

Adaptive Model-Based Diagnostic Mechanism Using a Hierarchical Model Scheme

Yoichiro Nakakuki Yoshiyuki Koseki Midori Tanaka

C&C Systems Research Laboratories

NEC Corporation

4-1-1 Miyazaki, Miyamae-ku, Kawasaki 216, JAPAN

nakakuki@btl.cl.nec.co.jp

Abstract

This paper describes an adaptive model-based diagnostic mechanism. Although model-based systems are more robust than heuristic-based expert systems, they generally require more computation time. Time consumption can be significantly reduced by using a hierarchical model scheme, which presents views of the device at several different levels of detail. We argue that in order to employ hierarchical models effectively, it is necessary to make economically rational choices concerning the trade-off between the cost of a diagnosis and its precision. The mechanism presented here makes these choices using a *model diagnosability criterion* which estimates how much information could be gained by using a candidate model. It takes into account several important parameters, including the level of diagnosis precision required by the user, the computational resources available, the cost of observations, and the phase of the diagnosis. Experimental results demonstrate the effectiveness of the proposed mechanism.

1 Introduction

Model-based diagnosis is an approach that uses a behavioral specification of a device [Davis 1984; de Kleer and Williams 1987; Genesereth 1984; Reiter 1987]. Although model-based systems are more robust than heuristic-based expert systems, they require more computation time. In general, the computational complexity of model-based diagnosis grows rapidly with the complexity of the device model. This paper proposes an efficient diagnostic mechanism using a hierarchical model scheme.

Other researchers have investigated several approaches for model-based device diagnosis. One successful approach is to use probabilistic information, e.g., the minimum entropy technique (GDE [de Kleer and Williams 1987], Sherlock [de Kleer and Williams 1989]) or the focusing technique [de Kleer 1991]. However, in order to deal with large scale problems, it is important to use not only those techniques but also a

hierarchical model scheme (XDE [Hamscher 1990]). A hierarchical model scheme can reason about the target device at multiple levels of abstraction: early in the diagnosis, an abstract level model can be used to eliminate parts of the device from consideration, while later a more detailed model can be used. Since diagnostic computation at more detailed levels is generally more complex and expensive, the selection of an appropriate level involves making trade-offs between diagnosis cost and diagnosis precision. In order to solve this problem, XDE uses a simple heuristic algorithm that tries to keep the level of the model as high as possible. This may not always be the most efficient strategy, as shown in the empirical comparison with this paper's more adaptive mechanism.

Consider as an example the problem of diagnosing an electronic device composed of several boards, each of which is composed of several chips. Sometimes a field service engineer may only want to know which board to replace, while at other times the faulty chip must be pinpointed. Diagnostic systems should be flexible enough to adapt to the required diagnosis precision.

Diagnostic systems should also minimize the total diagnosis cost, which we measure here in terms of time as the sum of the observation cost and the computation cost. The observation cost depends on the instruments being used. For example, the manual method of using a logic analyzer to capture a digital signal from a device is expensive, whereas an electron-beam tester can easily observe a signal anywhere within an LSI chip. In the manual case the number of observations taken will greatly affect the total diagnosis cost, while in the latter case the total cost will mainly be determined by the computation cost. Thus a diagnostic system should respond to both the observation cost and the computation cost.

This paper presents a diagnosis mechanism that takes into account four parameters: the phase of diagnosis, the computational environment, the cost of observing the target device, and the required diagnosis precision. Section 2 gives an illustrative example of a hierarchical model, and shows how information gain can be measured at various levels. Section 3 pro-

poses a model diagnosability criterion for estimating how much information could be gained from the various models. Section 4 proposes an adaptive diagnosis algorithm based on that criterion. Section 5 concludes with suggestions for several generalizations and extensions of the algorithm.

2 Diagnosis with Hierarchical Models

In most conventional hierarchical model-based approaches [Davis 1984; Hamscher 1990], the structure of a device is represented both as a physical hierarchy and as a logical (functional) hierarchy. The required diagnosis precision is usually represented as a level in the physical hierarchy. Taking the example of the device shown in Fig. 2-1, if the required diagnosis precision is the chip level, then a diagnostic system tries to find the faulty chip(s) among the three chips C_1 , C_2 , and C_3 .




Fig. 2-1 Hierarchical structure

In general, physical hierarchies and logical hierarchies have different structures [Davis 1984]. To simplify discussion, this paper assumes that they have the same structure, and also assumes that there is only a single fault in the target device. However, the proposed techniques can easily be extended to remove these assumptions.

