Bayesian Statistical Analysis for Performance Evaluation in Real-Time Control Systems

Boström, Pontus; Heikkilä, Mikko; Huova, Mikko; Walden, Marina; Linjama, Matti

Published: 01/01/2015

TUCS

Pontus Boström | Mikko Heikkilä | Mikko Huova | Marina Waldén | Matti Linjama

Bayesian statistical analysis for performance evaluation in real-time control systems

Turku Centre *for* Computer Science

TUCS Technical Report
No 1136, June 2015

Bayesian statistical analysis for performance evaluation in real-time control systems

Pontus Boström
Åbo Akademi University, Finland, pontus.bostrom@abo.fi
Mikko Heikkilä
Tampere University of Technology, Finland, mikko.heikkila@tut.fi
Mikko Huova
Tampere University of Technology, Finland, mikko.huova@tut.fi
Marina Waldén
Åbo Akademi University, Finland, marina.walden@abo.fi
Matti Linjama
Tampere University of Technology, Finland, matti.linjama@tut.fi

TUCS Technical Report
No 1136, June 2015

Abstract

This paper presents a method for statistical analysis of hybrid systems affected by stochastic disturbances, such as random computation and communication delays. The method is applied to the analysis of a computer controlled digital hydraulic power management systems, where such effects are present. Bayesian inference is used to perform parameter estimation and we use hypothesis testing based on Bayes factors to compare properties of different variants of the system to assess the impact of different random disturbances. The key idea is to use sequential sampling to generate only as many samples from the models as needed to achieve desired confidence in the result.

Keywords: Statistical model-checking, verification and validation, stochastic systems, hypothesis testing, parameter estimation, Bayes factors

TUCS Laboratory
Distributed systems laboratory

1 Introduction

Model-based design in e.g. Simulink is a popular way to design control software, where the discrete controller controls a continuous-time system. In this approach, the control algorithms are designed together with a simulation model of the system to be controlled. This allows use of system simulation for controller validation even before the system is built. Typically the models are synchronous, i.e. computation is assumed to take no time. Delays are often modeled deterministically with the worst case or average case. This simplifies modelling and simulation is also typically faster. However, random delays and computation times might have a large impact on the system behaviour in practise. Including these delays makes analysis of the models harder.

Even if safety is important, sufficient performance of the system is crucial. One important question to answer is how different modifications to a model impact performance. Here we are mainly interested in the impact of stochastic disturbances on the system compared to the ideal synchronous system. To compare models we rely on hypothesis testing. In [17] they use Bayesian hypothesis testing based on Bayes factors for Statistical Model Checking (SMC). They have developed a procedure to generate samples that for a BLTL property ϕ that holds with a probability p choose to either accept the hypothesis $H_0 : p \geq \theta$ or $H_1 : p < \theta$, where θ is a user defined bound. Compared to numerical model checking techniques, the advantage of using SMC is that it is fast and easy to implement for various modelling frameworks, as one only needs to sample traces. We extend the approach to hypothesis testing in [17] to compare parameters for statistical models derived from the different system models. Sequential sampling is used to draw samples until a desired confidence in the results has been achieved. As simulations can take a very long time this is very useful. Small sample sizes are also important when applying the methodology to the actual final system. We check hypotheses such as, e.g., is the mean of the mean square error over a time interval greater in one model or the other, or, is the rate of events greater in one model or the other. By using several different hypotheses we can build a comprehensive suite of checks to *automatically* compare how different versions of models behave in the presence of stochastic disturbances with desired confidence in the result.

We are also interested in estimation of parameters for different random variables, in particular for validating the statistical models used in hypothesis testing. Bayesian statistical model checking [17] provides methods to estimate probabilities for desirable properties expressed in some logic to hold in stochastic models. They consider the satisfaction of a Bounded Linear Time Logic (BLTL) formula as a Bernoulli random variable and estimate the probability that the property holds. The procedure ensures that the real probability is within given bounds with given probability chosen by the user.

The estimation can be carried out to an arbitrary level of precision. In this paper we use the same approach as [17] to do more general Bayesian inference. We use the same idea of sequential sampling to sample from random variables to estimate parameters of statistical models, in particular random variables that can be seen as approximately Normal and Poisson distributed. This is used to estimate properties such as the average mean square error of a signal over a time interval or the rate that events occur. The approach can be extended to other distributions.

In this paper we compare two different variants with different disturbances of a model of digital hydraulic power management system (DHPMS) [15]. This is done to get a better understanding of how this system will perform in real life under non-ideal conditions. Our contribution is:

- A significant extension of the approach in [17] to compare properties of different models and to analyse more general properties than the probability of a BLTL formula to hold. Both parameter estimation and hypothesis testing is discussed.

- Application of the methodology to a case study with a discussion on possible extensions and limitations. The case study demonstrates issues often encountered when running controllers under non-ideal conditions.

In Section 2, we present the case study. In Section 3 we briefly discuss Simulink models and SMC and in Section 4 we describe the performance metrics used in the paper. Section 5 presents the statistical techniques used and Section 6 shows the application to the case study. Sections 7 and 8 present related work and conclusions.

2 Case study

The example used in this paper is a model of a six-piston digital hydraulic power management system (DHPMS) with two independent outlets. The hydraulic diagram of the machine is shown in Figure 1 but, for simplicity, only three pumping pistons are presented. Each piston can be connected to either one of the outlets A or B, or to the tank line T via on/off control valves. Hence, the DHPMS can operate as a pump, motor, and transformer. Furthermore, it can provide independent supply line pressures for the actuators. The pressure levels are kept as close as possible to user-defined reference values by utilising a model-based control approach to select optimal pumping and motoring modes for each piston [12]. Furthermore, the valve timing must be accurate to avoid too high and too low pressures (cavitation) in the piston chambers which can damage the system. The control algorithm for the optimal control signals is described in [8]. An example simulation of

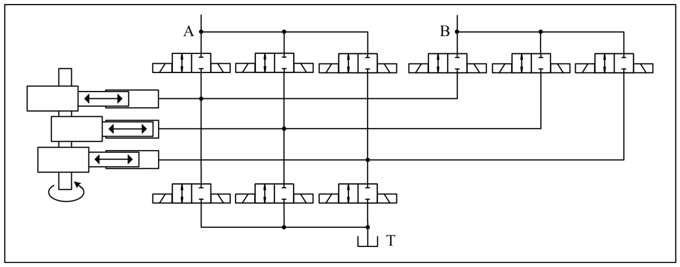

Figure 1: The hydraulic diagram of a DHPMS with three pistons. The figure shows three pistons connected to a rotating shaft on the left. Each piston is connected via valves to high pressure lines A and B, as well as a tank line T [15].

the complete synchronous system for 15 seconds is shown in Figure 2. Pressure tracking in the A and B-lines are shown to the left (two bottom-most graphs shows pressures p, the two topmost graphs shows fluid flows Q), while cylinder chamber pressures are shown to the right. As can be seen from the figures, the observed pressure (obs.) follows the reference (ref.) closely in both the A- and B-lines. The cylinder pressures are also below the high pressure limit of 20MPa.

2.1 Test systems

The goal is to investigate the impact of random time delays on system performance. This is used to assess how control algorithms constructed in an ideal model with deterministic timing will behave in practise. We compare two models:

System 1. In this model, the controller is assumed to be synchronous and all delays deterministic. Also the rotation speed of the motor is constant and known to the controller. This makes the model simpler and simulation is fast. The sampling period T_s is $50\mu s$, the delay d of opening and closing the valves are $1ms$. The rotation frequency of the motor is a fixed 25Hz.

System 2. In this model computation time is taken into account. The delay before reading inputs is assumed to be a uniformly distributed variable in the interval $[0.01T_s..0.1T_s]\mu s$. The computation time between the input is read until outputs are produced is assumed to be a uniformly distributed random variable in the interval $[0.1T_s..0.55T_s]\mu s$. The rotation speed of the motor is fixed at 25Hz, but the controller estimates the speed with an incremental rotary encoder [2]. Additionally, the valves open and close with a delay in the interval $[0.9d..1.1d]ms$.

