FAILURE PREDICTION AND DIAGNOSIS FOR SATELLITE MONITORING SYSTEMS USING BAYESIAN NETWORKS

Steven Bottone, Daniel Lee
DataPath, Inc., 13025 Danielson Street, Suite 200, Poway, CA 92064

Michael O'Sullivan
San Diego State University, Department of Mathematics and Statistics, San Diego, CA 92182

Mark Spivack
University of Cambridge, Centre for Mathematical Sciences, Wilberforce Road, Cambridge, UK CB3 0WA

ABSTRACT

Predicting failure in complex systems, such as satellite network systems, is a challenging problem. A satellite earth terminal contains many components, including high-powered amplifiers, signal converters, modems, routers, and generators, any of which may cause system failure. The ability to estimate accurately the probability of failure of any of these components, given the current state of the system, may help reduce the cost of operation. Probabilistic graphical models, in particular Bayesian networks, provide a consistent framework in which to address problems containing uncertainty and complexity. Building a Bayesian network for failure prediction in a complex system such as a satellite earth terminal requires a large quantity of data. Software monitoring systems have the potential to provide vast amounts of data related to the operating state of the satellite earth terminal. Measurable nodes of the Bayesian network correspond to states of measurable parameters in the system and unmeasurable nodes represent failure of various components. Nodes for environmental factors are also included. A description of Bayesian networks will be provided and a demonstration of inference on the Bayesian network, such as the calculation of the marginal probability of failure nodes given measurements and the maximum probability state of the system for failure diagnosis will be given. Using the data to learn local probabilities of the network will be covered. An interface between MaxView monitoring and control software and a Bayesian network API will also be described.

INTRODUCTION

Probabilistic graphical models, in particular Bayesian networks, provide a consistent framework in which to address problems containing uncertainty and complexity, such as failure prediction and diagnosis in a complex system [Ghahramani, 2001, Jensen, 2001, Neapolitan, 2004]. Probabilistic inference in high-dimensional problems only becomes tractable when the system can be made modular by imposing meaningful conditional independence assumptions [Cowell, et al., 1999]. Bayesian networks provide a natural way to accomplish this. As a combination of probability theory and graph theory, the probabilistic aspects of a graphical model provide a consistent way of connecting data to models, while graph theory provides an intuitively appealing interface to express independence assumptions as well as efficient computational algorithms [Jordan, 1998].

A Bayesian network consists of a set of variables, a graphical structure connecting the variables, and a set of local conditional probability distributions. A Bayesian network is commonly represented as a graph, which is a set of vertices and edges. The vertices, or nodes, represent the variables and the edges, or arcs, represent the conditional dependencies in the model. For a graphical model to be a Bayesian network, it must be a directed acyclic graph (DAG). The edges must be directed (the edges can be thought of as arrows) and there must be no cycles in the graph. It is often useful to think of the child nodes as being causally related to the parent nodes (the arrows are directed from parent nodes into a child node), although this does not necessarily have to be the case. In a Bayesian network, the joint probability distribution of all the nodes can be written as the product over all nodes of the conditional probability of each node given its parents. Let $V = \{V_1,\ldots,V_N\}$ be the set of N nodes comprising a Bayesian network. Then the joint probability distribution over all variables represented by the nodes in the graph is given by

$$P(V) = \prod_{i=1}^{N} P(V_i \mid \mathrm{pa}(V_i)),$$

where $\mathrm{pa}(V_i)$ is the set of all parent nodes of node V_i. This expression for the joint probability distribution can be written down by inspection of the Bayesian network. Given the joint probability distribution of a set of random variables, it is then possible, in principle, to determine all

978-1-4244-2677-5/08/$25.00 ©2008 IEEE

marginal and conditional probabilities over any subset of variables in the network.

