How to use noise to reduce complexity in quantization

Mark D. McDonnell*[a], Nigel G. Stocks[b], Charles E.M. Pearce[c] and Derek Abbott[a],

[a]Centre for Biomedical Engineering (CBME) and
School of Electrical & Electronic Engineering,
The University of Adelaide, SA 5005, Australia

[b]School of Engineering,
The University of Warwick, Coventry CV4 7AL, UK

[c]School of Mathematical Sciences
The University of Adelaide, SA 5005, Australia

ABSTRACT

Consider a quantization scheme which has the aim of quantizing a signal into $N+1$ discrete output states. The specification of such a scheme has two parts. Firstly, in the encoding stage, the specification of N unique threshold values is required. Secondly, the decoding stage requires specification of $N+1$ unique reproduction values. Thus, in general, $2N+1$ unique values are required for a complete specification. We show in this paper how noise can be used to reduce the number of unique values required in the encoding stage. This is achieved by allowing the noise to effectively make all thresholds independent random variables, the end result being a stochastic quantization. This idea originates from a form of stochastic resonance known as suprathreshold stochastic resonance. Stochastic resonance occurs when noise in a system is essential for that system to provide its optimal output and can only occur in nonlinear systems—one prime example being neurons. The use of noise requires a tradeoff in performance, however, we show that even very low signal-to-noise ratios can provide a reasonable average performance for a substantial reduction in complexity, and that high signal-to-noise ratios can also provide a reduction in complexity for only a negligible degradation in performance.

Keywords: stochastic resonance, suprathreshold stochastic resonance, quantization, optimal quantization, bifurcations, noise, analog-to-digital conversion, lossy compression

1. INTRODUCTION

We begin this paper in Section 2 by presenting the basic theory of quantization of a scalar random variable. This includes discussion of how the noise introduced by quantization is often measured by mean square error distortion, or by signal-to-quantization-noise ratio. Section 2 also describes an algorithm for finding the *optimal* quantization for a given source distribution, in the conventional case of deterministic quantization.

Next, Section 3 describes a form of quantization known as Suprathreshold Stochastic Resonance (SSR). This phenomenon occurs in arrays of *identical* threshold devices, which are all subject to independent noise.

Section 4 then examines extension of the SSR situation to arbitrary threshold values, and solves the problem of optimal quantization when each threshold is subject to independent noise, as the noise intensity varies.

Finally, Section 5 uses the results of Section 4 to discuss how complexity can be traded off against quantizer performance by increasing the input threshold noise intensity.

*mmcdonne@eleceng.adelaide.edu.au; phone +61 8 8303 6296; fax +61 8 8303 4360

Complex Systems, edited by Axel Bender, Proc. of SPIE Vol. 6039
60390E, (2006) · 0277-786X/06/$15 · doi: 10.1117/12.638476

2. QUANTIZATION THEORY

Quantization of a signal, or source quantity, consists of the partitioning of the signal into a discrete number of intervals, or cells.[1] A set of rules specifies which range of values of the signal get assigned to each cell. This part of a quantization is known as the *encoding* stage. If an estimate of the original signal is required to be made from this encoding, then each cell must also be assigned a *reproduction value*. This part is known as the *decoding* stage.

One of the main applications of quantization is in Analog-to-Digital Conversion (ADC) in an electronic circuit. However, quantization is also used elsewhere, and its theory is equally applicable. For example, quantization occurs whenever a rounding-off operation is made; in the representation of real numbers in a computer's architecture by floating-point or fixed-point representation; in the formation of histograms from real-valued data; and in numerical integration or numerical solutions to differential equations.

The remainder of this Section discusses the key quantization theory required for this paper.

2.1. Encoding

Consider a quantizer which uses N thresholds, with values $\theta_1, .., \theta_N$, to partition an input signal, x, into $N+1$ cells or output states. Suppose we label these states as $n = 0, .., N$, and consider this encoding as a discretely valued signal, y. Then this signal can be written as a function of x, in terms of the signum (sign) function as

$$y(x) = \frac{1}{2}\sum_{i=1}^{N} \text{sign}(x - \theta_i) + \frac{N}{2}. \quad (1)$$

2.2. Decoding and distortion of a quantizer

Each output state of a quantizer must be assigned a reproduction value in order to reconstruct the original signal. We define the reproduction value corresponding to output state $y = n$ as $z(n)$. Hence, the decoded output signal of a quantizer can be written as $z(y)$, to denote that the output is a function of the encoded signal, y.

Unlike sampling a bandlimited signal at the Nyquist rate, quantization of a signal will always cause some error in a reproduction of the original signal. This error is known as the distortion, or quantization noise, and is most commonly measured by the Mean Square Error (MSE) between the original signal, x, and the reproduced signal, $z(y(x))$.[1] Thus, the error signal is defined as a function of x as

$$\epsilon(x) = x - z(y(x)).$$

Suppose the input signal, x, is a sequence of *iid* (independent and identically distributed) samples from a continuously valued probability distribution, with Probability Density Function (PDF), $P(x)$. If $P(x)$ has support $x \in [a, b]$ then the MSE distortion is

$$\begin{aligned} D_{\text{ms}} &= \text{E}[(x - z)^2] \\ &= \sum_{i=0}^{N} \int_{\theta_i}^{\theta_{i+1}} (x - z(i))^2 P(x) dx, \end{aligned} \quad (2)$$

where we take $\theta_0 = a$ and $\theta_{N+1} = b$.

