

# Towards Software Health Management with Bayesian Networks

## Position Paper

Johann Schumann  
SGT, Inc., NASA Ames  
Johann.M.Schumann@nasa.gov

Ole J. Mengshoel CMU-SV  
Ole.J.Mengshoel@nasa.gov

Ashok N. Srivastava  
NASA ARC  
Ashok.N.Srivastava@nasa.gov

Adnan Darwiche UCLA  
darwiche@cs.ucla.edu

### ABSTRACT

More and more systems (e.g., aircraft, machinery, cars) rely heavily on software, which performs safety-critical operations. Assuring software safety though traditional V&V has become a tremendous, if not impossible task, given the growing size and complexity of the software.

We propose that iSWHM (Integrated SoftWare Health Management) can increase safety and reliability of high-assurance software systems. iSWHM uses advanced techniques from the area of system health management in order to *continuously* monitor the behavior of the software *during operation*, quickly detect anomalies and perform automatic and reliable root-cause analysis, while not replacing traditional V&V. Information provided by the iSWHM system can be used for automatic mitigation mechanisms (e.g., recovery, dynamic reconfiguration) or presented to a human operator. iSWHM's prognostic capabilities will further improve reliability and availability as it provides information about soon-to-occur failures or looming performance bottlenecks. In this paper, we will discuss challenges and future potential and describe how Bayesian networks (BN) could be used for iSWHM modeling.

### Categories and Subject Descriptors

D.2 [Software Engineering]: Distribution, Maintenance, Enhancement

## 1. INTRODUCTION

In modern aircraft and other complex machinery, important electrical, mechanical, and hydraulic components and systems are monitored by ISHM (Integrated System Health Management) systems. These can detect, diagnose, predict, and mitigate adverse events during the operation of the system. With the help of such diagnostics and prognostics techniques, appropriate mitigation strategies can be selected (re-

placement or repair, switch to redundant systems, etc.).

iSWHM (intelligent Software Health Management) extends this approach to software systems. An iSWHM system continuously monitors the behavior of the software and the interfacing hardware or sensor components. Using an abstract model of the software (e.g., a Bayesian statistical model), the iSWHM can detect unexpected behavior, reason about its root cause (failure identification) and trigger failure repair or mitigation actions. A similar mechanism can be used for prognostic purposes, trying to reliably predict possible software and/or system failures in the future.

In principle, an iSWHM system could operate similar to a traditional ISHM system, but would focus its attention on software instead of hardware. However, there are substantial differences between hardware and software systems:

- Software errors do not develop when the software is in operation but are introduced at some stage of the software development life cycle, as there is typically no wear-and-tear effects in software. Examples include requirements errors, design flaws, or coding errors, just to mention a few. If they are not detected and removed during testing, they remain (dormant) in the software system and can show up during operation.
- Failures in software often occur due to problematic interoperation with hardware. Hardware systems (and their sensors) might behave differently than expected, and thus could cause software failure. Such a different behavior could be by accident during development (e.g., a change in hardware is not reflected in the software design), as a result of hardware failure (e.g., a broken sensor cable, or a disabled sensor), or gradual degradation (e.g., signal noise increases beyond the specified level and causes the SW to behave erratically). Likewise, interoperation with human operators or other software systems can cause problems.
- Because many software systems interact with human operators, it is possible that the human can engage the system in an unexpected way due to the human misunderstanding the entire state of the system. In aeronautical applications on Flight Management Systems (FMS), this is called "Mode Confusion". If a pilot experiences mode confusion, he or she can interact with the system in unanticipated ways. Although verification and validation of the Flight Management System

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

FSE/DSP 2010 FSE/SDP Workshop

Copyright 2010 ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

Software can anticipate many such interactions, it is possible that the human operator could give a command which is unanticipated for a given configuration of the FMS.