Efficient algorithms for conducting inference on the graph, usually calculating the marginal probability of a node given evidence on some or all of the observable nodes, are available [see Jordan, 2004, for a description of various inference algorithms]. Inferences on a Bayesian network are guaranteed to be consistent, meaning that all probabilities calculated are indeed probabilities (nonnegative numbers summing to one). In a typical problem arising in failure prediction, there would be nodes on the graph representing the failures of the system as a whole and of each device, as well as nodes representing measurements obtained from each device.

There are three main problems associated with Bayesian networks [see Heckerman, 1998 and Neapolitan, 2004]: 1) Inference: given a model, to compute marginal probabilities on unobserved nodes, given evidence on some subset of other nodes. 2) Learning probabilities: given a model and some data, to estimate the unknown parameters for the local conditional probabilities. 3) Learning structure: given data, to estimate the unknown structure of the graph as well as the underlying local probabilities. These problems are listed in increasing order of difficulty. The remainder of this paper will concern itself mostly with inference on a Bayesian network for a specific monitoring and control system for a satellite earth terminal. We also briefly discuss parameter estimation for the local probabilities. The graphical structure is determined using expert knowledge and physical reasoning. Heckerman [1998] and Neapolitan [2004] discuss constructing Bayesian networks using data only.

SATELLITE EARTH TERMINALS

The basic building blocks of a typical satellite earth terminal are shown in Figure 1 [Inetdaemon, 2007]. The satellite system consists of two series of devices, called the *uplink chain* and the *downlink chain*. The uplink chain, shown on the left side of Figure 1, consists of the sequence of components that produces a radio frequency signal from data—typically digital—to be transmitted to the satellite. The description given here is not precise, as the exact configuration of a specific system can vary widely. The downlink chain, which refers to those components that convert the incoming radio frequency signal from a satellite to (digital) data, uses almost the same series of equipment as the uplink chain, except in reverse order. The downlink chain is shown on the right side of Figure 1.

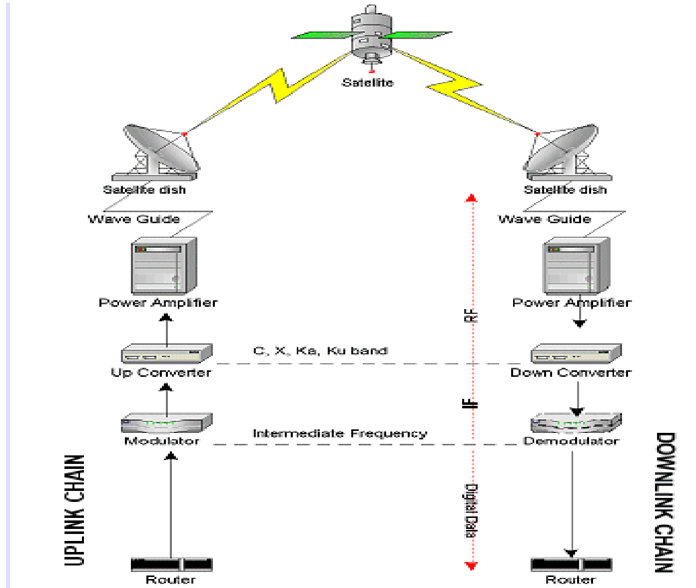

Figure 1. Basic components of a satellite earth terminal link.

A digital satellite uplink chain works as follows:

1. Digital data is sent to the modulator (often the same modem serves as both modulator and demodulator) where it is converted into a modulated signal at an intermediate frequency in the L band, centered on about 70–140 MHz.
2. The intermediate frequency signal is then sent to an upconverter, usually through a shielded coaxial cable, where the frequency of the signal is converted to the higher frequency that will eventually be transmitted to the satellite. These microwave frequencies are typically in the C, S, X, Ka, or Ku bands (frequencies above about 1000 MHz). Sometimes a block upconverter is used, which can accept a block of input frequencies that are then converted to the required carrier frequency.
3. Noise is removed from the converted signal using a band pass filter or similar means and is then sent to a high-powered amplifier (HPA), such as a traveling wave tube amplifier (TWTA), a klystron, or a solid state amplifier, where the signal is amplified to power levels strong enough for transmission to the satellite.
4. The final clean, amplified signal is sent down a waveguide to the satellite antenna dish.
5. The feed horn at the focal point of the dish emits the high-frequency radio transmission, which the dish focuses in the direction of the satellite. The direction of the dish is controlled by the antenna control unit (ACU).

