

Automated Diagnosis of Product-line Configuration Errors on Feature Models *

Jules White and Douglas C. Schmidt

Vanderbilt University,

Department of Electrical Engineering and Computer Science

Box 1679 Station B

Nashville, TN, 37235, USA

Email: {jules, schmidt}@dre.vanderbilt.edu

David Benavides and Pablo Trinidad and Antonio Ruiz–Cortés

University of Seville

Department of Computer Languages and Systems

ETSII - Av. de la Reina Mercedes S/N

41012 Seville - Spain

Email: {benavides, ptrinidad, aruiz}@us.es

Abstract

Feature models are widely used to model software product-line (SPL) variability. SPL variants are configured by selecting feature sets that satisfy feature model constraints. Configuration of large feature models involve multiple stages and participants, which makes it hard to avoid conflicts and errors. New techniques are therefore needed to debug invalid configurations and derive the minimal set of changes to fix flawed configurations.

This paper provides three contributions to debugging feature model configurations: (1) we present a technique for transforming a flawed feature model configuration into a Constraint Satisfaction Problem (CSP) and show how a constraint solver can derive the minimal set of feature selection changes to fix an invalid configuration, (2) we show how this diagnosis CSP can automatically resolve conflicts between configuration participant decisions, and (3) we present experiment results that evaluate our technique. These results show how our technique scales to models with over 5,000 features, which is beyond the size used to validate other automated techniques.

*This work has been partially supported by the European Commission (FEDER) and Spanish Government under CICYT project Web-Factories (TIN2006-00472).

1 Introduction

Current trends and challenges. Software Product-Lines (SPLs) are a technique for creating software applications composed from reusable parts that can be re-targeted for different requirement sets. For example, in the automotive domain, an SPL can be created that allows a car's software to provide Anti-lock Braking System (ABS) capabilities or simply standard braking. Each unique configuration of an SPL is called a *variant*.

SPL variants cannot be constructed arbitrarily. For example, a car cannot have both ABS and standard braking software controllers. A key step in building an SPL, therefore, is creating a model of the SPL's variability and the constraints on variant configuration. An effective technique for capturing these configuration constraints is *feature modeling* [11].

Feature models document SPL variability and configuration rules through the notion of *features*. Each feature represents a point of variability in the SPL, such as the type of braking system. A feature model can capture different types of variability, ranging from *SPL variability* (variations in customer or market requirements) to *software variability* (variations in software implementation)[13].

SPL variants can be specified as a selection or configuration of features. Feature models of these SPL variants are arranged in a tree-like structure where each successively deeper level in the tree corresponds to a more fine-grained

configuration option for the product-line variant, as shown by the feature model in Figure 1. The parent-child and cross-tree relationships capture the constraints that must be adhered to when selecting a group of features for a variant.

Existing research has focused on ensuring that the features chosen from feature models are correct and consistent with the SPL and variant requirements. For example, work has been done on automating the process of deriving a set of features that meet a requirement set using boolean circuit satisfiability techniques [12] or Constraint Satisfaction Problems (CSPs) [4]. Moreover, numerous tools have been developed, such as Big Lever Software Gears [6], Pure::variants [5], and the Feature Model Plug-in [7], to support the construction of feature models and correct selection of feature configurations. Yet other research has focused on performing some of the configuration implementation work by using model driven development [17].

Regardless of what tools and processes are used to configure SPL variants, however, there is always the possibility that mistakes will occur. For example, large SPLs often use *staged configuration* [8, 9], where features are selected in multiple stages to form a complete configuration iteratively, rather than choosing all features at once. At a late stage in the configuration process, developers may realize that a critically needed feature cannot be selected due to one or numerous decisions in some previous stages. It is hard to debug a configuration to figure out how to change decisions in previous stages to make the critical feature selectable.

Another challenging situation can arise when multiple participants are involved in the feature selection process and their desired feature selections conflict. For example, hardware developers for an automobile may desire a lower cost set of Electronic Control Units (ECUs) that cannot support the features needed by the software developer’s embedded controller code. In these situations, methods are needed to evaluate and debug conflicts between participants. Methods are also needed to recommend modifications to the participants feature selections to make them compatible.

Although prior research has shown how to identify flawed configurations [3, 12], conventional debugging mechanisms cannot pinpoint configuration errors and identifying corrective actions. More specifically, techniques are lacking that can take an arbitrary flawed configuration and produce the minimal set of feature selections and deselections to bring the configuration to a valid state. This paper focuses on addressing these gaps in existing research.

Solution approach → **Constraint-based feature configuration diagnosis.** Our approach to debugging feature model configurations transforms an invalid feature model configuration into a Constraint Satisfaction Problem (CSP) [18] and then uses a constraint solver to derive the minimal set of feature selection modifications that will bring the configuration to a valid state. performance. This

paper shows how this CSP-based diagnostic technique provides the following contributions to work on debugging errors in feature model configurations:

1. We provide a CSP-based [18] diagnostic technique, inspired by the diagnostic framework in [16], that can pinpoint conflicts and constraint violations in feature models
2. We show how this CSP diagnostic technique can be used to remedy a configuration error by automatically deriving the minimal set of features to select and deselect
3. We provide mechanisms for using this CSP technique to create cost-optimal mediation of conflicting configuration participant feature selection needs
4. We show how our CSP diagnostic technique allows stakeholders to debug a configuration error or conflict from different viewpoints
5. We provide empirical results showing that this CSP-based diagnostic technique is scalable enough to support industrial SPL feature models containing over 5,000 features.

