GAUSSIAN MIXTURE KALMAN PREDICTIVE CODING OF LSFS

Shaminda Subasingha†, *Manohar N. Murthi*†, *and Søren Vang Andersen*‡

† Dept. of Electrical and Computer Engineering
University of Miami, USA
s.subasingha@umiami.edu, mmurthi@miami.edu

‡ Department of Electronic Systems
Aalborg University, Denmark
sva@es.aau.dk

ABSTRACT

Gaussian Mixture Model (GMM)-based predictive coding of line spectral frequencies (lsf's) has gained wide acceptance. In such coders, each mixture of a GMM can be interpreted as defining a linear predictive transform coder. In this paper we optimize each of these linear predictive transform coders using Kalman predictive coding techniques to present GMM Kalman predictive coding. In particular, we show how suitable modeling of quantization noise leads to an adaptive a-posteriori GMM that defines a signal-adaptive predictive coder that provides superior coding of lsfs in comparison with the baseline GMM predictive coder. Moreover, we show how running the Kalman predictive coders to convergence can be used to design a stationary predictive coding system which again provides superior coding of lsfs but now with no increase in run-time complexity over the baseline.

***Index Terms*—** speech coding, vector quantization, Kalman filtering, Gaussian Mixture Models

1. INTRODUCTION

In recent years, Gaussian Mixture Model-based vector quantization (VQ) of speech signal parameters has received much attention. With a GMM providing an empirical probability density function (pdf) of vector data such as Line Spectral Frequencies (LSFs), various VQ methods are possible. In particular, [1] utilized high-rate theory to investigate recursive GMM-based VQ, while [2] presented a low-complexity recursive GMM-based VQ, in which each component Gaussian can be interpreted as a predictive transform coder.

Given that Kalman filtering principles [3] have previously been utilized within predictive coding of speech (e.g., [4, 5]), a natural question is whether Kalman filtering principles can be utilized within the context of GMM-based VQ. In this paper, we present GMM Kalman predictive coding (GMM-KF) of lsfs. With Kalman filtering providing a framework for quantization noise modeling and error covariance matrix updating, we demonstrate how an a-posteriori GMM based on past signal measurements can be utilized to perform predictive coding. In particular, the GMM-KF's ability to adapt the coding model to the previous signal measurements leads to better lsf modeling performance than that provided by a standard GMM recursive coder (GMM-RC) or memoryless coder. Moreover, we demonstrate a method for obtaining a low-complexity stationary GMM Kalman predictive coder (stationary GMM-KF) which also provides better lsf modeling performance.

The rest of this paper is organized as follows: Section 2 considers a state space model of GMM coding; Section 3 presents both the online and stationary GMM-KF; Section 4 presents simulations; Section 5 presents conclusions.

2. STATE SPACE MODEL OF A GMM CODER

Let X_k be a d-dimensional source vector at the time instant k, and $Y_{k,p}$ be the concatenation of p previous source vectors, i.e., $Y_{k,p} = [X'_{k-1}, ..., X'_{k-p}]'$.

Given a joint GMM of X_k and $Y_{k,p}$ one can write the conditional GMM of X_k given $Y_{k,p}$ as $p_{X|Y}(x|y) = \sum_{i=1}^{L} \alpha_i p_i(x|y)$. The component pdf $p_i(x|y)$ has the conditional mean $\mu_{i,k|k-1} = m_{i,X} + C_{i,XY}\, C^{-1}_{i,YY}(y_{k-1} - m_{i,Y})$, and conditional covariance $\mathcal{Q}_i = C_{i,XX} - C_{i,XY}\, C^{-1}_{i,YY} C_{i,YX}$ [2]. The conditional mixture probabilities α_i are given by, $\alpha_i = \frac{\rho_i p_i(y)}{p_Y(y)}$. In the GMM-RC, each Gaussian conditional density $p_i(x|y)$ defines a transform coder, with each coder competing to produce the best quantized value of a vector X_k given p past quantized values $Y_{k,p}$. A standard GMM-RC updates its a-posteriori coder probabilities α_i, but does not adapt the $p_i(x|y)$ themselves. To utilize Kalman filtering principles, which engender quantization noise modeling and coder adaptation, we now take each $p_i(x|y)$, and define a linear plant model, and measurement model which captures the transform coding process. Utilizing standard Kalman filtering notation [3], let us define the state vector x_k which consists of p source vectors, given by $x_k = \begin{bmatrix} X_k & . & . & . & X_{k-p+1} & 1 \end{bmatrix}'$. Given x_k, for each $p_i(x|y)$, we associate a linear model

$$x_{k+1} = F_i x_k + w_i \tag{1}$$

$$z_{i,k} = H_{i,k} x_k + v_{i,k}, \tag{2}$$

The state prediction matrix F_i of the i^{th} coder is selected such that the predicted state becomes the conditional mean

This work was supported in part by the National Science Foundation via Awards CCF-0347229, and CNS-0519933.

