

Continuously estimating persistent and intermittent failure probabilities.

Johan de Kleer, Lukas Kuhn, Juan Liu, Bob Price, Minh Do,
Rong Zhou

*Palo Alto Research Center, Palo Alto, CA 94304 USA,
{dekleer,lukas.kuhn,Juan.Liu,bprice,minh.do,rong.zhou}@parc.com*

Abstract:

Almost all previous work on model-based diagnosis has focused on persistent faults. However, some of the most difficult to diagnose faults are intermittent. It is very difficult to isolate intermittent faults which occur with low frequency but yet at high enough frequency to be unacceptable. For example, a printer which prints one blank page out of a 1000 or a computer that spontaneously reboots once per day is unacceptable. Accurate assessment of intermittent failure probabilities is critical to diagnosing and repairing equipment. This paper presents an overall framework for estimating component failure probabilities which includes both persistent and intermittent faults. These estimates are constantly updated while the equipment is running.

1. INTRODUCTION

Most work on model-based diagnosis addresses isolating single persistent faults in physical systems where only information (voltages, currents, pressures, etc.) are conveyed among system components. This paper extends model-based diagnosis to include intermittent faults and to physical machines where material is being transferred from one system component to another. Thus we extend model-based diagnosis to the very difficult diagnostic task of troubleshooting manufacturing lines and plants.

In this paper we draw our examples from printers which should be considered as a manufacturing line which runs continuously and changes paper from one state (blank) to another state (marked on, stapled, bound, etc.). Extending our group's work on developing self-aware printing machines Fromherz et al. (2003), we have designed and built the modular redundant printing engine illustrated in Figure 1. Such high-end reprographic machines operate more-or-less continuously providing a constant stream of observations and exception conditions. This paper addresses the challenge of estimating the probabilities of module failures from this data stream. These estimates are critical to avoiding modules which may fail (prognostics) as well as for sequential diagnosis.

To efficiently update the probabilities we need an approach which arrives at the correct posterior probabilities with minimal space and time. Recording all past itineraries and all possible hypothetical diagnoses is impractical, yet we do not want to throw away any valuable information. This paper describes a procedure based on maintaining a limited set of counters instead of the entire history without losing any information about past itineraries.

Figure 2 illustrates the basic software architecture of our system. The main task of the planner is to schedule sheets through the printer. The main task of the diagnoser is to estimate module failure probabilities and provide diagnostic guidance to the planner. Both the planner and




Fig. 1. Model of PARC's prototype highly redundant printer. It consists of two towers each containing 2 internal printers (large rectangles). Sheets enter on the left and exit on the right. Dark black edges with small rollers represent a bidirectional paper path. There are three main paper (horizontal) highways within the modular printer. The printer incorporates 2 types of media handling modules represented by small lighter edge rectangles. The motivation for this design is to continue printing even if some of the printers fail or some of the paper handling modules fail or jam.

diagnoser operate with a common model of the machine state.

The reprographic machines receive a continuous stream of print jobs. Each print job consists of a sequence of sheets of paper. The planner constructs an optimal itinerary for each sheet of paper which specifies the full trajectory each sheet travels through the machine. These plans can consist of dozens of modules. Failure is detected in two ways. First, a sheet arrives at a module while it is still handling a previous sheet. This will be detected by the module sensors and the module will immediately stop moving the paper

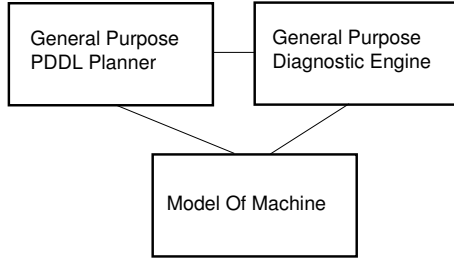


Fig. 2. Basic architecture.

(manifesting as a “jam”). Second, the system (Figure 1) has a scanner on the output so it can detect if the sheet has been damaged in any way.