The remainder of the paper is organized as follows: Section 2 describes the challenges of diagnosing configuration errors and conflicts in SPLs; Section 3 presents our CSP-based technique for diagnosing configuration errors and conflicts; Section 4 shows how our CSP-based technique can be extended to support conflict mediation, multi-viewpoint debugging, and faster diagnosis times; Section 5 presents empirical results demonstrating the ability of our technique to scale to feature models with thousands of features; Section 6 compares our work with related research; and Section 7 presents concluding remarks.

2 Challenges of Debugging Feature Model Configurations

This section evaluates different challenges that arise in realistic configuration scenarios; Section 3 describes our solutions to these challenges.

2.1 Challenge 1: Staged Configuration Errors

Staged configuration is a configuration process whereby developers iteratively select features to reduce the variability in a feature model until a variant is constructed. Czarnecki et al. [8, 9] use the context of software supply chains for embedded software in automobiles to demonstrate the need for staged configuration. In the first stage, software vendors provide software components that can be provided

in different configurations to actuate brakes, control infotainment systems, etc. In the second stage, hardware vendors of Electronic Control Units (ECUs) the software runs on must provide ECUs with the correct features and configuration to support the software components selected in the first stage.

The challenge with staged configuration is that feature selection decisions made at some point in time T have ramifications on the decisions made at all points in time $T' > T$. For example, it is possible for software vendors to choose a set of software component features for which there are no valid ECU configurations in the second configuration stage. Identifying the fewest number of configuration modifications to remedy the error is hard because there can be significant distance between T and T' .

This challenge also appears in larger models, such as those for software to control the automation of continuous casting in steel manufacture [14]. In large-scale models, configuration mimics staged configuration since developers cannot immediately understand the ramifications of their current decisions. At some later decision point, critical features that developers need may no longer be selectable due to some previous choice. Again, it is hard to identify the minimal set of configuration decisions to reverse in this scenario. Section 4.2 describes how we address this challenge using the pre-assignment of some variables with our CSP-based diagnosis technique.

2.2 Challenge 2: Mediating Conflicts

In many situations the desired features and needs of multiple stakeholders involved in configuring an SPL variant may conflict. For example, when configuring automotive systems, software developers may want a series of software component configurations that cannot be supported by the ECU configurations proposed by the hardware developers. To each party, their individual needs are critical and finding the middle ground to integrate the two is hard.

Another conflict scenario arises when configuration decisions made for an SPL variant must be reconciled with the constraints of the legacy environment in which it will run. For example, when configuring automotive software for the next year’s model of a car, a variant may initially be configured to provide the features most sought after by customers, such as a complex infotainment system. Typically, new model cars are not complete redesigns and thus, developers have to figure out how to run the new software configuration on the existing ECU configuration of the previous year’s model. If the new software configuration is not compatible with the legacy ECU configuration, developers need a way to derive the lowest cost set of modifications to the either the new software or the legacy ECU configuration, which is hard.

Section 4.3 describes how we address this challenge by diagnosing the superset of the desired conflicts and leveraging an alternate optimization goal for the CSP-diagnosis technique.

2.3 Challenge 3: Viewpoint-dependent Errors

The feature labeled as the source of an error in a feature model configuration may vary depending on the viewpoint used to debug it. In the feature model shown in Figure 1, for example, if a configuration is created that includes both *Non-ABS Controller* and *1 Mbit/s CAN Bus*, either feature can be viewed as the feature that is the source of the error. If we debug the configuration from the viewpoint that




Figure 1: Simple Feature Model for an Automobile

software trumps ECU hardware decisions, then the *1 Mbit/s CAN Bus* feature is the error. If we assume that ECU decisions precede software, however, then the *Non-ABS Controller* feature is the error.

A feature model may therefore require debugging from multiple viewpoints since the diagnosis of the feature that is the source of an error in a feature model depends on the viewpoint used to debug it. For small feature models, debugging from different points of view may be relatively simple. When feature models contain hundreds or thousands of features, the complexity of diagnosing a configuration from multiple viewpoints increases greatly. Section 4.2 describes how we address this challenge by specifying feature selections that cannot be modified by the solver during the diagnosis.

3 Configuration Error Diagnosis

The solution we propose is a technique for creating automated SPL variant diagnosis tools. Developers can use these tools to identify the minimal set of features that should be selected or deselected to transform an invalid configuration to a valid configuration. Moreover, depending on the input provided to diagnosis tools, a flawed configuration can be debugged from different viewpoints or mediate conflicts between multiple stakeholder decisions in a configuration process.

The key component of our automated SPL variant conflict diagnosis technique is the application of a CSP-based error diagnostic technique. In prior work, Benavides et al. [4] have shown how feature models can be transformed into CSPs to automate feature selection with a constraint solver [10]. Trinidad et al. [16] subsequently showed to extend this CSP technique to identify *full mandatory features*, *void features*, and *dead feature models* using Reiter’s theory of diagnosis [15]. This section presents an alternate diagnostic model for deriving the minimum set of features that should be selected or deselected to eliminate a conflict in a feature configuration.

