

A Toolchain for the Detection of Structural and Behavioral Latent System Properties*

Adam C. Jensen, Betty H.C. Cheng, Heather J. Goldsby

Michigan State University, East Lansing MI 48824, USA
{acj,chengb,hjg}@cse.msu.edu

Edward C. Nelson

Ford Research and Advanced Engineering, Dearborn, MI 48121, USA
enelson7@ford.com

May 19, 2011

Abstract

The cost to repair a requirements-based defect in software-based systems increases substantially with each successive phase of the software lifecycle in which the error is allowed to propagate. While tools exist to facilitate early detection of design flaws, such tools do not detect flaws in system requirements, thus allowing such flaws to propagate into system design and implementation. This paper describes an experience report using a toolchain that supports structural and behavioral analysis of UML state diagrams not currently available in commercial UML modeling tools. With the toolchain, models can be incrementally and systematically improved through syntax-based analysis, type checking, and detection of latent behavioral system properties, including feature interactions. This paper demonstrates with the analysis of industry-provided models that the proposed toolchain is an effective means for discovering unwanted latent system properties in the late requirements phase of development, thus reducing the number of defects propagated to successive phases.

1 Introduction

In software development, the cost to repair a defect increases substantially with each successive phase of the software lifecycle [1, 2]. When a defect is allowed to propagate into the design and implementation phases, the number of artifacts (e.g., models and documentation) that are affected by it also increases. Typically, during the requirements phase the system's stakeholders describe the key needs and problems that the system-to-be should address, usually using natural language. As a means to clarify and refine requirements that have been expressed in natural language, developers construct *domain models* that identify the key elements of the

*This work has been supported in part by NSF grants CCF-0541131, IIP-0700329, CCF-0750787, CCF-0820220, DBI-0939454, CNS-0854931, Army Research Office grant W911NF-08-1-0495, Ford Motor Company, and a Quality Fund Program grant from Michigan State University.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

| | | | |
|---|------------------------------------|---|----------------------------------|
| 1. REPORT DATE 19 MAY 2011 | 2. REPORT TYPE | 3. DATES COVERED 00-00-2011 to 00-00-2011 | |
| 4. TITLE AND SUBTITLE A Toolchain for the Detection of Structural and Behavioral Latent System Properties | | 5a. CONTRACT NUMBER | |
| | | 5b. GRANT NUMBER | |
| | | 5c. PROGRAM ELEMENT NUMBER | |
| 6. AUTHOR(S) | | 5d. PROJECT NUMBER | |
| | | 5e. TASK NUMBER | |
| | | 5f. WORK UNIT NUMBER | |
| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Michigan State University, East Lansing, MI, 48824 | | 8. PERFORMING ORGANIZATION REPORT NUMBER | |
| 9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) | | 10. SPONSOR/MONITOR'S ACRONYM(S) | |
| | | 11. SPONSOR/MONITOR'S REPORT NUMBER(S) | |
| 12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited | | | |
| 13. SUPPLEMENTARY NOTES | | | |
| 14. ABSTRACT The cost to repair a requirements-based defect in software-based systems increases substantially with each successive phase of the software lifecycle in which the error is allowed to propagate. While tools exist to facilitate early detection of design flaws, such tools do not detect flaws in system requirements thus allowing such flaws to propagate into system design and implementation. This paper describes an experience report using a toolchain that supports structural and behavioral analysis of UML state diagrams not currently available in commercial UML modeling tools. With the toolchain, models can be incrementally and systematically improved through syntax-based analysis, type checking, and detection of latent behavioral system properties, including feature interactions. This paper demonstrates with the analysis of industry-provided models that the proposed toolchain is an effective means for discovering unwanted latent system properties in the late requirements phase of development, thus reducing the number of defects propagated to successive phases. | | | |
| 15. SUBJECT TERMS | | | |
| 16. SECURITY CLASSIFICATION OF: | | | 17. LIMITATION OF ABSTRACT |
| a. REPORT unclassified | b. ABSTRACT unclassified | c. THIS PAGE unclassified | Same as Report (SAR) |
| | | | 18. NUMBER OF PAGES 17 |
| | | | 19a. NAME OF RESPONSIBLE PERSON |

system and their relationships to one another, as well as their relationships to external elements. In order to better understand the required behavior, developers often create prototypes or state-based representations based on the domain model. While simulations and executable prototypes enable *validation* of requirements, it is equally important to be able to *verify* requirements to identify inconsistencies, (invariant) property violations, etc. Thus, there is a need for tools that identify errors in requirements specifications based on analysis of early prototype models. This paper presents a toolchain that facilitates the detection of syntactic and semantic errors in state-based diagrams and also identifies properties that specify *latent behavior*, the unspecified and potentially unwanted behavior of the model.

Many tools have been developed to support model-driven engineering of software systems. Tools such as ArgoUML [3], Rational Software Architect [4], and Microsoft Visio [5] support visual modeling of software designs via the Unified Modeling Language (UML). IBM Rational Rhapsody [6] supports UML modeling as well as code generation and many consistency tests to ensure that the system under development is free of syntax errors. However, none of these tools performs syntax or type checking on state transition expressions in state diagrams. Particularly for applications involving complex logic and system behavior (e.g., embedded systems), transitions may contain complex guards and action statements that often define the core functionality of the system being modeled. Thus, tools that treat the transition expressions as uninterpreted strings allow subtle errors to propagate into the source code that is generated from the model, particularly in the context of model-driven engineering. Furthermore, while tools such as Rhapsody provide traceability from requirements to source code, to the best of the authors' knowledge no existing commercial or research tools provide the comprehensive automated identification of the collection of different types of errors covered by our toolchain for UML models.

