

Developing Large-Scale Bayesian Networks by Composition: Fault Diagnosis of Electrical Power Systems in Aircraft and Spacecraft

Ole J. Mengshoel

Intelligent Systems Division
NASA Ames Research Center
Moffett Field, CA 94035
Phone: (650) 604-4199
Ole.J.Mengshoel@nasa.gov

Scott Poll

Intelligent Systems Division
NASA Ames Research Center
Moffett Field, CA 94035
Phone: (650) 604-2143
Scott.Poll@nasa.gov

Tolga Kurtoglu

Intelligent Systems Division
NASA Ames Research Center
Moffett Field, CA 94035
Phone: (650) 604-1738
Tolga.Kurtoglu-1@nasa.gov

Abstract

In this paper, we investigate the use of Bayesian networks to construct large-scale diagnostic systems. In particular, we present a novel analytical and experimental approach to developing large-scale Bayesian networks by composition. This compositional approach reflects how (often redundant) subsystems are architected to form systems such as electrical power systems. We develop Bayesian networks and clique trees representing 24 different electrical power systems, including the real-world electrical power systems ADAPT. ADAPT is representative of electrical power systems deployed in aerospace, and is located at the NASA Ames Research Center. The largest among these 24 Bayesian networks contains over 1,000 random variables. Related work has used Bayesian networks to diagnose specific electrical power systems, however we are not aware of previous research that investigates a wide range of distinct electrical power systems as is done in this paper. While we consider diagnosis of power systems specifically, we believe this work has application to numerous health management problems, in particular in dependable systems such as aircraft and spacecraft.

Keywords: Bayesian networks; clique tree clustering; uncertainty; model-based diagnosis; knowledge engineering; electrical power systems; real-time systems; domain modeling; scalability; composition.

Track: Emerging Application or Methodologies.

Designation of the application domain(s): Aircraft; spacecraft; real-time systems; electrical power systems.

Identification of AI techniques employed or issues addressed: We investigate model-based diagnosis using probabilistic techniques. Specifically, we discuss the use of Bayesian networks to perform diagnosis and health management in electrical power systems in aircraft and spacecraft. We identify important issues that arise in engineering diagnostic applications in this area, namely the *modelling challenge*, the *real-time reasoning challenge*, and the *scalability challenge*. The *modelling challenge* concerns how to model

an EPS by means of Bayesian networks. The *real-time reasoning challenge* is associated with the embedding of AI components, including diagnostic reasoners, into hard real-time systems. The *scalability challenge*, which is the main focus of this paper, is concerned with how different EPSs, with varying architectures and components, can be represented in varying Bayesian networks for diagnosis, and how space requirements and computation times vary accordingly. The modelling challenge has been addressed by a high-level specification language from which Bayesian networks are auto-generated; the real-time reasoning challenge by off-line compilation of Bayesian networks into clique trees or arithmetic circuits which are used on-line. The scalability challenge is addressed by means of composition, in other words by considering the subsystems making up a system. For example, a electrical power systems can be made up by power storage and power distribution subsystems. We provide in this paper several novel analytical and experimental results that shed light on large-scale BNs developed for electrical power system diagnosis. The experimental part includes results for Bayesian networks and clique trees representing 24 different electrical power systems, including the Bayesian network and clique tree models of ADAPT.

Indication of application status (e.g., feasibility analysis, research prototype, operational prototype, deployed application, etc.): We discuss the development of diagnostic applications for 24 different electrical power systems, including the Advanced Diagnostics and Prognostics Testbed (ADAPT) (see also <http://ti.arc.nasa.gov/adapt/>). ADAPT, which has capabilities for power generation, power storage, and power distribution, is a fully operational electrical power system that is representative of such systems in aircraft and spacecraft. We have developed a diagnostic application that is an operational prototype working on real-world data from ADAPT. This paper takes the next step by considering how what we have learned from ADAPT can be applied to electrical power systems that are similar to ADAPT, but of different sizes and structures. Specifically, we discuss the composition of electrical power systems from power storage and power distribution subsystems, and how this composition is reflected in the Bayesian network and clique tree models of these EPSs.

Introduction

This paper is concerned with efficient probabilistic reasoning and diagnosis in particular. Our approach is based on developing a Bayesian network (Pearl 1988) model of a system, and then using it to efficiently compute answers to probabilistic queries. Bayesian networks and their inference engines provide a well-established approach to model-based diagnosis and monitoring (Lerner *et al.* 2000; Chien *et al.* 2002; Yongli *et al.* 2006; Mengshoel *et al.* 2008).

