

Discrete Multitone Signalling for Wireline Communication

Behraz Vatankhahghadim*, Nijwm Wary†, and Anthony Chan Carusone*

**Department of Electrical and Computer Engineering, University of Toronto, Toronto, Canada*

†*School of Electrical Science, Indian Institute of Technology Bhubaneswar, Bhubaneswar, India*

Abstract—For serial wireline communication beyond 56 Gb/s, bandwidth-efficient modulation is needed. Four-level pulse amplitude modulation (4-PAM) has become the standard technique at 56-64 and 112 Gb/s. However, whereas decision feedback equalization (DFE) with 5 taps or more was common for 2-PAM links at lower data rates, the speculative look-ahead techniques required to satisfy a DFE’s timing requirements at 56 Gb/s increase power consumption exponentially with the number of taps and the number of modulation levels. Thus, 4-PAM DFEs at 56 Gb/s are generally limited to 2 taps or fewer. This limitation, in turn, necessitates the use of long (10 taps or more) finite impulse response (FIR) feed-forward equalization (FFE) to accurately eliminate the intersymbol interference in long reach (LR) channels. Thus, LR receivers at 56-64 and 112 Gb/s comprise an analog-to-digital converter (ADC) followed by digital equalization, and the transmitters increasingly use a digital-to-analog converter (DAC) preceded by a digital filter. Discrete multitone (DMT) signalling obviates the need for a long FIR FFE and DFE, and has demonstrated better spectral efficiency than 4-PAM above 50 Gb/s. In this work, we consider the potential of DMT for wireline communication beyond 100 Gb/s. For example, a spectral efficiency of 2.5 bits/sample is achievable over an IEEE P802.3ck channel with 18 dB loss at 40 GHz, affording an aggregate data rate of 200 Gb/s at 80 GS/s with 150 fs of jitter and 1.26 mV of noise at the input to the receiver. Significant improvement in bit error rate (BER) is obtained by increasing DAC resolution to 8 bits.

Keywords—Wireline communication, Discrete multitone (DMT)

I. INTRODUCTION

With advances in computing and integrated circuit technologies, the serial data rate of state-of-the-art wireline links now exceeds 100 Gb/s. However, as the data rate increases, channel loss causes the complexity of equalization circuits to become prominent. Thus, pulse amplitude modulation (PAM) signalling with 4 levels has largely replaced non-return-to-zero (NRZ, i.e. 2-level PAM) signalling to improve spectral efficiency. The combination of multilevel modulation and complex channel characteristics necessitates transceivers that largely rely on digital-to-analog and analog-to-digital converters (DACs and ADCs) along with digital signal processing (DSP). However, to operate at such high speeds, parallel processing techniques are required in the DSP that cause the complexity of decision feedback equalization (DFE) to increase geometrically with increases in the number of PAM levels and the number of taps [1]. Moreover, heavy reliance on a DFE can also hinder the performance of forward error correction (FEC) due to the DFE’s propensity for error propagation [2]. A solution to this problem is using an analog-to-digital converter (ADC) based receiver followed by digital signal processing (DSP) [1], [3].




Fig. 1. DMT operation.

Having transitioned to architectures comprising data converters and DSP, wireline transceivers are now, essentially, modems. This evolution opens the door to even more efficient modulation schemes for data rates beyond 112 Gb/s. In particular, discrete multitone (DMT) signalling obviates the need for a long finite impulse response (FIR) feed-forward equalizer (FFE) and DFE, and has demonstrated better spectral efficiency than that of 4-PAM above 50 Gb/s [4]. DMT enables simple equalization without a feedback loop by subdividing a data stream in the frequency domain into multiple sub-channels, and it provides strong immunity to intersymbol interference (ISI) by the insertion of a cyclic prefix (CP). DMT is a popular scheme in optics [5], [6] and telecommunications [7], which have different channel characteristics and data rate requirements than high-speed wireline communication over copper wires; in copper wireline links, interest in DMT has increased only recently. Earlier attempts at implementing such systems in high-speed links cited limited power budgets [8] as an important obstacle. Recent work, however, has demonstrated low-power, high-sampling-rate data converters [9] and low-power DSP applicable to DMT systems [4], [10]. This paper begins with a tutorial on DMT modulation, focussing on the parameters and design choices relevant to high-speed wireline links. We then discuss the impairments that must be accounted for in simulations that accurately evaluate the potential for DMT in these applications including quantization noise and jitter. Finally, we report simulation results quantifying the potential for DMT over a modern standard wireline channel.

