

Constrained Kalman filtering via density function truncation for turbofan engine health estimation

Dan Simon^{a*} and Donald L. Simon^b

^aDepartment of Electrical Engineering, Cleveland State University, 2121 Euclid Avenue, Cleveland, OH 44115, USA;

^bNASA Glenn Research Center, Mail Stop 77-1, 21000 Brookpark Road, Cleveland, OH 44135, USA

(Received 7 August 2008; final version received 29 April 2009)

Kalman filters are often used to estimate the state variables of a dynamic system. However, in the application of Kalman filters some known signal information is often either ignored or dealt with heuristically. For instance, state variable constraints (which may be based on physical considerations) are often neglected because they do not fit easily into the structure of the Kalman filter. This article develops an analytic method of incorporating state variable inequality constraints in the Kalman filter. The resultant filter truncates the probability density function (PDF) of the Kalman filter estimate at the known constraints and then computes the constrained filter estimate as the mean of the truncated PDF. The incorporation of state variable constraints increases the computational effort of the filter but also improves its estimation accuracy. The improvement is demonstrated via simulation results obtained from a turbofan engine model. It is also shown that the truncated Kalman filter may provide a more accurate way of incorporating inequality constraints than other constrained filters (e.g. the projection approach to constrained filtering).

Keywords: Kalman filter; state constraints; estimation; probability density function; gas turbine engines; health monitoring; optimal filtering; constrained filtering

1. Introduction

For linear dynamic systems with white process and measurement noise, the Kalman filter is an optimal estimator. However, in the application of Kalman filters there is often known model or signal information that is either ignored or dealt with heuristically (Massicotte, Morawski and Barwicz 1995). This has resulted in recent efforts to incorporate constraints in the Kalman filter. For example, a projection method can be used to find the optimal way to incorporate hard inequality constraints on the states (Simon and Simon 2005). Another way of incorporating constraints is to use a regularisation method to enforce a soft limit on the changes of the state variables with respect to time (Simon and Simon 2006). Yet another approach is the use of ridge regression to bias estimates with low certainty towards their constraints (Dewallef, Leonard and Mathioudakis 2004). Other approaches to incorporating constraints in filtering algorithms include constraints on the mean squared values of the estimates (Vathsal and Sarma 1974) and constraints on the variance of the estimation error (Wang and Shu 2000). Non-linear ways of incorporating state constraints include moving horizon estimation (Rao, Rawlings and Mayne 2003; Goodwin, De Dona, Seron and Zhuo 2005) and particle filtering

(Simon 2006). These methods are more rigorous than Kalman filter modifications, but the trade-off is that they can require a lot of computational effort. Therefore they may not be feasible, depending on the application.

This article presents a way to generalise the Kalman filter such that known inequality constraints among the state variables are satisfied by the state variable estimates. The constraints that are imposed are hard constraints in that they are strictly enforced. However, in contrast to the projection method of constraint enforcement (Simon and Simon 2005), the state estimates are not projected onto the constraint surface. Rather, the probability density function (PDF) that is computed by the Kalman filter is truncated at the constraint edges, and the constrained state estimate becomes equal to the mean of the truncated PDF. This idea is based on a previously published method (Shimada, Shirai, Kuno and Miura 1998) but has been modified to handle two-sided inequality constraints.

The application considered in this article is aircraft turbofan engine health parameter estimation (Doel 1994a). Health parameters represent engine component efficiencies and flow capacities. The performance of gas turbine engines deteriorates over time. This deterioration reduces the fuel economy of the engine.

*Corresponding author. Email: d.j.simon@csuohio.edu

Airlines periodically collect engine data in order to evaluate the health of the engine and its components. The health evaluation can then be used to determine maintenance schedules. This offers the benefits of improved safety and reduced operating costs. The money-saving potential of such health evaluations is substantial, but only if the evaluations are reliable. The data used to perform health evaluations are typically collected during flight and later transferred to ground-based computers for post-flight analysis. Data are collected each flight at the same engine operating point and corrected to account for variability in ambient conditions. Various algorithms have been proposed to monitor engine health, such as weighted least squares (Doel 1994b), expert systems (DePold and Gass 1999), Kalman filters (Volponi, DePold, Ganguli and Daguang 2003), neural networks (Volponi et al. 2003), fuzzy logic (Ganguli 2003) and genetic algorithms (Kobayashi and Simon 2001).

It could be argued that static estimation methods (e.g. least squares) can be used to effectively estimate engine health, and that dynamic methods such as Kalman filtering are not needed. However, if we consider on-board real-time estimation, a dynamic estimator such as a Kalman filter gives better results than a static estimator. More accurate health parameter estimation gives more accurate on-board engine models, which are then used to process transient, dynamic input data.

This article develops the truncation method of constrained Kalman filtering, and then applies it to the estimation of engine health parameters. We use heuristic knowledge of the health parameter dynamics to constrain their estimate. For example, we know that health parameters generally do not improve with time. Engine health generally degrades over time, and we can incorporate this information into state constraints to improve our health parameter estimation. (This is assuming that no maintenance or engine overhaul is performed.) It should be emphasised that in this article we are confining the problem to the estimation of engine health parameters in the presence of degradation only. There are specific engine faults that can result in abrupt shifts in filter estimates, possibly even indicating an apparent improvement in some engine components. An actual engine performance monitoring system would need to include additional logic to detect and isolate such faults (Ganguli and Dan 2004).

Section 2 derives the constrained Kalman filter. Section 3 discusses the problem of turbofan health parameter estimation, along with the dynamic model that we use in our simulation experiments. Although the health parameters are not state variables of the model, the dynamic model is augmented in such a

way that a Kalman filter can estimate the health parameters following the approach described in previous publications (Friedland 1969; Lambert 1991). We show how this problem can be expressed in a way that is compatible with the constraints discussed in Section 2. Section 4 presents some simulation results based on a turbofan model linearized around a known operating point. We show (for our problem) that the truncated Kalman filter can estimate health parameters better than the unconstrained filter, and it can also estimate health parameters better than other constrained filters. Section 5 presents some concluding remarks and suggestions for further work.

2. Constrained Kalman filtering

Consider the discrete linear time-invariant system given by

$$\begin{aligned} x(k+1) &= Ax(k) + Bu(k) + w(k) \\ y(k) &= Cx(k) + e(k), \end{aligned} \quad (1)$$

where k is the time index, x is the state vector, u is the known control input and y is the measurement. The signals $\{w(k)\}$ and $\{e(k)\}$ are uncorrelated zero mean Gaussian noise input sequences with covariances

$$\begin{aligned} E[w(k)w^T(m)] &= Q\delta_{km} \\ E[e(k)e^T(m)] &= R\delta_{km} \\ E[w(k)e^T(m)] &= 0, \end{aligned} \quad (2)$$

where $E[\cdot]$ is the expectation operator and δ_{km} is the Kronecker delta function ($\delta_{km} = 1$ if $k = m$, $\delta_{km} = 0$ otherwise). The Kalman filter equations are given as follows (Simon 2006a):

$$\begin{aligned} K(k) &= A\Sigma(k)C^T(C\Sigma(k)C^T + R)^{-1} \\ \hat{x}(k+1) &= A\hat{x}(k) + Bu(k) + K(k)(y(k) - C\hat{x}(k)) \\ \Sigma(k+1) &= (A\Sigma(k) - K(k)C\Sigma(k))A^T + Q, \end{aligned} \quad (3)$$

where the filter is initialised with $\hat{x}(0) = E[x(0)]$ and $\Sigma(0) = E[(x(0) - \hat{x}(0))(x(0) - \hat{x}(0))^T]$. The Kalman filter estimate $\hat{x}(k)$ is a Gaussian random variable with a mean of $x(k)$ and a covariance matrix of $\Sigma(k)$.

