Analog to digital conversion using suprathreshold stochastic resonance

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

We present an analysis of the use of suprathreshold stochastic resonance for analog to digital conversion. Suprathreshold stochastic resonance is a phenomenon where the presence of internal or input noise provides the optimal response from a system of identical parallel threshold devices such as comparators or neurons. Under the conditions where this occurs, such a system is effectively a non-deterministic analog to digital converter. In this paper we compare the suprathreshold stochastic resonance effect to conventional analog to digital conversion by analyzing the the rate-distortion trade-off of each.

Keywords: Analog to Digital Conversion, ADC, noise, stochastic resonance, suprathreshold stochastic resonance, dithering, rate-distortion

1. INTRODUCTION

Analog to Digital Conversion (ADC) is a fundamental stage in the electronic storage and transmission of information. This process involves obtaining a sample of a signal, and its quantization to one of a finite number of levels. In this paper, we will consider theoretical measures of the performance of a stochastic quantization method, and compare with a conventional quantization scheme often used in ADCs.

We begin by describing some of the basic theory of quantization, in particular, the commonly used case of the uniform quantizer. We then briefly review a phenomenon known as stochastic resonance, and in particular, the case of Suprathreshold Stochastic Resonance (SSR), which we demonstrate can be viewed as a stochastic quantization scheme. We then compare the performance of SSR to the uniform quantizer, and show that conditions exist where SSR can outperform the uniform quantizer.

2. QUANTIZATION THEORY

Quantization of a signal or source consists of the partitioning of the signal into a discrete number of intervals, or cells. Certain rules specify which range of values of the signal get assigned to each cell. 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. The most common quantization scheme is the uniform scalar quantizer.[1] If the input signal, x, has an amplitude range between $[-m_p, m_p]$ then this quantizer 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

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

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

Smart Structures, Devices, and Systems II, edited by Said F. Al-Sarawi,
Proceedings of SPIE Vol. 5649 (SPIE, Bellingham, WA, 2005)
0277-786X/05/$15 · doi: 10.1117/12.582493

Thus, the output of the quantizer is the discrete signal, y, which can take on values between 0 and N. The reproduction value, $z(y)$ for the y-th cell is given by the midpoint of the interval number, $y \in 0, .., N$, in which an input sample lies. Thus,

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

Shannon's average mutual information measure[2] provides a measure of the rate of a quantizer, that is, the average number of bits per sample that the quantizer output contains about the input. In general, the rate will depend on the statistics of the input signal as well as the encoding. For a deterministic quantization, the rate is simply the entropy of the output encoding, which is given by

$$I(x, y) = H(y) = -\sum_{n=0}^{N} P_y(n) \log_2 P_y(n), \tag{3}$$

where $P_y(n)$ is the probability of output state n occurring. The maximum rate occurs when all output states are equally likely and is given by $I(x, y)_{\text{max}} = \log_2 (N + 1)$, that is, the number of bits at the output of the quantizer.

Information theory also tells us that the quantization of a signal will always cause some error in a reproduction of the original signal. This error is known as the distortion, and is most commonly measured by the mean square error between the original signal, and the reproduced signal.[1] Thus, the error is given by

$$\epsilon = x - z,$$

and the mean square distortion is

$$D_{\text{ms}} = \text{E}[(x - z)^2]. \tag{4}$$

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 mean square distortion power. If the input signal has power σ_x^2, then this can be expressed as

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

In general, design of a good quantizer for a given source involves a trade-off between rate and distortion. That is, decreasing the average distortion will always require an increase in the rate required. This trade-off is measured using the rate-distortion function,[3] often expressed as $R(D)$, where R is the rate, and D is the distortion measure. Shannon proved that a lower bound exists for $R(D)$ for a Gaussian source with power σ_x^2, and the mean square distortion measure,[3] and that this bound is given by

$$R(D_{\text{ms}}) = 0.5 \log_2 \left(\frac{\sigma_x^2}{D_{\text{ms}}} \right). \tag{5}$$

This equation says that no quantization scheme can achieve a distortion less than D_{ms} with a rate lower than $R(D_{\text{ms}})$ for a Gaussian source. Or in other words, a quantization scheme with rate R will provide a mean square distortion no smaller than D_{ms}. Note that (5) can be rearranged to give a lower bound on the SQNR in decibels of about $6.02R$. This is a well-known rule of thumb in quantizer design that states that a one bit increase in the number of output bits gives about a 6 dB increase in Signal to Noise Ratio (SNR).

