Improved Spread Spectrum: A New Modulation Technique for Robust Watermarking

Henrique S. Malvar, *Fellow, IEEE,* and Dinei A. F. Florêncio, *Member, IEEE*

***Abstract*—This paper introduces a new watermarking modulation technique, which we call *improved spread spectrum* (ISS). When compared with traditional spread spectrum (SS), the signal does not act as a noise source, leading to significant gains. In some examples, performance improvements over SS are 20 dB in signal-to-noise ratio (SNR) or ten or more orders of magnitude in the error probability. The proposed method achieves roughly the same noise robustness gain as quantization index modulation (QIM) but without the amplitude scale sensitivity of QIM. Our proposed ISS is as robust in practice as traditional SS.**

***Index Terms*—Data hiding, information embedding, spread spectrum, watermarking.**

I. Introduction

WATERMARKING—or embedding information in a way not immediately discernible but hard to reproduce—has been used as a way of reducing counterfeiting. These techniques have been used in documents, currency, and other applications for centuries. With the widespread use of digital representation for images, video, sounds, and other signals, copyright protection by using a "digital watermark" became a very active area of research (see [1] for an extensive bibliography). Watermarking in this new context is a complex problem, with issues that involve not only watermarking techniques but involving systems design, cryptography, and a series of economic and legal aspects as well. While we do appreciate the complexity of the problem, in this paper, we only deal with a single aspect of the problem: that of "hiding" or transmitting information under a signal, by adding or embedding an imperceptible signal (i.e., the watermark).

In many watermarking schemes, spread spectrum (SS) is the modulation technique used to embed the watermark [2]–[4]. In the simplest scheme, the bits composing the desired message (e.g., the name of the copyright owner) are modulated by an SS sequence and added to the signal. Since SS is robust to interfering noise, the amount of energy (or distortion) that has to be added to the watermarked signal to "erase" the watermark can be made very high. In fact, the signal itself is a source of interference. In more elaborate schemes, differences in the signal may be explored in order to reduce subjective distortion introduced by the watermark. Finally, other aspects of a complete watermarking system deal with aspects such as secure key distribution, resynching (to overcome malicious attacks to the watermark), and computational complexity.

In schemes using SS as the embedding technique, the signal itself is seen as a source of interference [2], [5]. In practical watermarking applications, the signal is generally much stronger than any interference the signal must endure, and therefore, the interference from the signal itself dominates the process. In fact, data hiding by low-bit(s) modulation (LBM) can be seen as an SS, where the interference from the signal was removed, and the SS sequence length can be reduced up to a single bit if channel noise is not present. Nevertheless, due to the fragility of LBM to attacks, other methods of reducing or eliminating the interference from the signal are necessary.

In [6], Chen and Wornell propose a new embedding method called quantization index modulation (QIM). Their method does reduce or eliminate the interference from the signal, achieving a much higher robustness to additive noise than SS. However, QIM obtains its gains from embedding the watermark in a lattice, making the watermark very sensitive to scaling of the signal, i.e., a simple change in the scale of the watermarked signal will practically erase the watermark. Although scaling may led to large mean-square errors, it is usually perceptually acceptable. Therefore, QIM is not applicable to watermarking signals whenever a malicious attack using scaling can take place.

In [5], Cox *et al.* present a framework where they indicate the need for removing the influence of the signal in the watermark detection process, but they come short of presenting a practical solution to the problem. More recently, three different practical solutions based on that framework have been proposed [8]. These solutions correspond to the cases of "maximizing correlation coefficient," "maximizing robustness," and requiring "constant robustness." Still, they do not handle the important case of how to insert the watermark to minimize the error rate at a fixed energy level (or more precisely, at a given average distortion level). Furthermore, the solution presented there is based on the assumption that detection is based on correlation coefficient. It is more common that simple correlation is used as the detection criteria and in that case, all solutions presented in [8] degenerate into traditional (i.e., blind) SS modulation.

In this paper, we propose a simple technique, which we call improved spread spectrum (ISS). This technique, in practice, removes the signal as source of interference, producing a dramatic improvement in the quality of the watermarking process. The gains for the ISS are similar to those obtained by QIM, but the method proposed herein does not suffer from the same sensitivity to amplitude scaling. ISS is essentially as insensitive to amplitude scaling as traditional SS.

Manuscript received February 4, 2002; revised November 21, 2002. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Ali N. Akansu.

The authors are with the Communication, Collaboration, and Signal Processing Group, Microsoft Research, Redmond, WA 98052 USA (e-mail: malvar@microsoft.com; dinei@microsoft.com).

Digital Object Identifier 10.1109/TSP.2003.809385

1053-587X/03$17.00 © 2003 IEEE

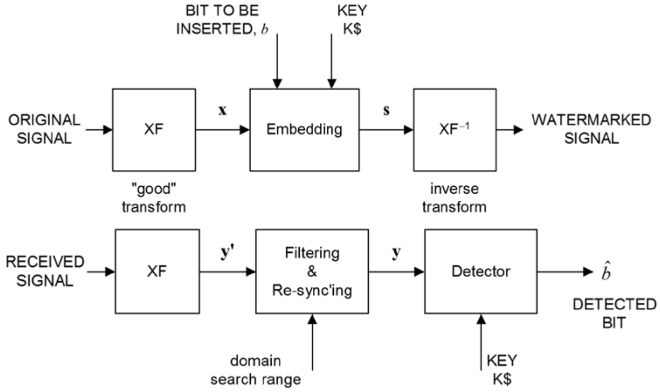

Fig. 1. Watermarking system. (a) Embedding. (b) Detection.

Practically, any watermarking system currently using SS would immediately profit from using the proposed scheme as a direct replacement for SS. Gains will vary according to signal-to-noise ratio (SNR) and operating error probability, but improvements of 20 dB in SNR resistance or reduction in error probabilities of 10 or more orders of magnitude are common.