Consider the hierarchical model scheme shown in Fig. 2-2. A full adder (a) is composed of five subcomponents, and an 8-bit ripple carry adder (b) comprises eight full adders.




Fig. 2-2 Hierarchical Model Scheme

There are 256 distinct models for an 8-bit ripple carry adder; of the three shown in Fig. 2-3, model X is the most abstract, model Z is the most detailed, and model Y lies between the two. In general, diagnosis from a more detailed model is more expensive, but it is also more specific. The selection of an appropriate level for a given diagnostic situation should take into account how much information can be gained at each different level of model detail. The next section presents a method of estimating this information gain.




Fig. 2-3 Example

2.1 Diagnosis Precision and Entropy

Several existing systems [de Kleer and Williams 1987; de Kleer and Williams 1989; Hamscher 1990; Koseki *et al.* 1990] use the entropy of a set of suspected components to estimate the expected information needed to complete a diagnosis. However, the expected information generally depends on the required diagnosis precision, as illustrated by the faulty 2-bit ripple carry adder shown in Fig. 2-4. Changes of diagnostic status are shown at two different levels: (a) the function-level, and (b) the gate-level. The suspected components are shown hatched; white components are no longer suspected. The fault probability $P(C)$ for each suspected component C is also shown. The figure gives the initial diagnostic status, and the status after each of two different sets of observations, A and B.

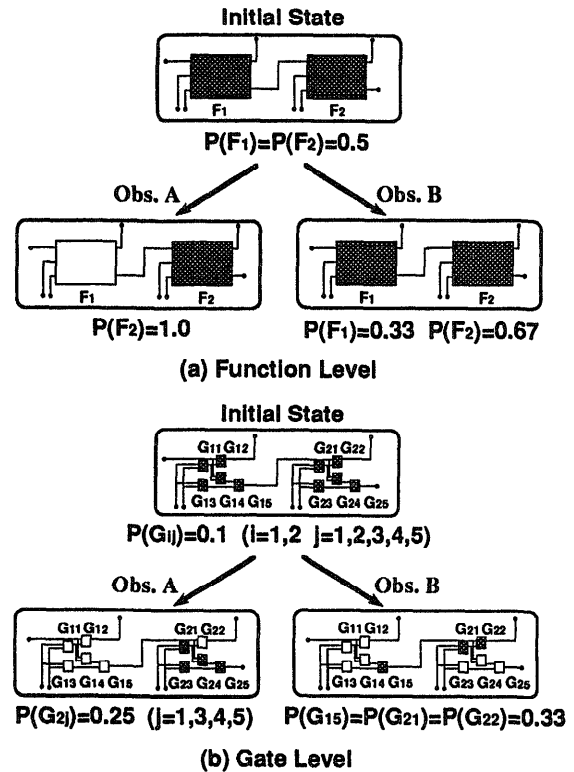


Fig. 2-4 Changes of Diagnostic Status

Fig. 2-4(a) illustrates the case where the required precision is the function-level. If we get observation A, showing that F_1 is normal, then the diagnosis is over. But even if we get observation B, both F_1 and F_2 are still suspected. At this level, observation A seems to be more informative than observation B. But in Fig. 2-4(b), where the required precision is the gate-level, observation A reduces the number of suspected components to four, whereas observation B reduces it to three. At this level observation B is more informative than A. This contrast illustrates the importance of taking into account the desired diagnosis precision when measuring the information gain.

In order to measure the information gain according to the given precision, we calculate the entropy for each level in the physical hierarchy. For instance, in the above example the entropy for the function-level (E_F) and for the gate-level (E_G) are defined as follows, and are expressed in terms of bits.

$$E_F = - \sum_i P(F_i) \log P(F_i)$$

$$E_G = - \sum_i \sum_j P(G_{ij}) \log P(G_{ij})$$

Fig. 2-5 summarizes the reduction of entropy achieved by the observations, in both levels of the above example.




Fig. 2-5 Changes of Entropy

The entropy at a given level is regarded as the remaining information required to complete a diagnosis at that level of precision; when it has been reduced to zero no further information from observations is required. The algorithm proposed here chooses the appropriate figure for entropy according to the level of precision required. For example, if the required precision is the function-level, it tries to reduce E_F , whereas if gate-level precision is required, it will use the gate-level entropy, E_G .

3 Model Diagnosability Criterion

This section introduces a model diagnosability criterion that provides an estimate (in terms of entropy) of the most detailed diagnosis that is achievable using a given model.

