





Figure 1: An object behavior model for the JAVA `Vector` class.

process makes our models mirror actual objects' behavior. Like all dynamic analyses, we build on the observation, that *common behavior is often correct behavior*; thus, our models are likely to represent *universal invariants*.

As an example, consider an object behavior model describing the behavior of the JAVA `Vector` class (Figure 1). The model has two states: one, where the vector is empty and one, where the vector holds at least one element. Transitions (and lack thereof) provide us with information about which method calls were, and which were not observed in a particular state during execution and how they changed the state of an object.

We can for instance see, that calling `Vector.add()` always ended up in a non-empty state and that calling `Vector.clear()` always caused the vector to be empty. We can also notice, that a call to `Vector.remove()` was never observed on an empty vector, and this is not without a reason. Calling `Vector.remove()` on an empty vector always fails with an exception being thrown.

On the other hand, some transitions may be non-deterministic. This is the case for `Vector.remove()` method call. Calling this method on a non-empty vector may either change the state of the vector or not. Knowledge in what states a particular method call was observed, and how the call changed that state allows us to extract pre- and postcondition for this method. For example, in case of `Vector.remove()`, it is necessary that the method is called on a non-empty vector, but the call can result in an empty vector as well as a non-empty one.

## 4 Related Work

Usage of finite state automata to abstract behavior of the program was investigated by many researchers in the past years, with Cook and Wolf (1998) being the seminal work about inferring finite state automata from event sequences.

In most approaches, anonymous states have been used (Cook and Wolf, 1998; Ammons et al., 2002). Some researchers have referred to implementation details when labeling states, like in the work of Whaley et al. (2002), who used variables for this purpose.

The work closest to ours is by Xie and Notkin (2004), but, unlike us, they need test cases to generate models and they do not restrict calls to only pure methods when getting the state of an object. Additionally, they do not abstract from concrete values to avoid state space explosion.

## 5 Conclusions and Future Work

Object behavior models capture essential properties of an object from the view of the object's client. Applying partitioning of class methods into *inspectors* and *mutators*, and using return values of calls to inspectors to represent an object's state distinguishes our work from prior work and is a central contribution of our approach.

As a future work, we plan to use object behavior models to check dynamically, whether executed code does not violate previously learned behavior. We also want to enhance models by extracting state of objects being returned from inspector method calls. This would allow us to express an object's state as the state of its constituents. Another idea is to check programs statically against mined models. Deviations from those models may point us to incorrect usage of an API.

**Acknowledgements** Christian Lindig provided valuable comments on earlier revisions of this paper.

## References

- Glenn Ammons, Rastislav Bodík, and Jim Larus. Mining specifications. In *Conference Record of POPL'02: The 29th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 4–16, Portland, Oregon, January 16–18, 2002.
- J. Cook and A. Wolf. Discovering Models of Software Processes from Event-Based Data. *ACM Transactions on Software Engineering and Methodology*, 7(3):215–249, July 1998.
- Alexandru Salcianu and Martin Rinard. Purity and side effect analysis for Java programs. In *Proceedings of the 6th International Conference on Verification, Model Checking and Abstract Interpretation*, number 3385 in LNCS, pages 199–215, January 2005.
- John Whaley, Michael Martin, and Monica Lam. Automatic extraction of object-oriented component interfaces. In Phyllis G. Frankl, editor, *Proceedings of the ACM SIGSOFT 2002 International Symposium on Software Testing and Analysis (ISSTA-02)*, volume 27(4) of *SOFTWARE ENGINEERING NOTES*, pages 221–231, New York, July 22–24 2002. ACM Press.
- Tao Xie and David Notkin. Automatic extraction of object-oriented observer abstractions from unit-test executions. In *Proceedings of the 6th International Conference on Formal Engineering Methods (ICFEM 2004)*, pages 290–305, November 2004.