

Toward Engineered, Useful Use Cases

Clay Williams, IBM T.J. Watson Research Center
Matthew Kaplan, IBM T.J. Watson Research Center
Tim Klinger, IBM T.J. Watson Research Center
Amit Paradkar, IBM T.J. Watson Research Center

Abstract


We explore common problems that exist in the practice of use case modeling: lack of consistency in defining use cases, misalignment between the UML metamodel and the textual representations of use cases expounded in the literature, and the lack of a semantics that allows use cases to be executable and analyzable. We propose an engineering approach to the issues that can provide a precise foundation for use case development. We next discuss four potential uses of such a foundation and identify the research problems that must be addressed to support these applications.

1 INTRODUCTION

Use cases were developed as a technique for capturing the required behavior of a software system. [Jacobson92] While they have been successful in having an impact on software development, their usage is not as prevalent as we believe it could be. Several potential factors account for this, including:

1. Confusion regarding the meaning of the term *use case*. Different approaches advocate graphical versus textual methods. Further confusion arises as levels of abstraction vary widely within and across projects using use cases.
2. Misaligned characterizations of use cases in the UML metamodel. In UML, use cases are characterized as BehavioredClassifiers, which does not correctly support the best practices for representing use cases that have been proposed. [Armour01, Cockburn01]
3. Use case semantics are poorly defined, limiting their usefulness to primarily being vehicles for communication. While this is an important use for them, wider industrial use requires that they support machine based processing and analysis.

This paper is a philosophical and conceptual exploration of each of these points in detail in which we identify potential solutions that support an engineering approach to creating use cases. However, engineering use cases is not an end in itself – the technique must



provide value across multiple dimensions to be adopted in an industrial setting. To support our approach, we discuss four potential uses of such a foundation (prototyping, estimation, refinement to design, and test generation). We close by discussing related work and future research.

2 USE CASE CONFUSION

Across projects, use cases exist at multiple levels of abstraction and are captured in many different notions. The leading practitioners [Armour01,Cockburn01] all advocate similar content that includes preconditions, main (success) sequences, alternative/exceptional sequences, and postconditions. We believe that this serves as a good basis for an engineered use case representation, but more detail is needed to define what the sequences contain.

Another issue is that it is not uncommon to see use cases used across widely different levels of abstraction, from high level usage of a system to extremely detailed descriptions of system behavior. This lack of consistency regarding the level of use case abstraction is directly influenced by the content supported within a use case model. We take the position that the most useful use cases are those that Cockburn calls “sea level” use cases [Cockburn01], which are defined at the user-goal level. Our desire is to make constructing good sea level use cases as intuitive and repeatable as possible. A second goal of our representation is to enable the capture of precise, analyzable, and executable use cases with as little modeling and information required as possible.

We start by defining a precise contextual foundation for use cases. Use cases are usually collected during the inception phases of a project. They allow the system requirements to be specified by saying *what* the system will do without saying *how* it does it. This requires the capability to talk about the key elements in the system in a very precise way. However, inception occurs prior to architecture definition and design, so we must carefully define the context against which use cases are developed in order to achieve this precision. This context is the domain model of the system.

A domain model is a model of the significant elements from the system’s application domain. These are the elements that will be created, modified, and used in the application. They are represented as UML classes with attributes. Also present in the domain model are domain rules (or invariants,) which capture constraints that must hold for instances of the domain objects to be valid. Figure 1 shows the domain model for a simple example ATM application.

The domain model plus the four primitive types (Boolean, float, integer, and string) serve as the type system for detailing use cases. Use cases are typically defined as sequences of actions that occur between the system and one or more actors, but the specifics of the action types that are supported are left undefined in UML. In our scheme, use cases support 4 basic actions and 4 flow-of-control actions, each of which has an associated statement for use in detailing the use case.

