

Exploiting Shared Resource Dependencies in Spectrum Based Plan Diagnosis

Shekhar Gupta

Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto CA 94304 USA
Email: shekhar.gupta@parc.com

Nico Roos

Masstricht University
The Netherlands
Email: roos@maastrichtuniversity.nl

Cees Witteveen

Department of Software Technology
Delft University of Technology, The Netherlands
Email: c.witteveen@tudelft.nl

Bob Price and Johan de Kleer

Palo Alto Research Center, USA
Email: {bob.price, dekleer}@parc.com

Abstract

In case of a plan failure, plan-repair is a more promising solution than replanning from scratch. The effectiveness of plan-repair depends on knowledge of which plan action failed and why. Therefore, in this paper, we propose an Extended Spectrum Based Diagnosis approach that efficiently pinpoints failed actions. Unlike Model Based Diagnosis (MBD), it does not require the fault models and behavioral descriptions of actions. Our approach first computes the likelihood of an action being faulty and subsequently proposes optimal probe locations to refine the diagnosis. We also exploit knowledge of plan steps that are instances of the same plan operator to optimize the selection of the most informative diagnostic probes. In this paper, we only focus on diagnostic aspect of plan-repair process.

Introduction

Classical planning assumes that the world is deterministic so that every action produces the intended effects. However, this assumption is not true in the real world where actions can fail because of unexpected events. Execution of a plan will lead to an unexpected goal state if one or more actions are behaving abnormally. When such incidents happen one possible way to achieve the desired goal state is repairing the original plan by adding/removing some actions.

In the case of partial observability, the idea of Model Based Diagnosis (MBD) can be further extended to diagnose faulty actions to detect the faulty action instances, in a plan where the plan can be seen as a system and the action can be understood as a component (Roos and Witteveen 2009). While powerful, model-based techniques require accurate fault models which are expensive to develop and in some cases the required data cannot be obtained at all. The proposed Spectrum Based Diagnosis (SBD) approach makes use of abstract frequency statistics to reveal possible causes of a problem without a fault model of the system. SBD has been successfully applied for software fault localization (Abreu et al. 2009) and hardware diagnosis (Arjan

Van Gemund and Abreu 2011). We use SBD to determine the health state of a plan step which infers the health state of corresponding action. In the planning domain, it is common for a single plan operator to be instantiated many times for different plan steps. All plan steps that are instantiated from the same operator will fail if resource associated with the plan operator fails. For instance, every attempt to schedule a shipment on a blocked road will fail. In the online replanning context, we are given the plan ahead of time, so we can exploit knowledge about the operator dependencies of actions within a plan. We propose Extended Spectrum Based Diagnosis which is able to exploit available information about such dependencies in the plan by elegantly extending the spectrum matrix. Diagnosis is further refined by proposing optimal probing locations. The probing strategy uses minimal information criteria to select a probing location.

Extended Spectrum Based Plan Diagnosis

Plan Execution and Monitoring A plan is a tuple $P = \langle \mathcal{O}, S, < \rangle$ where $S \subseteq inst(\mathcal{O})$ is a set of plan steps occurring in \mathcal{O} and $(S, <)$ is a partial order. The partial order relation $<$ specifies an *execution relation* between plan steps: for each $s \in S$ it holds that s is executed immediately after plan step s' such that $s' < s$ have been finished. Figure 1 illustrates a partially ordered plan.

Plan execution is validated by continuously monitoring the goal state. The difference in the observed value $\sigma_{goal}(v)^t$ of any variable v in the goal state from the expected value $\sigma_{goal}(v)$ implies the plan execution failure. Suppose plan shown in Figure 1 represents a multimodal transportation plan where five goods (v_1, \dots, v_5) need to be delivered from initial location to goal location using transportation modes (s_1, \dots, s_8) . In the final destination, it is observed that two goods $(v_2$ and $v_3)$ have not arrived which implies one or more plan steps are faulty. We apply SBD to isolate such faulty plan steps. The basic principle of SBD can be described as follows: if the value of a variable in the goal state is incorrect, then one or more plan steps involved in genera-




Figure 1: Plans and plan steps. Each state characterizes the values of five variables v_1, v_2, v_3, v_4 and v_5 . States are changed by application of plan steps s_i for $i = 1, 2, \dots, 8$. Plan steps having the same color (e.g. s_1 and s_7 , and s_2 and s_5) are instantiations of the same plan operator.

tion of that variable are faulty.

