



Technical report number 2003-04

A New Autocovariance Least-Squares Method for Estimating Noise Covariances *

Brian J. Odelson, Murali R. Rajamani and James B. Rawlings
Department of Chemical Engineering
University of Wisconsin-Madison
Madison, WI 53706

September 1, 2003
revised: August 18, 2004
re-revised: April 4, 2005

Abstract

Industrial implementation of model-based control methods, such as model predictive control, is often complicated by the lack of knowledge about the disturbances entering the system. In this paper, we present a new method (ALS) to estimate the variances of the disturbances entering the process using routine operating data. A variety of methods have been proposed to solve this problem. Of note, we compare ALS to the classic approach presented in [12]. This classic method, and those based on it, use a three-step procedure to compute the covariances. The method presented in this paper is a one-step procedure, which yields covariance estimates with lower variance on all examples tested. Furthermore, the formulation used in this paper provides necessary and sufficient conditions for uniqueness of the estimated covariances, previously not available in the literature. We show that the estimated covariances are unbiased and converge to the true values with increasing sample size. We also use semidefinite programming to add positive semidefiniteness constraints and maintain a convex program, which none of the classic methods achieve.

⁰This technical report is an expanded version of [19], and is also included in [18]

1 Introduction

Model-based control methods, such as model predictive control (MPC), have become popular choices for solving difficult control problems. Higher performance, however, comes at a cost of greater required knowledge about the process being controlled. Expert knowledge is often required to properly commission and maintain the regulator, target calculator, and state estimator of MPC, for example. This paper addresses the required knowledge for the state estimator, and describes a technique with which ordinary closed-loop data may be used to remove some of the information burden from the user. Consider the usual linear, time-invariant, discrete-time model

$$\begin{aligned}x_{k+1} &= Ax_k + Bu_k + Gw_k \\y_k &= Cx_k + v_k\end{aligned}$$

in which $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $G \in \mathbb{R}^{n \times g}$, $C \in \mathbb{R}^{p \times n}$, and $\{w_k\}_{k=0}^{N_d}$ and $\{v_k\}_{k=0}^{N_d}$ are uncorrelated zero-mean Gaussian noise sequences with covariances Q_w and R_v , respectively. The sequence $\{u_k\}_{k=0}^{N_d}$ is assumed to be a known input resulting from the actions of a controller. State estimates of the system are considered using a linear, time-invariant state estimator

$$\begin{aligned}\hat{x}_{k+1|k} &= A\hat{x}_{k|k} + Bu_k \\ \hat{x}_{k|k} &= \hat{x}_{k|k-1} + L[y_k - C\hat{x}_{k|k-1}]\end{aligned}$$

in which L is the estimator gain. The hat on the state estimate denotes the best linear least-squares estimate of the state using the gain L , which is not necessarily the optimal gain. We denote the residuals of the update equations as the L-innovations when calculated using state estimates calculated using a suboptimal gain. In order to use the optimal filter, we need to know the covariances of the disturbances, Q_w, R_v from which we can calculate the optimal estimator's error covariance and the optimal, Kalman filter gain. In most industrial process control applications, however, the covariances of the disturbances entering the process are not known. To address this requirement, estimation of the covariances from open-loop data has long been a subject in the field of adaptive filtering, and can be divided into four general categories: Bayesian [1, 8], maximum likelihood [5, 10], covariance matching [14], and correlation techniques. Bayesian and maximum likelihood methods have fallen out of favor because of their sometimes excessive computation times. They may be well suited to a multi-model approach as in [2]. Covariance matching is the computation of the covariances from the residuals of the state estimation problem. Covariance matching techniques have been shown to give biased estimates of the true covariances. The fourth category is correlation techniques, largely pioneered by Mehra [12, 13] and Carew and Bélanger [4, 7], which we consider further in this paper.

2 Innovations based correlation techniques and the ALS estimator

With the standard linear state estimator, the state estimation error, $\varepsilon_k = x_k - \hat{x}_{k|k-1}$, evolves according to

$$\varepsilon_{k+1} = \underbrace{(A - ALC)}_{\tilde{A}} \varepsilon_k + \underbrace{\begin{bmatrix} G, & -AL \end{bmatrix}}_{\tilde{G}} \underbrace{\begin{bmatrix} w_k \\ v_k \end{bmatrix}}_{\tilde{w}_k} \quad (3)$$

We define the state-space model of the L-innovations as

$$\begin{aligned} \varepsilon_{k+1} &= \tilde{A}\varepsilon_k + \tilde{G}\tilde{w}_k \\ \mathcal{Y}_k &= C\varepsilon_k + v_k \end{aligned}$$

in which

$$\mathcal{Y}_k = y_k - C\hat{x}_{k|k-1}$$

and we require subsequently that the system is detectable and the chosen filter is stable.

