Robust stochastic control based on imprecise probabilities ⋆

A. Benavoli [*], **L. Chisci** [**]

[*] *Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA), Lugano, Switzerland, alessio@idsia.ch.*
[**] *Dipartimento di Sistemi e Informatica (DSI), Università di Firenze, Firenze, Italy, chisci@dsi.unifi.it.*

Abstract: This paper deals with the optimal quadratic control problem for non Gaussian discrete-time linear stochastic systems from the perspective of *imprecise probabilities*. The adopted philosophy is to use a convex set of probability distributions to characterize the imprecision in the knowledge about the probabilistic relationships present in the system to be controlled. In particular, an uncertain system model, named *Linear Gaussian Vacuous Mixture (LGVM)*, in which disturbances and initial state uncertainty are described as convex combinations (mixtures) of nominal Gaussian distributions and unknown *vacuous* distributions, is adopted. A novel control approach is then derived, according to a worst-case paradigm, by minimizing the upper expectation of a finite-horizon quadratic cost functional with respect to all admissible probability distributions and exploiting a receding horizon strategy. Simulation experiments demonstrate its robustness in presence of large unexpected impulsive disturbances.

Keywords: Stochastic control; LQG; receding horizon control imprecise probability, Linear Gaussian Vacuous Mixture.

1. INTRODUCTION

Linear Quadratic Gaussian (LQG) control represents a well-established and effective optimal control design methodology. It assumes a linear system subject to additive white Gaussian disturbances under incomplete state information and aims at minimizing the expectation of a quadratic cost suitably weighting both the control effort and the state deviation from the desired trajectory. The key factor for the success of the LQG controller is that it enjoys the separation property, i.e. it consists of independently designed optimal state estimator (Kalman filter) and state feedback regulator (LQ controller). Unfortunately, however, the optimality of the LQG controller in the linear Gaussian context does not necessarily guarantee robustness with respect to real (non Gaussian) disturbances and/or modelling uncertainties. For this reason, several variants of the LQG problem like e.g. the risk-sensitive criterion (Whittle, 1981) and minimax LQG control (Petersen et al., 2000; Petersen, 2006) have been adopted. This paper tackles robust LQG control from the different perspective of *imprecise probabilities* (Walley, 1991), whose philosophy is to replace a single probability distribution with a convex set of probability distributions to take into account imprecision in the knowledge about the probabilistic relationships present in the system to be controlled. In (Walley, 1991), it is proved that a convex set of probability distributions can be equivalently characterized by the upper (or lower) expectation functional that it generates as the upper (lower) envelope of the expectations obtained from the distributions in such a set. Hence, the imprecision in the system model can equivalently be expressed in terms of lower/upper expectations. In this respect, recent work (Benavoli et al., 2010) derived the *imprecise probability* Kalman filter which essentially consists of propagating in time both the lower and upper state expectations over the set of assumed probability distributions. The objective of this paper is to extend the work in (Benavoli et al., 2010) to LQG control. In particular, an uncertain system model, named *Linear Gaussian Vacuous Mixture (LGVM)*, in which disturbances and initial state uncertainty are described as convex combinations (mixtures) of nominal Gaussian distributions and unknown *vacuous* distributions, is adopted. Further, the control problem is formulated, according to a worst-case paradigm, by minimizing the upper expectation of a finite-horizon quadratic cost functional with respect to all admissible probability distributions and exploiting a receding horizon strategy. The resulting controller, named *LQGVM*, is derived and simulation experiments demonstrate its robustness in presence of large unexpected impulsive disturbances. A worst-case (minimax) approach is also adopted in (Petersen et al., 2000; Petersen, 2006). The main difference with respect to the present paper is that, here, a larger set of distributions is used to model the imprecision in the system model. In (Petersen et al., 2000; Petersen, 2006), to describe the imprecision, the authors consider all distributions whose distance (measured in terms of the Kullback Leibler divergence) from a nominal distribution is less than a given constant. Conversely, in the present paper, the distributions of the disturbances are contaminations of a nominal Gaussian with arbitrary distributions whose distance from the nominal Gaussian

⋆ This work has been partially supported by the Swiss NSF grants n. 200020-121785/1 and by the "Hydronet" project.

can thus be infinite. Hence, the uncertainty considered here is more general.

The rest of the paper is organized as follows. Section 2 introduces the *linear Gaussian vacuous mixture (LQGVM)* model and the control problem statement. Section 3 extends LQG control to contaminating disturbances by using Walley's theory of coherent upper expectations. A new control strategy called LQGVM is then derived. In Section 4, numerical results showing the performance improvement of LQGVM over LQG in presence of large unexpected impulsive disturbances are presented. Finally Section 5 ends the paper.

2. LINEAR GAUSSIAN VACUOUS MIXTURE MODEL AND PROBLEM STATEMENT

Consider the stochastic discrete-time linear time-invariant dynamical system

$$\begin{cases} x_{t+1} = Ax_t + Bu_t + w_t, \\ y_t = Cx_t + v_t, \end{cases} \tag{1}$$

where: t is the time; $x_t \in \mathbb{R}^n$ is the state; $u_t \in \mathbb{R}^m$ is the control input; $y_t \in \mathbb{R}^p$ is the measured output; $w_t \in \mathbb{R}^n$ is a process disturbance; $v_t \in \mathbb{R}^p$ is a measurement disturbance; A, B and C are matrices of compatible dimensions. Furthermore, the initial state and disturbance signals are assumed to be modelled by:

$$\begin{aligned} w_t &= \epsilon_w w_t^1 + (1-\epsilon_w) w_t^2, \\ v_t &= \epsilon_v v_t^1 + (1-\epsilon_v) v_t^2, \\ x_0 &= \epsilon_x x_0^1 + (1-\epsilon_x) x_0^2, \end{aligned} \tag{2}$$

where: the scalars ϵ_w, ϵ_v and ϵ_x belong to $[0,1]$; w_t^1, v_t^1, x_0^1 are Gaussian, i.e.

$$w_t^1 \sim \mathcal{N}(0,Q),\ v_t^1 \sim \mathcal{N}(0,R),\ x_0 \sim \mathcal{N}(\hat{x}_0, \Pi_0); \tag{3}$$

$\mathcal{N}(\mu,\Sigma)$ denotes a Gaussian random variable with mean μ and variance Σ; the contaminating terms w_t^2, v_t^2, x_0^2 are assumed to have unknown arbitrary distributions. Note that the model which characterises w_t, v_t and x_0 is a special case of the so-called ϵ-*contamination* model (Berger, 1985). Varying the distributions of the *contaminating disturbances* w_t^2, v_t^2 and x_0^2, this model produces a family of distributions, i.e., the family of all convex combinations between a known nominal (Gaussian) distribution and an arbitrary distribution. This family can be used to address the fact that the noise model and a-priori state information (nominal distributions) are inexact and, thus, are contaminated to reflect their modelling uncertainty. The parameter ϵ represents, therefore, the degree of belief of the modeller that the system is Gaussian and $1-\epsilon$ the degree of imprecision in the knowledge of the involved probabilistic relationships.

It will be assumed throughout the paper that the triplet (A,B,C) is stabilizable and detectable. The objective is to design, via receding-horizon control, a nonlinear dynamic output feedback

$$u_t = h(\hat{x}_0, y_0, \dots, y_t), \tag{4}$$

so as to provide robust performance in presence of the contaminating disturbances.