In Section II, we present our framework and analyze traditional SS as it applies to watermarking. In Section III, we present our basic ISS technique. In Section IV, we introduce a simplified (linear) version of ISS and analyze and compare the performance in terms of noise immunity of our technique to that of traditional SS and QIM. In Section V, we analyze the variance introduced by our method in the signal distortion and introduce a few further enhancements, which control the variance, including a limited distortion version of the linear approximation introduced in Section IV. In Section VI, we derive the optimum ISS, and in Section VII, we present some conclusions.

II. Traditional Approach for SS-Based Watermarking

A general system for SS-based watermarking is shown in Fig. 1. Embedding is shown in Fig. 1(a) and detection in Fig. 1(b). In these figures, the box "good transform" is intended to represent a transform from the original signal domain to a domain where the data is more equally sensitive to tampering. Ideally, a "good transform" also removes any part of the data that is not perceptually significant. In the case of images, for example, the transform should be insensitive to translation, small contrast manipulations, lowpass filtering, and other common signal processing techniques. The idea is that after the transform, any *significant* change in the signal would significantly impair the image. Note that we include a box with the inverse transform, but the transform does not need to be strictly invertible since we can pass some side information from the original signal to the inverse transform.

Even though we recognize the difficulties involved in designing such a "good" transform, in this paper, we do not address this problem. Instead, we focus only on the next step, which is to actually insert the watermark after such transform. In our notation, the vector $\mathbf{x}$ is considered to be the original signal already in an appropriate transform domain to be marked. The vector $\mathbf{y}$ is the received vector, in the transform domain, after channel distortions.

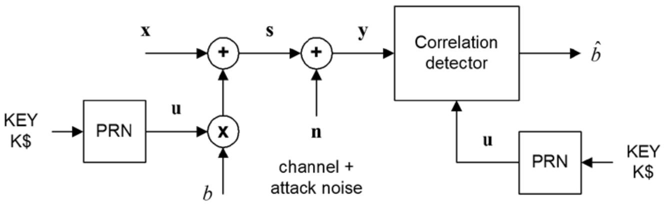

Fig. 2. Spread-spectrum-based watermarking. b is the bit to be embedded.

SS-based watermarking is shown in Fig. 2. A secret key $K\$$ is used by a pseudo random number generator (PRN) to produce a "chip sequence" $\mathbf{u}$ with zero mean and whose elements are equal to $+\sigma_u$ or $-\sigma_u$. The sequence $\mathbf{u}$ is then added to or subtracted from the signal $\mathbf{x}$ according to the variable b, where b assumes the values of $+1$ or -1 according to the bit (or bits) to be transmitted by the watermarking process. The signal $\mathbf{s}$ is the watermarked signal.

A simple analysis of SS-based watermarking leads to a simple formula for the probability of error. First, consider the definitions of inner product and norm:

$$\langle \mathbf{x}, \mathbf{u} \rangle \triangleq \frac{1}{N} \sum_{i=0}^{N-1} x_i u_i, \quad \text{and} \quad \|\mathbf{x}\| \triangleq \langle \mathbf{x}, \mathbf{x} \rangle \tag{1}$$

where N is the length of the vectors $\mathbf{x}$, $\mathbf{u}$, $\mathbf{s}$, $\mathbf{n}$, and $\mathbf{y}$ in Fig. 2.

Without loss of generality, we assume that we are embedding one bit of information in a vector $\mathbf{s}$ of N transform coefficients. Then, the bit rate is $1/N$ bits/sample. That bit is represented by the variable b, whose value is either -1 or $+1$. Embedding is performed by

$$\mathbf{s} = \mathbf{x} + b\mathbf{u}. \tag{2}$$

The distortion D in the embedded signal is defined by $\|\mathbf{s} - \mathbf{x}\|$. It is easy to see that for the embedding equation above, we have

$$D = \|b\mathbf{u}\| = \|\mathbf{u}\| = \sigma_u^2. \tag{3}$$

The channel is modeled as additive noise:

$$\mathbf{y} = \mathbf{s} + \mathbf{n}. \tag{4}$$

Detection is performed by first computing the (normalized) sufficient statistic r:

$$r \triangleq \frac{\langle \mathbf{y}, \mathbf{u} \rangle}{\langle \mathbf{y}, \mathbf{u} \rangle} = \frac{\langle b\mathbf{u} + \mathbf{x} + \mathbf{n}, \mathbf{u} \rangle}{\sigma_u^2} = b + x + n \tag{5}$$

and estimating the embedded bit by

$$\hat{b} = \text{sign}(r) \tag{6}$$

where $x \triangleq \langle \mathbf{x}, \mathbf{u} \rangle / \|\mathbf{u}\|$and $n \triangleq \langle \mathbf{n}, \mathbf{u} \rangle / \|\mathbf{u}\|$.

We assume simple statistical models for the original signal $\mathbf{x}$ and the attack noise $\mathbf{n}$. Namely, we assume both to be samples from uncorrelated white Gaussian random processes. Therefore

$$x_i \sim N\left(0, \sigma_x^2\right), \quad n_i \sim N\left(0, \sigma_n^2\right) \tag{7}$$

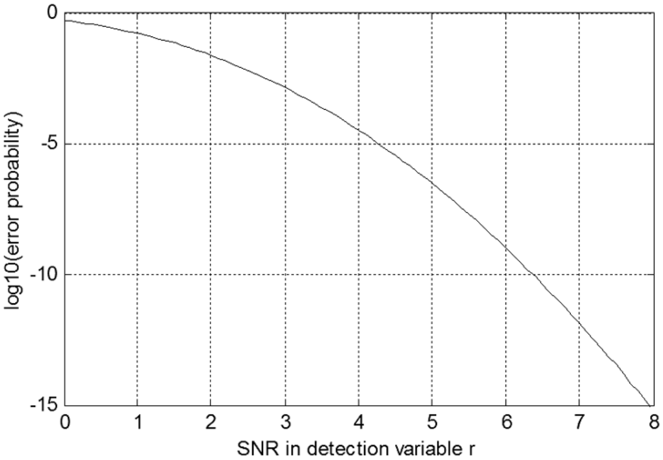

Fig. 3. Error probability for SS-based watermarking.