Obtaining Spectrum Matrix The spectrum matrix shows for every variable in σ_{goal} which plan steps are involved from the state σ_0 to σ_{goal} . It records, in the goal state, whether a particular variable v_i has the expected value or not. The spectrum matrix (A, e) , where $A = [a_{ij}]$ is the plan spectrum and e is the error vector can be constructed as follows: The plan spectrum A has N rows (one for each variable) and M columns (one for each plan step). We have $a_{ij} = 1$ if a plan step s_j is involved in the generation of variable v_i in σ_{goal} , else $a_{ij} = 0$. The vector e stores whether the outcome for variable v_i has the expected value ($e_i = +$) or not ($e_i = -$). The spectrum matrix for the partial order plan presented in Figure 1 is given as

	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8	e
v_1	1	0	1	0	0	1	0	0	+
v_2	1	0	1	0	0	1	0	0	-
v_3	1	1	1	1	0	0	1	0	-
v_4	1	1	0	1	1	0	0	1	+
v_5	1	1	0	1	1	0	0	1	+

The above spectrum matrix generates correct diagnosis for systems where components are assumed to be failed independently. However, in the planning domain this assumption is not true. In a plan, several plan steps can be instances of same plan operator, hence they are dependent, we define such plan steps as related plan steps. In our example shown in Figure 1, plan step s_1 and s_7 use the same truck and therefore, they are related in terms of a resource. Similarly, s_2 and s_5 are related. To take such related plan steps into account, the relations are encoded in the matrix A itself. Suppose that plan steps s and s' are related. If s is detected as faulty and $s < s'$, it seems reasonable to consider s' as faulty as well. We assume that failures can occur during plan execution and once a failure occurs, it persists. Formally, we calculate the extended spectrum matrix $A' = [a'_{ij}]$ from A as follows:

$$a'_{ij} = \bigvee_{j' < j, o(j')=o(j)} a'_{ij'} \vee a_{ij} \quad (1)$$

With respect to the plan depicted in Figure 1, the above definition results in the following spectrum matrix (new entries appear in bold face):

	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8	e
v_1	1	0	1	0	0	1	1	0	+
v_2	1	0	1	0	0	1	1	0	-
v_3	1	1	1	1	1	0	1	0	-
v_4	1	1	0	1	1	0	1	1	+
v_5	1	1	0	1	1	0	1	1	+

Probability and Mutual Information Calculation Application of minimal hitting set algorithm on extended matrix A' will generate diagnoses candidates (c_k) $\{c_1 = \langle s_1 \rangle, c_2 = \langle s_3 \rangle, c_3 = \langle s_7 \rangle, c_4 = \langle s_2, s_6 \rangle, c_5 = \langle s_4, s_6 \rangle, c_6 = \langle s_7, s_6 \rangle, c_7 = \langle s_5, s_6 \rangle\}$. Bayes' rule is applied on these candidates to derive a ranking in which all diagnosis are ranked according to their fault probabilities ($Pr(s_i)$). To identify a suitable location for a new probe the mutual information criterion can be used to evaluate and compare measurement choice based on their information contribution (Juan Liu and Zhou 2008), we have adapted same criterion in this paper. If X is a diagnostic state of a plan and Y is the measure value of a variable at a probing location then mutual information is represented as $I(X; Y)$. Table 1 summarizes results obtained for example shown in Figure 1.

s_i	$Pr(s_i)$	$Pr'(s_i)$	$I(X; Y)$	$I'(X; Y)$
s_1	0.200	0.160	0.512884	0.512884
s_2	0.002	0.002	0.008762	0.008762
s_3	0.800	0.762	0.016707	0.016707
s_4	0.002	0.002	0.264348	0.264348
s_5	0.000	0.002	0.004198	0.011041
s_6	0.007	0.008	0.000000	0.000000
s_7	0.003	0.160	0.000000	0.000000
s_8	0.000	0.000	0.074128	0.074128

Table 1: $Pr(s_i)$ and $I(X; Y)$ are derived for original matrix A . $Pr'(s_i)$ and $I'(X; Y)$ are derived for extended matrix A'

Analysis

In the plan described in Figure 1, s_3 has the strongest participation in the unexpected goal state outcomes for variables, v_2 and v_3 . In our results s_3 the highest probability of failure. The standard spectrum A assigns different probabilities to plan steps s_1 and s_7 . The extended spectrum, increased the fault probability of related plan steps s_7 and now s_7 and s_1 have equal probability. Without any ambiguity both the spectrum matrices suggest that s_1 is the most informative location to probe and that s_7 is the least. Since s_7 is in the goal state of the plan therefore no extra information can be gained which matches our mutual information computation. At the same time, extending the matrix reveals the information content at the output of plan step s_5 to the diagnoser. In this case, s_5 is closer to the middle of the plan than s_2 which means that it better splits the hypothesis space about possible causes of failure and therefore is more informative. Therefore, we conclude that the extended spectrum matrix opens up new options to increase the accuracy and decrease the cost of diagnosis in plans with related plan steps.

References

- Abreu, R.; Zoetewij, P.; Golsteijn, R.; and van Gemund, A. 2009. A practical evaluation of spectrum-based fault localization. *Journal of Systems and Software*.
- Arjan Van Gemund, S. G., and Abreu, R. 2011. The antares approach to automatic system diagnosis. In *Proceedings of the 22nd International Workshop on Principles of Diagnosis (DX-2011)*, 5–12.
- Juan Liu, Johan de Kleer, L. K., and Zhou, R. 2008. A unified information criterion for evaluating probe and test selection. In *PHM-2008*.
- Roos, N., and Witteveen, C. 2009. Models and methods for plan diagnosis. *Journal of Autonomous Agents and Multi-Agent Systems* 19(1):30–52.