Common kinds of failures are:

- A dog ear at one of the corners.
- Scuff marks on the paper caused by rollers (called nips) gripping the paper too tightly.
- The leading edge of the paper as it moves through the system may encounter a protrusion. (Leading edge damage.)
- Paper is crumpled or shredded inside the machine.

These systems have some striking differences from the commonly explored digital circuits analyzed in most of the model-based diagnosis literature:

- Most errors cannot be masked or cancelled out. A damaged sheet cannot be repaired by the machine.
- The sheet may be touched by the same module more than once.

We notate an itinerary and its outcome by the sequence of modules touched by the paper followed by Fail or Success. For example the itinerary in which a sheet passes through modules A,B,C,D,E,B,C and moved to the output tray without damage is represented as (A,B,C,D,E,B,C,Success). The itinerary in which a sheet passes through modules A and B and then jams in C is represented as (A,B,C, Fail).

2. OUTLINE AND ASSUMPTIONS

In this paper we provide solutions for all combinations of intermittent and persistent faults. Each itinerary consists of a sequence of modules. We adopt the counting convention from Abreu et al. (2006) and associate two counters with each module m :

- $c_{f,m}$: number of plans where m was used and failed.
- $c_{s,m}$: number of plans where m was used and succeeded.

The following simplifying assumptions apply for our reprographic engines:

- Every faulty module output is observable. (Catastrophic fault assumption.) Any damage to a piece of paper cannot be rectified by later modules. This assumption does not hold for digital systems where internal faulty signals can be masked to produce correct outputs. Our approach still applies for such systems but requires more reasoning to determine whether a faulty output is masked. (See de Kleer (2007).)

	o	Fail	Success
u			
Used		1	0
Not Used		0	1

$$p(O|M = m, U)$$

Fig. 3. Summary of the observation function in the single fault persistent case. Note that when diagnoses can have multiple faults, the test for whether a diagnosis is used generalizes to whether any of its models are used in the current itinerary.

- Fault probabilities are stationary. Our approach can be easily extended to accommodate slowly drifting probabilities through discounting.
- Paper cannot damage a module. Most applications of model-based diagnosis presume signals cannot damage the system. However this does not hold for production lines which transport heavy objects as a misrouted object could damage the machine itself. Fortunately, in reprographic machines the relatively fragile paper is always what gets damaged.
- Observations do not affect machine behavior. This assumption is made in most approaches to diagnosis.
- All faults are distinguishable. This is simply for exposition: as in digital circuits, indistinguishable faults are collapsed.

These assumptions hold in a broad range of systems. The only input our approach requires is the sequence of itinerary-outcome pairs where the itinerary is expressed as a set of modules. For example, printers, manufacturing plants, bottling plants, and packaging plants can exploit our approach.

3. SINGLE PERSISTENT FAULT

This case follows from GDE de Kleer and Williams (1987). Let $p(M)$ be the probability module M is faulted. Let U be whether the module was used in the plan that produced the observation O . The sequential Bayesian filter Berger (1995) is:

$$p_t(M|O, U) = \alpha p(O|M, U) p_{t-1}(M). \quad (1)$$

Where α is chosen so that the posterior probabilities sum to 1 (presuming we start with the knowledge there is a fault).

$p(O|M, U)$ is 1 in situations where $c_{f,m}$ are incremented, otherwise it is 0. Namely, if the module is not used in a failing itinerary it is exonerated by the single fault assumption, and if the module is used in a successful plan it is exonerated because we assume that every faulty output is observed. Figure 3 illustrates the possibilities.

Assume that at $t = 0$ all modules fail with prior probability $p_0 = 10^{-10}$. Consider the arrangement of mod-




Fig. 4. 6 module fragment with possible interaction paths.

ules in Figure 4. Consider the sequence of itineraries: (A,B,C,D,E,F,Fail), (A,B,C,Success), (E,F,Success). After the (A,B,C,D,E,F,Fail) itinerary, one of the 6 modules must be faulted. As the priors are all equal, each module must be faulted with probability $\frac{1}{6}$. As we assume faults are persistent and all faults are manifest, a successful itinerary exonerates all the modules of the itinerary. Thus the itinerary (A,B,C,Success) indicates that A,B and C are all working correctly. Finally, the itinerary (E,F,Success) exonerates modules E and F. Therefore, D is faulted with probability 1 (see Table 1).