We focus on NASA-relevant research problems that represent challenges in aircraft and spacecraft health management. In this paper, we take as our point of departure an electrical power system known as the Advanced Diagnostics and Prognostics Testbed (ADAPT). ADAPT is an electrical power system (EPS) developed at NASA Ames for supporting the development of diagnostic and prognostic models; for evaluating advanced warning systems; and for testing diagnostic tools and algorithms (Poll *et al.* 2007). ADAPT is representative of electrical power systems deployed in aerospace vehicles.

Progress in probabilistic model-based diagnosis is stimulated by real-world applications, and EPSs raise several challenges including the following: (1) The challenge of developing models that are capable of handling 100s or 1000s of different faults, many of which may occur at the same time; (2) The challenge of real-time diagnostic computing, especially on on-board avionics systems with limited processor and memory capacity; (3) The challenge of developing BNs (and in particular large-scale BNs) for a wide spectrum of EPS sizes while obtaining high performance.

To start addressing these challenges, we have developed a probabilistic approach to model-based diagnosis for ADAPT (Mengshoel *et al.* 2008). Our approach is based on developing Bayesian network models of aerospace vehicles or sub-systems of such vehicles. These models typically represent the health state of sensors and other system components explicitly by means of random variables. We have paid special attention to meeting two of the main challenges – (1) model development and (2) real-time reasoning – often associated with real-world application of model-based diagnosis technologies (Musliner *et al.* 1995; Mengshoel 2007a). To address the challenge of model development, we have developed a systematic approach to representing electrical power systems as Bayesian networks, supported by an easy-to-use specification language. To address the real-time reasoning challenge, we compile BNs into arithmetic circuits or clique trees. The evaluation of arithmetic circuits and clique trees supports real-time diagnosis by being predictable and fast. In experiments with the ADAPT BN, which currently contains 503 discrete nodes and 579 edges, the time taken to exactly compute the most probable explanation using an arithmetic circuits or a clique tree was in the order of 1-10 milliseconds.

Building on earlier results as discussed above, this paper investigates scalability issues associated with probabilistic methods and technologies. In particular, we consider challenge (3) above, and present a novel analytical and experimental approach to developing large-scale BNs by compo-

sition. This compositional approach reflects how (often redundant) subsystems are architected to form systems such as EPSs. Specifically, we consider 24 different EPS architectures including ADAPT, formed by the duplication and integration of a varying number of power storage and power distribution subsystems. Previous work has used BNs to diagnose specific EPSs (Chien *et al.* 2002; Yongli *et al.* 2006; Mengshoel *et al.* 2008), however we are not aware of other efforts that consider a range of distinct EPSs. While we consider EPS health management specifically, the work has application to numerous health management problems, including such problems in aircraft and spacecraft.

The remainder of this paper is structured as follows. Concepts related to Bayesian network are presented first, followed by a discussion of EPSs. We then present our scalability analysis and an EPS case study. We finally report on experimental results for 24 different EPSs including ADAPT, conclude, and outline future work.

Preliminaries

A Bayesian network (BN) structures a multi-variate probability distribution by using a directed acyclic graph (DAG). Our main emphasis will be on discrete BN nodes. A (discrete) BN node V is a discrete random variable with a mutually exclusive, exhaustive, and finite state space $\Omega_V = \Omega(V) = \{v_1, \dots, v_m\}$. We use the notation Π_V for the parents of a node V , Ψ_V for the children of V , and π_V for an instantiation of all parents Π_V of V . The notion of a Bayesian network can now be introduced (Pearl 1988).

Definition 1 (Bayesian network) *A Bayesian network is a tuple $(\mathbf{V}, \mathbf{E}, \mathbf{P})$, where (\mathbf{V}, \mathbf{E}) is a DAG with nodes $\mathbf{V} = \{V_1, \dots, V_n\}$, edges $\mathbf{E} = \{V_1, \dots, V_m\}$, and where $\mathbf{P} = \{\Pr(V_1 | \Pi_{V_1}), \dots, \Pr(V_n | \Pi_{V_n})\}$ is a set of conditional probability tables (CPTs). For each node $V_i \in \mathbf{V}$ there is one CPT, which defines a conditional probability distribution $\Pr(V_i | \Pi_{V_i})$.*

The independence assumptions induced by (\mathbf{V}, \mathbf{E}) in Definition 1 imply the following joint distribution:

$$\Pr(\mathbf{v}) = \Pr(V_1 = v_1, \dots, V_n = v_n) = \prod_{i=1}^n \Pr(v_i | \pi_{V_i}), \quad (1)$$

where $\Pi_{V_i} \subset \{V_{i+1}, \dots, V_n\}$.

