# Diagnosing Faults in Electrical Power Systems of Spacecraft and Aircraft

**Ole J. Mengshoel**
USRA/RIACS
NASA Ames Research Center
Moffett Field, CA 94035
Phone: (650) 604-4199
Ole.J.Mengshoel@nasa.gov

**Adnan Darwiche**
Computer Science Department
University of California
Los Angeles, CA 90095
Phone: (310) 206-5201
darwiche@cs.ucla.edu

**Keith Cascio**
Computer Science Department
University of California
Los Angeles, CA 90095
Phone: (310) 825-7564
keith@cs.ucla.edu

**Mark Chavira***
Computer Science Department
University of California
Los Angeles, CA 90095
Phone: (310) 825-7564
chavira@cs.ucla.edu

**Ann Patterson-Hine**
Intelligent Systems Division
NASA Ames Research Center
Moffett Field, CA 94035
Phone: (650) 604-4178
Ann.Patterson-Hine@nasa.gov

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

**Serdar Uckun**
Intelligent Systems Division
NASA Ames Research Center
Moffett Field, CA 94035
Phone: (650) 604-4996
Serdar.Uckun@nasa.gov

**Abstract**

Electrical power systems play a critical role in spacecraft and aircraft, and they exhibit a rich variety of failure modes. This paper discusses electrical power system fault diagnosis by means of probabilistic techniques. Specifically, we discuss our development of a diagnostic capability for an electrical power system testbed, ADAPT, located at NASA Ames. We emphasize how we have tackled two challenges, regarding modelling and real-time performance, often encountered when developing diagnostic applications. We carefully discuss our Bayesian network modeling approach for electrical power systems. To achieve real-time performance, we build on recent theoretically well-founded developments that compile a Bayesian network into an arithmetic circuit. Arithmetic circuits have low footprint and are optimized for embedded, real-time systems such as spacecraft and aircraft. We discuss our probabilistic diagnostic models developed for ADAPT along with successful experimental results.

**Keywords:** Bayesian networks; arithmetic circuits; uncertainty; model-based diagnosis; knowledge engineering; electrical power systems; real-time systems; domain modelling.

**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 perform model-based diagnosis using probabilistic techniques. Specifically, we discuss the use of Bayesian networks and arithmetic circuits to perform diagnosis and health management in electrical power systems in aircraft and spacecraft. Two of the main tools that we have used, SamIam and ACE, are available to the general public — see `http://reasoning.cs.ucla.edu/samiam/` and `http://reasoning.cs.ucla.edu/ace/` respectively for details. We address two important issues that arise in engineering diagnostic applications in this area, namely the *modelling challenge* and the *real-time reasoning 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.

**Indication of application status (e.g., feasibility analysis, research prototype, operational prototype, deployed application, etc.):** We discuss the development of a diagnostic application for 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. Our probabilistic diagnostic application is an operational prototype that works on real-world data from ADAPT.

*Mark Chavira is currently at Google.

## Introduction

Electrical power systems (EPS) are critical for the proper operation of aircraft and spacecraft (Button & Chicatelli 2005; Poll *et al.* 2007). EPS loads in an aerospace vehicle may include crucial subsystems such as avionics, propulsion, life support, and thermal management systems. In addition, there is a move towards all-electric aircraft and spacecraft designs, so reliance on EPSs in aerospace is increasing. Apart from their crucial role in spacecraft and aircraft, electrical power systems also play central roles in other parts of society, thus proper management of their health is important.

There are several challenges associated with EPS fault diagnosis. In this paper we discuss two challenges, which we call the modelling challenge and the real-time challenge respectively, that we have encountered while developing a diagnostic reasoner for a real-world EPS. We believe these challenges show up in a wide range of applications and are of general interest.

The *modelling challenge* concerns how to model an EPS by means of Bayesian networks. Our use of Bayesian networks is motivated by the combination of deterministic and stochastic behavior seen in such systems. For example, there is uncertainty regarding component and sensor failure. One part of the challenge is to construct an EPS diagnostic model that captures both types of behavior. Another part of the challenge is to model the EPS in sufficient detail to ensure high diagnostic accuracy. At the same time, the diagnostic model developed for a particular EPS should be robust, easy to extend and update (reconfigurable), and general enough that essentially the same approach can be used when modeling similar EPSs.

