

Low Complexity Carrier Recovery System Design for Baudrate Sampled Coherent Datacenter Interconnects

Jonas De Busscher, Axl Bomhals, Gertjan Coudyzer, *Member, IEEE*
Peter Ossieur, *Senior Member, IEEE*, Xin Yin, *Senior Member, IEEE*

Abstract—A novel low complexity carrier recovery system for baudrate coherent datacenter links is presented. By separating the phase contributions from residual frequency offset and laser phase noise, this system reduces the required number of complex multiplications for frequency recovery by 50% and for phase recovery by 81.2%. The design incorporates a multi-stage serialized blind phase search to estimate the average phase within a processing block, while another parallel branch estimates the total phase drift across a DSP block. This feed-forward structure enables highly parallelized DSP implementation while maintaining robustness against large carrier frequency offsets and cycle slips. The method is validated through both simulation and measurements for QPSK and 16QAM, demonstrating large power savings and improved tolerance for frequency offset compared to state-of-the-art techniques.

Index Terms—Coherent receiver, Carrier recovery, blind phase search, baudrate, low complexity, DSP

I. INTRODUCTION

DUE to their significantly higher power consumption compared to IM-DD systems, partially driven by increased digital signal processing (DSP) complexity, coherent transceivers have not yet achieved widespread adoption in datacenter environments. However, in short-reach applications, many power hungry DSP functions, such as chromatic dispersion and polarization mode dispersion compensation, can be eliminated or substantially simplified [1]. This shifts the primary DSP power burden to carrier recovery [2], which comprises carrier frequency recovery and carrier phase recovery. As datacenter networks continue to demand increasingly high data rates, short-reach coherent optical interconnects are the next logical step to achieve these targets [3] and minimizing the power consumption of these remaining DSP blocks becomes critical. Efficient carrier recovery architectures therefore play a key role in enabling next generation, energy efficient datacenter interconnects.

The analog front end typically operates at much higher speeds than the digital processing clock, introducing parallelisation as an additional constraint [4], [5]. Block-based processing increases delay in feedback loops, which can lead to implementation challenges. For this reason, feedforward solutions are generally preferred.

A commonly used approach for frequency recovery is the fourth power FFT method, which identifies the peak at four times the frequency offset in the spectrum of the fourth power of the input signal [6]. This technique offers

improved performance for a given block length at the cost of the computational complexity associated with fourth power operations and an FFT compared to other methods [7], while being easily implemented in parallel.

Carrier phase recovery must track the random walk of phase noise [8] and compensate for residual carrier frequency offset that remains due to the limited resolution of the frequency compensation block [9]. Multiple methods to tackle these problems have been proposed. Fourth power Viterbi-Viterbi [10] has been used primarily in QPSK systems, however, it has not been adapted by other formats due to weaker performance for higher order constellations [9]. The blind phase search proposed in [11], applies a set of possible test phases and the best match according to a cost metric is chosen, the main limitation of this is computational complexity. Every test angle requires a complex rotation, a cost metric and a decision, leading to increased power consumption if higher accuracy is needed. As only the best result is chosen, a lot of computation could be considered wasted. Moreover, introducing block-based processing instead of sliding window or symbol by symbol decision can decrease computational complexity by only using each symbol once, but results in significantly reduced performance due to fewer updates of the phase [12]. Blind algorithms are also susceptible to cycle slips when using rotationally symmetric constellations. Because these algorithms cannot distinguish between constellation points rotated by their symmetry angle, catastrophic failures can arise that cannot be corrected without a pilot tone [13].

Several techniques have been intensively researched to reduce the complexity of blind phase search. A two-stage BPS circuit was proposed in [14], performing both coarse and fine phase estimation using BPS. However, the approach is inherently limited to two stages due to the pattern effect, which can induce a non-monotonic cost function for certain input sequences when too few test phases are used per stage. As a result, a two-stage configuration is considered optimal [14]. In [15], [16], [17] multi-stage approaches using a maximum likelihood estimator as a second stage achieve a two to three fold reduction in complexity. Using this architecture, the first stage still requires a significant amount of test angles due to the pattern effect.

Quadratic interpolation enables an eight fold reduction in complexity [18], but requires feedback and symbol by symbol decisions, which complicates parallelisation. More recently, a feedforward interpolation based method achieved a two fold reduction in computational complexity [19].

In this work, we present a novel feed-forward multi-stage blind carrier phase recovery circuit using block-based process-

The authors are with IDLab, Department of Information Technology, Ghent University - imec, 9052 Gent, Belgium
(e-mail: jonas.debusscher@ugent.be)




Fig. 1. Top-level block diagram of proposed solution. A fraction of the input symbols are used in serialised blind phase search (SBPS), while another set is used in spread estimation. Both results are used to calculate the corrections needed for phase unwrapping (PU).

ing. This system significantly reduces the impact of the pattern effect for 16QAM, enabling a multi stage architecture with a small amount of test phases per stage as a result of a modified cost function. Intra-block phase compensation and extended frequency tracking mechanisms are introduced to reduce cycle slips, resulting in less stringent requirements for the frequency recovery block, reducing the rate at which fourth power and FFT operations need to be performed. This approach achieves a reduction of the amount of complex multipliers in blind phase search of 81.2% and 50% for frequency compensation.

The remainder of this manuscript is organized as follows. Section II describes the algorithm and its sub-blocks for QPSK modulated signals, while Section III details the adaptations required for 16QAM. Section IV shows the simulation results and section V compares the computational complexity with state-of-the-art methods. Measurements are discussed in section VI, followed by the conclusions in section VII.

