On the Relation of External and Internal Feature Interactions: A Case Study

Sergiy Kolesnikov
Sven Apel
University of Passau

Norbert Siegmund
Bauhaus-University Weimar

Christian Kästner
Carnegie Mellon University

Abstract—Detecting feature interactions is imperative for accurately predicting performance of highly-configurable systems. State-of-the-art performance prediction techniques rely on supervised machine learning for detecting feature interactions, which, in turn, relies on time consuming performance measurements to obtain training data. By providing information about potentially interacting features, we can reduce the number of required performance measurements and make the overall performance prediction process more time efficient. We expect that the information about potentially interacting features can be obtained by statically analyzing the source code of a highly-configurable system, which is computationally cheaper than performing multiple performance measurements. To this end, we conducted a qualitative case study in which we explored the relation between control-flow feature interactions (detected through static program analysis) and performance feature interactions (detected by performance prediction techniques using performance measurements). We found that a relation exists, which can potentially be exploited to predict performance interactions.

I. INTRODUCTION

A *feature* is an end-user-visible behavior or characteristic of a (software) product that satisfies a stakeholder's requirement [1]. Features are used to guide structure, reuse, and variation through the development of highly-configurable software systems [2]. While facilitating the development of highly-configurable software, reducing development costs, and improving product quality, making features optional introduces new challenges, such as the feature interaction problem [3]. A *feature interaction* occurs when the functionality of a feature or its non-functional properties (e.g., performance) are influenced by the presence or absence of one or more other features [2]. The presence of feature interactions hinders program comprehension and compositional reasoning about the functional and non-functional properties of features. That is, we cannot reason about the properties of a system configuration (i.e., a valid feature combination) in terms of a straightforward combination of the *individual* influences of the involved features on these properties. This is because we also have to consider the influences of possible interactions. A common practical scenario is searching for the best configuration of a system with respect to performance. To identify this configuration for a given operational environment, we need to know not only the individual influences of the involved features on performance, but also which interactions among these features exist and what influence on performance they have.

The problem of detecting feature interactions and quantifying their influence on performance has been addressed in the past by employing machine learning [4], [5], [6]. For building a training dataset and identifying interactions, these techniques rely on selecting a representative subset from all system configurations (i.e., sampling) and on measuring the performance of each configuration in this sample (Sec. III-C). The time needed to perform the measurements often makes up a substantial part of the overall time required by machine learning [7]. Therefore, reducing the measurement effort—by concentrating on system configurations that potentially have feature interactions—can make these techniques more time efficient and accurate [4].

The main question that we address in this paper is whether we can efficiently extract information about potentially existing feature interactions, which then can be used in performance prediction, which is, in turn, imperative to guide maintenance and evolution tasks. In our previous work [3], we described two types of interactions: (1) *external feature interactions*, which can be identified by observing the external behavior of a system, such as performance; and (2) *internal feature interactions*, which can be identified by analyzing or interpreting the source code of a system, for example, using control-flow analysis [3]. Our main hypothesis here is that there is a relation between internal and external interactions, and that we can make use of this relation to automatically identify external interactions by identifying internal interactions in a fast and efficient way. For example, multiple function calls from one feature to another (internal feature interactions) can result in a performance overhead. This performance overhead is present only if both—the caller and the callee features—are present in a configuration (external feature interaction). This way, the internal interaction is related to its external counterpart. This relation, if present, would give us hints about the existence of external feature interactions based on the internal ones. In this work, we follow up on this idea and report on at exploratory case study in which we investigated the control flow among features and its relation to performance feature interactions. We conjecture that by supplying the performance-prediction procedure with hints about which feature combinations are more likely or less likely to exhibit external feature interactions, the procedure can be made more focused on finding actual interactions.

Taking into account the *exploratory nature* of our study, the *qualitative character* of the expected results, as well as substantial technical challenges, we chose a *case-study approach* (see Sec. III-A) as our research method and two systems—the MBEDTLS encryption library and the SQLITE database engine—as non-trivial, real-world subject systems.[1] MBEDTLS and SQLITE are highly-configurable systems used

---

[1] https://tls.mbed.org/ https://www.sqlite.org/

by several large projects, such as OPENVPN and FIREFOX,[2] which makes our case study practice-oriented.

In a nutshell, using a recent machine learning technique (Sec. III-C), we learned the external (performance) feature interactions among the features of the subject systems. Furthermore, we manually inspected the code of the systems and confirmed that the learned performance interactions actually exist and that they are actually caused by the interplay of the corresponding features, and not just misinterpreted artefacts of measurement bias or environment noise.

Using a variability-aware control-flow analysis augmented by manual code inspection (Sec. III-B), we identified control-flow interactions among the features of MBEDTLS and SQLITE. That is, we identified the code locations where the features pass the control to one another.

Comparing the set of internal (control-flow) interactions with the set of external (performance) interactions revealed that those features that interact internally also interact externally (Sec. IV-B3b), which is in line with our expectation. Using the identified relation, we were able to substantially shrink the search space of performance feature interactions (Sec. IV-C). Furthermore, we made first steps towards developing an automated predictor for identifying features that are likely to interact externally based on the set of internal interactions, although, with mostly negative results (Sec. III-E). To the best of our knowledge, this is the first case study that analyzed both the external and the internal feature interactions for the same systems and investigated possible connections between these two types of interactions.

The contributions of this work are the following:

- We define a relation between internal and external interactions based on the features these interactions concern, and we discuss the plausibility of this relation.
- We define a conceptual framework for exploring the relation between internal and external interactions.
- In a first case study of this kind based on two real-world highly-configurable subject systems, we explore and confirm the relation between internal and external feature interactions.
- We discuss the implications of our findings for performance prediction of highly-configurable systems.

II. INTERNAL AND EXTERNAL FEATURE INTERACTIONS

To illustrate how features may interact internally and externally and how these interactions can be related, we use a simple example of an audio streaming system with five optional features: COMPRESS compresses the audio stream; ENCRYPT encrypts data; ADDMETADATA adds data about the stream quality, description of the audio content, information about its authors, etc., to the stream; LOGIP logs IPs of the users receiving the stream; RANKCONTENT ranks the audio content according to its popularity. The performance of the system is measured by the maximum number of users that can simultaneously receive an audio stream without the system becoming overloaded.

```
#ifdef ENCRYPT
void encrypt(payload_t *payload) {...}
#endif

#ifdef ADDMETADATA
void add_meatadata(packet_t *packet) {
  ...
#ifdef ENCRYPT
  encrypt(matadata);
#endif
  ...                    ADDMETADATA ∧ ENCRYPT
}
#endif

#ifdef LOGIP
void log(char *ip) {
  ...
#ifdef ENCRYPT
  encrypt(log_entry);
#endif
  ...                    LOGIP ∧ ENCRYPT
}
#endif

#ifdef COMRPRESS
void compress(payload_t *payload) {...}
#endif                   COMPRESS ∧ ENCRYPT

#ifdef RANKCONTENT
void rank() {...}
#endif
```

(a) Control-flow interactions in the audio streaming system.

$$100 - 15 \cdot \text{COMPRESS} - 15 \cdot \text{ENCRYPT} - 5 \cdot \text{ADDMETADATA} - 5 \cdot \text{LOGIP} - 5 \cdot \text{RANK}$$
$$-5 \cdot \boxed{\text{ADDMETADATA} \cdot \text{ENCRYPT}} \longleftarrow \text{feature interactions}$$
$$+10 \cdot \boxed{\text{COMPRESS} \cdot \text{ENCRYPT}} \longleftarrow$$

(b) A performance influence model with performance interactions.

