

2008

Generating Variation-point Obligations for Compositional Model Checking of Software Product Lines

Jing (Janet) Liu
Iowa State University

Samik Basu
Iowa State University, sbasu@iastate.edu

Robyn R. Lutz
Iowa State University, rlutz@iastate.edu

Follow this and additional works at: http://lib.dr.iastate.edu/cs_techreports

 Part of the [Software Engineering Commons](#)

Recommended Citation

Liu, Jing (Janet); Basu, Samik; and Lutz, Robyn R., "Generating Variation-point Obligations for Compositional Model Checking of Software Product Lines" (2008). *Computer Science Technical Reports*. 334.
http://lib.dr.iastate.edu/cs_techreports/334

This Article is brought to you for free and open access by the Computer Science at Iowa State University Digital Repository. It has been accepted for inclusion in Computer Science Technical Reports by an authorized administrator of Iowa State University Digital Repository. For more information, please contact digirep@iastate.edu.

Generating Variation-point Obligations for Compositional Model Checking of Software Product Lines

Abstract

Software product lines are widely used due to their advantageous reuse of shared features while still allowing optional and alternative features in the individual products. Especially for high-integrity product lines, we would like to use model checking to verify that key properties hold as each new product is built. However, this goal is currently hampered by the complexity of composing model-checking results for the features in a way that allows reuse for subsequent products. This paper presents an incremental and compositional model-checking technique that allows efficient reuse of model checking results associated with the features in a product line. It goes beyond related work in that it removes restrictions on how the features can be sequentially composed. This flexibility is important because it means that many more real-world systems can be model-checked. We have implemented the technique, and demonstrate and evaluate it on a medical device product line.

Keywords

Software product lines, compositional model checking, variation point, feature

Disciplines

Software Engineering

Generating Variation-point Obligations for Compositional Model Checking of Software Product Lines

Jing (Janet) Liu Samik Basu
Department of Computer Science
Iowa State University
Ames, IA, 50010
{janetlj, sbasu}@cs.iastate.edu

Robyn Lutz
Department of Computer Science
Iowa State University & JPL/Caltech
Ames, IA, 50010
rlutz@cs.iastate.edu

Abstract

Software product lines are widely used due to their advantageous reuse of shared features while still allowing optional and alternative features in the individual products. Especially for high-integrity product lines, we would like to use model checking to verify that key properties hold as each new product is built. However, this goal is currently hampered by the complexity of composing model-checking results for the features in a way that allows reuse for subsequent products. This paper presents an incremental and compositional model-checking technique that allows efficient reuse of model checking results associated with the features in a product line. It goes beyond related work in that it removes restrictions on how the features can be sequentially composed. This flexibility is important because it means that many more real-world systems can be model-checked. We have implemented the technique, and demonstrate and evaluate it on a medical device product line.

1 Introduction

Software product lines are widely used due to their advantageous reuse of shared elements, but this reuse across different products poses challenges for model checking of product lines. Especially for high-integrity product lines, we would like to use model checking to verify that key properties hold in each new product. However, model-based verification of software product lines is currently hampered by the complexity of composing model-checking results of the various features in a way that allows reuse when model-checking new products.

In a software product line, the products all share a common set of mandatory features but are differentiated one from the other by their variable (optional and alternative) features [20]. Each feature carries an increment of functionality for the system [2]. Typically, the set of variations

are selected and composed on top of the common base features to create each distinct, new product. The locations in the features (usually the common features) [11, 19] where other features can be added (usually the variable features) to construct the various products are called *variation points*.

Model checking [4, 10] takes a model of a given system’s design, and checks if it satisfies certain properties of the system, interpreted in terms of logic formulas. It is a powerful technique for enhancing the quality of software systems, e.g., by identifying flaws that would not have been caught otherwise [9, 12]. As such, model checking can play a vital role in verifying key properties of products in high-integrity product lines such as pacemakers, medical imaging systems, and avionics control systems.

However, formal reasoning about each product in isolation fails to exploit the fact that all the products in a product line share common features. Similarly, many products in a product line will share some of the variable features. Repeated verification of the same sets of features wastes resources and discourages industrial adoption of model-checking for product lines.

This paper presents an incremental and compositional model-checking technique that allows effective reuse of model checking results associated with the features in a product line. The contribution of the paper is that, in contrast with existing work on compositional model checking of features [3, 17, 18], we impose no restrictions (e.g., regarding the sequence, type of connection points, or number of connections) on how the features can be composed. Any type of sequential composition of features, not just pipelined composition, can be verified. Similarly, behavior of a feature that depends on another feature’s behavior (which may, in turn, depend on the first feature) can be verified.

To achieve these extensions, our technique generates obligations at the variation points such that the feature composition satisfies the desired property if and only if the

features added at variation points satisfy the corresponding obligations. These *variation-point obligations* guide the verification of features subsequently composed at the variation points as new products in the product line are built.

By allowing more kinds of interactions between features, our approach provides three important advantages for model-checking product lines. First, such flexibility means that many more real-world systems can be model-checked. This moves model checking closer to product-line development practice. Second, the implementation stores the variation-point obligations obtained for each feature during earlier model-checking runs, thus enabling reuse of previous model-checking results when a new product is composed. Re-verification is only performed when needed, providing savings in space and time over non-compositional model checking. Third, as a product line evolves, new variation points are typically introduced. The technique described in this paper accommodates such changes by identifying obligations at these new variation points from previous obligation computations done at those points of change. We have implemented the technique, and demonstrate and evaluate it on a medical device product line.

The rest of the paper is organized as follows. Section 2 provides a motivating example. Section 3 presents the preliminary information of this work. Section 4 gives an overview of this approach, and Section 5 illustrates each step in more detail. Section 6 demonstrates our technique on a simplified pacemaker product line and discusses test results. Finally, Section 7 describes related work, and Section 8 offers concluding remarks.

2 Illustrative Example

The work reported here was motivated by the difficulty of reusing model-checking results during the development and evolution of safety-critical product lines. Our effort is directed at enabling reuse of previous model-checking results so that system properties can be efficiently verified when a new product is built in the product line. The paper uses an example of a simplified pacemaker product line to evaluate performance of our approach (Sect. 6). A pacemaker [5] is an embedded medical device designed to monitor and regulate the beating of the heart when it is not beating at a normal rate. It is safety-critical because some failures can damage the patient’s health or even lead to loss of life [5, 14]. Figure 1 shows four products in the pacemaker product line [15, 16]:

BasePacemaker has the basic functionality shared by all pacemakers: generating a pulse if no heart beat is detected during the sensing interval. This mode of execution is called Inhibited Mode.

