Self Adaptive Software

Howard Shrobe & Robert Laddaga

Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

http://www.ai.mit.edu

The Problem: Can we find a way to build software systems which can autonomously adapt to unforeseen and changing circumstances? What are the organizing principles for such systems? Is there a common infrastructure that can support all self-adaptive software systems.

Motivation: We are concerned with the problem of building mission-critical systems that operate non-stop in a rapidly changing, and possibly hostile environment. It is unacceptable for these systems to "shatter" when confronted with change; they must instead fluidly adapt to circumstance by: figuring out what is wrong, identifying an alternative approach to getting the job done, establishing a consistent state from which to begin the recovery and then get the job done as well as possible. Reporting a bug and waiting for a new system build is simply not a viable option.

Previous Work: Our project involves many tasks for which there is prior work. In particular, work by Davis and Shrobe [1, 2, 3] and Williams and deKleer [7, 5, 4, 6, 8] on model-based troubleshooting is relevant to the diagnostic component of the task, work on the design of dynamic object oriented languages in the Lisp tradition also provide valuable insights. Work in the programmer's apprentice project by Shrobe, Rich and Waters [9] provides many of the underlying representations needed for this task.

Approach: Our approach to building adaptive software is based on the notion of a Dynamic Domain Architecture (DDA). Dynamic Domain Architectures structure an application domain into layers of common services where each service has a number of variant implementations tailored to different environmental conditions. The architectural level of description also provides "purpose links" which explain how the components of service achieve its overall goals. The DDA development environment synthesizes run-time sentinels to monitor the preconditions of the purpose links. The runtime services of the DDA are invoked when a sentinel signals the failure to achieve an expected condition; the runtime services are responsible for diagnosis of the failure and for selection and execution of a repair procedure. Since the DDA provides many alternative implementations of each component, a typical repair involves rolling back to a recovery point and invoking an alternative implementation.

One of the critical features of the runtime infrastructure of a Dynamic Domain Architecture is its ability to self diagnose. We are working on reformulating earlier work on model based troubleshooting so that it can be applied to failures in software systems. Model based troubleshooting reasons backwards from symptoms, the failure of a system to behave as expected, to an underlying cause, a description of how a specific component of the system failed to behave correctly. Such failures can be due to failures of the underlying operating system, network or computing hardware, or to an incompleteness in the code's coverage of the phenomenon (as is often the case in image understanding, for example). Reasoning about the behavior of a software system is notoriously difficult, but it can be made easier if the system gathers traces of its behaviors at many places. Runtime diagnosis is, therefore, partly enabled by code transformations in the development environment which synthesize data collectors, error detectors, condition handlers and restart locations. We are working on developing this infrastructure. Once a diagnosis is made, the system must reconfigure itself into a consistent state appropriate for the restart. This requires significant cleanup activity which would be much simpler if the memory of the main behaved transactionally; we are investigating techniques for making transactional memory and studying whether these techniques are appropriate for our needs.

Difficulty: This project is attempting to develop an infrastructure from an entirely new family of software systems. Although each of the components of this infrastructure has its precedents, the overall ensemble is entirely novel and represents a major challenge.

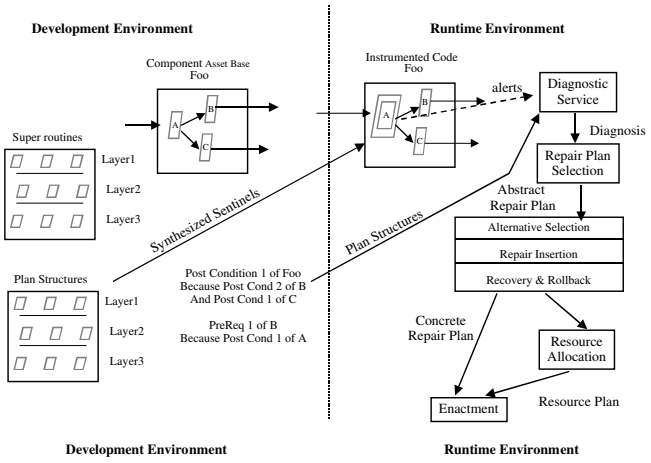

Figure 1: Dynamic Domain Architectures

Impact: Increasingly computer systems are designed to be embedded in and interact with their environments and to function over long periods of time. Such systems will experience variations of conditions which are impossible for the designers to explicitly anticipate, nevertheless the systems must be prepared to deal with cope with these circumstances. We currently have no general methodology or infrastructure adequate for this task. If our project succeeds, it will revolutionize the way such systems are designed and implemented.

Future Work: We are applying many of these ideas to resource rich, perceptually enabled environments such as the MIT Intelligent Room.

Research Support: This work is funded by DARPA under contract number F30602-97-2-0013, administered by Rome Laboratory.

References:

[1] R. Davis, H. Shrobe, W. Hamscher, K. Wieckert, M. Shirley, and S. Polit. Diagnosis based on descriptions of structure and function. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, pages 137–142, 1992.

[2] Randall Davis. Diagnostic reasoning based on structure and behavior. *AI*, 24:347–410, 1984.

[3] Randall Davis and Walter Hamscher. Model-based reasoning: Troubleshooting. In Howard Shrobe, editor, *Exploring Artificial Intelligence: Survey Talks from the National Conferences on Artificial Intelligence*, pages 297–346. Morgan Kaufman, 1988.

[4] J. de Kleer. Using crude probability estimates to guide diagnosis. *Artificial Intelligence*, 45:381–391, 1990.

[5] J. de Kleer and B. C. Williams. Diagnosis with behavioral modes. In *Proceedings of IJCAI-89*, pages 1324–1330, 1989.

[6] Johan de Kleer. Focusing on probable diagnoses. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, pages 842–848, 1991.

[7] Johan de Kleer and Brian C. Williams. Diagnosing multiple faults. *Artificial Intelligence*, 32:97–130, 1987.

[8] Johan de Kleer and Brian C. Williams. Diagnosing multiple faults. *AI*, 32:97–130, 1987.

[9] C. Rich, H. E. Shrobe, and R. C. Waters. An overview of the programmer's apprentice. In *IJCAI79*, 1979.