Decomposing Verification Around End-User Features*

Kathi Fisler[1] and Shriram Krishnamurthi[2]

[1] Department of Computer Science, WPI, Worcester, MA, USA
kfisler@cs.wpi.edu
[2] Computer Science Department, Brown University, Providence, RI, USA
sk@cs.brown.edu

Abstract. Practical program verification techniques must align with the software development methodologies that produce the programs. Numerous researchers have independently proposed models of program development in which modules encapsulate units of end-user functionality known as *features*. Such encapsulation reflects user concerns into a program's modular structure, which in turn promises to simplify program maintenance in the face of requirements evolution. The interplay between feature-oriented modules and verification raises some interesting challenges and opportunities. Such modules ameliorate some difficulties with conventional modular verification, such as property decomposition, while creating others, by contradicting assumptions that underlie most modular program verification techniques. This paper motivates the decomposition of systems by features and provides an overview of the promises and challenges it poses to verification.

1 A Notion of Software Development

For program verification to thrive, verification methodologies must align with software development methodologies. This goal imposes several requirements. First, verification tools should be able to handle program fragments of the style and granularity that programmers produce. Second, the effort to verify a program increment should bear some reasonable ratio to the effort to develop that increment. Third, the effort needed to reverify a program or fragment as it evolves should be proportional to the effort required to make the modification. Today's verification techniques fail to meet these goals, partly due to a misalignment between the models of software development and programming on which the techniques are built.

Our understanding of this problem is inspired by the picture in Figure 1 which Michael Jackson used in his presentation at ESEC/FSE 2001 (following his acceptance of the SIGSOFT Outstanding Research Award).[1] The box at

* This work is partially funded by NSF grants CCR-0305834, CCR-0132659, CCR-0447509 and CCR-0305950.

[1] We have transcribed this picture from our notes; a related version is in a paper [1].

B. Meyer and J. Woodcock (Eds.): Verified Software, LNCS 4171, pp. 74–81, 2008.

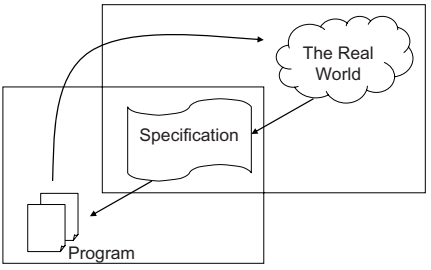

Fig. 1. The requirements-program feedback loop

the lower-left might be grossly characterized as the province of programming languages, proceeding from specifications to programs that, we hope, properly implement those specifications. Jackson calls this box the *solution space*. The box at the upper-right is the domain of requirements engineering: the collection of processes, many sociological (and imprecise!), that glean requirements for a system from its users and other stakeholders or, more broadly, from the fuzzy blob that is the "real world". In Jackson's terminology, the transactions in this world must remain in the *problem space*.

As soon as the program comes into existence, however, it itself becomes a part of the world. This invariably triggers a possibly new set of requirements. (As most requirements engineers and user interface designers will attest, a common user reaction is, "Oh, it does *that*?" followed by, "That is not what I meant at all. That is not it, at all" [2].) This cyclic dependency and directed flow of requirements is a source of many contemporary software development problems. This suggests that our techniques for software development should account for this by finding ways to be more pliable in the face of requirements and specification changes.

2 Program Development Styles and Verification

These ideas have significant implications for program verification. Research in property-based verification (which this paper considers distinct from correct-by-construction approaches) has often assumed a simplistic model of program development, in which the verifier has a complete program to analyze against established specifications. The body of work on modular verification relaxed that assumption to handle portions of programs that correspond to units of separate compilation [3]. This position paper argues that emerging forms of software development embody different assumptions from most current verification methodologies. It is therefore essential for verification to support these development techniques; better still would be if verification could *exploit* it.

Some researchers [4] believe verification should be organized around *software components*. Unfortunately, the term "component" seems to mean too many things (and often too little) in the literature. In particular, a component may or may not have a direct relationship to the specifications that inspired the program

in the first place. (Sometimes it might; in other cases, it may be a generic unit of reuse, such as a sorting routine or database interface.) We therefore believe the emphasis must shift from using terms like "modules" and "components" to discussing *what the modules encapsulate*.

