Automatica 67 (2016) 216–223

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Brief paper

Mean and variance of the LQG cost function ☆

Hildo Bijl [a], Jan-Willem van Wingerden [a], Thomas B. Schön [b], Michel Verhaegen [a]

[a] Delft Center for Systems and Control, Delft University of Technology, The Netherlands

[b] Department of Information Technology, Uppsala University, Sweden

ARTICLE INFO

Article history:
Received 13 May 2015
Received in revised form
30 September 2015
Accepted 17 December 2015
Available online 4 February 2016

Keywords:
Linear systems
Linear quadratic regulators
LQG control
Lyapunov equation
Probability density function
Matrix algebra

ABSTRACT

Linear Quadratic Gaussian (LQG) systems are well-understood and methods to minimize the expected cost are readily available. Less is known about the statistical properties of the resulting cost function. The contribution of this paper is a set of analytic expressions for the mean and variance of the LQG cost function. These expressions are derived using two different methods, one using solutions to Lyapunov equations and the other using only matrix exponentials. Both the discounted and the non-discounted cost function are considered, as well as the finite-time and the infinite-time cost function. The derived expressions are successfully applied to an example system to reduce the probability of the cost exceeding a given threshold.

1. Introduction

The Linear-Quadratic-Gaussian (LQG) control paradigm is generally well-understood in literature. (See for instance Anderson & Moore, 1990; Åström, 1970; Bosgra, Kwakernaak, & Meinsma, 2008; Skogestad & Postlethwaite, 2005.) There are many methods available of calculating and minimizing the expected cost $\mathbb{E}[J]$. However, much less is known about the resulting distribution of the cost function J. Yet in many cases (like in machine learning applications, in risk analysis and similar stochastic problems) knowledge of the full distribution of the cost function J, or at least knowledge of its variance $\mathbb{V}[J]$, is important. That is the focus of this paper. We derive analytical expressions for both the mean $\mathbb{E}[J]$ and the variance $\mathbb{V}[J]$ of the cost function distribution for a variety of cases. The expressions for the variance $\mathbb{V}[J]$ have not been published before, making that the main contribution of this paper.

The cost function J is usually defined as an integral over a squared non-zero-mean Gaussian process, turning its distribution into a generalized noncentral χ^2 distribution. This distribution does not have a known Probability Density Function (PDF), although its properties have been studied before in literature, for instance in Rice (1944), Sain and Liberty (1971) and Schwartz (1970) and methods to approximate it are discussed in Mathai and Provost (1992) and Davies (1980). No expressions for the variance of the LQG system cost function are given though.

In LQG control most methods focus on the expected cost $\mathbb{E}[J]$, but not all. For instance, Minimum Variance Control (MVC) (see Åström, 1970) minimizes the variance of the output $\boldsymbol{y}$, while Variance Constrained LQG (VCLQG) (see Collins & Selekwa, 1999; Conway & Horowitz, 2008) minimizes the cost function subject to bounds on the variance of the state $\boldsymbol{x}$ and/or the input $\boldsymbol{u}$. Alternatively, in Minimal Cost Variance (MCV) control (see Kang, Aduba, & Won, 2014; Won, Schrader, & Michel, 2008) the mean cost $\mathbb{E}[J]$ is fixed through an equality constraint and the cost variance $\mathbb{V}[J]$ (or alternatively the cost cumulant) is then minimized. However, expressions for the cost variance $\mathbb{V}[J]$ are still not given.

This paper is set up as follows. We present the problem formulation in Section 2 and derive the expressions that solve this problem in Section 3, also making use of the appendices. Section 4 then shows how the equations can be applied to LQG systems, which is subsequently done in Section 5. Finally, Section 6 contains the conclusions.

☆ This research is supported by the Dutch Technology Foundation STW, which is part of the Netherlands Organization for Scientific Research (NWO), and which is partly funded by the Ministry of Economic Affairs (Project number: 12173, SMART-WIND). The work was also supported by the Swedish research Council (VR) via the project Probabilistic modeling of dynamical systems (Contract number: 621-2013-5524). The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Valery Ugrinovskii under the direction of Editor Ian R. Petersen.

E-mail addresses: h.j.bijl@tudelft.nl (H.J. Bijl), j.w.vanwingerden@tudelft.nl (J.W. van Wingerden), thomas.schon@it.uu.se (T.B. Schön), m.verhaegen@tudelft.nl (M. Verhaegen).

http://dx.doi.org/10.1016/j.automatica.2016.01.030

Table 1
The theorems with which the mean and variance of J and J_T can be found, as well as the requirements for these theorems.