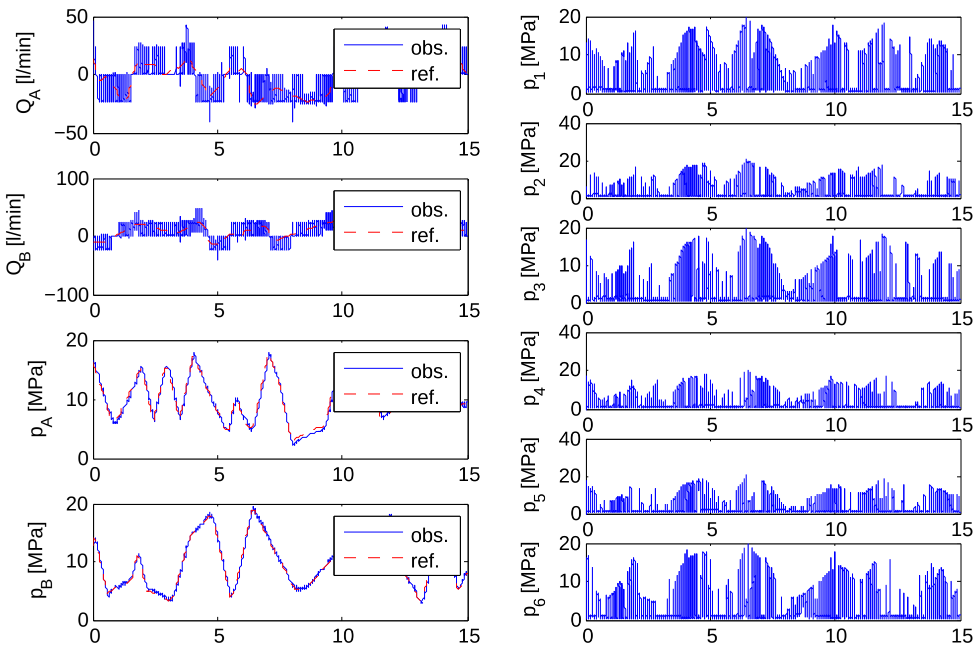

Figure 2: Results from a simulation of System 1. The figures to the left show flows Q and pressures p in the A and B-lines. The figures to the right shows pressures in the pumping cylinders 1..6.

System 3. This model has a slower sampling time than the other two models, $T_s = 100\mu s$. The valves are also slower with a delay of $4ms$. Computation times and valve delays are defined in the same way as in System 2.

3 Modelling and specification

We have used Simulink as our Modelling tool. A Simulink model can be assumed to be a probabilistic *discrete-time hybrid automaton* (DTHA) [17]. Hence, there is a well-defined probability measure over the trace space produced by simulating a model. Simulink/Stateflow is a complex language and a full formal semantics as a DTHA that accurately considers all features is difficult to define. However, we do not need a formal semantics, since we only sample traces. The sampling uses the built-in simulation capabilities of Simulink. Note that changing the simulation algorithm (differential equation solver) can change the semantics of the model, hence it is important to use the same configuration parameters for the simulation in all experiments. In [17] they use properties that can be checked on finite prefixes of simulations of a model, i.e. their truth-value depend only on a finite prefix. All statistics we compute apply only for finite prefixes of predefined length of simulation traces. Hence, we cannot draw conclusions about the behaviour of infinite traces from our statistical results. However, the considered simulation traces can be arbitrarily long. The models considered in this paper are open models, where the input signals are random signals with some properties, representing possible workloads of the system. The models are closed by also modelling the random input signals. Generating random input signals that are representative of real workloads in both the time and frequency domain can be challenging. However, a thorough investigation of this topic is outside the scope of this paper.

4 Definition of performance

We analyse performance using several different metrics. The goal of the metrics is to capture key performance properties of the system, where the metric for a model can be seen as a random variable. These metrics obtained from different versions of the system can then be used to compare them.

Mean square error. The DHPMS should provide pressure to the A- and B- lines as close as possible to the reference pressures. Hence, we need a characterisation of the deviation. One common way to characterise the size of the deviation is the mean square of the error signal v:

$$J = \lim_{T \to \infty} \int_0^T v(t)^2 dt \tag{1}$$

We cannot simulate the system for an infinite time and we therefore compute the mean square error for a finite time interval $[t_0, t_1]$, $J_n = \int_{t_0}^{t_1} v(t)^2 dt$ for each run n of the system. Then J_n is a random variable. We do not know the shape of $v(t)^2$, but if the system is time invariant then according to the *Central Limit Theorem* J_n is approximately normally distributed. This follows from the fact that J_n can be seen as the average of mean square errors of smaller intervals that then, if sufficiently long, are approximately independent and identically distributed (iid). Then if time invariance is again assumed, J is the average of all J_i with $E[J] = \lim_{N \to \infty} \frac{1}{N} \Sigma_{i=0}^{N} E[J_i] = E[J_i]$. Note that in special cases we can actually compute $E[J]$ exactly [14].

Safety properties. Safety properties can be formulated as BLTL properties and the probability that they hold can be directly estimated by the approach in [17]. One safety relevant performance property is to ensure that pressure peaks occur sufficiently rarely. We can analyse this property by computing the probability of a pressure peak during a certain time interval.

Rate of events. We are also interested in analysing how often good or bad events happen in the system. E.g. for the pressure in the pumping cylinders, we know that pressure peaks can occur. Furthermore, we are more interested in the event that the pressure becomes too high than the time the pressure stays too high. We can analyse this by analysing the rates of pressure peaks or low pressure (cavitation) events. If we assume the events are independent from each other, the number of events per time unit follows a Poisson distribution.

5 Statistical analysis

We cannot compute the properties of the random variables described earlier exactly due to system complexity and we instead use Bayesian statistics to analyse them. Bayesian inference is based on using the Bayes rule (2) [7] to fit a probability model to a set of data possibly using some prior information. The result is probability distribution on the parameters of the model.

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{p(y)} \tag{2}$$

Bayes rule gives the posteriori probability $p(\theta|y)$ where θ are the parameters of the model and y is the data. The notation $p(\theta|y)$ denotes the conditional probability for θ given y. The distribution $p(y|\theta)$ is called the sampling distribution and gives the probability distribution for observing y. The prior distribution $p(\theta)$ describes prior knowledge of θ, while $p(y)$ is a normalisation factor $p(y) = \int p(y|\theta)p(\theta)d\theta$ for continuous θ.

When data y has been observed we can use this to make predictions about an unknown observable $\hat{y}$ from the same process [7].

$$p(\hat{y}|y) = \int p(\hat{y}|\theta)p(\theta|y)d\theta \tag{3}$$

Hence, the prediction uses the posteriori probability density estimated for θ to estimate the probability for $\hat{y}$. Predictive distributions are used here to validate the models against the data.

The goal of parameter estimation is to estimate the probability of a certain property for any random simulation of a system. Each property can be evaluated on a prefix with fixed length of a simulation. When evaluating a BLTL formula we like to estimate the probability θ that the property holds in a random simulation. That the property ϕ holds for a simulation σ can be associated with a Bernoulli distributed random variable X. The conditional probability that the property holds is $p(x|\theta) = \theta^x(1-\theta)^{1-x}$ where $x = 1$ if the property holds ($\sigma \models \phi$) and $x = 0$ if it does not ($\sigma \not\models \phi$). The (unknown) probability for $x = 1$ is given by θ. For inference, we need a prior probability density for θ. If no information is available a non-informative prior can be used.

To simplify computation we can use a so called conjugate prior [7]. Using a conjugate prior means intuitively that the posteriori distribution has the same form as the prior distribution. The conjugate prior for the Bernoulli distribution is the Beta-distribution $Beta(\theta|\alpha, \beta) = \frac{1}{B(\alpha,\beta)}\theta^{\alpha-1}(1-\theta)^{\beta-1}$ where $B(\alpha, \beta)$ is the Beta function. If we have iid random variables then $p(x_1, \ldots, x_n|\theta) = \Pi_{i=1}^{n} p(x_i|\theta)$ where $x_1, \ldots, x_n$ is the results of n simulations. The posteriori distribution $p(\theta|x_1, \ldots, x_n)$ is then $p(\theta|x_1, \ldots, x_n) = Beta(\theta|\alpha + x, \beta + n - x)$ where x is the number of 1:s in the samples and n is the number of samples. α and β are given by the prior Beta-distribution. If $\alpha = \beta = 1$ the Beta-distribution equals the uniform distribution, which is considered a non-informative prior.