Now suppose that we are given the s scalar constraints

$$a_m(k) \leq \phi_m^T(k)x(k) \leq b_m(k) \quad m = 1, \dots, s, \quad (4)$$

where $a_m(k) < b_m(k)$. This is a two-sided constraint on the linear function of the state $\phi_m^T(k)x(k)$. If we have only a one-sided constraint, then we set $a_m(k) = -\infty$ or $b_m(k) = \infty$. Now suppose at time k that we have some estimate $\hat{x}(k)$ with covariance $\Sigma(k)$. The problem is to truncate the Gaussian PDF $N(x(k), \Sigma(k))$ at the s constraints given in (4), and then find the mean $\tilde{x}(k)$

and covariance $\tilde{\Sigma}(k)$ of the truncated PDF. These new quantities, $\tilde{x}(k)$ and $\tilde{\Sigma}(k)$, become the constrained state estimate and its covariance.

In order to make the problem tractable, we will define $\tilde{x}_i(k)$ as the state estimate after the first i constraints of (4) have been enforced, and $\tilde{\Sigma}_i(k)$ as the covariance of $\tilde{x}_i(k)$. We therefore initialise

$$\begin{aligned} i &= 0 \\ \tilde{x}_i(k) &= \hat{x}(k) \\ \tilde{\Sigma}_i(k) &= \Sigma(k). \end{aligned} \tag{5}$$

The sequence of steps that follows is intended to transform the state vector so that the constraints are decoupled. That is, we will obtain s transformed constraints, with each constraint involving only one transformed state. With this decoupling we will be able to easily enforce the constraints one transformed state at a time. We perform the transformation

$$z_i(k) = S_i W_i^{-1/2} T_i^T (x(k) - \tilde{x}_i(k)). \tag{6}$$

The reason for this transformation is that only one element of $z_i(k)$ is constrained, which makes the required PDF truncation tractable. T_i and W_i are obtained from the Jordan canonical decomposition of $\tilde{\Sigma}_i(k)$,

$$T_i W_i T_i^T = \tilde{\Sigma}_i(k). \tag{7}$$

We see that T_i is orthogonal and W_i is diagonal (therefore its square root is very easy to compute). S_i is obtained by using Gram–Schmidt orthogonalisation (Moon and Stirling 2000) to find the orthogonal S_i that satisfies

$$S_i W_i^{1/2} T_i^T \phi_i(k) = [(\phi_i^T(k) \tilde{\Sigma}_i(k) \phi_i(k))^{1/2} \quad 0 \quad \dots \quad 0]^T. \tag{8}$$

With these definitions we see that the upper bound (4) is transformed as

$$\begin{aligned} \phi_i^T(k) x(k) &\leq b_i(k) \\ \phi_i^T(k) T_i W_i^{1/2} S_i^T z_i(k) + \phi_i^T(k) \tilde{x}_i(k) &\leq b_i(k) \\ \frac{(\phi_i^T(k) T_i W_i^{1/2} S_i^T z_i(k))}{(\phi_i^T(k) \tilde{\Sigma}_i(k) \phi_i(k))^{1/2}} &\leq \frac{b_i(k) - \phi_i^T(k) \tilde{x}_i(k)}{(\phi_i^T(k) \tilde{\Sigma}_i(k) \phi_i(k))^{1/2}}, \\ [1 \quad 0 \quad \dots \quad 0] z_i(k) &\leq d_i(k) \end{aligned} \tag{9}$$

where $d_i(k)$ is defined by the above equation. Similarly we see that

$$\begin{aligned} [1 \quad 0 \quad \dots \quad 0] z_i(k) &\geq \frac{a_i(k) - \phi_i^T(k) \tilde{x}_i(k)}{(\phi_i^T(k) \tilde{\Sigma}_i(k) \phi_i(k))^{1/2}}, \\ &\geq c_i(k) \end{aligned} \tag{10}$$

where $c_i(k)$ is defined by the above equation. We therefore have the normalised scalar constraint

$$c_i(k) \leq [1 \quad 0 \quad \dots \quad 0] z_i(k) \leq d_i(k). \tag{11}$$

Lemma 1: $z_i(k)$ has an identity covariance matrix, and only its first element can have a non-zero mean.

Proof: See Appendix 1. □

Since $z_i(k)$ has an identity covariance matrix, its elements are statistically independent of each other. Only the first element of $z_i(k)$ is constrained, so the PDF truncation reduces to a one-dimensional truncation.

The constraint says that the first element of $z_i(k)$ must lie between $c_i(k)$ and $d_i(k)$. We therefore remove that part of the Gaussian PDF that is outside of the constraints and compute the area of the remaining portion of the PDF as

$$\begin{aligned} &\int_{c_i(k)}^{d_i(k)} \frac{1}{\sqrt{2\pi}} \exp(-\zeta^2/2) d\zeta \\ &= \frac{1}{2} \left[\operatorname{erf}(d_i(k)/\sqrt{2}) - \operatorname{erf}(c_i(k)/\sqrt{2}) \right], \end{aligned} \tag{12}$$

where $\operatorname{erf}(\cdot)$ is the error function, defined as

$$\operatorname{erf}(t) = \frac{2}{\sqrt{\pi}} \int_0^t \exp(-\tau^2) d\tau. \tag{13}$$

We normalise the truncated PDF so it has an area of one, and we find that the truncated PDF (i.e. the constrained PDF of the first element of $z_i(k)$) is given as

$$\begin{aligned} \text{PDF}(\zeta) &= \alpha_i \exp(-\zeta^2/2) \\ \alpha_i &= \frac{\sqrt{2}}{\sqrt{\pi} [\operatorname{erf}(d_i(k)/\sqrt{2}) - \operatorname{erf}(c_i(k)/\sqrt{2})]}. \end{aligned} \tag{14}$$

We can compute the mean and variance of the first element of $z_i(k)$ after constraint enforcement as

$$\begin{aligned} \mu_i &= \alpha_i \int_{c_i(k)}^{d_i(k)} \zeta \exp(-\zeta^2/2) d\zeta \\ &= \alpha_i [\exp(-c_i^2(k)/2) - \exp(-d_i^2(k)/2)] \\ \sigma_i^2 &= \alpha_i \int_{c_i(k)}^{d_i(k)} (\zeta - \mu_i)^2 \exp(-\zeta^2/2) d\zeta \\ &= \alpha_i \left[\exp(-c_i^2(k)/2) (c_i(k) - 2\mu_i) \right. \\ &\quad \left. - \exp(-d_i^2(k)/2) (d_i(k) - 2\mu_i) \right] + \mu_i^2 + 1 \end{aligned} \tag{15}$$

