Statistical Consistency With Dempster's Rule on Diagnostic Trees Having Uncertain Performance Parameters

Stephen D. Durham
Department of Statistics, University of South Carolina, Columbia, South Carolina

Jeffery S. Smolka
MEMC Electronic Materials, Inc., Spartanburg, South Carolina

Marco Valtorta
Department of Computer Science, University of South Carolina, Columbia, South Carolina

ABSTRACT

This paper defines statistical consistency, a property that we propose as a necessary characteristic of any calculus for evidence combination. Statistical consistency holds when the combination of repeated observations of a system in a given state leads to the indication that the system is in that state of nature. We show that, for a suitable choice of parameters, Dempster's rule has this desirable property, both for simple systems and for systems composed of a hierarchy of subsystems and described by diagnostic or fault trees, but for other parameter values the rule leads to the wrong conclusion. A necessary and sufficient condition for the existence of simple bpa's being statistically consistent is that $p_0 + q_1 > 1$*, where* p_0 *and* q_1 *are the reliability (or specificity) and sensitivity of individual sensors detecting malfunctions in components and (sub)systems. A sufficient condition for statistical consistency is that the reliability and sensitivity of each sensor be greater than* 0.5. *We show that statistical consistency is preserved under diagnostic*

Address correspondence to Marco Valtorta, Department of Computer Science, University of South Carolina, Columbia, SC 29208.

Received February 1, 1991; accepted June 3, 1991.

655 Avenue of the Americas, New York, NY 10010 0888-613X/92/$3.50

tree formation.

KEYWORDS: ***Dempster's rule, diagnostic trees, fault trees, statistical consistency, probabilistic reasoning, belief networks, expert system verification***

1. INTRODUCTION

This paper is concerned with the use of Dempster-Shafer theory for representing information and making decisions in a diagnostic setting. Dempster-Shafer theory represents pieces of evidence by belief functions and combines them according to Dempster's rule. There is still no general agreement on the interpretation of degrees of belief, which explains why (1) there does not exist a standard belief assessment methodology, and (2) the validity of Dempster's rule is an object of controversy (Shafer [1], Walley [2]).

With respect to this debate, we consider the attribution of beliefs to pieces of evidence and their combination by Dempster's rule without reference to any underlying interpretation or theoretical justification. As a practical matter, the fundamental question is: "Does it work?" In particular, we define an asymptotic property, which we call *statistical consistency*, and propose it as a modest necessary condition for the proper performance of Dempster's rule. Even though statistical consistency is a natural but rather weak requirement, it is not automatically achieved by Dempster-Shafer theory.

In the simplest case, we show that repeated independent readings of a sensor (or readings on multiple independent sensors) that has unknown reliability (specificity) p_0 and unknown sensitivity q_1 will finally identify the true state of a system if and only if the degrees of belief s_0 and s_1 associated with the readings "correct functioning" and "malfunctioning," respectively, satisfy

$$(1 - s_0)^{p_0/(1-p_0)} < 1 - s_1 < (1 - s_0)^{(1-q_1)/q_1}$$

The result provides constraints on the values of the beliefs to be attributed to the pieces of evidence in that case. It may be that s_0 and s_1 are chosen in such a way that these constraints are not satisfied, because p_0 and q_1 may be unknown, thus leading to statistically inconsistent calculations. Indeed, if $p_0 + q_1 < 1$, then there is no choice of probability masses s_0, s_1 that leads to statistically consistent results. If the values of p_0 and q_1 are not known exactly, but bounds for these quantities are known, the existence of belief values s_0 and s_1 giving statistically consistent results depends only on the bounds. In particular, if both the reliability and the sensitivity of the sensors are greater than 0.5, then the choice $s_0 = s_1$ will always work provided $0 < s_0 = s_1 < 1$. Normally, one may configure the sensors so that both p_0 and q_1 are greater than 0.5. However, unforeseen circumstances may cause one or both of these numbers to decline in such a way that $p_0 + q_1$ is less than

1, unbeknownst to the operator.

Furthermore, we show that local statistical consistency in a fault tree or diagnostic tree implies global statistical consistency.

We assume that the reader is familiar with the basics of the Dempster-Shafer theory of evidence (e.g., Shafer [3], Shenoy and Shafer [4], Smets [5]).

