Multicanonical Monte Carlo Study of the BER of an All-Optically 2R Regenerated Signal

Taras I. Lakoba

(Invited Paper)

***Abstract*—Using a numerical procedure based on the multicanonical Monte Carlo (MMC) algorithm, we compute the bit error rate (BER) at the output of a Mamyshev-type all-optical regenerator. For the specific device considered, this BER is degraded as compared to its value at the input to the regenerator, even though the Q-factor is improved. In absolute terms, the output BER decreases when the bandwidth of the optical filter placed before the regenerator to limit the amount of entering noise is decreased. We also present evidence that the degradation of the BER caused by the regenerator is due to high sensitivity of the output signal to small variations of the input signal shape. To our knowledge, this fact has not been pointed out in earlier studies of regenerators. In addition to the previous results, the modification of the MMC procedure that we proposed and used in this paper can be, in a certain sense, parallelized. When it is practical to do so, our procedure provides significant saving of the computational time over the standard MMC simulation.**

***Index Terms*—All-optical regeneration, Monte Carlo methods, nonlinear wave propagation, optical fiber communication, optical noise, optical signal detection, optical signal processing.**

I. Introduction

ALL-OPTICAL regeneration is being actively researched as it can increase the reach of fiber-optical transmission systems without expensive optical-to-electrical signal conversion. In one of the most studied setups, a chain of periodically placed regenerators would suppress the accumulation of amplified spontaneous emission (ASE) noise and/or other signal distortions. This occurs because each regenerator "squeezes" the probability density function (PDF) of both the ONE and ZERO levels of the incoming signal; then the noise accumulated between any two consecutive regenerators is added to a "cleaner" signal than it would be in the absence of the preceding regenerators. As a result, the "periodically cleaned-up" signal at the end of the transmission line with a chain of regenerators is degraded by the ASE less than a signal in a line without regenerators.

In most of the studies of signal quality improvement by a chain of regenerators (see, e.g., [1]–[4] and references therein), the effect of the regenerator is modeled by its *static* power transfer function. This means that the signal degraded by the ASE is modeled as having the shape and timing of an undistorted ONE, but with an appropriately scaled amplitude. An indication that this assumption may not be correct appeared recently in [5] and [6] in the context of semiconductor and all-optical 2R (reamplifying and reshaping) regenerators, respectively. Specifically, in [6], it was experimentally demonstrated that under certain conditions, a Mamyshev-type regenerator [7] can improve the bit error rate (BER) of a signal as compared to that measured immediately before the regenerator. If modeling the effect of the regenerator by its static transfer function were correct, such a BER improvement would not be possible since a device transforming the signal according to any, but the same for both regenerated ONEs compased to the input and ZEROs, static transfer function does not change the BER [8]. Therefore, modeling a regenerator's action by its static transfer function can be incorrect, and thus, an alternative treatment is needed.

In this paper, we propose a numerical procedure that allows one to compute the change of the BER at the regenerator's output as compared to that at the input. Our procedure uses the multicanonical Monte Carlo (MMC) algorithm [9], [10], [11], and hence, allows one to access BERs on the order of 10^{-10} and lower. For an application of this procedure, we focus exclusively on the Mamyshev-type regenerator. We observed the BER degradation, not improvement, as reported in [6], at the output of a single regenerator. However, a direct comparison between our results and those of [6] cannot be made since we had to use a setup different from that of [6], as we explain later. We also analyze mechanisms causing the BER degradation.

The main idea of *why* a regenerator can change the BER was stated in [8]: the static transfer functions for an average noisy ONE and an average noisy ZERO can be different. It was further clarified in [6] that a Mamyshev-type regenerator discriminates between noisy ONEs and ZEROs by their shape. Specifically, the authors of [6] suggested that at the input to the regenerator, these two types of symbols have, on average, different temporal widths. In what follows, we will examine this statement and confirm that its general idea, i.e., the fact that the BER can change because the noisy ONEs and ZEROs may have different width or other parameters, is correct. However, we will also show that, in general, the average noisy ZERO is not always wider than the average noisy ONE, as stated in [6]. More importantly, we will demonstrate that the observed difference in input pulsewidths does not fully account for the observed changes in the BER at the regenerator's output.

An open question, which is not considered in [6], is: how does the property of the regenerator to change the BER depend on the parameters of the regenerator? An attempt to answer that question was made in [12]. There, it was shown that the regenerator improves the Q-factor of the signal only if the signal

Manuscript received November 5, 2007; revised January 14, 2008. This work was supported by the National Science Foundation (NSF) under Grant DMS-0507429.

The author is with the Department of Mathematics and Statistics, University of Vermont, Burlington, VT 05401 USA (e-mail: lakobati@cems.uvm.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSTQE.2008.919279

is filtered by a bandpass filter (to reduce the amount of the ASE) before it enters the regenerator. Moreover, such an improvement in Q was observed only when the filter bandwidth had a value in a certain range; for other values of the filter bandwidth, the Q-factor was degraded by the regenerator. While these results suggest that the filter in front of the regenerator may change the BER at its output, they do not predict whether such a change will actually be an improvement. The well-known reason for this is that since the regenerator is a strongly nonlinear device, it significantly modifies the Gaussian statistics of the ASE, and hence, makes the Q-factor no longer related to the BER of the signal. We will show later that while improving the Q-factor, the regenerator can actually degrade the BER.

The main body of this paper is organized as follows. In Section II, we describe the numerical procedure that we developed to study signal transmission with low values of BER through a regenerator. This procedure uses the MMC algorithm, which allows one to access values of BER on the order of 10^{-10} and below with only $O(10^5)$ simulations [9], [10], [11]. However, in the context of our problem, this number of simulations is still time-prohibitive. We will describe a modification of the original algorithm that allows one to execute the simulations in parallel, and hence, makes them feasible on a node cluster. In Section III, we present our results on how the regenerator changes the BER and the Q-factor of the input signal for a system similar to that considered in [12]. In Section IV, we analyze mechanisms that contribute to the change of the BER by a regenerator. The main results of the paper are summarized in Section V. In Appendix A, we review the details of the standard MMC algorithm. In Appendix B, we discuss the accuracy of results obtained with our procedure.

It may be noted that the MMC algorithm was used in [4] to study the BER improvement by a chain of 2R regenerators distributed over a long transmission line. However, in that paper, the regenerators were modeled by their static transfer functions, and hence, the problem of the BER change by a single regenerator, which we study here, could not be addressed by the approach of [4]. Furthermore, modeling regenerators by their static transfer functions effectively replaced the solution of a partial differential equation that governs signal transmission by a small number of algebraic operations on the signal power. Therefore, the authors of [4] also did not need to be concerned with excessive computational time of their simulations.

II. Description of Numerical Procedure

The exposition in this section can be applied to any type of a regenerator or, more generally, any signal-transforming device or system, as long as it does not add noise to the input noisy signal. The parameters of the specific regenerator that we considered will be given in the following section.

A. Main Idea

We begin by explaining the main technical challenge of studying the BER of a regenerated signal by the MMC. A complete MMC simulation typically requires on the order of 10^5 individual runs [10], [11]. Each run through the regenerator takes about 5 seconds (for a relatively short bit sequence in a single channel). By the nature of the MMC algorithm, these runs are to be performed in series, i.e., cannot be parallelized. Thus, a direct MMC simulation would take somewhere between 3 and 5 days. Such a long time makes this conventional MMC-based approach hardly feasible if one needs to consider many sets of parameters (and also to allow for unavoidable mistakes in the exploratory phase of the research).