We are not limited to only Bernoulli distributed random variables. However, to avoid computational problems and thereby automate the approach, it is extremely useful if the sampling distribution has a conjugate prior. In this paper we use Poisson distribution as a sampling distribution where the Gamma distribution is a conjugate prior, as well as the normal distribution with unknown mean and variance. Given the prior distribution of the rate of events r, $Gamma(r|\alpha, \beta)$ where $\alpha, \beta > 0$ are user defined parameters and samples $k_1, \ldots, k_n$ drawn from a Poisson distribution with unknown rate r, $Poisson(k_i|r)$, the posteriori probability distribution for the rate r is [7]:

$$p(r|k_1, \ldots, k_n) = Gamma(r|\alpha + n\bar{k}, \beta + n) \tag{4}$$

Here $\bar{k}$ is the average of $k_1, \ldots, k_n$. In the case of normally distributed data $N(y_i|\mu, \sigma^2)$, where μ is the mean and σ^2 is the variance, the marginal posteriori distribution for the mean μ is given by the Student-t distribution:

$$p(\mu|y_1, \ldots, y_n) = t_{\nu_n}(\mu|\mu_n, \sigma_n^2/\kappa_n) \tag{5}$$

with ν_n degrees of freedom [7] where the parameters are:

$$\begin{array}{lll} \mu_n & = & \frac{\kappa_0}{\kappa_0+n}\mu_0 + \frac{n}{\kappa_0+n}\bar{y} \\ \kappa_n & = & \kappa_0 + n \\ \nu_n & = & \nu_0 + n \\ \nu_n\sigma_n^2 & = & \nu_0\sigma_0^2 + (n-1)s^2 + \frac{\kappa_0 n}{\kappa_0+n}(\bar{x} - \mu_0)^2 \end{array} \tag{6}$$

The parameters $\kappa_0, \nu_0, \mu_0, \sigma_0$ are chosen by the user for the prior distribution, $\bar{y}$ is the average and s^2 is the computed variance of the samples.

5.1 Bayesian estimation algorithm

The goal is to estimate parameters θ of a chosen statistical model. The idea here is to draw samples until the desired confidence in the result has been

Table 1: The interval $[t_0, t_1]$ for the three posteriori distributions used in the paper.

Distribution	lower bound t_0	upper bound t_1
Beta	$\hat{p} + \delta \leq 1 ? \max(0, \hat{p} - \delta) : 1 - 2\delta$	$\hat{p} - \delta \geq 0 ? \min(1, \hat{p} + \delta) : 2\delta$
Student-t	$\hat{\mu} - \delta\hat{\sigma}$	$\hat{\mu} + \delta\hat{\sigma}$
Gamma	$\hat{r} - \max(\delta, \delta\sqrt{\hat{r}}) \geq 0 ?$ $\hat{r} - \max(\delta, \delta\sqrt{\hat{r}}) : 0$	$\hat{r} - \max(\delta, \delta\sqrt{\hat{r}}) \geq 0 ?$ $\hat{r} + \max(\delta, \delta\sqrt{\hat{r}}) : 2\max(\delta, \delta\sqrt{\hat{r}})$

achieved. The confidence in the estimate is defined as an interval for $[t_0, t_1]$ in which the estimated parameter θ should be with the probability of at least c. The probability γ of θ being in the interval is given by

$$\gamma = \int_{t_0}^{t_1} p(\theta|y) d\theta \tag{7}$$

where $p(\theta|y)$ is the posteriori probability of θ given the data y. This integral can be computed efficiently numerically in the case of Beta, Gamma and Student-t distribution using e.g. MATLAB. However, for more complex distribution this is typically not the case [7]. The algorithm to estimate θ, which is a straightforward extension of the algorithm in [17], is shown in Figure 3. The first step in the loop is to draw a sample by simulating the model, then the function f is used to calculate the data needed in the parameter estimation from the generated simulation trace. Finally, the probability that θ is between t_0 and t_1 is computed. If the probability is high enough the loop terminates.

The interval $[t_0, t_1]$ used for the different distributions are given in Table 1, where δ is a user defined parameter determining the width of the interval. We use the mean of the posteriori distribution as the center of the interval. For the Poisson distributed variables and for the normally distributed random variables, the width of the interval is defined as a fraction δ of the estimated standard deviation of the sampling distributions. Note that for a Poisson process, the mean equals the variance. Note also the handling of boundary cases. One could also consider a (user-defined) static lower bound for the width $t_1 - t_0$ in the case of normally distributed data, since width now approaches zero when the variance approaches zero. Other choices of interval bounds are possible, as long as the width of the interval $t_1 - t_0$ converges to a strictly positive value.

5.2 Bayesian hypothesis testing

We compare models by using Bayesian hypothesis testing with Bayes factors [13]. Here we decide between two mutually exclusive hypotesis H_0 and H_1. For a parameter θ defined on two models M_1 and M_2 we will use the

Input :
  $f(\sigma, y)$ – A function that computes a statistic on a trace and adds it to existing statistics
  $p(\theta|y, \theta_0)$ – A posteriori probability density function
  $t_0(y, \theta_0)$ – A function that computes lower bound for the interval
  $t_1(y, \theta_0)$ – A function that computes upper bound for the interval
  $c \in (0.5, 1)$ – The interval coverage coefficent
Output :
  y – The statistic y needed for the posteriori distribution
repeat
  $\sigma :=$ Draw a sample trace from the system model
  $y := f(\sigma, y)$
  $t_0 := t_0(y, \theta_0)$
  $t_1 := t_1(y, \theta_0)$
  $\gamma := \int_{t_0}^{t_1} p(\theta|y, \theta_0) d\theta$
until $\gamma \geq c$
return y

Figure 3: The algorithm for parameter estimation to a desired precision

hypotheses:

$$H_0 : \theta_1 \geq \theta_2 \qquad\qquad H_1 : \theta_1 < \theta_2 \tag{8}$$
$$H_0 : |\theta_1 - \theta_2| \leq \epsilon \qquad\qquad H_1 : |\theta_1 - \theta_2| > \epsilon \tag{9}$$

The hypotheses in (8) are used to test if θ in model M_1 is greater or less than θ in M_2. The second set of hypotheses in (9) are used to check if θ in both models differ from each other with less than a user defined value ϵ.

For normally distributed data, θ will typically be the mean μ and for Poisson distributed data θ is the rate r. Based on the data d Bayes theorem then gives the posteriori probability hypothesis H_i

$$p(H_i|d) = \frac{p(d|H_i)p(H_i)}{p(d|H_0)p(H_0)+p(d|H_1)p(H_1)} \qquad i = 0, 1 \tag{10}$$

Here the prior probabilities must be strictly positive and $p(H_1) = 1 - p(H_0)$. The posteriori odds for hypothesis H_0 is

$$\frac{p(H_0|d)}{p(H_1|d)} = \frac{p(d|H_0)p(H_0)}{p(d|H_1)p(H_1)} \tag{11}$$

The Bayes factor B is then defined as $B = p(d|H_0)/p(d|H_1)$. When the priors are fixed, the Bayes factor is used to measure the confidence in the hypothesis H_0 against H_1. Guidelines for interpreting the Bayes factor are given in [13]. A Bayes factor $B \geq 100$ can be seen as strong evidence in favour for H_0 and a value of $B \leq 0.01$ as strong evidence in favour for H_1. We adapt the algorithm from [11, 17] to dynamically chose the number of samples so that either H_0 is accepted with a fixed threshold T or H_1 is accepted with a threshold $1/T$.