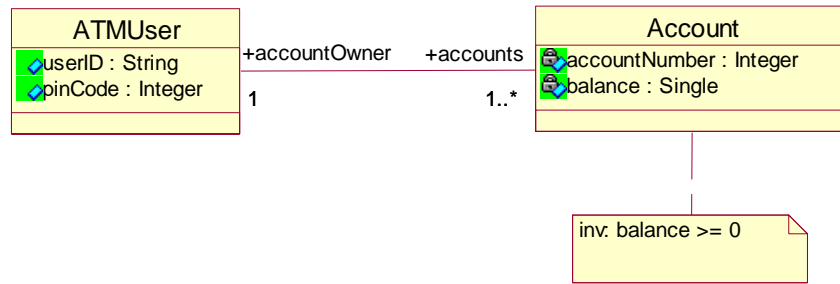


Fig. 1. The Domain Model for an ATM Application.

Basic Actions:

- Input – an actor provides input to the use case
- Output – the use case returns output to an actor
- Computation – a computation is performed using provided input and domain instance information. Computation entails creating, modifying, or deleting instances of domain classes.
- Exception Handling – the system responds to an issue with the input or state of the domain model instances.

Flow of Control Actions:

- Selection – allows the use case to conditionally execute actions
- Iteration – allows the use case to repeatedly execute an action sequence
- Inclusion – allows the use case to include the behavior of another use case
- Extension – defines the extension point for an extending use case

Use cases also have preconditions and postconditions, and both are written in terms of the domain model. Below is an example use case for the Withdraw Money capability for an ATM system.

Use Case: Withdraw Money

Precondition: ATM Customer must be logged onto system.


1. ATM Customer selects Withdraw.
2. System requests amount.
3. ATM Customer enters amount.
4. System dispenses cash.
5. System debits amount from the account balance.

Exceptions:

- 3.1 [amount greater than account balance]
 - 3.1.1 System displays balance exceeded warning
 - 3.1.2 If ATM Customer chooses "continue" rejoin at 3
 - 3.1.3 Else terminate use case

Postconditions:

Successful withdrawal:
 $account.balance = account.balance@pre - amount$
 Unsuccessful withdrawal:
 $account.balance$ is unchanged.



This example illustrates the four primary types of statements that can occur in a use case: input statements, computation statements, output statements, and exceptions. The example also illustrates selection actions and rejoining the main scenario from an exceptional alternative.

This representation (use cases, a high level domain model, and domain invariants) serves as a basis for engineered use cases. The use cases consist of very specific types of actions. They specify, in terms of the domain model, what occurs in an interaction with the system being specified without saying how it is done.

3 METAMODEL ALIGNMENT FOR ENGINEERED USE CASES

UML 2.0 Use Case Shortcomings

Figure 1 shows the current metamodel for UML 2.0 use cases [UML02].