ModeTransitivePacemaker has an additional feature called ModeTransition Extension that enables it to switch between Inhibited Mode and TriggeredMode during execu-

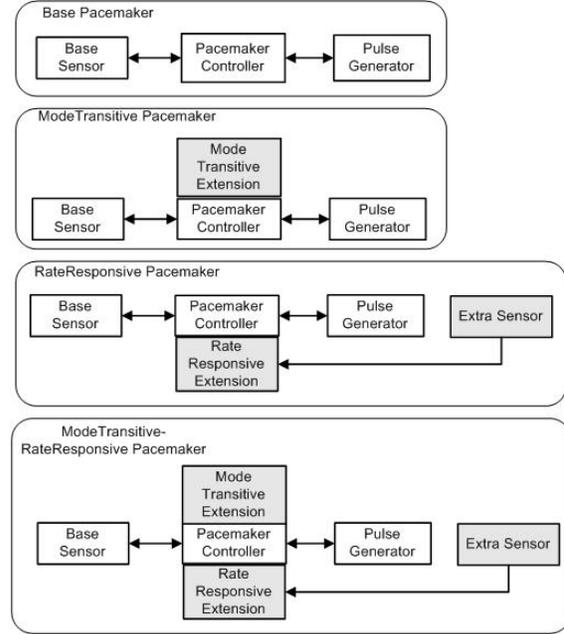


Figure 1. Pacemaker product line overview.

tion. In the TriggeredMode, a pulse follows every heartbeat to regulate the heartbeat.

RateResponsivePacemaker adds an Extra Sensor that can detect a patient’s activity level (i.e., respiration rate while resting vs. while exercising). This product has an additional feature, the RateResponsive Extension, that adjusts the sensing interval (to normal or upperRateLimit) according to the patient’s current activity level.

ModeTransitive-RateResponsivePacemaker combines the features of the ModeTransitivePacemaker and the RateResponsivePacemaker to provide both inhibited and triggered heartbeat regulation and adaptation to patient’s activity level.

Certain properties must be shown to be true for every product in the product line in order to assure patient safety. An example of such a property is: *In the InhibitedMode, the pacemaker shall always generate a pulse when no heartbeat is detected during the normal sensing interval.*

Since verification of this property involves BasePacemaker functionality common to all the products, we would like to avoid unnecessary, repeated checking of the same feature as each new product is built. We next describe our approach to achieving this through appropriate reuse of the results from previous model-checking runs. By generating and managing obligations at the variation points, the model-checking effort is aligned with the inherent variation points that a product-line development approach provides.

3 Preliminaries

We represent the functional behavior of features in a product line using finite state machines where states represent the configurations of the functional behavior and transitions from one state to another represent the evolution of the behavior between configurations. Formally, the model is defined as follows:

Definition 1 (Feature Behavioral Model). *Feature behavioral model* $FM = (S, S_0, V, T, L)$ where S is the set of states, $S_0 \subseteq S$ is the set of initial states, $V \subseteq S$ is the set of variation points, $T \subseteq S \times S$ is the transition relation, and $L : S \rightarrow 2^P$ is the labeling function which associates each state $s \in S$ with the set of propositions in P that are true in that state. We will denote $(s, s') \in T$ by $s \rightarrow s'$. \square

In the above, $s \in V$ acts as the variation point where one FM can be plugged into another, i.e., when two FMs are sequentially composed, new transitions are added from some variation points of one to states in the other. For each composition between FM_1 and FM_2 , we use $T_c^{FM_1, FM_2} \subseteq V^1 \times S^2$ where V^1 is the set of variation point of FM_1 and S^2 is the set of states in FM_2 . The relation $T_c^{FM_1, FM_2}$ denotes how the states in FM_2 are connected to the variation points of FM_1 . We define sequential composition as follows:

Definition 2 (Sequential Composition). *Given* $FM_1 = (S^1, S_0^1, V^1, T^1, L^1)$, $FM_2 = (S^2, S_0^2, V^2, T^2, L^2)$, $T_c^{FM_1, FM_2}$, and $T_c^{FM_2, FM_1}$, *the sequential composition* $Comp_{seq}(FM_1, FM_2) = (S^1 \cup S^2, S_0^1, V^{12}, T^{12}, L^{12})$,

1. $V^{12} = \{s \mid s \in V^1 \cup V^2\}$,
2. $T^{12} = T^1 \cup T^2 \cup T_c^{FM_1, FM_2} \cup T_c^{FM_2, FM_1}$, and
3. $L^{12}(s) = \begin{cases} L^1(s) & \text{if } s \in S^1 \\ L^2(s) & \text{otherwise} \end{cases}$ \square

Observe that the above definition allows FM_1 to be connected to FM_2 and vice versa, resulting in possible loops between the behaviors of the two features. FM_1 is the start feature at whose start states (S_0^1) we want a given property to be satisfied. The set of variation points V^{12} of $Comp_{seq}(FM_1, FM_2)$ includes the states in V^i ($i \in \{1, 2\}$) that may have been used in the composition. They are also the states that can be used as variation points for future additions of other features.

A *closed* FM is one which does not have any variation points ($V = \emptyset$). In other words, a closed FM cannot be augmented with new features. An open FM is one whose set of variation points is non-empty.

Temporal Logic CTL. Properties, in our setting, are described using CTL temporal logic [10]. We present a brief overview of the syntax and semantics of CTL formulas. The syntax of CTL can be defined as follows:

$$\phi \rightarrow \text{tt} \mid \text{ff} \mid P \mid \neg\phi \mid \phi \vee \phi \mid \text{EX}(\phi) \mid \text{E}(\phi \cup \phi) \mid \text{EG}(\phi)$$

The semantics of the CTL formulas are given in terms of the states of finite state systems (FM) that satisfy the formulas. The propositional constant tt is satisfied in all states while ff is not satisfied by any state. The proposition p ($\neg p$) is satisfied by state s such that $p \in L(s)$ ($p \notin L(s)$). $\neg\phi$ is satisfied by states where ϕ is not satisfied. $\phi_1 \vee \phi_2$ is satisfied by states that satisfy ϕ_1 or ϕ_2 . $\text{EX}(\phi)$ is satisfied by a state which has at least one transition to a state that satisfies ϕ . $\text{E}(\phi_1 \cup \phi_2)$ is satisfied by a state which has a path where ϕ_1 holds in every state in that path until a state satisfying ϕ_2 is reached. $\text{EG}(\phi)$ is satisfied by a state which has a path where every state in the path satisfies ϕ .

The above syntax forms the adequate set of CTL formula syntax. Some other widely used syntactic constructs such as $\text{EF}(\phi)$, $\text{AX}(\phi)$, $\text{AF}(\phi)$, $\text{A}(\phi_1 \cup \phi_2)$, $\text{AG}(\phi)$ can be obtained from the adequate set; for example: $\text{AX}(\phi) \equiv \neg\text{EX}(\neg\phi)$ and $\text{EF}(\phi) \equiv \text{E}(\text{tt} \cup \phi)$.

A state belonging to the semantics of ϕ implies that the state satisfies ϕ , denoted by $s \models \phi$. We say that a closed FM $FM = (S, S_0, \emptyset, T, L)$ satisfies a CTL formula ϕ , denoted by $FM \models \phi$, if and only if $\forall s \in S_0 : s \models \phi$. For a detailed discussion of model checking closed systems see [10].

4 Method Overview

As noted in Section 3, a closed FM can be verified to check whether or not it satisfies a desired CTL property. However, for an open FM such as ours, satisfiability of CTL properties may depend on the behavior of the features being connected to those variation points.

