Lifetime Information for Model-Based Diagnosis*

Ahmed Y. Tawfik and **Eric Neufeld**
Department of Computer Science, University of Saskatchewan
Saskatoon, Saskatchewan,Canada S7N 0W0
tawfik@cs.usask.ca and *eric@spr.usask.ca*

Abstract

In this paper, we propose an extension to model-based diagnosis that includes lifetime information. A probabilistic temporal knowledge representation consisting of a probabilistic ATMS with temporally changing probabilities represents lifetime information. Adding statistical survival models to such a knowledge representation gives a more powerful reasoning tool that can evaluate life expectancy under specific operating conditions. Not only does this extension speed up diagnosis but it also detects faults traditional diagnostic engines cannot.

1 Introduction

In his theory of diagnosis from first principles, Reiter [18] devises a systematic approach to diagnosis based on checking the discrepancy between model predictions and observations. A diagnosis is a set of components whose failure explains the inconsistencies between observed and predicted behaviors. In [6], de Kleer and Williams use component failure probabilities to guide diagnosis in Sherlock. Sherlock considers the component failure probability constant. The changes in this probability due to aging, operating conditions and decay are ignored.

Portinale ([16] and [17]) develops a Markov model with states corresponding to different device modes (normal, leaking and broken are possible modes for a pipe). Components are more likely to fail as their mean time to failure gets closer. This is particularly true for wear and fatigue failures. Measurements and observations are collected for the components expected to fail. In this model, component lifetimes are encoded implicitly in the mean time to absorption ($MTTA$). Expressing the lifetime as a single number, namely $MTTA$ does not provide information about the failure distribution. $MTTA$ may be a good approximation when considering components with narrow failure time centered around the mean (assuming a normal distribution). In many practical situations failure rates tend to have a large variance or even follow other distributions such as the bathtub distribution or the exponential distribution illustrated in Figure 1. In other cases the failure time depends on other factors and the failure time cannot be calculated in isolation. The probability of failure of each component at a particular time point is more useful in guiding the diagnosis process (than the mean time to failure for the component) because different components have different failure time distributions.

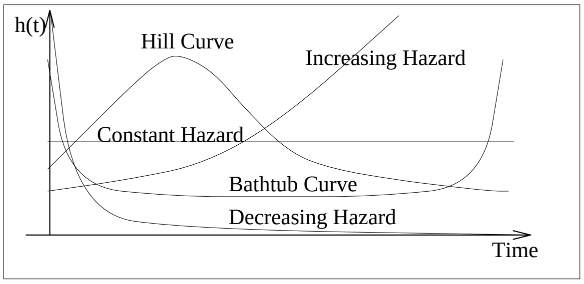

Figure 1: Some Typical Hazard Distributions

In this paper, we propose a technique that uses established statistical survival models and a probabilistic ATMS representation to reason about failures in physical systems. The knowledge base consists of two parts:

- The component specifications include the normal behavior, component failure modes, environmental and occupational factors affecting failure and the probability distribution of failure time.
- The trace is a timed list of known environmental and operational conditions.

The lifetime analysis module reads the trace and the component specifications to calculate a failure probability for each component at a given time. These probabilities are then fed into a Sherlock- like ATMS system. The ATMS uses these probabilities to guide search, assuming the components more likely to fail at this time would fail. It gradually expands the search to include less likely suspects until a set of components that explains all the observed behav-

*The authors acknowledge the support of the University of Saskatchewan and the Natural Sciences and Engineering Research Council of Canada NSERC.