The *real-time reasoning challenge* is associated with the embedding of AI components, including diagnostic reasoners, into hard real-time systems (Musliner *et al.* 1995). For NASA, decision support for manned missions and autonomous action for unmanned missions are both of great interest. The avionics of both manned and unmanned vehicles often utilize a hard real-time operating system (RTOS). An embedded diagnostic engine, which is part of a vehicle's avionics, should therefore be designed within the RTOS framework and within its resource bounds. For example, an RTOS task needs to declare a worst-case execution time (WCET). Unfortunately, it is also known that BN inference problems are inherently computationally hard (Cooper 1990; Shimony 1994; Park & Darwiche 2004). In addition, many inference algorithms are associated with high expectation and/or variance in their execution times, and their WCET is unknown. The real-time reasoning challenge is associated with integrating the computationally hard diagnosis problem into an RTOS setting, thereby achieving real-time diagnostic performance that is crucial in many aircraft and spacecraft applications.

In this paper we present our Bayesian network approach to EPS fault diagnosis. We discuss the development of a BN for the Advanced Diagnostics and Prognostics Testbed (ADAPT) developed and located at NASA Ames (Poll *et al.* 2007). The BN explicitly represents the health of sensors and components, and also contains random variables for other EPS parts. We emphasize the structure of the Bayesian network, and also briefly discuss the modeling process including support by semi-automatic BN generation based on a high-level system model.

We have experimentally evaluated our Bayesian network on a number of ADAPT fault scenarios. A key consideration in our experimentation is that the diagnostics process ultimately will be a set of RTOS real-time tasks (Mengshoel 2007a). In order enable such real-time embedding, the ADAPT BN was compiled off-line into an arithmetic circuit, which was then evaluated on-line (Darwiche 2003; Chavira & Darwiche 2007). A unique point compared to previous work (Chien, Chen, & Lin 2002; Yongli, Limin, & Jinling 2006) is how a complex diagnostic search problem is reduced to two simple components: an arithmetic circuit and a small-footprint arithmetic circuit evaluator. Compiling the ADAPT BN, which contains over 400 nodes representing over 100 EPS components, to an arithmetic circuit, and evaluating it using the ACE arithmetic circuit evaluator, turns out to give accurate diagnostic results as well as very fast and predictable inference times. The mean arithmetic circuit evaluation time is less than one millisecond for all our fault scenarios. This is a successful demonstration of our approach on a real-world problem of great importance to NASA (Button & Chicatelli 2005; Poll *et al.* 2007).

The rest of this paper is structured as follows. First we discuss challenges associated with the diagnosis of electrical power systems, and why this is an important problem. Second, we present our approach to diagnosis of electrical power systems by means of Bayesian networks and arithmetic circuits. Finally, we present empirical results for an electrical power system test bed developed at the NASA Ames research center.

## Diagnosis of Electrical Power Systems

In this section we discuss the crucial role of electrical power in aerospace and present an electrical power system testbed developed at the NASA Ames Research Center.

### The Role of Electrical Power Systems in Aerospace

EPS loads in an aerospace vehicle include crucial subsystems: avionics, propulsion, life support, and thermal management systems. Loss of electrical power to any of these subsystems could lead to serious consequences for personnel or the vehicle.

There are, from the point of view of vehicle health management, several technical challenges associated with electrical power systems in general. First, electrical power systems often have a large number of distinct modes due to mode-inducing components such as relays, circuit breakers, and loads. If an EPS has $m$ such components, and we conservatively assume 2 discrete states for each, there are potentially $2^m$ modes in the EPS. Second, EPSs are hybrid systems that combine discrete modes and continuous dynamics; switches between modes are both commanded (for relays and loads) and autonomous (for circuit breakers and health states). Third, there are timing issues including transient and delayed behavior. Fourth, there is both sensor and

