Design of an Extended Kalman Filter Frequency Tracker

Barbara F. La Scala and Robert R. Bitmead

Abstract— **The design of an extended Kalman filter for tracking a time-varying frequency is discussed. Its principal modes of failure are explained. The design tradeoff between balancing noise rejection and tracking at a maximal slew rate is discussed. The performance penalties for overdesign and underdesign of noise covariances are examined, and theoretically supported design guidelines are suggested.**

I. Introduction

The problem of determining the frequency of a signal is a common one; however, its inherent nonlinearity has also made it difficult. There is a long history of the application of extended Kalman filter (EKF) methods in frequency estimation and tracking going back at least to Snyder [1]. The analysis of design and performance has, however, been lacking. This correspondence examines a common EKF frequency tracker and considers the stability of its associated error system and what this indicates about design. The issue of stability of the errors of the EKF is the key point in any EKF design. Unlike the linear Kalman filter where general theoretical results have long been available, most EKF designs have relied on heuristic arguments to attempt to ensure stability. Previous approaches have been hampered by the lack of knowledge of the dynamics of the errors of the EKF. As a consequence, the design guidelines developed have only been applicable to the particular formulation considered and do not necessarily extend to other cases.

Recently, Song and Grizzle [2] have given conditions under which the EKF will be a locally asymptotic observer when applied to any undriven signal model. La Scala *et al.* [3] have extended this result to the case of general stochastic signals. The stability result in [3] quantifies the stability of the EKF in terms of the degree of nonlinearity of the system, noise covariances matrices, and bounds on the noise processes. These new results provide the first theoretical analysis of EKF performance for general nonlinear systems. The analysis of error dynamics provides insight into the behavior of EKF, which can be used to provide design guidelines for a particular application.

In this work, we focus on the choice of appropriate covariance matrices to balance noise rejection with tracking at a maximal slew rate for the frequency tracking problem. These choices are not straightforward. The nature of the performance penalty for overspecification and underspecification of noise covariances is shown. Simple extension of filter design methods for linear Kalman filters are shown to be inappropriate. Performance of the EKF tracker is illustrated via simulation results.

Manuscript received June 16, 1994; revised September 17, 1995. This work was supported by the Co-operative Research Centre for Robust and Adaptive Systems ($(\mathrm{CR})^2$ASys) and by the Australian Commonwealth Government under the Co-operative Research Centre Program. The associate editor coordinating the review of this paper and approving it for publication was Dr. Thomas F. Quatieri.

B. F. La Scala is with the Co-operative Research Centre for Sensor Signal and Information Processing, Department of Electrical and Electronic Engineering, University of Melbourne, Parkville, Australia (e-mail: bfls@ee.mu.oz.au).

R. R. Bitmead is with the Co-operative Research Centre for Robust and Adaptive Systems, Department of Systems Engineering, Research School of Information Sciences and Engineering, Australian National University, Canberra, Australia (e-mail: bob@syseng.anu.edu.au).

Publisher Item Identifier S 1053-587X(96)02385-7.

II. Signal Model

Kalman filter design and, by implication, EKF estimator design proceeds from a state space signal model of the process to be estimated. The signal model dynamics describe a mechanism for how the process may be evolving. Thus, the initial stage in the filter design process is to perform system identification. This process will not be considered here. Once the system has been identified, Kalman filter theory endeavors to construct an optimal estimator for the state, given the noise covariances Q and R, where optimality is measured in terms of error covariance. For linear Kalman filtering, the best design choice is to use the signal model determined by the system identification process and its associated estimated noise covariances. In the case of the EKF, the inclination is to use the same procedure. We will show here that this is, in fact, not necessarily a suitable design method for nonlinear signals.

In our EKF design, the signal model is a construct rather than an exact description of the measurement source. In this case, the matrices Q and R can no longer be regarded as the same values as the noise covariances driving the nonlinear system. Our focus here will be on the appropriate selection of design Q and R values.

