

Kalman Filter Constraint Switching for Turbofan Engine Health Estimation

Dan Simon *
Cleveland State University
Electrical Engineering
Stilwell Hall Room 332
2121 Euclid Avenue
Cleveland, Ohio, 44115

Donald L. Simon
US Army Research Laboratory
NASA Glenn Research Center
Mail Stop 77-1
21000 Brookpark Road
Cleveland, Ohio, 44135

Abstract

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. Recently published work has shown a new method for incorporating state variable inequality constraints in the Kalman filter. The resultant filter is a combination of a standard Kalman filter and a quadratic programming problem. The incorporation of state variable constraints has been shown to generally improve the filter's estimation accuracy. However, the incorporation of inequality constraints poses some risk to the estimation accuracy. After all, the Kalman filter is theoretically optimal, so the incorporation of heuristic constraints may degrade the optimality of the filter. This paper proposes a way to switch the filter constraints so that the state estimates follow the unconstrained (theoretically optimal) filter when the confidence in the unconstrained filter is high.

*Corresponding author. Email address d.j.simon@csuohio.edu. This work was supported by the NASA Aviation Safety and Security Program at NASA Glenn Research Center.

When confidence in the unconstrained filter is not so high, then we use our heuristic knowledge to constrain the state estimates. The confidence measure is based on the agreement of measurement residuals with their theoretical values. If some measurement residuals are low, and those residuals are highly sensitive to a given state, then we are confident that the unconstrained estimate of that state is correct. Otherwise, we incorporate our heuristic knowledge as state constraints. The algorithm is demonstrated on a linearized simulation of a turbofan engine to estimate engine health.

Key Words – Kalman Filter, State Constraints, Estimation, Quadratic Programming, Residuals, Gas Turbine Engines.

1 Introduction

For linear dynamic systems with white process and measurement noise, the Kalman filter is known to be 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 [14]. Previous work by the authors [20] resulted in a new method for incorporating state variable inequality constraints in the Kalman filter. This method is based on a generalization of [19], which dealt with the incorporation of state variable equality constraints in the Kalman filter. Constraints are enforced by projecting out-of-bound state estimates onto the constraint surface. Inequality constraints are inherently more complicated than equality constraints, but standard quadratic programming techniques can be used to solve the Kalman filter problem with inequality constraints. At each time step of the constrained Kalman filter, we solve a quadratic programming problem to obtain the constrained state estimate. It was shown earlier [20] that the constrained estimate has several important properties. For example, the constrained state estimate is unbiased and has a smaller error covariance than the unconstrained estimate. Also, the constrained estimate is always (i.e., at each time step) closer to the true state than the unconstrained estimate. The incorporation of state variable constraints was shown to improve the filter's estimation accuracy for turbofan health estimation.

However, these properties of the constrained filter hold true only if the state constraints that are enforced are correct. In practice, state constraints are often based on heuristic knowledge – that is, the constraints are more correctly viewed as “soft” constraints. The use of inequality constraints therefore poses some risk to the estimation accuracy. The Kalman filter is theoretically optimal, so the incorporation of heuristic constraints is a modification to the optimal filter. We want to be able to incorporate our heuristic knowledge into the filter, but we do not have absolute confidence in our heuristic knowledge.

The constrained filter is theoretically superior to the unconstrained filter, but only if the constraints are accurate. The incorporation of constraints is not always exact, and some judgment must be used in their definition. This paper proposes a way to switch the constraints so that the state estimate is equal to the unconstrained (theoretically optimal) estimate when the confidence in the unconstrained estimate is high. When confidence in the unconstrained filter is not so high, we use our heuristic knowledge to constrain the state estimates. The confidence measure is based on the agreement of measurement residuals with their theoretical values. If some measurement residuals are low, and the measurements corresponding to those residuals are highly sensitive to a given state, then we are confident that the unconstrained estimate of that state is correct. Otherwise, we incorporate our heuristic knowledge as state constraints.

The application considered in this paper is aircraft turbofan engine health parameter estimation [5]. Health parameters represent engine component efficiencies and flow capacities. The performance of a gas turbine engine deteriorates over time. This deterioration reduces the fuel economy of the engine. Airlines periodically collect engine data in order to evaluate the health of the engine and its components. The health evaluation is then used to determine maintenance schedules. Reliable health evaluations are used to anticipate future maintenance needs. 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 col-

lected each flight at the same engine operating points and corrected to account for variability in ambient conditions. Various algorithms have been proposed to monitor engine health, such as weighted least squares [6], expert systems [3], Kalman filters [23], neural networks [23], and genetic algorithms [12].

This paper applies constrained Kalman filtering, along with constraint switching on the basis of measurement residuals, to estimate engine health parameters. We use heuristic knowledge of the health parameter dynamics to constrain their estimate. For example, we know that health parameters never improve. Engine health always 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 emphasized that in this paper 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.

This paper is organized as follows. Section 2 presents a review of the constrained Kalman filter, along with a proposed method for how the residuals can be used for constraint switching. Section 2 also shows how a matrix quantifying the sensitivity of measurements to state variables can be obtained, and how the entries of that matrix can be used to quantify our confidence in the accuracy of the unconstrained Kalman filter estimates. Our confidence can then be used to decide whether or not to enforce heuristic constraints on the state variable estimates. 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 linearized dynamic model is augmented in such a way that a Kalman filter can estimate the health parameters following a previously published approach [8, 13]. We then show how this problem can be expressed in a way that is compatible with the constraints discussed in Section 2. Section 4 discusses the application of the sensitivity analysis and Kalman filter constraint switching technique to the turbofan engine health parameter estimation problem.