Then, it is easy to show that the sufficient statistic r is also Gaussian, i.e.,

$$r \sim N\left(m_r, \sigma_r^2\right), \quad m_r = E[r] = b\,\sigma_r^2 = \frac{\sigma_x^2+\sigma_n^2}{N\sigma_u^2}. \tag{8}$$

In particular, let us consider the case when $b = 1$. Then, an error occurs when $r < 0$, and therefore, the error probability p is given by

$$\begin{aligned} p =& \Pr\left\{\hat{b} < 0 | b = 1\right\} = \frac{1}{2}\mathrm{erfc}\left(\frac{m_r}{\sigma_r\sqrt{2}}\right) \\ =& \frac{1}{2}\mathrm{erfc}\left(\sqrt{\frac{\sigma_u^2 N}{2\left(\sigma_x^2+\sigma_n^2\right)}}\right) \end{aligned} \tag{9}$$

where erfc$(\cdot)$ is the complementary error function. The same error probability is obtained under the assumption that $b = -1$. A plot of that probability as a function of the SNR m_r/σ_r is shown in Fig. 3.

For example, from Fig. 3, we see that if we want an error probability better than 10^{-3}, then we need

$$\frac{m_r}{\sigma_r} > 3 \Rightarrow N\sigma_u^2 > 9\left(\sigma_x^2+\sigma_n^2\right) \tag{10}$$

or more generally, to achieve an error probability p, we need

$$N\sigma_u^2 > 2\left(\mathrm{erfc}^{-1}(p)\right)^2\left(\sigma_x^2+\sigma_n^2\right). \tag{11}$$

The equation above shows that we can trade the length of the chip sequence N with the energy of the sequence σ_u^2. It allows us to easily compute either N or σ_u^2, given the other variables involved.

III. New Approach via ISS

The main idea behind the ISS is that by using the encoder knowledge about the signal $\mathbf{x}$ (or more precisely, x, the projection of $\mathbf{x}$ on the watermark), we can enhance performance by modulating the energy of the inserted watermark to compensate for the signal interference. The new embedding approach is defined by a slight modification to the SS embedding (2), i.e., we vary the amplitude of the inserted chip sequence by a function $\mu(x,b)$:

$$\mathbf{s} = \mathbf{x} + \mu(x,b)\mathbf{u} \tag{12}$$

where, as before, $x \triangleq \langle \mathbf{x},\mathbf{u}\rangle/||\mathbf{u}||$. Note that the traditional SS is a particular case of ISS. In our notation, SS is a case of the ISS in which the function μ is made independent of x.

We now analyze a few small variations of the ISS approach. In particular, Section IV presents more detailed analysis of a linear approximation to μ, as this allows for a simpler mathematical analysis. Nevertheless, this linear approximation has a clear disadvantage in that the maximum distortion introduced by the watermark is not limited. In Section V, we propose a few ways in which the maximum distortion can be analyzed and controlled. Finally, in Section VI, we go back to the general problem of finding the optimum $\mu(x,b)$.

IV. Linear Approximation

A simpler version of the ISS is to restrict μ to be a linear function. Not only is this much simpler to analyze, it also provides a significant part of the gains in relation to traditional SS. In this case and due to the symmetry of the problem in relation to b and x, we have

$$\mathbf{s} = \mathbf{x} + (\alpha b - \lambda x)\mathbf{u}. \tag{13}$$

The parameters α and λ control the distortion level and the removal of the carrier distortion on the detection statistic. Traditional SS is obtained by setting $\alpha = 1$ and $\lambda = 0$.

With the same channel noise model as before, the receiver sufficient statistic is

$$r = \frac{\langle \mathbf{y},\mathbf{u}\rangle}{||\mathbf{u}||} = \alpha b + (1-\lambda)x + n. \tag{14}$$

Therefore, the closer we make λ to 1, the more the influence of x is removed from r. The detector is the same as in SS, i.e., the detected bit is sign(r).

The expected distortion of the new system is given by

$$\begin{aligned} E[D] =& E\left[||\mathbf{s}-\mathbf{x}||\right] \\ =& E\left[|\alpha b - \lambda x|^2\sigma_u^2\right] = \left(\alpha^2 + \frac{\lambda^2\sigma_x^2}{N\sigma_u^2}\right)\sigma_u^2. \end{aligned} \tag{15}$$

To make the average distortion of the new system to equal that of traditional SS, we force $E[D] = \sigma_u^2$, and therefore

$$\alpha = \sqrt{\frac{N\sigma_u^2 - \lambda^2\sigma_x^2}{N\sigma_u^2}}. \tag{16}$$

To compute the error probability, all we need is the mean and variance of the sufficient statistic r. They are given by

$$\begin{aligned} m_r =& \alpha b, \text{ and} \\ \sigma_r^2 =& \frac{\sigma_n^2 + (1-\lambda)^2\sigma_x^2}{N\sigma_u^2}. \end{aligned} \tag{17}$$

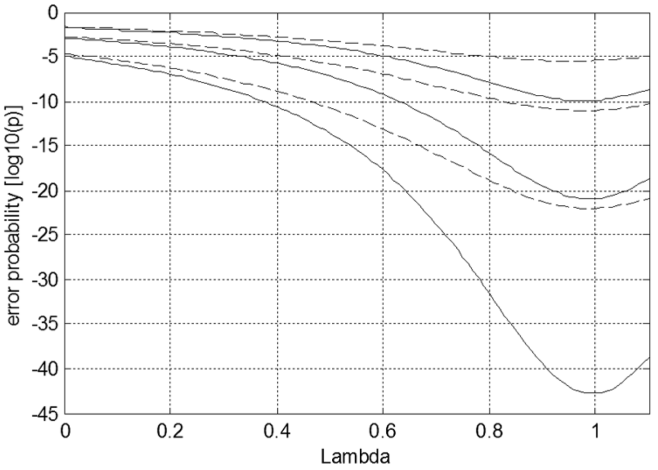

Fig. 4. Error probability as a function of λ. Solid lines represent a 10-dB SNR, and dashed lines represent a 7-dB SNR. The three lines correspond to values of equal to 5, 10, and 20 (with higher values having smaller error probability).

