Fluctuation and Noise Letters
Vol. 5, No. 3 (2005) L457–L468
© World Scientific Publishing Company

QUANTIZATION IN THE PRESENCE OF LARGE AMPLITUDE THRESHOLD NOISE

MARK D. MCDONNELL
Centre for Biomedical Engineering (CBME) and
School of Electrical & Electronic Engineering
The University of Adelaide, SA 5005, Australia
mmcdonne@eleceng.adelaide.edu.au

NIGEL G. STOCKS
School of Engineering
The University of Warwick, Coventry, CV4 7AL, UK
ngs@eng.warwick.ac.uk

CHARLES E. M. PEARCE
School of Mathematical sciences
The University of Adelaide, SA 5005, Australia
cpearce@maths.adelaide.edu.au

DEREK ABBOTT
Centre for Biomedical Engineering (CBME) and
School of Electrical & Electronic Engineering
The University of Adelaide, SA 5005, Australia
dabbott@eleceng.adelaide.edu.au

Received 1 June 2005
Revised 29 June 2005
Accepted 6 July 2005
Communicated by Riccardo Mannella

Signal quantization in the presence of independent, identically distributed, large amplitude threshold noise is examined. It has previously been shown that when all quantization thresholds are set to the same value, this situation exhibits a form of stochastic resonance known as *suprathreshold stochastic resonance.* This means the optimal quantizer performance occurs for a small input signal-to-noise ratio. Here we examine the performance of this stochastic quantization in terms of both mutual information and mean square error distortion. It is also shown that for low input signal-to-noise ratios that the case of all thresholds being identical provides the optimal mean square error distortion performance for the given noise conditions.

Keywords: Stochastic resonance; suprathreshold stochastic resonance; quantization; distortion; rate distortion; threshold noise; jitter; mutual information.

1. Introduction

Theoretical models of quantization usually consider only the case where ideal thresholding can be performed. Here, we examine the performance of scalar quantization when all thresholds are subject to *iid*, large amplitude noise, so that the thresholds become random variables. For a scalar quantizer with B output bits, $N = 2^B - 1$ thresholds are required. In the absence of noise, we denote the set of thresholds of a scalar quantizer as $\{\theta_n\}, n = 0, \ldots, N$. In the presence of *iid* zero-meaned threshold noise, with an even probability density function (PDF) given by $R(\eta)$, the nth threshold becomes a random variable with PDF $\Theta_n(\eta) = R(\eta - \theta_n)$ and mean θ_n.

Under such conditions, the model of the noisy quantizer can be described as follows. Assume a stationary signal source with PDF $P_x(x)$. This source is then quantized by N binary threshold elements, all of which are subject to *iid* additive input noise, η_n, with PDF $R(\eta)$. The output of each threshold element is denoted by y_n, where $y_n = 0.5 \text{ sign}(x + \eta_n) + 0.5$ so that $y_n \in \{0, 1\}$. The overall output of the quantizer is then $y = \sum_{n=0}^{N} y_n$. Thus, y is a discretely valued stochastic encoding of x, such that $y \in \{0, \ldots, N\}$. This model is shown in Fig. 1.

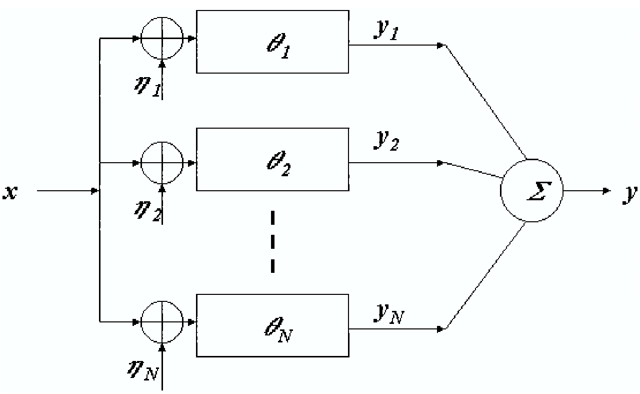

Fig. 1. Model of the stochastic quantizer. A sample, x, from the source distribution, $P_x(x)$, is quantized by N threshold elements, all of which are subject to *iid* additive noise, η_n. The output from the nth threshold, y_n, is unity if the sum of x and η_n is greater than the threshold value, θ_n, and zero otherwise. The overall output, y, is the sum of the y_ns.

