Decimation vs. quantization for data compression in TDOA systems

Mark L. Fowler†

Department of Electrical Engineering
State University of New York at Binghamton

ABSTRACT

The location of an electromagnetic emitter is commonly estimated by intercepting its signal and then sharing the data among several platforms. Doing this in a timely fashion requires effective data compression. Previous data compression efforts have focused on minimizing the mean-square error (MSE) due to compression. However, this criterion is likely to fall short because it fails to exploit how the signal's structure impacts the parameter estimates. Because TDOA accuracy depends on the signal's RMS bandwidth, compression techniques that can significantly reduce the amount of data while negligibly impacting the RMS bandwidth have great potential. We show that it is possible to exploit this idea by balancing the impacts of simple filtering/decimation and quantization and derive a criterion that determines an optimal balance between the amount of decimation and the level of quantization. This criterion is then used to show that by using a combination of decimation and quantization it is possible to meet requirements on data transfer time that can't be met through quantization alone. Furthermore, when quantization-alone approaches can meet the data transfer time requirement, we demonstrate that the decimation/quantization approach can lead to better TDOA accuracies. Rate-distortion curves are plotted to show the effectiveness of the approach.

Keywords: data compression, emitter location, time-difference-of-arrival, TDOA, quantization, decimation

1. INTRODUCTION

An effective way to locate electromagnetic emitters is to measure the time-difference-of-arrival (TDOA) between pairs of signals received at geographically separated sites[1,2,3]. The measurement of TDOA between these signals is done by coherently cross-correlating the signal pairs,[2,3] and requires that the signal samples of the two signals are available at a common site, which is generally accomplished by transferring the signal samples over a data link from one site to the other site. An important aspect of this that is not widely addressed in the literature is that often the available data link rate is insufficient to accomplish the transfer within the time requirement unless some form of lossy data compression is employed. For the case of Gaussian signals and noises, Matthiesen and Miller[4] established bounds on the rate-distortion performance for the TDOA problem and compared them to the performance achievable using scalar quantizers, where distortion is measured in terms of lost SNR due to the mean square error (MSE) of lossy compression. However, these results are not applicable when locating radar and communication emitters because the signals encountered are not Gaussian.

The two signals to be correlated are the complex envelopes of the received RF signals having RF bandwidth B. The complex envelopes can then be sampled at $F_s \geq B$ complex-valued samples per second; for simplicity here we will assume critical sampling, for which $F_s = B$. The signal samples are assumed to be quantized using $2b$ bits per complex sample (b bits for the real part, b bits for the imaginary part), where b is large enough to ensure fine quantization. The two noisy signals to be correlated are notated as

† Correspondence: mfowler@binghamton.edu

$$\begin{aligned}\hat{s}(k) &= s(k) + n(k) \\ &= [s_r(k) + js_i(k)] + [n_r(k) + jn_i(k)] \\ \hat{d}(k) &= d(k) + v(k) \\ &= [d_r(k) + jd_i(k)] + [v_r(k) + jv_i(k)]\end{aligned} \tag{1}$$

where $s(k)$ and $d(k)$ are the complex baseband signals of interest and $n(k)$ and $v(k)$ are complex white Gaussian noises, each with real and imaginary parts notated as indicated. The signal $d(k)$ is a delayed version of $s(k)$. The signal-to-noise ratios (SNR) for these two signals are denoted *SNR* and *DNR*, respectively[‡].

To cross correlate these two signals one of them (assumed to be $\hat{s}(k)$ here) is compressed, transferred to the other site, and then decompressed before cross-correlation, as shown in Figure 1. Signal $\hat{s}(k)$ has SNR of $SNR_q < SNR$ after lossy compression/decompression[5], and the output SNR after cross-correlation is given by

$$\begin{aligned}SNR_o &= \frac{WT}{\dfrac{1}{SNR_q} + \dfrac{1}{DNR} + \dfrac{1}{SNR_q DNR}} \\ &\stackrel{\Delta}{=} WT \times SNR_{eff}\end{aligned} \tag{2}$$

where *WT* is the time-bandwidth product (or coherent processing gain), with *W* being the noise bandwidth of the receiver and *T* being the duration of the received signal and SNR_{eff} is a so-called effective SNR[3]. The accuracies of the TDOA estimates are governed by the Cramer-Rao bound (CRB) for TDOA given by[3]

$$\sigma_{TDOA} \geq \frac{1}{2\pi B_{rms}\sqrt{SNR_o}}, \tag{3}$$

where B_{rms} is the signal's RMS (or Gabor) bandwidth in Hz given by

$$B_{rms}^2 = \frac{\int f^2 |S(f)|^2 df}{\int |S(f)|^2 df}$$

with $S(f)$ being the Fourier transform of the signal $s(k)$.

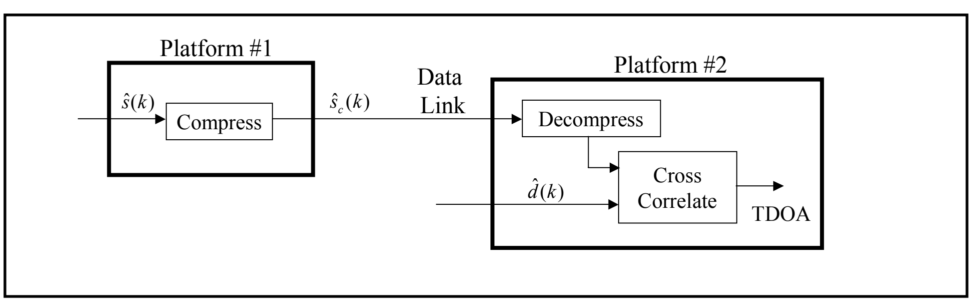

Figure 1: System Configuration for Compression

[‡] SNR (non-italic) represents an acronym for signal-to-noise ratio; *SNR* (italic) represents the SNR for $\hat{s}(k)$.