$\mu_{i,k|k-1}$ of the i^{th} coder, i.e.,

$$\begin{bmatrix} C_{i,XX}C^{-1}_{Y_Y\,(dxdp)} & m_{i,X} - C_{i,XX}\times \\ & C_{i,YY}m_{i,Y\,(dx1)} \\ I_{(d(p-1)xd(p-1))} \quad 0_{(d(p-1)xd)} & 0_{(d(p-1)x1)} \\ 0_{(1xdp)} & 1_{(1x1)} \end{bmatrix}. \quad (3)$$

The fixed '1' as the last element of x_k is used to accommodate the fixed mean components in $p_i(x|y)$. Noting that w_i represents the plant noise of the linear model (1). we define the covariance $E[w'_{i,k}, w_{i,k'}] = Q_i\delta_{k,k'}$, we set the top left $d \times d$ submatrix of Q_i to be $\lceil Q_i \rceil_{d\times d} = \mathcal{Q}_i$ with the other elements of Q_i set to zero.

Suppose that the i^{th} coder at time $k-1$ has a state vector $\hat{x}_{i,k-1|k-1}$. To predict the signal state x_k at time k based on the information at time $k-1$, the i^{th} coder forms $\hat{x}_{i,k|k-1} = F_i\hat{x}_{i,k-1|k-1}$. The elements of the prediction error $x_k - \hat{x}_{i,k|k-1}$ are transformed and quantized, and further processed to create the i^{th} coder's best representation of the current state. This is modeled by (2) in which the i^{th} coder observes noisy measurements of the true state x_k. To see this, we define the prediction error covariance matrix

$$\Sigma_{i,k|k-1} = E([x_k - \hat{x}_{i,k|k-1}]'[[x_k - \hat{x}_{i,k|k-1}]]) \quad (4)$$

which models the prediction error statistics. To be consistent with the standard GMM-RC, suppose that only the first d elements of $x_k - \hat{x}_{i,k|k-1}$ corresponding to the error between the current $d \times 1$ source vector and the prediction are quantized. Then denote the top left $d \times d$ submatrix of $\Sigma_{i,k|k-1}$ as $\lceil \Sigma_{i,k|k-1} \rceil_{d\times d}$. We perform an eignendecomposition $\lceil \Sigma_{i,k|k-1} \rceil_{d\times d} = U_{i,k}\Lambda_{i,k}U'_{i,k}$, and set $H_{i,k} = [U'_{i,k}\ \mathbf{0}_{d\times(d(p-1)+1)}]$. Then observe that $H_{i,k}(x_k - \hat{x}_{i,k|k})$ is the Karhunen-Loeve transform (KLT) of the first d values of the i^{th} coder's state prediction error at time k. Let the i^{th} coder's KLT be $\psi_{i,k} = H_{i,k}(x_k - \hat{x}_{i,k|k-1})$ while $\hat{\psi}_{i,k} = \psi_{i,k} + v_{i,k}$ is the quantized KLT where $v_{i,k}$ is the quantization noise.

Now reconsider the measurement equation (2). Through quantization, the i^{th} coder is observing noisy measurements of the true state x_k. Let us relate this to the standard GMM-RC. Suppose the i^{th} coder denotes its quantized state vector at time k as $\hat{x}_{i,k|k}$. Then see that in the standard GMM-RC, the i^{th} coder's candidate quantization of the source vector is given by $\hat{x}_{i,k|k} = \hat{x}_{i,k|k-1} + H'_{i,k}\hat{\psi}_{i,k}$ which we can re-write as $\hat{x}_{i,k|k} = \hat{x}_{i,k|k-1} + H'_{i,k}(z_{i,k} - H_{i,k}\hat{x}_{i,k|k-1})$. Then in a standard GMM-RC, each coder forms its competitive candidate $\hat{x}_{i,k|k}$, with the best candidate coder i^* selected, and the corresponding quantization information transmitted. Then $\hat{x}_{i^*,k|k}$ is used to update a common state $x_{k|k}$ which is utilized by each coder in its prediction step. In the standard GMM-RC, except for the coder conditional probabilities, the coder model parameters are not updated.