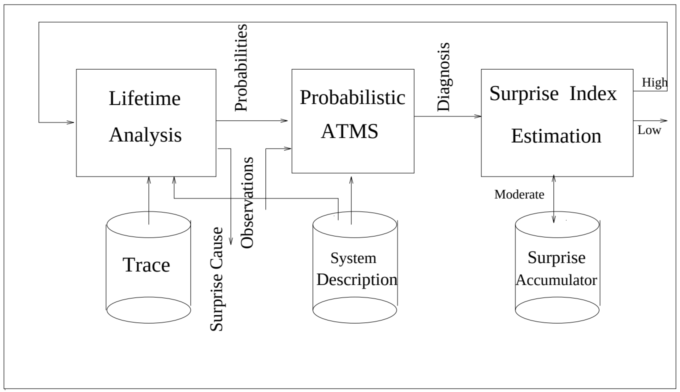

Figure 2: Diagnostic System Structure

ior is found (additional verification by measurements can be done at this stage). The surprise evaluation module is given a list of the failed components along with the failure probabilities of all components to calculate a surprise index. A low surprise index indicates that this failure is expected, a moderately surprising failure is noted for future follow-up and a highly surprising failure is immediately passed for further investigation. The investigation would start by checking possible environmental or operational factors that are not recorded in the trace and can cause this failure. If such a factor is identified, and it is controlled by a part of the system under diagnosis then the diagnosis process is repeated with this factor as an observation. If the factor could not be identified or not within the system being diagnosed, the operator is notified. Figure (2) shows the basic blocks of the proposed diagnosis system.

2 Time and Diagnosis

Nökel and Lamberti [12] classify temporal knowledge relevant to diagnosis into three categories:

- small scale phenomena,
- large-scale phenomena and
- very-large-scale phenomena.

The first of the three categories deals with very fast changes. The instantaneous value of such phenomena is not directly significant but are rather expressed in terms of some abstractions (e.g. the temporal abstraction 'frequency' is used for the physical phenomenon 'vibration'). Large-scale phenomena include changes that are significant and observable. Such changes are called events (e.g. the READ signal in a memory read cycle). The purpose of this paper is to deal with very-large-scale phenomena. The duration of such phenomena is longer than the diagnostic session itself and include long term effects such as device aging, material wear, corrosion and fatigue. Attempts to include this type of knowledge are limited to few works such as [17], discussed earlier.

3 The Probabilistic ATMS

Here we propose an ATMS that associates probabilities with its nodes. Each node corresponds to a component's mode. Conflicts are used to ensure consistency between assumptions (for example the same component cannot be in more than one mode at the same time). This particular knowledge representation is suitable for diagnostic reasoning [6]. Bayesian networks could be used instead of the probabilistic ATMS. Bayesian nets are used for diagnosis in [15], [14] and [5]. The advantage of the ATMS is that it can use focusing techniques. These techniques allow an ATMS to consider the components more likely to fail first.

4 Lifetimes: Why and How?

Intuitively, an older component is more likely to fail than a relatively new one [1]. Ignoring the notion of expected lifetime simply means losing valuable time especially in systems that use the probability of failure to guide the diagnosis. Human diagnosticians, especially in medicine, seem to consider age as a very significant factor. This is also true in the case of physical systems.

A better reason for including lifetime information is that some faults accelerate the failure of some other components. Replacing those components would rectify the problem temporarily. A diagnostic engine that does not take the component lifetime into account will continue replacing the failed components without ever detecting the real problem. In practice a light-bulb wired to a higher than nominal voltage burns-out quickly. A car alternator over-charging a battery would require more frequent battery fluid addition. The proposed diagnostic engine detects the repeated violations of lifetime expectation and makes additional measurements and tests. The device operating conditions such as temperature, humidity may also explain expected lifetime violations. To allow for such factors regressive lifetime models are used.

Lifetime is therefore part of the component description. The system description SD is supposed to include a reasonably accurate representation of component lifetimes as well as conditions known to affect the expected lifetime. Statistical survival models are a quantitative tool useful in making inferences regarding lifetimes. Lifetime inferences are not only limited to lifetime prediction but they also include finding the possible reasons for unexpected failures.

4.1 Getting the Lifetime Information

The lifetime of a component c_i is a random variable T_i defined as the time at which the component fails. The random variable T_i has a probability density function

[1] Reliability models suggest that a very new component may have a high initial failure probability due to burn-in faults.

$f_i(t)$ indicating how failure likelihood changes with time. The time scale can be clock time, utilization time or both. The failure of a chain used to hang a ceiling lamp does not depend on the utilization of the lamp and therefore the clock time is an appropriate scale. The failure of a light bulb depends mainly on the utilization time. Most devices however are affected by utilization and clock time (batteries for example). A bivariate model developed to represent the two time scales case is used to analyze the failure of cars during the warranty period [20].

The lifetime of components generally depends on three types of factors: inherent, environmental and operational. The inherent factors include the physical aspects of the device and the manufacturing process. Environmental factors are uncontrollable outer effects such as temperature and humidity. The specific job or function that the device is performing within the system constitutes an operational factor. The reasoning required is one that can account for such factors for each individual component. [2] Reliability engineers use statistical techniques to estimate lifetimes. They conduct experiments, usually under controlled conditions to find a mathematical model that explains the failure mode of each product as a function of time and other factors. The models usually used for this purpose are called survival models. Here we assume that these models are available and we are only going to use them to predict the failure time of each component.

