Robust/Game-Theoretic Methods in Filtering and Estimation

William M. McEneaney *

Abstract

Robust/Game Theoretic approaches to filtering and estimation are considered. Estimates are obtained where the errors are bounded by some measure of the cost to the opponent. Both continuous and discrete models are considered. In the continuous case, a rather fast algorithm is discussed for nonlinear problems, and it is noted that a Riccati equation update is sufficient in the linear–quadratic sub-case. Analogous results are presented in the discrete case, although the additional technicalities are not present there.

1 Introduction

We present a Robust/Game Theoretic approach to the filtering problem. In this approach, one obtains an estimate for the state of some system and a certain bound on the error in this estimate in terms of some cost associated with the disturbances. In situations where the largest disturbances are effected by an opposing player in a game, this is a natural approach in that this cost can represent some actual cost to the opposing player. However, this approach can be used in a more general context. In fact, the natural control analogue is H_∞ control, which has the game interpretation (also dissipation interpretation) of bounding the measure of the poorness of the control by a measure of the size of the disturbances in the system (see for instance [1], [10] among many notable others). This approach has been used with significant success in recent years. Further, it has recently been shown that the nonlinear H_∞ problem is equivalent to the risk–averse limit of a risk–sensitive stochastic formulation of the problem (cf. [3]) thus yielding a stochastic interpretation.

This approach is especially appropriate for systems where one desires to track the state of some process in the presence of an antagonistic player. In such a case, the traditional approach of assuming random disturbances generated by, say for instance a Gaussian process, would not be appropriate. It is also appropriate when the estimates will be utilized by a significantly risk–averse controller.

We will develop the Robust/Game Theoretic filtering approach first in the case of a continuous time/continuous state process with measurements at discrete times. This will lead to a PDE (partial differential equation) formulation for the information state between measurement times, and a simple subtraction for the measurement update. The information state [7] is the analogue of the conditional density in the stochastic approach. Some relatively fast methods have been developed for this problem in the fully nonlinear case, however these are still only appropriate for low-dimensional (sub-)systems. We will also briefly indicate the Riccati equation formulation which may be used in the linear–quadratic case; this is similar to the linear–quadratic Kalman filter update.

We will also discuss the development in the case of a discrete time/ discrete state–space model. Here, the technical difficulties will be significantly reduced. Readers who are interested only in the discrete case should be able to proceed directly to Section 4.

2 Continuous Models

Portions of this section are condensed from material in [9], but are needed here for the material to follow. Let the state dynamics be given by

$$\frac{dX}{dt} = f(X) + \sigma(X)w \tag{1}$$

where X is the state (taking values in $\mathbf{R}^n$), f represents the nominal dynamics, w is a deterministic, but (a priori) unknown, L_2 process, and σ is some (matrix–valued) multiplier on the disturbance. Assume that at each measurement time t_i, we receive a measurement which we model as

$$z_i = g(X(t_i)) + \rho(X(t_i))v_i \tag{2}$$

*Dept. of Mathematics, North Carolina State University, Raleigh, NC 27695–8205, USA, http://www4.ncsu.edu/~wmm/, wmm@eos.ncsu.edu, Research partially supported by DARPA grant F30602-99-2-0548, AFOSR grant F49620-98-1-0304 and NSF grant DMS-9971546.

where g represents the disturbance–free measurement model (taking values in $\mathbf{R}^l$), and v_i is the disturbance at time t_i. Here z_i may be vector–valued, and ρ may be matrix–valued.

We will assume that f, σ, g and ρ are all continuously differentiable. We will assume that f, g and σ are globally Lipschitz in x, that σ is bounded, and that $a \doteq \sigma\sigma^T$ is uniformly non–degenerate, i.e. that there exists $\delta > 0$ such that

$$\xi^T \sigma(x)\sigma^T(x)\xi \geq \delta|\xi|^2 \qquad \forall\, x \in \mathbf{R}^n,\ \forall\, \xi \in \mathbf{R}^n. \quad (3)$$

