Mathematical and Computer Modelling of Dynamical Systems: Methods, Tools and Applications in Engineering and Related Sciences

Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/nmcm20

Centralized Bayesian reliability modelling with sensor networks

K. Dedecius [a] & V. Sečkárová [a b]

[a] Institute of Information Theory and Automation, Academy of Sciences of the Czech Republic, Pod Vodárenskou věží 4, 182 08, Prague, Czech Republic

[b] Department of Probability and Mathematical Statistics, Faculty of Mathematics and Physics, Charles University in Prague, Sokolovská 83, 186 75, Prague, Czech Republic

Published online: 17 Apr 2013.

To cite this article: K. Dedecius & V. Sečkárová (2013): Centralized Bayesian reliability modelling with sensor networks, Mathematical and Computer Modelling of Dynamical Systems: Methods, Tools and Applications in Engineering and Related Sciences, DOI:10.1080/13873954.2013.789064

To link to this article: http://dx.doi.org/10.1080/13873954.2013.789064

Mathematical and Computer Modelling of Dynamical Systems, 2013
http://dx.doi.org/10.1080/13873954.2013.789064

Centralized Bayesian reliability modelling with sensor networks

K. Dedecius[a]* and V. Sečkárová[a,b]

[a]*Institute of Information Theory and Automation, Academy of Sciences of the Czech Republic, Pod Vodárenskou věží 4, 182 08 Prague, Czech Republic;* [b]*Department of Probability and Mathematical Statistics, Faculty of Mathematics and Physics, Charles University in Prague, Sokolovská 83, 186 75 Prague, Czech Republic*

(*Received 20 June 2012; final version received 19 March 2013*)

The article concerns reliability estimation in modern dynamic systems. It introduces a novel approach, exploiting a network of several independent spatially distributed sensors, actively probing the monitored system. A dedicated network element – the fusion centre – is then responsible for processing the information provided by sensors and evaluation of final reliability estimate. On the base of computational abilities of sensors, we propose two conceptually different reliability estimation scenarios: (1) the computationally cheaper dummy sensors scenario, in which the sensors send raw data to the fusion centre; and (2) the smart sensors scenario, when the data are processed locally by sensors, and the fusion centre subsequently merges their resulting information. The local processing allows to obtain 'low-level' reliability estimate from a particular sensor, which is of interest in large networks with communication constraints. In both cases, the emphasis is put on recursiveness, adaptivity and robustness of solutions. The Bayesian paradigm was adopted for consistent information representation, its adaptive dynamic processing and fusion.

Keywords: Bayesian modelling; sensor network; reliability; dynamic system monitoring

1. Introduction

Recent advances in networked communications and computing performance of electronic devices enabled rapid development of low-cost multifunctional sensor networks, in which the nodes effectively sense, communicate and process obtained data to evaluate the most accurate measurements. Unless there are special requirements (disaster management, space exploration, factory automation, etc.) [1], the topology of sensor networks is usually very flexible and the position of particular sensors need not be engineered nor predetermined [2]. Besides this, the sensor network provides more information about the observed system than a single sensor [3] and allows for fault tolerance and graceful degradation [4,5], robustness, etc. Several surveys are given, e.g. see [2,6,7]. The novelty of this article consists in the use of sensor network for reliability estimation.

We consider a network of spatially distributed sensors $S_1, S_2, \ldots, S_K$ with a fusion centre (Figure 1), devoted to monitoring a dynamic system (e.g. a computer cluster) and evaluating its reliability $\pi_t \in [0, 1]$ in discrete time instants $t = 1, 2, \ldots$. Generally, the International Organization for Standardization defines reliability as 'the characteristic

*Corresponding author. Email: dedecius@utia.cas.cz

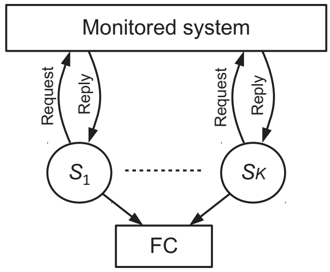

Figure 1. Sensor network: active probing of a system by a network of K sensors $S_1, \ldots, S_K$, sending their information to a fusion centre, FC.

of a product, or any component thereof, expressed as a probability that it would perform its required functions under defined conditions for specified operating periods' (ISO 9001:2008). In our conception (following the ISO definition), reliability is understood as a probability that the monitored system successfully responds to each request. That is,

$$\pi = \text{reliability} = \frac{\text{number of successess}}{\text{number of all requests}}. \tag{1}$$

From the variety of classical reliability estimation methods, e.g. based on queuing theory [8], combinatorial models and Monte Carlo methods [9], Markov models [10] and Bayesian approaches [11,12], we choose the Bayesian paradigm for its appealing consistency and versatility. The exploited monitoring principle is known as *active probing* [13,14]; the packet-probing technique is an example of its application in computer networks [15,16]. Active probing consists in periodic sending of short *request* messages to the monitored system, followed by awaiting *reply* messages to be received within predefined time periods.

