Testing of a Line Driver With Configurable Pre-Emphasis on Lossy Transmission Lines

Nicholas St. John, *Member, IEEE*, Soumyajit Mandal, *Senior Member, IEEE*, Edwin Villalpando, Piotr Maj, and Grzegorz W. Deptuch, *Senior Member, IEEE*

***Abstract*— Rare-event physics experiments such as the Deep Underground Neutrino Experiment (DUNE) or the next Enriched Xenon Observatory (nEXO) experiment search for rare, low-energy events, detected by sensitive detectors immersed in a cryogenic noble liquid (e.g., liquid argon or xenon). Readout electronics used within such detectors must consume minimal power while operating reliably in cryogenic environments. Furthermore, in the case of nEXO, maximizing the radiopurity of the environment is vital to minimize background noise, thus placing strict limits on the volume of dielectric materials, leading to high-loss data cables spanning distances up to 12 m. Such cables cause high attenuation and intersymbol interference (ISI), resulting in a high bit-error rate (BER). These issues were addressed by developing an integrated line driver with configurable pre-emphasis in a 65-nm CMOS process. The pre-emphasis parameters can be programmed to minimize BER for specific cables and data rates under power constraints. The driver was tested at both room and cryogenic temperatures. In both cases, the output BER was found to be strongly correlated with the pre-emphasis settings. Furthermore, analysis and simulation showed that adapting the pre-emphasis settings based on the incoming bit sequence can further improve performance with minimal changes to the current solution.**

***Index Terms*— Configurable pre-emphasis, cryogenic CMOS, line driver, wireline communications.**

I. Introduction

PHYSICS experiments designed to observe rare events, such as in the physics of neutrinos, require cryogenic environments [1], and some have strict radiopurity requirements [2]. These demands place stringent constraints on the material and power budgets for both the detectors and their readout circuits. Dielectric materials suffer from especially poor radiopurity, so transmitting digitized data from the detector over wired links requires the use of thin custom-designed cables, resulting in high insertion loss. In addition, the loss is frequency-dependent (dispersive), resulting in a low-pass cable transfer function (TF) shown as "Cable TF" in Fig. 1.

Manuscript received 7 June 2024; accepted 14 July 2024. Date of publication 12 August 2024; date of current version 18 September 2024. This work was supported by employees of Brookhaven Science Associates, LLC, through the U.S. Department of Energy, under Contract DE-SC0012704. *(Corresponding author: Nicholas St. John.)*

Nicholas St. John, Soumyajit Mandal, Piotr Maj, and Grzegorz W. Deptuch are with the Instrumentation Department, Brookhaven National Laboratory, Upton, NY 11973 USA (e-mail: nstjohn@bnl.gov).

Edwin Villalpando was with the Instrumentation Department, Brookhaven National Laboratory, Upton, NY 11973 USA. He is now with the Electrical and Computer Engineering Department, University of California, Berkeley, CA 94720 USA.

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TNS.2024.3430845.

Digital Object Identifier 10.1109/TNS.2024.3430845

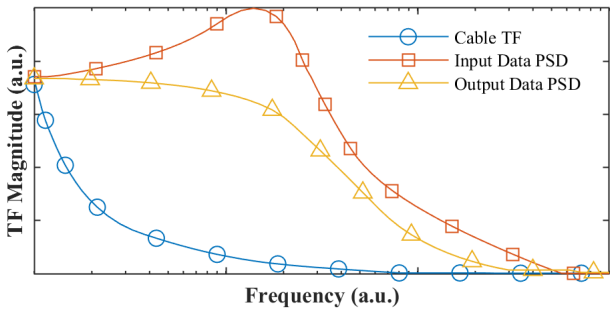

Fig. 1. Effects of pre-emphasis in the frequency domain. The plot shows a typical cable TF, the PSD of the input data after pre-emphasis, and the PSD of the output data.

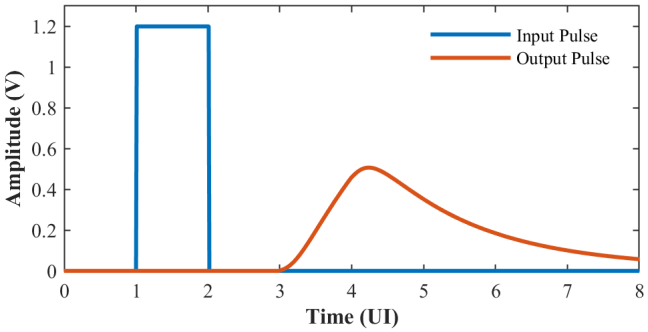

Fig. 2. Input pulse (blue) and output pulse after shaping by the cable impulse response (orange). The x-axis is normalized to the symbol period, which is also known as the UI.

In the time domain, the limited bandwidth of the cable results in pre- and post-cursor "tails" in the pulse response, as illustrated in Fig. 2. These tails cause intersymbol interference (ISI), which in turn leads to a rapid increase in bit-error rate (BER) with data rate.

Mitigating limited cable bandwidth requires reducing the data rate per cable to minimize ISI, which in turn increases the total number of parallel links and thus reduces the maximum allowable power consumption, $P_{\max}$, per link. The value of $P_{\max}$ is thus set both by the cryogenic environment, e.g., the need to avoid boiling, and the total output data rate of the detector. For example, the next Enriched Xenon Observatory (nEXO) experiment will require ~400 data links, each operating at speeds of 500 Mb/s at liquid xenon temperature (LXeT; $T = 165$ K), which imposes a limit of $P_{\max} \approx 15$ mW per link. Pre-emphasis or feed-forward equalization (FFE) is a valuable method for improving the BER of such power- and bandwidth-limited links [3], [4]. The concept involves the addition of high-frequency components to the power spectrum

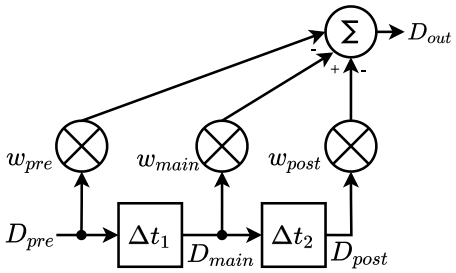

Fig. 3. Block diagram of a three-tap FIR filter used for pre-emphasis. The tap weights are denoted by w_{pre}, w_{main}, and w_{post}, respectively.

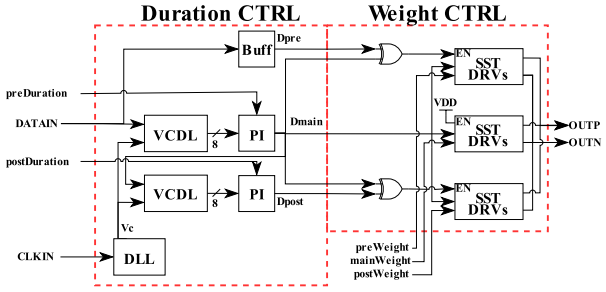

Fig. 4. Block diagram of a line driver with configurable pre-emphasis.