The mean and covariance of the transformed state estimate, after enforcement of the first constraint, are therefore given as

$$\begin{aligned} \tilde{z}_{i+1}(k) &= [\mu_i \quad 0 \quad \dots \quad 0]^T \\ \tilde{C}_{i+1}(k) &= \operatorname{diag}(\sigma_i^2, 1, \dots, 1). \end{aligned} \tag{16}$$

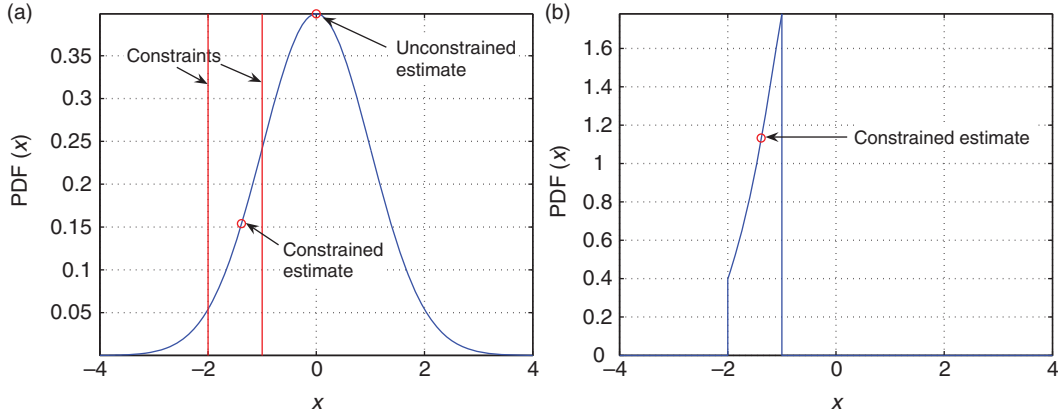


Figure 1. The unconstrained estimate violates the constraints. The constrained estimate is the centroid of the truncated PDF: (a) Unconstrained PDF; (b) constrained PDF.

We then take the inverse of the transformation (6) to find the mean and variance of the state estimate after enforcement of the first constraint.

$$\begin{aligned}\tilde{x}_{i+1}(k) &= T_i W_i^{1/2} S_i^T \tilde{z}_{i+1}(k) + \tilde{x}_i(k) \\ \tilde{\Sigma}_{i+1}(k) &= T_i W_i^{1/2} S_i^T \tilde{C}_{i+1}(k) S_i W_i^{1/2} T_i^T.\end{aligned}\quad (17)$$

We then increment i by one and repeat the process of (6)–(17) to obtain the state estimate after enforcement of the next constraint. Note that $\tilde{x}_0(k)$ is the unconstrained state estimate at time k , $\tilde{x}_1(k)$ is the state estimate at time k after the enforcement of the first constraint, $\tilde{x}_2(k)$ is the state estimate at time k after the enforcement of the first two constraints, etc. After going through this process s times (once for each constraint) we have the final constrained state estimate and covariance at time k .

$$\begin{aligned}\tilde{x}(k) &= \tilde{x}_s(k) \\ \tilde{\Sigma}(k) &= \tilde{\Sigma}_s(k).\end{aligned}\quad (18)$$

After this procedure, all of the original constraints of (4) will be satisfied, assuming that they are linearly independent.

Figure 1 shows an example of a one-dimensional state estimate before and after truncation. Before truncation the state estimate is outside of the state constraints. After truncation, the state estimate is set equal to the mean of the truncated PDF. This figure illustrates the motivation behind the constraint enforcement method presented in the preceding equations. The unconstrained state is assumed to be Gaussian. The Gaussian PDF is truncated at the constraint boundaries, and the constrained estimate is equal to the mean of the truncated PDF. An initial consideration of Figure 1(a) might indicate that the constrained estimate should lie on the constraint boundary. In fact, this is exactly the philosophy of

the projection approach to constrained filtering (Simon and Simon 2005). However, the PDF truncation approach considers both the constraints and the unconstrained Kalman filter's Gaussian distribution of the estimate. The resulting constrained estimate lies at a place within the constraint boundaries that is determined by both the information from the unconstrained filter and the constraints. Simulation results presented later in this article show that this PDF truncation approach outperforms the projection approach for turbofan engine health estimation.

Figure 2 shows another example of PDF truncation. In this case the unconstrained state estimate is inside the state constraints. However, truncation changes the PDF and so the constrained state estimate changes to the mean of the truncated PDF. It could be argued that the estimate should not be changed if it satisfies the constraints. In fact, the PDF truncation filter could be implemented either way. Whether to modify estimates that already satisfy the constraints (as shown in Figure 2), or leave those estimates unchanged, is an implementation decision that depends on the application and the engineer's judgement.

2.1. Bias

Probability density function truncation imposes a bias on the state estimate. The unconstrained Kalman filter has the property that the state estimate is the mean of the true state conditioned on the measurements. However, the truncated state estimate is biased. This is a drawback to this method of constraint enforcement, especially since other methods of constraint enforcement preserve unbiasedness (Simon and Simon 2005). However, if other features of the estimate are more important to the user than unbiasedness (e.g. root mean square (RMS) estimation error) then the truncation approach to constraint enforcement may still be

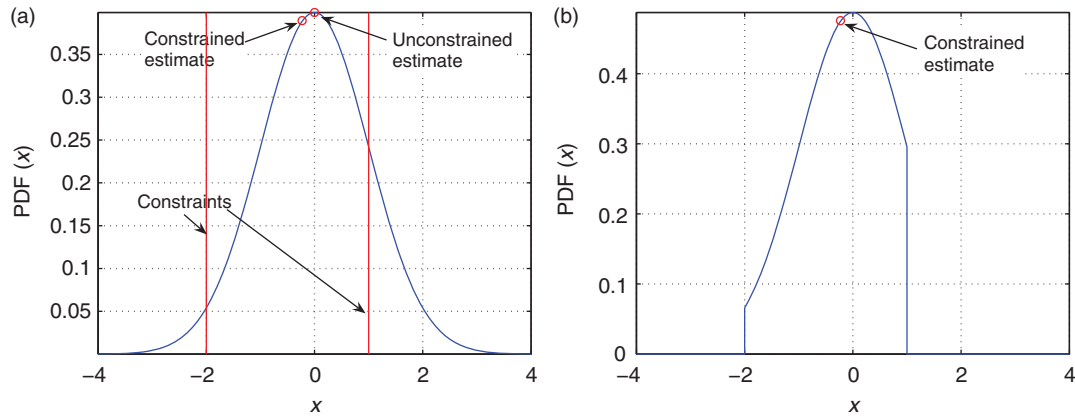


Figure 2. The unconstrained estimate satisfies the constraints. Nevertheless, the truncation approach to constrained estimation shifts the estimate to the centroid of the truncated PDF: (a) Unconstrained PDF; (b) constrained PDF.

attractive. It is shown later in Section 4 that the truncation approach outperforms the projection approach in our example in terms of RMS estimation error in spite of the biasedness of the truncated state estimate.

