On Optimum and Suboptimum Biasing Procedures for Importance Sampling in Communication Simulation

ROBERT J. WOLFE, MEMBER, IEEE, MICHEL C. JERUCHIM, FELLOW, IEEE, AND PETER M. HAHN, SENIOR MEMBER, IEEE

Abstract—**Importance sampling can significantly reduce simulation run time in estimating the bit error probability of digital systems if a suitable procedure is used. In this paper, we consider various aspects of optimum and suboptimum biasing procedures. We show that the strict optimum does not lead to a realistic procedure. We examine the suboptimal variations that are reasonably implementable, and for one such, we find the interesting result that for a linear system, the biasing scheme should have a time-varying mean proportional to the impulse response. Aspects of implementation are also discussed.**

I. Introduction

IMPORTANCE sampling has recently had a good deal of attention as a promising method of reducing run time in the estimation of error probability of digital systems where this estimate is obtained through simulation [1]–[9]. The essence of importance sampling is to alter the statistical properties of the noise processes[1] driving the system in such a way that many more errors occur per unit time. Since this artificial increase in error production is done in a known way, it can be corrected for. The net effect is to induce, under the proper circumstances, a decrease in the estimator variance or, equivalently, the run time, which can be dramatic. Assume, for the moment, that there is a single noise process whose M-dimensional pdf we denote as $f_N(\boldsymbol{n})$, where in the simulation context, $\boldsymbol{n} = [n(t_1), n(t_2), \cdots, n(t_M)]$ and $t_k - t_{k-1} = \Delta$, the simulation sampling interval. Implicit in this notation is the assumption that the output at any instant depends only on these M inputs. We call M the memory or dimensionality of the system. The altered process has a multidimensional density denoted by $f_N^*(\boldsymbol{n})$, which is referred to as the *biased* density. The ratio f_N^*/f_N is called the bias B_N and its reciprocal $w_N = B_N^{-1}$ is called the weight. The process of specifying and implementing f_N^* we refer to as the biasing procedure.

The most common biasing procedure, first discussed by Shanmugan and Balaban [1], amounts to increasing the variance (power) of the noise source. Recently, a variation of this approach for application to convolutionally coded systems has been suggested [7]. And, also quite recently, yet another approach has been offered by Lu and Yao [9]. This latter method consists of shifting the (nominally zero) mean of the noise distribution, while maintaining the variance at its original value. This technique will be shown to be close in some sense to the optimum procedure. In this context, we define the optimum procedure to be that which yields minimum estimator variance.

Our aim in this paper is to investigate certain issues surrounding the selection of a "good" biasing procedure. In particular, we first seek the optimum procedure, and then explore the nature of suboptimum (but more readily implemented) procedures that are suggested by the form of the optimum. What is, in fact, "optimum" is not unequivocal if we go beyond the strictly mathematical conditions that ensure minimum variance. We shall rigorously derive the form of the optimum biased density, but as we shall discuss, the mathematical optimum is, in a sense, a self-contradiction. That sense is that the solution implies knowledge of the system behavior which would preclude the need for simulation in the first place. If we constrain ourselves to a more "fuzzy" knowledge of this system, we get another optimal solution which will, in general, still be difficult to implement. We can, however, get an implementable suboptimum solution by approximating this last solution in a certain way. These various solutions imply a "block" approach to simulation, each block being the length of the system memory. Other things being equal, it is preferable to implement a "sequential" simulation. We thus also look at how the block approach can be implemented in a sequential fashion, although it turns out, unfortunately, that doing so results generally in a great loss of efficiency.

In the final section, we summarize and compare the different solutions, emphasizing the practical aspects of implementation in the simulation context.

Paper approved by the Editor for Signal Design, Modulation, and Detection of the IEEE Communications Society. Manuscript received March 29, 1988; revised November 18, 1988. This paper was presented in part at the 1988 Conference on Information Sciences and Systems, Princeton, NJ, March 1989.

The authors are with the General Electric Company, Philadelphia, PA 19101.

IEEE Log Number 9035432.

[1] One can also conceive of modifying the signal or source process. This is discussed later.

II. Problem Formulation

The error probability for a large class of communication systems can be written as

$$p = \int_{-\infty}^{\infty} H[g(\boldsymbol{s}, \boldsymbol{n})] f_S(\boldsymbol{s}) f_N(\boldsymbol{n})\, d\boldsymbol{s}\, d\boldsymbol{n} \tag{1}$$

where g is the system transfer characteristic, i.e., it specifies the output voltage at any time t, given the signal sequence $\boldsymbol{s} = [s(t), s(t-\Delta), \cdots, s(t-(M-1)\Delta)]$ and the noise sequence $\boldsymbol{n} = [n(t), n(t-\Delta), \cdots, n(t-(M-1)\Delta)]$. The system makes a hard decision, represented by the threshold function H:

$$H(x) = \begin{cases} 1, & x \in \text{error region} \\ 0, & \text{otherwise.} \end{cases}$$

Implicit in (1) is the conditioning upon a particular symbol (e.g., "1" or "0") upon which a decision is made. If different symbols are equally susceptible to error, then p is also the average symbol error rate.

Importance sampling is based on implementing the identity

$$p = \int_{-\infty}^{\infty} H[g(\boldsymbol{s}, \boldsymbol{n})] f_S(\boldsymbol{s}) w(\boldsymbol{n}) f_N^*(\boldsymbol{n})\, d\boldsymbol{s}\, d\boldsymbol{n} \tag{2}$$

where

$$w(\boldsymbol{n}) = f_N(\boldsymbol{n})/f_N^*(\boldsymbol{n}).$$

0090-6778/90/0500-0639$01.00

The empirical estimator suggested by (2) is

$$\hat{p} = \frac{1}{N}\sum_{i\in\mathcal{I}} H[g(\boldsymbol{s}_i, \boldsymbol{n}_i)]w(\boldsymbol{n}_i) \tag{3}$$

where the subscript i denotes symbol-spaced instants where decisions are made, $\mathcal{I}$ is the integer set describing the occurrences of any symbol under observation, and N is the cardinality of $\mathcal{I}$. In other words, (3) says to count the errors, weigh each by w, and average the result. The experiment in which (3) is embedded is one where f_N^* is the operative noise source.

Equations (2)–(3) are based on biasing the noise source only. The signal is not modified. In principle, one could also bias the signal, but doing so intelligently implies knowledge of the system behavior at a level of detail which is unrealistic and might well obviate the need for simulation in the first place. Consider the intuitive basis behind IS. We wish basically to throw most of the probability mass of the input processes in a region where errors are likely to occur. For noise processes, we can reasonably well imagine how to do this. For example, we know that very small values of noise will not cause errors, and sufficiently large values will. So, for example, increasing the noise variance seems a reasonable thing to start with. For the signal, the equivalent ideas are connected with sequences. Particular sequences are more prone to error than others. Hence, an importance sampling concept associated with the signal implies generating sequences with a relative frequency proportional to the degree of ISI. In order to do this, we would somehow have had to obtain the ISI for every sequence. This is impractical if the system memory is long, and if the system were linear, this knowledge would suffice to compute BER without simulation. Thus, it is customary to bias only the noise source(s). It is interesting to observe, however, that in principle, one could consider biasing the noise source in a sequence-dependent way, and this is equivalent to biasing the signal. We shall return to this in Section III.

The goodness of the estimator (3) is captured in its "time-reliability" product $\mathcal{C} = N\sigma^2(\hat{p})$ where $\sigma^2(\hat{p})$ is the variance of $\hat{p}$, and by its bias $p - E(\hat{p})$. If the true memory of the system is spanned by M samples, then $E(\hat{p}) = p$; otherwise, $E(\hat{p}) \neq p$ [2], [4], [6]. Obviously, for an IIR system, there will have to be truncation of the system response and some resulting estimator bias. This is a potentially serious source of error, but we assume here that M has been chosen to yield acceptable error, and concentrate only on $\mathcal{C}$ (further discussion on estimator bias is given in Section VI). One can show [6] that

$$\mathcal{C} = \int_{-\infty}^{\infty} H[g(\boldsymbol{s}, \boldsymbol{n})]f_S(\boldsymbol{s})w(\boldsymbol{n})f_N(\boldsymbol{n})\,d\boldsymbol{s}\,d\boldsymbol{n} - p^2 \tag{4}$$

when decisions are independent, and it is a good approximation (for the random error channel and reasonably small p) when that is not the case. We take (4) as our starting point for further development.