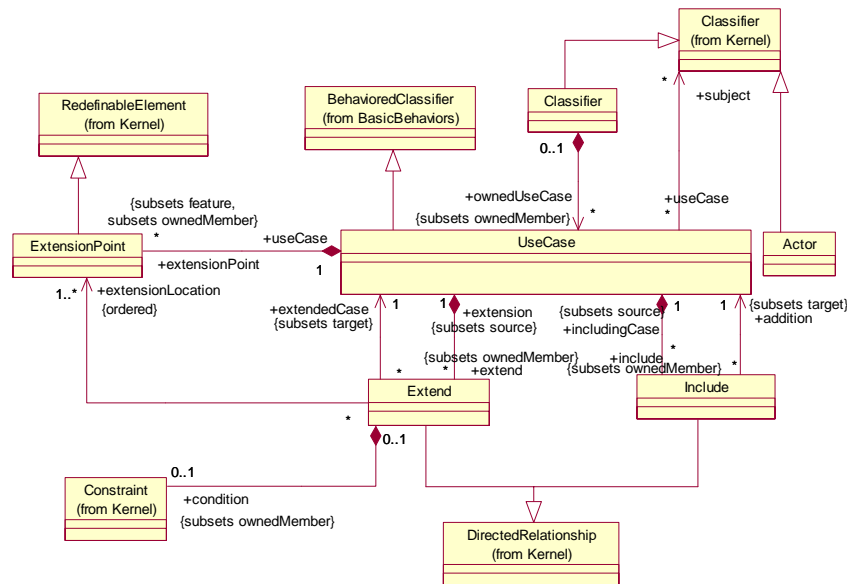



Fig. 2. The UML 2.0 Metamodel

Although the UML 2.0 metamodel improved over UML 1.x by adding the notion of subjects (what we call context in [Williams01]) to use cases, it does contain two main shortcomings. Given the current size of the UML 2.0 specification, the remedies to these issues are complex, which prevents a complete and formal rectification of them in this paper. In the subsection following this one, we sketch the approach that we believe needs to be used in order to remedy them.

In the UML 2.0 metamodel, the UseCase metaclass aligns with the BehavoredClassifier metaclass, which itself is a descendent of Classifier. The OMG standard for UML [UML02] states that "a classifier is a type." This is problematic for use



cases, as the essence of a use case is to describe an interaction between an actor and the system, not to classify types. One consequence of this misalignment is that generalization is poorly defined for use cases. In UML 2.0, generalization is defined as both feature inheritance and subtyping. It is unclear how these apply to use cases when they are understood to be actor-system interaction specifications. Another issue is that BehavioredClassifiers have ownedBehaviors, which misaligns with the typical textual based approaches for use cases.


There are two modifications that partially address the problems with the current use case alignment in the UML 2.0 metamodel. However, neither is sufficient to solve the problems without additional re-engineering of the metamodel. Use cases could subclass Behavior. In the current metamodel, this is problematic, as Behavior always exists in the context of a BehavioredClassifier and use cases should exist in the context of the domain model, which consists of several classifiers. Another approach would be to add constraints to UseCase to prohibit ownedBehaviors and attributes from existing on use cases. Then classifierBehavior could serve as the basis for specifying the use case behavior. However, this approach has the problem that it presumes that the classifierBehavior is updating the state of the use case instance rather than the state of domain class instances. To properly support alignment between the graphical syntax of UML and the textual details in a use case requires a careful re-examination of the UML metamodel.

We argued earlier that use cases should consist of four basic actions (input, output, computation, and exceptions) with supporting actions providing for control flow specification (inclusion, extension, sequencing, selection, and iteration.) Were it possible to align use cases more carefully with Behavior, refining the model to include specifics regarding contents of a use case would offer a degree of precision that is currently not available. Allowing use cases to support a (possibly enhanced) subset of the existing UML 2.0 actions defined for activities would provide specificity that is lacking in the current specification. As noted above, there is currently no good way to do this in the existing metamodel.

Toward a Metamodel Supporting Engineered Use Cases

As we cannot redesign the use case portion of the UML 2.0 metamodel in this paper without possibly causing problems in other portions of the specification, we instead outline an approach that can be used to address the concerns raised in the previous section.

The first area that must be addressed is which UML metaclass the UseCase metaclass specializes. As suggested earlier, we believe the best remedy is to have UseCase be a specialization of Behavior. This is consistent with the vision defined in section 2, where use cases have input parameters, return results (output), and have preconditions and postconditions. All of these concepts are already supported in the current Behavior metaclass. This requires that Behavior be modified or extended to allow association with multiple classifiers in order to support using the domain model as the use case context.



Were we to align UseCase with Behavior, the notion of supported actions is still missing. Like the Activities metaclass (which also aligns with Behavior,) the UseCase metaclass should specify a set of Actions which can be performed within a use case instance. A subset of the current actions defined in UML 2.0 would suffice as a basis for creating executable use cases. These include the *Read Write Actions* (Object Actions, Variable Actions, Link Actions, and Structural Feature Actions). These should be extended with specific *Computational Actions* (such as arithmetic, Boolean, and string related actions, and domain specific actions.) *Invocation Actions* are not required, as neither use cases nor their contexts support operations or signals. Given the alignments proposed in this section, we next explore the semantics required to support engineered use cases.