On this base, two novel network-based reliability estimation methods reflecting computational abilities of sensors are proposed:

- Dummy sensors setting – all sensors send raw measurements (numbers of failures) to the fusion centre. It evaluates the reliability estimates using the weighted likelihood principle [17]. This is computationally cheap and requires only low-cost sensors.
- Smart sensors setting – all sensors process their measurements locally and then provide the posterior information to the fusion centre. It builds and subsequently reduces a finite mixture of these posteriors [18]. Besides the global estimate, a (partial) information is available at each node and can be exploited, e.g. if the fusion centre or its links fail. This is of interest in mission critical applications. The computational burden in sensors is low and requires only slightly more expensive hardware than in the case of dummy sensors. The burden in the fusion centre is higher due to the mixture reduction.

In both cases, it is possible to process data in bursts instead of one-by-one and therefore to save sensor network's resources. It is also possible to formulate an intermediate setting: separate processing of dummy sensors' raw measurements in the fusion centre, followed by composing a mixture and its subsequent reduction. However, this special case only increases the computational burden in the fusion centre without providing the benefits of the smart sensors setting.

The consistent dynamic Bayesian treatment of available information makes both methods easily tunable – the only user-dependent setting parameters are forgetting factors, driving the effective number of past data that should be used for modelling. A higher number of data leads to more stable estimation, at the cost of slower response to changes (and vice versa). The inevitable fact that the sensors and their links are subjects to degradation and failure, both increasing the number of missing replies, is also properly reflected. This in turn allows easy addition and removal of sensors during the runtime. The proposed methods are the first steps towards more sophisticated setting, in which the replies to probes are categorized (e.g. fast/slow/no response) and probabilities of these categories are adaptively evaluated.

The organization of the article is as follows: Section 2 presents the main principle of dynamic Bayesian reliability modelling with a single sensor; Section 3 formulates two methods for adaptive data processing – for dummy and smart sensors, respectively. Algorithms are provided for both methods. Finally, Section 4 provides an illustrative example. Since the authors are not aware of any similar sensor network-based approach for reliability estimation, the example only compares the two proposed methods.

2. Bayesian setting of the problem

For the sake of clear exposition, the ensuing text first introduces the classical principles of Bayesian reliability modelling with a single sensor (e.g. [12]), standing also for information processing unit. With the necessary theory, we proceed to the more complicated multiple sensors case.

We emphasize that the time instants $t = 1, 2, \ldots$ relate to data processing. The monitoring period may be shorter, which means that the *number of requests* $n_t \geq 1$ may be processed at each t, yielding the *number of failures* $m_t \leq n_t$ and the *number of successes* $(n_t - m_t)$. These numbers are referred to as *measurements* at time t. After their processing, m_t is reset to 0.

2.1. Single sensor case

Based on the *active probing* monitoring scheme, we can model the number m_t of missing reply messages (regardless whether the monitored system effectively sent them or not) at time instant t with the binomial distribution, $m_t \sim \mathrm{Bi}(n_t, \pi_t)$, where $n_t \geq 1$ is the number of all *request* messages and $\pi_t \in [0, 1]$ is the probability of failure. If $n_t = 1$, then the binomial distribution is a Bernoulli distribution. The Bayesian paradigm then advocates the beta distribution as the natural conjugate prior to this type of model under unknown π_t, i.e. $\pi_t \sim \beta(r_t, s_t)$, where r_t and s_t are real positive *hyperparameters* at time t, accumulating the number of failures and successes, respectively. They drive the shape of the beta distribution, allowing for left- and right-skewness, high or low kurtosis and even a flat shape of the uniform distribution. Let $q(m_t|\pi_t, r_t, s_t)$ be the conditional probability density function (pdf) of m_t, i.e. its binomial model,

$$q(m_t|\pi_t, r_t, s_t) = \binom{n_t}{m_t} \pi_t^{m_t} (1 - \pi_t)^{n_t - m_t}. \tag{2}$$

