Automatica 50 (2014) 2546–2554

Contents lists available at ScienceDirect

Automatica

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

Brief paper

Continuous-time norm-constrained Kalman filtering[☆]

James Richard Forbes[a,1], Anton H.J. de Ruiter[b], David Evan Zlotnik[a]

[a] Department of Aerospace Engineering, University of Michigan, FXB Building, 1320 Beal Avenue, Ann Arbor, MI, 48109, USA

[b] Department of Aerospace Engineering, Ryerson University, 350 Victoria Street, Toronto, ON, Canada, M5B 2K3

ARTICLE INFO

Article history:
Received 28 July 2013
Received in revised form
24 December 2013
Accepted 12 May 2014
Available online 27 August 2014

Keywords:
State estimation
Constraints
Kalman filters
Extended Kalman filters

ABSTRACT

This paper considers continuous-time state estimation when part of the state estimate or the entire state estimate is norm-constrained. In the former case continuous-time state estimation is considered by posing a constrained optimization problem. The optimization problem can be broken up into two separate optimization problems, one which solves for the optimal observer gain associated with the unconstrained state estimates, while the other solves for the optimal observer gain associated with the constrained state estimates. The optimal constrained state estimate is found by projecting the time derivative of an unconstrained estimate onto the tangent space associated with the norm constraint. The special case where the entire state estimate is norm-constrained is briefly discussed. The utility of the filtering results developed are highlighted through a spacecraft attitude estimation example. Numerical simulation results are included.

1. Introduction

The control of a system often relies on an estimate of the system state. Moreover, the majority of real systems are nonlinear. For instance, estimates of position, velocity, attitude, and angular velocity are needed to control spacecraft, aircraft, and ground vehicles. As a result, the development of state estimators that can robustly and reliably provide a state estimate of a nonlinear process is paramount.

Broadly speaking, stochastic estimation methods can be divided into two main categories (Crassidis & Junkins, 2012; Jazwinski, 1970; Simon, 2006): batch methods and sequential methods. Batch methods, such as weighted-least-squares methods, sliding-window filters, and smoothers, use many or all measurements to estimate the state of the system over a range of time. Sequential methods, the most popular being the Kalman filter (Kalman, 1960), provide a state estimate in "one-step-ahead" fashion. Although batch methods can generally provide a better state estimate, for real-time and online applications, one-step-ahead methods are often preferred. Historically, the Kalman filter and its nonlinear variants (e.g., the extended Kalman filter (EKF) Simon, 2006, pp. 400–403, the unscented Kalman filter (UKF) Julier, Uhlmann, & Durrant-Whyte, 2000) have proven to be both computationally efficient and reliable. However, the traditional Kalman filter structure has no means to directly handle state constraints.

Various authors have considered discrete-time Kalman filtering while simultaneously accounting for linear or nonlinear state constraints. Inspiration for the present paper comes from Zanetti, Majji, Bishop, and Mortari (2009) where Kalman filtering in a discrete-time setting directly considering a norm constraint on all or part of the state is considered. The derivation of the discrete-time norm-constrained Kalman filter is accomplished by augmenting the objective function, that being the minimization of the error covariance, with the norm constraint. A particularly interesting result highlighted in Zanetti et al. (2009) is that normalizing the unconstrained estimate is in fact optimal.

Numerous other papers considering linear and nonlinear state constraints appear in the literature. For example, in Alouani and Blair (1993), Gupta (2007), Richards (1995), Tahk and Speyer (1990) and Wang, Chiang, and Chang (2002) linear equality state constraints are incorporated into the Kalman filter as pseudo-measurements. Doing so leads to a measurement noise covariance that is singular, which from a theoretical point of view is not problematic, but numerical issues may arise (Simon, 2010). In Gupta (2007) and Simon and Chia (2002) linear equality state constraints are enforced by projecting the unconstrained state estimate generated by the Kalman filter onto the constraint surface. The work of Simon and Chia (2002) is extended in Yang and

[☆] The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Giancarlo Ferrari-Trecate under the direction of Editor Ian R. Petersen.

E-mail addresses: forbesrj@umich.edu (J.R. Forbes), aderuiter@ryerson.ca (A.H.J. de Ruiter), dzlotnik@umich.edu (D.E. Zlotnik).