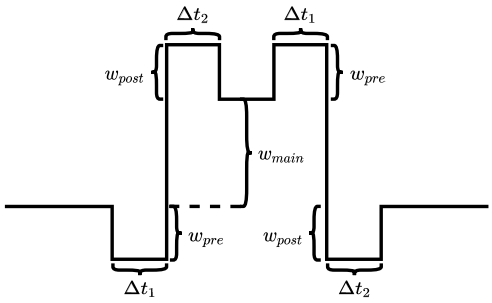

Fig. 5. Generic output waveform with configurable pre-emphasis for a single symbol (NRZ bit) [17].

of the incoming data signal which, when combined with the generally low-pass TF of the cable, results in a flatter in-band power spectral density (PSD), shown as "output PSD" in Fig. 1 and thus reduces ISI. Pre-emphasis is generally implemented via a finite-impulse response (FIR) filter clocked at the symbol period, which is also known as the unit interval (UI). Fig. 3 shows an example of a three-tap FIR filter, in which Δt_1 and Δt_2 are both hard-coded to an integer multiple of 1 UI, while the tap weights w_{pre}, w_{main}, and w_{post} are configurable.

It can be seen from Fig. 2 that, in the time domain, the goal of pre-emphasis is to subtract out the pre- and post-cursor tails seen at the receiving end of the cable, decreasing ISI. Decision-feedback equalization (DFE), a widely used nonlinear equalization method [5], performs a similar action at the receiver end. However, a DFE is not a predictive system, so it can only account for the postcursors, making pre-emphasis the only equalization method capable of canceling both pre- and post-cursors from the data signal [6].

To completely mitigate ISI, pre-emphasis must theoretically synthesize a high-pass TF that is the inverse of the cable TF over the frequency range of interest. However, in reality, the cable TF is dependent on a variety of environmental factors (e.g., signal reflections, cable length, data rate, and temperature). Thus, the ideal high-pass TF that must be generated can only be known via experimental measurements. Therefore, it would be best to maximize the flexibility of the possible TFs that can be generated by the FIR filter. However, previous designs [7], [8], [9], [10], [11], [12], [13], [14] have limited flexibility since the FFE runs at the symbol rate, i.e., the relative tap delays are hard-coded to integer multiples of 1 UI. Thus, allowing sub-UI relative tap delays, Δt_i, can improve the data link performance by enabling the output data to have a wider range of waveforms and power spectra. For example, postfabrication tuning of the FFE tap delays using variable time delay cells was shown to significantly reduce data-dependent jitter (DDJ) compared with a symbol-rate FFE [15]. The adjustment of symbol transition timings (known as phase pre-emphasis) using quadrature-rate FFEs with half-UI tap delays can also reduce DDJ [16]. However, in both cases, all the tap delays in the FFE were constrained to have the same nominal value. In contrast, here we consider a fully configurable FFE in which both the tap delays and weights are independently configurable. This line driver allows the generation of the generic output waveform shown in Fig. 5 [17]. We earlier reported the design of a line driver, shown in Fig. 4, implementing such fully configurable pre-emphasis in a 65-nm CMOS process. This circuit operates at data rates up to 1 Gb/s while consuming between 4 and 9 mW, which meets the power constraint. In addition, the output voltage swing is configurable from 0.2 to 1.2 V for generating the pre-emphasized waveform [4].

The rest of this article is organized as follows. Section II describes the design of the configurable line driver circuit being tested. Section III discusses the experimental setup both in terms of hardware and software. Experimental results are shown in Section IV and analyzed in Section V, while our conclusion is presented in Section VI.

II. Circuit Design

Fig. 4 shows the block diagram of the fully configurable line driver. The circuit, which operates from a 1.2-V power supply, is designed to transmit non-return-to-zero (NRZ) binary data streams over differential cables. The nominal data rate is 500 Mb/s, as expected for the nEXO experiment, but can be increased up to 1 Gb/s if required. An eight-stage delay-locked loop (DLL) provides a control voltage, V_c, to two voltage-controlled delay lines (VCDLs) identical to that within the DLL. These VCDLs create eight delayed copies of the input data waveform with relative delays of UI/8. The delayed copies can pass through or be mixed with adjacent taps via a phase interpolator (PI), resulting in 16 uniformly spaced relative delay options for each tap. The relative delay options are chosen via control inputs *preDuration* and *postDuration*, which correspond to Δt_1 and Δt_2, respectively, as shown in Fig. 5.

The design uses voltage-mode source-series terminated (SST) drivers to generate the output waveform due to their low power consumption, which is $\sim 4\times$ less than a current-mode driver for the same signal swing. This is illustrated in Fig. 6, where the top configuration corresponds to a current-mode logic (CML) driver and the bottom to an equivalent circuit of an SST driver. Analyzing these simplified models to solve

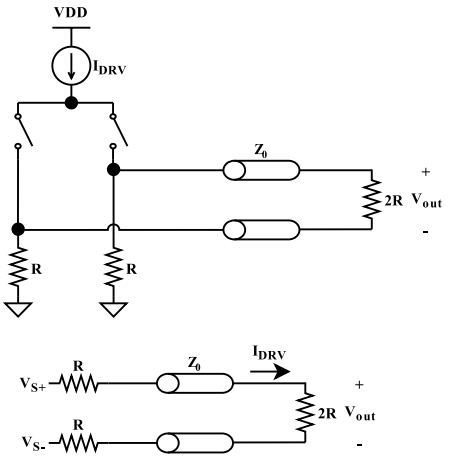

Fig. 6. Comparison between the current-mode data transmission (top) and voltage-mode data transmission (bottom).

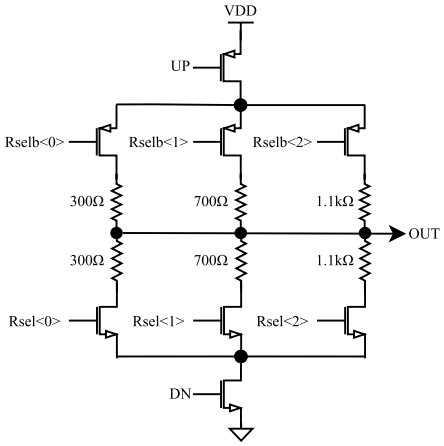

Fig. 7. SST driver cell with configurable source impedance.

for the output current, I_{DRV}, yields $I_{\mathrm{DRV}} = V_{\mathrm{out}}/8R$ for the current-mode case and $I_{\mathrm{DRV}} = V_{\mathrm{out}}/2R$ for the voltage-mode case. Thus, the value of I_{DRV} can be $4\times$ smaller in the voltage-mode case to get the same output voltage swing, allowing its power consumption to be reduced by the same factor.

The proposed line driver contains three separate arrays of SST drivers, with one array used per tap and the outputs summed together to generate the pre-emphasized waveform. Each array is composed of eight parallel driver slices with a configurable source impedance, as illustrated in Fig. 7. The data input for all three arrays is the main-data signal. However, the pre- and post-tap arrays are controlled by an XOR function between the main-data and pre- or post-data signals, respectively, which enables them during the predefined *preDuration* and *postDuration* periods, while the main-data SST array is always enabled. Within the SST array, the input data are inverted and used to drive both the UP and DN signals of each driver cell. Thus, the output of each cell is always either logic "1" or logic "0" and can only be in the high-Z state if the cell is not enabled. The number of cells that are enabled is determined via the control inputs *preWeight*, *mainWeight*, and *postWeight*, which correspond to w_{pre}, w_{main}, and w_{post}, respectively.