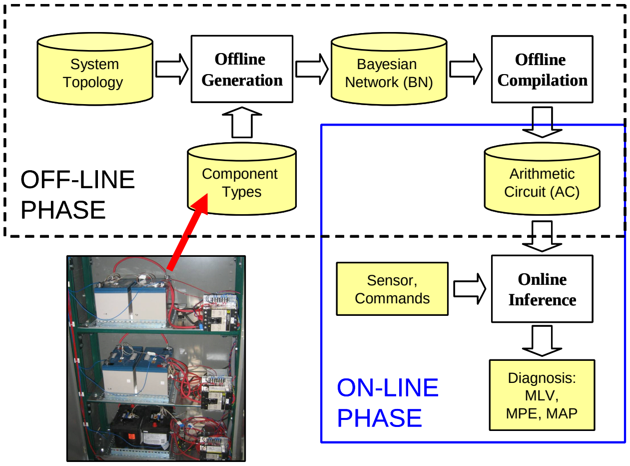


Figure 1: Our approach uses three distinct models that play different roles in the development process: a system model, a Bayesian network, and an arithmetic circuit. The system model has two components, namely the system topology and the component types.

system noise in EPSs. Fifth, the number and capabilities of sensors can be quite limited compared to what one would like to track, namely the health of all EPS components.

### ADAPT: An Electrical Power System Testbed

ADAPT provides: (i) a standard testbed for evaluating diagnostic algorithms and software; (ii) a capability for controlled insertion of faults, giving repeatable failure scenarios; and (iii) a mechanism for maturing and transitioning diagnostic technologies onto manned and unmanned vehicles (Poll *et al.* 2007). The EPS functions of ADAPT are as follows (see also `http://ti.arc.nasa.gov/adapt/`). For power generation, ADAPT currently uses utility power; there are plans to also investigate solar power generation. For power storage, ADAPT contains 3 sets of 24 VDC 100 Amp-hr sealed lead acid batteries. Power distribution is aided by electromechanical relays and two load banks with AC and DC outputs; there are also several circuit breakers. Our loads include pumps, fans, and light bulbs. There are sensors of several types, specifically for measuring voltage, current, relay position, temperature, light, and liquid flow. Control and monitoring of ADAPT takes place through programmable automation controllers. With the sensors included, ADAPT contains a few hundreds of components and is representative of EPSs used in aerospace.

## Meeting The Modelling Challenge

Bayesian networks (BNs) are used to represent multivariate probability distributions for the purpose of reasoning and learning under uncertainty (Pearl 1988). Random variables are in BNs represented by means of nodes in a directed acyclic graph. While a joint probability table's size is exponential in the number of discrete random variables, the BN provides a mechanism to compactly represent the joint probability table. BNs can contain both discrete and continuous random variables; our current EPS model contains discrete variables only.

The main points of our BN-based EPS modelling approach are as follows: (A) We use three different models during development and deployment. (B) We explicitly represent EPS health using BN nodes, thus supporting different diagnostic queries of interest. (C) Finally, we take a component-oriented and causal approach, where the BN structure reflects the components and causal structure of an EPS. In the rest of this section we discuss these three main points in turn.

(A) Figure 1 illustrates the three different EPS models we have used: the system model (consisting of a system topology and component types), the Bayesian network, and the arithmetic circuit. The justification for introducing a separate system model is the EPS structure present (as reflected in the distinction between the system topology and the component types) as well as the sheer size of BNs representing EPSs. Our current BN consists of well over 400 nodes, and it is easy to envision how a more detailed BN or a BN for a larger EPS could easily contain 1000 BN nodes or more. Unfortunately, developing such large BNs by hand, especially in the face of EPS change, is non-trivial. To meet this challenge, the Offline Generation process depicted in Figure 1 supports the automatic generation of an EPS Bayesian network from a high-level system model. This architecture, which distinguishes our work from previous work on EPS fault diagnosis using BNs (Chien, Chen, & Lin 2002; Yongli, Limin, & Jinling 2006), accommodates rapid changes in the EPS architecture as well as in individual EPS components. A key point here is that the system model is tailored to EPSs and is much more succinct than a Bayesian network, which again is much more succinct than an arithmetic circuit. Figure 1 shows how the diagnostic system developer is supported by a technique and tool pipeline for auto-generation of a Bayesian network from a high-level system model, and from a Bayesian network to an arithmetic circuit suitable for embedded reasoning. The Offline Compilation process generates an arithmetic circuit and is further discussed in our next section.

