Chapter 12

Conclusions and Future Work

Having reached the end of a dissertation of such length, it is probably best to conclude with some brief remarks on various directions for future research that are suggested by the methods and results presented in this thesis. For convenience, we have divided the discussion into two parts — the first deals with extensions of some of the results studied in the earlier chapters, and the second deals with what we believe is a very important and promising approach to estimation and control, namely the mixed H^2/H^∞ approach.

12.1 Various Extensions

We begin with various extensions suggested by the studies and results of the earlier chapters.

Continuous-time Krein Space Theory

Perhaps the first direction for future work that is suggested by this thesis is the development of a continuous-time counterpart to the discrete-time Krein space estimation theory given in Chapter 2. Although the author, and his coworkers, have not yet fully embarked on such an endeavor, preliminary investigations suggest that this should be possible without too much difficulty. Here the Krein space projections will be related to stationarizing certain indefinite quadratic integral costs, and the conditions for a minimizing solution will be given in terms of the positivity of certain integral operators.[1] When one has state-space structure, the invertibility (and positivity) of the integral operators can be checked via the existence (and properties) of solutions to certain Riccati differential equations. Moreover, the projections can be computed via the continuous-time Krein space Kalman filter, a natural generalization of the conventional continuous-time Kalman filter. We also remark that such studies should have ramifications to continuous-time H^∞ estimation and control, and to quadratic differential games, as well as other areas.

Suboptimal Recursive Total Least-Squares Algorithms

We mentioned in Chapter 4 that another application of Krein space estimation is in the development of recursive algorithms for suboptimal total least-squares problems. Preliminary studies towards this goal have been performed in [SHK96a], where it is shown that the total least-squares problem (or errors-in-variables method) can be reformulated as

$$\min_x \left[-\bar{\sigma}_{n+1}^2 + (b - Ax)^*(b - Ax)\right],$$

where $\bar{\sigma}_{n+1}^2$, b and A are known quantities that need not concern us here. The important fact is that the above problem is an indefinite quadratic minimization problem that is amenable to Krein space methods. Although it is very unlikely to be possible to give recursive solutions to the exact total least-squares problem (since the solution involves a singular value decomposition which cannot be performed recursively), the Krein space theory may suggest certain suboptimal variants of this exact solution that will allow for efficient recursive computations. Moreover, we believe that there are probably other areas of least-squares theory that may benefit from the Krein space approach.

[1] Of course, in dealing with integral operators, one must take care as to whether these operators are invertible, bounded, etc. However, the condition for the existence and uniqueness of continuous-time Krein space projections is that the corresponding integral operators be invertible (very much like the discrete-time case where the associated indefinite Gramian had to be invertible).

Numerical Analysis of H^∞ Square-Root and Chandrasekhar Algorithms

As mentioned in Chapter 5, the conventional square-root array algorithms are preferred because of their numerical stability, which is largely due to the fact that the dynamic range of the variables are reduced (thus also reducing the condition numbers), and because the computations involve unitary transformations that are well-known to not amplify numerical (or roundoff) errors. Since the H^∞ square-root array algorithms (and the H^∞ Chandrasekhar recursions) are the direct analogs of their conventional counterparts, it seems plausible that they may be more attractive for the numerical implementation of H^∞ filters (and controllers). However, since J-unitary, rather than unitary, operations are involved, further investigation is needed to determine what the numerical performances of these algorithms are, and what the best way to implement the J-unitary transformations is. This should be an area of much practical significance.

Further Study of the Indefinite DARE

In Chapter 7 we studied the discrete-time algebraic Riccati equation (DARE) with possibly indefinite coefficient matrices. We should mention that similar results can be obtained for the continuous-time algebraic Riccati equation (CARE), and indeed that it is quite easier to do so for the CARE than it is for the DARE.[2] Now in Chapter 7 the existence of solutions to the DARE were shown to be equivalent to the existence of certain proper factorizations of the associated Popov functions, and also equivalent to certain properties of the invariant subspaces of the associated Hamiltonian matrices. However, more explicit conditions on the Popov function for the existence of such factorizations would be desirable, and further research is necessary to see whether this can be done. More generally, some basic questions on the factorization of para-Hermitian rational transfer matrices, such as the existence of canonical, but possibly nonproper, factorizations, J-spectral factorizations, connections to algebraic Riccati

[2] In fact, the continuous-time algebraic Riccati equation is much easier to analyze than its discrete-time counterpart and, consequently, significantly more is known about the CARE than the DARE.

equations, or possibly some other form of algebraic equations, are open and worthy of further scrutiny.

We should also mention that the approach of Chapter 7 can be used to study the existence of solutions to algebraic Riccati equations with non-Hermitian coefficient matrices. In this framework, it can be shown that solutions to such equations exist if, and only if, certain proper factorizations of the associated Popov function (which is now a non-para-Hermitian rational transfer matrix) exist. However, it is not clear to this author, at this time, what benefit such an investigation might entail.