4.2 Tracing and Monitoring

To predict a component's lifetime, we need to know the age of the device, its utilization time and the conditions under which it was operated. An easy way to get information about a component's age is from maintenance records. It is also possible to get the diagnostic system to keep track of component replacement especially in the cases where the diagnostician is built into the system. In other situations, physical features may provide some clues about the newness of an element. This page will tend to become yellowish with time thus providing a clue about its recency to the reader. Contextual knowledge may be another source for component lifetime estimation. For example, assuming that a new car has a new engine, is usually a valid assumption.

The utilization time is sometimes measured by the system itself, the mileage on a car is a measure of its utilization. More often it can only be estimated. The estimation of utilization time is based on assumed utilization patterns. Regression analysis of utilization under various conditions can be used to capture such patterns. The system does not deduce such patterns and would ask the user to provide this information if it is important and not available.

[2] Approximations may be required to discard the less important factors.

The tracing subsystem, keeps track of the environmental conditions. To limit the storage requirements and the processing time, a time window is specified and older information is summarized and overwritten. The summarization process consists of providing an approximation for the changes that took place earlier using one of three techniques: 1) averaging over periods as long as there is no significant changes, 2) using the average in addition to a seasonal and trend parameters or 3) curve fitting.

4.3 Survival Analysis and Models

Survival analysis is a relatively new sub-area of statistics that emerged in response to the need to study events and event interaction. It differs from Markov models in that it uses distributions rather than single transition probabilities. It differs from life-tables methods in that it can account for time varying and time invariant factors affecting lifetimes. Survival analysis techniques, used in a wide variety of disciplines, are mostly the same whether the center of interest is the failure time of an engine or the success time in performing a given task in a learning process. The 'failure' is any event of interest that can occur at any point in time. The 'hazard' is the rate at which the event occurs or can occur and a 'risk' is a potential cause for failure. Here we assume continuous time, but the basic ideas are also valid for discrete time. Allison [1] provides a good introduction to survival analysis. Their use in probabilistic temporal reasoning is described in [21].

The temporal distribution of failure can be expressed as a survival function $S(t)$, a probability density function $f(t)$, a probability distribution $F(t)$ or a hazard function $h(t)$. The probability density function $f(t)$ is defined as

$$f(t) = \lim_{\Delta t \to 0} \frac{Pr(t \leq T < t + \Delta t)}{\Delta t}. \tag{1}$$

The probability distribution function $F(t)$ is the probability that the failure occurs before t and is defined in terms of $f(t)$ as

$$F(t) = Pr(T \leq t) = \int_0^t f(x)dx. \tag{2}$$

The survival function $S(t)$ is the probability that the failure occurs after time t and is defined as

$$S(t) = Pr(T \geq t) = \int_t^\infty f(x)dx. \tag{3}$$

The hazard function $h(t)$ is the rate at which the failure occurs and it is defined as

$$h(t) = \lim_{\Delta t \to 0} \frac{Pr(t \leq T < t + \Delta t | T \geq t)}{\Delta t}. \tag{4}$$

We can also deduce that $h(t) = \frac{f(t)}{S(t)}$. By substitution and integration we find that

$$S(t) = exp(-\int_0^t h(x)dx). \tag{5}$$

From the above equations, we can draw several useful conclusions. First, simple relations exist between the four functions and they can be deduced from one another. Actually knowing one of them fully specifies the others. Second, the survival function can take any shape based on the corresponding hazard function. For example, a constant hazard would result in an exponentially decreasing survival probability. Third, the analysis assumes that failure occurs once. This may be true in some cases, but does not take repair actions into consideration. Here, we shift in time the survival function to start at the repair time for replaced components. Repair is not necessarily a replacement procedure, special repair models apply in such situations [2]. Fourth, the equations described do not take into consideration the effect of variables other than time. The solution of this problem requires the use of models. These models replace $h(t)$ by a conditional version $h(t|X_1 \dots X_n)$ where X_i is in general a time varying factor affecting the survival. Those factors are sometimes called explanatory variables.