III. The Optimum Biased Density

We wish to minimize $\mathcal{C}$ given by (4) by suitable choice of f_N^*. The latter does not appear explicitly in (4), but of course, it is contained in $w(\boldsymbol{n})$. Since, by definition, $\mathcal{C} \geq 0$, it is sufficient to minimize the integral on the RHS of (4), namely,

$$I = \int_{-\infty}^{\infty} H[g(\boldsymbol{s}, \boldsymbol{n})]f_S(\boldsymbol{s})w(\boldsymbol{n})f_N(\boldsymbol{n})\,d\boldsymbol{s}\,d\boldsymbol{n}. \tag{5}$$

There are a number of ways to approach this problem, but one approach that we find particularly appropriate in the context of digital transmission is through the definition of an intuitively pleasing function that we call the system threshold characteristic [6] or STC. The STC, denoted T, is simply the inner integral in (5) over $\boldsymbol{s}$, viz.

$$T(\boldsymbol{n}) = \int_{-\infty}^{\infty} H[g(\boldsymbol{s}, \boldsymbol{n})]f_S(\boldsymbol{s})\,d\boldsymbol{s}. \tag{6}$$[2]

[2] Strictly speaking, the LHS should be denoted $T(g; \boldsymbol{n})$, but for brevity, we drop the g dependence.

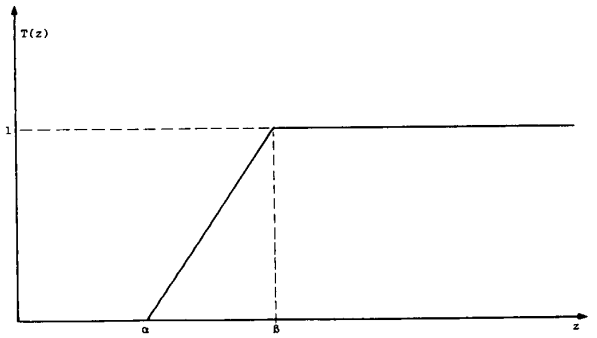

Fig. 1. Illustration of system threshold characteristic. Conditioning upon a symbol is implied. As drawn, Fig. 1 applies to a symbol which corresponds to negative voltage level. The linearity in the transition region is purely illustrative. The abscissa represents either the value of a single dimensional source ($M = 1$) or, in the case of a linear system, the output value $\boldsymbol{h} \cdot \boldsymbol{n}$.

The threshold characteristic can be looked upon as a conditional probability of error, given the noise vector $\boldsymbol{n}$. It thus embodies the ISI distribution, and even without knowing its specifics, we can visualize its general characteristics, which must be well behaved. For example, we must have $0 \leq T \leq 1$. Of course, T also depends on g. If the latter is linear, with impulse response $h(t)$, then $T(\boldsymbol{n}) = T_0(\boldsymbol{h} \cdot \boldsymbol{n})$; in other words, the STC depends only on the sampled value of the noise.[3] As an illustration, Fig. 1 shows a hypothetical threshold characteristic in one dimension (or, equivalently, for a linear system). Assuming a zero is sent, the point where T begins to be nonzero is always to the right of the origin if the eye is open.

Fig. 1 is idealized in the sense that, for a digital system with finite memory, the STC cannot be a continuous function (although it may be well approximated by one). Let m be the memory of the system in number of symbols.[4] If q is the number of different symbols, then given that the first symbol is fixed (the conditioning symbol), there are $L = q^{m-1}$ possible realizations of $\boldsymbol{s}$. Thus, for this case, we can rewrite (6) as

$$T(\boldsymbol{n}) = \sum_{i=1}^{L} H[g(\boldsymbol{s}_i, \boldsymbol{n})]f_S(\boldsymbol{s}_i) \tag{7a}$$

$$= \sum_{i=1}^{L} T(\boldsymbol{s}_i, \boldsymbol{n}) \tag{7b}$$

where i indexes the L possible M-dimensional vectors $\boldsymbol{s}$, and f_S is a discrete distribution, with weight $(1/L)$ at each $\boldsymbol{s}_i$ for a random source. Note that $T(\boldsymbol{s}_i, \boldsymbol{n})$ is the conditional probability of error and symbol sequence $\boldsymbol{s}_i$ given the noise vector $\boldsymbol{n}$. Hence, for each i, $T(\boldsymbol{s}_i, \cdot)$ is a step function in M-dimensional space, we call $T(\boldsymbol{s}_i, \boldsymbol{n})$ the sequence threshold characteristic. Depending upon whether we consider $T(\boldsymbol{n})$ to be a single system function, as in (6), or a sum of sequence-dependent functions, as in (7), we actually have two distinct optimization problems.

Indeed, in the first instance, we seek a single density $f^*(\boldsymbol{n})$ which minimizes the integral (5). In the second instance, once we allow ourselves the possibility of sequence-dependent solutions, we must generalize (5) appropriately. In particular, for each symbol sequence $\boldsymbol{s}_i$, we can seek an optimal density $f_i^*(\boldsymbol{n})$ which minimizes the variance conditioned on $\boldsymbol{s}_i$, $\sigma^2(\hat{p}_i)$ where $\hat{p}_i$ is the BER estimator for the symbol under observation, given $\boldsymbol{s} = \boldsymbol{s}_i$. From (4), this conditional variance is found from

$$N\sigma^2(\hat{p}_i) = \int_{-\infty}^{\infty} H[g(\boldsymbol{s}, \boldsymbol{n})]w_i(\boldsymbol{n})f_N(\boldsymbol{n})\,d\boldsymbol{n} - p_i^2 \tag{8}$$

[3] Throughout the paper, we use the dot notation, as in $\boldsymbol{h} \cdot \boldsymbol{n}$, to denote the scalar or inner product of vectors.

[4] Since Δ is the simulation sampling interval, the dimensionality M is related to m through $M = mT/\Delta$ where T is the symbol duration.

where N is the number of independent observations and p_i is the true BER given s_i. Then, the time–reliability product is given by the expectation over all sequences

$$\mathcal{C} = E\{N\sigma^2(\hat{p}_i)\} = \sum_{i=1}^{L} \sigma^2(\hat{p}_i) f_S(s_i). \tag{9}$$

Now, (9) is minimized if each summand is minimized. The ith summand is evidently (8) multiplied by $f_S(s_i)$, and will be minimized if the integral portion is minimized. Using (7), each such integral is expressible as

$$\int_{R^M} T(s_i, n) \frac{f^2(n)}{f_i^*(n)}\, dn \tag{10}$$

which, it is important to note, induces a set of L densities $\{f_i^*\}$ as the optimal solution. We now look at this case in the next subsection.

A. Optimum Solution Based on Sequence Threshold Characteristics

Let

$$\Omega_i = \{n \in \mathbb{R}^M \mid H[h(s_i, n)] = 1\} = \left\{ \begin{array}{l} \text{the set of noise vectors which, for given } s_i, \\ \text{are sufficient to cause an error} \end{array} \right\}.$$

Then, from the preceding discussion, we wish to minimize the functional[5]

$$J[f_1^*, f_2^*, \cdots, f_L^*] = \sum_i \int_{\Omega_i} f_S(s_i) \frac{f^2(n)}{f_i^*(n)}\, dn \tag{11a}$$

$$= \sum_i \int_{R^M} T(s_i, n) \frac{f^2(n)}{f_i^*(n)}\, dn \tag{11b}$$

subject to the constraints

$$K[f_i^*] = \int_{R^M} f_i^*(n)\, dn = 1. \tag{12}$$

As noted earlier, this formulation of the problem induces a set of L biased densities as the optimum solution. For brevity, we denote each such solution as $f_i^*(n)$. Hence, as discussed at the end of Section III, our objective is to find an f_i^* which minimizes (10).

