Reduced Complexity Equalization for Coherent Long-Reach Passive Optical Networks

Domaniç Lavery, Benn C. Thomsen, Polina Bayvel, and Seb J. Savory

(Invited Paper)

***Abstract*—Coherent receivers offer a potential solution for implementing a high capacity, long-reach (up to 100 km) passive optical network (LR-PON), due mainly to their high sensitivity, frequency selectivity and bandwidth efficiency. When using coherent receivers, received signals can be post-processed digitally to mitigate the specific impairments found in access networks and, additionally, relax the optical complexity requirements of the coherent receiver. However, the digital signal processing (DSP) must itself be low complexity in order to minimize the overall complexity and power consumption of the optical network unit (ONU).**

This paper focuses on the impact of reduced complexity equalization algorithms on receiver sensitivity in a LR-PON. It is found that a cascade of linear filters can be combined into a single, truncated, linear, adaptive filter with negligible impact on receiver sensitivity. Additionally, by utilizing a multiplier-free tap weight update algorithm, the overall complexity of a digital coherent receiver can be significantly reduced; making it attractive for use in an ONU. Matched filtering, chromatic dispersion compensation, and polarization tracking are all performed by the adaptive equalizer.

The performance of this low complexity, multiplier-free equalizer is experimentally verified for 3 GBd PDM-QPSK (12 Gbit/s) in both a back-to-back configuration, and in transmission over 100 km standard single mode fiber (SMF).

***Index Terms*—Passive optical networks; digital coherent receiver; backreflections; receiver sensitivity; Nyquist pulse shaping.**

I. Introduction

Coherent optical receivers have been revolutionary in core networks, enabling a significant increase in the capacity of long-haul communications. This has been primarily for two reasons. Firstly, phase- and polarization-diverse coherent detection enables the use of advanced modulation formats, such as polarization division multiplexed (PDM) quadrature phase shift keying (QPSK), which encode multiple bits of information per symbol, which in turn improves both optical and electronic bandwidth efficiency. Secondly, and perhaps more significantly for the recent resurgence of coherent optical communications, the use of digital signal processing (DSP) at the receiver can mitigate or completely compensate for many channel impairments (e.g., chromatic dispersion, polarization mode dispersion) without additional, potentially complicated, optical components [1].

Whether or not digital coherent receivers will be equally disruptive in passive optical networks (PON) will depend on the specific advantages of linear detection and DSP in a shorter reach, loss-limited scenario.

Manuscript received July 13, 2014.

The authors are with the Department of Electronic and Electrical Engineering, Optical Networks Group, University College London, London WC1E 7JE, U.K. (e-mail: d.lavery@ee.ucl.ac.uk; p.bayvel@ucl.ac.uk; s.savory@ucl.ac.uk).

One of the primary reasons for investigating coherent receivers for a PON, and the reason they were originally investigated for optical communications, is their excellent power sensitivity at high data rates when compared to envelope (i.e., direct) detection. Near shot noise-limited sensitivity can be achieved simply by including a high power local oscillator laser [2]. As a consequence, networks which are inherently loss-limited will be able to achieve greater capacity and reach when using coherent detection in place of direct detection.

It is not just the sensitivity of coherent receivers which can be exploited in passive networks. The linearity of detection preserves amplitude, phase and polarization information, enabling electronic or digital processing of the received signal. As in core networks, DSP (both at the transmitter and receiver) can be tailored to the channel, thus minimising or negating any sensitivity penalty due to channel impairments (including receiver impairments). Recent examples include compensation for aberrations in the optical front end (e.g., finite common mode rejection ratio [3], intra-receiver skew and quadrature imbalance [4], [5], high phase noise local oscillator lasers [6]), pulse shaping for mitigating reflections [7], [8], and high order modulation for reduced electrical/optical bandwidth requirements [7], [9], [10]. (See [11] for further details.)

Moreover, DSP and electronic signal processing can be used in ways which increase network capacity and flexibility. For example, in a wavelength division multiplexing (WDM) PON, where multiple access is achieved by allocating a wavelength to each optical network unit (ONU), when the local oscillator laser is tuned to the channel of interest, electronic low pass filtering can be used to select this channel [12]. This high frequency selectivity permits a WDM channel spacing on the order of Gigahertz. This high capacity architecture is known as an ultra dense WDM (UDWDM) PON.

For transmission using optical orthogonal frequency division multiplexing (OFDM), DSP can be used to select a subcarrier, or group of subcarriers, to allow dynamic bandwidth allocation with sub-wavelength granularity [13]. On the carrier side, the optical line terminal (OLT), it has been demonstrated that subcarriers can be jointly generated [14] and detected [15] to reduce and simplify network equipment requirements.

Finally, note that coherent receivers also permit the use of both polarization modes for transmission, which can further increase network capacity. However, in order to use both polarization modes, the signal state of polarization must be continuously tracked. Although optical polarization tracking methods exist, it is significantly simpler to use DSP for this purpose. Specifically, adaptive multiple input multiple output (MIMO) equalization has been used for polarization tracking in both core and access network scenarios [1], [12].

Although the use of DSP is a powerful technique for compensating channel impairments, it introduces an addi-

tional degree of complexity to the receiver. As noted in [11], although DSP can be used to reduce the optical complexity of the coherent receiver, it must itself be low complexity to be practical for implementation in a PON; where cost and power consumption are critical parameters.

Therefore, this paper seeks to use reduced complexity adaptive equalization algorithms to trade off receiver performance and complexity. We extend our earlier work in [11], [16], [17] to investigate the possibility of combining multiple linear filters into a single adaptive filter; namely as part of the equalizer. The trade off between receiver sensitivity and complexity is explored in the scenario where low complexity equalizers are used to track the signal state of polarization and simultaneously apply matched filtering.

The remainder of this paper is organized as follows. Section II outlines the PON architecture being considered in this work, and identifies the particular system parameters where MIMO equalization can be used to reduce overall receiver complexity. Section III details the equalization schemes considered herein; namely 2x2 complex MIMO (the conventional approach), 2x2 complex Sign-Data (signed data) MIMO, and 2x2 complex Sign-Sign (signed data and signed error) MIMO, with the latter two examples being reduced complexity implementations. Here, the complexity of each algorithm is compared (full details in Appendix A. Section IV outlines numerical simulations used to evaluate the performance of these equalizers under conditions which might be observed in a PON; shot noise, fast polarization rotations, and variable residual dispersion. This section also evaluates the impact of a reduced update rate on performance. Additionally, the performance of these algorithms is evaluated where they are used in lieu of a matched filter. This section concludes with an experimental verification of these algorithms in a loss-limited PON scenario. Section V describes the implications for a specific network scenario where these algorithms could be used; namely the long-reach coherent PON. Conclusions are drawn in section VI.

II. Coherent Long-Reach Passive Optical Networks

For economic reasons, it has been proposed that the reach of a conventional access network be extended (up to 100 km) to encompass a wider geographical area [18]. This would enable consolidation of real estate, and potentially simplify the network, overall, as more ONUs (and thus subscribers) can be served from a single OLT. Ideally, in such a network, active networking equipment would be removed from intermediate nodes.

Also for economic reasons, it is desirable for transmission in a PON to be bidirectional. In a conventional time-division-multiplexed (TDM) network architecture, such as the GPON [19], [20], the downstream and upstream transmission wavelengths are separated by approximately 200 nm (in the case of the GPON, the wavelengths are in the 1500 nm and 1300 nm bands, respectively). Transmitting bidirectionally in different transmission windows limits any interaction between upstream and downstream. It was recently proposed that, by taking advantage of WDM, the capacity of the access network could be further increased. This creates an issue for bidirectional transmission, as upstream (ONU to OLT) channels will be reflected into the downstream (OLT to ONU) transmission path, thus causing linear crosstalk at the ONU [8].

Consider the long-reach PON (LR-PON) network architecture shown in Fig. 1. Here it is assumed that the reach of the access network is extended in order to encompass the metropolitan area network (up to approximately 100 km) and therefore serve a greater number of ONUs from a single OLT.