A 1×1 mm^2 test structure, shown in Fig. 8(a), was used to test the proposed configurable line driver. The chip contains

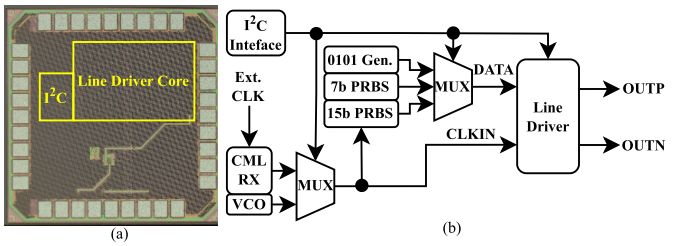

Fig. 8. (a) Die photograph of the 1-mm^2 line driver test chip in a 65-nm CMOS process and (b) block diagram of the chip.

the complete line driver and its test infrastructure, as shown in the block diagram in Fig. 8(b). The system clock is fed via either a CML receiver or an on-chip voltage-controlled oscillator (VCO), while the data are fed via on-chip generators that produce either a toggling ("0101, . . . ,") input or a 7- or 15-bit linear feedback shift register (LFSR) that generates a pseudorandom binary sequence (PRBS). The entire chip is programmed via an on-chip I^{2}C serial interface.

The line driver was tested on high-loss cables reflecting the electrical properties of the radiopure cables expected to be used in the nEXO experiment. The main experimental goal was to validate the expected improvement in data transmission efficiency provided by configurable pre-emphasis. Tests were run both at room temperature and in cryogenic conditions to confirm that the line driver can operate in the near-detector environment for rare-event physics experiments.

III. Experimental Setup

A. Hardware

The physical test setup used for verifying the line driver is shown in Fig. 9. The test chip was manually programmed over I^{2}C via a graphical user interface (GUI) that controls the FPGA on a single-board computer (National Instruments sbRIO-9629). The system clock was generated via an external low-jitter clock generator (Texas Instruments CDCM6208V2). This source is preferred to the on-chip VCO, which has significantly higher jitter. The differential output of the line driver is sent across the test cable, with each output terminated to the 50-Ω inputs of a high-speed digital oscilloscope (Tektronix MSO 72004C, 20-GHz bandwidth and 100 GS/s). The time-domain output signal, which spans at least 100 symbol periods (200 ns at 500 Mb/s), was captured for further analysis.

The cable under test is shown in Fig. 10. The design aimed to reflect the electrical properties of the radiopure cables expected to be used in the nEXO experiment. It consists of a differential stripline, i.e., a differential pair made of a thin layer of copper, which is sandwiched between two ground planes (also made of copper) using polyimide insulation. The prototype cable was fabricated in three different lengths, 3, 6, and 9 m. Based on calculations, the dc resistance should be approximately $R' = 15\ \Omega$/m with a differential characteristic impedance of $Z_0 = 100\ \Omega$. Ignoring shunt conductance, the expected attenuation coefficient of the cable at low frequencies is

$$\alpha_R \approx \frac{R'}{2Z_0} = 0.075\ \mathrm{m}^{-1} \tag{1}$$

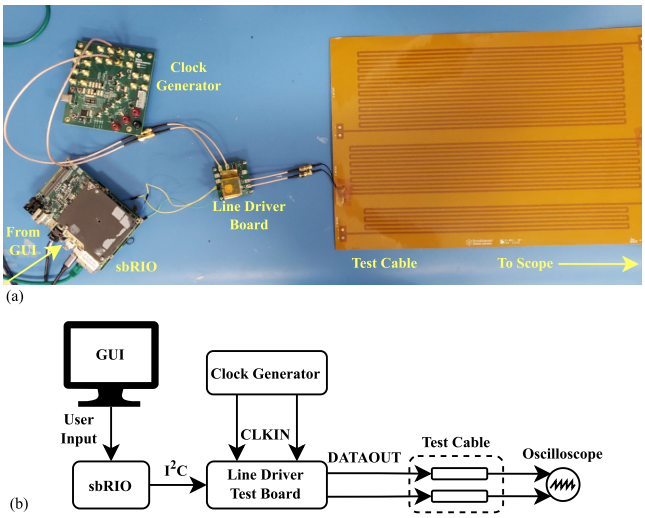

Fig. 9. (a) Experimental test setup and (b) block diagram.

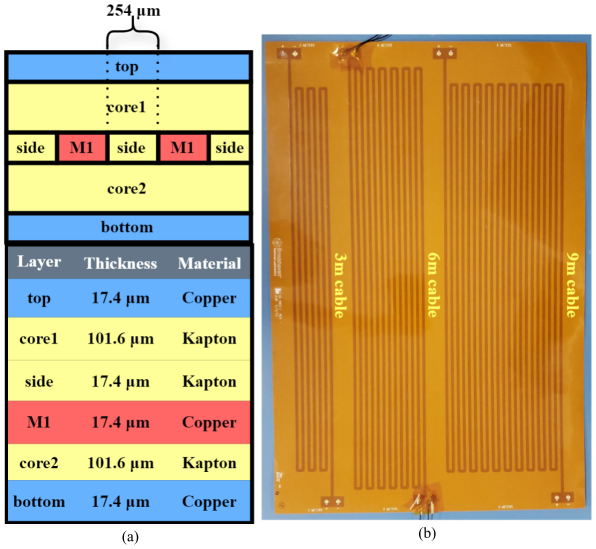

Fig. 10. Cable prototype used for testing. (a) Cross section (stack up) and (b) fabricated prototype.

corresponding to a signal loss of $20\alpha_R/\ln(10) \approx 0.65$ dB/m. However, the loss increases with frequency due to the impact of skin and proximity effects on R'.

B. Software

Due to the fact that the line driver has over 390 000 possible combinations of settings [18], an automated test system is required to run parameter sweeps. For this reason, a LabVIEW program was developed, performing the entire test flow autonomously for a user-defined set of parameter combinations. A flowchart of the software process is shown in Fig. 11. The process begins by initializing an array of settings, which is then loaded into the line driver chip over I^{2}C. The oscilloscope is programmed to capture the resulting time-domain waveforms at the receiving end of the cable. A custom MATLAB script (called from within LabVIEW) postprocesses the waveform data to estimate the eye diagram and then extract a figure-of-merit (FoM), defined as follows:

$$\text{FoM} = w_{\text{eye}} \times h_{\text{eye}} \tag{2}$$

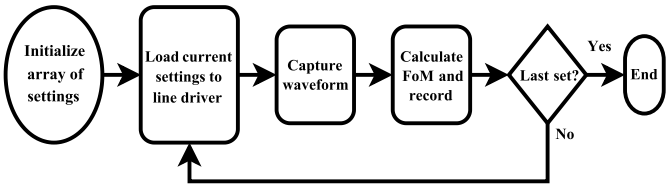

Fig. 11. Flowchart of the software used to automate line driver testing.

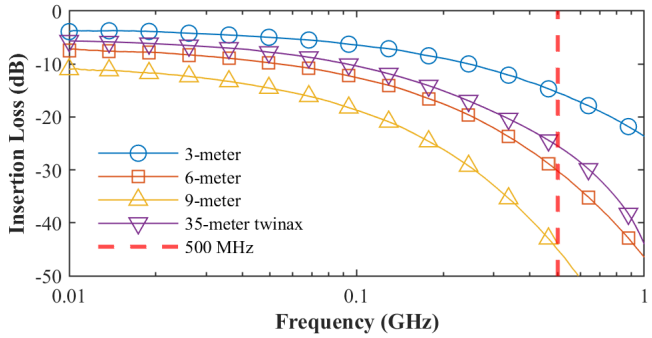

Fig. 12. Measured insertion loss of test cables with different lengths at RT. Loss at 500 MHz: −15 dB (3 m, radiopure), −30 dB (6 m, radiopure), −45 dB (9 m, radiopure), and −25 dB (35 m, twinax).

where w_{eye} and h_{eye} are the eye width and eye height, respectively. This FoM value is a worst-case measurement as it measures the eye opening for single transitions. This process is then repeated for a new array of settings to complete a parameter sweep. Each run only takes a few seconds, allowing large sweeps to be run within a reasonable time frame.

