

Probabilistic Health and Mission Readiness Assessment at System-Level

Michael Borth¹, Leonardo Barbini²

^{1,2}*ESI (TNO), Eindhoven, 5656 AE, The Netherlands*

michael.borth@tno.nl

leonardo.barbini@tno.nl

ABSTRACT

Predictive maintenance strategies which estimate remaining useful life of system components ensure reliable availability of assets and low total costs of ownership by preventing breakdowns and down-times with timely and well-scheduled maintenance. The focus on the components' life times falls short, however, to infer the system-level capability to achieve upcoming tasks, especially if these tasks vary either in the strain they cause for the system or in the environmental conditions in which the system needs to perform. Such an assessment of the health and mission readiness of a system is crucial for mobile assets like seafaring vessels undertaking long-term operations without the option to easily come in for repairs or for industrial assets that need to complete long production runs in one go under varying circumstances. We propose a multi-step methodology to achieve such assessments using both Bayesian reasoning for diagnosis and prognosis and physics-based simulation models. First, we construct an appropriate Bayesian network in an object-oriented way by fitting a pre-compiled library of network fragments to the system's schematics using generative techniques. We then parameterize the obtained network using a combination of expert knowledge and machine learning to fine tune system-level interactions between components and their link to the system's performance. The learning step uses past operational data that we augment or complement with synthetic data, created by a physics-based simulation model, where needed. Finally, we use the trained Bayesian network to assess the mission readiness of the system given the probabilistics of its diagnosed state, expected impact of possible maintenance interventions, and the estimated profile of the future use. We illustrate and verify our methodology on a cooling system with an active feedback control loop, but our approach for mission readiness assessment is domain-independent, universally applicable, and typically feasible where operational data and engineering knowledge can be brought together to solve its challenge.

Michael Borth et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

1. SYSTEM-LEVEL MISSION READINESS

In recent years, we saw improvements in the management of assets and their operational efficiency including reductions of the total cost of ownership based on predictive maintenance (PdM). Designed to estimate when maintenance should be performed based on the condition of the assets – and not on pre-determined schedules or in response to break-downs – PdM realizes cost-benefits by preventing asset down-times as well as unnecessary or ill-timed maintenance service (Mobley, 2002). Its successes are based on the novel fusion of Internet-of-Things (IoT) technologies, i.e., the extension of networked connectivity into physical devices and sensors, with advances in data science, leading to consistent listings of PdM as major market trend, e.g. by Gartner (Berthelsen, 2018), even though, as Mulders et al. (2017) found out, many companies are still in the early stages of implementing it.

Typically, PdM will include, among other steps, a failure prediction based on real-time asset performance monitoring and reliability modeling founded in, e.g., failure mode and effects analysis (FMEA) and mean-time between failure (MBT) models and statistics. This failure prediction often breaks down to an estimation of the assets' component's remaining useful life (RUL), for which various techniques exist, as Si et al. (2011) elaborate in their overview.

In our work, we find this estimate of the components' RUL to be of a too reductionist understanding of system health, as we face mobile assets like seafaring vessels undertaking long-term operations without the option to come in for repairs or industrial assets that must complete long production runs in one go under varying circumstances. For those systems, we need to assess their capability to achieve upcoming tasks: their mission-readiness. Especially for adaptive systems, which compensate for unfavorable environmental conditions, for sub-par performance of components, or for variations between tasks or missions, we see that this notion of system health is relevant and more expressive than the prediction of breakdowns and must therefore complement it. We illustrate this with a real-world case in this chapter before we describe our contribution on probabilistic health and mission readiness assessment at system-level in the remainder of the article.

1.1. Case Study: Adaptive Cooling in Vessels

Consider, for the purpose of illustration, the cooling system of seafaring vessels. Such a system's task is to provide a flow of cooling liquid of constant temperature to various systems, e.g., the ship's engines. The cooling liquid transports excess heat away from those systems, thus warmer liquid comes back to the heat exchanger of the cooling system, where sea water is used to bring it back to the setpoint temperature. As both the temperature of the sea water and the cooling liquid vary due to location, weather, load of the cooled systems and thus the amount of excess heat, the cooling system uses an adaptive control loop that sets various valves, pumps and a heater, as depicted in the figure below.

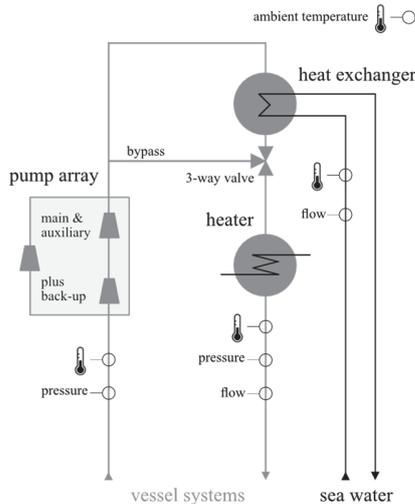


Figure 1. Schematic of a vessel's cooling system.

As a failure of the cooling system would severely hamper the vessel's ability to sail, it is usual for large vessels to safeguard its operation, e.g., with redundancy in the pump regime, but also with functionality to maintain performance, e.g., by the generation of pressure spikes that flush residue and other obstructions from the pipes. Switching pump configurations or the described intervention to flow characteristics change the behavior of the cooling system – which is why the control system is adaptive – but also the interpretation of the sensor data that is available for the purposes of control and health assessment.