and $p(\pi_t|r_t, s_t)$, the beta prior pdf of π_t,

$$p(\pi_t|r_t, s_t) = \frac{1}{B(r_t, s_t)} \pi_t^{r_t-1}(1-\pi_t)^{s_t-1},$$

where

$$B(r_t, s_t) = \int_0^1 z^{r_t-1}(1-z)^{s_t-1}\mathrm{d}z = \frac{\Gamma(r_t)\Gamma(s_t)}{\Gamma(r_t+s_t)}, \qquad r_t, s_t > 0$$

is the beta function. The initial values r_0, s_0 are either preset by the user or a flat noninformative beta distribution $\beta(1, 1)$ is chosen. The Bayes' theorem [19] recursively updates the prior beta distribution by new binomially distributed data, yielding again the beta distribution as the posterior pdf,

$$p(\pi_t|r_t, s_t) \propto q(m_t|\pi_{t-1}, r_{t-1}, s_{t-1})p(\pi_{t-1}|r_{t-1}, s_{t-1}), \tag{3}$$

where $\propto$ denotes proportionality, i.e. equality up to a normalizing factor. It is straightforward to check that the theorem simply evaluates hyperparameters r_t, s_t as follows:

$$r_t = r_{t-1} + m_t, \qquad s_t = s_{t-1} + (n_t - m_t), \tag{4}$$

i.e. r_t and s_t are incremented by the numbers of failures and successes, respectively. The actual estimate of reliability π_t is given by the expectation, defined for the beta-distributed variable as

$$\mathbb{E}\left[\pi_t|r_t, s_t\right] = \frac{s_t}{r_t + s_t}. \tag{5}$$

This directly coincides with the reliability definition (1). It is important to notice that the reliability estimates lie in interval (0, 1), with 0 and 1 as limit cases. This corresponds with reality: no real system can be considered absolutely perfect and there is no natural need to monitor completely failed systems.

2.2. Adaptivity of a single sensor

The incoming measurements accumulated in the beta distribution augment statistical knowledge of reliability under the situation when the sensor and communication links are perfectly reliable and the monitored system has either constant or very slowly varying reliability. Since these conditions are rather rare, most applications call for adaptive solutions. To make the given approach adaptively reflect the temporal evolution of reliability $\pi_t \to \pi_t^+$, we would need a probabilistic evolution model, e.g. in the form of a Markov model with a transition kernel $K(\pi_t, \pi_t^+)$. However, it is practically unreachable. The way around the issue consists in continuous discounting of old and potentially outdated information from posterior pdf $p(\pi_t|r_t, s_t)$. It turns reliability estimation into reliability tracking. We exploit the celebrated *exponential forgetting* [20,21] (for other methods, see, e.g. [22] and the overview therein), transforming the posterior pdf into a pdf with a higher variance,

$$p(\pi_t^+|r_t^+, s_t^+) \propto \left[p(\pi_t|r_t, s_t)\right]^{\lambda}, \tag{6}$$

where $\lambda \in (0, 1]$ is called forgetting factor.

Table 1. Number d of effective samples for given forgetting factor λ.

λ	0.999	0.998	0.995	0.99	0.98	0.95
d	1000	500	200	100	50	20

It is straightforward to see that the exponential forgetting (6) applied on the posterior pdf from (3) modifies hyperparameters of the underlying distribution by their multiplication,

$$r_t^+ = \lambda r_t, \qquad s_t^+ = \lambda s_t. \tag{7}$$

The resulting pdf is determined by hyperparameters from (7) and can successively enter the Bayes' theorem (3) as the prior pdf for the next time step.

Similarly to the fixed-size data window approaches, the choice of the forgetting factor λ determines the effective window size $d = (1 - \lambda)^{-1}$, e.g. see [21]. It expresses the number of past measurements stored in r_t and s_t used for subsequent modelling. Some selected pairs of values of λ and d are depicted in Table 1.

3. Sensor network

Suppose now the existence of a sensor network with K sensors $S_1, \ldots, S_K$. We will present two scenarios for information fusion, one suitable for dummy sensors providing the fusion centre with raw measurements (m_t and, if not fixed, n_t) and the other for smart sensors with computational abilities, sending pre-processed information. The fusion centre calculates own pdf (denoted by g), comprising weighted information from all involved sensors. The reliability estimate is then inferred from this pdf.

The quality of estimation with both dummy and smart sensors heavily depends on the ability of the fusion centre to prevent spoiling the reliability estimator by corrupted information. This is likely to happen if one or more sensors degrade or fail, or if a link between a sensor and the monitored system is affected by high noise or in any sense broken. Such cases call for adaptive suppression or complete elimination of related sensors' contribution.