2. NON-MSE DISTORTION CRITERIA

To ensure maximum performance it is necessary to employ a compression method that is designed specifically for this application. However, much of the past effort in developing general lossy compression methods has focused on minimizing the MSE due to compression; furthermore, even compression schemes developed for TDOA applications have also limited their focus to minimizing the MSE[4,5,6]. But when the goal is to estimate TDOA, the minimum MSE criterion is likely to fall short because it fails to exploit how the signal's structure impacts the parameter estimates. In such applications it is crucial that the compression methods minimize the impact on the TDOA estimation performance rather than stressing minimization of MSE as is common in many compression techniques. Achieving significant compression gains for the emitter location problem requires exploitation of how signal characteristics impact the TDOA accuracy. For example, the CRB in (3) shows that the TDOA accuracy depends on the signal's RMS bandwidth. Thus, compression techniques that can significantly reduce the amount of data while negligibly impacting the signal's RMS bandwidth have great potential. We show here that it is possible to exploit this idea through simple filtering and decimation together with quantization to meet requirements on data transfer time that can't be met through quantization-only approaches designed to minimize MSE. These results are encouraging because it is expected that non-MSE approaches more advanced than simple filtering and decimation will enable even larger improvements in performance.

It is desired to minimize σ_{TDOA} while adhering to a fixed data link rate constraint. One method for meeting the data rate constraint is to further quantize the digital samples to a small number of bits such that the data rate constraint is met[5]. This coarse quantization process reduces the SNR of the signal sent over the data link and can impact the performance of the cross-correlation when the number of bits per sample is made so small that the quantization noise power is comparable to the receiver noise power. Another method for meeting the data rate constraint would retain the full number of bits/sample, but would filter and decimate the signal to reduce the samples/second that are required to be sent over the data link. However, reducing the bandwidth of the signal also reduces the performance because the TDOA accuracy depends inversely on the signal's RMS bandwidth. As a third alternative, it is possible to operate somewhere between these two extremes by simultaneously quantizing and decimating, presumably doing each to a lesser degree than would be done in either of the two methods mentioned above. Thus, the goal here is to determine the proper trade off between quantization and decimation.

If we consider that the signals are collected for T seconds, then the total number of bits collected is $2bBT$. System requirements often specify a fixed length of time for the data transmission. Thus, if the transfer is constrained to occur within T_l seconds and the data link can transfer bits at the rate R_l bits/second then the total number of bits collected is constrained to satisfy $2bBT \leq R_l T_l$. Equivalently, if we define $R = R_l T_l / T$ as a fixed effective rate and assume equality in the constraint (i.e., fully utilize the allocated data link resources) we get

$$R = 2Bb. \tag{4}$$

This requirement may be achieved in various ways by selecting appropriate values of B and b, where different values of B would be obtained by filtering and decimating to a lower sample rate, and different values of b would be obtained through coarse quantization. Equation (4) links the values of B, b, and R through the data link constraint C specifying any two specifies the other. A subtle aspect of this relationship is that strict application of Equation (4) implies that bandwidth B is allowed to increase without bound as b decreases; however, the signal itself imposes an upper bound on this bandwidth, which together with the specified rate R sets a lower bound on the number of bits (quantizing any farther underutilizes the available rate).

Our interest lies in the scenario where the required effective rate R is so low that when Equation (4) is solved for b under the condition of using the signal's full bandwidth, the resulting number of bits is extremely small (say 1 or 2 bits per real/imaginary signal sample). In other words, to achieve the required rate while using the signal's full bandwidth would require unreasonably excessive quantization that could result in undesirable effects due to excessive nonlinearity of such a quantizer. For example, consider the case where the required effective rate is 32 kbps and the signal's bandwidth is 8 kHz. Equation (4) suggests that to achieve the required rate while using the signal's full bandwidth would require $b = 32 \times 10^3 / (2 \times 8000) = 2$ bits per sample. This level of quantization is very likely to be unacceptable due to excessive nonlinear effects, so we alternatively quantize <u>and</u> decimate. How do we find the best balance between quantization and decimation so as to optimize the accuracy of the TDOA estimate? We first answer this for the specific case where the system performance is limited by the quantization noise and then develop the general result. Prior to that, though, we present background results for quantization-only and decimation-only.

3. QUANTIZE-ONLY & DECIMATE-ONLY

1. Quantize-Only

As mentioned above, the number of bits used in the quantization can affect the SNR of the signal quantized and hence the output SNR. If we quantize the real and imaginary parts of $s(t)$ each using b bits then its SNR after quantization becomes[5]

$$SNR_q(b) = \frac{SNR}{1 + \alpha^2 \; SNR \left(\frac{2^{-2b}}{3} \right)}, \tag{5}$$

where α is the signal's peak factor (i.e., the ratio of the signal's peak value to its RMS value). Thus the impact of quantization on TDOA estimation can be assessed by using (5) in (2) and then using the result in (3).

Because the effective link rate constraint is often quite restrictive, a specific example of interest is the case where b is small and the quantization noise dominates the receiver noises. For that case (5) becomes $SNR_q(b) = 3\,2^{2b}/\alpha^2$, and when it is also much smaller than DNR, the output SNR becomes

$$SNR_o(b) = \frac{3\,BT}{\alpha^2} 2^{2b}. \tag{6}$$

This then impacts the TDOA accuracy as a function of b according to

$$\sigma_{TDOA}(b) \geq \frac{\alpha}{2\pi B_{rms} 2^b \sqrt{3BT}}. \tag{7}$$

For insight, we further consider the case where the signal spectrum is flat, for which case we have that $2\pi B_{rms} = 1.8B$, which then gives

$$\sigma_{TDOA}(b) \geq \frac{\alpha}{1.8 B^{3/2} 2^b \sqrt{3T}}. \tag{8}$$

If we use the constraint in (4) to fix b at the level needed to achieve the rate constraint, we write $2^b = 2^{R/2B}$ and get

$$\sigma_{TDOA}(R) \geq \frac{\alpha 2^{2B}}{1.8 B^{3/2} 2^R \sqrt{3T}}. \tag{9}$$

From this we can see the impact of R and B. For a given signal bandwidth B, σ_{TDOA} goes according to 2^{-R}. More interestingly, for a fixed effective rate R, σ_{TDOA} goes according to $2^B / B^{3/2}$, so increasing bandwidth (i.e., locating emitters with larger bandwidths) actually increases the estimation error, which is contrary to the expectation when there is no rate constraint that TDOA can be estimated more accurately for larger bandwidth signals. Thus, when the effective rate is constrained, the reduction of the signal's bandwidth may actually be helpful.