IV. Results

A. Cable Characterization

Prior to testing the line driver, each of the three test cables (with lengths of 3, 6, and 9 m, respectively) was characterized at room temperature (RT) using a four-port vector network analyzer (VNA). The measured four-port single-ended S-parameters of each cable were then converted into differential S-parameters. The results were analyzed to obtain insertion loss, which is plotted in Fig. 12. At 500 MHz, which is equal to 1/UI for the data transmission rate of 500 Mb/s used in these experiments, the loss of the cable was found to be 5 dB/m, which is about 8× larger than the loss at dc as calculated in (1). The result is signal losses of at least 30 dB over the distances required for the nEXO experiment (between 6 and 12 m). The figure also shows the measured insertion loss of a typical nonradiopure cable for comparison. The results show that a 35-m twinax cable, similar to that used by the Deep Underground Neutrino Experiment (DUNE), has the same insertion loss as only 5 m of radiopure cable.

The 6-m cable was further characterized in liquid nitrogen (LNT), and the insertion loss is compared with that at RT in Fig. 13. Note that LNT ($T = 77$ K) was used as opposed to LXeT ($T = 165$ K) since performance at LXeT can be interpolated from LNT and RT measurements.

As expected, the insertion loss decreases considerably at LNT due to the lower series resistance, R'. A detailed mathematical analysis of the cable characteristics with respect to temperature can be found in the Appendix. This analysis shows that the variation in cable loss with temperature at 500 MHz

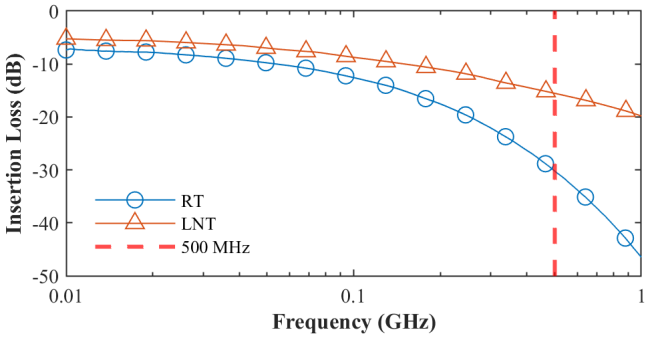

Fig. 13. Comparison of measured insertion losses of the 6-m test cable at RT and LNT. Loss at 500 MHz: −30 dB (at RT) and −17 dB (at LNT).

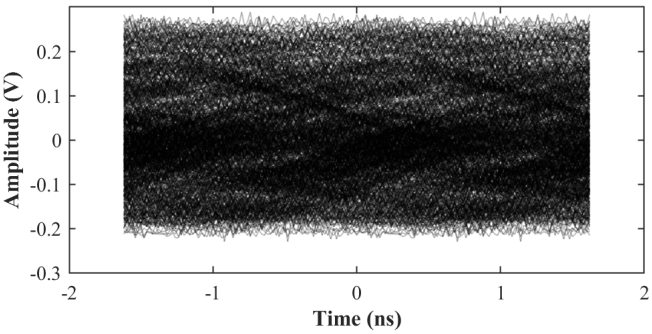

Fig. 14. Eye diagram of 7-b PRBS data with no pre-emphasis across 6-m test cable at 500 Mb/s.

(which lies firmly in the skin effect-dominated regime) is in agreement with the theory.

B. Testing of Configurable Pre-Emphasis

Prior to sending data with pre-emphasis across the cables, 7-bit PRBS data were sent across the 6-m test cable at 500 Mb/s with no pre-emphasis (both the pre- and post-tap SST drivers disabled, all seven of the main-tap drivers enabled). Fig. 14 shows that the eye at the receiving end is completely closed. Thus, pre-emphasis is a necessity in order to transmit data across this cable, which justifies the need for further testing.

Using the test automation program discussed earlier, a sweep of 2048 sets of parameters was run with the line driver sending 7-bit PRBS data over the 6-m test cable at 500 Mb/s. This subset of 2048 parameters was chosen to test each parameter across its allowable range while keeping the overall runtime reasonable. The values swept for each parameter were as follows.

1) *preWeight:* [0, 2, 5, 7] (# of SST slices enabled).
2) *mainWeight:* [1, 2, 5, 7] (# of SST slices enabled).
3) *postWeight:* [0, 2, 5, 7] (# of SST slices enabled).
4) *preDuration:* [4, 8, 12, 15] (#/16 of UI).
5) *postDuration:* [4, 8, 12, 15] (#/16 of UI).
6) *Rsel:* [1, 3] (1 = 300 Ω and 3 = 1.1 kΩ).

The FoM estimated for each run can be seen from Fig. 15(a). The measured FoM changes by an order of magnitude over the range of the sweep, illustrating that configurable pre-emphasis has a significant impact on data link performance. Fig. 15(b) shows the eye diagram for the optimal parameter set based on the FoM. Note that the signal swing is lower than without pre-emphasis, as shown in Fig. 14. However, the use of

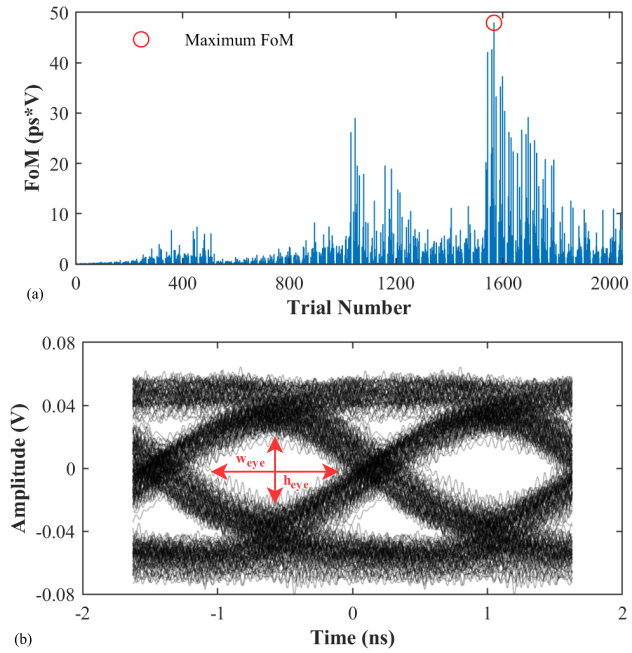

Fig. 15. Parameter sweep at RT. (a) FoM per trial and (b) eye diagram for the maximum FoM (eye height: 36 mV and eye width: 1.33 ns). Line driver settings for maximum FoM: preWeight = 0, mainWeight = 1, postWeight = 7, preDuration = 15, postDuration = 15, and Rsel = 1.