Fig. 1: Interactions in the audio streaming system.

A. Control-Flow Interactions (Internal)

In Figure 1a, we illustrate an excerpt of the implementation of the audio streaming system. The code of each feature is delimited using C preprocessor #ifdef annotations. We denote internal interactions among features with arrows. The boxes on the arrows contain *presence conditions* for the corresponding interactions [8], that is, which features must be enabled (or disabled) for the interaction to take place. For example, if both features ADDMETADATA and ENCRYPT are enabled, then metadata are encrypted along with the audio data. For this purpose, ADDMETADATA calls the encryption function of feature ENCRYPT (denoted by the solid red arrow). Consequently, there is a control-flow interaction between these two features.

Likewise, there is a control-flow interaction between features LOGIP and ENCRYPT (denoted by the dashed green arrow), since the log entries are encrypted if both features are enabled.

Finally, an internal interaction exists between features COMPRESS and ENCRYPT (denoted by the dotted blue arrow). This is a data-flow interaction, because both features operate on the same resource (i.e., the audio stream).

B. Performance Interactions (External)

In Figure 1b, we show a *performance influence model* [9] of the audio stream system. For a given system configuration, the model can predict the maximum number of users that

can simultaneously receive an audio stream without the system becoming overloaded.[3] To calculate the predicted value, we substitute 1 for the names of the enabled features and 0 for the names of all disabled features. Then, we evaluate the arithmetic expression. For example, for the configuration with feature COMPRESS enabled and the rest of the features disabled the system can reliably serve $100 - 15 \cdot 1 - 15 \cdot 0 + \cdots + 10 \cdot 0 = 85$ users.

The individual terms of the model (i.e., the summands) describe the influence of individual features *as well as* of their interactions on the performance of the system. The first term (100) describes the performance of the base configuration (with all features disabled). The second term ($-15 \cdot$COMPRESS) describes the influence of feature COMPRESS on the performance *relative* to the performance of the base system. Thus, the computationally expensive feature COMPRESS reduces the base performance by 15.

The terms containing more than one feature (denoted by boxes in Figure 1b) describe the influences of the interactions among the involved features on performance. For example, enabling both features ADDMETADATA and ENCRYPT makes the system encrypt not only the audio stream, but also the metadata that are added to the stream. This results in a computational overhead, which reduces the system's performance by 5 users that can be served.

In our example, we assume that encrypting a small string containing an IP address is so fast that this has no measurable effect on the performance of the system. Therefore, there is no a performance interaction between features LOGIP and ENCRYPT. Consequently, there is no a corresponding term in our performance influence model.

The last term in the model describes an interaction between features COMPRESS and ENCRYPT with a positive influence of the performance. Each of the two features individually has a negative influence of -15 on the system performance, but encryption is faster if the data were compressed before. Therefore, the combined influence of both features on performance is less than the sum of their individual influences: $-15 - 15 + 10 = -20$ and not -30.

Finally, feature RANKCONTENT as well as all other possible feature combinations have no measurable influence on performance and, therefore, they are not in the performance influence model.

C. Relating Control-Flow and Performance Interactions

Table I summarizes the control-flow and the corresponding performance interactions from our example (Fig. 1). The feature combinations (ADDMETADATA, ENCRYPT) and (COMPRESS, ENCRYPT) give rise to both control-flow and performance interactions. Based on our knowledge about the implementation, we can explain the causal relation between the control-flow and performance interactions captured by these feature combinations: The call to the computationally expensive encryption functionality (a control-flow interaction) leads to the performance decrease in the configurations containing the features that implement and use the encryption functionality (i.e., a performance interaction between these features occurs). Notice that the related control-flow and performance

---

TABLE I: A lists of interacting features from Figure 1. It illustrates which of the features interact internally (control-flow interaction), externally (performance interaction), or both.

Interacting Features	Control flow	Performance
ADDMETADATA, ENCRYPT	✓	✓
COMPRESS, ENCRYPT	✓	✓
LOGIP, ENCRYPT	✓	–

interaction involve exactly the same features, so we can also relate them based on the features they involve. However, the mere presence of control flow among features does not always indicate the presence of a performance interaction. For example, the control-flow interaction between features LOGIP and ENCRYPT has no corresponding performance interaction. So, it is an open question to what extent a presence of a control-flow interaction can be used as an *indicator* for a potentially existing performance interaction.

Also note that from 26[*] feature combinations possible in the audio streaming system only three combinations give rise to feature interactions. All remaining feature combinations can be ignored by an interaction detection technique, because features in these combinations do not interact.

In what follows, we investigate *to what extent* a relation between control-flow and performance interactions exists in a real-world setting. Furthermore, we define and evaluate a predictor that uses control-flow interactions to predict potential performance interactions. With such a predictor in place, we could make interaction detection more efficient and accurate, which would be a valuable contribution to research fields, such as optimization of non-functional properties, combinatorial testing, and sampling techniques.

III. CONCEPTUAL FRAMEWORK

Next, we describe the methods and tools used in our study and how we combined them in a conceptual framework to study relations among internal and external interactions.

A. Research Method

Our study is explorative in nature, because the way we study the relation between control-flow and performance interactions is novel and requires careful validation in itself. Our study involves bleeding-edge techniques for detecting control-flow and performance interactions that are technically challenging and, therefore, cannot be easily applied to a large number of non-trivial real-world systems. By focusing on two systems, we aim at increasing the internal validity of the study, because, this way, we can better identify and control confounding effects that may vary substantially from one subject system to another (e.g., architecture, size of features). Ultimately, we want to obtain deep insights in the nature of the relation between control-flow and performance feature interactions and not only report on shallow statistics. Taking these characteristics of our study into account, the literature on conducting empirical research in software engineering suggests a *case study* as a research method: Shull et al. describe case studies as "initial investigations of some phenomena" [10], Yin extends the description by stating that a case study is "an empirical

---

inquiry that investigates a contemporary phenomenon within its real-life context, especially when the boundaries between phenomenon and context are not clearly evident" [11], and Flyvbjerg adds that "case studies offer in-depth understanding of how and why certain phenomena occur, and can reveal the mechanisms by which cause–effect relationships occur" [12].