Given an FM and a desired property for its possible compositions with other FMs, our solution relies on generating a set of CTL formulas as obligations for each of its variation points. A composition satisfies the desired property if and only if the added features at each variation point satisfy the corresponding obligations. We refer to these obligations as *variation-point obligations*. As our definition of the feature composition allows loops between the features, such circular dependency is handled by recording in a global database, *answer set* (denoted by aSet), whether or not variation-point obligations are satisfied by a composition.

Thus, checking whether a sequential composition $Comp_{seq}(FM_1, FM_2, \dots, FM_m)$ satisfies a CTL formula ϕ amounts to checking whether all the start states of FM_1 satisfies ϕ and can be compositionally resolved as follows:

- Step 1 Generate the variation-point obligations for satisfying ϕ in all the variation points of FM_i (initially $i = 1$). Record the variation-point obligations in aSet .
- Step 2 Use $T_c^{FM_i, FM_k}$ to identify all the features FM_k s connected to the variation points of FM_i : if state t_k in FM_k is connected to variation point s_i of FM_i , where the variation-point obligation is ϕ_i , iterate from Step 1 (with $i = k$) to compute the variation-point obligations

for each FM_k , such that t_k satisfies φ_i . If t_k satisfies its obligation φ_i , then update the `aSet` entry for s_i in FM_i .

Step 3 If the `aSet` cannot be further updated from computing variation-point obligations, break from the iteration. Analyze `aSet` to identify loops between features and update `aSet` accordingly.

If the final `aSet` records that the start states of FM_1 satisfy φ , then the composition $Comp_{seq}(FM_1, FM_2, \dots, FM_m)$ satisfies φ .

5 Detailed Approach

In this section, we describe the generation of variation-point obligations, how to update the answer set to identify inter-feature loops, and our algorithm for compositional model checking.

Variation-point Obligations. Variation-point obligations are sets of formulas associated with the variation points that must be satisfied by the features connected to them. The obligations are also annotated with the boolean operators \wedge and \vee to handle cases where multiple features are connected to the same variation point. They are formally defined as follows:

Definition 3 (Variation-point Obligation). *Given an FM = (S, S_0, V, T, L) , an obligation at a variation point is a formula of the form: $\Psi \rightarrow (\phi, s, op) \mid \neg\Psi \mid \Psi \vee \Psi \mid \Psi \wedge \Psi$ where ϕ is a CTL formula, $s \in V$, $op \in \{\vee, \wedge, \perp\}$. \square*

The variation-point obligation (φ, s, \vee) states that one of the features added at the variation point (state s) must satisfy φ ; (φ, s, \wedge) means that any new feature added at state s must satisfy φ . A variation-point obligation of the form $(\varphi_1, s_1, \vee) \vee (\varphi_2, s_2, \vee)$ (resp. $(\varphi_1, s_1, \vee) \wedge (\varphi_2, s_2, \vee)$) states that (φ_1, s_1, \vee) or (resp. and) (φ_2, s_2, \vee) must be satisfied. Finally, $\neg(\varphi, s, \vee) \equiv (\neg\varphi, s, \wedge)$ is satisfied at the variation point if φ is not satisfied in any of the new features added at s .

We use $(\varphi, \epsilon, \perp)$ to indicate that φ is not an obligation at any variation point. We also use the following simplification rules: $(\text{tt}, \epsilon, \perp) \vee \psi \equiv (\text{tt}, \epsilon, \perp)$; $(\text{tt}, \epsilon, \perp) \wedge \psi \equiv \psi$; $\neg(\text{tt}, \epsilon, \perp) \equiv (\text{ff}, \epsilon, \perp)$; $(\text{ff}, \epsilon, \perp) \wedge \psi \equiv (\text{ff}, \epsilon, \perp)$; $(\text{ff}, \epsilon, \perp) \vee \psi \equiv \psi$; $\neg(\text{ff}, \epsilon, \perp) \equiv (\text{tt}, \epsilon, \perp)$; $\neg(\psi_1 \vee \psi_2) \equiv \neg\psi_1 \wedge \neg\psi_2$; $\neg(\psi_1 \wedge \psi_2) \equiv \neg\psi_1 \vee \neg\psi_2$. We use $\psi, \psi_1, \psi_2, \dots$ to denote variation-point obligation formulas while using $\varphi, \varphi_1, \varphi_2, \dots$ to represent CTL formulas.

Generation of variation-point obligation follows in similar fashion to local and on-the-fly CTL model checking [10] where, given a state and a CTL formula to be satisfied at that state, the algorithm proceeds by recursively exploring the reachable state space and by unfolding the CTL formula. We use the following equivalences of CTL formulas for formula unfolding: $E(\varphi_1 \cup \varphi_2) \equiv \varphi_2 \vee (\varphi_1 \wedge EX(E(\varphi_1 \cup \varphi_2)))$; $EG(\varphi) \equiv \varphi \wedge EX(EG(\varphi))$.

Step 1: Computing variation-point obligations. Given an FM and a CTL formula φ , we define for every state s in FM the functions `t_Obl` and `Obl`, which generate the obligations at the variation points of FM required for s to satisfy φ . The functions take five parameters: φ , the CTL formula that is required to be satisfied at s ; s , the current state of the FM; H , the history set recording the state-formula pairs that have been visited in the recursive definition of the functions (to handle loops in the FM); `aSetin` and `aSetout` (the answer sets before and after the invocation of the function).

The answer set contains elements of the form $(\varphi, s) \mapsto \psi$ where φ is a CTL formula and ψ is a variation-point obligation. We say that satisfiability of φ at s depends on the satisfiability of ψ . Specifically,

1. $(\varphi, s) \mapsto (\varphi', s', op')$ denotes that for s to satisfy φ , all (at least one of the) features connected via the variation point s' must satisfy φ' when op' is equal to \wedge (resp. \vee).
2. $(\varphi, s) \mapsto (pc, \epsilon, \perp)$ denotes that s satisfies (does not satisfy) φ when pc is a propositional constant equal to `tt` (resp. `ff`).
3. $(\varphi, s) \mapsto \psi_1 \wedge \psi_2$ denotes that s satisfies φ if both ψ_1 and ψ_2 are satisfied. Similarly, $(\varphi, s) \mapsto \psi_1 \vee \psi_2$ denotes that s satisfies φ if one of ψ_1 and ψ_2 is satisfied.

The `aSet` is necessary to handle loops across multiple features (see Step 3 below). It also allows us to reuse the previous results to remove redundant, repeated computations. The recursive definition of the functions `t_Obl` and `Obl` are presented in Figure 2.

Rule A corresponds to `t_Obl` (top-level call) which states that a variation-point obligation corresponding to state s and formula φ is equal to the result present in `aSetin` if `t_Obl` has been invoked on the same state-formula pair before. If the current invocation of `t_Obl` is the first-time call with the corresponding state-formula pair, then `Obl` is invoked, and its result ψ is used to update the answer set. Note that the call to `Obl` may update the `aSetin` to `aSettemp`. If the latter already contains an entry of the form $(\varphi, s) \mapsto \psi'$, then the mapping for (φ, s) is updated to $\psi \ op \ \psi'$ where op is decided on the basis of the formula being universal (e.g. AG, AU) or existential (e.g. EG, EU).

