Turning Fiction Into Non-fiction for Signal-to-Noise Ratio Estimation—The Time-Multiplexed and Adaptive Split-Symbol Moments Estimator

M. Simon[1] and S. Dolinar[1]

A means is proposed for realizing the generalized *split-symbol moments estimator (SSME) of signal-to-noise ratio (SNR), i.e., one whose implementation on the average allows for a number of subdivisions (observables), $2L$, per symbol beyond the conventional value of two, with other than an integer value of L. In theory, the generalized SSME was previously shown to yield optimum performance for a given true SNR, R, when $L = R/\sqrt{2}$ and thus, in general, the resulting estimator was referred to as the* fictitious *SSME. Here we present a time-multiplexed version of the SSME that allows it to achieve its optimum value of L as above (to the extent that it can be computed as the average of a sum of integers) at each value of SNR and as such turns fiction into non-fiction. Also proposed is an adaptive algorithm that allows the SSME to rapidly converge to its optimum value of L when in fact one has no a priori information about the true value of SNR.*

I. Introduction

Of the many measures that characterize the performance of a communication receiver, signal-to-noise ratio (SNR) is perhaps the most fundamental in that many of the other measures directly depend on its knowledge for their evaluation. As such, it is desirable that the estimation of SNR take place with as little known information as possible regarding other system parameters such as carrier phase and frequency, type/order of the modulation, data symbol stream, data format, etc. While the maximum-likelihood (ML) approach to the problem will result in the highest quality estimator, typically it results in a structure that is quite complex unless the receiver is provided with some knowledge of the data symbols typically obtained from data estimates made at the receiver (which themselves depend on knowledge of the SNR). An alternative approach is to use an ad hoc estimator that, although being simpler to implement, maintains a high level of performance (as measured by the variance of the estimator in comparison to the Cramer–Rao bound on this measure) and is completely independent of knowledge of the data.

---

[1] Communications Architectures and Research Section.

The research described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

One such ad hoc SNR estimator that has received considerable attention in the past is the so-called split-symbol moments estimator (SSME) [1–6] that, in its conventional implementation, forms its SNR estimation statistic from the sum and difference of information extracted from the first and second halves of each received data symbol. Implicit in this estimation approach, as is also the case for the ML estimators, is that the data rate and symbol timing are known or can be estimated. (With slight modification of the SNR estimation procedure, it is possible to somewhat relax these conditions.) The majority of investigations of the SSME have focused on the performance of the SSME for multiple phase-shift-keying (M-PSK) modulations—most often binary phase-shift-keying (BPSK)—with and without carrier frequency uncertainty. More recently, it has been shown [6] that by increasing the number of subdivisions (observables) per symbol beyond two (but still an even number), then beyond a certain critical value of SNR, one is able to improve the SNR estimator performance. In particular, it has been shown that the variance of the so-modified estimator tracks the Cramer–Rao bound (with a fixed separation from it) on the variance of an SNR estimator over the entire range of SNR values. Implicit in the derivation of the expression for the variance of the SNR estimator given in [6] was the assumption that the even number of subdivisions was the same for all symbols in the observation from which the SNR estimator was formed, and as such an optimum value of the number of subdivisions, denoted by $2L$, was determined for a given true SNR region, the totality of which spans the entire positive real line. Moreover, it was shown that if one ignores the requirement of having the number of subdivisions be an even integer and proceeds to minimize with respect to L the expression for the variance derived as mentioned above, an optimum value of L, denoted by $L_{\bullet}(R)$, can be determined for every value of true SNR, R. Specifically, it was shown in [6] that $L_{\bullet}(R) = R/\sqrt{2}$. Since, for arbitrary SNR, the value of $L_{\bullet}(R)$ is not necessarily integer, the resulting estimator was referred to as the *fictitious* SSME and resulted in a lower bound on the performance of the practical realizable SSME corresponding to integer L.