The next section will show how to design such a controller via a model predictive (receding-horizon) control approach based on *Coherent Lower/Upper Previsions (CLP/CUP).*

2.1 Stochastic Receding Horizon Control

Given the sequence $y^t = \{y_1, \dots, y_t\}$ of measurements and the sequence of control inputs $u^{\tau-1} = \{u_0, \dots, u_{\tau-1}\}$, with $\tau \geq t$, one can estimate any sequence of states $x^{t:\tau} = \{x_t, \dots, x_\tau\}$.[1] In the Bayesian framework, all relevant information on $x^{t:\tau}$ is represented by the posterior distribution $p(x^{t:\tau}|y^t, u^{\tau-1})$. Recall that the Markov structure of the model (1) implies the following independence conditions:

$$p(x_k|x^{k-1}, u^{k-1}) = p(x_k|x_{k-1}, u_{k-1}), \tag{5}$$

$$p(y^k|x^k, u^{k-1}) = \prod_{j=1}^{k} p(y_j|x_j). \tag{6}$$

Hence, it follows that

$$\begin{aligned} p(x^{t:\tau}|y^t, u^{\tau-1}) = \frac{p(x_t|y^{t-1}, u^{t-1})}{p(y_t|y^{t-1}, u^{t-1})}\, p(y_t|x_t) \\ \prod_{k=t+1}^{\tau} p(x_k|x_{k-1}, u_{k-1}). \end{aligned} \tag{7}$$

Once $p(x^{t:\tau}|y^t, u^{\tau-1})$ has been computed, it is possible to calculate the expected value

$$E_{X^{t:\tau}}[J(X^{t:\tau}, u^{t:\tau-1})|y^t, u^{\tau-1}]$$

for any cost function of interest $J(X^{t:\tau}, u^{t:\tau-1})$.[2]

Consider the following scalar function

$$\begin{aligned} J(x^{t:t+N}, u^{t:t+N-1}) := \\ \frac{1}{N}\Big[\sum_{k=t}^{t+N-1} \Big(||x_k||^2_{\Psi_x} + ||u_k||^2_{\Psi_u} \Big) + ||x_{t+N}||^2_{\Psi_f} \Big], \end{aligned} \tag{8}$$

where: $||x_k||^2_{\Psi_x} := x_k^T \Psi_x x_k$; $N \in \mathbb{N} := \{1, 2, \dots\}$; $\Psi_x = \Psi_x^T \geq 0$, $\Psi_f = \Psi_f^T \geq 0$ and $\Psi_u = \Psi_u^T \geq 0$ are given weight matrices. This quadratic cost function is commonly used to quantify the performance of a predictive controller whose aim is to steer the state x_k close to the origin. Fixed the control horizon $N \in \mathbb{N}$, i.e. $\tau = t+N$, and the past control sequence u^{t-1}, in stochastic Model Predictive Control (MPC), the objective is to find a control sequence $u^{t:t+N-1}$ that minimizes the expected value of the cost function $J(x^{t:t+N}, u^{t:t+N-1})$, i.e.

$$\min_{\substack{u_k \in \mathcal{U} \\ k=t,\dots,t+N-1}} E_{X^{t:t+N}} \left\{ J(X^{t:t+N}, u^{t:t+N-1}) \middle| y^t, u^{t+N-1} \right\}, \tag{9}$$

subject to the dynamics (1).

Let $\hat{u}^{t:t+N-1}$ be the optimal input sequence. According to the *receding horizon* strategy, only the first input $\hat{u}_t$ is applied to the system, i.e. $u_t = \hat{u}_t$ is set as the input to (1) at time t. Thus, at the next time instant $t+1$,

[1] In the following, we use the shorthand notation y^t for $y^{1:t}$, x^t for $x^{0:t}$ and u^t for $u^{0:t}$. We specify the initial time of the sequence only when it differs from 0 (for state and control inputs) or 1 (for the measurements) as for instance in $x^{t:\tau}$.

[2] In $E_{X^{t:\tau}}[J(X^{t:\tau}, u^{t:\tau-1})|y^t, u^{\tau-1}]$, we have used the capital letter $X^{t:\tau}$ in the cost $J(X^{t:\tau}, u^{t:\tau-1})$ to denote the fact that the expectation is w.r.t. the variable $X^{t:\tau}$, while the values of the measurements and control inputs are fixed. This notation is adopted for any expectation w.r.t. to generic functionals $g(\cdot)$, e.g., $E_X[g(X,u)|y,u]$ or $E_{X,Y}[g(X,Y,u)|u]$.

the optimization problem (9) is solved again and so on recursively. In absence of contaminating disturbances, i.e. $\epsilon_w = \epsilon_v = \epsilon_x = 1$, under the assumptions (5)-(6) it is well-known that the optimal receding horizon controller takes the form: $\hat{u}_t = -F_t\hat{x}_t$, where F_t can be obtained by the following backward recursion:

$$F_k = [B'S_{k+1}B + \Psi_u]^{-1}B'S_{k+1}A, \tag{10}$$

$$S_k = A'S_{k+1}[A - BF_k] + \Psi_x, \tag{11}$$

for $k = t+N, \ldots, t$ with $S_{t+N} = \Psi_f$ and $\hat{x}_t$ is the Kalman filter state estimate which can be obtained by the following forward recursion:

$$\begin{aligned} \hat{x}_k &= A\hat{x}_{k-1} + L_k[y_k - CA\hat{x}_{k-1}] \\ L_k &= P_kC'[CP_kC' + R]^{-1} \\ P_{k+1} &= A[I - L_kC]P_kA' + Q \end{aligned}$$

for $k = 0, \ldots, t$ and $P_0 = \Pi_0$. Furthermore, when Ψ_f is equal to the steady-state solution of the Riccati equation (11), the control problem reduces to the infinite-horizon *Linear Quadratic Gaussian (LQG)* control.

3. GENERALISATION TO CONTAMINATING DISTURBANCES WITH UNKNOWN DISTRIBUTION

Consider now the plant (1) in the case of presence of contaminating disturbances. For the contaminating disturbances, it has been assumed that w_t^2, v_t^2 and x_0^2 have an unknown probability distribution (which can also be time-varying). A way to probabilistically describe such disturbances is thus to consider all possible distributions they can assume at the various time instants. In other words, the contaminating disturbances can be modelled by considering sets of distributions. In the sequel, $\mathcal{P}_w$, $\mathcal{P}_v$, $\mathcal{P}_{x_0}$ will denote the sets of distributions for w_t^2, v_t^2 and, respectively, x_0^2. Given, for instance, the set of distributions $\mathcal{P}_w$ and any bounded utility function $g : \mathcal{W} \to \mathbb{R}$, one can define its lower and upper expectation w.r.t. $\mathcal{P}_w$ by

$$\begin{aligned} \underline{E}_W(g) &= \inf_{p_w \in \mathcal{P}_w} \int_W g(w)p_w(w)dw, \\ \overline{E}_W(g) &= \sup_{p_w \in \mathcal{P}_w} \int_W g(w)p_w(w)dw, \end{aligned} \tag{12}$$

where $g(\cdot)$ is assumed to be continuous almost everywhere in $\mathcal{W}$. In other words, the set $\mathcal{P}_w$ can then be characterized by the upper and lower expectations, $\overline{E}_W(g)$ and $\underline{E}_W(g)$, generated as the supremum and infimum of

$$E_W(g) = \int_W g(w)p_w(w)dw,$$

over the probability measures in $\mathcal{P}_w$. From (12), it can be verified that $\underline{E}_W(g) = -\overline{E}_W(-g)$ and, thus, $\overline{E}_W$ fully describes $\mathcal{P}_w$. It can also be verified that $\overline{E}_W$ satisfies the following properties:

(C1) $\overline{E}_W(g) \leq \sup_w g$,
(C2) $\overline{E}_W(\lambda g) = \lambda\overline{E}_W(g)$,
(C3) $\overline{E}_W(f+g) \leq \overline{E}_W(f) + \overline{E}_W(g)$,

for any $\lambda > 0$ and bounded functions $f(\cdot), g(\cdot)$.

