



CHALMERS

Chalmers Publication Library

Methodology for Power-Aware Coherent Receiver Design

This document has been downloaded from Chalmers Publication Library (CPL). It is the author's version of a work that was accepted for publication in:

Advanced Photonics Congress (SPPCom) 2013

Citation for the published paper:

Ahmad, T. ; Ai, Y. ; Muralidharan, P. (2013) "Methodology for Power-Aware Coherent Receiver Design". Advanced Photonics Congress (SPPCom) 2013 pp. SPT4D.4.

Downloaded from: <http://publications.lib.chalmers.se/publication/181057>

Notice: Changes introduced as a result of publishing processes such as copy-editing and formatting may not be reflected in this document. For a definitive version of this work, please refer to the published source. Please note that access to the published version might require a subscription.

Chalmers Publication Library (CPL) offers the possibility of retrieving research publications produced at Chalmers University of Technology. It covers all types of publications: articles, dissertations, licentiate theses, masters theses, conference papers, reports etc. Since 2006 it is the official tool for Chalmers official publication statistics. To ensure that Chalmers research results are disseminated as widely as possible, an Open Access Policy has been adopted. The CPL service is administrated and maintained by Chalmers Library.

(article starts on next page)

Methodology for Power-Aware Coherent Receiver Design

Tauseef Ahmad¹, Yun Ai², Pavithra Muralidharan¹, Naga V. Irukulapati², Pontus Johannisson³,
Henk Wymeersch², Erik Agrell², Per Larsson-Edefors¹, and Magnus Karlsson³

¹ Dept. of Computer Science and Engineering, ² Dept. of Signals and Systems, ³ Photonics Laboratory,
Chalmers University of Technology, SE-412 96 Göteborg, Sweden
pontus.johannisson@chalmers.se

Abstract: We describe a methodology to design and evaluate DSP hardware for a coherent receiver. Important parameters that can be assessed include DSP power consumption and chip area.
OCIS codes: (060.1660) Coherent communications; (060.4510) Optical communications

1. Introduction

In coherent optical transponders, digital signal processing (DSP) is an essential part which removes the need for optical dispersion compensation and analog phase/polarization tracking. While coherent systems clearly have many fundamental advantages, the corresponding system implementation is challenging as there are many implementation dimensions to explore and optimize, e.g., algorithms and their parameters, numerical resolution, hardware process technology, operating speed, and parallelism. To add to the implementation challenge, not only performance but also energy efficiency is rapidly becoming an important concern, which a number of recent studies bear testament to: Power estimates for entire networks have been computed, see, e.g., [1], while other studies have focused on parts such as the dispersion compensation [2].

The question then arises how to find a practically viable and reliable implementation process that can take algorithm choices, algorithm implementation in hardware, receiver performance, and receiver power consumption into account. This complex issue is trifold cross-disciplinary, requiring hardware implementation, algorithmic, and optical communication expertise. We describe an effort to find a methodology to perform application-specific integrated circuit (ASIC) design exploration. As a case study, we present the SNR requirement vs. DSP power consumption trade-off and give a division of the different parts of the receiver DSP in terms of power consumption and required chip area. Although the results describe a specific receiver design, the approach and methodology described in Section 2 should be of general interest. Even if results similar to the ones in Section 3 are likely internally available to a limited set of industry engineers, we believe this is the first openly published (albeit limited in scope) study of this kind.

2. Design Exploration

There are many implementation parameters, e.g., fixed-point word lengths and clock rate, which influence the final hardware implementation and its key metrics of chip area and power consumption. Unless it affects the overall system architecture adversely, minimizing the fixed-point data word length in all blocks leads to the lowest power consumption. Regarding clock rate, advanced CMOS circuits can be clocked at 10 GHz, but this requires an excessive amount of flip-flops and is not power efficient. Rather, a lower clock rate combined with DSP parallelization is preferable. The identification of an optimal combination of clock rate and parallelization requires simulations as it entails balancing static (leakage) and dynamic (switching) power consumption.

In order to enable cross-disciplinary collaboration in designing and evaluating the receiver DSP hardware we devised the following methodology, which is also illustrated in Fig. 1.

