A framework for continuously estimating persistent and intermittent failure probabilities

Johan de Kleer, Bob Price, Lukas Kuhn, Minh Do, Rong Zhou
Palo Alto Research Center
3333 Coyote Hill Road, Palo Alto, CA 94304 USA

Abstract

Almost all previous work on model-based diagnosis has focussed on persistent faults where the prior probability of component failure is provided by the manufacturer or estimated from fleet-wide service records. 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. This paper also extends model-based diagnosis to systems where material instead of information in the form of voltages, currents, pressures is conveyed from one component to another.

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, Bobrow, & de Kleer, 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.

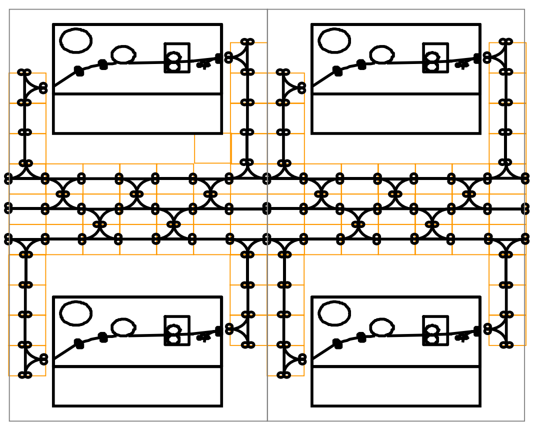

Figure 1: Model of It consists of two towers each containing 2 internal printers (large rectangles). Sheets enter on the left and exit on the right. Dark black edges with small rollers represents a bidirectional paper path. There are three main paper(horizontal) highways within the modular printer. The printer incorporates 2 types of media handling modules represented by small lighter edge rectangles (described in more detail in Section 8). 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.

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

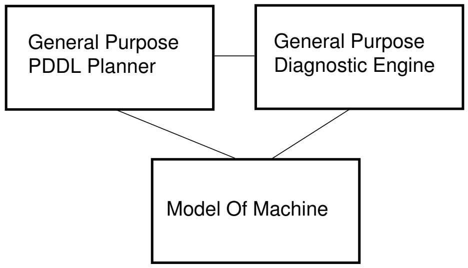

Figure 2: Basic architecture.

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 (mainfesting 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 multiple and persistent faults. Figure 3 illustrates the possibilities.

Each itinerary consists of a sequence of modules $m_1, \ldots, m_k$. We adopt the counting convention from [Abreu, Zoeteweij, & van Gemund, 2006] and associate four counters with each module m:

- m_{11}: number of plans where m was used and failed.
- m_{10}: number of plans where m was used and succeeded.
- m_{01}: number of plans where m was not used and failed.
- m_{00}: number of plans where m was not used and succeeded.

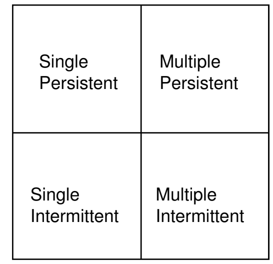

Figure 3: Fault combinations considered in this paper

	Fail	Success
Used	m_{11}	m_{10}
Not Used	m_{01}	m_{00}

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 outpust. Our approach still applies for such systems but requires more reasoning to determine whether a faulty output is masked. (See [de Kleer, 2007].)
- Fault probabilities are stationary. Our approach can be easily extended to accommodate slowly drifting probabilities through discounting.
- Bad module produces bad outputs. This holds in the reprographic domain but not in digital domains in which case more reasoning is required to estimate posterior probabilities.
- 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, indistinghuishable 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 & Williams, 1987]. Let $p(M)$ be the probability the module is faulted. The sequential bayesian filter [Berger, 1995] is:

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

Where α is chosen so that the posterior probabilities sum to 1 (presuming we start with the knowledge there is a fault). Let U be whether the module was used in the plan that produced the observation. $p(O|M,U)$ is 1 in situations where m_{00} or m_{11} 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 4 illustrates the possibilities.

u \ o	Fail	Success
Used	1	0
Not Used	0	1

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

Figure 4: 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.

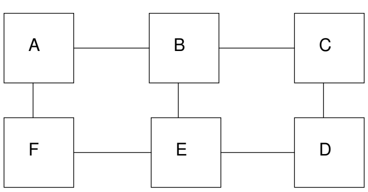

Figure 5: 6 module fragment with possible interaction paths.