A commonly used measure of a quantizer's performance is its signal-to-quantization-noise ratio (SQNR), which is the ratio of the input signal power to the MSE distortion power. If the input signal has power σ_x^2, then this can be expressed as

$$\text{SQNR} = 10 \log_{10}\left(\frac{\sigma_x^2}{D_{\text{ms}}}\right). \quad (3)$$

2.3. Uniform scalar quantizers

The most common quantization scheme is the *uniform scalar quantizer* (USC).[1] If the input signal, x, has an amplitude range between $[-m_p, m_p]$ then the USC partitions the signal uniformly into $N+1$ intervals of width $\Delta\theta = 2m_p/(N+1)$ by N thresholds. Hence the threshold values required are

$$\theta_i = -m_p\left(1 - \frac{2i}{N+1}\right) \quad i = 1, .., N. \tag{4}$$

For a USC, the reproduction value, $z(y)$, corresponding to the y-th cell, is usually given by the midpoint of the interval number, $y \in 0, .., N$, in which an input sample lies. This means that

$$z(y) = \frac{2m_p}{N+1}y - \frac{Nm_p}{N+1} \quad y = 0, .., N. \tag{5}$$

If we make the assumption that the PDF of the quantization noise, $\epsilon(x)$, is uniform, then it is straightforward to show that the MSE distortion is $(\Delta\theta)^2/12$, that is,

$$D_{\mathrm{ms}} = \frac{m_p^2}{3(N+1)^2}. \tag{6}$$

This assumption is only strictly true for a uniformly distributed input signal, however if $\Delta\theta$ is small—that is, N is reasonably large—and the signal PDF is reasonably smooth, then this is a reasonable approximation.[2] Substituting Eqn. (6) into Eqn. (3) gives

$$\mathrm{SQNR} = 3(N+1)^2\frac{\sigma_x^2}{m_p^2}. \tag{7}$$

2.4. Optimal quantization

For source distributions that are not uniform, the USC does not give the minimum possible MSE distortion. To achieve such a goal, it is necessary to find both the optimal thresholds, as well as the optimal reproduction values. For a source PDF, $P(x)$, with support, $x \in [a, b]$, it is well known that necessary conditions for optimal MSE quantization[2] are, (i) that the reproduction values are given by the *centroids* of the corresponding cell,

$$z(i) = \frac{\int_{\theta_i}^{\theta_{i+1}} xP(x)dx}{\int_{\theta_i}^{\theta_{i+1}} P(x)dx}, \quad i = 0, .., N, \tag{8}$$

where $\theta_0 = a$ and $\theta_N = b$, and (ii) that the optimal thresholds satisfy

$$\theta_i = \frac{z(i-1) + z(i)}{2}, \quad i = 1, .., N. \tag{9}$$

It can be shown that the MSE resulting from the reproduction values of Eqn. (8) is given by

$$D_{\mathrm{ms}} = \mathrm{E}[x^2] - \sum_{i=0}^{N} P_y(i)z(i)^2, \tag{10}$$

where $P_y(i) = \int_{\theta_i}^{\theta_{i+1}} P(x)dx$.

2.5. Lloyd Method I

Note that Eqns. (9) and (10) are mutually dependent on each other. However the thresholds and reproduction points that minimize D_{ms} can be found by a simple iterative procedure known as the *Lloyd Method I algorithm*,[2,3] which is a commonly used technique for finding the optimal quantization for a given source PDF. This algorithm begins with an initial guess for the reproduction values, and then finds the optimal thresholds for those values, from Eqn. (9). Given these new thresholds, a new set of reproduction values are found from Eqn. (8), and the new D_{ms} from Eqn (10). This new D_{ms} can be shown to be smaller than the previous D_{ms}. This iteration is repeated until D_{ms} no longer decreases, at which point the optimal thresholds and reproduction values have been found.

Fig. 1(a) shows the result of using the Lloyd Method I algorithm to find the optimal thresholds and reproduction values for a zero mean, unity variance Gaussian source. It is clear from this figure that each threshold is midway between each reproduction value, and that the size of each cell increases towards the tails of the source PDF.

Fig. 1(b) shows the SQNR that results from optimally quantizing the Gaussian source as N increases. Clearly, increasing N means that the SQNR increases, as we would expect from a finer quantization.

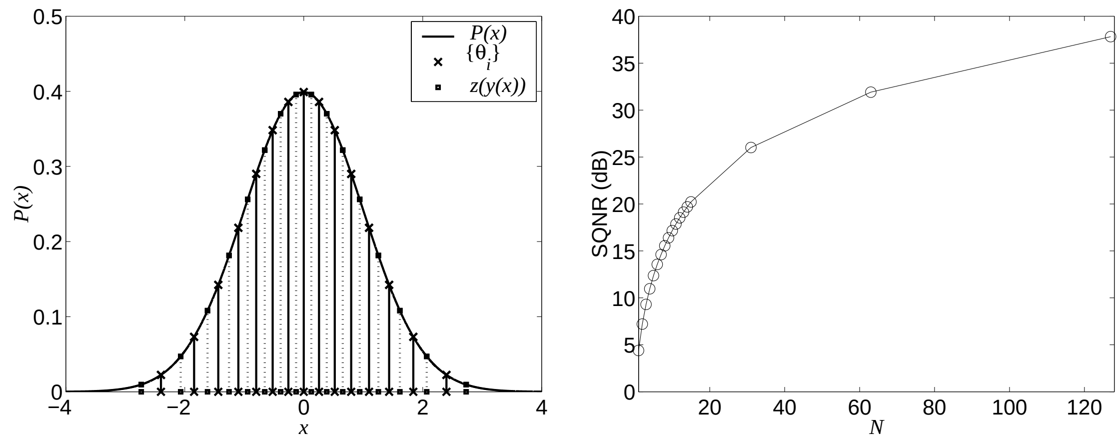

(a) Optimal quantization (b) SQNR vs N

Figure 1: Fig. 1(a) shows the optimal MSE distortion thresholds and reproduction values, as calculated by the Lloyd Method I algorithm, for $N = 15$ and a Gaussian source with zero mean, and unity variance. The values of the thresholds are shown on the x–axis with crosses, and the values of reproduction points are shown with squares. The PDF itself is plotted on the y–axis, against x, to indicate how the optimal thresholds and reproduction values vary as the height of $P(x)$ varies. Fig. 1(b) shows how the SQNR of an optimally scalar quantized unity variance Gaussian source increases with N. Values of N shown are $1, 2, .., 15$, as well as $31, 63$ and 127.