A BN can be provided *evidence* by setting or clamping some variables to known states. These nodes are called *evidence variables*. Taking into account the input on evidence variables, different probabilistic queries can be answered (Pearl 1988). These probabilistic queries include marginals, most probable explanation (MPE), and maximum a posteriori probability (MAP). While probabilistic queries can be used for many purposes, our focus in this paper is on diagnosis, where we query health variables representing the health of components, sensors, or both (Mengshoel *et al.* 2008).

It has been shown that exact MPE computation is NP-hard (Shimony 1994), and approximating an MPE to within a constant ratio-bound has also been proven to be NP-hard (Abdelbar and Hedetnieme 1998). There are two broad

classes of approaches to Bayesian network inference: Interpretation and compilation. In interpretation approaches, a Bayesian network is directly used for inference. In compilation approaches, such as the clique tree (Lauritzen and Spiegelhalter 1988; Shenoy 1989) and arithmetic circuit (Darwiche 2003; Chavira and Darwiche 2007) approaches, a Bayesian network is off-line compiled into a secondary data structure, and this secondary data structure is then used for on-line inference. In clique tree clustering, inference consists of propagation in a clique tree compiled from a Bayesian network. In arithmetic circuit evaluation, inference is performed in an arithmetic circuit that was compiled from a Bayesian network. Computation time depends on a number of structural and numerical factors associated with a BN and is not yet, despite recent progress, sufficiently understood.

Electrical Power Systems and ADAPT

Electrical power systems (EPSs) are crucial systems in aircraft and spacecraft (Button and Chicatelli 2005; Poll *et al.* 2007), and ADAPT has been developed to investigate health management technologies in a real-world setting. In this paper, we investigate ADAPT's power storage and distribution subsystems. Over a hundred sensors report their measurements to a health management system that monitors the status of the EPS. Typical sensor measurements of system variables include voltages, currents, temperatures, and switch positions. The ADAPT test bed provides a controlled environment to inject failures in a repeatable manner, and this makes it ideal for use in experiments with our novel techniques and models.

The physical hardware of the ADAPT EPS consists of battery chargers, batteries, relays, circuit breakers, inverters, wires, sensors, and loads. Most of the hardware is contained within equipment racks or cabinets, with the exception of the loads which are placed in the surrounding lab area. Three batteries may be interchangeably connected to two load banks. Each load bank can connect up to 6 alternating current (ac) loads and 2 direct current (dc) loads. The locations of the loads with respect to the load bank connection points are fixed for the purposes of any given experiment. Different configurations/modes of the EPS are commanded by opening and closing different combinations of relays between the batteries and the loads. As a consequence, ADAPT's system behavior is hybrid, consisting of discrete configuration changes and continuous behavior within the modes.

Composition and Scalability Analysis

We have developed a multi-variate Bayesian network model of the ADAPT EPS, currently containing over 500 random variables including over 100 health variables, where the health variables include components and sensors (Mengshoel *et al.* 2008). This BN supports the diagnosis of multiple sensor and/or component faults. Experiments in the ADAPT testbed have showed strong performance on scenarios with multiple faults as well as very fast and predictable inference times. We now consider the scalability

over a range of BNs representing different EPSs, including the ADAPT BN as described above as one data point.

Scalability, in terms of space requirement and computation time for clique tree evaluation, is determined by clique tree size (Lauritzen and Spiegelhalter 1988).

Definition 2 (Clique tree size) *Let Γ be the set of cliques in a clique tree compiled from a BN β . The (total) clique tree size is defined as*

$$\tau(\Gamma) = \sum_{\gamma \in \Gamma} \prod_{X \in \gamma} |\Omega_X|. \quad (2)$$

A number of interacting factors determine the number of cliques and the size of each clique in (2); we now discuss a few of them.

The Subsystem (or Composition) Factor: Suppose that we consider an EPS as a system that might be part of a larger system-of-systems (SoS) such as an aircraft. As we vary the size of the SoS, the size of its systems typically also need to vary. For example, as we vary the aircraft under consideration from a small UAV to a large commercial aircraft, the characteristics of the EPS also changes. Since a diagnostic BN needs to vary accordingly, we now consider this in terms of impact on clique tree size.

We partition a BN's nodes into subsystems $\Upsilon = \{1, \dots, v\}$, and identify subsystem types $\Theta = \{1, \dots, \theta\}$, with $\theta \leq v$. In EPSs, typical subsystem types are: power generation, power storage, and power distribution. ADAPT has, for example, 3 power storage and 2 power distribution subsystems. Hence, $\Upsilon = \{1, 2, 3, 4, 5\}$ and $\Theta = \{1, 2\}$ for the variant of ADAPT we investigate in this paper.