Assumption 1 (A, C) is detectable

Assumption 2 $\tilde{A} = A - ALC$ is stable

A stable filter gain L exists because of Assumption 1. In this formulation, the state and sensor noises are correlated:

$$\begin{aligned} E[\tilde{w}_k(\tilde{w}_k)^T] &\equiv \tilde{Q}_w = \begin{bmatrix} Q_w & 0 \\ 0 & R_v \end{bmatrix} \\ E[\tilde{w}_k v_k^T] &= \begin{bmatrix} 0 \\ R_v \end{bmatrix} \end{aligned}$$

Effect of initial condition. Assume the initial estimate error is distributed with mean m_0 and covariance P_0^- ,

$$E(\varepsilon_0) = m_0 \quad \text{cov}(\varepsilon_0) = P_0^-$$

Propagating the estimate error through the state evolution equation gives an explicit formula for the mean

$$E(\varepsilon_k) = \tilde{A}^k m_0$$

and the recursion for the covariance

$$\begin{aligned} \text{cov}(\varepsilon_k) &= P_k^- \\ P_j^- &= \tilde{A}P_{j-1}^- \tilde{A}^T + \tilde{G}\tilde{Q}_w\tilde{G}^T, \quad j = 1, \dots, k \end{aligned}$$

Because the filter is stable (Assumption 2), as k increases, the mean converges to zero and the covariance approaches a steady state given by the solution to the following Lyapunov equation

$$\begin{aligned} E(\varepsilon_k) &\rightarrow 0 \\ \text{cov}(\varepsilon_k) &\rightarrow P^- \\ P^- &= \bar{A}P^- \bar{A}^T + \bar{G}\bar{Q}_w\bar{G}^T \end{aligned} \quad (5)$$

We therefore assume that we have chosen k sufficiently large so that the effects of the initial conditions can be neglected, or, equivalently, we choose the steady-state distribution as the initial condition

Assumption 3 $E(\varepsilon_0) = 0 \quad \text{cov}(\varepsilon_0) = P^-$

Now consider the autocovariance, defined as the expectation of the data with some lagged version of itself [9],

$$\mathcal{C}_j = E[\mathcal{Y}_k \mathcal{Y}_{k+j}^T] \quad (6)$$

Using Equation 3 and the steady-state initial condition (Assumption 3) gives for the autocovariance

$$E(\mathcal{Y}_k \mathcal{Y}_k^T) = CP^-C^T + R_v \quad (7)$$

$$E(\mathcal{Y}_{k+j} \mathcal{Y}_k^T) = C\bar{A}^j P^- C^T - C\bar{A}^{j-1} ALR_v \quad j \geq 1 \quad (8)$$

which are independent of k because of our assumption about the initial conditions. The autocovariance matrix (ACM) is then defined as

$$\mathcal{R}(N) = \begin{bmatrix} \mathcal{C}_0 & \cdots & \mathcal{C}_{N-1} \\ \vdots & \ddots & \vdots \\ \mathcal{C}_{N-1}^T & \cdots & \mathcal{C}_0 \end{bmatrix} \quad (9)$$

The number of lags used in the ACM is a user-defined parameter, N . The off-diagonal autocovariances are not assumed zero, because we do not process the data with the optimal filter, which is unknown. The ACM of the L-innovations can be written as follows:

$$\begin{aligned} \mathcal{R}(N) &= \mathcal{O}P^- \mathcal{O}^T + \Gamma \left[\bigoplus_{i=1}^N \bar{G}\bar{Q}_w\bar{G}^T \right] \Gamma^T \\ &+ \Psi \left[\bigoplus_{i=1}^N R_v \right] + \left[\bigoplus_{i=1}^N R_v \right] \Psi^T + \bigoplus_{i=1}^N R_v \end{aligned} \quad (10)$$

in which

$$\mathcal{O} = \begin{bmatrix} C \\ C\bar{A} \\ \vdots \\ C\bar{A}^{N-1} \end{bmatrix} \quad \Gamma = \begin{bmatrix} 0 & 0 & 0 & 0 \\ C & 0 & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ C\bar{A}^{N-2} & \cdots & C & 0 \end{bmatrix} \quad \Psi = \Gamma \left[\bigoplus_{j=1}^N (-AL) \right]$$