Table 1. The resulting posterior probabilities $p(M = m|O, U)$ over one sequence of itineraries. One persistent fault.

t	$m = A$	$m = B$	$m = C$	$m = D$	$m = E$	$m = F$
0	10^{-10}	10^{-10}	10^{-10}	10^{-10}	10^{-10}	10^{-10}
1	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$	$\frac{1}{6}$
2	0	0	0	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$
3	0	0	0	1	0	0

4. SINGLE INTERMITTENT FAULT

This case extends the model for intermittent faults presented in de Kleer (2007) which was informed by Koren and Kohavi (1977). In the case of intermittent faults, $p(O|M, U)$ is still 0 in case m_{01} and 1 in case $c_{f,m}$. Otherwise, we need to estimate $p(O|M, U)$ using the counts. The probability that module m produces an incorrect output if faulted is calculated as follows:

$$\frac{c_{f,m}}{c_{f,m} + c_{s,m}}.$$

(The denominator can never be 0 as will be described later.) Let $p_0(M)$ be the prior probability that m is faulted. Given a particular observation O , Bayes rule gives:

$$p_1(M|O, U) = \alpha p(O|M, U) p_0(M).$$

The observation function $P(O|M, U)$, is estimated from the counts. If the observation is a Failure and m is used, then:

$$p(\text{Fail}|M = m, U) = \frac{c_{f,m}}{c_{f,m} + c_{s,m}},$$

and if is Success and m is used, then:

$$p(\text{Success}|M = m, U) = \frac{c_{s,m}}{c_{f,m} + c_{s,m}},$$

otherwise as m cannot affect o , if m good,

$$p(\text{Success}|M = m, U) = 1,$$

otherwise,

$$p(\text{Fail}|M = m, U) = 0,$$

(captures the single fault assumption). Figure 5 summarizes the 4 possibilities.

$i \backslash o$	Fail	Success
Used	$\frac{m_{11}}{m_{11} + m_{10}}$	$\frac{m_{10}}{m_{11} + m_{10}}$
Not Used	0	1

$$p(O|M = m, U)$$

Fig. 5. Summary of the observation function in the single fault intermittent case.

After many iterations of Bayes rule, intuitively,

$$p_t(M|O, U) = \alpha p(\text{good})^g p(\text{bad})^b p_0(M),$$

where there are g observations of m -used good behavior and b observations of m -used bad behavior. Formally:

$$p_t(M|O, U) = \begin{cases} 0 & \text{if } \exists U \in \text{Us.t. } U\text{Fail} \wedge m \notin U \\ \alpha w p_0(M) & \text{otherwise} \end{cases} \quad (2)$$

where,

$$w = \left[\frac{c_{s,m}}{c_{f,m} + c_{s,m}} \right]^{c_{s,m}} \left[\frac{c_{f,m}}{c_{f,m} + c_{s,m}} \right]^{c_{f,m}}. \quad (3)$$

Consider again the arrangement of modules in Figure 4 and 3 itineraries: (A,B,C,D,E,F,Fail), (A,B,C,Success), (E,F,Success). The probabilities are updated as follows: After the first observation all $c_{f,m}$ counters are 1 and the $c_{s,m}$ 0, therefore w 's are 1. After observing (A,B,C,Success) the counters ($c_{s,m}, c_{f,m}$) for $\{A, B, C\}$ are all 1, 1 and the counters for $\{D, E, F\}$ are all 0, 1. Therefore, $w = \frac{1}{4}$ for $\{A, B, C\}$ and 1 for the rest. We observe (E,F,Success) next. The counters for $\{A, B, C\}$ are all 1, 1. The counters for D are 0, 1 and the counters for $\{E, F\}$ are: 1, 1. Now suppose itinerary (A,B,C,D,E,F,Success) repeats for 7 iterations. Table 2 illustrates how the posterior probabilities evolve.