Conversely, given a functional $\overline{E}_W(g)$ that satisfies (C1)–(C3), it is possible to define a family $\mathcal{P}_w$ of probability distributions that generates the upper expectation $\overline{E}_W(g)$, for any $g(\cdot)$. This is proved in (Walley, 1991) and establishes a one-to-one correspondence between convex sets of probability distributions and upper expectations satisfying (C1)–(C3). A functional $\overline{E}_W(g)$ that satisfies (C1)–(C3) is called Coherent Upper Prevision (CUP) (Walley, 1991). Notice that, if the noise is modelled by a single distribution (i.e., $\mathcal{P}_w$ includes only a single distribution), then $\overline{E}_W[g] = \underline{E}_W[g] = E_W[g]$. This brings back to the standard expectation case. Conversely, if the distribution of the noise belongs to a set $\mathcal{P}_w$ (including more than one distribution), in general it holds that $\overline{E}_W[g] > \underline{E}_W[g]$ and, thus, standard expectation cannot be used but $\mathcal{P}_w$ can be equivalently characterized through its corresponding CUP.

For instance, in the contamination case described in (2), it results that (Benavoli et al., 2010):

$$\begin{aligned} \overline{E}_{W_k}[g] &= \sup_{p(w_k^2) \in \mathcal{P}_W} \Big[\epsilon_w \int_{w_k^1} g(w_k^1)\mathcal{N}(w_k^1; 0, Q)\ dw_k^1 \\ &\quad + (1-\epsilon_w) \int_{w_k^2} g(w_k^2)p(w_k^2)\ dw_k^2\Big] \\ &= \epsilon_w \int_{w_k} g(w_k)\ \mathcal{N}(w_k; 0, Q)\ dw_k \\ &\quad + (1-\epsilon_w) \sup_{w_k} g(w_k), \end{aligned} \tag{13}$$

Notice, in fact, that, since $\mathcal{P}_W$ includes all the possible distributions, the upper expectation of $g(\cdot)$ is obtained in correspondence of a delta function centred at the supremum of $g(\cdot)$. Since the contamination on w_k induces a contamination on x_{k+1}, i.e.,

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k + \epsilon_w w_k^1 + (1-\epsilon_w)w_k^2 \\ &= \epsilon_w(Ax_k + Bu_k + w_k^1) + (1-\epsilon_w)(Ax_k + Bu_k + w_k^2), \end{aligned}$$

it also follows that (Benavoli et al., 2010):

$$\begin{aligned} &\overline{E}_{X_k}(g|x_{k-1}, u_{k-1}) \\ &= \epsilon_w \int_{x_k} g(x_k)\mathcal{N}(x_k; Ax_{k-1} + Bu_{k-1}, Q)dx_k \\ &+ (1-\epsilon_w) \sup_{x_k} g(x_k). \end{aligned} \tag{14}$$

The CUP (14) is known as Linear Gaussian Vacuous Mixture (LGVM) model (Walley, 1991, Sec 2.9.2). Similarly for x_0 and v_k, one obtains:

$$\overline{E}_{X_0}(g) = \epsilon_x \int_{x_0} g(x_0)\ \mathcal{N}(x_0; \hat{x}_0, P_0)\ dx_0 + (1-\epsilon_x) \sup_{x_0} g(x_0), \tag{15}$$

$$\begin{aligned} \overline{E}_{Y_k}(g|x_k) &= \epsilon_v \int_{y_k} g(y_k)\ \mathcal{N}(y_k; Cx_k, R)\ dy_k \\ &\quad + (1-\epsilon_v) \sup_{y_k} g(y_k). \end{aligned} \tag{16}$$

Notice that the supremums in (14)–(16) make sense only if the functions $g(\cdot)$ are bounded. In the control case, the cost $J(x^{t:t+N}, u^{t:t+N-1})$ can in general be unbounded. Thus, to make things compatible with the theory of coherent upper previsions (Walley, 1991), in the following it is assumed that the cost J is a bounded real-valued function or, equivalently, that the worst-case value of the cost is finite. Since J is a function of states and control inputs, this

is also equivalent to assume that $x_k \in \mathcal{X}$ and $u_k \in \mathcal{U}$, where $\mathcal{X}$ and $\mathcal{U}$ are compact convex sets including the origin. The assumptions $x_k \in \mathcal{X}$ and $u_k \in \mathcal{U}$ can actually be interpreted as large constraints on state and control due to the physical system's limitations; notice that these constraints are never active along the desirable (nominal) trajectories of the system. [3]

From (9), it turns out that, in absence of contamination, the control inputs are determined by minimizing the expectation $E_{X^{t:t+N}}[J(X^{t:t+N}, u^{t:t+N-1})|y^t, u^{t+N-1}]$ or, equivalently, the joint expectation:

$$E_{X^{t+N},Y^t}[J(X^{t:t+N}, u^{t:t+N-1})|u^{t+N-1}]. \tag{17}$$

It is, in fact, well known that the joint and posterior expectations differ only by a positive constant that depends on the observations only and, thus, does not affect minimization.

In presence of contaminating disturbances, one must deal with sets of distributions or, equivalently, upper expectations. Then, a way to robustly tackle the presence of contaminating disturbances is by following a minimax approach and replace (9) by:

$$\min_{\substack{u_k \in \mathcal{U} \\ k=t,\dots,t+N-1}} \overline{E}_{X^{t:t+N}} \left\{ J(X^{t:t+N}, u^{t:t+N-1}) \middle| y^t, u^{t+N-1} \right\}, \tag{18}$$

i.e., minimizing the predictive posterior upper expectation of the cost. The question is now how to derive the above upper expectation from the knowledge of $\overline{E}_{X_0}$, $\overline{E}_{X_k}(\cdot|x_{k-1}, u_{k-1})$, $\overline{E}_{Y_k}(\cdot|x_k)$, the observations y^t and the control inputs u^{t+N-1}.