Section 5 presents some simulation results based on a turbofan model linearized around a known operating point. We show that the constrained Kalman filter can estimate health parameters better than the unconstrained filter, and the addition of constraint switching further improves estimation accuracy. Section 6 presents some concluding remarks and suggestions for further work.

2 Kalman Filtering with Constraint Switching

In this section we first summarize the standard Kalman filter equations. We then review constrained state estimation via the Kalman filter, and propose a method for residual-based constraint switching.

Consider the discrete linear time-invariant system given by

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

where k is the time index, x is the state vector, u is the known control input, y is the measurement, and $\{w(k)\}$ and $\{e(k)\}$ are noise input sequences. The problem is to find an estimate $\hat{x}(k+1)$ of $x(k+1)$ given the measurements $\{y(0), y(1), \dots, y(k)\}$. We will use the symbol $Y(k)$ to denote the column vector that contains the measurements $\{y(0), y(1), \dots, y(k)\}$. We assume that the following standard conditions are satisfied.

$$\begin{aligned} E[x(0)] &= \bar{x}(0) \\ E[w(k)] = E[e(k)] &= 0 \\ E[(x(0) - \bar{x}(0))(x(0) - \bar{x}(0))^T] &= \Sigma(0) \\ E[w(k)w^T(m)] &= Q\delta_{km} \\ E[e(k)e^T(m)] &= R\delta_{km} \\ E[w(k)e^T(m)] &= 0 \end{aligned} \tag{2}$$

where $E[\cdot]$ is the expectation operator, \bar{x} is the expected value of x , and δ_{km} is the Kronecker delta function ($\delta_{km} = 1$ if $k = m$, $\delta_{km} = 0$ otherwise). Q and R are

positive semidefinite covariance matrices. The Kalman filter equations are given by

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

where the filter is initialized with $\hat{x}(0) = \bar{x}(0)$, and $\Sigma(0)$ given above. It can be shown [1] that the Kalman filter has several attractive properties. For instance, if $x(0)$, $\{w(k)\}$, and $\{e(k)\}$ are jointly Gaussian, the Kalman filter estimate $\hat{x}(k+1)$ is the conditional mean of $x(k+1)$ given the measurements $Y(k)$, i.e., $\hat{x}(k+1) = E[x(k+1)|Y(k)]$. Even if $x(0)$, $\{w(k)\}$, and $\{e(k)\}$ are not jointly Gaussian, the Kalman filter estimate is the best *linear* estimator given the measurements $Y(k)$, i.e., of all estimates of $x(k+1)$ that are of the form $FY(k) + g$ (where F is a fixed matrix and g is a fixed vector), the Kalman filter estimate is the one that minimizes the variance of the estimation error. Also, the Kalman filter estimate $\hat{x}(k)$ is that value of ζ that maximizes the conditional probability density function $P(\zeta|Y(k))$. Finally, $\Sigma(k)$ is the covariance of the Kalman filter estimation error at time k .

2.1 Constrained Kalman Filtering

Now consider the system of (1) where we are given the additional constraint

$$D(k)x(k) \leq d(k) \tag{4}$$

where $D(k)$ is a known $s \times n$ matrix, s is the number of constraints, n is the number of state variables, and $s \leq n$. It is assumed in this paper that $D(k)$ is full rank, i.e., that $D(k)$ has rank s . This is an easily satisfied assumption. If $D(k)$ is not full rank that means we have redundant state constraints. In that case we can simply remove linearly dependent rows from $D(k)$ (i.e., remove redundant state constraints) until $D(k)$ is full rank. The time index k is omitted in the remainder of this section for ease of notation.

The problem of finding a constrained estimate for the state of the system (1) can be posed in three different ways [20]. Regardless of how we pose the problem,

we want to make sure that our constrained estimate \tilde{x} satisfies the constraint (4). That is,

$$D\tilde{x} \leq d \tag{5}$$

The solution to the constrained estimation problem turns out to be the solution to

$$\min_{\tilde{x}} (\tilde{x} - \hat{x})^T W (\tilde{x} - \hat{x}) \text{ such that } D\tilde{x} \leq d \tag{6}$$

where \hat{x} is the unconstrained (standard) Kalman filter estimate, and W is a symmetric positive definite weighting matrix. Note that if the unconstrained estimate satisfies the constraint, then the solution of the above equation is simply $\tilde{x} = \hat{x}$. That is, if the standard Kalman filter estimate satisfies the constraints, then the constrained estimate is equal to the unconstrained estimate.

Depending on the particular optimality criterion that is employed, W can take on several different values [20]. If a mean square error criterion is used then $W = I$. If a maximum probability criterion is used then $W = \Sigma^{-1}(k)$. If a projection method is used then W is an arbitrary positive definite matrix. The optimality of the constrained estimate does not depend on the conditional Gaussian nature of \hat{x} , i.e., $x(0)$, $\{w(k)\}$, and $\{e(k)\}$ in (1) are not assumed to be Gaussian.

The constrained estimate \tilde{x} has several attractive properties.

1. The solution \tilde{x} of the constrained state estimation problem given by (6) is an unbiased state estimator for the system (1) for any symmetric positive definite weighting matrix W .
2. The solution \tilde{x} of the constrained state estimation problem given by (6) with $W = \Sigma^{-1}$, where Σ is the covariance of the unconstrained estimate given in (3), has an error covariance that is less than or equal to that of the unconstrained state estimate.
3. Among all the constrained Kalman filters resulting from the solution of (6), the filter that uses $W = \Sigma^{-1}$ has the smallest estimation error covariance.
4. The solution \tilde{x} of the constrained state estimation problem given by (6) with $W = I$ satisfies the inequality

$$\|x(k) - \tilde{x}(k)\| \leq \|x(k) - \hat{x}(k)\| \text{ for all } k \tag{7}$$

where $\|\cdot\|$ is the vector two-norm.