II. IMPLEMENTATION FOR QPSK

To solve the problems mentioned above, we propose a novel multi-stage carrier phase recovery scheme, allowing a reduction of the computational complexity in the total carrier recovery system. Figure 1 shows a high level blockdiagram of the proposed system. At the input side, data symbols are presented in parallel blocks of blocksize B . Then, a subset of this block will be used in serialised blind phase search (SBPS) and another subset will be used for spread estimation. We distinguish two types of contributions to the phase of a received symbol. A deterministic contribution as a result of the finite resolution of the frequency recovery system and a stochastic contribution due to the finite laser linewidth, characterised by a Wiener random walk process [20]. The total phase shift ψ_i of a received symbol a_i can be written as the sum of both contributions

$$\psi_i = 2\pi\Delta f_{res}T_s i + \phi_i \quad (1)$$

With Δf_{res} the residual frequency offset post frequency recovery, T_s the symbol duration and ϕ_i the phase shift of the symbol due to phase noise. In this work, two parallel branches are used, each estimating one of the two contributions. The first branch performs serialised blind phase search, significantly reducing the computational cost compared to a parallel implementation as the estimated angle is successively approximated. This performs estimation of the phase shift in the middle of each processing block.




Fig. 2. Examples of possible choices for the number of test angles per stage (first stage in this example). Both show the set with odd or even number of test phases, meaning with and without 0.

A second branch estimates and tracks the phase shift over the full processing block. Depending on the resolution and update rate of the frequency recovery block, this can dominate the total phase shift. As the frequency offset is slowly varying compared to the update rate of the phase search circuit, the phase contributions caused by frequency offset are treated as deterministic.

The information of both branches is then combined in the phase unwrapper to extract the true phase of the symbols in this block. Knowledge of the residual frequency offset allows us to reduce cycle slip rate without adding complexity to the phase unwrapper. Reducing the cycle slip rate and applying fine phase corrections within a block, allows us to estimate the frequency offset sparsely. In the next sections, each sub-block will be discussed in detail.

A. Serialized blind phase search

The total phase space that needs to be searched stretches from $-\frac{\pi}{4}$ up to $\frac{\pi}{4}$ as $\frac{\pi}{2}$ is a rotational symmetry angle for all M-QAM constellations. This novel architecture uses a multi-stage design with N_s test angles per stage and $M = \log_{N_s}(N)$ stages. With each stage refining the phase search. N being the amount of test angles in a traditional approach as presented in [11]. Each stage rotates the received data by some angle and computes a cost metric based on the distance to the closest constellation point. The angle with the smallest cost metric is selected. The next stage will then add its angles to the outcome of the previous stage. An example of such a circuit is shown on figure 3.

N_s is selected such that maximum efficiency is achieved, which requires N_s to be even. This can be shown with the following reasoning. If N_s would be odd, that means that 0 would be added to the test angles. If N_s would be even, one can choose the test phases in such a way that they divide the positive and negative part of the search region equally as shown on figure 2, leading to a non uniform distribution of the test phases. The dashed arrow represents a test angle with value 0, including this one divides the entire search region in equal parts.

Now assume that 0 were to be the correct phase offset to be estimated and N_s is even. The first stage would have a 50% chance of selecting the smallest positive test phase, depending




Fig. 3. Serialized blind phase search block diagram. Amount of test phases per stage is 2.

on noise. The next stage then sees a positive error and would correct with the largest negative one. The next stage is always more refined than the previous, so all subsequent stages will correct for the initial positive guess. After the last stage, the error is again reduced to the resolution of the system, so the test angle 0 is implicit in a multi-stage design. The same reasoning holds when the first test angle would be chosen negative.

assuming that $N_s = 2K$, with test angle positive and negative search region in equal parts and M the total amount stages, the procedure to obtain the estimated phase ϕ using M stages, described above can be written more formally as:

$$\phi_M = \frac{\pi}{4} \frac{1}{K+1} - \sum_{m=2}^M \frac{K\pi}{4(K+1)^m} \quad (2)$$

With the first term being the selection of the smallest positive test angle and the other terms representing the choice of the largest negative test angle in all subsequent stages. Now in the ideal case let $M \rightarrow \infty$, and with some rearranging:

$$\begin{aligned} \phi_\infty &= \frac{\pi}{4} \frac{1}{K+1} - \frac{K\pi}{4} \left[\sum_{m=0}^{\infty} \left(\frac{1}{K+1} \right)^m - a_0 - a_1 \right] \\ a_0 &= 1 \\ a_1 &= \frac{1}{K+1} \end{aligned} \quad (3)$$

The series of m is now a converging geometric series and the total phase $\phi_\infty = 0$ can now be obtained. An implementable application requires M has to be finite, so some error will still be present. Assuming $N_s = 2K$ the test angles in the first stage can be written as

$$\phi_k = \pm \frac{\pi}{4} \frac{k}{K+1}, \quad \forall k \in \{1, \dots, K\} \quad (4)$$

resulting in a total maximum error after M stages of

$$e_{max}(K, M) = \frac{\pi}{4(K+1)^M} \quad K, M \geq 1 \quad (5)$$

While using $K \times M$ test angles worth of computation. Increasing the amount of test angles per half by one, leads to a reduction of the maximum error by a factor of

$$\begin{aligned} \frac{e_{max}(K, M)}{e_{max}(K+1, M)} &= \left(\frac{K+2}{K+1} \right)^M \\ &= \left(1 + \frac{1}{K+1} \right)^M \end{aligned} \quad (6)$$

The linear rise of computational effort and the inverse reduction of marginal gains in maximum error, leads us to conclude that fewer test angles per stage is better. In this work, N_s is chosen to be two ($K = 1$), implementing a binary search tree.