Theorem 1. *Assume that the measurements y_k for any $k = 1, \dots, t$ are idealisations of discrete events $\tilde{y}_k = B(y_k, \delta)$, where $B(y_k, \delta)$ are neighborhoods of y_k with positive probability and which converge to $\{y_k\}$ as their radius $\delta > 0$ decreases to zero. Furthermore, assume that*

$$\underline{E}_{X^t,Y^t}[I_{\{\tilde{y}^t\}}(Y^t)] > 0 \tag{19}$$

where $I_{\{\tilde{y}^t\}}(Y^t)$ denotes the indicator function [4] over the discretized observations $\tilde{y}^t$. Then, given u^{t+N-1}, the posterior CUP

$$\overline{E}_{X^{t:t+N}} \left\{ J(X^{t:t+N}, u^{t:t+N-1}) \middle| \tilde{y}^t, u^{t+N-1} \right\} \tag{20}$$

is equal to the unique value $\mu \in \mathbb{R}^+$ that satisfies the following equation:

$$0 = \overline{E}_{X^t,Y^t} \left[I_{\{\tilde{y}^t\}}(Y^t) \cdot (g'(X_t, u^{t:t+N-1}) - \mu) \middle| u^{t-1} \right], \tag{21}$$

where the above joint upper prevision is given by:

$$\overline{E}_{X_0} \Big[\overline{E}_{X_1} \Big[\overline{E}_{Y_1} \Big[\dots \overline{E}_{X_t} \Big[\overline{E}_{Y_t} \Big[I_{\{\tilde{y}^t\}}(Y^t) \cdot (g'(X_t, u^{t:t+N-1}) - \mu) \Big| X_t \Big] \Big| X_{t-1}, u_{t-1} \Big] \dots \Big| X_1 \Big] \Big| X_0, u_0 \Big] \Big], \tag{22}$$

and

$$g'(X_t, u^{t:t+N-1}) = \overline{E}_{X_{t+1}} \Big[\overline{E}_{X_{t+2}} \Big[\dots \overline{E}_{X_{t+N}} \Big[J(X^{t:t+N}, u^{t:t+N-1}) \Big| X_{t+N-1}, u_{t+N-1} \Big] \dots \Big| X_{t+1}, u_{t+1} \Big] \Big| X_t, u_t \Big]. \tag{23}$$

■

For a rigorous proof of Theorem 1, the reader is referred to (Benavoli et al., 2010). Hereafter, an intuitive justification is provided. Consider the joint CUP in (22), which is obtained as follows:

$$\sup_{p(x^t,y^t) \in \mathcal{P}_{X^t,Y^t}} \int_{x^t} \sum_{y^t} I_{\{\tilde{y}^t\}}(y^t) \cdot (g'(x_t, u^{t:t+N-1}) - \mu)\, p(x^t, y^t)\, dx^t \tag{24}$$

where $\mathcal{P}_{X^t,Y^t}$ is the set of distributions associated to $\overline{E}_{X^t,Y^t}$. Remember that the observations have been assumed discrete (from which the sum instead of the integral). The need of this assumption will be clarified in the sequel. By summing the Y^t variables, the argument of the supremum in (24) reduces to:

$$\int_{x^t} (g'(x_t, u^{t:t+N-1}) - \mu)\; p(x^t, \tilde{y}^t)\; dx^t \tag{25}$$

because of the indicator over $\tilde{y}^t$. By Bayes' rule, it holds that: $p(x^t|\tilde{y}^t) = p(x^t, \tilde{y}^t)/p(\tilde{y}^t)$, under the assumption that $p(\tilde{y}^t) > 0$ for all $\tilde{y}^t$ and $p(x^t, y^t) \in \mathcal{P}_{X^t,Y^t}$, where $p(x^t, y^t)$ is the joint PDF from which $p(y^t)$ has been marginalized. This constraint of positiveness can equivalently be rewritten as:

$$\inf_{p(x^t,y^t) \in \mathcal{P}_{X^t,Y^t}} E_{X^t,Y^t}[I_{\{\tilde{y}^t\}}(Y^t)] = \underline{E}_{X^t,Y^t}[I_{\{\tilde{y}^t\}}(Y^t)] > 0,$$

for all $\tilde{y}^t$. Thus, in Theorem 1, the hypothesis (19) is needed for the applicability of Bayes' rule. Notice that, in the continuous case, the probability that a random variable assumes a particular value is zero, which in our context means that we would have $\underline{E}_{X^t,Y^t}\left[I_{\{y^t\}}(Y^t)\right] = 0$. This is the reason for the assumption of discrete observations, which allows to impose that $\underline{E}_{X^t,Y^t}[I_{\{\tilde{y}^t\}}] > 0$. Notice that the assumption of discrete measurements makes sense in practice because of the finite resolution of sensors. Hereafter, this issue will be treated in a completely transparent way for the reader, assuming that the resolution δ of the measurement sensors can be considered arbitrarily small.

Under the assumption (19), (24) can then be rewritten as:

$$\sup_{p(x^t,y^t) \in \mathcal{P}_{X^t,Y^t}} \int_{x^t} \Big(g'(x_t, u^{t:t+N-1}) p(x^t|\tilde{y}^t) - \mu \Big) p(\tilde{y}^t) dx^t = 0 \tag{26}$$

Since $p(\tilde{y}^t) > 0$, it can be verified that the value μ solving (26) is

$$\begin{aligned} \mu &= \sup_{p(x^t,y^t) \in \mathcal{P}_{X^t,Y^t}} \int_{x^t} g'(x_t, u^{t:t+N-1}) p(x^t|\tilde{y}^t) dx^t \\ &= \sup_{p(x^t,y^t) \in \mathcal{P}_{X^t,Y^t}} \int_{x_t} g'(x_t, u^{t:t+N-1}) p(x_t|\tilde{y}^t) dx_t \end{aligned} \tag{27}$$

where the last equality follows from the fact that, since $g'(\cdot)$ does not depend on x^{t-1}, these variables can be marginalized out in the integration. By definition, (27) is equal to

[3] This assumption means that there are saturations in the system. In other words, the linearity assumption in (1) is just an approximation and holds for trajectories of the system that are far from the boundaries of $\mathcal{X}$ and $\mathcal{U}$. This assumption is not very restrictive for many practical control problems and allows to use the simple expressions (14)–(16) for the subsequent derivations.

[4] The indicator of the set $A \subseteq \mathcal{X}$ is the function $I_A : \mathcal{X} \to \{0, 1\}$ such that: $I_A(x) = 1$ if x belongs to A, and $I_A(x) = 0$ otherwise.

$\overline{E}_{X_t}\{g'(X_t, u^{t:t+N-1})|\tilde{y}^t, u^{t+N-1}\}$ and from (23) also to (20).

Notice that (21) is called Generalized Bayes Rule (GBR) (Walley, 1991, Ch. 6). From the above considerations, it can therefore be concluded that computing the posterior upper expectation $\overline{E}_{X_t}[g'|\tilde{y}^t, u^{t-1}]$ by solving the GBR equation is equivalent to computing the upper envelope of the conditional expectation $E_{X_t}[g'|\tilde{y}^t, u^{t-1}]$ that, in turn, can be performed by applying Bayes' rule to the set of distributions $\mathcal{P}_{X^t,Y^t}$ associated to the joint upper expectation $\overline{E}_{X^t,Y^t}$. The advantage of computing the posterior CUP through GBR is that, in this way, it is not necessary to work directly with the set of distributions $\mathcal{P}_{X^t,Y^t}$ but it is possible to work more conveniently with the upper expectation associated to it. Finally, notice that the joint upper expectation $\overline{E}_{X^t,Y^t}$ in (22) is obtained through the (nested) composition of the CUPs $\overline{E}_{X_0}$, $\overline{E}_{X_k}(\cdot|x_{k-1}, u_{k-1})$, $\overline{E}_{Y_k}(\cdot|x_k)$. Notice that, as in the standard expectation case, the nested structure of $\overline{E}_{X^t,Y^t}$ is a consequence of the Markov structure of the model (1). Similar comments hold for $g'(\cdot)$ defined in (23), i.e., the predictive part of the upper expectation in (18).