4 TOWARD PRECISE USE CASE SEMANTICS

In order to support an engineering approach to creating precise use cases, we must address semantic issues in addition to the structural/content issues discussed above. While space considerations prevent a formalized treatment of all relevant semantic issues, we provide informal details and in the following discussion. There are two areas for which we provide specific attention: execution semantics and generalization.


Execution Semantics

For use cases to be precise, we must be able to describe what they mean when they are executed. The execution of a use case is performed in the context of a state consisting of instances of domain model classes. Domain objects may be created, modified, or deleted. Links may be added, or deleted. Attributes may be modified. The following capabilities describe the execution semantics of use cases.

1. Input / Output statements – instantiate formal parameter(s) specified in the statement with value(s) of the appropriate type(s).
2. Computation – create/delete an instance of a domain object, create/delete a link between existing domain object instances, modify the attribute of a domain object.
3. Exception – check exception conditions and fires if true.

Given a use case, execution proceeds as follows. Each statement in the main scenario is executed. During its execution, the alternatives / exceptions defined for that statement are checked for execution eligibility. If one or more is eligible, we must determine which one takes priority and execute it. The semantics for determining priority must be defined in the language. If the alternative specifies where control returns to, that is the point where execution continues, otherwise, it resumes at the next statement in the main interaction course.

While this is only an informal explanation of execution semantics, our group is exploring surface representations for the use case body and will provide a formal semantics based upon that representation in the future.



Generalization Semantics

Generalization has traditionally been a very problematic area for use case semantics. UML 2.0 says nothing about use case generalization, leaving it unclear whether generalization is meant to be a sub-typing mechanism, or a mechanism for simply indicating that the scenarios specified by a parent use case can be replaced with the scenarios specified by its children. If sub-typing is the goal, and use cases are re-aligned to specialize the Behavior metaclass, the design-by-contract calculus can serve as a basis for formalizing generalization. This will ensure the capacity to substitute the child use case for the parent in a context. If scenario replacement is the goal, a careful analysis of the issues related to replacement is required. Of course, a full semantics of use case generalization must be worked out in the context of an updated metamodel that supports both alignment with UML and a precise textual specification for a use case.

5 INDUSTRIAL APPLICATIONS OF ENGINEERED USE CASES


Before engineered use cases will be accepted in an industrial setting, they must provide value that exceeds the effort required to create them. We believe that engineered use cases can provide value across at least four dimensions: prototyping, estimation, refinement to design, and test generation. The first three are in the early research stage in our organization, while the fourth is one in which we have considerable experience.

Prototyping

The first benefit of using precise, engineered use cases is the ability to use them as the foundation for an executable prototype of the system. This allows early feedback to the customer and provides a means for validating the requirements. The prototyping approach requires that a precise surface syntax for specifying the use case content exists, along with an environment capable of executing these use cases. We have developed beta versions of these elements, and have shown that use cases captured with them can be executed directly. It also requires a facility for prototyping UI elements and associating them with use cases. The steps toward building a useful prototype are as follows:

1. Build the domain and use case models.
2. Create mocked-up user interface (UI) elements for the input and output statements.
3. Associate the UI elements with the input and output statements.
4. Execute the prototype with all involved stakeholders, gathering important feedback.
5. Modify the model according to the feedback and repeat the exercise until the requirements are agreed upon by all stakeholders.

The advantage of prototyping based on engineered use cases is obvious – it requires almost no effort beyond normal analysis and design activities to have an executable prototype for validating user requirements.