(B) The ADAPT BN currently contains over 400 nodes, and models most of ADAPT from the batteries downstream. Since it is impossible to presents this BN in its totality here, Figure 2 presents the inputs and outputs of the BN along with a small example. Figure 3 shows the BN's conditional probability tables (CPTs) along with its corresponding arithmetic circuit. We now discuss the different BN node types that we have used to model an EPS. Let $\boldsymbol{X}$ denote all BN nodes. The EPS health nodes are $\boldsymbol{H}_E = \boldsymbol{H}_C \cup \boldsymbol{H}_S$, where $\boldsymbol{H}_E \subseteq \boldsymbol{X}$ and $\boldsymbol{H}_C \cap \boldsymbol{H}_S = \varnothing$. Here, $\boldsymbol{H}_C$ are the component health nodes and represent the health of an EPS excluding its sensors. $\boldsymbol{H}_S$ are the sensor health nodes, and represent the health of the EPS sensors, both their failure and nominal (healthy) modes. By introducing $\boldsymbol{H}_C$ and $\boldsymbol{H}_S$, we represent the health of EPS components and sensors explicitly in the BN. The BN also contains other types of nodes, representing other parts of an EPS subsystem. Specifically, we have input or evidence nodes $\boldsymbol{E}$, with $\boldsymbol{E} = \boldsymbol{E}_C \cup \boldsymbol{E}_S$,

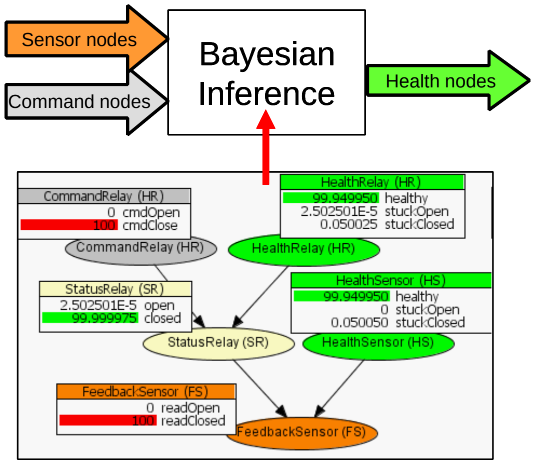


Figure 2: Our Bayesian diagnostic process has as *input* sensor readings for sensor nodes and observed commands for command nodes, and as *output* query nodes that provide the health status of sensors and EPS components.

where $\boldsymbol{E} \subseteq \boldsymbol{X}$ and $\boldsymbol{E}_C \cap \boldsymbol{E}_S = \emptyset$. Here, $\boldsymbol{E}_C$ are command nodes representing commands from a user to the EPS. $\boldsymbol{E}_S$ are the sensor nodes, which are used to input sensor readings — for example voltage, current, and temperature — from the EPS. We also have status nodes $\boldsymbol{S}$, with $\boldsymbol{S} \subseteq \boldsymbol{X}$, which are nodes that reflect the EPS structure but do not fit into any of the categories above. Finally, we have $\boldsymbol{X} = \boldsymbol{H}_E \cup \boldsymbol{E} \cup \boldsymbol{S}$, with $\boldsymbol{H}_E \cap \boldsymbol{E} = \emptyset$, $\boldsymbol{H}_E \cap \boldsymbol{S} = \emptyset$, and $\boldsymbol{E} \cap \boldsymbol{S} = \emptyset$.

(C) Figure 2 and Figure 3 provide a small example of our component-oriented and causal approach to EPS modelling. Here, $\boldsymbol{H}_C = \{\textit{HealthRelay}\}$, $\boldsymbol{H}_S = \{\textit{HealthSensor}\}$, $\boldsymbol{E}_C = \{\textit{CommandRelay}\}$, $\boldsymbol{E}_S = \{\textit{FeedbackSensor}\}$, and $\boldsymbol{S} = \{\textit{StatusRelay}\}$. This small BN with five nodes in fact represents an EPS component, namely a relay. Causally, the BN represents how the status of a relay (here *StatusRelay*) depends on the command given to it, *CommandRelay*, as well as its health, *HealthRelay*. In addition, the feedback message from the relay, *FeedbackSensor*, depends not only on the relay's status but also on the sensor's health, *HealthSensor*. The algorithm that creates a Bayesian network from the system model works as follows. Given small BNs representing different components, as presented above, an overall BN is composed according to the EPS system topology.