B. Identifying Control-Flow Interactions

To identify control-flow interactions, we use a variability-aware call-graph analysis [13] implemented in TYPECHEF[4] that identifies function calls among features implemented with preprocessor annotations (Fig. 1a). The central idea of a variability-aware analysis is to achieve efficiency by analyzing code parts that are shared by multiple system configurations only once. This is achieved by analyzing the source code of the system that still contains variability (e.g., the code with preprocessor annotations in Figure 1a), as opposed to analyzing the source code of individual configurations, which may be exponentially many in the number of features. A variability-aware call-graph analysis provides an efficient way to identify function calls among features of a highly-configurable system and makes the detection of internal interactions feasible.

The underlying data structure for the analysis is *the variable abstract syntax tree*. Similar to an abstract syntax tree (AST), a variable AST provides an abstraction of the source code that can be efficiently analyzed, but it also provides information on which part of the code belongs to which features (in the form of presence conditions). Using this information, a call-graph analysis can identify, for each function call, which feature is the caller and which feature is the callee. Furthermore, the analysis can identify a presence condition for each call, that is, which features must be enabled (or disabled) for the call to take place at runtime. For example, in Figure 1a, the call from feature ADDMETADATA to feature ENCRYPT (solid red arrow) occurs only if both features ADDMETADATA and ENCRYPT are enabled (denoted by the presence condition in the box under the arrow). Due to the static nature of the technique, the collected information about the calls may be an overapproximation, but this is a problem with any static analysis approach. The current implementation of the analysis also uses pointer analysis to increase the accuracy of the call graph [13].

C. Identifying Performance Interactions

For detecting performance feature interactions, we learn performance influence models (Fig. 1b). As discussed in Section II-B, a performance influence model captures the influences of individual features and their interactions on performance of a configurable system. We learn performance influence models using the tool SPL CONQUEROR,[5] which implements a state-of-the-art machine learning algorithm based on multivariate regression and forward feature selection [9]. The algorithm takes as input a sample of system configurations and corresponding performance measurements. The accuracy of the learned performance influence model depends, among other factors, on how representative the sampled configurations are for the entire configuration space. To get a performance

influence model of the highest possible accuracy, and, consequently, to detect feature interactions as precise as possible (i.e., to obtain the ground truth), we measured not a sample but all configurations of the subject system and used these measurements as the algorithm input. The performance measurements were done using a standard benchmark.

To build an influence model, SPL CONQUEROR starts with calculating a set of features and their combinations that can be included in the model to reduce the model's prediction error. For example, COMPRESS · ENCRYPT in Figure 1b is a feature combination that has been eventually included in the model during the learning process. The algorithm iterates over the set of features and their combinations and selects one element of the set that explains variations in the performance of the system best; that is, the element yielding the model's lowest prediction error, when incorporated into the model. The selection of candidates continues until either a predefined accuracy is reached or all features and feature combinations that could reduce the prediction error of the model have been considered. For a more in-depth description of the algorithm, we refer the reader to previous work [9].

D. Relating Control-Flow and Performance Interactions

After we have identified the internal (control-flow) interactions, the question is what we can learn from them regarding external (performance) interactions. To answer this question, we relate the control-flow interactions and performance interactions based on the features involved in them, as we explained it in our example in Section II-C. The goal is to find out if the features involved in performance interactions also occur in one or more internal interactions and vice versa. This is a feasibility check to see if the interactions can be related based on the features' occurrence at all. That is, if we find no interactions that can be related in this way, this would mean that it is impossible to define any relation between interactions based on the corresponding feature occurrences in these interactions.

We define a performance interaction i_p and a control-flow interaction i_c as related if $features(i_p) \subseteq features(i_c)$ or if $features(i_p) \supseteq features(i_c)$, where $features(i)$ is the set of features that contribute to the interaction i.

Furthermore, for each related pair of interactions, we determine how similar the interactions are (i.e., if they contain exactly the same features or if they also contain features that are present only in one of them). The similarity of the related interactions can be interpreted as the strength of their relation: the higher the similarity, the higher the strength of the relation. We calculate the similarity of interactions using the Jaccard index J [14]:

$$J(i_p, i_c) = \frac{features(i_p) \cap features(i_c)}{features(i_p) \cup features(i_c)}$$

where $features(i)$ is the set of features involved in the interaction i. The Jaccard index equals 1 if both interactions involve exactly the same features and is less than 1 otherwise.

E. Predicting Performance Interactions

If we find a relation between control-flow and performance feature interactions as defined in Section III-D, the question is whether we can use this relation to predict performance feature interactions.

One method is to build on our argumentation in Section II-C and to assume that every control-flow interaction corresponds to an existing performance interactions. Of course, we already know that there may be control-flow interactions without corresponding performance interactions. Nevertheless, it is an open question *how* accurate this simple method can be if applied to a real-world system.

We can also use a more advanced method based on reoccurring feature combinations in control-flow interactions: We argue that, if a set of features occurs in multiple control-flow feature interactions, then this set of features is also likely to give rise to one or more external interactions. The rationale behind this argument is that, if a set of feature is involved in many control-flow feature interactions, then chances are high that it is also involved in performance interactions, because the accumulated influence of the control-flow interactions on performance have a measurable effect.

The method that we use to identify such frequent feature sets is *frequent item set mining* [15], which has been successfully used as a general pattern mining method [16], [17]. In terms of frequent item set mining, we refer to a feature as an *item*. The set of all items (all features) is the *item base* B. A subset of the item base $I \subseteq B$ is an *item set* that corresponds to a feature combination. An item set (i.e., a feature combination) that denotes an internal interaction in a system is a *transaction* $t \in T$, where T is a set of transactions. Based on these definitions, we define the *support* (a.k.a. absolute frequency) s of an item set: $s = |\{t : t \in T \wedge I \subseteq t\}|$. The support value and a threshold $E \in [0, \infty)$ is used to decide which of the item sets are considered frequent: All item sets with the support value $s \geq E$ are *frequent* item sets. Based on our hypothesis, frequent item sets predict external feature interactions. We use an implementation of the Apriori algorithm from the ORANGE library[6] to calculate the support value.

IV. CASE STUDY

We address the following research questions in our study:
- **RQ1:** Do control-flow feature interactions and performance feature interactions relate (in terms of the definition of Section III-D)?
- **RQ2:** If the relation identified in RQ1 exists, what is the predictive power (i.e., precision and recall) of the predictors that we propose in Section III-E?

To answer these research questions, we search for control-flow and performance interactions in the subject systems, study their properties, identify relations between them, and evaluate predictors for performance interactions based on these relations.

A. Subject Systems

The case study was conducted using two real-world highly-configurable software systems: the MBEDTLS library implementing the transport security network protocol TLS/SSL and a SQL database engine SQLITE. The initial use case for the systems was embedded domain, but now they are also used in non-embedded projects, such as OPENVPN and FIREFOX.