Further Study of the Asymptotic Behaviour of the H^∞ Riccati Recursion

Chapter 8 studied the asymptotic behaviour of the Riccati recursion with possibly indefinite coefficient matrices. The main result was that, if certain inertia conditions are satisfied at all time time instants, then the solution to the Riccati recursion (exponentially) converges to the stabilizing solution of the associated DARE (assuming such a solution exists). Although the aforementioned inertia conditions need to be recursively checked, in the special case where the coefficient matrices of the Riccati recursion are positive semi-definite, they can be replaced by a single inertia condition, which results in a very explicit requirement on the initial condition. It would therefore be interesting to study whether it is possible to find other special cases where one can replace the (infinitely many) recursive inertia conditions, required for convergence, by a single inertia condition. In particular, it would be beneficial if one could do so for the special case that arises in H^∞ estimation and control.

H^∞ Adaptive Filtering

Chapter 11 provided a preliminary study of the use of the H^∞ criterion in adaptive filtering. We believe that further study of this approach, and especially issues such as the relationship between robustness and convergence and robustness and tracking (of time-variations), has great merit and may bear considerable fruit. Since its objectives are fully compatible with those of adaptive filtering, the H^∞ approach may, in the long run, compete with least-squares approaches to adaptive filtering in terms of their influence on the field.

12.2 Mixed H^2/H^∞ Estimation and Control

H^∞ optimal algorithms (be it in estimation, control, or signal processing) are most applicable in situations where the exact models and statistics of the underlying disturbances are not available (or are too expensive to obtain). Since they make no assumption about the disturbances, they have to accommodate for all conceivable disturbances, and are thus over-conservative. In many applications one may have some notion of what the statistics of the signals may be, and so one would somehow like to incorporate that knowledge while still guaranteeing the robustness of the algorithm. The mixed H^2/H^∞ formulation to estimation and control is an attempt in this direction.

To illustrate and motivate the approach, consider the adaptive filtering problem of Chapter 9 where the goal was to design an adaptive filtering algorithm for output prediction. Recall, in that case, that the H^2 and H^∞ optimal solutions were given by the RLS and LMS algorithms, respectively. Moreover, let us denote the transfer operators that map the disturbances to prediction errors for each of these algorithms by $\mathcal{T}_{p,rls}$ and $\mathcal{T}_{p,lms}$, respectively. Since RLS and LMS are linear filters, (for finite horizon problems) $\mathcal{T}_{p,rls}$ and $\mathcal{T}_{p,lms}$ will be finite matrices.

Fig. 12.2 shows the (squared) singular values of $\mathcal{T}_{p,rls}$ and $\mathcal{T}_{p,lms}$ for $N = 50$ (where N is the number of observed data points) and $\mu = .9$, for a simple one-dimensional (*i.e.*, single-tap) adaptive filtering problem. As can be seen the (squared) maximum singular value (which is just the maximum energy gain or (squared) H^∞ norm) for $\mathcal{T}_{p,lms}$ is unity, whereas for $\mathcal{T}_{p,rls}$ it is much larger. On the other hand, (under appropriate statistical assumptions) the RLS algorithm has the best average performance (which can be represented as the sum of the squared singular values and is just the area under the singular value curve), whereas the average performance of LMS is significantly worse. Thus the RLS algorithm will have better average performance than LMS, although its worst-case performance is significantly worse.

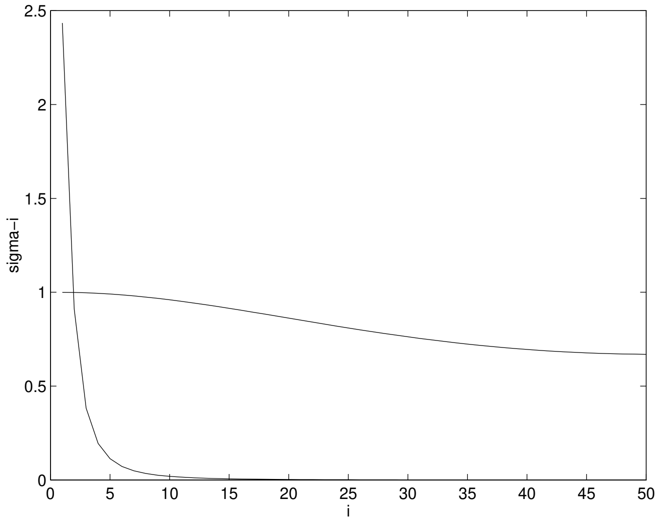

Figure 12.1: Singular values for $\mathcal{T}_{p,rls}$ and $\mathcal{T}_{p,lms}$ for $N = 50$ and $\mu = .9$.