This paper describes an experience report from using a newly-developed toolchain that supports syntax and type checking as well as detection of latent system properties. After requirements have been elicited for an embedded system, developers often build a domain model using class diagram syntax that describes the key elements of the system (including physical elements, such as sensors and actuators, and software elements, such as controllers) and elements in the environment with which the system interacts. A state diagram is created for each key element, resulting in a collection of interacting state diagrams. While such diagrams are useful for refining system requirements and may be used during the design phase, there is limited tool support for detecting errors in syntax and semantics, and to our knowledge there is no tool support for automatically identifying latent properties. The proposed toolchain has two key advantages over current approaches. First, all state transition expressions are parsed and type-checked, thus identifying many errors that existing tools do not address until the code generation phase. Second, automated detection of latent properties enables system developers to identify so-called *blind spots* in system requirements. Blind spots are missing or incomplete requirements that are overlooked by requirements engineers, and they are often discovered only after the system has been partially implemented or, worse yet, deployed to the field. By identifying these errors early in the development process and suggesting resolution strategies when possible, the proposed toolchain minimizes the number of subtle design defects and the cost of redesigning the system to correct the defects.

The proposed toolchain comprises three main tools: CYCLOPS, a model pre-processor that identifies common syntax and semantics errors in behavioral models specified in XMI (XML Metadata Interchange) format; HYDRA, a tool for translating UML behavioral models into Promela, the formal language for the

SPIN model checker [7]; and MARPLE, a tool for automatically generating properties that are satisfied by the model and may represent latent and potentially erroneous behavior. We apply this toolchain to an industrial software system from the automotive embedded systems domain. The software system comprises three subsystems: Lighting, Power Management, and Windshield Wipers. The Lighting subsystem handles all functionality related to interior lamps, headlights, and tail lights. The Power Management subsystem monitors and controls the ignition status, vehicle speed, door statuses, battery status, and other electronic features. The Windshield Wipers subsystem controls the movement and speed of the windshield wipers. The subsystems are sophisticated and interact with one another at run time, thus creating the potential for errors in modeling semantics, unintended behavior that spans multiple subsystems, and feature interactions.

Based on feedback from the developer of the model, it is clear that several of the detected errors would have been very difficult and time-consuming to detect and resolve without the use of the toolchain. The remainder of the paper is organized as follows. In Section 2, we discuss background concepts. We present the software model that was studied in this work in Section 3. Next, we describe the process of using the toolchain in Section 4. Section 5 describes related work. Our experience of applying the toolchain to an automotive embedded systems model is presented in Section 6. We discuss the results and consequences of applying the toolchain in Section 7. Finally, we present our conclusions and discuss future work in Section 8.

2 Background

In this section, we discuss background concepts and enabling technologies that support the proposed toolchain, including the Unified Modeling Language, the SPIN model checker, evolutionary computation, and novelty search. These enabling technologies are presented according to the tool(s) that leverage their capabilities.

2.1 CYCLOPS and HYDRA

CYCLOPS and HYDRA have been developed to support the construction of models in the Unified Modeling Language (UML) [8], the *de facto* standard in object-oriented software modeling. They enable developers to perform extensive error checking on UML models that describe system prototypes and support the translation of UML state diagrams into Promela for analysis with the SPIN model checker.

Unified Modeling Language. The Unified Modeling Language (UML) is a general-purpose visual modeling language that is used for modeling object-oriented software. It comprises several types of diagram notations, including support for class diagrams, interaction diagrams, state machine diagrams, and others. A UML model may contain many different diagrams that describe portions of the same system. For the purposes of this paper, we assume the use of UML version 1.5 and focus on state machine diagrams. A state machine diagram (hereafter, “state diagram”) describes the various states in which a system can be and the transitions between the states. Visually, a state diagram comprises rounded rectangles (representing states) and lines with arrows that indicate transitions between states. The lines are annotated with optional guards and trigger events that denote the events that trigger a transition and the actions generated as a result of the transition, respectively.

SPIN Model Checker. The SPIN model checker [7] is a tool for exhaustively verifying state-based models. It takes a model expressed in Promela and produces a model checker in C code. SPIN uses nondeterministic automata to check properties expressed in Linear Temporal Logic and performs exhaustive analysis of a system’s state space in order to identify error conditions and other undesirable system behaviors. It was originally developed to formally analyze distributed systems, and in recent years it has been used to analyze telecommunications protocols [9].

2.2 MARPLE

MARPLE is a tool that automatically discovers latent properties in UML state diagrams. It leverages novelty search, an evolutionary search technique, and formal model analysis to generate a list of properties that describe the behavior specified by the model.

