



Finding Broken Linux Configuration Specifications by Statically Analyzing the Kconfig Language

Jeho Oh*

University of Texas at Austin
Austin, TX, USA
jeho.oh@utexas.edu

Julian Braha

University of Central Florida
Orlando, FL, USA
julianbraha@knights.ucf.edu

Necip Fazıl Yıldıran*

University of Central Florida
Orlando, FL, USA
yildiran@knights.ucf.edu

Paul Gazzillo

University of Central Florida
Orlando, FL, USA
paul.gazzillo@ucf.edu

ABSTRACT

Highly-configurable software underpins much of our computing infrastructure. It enables extensive reuse, but opens the door to broken configuration specifications. The configuration specification language, Kconfig, is designed to prevent invalid configurations of the Linux kernel from being built. However, the astronomical size of the configuration space for Linux makes finding specification bugs difficult by hand or with random testing. In this paper, we introduce a software model checking framework for building Kconfig static analysis tools. We develop a formal semantics of the Kconfig language and implement the semantics in a symbolic evaluator called *kclause* that models Kconfig behavior as logical formulas. We then design and implement a bug finder, called *kismet*, that takes *kclause* models and leverages automated theorem proving to find unmet dependency bugs. *kismet* is evaluated for its precision, performance, and impact on kernel development for a recent version of Linux, which has over 140,000 lines of Kconfig across 28 architecture-specific specifications. Our evaluation finds 781 bugs (151 when considering sharing among Kconfig specifications) with 100% precision, spending between 37 and 90 minutes for each Kconfig specification, although it misses some bugs due to underapproximation. Compared to random testing, *kismet* finds substantially more true positive bugs in a fraction of the time.

CCS CONCEPTS

• **Software and its engineering** → **Software configuration management and version control systems; Automated static analysis; Software testing and debugging.**

KEYWORDS

software configuration, Kconfig, formal verification, static analysis

* Co-first authors.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ESEC/FSE '21, August 23–28, 2021, Athens, Greece

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8562-6/21/08...\$15.00

<https://doi.org/10.1145/3468264.3468578>

ACM Reference Format:

Jeho Oh, Necip Fazıl Yıldıran, Julian Braha, and Paul Gazzillo. 2021. Finding Broken Linux Configuration Specifications by Statically Analyzing the Kconfig Language. In *Proceedings of the 29th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '21)*, August 23–28, 2021, Athens, Greece. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3468264.3468578>

1 INTRODUCTION

Highly-configurable software product lines underpin much of our computing infrastructure, because configurability enables reuse without having to reprogram the software for new devices or applications. The Linux kernel is one such example of highly-configurable software that is used in billions of computing devices. With over 15,000 configuration options controlling everything from drivers, architecture, memory management, and more, there are over $2^{15,000}$ combinations, if only considering Boolean options. This extreme configurability makes its widespread use possible, but also opens the door to invalid configurations, which produce broken variations of the software.

To mitigate the chance of misconfiguration, developers provide *configuration specifications*, which define the intended combinations of configuration options. These specifications may be as simple as a text file describing configuration instructions, or as sophisticated as a machine-readable specification enforced at build time. These specifications, if only implicitly, define a software product line's *feature model*, i.e., the legal configurations of the software. In Linux and other systems software, such as BusyBox and coreboot, developers use the Kconfig language to specify configuration options, as well as their dependencies.

While Linux is effectively a software product line, its specification language, Kconfig, is unlike typical feature modeling languages [47]. Kconfig has complex semantics and additional language features, such as invisible variables, automated option selection, and user-interface support. With over 140,000 lines of Kconfig specifications in the Linux kernel, its complex behavior makes maintenance challenging. One example is the common pitfall highlighted in Kconfig's manual [28], the *unmet dependency bug*. These bugs lead to illegal configurations when developers unwittingly mix conflicting constructs in dependency specifications. With thousands of potentially vulnerable constructs and an ever-changing

specification, finding such bugs by hand is a practically impossible task.

Existing work on the analysis of Kconfig is focused on extracting a feature model, rather than checking for Kconfig bugs. Having a Linux feature model has been useful for applications outside of Kconfig, including measuring feature model hierarchies [10, 44, 45], supporting variability-aware analysis of C [11, 18, 20, 24, 26, 27, 31, 51, 56, 58], and dead code elimination [52, 53]. However, these tools make assumptions about Kconfig semantics that, while appropriate for their respective applications, make them less amenable to bug finding. For instance, KconfigReader explicitly omits modeling the language semantics that lead to unmet dependencies, leaving a checker as a separate problem [25, 30]. The other tools do the same and have less support for Kconfig semantics, omitting some parts of the Linux feature model altogether [16, 45]. Moreover, by focusing on feature modeling, prior tools bake in decisions about the analysis domain, i.e., a feature model, which limits the feasibility of repurposing the work for Kconfig bug finding and other source level tools.

In this paper, we introduce a software model checking framework for building Kconfig static analysis tools. Inspired by model checking frameworks for program code [1], we base our static analysis on our newly-developed formal semantics of the Kconfig language and leverage automated theorem proving to model Kconfig behavior and find bugs. Of existing descriptions of Kconfig semantics, all but one are either informal or example-based [15, 16] which, having no formalization, are not amenable to automated reasoning. The one prior formal semantics uses an idealized Kconfig syntax rather than Kconfig's actual grammar, is missing language behavior (including that leading to unmet dependencies), and uses an abstract domain designed for feature modeling [43]. In contrast, we develop a formal semantics of the concrete behavior of Kconfig when it checks a configuration file's validity. We define our semantics over the syntax derived from Kconfig's actual implementation, which contains a bison grammar specification.

This approach to Kconfig semantics has three key benefits over prior efforts. First, it is formal, making it possible to use automated reasoning tools. Second, it is concrete, which decouples the description of Kconfig semantics from any particular analysis objective. This leaves decisions about how to abstract Kconfig behavior to specific applications and should reduce future effort to design new Kconfig analyses. Third, it simplifies modeling Kconfig since, as we show, we can methodically derive an abstraction of Kconfig behavior from this concrete semantics.

To demonstrate the utility of our approach, we design and implement an analysis that finds the unmet dependency bugs highlighted in Kconfig's manual and is, to our knowledge, the first static analysis for finding such bugs. We first define the bug as a formal property in terms of the semantics, then show how a checker can be methodically derived from the semantics. We underapproximate non-Boolean options and use aggressive optimization to yield a bug-finder that is both fast and very precise. Moreover, it can also automatically localize and generate test cases for the unmet dependency bugs it finds. The trade-off is decreased recall due to false negatives, although we show that these are less common due to the rarity of non-Boolean options.

We implement the bug-finder and evaluate it on a recent version of the Linux kernel source, which contains over 140,000 lines of Kconfig describing 28 architecture-specific Kconfig specifications. Our results show that our bug finding is both precise and fast. The bug-finder finds 781 alarms (151 when considering sharing among Kconfig specifications) over all Linux kernel architectures' Kconfig specifications, all of which are verified true positives, for a precision of 100%. While such precision might be unusually high for a programming language analyzer, the Kconfig language has no iteration or recursion that would require overapproximation. With our optimizations, our bug finder takes an average of 40 minutes for one Kconfig specification, checking over 10,000 constructs.

While we are still in the process of reporting all bugs found by our tool, Linux maintainers have so far already confirmed 38 of our reports and committed 15 of our patches into the mainline Linux kernel repository, demonstrating the utility of our tooling. Committing patches is a manual process, requiring discussion with kernel maintainers, so investigating, reporting, and submitting patches for the alarms is ongoing.

Since, to our knowledge, no other static bug finder for unmet dependencies exists, we compare to Kconfig's built-in `randconfig` tool, the de facto standard random configuration testing tool used by Linux maintainers and the Intel 0-day test service [5]. Given the same amount of time to find bugs, `randconfig` yields only 98 alarms compared to our tool's 781. Even letting random testing run for over four days for each Kconfig specification, 135x longer than our tool's total time, the testing approach still only finds 175 bugs, far fewer than our tool. The random testing approach did find eight bugs missed by our tool, reflecting the tradeoff in performance gained by underapproximation. Even with this limitation, our tool finds many more bugs while taking far less time, a useful complement to testing.

In summary, this paper makes the following contributions:

- A formal semantics of Kconfig's concrete behavior (Section 3).
- An efficient design of a bug-finder and localizer for unmet dependency bugs (Section 4).
- An implementation of the bug finder, along with reusable components for creating Kconfig analyzers (Section 5).
- An experimental evaluation of the bug finder's precision, performance, and impact (Section 6).

2 OVERVIEW

In this section we introduce the Kconfig language, illustrate an unmet dependency bug, and summarize how our formal semantics enables the design of a static analysis to find such bugs.

2.1 Introduction to the Kconfig Language

Figure 1 is a simplified snippet of Kconfig from Linux v5.4.4. Configuration options are defined with the `config` construct (lines 1, 7, and 12). Inside each `config` declaration is a block of constructs that define the option's type (e.g. Boolean), constraints on its use, and text used by Kconfig to populate a user interface.