From the properties (C1)–(C3) of the CUP and since $J(x^{t:t+N}, u^{t:t+N-1})$ is a sum of positive terms, the cost g' in (23) can be decomposed as:

$$\begin{aligned} g' &= \ell_t + \overline{E}_{X_{t+1}}\Big[\ell_{t+1} + \overline{E}_{X_{t+2}}\Big[\ell_{t+2} + \dots \\ &+ \overline{E}_{X_{t+N}}\Big[\ell_{t+N}\Big|X_{t+N-1}, u_{t+N-1}\Big] \dots \Big|X_{t+1}, u_{t+1}\Big]\Big|X_t, u_t\Big], \end{aligned} \tag{28}$$

where $\ell_k = ||x_k||^2_{\Psi_x} + ||u_k||^2_{\Psi_u}$ for $k = t, \dots, t+N-1$ and $\ell_{t+N} = ||x_{t+N}||^2_{\Psi_f}$. Hence, from (14) it follows that

$$\begin{aligned} &\overline{E}_{X_{t+N}}\Big[\ell_{t+N}\Big|x_{t+N-1}, u_{t+N-1}\Big] = \\ &\epsilon_w \int_{x_{t+N}} ||x_{t+N}||^2_{\Psi_f} \mathcal{N}(x_{t+N}; Ax_{t+N-1} + Bu_{t+N-1}, Q) dx_{t+N} \\ &+ (1-\epsilon_w) \max_{x_{t+N}\in\mathcal{X}} ||x_{t+N}||^2_{\Psi_f}. \end{aligned} \tag{29}$$

Thus, one gets

$$\begin{aligned} &\overline{E}_{X_{t+N}}\Big[||X_{t+N}||^2_{\Psi_f}\Big|x_{t+N-1}, u_{t+N-1}\Big] \\ &= \epsilon_w\Big[Tr(\Psi_f Q) + ||Ax_{t+N-1} + Bu_{t+N-1}||^2_{\Psi_f}\Big] \\ &+ (1-\epsilon_w) \max_{x_{t+N}\in\mathcal{X}} ||x_{t+N}||^2_{\Psi_f} \\ &= \epsilon_w ||Ax_{t+N-1} + Bu_{t+N-1}||^2_{\Psi_f} + const, \end{aligned} \tag{30}$$

where Tr denotes the trace and *const* includes all the terms that do not depend on the state or the control inputs. Hence, adding $\ell_{t+N-1} = ||x_{t+N-1}||^2_{\Psi_x} + ||u_{t+N-1}||^2_{\Psi_u}$ to (30), one gets:

$$\begin{aligned} &||x_{t+N-1}||^2_{\Psi_x + \epsilon_w A^T\Psi_f A} + 2\epsilon_w x^T_{t+N-1} A^T \Psi_f B u_{t+N-1} \\ &+ ||u_{t+N-1}||^2_{\Psi_u + \epsilon_w B^T \Psi_f B} + const. \end{aligned} \tag{31}$$

Let us introduce the following notation: $\Lambda_{t+N-1} = \Psi_x + \epsilon_w A^T \Psi_f A$, $m_{t+N-1}(u_{t+N-1}) = 2\epsilon_w A^T \Psi_f B u_{t+N-1}$ and $f_{t+N-1}(u_{t+N-1})$ equal to the last two terms in (31). Then, by applying recursively the previous step to (28), at the end of the recursion one gets:

$$g' = ||x_t||^2_{\Lambda_t} + x_t^T m_t + f_t, \tag{32}$$

where the short notation f_t and m_t is introduced to denote $f_t(u^{t:t+N-1})$ and, respectively, $m_t(u^{t:t+N-1})$, which are obtained as follows:

$$\begin{aligned} \Lambda_k &= \Psi_x + \epsilon_w A^T \Lambda_{k+1} A \\ m_k &= 2\epsilon_w A^T \Lambda_{k+1} B u_k + A^T m_{k+1} \\ f_k &= f_{k+1} + ||u_k||^2_{\Psi_u + \epsilon_w B^T \Lambda_{k+1} B} \\ &+ \epsilon_w (Bu_k)^T m_{k+1} \\ &+ (1-\epsilon_w) \max_{x_{k+1}\in\mathcal{X}} \Big[||x_{k+1}||^2_{\Lambda_{k+1}} + \epsilon_w x^T_{k+1} m_{k+1}\Big] \end{aligned} \tag{33}$$

for $k = t+N-2, \dots, t$. Notice that the terms f_t and m_t in (32) do not depend on x_t, but depend on the control sequence $u^{t:t+N-1}$. By replacing (32) in (21), one gets

$$\begin{aligned} 0 &= \overline{E}_{X^t,Y^t}\left[I_{\{\tilde{y}^t\}}(Y^t)(||X_t||^2_{\Lambda_t} + X_t^T m_t - \mu')\middle|u^{t-1}\right] \\ &= \overline{E}_{X^t,Y^t}\left[I_{\{\tilde{y}^t\}}(Y^t)(||X_t||^2_{\Lambda_t} + X_t^T m_t - \mu')\middle|u^{t-1}\right], \end{aligned} \tag{34}$$

where $\mu' = -f_t + \mu$. Finally, the solution of (34) w.r.t. μ is given by

$$\mu = \mu' + f_t = \overline{E}_{X_t}\Big[||X_t||^2_{\Lambda_t} + X_t^T m_t\Big|\tilde{y}^t, u^{t-1}\Big] + f_t. \tag{35}$$

Hence, the control problem in (18) becomes

$$\min_{\substack{u_k\in\mathcal{U} \\ k=t,\dots,t+N-1}} \overline{E}_{X_t}\Big[||X_t||^2_{\Lambda_t} + X_t^T m_t\Big|\tilde{y}^t, u^{t-1}\Big] + f_t. \tag{36}$$

Notice that the value μ which solves (34) depends on the control inputs $u^{t:t+N-1}$. Thus, (34) and (36) must be jointly solved in order to find the optimal control sequence $\hat{u}^{t:t+N-1}$ and then the first sample $u_t = \hat{u}_t$ is actually applied to the plant according to the receding horizon strategy. The resulting receding horizon control algorithm, named LQGVM, is outlined below.

LQGVM control algorithm step (at time t)

I Initialize $u_k \in \mathcal{U}$ for $k = t, \dots, t+N-1$.

II repeat:

(1) Initialization: $\mu_M = \max_{x_t\in\mathcal{X}} g'(x_t, u^{t:t+N-1})$, $\mu_m = 0$.

(2) While $\mu_M - \mu_m > toll$ compute $\mu_c = (\mu_M + \mu_m)/2$.