A further argument for bandwidth reduction is when the rate constraint is so severe that it can't be met without reducing the signal's bandwidth. For example, if R = 24 kbps and B = 14 kHz, then the rate constraint requires that $b<1$, which is clearly not a feasible approach. Even when the rate constraint requires that b be on the order of 1 or 2 bits, the severe nonlinearity of such coarse quantization can seriously degrade the performance beyond what is indicated by the results developed above. In such cases it is necessary to consider some alternative to quantization. One such alternative is filtering and decimation.

2. Decimate-Only

Here we consider the impact of filtering and then decimating the two signals to be cross-correlated. Here, it is through decimation that we strive to meet the rate constraint. For simplicity we consider using ideal lowpass filters operating on the complex-valued baseband signal and we do not restrict the decimation factor to rational values, as would be done in practice. Thus, if we choose the filter such that the bandwidth is reduced by some factor γ with $0<\gamma<1$ then we can reduce the sampling rate by the factor γ also. Obviously, for practical signals, as we change the filter's cutoff we will change the signal's SNR and its RMS bandwidth; how these quantities change with the cutoff depends on the signal's spectral shape. Again for simplicity yet insight, we will assume that the signal's spectrum is flat, so that the effective SNR of the filtered signals will not change with the cutoff frequency. In practice though, since the signal's spectrum typically trails off at high frequency, the effective SNR will vary to some degree as the signal is filtered. If the filter has cutoff frequency $W_f\,/\,2$, then the RMS bandwidth becomes $2\pi B_{rms}=1.8W_f$, the time-bandwidth product becomes TW_f, and the output SNR becomes $SNR_o(W_f)=TW_f\times SNR_{eff}$. Then, the TDOA accuracy as a function of the filtered bandwidth becomes

$$\sigma_{TDOA}(W_f)\geq\frac{1}{1.8\,W_f^{3/2}\sqrt{T\;SNR_{eff}}}. \tag{10}$$

From this we see (as expected) that the accuracy becomes worse as we filter and decimate more. If, however, we put in the constraint on effective rate given in (4), namely that for a fixed number of bits $b_{\max}$ corresponding to the full amount available, the amount of filtering needed to meet the effective rate R is $W_f=R\,/\,2b_{\max}$. Then, if we filter to this rate-constrained bandwidth we have that

$$\sigma_{TDOA}(R)\geq\frac{\left(2b_{\max}\right)^{3/2}}{1.8\,R^{3/2}\sqrt{T\;SNR_{eff}\,(b_{\max})}}. \tag{11}$$

This shows that for the decimation-only case, σ_{TDOA} goes according to $R^{-3/2}$ (an inverse power law) rather than the much more rapid 2^{-R} (an exponential law) obtained above for the quantization-only case.

4. JOINT DECIMATION AND QUANTIZATION

Now we'd like to investigate the optimal trade-off between decimation and quantization. As before, let the received signals be filtered and decimated to a bandwidth of W_f and assume that the signals= spectra are flat so that the two SNRs don't depend on W_f. After filtering and decimation, the signal to be transmitted is quantized using 2b bits per complex sample (b bits for the real part and b bits for the imaginary part). The result is that the decimated and quantized signal has SNR given by (5). Now, the output SNR depends on the filtered bandwidth and the quantization level according to

$$\begin{aligned}SNR_o(W_f,b)&=\frac{W_fT}{\dfrac{1}{SNR_q(b)}+\dfrac{1}{DNR}+\dfrac{1}{SNR_q(b)\,DNR}}\\&=W_fT\;SNR_{eff}(b),\end{aligned} \tag{12}$$

where *DNR* is the SNR of $d(t)$, the signal that is not quantized. Using Equation (12) in Equation (3) gives a bound on TDOA accuracy that depends on the amounts of decimation and quantization, and is given by

$$\sigma_{TDOA}(W_f, b) \geq \frac{1}{1.8\, W_f^{3/2} \sqrt{T\, SNR_{eff}(b)}}. \tag{13}$$

This result has no constraint on the effective rate; it simply shows the impact of W_f and *b* on the TDOA accuracy. However, we wish to consider the rate constrained case, so the effective rate constraint gives $W_f = R/2b$, which after use in (13) removes the dependence on W_f and gives

$$\sigma_{TDOA}(b) \geq \frac{2^{3/2}}{1.8 R^{3/2} \sqrt{T}} \left[b^{3/2} \sqrt{\frac{1}{SNR_q(b)} + \frac{1}{DNR} + \frac{1}{SNR_q(b) DNR}} \right], \tag{14}$$

where it is really the bracketed term that is of interest here, since it shows the tradeoff between decimation and quantization, and can be considered as a decimation-quantization performance factor (for which smaller is better). For notational ease we introduce a notation for the bracketed term in (14), namely

$$\Gamma(b, SNR, DNR, \alpha) = b^{3/2} \sqrt{\frac{1}{SNR_q(b)} + \frac{1}{DNR} + \frac{1}{SNR_q(b) DNR}}. \tag{15}$$

It is important to remember that (14) includes the rate constraint, so for a fixed *R*, increasing *b* necessarily decreases W_f, and vice versa. The nonbracketed term in (14) just scales the result up or down depending on the values of the system parameters *R* and *T*. However, one important insight does come from the first term: the bound on σ_{TDOA} varies as the -3/2 power of the rate R; thus, if you double the allowable rate you get almost three times better accuracy, and if you quadruple the allowable rate you get eight times better accuracy. The reason that increasing the data rate improves the accuracy is because we have constrained the time available to transmit the data, so increasing the data rate allows an increase in the amount of information about the signal that can be transmitted. This is an important insight into the system design issues.