Lines 2, 8, and 13 are the type declarations. A `bool` option (line 8) has two possible settings, `y` or `n`. `y` means the feature is on and compiled into the kernel, and `n` means the feature is off and omitted

```

1 config TOUCHSCREEN_ADC
2     tristate
3     prompt "Generic ADC based touchscreen"
4     depends on IIO && INPUT_TOUCHSCREEN
5     select IIO_BUFFER_CB
6
7 config IIO_BUFFER
8     bool
9     prompt "Enable buffer support within IIO"
10    depends on IIO
11
12 config IIO_BUFFER_CB
13    tristate
14    prompt "IIO callback buffer"
15    depends on IIO && IIO_BUFFER

```

Figure 1: An example Kconfig specification that allows an unmet dependency violation and leads to a build error. Adapted from Linux source: `drivers/input/touchscreen/Kconfig`, `drivers/iio/Kconfig`, and `drivers/iio/buffer/Kconfig`.

from the kernel. A tristate option (lines 2 and 13) adds an additional setting, `m`. `m` is like `y` except that the build system compiles a loadable kernel module instead of linking to the main kernel binary [28].

tristate and bool are the most common configuration options, representing more than 95% of options in the Linux Kconfig specifications. The other possible types are string for strings, int for decimal integers, and hex for hexadecimal numbers.

Constraints on options are defined using `depends on` (lines 4, 10, and 15) and `select` (line 5), but the Kconfig language prohibits circular dependencies. `depends on` defines a *direct dependency*, which provides requirements that should hold before the option can be enabled. The dependency is expressed with a Boolean expression of other options. For instance, line 4 means that `TOUCHSCREEN_ADC` may not be enabled unless `IIO && INPUT_TOUCHSCREEN` is true, i.e., when both `IIO` and `INPUT_TOUCHSCREEN` are also enabled.

A *reverse dependency*, given by the `select` construct, inverts the dependency relationship by forcing the target of the `select` to be enabled. For instance, line 5 means that whenever `TOUCHSCREEN_ADC` is enabled, `IIO_BUFFER_CB` is forced to be enabled. A reverse dependency can only enable another option, not disable it, and only applies to bool or tristate options. *Kconfig permits reverse dependencies to override direct dependencies*, which can lead to unmet dependency bugs.

Options with a prompt are *visible* options that a user may enable. The prompt construct defines the prompt string for use in a user interface (lines 3, 9, and 14). Non-visible options, i.e., those with no prompt construct, can only be set by a `select` construct or take a specification-defined default value (not shown in this example). The visibility of an option affects the behavior of a `config` construct in subtle ways, which we describe in the formal semantics of Kconfig.

An Unmet Dependency Bug in the Wild. Figure 1 has an unmet dependency bug found by our automated analysis. All three of the configuration options defined in this example control specific C compilation units that are only built into the kernel when the options are enabled. `IIO_BUFFER_CB` controls `industrialio-buffer-cb.o` and `IIO_BUFFER` control `industrialio-buffer.o`.

`industrialio-buffer-cb.o` calls functions that are defined in `industrialio-buffer.o`, so the former cannot be built without the latter, otherwise there would be a build error. The developers capture this build dependency with a direct dependency in the definition of the `IIO_BUFFER_CB` option (line 15). This constraint, by itself, would prevent a user from giving a configuration that leads to the build error.

The `select IIO_BUFFER_CB` construct on line 5, however, can override this direct dependency under certain conditions. Specifically, if a user (or another `select`) enables `TOUCHSCREEN_ADC`, the `select` automatically force-enables `IIO_BUFFER_CB`. Kconfig permits such a configuration to proceed to build, albeit with a warning. Still, the build will fail, and the user will have to manually correct their configuration file in order to avoid the unmet dependency.

While the Kconfig documentation warns of `select`'s pitfalls and recommends not using it to override dependencies, it is difficult to check by hand whether any of its 17,000+ uses have an unmet dependency bug.

2.2 An Unmet Dependencies Bug Finder

We create a formal model of the unmet dependency bug according to the semantics of Kconfig. First, we model the space of valid Kconfig configurations in formal logic automatically with our symbolic evaluator `kc1ause`. Next, `kismet` generates verification conditions to prove the absence of an unmet dependency for each `select` construct in the Kconfig specification. Not all reverse dependencies can cause unmet dependencies, so `kismet` needs to consider constraints from all configuration options to rule out infeasible ones. The resulting verification conditions are discharged to the Z3 SMT solver [13]. When an unmet dependency cannot be ruled out, `kismet` raises an alarm. It then switches to test case generation, converting any counterexamples to Linux configuration files. `kismet` uses these tests on Kconfig and the build system to expose real bugs.

To see how `kc1ause` models dependencies, take Figure 1's definition of `TOUCHSCREEN_ADC` (line 1). Since it has no reverse dependencies, it can only be enabled when its direct dependencies hold, i.e., enabling `TOUCHSCREEN_ADC` implies `IIO` and `INPUT_TOUCHSCREEN` are also both enabled:

$$\text{TOUCHSCREEN_ADC} \rightarrow \text{IIO} \wedge \text{INPUT_TOUCHSCREEN}$$

When an option has reverse dependencies, its direct dependencies do not have to hold if its reverse dependencies already do. For instance, enabling `IIO_BUFFER_CB` (line 12) implies that its direct or reverse dependencies hold:

$$\begin{aligned} \text{IIO_BUFFER_CB} \rightarrow & (\text{IIO} \wedge \text{IIO_BUFFER} \\ & \vee \text{TOUCHSCREEN_ADC}) \end{aligned}$$

An unmet dependency happens when an option's reverse dependencies hold but its direct dependencies do not. For instance, an unmet dependency happens when `TOUCHSCREEN_ADC` force-enables

```

kconfig ::= statement+
statement ::= config | choice
config ::= config symbol type constrnts select*
choice ::= choice type constrnts config+ endchoice
type ::= bool | tristate
constrnts ::= prompt depends+ default+
prompt ::= prompt word if expr
default ::= default val if expr
depends ::= depends on expr
select ::= select symbol if expr
expr ::= expr && expr | expr || expr | ! expr | symbol
val ::= y | n

```

Figure 2: Formal syntax of a core fragment of Kconfig.

IIO_BUFFER_CB even though IIO_BUFFER_CB's direct dependencies are infeasible. This unmet dependency can be formalized as follows:

```

TOUCHSCREEN_ADC ∧ (IIO ∧ INPUT_TOUCHSCREEN)
∧ IIO_BUFFER_CB ∧ ¬(IIO ∧ IIO_BUFFER)

```

kismet tries to prove the *negation* of this condition, since it verifies the absence of unmet dependencies. If the proof fails, kismet raises an alarm and switches to test case generation.

3 THE SEMANTICS OF KCONFIG

The Kconfig language is a declarative configuration specification language. At its core, Kconfig takes a configuration file, which is a mapping from configuration options to their concrete values, and determines whether the configuration file is valid according to the developer's specifications. Developers use Kconfig language constructs to define configuration options, declaring their names, types, and any dependencies they have on other configuration options. Kconfig also supports user interfaces, and the language has additional constructs, such as `help`, to specify text elements of the interface. These do not affect the buildability of configuration files and act as comments.

We developed this formal semantics by studying the Kconfig manual, Kconfig's source code, as well as informal descriptions and examples from prior work [15, 16]. To check the fidelity of the semantics, we used new and existing benchmarks [15, 16], generated random configuration files fed to Kconfig as input, and evaluated this paper's bug-finder, which requires a correct semantics for its analysis to be precise. Given the size of the semantics, having dozens of rules and still more syntactic sugar rules, we highlight a core fragment of the language here, and document the remaining rules in an openly-archived formal semantics [40].

3.1 Configuration Declarations

Figure 2 shows the syntax of a core fragment of the Kconfig language for `bool` configuration options. A *kconfig* file contains a list of *statements*, which are either a configuration option declaration or a choice construct. A configuration option, *config*, has a *type*, constraints for direct dependencies, and zero or more *select* constructs for reverse dependencies.

Figure 3a defines the semantic valuation function *S* for statements. *S* functions take an immutable configuration file σ as input and return whether the configuration is valid or invalid, i.e., buildable or not. S_1 evaluates a Kconfig specification's list of statements by checking whether all statements are valid according to the input.

S_2 is the valuation function for config statements. The number of cases reflects the complexities of Kconfig's validity checking. The first covers reverse dependencies, using the valuation function *R*, which searches the entire *kconfig* file. (In practice, the Kconfig implementation memoizes reverse dependencies during parsing to avoid repeatedly traversing the syntax tree.) If a reverse dependency holds, that means the option must be enabled, i.e., $\sigma(sym) = y$, otherwise the configuration file does not match the specification.

The second case of S_2 handles a direct dependency when the reverse dependency does not hold. In this case, an option is valid regardless of its setting, because a user is free to enable or disable it. The third case covers non-visible configuration options, which have no prompt, so the option's value must match the specified default. The fourth case covers when none of the option's dependencies hold. Lastly, if none of these conditions are met, the configuration file is not valid.

Dependencies for non-Boolean types (string, int, and hex) behave similarly to `bool` and `tristate`, but there are additional constraints and expressions such as `range` and `inequalities`. The full semantics describes these differences [40].

3.2 Reverse Dependencies

To find any reverse dependencies for an option, Kconfig needs to search the entire specification for a `select` that can enable the option. This is partly why tracking down unmet dependencies is so difficult.