In this result and those to follow, we employ the standard definitions of the Kronecker product, Kronecker sum and the direct sum [6, 20]. In order to use the ACM relationship in a standard least-squares problem, we apply the “vec” operator, which is the column-wise stacking of a matrix into a vector [6]. If z_k is the k th column of an arbitrary Z matrix

$$\text{vec}(Z) = Z_s = \begin{bmatrix} z_1^T & \cdots & z_k^T \end{bmatrix}^T$$

Throughout this paper, we use the s subscript to denote the outcome of applying the vec operator. Applying the vec operator to Equation 10 and using the result of applying the vec operator on Equation 5

$$P_s^- = (\bar{A} \otimes \bar{A})P_s^- + (\bar{G}\bar{Q}_w\bar{G}^T)_s \quad (11)$$

yields

$$\begin{aligned} [\mathcal{R}(N)]_s &= [(\mathcal{O} \otimes \mathcal{O})(I_{n^2} - \bar{A} \otimes \bar{A})^{-1} \\ &+ (\Gamma \otimes \Gamma)\mathcal{I}_{n,N}](G \otimes G)(Q_w)_s \\ &+ \{[(\mathcal{O} \otimes \mathcal{O})(I_{n^2} - \bar{A} \otimes \bar{A})^{-1} \\ &+ (\Gamma \otimes \Gamma)\mathcal{I}_{n,N}](AL \otimes AL) \\ &+ [\Psi \oplus \Psi + I_{p^2N^2}]\mathcal{I}_{p,N}\}(R_v)_s \end{aligned} \quad (12)$$

in which $\mathcal{I}_{p,N}$ is a permutation matrix to convert the direct sum to a vector, i.e. $\mathcal{I}_{p,N}$ is the $(pN)^2 \times p^2$ matrix of zeros and ones satisfying

$$\left(\bigoplus_{i=1}^N R_v \right)_s = \mathcal{I}_{p,N}(R_v)_s$$

Ideally, we would like to compute the autocovariance as the expectation of the product $\mathcal{Y}_k \mathcal{Y}_{k+j}^T$. Practically, we approximate the expectation from the data using the time average, a valid procedure since the process is ergodic [9]. The estimate of the autocovariance is computed as

$$\hat{\mathcal{C}}_j = \frac{1}{N_d - j} \sum_{i=1}^{N_d - j} \mathcal{Y}_i \mathcal{Y}_{i+j}^T \quad (13)$$

which is the so-called unbiased autocovariance estimator. The estimated ACM, $\hat{\mathcal{R}}(N)$, is analogously defined using the computed $\hat{\mathcal{C}}_j$. At this point we can define a least-squares problem to estimate Q_w, R_v . We summarize Equation 12 as

$$\mathcal{A}x = b$$

in which

$$\begin{aligned} \mathcal{A} &= \left[\begin{array}{c|c} D(G \otimes G) & D(AL \otimes AL) + \\ & + [\Psi \oplus \Psi + I_{p^2N^2}]\mathcal{I}_{p,N} \end{array} \right] \\ D &= [(\mathcal{O} \otimes \mathcal{O})(I_{n^2} - \bar{A} \otimes \bar{A})^{-1} + (\Gamma \otimes \Gamma)\mathcal{I}_{n,N}] \end{aligned} \quad (14)$$

$$\mathbf{x} = \left[(Q_w)_s^T \quad (R_v)_s^T \right]^T \quad \mathbf{b} = \mathcal{R}(N)_s$$

We define the ALS estimate as

$$\hat{\mathbf{x}} = \arg \min_x \|\mathcal{A}\mathbf{x} - \hat{\mathbf{b}}\|_2^2 \quad (15)$$

in which

$$\hat{\mathbf{x}} = \left[(\hat{Q}_w)_s^T \quad (\hat{R}_v)_s^T \right]^T \quad \hat{\mathbf{b}} = \widehat{\mathcal{R}}(N)_s$$

The solution for the ALS estimate is the well known

$$\hat{\mathbf{x}} = \mathcal{A}^\dagger \hat{\mathbf{b}} \quad \mathcal{A}^\dagger = (\mathcal{A}^T \mathcal{A})^{-1} \mathcal{A}^T$$

The uniqueness of the estimate is a standard result of least-squares estimation [11]. The estimated covariances are symmetric due to the structure of the least-squares problem.