A schematic of the serialised blind phase search implementation is shown on figure 3, where $\Phi = \frac{\pi}{4}$, half of the search space for M-QAM constellations. The first stage performs a rotation of the data block over $\pm \frac{\pi}{8}$. A decision operation is performed and the distance of each rotated symbol to its decided symbol is calculated and averaged over all parallel input samples, reducing the influence of random noise.

The test angle resulting in the minimum deviation from a correct constellation is determined and is passed through to the next stage, where the testangle is added to the previous results. For $N_s = 2$, the test angles at stage m are

$$\phi_m = \pm \frac{\pi}{2^{m+2}} + \sum_{i=1}^{m-1} \phi_i \quad (7)$$

This way, the correct phase offset of a block is approximated better and better by each stage. As an example, a serialised implementation with 4 stages will test up to $\frac{\pi}{64}$, requiring 8 test angles. While a standard implementation would need 32 test angles [11], showing that only $2(\log_2(N) - 1)$ test angles are needed for N equivalent number of test angles.




Fig. 4. Phase spread and direction detection on fourth power of a block.

B. Spread detection

The blind phase search (BPS) branch estimates only the average phase over each processing block. However, in the presence of a large residual frequency offset, the phase may vary significantly from the first symbol of the block to the last. In block-based architectures, like this work, with a high degree of parallelization, this intra block phase variation can lead to performance degradation, as the single average phase estimate is no longer representative of all symbols in the block. To address this limitation, the spread detection subsystem detects and tracks the total phase shift between the first and last symbol of each block. This additional information contributes to the total phase compensation in two ways:

- correct for intra block phase differences
- estimate the total phase shift in a block to perform phase unwrapping based on extrapolated phases. Allowing us to increase the SBPS blocksize, further mitigating the pattern effect

For coherent optical communications, low linewidth lasers are used. At first the assumption will be made that the accumulated phase noise over one data block is much less than the accumulated phase due to frequency offset of the data and local oscillator (LO).

$$\sum_{i=0}^{B-1} \phi_i \ll 2\pi\Delta f T_s B \quad (8)$$

With B the data processing blocksize and ϕ_i the phase noise due to linewidth of symbol i . The total phase shift over a data block can then be assumed constant between neighboring blocks. To estimate this phase spread accurately, three symbols a_0, a_b, a_{B-1} are chosen from the current data block, where b is an index between 0 and $B-1$, here chosen to be halfway in the block. The symbols are raised to the fourth power, ordering them chronologically by removing the phase modulation. Figure 4 shows the resulting constellation plot.

From this the total phase difference between the first and last symbol can be estimated using:

$$\begin{aligned} \frac{\langle a_0^4, a_{B-1}^4 \rangle}{\|a_0^4\| \cdot \|a_{B-1}^4\|} &= \cos(4|\theta|) \\ \Rightarrow \frac{\arccos(\langle a_0^4, a_{B-1}^4 \rangle)}{4} &\approx |\theta| \end{aligned} \quad (9)$$

As the carrier recovery is located after the equaliser, it is assumed that the denominator of equation 9 is one. However, this assumption leads to a some inaccuracy for low SNR. This




Fig. 5. Block diagram of spread detection circuit.

observation together with a lookup table for $\arccos(x)$ leads to a low complexity estimation of θ . Using this method, the estimated angle is always positive. This is required as phase noise can have the total phase shift extend over π which in case of using $\arctan(x)$ to estimate the angle, would lead to the wrong sign. Causing many cycle slips and bit errors.

Estimating the direction of the phase shift will rely on a_b , which is chosen in the middle of a block. The direction of rotation can be written as

$$\begin{aligned} \text{sign}(\gamma) &= \text{sign}\left(\angle\left(a_b^4 (a_0^4)^*\right)\right) \\ &= \text{sign}\left(\arctan2\left(\frac{\Im\{a_b^4 (a_0^4)^*\}}{\Re\{a_b^4 (a_0^4)^*\}}\right)\right) \end{aligned} \quad (10)$$

Using the definition of $\arctan2(x)$ and taking into account the assumption that $\gamma \in [0, \pi)$

$$\text{sign}(\gamma) = -\text{sign}(\Re\{a_0^4\}\Im\{a_b^4\} - \Im\{a_0^4\}\Re\{a_b^4\}) \quad (11)$$

Only the sign is required, again leading to a low complexity calculation. Also note that choosing a_b and a_n to estimate the direction, analogous calculations can be used. Now using equations 11 and 9 the total phase shift over the data processing block, can be estimated as:

$$\hat{\theta} = \text{sign}(\gamma) \frac{\arccos(\langle a_0^4, a_{B-1}^4 \rangle)}{4} \quad (12)$$

This method heavily relies on reasoning on the constellation. As a consequence of this, low SNR will result in low estimation quality. Figure 6 shows the accuracy σ of the estimation calculated as:

$$\sigma = \frac{1}{B_{\text{sim}} - 1} \sum_{i=0}^{B_{\text{sim}}} 1 - \left| \frac{\hat{\theta}_i - \theta_i}{\theta_i} \right| \quad (13)$$

with B_{sim} the amount of simulated blocks. The intended applications use a high data rate and stable lasers, meaning that the assumption can be made that the laser frequency offset will be stable over multiple blocks. For example for a 32Gbaud transceiver using blocksize of 64, the total elapsed time from beginning of the block to the end would be $\frac{64}{32e9} = 2\text{ns}$. While commonly used DFB lasers have a coherence time that is orders of magnitude greater. Due to this large difference, a majority vote block can be added to the direction estimation and a filter to the spread estimation. This will cancel out any sudden estimation errors as long as the majority is correct. In order to keep the filter as simple as possible, it is implemented




Fig. 6. Accuracy increase by using filters, blocksize 64 and $\Delta f = 65\text{MHz}$.

as an average over a power of 2 number of estimations. Only adding 2 additions and a shift register to the complexity of the circuit. The size of the majority vote and averaging blocks should be chosen depending on the speed of laser frequency variations. For example, a 32Gbaud system with blocksize 64, has an update rate of 500MHz. Thus adding 512 times averaging will result in tracking capability of variations in the order of 1MHz.