II. PAM-4 vs DMT

DMT divides the communication channel into N narrow sub-channels as shown in Fig. 1. (Typically, the DC sub-channel is not used.) Each sub-channel carries an independent sequence of quadrature amplitude modulated (QAM) symbols within its




Fig. 2. (a) PAM4 symbol. (b) DMT symbol.

bandwidth S . Each DMT symbol simultaneously carries in-phase (I) and quadrature (Q) PAM links through the N sub-channels. The symbol rate is equal to the sub-channel bandwidth, S , so the symbol duration is $1/S$. The link operates at the Nyquist sampling rate for the total signal bandwidth, NS ; that is, $2NS$. Thus, there are $2N$ samples per DMT symbol.

A DMT symbol and 4-PAM symbol are compared in Fig. 2. For the same sampling rate, $2NS$, the PAM symbol duration is $1/(2NS)$, as shown in Fig. 2(a). Since each 4-PAM symbol carries two bits, the total data rate is $4NS$. By comparison, the DMT symbol is $2N$ times longer in duration, $1/S$. Assuming 16-QAM over each of the N sub-channels, the number of bits per symbol is $4N$, leading to the same bit rate as the 4-PAM case, $4NS$. If we use constellations larger than 16-QAM, we can achieve a higher data rate using DMT than 4-PAM with the same sampling rate. For example, in [4], 14 out of 15 non-DC sub-channels use 64-QAM to achieve a data rate equal to that of recent 4-PAM designs, while operating at a sampling rate 20% lower. In other words, the spectral efficiency of [4] is 1.25 times higher than that of 4-PAM (2.54 bits per symbol instead of 2 bits per symbol).

III. DMT SYSTEM DESCRIPTION

A block diagram for DMT signalling in wireline applications is shown in Fig. 3. Transmit data at a rate B arises in parallel at the transmitter. Using a bit loading algorithm, we allocate b_k bits to the k^{th} sub-channel depending on sub-channel signal-to-noise ratio (SNR), where $k \in \{1, 2, \dots, N\}$. Thus,

$$B = \sum_{k=1}^{k=N} b_k \quad (1)$$

The b_k bits are mapped to a 2^{b_k} -QAM symbol X_k . The symbols are converted to a time-domain vector using the inverse fast Fourier transform (IFFT). To ensure the resulting output is purely real-valued, we reverse the vector \mathbf{X} and take its conjugate prior to the IFFT. This new vector is then concatenated with the original one. The new, twice longer vector results in a real-valued signal at the IFFT output.

Note that the total number of time-domain samples (the length of the vector \mathbf{x}) is twice the number of sub-channels. To avoid intersymbol interference between consecutive DMT symbols, i.e. to reduce ISI, a cyclic prefix (CP) of length ν is inserted. The required CP length is generally determined by the length of the channel's pulse response (normalized to the sampling time).

Therefore, the total duration of each DMT symbol is $2N + \nu$ samples, and the total data rate is

$$D = \frac{B \cdot F_s}{2N + \nu}, \quad (2)$$

where F_s is the sampling rate of the link. For comparison, 4-PAM signalling achieves a data rate of $2F_s$. Note that in accordance with (1), B generally increases with N . Thus, for $N \gg \nu$, even very long channel pulse responses (large ν) will have a vanishingly small effect on the data rate (2).