2.2. Order of constraint application

Since the algorithm of (5)–(18) is non-linear, the constrained state estimate depends on the order in which the constraints are applied. The constraints will all be satisfied regardless of the order of constraint application, but the final state estimate may change if the order is changed. This shows that, in general, PDF truncation is not optimal. However, if the constraints are decoupled, as they are in the turbofan engine health estimation problem discussed in Section 3, then the order of constraint application does not matter.

2.3. Non-recursiveness of PDF truncation

If estimates that satisfy the constraints are modified via PDF truncation, as shown in Figure 2, it is important that the unconstrained Kalman filter be run independently of the truncation process. That is, the truncated state estimate and covariance $\tilde{x}(k)$ and $\tilde{\Sigma}(k)$ must not be used as the starting point for the unconstrained Kalman filter iteration in (3). This is because the information in the constraint at each time step should be limited to that specific time step, and should not be used (even implicitly) in succeeding time steps. The constraint at time k should be used only at time k , and should not have even an indirect effect on the filter at later time steps.

To illustrate this point, suppose that we use the constrained *a posteriori* state estimate to obtain the *a priori* estimate at the next time step. Consider the noise-free scalar system $x_{k+1} = x_k$ and the constraint

$x_k \geq 0$. Suppose that the PDF of the unconstrained *a priori* estimate is zero-mean with a variance of 1. If the next measurement is information-free (i.e. $R = \infty$) the *a posteriori* estimate is the same as the *a priori* estimate. The $x_k \geq 0$ constraint then gives a constrained estimate with a mean of $\sqrt{2/\pi}$ and a variance of $(\pi-2)/\pi$. After the next time update and information-free measurement, we truncate the PDF again. The mean increases and the variance decreases. Continuing this process gives a series of supposedly normal distributions with a mean that increases monotonically and a variance that decreases monotonically. The problem is that the information in the constraint was used recursively.

If the engineer chooses to not modify those state estimates that satisfy the constraints, then constrained state estimates could be recursively fed back into the Kalman filter algorithm without any problems. In that case only estimates that violate the constraints would be modified, which is similar to other approaches of constrained estimation (e.g. the projection approach (Simon and Simon 2005) and moving horizon estimation (Rao et al. 2003)).

2.4. Non-linear constraints

Although we have considered only linear state constraints, it is not conceptually difficult to extend this method to non-linear constraints. If the state constraints are non-linear they can be linearised as discussed in Simon and Chia (2002). The PDF truncation method has been extended to non-linear unscented Kalman filters in Teixeira, Chandrasekar, Torres, Aguirre and Bernstein (2009).

3. Turbofan engine health monitoring

Figure 3 shows a schematic representation of an aircraft turbofan engine (Parker and Guo 2003; Parker and




Figure 3. Schematic representation of a turbofan engine.

Melcher 2004). A single inlet supplies airflow to the fan. Air leaving the fan separates into two streams: one stream passes through the engine core, and the other stream passes through the annular bypass duct. The fan is driven by the low pressure turbine. The air passing through the engine core moves through the compressor, which is driven by the high pressure turbine. Fuel is injected in the main combustor and burned to produce hot gas for driving the turbines. The two air streams combine in the augmentor duct, where additional fuel is added to further increase the air temperature. The air leaves the augmentor through the nozzle, which has a variable cross-section area.

The simulation used in this article is a gas turbine engine simulation software package called Modular Aero Propulsion System Simulation (MAPSS) (Parker and Guo 2003; Parker and Melcher 2004). MAPSS is written using Simulink[®]. The MAPSS engine model is based on a low frequency, transient, performance model of a high-pressure ratio, dual-spool, low-bypass, military-type, variable cycle, turbofan engine with a digital controller. MAPSS does not simulate a specific engine, but instead provides a simulation environment for a generic engine. The controller update rate is 50 Hz, and the component level model balances the mass/energy equations of the system at a rate of 2500 Hz. The three state variables used in MAPSS are low-pressure rotor speed (XNL), high-pressure rotor speed (XNH) and the average hot section metal temperature (TMPC) (measured from aft of the combustor to the high pressure turbine). The discretised time invariant equations that model the turbofan engine can be summarised as follows:

$$\begin{aligned} x(k+1) &= f[x(k), u(k), p(k)] + w_x(k) \\ p(k+1) &= p(k) + w_p(k) \\ y(k) &= g[x(k), u(k), p(k)] + e(k) \end{aligned} \quad (19)$$

where k is the time index, x is the 3-element state vector, u is the 3-element control vector, p is the 10-element health parameter vector and y is the 11-element measurement vector. Note that the noise terms and health parameter degradation are not modelled in MAPSS but have been added to the model for the problem studied in this article. The health parameters change slowly over time. Between measurement times their deviations can be approximated by the zero mean noise $w_p(k)$ (although in our study the health parameters only changed once per flight). The noise term $w_x(k)$ represents inaccuracies in the system model and $e(k)$ represents measurement noise.

The states, controls, health parameters and measurements are summarised in Tables 1–4, along with their values at the nominal operating point considered in this article, which is a power lever angle of 21° at sea level static conditions (zero altitude and zero mach). Table 4 also shows typical signal-to-noise ratios for the measurements, based on NASA experience and previously published data (Merrill 1984). Sensor dynamics are assumed to be high enough bandwidth that they can be ignored in the dynamic equations. In Tables 1–4, *LPT* is used for *Low Pressure Turbine*, *HPT* is used for *High Pressure Turbine*, *LPC* is used for *Low Pressure Compressor* and *HPC* is used for *High Pressure Compressor*.

We linearise and augment (19) to obtain the system

$$\begin{aligned} \begin{bmatrix} x(k+1) \\ p(k+1) \end{bmatrix} &= \begin{bmatrix} A_1 & A_2 \\ 0 & I \end{bmatrix} \begin{bmatrix} x(k) \\ p(k) \end{bmatrix} + \begin{bmatrix} B \\ 0 \end{bmatrix} [u(k)] \\ &\quad + \begin{bmatrix} w_x(k) \\ w_p(k) \end{bmatrix} \\ y(k) &= [C_1 \quad C_2] \begin{bmatrix} x(k) \\ p(k) \end{bmatrix} + e(k) \\ &= C \begin{bmatrix} x(k) \\ p(k) \end{bmatrix} + e(k) \end{aligned} \quad (20)$$

Table 1. Modular aero propulsion system simulation turbofan model states and nominal values.

State	Nominal value
LPT rotor speed	7264 RPM
HPT rotor speed	12152 RPM
Average hot section metal temperature	1533°R

A Kalman filter can be applied to these linearised equations to estimate the state vector x and the health parameter vector p . The system matrices are given in Appendix 2.

Constraints can be incorporated in the state estimator by using heuristic knowledge of the behaviour of the health parameters. For example, it is known that health parameters generally do not improve with time. It is also expected that they degrade within a specific envelope.

$$\begin{aligned} p_m(k) &\leq p_m^{\max}(k), \quad m \in [1-10] \\ p_m(k) &\geq p_m^{\min}(k). \end{aligned} \quad (21)$$

This envelope constraint is in the linear form required in the constrained filtering problem statement (4) and is therefore amenable to the approach presented in this article. Note that this technique of constrained estimation does not take into account the possibility of abrupt changes in health parameters due to discrete damage events. That possibility must be addressed by some other means (e.g. residual checking (Doel 1994a)) in conjunction with the methods presented in this article.