To compute the performance factor Γ for a given set of *b*, *SNR*, *DNR*, and α values we first compute $SNR_q(b)$ using (5), and the result is used in (15). Plots of Γ versus *b*, parameterized by α, *SNR*, and *DNR* reveal the proper way to choose the optimal value of *b*; that is, how to tradeoff decimation and quantization. Note that the value of R does not affect these curves; therefore, the optimal level of quantization is *not* set by the allowable data rate. Instead, the optimal degree of quantization is set by the interplay between *SNR*, *DNR*, and the peak factor α of the signal to be quantized. Once this optimal number of bits *b* is determined, the appropriate amount of decimation is determined using $W_f = R/2b$, given the allowable effective rate *R*. To investigate the characteristics of this result we consider the following examples.

Example: High SNRs

Say we have the following signal scenario: *DNR* = 60 dB, *SNR* = 30 dB, α = 3.5, T = 1 s, and the signal's available bandwidth is *B* = 4 kHz. Also assume that the original signal samples were done with *b* = 10 bits. Consider the case where the data link rate is R_l = 2.4 kbps and the link time constraint is T_l = 10 s, then the effective rate is *R* = 24 kbps. We now consider three ways to achieve this desired effective rate: quantize only, decimate only, and quantize and decimate. We also compare the results for these three cases to the case when no decimation or quantization is needed to meet the limit on the rate.

Quantize Only: We use Equation (4) to determine the number of bits to which we must quantize in order to meet the effective rate requirement yet retain the full signal bandwidth:

$$\begin{aligned} b &= \frac{R}{2B} \\ &= \frac{24000}{2 \times 4000} \\ &= 3 \text{ bits} \end{aligned}$$

Now using $b = 3$ bits, $W_f = 4000$, and $2\pi B_{rms} = 1.8 \times 4000$ along with the other design parameters in Equations (5), (12), and (3) gives the lower bound of

$$\sigma_{TDOA} \geq 5.6 \times 10^{-7}$$

Decimate Only: Here we intend to use all of the available bits ($b = 10$ bits), but to decimate the signal to meet the data rate requirement. We use Equation (4) to determine the required reduced bandwidth:

$$\begin{aligned} W_f &= \frac{R}{2\,b} \\ &= 1.2 \text{ kHz} \\ &= 3B/10. \end{aligned}$$

Thus, we must decimate by a factor of 10/3. Now using $b = 10$ bits , $W_f = 1200$, and $2\pi B_{rms} = 1.8 \times 1200$ along with the other design parameters in Equations (5), (12), and (3) gives

$$\sigma_{TDOA} \geq 4.2 \times 10^{-7},$$

which is ¾ of the value achieved using only quantization.

Quantize and Decimate: Here we wish to find the optimum balance between quantization and decimation. We use the plot of Γ vs. b in Figure 2 for the case of DNR = 60 dB, SNR = 30 dB, and $\alpha = 3.5$ to determine that we should quantize to $b = 6$ bits to get the lowest value (optimal) for Γ. Then using $b = 6$ in Equation (4) gives the required filtered bandwidth to be

$$\begin{aligned} W_f &= \frac{R}{2\,b} \\ &= 2 \text{ kHz} \\ &= B/2. \end{aligned}$$

Thus the optimal trade-off is to decimate by a factor of 2 and quantize to 6 bits. Using these values gives

$$\sigma_{TDOA} \geq 2.8 \times 10^{-7},$$

which is ½ of the value achieved using quantization alone and 2/3 of the value achieved using decimation alone. Now if we could double the data rate to 48 kbps, then the number of bits we should use does not change, but we could double the useable bandwidth to 4 kHz, the full signal bandwidth. According to the dependence on R shown in (14), this will improve the performance by a factor of $2^{3/2} \approx 2.8$, so you get almost three times the accuracy for twice the data rate. More generally, the rate-distortion curve for the quantize/decimate method for this case is shown in Figure 3.

No Quantization or Decimation:
If you further increased the data rate to 80 kbps, then you would not need to perform any quantization or decimation. In that case you would compute $SNR_o(4000,10)$ using Equation (12) and then use it in Equation (2) with $2\pi B_{rms} = 1.8\times 4000$ to get

$$\sigma_{TDOA} \geq 0.7\times 10^{-7},$$

which is the best that can be done because there is no rate constraint. These results are summarized in Table 1.

Table 1: Summary of Results for High SNR Example

Method	Achieved R (kbps)	b (bits)	BW (kHz)	σ_{TDOA} (ns)
Quantize Only	24	3	4	560
Decimate Only	24	10	1.2	420
Quantize & Decimate	24	6	2	280
No Compression	80	10	4	70

Now, why did we not change the number of bits used when we increased the rate from 24 kbps to 48 kbps, but we did change it when we further increased the rate 80 kbps? Because when we hit 48 kbps we were using all of the signal's bandwidth; if we increase the rate beyond that we can no longer gain performance by increasing the useable bandwidth, so the only way to take advantage of the increased rate is to spend more bits per sample. When going from 24 kbps to 48 kbps we have a choice of how to use this extra rate: we can use more bandwidth or more bits C our analysis says that we are better off increasing the bandwidth and leaving the degree of quantization alone (however, once we increase the bandwidth to exploit the signal's full bandwidth, then we should increase the number of bits). These ideas are clearly seen in the rate-distortion curves shown in Figure 3 where for rates above *R* = 48 kbps it is seen for this case that the quantization & decimation method is identical to the quantization only method.