Oversimplified, the latter results in the situation that the same time-series of measurements taken in two different situations might have very different interpretations. A failing success in cooling could, e.g., indicate a reduced efficiency of the heat exchanger, but also point to the existence of residue hampering the flows or even sensor reading pollution, e.g., if temperature readings are off as their environment heats up. Consequently, both the control and the system's health assessment become a context-sensitive multi-variant problem that includes changing relationships between the system variables. It is, in the sense of Snowden's (2003) Cynefin framework, complex and not merely complicated.

1.2. Mission Readiness

The changing impacts of the cooling system's components lead to converse effects for our assessment of the system's health: a reduced usefulness of a component can be meaningless, if compensated for; compensating behavior can hide occurring loss of performance, but typically only up to a tipping point after which the overall system fails fast; fail-safe mechanisms like back-up components prevent loss of functionality, but their use can cause a strain to other parts of the system via ancillary effects.

We conclude that health indications that predict the remaining useful life of components will not predict the performance and capabilities of the overall system up to the point of a potential failure, as their static way of data interpretation is not suited for highly adaptive systems if the relevant dynamics include the way system components interact. Nonetheless, RUL indicators do provide a valuable time-corridor that cannot be ignored for maintenance planning, especially if drawn from data that measures wear and tear effects.

In our work, we use instead the above-mentioned notion of mission readiness. This term describes the lasting capability of a system to fulfill its function within a series of planned activities under a set of expected circumstances. With mission readiness, we express if the overall system will perform good enough and long enough to achieve set goals, as only this allows us to include the option of mission changes in maintenance considerations: We can, e.g., delay maintenance by sending another vessel of our fleet on the journey that the ship we originally considered could no longer do, thus keeping both vessels active.

Interestingly, we find that a system health notion like mission readiness, that investigates the ability to perform (i) a series of well-defined tasks under (ii) an expected range of conditions, addresses a challenge we see arising from the effects listed at the beginning of this section. Figure 2 illustrates how such effects can induce a switching behavior in health calculations. If we, e.g., use similarity-based RUL estimates in which the time-series of a condition indicator is mapped to previous observation to find the one most similar and thus most predictive, like e.g., Baru (2018) describes, we see that the systems we are interested in switch the trajectory they are on. This happens gradually, e.g., while adjusting to a new environment, or abruptly, e.g., in response to task changes. As mission readiness makes such circumstances explicit, we avoid the miscalculations or uncertainty that follow from this.

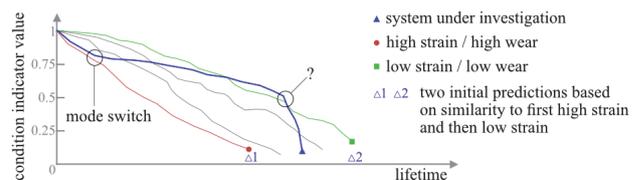


Figure 2. Switching behavior in RUL estimates.

2. PROBABILISTIC MISSION READINESS DIAGNOSIS

The assessment if a system can perform a given mission under the current circumstances is a diagnostic task. As we must deduce this likelihood from observations that are incomplete or imprecise (missing data and the limits of sensor accuracy), infer hidden states (health of components), and based on assumptions regarding the impact of the environment, we transpose this diagnosis as reasoning under uncertainty. Bayesian networks are a prime methodology for this purpose. Introduced by Pearl (1986), they have been shown to perform in a range of applications (Heckerman, 1995), especially as they are one of the few techniques that combine knowledge- and data-driven modeling, as showed, e.g., in (Jensen, 2007). This led to their even greater success in recent years, as the progress in data science improved the latter, while probabilistic programming techniques eased the former (Hardesty, 2015). We build our probabilistic mission readiness assessment with Bayesian networks in an approach that draws from all these concepts.

First, we generate a knowledge-based structure, i.e., an acyclic directed graph in which the nodes represent system variables, e.g., health states of components or sensor readings, and the edges encode probabilistic dependencies between variables and all their parents. We do this in so-called network fragments (Laskey, 1997) that, together with Pfeffer's work (1997), provide an object-oriented representation for probabilistic knowledge. Each network fragment renders a specific part of the system, e.g., a component, in a small set of nodes that follows distinct design patterns, e.g., for active and passive components or for sensors. This allows us to build a library of components and associated fragments that we use to compose the complete network without repetitive efforts using the generative techniques that probabilistic programming offers. In this step, we exploit the composability of individual network parts that results from adhering to engineering rules, as Borth illustrated in (2002) and Pearl (2009) founded in the principles of causality, as we trace the system's schematics like those shown in Figure 1 and build the links between fragments accordingly. We elaborate this in the upcoming section, but also refer to (Ricks, 2009) for a more detailed description of similar work.

A data-driven step is the next part of our approach to generate the Bayes net that provides the mission readiness assessment. Here, we use operational data of the system or, if necessary, similar systems or simulations, to determine the parameters of the network, i.e., the conditional probability distributions that define the links between the variables. This step, for which many techniques exist, see e.g., (Barber, 2012), individualizes the fragments according to their actual state and behavior and removes the need to capture these probabilities via knowledge engineering – which would be hard due to the sheer number of parameters that a large Bayes net holds.

Finally, we use probabilistic reasoning to diagnose the mission readiness, which we detail in the last section of this chapter.

2.1. Generation of the Probabilistic Model

As laid out above, we follow an object-oriented design flow to generate the diagnostic model for mission readiness assessment. In this flow, we analyze system components to match them to a set of network fragments that are in turn based on design patterns. Those patterns are illustrated in Figure 3 for active and passive components, e.g., the 3-way valve and the heat exchanger of the example cooling system. We see that all components are modeled around a core pattern that connects a component's age with its effect that in turn may be impacted by disturbances like temporary obstructions of moving parts. Next to the core, we use two extensions to model the positive effects of maintenance, if applicable, and, for active components, the relationship between the control setting and the overall effect of the component given its current setting and potential effect.