4. Simulation results

We simulated the methods discussed in this article using MATLAB[®]. We measured a steady state 3 s burst of engine data at 10 Hz during each flight. These routine data collections were performed over 100 flights at the single operating point shown in Tables 1, 2 and 4, except the engine's health parameters deteriorated a small amount each flight. We simulated a linear-plus-exponential degradation of the 10 health parameters over 100 flights. The simulated health parameter degradation was representative of turbofan performance data reported in the literature (Sasahara 1985). Figure 4 shows the degradations of all 10 health parameters that we used in our simulations.

In the Kalman filter we used a one-sigma state process noise equal to 0.005% of the nominal state values to allow the filter to be responsive to changes in the state variables. We also set the one sigma process noise for each component of the health parameter vector to 0.01% of the nominal parameter value.

Table 2. Modular aero propulsion system simulation turbofan model controls and nominal values.

Control	Nominal value
Main burner fuel flow	2454 lbm hr ⁻¹
Variable nozzle area	343 in ²
Rear bypass door variable area	154 in ²

Table 3. Modular aero propulsion system simulation turbofan model health parameters and nominal values.

Health parameter	Normalized value
Fan airflow	1
Fan efficiency	1
Booster tip airflow	1
Booster tip efficiency ^a	1
Booster hub airflow	1
Booster hub efficiency	1
High pressure turbine airflow	1
High pressure turbine efficiency	1
Low pressure turbine airflow	1
Low pressure turbine efficiency	1

Note: ^aThe fourth health parameter is not yet implemented in MAPSS.

Table 4. Modular aero propulsion system simulation turbofan model measurements, nominal values and signal-to-noise ratios.

Measurement	Nominal value	SNR
LPT exit pressure	19.33 psia	100
LPT exit temperature	1394°R	100
Percent low pressure spool rotor speed	63.47%	150
HPC inlet temperature	580.8°R	100
HPC exit temperature	965.1°R	200
Bypass duct pressure	20.66 psia	100
Fan exit pressure	17.78 psia	200
Booster inlet pressure	20.19 psia	200
HPC exit pressure	85.06 psia	100
Core rotor speed	12152 RPM	150
LPT blade temperature	1179°R	70

Note: SNR values are linear (not in decibels).

These values were obtained by manual tuning. They were small enough to give reasonably smooth estimates, and large enough to allow the filter to track slowly time-varying parameters. The Kalman filter gain was time varying, although it converged to steady state after about 50 flights. The health parameter estimates were updated once per time step (i.e. 30 times

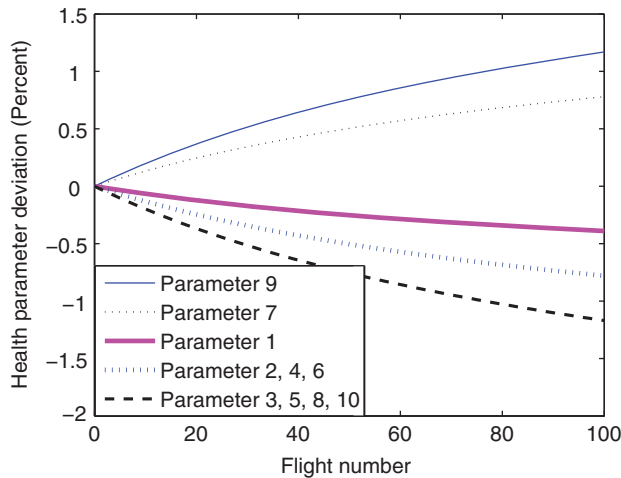


Figure 4. This figure shows the true health parameter deviations that were simulated for each of the 10 health parameters.

per flight). In the enforcement of the constraints in (21) we chose the constraint envelope as follows:

- (1) For the turbine airflow health parameters ($m \in [7, 9]$), whose values *increase* with time (i.e. an increase corresponds to a degradation), $p_m^{\max}(k)$ was set equal to a linear-plus-exponential degradation that was initialised to zero (i.e. $p_m^{\max}(0) = 0$) and reached a maximum of 6% after 500 flights, while $p_m^{\min}(k)$ was set equal to 0 for all k .
- (2) For the other health parameters ($m \in [1-6, 8, 10]$), whose values *decrease* with time (i.e. a decrease corresponds to a degradation), $p_m^{\min}(k)$ was set equal to a linear-plus-exponential degradation that was initialised to zero (i.e. $p_m^{\min}(0) = 0$) and reached a maximum magnitude of -6% after 500 flights, while $p_m^{\max}(k)$ was set equal to 0 for all k .

The constraint envelope was chosen on the basis of domain knowledge of the turbofan engine health estimation problem, and manual tuning. If the constraints are too loose then they do not provide any improvement over unconstrained filtering. If the constraints are too tight then they overly restrict the state estimates and do not make enough of an allowance for random variations in the health parameters.

Figure 5 shows a typical plot of health parameter deviation, along with the constraint envelope, the unconstrained estimate and the constrained estimate. The initial health parameter estimation errors were assumed to be zero. Note that the true health parameter changes once per flight, but the filter estimate is




Figure 5. In this example, constraint enforcement decreases the RMS estimation error from 12.2% to 9.2%.

updated each time step (i.e. 30 times per flight). It is seen that even though the unconstrained estimate lies within the constraint envelope, the constrained estimate is more accurate.

Figure 6 shows a different type of example where the true health parameter deviation is closer to the constraint envelope. In this case there are times when the unconstrained estimate lies outside of the constraint envelope, but the enforcement of constraints forces the constrained estimate to remain within the envelope.

Note that Figures 5 and 6 cannot be compared directly with our previous results in Simon and Simon (2006). First, the results in Simon and Simon (2006) were obtained with different turbine engine simulation software. Second, the estimation error percentages in Simon and Simon (2006) were calculated relative to nominal health parameter values, while the percentages in this article are calculated relative to health parameter degradation magnitudes. For example, suppose that nominal airflow is 100 kg s^{-1} , it degrades by 1 kg s^{-1} , and we estimate it to within 0.1 kg s^{-1} . That estimation accuracy would be reported in Simon and Simon (2006) as a 0.1% estimation error (i.e. $0.1/100$), while it is reported in this article as a 10% estimation error (i.e. $0.1/1$).