2.6. Complexity of a scalar quantizer

We define the complexity of a quantizer as the number of unique values required to completely describe its encoding and decoding operations. We have seen that specification of a scalar quantizer requires specification of N threshold values and $N + 1$ reproduction points. Therefore, a total of $2N + 1$ unique values are required for a given source PDF, and the quantizer's complexity increases linearly with N. We have also seen that the MSE distortion decreases with N. If it is assumed that the quantization noise distribution is uniform, then from Eqn. (6), the MSE distortion decreases in $\mathrm{O}(1/N^2)$.

Although this measure of complexity is not the same "kind" of complexity found in many natural systems, the number of unique values required to implement a quantizer is an important constraint in practical quantizer design, such as in an ADC implementation. As N increases, the difference between consecutive threshold values decreases, and due to non-idealities like voltage noise, it becomes more difficult to maintain the ideal values.

The next section now describes a form of *stochastic quantization*. It is necessary to begin with the historical background to this work, by giving a brief overview of *stochastic resonance*.

3. STOCHASTIC QUANTIZATION

3.1. Stochastic resonance

Stochastic Resonance (SR), although a term originally used in a very specific context, is now a term broadly applied to any phenomenon where the presence of internal noise or external input noise in a nonlinear system provides a better—according to some measure of performance—system response to a certain input signal than in the absence of noise.[4–7] The key term here is *nonlinear*. SR cannot occur in a linear system—linear in this sense means that the output of the system is a linear transformation of the input of the system.

SR has been widely observed throughout nature—it has been discovered in such diverse systems as climate models,[8] electronic circuits,[9] differential equations,[10] lasers,[11] neural models,[12,13] physiological neural populations,[14] simple threshold devices,[15,16] chemical reactions,[17] ion channels,[18] SQUIDs (Superconducting Quantum Interference Devices),[19] ecological models,[20] financial models,[21] and even social systems.[22]

SR is often described as a counter-intuitive phenomenon. This is largely due to its historical background, since in the first decade and a half since its discovery in 1980,[8] virtually all research into SR considered only systems driven by periodic—most often a sine-wave—input signals and broadband noise. In such systems, a natural measure of system performance is the output signal-to-noise ratio (SNR). In linear systems driven by periodic input signals—in fact, any linear system—it is well known that the output SNR is maximised in the absence of noise. When systems are analysed in terms of SNR, it is the norm to implicitly assume that noise is a problem, usually with good reason. Hence, observations of the presence of noise in a system providing the maximum output SNR are often seen to be highly counter-intuitive.

However, although not all noise can be described well as a random variable—it can, for example be constant or deterministic (even chaotic)—workers in the field have tended to focus on the stochastic case. The most common assumption is that the noise is Gaussian and white. When it is noted that there are many examples of systems or algorithms where randomness is of benefit, SR does not seem quite so counter-intuitive. Examples include Brownian ratchets[23]—mechanical applications of this idea include self-winding (batteryless) wristwatches[24]; dithering in analog-to-digital conversion[25,26]; Parrondo's games[27]; noise induced stabilisation[28,29]; noise induced synchronisation[30]; the use of mixed (probabilistic) optimal strategies in game theory[31]; neural network generalization[7]; and random search optimization algorithms, such as genetic algorithms[32] and simulated annealing.[33]

The distinguishing feature of SR that sets it apart from most of the above list is that SR is usually understood to occur in systems where there are well-defined input and output *signals* and the optimal output signal, according to some measure, occurs for some non-zero level and type of noise. As mentioned, SR was initially considered to be restricted to the case of periodic input signals. However, now it is used as an all encompassing term, whether or not the input signal is a periodic sine-wave, a periodic broadband signal, or aperiodic. An appropriate measure of output response depends on the task at hand, and the form of input signal. For example, for periodic signals and broadband noise, SNR is used, but for random aperiodic signals, mutual information[16] or correlation based measures[34] are more appropriate.

3.2. Suprathreshold stochastic resonance

One recent addition to the diverse range of SR results was the first observation of SR in a system consisting of many identical threshold devices subject to *iid* additive noise. In such a system, if each threshold device receives the same input signal, and the overall output is the sum of the individual threshold outputs, then SR occurs irrespective of whether the input signal is entirely subthreshold or not. This is in contrast to all prior studies on SR in simple threshold based systems, since for a single threshold device, and suprathreshold signals, SR effects disappear. Hence, the occurrence of SR in such a system was entitled Suprathreshold Stochastic Resonance (SSR).[35,36] The initial work considered the input to the array of thresholds to be *iid* samples from a Gaussian or uniform distribution, and the *iid* noise on each threshold to likewise be samples from a Gaussian or uniform distribution. The measure used was the mutual information between the input signal and the discretely valued output. The system in which SSR occurs is shown in Fig. 2, can be modelled mathematically as follows.