3.1. Dummy sensors

The dummy sensors scenario exploits the weighted likelihood principle [17], providing means for combination of statistical models, in our case of type (2). The fusion centre is responsible for evaluation of both the Bayes' rule (3) and the forgetting step (6). The sensors only provide raw measurements of the number m_t of failures (and if not preset, also n_t). In this setting, the Bayesian update (3) changes its form to

$$g(\pi_t|R_t, S_t, \mathcal{A}_t) \propto g(\pi_{t-1}|R_{t-1}, S_{t-1}, \mathcal{A}_{t-1}) \prod_{i=1}^{K} q(m_{i;t}|\pi_{t-1}, r_{i;t-1}, s_{i;t-1})^{\alpha_{i;t}}, \tag{8}$$

where $g(\pi_{t-1}|R_{t-1}, S_{t-1}, \mathcal{A}_{t-1})$ is the prior beta distribution with scalar real hyperparameters R_{t-1} and S_{t-1}. It is kept by the fusion centre. $\mathcal{A}_t = \{\alpha_{1;t}, \ldots, \alpha_{K;t}\}$ is a set of relative weights of sensors $S_1, \ldots, S_K$, taking values in $[0, 1]$ and summing to unity. They express the degree of belief of the fusion centre in information provided by each of the sensors. Their adaptive evaluation will be described below.

It is easy to see that the Bayesian update (8) recomputes the hyperparameters of the beta distribution local to the fusion centre by new measurements from the sensor network similarly to (4):

$$R_t = R_{t-1} + \sum_{i=1}^{K} \alpha_{i;t} m_{i;t}, \qquad S_t = S_{t-1} + \sum_{i=1}^{K} \alpha_{i;t}(n_t - m_{i;t}).$$

The exponential forgetting (6) preserving adaptivity of the reliability estimation takes an equivalent form,

$$R_t^+ = \lambda R_t, \qquad S_t^+ = \lambda S_t, \quad \lambda \in (0, 1].$$

With these hyperparameters, the resulting 'forgotten' pdf enters (8) as the prior pdf for the next update. The reliability estimate counterpart of (5) is then the mean value of the fusion centre's beta distribution,

$$\mathbb{E}\left[\pi_t | R_t, S_t, \mathcal{A}_t\right] = \frac{S_t}{R_t + S_t}. \tag{9}$$

Note that the reliability estimate follows the general definition (1). Here, the numbers of successes and requests are represented by convex combinations of sensors' measurements, reflecting the degree of belief in them.

This approach allows to use very simple sensors, effectively represented by incremental counters of failures and (potentially) successes for all trials.

3.1.1. Evaluation of weights $\alpha_{i;t}$

Each of the weighted likelihood components is assigned a weight $\alpha_{i;t} \in [0, 1]$. The Bayes' rule updating these weights reflects how the actual measurements fit the model kept by the fusion centre,

$$\alpha_{i;t} \propto \alpha_{i;t-1} q(m_{i,t} | R_{t-1}, S_{t-1}). \tag{10}$$

The model on the right-hand side of (10) is the predictive pdf of the fusion centre,

$$q(m_{i,t}|R_{t-1}, S_{t-1}) = \int q(m_{i,t}|\pi_{t-1}, R_{t-1}, S_{t-1})\, g(\pi_{t-1}|R_{t-1}, S_{t-1}, \mathcal{A}_{t-1}) \mathrm{d}\pi_{t-1}.$$

We exploit the fact that $m_{i,t}$ is conditionally independent of weights $\mathcal{A}_{t-1}$. Since the weights in (10) are updated according to actual measurements, the method is close to model switching.

Obviously, (10) updates the weights by the predictive pdfs with respect to the sensors' data. Up to here, the situation when all sensors work well is reflected and it remains to take their potential malfunctioning into account. The complete lack of a temporal evolution model for sensors' degradation immediately resembles that one in Section 2.2. Again, it is convenient to address the problem by exponential forgetting,

$$\alpha_{i;t}^+ = \alpha_{i;t}^{\omega}, \qquad \omega \in (0, 1].$$

Algorithm 1 Dummy sensors

Fusion centre initialization:
Set prior hyperparameters R_0 and S_0.
Set prior weights $\mathcal{A}_0$, e.g. uniformly $\alpha_{1;0} = \alpha_{2;0} = \ldots = \alpha_{K;0}$.
Set forgetting factors λ and ω, Table 1.
Online steps: for $t = 1, 2, \ldots$ **do**
    Sensors:
    for all the *sensors* $S_i \quad (i = 1, \ldots, K)$ **do**
        Set $m_{i;t} = 0$.
        Send n_t requests.
        Count $m_{i;t}$ failures.
    end
    Fusion centre:
    Input: $m_{i;t}, n_t (i = 1, \ldots, K)$.
    Update weights $\alpha_{i;t}$ by measurements, Equation (10).
    Update R_t and S_t by measurements.
    Estimate π_t, Equation (9).
    Forget R_t and S_t.
end

3.2. Smart sensors

In the smart sensors case, each sensor evaluates the Bayes' rule (3) and forgetting (6) internally. The fusion centre only merges their posterior distributions at each t. It exploits a convex combination of sensors' posterior beta distributions, i.e. a mixture, which is evaluated by the fusion centre [18]. The mixture represents a *consensus* of individual sources. Its use in reliability modelling with beta-binomial distributions in a different realm was proposed, e.g. in [23].