[1] Tel.: +1 734 764 3310; fax: +1 734 763 0578.

http://dx.doi.org/10.1016/j.automatica.2014.08.007

Blasch (2009) where nonlinear equality constraints are considered. As an alternative to the approach developed in Chen (2010), Ko and Bitmead (2007, 2010) and Simon and Chia (2002) use the linear equality state constraints to formulate a projected system, and then the Kalman filter is applied to the projected system to generate a state estimate. Unscented Kalman filtering accounting nonlinear equality state constraints is considered in Julier and La Viola (2007). The sigma points generated via the unscented transformation are projected onto the constraint surface. After the mean is computed (which does not necessarily satisfy the constraint), the mean is projected onto the constraint surface. For a survey of discrete-time Kalman filtering methods that account for linear and nonlinear state constraints, see Simon (2010).

This paper considers continuous-time Kalman filtering subject to a norm constraint on the state estimates. The main contribution of this work is the derivation of the continuous-time norm-constrained Kalman filter. This has not been previously considered in the literature. Estimating the state when only part of the state estimate is norm constrained and when the entire state estimate is norm constrained is investigated. A subtle feature of the filter presented is that, although a portion or the entire state estimate must satisfy a norm constraint, the true system state does not necessarily have to be constrained in the same way. Additionally, unlike Zanetti et al. (2009) a weight on the norm is incorporated into the filter formulation. Although inspiration for this work comes from Zanetti et al. (2009), the solution presented is different. Following the traditional continuous-time Kalman filter derivation, the time derivative of the error covariance is minimized. However, in order to force the state estimate to satisfy the norm constraint, the objective function is augmented not with the norm constraint directly, but with its time derivative. The solution to the optimization problem posed results in the time derivative of the unconstrained state estimate being projected onto the tangent space of the constraint surface. This projection is not forced upon the filter structure, but rather falls out naturally from the derivation. To showcase the utility of the continuous-time norm-constrained Kalman filter, the filter is used within an extended Kalman filter (EKF) framework to estimate the attitude of a rigid-body spacecraft. Spacecraft attitude estimation has been extensively considered in the literature; see Bar-Itzhack and Oshman (1985), Choukroun, Bar-Itzhack, and Oshman (2006), Shuster (1989) and Shuster and Oh (1981), as well as the survey paper Crassidis, Markley, and Cheng (2007).

The remainder of this paper is as follows. Preliminaries are reviewed in Section 2. Section 3.1 considers state estimation when only part of the state estimate is norm constrained. Norm-constrained Kalman filtering when the entire state estimate is constrained is briefly considered in Section 3.2. The role of a particular matrix, which is in fact a projection matrix, is discussed in Section 3.3. Spacecraft attitude estimation is considered in Section 4. The process and measurement models are presented in Sections 4.1 and 4.2. The EKF form of the estimator, resulting in the continuous-time norm-constrained EKF, is presented in Section 4.3. Numerical simulation results are presented in Section 4.4. The paper is drawn to a close in Section 5.

2. Preliminaries

Consider the continuous-time system

$$\dot{\mathbf{x}}(t) = \mathbf{A}(t)\mathbf{x}(t) + \mathbf{B}(t)\mathbf{u}(t) + \boldsymbol{\Gamma}_w(t)\mathbf{w}(t), \tag{1}$$

$$\mathbf{y}(t) = \mathbf{C}(t)\mathbf{x}(t) + \boldsymbol{\Gamma}_v(t)\mathbf{v}(t), \tag{2}$$