Evolutionary Computation. A family of optimization techniques modeled after Darwinian evolution by natural selection, evolutionary computation (EC) explores large solution spaces using biologically-inspired concepts such as mutation and selection [10]. EC is effective for finding solutions to problems that have large solution spaces that cannot be exhaustively explored in a reasonable amount of time. A specific solution to a given problem (e.g., finding a set of rules for steering a robotic vehicle on a roadway) is known as an *individual*. EC begins with a large population of randomly-generated individuals. This random population of individuals is known as the initial *generation*. Each individual is evaluated to determine its fitness for a given task. For a robotic vehicle, the fitness of an individual might be the percentage of time that the vehicle spent within the boundaries of the roadway. Once each individual’s fitness is computed, EC probabilistically selects a set of individuals that will represent the next generation. A individual that has a higher fitness is more likely to be selected than a individual with a low fitness. Before an individual is placed into the next generation, it undergoes a small change known as a *mutation*. For example, a single rule might be replaced or modified. It may also undergo a change known as *crossover* in which portions of two individuals are combined to form a new offspring individual that contains traits from each of the two ancestral individuals. This process of selection, mutation, and evaluation continues until a fixed number of generations have passed. If an individual representing an optimal solution exists and has been found, then it is returned as the output. Otherwise, the individual with the highest fitness is returned.

Novelty Search. One EC technique, known as *novelty search* [11], replaces the explicit fitness computation with a novelty function that measures how different each individual is from other individuals in the population and in an archive of previous individuals. Novelty search then selects individuals whose behavior is the most distant (i.e., the most novel), thus increasing the diversity in the population and exploring the solution space more efficiently than a random search. The specific measure of distance between individuals varies with the problem being solved, but a Euclidean distance is typically used when the behavior of an individual can be mapped to a numerical vector.

3 Body Subsystem Model

In this section, we describe the **Body Subsystem** model that was used in this study. The model describes and simulates embedded devices that control the electronic subsystems of a modern passenger automobile and was designed for the purposes of requirements elicitation and analysis.¹ The subsystems of the model include interior and exterior lighting, power management, and windshield wiper control. While the onboard electronics involves several more subsystems, these three were selected for our focus because they exhibit known, intended interactions. The objective was to investigate whether the subsystems also exhibit unknown interactions. The remainder of this section provides a brief description of each subsystem under study.

3.1 Lighting Subsystem

The **Lighting** subsystem comprises 16 classes and is responsible for managing interior lights, including map, vanity, trunk, and under-hood lamps; and exterior lights, including head lights (low- and high-beam) and tail lights. The subsystem also contains classes that monitor the intensity of ambient light in order to control day time running lights and activate the vehicle’s head lights and tail lights for night time driving. The domain model for the **Lighting Subsystem** is shown in Figure 1(a).

3.2 Power Management Subsystem

The **Power Management** subsystem comprises 25 classes and is responsible for monitoring ignition status, sleep mode status, battery voltage, and commands from remote key fobs. The subsystem responds to events such as the insertion of an ignition key, exceeding vehicle speed thresholds, and the firing of timers. A partial domain model for the **Power Management** subsystem is shown in Figure 1(b). The **PowerManagementFeature** class monitors the ignition status for the vehicle and coordinates power-saving. Its subclasses, such as **BatterySaver**, extend this functionality for specific components. The classes that are not shown due to space constraints include subclasses of **BatterySaver** and **SleepPowerModeControl** that are responsible for controlling the power status (e.g., awake or asleep) of electronic components, as well as a **Voltage_Range_Monitor** class that monitors battery voltage.

3.3 Windshield Wiper Subsystem

The **Windshield Wiper** subsystem comprises eight classes and is responsible for controlling wiper behavior. The classes represent hardware and software ranging from the low-level motor controller, the washer fluid pump, and a stall sensor that turns off the wiper motor if it detects that the wipers are not moving. The domain model for the **Windshield Wiper** subsystem is shown in Figure 2. Each of the arrows in the domain model indicates a read dependency between two classes.

¹The model was developed by the industrial partner as an example of an industrial-strength model with representative system elements and behavior. It does not contain any proprietary or specific configuration parameters of a deployed vehicle.




Figure 1: Lighting Subsystem and Power Management Subsystem Domain Models




Figure 2: Windshield Wiper Subsystem Domain Model

4 Process

In this section, we provide an overview of the process that was used to apply the proposed toolchain to the Body Subsystem model. A data flow diagram for the process is shown in Figure 3. The process begins with a system model in XMI (XML Metadata Interchange) format. It is assumed that this model has been produced by a visual modeling tool that has support for exporting models in XMI format. In this case, Rhapsody was used because of its support for requirements traceability, code generation capabilities, and support for state-based modeling.

4.1 Model Transformation

First, the XMI model is given to step **1: Model Transformation**. The CYCLOPS tool takes the XMI model and checks for common syntax and semantics errors. For example, it parses and checks each state transition expressions to ensure that they are well-formed and do not refer to undeclared classes, attributes, or operations. CYCLOPS produces specific error messages that indicate the nature of any errors that are discovered, and it makes suggestions when appropriate (e.g., when an attribute from another class is referenced as though it were declared in the current class). CYCLOPS supports an iterative process of analysis,

detection of errors, and model correction. This incremental error-correction cycle is shorter and more interactive than comparable techniques available in commercial tools. For example, using code generation to detect syntactic and semantic errors in a model would require at least one additional step for compilation and linking compared to our toolchain. Once the errors detected by CYCLOPS are resolved, it translates the XMI model into the Hydra Intermediate Language (HIL) that can then be processed by the HYDRA tool.