In the absence of noise, this model reduces to the familiar scalar quantizer model [1], if it is assumed that the thresholds are distributed across the signal distribution such that $\theta_1 < \theta_2 < \cdots < \theta_n$, in which case the set, $\{y_n\}$, provides a binary encoding of y.

Here however, although we also consider y as the number of thresholds elements that are "on," we do not require the set $\{y_n\}$ to provide a binary encoding. This is because if we discard the ordering of the threshold outputs and only consider the overall output value, y, to be the encoding, it turns out that effective quantization can still occur even if all thresholds are very noisy, and are all set to the same value. The fact that such a randomized quantization can be effective is not necessarily an intuitive result, especially when the source of the randomness is thought of as noise.

Such a situation was first demonstrated and described mathematically in 2000, and shown to be a form of stochastic resonance (SR), labeled as suprathreshold stochastic resonance (SSR) [2]. We shall see that in terms of an encoding, while far from the quality of the optimal encoding provided by noiseless scalar quantizers for the same N, this model does in fact provide an effective quantization. We have also previously shown that should a conventional scalar quantizer be subject to the same large amplitude threshold noise, the quantizer with all thresholds set equal to the same value provides superior performance for the same N [3].

2. Suprathreshold Stochastic Resonance

The initial formulation of the system described in Fig. 1 (with all thresholds equal to the same value) was motivated by studies of SR in threshold-based systems. SR is the term used to describe systems in which the presence of input or internal noise provides the optimal output of that system [4–8]. Although the term "stochastic resonance" was originally used to refer to the very specific case of nonlinear systems driven by periodic input signals, subject to additive white noise, and performance measured by the output signal-to-noise ratio (SNR), the name is now applied very broadly to any system in which some non-zero level of noise can provide a performance improvement. Examples of systems in which SR has been described include a Schmitt Trigger circuit [9], ring lasers [10], neurons [11], SQUIDS [12] and ion channels [13]. The literature on SR also contains many studies of systems consisting of a single threshold device (see for example [14–16]). SR occurs in such systems when the addition of noise to a subthreshold signal causes an output signal to occur that has some correlation with the input signal. In the absence of noise, there would be no output signal.

By contrast, the term SSR was coined for the multi-threshold system of Fig. 1, since SR effects occur regardless of whether the input signal is subthreshold or suprathreshold [2]. In contrast to dithering in analog-to-digital conversion — where a small amplitude noise signal is added to the input prior to thresholding [15, 17, 18] — all thresholds are set to the same value, and the noise signal is allowed to be very large. More importantly, unlike dithering, all thresholds are required to be subject to independent noise signals. This independence is the key feature that allows SSR to occur. In keeping with SR conventions, the system was analyzed by calculating the variation in a measure of system performance as the internal noise level increases from zero. Due to the input signal being taken to be *iid* samples from the distribution with PDF $P_x(x)$, mutual information was a natural measure to use. It was shown numerically that the maximum mutual information through the system occurred for the thresholds set equal to the signal mean, and some non-zero value of the noise variance, σ_η^2 [19].

To illustrate the means of calculating the mutual information, we commence by deriving a method of calculating the joint input–output PDF, $P_{xy}(x, n)$, for the general case of N arbitrary thresholds, and then simplifying to the SSR case, when all thresholds are equal to the signal mean. Now, $P_{xy}(x, n) = P(n|x)P_x(x)$. Integration with respect to x gives the probability that $y = n$, $P_y(n) = \int_{-\infty}^{\infty} P(n|x)P_x(x)dx$.

Assuming knowledge of $P_x(x)$ we can derive a method for calculating the transition probabilities, $P(n|x)$. Let $\hat{P}_n(x)$ be the probability of threshold element n being "on" (that is, signal plus noise exceeding the threshold θ_n), given the input signal x. Then

$$\hat{P}_n(x) = \int_{\theta_n - x}^{\infty} R(\eta) d\eta = 1 - F_R(\theta_n - x), \tag{1}$$

where F_R is the cumulative distribution function of the noise. In general, it is difficult to find an analytical expression for $P(n|x)$ and we will therefore rely on numerics. Given any arbitrary N, $R(\eta)$ and $\{\theta_n\}$, $\{\hat{P}_n(x)\}$ can be calculated exactly for any value of x from Eq. (1), from which $P(n|x)$ can be found using an efficient recursive formula [20]. For the particular case where the thresholds all have the same value, each $\hat{P}_n(x)$ has the same value $\hat{P}(x)$ and, as noted in [2] we have $P(n|x)$ given by the binomial distribution,