From a Kalman filtering perspective, the standard GMM-RC can be improved in several respects. In particular, by appropriately modeling the statistics of the quantization noise $v_{i,k}$ in the measurement equation, one can form the filtered state $\hat{x}_{i,k|k} = \hat{x}_{i,k|k-1} + K_{i,k}(z_{i,k} - H_{i,k}\hat{x}_{i,k|k-1})$ by utilizing the Kalman gain $K_{i,k}$. Moreover, one can utilize the noisy measurement $z_{i,k}$ to form the filtered error covariance matrix $\Sigma_{i,k|k}$ which is subsequently utilized to form $\Sigma_{i,k+1|k} = F_i\Sigma_{i,k|k}F'_i + Q_i$, which in turn is utilized to update the KLT matrix $H_{i,k+1}$. Therefore, the normal prediction, measurement, filtering, update cycle of a Kalman filter can be utilized to update the coder parameters of the L models based on the noisy measurements. This leads to more accurate tracking of the signal evolution and consequently more effective quantization. Let us consider how to apply Kalman filtering principles to this state-space model.

3. GMM-KALMAN PREDICTIVE CODER

Consider coding with the a-posteriori conditional GMM

$$p(x_k|Z^{ab}_k) = \sum_{j=1}^{L} P(M_{j,k}|Z^{ab}_k)p(x_k|M_{j,k}, Z^{ab}_k), \quad (5)$$

where $Z^{ab}_k = [z^{ab}_1, \cdots, z^{ab}_k]$ denotes the set of common *absolute* measurements available to all coders, $P(M_{j,k}|Z^{ab}_k)$ denotes the probability of coder j at time k given the past measurements Z^{ab}_k. Then $p(x_k|M_{j,k}, Z^{ab}_k)$ is the j^{th} coder's conditional Gaussian pdf of x_k given Z^{ab}_k, with mean $\hat{x}_{j,k|k}$, and covariance $\Sigma_{j,k|k}$. Given this conditional GMM, each conditional Gaussian $p(x_k|M_{j,k}, Z^{ab}_k)$ defines a particular transform coder. The GMM-KF follows a 1) competition 2)update cycle. In the competition cycle, each transform coder forms a candidate quantization of the current source vector, and in the update cycle, both the selected and non-selected coder parameters are updated using a common absolute measurement related to the selected coder's measurement. First, let us consider the competition in step 1.

The key for each coder is to model $v_{i,k}$. Suppose we let $R_{i,k} = E[v_{i,k}, v'_{i,k}]$. Recall that with $H_{i,k}$ obtained from an eigendecomposition of $\Sigma_{i,k|k-1}$, the KLT of the prediction error, $\psi_{i,k} = H_{i,k}(x_k - \hat{x}_{i,k|k})$, is quantized. Rather than assuming the quantization scheme of [2], let us assume that each element of the vector $\psi_{i,k}$ is divided by the square root of the corresponding eigenvalue from $\Lambda_{i,k}$, i.e., we form a vector $\Gamma_{i,k}$ with row l consisting of $\gamma^l_{i,k} = [\psi^l_{i,k}]/\sqrt{\lambda^l_{i,k}}$. Then $\Gamma_{i,k}$ can be viewed as a vector of uncorrelated Gaussian random variables with unit variance.

Consequently, each element of $\Gamma_{i,k}$ is independently quantized by a Max scalar quantizer [6] with element l quantized using $b^l_{i,k}$ bits. Then the quantization noise covariance is written as $E[v_{i,k}v'_{i,k}] = \Delta_{i,k}H_{i,k}\Sigma_{i,k|k-1}H'_{i,k}$. The noise factor matrix $\Delta_{i,k}$ is a diagonal matrix with elements $((\sigma^1_{i,k})^2, \cdots, (\sigma^d_{i,k})^2)$ where the terms $\sigma^l{}_{i,k}$ is determined by the noise factor of a $b^l_{i,k}$ bit Max quantizer [6].