Having discussed the DMT transmitter, we note that the reverse operations are required in the receiver. An ADC digitizes samples of the received waveform, which are then parallelized. We then remove the CPs from the DMT symbols and take an FFT to get Y_k . Note that only half of the FFT output is computed because the purely real-valued inputs y_n ensure complex-conjugate symmetry in Y_k . The received data is then equalized by multiplication of each sub-channel's symbol, Y_k , by the reciprocal of the channel response at the sub-channel's center frequency, H_k (i.e. by H_k^{-1}). Since DMT equalization is a simple scalar multiplication of the symbols in each sub-channel, it relies on the assumption that the channel response is relatively constant within each sub-channel. The accuracy of this assumption improves with narrower sub-channels, hence larger N . Finally, the noisy symbols at the equalizer output are quantized (demodulated) and mapped back to the recovered bits.

In summary, we have seen that a large number of sub-channels, N , can mitigate ISI at the expense of added computational complexity. The computation of a length- n (IFFT) (here n is almost twice N) is $O(n \log n)$ and must be performed every DMT symbol $2N + \nu \sim 2N$ samples in duration. Equalization requires only $O(n)$ scalar multiplications every DMT symbol. Thus, the computational complexity per sample increases slowly with N , $O(\log n)$. Considering the exponential growth of equalizer complexity in PAM [11], in some scenarios DMT may offer lower DSP complexity and, hence, power consumption.

IV. SIMULATION MODEL

In this section, we present simulation results using a behavioural model of a DMT wireline link including the finite resolutions and full-scale ranges of the data converters, jitter, and sampling phase errors in the receiver clock.

A. Data Converter Impairments

The finite full-scale ranges of both the transmitter DAC and receiver ADC are significant impairments in DMT links. Being the superposition of many sub-channels modulated by independent random data, DMT waveform samples tend towards a truncated Gaussian distribution with a high peak to average power ratio (PAPR). Avoiding saturation of the DAC and ADC entirely would, therefore, require relatively low root mean squared (rms) signal power and, hence, low SNR. Thus, it is preferable to subject the data converters to higher amplitude DMT waveforms and tolerate some occasional clipping [12]. To quantify the optimal waveform rms amplitude at the DAC and ADC, u_{rms} and v_{rms} , relative to DAC and ADC full-scale




Fig. 3. Block diagram of the DMT system.




Fig. 4. Illustration of the resolutions and full-scale ranges of the data converters.




Fig. 5. Frequency-domain channel response and the corresponding bit allocation. Fig. 6. Introducing clock jitter to the sampled waveform.

ranges, A_{FS} and D_{FS} , (see Fig. 4) we define the input backoff (IBO),

$$IBO_{DA} = 20 \log_{10} \frac{A_{FS}}{v_{rms}} \quad , \quad IBO_{AD} = 20 \log_{10} \frac{D_{FS}}{r_{rms}}$$

To find the optimal IBO_{DA} , the transmit waveform is scaled by a factor α_1 (as shown in Fig. 3). Similarly, a gain factor at the receiver, α_2 , allows for control over IBO_{AD} . Here, data converters are modelled as uniform quantizers, as in Fig. 4.

B. Bit Loading

Whereas [4] used the same 64-QAM constellation in all sub-channels, we perform bit-loading as in [13]. However, we do not employ the power-loading step of [13]; flat power allocation across all sub-channels achieves performance very close to that of optimal power-loading [14]. Fig. 5 shows the channel frequency response and resulting bit allocation for an exemplary channel from an IEEE standard body for 100 Gb/s wireline links (IEEE P802.3ck) [15].

The bit allocation depends on the link SNR, which we define as follows, with reference to the signal and noise rms levels defined in Fig. 3:

$$SNR = 10 \log_{10} \frac{\sigma_v^2}{\sigma_z^2} = 20 \log_{10} \frac{v_{rms}}{\sigma_z} \quad (3)$$

For example, with $SNR = 40$ dB and $v_{rms} = 0.126$ V, σ_z is 1.26 mV.