The expected number of failures $H(t)$ during a time period $[t, t+\delta t)$ is roughly the result of integration of the hazard function or $H(t) = \int_t^{t+\delta t} h(x)dx$. For small δt this integration can be approximated using the trapezoidal rule as $0.5(h(t)+h(t+\delta t)\delta t$. Here we use $H(t)$ as the probability of failure mode m_j at time t for component c_i (or $P_t(c_i, m_j)$). Technically speaking $h(t)$ is a conditional probability in the discrete time case only. In the case of continuous time $h(t)$ is not a probability nor is its integration. For continuous time $H(t)$ is the expected number of failures and therefore can take values greater than unity. When the value of $H(t)$ is much smaller than unity, it is a better estimate to base the diagnostic decision on. Values for $H(t)$ greater than one are not likely in this particular application because a component can fail only once in its lifetime and we are usually interested in the value of $H(t)$ over a short interval just before the time of failure (or from the last time the system was operational until the time or failure). The use of $h(t)$ and $H(t)$ in making repair and maintenance decisions is a common practice in engineering. Ascher and Feingold [2] discuss this issue further.

4.3.1 Causal Interactions

Attaching causal semantics to the statistical models makes it easier to integrate them into a knowledge based system. Causal models assume that events trigger effects either immediately or after a time lag. The effects produced by an event may be affected by other events. The causal ordering specifies a temporal relation between events and effects such that the effect of an event cannot precede the event itself. Events that do not interact at all, at neither the causal side nor on the effect side are completely independent or non-interacting.

Two causal relationships, namely $RESULT$ and $ENABLE$ are sufficient to describe the causal interactions underlying the models. $RESULT$ establishes a cause-effect relationship between an event and a state. A state is said to $ENABLE$ an event if it is a necessary condition for the event to take place [13]. The less formal notation 'X is a $RESULT$ of $C_1 \vee \cdots \vee C_n$' is used instead of '$RESULT(X, C_1 \vee \cdots \vee C_n)$' for improved readability whenever possible. Similarly for $ENABLE$, it is employed rather casually and 'a state $DISABLEs$ an event' is used to mean the complement of $ENABLing$ it. The terms *more likely* and *less likely* represent the effect on the probability of an event.

4.3.2 Competing Risks Model

As the name suggests, the competing risks model represents two or more potential risks racing to achieve a failure, but the success of one of them inhibits the others. $C_1 \vee \cdots \vee C_n$ $RESULTs$ in state S and state S $DISABLEs$ $C_1, \cdots, C_n$ from succeeding. The state S may be death and $C_1, \cdots, C_n$ are potential causes for death. S is not necessarily a final state, it may be one that just briefly blocks the other competing causes. For example, consider the case of two infections with the same virus. The state *anti-bodies present in blood* blocks *second infection*. Competition is a relation between events and in the statistical analysis of this model the nature of S does not affect the analysis. According to the underlying causal model, this statistical model applies to many situations. See [11] for possible application areas.

Berzuini [3] represents competing events using the networks of dates. In this representation two competing events C_1 and C_2 produce their effects at T_1 and T_2. T_1 and T_2 are temporal random variables of known distributions.

The result X occurs at $\min(T_1, T_2)$. To find the probability distribution for X which is the $RESULT$ of the competition of two causes with failure densities f_1 and f_2. Using survival analysis, the probability density of X is given by:

$$f(t) = f_1(t)S_2(t) + f_2(t)S_1(t) \qquad (6)$$

where the survival functions S_1 and S_2 are defined as above. The use of the survival analysis provides a compact and efficient way to represent and evaluate the overall effect of competing events.

4.3.3 Proportional Hazard Model

This model is one of the more widely used survival models. It is a parametric model that allows the effect of environmental and other factors to be taken into consideration. This model, first proposed by Cox [4] assumes that the natural logarithm of the ratio of the conditional hazard function (in the presence of explanatory variables) to the hazard $h_0(t)$ (in their absence) is a linear weighted sum of the risks or

$$h(t|X_1 \dots X_m) = h_0(t)e^{\sum_i^m \beta_i X_i(t)}. \qquad (7)$$

To overcome the limitations implied by the linearity assumption made by this model, nonlinear mapping functions may be used such that the hazard function remains a linear combination of the mapped factors. For example, an engine is more likely to fail when operated at an unusually high temperature. The hazard function therefore becomes

$$h(t|high - temperature) = h_0(t)e^{\beta z}$$

where z represents the temperature and β is a parameter that reflects the effect of temperature on the engine's failure probability. $h_0(t)$ is calculated in the absence of high temperature. This new function models the effect of *temperature* on *engine-lifetime*. The causal interpretation of this model can be viewed as in the previous cases through the introduction of a state. The *high-temperature RESULTs* in the engine state *less heat dissipation* and the event *engine-fails* is then more likely to happen quickly.