We use the posteriori probabilities $p(H_0|d)$ and $p(H_1|d)$ to compute the Bayes factor. Below we show the definition of the Bayes factor for normally distributed data and the hypotheses $H_0 : \mu_1 \geq \mu_2$ and $H_1 : \mu_1 < \mu_2$. In this case, the posteriori distribution of μ_1 and μ_2 are Student-t distributions. The probability for $\mu_1 \geq \mu_2$ is obtained by integrating joint probability distribution $p(\mu_1, \mu_2)$ over the area satisfying this condition. The same applies to $\mu_1 < \mu_2$. Assuming μ_1 and μ_2 are *independent*, the posteriori odds then becomes:

$$\frac{p(H_0|d)}{p(H_1|d)} = \frac{\int_{-\infty}^{\infty} \int_{-\infty}^{\mu_1} t_{\nu_{n1}}(\mu_1|\mu_{n_1}, \sigma_{n_1}^2/\kappa_{n_1}) t_{\nu_{n2}}(\mu_2|\mu_{n_2}, \sigma_{n_2}^2/\kappa_{n_2}) d\mu_2 d\mu_1}{\int_{-\infty}^{\infty} \int_{\mu_1}^{\infty} t_{\nu_{n1}}(\mu_1|\mu_{n_1}, \sigma_{n_1}^2/\kappa_{n_1}) t_{\nu_{n2}}(\mu_2|\mu_{n_2}, \sigma_{n_2}^2/\kappa_{n_2}) d\mu_2 d\mu_1} \tag{12}$$

where μ_n is the mean estimated from the sample average, σ_n^2 the variance estimated from the sample variance, ν_n and κ_n are the degrees of freedom. For H_0 and H_1 in (9) the posteriori odds become:

$$\frac{p(H_0|d)}{p(H_1|d)} = \frac{\int_{-\infty}^{\infty} \int_{\mu_1-\epsilon}^{\mu_1+\epsilon} t_{\nu_{n1}}(\mu_1|\mu_{n_1}, \sigma_{n_1}^2/\kappa_{n_1}) t_{\nu_{n2}}(\mu_2|\mu_{n_2}, \sigma_{n_2}^2/\kappa_{n_2}) d\mu_2 d\mu_1}{\begin{array}{c}\int_{-\infty}^{\infty} \int_{-\infty}^{\mu_1-\epsilon} t_{\nu_{n1}}(\mu_1|\mu_{n_1}, \sigma_{n_1}^2/\kappa_{n_1}) t_{\nu_{n2}}(\mu_2|\mu_{n_2}, \sigma_{n_2}^2/\kappa_{n_2}) d\mu_2 d\mu_1 + \\ \int_{-\infty}^{\infty} \int_{\mu_1+\epsilon}^{\infty} t_{\nu_{n1}}(\mu_1|\mu_{n_1}, \sigma_{n_1}^2/\kappa_{n_1}) t_{\nu_{n2}}(\mu_2, \mu_{n_2}, \sigma_{n_2}^2/\kappa_{n_2}) d\mu_2 d\mu_1\end{array}} \tag{13}$$

The Bayes factor then becomes $B = \frac{p(H_1)p(H_0|d)}{p(H_0)p(H_1|d)}$. For other forms of the posteriori distributions (Gamma distribution and Beta distribution in this paper) the Bayes factor for the hypothesis can be formulated in a similar manner. Note that the Bayes factor is notorious for being difficult to compute for arbitrary distributions [13]. In our cases numerical integration techniques in MATLAB work on the tests we have done when the number of samples are sufficiently large. However, more easily computed approximations, such as the Schwarz criterion, exists [13]. Markov Chain Monte-Carlo (MCMC) techniques [7] to compute approximate the integrals can also be used, but due to the large number of samples needed to accurately approximate the integrals, they can be very slow.

The Bayesian hypothesis testing algorithm in Figure 4 modifies the one in [17] to test hypotheses concerning two models. The algorithm draws iid sample traces from the models on which the desired hypothesis is tested. Samples are drawn until either the Bayes factor is greater than a predefined threshold indicating that H_0 should be accepted or the Bayes factor is smaller than $1/T$ indicating that H_1 should be accepted.

Note that it is useful to combine the estimation algorithm with the hypothesis testing algorithm. As the statistical models are approximations, the estimated θ is needed to validate that the statistical model approximates the data sufficiently well. Additionally, when deciding between hypotheses H_0 and H_1 in (8) and $\theta_0 \approx \theta_1$ then a huge number of samples might be needed to accept one of the hypothesis even if the values of the parameters are almost

Input :
$f(\sigma, y) -$ A function that computes a statistic on a trace and adds it to existing statistics
$p(\theta|y, \theta_0) -$ A posteriori probability density function
$T \geq 1 -$ The threshold to accept H_0
Output :
$H_0 : \theta_1 \geq \theta_2$ accepted or $H_1 : \theta_1 < \theta_2$ accepted
do
$\sigma_1 :=$ Draw a sample trace from the system model 1
$\sigma_2 :=$ Draw a sample trace from the system model 2
$y_1 := f(\sigma_1, y_1)$
$y_2 := f(\sigma_2, y_2)$
$B := \mathsf{BayesFactor}(p, y_1, y_2, \theta_0)$
if $(B > T)$ **then return** H_0 accepted
if $(B < 1/T)$ **then return** H_1 accepted
end do

Figure 4: The Bayesian hypothesis testing algorithm

the same. Hence, one can also stop the iteration when the parameters θ_0 and θ_1 has been estimated to desired precision, while the hypothesis testing is still inconclusive.

5.3 Algorithm analysis

We essentially use the same estimation and hypothesis testing algorithms as in [17], but use different probability distributions. They have proved termination (with high probability) of both parameter estimation and hypothesis testing. Termination of parameter estimation is straightforward to prove and the proof in [17] is straightforward to adapt here. Termination of hypothesis testing is more difficult, see [10] for a proof in the case of Bernoulli sampling distribution with a Beta prior. In the case of hypothesis testing, termination is in our case perhaps more of a theoretical interest, due to numerical problems when deciding between hypotheses where the evidence in favour of one or the other is weak.

The upper bound $1/T$ on the probability of making type I and type II errors in hypothesis testing has also been proved [17]. By type I error we mean that we reject the H_0 hypothesis even if it is true. A type II error is the error of accepting H_0 even if it is false. The bound on the probability on the type I/II error in hypothesis testing is straightforward to derive for the Poisson distribution (same proof as in [17]) and it works with minor modifications for the normal distribution. The probability of estimation errors [17] can be analysed in terms of Type I/II error by using the hypothesis $H_0 : \theta \in [t_0, t_1]$ and $H_1 : \theta \notin [t_0, t_1]$. The hypothesis H_0 then represents the case that θ is within the desired interval. The probability of a type I or II error is bounded above by $\frac{(1-c)\pi_0}{c(1-\pi)}$ where c is the coverage coefficient and π_0 is the prior probability of H_0. Note that this applies when we sample from

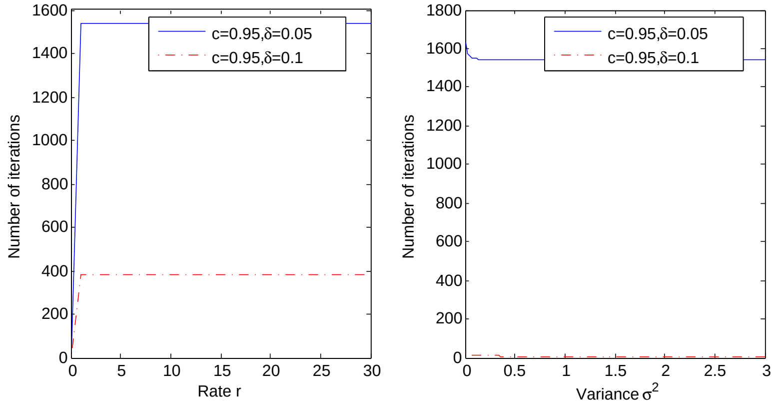

Figure 5: The number of iterations required to achieve desired precision using the parameter estimation algorithm. The number of iterations for Poisson distributed data is shown to the left and for normally distributed data to the right.

random variable with the assumed distribution. However, here we typically approximate the probability distribution of a random variable with an easy to use distribution. Hence, the error bounds provide only an idealised bound and the real bounds are unknown.

Figure 5 shows the number of iterations needed when estimating rate r for Poisson distributed data (left) and when estimating the mean (here the mean $\mu = 0$) of normally distributed data (right) using the interval bounds in Table 1. As the interval $[t_0, t_1]$ becomes smaller the number of iterations required increases rapidly. The increase in c does not have as strong impact, which is also noted in [17]. The number of iterations required for Poisson distributed data is almost independent of the rate r except for small rates, due to the fixed minimum size of the interval. In the case of normally distributed data, the number of samples is the same regardless of the variance, except when the variance becomes small when the number of samples required approaches infinity as the interval width approaches zero. Hence, proper scaling of the problem becomes essential.

6 Application

The testing methodology has been applied to the case study to evaluate 7 different properties in the tests *Test 1*, …, *Test 7* described below. The metrics in Section 4 are used. The results comparing System 1 and System 2 are summarised in Tables 2 and 3, while the results for comparing System 1 and System 3 are summarised in Tables 4 and 5. Recall the bounds on the probability of Type I and II errors and the bound on the probability of estimation errors discussed in Section 5.3.