We will also assume that Range$(\rho(x)) = \mathbf{R}^l$ for all $x \in \mathbf{R}^n$ which guarantees that for any z, x there exists some v satisfying (2) which should certainly be the case if our measurement model is properly constructed. Finally, we define ρ^{-1} by

$$\rho^{-1}(x)b = \operatorname{argmin}\{|v| : \rho(x)v = b\}. \quad (4)$$

Assume that ρ^{-1} is uniformly bounded, that is, there exists $C_\rho < \infty$ such that $|\rho^{-1}(x)b| \leq C_\rho|b|$ for all $x \in \mathbf{R}^n$ and $b \in \mathbf{R}^l$. Note also that these assumptions imply that if we view the integral version of (1)

$$x_T \doteq X(T) = x_0 + \int_0^T f(X(t)) + \sigma(X(t))w(t)\,dt$$

as a mapping from x_0 to x_T then this mapping is one–to–one and onto for any $w \in \mathrm{L}_2$, thereby ensuring that we may invert it.

Let $p_0(x_0)$ be a measure of our uncertainty about the initial state x_0, and let it have at most quadratic growth, so that $0 \leq p_0(x) \leq C(1 + |x|^2)$ for all $x \in \mathbf{R}^n$ for some $C < \infty$.

Suppose that we wish to estimate the state at time T, and that there have been N measurements in $(0, T]$. Consider the cost criterion

$$J(T, x_T, w(\cdot)) = -p_0(x_0) - \int_0^T |w(t)|^2\,dt \quad (5)$$
$$- \sum_{i=1}^N |\rho^{-1}(X(t_i))[y_i - g(X(t_i))]|^2,$$

and corresponding information state

$$P(T, x_T) = \sup_{w \in L_2} J(T, x_T, w).$$

Note that P is a generalization of the (additive inverse of the) minimum disturbance energy required for the state to be x_T at time T while producing the measurements observed up to time T.

Due to the discrete nature of the measurements, it is helpful recall the form of the dynamic programming principle for such a system. Let $P(t_j^+, x)$ be the value at time t_j just after measurement j, and $P(t_j^-, x)$ be the value at time t_j just before the measurement. If T is not a measurement time, and $t \in (t_{k-1}, t_k)$, we have

$$P(T, x) = \sup_{w \in L_2([t,T];\mathbf{R}^m)} \Big\{ P(t, X(t)) - \int_t^T |w(r)|^2\,dr$$
$$- \sum_{j=k}^N |\rho^{-1}(X(t_j))[z_j - g(X(t_j))]|^2 \Big\},$$

and when T occurs at the time of measurement j, we have

$$P(T^+, x) = P(T^-, x) - |\rho^{-1}(x)[z_j - g(x)]|^2.$$

The corresponding dynamic programming equations are

$$V(0, x) = -p_0(x) \qquad\qquad t = 0 \quad (6)$$
$$0 = V_t - \sup_{w \in \mathbf{R}^m} \{-[f(x) + \sigma(x)w]^T \nabla_x V - |w|^2\}$$
$$= V_t + f^T(x)\nabla_x V - \frac{1}{4}(\nabla_x V)^T a(x)\nabla_x V \quad (7)$$
$$V(t_j^+, x) = V(t_j^-, x) - |\rho^{-1}(x)[z_j - g(x)]|^2 \quad (8)$$

The following result is a standard type of dynamic programming result clarifying the (weak) sense in which the information state is a solution of PDE (7). However, as will be seen in the next section, there is a new, fast numerical technique which does not directly use the PDE interpretation as the basis for the algorithm. A proof of the theorem may be found in [9].

Theorem 2.1 *P is a continuous viscosity solution of (7) between measurement times.*

We will now indicate how the information state, P, may be used to obtain a robust/game theoretic state estimate. First we will indicate a direct approach which was originally proposed by Mortensen ([6], [11]), and then we will combine the above with a quadratic estimation criterion to obtain the game theoretic filter. For the direct Mortensen approach, note that $-P(T, x)$ represents, in some sense, the minimum disturbance energy needed for the target state at time T to be x for given measurements $z_\cdot$. The Mortensen estimate is any

$$\widehat{e}_T \in \operatorname*{argmax}_{x \in \mathbf{R}^n} P(T, x). \quad (9)$$

We can assert the existence of the argmax in (9) under reasonable assumptions as follows.