2. STATISTICAL CONSISTENCY

Dempster's rule is particularly useful in combining uncertain evidence or evidence that is contradictory. Consider the simple tree of diagnoses in Figure 1, which refers to the valve system of the auxiliary feedwater pump in Figure 2, found in a pressurized water nuclear reactor (Martin [6]) such as the one used at Three Mile Island. The auxiliary feedwater pump system serves two purposes. It removes heat from the pressurized water in the primary coolant loop, and, as it boils, it turns a turbine to produce electricity. The auxiliary system is kept at the desired equilibrium state by a set of valves and pumps. In the simplest version, two valves (entrance valve E and throttle valve T, composing the valve system V) and a pump must be properly coordinated for the system to work. The state of the root node in a tree of diagnoses would be determined by the possible states of each leaf hypothesis. For trees of diagnoses, a subset malfunctioning implies that a superset is malfunctioning, that is, the system is working in series. Note that although hypotheses concerning states of a system are usually represented as mutually exclusive possibilities, it seems more practical here to identify the status of separate systems as hypotheses without assuming mutual exclusiveness. It is always possible, of course, to create a set of mutually exclusive (logical) combinations to represent the same information. We will do that implicitly in the application of Dempster's rule. No additional independence assumptions are made with regard to the status of the various subsystems so that an arbitrary configuration of system states can be represented.

Now consider the evidence e_t collected over time. It is assumed that the pieces of evidence collected at different time periods (or from a collection of independent sensors) are independent vectors of sensor readings. (Note that we

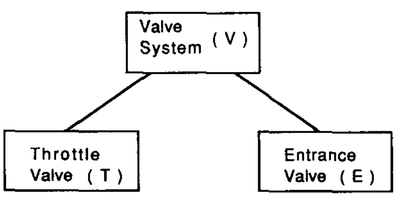

Figure 1. A tree of diagnoses.

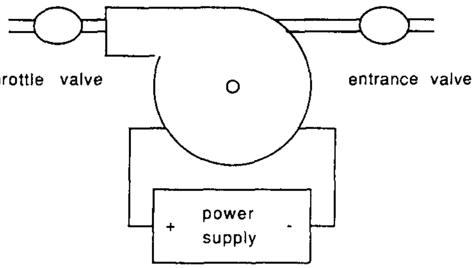

Figure 2. The auxiliary feedwater pump system.

assume only sensor reading independence, rather than subsystem independence.) The three sensor readings, $e_t = \{e_{t1}, e_{t2}, e_{t3}\}$, are taken with respect to the various subsystems. A knowledge engineer may define e_t as

$$e_{t1} = \begin{cases} 0 & \text{if sensor reading for the valve system indicates correct function} \\ 1 & \text{if the sensor reading for the valve system indicates malfunction} \end{cases}$$

$$e_{t2} = \begin{cases} 0 & \text{if the sensor reading for the throttle valve is open} \\ 1 & \text{if the sensor reading for the throttle valve is closed} \end{cases}$$

$$e_{t3} = \begin{cases} 0 & \text{if the sensor reading for the entrance valve is open} \\ 1 & \text{if the sensor reading for the entrance valve is closed} \end{cases}$$

The readings e_t do not necessarily follow the same hierarchical patterns as the possible states of the (system represented by the) tree of diagnoses. There, if either one or both subsystems are failing, then the entire system is failing. Altogether, there are four readings, (0,0,0), (1, 0, 1), (1, 1, 0), and (1, 1, 1), that constitute noncontradictory evidence corresponding to the four states of the tree. There are also four contradictory pieces of evidence: (0, 0, 1), (0, 1, 0), (1, 0, 0), and (0, 1, 1).

In a real system, the reliability and sensitivity of a sensor may degrade to an unacceptable level ($p_0 + q_1 < 1$) without warning, increasing the difficulty of reaching a correct diagnosis even further. As an important counterbalance to the likely presence of contradictory information in small samples of sensor readings, repeated observations should "average out" to give correct conclusions in the long run. This is the essence of our definition of statistical consistency. It is to be viewed as a "structural" criterion for the setup of the evidence-combining scheme, not an assertion that diagnostic trees must necessarily have unlimited access to repeated independent sensor readings.