4.3.4 Accelerated Time

A serious limitation of the proportional hazard model described above is that it assumes time invariant effect of the factors. This is not true in many cases. The effect produced by a certain factor depends on time. It is possible in such cases to use functions $\beta(t)$ instead of the parameters β in the model described above. Techniques that support such models have only recently been proposed by West [22] for continuous time and Singer and Willett [19] for discrete time. An alternative model for dealing with effects of explanatory variables on lifetime is to consider a different time scale t' and find a relation $m(t, X_1, \ldots, X_n)$ between a normal time unit, say a second and a time unit under the effect of the factors. This is done by assuming that the failure time depends on the explanatory variables. A commonly used function is $log(T) = \sum_i \beta_i X_i$. Kalbfleisch and Prentice [11] provide a more detailed description of this model. Substituting t' for t in the original survival function produces the new function in the presence of the factors. The effect of a higher speed on the probability that the car runs out of gas can be accounted for using this model. In this case, event *double the speed RESULTs* in the car *consuming more gas* and the event *out-of-gas* occurring sooner.

4.3.5 Additive Storage Model

Systems may fail due to the accumulation of stress. To model this type of failure, consider a fixed capacity reservoir with an incoming flow (in-flow) and an out-going flow (out-flow). The capacity of the reservoir is the maximum stress the system can tolerate without failure. The incoming flow is the additive stress applied to the system and the out-flow reflects the ability of the system to recover from stressful situations. Systems with adequate recovery can tolerate large stresses occurring over a long duration. They may fail however when the the same stress is concentrated during a shorter duration. The special case of unlimited recovery corresponds to systems that would only fail if the instantaneous stress applied to them exceeds the maximum stress they can tolerate. Systems with no recovery let the stress add up until failure. The use of different release rules (for the outflow) can be motivated by the fact that most systems exhibit different patterns. Metals have unlimited recovery to tensile stresses within the elastic region and have no recovery once the stress exceeds a critical value causing plastic deformation.

Some systems have the ability to learn how to deal with stress. These systems include adaptive systems and many biological systems. This behavior can be modeled in the storage context by reducing the amount of stress entering the reservoir resulting from an occurrence of this risk once the system has survived a similar risk. This process may be permanent (as in the case of measles) or temporary (as for the case of flu).

For storage models, the causal model assumes collaboration, the events have an additive *RESULT* but limited persistence (due to recovery). This model can be seen as a temporal variant of the Noisy OR-Gate model [15]. Storage models have been used to model dams, warehouses and similar systems. Glynn [8] gives the detailed analysis of a discrete-time storage process. The classification of systems response to stress is inspired by the treatment in [7].

4.3.6 Sequential Model

This model applies when a system has a set of spare /reserve components used to automatically replace the failed ones. The device fails when all the spare parts are used up. In the case of two such components with failure probability distributions $F_1(t)$ and $F_2(t)$. The total time to failure has a probability distribution $F(t)$ given by

$$F(t) = \int_0^\infty F_1(t)F_2(t-\tau)d\tau. \tag{8}$$

Equation (8) follows directly from the distribution of the sum of random numbers. In fact the total time to failure T is the sum of the random numbers T_1 and T_2 (the time to failure of the first and second components respectively). The underlying causal model consists of a single enabled event. The completion of this event results in a state *ENABLing* another event until the sequence is completed.

5 Reasoning about Lifetimes

Surprises [3] allow us to discover the causes of unanticipated failures. A surprise is the occurrence of an

[3] Some earlier works have discussed the applicability of surprises to rational agents. In [10] a rational agent can be surprised no matter what, this argument relies on a proof that uses Dempster- Shafer theory of evidence. Hsia argues that probability theory cannot capture the intuitive notion of surprise. The surprise index discussed above is based on work by Good [9].

event that is very unlikely to happen. The surprise index is a measure of the degree of surprise associated with the occurrence of an event [9]. A simple mathematical formulation of a surprise index is

$$Surprise = \frac{\sum_i (p_i)^2}{p_r}. \tag{9}$$

The summation is over all mutually exclusive possible outcomes and p_r is the probability of the event that actually occurs. The surprise index values ranges from zero to infinity. A value between zero and one corresponds to the case of a likely outcome occurring. Values greater than one indicate a surprise, the larger the index the more astonishing.