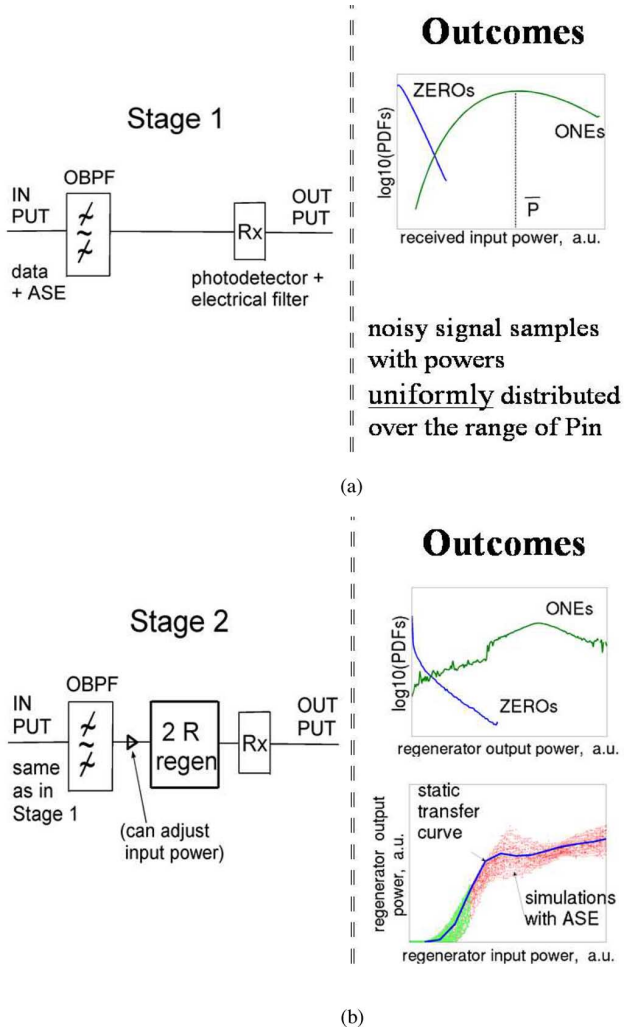

Fig. 1. Schematics of stages 1 and 2 of the proposed numerical procedure.

The key idea that we proposed and implemented to obtain the results reported later is to split the simulation into two stages. In the first stage, we run an MMC simulation to compute the BER of the original (i.e., nonregenerated) signal consisting of data and the ASE. Since one run through the receiver takes only a small fraction of a second, this entire stage typically takes no more than a couple of hours (on a Pentium IV-type machine), and can be performed independently for a ONE and a ZERO. The outcome of this stage contains two important parts, as schematically shown in Fig. 1(a). First, as usual, one obtains the PDFs of the received noisy ONE and ZERO. Those PDFs may extend to 10^{-10} and even lower. Second, for each small interval of values of the received signal power, $[P_{\rm in}, P_{\rm in} + \Delta P_{\rm in}]$, we collect a large number of *distinct* samples of noisy signal whose received

powers fall into this interval.[1] In what follows, we refer to the interval $[P, P+\Delta P]$ as a *bin*. By the nature of the MMC algorithm [13], one can obtain approximately equal number of such samples in all bins, even in those where the PDF is very low.

Then, in the second stage of our procedure, we send these noisy samples through the regenerator. As a result, for each ith bin $[P_{\mathrm{in},\,i},\ P_{\mathrm{in},\,i}+\Delta P_{\mathrm{in}}]$ containing *input* signals into the regenerator, we obtain a distribution, $D_i(P_{\mathrm{out}})$, of the *output* power values. From this distribution and the PDF of the input signal, $\mathrm{PDF}_{\mathrm{in}}$, we obtain the PDF at the output of the regenerator

$$\mathrm{PDF}_{\mathrm{out}}(P_{\mathrm{out},\,j}) = \sum_i \mathrm{PDF}_{\mathrm{in}}(P_{\mathrm{in},\,i})\ D_i(P_{\mathrm{out},\,j})\frac{\Delta P_{\mathrm{in}}}{\Delta P_{\mathrm{out}}}. \quad (1)$$

Here and later, indices i and j label input and output bins, respectively, and ΔP_{in} and ΔP_{out} are the respective bin widths. The last factor in (1) assures the normalization

$$\sum_j \mathrm{PDF}_{\mathrm{out}}(P_{\mathrm{out},\,j})\Delta P_{\mathrm{out}} = 1$$

where we have also assumed that the distributions $D_i(P_{\mathrm{out}})$ are normalized so that

$$\sum_j D_i(P_{\mathrm{out},\,j}) = 1 \qquad \text{for all } i. \quad (2)$$

The schematics of the second stage are shown in Fig. 1(b). What makes it possible to reduce the computational time of this stage of the procedure as compared to the direct MMC simulation is that now distinct signal samples can be processed (i.e., sent through the regenerator) *independently from one another*. Thus, this stage can be parallelized. The resulting gain in the computational time depends, obviously, on the number of available nodes in the computing cluster.

For the particular set of parameters considered in this work, the total time required to send all the input signal samples (for either ONE or ZERO) through the regenerator takes about 80 hours, i.e., four full days. When split among several nodes, this stage of the procedure can be easily turned into an overnight simulation.

Let us now comment on the connection between the original MMC algorithm [13] and our modification of it described above. If applied to the problem of obtaining the BER of a regenerated signal, the original MMC algorithm would obtain samples of noisy signal distributed approximately uniformly among all bins $[P_{\mathrm{out},\,j}, P_{\mathrm{out},\,j}+\Delta P_{\mathrm{out}}]$ of the powers at the *output* of the regenerator. (Not all of those samples would be distinct.) The algorithm would then proceed to obtain the PDF of the output signal, as described in [9] and [11]. In our procedure, we replaced the step of collecting (a large number of) samples with uniformly distributed output power by that of collecting samples with uniformly distributed *input* power. Then, in our case, the samples are not uniformly distributed among bins $[P_{\mathrm{out},\,j}, P_{\mathrm{out},\,j}+\Delta P_{\mathrm{out}}]$. However, the $\mathrm{PDF}_{\mathrm{out}}$ can still be computed as long as the number of samples in each of the latter bins remains large. This is indeed the case, as can be seen from the scatter plots in Fig. 6 below. In Appendix B, we will discuss the accuracy of our procedure in greater detail.

[1]The reason for the subscript "in" is explained in the next paragraph.

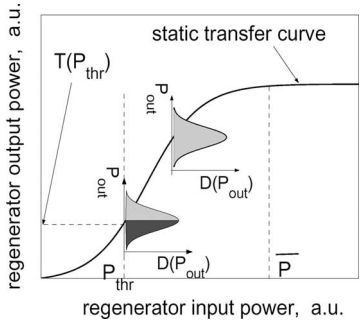

Fig. 2. Schematics of the expected outcome of the proposed numerical procedure.

Finally, before we turn to a more detailed description of out procedure, let us point out that it has an *additional advantage* over a direct MMC computation of the BER of a regenerated signal. Indeed, our procedure shows how the static transfer function of a regenerator is modified for the signal degraded by the ASE. This was one of the starting points of this work (see Introduction), and will be considered in more detail in the following sections. On the contrary, a direct MMC simulation of regenerated signals would not be able to yield such information.

B. Implementation Details

Here, we describe technical details of the numerical procedure presented earlier in this section. The reader who is not interested in such details may skip to Section III.

In stage 1 of the procedure, we first perform a *deterministic* simulation to obtain a static power transfer curve $P_{\mathrm{out}} = T(P_{\mathrm{in}})$ of the regenerator. (This takes only a couple of minutes.) Then, we select the input power of the undistorted ONE to have the corresponding P_{out} somewhere on the curve's plateau. Let us denote this input power $\overline{P_{\mathrm{in}}}$. This is schematically shown in Fig. 2.

Continuing with stage 1, we execute another code, which does not involve the regenerator. This code sends a signal, which consists of an undistroted symbol (either ONE or ZERO) and the ASE, through an optical bandpass filter (OBPF) followed by a receiver; the latter consists of a photodiode and an electrical low-pass filter [see Fig. 1(a)]. We specify a value of the optical signal-to-noise ratio (OSNR) before the OBPF, which along with $\overline{P_{\mathrm{in}}}$ then determines the *average* power density (per frequency interval) p_{ASE} of the ASE. For this p_{ASE}, we perform two separate, standard MMC simulations, one for a ONE and one for a ZERO, obtaining the PDF of the received powers of the respective symbol degraded by the ASE. In each of these two simulations, out of the entire bit sequence, we specify a bit with the respective symbol (ONE or ZERO) and record the power only at the midpoint of that bit.[2] A detailed account of the MMC algorithm can be found in [11] and [14]; we summarize

[2]In actuality, the detection instance is determined by the electronic circuit of the receiver to be an optimum for many consecutive bits simultaneously. Since most bits contain only a small amount of ASE, the optimal detection instance for them is around the midpoint of the bit, which justifies our choice of the detection instance.

its main steps in Appendix A, where we also emphasize certain aspects of it that are critical to our procedure. One of the advantages of the MMC simulation over a regular Monte Carlo simulation is that it generates a nearly uniform distribution of values $P_{\rm in}$ over a specified range. That is, for each bin $[P_{\rm in},\ P_{\rm in}+\Delta P_{\rm in}]$, one has approximately the same number of signal samples whose received powers fall within that interval. We save an equal number of *distinct* signal samples in each bin, as we will use them in stage 2 of the procedure. Finally, using the PDFs of ONEs and ZEROs obtained in stage 1, we compute the BER at the regenerator's input in the standard way (see, e.g., [15, Sec. 11-2]).

