

Feature Models, Grammars, and Propositional Formulas

Don Batory
Department of Computer Sciences
University of Texas at Austin
Austin, Texas 78712
`batory@cs.utexas.edu`

Abstract. Feature models are used to specify members of a product-line. Despite years of progress, contemporary tools provide limited support for feature constraints and offer little or no support for debugging feature models. We integrate prior results to connect feature models, grammars, and propositional formulas. This connection allows arbitrary propositional constraints to be defined among features and enables off-the-shelf satisfiability solvers to debug feature models. We also show how our ideas can generalize recent results on the staged configuration of feature models.

1 Introduction

A key technical innovation of software product-lines is the use of features to distinguish product-line members. A *feature* is an increment in program functionality [27]. A particular product-line member is defined by a unique combination of features. The set of all legal feature combinations defines the set of product-line members [22].

Feature models define features and their usage constraints in product-lines [10][18][21]. Current methodologies organize features into a tree, called a *feature diagram (FD)*, which is used to declaratively specify product-line members [3]. Relationships among FDs and grammars [19][11], and FDs and formal models/logic programming [7][24][25] have been noted in the past, but the potential of their integration is not yet fully realized.

Despite progress, tools for feature models seem ad hoc; they exhibit odd limitations and provide little or no support for debugging feature models. This is to be expected when a formal underpinning of feature models is lacking. In this paper, we integrate prior results to connect FDs, grammars, and propositional formulas. This connection enables general-purpose, light-weight, and efficient *logic truth maintenance systems (LTMSs)*[15] to propagate constraints as users select features so that inconsistent product specifications are avoided — much like syntax-directed editors guarantee compilable programs [26]. This connection also allows us to use off-the-shelf tools, called *satisfiability solvers* or *SAT solvers* [14], to help debug feature models by confirming compatible and incomplete feature sets. To our knowledge, the use of LTMSs and SAT solvers in feature modeling tools is novel.

Our approach is tutorial. We believe it is important that researchers and practitioners clearly see the fundamental underpinnings of feature models, and that light-weight and easy-to-build LTMS algorithms and easy-to-use SAT solvers can help address key weaknesses in existing feature model tools and theories.

2 Feature Models

A *feature model* is a hierarchically arranged set of features. Relationships between a *parent* (or *compound*) feature and its *child* features (or *subfeatures*) are categorized as:

- *And* — all subfeatures must be selected,
- *Alternative* — only one subfeature can be selected,
- *Or* — one or more can be selected,
- *Mandatory* — features that required, and
- *Optional* — features that are optional.

Or relationships can have $n:m$ cardinalities: a minimum of n features and at most m features can be selected [10]. More elaborate cardinalities are possible [11].

A *feature diagram* is a graphical representation of a feature model [22]. It is a tree where primitive features are leaves and compound features are interior nodes. Common graphical notations are depicted in Figure 1.




Figure 1 Feature Diagram Notations

Figure 2a is a feature diagram. It defines a product-line where each application contains two features r and s , where r is an alternative feature: only one of G , H , and I can be present in an application. s is a compound feature that consists of mandatory features A and C , and optional feature B .




Figure 2 A Feature Diagram and its Grammar

The connection between FDs and grammars is due to de Jong and Visser [19]. We will use iterative tree grammars. An *iterative grammar* uses iteration (e.g., one-or-more t^+ and zero-or-more t^* constructs) rather than recursion, to express repetition. A *tree grammar* requires every token to appear in exactly one pattern, and the name of every production to appear in exactly one pattern. The root production is an exception; it is not referenced in any pattern. More general grammars can be used, but iterative trees capture the minimum properties needed for our discussions.

Figure 3 enumerates the hierarchical relationships that can be expressed in a feature diagram. Each has a straightforward iterative tree grammar representation:

- Figure 3a is the production $s: e_1 e_2 \dots e_n$ assuming all subfeatures are mandatory. If a subfeature is optional (as is e_2), it is surrounded by [brackets]. Thus, the production for Figure 3a is $s: e_1 [e_2] \dots e_n$.
- Figure 3b is the production: $s: e_1 | e_2 | \dots | e_n$.

- Figure 3c corresponds to a pair of rules: $s:t+$; and $t:e_1 | e_2 | \dots | e_n$; meaning one or more of the e_i are to be selected.
- Figure 3d is the production: $s:e_1^*$, meaning zero or more e_1 can be selected,
- and Figure 3e is the production: $s:e_1+$, meaning one or more e_1 can be selected.