The solution $f_i^*(n)$ which minimizes (10) is based on a version of Holder's Inequality[6] that is not well known.

Theorem [10]: If $f \in L^{p+}$, $g \in L^{q+}$ where $0 < p < 1$ and $(1/p) + (1/q) = 1$, then[7]

$$\int f(x)g(x)\, dx \geq \left[\int [f(x)]^p\, dx\right]^{1/p} \left[\int [g(x)]^q\, dx\right]^{1/q} \tag{13}$$

where the class of functions L^{p+} is the class of pth power summable nonnegative functions

$$L^{p+} = \left\{ f\colon \mathbb{R}^M \to \mathrm{R} \mid f \text{ measurable}, \; f \geq 0 \text{ and } \int [f(x)]^p\, dx < \infty \right\}. \tag{14}$$

[5] In the following derivations, for notational simplicity, we shall suppress the subscript N from the noise densities.

[6] The same results can be obtained, although less cleanly, through the calculus of variations. The latter yields the same extremum, but does not show it to be the optimum, as does this approach.

[7] For all integrals, starting with (13)–(17), the integration is over $\mathbb{R}^M$.

To apply the theorem, we note that $T(s_i, n) f^2(n) \in L^{(1/2)+}$ because

$$\int [T(s_i, n) f^2(n)]^{1/2}\, dn = \int \sqrt{T(s_i, n)} f(n)\, dn \leq \int f(n)\, dn = 1 \tag{14}$$

since $T(s_i, n) \leq 1$.

We also observe that $[f_i^*(n)]^{-1} \in L^{(-1)+}$ because

$$\int \frac{dn}{f_i^*(n)^{-1}} = \int f_i^*(n)\, dn = 1.$$

Also, because $1/(1/2) + 1/(-1) = 2 - 1 = 1$, $p = 1/2$ and $q = -1$ are conjugate exponents. Therefore, the theorem applied to (11b) gives

$$\int \frac{T(s_i, n) f^2(n)}{f_i^*(n)}\, dn \geq \left[\int \sqrt{T(s_i, n)} f(n)\, dn\right]^2 \left[\int f_i^*(n)\, dn\right]^{-1} = \left[\int \sqrt{T(s_i, n)} f(n)\, dn\right]^2. \tag{15}$$

But the left-hand side of (15) can be made equal to the right-hand side if

$$f_i^*(n) = \lambda_i \sqrt{T(s_i, n)} f(n), \tag{16}$$

where λ_i is given by

$$\lambda_i^{-1} = \int \sqrt{T(s_i, n)} f(n)\, dn. \tag{17}$$

Note that λ_i also serves to properly normalize f_i^*. It can also be shown that the solutions (16) are unique among all measurable densities.

Hence, (16) is the optimum solution for every i. Equivalently, by definition of T, (16) is expressible as

$$f_i^*(n) = \begin{cases} (\lambda_i / \sqrt{L}) f(n), & n \in \Omega_i \\ 0, & \text{elsewhere.} \end{cases} \tag{16a}$$

Using the definition $w_i = f/f_i^*$, substituting (16) into (8) can be shown to yield $N\sigma^2(\hat{p}_i) = 0$. Hence, if we were able to implement a sequence-dependent optimum biasing, the expected variance or time–reliability product given by (9) would also be zero. This implies that a single experiment, one for each symbol sequence, is sufficient. Such an experiment would require running mL bits through the simulation, i.e., L sequences m bits long. Since $\sqrt{L}/\lambda_i$ is the probability of error given sequence s_i, we are presented here with a tautology of sorts because implicit in (16) is foreknowledge of the solution. This circularity manifests itself in two ways. First, as just noted, the scale factor in (16a) is really the reciprocal of Prob[error$|s_i$], and the set of such conditional probabilities constitutes the answer we seek. Second, knowledge of the regions Ω_i is necessary in order to generate $f_i^*(n)$. But, in principle, we could then simply integrate (1) over these regions to obtain p and thus obviate the need to generate $f_i^*(n)$ at all. Actually, the qualification "in principle" should not be taken lightly. In some cases, at least, this integration may, in fact, be intractable so that we may, after all, be reduced to evaluating the integral by Monte–Carlo means.

Wang and Bhargava [8] examined the question of the optimal density by considering jointly distributed signal and noise in the product space $\mathbb{R}^{M_1} \times \mathbb{R}^{M_2} = X$ where M_1 and M_2 are the dimensions (memory) associated with signal and noise, respectively. (Obviously, if the signal is two valued, the signal space consists of the vertices of a hypercube.) If $\Omega \subset X$ is the region where errors occur, Wang and Bhargava show that $f^*(s, n) = f(s, n)/p$ is the optimum solution for $(s, n) \in \Omega$ where $f(s, n)$ is the input joint density of signal and

noise. Note that this formulation is also sequence dependent by virtue of the space on which it is defined. It also induces a biasing of the signal as well as of the noise. Because the signal and noise are not necessarily independent in this formulation, it is not clear how one would actually implement this optimal density. If, however, the signal and noise are taken to be independent (which is usually the case), then one can show that $f^*(\boldsymbol{s},\boldsymbol{n})$ can be constructed from the set of sequence-dependent densities f_i^*, namely,

$$f^*(\boldsymbol{s},\boldsymbol{n}) = \sum_{i=1}^{L} (p_i/p) f_S(\boldsymbol{s}_i) f_i^*(\boldsymbol{n}) I_i(\boldsymbol{s})$$

where $I_i(\boldsymbol{s})$ is defined as unity if $\boldsymbol{s} = \boldsymbol{s}_i$ and zero otherwise and f_i^* is given by (16). In other words, the sequence-dependent solutions (16) give an explicit method of implementing $f^*(\boldsymbol{s},\boldsymbol{n})$. It is interesting to note that the sequence-dependent implementation is equivalent to biasing both signal and noise. Obviously, we have the same fundamental obstacle as before, namely, that if p were known, we would not need to implement IS in the first place. In order to break this logical circle, we need to look for suboptimum solutions which do not require the finest level of knowledge about the system that is implied by knowledge of the error regions.

In the preceding development, we dealt with ways of modifying the input distributions, which are those directly under the control of the simulator. (We call this the "input version" of IS [6].) The output version of IS refers to properties of the decision voltage $v = g(\boldsymbol{s},\boldsymbol{n})$, the output of the system. We noted in [2] that if $w(v) = p$ for v in $g(\Omega)$ where Ω is the error region in $(\boldsymbol{s},\boldsymbol{n})$ space, then $\mathcal{C} = 0$. This solution is the mapping of (16) to the output. We also noted in [2] the inherent contradiction in this "optimum" solution.

B. Optimum Solution Based on System Threshold Characteristic

The STC does not, per se, recognize individual sequences, but represents system behavior aggregated over all sequences. For this case, there is only one biased density $f^*(\boldsymbol{n})$, and from the preceding development, this can be immediately written as

$$f^*(\boldsymbol{n}) = \lambda \sqrt{T(\boldsymbol{n})} f(\boldsymbol{n}) \tag{18}$$

with λ given by an obvious modification of (17). Note that $T(\boldsymbol{n})$ is nonzero over $\Omega = \cup \Omega_i$. Although it is possible to make reasonable assumptions about T, analogous to that represented in Fig. 1, we are not much further ahead than before in constructing an easily implementable procedure. Such a procedure should have the property that, if $f(\boldsymbol{n})$ has independent coordinates, then $f^*(\boldsymbol{n})$ should also. This would allow us to generate $f^*(n)$ sequentially with an independent generator from sample to sample, which is a great convenience. We find, as discussed next, that (18) does lead to such a property if the system is linear. As usual, this linear case is not as much of interest in itself as it is in pointing to a procedure which may also be useful in the nonlinear case. It is important to note that, unlike the sequence-dependent case, $\mathcal{C}$ must be greater than zero. Essentially, this is because we have averaged information about the system in adopting a single biasing approach for all sequences.

