

A Qualitative Approach to Multiple Fault Isolation in Continuous Systems

Matthew Daigle and Xenofon Koutsoukos and Gautam Biswas

Institute for Software Integrated Systems (ISIS)
Department of Electrical Engineering and Computer Science
Vanderbilt University
Nashville, TN 37235, USA

Abstract

The multiple fault diagnosis problem is important, since the single fault assumption can lead to incorrect or failed diagnoses when multiple faults occur. It is challenging for continuous systems, because faults can mask or compensate each other's effects, and the solution space grows exponentially with the number of possible faults. We present a qualitative approach to multiple fault isolation in dynamic systems based on analysis of fault transient behavior. Our approach uses the observed measurement deviations and their temporal orderings to generate multiple fault hypotheses. The approach has polynomial space requirements and prunes diagnoses, resulting in an efficient online fault isolation scheme.

Introduction

Fault isolation is a key component of safety-critical systems. Quick fault isolation enables timely intervention so catastrophic situations can be avoided. In general, complex systems can fail in many different ways, and the likelihood of multiple faults occurring increases in harsh operating environments. Schemes that do not take into account multiple faults run the risk of generating incorrect diagnoses or even failing to find a diagnosis after faults occur.

We focus on multiple fault diagnosis in continuous systems. Early work (Reiter 1987; de Kleer & Williams 1987) was limited to static systems. We deal with dynamic systems with continuous behaviors. The work in (Gertler 1998) addresses continuous dynamics, but the parity relations approach is hard to apply to multiplicative faults and nonlinear systems. Our approach extends the TRANSCEND framework (Mosterman & Biswas 1999), which employs a qualitative approach for analysis of fault transient behavior. The diagnosis model predicts possible sequences of qualitative measurement deviations due to fault occurrences, and matching the predictions to observations helps isolate faults.

Multiple fault diagnosis is a difficult problem in dynamic systems because interactions among fault effects may obscure the resultant fault signatures. In this paper, we provide a systematic scheme for generation of multiple fault signatures from the single fault signatures. We treat multiple fault effects as the union of single fault effects constrained

by temporal orderings, and therefore, can be computed efficiently online, eliminating the need for precomputing all possible multiple fault effects. We define the notion of n -fault diagnosability, and use this to develop an extension to the TRANSCEND online fault isolation algorithm to find the consistent minimal fault sets of size $\leq n$. The approach requires only polynomial space complexity, and increases runtime efficiency of the multiple fault isolation algorithm by pruning the diagnoses as new measurements deviate.

Related Work

(Reiter 1987; de Kleer & Williams 1987) developed consistency based multiple fault diagnosis methods that were applied mostly to static systems. (Subramanian & Mooney 1996) extended the work to continuous systems with fault modes. The qualitative modeling framework quantizes the state space and specifies qualitative relations between the quantized states, which can result in a large number of states. In contrast, our approach uses a qualitative abstraction of the deviations in observed behavior from nominal behavior. Unlike previous work, our methodology also incorporates temporal information to increase the discriminatory power of the measurements.

(Puig *et al.* 2005) also uses temporal information to improve fault isolation, however, the approach is based on analytical redundancy relations, which are hard to construct for nonlinear systems and multiplicative faults. It is also developed only for single faults. In addition, in our approach the temporal information is generated systematically from the system model.

Sequential testing approaches have been explored in (Shakeri *et al.* 2000), but such strategies can be too slow for time-constrained diagnosis. In contrast, our approach is based on online analysis of measurement deviations, not tests. (Vedam & Venkatasubramanian 1997) presents a signed digraph (SDG) approach that assumes a limit on the number of faults, and considers smaller combinations of faults before larger ones to reduce the computational complexity of the diagnosis algorithm. In continuous systems, the effects of integration and time delays play a role in the observed effects, which most SDGs do not account for. Multiple fault diagnosis has also been investigated using fault propagation graphs (FGPs) (Tu *et al.* 2003). FPGs are at a much higher level of abstraction than our models.