Figure 3 Parent-Child Relationships in FDs

In general, each non-terminal node of a feature diagram is a production. The root is the start production; leaves are tokens. Figure 2b is the grammar of Figure 2a. An application defined by the feature diagram of Figure 2a is a sentence of this grammar.


Henceforth, we use the following notation for grammars. Tokens are **UPPERCASE** and non-terminals are **lowercase**. $r+$ denotes one or more instances of non-terminal r ; r^* denotes zero or more. $[r]$ and $[R]$ denote optional non-terminal r and optional token R . A pattern is a named sequence of (possibly optional or repeating) non-terminals and (possibly optional) terminals. Consider the production:

```

r : b+ A C      :: First
  | [D] E F     :: Second ;

```

The name of this production is r ; it has two patterns **First** and **Second**. The **First** pattern has one or more instances of b followed by terminals A and C . The **Second** pattern has optional token D followed by terminals E and F .



Grammars provide a graphics-neutral representation of feature models. For example, the grammar of Figure 2b could be displayed by the tree of Figure 2a or the GUI of Figure 4. (The GUI doesn't display features E and F , as they are mandatory — nothing needs to be selected). A popular Eclipse plug-in provides other possible graphical representations of FDs (tree and wizard-based), all of which are derived from a grammar-like specification [3].

In the next section, we show how iterative tree grammars (or equivalently feature diagrams) are mapped to propositional formulas.

3 Propositional Formulas

Mannion was the first to connect propositional formulas to product-lines [24]; we show how his results integrate with those of Section 2. A *propositional formula* is a set of boolean variables and a propositional logic predicate that constrains the values of these variables. Besides the standard \wedge , \vee , \neg , \Rightarrow , and \Leftrightarrow operations of propositional

logic, we also use $\text{choose}_1(e_1 \dots e_k)$ to mean at most one of the expressions $e_1 \dots e_k$ is true. More generally, $\text{choose}_{n,m}(e_1 \dots e_k)$ means at least n and at most m of the expressions $e_1 \dots e_k$ are true, where $0 \leq n \leq m \leq k$.

A grammar is a compact representation of a propositional formula. A variable of the formula is either: a token, the name of a non-terminal, or the name of a pattern. For example, the production:

$$\begin{array}{l} \mathbf{r} : \mathbf{A} \mathbf{B} \quad \quad \quad :: \mathbf{P1} \\ \quad \quad \quad \quad \quad \quad \quad | \mathbf{C} [\mathbf{r1}] \quad \quad :: \mathbf{P2} ; \end{array} \quad (1)$$

has seven variables: three $\{\mathbf{A}, \mathbf{B}, \mathbf{C}\}$ are tokens, two are non-terminals $\{\mathbf{r}, \mathbf{r1}\}$, and two are names of patterns $\{\mathbf{P1}, \mathbf{P2}\}$. Given these variables, the rules for mapping a grammar to a propositional formula are straightforward.

Mapping Productions. Consider production $\mathbf{r} : \mathbf{P}_1 | \dots | \mathbf{P}_n$, which has n patterns $\mathbf{P}_1 \dots \mathbf{P}_n$. Production \mathbf{r} can be referenced in one of three ways: \mathbf{r} (choose one), $\mathbf{r+}$ (choose one or more), and $\mathbf{r*}$ (choose zero or more). As $\mathbf{r*}$ can be encoded as $[\mathbf{r+}]$ (optionally choose one or more), there are only two basic references: \mathbf{r} and $\mathbf{r+}$. The propositional formulas for both are listed below.

Pattern	Formula
\mathbf{r}	$\mathbf{r} \Leftrightarrow \text{choose}_1(\mathbf{P}_1, \dots, \mathbf{P}_n)$
$\mathbf{r+}$	$\mathbf{r} \Leftrightarrow (\mathbf{P}_1 \vee \dots \vee \mathbf{P}_n)$

Mapping Patterns. A *basic term* is either a token or a production reference. A *pattern* is a sequence of one or more basic terms or optional basic terms. Let \mathbf{P} be the name of a pattern and let $\mathbf{t}_1 \dots \mathbf{t}_n$ be a sequence of basic terms. The formula for \mathbf{P} is:

$$\mathbf{P} \Leftrightarrow \mathbf{t}_1 \wedge \mathbf{P} \Leftrightarrow \mathbf{t}_2 \wedge \dots \wedge \mathbf{P} \Leftrightarrow \mathbf{t}_n \quad (2)$$