C. The Linear Case

In the linear case, we remarked earlier that

$$T(\boldsymbol{n}) = T_0(\boldsymbol{h} \cdot \boldsymbol{n}) \tag{19}$$

where T_0 can be looked upon as the conditional probability of error given $z = \boldsymbol{h} \cdot \boldsymbol{n}$, which is the noise variable as seen at the output of the system (input to decision device). Hence, T_0 is the output version of T and would look something like that shown in Fig. 1. Consequently, from (18),

$$f^*(\boldsymbol{n}) = \lambda \sqrt{T_0(\boldsymbol{h} \cdot \boldsymbol{n})} f(\boldsymbol{n}) \tag{20}$$

is the optimum M-dimensional input density. Unfortunately, as written, even if the components of $f(\boldsymbol{n})$ were independent, we could not generate $f^*(\boldsymbol{n})$ with independent components since $\sqrt{T_0(\boldsymbol{h} \cdot \boldsymbol{n})}$ is not factorizable in the given coordinates.

However, if we assume certain properties for the distribution of $\boldsymbol{n}$, then (20) can, in fact, be expressed as a product, through an appropriate transformation. Specifically, we assume that $\boldsymbol{n}$ is radially distributed in the sense that $f(\boldsymbol{n}) = \phi(\|\boldsymbol{n}\|)$ for some density ϕ. Let $\{\boldsymbol{b}_1, \boldsymbol{b}_2, \cdots, \boldsymbol{b}_M\}$ be an ordered orthonormal basis of $\mathbb{R}^M$ with $\boldsymbol{b}_1 = \boldsymbol{h}$ and $\|\boldsymbol{h}\| = 1$. Let $\mathcal{U}$ be the $M \times M$ matrix with columns $\boldsymbol{b}_1, \cdots, \boldsymbol{b}_M$, viz. $\mathcal{U} = [\boldsymbol{h}, \boldsymbol{b}_2, \cdots, \boldsymbol{b}_M]$. Set $\boldsymbol{n} = \mathcal{U}\boldsymbol{v}$. Then

$$I = \int_{\mathbb{R}^M} T(\mathcal{U}\boldsymbol{v}) \frac{f^2(\mathcal{U}\boldsymbol{v})}{f^*(\mathcal{U}\boldsymbol{v})} |\det \mathcal{U}| \, d\boldsymbol{v} \tag{21a}$$

$$= \int_{\mathbb{R}^M} T_0(\boldsymbol{h} \cdot \mathcal{U}\boldsymbol{v}) \frac{f^2(\boldsymbol{v})}{f_1^*(\boldsymbol{v})} \, d\boldsymbol{v} \tag{21b}$$

where, for simplicity, we set $f^*(\mathcal{U}\boldsymbol{v}) = f_1^*(\boldsymbol{v})$. Since f has radial symmetry, $f(\mathcal{U}\boldsymbol{v}) = f(\boldsymbol{v})$. Also, because $\boldsymbol{h} \cdot \mathcal{U}\boldsymbol{v} = \mathcal{U}^T h \cdot v = (1, 0, \cdots, 0) \cdot \boldsymbol{v} = v_1$ we have

$$I = \int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} T_0(v_1) \frac{f^2(v_1, \cdots, v_M)}{f_1^*(v_1, \cdots, v_M)} \, dv_1 \cdots dv_M \tag{22}$$

whence, using previous results, the optimal f_1^* is

$$f_1^*(v_1, \cdots, v_M) = \frac{\sqrt{T_0(v_1)} f(v_1, \cdots, v_M)}{\int_{-\infty}^{\infty} \sqrt{T_0(v_1)} m(v_1) \, dv_1} \tag{23}$$

where $m(v_1)$ is the marginal density of $f(v_1, \cdots, v_M)$.

Now, if $f(\boldsymbol{n})$ is Gaussian and has radial symmetry, then it has identical independent components, $f(\boldsymbol{n}) = \Pi f_i(\boldsymbol{n}_i)$, with equal variance σ^2. By definition, $f(\boldsymbol{n}) = f(\mathcal{U}_v) = f(\boldsymbol{v})$; hence, $f(\boldsymbol{v})$ has the same properties. Therefore, (23) is factorizable into M independent densities, the first of which is proportional to $\sqrt{T_0(v_1)} f_1(v_1)$ and the remaining $(M-1)$ of which are identical to one another. A point worthy of note is that v_1 is the $\boldsymbol{h}$ coordinate of $\boldsymbol{v}$ in the basis defined by $\mathcal{U}$. As will be seen, this direction is important in a suboptimum implementation of the procedure. The optimum procedure implied by the preceding development implies a block approach wherein M noise values are independently generated as just described, transformed through the matrix $\mathcal{U}$ into a new block of M values, and sent through the system. As noted elsewhere, only *one* symbol decision per block can be made.

If $f_i(v_i)$ is $N(0, \sigma^2)$, then $p(v_1) = \lambda \sqrt{T_0(v_1)} f_1(v_1)$ is a density that looks something like that in Fig. 2 where T_0 is taken to be of the same form as in Fig. 1. The exact shape of $p(v_1)$ depends on the relative values of α, β, and σ. The peak value can be shown to occur at $z = \alpha(1 + \sqrt{1 + (\sigma/\alpha)^2})$ if it occurs between α and β; otherwise, it occurs at β. The most significant aspect of $p(v_1)$ is the fact that it is a density with the mode shifted away from zero. This suggests that a practical alternative to actually attempting to generate $p(v_1)$ exactly is to approximate it with a nonzero-mean Gaussian having the same variance as $v_2, \cdots, v_M$. If we do this and then transform $(v_1, v_2, \cdots, v_M)$ back into the original coordinate system to get $(n_1, n_2, \cdots, n_M)$, the latter will all be normally distributed with the same variance, but generally different means, i.e., $n_i = N(\mu_i, \sigma^2)$. Indeed, if $\int v_1 p(v_1) \, dv_1 = a$, then $E(\boldsymbol{v}) = \mathcal{U}(a, 0, \cdots, 0)^T = a\boldsymbol{h}$, and so the mean vector is a multiple of $\boldsymbol{h}$. This leads us, therefore, to consider a biasing scheme where the biased density is normal with standard deviation $\sigma_*^2 = \gamma^2 \sigma^2$ and a (possibly) time-varying (or position-varying) mean μ_i, $i = 1, 2, \cdots, M$. This approach is discussed in more detail below.

IV. Biasing with Nonzero-Mean Gaussian Densities

We assume here that the input noise samples are independent.[8] At each time sample (coordinate), a normal r.v. is generated with

[8] This assumption is not central. If the initial coordinates were correlated, one could find an orthogonal transformation that would yield independent coordinates. However, the assumption simplifies the presentation. In addition, in the simulation context, one can, and usually does, generate initially independent variates.