Table 2. The resulting posterior probabilities $p(M = m|O, U)$ over one sequence of itineraries. One intermittent fault.

t	$m = A$	$m = B$	$m = C$	$m = D$	$m = E$	$m = F$
0	10^{-10}	10^{-10}	10^{-10}	10^{-10}	10^{-10}	10^{-10}
1	$\frac{1}{6}$.17	$\frac{1}{6}$.17	$\frac{1}{6}$.17	$\frac{1}{6}$.17	$\frac{1}{6}$.17	$\frac{1}{6}$.17
2	$\frac{1}{15}$.07	$\frac{1}{15}$.07	$\frac{1}{15}$.07	$\frac{4}{15}$.27	$\frac{4}{15}$.27	$\frac{4}{15}$.27
3	$\frac{1}{9}$.11	$\frac{1}{9}$.11	$\frac{1}{9}$.11	$\frac{4}{9}$.44	$\frac{1}{9}$.11	$\frac{1}{9}$.11
4	$\frac{16}{107}$.15	$\frac{16}{107}$.15	$\frac{16}{107}$.15	$\frac{27}{107}$.25	$\frac{16}{107}$.15	$\frac{16}{107}$.15
10	.16	.16	.16	.18	.16	.16

Working with the same 6 modules, consider a slightly more realistic example. Assume that the prior probabilities of intermittent failures are equal for all the modules. Consider the case in which module D is intermittently faulted and damages one out of every 1001 sheets (starting with sheet 1001). Suppose that the printer repeatedly executes the itineraries: (A,B,E,F), (C,B,E,D) (A,B,C) (F,E,D). After seeing 2000 itineraries the counts for A,F,C

and D are $c_{s,m} = 1000, c_{f,m} = 0$ and counts for B and E are $c_{s,m} = 1500, c_{f,m} = 0$. Suppose D damages the sheet during the itinerary (C,B,E,D). By the single fault assumption, modules A and F are exonerated and their posterior probability of failure is now 0. The w for modules B and E are now:

$$\left[\frac{1500}{1501}\right]^{1500} \frac{1}{1501} = .000245. \quad (4)$$

The term is higher for C and D as we have observed fewer samples of good behavior:

$$\left[\frac{1000}{1001}\right]^{1000} \frac{1}{1001} = .000368. \quad (5)$$

Normalizing, the posterior probability for B, and E failing are: 0.2 and for C, D are: 0.3. Suppose we see no errors in the next 2000 itineraries. Then, D damages a sheet in itinerary (D,E,F). By the single fault assumption, modules B and C are now exonerated. The values for w for D and E are now:

$$\left[\frac{2000}{2002}\right]^{2000} \left[\frac{2}{2002}\right]^2 = 1.352 \times 10^{-7}, \quad (6)$$

$$\left[\frac{3000}{3002}\right]^{3000} \left[\frac{2}{3002}\right]^2 = .601 \times 10^{-7}. \quad (7)$$

Normalizing $p(D|O) = 0.7, p(E|O) = 0.3$.

Table 3. The resulting posterior probabilities $p(M = m|O, U)$ over a more complex sequence of itineraries. One intermittent fault.

t	$m = A$	$m = B$	$m = C$	$m = D$	$m = E$	$m = F$
2001	0	.2	.3	.3	.2	0
4002	0	0	0	.7	.3	0
6003	0	0	0	.77	.23	0
8004	0	0	0	.83	.17	0
16008	0	0	0	.96	.04	0

In practice faults never occur with such regularity as in Table 3. Instead, every sequence of itineraries will yield different posterior probabilities.

As can be seen in this example, the restriction to single faults is a very powerful force for exoneration. All the modules not exonerated will have the same $c_{f,m}$ count. This results from the fact that under the single fault assumption, only modules that been used in every failing run remain suspect. Hence they have the same $c_{f,m}$. In our example, $c_{f,m} = 1$ in equations 4 and 5. After more observations, $c_{f,m} = 2$ in equations equations 6 and 7.