Suppose that there are n independent pieces of evidence collected corresponding to the basic probability assignments (bpa's) $m_1, m_2, \ldots, m_n$. Then define

$$M_n = m_1 \oplus m_2 \oplus \cdots \oplus m_n$$

M_n is the bpa representing the total information computed using Dempster's rule of combination. Let h_i represent the ith hypothesis in the frame of discernment Θ, which gives the various states of the system. Define μ_i to be the degenerate bpa assigning all mass to h_i. The total information combined according to Dempster's rule of combination, M_n, is defined to be *statistically consistent* if, for every i,

$$\lim_{n \to \infty} M_n \to \mu_i, \text{ almost surely, when } h_i \text{ is the true state of nature}$$

In this formulation Θ functions as a parameter space, and the randomness of M_n is due to the imperfect nature of the sensor system. Here the reliability is the probability that the sensor indicates that the system is working correctly given that it is. The sensitivity is the probability that the sensor indicates that the system is malfunctioning given that is the case. First, define the space $\Theta = \{h_0, h_1\}$, where h_0 is the hypothesis that a component is functioning properly, and h_1 is the hypothesis that a component is malfunctioning. For the sake of motivation, it is useful to compare the following coin-tossing problem with the reliability and sensitivity of a sensor reading within an automatic monitoring system. Imagine that there are two biased coins. Let h_0 be the hypothesis that the first coin is biased heads, with $P(H \mid h_0) = p_0$ and $P(T \mid h_0) = q_0$. Let h_1 be the hypothesis that the second coin is biased tails, with $P(H \mid h_1) = p_1$ and $P(T \mid h_1) = q_1$. Therefore the reliability of a sensor reading $P(0 \mid h_0)$ is equivalent to $P(H \mid h_0)$, the probability of observing the head for the biased heads coin. The sensitivity $P(1 \mid h_1)$ is equivalent to $P(T \mid h_1)$, which is the probability of observing a tails for a biased tails coin. (See Shafer [1] and Walley [2] for a broader discussion of how to represent statistical information using the Dempster–Shafer theory of evidence.)

Now consider statistical consistency for Dempster's rule of combination. If a coin is either biased heads or biased tails, how is Dempster–Shafer theory going to determine its bias? In standard statistical practice, the procedure is to toss the coin n times, observe the number of heads, then perform an appropriate test of hypotheses. To apply Dempster–Shafer theory to determine which coin is being tossed, let

$$n = \text{total number of coin tosses}$$

$$n_0 = \text{number of coin tosses resulting in a head}$$

$$n_1 = \text{number of coin tosses resulting in a tail}$$

If the jth toss of the coin results in a head, then assign the bpa $m_j(h_0) = s_0$, where $0 < s_0 < 1$, and assign the remainder of the mass to Θ, that is, $m_j(\Theta) = 1 - s_0$. The assignment of the mass $1 - s_0$ to Θ represents the measure of uncertainty of other possible alternatives. If the jth toss is tails, then assign the bpa $m_j(h_1) = s_1$, and assign the remaining mass to Θ, that is, $m_j(\Theta) = 1 - s_1$. Stated in other terms, for each piece of evidence, a simple support function will be defined that has as its focal element the hypothesis the evidence supports.

The strong law of large numbers (SLLN) implies that, with probability 1 for a large number of independent coin tosses, the eventual number of heads appearing face side up will be approximately equal to the probability of getting a head multiplied by the number of coin tosses (Feller [7]). Similarly, the eventual number of tails will be approximately equal to the number of coin tosses multiplied by the probability of getting a tail. We will analyze the limiting distribution of the bpa when there are two states of nature h_0 and h_1 to determine the values s_0 and s_1 for which M_n is statistically consistent. We will find the region of consistency for three cases: first, the case in which h_0 is true; second, the one in which h_1 is true; finally, the case in which it is unknown which state of nature is true.

We compute the combined bpa's using Barnett's algorithm [8], which is based on the permutational invariance of Dempster's rule. The first step of the algorithm is to calculate the combined evidence for each focal element. M^0 and M^1, the bpa's based on the combined evidence for h_0 and h_1, respectively, can be calculated as

$$M^0(h_0) = 1 - (1 - s_0)^{n0}, \qquad M^0(\Theta) = (1 - s_0)^{n0}$$

$$M^1(h_1) = 1 - (1 - s_1)^{n1}, \qquad M^1(\Theta) = (1 - s_1)^{n1}$$

The next step is to calculate the total bpa $M_n = M^0 \oplus M^1$. The orthogonal sum is

$$M_n(h_0) = \frac{[1 - (1 - s_0)^{n0}](1 - s_1)^{n1}}{(1 - s_0)^{n0} + (1 - s_1)^{n1} - (1 - s_0)^{n0}(1 - s_1)^{n1}}$$

$$M_n(h_1) = \frac{(1 - s_0)^{n0}[1 - (1 - s_1)^{n1}]}{(1 - s_0)^{n0} + (1 - s_1)^{n1} - (1 - s_0)^{n0}(1 - s_1)^{n1}}$$

$$M_n(\Theta) = \frac{(1 - s_0)^{n0}(1 - s_1)^{n1}}{(1 - s_0)^{n0} + (1 - s_1)^{n1} - (1 - s_0)^{n0}(1 - s_1)^{n1}}$$

