Distributed Monitoring and Diagnosis using Constraint Optimization Methods

Martin Sachenbacher and Tsoline Mikaelian

What: Our work focuses on the development of a distributed model-based mode estimation capability that will ensure robustness of aerospace systems by endowing them with intelligence that allows to collectively reason about faulty components and diagnose failures. Our approach is based on framing the mode estimation task as a constraint optimization problem, leveraging structural properties to decompose it into independent sub-problems, and distributing it among different physical units. The mode estimation capability is part of a model-based autonomy architecture that can exploit the results to take appropriate repair actions in order for the mission to proceed.

Why: Complex, multi-vehicle spacecraft systems, such as SPHERES (Fig. 1), are currently being developed in order to meet the challenging requirements of future aerospace missions. However, the increasing complexity of aerospace systems is also increasing the potential for mission failure due to malfunctions and unexpected component behavior. A capability for automated monitoring and diagnosis will reduce the impact of failures and increase the likelihood of mission success and science return. Engineering models of spacecraft components can be exploited to reason about possible component interactions and to identify and diagnose faults and take recovery actions. This provides a robust approach to fault management, where the task of the modeler is reduced to developing commonsense engineering models of spacecraft components. The challenge is to develop efficient algorithms that can reason through a large space of component interactions and find the most likely state consistent with the model and the observations.

How: The task of mode estimation is to find the most likely sequence of states that is consistent with the behavior model of the spacecraft and the sequence of commands and observations over time. This can be framed as constraint optimization problem where one must determine the most likely states at each time point and combine them into shortest path trajectories. Since the possible assignments can grow exponentially with the number of variables, the central problem is state explosion that leads to computational intractability if the number of variables or possible values is large.

The first part of our work aims at leveraging favorable properties of the mode estimation optimization problem to combat state explosion. Our research combines a range of different methods from different areas and investigates possible trade-offs between them:

- using soft constraints, such as semiring-CSPs [1], as a uniform representational framework for both constraints and uncertainty;
- decomposing constraint networks into trees [2, 3] as a basis for exploiting independence of subproblems through local consistency and dynamic programming;
- enumerating solutions in best-first order [6], allowing for anytime behavior and efficient focusing on leading solutions;
- encoding constraints symbolically [4] to exploit regularities and compactly represent very large belief spaces.

Figure 1: EO-1 earth observing satellite (left) and SPHERES formation flying test-bed (right)

The second part of our work aims at leveraging the performance improvements of the underlying constraint processing to introduce more accurate diagnostic algorithms, such as the n-step mode estimation algorithm. This algorithm uses a sliding window of n previous observations and commands in order to estimate and track the most likely trajectories of state evolution. Since it can exploit a more "global" view, n-step mode estimation will provide more accurate diagnostic results than current approaches to mode estimation that are based on maintaining k best estimates [5].

The third part of our work investigates ways to distribute the mode estimation task among different physical units, such as a fleet of spacecraft. We plan to accomplish this by exploiting the decomposition of constraint networks into hierarchies of independent sub-problems. Our approach to distributed mode estimation thus consists of two phases: an off-line compilation phase which compiles the constraint graph into a tree structure, and an on-line mode estimation phase which computes the solutions to the tree-structured problem using a distributed version of dynamic programming.

Progress: Up to now we have focussed on a centralized version of mode estimation. We developed a prototype for a constraint optimization engine that integrates structural decomposition, symbolic encoding, and best-first enumeration. This is based on our previous work in constraint optimization [6]. We also outlined a first version of an n-step mode estimation algorithm that extends the k-best mode estimation algorithm developed in [5] to get more accurate diagnosis results.

Future: Our next steps will address the issue of distributing the mode estimation sub-problems between different physical processors. This mapping must ensure minimal cost, processing power and communications among processors. As the distributed mode estimation capability will form a basic building block in the model-based autonomy architecture, it is also important to integrate the mode estimation capability with distributed pre-planning and reconfiguration. We will demonstrate and validate the distributed model-based executive on models of the EO-1 satellite and the SPHERES test-bed flying in the International Space Station (Fig. 1).

Research Support: This research is funded by the DARPA programs NEST and MOBIES under Air Force Research Lab Contract F33615-01-C-1896 (NEST) and Air Force Research Lab Contract F33615-00-C-1702 (MOBIES).

References:

[1] Stefano Bistarelli, Ugo Montanari, and Francesca Rossi. Semiring-based constraint satisfaction and optimization. *Journal of the ACM*, 44(2):201–236, 1997.

[2] Georg Gottlob, Nicola Leone, and Francesco Scarcello. A comparison of structural CSP decomposition methods. *Artificial Intelligence*, 124(2):243–282, 2000.

[3] Kalev Kask, Rina Dechter, Javier Larrosa, and Fabio Cozman. Unifying tree-decomposition schemes for automated reasoning. Technical report, University of California, Irvine, 2001.

[4] R.I. Bahar, E.A. Frohm, C.M. Gaona, G.D. Hachtel, E. Macii, A. Pardo, and F. Somenzi. Algebraic Decision Diagrams and Their Applications. In *IEEE /ACM International Conference on CAD*, pages 188–191, Santa Clara, California, 1993. IEEE Computer Society Press.

[5] B. Williams and P. Nayak. A model-based approach to reactive self-configuring systems. In *Proceedings of AAAI-96*, pages 971–978, 1996.

[6] Brian Williams and Robert Ragno. Conflict-directed A* and its role in model-based embedded systems. *Journal of Discrete Applied Mathematics*, 2003. To appear.