Reducing Sequential Diagnosis Costs by Modifying Reiter's Hitting Set Tree

Patrick Rodler* and **Manuel Herold**
Alpen-Adria Universität Klagenfurt
e-mail: firstname.lastname@aau.at

Extended Abstract

When systems such as software, physical devices or knowledge bases do not exhibit required properties, there is often a substantial number of competing fault explanations, called *diagnoses*. Sequential Diagnosis aims at suggesting a minimal-cost sequence of measurements to identify the root cause of a system failure, i.e. the diagnosis that pinpoints the actually faulty system components. Since the minimization of the overall measurement costs is NP-hard, sequential diagnosis methods rely on local optimizations to approximate the optimal sequence of measurements. To this end, an iteratively (re-)computed set of (e.g., the most probable) diagnoses usually serves as a decision basis for the selection of the best next measurement. The computation of diagnoses is often accomplished by hitting set algorithms, due to their favorable properties in terms of general applicability, independence of specific theorem provers, and ability to generate diagnoses in best-first order.

In this work, we argue that sequential diagnosis can be interpreted in two natural ways, as a static (*StatSD*) or dynamic (*DynSD*) problem, and that existing methods focus only on the latter. Both problem formulations assume a *diagnosis problem instance (DPI)*—consisting of knowledge about the system, its components, and the so-far made observations and measurements—being given as an input. The main difference between DynSD and StatSD is the way how new measurements are taken into account. In DynSD, they are directly used to extend (the measurements set of) the DPI, i.e., the relevant *DPI is dynamically changing* throughout the sequential process, each time shifting the focus to the *diagnoses space of the extended DPI*. StatSD, in contrast, assumes the (initial) *DPI is static* and does not add new measurements to it. Instead, it uses the measurements as constraints to prune the, quasi "frozen", *search space of diagnoses for the fixed DPI*.

We prove that, under certain conditions, solving StatSD leads to a lower expected number of measurements compared to DynSD. In general, however, StatSD in incomplete and gives no guarantee that the actual diagnosis will be found, as opposed to DynSD. Hence, DynSD and StatSD can be seen as two extremes of a *generalized view on sequential diagnosis*; on the one extreme, new information is *always* used to switch to a new DPI (implying completeness), whereas, on the other extreme, new information is *never* used to switch to a new DPI (implying less measurement costs).

To combine the benefits of both strategies, we introduce a novel variant of Reiter's hitting set algorithm that implements this generalized sequential diagnosis process where DPI-switches can take place anytime, while guaranteeing completeness. Supporting our theoretical results, empirical examinations using real-world problems reveal that the new algorithm—even with a fairly simple DPI-switching strategy based on the search tree size—substantially reduces the required effort to locate the actual diagnosis, compared to (sound and complete best-first) diagnoses computation algorithms that switch DPIs in every iteration. In concrete figures, an average of 20% and up to 65% of the user interaction cost to diagnose the system could be saved and the new strategy led to equal or lower measurement costs in 97% of the cases. Moreover, the proposed algorithm is as generally applicable as Reiter's method as well as fully compatible and able to synergize with other measurement selection approaches or optimization heuristics such as entropy.

*This work was supported by the Carinthian Science Fund (KWF), contract KWF-3520/26767/38701.