After the sensors process their measurements, they send the hyperparameters to the fusion centre to construct the mixture of posterior pdfs,

$$g(\pi_t | \mathcal{R}_t, \mathcal{S}_t, \mathcal{A}_t) = \sum_{i=1}^{K} \alpha_{i;t} p(\pi_t | r_{i;t}, s_{i;t}), \tag{11}$$

where the sets $\mathcal{R}_t = \{r_{1;t}, \ldots, r_{K;t}\}$ and $\mathcal{S}_t = \{s_{1;t}, \ldots, s_{K;t}\}$ formally represent the hyperparameters of all K sensors. Identically with the dummy sensors scenario, the set $\mathcal{A}_t = \{\alpha_{1;t}, \ldots, \alpha_{K;t}\}$ stands for relative weights of the sensors, such that $\alpha_{i;t} \in [0, 1]$ for all $i = 1, \ldots, K$ are summing to unity.

Now, the counterpart of the reliability estimates (5) or (9) is the merged point estimate. It is the expectation of the mixture (11),

$$\mathbb{E}\left[\pi_t | \mathcal{R}_t, \mathcal{S}_t, \mathcal{A}_t\right] = \sum_{i=1}^{K} \alpha_{i;t} \mathbb{E}\left[\pi_t | r_{i;t}, s_{i;t}\right] = \sum_{i=1}^{K} \alpha_{i;t} \frac{s_{i;t}}{r_{i;t} + s_{i;t}}. \tag{12}$$

In this scenario, the reliability estimate at each sensor follows the definition (1), more concretely (5). The result (12) is then a convex combination of all these estimates.

3.2.1. Evaluation of weights $\alpha_{i;t}$

Recall that the contribution of each sensor to the estimator (12) is driven by weights $\alpha_{i;t}$. Since, in this scenario, the fusion centre does not have access to raw measurements, it is impossible to use (10) for weights evaluation. Alternative approaches may consist in sensor elimination using a change-point detection, adaptive tuning based on time series analysis of hyperparameters and/or moments, or exploiting measures of dissimilarity between the ideal and true sensors' states. We focus on the last approach.

As the dissimilarity measure we exploit a member of the f-divergences family, namely the Kullback–Leibler divergence [24]. It is a natural choice in the Bayesian realm.

If $f(x)$ and $g(x)$ are two pdfs of a random variable X, acting on a common Borel set $\mathcal{X}$, their Kullback–Leibler divergence is defined as the functional

$$\mathcal{D}\,(f||g) = \int_{\mathcal{X}} f(x) \log \frac{f(x)}{g(x)} \mathrm{d}x.$$

The functional $\mathcal{D}\,(f||g)$ takes nonnegative values with equality if $f(x) = g(x)$ almost everywhere. It is a premetric, its lack of symmetry, i.e. $\mathcal{D}(f||g) \neq \mathcal{D}(g||f)$ and nonconformity with the triangle inequality prevents it from being a metric. Properties of the Kullback–Leibler divergence can be found, e.g. in [19].

The Kullback–Leibler divergence of two beta distributions $\beta(r,s)$ and $\beta(r',s')$ with pdfs f, g, respectively, is

$$\mathcal{D}(f||g) = \log \frac{B(r',s')}{B(r,s)} + (r'-r)\,[\psi(r+s) - \psi(r)] + (s'-s)\,[\psi(r+s) - \psi(s)] \tag{13}$$

where $\psi(\cdot)$ is the digamma function. The formula (13) is easy to prove. In terms of entropy $H(f)$ and cross-entropy $H(f,g)$, the divergence $\mathcal{D}(f||g) = H(f,g) - H(f)$ and from [25] it follows

$$H(f) = \log B(r,s) + (r-1)\,[\psi(r+s) - \psi(r)] + (s-1)\,[\psi(r+s) - \psi(s)]\,,$$

$$H(f,g) = \log B(r',s') + (r'-1)\,[\psi(r+s) - \psi(r)] + (s'-1)\,[\psi(r+s) - \psi(s)]\,.$$