The steps in the downlink chain are essentially the reverse of the uplink chain:

1. The satellite transmits a signal that contains the encoded data.
2. The signal is received at the satellite antenna dish.
3. The signal is amplified through a low noise amplifier (LNA) and fed to the downconverter.
4. The downconverter converts the high-frequency signal to an intermediate frequency. Sometimes the LNA and downconverter are combined into one device called a low noise block amplifier (LNB).
5. The intermediate frequency is fed into a demodulator and converted into a digital data signal.
6. The data is sent to the network through a router.

MONITORING AND CONTROL OF SATELLITE EARTH TERMINALS

It is possible to monitor a complicated network, like the satellite earth terminals just described, using appropriate software, such as MaxView network monitoring and control software available from DataPath, Inc. [DataPath, 2008]. Figure 2 shows a MaxView screenshot for the monitoring of some of the equipment in the uplink chain for a typical satellite earth station. Monitoring and control software allows the user to systematically monitor the state of each piece of equipment in both the uplink and downlink chains. Potentially, a wealth of data can be obtained that can assist in predictive maintenance and diagnostics of the system. The software can not only monitor various measurements observed on each piece of equipment in the system, including all faults and alarms, it can also record the time history of each measurement on each piece of equipment in every system that has been fielded. The amount of data that can be collected is potentially much more extensive than what can be obtained through typical reliability tests performed in the lab.

For example, a typical traveling wave tube high-power amplifier includes the following monitoring points:

Helix Voltage (kV)
Helix Current (mA)
Cathode Filament Current (A)
Tube Temp (C)
Power Supply Temp (C)

The TWTA consists of an elongated vacuum tube with a heated cathode that emits electrons. There is also a helix coil wrapped around the beam. Induced currents in this coil produce the amplified RF signal. Often, when the cathode filament is about to burn out, the cathode current shows a noticeable drop over time. Similarly, when the helix current begins to rise, it often indicates an imminent failure. These precursors have been observed on DataPath earth terminals. With enough data, estimates of the probability of failure of the HPA given the states of some of these measurable variables can be obtained. These estimated probabilities can be used in a Bayesian network for failure prediction and diagnosis of the satellite earth terminal system.

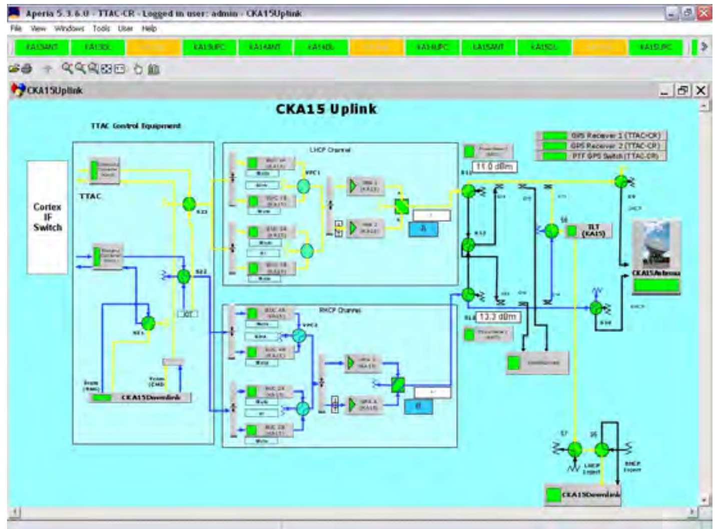

Figure 2. Monitoring the equipment in a satellite earth terminal.