We introduce a map f from nodes into subsystems: $f : \mathbf{V} \rightarrow \Upsilon$, and also a map g from subsystems into subsystem types: $g : \Upsilon \rightarrow \Theta$. Now, we can define different subsets of cliques from Γ , specifically $\Gamma_i = \{\gamma \in \Gamma \mid \text{for all } X \in \gamma, f(X) = i\}$, and obtain the following:

$$\tau(\Gamma_i) = \sum_{\gamma \in \Gamma_i} \prod_{X \in \gamma} |\Omega_X|. \quad (3)$$

In words, (3) provides the size of all cliques in a subsystem.

We define a set of interaction cliques Γ_0 as $\Gamma_0 = \Gamma - \cup_{i=1}^v \Gamma_i$. The set Γ_0 represents the interaction between different subsystems. We obtain the following alternative expression for total clique tree size:

$$\tau(\Gamma) = \sum_{i=0}^v \tau(\Gamma_i). \quad (4)$$

Now, instead of considering the subsystems individually as in (4), we make the assumption that each of them is identical (given its type). Formally, we let $i \in \Upsilon$ and assume $\tau(\Gamma_i) = \tau(\Gamma_{g(i)})$ as well as $c_0 = 1$ and obtain the following result:

$$\tau(\Gamma) = \sum_{i=0}^{\theta} c_i \times \tau(\Gamma_i), \quad (5)$$

where c_i represents the number of times a subsystem of type $i \in \Theta$ is found in a system. The significance of (5) is that it

enables us to analyze the impact (on clique tree size) of different systems, with different size and redundancy requirements, by taking a compositional approach. Specifically, if we know or can reliably estimate $\tau(\Gamma_i)$, we just need to count the number of times c_i a subsystem type i occurs, and then do this for all subsystem types in a given system. This aligns well with design methodologies that use redundancy and product-line approaches to support the development of EPSs for vehicles with different power requirements.

An important but non-trivial question to consider is the value of $\tau(\Gamma_0)$ in (5) as subsystems are composed in different ways to form a system. Based on (5), we can identify a few special cases and simplifications; further information is provided by our experiments. One simplification, which we call perfect compositionality, puts $c_0 = 0$ in (5) to ignore interactions and adds together the size of each subsystem. Clearly, this creates a lower bound that scales linearly with the number of subsystems c_i for a given Γ_i .

The State Space (or Discretization) Factor: In EPSs, continuous signals are often converted to discrete digital numbers by means of analog-to-digital (A/D) converters. A key parameter in A/D conversion is the number of bits in discretized signal, and how to map these discretized into BN node states. Fundamentally, there is a desire to maximize the fidelity of the BN to the underlying EPS, but at the same time the computation time cannot get too large, because then a diagnosis will not be computed in time. The cardinality of a node has a multiplicative effect in the cliques in which is an element, see (2), and hence one needs to carefully trade off the potential improvement in diagnostic accuracy (due to increased discretization) with the cost of increased computation time. Further, this factor may need to be taken into account multiple times according to c_i in (5).

The Interaction (or Ambiguity) Factor: Increased interaction or ambiguity in a BN has a detrimental effect on scalability. Consider bipartite BNs as an example (Mengshoel *et al.* 2006; Mengshoel 2007b). Example of low ambiguity is when each leaf node has (for example) $P = 1$ parent nodes. Example of high ambiguity is when each leaf node has (for example) $P = 5$ parent nodes. The higher the ambiguity, the faster cycles are induced in the moral graph, as a function of the ratio of leaf nodes to root nodes, thereby more quickly inducing cliques with many BN nodes in the clique tree. This factor is perhaps less of a concern in engineered systems including EPSs, since they are typically less ambiguous and often close to tree structured (see experimental results below). However, there may be some ambiguity in the interaction between subsystems, thus impacting the term $\tau(\Gamma_0)$ in (5).

Electrical Power System Case Study

We now consider a small EPS case study. The high-level specification for this EPS is shown in Table 1. Figure 1 shows the BN that is auto-generated from this high-level specification. Figure 2 shows the clique tree that this BN is compiled into, using clique tree clustering. We hypothesize that it is much easier for many, including people well-versed in probabilistic models, to provide information in the format illustrated in Table 1 compared to what is illustrated in

Part Name	Type of Part	Failure Probability	Upstream Part
Battery1	battery	0.0005	
Wire1	wire	0.0000	Battery1
Voltage1	sensorVoltage	0.0005	Wire1
Current1	sensorCurrent	0.0005	Wire1
Breaker1	breaker	0.0005	Wire1
Status1	sensorTouch	0.0005	Breaker1
Wire2	wire	0.0000	Breaker1
Relay1	relay	0.0005	Wire2
Feedback1	sensorTouch	0.0005	Relay1
Load1	load	0.0005	Relay1
Temp1	sensorCurrent	0.0005	Load1

Table 1: High-level specification of a small electrical power systems (EPS). The EPS consists of two subsystems, namely a battery subsystem (lines from Battery1 to Status1) and a load bank subsystem (lines from Wire2 to Temp1).