These bpa's provide statistically consistent results as long as

$$\lim M_n(h_0) = 1, \qquad \lim M_n(h_1) = 0, \qquad \lim M_n(\Theta) = 0$$

almost surely, when h_0 is the true state of nature. The region of consistency for M_n, given h_0 is the true state of nature, is

$$\frac{(1 - s_0)^{p0}}{(1 - s_1)^{q0}} < 1$$

For ease of visualization, the region of consistency is displayed as a function of $t_0 = 1 - s_0$ and $t_1 = 1 - s_1$ in Figure 3. In the second case h_1 is the true state of nature. The analysis is very similar, except now for M_n to be statistically consistent when h_1 is the true hypothesis, the limit for M_n should be (almost surely)

$$\lim M_n(h_0) = 0, \qquad \lim M_n(h_1) = 1, \qquad \lim M_n(\Theta) = 0$$

The region of consistency for M_n (cf. Figure 4), given h_1 is the true state of nature, is

$$\frac{(1 - s_0)^{p1}}{(1 - s_1)^{q1}} > 1$$

If it is unknown which is the true state of nature for the two-hypothesis case, Dempster's rule will be statistically consistent as long as the mass assigned to Θ for each simple support function lies in the intersection of the regions of both previous cases. This region of mutual consistency, described in Figure 5, is characterized by

$$(1 - s_0)^{p_0/(1-p_0)} < 1 - s_1 < (1 - s_0)^{(1-q_1)/q_1}$$

This region is clearly nonempty if and only if $p_0 > 1 - q_1$.

The preceding discussion is summarized in the following theorem.

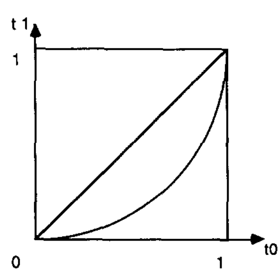

Figure 3. Region of consistency when h_0 is the true state of nature ($t_0 = 1 - s_0$, $t_1 = 1 - s_1$).

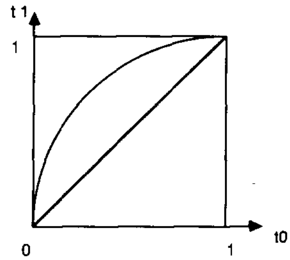

Figure 4. Region of consistency when h_1 is the true state of nature ($t_0 = 1 - s_0$, $t_1 = 1 - s_1$).

THEOREM 1 *Consider a simple two-state system with independent sensor readings indicating whether the system is functioning correctly or malfunctioning. There exist probability masses s_0, s_1 under which Dempster's rule is statistically consistent if and only if*

$$p_0 + q_1 > 1$$

where p_0 and q_1 are the sensor reliability and sensitivity, respectively.

It is quite possible that the reliability and sensitivity are not known precisely. The following is a convenient sufficient condition that applies when the reliability and sensitivity are bounded from below.

COROLLARY 1 *If the reliability p_0 and the sensitivity q_1 are both greater than 0.5, then Dempster's rule is statistically consistent for any bpa with $0 < s_0 = s_1 < 1$.*

If all we know is that $p_0 > a$ and $q_1 > b$ for known constants $0 < a$, $b < 1$, then it is still possible to use Corollary 1 along with grouped sensor observations. That is, if $p_0 > a$ and $q_1 > b$ are both guaranteed, then a series arrangement of k sensors will have reliability and sensitivity greater than 0.5

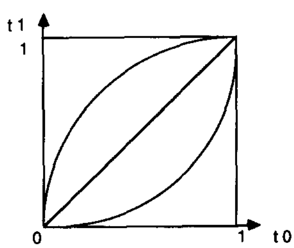

Figure 5. Region of consistency when the true state of nature is unknown ($t_0 = 1 - s_0$, $t_1 = 1 - s_1$).

as long as

$$k > \max\left[\frac{\log(1/2)}{\log(1-a)}, \frac{\log(1/2)}{\log(1-b)}\right]$$

If, on the other hand, both p_0 and q_1 are known to be less than 0.5, then it is still possible to use Corollary 1 by interchanging the interpretation of the sensor readings. (When the sensor says "OK," it means "not OK," and vice versa.)