The above properties all follow from the proofs presented in [20, 19].

2.2 Constraint Switching

Many times the constraints of (4) are more heuristic than exact. We have some confidence in the constraints, but we also have some confidence in the unconstrained Kalman filter estimates. We therefore need to somehow moderate our enforcement of the constraints.

In this subsection we analyze the sensitivity of the measurements to the states. We then propose using this information to decide if an unconstrained state variable estimate is reliable. We examine residuals that correspond to measurements that are highly sensitive to a given state. If those residuals are small, then we have a high confidence in the estimate of that state, and we relax the constraints. However, if those residuals are large, then we have a low confidence in the estimate of that state, and we enforce constraints. This process can be conceptually summarized as follows.

1. Run the standard, unconstrained Kalman filter.
2. Check the measurement residuals. If they agree with theory, then the unconstrained state estimate is reasonable and we have a good state estimate.
3. If the measurement residuals do not agree with theory, then use the constrained Kalman filter to obtain the state estimate.

In this rest of this section we will develop this idea in more detail.

Recall our system equations from (1).

$$\begin{aligned}x(k+1) &= Ax(k) + Bu(k) + w(k) \\ y(k) &= Cx(k) + e(k)\end{aligned}\tag{8}$$

We see that C can be interpreted as the sensitivity matrix of the measurements to the states. The element C_{ij} gives the sensitivity of the i th measurement to the j th

state. In practice we should normalize C by dividing each row by the corresponding measurement value. This gives a normalized sensitivity matrix $\Delta(k)$ as follows.

$$\Delta(k) = \begin{bmatrix} 1/y_1(k) & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & 1/y_q(k) \end{bmatrix} C \quad (9)$$

where q is the number of measurements. During the execution of the Kalman filter, the measurement residuals are given by

$$\nu(k) = y(k) - C\hat{x}(k) \quad (10)$$

The theoretical mean and covariance of the residuals are given as [1, 9]

$$\begin{aligned} E[\nu(k)] &= 0 \\ S(k) &= E[\nu(k)\nu^T(k)] \\ &= C\Sigma(k)C^T + R \end{aligned} \quad (11)$$

Therefore, if the measurement residuals satisfy their theoretical statistical properties, we can have confidence that the state estimates are reliable.

Residual based constraint switching proceeds as follows. We generate a list of the measurements that are most sensitive to each state. This can be obtained by sorting each column of the sensitivity matrix $\Delta(k)$ in descending order. Use the notation M_{ji} to denote the measurement number that has the j th largest sensitivity to the i th state. That is, M_{ji} is the row number in the i th column of $\Delta(k)$ that has the j th largest magnitude.

Now we take the first μ rows of the M matrix, where μ is a user defined threshold. This tells us the μ measurements that are most sensitive to each state. Looking at the first μ rows of each column of M , we see that if the residuals corresponding to those measurements are small, then we can have a high confidence in our unconstrained estimate of the state that corresponds to that column of M .

Notice that a second approach could also be taken to determining our confidence in the state estimates. For example, instead of seeing which residuals are most sensitive to the first state, we could see which states have the most effect on the first residual. The question of which of these two approaches to take remains an open

issue. In this paper we took the first approach, which consists of checking which residuals were the most sensitive to each state, one state at a time. This seems to be a more natural method (from an algorithmic point of view) since we can accomplish constraint switching one state at a time. The constraint switching algorithm can be summarized as follows.

1. We are given the following system with n states, q measurements, and s constraints.

$$\begin{aligned} x(k+1) &= Ax(k) + Bu(k) + w(k) \\ y(k) &= Cx(k) + e(k) \\ D(k)x(k) &\leq d(k) \end{aligned} \tag{12}$$

We initialize the Kalman filter quantities $\hat{x}(0)$, $\tilde{x}(0)$, and $\Sigma(0)$.

2. At each time step $k = 0, 1, \dots$, perform the following.

- (a) Run the unconstrained and constrained Kalman filters as follows.

$$\begin{aligned} K(k) &= A\Sigma(k)C^T(C\Sigma(k)C^T + R)^{-1} \\ \hat{x}(k+1) &= A\hat{x}(k) + Bu(k) + K(k)(y(k) - C\hat{x}(k)) \\ \Sigma(k+1) &= (A\Sigma(k) - K(k)C\Sigma(k))A^T + Q \\ \min_{\tilde{x}(k+1)} & [\tilde{x}(k+1) - \hat{x}(k+1)]^T W(k+1) [\tilde{x}(k+1) - \hat{x}(k+1)] \\ &\text{such that } D(k+1)\tilde{x}(k+1) \leq d(k+1) \end{aligned} \tag{13}$$

where $W(k)$ is our weighting matrix (see Section 2.1). This gives us an unconstrained estimate $\hat{x}(k+1)$ and a constrained estimate $\tilde{x}(k+1)$.

