

Prioritizing Configuration Relevance via Compiler-Based Refined Feature Ranking

Federico Bruzzone 

Università degli Studi di Milano, Milan, Italy

Walter Cazzola 

Università degli Studi di Milano, Milan, Italy

Luca Favini 

Università degli Studi di Milano, Milan, Italy

Abstract

Modern programming languages, most notably Rust, offer advanced linguistic constructs for building highly configurable software systems as aggregation of features—identified by a configuration. However, they pose substantial challenges for program analysis, optimization, and testing, as the combinatorial explosion of configurations often makes exhaustive exploration infeasible. In this manuscript, we present the first compiler-based method for prioritizing configurations. Our approach consists of four main steps: 1. extracting a tailored intermediate representation from the Rust compiler, 2. constructing two complementary graph-based data structures, 3. using centrality measures to rank features, and 4. refining the ranking by considering the extent of code they impact. A fixed number of most relevant configurations are generated based on the achieved feature ranking. The validity of the generated configurations is guaranteed by using a SAT solver that takes a representation of this graph in conjunctive normal form. We formalized this approach and implemented it in a prototype, RUSTYEX, by instrumenting the Rust compiler. An empirical evaluation on higher-ranked open source Rust projects shows that RUSTYEX efficiently generates user-specified sets of configurations within bounded resources, while ensuring soundness by construction. The results demonstrate that centrality-guided configuration prioritization enables effective and practical exploration of large configuration spaces, paving the way for future research in configuration-aware analysis and optimization.

2012 ACM Subject Classification Software and its engineering → Compilers; Software and its engineering → Software design engineering; Software and its engineering → Correctness

Keywords and phrases Highly Configurable Systems, Testing Variable Systems, Configuration Prioritization via Static Analysis, Rust.

Digital Object Identifier 10.4230/LIPIcs.CVIT.2016.23

Supplementary Material A replication package is available on Zenodo.

Software: <https://doi.org/10.5281/zenodo.17691776>

Funding This work was partially supported by the MUR project “T-LADIES” (PRIN 2020TL3X8X).

1 Introduction

Premise. Highly configurable software systems¹ often aim to satisfy a wide range of requirements. These systems provide a set of features that can be combined in different ways to create a variety of products [4]. Variability-rich software system development leverages principles from product line engineering, commonly referred to as *feature-oriented programming* [78]. Programming languages, such as C/C++ and Java, provide mechanisms to manage software variability via preprocessor directives, conditional compilation, and annotations.

¹These systems are also known as *product families* or *software product lines* (SPLs) [21]

Rust [59] has also been designed to support the development of highly configurable software systems via its *attribute system*²—specifically, the `cfg` attribute.

Problem Statement. Software variability can be complex, making it challenging to reason about all possible configurations [1]. This problem is further exacerbated by the fact that the number of possible configurations grows exponentially with the number of features, as demonstrated by Krueger [45]. For instance, the Linux kernel is a well-known highly configurable software system with a significant number of features [85] that exploits the native support for variability in C/C++ and, more recently, Rust. Over the years, the kernel and similar systems have been extensively studied, with numerous works highlighting the challenges posed by large configuration spaces [62, 91, 27, 58].

State of the Art Exhaustive approaches become infeasible due to the combinatorial explosion of configurations [5]. Therefore, a strategy is needed to reduce the number of configurations under analysis. Several approaches aim to achieve this goal. Sampling techniques [73, 2, 47], combinatorial testing [71, 54], and in particular t-wise testing [22, 75, 71] attempt to reduce the configuration space while ensuring coverage. Other works have proposed prioritization methods [29, 81, 72, 27], using criteria such as similarity [2], non-functional properties [80], or centrality-based measures [63, 74, 49, 6]. As reported by Classen *et al.* [20], many configurable systems are safety-critical, which makes prioritization even more important: the goal is to detect faults as early as possible while reducing the number of tests.

Limitations of Existing Approaches. Despite these efforts, the problem of finding a proper prioritization criterion remains open. Existing approaches 1. apply only to feature models [9, 77], without considering dependencies in the code, and 2. do not account for the extent of the code affected by each feature. Consequently, configurations that are critical to the system’s behavior may be omitted. Moreover, the need for principled configuration prioritization extends well beyond testing. Highly configurable systems must also address:

1. *compiler optimizations* and *performance analysis*, where either only a subset of variants can be feasibly evaluated, or configuration relevance can guide optimization strategies [94, 84, 90, 95],
2. *program comprehension* and *debugging*, where developers must reason about representative or critical configurations [37, 96],
3. *variability management* and *refactoring*, where structural changes should be guided by the actual impact of features in the code [53, 52], and
4. *regression analysis*, where prioritization helps identify which configurations are most likely to reveal behavioral differences after evolution [79, 88].

In all these contexts, treating features uniformly or relying on stochastic heuristics is insufficient.

Proposal. In this paper, we propose—to the best of our knowledge—the first general method for prioritizing the relevance of configurations via a compiler-based refined ranking of features. To this end, we present RUSTYEx, a fully automated tool designed to identify the *most relevant configurations* in highly configurable Rust software systems.³ We define

²Rust’s attributes are a form of syntactic metadata that can be attached to various parts of the code. Two kind of attributes are supported: *outer* and *inner*. The former is used to annotate *item* declarations (such as functions, structs, and enums), *expression* and *statement*. The latter is used to annotate *modules* and *block expressions* (in particular cases). We will refer to the *items*, *expressions*, and *statements* as *terms*. See <https://doc.rust-lang.org/reference> for more details.

³Although we use Rust in this work, our approach to configuration prioritization is language-agnostic and applicable to any language with native variability support, such as C/C++ or Java.

these configurations as those containing the most relevant features. A feature's relevance is measured by its centrality in the *feature dependency graph* and the extent of the code it impacts. As shown in Figure 1, we instrument the Rust compiler by performing an interprocedural static analysis [68, 69]⁴ to extract the *unified intermediate representation* (UIR). The UIR is obtained by removing irrelevant information from the Rust *abstract syntax tree* (AST) and by creating the *atoms*.⁵ RUSTYEX performs static analysis on the UIR to extract two main data structures: the *dependency graph* [60] and the *polytree* [25]. The former is a weighted directed graph, dubbed *feature dependency graph* (see ② and ② in Figure 2), that represents the dependencies between the features. The latter is an induced subgraph of the UIR, dubbed *atom dependency tree* (see ③ and ③ in Figure 2). This structure captures the *lexical scope-based* [23] dependencies between atoms, enriched with the extent of the code they affect. To address the previously mentioned limitations, we propose a method that combines both structures to identify the most relevant configurations. This method leverages centrality measures as structural metrics [30, 61, 65], and includes graph transformations into propositional and conjunctive normal form (CNF) [20] formulas, along with techniques for refining the ranking of features. RUSTYEX relies on a SAT solver to determine configurations that satisfy the CNF formula based on the refined feature ranking. This ensures that not only is the number of configurations reduced, but also that the most relevant configurations are prioritized. We validate our approach on higher-ranked open-source Rust projects by running RUSTYEX with a fixed number of configurations to generate. To provide a comprehensive evaluation, we report detailed metrics on the proposed structures for each project. Soundness is ensured by construction, guaranteeing that all generated configurations are valid.

Contributions. The main contributions of this paper are:

- the first general method for prioritizing configurations in highly configurable software systems via a compiler-based refined ranking of features,
- a formalization of the approach, from the extraction of the UIR to the algorithms for building the complementary data structures and prioritizing configurations,
- a detailed implementation of our approach in RUSTYEX, a fully automated tool for Rust software,
- an extensive evaluation of RUSTYEX on higher-ranked open-source Rust projects, demonstrating its effectiveness in identifying and prioritizing the most relevant configurations, and
- a formal proof of the soundness of our approach, ensuring that all generated configurations are valid.

Manuscript Structure. The remainder of the paper is organized as follows. Section 2 introduces the necessary background. Section 3 details the design of RUSTYEX, and Section 4 presents our evaluation. Section 6 discusses related work, and Section 7 concludes the paper.

2 Background

We provide background on Rust and its ownership system and on centrality measures.

⁴Inter-procedural analysis is a static analysis technique that analyzes the control and data flow of a program across multiple functions.

⁵From now on, we refer to an *atom* as a single *term* annotated with a given *configuration predicate*.

2.1 The Rust Programming Language.

Rust is a system programming language that focuses on safety, speed, and concurrency. It ensures memory safety without garbage collection, meaning pure Rust programs are free from null pointer dereferences and data races. Rust’s ownership system—integrated into its type system—is inspired by *linear logic* [35, 36] and *linear types* [98, 66], enforcing that each piece of memory (a value) has a single *owner* (the variable binding) at any time [19, 17]. When the owner goes out of scope, the memory is automatically deallocated, enabling user-defined destructors and supporting the *resource acquisition is initialization* (RAII) pattern [89]. Ownership can be transferred (*moved*) or temporarily shared (*borrowed*) through references. Rust supports two types of borrows: multiple *immutable* borrows or a single *mutable* borrow, but never both at the same time. These constraints, enforced by the compiler through *borrow checking*, guarantee memory safety and prevent dangling pointers by ensuring that reference lifetimes never outlive their owners. To support low-level operations, Rust provides **unsafe** blocks, where the compiler’s safety guarantees are suspended and the burden of avoiding undefined behavior falls on the programmer. Outside of these blocks, Rust enforces strict safety, making undefined behavior impossible in safe code. Rust supports software variability through its attribute system, particularly the `cfg` attribute, which allows conditional compilation based on specified configuration predicates. A feature in Rust is defined by all *terms* (see footnote 1) annotated with the same *config name* in a `cfg outer attribute`. The inclusion of a feature in a product depends on the evaluation of a *configuration predicate* at compile time, e.g., `#![cfg(any(unix, windows))]`. Cargo—Rust’s package manager—lets developers declare features using *config names* in the `Cargo.toml` file, combining them with logical operators into *configuration predicates*. Feature dependencies, i.e., *cross-tree constraints* of *feature models* [41, 83, 42], can also be specified in the same file.

2.2 Centrality Measures

Since the early days of social network analysis, *centrality measures* have been used in sociology and psychology to identify the key nodes in a network [82, 8, 43]. They have since become central to graph theory and network analysis [16, 24, 46]. A *network* is typically modeled as a graph with n nodes, indexed by $i \in \{1, 2, \dots, n\}$, and represented as an adjacency matrix $A \in \mathbb{R}^{n \times n}$, where $A_{ij} \neq 0$ indicates an edge between nodes i and j , and $A_{ij} = 0$ otherwise. As Bloch *et al.* [13] observe, not all centrality measures consider edge directions and weights, though these can be incorporated. Formally, a centrality measure is a function $\phi : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^n$, where $\phi_i(A)$ quantifies the relevance of node i in the network A . As a cardinal invariant, it enables ordinal ranking of nodes based on their scores ($\phi_i(A)$). Centrality measures are broadly categorized as: *geometry-based* and *spectral-based* [14]. In the remainder of this section, we introduce some of the most commonly used centrality measures.

Geometric Centrality Measures. The *geometric* centrality measures assign relevance by using a function of distances—i.e., the number of nodes at each distance from a given node. *Degree centrality* counts the number of edges connected to a node: $d^-(i)$ for in-degree and $d^+(i)$ for out-degree. *Closeness centrality*, introduced by Bavelas [8], defines a node’s relevance as inversely proportional to the sum of its distances node to all other nodes. Since the original formulation fails on disconnected graphs (due to infinite distances), a widely adopted variant is:

$$\text{Clo}_i(A) = \frac{1}{\sum_{\substack{j=1 \\ d(i,j) < \infty}}^n d(i,j)}$$

where $d(i, j)$ denotes the shortest path between nodes i and j . *Harmonic centrality* [57] is defined as the sum of the reciprocals of the shortest path lengths between a node and all others:

$$\text{Har}_i(A) = \sum_{j=1}^n \frac{1}{d(i, j)}$$

Unlike closeness centrality, it is well-defined for disconnected graphs; for isolated nodes, the sum evaluates to 0 since $d(i, j) = \infty$ for all $j \neq i$. *Betweenness centrality* [33] measures how often a node appears on the shortest paths between pairs of nodes:

$$\text{Bet}_i(A) = \sum_{s \neq i \neq t} \frac{\sigma_{st}(i)}{\sigma_{st}}$$

where σ_{st} is the number of shortest paths from s to t , and $\sigma_{st}(i)$ is the number of those paths passing through node i . Newman [64] extended *closeness centrality* to weighted networks using Dijkstra's algorithm to compute for shortest paths. Opsahl *et al.* [70] further generalized shortest-path measures by combining edge weights and counts, making them applicable to both *closeness* and *betweenness* centrality.

