



Circuit Implementation of Pilot-Based Dynamic MIMO Equalization for Coupled-Core Fibers

Downloaded from: <https://research.chalmers.se>, 2024-11-05 01:17 UTC

Citation for the original published paper (version of record):

Börjeson, E., Deriushkina, E., Karlsson, M. et al (2024). Circuit Implementation of Pilot-Based Dynamic MIMO Equalization for Coupled-Core Fibers. Optical Fiber Communication Conference, OFC 2024

N.B. When citing this work, cite the original published paper.

Circuit Implementation of Pilot-Based Dynamic MIMO Equalization for Coupled-Core Fibers

Erik Börjesson¹, Ekaterina Deriushkina¹, Mikael Mazur²,
Magnus Karlsson¹, and Per Larsson-Edefors¹

¹ Chalmers University of Technology, Gothenburg, Sweden

² Nokia Bell Labs, 600 Mountain Ave., Murray Hill, NJ 07974, USA

erikbor@chalmers.se

Abstract: We explore ASIC implementation for pilot-based MIMO equalizers for coupled-core transmission, considering chip area scaling trends and performance impact of time-dependent drift. For a system with 28-GBd subcarriers, an equalizer for 8×8 is 5.3 times larger than for 2×2 . © 2023 The Author(s)

1. Introduction

Space division multiplexing (SDM) using multi-core (MCF) or multi-mode fibers (MMF) is a promising way to overcome the emerging capacity crunch in fiber-optic networks by tapping into the spatial dimension [1]. Compared to single-mode fibers (SMF), the spatial density is increased by the number of cores or modes that share a common cladding. Uncoupled MCFs (UC-MCF) are directly compatible with today's systems and transponders designed for SMF and recently the first deployment of a submarine cable using UC-MCF was announced [2]. However, to avoid large penalties from cross-talk, UC-MCF offers only limited scalability. Coupled-core MCFs (CC-MCF) are MCFs with the cores purposely placed at an optimal distance to introduce strong coupling [3]. These fibers have a very compact impulse response compared to MMF, with a loss directly comparable to SMFs made from equivalent cores, which makes CC-MCF a strong candidate to further scale the spatial density in future submarine systems. These systems are highly space constrained and the fiber count using regular fiber diameters, which is preferred for reliability, is limited to 24 fiber pairs [4]. In addition, the cables are power limited and increasing the spatial dimensionality is therefore vital to maximize cable capacity [5]. However, while coupled-core fibers with up to 19 cores have been demonstrated [6], the inherent drawback of both CC-MCF and MMF is the requirement of larger multiple-input-multiple-output (MIMO) digital signal processing (DSP) systems to counteract the mixing between signals launched on different cores/modes. For a 4-core CC-MCF, the MIMO dimensionality must be increased to an 8×8 matrix compared to a 2×2 for SMF. While real-time feasibility has been shown using FPGAs over both field-deployed CC-MCFs [7] and transoceanic distances using a recirculating loop [8], no systematic studies of the required ASIC complexity have been performed to date.

In this work we report the first ASIC DSP implementation focused on scaling the spatial dimensionality beyond the well established 2×2 case. We focus on a DSP implementation using subcarriers and study the scaling of the MIMO equalizer needed for each of these carriers for dimensionality reaching 8×8 . Our implementation operates using 16QAM data symbols at a target bit-error rate (BER) of 10^{-3} . Synthesis using a 22-nm CMOS process allows us to study the impact of circuit parallelism etc. on area usage. Furthermore, we develop a model for the time-dependent drift of the coupled-core channel and use this, with parameters fit to previous long-term measurements of field-deployed CC-MCFs, to study equalizer performance. We show that for a total symbol rate of 28 GBd per subcarrier, an 8×8 equalizer occupies 5.3 times more chip area compared to a 2×2 implementation. Our results demonstrate the feasibility of ASIC implementations targeting higher-dimensionality SDM systems and establish area/complexity benchmarks for further DSP development.