In this article, we show how one can in practice turn the fictitious SNR estimator into a non-fictitious one. In particular, we demonstrate an implementation of the SSME that allows one to approach the unrestricted optimum value of L (to the extent that it can be computed as the average of a sum of integers) at every true SNR value. More specifically, the proposed approach, herein referred to as the *time-multiplexed* SSME, allows each symbol to possess its own number of subdivisions arranged in any way that, on the average (over all symbols in the observed sequence), achieves the desired optimum value of L. Furthermore, we propose an algorithm for *adaptively* achieving this optimum value of L when in fact one has no a priori information about the true value of SNR. As in the comparable analyses in [5] and [6], we consider the case wherein the symbol pulse shape is assumed to be rectangular and thus the observables from which the estimator is formed are the outputs of matched filters that are integrate-and-dump (I&D) filters. By comparison with the presentation in the above references, the presentation here will intentionally be brief. The reader is referred to [6] for a more detailed explanation and derivation of the various terms involved in the traditional operation of the SSME. The intent here is to focus strictly on the modification of the scheme to adaptively estimate SNR.

II. Signal Model and Formation of the Estimator

A block diagram of the complex baseband SSME is given in Fig. 1. Corresponding to an observation of N symbols, the complex baseband received signal received in the kth interval $(k-1)T \leq t \leq kT$ is given by

$$y(t) = md_k e^{j(\omega t + \phi)} + n(t), \quad (k-1)T \leq t \leq kT \tag{1}$$

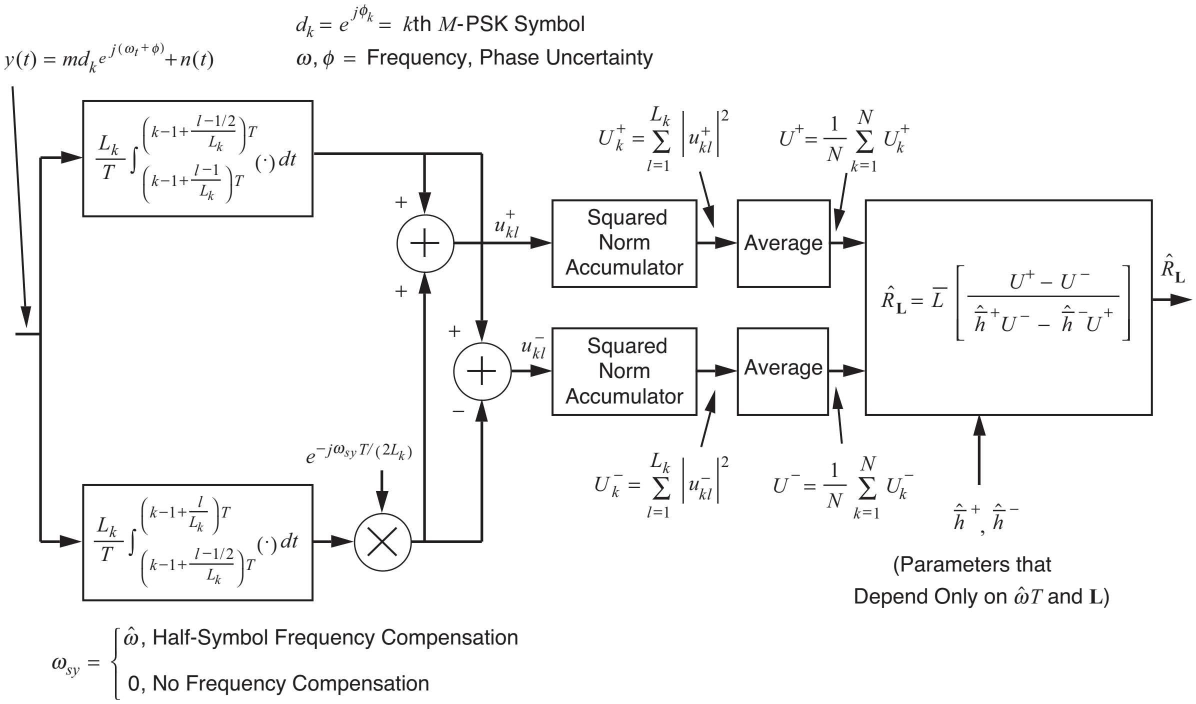

Fig. 1. Time-multiplexed split-symbol moments estimator of SNR for *M*-PSK modulation.