Theorem 2.2 *Suppose there exists $C_p > 0, \overline{x}_0 \in \mathbf{R}^n$ such that $p_0(x) \geq C_p|x - \overline{x}_0|^2$ for all $x \in \mathbf{R}^n$. Then $\widehat{e}_T$ exists.*

Proof. Let the global Lipschitz constant on f be K, and let the bound on σ be m_σ. Let $\overline{x}_T \in \mathbf{R}^n$ be such that $\overline{X}(0) = \overline{x}_0$ when $\frac{d\overline{X}}{dt} = f(\overline{X})$ and $\overline{X}(T) = \overline{x}_T$ (the existence of which is guaranteed by the assumptions). Let X satisfy the dynamics of (1) with $X(T) = x_T$ for any $w \in L_2[0,T]$ and any $x_T \in \mathbf{R}^n$. Let $x_0 \doteq X(0)$. Then

$$|X(t) - \overline{X}(t)| = |(x_0 - \overline{x}_0) + \int_0^t [f(X) - f(\overline{X}) + \sigma(X)w]\, dr|$$

which by the assumptions and Cauchy-Schwarz

$$\leq |x_0 - \overline{x}_0| + m_\sigma \sqrt{T}\|w\|_{L_2} + \int_0^t K|X - \overline{X}|\, dr$$

where K is the Lipschitz bound on f, and m_σ the bound on $\sigma(x)$. Employing Gronwall's inequality yields

$$|x_T - \overline{x}_T| \leq (|x_0 - \overline{x}_0| + m_\sigma\sqrt{T}\|w\|_{L_2})C_T$$

for appropriate $C_T < \infty$. Consequently,

$$C_p|x_0 - \overline{x}_0|^2 + \|w\|_{L_2}^2 \geq C_p\left[\frac{|x_T - \overline{x}_T|}{C_T} - m_\sigma\sqrt{T}\|w\|_{L_2}\right]^2 + \|w\|_{L_2}^2$$

which after some calculations one finds

$$\geq \frac{C_p}{C_T^2(1 + m_\sigma^2 T)}|x_T - \overline{x}_T|^2.$$

Therefore, by (5) and the assumptions

$$J(T, x_T, w) \leq -\frac{C_p}{C_T^2(1 + m_\sigma^2 T)}|x_T - \overline{x}_T|^2.$$

Since this is true for any $w \in L_2$, one has

$$P(T, x_T) \leq -\frac{C_p}{C_T^2(1 + m_\sigma^2 T)}|x_T - \overline{x}_T|^2 \qquad (10)$$

which yields the result. □

Although the direct Mortensen estimate is of some interest, the following robust/game theoretic estimate will have more desirable properties in this context. Specifically, the information state can be combined with a quadratic estimate error criterion to obtain the robust filter. Let the dynamics and measurement models be those given above in (1) and (2), and make the same assumptions. However, now let the cost criterion take the form

$$I(T, e, x_T, w) = \gamma^2 J(T, x_T, w) + |x_T - e|^2.$$

Let

$$\begin{aligned} W(T, e) &= \sup_{x_0 \in \mathbf{R}^n} \sup_{w \in \mathcal{W}} I(T, e, x_T, w) \\ &= \sup_{x_T \in \mathbf{R}^n} \sup_{w \in \mathcal{W}} I(T, e, x_T, w). \end{aligned}$$

Note that

$$W(T, e) = \sup_{x_T \in \mathbf{R}^n} \left[\gamma^2 P(T, x_T) + |x_T - e|^2\right]. \qquad (11)$$

It is clear here that γ must be large enough so that this supremum will be finite and achieved at some x_T. By (10), one obtains the following.

Theorem 2.3 *There exists $\gamma < \infty$ such that the supremum in (11) is finite.*

This lower bound on such γ is directly analogous to the optimal disturbance attenuation parameter in H_∞ control below which the supremum in that problem becomes unbounded. Note that in a computational system, given a P, one can choose a γ corresponding to this P such that the supremum is finite.

Now note that since $W(T, e)$ is a supremum of functions which are strictly convex in e, it is also strictly convex in e. Further, $W(T, e) \to \infty$ as $|e| \to \infty$. Consequently, the minimum over e is obtained at some point, and we define the robust filter estimate uniquely at time T as:

$$e_T \doteq \operatorname*{argmin}_{e \in \mathbf{R}^n} W(T, e).$$

Further, by (10) one may choose γ large enough so that $\min_{e \in \mathbf{R}^n} W(T, e) \leq 0$, in which case one obtains the error estimate

$$|x_T - e_T|^2 \leq \gamma^2 \left[p_0(x_0) + \|w\|^2 + \sum_{i=1}^N |v_i|^2\right]$$

for all $x_0 \in \mathbf{R}^n$ and all $w \in \mathrm{L}_2$ where $\|\cdot\|$ represents the L_2–norm over $[0, T]$. This is the robust bound on the estimate error in terms of the energy of the disturbance.