Less likely outcomes are more surprising than the more probable ones. However, the surprise index is a better indicator of the degree of surprise than the probability of the event by itself. The following example illustrates how the surprise index works. While it is not surprising at all that *someone* wins the lottery it is a big surprise if *I* win. Let there be n tickets sold to an equal number of persons and let us assume that the draw is among the sold tickets. The probability of winning for each individual is then $1/n$. Let X be a person unknown to us. The surprise index for the proposition *X won the lottery* is then

$$Surprise = \frac{n(1/n)^2}{1/n} = 1.$$

The proposition *I won the lottery* has only two possible outcomes actually winning or not with probabilities $1/n$ and $(1 - 1/n)$. For these two outcomes, the surprise associated with winning is

$$Surprise = \frac{(1/n)^2 + (1 - 1/n)^2}{1/n} = n - 2 + 2/n.$$

For a large n the surprise is very large, while there is no surprise at all if $n = 1$ (why be surprised to win if there is only one ticket sold?). The surprise index measures how unlikely an outcome is compared with other possible outcomes.

In diagnosis, a truth maintenance system may find n possible diagnoses for a particular malfunction. Each diagnosis assigns a failure mode to each failed component. Δ_i denotes the set of failed components for diagnosis i. A temporal failure probability $P_t(c_i, m_j)$ denotes the probability that at time t the component c_i be in the failure mode m_j. Assuming independent failures [4], the probability $P_t(\Delta_i)$ is given by the product:

$$P_t(\Delta_i) = \prod_{\forall (c_i, m_j) \epsilon \Delta_i} P_t(c_i, m_j).$$

[4]This assumption is common to many diagnostic engines. In some practical situations, it may not be reasonable. A joint probability distribution replaces the product in such case.

Moreover, the probability the failure of the set of components Δ_i has caused mlfunction M_j is given by

$$P_t(\Delta_i | M_j) = \frac{P_t(\Delta_i) P(M_j | \Delta_i)}{P_t(M_j)}.$$

A single fault assumption is equivalent to assuming that the failure of components is mutually exclusive. In this case the surprise associated with the failure of one component c_r is given by

$$Surprise = \frac{\sum_{\forall c_i} (P_t(c_i))^2}{P_t(c_r) \sum_{\forall c_i} P_t(c_i)}. \tag{10}$$

The summation in the denominator of Equation(10) is a normalization factor introduced because the sum of probabilities does not necessarily add up to one. The surprise associated with multiple faults is more easily expressed in terms of the logarithmic surprise index. The surprise associated with the failure of two components $r1$ and $r2$ is given by

$$Surprise = \sum_{c_i} P_t(c_i)(\log(P_t(c_i) - \log(P_t(c_{r1})) + \sum_{c_i} P_t(c_i)(\log(P_t(c_i) - \log(P_t(c_{r2})).$$

For any number of components the logarithmic surprise index is given by the sum of individual logarithmic surprise indexes. For the logarithmic surprise index, a value of zero or less indicates no surprise. A positive value indicates surprise.

Unsurprising failures need not be processed any further. For a moderately high surprise index, replacement of the failed component and monitoring its performance is a sensible decision. For example, if the probability of a new light bulb failing after ten hours of operation is ϵ (ϵ being a small number). The logarithmic surprise index associated with the failure of the bulb is s. The logarithmic surprise index resulting from n new light bulbs failing one after the other, having operated for the exact same duration and under the same conditions is then ns. The accumulation of surprises is done on a limited time span equal to the mean lifetime of the component. Two threshold surprise values S_{c_iL} and S_{c_iM} are chosen for each component c_i. Surprises higher than the low surprise threshold S_{c_iL} but less than the moderate surprise threshold S_{c_iM} are stored in the surprise accumulator. Surprises higher than S_{c_iM} prompt a more careful and detailed investigation into the causes of failures. Setting the surprise thresholds depends on the properties of the lifetime distribution of the component and some economic factors such as the replacement cost as opposed to the cost of additional measurements.