The initial proposals for long reach PONs included an amplifier at a remote node, immediately before the first passive power splitter, in order to overcome splitter loss. The transmission scheme employed was intensity modulation with direct detection at each ONU. To overcome the loss of the splitter, recent proposals have incorporated WDM, and an arrayed waveguide grating at the remote node to direct a single wavelength to each ONU. However, this then limits reconfigurability, as each ONU must operate bidirectionally at a specific wavelength. Having a “colored” ONU also creates problems for inventory and deployment, as all ONUs are different. A number of solutions to these issues have been proposed (e.g., [21]).

An alternative approach, and the approach considered here (Fig. 1), is to maintain a passive power splitter, and use a coherent receiver at each ONU. The frequency selectivity of the coherent receiver (through retuning the local oscillator laser to the wavelength of interest) allows the network itself to remain “colorless”, with filtering performed electronically at the ONU. Indeed, using a coherent receiver for wavelength selection permits the use of UDWDM. Coherent receivers offer the additional benefits of high sensitivity, and reduced optical/electronic bandwidth requirements due to the use of advanced modulation formats (e.g., PDM-QPSK, m-*ary*

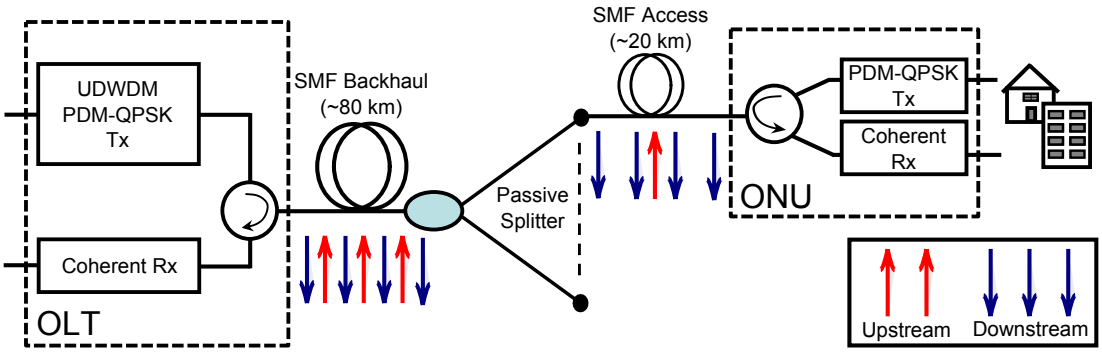

Fig. 1. Schematic of a coherent long-reach PON architecture. The arrows illustrate the upstream/downstream offset frequency plan which must be used to avoid crosstalk due to backreflections.

quadrature amplitude modulation (QAM)). As previously noted, the disadvantage of this approach is the additional optical complexity and electrical complexity requirements over, for example, a direct detection receiver.

There is the possibility to reduce the complexity of the coherent ONU by using only a single laser. The power from this laser can be divided, with one source used as the LO, and the other remodulated at an offset frequency for use as the upstream signal[1]. The reason for the frequency offset is to overcome crosstalk in the downstream signal path due to reflections and Rayleigh backscattering of the upstream channel [23] (see Fig. 1). It has previously been shown that applying a root-raised-cosine (RRC) filter to the transmitted signal can further mitigate the impact of crosstalk due to reflections [8], [7], [24], albeit for the additional DSP complexity required by matched filtering.

For matched filtering implemented as a finite impulse response (FIR) filter, the number of filter taps required to detect with negligible penalty grows with the inverse of the optical bandwidth occupied by the signal. However, even though the ideal filter response is infinite, practically the matched filter can be truncated for a small penalty [25][2]. Accepting a this small penalty, the optical bandwidth requirements and the DSP complexity can be significantly reduced.

It is noted that, in a network such as that shown in Fig. 1, an adaptive equalizer can be included to track the received signal state of polarization. In [5], [25], the blind adaptive equalizer was shown to perform the function of a matched filter. In this paper, the motivation is to extend this approach to include low complexity adaptive equalizer tap weight update algorithms, such that the matched filter can be omitted while still reducing the overall computational complexity of the DSP used in the ONU. These algorithms are described in the following section.

III. Low Complexity Adaptive Equalization

In order to reduce the overall power consumption of a digital coherent receiver, it is desirable to reduce the complexity of receiver DSP. Aside from forward error correction (FEC), the adaptive equalizer is potentially the most power consuming block of the DSP. This is because, in a direct implementation, the filter tap weight update algorithm requires approximately the same computational resources as the filter itself. This section examines some well-known algorithms for reduced complexity equalizer filter tap weight updates, including an implementation of multiplier-free tap weight updates, and details how these algorithms can be used for a MIMO finite impulse response (FIR) filter. The adaptive MIMO FIR filter algorithm considered (assumed hereafter to be T/2-spaced), Fig. 2, is well-documented (see for example, [26]), however it is summarised here for completeness. The first stage is to apply four parallel complex-valued filters to the signal as follows

$$\begin{aligned} x_{out}(k) &\leftarrow \mathbf{h}_{xx}^H \mathbf{x}_{in} + \mathbf{h}_{xy}^H \mathbf{y}_{in} \\ y_{out}(k) &\leftarrow \mathbf{h}_{yx}^H \mathbf{x}_{in} + \mathbf{h}_{yy}^H \mathbf{y}_{in} \end{aligned} \tag{1}$$

[1]If using a heterodyne receiver as in [22], then there is no requirement to modulate at an offset frequency as the LO frequency is already offset from the signal.

[2]Although, in this work, the motivation was to avoid this sensitivity penalty by including the matched filter as part of a many-tap chromatic dispersion compensating filter.

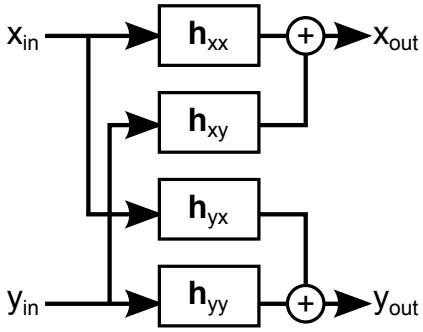

Fig. 2. Structure of the 2x2 MIMO equalizer.

where, in general, $\mathbf{h}_{out,in}$ is a filter, of length N taps, with subscripts indicating the input and output polarization states, superscript H indicates the element-wise complex conjugate transpose (Hermitian conjugate) of the equalizer tap weights in equation (1), and the signal input vectors are $\mathbf{x}_{in}$ and $\mathbf{y}_{in}$ for the X and Y polarizations, respectively. Finally, $x_{out}(k)$ and $y_{out}(k)$ are the k$^{\text{th}}$ filtered output samples for the X and Y polarizations, respectively. The filters are applied as a time-domain convolution, i.e.,

$$\mathbf{h}_{xx}^H \mathbf{x}_{in} = \sum_{n=0}^{N-1} h_{xx}^*(n) x_{in}(k-n) \tag{2}$$

The tap weight update algorithm considered is the constant modulus algorithm (CMA). Therefore, the output sample, x_{out} (k$^{\text{th}}$ sampling instant implied hereafter), is used for the update algorithm by calculating the deviation from a constant modulus

$$\begin{aligned} e_x &= 1 - |x_{\text{out}}^2| \\ e_y &= 1 - |y_{\text{out}}^2| \end{aligned} \tag{3}$$

where e_x and e_y are the errors on polarizations X and Y, respectively.

The filter and error calculation algorithms are common to each of the adaptive filters considered. The difference is in the tap weight update algorithm, which is detailed in the following subsections.

In order to analyze the complexity of the filtering algorithm, there are several repeated operations which are now highlighted. Consider the filter itself; $\mathbf{h}_{xx}^H \mathbf{x}_{in}$. Both the filter vector and input vector are complex numbers, and so computing this multiplication requires N complex multiplications for an N-tap filter. For each polarization there are two such filtering operations, leading to a computational complexity of 4N complex multiplications[3] (16N real multiplications).