BAYESIAN NETWORK FOR SATELLITE EARTH TERMINALS

A sample Bayesian network for failure prediction and diagnostics for a somewhat simplified satellite earth terminal system is shown in Figure 3. This Bayesian network represents the uplink chain of the system. Similar nodes can be added for the downlink chain, as well as any additional nodes for redundant components, such as the generator redundancy shown in the upper right of the figure. The general design of this particular Bayesian network has been developed using expert knowledge and physically reasonable assumptions. No attempt to use data-only methods to determine the graphical structure has been tried. Most links between nodes have a reasonable causal connection and each individual component has been treated as physically independent of the others.

For simplicity, each node on this Bayesian network is a discrete chance node, with a finite number of states. We have used two types of nodes in this network, failure nodes and measurement nodes. There are three types of failure nodes: component failure nodes, background failure nodes, and “intermediate failure” nodes. Each failure node has two states, which can be labeled *fail* and *no fail*, or *yes* and *no*. The probability of failure, that is, the probability of the state labeled *fail*, is the probability that the component in question, or the system as a whole for the "System Failure" node, will fail over a *fixed period of time*, say one week or one month. Clearly, the probability that the component does not fail in the fixed period of time is one minus the probability of failure. These probabilities depend on the evidence entered on the measurement nodes. The measurement nodes include the helix and cathode currents, trends on these currents, and the temperature nodes. If for some reason these measurements are not available, then prior probabilities on the states of these nodes will be used when propagating evidence in the network. Notice that the measurement node for a particular piece of equipment is a parent only of the failure node for that piece of equipment. The background failure nodes allow for failure due to unmeasurable or unknown random events and are typically are not instantiated (a node is *instantiated* when evidence becomes known that it is definitely in one of its possible states). The nodes labeled "Helix Current Failure" and the like are put in for calculational convenience. They facilitate the calculation of what is called the "noisy or" [see Jensen, 2001], which is discussed below.

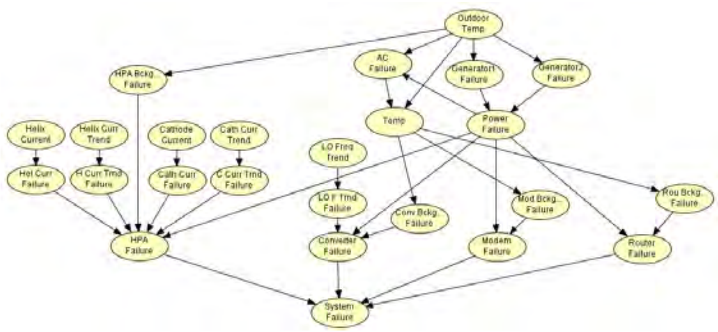

Figure 3. Bayesian network for failure prediction and diagnosis for a satellite earth terminal.

The "System Failure" node has four parent nodes, corresponding to the failure of the individual components for the HPA, upconverter, modem, and router. For these four components, the system will fail if any of the components fail. Therefore, the state *no fail* can only occur if the state of all individual component failure nodes is *no fail*. The conditional probability distribution for the "System Failure" node is then:

$$P(SF\text{=}no\ fail \mid HF\text{=}no\ fail, CF\text{=}no\ fail, MF\text{=}no\ fail, RF\text{=}no\ fail) = 1$$

$$P(SF\text{=}no\ fail \mid \text{otherwise}) = 0$$

The probability $P(SF\text{=}fail \mid HF,CF,MF,RF)$ is clearly $1 - P(SF\text{=}no\ fail \mid HF,CF,MF,RF)$, since there are only two failure states.