Spectral Centrality Measures. The *spectral* centrality measures evaluate node relevance using the dominant left eigenvector of the graph's adjacency matrix. *Eigenvector centrality* [15] assigns higher scores to nodes connected to other high-scoring nodes. It is defined as the eigenvector \mathbf{v} associated with the largest eigenvalue λ of the adjacency matrix A :

$$\lambda \mathbf{v} = A \mathbf{v}$$

If A is a *stochastic matrix*, the dominant eigenvalue is 1. However, the *eigenvector centrality* performs poorly on disconnected graphs [11]. *Katz centrality* [43] addresses this by considering the number of paths of varying lengths that connect a node to all other nodes. It is defined as:

$$\mathbf{k} = \mathbf{1}(I - \beta A)^{-1}$$

where $\mathbf{1}$ is a vector of ones, I is the identity matrix, and β is a damping factor satisfying $\beta < 1/\lambda$, with λ the dominant eigenvalue of A .

3 A Deep Dive into RustyEx

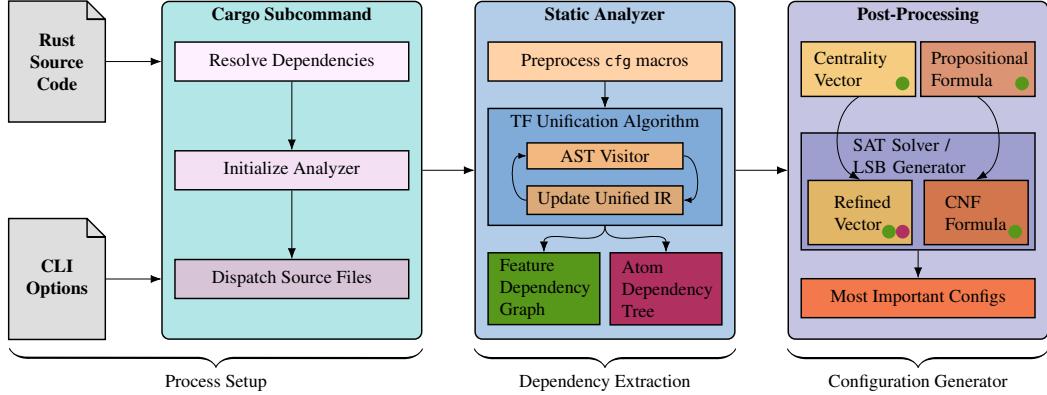
As introduced in Section 1, RUSTYEX is a fully automated tool that instruments Rust compiler to: 1. extract feature dependencies in Rust software, 2. assign feature weights based on their impact on the code, 3. apply centrality measures to rank features, and 4. prioritize configurations.

3.1 Process Overview

RUSTYEX performs its analysis in three main phases: 1. [process setup](#), 2. [dependency extraction](#), and 3. [configuration generation](#), as shown in Figure 1.

Process Setup. Rust projects make extensive use of `cfg` attributes for conditional compilation, both for *feature gating*—enabling or disabling features via `Cargo`—and for purposes like platform-specific code selection [3, 18, 50]. To ensure usability, RUSTYEX integrates seamlessly with the Rust toolchain. In this phase, it accepts a Rust project along with CLI

Figure 1 The phases of RUSTYEX (inspired from [51])



options that customize the analysis, such as the number (N) of configurations to generate and the chose centrality measure. The setup uses `Cargo` to resolve internal crate dependencies which are essential for the analysis.⁶ It then initializes the analysis in memory based on the provided options and configures the dispatch of source files to use the instrumented `rustc`.

Dependency Extraction. RUSTYEX instruments the Rust compiler to extract the *feature dependency graph* (see ②, and ② in Figure 2) and the *atom dependency tree* (see ③, and ③ in Figure 2) from the UIR derived from the AST. The details of these structures and their extraction are discussed in Section 3.2. This is integrated into the compiler by implementing the `rustc_driver_impl::Callback` trait, which provides four hooks: `config`, `after_crate_root_parsing`, `after_expansion`, and `after_analysis`.⁷ `rustc` performs the *eager* macro expansion⁸ after AST creation and before name resolution. Although, `after_crate_root_parsing` would be ideal for analysis before macro expansion, it lacks access to sub-modules that have not been resolved yet. Thus, we hook into `after_expansion`, which is invoked after *eager* macro expansion, name resolution, and AST validation. To avoid premature evaluation of `cfg` attributes, we register a custom file loader in the `config` callback. This loader renames `cfg` macros to delay their evaluation. In the `after_expansion` callback, we perform inter-procedural analysis on the AST to extract the UIR weighting its nodes with the *weighted* fixed-point algorithm shown in Algorithm 1 described in Section 3.2. These weights estimate each feature’s impact on the code.

Configuration Generation. Once the weighted feature dependency graph and the atom dependency tree are extracted, RUSTYEX uses them to identify the most relevant configurations. The weighted feature dependency graph serves two key tasks: 1. ranking features via *geometry-based* and/or *spectral-based* centrality measures (top-left box in the post-processing phase of Figure 1), and 2. generating a corresponding propositional formula (top-right box in the post-processing phase of Figure 1). These tasks are detailed in Section 3.3. The propositional formula is then converted to CNF. Next, the atom dependency tree refines the feature ranking (middle boxes in the post-processing phase of Figure 1). Finally, a SAT solver selects the top N configurations that satisfy the CNF formula, prioritizing those

⁶For performance, RUSTYEX analyzes only the target crate’s source files and compiles external dependencies with the standard Rust compiler. This reduces computational overhead but may slightly impact accuracy [97].

⁷See <https://doc.rust-lang.org/nightly/nightly-rustc/> for details.

⁸Eager expansion expands macro arguments as early as possible, regardless of the macro invocation.

with the highest refined ranking—i.e., the most relevant configurations (bottom box in the post-processing phase of Figure 1).

3.2 From AST to UIR

Overview. In the AST, the `cfg` attribute nodes and their associated *terms* appear as separate child nodes under a common ancestor. Building the UIR involves two key steps: *unification* and *weighting*. The term UIR stems from the *unification* step, which merges each `cfg` attribute with its associated *term* into a single node called an *atom*. UIR nodes are algebraic data types [48, 76, 93, 44, 10], specifically Σ -types—i.e., nodes can be either *atoms* or *relevant* plain AST nodes. Plain AST nodes are included in the UIR only if the *weighting* step assigns them non-zero relevance—e.g., generic bounds on parameters are usually excluded. For instance, in ① and ② of Figure 2, the function `foo` and its `#cfg(feature = "a")` attribute on line 1 are unified into a single *atom* node, centered on the annotated item `foo`. Conversely, the body nodes of the function `bar` remain plain AST nodes in the UIR because they contribute to its weight, helping quantify how much of the code is influenced by the involved `cfg` features.

Formalization. We define the AST as a pair (N_{ast}, E_{ast}) , where N_{ast} and E_{ast} are the sets of nodes and edges, respectively. Let $N_{rel} \subseteq N_{ast}$ denote the set of *relevant* nodes. A node is considered *relevant* if it contributes to the semantics of a feature-dependent element. For instance, `let` statements are relevant, while `extern` `crate` or `use` declarations are not.

The set of *atoms* is defined as:

$$A = \{(p, t) \mid ann(t) = p \wedge t \in T \wedge p \in P\}$$

where:

1. P is the set of all `cfg` predicates, each defined as a recursive Σ -type with the following structure:

$$p = \begin{cases} \text{single}(f) & \text{for } f \in F \\ \text{not}(f) & \text{for } f \in F \\ \text{any}(p_1, p_2) & \text{for } p_1, p_2 \in P \\ \text{all}(p_1, p_2) & \text{for } p_1, p_2 \in P \end{cases}$$

2. $T \subseteq N_{ast}$ is the set of all *terms*, and
3. $ann : T \rightarrow P$ is the surjective function that maps terms to `cfg` predicates.

The UIR is an *enhanced* induced subgraph of the AST, where nodes are enriched with `cfg` predicates and edge directions are reversed. Formally, the UIR is defined as $\mathcal{U} = (N, E, w_N)$, where:

- $N = A \cup N_{rel}$ with $A \cap N_{rel} = \emptyset$, contains both *atoms* and *relevant* AST nodes, forming a Σ -type structure,
- $E = \{(i, j) \mid (i, j) \in E'\}$, where $E' \subseteq E_{ast}$ is the set of reversed edges derived from the AST, and
- $w_N : N \rightarrow \mathbb{N}^+$ assigns weights to nodes based on the type of *term*.

Since UIR nodes may be compound Σ -types (e.g., atoms), $E' \subseteq E_{ast}$ alone does not capture all necessary structure. To preserve connectivity, if $(i, j) \in E_{ast}$ and there exists $a \in A \subseteq N$ s.t. $i = term(a)$ or $j = term(a)$, we extend E' as:

$$E' \leftarrow E' \cup \{ ((ann(i), i), j) \vee (i, (ann(j), j)) \},$$

Here $term : A \rightarrow T$ extracts the term from an $atom$. This ensures that when a node is unified into an $atom$, its original AST edges are preserved in the UIR, as if the node remained *relevant* on its own.

Unification Algorithm. The *unification algorithm* performs a *depth-first visit* traversal of the AST to identify `cfg` attributes and their associated *terms*, merging them into *atoms*. RUSTYEX implements the visitor pattern [34] via the `rustc_ast::visit::Visitor` trait. To track parent-child relationships during traversal, it maintains a *term stack*, where each visited term is pushed. Since `cfg` attributes always appear as leftmost children of terms, encountering one triggers the *unification* process. This involves:

1. extracting the configuration predicate from the `cfg` attribute,
2. parsing the predicate into a custom internal representation, dubbed `ComplexFeature`, and
3. popping the corresponding term from stack to *unify* it with the predicate into an *atom* node.

For example, the `cfg` attribute on line 3 of ① of Figure 2 is parsed as

```
ComplexFeature::Any(vec![
    ComplexFeature::Simple(
        Feature {
            name: "b".to_string(),
            not: false
        }
    ),
    ComplexFeature::Simple(
        Feature {
            name: "c".to_string(),
            not: false
        }
    )
])
```

and unified with the term `bar` to form the corresponding *atom* node. Once a term is fully visited, it is popped from the stack and an UIR edge is added from the child to its parent. Unlike the AST, UIR edges are reversed to reflect lexical scope-based relationships, capturing dependency flow—key for applying graph centrality measures in the feature dependency graph (Section 3.3). The `bar` node in the atom dependency tree (② in Figure 2) shows its unified features, highlighted by the ● and ● markers, emphasizing both unification and edge direction.

Weighting Algorithm. Each node in the UIR is assigned a positive integers, i.e., elements of \mathbb{N}^+ . During the *unification* phase, RUSTYEX assigns to every UIR node a *weight kind*, based solely on the kind of *term* it represents. These *weight kinds* determine how weights are computed in Algorithm 1 (lines 16–25), and include:

1. *no weight*: nodes irrelevant to the analysis that receive no weight (e.g., `extern crate` declarations),
2. *intrinsic weight*: nodes assigned a fixed weight of one (e.g., `let` statements),
3. *children weight*: nodes whose weight is the sum of their children weights (e.g., `foo` in ① of Figure 2), and
4. *reference weight*: nodes that refer to other nodes in the UIR and inherit their weight (e.g., the `bar()` call in `qux` in ① of Figure 2).