Assume that at $t = 0$ all modules fail with prior probability $p_0 = 10^{-10}$. Consider the arrangement of modules in Figure 5. 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 & Kohavi, 1977]. In the case of intermittent faults, $p(O|M,U)$ is still 0 in case m_{01} and 1 in case m_{11}. 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:

$$\frac{m_{11}}{m_{11}+m_{10}}.$$

(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:

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

U represents whether the module was used in the plan. The observation function $P(O|M,U)$, is estimated from the counts m_{ij}. If the observation is a Failure and m is used, then:

$$p(Fail|M=m,U) = \frac{m_{11}}{m_{11}+m_{10}},$$

and if is Success and m is used, then:

$$p(Success|M=m,U) = \frac{m_{10}}{m_{11}+m_{10}},$$

otherwise as m cannot affect o, if m good,

$$p(Success|M=m,U) = 1,$$

otherwise,

$$p(Fail|M=m,U) = 0,$$

(captures the single fault assumption). Figure 6 summarizes the 4 possibilities.

After many iterations of Bayes rule, intuitively,

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

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

$$p_t(M|\mathbf{O},\mathbf{U}) = \begin{cases} 0 & \text{if } m_{01} > 0 \\ \alpha w p_0(M) & \text{otherwise} \end{cases} \tag{2}$$

where,

$$w = [\frac{m_{10}}{m_{11}+m_{10}}]^{m_{10}}[\frac{m_{11}}{m_{11}+m_{10}}]^{m_{11}}. \tag{3}$$

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

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

Figure 6: Summary of the observation function in the single fault intermittent case.

Consider again the arrangement of modules in Figure 5 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 m_{11} counters are 1 and the rest 0, therefore w's are 1. After observing (A,B,C,Success) the counters $(m_{00}, m_{01}, m_{10}, m_{11})$ for $\{A, B, C\}$ are all $0, 0, 1, 1$ and the counters for $\{D, E, F\}$ are all $1, 0, 0, 1$. 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 $1, 0, 1, 1$. The counters for D are $2, 0, 0, 1$ and the counters for $\{E, F\}$ are: $1, 0, 1, 1$. 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 $m_{10} = 1000, m_{11} = 0$ and counts for B and E are $m_{10} = 1500, m_{11} = 0$. 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 u for modules B and E are now:

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

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

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

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:

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

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

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 m_{11} 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 m_{11}. In our example, $m_{11} = 1$ in equations 4 and 5. After more observations, $m_{11} = 2$ 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 utlize it. Both approaches converge to the same in the limit as the number of observations grow to infinity.

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, consider diagnoses d which assign 'good' or 'faulted' to every system module. 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 $p(Fail|D=d,U) = 1$. Table 7 summarizes these results.

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

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

Figure 7: Summary of the observation function in the multiple persistent case for an observation o of itinerary U.

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

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

6 Multiple Intermittent Faults

The generalization to multiple intermittent faults follows by replacing every occurrence of a single module fault with a candidate diagnosis. We apply Bayes rule as before:

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

Each tentative diagnosis has counters associated with it: $d_{00}, d_{01}, d_{10}, d_{11}$. We need a set of counters exponential in the maximum number of module faults we consider. The diagnosis counters are incremented as follows for every d. In a failing itinerary involving a bad module of d, d_{11} is incremented otherwise d_{10} is incremented. $p(good|d)$ and $p(bad|d)$ can now be computed directly in the same way as the single fault intermittent case.

Consider again the system of Figure 5 considering both single and double faults. Initially there are 22 possible diagnoses (1 no faults, 6 single faults, 15 double faults). The first itinerary is (A,B,C,D,E,F,Fail). Now $d_{11} = 1$ for all remaining 21 diagnoses. Now we see (A,B,C, Success). This increments d_{10} for 17 diagnoses (shortened for brevity sake): A,B,C,AB,AC,AD,AE, AF, BC, BD, BE, BF, CD, CE, CF, DE, DF. Table 4 lists the counts after repeating the same 4 itineraries as for the single intermittent fault case and Table 5 lists the posterior probabilities for the suspect diagnoses.

Table 4: d_{11}, d_{10} for multiple intermittent faults.

t	d_{11}	d_{10}	candidates
1	1	2	D
14	1	3	A AB AC AD B BC BD C CD DE DF E EF F
6	1	4	AE AF BE BF CE CF

Table 5: Posterior probabilities for multiple intermittent faults.

t	p	candidates
1	0.22	D
5	0.16	A B C E F
14	1.6×10^{-10}	AB AC AD BC BD CD DE DF EF
6	1.2×10^{-10}	AE AF BE BF CE CF

The posterior probability of any particular component failed is:

$$p(C|\mathbf{O},\mathbf{U}) = \sum_{n \in S} p(n).$$

Where S is the remaining set of suspect candidates (those for which $m_{01} = 0$). Suppose we see the itineraries (A, B, C, fail) followed by (D, E, F, fail). In the single fault case, this would produce an error. In our example, $p = 0.09$ for diagnoses AE AF BE BF CE CF and $p = 0.16$ for diagnoses AD BD CD. The individual component failure probabilities are: $p = 0.33$ for C, A, and B; $p = 0.26$ for E and F; and $p = 0.48$ for D. The probabilities sum to 2 because the system contains 2 faults.