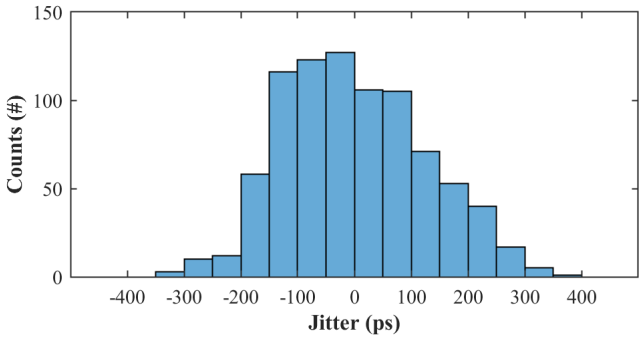

Fig. 16. Measured jitter histogram for the eye diagram with maximum FoM at RT. RMS jitter = 125.6 ps (6.3% of UI) and peak-to-peak jitter = 694 ps (34.7% of UI).

pre-emphasis reduces ISI, which results in a greatly improved eye diagram and more reliable data transmission. The jitter of the received waveform around the transition instants was also measured, and the results are shown in Fig. 16. The histogram is somewhat asymmetric but has a single peak, which suggests that random jitter (RJ) is dominant, i.e., DDJ due to ISI can be ignored. Collectively, these results show that it is possible to drive the 6-m cable at 500 Mb/s with no equalization at the receiving end, while still obtaining an acceptable eye diagram.

The output data waveform producing the highest FoM at RT is shown in Fig. 17. The waveform settings are maximum post-tap weight and duration (w_{post} and Δt_2), minimum main-tap weight (w_{main}), and zero pretap weight ($w_{\text{pre}} = 0$). In general, it was observed that nonzero pretap weights had an adverse affect on the output eye opening. The reasons for this behavior were further studied and will be discussed in Section V. Also note that the waveform shape resembles the derivative of the input data sequence. This arises from the fact that the TF of the 6-m cable at RT, as shown in Fig. 12, can be

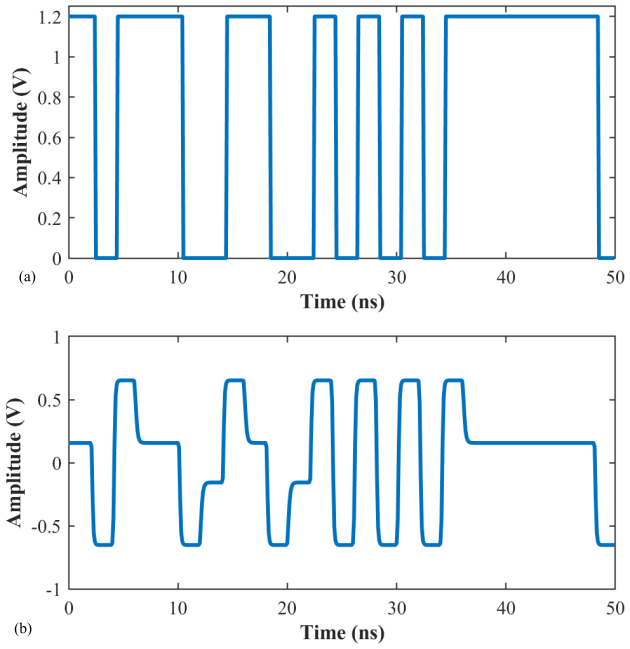

Fig. 17. (a) Input data and (b) output waveform with the pre-emphasis parameters that result in the maximum FoM at RT.

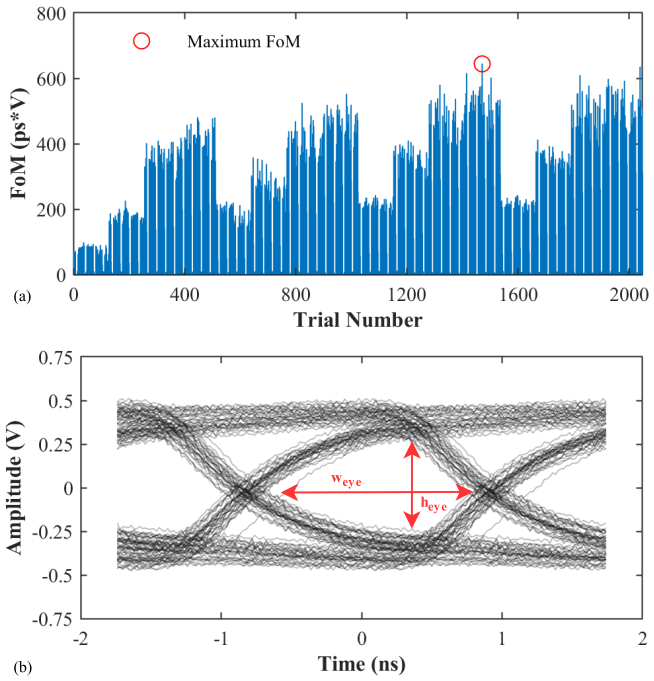

Fig. 18. Parameter sweep at LNT. (a) FoM per trial and (b) eye diagram for the maximum FoM (eye height: 528 mV and eye width: 1.22 ns). Line driver settings for maximum FoM: preWeight = 2, mainWeight = 7, postWeight = 5, preDuration = 15, postDuration = 15, and Rsel = 1.

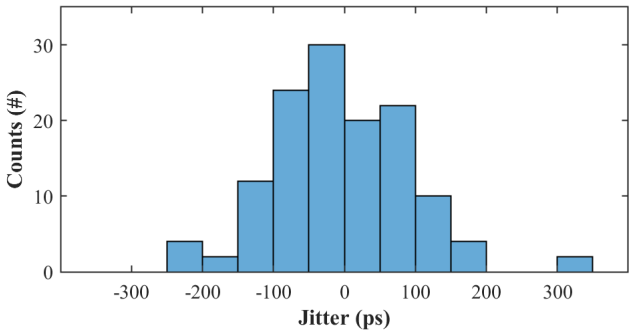

Fig. 19. Measured jitter histogram for the eye diagram with maximum FoM at LNT. RMS jitter = 95.57 ps (4.8% of UI) and peak-to-peak jitter = 540 ps (27% of UI).

approximated by single-pole roll-off (i.e., an integrator) over a relatively broad frequency range (100–500 MHz). Therefore, the inverse TF synthesized by the pre-emphasis FIR filter should mimic differentiation.

The same test was repeated with the cable at LNT. Since the cable loss is significantly lower at cryogenic temperatures, we expect the FoM to improve. This behavior is visible in Fig. 18(a), which shows that the maximum FoM is larger than that at RT by an order of magnitude. The optimal eye diagram, shown in Fig. 18(b), is clearly open. It also has less jitter compared with the case at RT as shown in Fig. 19, while RJ is again dominant. Therefore, the transmitted data can be easily recovered by a receiver with no equalization.

Comparison of Figs. 15(a) and 18(a) reveals that the FoM varies differently with parameter values at RT and LNT, i.e., the optimal pre-emphasis settings change with temperature. This fact can be further verified by plotting the transient waveform that provides the best FoM at LNT, as shown in Fig. 20. Note that this waveform is quite different from the optimum waveform at RT, which was plotted in Fig. 17. This result confirms the earlier claim that temperature can have a significant affect on the performance of the wired data link. Interestingly, the general trend seen at RT, in which nonzero pretap weights degrade performance, is not seen at LNT. The reasons for this observation are discussed in Section V. Finally, note that the results for LXeT are expected to lie between those for LNT and RT.

V. Analysis

From Fig. 17, we can see that the waveform that maximizes the FoM at RT uses both maximum weight and duration of the post-tap, while minimizing the main-tap weight and completely eliminating pretap operation. Thus, we can say that the FoM generally is directly proportional to the post-tap strength, and inversely proportional to the strength of the pretap and, surprisingly, the strength of the main-tap as well. This is not desirable, since it decreases the signal amplitude and results in a small eye height. In order to understand why this is the case, we refer to the output waveform in Fig. 21, which was obtained by sending data across the 6-m test cable with all three tap weights maximized.