Problem Formulation

For multiple fault diagnosis of continuous systems, our goal is to use system models to find consistent explanations for observed measurement deviations from nominal behavior. Given a system model \mathcal{S} , we define the faults of interest as $F = \{f_1, f_2, \dots, f_{\mathcal{F}}\}$, and the available measurements as $M = \{m_1, m_2, \dots, m_{\mathcal{M}}\}$. The measurements are time-varying signals obtained from the available sensors. One or more measurements deviating from their predicted nominal values at time t_f indicates fault occurrence. The model links fault hypotheses to measurement deviations and predicts fault effects after time t_f . The fault isolation procedure compares these predictions to actual system behavior to produce *candidates*.

Definition 1 (Candidate). *A candidate $c \subseteq F$ is a set of fault hypotheses. The set of all candidates is denoted as $C = \mathcal{P}(F)$, the power set of F .*

For example, the candidate $\{f_1, f_3\}$ implies both f_1 and f_3 have occurred. We wish to find candidates that are consistent with all observed deviating measurements. A candidate is consistent with the observations if its predicted effects (deviations) match the observed dynamic behavior after fault occurrence. A diagnosis is a collection of consistent candidates.

Definition 2 (Diagnosis). *A diagnosis $d \subseteq C$ is the set of candidates consistent with the observations for $t \geq t_f$.*

For example, the diagnosis $\{\{f_1\}, \{f_2, f_3\}\}$ (in shorthand, $\{f_1, f_2 f_3\}$) means that the occurrence of either f_1 or both f_2 and f_3 is consistent with the observations.

In multiple fault diagnosis of continuous systems, establishing candidates can be difficult, because the effects of a group of faults can be combined in different ways. Along with the inherent exponential space of possible diagnoses and the fact that smaller candidates are more likely, we focus only on finding *minimal diagnoses*.

Definition 3 (Minimal Diagnosis). *A diagnosis d is minimal if $(\forall c \in d) \neg (\exists c' \in d) c' \subset c$.*

A desirable property of a system is *diagnosability*. If a system is diagnosable then we should always obtain a unique candidate that is consistent with the observations.

Definition 4 (Diagnosability). *For a given set of faults and measurements, a system is diagnosable if, within finite time after the occurrence of one or more faults, the minimal diagnosis contains a single candidate.*

Since smaller candidates are more likely than larger candidates, and they represent the simplest explanation of the observed effects, we make a practical assumption and limit the maximum candidate size to n , i.e., no more than n faults will occur together in the system. Dropping the n -fault assumption does not limit our method because it is equivalent to setting n to $|F|$. These assumptions lead to the following definition of the multiple fault diagnosis problem.

Problem 1 (Multiple Fault Diagnosis). *Given a system model \mathcal{S} with a set of faults, F , a set of measurements, M , and a candidate size limit n , the multiple fault diagnosis problem is to find the diagnosis d such that d is minimal and $(\forall c \in d) |c| \leq n$.*

Background

Our diagnosis approach extends the TRANSCEND methodology (Mosterman & Biswas 1999), a model-based approach to continuous systems diagnosis, to multiple fault diagnosis. Faults are represented as persistent, abrupt parameter changes in the system, modeled as a bond graph (Karnopp, Margolis, & Rosenberg 2000). When faults occur, they produce transients causing measurements to deviate in time from nominal behavior that is defined by the system model. These deviations are analyzed as they occur to isolate faults in the system. The diagnosis model, the temporal causal graph (TCG), is derived from the system model. It captures the propagation of fault effects on measurements and, therefore, is used to compute predicted effects of faults on measurements. By comparing predicted and observed effects on measurements, we can obtain diagnoses.

Measurement deviations are represented as qualitative \pm values (above, below nominal), and are predicted as *fault signatures* using the TCG (Mosterman & Biswas 1999). A fault signature represents the qualitative value of zeroth-through k th-order derivative changes on a measurement due to a fault occurrence. Because only magnitude and slope can be reliably measured, we condense the signatures to the magnitude change symbol and the first nonzero derivative change, e.g., $00-+-$ becomes $0-$, and $+-+--$ becomes $+-$. We can do this because higher-order changes will eventually manifest as first-order changes, and only the first change on a measurement is useful for diagnosis (Mosterman & Biswas 1999). Therefore, we represent a fault signature for measurement m as an element of the set $\Sigma_m \triangleq \{m^{+-}, m^{-+}, m^{0+}, m^{0-}\}$ ¹. The superscript indicates the observed deviation. The first symbol represents the immediate direction of change (a discontinuity) at fault occurrence and the second symbol represents the slope of the change after fault occurrence.