Example: Low SNRs

We repeat the above example for the case when *SNR* = 10 dB and *DNR* = 20 dB; all other parameter values remain the same. Because the SNR values don't set the number of bits for quantize-only and the filtered bandwidth for decimate-only, those values do not change for this case (although the achieved accuracy does change because of the lower SNR values). For the quantize/decimate case, things do change; in particular, the optimal trade-off between quantization and decimation changes. The plot of performance factor Γ for this case is given in Figure 4, where it can be seen that it is monotonically increasing with the number of bits *b*. Thus it would appear that we should choose *b*=1; however, such a small value would surely lead to poor performance from the excessive nonlinearity due to such extreme quantization, where a reasonable minimum value on *b* is 3 or 4 bits. Choosing *b* = 4 bits then specifies that the filtered bandwidth should be

$$\begin{aligned} W_f &= \frac{R}{2b} \\ &= 3\,\text{kHz} \\ &= 3B/4, \end{aligned}$$

so we would decimate by a factor of 4/3. A plot of the rate-distortion curves are given in Figure 5, where for rates higher than *R* = 32 kbps it is seen that the quantize/decimate method is identical to the quantize only method.

Comparing these two examples provides much insight. From Figure 2 and Figure 4 we see that the optimal number of bits is larger at high SNRs than at low SNRs. This says that at high SNRs the impact of quantization on SNR_q becomes apparent at higher value of *b* (see (5)) and therefore the balance between decimation and quantization leans more toward decimation; however, at low SNRs the value of *b* can be made quite small (in theory) before its effect on SNR_q becomes apparent (again see (5)) so the balance between decimation and quantization leans more toward quantization. In fact, in the low SNR case the balance leans so much towards quantization that the value of Γ decreases monotonically as *b* decreases, indicating that the balance between quantization and decimation is never really met. However, due to excessive nonlinear effects, such low *b* values are not recommended and therefore the practical balance must lean more toward decimation than the theory states.

These characteristics are also seen in the rate-distortion curves shown in Figure 3 and Figure 5. For the high SNR case in Figure 3 we see that the quantize/decimate method uniformly outperforms both quantize-only and decimate-only (except at rates above 48 kbps where quantize-only and quantize/decimate are equivalent because the quantize/decimate method uses the full signal BW for those rates). It should also be observed that for the high SNR case, decimate-only is better than quantize-only at low rates but not at high rates. Thus, for the high SNR case we see that, both mathematically and practically, the quantize/decimate approach is favored at all effective rates. Figure 6 shows the ratio of σ_{TDOA} for quantize-only to σ_{TDOA} for quantize/decimate for the high SNR case; note that at R = 30 kbps (where from Figure 3 we see that quantize only can achieve the rate with a barely allowable b = 3 bits) that quantize/decimate approach gives about three times the accuracy as the quantize-only approach.

For the low SNR case, however, the rate-distortion curves in Figure 5 seem to indicate that quantize-only is nearly uniformly preferred over quantize/decimate and decimate-only. However, it is important to recognize that in this case the effective rates at which quantize-only is clearly better are precisely those rates at which the mathematics calls for excessive quantization to meet the rate constraint. Thus, from a practical viewpoint, quantize/decimate is preferred at these lower rates. Stated another way, quantize/decimate gives a viable means for meeting the lower rate constraints without suffering excessive nonlinearity effects from quantization. Thus, even in the low SNR case, the quantize/decimate approach is an effective way to meet the imposed rate constraint. These examples point out the following general principles:

- In a low SNR setting, quantize more and decimate less.
- In a high SNR setting, quantize less and decimate more.
- The quantize/decimate technique is well-suited when (i) SNR, DNR are high, or (ii) rate constraint imposes too small a number of bits for the quantization-only approach.
- The quantize/decimate approach provides a more flexible approach that results in a smooth rate-distortion curve whereas the quantize-only approach has a less flexible stepped rate-distortion curve.
- If the data rate is decreased, do not reduce the number of bits, but rather increase the decimation factor (i.e., reduce the bandwidth).
- If the data rate is increased, do not increase the number of bits, but rather decrease the decimation factor (i.e., increase the bandwidth) C until you are using the full bandwidth, then increase the number of bits used.

Of course, it should be remembered that the results presented above are for the case of flat signal spectra. For non-flat spectra the general conclusions still apply; however, it is likely that decimation would be favored even more than in the flat spectra case, because spectra that decay at high frequencies will lose RMS bandwidth at a slower rate as they are filtered to lower bandwidths. Also, the SNR is likely to change as the signals are filtered. The ideas presented above can be extended to handle these more general cases and provide algorithms that can properly trade-off quantization and decimation for the general case.

5. CONCLUSIONS

We have demonstrated that for TDOA-based emitter location, a combination of quantization and decimation is preferable to quantization alone when operating under a rate constraint. We have proposed a rate constraint in terms of an effective rate measure that we defined as $R = R_l T_l / T$. This constraint is well suited to this problem because it encapsulates the system parameters of link rate R_l, link time T_l, and signal collection time T into a single constraint. We then argued that unlike in past investigations, it is desirable to consider a non-MSE distortion criteria for the TDOA estimation problem and proposed the use of the Cramer-Rao bound for TDOA standard deviation. It was pointed out that this distortion criteria depends on both the compressed signal's SNR and its RMS bandwidth, from which we argued that a simple approach to exploiting this criterion is to simultaneously quantize and decimate the signal. Under the simplifying assumption of a flat signal spectrum we derived expressions that characterized the optimal trade-off between quantization and decimation. We showed that for the high SNR case, the quantize/decimate approach uniformly outperforms the quantize-only approach from an effective rate-distortion point of view. For the low SNR case we showed that there are effective rates below which the quantize-only approach can't work due to severe nonlinearity of the quantization; however, the quantize/decimate approach operates well at these low rates. Thus, the quantize/decimate approach provides an effective means for low-rate operation.

Obviously, lowpass filtering and decimation used here is the simplest way to exploit the RMS bandwidth's effect on TDOA accuracy. The results presented here point the way to more general filtering/decimation approaches for TDOA-only and the use of the wavelet transform to exploit the joint effect of RMS bandwidth and RMS duration for systems that supplement TDOA with frequency-difference-of-arrival (FDOA)[7].