Background: Feature Models and Configurations as CSPs. A CSP is a set of variables and a set of constraints over those variables. For example, $A + B \leq 3$ is a CSP involving the integer variables A and B . The goal of a constraint solver is to find a valid *labeling* (set of variable values) that simultaneously satisfies all constraints in the CSP. ($A = 1, B = 2$) is thus a valid labeling of the CSP.

To build the CSP for the error diagnosis technique, we construct a set of variables, F , representing the features in the feature model. Each configuration of the feature model is a set of values for these variables, where a value of 1 indicates the feature is present in the configuration and a value of 0 indicates it is not present. More formally, a configuration is a labeling of F , such that for each variable $f_i \in F$, $f_i = 1$ indicates that the i_{th} feature in the feature model is selected in the configuration. Correspondingly, $f_i = 0$ implies that the feature is not selected.

Given an arbitrary configuration of a feature model as a labeling of the F variables, developers need the ability to ensure the correctness of the configuration. To achieve this constraint checking ability, each variable f_i is associated with one or more constraints corresponding to the configuration rules in the feature model. For example, if f_j is a required subfeature of f_i , then the CSP would contain the constraint: $f_i = 1 \Leftrightarrow f_j = 1$.

Configuration rules from the feature model are captured in the constraint set C . For any given feature model configuration described by a labeling of F , the correctness of the configuration can be determined by seeing if the labeling satisfies all constraints in C . The steps of transforming a feature model to a CSP are described in [4].

3.1 Configuration Diagnostic CSP

When diagnosing configuration conflicts, developers need a list of features that should be selected or deselected to make an invalid configuration a valid configuration. The output of our CSP diagnostic architecture is this list of features to select and deselect, as shown in Figure 2. In Step 1 of Figure 2, the rules of the feature model and the current

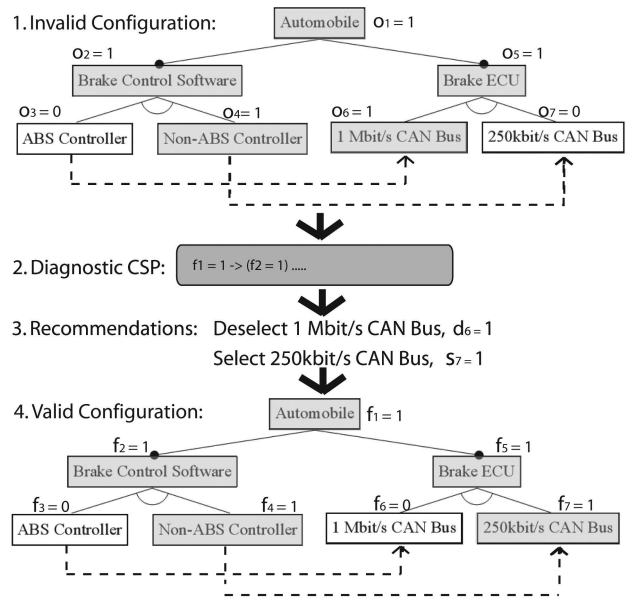


Figure 2: Diagnostic Technique Architecture

invalid configuration are transformed into a CSP. For example, $o_1 = 1$ because the *Automobile* feature is selected in the current invalid configuration. In Step 2, the solver derives a labeling of the diagnostic CSP. Step 3 takes the output of the CSP labeling and transforms it into a series of recommendations of features to select or deselect to turn the invalid configuration into a valid configuration. Finally, in Step 4, the recommendations are applied to the invalid configuration to create a valid configuration where each variable f_i equals 1 if the corresponding feature is selected in the new and valid configuration. For example, $f_7 = 1$, meaning that the *250 Kbit/s CAN Bus* is selected in the new valid configuration.

To enable the constraint solver to recommend features to select and deselect, two new sets of recommendation variables, S and D , are introduced to capture the features that need to be selected and deselected, respectively, to reach a valid configuration. For example, a value of 1 for variable $s_i \in S$ indicates that the feature f_i should be added to the current configuration. Similarly, $d_i = 1$ implies that the feature f_i should be removed from the configuration.

Thus, for each feature $f_i \in F$, there are variables $s_i \in S$ and $d_i \in D$. After the diagnosis CSP is labeled, the values of S and D serve as the output recommendations to the user as to what features to add or remove from the current configuration, as shown in Table 1.

Table 1 shows the complete inputs and outputs to diagnose the invalid configuration scenario shown in Figure 2.