- Many hardware systems have extremely complex physics-based models, which are used to predict the behavior of the system. These models can take the form of differential equations, difference equations, or other generative models based on physical laws. In general, software systems do not obey such physical laws. It is important to note, however, that the software systems themselves may be used to model or incorporate physical laws such as the case for motion dependent control systems. These models provide an interesting intermediate point between the two extremes of pure hardware systems with physical models and general software systems

All these differences (and commonalities) between ISHM of physical systems and software systems must be taken into account when developing novel techniques for an intelligent software health management system. Like all fault detection and monitoring systems, our iSWHM system is implemented as a piece of software itself. Safety analysis has to ask: “QUIS CUSTODIET IPSOS CUSTODES?” (Juvenal, “Who guards the guardians?”). This means that iSWHM systems must be at least as safe and dependable as the software they monitor (“host software”).

Some examples of major failures of safety-critical software systems can illustrate what kinds of software errors and problems an iSWHM system has to cope with. In Ariane V, several software modules from the smaller Ariane IV had been re-used [30]. However, the range of certain sensor values was larger (due to different physical dimensions and construction), which led to an uncaught overflow error, causing the rocket to behave erratically and required its destruction.

The recent incident with the NASA DART probe [4] was mainly caused by software problems. One major issue was that the GPS receiver was replaced just prior to launch with a different model, with different noise and bias without proper adaptation of the software. Automatic error correction in the navigation module caused wrong (biased) position and velocity values to be used as reference, causing the spacecraft to miss important trajectory points and to bump into the target satellite.

However, we do not claim that a SWHM system could have prevented all of these software-related mishaps.

## 2. ISWHM GOALS

While the overall goal of iSWHM is to extend and augment traditional V&V for a full lifecycle protection of software systems, thus ultimately enabling in-the-field assurance of composed software intensive systems, iSWHM needs to provide the following capabilities:

- iSWHM needs to *continuously monitor* the software under scrutiny. That software can be a compact piece of embedded software, or a huge, distributed software-rich system of systems, which might consist of heterogeneous components. It also must monitor the interactions of the software with the hardware, as many

software faults originate (or are triggered) by anomalies in software-hardware interactions.

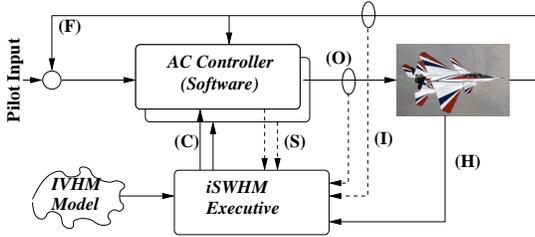
- iSWHM should provide model-based fault detection, fault identification (root cause analysis) and decision support (for mitigation systems or human operators).
- iSWHM should provide prognostic capabilities for enhanced system reliability, availability, and performance. iSWHM will not only react on problems that already occurred, but will be able to give future prognosis on performance (e.g., by predicting computational bottlenecks), availability, and reliability (prognosis of looming problems, e.g., memory leaks, overfull file systems, overloaded network connections).
- iSWHM will be capable of detecting environmental changes and *emerging behaviors* as those cannot be detected (by definition) during pre-deployment verification and validation (V&V).
- iSWHM will need to undergo rigorous V&V itself, as the iSWHM must be at least as reliable as the system it monitors.
- iSWHM models and reasoning capabilities must be able to minimize the number of false positives (spurious alarms) and false negatives (undetected failures).
- iSWHM must be integrated seamlessly with traditional V&V, as it is not intended to replace V&V but to augment it for in-the-field software assurance. In particular, pre-deployment V&V will provide verification credit. Also V&V information will be perused to improve iSWHM models.