To solve the EPS health monitoring problem, we dynamically update the BN model using sensor readings and user commands. We then pose a maximum a posteriori hypothesis query MAP($\boldsymbol{Q}$, $\boldsymbol{e}$) over nodes $\boldsymbol{Q}$ for evidence $\boldsymbol{e}$. Here, MAP($\boldsymbol{Q}$, $\boldsymbol{e}$) computes the joint explanation over $\boldsymbol{Q} \subseteq \boldsymbol{X} - \boldsymbol{E}$ with maximal probability, given $\boldsymbol{e}$ (Park & Darwiche 2004) (We may also approximate MAP as further discussed below.) Depending on $\boldsymbol{Q}$ we obtain three slightly different diagnostic queries, all of great interest:

- Diagnosis of components MAP($\boldsymbol{H}_C$,$\boldsymbol{e}$): Query regarding the health of the EPS components $\boldsymbol{H}_C$
- Diagnosis of sensors MAP($\boldsymbol{H}_S$,$\boldsymbol{e}$): Query regarding the health of the EPS sensors $\boldsymbol{H}_S$
- EPS diagnosis MAP($\boldsymbol{H}_E$,$\boldsymbol{e}$): Query regarding the health status of the entire EPS $\boldsymbol{H}_E$ (both components $\boldsymbol{H}_C$ and sensors $\boldsymbol{H}_S$)

While algorithms for efficiently computing MAP have been developed (Park & Darwiche 2004), it can be useful to approximate MAP using MPE (most probable explanation) or MLV (most likely value, which can easily be computed from marginals) (Pearl 1988). We say $\text{MAP}_{\text{MPE}}(\boldsymbol{Q},\boldsymbol{e})$ and $\text{MAP}_{\text{MLV}}(\boldsymbol{Q},\boldsymbol{e})$ respectively for these two approximations.

Returning to Figure 2, we consider $\boldsymbol{H}_E = \{\textit{HealthRelay}, \textit{HealthSensor}\}$ and $\boldsymbol{e} = \{\textit{CommandRelay} = \textit{cmdClose}, \textit{FeedbackSensor} = \textit{readClosed}\}$. Using computation of marginals, as illustrated in Figure 2, we obtain $\text{MAP}_{\text{MLV}}(\boldsymbol{H}_E,\boldsymbol{e}) = \{\textit{HealthRelay} = \textit{healthy}, \textit{HealthSensor} = \textit{healthy}\}$.

## Meeting The Real-Time Challenge

Musliner and his coauthors identified three approaches to real-time AI (Musliner *et al.* 1995); we employ what they call "embedding AI into a real-time system". Specifically, we consider the real-time operating systems (RTOSs) used in current aircraft and spacecraft avionics. These RTOSs are typically based on priority-based preemptive scheduling, where higher-priority tasks preempt lower-priority tasks. Each periodic RTOS task has a priority, a period, a deadline, and a worst-case execution time (WCET). The implication of using an RTOS as a platform for diagnostics is resource-boundedness. A periodic diagnostic task, when designed as a periodic RTOS task, needs to adhere to these hard real-time requirements (Musliner *et al.* 1995; Mengshoel 2007a).

At the same time, the computational hardness of most BN inference problems is well-known (Cooper 1990; Shimony 1994; Park & Darwiche 2004). In addition, empirical studies have established the difficulty of relatively small application BNs (Shwe *et al.* 1991) as well as synthetic BNs (Mengshoel, Wilkins, & Roth 2006; Mengshoel 2007b).

A designer of BN-based diagnostic systems must therefore, in the general case, carefully align BN resource consumption with the resource bounds imposed by the computational platform. The compilation approach to probabilistic inference is attractive in such resource-bounded settings. We mention two compilation paradigms, namely compilation to clique trees (Lauritzen & Spiegelhalter 1988; Andersen *et al.* 1989) and compilation to arithmetic circuits (Darwiche 2003; Chavira & Darwiche 2007). The arithmetic circuit paradigm is based on the observation that a BN may be represented as a multi-variate polynomial (MVP) in which terms consist of probabilities from the BN's CPTs and indicators take into account evidence. Unfortunately, an MVP grows exponentially with the size of a BN, hence one compiles a BN not into an MVP but instead into an equivalent and more compact arithmetic circuit. An example is shown in Figure 3. In many cases, sparse arithmetic circuits exist for BNs with 100s or 1000s of nodes. The arithmetic circuit's size depends on a BN's graphical and local structure: if BN has local structure, the arithmetic circuit may