The complexity of the error term calculation depends on the modulation format being equalized, however for PDM-QPSK the total complexity of this stage (both polarizations) is only four real multiplications, which is negligible for N $>$ 1. Therefore the complexity of the error term computation is neglected in the following (full analysis in Appendix A).

[3]For a large value of N, it is possible to reduce the required complex multiplications by using Fourier domain convolution, but there is no gain for N $<$ 8 [27]

A. Conventional Implementaiton of Tap Weight Updates

The conventional implementation of the tap weight update algorithm is as follows [1]

$$\begin{aligned}\mathbf{h}_{xx} &\leftarrow \mathbf{h}_{xx} + \mu e_x \mathbf{x}_{in} x^*_{out} \\ \mathbf{h}_{xy} &\leftarrow \mathbf{h}_{xy} + \mu e_x \mathbf{y}_{in} x^*_{out} \\ \mathbf{h}_{yx} &\leftarrow \mathbf{h}_{yx} + \mu e_y \mathbf{x}_{in} y^*_{out} \\ \mathbf{h}_{yy} &\leftarrow \mathbf{h}_{yy} + \mu e_y \mathbf{y}_{in} y^*_{out}\end{aligned} \tag{4}$$

where μ is a learning parameter that determines the speed of equalizer convergence, and the rate at which changes in the signal can be tracked.

The complexity of the tap weight update stage can be computed as follows. First, note that the perturbation to the tap updates are scaled by the learning parameter, μ. By restricting this parameter to be a power of 2, a hardware 2's complement implementation of multiplication by μ can be implemented with a simple bit shift [28]. Note also that the gradient term (the multiplication between the error and the output sample, e.g. $e_x x_{out}$) can be computed before multiplication with the input vector (a more costly operation). The error term is a real number while the output symbol is complex, thus computing the gradient for each filter requires 2 real multiplications (8 real multiplications in total). Finally, the gradient term can be multiplied by the input vector, requiring 4N complex multiplications. These 4N complex multiplications dominate the complexity of the tap weight update algorithm.

It is interesting to note that the order of magnitude complexity of the filtering and the filter adaptation are both[4] $\mathcal{O}(N)$, meaning that the tap updates contribute approximately half the adaptive equalizer complexity using the direct implementation of convolution and, potentially, more than half the equalizer complexity when implementing convolution in the frequency domain, where the complexity is $\mathcal{O}(\log_2(N))$ (for large N).

B. Signum Tap Weight Update Algorithms

It has been known for some time that equalizer tap weight updates can be simplified by discarding some, or all, of the information about the magnitude of the gradient term, keeping only the sign of the gradient [29][5]. In order to find the sign of the gradient, the signum function is required, defined as follows

$$\mathrm{sgn}(a) = \begin{cases} -1 & \text{if } a < 0 \\ 0 & \text{if } a = 0 \\ 1 & \text{otherwise} \end{cases} \tag{5}$$

where a is a real number. A complex signum function, based on signum by treating the real ($\Re$) and imaginary ($\Im$) components of a complex number, z, independently, is defined as follows

$$\mathrm{csgn}(z) = \mathrm{sgn}(\Re(z)) + j\,\mathrm{sgn}(\Im(z)) \tag{6}$$

The following two algorithms are for blind-adaptive filter tap weight updates, based on this principle.

[4]Big O notation indicating, in this case, that the complexity in terms of multiplications is linear.

[5]Note that similar algorithms have previously been used to simplify receiver implementations in hardware for both long-haul [30] and access networks [16], [31] research.

1) Sign-Data Tap Weight Update Algorithm: The initial simplification of the conventional tap weight updates can be made by first taking the complex sign of the output data symbol. Note that this simplification is justifiable because the expected value of the signal modulus is exactly unity. However, as in the conventional implementation of the update algorithm, the error term contains only amplitude information. Information about the quadrature of the data is still required; the csgn operation retains some of this information, as the signal is restricted to four equidistant points (in a Euclidean sense) on a unit circle: $\pi/4$,$3\pi/4$,$5\pi/4$ and $7\pi/4$. The algorithm is detailed below.

$$\begin{aligned}\mathbf{h}_{xx} &\leftarrow \mathbf{h}_{xx} + \mu e_x \mathbf{x}_{in}\,\mathrm{csgn}(x^*_{out}) \\ \mathbf{h}_{xy} &\leftarrow \mathbf{h}_{xy} + \mu e_x \mathbf{y}_{in}\,\mathrm{csgn}(x^*_{out}) \\ \mathbf{h}_{yx} &\leftarrow \mathbf{h}_{yx} + \mu e_y \mathbf{x}_{in}\,\mathrm{csgn}(y^*_{out}) \\ \mathbf{h}_{yy} &\leftarrow \mathbf{h}_{yy} + \mu e_y \mathbf{y}_{in}\,\mathrm{csgn}(y^*_{out})\end{aligned} \tag{7}$$

The complexity of this algorithm is now significantly reduced (see Table I) for two reasons. First, the multiplication between the error term and the input vector requires only $2N$ real multiplications, because the error term is a real number. Second, and most significantly, the product of this intermediate result and the sign of the output symbol requires only addition operations.

2) Sign-Sign Tap Weight Update Algorithm: This approach can be further extended by taking the sign of both the error term and the output data symbol, such that only the sign of the gradient is used to update the filter taps. This is possible because incrementally updating the filter taps using the sign of the gradient will still lead to minimisation of signal error.

$$\begin{aligned}\mathbf{h}_{xx} &\leftarrow \mathbf{h}_{xx} + \mu\,\mathrm{sgn}\,(e_x)\mathbf{x}_{in}\,\mathrm{csgn}(x^*_{out}) \\ \mathbf{h}_{xy} &\leftarrow \mathbf{h}_{xy} + \mu\,\mathrm{sgn}\,(e_x)\mathbf{y}_{in}\,\mathrm{csgn}(x^*_{out}) \\ \mathbf{h}_{yx} &\leftarrow \mathbf{h}_{yx} + \mu\,\mathrm{sgn}\,(e_y)\mathbf{x}_{in}\,\mathrm{csgn}(y^*_{out}) \\ \mathbf{h}_{yy} &\leftarrow \mathbf{h}_{yy} + \mu\,\mathrm{sgn}\,(e_y)\mathbf{y}_{in}\,\mathrm{csgn}(y^*_{out})\end{aligned} \tag{8}$$

Here, multiplications are not required. The gradient can be computed by inverting the sign of $\mathrm{csgn}(x_{out})$ based on the sign of $\mathrm{sgn}\,(e_x)$. Further, multiplying the gradient with the input vector requires only addition, as noted previously.

C. A Comparison of Implementation Complexity

A comparison of the complexity of each of the tap weight update algorithms is shown in Table I in terms of real (as opposed to complex) additions and multiplications. Interestingly, although the Sign-Data algorithm would reduce the number of total number of multiplications required in a hardware implementation, the complexity is still $\mathcal{O}(N)$ (in terms of multiplication operations). The Sign-Sign algorithm offers the lowest complexity tap weight update algorithm, assuming multiplications dominate computational complexity.

An extended computational complexity analysis is given in Appendix A. The next section describes simulations and experimental investigations of the simplified algorithms described above.

IV. Simulation and Experimental Investigation

For the following simulations, note that the results are presented independently of symbol rate, but the subsequent analysis is for a 3 GBd PDM-QPSK signal. This symbol rate assumes that the data rate is 10 Gbit/s plus 20% overhead for

TABLE I
COMPLEXITY (ADDITIONS AND MULTIPLIPLICATIONS) FOR EQUALIZER TAP WEIGHT UPDATE ALGORITHMS

	Conventional	Sign-Data	Sign-Sign
Real Multiplications	16N+4	8N	0
Real Additions	16N	16N	16N

FEC. The target bit error rate (BER) for this FEC overhead is 1.5×10^{-2}, assuming a hard decision FEC decoder [32].