$$P(n|x) = C_n^N \hat{P}^n(x)(1 - \hat{P}(x))^{N-n} \quad (0 \le n \le N).$$

The mutual information is that of a semi-continuous channel,

$$I(x, y) = -\sum_{n=0}^{N} P_y(n) \log_2 P_y(n) - \left(- \int_{-\infty}^{\infty} P_x(x) \sum_{n=0}^{N} P(n|x) \log_2 P(n|x) dx \right),$$

which can be calculated by numerical integration after applying the technique for calculating $P(n|x)$ described above.

It has also been demonstrated that in the case of the signal and noise having the same distribution, the mutual information is a function of the ratio of noise variance, σ_η^2, to the signal variance, σ_x^2 which we denote as σ^2. Examples of the variation of mutual information with σ are shown in Fig. 2 for a Gaussian signal and Gaussian noise. Instead of plotting the variation of the mutual information with σ as done in [2], here we have plotted against input SNR in decibels (dB), which is $-20 \log_{10}(\sigma)$. Note that the optimal mutual information occurs for a very low input SNR, and tends towards an SNR of zero as N increases.

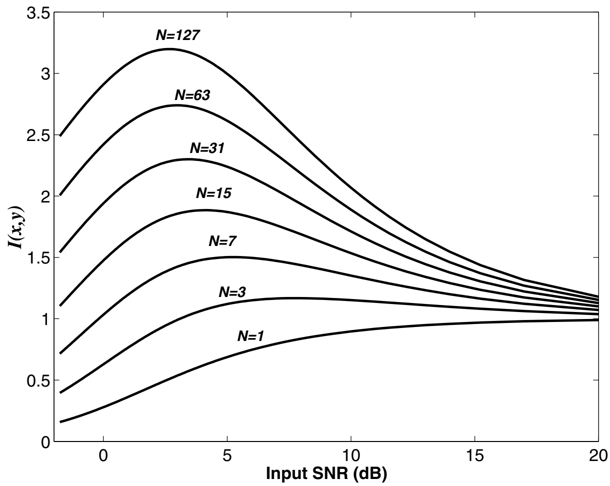

Fig. 2. Plot of mutual information against input SNR (i.e. $-20 \log_{10} \sigma$), for a zero mean Gaussian source with $\sigma_x = 1$, and Gaussian noise, with all thresholds set equal to the source mean of zero.

The reason for this is that in the absence of noise, only one bit per sample is available, since all thresholds are either on or off, and the output signal will only provide an indication of whether the input is above or below its mean. However, in the presence of the *iid* noise, all thresholds become random variables, and for any given signal sample will all have unique values. The distribution of these thresholds depends on the noise intensity. For small noise intensity most thresholds will still be close to the signal mean. Thus, although an increase in information will occur at the output, virtually no information about the intensity of the input will be available. However, as the noise intensity increases, the average threshold distribution will widen more and more, and the mutual information will increase, since on average, the output will provide a better indication of the intensity of the input signal. Eventually, as the noise gets too large, the threshold distribution will get wider than the signal dynamic range, and the mutual information will start to decrease. Thus, the maximum mutual information occurs for a quite small input SNR. That the optimal SNR decreases as N increases can be explained by noting that an increasing N means that more thresholds can, on average, be set to relatively large or small values without a loss in the number of thresholds closer to the signal mean. Hence, a larger noise intensity provides an improvement in the mutual information when compared to smaller N.

The same qualitative effect has also been shown to occur for various threshold non-idealities, for coupled and deterministic threshold noise [21], and for non-random input signals [22]. It has also been shown that the maximum mutual information occurs at a value that approaches $\sigma = 1$ as $N \to \infty$ for uniform signal and noise [23] and at a value that approaches $\sigma \simeq 0.60281$ for Gaussian signal and noise [24]. The SSR effect, which was also partly motivated by the model's similarity to populations of noisy sensory neurons, has subsequently been shown to occur in arrays of FitzHugh–Nagumo neuron models [25], and applied to cochlear implant encoding [26].