Definition 5 (Fault Signature). *A fault signature for a fault f and measurement m is the qualitative effect of the occurrence of f on m , and is denoted by $\sigma_{f,m} \in \Sigma_{f,m}$, where $\Sigma_{f,m} \subseteq \Sigma_m$. We denote the set of all fault signatures for fault f as Σ_f .*

Relative measurement orderings define, with respect to a given fault, a partial order of measurement deviations, and are based on the intuition that some measurements deviate before others due to a fault. These are predicted using the TCG based on common temporal subpaths (Daigle, Koutsoukos, & Biswas 2005).

Definition 6 (Relative Measurement Ordering). *Consider a fault f and measurements m_i and m_j . If f manifests in m_i before m_j then we define a relative measurement ordering between m_i and m_j for fault f , denoted as $m_i \prec_f m_j$. We denote the set of all measurement orderings for f as Ω_f .*

Throughout the paper we will illustrate the diagnosis methodology with a circuit example. Fig. 1(a) gives the schematic. The associated bond graph is given in Fig. 1(b). It models the elements of the circuit and the energy exchange

¹In general, $\sigma_{f,m}$ may not be unique if the direction of change cannot be determined by qualitative propagation.

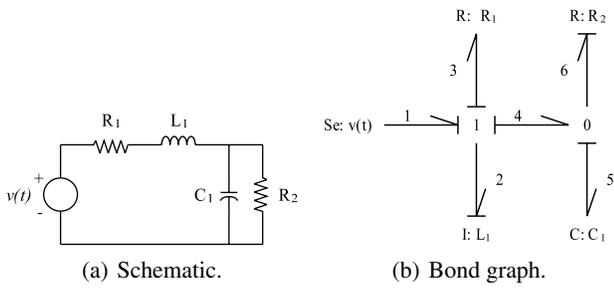


Figure 1: Circuit example.

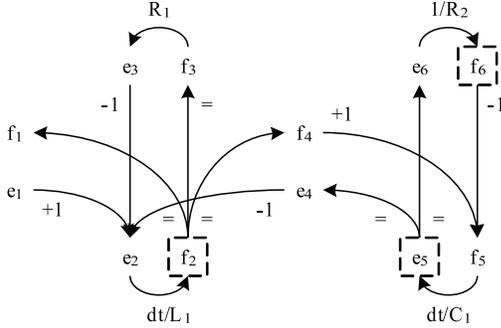


Figure 2: Temporal causal graph for the circuit.

between them (Karnopp, Margolis, & Rosenberg 2000). The derived TCG is given in Fig. 2. Relations between system variables are direct (+1) or inverse (-1) proportionality relations, component parameter values (e.g., R_1), or time-derivative effects (dt). For the circuit, the set of faults is assumed to be $F = \{R_1^-, R_2^+, C_1^+, L_1^-\}$, where the superscript indicates the direction of change of the parameter value. We define the measurement set as the current through L_1 , the voltage across C_1 , and the current through R_2 , or $M = \{f_2, e_5, f_6\}$ in the bond graph model.

The fault signatures and relative measurement orderings for the circuit system are given in Table 1. For example, consider L_1^- . A decrease in L_1 will cause an immediate increase in f_2 , because of the inverse relation implied in the TCG. Since all subsequent paths from f_2 to any other observed variable in the system contain some edge with a dt specifier (implying an integration), then deviations in these measurements will only be detected after f_2 deviates. Either e_5 or f_6 may deviate next. It cannot be determined which will deviate first because the path from e_5 to f_6 contains no integrals. The measurement deviations will not be abrupt because of the integration in the path from L_1 to the measurement, and the direction of change will be opposite that of f_2 because the -1 specifier in the path from f_2 to e_5 and f_6 indicates an inverse proportionality relationship.