Lemma 1 *The ALS estimate (Equation 15) exists and is unique if and only if \mathcal{A} has full column rank.*

We note that the ACM could be written in the output form as well, in which the ACM is computed from the outputs instead of the L-innovations. There are number of reasons why the output covariance estimator is inadequate in practical applications. In the output-based formulation, the control law would have to be specified in order to consider closed-loop data. Additionally, output autocovariance methods are not suitable for estimating integrated white noise disturbances, which are used widely in industrial MPC implementations to remove steady offset.

Comments on the initial filter gain. In principle, any stable filter gain, L , may be used to calculate the L-innovations. This initial gain simply parameterizes the L-innovations. The covariances of the underlying noise sequences are contained in the *outputs* of the process. While the choice of the initial filter may impact the number of data points required to find reliable estimates of the covariances, we show in Section 3 that the initial choice of the filter is irrelevant for large data sets. The autocovariance matrix in Equation 9 has nonzero off-diagonal elements for a suboptimal choice of L . Only when the true covariances (and optimal filter) are employed are the off-diagonal terms zero.

3 Properties of the ALS covariance estimates

In this section we evaluate the mean and variance of the ALS estimator. To this end, first we require the properties of the estimated autocovariance

Lemma 2 *The expectation of the estimated autocovariance ($\hat{\mathcal{C}}_j$) is equal to the autocovariance (\mathcal{C}_j) for all j , and the variance goes to zero inversely with sample size, N_d*

$$E[\hat{\mathcal{C}}_j] = \mathcal{C}_j \quad j = 0, \dots, N$$

$$\text{cov}(\hat{\mathcal{C}}_j) = O\left(\frac{1}{N_d - j}\right)$$

Proof. The expectation result follows from taking expectation of Equation 13 and the definition of the autocovariance, Equation 6. For brevity, the proof of the variance result is omitted. A derivation can be found in [3]. \square

Remark 1 *The unbiased result in finite sample size is due to the strong assumption we have made on the initial conditions, Assumption 3. If we weaken this assumption and allow nonzero expectation of initial error or covariance of initial error not equal to P^- , then the bias is nonzero with finite sample size, but decreases exponentially to zero with increasing sample size.*

Note that for the types of problems to be solved with this method, we choose $N_d \gg N$, and therefore $\text{cov}(\hat{\mathcal{E}}_j) \rightarrow 0$ as $N_d \rightarrow \infty$, for all j . We can choose large N_d because we require only routine operating data, not identification testing data with input excitation. The properties of the autocovariance estimate then imply the ALS estimates of the covariances are unbiased for all sample sizes, and converge to the true values with increasing sample size.

Theorem 1 *Given \mathcal{A} (Equation 14) has full column rank, the ALS noise covariance estimates (\hat{Q}_w, \hat{R}_v) (Equation 15) are unbiased for all sample sizes and converge asymptotically to the true covariances (Q_w, R_v) as $N_d \rightarrow \infty$.*

Proof. For compactness, we use the notation of the least-squares problem of Equation 15 in which $\mathcal{A}x = b$, and

$$\hat{x} = \arg \min_x \|\mathcal{A}x - \hat{b}\|_2^2$$

The expected value of the estimate is

$$\begin{aligned} E[\hat{x}] &= \mathcal{A}^\dagger E[\hat{b}] \\ &= \mathcal{A}^\dagger b \quad (\text{by Lemma 2}) \\ &= \mathcal{A}^\dagger \mathcal{A}x = x \end{aligned}$$

The covariance of the estimate is

$$\text{cov}(\hat{x}) = \mathcal{A}^\dagger \text{cov}(\hat{b}) (\mathcal{A}^\dagger)^T$$

From Lemma 2, $\text{cov}(\hat{b}) \rightarrow 0$ as $N_d \rightarrow \infty$. Therefore

$$\text{cov}(\hat{x}) \rightarrow 0 \text{ as } N_d \rightarrow \infty$$

\square

Remark 2 *Again, as in Lemma 2, the unbiased result in finite sample size is due to Assumption 3. If we remove this assumption, the bias is nonzero with finite sample size, but decreases exponentially to zero with increasing sample size.*