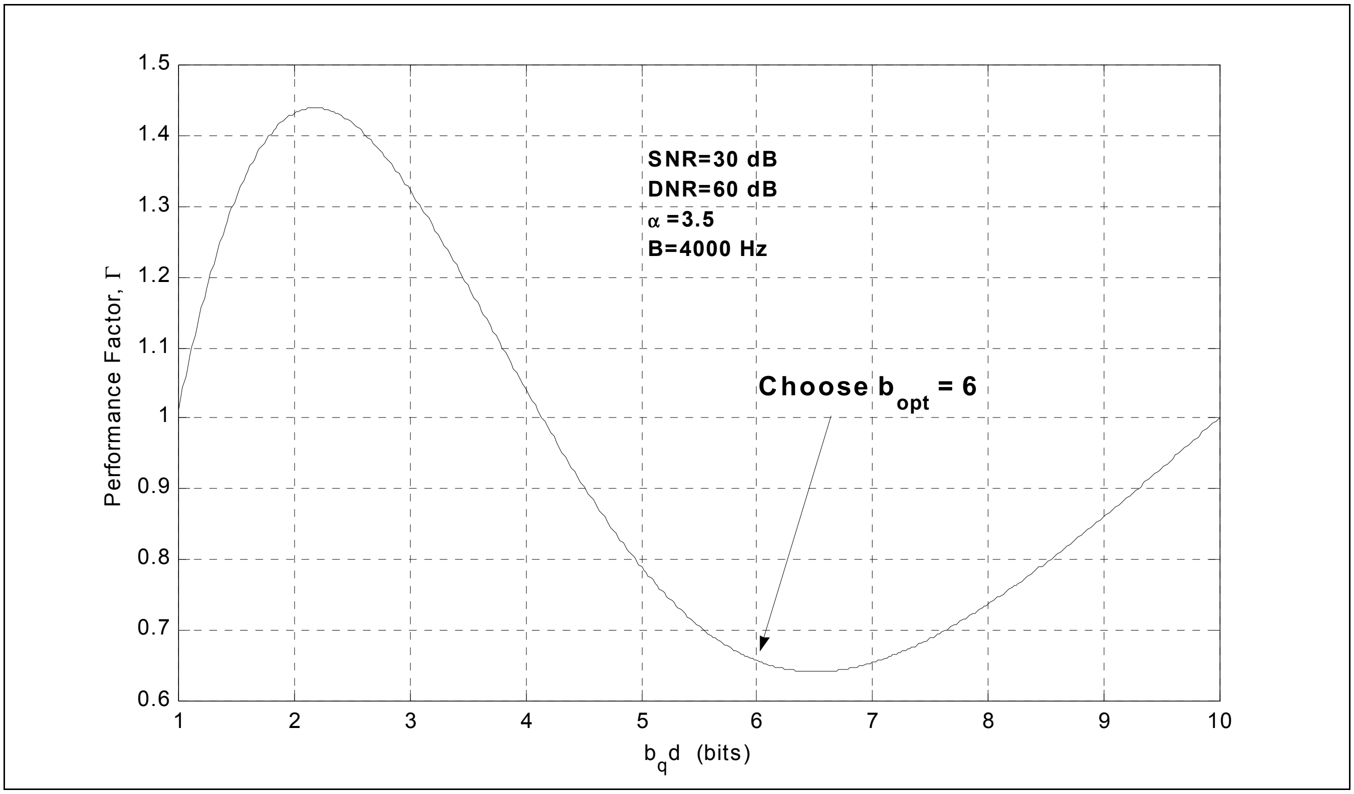

Figure 2: Performance Factor for High SNR Case

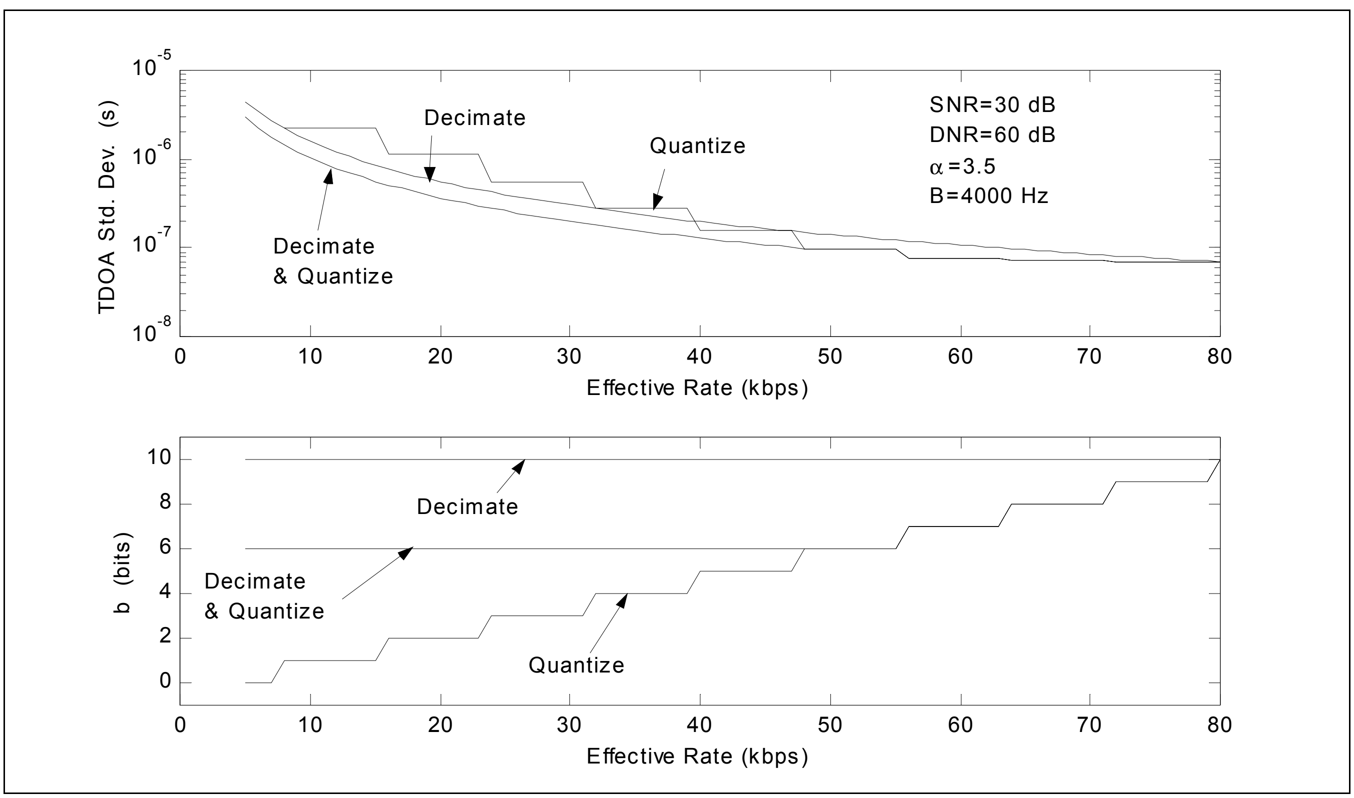

Figure 3: Rate-Distortion Curve for High SNR Case

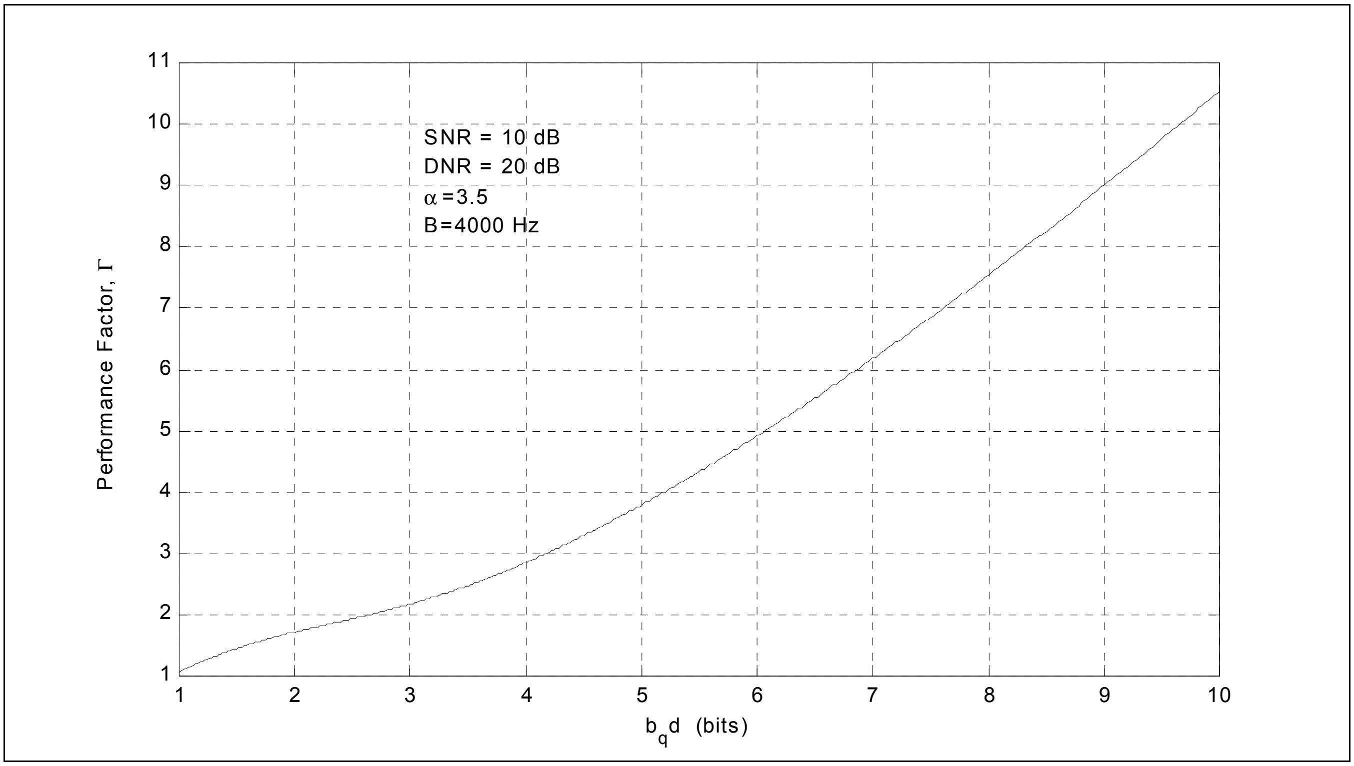

Figure 4: Performance Factor for Low SNR Case

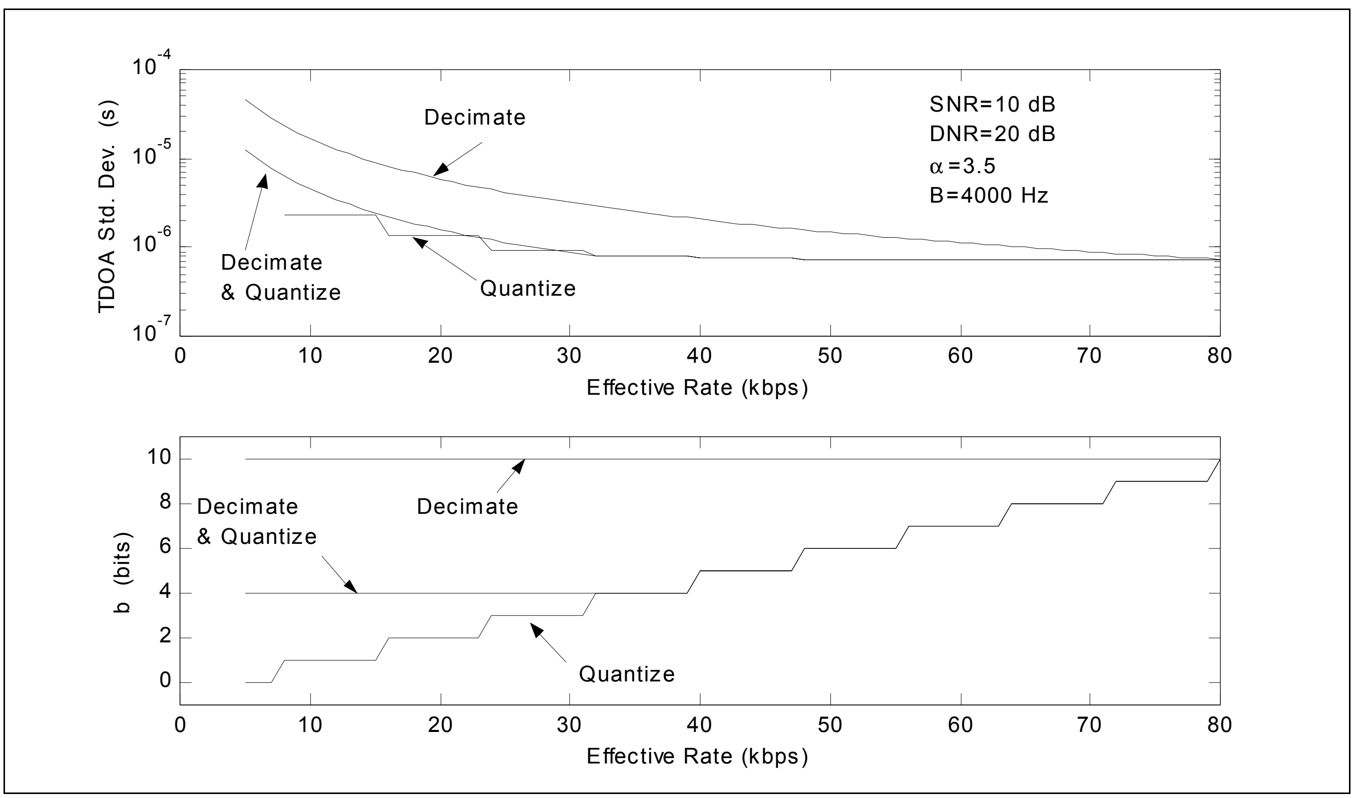

Figure 5: Rate-Distortion Curve for Low SNR Case

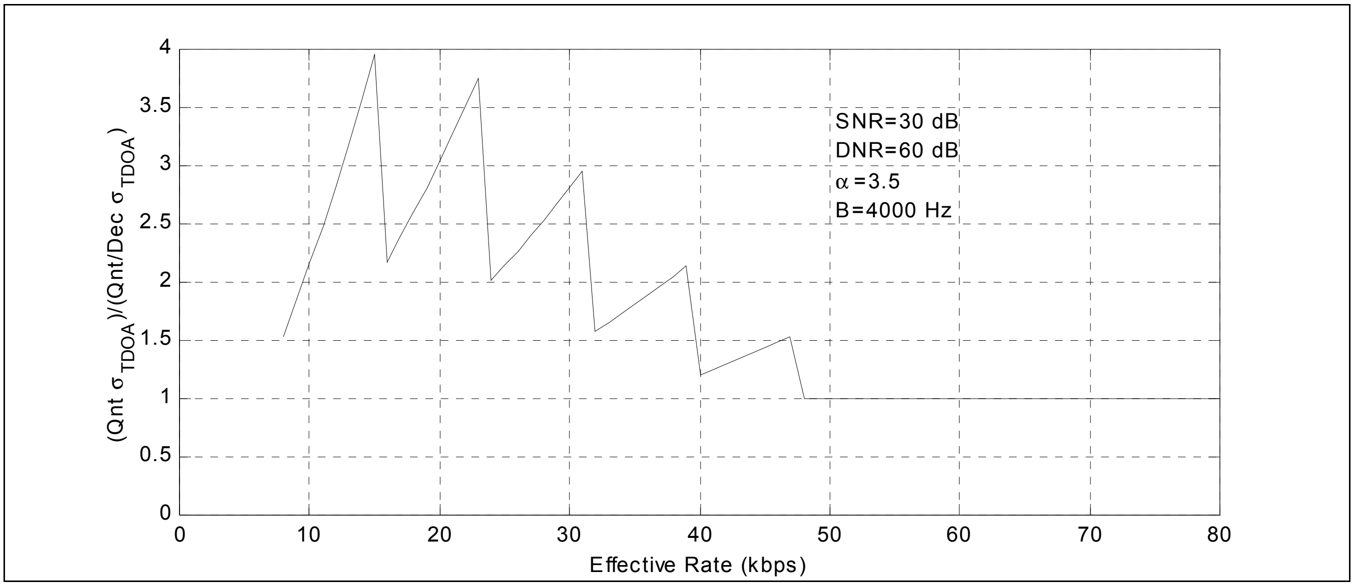

Figure 6: Ratio of Std. Dev. for Quantize Only and Quantize/Decimate

6. REFERENCES