That is, if \mathbf{P} is included in a design then terms $\mathbf{t}_1 \dots \mathbf{t}_n$ are also included, and vice versa. Consider pattern \mathbf{Q} whose second term is optional: $\mathbf{t}_1 [\mathbf{t}_2] \dots \mathbf{t}_n$. The formula for \mathbf{Q} is:

$$\mathbf{Q} \Leftrightarrow \mathbf{t}_1 \wedge \mathbf{t}_2 \Rightarrow \mathbf{Q} \wedge \dots \wedge \mathbf{Q} \Leftrightarrow \mathbf{t}_n \quad (3)$$

That is, if \mathbf{Q} is included in a design then terms \mathbf{t}_1 and \mathbf{t}_n are also included, and vice versa. In the case of optional term \mathbf{t}_2 , if \mathbf{t}_2 is selected, \mathbf{Q} is also selected; however, the converse is not true.

Using these rules, production (1) would be translated to the following formula:

$$\mathbf{r} \Leftrightarrow \text{choose}_1(\mathbf{P1}, \mathbf{P2}) \wedge \mathbf{P1} \Leftrightarrow \mathbf{A} \wedge \mathbf{P1} \Leftrightarrow \mathbf{B} \wedge \mathbf{P2} \Leftrightarrow \mathbf{C} \wedge \mathbf{r1} \Rightarrow \mathbf{P2}$$

Mapping Grammars. The propositional formula of a grammar is the conjunction of: the formula for each production, the formula for each pattern, and the predicate $\mathbf{root} = \mathbf{true}$, where \mathbf{root} is the grammar's start production. The propositional formula for the grammar of Figure 2b is:

$$\mathbf{e} = \mathbf{true} \wedge \mathbf{e} \Leftrightarrow \mathbf{r} \wedge \mathbf{e} \Leftrightarrow \mathbf{s} \wedge \mathbf{r} \Leftrightarrow \text{choose}_1(\mathbf{G}, \mathbf{H}, \mathbf{I}) \wedge \mathbf{s} \Leftrightarrow \mathbf{A} \wedge \mathbf{s} \Leftrightarrow \mathbf{C} \wedge \mathbf{B} \Rightarrow \mathbf{s} \quad (4)$$

Contrary to current literature, feature models are generally *not* context free grammars. There are often additional constraints, here called *non-grammar constraints*, that gov-

ern the compatibility of features. Current tools often limit non-grammar constraints to simple exclusion (choosing feature \mathbf{I} automatically excludes a given feature list) and inclusion (choosing feature \mathbf{I} includes a given feature list). We think exclusion and inclusion constraints are too simplistic. In earlier work [5], we implemented feature models as attribute grammars enabling us to write constraints of the form:

\mathbf{F} implies \mathbf{A} or \mathbf{B} or \mathbf{C}

Meaning if feature \mathbf{F} is chosen, one of features \mathbf{A} , \mathbf{B} , or \mathbf{C} must be chosen, as \mathbf{A} , \mathbf{B} , or \mathbf{C} satisfied a required property of \mathbf{F} . More often, we found that preconditions for feature usage were based not on a single property but on *sets* of properties which could be satisfied by *combinations* of features, leading to complex predicates of the form:

\mathbf{F} implies $(\mathbf{A}$ and $\mathbf{X})$ or $(\mathbf{B}$ and $(\mathbf{Y}$ or $\mathbf{Z}))$ or \mathbf{C}

Exclusion constraints had a similar generality. For this reason, we concluded that non-grammar constraints should be *arbitrary* propositional formulas. By mapping a grammar to a propositional formula, we now can admit arbitrary propositional constraints by conjoining them onto the grammar's formula. In this way, a feature model (grammar + constraints) *is* a propositional formula.

An immediate application of these ideas may help resolve a pesky problem in that feature models do not have unique representations as feature diagrams. (That is, there are multiple ways of expressing the same constraints). It is a daunting task to know if two FDs are equivalent; how tools handle redundant representations is left to tool implementors [12]. It is possible to show that two FDs are equivalent if their propositional formulas are equivalent. See [17] for details.