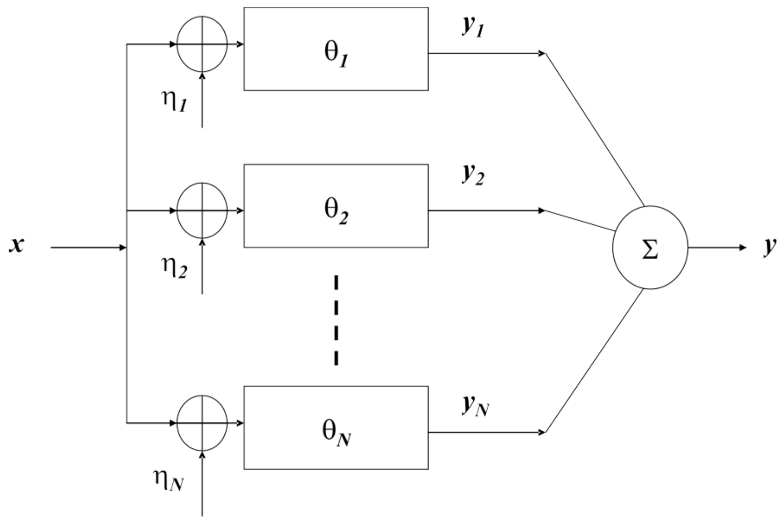

Figure 2: Array of N threshold devices. Each threshold device receives the same input signal, x, but is subject to independent additive noise, η_i. The output from threshold device i is unity if the sum of the signal and noise is greater than that devices's threshold, θ_i, and zero otherwise. The overall output, y, is the sum of the individual device outputs. For SSR, all thresholds are set to the same value, so that $\theta_i = \theta \ \forall \ i$. In the absence of noise, and for $\theta_1 < \theta_2 <, .., < \theta_N$, this model also describes a flash ADC.

Let the number of threshold devices be N. The i-th threshold device is subject to continuously valued *iid* additive noise, η_i ($i = 1, .., N$), drawn from a distribution with PDF $R(\eta)$. The output from each device, y_i, is unity if the input signal, x, plus the noise is greater than that device's threshold, θ_i, and zero otherwise. All outputs are summed to give the overall output signal, y. Hence, y is a discrete encoding of the input taking on integer values between 0 to N. This signal is therefore the number of devices that are "on" for each signal sample. For SSR, all thresholds are set to the same value, so that $\theta_i = \theta \ \forall \ i$. However, for arbitrary thresholds we have

$$y(x) = \frac{1}{2}\sum_{i=1}^{N} \mathrm{sign}[x + \eta_i - \theta_i] + \frac{N}{2}.$$

Note the similarity between this formula, and that of Eqn. (1). For the particular case of SSR, the difference is that now all threshold values are identical, rather than unique. Also, for arbitrary thresholds, due to the independent noise, $y(x)$ is now a random variable, rather than deterministic for a given input signal value, x.

Due to the nondeterministic encoding of an input value x to an output state n, the key function required to measure the system's performance is the joint PDF between the input and output signals, $P_{xy}(x, y)$. We commence by deriving a method of calculating $P_{xy}(x, y)$ for the general case of N arbitrary thresholds, and then simplify to the SSR case of all thresholds equal to the signal mean. We denote the PDF of the input signal as $P(x)$ and the probability mass function of the output as $P_y(n)$. Then we have, as a consequence of Bayes' theorem,

$$P_{xy}(x, y) = P(y = n|x)P(x).$$

We will refer to the set of conditional probabilities, $P(y = n|x)$ as the "transition probabilities," and abbreviate the notation used to $P(n|x)$. Integration of the joint probability with respect to x gives

$$P_y(n) = \int_{-\infty}^{\infty} P(n|x)P(x)dx \quad n = 0, .., N.$$

Recall we assume that the noise is *iid* at each comparator. Let $\hat{P}_i(x)$ be the probability of device i being "on" (that is, signal plus noise exceeding the threshold θ_i), given the input signal, x. Then

$$\hat{P}_i(x) = \int_{\theta_i - x}^{\infty} R(\eta)d\eta = 1 - F_R(\theta_i - x), \tag{11}$$

where $F_R(\cdot)$ is the cumulative distribution function of the noise and $i = 1, .., N$. If the noise has an even PDF then $\hat{P}_i(x) = F_R(x - \theta_i)$. In general, it is difficult to find an analytical expression for $P(n|x)$ and we will rely on numerics. Given any arbitrary N, $R(\eta)$, and $\{\theta_i\}$, $\{\hat{P}_i(x)\}$ can be calculated exactly for any value of x from Eqn. (11), from which $P(n|x)$ can be found using an efficient recursive formula.[37] For the particular case when the thresholds all have the same value, then each $\hat{P}_i(x)$ has the same value $\hat{P}(x)$ and, as noted by Stocks[35] we have the transition probabilities given by the binomial distribution as

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

Using the transition probabilities, it is possible to calculate the mutual information between the system input and output.[35, 37] However, in line with conventional quantization theory, we will find a decoding of y that approximates reconstructs the input signal, and examine the MSE distortion performance.[38]

Let the output of the array of comparators be decoded so that each value of $y = 0, .., N$ is decoded to $\hat{y} \in [\hat{y}_0.., \hat{y}_N]$, where $\hat{y}_n$ is the reproduction value for encoded value n. Thus, the MSE distortion is

$$D_{\text{ms}} = \int_{-\infty}^{\infty} \sum_{n=0}^{N} (x - \hat{y}_n)^2 P(n|x)P(x)dx. \tag{12}$$

Note that unlike the MSE distortion for a deterministic quantizer as given by Eqn. (2), for noise PDFs with infinite support, all output states are now achievable for any input value. Thus, the integral in Eqn. (12) is over all x, and we require knowledge of the transition probabilities, $P(n|x)$.