The next step is to allow developers to input their current

Variables	
Variable Explanations	$f_i \subset F$: feature variables for the valid configuration that will be transitioned to; $o_i \subset O$: the features selected ($o_i = 1$) in the current invalid configuration; $s_i \subset S$: features to select ($s_i = 1$) to reach the valid configuration; $d_i \subset D$: features to deselect ($d_i = 1$) to reach the valid configuration
Inputs	
Current Config.	$o_1 = 1, o_2 = 1, o_3 = 0, o_4 = 1, o_5 = 1, o_6 = 1, o_7 = 0$
Feature Model Rules	$f_1 = 1 \Leftrightarrow (f_2 = 1), f_1 = 1 \Leftrightarrow (f_5 = 1), f_2 = 1 \Rightarrow (f_3 = 1) \oplus (f_4 = 1), f_5 = 1 \Rightarrow (f_6 = 1) \oplus (f_7 = 1), (f_6 = 1) \vee (f_7 = 1) \Rightarrow (f_5 = 1), (f_3 = 1) \vee (f_4 = 1) \Rightarrow (f_2 = 1), f_3 = 1 \Rightarrow (f_6 = 1), f_4 = 1 \Rightarrow (f_7 = 1)$
Diagnostic Rules	$(f_i \subset F \mid \{(f_i = 1) \Rightarrow (o_i = 1 \oplus s_i = 1) \wedge (d_i = 0), (f_i = 0) \Rightarrow (o_i = 0 \oplus d_i = 1) \wedge (s_i = 0)\})$
Outputs	
Features to Select	$s_1 = 0, s_2 = 0, s_3 = 0, s_4 = 0, s_5 = 0, s_6 = 0, \mathbf{s_7 = 1}$
Features to Deselect	$d_1 = 0, d_2 = 0, d_3 = 0, d_4 = 0, d_5 = 0, \mathbf{d_6 = 1}, d_7 = 0$
New Valid Config.	$f_1 = 1, f_2 = 1, f_3 = 0, f_4 = 1, f_5 = 1, f_6 = 0, f_7 = 1$

Table 1: Diagnostic CSP Construction

configuration into the solver for diagnosis. Rather than directly setting values for the variables in F , developers use a special set of input variables called the *observations*. The observations are contained in the set of variables O . For each feature f_i present in the current flawed configuration, $o_i = 1$. Similarly, if f_i is not selected in the current invalid configuration, $o_i = 0$. As shown in Table 1, the observations capture the current invalid configuration that is provided as input to the solver.

To diagnose the CSP, we want to find an alternate but valid configuration of the feature model and suggest a series of changes to the current invalid configuration to reach the valid configuration. A valid configuration is a labeling of the variables in F (a configuration) such that all of the feature model constraints are satisfied. For each variable f_i , the value should be 1 if the feature is present in the new valid configuration that will be transitioned to. If a feature is not in the new configuration, f_i should equal 0. We always require $f_1 = 1$ to ensure that the root feature is always selected. It also must be noted that the technique only works with non-void feature models.

If feature f_i is included in the current configuration ($o_i = 1$), but must be removed to reach the new valid configuration, we want the solver to recommend that it be deselected ($d_i = 1$). Moreover, the solver should recommend that a feature f_i be selected ($s_i = 1$) if it is not currently selected ($o_i = 0$), but needs to be selected to reach a correct configuration. If a feature is not in the current configuration ($o_i = 0$), and is not needed to reach a correct state, then the solver should not recommend any changes to it ($s_i = 0$ and $d_i = 0$). To produce this desired behavior, we introduce two new constraints for each variable $f_i \subset F$: $(f_i = 1) \Rightarrow (o_i = 1 \oplus s_i = 1) \wedge (d_i = 0)$ and $(f_i = 0) \Rightarrow (o_i = 0 \oplus d_i = 1) \wedge (s_i = 0)$. The symbol \oplus represents *exclusive or* (XOR). These new constraints form the diagnostic rules, as shown in Table 1.

The behavior produced by these new constraints explains the following four cases that a feature modification can fall into:

1. **A feature is selected and does not need to be deselected.** If the i_{th} feature is in the current invalid configuration ($o_i = 1$), and also in the new valid configuration ($f_i = 1$), no changes need be made to it ($s_i = 0, d_i = 0$) $\Rightarrow \{(f_i = 1) \Rightarrow (o_i = 1 \oplus s_i = 0) \wedge (d_i = 0)\}$
2. **A feature is selected and needs to be deselected.** If the i_{th} feature is in the current invalid configuration ($o_i = 1$) but not in the new valid configuration ($f_i = 0$), it must be deselected ($d_i = 1$) $\Rightarrow \{(f_i = 0) \Rightarrow (o_i = 0 \oplus d_i = 1) \wedge (s_i = 0)\}$
3. **A feature is not selected and does not to be selected.** If the i_{th} feature is not in the current invalid configuration ($o_i = 0$) and is also not needed in the new configuration ($f_i = 0$) it should remain unchanged ($s_i = 0, d_i = 0$) $\Rightarrow \{(f_i = 0) \Rightarrow (o_i = 0 \oplus d_i = 1) \wedge (s_i = 0)\}$
4. **A feature is not selected and needs to be selected.** If the i_{th} feature is not selected in the current invalid configuration ($o_i = 0$) but is present in the new correct configuration $f_i = 1$, it must be selected ($s_i = 1$) $\Rightarrow \{(f_i = 1) \Rightarrow (o_i = 1 \oplus s_i = 0) \wedge (d_i = 0) = 1\}$

3.2 Optimal Diagnosis Method

The next step in the diagnosis process is to use the solver to find a labeling of the variables and produce a series of recommendations. For any given configuration with a conflict, there may be multiple possible ways to eliminate the problem. For example, in the automotive example from Section 2.3, the valid corrective actions were to either remove the *1 Mbit/s CAN Bus* and select the *250 Kbit/s CAN Bus* or to remove *Non-ABS Controller* and select *ABS Controller*. We therefore need to tell the solver how to select which of

the (many) possible corrective solutions to suggest to developers.