Similar to a large number of other real-world highly configurable systems, the subject systems are written in C using

[6]http://orange.biolab.si/

C-preprocessor directives to implement compile-time variable features. MBEDTLS comprises 50 K and SQLITE 195 K lines of code. Both systems have a highly modular architecture, which is thoroughly documented along with the corresponding preprocessor macro names allowing relatively easy matching of code to the corresponding modules and submodules.

a) MBEDTLS *Features and Feature Model:* At the top level, MBEDTLS consists of modules, such as *Cipher*, *Public Key*, *Hashing*. Each module implements the corresponding algorithms and protocols. For example, the *Cipher* module includes submodules that implement cipher algorithms, such as AES, DES, and ARC4. Submodules implement the features of the system. The cipher-algorithm features can be combined with other features, such as hash algorithms and public-key implementations, to provide an encryption protocol. We used the original documentation and manual code inspection to construct a feature model for MBEDTLS version 2.2.1, comprising 97 features and 1921 configurations.

b) SQLITE *Features and Feature Model:* SQLITE consists of a *Core* providing a C-language interface and being responsible for executing compiled SQL code, an *SQL Compiler*, and a *Backend* providing the low-level implementation of the database. A user can configure the operation of these modules by enabling or disabling their features through compile-time options. For example, *Core* can be configured to operate safely in a multithreaded environment by enabling the SQLITE_THREADSAFE feature. We studied the documentation and the source code of version 3.16.2 to construct a feature model comprising 12 features and 1533 configurations.

c) MBEDTLS *Performance Measurements:* The primary application of MBEDTLS is the encryption of data transmitted over a TCP/IP network. Ensuring fast and secure data transfer is commonly considered an important property of communication networks, such as the Internet. So, the time required to encrypt data and transfer them over the network is an important non-functional property of MBEDTLS. Measuring the time required by encryption alone is not representative, because different configurations may produce different amounts of payload (e.g., due to data compression and different amounts of generated metadata) influencing the transmission time. Therefore, we defined the performance measure for a configuration of MBEDTLS as the amount of time (in seconds) required to encrypt and successfully transmit a fixed amount of input data.

To detect performance feature interactions reliably based on performance benchmarks, it must be ensured that every feature included in a configuration is invoked during the benchmark of this configuration. Otherwise, the influence of features and their interactions on performance cannot be deduced from the benchmark results. The original automated test framework of MBEDTLS includes tests that check the library's functionality in a client-server environment and is suitable to serve as a typical benchmark suite. During the tests, the functionality of every feature in the configuration is tested, that is, every feature is actually invoked.

We used 2 GB of random data as input to ensure that the fastest configuration requires, at least, five seconds for transmission and to mitigate the influence of warm-up effects on the result. We repeated the benchmark 30 times to further

reduce the influence of measurement bias. To exclude the influence of network latencies, we ran the benchmark locally using the local network interface.

d) SQLITE *Performance Measurements:* The developers of SQLITE provide a performance benchmark that measures time required by the database to execute a set of queries.[7] The original benchmark is not compatible with the latest version of the system that we use, so we used it as guidance to create a new compatible benchmark. While constructing the benchmark we made sure that the features of SQLITE are actually invoked during the benchmarking process. Our benchmark measures the execution time in seconds. To reduce the influence of warm-up effects and measurement bias, the benchmark runs, at least, 25 seconds and every run is repeated 30 times.

The benchmarks for both systems were conducted on an Intel i5-4590, 16 GB RAM, 256 GB SSD, Ubuntu 16.04.

B. Results

1) Performance Interactions: We used SPL CONQUEROR and the performance benchmark results (cf. Sec. IV-A) as input data to identify performance interactions in MBEDTLS and SQLITE, as described in Section III-C. Table II lists for both systems the performance interactions and their influences on performance of the systems in seconds. The *negative values* in the influence column denote *positive influences* of the corresponding interactions on performance. That is, they denote how much less time a configuration that includes them would need to execute the benchmark.

The mean standard deviation for the performance measurements of MBEDTLS is 0.42 s. Therefore, we classified all interactions with the absolute influences less than this value as noise and discarded them. From the remaining 16 interactions, 11 are interactions between two features; and five are interactions among three features. The mean standard deviation for the performance measurements of SQLITE is 0.09 s. The influences of the three identified interactions for the system are much higher and, therefore, are unlikely to be noise. Two of the interactions are interactions between two features and one is an interaction between three features.

a) MBEDTLS*:* All identified interactions in MBEDTLS are among features implementing different ciphers, block cipher modes of operation (simply "modes"), and cryptographic hash functions. This is plausible, because these three types of algorithms work tightly together to implement an encryption protocol. Ciphers (e.g., AES) are used to encrypt data, modes (e.g., CBC) are used in combination with block ciphers to encrypt amounts of data larger than a block (i.e., a fixed amount of data a block cipher can operate on; 128 bit for AES), and cryptographic hash functions (e.g., SHA) are used with modes to implement authentication and to ensure data integrity.

To confirm that the identified performance interactions actually result from the interplay of the corresponding features, we manually inspected the source code of MBEDTLS. Next, we present the results of this code inspection.

Interaction 1 in Table II arises between a mode (CBC) and a hash function (SHA256). CBC uses hashing extensively to calculate keyed-hash message authentication code (HMAC). SHA256 is computationally more expensive than, for example,

[7]http://sqlite.org/speed.html

TABLE II: Performance interactions, their influences on performance of the systems in seconds, the number of the control-flow interactions related to them, and the mean value of the corresponding Jaccard indexes. (') marks the relations for manually added control-flow interactions.

	ID	Influence (sec)	Performance Interaction (features involved)	Rela-tions	Jaccard (mean)
MBEDTLS	1	10.73	CIPHER_MODE_CBC, SHA256_C	1'	1.00
	2	-9.71	AES_C, AESNI_C	10	0.53
	3	8.53	AESNI_C, SSL_CBC_RECORD_SPLITTING	2	0.38
	4	6.93	CIPHER_MODE_STREAM, AESNI_C	1'	1.00
	5	6.08	SHA256_C, CIPHER_MODE_STREAM	1'	1.00
	6	5.75	AES_C, AESNI_C, GCM_C	4	0.53
	7	3.49	CIPHER_MODE_CBC, SHA256_C, SHA256_SMALLER	1'	1.00
	8	3.45	SHA256_C, CIPHER_MODE_STREAM, SHA256_SMALLER	1'	1.00
	9	3.44	SHA256_C, AESNI_C, CIPHER_MODE_STREAM	1'	1.00
	10	3.14	CIPHER_MODE_CBC, RIPEMD160_C	1'	1.00
	11	-2.97	AES_C, GCM_C	13	0.40
	12	-2.84	CIPHER_MODE_STREAM, MD5_C	1'	1.00
	13	1.93	AESNI_C, CAMELLIA_C	4	0.35
	14	1.68	CIPHER_MODE_CBC, SHA1_C	1'	1.00
	15	1.60	CIPHER_MODE_STREAM, AESNI_C, MD5_C	1'	1.00
	16	1.51	RIPEMD160_C, CIPHER_MODE_STREAM	1'	1.00
SQLITE	1	1.50	DEFAULT_MEMSTATUS, THREADSAFE	1'	1.00
	2	1.47	MEMDEBUG, THREADSAFE	16	0.45
	3	1.41	DEFAULT_MEMSTATUS, MEMDEBUG, THREADSAFE	1'	1.00

To relate the influences to configuration run times, note that the fastest MBEDTLS configuration completed its benchmark in 6.7 seconds and the fastest SQLITE configuration completed its benchmark in 26.7 seconds.