In stage 2 of our procedure, we load the signal samples collected in stage 1 and send them through the OBPF (the same one as in stage 1), and then, through the regenerator.[3] For samples from a particular bin $[P_{\rm in},\ P_{\rm in}+\Delta P_{\rm in}]$, this will yield an interval of the output powers $P_{\rm out}$ centered somewhere around $T(P_{\rm in})$ (recall that $P_{\rm out}=T(P_{\rm in})$ is the static transfer curve of the regenerator). If the input samples were all of the same shape, then the interval of the output powers would be exactly $[T(P_{\rm in}),\ T(P_{\rm in}+\Delta P_{\rm in})]$. However, since the ASE causes the shapes of individual signal samples to be different, one obtains a broader distribution $D(P_{\rm out})$ of the powers of the regenerated signal, as schematically shown in Fig. 2 by the bell-like curves. [This is also illustrated by the lower-right subplot in Fig. 1(b).] Using these distributions and the PDFs of the input powers obtained at stage 1, one computes the PDFs at the output of the regenerator by (1). With those PDFs, the BER of the regenerated signal is computed.

Let us now use the schematics shown in Fig. 2 to explain how the regenerator can change the BER. Let $P_{\rm thr}$ be the threshold power that delineated between ONEs and ZEROs in the calculation of the BER in stage 1. If the action of the regenerator could be fully accounted for by its static transfer curve, then the threshold in the BER calculation for the regenerated signal would have been $T(P_{\rm thr})$, and the output BER would equal the input BER [8]. However, due to the finite width of the distribution $D(P_{\rm out})$, some of the ONEs that are above the threshold (i.e., do not produce any errors) before the regenerator may go below the threshold after the regenerator (i.e., result in errors), and vice versa. This is schematically shown by the black shading of the bottom half of the lower bell-like curve in Fig. 2. A similar statement applies to ZEROs as well. As a result, the BER at the regenerator's output can change as compared to its input value. What this change will be depends on the distributions $D_{\rm ONE}(P_{\rm out})$ and $D_{\rm ZERO}(P_{\rm out})$ in the vicinity of the threshold.

III. Numerically Obtained BER and Q-Factor Versus OBPF Bandwidth

In this section, we show the results of our numerical procedure applied to the regenerator considered in [12] whose parameters were scaled to correspond to 10 Gb/s operation. For example, all temporal (spectral) measurements listed below are four times greater (smaller) than the respective values used in [12], while the distance and power measurements are the same as in [12]. This scaling was done for purely technical reasons related to the particular code we used. Since the nondimensionalized propagation equation (see below) is unchanged by this scaling [16], [17], our conclusions remain valid for the 40-Gb/s setup considered in [12]. We did not consider the setup of [6], where the BER improvement by a regenerator was reported, because some of its parameters [e.g., the dispersion of the highly nonlinear fiber (HNLF)] were not listed.

A. Simulation Parameters

The following description pertains to the setup shown in Fig. 1. The signal is first sent through a dispersionless OBPF, which has a flat-top, fourth-order super-Gaussian transmission characteristic. The purpose of this OBPF is to limit the amount of ASE entering either the photodetector or the regenerator. We vary this OBPF's bandwidth from one simulation to another, as reported later. In stage 1 of our numerical procedure, the signal filtered by the OBPF goes into the receiver that consists of a photodetector followed by a fourth-order Bessel electrical filter with the bandwidth of 8 GHz. In stage 2 of the procedure, the signal filtered by the OBPF is sent through the regenerator, and then, detected by the same receiver as in stage 1. The regenerator consists of a 2.5-km section of lossless HNLF with $\gamma=8.4\,(\mathrm{W{\cdot}km})^{-1}$ and dispersion $D=-11.2$ ps/nm·km. The propagation of the signal in the HNLF is modeled by the nonlinear Schrödinger equation for one polarization and without any higher-order corrections. The broadened spectrum of the optical signal at the end of the HNLF is filtered by a Gaussian-shaped OBPF-2 (*not* shown in Fig. 1)[4] whose central frequency is offset from the signal's by 70 GHz to discriminate between ONEs and ZEROs [7]. The full-width at half-maximum (FWHM) of this OBPF-2 is 17.6 GHz, which is chosen so as to produce pulses of the same temporal width, 25 ps, as at the input of the first OBPF, shown in Fig. 1.

In this work, we simulated a sequence of only four bits, 0100. The extinction ratio in the ZERO bits was 20 dB. The peak power $\overline{P}$ of the undistorted ONE in front of the OBPF shown in Fig. 1 was 1.7 W.[5] Since this power is reduced after filtering by the OBPF, we introduced an option to adjust the peak power after the OBPF to the same value as before the OBPF. We will report results both without this power adjustment—since it was not used in [12] (see [18]), and with the adjustment—since it fixes the position of the operation point on the static transfer curve, and thereby, provides for a fairer comparison of the regenerator performance as we vary the OBPF bandwidth. In our simulations, the OSNR of the input signal in 0.1 nm and one polarization was 17 dB. We used this lower value than the authors of [12], since the computational time of stage 1 of our

[3]At this point, one may include an option to adjust the power of all samples by a factor $\max_t P_{0,\,\mathrm{before}}(t)/\max_t P_{0,\,\mathrm{after}}(t)$, where $P_{0,\,\mathrm{before\,(after)}}$ is the power of the undistorted ONE before (after) the OBPF. If such an adjustment is (is not) done, then the quoted input power into the regenerator will pertain to that before (after) the OBPF. This will be used in Section III.

[4]The suffix "-2" is used to distinguish this OBPF from the one in front of the regenerator, which *is* shown in Fig. 1 and plays a prominent role in this study.

[5]This yielded the same path-average power as in [12] when one accounted for the loss of the HNLF in that paper.

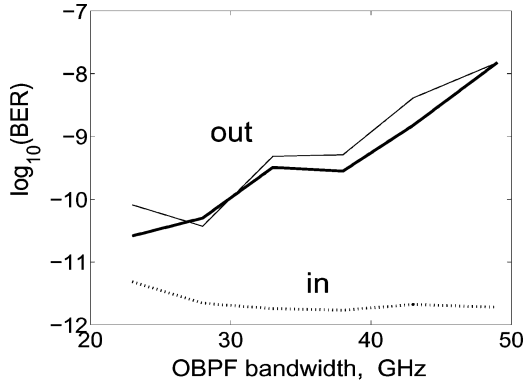

(a)

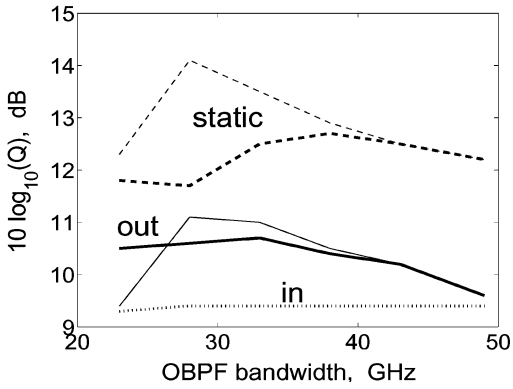

(b)

Fig. 3. BER and Q-factor before and after the regenerator. Dotted line: before; all other line types: after. Thick (thin) lines: the peak power after the OBPF is (is not) adjusted to equal its value before the OBPF. In (b), the dashed lines show the Q-factor obtained from the static transfer curves and the PDFs of the input signal.

procedure increases with the OSNR. We expect that the particular value of that parameter does not change the qualitative conclusions of our study. Having obtained the PDFs of ZEROs and ONEs, as explained in Section II-A, we computed the Q-factor as

$$Q = \frac{V_{(1)} - V_{(0)}}{\sigma_{(1)} + \sigma_{(0)}} \tag{3}$$

where for $k = 0$ or 1, $V_{(k)} = \langle P \rangle_{(k)}$, $\sigma_{(k)} = \sqrt{\langle P^2 \rangle_{(k)} - \langle P \rangle^2_{(k)}}$,

$$\langle P^n \rangle_{(k)} = \sum_i \mathrm{PDF}_{(k)}(P_i) P_i^n (\Delta P)_{(k)}, \qquad n = 1, 2 \tag{4}$$

and i is the index of the bin in the histogram. In the simulations of ONEs and ZEROs, we used 77 and 40 bins covering the power ranges $[0.15\overline{P},\ 1.65\overline{P}]$ and $[0,\ 0.4\overline{P}]$, respectively. The numbers of distinct signal samples collected in each of the ONE and ZERO bins were 800 and 2000. As noted in Section II-B, the power values P_i were detected always at the midpoint of the respective bit slot.