3 Programs as Collections of Features

An end-user of a system typically does not care which database interface or sorting routine is used in the implementation, or even whether they are used at all; rather, users describe how they would like to see the system behave (via methods such as use cases) as a collection of units of functionality that we call *features*. When requirements change, they often either add or remove features, or change a previously identified feature. Managing changes to features is therefore a key problem in software development. If each feature is implemented by a specific module, it becomes easier to identify where changes must be made. Furthermore, if these modules meet the criteria of components laid out by Szyperski [5], then creating the system is just a matter of component composition, while adding and removing features is simply a matter of writing a new composition. In other words, this style of program organization would respect the feedback loop in the Jackson figure, observing that if the program's shape mirrors the requirements and specifications, software evolution will become easier to manage.

Writing these identifiable increments in conventional programming languages is challenging because an increment may affect parts of a program across traditional module boundaries: such increments are said to be *cross-cutting*. This observation has led to a growing body of work on developing new forms of program modularity [6,7,8,9,10,11,12,13,14,15,16] that support modularization around features and composition to create a variety of individual systems (thus forming a software *product line* [17]). Some techniques are purely static, effectively manipulating the program's source, while others have dynamic elements, offering the ability to reflect on the state of the program's execution and then to modify it. There is now a growing awareness of this style of programming (especially as popularized by "aspect-oriented programming"), and several case-studies highlight its feasibility and observe its benefits.

4 Research Program on Verifying Feature Modules

This model of program organization offers a substantial benefit to verification as well. One of the main challenges in modular verification is the decomposition of properties to align with the program module's boundaries. In theory, feature-orientation should largely eliminate this hurdle. As the user's perspective frames both the features (modules) and the properties, the scope of each property largely matches the scope of some module (with the exception of global system properties). This correspondence should also make verification more scalable in the face of specification evolution, by localizing re-verification to the relevant parts of a

program. In return for these benefits, feature-orientation demands new theories of modular reasoning to support feature-based decompositions.

We have been working on techniques for modular model checking of feature-based designs since 2001 [18,19,20]. At a high level, this work shares the goals of other modular verification research, namely, to verify code fragments independently and derive some properties of the composed program from the properties of the fragments. The nature of feature-based design, however, adds some nuances to this problem.

1. Since features are added to programs to provide some user-defined functionality, we may need to prove that adding a feature (a) preserves established properties of programs, or (b) establishes a *new* property of the composed program. In some models, a feature might (c) establish a property that the original program should have satisfied, but did not.

   Performing these analyses modularly is important because we could potentially add many features to an existing program, at which point the cost of analyzing each possible product (a subset of features) becomes prohibitive. Feature-based design shifts the motivation for modular reasoning from making verification tractable in terms of computational resources to making it practical across the combinatorial number of products that can be built from designer-specified program modules.
2. Features can interact, causing properties of individual components to be violated in the composed program. In a telecommunications system, for example, a voice mail feature might be required to pick up an unanswered call after 4 rings and a call forwarding feature might be required to pass an unanswered call along after 4 rings. A system with both voice mail and call forwarding will respect only one of these requirements. Other examples are far more subtle. This *feature-interaction* problem is pervasive and not always amenable to formal analysis [21]. As model checking each combination of features for interactions is infeasible, modular analysis must support detecting those interactions that can be captured formally.

To date, our work has focused on property preservation (nuance 1a) with preliminary attention to feature interaction (nuance 2). We are able to model check CTL properties against individual features and perform lighter-weight checks to confirm that a feature's properties will be preserved when composed into an existing program. We are also able to modularly detect feature interactions that manifest as violations of properties expressed in CTL. Our work has identified several ways in which feature-based designs challenge the conventional assumptions of modular model checking:

1. Most modular verification work assumes parallel composition, while feature composition is largely (though not entirely) sequential. This in turn has interesting consequences. For example, module composition can create new paths through programs (parallel composition deletes but does not add paths). Although parallel composition can simulate sequential composition, it is not clear that doing so best exploits the advantages of sequential composition.

2. Modules are not closed because data from one module may persist into another. For instance, in an email system, one feature may encrypt a message and that attribute should persist into subsequent features, even though the models of those features do not mention encryption.
3. One module may refine the interpretations of propositions in another. E.g.: An email system may classify a message as anonymous if it has passed through an anonymous remailer, but adding a digital signing feature forces the interpretation of anonymity to also require the message to be unsigned. Reinterpreted or persistent propositions often preclude lifting properties proven of an individual module to the composed program.