MD5; therefore, this combination with the mode has a negative influence of 10.73 seconds on performance. Interactions 5, 7, 8, 10, 12, 13, 14, and 16 have a similar cause and explanation. In addition to a mode and a hash function, interactions 7 and 8 also include the feature SHA256_SMALLER, which denotes that an implementation of SHA256 with smaller binary footprint was used. However, this implementation also has a lower performance, which leads to the negative influence of this interaction on performance. Interaction 12 has a positive influence on performance of using a mode (stream mode, in this case) with a less computationally complex (but also less secure) MD5 hash function. In interaction 13, the AES cipher is used as a hash function in combination with the Camelllia cipher.

Interaction 2 arises from the usage of the AES cipher for encryption in combination with a native implementation of the AES algorithm in assembler (AESNI). The native implementation makes encryption faster, so this interaction has a positive influence of 9.71 seconds on performance.

Interaction 3 arises from the usage of the AES cipher for encryption in combination with an implementation of the CBC mode that includes a record splitting algorithm. This algorithm is a countermeasure against the BEAST attack on the SSL algorithm. The way record splitting is implemented increases the number of packets to be transmitted (compared to the number of packets without this countermeasure). The increased number of packets results, in turn, in a negative influence on performance.

Interactions 4, 6, 9, 11, and 15 arise from the influence of

further combinations of ciphers, modes, and hash functions on performance, similar to the first interaction.

b) SQLite: All performance interactions in SQLite include the feature THREADSAFE. This is plausible, because THREADSAFE is a crosscutting feature that adds the mutex and thread-safety logic to all unsafe regions in the code. This additional thread-safety code imposes a runtime overhead and makes the benchmarks for the configurations containing it run longer. We inspected the code of SQLite and confirmed that both features DEFAULT_MEMSTATUS and MEMDEBUG retrieve a mutex (i.e., use THREADSAFE feature) at a certain stage of operation that results in interaction among THREADSAFE and these features.

---

c) Summary: Overall, we identified 16 performance interactions in MBEDTLS. 11 of them occur between 2 features and 5 among 3 features. In SQLite, we identified 3 performance interactions. 2 interactions between 2 features and 1 among 3 features. Using domain knowledge and manual inspection of the source code, we identified the cause of all interactions and thereby confirmed that they actually exist in the systems and are caused by the interplay of the corresponding features.

---

2) Control-Flow Interactions: We used the variability-aware call-graph analysis implemented in TYPECHEF (Sec. III-B) to detect control-flow interactions in MBEDTLS and SQLite.

a) MBEDTLS: From 761 992 function calls in the system, we detected 575 560 control-flow feature interactions. This number of interactions includes duplicate interactions that appear if the corresponding function call between features occurs in multiple locations in the code. The number of unique control-flow interactions is 73.

Notably, among the unique control-flow interactions, there are interactions with up to 10 features, but most unique interactions involve only two features (Fig. 2a). If we also consider the duplicates (Fig. 2b), the overall picture stays largely the same: Only the number of interactions involving four features becomes larger than the number of those involving three features.

While manually exploring the source code of MBEDTLS, we found that cipher, mode, and hash algorithms call each other indirectly, using function pointers. This indirection was introduced by the designers of the library to decouple the algorithms and to make their concrete implementations interchangeable. TYPECHEF would need to be extended with a variability-aware, inter-procedural data-flow analysis to identify which features interact using indirect function calls. Being aware of this technical limitation of TYPECHEF, we added 11 indirect control-flow interactions that we collected while manually exploring the code to the set of interactions. Therefore, the total number of the identified unique control-flow interactions is 84 (73 were found using TYPECHEF and 11 manually). It would be infeasible to find manually all instances of indirect control-flow interactions, so their exact number (including duplicates) is unknown. We discuss the corresponding threats to validity in Section V.

*b) SQLite:* From over 14 587 337 function calls in the system, we detected 14 587 335 control-flow feature interac-

(a) Unique interactions only   (b) All interactions

Fig. 2: MBEDTLS: counts of features in control-flow interactions.


(a) Unique interactions only   (b) All interactions

Fig. 3: SQLite: counts of features in control-flow interactions.

tions. That is, all but two function calls involved more than one feature. The number of unique control-flow interactions is 37.

In contrast to MBEDTLS, most unique interactions involve 4 features, and there are interactions with up to 6 features (Fig. 3a). Although, if we also consider duplicates (Fig. 3b), the picture becomes similar to that in MBEDTLS: Interactions among 2 features prevail and the count of interactions decreases with the increasing number of involved features.

While manually inspecting the code of SQLite, we found that the option SQLITE_DEFAULT_MEMSTATUS (which is used by TYPECHEF to identify the code belonging to the feature DEFAULT_MEMSTATUS) is used to set a Boolean variable at compile-time. This variable is then used at runtime to check if feature DEFAULT_MEMSTATUS is enabled or disabled. This way, the feature can be enabled or disabled at runtime. Again, TYPECHEF would need a data-flow analysis to trace the connection the preprocessor macro to the corresponding Boolean variable to detect control-flow interactions in which feature DEFAULT_MEMSTATUS is involved. By further exploring the code, we identified two control-flow interactions of this kind and added them to the set of automatically detected interactions. Therefore, the total number of the identified unique control-flow interactions is 39.

---

c) Summary: Overall, we identified 575 571 control-flow interactions in MBEDTLS among which 84 were unique. Some interactions involve up to 10 features, but most interactions are between 2 features. For SQLite, we identified 14 587 335 control-flow interactions, with 39 unique. Due to technical limitations of TYPECHEF, indirect control-flow interactions in MBEDTLS and interactions induced by runtime variability in SQLite could not be detected by TYPECHEF. We manually inspected the source code to collect these interactions.