We can therefore compute the error probability p by

$$\begin{aligned} p =& \Pr\{r < 0|b = 1\} \\ =& \frac{1}{2}\text{erfc}\left(\frac{m_r}{\sigma_r\sqrt{2}}\right) \\ =& \frac{1}{2}\text{ erfc}\left(\sqrt{\frac{N\sigma_u^2 - \lambda^2\sigma_x^2}{2\left(\sigma_n^2 + (1-\lambda)^2\sigma_x^2\right)}}\right). \end{aligned} \tag{18}$$

We can also rewrite p as a function of the relative power of the SS sequence $N\sigma_u^2/\sigma_x^2$ and the SNR σ_x^2/σ_n^2

$$p = \frac{1}{2}\text{erfc}\left(\frac{1}{\sqrt{2}}\sqrt{\frac{\frac{N\sigma_u^2}{\sigma_x^2} - \lambda^2}{\frac{\sigma_n^2}{\sigma_x^2} + (1-\lambda)^2}}\right). \tag{19}$$

In Fig. 4, we plot p as a function of λ for various values of SNR and $N\sigma_u^2/\sigma_x^2$. Remember that $\lambda = 0$ corresponds to SS. Note that by proper selection of the parameter λ, the error probability in the proposed method can be made several orders of magnitude better than using traditional SS. For example, with a signal-to-interference ratio of 10 (i.e., 10 dB), we get a reduction in the error rate from $p_0 = 10^{-5}$ for traditional SS to $p = 1.55 \times 10^{-43}$ for the proposed method, which is a reduction of over 37 orders of magnitude in the error probability. Higher SNR values, which can happen in practical applications, lead to even higher gains.

As it can be inferred from Fig. 4, the error probability varies with λ, with the optimum value usually close to one. The expression for the optimum value for λ can be computed from the error probability p by setting $\partial p/\partial\lambda = 0$ and is given by

$$\lambda_{opt} = \frac{1}{2}\left(\left(1 + \frac{\sigma_n^2}{\sigma_x^2} + \frac{N\sigma_u^2}{\sigma_x^2}\right) - \sqrt{\left(1 + \frac{\sigma_n^2}{\sigma_x^2} + \frac{N\sigma_u^2}{\sigma_x^2}\right)^2 - 4\frac{N\sigma_u^2}{\sigma_x^2}}\right). \tag{20}$$

In addition, note from this expression that for N large enough, $\lambda_{opt} \to 1$ as SNR $\to \infty$.

A. *Improvement in Noise Immunity Over Traditional SS*

Until now, we considered the improvement in the error probability when using the ISS. We now try to answer the question of how much more noise (for the same error probability) can ISS stand compared with SS. For simplicity, we restrict our analysis to the linear ISS. Let us call σ_n^2 the noise level in our new improved SS system and σ_{no}^2 the noise system in the original SS system. Our goal now is to compute how much larger can ${\sigma_n}^2$ be compared with σ_{no}^2 for the same average distortion and same probability of error. Since we have the same error probabilities for the same SNR m_r/σ_r

$$\frac{N\sigma_u^2 - \lambda^2\sigma_x^2}{\sigma_n^2 + (1-\lambda)^2\sigma_x^2} = \frac{N\sigma_u^2}{\sigma_x^2 + \sigma_{no}^2} \tag{21}$$

or

$$\sigma_n^2 = \left(\sigma_x^2 + \sigma_{no}^2\right)\left(1 - \frac{\lambda^2\sigma_x^2}{N\sigma_u^2}\right) - (1-\lambda)^2\sigma_x^2. \tag{22}$$

For traditional SS, we have $\lambda = 0$, and thus, $\sigma_n^2 = \sigma_{no}^2$, as expected.

The optimal design for our improved SS system is obtained when we chose the parameter λ such that σ_n^2 is maximized, i.e., the improved system can tolerate the maximum amount of noise. Let us call λ_{opt} the optimal value of λ. Then, it is easy to show that

$$\lambda_{opt} = \frac{N\sigma_u^2}{N\sigma_u^2 + \sigma_x^2 + \sigma_{no}^2}. \tag{23}$$

For the optimal design, then, the allowable noise level can be written as

$$\sigma_n^2 = \sigma_{no}^2 + \lambda_{opt}\sigma_x^2. \tag{24}$$

If $\lambda_{opt} \approx 1$ and the channel SNR is high, then we have $\sigma_n^2 \gg \sigma_{no}^2$.

The proposed method improves the error ratio (and/or noise immunity) for any level of channel (attack) noise and for any level of desired error probability. We now select a more specific example to give an idea of the level of improvements that can be achieved with the proposed method. We recall that for the traditional SS system to work with a low probability of error, e.g., $p < 10^{-3}$, then we need $N\sigma_u^2 > 9(\sigma_n^2 + \sigma_x^2)$. That means $0.9 < \lambda_{opt} < 1$. In additon, let us call Q the signal-to-channel-noise ratio, i.e.,

$$Q \triangleq \frac{\sigma_x^2}{\sigma_{no}^2} \tag{25}$$

and let us call P the noise tolerance gain of our system when compared to traditional SS:

$$P \triangleq \frac{\sigma_n^2}{\sigma_{no}^2}. \tag{26}$$

Then, it is easy to show that the noise tolerance gain for our new system is given by

$$P = (1 + \lambda_o Q) > (1 + 0.9Q). \tag{27}$$

Therefore, for large Q, the improvement in noise tolerance is quite significant since it is approximately equal to $0.9Q$. In Fig. 5, we plot P as a function of Q, all in decibels.

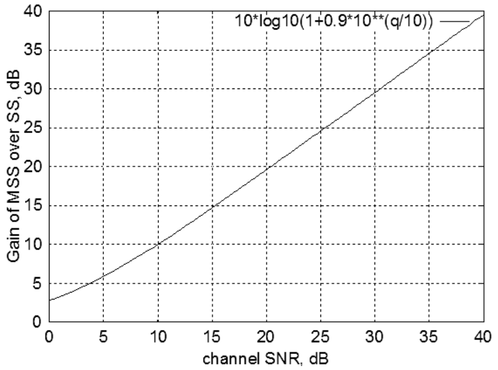

Fig. 5. Improvement in noise robustness of our new ISS over traditional SS, in decibels, as a function of the channel SNR. For a typical scenario of a channel SNR of 20 dB, the improvement is about 20 dB, i.e., our ISS system can tolerate 100 more times channel noise power than traditional SS.