3. Distortion Performance of SSR

The analysis of the variation in mutual information with increasing noise can in one sense be considered to be measuring the performance of the SSR model as a quantization scheme's encoder. In another sense, it can be considered as a measure of distortion. Some authors in computational neuroscience consider mutual information to be a suitable measure of distortion performance, that is, the "information distortion" [27]. For example, such work argues that neural sensory systems may have evolved to ensure that that the most information about a signal possible is transmitted from the sensory system to higher brain functions. In the SSR model, there is an easily calculated upper limit on the mutual information between input and output. This is simply the maximum noiseless mutual information, which is the entropy of the output signal. Thus, the maximum mutual information occurs when all output states are equally probable, and is given by $I_m = \log_2 (N + 1)$. We can thus formulate an information distortion measure, in absolute terms as the difference between the actual mutual information and I_m, or as a percentage of I_m. Such an analysis has been performed explicitly [28], and also considered in [19], where it was shown that for large N, the SSR model provides a maximum mutual information of about $0.5 I_m$.

However, in conventional quantization work, a quantizer's performance is usually analyzed in terms of how well its decoding reproduces the original, unquantized signal. The analysis is in terms of some distortion performance, most often the mean squared error between original signal and reconstruction [1]. To enable a reconstruction of the original signal, some reproduction point is required to be specified for each encoded output state. For conventional noiseless quantizers, the optimal mean square error reproduction points are known to be the centroids of the corresponding input partition cell. We apply the same principle, with the difference being that integrations are performed over the entire support of x. Specifically, the decoding of the nth value of y that gives the minimum possible mean square distortion is the decoding given by $\hat{x}_n = \mathrm{E}_x[x|n]$ [1]. This can be written as

$$\hat{x}_n = \int_x xP(x|n)dx, \quad n = 0, \dots, N.$$

Previous work has described a sub-optimal decoding, which does not vary with n [20, 29].

The distortion that results from this decoding is the minimum mean square error (MMSE). It is straightforward to show that

$$\mathrm{MMSE} = \mathrm{E}[x^2] - \sum_{n=0}^{N} \frac{\left(\int_x xP(n|x)P_x(x)dx\right)^2}{P_y(n)}.$$

Noting that for a zero mean $P_x(x)$, $\mathrm{E}[x^2]$ is simply the variance of the input signal, σ_x^2, and taking the MMSE as the output noise, an output SNR measure can also easily be constructed, just as in conventional quantization schemes [1], as

$$\mathrm{SNR} = 10 \log_{10} \left(\frac{\mathrm{E}[x^2]}{\mathrm{MMSE}} \right) \quad dB.$$

Thus, as with the mutual information, the MMSE and output SNR can easily be calculated numerically for given $P_x(x)$, $R(\eta)$ and N. Plots of the output SNR against input SNR are shown for Gaussian signal and Gaussian noise for various values of N in Fig. 3. It is clear that the same qualitative behavior occurs as for the mutual information. There is some very small value of input SNR that maximizes the output SNR. However the maximizing value is not the same as that which maximizes the mutual information. In fact, it can be seen that as N becomes larger, the SNR is optimized for a value of input SNR that gets close to zero decibels. Recalling the link between this work and neural coding, we point out that some authors have reported that input SNRs in sensory neural systems are typically of the order of zero decibels [30].

It is also of interest to examine the operational rate-distortion [31] performance of the SSR model and compare it to the theoretical rate-distortion curve for a Gaussian source of $R(D) = 0.5 \log_2 (\sigma_x^2/D)$, where D is the mean square distortion [31]. Figure 4 shows the mutual information plotted against the output SNR, where a given point on the curve corresponds to a particular value of input SNR. The *dotted line* in Fig. 4 shows the theoretical $R(D)$ in terms of the output SNR in decibels.

Observe that the plot for a single value of N starts at a rate of one bit per sample, then increases with both rate and output SNR, as the input SNR decreases. The

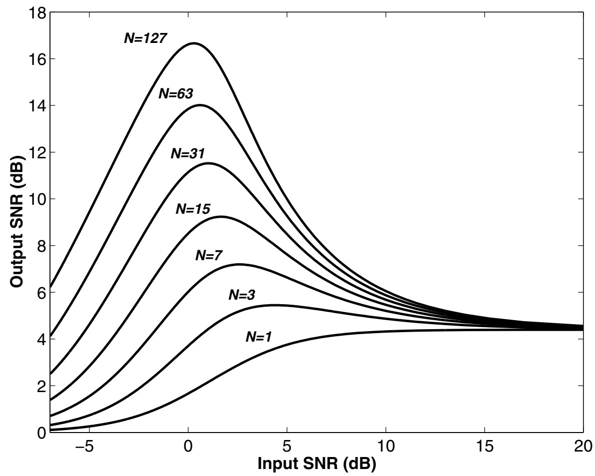

Fig. 3. Plot of output SNR against input SNR, for a zero mean Gaussian source with $\sigma_x = 1$, and Gaussian noise, with all thresholds set equal to the source mean of zero.