Once the surprise estimation module detects a surprise higher than the moderate surprise threshold either by itself or when combined with previous moderate surprises for the same component, the scope of the

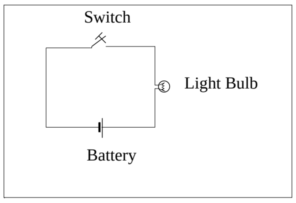

Figure 3: The Circuit

fault is modified to include this surprise as a symptom. This is done by the survival analysis module which will then try to find explanatory variables to justify the failure. The explanation is checked for consistency with the information available in the trace and the behavior of other elements.

6 Example

A simple torch (flashlight) circuit consists of a bulb, a switch and a battery connected as in Figure 3. The probability distributions for the lifetime of the bulb and the battery are normal with means of 1000 and 20 hours respectively. The variances are 200 and 5 for the bulb and the battery. The wiring and the switch rarely fail but their probability of failure is high initially due to burn-in faults. Then it drops as these defects usually affect the torch during the first few hours of operation. The failure probability finally rises again with aging. Figure 4 shows the hazard function distributions corresponding to these failure distributions.

To illustrate how this approach works, we consider three situations: expected failure, unexpected failure and preventive maintenance.

Scenario 1. A torch fails, the switch and the bulb have operated for 500 hours, the age of the battery is 16 hours. Using the hazard distribution given, the probability of failure of the switch during its 500^{th} hour is 0.0023. The probability of failure of the bulb during the same period is 0.000092. The probability of failure of the battery during its 16^{th} hour is 0.013. The probabilistic ATMS starts by considering a battery failure. This failure explains all observed symptoms and testing the battery verifies its failure. The surprise index in this case is equal to $\frac{(0.013)^2+(0.000092)^2+(0.0023)^2}{0.047} = 0.0136$. A normalization step is required here because the sum of probabilities is less than one. The normalized index is $0.013/(0.013+0.0023+0.000092) = 0.881$. This low surprise index indicates that the failure of the battery at this time point is not surprising at all.

Scenario 2. The torch, fails again four hours after replacing the battery.

The probability of failure of the bulb and that of the switch have not changed much but that of the battery is now practically nil (7×10^{-5}). The ATMS considers the switch first but it turns out that the

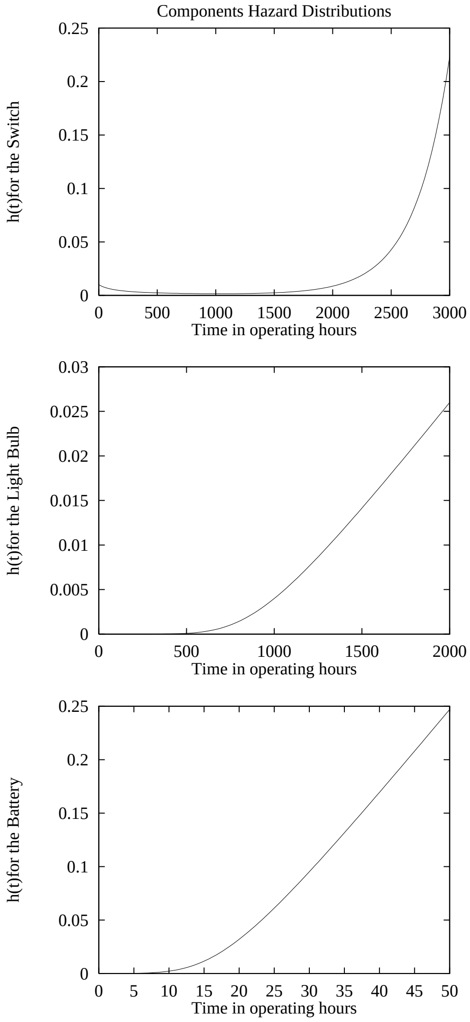

Figure 4: Hazard Distributions

bulb has failed. The normalized surprise index in this case is 24.07 . This value is high enough to prompt further investigations. Vibrations and high voltage are two risk factors for light bulbs. Measuring the voltage of the new battery indicates that it is defective because it is somehow generating 9 Volts instead of 1.5 Volts. An accelerated time failure model has predicted this condition.

Scenario 3. The torch with a new light bulb and battery is now to be used by a space exploration mission. They will be using it for five hours and because the weight is at premium, they cannot afford to take unnecessary spare parts. Any component that may fail during the next five hours is to be replaced by another with lower failure probability.