Event-Based Fault Modeling

We combine fault signatures and relative measurement orderings into an event-based framework. In this framework, measurement deviations constitute events. The temporal or-

Fault	f_2	e_5	f_6	Measurement Orderings
R_1^-	0+	0+	0+	$f_2 \prec e_5, f_2 \prec f_6$
R_2^+	0-	0+	-+	$e_5 \prec f_2, f_6 \prec f_2, f_6 \prec e_5$
C_1^+	0+	-+	-+	$e_5 \prec f_2, f_6 \prec f_2$
L_1^-	+-	0+	0+	$f_2 \prec e_5, f_2 \prec f_6$

Table 1: Fault signatures and relative measurement orderings for the circuit.

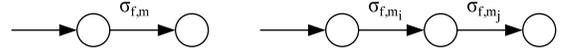


Figure 3: Fault signature LTS representation (left) and relative measurement ordering LTS representation (right).

ders of these deviations are constrained by the relative measurement orderings. For a specific fault, the combination of all fault signatures and relative measurement orderings yields all the possible ways a fault can manifest. We denote each of these possibilities as a *fault trace*.

Definition 7 (Fault Trace). A fault trace for a fault f , denoted by λ_f , is a string of length $\leq |M|$ that includes, for every $m \in M$ that will deviate due to f , a fault signature $\sigma_{f,m}$, such that the sequence of fault signatures satisfies Ω_f .

Consider C_1^+ . $\lambda_{C_1^+} = e_5^- f_6^- f_2^{0+}$ is a valid fault trace, but $\lambda_{C_1^+} = f_2^{0+} e_5^- f_6^-$ is not because the measurement deviation sequence does not satisfy $\Omega_{C_1^+}$. We group the set of all fault traces into a *fault language*, which can be represented concisely by a *labeled transition system (LTS)*.

Definition 8 (Fault Language). The fault language of a fault f , denoted by L_f , is the set of all fault traces for f .

Definition 9 (Labeled Transition System). A labeled transition system is a tuple $\mathcal{L} = (Q, q_0, \Sigma, \delta)$ such that: Q is a set of states, $q_0 \in Q$ is an initial state, Σ is a set of labels, and $\delta \subseteq Q \times \Sigma \times Q$ is a transition relation.

To systematically construct the LTS representation of a fault language, called a *fault model*, we can represent each fault signature and each relative measurement ordering as an LTS, and then compose all the information. Each fault signature $\sigma_{f,m}$ can be represented as an LTS, shown to the left of Fig. 3. It consists of only the single event corresponding to the fault signature². Also, each relative measurement ordering, $m_i \prec_f m_j$, with associated signatures σ_{f,m_i} and σ_{f,m_j} , can be represented as an LTS, shown to the right of Fig. 3. It consists of the two associated signatures in the determined ordering.

The following lemma describes how to construct the fault model for a fault language.

Lemma 1. The fault model of a fault language L_f for fault f , denoted by \mathcal{L}_f , is the synchronous product of the individual LTS for all $\sigma_{f,m} \in \Sigma_f$ and all $m_i \prec_f m_j \in \Omega_f$, where the alphabets for the LTS are taken to be the events contained in the LTS.

²If $\sigma_{m,f}$ is not unique, multiple edges for each possibility are needed going from the first state of the LTS to the final state.

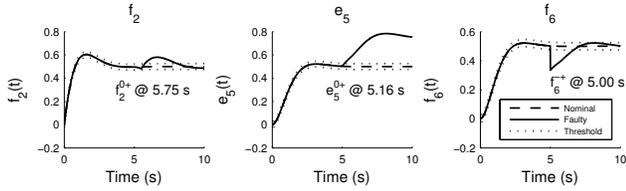


Figure 4: Faulty versus nominal behavior for the circuit measurements. R_2^+ is injected at 5.00 s and R_1^- at 5.50 s.