HYDRA is a model translator initially developed by McUmbler [12]. It takes a model in HIL format and produces an equivalent model in Promela (the PROcess MEta Language). Promela is a formal logic language that was developed to support analysis and exhaustive checking of concurrent systems of communicating processes [13]. Promela models are checked using the SPIN model checker [7], which identifies livelocks, deadlocks, error conditions, and other undesirable behavior. It also has support for verifying arbitrary properties specified using Linear Temporal Logic (LTL) [14].

4.2 Model Analysis

Goldsby and Cheng developed MARPLE, a novelty-search based tool for automatically discovering properties that represented unexpected, or *latent*, behavior in UML state diagrams [15]. Marple searches the space of properties for a given UML state diagram and identifies properties that meet two criteria: (1) the properties hold over all executions of the modeled system, and (2) the properties do not represent system requirements (i.e., known and acceptable properties). Marple uses Linear Temporal Logic (LTL), a modal temporal logic, to specify logical propositions that are interpreted as model properties. Specifically, MARPLE generates properties by instantiating the five most commonly occurring LTL specification patterns [16] with model specific elements. Each proposition is evaluated using the SPIN [7] model checker, and Marple stores a proposition in a data structure called the *archive* if it holds true. Once the archive contains at least one proposition, each new proposition that is evaluated is compared with a portion of the population as well as every proposition in the archive. By including the propositions in the archive in the comparison, the novelty search algorithm is able to “remember” the portions of the solution space that it has explored previously, thus ensuring that the algorithm does not stagnate or become “stuck” in a suboptimal portion of the space.

In step **2: Model Analysis**, we leverage the Promela model produced by HYDRA to search for latent system properties. The Promela model is given as input to the MARPLE tool along with a set of parameters and system-to-be requirements. Parameters for the MARPLE tool include the number of properties that should be returned, the size of the population that the novelty search algorithm should use, and the number of distinct classes that are mentioned in each property. The parameters may be tuned by the system developer according to the model being analyzed and the number of results that are desired. MARPLE then uses novelty search to explore the space of system properties to identify latent properties that are not explicitly identified as being part of the system requirements. As output, MARPLE produces a set of LTL properties, which are presented to the developer in natural language for readability purposes. To enable this natural language property representation, we use a previously developed set of specification patterns that translate between LTL properties and natural language [17].

4.3 Property Review

Finally, in step **3: Property Review**, the latent properties discovered by MARPLE are presented in natural language to the system developer for review. If a given property is desirable, then the developer may consider adding it to the list of explicit system requirements. If the property is undesirable, however, action must be taken to ensure that the property does not continue to hold. For example, the developer might examine the state diagrams for the classes that are mentioned in the property. If an error is discovered in the diagrams, then the model is revised and the toolchain is restarted at step **1**.




Figure 3: Data Flow Diagram

5 Related Work

Many commercial tools, such as ArgoUML [3], Rational Software Architect [4], Microsoft Visio [5], IBM Rational Rhapsody [6], and MATLAB (Simulink and StateFlow) [18], support the creation of UML models, syntax checking, simulation, and code generation capabilities. However, they do not support the automated detection of the full suite of syntactical and semantic error checks for state-based diagrams that we describe in this paper. Additionally, they do not support the identification of latent properties satisfied by the model. In this section, we explore research tools that have been created to address these two challenges.

5.1 Consistency Checking Among UML Class and State Diagrams

One key challenge that arises as the result of using multiple diagrams to provide different views of the same system is maintaining consistency among these different representations. As a result, researchers have developed a number of approaches to support various aspects of consistency checking among UML models (e.g., [19, 20, 21, 22, 23, 24, 25]). The toolchain described by this paper automatically detects inconsistencies in the syntax and semantics of UML class and state diagrams created as part of the late requirements engineering phase of development. Thus, we focus our attention on approaches that examine consistency

among these two diagram types. Simmonds *et al.* use rules presented in terms of description logic [26], a subset of first order predicate logic, to identify inconsistencies among UML class, sequence, and state diagrams during the design phase [24]. However, their approach does not check that the transitions within the state diagram use viable elements from the class diagram. Gomaa *et al.* present an approach to checking the consistency among use case diagrams, class diagrams, sequence diagrams, and state diagrams [22]. Their approach involves specifying consistency checking rules among the various types of diagrams, including class and state diagrams. Their manual approach involves specifying consistency-checking rules among the various types of diagrams, including class and state diagrams. Egyed proposes a Rational Rose plugin that can be used to detect and resolve inconsistencies that arise within UML models at the design phase [19]. Their approach relies upon the specification of consistency rules, which are periodically evaluated. To the best of our knowledge, these consistency rules can detect whether elements of the state diagrams are consistent with those that appear in the class diagram, but do not detect subtle errors, such as assignments that occur within transition guards. Schwarzl and Peischl propose an approach to statically analyzing state diagrams for syntax, existence, data type, communication, non-determinism, and transition hiding errors [23]. As part of this process, the transitions on the state diagrams are checked for well-formedness. The set of syntactical and semantic errors that they detect is a subset of the errors that CYCLOPS detects. However, the behavioral errors that they detect (e.g., deadlock conditions and circular messaging dependencies) are complementary to errors detected by the approach presented in this paper. These approaches enable developers to identify model inconsistencies, but they do not support the identification of latent properties. While it is possible to use the tools developed by Egyed and Schwarzl as a “frontend” to our toolchain, based on descriptions in the related publications it does not appear that they would detect the full suite of errors identified by CYCLOPS and HYDRA.