As the schematic of the cooling system shows, the valve and the heat exchanger connect in the flow of the cooling liquid. We use this connection to show the fragments and respective part of the network composition in Figure 4. The composition process eventually provides us with the structure of a full Bayesian network for the cooling system, as displayed in Figure 5 (we omitted pumps and heater sections for size reasons).

Due to the network structure mirroring the actual system, we find that domain experts have no difficulties working with the resulting Bayes net. They can, e.g., investigate and even debug individual parts that represent components in isolation but also examine how the parts act together in concert with the control system. The pattern-based object-oriented composition proved itself in greatly reducing our modeling efforts, as we saw a major benefit from front-loading: well-worked out patterns led to an easy generation of fragments for major component types (like pumps), which can then be easily instantiated to build a component library for consistent and effort-free re-use.

In practice, such a library holds, e.g., one Bayes net fragment for each pump model. The fragments of various pumps differ in the state space of their variables to reflect different product characteristics (flow, pressure) and their relationships between age, wear (reduced component effect) and tear (likelihood of disturbances and failures) that are expressed in the conditional probabilities associated with the fragments' nodes. Differences here relate to the pumps' quality, mean time between failures, and life expectancy, which often varies greatly between manufactures. Structural differences between fragments, on the other hand, originate from differences in sensory equipment as well as pump types if those come with different implemented functionalities, e.g., regarding the aforementioned generation of pressure spikes to flush obstructions out of the system.

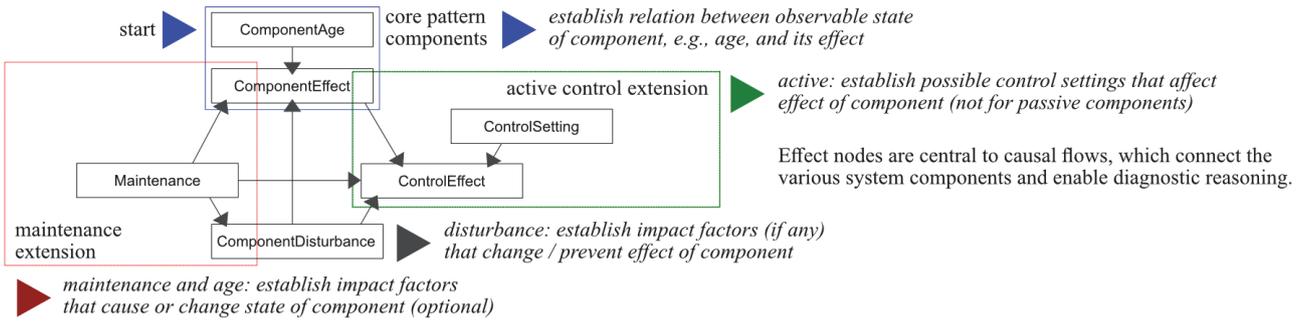


Figure 3. Pattern-based design of network fragments.