Fig. 21 shows that the output waveform exhibits a low-frequency common-mode shift, which is more commonly referred to as dc wandering [19]. This behavior arises from the fact that, for high enough losses, a full transition cannot occur within the period of a single bit (i.e., 1 UI). Therefore, when longer sequences of repeating bits are seen, the signal amplitude changes beyond that of a single bit, shifting the common-mode voltage over time. Data links usually account for this affect using line codes which ensure that there can only be a certain number of repeating bits without a transition, such as in 8-b/10-b encoding [20]. However, the well-known benefits of such encoding come at the cost of reduced symbol

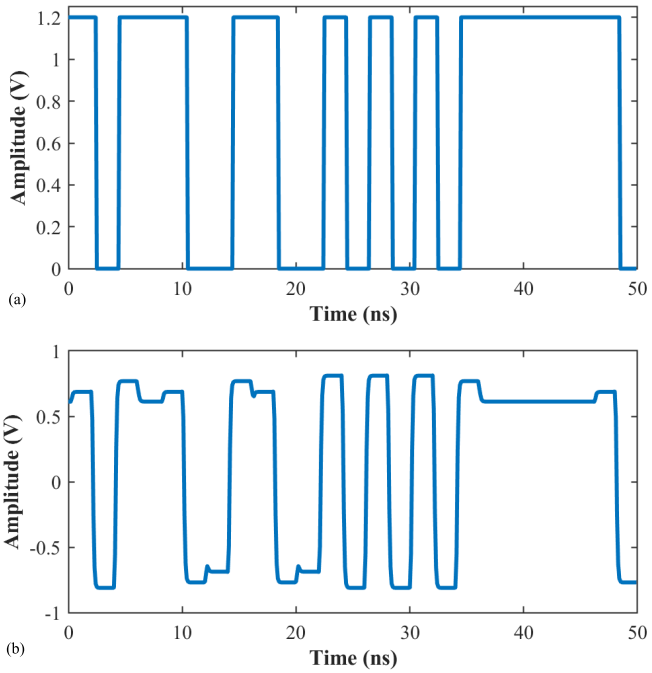

Fig. 20. (a) Input data and (b) output waveform with the pre-emphasis parameters that result in the maximum FoM at LNT.

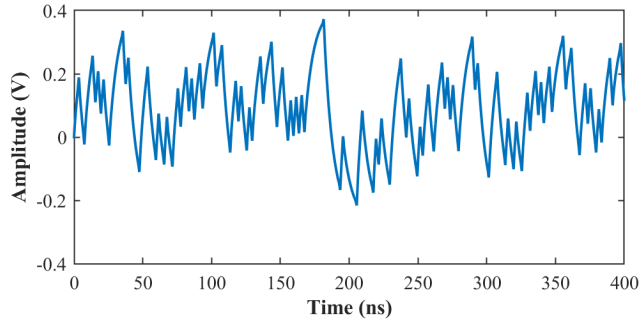

Fig. 21. Simulated output data across 6-m test cable with all three tap weights (w_{pre}, w_{main}, and w_{post}) maximized, showing the effects of dc wandering.

period and increased bandwidth for a given data rate, which in turn can significantly increase the amount of ISI introduced by a high-loss cable TF. For example, consider the 6-m cable at RT for the nominal data rate of 500 Mb/s. The 25% bandwidth overhead of an 8-b/10-b code would introduce ~7 dB of additional loss at a frequency of 1/UI. Thus, adaptive pre-emphasis and line coding can be considered as complementary techniques.

Another possible way to mitigate this dc wandering affect would be to eliminate the pretap and reduce the main-tap weight *only* if a repeating bit occurs. This would equalize the amplitude of sequences of repeating bits to that of a single bit, without lowering the amplitude of the transmitted signal.

This concept, which we refer to as *adaptive pre-emphasis*, was studied to see if it can further improve data link performance for lossy cables as compared to the configurable (but nonadaptive) pre-emphasis we have discussed so far. Fig. 22 illustrates the adaptive pre-emphasis concept. The pretap is enabled for single bits. However, when a repeating bit is seen, the main weight is reduced by a factor of $k > 1$, and the pretap weight is set to 0 prior to the next transition.

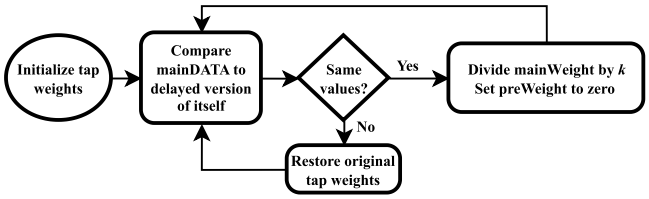

Fig. 22. Flowchart of the proposed adaptive pre-emphasis method.

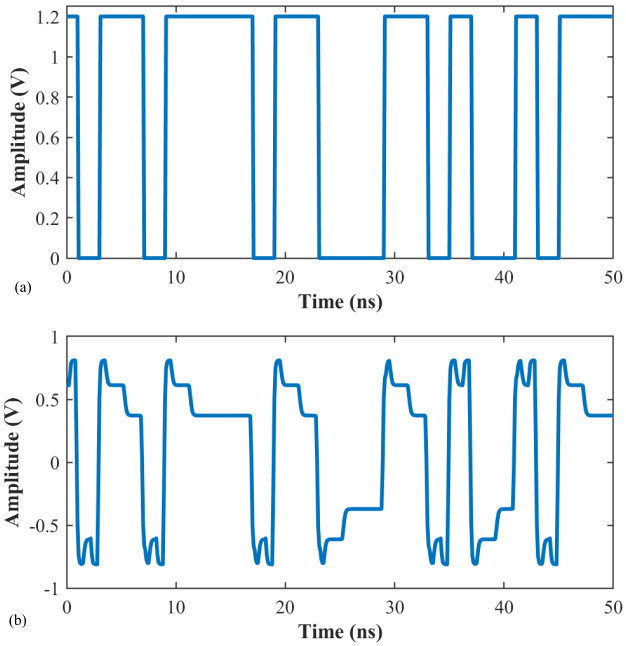

Fig. 23. (a) Simulated input data and (b) output waveform with adaptive pre-emphasis. Typical tap weights and durations were assumed, while the weight adaptation factor was set to $k = 2$.

Fig. 23 illustrates a typical PRBS data stream with adaptive pre-emphasis enabled.

In order to determine if adaptive pre-emphasis actually improves performance with respect to configurable (but nonadaptive) pre-emphasis, a parameter sweep was run in simulation for both pre-emphasis methods (assuming a 6-m cable at RT) over the same parameter ranges as in the experiments. The best eye diagrams for each case were then compared with one another. Rather than relying on the intuitive FoM defined in (2) (i.e., the product of eye height and eye width), here we explicitly define "best" to mean "lowest BER." The BER, P_e, of a binary NRZ data stream arises from both amplitude and timing errors, i.e., from voltage noise and jitter, respectively. Assuming that the voltage noise is Gaussian, the corresponding BER for a given sampling instant, t_s, is given by

$$P(\epsilon|t_s) = \sum_{i=\{0,1\}} \frac{p_i}{4} \text{erfc}\left(\frac{|v_i(t_s) - V_D|}{\sigma_i\sqrt{2}}\right) \tag{3}$$

where erfc$(\cdot)$ is the complementary error function, the values of p_i $(i = \{0, 1\})$ are the bit probabilities, the values of$v_i(t_s)$ are the mean voltage levels at the sampling instants, the values of σ_i are the corresponding rms noise levels, and V_D is the decision threshold. Note that the latter includes the effects of both fundamental noise sources (thermal and shot noise) as well as ISI. The argument of the complementary error function