After *unification*, RUSTYEX runs a fixed-point algorithm (Algorithm 1) to compute the final weight function w_N over UIR nodes. In Rust, multiple *terms* can share the same identifier only if their `cfg` attributes are mutually exclusive. Let o_n denote the number of distinct *atoms* a node $n \in N$ appears in. The algorithm initializes:

Algorithm 1 Calculate The Weight of the UIR Nodes

```

1:  $m_w \leftarrow \emptyset$ 
2:  $Q \leftarrow \emptyset$ 
3:  $r \leftarrow \text{root}(\mathcal{U})$ 
4:  $\mathcal{U}' \leftarrow \mathcal{U}^T$  ▷ transpose
5:  $\text{CALC\_WEIGHT}(\mathcal{U}', r, Q, m_w)$ 
6:  $\text{RESOLVE\_QUEUE}(\mathcal{U}', Q, m_w)$ 
7:  $\mathcal{U} \leftarrow \mathcal{U}'^T$  ▷ transpose back

8: function  $\text{CALC\_WEIGHT}(\mathcal{U}', n, Q, m_w)$ 
9:   for all  $adj \in \text{ADJ}(\mathcal{U}', n)$  do ▷ adjacency nodes
10:     $\text{CALC\_WEIGHT}(\mathcal{U}', adj, Q, m_w)$ 
11:   end for
12:    $ch\_w \leftarrow 0$  ▷ children weight
13:   for all  $adj \in \text{ADJ}(\mathcal{U}', n)$  do
14:     if  $adj.\text{status} = \text{WAIT}$  then ▷ check the status
15:        $n.\text{status} \leftarrow \text{WAIT}; \text{return}$ 
16:     end if
17:      $ch\_w \leftarrow ch\_w + adj.\text{weight}$ 
18:   end for
19:   match  $n.\text{weight\_kind}$  with
20:     case NO:  $n.\text{weight} \leftarrow 0$ 
21:     case INTRINSIC:  $n.\text{weight} \leftarrow 1 + ch\_w$ 
22:     case CHILDREN:  $n.\text{weight} \leftarrow ch\_weight$ 
23:     case REFERENCE( $called$ ):
24:       if  $m_w[called] = \emptyset$  then ▷ no defs found
25:          $n.\text{status} \leftarrow \text{WAIT}$ 
26:          $Q \leftarrow Q \cup \{n\}$ 
27:         return
28:       else
29:          $t.\text{weight} \leftarrow \text{AVG}(m_w[called])$ 
30:       end if
31:     end match
32:      $n.\text{status} \leftarrow \text{WEIGHTED}$ 
33:      $m_w[n] \leftarrow m_w[n] \cup n.\text{weight}$ 
34:   end function

35: function  $\text{RESOLVE\_QUEUE}(\mathcal{U}', Q, m_w)$ 
36:    $S \leftarrow \emptyset$  ▷ set of seen nodes
37:   while  $Q \neq \emptyset$  do
38:      $c \leftarrow Q \downarrow$ 
39:      $Q \leftarrow Q - \{c\}$  ▷ dequeue
40:     if  $S \cap \{c, \text{LEN}(Q)\} \neq \emptyset$  then ▷ recovery mechanism
41:        $c.\text{weight} \leftarrow \text{DEF\_W}$ 
42:       continue ▷ assign a default weight
43:     end if
44:      $S \leftarrow S \cup \{c, \text{LEN}(Q)\}$ 
45:      $\text{CALC\_WEIGHT}(\mathcal{U}', c, Q, m_w)$ 
46:     if  $c.\text{status} = \text{WAIT}$  then
47:        $Q \leftarrow Q \cup \{c\}$ 
48:     end if
49:   end while
50: end function

```

- $m_w : N \rightarrow \mathbb{N}^{+o_n}$, a map storing all weights associated to the node identifier,
- an empty queue Q ,
- the root node $r = \text{root}(\mathcal{U})$, and
- the transposed UIR $\mathcal{U}' = \mathcal{U}^T$, since the original UIR is built bottom-up and transposing it reestablishes parent-child relationships.

The *weighting* algorithm starts at the root node r , invoking the recursive `CALC_WEIGHT` function (line 17), which computes the weight of each node $n \in N$ by traversing its children in \mathcal{U}' and applying the rules associated with its *weight kind*. Computed weights are stored both in the node and in the map m_w for reuse. For each node $n \in N$ of kind *reference*, such as function calls, the algorithm checks whether the referenced node already has a computed weight in m_w . If the weight is not yet available, the node is marked as *wait* and added to the queue Q . The fixed-point `RESOLVE_QUEUE` function (line 30) then processes Q , computing weights for the queued nodes until the queue is empty or a cycle is detected—e.g., due to direct or mutual recursion [67, 55]. A cycle is detected when Q remains unchanged between two complete iterations. To prevent infinite loops, the algorithm assigns a default fallback weight to nodes involved in cycles and continues processing the remaining queue.⁹ Finally, the UIR is transposed back to its original direction (undoing the earlier transformation), and the finalized weights are stored in w_N .

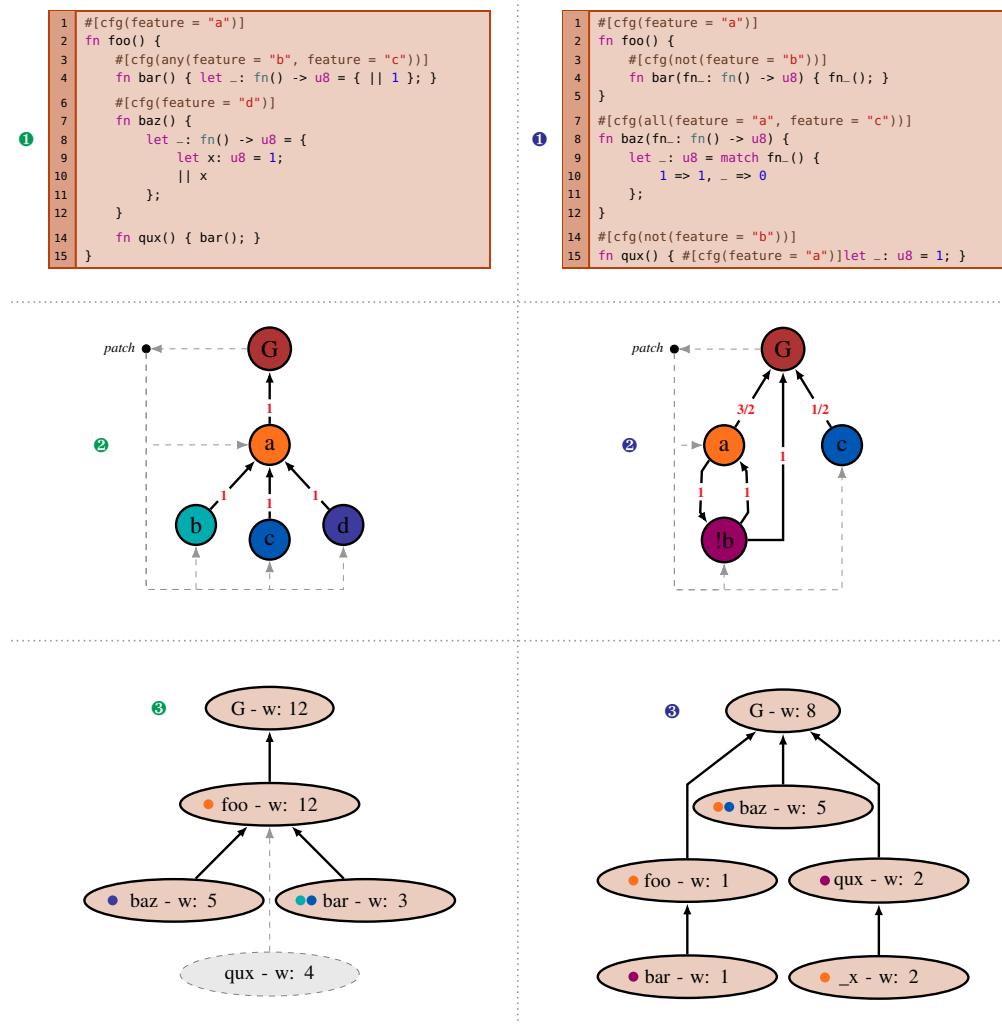
3.3 From the UIR to the Feature Dependency Graph

Overview. Depicted in Figure 2 (❷, ❸), the feature dependency graph is a weighted directed graph that represents dependencies between software features. It is built by applying the *feature-extraction* algorithm to the UIR. This algorithm defines rules for assigning edge weights based on the *configuration predicates* of the `cfg` attributes in which the features appear. At a high level, the graph reflects lexical-scope-based relationships among features. Specifically, if f_i is used within the lexical scope of an *atom* annotated with f_j , then f_i depends on f_j and a directed edge is created from f_i to f_j . For instance, in line 3 of Figure 2 ❷, `feature = "b"` and `feature = "c"` appear in an *any configuration predicate*, within the lexical scope of an *atom* unified with the `feature = "a"` (line 1) and enclosing the `foo` function (line 2). According to the *feature-extraction* algorithm, two edges of weight 1 are created: from `feature = "b"` and from `feature = "c"` to `feature = "a"`. In contrast, at line 7 of ❸, `feature = "c"` is used in an *all configuration predicate* within the *global* lexical scope (denoted as ❸). Consequently, an edge of weight 1/2 is added from `feature = "c"` to ❸. As in the UIR, edge directions in the feature dependency graph are reversed compared to the original AST. The novel insight lies in preserving this natural edge direction: it allows us to interpret a feature's *reputation* (or relevance) based on the *state* (i.e., *configuration predicate*) of the features that depend on it. This enables the application of graph centrality measures on the feature dependency graph, where the edge weights are designed to support more accurate feature ranking.

Formalization. Until now, we referred to the feature dependency graph as a graph. However, during its construction, it is initially a *multigraph*, since a feature f_i may depend on another feature f_j in multiple parts of the code. Formally, this multigraph is denoted as $\mathcal{F} = (F, D, w_R)$, where:

⁹Another approach has been proposed by instrumenting LLVM-IR [97], attempting to solve a system of linear equations that, by definition, could not admit solutions.

Figure 2 RUSTYEX process demonstrated on two scenarios



1. $F = \{(f_i, p_i) \mid f_i \in CN \wedge p_i \subseteq P\}$ is the set of nodes, where CN is the set of all *feature names* and p_i is the set of all *cfg* predicates in which f_i is used;
2. $D = (U, m)$ is the *multiset* of directed edges representing feature dependencies, where $U = \{(i, j) \mid i, j \in F\}$ is the set of edges and $m : U \rightarrow \mathbb{N}^+$ is the *multiplicity* function that counts how many times a dependency occurs; and
3. $w_R : R \rightarrow \mathbb{R}$, where

$$R = \{((i, j), k) \mid (i, j) \in U, 1 \leq k \leq m((i, j))\},$$

is the weight function that assigns a weight to each individual edge instance, indexed by the multiplicity index k , with weights determined by the *configuration predicates* in which the dependency arises.

The *feature-extraction* algorithm then squashes these multiple edges into single ones by aggregating their weights, effectively transforming the multigraph \mathcal{F} into a simple graph $\mathcal{F}' = (F, D', w'_R)$, where:

$$D' = \{(i, j) \mid (i, j) \in U\}$$

is a set of unique edges obtained by removing the multiplicities from D , and the aggregated weight function is defined as:

$$w'_R((i, j)) = \sum_{k=1}^{m((i, j))} w_R((i, j), k), \forall (i, j) \in D'.$$

Here, the sum aggregates the weights of all instances of the edge (i, j) , with $m(i, j)$ denoting the total number of such instances. Each individual edge instance (i, j, k) contributes to the total weight assigned to the edge (i, j) in the final graph \mathcal{F}' .