Just as with the conventional deterministic scalar quantizer, we can find the optimal MSE distortion decoding. We label the output signal resulting from this decoding as $\hat{x}$. Each value of this decoding is the expected value of x given y,[2] which is

$$\hat{x}_n = \text{E}[x|n] = \frac{\int_{-\infty}^{\infty} xP(x)P(n|x)dx}{\int_{-\infty}^{\infty} P(x)P(n|x)dx} \quad n = 0, .., N. \tag{13}$$

Again, note the similarity between Eqn. (13) and Eqn. (8). Just as with the distortion, the integrations are now over all x, and the reproduction values depend on the transition probabilites. Using this decoding, then we have the Minimum Mean Square Error (MMSE) distortion,

$$\text{MMSE} = \text{E}[x^2] - \text{E}[\hat{x}_n^2]. \tag{14}$$

Noting that $\text{E}[x] = \text{E}[\hat{x}] = 0$, the MMSE is the difference between the variance of the input signal and the variance of the optimally decoded output signal, just as with the deterministic scalar quantizer.

In line with the original SSR work,[35] we define the input noise intensity as the square root of the ratio of noise power to signal power, $\sigma^2 = \sigma_\eta^2/\sigma_x^2$. This is the inverse of the input SNR.

Fig. 3(a) shows for various N, the MMSE distortion for a unity variance Gaussian source, for Gaussian noise, as a function of noise intensity, σ. It is clear that SR occurs, since the distortion is minimized for a nonzero value of σ. As N increases the distortion decreases for all σ, and the value of σ at which the minimum distortion occurs increases towards unity. This value corresponds to an input SNR of 0 dB.

Fig. 3(b) shows the SQNR performance of SSR, at the optimal value of σ, as a function of N, compared with the optimal noiseless SQNR as a function of N. We can see that SSR can achieve the same SQNR as an optimal noiseless scalar quantizer, but to do this, requires a larger value of N. For example, SSR with $N = 63$ (6 bit output), provides about the same SQNR performance as a noiseless quantizer with $N = 7$ (3 bit output).

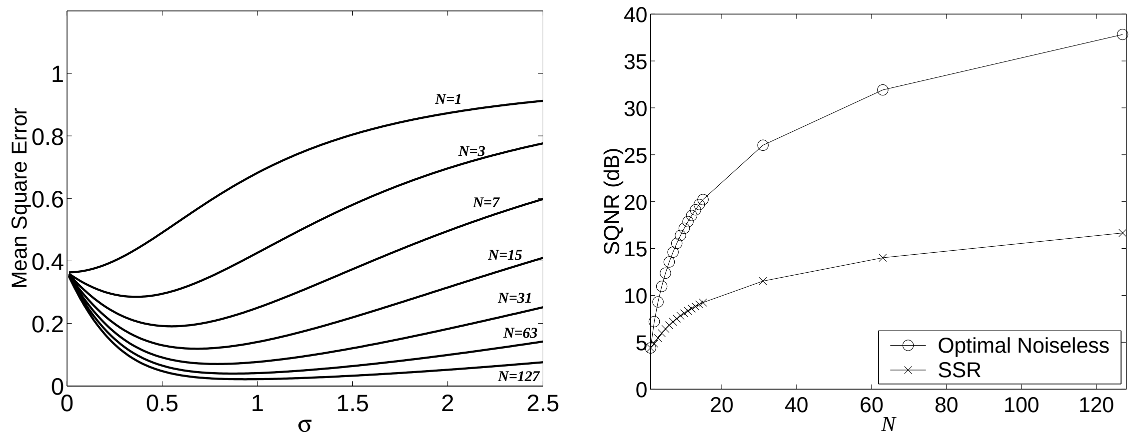

(a) SSR for various N (b) SQNR vs N for SSR

Figure 3: Fig. 3(a) shows for various N, the MMSE distortion for a unity variance Gaussian source, for Gaussian noise, as a function of noise intensity, σ. As N increases the distortion decreases for all σ, and the value of σ at which the minimum distortion occurs increases towards unity. This value corresponds to an input SNR of 0 dB. Fig. 3(b) shows the SQNR performance of SSR, at the optimal value of σ, as a function of N, compared with the optimal noiseless SQNR as a function of N.

4. OPTIMAL STOCHASTIC QUANTIZATION

We now aim to find the threshold values that minimize the MSE distortion in the quantization model shown in Fig. 2, as a function of input noise intensity. This goal can be formulated as the following nonlinear optimization problem:

$$\begin{aligned} \text{Find:} \quad & \min_{\{\theta_i\}} \text{MMSE} \\ \text{subject to:} \quad & \{\theta_i\} \in \mathbb{R}^N. \end{aligned} \tag{15}$$

The solution to Problem (15) in the absence of noise is exactly the optimal solution found by the Lloyd Method I algorithm. However, for nonzero noise, although we have Eqn. (13) for the optimal reproduction values, which implicitly depend on the threshold values, via the transition probabilities, we do not have a necessary condition for the optimal thresholds. Hence, it is not possible to solve Problem (15) using an iterative descent method like the Lloyd Method I algorithm.

However, it is possible to solve Problem (15) numerically using standard unconstrained non-linear optimization methods such as the conjugate gradient method.[39] The main problem with such an approach is that the objective function is not convex in $\{\theta_i\}$, and there exist many local optima. This problem can be overcome by employing random search techniques such as simulated annealing,[33] or by solving for many different initial threshold values.

4.1. Results

We present here, as an example, results for the optimal quantization obtained by solving Problem (15) for Gaussian signal with unity variance, and Gaussian noise, for $N = 5$. Fig. 4(a) shows the optimal thresholds plotted against increasing noise intensity, σ. The optimal noiseless thresholds, as calculated by the Lloyd method I are shown with crosses at $\sigma = 0$. Fig. 4(a) also shows the optimal reconstruction values as a function of σ. Fig. 4(b) shows the MMSE distortion obtained with the optimal thresholds, and for comparison, the MMSE for SSR.