	If $\alpha \neq 0$	If $\alpha = 0$	Requirements
$\mathbb{E}[J_T]$	Theorem 1	Theorem 3	A and A_α Sylvester
$\mathbb{E}[J]$	Theorem 2		$\alpha < 0$ and A_α stable
$\mathbb{V}[J_T]$	Theorem 4	Theorem 6	$A_{-\alpha}$, A, A_α and $A_{2\alpha}$ Sylvester
$\mathbb{E}[J]$	Theorem 5		$\alpha < 0$ and A_α stable

2. Problem formulation

We consider continuous linear systems subject to stochastic process noise. Formally, we write these as

$$d\boldsymbol{x}(t) = A\boldsymbol{x}(t)\,dt + d\boldsymbol{w}(t), \tag{1}$$

where $\boldsymbol{w}(t)$ is a vector of Brownian motions. (Note that (1) is not an LQG system, because it is lacking input. The extension to LQG systems will be discussed in Section 4.) As a result, $d\boldsymbol{w}(t)$ is a Gaussian random process with zero-mean and an (assumed constant) covariance of $V\,dt$. Within the field of control (see for instance Skogestad & Postlethwaite, 2005) this system is generally rewritten according to

$$\dot{\boldsymbol{x}}(t) = A\boldsymbol{x}(t) + \boldsymbol{v}(t), \tag{2}$$

where $\boldsymbol{v}(t)$ is zero-mean Gaussian white noise with intensity V. That is, $\mathbb{E}[\boldsymbol{v}(t)\boldsymbol{v}^T(\tau)] = V\delta(t-\tau)$, with $\delta(.)$ the Kronecker delta function. From a formal mathematical perspective this simplification is incorrect, because $\boldsymbol{v}(t)$ is not measurable with nonzero probability. However, since this notation is common in the control literature, and since it prevents us from having to evaluate the corresponding Itô integrals, we will stick with it, although the reader is referred to Øksendal (1985) for methods to properly deal with stochastic differential equations.

We assume that the initial state $\boldsymbol{x}(0) = \boldsymbol{x}_0$ has a Gaussian distribution satisfying

$$\boldsymbol{\mu}_0 \equiv \mathbb{E}[\boldsymbol{x}_0] \quad \text{and} \quad \Sigma_0 \equiv \mathbb{E}[\boldsymbol{x}_0\boldsymbol{x}_0^T]. \tag{3}$$

Note that the variance of $\boldsymbol{x}_0$ is *not* Σ_0, but actually equals $\Sigma_0 - \boldsymbol{\mu}_0\boldsymbol{\mu}_0^T$. We will use two different cost functions in this paper: the infinite-time cost J and the finite-time cost J_T, respectively defined as

$$J \equiv \int_0^\infty e^{2\alpha t}\boldsymbol{x}^T(t)Q\boldsymbol{x}(t)\,dt, \tag{4}$$

$$J_T \equiv \int_0^T e^{2\alpha t}\boldsymbol{x}^T(t)Q\boldsymbol{x}(t)\,dt, \tag{5}$$

where Q is a user-defined symmetric weight matrix. The parameter α can be positive or negative. If it is positive, it is known as the prescribed degree of stability (see Anderson & Moore, 1990 or Bosgra et al., 2008), while if it is negative (like in Reinforcement Learning applications) it is known as the discount exponent.

3. Mean and variance of the LQG cost function

In this section we derive expressions for $\mathbb{E}[J]$, $\mathbb{E}[J_T]$, $\mathbb{V}[J]$ and $\mathbb{V}[J_T]$. An overview of derived theorems, as well as the corresponding requirements, is shown in Table 1.

3.1. Notation and terminology

Concerning the evolution of the state, we define $\boldsymbol{\mu}(t) \equiv \mathbb{E}[\boldsymbol{x}(t)]$, $\Sigma(t) \equiv \mathbb{E}[\boldsymbol{x}(t)\boldsymbol{x}^T(t)]$ and $\Sigma(t_1,t_2) \equiv \mathbb{E}[\boldsymbol{x}(t_1)\boldsymbol{x}^T(t_2)]$. These quantities can be found through the theorems of Appendix A.