5.2 Detection of Latent Properties

Several approaches generate temporal logic properties that specify the behavior of systems [27, 28, 29, 30, 31]. Because the objective of our approach is to automatically identify obscure latent properties that might not otherwise be discovered, we focus on how the approaches blend developer knowledge and automation to identify properties. Perracotta [31] is a dynamic inference approach that infers properties from imperfect execution traces, which have been generated by running the program code. To produce these execution traces, the developer must instrument the program to monitor events and states of interest; these are used to form the possible propositions. Perracotta then creates properties by instantiating eight variations of the temporal logic response pattern with the propositions. Weimer and Necula proposed a static inference approach [30], which analyzes program text and generates properties. These properties specify potentially erroneous behavior of the error-handling portions of the source code. Lastly, Chang *et al.* [28] proposed a dynamic inference approach that generates properties from program event traces. The program traces are created during the execution of the program and track developer-specified events. Chang’s approach involves refining the inference templates built using the Propel patterns [32] to eliminate properties that are not satisfied by the program’s event traces.

These approaches differ from our toolchain-based approach in two key ways. First, they focus on automatically generating properties that describe the behavior of the code, rather than the model. As such, the cost of correcting errors in the later development phase is likely to be more expensive. Second, in general,

these approaches rely on the developer to select portions of the code to explore for properties. This limits the ability of the approaches to discover properties that represent unwanted latent behavior in blind spots. These notable differences mean that our approach can be used in a complementary fashion. Specifically, as part of the model-driven development process our toolchain can be used to automatically discover properties that may represent unwanted latent behavior within the UML model. Once the UML model has been translated to code, the other approaches could be used to ensure that no errors have been introduced.

6 Applying the Toolchain

This section describes our experience of applying the proposed toolchain to the **Body Subsystem** model that was presented in Section 3. We present the types of errors that were discovered, the mitigation strategy that was used for each error, and the consequences of correcting the error. For clarity, we present the errors according to the stages of the toolchain. That is, we begin with a discussion of syntax and consistency errors that CYCLOPS detected. Next, we discuss the errors in types and semantics that CYCLOPS also detected. Finally, we describe how the model was translated into the Promela language and discuss the latent properties that MARPLE discovered.

6.1 Preliminaries

Models of software-based systems are typically constructed using sophisticated modeling tools. Such tools enable a software engineer to design models in a visual manner using a language such as the Unified Modeling Language (UML). The **Body Subsystem** model was developed using IBM Rational Rhapsody [6]. Rhapsody uses UML as its visual design language and provides full support for both behavioral (e.g., sequence or state) diagrams and structural (e.g., class) diagrams. The model comprises class diagrams, sequence diagrams, and state diagrams, thus providing a rich domain vocabulary (i.e., class, operation, and attribute names) as well as a complete set of states and transitions that represent the behavior of the system-to-be. The **Body Subsystem** model contains 52 classes, 37 state diagrams, 255 states (including composite states), and 400 state transitions. There are fewer state diagrams than classes because several of the classes are abstract superclasses or static classes that serve as structures. The model generated approximately 38,000 lines of C++ code.

6.2 Phase I: Syntax and Consistency Check (CYCLOPS)

We begin by applying CYCLOPS to the model, which comprises class and state diagrams. CYCLOPS performs a battery of checks on the input model before it is passed to HYDRA to be translated into Promela. It examines each class, attribute, and operation reference and verifies that the referenced element exists. CYCLOPS also checks for unmatched or missing parentheses, missing semicolons between action statements, and ensures that attributes and operations do not have the same name as their owning class. It also ensures that each state transition expression is well-formed. CYCLOPS identified a wide range of errors in our model, including references to undeclared variables and typographical mistakes. The categories of errors that were discovered, and their frequency of occurrence, are shown in Table 1.

| Error Description | Frequency |
|--|------------------|
| Mangled transition expression (incorrect ordering) | 12 |
| “==” operator used in action list | 9 |
| “=” operator used in guard expression | 4 |
| Missing/unmatched parentheses | 16 |
| Missing semicolons | 19 |
| Attribute has same name as an existing class | 2 |

Table 1: Phase I Errors

6.2.1 Error Mitigation.

Defects that are discovered during Phase I are typically inconsistencies that result from typographical errors. Automated tools cannot make reliable suggestions for resolving most defects of this type, and therefore CYCLOPS must rely on software engineers who are familiar with the model to correct the problem. Once each defect has been corrected, the revised model is given again as input to CYCLOPS, and the Phase I analysis is reapplied. It takes less than one second to parse and check the **Body Subsystem**, thus providing an interactive experience. This incremental defect resolution process proceeds until no further syntax errors are found in the model.

