create a software product line from scratch or modify an existing one by
1222 incrementally adding features and behavior to its models. For the feature
1223 model generation (i.e., to create a new feature model or change an existing
1224 one), the tool relies on the SPLAR tool [33]. The desired characteristics
1225 of the resulting feature model are obtained by defining accordingly the set
1226 of parameters provided by SPLAR. Examples of such parameters are the
1227 number of features to be created, the amount in percentage for each kind of
1228 feature (mandatory, optional, OR-inclusive and OR-exclusive), and the num-
1229 ber of cross-tree constraints. As our SPL-Generator tool intends to create
1230 product lines that resemble real-world product lines, it produces only con-
1231 sistent feature-models (i.e., the SPLAR’s parameter for creating consistent
1232 feature-models is always set to `true`).

1233 To create behavioral models, the SPL-Generator tool considers the UML
1234 behavioral diagrams and follows the refinement of activity diagrams into se-
1235 quence diagrams presented in Section 2.3. For creating activity and sequence
1236 diagrams, the generator tool is also guided by a set of parameters for each
1237 kind of behavioral diagram. For an activity diagram, it is possible to define
1238 how many activities it will comprise, the number of decision nodes, and how
1239 many sequence diagrams will refine each created activity. For a sequence
1240 diagram, it is possible to define its size in terms of numbers of behavioral
1241 fragments, the size of each behavioral fragment in terms of the number of
1242 messages, the number of lifelines, the number of different reliability values
1243 (such that each lifeline will randomly assume only one value) and the range
1244 for them. Thus, one possibly generated sequence diagram would have 5 be-
1245 havioral fragments, each one containing 8 messages between 3 lifelines, whose
1246 reliability values are within the range [0.99, 0.999].

1247 Finally, the SPL-Generator tool also provides a parameter to define how
1248 the feature model and the behavioral models will be related. The allocation
1249 of a behavioral fragment (implementing a feature’s behavior) can be fully
1250 *randomized* within the set of created sequence diagrams, or it can be *topo-*
1251 *logical*, which means the relations between the behavioral fragments mimic

¹⁰<https://github.com/SPLMC/spl-generator/>

1252 the relations between the corresponding features. In the latter, we assume a
1253 child feature refines its parent, so its behavioral fragment is nested into its
1254 parent's behavioral fragment.

1255 **References**

- 1256 [1] Apel, S., Batory, D., Kästner, C., Saake, G. (Eds.), 2013. Feature-
1257 oriented software product lines: concepts and implementation. Springer,
1258 Berlin.
- 1259 [2] Bahar, R., Frohm, E., Gaona, C., Hachtel, G., Macii, E., Pardo, A.,
1260 Somenzi, F., 1997. Algebraic decision diagrams and their applications.
1261 *Formal Methods in System Design* 10 (2), 171–206.
- 1262 [3] Baier, C., Katoen, J.-P., 2008. *Principles of Model Checking (Representa-*
1263 *tion and Mind Series)*. The MIT Press.
- 1264 [4] Bodden, E., Tolêdo, T., Ribeiro, M., Brabrand, C., Borba, P., Mezini,
1265 M., 2013. SPLLIFT: Statically Analyzing Software Product Lines in
1266 Minutes Instead of Years. In: *Proceedings of the 34th ACM SIG-*
1267 *PLAN Conference on Programming Language Design and Implemen-*
1268 *tation. PLDI '13*. ACM, New York, NY, USA, pp. 355–364.
- 1269 [5] Clarke, E. M., Grumberg, O., Peled, D. A., 1999. *Model checking*. MIT
1270 Press, Cambridge, Mass.
- 1271 [6] Classen, A., Cordy, M., Heymans, P., Legay, A., Schobbens, P.-Y., Oct
1272 2012. Model checking software product lines with SNIP. *International*
1273 *Journal on Software Tools for Technology Transfer* 14 (5), 589–612.
- 1274 [7] Classen, A., Cordy, M., Heymans, P., Legay, A., Schobbens, P.-Y., Feb.
1275 2014. Formal semantics, modular specification, and symbolic verification
1276 of product-line behaviour. *Science of Computer Programming* 80, Part
1277 B, 416–439.
- 1278 [8] Classen, A., Cordy, M., Schobbens, P.-Y., Heymans, P., Legay, A.,
1279 Raskin, J.-F., 2013. Featured Transition Systems: Foundations for Veri-
1280 fying Variability-Intensive Systems and Their Application to LTL Model
1281 Checking. *IEEE Transactions on Software Engineering* 39 (8), 1069–
1282 1089.