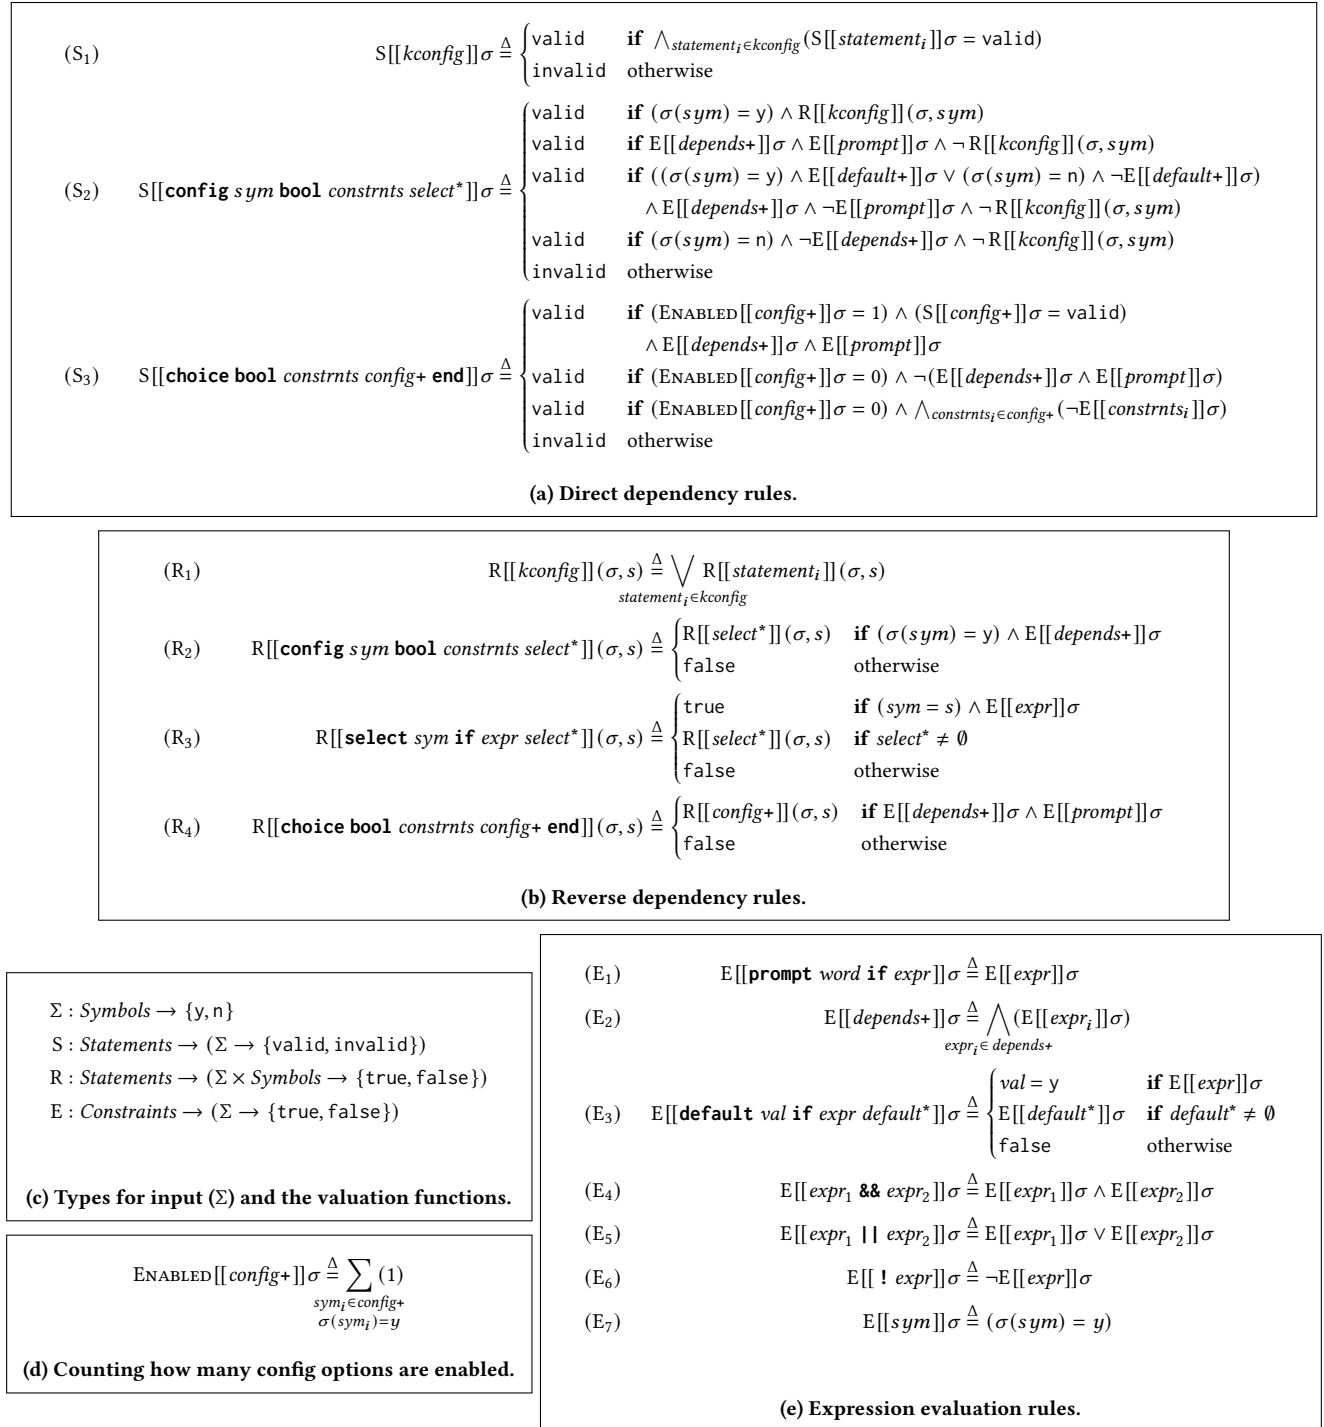
Figure 3b defines the valuation function *R* for reverse dependencies. It takes both the configuration file σ and a symbol name *s* and returns a Boolean value: `true` if that symbol is selected by some option or `false` if not. R_1 's disjunction reflects the need for only one `select` to force-enable an option.

R_2 checks to see if an option is enabled and its dependencies are met, then calls R_3 to evaluate any `select` constructs. R_3 checks whether there is a `select` for the input symbol *s*. R_4 checks whether any configuration option within a choice block selects symbol *s*.

Options other than `tristate` and `bool` cannot be the selector or selectee of a reverse dependency.

3.3 Choice Constructs

A choice construct defines a mutually-exclusive set of configuration options. Choices are useful in configuration specifications, because expressing them with Boolean logic alone is verbose. Figure 4 is an example of a choice from the Linux kernel that allows only one of several compression algorithms for a file system to be enabled. A choice block starts with a choice keyword (line 1) and ends with an `endchoice` (line 10). It contains a list of configuration options which, besides the mutual-exclusion rule, behave mostly like any other options, except that they cannot have reverse dependencies or default values.




```

1 choice
2   prompt "Decompressor parallelisation options"
3   depends on SQUASHFS
4   config SQUASHFS_DECOMP_SINGLE
5     bool "Single threaded compression"
6   config SQUASHFS_DECOMP_MULTI
7     bool "Use multiple decompressors"
8   config SQUASHFS_DECOMP_MULTI_PERCPU
9     bool "Use percpu multiple decompressors"
10 endchoice

```

Figure 4: An example of a choice construct.

The S_3 function in Figure 3a describes the choice block's semantics. The first case covers the mutual exclusion property, requiring that only one of the configuration options is enabled. This condition also recursively checks that all the nested configs' dependencies are valid.

Choice constructs also have their own direct dependencies, so Kconfig permits no options to be enabled when the choice dependencies are not met. The second case of S_3 covers this situation. The third case covers the situation when none of the nested config options' dependencies are met, in which case Kconfig also permits the choice to have no options enabled. choice constructs can also take the optional keyword to allow for no options to be enabled even if its direct dependencies are met. The rules are slightly different from a regular choice, and we present them in the full semantics [40].

The choice statement described above has bool type. The only other type a choice can take is tristate, in which case its behavior is the same as bool, except that more than one choice may be set to m.

3.4 Constraint Expressions

Figure 3e defines rules for evaluating constraint expressions, which return a Boolean true or false. E_1 , E_2 , and E_3 are the prompt, depends on, and default constructs, respectively. Each is a carrier for a logical expression, and it is their interaction with config and choice that gives them distinct meaning. The rest of the rules are typical Boolean operators (E_4 – E_7).

3.5 Syntactic Sugar

Kconfig has three additional statements that can be desugared to config and choice: if, menu, and menuconfig. Unlike the control-flow construct in programming languages, Kconfig's if is just syntactic sugar for adding constraints in bulk to its nested statements. The menu statement behaves like an if block, but also adds text to the user interface. menuconfig is a combination of config and menu. The full semantics [40] contains the desugaring rules for these.

The Kconfig language also has a great deal of flexibility in its syntax. Most of the behavior of a Kconfig specification is insensitive to the ordering of options and constraints. Therefore, our syntax defines ordering on constraints to reduce the number of syntactic sugars rules needed.

```

1 config X
2   select A if D_X
3   K_X // other constraints for X
4
5 config A
6   depends on D_A
7   K_A // other constraints for A
8 K_other // constraints from other configuration options

```

Figure 5: Components of an unmet dependency condition.

Kconfig has limited metaprogramming facilities via preprocessor constructs for file inclusion and macro expansion [2], which we do not model. Our implementation runs the preprocessor before symbolic evaluation to ensure that all files are included and macros are expanded.