Consider the three models for a 2-bit ripple carry adder shown in Fig. 3-1. Assume that the required diagnosis precision is the gate-level and that each of the ten gates, G_{ij} ($i = 1, 2; j = 1, 2, 3, 4, 5$), has the same fault probability of 0.1. (Thus under the single fault assumption the probability failure for each function-level component, F_i , is 0.5.) This section works through the calculation of the minimal entropy achievable by models A, B and C.

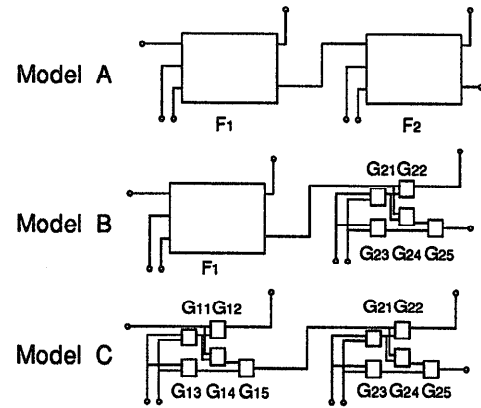


Fig. 3-1 Models for a 2-bit Adder

Using model A, if enough observations are given, it is possible to find a faulty function (F_1 or F_2), but the faulty gate in the function can never be pinpointed. If, for instance, the faulty component is G_{11} , the system can only conclude that the faulty component is $G_{11}, G_{12}, G_{13}, G_{14}$, or G_{15} with probability 0.2 each. With model A the gate-level entropy can never be reduced below $5 \cdot (-0.2 \log 0.2) = 2.32$. No matter how many further observations are given, the system still cannot obtain the additional 2.32 bits information needed.

Using model B the system can (given enough observations) find the faulty gate provided it is one of the five gates, G_{2j} ($j = 1, 2, 3, 4, 5$). If not, the entropy can not be reduced to less than $5 \cdot (-0.2 \log 0.2) = 2.32$. Thus the expected lower bound for the entropy reduction is:

$$0.5 \cdot 2.32 + 0.5 \cdot 0 = 1.16$$

Using model C, the faulty gate can always be found (given enough observations), so the expected lower bound for the entropy reduction is 0.

As an estimate of the completeness of the diagnosis achievable by a model, we define the model diagnosability $D(M)$ for a model M . The maximum value of 1.0 indicates that complete diagnosis always achievable.

$$D(M) = \frac{E_{cur} - E_{low}}{E_{cur}}$$

Here, E_{cur} is the current entropy and E_{low} is the expected lower bound for the entropy. The current entropy expresses the expected information needed to

complete a diagnosis. The numerator ($E_{cur} - E_{low}$) indicates how much information is expected to be gained by using model M .

In the example above, current entropy is $10 \cdot (-0.1 \log 0.1) = 3.32$ at the initial stage of a diagnosis (Fig. 3-2(a)). Therefore, the $D(M)$ for each model is calculated as follows:

$$D(\text{model A}) = \frac{3.32 - 2.32}{3.32} = 0.30$$

$$D(\text{model B}) = \frac{3.32 - 1.16}{3.32} = 0.65$$

$$D(\text{model C}) = \frac{3.32 - 0.00}{3.32} = 1.00$$

Fig. 3-2(b) summarizes these results. It shows that a diagnosis with model A can gain at most 30% of necessary information, but model C is powerful enough to gain all the necessary information.

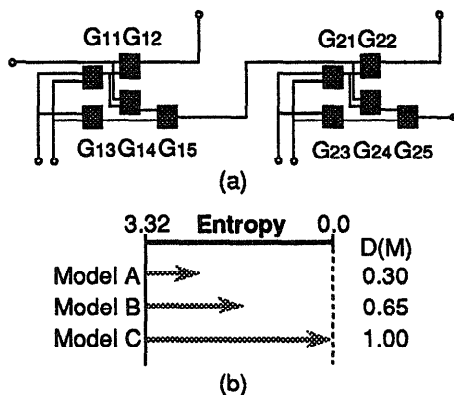


Fig. 3-2 Initial Stage

Next assume that the set of suspected components has been narrowed down by some observations to those hatched in Fig. 3-3(a). Then the values for $D(M)$ change to those shown in Fig. 3-3(b). Now no information can be gained if model A is used, but model B and C have the ability to gain the all information needed to pinpoint a faulty gate.




Fig. 3-3 Later Stage

4 Adaptive Diagnosis Mechanism

The previous section introduced the model diagnosability criterion; this section presents an adaptive diagnostic algorithm called HIMA that uses that criterion to select an appropriate model at each stage of a diagnosis.