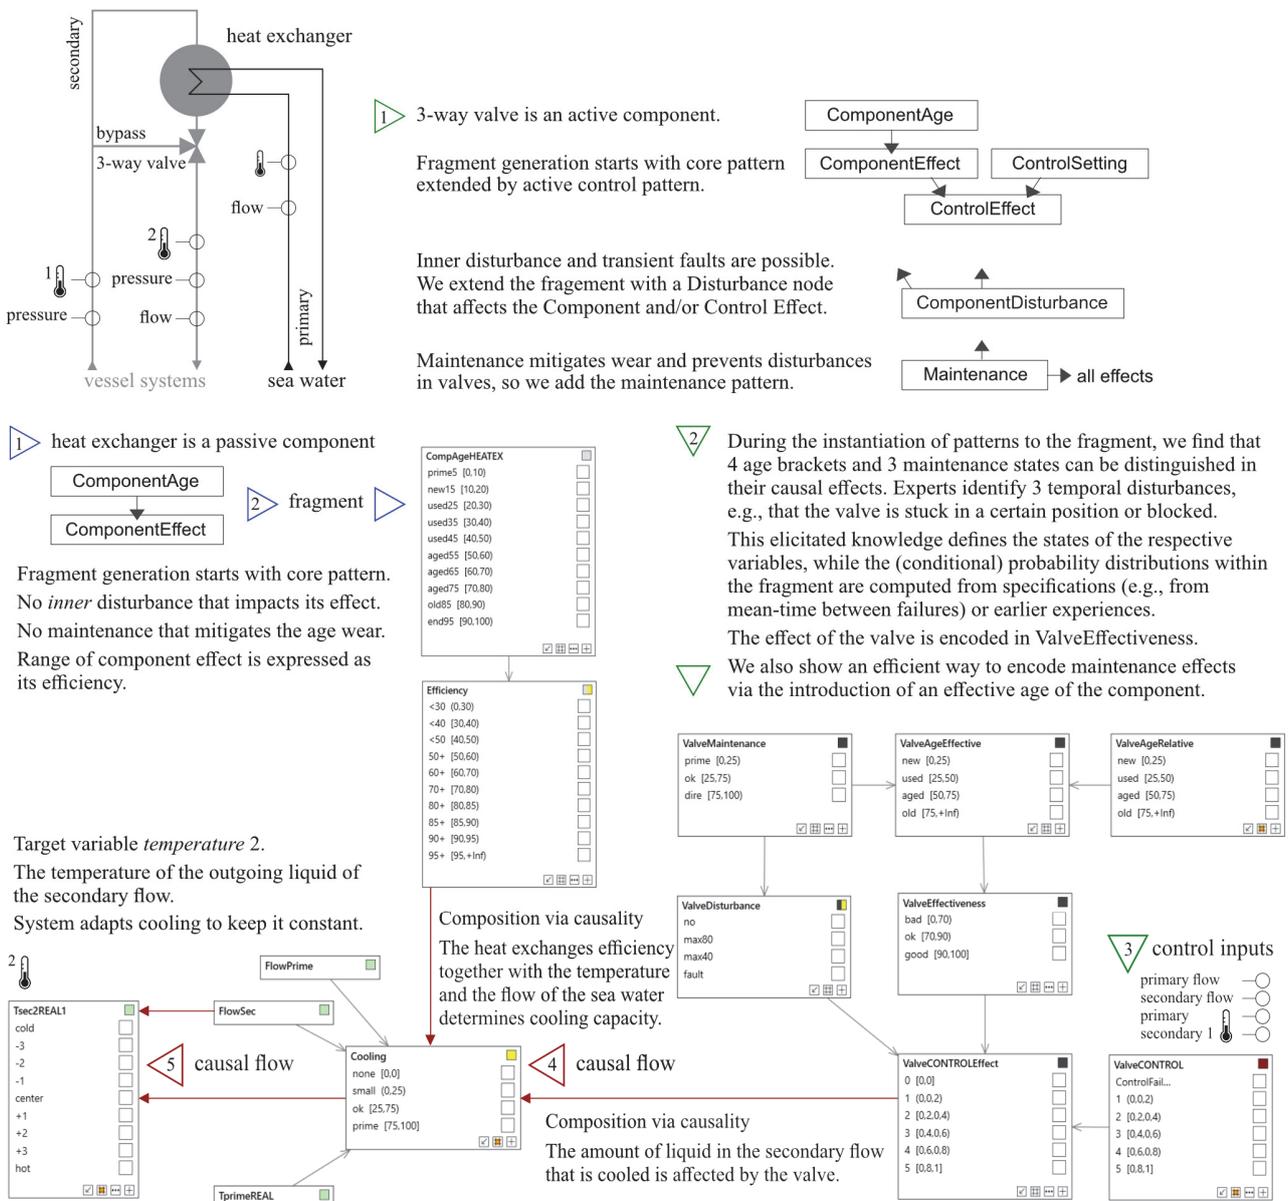


Figure 4. Network composition.