The following nonlinear signal model describes the evolution of noisy quadrature data with a slowly time-varying frequency, where

$\{y(k)\}$ two-vector received signal,
x_3 unknown time-varying frequency,
x_1 and x_2 in-phase and quadrature signals.

$$\begin{bmatrix} x_1(k+1) \\ x_2(k+1) \\ x_3(k+1) \end{bmatrix} = \begin{bmatrix} \cos x_3(k) & -\sin x_3(k) & 0 \\ \sin x_3(k) & \cos x_3(k) & 0 \\ 0 & 0 & 1-\epsilon^a \end{bmatrix} \begin{bmatrix} x_1(k) \\ x_2(k) \\ x_3(k) \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ w_3(k) \end{bmatrix} \quad (1)$$

$$\begin{bmatrix} y_1(k) \\ y_2(k) \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} x_1(k) \\ x_2(k) \\ x_3(k) \end{bmatrix} + \begin{bmatrix} v_1(t) \\ v_2(t) \end{bmatrix}. \quad (2)$$

The parameter $\epsilon^a \in (0,1)$ determines the rate of time variation of x_3 and is chosen so that the frequency varies slowly enough that the signal appears periodic over several cycles. The signals $\{v(k)\}$ and $\{w(k)\}$ are zero mean, independent noise processes with $E[w(k)w(k)^T] = Q^a$ and $E[v(k)v(k)^T] = R^a$, $Q^a \geq 0$, and $R^a > 0$. The frequency of the signal x_3 represents the state that we wish to recover. The x_1 and x_2 components are noiseless transformations of the x_3 component, and hence, all the entries of Q^a are zero except for $Q^a(3,3) = q^a$.

This simple model has properties that make it useful for designing filters for quasiperiodic signals generated from a variety of applications. The Cartesian formulation of the model makes it close to linear, thus reducing errors due to linearization effects when filtering. Furthermore, in the absence of prior knowledge of systematic variation in the frequency of the signal (other than that it is slowly varying), modeling the variation in frequency via a random walk is a reasonable choice, as this simple model can be tuned in a straightforward manner to allow for varying rates of change via appropriate design choices for Q^a and ϵ^a. These choices will depend on estimates of the true signal's signal-to-noise ratio (SNR) and slew rate.

Note that while it is not possible to achieve global observability for any formulation of the discrete-time frequency tracking problem due to aliasing effects, the system given by (1) and (2) is strongly locally observable [3], [1]. Since it is possible to resolve the state locally without ambiguity in the deterministic case of our formulation, it is reasonable to expect that an observer could be constructed for this model in the stochastic environment that could successfully track the frequency of the received signal.

III. EKF Observer

The EKF is derived by linearizing the signal model about the current predicted state estimate and then using the Kalman filter on this linearized system to calculate a gain matrix. This gain matrix, along with the nonlinear signal model and new signal measurement, is used to produce the filtered state estimate and, then, an estimate of the state at the next time instant. The equations for the EKF can be found in [4, ch. 6]. The EKF gain matrix $K(k)$ is calculated according to the Kalman filter equations for an associated linearized system whose noise processes are a combination of the original noise processes in the nonlinear signal model w and v and the neglected terms in the Taylor series expansion of the nonlinear signal model.

When constructing our EKF filter for the frequency tracking problem, we will denote the values used for the AR model and noise covariances by ϵ^d, Q^d, and R^d. Note that these values are design parameters and are not necessarily equal to the values that would be obtained via system identification, although the usual implication (which we seek to dispel here) would have them so. When choosing the filter parameters Q^d and ϵ^d, it is necessary to allow not only for the expected maximum slew rate of the frequency but also linearization errors. The choice of R^d and Q^d will affect the stability properties and rate of response of the filter. If these design issues are neglected and these values are set to the values suggested by the signal's SNR and slew rate, the resulting filter may have undesirable properties.

The design compromises in developing extended Kalman filters (and indeed Kalman filters) based on signal models are to balance 1) filter divergence, where the filter becomes increasingly confident of increasingly bad estimates through the computed error covariance becoming too small, leading to loss of tracking capabilities and 2) sensitivity to noise, where the parameter estimate fails to reject enough of the measurement noise process. These effects can be tied to questions of the magnitude of the Kalman gain $K(k)$ being either too small or too large, respectively. In the following sections, we will discuss the effect of suitable choices of Q^d, ϵ^d, and R^d and how these problems might be avoided.

IV. Choice of the Design Parameters

For linear Kalman filters, optimality and dual optimal control arguments can be used to determine the effect of varying choices of Q^d and R^d [4], [5]. It can be shown that if (Q^d, R^d) are selected to be greater than the actual values, then the achieved performance is bounded above by the performance computed as the solution of the Riccati equation. Underdesign possesses no such guarantees, and there is no bound on the level of the actual state error covariance. Accordingly, overdesign is to be preferred with the Kalman filter and, hence, with the EKF. For the EKF, there are still reasons to pursue even larger Q^d.