2. Equalizer Design

The equalizer is based on the least mean square (LMS) algorithm and uses pilot symbols to update the filters. For N I/Q input signals, our adaptive equalizer is constructed using a filter bank of $N \times N$ complex finite-impulse response (FIR) filters, $\mathbf{w}_{i,j}$, where i and j are the input and output indices, respectively. Pilots symbols are inserted periodically in all N symbol streams, to enable calculation of a complex estimation error at sampling time k as

$$e_j(k) = r_j(k) - p_j(k), \quad (1)$$

where p_j is the received pilot after filtering and r_j is the expected value. The estimation error is then used to update the filter coefficients as

$$\mathbf{w}_{i,j}(k+1) = \mathbf{w}_{i,j}(k) + \mu e_j(k) \mathbf{x}_i^*(k), \quad (2)$$

where \mathbf{x}_i^* is the complex conjugate of the input symbols and μ is the step size.

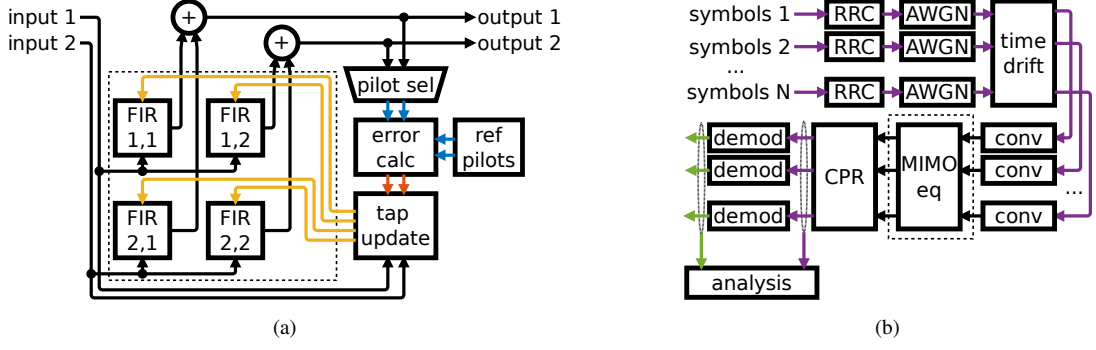


Fig. 1: Block diagrams of (a) a 2×2 equalizer architecture, where the bank containing 2×2 complex filters is marked with a dashed square, and (b) the simulation system model, where the circuit's HDL model is marked with a dashed square.

Our circuit implementation of the LMS equalizer is parameterized and allows the number of inputs (N) and filter taps, pilot-insertion ratio, step size, circuit parallelism P , and signal wordlengths to be adjusted for different systems. A simplified block diagram of a 2×2 version of our circuit is shown in Fig. 1a, with $P = 1$ and all pipelining registers removed for clarity. The input signals are 2X-oversampled complex symbols which are used as inputs to the parallel complex FIR filters. The filter outputs are 1 sample-per-symbol signals and since the parallelism and the pilot-insertion ratio do not necessarily match, the *pilot sel* sub-unit is used to extract the filtered pilot symbols from the correct lane. These are single-lane signals, marked in blue. After error calculation is performed as in (1), the estimated error (red), which has a one bit longer wordlength than the symbols to avoid overflow, is fed to the *tap update* sub-unit. In this unit, the original input symbols are delayed to match the pipelining of the design and used to update the filter according to (2). The step size (μ) is implemented as a right shift, to avoid additional multiplications, limiting the step-size parameter to 2^{-m} , where m is a positive integer. The coefficient values stored in the *tap update* sub-unit use a longer wordlength than the FIR coefficients used for filtering, to allow for smaller step sizes without losing precision.

3. Channel Model

Our simulation setup for an $N \times N$ system is shown in Fig. 1b. The randomized 16QAM symbols with inserted QPSK pilots are fed to a 2X-upsampling root-raised cosine (RRC) pulse-shaping filter. After addition of additive-white Gaussian noise (AWGN), the signals are converted to fixed-point representation, parallelized, and evaluated using logic simulation of a hardware description (HDL) model of the MIMO equalizer circuit. The resulting signals go through pilot-based carrier phase recovery (CPR) before demodulation and analysis.