We note that the performance of our new ISS system is quite close to that of QIM [6]. The noise tolerance improvement of QIM over SS is slightly above Q, whereas in our system, it is slightly below Q. However, our ISS system is not sensitive to amplitude scaling of the received signal $\mathbf{y}$, like QIM. Therefore, ISS can be more robust in practical applications.

B. Comparison of Required Watermark Energy

In many applications, a desired error probability and a certain SNR are specified. In such cases, the objective is to minimize the energy of the watermark, i.e., the signal distortion. We now use this situation to compare the linear ISS to traditional SS to STDM and to a theoretical bound.

For a given signal and noise energy and a desired error probability, (11) gives us the necessary energy in the watermark for traditional SS. A similar equation for the linear ISS can be obtained by inverting (18). The objective of ISS is to reduce the influence of the signal as a source of interference. A natural performance bound is therefore the result that could be achieved if the decoder had knowledge of the signal (and therefore could remove any influence from the detection statistic). This is equivalent to the problem of communication in presence of noise, where it is shown in [9] that if the encoder knows the channel noise (which, in our case, is the signal to be marked), then the achievable capacity is the same as if the decoder knew the channel noise. In other words, the encoder can precompensate for the channel noise. That is exactly what we are attempting to achieve with ISS.

In [6], Chen and Wornel show that the performance of STDM is only 1.25 dB above this bound. Fig. 6 shows a plot of these numbers for attacks corresponding to 5 , 10, and 20 dB SNR. In each figure, the solid line represents the theoretical bound, the dash-doted line represents the performance of traditional SS, the dashed line represents the performance of STDM, and the two dotted lines represent two versions of linear ISS: the simplest one ($\lambda = 1$) and with λ optimized according to (20). Note that in each case, traditional SS requires around the same extra energy in the watermark as the attack SNR. For error probabilities below 10^{-5} with attacks over 10 dB, the ISS performance is within 2 dB of the theoretical bound, and it even outperforms STDM for error probabilities below 10^{-8} (below 10^{-3} for a 5-dB attack). We note that other QIM methods that are more elaborate than STDM would help reduce the gap to the theoretical bound. Nevertheless, all QIM methods suffer from the scale sensitivity problem. In summary, ISS is much simpler, robust to scaling attacks, and does not require modifying the decoder, and its performance is similar to that of QIM.

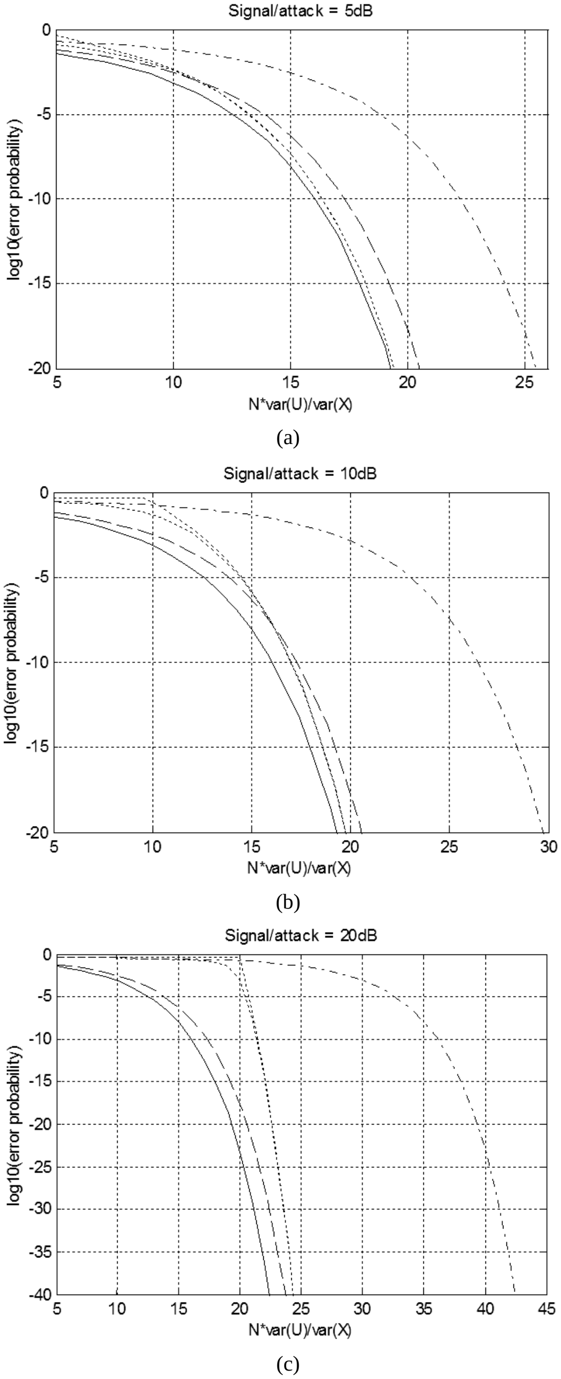

(a)

(b)

(c)

Fig. 6. Error probability as a function of watermark energy. Error probability for ISS (dotted lines) compared with SS (dash-dot lines), STDM (dashed line), and a theoretical bound (solid line). The SNR is 5, 10, and 20 dB.

V. Analysis and Control of Maximum Distortion

The choice of λ and α we have used in our computations guarantees the same level of average distortion D as in SS. Nevertheless, with the proposed method, the distortion is not constant, as in SS. Indeed, in the linear approximation described until now, D is not even limited, and that may be a problem in some applications. In particular, up to now, we have considered λ to be a constant. While this has allowed a simplified mathematical analysis, it has introduced the undesired effect that the distortion is not limited. In this section, we propose a slightly different version of the ISS, which limits the distortion to a desired maximum, but still gets most of the gains of the ISS. We first show in Section V-A that by introducing a limit on the distortion, we can actually obtain further improvements in the results we have shown. A disadvantage of this limit is that it is not under our control. In Section V-B, we address the case of applications that need a more strict control on the maximum distortion.