4 Logic Truth Maintenance Systems

Feature models are the basis for declarative domain-specific languages for product specifications. As users select features for a particular application, we would like implications of these selections to be propagated, so that incorrect specifications can be avoided. A *Logic-Truth Maintenance Systems (LTMS)* can be used for this purpose.

A LTMS is a classic AI program that maintains the consequences of a propositional formula. An LTMS application is defined by:

- a set of boolean *variables*,
- a set of *propositional logic predicates* to constrain the values of these variables,¹
- *premises* (assignments to variables that hold universally),
- *assumptions* (assignments to variables that hold for the moment, but may be later retracted), and
- *inferences* (assignments to variables that follow from premises and assumptions).

The activities of an LTMS are to:

- compute inferences,

1. Equivalently, a single predicate can be used which is the conjunction of the input predicates.

- provide a rationale for variable assignments,
- detect and report contradictions,
- retract and/or make new assumptions, and
- maintain a database of inferences for efficient backtracking.

A *SAT* (*propositional satisfiability*) solver relies on an LTMS to help it search the combinatorial space for a set of variable assignments that satisfy all predicates. The efficiency of SAT solvers relies on a database of knowledge of previously computed inferences to avoid redundant or unnecessary searches [15].

What makes an LTMS complicated is (a) how it is to be used (e.g., a SAT solver requires considerable support) and (b) the number of rules and variables. If the number is large, then it is computationally infeasible to recompute inferences from scratch; retractions and new assumptions require incremental updates to existing assignments. This requires a non-trivial amount of bookkeeping by an LTMS.

Fortunately, a particularly simple LTMS suffices for our needs. First, the number of rules and variables that arise in feature models isn't large enough for performance to be an issue. (Inferences can be recomputed from scratch in a fraction of a second). Second, searching the space of possible variable assignments is performed *manually* by feature model users as they select and deselect features. Thus, an LTMS that supports only the first three activities previously listed is needed. Better still, standard algorithms for implementing LTMSs are well-documented in AI texts [15]. The challenge is to adapt these algorithms to our needs.

The mapping of LTMS inputs to feature models is straightforward. The variables are the tokens, production names, and pattern names of a grammar. The propositional formula is derived from the feature model (grammar + constraints). There is a single premise: **root=true**. Assumptions are features that are manually selected by users. Inferences are variable assignments that follow from the premise and assumptions.

In the next sections, we outline the LTMS algorithms that we have used in building our feature modeling tool **guids1**, whose use we illustrate in Section 5.

4.1 LTMS Algorithms

The *Boolean Constraint Propagation* (*BCP*) algorithm is the inference engine of an LTMS. Inputs to a BCP are a set of variables $\{v_1 \dots v_m\}$ and a set of arbitrary propositional predicates $\{p_1 \dots p_n\}$ whose conjunction $p_1 \wedge \dots \wedge p_n$ defines the *global constraint* (*GC*) on variable assignments (i.e., the formula of a feature model). BCP algorithms require the GC to be in *conjunctive normal form* (*CNF*) [15]. Simple and efficient algorithms convert arbitrary p_j to a conjunction of clauses, where a *clause* is a disjunction of one or more *terms*, a term being a variable or its negation [15].

BCP uses three-value logic (**true**, **false**, **unknown**) for variable assignments. Initially, BCP assigns **unknown** to all variables, except for premises which it assigns **true**. Given a set of variable assignments, each clause **c** of GC is either:

- *satisfied*: some term is **true**.

- *violated*: all terms are **false**.
- *unit-open*: one term is **unknown**, the rest are **false**.
- *non-unit open*: more than one term is **unknown** and the rest are **false**.

A unit-open term enables the BCP to change the **unknown** assignment to **true**. Thus, if clause C is $x \vee \neg y$ and x is **false** and y is **unknown**, BCP concludes y is **false**.