Figure 1 or Figure 2. On the other hand, the high-level specification language is restricted to represent a certain class of BNs and not BNs in general.

Each line in a high-level specification represents one part of an EPS, and also contains information about its type, failure probability, and location within the overall system. For example, the line *Breaker1 breaker 0.0005 Wire1* in Table 1 communicates that *Breaker1* is a circuit *breaker*; has failure probability *0.0005*; and is downstream of *Wire1*. Broadly speaking, this specification is for an EPS with a single battery, *Battery1*, powering a single load *Load1*, and containing a few sensors and components. Specifically, *Battery1* has a wire *Wire1* downstream of it. *Wire1* has three parts connected to it, namely a voltage sensor *Voltage1*, a current sensor *Current1*, and a circuit breaker *Breaker1*. *Breaker1* has a feedback sensor *Status1*, which reports whether the breaker is open or closed, attached to it. *Wire2*, which is the first part that we consider to be part of the load bank subsystem (as opposed to the battery subsystem), is downstream of *Breaker1* and has feedback sensor *Feedback1* as well as *Relay1* attached to it. *Relay1* controls power flow into *Load1*, which has a sensor *Temp1* attached to it.

The auto-generated BN in Figure 1 contains one or more BN nodes for each part (or line) in the specification, and in addition nodes that glue the parts together into an overall EPS model. A key point here is that nodes can be partitioned, as indicated in the figure, into nodes that belong to the battery subsystem or the load bank subsystem. Formally, we have $\Upsilon = \{1, 2\}$ and $\Theta = \{1, 2\}$, with the map f as indicated by the coloring in Figure 1 and the map g simply $g(i) = i$ for $i \in \{1, 2\}$. Roughly speaking, the BN reflects both the “push” of power from the battery to the load as well as the “pull” of current by the load. For example, *Voltage1 Battery1* is — subject to *Health Battery1* (whether *Battery1* is operational or not) and *Closed Wire1* (whether *Wire1* is open or closed) — propagated downstream to *Voltage1 Wire1*, and so forth.

Figure 2 shows a clique tree resulting from the compilation of the BN in Figure 1. Here, cliques in Γ_1 represent

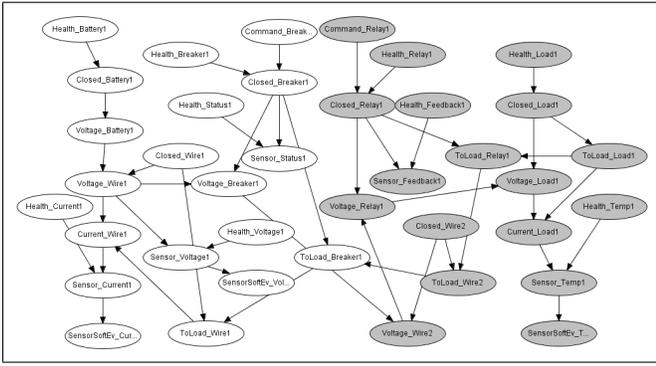


Figure 1: The BN auto-generated from a high-level specification (see Table 1) of a small electrical power systems. The BN represents two subsystems, namely a battery subsystem (white nodes) and a load bank subsystem (grey nodes).

the battery subsystem, those in Γ_2 the load bank subsystem, while cliques in Γ_0 represent the interaction between the two subsystems. Clique tree size is $\tau(\Gamma) = 264$, with $\tau(\Gamma_0) = 48$, $\tau(\Gamma_1) = 98$, and $\tau(\Gamma_2) = 118$.

Experiments

While our discussion earlier in this article has clarified a number of issues related to the development of large BNs by means of composition of subsystems, we made a few simplifying assumptions. In particular, our analysis did not fully account for: the pruning of BN nodes during auto-generation; possible interaction between subsystems; and the result of different clique tree decompositions (due to varying BN structures). To complement our analysis earlier in this article as well as previous accuracy results for ADAPT (Mengshoel *et al.* 2008), we have performed scalability experiments as reported in this section.