(a) For $\mu \in \{\mu_m, \mu_M, \mu_c\}$, do:

(i) define $g_t(x_t, \mu) = ||x_t||^2_{\Lambda_t} + x_t^T m_t + f_t - \mu$.

(ii) For $k = t, \dots, 1$, compute:

$$\begin{aligned} &g_{k-1}(x_{k-1}, \mu) \\ &= \overline{E}_{X_k}\left[\overline{E}_{Y_k}\left[I_{\{\tilde{y}_k\}}(Y_k) g_k(X_k, \mu)\middle|X_k\right]\middle|x_{k-1}, u_{k-1}\right] \\ &= \epsilon_w \int_{x_k} g_k(x_k, \mu)\mathcal{N}(x_k; Ax_{k-1} + Bu_{k-1}, Q) \\ &\cdot\Big[\epsilon_v \mathcal{N}(y_k; Cx_k, R) + (1-\epsilon_v) I_{\{g_k(x_k,\mu)>0\}}(x_k)\Big] dx_k \\ &+(1-\epsilon_w) \max_{x_k\in\mathcal{X}} g_k(x_k, \mu)\Big[\epsilon_v \mathcal{N}(y_k; Cx_k, R) \\ &+(1-\epsilon_v) I_{\{g_k(x_k,\mu)>0\}}(x_k)\Big]. \end{aligned} \tag{37}$$

(iii) Compute $\overline{E}_{X_0}\Big[g_0(X_0, \mu)\Big]$, where $\overline{E}_{X_0}$ is defined in (15).

(b) If $\overline{E}_{X_0}\Big[g_0(X_0,\mu_m)\Big]\cdot\overline{E}_{X_0}\Big[g_0(X_0,\mu_c)\Big] < 0$, then $\mu_M=\mu_c$.
(c) Else if $\overline{E}_{X_0}\Big[g_0(X_0,\mu_M)\Big]\cdot\overline{E}_{X_0}\Big[g_0(X_0,\mu_c)\Big] < 0$, then $\mu_m=\mu_c$.
(d) Else go to step 3.
(3) Return μ_c.

III Check if the optimality condition for
$\overline{E}_{X_t}\Big[||X_t||^2_{\Lambda_t}+X_t^T m_t(u^{t:t+N-1})\Big|\tilde{y}^t,u^{t-1}\Big]+f_t(u^{t:t+N-1})$,
(the function to be minimized) is satisfied, where
$\overline{E}_{X_t}\Big[||X_t||^2_{\Lambda_t}+X_t^T m_t(u^{t:t+N-1})\Big|\tilde{y}^t,u^{t-1}\Big]=\mu_c$.

IV If the optimality condition is satisfied return $\hat{u}^{t:t+N-1}=u^{t:t+N-1}$ and exit; else determine the next iteration control sequence $u^{t:t+N-1}$ and go back to step II.

V Apply the control input $u_t=\hat{u}_t$ to the plant.

Steps I–III solve the optimization in (36). In the numerical examples in Section 4, we have used the SNOPT solver (Gill et al., 2002), an SQP (Sequential Quadratic Programming) algorithm for large-scale constrained optimization [5], to implement steps I–III.

In particular, step II implements the bisection method for finding $\overline{E}_{X_t}\Big[||X_t||^2_{\Lambda_t}+X_t^T m_t\Big|\tilde{y}^t,u^{t-1}\Big]$ for a given control sequence $u^{t:t+N-1}$. Notice that equation (37) has been obtained from (14) and (16). The indicator $I_{\{g_k(x_k,\mu)>0\}}(x_k)$ in (37) comes out from the relation $\max_{y_k} I_{\{\tilde{y}_k\}}(y_k)g_k(x_k,\mu)=\max(0,g_k(x_k,\mu))$, where the latter member is equal to $g_k(x_k,\mu)$ whether $g_k(x_k,\mu)>0$.

There is no closed-form solution for the integral of the indicator in (37) and, thus, numerical methods (e.g., particle filter) must be employed. The computations, however, are much simpler in the case $\epsilon_v=1$ (i.e., no contamination on the measurement noise). In fact, since $g_k(x_k,\mu)$ is quadratic in x_k, in this case the first integral in (37) can be solved analytically by exploiting the following property of the Gaussian densities:

$$\begin{aligned}&\mathcal{N}(x_i;\zeta_{i-1},Q)\;\mathcal{N}(y_i;Cx_i,R)=\mathcal{N}(y_i;C\zeta_{i-1},W_a)\\&\cdot\mathcal{N}(x_i;W_bQ^{-1}\zeta_{i-1}+W_bC'R^{-1}y_i,W_b),\end{aligned}\tag{38}$$

where $\zeta_{i-1}=Ax_{i-1}+Bu_{i-1}$, $W_a=R+CQC^T$ and $W_b^{-1}=Q^{-1}+C^TR^{-1}C$. Then at time t, one thus has:

$$\begin{aligned}&\int_{x_t}(||x_t||^2_{\Lambda_t}+x_t^Tm_t)\mathcal{N}(x_t;\zeta_{t-1},Q)\;\mathcal{N}(y_t;Cx_t,R)dx_t\\&=\mathcal{N}(y_t;C\zeta_{t-1},W_a)\left(Tr(\Lambda_tW_b)+||\rho_{t-1}||_{\Lambda_t}+(\rho_{t-1})^Tm_t\right),\end{aligned}$$

where $\rho_{t-1}=W_bQ^{-1}\zeta_{i-1}+W_bC'R^{-1}y_i$. This step can be repeated recursively up to the initial time $t=0$, as described with more details in (Benavoli et al., 2010, Sec. VI).

As it is evident from the above algorithm, the posterior upper prevision cannot be computed recursively (Benavoli et al., 2010). In fact, since CUPs are nonlinear (compare Axiom (C3) with the linearity property of standard expectation), in order to compute $\overline{E}_{X_t}[g_t(X_t,\mu)|\tilde{y}^t,u^{t-1}]$, it becomes necessary to go through the joint upper prevision $\overline{E}_{X^t,Y^t}$, i.e., to propagate back in time the functional $g_t(X_t,\mu)$ until the initial state is reached, and then to find the value of μ which satisfies $\overline{E}_{X_0}[g_t(X_0,\mu)]=0$. In the contamination case, this means that, at the end, the quantity $g_0(X_0,\mu)$ will have 2^t terms (in the case $\epsilon_v=1$). A possible way to overcome this computational issue is to truncate the recursion after L steps in the past by finding a CUP which approximates $\overline{E}_{X_{t-L}}[g|\tilde{y}^{1:t-L},u^{t-L-1}]$ and use X_{t-L} as new initial state. In the numerical example of Section 4, we have used a variable time-window L and approximated $\overline{E}_{X_{t-L}}[g|\tilde{y}^{1:t-L},u^{t-L-1}]$ with $E_{X_{t-L}}[g|\tilde{y}^{1:t-L},u^{t-L-1}]$ (obtained via the KF), whenever $\overline{E}_{X_{t-L}}[X'_{t-L}\cdot X_{t-L}|\tilde{y}^{1:t-L},u^{t-L-1}]-\underline{E}_{X_{t-L}}[X'_{t-L}\cdot X_{t-L}|\tilde{y}^{1:t-L},u^{t-L-1}]\leq\gamma$, where γ is a suitable threshold. The idea is that, whenever $\overline{E}_{X_{t-L}}[X'_{t-L}\cdot X_{t-L}|\tilde{y}^{1:t-L},u^{t-L-1}]$ is suitably close to $\underline{E}_{X_{t-L}}[X'_{t-L}\cdot X_{t-L}|\tilde{y}^{1:t-L},u^{t-L-1}]$, the dynamic system (1) is almost Gaussian and, thus, the CUP $\overline{E}_{X_{t-L}}[g|\tilde{y}^{1:t-L},u^{t-L-1}]$ can be approximated by $E_{X_{t-L}}[g|\tilde{y}^{1:t-L},u^{t-L-1}]$ obtained via the KF.
In the sequel, the control strategy obtained according to the above algorithm will be referred to as the *Linear Quadratic Gaussian-Vacuous Mixture* (LQGVM) control.