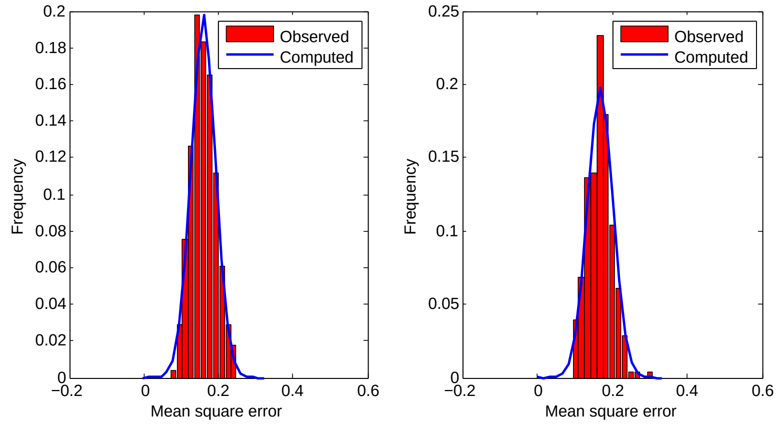

Figure 6: The estimated and observed distribution of the mean square error in System 1 (left) and System 2 (right).

Table 2: Results from parameter estimation in the five tests for System 2

Test	Post. dist.	c	δ	γ	iterations
Test 1	Student-t	0.9	0.1	0.9	274
Test 2	Beta	0.9	0.02	0.9	539
Test 3	Gamma	0.9	0.1	0.9	271
Test 4	Gamma	0.9	0.1	0.99	73
Test 5	Student-t	0.9	0.1	0.9	273
Test 6	Gamma	0.9	0.1	0.9	271
Test 7	Beta	0.9	0.01	0.9	113

6.1 Comparision of System 1 with System 2

Mean square error of pressure tracking (Test 1). When analysing pressure tracking performance, we focus on the pressure in the A-line of the system. We are interested in the difference $p_A - p_{A,ref}$, where p_A is the (continuous) pressure signal and $p_{A,ref}$ is the (discrete) reference pressure.

This test consider the mean square error (1) of the difference $p_A - p_{A,ref}$ over a time window of 5-15s. This is a performance property. The square error is scaled by a factor of 10^{-6} to avoid numerical problems. The time windows have been chosen to avoid transients at the start of the system in order to focus on steady-state behaviour.

Figure 6 shows the observed and estimated distributions of mean square errors for System 1 and System 2, respectively. As can be seen from the figures, the normal distribution is a fairly good approximation of the observed distribution. A surprising result is that the control quality does not decrease even with the delays. The hypothesis H_0 in (9) stating that the mean square errors differ at most by ϵ, where ϵ is the estimated standard deviation of System 1, is accepted with the threshold $T = 100$ in 21 iterations.

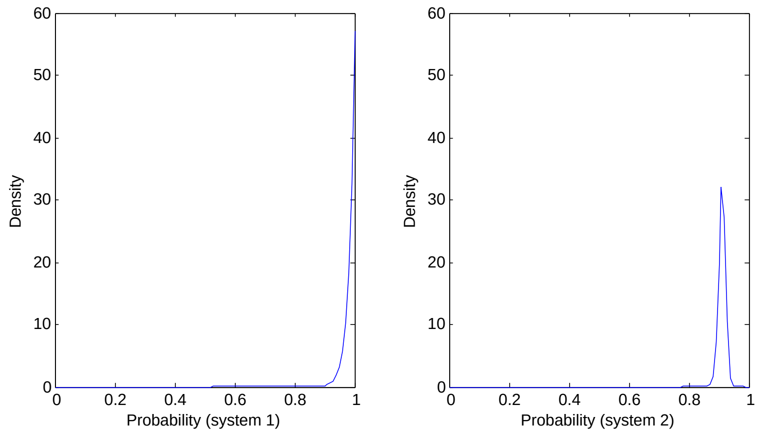

Figure 7: The estimated the probability that the that property $\mathbf{G}^{15s}(p_i(1) \leq 21MPa)$ holds in System 1 (left) and System 2 (right).

Probability of pressure peaks (Test 2). Absence of peaks is described as the BLTL property, $\mathbf{G}^{15s}(p_i(1) \leq 21MPa)$. The probability of a property violation is assumed to be a Bernoulli distributed random variable. The estimated probability for the property to hold for System 1 and System 2 is shown in Figure 7. The hypothesis that the property holds with greater probability in System 1 than in System 2 is accepted after 76 iterations.

The rate of cylinder pressure peaks (Test 3) and low pressure (Test 4). A pressure peak event is defined as a pressure rising above 21MPa and a low pressure event is defined as the pressure falling below 0.1MPa. All cylinders are analysed separately and events are counted for each type of violation (low pressure, pressure peak) separately. A Poisson distribution is assumed for the rate of events.

Figure 8 shows estimated and observed distribution of low pressure events for the first pumping cylinder in System 1 and System 2, respectively. The rates for all 6 cylinders are very similar in both cases. In System 1 the events occur rarely, while in System 2 they occur with an average rate of above 20

Table 3: Results from hypothesis testing comparing System 1 and System 2

Test	Post. dist.	T	Result (iterations)
Test 1, hypotheses (8)	Student-t	100	H1 accepted (233)
Test 1, hypotheses (9), $\epsilon = 0.033$	Student-t	100	H0 accepted (21)
Test 2, hypotheses (8)	Beta	100	H0 accepted (76)
Test 4, hypotheses (8)	Gamma	100	H1 accepted (50)
Test 5, hypotheses (8)	Student-t	100	H1 accepted (9)
Test 5, hypotheses (9), $\epsilon = 82$	Student-t	100	H0 accepted (46)
Test 7, hypotheses (8)	Beta	100	H1 accepted (4)

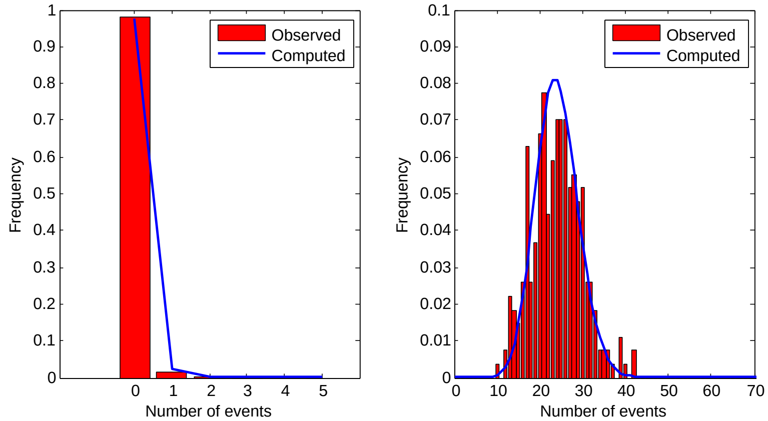

Figure 8: The rate at which low pressure events occur in System 1 (left) and in System 2 (right).

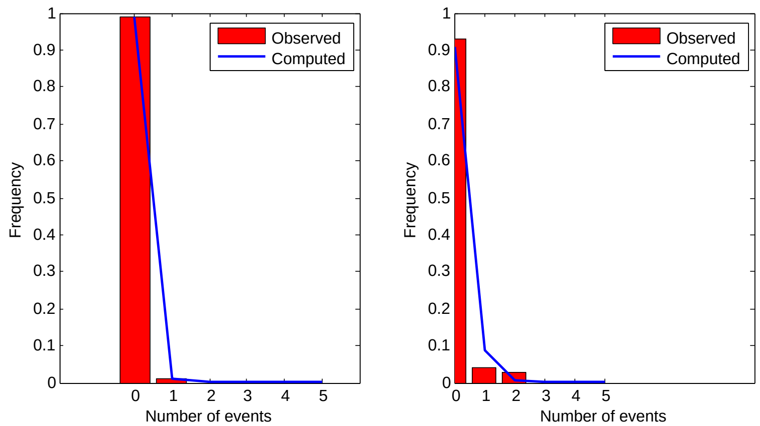

Figure 9: The rate at which high pressure events occur in System 1 (left) and in System 2 (right).

events per 10s. The hypothesis that the rate in System 2 is greater than in System 1 is accepted after 50 iterations with threshold $T = 100$. This is actually the main difference between System 1 and System 2. For example, in System 1 the probability that there are 0 cavitation events during one minute assuming Poisson distributed events and that the rate is constant over time is 0.91.