Experimental Design

The goal of the experiments is to study BNs representing different EPSs with varying number of subsystems of different types. Different EPS BN models were created using the high-level specification language. Clearly, we do not intend to implement real-world testbeds (similar to ADAPT) for all of these, and our goal here is rather to study the sizes of the generated BNs and clique trees (which determines computation time), which is one important design parameter when developing EPSs and their diagnostic capabilities. Following this approach, we developed 24 different EPS architectures using the high-level specification language, giving 24 auto-generated BNs, which were compiled into 24 clique trees which can be used for EPS diagnosis. In Table 2, the notation $EPS(x,y)$ is used to represent an EPS in which x represents the number of battery subsystems and y represents the number of load bank subsystems (see (Poll *et al.* 2007) for details on these subsystems).

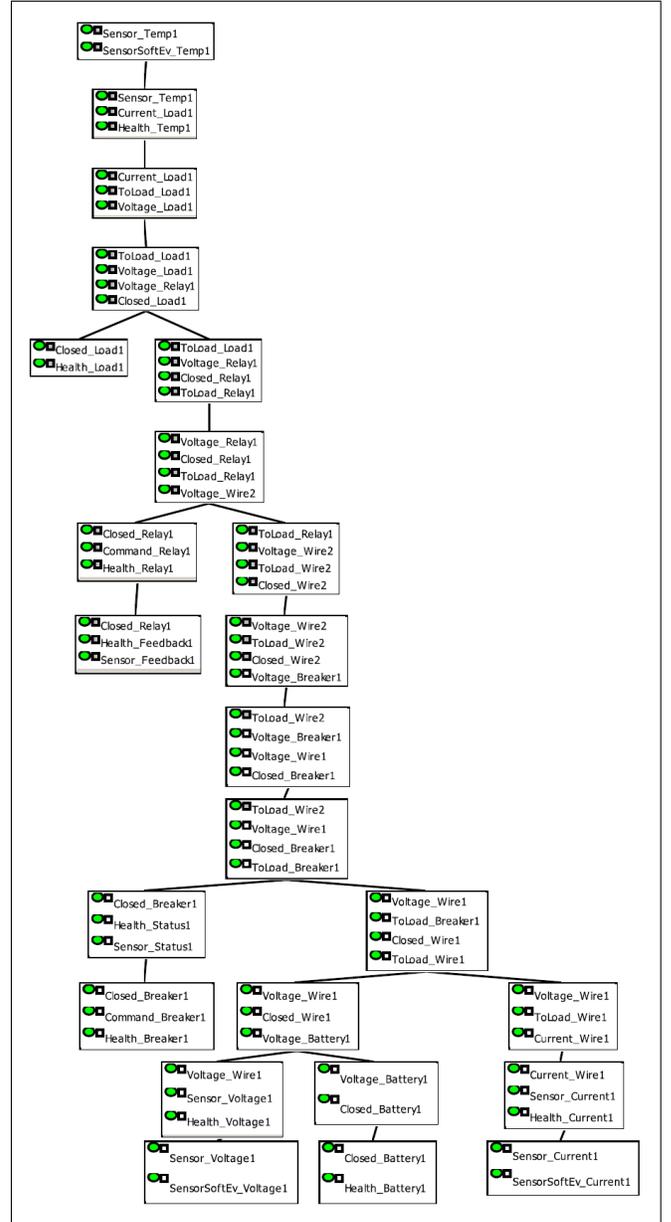


Figure 2: The clique tree compiled from a BN (see Figure 1) representing a small electrical power systems. These cliques can be partitioned into those that contain battery subsystem nodes (white nodes only) load bank subsystem nodes (grey nodes), and both nodes (white and grey nodes). The 11 cliques at the bottom are battery subsystem cliques. Among the cliques at the top, 9 are load bank subsystem cliques. In the middle, there are 3 cliques representing the interface between the two subsystems.

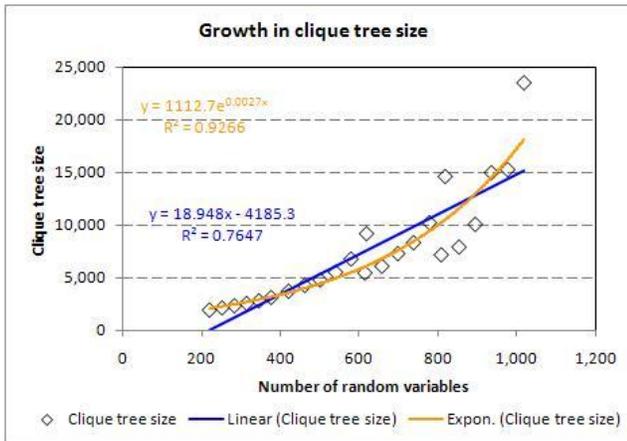


Figure 3: This figure shows how clique tree size varies as a function of the number of random variables (or BN nodes). Clique tree size determines computation time, while the number of random variables varies from EPS to EPS. Each data point, of which there are 24, represents an EPS.