C. Correction and phase unwrapping

After estimating the phase offset of the middle symbol $\langle\phi\rangle$ using blind phase search and estimating the total phase spread θ , this can be combined to estimate the absolute phase shift from the first symbol to the last symbol of the data block. Due to the rotational symmetry of the constellation, this needs to be unwrapped in order to prevent cycle slips. The phase of the first and last symbol are calculated as:

$$\begin{aligned} \angle a_0 &\approx \langle\phi\rangle - \frac{\theta}{2}, & \text{lagging phase} \\ \angle a_{B-1} &\approx \langle\phi\rangle + \frac{\theta}{2}, & \text{leading phase} \end{aligned} \quad (14)$$

The leading and lagging phase are used in the phase unwrapper to check whether a multiple of $\frac{\pi}{2}$ should be added. The leading phase of the previous block is compared to the lagging phase of the current block. When using only estimated phases by SBPS, the time elapsed between 2 phases to check for wrapping is BT_s . Using these extrapolated estimations for the first and last phase in a block allows us to reduce this elapsed time to the symbol time T_s , improving our detection of phase wrapping when the residual frequency offset is high. Whenever the absolute value of the difference is larger than $\frac{\pi}{4}$, a counter is increased or decreased. Based on this counter a multiple of $\frac{\pi}{2}$ is added to the phase calculated by the serialized blind phase search. Figure 8 qualitatively illustrates the estimated phase and the extrapolated phases being used to detect phase wrapping. The estimated phases by SBPS (red dots) are not sufficient to detect phase wrapping due to the large frequency offset present. The crosses above and below each estimation point indicate the estimated phase range that was calculated

in the phase spread block. The use of the extrapolated phases allows us to detect that phase wrapping has not occurred in the left case and has occurred in the right case.




Fig. 7. Phase unwrapper using the leading and lagging phase.

D. Phase compensation

The phase unwrapping and correction prevents cycle slips, even at larger residual frequency offsets. However, the symbols in a block are still spread out over a large phase range. To compensate for the phase error at multiple locations along the data block, the middle phase error ϕ_e and phase spread θ are used to emulate more frequent estimation by interpolation. The data is divided in N sections and θ is divided by $2N$, each of the N subsections is then compensated for by $(2k+1)\frac{\theta}{2N}$. This way, all subsections are compensated for by the interpolated phase of the middle symbol of the subsection. This results in less phase noise after compensation when the frequency offset is high as shown on figure 9. Here a blocksize of 64 is used, window size 32 and a frequency offset of 32.5MHz at a baudrate of 32GBaud.

E. Considerations for low frequency offset

Until now, the frequency offset was assumed to be large. However, when the frequency offset is low, the phase noise can be dominant over the frequency offset. This results in low spread and direction detection accuracy. Whenever the frequency offset is low, due to the presence of phase noise, the accuracy of the total phase shift over a block will be




Fig. 8. Use of extrapolated phases for phase unwrapping. Phase wrapping is determined based on the leading and lagging phase (crosses), instead of using the SBPS output (dots).




Fig. 9. Comparison between output of phase compensation without (a) and with (b) spread detection, for QPSK. Using blocksize $B=64$, window size $W=32$ and frequency offset $\Delta f = 32.5\text{MHz}$.

limited. Due the averaging, as described in subsection II-B, the estimated error will be small so it will not contribute much when calculating the leading and lagging phase for unwrapping or the subblock phases for compensation. To fully mitigate the effect, a deadzone of 0.1 radians is used for the frequency tracking branch, it will have no contribution when frequency offset is low. In this scenario, there are multiple estimations per $\frac{\pi}{2}$ of phase rotation performed by SBPS, rendering the extrapolated phases unnecessary for cycle slip detection. This is illustrated further in section IV where the simulation results are discussed.

III. EXTENSION TO 16QAM

The section above describes how the system operates for QPSK constellations. The solution is constellation aware, so some adaptation is needed for it to work for 16QAM.

A. Serialised blind phase search

Blind phase search algorithms are known to suffer from the pattern effect when using them for M-QAM [14], [18]. Leading to an input dependent cost function that can be non-monotonic with local minima. When using a large amount of test phases, this problem can be avoided as they cover the entire search space, always choosing the global minimum. When only two test phases are used, the algorithm can easily converge to a local minimum, resulting in incorrect estimation and compensation. These local minima arise due to incorrect decision of the symbols during the calculation of the cost function. Figure 10 shows a resulting constellation diagram when using two test angles. It is clear that the BPS will converge to multiple local minima, resulting in many symbol errors.

Although the cost function is input dependent, some insight can still be gained by considering the average cost function. Figure 11 shows the result of rotating the full constellation from 0 to $\frac{\pi}{2}$. The angle at which the error function enters the non-monotonic region is 0.293 and 1.28 radians. The total non-monotonic region is larger than the desired region.

We propose a modification to the cost function that significantly reduces the impact of the pattern effect. At the input of the SBPS, the radius of each symbol will be checked. If it exceeds a threshold, the decision is performed using the outer




Fig. 10. Comparison between output of serialized blind phase search using standard decision (a) and radius directed decision (b) for 16QAM using blocksize $B=64$, window size $W=32$ and frequency offset $\Delta f = 32.5\text{MHz}$.