Finally, notice that, in the case $\epsilon_w=\epsilon_v=\epsilon_x=1$, all CUPs become standard expectations. The plant (1) is linear, time-invariant and Gaussian and, thus, when Ψ_f is equal to the steady-state solution of the Riccati equation (11), the LQGVM control reduces to the steady-state LQG control. Hence, the control is optimal and stable in the LQG sense.
Conversely in the case in which $\epsilon_v,\epsilon_w,\epsilon_x<1$, the plant (1) is still linear but non stationary (since the distributions of the contaminating noises can be non stationary) and non Gaussian. Furthermore, since the distributions of the contaminating disturbances can be arbitrary (unbounded support, variance, etc.), stability cannot be guaranteed.
A question to be addressed in future work is whether LQGVM provides a stabilizing (in the LQG sense) control law in the case $\epsilon_v,\epsilon_w,\epsilon_x<1$ but in the absence of contaminating disturbances. In this case, one could take advantage of the robustness of LQGVM in presence of sporadic contaminating disturbances but also with the guarantee of stability when the system returns to be Gaussian. Numerical simulations seem to confirm that this is the case as shown in the next section.

4. NUMERICAL EXAMPLE

Monte Carlo simulations have been carried out in order to assess the performance of the LQGVM controller proposed in the previous section. These simulations compare the LQGVM and classical LQG controllers, considering non Gaussian situations. In particular, the following system has been considered:

$$\begin{aligned}&\begin{cases}x_{t+1}=Ax_t+Bu_t+w_t\\y_t=Cx_t+v_t\end{cases}\\&A=\begin{bmatrix}1&1\\0&1\end{bmatrix},\quad B=\begin{bmatrix}0\\1\end{bmatrix},\quad C=[\,1\;0\,],\end{aligned}\tag{39}$$

where $w_t=\epsilon_w w_t^1+(1-\epsilon_w)w_t^2$, $w_t^1\sim\mathcal{N}(0,Q)$, $x_0=\hat{x}_0$ (i.e., $\epsilon_x=1$), $\hat{x}_0\sim N(0,P_0)$, $v_t\sim N(0,R)$ (i.e. $\epsilon_v=1$),

[5] Notice that no constraints violations have occurred in all the simulations since, as already discussed, $x_k\in\mathcal{X}$ and $u_k\in\mathcal{U}$ must be interpreted as large constraints on state and control, which are never active along the nominal trajectories of the system.

Case	ϵ_w	w_t^2	q	r/q	p_0
1	0.999	$7\delta_7/(1-\epsilon_w)$	0.1	1	0.2
2	0.999	$7\delta_7/(1-\epsilon_w)$	0.1	0.1	0.2
3	0.999	$[0,\ 4\delta_7/(1-\epsilon_w)]^T$	0.1	0.1	0.2

Table 1. Simulation parameters. δ_k is the Kronecker delta.

Case	J_{LQGVM}	J_{LQG}	% reduction
1	5.6	6.2	10%
2	5.3	9.4	42%
3	7.8	11	30%

Table 2. Cost comparison between LQGVM and LQG control

$$P_0 = \begin{bmatrix} p_0 & 0 \\ 0 & p_0 \end{bmatrix}, \quad Q = \begin{bmatrix} q & 0 \\ 0 & q \end{bmatrix}, \quad R = r,$$

with $p_0, q, r > 0$. The physical constraints $x_1 \in [-150, 150]$ and $x_2 \in [-30, 30]$ have been assumed for the two state components. Further, in (8) the control horizon has been fixed to $N = 5$ and the weight matrices have been selected as

$$\Psi_x = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \Psi_u = 1,$$

and Ψ_f is chosen equal to the unique non negative solution of the associated algebraic Riccati equation, so that LQGVM reduces to LQG whenever the contaminating disturbances are not present. Furthermore, the following physical constraint $u \in [-20, 20]$ has been assumed for the control input.

Simulations have been performed considering both the system in (39) and its one-dimensional restriction, i.e., $A = B = C = 1$ etc. Note that, in all simulations, both LQGVM and LQG were designed without assuming the knowledge of the contaminating terms w_t^2. The aim is to investigate the relative sensitivity of both LQG and LQGVM to (heavy tailed) disturbances of the nominal Gaussian density. A trajectory of 16 time steps and 100 independent Monte Carlo runs have been considered. For the one-dimensional system, cases 1 and 2 detailed in table 1 have been simulated.

Notice that in both cases the disturbance realization undergoes a jump of 7 units at the time instant $t = 7$. This can be interpreted as an unmodelled disturbance. The choice of such a large value of ϵ is motivated from the fact that normally the system's behavior is expected to be Gaussian apart from sporadic instants in which non Gaussian disturbances can be present.

The corresponding simulation results are shown in figs. 1–2 where the plotted trajectories are averaged over the 100 Monte Carlo trials. Notice that, since the performance of LQG depends on the quality of the estimate provided by the Kalman filter, the estimated states, i.e. $E_{X_t}[X_t|y^t, u^t]$, are also reported. Similarly, for the LQGVM controller, the lower and upper means, i.e. $\underline{E}_{X_t}[X_t|y^t, u^t]$ and $\overline{E}_{X_t}[X_t|y^t, u^t]$, are also displayed in order to show how the better performance of LQGVM is mainly due to the higher robustness to modelling uncertainties.

In the upper plot, figs. 1–2 report the true trajectory of the system when controlled by LQG, the true trajectory of the system when controlled by LQGVM , the Kalman filter's state estimate (KF) for the system controlled by LQG and the lower/upper state estimates (lower/upper) for the system controlled by LQGVM . Conversely, the lower plot displays the control input trajectories obtained via LQGVM and, respectively, LQG. From figs. 1-2, it can be seen that whenever no contaminating disturbances are present (i.e. in the time interval $[1, 7]$): (1) the LQGVM - and LQG-controlled state trajectories coincide; (2) the KF estimate coincides with both the lower and upper means. This happens, of course, due to system's Gaussianity in the time interval $[1, 7]$.

