ConfigFix: Interactive Configuration Conflict Resolution for the Linux Kernel

Experiences from a Decade of Reverse-Engineering the Semantics from
the Linux Kernel Configurator and Realizing Intelligent Configuration Support

Patrick Franz*, Thorsten Berger[†*], Ibrahim Fayaz[‡], Sarah Nadi[§], Evgeny Groshev*

*Chalmers | University of Gothenburg [†]Ruhr University Bochum [‡]VecScan AB (Vector Sweden) [§]University of Alberta

Abstract—Highly configurable systems are highly complex systems, with the Linux kernel arguably being one of the most well-known ones. Since 2007, it has been a frequent target of the research community, conducting empirical studies and building dedicated methods and tools for analyzing, configuring, testing, optimizing, and maintaining the kernel in the light of its vast configuration space. However, despite a large body of work, mainly bug fixes that were the result of such research made it back into the kernel's source tree. Unfortunately, Linux users still struggle with kernel configuration and resolving configuration conflicts, since the kernel largely lacks automated support. Additionally, there are technical and community requirements for supporting automated conflict resolution in the kernel, such as, for example, using a pure C-based solution that uses only compatible third-party libraries (if any).

With the aim of contributing back to the Linux community, we present CONFIGFIX, a tooling that we integrated with the kernel configurator, that is purely implemented in C, and that is finally a working solution able to produce fixes for configuration conflicts. In this experience report, we describe our experiences ranging over a decade of building upon the large body of work from research on the Linux kernel configuration mechanisms as well as how we designed and realized CONFIGFIX while adhering to the Linux kernel's community requirements and standards. While CONFIGFIX helps Linux kernel users obtaining their desired configuration, the sound semantic abstraction we implement provides the basis for many of the above techniques supporting kernel configuration, helping researchers and kernel developers.

Index Terms—software configuration, semantic abstraction, conflict resolution, Linux kernel

I. INTRODUCTION

The Linux kernel is the world's largest software development project [19] by the number of its contributors. Being highly versatile, the kernel operates in a diversity of environments, ranging from Android phones to large supercomputer clusters. As such, it is not only a highly successful operating-system kernel, but also a *highly configurable system* [72]—nowadays boasting 28 million lines of code [49] and over 15,000 configuration options (a.k.a., *features* [9], [61]). To this end, the kernel relies on mechanisms known from the fields of software product lines [6], [17], model-driven engineering [24], and software configuration [11]. Specifically, the Linux kernel includes a configurable build system [14], preprocessor-enabled variation points, a model-based representation of configuration options and their constraints (a.k.a., *variability model*) [16], [55], and an interactive configurator tool [72].

Being completely open source, with a vast evolution history available, researchers have studied many different aspects of it, including software evolution [5], [36], [61], [80] and software maintenance [3], [35], [37], [77] aspects, as well as its configuration mechanisms—the focus of this paper.

Studies of the Linux kernel's configuration mechanisms started back in 2007 [71], [72], [76], followed in 2010 by our and other researchers' studies of its variability modeling language Kconfig and its variability model [15], [16], [69]. Examples include the evolution of this model [52], the co-evolution and consistency of variation points [41], [56], [57], [61], [62], [75], as well as the synthesis of variability models from code [54], [55], [70]. Despite all the above research efforts related to configuration, users of the kernel did not benefit directly yet. They still struggle with creating their desired configuration [34], given the huge configuration space and intricate constraints among features. Beyond a simple and very limited support for choice propagation, the configurator does not offer any intelligent support for resolving configuration conflicts—for instance, enabling a feature requires transitively changing many other features. As such, achieving the desired configuration can be laborious and error-prone, which is unfortunate given all the work in the research community, which never made it back into the kernel.

To the best of our knowledge, only few contributions originating from research made it back into the kernel, and mainly in the form of bug fixes [56], [60], [75]. While there are tools such as Coccinelle [50] and Undertaker [74] that become known and often individually adopted among kernel developers, none were formally integrated into the kernel codebase or are listed as a necessary kernel tool. In 2015, the situation was about to change with the Kconfig-sat initiative [42], where kernel developers recognized the need and got in touch with researchers working on kernel configuration studies, including us. Given our experience with studying the Linux kernel's configuration information and developing tools to analyze it, we decided it was time to give back to the Linux community and try to practically integrate these techniques into the official Linux kernel configurator. In fact, implementing a sound translation of the variability model to propositional logics (which is a semantic abstraction) given the expressiveness and intricate semantics of the the Kconfig language (explained shortly), has long been an open problem, with multiple translations
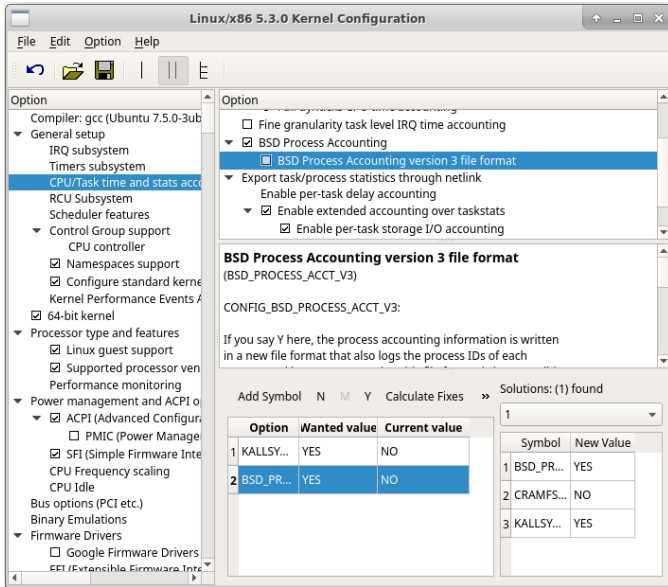
Fig. 1. The kernel configurator extended with CONFIGFIX

proposed [1], [40], [68], each with its own shortcomings [28].

We describe experiences covering more than a decade of different efforts on reverse-engineering the formal semantics from the Linux kernel configurator and implementing sound semantic abstractions. These are the prerequisite for many techniques and use cases supporting software configuration. After providing an overview on the Linux kernel's configuration tooling, research efforts on it, and their historical perspective, we introduce our tool CONFIGFIX, which offers intelligent configuration support for the Linux kernel configurator and realizes an evaluated, pure-C implementation of a semantic abstraction and a technique to resolve configuration conflicts. CONFIGFIX takes a current configuration (i.e., a selection of features with concrete values), a set of configuration constraints declared in the kernel's variability model, and a set of configuration options whose values the user wants to change, in order to calculate a set of *fixes* to reach the desired configuration. We built on our and others' prior work in the field to realize a translation of the kernel's variability model into propositional logics, to implement a configuration-conflict resolution algorithm relying on SAT solving, and to integrate both into the graphical configurator tool *xconfig*. Our work led to finally obtaining a viable solution integratable in the kernel's source tree. CONFIGFIX is freely available [2], together with details about its evaluation.

II. SOFTWARE CONFIGURATION AND THE LINUX KERNEL

We briefly introduce the field of software configuration, the Linux kernel's configuration facilities, and the notion of configuration conflict.

A. Software Configuration

Software configuration is concerned with methods and tools to configure software, originally stemming from the field of product configuration, a subfield of AI [33], [63]. The challenge is to obtain a configuration that meets end-user requirements, considering all constraints among the configuration options

(i.e., *features*). The configuration process is typically supported with an interactive configurator tool, offering support for propagating choices and resolving configuration conflicts. Software configurators [7], [8] have been studied in many domains [12], [17], [81], including configurable systems software (e.g., Linux kernel, eCos embedded operating system or 3D printer firmware [13], [16], [47], [68], [71]), automotive [31], avionics [51], [66], and telecommunication systems [73], embedded and safety-critical software [10], [45], [65], [78], as well as web-based configuration [4].

B. The Linux Kernel and its Configurator

The Linux kernel's configurability aims at customizing the kernel beyond its core functionality of CPU & memory management towards many different hardware architectures (ranging from embedded devices to supercomputer architectures) and including optional functionality, such as device or filesystem drivers. Currently, over 15,000 configuration options (henceforth called *features* [9], [61]), which come with intricate constraints among them, control variation points in C source files using conditional compilation directives (e.g., #if) of the C preprocessor, as well as they control the inclusion of individual files in the build process. In addition to this static mechanism, many features also control *loadable kernel modules* (e.g., network or USB drivers) that can be loaded dynamically at runtime.