- (b) Compute the theoretical residual covariance $S(k+1)$ from (11).
- (c) For $i = 1, \dots, n$, perform the following.
 - i. Find the rows with the μ largest magnitudes in the i th column of the $\Delta(k)$ matrix. Label these row numbers M_{ji} ($j = 1, \dots, \mu$).
 - ii. Examine the μ residuals that correspond to measurement numbers M_{ji} ($j = 1, \dots, \mu$). If all μ of these residuals have been smaller than

$\alpha S_{rr}(k+1)$ (where $r = M_{ji}$) for κ consecutive time steps, then use $\hat{x}_i(k+1)$ as the estimate of the i th state. Otherwise, use $\tilde{x}_i(k+1)$ as the estimate of the i th state

In practice, the decision of how many residuals to use for each state variable (the value of μ), what relative threshold values to use for those residuals (the value of α), and how long a residual must remain “small” before we trust the unconstrained estimate (the value of κ) are open questions. In addition, instead of using only the diagonal elements of $S(k+1)$, we could algebraically transform the residuals in such a way that their covariance becomes a diagonal matrix, and correspondingly modify step 2(c)ii in the above algorithm. This might result in more intelligent switching between the unconstrained and constrained state estimates. In spite of these open issues, the theory presented in this section gives a general approach for deciding when to relax constraints and when to enforce constraints.

As we switch between the unconstrained and constrained estimates, a question arises regarding the properties of this switched filter as compared to the properties of the constrained filter that were described in Section 2.1. First we note that the first property of the constrained filter carries over to the switched filter. That is, since the unconstrained and constrained filters are both unbiased, the switched filter is also unbiased. The last three properties, however, do not in general carry over to the switched filter. In theory we would like those properties to apply to the switched filter. However if our constraints are heuristic rather than rigorous, we cannot guarantee that those properties hold even for the constrained filter, much less for the switched filter. The philosophy of the switched filter proposed here is to examine residuals in order to see at each time step whether or not our heuristic constraints are believable. If the residuals indicate that the heuristic constraints are believable, then we enforce the constraints. If the heuristic constraints in addition reflect mathematical reality, then all four properties from Section 2.1 will carry over to the switched filter.

3 Turbofan Engine Health Monitoring

Figure 1 shows a schematic representation of a turbofan engine [16]. 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.




Figure 1: Schematic representation of a turbofan engine.

The simulation used in this paper is a gas turbine engine simulation software package called MAPSS (Modular Aero Propulsion System Simulation) [16]. 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. 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, high-pressure rotor speed, and the

average hot section metal temperature (measured from aft of the combustor to the high pressure turbine).

The discretized time invariant equations that model the turbofan engine can be summarized as follows.

$$\begin{aligned}
 x(k+1) &= f[x(k), u(k), p(k)] + w_x(k) \\
 p(k+1) &= p(k) + w_p(k) \\
 y(k) &= g[x(k), u(k), p(k)] + e(k)
 \end{aligned} \tag{14}$$

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 modeled in MAPSS but have been added to the model for the problem studied in this paper. 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. A Kalman filter can be used with (14) to estimate the state vector x and the health parameter vector p .

The states, controls, health parameters, and measurements are summarized in Tables 1–4, along with their values at the nominal operating point considered in this paper, 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 [15]. 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*.

Constraints can be incorporated in the state estimator by using heuristic knowledge of the behavior of the health parameters. For example, it is known that health parameters do not improve over time.

$$p_m(k) \leq p_m(k-1), \quad m \in [1-6, 8, 10] \tag{15}$$

| State | Nominal Value |
|---------------------------------------|---------------|
| LPT Rotor Speed | 7264 RPM |
| HPT Rotor Speed | 12152 RPM |
| Average Hot Section Metal Temperature | 1533 °R |

Table 1: MAPSS turbofan model states and nominal values.

| Control | Nominal Value |
|--------------------------------|---------------------|
| Main Burner Fuel Flow | 2454 lbm / hr |
| Variable Nozzle Area | 343 in ² |
| Rear Bypass Door Variable Area | 154 in ² |

Table 2: MAPSS turbofan model controls and nominal values.

$$p_m(k) \geq p_m(k-1), \quad m \in [7, 9]$$

In addition, it is known that the health parameters vary slowly with time. For example, since $\tilde{p}_1(k)$ is the constrained estimate of $p_1(k)$, we can enforce the following constraints on $\tilde{p}_1(k)$.

$$\begin{aligned} \tilde{p}_1(k) &\leq p_1(0) \\ \tilde{p}_1(k) &\leq \tilde{p}_1(k-1) + \gamma_1^+ \\ \tilde{p}_1(k) &\geq \tilde{p}_1(k-1) - \gamma_1^- \end{aligned} \tag{16}$$

where γ_1^+ and γ_1^- are nonnegative factors chosen by the user. These factors allow the health parameter estimate to vary only within prescribed limits from one time step to the next. Typically we choose $\gamma_1^- > \gamma_1^+$ so that the parameter estimate can change more in the negative direction than in the positive direction. This is in keeping with our *a priori* knowledge that this particular health parameter never increases with time. Ideally we would have $\gamma_1^+ = 0$ since $p_1(k)$ never increases with time. However, since the health parameter estimate varies around the true value of the health parameter, we choose $\gamma_1^+ > 0$. This allows some time-varying increase in the health parameter estimate to compensate for a previous estimate that was smaller than the true value.

| Health Parameter | Normalized Value |
|----------------------------------|------------------|
| Fan airflow | 1 |
| Fan efficiency | 1 |
| Booster tip airflow | 1 |
| Booster tip efficiency* | 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 |

Table 3: MAPSS turbofan model health parameters and nominal values. (*) The fourth health parameter is not yet implemented in MAPSS.