4 Discussion and comparison to previous approaches

4.1 Comparison to correlation based methods

The pioneering work of Mehra [12, 13] and Bélanger [7] has seen successful application using open-loop data and remains highly cited. Mehra employs a three-step procedure to estimate (Q_w, R_v) : (i) Solve a least squares problem to estimate P^-C^T from the estimated autocovariances using Equations 7 and 8. (ii) Use Equation 7 and the estimated P^-C^T to solve for R_v . (iii) Solve a least-squares problem to estimate Q_w from the estimated P^-C^T and R_v using Equation 11. We offer two criticisms of the classic Mehra approach. Our first comment concerns the conditions for uniqueness of (\hat{Q}_w, \hat{R}_v) in Mehra's approach. These conditions were stated (without proof) as

1. (A, C) observable
2. A full rank
3. The number of unknown elements in the Q_w matrix, $g(g + 1)/2$, is less than or equal to np

These conditions were also cited by Bélanger [4]. As a counterexample, consider

$$A = \begin{bmatrix} 0.9 & 0 & 0 \\ 1 & 0.9 & 0 \\ 0 & 0 & 0.9 \end{bmatrix} \quad C = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad G = I$$

The Mehra conditions predict that unique covariances exist, but the \mathcal{A} matrix in Equation 14 does not have full column rank for this case. Thus these conditions are not sufficient. The problem here is that although P^-C^T and R_v are uniquely estimatable from the data, Q_w is not. Examining the null space of the stacked version of the P^-C^T equation shows that any multiple of the following matrix can be added to an estimate of \hat{Q}_w without changing the fit to the autocovariance data

$$Q = \begin{bmatrix} 0.117 & -0.552 & 0 \\ -0.552 & -0.613 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

Consider a second counterexample,

$$A = \begin{bmatrix} 0.1 & 0 \\ 0 & 0.2 \end{bmatrix} \quad G = \begin{bmatrix} 1 \\ 2 \end{bmatrix} \quad C = [1 \quad \delta]$$

When $\delta = 0$, this system is not observable, and thus does not meet Mehra's conditions. But \mathcal{A} has full column rank for $\delta = 0$, the ALS method estimates unique covariances, and thus Mehra's conditions are also not necessary. In this example, one can use just state x_1 to uniquely distinguish the process disturbance from the output disturbance. It makes no difference whether or not the second state is observable.

Our second comment concerns the large variance associated with Mehra's method. This point was first made by Neethling and Young [15], and seems to have been largely overlooked. First, step (ii) above is inappropriate because the zero-order lag autocorrelation estimate in Equation 7 is not known perfectly. Second, breaking a single-stage estimation of Q_w and R_v into two stages by first finding P^-C^T and R_v and then using these estimates to estimate Q_w in steps (i) and (iii) also increases the variance in the estimated Q_w .

To quantify the size of the variance inflation associated with Mehra's method, consider a third example, which has a *well-conditioned* observability matrix

$$A = \begin{bmatrix} 0.1 & 0 & 0.1 \\ 0 & 0.2 & 0 \\ 0 & 0 & 0.3 \end{bmatrix} \quad G = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \quad C = [0.1 \quad 0.2 \quad 0]$$

Data are generated using noise sequences with covariance $Q_w = 0.5, R_v = 0.1$. The L-innovations are calculated with a filter gain corresponding to incorrect noise variances $Q_w = 0.2$ and $R_v = 0.4$. Mehra's method and the ALS method are run using $N_d = 1000$ data points, $N = 15$. The simulation is repeated 200 times to illustrate the mean and variances of the estimators. In Figure 1, the estimates of (Q_w, R_v) using Mehra's method are plotted. The variance of the estimates is large, and many of the estimates are negative, which is aphysical. In Figure 2, the ALS estimates of (Q_w, R_v) are plotted, on much tighter axes. The variance of the ALS estimates is much smaller than in Mehra's method, and none of the estimates are negative. Note that Neethling and Young[15] discuss other examples with behavior similar to this one.