Feature-Extraction Algorithm. The *feature-extraction* algorithm builds the feature dependency graph \mathcal{F}' from the UIR. It begins by iterating over the UIR atoms $A \subseteq N$. For each atom $a \in A$, it retrieves its parent node $\hat{a} \in \{x \mid (\hat{a}, x) \in E \wedge x \in A\}$. The existence of \hat{a} and uniqueness are guaranteed for all $a \in A$ such that $a \neq \mathbf{G}$, since the UIR is an induced subgraph of the AST, and every node in a tree has a unique parent. This follows from the fact that the UIR is an *enhanced* induced subgraph of the AST. Concretely, \hat{a} is the nearest ancestor of a that is an *atom*. The algorithm creates a directed edge from each feature $f \propto \text{pred}(a) = p_a$ to each feature $\hat{f} \propto \text{pred}(\hat{a}) = p_{\hat{a}}$, where \propto is read as “*involved in*” and $\text{pred} : A \rightarrow P$ returns the configuration predicate that annotates each *atom*. The predicate sets of the *feature nodes* f and \hat{f} are updated as:

$$(f, p) \leftarrow (f, p \wedge \text{pred}(a)) \text{ and } (\hat{f}, p') \leftarrow (\hat{f}, p' \wedge \text{pred}(\hat{a})).$$

The weights of the *multi-edges* are derived from the nested `cfg` predicate of a . For each feature node \hat{f} and each $(f, w) \in \chi(p_a, 1)$, the algorithm creates a weighted edge from f to \hat{f} , using the recursively defined function χ :

$$\chi(p_a, w) = \begin{cases} [(f, w)], & \text{if } p_a \equiv \text{single}(f) \vee \text{not}(f) \\ \chi(p_1, 1) \oplus \chi(p_2, 1), & \text{if } p_a \equiv \text{any}(p_1, p_2) \\ \chi(p_1, d_{p_1}^w) \oplus \chi(p_2, d_{p_2}^w), & \text{if } p_a \equiv \text{all}(p_1, p_2) \end{cases}$$

where $d_p^w = (w \times |\{f \propto p\}|)^{-1}$ and the \oplus infix operator concatenates two lists. For example, in ① of Figure 2, the `feature = "b"` (line 3) appears in an `any` predicate nested under both `feature = "a"` (line 1) and `foo` (line 2), so an edge with weight 1 is added from `feature = "b"` to the `feature = "a"`. As discussed in Section 3.3, this weighting scheme prioritizes features in `any` predicates over those in `all` predicates, reflecting the fact that `any` predicates are less restrictive (only one feature needs to be enabled), whereas `all` predicates require all involved features to be enabled simultaneously. Finally, the graph is simplified by summing the weights of all edges with the same source and target nodes.

Centrality Measures. Once the feature dependency graph is built, RUSTYEX ranks features by relevance using a centrality measure specified via command-line argument (see Figure 1, marked with a ● inside the *Centrality Vector* box). Since \mathcal{F} is a disconnected, weighted, directed graph, some centrality measures are either undefined or behave unpredictably [11]. For example, *eigenvector centrality* requires a connected graph [11] and both *closeness centrality* and *betweenness centrality* need adaptations to handle weighted directed graphs [64, 70].¹⁰ To ensure connectivity, we introduce a synthetic `patch` node (see ② and ② in Figure 2). This node has a single incoming edge from \mathbf{G} and outgoing edges to all other nodes, each with weight 1. This transformation guarantees that \mathcal{F} is fully connected, allowing the following centrality measures to be applied without modification:

¹⁰RUSTYEX uses `rustworkx` [92] to perform centrality measures.

- Newman Closeness Centrality [64]
- Opshal Betweenness Centrality [70]
- Eigenvector Centrality [15]
- Katz Centrality [43]

After computing centrality scores, a vector \vec{v} is produced where each entry v_i represents the centrality of feature f_i . The *patch* nodes are excluded from the vector, as they are not *extracted* features. This approach aligns with the intuition behind our graph construction: a feature's reputation is defined by the features that depend on it. Centrality measures serve as a principled way to rank features by structural relevance.

Propositional Formula and CNF. From the feature dependency graph, RUSTYEX builds a *propositional formula* by iterating over each feature node $f_i \in F$. For each node, it generates a clause $\varphi_i = \ell(\text{pred}(f_i))$, where ℓ maps `cfg` predicates to *logical* expressions. It then adds an implication $f_i \Rightarrow f_j$ for each edge $(f_i, f_j) \in D'$, capturing feature dependencies. The overall formula φ is defined as:

$$\varphi = \bigwedge_{f_i \in F} \left(\varphi_i \wedge \bigwedge_{f_j \in \text{succ}(f_i)} (f_i \Rightarrow f_j) \right)$$

where $\text{succ}(f_i)$ is the set of successors of f_i in \mathcal{F}' . Finally, φ is converted to CNF formula φ_{cnf} [38]. Duplicate clauses may be removed before or after the CNF conversion.

3.4 From the UIR to the Atom Dependency Tree

Overview. Figure 2 (boxes ❸ and ❹) shows the atom dependency tree derived from the UIR. The name highlights its selective retention of UIR Σ -type nodes: only the *atom* variant is preserved, while *plain* AST nodes are removed. Specifically, only UIR nodes with `cfg` attributes are kept. For instance, in Figure 2, box ❸, the gray `qux` node (representing the `qux` function from line 14 in Figure 2, box ❶) is excluded because it lacks any `cfg` attribute. The atom dependency tree captures lexical scope-based dependencies among UIR *atoms*. Unlike the feature dependency graph, which flattens UIR structure, this tree preserves parent-child relationships between *atoms* and, by extension, between *features*. Each node is weighted by the amount of code affected by that atom's `cfg` condition, reflecting the feature's impact. To avoid losing information when *plain* AST nodes are discarded, the *atom dependency extraction* algorithm aggregates their weights into their nearest ancestor *atom* node (see `foo` node in Figure 2, box ❹). Finally, the centrality vector \vec{v} is refined using the structure of the atom dependency tree to get a better ranking of feature relevance.

Formalization. Given the UIR definition $\mathcal{U} = (N, E, w_N)$, we define the atom dependency tree as a directed graph $\mathcal{A} = (A, E_A, w_A)$, where:

- $A = N \setminus N_{\text{rel}}$ is the set of *atom* nodes from the UIR,
- $E_A = E \setminus \{(i, j) \mid (i, j) \in E \wedge (i \in N_{\text{rel}} \vee j \in N_{\text{rel}})\}$ is the set of edges only connecting *atom* nodes, and
- $w_A : A \rightarrow \mathbb{N}^+$ is the weight function that assigns to each atom a weight reflecting the amount of code affected by the `cfg` attributes in which its features appear.

The atom dependency tree (\mathcal{A}) is an induced subgraph of \mathcal{U} , containing all *atoms* and their dependencies, but excludes *plain* AST nodes. This holds trivially since $A \subseteq N$ and $E_A \subseteq E$.

Atom Dependency Extraction Algorithm. The *atom dependency extraction* algorithm builds the atom dependency tree from the UIR. It iterates over the UIR atoms $A \subseteq N$. For each atom $a \in A$, it identifies its parent $\hat{a} \in \{x \mid (\hat{a}, x) \in E \wedge x \in A\}$. As before,

the existence of \hat{a} is trivial to prove. The algorithm then distinguishes two cases: 1. if $\text{pred}(a) \neq \emptyset$, it creates a edge from a to \hat{a} and updates \hat{a} 's weight by adding a 's weight, 2. if $\text{pred}(a) = \emptyset$, it *does not* create an edge from a to \hat{a} but still updates \hat{a} 's weight. For example, in Figure 2, box ③, the `qux` node contributes to the weight of `foo` but does not create an edge. The parent weight update is performed as:

$$w_A(\hat{a}) \leftarrow w_A(\hat{a}) + w_N(a).$$

Refinement Algorithm. The *refinement algorithm* refines the centrality vector \vec{v} using the atom dependency tree \mathcal{A} (see Figure 1). It iterates over each node $a \in A$ in \mathcal{A} , extracts every feature $f \propto \text{pred}(a)$, and updates its centrality value in \vec{v} as:

$$\vec{v}_f \leftarrow \vec{v}_f + \begin{cases} \overline{w_A(a)}, & \text{if } f \propto \text{single}(f) \vee f \propto \text{not}(f) \\ \overline{w_A(a)}, & \text{if } f \propto \text{any}(p_1, p_2) \\ \frac{\overline{w_A(a)}}{|\{f \propto \text{pred}(a)\}|}, & \text{if } f \propto \text{all}(p_1, p_2) \end{cases}$$

Here $f \propto p$, with $p \in P$, means that f appears in the p -variant of a 's `cfg` predicate, and $\overline{w_A(a)} \in [0, 1]$ is the normalized weight of the node a . The normalization prevents the inclusion of the existing centrality values in \vec{v} . Although \mathcal{F} captures dependencies between features, it may fall short of expressing their actual relevance. The atom dependency tree \mathcal{A} refines centrality values by incorporating the *extent* of code affected by each feature, yielding \vec{v}' a new permutation of the vector \vec{v} .

3.5 Configuration Generation

Configuration generation is the final step of RUSTYEx. It produces the most relevant configurations based on the centrality vector \vec{v}' and the CNF formula φ_{cnf} (see Figure 1). While more advanced lexical-scope-based strategies could be considered, RUSTYEx uses SAT solvers—both incremental (e.g., `MiniSat` [26, 87] and `CaDiCaL` [12]) and non-incremental (e.g., `Kissat` [12]). Given a fixed number K of configurations to generate, the algorithm iterates over the values in \vec{v}' . For each feature f_i , it updates the formula as:

$$\varphi_{cnf} \leftarrow \varphi_{cnf} \wedge f_i$$

Then, it uses a SAT solver to generate all the configurations satisfying φ_{cnf} . A SAT solver is then invoked to produce all satisfying configurations for the updated formula. If a satisfying configuration is found, it is added to the result set. This process repeats until K configurations are obtained, negating each newly found configuration to favor the discovery of new ones.

3.6 Applications and Clarifications

The generation of configurations in RUSTYEx is performed incrementally and in a *lazy* fashion. The tool ranks all features according to their refined centrality in \vec{v}' and then generates only the top K configurations most likely to cover critical feature interactions. In practice, this means prioritizing configurations that include the highest-ranked features in \vec{v}' , starting with \vec{v}'_0 , the most central feature.

As discussed earlier, this prioritization strategy supports efficient exploration of the configuration space by focusing on variants that are most representative and impactful in terms of feature relevance. By analyzing these configurations first, developers can uncover

important interactions and behaviors in the most relevant scenarios, improving overall system understanding and validation. Moreover, this targeted approach significantly reduces the computational cost and resource consumption associated with exhaustive analysis of all configurations, making it a practical solution for large-scale, highly configurable systems. For instance, these prioritized configurations could also serve as candidates for targeted performance profiling or static analysis, ensuring that critical feature combinations are examined before less relevant ones.

This configuration-aware pipeline highlights the potential of centrality-guided heuristics in feature-oriented analysis of configurable systems and paves the way for future work in optimizing configuration sampling based on structural and semantic program properties. It suggests a broader paradigm where program structure and semantics guide systematic exploration of the configuration space, potentially benefiting tasks such as performance evaluation or security analysis.

4 Evaluation

The evaluation of our approach is twofold: 1. we prove the soundness of RUSTYEx by showing that the generated configurations are valid, and 2. we evaluate the performance of RUSTYEx on a set of 40 real-world, high-ranking open-source Rust projects.

4.1 Soundness

Premises. To prove correctness, we show that: 1. the CNF encoding faithfully represents all configuration constraints, 2. the SAT solver returns only satisfying assignments, and 3. no invalid configuration is ever generated.

Construction of the CNF Formula. *Clause formation*—for each feature node $f_i \in F$, a clause $\varphi_i = \ell(\text{pred}(f_i))$ is created to enforce the logical interpretation of the feature’s configuration predicate. That is, if a feature is enables in some configuration, then its predicate must be satisfied. *Implication formation*—for every outgoing edge to f_j from f_i (i.e., for each f_j such that $(f_i, f_j) \in D'$), an implication $f_i \Rightarrow f_j$ is added to capture the dependency relations. This ensures that if feature f_i is active, then all its dependent f_j are also active. *Extension*—the formula φ is extended to include mandatory features specified in `Cargo.toml`, guaranteeing their presence in all generated configurations. *Equivalence Preservation*—since the CNF conversion uses standard techniques that preserve logical equivalence, any assignment that satisfying the CNF formula φ_{cnf} also satisfies the original formula φ .

Correctness of the SAT Solution. Given a SAT solver sound and complete,¹¹ any assignment α returned for φ_{cnf} satisfies every clause in φ_{cnf} , and thus also in φ . This implies:

1. for each feature f_i with clause φ_i , its activation under α satisfies its configuration predicate: $\alpha \models \varphi_i \quad \forall f_i \in F$ —that is, α models the predicate;
2. for each dependency $f_i \Rightarrow f_j$, the implication holds under the assignment α if $\alpha(f_i) = \text{true} \Rightarrow \alpha(f_j) = \text{true}$;
3. all mandatory features from `Cargo.toml` are active in α : $\forall f_i \in F_{\text{mandatory}}, \alpha(f_i) = \text{true}$.