After the jump's instant, $t = 8$, the KF underestimates the value of the state. Conversely, looking at the lower and upper means, it is evident that the filter based on CUPs correctly detects the jump and is able to enlarge the difference between the upper and lower means in order to include the true value of the state. The difference $\overline{E}_{X_t}[X_t|y^t, u^t] - \underline{E}_{X_t}[X_t|y^t, u^t]$ is related to the model imprecision present in the system during the jump. The higher robustness of the CUP-based filter to detect unmodelled behaviors reflects also into a better control performance. Notice, in fact, that in both figs. 1-2 (upper plots), the LQGVM -controlled system converges to zero, after the jump, faster than the LQG-controlled system. This is obtained by the application of a stronger control input at time $t = 8$, as evident from the lower plots in figs. 1-2. In the case 2, the advantage of LQGVM over LQG is enhanced due to the slower convergence rate of the KF estimate to the true state, as evident from fig. 2. The cost (8), averaged over the 100 MC runs and the whole trajectory, can be used to numerically compare the LQG and LQGVM performances The results of this comparison are reported in table 2 (cases 1 and 2). It is evident that LQGVM provides a lower average cost and, thus, a better performance than LQG in non Gaussian situations.

Consider now the two-dimensional system and the case 3 in table 1. In this case, w_t^2 can be interpreted as an unmodelled disturbance which acts only on the second component of the state. For this case, a trajectory of 20 time steps has been considered but only the time steps from 6 to 16 have been shown in fig. 3 (the trajectories of systems controlled by LQG and LQGVM are similar in the time intervals $[0, 6]$ and $[16, 20]$). From fig. 3 and the average cost in table 2 (case 3), it can be noticed that the behaviour of the systems controlled by LQG and LQGVM is similar to that discussed above for the one-dimensional system. The only difference w.r.t. the one-dimensional system is that the unmodelled disturbance is detected with a delay of one time instant (i.e., at time $t = 9$); this is because the disturbance acts on the second component of the state, which is not directly measured. It is perhaps more interesting to remark the difference between KF estimate and lower/upper means. We can see from the second plot in fig. 3 that at times $t = 9, 10$ the upper mean goes close to 30 which is the assumed upper bound for the second component of the state. This means that the upper mean is almost "vacuous", i.e., $\overline{E}_{X_t}[X_t|y^t, u^t] \approx \max_{x_t \in \mathcal{X}} x_t$. This behaviour is due to the lack of observability for the second component of the state. In fact, the disturbance w_t^2 is so strong that the maximization term in the last equation in (37) becomes dominant. Since this term depends only on the measurement equation, the second component of

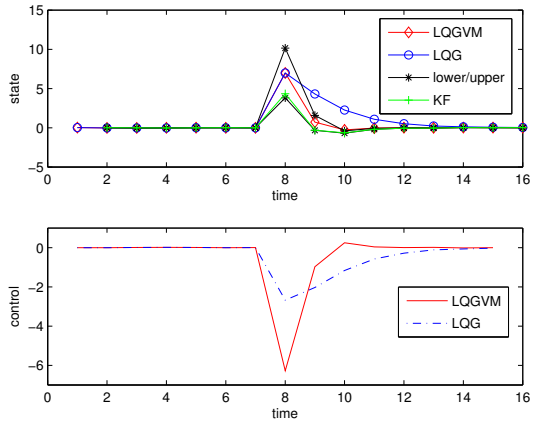

Fig. 1. Case 1: 1D system with $r/q = 1$.

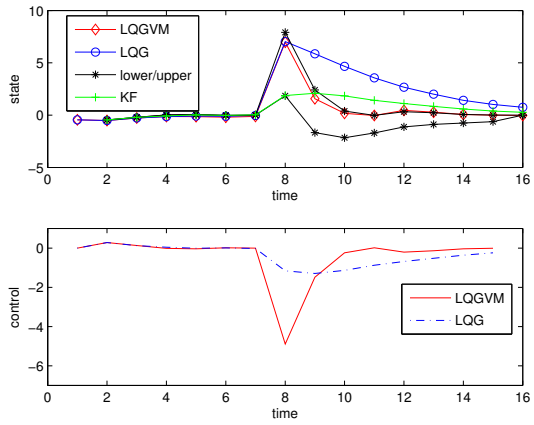

Fig. 2. Case 2: 1D system with $r/q = 10$.

the state is unobservable and, thus, free to vary during the optimisation. In practice, because of the disturbance w_t^2, the information on the second component carried by the prior estimate $\hat{x}_0$ is partially lost at time $t = 9$, and the CUP-based filter has to estimate it again from the measurements. Thus, in some sense, the filter performs a re-initialisation after the jump. This problem does not affect the determination of the control input, because the control cost J takes into account jointly both state components.

5. CONCLUSIONS

In this paper, we have proposed a robust reformulation of LQG control that allows for imprecision in the knowledge about the probabilistic distributions of disturbances and initial state. More specifically, an uncertainty description known as *Gaussian Vacuous Mixture (GVM)* has been adopted in place of the standard Gaussian assumption, and the control objective has been formulated in terms of the minimization of the worst case expectation of a quadratic cost with respect to all admissible probability distributions defined by the GVM uncertainty description. It has been demonstrated by simulation experiments how the resulting controller, named LQGVM, can outperform the LQG controller in non Gaussian situations. Future work will address the following issues: (1) comparison

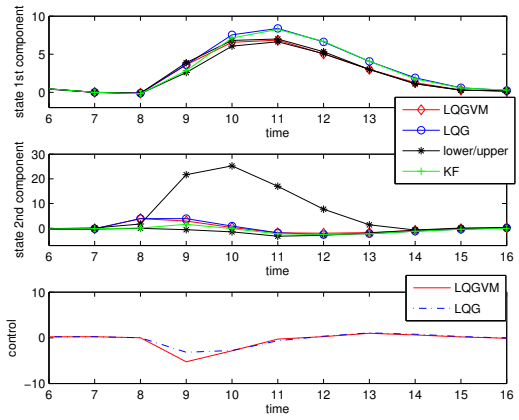

Fig. 3. Case 3: 2D system with $r/q = 1$.

of LQGVM with risk-sensitive LQG (Whittle, 1981) and minimax LQG (Petersen et al., 2000; Petersen, 2006) in terms of performance and computational cost; (2) stochastic stability analysis of LQGVM control.

ACKNOWLEDGEMENTS

This work has been partially supported by the Swiss NSF grants n. 200020-121785/1 and Swiss CTI-project n. 9171.1 ("Hydronet").

REFERENCES

Benavoli, A., Zaffalon, M., and Miranda, E. (2010). Robust filtering through coherent lower previsions. *Accepted for Automatic Control, IEEE Transactions on.* doi: 10.1109/TAC.2010.2090707.

Berger, J.O. (1985). *Statistical Decision Theory and Bayesian Analysis.* Springer Series in Statistics, New York.

Gill, P., Murray, W., and Saunders, M. (2002). SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Journal on Optimization*, 12(4), 979–1006.

Petersen, I. (2006). Minimax LQG control. *International Journal of Applied Mathematics and Computer Science*, 16(3), 309.

Petersen, I., James, M., and Dupuis, P. (2000). Minimax optimal control of stochastic uncertain systems with relative entropy constraints. *Automatic Control, IEEE Transactions on*, 45(3), 398–412.

Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities.* Chapman and Hall, New York.

Whittle, P. (1981). Risk-sensitive linear/quadratic/Gaussian control. *Advances in Applied Probability*, 13(4), 764–777.