4 DESIGNING THE BUG FINDER

Our bug-finder, called *kismet* works by generating a formula for each select describing the configurations under which it triggers an unmet dependency bug. This requires both syntax analysis, to identify select constructs, as well as semantic analysis, to construct a formal model of the bug automatically. *kismet* discharges the formal conditions to an SMT solver to check satisfiability. The Kconfig language allows developers to define constraints using symbolic Boolean formulas. Since our goal is to analyze all solutions to these constraints simultaneously, the analysis problem is at least as hard as Boolean satisfiability in general. The main challenge to designing *kismet* is ensuring scalability while preserving enough precision to identify the exact constructs causing the unmet dependency alarm.

4.1 Identifying Select Constructs

The first challenge for *kismet* is to identify select constructs in the Kconfig specifications. Walking over each config construct syntactically, it records all pairs of options involved in a select operation. For instance, in Figure 1, *kismet* identifies the pair (TOUCHSCREEN_ADC, IIO_BUFFER_CB) which contains the *selector* and *selectee*, respectively. In order to verify whether the select is free from an unmet dependency bug, *kismet* needs to account for all of the dependencies that constrain both the selector and the selectee.

The schematic in Figure 5 highlights what conditions *kismet* uses from the Kconfig specification to construct a verification condition. The configuration option X (line 1) selects A (line 5) with the select construct on line 2. The select construct itself is constrained by some if dependency, captured by a logical formula D_X (line 2). Additionally, X has its own dependencies K_X controlling when it can be enabled (line 3). A's direct dependencies are D_A (line 6), while K_A (line 7) represents any prompt or default constraints. K_{other} represents the constraints from all other configuration options.

4.2 Modeling Unmet Dependency Bugs

X's select construct only causes an unmet dependency if the select overrides A's direct dependencies, i.e., when A's dependencies are unsatisfied. If we just consider the constraints of the selector and the selectee, the formula for unmet dependency is as follows:

$$\phi_{\text{unmet}} = X \wedge D_X \wedge K_X \quad (1)$$

$$\wedge A \wedge \neg(D_A \wedge K_A) \quad (2)$$

ϕ_{unmet} means the following: the *selector* option X is enabled and its select and other constraints $D_X \wedge K_X$ are met (subexpression 1); and the *selectee* option A is enabled while its dependencies $D_A \wedge K_A$ are *not* met (subexpression 2).

ϕ_{unmet} is an overapproximation, however, because it only accounts for the constraints from two configuration options, the selector and selectee. Constraints from other configuration options (K_{other}) can make ϕ_{unmet} unsatisfiable. Without accounting for those, we can expect more false positive alarms. A precise condition would contain these constraints as well:

$$\phi_{\text{unmet}}(\text{precise}) = \phi_{\text{unmet}} \wedge K_{\text{other}}$$

Optimizing the bug-finder. The Linux Kconfig specification has thousands of configuration options, so $\phi_{\text{unmet}}(\text{precise})$ is a substantially more expensive formula to solve; it has the constraints from thousands of configuration options instead of just the two in ϕ_{unmet} . To make solving more efficient, we use two techniques. First, if a selectee option has no direct dependencies, then an unmet dependency bug is not possible. Second, we first check the ϕ_{unmet} condition and only check $\phi_{\text{unmet}}(\text{precise})$ if the first check is satisfiable. This optimization is safe, because if ϕ_{unmet} is unsatisfiable, we know that $\phi_{\text{unmet}}(\text{precise})$ is also unsatisfiable, i.e. $\neg\phi_{\text{unmet}}$ entails $\neg\phi_{\text{unmet}}(\text{precise})$:

$\phi_{\text{unmet}} \rightarrow \phi_{\text{unmet}}$	tautology
$\phi_{\text{unmet}} \wedge K_{\text{other}} \rightarrow \phi_{\text{unmet}}$	strengthening
$\phi_{\text{unmet}}(\text{precise}) \rightarrow \phi_{\text{unmet}}$	substitution
$\neg\phi_{\text{unmet}} \rightarrow \neg\phi_{\text{unmet}}(\text{precise})$	contrapositive

This simple optimization has a substantial impact on precision and performance as we show in the evaluation section.

4.3 Modeling Kconfig Semantics

Until now, we have described ϕ_{unmet} with placeholders for configuration option constraints. But interpreting these constraints as logical formulas requires modeling Kconfig's semantics. In this section we show how we methodically derive these from our formal semantics (Section 3).

Recall that Kconfig takes a configuration file as input and determines its validity according to the specifications. As with prior Kconfig feature modeling tools, we represent configuration options as symbolic Boolean options, collapsing tristate option's y and m to true. While this underapproximates tristate, it greatly reduces the space of possible configurations, improving solver performance. Similarly, we approximate non-Booleans with a finite range of options, as in prior work [30]. Less than 5% of options are non-Boolean in Linux Kconfig specifications.

To derive the model from Kconfig semantics, recall that our concrete semantics describes each case in which a Kconfig statement

describes a valid configuration given an input configuration file. For each Kconfig syntactic construct in the specification under analysis, our bug finder automatically constructs a symbolic formula ϕ_i corresponding to its valuation function from the semantics in Figure 3. The formula is the disjunction of each condition leading to a valid result. For instance, a config statement, defined by semantic rule S₂ in Figure 3, has four valid cases. kclause constructs ϕ_{config} as a disjunction of each of these case conditions, where C is the symbolic value of the option:

$$\begin{aligned} \phi_{\text{config}} = & (C \wedge \phi_{\text{reverse}}) \\ & \vee (\phi_{\text{depends}} \wedge \phi_{\text{prompt}} \wedge \neg\phi_{\text{reverse}}) \\ & \vee ((C \wedge \phi_{\text{default}} \vee \neg C \wedge \neg\phi_{\text{default}}) \\ & \quad \wedge \phi_{\text{depends}} \wedge \neg\phi_{\text{prompt}} \wedge \neg\phi_{\text{reverse}}) \\ & \vee (\neg C \wedge \neg\phi_{\text{depends}} \wedge \neg\phi_{\text{reverse}}) \end{aligned}$$

Each of the four disjunctive terms corresponds to each of the four valid conditions from S₂. Accesses to the concrete configuration option value $\sigma(\text{sym})$ are replaced with a symbolic Boolean value C. Calls to other valuation functions are replaced with the symbolic formulas for that syntax, e.g., ϕ_{depends} for the depends on construct.

For the configuration option IIO_BUFFER_CB in Figure 1, ϕ_{config} is constructed from the following symbolic formulas:

$$\begin{aligned} C &= \text{IIO_BUFFER_CB} \\ \phi_{\text{reverse}} &= \text{TOUCHSCREEN_ADC} \wedge \text{IIO} \wedge \text{INPUT_TOUCHSCREEN} \\ \phi_{\text{depends}} &= \text{IIO} \wedge \text{IIO_BUFFER} \\ \phi_{\text{prompt}} &= \text{true} \\ \phi_{\text{default}} &= \text{false} \end{aligned}$$

ϕ_{reverse} and ϕ_{direct} are the reverse and direct dependencies respectively. The option is always visible since it has an unconditional prompt (ϕ_{prompt}), and the default value is false since it has no default specification (ϕ_{default}).

Substituting the symbolic formulas into ϕ_{config} and simplifying the disjunctive terms, we get the following:

$$(\text{IIO_BUFFER_CB} \quad (3)$$

$$\wedge \text{TOUCHSCREEN_ADC} \wedge \text{IIO} \wedge \text{INPUT_TOUCHSCREEN})$$

$$\vee (\text{IIO} \wedge \text{IIO_BUFFER} \quad (4)$$

$$\wedge \neg(\text{TOUCHSCREEN_ADC} \wedge \text{IIO} \wedge \text{INPUT_TOUCHSCREEN}))$$

$$\vee (\text{false}) \quad (5)$$

$$\vee (\neg\text{IIO_BUFFER_CCB} \wedge \neg(\text{IIO} \wedge \text{IIO_BUFFER}) \quad (6)$$

$$\wedge \neg(\text{TOUCHSCREEN_ADC} \wedge \text{IIO} \wedge \text{INPUT_TOUCHSCREEN}))$$

In summary, this formula means that IIO_BUFFER_CB is legal to enable if its reverse dependency holds (subexpression 3), is legal to either enable or disable if its direct dependency holds (subexpression 4), never takes a default value (subexpression 5), and can otherwise only be disabled when its direct and reverse dependencies do not hold (subexpression 6). The rest of the symbolic evaluator's valuation functions are similarly derived from the formal semantics and can be found with the openly-archived formal semantics [40].

The benefit of this approach is that it removes guesswork from designing Kconfig analysis tools. Instead, tool writers can rely on a common semantics to mechanically derive an analysis for whatever abstraction they would like to use for analysis, tailoring the choice of formalism for configuration options based on their specific application. We demonstrate just one possible set of choices for deriving




Figure 6: The components of the infrastructure and how they work together for unmet dependency bug-finding.

the Kconfig analysis. Moreover, future extensions to Kconfig by developers can be captured by updates to the semantics, easing adoption for Kconfig tools that mechanically derive their analyses from the semantics.