Note that the information state computations are recursive. That is, having obtained an estimate at time T, one can obtain an estimate at time $\widehat{T} > T$ by extending the solution of (9) from T to $\widehat{T}$.

3 Computational Issues for the Continuous Model

The chief difficulty in a general method for the computation of this nonlinear filter is the propagation of the information state, P, between measurement times. In the stochastic analogue, this is similar to the propagation of a solution to the Zakai or Kushner equation. The form of the above problem leads to a unique method for the propagation of the information state between measurement times. This method allows for greatly increased computational speed. The following discussion outlines the analytical basis of the algorithm; the details may be found in [2].

Let S_t denote the semi-group associated with PDE (7). That is, given the solution at time t, $P(t,\cdot)$, the solution at time $t+\tau$ is given by $S_\tau[P(t,\cdot)](x)$. Consider the max–plus algebra over $\mathbf{R}\cup\{-\infty\}$ which is given by

$$\begin{aligned} a \oplus b &= \max\{a,b\} \\ a \otimes b &= a+b. \end{aligned}$$

An important result is that this semi-group is linear over the max–plus algebra. That is, given functions $\phi_1(x), \phi_2(x)$ and constants a_1, a_2, one has

$$\begin{aligned} S_t[(a_1\otimes\phi_1(\cdot))\oplus(a_2\otimes\phi_2(\cdot))](x) &= \big(a_1\otimes S_t[\phi_1(\cdot)](x)\big) \\ &\quad \oplus\big(a_2\otimes S_t[\phi_2(\cdot)](x)\big). \end{aligned}$$

Using a max–plus basis expansion of the space of solutions, so that $P(t_i,x) = \bigoplus_j [a_j(t_i)\otimes\psi_j(x)]$ with basis functions ψ_j, one may obtain

$$S_\tau\Big[\bigoplus_j [a_j(t_i)\otimes\psi_j]\Big](x) = \bigoplus_j [a_j(t_i)\otimes S_\tau[\psi_j]](x).$$

Precomputing the $S_\tau[\psi_j]$, and expanding them in the form $S_\tau[\psi_j] = \bigoplus_k [B_{j,k}\otimes\psi_k]$, one obtains the following max–plus linear propagation algorithm

$$A(t_i+\tau) = B\otimes A(t_i)$$

where $A(t)$ is the vector of coefficients in the basis expansion at time t and B is the matrix $[B_{j,k}]$. Note that the matrix B can be precomputed, and consequently, the (approximate) propagation of P between measurement times reduces to a max–plus matrix–vector multiplication. See [2] for the details.

Some snapshots of P in a simple one–dimensional example are depicted in Figures 1–4 to indicate the possible nonsmooth and multi-lobe nature of the information state. (Each figure actually depicts three closely spaced snapshots.)

Even with the significant advance described just above, one cannot hope to filter high-dimensional nonlinear problems in real–time. Of course, in the linear–quadratic case, the computations reduce to propagation of the Riccati equations as usual, and so are quite computable. To see this, suppose that $f(x) = Ax$, $g(x) =$