The most basic suggestion selection criteria that developers can use to guide the solver’s diagnosis is to tell it to minimize the number of changes that to make to the current configuration, *i.e.*, developers prefer suggestions that require changing as few things as possible in the current invalid configuration. To achieve this, we solve for a labeling of the CSP that minimizes the sum of the variables in $S \cup D$. The sum of $S \cup D$ is the total number of changes that the solution requires the developer to make. By minimizing this sum, therefore, we minimize the total number of required changes.

Each labeling of the diagnostic CSP will produce two sets of features corresponding to the features that should be selected (S) and deselected (D) to reach the new valid configuration. Furthermore, the labeling also causes the solver to backtrack and create new values for F corresponding to the proposed solution to transition to. Thus, developers can cycle through the different potential labelings of the diagnostic CSP to evaluate potential remedies and the configurations they will produce.

Table 1 shows a complete set of inputs and output suggestions for diagnosing the automotive software example from Section 2.3. If there are multiple labelings of the CSP, initially only one will be returned. After the first solution has been found, however, the solver can much more efficiently cycle through the other equally ranked sets of corrective suggestions.

4 Solution Extensibility and Benefits

This section presents different benefits of our CSP-based diagnostic approach and possible ways of extending it.

4.1 Bounding Diagnostic Method

For extremely large feature models, finding the optimal number of changes may not be possible due to time constraints. In these cases, an alternate, more scalable, approach is to attempt to find any suggestion that requires fewer than K changes or with a cost less than K . Rather than directly asking for an optimal answer, we instead add the constraint $\sum_{i=1}^n s_i + d_i \leq K$ to the CSP and ask the solver for any solution.

The sum of all variables $s_i \in S$ and $d_i \in D$ represents the total number of feature selections and deselections that need to be made to reach the new valid configuration. Therefore, the sum of both of these sets is the total number of modifications that must be made to the original invalid configuration. The new constraint, $\sum_{i=1}^n s_i + d_i \leq K$, ensures that the solver only accepts diagnosis solutions that require the developer to make K or fewer changes to the invalid solution.

The solver is asked for **any** answer that meets the new constraints. In return, the solver will provide a solution that is not necessarily perfect, but which fits our tolerance for change. If no solution is found, we can increment K by a factor and re-invoke the solver or reassess our requirements. As is shown in Section 5.4, searching for a bounded solution rather than an optimal solution is significantly faster.

If the solver cannot find a diagnosis that makes fewer than K modifications, it will state that there is no valid solution that fits a K change budget. From the experiments performed in Section 5, we found that in these cases where no solution can be found, the solver was able to diagnose the model much faster than if there was a viable solution under K . Furthermore, the more cross-tree constraints that were within the feature model, the faster the solver was able to determine if there was no solution within the bound K .

4.2 Debugging from Different Viewpoints

As we discussed in Section 2.3, we need the ability to debug the configuration from different viewpoints. Each viewpoint represents a set of features that the solver should avoid suggesting to add or remove from the current configuration. For example, using the automobile scenario from Section 2.3, the solver can debug the problem from the point of view that audio trumps storage by telling the solver not to suggest selecting or deselecting any audio features.

Debugging from a viewpoint works by pre-assigning values for a subset of the variables in F and O . For example, to force the feature f_i currently in the configuration to remain unaltered by the diagnosis, the values $f_i = 1$ and $o_i = 1$ are provided to the solver. Since $(f_i = 1) \Rightarrow (o_i = 1 \oplus s_i = 1) \wedge (d_i = 0)$, pre-assigning these values will force the solver to label $s_i = 0$ and $d_i = 0$.

To debug from a given point of view, for each feature f_v , in that viewpoint, we first add the constraints, $f_v = 1$, $o_v = 1$, $s_v = 0$, and $d_v = 0$. The solver then derives a diagnosis that recommends alterations to other features in the configuration and maintains the state of each feature f_v . The diagnostic model can therefore be used to debug from different viewpoints and address Challenge 3 from Section 2.3.

Pre-assigning values for variables in F and O can also be used to debug staged configuration errors from Challenge 1, Section 2.1. With staged configuration errors, at some point in time T' , developers attempt to select a feature that is in conflict with another feature selected at time $T < T'$. To debug this type of conflict, developers pre-assign the desired (but currently unselectable) feature at time T' the value of 1 for its o_i and f_i variables. Developers can also pre-assign values for one or more other features decisions from previous stages of the configuration that must not be altered. The solver is then invoked to find a configuration that includes the desired feature at T' and minimizes the

number of changes to feature configuration decisions that were made at all points in time $T < T'$.

4.3 Cost Optimal Conflict Resolution

As shown in Section 2.2, conflicts can occur when multiple stakeholders in a configuration process pull the solution in different directions. In these situations, debugging tools are needed to mediate the conflict in a cost conscious manner. For example, when the software configuration of a car is not compatible with the legacy ECU configuration, it is (probably) cheaper to recommend changes to the software configuration than to change the ECU configuration and possibly the assembly process of the car. Therefore, the solver should try to minimize the overall cost of the changes.