Fig. 11. Standard error function showing local minima and a monotonic radius directed error function. The non-monotonic region lies between 0.293 and 1.28 radians.

four QPSK like symbols. Otherwise the decision is performed using the inner twelve as a decision set. This adds a radius calculation for each symbol that is used for phase estimation. This works because the outer symbols have a greater impact on the cost function for a given rotation due to their distance from the origin. Adding a radius directed decision will avoid an erroneous decision involving these outer symbols, adding a large monotone contribution. Figure 10 shows the obtained constellation when using this radius directed approach. Figure 11 shows the difference between the standard cost function and the modified version when using 2 test angles per stage.

B. Spread detection

16QAM constellations have additional phase and amplitude modulation compared to QPSK. When raising such a constellation to the fourth power, four points are obtained at $-4.24 + 0j, -0.04 + 0j, 0.28 + 0.96j, 0.28 - 0.96j$. Forming three concentric circles with radii 4.24, 0.04, and 1 when a frequency offset is present, depicted on figure 12.

The angle between the two symbols on the center circle is $\alpha = 147.5^\circ$. When these point would be used to estimate the total phase spread over a data block, it would lead to 147.5° of ambiguity. Resulting in incorrect estimations.




Fig. 12. 16QAM symbols to the fourth power.

To solve this, only points from the outer and inner QPSK-like rings are used as the fourth power operation removes the phase modulation completely.

The direction estimation candidates and spread estimation candidates can be chosen separately as it is assumed that frequency offset is constant in subsequent blocks. The candidate symbols are then divided by their radius, after which the same operations as for QPSK are performed. Only half the symbols in 16QAM are suitable for spread estimation. So in order to increase the frequency of estimations, the outer L symbols are searched instead of only the ones located at the edge of a block. For estimation purposes, the outer L symbols will be treated as if they would be located at the edge of the block. If L is larger, there will be more frequent estimation at the expense of accuracy as the symbols can be removed further from the outside of a block. To evaluate this, we first look at the probability of finding a good pair of symbols in the first and last L symbols of a block. Taking into account that the probability of receiving a symbol that either lies on the inner or outer circle in 16QAM is $\frac{1}{2}$, it can be shown that the probability of finding matching pairs in the two edges of the data block can be expressed as

$$\Pr(\text{hit}) = 1 - \Pr(\text{no hit}) \quad (15)$$

$$\Pr(\text{hit}) = 1 - \left[\left(\frac{1}{2L} \right)^2 + 2 \left(\frac{1}{2L} \right) \left(1 - \frac{1}{2L} \right) \right] \quad (16)$$

This rises rapidly when expanding the search subset. In this work $L = 3$ has been chosen as a compromise, yielding a probability of 76.5% finding a pair in any given 2 sets of three symbols.

This adds a radius calculation to the required operations, which is equivalent to a complex multiplication. When a symbols from the inner circle would be chosen to estimate the spread, this always leads to an incorrect estimation. For this reason, the averaging following the estimation is required for 16QAM for better performance at low SNR where it was optional for QPSK. This way the erroneous estimations get averaged out with the correct ones. When adding the above considerations to the circuit, it also works for 16QAM.

IV. SIMULATION RESULTS

A. Simulation model

For all simulation results, a PRBS15 sequence is generated and modulated as a 32Gbaud QPSK or 16QAM signal at one sample per symbol. Then AWGN is added, followed by Wiener phase noise [20] and frequency offset, total phase noise was added as $\text{PN} = \exp(j(2\pi\Delta f t + \phi(t)))$. A Simulink model and testbench of the proposed architecture was made. It is assumed that the only impairment left in the signal is due to laser frequency offset and laser phase noise as carrier recovery is the final block in a standard coherent DSP chain.

B. Cycle slip rate

During phase unwrapping, a phase discontinuity can be missed, leading to a cycle slip. For M-QAM constellations, this shows up as $\frac{\pi}{2}$ ambiguity. If no pilot tone or additional coding is applied, the bit error rate would eventually rise to 0.5. However, this requires extra overhead and for non gray coded constellations worse performance whenever no cycle slip occurs. in accordance with [21] the cycle slip rate (CSR) is computed as:

$$\text{CSR} = \frac{1}{B_{\text{sim}} - 1} \sum_{b=2}^{B_{\text{sim}}} |\tilde{\phi} - \tilde{\phi}_{-1}| \quad (17)$$

With B_{sim} the total number of blocks simulated and $\tilde{\phi}$ defined as:

$$\tilde{\phi} = \text{round} \left[\frac{2(\hat{\phi} - \bar{\phi})}{\pi} \right] \quad (18)$$

A total of 2^{15} blocks were simulated each containing 64 symbols, generated at one sample per symbol for a 16QAM constellation. As the mechanism for reducing cycle slips is the same as for QPSK, they both give similar results. The systems are tested using the same equivalent amount of test angles, being 64 and blocksize $B = 32$. The main difference between the methods being the extended frequency tracking, made possible by the additional phase spread calculation circuit. Figures 13a and 13b show a comparison between the proposed solution and the standard BPS. The CSR is calculated over different linewidth symbol duration products ($\Delta\nu T_s$) and frequency symbol duration products ($\Delta f T_s$). One can see that for larger frequency offsets, the proposed solution has a significantly higher tolerance for cycle slips. This further supports our claim that the frequency recovery block can be simplified and its update interval increased, resulting in even more efficiency gains due to the high computational effort needed to compute fourth powers of complex numbers.