Even more fundamental, however, is our growing understanding that *verification may be the wrong problem to solve*. Traditional verification methods (especially model checking) are primarily designed to authoritatively determine the truth or falsity of properties over models. However, most of the property violations we observe arise only upon composition, because some compositions satisfy the properties while others fail them. Most verification runs over individual features are invariably inconclusive because there isn't enough information in the module to entirely satisfy or fail the property. This pushes the verification decision onto the composition step, which needs appropriate information that it can then use to perform a lightweight check that (a) is not too expensive, and (b) does not involve re-examining the innards of the individual modules. The appropriate analysis on individual modules is therefore some form of constraint generation, rather than outright verification. In our work on this approach [18] the constraints have been propositional and temporal, reflecting their foundation in model checking.

The need for constraint generation in practice does not contradict our earlier claim that properties align naturally with features. The truth of a property in a feature often depends on two specific pieces of information from the feature's environment: (1) whether control paths exiting the feature reach particular states in the original program, and (2) the values of persistent and reinterpreted data propositions determined in earlier features. Our proposed constraint generation is akin to generating environmental constraints; this step is feasible in our work because we generate environments specific to the program properties that a feature must preserve. The checks required to discharge these constraints at feature-composition time tend to be lightweight (simple propositional or reachability checks) because the feature that aligns with a property discharges the bulk of that property's obligations.

5 Some Challenges

Our observations leave open a large collection of interesting research problems. Some questions that need to be addressed include:

– Theoretical foundations for richer notions of composition. Most verification research is built on purely sequential composition or variations (synchronous, interleaved, etc) on parallel composition. The composition models in many

feature-oriented programs lie between these extremes. We have studied programs that employ what we call *quasi-sequential composition* [19]: at the highest level, modules compose sequentially, but each module is formed of components that compose in parallel. This form of composition is interesting because it leads to states in the global state space that span different modules, but in controlled and predictable ways. Quasi-sequentiality is reminiscent of the pattern of execution in Valiant's Bulk Synchronous Parallel model [22], but that model has no notion of program modularity and does not consider verification.
- Techniques for generating temporal constraints rich enough to support modular feature verification.
- Techniques for determining whether a code fragment introduces a desirable though previously untrue property over a program.
- Theories of compositional reasoning that are tuned for predicting feature interaction errors, instead of checking success or failure of known properties.

6 Perspective

We find it telling that several researchers, working independently and in entirely distinct areas (often without any knowledge of each other), have within a few years proposed extremely similar models of software development centered around features. We believe that Jackson's picture explains why this model is not accidental, but rather fundamental to the way programs originate and evolve. Without lapsing into thoughts of silver bullets, we should take this model seriously, especially as formal programming language research is beginning to catch up with these less formal approaches (several papers [23,24,25,26] offer a representative sampling). Any attempt to lay the foundations for a practical program verifier must look ahead to how programs will be developed in the future, not only at programs written using antediluvian methods in legacy languages.

Acknowledgements. Many colleagues and students have shaped our thoughts on this problem. Matthew Flatt, Matthias Felleisen, and Robby Findler helped us understand program modularity. Don Batory taught us about scaling modularity and gave us invaluable help with examples. A series of Brown undergraduates—Harry Li, Colin Blundell, and Michael Greenberg—collaborated with us on much of the verification work, greatly enhancing both our knowledge and enjoyment.

References

1. Jackson, M.: Why software writing is difficult and will remain so. Information Processing Letters 88, 13–25 (2003)
2. Eliot, T.S.: The love song of J. Alfred Prufrock. In: Prufrock and Other Observations, The Egoist, Ltd. London (1917)
3. Abadi, M., Lamport, L.: Conjoining specifications. ACM Transactions on Programming Languages and Systems 17, 507–534 (1995)