rate then reaches its maximum before the output SNR does. Then with continuing decreasing input SNR, the curve reaches its output SNR maximum, before curling back down towards the $R(D)$ curve. Note that this means that (except for very low input SNRs) there are two values of input SNR for which the same distortion can occur, corresponding to two different rates. If the main goal of a quantizer is to operate with maximum output SNR, this observation indicates that the optimal value of input SNR to use is the one which achieves the maximum SNR, rather than the maximum rate.

A further observation is the fact that for very low input SNRs, the SSR quantization scheme provides an output that is very close to the theoretical rate-distortion curve. However, this is probably irrelevant for quantizer design, because the main constraint on design would be the number of output bits, rather than the mutual information. For example in the case of $N = 127$ (i.e. a 7-bit output) the maximum output SNR that can be achieved is about 16.66-dB. By contrast, for the same value of N, a standard noiseless uniform scalar quantizer provides an output SNR of over 30-dB for an Gaussian source that is not companded. Alternatively, note that for a 13-dB output SNR, the uniform quantizer requires a 3-bit output whereas SSR requires 7 bits.

It is possible however, that in a practical quantization scheme, it may be an acceptable tradeoff to use the far many more thresholds required in the case of SSR to provide the same performance as a uniform scalar quantizer, to achieve a lesser complexity. In the case of SSR, all thresholds have the same value, whereas the uniform scalar quantizer requires N different threshold values.

A further scenario in which SSR could be usefully employed in a practical quantizer is under conditions where large threshold noise is unavoidable. If it is assumed that the *iid* additive noise present on each threshold in SSR is also present on the thresholds of a uniform scalar quantizer, then the distortion performance of SSR can match that of the uniform quantizer when each have the same value of N [3].

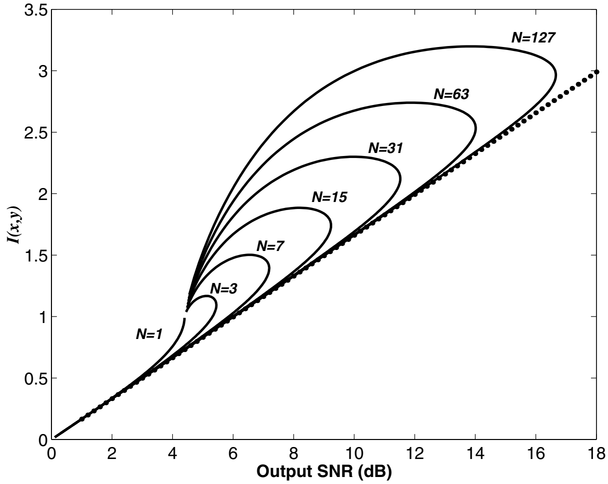

Fig. 4. Plot of mutual information against output SNR (in decibels) for a zero mean Gaussian source with unity variance, Gaussian noise, and a number of values of N. The *dotted line* shows the theoretical lower bound of the rate required for a given distortion. Note that for small input SNRs, a given output SNR can be achieved by two different values of mutual information. This is due to the concave nature of the plots of mutual information and output SNR against input SNR.

4. Optimal Thresholds

The previous sections studied scalar quantization in the form of SSR, that is, all thresholds were required to have the same value. Here we relax this constraint, and consider the problem of finding an optimal quantization in the presence of the same noise conditions as those present in the SSR model. Hence, we aim to find the threshold settings that maximize the mutual information, or minimize the mean square distortion, as the input SNR varies. These aims can be formulated as the following two nonlinear optimization problems,

$$\begin{aligned} &\text{Find:} && \max_{\{\theta_n\}} I(x,y) \\ &\text{subject to:} && \{\theta_n\} \in \mathbf{R}^{\mathbf{N}}, \end{aligned} \tag{2}$$

and

$$\begin{aligned} &\text{Find:} && \min_{\{\theta_n\}} \text{MMSE} \\ &\text{subject to:} && \{\theta_n\} \in \mathbf{R}^{\mathbf{N}}. \end{aligned} \tag{3}$$

Since the free variables are the set $\{\theta_n\}$, this is simply an N-dimensional maximization problem, and is therefore amenable to standard unconstrained optimization techniques. However, the objective function is not convex, and there exist a number of local optima. This problem can be overcome by employing random search techniques such as simulated annealing. We present here results for the optimal quantization obtained by solving problem (2) and problem (3) for Gaussian signal and Gaussian noise, $\sigma_x = 1$ and $N = 5$. Figure 5 shows the optimal thresholds for problem (2), and Fig. 6 shows the optimal thresholds for problem (3).