C. Bit error rate

Figure 14 shows the cycle slip rate (CSR) and bit error rate (BER) for the proposed method and a standard solution using QPSK modulation. In these simulations, the DSP block size is set to 64 and window size to 32, receiving a PRBS15 sequence of 2^{20} symbols. A remark should be made when interpreting the BER curve. For the proposed method, a BER




Fig. 13. Cycle slip rate over frequency offset and laser linewidth for standard blind phase search (a) and this work (b). 64 equivalent amount of test angles and blocksize 64.

of approximately 0.5 is observed until the signal reaches a certain SNR threshold. Beyond this point, all frequency offsets follow the same BER trend. This behavior can be explained by the CSR curve as long as cycle slips can occur, they will eventually happen over a sufficiently long observation time, resulting in a BER value of 0.5 or a non representative BER. For example, the green curve shows a BER of 0.1 at an SNR of 12dB. However, this is an artifact of the finite simulation length. In this particular simulation run, the cycle slip occurred relatively late, leading to a large number of correctly detected bits before the slip occurred. If the simulation were extended indefinitely, the BER would converge to 0.5. The results demonstrate that the proposed method produces fewer cycle slips for a given frequency offset while maintaining performance comparable to the standard method for smaller frequency offsets. The standard method is not able to compensate for residual frequency offsets up to $1e - 3\Delta fT_s$.

A similar trend can be observed in Fig. 15. For the 16QAM

case, however, a BER floor appears even when no cycle slips occur. This effect arises from symbol errors in the inner ring of the constellation. The accumulated intra-block phase shift can lead to incorrect symbol decisions, depending on the residual frequency offset presented to the phase recovery block. Despite this limitation, the proposed method still achieves improved performance compared to the case where BPS is used.

V. COMPLEXITY COMPARISON

Table I shows how the complexity of the proposed method compares to the current state-of-the-art implementations of feed-forward blind phase search related methods in function of the blocksize B , amount of test angles N and the summing window W . In most entries, the summing window is not mentioned as it is the same as the blocksize. In our proposed solution, the window size is a fraction of the total block size for SBPS phase estimation.

In addition to the complexity, it is important to note that the degree of parallelisation differs in other architectures. In this work all samples are processed block by block instead of using a sliding window. As only the amount of operations is listed, this does not take into account how often it is executed. Here blocksize 64 is used with window size 32, this degree of parallelism is not achievable with standard BPS for the same complexity as has been shown in the simulation section IV.

VI. MEASUREMENTS

Figure 16 shows a schematic of the measurement setup. PRBS generation, symbol mapping, and pulse shaping are performed in MATLAB, and the waveform is produced by a Keysight M8196A 92GSa/s AWG with 0.15V output amplitude. A PRBS10 sequence is used together with RRCOS pulse shaping with a roll-off factor of 0.35. The in-phase and quadrature signals are amplified by SHF S807C amplifiers providing 23dB gain to drive the MZM. An IXBLUE TFLN MZM modulates the electrical signal onto a 1550nm optical carrier generated by a Keysight N7714A laser source. A 99/1 splitter diverts a small portion of the optical power to the ABC to maintain operation at the minimum-transmission point. After 1km of SMF-28 fiber, a Keysight N7752A tunable attenuator controls the received signal power. To obtain input powers between -5dBm and -10dBm, an EDFA is inserted before the attenuator. The 90° optical hybrid is polarization sensitive, requiring a polarization controller at its input. The LO is generated from another output port of the same Keysight N7714A, with the frequency offset between the two outputs nominally set to 0GHz. Coherent detection is performed using BPD3320R-VF-FP balanced photodiodes with 70GHz EO bandwidth and 0.5 responsivity at 1550nm when biased at $\pm 2.8V$. Two SHF M827A amplifiers are placed before the Lecroy 10-65Zi-A 160GSa/s RTO to ensure sufficient input amplitude. The RTO records 0.5ms of data to accommodate laser frequency drift, and offline DSP is performed in MATLAB and Simulink. The signal is resampled to 2sps and passed through a matched filter. Due to the long fiber span between the hybrid and the separate balanced photodetector modules, IQ skew at the receiver is estimated and compensated. After timing recovery,




Fig. 14. Cycle slip rate (CSR) and (bit error rate) BER Vs signal to noise ratio (SNR) for different frequency offsets, using QPSK. Total linewidth set to 0.5MHz. Lower figures use standard BPS, top figures use proposed solution.




Fig. 15. Cycle slip rate (CSR) and (bit error rate) BER Vs signal to noise ratio (SNR), using 16QAM. Total linewidth set to 0.5MHz. Lower figures use standard BPS, top figures use proposed solution.




Fig. 16. Schematic of the measurement setup.

the signal is downsampled to 1sps and equaliser using an 11 tap FFE. A Simulink model is developed to demonstrate the frequency and phase recovery. Data is streamed in from a file and processed in parallel using a blocksize of 64. To compare the proposed method with standard BPS, models with identical computational complexity are implemented. For QPSK, a five-stage serialized BPS is used, and for 16QAM, a six-stage version. The frequency tracking branch performs a 512 times averaging. The DSP block size is set to 64, and the SBPS summing window uses the 32 central samples. Frequency offset estimation and compensation are implemented using a fourth-power and FFT-based technique. Due to the increased frequency-tracking capability of the proposed approach, the frequency is estimated only once every 100 μ s using a 2^{10} -point FFT. This interval is chosen such that the residual frequency offset remains within the trackable region of the phase-compensation stage, preventing cycle slips. Figure 17 shows the output of the spread estimation and frequency estimation block over time. Whenever a new estimation of the frequency offset is performed, the residual frequency offset is low, meaning that the spread estimation block can be reset, removing all memory effects. This allows us to achieve higher degree of parallelisation without increasing the update rate of the frequency recovery block.