Let $D(M)$ be the diagnosability for model M , and let C be the average cost of an observation (in terms of the time required to make it). The diagnostic process consists of several observation/computation cycles (also called phases), so $T(M) + C$ is the expected cost for a cycle, where $T(M)$ is the expected time to calculate the suspects (given an observation) under model M . We assume that $T(M)$ can be estimated empirically or analytically, and that C is a model-independent constant. To choose an appropriate model, we evaluate each model by using the following criterion:

$$E(M) = \frac{D(M)}{T(M) + C}$$

At each diagnostic cycle the model with the greatest value for $E(M)$ is selected as the best one.

This diagnostic mechanism adapts its choice of level according to several factors: the phase of the diagnosis, the given diagnosis precision, and the costs of observation and computation.

The remainder of this section is a worked example illustrating the algorithm's behavior under two different economic situations: first where the cost of observations is very low relative to computation time, and then when it is relatively high.

Returning to the three models of 8-bit ripple carry adder of Fig. 2-3, assume that the required diagnosis precision is the gate-level and that the expected computation time for each model is as follows. (These values are derived empirically using our diagnostic engine [Koseki *et al.* 1990].)

$$T(\text{model X}) = 0.30 \text{ (sec)}$$

$$T(\text{model Y}) = 0.39 \text{ (sec)}$$

$$T(\text{model Z}) = 1.22 \text{ (sec)}$$

First consider the case where the cost of observations is relatively low, i.e. $C \ll T(M)$ for each model M . Then $E(M)$ can be approximated as:

$$E(M) \approx D(M)/T(M).$$




Fig. 4-1 Examples of Diagnostic Stages

Initial stage Fig. 4-1(a) shows the initial stage of the diagnosis, when all gates are suspected with the same probability (1/40). At this stage model X is selected because it has the largest value for $E(M)$, as shown in Table 4-1(a). Model X can gain about 56% of the necessary information at a low cost.

Later stage Suppose that after some observations, the diagnosis has proceeded to the state shown in Fig. 4-1(b). The values for $E(M)$ in Table 4-1(b) show that model Y should now be selected. Model X can gain at most only 30% of necessary information, and model Z is relatively costly.

Table 4-1 $E(M)$ for each model ($C \ll T(M)$)

(a) Initial Stage			
Model	Diagnosability	Cost	$E(M)$
X	0.56	0.30	1.86
Y	0.62	0.39	1.58
Z	1.00	1.22	0.82

(b) Later Stage			
Model	Diagnosability	Cost	$E(M)$
X	0.30	0.30	1.00
Y	0.65	0.39	3.42
Z	1.00	1.22	0.82

This example shows the HIMA algorithm's ability to change the level of the model appropriately at each stage.

Now consider the opposite case, where the cost of observations is high relative to the computation cost, i.e. $T(M) \ll C$ for each model M . In this case the diagnosis cost is barely affected by the observation cost C . For example, if $C = 100.0$ sec in the above example, then the diagnosis cost (the sum of the observation cost and the computation cost) for models X, Y and Z are 100.30, 100.39 and 101.22, respectively. Table 4-2 shows that model Z will be selected in both stages of

the diagnosis. Intuitively this shows that when observations are expensive, it is worth maintaining a very detailed model at all times, whereas if observations are cheap, this detail is needed only in the later stages.

Table 4-2 $E(M)$ for each model ($T(M) \ll C$)

(a) Initial Stage			
Model	Diagnosability	Cost	$E(M)$
X	0.56	100.3	0.0056
Y	0.62	100.4	0.0061
Z	1.00	101.2	0.0099

(b) Later Stage			
Model	Diagnosability	Cost	$E(M)$
X	0.30	100.3	0.0030
Y	0.65	100.4	0.0065
Z	1.00	101.2	0.0099

This contrast illustrates the HIMA algorithm's ability to adapt to the dynamic economics of the observation and computation. Computation cost obviously depends on the computing machinery available; a 1-MIPS computer will require far more time than a 100-MIPS machine, so it is important that this factor is specified as an input to the algorithm.

Finally, suppose that the required diagnosis precision were changed from the gate level to the function-level in each of the two examples above. All of the three models have enough diagnosability, so model X would be selected because it has the least expected cost among the three. This illustrates the adaptability of the HIMA algorithm to the required diagnosis precision.