- 1283 [9] Classen, A., Heymans, P., Schobbens, P., Legay, A., 2011. Symbolic
1284 model checking of software product lines. In: 2011 33rd International
1285 Conference on Software Engineering (ICSE). pp. 321–330.
- 1286 [10] Classen, A., Heymans, P., Schobbens, P.-Y., Legay, A., Raskin, J.-F.,
1287 2010. Model Checking Lots of Systems: Efficient Verification of Tempo-
1288 ral Properties in Software Product Lines. In: Proceedings of the 32Nd
1289 ACM/IEEE International Conference on Software Engineering - Volume
1290 1. ICSE '10. ACM, New York, NY, USA, pp. 335–344.
- 1291 [11] Clements, P., Northrop, L., 2002. Software product lines: practices
1292 and patterns. The SEI series in software engineering. Addison-Wesley,
1293 Boston.
- 1294 [12] Cormen, T. H., Stein, C., Rivest, R. L., Leiserson, C. E., 2001. Intro-
1295 duction to Algorithms, 2nd Edition. McGraw-Hill Higher Education.
- 1296 [13] Czarnecki, K., Eisenecker, U., 2000. Generative programming: methods,
1297 tools, and applications. Addison Wesley, Boston.
- 1298 [14] Czarnecki, K., Pietroszek, K., 2006. Verifying feature-based model tem-
1299 plates against well-formedness OCL constraints. ACM Press, p. 211.
- 1300 [15] Daws, C., sep 2005. Symbolic and Parametric Model Checking of
1301 Discrete-time Markov Chains. In: Liu, Z., Araki, K. (Eds.), Proceed-
1302 ings of the First International Conference on Theoretical Aspects of
1303 Computing. Vol. 3407 of Lecture Notes in Computer Science. Springer
1304 Berlin Heidelberg, Berlin, Heidelberg, pp. 280–294.
- 1305 [16] Dimovski, A. S., Al-Sibahi, A. S., Brabrand, C., Wąsowski, A., 2015.
1306 Family-Based Model Checking Without a Family-Based Model Checker.
1307 Springer International Publishing, Cham, pp. 282–299.
- 1308 [17] Dubslaff, C., Baier, C., Kluppelholz, S., 2015. Probabilistic Model
1309 Checking for Feature-Oriented Systems. In: Chiba, S., Tanter, E., Ernst,
1310 E., Hirschfeld, R. (Eds.), Transactions on Aspect-Oriented Software De-
1311 velopment XII. No. 8989 in Lecture Notes in Computer Science. Springer
1312 Berlin Heidelberg, pp. 180–220, DOI: 10.1007/978-3-662-46734-3_5.