1. P. C. Chestnut, "Emitter location accuracy using TDOA and differential doppler ," *IEEE Trans. Aero. and Electronic Systems*, vol. AES-18, pp. 214-218, March 1982.
2. S. Stein, "Differential delay/doppler ML estimation with unknown signals," *IEEE Trans. Sig. Proc.*, vol. 41, pp. 2717 - 2719, August 1993.
3. S. Stein, "Algorithms for ambiguity function processing," *IEEE Trans. Acoust., Speech, and Signal Processing*, vol. ASSP-29, pp. 588 - 599, June 1981.
4. D. J. Matthiesen and G. D. Miller, "Data transfer minimization for coherent passive location systems," Report No. ESD-TR-81-129, Air Force Project No. 4110, June 1981.
5. M. L. Fowler, "Coarse quantization for data compression in coherent location systems," to appear in *IEEE Transactions on Aerospace and Electronic Systems*.
6. M. L. Fowler, "Data compression for TDOA/DD-based location system," US Patent #5,991,454 issued Nov. 23, 1999, held by Lockheed Martin Federal Systems.
7. M. L. Fowler, "Exploiting RMS time-frequency structure for data compression in emitter location systems," accepted to National Aerospace and Electronics Conference (NAECON), Dayton, Ohio, 10-12 October, 2000.