The subgraph that includes the "HPA Failure" node implements the "noisy or" condition. The assumption here is that any of the causes of the HPA failing are independent of one another. Therefore, if, for example, the probability is p_1 that the HPA will fail given a high value of the helix current independent of knowledge of the other possible causal factors, and p_2 is the probability that it will fail given a trend in the cathode current, then the noisy or assumption is that the probability of failure given that both have occurred is simply $1 - (1 - p_1)(1 - p_2)$. A similar argument holds for any number of independent possible causes of failure. The configuration that includes the intermediate failure nodes implements this condition. One only needs to know the probability of failure for each cause separately. Without these intermediate nodes, one would have to calculate the entire probability distribution table for all combinations of states of each parent node. The uninstantiated background failure node assures that the probability of failure, given that no known causes have been observed, is a constant.

The subgraph of nodes in the upper right of Figure 3 can be thought of as the environmental nodes. The states of these nodes potentially affect each component in the system. The node labeled "Temp" is the indoor temperature of the system container and directly affects all components that are in the container through their background failure nodes (high temperatures do not affect in a known way the measurable quantities that might lead to failure). Since the high-power amplifiers and generators are outside, only the "Outdoor Temp" is an ancestor for these nodes. There are two generators, a main one and a backup. Both must fail for the power to fail. The generators are more likely to fail when the outdoor temperature is hot. The indoor temperature depends on the outdoor temperature and whether the air conditioning unit is working. The state of the air conditioner depends on the how hot it is outside and whether the power is working.

ESTIMATION OF PARAMETERS

The problem of estimating the value of the parameters for the local conditional probability distributions in a Bayesian network, also referred to as learning the probabilities or parameter learning, is relatively straightforward if enough complete data samples, sometimes called cases, are available. If some of the data samples are incomplete, or if there are some completely unobserved nodes, then this task becomes much more difficult [Heckerman, 1998]. Another difficulty that arises is when there are insufficient data samples to cover all possible combinations of parent states to estimate all of the parameters in some of the conditional probability tables. When there is a lack of data, one often has to resort to using reasonable physical assumptions and prior knowledge to estimate the probabilities, or other subjective methods.

For example, the node labeled "HPA Failure" in Figure 3 has six parent nodes, corresponding to the value of the helix current, the trend in the helix current, the value of the cathode current, the trend in the cathode current, and a power failure and background failure node. Assuming for the moment that there are no intermediate failure nodes and that each of the parent nodes has only two states, then the conditional probability table would have $2^6 = 64$ independent parameters to estimate. For instance, to estimate the probability that the HPA would fail given that the helix current is high and the cathode current is trending and that the other four nodes are in states that usually do not lead to failure, we would have to observe the frequency that the HPA fails under these exact conditions certainly over 10 times to even begin to get a good estimate of this parameter. However, it is very unlikely that the HPA fails because just these two conditions are present, so it is extremely unlikely we ever observe enough cases where these precise conditions are the case. So the assumption we make, in order to fill out all 64 independent entries in the table, is that any known cause of failure for this particular piece of equipment is independent of any other cause. This is the "noisy or" assumption, and it is incorporated into the Bayesian network by using the intermediate failure nodes. To fill out the table, we need only to estimate the probability that the HPA fails given any one of the known possible causes. For example, if we have observed the helix current being high 100 times and the HPA failed 50 of those times in the next month, then we would set the probability that the HPA fails given the current is high and all other nodes are normal to be 0.5 (we might change this a bit if we had differing prior information on this occurrence from some other source, such as a lab test). Similarly, if we observed that the HPA fails 20% of the time when the cathode current is trending, we would set

$$P(HF{=}fail \mid HC{=}normal, HCT{=}normal, CCF{=}normal, CCT{=}yes, BF{=}no, PF{=}no) = 0.2$$

The noisy or assumption now allows us to set the probability that the HPA fails given both the helix current is high and the cathode current is trending by assuming that the probability that the HPA does not fail is just the product of the probabilities that each condition does not cause a failure. That is, if

$$p_1 = P(HF{=}fail \mid HC{=}high, HCT{=}normal, CCF{=}normal, CCT{=}no, BF{=}no, PF{=}no)$$

and

$$p_2 = P(HF{=}fail \mid HC{=}normal, HCT{=}normal, CCF{=}normal, CCT{=}yes, BF{=}no, PF{=}no)$$

then

$$P(HF{=}no \mid HC{=}high, HCT{=}normal, CCF{=}normal, CCT{=}yes, BF{=}no, PF{=}no) = (1-p_1)(1-p_2)$$