4.1 Incorporating prior counts

So far we presume nothing is known about the counts prior to making observations. If counts are initially 0, then the denominator of equation 3 will be 0. One possible approach to avoid this is Laplace's adjustment: make all initial counts 1, which is equivalent to assuming a uniform prior over $p(m)$. Another approach which we utilize in this paper is to observe that equation 3 need never be evaluated until an observation is made. The current observation is always included in counts, thus the denominator of equation 3 will never be 0 whenever we want to utilize it. Both approaches converge to the same in the limit as the number of observations grow to infinity.

$d \backslash o$	Fail	Success
$bad(d) \cap U \neq \emptyset$	1	0
$bad(d) \cap U = \emptyset$	0	1

$p(O|D = d, U)$

Fig. 6. Summary of the observation function in the multiple persistent case for an observation o of itinerary U .

One important detail we leave out of the examples for conciseness is that if a module has operated perfectly for very large counts it takes too many failing samples before its posterior probability rises sufficiently to be treated as a leading candidate diagnosis. Therefore, for our application, we apply a small exponential weighting factor λ at every increment such that counts 100,000 in the past will have only half the weight of new samples ($\lambda = 0.99999$).

5. MULTIPLE PERSISTENT FAULTS

Instead of modules, we consider diagnosis hypotheses D . Let $good(D)$ be all the good modules of D and $bad(D)$ all the bad modules of D . The number of possible diagnoses will be exponential in the number of modules. In practice, we only consider the more probable diagnoses, but for the moment consider the general case.

Analogous to the single persistent fault case:

$$p_t(D|O, U) = \alpha p(O|D, U) p_{t-1}(D).$$

To determine the prior probability of a diagnosis $p_0(D)$ we presume modules fail independently:

$$p_0(D) = \prod_{g \in good(D)} p_0(g) \prod_{b \in bad(D)} (1 - p_0(b)).$$

It remains to determine $p(O|D, U)$. If all the modules used in an itinerary are a subset of the good modules of a diagnosis d , then $p(Fail|D = d, U) = 0$ and $p(Success|D = d, U) = 1$. In every remaining case (i.e., if any of the used modules are bad in d), then $p(Success|D = d, U) = 0$ and $p(Fail|D = d, U) = 1$. Figure 6 summarizes these results.

For diagnostic purposes we need to compute the posterior probability that a particular module is faulted:

$$p(m|o_1, \dots, o_t) = \sum_{d \text{ s.t. } m \in bad(d)} p(d|o_1, \dots, o_t).$$

6. MULTIPLE INTERMITTENT FAULTS

We apply Bayes rule as before:

$$p_t(D|O, U) = \alpha p(O|D, U) p_{t-1}(D).$$

Let $p_b(m)$ be the probability that a module m produces an incorrect output when faulted. Let us first assume that $p_b(m)$ is given. In this case $P_t(D|O, U)$ is given by:

$$p_t(M|\mathbf{O}, \mathbf{U}) = \alpha w p_0(M) \quad (8)$$

$$w = \begin{cases} 1 - \prod_{m \in \text{bad}(D) \cap U} (1 - p_b(m)) & \text{If Fail} \\ \prod_{m \in \text{bad}(D) \cap U} (1 - p_b(m)) & \text{If Success} \end{cases} \quad (9)$$

As in the previous analyses, the posterior probabilities of the diagnoses is obtained by repeatedly applying Bayes rule. Before introducing a more direct method consider the result of applying Bayes rule repeatedly. Iterating Bayes rule results in:

$$w = \prod_{U \text{ fails}} [1 - \prod_{m \in \text{bad}(D) \cap U} (1 - p_b(m))] \times \prod_{U \text{ succeeds}} \prod_{m \in \text{bad}(D) \cap U} (1 - p_b(m)) \quad (10)$$

We consider both terms separately. The second term, success, can be computed simply by maintaining the counter (as in the single fault case) $c_{s,m}$ for each module:

$$\prod_{U \text{ succeeds}} \prod_{m \in \text{bad}(d) \cap U} (1 - p_b(m)) = (1 - p_b(m))^{c_{s,m}}.$$