We define the matrices $A_\alpha \equiv A+\alpha I$ and similarly $A_{k\alpha} \equiv A+k\alpha I$ for any number k. We also define $X^Q_{k\alpha}$ and $\bar{X}^Q_{k\alpha}$ to be the solutions of the Lyapunov equations

$$A_{k\alpha}X^Q_{k\alpha} + X^Q_{k\alpha}A^T_{k\alpha} + Q = 0, \tag{6}$$

$$A^T_{k\alpha}\bar{X}^Q_{k\alpha} + \bar{X}^Q_{k\alpha}A_{k\alpha} + Q = 0. \tag{7}$$

We often have $\alpha = 0$. In this case A_0 equals A, and we similarly shorten X^Q_0 to X^Q. The structure inherent in the Lyapunov equation induces interesting properties in its solutions $X^Q_{k\alpha}$, which are outlined in Appendix B.

We define the time-dependent solution $X^Q_{k\alpha}(t_1,t_2)$ as

$$X^Q_{k\alpha}(t_1,t_2) = \int_{t_1}^{t_2} e^{A_{k\alpha}t}Qe^{A^T_{k\alpha}t}\,dt. \tag{8}$$

This integral can be calculated efficiently by solving a Lyapunov equation. (See Theorem 14.) Often it happens that the lower limit t_1 of $X^Q_{k\alpha}(t_1,t_2)$ equals zero. To simplify notation, we then write $X^Q_{k\alpha}(t) \equiv X^Q_{k\alpha}(0,t)$. Another integral solution $\tilde{X}^Q_{k_1\alpha,k_2\alpha}(T)$ is defined as

$$\tilde{X}^Q_{k_1\alpha,k_2\alpha}(T) \equiv \int_0^T e^{A_{k_1\alpha}(T-t)}Qe^{A_{k_2\alpha}t}\,dt. \tag{9}$$

This quantity can be calculated (see van Loan, 1978) through

$$\tilde{X}^Q_{\alpha_1,\alpha_2}(T) = \begin{bmatrix} I & 0 \end{bmatrix} \exp\left(\begin{bmatrix} A_{\alpha_1} & Q \\ 0 & A_{\alpha_2} \end{bmatrix} T\right) \begin{bmatrix} 0 \\ I \end{bmatrix}. \tag{10}$$

Considering terminology, we say that a matrix A is *stable* (Hurwitz) if and only if it has no eigenvalue λ_i with a real part equal to or larger than zero. Similarly, we say that a matrix A is *Sylvester* if and only if it has no two eigenvalues λ_i and λ_j (with possibly $i = j$) satisfying $\lambda_i = -\lambda_j$. This latter definition is new in literature, but to the best of our knowledge, no term for this matrix property has been defined earlier.

3.2. The expected cost

We now examine the expected costs $\mathbb{E}[J]$ and $\mathbb{E}[J_T]$. Expressions for these costs are already known for various special cases. (See for instance Åström, 1970; Bosgra et al., 2008.) To provide a complete overview of the subject, we have included expressions which are as general as possible.

Theorem 1. *Consider system (2). Assume that $\alpha \neq 0$ and that A and A_α are both Sylvester. The expected value $\mathbb{E}[J_T]$ of the finite-time cost J_T (5) then equals*

$$\mathrm{tr}\left(\left(\Sigma_0 - e^{2\alpha T}\Sigma(T) + \left(1-e^{2\alpha T}\right)\left(\frac{-V}{2\alpha}\right)\right)\bar{X}^Q_\alpha\right). \tag{11}$$

Proof. From (5) follows directly that

$$\mathbb{E}[J_T] = \mathrm{tr}\left(\int_0^T e^{2\alpha t}\Sigma(t)\,dt\,Q\right) = \mathrm{tr}\left(Y(T)Q\right), \tag{12}$$

where $Y(T)$ is defined as the above integral. To find it, we multiply (A.7) by $e^{2\alpha t}$ and integrate it to get

$$\int_0^T e^{2\alpha t}\dot{\Sigma}(t)\,dt = AY(T) + Y(T)A^T + \int_0^T e^{2\alpha t}V\,dt. \tag{13}$$

The left part, through integration by parts, must equal

$$\int_0^T e^{2\alpha t}\dot{\Sigma}(t)\,dt = \left(e^{2\alpha T}\Sigma(T) - \Sigma_0\right) - 2\alpha Y(T). \tag{14}$$