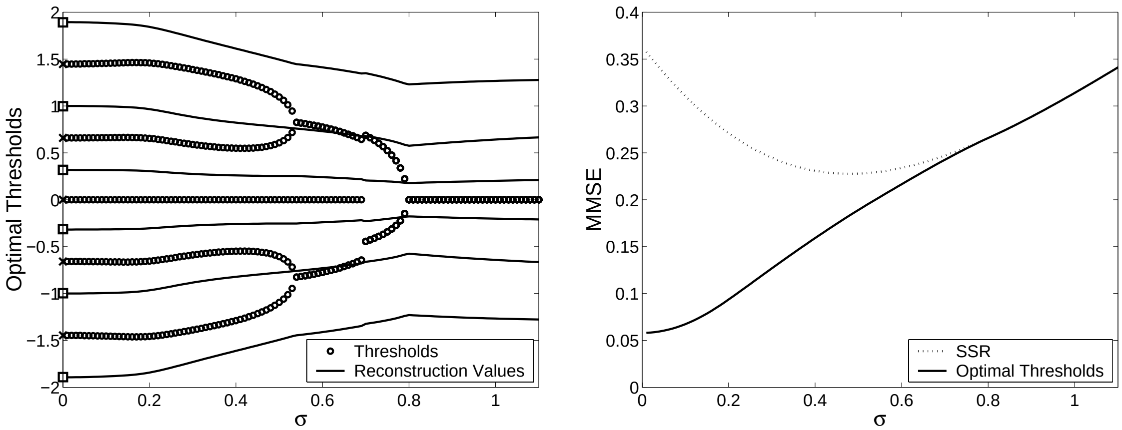

(a) Optimal quantization

(b) MMSE from optimal quantization

Figure 4: Fig. 4(a) shows the thresholds (circles) and reconstruction values (thick lines) that minimize the MMSE, against increasing σ, for $N = 5$ and a zero mean Gaussian source with unity variance, and Gaussian noise. The optimal noiseless thresholds are shown with crosses, and the optimal noiseless reproduction values with squares. Fig. 4(b) shows the MMSE that results from the optimal quantization of Fig. 4(a), compared with the MMSE from the SSR situation.

4.2. Discussion

Fig. 4(b) shows that the MMSE distortion obtained with the optimal thresholds is strictly increasing with increasing σ. This means that no SR effect is seen for optimized thresholds. The MMSE for SSR with $N = 5$ is shown for comparison. As σ increases, the difference in MMSE between the two cases decreases, until for $\sigma \simeq 0.8$, SSR becomes optimal.

From Fig. 4(a), we observe, firstly, that for very small noise, the optimal thresholds are consistent with the optimal noiseless values. There does not appear to be a discontinuity in the optimal thresholds as the noise intensity increases from zero to some small nonzero value. However, the most striking feature is the fact that *bifurcations* are present. For σ between zero and some critical value greater than zero, the optimal placement of the 5 thresholds are at unique values. However, for larger σ, a bifurcation occurs, and some of the optimal thresholds coincide to the same value. As σ increases, more bifurcations occur, and increasing fractions of the total number of thresholds tend to cluster to fewer identical values.

The bifurcation structure is quite surprising, but appears to be fundamental to the problem type, as we have found similar bifurcation structures in the optimal threshold settings for measures other than MMSE, including mutual information,[40] and other signal and noise distributions. Finally, it is evident that above a certain value of σ the SSR situation is optimal. That is, the optimal quantization for large noise is to set all thresholds to the signal mean.

The discontinuous bifurcations in the optimal threshold diagram are due to the presence of many locally optimal threshold configurations. In fact, numerical experiments find that for every value of σ, there is at least one locally optimal solution—that is a set of threshold values giving a gradient vector of the MMSE distortion with respect to $\{\theta_i\}$ of zero—corresponding to every possible *integer partition* of N. For each partition, there are as many locally optimal solutions as there are unique orderings of that partition. For small σ, all of these local optima are unique. As σ increases, more and more of these local optima bifurcate continuously to become coincidental with other local optima.

Due to this structure, it is very difficult to find the exact optimal solution for larger N. However, we have found that many of the local optima provide a MMSE distortion very close to optimal. Therefore, if we are only interested in finding solutions that are close to optimal, we are able to approximately solve Problem (15) for larger N. If we do this, we find a bifurcation structure like that shown in Fig. 4(a) persists, but that the value of σ for which the first bifurcation occurs decreases, and that there are many more bifurcations. This means that even for small values of noise intensity, it is not optimal to have N unique threshold values—i.e. some of the threshold values can be identical. We can exploit this fact to reduce the complexity of quantization. An example of this is given in Section 5.

5. NOISE VS COMPLEXITY TRADEOFF

We now comment on how, even very low input SNRs, a reasonable average performance for a substantial reduction in complexity can be achieved in a quantizer. For large input SNRs, we can also provide a reduction in complexity for only a negligible degradation in performance.

5.1. SSR

We saw in section 3 that for large σ—that is, a small input SNR—the same MSE distortion can be achieved by SSR as for the optimal thresholds in the absence of noise. However, this requires that we use a larger value of N in the SSR situation. Alternatively, we can tradeoff complexity for performance, since, for the same value of N, although SSR does not achieve as good a distortion as an optimal quantizer in the absence of noise, it requires only a single threshold value to be specified, instead of N.

5.2. Arbitrary thresholds

Since we would always desire as small a distortion as possible, it is preferable to try to operate a quantizer under small input SNR conditions. Our results for SSR show that for small input SNRs, the distortion performance degrades rapidly with decreasing σ. However, our observations in Section 4 indicate that the performance for small σ can be improved from the SSR situation without resorting to N unique thresholds, simply by clustering a number of thresholds at a small number of unique values.