The choice of op can be explained as follows. ψ and ψ' are the variation-point obligations that need to be satisfied for φ to hold at s . If φ is an universal (resp. existential) formula, the obligation at the variation point will also require all (resp. at least one) features connected to that variation point to satisfy that obligation. Accordingly, the result is obtained via conjunction or disjunction operation(s).

Rules 1–8 correspond to the `Obl` function. Observe that `Obl` invokes `t_Obl` to appropriately use the `aSet`. The first three rules in Figure 2 state that for propositional constants and propositions, there is no obligation at the variation points; satisfiability of these types of CTL formulas

$$\begin{aligned}
\text{A. } \quad t_Obl(\varphi, s, H, aSet_{in}, aSet_{out}) &:= \begin{cases} \psi & \text{if } (\varphi, s) \mapsto \psi \in aSet_{in}; \text{ where } aSet_{out} := aSet_{in} \\ Obl(\varphi, s, H, aSet_{in}, aSet_{temp}) & \text{otherwise} \\ \text{where } \psi := Obl(\varphi, s, H, aSet_{in}, aSet_{temp}) \\ aSet_{out} := \begin{cases} aSet_{temp}[(\varphi, s) \mapsto \psi' / (\varphi, s) \mapsto \psi \text{ op } \psi'] \\ \text{if } (\varphi, s) \mapsto \psi' \in aSet_{temp} \\ \text{where } op := \begin{cases} \wedge & \text{if } \varphi \text{ is a universal} \\ \vee & \text{otherwise} \end{cases} \end{cases} \\ aSet_{temp} \cup \{(\varphi, s) \mapsto \psi\} & \text{otherwise} \end{cases} \\
1. \quad Obl(tt, s, H, aSet, aSet) &:= (tt, \epsilon, \perp) \\
2. \quad Obl(ff, s, H, aSet, aSet) &:= (ff, \epsilon, \perp) \\
3. \quad Obl(p, s, H, aSet, aSet) &:= \begin{cases} (tt, \epsilon, \perp) & \text{if } p \in L(s) \\ (ff, \epsilon, \perp) & \text{otherwise} \end{cases} \\
4. \quad Obl(\neg\varphi, s, H, aSet_{in}, aSet_{out}) &:= \neg t_Obl(\varphi, s, H, aSet_{in}, aSet_{out}) \\
5. \quad Obl(\varphi_1 \vee \varphi_2, s, H, aSet_{in}, aSet_{out}) &:= t_Obl(\varphi_1, s, H, aSet_{in}, aSet_{temp}) \vee \\ & \quad t_Obl(\varphi_2, s, H, aSet_{temp}, aSet_{out}) \\
6. \quad Obl(E(\varphi_1 \cup \varphi_2), s, H, aSet_{in}, aSet_{out}) &:= \begin{cases} (ff, \epsilon, \perp) & \text{if } (E(\varphi_1 \cup \varphi_2), s) \in H; \text{ where } aSet_{out} := aSet_{in} \\ t_Obl(\varphi_2 \vee (\varphi_1 \wedge EX(E(\varphi_1 \cup \varphi_2))), s, \\ H \cup \{(E(\varphi_1 \cup \varphi_2), s)\}, aSet_{in}, aSet_{out}) & \text{otherwise} \end{cases} \\
7. \quad Obl(EG(\varphi), s, H, aSet_{in}, aSet_{out}) &:= \begin{cases} (tt, \epsilon, \perp) & \text{if } (EG(\varphi), s) \in H; \text{ where } aSet_{out} := aSet_{in} \\ t_Obl(\varphi \wedge EX(EG(\varphi)), s, H \cup \{(EG(\varphi), s)\}, aSet_{in}, aSet_{out}) & \text{otherwise} \end{cases} \\
8. \quad Obl(EX(\varphi), s, H, aSet_{in}, aSet_{out}) &:= \bigvee_{s \rightarrow s'} t_Obl(\varphi, s', H, aSet_{in}, aSet_{out}) \vee \begin{cases} (\varphi, s, \vee) & \text{if } s \in V \\ (ff, \epsilon, \perp) & \text{otherwise} \end{cases}
\end{aligned}$$

Figure 2. Variation-point obligations

can be decided at the current state s . As these function rules do not update the answer set, $aSet_{in}$ and $aSet_{out}$ remain unchanged. In Rule 5, the answer set updates are chained from one t_Obl call to the other.

Rules 6 and 7 use H to decide the satisfiability of EU and EG properties in the presence of a loop (in the same feature model). If the state-formula pair is present in the history set, this shows circular dependency in a path. Thus, for the least fixed point formula EU, the result is (ff, ϵ, \perp) . For the greatest fixed point formula EG, the result is (tt, ϵ, \perp) . On the other hand, if the state-formula pair is not present in the history set, the formula is expanded to its equivalent form with t_Obl being invoked, and the history set is updated.

Finally, Rule 8 deals with $EX(\varphi)$ formulas. The obligation is computed by expanding or moving to all possible next states of the current state s . There are two disjuncts in the result. The first disjunct shows that for each $s \rightarrow s'$, t_Obl is computed using s' and φ , and the results are OR-ed. This is because $EX(\varphi)$ is satisfied at s if there exists one next state that satisfies φ . The second disjunct states that if s is a variation point, then one of its future next states (there could be one or several), which is a state of a new feature connected to it, will have the obligation to satisfy φ .

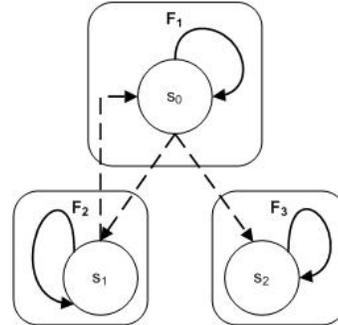


Figure 3. Feature composition example.

Example. Figure 3 shows three features with the behavior of each represented by a state with a self-loop. Inter-feature transitions are shown as broken lines. All states in the example are variation points. Given the CTL property $\varphi = E(tt \cup p)$ to verify over the three-feature composition $Comp_{seq}(FM_1, FM_2, FM_3)$, we firstly compute the variation-point obligation for F_1 , shown in Figure 4. The downward arrows in Figure 4 show the invocation of the t_Obl and