Pressure peaks occur with approximately the same rate in both System 1 and System 2. The rate of high pressure events in System 1 and System 2 are shown in Figure 9. However, the high pressure peaks in System 2 do not follow a Poisson distribution very well, which indicates that the events are not independent. Additional tests have also confirmed this result. This is actually an interesting discovery. Hence, that a probabilistic model does not fit the data is here an interesting result in itself. However, the hypothesis testing approach cannot be used to compare models in this case. A more accurate approximation, recommended in [7] when events are clustered, is a negative binomial distribution.

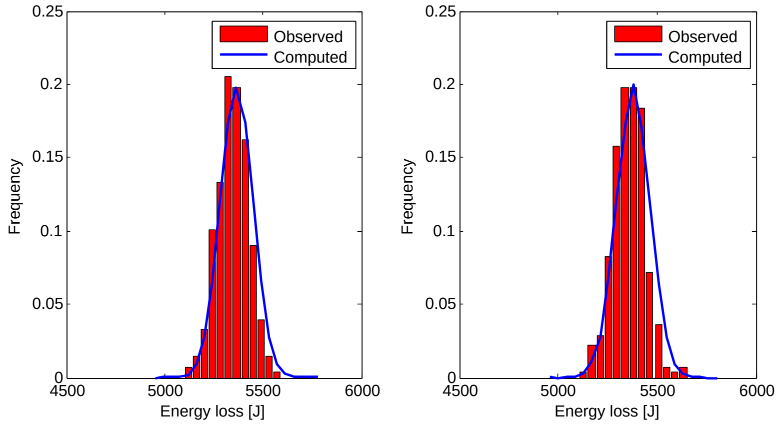

Figure 10: The estimated and observed distribution of the energy loss in System 1 (left) and System 2 (right).

Energy losses (Test 5). We can estimate the energy losses in the pumping process. Test 5 compares the energy losses in the two systems. As the total energy losses are the sum of smaller random losses, the central limit theorem motivates the assumption that the energy losses over a fixed time interval is a normally distributed random variable if the system is time invariant. The energy loss during the time interval 0-15s is similar in both systems (see Figure 10), the estimated mean in System 1 is $\hat{\mu}_1 = 5370J$ and $\hat{\mu}_2 = 5380J$ in System 2. The hypothesis that the means are within one standard deviation $\hat{\sigma} = 87$ from each other is accepted in 46 iterations.

The rate at which the valve delay is further that 15% from the expected value (Test 6). The delays in closing and opening valves is an important factor for system performance. Each valve has a pre-determined average delay. This delay can be compensated for in the control software. However, the actual delay is a random variable with unknown distribution. We analyse the rate of the delay of opening or closing of a valve will be outside the limits given by the interval $d \in [d_a - 0.15d_a, d_a + 0.15d_a]ms$ where d_a is the average delay of the valve. As we assume the length of delays in different valve events are independent from each other, we can assume the number of events in a time interval follows a Poisson distribution. The rate gives information of how often a property becomes false (in violations/10s). It does not provide information about the fraction of time it is false. This gives a more fine grained information than just a probability that a property violated in a given time.

The expected valve delay in System 2 is $d_a = 1.45ms$. In this test the rate of violations of these bounds are assessed. The unit is number of violations in 10s. The estimation gives the mode of the rate of $\hat{r} = 2.55$ for the Poisson distribution. Using a predictive distribution, we have validated that this

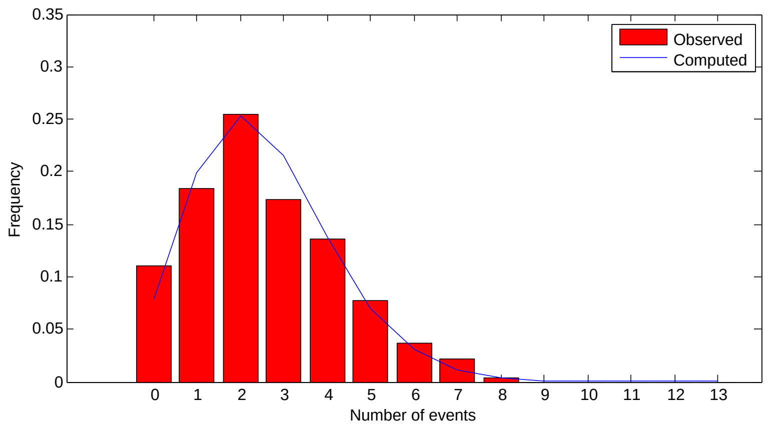

Figure 11: The rate at which the delay of valves is more then 15% from the expected in System 2.

distribution fits the data closely (see Figure 11).

Probability of cavitation (Test 7). Cavitation is defined as the pressure $p_i = 0$ in a cylinder, in BLTL the property to check is $\mathbf{F}^{15s} p_i = 0$. In System 1 cavitation is never observed in the time interval 5-15s while in System 2, it is observed every time. The probability distribution of cavitation in System 1 and System 2 is shown in Figure 12. The hypothesis that the probability is smaller in System 1 than System 2 is accepted after only four iterations.

Conclusions. The control performance given by the mean square error and the energy consumption are not significantly negatively affected by the additional disturbances in System 2. The only significant effect is the increase of low pressure events. However, when testing several hypotheses, the probability to have at least one type I/II error is $p(n_{error} \geq 1) = 1 - p(n_{error} = 0)$. If hypothesis i has threshold T_i and the errors are independent, then this probability is bounded above, $p(n_{error} \geq 1) \leq 1 - \Pi(1 - \frac{1}{T_i})$. Hence, the probability of type I and II errors can be significant, if we have many hypotheses and a relatively small T.

6.2 Comparision of System 1 with System 3

We are also interested how much slower (and thereby cheaper) valves than in System 2 can be used. Also a slower sampling time allows cheaper processors

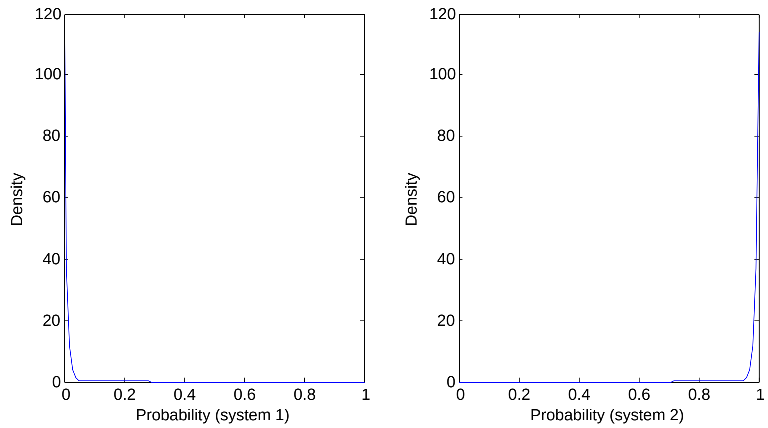

Figure 12: The estimated probability density that cavitation occurs in System 1 (left) and System 2 (right).

to be used or it allows running also other applications on the same platform. Here we use a model that has a slower sampling time, $T_s = 100\mu s$. The valves are also slower with a delay of $4ms$. Computation times and valve delays are defined in the same way as in System 2. We use the same tests *Test 1*,..., *Test 7* as in the comparision with System 2. The results are summarised in Tables 4 and 5.

Mean square error of pressure tracking (Test 1). This test consider the mean square error (1) of the pressure difference $(p_{A,ref} - p_A)$ over a time window of 5-15s. The square error is scaled by a factor of 10^{-6} to avoid numerical problems.

Figure 13 shows the observed and estimated distributions of mean square errors for System 1 and System 3 respectively. The hypotesis that the mean of mean square errors is greater in System 3 than in System 1 is accepted in 86 iterations. The hypothesis that the difference between the means is greater than one standard deviation is also accepted in 162 iterations.

Probability of pressure peaks (Test 2). The probability of pressure peaks in the pumping cylinders is estimated in this test. Absence of peaks is described as the BLTL property, $\mathbf{G}^{15s}(p_i \leq 21MPa)$. The estimated probabilites for the property to hold in System 1 and System 3 are shown in Figure 14. The hypothesis that the property holds with greater probability in System 1 than in System 3 is accepted after 4 iterations.