3. STATISTICAL CONSISTENCY IN DIAGNOSTIC TREES

Dempster's rule of combination can determine, in the limit, the status of a single system given that the evidence is collected as a simple support function and bpa's are assigned within the region of mutual consistency. However, can the Dempster–Shafer theory produce statistically consistent results for a system organized in a diagnostic tree structure? In this section, we show that the answer is yes, if the bpa's for each of the subsystems are assigned within the corresponding region of mutual consistency.

Assume that each node in the tree has its own independent sensor reading to collect evidence in the form of simple support functions for the hypothesis that it confirms. In practice, we would like to combine the information derived from multiple sensor readings on each node* of the diagnostic tree to obtain a complete assessment of the state of the system or a part of it. Is Dempster's rule statistically consistent when the information from all the sensors is combined?

Consider the valve system example described in Figures 1 and 2. If we identify each node with the corresponding (sub)system, the frame of possible states for the valve system in the auxiliary feedwater pump is $\Theta = \{ET, ET^c, E^cT, E^cT^c\}$, represented by the tree of Venn diagrams in Figure 6. Let m_v, m_t, m_e be the bpa's computed by combining the evidence from n independent observations at each of the three nodes in isolation. (The dependence on n, the number of sensor readings, will not be indicated from here on.) Let us assume that all sensors are sufficiently sensitive and reliable to satisfy the conditions of Theorem 1. Then the analysis of the previous section applies for each node and

$$\lim_{n\to\infty} m_j = \mu_{ij}, \text{ almost surely, for every } i,$$

* We use "node" for the component or subsystem that it indicates.

where μ_{ij} is the degenerate bpa assigning all mass to h_{ij}, and h_{ij} indicates that node j is in state i for j in $\{v, t, e\}$. Consider now $M_v = m_v \oplus m_e \oplus m_t$, where the combination takes place on $\{V, V^c, \Theta\}$. Also, consider $M_e = m_v \oplus m_e \oplus m_t$, where the combination takes place on $\{E, E^c, \Theta\}$, and $M_t = m_v \oplus m_e \oplus m_t$, where the combination takes place on $\{T, T^c, \Theta\}$. (Combinations on coarser or more refined frames require minimal extension or projection operations, as described, for example, by Shafer et al. [9].) The total information combined on the diagnostic tree of Figure 6 is statistically consistent if

$$\lim_{n\to\infty} M_j = \mu_{ij}, \text{ almost surely, for every } i,$$

where h_{ij} is the true state of nature, for j in $\{v, t, e\}$.

Before stating and proving our result, we give three definitions.

Global bpa (gb): The bpa on the state of a subsystem or on the whole system computed by combining all available evidence in the tree using Dempster's rule

Local bpa (lb): The bpa on the state of a subsystem derived from evidence (e.g., sensor readings) pertaining to the subsystem only

Subtree global bpa (sgb): The bpa on the state of a subsystem derived from evidence (e.g., sensor readings) pertaining to the subsystem and to its components only

THEOREM 2 *Let a binary diagnostic tree be given. If the local bpa collected at each node is statistically consistent, then each global bpa is statistically consistent.*

Proof First we shall show that the global bpa at the root node n is statistically consistent for the evidence collected at the root node and the projected coarsenings of the evidence collected at each of its first-generation nodes. Second, we shall show how to compute statistically consistent sgb's for

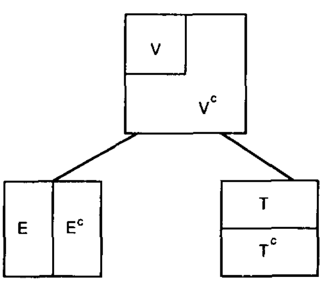

Figure 6. Diagnostic tree for the valve system.

all nodes in the fault tree. Third, we give an algorithm to compute the gb at each node of the tree and show that each gb is statistically consistent.

CLAIM 1 The global bpa at the root node n of the diagnostic tree is statistically consistent.