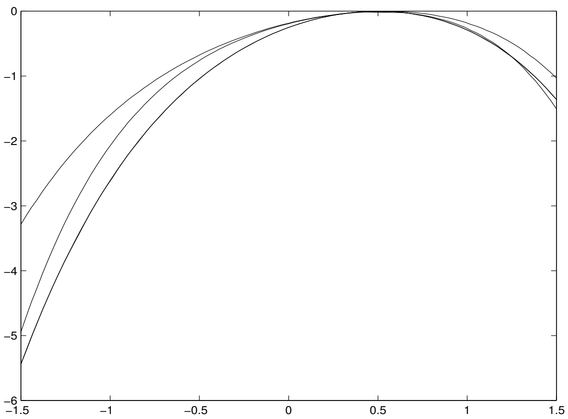

Figure 1: $(t = 0)$

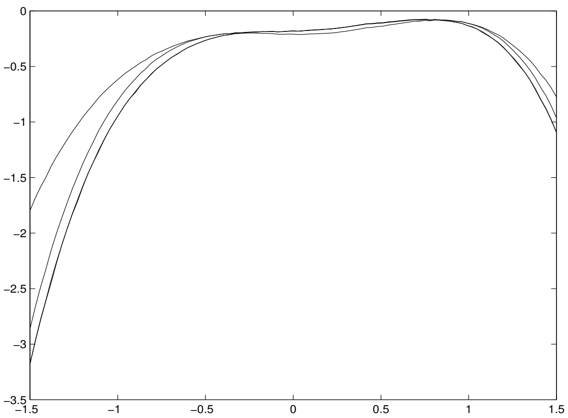

Figure 2: $(t = 0.2)$

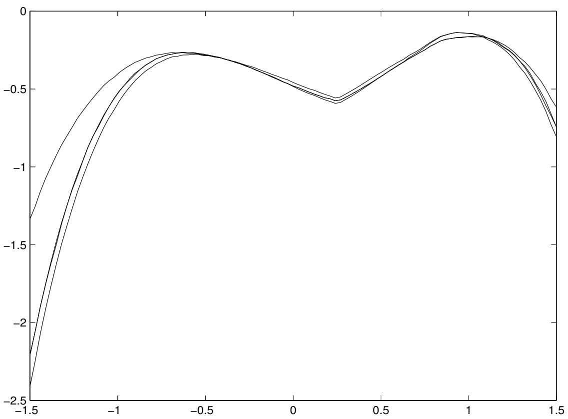

Figure 3: $(t = 0.4)$

Bx, σ, ρ are constant, and $p_0(x) = -(x-\overline{x}_0)^T\widetilde{P}_0(x-\overline{x}_0)$ for some $\overline{x}_0$ and positive definite symmetric $\widetilde{P}_0$. Then, one finds that $P(t,\cdot)$ is quadratic. Specifically, the information state may be written in the form $P(t,x) = -(x-\overline{x}(t))^T\widetilde{P}(t)(x-\overline{x}(t)) + r(t)$ for proper choice of

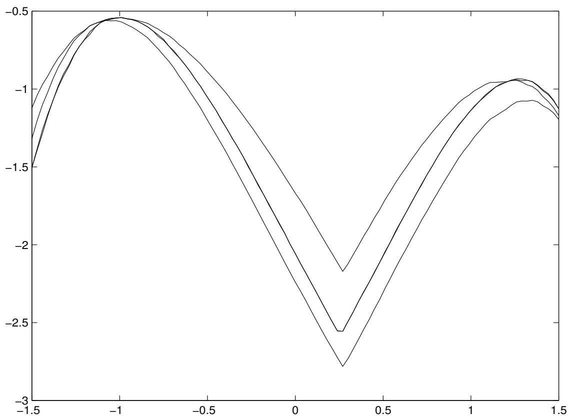

Figure 4: ($t = 0.6$)

$\widetilde{P}, \overline{x}, r$. In fact, between measurement times, $\widetilde{P}$ satisfies

$$\dot{\widetilde{P}} = -(A^T\widetilde{P} + \widetilde{P}A) + \widetilde{P}^T\sigma\sigma^T\widetilde{P},$$

while at measurement times, the update is given by

$$\widetilde{P}^{+} = \widetilde{P}^{-} - B^T\rho^{-1}\rho^{-T}B.$$

There are corresponding updates for $\overline{x}(t)$ and $r(t)$. Further this estimate is identical to that generated by the (now linear) Kalman filter for the analogous (linear) stochastic model with initial covariance $\widetilde{P}_0^{-1}$, and in fact, $\widetilde{P}^{-1}(t)$ is the covariance matrix for $t > 0$.

4 Discrete Time/Discrete State

The previous sections described some of the current theory behind robust/game-theoretic estimation in the case of a system where the variables were continuous. Of course in many situations such as the "command and control" situation, many of the variables are expected to be discrete, say from some finite set. Here we will suppose the state variables take values in a discrete set, say $\mathcal{M}$. The typical approach to such problems is the Markov chain approach. In contradistinction to this, we will indicate a robust/game-theoretic approach to the estimation problem with state taking values in $\mathcal{M}$. It is not difficult to imagine applications for which this approach would be appropriate. First of all, there are situations where the dynamics and measurements may be corrupted by an antagonistic player. For instance, in battlefield management, one may have an opposing player who will attempt to corrupt our observations so as to give a false picture of the true situation. Secondly, there are problems where one wishes to be very risk-averse such as certain portfolio optimization problems.