where $d_k = e^{j\phi_k}, k = 1, 2, \cdots, N$ is the kth transmitted M-PSK symbol, m is an amplitude parameter for this symbol, $n(t)$ is a complex zero-mean additive white Gaussian noise (AWGN) process, ϕ and ω are the carrier phase and frequency uncertainties (offsets), and T is the symbol time. Consider uniformly subdividing the kth symbol interval into $2L_k$ (L_k integer) subdivisions each of length $T_k/2 = (T/L_k)/2$. In each of these L_k pairs of split-symbol intervals, we apply the signal in Eq. (1) to first- and second-half normalized (by the integration interval) I&Ds, the outputs of which are summed and differenced to form the signals $\left\{u_{kl}^{\pm} \triangleq s_{kl}^{\pm} + n_{kl}^{\pm}, \, l = 1, \cdots, L_k\right\}$, where $s_{kl}^{\pm}$ and $n_{kl}^{\pm}$, respectively, denote their signal and noise components. For each k, the $u_{kl}^{\pm}$'s are independent and identically distributed (i.i.d.); however, their statistics vary from symbol to symbol. Denote the relevant symbol-dependent parameters of the signal and noise components of $u_{kl}^{\pm}$ as $m_k^2, \sigma_k^2, h_k^{\pm}$, and the SNR in the kth symbol as $R_k = m_k^2/\left(2\sigma_k^2\right)$. In particular, $\sigma_k^2 = \sigma^2 L_k$ is the variance per component (real and imaginary)[2] of $n_{kl}^{\pm}$, and analogous to the analysis in [6], the mean-squared value of $s_{kl}^{\pm}$ can be expressed as

$$\left|s_{kl}^{\pm}\right|^2 = m_k^2 h_k^{\pm} \tag{2}$$

where, because of the normalization of the I&Ds, $m_k^2 = m^2$ independent of k, and $h_k^{\pm}$ is a parameter that reflects the amount of frequency offset and the degree to which it is compensated for. Specifically,

$$h_k^{\pm} = \operatorname{sinc}^2\left(\frac{\delta_k}{4}\right)\frac{1 \pm \cos\left(\delta_{k_{sy}}/2\right)}{2} \tag{3}$$

[2] Note that σ^2 is the variance per component of the $u_k^{\pm}$ samples in the conventional SSME corresponding to $L = 1$ in each symbol interval.

where sinc $x \triangleq \sin x/x$, $\delta_k \triangleq \omega T_k$, and $\delta_{k_{sy}} \triangleq \delta_k - \omega_{sy} T_k = (\omega - \omega_{sy}) T_k$ with ω_{sy} the compensation frequency applied to the second-half I&D outputs. The various options for compensating the phase variations across a given symbol due to the frequency offset, i.e., the choices of ω_{sy}, are discussed in [5] and are not repeated here. For example, in the case of no frequency error ($\omega = 0$) and thus no compensation ($\omega_{sy} = 0$), we would have $h_k^+ = 1$, $h_k^- = 0$.

Based on the above, each $\left|u_{kl}^{\pm}\right|^2 = \sigma_k^2 \chi_2^2 \left(2h_k^{\pm} R_k\right)$, where $\chi_n^2(\mu)$ denotes a (generally non-central) chi-squared random variable (RV) with n degrees of freedom, non-centrality parameter μ, and unit variances for each degree of freedom. In general, we know that $E\left\{\chi_n^2(\mu)\right\} = n + \mu$ and var $\left\{\chi_n^2(\mu)\right\} = 2n + 4\mu$ for all n and μ. Furthermore, using [7, Eq. (2.39)] for the inverse moments of central chi-squared RVs, we have $E\left\{[\chi_n^2(0)]^{-1}\right\} = (n-2)^{-1}$ and $E\left\{[\chi_n^2(0)]^{-2}\right\} = [(n-2)(n-4)]^{-1}$ for even n and $\mu = 0$. Expressions for higher-order moments of $\chi_n^2(\mu)$ or its reciprocal can be determined using [7, Eq. (2.47)].