Users configure the kernel interactively via its configurator, which exists in three variants. Figure 1 shows the graphical configurator *xconfig*. The other two variants target shell users. All features come with default values, and users can then assign values to the individual features according to their types and constraints, establishing a configuration. The features, their organization in a hierarchy, and their constraints are declared in files using the Kconfig language, which are input to the configurator.

Xconfig supports basic validity checking of configuration choices (to prevent some constraint violation) as well as simple imperative choice propagation. The latter, given the absence of an intelligent reasoner, needs to be encoded with a dedicated imperative mechanism in Kconfig, which is error-prone and, given its imperative nature, cannot be used to resolve configuration conflicts. In contrast, various open-source (e.g., FeatureIDE [53], Dopler [27], eCos' configtool [13]) and commercial configurators (e.g., pure::variants [18], Gears [46]), come with a reasoner.

C. The Kconfig Language

At the core of the Linux kernel configurator is the language Kconfig—a domain-specific language for variability modeling. Originally created for the Linux kernel, it has since been adopted by at least ten other open-source projects, such as BusyBox [16]. A core challenge in the community was obtaining a sound logical representation of the main semantics of Kconfig as a prerequisite to develop analysis and configuration techniques. However, as we will illustrate, Kconfig is surprisingly expressive with exceptionally intricate semantics.

Language Concepts. Kconfig comes with a textual syntax and concepts known from feature modeling (a popular kind of variability modeling language) [26], [39], [58]: a *hierarchy of features*, different *feature types*, *feature groups* (e.g., OR, XOR or MUTEX groups), and cross-tree constraints [16]. A feature model describes the set of all possible configurations as its main semantics. Since feature modeling languages are typically limited to Boolean features and propositional constraints, they can easily be converted into propositional logics by implementing their semantics.

Kconfig's syntax and semantics go well beyond feature modeling. For scaling the variability model and configuration process, Kconfig incorporates concepts such as visibility conditions (to conditionally show whole subtrees), modularization concepts, derived defaults / derived features, hierarchy manipulation), and an expressive constraint language including comparison, arithmetic, and String operators. Interestingly, features can also inherit constraints of their parents in non-transparent ways. Furthermore, Kconfig has a domain-specific vocabulary (main keywords) that fosters comprehension among Linux developers.

These concepts substantially complicate Kconfig's syntax and semantics. In fact, many intricate semantic interactions between different language elements exist—most notably between seven (sic!) language constructs to express constraints (`prompt`, `default`, `depends on`, `select`, `imply`, `visible if`, and `range`). For instance, a default value becomes a constraint when the feature is not visible, as determined by other constraints. For further details, we refer to Kconfig's official documentation [85] and our prior work [16], [68], which includes seven pages of reverse-engineered denotational semantics for Kconfig. Finally, Kconfig is continuously extended with language constructs, such as recently with the statement `imply`[1] as a special case of the `select` statement used for imperative choice propagation.

Imperative Choice Propagation

A frequent issue driving the complexity of the Kconfig semantics is that the developers incorporated imperative choice propagation—mainly through the `select` statement, which interacts with other Kconfig elements in intricate ways. This issue frequently complicated our and others' efforts realizing a sound semantic abstraction. It would have been better to keep the language cleaner and separately implement choice propagation via a reasoner, as CONFIGFIX does, but now accounting for the imperative choice propagation.

Features can be of different types: `bool`, `tristate`, `string`, `hex`, and `int`. Tristate features are used to control the binding mode of features and can have three values: **y** (yes, compile feature into kernel), **n** (no, do not compile feature) or **m** (mod, compile as loadable kernel module). Kconfig offers this feature type together with three-state logics that follows Kleene's rules for three-state logics [44]. Intuitively, the value of a tristate feature is encoded as the number 0, 1 or 2. The

logical operators are then defined over numbers: && returns the minimum and || the maximum of the two operands, and ! returns 2 minus the operand.

In our experience, the complexity of Kconfig is a result of the design of its configurator tooling. Instead of performing expensive language engineering [23], [25], [29], [48], [64] and adopting a configurator that comes with more intelligent reasoning capabilities [7], we learned that the community prefers transparent and easily scriptable solutions as opposed to heavy machinery, such as off-the-shelf reasoners that are difficult to understand. We learned this preference from the discussion on the kernel mailing list preceding the introduction of Kconfig and its tooling. An alternative candidate was a a configurator tool and language with built-in conflict-resolution support, which the community explicitly decided against, expressing these reasons.

Kconfig Language and Configurator Design

Kconfig is a popular language, but surprisingly expressive, coming with intricate syntax and semantics. We learned that the kernel community preferred the script-style xconfig and Kconfig over more systematically engineered tooling, mainly to be able to fully control and evolve the tooling.

Language Semantics and Abstractions. Motivated by the prospect of interesting empirical insights and being able to evaluate configuration- and variability-related research prototypes, researchers started looking into Kconfig and its tooling back in 2007 [71], [72], [76]. This was followed up with holistic studies of Kconfig, its tooling, and its models in 2010 [15], [16], [69]. For instance, we were the first to formally describe the semantics of Kconfig, which allowed its translation into propositional logic both by ourselves and other researchers [68], which resulted in the first translation tool called LVAT [67] (Linux Variability Analysis Tools), a tool suite written in Scala for analyzing Kconfig models. Around the same time, Zengler and Küchlin also provided a translation of Kconfig into propositional logics [28], [84]. Furthermore, the tool Undertaker [1] analyzes `#ifdef` code in the Linux kernel to identify dead code blocks—variation points whose constraints conflict with the variability model. To this end, Undertaker came with a translation into propositional logics. As part of the TypeChef infrastructure [41], Kästner implemented Kconfigreader [40] in Scala. Most recently, Fernandez-Amoros et al. also proposed yet another translation of Kconfig into propositional logic [30] which, however, omits its three-state logic and, therefore, a large part of the semantics.

A translation into an SMT representation was created by Xiong et al. [83], who presented RangeFix—a technique to generate configuration conflict fixes specifically for non-propositional configuration spaces. For non-propositional features, it provides ranges to which the value needs to be changed by the user to resolve a conflict. Our fix generation is conceptually based on RangeFix, but through simplifications since we do not need all of its computation steps.

A translation not originating from researchers exist as

---

[1] https://gitlab.freedesktop.org/panfrost/linux/commit/237e3ad0f195d8fd34f1299e45f04793832a16fc

well. As part of a Google Summer of Code project, Vegard Nossum contributed Satconfig [59], which comes with a translation implemented in C and allows reasoning via the SAT solver PicoSAT [22], for instance, completing a configuration based on an initial, partial configuration. We investigated the translation of Satconfig as well [38], but found shortcomings in the handling of tristate features, leading to incorrect fixes; there was also limited documentation for the translation and implementation. Despite the limitations, Satconfig was fast, indicating that a C-based solution can be scaled to the size of the kernel's varibility model. It also showed the advantages of running a SAT solver directly in the configurator tool, as well as the feasibility of implementing the translation into propositional logics in C. Its design inspired our data structures.

With the exception of one tool [30], all produce propositional formulas in conjunctive normal form (CNF), a prerequisite for SAT solvers, and typically apply a Tseitin transformation [79], which introduces auxiliary variables to avoid formula explosion.

A systematic comparison by ElSharkawy et al. [28] of LVAT [67], Undertaker [1], and KconfigReader [40] showed that all three tools had shortcomings in their sound encoding and abstraction of the Kconfig semantics. Notably, they found that KconfigReader could correctly handle the majority of the current Kconfig semantics, which steered our decision to re-implement KconfigReader's translation in C with some deviations (explained shortly).

> **Kconfig Semantic Abstraction**
>
> *Over the last decade, multiple researchers and one practitioner implemented propositional abstractions for Kconfig— to provide the basis for SAT-based reasoning, system analysis techniques, and guiding users configuring the kernel. None of these abstractions was fully sound and complete, given the complexity of the Kconfig language.*

The Kconfig-SAT Initiative. In October 2015, the Kernel developer Luis R. Rodriguez contacted some researchers including us who have worked on SAT-based configuration support for the Linux kernel and Linux package management. After discussing configuration issues related to Kconfig, and after being made aware of our research, he launched the Kconfig-SAT initiative, among others, with a project wiki page [42] and a mailing list [32]. Kconfig-SAT is now also described in the Kconfig documentation [43].