To illustrate the application of this formula, consider the use of a uniform scalar quantizer to quantize a Gaussian source with probability density function (pdf) $P_x(x)$. Let the source have mean zero and variance σ_x^2. Let $m_p = 3\sigma_x$ and $N = 15$ so we have a four bit output. Then

$$P_y(n) = \int_{x=\theta_n}^{x=\theta_{n+1}} P_x(x) dx \quad n \in 0, .., N, \tag{6}$$

where here we have let $\theta_0 = -\infty$ and $\theta_{N+1} = \infty$.

The rate can now be calculated numerically from (6) and (3). To calculate the mean square distortion, note from (4) and (2) that for a given m_p and N, the distortion will depend only on the quantities $\mathrm{E}[y^2]$ and $\mathrm{E}[xy]$. These are given by

$$\mathrm{E}[y^2] = \sum_{n=0}^{N} n^2 P_y(n),$$

and

$$\mathrm{E}[xy] = \sum_{n=0}^{N} n \frac{\int_{x=\theta_n}^{x=\theta_{n+1}} x P_x(x) dx}{P_y(n)}.$$

We are now in a position to numerically calculate the rate and mean square distortion (or SQNR) for such a quantization, for any value of N. The result is plotted in Figure 1 where the rate (i.e. the mutual information) is shown on the y-axis (in bits per sample) and the SQNR on the x-axis (in dB). The theoretical lower bound is indicated by the thick line. Note that as expected, the actual rate for a given value of N and corresponding SQNR is always above the theoretical lower bound. Note that as N increases, the performance of the uniform quantizer gets worse and worse when compared with the lower bound. This is due to the fact that we have set $m_p = 3\sigma_x$. For larger N, a better quantization scheme is one that allows its maximum and minimum thresholds to be placed further along the Gaussian pdf's tails.

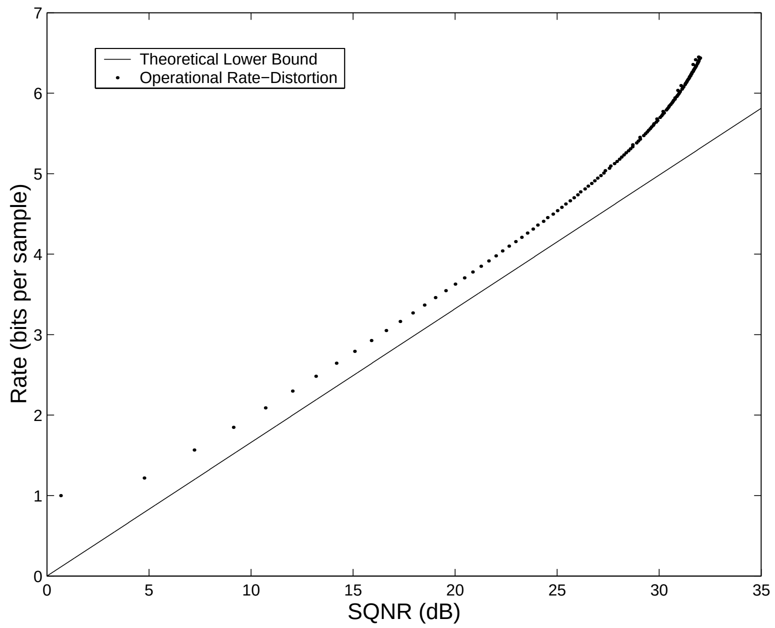

Figure 1. Plot of SQNR against corresponding rate (mutual information) for N between 1 and 127 for a Gaussian source with unity variance. The dots give the operational rate-distortion tradeoff for each value of N and the solid line gives the theoretical lower bound on the rate required for each value of SQNR.