We obtained estimates of the health parameters using unconstrained (standard) Kalman filtering, constrained Kalman filtering using the projection approach (Simon and Simon 2005) and constrained Kalman filtering using the PDF truncation approach. We ran 100 Monte Carlo simulations, each simulation consisting of 100 flights and the same health parameter degradation, but different realisations of the

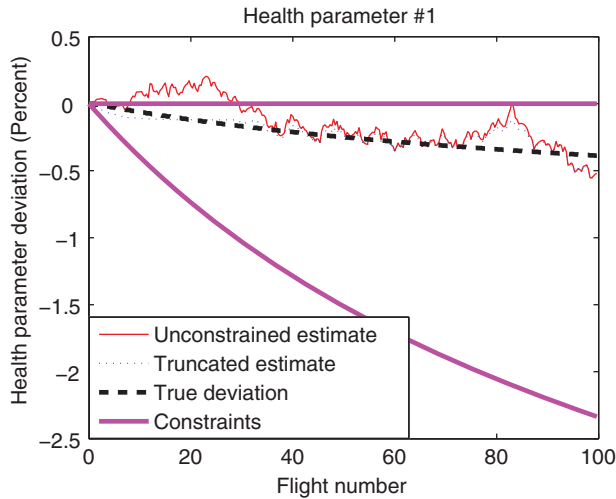


Figure 6. In this example, constraint enforcement decreases the RMS estimation error from 13.4% to 6.6%.

measurement noise. Table 5 shows the performance of the filters averaged over 100 Monte Carlo simulations of 100 flights each. The standard Kalman filter estimates the health parameters to within 7.3% of their final degradations. The projection-based constrained filter estimates the health parameters to within 6.5% of their final degradations. The truncation approach to constrained filtering estimates the parameters to within 6.0% of their final degradations. These numbers show the improvement that is possible with the truncation approach to constrained Kalman filtering.

These examples involve processing 3 s of data each flight, which is similar to standard ground-based estimation based on snap-shot measurements. In this scenario, it could be argued that static estimators (e.g. least squares) would work just as well as a Kalman filter. However, when we consider on-board real-time estimation, a Kalman filter will give better performance than a static estimator. Dynamic estimators can be used to update health estimation in flight, which could then be input to a flight control algorithm.

A comparison between other constrained approaches to filtering for aircraft turbofan engine health estimation can be found in the literature. These methods include the projection approach (which is compared with PDF truncation here) and the switched projection approach (Simon and Simon 2006), the use of soft constraints (Simon and Simon 2006), constrained H_∞ filtering (Simon 2006b) and ridge regression (Dewallef et al. 2004). A general comparison of different methods of constrained Kalman filtering, including the PDF truncation method presented here, can be found in Simon (2008b).

Table 5. Percent RMS health parameter estimation errors of the Kalman filters averaged over 100 Monte Carlo simulations.

Health parameter	Estimation error (%)		
	Unconstrained filter	Projection filter	Truncated filter
Fan airflow	12.6	8.7	7.5
Fan efficiency	7.1	6.4	6.5
Booster tip airflow	11.3	11.1	9.3
Booster tip efficiency ^a	N/A	N/A	N/A
Booster hub airflow	7.2	6.6	7.0
Booster hub efficiency	3.9	3.2	3.8
High pressure turbine airflow	4.2	3.2	3.9
High pressure turbine efficiency	4.3	4.0	3.9
Low pressure turbine airflow	3.8	3.5	3.5
Low pressure turbine efficiency	11.4	11.3	8.6
Average	7.3	6.5	6.0

Note: The estimation error is measured as $|(p - \hat{p})/p_f|$, where p is the true health parameter value, \hat{p} is the estimated health parameter value, and p_f is the health parameter value at the end of the simulation.

^aThe fourth health parameter is not yet implemented in MAPSS.

4.1. Determination of constraints

The $p_m^{\max}(k)$ and $p_m^{\min}(k)$ constraints are system dependent and were obtained in this example using domain knowledge and tuning. Constraints that are too loose or too tight will result in worse performance as shown in Table 5. The results in Table 5 were obtained with constraints of 0% and 6%, as discussed at the beginning of this section. If the constraints are relaxed to -3% and 9%, the average constrained performance degrades to 7.3%, which is the same as the unconstrained filter, showing that loose constraints do not provide any improvement relative to unconstrained filtering. If the constraints are tightened to 0% and 2%, the average constrained performance degrades to 10.6%, which is worse than the unconstrained filter. Constraint tuning affects the performance of both the projection approach and the PDF truncation approach to constrained filtering. However, for all reasonable constraint envelopes that we investigated, the PDF truncation approach outperformed projection by an amount similar to that seen in Table 5.

4.2. Estimation of slower health degradations

The robustness of constrained Kalman filtering to varying rates of health parameter degradation is suggested by further simulation results. The

simulations discussed above were repeated except that the health parameters degraded only 10% as fast as shown in Figure 4 but over 1000 flight cycles. In this case the standard unconstrained Kalman filter estimated the health parameters with an average RMS error of 6.7%, while the PDF truncation filter estimated the health parameters with an average RMS error of 5.0% (averaged over 100 Monte Carlo simulations). Comparing these results with Table 5 shows that both standard unconstrained Kalman filtering and PDF truncation filtering perform better when the health parameter degradation is slower.

4.3. Computational effort

The improved performance of the constrained filter comes with a price, and that price is computational effort. The algorithm outlined in (6)–(17) requires Jordan decomposition and Gram–Schmidt orthogonalisation. However, if the constraints of (4) are decoupled (as they are in our example) then the computational effort can be largely reduced by ignoring the cross-covariance terms in the state estimator and hence avoiding these matrix computations. In any case, computational effort is not a critical issue for turbofan health estimation since the filtering is performed on ground-based computers after each flight.

In this example, the projection approach to constrained filtering requires a negligible amount of additional computational effort relative to unconstrained filtering. This is because the constraints are decoupled. However, the PDF truncation approach discussed in this article requires a noticeable amount of additional computational effort. This is primarily because of the error function and exponential calculations in (14) and (15). (If the constraints were coupled, then the computational effort would be primarily due to the Jordan decomposition and Gram–Schmidt orthogonalisation.) For the 100-flight simulations performed for this article on a 1.5 GHz PC with 1 GB of RAM, unconstrained Kalman filtering required 1.8 s of CPU time and the PDF truncation approach required 5.6 s of CPU time. Even with its increased complexity, CPU time will probably not be a primary concern for the truncation approach to constrained filtering. For example, extrapolating the CPU times obtained above, we see that processing data for 1000 flights of 1000 engines would require 16 h of CPU time using the truncation approach to constrained filtering. Also, since the Kalman filter is applied recursively, the health parameter estimates are updated with only a small number of new data points each flight (30 measurements per flight in this article).

However, a relinearisation process should be performed every few flights in order to obtain the best possible estimation performance (Simon 2008a). If CPU time is a consideration, fast table lookups for the error function and exponential calculations in (14) and (15) could be implemented.

4.4. Abrupt faults

The Kalman filter works well only if the assumed system model matches reality fairly closely. This implies that the system model should be relinearised periodically in order to maximise estimation performances (Simon 2008a). The method presented in this article will not work well by itself if there are large sensor biases or hard faults due to severe component failures. A mission-critical implementation of a Kalman filter should always include some sort of additional residual check to verify the validity of the Kalman filter results (Gelb 1974), particularly for the application of turbofan engine health estimation considered in this article (Doel 1994a).

5. Conclusion

We have presented a PDF truncation method for incorporating constraints into a Kalman filter. If the system whose state variables are being estimated has known state variable constraints, then those constraints can be incorporated into the Kalman filter as shown in this article. For the aircraft turbofan engine health estimation problem, the use of constraints generally improves the accuracy of health estimation.