We can extend the CSP model to perform cost-based feature selection and deselection optimization. First, we extend the model to associate a cost variable, $b_i \in B$, with each feature in the feature model. Each cost variable represents how expensive (or conversely how beneficial) it is for the solver to recommend that the state of that feature be changed. Before each invocation of the debugger, the stakeholders provide these cost variables to guide the solver in its recommendations of features to select or deselect.

Next, we construct the superset of the features that the various stakeholders desire. The superset represents the ideal, although incorrect, configuration that the stakeholders would like to have. The goal is to find a way to reach a correct configuration from this superset of features that involves the lowest total cost for changes. The superset is input to the solver as values for the variables in O .

Finally, we alter our original optimization goal so that the solver will attempt to minimize (or maximize if we prefer) the cost of the features it suggests selecting or deselecting. We define a new global cost variable G and let G capture the sum of the costs of the changes that the solver suggests: $G = \sum_{i=1}^n (d_i * b_i) + (s_i * b_i)$. G is thus equal to the sum of the costs of all the features that the solver either recommends to select or deselect. Rather than simply instructing the solver to minimize the sum of $S \cup D$, we instead ask for the minimization or maximization of G .

The result of the labeling is a series of changes needed to reach a valid configuration that optimally integrates the desires and decisions of the various stakeholders. Of course, one particular stakeholder may have to incur more cost than another in the interest of reaching a globally better solution. Further constraints, such as limiting the maximum difference between the cost incurred by any two stakeholders, could also be added. The mediation process can be tuned to provide numerous types of behavior by providing different optimization goals. This CSP diagnostic method enables us to address Challenge 2 from Section 2.2.

5 Empirical Results

Automated diagnostic methods are of little use if they cannot scale to handle non-trivial feature models of production systems. This section presents empirical results from experiments we performed to evaluate the scalability of our CSP-based diagnosis technique. We compare both the optimal and bounding methods for diagnosing errors and conflicts.

5.1 Experimental Platform

To perform our experiments, we used the implementation of the diagnosis algorithm which is provided by the Model Intelligence libraries from the Eclipse Foundation's Generic Eclipse Modeling System (GEMS) project [2]. Internally, the GEMS Model Intelligence implementation uses the Java Choco Constraint Solver [1] to derive labelings of the diagnostic CSP. The experiments were performed on a computer with an Intel Core DUO 2.4GHZ CPU, 2 gigabytes of memory, Windows XP, and a version 1.6 Java Virtual Machine (JVM). The JVM was run in client mode using a heap size of 40 megabytes (-Xms40m) and a maximum memory size of 256 megabytes (-Xmx256m).

A challenging aspect of the scalability analysis is that CSP-based techniques can vary in solving time based on individual problem characteristics. In theory, CSP's have exponential worst case time complexity, but in practice they are often much faster. To demonstrate a high confidence in our technique, it was necessary to apply our technique to as many models as possible. The key challenge with this approach is that hundreds or thousands of real feature models are not readily available and manually constructing them is not practical.

To provide the large numbers of feature models needed for our experiments, therefore, we built a feature model generator. This generator can randomly generate feature models with the desired branching and constraint characteristics. We also imbued the generator with the capability to generate feature selections from a feature model and probabilistically insert a bounded number of errors/conflicts into the configuration. The feature model generator is available in open-source form from [].

From preliminary feasibility experiments we conducted, we observed that the branching factor of the tree had little effect on the algorithm's solving time. We also compared diagnosis time using models with 0%, 10%, and 50% cross-tree constraints and saw no difference in diagnosis time. The key indicator of the solving complexity was the number of XOR- or cardinality-based feature groups in a model. XOR and cardinality-based feature groups are features that require the set of their selected children to satisfy a cardinality constraint (the constraint is 1..1 for XOR).

For our tests, we limited the branching factor to at most five subfeatures per feature. We also set the probability of XOR- or cardinality-based feature groups being generated to 1/3 at each feature with children. We chose 1/3 since most feature models we have encountered contain more required and optional relationships than XOR- and cardinality-based feature groups.

To generate feature selections with errors, we used a probability of 1/50 that any particular feature would be configured incorrectly. For each model, we bounded the total errors at 5. In our initial experiments, the solving time was not affected by the number of errors in a given feature model. Again, the prevalence of XOR- or cardinality-based feature groups was the key determiner of solving time.

5.2 Bounding Method Scalability

First, we tested the scalability of the less computationally complex bounding diagnosis method. The speed of our technique allowed us to test 2,000 feature models at each data point (2,000 different variations of each size feature model) and test the bounding method’s scalability for feature models up to 500 features. With models above 500 features, we had to reduce the number of samples at each size to 200 models due to time constraints. Although these samples are small, they demonstrate the general performance of our technique. Moreover, the results of our experiments with feature models up to 500 features were nearly identical with sample sizes between 100 and 2,000 features.

Figure 3 shows the time required to diagnose feature models ranging in size from 50 to 500 features using the bounded method. The figure captures the worst, average,




Figure 3: Diagnosis Time with Bounding Method

and best solving time in the experiments. As seen from the results, our technique could diagnose 500 feature models in an average of ≈ 600 ms.