Now, for each k, define $U_k^{\pm} = \sum_{l=1}^{L_k} \left|u_{kl}^{\pm}\right|^2 / L_k$. Then, based on the above chi-squared characterization of $\left|u_{kl}^{\pm}\right|^2$, and recognizing that

$$R = R_k L_k = \frac{m^2}{2\sigma^2} \tag{4}$$

is the true SNR to be estimated, we have $U_k^{\pm} = \left(\sigma_k^2/L_k\right) \chi_{2L_k}^2 \left(2h_k^{\pm} R_k L_k\right) = \sigma^2 \chi_{2L_k}^2 \left(2h_k^{\pm} R\right)$ with first mean and variance

$$\left.\begin{aligned} E\left\{U_k^{\pm}\right\} &= 2\sigma^2\left(L_k + h_k^{\pm} R\right) \\ \text{var}\left\{U_k^{\pm}\right\} &= 4\sigma^4\left(L_k + 2h_k^{\pm} R\right) \end{aligned}\right\} \tag{5}$$

Solving for R in terms of $E\left\{U_k^{\pm}\right\}$ from the first equation in Eq. (5), we obtain

$$R = L_k \left[\frac{E\left\{U_k^+\right\} - E\left\{U_k^-\right\}}{h_k^+ E\left\{U_k^-\right\} - h_k^- E\left\{U_k^+\right\}}\right] \tag{6}$$

At this point, we could proceed as we did in [5] by replacing expected values of $U_k^{\pm}$ with their sample values to obtain estimates of R from each symbol, and then average over the N estimates obtained from the N symbols, resulting in the ad hoc estimator

$$\hat{R}'_{\mathbf{L}} = \frac{1}{N} \sum_{k=1}^{N} L_k \left[\frac{U_k^+ - U_k^-}{h_k^+ U_k^- - h_k^- U_k^+}\right] \tag{7}$$

where $\mathbf{L} = (L_1, L_2, \cdots, L_N)$ denotes the oversampling vector for the N-symbol observation. Unfortunately, this has the potential of being a very bad estimator, because from our previous analyses we have observed that both the bias and the variance of the split-symbol estimate become unbounded if it is based on only a single symbol, i.e., $N = 1$. If $\{L_k\}$ takes on only a few discrete values, we can avoid this singularity by grouping symbols with the same L_k, obtaining an estimate from each group, and then averaging the estimates from all the groups. A better approach is to average the $U_k^{\pm}$ samples prior to forming them into an ad hoc estimator. Specifically, we form $U^{\pm} = (1/N) \sum_{k=1}^{N} U_k^{\pm}$, which has the chi-squared characterization $U^{\pm} = (\sigma^2/N) \chi_{2\bar{L}N}^2 \left(2\bar{h}^{\pm} N R\right)$, where $\bar{L} = (1/N) \sum_{k=1}^{N} L_k$ and $\bar{h}^{\pm} = (1/N) \sum_{k=1}^{N} h_k^{\pm}$. The mean and variance of $U^{\pm}$ are immediately given by

$$\left.\begin{aligned} E\left\{U^{\pm}\right\} &= 2\sigma^2\left(\bar{L}+\bar{h}^{\pm}R\right) \\ \operatorname{var}\left\{U^{\pm}\right\} &= \left(\frac{4\sigma^4}{N}\right)\left(\bar{L}+2\bar{h}^{\pm}R\right) \end{aligned}\right\} \tag{8}$$

Solving for R in terms of $E\left\{U^{\pm}\right\}$, we obtain

$$R = \bar{L}\left[\frac{E\left\{U^{+}\right\}-E\left\{U^{-}\right\}}{\bar{h}^{+}E\left\{U^{-}\right\}-\bar{h}^{-}E\left\{U^{+}\right\}}\right] \tag{9}$$

Now we replace expected values with sample values and $\bar{h}^{\pm}$ with estimates $\hat{\bar{h}}^{\pm}$ based on an estimate $\hat{\omega}$ of the frequency offset ω in this single equation to get our SNR estimate,

$$\hat{R}_{\mathbf{L}} = \bar{L}\left[\frac{U^{+}-U^{-}}{\hat{\bar{h}}^{+}U^{-}-\hat{\bar{h}}^{-}U^{+}}\right] \tag{10}$$