Whereas ISHM is a mature field, research on the specific topic of software health management is still in its infancy. The 2009 SHM workshop [16] held during the Conference on Space Mission Challenges for Information Technology (SMC-IT 2009) gives an overview of some of the state-of-the-art approaches. Obviously, monitoring of software, while it is in operation, is an important topic of research. Extending the notion of runtime monitors for runtime verification enables the designer to explore possible fault states of the software in advance (see [31, 8]). The dynamic monitoring of highly reliable and redundant software poses its own challenges (e.g., [12]). Other SWHM research focuses on specific software architectures that are particularly amenable for SWHM (e.g., [9]). Some research describes SWHM in architectures that conform to the ARINC 653 standard [1] while others discuss the automated generation of fault trees [18]. A process-oriented approach to regularly check on the health of a (large) software system has been discussed in [26]. Here, the goal is that regular (non-automated) health checks improve the technical condition of the software and has positive economic effectiveness.

## 3. TECHNOLOGY

The principle architecture of an iSWHM system (Figure 1) consists of four components: the system (SW and hardware) to be monitored; a health model of the system; the iSWHM reasoning engine that performs failure detection; identification and prognostics; and components for failure announcement or failure mitigation. In the following, we will focus on

iSWHM modeling and reasoning aspects, as their scalability is particularly important for large-scale and heterogeneous software systems. Under the assumption that modeling and reasoning is done using Bayesian networks and arithmetic circuits [6, 25], we will also briefly discuss existing techniques and research needed to improve V&V of iSWHM.



**Figure 1: Principal architecture for iSWHM for an aircraft controller: pilot input and feedback (F) produce actuator output (O). Hardware health sensors (H), signal quality data (I), and software quality data (S) go into the iSWHM system, which produces a recovery/mitigation signal (C).**

### 3.1 Advanced iSWHM Modeling

Many software systems include a wide range of different and heterogeneous components along many dimensions, e.g., embedded vs. ground SW, autonomous vs. human-in-the-loop. As a consequence of the uncertain, heterogeneous, and interacting nature of these systems, as well as their environments, there is a need for supporting probabilistic modelling and analysis paradigms, techniques, and tools. We believe that it is important to pursue research regarding, on the one hand, probabilistic modeling and reasoning approaches (including dynamic Bayesian networks), and on the other hand, probabilistic model checking and more generally formal methods that involve probabilities. In particular, there is a need to integrate dynamic Bayesian networks research and research on probabilistic model checking based on Markov chains, as previous research efforts have largely been pursued independently and in different research communities.

We emphasize large-scale probabilistic graphical models, in particular hierarchical and compositional Bayesian networks, and algorithms for probabilistic inference and machine learning using graphical models to ensure scalability to large software systems. These Bayesian graphical models allow the designer to specify large models (for large software systems) in a hierarchical and structured way and use different levels of abstraction for the individual SW components.

Great progress has been made over the last decade, in learning and reasoning using probabilistic graphical models, including Bayesian networks (BN) and Markov networks. In addition to being well-suited to automated analysis, these graphical models are also amenable to visualization<sup>1</sup>. While most BN inference problems are computationally hard in the general case, efficient algorithms have been developed [6, 17] that open the path toward successful applications in a wide range of automated reasoning areas, for example in model-based diagnosis (e.g., [21]), sensor validation [2, 22], or intelligent data analysis [15, 27].

<sup>1</sup>e.g., <http://reasoning.cs.ucla.edu/samiam>

### 3.2 iSWHM with Arithmetic Circuits

Based upon sensor signals and health monitor signals, the iSWHM engine tries to disambiguate the information and locate the failures that have occurred, using a representation of a model of the host software. The reasoning engine often also makes decisions on how to overcome the failure or to recover the system. Because most inference problems needed for fault detection and identification are computationally expensive, we propose to use techniques that *compile* BNs into a data-structure, which allows highly optimized and efficient processing. For Bayesian networks, the fourth author has developed a translation of BNs into arithmetic circuits [3, 5, 6]. Powerful optimizations keep these data structures compact, so that reasoning can be performed over large models. Future research will investigate the compilation of hierarchical and heterogeneous Bayesian network models.

### 3.3 V&V of iSWHM Systems

The iSWHM system is a piece of software itself. Based on our discussion above, it is obvious that the iSWHM system is highly safety critical. False alarms or undetected faults can have severe consequences, ranging from unnecessary switching to redundant components to potential loss of life.