6.3 Phase II: Semantics and Type Check (CYCLOPS)

Next, we used CYCLOPS to check the semantics of each state transition in the model’s state diagrams. CYCLOPS ensures, for example, that each reference to an attribute, operation, or class is valid with respect to the model being analyzed. Furthermore, CYCLOPS verifies that boolean comparisons and assignments are between compatible data types. The errors that were detected in this phase are shown in Table 2.

| Error Description | Frequency |
|--|------------------|
| Misspelling of class/attribute/operation names | 18 |
| Reference to undefined variable | 32 |
| Reference to undefined operation | 4 |
| Reference to undefined enumerated type | 11 |
| Comparison or assignment of incompatible type | 7 |
| Using an attribute as a message or event | 2 |

Table 2: Phase II Errors

6.3.1 Error Mitigation.

Defects that are discovered during Phase II are more subtle, and therefore more difficult to detect, than those discovered during Phase I. The principal focus of Phase II is on parsing and verifying the contents of state transition expressions. A state transition expression specifies the conditions under which the modeled system will move from the current state to the next state and what actions (e.g., variable assignments or calls to operations) will be taken as a result of the transition. Each expression comprises an optional triggering

event, a set of expressions that form a *guard*, a set of actions to perform, and finally a set of messages (events) to raise. A sample state transition expression is as follows:

```
triggerEvent[counter>counter_thresh]done=1;counter=0~nextStep()
```

When the event `triggerEvent` fires, the transition is activated. If the conditions in the guard statement are all satisfied (i.e., if the value of `counter` is greater than the value of `counter_thresh`, then the list of assignments is executed, and finally the event `nextStep` is fired. In a complete model, there would be another transition defined that would be triggered by the `nextStep` event, and the flow of execution would continue.

Errors in state transition expressions can be difficult to detect by visual inspection. For example, it is easy to overlook a missing semicolon between statements in an action list or an assignment (equality) operator that should be a equality (assignment) operator. In the `Body Subsystem` model, many statements in action lists were separated by a line-break rather than a semicolon. There were multiple instances of equality operators being used in assignment statements and assignment operators being used in guards. Neither of these incorrect usages has any effect on code execution. That is, an assignment statement is trivially true, and thus does not affect the satisfaction of a guard. Likewise, a boolean comparison has no side effects and therefore does not belong in an action list. However, both of these constructs still produce valid, executable code in many programming languages that are used for embedded systems (e.g., C). There is therefore a disconnect between the intent of the code and its actual behavior when the system is executed, thus making this class of subtle defects potentially very serious.

6.4 Phase III: Model Translation

In the third phase, the model is free of syntactic and semantic errors and is ready to be translated by HYDRA into the formal language Promela. HYDRA begins by translating the model into the Hydra Intermediate Language (HIL). This intermediate step enables us to build new front-end translators for successive versions of XMI, whose formats evolve over time, without needing to modify the core translation code in HYDRA. The HIL code is then translated into Promela. By constructing an equivalent model in Promela, we are able to conduct formal analysis of the model and to verify model properties specified in Linear Temporal Logic. Each state diagram in the model is treated as a distinct Promela process, thus facilitating the interleaved execution that often reveals unexpected interactions among system components. The translation phase completes within two seconds for the `Body Subsystem` model.

6.5 Phase IV: Discovery of Latent System Properties (Marple)

In the fourth and final phase, the Promela model that was produced by HYDRA is given as input to the third tool, MARPLE. MARPLE uses novelty search to explore the space of properties for a given model and to identify latent properties that are not system requirements. If a latent property is desirable, then it is added to the system requirements. A latent property that is undesirable must be addressed by the system's developer. Potential problems created by unwanted latent properties include incorrect functional behavior, feature interactions, distributed behavior problems, and behavioral inconsistencies. This phase takes on

the order of six hours to complete on a 1.8 GHz PC with 16 gigabytes of memory running Linux. In our experiments, MARPLE was configured to return 25 properties.

Next, we present a sample set of latent properties that were discovered in the `Body Subsystem`. We provide a natural language representation of each property along with a brief discussion of the property, its consequences, and the mitigation strategy that was used.

Property 1:

Globally, `WiperModes.WiperMaster != RSMODE` eventually holds.

Property 1 states that the `WiperMaster` attribute in the `WiperModes` class must eventually have a value that is not `RSMODE`. Through correspondence with the model developer, we learned that this property should not hold, and the developer determined that one of the state transitions in the `WiperModes` state diagram was missing a guard, and therefore the transition was always available to be executed. Once the missing guard was added, we verified that the property no longer held.

Property 2:

Globally, it is always the case that
if `DrvrDrSwitch.Switch == 1` holds, then
`VoltageRangeMonitor.VBattRaw != 18` previously held.

Property 2 states that if the `Switch` attribute in the `DrvrDrSwitch` (Driver Door Switch) class has a value of 1 then the value `VBattRaw` attribute of the `VoltageRangeMonitor` class must not have been 18 in the previous state. Once the property was identified, the model developer was able to identify a missing assignment statement for the `INITIAL` state in the `VoltageRangeMonitor` state diagram. After the missing assignment statement was added, the property no longer held.

Property 3:

Globally, it is always the case that
if `WiperModes.Command == 5` holds, then
`AmbientLightSensorInput.lightLevel != 4`
previously held.

Property 3 states that if the value of the `Command` attribute in the `WiperModes` class is `HALT`, then the value of the `lightLevel` attribute in the `AmbientLightSensorInput` class must not have been `TWILIGHT` in the previous state. Despite the different set of classes and attributes in this property as opposed to property 2, the model developer discovered that Property 2 and Property 3 held because of the same missing assignment statement in the `VoltageRangeMonitor` state diagram. After the statement was added to remedy Property 2, Property 3 no longer held.

Property 4:

Globally, it is always the case that
`WiperModes.Command != HALT`.