And, therefore,

$$P(HF{=}fail \mid HC{=}high, HCT{=}normal, CCF{=}normal, CCT{=}yes, BF{=}no, PF{=}no) = 1 - (1-p_1)(1-p_2)$$

A similar calculation can be made to fill in the rest of the table, although it gets slightly more complicated when considering the case when more than two possible modes of failure have occurred. This complication in filling out a noisy or table can be avoided by placing intermediate failure nodes, as described above and shown in Figure 3, between the measurable nodes and the component failure node. The conditional probability table for the "Helix Current Failure" node, for example, is given by

$$P(HCF{=}fail \mid HC{=}normal) = 0$$
$$P(HCF{=}fail \mid HC{=}high) = 0.5$$
$$P(HCF{=}no \mid HC{=}normal) = 1$$
$$P(HCF{=}no \mid HC{=}high) = 0.5$$

and similarly for the other intermediate failure nodes.

As mentioned earlier, the Bayesian network in Figure 3 computes the probability of failure over a fixed period of time, say one month. Failure data, however, is usually not given in terms of a number of failures over a fixed period of time. Rather, the data is usually in the form of the times to failure from a given start time for a number of presumably identical components. Often there are components that have not failed up to the end of the testing period. There are well-known methods in the theory of reliability engineering that can be used to estimate the probability of failure over a fixed period of time, given a set of data consisting of times of failure [Lawless, 2002]. Typically, it is assumed that the failure times are samples from a lifetime distribution, such as a Weibull or lognormal distribution [ReliaSoft, 2005], with unknown parameters to be determined from data. In principle, there would be a different distribution (it may be the same type of distribution with a different set of parameters) for each parent state of each failure node. Once the parameters have been estimated using sample data for an assumed distribution type, it is then a straightforward matter of determining the probability of failure over a fixed period of time, given that the piece of equipment under consideration has not yet failed. In the general case, this probability of failure will depend on how long the equipment has run without a failure, so this probability is a function of time. Most components wear out over time and so the probability of failure over the next fixed period of time is generally an increasing function of time. The Bayesian network would then have conditional probability tables whose parameters change with time, reflecting the characteristics of the particular distribution of the component lifetimes.

PROPAGATION OF PROBABILITIES IN THE BAYESIAN NETWORK

As an example of how probabilities are propagated in a Bayesian network, let us assume that we have observed that the indoor temperature is hot, so the state of the "Temp" node is *hot* (we have assumed for simplicity that the "Temp" node has two states, *normal* and *hot*). If no nodes are instantiated, then the marginal probability that the system will fail in the next month is 12.7% and the probability that the HPA will fail is 6.4%. (These probabilities are calculated using Hugin Expert graphical modeling software [Hugin, 2007]. The conditional probability tables have been estimated using available field data. The topic of parameter estimation for failure prediction models was briefly discussed above.) Although the indoor temperature is not a direct cause to the HPA failing (the "Temp" node is not an ancestor of the "HPA Failure" node), we nevertheless expect the probability of HPA failure to increase if the temperature goes up in the system container. The Bayesian network will quantify this change in probabilities. When the evidence that the indoor temperature is hot is entered into the Bayesian network and the evidence is propagated, the marginal probability that the HPA will fail in the next month increases to 11.2% and the probability that the system will fail increases to 21.7%. If the indoor temperature is hot, then the probability that the outdoor temperature is hot and that the AC has failed has increased. If the outdoor temperature is more likely to have risen, then the probability that there is a background failure due to the high outdoor temperature increases, which increases the likelihood of an HPA failure. In addition, if it is more likely that it is hot outside, then it is more likely that one or both generators will fail, which increases the probability that the power has failed, which then increases the likelihood that the HPA will fail. These two possible ways contribute to the increased probability of failure of the HPA, knowing just that the indoor temperature is hot. The Bayesian network quantifies all these contributions to the calculation of the marginal probability of the "HPA Failure" node. If the probability that the HPA will fail in the next month is high enough, then a decision can be made whether it is best to replace the part before the failure occurs.