B. Results

Fig. 3(a) and Fig. 3(b) shows the BER and Q-factor measured at the input and the output of the regenerator plotted versus the FWHM of the OBPF in Fig. 1. Since these quantities are determined by statistical simulations, they have a certain error, which we discuss in detail in Appendix B. The results in Fig. 3 are shown both when the input power after the OBPF was and was not adjusted to equal its value before the OBPF. Without the power adjustment, which was the case considered in [12], the filtering of the pulses entering the regenerator leads to the decrease of the power of the undistorted ONE by the corresponding value of the adjustment factor (see Fig. 4). For the 23-GHz-wide OBPF, such decreased power is about 1.0 W, which is near the "knee" of the static transfer curve (see Fig. 6(a)). From the same figure, one sees that the output power fluctuations around the "knee" are significantly greater than those at the flatter part of the transfer curve. This explains the abrupt downturn of the thin solid line in Fig. 3(b).[6] On the other hand, when we adjust the power after the OBPF so as to maintain $\overline{P} = 1.7$ W, the output Q-factor does not decrease appreciably with the OBPF bandwidth. However, that could change in the presence of intersymbol interference, which in our simulations with the bit sequence 0100 did not occur. In what follows, we present and discuss results only for the case where the power after the OBPF was adjusted.

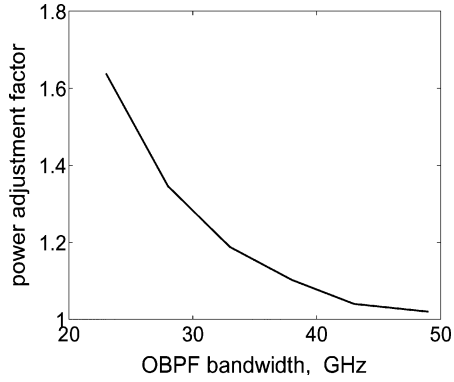

Fig. 4. Adjustment factor that can be applied after the OBPF to restore the pulse peak power; see Fig. 1(b).

The most important conclusion that can be drawn from Fig. 3 is that although the regenerator may improve the Q-factor, it does not improve, and actually degrades, the BER. Another conclusion is that this BER degradation is greater for a wider OBPF. The output Q-factor also decreases with widening the OBPF, as was originally observed in [12].

From Fig. 5, which shows the PDFs of the regenerated ONEs and ZEROs,[7] it is immediately obvious that the BER degradation occurs primarily from the regenerated ONEs. Fig. 6 provides supporting details to that statement, as well as explains why the BER degradation diminishes with the decrease of the OBPF bandwidth. Indeed, expectedly enough, as the spectrum of the ASE allowed into the regenerator becomes narrower, the temporal shape of the noisy signal becomes closer to that of the undistorted ONE, and the deviation of output power from the value dictated by the static transfer curve decreases. Note that the noise increase in regenerated ONEs compared to the input

[6]The reason why the corresponding drop of Q with the OBPF bandwidth is much greater in [12] may be both due to their OSNR being higher, and also, due to the presence of intersymbol interference in their simulations.

[7]The PDFs of *input* ONEs and ZEROs are well known, and hence, shown only in the upper-right subplot in Fig. 1(a).

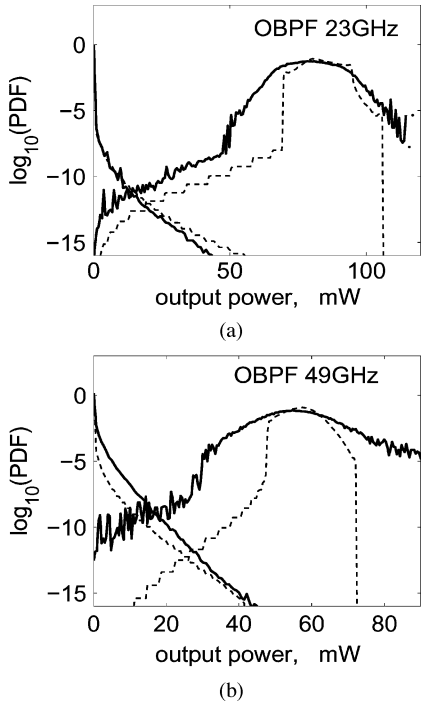

Fig. 5. PDFs of the signal at the output of the regenerator. Solid lines: calculated with the numerical procedure of Section II; Dashed lines: calculated using the static transfer curves and the PDFs of the input signal. The OBPF bandwidth is indicated in the figure. Here and in Figs. 6 through 8, the power is measured by the receiver at the midpoints of bit 2 for ONEs and bit 4 for ZEROs. The actual peak power of an individual pulse from a bin labeled with a given value of power may be slightly greater than that value due to the pulse center being randomly shifted from the bit's midpoint.

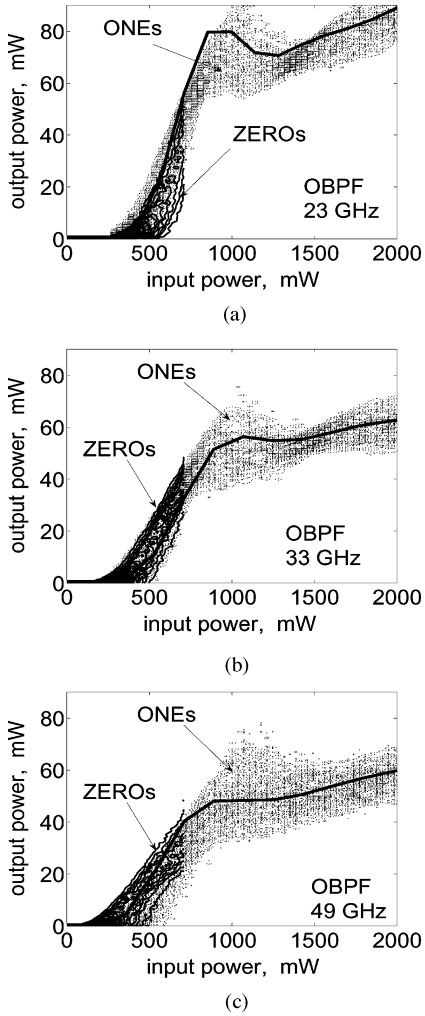

Fig. 6. Contours of the output vs input power scatter plots of the regenerated signal. Thick line: the static transfer curve; thinner contours: ZEROs; dots: ONEs. The OBPF bandwidth is indicated in the figure. See caption of Fig. 5 regarding the axis labels. In addition, the input power is scaled from its value after the receiver's electrical filter to correspond to the power entering the regenerator.

qualitatively agrees with a similar result for supercontinuum generation; see, e.g., the data shown by circles (corresponding to the normal fiber dispersion) in Fig. 2(b) of [19].

Another observation that can be made from Fig. 6 is that the regenerated ZEROs *do*, on average, have different output powers from the regenerated ONEs with the same input power. Specifically, the average output power of ZEROs decreases as compared to the average output power of ONEs as the bandwidth of the OBPF decreases. In other words, the regenerator *can discriminate* between noisy ONEs and ZEROs, as was stated in [6]. However, even when the average output power of ZEROs is less than that of ONEs [as in Fig 6(a)], the BER is still degraded by the regenerator. This occurs because the detrimental effect of the decreased output power of ONEs as compared to the static transfer curve outweighs the positive effect of the decreased output power of ZEROs relative to the same curve.