3. SUPRATHRESHOLD STOCHASTIC RESONANCE

Stochastic Resonance[4–7](SR) is a term used to describe a system in which the presence of input or internal noise is required to achieve an optimal output from that system. Initial studies of SR in the 1980s appeared mainly in the physics literature, and generally consisted of the observation of SR in nonlinear systems driven by periodic signals subject to additive white noise. The measure used was the output SNR, which was usually shown to possess a maximum for some non zero value of input SNR. Examples of systems in which SR was described include a Schmitt Trigger circuit,[8–10] ring lasers,[11] neurons,[12,13] SQUIDS (super conducting quantum interference devices)[14] and ion channels.[15]

Since the early studies, the definition of SR has been extended to virtually any sort of phenomenon where noise or randomness can be useful in some way to improving a system's or algorithm's performance. Of particular note was the extension in 1995 to aperiodic input signals.[16] Such an extension generally requires the use of measures other than SNR, such as correlation coefficient[16,17] or mutual information.[18] Also of relevance here is the presentation of many studies of SR in systems consisting of a single threshold device.[19,20] SR can be easily understood in such a system by considering an entirely subthreshold signal subject to additive noise. If the output of the threshold device is a pulse every time the threshold is crossed, then no pulse will ever occur in the absence of noise. However if the right amount of noise is present, then output pulses will occur with some correlation to the amplitude of the input signal. Hence, the presence of noise is better than its absence.

One recent extension of SR was to consider the case of many identical threshold devices subject to *iid* additive noise. If each threshold device receives the same input signal, it was shown that if the overall output is the sum of the individual outputs, then SR occurs irrespective of whether the input signal is entirely subthreshold or not. Hence, the occurrence of SR in such a system was entitled Suprathreshold Stochastic Resonance (SSR).[21,22] 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 (shown in Fig. 2) can be modelled mathematically as follows.

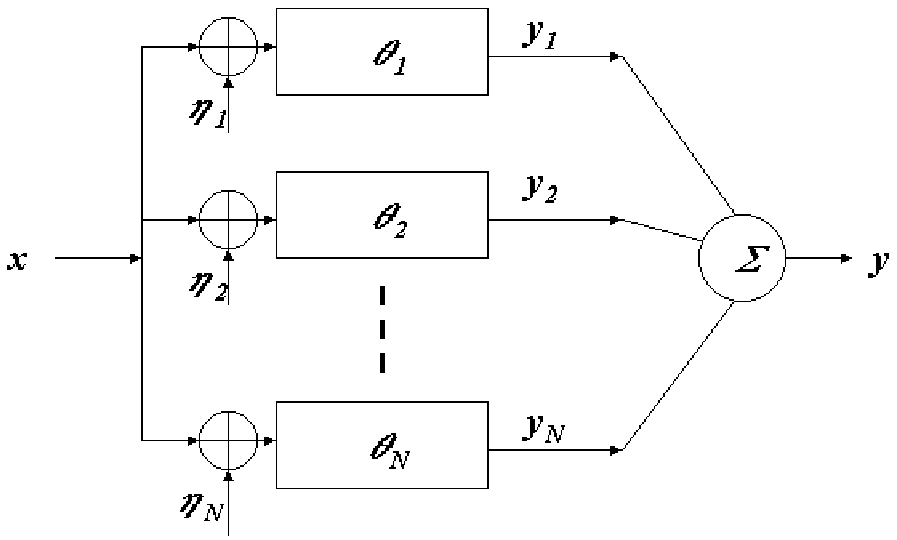

Figure 2. Array of N comparators. Each comparator receives the same input signal, x, and is subject to independent additive noise, η_n. The output from comparator n is unity if the sum of the signal and noise is greater than that comparator's threshold, θ_n, and zero otherwise. The overall output, y, is the sum of the individual comparator outputs.

Let the number of threshold devices be N. The n-th device is subject to continuously valued *iid* additive noise, η_n $(n = 1, .., N)$ and the output from each, y_n, is unity if the input signal, x, plus the noise is greater than its threshold, θ_n, 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. It can be considered as the number of devices that are currently "on." Thus

$$y = \frac{1}{2}\sum_{n=1}^{N} \text{sign}[x + \eta_n - \theta_n] + \frac{N}{2}.$$

This formula can also be applied to the uniform quantizer, where the thresholds are given by (1).