¹¹A SAT solver is *sound* if it only returns assignments that satisfy the formula, and *complete* if it finds a solution whenever one exists.

Table 1 Results of the experiments conducted on 40 Rust projects. All columns are aggregated per project, summing the values of all crates within the workspace, except for *Peak Memory Usage* and *UIR Height*, which are the maximum values.

Project Name	GitHub Stars	Crates.io Downloads	Lines of Code	Workspace Members	Failed Members	Dependent Features	UIR Nodes	UIR Edges	UIR Height	Feat. Nodes	D.G. Edges	Feat. D.G. Edges	Squashed D.G. Edge	Atom D.T. Nodes	Atom D.T. Edges	Execution Time	Memory Usage
rustdesk	81965	2826	108908	8	1	67	8	1353	1346	13	24	41	27	29	22	1 s	53 MB
gitoxide	9490	76429	223296	79	2	554	139	180166	180089	53	321	910	473	587	510	817 s	999 MB
deno	101658	419085	284483	36	3	561	8	123883	123850	31	124	1220	153	570	537	320 s	945 MB
tauri	89324	3840655	82089	26	2	255	65	26175	26151	30	78	193	89	144	120	108 s	902 MB
sway	62453	1092	210721	28	4	369	7	119724	119700	33	75	118	78	91	67	1156 s	1910 MB
fuel-core	57814	265047	148682	36	4	370	69	85528	85496	31	132	437	188	351	319	153 s	488 MB
alacritty	57640	192714	32812	5	1	39	6	27244	27240	24	17	47	25	36	32	117 s	513 MB
zed	54095	54231	591481	178	2	2458	91	23791	23615	28	381	268	237	231	55	91 s	497 MB
bat	51048	1496904	14971	1	0	30	9	675	674	19	5	14	7	8	7	26 s	339 MB
ripgrep	50241	938719	50285	10	2	40	8	60181	60173	27	43	158	69	124	116	320 s	887 MB
meilisearch	49152	1288	166851	19	2	196	30	44000	43983	30	46	69	41	55	38	404 s	2477 MB
fuels-rs	43907	5611	42484	24	2	90	22	43578	43556	28	70	165	78	138	116	161 s	691 MB
typst	37358	57580	112435	20	4	185	15	48913	48897	27	49	79	50	57	41	81 s	465 MB
helix	35745	103645	93042	14	1	143	12	55807	55794	28	59	175	89	133	120	595 s	980 MB
ruff	35669	4764	370658	38	3	402	23	124392	124357	51	130	237	161	176	141	462 s	504 MB
lapce	34961	35019	67802	5	2	73	3	20525	20522	41	16	44	25	32	29	195 s	948 MB
nushell	33821	975	291737	36	6	213	26	126096	126066	28	108	303	138	232	202	662 s	956 MB
polars	31805	2400845	376212	26	3	362	789	58368	58345	32	103	448	161	354	331	277 s	529 MB
swc	31649	1767375	648247	116	15	957	230	412313	412212	258	326	1122	369	722	621	2016 s	3425 MB
influxdb	29458	283996	54304	19	2	352	15	62082	62065	27	58	116	67	92	75	164 s	504 MB
tabby	29742	5637	41457	19	1	270	15	52972	52954	21	55	134	56	109	91	442 s	1827 MB
servo	29222	8792	367323	8	1	45	19	2670	2663	10	19	23	17	17	10	1 s	64 MB
wasmer	19335	4820432	263877	35	8	350	153	72860	72833	26	121	318	164	235	208	246 s	435 MB
diem	16702	16759	428331	186	10	2239	130	427311	427135	65	477	1207	475	757	581	1359 s	968 MB
texture-synthesis	1768	61696	4735	3	0	7	2	13211	13208	23	15	30	22	21	18	11 s	295 MB
kajiy	5006	1060	26638	14	3	97	3	28922	28911	29	25	22	17	19	8	68 s	1395 MB
rust-gpu	7412	2135	44422	19	1	66	24	7581	7563	34	46	41	38	31	13	12 s	360 MB
substrate	8381	1938	595987	270	8	2986	554	516169	515907	43	863	2172	1002	1681	1419	1135 s	1801 MB
tantivy	12662	5253776	118896	9	1	67	11	34915	34907	51	32	84	41	66	58	104 s	478 MB
tonic	10477	93238823	41456	29	0	166	42	37584	37555	25	88	330	104	254	225	270 s	545 MB
sendme	357	15473	1031	1	0	20	0	1577	1576	18	2	1	1	1	0	52 s	525 MB
komodo	2766	822	63829	12	3	57	2	12443	12434	33	20	33	13	31	22	80 s	685 MB
quiche	9841	541561	84162	9	2	28	0	23455	23448	67	25	62	30	50	43	28 s	352 MB
rolldown	10119	941	49343	35	1	256	12	35014	34980	32	89	149	76	126	92	64 s	354 MB
iroh	3892	96831	39197	7	1	129	14	15758	15752	20	22	70	28	43	37	100 s	381 MB
spl	1194	1238	104219	25	1	363	44	70776	70752	53	77	183	87	147	123	248 s	549 MB
union	22086	11795	387611	148	2	1562	248	134318	134172	33	518	843	605	605	459	212 s	563 MB
hyperswitch	13381	13567	575328	33	6	450	116	107187	107160	29	123	754	183	578	551	599 s	3800 MB
egui	23706	4248075	101376	39	4	138	44	39574	39539	44	111	235	126	181	146	88 s	489 MB
pueue	5297	72209	18340	3	2	24	0	15463	15462	23	5	18	8	15	14	64 s	506 MB
Total	1212599	120362360	7329058	1628	116	17036	3008	3294554	3293042	258	4898	12873	5618	9129	7617	13328 s	3800 MB
Average	30314	3009059	183226	40	2	425	75	82363	82326	37	122	321	140	228	190	333 s	885 MB

By definition, a configuration is valid if and only if:

- it satisfies all the configuration constraints including feature predicates and dependencies, and
- it includes all mandatory features.

Therefore, since any assignment α produced by the SAT solver satisfies all clauses in φ_{cnf} , every configuration generated by RUSTYEX is valid.

Reinforcement by Contradiction. Assume, for contradiction, that the SAT solver returns a configuration α' that is invalid—that is, it violates one or more constraints. Then there exists at least one clause in φ_{cnf} that is *not* satisfied by α' . By definition, a *satisfying assignment* satisfies all clauses in the formula, and the soundness of the SAT solver guarantees that α' is a satisfying assignment for φ_{cnf} . This contradiction implies that our assumption is false: α' cannot be invalid. Hence, all configurations produced by the SAT solver are valid.

4.2 Performance Evaluation

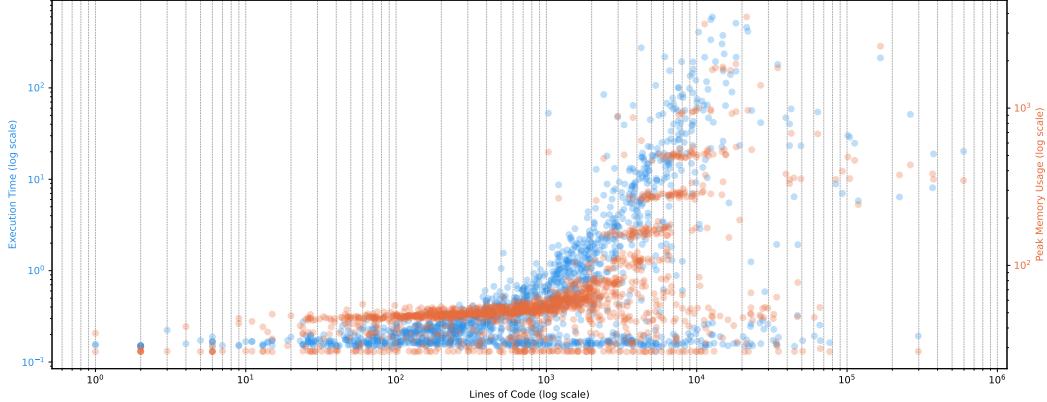
Experimental Setup. All experiments were conducted on a commodity laptop equipped with an Intel Core i5-1135G7@2.40 GHz CPU and 16 GB of RAM. This choice highlights the lightweight nature of our approach: unlike many program analysis and variability-management tools, which often require dedicated servers or high-performance hardware, RUSTYEX can be executed on standard developer workstations. Most of the analyzed projects are organized as `Cargo` workspaces, which naturally decompose the system into multiple crates. For each crate, we executed RUSTYEX with a timeout of 10 minutes, which is a realistic bound for integration in a continuous integration (CI) pipeline. Various metrics were collected and then aggregated at the project level, including counts of features and dependencies, statistics on UIR nodes and edges, and summaries of the feature dependency graph and the atom dependency tree. We also monitored peak memory usage and execution time. Overall, these choices demonstrate that RUSTYEX can be seamlessly adopted in everyday development workflows, where fast turnaround and modest resource requirements are essential. A replication package with all experimental data and scripts is available on Zenodo:

<https://doi.org/10.5281/zenodo.17691776>.

Execution Time and Scalability. Figure 3—highlighted in ●—shows the correlation between lines of code and execution time on a logarithmic scale. RUSTYEX successfully completed the analysis of about 93% of the projects, showing strong scalability across a wide range of sizes and domains. Execution time scaled smoothly with code size: for small to medium projects, the analysis often completed within seconds, while for the largest ones (e.g., `hyperswitch` and `swc`) execution time peaked at 33 minutes. The average execution time per project was 333 seconds, with individual crates requiring about 8 seconds on average. These results confirm that our tool is not only efficient in practice, but also robust enough to handle large modular codebases with hundreds of thousands of lines of code and thousands of features. Importantly, the logarithmic trend visible in the plot suggests that the analysis cost grows sub-linearly compared to project size, which indicates that RUSTYEX is suitable for long-term scalability.

Memory Usage. Figure 3—highlighted in ●—shows the correlation between lines of code and peak memory usage, also on a logarithmic scale. Memory consumption was generally modest: most projects required less than 1 GB of RAM, while only a handful of large codebases such as `hyperswitch` and `swc` reached a peak of 3.8 GB. The average peak usage was 885 MB

Figure 3 Correlation between lines of code (x-axis), execution time (s) ●, and peak memory usage (MB) ●.



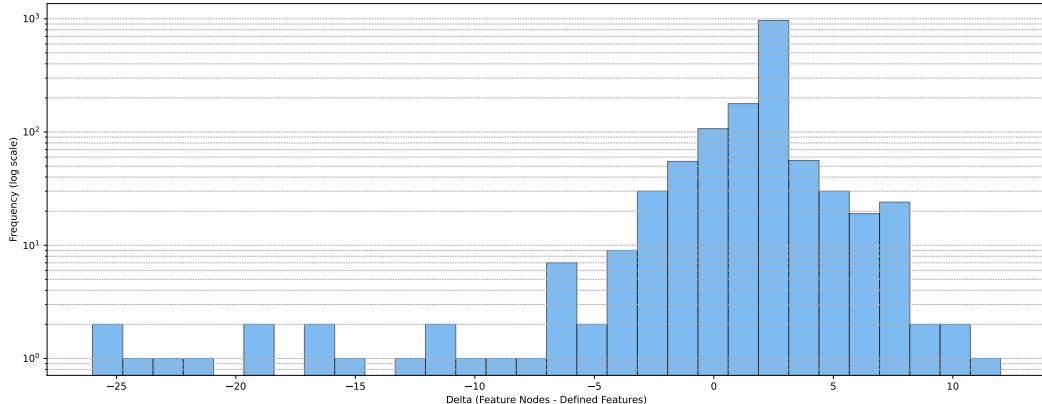
per project, which fits comfortably within the capacity of mainstream laptops. At the crate level, average peak memory dropped to 105 MB, indicating that memory requirements depend more on the internal complexity of each crate than on the overall project size. This stability across a diverse set of projects confirms that RUSTYEX can be executed in environments with limited resources, such as cloud-based CI/CD pipelines or developer laptops.