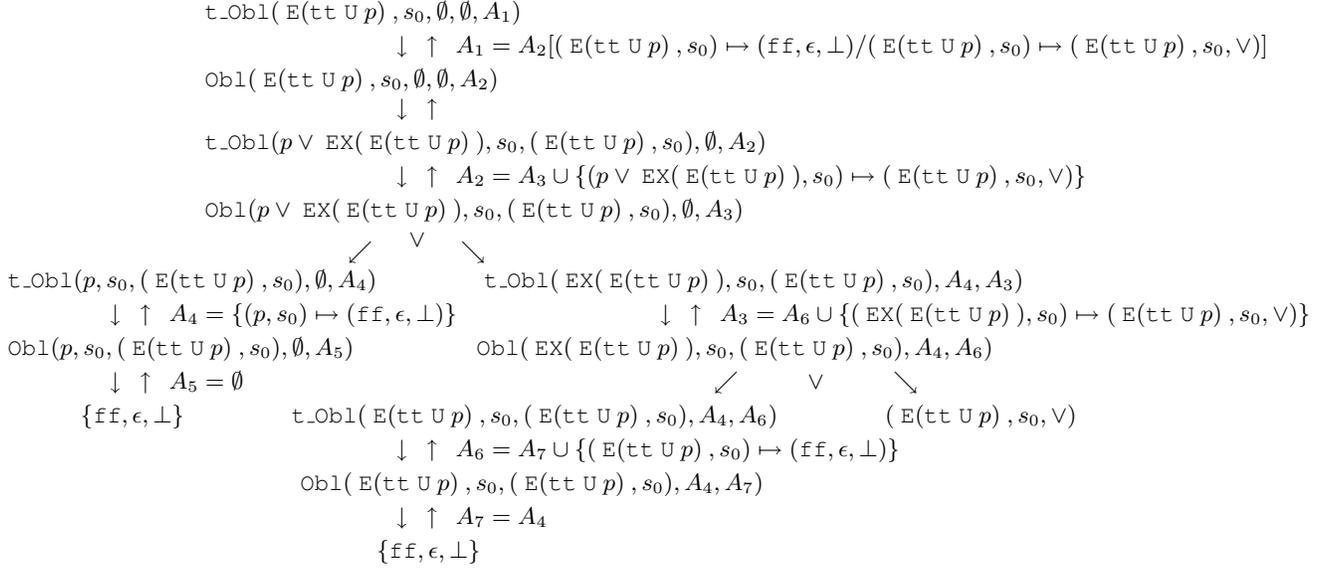


Figure 4. Example of Computing Variation-point Obligations

Obl functions, while the upward arrows show the updates to *aSet*. The variation-point obligations for F_2 and F_3 are computed in a similar fashion.

Step 2: Updating *aSet*. After the computation of variation-point obligation terminates for one FM, *aSet* is updated with new results $(\text{tt}, \epsilon, \perp)$ and $(\text{ff}, \epsilon, \perp)$ in order to incorporate information regarding whether the state s satisfies a formula:

$\text{update}(\text{aSet}) := \text{aSet}[(\varphi, s) \mapsto \psi / (\varphi, s) \mapsto \psi[\psi_i/\psi'_i]]$
 where $\psi_i := (\varphi_i, s_i, \text{op}_i)$ and $s_i \rightarrow t_i \in T_c^{\text{FM}_m, \text{FM}_n}$ and
 $(\varphi_i, t_i) \mapsto (pc, \epsilon, \perp) \in \text{aSet} \wedge pc \in \{\text{tt}, \text{ff}\}$ and

$$\psi'_i = \begin{cases} \psi_i & \text{if } (pc = \text{tt}) \wedge (\text{op}_i = \wedge) \\ \psi_i & \text{if } (pc = \text{ff}) \wedge (\text{op}_i = \vee) \\ (\text{tt}, \epsilon, \perp) & \text{if } (pc = \text{tt}) \wedge (\text{op}_i = \vee) \\ (\text{ff}, \epsilon, \perp) & \text{otherwise} \end{cases}$$

The function states that the entry $(\varphi, s) \mapsto \psi$ in *aSet* is updated to $(\varphi, s) \mapsto \psi[\psi_i/\psi'_i]$ (ψ_i , a subformula of ψ , is replaced by ψ'_i). ψ_i is a variation-point obligation of the form $(\varphi_i, s_i, \text{op}_i)$ that was computed for a FM. If the state t_i of another FM is connected to the variation point s_i and there exists an entry $(\varphi_i, t_i) \mapsto (pc, \epsilon, \perp)$ ($pc \in \{\text{tt}, \text{ff}\}$) in the *aSet*, then we can use (pc, ϵ, \perp) to update ψ_i in ψ . For example, if $\text{op}_i = \wedge$, indicating that all next states of s_i should satisfy φ_i , then in the case $pc = \text{tt}$, ψ_i remains unaltered since the satisfiability of φ_i in one next state does not prove that φ_i is satisfied in all next states; on the other hand, if $pc = \text{ff}$, then it can be concluded that the variation-point obligation has not been satisfied at s_i .

Example. Continuing our example, after variation-point

obligations for F_1 , F_2 and F_3 are computed,

$\text{aSet} = \{(\varphi, s_0) \mapsto (\varphi, s_0, \vee),$
 $(p \vee \text{EX}(\varphi), s_0) \mapsto (\varphi, s_0, \vee), (p, s_0) \mapsto (\text{ff}, \epsilon, \perp),$
 $(\text{EX}(\varphi), s_0) \mapsto (\varphi, s_0, \vee), (\varphi, s_1) \mapsto (\varphi, s_1, \vee),$
 $(p \vee \text{EX}(\varphi), s_1) \mapsto (\varphi, s_1, \vee), (p, s_1) \mapsto (\text{ff}, \epsilon, \perp),$
 $(\text{EX}(\varphi), s_1) \mapsto (\varphi, s_1, \vee), (\varphi, s_2) \mapsto (\text{ff}, \epsilon, \perp),$
 $(p \vee \text{EX}(\varphi), s_2) \mapsto (\text{ff}, \epsilon, \perp), (p, s_2) \mapsto (\text{ff}, \epsilon, \perp),$
 $(\text{EX}(\varphi), s_2) \mapsto (\text{ff}, \epsilon, \perp)\}.$

Applying update to the above *aSet* does not change any obligations.

Step 3: Identifying Inter-Feature Loops from *aSet*. To summarize, once the variation-point obligations have been computed for all FMs (step 1-2) and for every $(\varphi, s) \mapsto \psi$ in *aSet*, each subformula of ψ , $(\varphi_i, s_i, \text{op}_i)$, has a corresponding $(\varphi_i, s_i) \mapsto \psi'$ in *aSet*, we can conclude that no further updates to *aSet* can be computed.

We can now search for any chain of variation-point obligations from *aSet* to identify loops between features. An example of such a chain is the circular dependency between FM_1 and FM_2 in Fig 3: $(\varphi, s_0) \mapsto (\varphi, s_0, \vee)$ and $(\varphi, s_1) \mapsto (\varphi, s_1, \vee)$, where $s_0 \rightarrow s_1$ and $s_1 \rightarrow s_0$. The search for the inter-feature loops is done by applying the $\text{update}_F(\text{aSet})$ function on each element in *aSet*.