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, $P_{xy}(x,n)$. We commence by deriving a method of calculating $P_{xy}(x,n)$ for the general case of N arbitrary thresholds, and then simplify to the SSR case of all thresholds equal to the signal mean. Denote the pdf of the input signal as $P_x(x)$ and the probability mass function of the output as $P_y(n)$. Then we have, as a consequence of Bayes' theorem,

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

Integration of the the joint probability with respect to x gives

$$P_y(n) = \int_{-\infty}^{\infty} P(n|x)P_x(x)dx.$$

We will assume knowledge of $P_x(x)$ and derive a method for calculating $P(n|x)$. Recall we assume that the noise (with pdf $R(\eta)$), is *iid* at each comparator. Let $\hat{P}_n$ be the probability of device n being "on" (that is, signal plus noise exceeding the threshold θ_n), given the input signal x. Then

$$\hat{P}_n = \int_{\theta_n - x}^{\infty} R(\eta)d\eta = 1 - F_R(\theta_n - x), \tag{7}$$

where F_R is the cumulative distribution function of the noise and $n = 1,..,N$. If the noise has an even pdf then $\hat{P}_n = F_R(x - \theta_n)$. 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 $\{\theta_n\}$, $\{\hat{P}_n(x)\}$ can be calculated exactly for any value of x from (7), from which $P(n|x)$ can be found using an efficient recursive formula.[23] For the particular case when the thresholds all have the same value, then each $\hat{P}_n$ has the same value $\hat{P}$ and, as noted by Stocks[21] we have the binomial distribution

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

The mutual information, $I(x,y)$, between x and y is given by the entropy of the output, $H(y)$, less the conditional entropy of the output given the input, $H(y|x)$. As noted by Stocks,[21] the system can be interpreted as the transmission of information through a channel. $H(y)$ describes the maximum rate of the transmission, and $H(y|x)$ can be interpreted as the amount of encoded information about the input signal lost through the channel. Since the input to the array is continuously valued and the output is discretely valued, the channel can be considered to be semi-continuous. The mutual information through such a channel is given by

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

which can be calculated by numerical integration after applying the technique for calculating $P(n|x)$ described above. Unlike a conventional deterministic quantization scheme, there is some uncertainty about which output state a given input value x will be encoded to. Indeed, if the noise pdf has infinite support, then all output states can be achieved by any input value.

The mutual information for a Gaussian signal with variance of unity and a number of values of N is shown in Fig. 3 for a range of input SNR's. The threshold noise is also Gaussian. This plot replicates the results obtained by Stocks.[21]

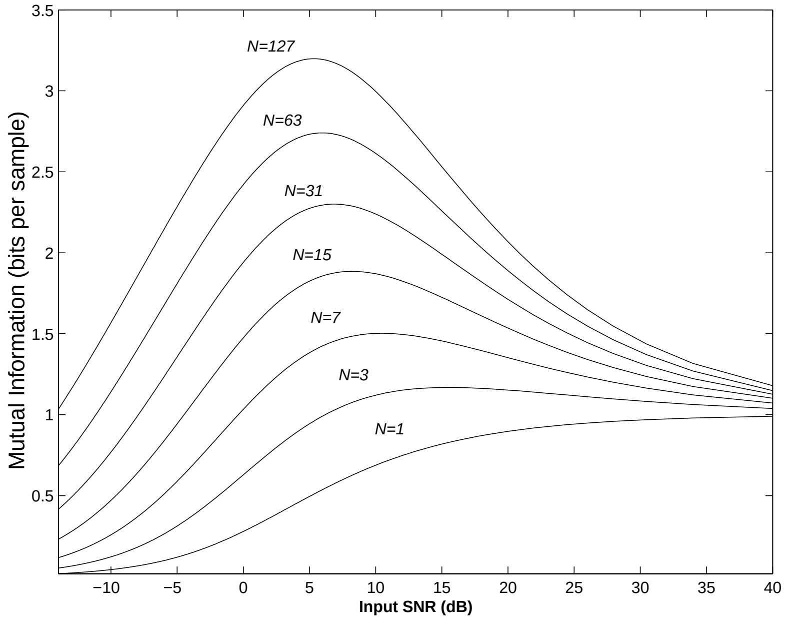

Figure 3. Plot of Mutual information against input SNR for Gaussian signal and noise, and a number of values of N. Note that the maximum of the mutual information occurs for a very low input SNR, but increases with increasing N.