On figures 18 and 19, 3 different architectures are shown, this work, BPS₁ and BPS₂. BPS₁ is a standard BPS implementation and is equal in complexity to this work. BPS₂ also is a standard BPS implementation, but achieves the same degree of parallelism with comparable BER performance at the cost of higher computational complexity. As stated before, the proposed solution updates the frequency offset every 100 μ s, while the standard BPS solution with comparable




Fig. 17. Frequency estimation every 100 μ s and residual frequency offset tracking.

BER performance, updates the frequency offset every 50 μ s. This shows 50% power savings for this work compared to the standard solution.

The QPSK curve for this work matches the one obtained from simulations, seen on figure 14. When the signal has sufficient SNR, the proposed system is capable of tracking extended frequency offset, leading to a discontinuity in the BER curve at -34dBm. The standard blind phase search implementation is not capable of compensating for the present frequency offset and phase noise when using the same computational complexity as this work, showing a BER close to 0.5 for all input powers.

TABLE I
COMPLEXITY COMPARISON BETWEEN SEVERAL BLIND PHASE SEARCH ARCHITECTURES

architecture	multiplication	additions	comparisons	decisions	$\sqrt{x}$
BPS[11]	$6NB + 4B$	$6NB + 2B + 2$	$N - 1$	NB	0
MB-BPS [19]	$6NB + 10B + N$	$6NB + 6B + N + 2$	$N + 1$	NB	0
BPS-ML [17]	$6(N + 2)B + 4B$	$6(N + 2)B + 2B + 2$	$N + 1$	$(N + 2)B$	0
PCPE-BPS [21]	$6NB + 14B + 8$	$5NB + 8N + 3$	$N - 1$	NB	1
this work	$12W(\log_2(N) - 1) + 4(W + B + 9) + 22$	$12W(\log_2(N) - 1) + 2(W + B + 9) + 8$	N	$(\log_2(N) - 1)W$	0




Fig. 18. Measurements of 0.5ms using QPSK modulation. This work achieves far greater performance than standard BPS for the same computational complexity, while updating the frequency offset every 100 μ s.




Fig. 19. Measurements of 0.5ms using 16QAM modulation. This work achieves far greater performance than standard BPS for the same computational complexity, while updating the frequency offset every 100 μ s.

The 16QAM curves show similar behaviour to the simulations. The obtained BER floor can be explained by looking at the simulation results on figure 15. For each simulated residual frequency offset, a different BER floor is obtained. During the measurements, the frequency offset and residual offset changes continuously, meaning that the obtained BER floor is dependent on how much residual frequency offset is tolerated.

Both measurements show a great improvement with respect to the standard BPS algorithm, achieving pre-FEC BER values below the 6.7% HD FEC threshold in both cases, while using 50% less resources for frequency recovery and 81.2% less multipliers and 82.8% less adders for 16QAM phase recovery

TABLE II
RELATIVE AMOUNT OF COMPLEX MULTIPLICATIONS USED FOR PHASE ESTIMATION IN THIS WORK AND STANDARD BPS, RESULTING IN AN EQUIVALENT BER CURVE. EQUIVALENT NUMBER OF TEST ANGLES IS SET TO 64, BLOCKSIZE IS SET TO 32

architecture	multipliers	relative	additions	relative
BPS ₂	12544	100%	12418	100%
this work	2362	18.8%	2140	17.2%

as shown in table II.

VII. CONCLUSION

A novel low complexity carrier recovery system for baudrate sampled coherent interconnects was introduced in which a part of frequency tracking is offloaded to the phase recovery system, allowing increased parallelism and summing window. The cost function of BPS has been modified to significantly reduce the impact of the pattern effect. Simulation and measurement of a 32GBaud system confirm large power savings while also showing lower cycle slip rate and improved BER performance with respect to current state-of-the-art when using a solution that is identical in computational complexity. A pre-FEC BER below the 6.7% HD-FEC was demonstrated while reducing the amount of complex multipliers needed for phase recovery by 81.2% and the amount of adders by 82.8%, while reducing the complexity of frequency recovery by 50%.

REFERENCES

- [1] R. Nagarajan, I. Lyubomirsky, and O. Agazzi, "Low Power DSP-Based Transceivers for Data Center Optical Fiber Communications (Invited Tutorial)," en, *Journal of Lightwave Technology*, vol. 39, no. 16, pp. 5221–5231, Aug. 2021. DOI: 10.1109/JLT.2021.3089901.
- [2] Z. Jia and L. A. Campos, "Coherent Optics Ready for Prime Time in Short-Haul Networks," en, *IEEE Network*, vol. 35, no. 2, pp. 8–14, Mar. 2021, power consumption breakdown. DOI: 10.1109/MNET.011.2000612.
- [3] P. Ossieur, B. Moeneclaey, G. Coudyzer, J. Lambrecht, J. Craninckx, E. Martens, J. Van Driessche, C. Bruynsteen, J. De Busscher, J. Declercq, Y. Gu, S. Niu, T. Pannier, N. Singh, L. Van Severen, L. Bogaert, D. Maes, T. Vanackere, E. Vissers, J. Zhang, S. Kumari, R. Oldenbeuving, X. Zhang, G. Van Elzakker, J. Van Campenhout, P. Absil, B. Kuyken, X. Yin, G. Torfs, J. Missinne, G. Van Steenberge, G. Roelkens, and J. Bauwelinck, "Integrated Photonics and Electronics for Optical Transceivers Supporting AI/ML Applications," en, *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 31, no. 3: AI/ML Integrated Opto, pp. 1–16, May 2025. DOI: 10.1109/JSTQE.2025.3539379.
- [4] C. Fougstedt, O. Gustafsson, C. Bae, E. Börjeson, and P. Larsson-Edefors, "ASIC Design Exploration for DSP and FEC of 400-Gbit/s Coherent Data-Center Interconnect Receivers," en, in *Optical Fiber Communication Conference (OFC) 2020*, San Diego, California: Optica Publishing Group, 2020, Th2A.38. DOI: 10.1364/OFC.2020.Th2A.38.
- [5] D. A. Moreo, M. A. Castrillon, A. Aguirre, M. R. Hueda, and O. E. Agazzi, "Design Tradeoffs and Challenges in Practical Coherent Optical Transceiver Implementations," en, *Journal of Lightwave Technology*, vol. 34, no. 1, pp. 121–136, Jan. 2016. DOI: 10.1109/JLT.2015.2470114.