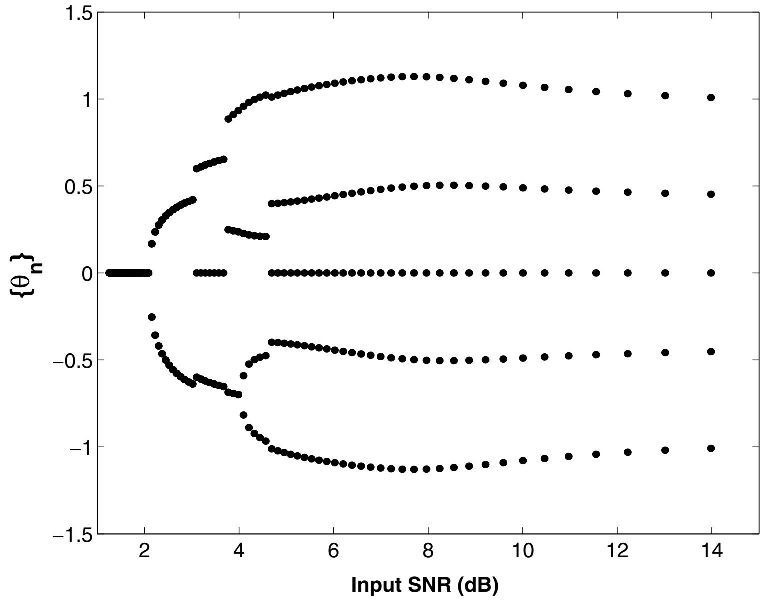

Fig. 5. Plot of thresholds that maximize the mutual information, against increasing input SNR, for $N = 5$ and a zero mean Gaussian source with unity variance, and Gaussian noise.

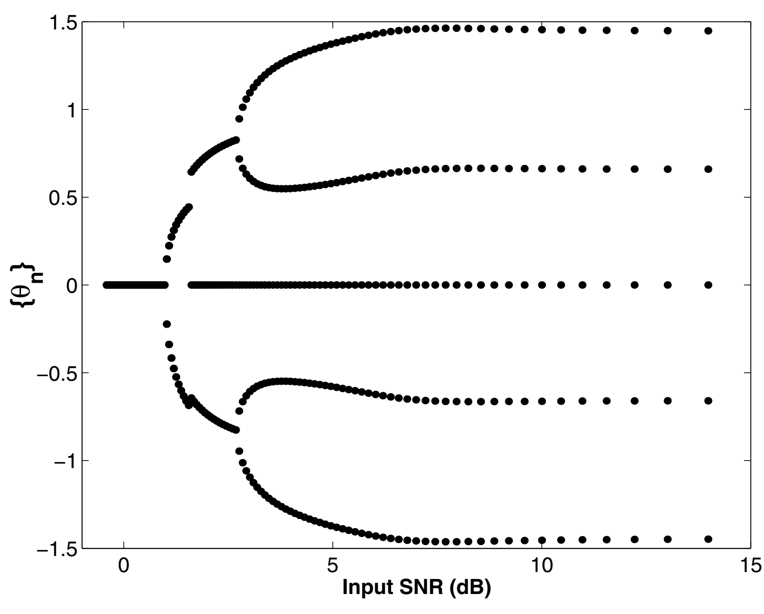

Fig. 6. Plot of thresholds that minimize the Minimum Mean Square Error (MMSE), against increasing input SNR, for $N = 5$ and a zero mean Gaussian source with unity variance, and Gaussian noise.