IV. Causes of Deviation of Regenerated Signal's Power From Static Transfer Curve

The authors of [6] argued that the regenerator can discriminate between noisy ONEs and ZEROs because the latter are, on average, wider (in the temporal domain) than the former. We will show later that although this statement about the relative widths may be correct under certain conditions, its explanation given in [6] was not. We will then show that the previous statement about the sensitivity of the regenerator output to the width of the input pulse is a particular case of a more general phenomenon.

The reason that the explanation of [6] about the widths of ONEs and ZEROs is not correct is the following. The ONEs and ZEROs examined there were taken *not* from the "tails" of the distributions where the PDFs of these symbols intersect and where the decision is made by the receiver, but near the maxima of those distributions, which play no role in the detection. The regular Monte Carlo simulations used in [6] cannot yield a sufficient number of samples near the decision threshold. On the contrary, our procedure based on the MMC algorithm collects sufficient number of samples (see Section III-A) in the entire range of the input power values.

In Fig. 7, we plot the measured widths of pulses after the OBPF and before the regenerator that are averaged over each bin of the histogram. The standard deviations of these widths near the decision threshold (i.e., where the normalized input power is 0.2–0.3) decreases from about 6 to 3 ps as the OBPF bandwidth decreases from 49 to 23 GHz, and are slightly larger for ZEROs

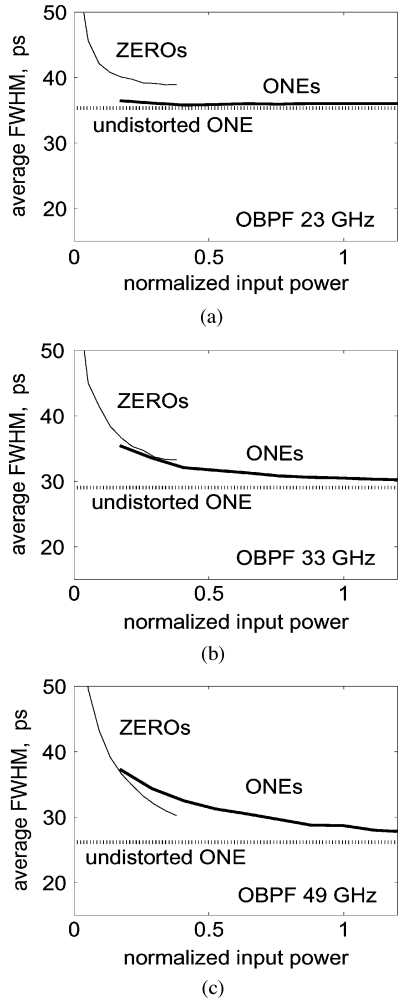

Fig. 7. Per-bin averages of input pulsewidths. We first compute the root-mean-square width of each pulse, and then, multiply it by $2\sqrt{2\ln 2}$ to obtain the FWHM, where the last relation holds for Gaussian pulses. Thick, thin, and dotted lines correspond to noisy ONEs, noisy ZEROs, and the undistorted ONE, respectively. The OBPF bandwidth is indicated in the figure. The input power is normalized to the power $\overline{P} = 1.7$W of the undistorted ONE.

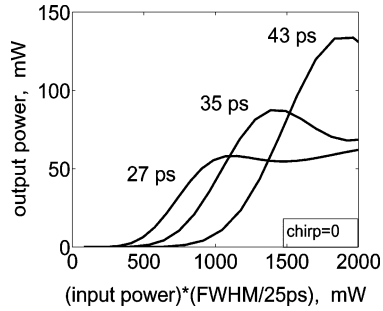

Fig. 8. Static transfer curves corresponding to different widths of unchirped input pulses. Unlike in the previous figures, here, we plot the output power vs a quantity proportional to the *total energy* rather than the peak power (see caption to Fig. 6) of the input pulse. This is done because a given bin contains pulses with the same power at the receiver, and the action of a narrow-band electrical filter at the receiver is similar to that of an integrator.

than for ONEs. A conclusion that follows from Fig. 7 is that the relation between the widths of input ONEs and ZEROs *near the decision threshold* depends on the OBPF bandwidth. Moreover, this relation is correlated with the relative behavior of the average output powers of the symbols that is seen in Fig. 6. Namely, when the average width of input ONEs is less (greater) than the average width of input ZEROs, the average output power of ONEs is greater (less) than the average output power of ZEROs. As was pointed out in [6], this is explained by the dependence of the static transfer curve on the width of the input pulse; this dependence for the regenerator under study is shown in Fig. 8. Similar results were also reported in [17]. However, as pointed out at the end of Section III, the aforementioned ability of the regenerator to discriminate between input ONEs and ZEROs by their widths is insufficient to achieve the BER improvement.

Parenthetically, let us note that for the 33-GHz wide OBPF, the average widths of ZEROs and ONEs are almost the same near the decision threshold. This indicates that in the absence of the regenerator, an OBPF with this bandwidth approximates the matched optical filter (see, e.g., [15, Secs. 5-4 and 14-3]). Indeed, such a filter, whose use in front of the photodetector results in the minimum BER, *cannot discriminate* between ONEs and ZEROs. For if it could, then the filtered spectra of those symbols would have had to be different, and then, using a different OBPF would lower the BER. For example, suppose that after a given OBPF, ZEROs are spectrally narrower than ONEs. Then, sending those signals through a second OBPF that has a "dip" at the central frequency would improve the detection by filtering out more noise than the data. Thus, the combination of these two OBPFs is "better" than the first OBPF alone.

We also examined the standard deviations[8] of the temporal center, central frequency, and chirp of pulses after the OBPF in order to determine their possible effect on the regenerator output. The first two of these quantities can be shown to have negligible effect for all the OBPF bandwidths considered. The standard deviation of chirp near the decision threshold is about 150–200 ps/nm for ZEROs and about two-thirds of that value for ONEs when the OBPF is wider than 40 GHz. Static transfer curves[9] for input pulses with chirp of ± 200 ps/nm predict significantly lower (by a factor of 2 or so) output power near the decision threshold (but not around the plateau of the curve) as compared to the unchirped case. However, when the chirp is reduced to ± 100 ps/nm, the curves for the chirped and unchirped cases predict approximately the same output powers. For the OBPF bandwidth narrower than 35 GHz, the standard deviations of chirp of both ZEROs and ONEs are below 100 ps/nm, and hence, should have only little effect on the power of the regenerated signal.

Next, we visually inspected the correspondence between input noisy pulses and the regenerated ones in order to verify the earlier predictions regarding the effect of pulsewidth and possibly chirp on the output powers. Surprisingly, we found that while there is some correlation between the input pulsewidth and the output power, a stronger effect on the output can often come from relatively small variations of the input pulse shape

[8]The means of these quantities vanish.

[9]We do not depict these curves here; the reader can find qualitatively similar curves in Fig. 7(b) of [17].

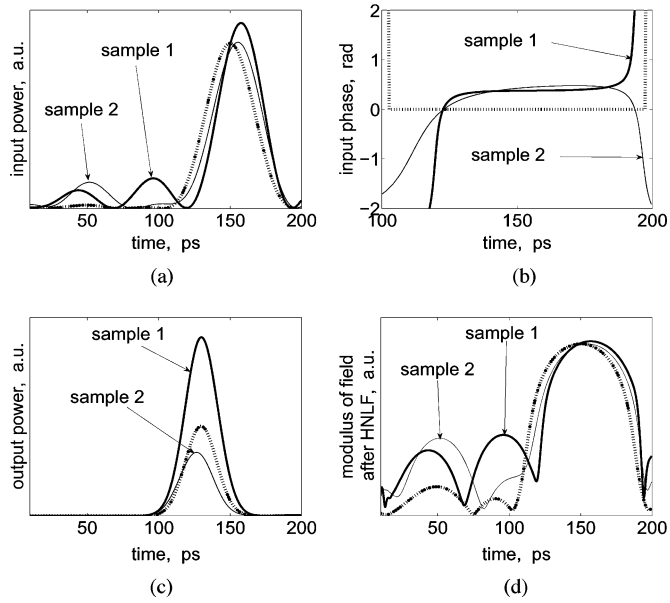

Fig. 9. Typical situation where two close inputs to the regenerator produce greatly different outputs. Thick, thin, and dotted lines correspond to the two sample input pulses and the undistorted ONE with a properly scaled input power. Relevant parameters are found in the text. In (b), we chose to present the phases of pulses in a smaller temporal window than quantities in (a), (c), and (d) in order to emphasize that these phases differ only slightly near the pulse centers. In (d), note that we depicted field amplitudes, not the powers, because it is from those (complex) amplitudes that the spectrum is obtained. The phases of these fields appear to be more similar than the amplitudes, and are not shown.