Among the kernel community, the awareness for needing intelligent configuration support rose, despite some skepticism about SAT solvers in general by Linus Torvalds: *"The SAT solver will only hurt, because it will bring in all those irrelevant people who are interested in SAT solving, not in making things easy for users."*[2] Nevertheless, our interaction via the mailing list was insightful, as it helped us understanding the kernel community, as well as it provided requirements for a solution. For instance, in addition to providing sound fixes, a solution should be fast for user acceptance [38].
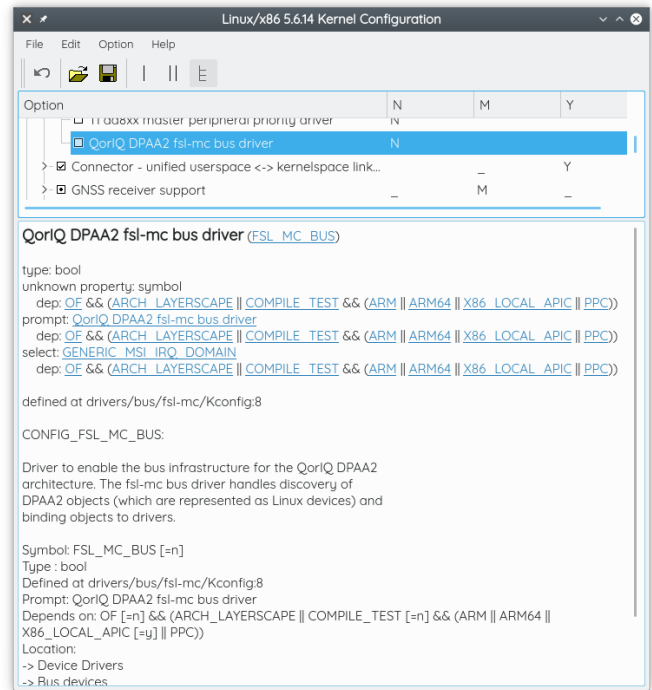
Fig. 2. A configuration conflict in *xconfig*

> **Interacting with the Linux Kernel Community**
>
> *For researchers, it is not clear how the Linux community works and how to effectively interact with it. While we were approached by a kernel developer, for other disseminations, we encourage the community to provide guidelines and requirements for solutions to eventually be integrated, as well as to actively approach researchers for mutual benefit.*

III. CONFIGFIX

We first introduce the notion of configuration conflict, then give a practical introduction into our tool CONFIGFIX, followed by discussing its architecture and major design decisions.

A. Configuration Conflicts

A conflict arises when changing the value of one or multiple features violates a constraint. To resolve conflicts, users need to (transitively) follow the dependencies, which is laborious and error-prone. A survey [34] found that users are commonly challenged with conflict resolution in the kernel configurator, with 20 % of the survey respondents needing roughly "a few dozen minutes" to resolve a conflict. The feature descriptions only provide incomplete and sometimes hard to understand (or even incorrect) advice, leading to users blindly choosing default or recommended values without grasping the consequences. Furthermore, default values sometimes contradict with the advice, for instance, when the description recommends enabling a feature, but the default value is "no" (disabled). As such, resolving conflicts can be particularly challenging and frustrating for inexperienced users, who lack the experience and resort to guessing.

Let us illustrate configuration conflicts with a feature that has particularly complex constraints. Figure 2 shows the

feature "QorIQ DPAA2 fsl-mc bus driver" (a bus infrastructure driver) in xconfig. Currently, it cannot be enabled (checkbox not selectable, indicated by missing underscores in columns "M" and "Y") due to the unmet constraint (current values of those features shown at the bottom of Fig. 2):

```
depends on OF && (ARCH_LAYERSCAPE ||
            (COMPILE_TEST && (ARM || ARM64 ||
            X86_LOCAL_APIC || PPC)))
```

Xconfig does not even show which parts of the constraint are unmet. The user needs to manually resolve the conflict by transitively looking at the features' dependencies, taking the full richness of its constraint language into account, which might mean needing to enable and disable a set of many features at the same time. Doing so, however, might have ripple effects by triggering the imperative choice propagation (select statement) in *xconfig*, which might in turn invalidate the user's resolution. The user might not even realize what other features the imperative choice propagation is enabling or disabling. Depending on constraints, the respective feature might not even be visible to the users in the configurator, further challenging the configuration process. The goal of CONFIGFIX is to facilitate this conflict resolution process by automatically finding the needed feature values to reach a desired configuration.

B. CONFIGFIX *Overview*

Workflow. Figure 3 shows the overall workflow of CONFIGFIX from the end user's perspective. (1) shows the view the user sees once they launch *xconfig* and have chosen "show all options" from the Options menu. This option shows features that would have normally been hidden, because they have unmet dependencies. The user can then identify the features they wish to enable, and click on "Add Symbol" which will add the feature to the bottom left pane, as shown in (2). At that point, the user can change the value of the feature to "Y", "N" or "M"—since the feature can be compiled as a loadable kernel module ("M"=module), otherwise only compiled directly into the kernel binary ("Y"=yes) or not at all ("N"=no). Once the user has added all the features they would like to change the values for, they can click on "Calculate Fixes," which will call CONFIGFIX's internal conflict resolution algorithm with the list of features specified in the bottom left pane. The returned solutions will be displayed in the bottom right pane as shown in (3). Each solution is a set of feature values that need to be made in order to set the wanted value(s) for the feature(s) indicated in (2). As there can be multiple ways to satisfy the user constraints, each specific solution can be viewed by selecting it from the Solutions combobox. In later versions, we might add feature drag and drop. The user can then apply any of the solutions, which would allow them to change the values of their desired feature(s).

DIMACS Export. Given the demand for accurate translations of variability models expressed in the Kconfig language, CONFIGFIX offers an export into a DIMACS file. DIMACS is
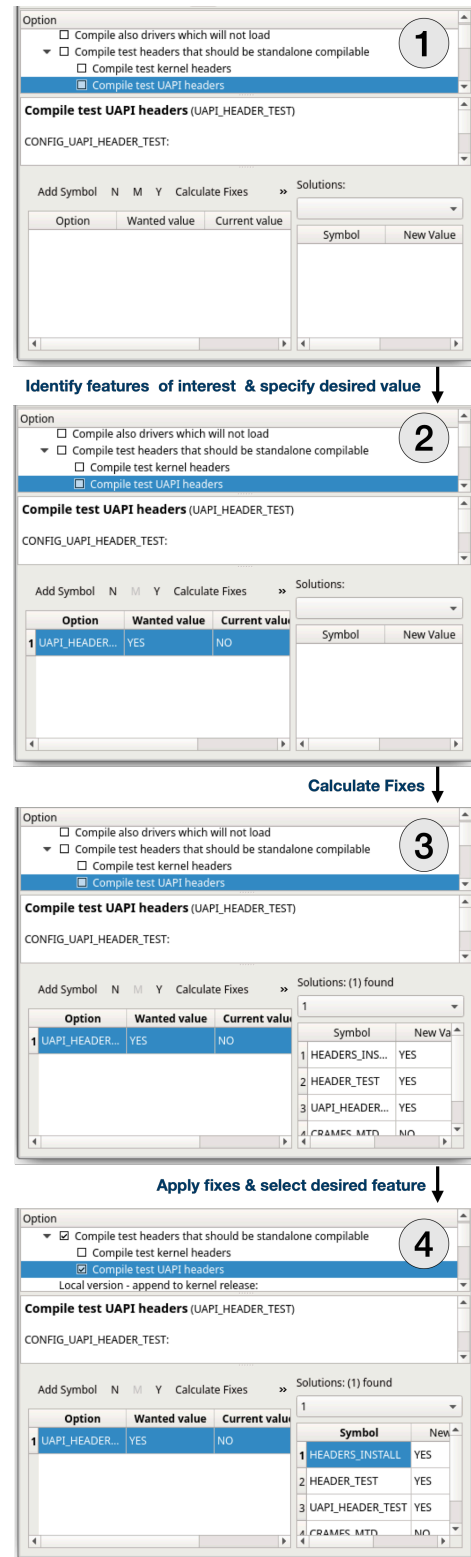

Fig. 3. User's workflow using CONFIGFIX inside *xconfig*

a common format accepted by many SAT solvers.[3] The export is launched via make cfoutconfig.

