Model-Based Hybrid Systems

Brian C. Williams, Michael W. Hofbaur & Melvin M. Henry

Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

http://www.ai.mit.edu

The Problem: The ability to monitor and diagnose complex physical devices is critical for building highly autonomous artifacts that can operate robustly in harsh environments over a long period of time. We propose a hybrid monitoring, diagnosis and model learning capability for physical devices that exhibit complex discrete and continuous behaviors.

Motivation: The year 2000 was kicked off with two missions to Mars, following on the heals of the highly successful Mars Pathfinder mission. Just before the first mission, Mars Climate Orbiter, reached Mars orbital insertion, the operations team identified contradictory attitude estimates, one suggesting that the vehicle was coming in at an elevation too low to successfully achieve insertion. Unfortunately time did not allow the team to resolve this inconsistency or to plan a reliable course correction before insertion. The Climate Orbiter proceeded until it burned up in the Martian atmosphere. After extensive investigation it was found that a table used by the navigation system, which describes the small forces impinging upon the space vehicle, was mistranslated into the wrong units. This bug introduced a small, but indiscernible failure that over a lengthy time period produced the loss of the orbiter.

The problem of misinterpreting a system's dynamics was punctuated later in the year when the Orbiter's sibling, Mars Polar Lander, vanished without a trace. After months of analysis the failure investigation team concluded that the vehicle most likely crashed into Mars because it incorrectly shutdown its engine at 50 meters above the surface. This failure, like the orbiter, resulted from a misinterpretation of the vehicle's dynamics, in this case due to a faulty software monitor.

Figure 1: Mars Polar Lander

The above case study is a dramatic instance of a common problem – increasingly complex systems are being developed whose failure symptoms are nearly indiscernible up until a catastrophic result occurs. To tackle this problems we must address two issues. First, these failures are manifest through a coupling between a system's continuous dynamics and its evolution through different behavior modes. Hence to address this problem we need hybrid monitoring and diagnosis capabilities that are be able to track a system's behavior along both its continuous state changes and its discrete mode changes. Second, failures may generate symptoms that are initially on the same scale as sensor and actuator noise. To discover these symptoms statistical methods need to be applied to separate the noise from the true dynamics.

Previous Work: Our work brings together methods developed in the fields of common sense modeling, model-based diagnosis, hybrid systems, and control. For work on common sense modeling we look to research in qualitative physics[7]. For work on diagnosis and monitoring we look to research in model-based diagnosis[4][2][10] and belief state update using Hidden Markov Models[5]. Finally, hybrid systems and control theory provides important foundations for hybrid modeling[1] and estimation[6].

Approach: The starting point for our approach is a modeling formalism, called *concurrent probabilistic hybrid automata* (cPHA), that merges hidden Markov models with continuous dynamical system models. These models capture both, the discrete and continuous behaviors, that are exhibited by complex physical devices. The hybrid modeling paradigm extends the modeling paradigm used in previous model-based reasoning and diagnosis systems, such as the Livingstone System[10].

System monitoring and diagnosis, i.e. estimating the current mode of the system, will be formulated by extending the mode identification capability of Livingstone to handle this quantitatively enriched representation. To achieve efficiency, and to focus the systems attention on the most plausible diagnoses, we will formulate the hybrid mode estimation algorithm as a generalization of conflict directed A* search, used in Livingstone and its predecessor Sherlock.

Finally, we address the challenge of learning and refining models of mixed, software hardware behaviors, as cPHA. Large-scale model learning has been highlighted as one of the grand challenges of machine learning. Progress has been particularly significant in the area of Bayes net learning. Learning of cPHA is a new open challenge for machine learning with broad application. We address the challenge of large-scale cPHA model learning by substantially generalizing our previous work on decompositional model-based learning of purely continuous systems[9].

Impact: Through the development of hybrid mode estimation and hybrid model learning capabilities, we will extend our ability to build highly robust (fault-tolerant) systems that can operate robustly in harsh environments over a long period of time.

Future Work: We will generalize our hybrid modeling paradigm by formulating an expressive compositional hybrid model-based programming language (HMPL), which will provide a rich set of primitives that can be used to describe programs, formulate uncertainty and represent continuous dynamics. HMPL will unify past work on Hybrid CC[3] and the reactive model-based programming language RMPL[8]. HMPL will draw upon HCC's ability to describe continuous dynamics and RMPL's ability to describe uncertainty and cost of actions and will provide a rich set of primitives that can be used to describe programs, formulate uncertainty and represent continuous dynamics. We will compile this expressive model into a concurrent probabilistic hybrid automaton (cPHA), which serve as the basis for applying our proposed reasoning methods.

Research Support: This research is supported by NASA under contract NAG2-1388.

References:

[1] M. Branicky, V. Borkar, and S. Mitter. Unified framework for hybrid control: Model and optimal control theory. *IEEE Transactions on Automatic Control*, 43:31–45, 1998.

[2] J. de Kleer and B. Williams. Diagnosing multiple faults. *Artificial Intelligence*, 32(1):97–130, 1987.

[3] V. Gupta, R. Jagadeesan, and V. Saraswat. Computing with continuous change. *Science of Computer Programming*, 30(1–2):3–49, 1998.

[4] W. Hamscher, L. Console, and J. de Kleer, editors. *Readings in Model-Based Diagnosis*. Morgan Kaufmann, San Mateo, CA, 1992.

[5] L. Kaebling, M. Littman, and A. Cassandra. Planing and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1-2):99–134, 1998.

[6] X.R. Li and Y. Bar-Shalom. Multiple-model estimation with variable structure. *IEEE Transactions on Automatic Control*, 41:478–493, 1996.

[7] D. Weld and J. de Kleer, editors. *Qualitative Reasoning about Physical Systems*. Morgan Kaufmann, San Mateo, CA, 1990.

[8] B. Williams, S. Chung, and V. Gupta. Mode estimation of model-based programs: Monitoring systems with complex behavior. In *Proceedings of the International Jount Conference on Artificial Intelligence (IJCAI'01)*, pages 579–585, 2001.

[9] B. Williams and B. Millar. Decompositional, model-based learning and its analogy to diagnosis. In *Proceedings of the 15th National Conference on Artificial Intelligence (AAAI-98)*, 1998.

[10] B. Williams and P. Nayak. A model-based approach to reactive self-configuring systems. In *Proceedings of the 13th National Conference on Artificial Intelligence (AAAI-96)*, 1996.