4.2 Enforcing semidefinite constraints

When dealing with a finite set of measurements and inputs from a closed-loop plant, the ALS estimate of the covariances from Equation 15 may not be positive semidefinite, although the variance of the estimates may be smaller than the two-step procedure. Such estimates are physically meaningless. The constrained ALS estimation problem is then given by:

$$\Phi = \min_{Q_w, R_v} \left\| \mathcal{A} \begin{bmatrix} (Q_w)_s \\ (R_v)_s \end{bmatrix} - \hat{\mathbf{b}} \right\|_2^2 \quad (16)$$

subject to, $Q_w \geq 0 \quad R_v \geq 0$

A recent *ad hoc* method of imposing positive semidefiniteness on only the estimates of R_v is given in [17]. The constraints in Equation 16 are convex in Q_w, R_v and the optimization is in the form of a semidefinite programming (SDP) problem [21]. The matrix inequalities $Q_w \geq 0, R_v \geq 0$ can then be handled by adding a logarithmic barrier function to the objective. The optimization in Equation 16 becomes:

$$\Phi = \min_{Q_w, R_v} \left\| \mathcal{A} \begin{bmatrix} (Q_w)_s \\ (R_v)_s \end{bmatrix} - \hat{\mathbf{b}} \right\|_2^2 - \mu \log \begin{vmatrix} Q_w & 0 \\ 0 & R_v \end{vmatrix} \quad (17)$$

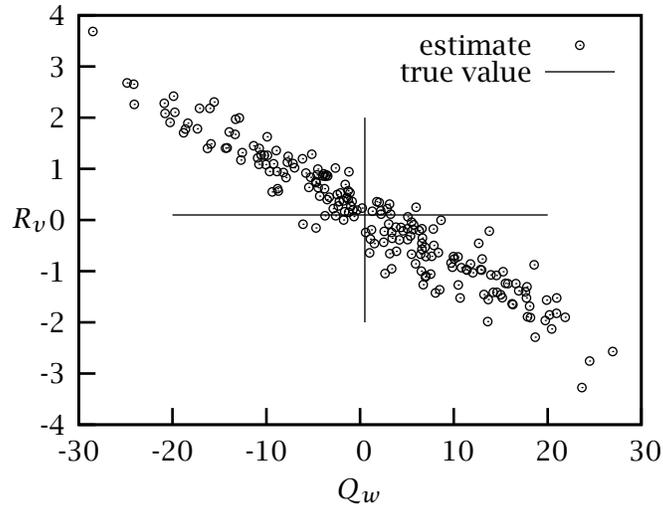


Figure 1: Estimates of Q_w and R_v using Mehra's method.

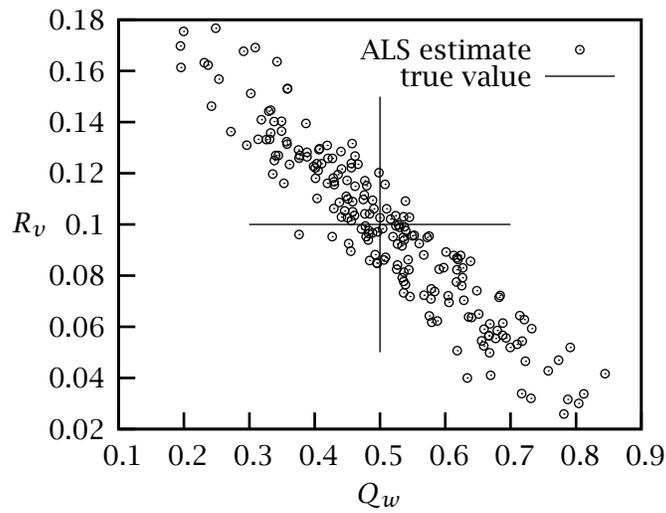


Figure 2: Estimates of Q_w and R_v using proposed ALS method. Notice the axes have been greatly expanded compared to Fig. 1.

in which, μ is called the *barrier parameter* and $|\cdot|$ denotes the determinant of the matrix [16]. The optimization in Equation 17 is convex with analytical gradients. A simple path following algorithm based on Newton steps is then enough to find the global optimum. The details of the algorithm can be found in [22, chap. 10]. The convexity of the optimization in Equation 16 and Lemma 1 ensure uniqueness of the covariance estimates. The algorithm generalizes efficiently for large dimension problems.