4. DISTORTION AND SSR

Another way of interpreting the SSR model is that of a lossy source encoder or quantizer.[1,3] Due to its nondeterministic nature, the quantization achieved by SSR is stochastic. In conventional quantization theory or ADCs, dithering is a well known case of such a non-deterministic quantization.[24,25] Dithering is the addition of a random signal to a signal to be quantized, with the aim of controlling the statistical properties of the quantization noise. However SSR differs from dithering, as the term is usually applied, in several key ways. Like dithering, SSR can be considered as a quantizer in which all thresholds are random variables. However, since dithering is usually applied by adding a small random signal to the signal to be quantized, the net result is that the set of thresholds retains the same fixed separations for any given signal sample. In SSR, since there is independent noise added to each threshold, if this noise is large enough, it can be shown that the optimal thresholds all have the same value,[26] which is equal to the signal mean.[22] This is in contrast to conventional quantization in which noise free thresholds are always distributed deterministically across the signal dynamic range.

To consider SSR as a quantizer, we need to specify some means of reconstructing the input signal from an encoding, as in the uniform scalar quantizer. Consider a linear decoding of the output of the array, y, that is, one of the form

$$\hat{y} = ay + b, \quad y \in 0, .., N,$$

where a and b are constants for all values of y. Such a linear decoding will give a set of outputs that are evenly spaced. Consider only signals with an even pdf and a mean of zero. If we impose the condition that the mean of the output is also zero, this requires that $\mathrm{E}[\hat{y}] = a\mathrm{E}[y] + b = 0$. It is straightforward to show that $\mathrm{E}[y] = N/2$ and that therefore setting $a = -2b/N$ is sufficient to achieve a zero-mean decoding. Changing notation by letting $c = -b$ gives

$$\hat{y} = \frac{2c}{N}y - c, \quad y \in 0, .., N. \tag{8}$$

Thus, the mean square value of y is also its variance, since the mean is zero. It is also easy to show that the variance of the error, $\epsilon = \hat{y} - x$, given x is equal to the variance of $\hat{y}$ given x. Thus

$$\mathrm{var}[\epsilon|x] = \mathrm{var}[\hat{y}|x] = \frac{4c^2}{N^2}\mathrm{var}[y|x].$$

Therefore the conditional mean square distortion is

$$\begin{aligned} D_{\mathrm{ms}}(x) &= \mathrm{var}[\epsilon|x] + \mathrm{E}[\epsilon|x]^2 \\ &= \mathrm{var}[\hat{y}|x] + \mathrm{E}[\hat{y} - x|x]^2 \\ &= \mathrm{var}[\hat{y}|x] + \mathrm{E}[\hat{y}|x]^2 - 2x\mathrm{E}[y|x] + x^2 \\ &= \mathrm{E}[\hat{y}^2|x] - 2x\mathrm{E}[\hat{y}|x] + x^2. \end{aligned}$$

The total average mean square distortion is the expected value of the conditional distortion,

$$D_{\mathrm{ms}} = \mathrm{E}[\hat{y}^2] - 2\mathrm{E}[x\hat{y}] + \mathrm{E}[x^2]. \tag{9}$$

Note that $\mathrm{E}[\hat{y}^2]$ is the mean square value of the output, and $\mathrm{E}[x\hat{y}]$ is the correlation between x and $\hat{y}$. Substituting (8) into (9) and simplifying gives

$$D_{\mathrm{ms}} = \frac{4c^2}{N^2}\mathrm{E}[y^2] - \frac{4c}{N}\mathrm{E}[xy] + \mathrm{E}[x^2] - c^2. \tag{10}$$