A. Approximating the Matched Filter Response

The performance for all three algorithms was first evaluated in simulation. Ideal root-raised-cosine (RRC) pulse shapes (with a roll-off factor, α, between 0 and 1, inclusive) were applied to PDM-QPSK signals. For each data point, 2^{20} symbols were simulated, where each symbol was drawn from a random distribution. The receiver sensitivity was then determined when using each of the aforementioned adaptive equalizers to adaptively converge to the matched filter response. The ideal matched filter has an infinite response, and so the adaptive equalizers necessarily produce a truncated matched filter. The performance of the adaptive equalizers is shown in Fig. 3, for different roll-off factors and different numbers of filter taps.

It can be seen from Fig. 3 that the matched filter response is approximated by the adaptive equalizer with very few taps relative to the ideal matched filter. Indeed, when using more than 5 equalizer taps, the sensitivity penalty is below 1 dB for all roll-off factors; including a "brick wall" filter ($\alpha = 0$). Notably, this result is independent of the equalizer tap weight update algorithm.

In order to maintain a sensitivity penalty below 0.5 dB when using a 5-tap equalizer, the maximum RRC filter roll-off is $\alpha = 0.1$. This is an acceptably small sensitivity penalty, and so this is the scenario which will be considered in the following sections.

B. Tracking Fast Polarization Rotations

It is conventional to use polarization rotations to evaluate the equalizer's tracking performance [1]. Indeed, as polarization rotations will always occur in transmission over optical fiber, it is important to estimate the maximum polarization rotation rate which can be tracked.

To investigate this, continuous polarization rotations were added to the simulated RRC filtered signal, prior to a balanced, phase- and polarization-diverse coherent receiver. In order to make a fair comparison between the equalizers, each was initialised using ideal tap weights, and each equalizer filter length was 5-taps. The results for the three equalizers are shown in Fig. 4.

As shown in Fig. 4(a), all three algorithms can track polarization rotations 5×10^{-4} rad/s/Hz. For a symbol rate of 3 GBd, the corresponding maximum polarization rotation rates were 3.3×10^6 rad/s (Conventional), 2.8×10^6 rad/s (Sign-Data), and 2.0×10^6 rad/s (Sign-Sign).

Note that the signum operation applied to the output data and/or the error term increases the effective magnitude of the learning parameter, μ, leading to the maximum permissible value of μ to be lower for both the Sign-Data and the Sign-Sign algorithms.

In a hardware implementation, it is desirable to implement the equalization in parallel (i.e., M parallel FIR filters

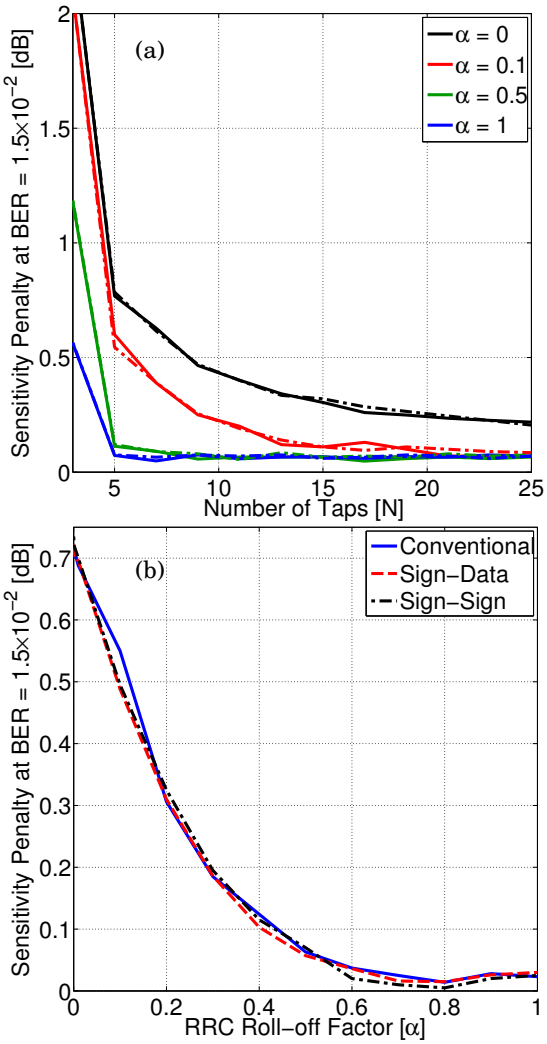

Fig. 3. Receiver sensitivity (at a BER of 1.5×10^{-2}) evaluated through numerical simulations when using an adaptive equalizer to estimate the matched filter response to a root-raised-cosine filtered PDM-QPSK signal. The receiver sensitivity is evaluated for different filter lengths, and three update algorithms: Conventional, Sign-Data, and Sign-Sign. (a) Sensitivity penalty at as a function of number of taps for different RRC roll-off factors. Conventional (solid line), Sign-Sign (dashed line), Sign-Data (equivalent performance so not shown for clarity). (b) Sensitivity penalty relative to performance using an ideal matched filter for different RRC roll-off factors.

for an M symbol bus), meaning that the equalizer can only be updated once every bus. If the equalizer is updated using the gradient term from only one of the parallel filters, then the maximum polarization tracking rate is simply reduced by the degree of parallelism. However, if the equalizer is updated using the average (or sum) of the gradients over a single bus, then some of the tracking efficacy can be recovered. (These two parallel architectures are considered in [16].)

Fig. 4(b) shows the impact of a parallel update algorithm on tracking efficacy, where the tap weight update is the sum of tap weight updates computed over a single bus. It can be seen that, with increasing parallelism, there is not a significant degradation in maximum tolerable polarization rotation rate. For example, for a 3 GBd PDM-QPSK signal, the best and worst performance for these algorithms is 3.3 Mrad/s and 2 Mrad/s (symbol-by-symbol update, $M = 1$)

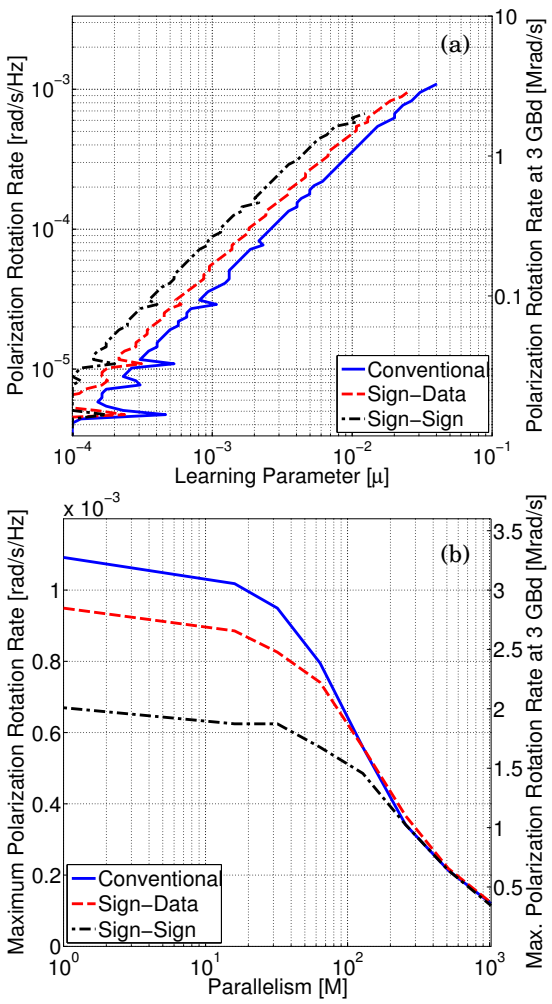

Fig. 4. Simulation results for a 5-tap adaptive equalizer tracking polarization rotations applied to a Nyquist-shaped (RRC $\alpha = 0.1$) PDM-QPSK signal. The signal was noise loaded to an SNR 1 dB above the sensitivity limit (shot noise limit plus 0.5 dB as shown in Fig. 3). Polarization rotation rate is normalized to the symbol rate, and evaluated for 3 GBd on the second y-axis. (a) Ability to track polarization rotations (performance measured at a 1 dB receiver sensitivity penalty). (b) Maximum tolerable polarization rotation rate when equaliser is updated on a bus-by-bus basis (degree of parallelism is equivalent to number of symbols in a bus).

and 1.1 Mrad/s and 1.0 Mrad/s (parallelism $M = 256$). For $M = 256$, this outperforms, by two orders of magnitude, the scenario where only one filter output is used to update the equalizer tap weights.