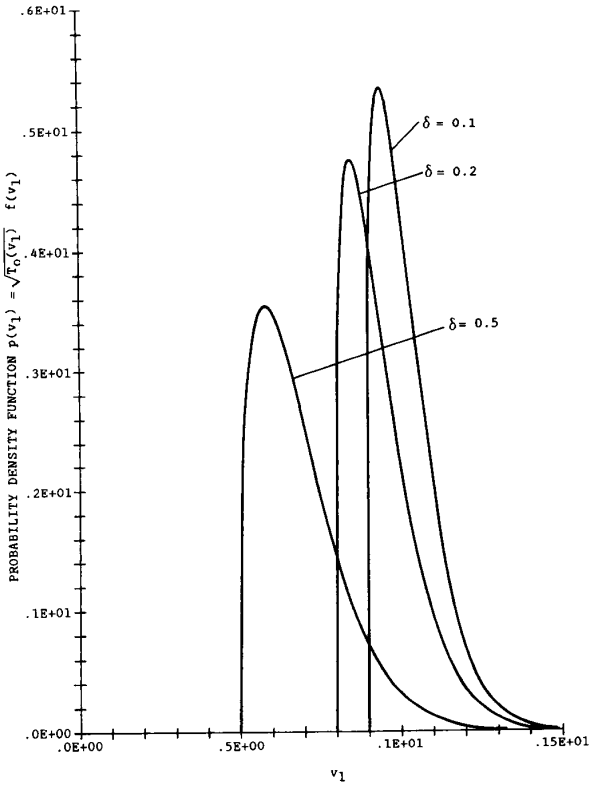

Fig. 2. Probability density function $p(v_1)$ of component of biased density in the "direction" of the impulse response. For this figure, $\rho = 10$ dB [see definitions of ρ and δ following (27)].

distribution $N(\mu_i, \sigma_*^2)$, $i = 1, 2, \cdots, M$ where $\sigma_*^2 = \gamma^2\sigma^2$ and σ^2 is the variance corresponding to the operating point of interest. Elsewhere [6], we referred to γ as the σ multiplier. The case where μ_i is nonzero but $\gamma = 1$ was originally considered by Lu and Yao [9]. For the conditions given (keeping in mind the linear assumptions), we can show that (7) can be expressed as

$$I(r, \boldsymbol{C}) = \left(\frac{r^2}{\sqrt{2r^2 - 1}}\right)^M e^{||\boldsymbol{C}||^2(2r^2-1)/\sigma^2} \cdot \int_{-\infty}^{\infty} \frac{T_0(x)}{\sqrt{2\pi\sigma^2 r^2}} e^{-[x-(1-2r^2)(\boldsymbol{h}\cdot\boldsymbol{C})]^2/2\sigma^2 r^2}\, dx \quad (24)$$

where (for notational convenience) we define

$$r^2 = \gamma^2/(2\gamma^2 - 1) \quad (25a)$$

$$\boldsymbol{C} = (\mu_1, \mu_2, \cdots, \mu_M) \quad (25b)$$

and T_0 is the STC. For the case $\boldsymbol{C} = \boldsymbol{0}$, (24) reduces to the traditional method of biasing where only the noise variance is modified.

We are looking for the minimum of I. It turns out that the expression for the optimum $(r, \boldsymbol{C})$ is uninstructively complicated. It is more illuminating to ask what is the optimum $\boldsymbol{C}$, given a value of r. In the Appendix, we show that for any r and T, $I(r, \boldsymbol{C})$ is minimized if $\boldsymbol{C} = c\boldsymbol{h}$ where c is a scalar constant. In other words, the mean μ_i at each instant is proportional to the value of the impulse response at that instant. In geometrical terms, this means that $\boldsymbol{C}$ is in the "direction" of $\boldsymbol{h}$ in $\mathbb{R}^M$. This can be heuristically illustrated, as in the sketch of Fig. 3. We note that in a linear system $T(\boldsymbol{n}) = T_0(\boldsymbol{h}\cdot\boldsymbol{n})$. Hence, T is constant on hyperplanes perpendicular to $\boldsymbol{h}$. If $T = 0$ for $\boldsymbol{h}\cdot\boldsymbol{n} \leq \alpha$ and $T = 1$ for $\boldsymbol{h}\cdot\boldsymbol{n} \geq \beta$, then the corresponding hyperplanes L_0, L_1 bound the transition region in a manner analogous to the one-dimensional view of Fig. 1.

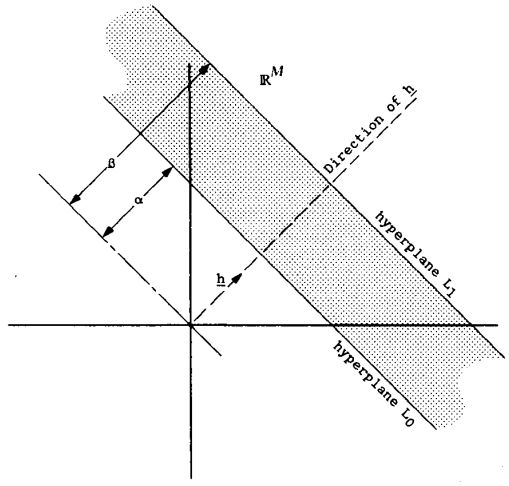

Fig. 3. Illustration of M-dimensional threshold characteristic. Impulse response $\boldsymbol{h}$ is perpendicular to hyperplanes L_0 at distance α and L_1 at distance β from the origin.

It is significant to note that if $r = 1$ (i.e., the noise variance is unchanged) and $\boldsymbol{C} = c\boldsymbol{h}$, then $I(r, \boldsymbol{C})$ becomes *independent* of M irrespective of the form of T. It is interesting to observe that the direction of $\boldsymbol{h}$ is also optimal in a different sense. Davis [4] has shown that in the conventional implementation ($\mu_i = 0$, all i), an orthogonal transformation (e.g., the matrix $\mathcal{U}$ earlier defined) with the first column equal to $\boldsymbol{h}$ creates an equivalent linear system with unit memory, which is the best situation in that case. We might perhaps coin a "matched filtering" principle in this context!

The best value of the scalar c that multiplies $\boldsymbol{h}$ is a function of T and can be found from the equation $\partial I/\partial c(r, c\boldsymbol{h}) = 0$. However, the solution is not obvious, and it is more instructive to look at a particular case, which we shall do presently. A point to be made in this regard is that, as in the formulation of the optimum, we can interpret T as a single system function or as a sum of sequence-dependent functions as in (9). In the former instance, the scalar c is a single number, while in the latter case, there is a different scalar corresponding to each distinct sequence. This last approach, using a sequence-dependent mean for the biased density, is considered by Lu and Yao [9].

In either of these two approaches, a block implementation is clearly implied since the sequence of means is not constant. If, in addition, the sequence of means is the set $\{c_i\boldsymbol{h}\}$ where c_i is the scalar corresponding to the ith sequence, there is a further implication concerning the complexity of implementation as well as the realism of the approach. The complexity issue is not overriding. But we would certainly have to know the sampled signal value for each possible sequence—this is implied by knowledge of T. For a 10 bit memory, say, this means knowing 2^{10} values. Aside from the effort in gathering (computing) these values, once we knew them, there would be no point in proceeding with the simulation. We could directly compute the error probability! We are thus in the same kind of logical impasse as in the sequence-dependent optimal formulation.

From an implementation standpoint, it would certainly be simpler if the sequence of means were constant. Among other things, it would permit a sequential implementation, which is generally more desirable. We thus look into this case in the next section.

A. Biasing with Fixed-Mean Gaussian Densities: Linear STC

We now assume $\boldsymbol{C} = c(1, 1, \cdots, 1)$ i.e., the mean at each sample time is a fixed amount. Clearly, there is then no need to block the data. Further, in order to obtain computable results, we need to make a specific assumption concerning T. We shall assume that we

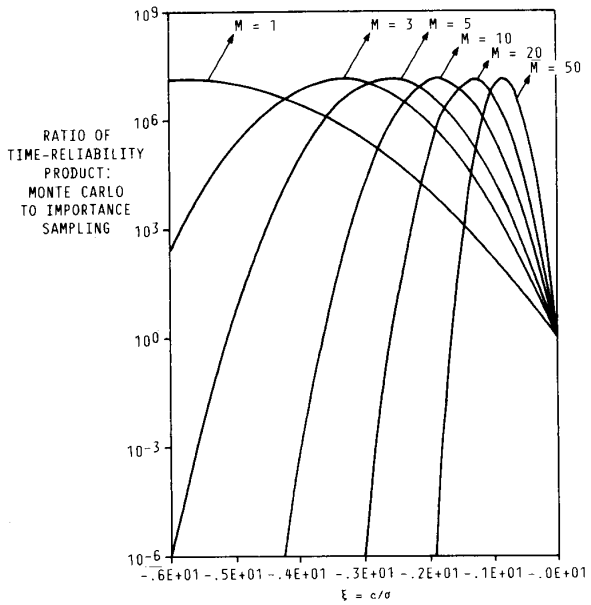

Fig. 4. Importance sampling improvement versus $\zeta = c/\sigma$ where c is the mean of the Gaussian distribution; M is the system memory and $p = 10^{-8}$.