PROOF OF CLAIM 1 This proof is by induction in the height of the tree.

Basis. The base case is a tree of three nodes (height 1), as shown in Figure 6. We can show that the base case holds by computing M_v, M_e, and M_t and showing that their limit behaves as required by the definition of statistical consistency when the bpa's and performance parameters for every node in the tree satisfy the conditions of Theorem 1. This is a straightforward computation.

Inductive step. Take any tree of height $d + 1$, $d > 1$. Let n be the root of the tree. There are two subcases. In the first case, n has two children. (See Figure 7.) Call the children p and q. Since $d > 1$, p and q are roots of nontrivial subtrees. Call the subtrees $S1$ and $S2$. Detach $S1$ and $S2$ from the tree; the induction hypothesis is that the gb's at p and q for subtrees $S1$ and $S2$ in isolation are statistically consistent. By definition of sgb, the sgb's at p and q are statistically consistent. By the hypothesis of the theorem, the lb at n is statistically consistent. The gb at n is therefore the combination of three locally consistent belief functions in a tree of height 1. By computations totally analogous to those of the base case, gb(n) is statistically consistent.

In the second case n has only one child. (See Figure 8.) Call the child p. Since $d > 1$, p is the root of a nontrivial subtree. Call the subtree $S1$. Detach $S1$ from the tree; the induction hypothesis is that the gb at p for subtree $S1$ in isolation is statistically consistent. By the hypothesis of the theorem, the lb at n is statistically consistent. The gb at m is therefore the combination of two locally consistent belief functions in a tree of height 1. The reader can easily verify that such a combination is statistically consistent. (End of proof of claim 1.)

The inductive proof of claim 1 can be used as the specification of an algorithm to compute the global bpa at the root of the tree and the sgb's at all

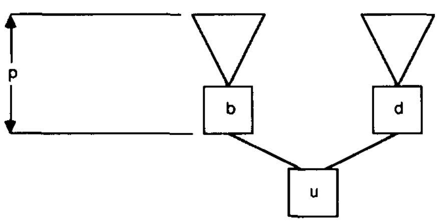

Figure 7. A tree of height $d + 1$ with two subtrees.

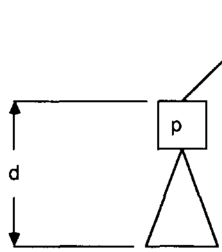

Figure 8. A tree of height $d + 1$ with one subtree.

internal nodes of the tree. Now we show how to compute the global bpa at all nodes in the tree while preserving statistical consistency.

Basis. We have already computed gb of the root node n, and it is statistically consistent.

Inductive step for the theorem. Consider the situation described in Figure 7. We have gb(n), sgb(p), and sgb(q). All of these bpa's are statistically consistent. Now we can compute gb(p) by combining gb(n) and sgb(p) using Dempster's rule. A simple computation shows that gb(p) is statistically consistent. Similarly, gb(q) is statistically consistent.

We remark that we can substitute the fault tree (Barlow et al. [10], Dempster and Kong [11]) of Figure 9 for the diagnostic tree of Figure 6 with no gain or loss of information. Similarly, the fault tree of Figure 11 corresponds to the diagnostic tree of Figure 10. (A convention concerning the relative size of complemented versus noncomplemented partitions indicates whether a node is an AND or an OR node.) As an illustration, the fault tree of Figure 11 could be interpreted as representing a power supply system (A) for the throttle valve. Assume that there are a primary (P) and a secondary (Q) power supply. When the primary supply fails, the secondary automatically takes over, so that the power supply system is faulty if and only if both the primary and secondary systems are faulty. The reader can show that Dempster's rule is statistically

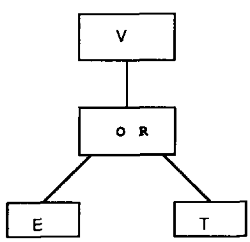

Figure 9. Fault tree equivalent to the diagnostic tree of Figure 6.

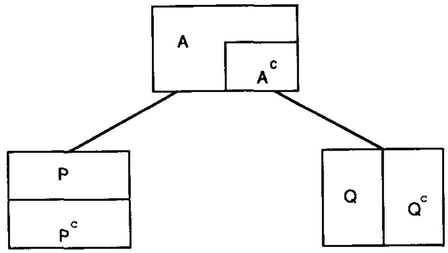

Figure 10. Diagnostic tree for the power system.

consistent on the tree of Figure 10 under the same conditions as for the tree of Figure 6. More generally, Theorem 2 holds for fault trees with AND and OR nodes and where nodes may have more than two children; that is, each (sub)system is composed of several components in series (OR) or parallel (AND). The proof of this general result is similar to that of Theorem 2.