The upper bound used for this experiment was a maximum of 6 feature selection changes. When the 6 feature bound was too tight for the diagnosis (*i.e.*, more than 6 changes were needed to reach a correct state) the solver would declare that there was no valid solution in a very short amount of time. We therefore discarded all instances where

the bound was too tight to avoid skewing the results towards shorter solving times.

Figure 4 shows the results of testing the solving time of the bounding method on feature models ranging in size from 750 to 5,000 features. Models of this size were sufficient to

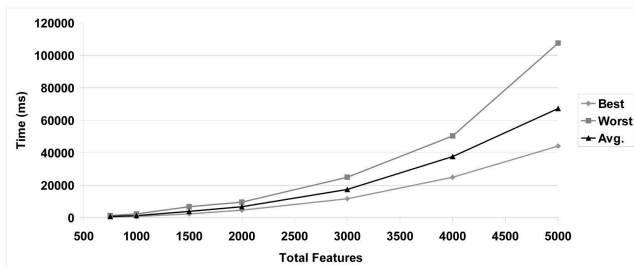


Figure 4: Diagnosis Time Using Bounding Method for Large Feature Models

demonstrate scalability for production systems. The results show that for a 5,000 feature model, the average diagnosis time was ≈ 1 minute.

Another key variable that we tested was how the tightness of the bound on the maximum number of feature changes affected the solving time of the technique. We took a set of 200 feature models and applied varying bounds to see how the bound tightness affected solution time. Figure 5 shows that tighter bounds produced faster solution times. We hypothesize that the tighter bound allows the solver to

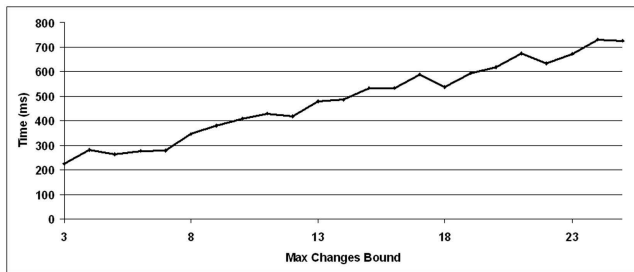


Figure 5: 500 Feature Diagnosis Time with Bounding Method and Varying Bounds

discard infeasible solutions more quickly and thus arrive at a solution faster.

5.3 Optimal Method Scalability

Next, we tested the scalability of the optimal diagnosis method. For the tests of the optimal method’s scalability for feature models up to 2,000 features, the slower speed of the technique allowed us to test 200 feature models at each data point. We stopped at 2,000 features because by that data point, each diagnosis required an average of ≈ 7.5 minutes and the 200 samples required over 24 hours to test.

The results from feature models up to 500 features can be seen in Figure 6. At 500 features, the optimal method




Figure 6: Diagnosis Time with Optimal Method

required an average of roughly 8 seconds to produce a diagnosis. The tests from larger models ranging in size up to 2,000 features are shown in Figure 7. For a model with




Figure 7: Diagnosis Time with Optimal Method for Large Feature Models

2,000 features, the solver required an average of approximately 7.5 minutes per diagnosis.

5.4 Comparative Analysis of Optimal and Bounding Methods

Finally, we compared the scalability and quality of the results produced with the two methods. As seen by comparing Figures 3 and 6 the bounding method performs and scales significantly better than the optimal method. This result raises the key question of how much of a tradeoff in solution quality for speed is made when the bounding method is used over the optimal method.

The bound that is chosen determines the quality of the solution that is produced by the solver. We can state the optimality of a diagnosis given by the bounding method as the number of changes suggested by the bounding method, $Bounded(S \cup D)$, divided by the optimal number of changes, $Opt(S \cup D)$, which yields $\frac{Bounded(S \cup D)}{Opt(S \cup D)}$. Since the bounding method uses the constraint $(S \cup D) \leq K$ to ensure that at most K changes are suggested, we can state the worst case optimality of the bounded method as $\frac{K}{Opt(S \cup D)}$. The closer

our bound, K , is to the true optimal number of changes to make, the better the diagnosis will be.

Since tighter bounds produce faster solving times *and* better results, debuggers should start with very small bounds and iteratively increase them upward as needed. One approach is to layer an adaptive algorithm on top of the diagnosis algorithm to move the bound by varying amounts each time the bound proved too tight. Another approach is to employ a form of binary search to hone in on the ideal bound. We will be investigating both of these techniques in future work.

6 Related Work

In prior work [16], Trinidad et al. have shown how feature models can be transformed into diagnosis CSPs and used to identify *full mandatory features*, *void features*, and *dead feature models* [16]. Developers can use this diagnostic capability to identify feature models that do not accurately describe their products and to understand why not. The technique we described in this paper builds on this idea of using a CSP for automated diagnosis. Whereas Trinidad focuses on diagnosing feature models that do not describe their products, however, we build an alternate diagnosis model to identify conflicts in feature configurations. Moreover, we provide specific recommendations as to the minimal set of features that can be selected or deselected to eliminate the error.

Batory et al. [3] also investigated debugging techniques for feature models. Their techniques focus on translating feature models into propositional logic and using SAT solvers to automate configuration and verify correctness of configurations. In general, their work touches on debugging feature models rather than individual configurations. Our approach focuses on another dimension of debugging, the ability to pinpoint errors in individual configurations and to specify the minimal set of feature selections and deselections to remove the error. Moreover, we show how automated debugging can be combined with probing to help understand and integrate bottom-up forces into the configuration process.