where $\mathbf{x} \in \mathbb{R}^n$ is the system state, $\mathbf{u} \in \mathbb{R}^{n_u}$ is the known control input, $\mathbf{y} \in \mathbb{R}^{n_y}$ is the measurement, $\mathbf{w} \in \mathbb{R}^{n_w}$ is the process noise/disturbance, and $\mathbf{v} \in \mathbb{R}^{n_v}$ is the measurement noise. The time-varying matrices $\mathbf{A}(\cdot)$, $\mathbf{B}(\cdot)$, $\mathbf{C}(\cdot)$, $\boldsymbol{\Gamma}_w(\cdot)$, and $\boldsymbol{\Gamma}_v(\cdot)$ are of appropriate dimension and piecewise continuous, and $\boldsymbol{\Gamma}_v(\cdot)$ has full row rank. The process and measurement noise are assumed to be zero-mean and white with autocovariances $E\left[\mathbf{w}(t)\mathbf{w}^{\mathsf{T}}(\tau)\right] = \mathbf{Q}(t)\delta(t-\tau)$ and $E\left[\mathbf{v}(t)\mathbf{v}^{\mathsf{T}}(\tau)\right] = \mathbf{R}(t)\delta(t-\tau)$, respectively, where $\mathbf{Q}(\cdot) \geq 0$ and $\mathbf{R}(\cdot) > 0$ are piecewise continuous. Additionally, $\mathbf{x}(\cdot)$, $\mathbf{w}(\cdot)$, and $\mathbf{v}(\cdot)$ are assumed to be independent for all time. To be concise, the temporal argument of functions and matrices will no longer be written unless clarity is required.

3. Norm-constrained Kalman filtering

3.1. Norm-constraining part of the state

Consider (1) and (2) partitioned in the following way:

$$\begin{bmatrix} \dot{\mathbf{z}} \\ \dot{\mathbf{q}} \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{A}_{zz} & \mathbf{A}_{zq} \\ \mathbf{A}_{qz} & \mathbf{A}_{qq} \end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{bmatrix} \mathbf{z} \\ \mathbf{q} \end{bmatrix}}_{\mathbf{x}} + \underbrace{\begin{bmatrix} \mathbf{B}_z \\ \mathbf{B}_q \end{bmatrix}}_{\mathbf{B}} \mathbf{u} + \underbrace{\begin{bmatrix} \boldsymbol{\Gamma}_{w,z} \\ \boldsymbol{\Gamma}_{w,q} \end{bmatrix}}_{\boldsymbol{\Gamma}_w} \mathbf{w}, \tag{3}$$

$$\mathbf{y} = \underbrace{\begin{bmatrix} \mathbf{C}_z & \mathbf{C}_q \end{bmatrix}}_{\mathbf{C}} \begin{bmatrix} \mathbf{z} \\ \mathbf{q} \end{bmatrix} + \boldsymbol{\Gamma}_v \mathbf{v}, \tag{4}$$

where $\mathbf{z} \in \mathbb{R}^{n_z}$, $\mathbf{q} \in \mathbb{R}^{n_q}$, and $n = n_z + n_q$. The matrices $\mathbf{A}_{zz}$, $\mathbf{A}_{zq}$, $\mathbf{A}_{qz}$, $\mathbf{A}_{qq}$, $\mathbf{B}_z$, $\mathbf{B}_q$, $\boldsymbol{\Gamma}_{w,z}$, $\boldsymbol{\Gamma}_{w,q}$, $\mathbf{C}_z$, and $\mathbf{C}_q$ are dimensioned appropriately.

Consider the following linear estimator dynamics:

$$\begin{bmatrix} \dot{\hat{\mathbf{z}}} \\ \dot{\hat{\mathbf{q}}} \end{bmatrix} = \begin{bmatrix} \mathbf{A}_{zz} & \mathbf{A}_{zq} \\ \mathbf{A}_{qz} & \mathbf{A}_{qq} \end{bmatrix} \underbrace{\begin{bmatrix} \hat{\mathbf{z}} \\ \hat{\mathbf{q}} \end{bmatrix}}_{\hat{\mathbf{x}}} + \begin{bmatrix} \mathbf{B}_z \\ \mathbf{B}_q \end{bmatrix} \mathbf{u} + \underbrace{\begin{bmatrix} \bar{\mathbf{K}}_z \\ \bar{\mathbf{K}}_q \end{bmatrix}}_{\bar{\mathbf{K}}} \mathbf{r}, \tag{5}$$