An issue to consider is that for the EKF, the solution of the Riccati equation is only a first-order approximation to the true error covariance. It has been shown [6] that this approximation is an underestimate, and the error between the first- and second-order approximations has been derived. This illustrates how the linearization itself introduces effects that increase the effective Q value. Thus, the associated linearized system used to calculate the gain is more "noisy" than the original nonlinear system.

A. Choice of Q^d

The choice of Q^d is crucial to the performance of any EKF observer. From the equation for $P(k+1|k)$, it can be seen that the minimum value of the solution of this Riccati equation is given by Q^d. Too small a value of Q^d leads to overconfidence in the accuracy of the estimates ($P(k+1|k)$ too small) and, consequently, to the filter paying insufficient attention to new data ($K(k)$ too small). This will cause filter divergence.

In the past, more precise statements on the appropriate choice of Q^d have not been possible. The recent result of [3] provides the following insights. If Q^d is chosen to be positive definite, then the controllability Gramian of $[F, Q^d]$ will be bounded and positive definite. It is shown in [3] that this property is one of a set of sufficient conditions for the errors of the EKF for any nonlinear system to be bounded, provided the initial error is small enough. In addition, from the theorems shown in [3], the following result for the frequency tracking problem can be derived.

Lemma 1: For all matrices of the form $Q^d = diag(q_1^d, q_1^d, q_2^d)$, where $0 < q_1^d \leq q_2^d < \infty$, $Q^d = q_2^d I$ maximizes the size of the region of admissible initial errors. □

B. Tradeoff Between Q^d and ϵ^d

For the frequency tracking problem, a key design issue is that of the maximum rate of change of the frequency (slew rate) that the filter can track. As we wish to be able to track either an increasing or decreasing frequency, we measure the slew rate with the variance of the frequency. For the signal model (1)–(2), the variance of the frequency is given by $\sigma^2 = \frac{q_2^d}{1-(1-\epsilon^d)^2}$, and therefore, the maximal slew rate of the signal that the EKF can track is determined by q_2^d and ϵ^d. Hence, a large Q^d allows the EKF to track a signal with a potentially large slew rate. The drawback of a large Q^d is that even when the EKF is tracking well, if the slew rate is low, the filter is overly sensitive to noise, causing estimates to fluctuate widely around the true frequency value. Setting $\epsilon^d < \epsilon^a$ alleviates this problem by allowing the value of Q^d to be decreased while still retaining the desired range for the frequency estimate. This, along with the design of Q^d and R^d, is related to measures to achieve guarantees of the degree of stability of Kalman filters [4, sect. 6.2].

C. Choice of R^d

The choice of R^d for the frequency tracking problem is less critical than that of Q^d. The tracking ability of the EKF in this case is relatively insensitive to the value of R. Instead, the value of R^d affects the degree of variation in the state estimates. This can be shown by considering once again the dual optimal control problem. It is well known in the control literature, [7, ch. 6] for example, that increasing the cost of control (i.e., increasing R^d) reduces the feedback gain and the speed of response of the controlled system. Since the feedback gain in the control problem is equal to $-K(k)^T$, reducing the magnitude of the feedback gain is equivalent to reducing the gain in the estimation problem as well. If control is cheap (i.e., R^d is small), this increases the speed of response of the system by allowing higher feedback gains.

V. Simulation Results

The following simulation results illustrate the importance of following the design guidelines discussed in the previous section. For these simulations, the signal was not generated by the model used to

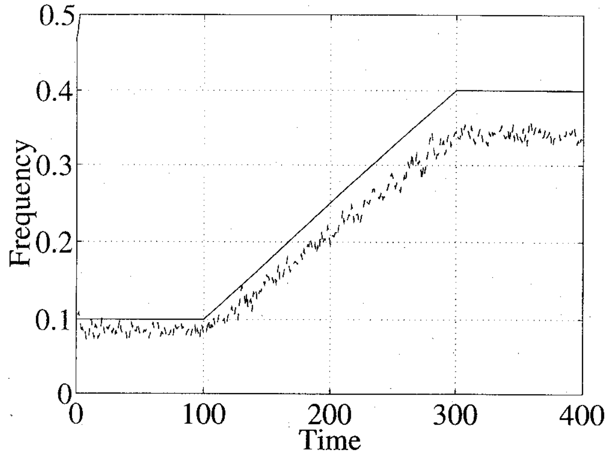

Fig. 1. EKF frequency estimate when $Q^d = Q^a$, $\epsilon^d = \epsilon^a$, and $R^d = R^a$. The target signal is given by the solid line.