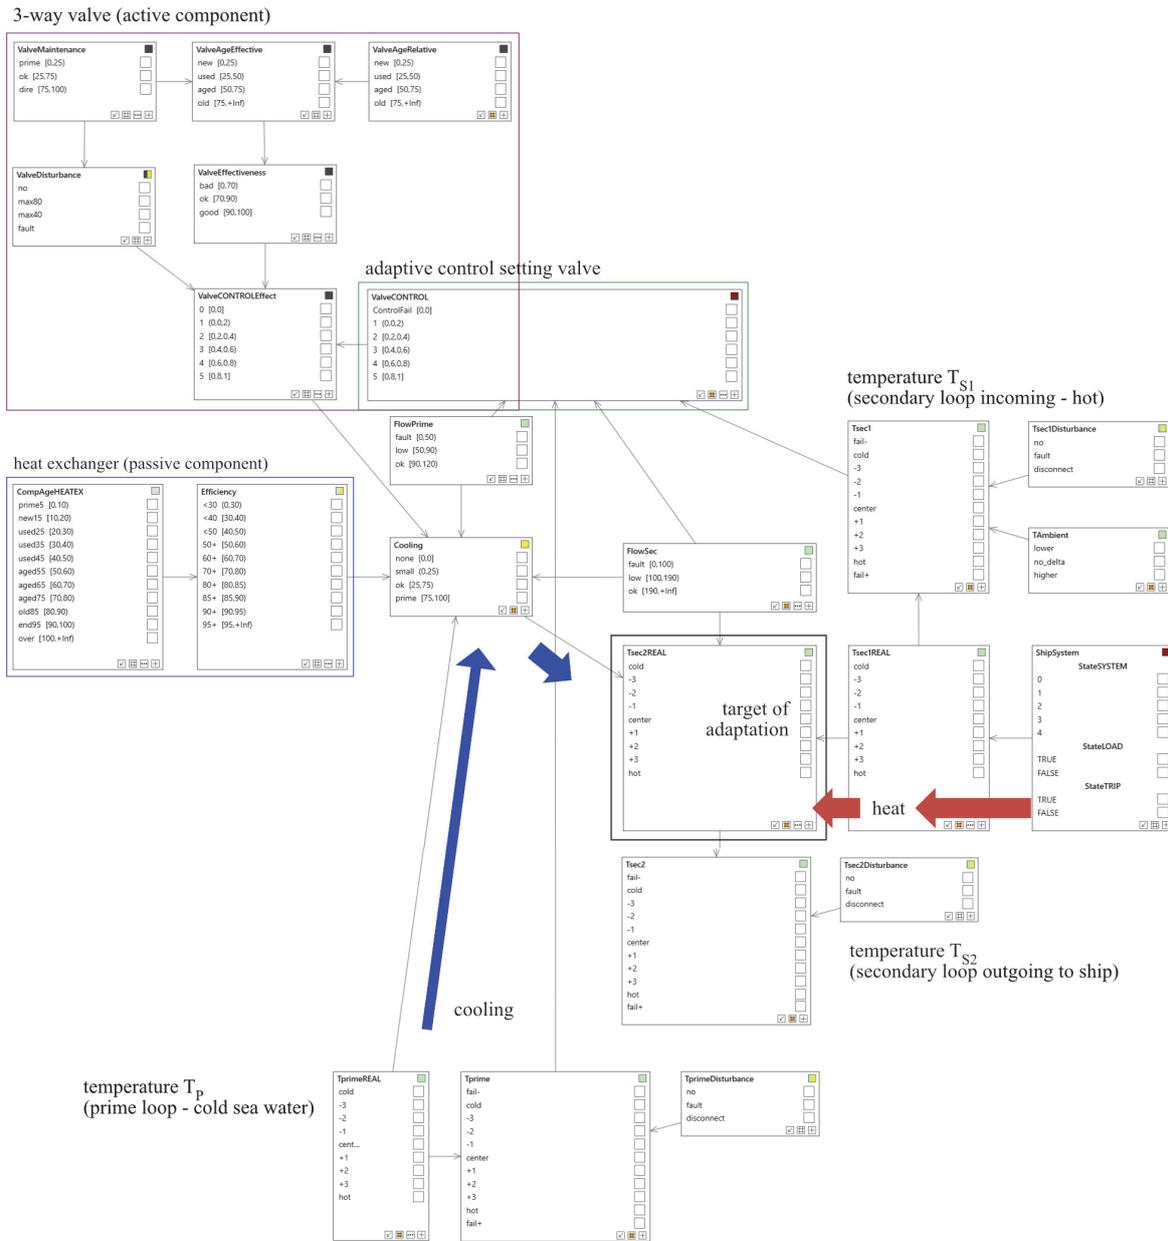


Figure 5. Diagnostic Bayes net for the cooling system (heater and pumps omitted). The figure shows parts of a screenshot from the BayesServer software which provides the Bayes nets modeling environment together with the necessary reasoning and machine learning capabilities.

The Bayes net has two main points in which the fragments connect: (i) the (real) temperature T_{S2} , i.e., the temperature of the secondary loop going towards the vessel systems, which is determined by the cooling effect, the secondary flow, and the temperature which the liquid had coming from the vessel’s systems; and (ii) the valve control that holds the adaptive regulation of the 3-way valve that sets the mixture of warm and cold water to stabilize the outgoing liquid’s temperature.

The Bayes net also accounts for sensor reading inaccuracies and sensor failures, as it models each temperature sensor with three nodes: one for the actual temperature (real), one for the sensor reading, and one for possible sensor disturbances.

2.2. Data-based Instantiation of the Probabilistic Model

We already explained in the introduction of this chapter that we individualize all fragments according to their actual state and behavior, i.e., we set the a priori conditional probability distributions such that they reflect the system under investigation as closely as possible. This ensures, e.g., that the diagnosis for two systems differs even though both show a similar performance, but one shows a history of issues with a badly maintained pump while the other uses a heat exchanger of lower capacity or quality.

Given the large number of network parameters (several hundred thousand for a network like the one we display here), this is done best by machine learning from data. As mentioned above, several techniques and algorithms exist for this task and Bayes net tools typically come with one or more readily available. In our experience, all of them provide suitable results given adequate data – and the latter is the challenge.

For an effective data-based instantiation, machine learning algorithms require a dataset that captures how the system under investigation operates in the largest possible health and behavior space. For instance, we require data from faulty or deteriorating equipment to observe the impact and range of effects it shows to provide the Bayes net with the ability to reason backwards from observable effects to root causes. However, as it happens with real world systems, large data sets are often only available for times of without issues while little or no data is available that illustrates malfunctions or other unwanted effects (Sobie, 2018).

In our work, we address this unbalance by using a physical model to produce synthetic data that complements real-world operational data with simulated observations of rare events. Our physical model is based on the thermal equilibrium between the primary and secondary liquids of the vessel's cooling system. In the 1D geometry shown in Figure 6, this equilibrium is described by Eq. (1)

$$\begin{cases} f_p \frac{\partial T_p}{\partial x} = \gamma(T_s - T_p) \\ f_s \frac{\partial T_s}{\partial x} = \gamma(T_s - T_p) \end{cases} \quad (1)$$

where T_s is the temperature of the water in the vessel system, T_p of the sea water, $f_{s,p}$ the respective flow rates and γ is a degradation variable, for example describing the degrading efficiency of the heat exchanger.

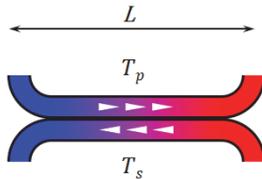


Figure 6. Scheme of 1D geometry.

By solving Eq. (1) with boundary conditions

$$\begin{aligned} T_p(0) &= T_{p1} \\ T_s(0) &= T_{s1} \\ T_s(L) &= T_{s2} \end{aligned}$$

we obtain the following expression for the degradation variable:

$$\gamma L = -\frac{f_s f_p}{f_p - f_s} \log \left[1 + \left(1 - \frac{f_s}{f_p} \right) \frac{T_{s1} - T_{s2}}{T_{s2} - T_{p1}} \right] \quad (2)$$

Equation (2) then allows us to simulate for each time instant, the effect of the degradation by solving on a variable of interest on the right-hand side for a chosen value of γL . This is shown for T_{s2} in Figure 7.

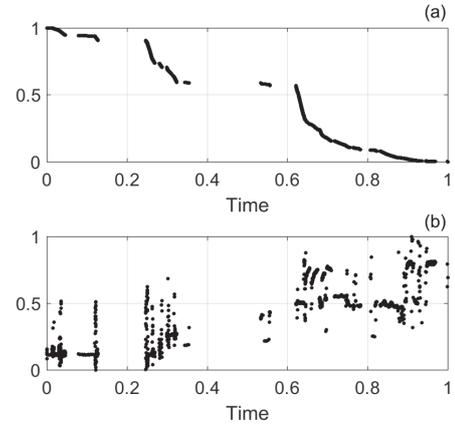


Figure 7. Simulated synthetic data. (a) The degradation. (b) The temperature of the water in the vessel. Values are rescaled between 0 and 1.