4.1 Algorithm

In general, there are huge number of possible models for a given hierarchical model scheme and target device, and this number grows exponentially with the number of components in the device. Since it is impracticable to test all possible models, we aim to reduce the number of models considered. The following pseudocode for the HIMA algorithm introduces a heuristic search mechanism to achieve this end.

```

M ← Most abstract level model;
while {there is an expandable component  $c_i$  in M
and  $E(\text{expand}(M, c_i)) > E(M)$ }
M ←  $\text{expand}(M, c_i)$ ;

```

Expandable components are those that have models at a more detailed level, and $\text{expand}(M, c_i)$ is a model obtained from model M by replacing component c_i with the components that comprise c_i at the next level below. For example, in Fig. 2-3,

model $Y = \text{expand}(\text{Model } X, f_7)$ (f_7 is the seventh full-adder from the left). No component of model Z is expandable. The algorithm requires time linearly proportional to the number of components in the model scheme.

4.2 Experimental Results

We evaluated the performance of the HIMA algorithm on a 16-bit adder represented by a three-level hierarchical model scheme. The number of components at successive levels was 2, 32 and 160. The required diagnosis precision was set to the most detailed level.

The performance of the HIMA algorithm was compared to the two obvious "strawman" algorithms that could be used in two extremely different diagnostic environments. Results show that the HIMA algorithm outperforms both strawman algorithms, even under the conditions most favorable to each.

The first strawman, FIX, uses a fixed model throughout the diagnostic process, determined by the required precision (so in this case, the most detailed level model is always used). The other strawman algorithm, AHAP, keeps the level of the model to use as high as possible: it changes to a more detailed level only if there is no possibility of gaining information using the current model.

Single faults were generated randomly, and the average cost (the sum of computation cost and observation cost) of pinpointing the faulty component was measured. The experiments were performed with three different expected observation costs (1msec, 1sec and 100sec). The results shown in Table 4-3 show that the HIMA algorithm performs best in all cases.

Table 4-3 Average Diagnosis Cost (sec)

Algorithm	Expected Observation Cost		
	1 msec	1 sec	100 sec
FIX	4.0	11.3	734.0
AHAP	1.3	21.7	2041.3
HIMA	0.8	8.1	730.8

5 Discussion

The technique described in this paper adapts to several factors: the required precision, the given computation power, the observation cost, and the phase of diagnosis. Although some simplifying assumptions were made to the diagnosis problem, the proposed mechanism can naturally be extended to more general cases, which have natural justifications in the real-world diagnosis tasks. First, the diagnosis precision need not be restricted to a fixed level in the physical hierarchy. For example, according to the availability of spare parts, chip-level precision may be required for some parts of

the target device, and board-level precision may be required for others. Second, the physical hierarchy and the logical hierarchy need not be identical. Third, whereas we assumed that the observation cost is model independent, this need not be the case: the output signal of a whole board may be cheaper to observe than the output signal of an intermediate chip.

The HIMA algorithm can be extended in several other ways. First, in some domains it may be preferable to modify the model diagnosability criterion, because it does not estimate the number of required observations. For example, even if $D(M) = 0.9$ for a certain model M , a diagnosis with the model may require dozens of observations to gain this 90% of the necessary information, so the criterion does not always estimate the diagnosability exactly. Second, the algorithm requires estimates of the computation cost and the fault probability for each component. Inductive learning techniques [Nakakuki *et al.* 1990] or analytical methods can provide this.

Acknowledgment

This research has been carried out as a part of the Fifth Generation Computer Project. The authors would like to thank Katsumi Nitta of the Institute for New Generation Computer Technology for his support. The authors also would like to thank Masahiro Yamamoto and Takeshi Yoshimura for their encouragement in this work, and Jason Catlett for his valuable suggestions.

References

- Davis, R. 1984. Diagnostic reasoning based on structure and behavior. *Artificial Intelligence* 24:347-410.
- de Kleer, J. and Williams, B. C. 1987. Diagnosing multiple faults. *Artificial Intelligence* 32:97-130.
- de Kleer, J. and Williams, B. C. 1989. Diagnosis with behavioral modes. *Proc. IJCAI-89* 2:1324-1330.
- de Kleer, J. 1991. Focusing on probable diagnosis. *Proc. AAAI-91* 2:842-848.
- Genesereth, M. R. 1984. The use of design descriptions in automated diagnosis. *Artificial Intelligence* 24:411-436.
- Hamscher, W. 1990. XDE: Diagnosing devices with hierarchic structure and known component failure modes. *Proc. CAIA-90* 1:48-54.
- Koseki, Y.; Nakakuki, Y.; and Tanaka, M. 1990. An adaptive model-Based diagnostic system. *Proc. PRICAI'90* 1:104-109.
- Nakakuki, Y.; Koseki, Y.; and Tanaka, M. 1990. Inductive learning in probabilistic domain. *Proc. AAAI-90* 2:809-814.
- Reiter, R. 1987. A theory of diagnosis from first principles. *Artificial Intelligence* 32:57-95.