1. The receiver DSP is defined and the performance is evaluated in numerical simulations using a floating-point representation in MATLAB.
2. The receiver DSP blocks are modified to operate on fixed-point numbers. This is done with the MATLAB fixed-point toolbox and this model is used for two tasks:
 - (a) To generate the test vectors that are needed to verify the hardware and to calculate the power consumption.
 - (b) To analyze the impact of the limited resolution on the system's bit error rate (BER) and the trade-off between required SNR and receiver data resolution which, in turn, affects the ASIC power consumption.
3. The DSP blocks are coded in a hardware description language (VHDL) such that arbitrary fixed-point word lengths can be used. We emphasize that this step cannot be automated in any reliable way. All blocks are verified with the MATLAB system model using the Cadence Incisive Enterprise Simulator (IES) [3].
4. The VHDL blocks are synthesized in the Cadence RTL Compiler [4] for a 65 nm 0.9 V general-purpose standard- V_T cell library assuming nominal transistor parameters (process corner) and 25 °C. Area and circuit

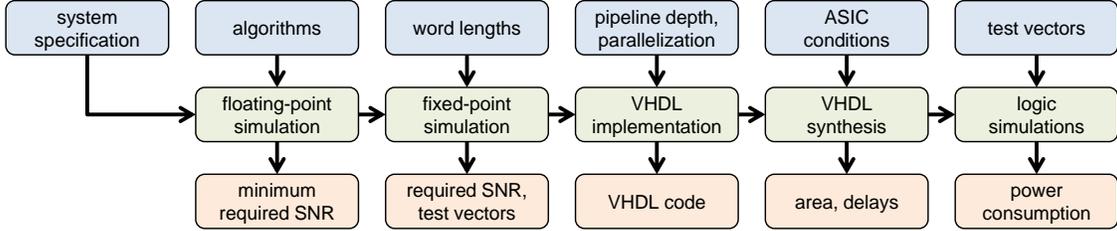


Fig. 1: A flow chart outlining the proposed design methodology. The top row shows the input to the different process steps in the middle row and the bottom row is the output.

delay (timing) information are obtained for the resulting gate netlists.

- Using the block netlists and the test vectors, IES logic simulations are performed to generate signal switching information (VCD), which is back-annotated to step 4 from which power information is extracted using the RTL Compiler.

This design process is a stable design flow based on commercial software, but we have identified three practical limitations of our study.

- The only optimization of the hardware we have performed is the tuning of the word lengths. This means that power and area improvements are possible for all blocks.
- All results are for synthesized netlists, i.e., no physical implementation (place and route) has been performed. The physical implementation generally leads to somewhat higher area and power values, since signal buffers (including clock tree) need to be inserted to compensate for RC delays.
- When power consumption is a critical metric, ASICs are often implemented using custom circuit design, which potentially reduces power consumption significantly but leads to extremely high development costs [5].

3. Case Study

We study the receiver of a fiber-optic communication system operating at 28 Gbaud, using polarization-multiplexed quadrature phase-shift keying (PM-QPSK), corresponding to a bit rate of 112 Gb/s. The optical front-end is followed by the analog-to-digital converter (ADC) and the subsequent DSP blocks: chromatic dispersion (CD) compensation, symbol timing recovery (STR), and a dynamic feedback equalizer with polarization and phase tracking. The final output corresponds to symbol decisions.

Chromatic dispersion compensation: CD, modeled as an all-pass filter, is one of the dominant linear impairments, causing inter-symbol interference over many symbol slots. In this model, the CD for a fiber length of 800 km is compensated in the time domain using a 341-tap static FIR filter [6].

Symbol timing recovery: The timing offset has to be estimated and compensated for the proper extraction of symbols. In this paper, timing estimation is done using the Gardner algorithm [7], operating at 2 samples per symbol, while interpolation is performed using a linear interpolator [8].

Dynamic equalization: To compensate for the dynamic effects, including polarization mixing, polarization-mode dispersion, and phase noise, we employ a dynamic equalizer with a phase-feedback loop; see [9] for implementation details. For the FIR filters, we use the typical butterfly structure. Assuming steady-state operation, we rely on a decision-directed method [6]. Before making the final decision on the transmitted symbol, frequency offset and laser phase noise are estimated and compensated for [10].

3.1. Results and Discussion

Fig. 2 shows the DSP power consumption as a function of the required input SNR for the selected system configurations at a fixed BER of 10^{-3} . With the parameters mentioned above, the CD compensation filter dissipates a large fraction of the total power and we have used the minimum word length (5 bits) in the CD filter that can reach the target BER. We have then varied the parameters for the dynamic equalizer, which is the second largest DSP block, and observed the effect of the parameter changes.

The theoretical minimum SNR¹ in an additive white Gaussian noise (AWGN) channel when using QPSK at a BER of 10^{-3} is 9.8 dB, but in the floating-point simulations in MATLAB, the required SNR was 10.3 dB. This SNR penalty

¹SNR is defined as E_s/N_0 , where E_s is the average symbol energy and N_0 is the single-sided noise spectral density.