Hence, we solve Problem (15) for a given N, under the constraint that only some fraction of the N threshold values are allowed to be unique. We consider, as an example, $N = 16$, and the cases of (i) 2 unique threshold values, with two clusters of size 8; and (ii) 4 unique threshold values, with clusters of size 4.

The optimal thresholds resulting from solving these problems are shown in Figs. 5(a) and 5(b). The resultant MMSE distortion is shown in Fig. 5(c).

We can see from Fig. 5(c), that for small σ, the MMSE distortion is decreased by increasing the number of clusters of unique threshold values, while for larger σ, it is optimal to have less clusters. In each situation we can see that SR occurs, since the optimal noise intensity is nonzero. However, as the number of clusters increases, this noise intensity decreases, and the optimal MMSE distortion at this noise intensity improves.

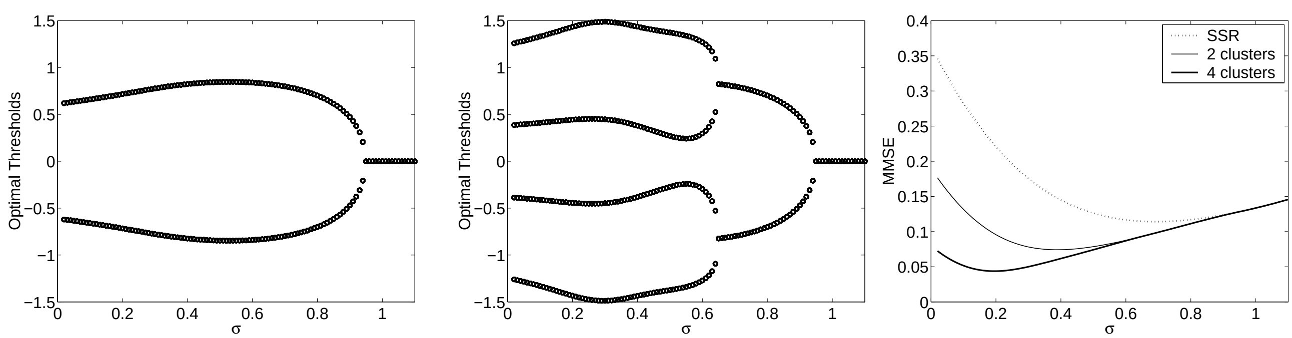

(a) $N = 16$, two clusters (b) $N = 16$, four clusters (c) $N = 16$, MMSE comparison

Figure 5: Figs. 5(a) and 5(b) show the optimal values of two threshold clusters and four threshold clusters for $N = 16$, and Gaussian signal and noise, as a function of σ. The resultant MMSE distortion for these two cases, as well as for SSR, is shown in Fig. 5(c).

6. CONCLUSIONS

We have demonstrated in this paper how noise can be used to tradeoff performance for complexity in a quantization scheme. This phenomenon relies on the assumption of *independent* threshold noise. Such a situation could occur when quantization of data is required to be made by widely spatially separated sensors, and therefore may have applications for distributed sensor networks. Alteratively if each threshold is required to operate sequentially in time for every signal sample, independent noise may result. This work may also be relevant to our understanding of how sensory neurons encode information. It has previously been shown that SSR can occur in arrays of neuron models,[41] and SSR has been applied to cochlear implant encoding.[42]

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 and we thank them for their support.

REFERENCES

1. R. M. Gray and D. L. Neuhoff, "Quantization," *IEEE Transactions on Information Theory* **44**(6), pp. 2325–2383, 1998.
2. A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*, Kluwer Academic Publishers, 1992.
3. S. P. Lloyd, "Least squares quantization in PCM," *IEEE Transactions on Information Theory* **IT-28**(2), pp. 129–137, 1982.
4. L. Gammaitoni, P. Hänggi, P. Jung, and F. Marchesoni, "Stochastic resonance," *Reviews of Modern Physics* **70**(1), pp. 223–287, 1998.
5. V. S. Anishchenko, A. B. Neiman, F. Moss, and L. Schimansky-Geier, "Stochastic resonance: Noise enhanced order," *Uspekhi Fizicheskikh Nauk* **169**(1), pp. 7–38, 1999.
6. S. Kay, "Can detectability be improved by adding noise?," *IEEE Signal Processing Letters* **7**(1), pp. 8–10, 2000.
7. G. P. Harmer, B. R. Davis, and D. Abbott, "A review of stochastic resonance: Circuits and measurement," *IEEE Transactions on Instrumentation and Measurement* **51**(2), pp. 299–309, 2002.
8. R. Benzi, A. Sutera, and A. Vulpiani, "The mechanism of stochastic resonance," *Journal of Physics A: Mathematical and General* **14**, pp. L453–L457, 1981.
9. S. Fauve and F. Heslot, "Stochastic resonance in a bistable system," *Physics Letters A* **97A**(1-2), pp. 5–7, 1983.
10. R. Benzi, A. Sutera, and A. Vulpiani, "Stochastic resonance in the Landau-Ginzburg equation," *Journal of Physics A: Mathematical and General* **18**, pp. 2239–2245, 1985.
11. B. McNamara, K. Wiesenfeld, and R. Roy, "Observation of stochastic resonance in a ring laser," *Physical Review Letters* **60**(25), pp. 2626–2629, 1988.
12. A. Bulsara, E. W. Jacobs, and T. Zhou, "Stochastic resonance in a single neuron model—theory and analog simulation," *Journal of Theoretical Biology* **152**, pp. 531–555, 1991.
13. A. Longtin, A. Bulsara, and F. Moss, "Time-interval sequences in bistable systems and the noise-induced transmission of information by sensory neurons," *Physical Review Letters* **67**(5), pp. 656–659, 1991.
14. J. K. Douglass, L. Wilkens, E. Pantazelou, and F. Moss, "Noise enhancement of information transfer in crayfish mechanoreceptors by stochastic resonance," *Nature* **365**, pp. 337–339, 1993.
15. K. Wiesenfeld, D. Pierson, E. Pantazelou, C. Dames, and F. Moss, "Stochastic resonance on a circle," *Physical Review Letters* **72**(14), pp. 2125–2129, 1994.
16. A. Bulsara and A. Zador, "Threshold detection of wideband signals: A noise induced maximum in the mutual information," *Physical Review E* **54**(3), pp. R2185–R2188, 1996.
17. D. S. Leonard and L. E. Reichl, "Stochastic resonance in a chemical-reaction," *Physical Review E* **49**, pp. 1734–1737, 1994.