and even from the field in adjacent bits. That is, variations of input pulse shapes that *are not prominently reflected* by "average" (over the temporal bit) quantities of a pulse, such as its width, can cause as large an effect on the output pulse power as fluctuations of those bit-average quantities. A typical example is shown in Fig. 9, where we depict two samples of an input ZERO in the bit $[100\text{ ps}, 200\text{ ps}]$ with the received power of about $0.25\overline{P} \approx 430$ mW. We chose to present this example for the OBPF bandwidth of 23 GHz to emphasize that the high sensitivity of the regenerator's output with respect to small variations of the shape of the input signal exists even for such a narrow bandwidth. The temporal widths (computed as explained in the caption to Fig. 7) for samples 1 and 2 are, respectively, 38.3 and 36.7 ps, and the chirps are -0.9 and 24.5 ps/nm (such small values of chirp did not affect the output power). Even though the samples' input profiles look quite similar [see Fig. 9(a) and Fig. 9(b)], the output pulses have vastly different powers [see Fig. 9(c)]. Note that the slightly *wider* input sample has the *higher* output power than the other sample, contrary to the dependence characteristic of static transfer curves, shown in Fig. 8. Thus, small variations in the input pulse profile can cause large variations of the output, and, moreover, their effect can outweigh the effect of bit-average quantities (such as width) of the pulse field. High sensitivity of the output power to small variations of the input is well known in the context of supercontinuum generation (see, e.g., [20]). However, supercontinuum generation occurs in a much more nonlinear regime than the one considered here, which is manifested, e.g., by the soliton fission that is observed in [20] and absent in our study. High sensitivity of the output in a setup that is more similar to ours, although being in a different physical context, is reported in [21].

Let us emphasize that the high sensitivity of the power of the regenerated signal is the very property that allows the enhanced discrimination of ZEROs that are *well* below the threshold from the ONEs that are *well* above the threshold. Indeed, this discrimination is provided by the steep part of the static transfer curve, where the output power is highly sensitive to the input one. However, there is no reason to expect that such sensitivity should be restricted only to few parameters (such as peak power and width) of the input pulse and not pertain to other pulse shape variations. This is precisely the phenomenon illustrated by Fig. 9 and discussed in the previous paragraph. Furthermore, consider the flatter regions of static transfer curves in Fig. 6; they correspond to input ONEs relatively slightly contaminated by the ASE. The rather "fat" scatter plots of the output vs input power in those regions suggest that while in the flatter part of the transfer curve, the output power ceases to be highly sensitive to the input power, it apparently *remains* highly sensitive to other small changes of the pulse shape.

V. Conclusion

In this paper, we made the following contributions.

First, we proposed (see Section II) a method by which the MMC simulations can be effectively parallelized, which reduces its computational time. Our method applies not just to a regenerator, but whenever the problem in question has the following setup. Suppose that one needs to measure the PDF of the output of a random quantity X when this quantity has been transformed by some nonlinear system $\mathcal{N}$. The only restriction on $\mathcal{N}$ is that it not add noise to X. Simulating the required number of samples passing through $\mathcal{N}$ would make the standard MMC simulations time-prohibitive (and the regular, i.e., non-multicanonical, Monte Carlo simulations simply not realizable). On the other hand, suppose that one can pass X through another system $\mathcal{M}$ such that: 1) the MMC simulations with $\mathcal{M}$ are relatively quick and 2) there is some "relatively smooth" (yet, of course, unknown) dependence between the outputs of $\mathcal{M}$ and $\mathcal{N}$. Then, the modified MMC-based procedure is run in two stages. In stage 1, perform the standard MMC simulations with $\mathcal{M}$ and collect samples in the range of interest. In stage 2, pass those samples through $\mathcal{N}$ and obtain the desired PDF, as explained earlier in Section II. Processing different samples in stage 2 can be done independently, and hence, this stage can be performed on several computers simultaneously, thereby increasing the computational speed by an order of magnitude. A minor drawback of our procedure as compared to the standard MMC procedure of [9] is that the resulting PDFs in ours are more jagged, and hence, have somewhat larger error, as demonstrated in Appendix B.

Second, using our numerical procedure, we analyzed (in Section III) the BER and the Q-factor of the signal at the output of a Mamyshev-type regenerator as a function of the bandwidth of the OBPF, where the latter limits the amount of noise entering the regenerator. We confirmed the result of [12] that the Q-factor is improved by the regenerator only when the OBPF bandwidth is in a certain range. However, we showed that

even when the Q-factor is improved, the BER is still degraded by the specific regenerator we considered. Thus, although the Q-factor may find some use in predicting the performance of a *chain* of regenerators, it is not relevant to predicting the output BER of a single device. This conclusion, of course, is not unexpected, since the PDF of the signal at the output of the highly nonlinear regenerator is not Gaussian. More importantly, we confirmed a *general observation* made in [12] that the bandwidth of the OBPF in front of the regenerator is one of its important control parameters. Namely, we showed that the BER of the regenerated signal decreases in absolute terms when this OBPF's bandwidth is decreased.

As was mentioned in Section III, we had not had sufficient information to model the setup of [6], where BER improvement by a single regenerator was reported. However, since our results apparently contradict those of [6], we varied parameters of our regenerator in order to find a regime where the BER would improve. The two main alterations we considered were: 1) widening the OBPF-2 at the output of the regenerator (see Section III-A), since this corresponded to the situation considered in [6] and 2) shifting the power $\overline{P}$ from 1.7 to 3.0 W, i.e., to the gently increasing part of the static transfer curve. (The gradual shifting of $\overline{P}$ in the opposite direction was already considered in Section III-B via not adjusting the power after the OBPF.) Neither of these alterations changed our main conclusion that a single regenerator degraded, but not improved, the BER. It may be useful to analyze the setup considered in [6] (which would require the knowledge of all of its relevant parameters) with either our numerical procedure or with the standard MMC simulation, in order to understand what exactly made the BER improvement observed there possible.

The third main contribution of this work was discussed in Section IV. We demonstrated that the output of the regenerator is highly sensitive not only to the width or amplitude of the input pulse, but also to small variations of the input pulse shape that are not captured by those bit-average quantities. To our knowledge, this phenomenon was not pointed out in earlier studies of regenerators. It is that type of sensitivity that largely contributes to the broadness of the output vs input power scatter plots (see Fig. 6). This broadness, in its turn, "raises the tail" of the PDFs, as illustrated in Fig. 5, which leads to the degradation of the BER at the output of a single regenerator.

Parenthetically, let us point out a curious difference between the effects of a receiver (consisting of a photodiode and an electrical filter) and a regenerator on the signal containing optically filtered ASE. The receiver smoothens out shape variations of such a signal, while the regenerator often amplifies them. Stated differently, while the receiver's action is integration, which is known to smoothen jagged input, the regenerator's action could be likened to that of a multidimensional differentiator, which enhances the jaggedness along some of the dimensions.

APPENDIX A

OVERVIEW OF THE MMC ALGORITHM

Here, we mostly follow the exposition of the MMC algorithm found in [11] and [14] and emphasize certain aspects of it that are critical for out case.

Statistical computation of the PDF of a quantity X up to very low values requires generation of random realizations of X that are very rare. To collect sufficiently many such rare realizations so as to obtain a smooth histogram [i.e., the approximation to the PDF(X)] is often not feasible when using the regular Monte Carlo simulations. The MMC is an iterative procedure that dramatically increases the number of such rare realizations by biasing them using the PDF determined at the previous iteration.