Of greater significance is the relative performance of the equalizer tap weight update algorithms when the parallelism is increased. With a parallelism of $M = 256$, the performance of the algorithms has converged, indicating that averaging (or in this case, summing) the parallel equalizer update gradients can mitigate against the uncertainty introduced by only using the sign, rather than the magnitude, of the error and/or the received signal.

The rate of polarization rotations in a 100 km span will be strongly dependent on the physical environment of the fiber. However, a worst case scenario is the polarization rotations induced when the fiber is exposed to physical shock. In [33], a 20 km dispersion compensating module was exposed to physical shock, which was found to induce a maximum polarization rotation rate on the order of 10^5 rad/s (see also [34]). The results in Fig. 4 indicate that, for a system operating at 3 GBd, rotations up to 3.5×10^5 rad/s could be tracked, for all equalizer tap weight update algorithms considered herein, for a degree of parallelism up to and including the highest degree of parallelism considered, $M = 1024$. Therefore, up to a parallelism of $M = 1024$, it is not expected that fast polarization rotations would present an issue for adaptive equalization.

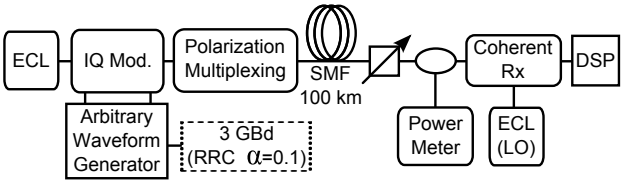

Fig. 5. Experimental configuration to investigate DSP performance in a LR-PON. The variable optical attenuator immediately after the 100 km fiber span is used to emulate passive splitter loss.

C. Experimental Verification

The numerical simulations presented thus far are quite general; there is no assumption about symbol rate and, consequently, no consideration of transmission. Therefore, to validate the simulation results in the context of a PON, these equalizers were also investigated experimentally; firstly in a back-to-back configuration to evaluate the performance of each equalizer and, subsequently, in transmission over 100 km SMF, to evaluate the impact of the residual chromatic dispersion typical of an access network.

The experimental configuration is shown in Fig. 5. To generate the QPSK signal, $2^{15}-1$ symbols were generated using two decorrelated $2^{15}-1$ pseudorandom binary sequences, and applied using an arbitrary waveform generator (AWG) to an optical carrier at 1550 nm (20 kHz linewidth) via a nested Mach-Zehnder Modulator (denoted 'IQ Mod.' in Fig. 5). The AWG was operated at 12 Gsamples/s, providing an oversampling rate of 4 samples/symbol. Where RRC filtering (roll-off $\alpha = 0.1$) was applied, the signal was precomputed offline using an ideal frequency domain filter.

Dual polarization transmission was emulated using a polarization multiplexing emulation stage. This consisted of a 3 dB splitter, to divide the modulated signal into two paths with a relative delay of 55 symbols. The signals were then recombined in orthogonal polarization states using a polarization controller in each arm, followed by a polarization beam combiner.

Where transmission was investigated, the signal was launched with a power of 0 dBm, and propagated over 100 km SMF, with a total loss of 20 dB (0.2 dB/km). In both the transmission and back-to-back configurations, a variable optical attenuator (VOA) was used to emulate the power loss associated with the passive splitter in the PON. Finally, the optical signal was detected using an integrated phase- and polarization-diverse coherent receiver with balanced photoreceivers (nominal signal common mode rejection ratio 16 dB). The local oscillator laser, an ECL with linewidth 100 kHz, had its optical output power set to +15 dBm measured at the input to the coherent receiver. The coherent receiver output signal was sampled using a digital sampling oscilloscope at 50 Gsamples/s and subsequently resampled to 2 samples/symbol.

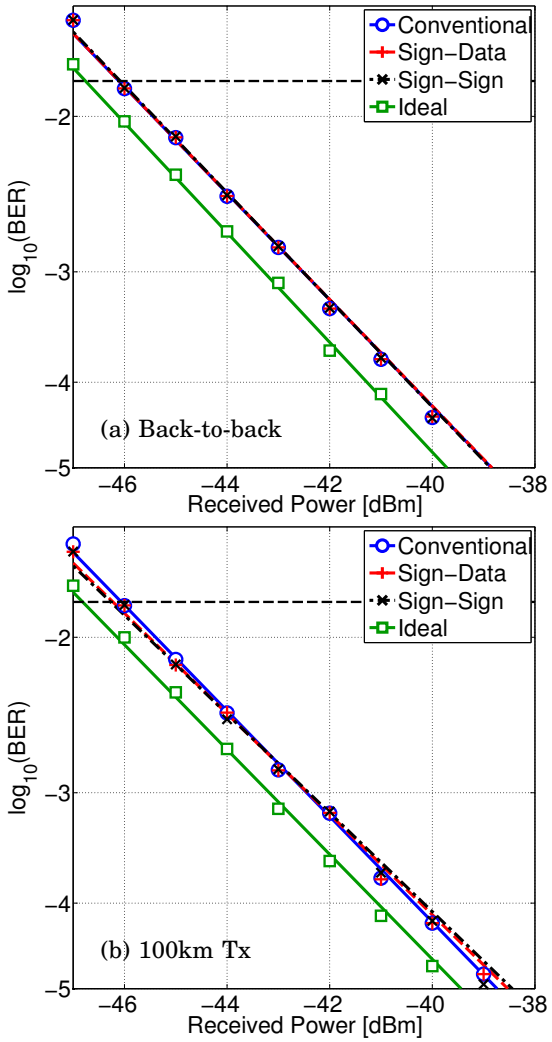

Fig. 6. Experimental investigation of three equalizer tap weight update algorithms: conventional, sign of output data (Sign-Data), and sign of output data and error term (Sign-Sign). The transmitted signal was Nyquist shaped (RRC roll-off $\alpha = 0.1$) 3 GBd PDM-QPSK. In the Ideal case, an ideal matched filter was applied at the receiver. In all other cases, the 5-tap equalizer was used to approximate the matched filter response. The horizontal dashed line indicates the assumed FEC threshold; a BER of 1.5×10^{-2}. (a) Back-to-back configuration, and (b) transmission over 100 km SMF.

The DSP chain used in the experimental work was as follows. The signal was normalised to unit power prior to equalization using the algorithms detailed herein. A 4^{th} power frequency domain algorithm was used to determine the frequency offset between signal and LO [35], and the Viterbi and Viterbi algorithm was used for carrier phase estimation [36]. Hard decisions were made prior to bit error rate estimation.

The results of the back-to-back experimental investigation are shown in Fig. 6(a). In order to determine the benchmark receiver sensitivity, the signal was first processed using an ideal matched filter (in this case RRC roll-off $\alpha = 0.1$). The achieved receiver sensitivity was -46.8 dBm at a BER of 1.5×10^{-2}. As expected from simulations, there is a sensitivity penalty of approximately 0.5 dB at this BER when using a 5-tap adaptive equalizer to approximate the matched filter response.

Comparing Figs. 6(a) and (b), it can be seen that all three algorithms exhibit identical performance in both the back-to-back and transmission scenarios, meaning that chromatic dispersion is also accurately compensated by a 5-tap equalizer for this transmission distance and symbol rate.

Though simulations showed that the conventional and Sign-Data algorithms offer superior performance when tracking fast rotations, the Sign-Sign algorithm is sufficient for the impairments encountered in 100 km transmission over spooled SMF. Indeed, for a practical (parallel) implementation in hardware, the performance difference between the algorithms is negligible. The rate of polarization rotations could be greater in installed fiber although, when taken with the simulation results in Fig. 4, it would appear that, of the algorithms considered, the lowest complexity algorithm of the three considered, Sign-Sign, is the better candidate for equalization in a digital coherent receiver in an access network scenario.