Given this modeling of the quantization noise, how are Kalman filtering recursions performed? If we assume that in the quantization of $\Gamma_{i,k}$, that the quantization noise values along each dimension are uncorrelated, then by the centroid condition of scalar quantizers, we can show that

$$E(\psi_{i,k}v'_{i,k}) = -E(v_{i,k}v'_{i,k}) \tag{6}$$

With this correlated noise assumption, the candidate Kalman Gain and filtered error covariance update equations are [5]

$$K_{i,k} = \Sigma_{i,k/k-1}H'_{i,k}(H_{i,k}\Sigma_{i,k/k-1}H'_{i,k})^{-1} \tag{7}$$

$$\Sigma_{i,k|k} = \Sigma_{k|k-1} - K_{i,k}(I - \Delta_{i,k})H_{i,k}\Sigma_{i,k|k-1}. \tag{8}$$

Therefore, each coder forms candidate quantized values $\hat{x}_{i,k|k} = \hat{x}_{i,k|k-1} + K_{i,k}(z_{i,k} - H\hat{x}_{i,k|k-1})$. The coder i^*, which minimizes $\|x_k - \hat{x}_{i,k|k}\|^2$ (or alternatively log spectral distortion (LSD)) is chosen as the selected coder. Thus, only the quantization information of the selected coder i^* is transmitted. Then the selected coder can form its predicted covariance matrix $\Sigma_{i^*,k+1|k}$ from its $\Sigma_{i^*,k|k}$ in (8), and prepare to enter the competition to quantize x_{k+1} by computing $H_{i^*,k+1}$. However, the coders $j \neq i^*$ that were not selected cannot use the updates in (7) and (8), as these equations are posited on the coder j actually being selected. Therefore, we need to consider two problems: (1) How are the non-selected coder parameters of $p(x(k)|M_{j,k}, Z_k^{ab})$ updated? (2) How are the coder probabilities $P(M_{j,k}|Z_k^{ab})$ updated?

3.1. Updating Non-Selected Coders

The selected coder i^* observes the measurement $z^*_{i^*,k} = {H^*}_{i^*,k}x_k + v_{i^*,k}$ at time k within the frame of reference of its transform coder. To proceed, we convert $z^*_{i^*,k}$ into an *absolute* frame of reference by forming the common absolute measurement z_k^{ab} observed by all the non-selected coders,

$$z_k^{ab} = {H^*}'_{i^*,k}z^*_{i^*,k} = x_k + {H^*}'_{i^*,k}v_{i^*,k}. \tag{9}$$

Then the $j \neq i^*$ coder can calculate its $K_{j,k}$ and $\Sigma_{j,k|k}$ through z_k^{ab}, in which we view ${H^*}'_{i^*,k}v_{i^*,k}$ as an uncorrelated noise [5] since $v_{*,k}$ is from coder $i^* \neq j$. Now suppose that each $j \neq i^*$ converts z_k^{ab} to its own frame of reference by forming local measurements

$$\tilde{z}_{j,k} = H_{j,k}z_k^{ab} = H_{j,k}(x_k + {H^*}'_{i^*,k}v_{i^*,k}). \tag{10}$$

We can easily show that for coder $j \neq i^*$ the use of its local measurement $\tilde{z}_{j,k}$ instead of z_k^{ab} results in the same $\Sigma_{j,k|k}$ and $K_{j,k}$. To find these parameters, we follow the standard procedures in [3]. After finding the joint conditional statistics of $[x'_k, \tilde{z}_{j,k}]$, we can use properties of conditional Gaussian random variables to find

$$K_{j,k} = \Sigma_{j,k|k-1}H'_{j,k}(\Omega_{j,k})^{-1}$$
$$\Sigma_{j,k|k} = \Sigma_{j,k|k-1} - \Sigma_{j,k|k-1}H'_{j,k}(\Omega_{j,k})^{-1}H_{j,k}\Sigma_{j,k|k-1}$$
$$\Omega_{j,k} = H_{j,k}\Sigma_{j,k|k-1}H'_{j,k} + H_{j,k}{H^*}'_{i^*,k}R_{i^*,k}{H^*}_{i^*,k}H'_{j,k}$$

Then $\hat{x}_{j,k|k} = \hat{x}_{j,k|k-1} + K_{j,k}(\tilde{z}_{j,k} - H_{j,k}\hat{x}_{j,k|k-1})$ results in a different filtered state for each of the L coders.