A. Maximum Useful Distortion for ISS

Looking at the extremes of the distortion in the ISS, we realize that large values of distortion occur in two different situations: In one of them, the sign of x is the same of b. For this case, it is easy to note that if $\lambda x/\alpha b > 1$, then the effect of λ is, in fact, to reduce the strength of the watermark. In other words, by setting $(\alpha b - \lambda x) = 0$ whenever $\lambda x/\alpha b > 1$, we can make the distortion zero and obtain even better error rates since the watermark will in fact be stronger.

The above procedure takes cares of limiting the distortion to one side, i.e., the side where x and b have the same sign. When x and b have opposite signs, the presence of the signal is, in fact, reducing the energy of the watermark. In this case, λ is helping restore the watermark to an ideal level. In fact, if $\lambda = 1$, the watermark would always be present at exactly the same energy level, regardless of the value of x. Nevertheless, as we have shown before, the choice of $\lambda = 1$ is not the optimum choice. For any value $\lambda < 1$, there is a value of x above (below) which the watermark is not correctly detected, even in the absence of noise. More precisely, if

$$\frac{x}{b} < \frac{-\alpha}{(1-\lambda)} \tag{28}$$

then we would have erroneous detection even in the absence of noise. Yet, these are large values of x, which therefore imply in large values of distortion. Since this watermark is not going to be detected anyway, again setting $(\alpha b - \lambda x) = 0$ (i.e., setting distortion to zero) is the most reasonable choice.

We have therefore shown that even though we assumed an unbounded distortion model to simplify our mathematical analysis, a more natural choice of bounds for λ in the ISS leads to a limited distortion algorithm. We can summarize the choice of λ for this algorithm as

$$\lambda = \begin{cases} \lambda_0 & \text{if, } -\frac{\alpha}{(1-\lambda_0)} < bx < \frac{\alpha}{\lambda_0} \\ \frac{\alpha b}{x}, & \text{otherwise} \end{cases} \tag{29}$$

where λ_0 is the precomputed value for λ as optimized by the previously described methods. Again, note that this second choice for λ for when x is outside the given interval implies zero distortion.

B. ISS With Limited Distortion

In the previous section, we showed that limiting the distortion does not necessarily affect the detection error rate. This modification can be seen as "giving up" whenever x is too large in the opposite direction of b and allowing a stronger watermark whenever x is too strong but in the same direction as b. In these cases, we do not transmit the watermark as the distortion necessary to allow reception would be too high. The only problem with this way of limiting the distortion is that we have no control over what that limit is: The limit depends only on the values established for α and λ. However, using the same general principles, we could limit distortion to whatever level we want. One way of expressing this would be to introduce a window function w, which limits the region where we introduce the watermark. We can express this by

$$\mathbf{s} = \mathbf{x} + w(x)(\alpha b - \lambda x)\mathbf{u} \tag{30}$$

where

$$w(x) = \begin{cases} 0, & \text{if} \quad bx > \frac{\alpha}{\lambda} \\ 0, & \text{if} \quad bx < -K \\ 1, & \text{otherwise} \end{cases} \tag{31}$$

where K is the "give up" parameter. Note that the first 0 is for the case where the watermark would be correctly detected even without inserting any additional signal (and therefore there is no need to increase distortion). The second 0 is for the "give up" case: If x is too strong and in the wrong direction, we just give up (and allow an error to occur). This allows us to guarantee a maximum distortion.

We can compute the new expected distortion as

$$\begin{aligned} E[D] =& E\left[\|\mathbf{s} - \mathbf{x}\|\right] \\ =& \sigma_u^2 E\left[w^2(x)|\alpha b - \lambda x|^2\right] \\ =& \sigma_u^2 \int_{-K}^{\alpha} \left(\alpha^2 + \lambda^2 x^2\right) p(x) dx. \end{aligned} \tag{32}$$

Therefore, to have the same D as traditional SS, we need to make the last integral equal to one. A simpler approach is to obtain an upper bound by extending the integrals to infinity. This is then the same as we had for the original ISS, i.e., we can guarantee $E[D] < \sigma_u^2$, by making

$$\alpha = \sqrt{\frac{N\sigma_u^2 - \lambda^2 \sigma_x^2}{N\sigma_u^2}}\,. \tag{33}$$

In this case, the error probability p is

$$\begin{aligned} p =& \Pr\left\{r < 0 | b = 1\right\} \\ =& \Pr\left\{(x + \mu(x)\left(\alpha - \lambda x\right) + n < 0)\right\} \\ \leq& \Pr\left\{x < -K\right\} \\ &+ \Pr\left\{-n > \alpha + (1-\lambda)\, x \,|x \geq -K\right\} \Pr\left\{x \geq -K\right\}. \end{aligned} \tag{34}$$

The analysis can be somewhat simplified if we make $\lambda = 1$, and in this case, we have

$$p \leq \Pr\left\{x < -K\right\} + \Pr\left\{n > \alpha\right\} \Pr\left\{x \geq -K\right\} \tag{35}$$

while for SS, we would have

$$p \cong \Pr\left\{x < -K\right\} + \Pr\left\{n > 1 + x\right\} \Pr\left\{x \geq -K\right\}. \tag{36}$$

In other words, we provide the same gains as before by removing the influence of x from the detection, except that now, we only do that for a predefined range of x.