V. Implications for a Long-Reach Passive Optical Network

The access network architecture outlined in section II is fundamentally loss limited; that is, the shot noise introduced at the receiver is normally the dominant noise source. However, a sensitivity penalty can be incurred when using low complexity optical components [3], when launching a large signal power due to fiber nonlinearity [37], when operating the receiver in the presence of additional noise sources (e.g., upstream channel reflections [7]) and, indeed, reduced complexity DSP. Thus, considering the above, this section discusses the implications of low complexity DSP on the sensitivity of the coherent receiver, and the implications on the capacity and reach of a LR-PON.

Consider, first, a unidirectional implementation of the network shown in Fig. 1 (i.e., where a fiber pair is used for bidirectional transmission). This network offers a best case scenario in terms of downstream capacity, as all bandwidth is available for modulation. By using the experimental sensitivity limit obtained from Fig. 6, the maximum split ratio and reach for a PON can be estimated. This result is shown in Fig. 7.

Assuming a launch power of 0 dBm per WDM channel, the power budget with an ideal receiver would be sufficient to support transmission over 100 km with a 1:256-way passive split; this conclusion does not change if a 0.5 dB sensitivity penalty is accepted for using reduced complexity DSP. If a further 1 dB margin is required to tolerate fast polarization rotations, section IV, then the achievable split ratio would be reduced to 1:128. If the launch power per channel were increased to a more moderate 5 dBm, then the achievable split ratio at 100 km would be 1:1024 in all three scenarios[6].

As previously noted, fiber nonlinearity can limit the achievable receiver sensitivity, and this effect is power dependent. In a UDWDM PON, the launch power per channel will be limited [7], and so this analysis provides an upper bound on performance. Nevertheless, this result indicates that the relative penalty for reduced complexity DSP will not be significant.

[6]Note that a moderate signal launch power can be achieved using lasers monolithically integrated with semiconductor optical amplifiers, for example [38]. Such lasers are suitable for volume production are therefore have the potential to satisfy the low cost requirement of access networks.

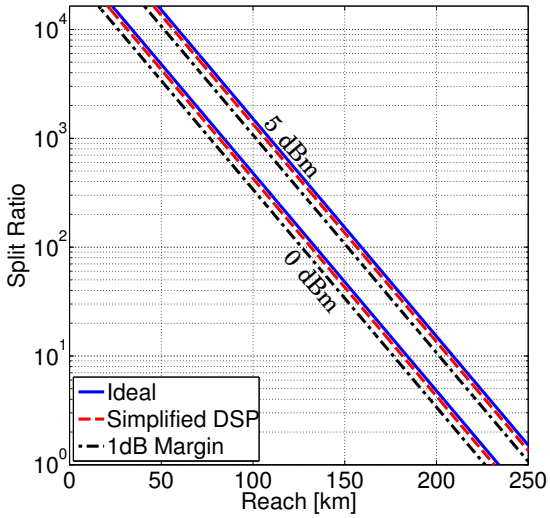

Fig. 7. Estimated achievable reach and split ratio based on the achieved receiver sensitivity for the experimental investigation presented herein. The assumed receiver sensitivities are -46.8 dBm and -46.3 dBm for the ideal matched filter and reduced complexity (5-tap equalizer only) scenarios. The reach for an additional 1 dB sensitivity margin is also indicated, which allows for polarization rotations up to 1 Mrad/s. Assumed launch powers are 0 dBm and 5 dBm as indicated.

These results also have an implication in bidirectional (full duplex) transmission over a single optical fiber. Assuming the channels can be ideally filtered with a RRC roll-off factor of $\alpha = 0.1$, the minimum channel spacing to avoid linear crosstalk is 110% of the symbol rate; in this case 3.3 GHz. It has previously been shown that the receiver sensitivity penalty due to reflections can be completely mitigated when using Nyquist pulse shaping with sufficient offset between upstream and downstream channels [8], [10]. As a consequence, the upstream/downstream channel spacing could be reduced to be asymptotically close to 3.3 GHz, with no additional receiver DSP complexity, and no sensitivity penalty. The implication for other pulse shapes and equalizer lengths can be determined from Fig. 3.

VI. Conclusion

Transceiver complexity is a key concern for access networks, particularly for the ONU. In access networks incorporating digital coherent receivers, the optical, electrical and digital complexity must be minimised in order to reduce both manufacturing costs and operational power consumption. To this end, this paper focused on reducing the complexity of the receiver DSP.

For the specific case of an access network, the adaptive equalizer was identified as being a source of significant complexity, the majority of which was confined to the adaptive equalizer tap weight update algorithm. Two reduced complexity tap weight update algorithms were identified, one of which (Sign-Sign) was multiplier-free. They were tested for a PDM-QPSK signal, and found to operate without sensitivity penalty, both in simulations and experimentally in loss-limited transmission over 100 km standard single mode fiber; a distance commensurate with a LR-PON.

It was noted that restricted bandwidth channels are of interest in UDWDM PON scenarios, both for minimising crosstalk due to backreflections and also for maximising the transmission capacity of the PON. Conventionally, this is achieved by applying a root-raised-cosine pulse shape to the transmitted signal, and applying an equivalent matched filter at the receiver. The filters are applied using DSP. Here, a further simplification to the DSP was proposed, where the reduced complexity blind adaptive equalizer was used to simultaneously track the signal state of polarization and apply the matched filter response.

It was found that, accepting a 0.5 dB sensitivity penalty, the matched filter could be omitted when the signal was shaped using a root-raised cosine filter with a roll-off $\alpha = 0.1$ and a 5-tap adaptive equalizer was employed at the receiver.

The equalizer was also evaluated in terms of fast tracking efficacy. For each of the tap weight update algorithms considered, the rate at which polarization rotations could be tracked was determined, allowing a 1 dB sensitivity penalty. It was found that, even if the algorithms were implemented with a high degree of parallelism, the tracking ability of each algorithm exceeds the rotation rate expected when the optical fiber is subject to mechanical disturbance.

Experimentally, for 3 GBd (12 Gbit/s) PDM-QPSK signals, the reduced complexity receiver sensitivity was determined to be -46.3 dBm. Allowing a 1 dB sensitivity margin for tracking fast polarization rotations, and assuming 0 dBm per channel launch power and linear transmission, the receiver sensitivity would be sufficient to allow 100 km transmission with a 1:128-way passive split. Assuming an increased launch power of 5 dBm per channel, the split ratio could be increased to 1:1024.

These results indicate that a reduced complexity adaptive equalizer (with a multiplier-free tap weight update algorithm) could be used to achieve the long transmission distance and high split ratio required in a LR-PON, while operating with a significantly reduced complexity compared with the conventional approach.

Appendix A
Computational Complexity of Equalizer Tap Weight Update Algorithms

The implicit assumption made in the above is that the multiplications and additions dominate the overall computational complexity of the equalizer update algorithms. Of these, the multiplications are assumed to be more computationally intensive than the additions. This assumption provides a basis for estimating the complexity of an algorithm without implementing and simulating the operation of a specific hardware design.

Nevertheless, a power consumption analysis of a specific implementation of the signed update algorithm exists, and the reader is referred to [16]. Here, the computational complexity of all computational logic is considered; see Table II [17]. It should be noted that in [16], the algorithm was implemented using only 4-bits of precision and that, under this constraint, the power consumption of the additions will be comparable to the multiplications. As the signal inverters (multiplication by -1) are also implemented using addition operations, this leads to a non-negligible contribution of the additions to overall power consumption. Even so, the total power consumption of the multiplier-free update algorithm is still significantly lower than the conventional algorithm; thus, it can be assumed that the analysis presented herein holds for all hardware processing with precision greater than or equal to 4-bits.