---

3) Relating Interactions:

a) Performance Interactions → Control-Flow Interactions.: Using the relation definition $features(i_p) \subseteq$

$features(i_c)$ (Section III-D), for each performance interaction, we identified all unique related control-flow interactions (i.e., all control-flow interactions involving exactly the same features as the performance interaction). Furthermore, for each pair of related interactions, we calculated the Jaccard index (Section III-D), which denotes how similar the interactions are (the index equals 1 if both interactions involve exactly the same features and is less than 1 otherwise).

Table II summarizes the results. For each performance interaction, it shows the number of the related control-flow interactions and the mean of all Jaccard indexes calculated for these relations. The apostrophe (') denotes the performance interactions that are related to the manually identified indirect control-flow interactions for which we were not able to establish the exact number of occurrences (cf. Section IV-B2). The numbers show that there is a relation between every performance interaction and, at least, one control-flow interaction. The Jaccard indexes show further that the related control-flow interactions, which were automatically detected by TYPECHEF, involve, on average, twice as many or even more features than there are in the corresponding performance interactions.

b) Control-Flow Interactions → Performance Interactions.: Using the relation definition $features(i_p) \supseteq features(i_c)$ (Section III-D), for each control-flow interaction, we identified all related performance interactions (i.e., all performance interactions involving exactly the same features as the control-flow interaction).

Table III summarizes the results. For MBEDTLS, among the 84 unique control-flow interactions, we found 4 interactions that have one or more related performance interactions. For SQLITE, among the 39 unique control-flow interactions, we found 2 interactions that have one or more related performance interactions. The Jaccard indexes show that the related performance interactions that were automatically detected by TYPECHEF involve mostly the same features as the corresponding control-flow interactions. The manually added control-flow interactions match exactly the related performance interactions.

---

c) Summary: We found a relation between every of the 16 identified performance interactions and one or more control-flow interactions. The Jaccard indexes for the related interactions show that the interactions do not generally contain exactly the same features and that the related control-flow interactions involve, on average, twice as many features than there are in the corresponding performance interactions.

---

4) Predicting Performance Interactions:

a) MBEDTLS Direct Matching: As we describe in Section III-E, one prediction method is to assume that every control-flow interaction induces a performance interaction that involves exactly the same features. In MBEDTLS, from the 73 automatically identified unique control-flow interactions there are three—interactions 1, 3, and 4 in Table III—that have exactly the same features as the related performance interactions 2, 6, and 11 in Table II. That is, three of the 16 performance interactions could be predicted by the direct matching. Therefore, the precision of the direct matching is 4.11 % and the recall is 18.75 %. If we also incorporate the 11 indirect control-flow interactions, which we identified by manually inspecting the code, the total number of match-

TABLE III: Control-flow interactions, the number of the related performance interactions, and the mean value of the corresponding Jaccard indexes. Control-flow interactions without related performance interactions are not listed. (') marks the relations for manually added control-flow interactions.

	ID	Control-Flow Interaction (features involved)	Relations	Jaccard (mean)
MBEDTLS	1	AES_C, AESNI_C	2	0.83
	2	GCM_C, AESNI_C	1	0.67
	3	GCM_C, AES_C	2	0.83
	4	GCM_C, AES_C, AESNI_C	1	1.00
	5	CIPHER_MODE_CBC, SHA256_C	1'	1.00
	6	CIPHER_MODE_STREAM, AESNI_C	1'	1.00
	7	SHA256_C, CIPHER_MODE_STREAM	1'	1.00
	8	CIPHER_MODE_CBC, SHA256_C, SHA256_SMALLER	1'	1.00
	9	SHA256_C, CIPHER_MODE_STREAM, SHA256_SMALLER	1'	1.00
	10	SHA256_C, AESNI_C, CIPHER_MODE_STREAM	1'	1.00
	11	CIPHER_MODE_CBC, RIPEMD160_C	1'	1.00
	12	CIPHER_MODE_STREAM, MD5_C	1'	1.00
	13	CIPHER_MODE_CBC, SHA1_C	1'	1.00
	14	CIPHER_MODE_STREAM, AESNI_C, MD5_C	1'	1.00
	15	RIPEMD160_C, CIPHER_MODE_STREAM	1'	1.00
SQLITE	1	DEFAULT_MEMSTATUS, THREADSAFE	1'	1.00
	2	DEFAULT_MEMSTATUS, MEMDEBUG, THREADSAFE	1'	1.00

ing control-flow interactions becomes 14. Including indirect control-flow interactions increases the precision and recall to 16.7 % and 51.85 % respectively.

b) SQLITE Direct Matching: In SQLITE, there are no automatically identified unique control-flow interactions that match exactly any of the performance interactions. Including the manually added control-flow interactions gives the prediction precision of 5.13 % and the recall of 67 %.

c) MBEDTLS Frequent Item Sets: Using frequent item set analysis (cf. Sec. III-E) on the set of control-flow interactions for MBEDTLS, we found 44 item sets, of which we calculated the support values. The support values range from 11 % to 34 %, meaning that there are item sets occurring in 11 % to 34 % of all control-flow interactions.

Two of the found item sets match exactly the performance interactions 2 and 11 of Table II. Notice that we ran the frequent item set analysis only on the automatically detected control-flow interactions. We were not able to run it on the indirect control-flow interactions, because then we would have to find every instance of such interaction manually, which is infeasible. Nevertheless, we incorporated the indirect control-flow interactions into further analysis by approximating their support values based on the distribution of support values for similar indirect interactions (see Sec. V, for threats to validity). Among the 44 detected item sets, there are 33 item sets capturing interactions among ciphers, modes, and hash functions. We assigned support values to the indirect control-flow interactions according to the distribution of the support values of these 33 item sets. That is, 6 % of the interactions were assigned a support value of 11 %, 3 % were assigned a support value of 12 %, and so on.

By varying the threshold E, as described in Section III-E, we are able to decide which of the identified item sets are considered frequent. By setting the threshold to 0, we consider

all identified item sets as frequent. When the threshold is increased the item sets with lower support values are not considered frequent anymore. For example, if we set the threshold to 15 % only 25 % of the identified item sets will have a higher support value and will be considered frequent. Changing the threshold this way allows us to observe its influence on the predictive power (i.e., precision and recall) of the frequent item sets.

To calculate how good the item sets are in predicting performance interactions, we compared how many of them denote the actually identified performance interactions (i.e., contain exactly the same features as the performance interactions). The low precision and recall values for MBEDTLS summarized in Table IV show that our predictor based on the frequent item sets has only a low predictive power. Increasing the threshold value decreases the precision and recall of the predictor.

d) SQLITE *Frequent Item Sets:* Applying the same frequent item set method to the control-flow interactions of SQLITE resulted in four frequent item sets with support values ranging from 20 % to 100 %. None of these frequent item sets matched the performance interactions. We could not approximate the distribution of the support values for the manually detected control-flow interactions, because they do not exhibit any commonalities with the calculated frequent item sets as it was the case for MBEDTLS.