The BCP algorithm maintains a stack S of clauses to examine. Whenever it encounters a violated clause, it signals a contradiction (more on this later). Assume for now there are no contradictions. The BCP algorithm is simple: it marches through S finding unit-open clauses and setting their terms.

```

while (S is not empty) {
  c = S.pop();
  if (c.is_unit_open) {
    let t be term of c whose value is unknown;
    set(t);
  }
}

```

`set(t)` — setting a term — involves updating the term's variable's assignment (e.g., if t is $\neg y$ then y is assigned **false**), pushing unit-open terms onto S , and signalling contradictions:

```

set variable of t so that t is true;
for each clause C of GC containing  $\neg t$  {
  if (C.is_unit_open) S.push(C);
  else
  if (C.is_violated) signal_contradiction();
}

```

Invoking BCP on its initial assignment to variables propagates the consequences of the premises. For each subsequent assumption, the variable assignment is made and BCP is invoked. Let L be the sequence of assumptions (i.e., user-made variable assignments). The consequences that follow from L are computed by:

```

for each variable l in L {
  set(l);
  BCP();
}

```

If an assumption is retracted, it is simply removed from L .

A contradiction reveals an inconsistency in the feature model. When contradictions are encountered, they (and their details) must be reported to the feature model designers for model repairs.

Example. Suppose a feature model has the contradictory predicates $x \Rightarrow y$ and $y \Rightarrow \neg x$. If $x = \text{true}$ is a premise, BCP infers $y = \text{true}$ (from clause $x \Rightarrow y$), and discovers clause $(y \Rightarrow \neg x)$ to be violated, thus signalling a contradiction.

Explanations for why a variable has a value (or why a contradiction occurs) requires extra bookkeeping. Each time BCP encounters a unit-open clause, it keeps a record of

its conclusions by maintaining a 3-tuple of its actions $\langle \text{conclusion}, \text{reason}, \{\text{antecedents}\} \rangle$ where *conclusion* is a variable assignment, *reason* is the predicate (or clause) that lead to this inference, and *antecedents* are the 3-tuples of variables whose values were referenced. By traversing antecedents backwards, a justification for a conclusion can be presented in a human-understandable form.

Example. The prior example generates a pair of tuples: $\#1: \langle \mathbf{x}=\mathbf{true}, \text{premise}, \{\} \rangle$ and $\#2: \langle \mathbf{y}=\mathbf{true}, \mathbf{x} \Rightarrow \mathbf{y}, \{\#1\} \rangle$. The explanation for $\mathbf{y}=\mathbf{true}$ is: $\mathbf{x}=\mathbf{true}$ is a premise and $\mathbf{y}=\mathbf{true}$ follows from $\mathbf{x} \Rightarrow \mathbf{y}$.

4.2 A Complete Specification

A product specification (or equivalently, a variable assignment) is *complete* if the GC predicate is satisfied. What makes this problem interesting is how the GC predicate is checked. Assume that a user specifies a product by selecting features from a GUI or FD. When a feature is selected, the variable for that feature is set to **true**; a deselection sets it to **unknown**. (Inferencing can set a variable to **true** or **false**). Under normal use, users can only declare the features that they want, *not what they don't want*.

At the time that a specification is to be output, all variables whose values are **unknown** are assumed **false** (i.e., these features are not to be in the target product). The GC is then evaluated with this variable assignment in mind. If the GC predicate is satisfied, a valid configuration of the feature model has been specified. However, if a clause of GC fails, then either a complete sentence has not yet been specified or certain non-grammar constraints are unsatisfied. In either case, the predicate that triggered the failure is reported thus providing guidance to the user on how to complete the specification. This guidance is usually helpful.

5 An Example

We have built a tool, called **guids1**, that implements the ideas in the previous sections. **guids1** is part of the AHEAD Tool Suite [6] a set of tools for product-line development that support feature modularizations and their compositions. In the following section, we describe a classical product line and the **guids1** implementation of its feature model.

5.1 The Graph Product Line (GPL)

The *Graph Product-Line (GPL)* is a family of graph applications that was inspired by early work on modular software extensibility [27]. Each GPL application implements one or more graph algorithms. A **guids1** feature model for GPL (i.e., its grammar + constraints) is listed in Figure 5.

The semantics of the GPL domain are straightforward. A graph is either **Directed** or **Undirected**. Edges can be **Weighted** with non-negative numbers or **Unweighted**. A graph application requires at most one search algorithm: depth-first search (**DFS**) or breadth-first search (**BFS**), and one or more of the following algorithms:


```

// grammar

GPL : Driver Alg+ [Src] [Wgt] Gtp :: MainGpl ;
Gtp : Directed | Undirected ;
Wgt : Weighted | Unweighted ;
Src : BFS | DFS ;
Alg : Number | Connected | Transpose StronglyConnected :: StrongC
    | Cycle | MSTPrim | MSTKruskal | Shortest ;
Driver : Prog Benchmark :: DriverProg ;

%% // constraints

Number implies Src ;
Connected implies Undirected and Src ;
StrongC implies Directed and DFS ;
Cycle implies DFS ;
MSTKruskal or MSTPrim implies Undirected and Weighted ;
MSTKruskal or MSTPrim implies not (MSTKruskal and MSTPrim) ;
Shortest implies Directed and Weighted ;