4. Xie, F., Browne, J.C.: Verified systems by composition from verified components. In: Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, New York, NY, USA, pp. 277–286. ACM Press, New York (2003)
5. Szyperski, C.: Component Software: Beyond Object-Oriented Programming. Addison-Wesley, Reading (1998)
6. Aßmann, U.: Invasive Software Composition. Springer, Heidelberg (2003)
7. Batory, D.: Feature-oriented programming and the AHEAD tool suite. In: International Conference on Software Engineering (2004)
8. Batory, D., O'Malley, S.: The design and implementation of hierarchical software systems with reusable components. ACM Transactions on Software Engineering and Methodology 1, 355–398 (1992)
9. Findler, R.B., Flatt, M.: Modular object-oriented programming with units and mixins. In: ACM SIGPLAN International Conference on Functional Programming, pp. 94–104 (1998)
10. Harrison, W., Ossher, H.: Subject-oriented programming: a critique of pure objects. In: ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages & Applications, pp. 411–428 (1993)
11. Jackson, M., Zave, P.: Distributed feature composition: A virtual architecture for telecommunications services. IEEE Transactions on Software Engineering 24, 831–847 (1998)
12. Kiczales, G., Lamping, J., Mendhekar, A., Maeda, C., Lopes, C.V., Loingtier, J.M., Irwin, J.: Aspect-oriented programming. In: European Conference on Object-Oriented Programming (1997)
13. Lieberherr, K.J.: Adaptive Object-Oriented Programming. PWS Publishing, Boston (1996)
14. Mezini, M., Lieberherr, K.: Adaptive plug-and-play components for evolutionary software development. In: ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages & Applications, pp. 97–116 (1998)
15. Smaragdakis, Y., Batory, D.: Implementing layered designs and mixin layers. In: European Conference on Object-Oriented Programming, pp. 550–570 (1998)
16. van Ommering, R.: Building Product Populations with Software Components. PhD thesis, Rijksuniversitat Groningen (2004)
17. Clements, P., Northrop, L.: Software Product Lines: Practices and Patterns. Addison-Wesley, Reading (2002)
18. Blundell, C., Fisler, K., Krishnamurthi, S., Hentenryck, P.V.: Parameterized interfaces for open system verification of product lines. In: IEEE International Conference on Automated Software Engineering (2004)
19. Fisler, K., Krishnamurthi, S.: Modular verification of collaboration-based software designs. In: Symposium on the Foundations of Software Engineering, pp. 152–163. ACM Press, New York (2001)
20. Li, H.C., Krishnamurthi, S., Fisler, K.: Modular verification of open features through three-valued model checking. Automated Software Engineering 12, 349–382 (2005)
21. Keck, D.O., Kuehn, P.J.: The feature and service interaction problem in telecommunications systems: A survey. IEEE Transactions on Software Engineering 24, 779–796 (1998)
22. Valiant, L.G.: A bridging model for parallel computation. Communications of the ACM 33, 103–111 (1990)
23. Ancona, D., Lagorio, G., Zucca, E.: Jam—designing a Java extension with mixins. ACM Transactions on Programming Languages and Systems 25, 641–712 (2003)

24. Flatt, M., Krishnamurthi, S., Felleisen, M.: Classes and mixins. In: ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, pp. 171–183 (1998)
25. Odersky, M., Altherr, P., Cremet, V., Emir, B., Maneth, S., Micheloud, S., Mihaylov, N., Schinz, M., Stenman, E., Zenger, M.: An overview of the Scala programming language. Technical Report IC/2004/64, EPFL Lausanne, Switzerland (2004)
26. Schärli, N., Ducasse, S., Nierstrasz, O., Black, A.: Traits: Composable units of behavior. In: European Conference on Object-Oriented Programming, pp. 248–274 (2003)

A Discussion on Kathi Fisler's Presentation

Greg Nelson

Kathi, I have a question about your "New challenges" slide. It sounds at first very daunting. But I wonder if you would agree that the fundamental problem with that multilingual issue is... (sentence left incomplete). I mean, a simple answer would be that for each language in which you program, you develop a verifier, and the problem with that simple answer is, that the interfaces between two components may not be in the language of either component. But I only see a small, bounded number of possible interfaces: There is the procedure call interface as when a Modula program calls a C program, there is the byte stream interface as when you pipe the output of grep into sort or when you open a TCP-connection to a file server, and there are the remote procedure call kinds of interfaces as in an HTTP call or a Java RMI call, where it is like a procedure call interface, but it crosses machine boundaries. And if you can handle those kinds of interfaces, then it should not matter that the components are in different languages, I hope.

Kathi Fisler

Well, it is not that you have a large number of interfaces. The problem is, you have an awful lot of interaction between the components at those interfaces, okay? So, it is fairly small to say: "I make a call to my access control policy and say: 'Is this action permitted?'" But these kinds of interactions between language do not happen in a small number of places. So that is really where this problem is that is going to be fairly challenging. I hope, this answers the question.