We have seen that the constrained filter requires more computational effort than the standard Kalman filter. This is due to the addition of s matrix decompositions and associated computations that must be performed at each time step (one for each constraint). The engineer must therefore perform a tradeoff between computational effort and estimation accuracy.

The Kalman filter works well only if the assumed system model matches reality fairly closely. The constraint enforcement method presented in this article will not work well if there are large sensor biases or hard faults due to severe component failures. A mission-critical implementation of a Kalman filter should always include some sort of residual check to verify the validity of the Kalman filter results, particularly for the application of turbofan engine health estimation considered in this article.

The method presented in this article changes the unconstrained estimate even if the unconstrained

estimate satisfies the constraints. An improvement might be seen if PDF truncation is performed only for those estimates that violate the constraints. The desired approach depends on if the engineer is more confident in the unconstrained state estimate or more confident in the constraint information.

Acknowledgement

This work was supported by the NASA Aviation Safety and Security Program at NASA Glenn Research Center.

Notes on contributors



Dan Simon received his BS degree from Arizona State University (1982), an MS degree from the University of Washington (1987), and a PhD degree from Syracuse University (1991), all in electrical engineering. He worked in industry for 14 years at Boeing, TRW, and several small companies. His industrial experience includes work in the aerospace, automotive, agricultural, GPS, biomedical, process control and software fields. In 1999 he moved from industry to academia, where he is now a professor in the Electrical and Computer Engineering Department at Cleveland State University. His teaching and research involves embedded systems, control systems and computer intelligence. He has published over 60 refereed conference and journal papers, and is the author of the text 'Optimal State Estimation' (John Wiley & Sons, 2006).



Donald L. Simon received his BS degree from Youngstown State University (1987) and an MS degree from Cleveland State University (1990), both in electrical engineering. During his career as an employee of the US Army Research Laboratory (1987–2007), and the NASA Glenn Research Center (from 2007 to present), he has focused on the development of advanced control and health management technologies for current and future aerospace propulsion systems. His specific research interests are in aircraft gas turbine engine performance diagnostics and performance estimation. He currently leads the propulsion gas path health management research effort ongoing under the NASA Aviation Safety Program, Integrated Vehicle Health Management Project.

References

DePold, H., and Gass, F. (1999), 'The Application of Expert Systems and Neural Networks to Gas Turbine Prognostics and Diagnostics', *ASME Journal of Engineering for Gas Turbines and Power*, 121, 607–612.

- Dewallef, P., Leonard, O., and Mathioudakis, K. (2004), 'On-line Aircraft Engine Diagnostic using a Soft-constrained Kalman Filter', *ASME Turbo Expo 2004*, Vienna, Austria. Paper GT2004-53539.
- Doel, D. (1994a), 'TEMPER – A Gas-path Analysis Tool for Commercial Jet Engines', *ASME Journal of Engineering for Gas Turbines and Power*, 116, 82–89.
- Doel, D. (1994b), 'An Assessment of Weighted-least-squares-based Gas Path Analysis', *ASME Journal of Engineering for Gas Turbines and Power*, 116, 366–373.
- Friedland, B. (1969), 'Treatment of Bias in Recursive Filtering', *IEEE Transactions on Automatic Control*, AC14, 359–367.
- Ganguli, R. (2003), 'Application of Fuzzy Logic for Fault Isolation of Jet Engines', *ASME Journal of Engineering for Gas Turbines and Power*, 125, 617–623.
- Ganguli, R., and Dan, B. (2004), 'Trend Shift Detection in Jet Engine Gas Path Measurements using Cascaded Recursive Median Filter with Gradient and Laplacian Edge Detector', *ASME Journal of Engineering for Gas Turbines and Power*, 126, 55–61.
- Gelb, A. (1974), *Applied Optimal Estimation*, Cambridge, MA: MIT Press.
- Goodwin, G., De Dona, J., Seron, M., and Zhuo, X. (2005), 'Lagrangian Duality Between Constrained Estimation and Control', *Automatica*, 41, 935–944.
- Kobayashi, T., and Simon, D.L. (2001), 'A Hybrid Neural Network-genetic Algorithm Technique for Aircraft Engine Performance Diagnostics', in *37th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, pp. 2822–2827.
- Lambert, H. (1991), 'A Simulation Study of Turbofan Engine Deterioration Estimation using Kalman Filtering Techniques', NASA Technical Memorandum 104233.
- Massicotte, D., Morawski, R., and Barwicz, A. (1995), 'Incorporation of a Positivity Constraint into a Kalman-filter-based Algorithm for Correction of Spectrometric Data', *IEEE Transactions on Instrumentation and Measurement*, 44, 2–7.
- Merrill, W. (1984), 'Identification of Multivariable High-performance Turbofan Engine Dynamics from Closed-loop Data', *AIAA Journal of Guidance, Control, and Dynamics*, 7, 677–683.
- Moon, T., and Stirling, W. (2000), *Mathematical Methods and Algorithms for Signal Processing*, New Jersey: Prentice Hall, Upper Saddle River.
- Parker, K., and Guo, T. (2003), 'Development of a Turbofan Engine Simulation in a Graphical Simulation Environment', NASA TM 2003-212543.
- Parker, K., and Melcher, K. (2004), *The Modular Aero-propulsion Systems Simulation (MAPSS) Users' Guide*, NASA TM 2004-212968.
- Rao, C., Rawlings, J., and Mayne, D. (2003), 'Constrained State Estimation for Nonlinear Discrete-time Systems: Stability and Moving Horizon Approximations', *IEEE Transactions on Automatic Control*, 48, 246–258.
- Sasahara, O. (1985), 'JT9D Engine/Module Performance Deterioration Results from Back to Back Testing',

- International Symposium on Air Breathing Engines*, 528–535.
- Shimada, N., Shirai, Y., Kuno, Y., and Miura, J. (1998), ‘Hand Gesture Estimation and Model Refinement using Monocular Camera – Ambiguity Limitation by Inequality Constraints’, in *Third IEEE International Conference on Automatic Face and Gesture Recognition*, Nara, Japan, pp. 268–273.
- Simon, D. (2006a), *Optimal State Estimation*, New York: John Wiley & Sons.
- Simon, D. (2006b), ‘H-infinity Filtering with Inequality Constraints for Aircraft Turbofan Engine Health Estimation’, in *IEEE Conference on Decision and Control*, San Diego, pp. 3291–3296.
- Simon, D. (2008a), ‘A Comparison of Filtering Approaches for Aircraft Engine Health Estimation’, *Aerospace Science and Technology*, 12, 276–284.
- Simon, D. (2009), ‘Kalman Filtering with State Constraints: A Survey of Linear and Nonlinear Algorithms’, *IET Control Theory & Applications*, accepted for publication.
- Simon, D., and Chia, T. (2002), ‘Kalman Filtering with State Equality Constraints’, *IEEE Transactions on Aerospace and Electronic Systems*, 39, 128–136.
- Simon, D., and Simon, D.L. (2005), ‘Aircraft Turbofan Engine Health Estimation using Constrained Kalman Filtering’, *ASME Journal of Engineering for Gas Turbines and Power*, 127, 323–328.
- Simon, D., and Simon, D.L. (2006), ‘Kalman Filter Constraint Switching for Turbofan Engine Health Estimation’, *European Journal of Control*, 12, 331–343.
- Simon, D., and Simon, D.L. (2006), ‘Kalman Filtering with Inequality Constraints for Turbofan Engine Health Estimation’, *IEE Proceedings Control Theory & Applications*, 153, 371–378.
- Teixeira, B., Chandrasekar, J., Torres, L., Aguirre, L., and Bernstein, D. (2009), ‘State Estimation for Linear and Non-linear Equality-constrained Systems’, *International Journal of Control*, 82, 918–936.
- Vathsal, S., and Sarma, I. (1974), ‘Constrained Linear Minimax Filter for Systems with Large Plant Uncertainties’, *International Journal of Systems Science*, 5, 479–492.
- Volponi, A., DePold, H., Ganguli, R., and Daguang, C. (2003), ‘The use of Kalman Filter and Neural Network Methodologies in Gas Turbine Performance Diagnostics: A Comparative Study’, *ASME Journal of Engineering for Gas Turbines and Power*, 125, 917–924.