Experimental Results

We now turn to the experimental results for the diagnostic BN models of 24 EPSs including ADAPT. Table 2 and Figure 3 summarize the experimental results; a few key observations follow:

- In Table 2, $\min(m/n) = 1.13$, while $\max(m/n) = 1.17$. This shows that our auto-generated BNs are fortunately quite sparse, given that $m/n = 1$ for trees.
- There is an approximately 5-time increase in BN size from EPS(1,1) to EPS(6,4), and a little over 12-time increase in clique tree size. We believe that this is quite promising, given the inherent hardness of BN computation. Further, if we consider EPS(5,4) instead of the outlier EPS(6,4), we have 4.4 times as many BN nodes compared to EPS(1,1) and only an 8-time increase in clique tree size.
- Generally speaking, n/m and the number of CPT parameters increase with increasing EPS size in Table 2, and both of these factors may suggest a harder inference problem. The product of the ratio n/m and the number of CPT parameters is maximal for EPS(6,4), perhaps giving some explanation for why this architecture is an outlier.
- The regression results in Figure 3 exhibit better fit for the exponential model ($R^2 = 0.9266$) than for the linear model ($R^2 = 0.7647$), pointing to the importance of the potentially nonlinear term $\tau(\Gamma_0)$ in (5). However, and in particular if the outlier EPS(6,4) is excluded, both models are quite reasonable.

Conclusion and Future Work

Due to their high level of predictability and fast execution times, Bayesian network compilation approaches are well-suited to automated diagnosis in the setting of on-board

resource-bounded reasoning and real-time systems of interest to NASA (Mengshoel *et al.* 2008). This paper improves the understanding of the scaling behavior of clique tree clustering in the context of composing large-scale BNs. A designer of model-based diagnostic systems, specifically using Bayesian networks, can use our novel approach to determine the impact of varying EPS architectures consisting of repeated subsystems on the computation time of diagnostic queries.

This work has been performed in the context of NASA’s ADAPT electrical power system testbed. While ADAPT is not a replica of an EPS that has been deployed on aircraft or spacecraft, it is representative of EPSs deployed on NASA missions. In this paper we have investigated how the BN-based approach to probabilistic diagnosis for ADAPT scales to other electrical power systems composed in a similar manner from power storage and power distribution subsystems.

This work enables the transition of diagnostic and health management technologies to NASA’s mission systems. In particular, it appears that Bayesian networks, techniques, and algorithms for diagnosis can be applied to distinguish between sensor failures and component failures, a problem of great interest to NASA. Sensor failures have been a costly problem for the Shuttle program, for example, with sensor failures causing launch delays in several cases. Future work will aim to help NASA in developing model-based diagnostic and sensor validation approaches that take into account the limited resources available on varying mission hardware.

References

- A. M. Abdelbar and S. M. Hedetnieme. Approximating MAPs for belief networks is NP-hard and other theorems. *Artificial Intelligence*, 102:21–38, 1998.
- R. M. Button and A. Chicatelli. Electrical power system health management. In *Proceedings of the 1st International Forum on Integrated System Health Engineering and Management in Aerospace*, Napa, CA, 2005.
- M. Chavira and A. Darwiche. Compiling Bayesian networks using variable elimination. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)*, pages 2443–2449, Hyderabad, India, 2007.
- C.-F. Chien, S.-L. Chen, and Y.-S. Lin. Using Bayesian network for fault location on distribution feeder. *IEEE Transactions on Power Delivery*, 17:785–793, 2002.
- A. Darwiche. A differential approach to inference in Bayesian networks. *Journal of the ACM*, 50(3):280–305, 2003.
- S. Lauritzen and D. J. Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems (with discussion). *Journal of the Royal Statistical Society series B*, 50(2):157–224, 1988.
- U. Lerner, R. Parr, D. Koller, and G. Biswas. Bayesian fault detection and diagnosis in dynamic systems. In *Proceedings of the Seventeenth national Conference on Artificial Intelligence (AAAI-00)*, pages 531–537, 2000.