3.2. Updating coder probabilities

Now consider updating $P(M_{j,k}|Z_k^{ab})$. By Bayes' rule,

$$p(x(k)|Z_k^{ab}) = \sum_{j=1}^{L} p(x(k)|M_{j,k}, Z_k^{ab}) \times \frac{p(z_k^{ab}|Z_{k-1}^{ab}, M_{j,k})P(M_{j,k}|Z_{k-1}^{ab})}{\sum_{i=1}^{L} p(z_k^{ab}|Z_{k-1}^{ab}, M_{i,k})P(M_{i,k}|Z_{k-1}^{ab})} \tag{11}$$

For the likelihood value $p(z_k^{ab}|Z_{k-1}^{ab}, M_{j,k})$, we can show

$$p(z_k^{ab}|Z_{k-1}^{ab}, M_{j,k}) = \mathcal{N}[\tilde{z}_{j,k} - H_{j,k}\hat{x}_{j,k|k-1}, \Omega_{j,k}] \tag{12}$$

where $\mathcal{N}[\tilde{z}_{j,k} - \varrho, \Xi]$ denotes a Gaussian pdf with mean ϱ and covariance Ξ. Now, by substituting (12) into (11) we get,

$$p(x(k)|Z_k^{ab}) = \sum_{j=1}^{L} p(x(k)|M_{j,k}, Z_k^{ab}) \times \frac{\alpha_{j,k-1}\,\mathcal{N}[\tilde{z}_{j,k} - H_{j,k}\hat{x}_{j,k|k-1}, \Omega_{j,k}]}{\sum_{j=1}^{L}\alpha_{j,k-1}\,\mathcal{N}[\tilde{z}_{j,k} - H_{j,k}\hat{x}_{i,k|k-1}, \Omega_{j,k}]}$$

Defining $\alpha_{j,k} = P(M_{j,k}|Z_k^{ab})$, we then see that

$$\alpha_{j,k} = \frac{\alpha_{j,k-1}\,\mathcal{N}[\tilde{z}_{j,k} - H_{j,k}\hat{x}_{j,k|k-1}, \Omega_{j,k}]}{\sum_{j=1}^{L}\alpha_{j,k-1}\,\mathcal{N}[\tilde{z}_{j,k} - H_{j,k}\hat{x}_{i,k|k-1}, \Omega_{j,k}]}. \tag{13}$$

To avoid coder probabilities from becoming unrealistically skewed, we adopt a rescaling strategy in which we set $\alpha_{j,k} = \rho_j$, (the initial probabilities) if the smallest coder probability is less than a constant, i.e., $\alpha_{j_{min},k} < \alpha_{min}$.

Therefore, in the online GMM-KF, each model i uses its filtered state $\hat{x}_{i,k-1|k-1}$ to form a candidate quantization of the current source vector. The selected coder i^* transmits the quantization information, and the corresponding model decoder i^* forms the noisy measurement which is used with the Kalman recursions to form $\hat{x}_{i^*,k|k}$. The non-selected coders at both the sender and destination update their models as illustrated above, and then the $\alpha_{j,k-1}$ values are updated, along with the bit allocations that follow [2]. Therefore, the online GMM-KF allows for an a-posteriori GMM to adaptively track the signal evolution, leading to more effective quantization.

3.3. Stationary GMM Kalman Predictive Coder

The online GMM-KF is computationally expensive in that for each iteration, each coder must perform eigendecompositions to obtain parameters for the KLT and bit allocations. However, for a fixed total bit-rate, we can run the online GMM-KF over a representative test file for a number of iterations (1000 in this paper) until the Kalman gain and covariance matrices for each coder converge to stationary values, leading to a stationary GMM-KF. In contrast to the online coder, the stationary coder utilizes a fixed KLT since the error covariance matrix is fixed. Only the a-posteriori coder probabilities and the new bit-allocations need to be calculated in each iteration. Therefore, the stationary GMM-KF is essentially of the same computational complexity as the standard GMM-RC.

Table 1. GMM-KF Coder performance comparison (lossless channel)

	GMM-KF		GMM-RC		NRC	
Bit	SNR	LSD	SNR	LSD	SNR	LSD
20	39.98	0.8296	38.65	0.9254	37.54	1.1020
25	42.63	0.5731	40.94	0.6756	40.11	0.8290
30	45.23	0.4222	44.64	0.4732	42.99	0.6012
35	48.02	0.3209	47.64	0.3643	45.71	0.4401

Table 2. Stationary GMM-KF Coder performance comparison (Lossless Channel)

	GMM-KF		GMM-RC		NRC	
Bit	SNR	LSD	SNR	LSD	SNR	LSD
20	39.63	0.8315	38.65	0.9254	37.54	1.1020
25	42.46	0.6052	40.94	0.6756	40.11	0.8290
30	45.16	0.4443	44.64	0.4732	42.99	0.6012
35	47.93	0.3207	47.64	0.3643	45.71	0.4401

Table 3. Stationary GMM-KF Coder performance comparison (channel loss 5% without feedback)