where $\hat{\mathbf{z}} \in \mathbb{R}^{n_z}$ is the estimate of $\mathbf{z}$, $\hat{\mathbf{q}} \in \mathbb{R}^{n_q}$ is the estimate of $\mathbf{q}$, $\mathbf{r} = \mathbf{y} - \hat{\mathbf{y}}$ is the measurement residual, and $\hat{\mathbf{y}} = \mathbf{C}_z\hat{\mathbf{z}} + \mathbf{C}_q\hat{\mathbf{q}}$ is the predicted measurement. The observer gain $\bar{\mathbf{K}} \in \mathbb{R}^{n \times n_y}$ has been partitioned into $\bar{\mathbf{K}}_z \in \mathbb{R}^{n_z \times n_y}$ and $\bar{\mathbf{K}}_q \in \mathbb{R}^{n_q \times n_y}$. The estimate $\hat{\mathbf{z}} \in \mathbb{R}^{n_z}$ is not constrained, however, $\hat{\mathbf{q}} \in \mathbb{R}^{n_q}$ is constrained in the following way:

$$\hat{\mathbf{q}}^{\mathsf{T}}\mathbf{W}\hat{\mathbf{q}} = \ell, \quad \forall t \in \mathbb{R}^+, \tag{6}$$

where $\mathbf{W} \in \mathbb{R}^{n_q \times n_q}$, $\mathbf{W} = \mathbf{W}^{\mathsf{T}} > 0$ is a constant weighting matrix. The constraint (6) can be equivalently written as $\left\| \sqrt{\mathbf{W}}\hat{\mathbf{q}} \right\| = \sqrt{\ell}$ where $\sqrt{\mathbf{W}}$ is the square root of the matrix $\mathbf{W}$. Differentiating (6) gives

$$2\hat{\mathbf{q}}^{\mathsf{T}}\mathbf{W}^{\mathsf{T}}\dot{\hat{\mathbf{q}}} = 0, \quad \forall t \in \mathbb{R}^+. \tag{7}$$

The initial state estimates are $\hat{\mathbf{z}}(0)$ and $\hat{\mathbf{q}}(0)$ where $\hat{\mathbf{q}}^{\mathsf{T}}(0)\mathbf{W}\hat{\mathbf{q}}(0) = \ell$. The objective at hand is to find $\bar{\mathbf{K}}$ in an optimal way so that $2\hat{\mathbf{q}}^{\mathsf{T}}\mathbf{W}^{\mathsf{T}}\dot{\hat{\mathbf{q}}} = 0, \forall t \in \mathbb{R}^+$, meaning that $\dot{\hat{\mathbf{q}}}$ must be perpendicular to $\mathbf{W}\hat{\mathbf{q}}$ for all time.

It is worth mentioning that although $\hat{\mathbf{q}}$ must satisfy (6) for all time, the true state $\mathbf{q}$ is not required to satisfy $\mathbf{q}^{\mathsf{T}}\mathbf{W}\mathbf{q} = \ell$. Such a situation may occur when a real system only approximately satisfies $\mathbf{q}^{\mathsf{T}}\mathbf{W}\mathbf{q} = \ell$ due to physical limitations, inaccuracies, or deliberate simplification of a more complicated process.

The estimation error is defined as $\mathbf{e} = \mathbf{x} - \hat{\mathbf{x}}$. Using (3) and (5), along with the definition of the estimation error, the error dynamics are $\dot{\mathbf{e}} = (\mathbf{A} - \bar{\mathbf{K}}\mathbf{C})\mathbf{e} + \boldsymbol{\Gamma}_w\mathbf{w} - \bar{\mathbf{K}}\boldsymbol{\Gamma}_v\mathbf{v}$. Defining the estimation-error covariance to be $\mathbf{P}(t) = E\left[\mathbf{e}(t)\mathbf{e}^{\mathsf{T}}(t)\right]$, and assuming that $\bar{\mathbf{K}}$ is non-random, it is straightforward to show that (Crassidis & Junkins, 2012, p. 170)

$$\dot{\mathbf{P}} = (\mathbf{A} - \bar{\mathbf{K}}\mathbf{C})\mathbf{P} + \mathbf{P}(\mathbf{A} - \bar{\mathbf{K}}\mathbf{C})^{\mathsf{T}} + \boldsymbol{\Gamma}_w\mathbf{Q}\boldsymbol{\Gamma}_w^{\mathsf{T}} + \bar{\mathbf{K}}\boldsymbol{\Gamma}_v\mathbf{R}\boldsymbol{\Gamma}_v^{\mathsf{T}}\bar{\mathbf{K}}^{\mathsf{T}}. \tag{8}$$