Constraints (16) are linear and can therefore easily be incorporated into the form $D(k)\tilde{x}(k) \leq d(k)$ as required in the constrained filtering problem statement (4). Note that the constrained filtering approach presented here 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 [5]) in conjunction with the methods presented in this paper.

4 Turbofan Engine Health Parameter Sensitivity Analysis

In this section we apply the constrained Kalman filtering constraint switching procedure introduced in Section 2.2 to the turbofan engine health parameter estimation problem. This includes analyzing the sensitivity of the measurements to the health parameter values as was done in reference [2]. As discussed in Section 2.2, we then use this information to decide if an unconstrained health parameter estimate is reliable. If measurement residuals that are highly sensitive to a given health parameter are near zero, then we have a high confidence in the estimate of that health parameter, and we relax the constraints. However, if the measurement residuals are large, then we have a low confidence in the estimate of that health parameter, and we

| 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 |

Table 4: MAPSS turbofan model measurements, nominal values, and signal-to-noise ratios. SNR is defined here as the nominal measurement value divided by one standard deviation of the measurement noise.

enforce constraints that correspond to our heuristic knowledge of health parameter behavior.

The nonlinear MAPSS model of (14) exists as hundreds of MATLAB[®] routines, Simulink[®] blocks, and table lookups. The model was linearized using finite differences and then augmented to obtain the system

$$\begin{aligned}
 \begin{bmatrix} x(k+1) \\ p(k+1) \end{bmatrix} &= \begin{bmatrix} A_1 & A_2 \\ 0 & I \end{bmatrix} \begin{bmatrix} x(k) \\ p(k) \end{bmatrix} + \begin{bmatrix} B \\ 0 \end{bmatrix} \begin{bmatrix} u(k) \end{bmatrix} + \begin{bmatrix} w_x(k) \\ w_p(k) \end{bmatrix} \quad (17) \\
 y(k) &= \begin{bmatrix} C_1 & C_2 \end{bmatrix} \begin{bmatrix} x(k) \\ p(k) \end{bmatrix} + e(k) \\
 &= C \begin{bmatrix} x(k) \\ p(k) \end{bmatrix} + e(k)
 \end{aligned}$$

If we followed the approach given in Section 2.2 we would use C_2 as the sensitivity of the measurements to the health parameters. However, if the system is operating in steady state so that $x(k+1) = x(k)$ and $p(k+1) = p(k)$, then the coupling between $x(k)$ and $p(k)$ can be exploited to obtain more complete sensitivity information. In this case (17) can be solved for $y(k)$ as

$$y(k) = [C_1(I - A_1)^{-1}A_2 + C_2]p(k) + C_1(I - A_1)^{-1}w_x(k) + \quad (18)$$

$$\begin{aligned}
& C_1(I - A_1)^{-1}Bu(k) + e(k) \\
= & \Delta p(k) + C_1(I - A_1)^{-1}(Bu(k) + w_x(k)) + e(k)
\end{aligned}$$

where Δ , defined by the above equation, is the sensitivity matrix of the measurements to the health parameters. The element Δ_{ij} gives the sensitivity of the i th measurement to the j th health parameter. In practice we normalize Δ by dividing each row by the corresponding nominal measurement value given in Table 4.

The difference between the sensitivity matrix obtained using the approach of Section 2.2 (C_2), and the sensitivity matrix obtained here (Δ), is analogous to the difference between a partial derivative and a total derivative. Δ is a more accurate measure of the sensitivity (assuming that the system is in steady state).

During the execution of the Kalman filter, the measurement residuals are given by

$$\nu(k) = y(k) - [C_1\hat{x}(k) + C_2\hat{p}(k)] \quad (19)$$

The theoretical mean and variance of the residuals are given in (11). Therefore, if the measurement residuals satisfy their theoretical statistical properties, we can have confidence that the state and health parameter estimates are reliable. We generate a list of the measurements that are most sensitive to each health parameter. This is obtained by sorting each column of the sensitivity matrix Δ in descending order. In the case of MAPSS at the operating point used in this paper, the normalized sensitivity matrix is given as

$$\Delta = \begin{bmatrix}
0.01 & 0.06 & \mathit{0.12} & 0.00 & \mathit{0.27} & 0.39 & \mathit{0.06} & \mathit{0.27} & 0.14 & 0.02 \\
\mathit{0.43} & \mathit{0.33} & 0.09 & 0.00 & 0.16 & 0.17 & 0.04 & 0.15 & \mathit{0.45} & \mathit{0.19} \\
\mathit{0.14} & \mathit{0.21} & 0.10 & 0.00 & 0.09 & 0.07 & 0.01 & 0.02 & \mathit{0.35} & 0.11 \\
0.05 & \mathit{0.25} & 0.11 & 0.00 & 0.12 & 0.01 & 0.02 & 0.02 & \mathit{0.41} & \mathit{0.14} \\
0.04 & 0.03 & 0.02 & 0.00 & 0.03 & 0.01 & 0.01 & 0.02 & 0.11 & 0.04 \\
0.01 & 0.07 & 0.09 & 0.00 & 0.01 & 0.16 & \mathit{0.21} & 0.18 & 0.03 & 0.04 \\
0.00 & 0.19 & 0.04 & 0.00 & 0.07 & 0.08 & 0.03 & 0.08 & 0.28 & \mathit{0.11} \\
\mathit{0.08} & 0.10 & \mathit{0.23} & 0.00 & \mathit{0.17} & \mathit{0.64} & \mathit{1.22} & \mathit{0.52} & 0.13 & 0.09 \\
0.03 & 0.10 & 0.04 & 0.00 & 0.10 & \mathit{0.44} & 0.05 & 0.18 & 0.00 & 0.05 \\
0.06 & 0.12 & \mathit{0.13} & 0.00 & \mathit{0.16} & \mathit{0.64} & 0.05 & \mathit{0.43} & 0.02 & 0.10 \\
0.01 & 0.12 & 0.03 & 0.00 & 0.07 & 0.17 & 0.03 & 0.14 & 0.10 & 0.07
\end{bmatrix} \quad (20)$$

The three largest sensitivities in each column are italicized. (The fourth health parameter is not yet implemented in MAPSS, so the fourth column of Δ is zero.)