To proceed, we need the following notations. Let $\hat{\mathbf{z}} = (\hat{z}_1, \dots \hat{z}_L)$ be the discrete Fourier spectrum of the numerically generated ASE entering the OBPF shown in Fig. 1. The width of the ASE spectrum should only slightly exceed the full bandwidth of the OBPF, as we will explain later. The $\hat{z}_i$'s are independent random quantities having identical PDFs $\rho(z_i)$ with zero mean and the variance σ_z^2 such that for all i

$$\sigma_z^2 \equiv \langle \hat{z}_i^* \hat{z}_i \rangle = p_{\mathrm{ASE}} \Delta\omega \tag{5}$$

where $*$ denotes the complex conjugate, p_{ASE} is the ASE power density, and $\Delta\omega$ is the spacing between the discrete frequencies. The ASE field $\mathbf{z}(t)$ is added to the field of the data $\mathbf{s}(t)$, and then, the power $P(\mathbf{z}) = \int G_{\mathrm{el}}(t-t')|\int G_{\mathrm{opt}}(t'-t'')[\mathbf{s}(t'') + \mathbf{z}(t'')]dt''|^2 dt'$ is detected by the receiver [see Fig. 1(a)], where G_{opt} and G_{el} are the impulse response functions of the OBPF and the electrical filter, respectively. Let us subdivide the entire range $[\min P(\mathbf{z}),\ \max P(\mathbf{z})]$ into small bins $[P_k,\ P_k + \Delta P]$, $k = 1, \dots, K$.

The MMC simulation contains two loops. The outer loop is over subsequent iterations. The inner loop within each MMC iteration is over individual realizations of $\mathbf{z}$. Before the beginning of each iteration, reset the histogram $\mathbf{H} = (H_1, \dots H_K)$ to zero. The kth entry of this histogram at the end of the iteration will then indicate how many times the received power P have been *accepted* (see below) into bin k. For the first sample in each iteration, compute $P(\mathbf{z})$; if that value falls into bin k, accept it into that bin (and hence, let $H_k = 1$). Denote the random vector of the ASE spectrum at this step by $\hat{\mathbf{z}}_a$ and the corresponding ASE field by $\mathbf{z}_a$. For each subsequent realization of $\mathbf{z}$, do the following.

1) Generate a random vector

$$\hat{\mathbf{z}}_b^{\mathrm{prov}} = \hat{\mathbf{z}}_a + \epsilon \Delta\hat{\mathbf{z}} \tag{6}$$

where each component of $\Delta\hat{\mathbf{z}}$ is a symmetrically distributed random variable independent of all the other components. The magnitude ϵ of the ASE increment plays more important role in our procedure than in the original MMC, and therefore, we will comment on it in more detail later. Accept the lth component of $\hat{\mathbf{z}}_b^{\mathrm{prov}}$ with the probability $\min(\rho(\hat{z}_{b,l}^{\mathrm{prov}})/\rho(\hat{z}_{a,l}),\ 1)$. If $\hat{z}_{b,l}^{\mathrm{prov}}$ has been accepted, set $\hat{z}_{b,l} = \hat{z}_{b,l}^{\mathrm{prov}}$; otherwise, set $\hat{z}_{b,l} = \hat{z}_{a,l}$. Note that, as announced earlier, the number L of the frequency components should only slightly exceed a number of frequencies that can pass through the OBPF placed in front of the regenerator. Otherwise, the algorithm may accept a $\hat{\mathbf{z}}_b^{\mathrm{prov}}$ which differs from $\hat{\mathbf{z}}_a$ only in frequencies lying outside the OBPF's passband, while both of these noise realizations

result in essentially the same signal entering into the regenerator. In this work, we took the spectral width $L\Delta\omega$ of $\hat{\mathbf{z}}$ to be 50% greater than the OBPF's FWHM.

2) For the $\mathbf{z}_b$ constructed in step 1), compute the received power $P(\mathbf{z}_b)$; suppose that it falls into bin m. Accept this random realization with the probability $\min(\,\mathrm{PDF}(P(\mathbf{z}_a))/\mathrm{PDF}(P(\mathbf{z}_b)),\,1)$, where PDF (P), an approximation to the true PDF, is computed from the histogram $\mathbf{H}$ at the *previous iteration* (see below). At the first iteration, some initial profile of the PDF needs to be assumed. This profile turns out to be more important in our procedure than in the standard MMC, as we will discuss in detail later. If $P(\mathbf{z}_b)$ is accepted, increment H_m by 1, and then, rename $\mathbf{z}_a = \mathbf{z}_b$ and $k = m$. If $P(\mathbf{z}_b)$ is rejected, increment H_k by 1 and keep $\mathbf{z}_a$ and k the same as in the previous realization.
3) In addition, if $P(\mathbf{z}_b)$ is accepted, save the signal realization $\mathbf{s} + \mathbf{z}_b$ so as to use it in stage 2 of the procedure. *This is the step that is not needed in the standard MMC*.
4) Repeat the previous steps a number of times specified for each iteration. In our simulations, each iteration produced 100 times more realizations than the number of bins. At the end of each iteration, use the accumulated histogram $\mathbf{H}$ to update the PDF (see, e.g., (6) in [11]).
5) Terminate the iterations when the number of collected distinct random realizations of the signal in *all* bins reaches a specified number, $N_{\mathrm{collected}}$. (This termination criterion is different from that used in previous applications of the MMC [10], [11], [14].) To save the disk space, we also recorded only the first $N_{\mathrm{collected}}$ realizations in each bin and discarded all others.

We now discuss the importance of properly choosing the initial PDF, $\mathrm{PDF}_{\mathrm{initial}}$, and the magnitude of the ASE increment, ϵ in (6). Previous researchers who used the MMC started with a constant $\mathrm{PDF}_{\mathrm{initial}}$. In such a case, the entries of histogram $\mathbf{H}$ corresponding to the very rare signal realizations fill out not because many new such realizations are generated, but because the "old" ones are counted repeatedly (see the second step of the description above). This slows down the collection of distinct signal realizations. We found that starting close to the analytic approximation for the PDF [see, e.g., (13) and either (24) or (A3) in [22]] considerably accelerates the collection of distinct realizations.

The effect of the parameter ϵ on the simulation's ability to collect *distinct* signal realizations is more subtle. Namely, if ϵ is "too small," then the successive realizations differ very little from one another, and hence, are not sufficiently independent to produce a good approximation to the BER at the output of the regenerator. On the other hand, if ϵ is "too large," only a small percentage of new realizations will be accepted by the MMC algorithm. Moreover, some of those realizations $\mathbf{z}_b$ may be still rather close to $\mathbf{z}_a$ because, although ϵ is not small, the random quantity $\epsilon\Delta\mathbf{z}$ can still be small. It may be better to use "intermediate" values of ϵ so as to allow the ASE realization to *gradually* drift sufficiently far away from its original state during one iteration of the MMC. We empirically found that using $\epsilon \sim 1$ produces sufficiently distinct signal realizations.

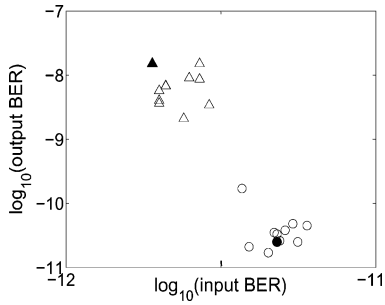

Fig. 10. Output vs input BER values computed by the procedure of Section II with different seeds of the random number generator. The OBPF bandwidths are 23 GHz (circles) and 49 GHz (triangles).The filled symbols are the corresponding data points from Fig. 3(a).

The need to empirically find a "good" combination of $\mathrm{PDF}_{\mathrm{initial}}$ and ϵ, which may vary with the OBPF's and the electrical filter's bandwidths, as well as with the OSNR and possibly other parameters, is currently a minor drawback of our procedure. Employing biased multicanonical sampling [23], [24] may improve the collection of distinct signal realization; this issue is left for future investigation.

Appendix B

Accuracy of Numerical Procedure Described in Section II

It is known [14], [24] that in its "tail" region, the PDF found by the MMC algorithm has a standard deviation that is on the order of the PDF value itself. Consequently, the BER determined by such a PDF can easily fluctuate by a factor of 2 or so from one MMC simulation to another. In addition to this variation that is intrinsic to the MMC algorithm, our procedure introduces jaggedness of the PDF curves, which is especially pronounced in ONEs for low values of the output power [see Fig. 5(b)]. This jaggedness, which should decrease with the increase of the number of collected samples, leads to additional variation in the BER calculated by the procedure of Section II as compared to what would be calculated by the standard MMC algorithm.