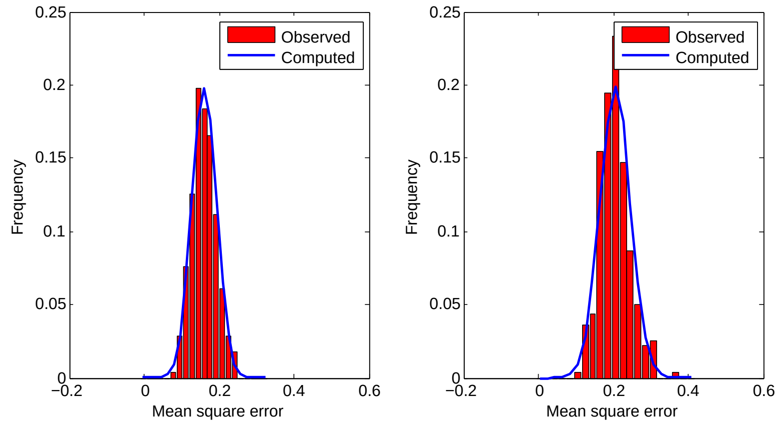

Figure 13: The estimated and observed distribution of the mean square error in System 1 (left) and System 3 (right) .

Table 4: Results from parameter estimation in the five tests for System 3

Test	Post. dist.	c	δ	γ	iterations
Test 1	Student-t	0.9	0.1	0.9	273
Test 2	Beta	0.9	0.01	0.9	339
Test 3	Gamma	0.9	0.1	0.9	271
Test 4	Gamma	0.9	0.1	0.9	271
Test 5	Student-t	0.9	0.1	0.9	273
Test 7	Beta	0.9	0.01	0.9	113

Table 5: Results from hypothesis testing comparing System 1 and System 3

Test	Post. dist.	T	Result (iterations)
Test 1, hypotheses (8)	Student-t	100	H1 accepted (86)
Test 1, hypotheses (9), $\epsilon = 0.033$	Student-t	100	H1 accepted (162)
Test 2, hypotheses (8)	Beta	100	H0 accepted (4)
Test 5, hypotheses (8),	Student-t	100	H1 accepted (20)
Test 5, hypotheses (9), $\epsilon = 82$	Student-t	100	H1 accepted (10)
Test 7, hypotheses (8)	Beta	100	H1 accepted (4)

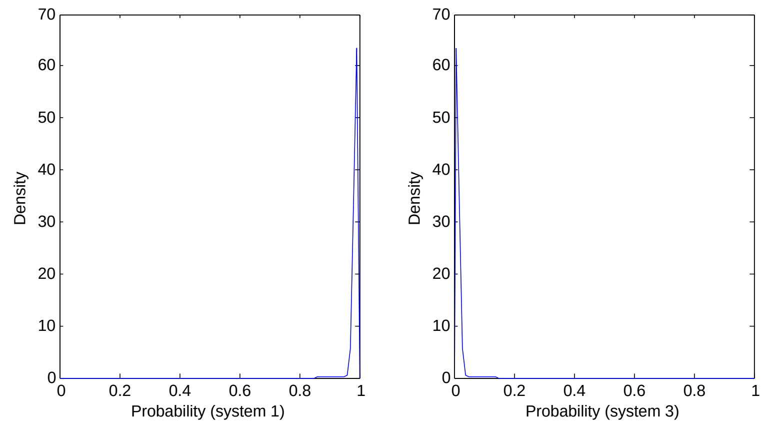

Figure 14: The estimated the probability that the that property $\mathbf{G}^{15s}(p_i(1) \leq 21MPa)$ holds in System 1 (left) and System 3 (right).

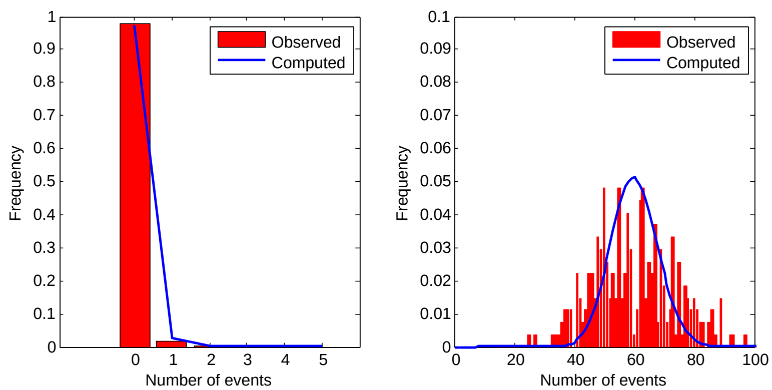

Figure 15: The rate at which low pressure events occur in System 1 (left) and in System 3 (right).

The rate of cylinder pressure peaks (Test 3) and low pressure (Test 4). A pressure peak event is defined as a pressure rising above 21MPa and a low pressure event is defined as the pressure falling below 0.1MPa. All cylinders are analysed separately and events are counted for each type of violation (low pressure, pressure peak) separately.

Figure 15 shows estimated and observed distribution of low pressure events for the first pumping cylinder in System 1 and System 3, respectively. The rates for all 6 cylinders are very similar in both cases. In System 1 the events occur rarely, while in System 3 they occur regularly. Note that the low pressure events do not follow a Poisson distribution in System 3. This is probably due to the impact of the reference pressure becomes visible. The rate of high pressure events are shown in Figure 16. Just as in System 2, the high pressure peaks do not follow a Poisson distribution, which indicates that the events are not independent.

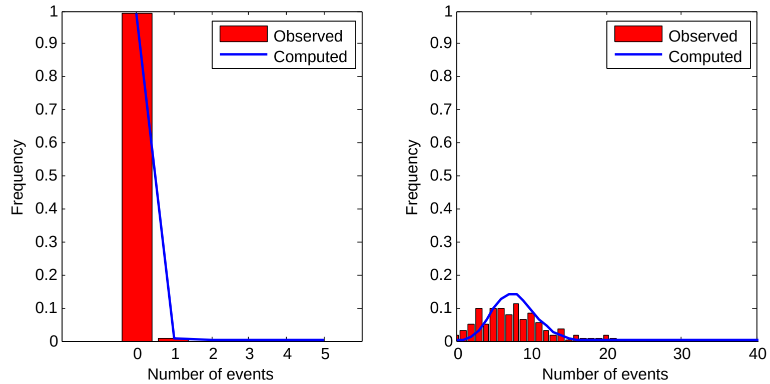

Figure 16: The rate at which high pressure events occur in System 1 (left) and in System 3 (right).

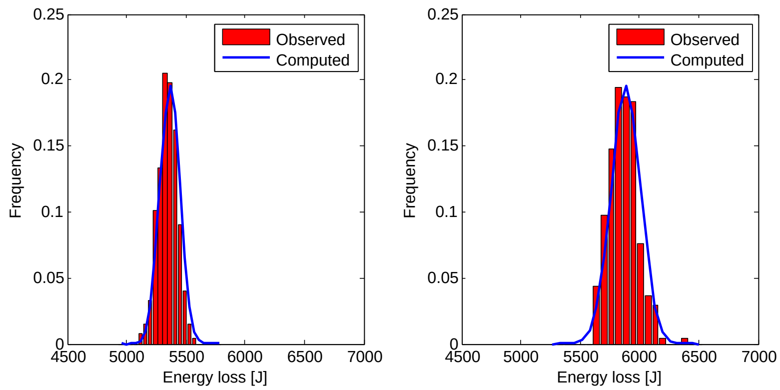

Figure 17: The estimated and observed distribution of the energy loss in System 1 (left) and System 3 (right).

Energy losses (Test 5). The energy losses during the time interval 0-15s are similar in both systems as shown in Figure 10, the estimated mean in System 1 is $\hat{\mu}_1 = 5370J$ and $\hat{\mu}_3 = 5890J$ in System 3, indicating a 10% increase in energy consumption in System 3 compared to System 1. The hypothesis that the means are more than one standard deviation $\hat{\sigma} = 10$ from each other is accepted in 50 iterations.

The rate at which the valve delay is further that 15% from the expected value (Test 6). This test was skipped for System 3, since the computation time becomes very long and it is not expected to reveal anything interesting not already shown in the corresponding test for System 2. The reason is that these delays depend directly on the computation times and valve delays and there is no feedback involved.