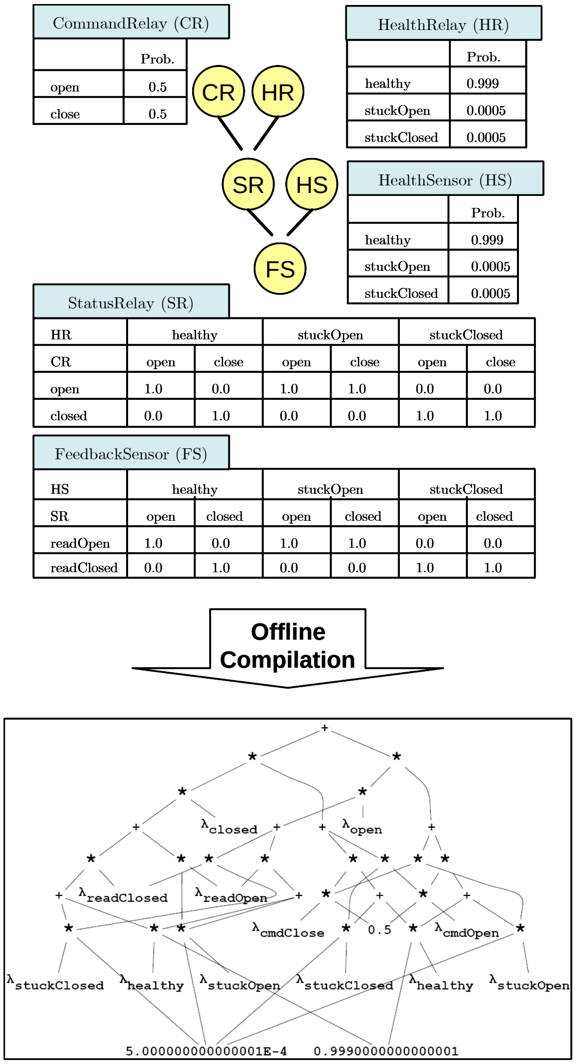


Figure 3: Compilation of a Bayesian network representing an electrical power system relay (top) into an equivalent arithmetic circuit (bottom).

be small despite large treewidth. A range of probabilistic queries — including MAP, MPE, and marginals/MLVs — can be computed using an arithmetic circuit.

The arithmetic circuit evaluator (ACE) was used to compile the ADAPT BN into an arithmetic circuit (see `http://reasoning.cs.ucla.edu/ace/` for details on ACE). ACE is a new technology that exploits a diagnostic model's behavioral aspects, and specifically its local structure. By being sensitive to local structure, ACE has been able to compile BNs with prohibitively large treewidths (Chavira & Darwiche 2007). For ADAPT, we have modelled the EPS such that its structure is maintained to a great degree; in addition the ADAPT BN has a large number of deterministic nodes. What do the arithmetic circuits look like for ADAPT, when compiled using ACE? Table 1 summarizes the results of ACE compilation, using the tabular

| Parameter | MPEs | Marginals |
|---|---|---|
| Network Read Time (sec.) | 0.328 | 0.313 |
| Initialization Time (sec.) | 0.25 | 0.25 |
| Compile Time (sec.) | 0.063 | 0.047 |
| Write Time (sec.) | 0.031 | 0.047 |
| Total Time (sec.) | 0.672 | 0.657 |
| Maximum Cluster Size | 256 | 256 |
| Number of Nodes | 3819 | 3854 |
| Number of Edges | 5660 | 5730 |
| Number of Variables | 685 | 685 |
| Collisions | 0.5845 | 0.5788 |

Table 1: Statistics for the compilation of the ADAPT BN into two arithmetic circuits for computation of MPEs and marginals respectively.

compilation option, into two arithmetic circuits for computation of MPEs and marginals respectively. The measurements reported in Table 1 (as well as in Figure 4) were made on a PC with an Intel 4 1.83 Ghz processor, 1 GB RAM, and Windows XP. For both arithmetic circuits generated here, we see from Table 1 that the total compilation time is less than one second and the arithmetic circuits, measured in number of nodes and edges, are relatively small. The small sizes of the arithmetic circuits enables real-time computation, as we investigate in the next section.