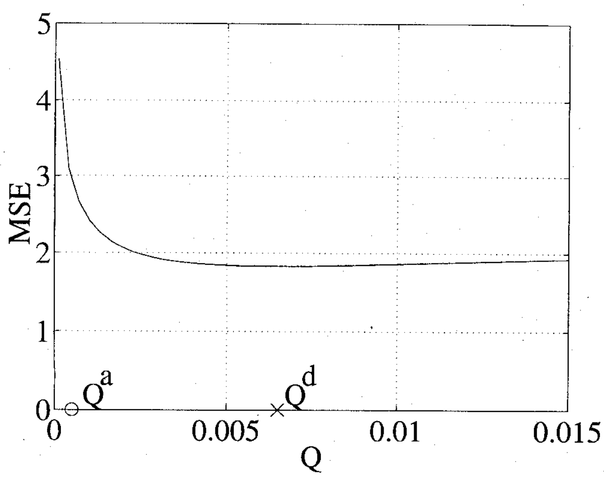

Fig. 2. MSE of EKF frequency estimate versus Q^d.

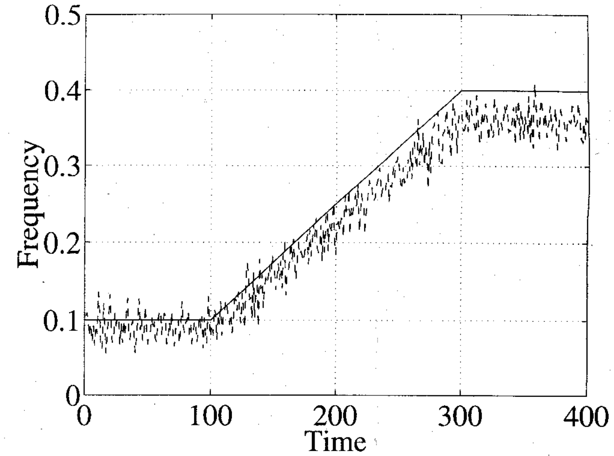

Fig. 3. EKF frequency estimate when $Q^d > Q^a$, $\epsilon^d = \epsilon^a$, and $R^d = R^a$. The target signal is given by the solid line.

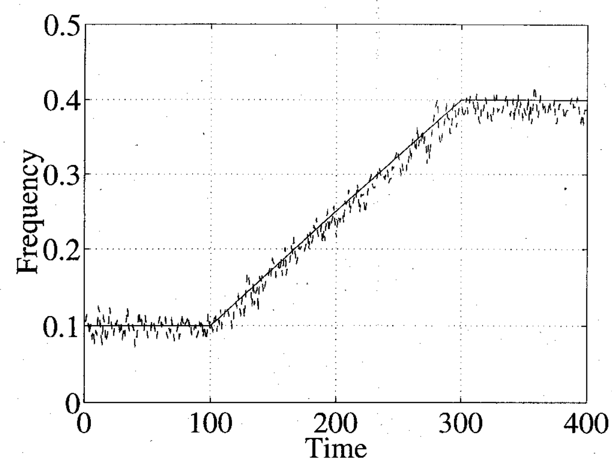

Fig. 4. EKF frequency estimate when $Q^d > Q^a$, $\epsilon^d < \epsilon^a$, and $R^d = R^a$. The target signal is given by the solid line.

construct the EKF tracker. This was to emphasize the fact that the signal model is used purely for the construction of the filter.

In the first set of examples, the true signal had an SNR of 17 dB, and the rate of change in frequency varied in the range [0, 0.04]. Such a signal could be generated by the signal model with the parameters $Q^a = 0.0005$, $\epsilon^a = 0.05$, and $R^a = 0.01$. Fig. 1 shows the effect of designing an EKF frequency tracker using these effective actual values. A common assumption is that using these design values would yield good performance. As predicted by [3] and illustrated in Fig. 1, this is not the case. Mean squared tracking error of the frequency versus Q^d is plotted in Fig. 2. The value of Q^a in the nonlinear signal model is shown as is the best Q^d value. It shows the large penalty for too small a Q^d and the modest degradation for too large a value. That higher Q^d does permit more accurate tracking is shown by Fig. 3, where $Q^d = 0.0065$ as the linearization errors in the associated linearized signal model are now taken in account.