---

[3]http://www.satcompetition.org/2009/format-benchmarks2009.html

C. Solution Overview

Figure 4 illustrates CONFIGFIX's fix calculation. First, CONFIG-FIX translates the Kconfig model into a propositional formula according to the Kconfig semantics we reverse-engineered and re-implemented (and discuss shortly in Sec. IV-A). It then translates this formula into conjunctive normal form (CNF) to be able to later process it with a SAT solver. The formula encoding the Kconfig variability model represents the *hard constraints* that must always be satisfied. In other words, CONFIGFIX cannot violate any of these constraints during its fix generation. As input, CONFIGFIX also takes in the current configuration (a list of features and their corresponding values) and the user's configuration goal (a list of features to change and their desired values). CONFIGFIX considers the current configuration as *soft constraints*, since some of the current feature values in the configuration will need to change. As the next step, CONFIGFIX creates a single formula that is a conjunction of all hard and soft constraints and then queries the SAT solver for satisfiability. If it is satisfiable, then there is nothing to find fixes for, and the desired values can be applied. Otherwise, CONFIGFIX triggers our C-based RangeFix implementation to calculate fixes.

D. Fix Generation

Given a conflict, we want to find a fix that requires a minimal number of changes to the configuration. In the literature, various conflict-resolution algorithms exist [33], but many are not applicable here. They either produce only one fix, a long list of fixes (challenging users to identify/apply the most suitable one), or they only offer limited support for non-Boolean constraints. We selected *RangeFix* [83], which was designed to mitigate these shortcomings. Its fixes offer a range of values for certain features, and it supports non-Boolean features and constraints.

Rangefix's fixes adhere to three main properties:

- *correctness*: any configuration resulting from a fix must be correct, i.e. satisfy the violated constraints;
- *maximality of ranges*: when ranges overlap, the fix shall contain a maximum range;
- *minimality of changes*: the number of features that need to be changed should be reasonably minimal (realized using heuristics we defined) to avoid unnecessarily breaking features values set by the user before.

RangeFix generates fixes in three stages. Its input is an abstraction of a variability model (e.g., CNF formula), where features are represented by variables. In the first stage, all minimal sets of variables, which have to be changed, are generated. These sets are called diagnoses. In stages 2 and 3, the new values for each variable in a diagnosis are calculated.

We now explain the algorithm using an example. Let us define the tuple (V, e, c) as a *constraint violation*, where V is a set of variables, e a configuration with a value defined for each variable and c a set of constraints which is violated. The goal is to find a new configuration e', such that c is satisfied. We define our set of variables V as:

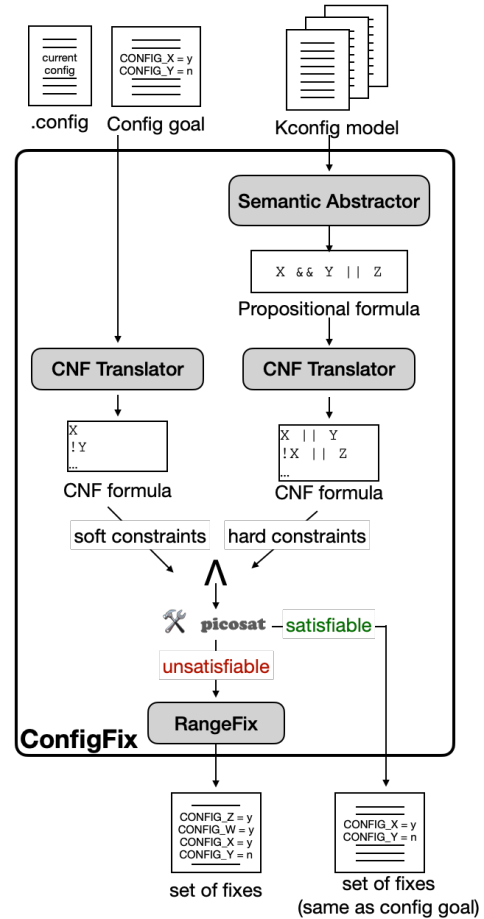$$\{m : Boolean, a : Integer, b : Integer\}$$

Fig. 4. Overview of ConfigFix's components

We define a set of constraints as:

$$(m \rightarrow a > 10) \ \wedge (\neg m \rightarrow b > 10) \ \wedge (a < b)$$

Finally, a configuration e with values for the three variables is required: $\{m := True, a := 6, b := 5\}$. This configuration violates first and the third constraint.

RangeFix generates the diagnoses during the first stage by using the constraint solver's ability to find unsatisfiable cores. During each iteration in the first step, a constraint from an unsatisfiable core is removed and the diagnoses are extended until no more unsatisfiable core is found. Applying this to our example yields the following two diagnoses: $\{m, b\}$ and $\{a, b\}$. At this point, we know that we need to change either $\{m, b\}$ or $\{a, b\}$ to resolve the conflict.

The stages 2 and 3 are subsequently performed together for each diagnosis. First, all unchanged variables are replaced by their current values during stage 2. Finally, in stage 3, the violated constraints are minimized via heuristic rules we defined and split into minimal clauses to generate the fixes. This leads to the following two fixes for the conflict:

- $[m := False, b : b > 10]$
- $[(a, b) : a > 10 \wedge a < b]$

If any of these two fixes is applied, all previously violated constraints will be satisfied again. If the first fix is chosen, then m needs to be set to $False$ and b simply to any value larger

than 10. If the second fix is chosen, then a needs to be set to a value larger than 10, and b must simply be larger than a. This illustrates the advantages of RangeFix compared to existing conflict-resolution algorithms. While there are infinitely many possible solutions, the user is presented with only the minimal set. Other approaches might have presented only one of the two solutions or a very long list of possible solutions, since there are infinitely many possible combinations that satisfy a being smaller than b in the second fix. As such, CONFIGFIX returns a maximum of three fixes for each conflict even if more fixes exist. In most cases, presenting three fixes for a user is sufficient to choose a suitable fix and generating more fixes can take significantly more time.

E. Example Fix

We now demonstrate how CONFIGFIX finds different fixes for a conflict, and how it deals with the intricacies of Kconfig's semantics, especially its various language constructs to define constraints (cf. Sec. II-C). Given a default configuration, assume a user wants to enable the feature `MEDIA_TUNER_SIMPLE`, which provides support for various media tuners. The feature has dependencies, and its parent feature depends on other features as it has visibility constraints. In the default configuration, the parent is hidden and the dependencies are not met, so it is not configurable by the user.

For this conflict, eight possible fixes exist and, each fix changes the values of five features, although some values change implicitly through Kconfig's imperative choice propagation (`select` statements). Generally, the feature can be enabled through two different means: The parent feature can be made visible and then the user can explicitly set a value for `MEDIA_TUNER_SIMPLE` after its dependencies have been met. The other possibility is to enable the module through a `select` statement while keeping the module invisible.

In the first case, the following fix first makes the feature visible, and then the user can set an explicit value:

- `MEDIA_SUPPORT => yes`
- `MEDIA_DIGITAL_TV_SUPPORT => yes`
- `MEDIA_SUBDRV_AUTOSELECT => no`
- `MEDIA_TUNER_SIMPLE => yes`

In the latter case, the following fix ensures that the module is selected and, therefore, enabled:

- `MEDIA_SUPPORT => yes`
- `MEDIA_ANALOG_TV_SUPPORT => yes`

The feature is invisible, though, and it might not be obvious to the user why the feature has been enabled. In the end, all eight fixes calculated are able to enable `MEDIA_TUNER_SIMPLE` in one way or another.

IV. EXPERIENCES AND CHALLENGES

We now discuss experiences and challenges faced when realizing CONFIGFIX. Overall, our first attempt [38] was much closer to the RangeFix implementation. We experimented with the C-based translation by Vegard Nossum (cf. Sec. II-C) and an existing Scala-based RangeFix implementation prototypically integrated into *xconfig*, but which only covered the first of the three stages of RangeFix. We surveyed kernel developers and Kconfig-SAT members about the user interface we implemented, and provided screencasts illustrating the solution [38]. The results contributed to the present attempt, where we focused on a purely C-based translation and fix generation, the C-based SAT solver Picosat [21], and an improved integration into *xconfig*.