The equations defining both the estimator $\hat{R}_{\mathbf{L}}$ and the underlying observables $U^{\pm}$ in terms of standard chi-squared random variables are identical in form to those obtained for the special case of uniform subsampling of all the symbols, $\mathbf{L}=(L,L,\cdots,L)$. The parameters $\bar{L}$, $\bar{h}^{\pm}$ for the general case reduce to the constants L, $h^{\pm}$ for the special case. The special case $\mathbf{L}=(L,L\cdots,L)$ produces the estimator $\hat{R}_L$ defined in [6], where we assumed constant L for all symbols. Thus, we can apply our previous performance calculations for $\operatorname{var}\left\{\hat{R}_L\right\}$ to obtain the corresponding expressions for $\operatorname{var}\left\{\hat{R}_{\mathbf{L}}\right\}$ by simply replacing L and $h^{\pm}$ in those expressions with $\bar{L}$ and $\bar{h}^{\pm}$, respectively. In the case of zero frequency offset, analyzed in Sections 1 through 5 of [6], we now can achieve the variance expression for any value of $\bar{L}$ achievable by *averaging* integers, not just integer values of L themselves. For large N, this means that we can achieve the performance of our fictitious estimator $\hat{R}_{\bullet}$ for a very dense set of values of R satisfying $L_{\bullet}(R)=R/\sqrt{2}\geq 1$. Of course, the fictitious estimator remains fictitious for $L_{\bullet}(R)=R/\sqrt{2}<1$ (i.e., the region of R where we did not attempt to use it as a benchmark).

III. An Adaptive SSME

Given that $\hat{R}_{\mathbf{L}}$ achieves the performance of $\hat{R}_{\bar{L}}$, we now have a method for adaptively selecting the oversampling factor L. We can start with an initial guess, and then increase or decrease L in response to intermediate SNR estimates $\hat{R}_{\mathbf{L}}$ based on the symbols observed up to now. The key point is that the estimator $\hat{R}_{\mathbf{L}}$ at any point in time achieves exactly the same performance as the estimator $\hat{R}_L$ with $L=\bar{L}$, based on the same cumulative number of symbols. Thus, no symbols are wasted if an adaptive SNR estimation algorithm starts out with a non-optimum value of L but adapts over time to generate a vector sequence $\mathbf{L}$ for which the average $\bar{L}$ approaches the optimum value of L, namely, $L=L_{\bullet}(R)=R/\sqrt{2}$. Figure 2 is a flow diagram of such an adaptive scheme modeled after the robust version of the generalized SSME [6] wherein the integer values of L are restricted to the set $b^l,\ l=0,1,2,3,\cdots$ for some integer base b. The operation of the scheme is described as follows.

Initially, consider an observation of n symbols and set $L_k=L=1,\ k=1,2,\cdots,n$. Next, evaluate the sum and difference accumulated variables $U^{\pm}$ for the n symbol observation. Proceed to evaluate the SNR estimator $\hat{R}=\hat{R}_{\mathbf{L}}(U^{\pm})$ in accordance with Eq. (10), taking note of the fact that, for this choice of $\mathbf{L}$, $\bar{L}=1$. Next, we compare the current value of L, namely $L=1$, to the desired optimum value of L based, however, on the current estimate of R, i.e., $L_{\bullet}\left(\hat{R}\right)=\hat{R}/\sqrt{2}$, to get an indication of how close we are to where we are headed, namely, the determination of the optimum value of L. If $L_{\bullet}\left(\hat{R}\right)$ exceeds