C. Clock jitter

Jitter was included in the model by randomly perturbing the receiver sample times from their uniformly-spaced ideal times during transmission. The variance of the jitter is denoted σ_j^2 . To accurately incorporate jitter, the simulation must produce the DMT waveform with high time resolution. Since real channel measurements are only available up to 50 GHz [15], it is necessary to upsample the channel's pulse response $10\times$. This was done using an 8th order Butterworth reconstruction filter. A higher upsampling factor would afford more accuracy, but slow the simulations. Instead, finer time steps were achieved by linearly interpolating between neighbouring samples of the $10\times$ upsampled waveforms. The procedure is illustrated in Fig. 6, where an exaggerated σ_j of 0.1 UI¹ (1.25 ps) is used to better show the deviation.

D. Sampling Phase

We also sweep the receiver's nominal sampling phase (around which we introduce jitter). Whereas sampling phase offsets in a PAM receiver hurt timing margin, an offset in the sampling phase of a DMT waveform can be thought of as a rotation of each sub-channel's QAM constellation. Therefore, its impact can be mitigated by appropriate choice of the complex-valued multiplicative constants in the receiver's per-tone equalizer. Note that whereas the waveforms on the channel are upsampled to capture the impact of jitter and sampling phase offset, the model performs all transmit and receive DSP operations on the up- and then down-sampled data to reduce simulation times.

V. SIMULATION RESULTS

This section summarizes some behavioural simulation results of a DMT link using the model described in Section IV. With $N = 255$, each DMT symbol was 532 samples long (512-point IFFT, plus a CP of length $\nu = 20$). The bit-loading algorithm assigned 1321 bits to each symbol, yielding a spectral efficiency of 2.48 bits/sample. With a sampling rate of 80 GS/s, this spectral efficiency results in 198.6 Gb/s. Furthermore, the DAC and ADC full-scale range limits are, respectively, ± 0.5 V and ± 0.2 V. Fig. 7 shows the channel pulse response (sampled at 80 GS/s) in the time and frequency domains. The channel is selected from the IEEE P802.3ck Task Force website [16].

With $SNR = 40$ dB, σ_j was increased from 0 to 0.06 UI and 0.12 UI, and different combinations of transmitter IBO and receiver IBO (IBO_{DA} and IBO_{AD}) were used. The result is shown in Fig. 8. As the figure shows, and since simulations with

¹Unit interval (UI) in this paper refers to the sampling period, i.e. the reciprocal of the sampling rate of the link.




Fig. 7. Channel response. (a) Pulse response (sampling rate 80 GS/s) with the $\nu = 20$ largest terms highlighted in red. (b) Frequency response.




Fig. 8. Changing the IBO at the transmitter and receiver. SNR = 40 dB, DAC and ADC are 7-bit. The three layers of the plot, from top to bottom, correspond to $\sigma_j = 0.0$ UI, 0.006 UI and 0.012 UI, respectively.

other SNR conditions and sampling rates generate similar “bowl-shaped” plots, the combination (12 dB, 12 dB) is a suitable choice for $(\text{IBO}_{\text{DA}}, \text{IBO}_{\text{AD}})$.

With the IBO values fixed at the aforementioned levels and σ_j set to 0.012 UI, the sampling phase was varied, yielding Fig. 9. Also included in this figure is the effect of increasing the sampling rate. For a better comparison, the overall data rates were kept close to 200 Gb/s by reducing the bits per symbol to 1182 and 1064 at sampling rates of 90 and 100 GS/s, respectively. This reduction leads to lower spectral efficiency.

All the results thus far are based on a 7-bit DAC and ADC. Fig. 10 shows the effect of increasing or decreasing the data converter resolutions with a sampling rate of 80 GS/s, SNR = 40 dB, and $\sigma_j = 0.012$ UI. Note that significant improvement in BER is achieved by increasing data converter resolution to 8-bits, particularly in the DAC.