DIAGNOSIS

If the system actually fails, it would be useful to know which component is most likely to have failed. One way to diagnose which component is the one to have most likely failed is to compute the most probable configuration of all uninstantiated nodes. The most probable configuration determines the most likely state of each uninstantiated node. If the system has failed, then the evidence that the "System Failed" node is in state *fail* can be entered into the Bayesian network. Once the most probable configuration of the network has been found, it is often the case that just one component has *fail* as its most likely state. It may be the case that no components have *fail* as its most likely state. In that case one can still choose the most likely component to have failed by looking at how likely are the relative likelihoods of the *fail* and *no fail* states of each component. It is also possible that more than one component has *fail* as its most likely state. In this case one can use the relative likelihoods to choose the most likely component to have failed. In a realistic situation, the results serve as a guide for either choosing an order of replacing the pieces of equipment or checking the equipment that can be fixed. If there is a greater cost for replacing one part compared to another part, then the order of replacement may change. One would then attempt to maximize the utility.

As a simple example, let us say that we know it is hot outdoors and have entered the evidence that the "Outdoor Temp" node is *hot*. Now, if the system fails, we also enter the evidence that the "System Failed" node is in state *fail*. We can calculate the most probable configuration using the "Max Propagation" feature in Hugin Expert. It gives that the most likely state of the "HPA Failure" node is *fail*, while the most likely state of the other three component failure nodes is *no fail*. This result can then serve as a guide to determining which component has actually failed.

MAXVIEW – BAYESIAN NETWORK INTERFACE

Datapath, Inc. has developed software that provides an interface between the MaxView monitoring and control software and a probabilistic model, such as the API provided for Hugin Bayesian networks. A top-level diagram of this interface is shown in Figure 4. MaxVew collects various faults, alarms, and operating data from all devices comprising one or more satellite earth terminals. The interface consolidates this information in the form of the inputs expected by the probabilistic model and then sends these inputs to the model. Results of the model, usually in the form of probability distributions of the unobserved variables, are made available to the MaxView graphical user interface (GUI) to be used in monitoring the health of the system.

A more detailed diagram of the interface is shown in Figure 5. Each variable monitored by MaxView has a unique identifier. A time history of each of these variables is collected by the Probabilistic Model Client. An XML file called the Device Model Map provides a correspondence between the unique identifier for each variable and the corresponding measurement node in the probabilistic model (Bayesian network). A tool has been built with which this correspondence can be made using a GUI (this is indicated by the dashed arrow from the model box to the Device Model Map box).

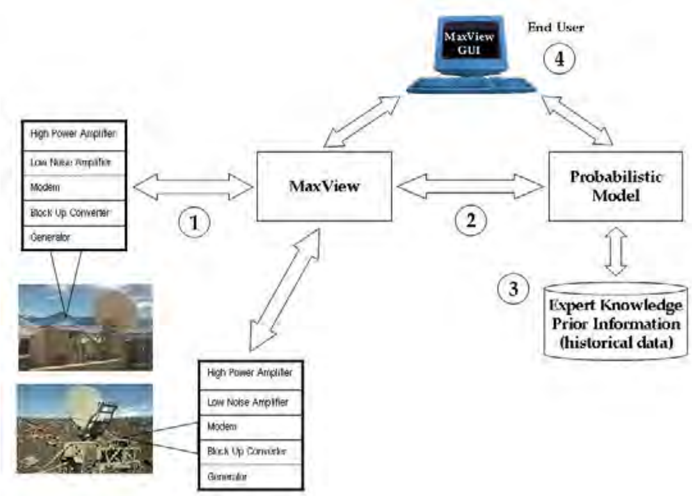

Figure 4. Interface between MaxView monitoring and control software and the Bayesian network.