We see that the measurements that are most sensitive to the first health parameter are measurement numbers 2, 3, and 8; the measurements that are most sensitive to the second health parameter are measurement numbers 2, 3, and 4; and so on. This tells us that if residuals 2, 3, and 8 are small, then we can have a high confidence in our unconstrained estimate of health parameter 1; if residuals 2, 3, and 4 are small, then we can have a high confidence in our unconstrained estimate of health parameter 2; and so on. In practice, the decision of how many residuals to use for each health parameter, and what threshold values to use for those residuals, is an open question. Nevertheless, the theory presented in this section gives a general approach for deciding when to relax constraints and when to enforce constraints.

5 Simulation Results

We simulated the methods discussed in this paper using MATLAB[®]. We measured a 3 second burst of open-loop engine data at 10 Hz during each flight (while the engine was operating at steady state). These routine data collections were performed over 100 flights at the single operating point shown in Tables 1, 2, and 4. The engine's health parameters were initialized to the values shown in Table 3 and then deteriorated a small amount once each flight (i.e., once every 30 time steps). The signal-to-noise ratios were determined on the basis of NASA experience and previously published data [15] and are shown in Table 4. 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 to 0.01% of the nominal parameter value. 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. In the enforcement of constraints we chose the γ variables in (16) such that the maximum allowable change in \tilde{p} was a linear-plus-exponential function of time that reached a maximum of 9% after 500 flights in the direction of expected change, and 3% after 500 flights in the opposite direction. The true health parameter values never change in a direction opposite

to the expected change. However, we allow the estimate to change in the opposite direction to allow the Kalman filter to compensate for the fact that the previous estimate might be either too large or too small. The constraint boundaries are illustrated in Figure 2. In our simulations the true health parameters changed once per flight. However, the constraints that we imposed on our estimates varied with each time step.




Figure 2: The constraints are determined by allowing the state estimate to change a maximum of 3.0 times the expected magnitude in the expected direction of health parameter change, and 1.0 times the expected magnitude in the opposite direction.

We set the weighting matrix W in (6) equal to the identity matrix. Although Σ^{-1} is the optimal value of W in terms of the error covariance, we found from experience that setting $W = I$ results in only a small loss of performance, but it generates a significant savings in computational effort. This is because we avoid inverting the

13×13 covariance matrix at each time step. Although this extra computational effort would not be an issue for health parameter estimation in a postprocessing scenario, it could be important if the health parameter estimates were used in a real time control algorithm. Also note that in our system the health parameters that we estimate are all normalized to be equal to one. In another system with states or parameters that differ by an order of magnitude or more, it may be important to use $W = \Sigma^{-1}$ in order to account for varying scales.

We simulated a linear-plus-exponential degradation of the 10 health parameters over 100 flights. The initial health parameter estimation errors were assumed to be zero. The simulated health parameter degradations were representative of turbofan performance data reported in the literature [18]. Figure 3 shows a typical plot of the true deviation of health parameters 1 and 5, along with the deviations estimated by the unconstrained Kalman filter. We can see from the plot that shortly after flight 83 the unconstrained estimates are quite good. In this case we would not want to enforce constraints on the health parameter estimates at this particular time (although we may want to enforce constraints again later). But how can we *know* that the unconstrained health parameter estimates are good?

If we look at (20) to find the three measurements that are most sensitive to health parameters 1 and 5, we come up with measurements 1, 2, 3, 8, and 10. A closeup of the normalized residuals of these measurements, shown in Figure 4, indicates that they are indeed small between flight 83 and 84. This indicates that we can have a high confidence in the unconstrained estimates and relax our constraints at that moment.

Figure 5 shows what happens if we do *not* relax our constraints. A comparison of Figures 3 and 5 shows that the constrained estimate is better than the unconstrained estimate overall. Table 5 shows the means and standard deviations of the constrained and unconstrained estimates of the two health parameters shown in the figures. The mean errors of the unconstrained and constrained estimates are so close as to be essentially identical, but the constrained estimates clearly have smaller standard deviations than the unconstrained estimates.

At flight 83 the unconstrained estimate is good, but the enforcement of con-




Figure 3: Unconstrained Kalman filter estimates of health parameters 1 and 5. The estimates are noisy, but between flight 83 and 84 the estimates are quite accurate.

straints does not allow the constrained estimate to “catch up” to the unconstrained estimate. This is because our constrained estimator does not allow the estimates to change as quickly as the unconstrained estimator. This smoothing effect is why, overall, the constrained estimates in Figure 5 are more accurate than the unconstrained estimates in Figure 3. However, this is also why, in Figure 5, the constrained estimate cannot catch up to the unconstrained estimate at flight 83, even though the unconstrained estimate is better at that point in time.