TABLE II
IMPLEMENTATION COMPLEXITY FOR CMA USING DIFFERENT TAP WEIGHT UPDATE ALGORITHMS

	Conventional Implementation	Conventional (Gauss' Method)	Sign-Data	Sign-Sign	Fixed Cost for Error Computation	Fixed Cost for FIR Filter: Direct (Gauss)
Real Multiplications	16N+4	12N+4	8N	0	4	16N (12N)
Real Additions	16N	28N	16N	16N	4	8N+2 (20N+2)
Bit Shifts	4	4	4	4	0	0 (0)
Inverters	6	10	12N+2	12N+6	0	0 (0)
Comparators	0	0	4	6	0	0 (0)

There are several known ways to implement a complex multiplication in hardware. For a multiplication between two complex numbers, wz, where $w = a + jb$ and $z = c + jd$ and $a, b, c, d \in \Re$, the direct multiplication is computed as

$$wz = (ac - bd) + j(ad + bc) \tag{9}$$

which requires four real multiplications and two real additions. The most efficient implementation of this operation (assuming multiplication is more costly than addition) is [39]

$$wz = [d(a - b) + a(c - d)] + j[d(a - b) + b(c + d)] \tag{10}$$

which is sometimes called Gauss' method. Due to the repetition of the term $d(a - b)$, only three real multiplications are required in this implementation, although at the cost of an additional three real addition operations (five in total). The computational complexity analysis using this multiplication algorithm is also given in Table II.

ACKNOWLEDGMENT

This work was supported by the U.K. EPSRC Programme Grant UNLOC (UNLocking the Capacity of Optical Communications) EP/J017582/1 and the EU CRITICAL project.

The authors wish to thank Milen Paskov, Masaki Sato and Dr. Daniel Cardenas for valuable discussions in relation to this work.

REFERENCES

[1] S. J. Savory, "Digital filters for coherent optical receivers," *Opt. Express*, vol. 16, no. 2, pp. 804–817, Jan. 2008.

[2] K. Kikuchi and S. Tsukamoto, "Evaluation of sensitivity of the digital coherent receiver," *J. Lightwave Technol.*, vol. 26, no. 13, pp. 1817–1822, Jul. 2008.

[3] D. Lavery, R. Maher, D. Millar, B. C. Thomsen, P. Bayvel, and S. J. Savory, "Digital coherent receivers for long-reach optical access networks," *J. Lightwave Technol.*, vol. 31, no. 4, pp. 609–620, Feb. 2013.

[4] M. Faruk and K. Kikuchi, "Compensation for in-phase/quadrature imbalance in coherent-receiver front end for optical quadrature amplitude modulation," *Photonics Journal, IEEE*, vol. 5, no. 2, p. 7800110, April 2013.

[5] M. Paskov, D. Lavery, and S. Savory, "Blind equalization of receiver in-phase/quadrature skew in the presence of nyquist filtering," *Photonics Technology Letters, IEEE*, vol. 25, no. 24, pp. 2446–2449, Dec 2013.

[6] I. Fatadin, D. Ives, and S. J. Savory, "Differential carrier phase recovery for QPSK optical coherent systems with integrated tunable lasers," *Opt. Express*, vol. 21, no. 8, pp. 10 166–10 171, Apr. 2013.

[7] A. Shahpari, J. D. Reis, R. Ferreira, D. M. Neves, M. Lima, and A. Teixeira, "Terabit+ (192 × 10 Gb/s) Nyquist shaped UDWDM coherent PON with upstream and downstream over a 12.8 nm band," in *Proceedings of Optical Fiber Communication Conference*, 2013, paper PDP5B.3.

[8] D. Lavery, M. Paskov, and S. J. Savory, "Spectral shaping for mitigating backreflections in a bidirectional 10 Gbit/s coherent WDM-PON," in *Proceedings of Optical Fiber Communication Conference*, 2013, paper OM2A.6.

[9] D. Lavery, C. Behrens, and S. J. Savory, "A comparison of modulation formats for passive optical networks," *Opt. Express*, vol. 19, no. 26, pp. B836–B841, Dec. 2011.

[10] J. D. Reis, A. Shahpari, R. M. Ferreira, D. M. Neves, M. Lima, and A. L. Teixeira, "Nyquist signaling for spectrally-efficient optical access networks," in *Optical Fiber Communication Conference*, 2014, paper. W3G.3.

[11] D. Lavery and S. J. Savory, "Digital coherent technology for long-reach optical access," in *Optical Fiber Communication Conference*, 2014, paper Tu2F.1.

[12] D. Lavery, M. Ionescu, S. Makovejs, E. Torrengo, and S. J. Savory, "A long-reach ultra-dense 10 Gbit/s WDM-PON using a digital coherent receiver," *Opt. Express*, vol. 18, no. 25, pp. 25 855–25 860, Dec. 2010.

[13] N. Cvijetic, M. Cvijetic, M.-F. Huang, E. Ip, Y.-K. Huang, and T. Wang, "Terabit optical access networks based on WDM-OFDMA-PON," *J. Lightwave Technol.*, vol. 30, no. 4, pp. 493–503, Feb 2012.

[14] H. Rohde, E. Gottwald, P. Alves, C. Oliveira, I. Dedic, and T. Drenski, "Digital multi-wavelength generation and real time video transmission in a coherent ultra dense WDM PON," in *Optical Fiber Communication Conference*, 2013.

[15] C. Kottke, J. K. Fischer, R. Elschner, F. Frey, J. Hilt, C. Schubert, D. Schmidt, Z. Wu, and B. Lankl, "Coherent UDWDM PON with joint subcarrier reception at OLT," *Opt. Express*, vol. 22, no. 14, pp. 16 876–16 888, Jul 2014.

[16] D. Cardenas, D. Lavery, P. Watts, and S. J. Savory, "Reducing the power consumption of the CMA equalizer update for a digital coherent receiver," in *Optical Fiber Communication Conference*, 2014, paper Th4D.5.

[17] D. J. P. Lavery, "Digital coherent receivers for passive optical networks," Ph.D. dissertation, Dept. Elect. Eng., Univ. College London, London, UK, Oct. 2013.

[18] D. Shea, A. Ellis, D. Payne, R. Davey, and J. Mitchell, "10 Gbit/s PON with 100 km reach and x1024 split," in *Proceedings of European Conference on Optical Communication*, 2003, paper We.P.147.

[19] *APON and BPON family of recommendations ITU-T G.983*, Telecommunication Standardization Sector Std., 2001-2005.

[20] *GPON family of recommendations ITU-T G.984*, Telecommunication Standardization Sector Std., 2003-2008.

[21] K. Y. Cho, U. H. Hong, S. P. Jung, Y. Takushima, A. Agata, T. Sano, Y. Horiuchi, M. Suzuki, and Y. C. Chung, "Long-reach 10-Gb/s RSOA-based WDM PON employing QPSK signal and coherent receiver," *Opt. Express*, vol. 20, no. 14, pp. 15 353–15 358, Jul. 2012.

[22] S. Smolorz, E. Gottwald, H. Rohde, D. Smith, and A. Poustie, "Demonstration of a coherent UDWDM-PON with real-time processing," in *Proceedings of Optical Fiber Communication Conference*, 2011, paper PDPD4.

[23] D. Lavery, C. Behrens, and S. Savory, "On the impact of back-reflections in a bidirectional 10 Gbit/s coherent WDM-PON," in *Proceedings of Optical Fiber Communication Conference*, 2012, paper OTh1F.3.

[24] H. Rohde, E. Gottwald, A. Teixeira, J. Reis, A. Shahpari, K. Pulverer, and J. Wey, "Coherent ultra dense WDM technology for next generation optical metro and access networks," *J. Lightwave Technol.*, vol. 32, no. 10, pp. 2041–2052, May 2014.

[25] J. Wang, C. Xie, and Z. Pan, "Matched filter design for rrc spectrally shaped nyquist-wdm systems," *IEEE Photonics Technology Letters*, vol. 25, no. 23, pp. 2263–2266, Dec 2013.