Estimation

Once the requirements have been validated using the prototyping approach described above, the next step is to determine how much time and cost are involved in realizing the actual system modeled by the use cases. Use cases alone cannot serve as a basis for estimation, but are an important element for determining the functional complexity of the application being developed. Non-functional requirements also play a significant role in determining time and effort.

We believe that our use case approach will serve as a sound basis for estimation, and we are in the very early stages of exploring this idea further. The basic steps in developing such an estimation approach are:

1. Examine existing projects using the engineered use case methods to gather baseline information for an estimation model.
2. Construct the initial estimation model, and apply it to new projects.
3. Gather feedback from the projects, and refine the estimation model.
4. Release the model for general use.


Refinement to Design

Once the estimate has been accepted and the work authorized, the next two activities occur in parallel. The first is to refine the use cases to a detailed design of the system. This is done as follows:

1. The domain model is refined to a full analysis model. This is done by identifying boundary and controller classes that serve as additions to the model classes already present in the domain model. The boundary classes can be inferred from the UI elements discussed in the section on prototyping.
2. Operations for all of the classes in the analysis model are identified. These are informed by the computation actions in the use case model, as well as required controller and UI computations for the system.
3. The input statements are mapped to operation invocations on the boundary classes.
4. The output statements are mapped to returned parameters from or operation invocations on the boundary classes.
5. Interaction diagrams are used to show how use cases are realized in the system using the mappings defined in steps 3-4 as a basis.
6. Further refinement to a detailed design and implementation can now proceed from the analysis model and use case realizations.

Test Case Creation/Generation

Finally, we must validate that the system constructed from the refinement described in the above section actually meets the requirements specified by the use cases. This is done by creating precise test cases from the use cases, domain model, and domain invariants. Our group has created test generation tools that automatically produce test cases from




these models. We have shown a decrease in the effort required to execute the test suite and an improvement in the coverage of the functionality of the system. Table 1 shows the results of the initial pilots performed with an early prototype version of this approach.

Pilot	Properties	Outcome
1	<ul style="list-style-type: none">• Test cases developed by Research• Provided to system owner for execution• Compared with existing test suite developed manually	<ul style="list-style-type: none">• Five-fold reduction in test suite size (30 generated test cases vs. 150 manually created) without a loss of functional coverage• 68% decrease in time to setup/run/verify test suite• Major defect detected by new suite that escaped detection by manually developed suite.
2	<ul style="list-style-type: none">• Modeling and generation done in testing organization.• Test expert created model and generated test cases• Comparison against a manually developed suite created independently by a different expert.	<ul style="list-style-type: none">• Coverage of use case functionality improved 30% in the model-based suite. This was achieved with 10% fewer use case invocations.• Testing team estimates achieving comparable coverage to model-based suite would require twice the effort.
3	<ul style="list-style-type: none">• Modeling and generation done in testing organization.• UML use case model built during system design, used by testing organization.• Testers analyzed the suite for usability in the project.	<ul style="list-style-type: none">• >90% coverage of the functionality with a suite of 60 test cases.• The suite tested 40 use cases. In total, there were 145 input parameters, 335 partitions, and 117 flow edges in the model.

Table 1. Pilot Results for Test Generation Tool

[Williams01] has specific details for these early results. The major conceptual steps in test generation are as follows:

1. For each use case, ensure that there are test cases for each path through the use case. This can be done using a technique such as the basis paths method [Poole95].
2. For each condition in a use case (either selection or exception), ensure that condition coverage is achieved using a condition coverage technique [Beizer90].



3. For each domain rule in the domain model, determine all use cases which update any variable in the domain rule. Determine the interesting ways in which the implementation could be tested to ensure the rule holds, and perform those tests.

6 RELATED WORK

Two recent papers explore extensions or alternatives to the standard UML metamodel for use cases. These include [Nakatana01], which integrates the activity graph metamodel and use case metamodel. In this view, use cases are still stateful classifiers, and the tight link between the actions comprising the use case and the domain model is missing. [Rui03] presents a metamodel that supports refactoring of use case models. This is interesting work, although the metamodel is not aligned with UML, nor does it pay careful attention to the context in which use cases are written.