The component with the highest probability of failure during the five hours mission is the switch. Unfortunately, replacing it will raise the probability of failure about 10 times. It is therefore safer to leave it. Vibrations are a risk factor for the light bulb as mentioned earlier. For this reason it may be necessary to take a spare bulb and try to avoid exposing the torch to vibrations. The new battery stands a chance of 3 x 10^{-4} of failing during this period.

7 Conclusions

This work shows the importance of lifetime information in diagnostic problem solving. It introduces some statistical models as tools for reasoning about the failure probability of a device under given operating conditions and as a guide in explaining unexpected failures. The notion of surprise is used to point the system to unusual failures. Finding out the reasons for a surprising failure is part of the diagnosis.

References

[1] P. Allison. *Event History Analysis.* Sage, Beverly Hills, 1984.

[2] H. Ascher and H. Feingold. *Repairable Systems Reliability.* Marcel Dekker, Lecture Notes in Statistics, New York, 1984.

[3] C. Berzuini. Representing time in causal probabilistic networks. In *Uncertainty in Artificial Intelligence 5*, pages 15–28. Elsevier Science Publishers B.V., 1990.

[4] D. Cox. Regression models and lifetables. *Journal of the Royal Statistical Society*, B34:187–220, 1972.

[5] J. Breese D. Heckerman and K. Rommelse. Decision-theoretic troubleshooting. *Communications of the ACM*, 38(3):49–57, 1995.

[6] J. de Kleer and B. Williams. Diagnosis with behavioral modes. In *Proceedings of IJCAI-89, Detroit, MI*, pages 1324–1330, 1989.

[7] I. Chen F. Bastani and T. Tsao. Reliability of systems with fuzzy failure criterion. In *IEEE Annual Reliability and Maintainability Symposium*, pages 442–448, Anaheim, California, 1994.

[8] J. Glynn. A discrete-time storage process with a general release rule. *Journal of Applied Probability*, 26:566–583, 1989.

[9] I.J. Good. *Good Thinking: The Foundations of Probability and Its Applications.* University of Minnesota Press, Minneapolis, 1983.

[10] Y. Hsia. A rational agent can be surprised no matter what. In *Proc. of the ninth Canadian Conference on Artificial Intelligence AI'92, Vancouver, Canada*, pages 106– 112, 1992.

[11] J. Kalbfleisch and R. Prentice. *The Statistical Analysis of Failure Time Data.* John Wiley and Sons, New York, 1980.

[12] K. Nokel and H. Lamberti. Temporally distributed symptoms in a diagnostic application. *Artificial Intelligence in Engineering*, 6(4):196, 1991.

[13] M. Pazzani. *Creating a Memory of Causal Relationships : An Integration of Empirical and Explanation-Based Learning Methods.* LEA Publishers, Hillsdale, NJ, 1990.

[14] J. Pearl. Distributed revision of composite beliefs. *Artificial Intelligence*, 33:173–215, 1988.

[15] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kauffmann, San Mateo, CA, 1988.

[16] L. Portinale. Modelling uncertain temporal evolutions in model-based diagnosis. In *Proceedings of the* 8^{th} *International Workshop on Uncertainty in Artificial Intelligence*, pages 244–251, Stanford, CA, 1992.

[17] L. Portinale. Selecting observation time in the monitoring and interpretation of time-varying data. In *Advances in AI, Third Congress of the Italian Association for AI, AI*AI'93*, pages 314–325. Springer-Verlag Lecture Notes in AI, 1993.

[18] R. Reiter. 'a theory of diagnosis from first principles. *Artificial Intelligence*, 32:57–95, 1987.

[19] J. Singer and J. Willett. It's about time: Using discrete-time survival analysis to study duration and the timing of events. *Journal of educational statistics*, 18(2), 1993.

[20] N. Singpurwalla. Survival under multiple time scales in dynamic environment. In J. Klein and P. Goel, editors, *Survival Analysis: State of the Art.* Nato ASI Series, Kluwer Academic Publisher, Dordrecht, Netherlands, 1991.

[21] A. Tawfik and E. Neufeld. Temporal bayesian networks. In *Proceedings of Time-94:International Workshop on Temporal Knowledge Representation and Reasoning*, pages 85–92, Pensacola, Florida, 1994.

[22] M. West. Modelling time-varying hazards and covariate effects. In J. Klein and P. Goel, editors, *Survival Analysis: State of the Art.* Nato ASI Series, Kluwer Academic Publisher, Dordrecht, Netherlands, 1991.