Thus, there arise two questions. 1) How large is this additional variation as compared to that intrinsic to the standard MMC? 2) Is the overall variation of the BER calculated by our procedure "acceptable"? We will answer question (1) in the next paragraph. An answer to question (2) depends, obviously, on the goal of a particular study. Our main goal was to compare the BERs at the input and output of a regenerator. As we demonstrate in Fig. 10, variations in the calculated BER do not compromise the qualitative trend observed in Fig. 3(a), and thus, the results of our procedure are indeed acceptable for our purposes. Moreover, it should be noted that the BER measured in a physical experiment is also a statistical quantity and fluctuates in time by a factor of order 2. Therefore, having similar variations in a numerical experiment do not compromise the relevance of the numerical results to the physical ones that they model.

To quantitatively validate our main results, shown in Fig. 3(a), we performed two sets of ten simulations each using the procedure of Section II, where we varied the seed of the random

number generator. The bandwidths of the OBPF in front of the regenerator in these two sets were 23 and 49 GHz. The results of these simulations are shown in Fig. 10, where we also plotted the corresponding two data points from Fig. 3(a). For each set, we used $\epsilon = 1.1$ and the $\mathrm{PDF}_{\mathrm{initial}}$ computed from the formulas in [22] (see Appendix A). In those formulas, we set the OSNR to 80% and 100% of its actual value of 17 dB (see Section III-A) for ONEs and ZEROs, respectively. This was done to optimize the rate at which distinct signal samples were collected, but did not significantly affect the final PDFs. One can see that the variations of the output BER (computed by our procedure) do not significantly exceed those of the input BER (computed by the standard MMC algorithm in stage 1 of our procedure). The outlier data point for the set corresponding to the 23-GHz OBPF occurred because the $\mathrm{PDF}_{\mathrm{ONE}}$ had a spike right at the decision threshold. Even with those variations, the main conclusion of Fig. 3(a)—that a single Mamyshev-type regenerator degrades the BER—remains unchanged.

ACKNOWLEDGMENT

The author would like to thank Prof. W. Kath for the discussion on MMC algorithm, Prof. M. Vasilyev for the discussion on matched filter, and J. Lawson for his assistance with logistical issues of running simulations at the Vermont Advanced Computing Center.

REFERENCES

[1] P. Öhlén and E. Berglind, "Noise accumulation and BER estimates in concatenated nonlinear optoelectronic repeaters," *IEEE Photon. Technol. Lett.*, vol. 9, no. 7, pp. 1011–1013, Jul. 1997.

[2] F. Öhman and J. Mørk, "Modeling of bit error rate in cascaded 2R regenerators," *J. Lightw. Technol.*, vol. 24, pp. 1057–1063, 2006.

[3] M. R. G. Leiria and A. V. T. Cartaxo, "On the optimization of regenerator parameters in a chain of 2R all-optical regenerators," *IEEE Photon. Technol. Lett.*, vol. 18, no. 16, pp. 1711–1713, Aug. 2006.

[4] I. Nasieva, A. Kaliazin, and S. K. Turitsyn, "Multicanonical Monte Carlo modelling of BER penalty in transmission systems with optical regeneration," *Opt. Commun.*, vol. 262, pp. 246–249, 2006.

[5] F. Öhman, B. Tromborg, J. Mørk, A. Aurelius, A. Djupsjöbacka, and A. Berntson, "Measurements and simulations of nonlinear noise redistribution in an SOA," *IEEE Photon. Technol. Lett.*, vol. 17, no. 1, pp. 16–18, Jan. 2005.

[6] M. Rochette, L. Fu, V. Ta'eed, D. J. Moss, and B. J. Eggleton, "2R optical regeneration: An all-optical solution for BER improvement," *IEEE J. Sel. Topics Quantum Electron*, vol. 12, no. 4, pp. 736–744, Jul./Aug. 2006.

[7] P. V. Mamyshev, "All-optical regeneration based on self-phase modulation effect," in *Proc. 24th Eur. Conf. Opt. Commun. (ECOC)* Sep. 20–24, 1998, vol. 1, pp. 475–476.

[8] M. Rochette, J. N. Kutz, J. L. Blows, D. Moss, T. J. Mok, and B. J. Eggleton, "Bit-error-ratio improvement with 2R optical regenerators," *IEEE Photon. Technol. Lett.*, vol. 17, no. 4, pp. 908–910, Apr. 2005.

[9] B. A. Berg, "Algorithmic aspects of multicanonical simulations," *Nucl. Phys. B (Proc. Suppl.)*, vol. 63 A–C, pp. 982–984, 1998; also Available at http://www.arxiv.org, paper hep-lat/9708003.

[10] D. Yevick, "Multicanonical communication system modeling—Application to PMD statistics," *IEEE Photon. Technol. Lett.*, vol. 14, no. 11, pp. 1512–1514, Nov. 2002.

[11] R. Holzlöhner and C. R. Menyuk, "Use of multicanonical Monte Carlo simulations to obtain accurate bit error rates in optical communication systems," *Opt. Lett.*, vol. 28, pp. 1894–1896, 2003.

[12] T. N. Nguyen, M. Gay, L. Bramerie, T. Chartier, J.-C. Simon, and M. Joindot, "Noise reduction in 2R-regeneration technique utilizing self-phase modulation and filtering," *Opt. Exp.*, vol. 14, pp. 1737–1747, 2006.

[13] B. A. Berg and T. Neuhaus, "Multicanonical ensemble: A new approach to simulate first-order phase transitions," *Phys. Rev. Lett.*, vol. 68, pp. 9–12, 1992.

[14] A. O. Lima, I. T. Lima, and C. R. Menyuk, "Error estimation in multicanonical Monte Carlo simulations with applications to polarization-mode-dispersion emulators," *J. Lightw. Technol.*, vol. 23, pp. 3781–3789, 2005.

[15] A. B. Carlson, *Communication Systems: An Introduction to Signals and Noise in Electrical Communication*, 3rd ed. New York: McGraw-Hill, 1986.

[16] M. Vasilyev and T. I. Lakoba, "All-optical multichannel 2R regeneration in a fiber-based device," *Opt. Lett.*, vol. 30, pp. 1458–1460, 2005.

[17] L. Provost, C. Finot, P. Petropoulos, K. Mukasa, and D. J. Richardson, "Design scaling rules for 2R-optical self-phase modulation-based regenerators," *Opt. Exp.*, vol. 15, pp. 5100–5113, 2007.

[18] T. Chartier, private communication, 2007.

[19] K. L. Corwin, N. R. Newbury, J. M. Dudley, S. Coen, S. A. Diddams, K. Weber, and R. S. Windeler, "Fundamental noise limitations to supercontinuum generation in microstructure fiber," *Phys. Rev. Lett.*, vol. 90, pp. 113904-1–113904-4, 2003.

[20] A. L. Gaeta, "Nonlinear propagation and continuum generation in microstructured optical fibers," *Opt. Lett.*, vol. 27, pp. 924–926, 2002.

[21] D. R. Solli, C. Ropers, P. Koonath, and B. Jalali, "Optical rogue waves," *Nature*, vol. 450, pp. 1054–1057, 2007.

[22] D. Marcuse, "Derivation of analytical expressions for the bit-error probability in lightwave systems with optical amplifiers," *J. Lightw. Technol.*, vol. 8, pp. 1816–1823, 1990.

[23] T. Lu and D. Yevick, "Biased multicanonical sampling," *IEEE Photon. Technol. Lett.*, vol. 17, no. 7, pp. 1420–1422, Jul. 2005.

[24] D. Yevick and T. Lu, "Improved multicanonical algorithms," *J. Opt. Soc. Am. A*, vol. 23, pp. 2912–2918, 2006.

Taras I. Lakoba received the Diploma in physics from Moscow State University, Moscow, Russia, in 1989, and the Ph.D. degree in applied mathematics from Clarkson University, Potsdam, NY, in 1996.

In 2000, he joined the Optical Networking Group, Lucent Technologies, where he was engaged in the development of a ultralong-haul dense wavelength division multiplexing (DWDM) terrestrial transmission system. Since 2003, he has been with the Department of Mathematics and Statistics, University of Vermont, Burlington, where he is currently an Assistant Professor. His current research interests include multichannel all-optical 2R regeneration, the effect of noise in optical communication systems, numerical methods for obtaining stationary solutions of nonlinear wave equations, and perturbation techniques.