Available online at www.sciencedirect.com
SCIENCE DIRECT®
Future Generation Computer Systems 21 (2005) 1047–1055
FGCS FUTURE GENERATION COMPUTER SYSTEMS
www.elsevier.com/locate/fgcs

On the construction of a reduced rank square-root Kalman filter for efficient uncertainty propagation

Dimitri Treebushny [a,*], Henrik Madsen [b]

[a] *Institute of Mathematical Machines and System Problems, NAS Ukraine, Prospekt Glushkova 42, Kiev, Ukraine*
[b] *DHI Water and Environment, Agern Allé 11, DK-2970 Horsholm, Denmark*

Available online 2 April 2004

Abstract

The Kalman filter is a sequential estimation procedure that combines a stochastic dynamical model with observations in order to update the model state and the associated uncertainty. In the situation where no measurements are available the filter works as an uncertainty propagator. The most computationally demanding part of the Kalman filter is to propagate the covariance through the dynamical system, which may be completely infeasible in high-dimensional models. The reduced rank square-root (RRSQRT) filter is a special formulation of the Kalman filter for large-scale applications. In this formulation, the covariance matrix of the model state is expressed in a limited number of modes M. In the classical implementation of the RRSQRT filter the computational costs of the truncation step grow very fast with the number of modes ($>M^3$). In this work, a new approach based on the Lanzcos algorithm is formulated. It provides a more cost-efficient scheme and includes a precision coefficient that can be tuned for specific applications depending on the trade-off between precision and computational load.

Keywords: Kalman filter; Reduced rank square-root; RRSQRT; Lanczos algorithm

1. Introduction

The Kalman filter [6] is a sequential estimation procedure that combines a stochastic dynamical model with observations in order to update the model state and the associated uncertainty. The filter works with the mean and the covariance of the state vector and contains two steps: (i) model propagation and (ii) measurement update. If no measurements are available, the filter can be treated as an uncertainty propagator. The most computationally demanding part of the Kalman filter is to propagate the covariance through the dynamical system.

In absolute precision mathematics for low-dimensional, linear, Gaussian, dynamical systems the classical formulation of the Kalman filter works fine. Problems appear when we are interested in real-life applications, which involve finite computer precision, high-dimensional, non-linear system dynamics, and non-additive, non-Gaussian system noise.

* Corresponding author.
E-mail addresses: dima@env.com.ua (D. Treebushny), hem@dhi.dk (H. Madsen).

0167-739X/$ – see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.future.2004.03.005

The question of development, choice and implementation of suboptimal Kalman filtering [12] in engineering applications has been investigated since the first work of Kalman [6]. Suboptimal Kalman filter procedures should fulfill the following requirements:

1. to be robust in finite precision arithmetic;
2. to effectively reduce the computational burden compared to a straightforward implementation of the classical Kalman filter, which may be completely infeasible in high-dimensional models;
3. to be able to deal with non-linear systems;
4. to be able to deal with non-additive and non-Gaussian noise.

An important step towards filter robustness was made by dealing with a square-root of the covariance matrix S, $S{\cdot}S^{\mathrm{T}} = P$ instead of the full covariance matrix itself P. This approach automatically preserves the property of P to be a covariance matrix:

$$P \geq 0 \quad \text{or} \quad \lambda_i \geq 0, \ \ \lambda_i \text{ is eigenvalue of } P. \tag{1}$$

Calculation of a square-root of a positive definite matrix P can be made by any known algorithm such as Cholesky decomposition or SVD decomposition of a symmetric matrix. Verlaan and Heemink in Ref. [16] remarked that the square-root approach has an additional important feature. Because the square-root matrix has a much smaller range of eigenvalues, or in other words a much smaller condition number, square-root-based algorithms are numerically more stable than the classical Kalman filter.

The second requirement, the filter being feasible for high-dimensional systems, leads to the logical construction:

$$P \approx S{\cdot}S^{\mathrm{T}}, \tag{2}$$

where a huge matrix $P \sim N \times N$ is approximated with its small rank square-root $S \sim N \times M$, $M \ll N$.

There are two known approaches for construction of a small rank square-root approximation: either a stochastic or a deterministic approach. The former is realized in the ensemble Kalman filter [3]. It is based on the fundamental theorem of statistical estimation theory. Let $\{\xi_i;\ i \geq 1\}$ be independent realizations of the same random variable ξ with mean $a = E\xi$ and covariance $P = E\lfloor(\xi - a)(\xi - a)^{\mathrm{T}}\rfloor$. Let

$$a_M = \frac{1}{M}\sum_{i=1}^{M}\xi_i,$$

$$S_M = \frac{1}{\sqrt{M-1}}[\,\xi_1 - a_M \quad \xi_2 - a_M \quad \cdots \quad \xi_M - a_M\,], \tag{3}$$

then with probability 1

$$\begin{aligned} P_M &= S_M \cdot (S_M)^{\mathrm{T}} \\ &= \frac{1}{M-1}\sum_{i=1}^{M}(\xi_i - a_M)\cdot(\xi_i - a_M)^{\mathrm{T}} \to P, \\ &\quad M \to +\infty. \end{aligned} \tag{4}$$

With the ensemble Kalman filter approach all the requirements outlined above are fulfilled:

1. The ensemble KF uses a square-root approximation of the covariance matrix, and hence it is robust in finite precision arithmetic.
2. The associated computational cost for propagation of the covariance matrix is in the order of 100–500 model runs [8,13] which can be easily parallelized. Thus, it has the potential to reduce drastically the computational costs as compared to the classical Kalman filter implementation.
3. It works with any system propagator, linear or non-linear. In this respect it is able to interface the model in a highly transparent way (the model is just a black box for the filter), and hence is feasible for most real-life applications.
4. The filter has no specific restrictions with respect to the definition of system noise. The noise can be introduced in any part of the model.

The only drawback of this approach is the relatively slow convergence of the covariance estimate:

$$P_M - P \sim \frac{1}{\sqrt{M}}, \quad M \to +\infty. \tag{5}$$

In the other approach, a deterministic low rank square-root approximation is constructed in the following way. If P is represented according to its SVD decomposition

$$\begin{aligned} &P = LDL^{\mathrm{T}}, \qquad D = \mathrm{diag}\{\lambda_i\}, \quad \lambda_i \geq \lambda_{i+1} \geq 0, \\ &LL^{\mathrm{T}} = L^{\mathrm{T}}L = I, \end{aligned} \tag{6}$$