Intermediate Structures. The intermediate representations produced by RUSTYEX remained compact and efficient to process. On average, the UIR contained approximately 82,000 nodes and a similar number of edges, which is relatively small compared to the size of the original ASTs. The *feature dependency graphs* were significantly more compact, typically comprising between 100 and 500 edges after multi-edge squashing. Even smaller were the atom dependency trees, which were on average 95% smaller than the UIR itself. This compactness is particularly important because it directly impacts the cost of subsequent analyses and ranking procedures: smaller graphs and trees allow for faster computation of centrality measures and logical transformations. These results highlight that our abstraction strategy, centered on atoms and UIR, strikes a good balance between preserving semantic detail and reducing analysis overhead.

Feature Complexity. Most projects defined a substantial number of features, with a median of 425 per project. On average, RUSTYEX detected 122 nodes in the feature dependency graph, often revealing cross-feature dependencies that underscore the challenges of analyzing highly configurable systems. As shown in Figure 4, the number of detected features is generally greater than the number explicitly declared in `Cargo.toml` files. This discrepancy arises mainly from two factors: 1. feature overlaps across different crates within the same workspace, which are not redefined in the manifest, and 2. the handling of the `not` predicate, which introduces logically independent variants and, in the worst case, doubles the number of detected features.

Despite the substantial number of features, their utilization remained modest: on average, only 228 feature-related artifacts were identified, which means that each feature was used fewer than two times on average. This finding suggests that while projects expose a rich configuration space at the feature level, many features have limited practical impact on the actual codebase. In turn, this motivates the need for principled prioritization: if most features are rarely used, developers and testers should concentrate on the ones that have the greatest influence on the system.

■ **Figure 4** Delta between the number of *declared features* and the number of *detected feature*.



Robustness. Finally, we assessed the robustness of RUSTYEx. Out of more than 1,600 crates analyzed (roughly 40 per project), only 116 failed, yielding a 93% success rate. Most failures were attributable to missing system dependencies during compilation, which are unrelated to the analysis itself. Timeouts were rare, affecting fewer than 2% of the crates. Importantly, no major crashes or incorrect results were observed. Even large modular workspaces such as `diem` and `deno` were analyzed successfully without manual intervention. These results confirm that RUSTYEx is mature and reliable enough to be integrated into the toolchains of both researchers and practitioners.

5 Threats to Validity

We organize our discussion following Wohlin *et al.* [99]’s taxonomy.

5.1 Construct Validity

Proxy metrics for feature relevance. We rely on *geometric* and *spectral* centrality computed on the feature dependency graph to approximate feature relevance. These metrics are not neutral: each has biases and limitations, especially regarding disconnected components and local vs. global influence.

Mitigation. RUSTYEx allows users to select the centrality measure that best fits their project domain, thus reducing the risk of systematic misinterpretation. To avoid excluding isolated features, the tool introduces a *patch node* that connects otherwise disconnected components, ensuring that all features are represented in the ranking process.

Weighting of UIR nodes. In cases where definitions are missing or cycles are detected, nodes of type *reference* are assigned a default weight. This approximation may underestimate the role of unresolved external calls, potentially distorting prioritization.

Mitigation. RUSTYEx implements a recovery mechanism that assigns fallback weights, preventing infinite loops and limiting distortion. This ensures that unresolved dependencies do not cancel the influence of otherwise relevant nodes.

5.2 Internal Validity

Assumptions about SAT solver correctness. Our configuration generation relies on the soundness and completeness of the underlying SAT solver. If the solver produced incorrect

results, the validity of generated configurations would be at risk.

Mitigation. To minimize this threat, we use widely adopted and thoroughly tested SAT solvers, reducing the likelihood of errors and strengthening the robustness of our approach.

Timeouts and analysis failures. A fraction of crates ($\sim 7\%$) failed due to timeouts or missing system dependencies. Such failures may introduce selection bias, as not all crates are equally represented in the analysis.

Mitigation. We adopt a fixed 10-minute timeout per crate to ensure fairness across the dataset. Failed crates are skipped, but since the majority of failures were due to missing dependencies rather than intrinsic tool limitations, the overall validity of the evaluation remains preserved.

5.3 External Validity

Focus on open-source Rust projects. Our evaluation is based on 40 open-source projects from GitHub and `crates.io`. While these projects cover a broad spectrum of real-world software, they may not capture the full variability of proprietary or industrial codebases.

Mitigation. Many of the analyzed projects are widely used in production and serve as dependencies for industrial systems. This increases confidence that the results generalize beyond purely academic or hobbyist software. Furthermore, since our method is language-agnostic, future work will investigate applications in other ecosystems such as C/C++.

Fixed configuration budget. We generated a fixed number K of configurations per project. In practice, real-world scenarios may require adaptive strategies that vary the number of generated configurations depending on project size, release stage, or available resources.

Mitigation. RUSTYEX supports parameterized configuration policies, enabling users to adjust K according to their specific needs and constraints. This flexibility ensures that the tool can be adapted to dynamic development and testing scenarios.

5.4 Conclusion Validity

Dependence on feature ranking. The order in which configurations are generated depends on both the selected centrality metric and the refinements applied to the atom dependency tree. Different choices may therefore lead to different prioritized sets.

Mitigation. RUSTYEX makes it possible to switch among centrality measures and disable refinements. This allows users to conduct sensitivity analyses and evaluate how rankings vary under different assumptions.

No ground truth for configurations. There is no universally accepted reference set of “correct” or “important” configurations. As a result, it is challenging to directly evaluate the relevance of the configurations produced by our method.

Mitigation. To address this, RUSTYEX enforces structural consistency checks and ensures compliance with `Cargo.toml` constraints, such as required features and valid predicate logic. Additionally, it verifies that all generated configurations are satisfiable with respect to the derived CNF formula. These safeguards ensure that generated configurations are both valid and representative of critical execution paths.

6 Related Work

Prioritizing configurations in highly configurable systems has been extensively studied, though from different perspectives and often under different assumptions. Several comprehensive

surveys provide an overview of the field, such as [1] (see in particular Sect. 4.4), [29], and [39], which systematically map research trends and highlight open challenges. Building on these, we briefly summarize the most relevant contributions along three main dimensions.

Static Analysis and Preprocessing. Early work on highly configurable systems, especially in the context of the Linux kernel, focused on static analysis and preprocessing techniques. El Sharkawy *et al.* [28] developed methods to process `#ifdef` directives for extracting software metrics, making configuration-specific complexity more manageable. Similarly, Sincero *et al.* [86] investigated preprocessing of C macros to detect dead code, thus helping to reduce variability-induced maintenance effort. These approaches are closely related to our idea of analyzing variability from the source, but they remain tied to metrics extraction or defect detection in specific ecosystems. By contrast, RUSTYEx generalizes the notion of variability-aware static analysis beyond macros or directives, targeting the identification of *relevant configurations* in Rust projects without assuming language-specific preprocessing artifacts.

Feature Model-based Prioritization. Another major line of research builds on explicit feature models, where features and their relationships are captured in structured diagrams. Bagheri *et al.* [6] proposed the *stratified analytic hierarchy process* to prioritize and select features by decomposing the decision process into manageable layers. Peng *et al.* [74] employed a directed, weighted acyclic graph to assess feature importance via *weighted degree centrality*, highlighting asymmetries in feature influence. Mannion *et al.* [56] explored weighting strategies based on variability types, assuming a graph with explicit sources and sinks. Beyond prioritization, Bagheri *et al.* [7] also assessed the maintainability of feature models using structural metrics such as cyclomatic complexity [61], network centrality [65], and classical software engineering metrics [30]. While these approaches are rigorous and effective, they presuppose the existence of a formal feature model—a strong assumption in practice, since many modern systems (including Rust projects) rely on decentralized, implicit forms of variability encoded directly in build manifests and conditional compilation. RUSTYEx differs in this respect, as it does not require a predefined feature model, but instead reconstructs variability information directly from source and build metadata.

Centrality Measures for Prioritization. A complementary strand of research applies graph-theoretic concepts, especially centrality measures, to variability management and testing. Mohammed *et al.* [63] ranked configurations within a single feature model using a variety of centrality metrics, thus identifying configurations with disproportionately high influence. Levasseur *et al.* [49] used centrality, object-oriented metrics, and machine learning to prioritize unit tests, showing that structural graph properties can guide resource allocation in testing. Ferreira *et al.* [32] went further by proposing *variational* call graphs [31, 40] enriched with centrality information to pinpoint functions that may become vulnerabilities under certain feature combinations. While all these works demonstrate the value of centrality in highlighting critical elements, they remain tied either to explicit feature models, to testing strategies, or to specialized tasks such as vulnerability detection. Our approach builds on the same intuition—that centrality can capture relevance—but applies it to a new abstraction level: the atom dependency tree derived from Rust’s variability constructs. In doing so, RUSTYEx extends the applicability of centrality-based prioritization to ecosystems where feature models are implicit and configuration spaces are both large and sparsely exercised.

7 Conclusion

In this paper, we presented the first general method for prioritizing configurations in highly configurable software systems via a compiler-based refined ranking of features. Unlike previous approaches, our method does not rely on pre-existing feature models and explicitly accounts for feature dependencies and the extent of code affected by each feature.

The method combines inter-procedural static analysis to extract a *unified intermediate representation* (UIR), construction of the *feature dependency graph* and the *atom dependency tree*, centrality-based ranking, and CNF-based SAT solving. This combination enables the identification and generation of the most relevant configurations while reducing the total number of configurations to consider.

To demonstrate the practicality of our method, we implemented it in RUSTYEX, a fully automated tool for Rust software. Extensive evaluation on high-profile open-source Rust projects shows that RUSTYEX is scalable, robust, and efficient: it handles large codebases with thousands of features, maintains modest memory and runtime requirements, and achieves a high success rate across crates. The approach is language-agnostic and can be applied to other ecosystems with native variability support, such as C/C++ and Java.

By explicitly prioritizing the most relevant features and configurations, our method—and its RustyEx implementation—supports efficient testing, compiler optimizations, program comprehension, variability management, and regression analysis. Generated configurations are valid, sound, and representative of critical execution paths, providing a principled alternative to stochastic or uniform strategies.

Overall, this work introduces the first general, formally sound, and practical approach to configuration prioritization, improving efficiency and fault detection in highly configurable systems and laying the foundation for future extensions to other languages and domains.