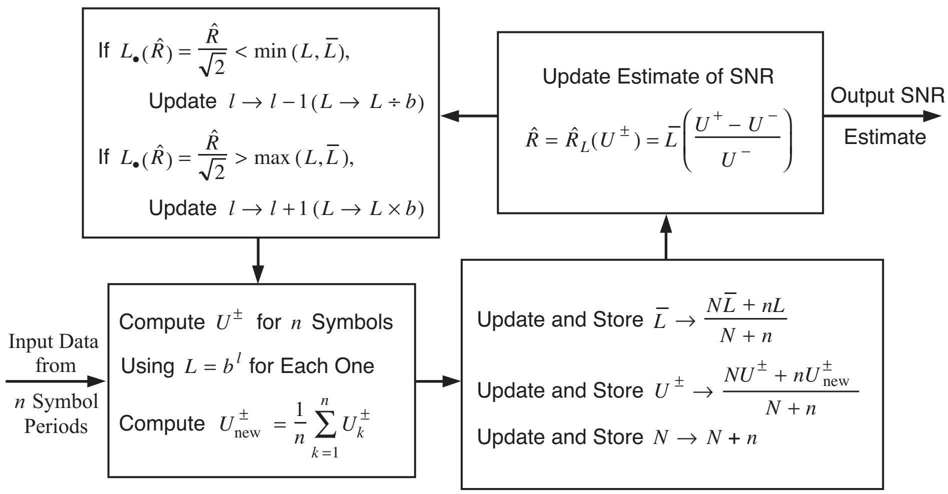

Fig. 2. A robust adaptive SSME scheme.

unity, which on the average is likely to be the case if the true SNR is greater than $\sqrt{2}$, increment L by multiplying it by b and proceed to process the next n input symbols, as will be described momentarily. On the other hand, if $L_\bullet(\hat{R})$ is less than or equal to unity, which on the average is likely to be the case if the true SNR is less than or equal to $\sqrt{2}$, then leave L unchanged[3] and again proceed to process the next n input symbols. Moving on to the next set of n symbols, compute new values of $U^{\pm}$, denoted by $U^{\pm}_{\text{new}}$, using the updated value of L as determined above for all $L_k,\ k = n+1, n+2, \cdots 2n$. Let N denote the running average of the number of symbols. (Assume that N initially was set equal to n, corresponding to the first set of observed symbols.) Update the current values of $U^{\pm}$ with the new $U^{\pm}_{\text{new}}$ values according to the weighted average $\left(NU^{\pm} + nU^{\pm}_{\text{new}}\right)/(N+n)$ and store these as $U^{\pm}$. Update the running average of L in accordance with $\left(N\bar{L} + nL\right)/(N+n)$ and store this as $\bar{L}$. Finally, update the value of N to $N+n$ and store this new value. Using the updated $U^{\pm}$, compute an updated SNR estimate $\hat{R} = \hat{R}_{\mathbf{L}}(U^{\pm})$ in accordance with Eq. (10). Next, using this updated SNR estimate, compute the updated estimate of the optimum L, namely, $L_\bullet(\hat{R}) = \hat{R}/\sqrt{2}$, and use it to update the current value of L in accordance with the following rule:

$$\left.\begin{array}{l}
\text{If } L_\bullet\left(\hat{R}\right) < \min\left(L, \bar{L}\right), \text{then divide } L \text{ by } b \\[2ex]
\text{If } L_\bullet\left(\hat{R}\right) > \max\left(L, \bar{L}\right), \text{then multiply } L \text{ by } b \\[2ex]
\text{If} \min\left(L, \bar{L}\right) \le L_\bullet\left(\hat{R}\right) \le \max\left(L, \bar{L}\right), \ \text{do not change } L
\end{array}\right\} \tag{11}$$

Finally, using the updated value of L, proceed to process the next n symbols, whereupon the algorithm repeats as described above.

To illustrate the behavior of the robust adaptive SSME scheme, simulations were conducted to demonstrate the rate at which $\bar{L}$ converges to the true optimum L and also the manner in which this convergence takes place. The first simulation illustrated in Fig. 3 demonstrates the ideal performance of the scheme

---

[3] As we shall see shortly, in all other circumstances of this nature, we would proceed to decrement L by dividing it by b. However, since the current value of L is already equal to unity, which is the smallest nonzero integer, we cannot reduce it any further.