Therefore, all iSWHM system components have to undergo rigorous V&V as well as certification. In general, this involves two major parts: (a) V&V of the iSWHM model, i.e., assuring that the model reflects the software system and its failures/faults correctly and sufficiently, and (b) V&V of the algorithm and implementation of the proper iSWHM system, i.e., the reasoning engine and executive that will be running during the operation of the overall system.

As discussed earlier, our iSWHM models are represented as Bayesian networks, which can be compiled into arithmetic circuits of bounded size [6], which enables the iSWHM engine to perform efficient reasoning. However, this executive is a highly non-standard algorithm, which means that specific V&V techniques are needed. In particular, the following research questions need to be addressed:

- Correctness and completeness of model compilation: will reasoning with arithmetic circuits yield the same results as reasoning over the BN?
- Functional correctness of the iSWHM reasoning executive: does the implementation of the executive perform the right kinds of reasoning operations on the compiled model?
- Runtime and memory limitations: can the run-time for reasoning be limited (real-time guarantee)? Can the memory requirements for reasoning be limited upfront, such that no dynamic memory handling is necessary?

For V&V, advanced verification and validation tools, like the Java PathFinder model checker<sup>2</sup>, automatic generation of test-cases with symbolic PathFinder [24], compositional verification [11], and parametric testing [13, 28] will provide a basis. It is expected that techniques for the verification and validation of system health management software (e.g., [19, 29]) can be adapted for iSWHM.

<sup>2</sup><http://javapathfinder.sourceforge.net>

## 4. CONCLUSIONS

iSWHM is a key technology for detecting, diagnosing, predicting, and mitigating the adverse events during the operation of safety-critical software systems. The use of Bayesian networks for modeling and model-compilation into arithmetic circuits offers advantages (like well-defined semantics, wide range of techniques, tools, and algorithms, as well as exact compilation), which makes their use in monitoring the health of safety-critical and embedded software possible. Approaches for V&V and certification of iSWHM can be based upon advanced verification tools (like model checking), but substantial research will be necessary to address these issues.