VI. ESTIMATION OF POWER CONSUMPTION

We can estimate the power consumption of such a DMT receiver by extrapolating from other recently-reported works. For example, the ADC in [9] consumes approximately 480 mW when operated at 80 GS/s. In [4], 68 mW of DSP power is reported for a 32-point FFT and frequency-domain equalization. Although the FFT complexity increases in our system, the number of clock cycles allowed for each FFT calculation is also higher. Extrapolating the energy per FFT computation of [4] to a 512-point FFT assuming $\mathcal{O}(n \log n)$ complexity for both the




Fig. 9. Changing the sampling phase of the receiver ADC and the sampling rate. SNR = 40 dB, $\sigma_j = 0.012$ UI, DAC and ADC resolutions are 7-bit.

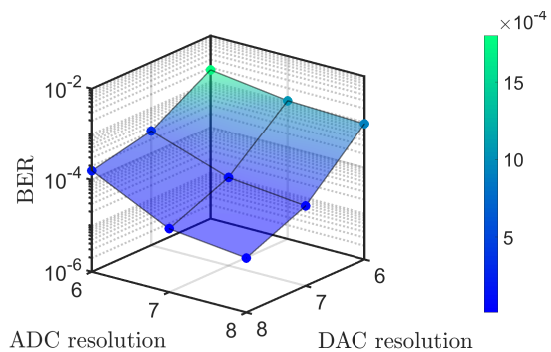


Fig. 10. Changing the resolutions of the transmitter and receiver data converters. SNR = 40 dB and $\sigma_j = 0.012$ UI.

FFT and equalizer yields 472 mW (2.37 pJ/b). Including the ADC, we estimate the energy efficiency of the receiver in our system at 4.8 pJ/b. Note that [4] and [9] use, respectively, 14nm FinFET and 32nm technologies. Further reductions in power may be possible with technology scaling to, for example, the 7nm node.

The estimated 4.8 pJ/b compares reasonably to 3.4 pJ/b in the receiver of [17], 2.9 pJ/b in [4] and 6.6 pJ/b in the receiver (DSP power not included) of [18]. The overall data rate of this work is 199 Gb/s with a channel loss of 21 dB at $1/4$ the bit rate. By comparison, [17], [4] and [18] operate at 56 Gb/s at channel losses of 42, 28 and 31 dB, respectively, at $1/4$ their corresponding bit rates.

VII. CONCLUSION

This paper presented an overview of DMT modulation for high-speed wireline applications. Having identified the potential obstacles to earlier implementations and discussed recent developments, we described a DMT transceiver architecture and behavioural model. Simulation results demonstrate operation at a BER of 10^{-4} or lower with a 7-bit DAC and ADC operating at 80 GS/s, 150 fs of jitter and 1.26 mV of noise at the input to the receiver. Lower BER values were achieved with a lower spectral efficiency at a sampling rate of 100 GS/s, or by increasing data converter resolution to 8 bits, particularly in the DAC.