The previously described divergence will be exploited in the following way: as a step towards metric, we first symmetrize the Kullback–Leibler divergence as proposed in the original paper [24],

$$\tilde{\mathcal{D}}\,(f||g) = \mathcal{D}(f||g) + \mathcal{D}(g||f). \tag{14}$$

With the help of (14), the actual information carried by sensors, represented by posterior pdf $p(\pi_t|r_{i;t}, s_{i;t})$, is then compared to ideal time-invariant pdf $p(\pi_I|r_I, s_I)$. Its fixed hyperparameters r_I and s_I represent the ideal scenario when there is no failure ($r_I = 1$) in the data window of length d, i.e. $s_I = d + 1$. The rule for $\alpha_{i;t}$ then reads as

$$\alpha_{i;t} \propto \frac{1}{\tilde{\mathcal{D}}\left(p(\pi_t|r_{i;t}, s_{i;t})||p(\pi_I|r_I, s_I)\right)}. \tag{15}$$

Algorithm 2 Smart sensors

Initialization of sensors:
for all the *sensors* $S_i \quad (i = 1, \ldots, K)$ **do**
    Set prior hyperparameters $r_{i;0}$ and $s_{i;0}$.
    Set forgetting factor λ, Table 1.
end
Initialization of fusion centre:
Set prior weights $\mathcal{A}_0$, e.g. uniformly $\alpha_{1;0} = \alpha_{2;0} = \ldots = \alpha_{K;0}$.
Online steps: for $t = 1, 2, \ldots$ **do**
    Sensors:
    for all the *sensors* $S_i \quad (i = 1, \ldots, K)$ **do**
        Set $m_{i;t} = 0$.
        Send n_t requests.
        Count $m_{i;t}$ failures.
        Update $r_{i;t}$ and $s_{i;t}$, Equation (4).
        Output: $r_{i;t}, s_{i;t}$
        Forget $r_{i;t}, s_{i;t}$, Equation (7).
    end
    Fusion centre:
    Input: $r_{i;t}, s_{i;t} \quad (i = 1, \ldots, K)$.
    Update weights $\alpha_{i;t}$, Equation (15).
    Estimate π_t, Equation (12).
end

Since the pdfs $p(\pi_t | r_{i;t}, s_{i;t})$ carry the information brought by the effective data window, there is no need of any form of additional forgetting. Rather than model switching, this scenario is close to model averaging. The drawback of this method follows from properties of the Kullback–Leibler divergence, preventing much stronger suppression of partially failing sensors (as shown in Section 4). We conjecture that other divergence measure could yet improve the estimator properties.

4. Example

In this short example, we model reliability of a system using three distributed sensors prone to errors, processing $n_t = 5$ requests each step of t. The length of the data is 1500 samples, $\pi_t = 0.99$ for $t = 1, \ldots, 300$ and $\pi_t = 0.6$ for the rest. The first two sensors work well while the third sensor degrades at $t = 150$, causing the measurements to indicate false reliability corresponding to half of the true value. To show recovery, the third sensor starts working well at $t = 900$. The simulation starts from informative prior pdf corresponding to the state that everything works well; all forgetting factors are 0.98. The evolution of estimated reliability is depicted in Figure 2, statistics of the estimation error $(\mathbb{E}[\pi_t | \cdot] - \pi_t)$ are given in Table 2.

The dummy sensors setting leads to an estimator with smaller bias compared to smart sensors. This is connected with the model switching property of the weights evaluation method. The smart sensors setting, measuring the divergence of the actual pdf from the ideal pdf, has faster response to abrupt changes, at the cost of higher bias. In both cases, the response is good and can be further tuned by forgetting factors. The degradation of a single sensor at $t = 150$ is quickly reflected by both approaches; the estimation is stable.

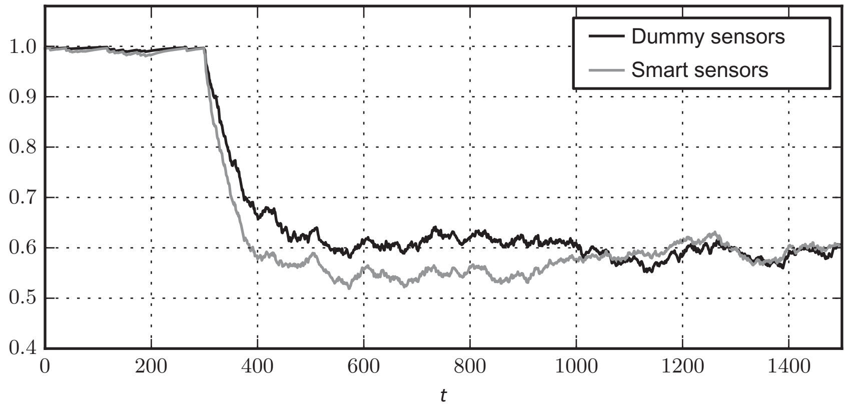

Figure 2. Estimated reliability.