$\text{update}_F(\text{aSet}) = \text{aSet}[(\varphi, s) \mapsto \psi / (\varphi, s) \mapsto (pc, \epsilon, \perp)]$
 where $pc = \text{INTERP}((\varphi, s), \emptyset)$

Algorithm 1 computes INTERP, taking as input parameters (φ, s) and the set *Dep* which records the elements on which the mapping result of (φ, s) depends. If $(\varphi, s) \mapsto (pc, \epsilon, \perp)$, then the mapping result does not depend on other

Algorithm 1 Analysis for Inter-Feature Loops

```
1: procedure INTERP( $(\varphi, s), \text{Dep}$ )
2:   if  $(\varphi, s) \mapsto (pc, \epsilon, \perp)$  then
3:     return  $pc$ 
4:   end if
5:   if  $(\varphi, s) \in \text{Dep}$  then
6:     if  $\varphi$  is gfp then
7:       return true
8:     else return false
9:     end if
10:  end if
11:   $(\varphi, s) \mapsto (\varphi_1, s_1, op) \dots op(\varphi_k, s_k, op) \in \text{aSet}$ 
12:   $\text{Next} := \bigcup_{1 \leq i \leq k} \{(\varphi_i, t_i) \mid s_i \rightarrow t_i \in T_c^{\text{FM}_m, \text{FM}_n}\}$ 
13:  if  $(op = \wedge)$   $\text{res} = \text{tt}$  else  $\text{res} = \text{ff}$ 
14:  for  $(\varphi', t') \in \text{Next}$  do
15:     $\text{res} = \text{res op INTERP}((\varphi', t'), \text{Dep} \cup \{(\varphi, s)\})$ 
16:  end for
17:  return  $\text{res}$ 
18: end procedure
```

elements, and the procedure immediately returns pc . In this case the answer has been resolved. Otherwise (Lines 5–10), if the mapping result is in the set Dep , a circular dependency is identified, and the return result is computed based on whether or not φ is a greatest or a least fixed point formula (similar to the way we detected intra-feature loops during variation-point obligation computation using the history set, see Figure 2). Finally, in Lines 11–17, the algorithm computes the dependency of results across features. This is firstly performed in Lines 11 and 12, where for each (φ_i, s_i, op) , t_i (the state connected to s_i) is identified and collected in the set Next . Then INTERP is recursively invoked on each element of Next and the result is aggregated based on whether op is \wedge or \vee (Line 13–17).

Example. We perform $\text{update}_F(\text{aSet})$ on each element in the previously-computed aSet (the sequence does not matter). For example, if the following element is picked first: $(\varphi, s_1) \mapsto (\varphi, s_1, \vee)$, we compute $\text{INTERP}((\varphi, s_1), \emptyset) = \text{ff} \vee \text{INTERP}((\varphi, s_0), \{(\varphi, s_1)\}) = \text{ff} \vee \text{INTERP}((\varphi, s_1), \{(\varphi, s_1), (\varphi, s_0)\}) = \text{ff}$. We then replace $(\varphi, s_1) \mapsto (\varphi, s_1, \vee)$ with $(\varphi, s_1) \mapsto (\text{ff}, \epsilon, \perp)$ in the aSet . Other elements of aSet are updated in a similar fashion until no change can be done to aSet . This means that all circular dependencies between these features have been resolved.

Summary: Compositional Algorithm. Algorithm 2 presents our compositional algorithm introduced in Section 4 using the functions described above. This algorithm model checks the current product (a composition of features), reusing results from previous model-checking of those features for other products in the product line.

To summarize, given a composition

Algorithm 2 Compositional Model Checking

```
1: procedure COMPOSE( $\text{Comp}_{seq}(\text{FM}_1, \text{FM}_2, \dots, \text{FM}_m), \varphi$ )
2:    $\bigwedge_{s_0 \in S_0^1} \text{t\_Obl}(\varphi, s_0, \emptyset, \emptyset, \text{aSet})$ 
3:   repeat
4:      $\text{tmp} := \text{aSet}$ 
5:     for each  $((\varphi', s') \mapsto \psi) \wedge (\psi \neq (pc, \epsilon, \perp))$  do
6:        $\psi := \text{getObl}(\psi, \text{aSet}, \text{aSet}_{\text{new}})$ 
7:        $\text{aSet} := \text{update}(\text{aSet}_{\text{new}})$ 
8:     end for
9:   until  $(\text{aSet} = \text{tmp})$ 
10:  return  $\text{update}_F(\text{aSet})$ 
11: end procedure
```

$\text{Comp}_{seq}(\text{FM}_1, \text{FM}_2, \dots, \text{FM}_m)$ and a formula φ , the algorithm first obtains the variation-point obligations for FM_1 such that all its start states can satisfy φ (Line 2). In Lines 5 and 6, the variation-point obligations of the other features connected to the variation points of FM_1 are computed using the function getObl , defined as follows:

$$\begin{aligned} \text{getObl}((\varphi, s, op), \text{aSet}_{in}, \text{aSet}_{out}) &:= \\ \text{Op}_{s \rightarrow s'}[\text{t_Obl}(\varphi, t, \emptyset, \text{aSet}_{in}, \text{aSet}_{out})] & \\ \text{where } s \rightarrow t \in T_c^{\text{FM}_m, \text{FM}_n}, op \in \{\vee, \wedge\} & \end{aligned}$$

$$\begin{aligned} \text{getObl}(\psi_1 \text{ op } \psi_2, \text{aSet}_{in}, \text{aSet}_{out}) &:= \\ \text{getObl}(\psi_1, \text{aSet}_{in}, \text{aSet}_{temp}) \text{ op} & \\ \text{getObl}(\psi_2, \text{aSet}_{temp}, \text{aSet}_{out}) & \end{aligned}$$

The process of computing the variation-point obligation is iterated (Line 3–9) until no more updates on the aSet can be made (Line 9). At this point, the function $\text{update}_F(\text{aSet})$ is invoked to identify loops between features and infer results from variation-point obligations represented in greatest and least fixed point formulas in the aSet . We say that $\text{Comp}_{seq}(\text{FM}_1, \text{FM}_2, \dots, \text{FM}_m) \models \varphi$ when for all start states s_0 of FM_1 , $((\varphi, s_0) \mapsto (\text{tt}, \epsilon, \perp)) \in \text{update}_F(\text{aSet})$. At this point, satisfaction of the property of interest in the composed product has been determined. The remark below follows from the above discussion.

Remark 1. $\text{Comp}_{seq}(\text{FM}_1, \text{FM}_2, \dots, \text{FM}_m) \models \varphi \Leftrightarrow \forall s_0 \in S_0^1 : (\varphi, s_0) \mapsto (\text{tt}, \epsilon, \perp) \in \text{COMPOSE}(\text{Comp}_{seq}(\text{FM}_1, \text{FM}_2, \dots, \text{FM}_m), \varphi)$.

Example. So far, we have performed every step of Algorithm 2. Since $(\varphi, s_0) \mapsto (\text{ff}, \epsilon, \perp) \in \text{update}_F(\text{aSet})$, the verification result is that $\text{Comp}_{seq}(\text{FM}_1, \text{FM}_2, \text{FM}_3) \not\models \varphi$.

6 Case Study

We have implemented our algorithm in a research prototype model checker. Detailed information is provided in the following link: <http://www.cs.iastate.edu/~janetlj/ase08/>.

