



## High-Performance Algorithm for Solving the Diagnosis Problem

Computation time is reduced substantially.

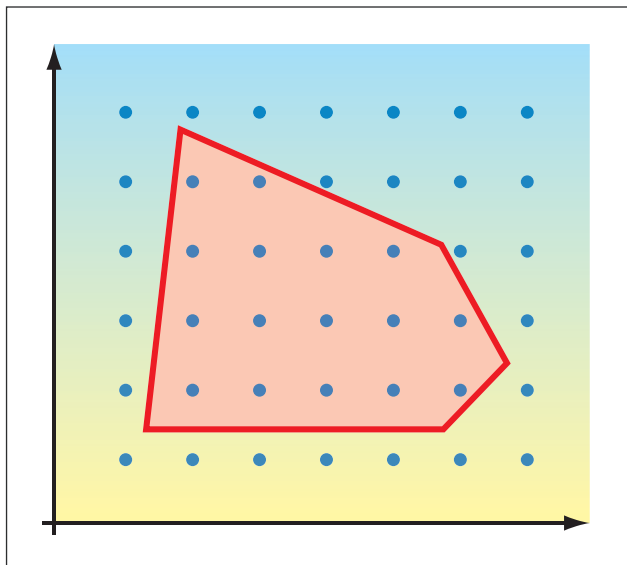
NASA's Jet Propulsion Laboratory, Pasadena, California

An improved method of model-based diagnosis of a complex engineering system is embodied in an algorithm that involves considerably less computation than do prior such algorithms. This method and algorithm are based largely on developments reported in several NASA Tech Briefs articles: "The Complexity of the Diagnosis Problem" (NPO-30315), Vol. 26, No. 4 (April 2002), page 20; "Fast Algorithms for Model-Based Diagnosis" (NPO-30582), Vol. 29, No. 3 (March 2005), page 69; "Two Methods of Efficient Solution of the Hitting-Set Problem" (NPO-30584), Vol. 29, No. 3 (March 2005), page 73; and "Efficient Model-Based Diagnosis Engine" (NPO-40544), on the following page.

Some background information from the cited articles is prerequisite to a meaningful summary of the innovative aspects of the present method and algorithm. In model-based diagnosis, the function of each component and the relationships among all the components of the engineering system to be diagnosed are represented as a logical system denoted the system description (SD). Hence, the expected normal behavior of the engineering system is the set of logical consequences of the SD. Faulty components lead to inconsistencies between the observed behaviors of the system and the SD. Diagnosis — the task of finding faulty components — is reduced to finding those components, the abnormalities of which could explain all the inconsistencies. The solution of the diagnosis problem should be a minimal diagnosis, which is a minimal set of faulty components.

The calculation of a minimal diagnosis is inherently a hard problem, the solution of which requires amounts of computation time and memory that increase exponentially with the number of components of the engineering system.

Among the developments to reduce the computational burden, as reported in the cited articles, is the mapping of the diagnosis problem onto the integer-programming (IP) problem. This mapping makes it possible to utilize a variety of algorithms developed previously for IP to solve the diagnosis problem. In the IP approach, the diagnosis problem can be formulated as a linear integer optimization



The Sets of Solutions of the LP relaxation and IP problems are represented by (1) the polygon and (2) the grid points inside the polygon, respectively.

problem, which can be solved by use of well-developed integer-programming algorithms. This concludes the background information.

The point of departure for the development of the present improved method and algorithm is the discovery that the mapping of the diagnosis problem onto the IP problem makes it possible to calculate lower and upper bounds on the number of faulty components before solving the diagnosis problem. A solution window for the problem is introduced, based on these bounds.

One of the algorithms heretofore commonly used to solve the IP problem is known in the art as the branch-and-bound algorithm. In terms that are necessarily oversimplified for the sake of

brevity, the traditional version of the branch-and-bound algorithm includes the following:

- Using non-integer and integer solutions of the corresponding linear-programming (LP) relaxation problem (see figure), in which the variables of the IP problem are no longer constrained to be integers, to define lower and upper bounds, respectively;
- Successively dividing (branching) the IP problem into subproblems on the basis of the bounds; and
- Eliminating any subproblem, the LP lower bound of which exceeds the upper bound of another subproblem.

The branching, bounding, and elimination are repeated until all subproblems are eliminated or an integer solution is found.

The present algorithm is an alternative version of the branch-and-bound algorithm. In this algorithm, the amount of computation required for execution is reduced, relative to that of the traditional version, by exploiting the structure of the problem. The amount of computation is further reduced by using the solution window to effect a massive pruning of branches. Yet another reduction is effected by dynamically updating and narrowing the window to enable further pruning. In simulation tests on systems of 40 components, the present algorithm was found to solve the diagnosis problem at an average speed about 300 times faster than that of the traditional branch-and-bound algorithm.

This work was done by Amir Fijany and Farrokh Vatan of Caltech for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP (see page 1).

The software used in this innovation is available for commercial licensing. Please contact Karina Edmonds of the California Institute of Technology at (626) 395-2322. Refer to NPO-41456.