	GMM-KF		GMM-RC		NRC	
Bit	SNR	LSD	SNR	LSD	SNR	LSD
20	33.99	1.3938	30.33	1.8056	33.42	1.3609
25	34.73	1.1605	29.22	1.8731	33.42	1.1551
30	35.16	0.9706	28.96	1.7974	34.09	0.9410
35	35.38	0.8652	27.63	1.9518	32.81	0.8788

4. SIMULATIONS

We compare the performance of both the online and stationary GMM-KFs to a baseline GMM-RC as well as a memoryless GMM non-recursive coder (NRC) in lsf quantization. A 16 mixture GMM was trained on 400,000, dimension $d = 10$ lsf vectors from the DARPA TIMIT database. Testing was performed on a separate set of 30,000 lsf vectors. For both predictive coders, $p = 1$, meaning that the state vector is only of dimension d.

4.1. Online GMM Kalman Predictive Coder

Table 1 shows performance under lossless channel conditions, with both LSD and SNR shown. In the figures, one can see that the online GMM-KF consistently provides both better LSD and SNR performance than either the GMM-RC or NRC cases. The improvement in SNR is noteworthy, as the GMM Kalman predictive coding principles described in this paper are applicable to a wide variety of signals.

4.2. Stationary GMM-Kalman Predictive Coder

The comparisons between the stationary GMM-KF and the baseline GMM-RC and NRC can be seen in Table 2 for a lossless channel. The stationary GMM-KF provides better performance both in terms of SNR and LSD than the baseline coders. This is particularly noteworthy as the stationary coder has roughly the same complexity as the baseline GMM-RC.

Table 4. Stationary GMM-KF Coder performance comparison (channel loss 10% without feedback)

	GMM-KF		GMM-RC		NRC	
Bit	SNR	LSD	SNR	LSD	SNR	LSD
20	31.97	1.7136	27.99	2.4975	31.47	1.6117
25	32.40	1.4990	27.32	2.6045	31.12	1.4686
30	32.59	1.3525	26.95	2.5704	31.56	1.2627
35	32.72	1.2653	25.93	2.8444	30.37	1.2570

The Stationary GMM-KF and the baseline coders are tested with packet loss conditions without feedback. With packet loss, the GMM-KF propagates its predictions as the measurements, and combines all the coder measurements for packet loss concealment (PLC). The baseline coders utilize suitably combined predictions from the RC models to perform PLC. The comparisons are given in Tables 3 and 4 for packet loss rates of 5%and 10% respectively. In both cases, the stationary GMM-KF outperformed the baseline GMM-RC as the GMM-KF's a-posteriori coder probabilities are conditioned on all measurements, while the GMM-RC probabilities are only conditioned on the last quantized value. The GMM-KF and the GMM NRC provided nearly the same performance in terms of LSD for the packet loss case, with the NRC providing a very slight advantage. However, in terms of SNR, the stationary GMM-KF provided better performance than the NRC. Therefore, the stationary GMM-KF presents attractive lsf quantization abilities at low complexity.

5. CONCLUSION

In this paper, we have shown how suitable modeling of quantization noise within a state-space framework leads to an adaptive a-posteriori GMM that defines a signal-adaptive predictive coder that provides superior coding of lsfs in comparisons with a baseline GMM-RC. Moreover, we show how one can obtain a stationary GMM-KF coder which also provides better coding of lsfs than a regular GMM-RC without any increase in complexity.

6. REFERENCES

[1] J. Samuelsson and P. Hedelin, "Recursive coding of spectrum parameters," *IEEE Trans. on Speech and Audio Proc.*, vol. 9,(5), pp. 492 – 503, 2001.

[2] W.R. Gardner A.D. Subramaniam and B.D. Rao, "Low complexity recursive coding of spectrum parameters," ICASSP '02, 2002, vol. 1, pp. 637–640.

[3] B.D.O. Anderson and J.B.Moore, *Optimal filtering*, Prentice-Hall, 1979.

[4] Gibson J.D. Fischer, T.R. and B. Koo, "Estimation and noisy source coding," *IEEE Trans. on Acoustics, Speech, and Signal Proc.*, vol. 38,(1), pp. 23–34, 1990.

[5] S.H. Jensen S.V. Andersen and E. Hansen, "Quantization noise modeling in low-delay speech coding," IEEE Speech Coding Workshop, Sept. 1997, pp. 65 – 66.

[6] J. Max, "Quantizing for minimum distortion," *IEEE Trans. on Info. Theory*, vol. 6,(1), pp. 7–12, 1960.