A. Semantic Abstractor

The largest challenge was obtaining a sound and stable logical abstraction of Kconfig. To obtain requirements, we interacted with the community, specifically via the kconfig-sat mailing list. We recognized a strong preference for SAT solvers, as opposed to more expressive solvers, such as SMT. Even though, the latter could support a larger part of the semantics, it was pointed out that integrating a SAT solver could also help at other places in the kernel, especially CPU scheduling support. Furthermore, SMT solvers are typically slower, and not many come with a GPL-compatible license, as required for integration into the kernel's codebase. For practical matters, we learned it should be possible to use or compile the SAT solver with the tools needed to compile the Linux kernel, such as *gcc*. This requirement excluded some modern and very fast solvers, including CaDiCaL [20].

In general, our options for realizing a propositional semantic abstraction for a SAT solver were to (i) develop a new translation from scratch, to (ii) build on *Satconfig* (cf. Sec. II-C) or to (iii) investigate other existing alternatives. Strategy (i) has the disadvantage of lacking a reference, when the translation is checked for correctness. Since correctness was essential, we disregarded this strategy. As explained above, Satconfig showed deficiencies in handling `tristate` features and in code documentation. Based on others' systematic comparison of translations [28] (also cf. Sec. II-C), we chose to re-implement and extend the Scala-based *KconfigReader* [40].

Our overall strategy was to inspect the translation in Kconfig-Reader, to re-implement it freely in C, being inspired by data structures from Satconfig, and to test the translation incrementally with smaller, hand-crafted models, thereby also debugging and fixing our implementation as well as remaining deficiencies of KconfigReader. Finally, we tested with example conflicts in the full kernel model (one of these conflicts we described in Sec. III-E). KconfigReader works accurately for `bool` and `tristate` features, as well as for many non-Boolean properties, but with some remaining smaller limitations.

Adapting and re-implementing KconfigReader in C allowed the direct integration into the kernel configurator in the kernel source tree. It also allowed us to use the configurator's parser to parse the variability model. In fact, parsing Kconfig is another general challenge acknowledged by the researchers who produced translation. So, relying on the maintained parser provides robustness. We traversed the internal, AST-based representation and stored intermediate results in our own C data structures. A specific challenge was to implement scalable representations of the propositional formula in C, in a way that it can be easily traversed and transformed into conjunctive normal form by applying the typical logical laws and a Tseitin transformation.

As a consequence, some parts of the translation had to be implemented in a completely different way than in KconfigReader.

> **Propositional Transformation in C**
>
> *C was not ideal to model propositional logic formulas. An OO language, especially one with functional-programming constructs, would have likely led to much cleaner code.*

To account for limitations, we needed to deviate from the semantics realized in KconfigReader. One example was the translation of the imperative choice propagation—over 10,000 `select` statements exist in the whole variability model. So, how to model this type of constraint has substantial impact on the performance. KConfigReader models the `select`-constraints under the constraints for the selected feature and not the selecting feature. An advantage of this behavior is the number of constraints, since there will only be one constraint for each selected feature independently by how many other features it is selected. The disadvantage is that the constraints can potentially become very large formulas, when a feature is selected by many other features and depending on how many constraints these features have.

> **Challenges of Achieving a Sound Semantic Abstraction**
>
> *As known for over a decade, the Kconfig syntax and semantics are intricate. Especially, the semantics are more expressive than logical representations of off-the-shelf reasoners, including SMT and of course SAT. In addition, the continuous evolution of Kconfig, lack in documentation, and non-obvious semantic abstractions, challenged providing a sound abstraction. The kernel configurator's code is also not well documented, only some of its data structures, no functions. Furthermore, a naive CNF translation explodes without Tseitin transformation and incorporating domain/-expert knowledge for effectively splitting clauses.*

B. Choosing a SAT Solver

Choosing PicoSAT [21] was inspired by Satconfig (cf. Sec. II-C), which showed that PicoSAT can be easily integrated into the kernel. More importantly, it: (i) is written in C and can be compiled with `gcc`; (ii) has a C-API, so can be called called directly within *xconfig* without needing external calls; (iii) has a Linux-compatible license (MIT); (iv) can identify and return unsatisfiable cores; and (v) is reasonably fast [22]. This made PicoSAT our best candidate. However, a downside was that the solver is outdated and not actively maintained anymore, despite still being use in industry.

> **Choice of SAT Solver**
>
> *While faster SAT solvers exist, the need to use a pure C-based SAT solver with a C API and Linux-kernel-compatible license, restricted our options to PicoSAT. While it is not actively maintained anymore, it has a well-structured implementation and turned out to be fast enough for interactive fix generation in* CONFIGFIX.

C. GUI Integration

We extended the graphical configurator *xconfig* to provide an intuitive interface for entering the desired feature values, for observing the proposed fixes, and for applying the desired ones. We decided on having the conflict resolution integrated within the same window to follow conflict resolution interfaces found in other configurators, specifically that of ECOS [15], [34], [82]. We added a new pane at the bottom of *xconfig* to collect the features desired by the user. The view is divided into two parts: (i) the collected feature list and (ii) the solutions obtained from CONFIGFIX. We initially planned to support drag and drop, left it for later in favor of obtaining a working solution for now.

After realizing the initial UI, we started realizing the interfaces between the UI, our RangeFix implementation, and the SAT solver PicoSAT via its C API. *xconfig* is written in C++ with the Qt toolkit, while we implemented CONFIGFIX in pure C and call it within *xconfig*. The Glib library is used in CONFIGFIX and *xconfig* for its variety of available data structures resulting.

The challenge with implementing and integrating the user interface was in understanding the libraries Glib and Qt, specifically Qt slots and signals, and the type-checked event signalling mechanism in Qt. Furthermore, we needed to figure out how to interface C++ with C code (name unmangling), and how to pass data around between user interface elements (parent-child and sibling in the interface hierarchy).

D. Scalability and Performance Improvements

The main challenge was translating the huge formulas of the full kernel variability model with over 15,000 features into CNF. Notably, the encoding of certain Kconfig aspects had a substantial impact on the resulting CNF and SAT solver performance.

To improve the performance of the SAT solver, we changed the encoding of the `select` statement by splitting up the various statements. Instead of a single constraint for a selected feature, each `select` statement now creates constraints on its own. But, since this also had an affect on we encoded the dependencies of a feature, introducing a new variable was needed, which indicates whether a feature has been selected. This significantly simplified the constraint encoding.

While we have increased the number of constraints in total by several thousand constraints, we were able to reduce the number of CNF clauses and auxiliary variables significantly As a consequence, a single run of PicoSAT became more than 65 % faster than it was before. A disadvantage of this decision is, that we lost the ability to syntactically check our constraints for equivalence against the constraints produced by KConfigReader. Still, the gain in performance justified this decision, and we conceived an alternative evaluation, explained shortly.

Achieved Performance

We achieved a final translation time of the entire Linux kernel variability model into a CNF formula of around 1.5 seconds on an Intel i7 laptop. The initial run of PicoSAT to check for satisfiability takes about 2.5 to 3 seconds. Finally, finding fixes for a conflict can be achieved in as little as 1 second in some cases, although the number depends heavily on the conflict and the number of enabled features. The most impact on improving the performance came from incorporating our domain/expert knowledge into the translation, including effective formula splitting and using a Tseitin transformation.

V. EVALUATION

We now briefly discuss our evaluation of CONFIGFIX.

A. Conflict and Fix Generation

With over 15,000 features, the configuration space for the Linux kernel is huge and, therefore, crafting a small number of examples to be evaluated is not sufficient. Instead, a more systematic approach is needed. We made use of the kernel's ability to randomly sample a configuration and then we created random conflicts (that are definitely resolvable, see below) for each sampled configuration.

To obtain sufficiently diverse configurations, we used the kernel's `randconfig` tool to generate configuration samples for three of the more popular available architectures in the kernel: `x86_64`, `arm64`, and `openrisc`. Randconfig allows to skew the probabilities for features of type `bool` and `tristate` to be enabled. Our probabilities ranged from 10 % to 90 % for a feature to be set to `no`. This setting affects the number of enabled features and, consequently, the fix sizes (number of features to be changed), since a low number of enabled features tends to lead to larger fixes regardless of the conflict size. This way, we can evaluate the performance of CONFIGFIX for varying fix sizes.