Proof. Since the synchronous product must obey all individual ordering constraints and includes all measurement deviation events for the fault, it produces all valid measurement deviation sequences and no others. \square

Multiple Fault Diagnosis

Multiple fault diagnosis is more complex than single fault diagnosis due to (i) fault masking and compensation and (ii) relative time of individual fault occurrences. For example, say R_1^- and R_2^+ both occur, and R_2^+ happens before R_1^- , as shown in Fig. 4. This yields the fault trace $f_6^{-+} e_5^{0+} f_2^{0+}$. R_1^- compensates the effect of R_2^+ on f_2 , which is why f_2^{0+} manifests and not f_2^{0-} . However, if R_1^- had occurred later in time, or the selected thresholds were smaller, then the effect on f_2 from R_2^+ may be detected first. That is, $f_6^{-+} e_5^{0+} f_2^{0-}$ is also a valid fault trace for $\{R_2^+ R_1^-\}$ and $\{R_2^+\}$. Which changes are initially observed depends on masking and compensation characteristics of the faults and when they occur relative to one another in the system. The resultant diagnosis can then only be considered as a *best effort* diagnosis, e.g., R_2^+ may occur without observing any direct evidence, so it would not be included in the diagnosis.

To accommodate all possible effects of multiple faults, we derive the combined fault signatures as the union of the single fault signatures, constrained by relative measurement orderings, thus representing all possible physically valid combinations of effects. We assume that when multiple faults occur, they obey the relative measurement orderings of each of the faults. For example, if both R_2^+ and C_1^+ occur, we cannot observe $f_6^{-+} e_5^{-+}$ followed by f_2^{0-} . The reason is that f_2^{0-} is consistent with R_2^+ , but only if preceded by e_5^{0+} . The physical reasoning behind this is that the ordering $e_5^{0+} \prec_{R_2^+} f_2^{0-}$ implies that the fastest way for R_2^+ to affect f_2 is through e_5 . So if the C_1^+ effect reaches e_5 first, it will traverse this same path to f_2 , and cause f_2 to deviate from the effect propagating on this path. Therefore when the R_2^+ effect reaches e_5 , it cannot propagate to f_2 any faster than the C_1^+ effect, so we will not observe its effect on f_2 as the first change on the measurement. By ensuring that we respect single fault measurement orderings when computing multiple fault effects, we will not violate this property.

We also assume that (i) a single fault parameter will change only once and (ii) at most n faults will occur together in the system. By faults occurring together, we mean that the measurement deviations caused by the faults will

be interleaved in time. If all measurements deviate due to one fault before the next fault occurs, diagnosis reduces to the single fault case (Mosterman & Biswas 1999). The first assumption implies a candidate cannot have two different deviations of the same fault parameter. For example, a candidate including both R_1^+ and R_1^- is rejected. The second assumption implies candidates of larger size are less likely than candidates of smaller size. We implement this by assuming the maximum candidate size is n , and this keeps our diagnosis procedure efficient.

We introduce the notion of n -fault diagnosability. Our diagnosis algorithm will always obtain a unique result if the system is n -fault diagnosable and at most n faults occur.

Definition 10 (n -Fault Diagnosability). *For a given set of faults and measurements, a system is n -fault diagnosable if, within finite time after the occurrence of at most n faults, the minimal diagnosis contains a single candidate c such that $|c| \leq n$.*

Diagnoser Design

Our *diagnoser* is designed to trace measurement deviations and output the diagnosis assuming candidate size limit n .

Definition 11 (Diagnoser). *A diagnoser is a tuple $\mathcal{D} = (Q, q_o, \Sigma, \delta, D, Y)$ such that: Q is a set of states, $q_o \in Q$ is an initial state, Σ is a set of labels, $\delta \subseteq Q \times \Sigma \times Q$ is a transition relation, $D \subseteq \mathcal{P}(C)$ is a set of diagnoses, and $Y : Q \rightarrow D$ is a diagnosis map.*

A diagnoser is an LTS extended by a set of diagnoses and a diagnosis map. Similar to a fault model, the labels correspond to measurement deviations. A diagnoser associates each state with a diagnosis, i.e., the set of candidates consistent with the measurement deviations seen thus far.