5 IMPLEMENTATION

The analysis framework is implemented in about two thousand source lines of Python, and about one thousand source lines of C. It consists of four components, shown in Figure 6. The `kextract` tool wraps the parser from the Linux implementation of Kconfig [54] with a C extension to desugar the Kconfig specification into a desugared version of the Kconfig language. The `kclause` tool, written in Python, reads in desugared Kconfig and constructs logical formulas for each configuration options’ constraints, outputting them in the SMTLIB2 [9] format. `kismet`, also written in Python, finds each select construct from the `kextract` output and uses the logical models from `kclause` to generate the unmet dependency condition for each construct. `kismet` finally passes this condition in SMTLIB2 format to the `klocalizer` tool, which uses the Z3 SMT solver [13] to check for the satisfiability of the bug condition. For satisfiable conditions, `klocalizer` can also generate solutions to the condition in the Linux `.config` file format, which we use to test the solution against the actual Kconfig implementation. All source code is available online as free and open-source software¹ as well as in an openly-archived artifact [38].

6 EXPERIMENTAL EVALUATION

We evaluate our bug finding approach for precision, performance, and impact on real-world code.

¹<https://github.com/paulgazz/kmax>

6.1 Experimental Setup

We use Kconfig specifications from a recent version (v5.4.4) of the Linux kernel source code² as the target of our study. With over 140,000 lines of specifications and over 15,000 configuration options, Linux represents, to our knowledge, the largest user of Kconfig.

The Linux kernel not only provides a large Kconfig specification, but multiple ones as well, due to its support for multiple hardware platforms. Each of its 28 architecture families³ has its own Kconfig specification, effectively providing 28 separate Kconfig specifications to use for evaluation. Because of the hardware abstraction layer, however, these architectures share at least some portion of the codebase in common, and therefore also share a large portion of their Kconfig specifications; about 100,000 lines, two-thirds, are architecture-independent. Each architecture has between 10,014 and 12,744 select constructs for a total of 289,202. Deduplicating these, there are 17,006 unique select constructs, although the constraints due to architecture-specific Kconfig files may differ. Due to this sharing, we not only report results for each architecture’s Kconfig specifications but also the aggregate and deduplicated alarms across architectures.

All experiments were executed on a server with an AMD EPYC 7401 24-Core Processor with 512GB of RAM running Ubuntu 18.04, where we measured performance using the UNIX `time` utility. Since this machine allows for high parallelism, we ran the experiments for the 28 architectures’ Kconfig specifications in parallel on separate copies of the Linux kernel source code. Replication scripts are available with the source code repository¹.

6.2 Data Availability

All experimental data are available as archived open data [39].

6.3 Research Questions

Our evaluation seeks to answer the following research questions:

RQ1 (Precision) How precise is our analysis when finding unmet dependencies? To measure bug-finding effectiveness, we run our tool on all 28 Linux Kconfig specifications and collect the alarms reported. We also automatically validate whether the alarms are true positives by generating and building test cases automatically. We expect that, if our semantics reflect real Kconfig behavior, that our symbolic model of unmet dependencies and Kconfig behavior should yield high precision, i.e., few false positives.

RQ2 (Performance) How fast is bug-finding? We record the running time of our bug-finder when applied to all 28 Linux Kconfig specifications, i.e., the experiment from RQ1. We report the distribution of running times per architecture, the aggregate time, as well as the breakdown between desugaring, generating bug conditions, and solving. We expect that our design choices and optimization will yield a fast enough analysis to make running `kismet` feasible for developers to use regularly.

RQ3 (Impact) How useful are the resulting alarms to developers? We evaluate the impact of our bug-finding approach by manually submitting some reports and patches to the kernel

²<https://cdn.kernel.org/pub/linux/kernel/v5.x/linux-5.4.4.tar.xz>

³alpha, arc, arm, arm64, c6x, csky, h8300, hexagon, i386, ia64, m68k, microblaze, mips, nds32, nios2, openrisc, parisc, powerpc, riscv, s390, sh, sh64, sparcc, sparcc64, um, unicore32, x86_64, and xtensa

Table 1: kismet’s bug-finding results across all 28 architecture Kconfig specifications.

Metric	Percentiles				
	Max	75th	50th	25th	Min
Constructs	12,744	10,386	10,108	10,044	10,014
Alarms raised	53.00	31.25	25.00	22.75	10.00
Precision	100%	100%	100%	100%	100%

maintainers. We expect that, if the resulting alarms are correct and provide value to the kernel maintainers, they will confirm the reports and accept our patches.

RQ4 (Comparison) How does our approach compare to random configuration testing? To our knowledge, no related tool for finding unmet dependencies in Kconfig exists. To provide a baseline time to search for bugs, we use random configuration testing with Kconfig’s built-in `randconfig` tool. We compare the bugs found, given the same amount of time as kismet and also allow `randconfig` generation to run for several days. We expect that our static approach will perform better, given the enormity of the search space of configurations, but we also expect to find new bugs missed by kismet’s underapproximation of non-Booleans.

6.4 RQ1: Precision

We run kismet on each of the 28 architectures’ Kconfig specifications and collect the resulting alarms. kismet reports the pair of configuration options involved in the unmet dependency, i.e., the *selector* and the *selectee*. Finally, we validate whether the alarm is a true positive by generating a test case. This works by querying the Z3 SMT solver for a satisfying solution to ϕ_{unmet} (precise), the bug’s logical formula, then converting the solution into the Linux `.config` configuration file format.

Table 1 summarizes the analysis results of our experiments. The rows list the number of constructs analyzed, the number of alarms raised by kismet, and the precision, i.e., the percent of all alarms that are true positives. The columns show the distribution of these metrics across the 28 architectures’ Kconfig specifications as percentiles. kismet checks between 10,014 and 12,744 select constructs for each architecture, finding between 10 and 53 alarms per Kconfig specification, for a total of 781 alarms over 289,202 constructs. All alarms are confirmed to be true positives by generating test cases that trigger the alarm, for a precision of 100%. While such high precision would be unusual for static analysis, the core fragment of Kconfig that we model requires no over-approximation that could lead to false positives. Since the ground truth number of bugs in real-world Linux Kconfig specifications is unknown, we do not compute recall, but we address false negatives in RQ4.

Although each architecture has its own Kconfig specification, they all share a large common set of Kconfig files. The consequence is that fixing a bug in one architecture’s Kconfig specification can fix it for several others. Deduplicating these bug yields 151 total alarms for unique select constructs across all architectures. In some cases, the same select construct was a true unmet dependency in one

Table 2: kismet’s bug-finding time in minutes for all 28 Kconfig specifications, broken down by each phase of analysis.

Analysis Phase	Time Percentiles (minutes)				
	Max	75th	50th	25th	Min
1. <code>kc1ause</code>	7.21	5.52	5.35	5.21	5.03
2. Syntax check	0.15	0.12	0.12	0.11	0.11
3. ϕ_{unmet}	2.35	2.00	1.94	1.88	1.62
4. ϕ_{unmet} (precise)	79.31	33.79	32.16	31.23	29.08
5. Confirmation	2.01	1.06	0.81	0.72	0.38
Total Time	90.21	42.12	40.30	39.41	37.13

architecture’s Kconfig specification but not others, which is possible because of architecture-specific constraints. In these cases, we counted the construct as a true alarm in the deduplicated set.

Summary: our approach is precise, yielding 100% precision on Linux’s very large, real-world Kconfig specification, and finds many new bugs: 781 true positive bugs or 151 if we deduplicate common constructs across architectures.

6.5 RQ2: Performance

To evaluate performance, we measure kismet’s running time, broken down by each phase of its analysis. Table 2 is the distribution of running times across each of the 28 architecture-specific Kconfig specifications. Each row is the phase of analysis, with the total time in the last row, while each column is percentiles in the distribution of running times.

kismet takes between 37 and 90 minutes on one Kconfig specification file, for a total of 20 hours in all, including the time spent generating a test case to automatically confirm true positives. We break down the timing into five phases: (1) *kc1ause* is the time spent modeling Kconfig constructs, which we perform at the beginning of analysis to cache the results. (2) *Syntax check* includes both identifying each select construct and the optimization that rules out selectees with no dependencies. As discussed in Section 4.2 on optimization, (3) ϕ_{unmet} is the time spent checking the imprecise bug formula, and (4) ϕ_{unmet} (precise) is the time spent checking the precise bug formula, if the imprecise one does not rule out the bug. (5) *Confirmation* is the time spent generating a test case for the bug and checking it against the actual Kconfig implementation; this is not part of the static analysis, per se, but it only takes a comparatively small amount of time.

In most cases, kismet takes less than hour for an architecture, making it fast enough for use on each commit of the Kconfig specification. The largest amount of time is spent on the precise formula check, which shows the importance of our optimization in avoiding making that check. Checking ϕ_{unmet} is fast: it takes less than an hour for hundreds of thousands of select constructs, albeit with low precision (less than 2%). 85% of the constructs are ruled out, however, reducing the time needed to solve the precise condition.

Summary: kismet is fast, taking between 37 and 90 minutes to analyze between 10,014 to 12,744 select constructs in a Kconfig specification, enabling frequent bug finding runs.

6.6 RQ3: Impact

We evaluate the impact of our bug-finder, and the semantics on which it is based, by reporting alarms to the kernel developers and submitting patches to the mainline Linux repository, specifically via the Linux kernel mailing list [50] and the kernel.org Bugzilla website [3]. Developer confirmation of bugs provides confidence in the utility of the alarms, beyond precision. Moreover, acceptance of patches by official maintainers reflects the beneficial impact of the results on this prevalent and frequently used codebase.

While our bug-finder is fully automated, submitting reports and patches is a manual process, requiring time to create them and communicate with human Linux maintainers. Moreover, maintainers may opt to not patch even true alarms, may not respond immediately, or may request different changes than what we proposed in the patch. Since the Kconfig specification gradually changes over time with the rest of the codebase, prior bugs may no longer occur, due to manual fixes, removal of options, etc. We believe it is feasible to use kismet in continuous integration, but we leave such infrastructure development as future work. For these reasons, we have not yet submitted all alarms; repairing all is an ongoing process, and we report the current state of the bug repairs in progress.