To compute the first term we associate a single counter $c_{f,s}$ with each set of modules s utilized in a failing itinerary i . Since the approach is not dependent on the order of modules within an itinerary, we keep a counter per module set instead of for each itinerary to reduce the number of counters. For example we capture the two itineraries (A,B,C) and (C,B,A) by the same counter $c_{f,s}$ for module set $s = \{A,B,C\}$. Let S be the set of all such sets which have failed at least once. The first term is then:

$$\prod_{U \text{ fails}} [1 - \prod_{m \in \text{bad}(D) \cap U} (1 - p_b(m))] = \prod_{s \in S} [1 - \prod_{m \in \text{bad}(D) \cap s} (1 - p_b(m))]^{c_{f,s}} \quad (11)$$

Notice that we need not store the module sets of successful itineraries (by far the dominant case).

6.1 Learning the Intermittency Rate

As in the single fault case, we can learn the intermittency parameters of module failure $\{p_b(m)\}$, which we denote as q for brevity. In practice, it could be a single scalar (assuming that all the modules have the same intermittency parameter) or a vector (allowing the modules to have different intermittency). Here we present a general methodology for the estimation of q .

The goal of learning is to estimate the value of q to best match the observation O . To achieve this goal, we treat q as a deterministic unknown parameter, and formulate the learning problem as a maximum-likelihood estimation problem:

$$\hat{q}_{ML} = \arg \max_q p_q(O). \quad (12)$$

Here O is the observation history, i.e., the itineraries and their corresponding outputs. The question now is how to evaluate $p_q(O)$. For this, we have:

$$p_q(O) = \sum_D p_q(O|D)p(D) \quad (13)$$




Fig. 7. Learning of intermittency q . ($\ln p_q(O)$ vs. q)

where $p(D)$ is the prior probability, initially all equal for all hypotheses. The observation likelihood $p_q(O|D)$ is the $p(O|D,U)$ given by (9) in the previous section; the itinerary U is known, hence we remove U to simplify notation. Given any intermittency parameter q (which equivalently specifies all the module mis-operation probability $p_b(m)$), we can evaluate the observation likelihood $p_q(O)$ by the marginalization (13). The optimal estimate is then obtained by search over the space for maximal $p_q(O)$.

Example. Assume all faulty modules have the same intermittency parameter, i.e., $p_b(m) = q$ for all modules m . In this case, given any itinerary, the probability of observing a success or a failure is:

$$p_q(O|D,U) = \begin{cases} 1 - (1 - q)^{n(D,U)} & \text{if Fail} \\ (1 - q)^{n(D,U)} & \text{if Success} \end{cases} \quad (14)$$

where the exponent $n(D,U) \triangleq |\text{bad}(D) \cap U|$ denotes the number of bad modules in the hypothesis D that are involved in the itinerary U . For any given D and U , $n(D,U)$ is easy to evaluate. This enables us to express $p_q(O)$ as simple polynomial function of q . We then search for optimal $q \in [0, 1]$.

Figure 7 shows the learning for a simple system consisting of five modules, among which two modules have faults with an intermittency rate of 0.2. Learning is done based on 100, 200, and 500 randomly simulated trials; we show the results as the green, blue, and red curves respectively. The curves plot the computed observation likelihood $\ln p_q(O)$ as a function of q . The maximum likelihood estimates are marked with circles. Here we see that with more trials, the estimated q is closer to the underlying true value. For example, with 500 trials, $\hat{q}_{ML} = 0.19$ (ground truth is 0.2). As more trials are incorporated, the likelihood $\ln p_q(O)$ has a more prominent optimal q estimate. This is expected.

This algorithm computes $p(D)$ and q simultaneously and converges rapidly. The generalization to the situation where all q 's are different requires a multi-dimensional optimization.

7. CAPABILITIES

The catastrophic fault assumption is not correct for complex modules. Each module type has a set of actions it can perform. One of those actions may be faulty, but the




Fig. 8. A more detailed figure of a three way module. The 6 possible paper movements (capabilities) are indicated on the diagram.