do not have knowledge of sequence-level detail. Hence, T_0 is the *system* threshold characteristic, and we assume it to be "linear," i.e., a straight line in the transition region, as shown in Fig. 1. This assumption is for computational ease, and has no connection with the linearity of the system. Thus,

$$T_0(x) = \begin{cases} 0, & x \le \alpha \\ \dfrac{x-\alpha}{\beta-\alpha}, & \alpha < x < \beta \\ 1, & x \ge \beta. \end{cases} \tag{26}$$

Given arbitrary $\boldsymbol{C} \in \mathbb{R}^M$, (24) now takes the form

$$I(r, \boldsymbol{c}) = \left(\frac{r^2}{\sqrt{2r^2-1}}\right)^M \frac{e^{\|\boldsymbol{C}\|^2(2r^2-1)/\sigma^2}}{B-A} \cdot [BQ(B) - AQ(A) + \phi(A) - \phi(B)] \tag{27}$$

where

$$A = \frac{\sqrt{\rho}}{r}\left[1 - \delta + \frac{(\boldsymbol{h}\cdot\boldsymbol{C})(2r^2-1)}{\sigma\sqrt{\rho}}\right]$$

$$B = \frac{\sqrt{\rho}}{r}\left[1 + \delta + \frac{(\boldsymbol{h}\cdot\boldsymbol{C})(2r^2-1)}{\sigma\sqrt{\rho}}\right]$$

$$\delta = (\beta-\alpha)/(\beta+\alpha)$$

$$\rho = [0.5(\beta+\alpha)]^2/\sigma^2 = SNR$$

$$Q(x) = \int_{-\infty}^{x} \phi(y)\,dy$$

$$\phi(x) = (\sqrt{2\pi})^{-1} \exp(-x^2/2)$$

and in the preceding, $\sum h_i^2 = 1$ is implicit. When $\boldsymbol{C} = c(1, 1, \cdots, 1)$, (27) and the subsequent definitions apply simply by replacing $\boldsymbol{h}\cdot\boldsymbol{C}$ by $c\eta$ where $\eta = \sum h_i$, and $\|\boldsymbol{C}\|^2$ becomes Mc^2.

In the case that a system is characterizable as an I&D filter, (27) becomes independent of M if $r = 1$. This can be seen from the fact that $\sum h_i^2 = 1$ implies $h_i = 1/\sqrt{M}$, whence $\eta = \sqrt{M}$. Hence, setting $x = c\sqrt{M}$, say, we see that (27) depends only on the variable x; hence, the maximum value of I will be independent of M. For $p = 10^{-8}$, Fig. 4 shows the ratio $\mathcal{C}$ in the Monte-Carlo case to that in the present (I&D) case with $r = 1$ and $\delta = 0$.[9] The ordinate gives the run time improvement directly for the same estimator variance (roughly, the same confidence). The abscissa $\zeta = c/\sigma$. As can be seen, tremendous improvement is potentially available, independent of M. Although this improvement must be somewhat dependent on T_0, it is not likely that the T_0 in any system would depart radically from (26); hence, Fig. 4 is probably representative. (This is also supported by the numerical results presented in [9].) Unfortunately, this improvement will rarely be attained in practice (using the approach of this section). This is because practical systems generally do not have an $\boldsymbol{h}$ corresponding to an I&D filter. There may be an I&D filter in the receiver, but the overall transfer function represented by $\boldsymbol{h}$ will not be so idealized. If, for example, we take the family of Butterworth filters as more representative of reality, we compute (for two–six poles) a *maximum* improvement of about 2!

It may seem surprising that the improvement is so sensitive to the nature of $\boldsymbol{h}$ when we use a constant mean shift. This is really a manifestation of a more general property of $I(r, \boldsymbol{C})$, namely, that the improvement decreases exponentially with the difference between $\boldsymbol{C}$ and $\boldsymbol{h}$. (This statement is made more precise in the Appendix.)

V. Multiple Noise Sources

In perhaps most problems of interest, we have to deal with more than one noise source. Certainly, for bandpass systems, even if there is only one physical noise source, the simulation of such a source necessitates the activation of two (in-phase and quadrature) sources. If there are multiple independent sources, all that we have said earlier carries over in a natural way. The optimal formulations are just as (if not less) unrealizable as in the single-source case. The most interesting generalization of the single-source case is when the kth source is M_k-dimensional zero-mean Gaussian, and we bias by shifting the mean of each source. If each source has impulse response $\boldsymbol{h}_k$ associated with it, then one can show that the optimal (block) procedure consists of shifting the (position-dependent) mean of the kth source by $c_k\boldsymbol{h}_k$ where c_k is a constant.

Summary and Discussion

A good biasing procedure for importance sampling should provide a usefully large run-time improvement and (generally less important) should be relatively straightforward to implement. By this, we do not mean (necessarily) fewer lines of code. Rather, we believe it is preferable for a biasing procedure to entail no change in the simulation structure that would be used for a Monte-Carlo simulation. One would simply "turn on" the biased noise generator(s). Generally, a Monte-Carlo simulation is "sequential," i.e., it evolves in time. The implementation can be "time domain" where a new sample is generated at every simulation clock tick at every node or it can use blocks of samples for FFT processing. However, provided the blocks are properly spliced, this is equivalent to sequential processing. The main desirability of sequential time is the ability to simulate long-memory feedback devices such as phase-locked loops. We should point out, however, that there are many cases where explicit simulation of such devices need not be done. In such cases, the appeal of a sequential simulation is perhaps more esthetic than substantive because it mirrors reality, but after all, analytic techniques do not generally emulate reality and can still provide perfectly good results. Hence, although we favor a sequential procedure, a nonsequential one can be applicable, and may even be the method of choice if it provides superior run time improvement.

The optimum and suboptimum procedures (except for the one of Section IV-A) all imply a nonsequential implementation. We called these "block" approaches, but with a different sense than that used

[9] The condition $\delta = 0$ is necessary in order to be consistent with the I&D assumption.

just above. Here, the blocks must be m bits long where m is the memory. The choice of m presents a tradeoff between estimator bias and run time improvement (equivalently, variance). This is because only one bit decision is valid in any block, namely, that for the most recent bit since it is the only one in the block that has experienced the full effect of the past. Thus, the run-time improvement for these methods must be reduced by the factor $(1/m)$. If the system were IIR, there would, of course, be no choice but to truncate the response, but even if the system response were truly time limited, we might still wish to truncate it so as to make m not overly large. In either case, the truncation implies estimator bias, but the manner in which this bias arises is quite different in the block processing case than in the sequential implementation. In the latter situation, the system response itself is not (or need not be) truncated when running the simulation. What is truncated in conventional importance sampling is the extent of the past taken into account in the weighting procedure of (3). Such truncation can be shown to yield an estimate $\hat{p}$ such that $E(\hat{p}) > p$ [6], but quantitative evaluation of $E(\hat{p}) - p$ generally cannot be done without specific and simplifying assumptions on the system. For the block implementation, on the other hand, truncation of the system response results in bias, even without importance sampling. The extent of this bias can be set up symbolically using (1), but again, cannot be evaluated without rather specific assumptions on the system. Unlike the sequential case, however, this bias can be either positive or negative. The biasing scheme for IIS-type importance sampling specifies that the mean of the noise density be proportional to the system response. The scheme, in effect, *requires* that the weighting not be truncated, and in fact, there is nothing to be gained by doing so. If we adhere to this requirement, IIS, per se, does not additionally bias the estimator beyond what may already exist by virtue of the block implementation. In order to choose an m which effects a good balance between run-time and estimator bias, it may be necessary for the simulator to do a side experiment in order to see what is the smallest m that can be used without inducing unacceptable bias.