Another common area for focus is the correct representation of use case structure. [Metz03] focuses on the necessity of being able to correctly and concisely specify alternative courses in a use case. This paper is related to the need to support alternative sequences of statements and exceptions, and defines the differences between these precisely. [Li00] offers a structured representation for use case content, and a semi-automated approach to translating use cases into sequence diagrams. This contribution is interesting for both representation and refinement purposes.

There is a significant body of work seeking to formalize the content or manner in which use cases are captured. Both [Lee98] and [Jorgensen04] seek to formalize use cases using concepts from Petri Nets. The first paper seeks to represent use cases using the formalism of Constraints-based Modular Petri Nets, while the second uses Colored Petri Nets. While both papers provide an elegant formal basis for use cases and both could be used for use case execution, neither Petri Net approach has been accepted in industry beyond the embedded/real-time space. Finally, [Rysavy04] seeks to represent use cases using higher-order logic. The benefits of this approach are an architecture that supports both static and behavioral information. Again, the downside is lack of industry acceptance of formal representations like HOL.

Finally, in the proposed application areas, there is significant work in the testing area [Basanieri02,Briand01] and the estimation area [Smith03]. The testing work typically uses refinements of use cases to interaction diagrams for generation, which differs from our approach [Williams01]. Estimation work to date has been somewhat disappointing in its accuracy, primarily due to the wide variation in the abstraction levels at which use cases are captured. We hope that our more defined approach will help mitigate these problems.



7 CONCLUSIONS

We have argued that for use cases to be truly useful for industrial software development an engineering approach to developing them must be used. Such an approach requires first that use cases be well defined, both structurally and behaviorally. To facilitate this definition, we proposed a method where use cases are developed in the context of a domain model and domain rules. The content of the use cases is carefully elaborated using this context.


Next, we proposed an approach for realigning the UML 2.0 metamodel for use cases to reflect the structure and capabilities described above, followed by a discussion of semantics issues associated with use cases. This focused on execution and generalization.

Finally, we presented four application areas that place engineered use cases in the context of the software lifecycle. These are prototyping, estimation, refinement to design, and test creation. We presented a general outline for using engineered use cases in each area.


What we have presented in this paper is a philosophical and conceptual framework for developing engineered use cases. What is lacking is the formalized machinery to ensure that the approach is well grounded and the complete tooling to support the endeavor. Our team is currently exploring these areas, as well as they key application areas defined above. We believe that precise, engineered use cases can be useful across the lifecycle and result in cost effective, reliable software systems.

REFERENCES

- [Armour01] F. Armour and G. Miller: *Advanced Use Case Modeling: Software Systems* Addison-Wesley, 2001.
- [Basanieri02] F. Basanieri, A. Bertolino, and E. Marchetti: “The Cow_Suite Approach to Planning and Deriving Test Suites in UML Projects”, *Proceedings of the 5th International Conference on UML (<<UML>> 2002)*, pp. 383-397, Sept. 30-Oct. 4, 2002.
- [Beizer90] B. Beizer: *Software Testing Techniques*, Van Nostrand Reinhold, 1990.
- [Briand01] L. Briand and Y. Labiche: “A UML-Based Approach to System Testing”, *Proceedings of the 4th International Conference on UML (<<UML>> 2004)*, pp. 194-208, October, 2001.
- [Cockburn01] A. Cockburn: *Writing Effective Use Cases*, Addison-Wesley, 2001.
- [Jacobson92] I. Jacobson: *Object-Oriented Software Engineering: A Use Case Driven Approach*, Addison-Wesley, 1992.