To further clarify the problem formulation here, let us suppose that the time variable is discrete (although it will be clear that continuous time problems have an analogous approach). In fact, let us simply suppose that the time variable takes integer values starting from initial time $k = 0$. In this section, we will ignore the technical details in order to focus simply on the general picture.

Let the state at time k be given by α_k where $\alpha_k \in \mathcal{M}$ for all k. Let up suppose that the cardinality of $\mathcal{M}$ is m, and let the possible values of the state correspond to the elements of $\{1, 2, 3, ..., m\}$ which we will henceforth refer to as $\mathcal{M}$ (i.e. we will simply take $\mathcal{M} = \{1, 2, 3, ..., m\}$).

By analogy with the continuous case, one could suppose a dynamic model of the form

$$\alpha_{k+1} = f(k, \alpha_k, w_k)$$

where w_k is the disturbance (due to an antagonistic player) at time k. Then, one could associate a cost with each value of w_k of the form say $D(w_k)$. However, it appears conceptually and numerically simpler to develop a model more analogous to the Markov chain approach. For instance, by analogy with the Markov chain transition **probability** matrix, we will formulate the above dynamics and cost in terms of a transition **cost** matrix. more specifically, suppose the cost for transitioning (in one time-step) from state $i \in \mathcal{M}$ to $j \in \mathcal{M}$ is denoted by $c_f(i, j)$ where we are further simplifying the notation by assuming that this is time–independent. Then, we form the matrix C^f as the matrix whose elements are given by

$$C^f_{i,j} = c_f(i, j).$$

Note that one may restrict the state from transitioning from say state i_0 to state j_0 in one time step by setting $C^f_{i_0,j_0} = +\infty$.

Similarly, suppose that we have a fixed, finite set of possible measurement values. Let us suppose that there are n possible values, and simply denote them by $\mathcal{N} = \{1, 2, 3, ..., n\}$. Then for each time k, the k^{th} measurement takes its value in $\mathcal{N}$, i.e. $y(k) \in \mathcal{N}$ for all k. Let the cost to the disturbance for producing measurement j when the state is i be $c_g(i, j)$. Again, we can organize this data in the form of an $m \times n$ matrix C^g where

$$C^g_{i,j} = c_g(i, j).$$

Finally, let the vector p^0 of length m correspond to our initial information (at time $k = 0$).

It is not difficult to see how one would compute the worst–case cost for being in state j at time $k = 1$ given measurement $y(1) \in \mathcal{N}$. It would be

$$\sup_{i \in \mathcal{M}} \left\{ -\left[p^0_i + C^f_{i,j} + C^g_{j,y(1)} \right] \right\}.$$

To generalize this to later times, let $\widehat{\alpha}^k$ be a vector of length $k+1$ of elements of $\mathcal{M}$ corresponding to a state transition path from time 0 to time k. Define the cost to be in state j at time k with past history $\widehat{\alpha}^{k-1}$, given measurements $y(l)$ for $l = 1, 2, ...k$ to be

$$J(k, j, \widehat{\alpha}^{k-1}) = -\Bigg[p^0_{\widehat{\alpha}_0^{k-1}} + \sum_{l=1}^{k-1} \Big(C^f_{\widehat{\alpha}_{l-1}^{k-1}, \widehat{\alpha}_l^{k-1}} + C^g_{\widehat{\alpha}_l^{k-1}, y(l)} \Big) + C^f_{\widehat{\alpha}_{k-1}^{k-1}, j} + C^g_{j, y(k)} \Bigg].$$

As in the continuous case, one defines the "information state" at time k for state j in the worst–case cost sense as

$$P(k, j) = \sup_{\widehat{\alpha}^{k-1} \in \mathcal{A}^{k-1}} J(k, j, \widehat{\alpha}^{k-1})$$

where $\mathcal{A}^k$ is the set of sequences of length $k+1$ taking values in $\mathcal{M}$.