For each random configuration we then introduce conflicts by randomly choosing target features that: (i) have a prompt (i.e., are not completely invisible, like derived features [16]), (ii) are of type `bool` or `tristate`, (iii) are not a choice group, and (iv) have at least one other possible value different from its current value that cannot be selected at the moment. In the case of `tristate` options with two values that cannot be selected at the moment, the target value is randomly chosen from these two values.

While finding fixes, CONFIGFIX may not return results for two reasons. The first is that a fix may not exist. Some features depend on other architecture-specific options; therefore, they can only be configured for certain architectures. Thus, including such architecture-dependent features into conflicts to be resolved on a different architecture by definition has no fix. In this case, CONFIGFIX not returning a fix is the expected behavior. On the other hand, CONFIGFIX may also not return fixes due to bugs in our implementation or design. In order to objectively evaluate CONFIGFIX, we, therefore, want to ensure that our chosen target features can be configured (i.e., the conflict can be resolved) for the used architecture to rule out the first possibility. Therefore, we generate a base configuration for each architecture using `randconfig` with the probability of 100% for a feature to be enabled. Such base configuration will have as many enabled features as possible. Before including a feature in a conflict, we compare its target value with the base configuration to see whether this value can indeed be obtained for that architecture.

For each configuration sample, we restrict the selection of conflicting features to those that can receive their target values on the chosen architecture, as witnessed by the base configuration. We create such resolvable conflicts containing between 1–10 features and let CONFIGFIX generate fixes for each conflict.

B. Analysis

When applying each produced fix for a given conflict, we analyze whether (i) the fix resolves the conflict, i.e. whether our randomly chosen features obtain their target values and whether (ii) every feature in a fix obtains its target value if the fix resolves the conflict after applying the fix. This results in three possible outcomes for each fix:

- **Applicable and Resolves Conflict:** The fix is fully applicable (i.e., we can apply all the specified values in the fix) and resolves the conflict. This is the optimal outcome.
- **Not Applicable, but Resolves Conflict:** The fix is not fully applicable (i.e., some of the values specified in the fix cannot be applied), but it resolves the conflict nonetheless. While the fix is not optimal, since it contains invalid or redundant feature values, it is an acceptable outcome, since the conflict is resolved. Specifically, the features to change, which are part of the fix, can be changed.
- **Does Not Resolve Conflict:** The proposed fix does not resolve the conflict, since the target features cannot be changed by applying the conflict.

C. Results

We generated a total of 27 random sample configurations with a varying number of enabled features for the three target architectures (`x86_64`, `arm64`, and `openrisc`). For each configuration, we created 50 conflicts containing between 1–10 features to change. This resulted in 1350 conflicts for CONFIGFIX to solve. We summarize the results in Table I.

Out of the 1,350 conflicts, CONFIGFIX returned at least one fix for 1,055 (78.2 %) conflicts, of which 723 (53.6%) were resolved. The remaining 295 conflicts did not receive any fixes, mostly due to a limit on the CONFIGFIX running time, which we used to make testing feasible. For the 1,055 conflicts that received at least one fix, a total of 2,482 fixes were returned (recall our intentional limitation to a maximum of three fixes per conflict, as discussed in Sec. III-D).

As seen in Fig. 5, the obtained fixes comprise 2–175 features that need to be changed, although the typical range is much smaller. The distribution is left-skewed, with an average fix size of 23 features, and the majority of the fixes lying around the
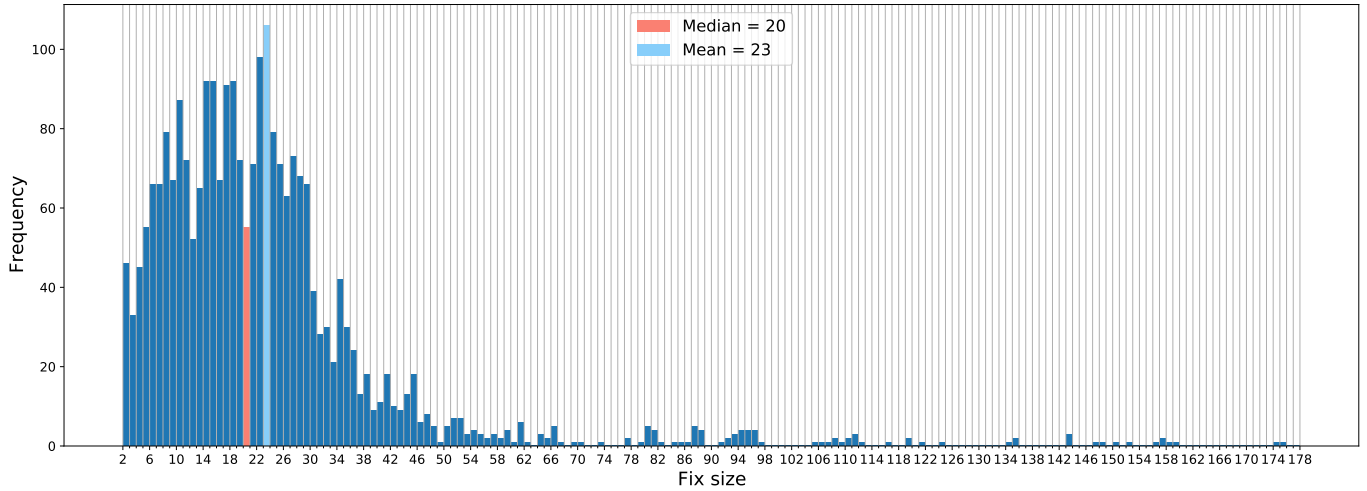
Fig. 5. Distribution of fix sizes (number of features that need to be changed to resolve the conflict)

fixes returned by CONFIGFIX resolved the conflict.

┌─ **Evaluation Summary** ─────────────────────────

We evaluated CONFIGFIX *with 1,350 conflicts of different sizes that were randomly generated for 27 configuration samples on three hardware architectures. 65 % of the fixes generated by* CONFIGFIX *resolved the conflict.*
└──

While the ideal outcome would have been to resolve all conflicts, we believe that this percentage is still acceptable given that the semantic abstraction is sound, but cannot be complete. While perhaps an encoding into SMT could yield a slightly better result, recall that there was a strong preference amongst the Linux community for SAT solving using a C-based solver.

TABLE I
CONFIGFIX EVALUATION RESULTS FOR RANDOM CONFIGURATIONS AND CONFLICTS

metric	value
number of sampled configurations	27
conflict sizes	1–10
generated conflicts	1,350[1] (100.0 %)
conflicts with at least one generated fix	1,055 (78.2 %)
number of resolved conflicts	723 (53.6 %)
total number of generated fixes	2,482 (100.0 %)
fixes that resolve the conflict	**1,609 (65.0 %)**
fully applicable and resolve the conflict	1,317 (53.0 %)
not fully applicable, but resolve the conflict	292 (12.0 %)
do not resolve the conflict	868 (35.0 %)

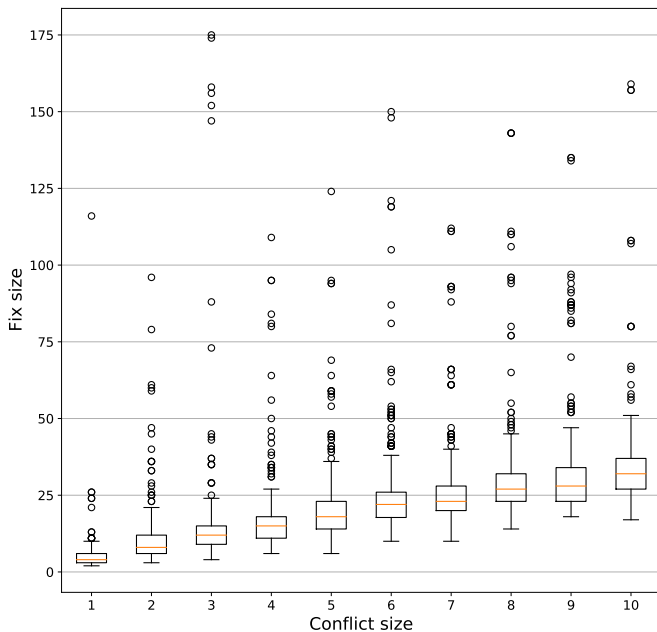[1] For each configuration sample, five conflicts of each size.