Table 2. Error statistics of dummy and smart sensors setting, respectively.

	Dummy	Smart
Mean of errors	0.016	−0.011
Median of errors	0.005	−0.015
Standard deviation	0.058	0.053
Maximum negative error	−0.048	−0.081
Maximum positive error	0.393	0.386

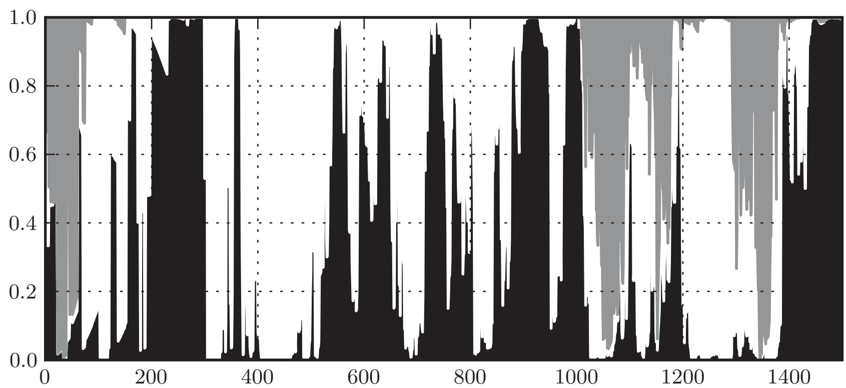

Figure 3. Dummy sensors setting: evolution of weights $\alpha_{i;t}$. Degrading node's weight is grey.

Also, the recovery of this sensor at $t = 900$ causes no instabilities. The switching and averaging properties of dummy and smart sensor settings are evident from Figures 3 and 4, respectively.

5. Conclusion and future work

We have presented a novel method for dynamical system monitoring with spatially distributed sensors, developed within the Bayesian statistical framework. It belongs to a class of centralized methods, in which each sensor collects and potentially processes data. The

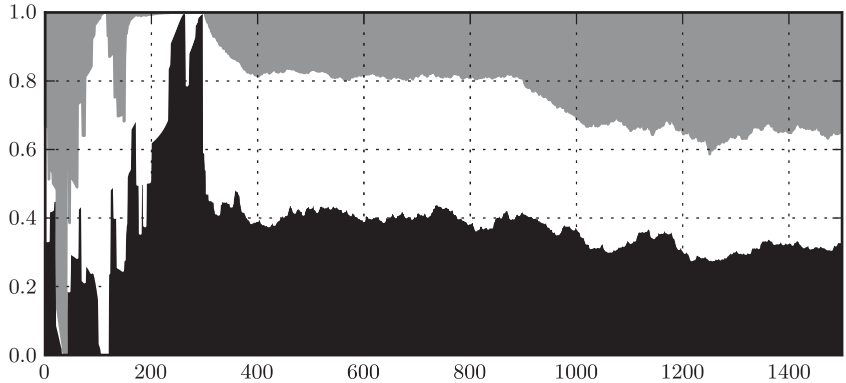

Figure 4. Smart sensors setting: evolution of weights $\alpha_{i;t}$. Degrading node's weight is grey.

data or the processed information is then transmitted to a dedicated network element, called fusion centre, responsible for final computations and reliability estimation. The method is dynamic; it allows to recursively incorporate information carried by new data and adaptively reflect changes of the observed reality.

The fact that the sensors and their links to the monitored system are subject to degradation and failure is properly reflected. An unsolved question is how to reflect potential delays in communication between the sensors and the fusion centre. A suitable solution of this issue would further improve robustness of the method and broaden its applicability in large global networks. An alternative solution, consisting in partial or full decentralization of reliability estimation is planned as well.

The proposed methods are the first steps towards more complicated setting, in which the responses to probes are categorized (e.g. fast/slow/no response) and the probabilities of these categories are adaptively estimated.

Acknowledgement

This work has been supported by the grants MEYS 7D12004 (E!7662 ProDisMon) and SVV 265 315.