The diagnoser construction procedure is shown as Algorithm 1. It is described as combining two diagnosers, but can be easily modified to combine k diagnosers simultaneously. Diagnosers are constructed by incrementally composing subdiagnosers, i.e., a diagnoser for a set of faults F_i is composed with a diagnoser for a set of faults F_j to create a new diagnoser for $F_i \cup F_j$. Initially, we begin with diagnosers for singleton fault sets. These are constructed using the individual fault models. For a single fault f , we augment \mathcal{L}_f to form \mathcal{D}_f by constructing the diagnosis map as mapping every state except the initial state to $\{f\}$. The initial state is mapped to the empty diagnosis \emptyset , because until a measurement deviation is observed, we assume the system is operating nominally. The diagnosers corresponding to the individual faults of the circuit are shown in Fig. 5.

The construction algorithm operates by tracing paths in the two given diagnosers. If the same event label is available in both current states, then we advance in both machines, i.e., $(q_1, q_2) \xrightarrow{\sigma} (\delta(q_1, \sigma), \delta(q_2, \sigma))$. Otherwise, we advance in only one, e.g., if σ can only be taken from q_1 , then, $(q_1, q_2) \xrightarrow{\sigma} (\delta(q_1, \sigma), q_2)$. However, if the measurement associated with σ has already deviated along the current path (tracked using H), $\delta((q_1, q_2), \sigma)$ is set to \emptyset , because only the initial change in a measurement is used for isolation. This also occurs if the computed diagnosis for the

Algorithm 1 $\mathcal{D} \leftarrow \text{CreateDiagnoser}(\mathcal{D}_1, \mathcal{D}_2)$

```
 $Q \leftarrow \emptyset, \delta \leftarrow \emptyset, D \leftarrow \emptyset, \Sigma \leftarrow \Sigma_1 \cup \Sigma_2, q_o \leftarrow (q_{o1}, q_{o2}),$   
 $Y(q_o) \leftarrow \emptyset, Q_{\text{pend}} \leftarrow \{q_o\}$   
while  $Q_{\text{pend}} \neq \emptyset$  do  
   $(q_1, q_2) \leftarrow \text{pop}(Q_{\text{pend}})$   
  for all  $\sigma_m \in \Sigma$  do  
    if  $m \notin H((q_1, q_2))$  then  
      if  $\delta_1(q_1, \sigma_m)$  and  $\delta_2(q_2, \sigma_m)$  then  
         $q' \leftarrow (\delta_1(q_1, \sigma_m), \delta_2(q_2, \sigma_m))$   
         $h \leftarrow Y(\delta_1(q_1, \sigma_m)) \cup Y(\delta_2(q_2, \sigma_m))$   
      else if  $\delta_1(q_1, \sigma_m)$  then  
         $q' \leftarrow (\delta_1(q_1, \sigma_m), q_2)$   
         $h \leftarrow Y(\delta_1(q_1, \sigma_m))$   
      else if  $\delta_2(q_2, \sigma_m)$  then  
         $q' \leftarrow (q_1, \delta_2(q_2, \sigma_m))$   
         $h \leftarrow Y(\delta_2(q_2, \sigma_m))$   
      else  
         $q' \leftarrow \emptyset$   
         $h \leftarrow \emptyset$   
      if  $q' \neq \emptyset$  then  
        if  $Y((q_1, q_2)) = \emptyset$  then  
           $d \leftarrow h$   
        else  
           $d \leftarrow \text{AndDiagnoses}(Y((q_1, q_2)), h)$   
        if  $d \neq \emptyset$  then  
           $Q \leftarrow Q \cup \{q'\}$   
           $H(q') \leftarrow H((q_1, q_2)) \cup \{m\}$   
           $\delta((q_1, q_2), \sigma_m) \leftarrow q'$   
           $D \leftarrow D \cup \{d\}$   
           $Y(q') \leftarrow d$   
          if  $q' \notin Q_{\text{pend}}$  then  
             $\text{push}(Q_{\text{pend}}, q')$ 
```

new state, d , is empty, because this means the current sequence of measurement deviations cannot be explained by a candidate of size $\leq n$.