Figure 6 shows what happens when we momentarily relax the constraints on the estimates of health parameters 1 and 5. Since the highly sensitive measurement residuals in Figure 4 are small at flight 83, we relax the constraints momentarily, allowing the constrained estimate to change abruptly for one time instant. We




Figure 4: Normalized residuals of measurements 1, 2, 3, 8, and 10 between flight 83 and 84. The residuals are less than one standard deviation for several time steps. This indicates that we can have a high confidence in our estimate of the health parameters to which those measurements are highly sensitive.

reset the constrained estimates to the unconstrained estimate values, and reapply the constraints for future estimates. The overall effect is an improvement in the accuracy of the constrained health parameter estimate.

In this example we chose to look at health parameters 1 and 5, and we chose to look at the three most sensitive residuals to each health parameter, which corresponded to measurements 1, 2, 3, 8, and 10. This example is only for illustrative purposes. In general we will not look at combinations of health parameters; we will rather look at each health parameter individually, and a certain number of residuals that are the most sensitive to each health parameter, where the number of residuals per health parameter is obtained by manual tuning.

| | Unconstrained Estimate | | Constrained Estimate | |
|------------------|------------------------|-----------|----------------------|-----------|
| | Mean | Std. Dev. | Mean | Std. Dev. |
| Health Param. #1 | -0.024 | 0.166 | 0.046 | 0.077 |
| Health Param. #5 | -0.172 | 0.193 | -0.151 | 0.077 |

Table 5: Means and standard deviations of the unconstrained and constrained state estimation errors shown in Figures 3 and 5. The top trace in each figure is health parameter #1 and the bottom trace is health parameter #5. The mean errors of the unconstrained and constrained estimates are so close as to be essentially identical, but the constrained estimates clearly have smaller standard deviations than the unconstrained estimates.

We ran 20 Monte Carlo simulations like this, each simulation consisting of 100 flights and the same health parameter degradation, but different measurement noise. Table 6 shows the performance of the filters averaged over 100 flights and 20 simulations. The standard Kalman filter estimates the health parameters to within 7.4% of their final degradations. The constrained filter estimates the health parameters to within 6.5% of their final degradations. The constrained filter with the addition of residual based constraint switching estimates the health parameters to within 6.2% of their final degradations. These numbers show the improvement that is possible with the constrained Kalman filter, and with residual based switching of the constraints.

The improved performance of the constrained filter comes with a price, and that price is computational effort. The constrained filter requires more computational effort than the unconstrained filter, and the incorporation of residual based switching requires more effort yet. However, computational effort is not a critical issue for turbfan health estimation since the filtering is performed on ground-based computers after each flight.

Note that the Kalman filter works well only if the assumed system model matches reality fairly closely. The method presented in this paper, by itself, 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




Figure 5: Constrained Kalman filter estimates of health parameters 1 and 5. The estimates are smooth and more accurate than the unconstrained estimates, but between flight 83 and 84 the estimates are less accurate than the unconstrained estimates (see Figure 3).

of additional residual check to verify the validity of the Kalman filter results [9], particularly for the application of turbofan engine health estimation considered in this paper [5].

6 Conclusion and Discussion

We have presented a residual based method for switching the constraints of a Kalman filter. The constrained Kalman filter uses a projection method to maintain the state variable estimates within a user-defined envelope. However, the constraints for many problems, including the turbofan health estimation problem investigated in this




Figure 6: Constrained Kalman filter estimates of health parameters 1 and 5. The estimates are set equal to the unconstrained estimates between flight 83 and 84 due to the small measurement residuals at this time.

paper, are heuristic. Therefore the engineer incurs some risk when implementing constraints. Although the use of constraints generally improves the accuracy of health estimation, there may be times when the constrained estimate is worse than the unconstrained estimate. If the unconstrained Kalman filter estimate is accurate then the incorporation of constraints can degrade the estimate. The use of residuals can quantify our confidence in the accuracy of the unconstrained estimate and tell us whether or not constraints should be incorporated.

The residual based method presented here measures residuals that are highly sensitive to given health parameters to decide whether or not constraints should be enforced on that health parameter. In practice there are several questions that need to be answered in the implementation of this theory. For instance, how many

| Health Parameter | Estimation Error (%) | | |
|----------------------------------|----------------------|--------------------|------------------|
| | Unconstrained Filter | Constrained Filter | Switching Filter |
| Fan airflow | 12.9 | 9.2 | 8.2 |
| Fan efficiency | 6.9 | 6.2 | 5.0 |
| Booster tip airflow | 10.9 | 10.6 | 10.0 |
| Booster tip efficiency* | N/A | N/A | N/A |
| Booster hub airflow | 7.4 | 6.8 | 6.2 |
| Booster hub efficiency | 3.8 | 3.1 | 3.0 |
| High pressure turbine airflow | 4.3 | 3.3 | 3.2 |
| High pressure turbine efficiency | 4.2 | 3.8 | 3.7 |
| Low pressure turbine airflow | 3.6 | 3.3 | 3.2 |
| Low pressure turbine efficiency | 11.3 | 11.2 | 11.0 |
| Average | 7.4 | 6.5 | 5.9 |

Table 6: Health parameter estimation errors (percent) of the Kalman filters, averaged over all flights. The estimation error is measured as $|(p-\tilde{p})/p_f|$, where p is the true health parameter value, \tilde{p} is the estimated health parameter value, and p_f is the health parameter value at the end of the simulation.

(*) The fourth health parameter is not yet implemented in MAPSS.

residuals should be used to decide whether or not to relax the constraints? How small should the residuals be (and for how many time steps) before constraints are relaxed? Or, using another approach, how *large* should the residuals be (and for how many time steps) before constraints are *enforced*? For this paper, the answers to these questions were found by manual adjustments, but further work could focus on a more theoretically and statistically rigorous analysis of the optimal answers to these questions.