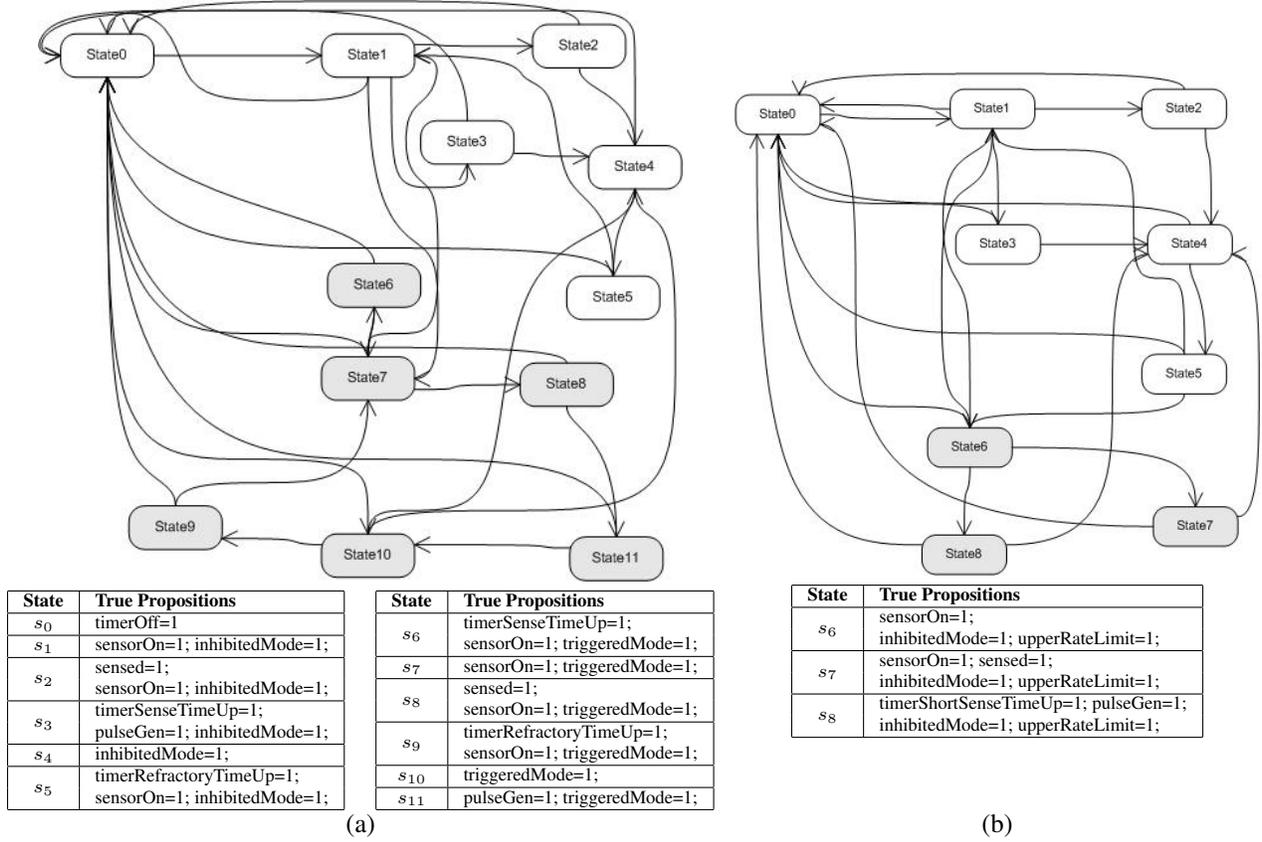


Figure 5. (a) Base Controller and Mode-Transitive Extension. (b) Base Controller and Rate-Responsive Controller Extension.

We evaluated our technique by conducting experiments on the pacemaker product line discussed in Section 2. Figures 5(a) and (b) depict how the Mode Transitive Extension (FM_1) and Rate Responsive Extension (FM_2) are sequentially composed with BasePacemaker’s controller (FM_0). The states in the extensions are shown in grey. The variation points are the states which have outgoing transitions leading to another model in the composition. The propositions satisfied at each state are shown in the corresponding tables below each figure.

It is worth mentioning that although the mode-transitive feature in Fig. 5(a) extends the base controller functionality in a sequential manner, the rate-responsive feature does so in a more complicated manner. This is because it also introduces a new component (extra sensor). We model its effect on the extension to the base controller’s functionality by introducing an abstract event in the extension (“upperRateLimit=1” in Figure 5(b)’s Table). This modeling strategy works because the CTL property to be verified does not depend on the behavior of the extra sensor.

In the controller for the fourth product (not shown), ModeTransitive-RateResponsive Pacemaker, FM_1 and FM_2 are sequentially composed with FM_0 , as in Figure 5 (a) and (b). FM_1 and FM_2 do not connect directly; a pacemaker controller can be in either the Triggered Mode or Inhibited Mode, but not in both at the same time.

The following CTL formula describes the required property for the product line that was textually introduced in Section 2:

$$AG((sensed=0 \wedge timerSenseTimeUp=1 \wedge inhibitedMode=1) \Rightarrow EF(pulseGen=1 \wedge inhibitedMode=1))$$

This formula was used in the following evaluation.

To evaluate the space and time performance of our approach, we compared compositional model checking (CMC) with non-compositional model checking (NMC) for the four products in the product line and recorded the experimental results in Table 1. In the table, MT denotes ModeTransitive, while RR denotes RateResponsive. (t_Ob1) and ($Ob1$) record the number of times the t_Ob1

Table 1. Test Data for Pacemaker Product Line

# of Invocations	Non-Compositional Model Checking (NMC)				Compositional Model Checking (CMC)			
	Base	MT	RR	MT-RR	Base	MT	RR	MT-RR
	Pacemaker	Pacemaker	Pacemaker	Pacemaker	Pacemaker	Pacemaker	Pacemaker	Pacemaker
(t_Obl)	33	68	52	87	40	51	30	28
(Obl)	20	38	29	47	22	22	12	5
(Obl _{top_test})	65	125	94	154	71	76	39	20
(Obl _{test})	64	124	93	153	66	65	32	6
(add _{obl})	91	179	135	222	102	126	72	61
# of State Visits	13	30	23	40	16	23	13	14

and `Obl` functions are performed, respectively. Similarly, `(Obltop_test)` and `(Obltest)` record the times these two assistant functions in our implementation have been invoked. The assistant functions serve to conduct lightweight tests (by "lightweight" we mean that they check only the current state), e.g., to see if one branch of a disjunct formula is true during obligation computation. `(addobl)` records the number of times an obligation is added to our implementation of `aSet`. Finally, # of state visits records the number of times the states in a model are visited.

The results of NMC are obtained from checking each of the four products in its entirety, without breaking it into separate features. No variation points are specified, and the start state is always resolved to a true or false result at the end of the checking.

The results of CMC are obtained from calculating the test data for the added features and connections in each product. For example, verifying the `BasePacemaker` involves checking FM_0 with states 1, 4, and 5 as its variation points (i.e., the union set of the variation points needed by the MT and RR extensions), plus applying the `updateF(aSet)` to get the final result. We found that, as expected, generating obligations at the variation points and performing `updateF(aSet)` introduced a slight overhead for the base product that was amortized over its subsequent reuse. Verifying the `RateResponsivePacemaker` involves checking the FM_2 with states 6, 7, and 8 as its variation points, plus checking the connections from FM_2 to FM_0 , and applying the `update` and `updateF(aSet)` functions. Test results were similarly collected for the other two products.

