Improving Model-based Mode Estimation through Offline Compilation

Seung H. Chung, John M. Van Eepoel, Brian C. Williams

MIT Space Systems Laboratory
MIT Artificial Intelligence Laboratory
Cambridge, MA, 02139
{chung, vanny, williams}@mit.edu

Many recent and future space missions point to the need for more capable fault diagnostic systems. Many fault management systems are based on expert systems in which a rule-based diagnostic engine is used to detect faults. To create these rules, engineers must reason through system wide interactions, consequently, the set of rules is limited by the faults that engineers can recognize. This lack of robustness can be detrimental to the spacecraft. In contrast, model-based fault protection systems eliminate the need to enumerate all rules by reasoning on the common-sense model of a system. Model-based fault protection systems such as GDE [1] and Livingstone [2], however, are limited by the uncertainty in the time required to perform onboard reasoning.

As a solution to this problem, a Miniature Mode Estimation system (Mini-ME) is introduced. Mini-ME estimates the current mode of the spacecraft and detects any failures through model-based reasoning. Mini-ME is composed of offline and online steps.

The first offline task in Mini-ME is to compile a commonsense model specified by engineers into *dissents*. Each dissent may be viewed as a rule that identifies a conflict between observations and component modes. The advantage of compiling a model into a set of dissents is that all irrelevant information in the model is removed and only the information necessary for diagnosis is kept. Since the dissents are generated offline, this does not factor in to the real time performance of the system.

Mini-ME rewrites the dissents into the partial diagnosis rules that are more useful for fault diagnosis. In diagnosis the objective is to determine the modes of components that agree with the current observations. A partial diagnosis resolves a conflict of some specified observations.

The first online step in Mini-ME is triggering rules. If the observation in a partial diagnosis rule corresponds to the current observation, then the component modes the rule implies are the relevant partial diagnoses. The advantage of Mini-ME over other model-based fault diagnosis systems is that the NP-complete satisfiability problem is removed from online computation. Instead of requiring possibly an exponential satisfiability search, Mini-ME generates partial diagnoses in time that is linear in the number of rules associated with a set of observations.

First introduced in GDE [1], the final step in Mini-ME is to generate the most likely kernel diagnosis from the partial diagnoses. A kernel diagnosis is a minimal set of component modes that resolves a full set of conflicts. The objective is to find the minimal set covering in the best-first order such that the search time is minimized.

In rule-based systems, generating the mapping from symptoms to recovery action is not as intuitive, yet the engineer must enumerate the rules manually. Mini-ME, however, automatically enumerates these rules from the commonsense models specified by the engineer. While rules are not inherently intuitive, generating a system model is apparent. Mini-ME marries the benefits of fast responsiveness of rule-based systems, and the use of models to specify the behavior of a system.

References

[1] J. de Kleer and B. Williams. "Diagnosing Multiple Faults," Artificial Intelligenc, 32:100-117, 1987.

[2] B. Williams and P. Nayak, "A Model-based Approach to Reactive Self-Configuring Systems," In Proceedings of AAAI-98, 971-978, 1996.

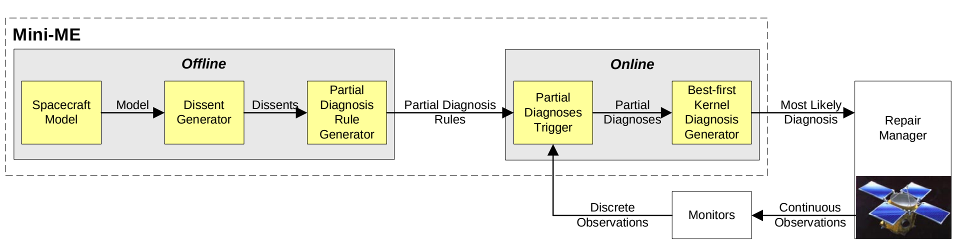

Figure 1: Mini-ME Architecture