## Experimental Results

We now turn to experiments using the ADAPT Bayesian network and arithmetic circuits. For experimentation, EPS failure scenarios were generated using the ADAPT EPS at NASA Ames. These scenarios cover both component failures (experiments 304, 306, 309, and 310 in Table 2) and sensor failures (experiments 305, 308, and 311); many previous efforts have only considered one type of failure. In each of these experiments, ADAPT's initial state was as follows: Circuit breakers were commanded closed; they had evidence $\boldsymbol{e}$ clamped to *cmdClose*. Relays were commanded open; they had evidence clamped to *cmdOpen* in $\boldsymbol{e}$. In this initial state, all health nodes $\boldsymbol{H}_E$ are deemed healthy when computing MAP, $\text{MAP}_{\text{MPE}}$, and $\text{MAP}_{\text{MLV}}$. After ADAPT system reconfigurations and fault insertion (for example insertion of "Relay EY260 failed open" – see ID 304 in Table 2), the ADAPT BN or an arithmetic circuit compiled from it is used to compute a diagnosis.

The Bayesian network developed contains over 400 variables. We executed probabilistic queries over the health variables $\boldsymbol{H}_E$ in order to find out which components or sensors, if any, were in non-healthy states. Here, ACE was used to compute MPEs and marginals/MLVs. We report here on the queries $\text{MAP}_{\text{MPE}}(\boldsymbol{H}_E, \boldsymbol{e})$ and $\text{MAP}_{\text{MLV}}(\boldsymbol{H}_E, \boldsymbol{e})$ computed by ACE. To compute $\text{MAP}(\boldsymbol{H}_E, \boldsymbol{e})$, SamIam was used (see `http://reasoning.cs.ucla.edu/samiam/` for details).

The results of the ADAPT experiments are provided in Table 2 and in Figure 4. Since $\boldsymbol{H}_E$ contains over 120 nodes, we only show the variables deemed to be non-healthy

| ID | Fault Description | Diagnosis: MAP, $\text{MAP}_{\text{MPE}}$, and $\text{MAP}_{\text{MLV}}$ | Correct |
|---|---|---|---|
| 304 | Relay EY260 failed open | *Health_relay_ey260_cl = stuckOpen* | Yes |
| 305 | Relay feedback sensor ESH175 failed | *Health_relay_ey175_cl = stuckOpen* | Yes |
| 306 | Circuit breaker ISH262 tripped | *Health_breaker_ey262_op = stuckOpen* | Yes |
| 308 | Voltage sensor E261 failed | *Health_e261 = stuckVoltageLo* | Yes |
| 309 | Battery BATT1 voltage low | *Health_battery1 = stuckDisabled* | Yes |
| 310 | Inverter INV1 failed off | *Health_inv1 = stuckOpen* | Yes |
| 311 | Load sensor LT500 failed | *Health_LT500 = stuckLow* | Yes |

Table 2: Diagnostic results for different fault scenarios (with IDs 304, 305, . . . ) for the electrical power system testbed ADAPT.

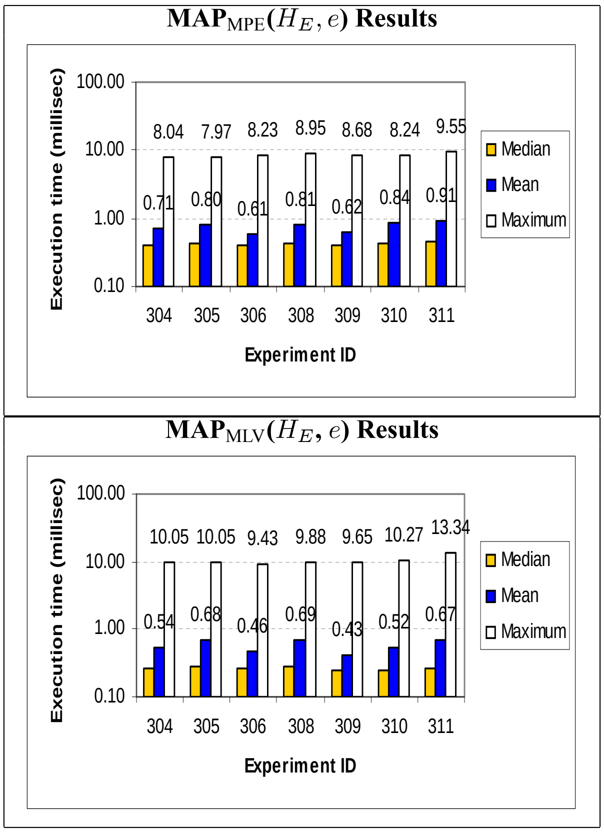


Figure 4: Execution time results for ACE for the ADAPT testbed. *Top:* Results for the most probable explanation (MPE); *Bottom:* Results for the most likely value (MLV).