There are several notable features in these two plots, which both show the same qualitative behavior. Firstly, for large input SNRs, the optimal thresholds are all uniquely valued, and widely distributed across the input distribution's support. For very low input SNRs, the optimal thresholds are all equal to the same value of zero. This is precisely the situation that we impose in the SSR model. This shows that the SSR situation of all thresholds identical is an optimal quantization for sufficiently low input SNRs, when there is *iid* additive threshold noise. A further observation is the existence of bifurcations in the optimal threshold settings. Note that as the input SNR decreases, there are values of SNR at which the number of unique threshold values decreases. This occurs several times with decreasing SNR until the point where SSR becomes optimal. We have found such behavior to persist for larger N [32], for other source and noise distributions, and other measures.

Although to date we do not have an analytical explanation for the occurrence of these bifurcations, our optimization formulation is similar to previous work on clustering and neural coding problems solved using a method known as *deterministic annealing* [27, 33] in which similar bifurcations have occurred. In particular, the formulation reached in [27] can be expressed in a fashion identical to problem (2), with one exception. Here, in problem (2), the set $\{\theta_n\}$, imposes the set of transition probabilities, $\{P(n|x)\}$. In [27], there are no such structural constraints imposed on how $P(n|x)$ is obtained, and the set $\{P(n|x)\}$ is considered to be the set of variables to optimize.

5. Conclusions

We have shown in this paper how a phenomenon firstly described in the statistical physics literature, known as *suprathreshold stochastic resonance*, may be described in terms of lossy source coding and quantization theory. In particular, it has been demonstrated that SSR is in fact equivalent to scalar quantization in the presence of *iid* additive threshold noise. The independence of the noise at each threshold acts to randomly distribute the quantizer's thresholds across the source support, and effectively provide a stochastic quantization. We have shown that this quantization can be analyzed in terms of mutual information, mean square distortion, and operation rate-distortion measures. Furthermore, we have relaxed the SSR constraint of all thresholds being set to the same value, and optimized the thresholds for varying input SNRs. This optimization has shown that for sufficiently small input SNRs, the SSR model provides the optimal response. A side product of such optimizations shows that bifurcations occur in the optimal quantization as the input SNR decreases, so that more and more of the optimal thresholds coincide to the same values.

Acknowledgments

This work was funded by the Australian Research Council, the Australian Academy of Science, the Cooperative Research Centre for Sensor Signal and Information Processing (Australia), and the Leverhulme Trust, grant number F/00 215/J. We thank them for their support.

References

[1] R. M. Gray and D. L. Neuhoff, *Quantization*, *IEEE Trans. Info. Theor.* **44** (1998) 2325–2383.

[2] N. G. Stocks, *Suprathreshold stochastic resonance in multilevel threshold systems*, *Phys. Rev. Lett.* **84** (2000) 2310–2313.

[3] M. D. McDonnell, N. G. Stocks, C. Pearce and D. Abbott, *Neural mechanisms for analog to digital conversion*, in *BioMEMS and Nanoengineering*, ed. Dan V. Nicolau, *Proc. SPIE* **5275** (2003) 278–286.

[4] M. I. Dykman, D. G. Luchinsky, R. Mannella, P. V. E. McClintock, N. D. Stein and N. G. Stocks, *Stochastic resonance in perspective*, *Il Nuovo Cimento* **17** (1995) 661–683.

[5] L. Gammaitoni, P. Hänggi, P. Jung and F. Marchesoni, *Stochastic resonance*, *Rev. Mod. Phys.* **70** (1998) 223–287.

[6] S. Kay, *Can detectability be improved by adding noise?*, *IEEE Signal Process. Lett.* **7** (2000) 8–10.

[7] D. Abbott, *Overview: Unsolved problems of noise and fluctuations*, *Chaos* **11** (2001) 526–538.

[8] G. P. Harmer, B. R. Davis and D. Abbott, *A review of stochastic resonance: Circuits and measurement*, *IEEE Trans. Instrum. Meas.* **51** (2002) 299–309.

[9] S. Fauve and F. Heslot, *Stochastic resonance in a bistable system*, *Phys. Lett. A* **97** (1983) 5–7.

[10] B. McNamara, K. Wiesenfeld and R. Roy, *Observation of stochastic resonance in a ring laser*, *Phys. Rev. Lett.* **60** (1988) 2626–2629.

[11] A. Longtin, *Stochastic resonance in neuron models*, *J. Statist. Phys.* **70** (1993) 309–327.

[12] K. Wiesenfeld and F. Moss, *Stochastic resonance and the benefits of noise: From ice ages to crayfish and SQUIDS*, *Nature* **373** (1995) 33–36.