[26] S. Savory, "Digital coherent optical receivers: Algorithms and

subsystems," *IEEE J. Sel. Top. Quantum Electron.*, vol. 16, no. 5, pp. 1164–1179, Sep. 2010.
[27] M. S. Faruk and K. Kikuchi, "Adaptive frequency-domain equalization in digital coherent optical receivers," *Opt. Express*, vol. 19, no. 13, pp. 12 789–12 798, Jun. 2011.
[28] H.-M. Chin, D. Millar, and S. Savory, "Fixed point precision requirements of the CMA for digital coherent access," in *Proceedings of IEEE Photonics Conference*, 2012, pp. 453–454.
[29] A. D. Poularikas and Z. M. Ramadan, *Adaptive Filtering Primer with MATLAB*. Taylor and Francis, 2006.
[30] D. Ogasahara, K. Fukuchi, M. Arikawa, and E. L. T. de Gabory, "Real-time evaluation of optical nonlinear effects on 112Gbps PM-QPSK signal in dispersion managed links," in *Optical Fiber Communication Conference*, 2011, paper OMR3.
[31] J. K. Fischer, R. Elschner, F. Frey, J. Hilt, C. Kottke, C. Schubert, Z. Wu, D. Schmidt, and B. Lankl, "Digital signal processing for coherent UDWDM passive optical networks," in *Proceedings of Photonic Networks; 15. ITG Symposium*, May 2014.
[32] F. Chang, K. Onohara, and T. Mizuochi, "Forward error correction for 100 G transport networks," *IEEE Commun. Mag.*, vol. 48, no. 3, pp. S48 –S55, Mar. 2010.
[33] M. Reimer, D. Dumas, G. Soliman, D. Yevick, and M. O'Sullivan, "Polarization evolution in dispersion compensation modules," in *Optical Fiber Communication Conference*, 2009.
[34] P. Krummirich and K. Kotten, "Extremely fast (microsecond timescale) polarization changes in high speed long haul WDM transmission systems," in *Optical Fiber Communication Conference*, 2004, vol. 2.
[35] S. J. Savory, G. Gavioli, R. I. Killey, and P. Bayvel, "Electronic compensation of chromatic dispersion using a digital coherent receiver," *Opt. Express*, vol. 15, no. 5, pp. 2120–2126, Mar. 2007.
[36] A. Viterbi and A. Viterbi, "Nonlinear estimation of PSK-modulated carrier phase with application to burst digital transmission," *IEEE Trans. Inf. Theory*, vol. 29, no. 4, pp. 543–551, Jul. 1983.
[37] H. Rohde, S. Smolorz, E. Gottwald, and K. Kloppe, "Next generation optical access: 1 Gbit/s for everyone," in *Proceedings of European Conference on Optical Communication*, 2009, paper 10.5.5.
[38] S. C. Davies, R. A. Griffin, A. J. Ward, N. D. Whitbread, I. Davies, L. Langley, S. Fourte, J. Mo, Y. Xu, A. Carter, "Narrow linewidth, high power, high operating temperature digital supermode distributed Bragg reflector laser," in *Proceedings of European Conference on Optical Communication*, 2013, paper Th.1.B.3.
[39] A. Wenzler and E. Luder, "New structures for complex multipliers and their noise analysis," in *IEEE International Symposium on Circuits and Systems*, vol. 2, 1995, pp. 1432–1435 vol.2.

Domaniç Lavery (S'09-M'13) received the MPhys degree in theoretical physics from the University of Durham, Durham, U.K. in 2009, and the Ph.D degree in Electronic and Electrical Engineering from University College London (UCL), London, U.K. in 2013.

His doctoral research focused on the use of digital coherent transceivers and their application to spectrally efficient, high capacity passive optical networks. He is currently continuing his work with the Optical Networks Group at UCL as a Research Associate, investigating techniques for maximizing channel capacity in nonlinear fiber transmission.

Dr. Lavery is a member of the Optical Access Systems and Wireless Backhaul Networks subcommittee for OFC 2015. He was awarded the IEEE Photonics Society Graduate Student Fellowship in 2012, and the Marconi Society's Paul Baran Young Scholar Award in 2013. He is an associate member of the institute of physics (IOP) as well as a member of the IEEE, and has previously acted as a reviewer for IEEE and OSA publications.

Benn C. Thomsen (M06) received the B.Tech. degree in optoelectronics and M.Sc. and Ph.D. degrees in physics from The University of Auckland, Auckland, New Zealand.

His doctoral research involved the development and characterization of short optical pulse sources suitable for high-capacity optical communication systems. He then joined the Optoelectronics Research Centre, Southampton University, U.K., as a Research Fellow in 2002, where he carried out research on ultrashort optical pulse generation and characterization, optical packet switching based on optically coded labels, and all-optical pulse processing. He joined the Optical Networks Group, University College London, London, U.K., in 2004, and held an EPSRC Advanced Fellowship from 2006 to 2011 and was appointed as a lecturer in 2007. He is currently a senior lecturer at UCL and his research focuses on optical transmission, physical-layer implementation of dynamic optical networking technology and the development of high capacity multimode fibre systems exploiting MIMO DSP.

Polina Bayvel (S'87-M'89-SM'00-F'10) received her B.Sc. (Eng) and Ph.D. degrees in Electronic and Electrical Engineering from the University of London, UK, in 1986 and 1990, respectively. In 1990, she was with the Fiber Optics Laboratory, General Physics Institute, Moscow (Russian Academy of Sciences), under the Royal Society Postdoctoral Exchange Fellowship. She was a Principal Systems Engineer with STC Submarine Systems, Ltd., London, UK, and Nortel Networks (Harlow, UK, and Ottawa, ON, Canada), where she was involved in the design and planning of optical fibre transmission networks. During 1994-2004, she held a Royal Society University Research Fellowship at University College London (UCL), and in 2002, she became a Chair in Optical Communications and Networks. She is currently the Head of the Optical Networks Group, UCL. She has authored or co-authored more than 300 refereed journal and conference papers. Her research interests include wavelength-routed optical networks, high-speed optical transmission, and the mitigation of fibre nonlinearities.

Prof. Bayvel is a Fellow of the Royal Academy of Engineering (FREng.), the Optical Society of America, the UK Institute of Physics, and the Institute of Engineering and Technology. She was the recipient of the Royal Society Wolfson Research Merit Award (2007-2012), 2013 IEEE Photonics Society Engineering Achievement Award and 2014 Royal Society Clifford Patterson Prize Lecture and Medal.

Seb J. Savory (M'07-SM'11) received the M.Eng., M.A., and Ph.D. degrees in engineering from the University of Cambridge, U.K., in 1996, 1999, and 2001, respectively, and the M.Sc. (Maths) degree in mathematics from the Open University, Milton Keynes, U.K., in 2007.

His interest in optical communications began in 1991, when he joined Standard Telecommunications Laboratories, Harlow, U.K., prior to being sponsored through his undergraduate and postgraduate studies, after which he rejoined Nortels Harlow Laboratories in 2000. In 2005, he joined the Optical Networks Group at University College London (UCL) where he held a Leverhulme Trust Early Career Fellowship from 2005 to 2007, before being appointed as a Lecturer in 2007 and subsequently Reader in Optical Fibre Communication in 2012. From June 2009 to June 2010, he was also a Visiting Professor at the Politecnico di Torino, Torino, Italy and in 2014 he was awarded a Royal Academy of Engineering / Leverhulme Trust Senior Research Fellowship to conduct research into Realising the Capacity in Fibre-Optic Networks with Uncertainty and Nonlinearity.

Dr. Savory is the Editor-in-Chief of IEEE Photonics Technology Letters, an IEEE Photonics Society representative on the Steering Committee of OFC and will serve as a General Chair for OFC in 2015. He previously served as a Program Chair for OFC 2013, Signal Processing in Photonic Communications (SPPCom) 2012 and SPPCom 2013. In addition, he serves on the Editorial Board for IET Optoelectronics and the technical program committee for the European Conference on Optical Communication.