The dynamic (temporally drifting) channel model used to test the equalizer is a unitary $N \times N$ transfer matrix $\mathbf{U}(k)$ updated at discrete times k , with additive white Gaussian noise. The update is given by $\mathbf{U}(k+1) = \exp(j\mathbf{H}(k)\gamma)\mathbf{U}(k)$, where $\mathbf{H}(k)$ is a random $N \times N$ Hermitian matrix whose nonzero real/imaginary parts are random iid as $\mathcal{N}(0, 1)$, and γ is a numerical parameter controlling how correlated $\mathbf{U}(k+1)$ and $\mathbf{U}(k)$ are. We can show that $E[\mathbf{U}(k)] \approx \exp(-\gamma^2 k)\mathbf{U}(0)$, enabling us to relate γ to a characteristic drift time T_D via $k\gamma^2 = k/(F_{\text{samp}}T_D)$, where F_{samp} is the sampling rate.

4. Results

We assume a pilot insertion ratio of $1/64$, for which the QPSK pilots are inserted simultaneously in all 16QAM input streams. We use 9 taps for the FIR filters and a signal-to-noise ratio of 15 dB for the upsampled symbols, resulting in a BER $\approx 10^{-3}$. The wordlengths of all signals are selected to strike a good balance between performance and area; subsequently 8 bits were used for the input symbols.

The convergence behavior of our MIMO equalizer circuit is shown in Fig. 2, where the step size (μ) was selected to minimize the root mean square of the error-vector magnitude of the received symbols, once the equalizer has converged. The convergence is slower for a high T_D , as μ can be smaller while still allowing for good tracking of the temporal drift, resulting in a lower error after convergence. Significant error levels emerge for $T_D = 0.01$ ms, where the equalizer can no longer keep up with the faster drift. This is especially pronounced for the more sensitive 8×8 circuit. For $T_D = 1$ ms the performance is similar to a system with no time drift. This value is a reasonable assumption based on [7], where a 69-km CC-MCF placed in a stable tunnel provides a good benchmark for underground or submarine cables. The higher error floor for 8×8 can be reduced to the level of 4×4 by increasing all signal wordlengths by 1 bit, but this will incur an area penalty.

To estimate equalizer area, HDL models of the equalizer component (Fig. 1a) are synthesized to a 22-nm FD-SOI CMOS process, for a clock rate of 1.17 GHz. By increasing the circuit parallelism (P), the symbol rate (F_{symb}) can be adjusted upwards at the cost of larger area. Fig. 3a shows the cell area as a function of the total symbol rate for one subcarrier (NF_{symb}), using different number of cores. The encircled data points show different circuit

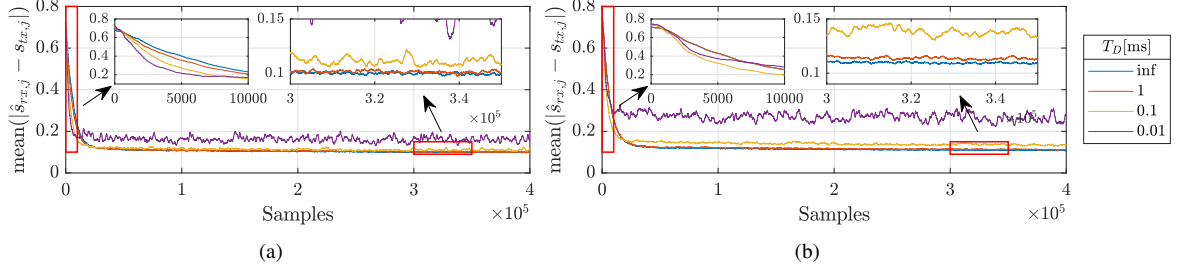


Fig. 2: Convergence behavior of (a) a 4×4 and (b) an 8×8 circuit, where the y-axes show the absolute value of the difference between the transmitted and received signals, averaged over all inputs using a moving average filter.