[13] S. M. Bezrukov and I. Voydanoy, *Noise-induced enhancement of signal transduction across voltage-dependent ion channels*, *Nature* **378** (1995) 362–364.

[14] Z. Gingl, L. B. Kiss and F. Moss, *Non-dynamical stochastic resonance: Theory and experiments with white and arbitrarily coloured noise*, *Europhys. Lett.* **29** (1995) 191–196.

[15] L. Gammaitoni, *Stochastic resonance and the dithering effect in threshold physical systems*, *Phys. Rev. E* **52** (1995) 4691–4698.

[16] A. Bulsara and A. Zador, *Threshold detection of wideband signals: A noise induced maximum in the mutual information*, *Phys. Rev. E* **54** (1996) R2185–R2188.

[17] R. M. Gray and T. G. Stockham, *Dithered quantizers*, *IEEE Trans. Info. Theor.* **39** (1993) 805–812.

[18] R. A. Wannamaker, S. P. Lipshitz and J. Vanderkooy, *Stochastic resonance as dithering*, *Phys. Rev. E* **61** (2000) 233–236.

[19] N. G. Stocks, *Information transmission in parallel threshold arrays: Suprathreshold stochastic resonance*, *Phys. Rev. E* **63** (2001) 041114.

[20] M. D. McDonnell, D. Abbott and C. E. M. Pearce, *An analysis of noise enhanced information transmission in an array of comparators*, *Microelectron. J.* **33** (2002) 1079–1089.

[21] F. Martorell, M. D. McDonnell, D. Abbott and A. Rubio, *Generalized noise resonance: Using noise for signal enhancement*, in *Fluctuations and Noise in Biological, Biophysical, and Biomedical Systems II*, eds. D. Abbott, S. M. Bezrukov, A. Der and A. Sánchez, *Proc. SPIE* **5467** (2004) 163–174.

[22] D. Rousseau and F. Chapeau-Blondeau, *Suprathreshold stochastic resonance and signal-to-noise ratio improvement in arrays of comparators*, *Phys. Lett. A* **321** (2004) 280–290.

[23] N. G. Stocks, *Suprathreshold stochastic resonance: An exact result for uniformly distributed signal and noise*, *Phys. Lett. A* **279** (2001) 308–312.

[24] T. Hoch, G. Wenning and K. Obermayer, *Optimal noise-aided signal transmission through populations of neurons*, *Phys. Rev. E* **68** (2003) 011911.

[25] N. G. Stocks and R. Mannella, *Generic noise enhanced coding in neuronal arrays*, *Phys. Rev. E* **64** (2001) 030902(R).

[26] N. G. Stocks, D. Allingham and R. P. Morse, *The application of suprathreshold stochastic resonance to cochlear implant coding*, *Fluct. Noise Lett.* **2** (2002) L169–L181.

[27] A. G. Dimitrov and J. P. Miller, *Neural coding and decoding: Communication channels and quantization*, *Netw.: Comput. Neural Syst.* **12** (2001) 441–472.

[28] N. G. Stocks, *Information transmission in parallel threshold networks: Suprathreshold stochastic resonance and coding efficieny*, in *Noise in Physical Systems and 1/f Fluctuations*, *ICNF 2001*, ed. G. Bosman, World Scientific, Singapore (2001) 594–597.

[29] M. D. McDonnell, D. Abbott and C. E. M. Pearce, *A characterization of suprathreshold* stochastic resonance in an array of comparators by correlation coefficient, *Fluct. Noise Lett.* **2** (2002) L205–L220.

[30] W. Bialek, M. DeWeese, F. Rieke and D. Warland, *Bits and brains: Information flow in the nervous system*, *Physica A* **200** (1993) 581–593.

[31] T. Berger and J. D. Gibson, *Lossy source coding*, *IEEE Trans. Info. Theor.* **44** (1998) 2693–2723.

[32] M. D. McDonnell, N. G. Stocks, C. E. M. Pearce and D. Abbott, *Point singularities and suprathreshold stochastic resonance in optimal coding*, (2004) Preprint arXiv:cond-mat/0409528.

[33] K. Rose, *Deterministic annealing for clustering, compression, classification, regression and related optimization problems*, *Proc. IEEE* **86** (1998) 2210–2239.

Copyright of Fluctuation & Noise Letters is the property of World Scientific Publishing Company. The copyright in an individual article may be maintained by the author in certain cases. Content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.