As of writing, we have submitted 38 reports or patches, 19 have been confirmed with the remainder pending, and 15 of our patches have already been committed to the Linux kernel codebase. Up-to-date information about the reporting and patching effort can be found in the source code repository⁴.

Knowing the effect of unmet dependencies on the kernel is difficult to measure. Such a configuration is not supposed to be feasible, and developers have been so far highly receptive to patches of unmet dependency bugs. While we do not know all the effects of an unmet dependency, one common result is a broken build, e.g., Figure 1, which is undesirable for any software product. We measured how often a broken build results from the bugs we found by attempting to build the generated .config from kismet and hand-checking the reason for the broken build. Build errors account for 68% of all tests. 29% of configuration files trigger build errors whose root cause is the unmet dependency bug from which the configuration file was generated. 27% fail due to bugs other than the one used to generate the test case. Since a build error halts the build process, we cannot easily determine whether the build would have encountered an error related to the unmet dependency, so we conservatively assume these are not caused by unmet dependencies.

Summary: *The bug finding results have resulted in 38 reports and 15 committed patches to the Linux kernel so far, with further patch submission and discussion ongoing.*

6.7 RQ4: Comparison

While kismet is 100% precise for its fragment of the Kconfig semantics, its underapproximation of non-Boolean leaves it susceptible to false negatives. To gather a set of unmet dependency benchmarks that include bugs not findable by kismet, we use a built-in Kconfig utility for generating random configurations. Generating random configurations for over four days for each architecture in parallel (a combined time of more than three months), we generated over 11,000,000 configuration files, which raised 2,857,938

⁴https://github.com/paulgazz/kmax/blob/master/docs/bugs_found.md

Table 3: Percent of the bugs found by kismet compared to randconfig given both the same amount of time as and 135x more time than kismet.

Tool	Percentiles				
	Max	75th	50th	25th	Min
kismet	100.00%	100.00%	100.00%	100.00%	87.10%
randconfig					
Same time	62.86%	12.94%	6.80%	2.68%	0.00%
135x time	77.14%	22.55%	17.42%	10.54%	0.00%

unmet dependency alarms, yielding 175 unique unmet dependency bugs. Comparing these to kismet’s results, kismet adds 614 unique unmet dependencies not found in this random testing.

Since no other tools to our knowledge analyze unmet dependencies, we compare the performance of kismet against a random testing approach, to see whether there is a benefit in running time and bugs found to using kismet. Using the combined set of bugs from months of randconfig and kismet’s results, we compare the percent of bugs found given the same amount of time. Table 3 shows the results of this comparison of the percentage of bugs found from the benchmark set. The columns show the distribution of these percentages across all architectures’ Kconfig specifications. kismet finds 100% for almost all architectures, reflecting the fact that even after months of compute time, very few additional bugs were found by random testing compared with kismet. randconfig (row “Same time”), given the same amount of time that kismet took, finds on average only a small fraction of the set of bugs, 6.80%, with a maximum of only 62.86%. Even given several days to run (row “135x time”), randconfig still only finds a fraction of the benchmark bugs. In contrast, there were only eight bugs not found by kismet, leading to a worst-case of 87.10% benchmark coverage by kismet.

While our benchmark is not the ground truth of Linux’s complete set of bugs, which is not feasible to find by hand given the months of compute time to generate configuration files, it provides at least an estimate of the relative performance of kismet versus random testing. The results show the large performance benefit of using kismet compared to random testing. In the same amount of time, kismet finds many more bugs than random testing, providing a fast and precise complement to random testing that can be run regularly against new commits to the Kconfig specification.

Summary: *kismet finds many more true positives bugs in far less time than random testing, although there are also false negatives as expected by deliberate underapproximation.*

7 THREATS TO VALIDITY

Internal Threats. Our formal semantics needs to match the actual behavior of Kconfig, otherwise, any analyses based on it may yield incorrect results. We mitigated this using the Kconfig documentation, reviewing its actual C implementation, and collecting a Kconfig test suite. Moreover, the 100% precision of the bug-finder, validated with generated test cases and some developer confirmation, testifies to the accuracy of the semantics. kismet is deliberately underapproximate for non-Boolean options, however, so this part of the

semantics is not supported by the bug-finding results, but by the documentation, implementation, and test suite only.

External Threats. While Kconfig is used by several popular, low-level systems software (BusyBox, coreboot, etc), our evaluation only applies to Linux. Linux, however, is the largest user of Kconfig that we know of, and has multiple Kconfig specifications. We evaluate our bug-finder on one recent version of the Linux source code, but Kconfig specifications change gradually with each kernel version. Different versions may yield different numbers of alarms. We leave a long-term study of Kconfig bugs across versions and projects as future work. Our bug-finder currently checks for one kind of bug. The performance of the bug-finder could vary for different bug types or analysis tasks. Our work is specific to the Kconfig specification language, so we do not show applicability to other specification languages. Given the large time investment in creating and evaluating accurate formal semantics and a corresponding analysis infrastructure, we leave generalizing the approach to other specification languages as future work.

8 RELATED WORK

Modeling Kconfig specifications. There are several prior efforts that convert Kconfig to logical formulas for various applications. Zengler et al. and Walch et al. modeled Kconfig in the DIMACS SAT format with the goal of finding Kconfig language metrics, including the number of options, types, and mandatory configuration options [57, 61]. She et al. describe a formal semantics [43] and a tool called LVAT that converts Kconfig specifications to the DIMACS SAT solver format [10, 44, 45]. It was designed for collecting statistics about the Kconfig language such as the number of options, the hierarchy of dependencies, and other metrics [10], rather than for precise formal verification of configuration specifications. Tool development appears to have stopped for LVAT in 2013 [42]. The undertaker project has a tool to convert Kconfig's `dumpconf` output to the DIMACS SAT format for use in identifying dead code blocks in unprocessed C code [4, 53]. The `kconfigreader` tool converts the output of a Kconfig tool called `dumpconf`, which dumps each configuration options' constraint expressions, into the DIMACS SAT solver format [25, 30]. El-Sharkawy et al., describes an informal semantics of Kconfig, provides illustrative examples, and evaluates the limitations of other tools [15]. Fernandez et al. described informal semantics for Kconfig constructs that they identified as incorrectly supported in prior conversion tools [16]. They provide a set of example Kconfig constructs that illustrate these limitations, which we have incorporated into `kc` clause's test suite. Fernandez et al. also describe a new conversion tool that produces Binary Decision Diagrams but has not been evaluated on Linux Kconfig specifications.

Analyses of other configuration languages. Shambaugh et al. [41] perform formal verification of the Puppet deployment configuration language to detect non-deterministic system state updates and other undesirable system configurations. Weiss et al. [59] automate Puppet configuration repair using formal reasoning over a propositional model of the language. Anderson et al. [8] formally verify the SmartFrog infrastructure deployment language to prove properties such as termination of compilation, comparing multiple

implementations of SmartFrog compilers. Sotiropoulos et al. formally modeled the system call trace of the Puppet tool to find faults from ordering violations on resource usage [49]. Horton and Parnin infer system dependencies from Python code in order to generate Docker specification files [22]. They also inferred dependencies from Python code snippets to check if their package dependencies are out of date [23]. Bouchet et al. use formal verification to check for inadvertent public access to Amazon S3 instances [12]. Chenygyuan et al. mined frequently used dependencies between entities from deployment descriptors for Java-EE-platform-based applications to validate if a new deployment descriptor is violating mined dependencies [60]. Hanappi et al. formally modeled configuration scripts and resource usage to test if a system can recover from failures such as network outages and reach a stable state [21].

Studies on variability bugs. Some prior work extracted variability information from Makefiles and source code for finding bugs, dead code blocks, or inconsistencies between variability specification and implementation. [11, 14, 19, 36, 37, 48, 52]. Prior work also analyzed bugs or warnings raised from sampled configurations to classify them and understand how they are introduced [33, 34]. Similar analyses were performed on the bugs or vulnerabilities reported in the bug database or source commits [6, 7, 17, 35]. Others studied configuration sampling algorithms to find more variability bugs with fewer samples [29, 32, 46, 55].

9 CONCLUSION

We have introduced a new formal semantics and model checking infrastructure for analyzing Kconfig specification files and methodically derived a bug-finder, called `kismet`, for unmet dependencies, a common pitfall for Kconfig maintainers. Our results show that our bug-finder is precise, fast, and has resulted in patches to the mainline Linux kernel source code confirmed and accepted by maintainers. Future work includes continuing to repair all bugs found by `kismet`, applying it to ongoing kernel development and other software, and applying our analysis framework to other maintenance challenges.

We also plan to explore applying these model checking principles to other configuration specification languages to further improve the state of language tooling for software operations at large. As software operations are further automated, the languages used for configuring, building, and deploying software become an increasingly large component of the source code. These languages introduce new opportunities for less traditional software vulnerabilities, such as security misconfiguration. As our work demonstrates, these languages lend themselves to automated analysis, suggesting the future benefits of applying rigorous design and automated reasoning to software operations languages in general.

ACKNOWLEDGMENTS

We would like to thank the anonymous referees for their valuable comments and helpful suggestions, Julia Lawall for advice and input into the work, and Elaine Weyuker for her feedback. This work is supported by the National Science Foundation under CCF-1941816 and CCF-1840934.