```

Figure 5. GPL Model

- **Vertex Numbering (Number):** A unique number is assigned to each vertex.
- **Connected Components (Connected):** Computes the *connected components* of an undirected graph, which are equivalence classes under the reachable-from relation. For every pair of vertices \mathbf{x} and \mathbf{y} in a component, there is a path from \mathbf{x} to \mathbf{y} .
- **Strongly Connected Components (StrongC):** Computes the *strongly connected components* of a directed graph, which are equivalence classes under the reachable relation. Vertex \mathbf{y} is reachable from vertex \mathbf{x} if there is a path from \mathbf{x} to \mathbf{y} .
- **Cycle Checking (Cycle):** Determines if there are cycles in a graph. A cycle in directed graphs must have at least 2 edges, while in undirected graphs it must have at least 3 edges.
- **Minimum Spanning Tree (MSTPrim, MSTKruskal):** Computes a *Minimum Spanning Tree (MST)*, which contains all the vertices in the graph such that the sum of the weights of the edges in the tree is minimal.
- **Single-Source Shortest Path (Shortest):** Computes the shortest path from a source vertex to all other vertices.

The grammar that defines the order in which GPL features are composed is shown in Figure 5. Not all combinations of features are possible. The rules that govern compatibilities are taken directly from algorithm texts [9] and are listed in Figure 6. These constraints are listed as additional propositional formulas (below the %% in Figure 5). When combined with the GPL grammar, a feature model for GPL is defined.

Algorithm	Required Graph Type	Required Weight	Required Search
Vertex Numbering	Any	Any	BFS, DFS
Connected Components	Undirected	Any	BFS, DFS
Strongly Connected Components	Directed	Any	DFS
Cycle Checking	Any	Any	DFS
Minimum Spanning Tree	Undirected	Weighted	None
Shortest Path	Directed	Weighted	None

Figure 6. Feature Constraints in GPL

The GUI that is generated from Figure 5 is shown in Figure 7. The state that is shown results from the selection of **MSTKruskal** — the **Weighted** and **Undirected** features are automatically selected as a consequence of constraint propagation. Further,

Shortest, **MSTPrim**, **StrongC**, **Unweighted**, and **Directed** are greyed out, meaning that they are no longer selectable as doing so would create an inconsistent specification. Using an LTMS to propagate constraints, users can only create correct specifications. In effect, the generated GUI is a declarative domain-specific language that acts as a syntax-directed editor which prevents users from making certain errors.



Figure 7 Generated GUI for the GPL Model

Although not illustrated, `guids1` allows additional variables to be declared in the constraint section to define properties. Feature constraints can then be expressed in terms of properties, like that in [5], to support our observations in Section 3.

Another useful capability of LTMSs is to provide a justification for automatically selected/deselected features. We have incorporated this into `guids1`: placing the mouse over a selected feature, a justification (in the form of a proof) is displayed. In the example of Figure 7, the justification for **Undirected** being selected is:

```
MSTKruskal because set by user
Undirected because ((MSTKruskal or MSTPrim)) implies
((Undirected and Weighted))
```

Meaning that **MSTKruskal** was set by the user, and **Undirected** is set because the selection of **MSTKruskal** implies **Undirected** and **Weighted**. More complex explanations are generated as additional selections are made.

5.2 Debugging Feature Models

Debugging a feature model without tool support is notoriously difficult. When we debugged feature models prior to this work, it was a laborious, painstaking, and error-prone effort to enumerate feature combinations. By equating feature models with propositional formulas, the task of debugging is substantially simplified.

Given a propositional formula and a set of variable assignments, a SAT solver can determine whether there is a value assignment to the remaining variables that will satisfy the predicate. Thus, debugging scripts in `guids1` are simply statements of the form `<S,L>` where **L** is a list of variable assignments and **S** is **true** or **false**. If **S** is **true**, then the SAT solver is expected to confirm that **L** is a compatible set of variable assignments; if **S** is **false**, the solver is expected to confirm that **L** is an incompatible set of assignments. Additional simple automatic tests, not requiring a SAT solver, is to

verify that a given combination of features defines a product (i.e., a legal and complete program specification). Both the SAT solver and complete-specification-tests were instrumental in helping us debug the GPL feature model.