Another point to remember in the implementation is that the initial formulation was conditioned upon the decisioning of a given symbol. If the error probability is different for different symbols, as it will often be due to asymmetries in the waveform, then the procedure must be replicated for each type of symbol as the most recent one in the sequence. The effect of the symbol to be detected on the procedure is to specify a different region of $\mathbb{R}^M$ as the error region. The net effect on the procedures of Section IV is that the sign of c must be coupled with the symbol. The procedure of Section IV-A is sequential, and hence does not suffer from the $(1/m)$ reduction. However, it too is symbol sensitive, which implies that when a given value of c is used, symbols of the "opposite" type must be ignored. On the average, then, a factor of two reduction in efficiency would be experienced by this technique. Unfortunately, even though this method (of Section IV-A) is sequential, it cannot be relied upon to yield appreciable improvement for arbitrary $\boldsymbol{h}$. It should be noted that the conventional technique of increasing the noise variance maintains the noise pdf symmetry with respect to the symbols, and hence is symbol neutral.

Within the set of block approaches, we also noted that there is, in principle, an option to make the procedure sequence dependent or not. However, as noted earlier, the sequence-dependent approach may not be practical for systems with large M, but there is still a very large improvement obtainable, even if we restrict ourselves to using a system threshold characteristic, as implied by the example of Section IV-A.

It has been pointed out that the optimum solution requires an unrealistically detailed level of knowledge about the system. The suboptimum solutions also imply some knowledge, but of a more readily obtainable form. In particular, the method of Section IV stipulates knowledge of $\boldsymbol{h}$, the system "impulse response." It appears that any refined approach to IS requires some form of system identification. The method proposed by Davis [4] also requires knowledge of $\boldsymbol{h}$, as does another form of "efficient" IS suggested by the authors [6] which can be sequentially implemented. This is obtainable by different means within the simulation itself. The very definition of $\boldsymbol{h}$ implies a linear system or a linear approximation of a nonlinear system. What we cannot answer at present is the sensitivity of various approaches to the linear equivalent embodied in $\boldsymbol{h}$. The answer to this question may have a significant influence on the biasing procedure of choice. Otherwise, the preferred technique is not necessarily obvious, as some balance between ease of implementation and run time improvement must be struck, and this may well be application dependent.

Appendix

We first show that $I(r, \boldsymbol{C})$ given by (24) is minimized when $\boldsymbol{C} = c\boldsymbol{h}$ for some $c \in \mathbb{R}$. Let

$$\boldsymbol{h}^{\perp} = \{\boldsymbol{K} \in \mathbb{R}^M \mid \boldsymbol{h} \cdot \boldsymbol{K} = 0\}$$

be the orthogonal complement of $\boldsymbol{h}$.

Lemma A.1: Every $\boldsymbol{C} \in \mathbb{R}^M$ admits a unique decomposition $\boldsymbol{C} = a\boldsymbol{h} + \boldsymbol{K}$ where $a \in \mathbb{R}$ and $\boldsymbol{K} \in \boldsymbol{h}^{\perp}$. Moreover,

$$a = \frac{\boldsymbol{h} \cdot \boldsymbol{C}}{\|\boldsymbol{h}\|} \tag{A.1}$$

and

$$\boldsymbol{K} = \boldsymbol{C} - \frac{(\boldsymbol{h} \cdot \boldsymbol{C})\boldsymbol{h}}{\|\boldsymbol{h}\|}. \tag{A.2}$$

The proof is straightforward and is omitted.

Let

$$\psi\colon \mathbb{R}^M \to \operatorname{sp}(\boldsymbol{h}) \triangleq \{t\boldsymbol{h} \mid t \in \mathbb{R}\} \tag{A.3}$$

be the orthogonal projection of $\mathbb{R}^M$ onto the subspace spanned by $\boldsymbol{h}$. By the lemma above,

$$\psi(\boldsymbol{C}) = (\boldsymbol{h} \cdot \boldsymbol{C})\boldsymbol{h} \tag{A.4}$$

if $\|\boldsymbol{h}\| = 1$, which we assume.

Theorem A.1:[10] For $\boldsymbol{C} \notin \operatorname{sp}(\boldsymbol{h})$, $r > \sqrt{2}/2$,

$$I(\psi(\boldsymbol{C})) < I(\boldsymbol{C}).$$

Remark: The theorem says that variance is always reduced by projecting onto the subspace spanned by $\boldsymbol{h}$, i.e., minimum variance occurs by translation along $\boldsymbol{h}$.

The proof depends on a simple lemma.

Lemma A.2: $I(a\boldsymbol{h} + \boldsymbol{K}) = \exp[\|\boldsymbol{K}\|^2(2r^2 - 1)/\sigma^2] I(a\boldsymbol{h})$ for $a \in \mathbb{R}$ and $\boldsymbol{K} \in \boldsymbol{h}^{\perp}$.

Proof of Lemma A.2: By definition [from (24)]

$$I(a\boldsymbol{h} + \boldsymbol{K}) = G(\sigma, r, a\boldsymbol{h} + \boldsymbol{K}) \int_{-\infty}^{\infty} \frac{T_0(x)}{\sqrt{2\pi\sigma^2 r^2}} \cdot \exp\left(-\frac{[x - (1 - 2r^2)(a\boldsymbol{h} + \boldsymbol{K}) \cdot \boldsymbol{h}]^2}{2\sigma^2 r^2}\right) dx \tag{A.5}$$

where $G(\sigma, r, \boldsymbol{C}) = \left(\frac{r^2}{\sqrt{2r^2 - 1}}\right)^M \exp(\|C\|^2(2r^2 - 1)/\sigma^2)$.

Now,

$$(a\boldsymbol{h} + \boldsymbol{k}) \cdot \boldsymbol{h} = a(\boldsymbol{h} \cdot \boldsymbol{h}) + \boldsymbol{K} \cdot \boldsymbol{h} = a \tag{A.6}$$

since $(\boldsymbol{K} \cdot \boldsymbol{h}) = 0$ and $\|\boldsymbol{h}\| = 1$. Also,

$$\|a\boldsymbol{h} + \boldsymbol{K}\|^2 = (a\boldsymbol{h} + \boldsymbol{K}) \cdot (a\boldsymbol{h} + \boldsymbol{K}) = a^2 + \|\boldsymbol{K}\|^2. \tag{A.7}$$

[10] For notational simplicity, we henceforth suppress the explicit dependence of I on r.

Substituting (A.6) and (A.7) into (A.5) yields

$$I(a\boldsymbol{h}+\boldsymbol{K}) = e^{|\boldsymbol{K}|^2(2r^2-1)/\sigma^2} I(a\boldsymbol{h}). \qquad \text{Q.E.D.}$$

Proof of Theorem: From Lemma A.1 and the definition of ψ, it follows that we can write

$$\boldsymbol{C} = \psi(\boldsymbol{C}) + \boldsymbol{K}.$$

Furthermore, Lemma A.2 shows that

$$I(\boldsymbol{C}) = e^{|\boldsymbol{K}|^2(2r^2-1)/\sigma^2} I(\psi(\boldsymbol{C})).$$

Since $\boldsymbol{C} \notin \text{sp}(\boldsymbol{h})$ implies $\boldsymbol{K} \neq 0$, we have

$$I(\boldsymbol{C}) > I(\psi(\boldsymbol{C})). \qquad \text{Q.E.D.}$$

Now, we compare the variance in the $\mathbf{1}$ direction to the optimal translation (along $\boldsymbol{h}$).