Mannion et al. [12] present a method for encoding feature models as propositional formulas using first-order logic. These propositional formulas can then be used to check the correctness of a configuration. Mannion, however, does not touch on how incorrect configurations are debugged. In contrast, our technique provides this capability and can therefore recommend the minimal feature modifications to rectify the problem.

Pure::variants [5], Feature Modeling Plugin (FMP) [7], and Big Lever Software Gears [6] are tools developed to help developers create correct configurations of SPL feature models. These tools enforce constraints on modelers

as the features are selected. None of these tools, however, addresses cases where feature models with incorrect configurations are created and require debugging. The technique described in this paper provides this missing capability. These tools and our approach are complementary since the tools help to ensure that correct configurations are created and our technique diagnoses incorrect configurations that are built.

7 Concluding Remarks

It is hard to debug conflicts and errors in large feature models created through staged or multi-stakeholder configuration. This paper described how CSPs can be built to diagnose errors and conflicts in configurations. The diagnosis provided by a CSP can specifically recommend the minimum or cost optimal set of features that should be selected or deselected in a faulty configuration.

The following are lessons learned from our efforts thus far:

- By applying bounding and optimal diagnosis, our technique can scale to feature models with several thousand features.
- In future work, we plan to investigate different methods of honing the boundary used in the bounding method. We also intend to investigate the diagnosis of configuration errors in SPLs that use multiple models to capture different viewpoints of the SPL.

References

- [1] Choco constraint programming system. <http://choco.sourceforge.net/>.
- [2] Generic eclipse modeling system (gems) <http://eclipse.org/gmt/gems>.
- [3] D. Batory. Feature Models, Grammars, and Propositional Formulas. *Software Product Lines: 9th International Conference, SPLC 2005, Rennes, France, September 26-29, 2005: Proceedings*, 2005.
- [4] D. Benavides, P. Trinidad, and A. Ruiz-Cortés. Automated Reasoning on Feature Models. *17th Conference on Advanced Information Systems Engineering (CAiSE05, Proceedings)*, LNCS, 3520:491–503, 2005.
- [5] D. Beuche. Variant Management with Pure::variants. Technical report, Pure-Systems GmbH, <http://www.pure-systems.com>, 2003.
- [6] R. Buhrdorf, D. Churchett, and C. Krueger. Salion’s Experience with a Reactive Software Product Line Approach. In *Proceedings of the 5th International Workshop on Product Family Engineering*, Siena, Italy, November 2003.
- [7] K. Czarnecki, M. Antkiewicz, C. Kim, S. Lau, and K. Pietroszek. In *FMP and FMP2RSM: Eclipse Plug-ins for Modeling Features Using Model Templates*, pages 200–201. ACM Press New York, NY, USA, October 2005.
- [8] K. Czarnecki, S. Helsen, and U. Eisenecker. Staged Configuration Using Feature Models. *Software Product Lines: Third International Conference, SPLC 2004, Boston, MA, USA, August 30-September 2, 2004: Proceedings*, 2004.
- [9] K. Czarnecki, S. Helsen, and U. Eisenecker. Staged configuration through specialization and multi-level configuration of feature models. *Software Process Improvement and Practice*, 10(2):143–169, 2005.
- [10] J. Jaffar and M. Maher. Constraint Logic Programming: A Survey. *constraints*, 2(2):0.
- [11] K. C. Kang, S. Kim, J. Lee, K. Kim, E. Shin, and M. Huh. FORM: A Feature-Oriented Reuse Method with Domain-specific Reference Architectures. *Annals of Software Engineering*, 5(0):143–168, January 1998.
- [12] M. Mannion. Using first-order logic for product line model validation. *Proceedings of the Second International Conference on Software Product Lines*, 2379:176–187, 2002.
- [13] A. Metzger, K. Pohl, P. Heymans, P.-Y. Schobbens, and G. Saval. Disambiguating the documentation of variability in software product lines: A separation of concerns, formalization and automated analysis. In *Requirements Engineering Conference, 2007. RE ’07. 15th IEEE International*, pages 243–253, 2007.
- [14] R. Rabiser, P. Grunbacher, and D. Dhungana. Supporting Product Derivation by Adapting and Augmenting Variability Models. *Software Product Line Conference, 2007. SPLC 2007. 11th International*, pages 141–150, 2007.
- [15] R. Reiter. A theory of diagnosis from first principles. *Artificial Intelligence*, 32(1):57–95, 1987.
- [16] P. Trinidad, D. Benavides, A. Durán, A. Ruiz-Cortés, and M. Toro. Automated error analysis for the agilization of feature modeling. *Journal of Systems and Software*, in press, 2007.
- [17] S. Trujillo, D. Batory, and O. Diaz. Feature Oriented Model Driven Development: A Case Study for Portlets. *29th International Conference on Software Engineering (ICSE 2007)*, Minneapolis, Minnesota, USA, May 20, 26:61, 2007.
- [18] P. Van Hentenryck. *Constraint Satisfaction in Logic Programming*. MIT Press Cambridge, MA, USA, 1989.