Specifically, Figure 7 (a) shows for each time stamp the chosen value of a degrading γL and Figure 7 (b) shows the corresponding T_{s2} calculated using Eq. (2) for given measurement data f_s, f_p, T_{s1}, T_{p1} and γL . The time span is of around 3 months while the temperature span for T_{s2} is of about 3°C. Figure 7 shows, as expected, that a degradation in the cooling system manifests as an increase in the temperature of the water entering back in the vessel. Discontinuities in the time are due to absence of measurement data but pose no issue for the Bayesian techniques we use.

Following this approach, we produce data of degradation scenarios for the variables in the right-hand side of Eq. (2) and construct a data set that we use to learn the conditional probability distributions of the Bayes net with maximum likelihood estimation (Barber, 2012), thus instantiating it.

Further information on the possible need to complement operational data can be found in (Borth, 2018), where the authors illustrate when even large data sets are insufficient. (Mas, 2017), e.g., illustrates another successful case where simulations solved this challenge.

2.3. Diagnostic Reasoning

The diagnostic ability of the created Bayes net is two-fold: First, it determines the system's ability to bring the target variable, i.e., the temperature of the cooling water into the intended range, which addresses mission-readiness. Furthermore, it resolves possible causes if this is not the case, e.g., by reasoning that the heat exchanger is not effective or that the pump does not work.

Understood and used like this, diagnostic reasoning centers on one specific situation and one point in time. We will relax these boundaries in the next chapter where we address the mission readiness assessment for future scenarios. However, as the diagnosis process lays the foundation for that, we detail it first in this section.

Diagnosis with a Bayesian network like ours works by an instantiation of observable variables; one provides evidence to the network, e.g., on the measured temperatures and the age of the components. Given the conditional probability distributions that specify the relationship between variables, the algorithms for Bayesian reasoning then update the other variables by computing their marginal probabilities, i.e., the likelihoods of each possible state. Mathematically, the Bayes net encodes the joint probability distribution over all modeled variables and an update process like this conditions said distribution with the evidence according to the Bayes rule, ensuring consistent probabilities for all states.

Practically, we see a bi-directional reasoning process from cause to effects and from observed effects to possible causes. For the cooling system, e.g., the Bayes net, raises the likelihood of the states that encode lower efficiency if we provide the information that the heat exchanger is old. If we furthermore provide evidence on the observed temperatures, it reasons about the cooling effect: If the liquid leaving the heat exchanger is warmer than intended, the reasoning finds the hypothesis of an inefficient component confirmed, it reduces the uncertainty of this diagnosis. If, on the other hand, the cooling effect is strong, the diagnosis becomes uncertain, an efficient heat exchanger is no longer ruled out, but the reasoner also checks for alternate explanations, like a high flow of water and liquid that would compensate for the expected loss of efficiency.

Given that the Bayes net can reason simultaneously from cause to effect and from effect to cause, it covers both diagnosis and prognosis. In our case, it can, e.g., predict the outgoing cooling water's temperature if we provide it with information on the component's age and maintenance status together with the observations on the incoming sea water's temperature and state information of the vessel's systems that generate heat. This set of information items summarizes a specific situation for the cooling system, encompassing both the circumstances (sea water temperature) and the task (cooling down the ship given a specific load). We use this in the next chapter, where it becomes part of a scenario.

3. SCENARIO-BASED MISSION READINESS

With the probabilistic model and the associated diagnostic reasoning in place, we can now investigate the system's readiness to perform a mission in a given scenario.

Scenarios describe concrete and easy to understand situations or episodes in which a sequence of activities resolve a problematic task which is often of a transformative nature. They can be defined as an account of a possible set of events that might reasonably take place (Jarke, 1998) and they fit well with our notion of mission readiness given their mesh into requirements engineering that Filippidou (1998) summarizes.

In our simplified example, we started to describe this at the end of the previous section: all components' ages and maintenance statuses together with the incoming sea water's temperature form the circumstances for the task, which is to cope with the heat that the vessel's systems generate to fulfill the set mission goals.

Using the probabilistic reasoning, we compute likelihoods of all possible states of variables that represent the system's ability to do so, i.e., the outgoing cooling liquid's temperature, as a failure to keep that one within acceptable ranges translates into a failure to perform the mission. Figure 8 shows such an investigation of four related scenarios that we use for decision support regarding the need to maintain a system prior to its next deployment:

Ideally, we want a probability of 1 for the target temperature (center), which translates to the statement that the system will always, i.e., under any circumstances, keep the temperatures low. Realistically, we require that the probability for the target temperature is high, and that high temperatures, (+3, hot) are very unlikely. This translates to thresholds for the probabilities that are depicted in the figure as two dotted lines. Such thresholds are application-dependent certainty-levels, expressing acceptable risks which stem from a calculation of expected costs in the event of failure.

In the first scenario (a), which describes the current status and task of our system, the center probability is above the upper dotted line, all others are below the lower – the risks are acceptable. The second scenario (b), however, projects the system's performance for the same task in the future, i.e., reasons with a higher age of all components.