VI. Optimum ISS

We now analyze the more generic case, where the function $\mu(x,b)$ is not restricted to be linear. We then express $\mathbf{s}$ as

$$\mathbf{s} = \mathbf{x} + \mu(x,b)\mathbf{u}. \tag{37}$$

We can then find the optimum solution for $\mu(x,b)$. We first note that since $\mathbf{x}$, $\mathbf{n}$ and b are independent, $\mu(x,b)$ is odd symmetric in the sense that $\mu(x,b) = -\mu(-x,-b)$. For simplicity and without loss of generality, we assume from now on that $b = 1$ and write simply $\mu(x)$, i.e.,

$$\mathbf{s} = \mathbf{x} + \mu(x)\mathbf{u}\ . \tag{38}$$

The distortion for a certain value of x is

$$d(x) = (\mu(x))^2\,\sigma_u^2 \tag{39}$$

and, as before, our sufficient statistic r, which is computed from $\mathbf{y} = \mathbf{s} + \mathbf{n}$, is

$$r(x) = x + \mu(x) + n. \tag{40}$$

We want to find the function $\mu(x)$ that minimizes the expected detection error probability $p = E\{pe(x)\}$ for a given expected distortion $D = E\{d(x)\}$. We can compute $pe(x)$ as

$$\begin{aligned} pe(x) &= \Pr\{r<0\} \\ &= \Pr\{(x+\mu(x)+n)<0\} = \frac{1}{2}\text{erfc}\left(\frac{x+\mu(x)}{\sigma_n\sqrt{2}}\right). \end{aligned} \tag{41}$$

To be optimum, $\mu(x)$ must be such that it satisfies

$$\frac{\partial}{\partial d}pe(x) = K' \tag{42}$$

for some constant K'. Therefore, an optimum solution $\mu(x)$ has to satisfy

$$\frac{\partial}{\partial \mu}d(x) = 0, \text{ or } \frac{\frac{\partial}{\partial \mu}pe(x)}{\frac{\partial}{\partial \mu}d(x)} = K'. \tag{43}$$

Since $d(x)$ is simply the (scaled) square of $\mu(x)$, the first condition is satisfied only for $\mu(x) = 0$. The second condition can be rewritten as

$$\frac{\frac{\partial}{\partial \mu}\text{erfc}\left(\frac{x+\mu(x)}{\sigma_n\sqrt{2}}\right)}{2\mu(x)} = K' \tag{44}$$

$$\Rightarrow e^{-\left((x+\mu(x))/(\sigma_n\sqrt{2})\right)^2} = K\mu(x) \tag{45}$$

where K is another constant. Unfortunately, there is no closed-form solution for $\mu(x)$, but we can solve the above equation numerically.

The expected error probability depends on the variance of the noise and on the constant K. We can, therefore look at K as a parameter that determines the final balance between distortion and error probability. Depending on the values of K, σ_n, and x, the equation has one, two, or three solutions. We illustrate this in Fig. 7, where we have plotted the right side of the equation (a straight line) for a certain K and the left side (a normal curve) for four different values of x (i.e., 0, -3, -6, and -9). For positive values of x, the peak of curve is even more to the left. The solution is one of the points where the normal curve intersects the line. It is clear that for positive values of x, we must choose the only intersection point. As the value of x becomes more and more negative (i.e., as the normal curve goes more and more to the right), there will be two extra intersection points, both to the left of the peak. Reasoning about the increase in distortion and the increase in error probability, we can discard the middle intersection point. To decide between the first and third point, we should compute the ratio between the difference in error probability and the difference in distortion at each point. Graphically, this can be interpreted in Fig. 7 as balancing the areas between the straight and the normal curves. In other words, we select the point closest to the origin whenever the area between the curves from this point to the second intersection point (i.e., the area in which the normal is below the straight line) is higher than the area between the curves from the second to the third intersection points (i.e., the area in which the normal is above the straight line). Otherwise, we select the third intersection point.

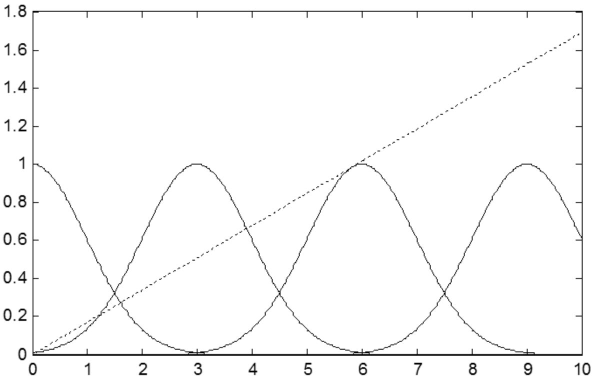

Fig. 7. Solving the equation for the optimum μ. Each normal curve corresponds to a different value of x. Intersection points are "candidates" to optimum solutions of $\mu(x)$.

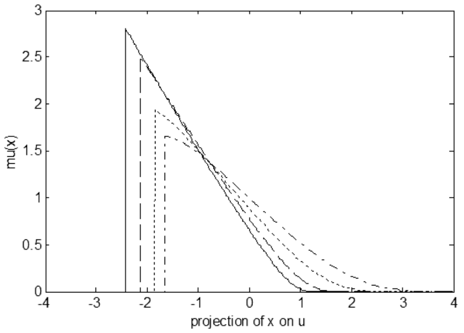

Fig. 8. Optimal $\mu(x)$ for several SNRs. From left to right, SNR is 10, 7, 3, and 0 dB. In all cases, distortion is 20 dB below the signal, and $N = 100$.

Fig. 8 shows a plot of the optimum $\mu(x)$ for some different situations. In particular, we vary the SNR ratio while keeping the average distortion constant. As it can be noted in the figure, the approximation of $\mu(x)$ by a straight line segment is a reasonable approximation for a large number of situations. In particular, the higher the SNR (or the stronger the watermark), the more appropriate the approximation. Again, remember that all previous analyses refer to the case where $b = 1$. The case where $b = -1$ behaves in the same way because μ is odd symmetric, as we mentioned before.