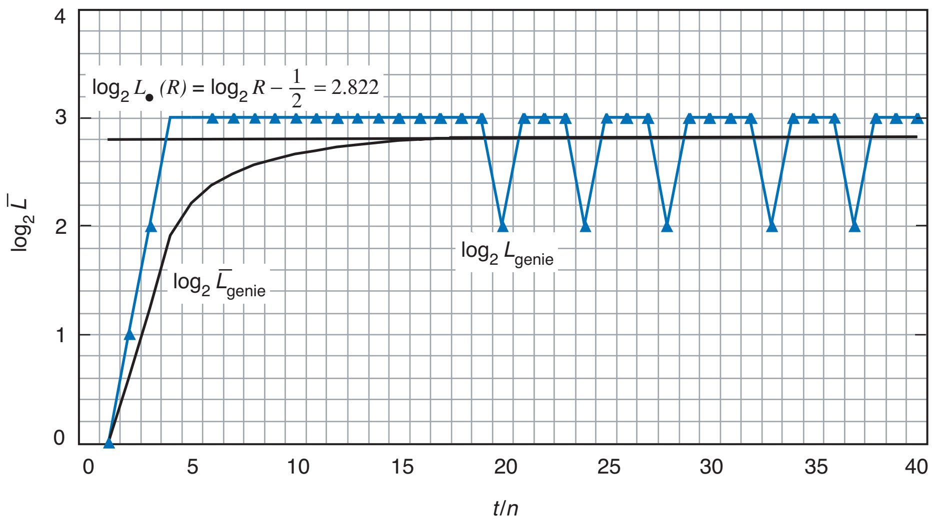

Fig. 3. Ideal performance of the robust adaptive SSME scheme. (Adaptive SSME with magic genie estimate of true SNR, $R = 10$.)

assuming no frequency error, i.e., $\hat{\hat{h}}^+ = 1, \hat{\hat{h}}^- = 0$, and the following parameters: $R = 10, b = 2$, and $n = 10$. By "ideal" is meant that the same adaptive feedback rule for updating L as in Eq. (11) is used except that a magic genie is assumed to be available to provide the true SNR, R, to the update rule rather than using the estimate of R. That is, the update of L in accordance with Eq. (11) is carried out using $L_\bullet(R)$ rather than $L_\bullet(\hat{R})$. The horizontal axis in Fig. 3 is measured in discrete units of time corresponding to the cumulative number of n-symbol batches processed, each with a fixed value of L. The vertical axis represents two different indicators of the performance, corresponding to the behavior of $\log_2 L$ and $\log_2 \bar{L}$ as they are updated in each cycle through the feedback loop. For the assumed parameters, the optimum value of L to which the scheme should adapt is given in logarithmic terms by $\log_2 L_\bullet(10) = \log_2 \left(10/\sqrt{2}\right) = 2.822$. From the illustrations in Fig. 3, we observe that $\log_2 L$ quickly rises (in three steps) from its initial value of $\log_2 1 = 0$ to $\log_2 8 = 3$ and then eventually fluctuates between $\log_2 8 = 3$ and $\log_2 4 = 2$ with a 3:1 or 4:1 duty cycle. At the same time, $\log_2 \bar{L}$ smoothly rises toward the optimum $\log_2 \bar{L}_\bullet$, converging asymptotically to this limit (with indistinguishable difference) in fewer than 20 cycles of the feedback loop or, equivalently, 200 symbol intervals.

Figure 4 is an illustration of the actual performance of the scheme as illustrated in Fig. 2, i.e., in the absence of a magic genie to provide the true SNR. The same parameter values as in Fig. 3 were assumed, and 10 different trials were conducted. Also superimposed on this figure for the purpose of comparison is the $\log_2 \bar{L}$ magic genie performance obtained from Fig. 3. For 6 out of the 10 trials, the actual performance was indistinguishable from that corresponding to the magic genie. For the remaining 4 trials, $\log_2 \bar{L}$ overshoots its target optimum value but still settles toward this value within 20 cycles of the algorithm. For all 10 trials, there is a small dispersion from the optimum level even after 40 cycles. This is due to residual error in estimating R after $N = 400$ symbols since the variance only decreases as $1/N$.

IV. Conclusions

It is possible to modify the implementation of the SSME of SNR to allow one to approach the unrestricted (not requiring it to be integer) optimum value of the number of subdivisions, L, of the symbol interval (to the extent that it can be computed as the average of a sum of integers) at every true SNR value. The proposed approach allows each symbol to possess its own number of subdivisions arranged in any way that, on the average (over all symbols in the observed sequence), achieves the desired optimum value of L. Furthermore, it has been shown that there exists a rapidly converging algorithm for *adaptively* achieving this optimum value of L when in fact one has no a priori information about the true value of SNR.