Fig. 6. Distribution of fix sizes (number of features that need to be changed) by conflict size (number of unchangeable features that a user wants to change)

median of 20 features. Figure 6 shows that outliers are common for all conflict sizes, but in general, fix sizes demonstrate a seemingly linear dependence on the conflict size. This indicates that CONFIGFIX can be used for resolving conflicts of varying sizes without the risk of fix size explosion.

We also investigated the outcome of each fix, summarized in Table I. The optimal outcome (**Applicable and Resolves Conflict**) was achieved for 1,317 of these fixes (53 %), while 292 fixes (12 %) still resolved the conflict despite not being fully applicable (**Not Applicable, but Resolves Conflict**). Finally, 868 (35 %) returned fixes that did not resolve the conflict (**Does Not Resolve Conflict**). So, in summary, 65 % of the

VI. CONCLUSION

We reported our experience of leveraging results from 13 years of academic research. We created the tool CONFIGFIX, which we integrated in the Linux kernel configurator by adhering to all requirements coming from the Linux community in the context of the Kconfig-SAT initiative [42]. CONFIGFIX realizes a stable and sound transformation of the configurator's underlying variability-modeling language Kconfig into a propositional abstraction, as well as it provides a tested

configuration-conflict resolution technique that can guide users achieving their desired Linux kernel configuration. We believe that our tool [2] helps the Linux kernel community not only supporting the configuration process, but also conduct further analyses or support based on the stable translation. Likewise, we invite researchers evaluating their own and novel techniques upon the translation, as well as to improve the fixes, for instance, providing optimization support based on quality attributes.

REFERENCES

[1] "The VAMOS project, undertaker," https://vamos.informatik.uni-erlangen.de/trac/undertaker.

[2] "ConfigFix," https://bitbucket.org/easelab/configfix, 2020.

[3] I. Abal, C. Brabrand, and A. Wasowski, "42 variability bugs in the linux kernel: A qualitative analysis," in *ASE*, 2014.

[4] E. Abbasi, A. Hubaux, M. Acher, Q. Boucher, P. Heymans, A. Heymans, F. FSR, and W. Region, "What's in a web configurator? empirical results from 111 cases," University of Namur, Tech. Rep. P-CS-TR CONF-000001, 2012.

[5] G. Antoniol, U. Villano, E. Merlo, and M. Di Penta, "Analyzing cloning evolution in the linux kernel," *Information and Software Technology*, vol. 44, no. 13, pp. 755–765, 2002.

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

[7] R. Bashroush, M. Garba, R. Rabiser, I. Groher, and G. Botterweck, "Case tool support for variability management in software product lines," *ACM Comput. Surv.*, vol. 50, no. 1, Mar. 2017.

[8] D. Benavides, S. Segura, and A. Ruiz-Cortés, "Automated analysis of feature models 20 years later: A literature review," *Information Systems*, vol. 35, no. 6, pp. 615–636, 2010.

[9] T. Berger, D. Lettner, J. Rubin, P. Grünbacher, A. Silva, M. Becker, M. Chechik, and K. Czarnecki, "What is a feature? a qualitative study of features in industrial software product lines," in *SPLC*, 2015.

[10] T. Berger, D. Nair, R. Rublack, J. M. Atlee, K. Czarnecki, and A. Wasowski, "Three cases of feature-based variability modeling in industry," in *International Conference Model-Driven Engineering Languages and Systems (MODELS)*, 2014, pp. 302–319.

[11] T. Berger, R.-H. Pfeiffer, R. Tartler, S. Dienst, K. Czarnecki, A. Wasowski, and S. She, "Variability mechanisms in software ecosystems," *Information and Software Technology*, vol. 56, no. 11, pp. 1520–1535, 2014.

[12] T. Berger, R. Rublack, D. Nair, J. M. Atlee, M. Becker, K. Czarnecki, and A. Wasowski, "A survey of variability modeling in industrial practice," in *VaMoS*, 2013.

[13] T. Berger and S. She, "Formal semantics of the CDL language," 2010, technical Note. Available at http://thorsten-berger.net/cdl_semantics.pdf.

[14] T. Berger, S. She, R. Lotufo, K. Czarnecki, and A. Wasowski, "Feature-to-code mapping in two large product lines," in *SPLC*, 2010.

[15] T. Berger, S. She, R. Lotufo, A. Wasowski, and K. Czarnecki, "Variability modeling in the real: A perspective from the operating systems domain," in *ASE*, 2010.

[16] ——, "A Study of Variability Models and Languages in the Systems Software Domain," *IEEE Trans. Softw. Eng.*, vol. 39, no. 12, pp. 1611–1640, Dec. 2013.

[17] T. Berger, J.-P. Steghöfer, T. Ziadi, J. Robin, and J. Martinez, "The state of adoption and the challenges of systematic variability management in industry," *Empirical Software Engineering*, vol. 25, pp. 1755–1797, 2020.

[18] D. Beuche, "Variants and variability management with pure::variants," in *SPLC*, 2004.

[19] S. Bhartiya, "Linux is the largest software development project on the planet: Greg Kroah-Hartman," http://cio.com/article/3069529.

[20] A. Biere, "CaDiCaL Simplified Satisfiability Solver," http://fmv.jku.at/cadical.

[21] ——, "PicoSAT," http://fmv.jku.at/picosat.

[22] ——, "Adaptive Restart Strategies for Conflict Driven SAT Solvers," in *SAT*, 2008.

[23] M. Brambilla, J. Cabot, and M. Wimmer, *Model-Driven Software Engineering in Practice*. Morgan & Claypool Publishers, 2012.

[24] ——, *Model-driven software engineering in practice*. Morgan & Claypool, 2017.

[25] B. Combemale, R. France, J.-M. Jézéquel, B. Rumpe, J. Steel, and D. Vojtisek, *Engineering modeling languages: Turning domain knowledge into tools*. CRC Press, 2016.

[26] K. Czarnecki and U. W. Eisenecker, *Generative Programming: Methods, Tools, and Applications*, 2000.

[27] D. Dhungana, P. Grünbacher, and R. Rabiser, "The DOPLER Meta-Tool for Decision-Oriented Variability Modeling: a Multiple Case Study," *J. ASE*, vol. 18, no. 1, pp. 77–114, Mar. 2011.

[28] S. El-Sharkawy, A. Krafczyk, and K. Schmid, "Analysing the kconfig semantics and its analysis tools," in *GPCE*, 2015.

[29] S. Erdweg, T. van der Storm, M. Völter, M. Boersma, R. Bosman, W. R. Cook, A. Gerritsen, A. Hulshout, S. Kelly, A. Loh *et al.*, "The State of the Art in Language Workbenches," in *Software Language Engineering*. Springer, 2013.

[30] D. Fernandez-Amoros, R. Heradio, C. Mayr-Dorn, and A. Egyed, "A kconfig translation to logic with one-way validation system," in *SPLC*, 2019.

[31] R. Flores, C. Krueger, and P. Clements, "Mega-Scale Product Line Engineering at General Motors," in *Proc. SPLC*, 2012.

[32] Google Groups, "kconfig-sat," https://groups.google.com/forum/#!forum/kconfig-sat.

[33] A. Hubaux, D. Jannach, C. Drescher, L. Murta, T. Männistö, K. Czarnecki, P. Heymans, T. Nguyen, and M. Zanker, "Unifying software and product configuration: A research roadmap," in *ConfWS*, 2012.

[34] A. Hubaux, Y. Xiong, and K. Czarnecki, "A User Survey of Configuration Challenges in Linux and eCos," in *VaMoS*, 2012.

[35] A. Israeli and D. G. Feitelson, "Characterizing software maintenance categories using the linux kernel," The Hebrew University of Jerusalem, Tech. Rep. 2009–10, 2009.

[36] ——, "The linux kernel as a case study in software evolution," *Journal of Systems and Software*, vol. 83, no. 3, pp. 485–501, 2010.

[37] Y. Jiang, B. Adams, and D. M. German, "Will my patch make it? and how fast? case study on the linux kernel," in *MSR*, 2013.