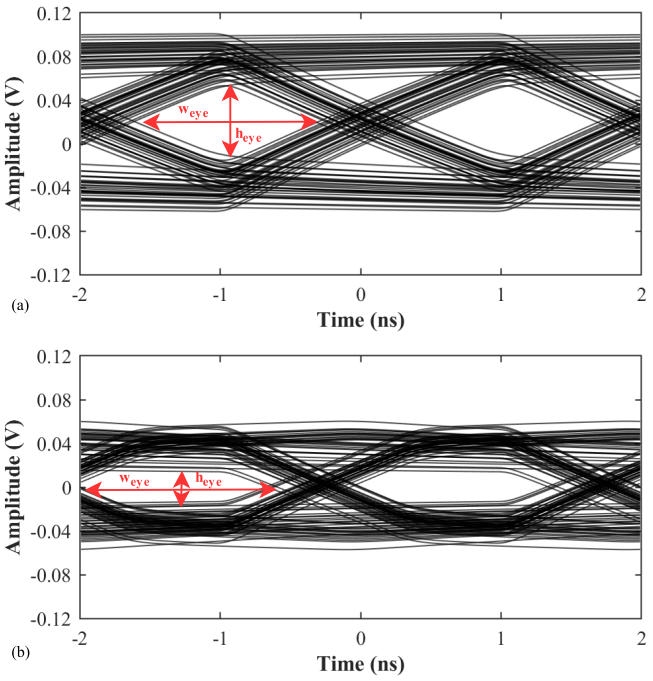

Fig. 24. Simulated eye diagrams for (a) nonadaptive pre-emphasis (eye width: 1.347 ns and eye height: 69.99 mV) and (b) adaptive pre-emphasis (eye width: 1.327 ns and eye height: 28.82 mV) assuming a 6-m-long cable at RT.

in (3) is known as the signal-to-noise ratio (SNR) of the eye diagram and denoted by $\mathrm{SNR}(t_s)$.

In addition, the sampling instant, t_s, varies randomly due to the clock and data jitter. Denoting the probability distribution of t_s by $p(t_s)$, the expected BER is given by

$$P_e = \int_{-\infty}^{\infty} P(\epsilon|t_s)p(t_s)dt_s. \tag{4}$$

Since P_e is a monotonically decreasing function of $\mathrm{SNR}(t_s)$ for a given sampling instant, this quantity can be used as an FoM for comparing eye diagrams. The rms jitter of the clock generator used in the experiments (Texas Instruments CDCM6208V2) was $\sigma_t \approx 0.3$ ps, which is $<10^{-3}\times$ UI and so has a negligible impact on BER. Thus, it is sufficient to use the maximum value of $\mathrm{SNR}(t_s)$, which corresponds to the optimum sampling instant, as the FoM.

The top performing eye diagrams for both cases are shown in Fig. 24 for an adaptation factor of $k = 2$. The measured SNR improved by 5% on applying adaptive pre-emphasis as compared with pre-emphasis with no adaptability. Furthermore, power consumption decreased by over 20% (from 4.9 to 3.8 mW) due to the reduced rms output voltage. This result suggests that adaptive pre-emphasis will provide significantly higher SNR than nonadaptive pre-emphasis for the same power consumption.

Note that the eye height obtained using adaptive pre-emphasis is actually smaller than with nonadaptive pre-emphasis, due to the fact that the amplitude is lowered for repeating bit sequences. This places more stringent requirements on the receiver in terms of amplitude noise. Nevertheless, it has higher SNR due to reduced ISI, which in turn reduces the values of σ_0 and σ_1. In addition, the peak input amplitude can be increased for a given power consumption,

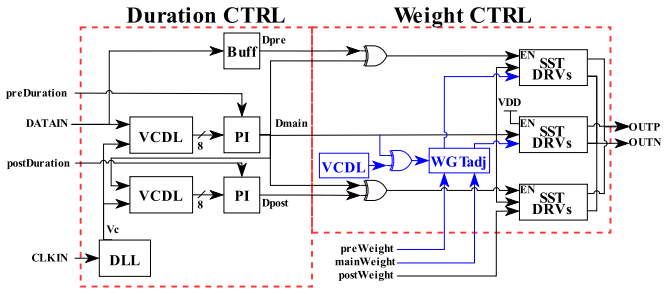

Fig. 25. Block diagram of the modified line driver, with the added circuits required to implement adaptive pre-emphasis shown in blue.

which mitigates the eye amplitude issue and thus improves the BER. Furthermore, the SNR is also relatively constant with the position of the sampling instant, t_s, thus increasing the tolerance to jitter. In contrast, nonadaptive pre-emphasis results in a diamond-like eye opening, such that SNR decreases rapidly away from the ideal sampling instant. Fundamentally, this is because of the limited bandwidth of the cable TF, which results in transition times that are comparable with the symbol period (i.e., 1 UI).

Given the conclusion that the adaptive pre-emphasis improves data link performance, it is currently being implemented within the line driver. Fortunately, this requires minimal hardware changes, as shown in Fig. 25, where the required additions are shown in blue. Repeating bits are detected by delaying the main tap via an additional VCDL and comparing it with the main tap data. If these match, a repeating bit is detected and the pretap and main-tap weights (w_{pre} and w_{main}, respectively) are adjusted as follows: the pretap weight is zeroed completely (resulting in $w_{\mathrm{pre}} = 0$), while the main-tap weight is divided by a user-defined value (resulting in w_{main}/k, where $k > 1$). The value of k can be programmed over the range 1–8.

VI. Conclusion

A line driver with configurable pre-emphasis was successfully tested on different high-loss cables reflecting the electrical properties of the radiopure cables expected to be used in the nEXO experiment, both at RT and LNT. The inclusion of configurable pre-emphasis was shown to have a highly beneficial impact on data transmission performance. Furthermore, the concept of varying the pre-emphasis parameters based on the input bit sequence, known as adaptive pre-emphasis, was also proposed and simulated. The proposed adaptive pre-emphasis method implements a simple, single-point weight adjustment based on the detection of repeating bit sequences. Simulations reveal that this technique, which only requires minor circuit modifications to the original line driver, can both improve BER and lower power consumption compared with nonadaptive pre-emphasis. The scheme is currently being implemented within a modified version of the line driver circuit, which is promising for use in the readout circuits for rare-event search experiments. Further work will explore automated on-chip optimization of the pre-emphasis settings for a given cable using artificial intelligence (AI) processing circuitry.

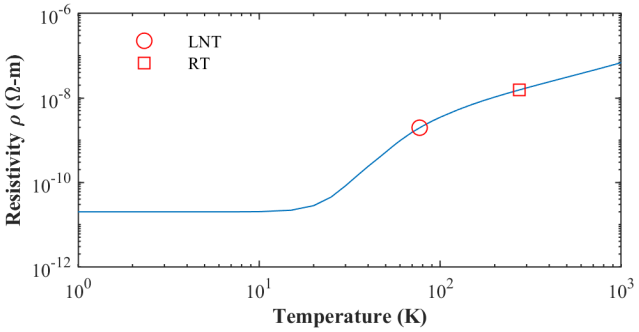

Fig. 26. Typical resistivity of annealed copper as a function of temperature, with the intrinsic resistivity $\rho_0 = 2 \times 10^{-11}$ Ω·m. Data from [21]. Resistivity at RT and LNT are labeled.