**Acknowledgements** This work is in part supported by the NASA Aviation Safety Program IVHM project (NRA NNX08AY50A).

## 5. REFERENCES

- [1] M. Barry and G. Horvath. Goal-based Flight Software Health Management Services. In *SHM 2009*, 2009.
- [2] T. W. Bickmore. A probabilistic Approach to Sensor Data Validation. In *AIAA, SAE, ASME, and ASEE 28th Joint Propulsion Conf. and Exhibit*, 1992.
- [3] M. Chavira and A. Darwiche. Compiling Bayesian networks using variable elimination. In *Proc. IJCAI-07*, pages 2443–2449, 2007.
- [4] [http://www.nasa.gov/pdf/148072main\\_DART\\_mishap\\_overview.pdf](http://www.nasa.gov/pdf/148072main_DART_mishap_overview.pdf)
- [5] A. Darwiche. A differential Approach to Inference in Bayesian Networks. *JACM*, 50(3):280–305, 2003.
- [6] A. Darwiche. *Modeling and Reasoning with Bayesian Networks*. Cambridge University Press, 2009.
- [7] M. desJardins, P. Rathod, and L. Getoor. Bayesian Network Learning with Abstraction Hierarchies and context-specific Independence. In *Proc. ECML-2005*, volume 16, 2005.
- [8] W. Dong, M. Leucker, and C. Schallhart. Impartial anticipations in runtime verification. In *6th Int. Symp. on Automated Technology for Verification and Analysis (ATVA '08)*, vol 5311 LNCS. Springer, 2008.
- [9] A. Dubey, G. Karsai, R. Kereskenyi, and M. Mahadevan. A Real-Time Component Framework: Experience with CCM and ARINC-653. *IEEE International Symposium on Object-Oriented Real-Time Distributed Computing*, 2010.
- [10] N. Friedman and M. Goldszmidt. Learning Bayesian Networks with local Structure. In *Proc. UAI-96*, pages 252–262, 1996.
- [11] D. Giannakopoulou and C. S. Pasareanu. Interface Generation and compositional Verification in Java PathFinder. In *FASE*, vol 5503 of LNCS, pages 94–108. Springer, 2009.
- [12] A. Goodlow and L. Pike. Toward Monitoring fault-tolerant embedded Systems. In *SHM-2009*, 2009.
- [13] K. Gundy-Burlet, J. Schumann, T. Menzies, and T. Barrett. Parametric Analysis of ANTARES Re-entry Guidance Algorithms using advanced Test Generation and Data Analysis. In *iSAIRAS*, 2008.
- [14] Y. Guo, D. Wilkinson, and D. Schuurmans. Maximum Margin Bayesian Networks. In *Proc. UAI*, page 233, 2005. AUAI Press.
- [15] P. Jones, C. Hayes, D. Wilkins, R. Bargar, J. Sniezek, P. Asaro, O. J. Mengshoel, D. Kessler, M. Lucenti, I. Choi, N. Tu, and J. Schlabach. CoRAVEN: Modeling and Design of a Multimedia intelligent Infrastructure for collaborative Intelligence Analysis. In *Proceedings of the International Conference on Systems, Man, and Cybernetics*, pages 914–919, 1998.
- [16] G. Karsai, editor. *1st International Workshop on Software Health Management (SHM 2009)*. ISIS, Vanderbilt University, 2009.
- [17] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles And Techniques*. MIT Press, 2009.
- [18] T. Kurtoglu, R. Lutz, and A. Patterson-Hine. Using auto-generated Diagnostic Trees for optimized Fault Handling. In *SHM 2009*, 2009.
- [19] A. E. Lindsey and C. Pecheur. Simulation-based Verification of autonomous Controllers via Livingstone pathfinder. In *TACAS 2004*, vol 2988 of LNCS, pages 357–371. Springer, 2004.
- [20] O. J. Mengshoel. Understanding the Role of Noise in Stochastic Local Search: Analysis and Experiments. *Artificial Intelligence* 172:8–9, pages 955–990, 2008.
- [21] 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 *Proc. IAAI-08*, pages 1699–1705, 2008.
- [22] O. J. Mengshoel, A. Darwiche, and S. Uckun. Sensor Validation using Bayesian Networks. In *iSAIRAS*, 2008.
- [23] O. J. Mengshoel, D. Roth, and D. C. Wilkins. Portfolios in Stochastic Local Search: Efficiently Computing Most Probable Explanations in Bayesian Networks. *JAR* (accepted), 2010.
- [24] C. S. Pasareanu and W. Visser. Symbolic Execution and Model Checking for Testing. In *Haifa Verification Conf.*, vol 4899 of LNCS, pages 17–18. Springer, 2007.
- [25] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, 1988.
- [26] M. Pizka and T. Panas. Establishing economic Effectiveness through Software Health Management. In *SHM 2009*, 2009.
- [27] C. C. Ruokangas and O. J. Mengshoel. Information filtering using Bayesian networks: Effective user interfaces for aviation weather data. In *Int. Conf on intell. User Interfaces*, pages 280–283, 2003.
- [28] J. Schumann, A. Bajwa, and P. Berg. Parametric Testing of Launch Vehicle FDDR Models. In *AIAA Space*, 2010.
- [29] J. Schumann, A. Srivastava, and O. Mengshoel. Who guards the Guardians? — toward V&V of Health Management Software (short paper). In *Runtime Verification 2010 (submitted)*, 2010.
- [30] Wired.com: "History's Worst Software Bugs", 2009.
- [31] C. Zhao, W. Dong, J. Wang, P. Sui, and Z. Qi. Software active online Monitoring under anticipatory Semantics. In *SHM 2009*, 2009.