REFERENCES

- [1] 2008. Boogie: An Intermediate Verification Language. <https://www.microsoft.com/en-us/research/project/boogie-an-intermediate-verification-language/>.
- [2] 2020. Kconfig macro language. <https://www.kernel.org/doc/html/latest/kbuild/kconfig-macro-language.html>, last accessed on 11/19/20.
- [3] 2020. Kernel.org Bugzilla page. <https://bugzilla.kernel.org/>, last accessed on 11/19/20.
- [4] 2020. Undertaker Project Page. <https://vamos.informatik.uni-erlangen.de/trac/undertaker>, last accessed on 11/19/20.
- [5] 2021. 0-Day Test Service. <https://01.org/lkp/documentation/0-day-test-service>.
- [6] Iago Abal, Claus Brabrand, and Andrzej Wasowski. 2014. 42 Variability Bugs in the Linux Kernel: A Qualitative Analysis. In *Proceedings of the 29th ACM/IEEE International Conference on Automated Software Engineering (Vasteras, Sweden) (ASE '14)*. Association for Computing Machinery, New York, NY, USA, 421–432. <https://doi.org/10.1145/2642937.2642990>
- [7] Iago Abal, Jean Melo, Ștefan Stănculescu, Claus Brabrand, Márcio Ribeiro, and Andrzej Wasowski. 2018. Variability Bugs in Highly Configurable Systems: A Qualitative Analysis. *ACM Trans. Softw. Eng. Methodol.* 26, 3, Article 10 (Jan. 2018), 34 pages. <https://doi.org/10.1145/3149119>
- [8] Paul Anderson and Herry Herry. 2016. A formal semantics for the SmartFrog configuration language. *Journal of Network and Systems Management* 24, 2 (2016), 309–345. <https://doi.org/10.1007/s10922-015-9351-y>
- [9] Clark Barrett, Aaron Stump, and Cesare Tinelli. 2010. The SMT-LIB Standard: Version 2.0. In *Proceedings of the 8th International Workshop on Satisfiability Modulo Theories (Edinburgh, UK)*. A. Gupta and D. Kroening (Eds.).
- [10] Thorsten Berger, Steven She, Rafael Lotufo, Andrzej Wasowski, and Krzysztof Czarnecki. 2013. A Study of Variability Models and Languages in the Systems Software Domain. *IEEE Transactions on Software Engineering* 39, 12 (2013), 1611–1640. <https://doi.org/10.1109/TSE.2013.34>
- [11] Eric Bodden, Társlis Tolédó, Márcio Ribeiro, Claus Brabrand, Paulo Borba, and Mira Mezini. 2013. SPL^{sup>LIFT}: Statically Analyzing Software Product Lines in Minutes Instead of Years. (2013), 355–364. <https://doi.org/10.1145/2491956.2491976>
- [12] Malik Bouchet, Byron Cook, Bryant Cutler, Anna Druzkina, Andrew Gacek, Liana Hadarean, Ranjit Jhala, Brad Marshall, Dan Peebles, Neha Rungta, Cole Schlesinger, Chriss Stephens, Carsten Varming, and Andy Warfield. 2020. Block Public Access: Trust Safety Verification of Access Control Policies. In *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (Virtual Event, USA) (ESEC/FSE 2020)*. Association for Computing Machinery, New York, NY, USA, 281–291. <https://doi.org/10.1145/3368089.3409728>
- [13] Leonardo De Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *Proceedings of the Theory and Practice of Software, 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (Budapest, Hungary) (TACAS'08/ETAPS'08)*. Springer-Verlag, Berlin, Heidelberg, 337–340.
- [14] Christian Dietrich, Reinhard Tartler, Wolfgang Schröder-Preikschat, and Daniel Lohmann. 2012. A Robust Approach for Variability Extraction from the Linux Build System. In *Proceedings of the 16th International Software Product Line Conference - Volume 1 (Salvador, Brazil) (SPLC '12)*. Association for Computing Machinery, New York, NY, USA, 21–30. <https://doi.org/10.1145/2362536.2362544>
- [15] Sascha El-Sharkawy, Adam Krafczyk, and Klaus Schmid. 2015. Analysing the Kconfig Semantics and its Analysis Tools. In *Proceedings of the 2015 ACM SIGPLAN International Conference on Generative Programming: Concepts and Experiences (Pittsburgh, PA, USA) (GPCE 2015)*. Association for Computing Machinery, New York, NY, USA, 45–54. <https://doi.org/10.1145/2814204.2814222>
- [16] David Fernandez-Amoros, Ruben Heradio, Christoph Mayr-Dorn, and Alexander Egyed. 2019. A Kconfig Translation to Logic with One-Way Validation System. In *Proceedings of the 23rd International Systems and Software Product Line Conference - Volume A (Paris, France) (SPLC '19)*. Association for Computing Machinery, New York, NY, USA, 303–308. <https://doi.org/10.1145/3336294.3336313>
- [17] Gabriel Ferreira, Momin Malik, Christian Kästner, Jürgen Pfeffer, and Sven Apel. 2016. Do #Ifdefs Influence the Occurrence of Vulnerabilities? An Empirical Study of the Linux Kernel. In *Proceedings of the 20th International Systems and Software Product Line Conference (Beijing, China) (SPLC '16)*. ACM, New York, NY, USA, 65–73. <https://doi.org/10.1145/2934466.2934467>
- [18] Alejandra Garrido and Ralph Johnson. 2005. Analyzing Multiple Configurations of a C Program. In *ICSM*. 379–388.
- [19] Paul Gazzillo. 2017. Kmax: Finding All Configurations of Kbuild Makefiles Statically. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering (Paderborn, Germany) (ESEC/FSE 2017)*. ACM, New York, NY, USA, 279–290. <https://doi.org/10.1145/3106237.3106283>
- [20] Paul Gazzillo and Robert Grimm. 2012. SuperC: Parsing All of C by Taming the Preprocessor. In *Proceedings of the 33rd ACM SIGPLAN Conference on Programming Language Design and Implementation (Beijing, China) (PLDI '12)*. ACM, New York, NY, USA, 323–334. <https://doi.org/10.1145/2254064.2254103>
- [21] Oliver Hanappi, Waldemar Hummer, and Schahram Dustdar. 2016. Asserting Reliable Convergence for Configuration Management Scripts. In *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications (Amsterdam, Netherlands) (OOPSLA 2016)*. Association for Computing Machinery, New York, NY, USA, 328–343. <https://doi.org/10.1145/2983990.2984000>
- [22] Eric Horton and Chris Parnin. 2019. DockerizeMe: Automatic Inference of Environment Dependencies for Python Code Snippets. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. 328–338. <https://doi.org/10.1109/ICSE.2019.00047>
- [23] Eric Horton and Chris Parnin. 2019. V2: Fast Detection of Configuration Drift in Python. In *2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. 477–488. <https://doi.org/10.1109/ASE.2019.00052>
- [24] Alexandru Florin Iosif-Lazar, Jean Melo, Aleksandar S. Dimovski, Claus Brabrand, and Andrzej Wasowski. 2017. Effective Analysis of C Programs by Rewriting Variability. *CoRR* (2017).
- [25] Christian Kästner. 2020. kconfigreader. <https://github.com/ckaestne/kconfigreader>, last accessed on 11/19/20.
- [26] Christian Kästner, Paolo G. Giarrusso, Tillmann Rendel, Sebastian Erdweg, Klaus Ostermann, and Thorsten Berger. 2011. Variability-Aware Parsing in the Presence of Lexical Macros and Conditional Compilation. In *Proceedings of the 2011 ACM International Conference on Object Oriented Programming Systems Languages and Applications (Portland, Oregon, USA) (OOPSLA '11)*. Association for Computing Machinery, New York, NY, USA, 805–824. <https://doi.org/10.1145/2048066.2048128>
- [27] Christian Kästner, Klaus Ostermann, and Sebastian Erdweg. 2012. A Variability-Aware Module System. In *Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages and Applications (Tucson, Arizona, USA) (OOPSLA '12)*. Association for Computing Machinery, New York, NY, USA, 773–792. <https://doi.org/10.1145/2384616.2384673>
- [28] The kernel development community. 2020. Kconfig Language. <https://www.kernel.org/doc/html/latest/kbuild/kconfig-language.html>, last accessed on 11/19/20.
- [29] Chang Hwan Peter Kim, Don S. Batory, and Sarfraz Khurshid. 2011. Reducing Combinatorics in Testing Product Lines. In *Proceedings of the Tenth International Conference on Aspect-Oriented Software Development (Porto de Galinhas, Brazil) (AOSD '11)*. Association for Computing Machinery, New York, NY, USA, 57–68. <https://doi.org/10.1145/1960275.1960284>
- [30] Christian Kästner. 2017. Differential Testing for Variational Analyses: Experience from Developing KConfigReader. arXiv:1706.09357 [cs.SE]
- [31] Jörg Liebig, Alexander von Rhein, Christian Kästner, Sven Apel, Jens Dörre, and Christian Lengauer. 2013. Scalable Analysis of Variable Software. In *Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering (Saint Petersburg, Russia) (ESEC/FSE 2013)*. Association for Computing Machinery, New York, NY, USA, 81–91. <https://doi.org/10.1145/2491411.2491437>
- [32] Flávio Medeiros, Christian Kästner, Márcio Ribeiro, Rohit Gheyi, and Sven Apel. 2016. A Comparison of 10 Sampling Algorithms for Configurable Systems. In *Proceedings of the 38th International Conference on Software Engineering (Austin, Texas) (ICSE '16)*. Association for Computing Machinery, New York, NY, USA, 643–654. <https://doi.org/10.1145/2884781.2884793>
- [33] Jean Melo, Elvis Flesborg, Claus Brabrand, and Andrzej Wasowski. 2016. A Quantitative Analysis of Variability Warnings in Linux. In *Proceedings of the Tenth International Workshop on Variability Modelling of Software-intensive Systems (Salvador, Brazil) (VaMoS '16)*. ACM, New York, NY, USA, 3–8. <https://doi.org/10.1145/2866614.2866615>
- [34] Austin Mordahl, Jeho Oh, Ugur Koc, Shiyi Wei, and Paul Gazzillo. 2019. An Empirical Study of Real-World Variability Bugs Detected by Variability-Oblivious Tools. In *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (Tallinn, Estonia) (ESEC/FSE 2019)*. Association for Computing Machinery, New York, NY, USA, 50–61. <https://doi.org/10.1145/3338906.3338967>
- [35] Raphael Muniz, Larissa Braz, Rohit Gheyi, Wilkerson Andrade, Balduino Fonseca, and Márcio Ribeiro. 2018. A Qualitative Analysis of Variability Weaknesses in Configurable Systems with #Ifdefs. In *Proceedings of the 12th International Workshop on Variability Modelling of Software-Intensive Systems (Madrid, Spain) (VAMOS 2018)*. ACM, New York, NY, USA, 51–58. <https://doi.org/10.1145/3168365.3168382>
- [36] Sarah Nadi, Thorsten Berger, Christian Kästner, and Krzysztof Czarnecki. 2015. Where Do Configuration Constraints Stem From? An Extraction Approach and an Empirical Study. *IEEE Transactions on Software Engineering* 41, 8 (2015), 820–841. <https://doi.org/10.1109/TSE.2015.2415793>
- [37] Sarah Nadi and Ric Holt. 2012. Mining Kbuild to Detect Variability Anomalies in Linux. In *Proceedings of the 2012 16th European Conference on Software Maintenance and Reengineering (CSMR '12)*. IEEE Computer Society, USA, 107–116. <https://doi.org/10.1109/CSMR.2012.21>
- [38] Jeho Oh, Necip Fazıl Yildiran, Julian Braha, and Paul Gazzillo. 2021. Artifact from "Finding Broken Linux Configuration Specifications by Statically Analyzing the Kconfig Language". <https://doi.org/10.5281/zenodo.4885001>
- [39] Jeho Oh, Necip Fazıl Yildiran, Julian Braha, and Paul Gazzillo. 2021. Experimental data from "Finding Broken Linux Configuration Specifications by Statically Analyzing the Kconfig Language". <https://doi.org/10.5281/zenodo.4563310>