Name of EPS	Number of Batteries	Number of Load Banks	n - BN Nodes	m - BN Edges	Ratio m/n	CPT Params.	Product	Clique Tree Size
EPS(6,4)	6	4	1,018	1,194	1.17	7.75	9.09	23,428
EPS(5,4)	5	4	977	1,143	1.17	7.54	8.82	15,202
EPS(4,4)	4	4	936	1,092	1.17	7.45	8.69	14,912
EPS(3,4)	3	4	895	1,041	1.16	7.43	8.64	10,014
EPS(2,4)	2	4	854	990	1.16	7.43	8.61	7,868
EPS(1,4)	1	4	809	935	1.16	7.48	8.64	7,130
EPS(6,3)	6	3	819	954	1.16	7.66	8.92	14,550
EPS(5,3)	5	3	779	906	1.16	7.47	8.69	10,164
EPS(4,3)	4	3	739	858	1.16	7.40	8.59	8,274
EPS(3,3)	3	3	699	810	1.16	7.38	8.55	7,248
EPS(2,3)	2	3	659	762	1.16	7.40	8.56	6,046
EPS(1,3)	1	3	616	711	1.15	7.46	8.61	5,404
EPS(6,2)	6	2	620	714	1.15	7.50	8.64	9,128
EPS(5,2)	5	2	581	669	1.15	7.35	8.46	6,726
EPS(4,2)	4	2	542	624	1.15	7.30	8.40	5,476
EPS(3,2)	3	2	503	579	1.15	7.30	8.40	4,738
EPS(2,2)	2	2	464	534	1.15	7.34	8.45	4,224
EPS(1,2)	1	2	423	487	1.15	7.43	8.55	3,678
EPS(6,1)	6	1	379	426	1.12	7.22	8.12	3,082
EPS(5,1)	5	1	348	392	1.13	7.13	8.03	2,768
EPS(4,1)	4	1	317	358	1.13	7.11	8.03	2,518
EPS(3,1)	3	1	286	324	1.13	7.15	8.10	2,300
EPS(2,1)	2	1	255	290	1.14	7.22	8.21	2,098
EPS(1,1)	1	1	223	255	1.14	7.35	8.40	1,896

Table 2: Electrical power systems (EPSs) are in aerospace often developed by composition or aggregation of subsystems. In this table we consider the effect of varying the number of two types of EPS subsystems, namely a battery subsystem and a load bank subsystem, on its BN representation and clique tree size. Total clique tree size is of interest because it determines the computation time for a wide range of probabilistic queries. The number of batteries is varied from 1 to 6. The number of load banks is varied from 1 to 4. The table also shows the number of BN nodes, the number of BN edges, the ratio of BN edges to BN nodes, the number of CPT parameters, and the product of this ratio and the number of CPT parameters. Here, the ADAPT BN corresponds to the highlighted EPS(3, 2) model.

O. J. Mengshoel, D. C. Wilkins, and D. Roth. Controlled generation of hard and easy Bayesian networks: Impact on maximal clique tree in tree clustering. *Artificial Intelligence*, 170(16-17):1137–1174, 2006.

O. J. Mengshoel, A. Darwiche, K. Cascio, M. Chavira, S. Poll, and S. Uckun. Diagnosing faults in electrical power systems of spacecraft and aircraft. In *Proceedings of the Twentieth Innovative Applications of Artificial Intelligence Conference (IAAI-08)*, pages 1699–1705, Chicago, IL, 2008.

O. J. Mengshoel. Designing resource-bounded reasoners using Bayesian networks: System health monitoring and diagnosis. In *Proceedings of the 18th International Workshop on Principles of Diagnosis (DX-07)*, pages 330–337, Nashville, TN, 2007.

O. J. Mengshoel. Macroscopic models of clique tree growth for Bayesian networks. In *Proceedings of the Twenty-Second National Conference on Artificial Intelligence (AAAI-07)*, pages 1256–1262, Vancouver, British Columbia, 2007.

D. Musliner, J. Hendler, A. K. Agrawala, E. Durfee, J. K. Strosnider, and C. J. Paul. The challenges of real-time AI.

IEEE Computer, 28:58–66, January 1995.

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

S. Poll, A. Patterson-Hine, J. Camisa, D. Garcia, D. Hall, C. Lee, O. J. Mengshoel, C. Neukom, D. Nishikawa, J. Ossensfort, A. Sweet, S. Yentus, I. Roychoudhury, M. Daigle, G. Biswas, and X. Koutsoukos. Advanced diagnostics and prognostics testbed. In *Proceedings of the 18th International Workshop on Principles of Diagnosis (DX-07)*, pages 178–185, Nashville, TN, 2007.

P. P. Shenoy. A valuation-based language for expert systems. *International Journal of Approximate Reasoning*, 5(3):383–411, 1989.

E. Shimony. Finding MAPs for belief networks is NP-hard. *Artificial Intelligence*, 68:399–410, 1994.

Z. Yongli, H. Limin, and L. Jinling. Bayesian network-based approach for power system fault diagnosis. *IEEE Transactions on Power Delivery*, 21:634–639, 2006.