The information state captures all the information needed to produce the robust estimate at the current time. (It is analogous to the conditional probability distribution in stochastic models.) In some sense, the information state, $P(k, j)$, represents some generalized notion of the cost to the opposing player for the state to be in state j at time k while having produced the measurements that have been seen. Given the information state at any time, one chooses the parameter γ sufficiently large so that

$$W(k, e) \doteq \sup_{j \in \mathcal{M}} \left[|j - e|^2 + \gamma^2 P(k, j) \right]$$

has a minimum (in e) which is non-positive. (Note that the existence of such a γ actually implies assumptions on the choice of C^f, C^g, p^0. In the continuous case of Section 2, there were standard simple assumptions which guarantee this. However, in the more "free–form" description of this section, there are no obvious choices for these assumptions, and one would need to be more specific about the form of the model before making such assumptions. Also, one might choose something other than the quadratic form $|j-e|^2$ since there is no obvious reason why that would necessarily be the most appropriate form.) Let e_k be the minimizer. This implies that

$$W(k, e_k) = \sup_{j \in \mathcal{M}} \left[|j - e_k|^2 + \gamma^2 P(k, j) \right] \le 0.$$

In other words,

$$|j - e_k|^2 \le -\gamma^2 P(k, j) \qquad \forall\, j \in \mathcal{M}.$$

Using the definition of P, this yields

$$|j - e_k|^2 \le \gamma^2 \Bigg[p^0_{\widehat{\alpha}_0^{k-1}} + \sum_{l=1}^{k-1} \Big(C^f_{\widehat{\alpha}_{l-1}^{k-1}, \widehat{\alpha}_l^{k-1}} + C^g_{\widehat{\alpha}_l^{k-1}, y(l)} \Big) + C^f_{\widehat{\alpha}_{k-1}^{k-1}, j} + C^g_{j, y(k)} \Bigg]$$

for any initial state $\widehat{\alpha}_0^{k-1}$ and any choice of hostile counteractions. That is the squared error in our estimate (e_k) of the state at time k is bounded above by the right-hand side where the right-hand side represents some measure of the costs to the opposing player for the disturbances it employs. Note that the larger one needs to take γ, the less one can attenuate the effects of the counteractions on our squared estimate error.

Let us also write down the dynamic programming equation that one would solve to update P from one time to the next.

$$P(k+1, j) = \sup_{i \in \mathcal{M}} \Big\{ C^f_{i,j} + C^g_{j, y(k+1)} + P(k, i) \Big\}.$$

Of course, the right-hand side is a max–plus linear operator on $P(k, \cdot)$, and this may be written more abstractly as

$$P(k+1, j) = \mathcal{S}_1 \big[P(k, \cdot) \big](j) \qquad \forall\, j \in \mathcal{M}$$

where $\mathcal{S}_1$ is max–plus linear.

We illustrate this with a very simple example. Let the state space be $\mathcal{M} = \{0, 1, 2, 3, 4, 5\}$. Let the transition cost matrix be

$$C^f = \begin{bmatrix} 0 & 1 & +\infty & +\infty & +\infty & +\infty \\ 0 & 0 & 1 & +\infty & +\infty & +\infty \\ 0 & 0 & 0 & 1 & +\infty & +\infty \\ 0 & 0 & 0 & 0 & 1 & +\infty \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

Let the measurement cost matrix be

$$C^g = \begin{bmatrix} 0 & 1 & 2 & 3 & 4 & 5 \\ 1 & 0 & 1 & 2 & 3 & 4 \\ 2 & 1 & 0 & 1 & 2 & 3 \\ 3 & 2 & 1 & 0 & 1 & 2 \\ 4 & 3 & 2 & 1 & 0 & 1 \\ 5 & 4 & 3 & 2 & 1 & 0 \end{bmatrix}$$

Let the initial information be $p^0 = (1, 0, 1, 2, 3, 4)$. A sequence of measurements was generated as depicted (by the $*$'s) in Figure 5 below. A robust estimator also appears (solid lines) in Figure 5. (This was just one possible robust estimator since we did not restrict ourselves to the quadratic form on the right-hand side in the definition of W.) The information state is depicted over the time period in Figure 6.