Expressions for the mean square value of y given x and the expected value of y given x in terms of $\hat{P}$ are given in .[23] Substituting these into (10) gives

$$D_{\mathrm{ms}} = \frac{4c^2}{N}(\mathrm{E}[\hat{P}] + (N-1)\mathrm{E}[\hat{P}^2]) - 4c\mathrm{E}[x\hat{P}] + \mathrm{E}[x^2] - c^2.$$

What is the optimal linear decoding? The only unknown in (8) is c. The optimal value of c can be found by differentiating (10) with respect to c, setting the result to zero, and solving for c. The result is

$$c = \frac{N\mathrm{E}[xy]}{2\mathrm{var}[y]}. \tag{11}$$

Such a decoding is known as a Wiener decoding, and is the optimal linear decoding for the mean square error distortion.

Substituting (11) into (10) and simplifying gives

$$\begin{aligned} D_{\mathrm{ms}} &= \mathrm{E}[x^2] - \frac{\mathrm{E}[xy]^2}{\mathrm{var}[y]} \\ &= \mathrm{E}[x^2]\left(1 - \frac{\mathrm{E}[xy]^2}{\mathrm{E}[x^2]\mathrm{var}[y]}\right) \\ &= \mathrm{E}[x^2](1 - \rho_{xy}^2), \end{aligned}$$

where ρ_{xy} is the correlation coefficient[17] between x and y. Hence, the optimal mean square distortion is entirely dependant on the correlation coefficient.

Noting that $\mathrm{E}[x^2]$ is the input signal power, the SQNR can be written as

$$\mathrm{SQNR} = -10\log_{10}(1 - \rho_{xy}^2).$$

The SQNR for a Gaussian input signal with unit variance is shown in Fig. 4 for Gaussian noise, various values of N and a range of input SNRs. Note that the same qualitative behavior occurs as for the mutual information measure, with the difference being that the optimal input SNRs are slightly different.

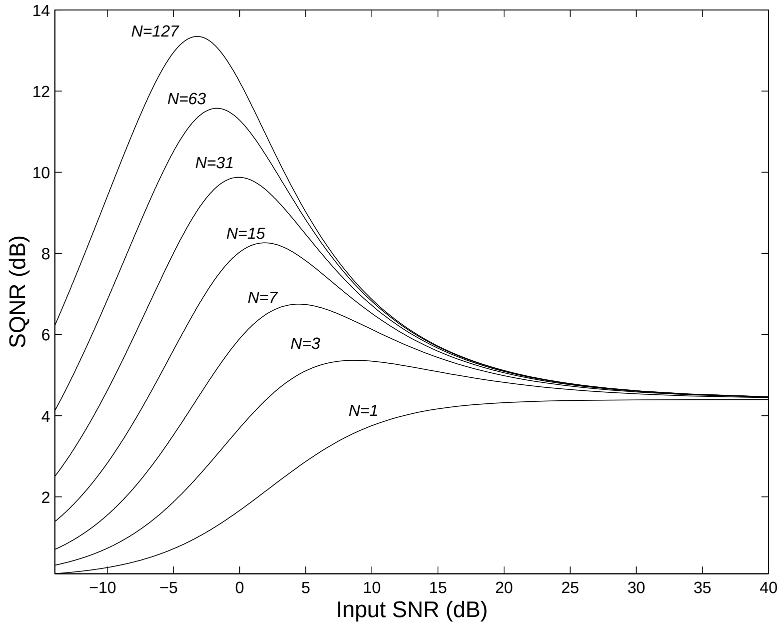

Figure 4. Plot of SQNR (in dB) against input SNR (in dB) for a Gaussian signal with unity variance, Gaussian noise, and a number of values of N. Note that the maximum of the SQNR increases with increasing N and occurs for a very low input SNR. The SQNR shows the same qualitative behavior as the mutual information, however the optimal values of input SNR are not the same.