REFERENCES

- [1] S. Kiran, S. Cai, Y. Luo, S. Hoyos, and S. Palermo, "A 52-Gb/s ADC-based PAM-4 receiver with comparator-assisted 2-bit/stage SAR ADC and partially unrolled DFE in 65-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 54, no. 3, pp. 659–671, Mar. 2019.
- [2] M. Yang, S. Shahramian, H. Shakiba, H. Wong, P. Krotnev, and A. Chan Carusone, "Statistical BER analysis of wireline links with non-binary linear block codes subject to DFE error propagation," *IEEE Trans. Circuits Syst. I, Reg. Papers*, 2019.
- [3] L. Wang, Y. Fu, M. LaCroix, E. Chong, and A. Chan Carusone, "A 64-Gb/s 4-PAM transceiver utilizing an adaptive threshold ADC in 16-nm FinFET," *IEEE J. Solid-State Circuits*, vol. 54, no. 2, pp. 452–462, Feb. 2019.
- [4] G. Kim, L. Kull, D. Luu, M. Braendli, C. Menolfi, P. Francese *et al.*, "A 161mW 56Gb/s ADC-based discrete multitone wireline receiver data-path in 14nm FinFET," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2019, pp. 476–477.
- [5] J. Armstrong, "OFDM for optical communications," *J. Lightw. Technol.*, vol. 27, no. 3, pp. 189–204, Feb. 2009.
- [6] N. Eiselt, D. Muench, A. Dochhan, H. Griesser, M. Eiselt, J. J. Vegas Olmos *et al.*, "Performance comparison of 112-Gb/s DMT, Nyquist PAM4, and partial-response PAM4 for future 5G ethernet-based fronthaul architecture," *J. Lightw. Technol.*, vol. 36, no. 10, pp. 1807–1814, May 2018.
- [7] G. Ginis, "Multi-line coordinated communication for broadband access networks," Ph.D. dissertation, Stanford Univ., Stanford, CA, 2002.
- [8] A. Amirkhany, A. Abbasfar, V. Stojanović, and M. A. Horowitz, "Analog multi-tone signaling for high-speed backplane electrical links," in *IEEE Global Communications Conf. (GlobeCom)*, 2006, pp. SPC03–5:1–6.
- [9] L. Kull, T. Toifl, M. Schmatz, P. A. Francese, C. Menolfi, M. Braendli *et al.*, "A 90GS/s 8b 667mW 64× interleaved SAR ADC in 32nm digital SOI CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2014, pp. 378–379.
- [10] G. Kim, L. Kull, D. Luu, M. Braendli, C. Menolfi, P. A. Francese *et al.*, "Parallel implementation technique of digital equalizer for ultra-high-speed wireline receiver," in *IEEE Int. Symp. Circuits and Systems (ISCAS)*, 2018.
- [11] S. Palermo, S. Hoyos, S. Cai, S. Kiran, and Y. Zhu, "Analog-to-digital converter-based serial links: an overview," *IEEE Solid-State Circuits Mag.*, vol. 10, no. 3, pp. 35–47, Aug. 2018.
- [12] C. H. Azolini Tavares, J. C. Marinello Filho, C. M. Panazio, and T. Abrão, "Input back-off optimization in OFDM systems under ideal pre-distorters," *IEEE Wireless Commun. Lett.*, vol. 5, no. 5, pp. 464–467, Oct. 2016.
- [13] P. S. Chow, J. M. Cioffi, and J. A. C. Bingham, "A practical discrete multitone transceiver loading algorithm for data transmission over spectrally shaped channels," *IEEE Trans. Commun.*, vol. 43, no. 2/3/4, pp. 773–775, Feb. 1995.
- [14] P. S. Chow, "Bandwidth optimized digital transmission techniques for spectrally shaped channels with impulse noise," Ph.D. dissertation, Stanford Univ., Stanford, CA, May 1993.
- [15] R. Mellitz, "100 GEL C2M flyover host files: Tp0 to Tp2, with and without manufacturing variations, for losses 9, 10, 11, 12, 13, and 14 dB," IEEE 802.3 100GEL Study Group - Tools and Channels, 2018. [Online]. Available: http://www.ieee802.org/3/ck/public/18_05/
- [16] IEEE P802.3ck Task Force - Tools and Channels, 2019. [Online]. Available: <http://www.ieee802.org/3/ck/public/tools/index.html>
- [17] M. Pisati, F. D. Bernardinis, P. Pascale, C. Nani, M. Sosio, E. Pozzati *et al.*, "A sub-250mW 1-to-56Gb/s continuous-range PAM-4 42.5dB IL ADC/DAC-based transceiver in 7nm FinFET," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2019, pp. 116–117.
- [18] Y. Frans, J. Shin, L. Zhou, P. Upadhyaya, J. Im, V. Kireev *et al.*, "A 56-Gb/s PAM4 wireline transceiver using a 32-way time-interleaved SAR ADC in 16-nm FinFET," *IEEE J. Solid-State Circuits*, vol. 52, no. 4, pp. 1101–1110, Apr. 2017.