- [40] Jeho Oh, Necip Fazıl Yildiran, Julian Braha, and Paul Gazzillo. 2021. Formal Semantics of Kconfig for "Finding Broken Linux Configuration Specifications by Statically Analyzing the Kconfig Language". <https://doi.org/10.5281/zenodo.4950763>
- [41] Rian Shambaugh, Aaron Weiss, and Arjun Guha. 2016. Rehearsal: A Configuration Verification Tool for Puppet. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Santa Barbara, CA, USA) (PLDI '16). Association for Computing Machinery, New York, NY, USA, 416–430. <https://doi.org/10.1145/2908080.2908083>
- [42] Steven She. 2013. LVAT Archive. <https://code.google.com/archive/p/linux-variability-analysis-tools/>, last accessed on 11/19/20.
- [43] Steven She and Thorsten Berger. 2010. Formal semantics of the Kconfig language. *Technical note, University of Waterloo* 24 (2010).
- [44] Steven She, Rafael Lotufo, Thorsten Berger, Andrzej Wąsowski, and Krzysztof Czarnecki. 2011. Reverse Engineering Feature Models. In *Proceedings of the 33rd International Conference on Software Engineering* (Waikiki, Honolulu, HI, USA) (ICSE '11). Association for Computing Machinery, New York, NY, USA, 461–470. <https://doi.org/10.1145/1985793.1985856>
- [45] She, Steven. 2013. *Feature Model Synthesis*. Ph.D. Dissertation. <http://hdl.handle.net/10012/7834>
- [46] Jiangfan Shi, Myra B. Cohen, and Matthew B. Dwyer. 2012. Integration Testing of Software Product Lines Using Compositional Symbolic Execution. In *Proceedings of the 15th International Conference on Fundamental Approaches to Software Engineering* (Tallinn, Estonia) (FASE '12). Springer-Verlag, Berlin, Heidelberg, 270–284. https://doi.org/10.1007/978-3-642-28872-2_19
- [47] J. Sincero, H. Schirmeier, W. Schröder-Preikschat, and O. Spinczyk. 2007. Is the linux kernel a software product line?. In *Proceedings of the International Workshop on Open Source Software and Product Lines* (Kyoto, Japan) (SPLC-OSSPL). 134–140.
- [48] Julio Sincero, Reinhard Tartler, Daniel Lohmann, and Wolfgang Schröder-Preikschat. 2010. Efficient Extraction and Analysis of Preprocessor-Based Variability. In *Proceedings of the Ninth International Conference on Generative Programming and Component Engineering* (Eindhoven, The Netherlands) (GPCE '10). Association for Computing Machinery, New York, NY, USA, 33–42. <https://doi.org/10.1145/1868294.1868300>
- [49] Thodoris Sotiropoulos, Dimitris Mitropoulos, and Diomidis Spinellis. 2020. Practical Fault Detection in Puppet Programs. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering* (Seoul, South Korea) (ICSE '20). Association for Computing Machinery, New York, NY, USA, 26–37. <https://doi.org/10.1145/3377811.3380384>
- [50] Jasper Spaans. 2020. Linux Kernel Mailing List. <https://lkml.org/>, last accessed on 11/19/20.
- [51] Reinhard Tartler, Christian Dietrich, Julio Sincero, Wolfgang Schröder-Preikschat, and Daniel Lohmann. 2014. Static Analysis of Variability in System Software: The 90,000 #ifdefs Issue. In *Proceedings of the 2014 USENIX Conference on USENIX Annual Technical Conference* (Philadelphia, PA) (USENIX ATC '14). USENIX Association, USA, 421–432.
- [52] Reinhard Tartler, Daniel Lohmann, Julio Sincero, and Wolfgang Schröder-Preikschat. 2011. Feature Consistency in Compile-Time-Configurable System Software: Facing the Linux 10,000 Feature Problem. In *Proceedings of the Sixth Conference on Computer Systems* (Salzburg, Austria) (EuroSys '11). Association for Computing Machinery, New York, NY, USA, 47–60. <https://doi.org/10.1145/1966445.1966451>
- [53] Reinhard Tartler, Julio Sincero, Christian Dietrich, Wolfgang Schröder-Preikschat, and Daniel Lohmann. 2012. Revealing and repairing configuration inconsistencies in large-scale system software. *International Journal on Software Tools for Technology Transfer* 14, 5 (2012), 531–551.
- [54] Linux Tovalds. 2020. Linux Kconfig Source Code. <https://github.com/torvalds/linux/tree/master/scripts/kconfig>, last accessed on 11/19/20.
- [55] Mahsa Varshosaz, Mustafa Al-Hajjaji, Thomas Thüm, Tobias Runge, Mohammad Reza Mousavi, and Ina Schaefer. 2018. A classification of product sampling for software product lines. In *Proceedings of the 22nd International Systems and Software Product Line Conference-Volume 1*. 1–13.
- [56] Alexander von Rhein, Jörg Liebig, Andreas Janker, Christian Kästner, and Sven Apel. 2018. Variability-Aware Static Analysis at Scale: An Empirical Study. *ACM Transactions on Software Engineering and Methodology* 27, 4 (2018), Article No. 18. <https://doi.org/10.1145/3280986>
- [57] Martin Walch, Rouven Walter, and Wolfgang Küchlin. 2015. Formal analysis of the Linux kernel configuration with SAT solving. In *Configuration Workshop*. 131–138.
- [58] Eric Walkingshaw, Christian Kästner, Martin Erwig, Sven Apel, and Eric Bodden. 2014. Variational Data Structures: Exploring Tradeoffs in Computing with Variability. In *Proceedings of the 2014 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming & Software* (Portland, Oregon, USA) (Onward! 2014). Association for Computing Machinery, New York, NY, USA, 213–226. <https://doi.org/10.1145/2661136.2661143>
- [59] Aaron Weiss, Arjun Guha, and Yuriy Brun. 2017. Tortoise: Interactive System Configuration Repair. In *Proceedings of the 32nd IEEE/ACM International Conference on Automated Software Engineering* (Urbana-Champaign, IL, USA) (ASE 2017). IEEE Press, 625–636.
- [60] Chengyuan Wen, Yaxuan Zhang, Xiao He, and Na Meng. 2020. Inferring and Applying Def-Use like Configuration Couplings in Deployment Descriptors. In *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering* (Virtual Event, Australia) (ASE '20). Association for Computing Machinery, New York, NY, USA, 672–683. <https://doi.org/10.1145/3324884.3416577>
- [61] Christoph Zengler and Wolfgang Küchlin. 2010. Encoding the Linux kernel configuration in propositional logic. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI 2010) Workshop on Configuration*, Vol. 2010. 51–56.