Probability of cavitation (Test 7). In this test the probability of cavitation in the cylinders is assessed in the two systems. The property in BLTL is, $\mathbf{F}^{15s} p_i = 0$. In System 1 cavitation is never observed in the time interval 5-15s while in System 3, it is observed every time. The hypothesis that the probability is smaller in System 1 than System 3 is accepted after only four iterations.

Conclusions. The control performance does not deteriorate very much due to the large delays indicated by the relatively small increase in the mean square error. The system also becomes less energy efficient, but only by an estimated 10%. The main problem is that pressure peaks and cavitation occur much more often, and both situations are problematic. Overall, the problems already partly visible in System 2 becomes worse here due to longer delays with more variation. However, this evaluation pinpoints the problems and provides means to evaluate if potential new solutions improve the situation.

7 Related work

Statistical model checking has been an active area of research. The focus has been on checking various kinds of (bounded) temporal logic formula on stochastic models. Here we directly extend the methodology in [17] to more general properties than can be expressed by temporal logic formulas using well-known Bayesian techniques. Younes and Simmons [16] have used hypothesis testing to determine if the probability that a temporal logic property holds is above a desired limit. They use the *sequential probability ratio test* (SPRT) to adaptively sequentially sample only as many times as needed. Additionally, they can also bound the probability of Type I and II errors. However, they perform no parameter estimation, which is here important for model validation. For parameter estimation, techniques based on the Chernoff-Hoeffding bound have been used by Hérault et.al. [9]. According to experiments by Zuliani et. al. [17], this estimation approach can be significantly slower than the Bayesian approach used here. Both hypothesis testing using SPRT and parameter estimation using the Chernoff-Hoeffding bound has been implemented in UPPAAL-SMC [6, 3]. Additionally they implement comparisons of probabilities from different models by hypothesis testing based on SPRT. David et.al. [5] have also used ANOVA for model comparison. The goal is similar to how we compare parameters in different models using Bayesian hypothesis testing. However, an in-depth comparison is further work.

Jitterbug [4] can be used to compute the mean square error (1) numerically for a special kind of systems. However, in our case the systems do not fulfil the requirements (the system dynamics is not linear). We have also analysed worst case timing for a version of the case study system [2] by applying model-checking of timed automata in the TIMES tool [1]. That paper focuses on the improvements obtained by using interrupt driven tasks instead of periodic tasks. Here the focus is on the impact of random delays on system performance.

8 Conclusions

This paper presents an approach to statistically analyse and compare stochastic models, based on Bayesian parameter estimation and hypothesis testing. We demonstrated the approach to compare two versions of a model of a digital hydraulics power management system. The differences in the models concern stochastic delays and disturbances. The results from the hypothesis testing show that we can gain confidence in the result with relatively few samples. The parameter estimation was used to validate that data fits the used probability model. However, comprehensive model validation [7] is outside the scope of the paper.

In [17], they give bounds on probability of Type I and Type II errors in hypothesis testing and bounds on the error of estimated probabilities. Although the same results apply here, the Poisson and Normal distributions are approximations of the real processes and, hence, model validation becomes an essential step in order to draw valid conclusions. Therefore, parameter estimation is an important complement to hypothesis testing. To automate the analysis there need to be efficient ways to accurately compute the needed integrals. The need for certain types of statistical models then also limits the approach to certain properties, where the models are expected to provide a good fit or where a poor fit provides some insight into the behaviour of the process.

The approach in the paper is not limited to analysing impact of random delays, but can be used for other purposes as well. The key advantage is to generate only as many samples as needed to gain a desired confidence. Additionally, prior information can be included by using the prior distributions. Future work include using more complex statistical models. Incremental model validation would also be useful to ensure that the checking process returns an accurate result.

Acknowledgment. The work has been partially funded by the EDiHy project (no. 139540 and no. 140003) funded by the Academy of Finland. We also like to thanks Anton Tarasyuk for discussions on probability.

References

[1] Amnell, T., Fersman, E., Mokrushin, L., Pettersson, P., Yi, W.: TIMES: a tool for schedulability analysis and code generation of real-time systems. In: Larsen, K.G., Niebert, P. (eds.) FORMATS. LNCS, vol. 2791. Springer, Heidelberg (2003)

[2] Boström, P., Alexeev, P., Heikkilä, M., Huova, M., Waldén, M., Linjama, M.: Analysis of real-time properties of a digital hydraulic power management system. In: Lang, F., Flammini, F. (eds.) FMICS. LNCS, vol. 8718. Springer (2014)

[3] Bulychev, P., David, A., Larsen, K.G., Legay, A., Mikučionis, M., Poulsen, D.B., van Vliet, J., Wang, Z.: Statistical model checking for networks of priced timed automata. In: Fahrenberg, U., Tripakis, S. (eds.) FORMATS. LNCS, vol. 6919. Springer, Heidelberg (2011)

[4] Cervin, A., Henriksson, D., Lincoln, B., Eker, J., Årzén, K.E.: How does control timing affect performance? - analysis and simulation of timing using Jitterbug and TrueTime. IEEE Control Systems Magazine 23(3) (2003)

[5] David, A., Du, D., Larsen, K.G., Legay, A., Mikučionis, M.: Optimizing control strategy using statistical model checking. In: Brat, G., Rungta, N., Venet, A. (eds.) NFM. LNCS, vol. 7871. Springer, Heidelberg (2013)

[6] David, A., Larsen, K.G., Legay, A., Mikučionis, M., Wang, Z.: Time for statistical model checking of real-time systems. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV. LNCS, vol. 6806. Springer, Heidelberg (2011)

[7] Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B.: Bayesian Data Analysis. Chapman & Hall/CRC, 2nd edn. (2004)

[8] Heikkilä, M., Tammisto, J., Huova, M., Huhtala, K., Linjama, M.: Experimental evaluation of a piston-type digital pump-motor-transformer with two independent outlets. In: Bath/ASME Symposium on Fluid Power and Motion Control (2010)

[9] Hérault, T., Lassaigne, R., Magniette, F., Peyronnet, S.: Approximate probabilistic model checking. In: B. Steffen, G.L. (ed.) VMCAI. LNCS, vol. 2937. Springer (2004)

[10] Jha, S.K., Clarke, E.M., Langmead, C.J., Legay, A., Platzer, A.: Statistical model checking for complex stochastic models in systems biology. Tech. Rep. CMU-CS-09-110, School of Computer Science, Carnegie Mellon University (2009)

[11] Jha, S.K., Clarke, E.M., Langmead, C.J., Legay, A., Platzer, A., Zuliani, P.: A bayesian approach to model checking biological systems. In: Degano, P., Gorrieri, R. (eds.) CMSB. LNCS, vol. 5688. Springer (2009)

[12] Karvonen, M., Heikkilä, M., Huova, M., Linjama, M.: Analysis by simulation of different control algorithms of a digital hydraulic two-actuator system. International Journal of Fluid Power 15(1), 33–44 (Mar 2014)

[13] Kass, R.E., Raftery, A.E.: Bayes factors. Journal of American Statistical Association 90(430), 773–795 (1995)

[14] Lincoln, B., Cervin, A.: Jitterbug: A tool for analysis of real-time control performance. In: 41st IEEE Conference on Decision and Control. IEEE (2002)

[15] Linjama, M., Huhtala, K.: Digital pump-motor with independent outlets. In: The 11th Scandinavian International Conference on Fluid Power, SICFP'09. Linköping, Sweden (2009)

[16] Younes, H.L.S., Simmons, R.G.: Statistical probabilistic model checking with focus on time-bounded properties. Information and Computation 204 (2006)

[17] Zuliani, P., Platzer, A., Clarke, E.M.: Bayesian statistical model checking with application to Stateflow/Simulink verification. Formal methods in system design 43 (2013)

TURKU
CENTRE *for*
COMPUTER
SCIENCE

Joukahaisenkatu 3-5 A, 20520 TURKU, Finland | www.tucs.fi

University of Turku
Faculty of Mathematics and Natural Sciences
- Department of Information Technology
- Department of Mathematics

Turku School of Economics
- Institute of Information Systems Sciences

Åbo Akademi University
- Department of Computer Science
- Institute for Advanced Management Systems Research

ISBN 978-952-12-3210-7
ISSN 1239-1891