Property 4 states that the value of the `Command` attribute in the `WiperModes` class will never be `HALT`. From this property, the model developer determined that a triggering event in the `RelayControl` class (part of the Windshield Wipers subsystem) never occurs, and thus the state machine remains in the `WAIT` state indefinitely. Figure 4 shows partial state diagrams from the `RelayControl` and `WiperModes` classes. There was a missing call to the event `RlyCtlActive` (shown in **bold**) in the transition expression for the initial state in `RelayControl` (Figure 4(a)). Since the transition expression for `WiperModes` (Figure 4(b)) is waiting

for the event to be fired, it will wait indefinitely. After adding a call to the missing event in the appropriate state transition in `RelayControl`, the property no longer held.




Figure 4: Partial State Diagrams for Classes Affected by Property 4

7 Discussion

In this section, we present a discussion of the results of applying the proposed toolchain and consider the consequences of its use in an industrial development setting. As in previous sections, we present the discussion in terms of each phase of the toolchain.

7.1 Syntax and Semantics Defects

We had access to two major revisions of the `Body Subsystem` model for this work: an early revision that had not been used to generate source code and thus contained syntax errors and type inconsistencies, and a subsequent revision that had undergone source code generation and compilation. In order to assess `CYCLOPS`'s ability to detect syntax and semantics errors, we applied it to the earlier model revision. `CYCLOPS` detected all of the errors that the compilers had detected during source code generation and compilation, and it also identified additional errors that were subtle and would be difficult to locate by manual inspection. For example, an assignment statement that was mistyped as a boolean comparison would not be detected by a compiler, but such a mistake may have an adverse effect on system behavior. The developer of the `Body Subsystem` model stated that without the use of a tool such as `CYCLOPS`, these subtle errors would have been allowed to propagate into generated source code and, perhaps, into the design and implementation of the system. Since system models are typically small during the late requirements stage of the software lifecycle, such defects are straightforward to resolve once they have been identified. Identifying and resolving these subtle defects in the requirements stage reduces the amount of time spent debugging and reengineering the system at later stages of development.

7.2 Latent Property Detection

Once the syntactic and semantic errors in the `Body Subsystem` model had been resolved, we discovered several interesting properties that indicated deeper, and more subtle, problems with the model. First, a subset of the numerical constants in the model had not been initialized, and thus any behavior that depended on the

constant values was affected. The set of variables and constants in the model is given to MARPLE as input, and each variable has an associated range of feasible values. In initial experiments, we used a small range (e.g., (0, 10)) for each variable, and MARPLE explored the space of properties using that range. However, such a small and fixed range was neither sufficient nor accurate for the model. In a subsequent iteration, the model developer provided a range of feasible values for each variable, and subsequent experiments yielded more sensible properties, including those that were presented in Section 6.

While the proposed toolchain detects several types of model errors, the developer of the **Body Subsystem** told us that the toolchain is most useful for identifying portions of the model or system requirements that are missing. The toolchain identified a set of missing constant initializations, transition guards, and transition action statements. The discovered properties did not always point directly to the missing model components (e.g., properties 2 and 3 in Section 6), but they yielded enough information for a developer with knowledge of the system and model to make inferences about the possible causes of the defect, to identify the cause, and to revise the model accordingly. In the absence of the proposed toolchain, such defects would most likely be discovered during integration testing after the source code has been completed, thus making the cost to repair the defects significantly greater than if they had been discovered in the requirements stage.

MARPLE uses an evolutionary search technique to explore the space of properties for a given model. Due to inherent randomness in the search process, it is unlikely that MARPLE will revisit the same property in independent executions. However, it is straightforward to make note of any interesting properties and to re-examine them at a later time to monitor for regressions, even if the model has been revised since the properties were discovered. We leveraged this capability to track flaws in the model and to verify that unwanted properties no longer held once a model refinement had been deployed. This ability to track defects over time facilitated a step-wise, iterative model refinement process that enabled us to work with the model developer, who works with us remotely, to incrementally resolve the problems that our toolchain identified. The phases of the proposed toolchain provide a natural and effective workflow for identifying defects, deciding on a resolution, and analyzing the revised model for subsequent defects.

8 Conclusions

In this paper, we presented an experience report describing the use of a toolchain for detecting syntactic and semantics errors in behavioral system models, as well as detecting latent system properties during the early requirements phase of the software lifecycle. We demonstrated that the proposed toolchain is an effective means for identifying syntax errors, resolving ambiguous references, and discovering unwanted latent system properties.

We are considering several avenues for future work. First, we plan to integrate metamodel-level consistency checking into the CYCLOPS tool, thus enabling flexible and robust error detection that is grounded in a formal semantics for UML state diagrams. Next, we are investigating patterns within the discovered latent properties and to leverage their key features to fine-tune parameters for the MARPLE tool. Finally, we are exploring several strategies for reconfiguring the toolchain to detect situations in which two system features interact and lead to system failures or other unexpected behavior.