7 Replacement

The posterior module probabilities computed in our approach are the probabilities of misdiagnosis: that a replacement of the module will not provide any improvement in system performance. The decision whether to replace a module or not depends on the cost of replacement and that is out of scope of this paper. When a module is replaced, its prior is set using past experience with these types of modules, and all its counters m_{ij} are set to 0.

u \ o	Fail	Success
$bad(d) \cap U \neq \emptyset$	$\frac{d_{11}}{d_{11}+d_{10}}$	$\frac{d_{10}}{d_{11}+d_{10}}$
$bad(d) \cap U = \emptyset$	0	1

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

Figure 8: Summary of the observation function in the multiple intermittent case for an observation o of itinerary U.

8 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 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 9 illustrates a 2 way module with 2 capabilities. Figure 10 illustrates a 3 way module with 6 capabilities. Figure 11 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.

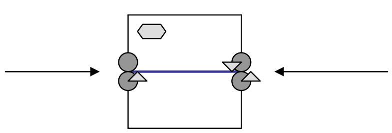

Figure 9: A more detailed figure of a two way module. The 2 possible paper movements (capabilities) are indicated with arrows on the diagram (paper can enter from the right and exit left, or enter from the left and exit right).

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 if 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.

9 Initial Results

The acid success test is whether the posterior probabilities calculated by our approach, when incorporated into a larger

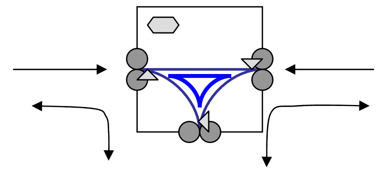

Figure 10: A more detailed figure of a three way module. The 6 possible paper movements (capabilities) are indicated on the diagram.

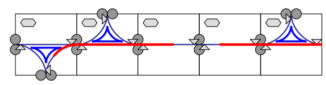

Figure 11: A more detailed figure of 5 modules connected together moving two sheets of paper.

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. We presumed single faults and both experiments used the single-fault exoneration rule. The only difference was the posterior probabilities assigned to the unexonerated single faults. We measured the number of sheets needed to isolate the module once a fault is detected. Table 6 lists initial results. The first column is a single 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 (with exoneration) approach. The table shows our approach requires far fewer wasted sheets to isolate a fault. The description of the overall system and more analysis of performance can be found in [Kuhn *et al.*, 2008].

Table 6: Initial results over random with our approach.

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

10 Conclusions

This paper laid out a framework for diagnosing any combination of persistent and intermittent faults in a continously operating piece of machinery where objects, not signals, are passed from one module to the next. With this extension to model-based diagnosis we have applied on-line diagnosis to modular reprographic equipment. More importantly, it extends model-based diagnosis to the real challenges faced in diagnosing manufacturing plants, packaging equipment, laboratory test equipment, etc.

References

[Abreu, Zoeteweij, & van Gemund, 2006] Abreu, R.; Zoeteweij, P.; and van Gemund, A. 2006. An evaluation of similarity coefficients for software fault localization. In *In Proceedings of the 12th IEEE Pacific Rim Symposium on Dependable Computing (PRDC'06)*.

[Berger, 1995] Berger, J. O. 1995. *Statistical Decision Theory and Bayesian Analysis*. Springer Verlay, New York.

[de Kleer & Williams, 1987] de Kleer, J., and Williams, B. C. 1987. Diagnosing multiple faults. *Artificial Intelligence* 32(1):97–130. Also in: *Readings in NonMonotonic Reasoning*, edited by Matthew L. Ginsberg, (Morgan Kaufmann, 1987), 280–297.

[de Kleer, 2007] de Kleer, J. 2007. Diagnosing intermittent faults. In *18th International Workshop on Principles of Diagnosis*, 45–51.

[Fromherz, Bobrow, & de Kleer, 2003] Fromherz, M.; Bobrow, D.; and de Kleer, J. 2003. Model-based computing for design and control of reconfigurable systems. *The AI Magazine* 24(4):120–130.

[Koren & Kohavi, 1977] Koren, I., and Kohavi, Z. 1977. Diagnosis of intermittent faults in combinational networks. *IEEE Trans. Computers* 26(11):1154–1158.

[Kuhn *et al.*, 2008] Kuhn, L.; Price, B.; de Kleer, J.; Do, M.; and Zhou, R. 2008. Pervasive diagnosis: the integration of active diagnosis into production plans. In *Proc. AAAI 2008*.