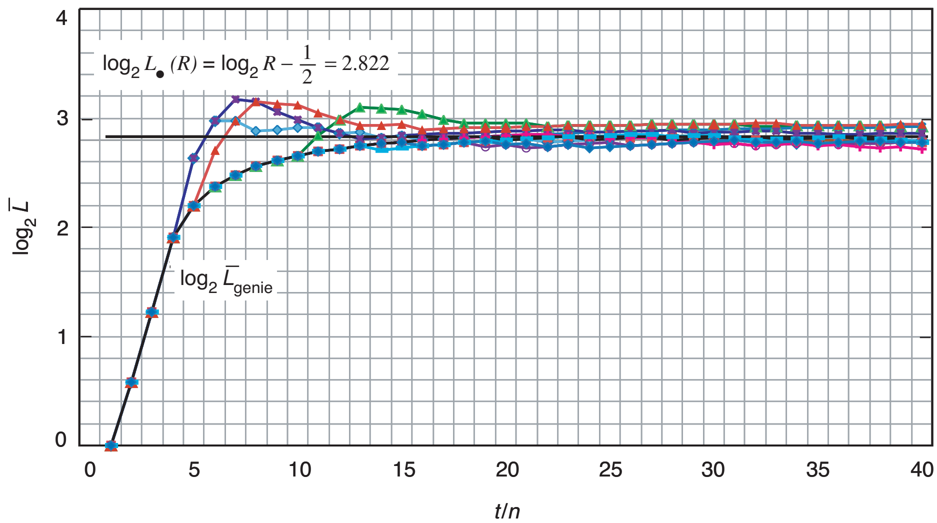

Fig. 4. Actual performance of the robust adaptive SSME scheme. (Adaptive SSME: 10 trials with $N = 400$, $n = 10$, true $R = 10$.**)**

References

[1] M. K. Simon and A. Mileant, "SNR Estimation for the Baseband Assembly," *The Telecommunications and Data Acquisition Progress Report 42-85, January–March 1986*, Jet Propulsion Laboratory, Pasadena, California, pp. 118–126, May 15, 1986. http://tmo.jpl.nasa.gov/tmo/progress_report/42-85/85K.PDF

[2] B. Shah and S. Hinedi, "The Split Symbol Moments SNR Estimator in Narrow-Band Channels," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-26, pp. 737–747, September 1990.

[3] S. Dolinar, "Exact Closed-Form Expressions for the Performance of the Split-Symbol Moments Estimator of Signal-to-Noise Ratio," *The Telecommunications and Data Acquisition Progress Report 42-100, October–December 1989*, Jet Propulsion Laboratory, Pasadena, California, pp. 174–179, February 15, 1990. http://tmo.jpl.nasa.gov/tmo/progress_report/42-100/100N.PDF

[4] Y. Feria, "A Complex Symbol Signal-to-Noise Ratio Estimator and Its Performance," *The Telecommunications and Data Acquisition Progress Report 42-116, October–December 1993*, Jet Propulsion Laboratory, Pasadena, California, pp. 232–245, February 15, 1994. http://tmo.jpl.nasa.gov/tmo/progress_report/42-116/116s.pdf

[5] M. K. Simon and S. Dolinar, "Signal-to-Noise Ratio Estimation for Autonomous Receiver Operation," *GLOBECOM 2004 Conference Record*, Dallas, Texas, November 29–December 3, 2004.

[6] M. Simon and S. Dolinar, "Improving Signal-to-Noise Ratio Estimation for Autonomous Receivers," *The Interplanetary Network Progress Report*, vol. 42-159, Jet Propulsion Laboratory, Pasadena, California, pp. 1–19, November 15, 2004. http://ipnpr.jpl.nasa.gov/tmo/progress_report/42-159/159D.pdf

[7] M. K. Simon, *Probability Distributions Involving Gaussian Random Variables: A Handbook for Engineers and Scientists*, Norwell, Massachusetts: Kluwer Academic Publishers, 2002.