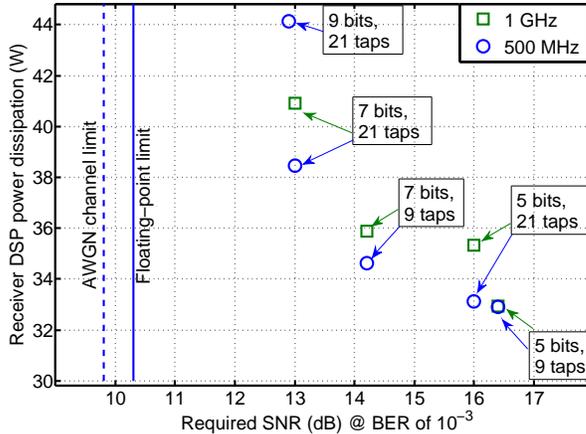


Fig. 2: Power consumption as a function of the required SNR for different dynamic equalizer filter lengths (in taps) and word lengths (in bits). Two different clock rates, 1 GHz (squares) and 500 MHz (circles), were evaluated.

500 MHz	Area (mm ²)	Power (W)
CD filter	47.92	34.63
STR	0.27	0.13
Dyn. eq.	6.97	3.82
Total	55.16	38.58

1 GHz	Area (mm ²)	Power (W)
CD filter	27.72	35.89
STR	0.16	0.14
Dyn. eq.	4.25	5.29
Total	32.13	41.32

Table 1: Area and power consumption for different receiver blocks with a configuration of 5 bits and 341 taps for the CD FIR filter, 6 bits for the STR block, and 7 bits and 9 taps for the dynamic equalizer.

is due to the above described receiver algorithms, which were chosen in order to have a moderate implementation complexity. We can view 10.3 dB as the minimum required SNR when using long word lengths.

The impact of the dynamic equalizer word length, the number of FIR filter taps, and the clock rate on the power consumption is shown in Fig. 2. (Since the ASIC design software is based on heuristic algorithms, the power estimate results exhibit some uncertainties.) It is seen that it is possible to achieve a lower required SNR at the cost of increased power consumption and that the power consumption is not significantly affected by the clock rate. This means that the doubling of the area for the lower clock rate is compensated for by the lower switching frequency.

We see that with 7 bits word length, the required SNR is around 13 dB. If we allow an additional penalty of around 1 dB, we can reduce the power consumption with approximately 5 W by using a shorter dynamic equalizer filter. Additional (smaller) reductions in power consumption are possible by reducing the word length to 5 bits, but at an additional 2–3 dB cost of required SNR. Fig. 2 shows that by increasing the resolution/tap count of the dynamic equalizer, the required SNR will not reach the floating-point limit and the reason is our choice to limit the resolution in the CD filter. This penalty can be avoided at the cost of an increased DSP power consumption and an ADC with higher resolution.

Table 1 shows the area requirement and power consumption estimates at two different clock frequencies for a configuration of 5 bits and 341 taps for the CD FIR filter, 6 bits for the STR block, and 7 bits and 9 taps for the dynamic equalizer. As the CD compensation is dominating the power consumption, it is important to optimize this part. On the other hand, since the timing estimation and interpolation is such a small part of the total complexity, it is reasonable to design this block to have high performance.

In conclusion, we have presented a hardware design and evaluation methodology that allows estimation of the hardware requirements of the different receiver DSP parts. Using this model, we have performed a case study of the ASIC power consumption and the trade-off between algorithm performance and hardware complexity. Although these results are based on a number of specific design choices, we believe that the described methodology will be useful in future more general studies.

References

- [1] D. Kilper, K. Guan, K. Hinton, and R. Ayre, Proc. IEEE **100**, 1168–1187 (2012).
- [2] B. S. G. Pillai, B. Sedighi, W. Shieh, and R. S. Tucker, in “Optical Fiber Communication Conference (OFC),” (2012), p. OM3A.8.
- [3] Cadence, http://www.cadence.com/rl/Resources/datasheets/incisive_enterprise_specman.pdf.
- [4] Cadence, http://www.cadence.com/rl/Resources/datasheets/encounter_rtlcompiler.pdf.
- [5] W. J. Dally and A. Chang, in “Proceedings of the 37th Annual Design Automation Conference,” (2000), pp. 643–647.
- [6] S. J. Savory, Opt. Express **16**, 804–817 (2008).
- [7] F. M. Gardner, IEEE Trans. Commun. **COM-34**, 423–429 (1986).
- [8] T. Ahmed, Master’s thesis, Chalmers University of Technology (2012), <http://publications.lib.chalmers.se/records/fulltext/161391.pdf>.
- [9] F. Toft *et al.*, in “IEEE Int. Symp. on Circuits and Systems (ISCAS),” (2012), pp. 3234–3237.
- [10] B.-E. Olsson, J. Mårtensson, A. Kristiansson, and A. Alping, in “Optical Fiber Communication Conference (OFC),” (2010), p. OMK5.