Let $\mathbf{1} \triangleq (1, 1, \cdots, 1)$, whence $c\mathbf{1} = (c, c, \cdots, c)$, and write

$$c\mathbf{1} = (c\mathbf{1}\cdot\boldsymbol{h})\boldsymbol{h} + (c\mathbf{1} - (c\mathbf{1}\cdot\boldsymbol{h})\boldsymbol{h})$$

$$= \eta c\boldsymbol{h} + \boldsymbol{K} \quad \text{since } \eta = \sum h_i. \tag{A.8}$$

Then, since $\|\boldsymbol{h}\| = 1$,

$$I(c\mathbf{1}) = e^{c^2|1-\eta h|^2(2r^2-1)/\sigma^2} I(c\eta\boldsymbol{h}). \tag{A.9}$$

But

$$\sum_{i=1}^{M}(1-\eta h_i)^2 = M - \eta^2 = \|1-\eta h\|^2; \tag{A.10}$$

hence,

$$I(c\mathbf{1}) = e^{c^2(M-\eta^2)(2r^2-1)/\sigma^2} I(c\eta\boldsymbol{h}) \tag{A.11}$$

or

$$I(c\mathbf{1}) = e^{Mc^2(2r^2-1)\sin\theta/\sigma^2} I(c\eta\boldsymbol{h})$$

where θ, $0 \le \theta \le \pi$ is the angle between $\mathbf{1}$ and $\boldsymbol{h}$, and we see from (A.11) that (even with $r = 1$) the variance is exponentially increasing with M, unless $M = \eta^2$. But this is precisely the case for an I&D filter for which the normalization $\sum h_i^2 = 1$ holds. In other words, for this filter, the $\mathbf{1}$'s direction *is* the direction of $\boldsymbol{h}$. Otherwise, as noted in the body, the exponential dependence on M induces a major degrading effect on the improvement.

To see how system identification can affect variance, let $\boldsymbol{C}$ be an approximation to $\boldsymbol{h}$, the true finite-dimensional impulse response. Then we can write

$$\boldsymbol{C} = \psi(\boldsymbol{C}) + \boldsymbol{K}$$

as before and apply Lemma A.2 to find that for any $r > \sqrt{2}/2$

$$I(r, \boldsymbol{C}) = \exp(\|\boldsymbol{K}\|^2(2r^2-1)/\sigma^2) I(r, \psi(\boldsymbol{C}))$$

$$\ge p^2 \exp(\|\boldsymbol{K}\|^2(2r^2-1)/\sigma^2).$$

But $\|\boldsymbol{K}\|$ is the distance between $\boldsymbol{C}$ and the subspace $\text{sp}(\boldsymbol{h})$ defined in (A.3), and so variance grows exponentially in the distance between $\text{sp}(\boldsymbol{h})$ and the approximation $\boldsymbol{C}$. Therefore, careful and accurate system identification needs to be implemented for the technique to achieve full potential.

Acknowledgment

The authors are indebted to M. Schiavoni for programming the equations leading to Figs. 2 and 4 of the text.

References

[1] K. S. Shanmugan and P. Balaban, "A modified Monte Carlo simulation technique for the evaluation of error rate in digital communications systems," *IEEE Trans. Commun.*, vol. COM-28, Nov. 1980.

[2] M. C. Jeruchim, "Techniques for estimating the bit error rate in the simulation of digital communications systems," *IEEE J. Select. Areas Commun.*, vol. SAC-2, pp. 153–170, Jan. 1984.

[3] ——, "On the application of importance sampling to the simulation of digital satellite and multihop links," *IEEE Trans. Commun.*, vol. COM-32, Oct. 1984.

[4] B. R. Davis, "An improved importance sampling method for digital communication system simulations," *IEEE Trans. Commun.*, vol. COM-34, pp. 715–719, July 1986.

[5] P. M. Hahn, M. C. Jeruchim, and T. J. Klandrud, "Implementation of importance sampling in multihop communication simulation," in *Proc. IEEE GLOBECOM '86*, vol. 1, Houston, TX, Dec. 1986, pp. 4.1.1–4.1.5.

[6] P. M. Hahn and M. C. Jeruchim, "Developments in the theory and application of importance sampling," *IEEE Trans. Commun.*, vol. COM-35, pp. 706–714, July 1987.

[7] M. A. Herro and J. M. Nowack, "The use of importance sampling in coded digital communication simulation," presented at the 5th Int. Conf. Syst. Eng., Dayton, OH, Sept. 1987.

[8] Q. Wang and V. K. Bhargava, "On the application of importance sampling to BER estimation in the simulation of digital communication systems," *IEEE Trans. Commun.*, vol. COM-35, pp. 1231–1233, Nov. 1987.

[9] D. Lu and K. Yao, "Improved importance sampling technique for efficient simulation of digital communication systems," *IEEE J. Select. Areas Commun.*, vol. 6, Jan. 1988.

[10] E. Hewitt and K. Stromberg, *Real and Abstract Analysis*. New York: Springer-Verlag, 1965, p. 191.

★

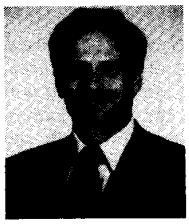

Robert J. Wolfe (M'87) was born in New York, NY, in 1947. He received the B.S. degree in mathematics from Long Island University, New York, NY, in 1970 and the M.A. and Ph.D. degrees in mathematics from the State University of New York at Buffalo, Amherst, NY, in 1972 and 1976, respectively.

From 1976 to 1980 he was an Assistant Professor of Mathematics and Statistics at Case Western Reserve University, Cleveland, OH. From 1980 to 1986 he worked for the Department of Defense, Fort Meade, MD, in the area of secure communications. Since 1986 he has been with General Electric Aerospace, King of Prussia, PA, working on the development of simulation tools and techniques for the performance evaluation of digital satellite communications systems.

★

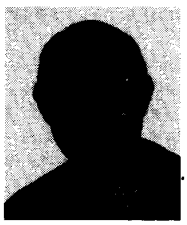

Michel C. Jeruchim (S'60–M'61–SM'81–F'86) was born in Paris, France, in 1937. He received the B.E.E. degree (cum laude) from the City College of New York, New York, NY, in 1961 and the M.S.E.E. and Ph.D. degrees from the University of Pennsylvania, Philadelphia, in 1963 and 1967, respectively.

Since 1961 he has been with General Electric Aerospace, King of Prussia, PA, working in a variety of communication-related disciplines as applied to satellite communications. One of his main activities has been the development of simulation tools for the performance evaluation of digital satellite communications systems.

Dr. Jeruchim has published or presented more than two dozen papers on various topics in communications, and is coauthor of the book, *Communication Satellites in the Geostationary Orbit* (Artech, 2nd ed., 1987). He has served as a member of a number of U.S. delegations to international technical standards and radio regulatory conferences. He is currently Chairman of the IEEE Communications Society's Subcommittee on Computer-Aided Modeling, Analysis, and Design of Communication Systems.

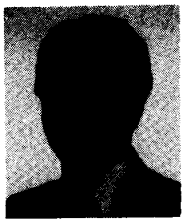

Peter M. Hahn (M'60-SM'71) was born in Vienna, Austria, in 1937. He received the B.E.E. degree from the City College of New York, New York, NY, in 1958 and the M.S. and Ph.D. degrees in electrical engineering from the University of Pennsylvania, Philadelphia, in 1962 and 1968, respectively.

He is Manager, Technology Development and Production at General Electric Aerospace, Blue Bell, PA. He is responsible for technology development, producibility assurance, and survivability/vulnerability analysis for the Strategic Defense System Systems Engineering Program. After graduation in 1958, he was employed by Philco Research for three years, attended graduate school, and worked at RCA until 1967. From 1967 to 1976 he worked on communication and postal systems at Philco-Ford Corporation (later Ford Aerospace and Communications Corporation). His last position at Ford was an Engineering Section Head for packet- and message-switching systems. From 1976 to 1985 he worked at RCA Government Systems Division as an Advanced Programs Manager and as a staff Technical Advisor. From 1985 to 1988 he supervised communications systems analysis and simulation at General Electric. In addition, he is an Adjunct Assistant Professor at the University of Pennsylvania, and an Adjunct Professor at Drexel University, Philadelphia, PA.

Dr. Hahn has published 14 papers in IEEE TRANSACTIONS or conference proceedings, and has been appointed by the IEEE Educational Activities Board as an ABET Program Accreditor. He is also a member of the IEEE Technology Transfer Committee.