5 Conclusions

In this paper we have developed a new method (ALS) for using the autocovariance of the L-innovations to estimate the covariances of the system disturbances, which are required for optimal state estimation. We have shown that the ALS covariance estimates are unbiased and converge asymptotically to the true system covariances with increasing sample size. Using the standard properties of least-squares estimation, necessary and sufficient conditions for unique covariance estimates have been provided. Examples are provided to show that conditions previously stated in the literature are neither necessary nor sufficient for uniqueness of these estimates. Previously reported methods for this estimation have been shown to have unnecessarily large variance. We also add semidefiniteness constraints on the covariances by solving a convex semidefinite programming problem.

6 Acknowledgments

Financial support came from the industrial members of the Texas-Wisconsin Modeling and Control Consortium and NSF through grant #CTS-0105360. The authors thank the anonymous reviewers for many helpful comments and Professor John F. MacGregor for useful discussions.

References

- [1] D. Alspach. A parallel filtering algorithm for linear systems with unknown time-varying noise statistics. *IEEE Trans. Auto. Cont.*, 19(5):552-556, 1974.
- [2] A. Averbuch, S. Itzikowitz, and T. Kapon. Radar target tracking - viterbi versus IMM. *IEEE Trans. Aero. Elec. Sys.*, 27(3):550-563, 1991.
- [3] M. Bartlett. On the theoretical specification of sampling properties of autocorrelated time series. *J. Roy. Stat. Soc. Suppl.*, 8:27-41, 1946.
- [4] P. Bélanger. Estimation of noise covariance matrices for a linear time-varying stochastic process. *Automatica*, 10:267-275, 1974.

- [5] T. Bohlin. Four cases of identification of changing systems. In R. Mehra and D. Lainiotis, editors, *System Identification: Advances and Case Studies*. Academic Press, 1st edition, 1976.
- [6] J. Brewer. Kronecker products and matrix calculus in system theory. *IEEE Trans. Circuits and Systems*, 25(9):772-781, 1978.
- [7] B. Carew and P. Bélanger. Identification of optimum filter steady-state gain for systems with unknown noise covariances. *IEEE Trans. Auto. Cont.*, 18(6):582-587, 1973.
- [8] C. Hilborn and D. Lainiotis. Optimal estimation in the presence of unknown parameters. *IEEE Trans. Systems, Science, and Cybernetics*, 5(1):38-43, 1969.
- [9] G. Jenkins and D. Watts. *Spectral Analysis and its Applications*. Holden-Day, 500 Sansome Street, San Fransisco, California, 1968.
- [10] R. Kashyap. Maximum likelihood identification of stochastic linear systems. *IEEE Trans. Auto. Cont.*, 15(1):25-34, 1970.
- [11] C. L. Lawson and R. J. Hanson. *Solving Least Squares Problems*. SIAM, Philadelphia, PA, 1995.
- [12] R. Mehra. On the identification of variances and adaptive Kalman filtering. *IEEE Trans. Auto. Cont.*, 15(12):175-184, 1970.
- [13] R. Mehra. Approaches to adaptive filtering. *IEEE Trans. Auto. Cont.*, 17:903-908, 1972.
- [14] K. Myers and B. Tapley. Adaptive sequential estimation with unknown noise statistics. *IEEE Trans. Auto. Cont.*, 21:520-523, 1976.
- [15] C. Neethling and P. Young. Comments on "Identification of optimum filter steady-state gain for systems with unknown noise covariances". *IEEE Trans. Auto. Cont.*, 19(5):623-625, 1974.
- [16] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer-Verlag, New York, 1999.
- [17] G. Noriega and S. Pasupathy. Adaptive estimation of noise covariance matrices in real-time preprocessing of geophysical data. *IEEE Trans. Geosci. Remote Sensing*, 35(5):1146-1159, 1997.
- [18] B. J. Odelson. *Estimating Disturbance Covariances From Data for Improved Control Performance*. PhD thesis, University of Wisconsin-Madison, 2003.
- [19] B. J. Odelson, M. R. Rajamani, and J. B. Rawlings. A new autocovariance least-squares method for estimating noise covariances. Submitted for publication in *Automatica*, April 2005.

- [20] S. Searle. *Matrix algebra useful for statistics*. John Wiley & Sons, New York, 1982.
- [21] L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM Rev.*, 38(1):49–95, March 1996.
- [22] H. Wolkowicz, R. Saigal, and L. Vandenberghe, editors. *Handbook of Semidefinite Programming: Theory, Algorithms and Applications*. Kluwer Academic Publishers, Massachusetts, 2000.