References

[1] M. Younis and K. Akkaya, *Strategies and techniques for node placement in wireless sensor networks: a survey*, Ad Hoc Netw. 6 (4) (2008), pp. 621–655.
[2] I.F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, *A survey on sensor networks*, IEEE Commun. Mag. 40 (8) (2002), pp. 102–114.
[3] S.S. Ram, V.V. Veeravalli, and A. Nedić, *Distributed and recursive parameter estimation*, in *Sensor Networks, Signals and Communication Technology*, G. Ferrari, ed., Springer, Berlin, Heidelberg, 2009, Ch. 2, pp. 17–38.
[4] L. Paradis and Q. Han, *A survey of fault management in wireless sensor networks*, J. Netw. Syst. Manag. 15 (2007), pp. 171–190, doi:10.1007/s10922-007-9062–0.
[5] I. Dietrich and F. Dressler, *On the lifetime of wireless sensor networks*, ACM Trans. Sen. Netw. 5 (1) (2009), pp. 5:1–5:39.
[6] I.F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, *Wireless sensor networks: a survey*, Comput. Netw. 38 (4) (2002), pp. 393–422.
[7] C. Raghavendra, K. Sivalingam, and T. Znati, *Wireless Sensor Networks*, Springer, New York, 2004.

[8] S. Minkevičius, *The impact of overload conditions on computer network reliability*, Inform. Proc. Lett. 104 (1) (2007), pp. 32–35.
[9] I. Gertsbakh and Y. Shpungin, *Models of Network Reliability: Analysis, Combinatorics, and Monte Carlo*, CRC Press, Boca Raton, FL, 2010.
[10] J. Chiquet, M. Eid, and N. Limnios, *Modelling and estimating the reliability of stochastic dynamical systems with Markovian switching*, Reliab. Eng. Syst. Safe. 93 (12) (2008), pp. 1801–1808.
[11] P. Weber and L. Jouffe, *Complex system reliability modelling with dynamic object oriented Bayesian networks*, Reliab. Eng. Syst. Safe. 91 (2) (2006), pp. 149–162.
[12] M.S. Hamada, A.G. Wilson, C.S. Reese, and H.F. Martz, *Bayesian Reliability*, Springer, Berlin, 2008.
[13] M. Natu and A. Sethi, *Active probing approach for fault localization in computer networks*, 4th IEEE/IFIP Workshop on End-to-End Monitoring Techniques and Services, Institute of Electrical and Electronics Engineers (IEEE), Vancouver, BC, 3 April 2006, pp. 25–33.
[14] M. Brodie, I. Rish, S. Ma, G. Grabarnik, and N. Odintsova, *Active probing*, Tech. Rep. RC22817, IBM, Armonk, 2002.
[15] M.J. Luckie, A.J. McGregor, and H.-W. Braun, *Towards improving packet probing techniques*, Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement, IMW '01, ACM, New York, 1–2 November 2001, pp. 145–150.
[16] R. Mondragon, A. Moore, J. Pitts, and J. Schormans, *Analysis, simulation and measurement in large-scale packet networks*, IET Commun. 3 (6) (2009), pp. 887–905.
[17] X. Wang and J. Zidek, *Derivation of mixture distributions and weighted likelihood function as minimizers of KL-divergence subject to constraints*, Ann. Inst. Stat. Math. 57 (2005), pp. 687–701.
[18] S. Frühwirth-Schnatter, *Finite Mixture and Markov Switching Models* (Springer Series in Statistics), 1st ed., Springer, New York, 2006.
[19] J.M. Bernardo and A.F.M. Smith, *Bayesian Theory*, 1st ed., Wiley, New York, 1994.
[20] A.H. Jazwinski, *Stochastic Processes and Filtering Theory*, Academic Press, New York, 1970.
[21] V. Peterka, *Bayesian approach to system identification*, in *Trends and Progress in System Identification*, P. Eykhoff, ed., Pergamon Press, Oxford, 1981, pp. 239–304.
[22] K. Dedecius, I. Nagy, and M. Kárný, *Parameter tracking with partial forgetting method*, Int. J. Adapt. Control Signal Process. 26 (1) (2011), pp. 1–12.
[23] V. Savchuk and H. Martz, *Bayes reliability estimation using multiple sources of prior information: binomial sampling*, IEEE Trans. Reliab. 43 (1) (1994), pp. 138–144.
[24] S. Kullback and R.A. Leibler, *On information and sufficiency*, Ann. Math. Stat. 22 (1) (1951), pp. 79–86.
[25] C. Alexander, G.M. Cordeiro, E.M.M. Ortega, and J.M. Sarabia, *Generalized beta-generated distributions*, Comput. Stat. Data Anal. 56 (6) (2012), pp. 1880–1897.