- 1313 [18] Dubslaff, C., Klüppelholz, S., Baier, C., 2014. Probabilistic Model
1314 Checking for Energy Analysis in Software Product Lines. In: Proceed-
1315 ings of the 13th International Conference on Modularity. MODULAR-
1316 ITY '14. ACM, New York, NY, USA, pp. 169–180.
- 1317 [19] Ferreira Leite, A., Alves, V., Nunes Rodrigues, G., Tadonki, C., Eisen-
1318 beis, C., Magalhaes Alves de Melo, A., Jun. 2015. Automating Resource
1319 Selection and Configuration in Inter-clouds through a Software Product
1320 Line Method. In: 2015 IEEE 8th International Conference on Cloud
1321 Computing (CLOUD). pp. 726–733.
- 1322 [20] Filieri, A., Ghezzi, C., 2012. Further steps towards efficient runtime
1323 verification: Handling probabilistic cost models. In: Proceedings of the
1324 First International Workshop on Formal Methods in Software Engineer-
1325 ing: Rigorous and Agile Approaches. FormSERA '12. IEEE Press, Pis-
1326 cataway, NJ, USA, pp. 2–8.
- 1327 [21] Ghezzi, C., Sharifloo, A. M., Mar. 2013. Model-based verification of
1328 quantitative non-functional properties for software product lines. *Infor-*
1329 *mation and Software Technology* 55 (3), 508–524.
- 1330 [22] Grunske, L., 2008. Specification patterns for probabilistic quality prop-
1331 erties. In: ICSE '08. ACM, New York, NY, USA, pp. 31–40.
- 1332 [23] Hahn, E., Hermanns, H., Zhang, L., 2010. Probabilistic reachability for
1333 parametric markov models. *STTT*, 1–17.
- 1334 [24] Hahn, E. M., Hermanns, H., Wachter, B., Zhang, L., 2010. Param: A
1335 model checker for parametric markov models. In: CAV. pp. 660–664.
- 1336 [25] Hao, Y., Foster, R., 2008. Wireless body sensor networks for health-
1337 monitoring applications. *Physiological Measurement* 29 (11), 27–56.
- 1338 [26] Heradio, R., Perez-Morago, H., Fernandez-Amoros, D.,
1339 Javier Cabrerizo, F., Herrera-Viedma, E., Apr. 2016. A biblio-
1340 metric analysis of 20 years of research on software product lines.
1341 *Information and Software Technology* 72, 1–15.
- 1342 [27] Iris, R., Erica, B., Frohm, A., Gaona, C. M., Hachtel, G. D., Macii, E.,
1343 Pardo, A., Somenzi, F., 1993. Algebraic Decision Diagrams and Their
1344 Applications. In: Proceedings of the 1993 IEEE/ACM International

- 1345 Conference on Computer-aided Design (ICCAD '93). IEEE Computer
1346 Society Press, Santa Clara, California, USA, pp. 188–191.
- 1347 [28] Kowal, M., Tschaikowski, M., Tribastone, M., Schaefer, I., Nov. 2015.
1348 Scaling Size and Parameter Spaces in Variability-Aware Software Perfor-
1349 mance Models (T). In: 2015 30th IEEE/ACM International Conference
1350 on Automated Software Engineering (ASE). pp. 407–417.
- 1351 [29] Kramer, J., Magee, J., Sloman, M., Lister, A., Jan. 1983. CONIC: an in-
1352 tegrated approach to distributed computer control systems. *Computers
1353 and Digital Techniques*, IEE Proceedings E 130 (1), 1–.
- 1354 [30] Liang, J. H., Ganesh, V., Czarnecki, K., Raman, V., 2015. SAT-based
1355 Analysis of Large Real-world Feature Models is Easy. In: *Proceedings
1356 of the 19th International Conference on Software Product Line. SPLC
1357 '15*. ACM, New York, NY, USA, pp. 91–100.
- 1358 [31] Linden, F. v. d., Schmid, K., Rommes, E., 2007. *Software product
1359 lines in action: the best industrial practice in product line engineering*.
1360 Springer, Berlin ; New York.
- 1361 [32] Machado, I. d. C., McGregor, J. D., Cavalcanti, Y. C., de Almeida, E. S.,
1362 Oct. 2014. On strategies for testing software product lines: A systematic
1363 literature review. *Information and Software Technology* 56 (10), 1183–
1364 1199.
- 1365 [33] Mendonca, M., Branco, M., Cowan, D., 2009. S.p.l.o.t.: Software prod-
1366 uct lines online tools. In: *Proceedings of the 24th ACM SIGPLAN
1367 Conference Companion on Object Oriented Programming Systems Lan-
1368 guages and Applications. OOPSLA '09*. ACM, New York, NY, USA, pp.
1369 761–762.
- 1370 [34] Nunes, V., Fernandes, P., Alves, V., Rodrigues, G., Sep. 2012. Variabil-
1371 ity Management of Reliability Models in Software Product Lines: An
1372 Expressiveness and Scalability Analysis. In: *2012 Sixth Brazilian Sym-
1373 posium on Software Components Architectures and Reuse (SBCARS)*.
1374 pp. 51–60.
- 1375 [35] Object Management Group, 2011. *The UML profile for MARTE: Mod-
1376 eling and analysis of real-time and embedded systems*. Version 1.1.