References

- 1 Halimeh Agh, Aidin Azamnour, and Sefan Wagner. Software Product Line Testing: A Systematic Literature Review. *Empirical Software Engineering*, 29(146), September 2024.
- 2 Mustafa Al-Hajjaji, Thomas Thüm, Malte Lochau, Jens Meinicke, and Gunter Saake. Effective Product-Line Testing Using Similarity-Based Product Prioritization. *Software and Systems Modeling*, 18:499–521, February 2019.
- 3 Brian Anderson, Lars Bergstrom, Manish Goregaokar, Josh Matthews, Keegan McAllister, Jack Moffitt, and Simon Sapin. Engineering the Servo Web Browser Engine Using Rust. In Tao Xie and Dongmei Zhang, editors, *Proceedings of the 38th International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP'16)*, pages 81–89, Austin, TX, USA, May 2016. IEEE.
- 4 Sven Apel, Don Batory, Christian Kästner, and Gunter Saake. *Feature-Oriented Software Product Lines*. Springer, April 2013.
- 5 Sven Apel, Hendrik Speidel, Philipp Wendler, Alexander von Rhein, and Dirk Beyer. Detection of Feature Interactions Using Feature-Aware Verification. In Corina Păsăreanu and John Hosking, editors, *Proceedings of the 26th International Conference on Automated Software Engineering (ASE'11)*, pages 372–375, Lawrence, KS, USA, November 2011. IEEE.
- 6 Ebrahim Bagheri, Mohsen Asadi, Dragan Gasevic, and Samaneh Soltani. Stratified Analytic Hierarchy Process: Prioritization and Selection of Software Features. In Jan Bosch and Jaejoon Lee, editors, *Proceedings of the 14th International Software Product Line Conference (SPLC'10)*, Lecture Notes on Computer Science 6287, pages 300–315, Jeju Island, South Korea, September 2010. Springer.
- 7 Ebrahim Bagheri and Dragan Gasevic. Assessing the Maintainability of Software Product Line Feature Models Using Structural Metrics. *Software Quality Journal*, 19(3):576–612, January 2011.
- 8 Alex Bavelas. Communication Patterns in Task-Oriented Groups. *The Journal of the Acoustical Society of America*, 22(6):725–730, November 1950.
- 9 David Benavides, Sergio Segura, and Antonio Ruiz-Cortés. Automated Analysis of Feature Models 20 Years Later: A Literature Review. *Information Systems*, 35(6):615–636, September 2010.
- 10 Jan A. Bergstra and John V. Tucker. Equational Specifications, Complete Term Rewriting Systems, and Computable and Semicomputable Algebras. *Journal of ACM*, 42(6):1194–1230, November 1995.
- 11 Abraham Berman and Robert J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences*. Academic Press, January 1979.
- 12 Armin Biere, Katalin Fazekas, Mathias Fleury, and Maximilian Heisinger. CaDiCaL, Kissat, Paracooba, Plingeling and Treengeling Entering the SAT Competition 2020. In Tomáš Balyo, Nils Froleyks, Marijn J. H. Heule, Markus Iser, Matti Järvisalo, and Martin Suda, editors, *Proceedings of the SAT Competition 2020*, pages 50–53, Alghero, Italy, July 2020. University of Helsinki.
- 13 Francis Bloch, Matthew O. Jackson, and Pietro Tebaldi. Centrality Measures in Networks. *Social Choice and Welfare*, 61(2):413–453, April 2023.
- 14 Paolo Boldi and Sebastiano Vigna. Axioms for Centrality. *Internet Mathematics*, 10(3–4):222–262, September 2014.
- 15 Phillip Bonacich. Factoring and Weighting Approaches to Status Scores and Clique Identification. *Journal of Mathematical Sociology*, 2(1):113–120, 1972.
- 16 Stephen P. Borgatti. Centrality and Network Flow. *Social Networks*, 25(1):55–71, January 2005.
- 17 Chandrasekhar Boyapati, Robert Lee, and Martin Rinard. Ownership Types for Safe Programming: Preventing Data Races And Deadlocks. In Satoshi Matsuoka, editor, *Proceedings of the 17th Annual ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA'02)*, pages 211–230, Seattle, WA, USA, November 2002. ACM Press.

18 Shao-Fu Chen and Yu-Sung Wu. Linux Kernel Module Development with Rust. In Yousra Aafer, Shujun Li, and Yulei Wu, editors, *Proceedings of the Conference on Dependable and Secure Computing (DSC'22)*, pages 1–2, Edinburgh, United Kingdom, June 2022. IEEE.

19 David G Clarke, John M Potter, and James Noble. Ownership types for flexible alias protection. In Craig Chambers, editor, *Proceedings of 13th International Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA '98)*, pages 48–64, Vancouver, BC, Canada, October 1998. ACM.

20 Andreas Classen, Maxime Cordy, Pierre-Yves Schobbens, Patrick Heymans, Axel Legay, and Jean-François Raskin. Featured Transition Systems: Foundations for Verifying Variability-Intensive Systems and Their Application to LTL Model Checking. *IEEE Transactions on Software Engineering*, 39(8):1069–1089, August 2013.

21 Paul Clements and Linda Northrop. *Software Product Lines: Practices and Patterns*. Addison-Wesley, August 2001.

22 Myra B. Cohen, Matthew B. Dwyer, and Jiangfan Shi. Interaction Testing of Highly-Configurable Systems in the Presence of Constraints. In David S. Rosenblum and Sebastian G. Elbaum, editors, *Proceedings of the 16th International Symposium on Software Testing and Analysis (ISSTA'07)*, pages 129–139, London, United Kingdom, July 2007. ACM.

23 Keith D. Cooper and Linda Torczon. *Engineering a Compiler*. Morgan Kaufmann, November 2022.

24 Kousik Das, Sovan Samanta, and Madhumangal Pal. Study on Centrality Measures in Social Networks: A Survey. *Social Network Analysis and Mining*, 8:13:1–13:11, February 2018.

25 Sanjoy Dasgupta. Learning Polytrees. In Kathryn B. Laskey and Henri Prade, editors, *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence (UAI'99)*, pages 134–141, Stockholm, Sweden, July 1999. Morgan Kaufmann Publishers Inc.

26 Niklas Eén and Niklas Sörensson. An Extensible SAT-Solver. In Enrico Giunchiglia and Armando Tacchella, editors, *Proceedings of the 6th International Conference on Theory and Applications of Satisfiability Testing (SAT'03)*, LNCS 2919, pages 502–518, Santa Margherita Ligure, Italy, May 2003. Springer.

27 Sascha El-Sharkawy, Adam Krafczyk, and Klaus Schmid. An Empirical Study of Configuration Mismatches in Linux. In Myra Cohen and Mathieu Acher, editors, *Proceedings of the 21st International Systems and Software Product Line Conference (SPLC'17)*, pages 19–28, Sevilla, Spain, September 2017. ACM.

28 Sascha El-Sharkawy, Adam Krafczyk, and Klaus Schmid. Fast Static Analyses of Software Product Lines: An Example With More Than 42,000 Metrics. In Maxime Cordy and Mathieu Acher, editors, *Proceedings of the 14th International Working Conference on Variability Modelling of Software-Intensive Systems (VaMoS'20)*, pages 1–9, Magdeburg Germany, February 2020. ACM.

29 Sascha El-Sharkawy, Nozomi Yamagishi-Eichler, and Klaus Schmid. Metrics for Analyzing Variability and Its Implementation in Software Product Lines: A Systematic Literature Review. *Information and Software Technology*, 106:1–30, February 2019.

30 Norman E. Fenton. *Software Metrics: a Rigorous Approach*. London: Chapman & Hall, 1991.

31 Gabriel Ferreira, Christian Kästner, Jürgen Pfeffer, and Sven Apel. Characterizing Complexity of Highly-Configurable Systems with Variational Call Graphs: Analyzing Configuration Options Interactions Complexity in Function Calls. In *Proceedings of the Symposium and Bootcamp on the Science of Security (HotSoS'15)*, pages 1–2, Urbana-Champaign, IL, USA, April 2015. ACM.

32 Gabriel Ferreira, Momin Malik, Christian Kästner, Jürgen Pfeffer, and Sven Apel. Do #ifdefs Influence the Occurrence of Vulnerabilities? An Empirical Study of the Linux Kernel. In Rick Rabiser and Bing Xie, editors, *Proceedings of the 20th International Systems and Software Product-Line Conference (SPLC'16)*, pages 65–73, Beijing, China, September 2016. ACM.

33 Linton C. Freeman. A Set of Measures of Centrality Based on Betweenness. *Sociometry*, 40(1):35–41, March 1977.

- 34 Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Professional Computing Series. Addison-Wesley, Reading, MA, USA, 1995.
- 35 Jean-Yves Girard. Linear Logic. *Theoretical Computer Science*, 50(1):1–101, 1987.
- 36 Jean-Yves Girard, Yves Lafont, and Laurent Regnier. *Advances in Linear Logic*. Cambridge University Press, July 1995.
- 37 Axel Halin, Alexandre Nuttinck, Mathieu Acher, Xavier Devroey, Gilles Perrouin, and Benoit Baudry. Test Them All, Is It Worth It? Assessing Configuration Sampling on the JHipster Web Development Stack. *Empirical Software Engineering*, 24(2):674–717, July 2019.
- 38 Colin Howson. *Logic with Trees: An Introduction to Symbolic Logic*. Routledge, February 1997.
- 39 Muhammad Idham, Shahliza Abd Halim, Dayang Norhayati Abang Jawawi, Zalmiyah Zakaria, Aldo Erianda, and Nachnoer Arss. Test Case Prioritization for Software Product Line: A Systematic Mapping Study. *Journal on Informatics Visualization*, 7(3-2):2126–2134, November 2023.
- 40 Neil F. Johnson. *Simply Complexity: A Clear Guide to Complexity Theory*. Oneworld Publications, October 2009.
- 41 Kyo C. Kang, Sholom G. Cohen, James A. Hess, William E. Novak, and A. Spencer Peterson. Feature-Oriented Domain Analysis (FODA) Feasibility Study. Technical Report CMU/SEI-90-TR-21, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, November 1990.
- 42 Ahmet Serkan Karataş, Halit Oğuztüzün, and Ali Doğru. Global Constraints on Feature Models. In David Cohen, editor, *Proceedings of the 16th International Conference on Principles and Practice of Constraint Programming (CP'10)*, LNCS 6308, pages 537–551, St. Andrews, Scotland, September 2010. Springer.
- 43 Leo Katz. A New Status Index Derived From Sociometric Analysis. *Psychometrika*, 18(1):39–43, March 1953.
- 44 Andrew Kennedy and Claudio V. Russo. Generalized Algebraic Data Types and Object-Oriented Programming. In Richard P. Gabriel, editor, *Proceedings of 19th ACM International Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA '05)*, pages 21–40, San Diego, CA, USA, October 2005. ACM.
- 45 Charles W. Krueger. New Methods in Software Product Line Practice. *Communications of the ACM*, 49(12):37–40, December 2006.
- 46 Andrea Landherr, Bettina Friedl, and Julia Heidemann. A Critical Review of Centrality Measures in Social Networks. *Business and Information Systems Engineering*, 2(5):371–385, October 2010.
- 47 Jihyun Lee and Sunmyung Hwang. Combinatorial Test Design Using Design-Time Decisions for Variability. *Journal of Software Engineering and Knowledge Engineering*, 29(8):1141–1158, August 2019.
- 48 Daniel J. Lehmann and Michael B. Smyth. Algebraic Specification of Data Types: A Synthetic Approach. *Journal of Mathematical Systems Theory*, 14(2):97–139, December 1981.
- 49 Marc-Antoine Levasseur and Mourad Badri. Prioritizing Unit Tests Using Object-Oriented Metrics, Centrality Measures, and Machine Learning Algorithms. *Innovations in Systems and Software Engineering*, pages 1–27, February 2024.
- 50 Hongyu Li, Liwei Guo, Yexuan Yang, Shangguang Wang, and Mengwei Xu. An Empirical Study of Rust-for-Linux: The Success, Dissatisfaction, and Compromise. In Saurabh Bagchi and Yiying Zhang, editors, *Proceedings of the USENIX Annual Technical Conference (USENIX'24)*, pages 425–443, Santa Clara, CA, USA, July 2024. Curran Associates, Inc.
- 51 Zhuohua Li, Jincheng Wang, Mingshen Sun, and John C. S. Lui. MirChecker: Detecting Bugs in Rust Programs via Static Analysis. In Giovanni Vigna and Elaine Shi, editors, *Proceedings of the 2021 Conference on Computer and Communications Security (CCS'21)*, pages 2183–2196, Virtual, South Korea, November 2021. ACM.

52 Jörg Liebig, Sven Apel, Andreas Janker, Florian Garbe, and Sebastian Oster. Handling Static Configurability in Refactoring Engines. *Computer*, 50(7):44–53, 2017.

53 Jörg Liebig, Andreas Janker, Florian Garbe, Sven Apel, and Christian Lengauer. Morpheus: Variability-Aware Refactoring in the Wild. In Gerardo Canfora and Sebastian Elbaum, editors, *Proceedings of the 37th International Conference on Software Engineering (ICSE'15)*, pages 380–391, Florence, Italy, May 2015. IEEE.

54 Malte Lochau, Sebastian Oster, Ursula Goltz, and Andy Schürr. Model-Based Pairwise Testing for Feature Interaction Coverage in Software Product Line Engineering. *Software Quality Journal*, 20(3-4):567–604, September 2012.

55 Eva Magnusson and Görel Hedin. Circular Reference Attributed Grammars—Their Evaluation and Applications. *Science of Computer Programming*, 68(1):21–37, August 2007.

56 Mike Mannion and Hermann Kaindl. Using Binary Strings for Comparing Products From Software-Intensive Systems Product Lines. In Ina Schaefer and Maurice H. ter Beek, editors, *Proceedings of the 25th International Software Product Line Conference (SPLC'21)*, pages 257–266, Leicester, United Kingdom, September 2021. ACM.

57 Massimo Marchiori and Vito Latora. Harmony in the Small-World. *Physica A: Statistical Mechanics and its Applications*, 285(3-4):539–546, October 2000.