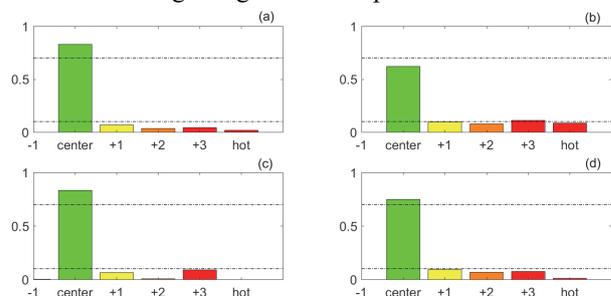


Figure 8. Mission readiness for four scenarios.

We see that sufficient cooling is no longer given. Scenario (c) checks if the aged system can cope with a lower heat generation, which is the case; we have the option to send the vessel on a lighter mission. Alternatively, we can run the high load mission after maintenance, as scenario (d) shows with acceptable probabilities.

3.1. Concurrent Mission Readiness Assessment

The scenario-based mission readiness computation above works best if (i) the Bayesian reasoning has consistent and accurate information for the technical diagnosis together with a specific instantiation of the mission and its settings and (ii) the degradation of the system depends primarily on its age.

The first of these two points translates to minimal uncertainty within the interpretation of observations; the second allows for an easy forecast into a foreseeable future, e.g., for us to predict the mission readiness at the end of an extended mission by setting the appropriate age. Both points oversimplify the underlying complexity. Sensor data for diagnosis, e.g., is often missing or inaccurate, especially for temperature readings that face heat transfer across components or environmental effects. Furthermore, there are causal noise factors, like drip leaks that affect pressures. Wear and tear over several months is, on the other hand, often dependent on additional factors, including said environmental effects but also the strain due to the nature and intensity of the system’s task.

Even in relatively simple adaptive systems, like the one we investigate here, such extra factors and their interdependencies are typically neither well understood nor captured in diagnostic models. The interdependencies, including their history, between different component of the system introduce additional feedback cycles that are relevant for prediction of mission readiness over longer time-spans. We address this here by linking two data-driven parts: the machine learning part of our model generation which we described in Section 0 and a concurrent mission readiness assessment.

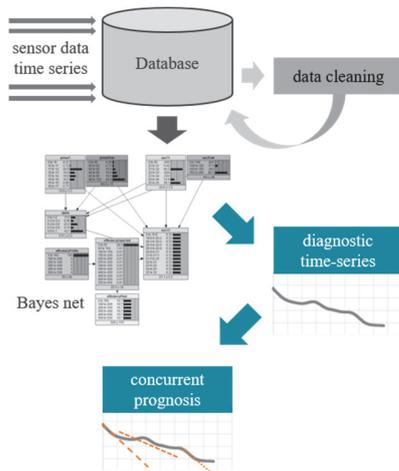


Figure 9. Concurrent mission readiness assessment.

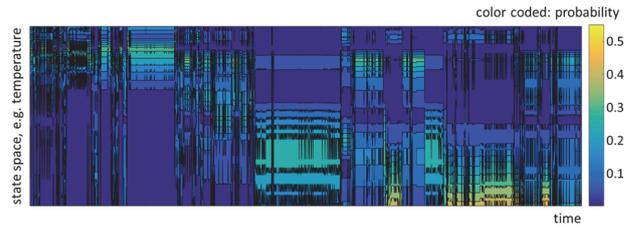


Figure 10. Health assessment probability field. It shows the color-coded probability of all investigated states over time. The probabilities of all states at a point in time add up to 1.

Our approach is schematically shown in Figure 9. We use the Bayes nets to create a diagnostic time-series for our mission readiness and use this in a forecast. As the Bayes net is instantiated via the machine learning step that captures many of the implicit interdependencies, we thus generate a concurrent prognosis that responds to the effects of age, task, strain, wear and tear, and the resulting adaptation.

Each prognosis, which we can run at arbitrary time intervals, provides us with a probabilistic assessment forecast. Specifically, the underlying data that our method provides is a probability field over time, like the one in Figure 10, where the vertical cross sections at each time stamp represents a probability distribution over the states, like those of the subgraphs of Figure 8.

While this is more complex than a scalar time-series (like the related Figures 7 (a) and 11, see below), the probability field offers a wealth of information: it supports, e.g., prognosis and health assessments that are fine-tuned to specific needs, e.g., to a risk averse strategy that requires an absence of any substantial likelihood of mission failure.

The prognosis part for the concurrent mission readiness assessment as shown in Figure 9 is detailed in Figure 11. Here, as an example, we reduce the probability field of Figure 10 to a scalar time-series by calculating the weighted mean for each time stamp, with weights the probability per state. For prognosis we then use linear regression, an analysis technique widely adopted for prediction and forecasting (Freedman, 2009). Figure 11 shows three prognosis results at different time-steps, black lines, together with their uncertainty interval, red lines, which grows subsequently smaller.

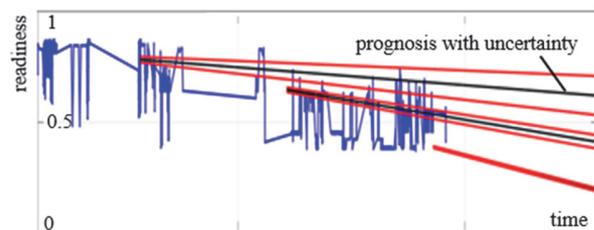


Figure 11. Mission readiness prognosis.

4. CONCLUSION

The concurrent mission readiness assessment we introduced in the last section concludes our method to predict the ability of a system to perform long-running missions. Based on its mission readiness diagnosis for which we detailed the generation process of Bayes nets for probabilistic reasoning, we see if the system requires maintenance prior to the mission or if the system's adaptive capabilities will allow it to cope with both existing and future degradation. In this, our approach integrates over a system's adaptive behavior, circumstances, as well as the technical state of the system's components. Together with the subsequent assessment of this readiness within given scenarios, it thus exceeds known work that determines health and maintenance needs solely based on estimates for the remaining useful life of components.