5. RATE DISTORTION RESULTS

We are now in a position to examine the rate-distortion trade-off for SSR and compare it to the case of the uniform scalar quantizer. Fig. 5 shows the result of plotting the rates shown in Fig. 3 against the distortion shown in Fig. 4. 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 SQNR, as the input SNR decreases. The rate then reaches its maximum before the SQNR does. Then with continuing decreasing input SNR, the curve reaches its SQNR maximum, before curling back down towards the theoretical lower bound. 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. Since the main goal of a quantizer is to operate with maximum SQNR, this observation would indicate that the optimal value of input SNR to use, is the one which achieves the maximum SQNR, rather than the maximum rate.

A further observation is the fact that for very low input SNRs, this 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 7 bits) the maximum SQNR that can be achieved is about 13 dB. By contrast, for the same value of N, the uniform scalar quantizer provides an SQNR of over 30 dB. Alternatively, note that for a 13 dB SQNR, the uniform quantizer requires a 3 bit output whereas SSR requires 7 bits.

It is possible however, that the need to have far many more thresholds in the case of SSR to provide the same performance as a uniform scalar quantizer, is offset by 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 an ADC 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.[27]

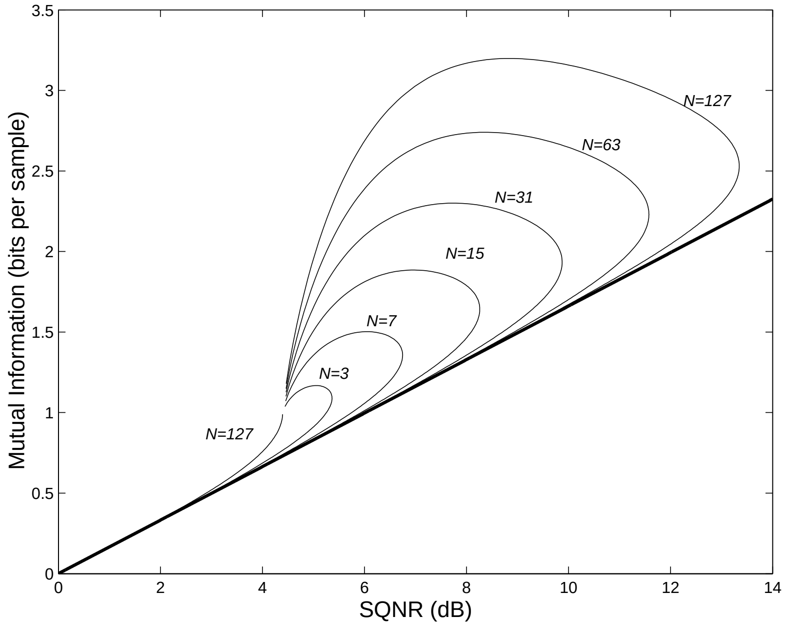

Figure 5. Plot of mutual information against SQNR (in dB) for a Gaussian signal with unity variance, Gaussian noise, and a number of values of N. The thick solid line shows the theoretical lower bound on the rate required for a given distortion. Note that there are in general two values of mutual information that achieve the same SQNR and two values of SQNR that achieve the same mutual information. This is due to the concave nature of the plots of mutual information and SQNR against input SNR.

6. CONCLUSIONS

We have demonstrated that suprathreshold stochastic resonance can be described as a stochastic quantization method. Quantization methods can be characterized by their rate-distortion trade-off, and we have shown here that SSR is no exception. For a Gaussian source, and the optimal linear decoding, we have found that the best operating input SNR is the one that provides the minimum mean square distortion, rather than the SNR that provides the maximum mutual information. When compared with a uniform scalar quantizer, SSR requires more than twice as many output bits to achieve the same average distortion performance. However this could be worthwhile if implemented in practise, as ADC's such as flash ADC's require N unique threshold values, whereas in SSR all N thresholds have the same value. Furthermore, if implemented in a situation where large amplitude threshold noise is unavoidable, the performance of SSR is on a par with the uniform quantizer when each have the same number of output bits.

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**, pp. 2325–2383, 1998.
2. T.M. Cover and J.A. Thomas, *Elements of Information Theory*, John Wiley and Sons, New York, 1991.
3. T. Berger and J.D. Gibson, "Lossy source coding," *IEEE Transactions on Information Theory* **44**, pp. 2693–2723, 1998.