[38] D. Jonsson, "A case study of interactive conflict-resolution support in software configuration," Master's thesis, Chalmers University of Technology, 2016. [Online]. Available: https://hdl.handle.net/20.500.12380/238168

[39] K. Kang, S. Cohen, J. Hess, W. Nowak, and S. Peterson, "Feature-oriented domain analysis (FODA) feasibility study," Carnegie-Mellon University, Pittsburgh, PA, USA, Tech. Rep., 1990.

[40] C. Kästner, "KConfig Reader," https://github.com/ckaestne/kconfigreader.

[41] C. Kästner, P. G. Giarrusso, T. Rendel, S. Erdweg, K. Ostermann, and T. Berger, "Variability-aware parsing in the presence of lexical macros and conditional compilation," in *OOPSLA*, 2011.

[42] Kernelnewbies, "Linux kconfig SAT integration," https://kernelnewbies.org/KernelProjects/kconfig-sat, accessed: 2019-12-05.

[43] Kernel.org, "Kconfig language," https://www.kernel.org/doc/Documentation/kbuild/kconfig-language.txt, accessed: 2019-12-05.

[44] S. C. Kleene, "On notation for ordinal numbers," *The Journal of Symbolic Logic*, vol. 3, no. 4, pp. 150–155, 1938.

[45] C. W. Krueger, D. Churchett, and R. Buhrdorf, "Homeaway's transition to software product line practice: Engineering and business results in 60 days," in *SPLC*, 2008.

[46] C. W. Krueger and P. C. Clements, "Systems and software product line engineering with biglever software gears," in *SPLC*, 2013.

[47] J. Krueger, W. Gu, H. Shen, M. Mukelabai, R. Hebig, and T. Berger, "Towards a better understanding of software features and their characteristics: A case study of marlin," in *VaMoS*, 2018.

[48] R. Lämmel, *Software languages: Syntax, semantics, and metaprogramming*. Springer, 2018.

[49] M. Larabel, "The Linux Kernel Enters 2020 At 27.8 Million Lines In Git But With Less Developers For 2019," https://www.phoronix.com/scan.php?page=news_item&px=Linux-Git-Stats-EOY2019.

[50] J. Lawall and G. Muller, "Coccinelle: 10 years of automated evolution in the linux kernel," in *USENIX ATC*, 2018.

[51] R. Lindohf, J. Krueger, E. Herzog, and T. Berger, "Software product-line evaluation in the large," *Empirical Software Engineering*, 2020, preprint.

[52] R. Lotufo, S. She, T. Berger, K. Czarnecki, and A. Wasowski, "Evolution of the Linux Kernel Variability Model," in *SPLC*, 2010.

[53] J. Meinicke, T. Thüm, R. Schröter, F. Benduhn, T. Leich, and G. Saake, *Mastering Software Variability with FeatureIDE*. Springer, 2017.

[54] S. Nadi, T. Berger, C. Kästner, and K. Czarnecki, "Mining configuration constraints: Static analyses and empirical results," in *ICSE*, 2014.

[55] ——, "Where do configuration constraints stem from? an extraction approach and an empirical study," *IEEE Trans. Softw. Eng.*, vol. 41, no. 8, pp. 820–841, 2015.

[56] S. Nadi and R. C. Holt, "Make it or break it: Mining anomalies from linux kbuild," in *WCRE*, 2011, pp. 315–324.

[57] ——, "Mining kbuild to detect variability anomalies in linux," in *CSMR*, 2012.

[58] D. Nešić, J. Krüger, S. Stănciulescu, and T. Berger, "Principles of Feature Modeling," in *FSE*, 2019.

[59] V. Nossum, "satconfig," https://github.com/vegard/linux-2.6/tree/v4.7+kconfig-sat, accessed: 2019-12-05.

[60] Y. Padioleau, J. Lawall, R. R. Hansen, and G. Muller, "Documenting and automating collateral evolutions in linux device drivers," *SIGOPS Oper. Syst. Rev.*, vol. 42, no. 4, p. 247–260, Apr. 2008.

[61] L. Passos, R. Queiroz, M. Mukelabai, T. Berger, S. Apel, K. Czarnecki, and J. Padilla, "A study of feature scattering in the linux kernel," *IEEE Trans. Softw. Eng.*, 2018, preprint.

[62] L. Passos, L. Teixeira, N. Dintzner, S. Apel, A. Wasowski, K. Czarnecki, P. Borba, and J. Guo, "Coevolution of variability models and related software artifacts," *Empirical Softw. Engg.*, vol. 21, no. 4, pp. 1744–1793, Aug. 2016.

[63] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed., ser. Series in Artificial Intelligence. Upper Saddle River, NJ: Prentice Hall, 2010. [Online]. Available: http://aima.cs.berkeley.edu/

[64] S. Schauss, R. Lämmel, J. Härtel, M. Heinz, K. Klein, L. Härtel, and T. Berger, "A chrestomathy of dsl implementations," in *International Conference on Software Language Engineering (SLE)*. ACM, 2017.

[65] K. Schmid, I. John, R. Kolb, and G. Meier, "Introducing the pulse approach to an embedded system population at testo ag," in *ICSE*, 2005.

[66] D. C. Sharp, "Reducing avionics software cost through component based product line development," in *17th DASC. AIAA/IEEE/SAE. Digital Avionics Systems Conference. Proceedings (Cat. No. 98CH36267)*, 1998.

[67] S. She, "LVAT," https://code.google.com/archive/p/linux-variability-analysis-tools/, 2013.

[68] S. She and T. Berger, "Formal semantics of the kconfig language," 2010, technical Note. Available at http://www.eng.uwaterloo.ca/~shshe/kconfig_semantics.pdf.

[69] S. She, R. Lotufo, T. Berger, A. Wasowski, and K. Czarnecki, "The variability model of the linux kernel," in *VaMoS*, 2010.

[70] ——, "Reverse engineering feature models," in *ICSE*, 2011.

[71] J. Sincero, H. Schirmeier, W. Schröder-Preikschat, and O. Spinczyk, "Is The Linux Kernel a Software Product Line?" in *SPLC-OSSPL*, 2007.

[72] J. Sincero and W. Schröder-Preikschat, "The linux kernel configurator as a feature modeling tool," in *ASPL*, 2008.

[73] M. Svahnberg and J. Bosch, "Evolution in software product lines: Two cases," *Journal of Software Maintenance*, vol. 11, no. 6, pp. 391–422, Nov. 1999.

[74] R. Tartler, C. Dietrich, J. Sincero, W. Schröder-Preikschat, and D. Lohmann, "Static analysis of variability in system software: The 90,000 #ifdefs issue," in *USENIX ATC*, 2014.

[75] R. Tartler, D. Lohmann, J. Sincero, and W. Schröder-Preikschat, "Feature consistency in compile-time-configurable system software: Facing the linux 10,000 feature problem," in *EuroSys*, 2011.

[76] R. Tartler, J. Sincero, W. Schröder-Preikschat, and D. Lohmann, "Dead or alive: Finding zombie features in the linux kernel," in *FOSD*, 2009.

[77] Y. Tian, J. Lawall, and D. Lo, "Identifying linux bug fixing patches," in *ICSE*, 2012.

[78] P. Toft, D. Coleman, and J. Ohta, "A cooperative model for cross-divisional product development for a software product line," in *SPLC*, 2000.

[79] G. Tseitin, "On the complexity of derivation in propositional calculus," *Zapiski Nauchnykh Seminarov LOMI*, vol. 8, 01 1983.

[80] Q. Tu and M. W. Godfrey, "Evolution in open source software: A case study," in *ICSM*, 2000.

[81] F. J. van der Linden, K. Schmid, and E. Rommes, *Software Product Lines in Action: The Best Industrial Practice in Product Line Engineering*, 2007.

[82] B. Veer and J. Dallaway, "The eCos component writer's guide," the eCos component writer's guide, Available from http://ecos.sourceware.org/docs-2.0/cdl-guide/cdl-guide.html.

[83] Y. Xiong, H. Zhang, A. Hubaux, S. She, J. Wang, and K. Czarnecki, "Range fixes: Interactive error resolution for software configuration," *IEEE Trans. Softw. Eng.*, vol. 41, no. 6, pp. 603–619, June 2015.

[84] C. Zengler and W. Küchlin, "Encoding the linux kernel configuration in propositional logic," in *ECAI*, 2010.

[85] R. Zippel and contributors, "kconfig-language.txt," available in the kernel tree at kernel.org.