APPENDIX

The resistivity, ρ, of metals such as copper is dependent on the presence of impurities, which result in scattering of the carriers [21]. The contribution of such impurities to resistivity is approximately temperature-independent, such that the total resistivity versus temperature can be written as follows:

$$\rho(c, T) = \rho_0(c) + \rho_i(T) \tag{5}$$

where c is the impurity concentration and ρ_0 is the residual resistivity at 0 K due to the impurities. For example, a good-quality annealed copper wire has $\rho_0 \approx 2 \times 10^{-11}$ Ω·m. In addition, ρ_i is the intrinsic resistivity, which goes 0 at 0 K approximately $\propto T^5$ and becomes $\propto T$ at high temperatures, i.e., as $T \to \infty$ [21]. The resulting total resistivity $\rho(c, T)$ for annealed copper wire is shown in Fig. 26; it is approximately proportional to temperature (i.e., $\rho(T) \propto T$) above ~55 K.

The skin depth in a conductor is given by the well-known expression

$$\delta(T) = \sqrt{\frac{2\rho(T)}{\omega\mu}} \tag{6}$$

where $\omega = 2\pi f$ is the angular frequency, μ is the permeability (equal to $\mu_0 = 4\pi \times 10^{-7}$ H/m for nonmagnetic materials such as copper), and we have explicitly included the temperature dependence of ρ. The series resistance of the cable, R', is dominated by that of the center conductors. Denoting the thickness and width of these conductors by t and w, respectively, we get

$$\begin{aligned} R'(T) &\approx \frac{\rho(T)}{wt} \quad \text{if } \delta > t/2 \\ &\approx \frac{\rho(T)}{2w \cdot \delta(T)} \quad \text{if } \delta < t/2. \end{aligned} \tag{7}$$

The transition between these two regimes is temperature-dependent, but occurs at a few megahertz for our cable (which has $t = 17.4$ μm). For example, it decreases from 60 to 15 MHz between RT and LNT.

Combining (5)–(7), we see that $R' \propto T$ at low frequencies (where $\delta > t/2$). For example, the measured dc resistance for the 6-m cable was estimated to be 67 and 27 Ω for RT and LNT respectively, in reasonable agreement with our analysis. In contrast, $R' \propto (T)^{1/2}$ in the skin effect-dominated regime (where $\delta < t/2$). In this case, we expect the cable loss to decrease by a factor of approximately 2 (in dB) as the temperature decreases by a factor of 4, e.g., from RT to LNT.

ACKNOWLEDGMENT

The authors would like to thank Veljko Radeka, Sergio Rescia, and Eric Raguzin of the Instrumentation Division, Brookhaven National Laboratory, Upton, NY, USA, for their help and guidance in our work.

REFERENCES

[1] F. Capozzi, S. W. Li, G. Zhu, and J. F. Beacom, "DUNE as the next-generation solar neutrino experiment," *Phys. Rev. Lett.*, vol. 123, no. 13, Sep. 2019, Art. no. 131803.

[2] J. Albert, "Sensitivity and discovery potential of the proposed nEXO experiment to neutrinoless double-β decay," *Phys. Rev. C*, vol. 97, no. 6, 2018, Art. no. 065503.

[3] J. F. Bulzacchelli, "Equalization for electrical links: Current design techniques and future directions," *IEEE Solid-State Circuits Mag.*, vol. 7, no. 4, pp. 23–31, Fall. 2015.

[4] N. St. John et al., "A low-power 1 Gb/s line driver with configurable pre-emphasis for lossy transmission lines," *J. Instrum.*, vol. 18, no. 4, Apr. 2023, Art. no. C04009.

[5] B. Razavi, "The decision-feedback equalizer [a circuit for all seasons]," *IEEE Solid-State Circuits Mag.*, vol. 9, no. 4, pp. 13–132, Fall 2017.

[6] F. Yuan, A. R. Al-Taee, A. Ye, and S. Sadr, "Design techniques for decision feedback equalisation of multi-giga-bit-per-second serial data links: A state-of-the-art review," *IET Circuits, Devices Syst.*, vol. 8, no. 2, pp. 118–130, Mar. 2014.

[7] Y. Fu, Z. Wen, and L. Chen, "A 3.75 Gb/s CML output driver with configurable pre-emphasis in 65nm CMOS technology," in *Proc. 2nd Int. Conf. Commun. Inf. Process.*, Nov. 2016, pp. 200–205.

[8] C. Menolfi et al., "A 16Gb/s source-series terminated transmitter in 65 nm CMOS SOI," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2007, pp. 446–614.

[9] V. Gromov et al., "Development of a low power 5.12 Gbps data serializer and wireline transmitter circuit for the VeloPix chip," *J. Instrum.*, vol. 10, no. 1, Jan. 2015, Art. no. C01054.

[10] K. Iniewski et al., "3.125 Gb/s power efficient line driver with 2-level pre-emphasis and 2kV HBM ESD protection," in *Proc. IEEE Int. Symp. Circuits Syst.*, 2005, pp. 1154–1157.

[11] K. Huang, Z. Wang, X. Zheng, C. Zhang, and Z. Wang, "A 80 mW 40 Gb/s transmitter with automatic serializing time window search and 2-tap pre-emphasis in 65 nm CMOS technology," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 62, no. 5, pp. 1441–1450, May 2015.

[12] P. Moreira et al., "IpGBT documentation: Release," CERN, Tech. Rep. 2809058, 2022. [Online]. Available: https://cds.cern.ch/record/2809058

[13] J. L. Casall, "Driver pre-emphasis for data transmission," Ph.D. dissertation, Dept. Elect. Comput. Eng., Texas Tech Univ., Lubbock, TX, USA, 2001.

[14] K. Jung, Y. Lu, and E. Alon, "Power analysis and optimization for high-speed I/O transceivers," in *Proc. IEEE 54th Int. Midwest Symp. Circuits Syst. (MWSCAS)*, Aug. 2011, pp. 1–4.

[15] M. Bichan and A. Chan Carusone, "A 6.5 Gb/s backplane transmitter with 6-tap FIR equalizer and variable tap spacing," in *Proc. IEEE Custom Integr. Circuits Conf.*, Sep. 2008, pp. 611–614.

[16] J. F. Buckwalter, M. Meghelli, D. J. Friedman, and A. Hajimiri, "Phase and amplitude pre-emphasis techniques for low-power serial links," *IEEE J. Solid-State Circuits*, vol. 41, no. 6, pp. 1391–1399, Jun. 2006.

[17] G. D. N. S. John and S. Mandal, "User-configurable high-speed line driver," U.S. Patent 2023 0 231 746 720, Jul. 20, 2023.

[18] N. St. John, S. Mandal, E. Raguzin, P. Maj, and G. W. Deptuch, "Self-adaptation of line driver pre-emphasis parameters via embedded line loss sensing," in *Proc. 18th Conf. Ph.D Res. Microelectron. Electron. (PRIME)*, Jun. 2023, pp. 249–252.

[19] N. Na, D. M. Dreps, and J. A. Hejase, "DC wander effect of DC blocking capacitors on PCIe Gen3 signal integrity," in *Proc. IEEE 63rd Electron. Compon. Technol. Conf.*, May 2013, pp. 2063–2068.

[20] A. X. Widmer and P. A. Franaszek, "A DC-balanced, partitioned-block, 8B/10B transmission code," *IBM J. Res. Develop.*, vol. 27, no. 5, pp. 440–451, Sep. 1983.

[21] R. A. Matula, "Electrical resistivity of copper, gold, palladium, and silver," *J. Phys. Chem. Reference Data*, vol. 8, no. 4, pp. 1147–1298, Oct. 1979.