- Wang, Z., and Shu, H. (2000), ‘Variance-constrained Multiobjective Filtering for Uncertain Continuous-time Stochastic Systems’, *International Journal of Systems Science*, 31, 1175–1183.

Appendix 1

This appendix contains a proof of Lemma 1. First we prove the covariance part of the lemma. Note from (6) to (7) that

$$\begin{aligned} E[(z_i(k) - \bar{z}_i(k))(z_i(k) - \bar{z}_i(k))^T] &= SW^{-1/2}T^T\tilde{\Sigma}_i(k)TW^{-1/2}S^T \\ &= SW^{-1/2}T^T(TWT^T)T^{-1/2}S^T \\ &= I. \end{aligned} \quad (22)$$

We have proved that the covariance of $z_i(k)$ is the identity matrix.

Next consider the mean of $z_i(k)$. When $i=0$, the mean of $z_i(k)$ is given as

$$\begin{aligned} E(z_0(k)) &= SW^{-1/2}T^TE(x(k) - \tilde{x}_0(k)) \\ &= SW^{-1/2}T^TE(x(k) - \hat{x}(k)) \\ &= 0. \end{aligned} \quad (23)$$

Now we use induction to show that only the first component of the mean of $z_i(k)$ can be non-zero. For $i > 0$, suppose that $\tilde{z}_i(k) = [\mu \ 0 \ \dots \ 0]^T$ and $E(z_{i-1}(k)) = [\nu \ 0 \ \dots \ 0]^T$. Then the mean of $z_i(k)$ is given as

$$\begin{aligned} E(z_i(k)) &= SW^{-1/2}T^TE(x(k) - \tilde{x}_i(k)) \\ &= SW^{-1/2}T^TE(x(k) - TW^{1/2}S^T\tilde{z}_i(k) - \tilde{x}_{i-1}(k)) \\ &= \tilde{z}_i(k) + SW^{-1/2}T^TE(x - \tilde{x}_{i-1}(k)) \\ &= [\mu \ 0 \ \dots \ 0]^T + E(z_{i-1}(k)) \\ &= [\mu \ 0 \ \dots \ 0]^T + [\nu \ 0 \ \dots \ 0]^T. \end{aligned} \quad (24)$$

We have proved that only the first element of $z_i(k)$ can have a non-zero mean.

Appendix 2

This appendix gives the matrices that are used in (20).

$$A_1 = \begin{bmatrix} 0.9029 & 0.0411 & 0.0381 \\ -0.0069 & 0.9088 & 0.0432 \\ -0.0001 & -0.0004 & 0.9924 \end{bmatrix}$$

$$A_2 = \begin{bmatrix} -294.2884 & 254.1007 & -8.8834 & 0 & 241.4922 & -47.2503 & 10.0174 & -19.6286 & -365.4922 & 116.1297 \\ -17.6391 & -38.4735 & -131.9580 & 0 & -274.2797 & 449.5393 & -61.9914 & 302.5021 & 126.9419 & 35.4716 \\ 0.1598 & -1.2844 & 0.0262 & 0 & -2.3248 & -3.3278 & -1.4288 & -0.1425 & -0.0200 & -0.0275 \end{bmatrix}$$

$$B = \begin{bmatrix} 0.0805 & 0.4928 & -0.1557 \\ 0.0910 & 0.1678 & -0.0341 \\ 0.0018 & -0.0003 & -0.0001 \end{bmatrix}, \quad C_1 = \begin{bmatrix} -0.0034 & 1 & 0.0237 \\ 0.0087 & 0.0002 & 0.0002 \\ 0.0016 & -0.0006 & 0.0001 \\ 0.0022 & -0.0005 & 0.0001 \\ 0.0181 & -0.0024 & 0.0008 \\ 0.0148 & 0.0493 & 0.0094 \\ 0.0018 & 0.0000 & 0.0002 \\ 0.0030 & 0.0127 & 0.0048 \\ -0.0012 & -0.0302 & 0.0656 \\ -0.0172 & -0.1098 & 0.1218 \\ 0.0010 & 0.0007 & 0.0004 \end{bmatrix}$$

$$C_2 = \begin{bmatrix} -2.5330 & -21.8759 & -64.9562 & 0.0000 & -161.0030 & 240.9109 & -34.5505 & 160.8764 & 66.6444 & 18.8739 \\ -1.3496 & 1.1914 & 0.0113 & 0.0000 & 1.1016 & -0.2965 & 0.0416 & -0.1417 & -1.7128 & 0.5402 \\ 2.2717 & -0.2729 & -1.4787 & 0.0000 & -1.9038 & -0.5124 & -0.0734 & -0.3334 & -0.2629 & 0.0053 \\ 5.6643 & -0.3723 & -1.4596 & 0.0000 & -1.6750 & -0.4341 & -0.0741 & -0.2731 & -0.4132 & 0.0709 \\ 30.3031 & -62.3520 & -6.3313 & 0.0000 & -7.7714 & -2.6506 & -0.3207 & -1.6537 & -3.3329 & 0.7082 \\ 27.3703 & -65.9505 & -8.5258 & 0.0000 & 128.5380 & -391.3869 & -161.7529 & 6.3074 & 0.5187 & 1.4964 \\ 5.4493 & -0.0209 & 1.9711 & 0.0000 & -0.4037 & -0.4561 & 0.0269 & -0.3159 & -0.1625 & -0.0255 \\ 1.4155 & 10.5671 & -1.0475 & 0.0000 & 52.2323 & -2.7962 & -93.1014 & 1.9595 & 0.1496 & 0.4219 \\ 31.0947 & -124.3117 & 3.6502 & 0.0000 & -201.6218 & -340.3074 & -76.2822 & -101.8336 & 52.7272 & -43.3964 \\ 34.8192 & -182.6730 & 15.8262 & 0.0000 & -531.7816 & -289.2280 & 9.1885 & -209.9746 & 104.9357 & -88.2893 \\ 2.7499 & 0.4834 & 0.8791 & 0.0000 & 2.3038 & -0.5266 & 0.2287 & -0.4034 & 0.1298 & -0.1472 \end{bmatrix}.$$

(25)

Downloaded At: 00:55 20 January 2010