Note, moreover, that although the LMS algorithm does not allow any amplification of the disturbances, it does not provide significant suppression of the disturbances, either. (The smallest squared singular value for $\mathcal{T}_{p,lms}$ which represents the minimum energy gain is roughly 0.65.) Since the H^∞ optimal filters are not unique (LMS is only the central solution), it is very interesting to study the possibility of choosing other H^∞ optimal filters to further reduce the sum of the (squared) singular values of the transfer operator. This will result in algorithms that have the best possible average behaviour, while at the same time having the best possible worst-case performance. This is the essence of the mixed H^2/H^∞ framework.

Problem Formulation

We should note that classical methods in estimation theory (such as least-squares, maximum-likelihood, and maximum entropy) and the more recent robust methods in estimation theory (such as H^∞) can be regarded as two extremes in terms of their requirements regarding the statistical properties of the exogenous signals, as well as in terms of their goals. In classical estimation methods optimality of the average (or expected) performance of the estimator under some assumptions regarding the statistical nature of the signals is the key issue and hence their performance heavily depends upon the validity of these assumptions. On the other hand, robust estimation methods, or so-called minimax estimation strategies, safeguard against the worst-case disturbances and therefore make no assumptions on the (statistical) nature of the signals.

The mixed estimation (control) problem was introduced as a compromise between these two extreme point of views [BH89, YBC92, KR92, LA94, ZGBD94a, ZGBD94b, Meg94, FFT94, FFL95, HHK96]. The mixed H^2/H^∞ problem allows one to trade off between the best average performance of the H^2 estimator (controller) and the best guaranteed worst-case performance of the H^∞ estimator (controller). As a result, the optimal mixed H^2/H^∞ estimators (controllers) achieve the best average performance, not over the set of all estimators (controllers), but over a restricted set of estimators (controllers) that achieve a certain worst case performance bound. We note that the suboptimal (and even optimal) H^∞ estimators (controllers) are highly non-unique and the mixed H^2/H^∞ approach attempts to exploit this non-uniqueness by choosing the one that has the best average performance. Unlike optimal H^2 and suboptimal H^∞ problems, the question of finding the optimal mixed solution is still open.

Recall from Chapter 1 that the objectives in H^2 estimation and control are to solve the following problems,

$$\min_{\text{causal } \mathcal{K}} \|\mathcal{T}_{\mathcal{K}}\|_2 \quad \text{and} \quad \min_{\text{causal } \mathcal{K}} \|\mathcal{T}^c_{\mathcal{K}}\|_2 \tag{12.2.1}$$

where $\mathcal{T}_{\mathcal{K}}$ is the transfer operator from the disturbances to the estimation errors, $\mathcal{T}^c_{\mathcal{K}}$ is the transfer operator from the disturbances to the regulated and control signals, and $\mathcal{K}$ is, depending on the problem, the causal estimator or controller.

Likewise, in (suboptimal) H^∞ estimation and control, for a given $\gamma > 0$, we are required to find causal estimator and controllers that achieve

$$\|\mathcal{T}_{\mathcal{K}}\|_\infty < \gamma \quad \text{and} \quad \|\mathcal{T}^c_{\mathcal{K}}\|_\infty < \gamma. \tag{12.2.2}$$

In view of the above arguments, the mixed H^2/H^∞ problems can be formulated as follows,

$$\left\{ \begin{array}{c} \min_{\text{causal } \mathcal{K}} \|\mathcal{T}_{\mathcal{K}}\|_2 \\ \text{subject to } \|\mathcal{T}_{\mathcal{K}}\|_\infty < \gamma \end{array} \right. \quad \text{and} \quad \left\{ \begin{array}{c} \min_{\text{causal } \mathcal{K}} \|\mathcal{T}^c_{\mathcal{K}}\|_2 \\ \text{subject to } \|\mathcal{T}^c_{\mathcal{K}}\|_\infty < \gamma \end{array} \right. \tag{12.2.3}$$

Recall that $\mathcal{T}_\mathcal{K}$ and $\mathcal{T}_\mathcal{K}^c$ are afine in $\mathcal{K}$. For example, in estimation

$$\mathcal{T}_\mathcal{K} = \left[\begin{array}{cc} \mathcal{L} - \mathcal{K}\mathcal{H} & -\mathcal{K} \end{array} \right], \tag{12.2.4}$$

where $\mathcal{L}$ and $\mathcal{H}$ are known.

The mixed H^2/H^∞ problem (12.2.3) has turned out to be surprisingly difficult, and satisfactory solutions are not yet known. Some recent results include [BH89, YBC92, KR92, LA94, ZGBD94a, ZGBD94b, Meg94, FFT94, FFL95, HHK96]. Some preliminary studies show that there may be some merit in departing from linear estimators and controllers in (12.2.3), and instead trying to solve the mixed problem with (possibly) nonlinear ones [HK96].