- 1377 [36] Pessoa, L., Fernandes, P., Castro, T., Alves, V., Rodrigues, G. N., Car-
1378 valho, H., 2017. Building reliable and maintainable dynamic software
1379 product lines: An investigation in the body sensor network domain.
1380 *Information and Software Technology* 86, 54 – 70.
- 1381 [37] Plath, M., Ryan, M., Sep. 2001. Feature integration using a feature
1382 construct. *Science of Computer Programming* 41 (1), 53–84.
- 1383 [38] Pohl, K., Böckle, G., Linden, F. v. d., 2010. *Software product line en-
1384 gineering: foundations, principles, and techniques*. Springer, New York,
1385 NY.
- 1386 [39] Rodrigues, G. N., Alves, V., Nunes, V., Lanna, A., Cordy, M.,
1387 Schobbens, P.-Y., Sharifloo, A. M., Legay, A., Jan. 2015. Modeling and
1388 Verification for Probabilistic Properties in Software Product Lines. In:
1389 2015 IEEE 16th International Symposium on High Assurance Systems
1390 Engineering (HASE). pp. 173–180.
- 1391 [40] Siegmund, N., von Rhein, A., Apel, S., 2013. Family-based Performance
1392 Measurement. In: *Proceedings of the 12th International Conference on
1393 Generative Programming: Concepts & Experiences*. GPCE '13. ACM,
1394 New York, NY, USA, pp. 95–104.
- 1395 [41] Thüm, T., Kästner, C., Benduhn, F., Meinicke, J., Saake, G., Leich,
1396 T., Jan. 2014. Featureide: An extensible framework for feature-oriented
1397 software development. *Sci. Comput. Program.* 79, 70–85.
- 1398 [42] Thüm, T., Apel, S., Kästner, C., Schaefer, I., Saake, G., Jun. 2014. A
1399 Classification and Survey of Analysis Strategies for Software Product
1400 Lines. *ACM Comput. Surv.* 47 (1), 6:1–6:45.
- 1401 [43] University of Magdeburg, O. v. G., 2011. SPL2go. Available at [http:
1402 //spl2go.cs.ovgu.de/](http://spl2go.cs.ovgu.de/), accessed: 2016-01-27.
- 1403 [44] Varshosaz, M., Khosravi, R., Dec. 2014. Model Checking of Soft-
1404 ware Product Lines in Presence of Nondeterminism and Probabilities.
1405 In: *Software Engineering Conference (APSEC), 2014 21st Asia-Pacific*.
1406 Vol. 1. pp. 63–70.
- 1407 [45] Walkingshaw, E., Kästner, C., Erwig, M., Apel, S., Bodden, E., 2014.
1408 *Variational Data Structures: Exploring Tradeoffs in Computing with*

- 1409 Variability. In: Proceedings of the 2014 ACM International Symposium
1410 on New Ideas, New Paradigms, and Reflections on Programming & Soft-
1411 ware. Onward! 2014. ACM, New York, NY, USA, pp. 213–226.
- 1412 [46] Weiss, D. M., 2008. The product line hall of fame. In: Proceedings of
1413 the 12th International Software Product Line Conference (SPLC). IEEE
1414 Computer Society, Washington, DC, USA, p. 395.