4. L. Gammaitoni, P. Hänggi, P. Jung, and F. Marchesoni, "Stochastic resonance," *Reviews of Modern Physics* **70**, 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**, pp. 7–38, 1999.
6. S. Kay, "Can detectability be improved by adding noise?," *IEEE Signal Processing Letters* **7**, pp. 8–10, 2000.
7. D. Abbott, "Overview: Unsolved problems of noise and fluctuations," *Chaos* **11**, pp. 526–538, 2001.
8. S. Fauve and F. Heslot, "Stochastic resonance in a bistable system," *Physics Letters A* **97A**, pp. 5–7, 1983.
9. V.I. Melnikov, "Schmitt trigger: a solvable model of stochastic resonance," *Physical Review E* **48**, pp. 2481–2489, 1993.
10. G.P. Harmer and D. Abbott, "Simulation of circuits demonstrating stochastic resonance," *Microelectronics Journal* **31**, pp. 553–559, 2000.
11. B. McNamara, K. Kiesenfeld, and R. Roy, "Observation of stochastic resonance in a ring laser," *Physical Review Letters* **60**, pp. 2626–2629, 1988.
12. A. Longtin, "Stochastic resonance in neuron models," *Journal of Statistical Physics* **70**, pp. 309–327, 1993.
13. 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.
14. K. Wiesenfeld and F. Moss, "Stochastic resonance and the benefits of noise: from ice ages to crayfish and SQUIDS," *Nature* **373**, pp. 33–36, 1995.
15. S.M. Bezrukov and I. Voydanoy, "Noise-induced enhancement of signal transduction across voltage-dependent ion channels," *Nature* **378**, pp. 362–364, 1995.
16. J.J. Collins, C.C. Chow, and T.T. Imhoff, "Aperiodic stochastic resonance in excitable systems," *Physical Review E* **52**, pp. R3321–R3324, 1995.
17. M.D. McDonnell, D. Abbott, and C.E.M. Pearce, "A characterization of suprathreshold stochastic resonance in an array of comparators by correlation coefficient," *Fluctuation and Noise Letters* **2**, pp. L205–L220, 2002.
18. C. Heneghan, C.C. Chow, J.J. Collins, T.T. Imhoff, S.B. Lowen, and M.C. Teich, "Information measures quantifying aperiodic stochastic resonance," *Physical Review E* **54**, pp. R2228–R2231, 1996.
19. P. Jung and G. Mayer-Kress, "Stochastic resonance in threshold devices," *Il Nuovo Cimento* **17**, pp. 827–834, 1995.
20. Z. Gingl, L.B. Kiss, and F. Moss, "Non-dynamical stochastic resonance: Theory and experiments with white and arbitrarily coloured noise," *EuroPhysics Letters* **29**, pp. 191–196, 1995.
21. N.G. Stocks, "Suprathreshold stochastic resonance in multilevel threshold systems," *Physical Review Letters* **84**, pp. 2310–2313, 2000.
22. N.G. Stocks, "Information transmission in parallel threshold arrays: Suprathreshold stochastic resonance," *Physical Review E* **63**, 041114, 2001.
23. 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**, pp. 1079–1089, 2002.
24. R.M. Gray and T.G. Stockham, "Dithered quantizers," *IEEE Transactions on Information Theory* **39**, pp. 805–812, 1993.
25. R.A. Wannamaker, S.P. Lipshitz, J. Vanderkooy, and J.N. Wright, "A theory of non-subtractive dither," *IEEE Transactions on Signal Processing* **48**, pp. 499–516, 2000.
26. M.D. McDonnell, N.G. Stocks, C.E.M. Pearce, and D. Abbott, "Point singularities and suprathreshold stochastic resonance in optimal coding." Preprint arXiv:cond-mat/0409528, 2004.
27. M.D. McDonnell, N.G. Stocks, C.E.M. Pearce, and D. Abbott, "Neural mechanisms for analog to digital conversion," in *BioMEMS and Nanoengineering*, D. V. Nicolau, ed., *Proceedings of SPIE* **5275**, pp. 278–286, 2003.