- [Jorgensen04] J.B. Jorgensen and C. Bossen: “Executable Use Cases: Requirements for a Pervasive Health Care System”, *IEEE Software*, vol. 21, no. 2, pp. 34-41, March-April, 2004.
- [Lee98] W.J. Lee, S.D. Cha, and Y.R. Kwon: “Integration and Analysis of Use Cases Using Modular Petri Nets in Requirements Engineering”, *IEEE Transactions on Software Engineering*, vol. 24, no. 12, December, 1998, 1115-1130.
- [Li00] L Li: “Translating Use Cases to Sequence Diagrams”, *Proceedings of the Fifteenth IEEE International Conference on Automated Software Engineering (ASE'00)*, pp. 293-296, September 11 - 15, 2000.
- [Metz03] P. Metz, J. O'Brien, and W. Weber: „Specifying Use Case Interaction: Types of Alternative Courses”, *Journal of Object Technology*, vol. 2, no. 2, pp. 111-131, March-April 2003, http://www.jot.fm/issues/issue_2003_03/article1.
- [Nakatana01] T. Nakatana, T. Urai, S. Ohmura, and T. Temai: “A Requirements Description Metamodel for Use Cases”, *Proceedings of the Eighth Asia-Pacific Software Engineering Conference (APSEC'01)*, pp. 251-258, December 04 - 07, 2001.
- [Poole95] J. Poole: *A Method to Determine a Basis Set of Paths to Perform Program Testing*. National Institute of Standards and Technology, Report #5737, November 1995.
- [Rui03] K. Rui and G. Butler: “Refactoring Use Case Models: The Metamodel”, *Proceedings of the 25th Australasian Computer Society Conference (ACSC 2003)*, pp. 301-308, 2003.
- [Rysavy04] O. Rysavy and F. Bures: “Formal Abstract Architecture for Use Case Specifications”, *Proceedings of the 11th IEEE International Conference and Workshop on the Engineering of Computer-Based Systems (ECBS'04)*, pp. 203-210, May 24-27, 2004.
- [Smith03] J. Smith: “The Estimation of Effort and Size Based on Use Cases”, IBM/Rational report available at <http://www3.software.ibm.com/ibmdl/pub/software/rational/web/whitepapers/2003/finalTP171.pdf>
- [UML02] *UML 2.0 Superstructure Specification*. OMG Adopted Specification ptc/04-10-02.
- [Williams01] C. Williams: “Toward a test-ready meta-model for use cases”, *Proceedings of the Workshop on Practical UML-based Rigorous Development Methods*, pp. 270-287, October, 2001.



About the authors

Clay Williams is a Research Staff Member and the Manager of the Software Quality and Testing Research Group at the IBM Watson Research Center. His group primarily focuses on issues related to modeling, testing, and measurement. His personal research focuses on effective use of modeling technology to enable the construction of high quality software. He has a Ph.D. in computer science from Texas A&M University. E-Mail: clayw@us.ibm.com

Matthew Kaplan is an Advisory Software Engineer in the Software Quality and Testing Research Group at the IBM Watson Research Center. His current work focuses on the use of modeling technology to improve testing, as well as how to model and use domain specific information (e.g. security information) to improve the quality of software. He has an M.S. in computer science from Brown University. E-Mail: mmk@us.ibm.com

Tim Klinger is an Advisory Software Engineer in the Software Quality and Testing Research Group at the IBM Watson Research Center. He primarily focuses on the application of novel techniques from AI to complex problems across the software engineering lifecycle, from design and construction to testing and validation. He has a Ph.D. from New York Univeristy. E-Mail: tklinger@us.ibm.com.

Amit Paradkar is a Research Staff Member in the Software Quality and Testing group at IBM T.J. Watson Research Center, Yorktown, NY. His primary research interest is in the area of software testing - particularly model based test generation using UML. He is also interested in software testability, and applications of AI techniques to Software Engineering problems. He has received Ph.D. in Computer Science, MS in Computer Science, and MS in Textile Engineering and Sciences - all from North Carolina State University in Raleigh. E-Mail: paradkar@us.ibm.com