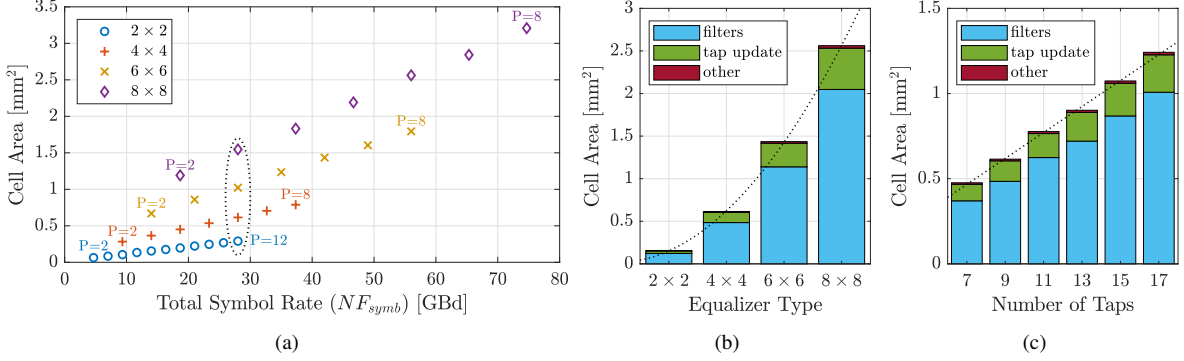


Fig. 3: Area based on synthesis using 8 data bits, 9 filter taps and 9 bits for the filter coefficients. In (a) we vary dimensionality and circuit parallelism (P), resulting in different total symbol rates. The implementation area, divided into sub-units, for $P = 6$ (7 GBd/input) is shown in (b), while (c) shows the area for 4×4 implementations with $P = 6$ using different filter tap counts.

configurations which result in identical NF_{symb} but which use different number of cores and P , clearly illustrating the challenge of equalization for CC-MCF with higher core counts. The quadratic scaling of the area to higher dimensionality is emphasized in Fig. 3b, where the area distribution between different equalizer sub-units is also shown. The majority of the area is used by the *filters*, while the significance of the *tap update* sub-unit becomes larger for higher number of cores. Because its single-lane operations are straightforward, the area of the error calc sub-unit is negligible. Thus, this sub-unit has been included in the *other* category, along with e.g. pipelining registers. The scaling with additional filter taps is linear for all configurations as shown in Fig. 3c for the 4×4 case. There are many potential optimizations of this baseline implementation, especially when also considering aspects of DSP power dissipation. One such example is clock gating of the feedback loop, as suggested in [9].

5. Conclusion

By performing ASIC implementations of pilot-based MIMO equalizers for coupled-core fibers, we have been able to quantify how circuit area scales with increased dimensionality and other parameters. Furthermore, using a new model for the temporal drift of a coupled-core channel, we were able to analyze MIMO equalizer convergence. Our results clearly show that an increasing MIMO dimensionality significantly impacts the equalizer circuit area, mainly because of the increased filter complexity but to some extent because of longer wordlengths.

We thank GlobalFoundries for the design kit provided through the University Partnership Program.

References

1. B. J. Puttnam *et al.*, “Space-division multiplexing for optical fiber communications,” *Optica* **8**, 1186–1203 (2021).
2. B. Quigley *et al.*, “Boosting subsea cables with multi-core fiber technology,” <https://cloud.google.com/blog/products/infrastructure/delivering-multi-core-fiber-technology-in-subsea-cables>.
3. T. Hayashi *et al.*, “Record-low spatial mode dispersion and ultra-low loss coupled multi-core fiber for ultra-long-haul transmission,” *IEEE JLT* **35**, 450–457 (2016).
4. G. Mohs *et al.*, “High-capacity submarine cables - past, present and future,” in *OFC*, (2023), p. Tu3G.1.
5. R. Dar *et al.*, “Cost-optimized submarine cables using massive spatial parallelism,” *IEEE JLT* **36**, 3855–3865 (2018).
6. G. Rademacher *et al.*, “Randomly coupled 19-core multi-core fiber with standard cladding diameter,” in *OFC*, (2023), p. Th4A.4.
7. M. Mazur *et al.*, “Real-time MIMO transmission over field-deployed coupled-core multi-core fibers,” in *OFC*, (2022), p. Th4B.8.
8. S. Bepu *et al.*, “Long-haul coupled 4-core fiber transmission over 7,200 km with real-time MIMO DSP,” *IEEE JLT* **40**, 1640–1649 (2022).
9. C. Fougstedt *et al.*, “Dynamic equalizer power dissipation optimization,” in *OFC*, (2016), p. W4A.2.