18. S. M. Bezrukov and I. Voydanoy, “Noise-induced enhancement of signal transduction across voltage-dependent ion channels,” *Nature* **378**(6555), pp. 362–364, 1995.
19. A. D. Hibbs, A. L. Singsaas, E. W. Jacobs, A. R. Bulsara, J. J. Bekkedahl, and F. Moss, “Stochastic resonance in a superconducting loop with a Josephson-junction,” *Journal of Applied Physics* **77**, pp. 2582–2590, 1995.
20. A. Blarer and M. Doebeli, “Resonance effects and outbreaks in ecological time series,” *Ecology Letters* **2**, pp. 167–177, 1999.
21. X. M. Mao, K. Sun, and Q. Ouyang, “Stochastic resonance in a financial model,” *Chinese Physics* **11**, pp. 1106–1110, 2002.
22. R. Wallace, D. Wallace, and H. Andrews, “AIDS, tuberculosis, violent crime, and low birthweight in eight us metropolitan areas: Public policy, stochastic resonance, and the regional diffusion of inner-city markers,” *Environment and Planning A* **29**, pp. 525–555, 1997.
23. C. R. Doering, “Randomly rattled ratchets,” *Il Nuovo Cimento* **17**(7-8), pp. 685–697, 1995.
24. J. A. Paradiso and T. Starner, “Energy scavenging for mobile and wireless electronics,” *IEEE Pervasive Computing* **4**(1), pp. 18–27, 2005.
25. R. M. Gray and T. G. Stockham, “Dithered quantizers,” *IEEE Transactions on Information Theory* **39**(3), pp. 805–812, 1993.
26. R. A. Wannamaker, S. P. Lipshitz, J. Vanderkooy, and J. N. Wright, “A theory of non-subtractive dither,” *IEEE Transactions on Signal Processing* **48**(2), pp. 499–516, 2000.
27. G. P. Harmer and D. Abbott, “Losing strategies can win by Parrondo’s paradox,” *Nature* **402**, p. 864, 1999.
28. R. Toral, C. R. Mirasso, E. Hernandez-Garcia, and O. Piro, “Synchronization of chaotic systems by common random forcing,” in *Second International Conference on Unsolved Problems of Noise and Fluctuations, Adelaide, Australia*, D. Abbott and L. B. Kish, eds., American Institute of Physics, Melville, New York, 1999.
29. G. K. Basak, “Stabilization of dynamical systems by adding a colored noise,” *IEEE Transactions on Automatic Control* **46**(7), pp. 1107–1111, 2001.
30. A. Neiman, “Synchronizationlike phenomena in coupled stochastic bistable systems,” *Physical Review E* **49**(4), pp. 3484–3487, 1994.
31. J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behaviour*, Princeton University Press, 1944.
32. N. Gershenfeld, *The Nature of Mathematical Modeling*, Cambridge University Press, Cambridge, UK, 1999.
33. S. Kirkpatrick, C. D. Gelatt Jr, and M. P. Vecchi, “Optimisation by simulated annealing,” *Science* **220**, pp. 671–680, 1983.
34. J. J. Collins, C. C. Chow, and T. T. Imhoff, “Aperiodic stochastic resonance in excitable systems,” *Physical Review E* **52**(4), pp. R3321–R3324, 1995.
35. N. G. Stocks, “Suprathreshold stochastic resonance in multilevel threshold systems,” *Physical Review Letters* **84**(11), pp. 2310–2313, 2000.
36. N. G. Stocks, “Information transmission in parallel threshold arrays: Suprathreshold stochastic resonance,” *Physical Review E* **63**(4), p. 041114, 2001.
37. M. D. McDonnell, D. Abbott, and C. E. M. Pearce, “An analysis of noise enhanced information transmission in an array of comparators,” *Microelectronics Journal* **33**(12), pp. 1079–1089, 2002.
38. M. D. McDonnell, N. G. Stocks, C. E. M. Pearce, and D. Abbott, “Quantization in the presence of large amplitude threshold noise,” *Fluctuation and Noise Letters* **5**(3), pp. L457–L468, 2005.
39. J. Nocedal and S. J. Wright, *Numerical Optimization*, Springer-Verlag, New York, 1999.
40. M. D. McDonnell, N. G. Stocks, C. E. Pearce, and D. Abbott, “Optimal information transmission in nonlinear arrays through suprathreshold stochastic resonance.” To Appear in Physics Letters A (accepted 16 Nov. 2005), 2005.
41. N. G. Stocks and R. Mannella, “Generic noise enhanced coding in neuronal arrays,” *Physical Review E* **64**, p. 030902(R), 2001.
42. N. G. Stocks, D. Allingham, and R. P. Morse, “The application of suprathreshold stochastic resonance to cochlear implant coding,” *Fluctuation and Noise Letters* **2**(3), pp. L169–L181, 2002.