Table 1 shows that the compositional model checking approach does provide savings in the product line. For example, in NMC the cost for checking the product line (measured in `(t_Obl)`) was $33+68+52+87=240$, while the cost in CMC was $40+51+30+28=149$. This is because the common features (e.g., FM_0) were checked repeatedly in NMC. If no prior checking for any of the features had been done, the cost for checking the `RateResponsivePacemaker` (measured in `(t_Obl)`) would have been $40+30=70$, which is more than the value of 52 for the NMC. This difference is due to the cost of generating assets for reuse, i.e., generating

obligations at the variation points and maintaining a different `aSet` (for applying the `update` and `updateF(aSet)` functions) for each composition. To summarize, as with the product-line approach itself, CMC shows savings when features are reused.

We briefly note here the suitability of this approach for *product-line evolution*. Since any state in a model can become a variation point, when structural changes to a feature model occur, the affected states can be treated as variation points. For example, if the state s that was not a variation point becomes a new variation point, the answer set elements involving s identify temporal obligations for any features composed at s . This allows us to model check whether the new variation-point obligations are preserved after the change. Because product lines routinely experience significant change over their lifetimes, the continued usefulness of previous model-checking results to the product line development contributes to the practicality of this technique.

7 Related Work

Several recent works have investigated representations of variability within a product line in behavioral models. Fantechi and Gnesi identify whether a product belongs to a product line [6]. Fischbein, Uchitel and Braberman propose a technique to determine whether a variability undermines a product-line property [7]. Lauenroth and Pohl describe how variability complicates the consistency checking of a product in the product line [13]. This line of work does not treat the common and variable functionality as equal units to be composed. Instead, they have a well-defined base with relatively small variations (e.g., transitions in a finite state machine) that may be verified through techniques like behavioral conformance [7]. The variations that they can verify do not cover all the ones that exist in our case.

We now compare our proposed technique with techniques whose main objective is to effectively use compositional verification in a product line setting. In the context of open-system verification where features are added in a sequential fashion, our work is closely related to work by Blundell, Fisler, Krishnamurthi and Hentenryck [3]. The authors propose a framework in which interface obligations

are generated as temporal properties. Our technique differs from theirs in that [3] requires interface states (here, the variation points) to be terminal states with no outgoing transitions, while our approach does not have this restriction. Secondly, we permit features to be added in ways that allow intra and inter-feature loops, a flexibility that was needed to accurately model the pacemaker product line. Wang extends [3] in [18] and allows inter-feature loops. However, that work assumes that interface states are sufficient for composing different features and does not re-explore the non-interface states. This implicitly puts a restriction on the type of inter-feature loops that can be verified.

In [17], Thang presents the necessary conditions which, when satisfied by the base and its extension feature(s), ensure that the property verification results hold before and after the base is extended with the corresponding features. Though the work allows loops between the base and the extensions, it does not provide insights into the cases where the necessary conditions are violated.

Our work falls into the category of compositional verification [1]. We use sequential composition (Def. 2) rather than parallel composition, as in, e.g., Giannakopoulou, Pasareanu and Barringer [8], because it would add unnecessary complexity to the state space and obscure the interfaces among features that we want to maintain in a product-line setting for effective reuse.

8 Conclusion

This paper presents an incremental and compositional model-checking technique for performing sequential composition of different features in a product-line setting. By computing and managing variation-point obligations, we enable reuse of previous verification results when a new product is composed. Re-checking occurs only when and as needed. Additionally, this approach provides more flexibility in how features interact than existing techniques, bringing models more in line with real-world product lines. Evaluation done using a prototype implementation to model check a simplified pacemaker product line shows that this technique can reduce the amount of re-verification needed to assure that a required property holds for each new product in the product line.

Acknowledgements This research was supported by the National Science Foundation under grants 0541163, 0702758 and 0709217.

References

[1] M. Abadi and L. Lamport. Conjoining specifications. *ACM Trans. Program. Lang. Syst.*, 17(3):507–535, 1995.
 [2] D. Batory, D. Benavides, and A. Ruiz-Cortes. Automated analysis of feature models: challenges ahead. *Commun. ACM*, 49(12):45–47, 2006.

[3] C. Blundell, K. Fisler, S. Krishnamurthi, and P. V. Hentenryck. Parameterized interfaces for open system verification of product lines. In *ASE '04*, pages 258–267, Washington, DC, USA, 2004. IEEE Computer Society.
 [4] E. M. Clarke, O. Grumberg, and D. A. Peled. *Model Checking*. MIT Press, 2000.
 [5] K. A. Ellenbogen and M. A. Wood. *Cardiac Pacing and ICDs (fourth edition)*. Blackwell Publishing, Inc., 2005.
 [6] A. Fantechi and S. Gnesi. A behavioural model for product families. In *ESEC-FSE '07*, pages 521–524, New York, NY, USA, 2007. ACM.
 [7] D. Fischbein, S. Uchitel, and V. Braberman. A foundation for behavioural conformance in software product line architectures. In *ROSATEA '06*, pages 39–48, New York, NY, USA, 2006. ACM.
 [8] D. Giannakopoulou, C. S. Pasareanu, and H. Barringer. Assumption generation for software component verification. In *ASE '02*, page 3, Washington, DC, USA, 2002. IEEE Computer Society.
 [9] K. Havelund, M. R. Lowry, and J. Penix. Formal analysis of a space-craft controller using SPIN. *TSE*, 27(8):0098–5589, 2001.
 [10] M. Huth and M. Ryan. *Logic in Computer Science: modelling and reasoning about systems (second edition)*. Cambridge University Press, 2004.
 [11] I. Jacobson, M. Griss, and P. Jonsson. *Software Reuse: Architecture, Process and Organization for Business Success*. Addison-Wesley, 1997.
 [12] R. Kaivola. Formal verification of pentium 4 components with symbolic simulation and inductive invariants. In K. Etessami and S. K. Rajamani, editors, *CAV*, volume 3576 of *Lecture Notes in Computer Science*, pages 170–184. Springer, 2005.
 [13] K. Lauenroth and K. Pohl. Towards automated consistency checks of product line requirements specifications. In *ASE '07*, pages 373–376, New York, NY, USA, 2007. ACM.
 [14] B. Littlewood and L. Strigini. Validation of ultrahigh dependability for software-based systems. *Commun. ACM*, 36(11):69–80, 1993.
 [15] J. Liu, J. Dehlinger, and R. Lutz. Safety analysis of software product lines using state-based modeling. *J. Syst. Softw.*, 80(11):1879–1892, 2007.
 [16] J. Liu, J. Dehlinger, H. Sun, and R. Lutz. State-based modeling to support the evolution and maintenance of safety-critical software product lines. In *ECBS '07*, pages 596–608, Washington, DC, USA, 2007. IEEE Computer Society.
 [17] N. T. Thang. *Incremental Verification of Consistency in Feature-Oriented Software*. PhD thesis, Japan Advanced Institute of Science and Technology, 2005.
 [18] X. Wang. A modular model checking algorithm for cyclic feature compositions. Master's thesis, Worcester Polytechnic Institute, 2005.
 [19] D. L. Webber and H. Goma. Modeling variability in software product lines with the variation point model. *Sci. Comput. Program.*, 53(3):305–331, 2004.
 [20] D. M. Weiss and R. Lai. *Software Product Line Engineering: A Family-Based Software Development Process*. Addison-Wesley, 1999.