We have seen that the constrained filter requires more computational effort than the standard Kalman filter. The incorporation of constraint switching requires yet more computational effort. This is due to the addition of the quadratic programming problem that must be solved in the constrained Kalman filter, and the residual checking logic that must be performed in the constraint switching process. The engineer must therefore perform a tradeoff between computational effort and estimation accuracy. For real time applications the improved estimation accuracy may not be

worth the increase in computational effort.

Although we have considered only linear state constraints, it is not conceptually difficult to extend this paper to nonlinear constraints. If the state constraints are nonlinear they can be linearized as discussed in [19]. Further work could explore the incorporation of state constraints for optimal smoothing, or the use of constraint switching in H_∞ filtering [21].

Note that methods other than the projection approach discussed in this paper can be used for incorporating inequality constraints into state estimation problems. For example, inequality constraints can be incorporated in a less rigorous way as soft constraints [4, 22]. Another more rigorous approach is to treat a general nonlinear constrained state estimation problem as a constrained parameter estimation problem. This results in computational effort that increases without bound as the time index of the system increases, but the computational effort can be bounded by limiting the cost function to a fixed number of time steps. This is called moving horizon estimation [11, 17]. Further work could explore the relationship between projection based constrained estimation and moving horizon estimation.

References

- [1] B. Anderson and J. Moore, Optimal Filtering (Prentice Hall, Englewood Cliffs, New Jersey, 1979).
- [2] B. Brunell, D. Viassolo, and R. Prasanth, Model Adaptation and Nonlinear Model Predictive Control of an Aircraft Engine, ASME Turbo Expo 2004, Vienna, Austria, Paper GT2004-53780, June 2004.
- [3] H. DePold and F. Gass, The application of expert systems and neural networks to gas turbine prognostics and diagnostics, ASME Journal of Engineering for Gas Turbines and Power (121), pp. 607-612, Oct. 1999.
- [4] P. Dewallef, O. Leonard, and K. Mathioudakis, On-line aircraft engine diagnostic using a soft-constrained Kalman filter, ASME Turbo Expo, Paper GT2004-53539 (2004).

- [5] D. Doel, TEMPER – A gas-path analysis tool for commercial jet engines, ASME Journal of Engineering for Gas Turbines and Power (116), pp. 82-89, Jan. 1994.
- [6] D. Doel, An assessment of weighted-least-squares-based gas path analysis, ASME Journal of Engineering for Gas Turbines and Power (116), pp. 366-373, April 1994.
- [7] R. Fletcher, Practical Methods of Optimization – Volume 2: Constrained Optimization (John Wiley & Sons, New York, 1981).
- [8] B. Friedland, Treatment of bias in recursive filtering, IEEE Transactions on Automatic Control (AC14), pp. 359-367, Aug. 1969.
- [9] A. Gelb, Applied Optimal Estimation (MIT Press, Cambridge, Massachusetts, 1974).
- [10] P. Gill, W. Murray, and M. Wright, Practical Optimization (Academic Press, New York, 1981).
- [11] G. Goodwin, J. De Dona, M. Seron, and X. Zhuo, Lagrangian duality between constrained estimation and control, Automatica (41), pp. 935-944, June 2005.
- [12] T. Kobayashi and D. L. Simon, A hybrid neural network-genetic algorithm technique for aircraft engine performance diagnostics, 37th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, July 2001.
- [13] H. Lambert, A simulation study of turbofan engine deterioration estimation using Kalman filtering techniques, NASA TM 104233, June 1991.
- [14] D. Massicotte, R. Morawski, and A. Barwicz, Incorporation of a positivity constraint into a Kalman-filter-based algorithm for correction of spectrometric data, IEEE Transactions on Instrumentation and Measurement (44), pp. 2-7, February 1995.
- [15] W. Merrill, Identification of multivariable high-performance turbofan engine dynamics from closed-loop data, AIAA Journal of Guidance, Control, and Dynamics (7), pp. 677-683, Nov. 1984.

- [16] K. Parker and K. Melcher, The modular aero-propulsion systems simulation (MAPSS) users' guide, NASA TM 2004-212968, March 2004.
- [17] C. Rao, J. Rawlings, and D. Mayne, Constrained state estimation for nonlinear discrete-time systems: Stability and moving horizon approximations, *IEEE Transactions on Automatic Control* (48), pp. 246-258, Feb. 2003.
- [18] O. Sasahara, JT9D engine/module performance deterioration results from back to back testing, *International Symposium on Air Breathing Engines*, pp. 528-535, 1985.
- [19] D. Simon and T. Chia, Kalman filtering with state equality constraints, *IEEE Transactions on Aerospace and Electronic Systems* (39), pp. 128-136, Jan. 2002.
- [20] D. Simon and D. L. Simon, Aircraft turbofan engine health estimation using constrained Kalman filtering, *ASME Journal of Engineering for Gas Turbines and Power* (127), pp. 323-328, April 2005.
- [21] D. Simon, A game theory approach to constrained minimax state estimation, *IEEE Transactions on Signal Processing*, in print.
- [22] D. Simon and D. L. Simon, Kalman filtering with inequality constraints for turbofan engine health estimation, *IEE Proceedings – Control Theory and Applications*, in print.
- [23] A. Volponi, H. DePold, and R. Ganguli, The use of Kalman filter and neural network methodologies in gas turbine performance diagnostics: a comparative study, *ASME Turbo Expo 2000*, Munich, Germany, Paper 2000-GT-547, May 2000.