We consider the introduced health computations that account for a system's adaptation to circumstances and the technical state of its components to be crucial for modern, data-driven systems that constantly change their mode of operations. Seeing how complex relations within such systems complicate a forecast of their behavior, we cover prognosis with a dual approach that first computes in the relationship of age to degradation, allowing for an assessment at a future point in time, and second integrates a concurrent prognosis that updates given recent operational data.

This last part of the prognosis allows us to respond to recent developments, e.g., high strain operations that lead to a component's early degradation. It is, however, also an area of future work: We believe it to be possible to model the effects of such feedback loops in more detail regarding their effects over time. We can, e.g., extend our diagnostic network, which just covers them as part of ongoing effects, towards a time-sliced network that separates effects as per their duration. This should further increase the accuracy of our prognosis of mission readiness for systems that face abnormal spikes in their wear and tear due to periods of extreme and non-sustainable usage. Furthermore, we intend to link diagnosis and prognosis to an active monitoring that detects deviations from expectations, e.g., as a result from unforeseen extreme usage, as this will allow us to trigger a timely re-evaluation of maintenance needs.

ACKNOWLEDGEMENT

The research is carried out as part of the SFS project under the responsibility of ESI with Thales Nederland B.V. as the carrying industrial partner.

The SFS research is supported by the Netherlands Ministry of Economic Affairs.

REFERENCES

- Barber, D. (2012). *Bayesian reasoning and machine learning*. Cambridge University Press.
- Baru, A. (2018). *Three Ways to Estimate Remaining Useful Life for Predictive Maintenance*. Technical Article. MathWorks. nl.mathworks.com/company/newsletters/articles/three-ways-to-estimate-remaining-useful-life-for-predictive-maintenance.html
- Berthelsen, E. (2018). *Market Trends: Predictive Maintenance Drives IoT in Manufacturing Operations*. Report. Gartner. www.gartner.com/doc/3856379/market-trends-predictive-maintenance-drives
- Borth, M., & von Hasseln, H. (2002). Systematic Generation of Bayesian Networks from Systems Specifications. *Intelligent Information Processing* (pp. 155-166), Montreal. doi: 10.1007/978-0-387-35602-0_14
- Borth, M., & van Gerwen, E. (2018). Data-driven Aspects of Engineering. *13th Annual Conference on System of Systems Engineering* (pp. 219-224), Paris. doi: 10.1109/SYSOSE.2018.8428711
- Filippidou, D. (1998). Designing with scenarios: a critical review of Current Research and Practice. *Requirements Engineering*, 3, pp. 1-22. doi=10.1007/BF02802918.
- David A. Freedman (2009). *Statistical Models: Theory and Practice*. Cambridge University Press.
- Hardesty, L. (2015). *Probabilistic programming does in 50 lines of code what used to take thousands*. phys.org/MIT.phys.org/news/2015-04-probabilistic-lines-code-thousands.html
- Heckerman, D., Mamdani, A., & Wellman, M.P. (1995). Real-world applications of Bayesian networks. *Communications of the ACM*, vol. 38, 3, pp. 24-26. doi=10.1145/203330.203334
- Jarke, M., & Bui, T., & Carroll, J. (1998). Scenario Management: An Interdisciplinary Approach. *Requirements Engineering*, 3, pp. 155-173. doi=10.1007/s007660050002.
- Jensen, F.V. (2007). *Bayesian Networks and Decision Graphs*. Ney York: Springer.
- Mas, P. Deep Learning & deep analysis: The interplay between neural networks and complex system simulations. SiA, 2017-AT-02-04. www.sia.fr/publications527-deep-learning-deep-analysis-the-interplay-between-neural-networks-and-complex-system-simulations
- Mobley, R. K. (2002). *An Introduction to Predictive Maintenance*. 2nd Edition. Oxford: Butterworth-Heinemann.
- Mulders, M., & Haarman, M. (2017). *Predictive Maintenance 4.0*. Report. PcW, The Netherlands. www.pwc.nl/nl/assets/documents/pwc-predictive-maintenance-4-0.pdf
- Pearl, J. (1986). Fusion, propagation, and structuring in belief networks. *Artificial Intelligence*, 29 (3): pp. 241-288. doi:10.1016/0004-3702(86)90072-X.
- Pearl, J. (2009). *Causality*. Cambridge University Press.

- Pfeffer, A., & Koller, D. (1997). Object-oriented Bayesian networks. In *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence* (pp. 302-313). Morgan Kaufmann Publishers.
- Ricks, B. W., & Mengshoel, O. J. (2009). Methods for probabilistic fault diagnosis: An electrical power system case study. Annual Conference of the Prognostics and Health Management Society.
- Si, X., Wang, W., Hu, C., & Zhou, D. (2011). Remaining useful life estimation – A review on the statistical data driven approaches. *European Journal of Operational Research*, vol. 213 (1), pp. 1-14.
doi: 10.1016/j.ejor.2010.11.018
- Snowden, D. J., & Kurtz, C. F. (2003). The new dynamics of strategy: Sense-making in a complex and complicated world. *IBM Systems Journal*, 42 (3), pp. 462-483.
doi:10.1147/sj.423.0462.
- Sobie, C., Freitas, C., & Nicolai, M. (2018). Simulation-driven machine learning: Bearing fault classification. *Mechanical Systems and Signal Processing*, 99, pp. 403-419.