Fig. 9. A more detailed figure of 5 modules connected together moving two sheets of paper.

module may always succeed at other actions. Therefore, we apply the framework we have developed to actions, not modules. Each capability fails approximately independently. Figure 8 illustrates a 3 way module with 6 capabilities. Figure 9 illustrates 5 modules of the two types connected together. Circles indicate rollers, triangles indicate sensors, and two sheets of paper are indicated in red. Note that three modules can be acting on the same sheet of paper at one time.

It is possible to design machine configurations where a failure in the output capability of one module cannot be distinguished from a failure in the input capability of the connected module. In our framework, this will show up as a double fault when in fact only one of the two modules is faulted. We avoid this confusion by applying an idea from digital circuits to collapse indistinguishable faults. In addition, we always allow multiple faults: we have found most equipment always contains multiple, low frequency, intermittent faults.

8. INITIAL RESULTS

The acid success test is whether the posterior probabilities calculated by our approach, when incorporated into a larger system, improve overall performance (including planning, diagnosing and production). We ran two experiments: (1) posterior probabilities were assigned randomly, and (2) posterior probabilities computed using the approach of this paper. The only difference was the posterior probabilities assigned to the faults. We measured the number of sheets needed to isolate the module(s) once a fault is detected. Table 4 lists initial results. The first column is a fault intermittent rate. The second column is the number of sheets needed using our approach and the third column is the number sheets needed using a random approach. The table shows our approach requires far fewer wasted sheets (and therefore downtime) to isolate a fault.

The description of the overall system and more analysis of performance can be found in Kuhn et al. (2008).

9. CONCLUSIONS

This paper lays out a framework for continuously diagnosing any combination of persistent and intermittent faults. Furthermore we introduced an extension to learn

Table 4. Initial results over random with our approach.

q	this paper	random	improvement
0.01	47	202	430%
0.1	18	30	167%

the intermittency rate q simultaneously while we compute the posterior probabilities. We showed that it is possible to update the probabilities and learn the intermittency rate without recording all past itineraries. We describe a procedure based on maintaining a limited set of counters instead of the entire itinerary history without losing any valuable diagnosis information. Table 5 shows the granularity of the counters needed to store the entire diagnostic history. Note: the number of modules $|M|$ is much smaller than the number of module sets $|S|$ and the number of module sets $|S|$ is much smaller than the number of itineraries $|I|$. With this extension to model-based diagnosis we have

Table 5. Table shows the granularity of counters needed to store the entire diagnostic history.

	c_s	c_f
single fault	per module	per module
multiple fault	per module	per module set

applied on-line diagnosis to modular reprographic equipment. More importantly, it extends model-based diagnosis to the real challenges, such as efficient diagnosis of intermittent multiple faults, faced in diagnosing manufacturing plants, packaging equipment, laboratory test equipment, etc.

REFERENCES

- R. Abreu, P. Zoetewij, and A.J.C. van Gemund. An evaluation of similarity coefficients for software fault localization. In *In Proceedings of the 12th IEEE Pacific Rim Symposium on Dependable Computing (PRDC'06)*, Riverside, CA, USA, December 2006.
- J. O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Springer Verlay, New York, 1995.
- J. de Kleer. Diagnosing intermittent faults. In *18th International Workshop on Principles of Diagnosis*, pages 45–51, Nashville, USA, 2007.
- J. de Kleer and B. C. Williams. Diagnosing multiple faults. *Artificial Intelligence*, 32(1):97–130, April 1987. Also in: *Readings in NonMonotonic Reasoning*, edited by Matthew L. Ginsberg, (Morgan Kaufmann, 1987), 280–297.
- M.P.J. Fromherz, D.G. Bobrow, and J. de Kleer. Model-based computing for design and control of reconfigurable systems. *The AI Magazine*, 24(4):120–130, 2003.
- Israel Koren and Zvi Kohavi. Diagnosis of intermittent faults in combinational networks. *IEEE Trans. Computers*, 26(11):1154–1158, 1977.
- Lukas Kuhn, Bob Price, Johan de Kleer, Minh Do, and Rong Zhou. Pervasive diagnosis: the integration of active diagnosis into production plans. In *Proc. AAAI 2008*, Chicago, 2008.