

Summary of “Sequence-Oriented Diagnosis of Discrete-Event Systems”

Gianfranco Lamperti ✉ 


Department of Information Engineering, University of Brescia, Italy

Stefano Trerotola ✉

Department of Information Engineering, University of Brescia, Italy

Marina Zanella ✉ 

Department of Information Engineering, University of Brescia, Italy

Xiangfu Zhao ✉ 

School of Computer and Control Engineering, Yantai University, China

Abstract

What follows is an extended abstract of a paper recently published in the *Journal of Artificial Intelligence Research* [1]. Broadly speaking, the approaches to model-based diagnosis of discrete-event systems (DESs) can be accommodated within a spectrum ranging from no to total knowledge-compilation. That paper presents three approaches, one for either end of the spectrum, and one falling in between, in order to compute diagnoses that are sequences of faults.

In the literature, model-based diagnosis has always been conceived as *set-oriented*, meaning that a candidate is a *set* of faults, or faulty components, that explains a collection of observations. This perspective applies to both static and dynamical systems. Diagnosis of DESs is no exception: a candidate is traditionally a set of faults occurring in a trajectory of the DES that conforms with a given sequence of observations, namely a *temporal observation*. To improve diagnostic explanation and support decision making, a *sequence-oriented* perspective to diagnosis of DESs is presented, where a candidate is a *sequence* of faults occurring in a trajectory of the DES that produces the temporal observation. That candidate, called a *fault sequence*, differs from a classical candidate mainly in three ways: (i) it includes the *multiset* of faults occurred in the trajectory, where all the occurrences of the same fault are encompassed, (ii) it provides a total temporal order of faults, and (iii) its length may be unbounded, as the same fault may occur an unlimited number of times in the trajectory. Therefore, in a sequence-oriented perspective, the diagnosis output is the (possibly infinite) set of all distinct fault sequences, each relevant to a (possibly infinite) set of trajectories of the DES that produce a given temporal observation. Having a possibly unbounded set of candidates contrasts with set-oriented diagnosis, where the set of candidates is bounded by the powerset of the domain of faults. Still, a possibly unbounded set of fault sequences is shown to be a regular language, which can be defined as a regular expression over the domain of faults: this regular expression represents the output of sequence-oriented diagnosis.

The additional information on the chronological order of faults and their multiple occurrences embedded in a fault sequence may be important for ranking purposes, as well as for supporting diagnosticians and/or troubleshooting algorithms to make decisions in critical scenarios and, possibly, to take some sequential diagnosis steps. Since the output of sequence-oriented diagnosis includes additional pieces of information with respect to the output of set-oriented diagnosis, two questions raise quite naturally: *how can the sequence-oriented task be performed, and which is its performance?* The paper addresses the former question for a class of DESs usually considered by the authors (*active systems*) in three ways, and shows some experimental evidence to answer the latter.

Specifically, three sound and complete techniques to perform the task of *monitoring-based* diagnosis are described and compared, where a new candidate set is generated at the occurrence of each observation, namely: (1) *blind* diagnosis, with *no* compiled knowledge, (2) *greedy* diagnosis, with *total* knowledge compilation, and (3) *lazy* diagnosis, with *partial* knowledge compilation. By *knowledge* we mean a data structure, slightly similar to a classical DES diagnoser, which is called an *explainer*, that can be generated (compiled) either entirely offline (greedy diagnosis) or incrementally online (lazy diagnosis). Among them, only the blind diagnosis engine relies solely on the model of the DES, without producing any data structure offline nor performing any offline processing.



© Gianfranco Lamperti, Stefano Trerotola, Marina Zanella, and Xiangfu Zhao; licensed under Creative Commons License CC-BY 4.0

35th International Conference on Principles of Diagnosis and Resilient Systems (DX 2024).

Editors: Ingo Pill, Avraham Natan, and Franz Wotawa; Article No. 34; pp. 34:1–34:2

OpenAccess Series in Informatics



OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

It constructs and iteratively updates the portion of a DES behavioral space that conforms with the temporal observation perceived so far. To compute the resulting regular expression, a state-elimination algorithm has been exploited from the literature, which takes as input a non-deterministic automaton and generates the regular expression of the accepted language. In greedy diagnosis, the explainer generated offline has a size that is smaller than the size of the classical diagnoser that enables the online computation of set-oriented candidates. However, its smaller size comes with a disadvantage: the explainer is a non-deterministic finite automaton, whereas the diagnoser is deterministic. Hence, in greedy diagnosis, whenever a new observable event is perceived, the time required by the online search within the explainer may be longer than the search time in the diagnoser. Moreover, building a complete explainer is out of the question for real size DESs, owing to the exponential explosion of the states involved. The building blocks of the explainer are called *fault spaces*. Each fault space is generated by a variant of the state-elimination algorithm in the literature, whose pseudocode is specified in the paper. In the construction of a complete explainer, the generation of the regular expressions is very expensive since there is a trade-off between the richness of the information embedded in the diagnosis output and the time required for its computation. The paper includes some hints about the asymptotic time complexity of the algorithms. A possible mitigation of the complexity comes from lazy diagnosis, which performs only partial knowledge compilation, thus obtaining a *partial explainer*. The three diagnosis engines have been implemented, with the software code being open source (see [1] for the link). The experimental activity recorded in the paper indicates that only lazy diagnosis may be viable in non-trivial application domains, both for the construction of the partial explainer and for the online computation of the candidates. A partial explainer can be progressively extended, the upgrade being relevant either just to a single observable event (the latest one perceived during monitoring, in case it is not encompassed yet by the current partial explainer) or to the observation pattern corresponding to a behavioral scenario. Typically, behavioral scenarios can be defined to represent frequent and/or critical evolutions of the DES that we want to detect quickly during monitoring via compiled knowledge.

A sample application inspired by a small real-world example in the literature is described in the paper. The aim is to show that in the considered domain of Labour Market, as well as in other similar domains, sequence-oriented diagnosis is actually more convenient than set-oriented diagnosis.

2012 ACM Subject Classification Computing methodologies → Causal reasoning and diagnostics

Keywords and phrases model-based reasoning, monitoring-based diagnosis, discrete-event systems, sequence-oriented diagnosis, active systems, knowledge-compilation, laziness

Digital Object Identifier 10.4230/OASIS.DX.2024.34

Category Extended Abstract

Funding This work was supported in part by the National Natural Science Foundation of China (grant number 61972360).

References

- 1 G. Lamperti, S. Trerotola, M. Zanella, and X. Zhao. Sequence-oriented diagnosis of discrete-event systems. *Journal of Artificial Intelligence Research*, 78:69–141, 2023. doi:10.1613/jair.1.14630.