It is straightforward to list a large number of tests to validate a model; test suites can be run quickly. (SAT solvers have become very efficient, finding variable assignments for tens of thousands of variables in minutes). Although we cannot prove a model is correct, we are comforted by the fact that we can now run a *much* more thorough set of tests on our models automatically than we could have performed previously.

6 Staged Configuration Models

Staged configuration of has recently been proposed as an incremental way to progressively specialize feature models [11][12]. At each stage, different groups or developers make product configuration choices, rather than a configuration being specified by one person at one time. Specializations involve the selection or deselection of features and adding more constraints (e.g., converting a one-or-more selection to single selection).

Staged configuration is accomplished by (1) simplifying the grammar by eliminating choices or making optional choices mandatory, and (2) simplifying the non-grammar constraints. Both are required ([12] addresses grammar simplification). By limiting changes only to grammars, it is possible to preselect **MSTKruskal** and deselect **Unweighted** in a staged configuration and adjust the GPL grammar (making **MSTKruskal** mandatory and removing **Unweighted**). But the resulting model is unsatisfiable, as **MSTKruskal** requires **Unweighted**.

A generalization of the GUI presented earlier could be used to accomplish staged specifications. Each selectable feature will require a toggle that allows a feature to be selected (**true**), deselected (**false**), or to postpone its choice to a later stage (**unknown**). In this way, designers can distinguish features that are preselected from those that are permanently removed. The LTMS algorithm remains unchanged; constraints are propagated as before guaranteeing that the resulting model is consistent. Inferred feature selections and deselections can be used to further simplify the grammar and its non-grammar constraints.

More generally, where constraints on non-boolean variables (e.g. performance constraints) are part of a feature model, a more general logic, constraint propagation algorithms, and predicate simplification algorithms will be needed [13]. However, our work applies to many existing feature models, and we believe that current results on staged configuration can be improved for these cases with our suggestions.

7 Related Work

There is a great deal of prior work on feature modeling. For brevity, we focus on the key papers that are relevant. We are aware that technologies similar to that presented in this paper may be known to pockets of researchers in industry (e.g., [1][2][7][16]), but the basic relationship of feature models, attribute grammars, and propositional formulas does not seem to be widely appreciated or understood.

As mentioned earlier, the connection of feature models to grammars is not new. In 1992, Batory and O'Malley used grammars to specify feature models in GenVoca [4], and in 1997 showed in [5] how attribute grammars could express non-grammar constraints. In 2002, de Jonge and Visser recognized that feature diagrams were context free grammars. Czarnecki, Eisenecker, et al. have since used grammars to simplify feature models during staged configuration [10].

The connection of product-line configurations with propositional formulas is due to Mannion [24]. The next closest work to that presented in this paper is that of Beuche [7] and Pure::Variants [25], where feature models are translated into Prolog. Prolog is used as a constraint inference engine to accomplish the role of an LTMS. Non-grammar constraints are expressed by inclusion and exclusion predicates; while user-defined constraints (i.e., Prolog programs) could be arbitrary. User-defined constraints were not recommended as they could be arbitrarily complex or undecidable, thus causing problems in automatic problem solving [7]. While true for prolog programs, propositional logic is decidable, and arbitrarily complex propositional logic formulas (in theory) should be no problem. We are unaware of tools that follow from [7] to debug feature models.

8 Conclusions

In this paper, we integrated existing results to clearly expose a fundamental connection between FDs, grammars, and propositional formulas. This connection has enabled us to leverage light-weight, efficient, and easy-to-build LTMSs and off-the-shelf SAT solvers to bring useful new capabilities to feature modeling tools. LTMSs provide a simple way to propagate constraints as users select features in product specifications. SAT solvers provide automated support to help debug feature models. We believe that the use of LTMSs and SAT solvers in feature model tools is novel. Further, we explained how work on staged configuration models could be improved by integrating non-grammar constraints into a staging process.

We believe that the foundations presented in this paper will be useful in future tools for product-line development.