58 Hugo Martin, Mathieu Acher, Juliana Alves Pereira, and Jean-Marc Jézéquel. A Comparison of Performance Specialization Learning for Configurable Systems. In Ina Schaefer and Maurice H. ter Beek, editors, *Proceedings of the 25th International Software Product Line Conference (SPLC'21)*, pages 46–57, Leicester, United Kingdom, September 2021. ACM.

59 Nicholas D. Matsakis and Felix S. Klock. The Rust Language. *ACM SIGAda Letters*, 34(3):103–104, October 2014.

60 Antoni Mazurkiewicz. Introduction to Trace Theory. In Volker Diekert and Grzegorz Rozenberg, editors, *The Book of Traces*, chapter 1, pages 3–41. World Scientific Publishing, March 1995.

61 Thomas J. McCabe. A Complexity Measure. *IEEE Transactions on Software Engineering*, 2(4):308–320, December 1976.

62 Jean Melo, Elvis Flesborg, Claus Brabrand, and Andrzej Wąsowski. A Quantitative Analysis of Variability Warnings in Linux. In Ina Schaefer and Vander Alves, editors, *Proceedings of the 10th International Workshop on Variability Modelling of Software-Intensive Systems (VaMoS'10)*, pages 3–8, Salvador, Brazil, January 2016. ACM.

63 Fathiya Mohammed, Mike Mannion, Hermann Kaindl, and James Patterson. Evaluating the Relative Importance of Product Line Features using Centrality Metrics. In Massimo Mecella and Arend Rensink, editors, *Proceedings of the 19th International Conference on Software Technologies (ICSOFT 2024)*, pages 469–476, Dijon, France, July 2024. SciTe Press.

64 Mark E. J. Newman. Scientific Collaboration Networks. II. Shortest Paths, Weighted Networks, and Centrality. *Physical Review E*, 64(1):016132, June 2001.

65 Mark E. J. Newman. *Networks: An Introduction*. Oxford University Press, first edition, March 2010.

66 Martin Odersky. Observers for Linear Types. In Bernd Krieg-Brückner, editor, *Proceedings of the 4th European Symposium on Programming (ESOP'92)*, LNCS 582, pages 390–407, Rennes, France, February 1992. Springer.

67 David von Oheimb. Hoare Logic for Mutual Recursion and Local Variables. In C. Pandu Rangan, Venkatesh Raman, and Ramaswamy Ramanujam, editors, *Proceedings of the 19th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'99)*, LNCS 1738, pages 168–180, Chennai, India, December 1999. Springer.

68 Kurt M. Olender and Leon J. Osterweil. Cecil: A Sequencing Constraint Language for Automatic Static Analysis Generation. *IEEE Transactions on Software Engineering*, 16(3):268–280, March 1990.

69 Kurt M. Olender and Leon J. Osterweil. Interprocedural Static Analysis of Sequencing Constraints. *Transactions on Software Engineering and Methodology*, 1(1):21–52, January 1992.

70 Tore Opsahl, Filip Agneessens, and John Skvoretz. Node Centrality in Weighted Networks: Generalizing Degree and Shortest Paths. *Social Networks*, 32(3):245–251, July 2010.

71 Sebastian Oster, Florian Markert, and Philipp Ritter. Automated Incremental Pairwise Testing of Software Product Lines. In Jan Bosch and Jaejoon Lee, editors, *Proceedings of the 14th International Software Product Line Conference (SPLC'10)*, Lecture Notes on Computer Science 6287, pages 196–210, Jeju Island, South Korea, September 2010. Springer.

72 José A. Parejo, Ana B. Sánchez, Sergio Segura, Antonio Ruiz-Cortés, Roberto E. Lopez-Herrejon, and Alexander Egyed. Multi-Objective Test Case Prioritization in Highly Configurable Systems: A Case Study. *Journal of Systems and Software*, 122:287–310, December 2016.

73 Sachin Patel, Priya Gupta, and Vipul Shah. Combinatorial Interaction Testing with Multi-Perspective Feature Models. In Yang Liu and Pang, editors, *Proceedings of the 6th International Conference on Software Testing, Verification and Validation Workshops (ICSTW'13)*, pages 321–330, Luxembourg City, Luxembourg, March 2013. IEEE.

74 Zhenlian Peng, Jian Wang, Keqing He, and Hongtao Li. An Approach for Prioritizing Software Features Based on Node Centrality in Probability Network. In Georgia Kapitsaki and Eduardo Almeida, editors, *Proceedings of the 15th International Conference on Software Reuse (ICSR'15)*, LNCS 9679, pages 106–121, Limassol, Cyprus, June 2015. Springer.

75 Gilles Perrouin, Sagar Sen, Jacques Klein, Benoît Baudry, and Yves le Traon. Automated and Scalable T-Wise Test CASE Generation Strategies for Software Product Lines. In Ana Rosa Cavalli and Sudipto Ghosh, editors, *Proceedings of the 3rd International Conference on Software Testing, Verification and Validation (ICST'10)*, pages 459–468, Paris, France, April 2010. IEEE.

76 Benjamin C. Pierce. *Types and Programming Languages*. MIT Press, February 2002.

77 Richard Pohl, Kim Lauenroth, and Klaus Pohl. A Performance Comparison of Contemporary Algorithmic Approaches for Automated Analysis Operations on Feature Models. In Corina păsăreanu and John Hosking, editors, *Proceedings of the 26th International Conference on Automated Software Engineering (ASE'11)*, pages 313–322, Lawrence, KS, USA, November 2011. IEEE.

78 Christian Prehofer. Feature-Oriented Programming: A Fresh Look at Objects. In Mehmet Akşit and Satoshi Matsuoka, editors, *Proceedings of the 11th European Conference on Object-Oriented Programming (ECOOP'97)*, Lecture Notes in Computer Science 1241, pages 419–443, Helsinki, Finland, June 1997. Springer.

79 Xiao Qu, Myra B. Cohen, and Gregg Rothermel. Configuration-Aware Regression Testing: An Empirical Study of Sampling and Prioritization. In Andreas Zeller, editor, *Proceedings of the 17th International Symposium on Software Testing and Analysis (ISSTA'08)*, pages 75–86, Seattle, WA, USA, July 2008. ACM.

80 Ana B. Sánchez, Sergio Segura, José A. Parejo, and Antonio Ruiz-Cortés. Variability Testing in the Wild: The Drupal CASE Study. *Software and Systems Modeling*, 16(2):173–194, April 2017.

81 Ana B. Sánchez, Sergio Segura, and Antonio Ruiz-Cortés. A Comparison of Test Case Prioritization Criteria for Software Product Lines. In Lauric Williams and Claes Wohlin, editors, *Proceedings of 7th International Conference on Software Testing, Verification and Validation (ICST'14)*, pages 41–50, Cleveland, OH, USA, March 2014. IEEE.

82 John R. Seeley. The Net of Reciprocal Influence; A Problem in Treating Sociometric Data. *Canadian Journal of Psychology*, 3(4):234–240, December 1949.

83 Christoph Seidl, Tim Winkelmann, and Ina Schaefer. A Software Product Line of Feature Modeling Notations and Cross-Tree Constraint Languages. In Andreas Oberweis and Ralf H. Reussner, editors, *Proceedings of the 13th Edition of Modellierung (Modellierung'16)*, LNI P-254, pages 157–172, Karlsruhe, Germany, March 2016. Springer.

84 Norbert Siegmund, Alexander Grebhanhn, Sven Apel, and Christian Kästner. Performance-Influence Models for Highly Configurable Systems. In Mark Harman and Patrick Heymans,

editors, *Proceedings of the 10th Joint Meeting on Foundations of Software Engineering (ES-EC/FSE'15)*, pages 284–294, Bergamo, Italy, September 2015. ACM.

85 Julio Sincero, Horst Schirmeier, Wolfgang Schröder-Preikschat, and Olaf Spinczyk. Is the Linux Kernel a Software Product Line? In Frank van der Linden and Björn Lnnell, editors, *Proceedings of the 2nd International Workshop on Open Source Software and Product Lines (SPLC-OSSPL'07)*, Kyoto, Japan, September 2007.

86 Julio Sincero, Reinhard Tartler, Daniel Lohmann, and Wolfgang Schröder-Preikschat. Efficient Extraction and Analysis of Preprocessor-Based Variability. In Jaakko Järvi, editor, *Proceedings of the 9th International Conference on Generative Programming and Component Engineering (GPCE'10)*, pages 33–42, Eindhoven, The Netherlands, October 2010. ACM.

87 Niklas Sörensson and Niklas Eén. MiniSat v1.13 – A SAT Solver with Conflict-Clause Minimization. In Fahiem Bacchus, editor, *Proceedings of the 8th International Conference on Theory and Applications of Satisfiability Testing (SAT'05)*, LNCS 3569, pages 1–2, St. Andrews, Scotland, June 2005. Springer. Poster.

88 Sabrina Souto and Marcelo d'Amorim. Time-space efficient regression testing for configurable systems. *Journal of Systems and Software*, 137:733–746, March 2018.

89 Bjarne Stroustrup. *The Design and Evolution of C++*. Addison-Wesley, first edition, March 1994.

90 Reinhard Tartler, Christian Dietrich, Julio Sincero, Wolfgang Schröder-Preikschat, and Daniel Lohmann. Static Analysis of Variability in System Software: The 90,000 #ifdefs Issue. In Garth Gibson and Nickolai Zeldovich, editors, *Proceedings of the USENIX Annual Technical Conference (USENIX ATC'14)*, pages 421–432, Philadelphia, PA, USA, June 2014. USENIX Association.

91 Reinhard Tartler, Daniel Lohmann, Julio Sincero, and Wolfgang Schröder-Preikschat. Feature Consistency in Compile-Time-Configurable System Software: Facing the Linux 10,000 Feature Problem. In Gernot Heiser, editor, *Proceedings of the 6th Conference on Computer Systems (EuroSys'11)*, pages 47–60, Salzburg, Austria, April 2011. ACM.

92 Matthew Treinish, Ivan Carvalho, Georgios Tsilimigounakis, and Nahum Sá. rustworkx: A High-Performance Graph Library for Python. *Journal of Open Source Software*, 7(79):3968:1–3968:32, November 2022.

93 David A. Turner. Miranda: A Non-Strict Functional Language with Polymorphic Types. In Jean-Pierre Jouannaud, editor, *Proceedings of the 1st International Conference on Functional Programming Languages and Computer Architecture (FPCA'85)*, LNCS 201, pages 1–16, Nancy, France, September 1985. Springer.

94 Miguel Velez, Pooyan Jamshidi, Florian Sattler, Norbert Siegmund, Sven Apel, and Christian Kästner. ConfigCrusher: Towards White-Box Performance Analysis for Configurable Systems. *Automated Software Engineering*, 27:265–300, August 2020.

95 Miguel Velez, Pooyan Jamshidi, Norbert Siegmund, Sven Apel, and Christian Kästner. White-Box Analysis Over Machine Learning: Modeling Performance of Configurable Systems. In *Proceedings of the 43rd International Conference on Software Engineering (ICSE'21)*, pages 1072–1084, Madrid, Spain, May 2021. IEEE.

96 Miguel Velez, Pooyan Jamshidi, Norbert Siegmund, Sven Apel, and Christian Kästner. On Debugging the Performance of Configurable Software Systems: Developer Needs and Tailored Tool Support. In Daniela Damian and Andreas Zeller, editors, *Proceedings of the 44th International Conference on Software Engineering (ICSE'22)*, pages 1571–1583, Pittsburgh, USA, May 2022. ACM.

97 S. Venkatakeerthy, Rohit Aggarwal, Shalini Jain, Maunendra Sankar Desarkar, and Ramakrishna Updrasta. IR2Vec: LLVM IR Based Scalable Program Embeddings. *ACM Transactions on Architecture and Code Optimization*, 17(4):32:1–32:27, December 2020.

98 Philip Wadler. Linear Types Can Change the World! In Manfred Broy and Cliff B. Jones, editors, *Proceedings of the 2nd Working Conference on Programming Concepts and Methods (IFIP'90)*, pages 561–582, Sea of Galilee, Israel, April 1990. North-Holland.

99 Claes Wohlin, Per Runeson, Martin Höst, Magnus C. Ohlsson, Björn Regnell, and Anders Wesslén. *Experimentation in Software Engineering*. Springer, 2012.