Some of the inputs to the model are not simply the current value of one of the measured variables, but rather some function of the time history of a variable. The trend nodes in Figure 3 are examples of this. The existence of a significant trend is a function of the time history of the variable. The number of values of the time history to use in computing the trend depends two quantities: 1) the time span for which the trend is considered relevant and 2) the significance level of test. This information is provided by an XML file called the Probabilistic Model Configuration.

CONCLUSION

Monitoring and control software, such as MaxView, has the potential of accumulating and storing vast amounts of data representing the time-varying state of extremely complex systems. This wealth of data can be exploited to assist in the building of models, like the Bayesian networks discussed here, which may be used for such applications as failure prediction and diagnosis and other inference problems.

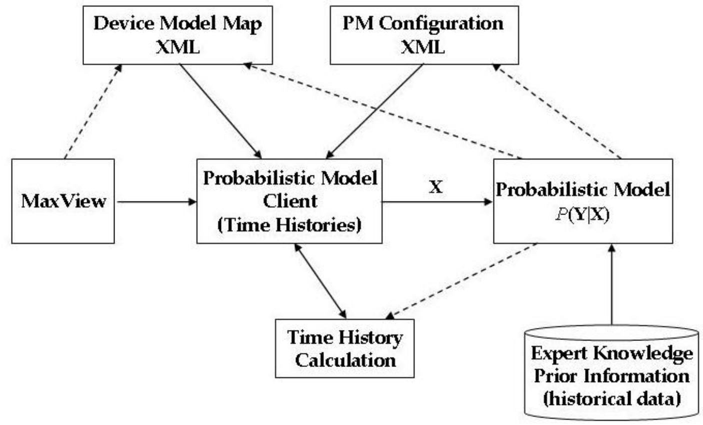

Figure 5. A more detailed diagram of the MaxView – Bayesian network interface.

REFERENCES

1. Cowell, R. G., A. P. David, S. L. Lauritzen, and D. J. Spiegelhalter, *Probabilistic Networks and Expert Systems*, Springer, 1999.
2. DataPath, www.datapath.com/ maxviewsoftwaresolutions-home, 2008.
3. Ghahramani, Z., An Introduction to Hidden Markov Models and Bayesian Networks, *International Journal of Pattern Recognition and Artificial Intelligence*, **15**(1):9–42, 2001.
4. Heckerman, D., A Tutorial on Learning With Bayesian Networks, in *Learning in Graphical Models*, M. I. Jordan, editor, MIT Press, 1998.
5. Hugin, www.hugin.com, 2007.
6. Inetdaemon, www.inetdaemon.com/tutorials/telecom/ satellite/uplink_chain.shtml, 2007.
7. Jensen, F. V., *Bayesian Networks and Decision Graphs*, Springer, 2001.
8. Jordan, M. I., Graphical Models, *Statistical Science*, Vol. 19, No. 1, 140–155, 2004.
9. Jordan, M. I., Preface, in *Learning in Graphical Models*, M. I. Jordan, editor, MIT Press, 1998.
10. Kjærulff, U. B. and A. L. Madsen, *Probabilistic Networks for Practitioners — A guide to Construction and Analysis of Bayesian Networks and Influence Diagrams*, Hugin Expert, unpublished, 2006.
11. Lawless, J. F., *Statistical Models and Methods for Lifetime Data*, Wiley-Interscience, 2002.
12. Neapolitan, R. E., *Learning Bayesian Networks*, Pearson Prentice Hall, 2004.
13. ReliaSoft, *Life Data Analysis Reference*, ReliaSoft publishing, 2005.