These estimates would then be used by a robust/game–theoretic controller. See [5],[4] for examples of the manner in which robust filters are sometimes used as controller inputs in the continuous case. We will be employing these discrete robust filter estimates in a simpler game–theoretic method to be discussed in fu

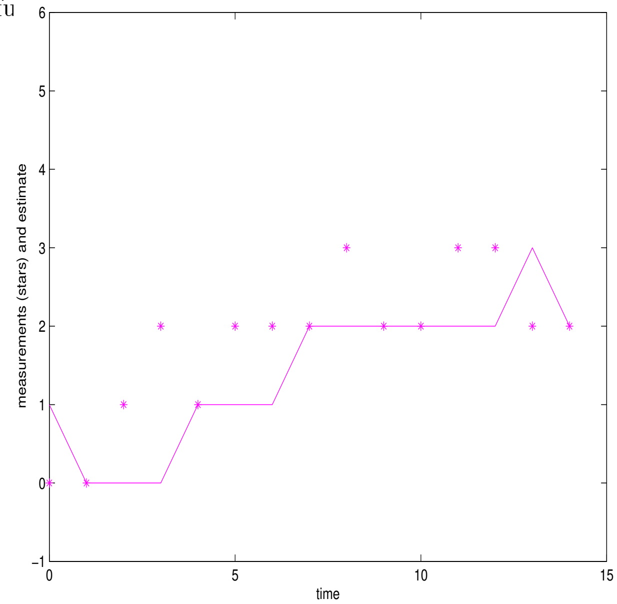

Figure 5: Measurements and Robust/Game–Theoretic Estimates

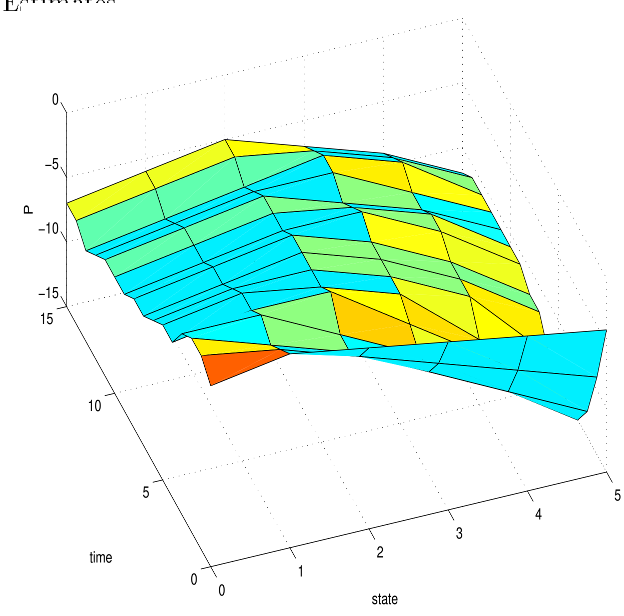

Figure 6: Information State ($P(k,j)$)

References

[1] T. Basar and P. Bernhard, **H$_\infty$ –Optimal Control and Related Minimax Design Problems**, Birkhäuser, Boston (1991).

[2] W.H. Fleming and W.M. McEneaney, "A max–plus based algorithm for an HJB equation of nonlinear filtering", SIAM J. Control and Optim., (to appear).

[3] W. H. Fleming and W. M. McEneaney, Risk sensitive control on an infinite time horizon, SIAM J. Control and Optim., **33** (1995), 1881–1915.

[4] E. Gallestey, M.R. James and W.M. McEneaney, "Max–plus approximation methods in partially observed H$_\infty$ control", 38th IEEE Conf. on Decision and Control (to appear).

[5] J.W. Helton and M.R. James, **Extending H_∞ Control to Nonlinear Systems: Control of Nonlinear Systems to Achieve Performance Objectives**, to appear, SIAM 1999.

[6] O. J. Hijab, Minimum Energy Estimation, University of California, Berkeley, PhD Thesis, 1980.

[7] M. R. James and J. S. Baras, Partially observed differential games, infinite dimensional HJI equations, and nonlinear H$_\infty$ control, SIAM J. Control and Optim., **34** (1996), 1342–1364.

[8] V.P. Maslov, "On a new principle of superposition for optimization problems", Russian Math. Surveys, **42** (1987) 43–54.

[9] W. M. McEneaney, "Robust/H$_\infty$ filtering for nonlinear systems", Systems and Control Letters, Vol. 33 (1998), 315–325.

[10] W. M. McEneaney, Robust control and differential games on a finite time horizon, Math. of Controls Signals and Systems, **8** (1995), 138–166.

[11] R. E. Mortensen, Maximum Likelihood Recursive Nonlinear Filtering, J. Optim. Theory Appl., **2** (1968) 386–394.