- [6] M. Selmi, "Accurate digital frequency offset estimator for coherent PolMux QAM transmission systems," en, *Proceedings of the 35th European Conference on Optical Communication*, 2009.
- [7] M. S. Faruk and S. J. Savory, "Digital Signal Processing for Coherent Transceivers Employing Multilevel Formats," en, *Journal of Lightwave Technology*, vol. 35, no. 5, pp. 1125–1141, Mar. 2017. DOI: 10.1109/JLT.2017.2662319.
- [8] X. Du, Q. Wang, and P.-Y. Kam, "Maximum Likelihood Estimation of Wiener Phase Noise Variance in Coherent Optical Systems," en, *Journal of Lightwave Technology*, vol. 42, no. 9, pp. 3163–3173, May 2024. DOI: 10.1109/JLT.2024.3357289.
- [9] X. Zhou, "Efficient Clock and Carrier Recovery Algorithms for Single-Carrier Coherent Optical Systems: A systematic review on challenges and recent progress," en, *IEEE Signal Processing Magazine*, vol. 31, no. 2, pp. 35–45, Mar. 2014. DOI: 10.1109/MSP.2013.2281071.
- [10] A. J. Viterbi and A. M. Viterbi, "Nonlinear estimation of PSK-modulated carrier phase with application to burst digital transmission," en, *IEEE Transactions on Information Theory*, vol. 29, no. 4, pp. 543–551, Jul. 1983. DOI: 10.1109/TIT.1983.1056713.
- [11] T. Pfau, S. Hoffmann, and R. Noe, "Hardware-Efficient Coherent Digital Receiver Concept With Feedforward Carrier Recovery for M-QAM Constellations," *Journal of Lightwave Technology*, vol. 27, no. 8, pp. 989–999, Apr. 2009. DOI: 10.1109/JLT.2008.2010511.
- [12] X. Li, Y. Cao, S. Yu, W. Gu, and Y. Ji, "A Simplified Feedforward Carrier Recovery Algorithm for Coherent Optical QAM System," en, *Journal of Lightwave Technology*, vol. 29, no. 5, pp. 801–807, Mar. 2011. DOI: 10.1109/JLT.2011.2106197.
- [13] E. Börjeson and P. Larsson-Edefors, "Cycle-Slip Rate Analysis of Blind Phase Search DSP Circuit Implementations," in *2020 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2020, pp. 1–3.
- [14] J. Li, L. Li, Z. Tao, T. Hoshida, and J. C. Rasmussen, "Laser-Linewidth-Tolerant Feed-Forward Carrier Phase Estimator With Reduced Complexity for QAM," *Journal of Lightwave Technology*, vol. 29, no. 16, pp. 2358–2364, Aug. 2011. DOI: 10.1109/JLT.2011.2159580.
- [15] Y. Gao, A. P. T. Lau, C. Lu, J. Wu, Y. Li, K. Xu, W. Li, and J. Lin, "Low-complexity two-stage carrier phase estimation for 16-QAM systems using QPSK partitioning and maximum likelihood detection," in *2011 Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference*, Mar. 2011, pp. 1–3.
- [16] Y. Chen and X. G. Huang, "A Linewidth-Tolerant Two-Stage CPE Using a New QPSK-Partitioning Approach and an Enhanced Maximum Likelihood Detection for 64-QAM Coherent Optical Systems," *Journal of Lightwave Technology*, vol. 33, no. 18, pp. 3883–3889, Sep. 2015. DOI: 10.1109/JLT.2015.2448113.
- [17] X. Zhou, "An Improved Feed-Forward Carrier Recovery Algorithm for Coherent Receivers With -QAM Modulation Format," en, *Photonics Technology Letters*, pp. 1051–1053, Jul. 2010.
- [18] H. Sun, K.-T. Wu, S. Thomson, and Y. Wu, "Novel 16QAM carrier recovery based on blind phase search," en, in *2014 The European Conference on Optical Communication (ECOC)*, Cannes, France: IEEE, Sep. 2014, pp. 1–3. DOI: 10.1109/ECOC.2014.6963896.
- [19] W. Pang, F. Wang, Q. Zhang, R. Gao, Z. Li, C. Wang, Y. Cui, Q. Xu, X. Yuan, H. Chang, and X. Xin, "Low Complexity Blind Phase Search for Coherent Optical Communication," in *2024 22nd International Conference on Optical Communications and Networks (ICOON)*, ISSN: 2771-3059, Jul. 2024, pp. 1–3. DOI: 10.1109/ICOON63276.2024.10648301.
- [20] F. Munier, E. Alpman, T. Eriksson, A. Svensson, and H. Zirath, "Estimation of Phase Noise for QPSK Modulation over AWGN Channels," in *Symp., Linköping, Sweden*, Nov. 4–5, 2003.
- [21] J. C. M. Diniz, Q. Fan, S. M. Ranzini, F. N. Khan, F. Da Ros, D. Zibar, and A. P. T. Lau, "Low-complexity carrier phase recovery based on principal component analysis for square-QAM modulation formats," en, *Optics Express*, vol. 27, no. 11, p. 15 617, May 2019. DOI: 10.1364/OE.27.015617.