Fig. 9 compares the performance (in terms of error probability) for several of the ISS variations we have discussed so far. The particular data in this plot is for a SNR of 20 dB. The continuous line represents the error rate for traditional SS. The dash-dot line refers to the linear ISS and the dashed line to the

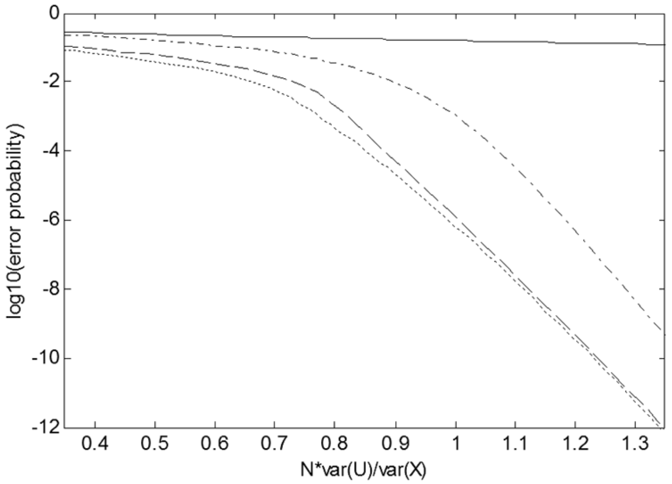

Fig. 9. Error probability for the SS and several versions of ISS. The SNR is 20 dB. The continuous line refers to traditional SS. The dash-dot line refers to the linear ISS and the dashed line to the limited distortion version of the linear ISS. The dotted line represents the fully optimized ISS.

limited distortion version of the linear ISS. The dotted line represents the fully optimized ISS. We can note that for most of the values of $N\sigma_u^2/\sigma_u^2$, the limited distortion linear approximation performs practically as well as the optimum choice of $\mu(x)$. Yet, if operating in the knee region, where the curve differs most, the difference of 0.6 in $\log(pe)$ means that the error probability is four times higher. Note also how adding the limitation on the distortion (or using the optimum solution) reduces the error probability by almost two orders of magnitude.

VII. Conclusion

In this paper, we have proposed a new spread spectrum (SS) technique, which we refer to as improved spread spectrum (ISS), for use in watermarking applications. We have shown that the ISS provides an exceptional improvement over traditional SS, with improvements in the error probability of several orders of magnitude for most typical scenarios. SS is currently used by many watermarking schemes as the information embedding (or modulation) technology. The proposed ISS technique can be readily applied to practically any watermarking technique currently using SS, taking immediate advantage of the gains. Furthermore, ISS does not require any change in the detection scheme, and in some cases, it could be applied even to systems that are already deployed, as far as we still have access to the encoders (this is often the case in media distribution schemes).

References

[1] R. J. Anderson and F. A. P. Petitcolas. (1999) Information hiding: An annotated bibliography. [Online]. Available: http://www.cl.cam.ac.uk/fapp2/steganography/bibliography

[2] A. Z. Tirkel, C. F. Osborne, and R. G. van Schyndel, "Image watermarking—A spread spectrum application," in *Proc. IEEE 4th Int. Symp. Spread Spectrum Techn. Applicat.*, Mainz, Germany, 1996, pp. 785–789.

[3] F. Hartung, J. K. Su, and B. Girod, "Spread spectrum watermarking: Malicious attacks and counterattacks," *Proc. SPIE*, vol. 3657, pp. 147–158, Jan. 1999.

[4] D. Kirovski and H. Malvar, "Robust spread-spectrum audio watermarking," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Salt Lake City, UT, May 2001.

[5] I. Cox, M. Miller, and A. McKellips, "Watermarking as communications with side information," *Proc. IEEE*, vol. 87, pp. 1127–1141, July 1999.

[6] B. Chen and G. Wornell, "Achievable performance of digital watermarking systems," in *Proc. Int. Conf. Multimedia Comput. Syst.*, Florence, Italy, June 1999, pp. 13–18.

[7] ——, "Quantization index modulation: A class of provably good methods for digital watermarking and information embedding," *IEEE Trans. Inform. Theory*, vol. 47, pp. 1423–1443, May 2001.

[8] M. Miller, I. Cox, and J. Bloom, "Informed embedding: Exploiting image and detector information during watermark insertion," in *Proc. IEEE Int. Conf. Image Process.*, 2000.

[9] M. H. M. Costa, "Writing on dirty paper," *IEEE Trans. Inform. Theory*, vol. IT-29, pp. 439–441, May 1983.

Henrique S. Malvar (M'79–SM'91–F'97) received the B.S. degree in 1977 from Universidade de Brasília, Brasília, Brazil, the M.S. degree in 1979 from Universidade Federal do Rio de Janeiro, Rio de Janeiro, Brazil, and the Ph.D. degree in 1986 from the Massachusetts Institute of Technology (MIT), Cambridge, all in electrical engineering.

From 1979 to 1993, he was with the faculty of the Universidade de Brasília. From 1986 to 1987, he was a Visiting Assistant Professor of electrical engineering at MIT and a senior researcher at PictureTel Corporation, Andover, MA. In 1993, he rejoined PictureTel, where he stayed until 1997 as Vice President of Research and Advanced Development. Since 1997, he has been a Senior Researcher at Microsoft Research, Redmond, WA, where he heads the Communication, Collaboration, and Signal Processing Research Group. His research interests include multimedia signal compression and enhancement, fast algorithms, multirate filterbanks, and wavelet transforms. He has several publications in these areas, including the book *Signal Processing with Lapped Transforms* (Boston, MA: Artech House, 1992). He is an Associate Editor for the journal *Applied and Computational Harmonic Analysis*.

Dr. Malvar is an Associate Editor of the IEEE Transactions on Signal Processing and a member of the Signal Processing Theory and Methods Technical Committee of the IEEE Signal Processing Society. He received the Young Scientist Award from the Marconi International Fellowship and Herman Goldman Foundation in 1981. He also received the Senior Paper Award in Image Processing in 1992 and the Technical Achievement Award in 2002, both from the IEEE Signal Processing Society.

Dinei A. F. Florêncio (S'88–M'96) received the B.S. and M.S. degrees from Universidade de Brasilia, Brasilia, Brazil, in 1983 and 1991, respectively, and the Ph.D. degree from the Georgia Institute of Technology, Atlanta, in 1996, all in electrical engineering.

Since 1999, he has been a Researcher at Microsoft Research, Redmond, WA. From 1996 to 1999, he was a Member of the Research Staff at the David Sarnoff Research Center, Princeton, NJ. He was also a summer intern with Interval Research, Palo Alto, CA, in 1994 and an Associated Researcher with NCR's Human Interface Lab, Atlanta, GA, from 1994 to 1996. His current research interests include signal compression and enhancement and real-time communications.

Dr. Florêncio received received the 1998 Sarnoff Achievement Award.