References

- [1] R.R. Lutz. Analyzing software requirements errors in safety-critical, embedded systems. In *Requirements Engineering, 1993., Proceedings of IEEE International Symposium on*, pages 126–133. IEEE, 1993.
- [2] R. Pressman. *Software Engineering: A Practitioner’s Approach*. 2007.
- [3] Tigris.org. ArgoUML Modeling Tool. Available at <http://argouml.tigris.org/>, April 2011.
- [4] International Business Machines Corporation. Rational Software Architect. Available at <http://www.ibm.com/developerworks/rational/products/rsa/>, April 2011.
- [5] Microsoft Corporation. Available at <http://office.microsoft.com/en-us/visio/>, April 2011.
- [6] International Business Machines Corporation. Rational Rhapsody. Available at <http://www-01.ibm.com/software/awdtools/rhapsody/>, April 2011.
- [7] G.J. Holzmann. The model checker SPIN. *Software Engineering, IEEE Transactions on*, 23(5):279–295, 2002.
- [8] G. Booch, J. Rumbaugh, and I. Jacobson. *Unified Modeling Language User Guide, The (Addison-Wesley Object Technology Series)*. Addison-Wesley Professional, 2005.
- [9] B. Long, J. Dingel, and TC Graham. Experience applying the SPIN model checker to an industrial telecommunications system. In *Proceedings of the 30th international conference on Software engineering*, pages 693–702. ACM, 2008.
- [10] T. Bäck, D.B. Fogel, and Z. Michalewicz. *Handbook of evolutionary computation*. Taylor & Francis, 1997.
- [11] J. Lehman and K.O. Stanley. Exploiting open-endedness to solve problems through the search for novelty. *Artificial Life*, 11:329, 2008.
- [12] W.E. McUmber and B.H.C. Cheng. A general framework for formalizing UML with formal languages. In *Proceedings of the 23rd International Conference on Software Engineering*, pages 433–442. IEEE Computer Society, 2001.
- [13] G.J. Holzmann, American Telephone, and Telegraph Company. *Design and validation of computer protocols*, volume 94. Prentice Hall New Jersey, 1991.
- [14] A. Pnueli. The temporal logic of programs. In *18th Annual Symposium on Foundations of Computer Science*, pages 46–57. IEEE, 1977.
- [15] H. Goldsby and B. Cheng. Automatically Discovering Properties That Specify the Latent Behavior of UML Models. *Model Driven Engineering Languages and Systems*, pages 316–330, 2010.
- [16] Matthew B. Dwyer, George S. Avrunin, and James C. Corbett. Patterns in property specifications for finite-state verification. In *Proceedings of the 21st International Conference on Software Engineering*, pages 411–420. IEEE Computer Society Press, 1999.

- [17] Sascha Konrad and Betty H. C. Cheng. Real-time specification patterns. In *Proceedings of the International Conference on Software Engineering (ICSE05)*, St Louis, MO, USA, May 2005.
- [18] The MathWorks, Inc. MATLAB and Simulink. Available at <http://www.mathworks.com/>, April 2011.
- [19] A. Egyed. Fixing inconsistencies in UML design models. In *Proceedings of the 29th international conference on Software Engineering*, pages 292–301. IEEE Computer Society, 2007.
- [20] A. Egyed. Scalable consistency checking between diagrams-The VIEWINTEGRA Approach. In *ase*, page 387. Published by the IEEE Computer Society, 2001.
- [21] G. Engels, J.M. Küster, R. Heckel, and L. Groenewegen. A methodology for specifying and analyzing consistency of object-oriented behavioral models. 26(5):186–195, 2001.
- [22] H. Gomaa and D. Wijesekera. Consistency in multiple-view UML models: a case study. In *Workshop on "Consistency Problems in UML-based Software Development II"*, page 1. Citeseer, 2003.
- [23] C. Schwarzl and B. Peischl. Static-and dynamic consistency analysis of UML state chart models. *Model Driven Engineering Languages and Systems*, pages 151–165, 2010.
- [24] J. Simmonds, R. Van Der Straeten, V. Jonckers, and T. Mens. Maintaining consistency between uml models using description logic. *Série Lobjet-logiciel, base de données, réseaux*, 10(2-3):231–244, 2004.
- [25] R. Wagner, H. Giese, and U. Nickel. A plug-in for flexible and incremental consistency management. In *Proc. of the International Conference on the Unified Modeling Language 2003 (Workshop 7: Consistency Problems in UML-based Software Development)*, San Francisco, USA, 2003.
- [26] F. Baader. *The description logic handbook: theory, implementation, and applications*. Cambridge Univ Pr, 2003.
- [27] William Chan. Temporal-logic queries. In *CAV '00: Proceedings of the 12th International Conference on Computer Aided Verification*, pages 450–463, London, UK, 2000. Springer-Verlag.
- [28] Richard M. Chang, George S. Avrunin, and Lori A. Clarke. Property inference from program executions. Technical Report UM-CS-2006-26, University of Massachusetts, 2006.
- [29] Arie Gurfinkel, Marsha Chechik, and Benet Devereux. Temporal logic query checking: A tool for model exploration. *IEEE Transactions on Software Engineering*, 29(10):898–914, 2003.
- [30] Westley Weimer and George C. Necula. Mining temporal specifications for error detection. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pages 461–476, 2005.
- [31] Jinlin Yang, David Evans, Deepali Bhardwaj, Thirumalesh Bhat, and Manuvir Das. Perracotta: mining temporal API rules from imperfect traces. In *ICSE '06: Proceedings of the 28th international conference on Software engineering*, pages 282–291, New York, NY, USA, 2006. ACM.
- [32] Rachel L. Smith, George S. Avrunin, Lori A. Clarke, and Leon J. Osterweil. Propel: an approach supporting property elucidation. In *ICSE '02: Proceedings of the 24th International Conference on Software Engineering*, pages 11–21, New York, NY, USA, 2002. ACM.