Acknowledgements. I thank Ben Kuipers and Ray Mooney for directing me to the literature on LTMSs. I gratefully acknowledge the helpful comments from Mark Grechanik and Jack Sarvela, and Daniel Le Berre for his help with the SAT4J solver. I also thank Jun Yuan, Karim Jamal, and Maria Zolotova for their help in building `guids1`.

9 References

- [1] American Standard, <http://www.americanstandard-us.com/planDesign/>
- [2] Annenberg/CPB, <http://www.learner.org/exhibits/parkphysics/coaster/>
- [3] M. Antkiewicz and K. Czarnecki, "FeaturePlugIn: Feature Modeling Plug-In for Eclipse", *OOPSLA'04 Eclipse Technology eXchange (ETX) Workshop*, 2004.
- [4] D. Batory and S. O'Malley, "The Design and Implementation of Hierarchical Software Systems with Reusable Components", *ACM TOSEM*, October 1992.

- [5] D. Batory and B.J. Geraci, "Composition Validation and Subjectivity in GenVoca Generators", *IEEE TSE*, February 1997, 67-82.
- [6] D. Batory, *AHEAD Tool Suite*, www.cs.utexas.edu/users/schwartz/ATS.html
- [7] D. Beuche, "Composition and Construction of Embedded Software Families", Ph.D. thesis, Otto-von-Guericke-Universitaet, Magdeburg, Germany, 2003.
- [8] BMW, <http://www.bmwusa.com/>
- [9] T.H. Cormen, C.E. Leiserson, and R.L.Rivest. *Introduction to Algorithms*, MIT Press, 1990.
- [10] K. Czarnecki and U. Eisenecker. *Generative Programming Methods, Tools, and Applications*. Addison-Wesley, Boston, MA, 2000.
- [11] K. Czarnecki, S. Helsen, and U. Eisenecker, "Formalizing Cardinality-based Feature Models and their Specialization", *Software Process Improvement and Practice*, 2005 10(1).
- [12] 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), 2005.
- [13] K. Czarnecki, private correspondence, 2005.
- [14] N. Eén and N. Sörensson, "An extensible SAT solver". *6th International Conference on Theory and Applications of Satisfiability Testing*, LNCS 2919, pp 502-518, 2003.
- [15] K.D. Forbus and J. de Kleer, *Building Problem Solvers*, MIT Press 1993.
- [16] Gateway Computers. <http://www.gateway.com/index.shtml>
- [17] M. Grechanik and D. Batory, "Verification of Dynamically Reconfigurable Applications", in preparation 2005.
- [18] J. Greenfield, K. Short, S. Cook, S. Kent, and J. Crupi, *Software Factories: Assembling Applications with Patterns, models, Frameworks and Tools*, Wiley, 2004.
- [19] M. de Jong and J. Visser, "Grammars as Feature Diagrams".
- [20] D. Streitferdt, M. Riebisch, I. Philippow, "Details of Formalized Relations in Feature Models Using OCL". *10th IEEE Symposium and Workshops on Engineering of Computer-Based Systems (ECBS'03)*, Huntsville Alabama, USA, April 7-11, 2003. IEEE Computer Society, 2003, pp. 297-304.
- [21] A. Hein, M. Schlick, R. Vinga-Martins. Applying Feature Models in Industrial Settings. In *Software Product Lines: Experience and Research Directions*. (SPLC1). Kluwer Academic Publishers, 2000.
- [22] K. Kang, S. Cohen, J. Hess, W. Nowak, and S. Peterson. "Feature-Oriented Domain Analysis (FODA) Feasibility Study". Technical Report, CMU/SEI-90TR-21, November 1990.
- [23] R.E. Lopez-Herrejon and D. Batory, "A Standard Problem for Evaluating Product-Line Methodologies", *GCSE 2001*, September 9-13, 2001 Messe Erfurt, Erfurt, Germany.
- [24] M. Mannion, "Using first-order logic for product line model validation". *2nd Software Product Line Conf. (SPLC2)*, #2379 in LNCS, 176-187, 2002.
- [25] Pure-Systems, "Technical White Paper: Variant Management with pure::variants", www.pure-systems.com, 2003.
- [26] T. Teitelbaum and T. Reps, "The Cornell Program Synthesizer: a Syntax-Directed Programming Environment", *CACM*, v.24 n.9, p.563-573, Sept. 1981.
- [27] P. Zave, "FAQ Sheet on Feature Interactions", www.research.att.com/~pamela/faq.html