It is clear from Fig. 3 that the EKF with $Q^d > Q^a$, $\epsilon^d = \epsilon^a$, and $R^d = R^a$ is still not ideal. The filter remains unable to track when the slew rate is at its maximum, and the degree of variability in the state estimates is large. Setting $\epsilon^d = 0.01$ and then adjusting Q^d down to $Q^d = 0.0013$ to maintain the same maximum slew rate permits the filter to maintain a similar tracking performance with less sensitivity to noise. This is illustrated in Fig. 4. Finally, setting $R^d = 0.05$ and increasing Q^d to $Q^d = 0.0023$ produces a filter that is able to track the frequency with less peaking. This is illustrated in Fig. 5. The MSE for this filter was less than one tenth of that for the filter using the effective actual values.

In addition to the previous, single example, further simulations were performed and compared with the performance of an adaptive notch filter (ANF) [8] for the same underlying signal and four different SNR's. For each SNR, the experiment was repeated 1000 times with an initial error in the frequency estimate varying uniformly with an error of up to ±20%. Table I summarizes the results for the ANF and EKF with an optimal choice of parameters and the EKF using the effective actual signal values. It gives the mean and standard deviation of the mean squared error in the frequency estimate.

VI. Conclusion

In this correspondence, we have given theoretically supported guidelines for the design of stable extended Kalman filters, and this has been applied to an EKF frequency tracker. We have discussed the

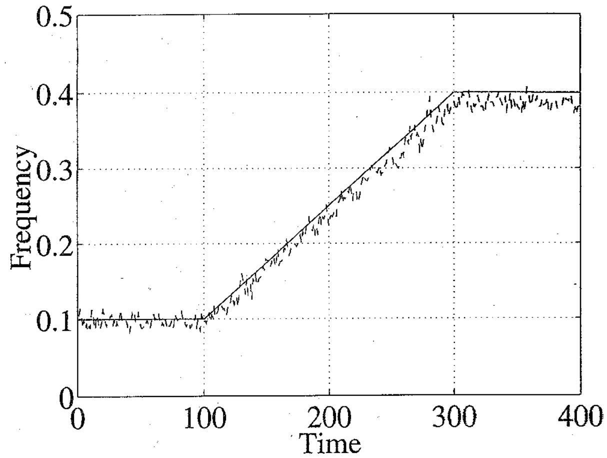

Fig. 5. EKF frequency estimate when $Q^d > Q^a$, $\epsilon^d < \epsilon^a$, and $R^d > R^a$. The target signal is given by the solid line.

TABLE I
MEAN AND STANDARD DEVIATION OF FREQUENCY ESTIMATE ERROR

SNR	ANF - "best"		EKF - "best"		EKF - "actual"	
	mean(MSE)	std(MSE)	mean(MSE)	std(MSE)	mean(MSE)	std(MSE)
17dB	0.4185	0.0342	0.3252	0.0255	2.9740	0.0329
13dB	0.4497	0.0534	0.4493	0.0290	3.9001	0.0765
10dB	0.5287	0.0506	0.4904	0.0374	4.7783	0.1572
7dB	0.6361	0.0908	0.7268	0.0517	25.782	17.868

design tradeoffs between balancing noise rejection with tracking at a maximal slew rate. The nonlinearity inherent in this problem makes this tradeoff nonobvious, critical to the performance of the filter, and counter to the usual optimality guidelines for Kalman filtering. The performance penalties for overestimation and underestimation of the noise covariances are further illustrated via simulation results. These demonstrate the importance of designing a sufficiently conservative filter and might explain why so few successful applications of extended Kalman filters have been reported.

REFERENCES

[1] D. L. Snyder, *The State-Variable Approach to Continuous Estimation with Applications to Analog Communications Theory.* Boston, MA: MIT Press, 1969.

[2] Y. Song and J. W. Grizzle, "The extended Kalman filter as a local asymptotic observer for nonlinear discrete-time systems," *J. Math. Syst., Estimation Contr.*, vol. 5, no. 1, pp. 59–78, 1995.