4. CONCLUSION

We derive necessary and sufficient conditions for the statistical consistency of Dempster's rule on diagnostic trees (and fault trees). The conditions are described by a region of consistency for the degrees of belief associated with the readings "correct functioning" or "malfunctioning." This region is determined by the reliability and sensitivity of sensors and does not depend on the particular series/parallel arrangement of components. In particular, the result shows that if both the reliability and sensitivity of a sensor are greater than 0.5, then Dempster's rule will be statistically consistent for any assignment of masses $0 < s_0 = s_1 < 1$. On the other hand, if $p_0 + q_1 < 1$, then the region of mutual consistency is empty and Dempster's rule may well lead us to the wrong conclusion.

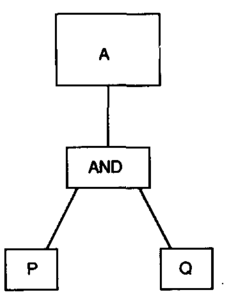

Figure 11. Fault tree equivalent to the diagnostic tree of Figure 10.

One practical implication of these results is that the sensitivity and reliability of a sensor do not need to be extremely close to 1 to determine the true state of a system. Dempster's rule of combination will provide a total bpa that can accurately determine the status of a system as long as sufficiently many independent observations can be obtained. Even if the reliability and/or sensitivity is less than 0.5, combinations of sensor readings could replace a single reading to provide values in excess of 0.5. The important thing is that (lower) bounds must be known for these quantities.

The sampling scheme used here is the most direct for our analysis. However, a variety of alternative sampling plans exist that will maintain the same balance of sensor readings as above, that is, those that obey the law of large numbers. For example, instead of a complete sample of all nodes each time, sensor data may arrive sequentially and at random, or sampling may occur according to a finite ergodic Markov chain over states consisting of nodes of the tree together with sensor readings at those nodes. Since statistical consistency is preserved on the fault tree structure, it is sufficient that the local bpa's be consistent. Thus any sampling scheme that samples each node infinitely often can be expected to provide consistency.

ACKNOWLEDGMENTS

We thank the referees for many constructive suggestions on the presentation of this material. Research was supported by U.S. Army Office Grant DAAL-03-87-K-0101 to S.D.D. and a University of South Carolina Summer Grant to M.V.

References

1. Shafer, G., Belief functions and parametric models, *J. Roy. Stat. Soc. B* **44**(3), 322–352, 1982.

2. Walley, P., Belief function representations of statistical evidence, *Ann. Stat.* **15**(4), 1439–1465, 1987.

3. Shafer, G., Perspectives on the theory and practice of belief functions, *Int. J. Approx. Reasoning* **4**, 323–362, 1990.

4. Shenoy, P. P., and Shafer, G., An axiomatic framework for Bayesian and belief function propagation, *Proceedings of the AAAI Workshop on Uncertainty in Artificial Intelligence*, 249–260, 1988.

5. Smets, P., Belief functions, in *Non-Standard Logics for Automated Reasoning* (P. Smets, E. H. Mamdani, D. Dubois, and H. Prade, Eds.), Academic, London, 254–286, 1988.

6. Martin, D., *Three Mile Island: Prologue or Epilogue?*, Ballenger, Cambridge, Mass., 1980.

7. Feller, W., *An Introduction to Probability Theory and Its Applications*, Vol. I, 3rd ed., Wiley, New York, 1968.

8. Barnett, J. A., Computational models for a mathematical theory of evidence, *Proceedings of the 7th International Joint Conference on Artificial Intelligence* (IJCAI-81), 868–875, 1981.

9. Shafer, G., Shenoy, P. P., and Mellouli, K., Propagating belief functions in qualitative Markov trees, *Int. J. Approx. Reasoning* **1**, 349–400, 1987.

10. Barlow, R. E., J. B. Fussell, and N. D. Singpurwalla (Eds.), *Reliability and Fault Tree Analysis*, SIAM, Philadelphia, 1975.

11. Dempster, A. P., and Kong, A., Uncertain evidence and artificial analysis, *J. Stat. Planning Inference* **20**, 355–368, 1988.