in Table 2. Further, the diagnostic results of the queries $\text{MAP}_{\text{MPE}}(\boldsymbol{H}_E, \boldsymbol{e})$, $\text{MAP}_{\text{MLV}}(\boldsymbol{H}_E, \boldsymbol{e})$, and $\text{MAP}(\boldsymbol{H}_E, \boldsymbol{e})$ turned out to be the same, hence we consolidate them into one column in Table 2. ADAPT uses a 2 Hz sampling rate, and a probabilistic query was posed to ACE after each sample in an experimental run. The execution time statistics displayed in Figure 4 are based on the execution times for all probabilistic queries during an experimental run. Each execution time is for an entire inference step, i.e. translating measurements to evidence, committing evidence to the arithmetic circuit, and evaluating the arithmetic circuit.

Our main observations regarding these experiments are as follows. First, we see in Table 2 that the different diagnostic queries correctly diagnose all these component and sensor failure scenarios. Second, we would like to emphasize the fast and predictable inference times for the arithmetic circuits – see Figure 4. These are very important factors for real-time applications including electrical power system health management.

## Conclusion

Electrical power systems are crucially important in spacecraft and aircraft, thus motivating our interest in diagnosis of faults in such systems. In this paper we have presented a probabilistic approach to fault diagnosis in electrical power systems. Specifically, we have discussed how ADAPT, an electrical power system testbed at NASA Ames, can be represented as a Bayesian network which is the basis for answering diagnostic queries. We have highlighted two challenges, the modelling and real-time reasoning challenges, often associated with the development of model-based diagnostic engines for spacecraft and aircraft, and shown how they are overcome in our setting.

Meeting the modelling challenge, we have discussed how the EPS BN is structured in a component-based and causal manner. We have also considered how to meet the real-time (or resource-bounded) challenge associated with the real-time operating systems (RTOSs) used in spacecraft and aircraft. Our approach meets this challenge by compilation into an arithmetic circuit, where inference is fast and predictable, thereby enabling embedding into real-time tasks of RTOSs. Our BN-based fault diagnosis methodology has been successfully evaluated through experiments using real-world data from the ADAPT EPS.

**Acknowledgments** This material is based upon work supported by NASA under awards NCC2-1426 and NNA07BB97C.

## References

Andersen, S. K.; Olesen, K. G.; Jensen, F. V.; and Jensen, F. 1989. HUGIN—a shell for building Bayesian belief universes for expert systems. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, volume 2, 1080–1085.

Button, R. M., and Chicatelli, A. 2005. Electrical power system health management. In *Proceedings of the 1st International Forum on Integrated System Health Engineering and Management in Aerospace*.

Chavira, M., and Darwiche, A. 2007. Compiling Bayesian networks using variable elimination. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)*, 2443–2449.

Chien, C.-F.; Chen, S.-L.; and Lin, Y.-S. 2002. Using Bayesian network for fault location on distribution feeder. *IEEE Transactions on Power Delivery* 17:785–793.

Cooper, F. G. 1990. The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence* 42:393–405.

Darwiche, A. 2003. A differential approach to inference in Bayesian networks. *Journal of the ACM* 50(3):280–305.

Lauritzen, S., and Spiegelhalter, D. J. 1988. 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.

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

Mengshoel, O. J. 2007a. 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)*, 330–337.

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

Musliner, D.; Hendler, J.; Agrawala, A. K.; Durfee, E.; Strosnider, J. K.; and Paul, C. J. 1995. The challenges of real-time AI. *IEEE Computer* 28:58–66.

Park, J. D., and Darwiche, A. 2004. Complexity results and approximation strategies for MAP explanations. *Journal of Artificial Intelligence Research (JAIR)* 21:101–133.

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

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

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

Shwe, M.; Middleton, B.; Heckerman, D.; Henrion, M.; Horvitz, E.; Lehmann, H.; and Cooper, G. 1991. Probabilistic diagnosis using a reformulation of the INTERNIST-1/QMR knowledge base: I. The probabilistic model and inference algorithms. *Methods of Information in Medicine* 30(4):241–255.

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