[3] B. F. La Scala, R. R. Bitmead, and M. R. James, "Conditions for stability of the extended Kalman filter and their application to the frequency tracking problem," *Math. Contr., Signals Syst.*, vol. 8, no. 1, pp. 1–27, 1996.

[4] B. D. O. Anderson and J. B. Moore, *Optimal Filtering.* Englewood Cliffs, NJ: Prentice-Hall, 1979.

[5] R. R. Bitmead, M. Gevers, and V. Wertz, *Adaptive Optimal Control: The Thinking Man's GPC.* Sydney, Australia: Prentice-Hall, 1990.

[6] A. H. Jazwinski, *Stochastic Processes and Filtering Theory.* London: Academic, 1970.

[7] B. D. O. Anderson and J. B. Moore, *Optimal Control.* Englewood Cliffs, NJ: Prentice-Hall, 2nd ed., 1990.

[8] S.-C. Pei and C.-C. Tseng, "Adaptive IIR notch filter based on least mean p-power error criterion," *IEEE Trans. Circuits Syst.–II*, vol. 40, no. 8, pp. 525–529, 1993.

Analysis and Optimization of Subset Averaged Median Filters

Dong Hee Kang, Jongkwan Song, and Yong Hoon Lee

***Abstract*—By analyzing the subset averaged median (SAM) filter based on threshold decomposition, we show that the class of SAM filters is identical to the class of extended threshold Boolean filters (ETBF's) with the *extended self-dual* property. This result indicates that the class of SAM filters encompasses a variety of digital filters such as linear finite impulse response (FIR), weighted median, symmetric L-filters, and any filter defined by a linear combination of these filters. A procedure for determining an optimum SAM filter in the mean square error (MSE) sense is developed. It is shown that the optimization of SAM filters may result in a FIR Wiener filter when the input is Gaussian and in a median-type filter for non-Gaussian inputs.**

I. INTRODUCTION

The subset averaged median (SAM) filter is a multistage digital filter [1], [2] containing median subfilters. In this filter, the final output is a weighted average calculated over the outputs of median subfilters. SAM filters[1] were introduced in [3] and [4] as an extension of median filters.

In this correspondence, we shall show that the class of ETBF's [5], [6] with an *extended self-dual* property is identical to the SAM class filter. This result indicates that the class of SAM filters encompasses a variety of filters such as linear FIR, weighted median [7], symmetric L-filters [8], and any filter defined by a linear combination of these filters. In addition, we shall develop a procedure for determining the best SAM filter in the mean square error (MSE) sense. It will be observed through computer simulation that the optimization of SAM filters may result in an optimal FIR Wiener filter when the input is zero-mean Gaussian and in a nonlinear filter that outperforms the Wiener filter for non-Gaussian inputs.

II. THE SAM FILTER AND THE ETBF

In this section, we review the definitions of SAM filters and ETBF's, and introduce some subclasses of these filters such as *full*-SAM filters and extended self-dual ETBF's.

A. The SAM Filter

Consider a nonrecursive filter that evaluates its output from the input sequence $\boldsymbol{X}(n) = (X_1(n), X_2(n), \cdots, X_N(n))$ taken from the window at time n where $X_j(n)$ is the jth input sample from the left of the window and N is the window size. We define the ith median subfilter denoted by $F_i(\boldsymbol{X})$ as follows:

$$F_i(\boldsymbol{X}) = \text{med}(\boldsymbol{X}_i) \qquad (1)$$

where $\boldsymbol{X}_i$ is a subsequence of $\boldsymbol{X}$ and $F_i(\boldsymbol{X})$ takes the median value of the inputs in $\boldsymbol{X}_i$. Here and in the rest of this correspondence, the time index n is dropped from $\boldsymbol{X}(n)$ and $X_j(n)$ to simplify the notation. It is assumed that the number of inputs in $\boldsymbol{X}_i$ is an odd

Manuscript received August 12, 1994; revised September 21, 1995. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Thomas F. Quatieri.

The authors are with the Department of Electrical Engineering, Korea Advanced Institute of Science and Technology, Taejon, Korea (e-mail: yohlee@eekaist.kaist.ac.kr)

Publisher Item Identifier S 1053-587X(96)02386-9.

[1] SAM filters in [3] calculate the simple average of median subfilter outputs.