> *e) Summary:* We defined two predictors for performance interactions based on their relation with control-flow interactions. The first predictor is based on the assumption that every control-flow interaction induces a performance interaction that involves exactly the same features. The second predictor is based on the assumption that the recurring feature combinations in control-flow interactions capture the related performance interactions. The evaluation showed that both predictors have only low precision and recall values.

C. Discussion

Based on our results, we conclude that there is indeed a quantifiable relation between control-flow and performance interactions. We confirmed this by manually inspecting the code and by comparing which features are involved in the detected performance interactions and how these features interact at the control-flow level. We found that features involved in performance interactions work closely together to implement the systems' functionality and also interact at the control-flow level. That is, the same features that are involved in performance interactions are also involved in control-flow interactions. Therefore, we can positively answer research question RQ1.

The relation we found among control-flow and performance feature interactions has implications for performance prediction techniques for highly-configurable systems. As we discussed in Section IV-B3, the identified control-flow interactions capture the features that are involved in the performance interactions. Of course, we cannot identify these features precisely, because the same control-flow interactions also involve other features that are not involved in performance interactions (this is also a reason for direct matching prediction having low precision and recall; cf. Sec IV-B4). Nevertheless, assuming that only the features from the identified control-flow interac-

TABLE IV: Precision and recall values for the item sets as predictors for the performance interactions in MBEDTLS. (*) marks the precision and recall values for the item sets with incorporated indirect control-flow interactions.

Threshold	Precision	Recall	Precision*	Recall*
0	4.5	12.5	23.6	48.1
15	2.3	6.3	5.5	11.1
20	0	0	1.8	3.7

tions can give rise to a performance interaction considerably reduces the search space of the potential performance feature interactions, because otherwise we have to assume that any (valid) feature combination may give rise to a performance interaction. MBEDTLS has 134 057 valid feature combinations of two and three features, but the 84 identified unique control-flow interactions (Sec. IV-B2) result in only 452 *potential* performance interactions (among two and three features). Notice that these include all 16 actually existing performance feature interactions that we identified. That is, we are able to shrink the search space of performance feature interactions by almost 300 times (452 instead of 134 057) without losing any of the actually existing performance feature interactions. SQLITE has 524 valid feature combinations of two and three features and (based on the 39 identified unique control-flow interactions) only 131 *potential* performance interactions (among two and three features). These potential performance interactions also include all 3 actually existing performance interactions that we identified. That is, the search space shrinks by 4 times. These results have immediate consequences for performance prediction techniques based on machine learning and relying on sampling for building a training dataset: By exploiting our findings they can make sampling more focused on the configurations that potentially include interacting features, which may improve their prediction accuracy.

With respect to RQ2, we conclude that predictors based on direct matching and frequent item sets have low precision and recall values. One possible reason is that the predictors rely solely on control-flow data, but features can also interact via data flow. For example, they can exchange data through shared data structures. This interplay at the data-flow level can be interpreted as data-flow feature interactions (much like control-flow feature interactions), which may also induce performance interactions. For example, a feature may block other features by locking a shared data structure, which may have a negative influence on the performance of the system. Therefore, enriching the data used by the predictors with the information about data-flow interactions may increase their predictive power. So, a takeaway message here is that predictors should consider the interplay of features not only on the control-flow level, but also at the data-flow level, and other levels. Another reason may be that not all features involved in a control-flow interaction are also involved in a related performance interaction. The Jaccard index values in Table II show that only about half of the features in a control-flow interaction are also present in the related performance feature interaction. For example, the interaction (AES_C, AESNI_C) has the average Jaccard index of 0.46. This means that, on average, a related control-flow interaction has two other features additionally involved,

in addition to features AES_C and AESNI_C. Both predictors for a given control-flow interaction are not able to distinguish among features that are involved in a related performance interaction and those that are not.

a) Further Observations.: A further observation is related to the distribution of the number of features involved in the control-flow and performance interactions. For MBEDTLS, in most cases, interactions (both, control-flow and performance) involve two or three features. For SQLITE, in most cases, control-flow interactions involve four features, but this is only the case because every single control-flow interaction involves the two crosscutting features THREADSAFE and ENABLE_API_ARMOR. If we ignore these crosscutting features, the pictures becomes similar to MBEDTLS. The performance interactions in SQLITE involve two or three features as in MBEDTLS. From these data, we conclude that the frequency of interactions decreases with the growing number of the involved features. This shows that features tend to interact at the same rate (two or three features per interaction) independently of the type of the interaction (control-flow or performance). This is another indication for a relation between control-flow and performance interactions.

Finally, for MBEDTLS, we found that most of the frequent item sets that we identified in the control-flow interactions contain features from three groups of algorithms: ciphers, modes, and hashes. Even though most of the frequent item sets do not resemble existing performance interactions, they still capture the general pattern of the detected performance interactions, namely, that these interactions involve features from these three groups of algorithms. For SQLITE the frequent item sets capture the crosscutting features, such as THREADSAFE and ENABLE_API_ARMOR. The crosscutting feature THREADSAFE was involved in all identified performance interactions.

V. THREATS TO VALIDITY

a) Internal Validity.: Due to technical limitations of TYPECHEF, we were unable to identify the exact number of indirect function calls between features (i.e., calls made using function pointers) and, consequently, the exact support values for the corresponding item sets (Sec. IV-B4). We approximated these support values based on the distribution of the support values for the item sets calculated from direct function calls. Our approximation method may result in an inaccurate calculation of the precision and recall values of the frequent item set predictor. Nevertheless, we expect that improving the approximation would rather improve the precision and recall of the predictor.

b) External Validity.: Our study is exploratory in nature and aims at the initial investigation of the relation between control-flow and performance feature interactions. Since we focused on analyzing two systems, our results may not hold for other highly-configurable systems. Our study setup can serve as a blueprint for further studies that can rely on our conceptual framework for studying relations among external and internal interactions.

VI. RELATED WORK

In recent years, a number of papers aimed at detecting feature interactions in highly-configurable systems. We summarize and subdivide them according to our classification [3] into those considering internal feature interactions and those considering external feature interactions. To our best knowledge, there is no work that study these two types of interactions in combination, as we do it in this study.

a) Internal Feature Interactions: Detection of internal feature interactions is often used by techniques that aim at minimizing test-suite and test-effort for highly configurable systems. Reisner et al. [18], Nguyen et al. [19], Tartler et al. [20] apply symbolic evaluation, dynamic and static program analysis respectively to infer minimal sets of features responsible for a given code coverage. Kim et al. [21] apply static program analysis to identify features that do not interact with other features w.r.t. to the test-suite. Garvin et al. explores a connection between feature interactions and interaction faults [22]. Lillack et al. extends static taint analysis to automatically identify interactions among load-time configuration options [23].

b) External Feature Interactions: A number of recently proposed performance prediction techniques for highly configurable systems by Guo et al. [5], Siegmund et al. [4], Sarkar et al. [24], Thereska et al. [25], Westermann et al. [26], and Zhang et al. [6] use machine-learning techniques, such as, CART, multivariate regression, and Fourier learning, for learning a performance function based on the performance measurements of a configuration sample. These techniques learn performance (external) feature interactions as an integral part of the overall black-box learning process, that is, without considering the internal feature interactions.