The diagnosis for the new state is formed using Algorithm 2, by combining the current diagnosis with the hypothesis set. The hypothesis set, h , is the set of candidates consistent with the current event. It is formed as the union of the diagnoses of the diagnoser states advanced to via σ . Each candidate of the given diagnosis (the diagnosis of the previous diagnoser state) is augmented with each candidate of the hypothesis set. Essentially, this is an *and* operation. For example, if the previous diagnosis is $\{C_1^+, R_2^+\}$ (meaning that C_1^+ or R_2^+ occurred), and the hypothesis set

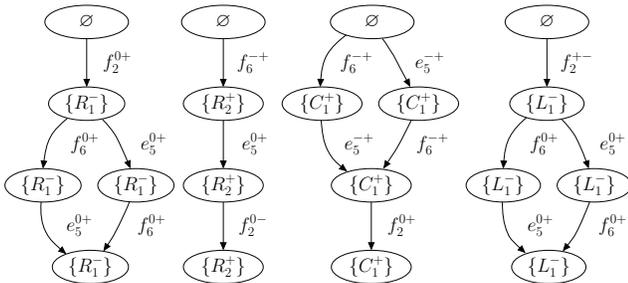


Figure 5: Diagnoser for the individual faults of the circuit.

Algorithm 2 $d' \leftarrow \text{AndDiagnoses}(d_1, d_2)$

```
 $d' \leftarrow \emptyset$   
for all  $c_i \in d_1$  do  
  for all  $c_j \in d_2$  do  
     $c' \leftarrow c_i \cup c_j$   
     $d' \leftarrow d' \cup c'$   
for all  $c \in d'$  do  
  if  $(\exists c' \in d') c' \subset c$  or  $|c| > n$  then  
     $d' \leftarrow d' - \{c\}$ 
```

is $\{L_1^-\}$ (meaning L_1^- occurred), then the new diagnosis is $\{C_1^+ L_1^-, R_2^+ L_1^-\}$ (meaning that either C_1^+ and L_1^- or R_2^+ and L_1^- occurred). The second loop of the procedure prunes the diagnosis by removing supersets and candidates of size above n . Therefore, the resultant diagnosis is guaranteed to be minimal.

The diagnoser for candidate sets of size ≤ 2 and the selected fault set is shown in Fig. 6. It illustrates certain properties of the system. Since not all the leaves have singleton diagnoses, then the system is not double-fault diagnosable. For example, taking the leftmost path, the diagnosis is $\{C_1^+ L_1^-, L_1^- R_2^+\}$. We will at least know that L_1^- has occurred, but will not be able to distinguish whether C_1^+ or R_2^+ is the second fault that occurred. The system is single-fault diagnosable, however, because for any fault trace that can be explained by a single fault, the diagnosis is a singleton.

Online Diagnoser Implementation

Even for a small number of faults, diagnoser size can quickly grow as n increases. For large numbers of faults and measurements, computing the diagnoser for use in online diagnosis is not space-efficient. Alternatively, we could create single diagnosers for each fault, run them simultaneously, and combine the diagnoses. Individual diagnosers may be large, however, if there are few measurement orderings for the fault. To address the space complexity, we instead store only the single fault effects, i.e., for each fault we store its fault signatures and relative measurement orderings (Table 1). As measurement deviations occur, we can check consistency using this stored information to generate hypothesis sets and refine diagnoses.

Given a current diagnosis d_{i-1} , and an event σ_i occurs, we can check which faults are consistent with σ_i . Consistent faults will match the measurement deviation encoded in σ_i and its temporal order with respect to previous deviations, forming the hypothesis set h_i . If $i = 1$, then the new diagnosis d_i is simply h_i , otherwise, the new diagnosis must be consistent with d_{i-1} and with the new information, i.e., $d_i = \text{AndDiagnoses}(d_{i-1}, h_i)$.

Thus, we are only constructing the *path* of the diagnoser corresponding to the particular fault trace observed. This is more space-efficient than using the complete diagnoser for online diagnosis, which, in the worst case, has $O(|M|!)$ fault traces and $O(|M|!)$ states. Storing only fault signatures and relative measurement orderings for single faults, on the other hand, takes $O(|F||M|^2)$ space. Single fault information is composed to obtain multiple fault information, so we do not