VII. CONCLUSION

In our case study we explored the relation among control-flow and performance feature interactions that occur in highly configurable systems. Using the encryption library MBEDTLS and the database engine SQLITE as real-world subject systems, we identified control-flow and performance feature interactions using static program analysis and machine learning. Analyzing the interactions, we found that they can be related based on the involved features. By manually inspecting the code, we confirmed the causal relation between the interplay of features at the control-flow level and the identified performance interactions among the same features. Furthermore, based on the identified relation, we defined two predictors for performance feature interactions and conducted a preliminary evaluation of these predictors. The evaluation showed that the predictors have low precision and recall, presumably, because features also interact at the data-flow level. Future predictors based on the internal feature interactions should consider both control-flow and data-flow interactions to improve their predictive power.

Beside this negative result, using the identified relation among control-flow and performance feature interactions, we are still able to shrink the search space of performance feature interactions (by almost 300 times for MBEDTLS and by 4 times for SQLITE) without losing any of the performance feature interactions actually existing in our subject systems. Performance prediction techniques that rely on sampling can use our results to make their sampling more focused on configurations with potential performance interactions.

REFERENCES

[1] K. Kang, S. Cohen, J. Hess, W. Novak, and A. Peterson, "Feature-Oriented Domain Analysis (FODA) Feasibility Study," Carnegie Mellon University, Tech. Rep. CMU/SEI-90-TR-21, 1990.

[2] S. Apel, D. Batory, C. Kästner, and G. Saake, *Feature-Oriented Software Product Lines*. Springer, 2013.

[3] S. Apel, S. Kolesnikov, N. Siegmund, C. Kästner, and B. Garvin, "Exploring feature interactions in the wild: The new feature-interaction challenge," in *Proc. FOSD Workshop*. ACM, 2013, pp. 1–8.

[4] N. Siegmund, S. Kolesnikov, C. Kästner, S. Apel, D. Batory, M. Rosenmüller, and G. Saake, "Predicting performance via automated feature-interaction detection," in *Proc. ICSE*. IEEE, 2012, pp. 167–177.

[5] J. Guo, K. Czarnecki, S. Apel, N. Siegmund, and A. Wasowski, "Variability-aware performance prediction: A statistical learning approach," in *Proc. ASE*. IEEE, 2013, pp. 301–311.

[6] Y. Zhang, J. Guo, E. Blais, and K. Czarnecki, "Performance prediction of configurable software systems by fourier learning (t)," in *Proc. ASE*. IEEE, 2015, pp. 365–373.

[7] N. Siegmund, A. von Rhein, and S. Apel, "Family-based performance measurement," in *Proc. GPCE*. ACM, 2013, pp. 95–104.

[8] A. von Rhein, A. Grebhahn, S. Apel, N. Siegmund, D. Beyer, and T. Berger, "Presence-condition simplification in highly configurable systems," in *Proc. ICSE*, vol. 1. IEEE, 2015, pp. 178–188.

[9] N. Siegmund, A. Grebhahn, S. Apel, and C. Kästner, "Performance-influence models for highly configurable systems," in *Proc. ESEC/FSE*. ACM, 2015, pp. 284–294.

[10] F. Shull, J. Singer, and D. Sjøberg, *Guide to Advanced Empirical Software Engineering*. Secaucus, NJ, USA: Springer, 2007.

[11] R. Yin, *Case Study Research–Design and Methods*. Thousand Oaks, CA: Sage, 2003.

[12] B. Flyvbjerg, "Five misunderstandings about case-study research," *Qualitative Inquiry*, vol. 12, no. 2, pp. 219–245, 2006.

[13] G. Ferreira, C. Kästner, J. Pfeffer, and S. Apel, "Characterizing complexity of highly-configurable systems with variational call graphs: Analyzing configuration options interactions complexity in function calls," in *Proc. HotSoS*. ACM, 2015, p. 17:1–2.

[14] P. Jaccard, "The distribution of the flora in the alpine zone." *New phytologist*, vol. 11, no. 2, pp. 37–50, 1912.

[15] C. Borgelt, "Frequent item set mining," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 2, no. 6, pp. 437–456, 2012.

[16] O. Maqbool and H. Babri, "Hierarchical clustering for software architecture recovery," *IEEE Transactions on Software Engineering*, vol. 33, no. 11, pp. 759–780, 2007.

[17] Y. Qiao, J. He, Y. Yang, and L. Ji, "Analyzing malware by abstracting the frequent itemsets in api call sequences," in *Proc. TrustCom*. IEEE, 2013, pp. 265–270.

[18] E. Reisner, C. Song, K. Ma, J. S. Foster, and A. Porter, "Using symbolic evaluation to understand behavior in configurable software systems," in *Proc. ICSE*. ACM, 2010, pp. 445–454.

[19] T. Nguyen, U. Koc, J. Cheng, J. S. Foster, and A. Porter, "igen: Dynamic interaction inference for configurable software," in *Proc. FSE*. ACM, 2016, pp. 655–665.

[20] R. Tartler, D. Lohmann, C. Dietrich, C. Egger, and J. Sincero, "Configuration coverage in the analysis of large-scale system software," *SIGOPS Oper. Syst. Rev. (ACM OSR)*, vol. 45, no. 3, pp. 10–14, 2012.

[21] C. Kim, D. Batory, and S. Khurshid, "Reducing combinatorics in testing product lines," in *Proc. AOSD*. ACM, 2011, pp. 57–68.

[22] B. J. Garvin and M. Cohen, "Feature interaction faults revisited: An exploratory study," in *Proc. ISSRE*. IEEE, 2011, pp. 90–99.

[23] M. Lillack, C. Kästner, and E. Bodden, "Tracking load-time configuration options," in *Proc. ASE*. ACM, 2014, pp. 445–456.

[24] A. Sarkar, J. Guo, N. Siegmund, S. Apel, and K. Czarnecki, "Cost-efficient sampling for performance prediction of configurable systems," in *Proc. ASE*. IEEE, 2015, pp. 342–352.

[25] E. Thereska, B. Doebel, A. Zheng, and P. Nobel, "Practical performance models for complex, popular applications," *SIGMETRICS Perform. Eval. Rev.*, vol. 38, no. 1, pp. 1–12, 2010.

[26] D. Westermann, J. Happe, R. Krebs, and R. Farahbod, "Automated inference of goal-oriented performance prediction functions," in *Proc. ASE*. ACM, 2012, pp. 190–199.